Chapter 8

Eigenvalues and Singular Values

So far, our physical applications of linear algebra have concentrated on statics: unchanging equilibrium configurations of mass–spring chains, circuits, and structures, all modeled by linear systems of algebraic equations. It is now time to set the universe in motion. In general, a (continuous) *dynamical system* refers to the (differential) equations governing the time-varying motion of a physical system, be it mechanical, electrical, chemical, fluid, thermodynamical, biological, financial, ... Our immediate goal is to solve the simplest class of continuous dynamical models, which are first order autonomous linear systems of ordinary differential equations.

We begin with a very quick review of the scalar case, whose solutions are simple exponential functions. This inspires us to try to solve a vector-valued linear system by substituting a similar exponential solution formula. We are immediately led to the system of algebraic equations that define the eigenvalues and eigenvectors of the coefficient matrix. Thus, before we can make any progress in our study of differential equations, we need to learn about eigenvalues and eigenvectors, and that is the purpose of the present chapter. Dynamical systems are used to motivate the subject, but serious applications will be deferred until Chapter 10. Additional applications of eigenvalues and eigenvectors to linear iterative systems, stochastic processes, and numerical solution algorithms for linear algebraic systems form the focus of Chapter 9.

Each square matrix possesses a collection of one or more complex scalars, called eigenvalues, and associated vectors, called eigenvectors. From a geometrical viewpoint, the matrix defines a linear transformation on Euclidean space; the eigenvectors indicate the directions of pure stretch and the eigenvalues the extent of stretching. We will introduce the non-standard term "complete" to describe matrices whose (complex) eigenvectors form a basis of the underlying vector space. A more common name for such matrices is "diagonalizable", because, when expressed in terms of its eigenvector basis, the matrix representing the corresponding linear transformation assumes a very simple diagonal form, facilitating the detailed analysis of its properties. A particularly important class consists of the symmetric matrices, whose eigenvectors form an orthogonal basis of $\mathbb{R}^n$; in fact, this is how orthogonal bases most naturally appear. Most matrices are complete; incomplete matrices are trickier to deal with, and we discuss their non-diagonal Schur decomposition and Jordan canonical form in Section 8.6.

A non-square matrix A does not possess eigenvalues. In their place, one studies the eigenvalues of the associated square Gram matrix $K = A^T A$, whose square roots are known as the singular values of the original matrix. The corresponding singular value decomposition (SVD) supplies the final details for our understanding of the remarkable geometric structure governing matrix multiplication. The singular value decomposition is used to define the pseudoinverse of a matrix, which provides a mechanism for "inverting" non-square and singular matrices, and an alternative construction of least squares solutions to general linear systems. Singular values underlie the powerful method of modern statistical data analysis known as principal component analysis (PCA), which will be developed in the

© Springer International Publishing AG, part of Springer Nature 2018

P. J. Olver, C. Shakiban, *Applied Linear Algebra*, Undergraduate Texts in Mathematics,
https://doi.org/10.1007/978-3-319-91041-3_8

final section of this chapter and appears in an increasingly broad range of contemporary applications, including image processing, semantics, language and speech recognition, and machine learning.

Remark. The numerical computation of eigenvalues and eigenvectors is a challenging issue, and must be deferred until Section 9.5. Unless you are prepared to consult that section in advance, solving the computer-based problems in this chapter will require access to computer software that can accurately compute eigenvalues and eigenvectors.

8.1 Linear Dynamical Systems

Our new goal is to solve and analyze the simplest class of dynamical systems, namely those modeled by first order linear systems of ordinary differential equations. We begin with a thorough review of the scalar case, including a complete investigation into the stability of their equilibria — in preparation for the general situation to be treated in depth in Chapter 10. Readers who are not interested in such motivational material may skip ahead to Section 8.2 without incurring any penalty.

Scalar Ordinary Differential Equations

Consider the elementary first order scalar ordinary differential equation

$$\frac{du}{dt} = a\,u. \tag{8.1}$$

Here $a \in \mathbb{R}$ is a real constant, while the unknown $u(t)$ is a scalar function. As you no doubt already learned, e.g., in [**7**, **22**, **78**], the general solution to (8.1) is an exponential function

$$u(t) = c\,e^{a\,t}. \tag{8.2}$$

The integration constant c is uniquely determined by a single *initial condition*

$$u(t_0) = b \tag{8.3}$$

imposed at an initial time t_0. Substituting $t = t_0$ into the solution formula (8.2),

$$u(t_0) = c\,e^{a\,t_0} = b, \qquad \text{and so} \qquad c = b\,e^{-a\,t_0}.$$

We conclude that

$$u(t) = b\,e^{a(t-t_0)} \tag{8.4}$$

is the unique solution to the scalar initial value problem (8.1, 3).

Example 8.1. The radioactive decay of an isotope, say uranium-238, is governed by the differential equation

$$\frac{du}{dt} = -\,\gamma\,u. \tag{8.5}$$

Here $u(t)$ denotes the amount of the isotope remaining at time t; the coefficient $\gamma > 0$ governs the decay rate. The solution is an exponentially decaying function $u(t) = c\,e^{-\gamma\,t}$, where $c = u(0)$ is the amount of radioactive material at the initial time $t_0 = 0$.

The isotope's *half-life* $t_\star$ is the time it takes for half of a sample to decay, that is, when $u(t_\star) = \frac{1}{2}\,u(0)$. To determine $t_\star$, we solve the algebraic equation

$$e^{-\gamma\,t_\star} = \frac{1}{2}, \qquad \text{so that} \qquad t_\star = \frac{\log 2}{\gamma}. \tag{8.6}$$

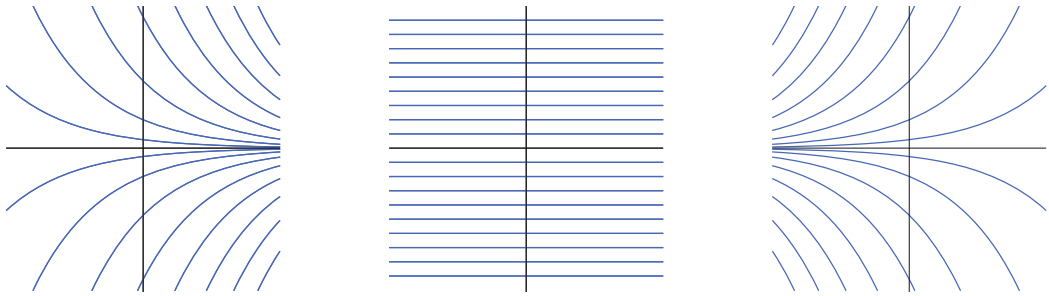

Figure 8.1. Solutions to $\dot{u} = a\,u$.

Thus, after an elapsed time of $t_\star$, half of the original material has decayed. After a further $t_\star$, half of the remaining half has decayed, leaving a quarter of the original radioactive material. And so on, so that at each integer multiple $n\,t_\star$, $n \in \mathbb{N}$, of the half-life, the remaining amount of the isotope is $u(n\,t_\star) = 2^{-n}\,u(0)$.

Returning to the general situation, let us make some elementary, but pertinent, observations about this simplest linear dynamical system. First of all, since the equation is homogeneous, the zero function $u(t) \equiv 0$ — corresponding to $c = 0$ in the solution formula (8.2) — is a constant solution, known as an *equilibrium solution*, or *fixed point*, since it does not depend on t. If the coefficient $a > 0$ is positive, then the solutions (8.2) are exponentially growing (in absolute value) as $t \to +\infty$. This implies that the zero equilibrium solution is *unstable*. The initial condition $u(t_0) = 0$ produces the zero solution, but if we make a tiny error (either physical, numerical, or mathematical) in the initial data, say $u(t_0) = \varepsilon$, then the solution $u(t) = \varepsilon\, e^{a(t-t_0)}$ will eventually be far away from equilibrium. More generally, any two solutions with very close, but not equal, initial data will eventually become arbitrarily far apart: $|\,u_1(t) - u_2(t)\,| \to \infty$ as $t \to \infty$. One consequence is an inherent difficulty in accurately computing the long-time behavior of the solution, since small numerical errors may eventually have very large effects.

On the other hand, if $a < 0$, the solutions are exponentially decaying in time. In this case, the zero equilibrium solution is *stable*, since a small change in the initial data will have a negligible effect on the solution. In fact, the zero solution is *globally asymptotically stable*. The phrase "asymptotically stable" indicates that solutions that start out near the equilibrium point approach it in the large time limit; more specifically, if $u(t_0) = \varepsilon$ is small, then $u(t) \to 0$ as $t \to \infty$. "Globally" implies that *all* solutions, no matter how large the initial data, eventually approach equilibrium. In fact, for a linear system, the stability of an equilibrium solution is inevitably a global phenomenon.

The borderline case is $a = 0$. Then all the solutions to (8.1) are constant. In this case, the zero solution is *stable* — indeed, *globally stable* — but not asymptotically stable. The solution to the initial value problem $u(t_0) = \varepsilon$ is $u(t) \equiv \varepsilon$. Therefore, a solution that starts out near equilibrium will remain nearby, but will not asymptotically approach it. The three qualitatively different possibilities are illustrated in Figure 8.1. A formal definition of the various notions of stability for dynamical systems can be found in Definition 10.14.

Exercises

8.1.1. Solve the following initial value problems:

(a) $\dfrac{du}{dt} = 5u, \;\; u(0) = -3,$ (b) $\dfrac{du}{dt} = 2u, \;\; u(1) = 3,$ (c) $\dfrac{du}{dt} = -3u, \;\; u(-1) = 1.$

8.1.2. Suppose a radioactive material has a half-life of 100 years. What is the decay rate γ? Starting with an initial sample of 100 grams, how much will be left after 10 years? after 100 years? after 1,000 years?

8.1.3. Carbon-14 has a half-life of 5730 years. Human skeletal fragments discovered in a cave are analyzed and found to have only 6.24% of the carbon-14 that living tissue would have. How old are the remains?

8.1.4. Prove that if $t_\star$ is the half-life of a radioactive material, then $u(n\,t_\star) = 2^{-n}\,u(0)$. Explain the meaning of this equation in your own words.

8.1.5. A bacteria colony grows according to the equation $du/dt = 1.3\,u$. How long until the colony doubles? quadruples? If the initial population is 2, how long until the population reaches 2 million?

8.1.6. Deer in northern Minnesota reproduce according to the linear differential equation $\dfrac{du}{dt} = .27\,u$ where t is measured in years. If the initial population is $u(0) = 5{,}000$ and the environment can sustain at most 1,000,000 deer, how long until the deer run out of resources?

♢ 8.1.7. Consider the inhomogeneous differential equation $\dfrac{du}{dt} = a\,u + b$, where a, b are constants.

(a) Show that $u_\star = -b/a$ is a constant equilibrium solution. (b) Solve the differential equation. *Hint*: Look at the differential equation satisfied by $v = u - u_\star$.

(c) Discuss the stability of the equilibrium solution $u_\star$.

8.1.8. Use the method of Exercise 8.1.7 to solve the following initial value problems:

(a) $\dfrac{du}{dt} = 2u - 1, \; u(0) = 1,$ (b) $\dfrac{du}{dt} = 5u + 15, \; u(1) = -3,$ (c) $\dfrac{du}{dt} = -3u + 6, \; u(2) = -1.$

8.1.9. The radioactive waste from a nuclear reactor has a half-life of 1000 years. Waste is continually produced at the rate of 5 tons per year and stored in a dump site.
(a) Set up an inhomogeneous differential equation, of the form in Exercise 8.1.7, to model the amount of radioactive waste. (b) Determine whether the amount of radioactive material at the dump increases indefinitely, decreases to zero, or eventually stabilizes at some fixed amount. (c) Starting with a brand new site, how long will it be until the dump contains 100 tons of radioactive material?

♡ 8.1.10. Suppose that hunters are allowed to shoot a fixed number of the northern Minnesota deer in Exercise 8.1.6 each year. (a) Explain why the population model takes the form $\dfrac{du}{dt} = .27\,u - b$, where b is the number killed yearly. (Ignore the seasonal aspects of hunting.)

(b) If $b = 1{,}000$, how long until the deer run out of resources? *Hint*: See Exercise 8.1.7.
(c) What is the maximal rate at which deer can be hunted without causing their extinction?

8.1.11. (a) Write down the exact solution to the initial value problem $\dfrac{du}{dt} = \dfrac{2}{7}\,u, \;\; u(0) = \dfrac{1}{3}.$

(b) Suppose you make the approximation $u(0) = .3333$. At what point does your solution differ from the true solution by 1 unit? by 1000 units? (c) Answer the same question if you also approximate the coefficient in the differential equation by $\dfrac{du}{dt} = .2857\,u$.

♢ 8.1.12. Let a be complex. Prove that $u(t) = ce^{at}$ is the (complex) solution to our scalar ordinary differential equation (8.1). Describe the asymptotic behavior of the solution as $t \to \infty$, and the stability properties of the zero equilibrium solution.

♢ 8.1.13. (*a*) Prove that if $u_1(t)$ and $u_2(t)$ are any two distinct solutions to $\dfrac{du}{dt} = a\,u$ with $a > 0$, then $|\,u_1(t) - u_2(t)\,| \to \infty$ as $t \to \infty$. (*b*) If $a = .02$ and $u_1(0) = .1$, $u_2(0) = .05$, how long do you have to wait until $|\,u_1(t) - u_2(t)\,| > 1,000$?

First Order Dynamical Systems

The simplest class of *dynamical system* is a coupled system of n first order ordinary differential equations

$$\frac{du_1}{dt} = f_1(t, u_1, \ldots, u_n), \qquad \ldots \qquad \frac{du_n}{dt} = f_n(t, u_1, \ldots, u_n).$$

The unknowns are the n scalar functions $u_1(t), \ldots, u_n(t)$ depending on the scalar variable $t \in \mathbb{R}$, which we usually view as time, whence the term "dynamical". We will often write the system in the equivalent vector form

$$\frac{d\mathbf{u}}{dt} = \mathbf{f}(t, \mathbf{u}), \tag{8.7}$$

in which $\mathbf{u}(t) = (u_1(t), \ldots, u_n(t))^T$ is the vector-valued solution, which serves to parameterize a curve in $\mathbb{R}^n$, and $\mathbf{f}(t, \mathbf{u})$ is a vector-valued function with components $f_i(t, u_1, \ldots, u_n)$ for $i = 1, \ldots, n$. A dynamical system is called *autonomous* if the time variable t does not appear explicitly on the right-hand side, and so has the form

$$\frac{d\mathbf{u}}{dt} = \mathbf{f}(\mathbf{u}). \tag{8.8}$$

Dynamical systems of ordinary differential equations appear in an astonishing variety of applications, including physics, astronomy, chemistry, biology, weather and climate, economics and finance, and have been intensely studied since the very first days following the invention of calculus.

In this text, we shall concentrate most of our attention on the very simplest case: a homogeneous, linear, autonomous dynamical system, in which the right-hand side $f(\mathbf{u})$ is a linear function of $\mathbf{u}$ that is independent of the time t, and hence given by multiplication by a constant matrix. Thus, the system takes the form

$$\frac{d\mathbf{u}}{dt} = A\,\mathbf{u}, \tag{8.9}$$

in which A is a constant $n \times n$ matrix. Writing out the system (8.9) in full detail produces

$$\begin{aligned}
\frac{du_1}{dt} &= a_{11}\,u_1 + a_{12}\,u_2 + \;\cdots\; + a_{1n}\,u_n, \\
\frac{du_2}{dt} &= a_{21}\,u_1 + a_{22}\,u_2 + \;\cdots\; + a_{2n}\,u_n, \\
&\;\;\vdots \qquad\qquad\qquad\qquad \vdots \\
\frac{du_n}{dt} &= a_{n1}\,u_1 + a_{n2}\,u_2 + \;\cdots\; + a_{nn}\,u_n,
\end{aligned} \tag{8.10}$$

and we seek the solution $\mathbf{u}(t) = (\, u_1(t), \ldots, u_n(t)\,)^T$. In the autonomous case, which is the only type to be treated in depth here, the coefficients a_{ij} are assumed to be (real) constants. We are interested not only in the formulas for the solutions, but also in understanding their qualitative and quantitative behavior.

Drawing our inspiration from the exponential solution formula (8.2) for the scalar version, let us investigate whether the vector system admits any solutions of a similar exponential form

$$\mathbf{u}(t) = e^{\lambda t}\,\mathbf{v}. \tag{8.11}$$

We assume that λ is a constant scalar, so $e^{\lambda t}$ is the usual scalar exponential function, while $\mathbf{v} \in \mathbb{R}^n$ is a constant vector. In other words, the components $u_i(t) = e^{\lambda t} v_i$ of our desired solution are assumed to be constant multiples of the *same* exponential function. Since $\mathbf{v}$ is constant, the derivative of $\mathbf{u}(t)$ is easily found:

$$\frac{d\mathbf{u}}{dt} = \frac{d}{dt}\left(\, e^{\lambda t}\,\mathbf{v}\,\right) = \lambda\, e^{\lambda t}\,\mathbf{v}.$$

On the other hand, since $e^{\lambda t}$ is a scalar, it commutes with matrix multiplication, and so

$$A\,\mathbf{u} = A\, e^{\lambda t}\,\mathbf{v} = e^{\lambda t} A\,\mathbf{v}.$$

Therefore, $\mathbf{u}(t)$ will solve the system (8.9) if and only if

$$\lambda\, e^{\lambda t}\,\mathbf{v} = e^{\lambda t} A\,\mathbf{v},$$

or, canceling the common scalar factor $e^{\lambda t}$,

$$\lambda\,\mathbf{v} = A\,\mathbf{v}.$$

The result is a system of n algebraic equations relating the vector $\mathbf{v}$ and the scalar λ. Analysis of this system and its ramifications will be the topic of the remainder of this chapter. A broad range of significant applications will appear in the subsequent two chapters.

8.2 Eigenvalues and Eigenvectors

In view of the preceding motivational section, we hereby inaugurate our discussion of eigenvalues and eigenvectors by stating the basic definition.

Definition 8.2. Let A be an $n \times n$ matrix. A scalar λ is called an *eigenvalue* of A if there is a *non-zero* vector $\mathbf{v} \neq \mathbf{0}$, called an *eigenvector*, such that

$$A\,\mathbf{v} = \lambda\,\mathbf{v}. \tag{8.12}$$

In geometric terms, the matrix A has the effect of stretching the eigenvector $\mathbf{v}$ by an amount specified by the eigenvalue λ.

Remark. The odd-looking terms "eigenvalue" and "eigenvector" are hybrid German–English words. In the original German, they are *Eigenwert* and *Eigenvektor*, which can be fully translated as "proper value" and "proper vector". For some reason, the half-translated terms have acquired a certain charm, and are now standard. The alternative English terms *characteristic value* and *characteristic vector* can be found in some (mostly older) texts. Oddly, the terms *characteristic polynomial* and *characteristic equation*, to be defined below, are still used rather than "eigenpolynomial" and "eigenequation".

The requirement that the eigenvector $\mathbf{v}$ be nonzero is important, since $\mathbf{v} = \mathbf{0}$ is a trivial solution to the eigenvalue equation (8.12) for *every* scalar λ. Moreover, as far as solving

linear ordinary differential equations goes, the zero vector $\mathbf{v} = \mathbf{0}$ gives $\mathbf{u}(t) \equiv \mathbf{0}$, which is certainly a solution, but one that we already knew.

The eigenvalue equation (8.12) is a system of linear equations for the entries of the eigenvector $\mathbf{v}$ — provided that the eigenvalue λ is specified in advance — but is "mildly" nonlinear as a combined system for λ and $\mathbf{v}$. Gaussian Elimination per se will not solve the problem, and we are in need of a new idea. Let us begin by rewriting the equation in the form†

$$(A - \lambda\,\mathrm{I}\,)\,\mathbf{v} = \mathbf{0}, \tag{8.13}$$

where I is the identity matrix of the correct size, whereby $\lambda\,\mathrm{I}\,\mathbf{v} = \lambda\mathbf{v}$. Now, for given λ, equation (8.13) is a homogeneous linear system for $\mathbf{v}$, and always has the trivial zero solution $\mathbf{v} = \mathbf{0}$. But we are specifically seeking a nonzero solution! According to Theorem 1.47, a homogeneous linear system has a nonzero solution $\mathbf{v} \neq \mathbf{0}$ if and only if its coefficient matrix, which in this case is $A - \lambda\,\mathrm{I}$, is singular. This observation is the key to resolving the eigenvector equation.

Theorem 8.3. A scalar λ is an eigenvalue of the $n \times n$ matrix A if and only if the matrix $A - \lambda\,\mathrm{I}$ is singular, i.e., of rank $< n$. The corresponding eigenvectors are the nonzero solutions to the eigenvalue equation $(A - \lambda\,\mathrm{I}\,)\,\mathbf{v} = \mathbf{0}$.

We know a number of ways to characterize singular matrices, including the vanishing determinant criterion given in (1.84). Therefore, the following result is immediate:

Corollary 8.4. A scalar λ is an eigenvalue of the matrix A if and only if λ is a solution to the *characteristic equation*

$$\det(A - \lambda\,\mathrm{I}\,) = 0. \tag{8.14}$$

In practice, when computing eigenvalues and eigenvectors by hand using exact arithmetic, one first solves the characteristic equation (8.14) to obtain the set of eigenvalues. Then, for each eigenvalue λ one uses standard linear algebra methods, e.g., Gaussian Elimination, to solve the corresponding linear system (8.13) for the associated eigenvector $\mathbf{v}$.

Example 8.5. Consider the 2×2 matrix

$$A = \begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix}.$$

We compute the determinant in the characteristic equation using formula (1.38):

$$\det(A - \lambda\,\mathrm{I}\,) = \det\begin{pmatrix} 3-\lambda & 1 \\ 1 & 3-\lambda \end{pmatrix} = (3-\lambda)^2 - 1 = \lambda^2 - 6\lambda + 8.$$

Thus, the characteristic equation is a quadratic polynomial equation, and can be solved by factorization:

$$\lambda^2 - 6\lambda + 8 = (\lambda - 4)(\lambda - 2) = 0.$$

We conclude that A has two eigenvalues: $\lambda_1 = 4$ and $\lambda_2 = 2$.

For each eigenvalue, the corresponding eigenvectors are found by solving the associated homogeneous linear system (8.13). For the first eigenvalue, the eigenvector equation is

$$(A - 4\,\mathrm{I}\,)\,\mathbf{v} = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \qquad \text{or} \qquad \begin{aligned} -x + y &= 0, \\ x - y &= 0. \end{aligned}$$

† Note that it is not legal to write (8.13) in the form $(A - \lambda)\mathbf{v} = \mathbf{0}$ since we do not know how to subtract a scalar λ from a matrix A. Worse, if you type $A - \lambda$ in MATLAB or MATHEMATICA, the result will be to subtract λ from *all* the entries of A, which is *not* what we are after!

The general solution is

$$x = y = a, \qquad \text{so} \qquad \mathbf{v} = \begin{pmatrix} a \\ a \end{pmatrix} = a \begin{pmatrix} 1 \\ 1 \end{pmatrix},$$

where a is an arbitrary scalar. Only the nonzero solutions† count as eigenvectors, and so the eigenvectors for the eigenvalue $\lambda_1 = 4$ must have $a \neq 0$, i.e., they are all nonzero scalar multiples of the basic eigenvector $\mathbf{v}_1 = (\,1, 1\,)^T$.

Remark. In general, if $\mathbf{v}$ is an eigenvector of A for the eigenvalue λ, then so is every nonzero scalar multiple of $\mathbf{v}$. In practice, we distinguish only linearly independent eigenvectors. Thus, in this example, we shall say "$\mathbf{v}_1 = (\,1, 1\,)^T$ is *the* eigenvector corresponding to the eigenvalue $\lambda_1 = 4$", when we really mean that the set of eigenvectors for $\lambda_1 = 4$ consists of all nonzero scalar multiples of $\mathbf{v}_1$.

Similarly, for the second eigenvalue $\lambda_2 = 2$, the eigenvector equation is

$$(A - 2\,\mathrm{I}\,)\,\mathbf{v} = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

The solution $(\,-a, a\,)^T = a\,(\,-1, 1\,)^T$ is the set of scalar multiples of the eigenvector $\mathbf{v}_2 = (\,-1, 1\,)^T$. Therefore, the complete list of eigenvalues and eigenvectors (up to scalar multiple) for this particular matrix is

$$\lambda_1 = 4, \qquad \mathbf{v}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \qquad\qquad \lambda_2 = 2, \qquad \mathbf{v}_2 = \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

Example 8.6. Consider the 3×3 matrix

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

Using the formula (1.88) for a 3×3 determinant, we compute the characteristic equation

$$\begin{aligned} 0 = \det(A - \lambda\,\mathrm{I}\,) &= \det \begin{pmatrix} -\lambda & -1 & -1 \\ 1 & 2-\lambda & 1 \\ 1 & 1 & 2-\lambda \end{pmatrix} \\ &= (-\lambda)(2-\lambda)^2 + (-1)\cdot 1 \cdot 1 + (-1)\cdot 1 \cdot 1 \\ &\qquad - 1 \cdot (2-\lambda)(-1) - 1 \cdot 1 \cdot (-\lambda) - (2-\lambda)\cdot 1 \cdot (-1) = -\lambda^3 + 4\lambda^2 - 5\lambda + 2. \end{aligned}$$

The resulting cubic polynomial can be factored:

$$-\lambda^3 + 4\lambda^2 - 5\lambda + 2 = -(\lambda - 1)^2\,(\lambda - 2) = 0.$$

---

† If, at this stage, you end up with a linear system with only the trivial zero solution, you've done something wrong! Either you don't have a correct eigenvalue — maybe you made a mistake setting up and/or solving the characteristic equation — or you've made an error solving the homogeneous eigenvector system. On the other hand, if you are working with a numerical approximation to the eigenvalue, then the resulting numerical homogeneous linear system will almost certainly not have a nonzero solution, and therefore a completely different approach must be taken to finding the corresponding eigenvector; see Sections 9.5 and 9.6.

Most 3×3 matrices have three different eigenvalues, but this particular one has only two: $\lambda_1 = 1$, which is called a *double eigenvalue*, since it is a double root of the characteristic equation, along with a *simple eigenvalue* $\lambda_2 = 2$.

The eigenvector equation (8.13) for the double eigenvalue $\lambda_1 = 1$ is

$$(A - \mathrm{I})\,\mathbf{v} = \begin{pmatrix} -1 & -1 & -1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

The general solution to this homogeneous linear system

$$\mathbf{v} = \begin{pmatrix} -a-b \\ a \\ b \end{pmatrix} = a \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix} + b \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix}$$

depends upon two free variables: $y = a$ and $z = b$. Every nonzero solution forms a valid eigenvector for the eigenvalue $\lambda_1 = 1$, and so the general eigenvector is any non-zero linear combination of the two "basis eigenvectors" $\mathbf{v}_1 = (\,-1, 1, 0\,)^T$, $\widehat{\mathbf{v}}_1 = (\,-1, 0, 1\,)^T$.

On the other hand, the eigenvector equation for the simple eigenvalue $\lambda_2 = 2$ is

$$(A - 2\,\mathrm{I})\,\mathbf{v} = \begin{pmatrix} -2 & -1 & -1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

The general solution

$$\mathbf{v} = \begin{pmatrix} -a \\ a \\ a \end{pmatrix} = a \begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix}$$

consists of all scalar multiples of the eigenvector $\mathbf{v}_2 = (\,-1, 1, 1\,)^T$.

In summary, the eigenvalues and (basis) eigenvectors for this matrix are

$$\begin{aligned} \lambda_1 &= 1, \qquad \mathbf{v}_1 = \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}, \qquad \widehat{\mathbf{v}}_1 = \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix}, \\ \lambda_2 &= 2, \qquad \mathbf{v}_2 = \begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix}. \end{aligned} \tag{8.15}$$

This means that every eigenvector for the simple eigenvalue $\lambda_2 = 2$ is a nonzero scalar multiple of $\mathbf{v}_2$, while every eigenvector for the double eigenvalue $\lambda_1 = 1$ is a nontrivial linear combination $a\,\mathbf{v}_1 + b\,\widehat{\mathbf{v}}_1$ of the two linearly independent eigenvectors $\mathbf{v}_1, \widehat{\mathbf{v}}_1$.

In general, given a real eigenvalue λ, the corresponding *eigenspace* $V_\lambda \subset \mathbb{R}^n$ is the subspace spanned by all its eigenvectors. Equivalently, the eigenspace is the kernel

$$V_\lambda = \ker(A - \lambda\,\mathrm{I}). \tag{8.16}$$

Thus, $\lambda \in \mathbb{R}$ is an eigenvalue if and only if $V_\lambda \neq \{\mathbf{0}\}$ is a nontrivial subspace, and then every nonzero element of V_λ is a corresponding eigenvector. The most economical way to indicate each eigenspace is by writing out a basis, as in (8.15), with $\mathbf{v}_1, \widehat{\mathbf{v}}_1$ giving a basis for the eigenspace V_1, while $\mathbf{v}_2$ is a basis for the eigenspace V_2. In particular, 0 is an eigenvalue if and only if $\ker A \neq \{\mathbf{0}\}$, and hence A is singular.

Proposition 8.7. A matrix is singular if and only if it has a zero eigenvalue.

Example 8.8. The characteristic equation of the matrix $A = \begin{pmatrix} 1 & 2 & 1 \\ 1 & -1 & 1 \\ 2 & 0 & 1 \end{pmatrix}$ is

$$0 = \det(A - \lambda\,\mathrm{I}) = -\lambda^3 + \lambda^2 + 5\lambda + 3 = -(\lambda+1)^2(\lambda-3).$$

Again, there is a double eigenvalue $\lambda_1 = -1$ and a simple eigenvalue $\lambda_2 = 3$. However, in this case the matrix

$$A - \lambda_1\,\mathrm{I} = A + \mathrm{I} = \begin{pmatrix} 2 & 2 & 1 \\ 1 & 0 & 1 \\ 2 & 0 & 2 \end{pmatrix}$$

has a one-dimensional kernel, spanned by $\mathbf{v}_1 = (2, -1, -2)^T$. Thus, even though λ_1 is a double eigenvalue, it admits only a one-dimensional eigenspace. The list of eigenvalues and eigenvectors is, in a sense, incomplete:

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

Example 8.9. Finally, the matrix $A = \begin{pmatrix} 1 & 2 & 0 \\ 0 & 1 & -2 \\ 2 & 2 & -1 \end{pmatrix}$ has characteristic equation

$$0 = \det(A - \lambda\,\mathrm{I}) = -\lambda^3 + \lambda^2 - 3\lambda - 5 = -(\lambda+1)(\lambda^2 - 2\lambda + 5).$$

The linear factor yields the eigenvalue -1. The quadratic factor leads to two complex roots, $1 + 2\,\mathrm{i}$ and $1 - 2\,\mathrm{i}$, which can be obtained via the quadratic formula. Hence A has one real and two complex eigenvalues:

$$\lambda_1 = -1, \qquad \lambda_2 = 1 + 2\,\mathrm{i}, \qquad \lambda_3 = 1 - 2\,\mathrm{i}.$$

On solving the associated linear system $(A + \mathrm{I})\,\mathbf{v} = \mathbf{0}$, the real eigenvalue $\lambda_1 = -1$ is found to have corresponding eigenvector $\mathbf{v}_1 = (-1, 1, 1)^T$.

Complex eigenvalues are as important as real eigenvalues, and we need to be able to handle them too. To find the corresponding eigenvectors, which will also be complex, we need to solve the usual eigenvalue equation (8.13), which is now a complex homogeneous linear system. For example, the eigenvectors for $\lambda_2 = 1 + 2\,\mathrm{i}$ are found by solving

$$\big[\,A - (1 + 2\,\mathrm{i})\,\mathrm{I}\,\big]\mathbf{v} = \begin{pmatrix} -2\,\mathrm{i} & 2 & 0 \\ 0 & -2\,\mathrm{i} & -2 \\ 2 & 2 & -2 - 2\,\mathrm{i} \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

This linear system can be solved by Gaussian Elimination (with complex pivots). In this case, a simpler strategy is to work directly: the first equation $-2\,\mathrm{i}\,x + 2\,y = 0$ tells us that $y = \mathrm{i}\,x$, while the second equation $-2\,\mathrm{i}\,y - 2\,z = 0$ says $z = -\mathrm{i}\,y = x$. If we trust our calculations so far, we do not need to solve the final equation $2\,x + 2\,y + (-2 - 2\,\mathrm{i})z = 0$, since we know that the coefficient matrix is singular and hence this equation must be a consequence of the first two. (However, it does serve as a useful check on our work.) So, the general solution $\mathbf{v} = (x, \mathrm{i}\,x, x)^T$ is an arbitrary constant multiple of the complex eigenvector $\mathbf{v}_2 = (1, \mathrm{i}, 1)^T$. The eigenvector equation for $\lambda_3 = 1 - 2\,\mathrm{i}$ is similarly solved for the third eigenvector $\mathbf{v}_3 = (1, -\mathrm{i}, 1)^T$.

Summarizing, the matrix under consideration has one real and two complex eigenvalues, with three corresponding eigenvectors, each unique up to (complex) scalar multiple:

$$\lambda_1 = -1, \qquad \lambda_2 = 1 + 2\,\mathrm{i}\,, \qquad \lambda_3 = 1 - 2\,\mathrm{i}\,,$$

$$\mathbf{v}_1 = \begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix}, \qquad \mathbf{v}_2 = \begin{pmatrix} 1 \\ \mathrm{i} \\ 1 \end{pmatrix}, \qquad \mathbf{v}_3 = \begin{pmatrix} 1 \\ -\mathrm{i} \\ 1 \end{pmatrix}.$$

Note that the third complex eigenvalue is the complex conjugate of the second, and the eigenvectors are similarly related. This is indicative of a general fact for real matrices:

Proposition 8.10. If A is a real matrix with a complex eigenvalue $\lambda = \mu + \mathrm{i}\,\nu$ and corresponding complex eigenvector $\mathbf{v} = \mathbf{x} + \mathrm{i}\,\mathbf{y}$, then the complex conjugate $\overline{\lambda} = \mu - \mathrm{i}\,\nu$ is also an eigenvalue with complex conjugate eigenvector $\overline{\mathbf{v}} = \mathbf{x} - \mathrm{i}\,\mathbf{y}$.

Proof: First take complex conjugates of the eigenvalue equation (8.12):

$$\overline{A}\,\overline{\mathbf{v}} = \overline{A\,\mathbf{v}} = \overline{\lambda\,\mathbf{v}} = \overline{\lambda}\,\overline{\mathbf{v}}.$$

Using the fact that a real matrix is unaffected by complex conjugation, $\overline{A} = A$, we conclude $A\,\overline{\mathbf{v}} = \overline{\lambda}\,\overline{\mathbf{v}}$, which is the equation for the eigenvalue $\overline{\lambda}$ and eigenvector $\overline{\mathbf{v}} \neq \mathbf{0}$. *Q.E.D.*

As a consequence, when dealing with real matrices, we need to compute the eigenvectors for only *one* of each complex conjugate pair of eigenvalues. This observation effectively halves the amount of work in the unfortunate event that we are confronted with complex eigenvalues.

The *eigenspace* associated with a complex eigenvalue λ is the subspace $V_\lambda \subset \mathbb{C}^n$ spanned by the corresponding (complex) eigenvectors. One might also consider complex eigenvectors associated with a real eigenvalue, but this doesn't add anything to the picture — they are merely complex linear combinations of the real eigenvectors. Thus, we need to introduce complex eigenvectors only when dealing with genuinely complex eigenvalues.

Remark. The reader may recall that we said that one should never use determinants in practical computations. So why have we reverted to using determinants to find eigenvalues? The truthful answer is that the practical computation of eigenvalues and eigenvectors *never* resorts to the characteristic equation! The method is fraught with numerical traps and inefficiencies when (a) computing the determinant leading to the characteristic equation, then (b) solving the resulting polynomial equation, which is itself a nontrivial numerical problem†, [**8**, **66**], and, finally, (c) solving each of the resulting linear eigenvector systems. Even worse, if we know only an approximation $\widetilde{\lambda}$ to the true eigenvalue λ, the approximate eigenvector system $(A - \widetilde{\lambda}\,\mathrm{I}\,)\,\mathbf{v} = \mathbf{0}$ will almost certainly have a nonsingular coefficient matrix, and hence admits only the trivial solution $\mathbf{v} = \mathbf{0}$ — which does not even qualify as an eigenvector!

Nevertheless, the characteristic equation does give us important theoretical insight into the structure of the eigenvalues of a matrix, and can be used when dealing with very small matrices, e.g., 2×2 and 3×3, presuming exact arithmetic is employed. Numerical

† In fact, one effective numerical strategy for finding the roots of a polynomial is to turn the procedure on its head, and calculate the eigenvalues of a matrix whose characteristic equation is the polynomial in question! See [**66**] for details.

algorithms for computing eigenvalues and eigenvectors are based on completely different ideas, and will be deferred until Sections 9.5 and 9.6.

Exercises

8.2.1. Find the eigenvalues and eigenvectors of the following matrices:

(a) $\begin{pmatrix} 1 & -2 \\ -2 & 1 \end{pmatrix}$, (b) $\begin{pmatrix} 1 & -\frac{2}{3} \\ \frac{1}{2} & -\frac{1}{6} \end{pmatrix}$, (c) $\begin{pmatrix} 3 & 1 \\ -1 & 1 \end{pmatrix}$, (d) $\begin{pmatrix} 1 & 2 \\ -1 & 1 \end{pmatrix}$,

(e) $\begin{pmatrix} 3 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 3 \end{pmatrix}$, (f) $\begin{pmatrix} -1 & -1 & 4 \\ 1 & 3 & -2 \\ 1 & 1 & -1 \end{pmatrix}$, (g) $\begin{pmatrix} 1 & -3 & 11 \\ 2 & -6 & 16 \\ 1 & -3 & 7 \end{pmatrix}$, (h) $\begin{pmatrix} 2 & -1 & -1 \\ -2 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix}$,

(i) $\begin{pmatrix} -4 & -4 & 2 \\ 3 & 4 & -1 \\ -3 & -2 & 3 \end{pmatrix}$, (j) $\begin{pmatrix} 3 & 4 & 0 & 0 \\ 4 & 3 & 0 & 0 \\ 0 & 0 & 1 & 3 \\ 0 & 0 & 4 & 5 \end{pmatrix}$, (k) $\begin{pmatrix} 4 & 0 & 0 & 0 \\ 1 & 3 & 0 & 0 \\ -1 & 1 & 2 & 0 \\ 1 & -1 & 1 & 1 \end{pmatrix}$.

8.2.2. (a) Find the eigenvalues of the rotation matrix $R_\theta = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}$. For what values of θ are the eigenvalues real?

(b) Explain why your answer gives an immediate solution to Exercise 1.5.7c.

8.2.3. Answer Exercise 8.2.2a for the reflection matrix $F_\theta = \begin{pmatrix} \cos\theta & \sin\theta \\ \sin\theta & -\cos\theta \end{pmatrix}$.

8.2.4. Write down (a) a 2×2 matrix that has 0 as one of its eigenvalues and $(1, 2)^T$ as a corresponding eigenvector; (b) a 3×3 matrix that has $(1, 2, 3)^T$ as an eigenvector for the eigenvalue -1.

8.2.5. (a) Write out the characteristic equation for the matrix $\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ \alpha & \beta & \gamma \end{pmatrix}$.

(b) Show that, given any 3 numbers a, b, and c, there is a 3×3 matrix with characteristic equation $-\lambda^3 + a\lambda^2 + b\lambda + c = 0$.

8.2.6. Find the eigenvalues and eigenvectors of the cross product matrix $A = \begin{pmatrix} 0 & c & -b \\ -c & 0 & a \\ b & -a & 0 \end{pmatrix}$.

8.2.7. Find all eigenvalues and eigenvectors of the following complex matrices:

(a) $\begin{pmatrix} \mathrm{i} & 1 \\ 0 & -1+\mathrm{i} \end{pmatrix}$, (b) $\begin{pmatrix} 2 & \mathrm{i} \\ -\mathrm{i} & -2 \end{pmatrix}$, (c) $\begin{pmatrix} \mathrm{i}-2 & \mathrm{i}+1 \\ \mathrm{i}+2 & \mathrm{i}-1 \end{pmatrix}$, (d) $\begin{pmatrix} 1+\mathrm{i} & -1-\mathrm{i} & 1-\mathrm{i} \\ 2 & -2-\mathrm{i} & 2-2\mathrm{i} \\ -1 & 1+\mathrm{i} & -1+2\mathrm{i} \end{pmatrix}$.

8.2.8. Find all eigenvalues and eigenvectors of

(a) the $n \times n$ zero matrix O; (b) the $n \times n$ identity matrix I.

8.2.9. Find the eigenvalues and eigenvectors of an $n \times n$ matrix with every entry equal to 1. *Hint*: Try with $n = 2, 3$, and then generalize.

♢ 8.2.10. Let A be a given square matrix. (a) Explain in detail why every nonzero scalar multiple of an eigenvector of A is also an eigenvector. (b) Show that every nonzero linear combination of two eigenvectors $\mathbf{v}, \mathbf{w}$ corresponding to the *same* eigenvalue is also an eigenvector. (c) Prove that a linear combination $c\mathbf{v} + d\mathbf{w}$, with $c, d \neq 0$, of two eigenvectors corresponding to *different* eigenvalues is never an eigenvector.

♢ 8.2.11. Let λ be a real eigenvalue of the real $n \times n$ matrix A, and $\mathbf{v}_1, \ldots, \mathbf{v}_k$ a basis for the associated eigenspace V_λ. Suppose $\mathbf{w} \in \mathbb{C}^n$ is a complex eigenvector, so $A\mathbf{w} = \lambda\mathbf{w}$. Prove that $\mathbf{w} = c_1\mathbf{v}_1 + \cdots + c_k\mathbf{v}_k$ is a complex linear combination of the real eigenspace basis. *Hint*: Look at the real and imaginary parts of the eigenvector equation.

8.2.12. *True or false*: If $\mathbf{v}$ is a real eigenvector of a real matrix A, then a nonzero complex multiple $\mathbf{w} = c\,\mathbf{v}$ for $c \in \mathbb{C}$ is a complex eigenvector of A.

♢ 8.2.13. Define the *shift map* $S\colon \mathbb{C}^n \to \mathbb{C}^n$ by $S(v_1, v_2, \ldots, v_{n-1}, v_n)^T = (v_2, v_3, \ldots, v_n, v_1)^T$.
(*a*) Prove that S is a linear map, and write down its matrix representation A.
(*b*) Prove that A is an orthogonal matrix. (*c*) Prove that the sampled exponential vectors $\boldsymbol{\omega}_0, \ldots, \boldsymbol{\omega}_{n-1}$ defined in (5.102) form an eigenvector basis of A. What are the eigenvalues?

Basic Properties of Eigenvalues

If A is an $n \times n$ matrix, then its *characteristic polynomial* is defined to be

$$p_A(\lambda) = \det(A - \lambda\,\mathrm{I}) = c_n\,\lambda^n + c_{n-1}\,\lambda^{n-1} + \cdots + c_1\,\lambda + c_0. \tag{8.17}$$

The fact that $p_A(\lambda)$ is a polynomial of degree n is a consequence of the general determinantal formula (1.87). Indeed, every term is prescribed by a permutation π of the rows of the matrix, and equals plus or minus a product of n distinct matrix entries including one from each row and one from each column. The term corresponding to the identity permutation is obtained by multiplying the diagonal entries together, which, in this case, is

$$(a_{11} - \lambda)\,(a_{22} - \lambda) \cdots (a_{nn} - \lambda) = (-1)^n \lambda^n + (-1)^{n-1}\big(a_{11} + a_{22} + \cdots + a_{nn}\big)\,\lambda^{n-1} + \cdots. \tag{8.18}$$

All of the other terms have at most $n-2$ diagonal factors $a_{ii} - \lambda$, and so are polynomials of degree $\leq n-2$ in λ. Thus, (8.18) is the only summand containing the monomials λ^n and λ^{n-1}, and so their respective coefficients are

$$c_n = (-1)^n, \qquad c_{n-1} = (-1)^{n-1}(a_{11} + a_{22} + \cdots + a_{nn}) = (-1)^{n-1}\,\mathrm{tr}\,A, \tag{8.19}$$

where $\mathrm{tr}\,A$, the sum of its diagonal entries, is called the *trace* of the matrix A. The other coefficients $c_{n-2}, \ldots, c_1, c_0$ in (8.17) are more complicated combinations of the entries of A. However, setting $\lambda = 0$ implies

$$p_A(0) = c_0 = \det A, \tag{8.20}$$

and hence the constant term in the characteristic polynomial equals the determinant of the matrix. In particular, if $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ is a 2×2 matrix, its characteristic polynomial has the explicit form

$$\begin{aligned} p_A(\lambda) = \det(A - \lambda\,\mathrm{I}) &= \det \begin{pmatrix} a - \lambda & b \\ c & d - \lambda \end{pmatrix} \\ &= \lambda^2 - (a+d)\lambda + (a\,d - b\,c) = \lambda^2 - (\mathrm{tr}\,A)\lambda + (\det A). \end{aligned} \tag{8.21}$$

According to the Fundamental Theorem of Algebra, [**26**], every (complex) polynomial of degree $n \geq 1$ can be completely factored, and so we can write the characteristic polynomial (8.17) in factored form:

$$p_A(\lambda) = (-1)^n(\lambda - \lambda_1)\,(\lambda - \lambda_2)\,\cdots\,(\lambda - \lambda_n). \tag{8.22}$$

The complex numbers $\lambda_1, \ldots, \lambda_n$, some of which may be repeated, are the *roots* of the characteristic equation $p_A(\lambda) = 0$, and hence the eigenvalues of the matrix A. Therefore, we immediately conclude:

Theorem 8.11. An $n \times n$ matrix possesses at least one and at most n distinct complex eigenvalues.

Most $n \times n$ matrices — meaning those for which the characteristic polynomial factors into n *distinct* factors — have *exactly* n complex eigenvalues. More generally, an eigenvalue λ_j is said to have *multiplicity* k if the factor $(\lambda - \lambda_j)$ appears exactly k times in the factorization (8.22) of the characteristic polynomial. An eigenvalue is *simple* if it has multiplicity 1, *double* if it has multiplicity 2, and so on. In particular, A has n distinct eigenvalues if and only if all its eigenvalues are simple. In all cases, when the repeated eigenvalues are counted in accordance with their multiplicity, every $n \times n$ matrix has a total of n complex eigenvalues.

An example of a matrix with just one eigenvalue, of multiplicity n, is the $n \times n$ identity matrix I, whose only eigenvalue is $\lambda = 1$. In this case, *every* nonzero vector in $\mathbb{R}^n$ is an eigenvector of the identity matrix (why?), and so the eigenspace V_1 is all of $\mathbb{R}^n$. At the other extreme, the $n \times n$ "bidiagonal" *Jordan block matrix*†

$$J_{a,n} = \begin{pmatrix} a & 1 & & & & \\ & a & 1 & & & \\ & & a & 1 & & \\ & & & \ddots & \ddots & \\ & & & & a & 1 \\ & & & & & a \end{pmatrix} \tag{8.23}$$

also has only one eigenvalue, $\lambda = a$, again of multiplicity n. But in this case, $J_{a,n}$ has only one eigenvector (up to scalar multiple), which is the first standard basis vector $\mathbf{e}_1$, and so its eigenspace is one-dimensional.

Remark. If λ is a complex eigenvalue of multiplicity k for the real matrix A, then its complex conjugate $\overline{\lambda}$ also has multiplicity k. This is because complex conjugate roots of a real polynomial necessarily appear with identical multiplicities.

Remark. If $n \leq 4$, then one can, in fact, write down an explicit formula for the solution to a polynomial equation of degree n, and hence explicit (but rather complicated and not particularly helpful) formulas for the eigenvalues of general 2×2, 3×3, and 4×4 matrices. As soon as $n \geq 5$, there is no explicit formula (at least in terms of radicals), and so one must usually resort to numerical approximations. This remarkable and deep algebraic result was proved in the early nineteenth century by the young Norwegian mathematician Niels Henrik Abel, [**26**].

Proposition 8.12. A square matrix A and its transpose A^T have the same characteristic equation, and hence the same eigenvalues with the same multiplicities.

Proof: This follows immediately from Proposition 1.56, that the determinant of a matrix and its transpose are identical. Thus,

$$p_A(\lambda) = \det(A - \lambda\,\mathrm{I}) = \det(A - \lambda\,\mathrm{I})^T = \det(A^T - \lambda\,\mathrm{I}) = p_{A^T}(\lambda). \qquad Q.E.D.$$

Remark. While A^T has the same eigenvalues as A, its eigenvectors are, in general, different. An eigenvector $\mathbf{v}$ of A^T, satisfying $A^T\mathbf{v} = \lambda\mathbf{v}$, is sometimes referred to as a *left eigenvector* of A, since it satisfies $\mathbf{v}^T A = \lambda\mathbf{v}^T$. A more apt, albeit rather non-conventional, name for $\mathbf{v}$ that conforms with our nomenclature conventions would be *co-eigenvector*.

† All non-displayed entries are zero.

If we explicitly multiply out the factored product (8.22) and equate the result to the characteristic polynomial (8.17), we find that its coefficients $c_0, c_1, \ldots c_{n-1}$ can be written as certain polynomials of the roots, known as the *elementary symmetric polynomials*. The first and last are of particular importance:

$$c_0 = \lambda_1 \lambda_2 \cdots \lambda_n, \qquad c_{n-1} = (-1)^{n-1} (\lambda_1 + \lambda_2 + \cdots + \lambda_n). \tag{8.24}$$

Comparison with our previous formulas (8.19, 20) for the coefficients c_0 and c_{n-1} leads to the following useful result.

Proposition 8.13. The sum of the eigenvalues of a square matrix A equals its trace:

$$\lambda_1 + \lambda_2 + \cdots + \lambda_n = \operatorname{tr} A = a_{11} + a_{22} + \cdots + a_{nn}. \tag{8.25}$$

The product of the eigenvalues equals its determinant:

$$\lambda_1 \lambda_2 \cdots \lambda_n = \det A. \tag{8.26}$$

Keep in mind that, in evaluating (8.25, 26), one must add or multiply repeated eigenvalues according to their multiplicity.

Example 8.14. The matrix $A = \begin{pmatrix} 1 & 2 & 1 \\ 1 & -1 & 1 \\ 2 & 0 & 1 \end{pmatrix}$ considered in Example 8.8 has trace and determinant

$$\operatorname{tr} A = 1, \qquad \det A = 3,$$

which fix, respectively, the coefficient of λ^2 and the constant term in its characteristic equation. This matrix has two distinct eigenvalues: -1, which is a double eigenvalue, and 3, which is simple. For this particular matrix, (8.25, 26) become

$$1 = \operatorname{tr} A = (-1) + (-1) + 3, \qquad 3 = \det A = (-1)(-1)\,3.$$

Note that the double eigenvalue contributes twice to both the sum and the product.

Exercises

8.2.14. (*a*) Compute the eigenvalues and corresponding eigenvectors of $A = \begin{pmatrix} 1 & 4 & 4 \\ 3 & -1 & 0 \\ 0 & 2 & 3 \end{pmatrix}$.

(*b*) Compute the trace of A and check that it equals the sum of the eigenvalues. (*c*) Find the determinant of A and check that it is equal to to the product of the eigenvalues.

8.2.15. Verify the trace and determinant formulas (8.25, 26) for the matrices in Exercise 8.2.1.

8.2.16. (*a*) Find the explicit formula for the characteristic polynomial $\det(A - \lambda \mathrm{I}) = -\lambda^3 + a\lambda^2 - b\lambda + c$ of a general 3×3 matrix. Verify that $a = \operatorname{tr} A$, $c = \det A$. What is the formula for b? (*b*) Prove that if A has eigenvalues $\lambda_1, \lambda_2, \lambda_3$, then $a = \operatorname{tr} A = \lambda_1 + \lambda_2 + \lambda_3$, $b = \lambda_1 \lambda_2 + \lambda_1 \lambda_3 + \lambda_2 \lambda_3$, $c = \det A = \lambda_1 \lambda_2 \lambda_3$.

8.2.17. Prove that the eigenvalues of an upper triangular (or lower triangular) matrix are its diagonal entries.

♢ 8.2.18. Let $J_{a,n}$ be the $n \times n$ Jordan block matrix (8.23). Prove that its only eigenvalue is $\lambda = a$ and the only eigenvectors are the nonzero scalar multiples of the standard basis vector $\mathbf{e}_1$.

♢ 8.2.19. Suppose that λ is an eigenvalue of A. (*a*) Prove that $c\lambda$ is an eigenvalue of the scalar multiple cA. (*b*) Prove that $\lambda + d$ is an eigenvalue of $A + d\,\mathrm{I}$. (*c*) More generally, $c\lambda + d$ is an eigenvalue of $B = cA + d\,\mathrm{I}$ for scalars c, d.

8.2.20. Show that if λ is an eigenvalue of A, then λ^2 is an eigenvalue of A^2.

8.2.21. *True or false*: (*a*) If λ is an eigenvalue of both A and B, then it is an eigenvalue of the sum $A + B$. (*b*) If $\mathbf{v}$ is an eigenvector of both A and B, then it is an eigenvector of $A + B$.

8.2.22. *True or false*: If λ is an eigenvalue of A and μ is an eigenvalue of B, then $\lambda\mu$ is an eigenvalue of the matrix product $C = AB$.

♢ 8.2.23. Let A and B be $n \times n$ matrices. Prove that the matrix products AB and BA have the same eigenvalues. *Hint*: How should the eigenvectors be related?

♢ 8.2.24. (*a*) Prove that if $\lambda \neq 0$ is a nonzero eigenvalue of the nonsingular matrix A, then $1/\lambda$ is an eigenvalue of A^{-1}. (*b*) What happens if A has 0 as an eigenvalue?

♢ 8.2.25. (*a*) Prove that if $|\det A| > 1$, then A has at least one eigenvalue with $|\lambda| > 1$.
(*b*) If $|\det A| < 1$, are all eigenvalues $|\lambda| < 1$? Prove or find a counterexample.

8.2.26. Prove that A is a singular matrix if and only if 0 is an eigenvalue.

8.2.27. Prove that *every* nonzero vector $\mathbf{0} \neq \mathbf{v} \in \mathbb{R}^n$ is an eigenvector of A if and only if A is a scalar multiple of the identity matrix.

8.2.28. How many unit (norm 1) eigenvectors correspond to a given eigenvalue of a matrix?

8.2.29. *True or false*: (*a*) Performing an elementary row operation of type #1 does not change the eigenvalues of a matrix. (*b*) Interchanging two rows of a matrix changes the sign of its eigenvalues. (*c*) Multiplying one row of a matrix by a scalar multiplies one of its eigenvalues by the same scalar.

8.2.30. (*a*) *True or false*: If $\lambda_1, \mathbf{v}_1$ and $\lambda_2, \mathbf{v}_2$ solve the eigenvalue equation (8.12) for a given matrix A, so does $\lambda_1 + \lambda_2$, $\mathbf{v}_1 + \mathbf{v}_2$. (*b*) Explain what this has to do with linearity.

8.2.31. As in (4.35), an elementary reflection matrix has the form $Q = \mathrm{I} - 2\mathbf{u}\mathbf{u}^T$, where $\mathbf{u} \in \mathbb{R}^n$ is a unit vector. (*a*) Find the eigenvalues and eigenvectors of the elementary reflection matrices for the unit vectors (*i*) $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, (*ii*) $\begin{pmatrix} \frac{3}{5} \\ \frac{4}{5} \end{pmatrix}$, (*iii*) $\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$, (*iv*) $\begin{pmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ -\frac{1}{\sqrt{2}} \end{pmatrix}$.
(*b*) What are the eigenvalues and eigenvectors of a general elementary reflection matrix?

♢ 8.2.32. Let A and B be similar matrices, so $B = S^{-1}AS$ for some nonsingular matrix S.
(*a*) Prove that A and B have the same characteristic polynomial: $p_B(\lambda) = p_A(\lambda)$.
(*b*) Explain why similar matrices have the same eigenvalues. (*c*) Do they have the same eigenvectors? If not, how are their eigenvectors related? (*d*) Prove that the converse to part (*c*) is false by showing that $\begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}$ and $\begin{pmatrix} 1 & 1 \\ -1 & 3 \end{pmatrix}$ have the same eigenvalues, but are *not* similar.

8.2.33. Let A be a nonsingular $n \times n$ matrix with characteristic polynomial $p_A(\lambda)$.
(*a*) Explain how to construct the characteristic polynomial $p_{A^{-1}}(\lambda)$ of its inverse directly from $p_A(\lambda)$. (*b*) Check your result when $A =$ (*i*) $\begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$, (*ii*) $\begin{pmatrix} 1 & 4 & 4 \\ -2 & -1 & 0 \\ 0 & 2 & 3 \end{pmatrix}$.

♢ 8.2.34. Prove that the only eigenvalue of a nilpotent matrix, cf. Exercise 1.3.12, is 0. (The converse is also true; see Exercise 8.6.20.)

8.2.35. Given an idempotent matrix, so that $P = P^2$, find all its eigenvalues and eigenvectors.

8.2.36. (*a*) Prove that every real 3×3 matrix has at least one real eigenvalue.
(*b*) Find a real 4×4 matrix with no real eigenvalues.
(*c*) Can you find a real 5×5 matrix with no real eigenvalues?

8.2.37. (*a*) Show that if A is a matrix such that $A^4 = \mathrm{I}$, then the only possible eigenvalues of A are $1, -1, \mathrm{i}$, and $-\mathrm{i}$.
(*b*) Give an example of a real matrix that has all four numbers as eigenvalues.

8.2.38. (*a*) Prove that if λ is an eigenvalue of A, then λ^n is an eigenvalue of A^n. (*b*) State and prove a converse.

8.2.39. *True or false*: All the eigenvalues of an $n \times n$ permutation matrix are real.

8.2.40. (*a*) Show that if all the row sums of A are equal to 1, then A has 1 as an eigenvalue.
(*b*) Suppose all the column sums of A are equal to 1. Does the same result hold?

♢ 8.2.41. Prove that if $\mathbf{v}$ is an eigenvector of A with eigenvalue λ and $\mathbf{w}$ is an eigenvector of A^T with a different eigenvalue $\mu \neq \lambda$, then $\mathbf{v}$ and $\mathbf{w}$ are orthogonal vectors with respect to the dot product. Illustrate this result when (*i*) $A = \begin{pmatrix} 0 & -1 \\ 2 & 3 \end{pmatrix}$, (*ii*) $A = \begin{pmatrix} 5 & -4 & 2 \\ 5 & -4 & 1 \\ -2 & 2 & -3 \end{pmatrix}$.

8.2.42. Let Q be an orthogonal matrix. (*a*) Prove that if λ is an eigenvalue, then so is $1/\lambda$.
(*b*) Prove that all its eigenvalues are complex numbers of modulus $|\lambda| = 1$. In particular, the only possible real eigenvalues of an orthogonal matrix are ± 1. (*c*) Suppose $\mathbf{v} = \mathbf{x} + \mathrm{i}\,\mathbf{y}$ is a complex eigenvector corresponding to a non-real eigenvalue. Prove that its real and imaginary parts are orthogonal vectors having the same Euclidean norm.

♢ 8.2.43. (*a*) Prove that every 3×3 proper orthogonal matrix has $+1$ as an eigenvalue.
(*b*) *True or false*: An improper 3×3 orthogonal matrix has -1 as an eigenvalue.

♢ 8.2.44. (*a*) Show that the linear transformation defined by a 3×3 proper orthogonal matrix corresponds to rotating through an angle around a line through the origin in $\mathbb{R}^3$ — the axis of the rotation. *Hint*: Use Exercise 8.2.43(*a*).
(*b*) Find the axis and angle of rotation of the orthogonal matrix $\begin{pmatrix} \frac{3}{5} & 0 & \frac{4}{5} \\ -\frac{4}{13} & \frac{12}{13} & \frac{3}{13} \\ -\frac{48}{65} & -\frac{5}{13} & \frac{36}{65} \end{pmatrix}$.

♢ 8.2.45. Prove that every *proper* affine isometry $F(\mathbf{x}) = Q\,\mathbf{x} + \mathbf{b}$ of $\mathbb{R}^3$, where $\det Q = +1$, is one of the following: (*a*) a *translation* $\mathbf{x} + \mathbf{b}$, (*b*) a *rotation* centered at some point of $\mathbb{R}^3$, or (*c*) a *screw motion* consisting of a rotation around an axis followed by a translation in the direction of the axis. *Hint*: Use Exercise 8.2.44.

8.2.46. Suppose Q is an orthogonal matrix. (*a*) Prove that $K = 2\,\mathrm{I} - Q - Q^T$ is a positive semi-definite matrix. (*b*) Under what conditions is $K > 0$?

♡ 8.2.47. Let M_n be the $n \times n$ tridiagonal matrix whose diagonal entries are all equal to 0 and whose sub- and super-diagonal entries all equal 1. (*a*) Find the eigenvalues and eigenvectors of M_2 and M_3 directly. (*b*) Prove that the eigenvalues and eigenvectors of M_n are explicitly given by

$$\lambda_k = 2\cos\frac{k\pi}{n+1}, \quad \mathbf{v}_k = \left(\sin\frac{k\pi}{n+1}, \ \sin\frac{2k\pi}{n+1}, \ \ldots \ \sin\frac{nk\pi}{n+1}\right)^T, \quad k = 1, \ldots, n.$$

How do you know that there are no other eigenvalues?

♢ 8.2.48. Let $a, b \in \mathbb{R}$. Determine the eigenvalues and eigenvectors of the $n \times n$ tridiagonal matrix with all diagonal entries equal to a and all sub- and super-diagonal entries equal to b. *Hint*: See Exercises 8.2.19 and 8.2.47.

♡ 8.2.49. Find a formula for the eigenvalues of the tricirculant $n \times n$ matrix Z_n that has 1's on the sub- and super-diagonals as well as its $(1, n)$ and $(n, 1)$ entries, while all other entries are 0. *Hint*: Use Exercise 8.2.47 as a guide.

♢ 8.2.50. Let A be an $n \times n$ matrix with eigenvalues $\lambda_1, \ldots, \lambda_k$, and B an $m \times m$ matrix with eigenvalues $\mu_1, \ldots, \mu_l$. Show that the $(m+n)\times(m+n)$ block diagonal matrix $D = \begin{pmatrix} A & \mathrm{O} \\ \mathrm{O} & B \end{pmatrix}$ has eigenvalues $\lambda_1, \ldots, \lambda_k$, $\mu_1, \ldots, \mu_l$ and no others. How are the eigenvectors related?

♡ 8.2.51. *Deflation*: Suppose A has eigenvalue λ and corresponding eigenvector $\mathbf{v}$. (*a*) Let $\mathbf{b}$ be any vector. Prove that the matrix $B = A - \mathbf{v}\mathbf{b}^T$ also has $\mathbf{v}$ as an eigenvector, now with eigenvalue $\lambda - \beta$, where $\beta = \mathbf{v} \cdot \mathbf{b}$. (*b*) Prove that if $\mu \neq \lambda - \beta$ is any other eigenvalue of A, then it is also an eigenvalue of B. *Hint*: Look for an eigenvector of the form $\mathbf{w} + c\mathbf{v}$, where $\mathbf{w}$ is an eigenvector of A. (*c*) Given a nonsingular matrix A with eigenvalues $\lambda_1, \lambda_2, \ldots, \lambda_n$ and $\lambda_1 \neq \lambda_j$ for all $j \geq 2$, explain how to construct a *deflated* matrix B whose eigenvalues are $0, \lambda_2, \ldots, \lambda_n$. (*d*) Try out your method on the matrices $\begin{pmatrix} 3 & 3 \\ 1 & 5 \end{pmatrix}$ and $\begin{pmatrix} 3 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 3 \end{pmatrix}$.

♡ 8.2.52. Let $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ be a 2×2 matrix. (*a*) Prove that A satisfies its own characteristic equation, meaning $p_A(A) = A^2 - (\operatorname{tr} A)\, A + (\det A)\, \mathrm{I} = \mathrm{O}$. **Remark.** This result is a special case of the *Cayley–Hamilton Theorem*, to be developed in Exercise 8.6.22.

(*b*) Prove the inverse formula $A^{-1} = \dfrac{(\operatorname{tr} A)\, \mathrm{I} - A}{\det A}$ when $\det A \neq 0$.

(*c*) Check the Cayley–Hamilton and inverse formulas when $A = \begin{pmatrix} 2 & 1 \\ -3 & 2 \end{pmatrix}$.

The Gershgorin Circle Theorem

In general, precisely computing the eigenvalues of a matrix is not easy, and, in most cases, must be done through a numerical eigenvalue procedure; see Sections 9.5 and 9.6. In certain applications, though, we may not require exact numerical values, but only their approximate locations. The *Gershgorin Circle Theorem*, due to the early-twentieth-century Russian mathematician Semyon Gershgorin, serves to restrict the eigenvalues to a certain well-defined region in the complex plane.

Definition 8.15. Let A be an $n \times n$ matrix, either real or complex. For each $1 \leq i \leq n$, define the ith *Gershgorin disk*

$$D_i = \{\, |\, z - a_{ii}\, | \leq r_i \;|\; z \in \mathbb{C} \,\}, \qquad \text{where} \qquad r_i = \sum_{\substack{j=1 \\ j \neq i}}^{n} |\, a_{ij}\, |. \tag{8.27}$$

The *Gershgorin domain* $D_A = \bigcup_{i=1}^{n} D_i \subset \mathbb{C}$ is the union of the Gershgorin disks.

Thus, the ith Gershgorin disk D_i is centered at the ith diagonal entry a_{ii} of A, and has radius r_i equal to the sum of the absolute values of the off-diagonal entries that are in its ith row. We can now state the Gershgorin Circle Theorem.

Theorem 8.16. All real and complex eigenvalues of the matrix A lie in its Gershgorin domain $D_A \subset \mathbb{C}$.

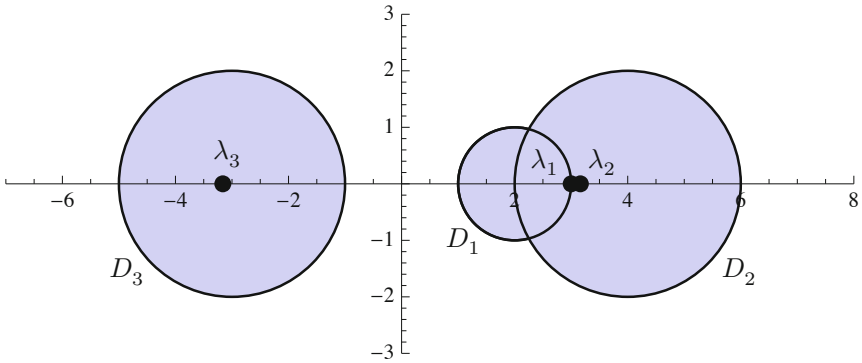

Figure 8.2. Gershgorin Disks and Eigenvalues.

Example 8.17. The matrix $A = \begin{pmatrix} 2 & -1 & 0 \\ 1 & 4 & -1 \\ -1 & -1 & -3 \end{pmatrix}$ has Gershgorin disks

$$D_1 = \{\, |\, z - 2\,| \le 1 \,\}, \qquad D_2 = \{\, |\, z - 4\,| \le 2 \,\}, \qquad D_3 = \{\, |\, z + 3\,| \le 2 \,\},$$

which are plotted in Figure 8.2. The eigenvalues of A are

$$\lambda_1 = 3, \qquad \lambda_2 = \sqrt{10} = 3.1622\ldots, \qquad \lambda_3 = -\sqrt{10} = -3.1622\ldots.$$

Observe that λ_1 belongs to both D_1 and D_2, while λ_2 lies in D_2, and λ_3 is in D_3. We thus confirm that all three eigenvalues are in the Gershgorin domain $D_A = D_1 \cup D_2 \cup D_3$.

Proof of Theorem 8.16: Let $\mathbf{v}$ be an eigenvector of A with eigenvalue λ. Let $\mathbf{u} = \mathbf{v}/\|\mathbf{v}\|_\infty$ be the corresponding unit eigenvector with respect to the ∞ norm, so

$$\|\mathbf{u}\|_\infty = \max\{\, |\, u_1\,|, \ \ldots\ , |\, u_n\,| \,\} = 1.$$

Let u_i be an entry of $\mathbf{u}$ that achieves the maximum: $|\, u_i\,| = 1$. Writing out the ith component of the eigenvalue equation $A\mathbf{u} = \lambda\,\mathbf{u}$, we obtain

$$\sum_{j=1}^{n} a_{ij}\, u_j = \lambda\, u_i, \qquad \text{which we rewrite as} \qquad \sum_{j \ne i} a_{ij}\, u_j = (\lambda - a_{ii})\, u_i.$$

Therefore, since all $|\, u_j\,| \le 1$, while $|\, u_i\,| = 1$,

$$|\, \lambda - a_{ii}\,| = |\, \lambda - a_{ii}\,|\,|\, u_i\,| = |\, (\lambda - a_{ii})\, u_i\,| = \left| \sum_{j \ne i} a_{ij}\, u_j \right| \le \sum_{j \ne i} |\, a_{ij}\,|\,|\, u_j\,| \le \sum_{j \ne i} |\, a_{ij}\,| = r_i.$$

This immediately implies that $\lambda \in D_i \subset D_A$ belongs to the ith Gershgorin disk. *Q.E.D.*

According to Proposition 8.7, a matrix A is singular if and only if it admits zero as an eigenvalue. Thus, if its Gershgorin domain does not contain 0, it cannot be an eigenvalue, and hence A is necessarily invertible. The condition $0 \notin D_A$ requires that the matrix have large diagonal entries, as quantified by the following definition.

Definition 8.18. A square matrix A is called *strictly diagonally dominant* if

$$|\, a_{ii}\,| \ > \sum_{j \ne i} |\, a_{ij}\,|, \qquad \text{for all} \qquad i = 1, \ldots, n. \tag{8.28}$$

In other words, strict diagonal dominance requires each diagonal entry to be larger, in absolute value, than the sum of the absolute values of *all* the other entries in its row. For example, the matrix $\begin{pmatrix} 3 & -1 & 1 \\ 1 & -4 & 2 \\ -2 & -1 & 5 \end{pmatrix}$ is strictly diagonally dominant since

$$|\,3\,| > |\,-1\,| + |\,1\,|, \qquad |\,-4\,| > |\,1\,| + |\,2\,|, \qquad |\,5\,| > |\,-2\,| + |\,-1\,|.$$

Diagonally dominant matrices appear frequently in numerical solution methods for both ordinary and partial differential equations.

Theorem 8.19. A strictly diagonally dominant matrix is nonsingular.

Proof: The diagonal dominance inequalities (8.28) imply that the radius of the ith Gershgorin disk is strictly less than the modulus of its center: $r_i < |\,a_{ii}\,|$. This implies that the disk cannot contain 0; indeed, if $z \in D_i$, then, by the triangle inequality,

$$r_i > |\,z - a_{ii}\,| \geq |\,a_{ii}\,| - |\,z\,| > r_i - |\,z\,|, \qquad \text{and hence} \qquad |\,z\,| > 0.$$

Thus, 0 does not lie in the Gershgorin domain D_A, and so cannot be an eigenvalue. *Q.E.D.*

Warning. The converse to this result is obviously not true; there are plenty of nonsingular matrices that are not strictly diagonally dominant.

Exercises

8.2.53. For each of the following matrices,
(i) find all Gershgorin disks; (ii) plot the Gershgorin domain in the complex plane;
(iii) compute the eigenvalues and confirm the truth of the Circle Theorem 8.16:

(a) $\begin{pmatrix} 1 & -2 \\ -2 & 1 \end{pmatrix}$, (b) $\begin{pmatrix} 1 & -\frac{2}{3} \\ \frac{1}{2} & -\frac{1}{6} \end{pmatrix}$, (c) $\begin{pmatrix} 2 & 3 \\ -1 & 0 \end{pmatrix}$, (d) $\begin{pmatrix} 3 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 3 \end{pmatrix}$,

(e) $\begin{pmatrix} -1 & 1 & 1 \\ 2 & 2 & -1 \\ 0 & 3 & -4 \end{pmatrix}$, (f) $\begin{pmatrix} \frac{1}{2} & 0 & 0 \\ 0 & 0 & \frac{1}{3} \\ \frac{1}{4} & \frac{1}{6} & 0 \end{pmatrix}$, (g) $\begin{pmatrix} 0 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & -1 & 1 \end{pmatrix}$, (h) $\begin{pmatrix} 3 & 2 & 0 & 0 \\ 1 & 2 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 2 & 1 \end{pmatrix}$.

8.2.54. *True or false*: The Gershgorin domain of the transpose of a matrix A^T is the same as the Gershgorin domain of the matrix A, that is, $D_{A^T} = D_A$.

♢ 8.2.55. (i) Explain why the eigenvalues of A must lie in its *refined Gershgorin domain* $D_A^* = D_{A^T} \cap D_A$. (ii) Find the refined Gershgorin domains for each of the matrices in Exercise 8.2.53 and confirm the result in part (i).

8.2.56. *True or false*: (a) A positive definite matrix is strictly diagonally dominant.
(b) A strictly diagonally dominant matrix is positive definite.

♢ 8.2.57. Prove that if K is symmetric, strictly diagonally dominant, and each diagonal entry is positive, then K is positive definite.

8.2.58. (a) Write down an invertible matrix A whose Gershgorin domain contains 0.
(b) Can you find an example that is also strictly diagonally dominant?

8.3 Eigenvector Bases

Most of the vector space bases that play a distinguished role in applications are assembled from the eigenvectors of a particular matrix. In this section, we show that the eigenvectors of a "complete" matrix automatically form a basis for $\mathbb{R}^n$, or, in the complex case, $\mathbb{C}^n$. In the following subsection, we use the eigenvector basis to rewrite the linear transformation determined by the matrix in a simple diagonal form, hence the alternative more common term "diagonalizable" for such matrices. The most important cases — symmetric and positive definite matrices — will be treated in the following section.

The first task is to show that eigenvectors corresponding to distinct eigenvalues are automatically linearly independent.

Lemma 8.20. If $\lambda_1, \ldots, \lambda_k$ are *distinct* eigenvalues of a matrix A, so $\lambda_i \neq \lambda_j$ when $i \neq j$, then the corresponding eigenvectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ are linearly independent.

Proof: The result is proved by induction on the number of eigenvalues. The case $k = 1$ is immediate, since an eigenvector cannot be zero. Assume that we know that the result is valid for $k - 1$ eigenvalues. Suppose we have a vanishing linear combination:

$$c_1 \mathbf{v}_1 + \cdots + c_{k-1} \mathbf{v}_{k-1} + c_k \mathbf{v}_k = \mathbf{0}. \tag{8.29}$$

Let us multiply this equation by the matrix A:

$$\begin{aligned} A\big(c_1 \mathbf{v}_1 + \cdots + c_{k-1} \mathbf{v}_{k-1} + c_k \mathbf{v}_k\big) &= c_1 A \mathbf{v}_1 + \cdots + c_{k-1} A \mathbf{v}_{k-1} + c_k A \mathbf{v}_k \\ &= c_1 \lambda_1 \mathbf{v}_1 + \cdots + c_{k-1} \lambda_{k-1} \mathbf{v}_{k-1} + c_k \lambda_k \mathbf{v}_k = \mathbf{0}. \end{aligned}$$

On the other hand, if we multiply the original equation (8.29) by λ_k, we also have

$$c_1 \lambda_k \mathbf{v}_1 + \cdots + c_{k-1} \lambda_k \mathbf{v}_{k-1} + c_k \lambda_k \mathbf{v}_k = \mathbf{0}.$$

On subtracting this from the previous equation, the final terms cancel, and we are left with the equation

$$c_1(\lambda_1 - \lambda_k)\mathbf{v}_1 + \cdots + c_{k-1}(\lambda_{k-1} - \lambda_k)\mathbf{v}_{k-1} = \mathbf{0}.$$

This is a vanishing linear combination of the first $k-1$ eigenvectors, and so, by our induction hypothesis, can happen only if all the coefficients are zero:

$$c_1(\lambda_1 - \lambda_k) = 0, \qquad \ldots \qquad c_{k-1}(\lambda_{k-1} - \lambda_k) = 0.$$

The eigenvalues were assumed to be distinct, and consequently $c_1 = \cdots = c_{k-1} = 0$. Substituting these values back into (8.29), we find that $c_k \mathbf{v}_k = \mathbf{0}$, and so $c_k = 0$ also, since the eigenvector $\mathbf{v}_k \neq \mathbf{0}$. Thus we have proved that (8.29) holds if and only if $c_1 = \cdots = c_k = 0$, which implies the linear independence of the eigenvectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$. This completes the induction step. *Q.E.D.*

The most important consequence of this result concerns when a matrix has the maximum allotment of eigenvalues.

Theorem 8.21. If the $n \times n$ real matrix A has n distinct real eigenvalues $\lambda_1, \ldots, \lambda_n$, then the corresponding real eigenvectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$ form a basis of $\mathbb{R}^n$. If A (which may now be either a real or a complex matrix) has n distinct complex eigenvalues, then the corresponding eigenvectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$ form a basis of $\mathbb{C}^n$.

For instance, the 2×2 matrix in Example 8.5 has two distinct real eigenvalues, and its two independent eigenvectors form a basis of $\mathbb{R}^2$. The 3×3 matrix in Example 8.9 has

three distinct complex eigenvalues, and its eigenvectors form a basis for $\mathbb{C}^3$. If a matrix has multiple eigenvalues, then there may or may not be an eigenvector basis of $\mathbb{R}^n$ (or $\mathbb{C}^n$). The matrix in Example 8.6 admits an eigenvector basis, whereas the matrix in Example 8.8 does not. In general, it can be proved† that the dimension of an eigenspace is less than or equal to the corresponding eigenvalue's multiplicity. In particular, every simple eigenvalue has a one-dimensional eigenspace, and hence, up to scalar multiple, only one associated eigenvector.

Definition 8.22. An eigenvalue λ of a matrix A is called *complete* if the corresponding eigenspace $V_\lambda = \ker(A - \lambda\,\mathrm{I})$ has the same dimension as its multiplicity. The matrix A is said to be *complete* if all its eigenvalues are complete.

Note that a simple eigenvalue is automatically complete, since its eigenspace is the one-dimensional subspace or eigenline spanned by the corresponding eigenvector. Thus, only multiple eigenvalues can cause a matrix to be incomplete.

Remark. The multiplicity of an eigenvalue λ_i is sometimes referred to as its *algebraic multiplicity*. The dimension of the eigenspace V_λ is its *geometric multiplicity*, and so completeness requires that the two multiplicities be equal. The word "complete" is not standard, but has been chosen because it can be used to describe both matrices and their individual eigenvalues. Alternative terms used to describe complete matrices include *perfect*, *semi-simple*, and, as we discuss shortly, *diagonalizable*.

Theorem 8.23. An $n \times n$ real or complex matrix A is complete if and only if its eigenvectors span $\mathbb{C}^n$. In particular, an $n \times n$ matrix that has n distinct eigenvalues is complete.

Or, stated another way, a matrix is complete if and only if its eigenvectors form a basis of $\mathbb{C}^n$. Most matrices are complete. Incomplete $n \times n$ matrices, which have fewer than n linearly independent complex eigenvectors, are less pleasant to deal with, and we relegate most of the messy details to Section 8.6.

Remark. We already noted that complex eigenvectors of a real matrix always appear in conjugate pairs: $\mathbf{v} = \mathbf{x} \pm \mathrm{i}\,\mathbf{y}$. If the matrix is complete, then it can be shown that its real eigenvectors combined with the real and imaginary parts of its complex conjugate eigenvectors form a real basis for $\mathbb{R}^n$. (See Exercise 8.3.12 for the underlying principle.) For instance, the complex eigenvectors of the 3×3 matrix appearing in Example 8.9 are $\begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} \pm \mathrm{i} \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$. The vectors $\begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix}$, $\begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$, $\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$, consisting of its real eigenvector and the real and imaginary parts of its complex eigenvectors, form a basis for $\mathbb{R}^3$.

Exercises

8.3.1. Which of the following are complete eigenvalues for the indicated matrix? What is the dimension of the associated eigenspace? (*a*) 3, $\begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$; (*b*) 2, $\begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix}$;

† This follows from the Jordan Canonical Form Theorem 8.51.

$(c)\ 1,\ \begin{pmatrix} 0 & 0 & -1 \\ 1 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}$; $(d)\ 1,\ \begin{pmatrix} 1 & 1 & -1 \\ 1 & 1 & 0 \\ 1 & 1 & 2 \end{pmatrix}$; $(e)\ -1,\ \begin{pmatrix} -1 & -4 & -4 \\ 0 & -1 & 0 \\ 0 & 4 & 3 \end{pmatrix}$;

$(f)\ -\mathrm{i},\ \begin{pmatrix} -\mathrm{i} & 1 & 0 \\ -\mathrm{i} & 1 & -1 \\ 0 & 0 & -\mathrm{i} \end{pmatrix}$; $(g)\ -2,\ \begin{pmatrix} 1 & 0 & -1 & 1 \\ 0 & 1 & 0 & 1 \\ -1 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix}$; $(h)\ 1,\ \begin{pmatrix} 1 & -1 & -1 & -1 & -1 \\ 0 & 1 & -1 & -1 & -1 \\ 0 & 0 & 1 & -1 & -1 \\ 0 & 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$.

8.3.2. Find the eigenvalues and a basis for the each of the eigenspaces of the following matrices. Which are complete?

$(a)\ \begin{pmatrix} 4 & -4 \\ 1 & 0 \end{pmatrix}$, $(b)\ \begin{pmatrix} 6 & -8 \\ 4 & -6 \end{pmatrix}$, $(c)\ \begin{pmatrix} 3 & -2 \\ 4 & -1 \end{pmatrix}$, $(d)\ \begin{pmatrix} \mathrm{i} & -1 \\ 1 & \mathrm{i} \end{pmatrix}$, $(e)\ \begin{pmatrix} 4 & -1 & -1 \\ 0 & 3 & 0 \\ 1 & -1 & 2 \end{pmatrix}$,

$(f)\ \begin{pmatrix} -6 & 0 & -8 \\ -4 & 2 & -4 \\ 4 & 0 & 6 \end{pmatrix}$, $(g)\ \begin{pmatrix} -2 & 1 & -1 \\ 5 & -3 & 6 \\ 5 & -1 & 4 \end{pmatrix}$, $(h)\ \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ -1 & 1 & -1 & 0 \\ 1 & 0 & -1 & 0 \end{pmatrix}$, $(i)\ \begin{pmatrix} -1 & 0 & 1 & 2 \\ 0 & 1 & 0 & 1 \\ -1 & -4 & 1 & -2 \\ 0 & 1 & 0 & 1 \end{pmatrix}$.

8.3.3. Which of the following matrices admit eigenvector bases of $\mathbb{R}^n$? For those that do, exhibit such a basis. If not, what is the dimension of the subspace of $\mathbb{R}^n$ spanned by the eigenvectors? $(a)\ \begin{pmatrix} 1 & 3 \\ 3 & 1 \end{pmatrix}$, $(b)\ \begin{pmatrix} 1 & 3 \\ -3 & 1 \end{pmatrix}$, $(c)\ \begin{pmatrix} 1 & 3 \\ 0 & 1 \end{pmatrix}$, $(d)\ \begin{pmatrix} 1 & -2 & 0 \\ 0 & -1 & 0 \\ 4 & -4 & -1 \end{pmatrix}$,

$(e)\ \begin{pmatrix} 1 & -2 & 0 \\ 0 & -1 & 0 \\ 0 & -4 & -1 \end{pmatrix}$, $(f)\ \begin{pmatrix} 2 & 0 & 0 \\ 1 & -1 & 1 \\ 2 & 1 & -1 \end{pmatrix}$, $(g)\ \begin{pmatrix} 0 & 0 & -1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}$, $(h)\ \begin{pmatrix} 0 & 0 & -1 & 1 \\ 0 & -1 & 0 & 1 \\ 1 & 0 & -1 & 0 \\ 1 & 0 & 1 & 0 \end{pmatrix}$.

8.3.4. Answer Exercise 8.3.3 with $\mathbb{R}^n$ replaced by $\mathbb{C}^n$.

8.3.5. (a) Give an example of a 3×3 matrix with 1 as its only eigenvalue, and only one linearly independent eigenvector. (b) Find one that has two linearly independent eigenvectors.

8.3.6. *True or false*: (a) Every diagonal matrix is complete.
(b) Every upper triangular matrix is complete.

8.3.7. Prove that if A is a complete matrix, then so is $cA + d\,\mathrm{I}$, where c, d are any scalars.

8.3.8. (a) Prove that if A is complete, then so is A^2.
(b) Give an example of an incomplete matrix A such that A^2 is complete.

8.3.9. Let U be an upper triangular matrix with all its diagonal entries equal. Prove that U is complete if and only if U is a diagonal matrix.

♢ 8.3.10. Suppose $\mathbf{v}_1, \ldots, \mathbf{v}_n$ is an eigenvector basis for the complete matrix A, with $\lambda_1, \ldots, \lambda_n$ the corresponding eigenvalues. Prove that *every* eigenvalue of A is one of the $\lambda_1, \ldots, \lambda_n$.

♢ 8.3.11. Show that if A is complete, then every similar matrix $B = S^{-1}AS$ is also complete.

♢ 8.3.12. (a) Prove that if $\mathbf{x} \pm \mathrm{i}\,\mathbf{y}$ is a complex conjugate pair of eigenvectors of a real matrix A corresponding to complex conjugate eigenvalues $\mu \pm \mathrm{i}\,\nu$ with $\nu \neq 0$, then $\mathbf{x}$ and $\mathbf{y}$ are linearly independent real vectors. (b) More generally, if $\mathbf{v}_j = \mathbf{x}_j \pm \mathrm{i}\,\mathbf{y}_j$, $j = 1, \ldots, k$, are complex conjugate pairs of eigenvectors corresponding to *distinct* pairs of complex conjugate eigenvalues $\mu_j \pm \mathrm{i}\,\nu_j$, $\nu_j \neq 0$, then the real vectors $\mathbf{x}_1, \ldots, \mathbf{x}_k, \mathbf{y}_1, \ldots, \mathbf{y}_k$ are linearly independent. (c) Prove that if A is complete, then there exists a basis of $\mathbb{R}^n$ consisting of its real eigenvectors and real and imaginary parts of its complex eigenvectors.

Diagonalization

Let $L\colon \mathbb{R}^n \to \mathbb{R}^n$ be a linear transformation on n-dimensional Euclidean space. As we know, cf. Theorem 7.5, $L[\mathbf{x}] = A\,\mathbf{x}$ is prescribed by multiplication by an $n \times n$ matrix A. However, the matrix representing a given linear transformation will depend on the choice of basis for the underlying vector space $\mathbb{R}^n$. Linear transformations having a complicated matrix representation in terms of the standard basis $\mathbf{e}_1, \ldots, \mathbf{e}_n$ may be considerably simplified by choosing a suitably adapted basis $\mathbf{v}_1, \ldots, \mathbf{v}_n$. We are now in a position to understand how to effect such a simplification.

For example, the linear transformation $L\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x - y \\ 2x + 4y \end{pmatrix}$ studied in Example 7.19 is represented by the matrix $A = \begin{pmatrix} 1 & -1 \\ 2 & 4 \end{pmatrix}$ — when expressed in terms of the standard basis of $\mathbb{R}^2$. In terms of the alternative basis $\mathbf{v}_1 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$, $\mathbf{v}_2 = \begin{pmatrix} 1 \\ -2 \end{pmatrix}$, it is represented by the diagonal matrix $\begin{pmatrix} 2 & 0 \\ 0 & 3 \end{pmatrix}$, implying that it has a simple stretching action on the new basis vectors: $A\mathbf{v}_1 = 2\mathbf{v}_1$, $A\mathbf{v}_2 = 3\mathbf{v}_2$. Now we can understand the reason for this simplification. *The new basis consists of the two eigenvectors of the matrix A.* This observation is indicative of a general fact: representing a linear transformation in terms of an eigenvector basis has the effect of changing its matrix representative into a simple diagonal form — thereby *diagonalizing* the original coefficient matrix.

According to (7.31), if $\mathbf{v}_1, \ldots, \mathbf{v}_n$ form a basis of $\mathbb{R}^n$, then the corresponding matrix representative of the linear transformation $L[\mathbf{v}] = A\mathbf{v}$ is given by the similar matrix $B = S^{-1}A\,S$, where $S = (\,\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n\,)$ is the matrix whose columns are the basis vectors. In the preceding example,

$$S = \begin{pmatrix} 1 & 1 \\ -1 & -2 \end{pmatrix}, \quad \text{and hence} \quad S^{-1}A\,S = \begin{pmatrix} 2 & 1 \\ -1 & -1 \end{pmatrix}\begin{pmatrix} 1 & -1 \\ 2 & 4 \end{pmatrix}\begin{pmatrix} 1 & 1 \\ -1 & -2 \end{pmatrix} = \begin{pmatrix} 2 & 0 \\ 0 & 3 \end{pmatrix}.$$

Definition 8.24. A square matrix A is called *diagonalizable* if there exists a nonsingular matrix S and a diagonal matrix $\Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_n)$ such that

$$S^{-1}A\,S = \Lambda, \qquad \text{or, equivalently,} \qquad A = S\,\Lambda\,S^{-1}. \tag{8.30}$$

A diagonal matrix represents a linear transformation that simultaneously stretches[†] in the direction of the basis vectors. Thus, every diagonalizable matrix represents an elementary combination of (complex) stretching transformations.

To understand the diagonalization equation (8.30), we rewrite it in the equivalent form

$$A\,S = S\,\Lambda. \tag{8.31}$$

Using the columnwise action (1.11) of matrix multiplication, one easily sees that the kth column of this $n \times n$ matrix equation is given by

$$A\mathbf{v}_k = \lambda_k \mathbf{v}_k,$$

where $\mathbf{v}_k$ denotes the kth column of S. Therefore, the columns of S are necessarily eigenvectors, and the entries of the diagonal matrix Λ are the corresponding eigenvalues. And,

---

[†] A negative diagonal entry represents the combination of a reflection and stretch. Complex entries indicate complex "stretching" transformations. See Section 7.2 for details.

as a result, a diagonalizable matrix A must have n linearly independent eigenvectors, i.e., an eigenvector basis, to form the columns of the nonsingular diagonalizing matrix S. Since the diagonal form Λ contains the eigenvalues along its diagonal, it is uniquely determined up to a permutation of its entries.

Now, as we know, not every matrix has an eigenvector basis. Moreover, even when it exists, the eigenvector basis may be complex, in which case S is a complex matrix, and the entries of the diagonal matrix Λ are the complex eigenvalues. Thus, we should distinguish between complete matrices that are diagonalizable over the complex numbers and the more restrictive class of real matrices that can be diagonalized by a real matrix S.

Theorem 8.25. A matrix is complex diagonalizable if and only if it is complete. A real matrix is real diagonalizable if and only if it is complete and has all real eigenvalues.

Example 8.26. The 3×3 matrix $A = \begin{pmatrix} 0 & -1 & -1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{pmatrix}$ considered in Example 8.5 has eigenvector basis

$$\mathbf{v}_1 = \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}, \qquad \mathbf{v}_2 = \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix}, \qquad \mathbf{v}_3 = \begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix}.$$

We assemble these to form the eigenvector matrix

$$S = \begin{pmatrix} -1 & -1 & -1 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}, \qquad \text{whereby} \qquad S^{-1} = \begin{pmatrix} -1 & 0 & -1 \\ -1 & -1 & 0 \\ 1 & 1 & 1 \end{pmatrix}.$$

The diagonalization equation (8.30) becomes

$$S^{-1} A\, S = \begin{pmatrix} -1 & 0 & -1 \\ -1 & -1 & 0 \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} 0 & -1 & -1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{pmatrix} \begin{pmatrix} -1 & -1 & -1 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{pmatrix} = \Lambda,$$

with eigenvalues of A appearing on the diagonal of Λ, in the same order as the eigenvectors.

Remark. If a matrix is not complete, then it cannot be diagonalized. A simple example is a matrix of the form $A = \begin{pmatrix} 1 & c \\ 0 & 1 \end{pmatrix}$ with $c \neq 0$, which represents a shear in the direction of the x-axis. Incomplete matrices represent generalized shearing transformations, and will be the subject of Section 8.6.

Exercises

8.3.13. Diagonalize the following matrices: (a) $\begin{pmatrix} 3 & -9 \\ 2 & -6 \end{pmatrix}$, (b) $\begin{pmatrix} 5 & -4 \\ 2 & -1 \end{pmatrix}$, (c) $\begin{pmatrix} -4 & -2 \\ 5 & 2 \end{pmatrix}$,

(d) $\begin{pmatrix} -2 & 3 & 1 \\ 0 & 1 & -1 \\ 0 & 0 & 3 \end{pmatrix}$, (e) $\begin{pmatrix} 8 & 0 & -3 \\ -3 & 0 & -1 \\ 3 & 0 & -2 \end{pmatrix}$, (f) $\begin{pmatrix} 3 & 3 & 5 \\ 5 & 6 & 5 \\ -5 & -8 & -7 \end{pmatrix}$, (g) $\begin{pmatrix} 2 & 5 & 5 \\ 0 & 2 & 0 \\ 0 & -5 & -3 \end{pmatrix}$,

(h) $\begin{pmatrix} 1 & 0 & -1 & 1 \\ 0 & 2 & -1 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -2 \end{pmatrix}$, (i) $\begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}$, (j) $\begin{pmatrix} 2 & 1 & -1 & 0 \\ -3 & -2 & 0 & 1 \\ 0 & 0 & 1 & -2 \\ 0 & 0 & 1 & -1 \end{pmatrix}$.

8.3.14. Diagonalize the *Fibonacci matrix* $F = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$.

8.3.15. Diagonalize the matrix $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ of rotation through 90°. How would you interpret the result?

8.3.16. Diagonalize the rotation matrices (*a*) $\begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$, (*b*) $\begin{pmatrix} \frac{5}{13} & 0 & \frac{12}{13} \\ 0 & 1 & 0 \\ -\frac{12}{13} & 0 & \frac{5}{13} \end{pmatrix}$.

8.3.17. Which of these matrices have real diagonal forms? (*a*) $\begin{pmatrix} -2 & 1 \\ 4 & 1 \end{pmatrix}$, (*b*) $\begin{pmatrix} 1 & 2 \\ -3 & 1 \end{pmatrix}$,

(*c*) $\begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}$, (*d*) $\begin{pmatrix} 0 & 3 & 2 \\ -1 & 1 & -1 \\ 1 & -3 & -1 \end{pmatrix}$, (*e*) $\begin{pmatrix} 3 & -8 & 2 \\ -1 & 2 & 2 \\ 1 & -4 & 2 \end{pmatrix}$, (*f*) $\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -1 & -1 & -1 & -1 \end{pmatrix}$.

8.3.18. Diagonalize the following complex matrices:

(*a*) $\begin{pmatrix} \mathrm{i} & 1 \\ 1 & \mathrm{i} \end{pmatrix}$, (*b*) $\begin{pmatrix} 1-\mathrm{i} & 0 \\ \mathrm{i} & 2+\mathrm{i} \end{pmatrix}$, (*c*) $\begin{pmatrix} 2-\mathrm{i} & 2+\mathrm{i} \\ 3-\mathrm{i} & 1+\mathrm{i} \end{pmatrix}$, (*d*) $\begin{pmatrix} -\mathrm{i} & 0 & 1 \\ -\mathrm{i} & 1 & -1 \\ 1 & 0 & -\mathrm{i} \end{pmatrix}$.

8.3.19. Write down a real matrix that has

(*a*) eigenvalues $-1, 3$ and corresponding eigenvectors $\begin{pmatrix} -1 \\ 2 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \end{pmatrix}$,

(*b*) eigenvalues $0, 2, -2$ and associated eigenvectors $\begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 3 \end{pmatrix}$;

(*c*) an eigenvalue of 3 and corresponding eigenvectors $\begin{pmatrix} 2 \\ -3 \end{pmatrix}, \begin{pmatrix} 1 \\ 2 \end{pmatrix}$;

(*d*) an eigenvalue $-1 + 2\,\mathrm{i}$ and corresponding eigenvector $\begin{pmatrix} 1+\mathrm{i} \\ 3\,\mathrm{i} \end{pmatrix}$;

(*e*) an eigenvalue -2 and corresponding eigenvector $\begin{pmatrix} 2 \\ 0 \\ -1 \end{pmatrix}$;

(*f*) an eigenvalue $3 + \mathrm{i}$ and corresponding eigenvector $\begin{pmatrix} 1 \\ 2\,\mathrm{i} \\ -1-\mathrm{i} \end{pmatrix}$.

8.3.20. A matrix A has eigenvalues -1 and 2 and associated eigenvectors $\begin{pmatrix} 1 \\ 2 \end{pmatrix}$ and $\begin{pmatrix} 2 \\ 3 \end{pmatrix}$. Write down the matrix form of the linear transformation $L[\mathbf{u}] = A\mathbf{u}$ in terms of (*a*) the standard basis $\mathbf{e}_1, \mathbf{e}_2$; (*b*) the basis consisting of its eigenvectors; (*c*) the basis $\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 3 \\ 4 \end{pmatrix}$.

♢ 8.3.21. Prove that two complete matrices A, B have the same eigenvalues (with multiplicities) if and only if they are similar, i.e., $B = S^{-1} A S$ for some nonsingular matrix S.

8.3.22. Let B be obtained from A by permuting both its rows and columns using the same permutation π, so $b_{ij} = a_{\pi(i),\pi(j)}$. Prove that A and B have the same eigenvalues. How are their eigenvectors related?

8.3.23. *True or false*: If A is a complete upper triangular matrix, then it has an upper triangular eigenvector matrix S.

8.3.24. How many different diagonal forms does an $n \times n$ diagonalizable matrix have?

8.3.25. Characterize all complete matrices that are their own inverses: $A^{-1} = A$. Write down a non-diagonal example.

♡ 8.3.26. Two $n \times n$ matrices A, B are said to be *simultaneously diagonalizable* if there is a nonsingular matrix S such that both $S^{-1}AS$ and $S^{-1}BS$ are diagonal matrices.
(a) Show that simultaneously diagonalizable matrices commute: $AB = BA$.
(b) Prove that the converse is valid, provided that one of the matrices has no multiple eigenvalues. (c) Is every pair of commuting matrices simultaneously diagonalizable?

8.4 Invariant Subspaces

The notion of an invariant subspace of a linear map plays an important role in dynamical systems, both finite- and infinite-dimensional, as well as in linear iterative systems, and in linear control systems. With the theory of eigenvalues and eigenvectors in hand, we are now able to completely characterize them.

Definition 8.27. Let $L: V \to V$ be a linear transformation on a vector space V. A subspace $W \subset V$ is said to be *invariant* if $L[\mathbf{w}] \in W$ whenever $\mathbf{w} \in W$.

Trivial examples of invariant subspaces, valid for any linear map, are the entire space $W = V$ and the zero subspace $W = \{\mathbf{0}\}$. If $L = \mathrm{I}$ is the identity transformation, then *every* subspace $W \subset V$ is invariant. More interestingly, both the kernel and image of L are invariant subspaces. Indeed, when $\mathbf{w} \in \ker L$, then $L[\mathbf{w}] = \mathbf{0} \in \ker L$, proving invariance. Similarly, if $\mathbf{w} \in \operatorname{img} L$, then $L[\mathbf{w}]$ also lies in $\operatorname{img} L$ by the definition of image.

Example 8.28. Let $V = \mathbb{R}^2$. Let us find all invariant subspaces of the scaling transformation $L(x, y) = (2x, 3y)^T$. If W is the line spanned by a vector $\mathbf{w} = (a, b)^T \neq \mathbf{0}$ then $L[\mathbf{w}] = (2a, 3b)^T \in W$ if and only if $(2a, 3b)^T = c\mathbf{w} = (ca, cb)^T$ for some scalar c. This is clearly possible if and only if either $a = 0$ or $b = 0$. Thus, the only one-dimensional invariant subspaces of this scaling transformation are the x- and y-axes.

Next, consider the linear transformation $L(x, y) = (x + 3y, y)^T$ corresponding to a shear in the direction of the x-axis. By the same reasoning, the one-dimensional subspace spanned by $\mathbf{w} = (a, b)^T \neq \mathbf{0}$ is invariant if and only if $(a + 3b, b)^T = (ca, cb)^T$ for $c \in \mathbb{R}$, which is possible only if $b = 0$. Thus, the only one-dimensional invariant subspace of this shearing transformation is the x-axis itself.

Finally, consider the linear transformation $L(x, y) = (-y, x)^T$ corresponding to counterclockwise rotation by $90°$. It is easy to see, either geometrically or algebraically, that L has no nontrivial invariant subspaces. On the other hand, if we view L as a map on $\mathbb{C}^2$, then one can show that there are two one-dimensional complex invariant subspaces, namely those spanned by its eigenvectors $\mathbf{w}_1 = (1, \mathrm{i})^T$ and $\mathbf{w}_2 = (1, -\mathrm{i})^T$.

From here on, we restrict our attention to the finite-dimensional case, and so the linear transformation L on either $\mathbb{R}^n$ or $\mathbb{C}^n$ is given by matrix multiplication: $L[\mathbf{x}] = A\mathbf{x}$ for some $n \times n$ matrix A.

Proposition 8.29. A one-dimensional subspace is invariant under the linear transformation $L[\mathbf{x}] = A\mathbf{x}$ if and only if it is an eigenline spanned by an eigenvector of A.

Proof: Let W be spanned by the single non-zero vector $\mathbf{w} \neq \mathbf{0}$. Then $A\mathbf{w} \in W$ if and only if $A\mathbf{w} = \lambda\mathbf{w}$ for some scalar λ. But this means that $\mathbf{w}$ is an eigenvector of A with eigenvalue λ. *Q.E.D.*

Thus, if A has no multiple eigenvalues, it has a finite number of one-dimensional invariant subspaces, namely the eigenspaces (eigenlines) associated with each eigenvalue. On the other hand, if λ is a multiple eigenvalue, then every one-dimensional subspace of its eigenspace $W \subset V_\lambda = \{ \mathbf{v} \mid A\mathbf{v} = \lambda \mathbf{v} \}$ is invariant.

We already observed that if $A = \lambda \mathrm{I}$, then every subspace $V \subset \mathbb{C}^n$ is invariant. On the other hand, if A is diagonal, with all distinct entries, then the invariant subspaces are necessarily spanned by a finite collection of standard basis vectors $\mathbf{e}_{l_1}, \ldots, \mathbf{e}_{l_k}$, which are recognized as its eigenvectors. This is a special case of the following general characterization of invariant subspaces of complete matrices. The incomplete case will be dealt with in Section 8.6.

Theorem 8.30. If A is a complete matrix, then every k-dimensional complex invariant subspace is spanned by k linearly independent eigenvectors of A.

Proof: Let $W \neq \{\mathbf{0}\}$ be a nontrivial invariant subspace. Thanks to completeness, we can express every nonzero vector $\mathbf{0} \neq \mathbf{w} \in W$ as a linear combination, $\mathbf{w} = c_1\mathbf{v}_1 + \cdots + c_j\mathbf{v}_j$, where $\mathbf{v}_1, \ldots, \mathbf{v}_j$ are eigenvectors associated with *distinct* eigenvalues $\lambda_1, \ldots, \lambda_j$ of A and all coefficients $c_i \neq 0$. We claim that this implies that each represented eigenvector $\mathbf{v}_i \in W$ for $i = 1, \ldots, j$, which clearly establishes the result. To prove the claim, we write

$$A\mathbf{w} - \lambda_j\mathbf{w} = c_1(\lambda_1 - \lambda_j)\mathbf{v}_1 + \cdots + c_{k-1}(\lambda_{j-1} - \lambda_j)\mathbf{v}_{j-1} \in W,$$

since, by the assumption of invariance, both terms on the left hand side belong to W. Moreover, since the eigenvalues are distinct, we must have $c_i(\lambda_i - \lambda_j) \neq 0$ for $i = 1, \ldots, j-1$. Iterating this process, we eventually conclude that a nonzero multiple of $\mathbf{v}_1$ and hence $\mathbf{v}_1$ itself belongs to W. This result is independent of the ordering of the eigenvectors, and hence all $\mathbf{v}_1, \ldots, \mathbf{v}_j \in W$. *Q.E.D.*

If A is a complete real matrix that possesses all real eigenvalues, then the same proof shows that every real invariant subspace has the form given in Theorem 8.30. If A is real and complete, with complex conjugate eigenvalues, Theorem 8.30 describes its complex invariant subspaces. Its real invariant subspaces are obtained from the real and imaginary parts of the eigenvectors. For example, if $\mathbf{v}_\pm = \mathbf{x} \pm \mathrm{i}\,\mathbf{y}$ are a complex conjugate pair of eigenvectors, then they individually span one-dimensional complex invariant subspaces. However, the smallest corresponding real invariant subspace is the two-dimensional subspace spanned by $\mathbf{x}$ and $\mathbf{y}$.

Example 8.31. Consider the three-dimensional rotation (permutation) matrix $A = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$. It has one real eigenvalue, $\lambda_1 = 1$, and two complex conjugate eigenvalues, $\lambda_2 = \frac{1}{2} + \frac{\sqrt{3}}{2}\,\mathrm{i}$ and $\lambda_3 = \frac{1}{2} - \frac{\sqrt{3}}{2}\,\mathrm{i}$. The complex invariant subspaces are spanned by 0, 1, 2, or 3 of the corresponding complex eigenvectors $\left(-\frac{1}{2}, -\frac{1}{2}, 1\right)^T \pm \mathrm{i}\left(\frac{\sqrt{3}}{2}, -\frac{\sqrt{3}}{2}, 0\right)^T$. There is a single one-dimensional real invariant subspace, spanned by the real eigenvector $(1, 1, 1)^T$, and a single two-dimensional real invariant subspace, which is the orthogonal complement spanned by the real and imaginary parts of the complex conjugate eigenvectors. Indeed, this indicates a general property satisfied by all 3×3 rotation matrices, with the exception of the trivial identity matrix. The unique one-dimensional real invariant subspace is the axis of the rotation, and the matrix reduces to a two-dimensional rotation on its orthogonal complement. See Exercise 8.2.44 for further details.

Exercises

8.4.1. Find all invariant subspaces $W \subset \mathbb{R}^2$ of the following linear transformations $L: \mathbb{R}^3 \to \mathbb{R}^3$: (*a*) the scaling transformation $(2x, 3y, 4z)^T$; (*b*) the shear $(x + 3y, y, z)^T$; (*c*) counterclockwise rotation by a 45° angle around the x-axis.

8.4.2. Find all invariant subspaces of the following matrices: (*a*) $\begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$, (*b*) $\begin{pmatrix} 3 & 0 \\ -2 & 3 \end{pmatrix}$, (*c*) $\begin{pmatrix} 0 & 0 & 2 \\ 0 & 1 & 0 \\ 2 & 0 & 0 \end{pmatrix}$, (*d*) $\begin{pmatrix} -6 & 0 & -8 \\ -4 & 2 & -4 \\ 4 & 0 & 6 \end{pmatrix}$, (*e*) $\begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}$, (*f*) $\begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}$.

8.4.3. Find all complex invariant subspaces and all real invariant subspaces of

$$(a)\ \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \quad (b)\ \begin{pmatrix} -1 & 2 \\ 1 & 1 \end{pmatrix}, \quad (c)\ \begin{pmatrix} 3 & 3 & 5 \\ 5 & 6 & 5 \\ -5 & -8 & -7 \end{pmatrix}, \quad (d)\ \begin{pmatrix} 2 & 5 & 5 \\ 0 & 2 & 0 \\ 0 & -5 & -3 \end{pmatrix}.$$

8.4.4. Prove that if W is an invariant subspace for A, then it is also invariant for A^2. Is the converse to this statement valid?

♢ 8.4.5. Let $V \subset \mathbb{R}^n$ be an invariant subspace for the $n \times n$ matrix A. Explain why every eigenvalue and eigenvector of the linear map obtained by restricting A to V are also eigenvalues and eigenvectors of A itself.

8.4.6. *True or false*: If V and W are invariant subspaces for the matrix A, then so is

$$(a)\ V + W; \quad (b)\ V \cap W; \quad (c)\ V \cup W; \quad (d)\ V \setminus W.$$

8.4.7. *True or false*: If V is an invariant subspace for the $n \times n$ matrix A and W is an invariant subspace for the $n \times n$ matrix B, then $V + W$ is an invariant subspace for the matrix $A + B$.

8.4.8. *True or false*: If W is an invariant subspace of the matrix A, then it is also an invariant subspace of A^T.

8.4.9. *True or false*: If W is an invariant subspace of the nonsingular matrix A, then it is also an invariant subspace of A^{-1}.

8.4.10. Which 2×2 orthogonal matrices have a nontrivial real invariant subspace?

8.4.11. *True or false*: If $Q \neq \pm \mathrm{I}$ is a 4×4 orthogonal matrix, then Q has no real invariant subspaces.

♢ 8.4.12. (*a*) Let A be an $n \times n$ symmetric matrix, and let $\mathbf{v}$ be an eigenvector. Prove that its orthogonal complement under the dot product, namely, $V^\perp = \{\, \mathbf{w} \in \mathbb{R}^n \mid \mathbf{v}_1 \cdot \mathbf{w} = 0 \,\}$, is an invariant subspace. (*b*) More generally, prove that if $W \subset \mathbb{R}^n$ is an invariant subspace, then its orthogonal complement $W^\perp$, is also invariant.

8.5 Eigenvalues of Symmetric Matrices

Fortunately, the matrices that arise in most applications are complete and, in fact, possess some additional structure that ameliorates the calculation of their eigenvalues and eigenvectors. The most important class is that of the symmetric, including positive definite, matrices. In fact, not only are the eigenvalues of a symmetric matrix necessarily real, the eigenvectors always form an *orthogonal basis* of the underlying Euclidean space, enjoying all the wonderful properties we studied in Chapter 4. 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 8.32. 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 real symmetric matrices are complete and real diagonalizable.

Remark. Orthogonality is with respect to the standard dot product on $\mathbb{R}^n$. As we noted in Section 7.5, the transpose is a particular case of the adjoint operation when we use the Euclidean dot product. An analogous result holds for more general self-adjoint linear transformations under more general inner products on $\mathbb{R}^n$; see Exercise 8.5.10.

Example 8.33. The 2×2 matrix $A = \begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix}$ considered in Example 8.5 is symmetric, and so has real eigenvalues $\lambda_1 = 4$ and $\lambda_2 = 2$. You can easily check that the corresponding eigenvectors $\mathbf{v}_1 = (\,1, 1\,)^T$ and $\mathbf{v}_2 = (\,-1, 1\,)^T$ are orthogonal: $\mathbf{v}_1 \cdot \mathbf{v}_2 = 0$, and hence form an orthogonal basis of $\mathbb{R}^2$. The orthonormal eigenvector basis promised by Theorem 8.32 is obtained by dividing each eigenvector by its Euclidean norm:

$$\mathbf{u}_1 = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}, \qquad \mathbf{u}_2 = \begin{pmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}.$$

Example 8.34. 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:

$$\lambda_1 = 9, \qquad \lambda_2 = 3, \qquad \lambda_3 = -3,$$

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

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

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

The eigenvalues of a symmetric matrix can be used to test its positive definiteness.

Theorem 8.35. A symmetric matrix $K = K^T$ is positive definite if and only if all of its eigenvalues are strictly positive.

Proof: First, if $K > 0$, then, by definition, $\mathbf{x}^T K \mathbf{x} > 0$ for all nonzero vectors $\mathbf{x} \in \mathbb{R}^n$. In particular, if $\mathbf{x} = \mathbf{v} \neq \mathbf{0}$ is an eigenvector with (necessarily real) eigenvalue λ, then

$$0 < \mathbf{v}^T K \mathbf{v} = \mathbf{v}^T (\lambda \mathbf{v}) = \lambda \mathbf{v}^T \mathbf{v} = \lambda \, \| \mathbf{v} \|^2, \tag{8.32}$$

which immediately proves that $\lambda > 0$.

Conversely, suppose K has all positive eigenvalues. Let $\mathbf{u}_1, \ldots, \mathbf{u}_n$ be the orthonormal eigenvector basis guaranteed by Theorem 8.32, with $K\mathbf{u}_j = \lambda_j \mathbf{u}_j$ with $\lambda_j > 0$. Writing

$$\mathbf{x} = c_1 \mathbf{u}_1 + \cdots + c_n \mathbf{u}_n, \qquad \text{we obtain} \qquad K\mathbf{x} = c_1 \lambda_1 \mathbf{u}_1 + \cdots + c_n \lambda_n \mathbf{u}_n.$$

Therefore, using the orthonormality of the eigenvectors,

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

whenever $\mathbf{x} \neq \mathbf{0}$, since only $\mathbf{x} = \mathbf{0}$ has coordinates $c_1 = \cdots = c_n = 0$. This establishes the positive definiteness of K. *Q.E.D.*

Remark. The same proof shows that K is positive semi-definite if and only if all its eigenvalues satisfy $\lambda \geq 0$. A positive semi-definite matrix that is not positive definite admits a zero eigenvalue and one or more *null eigenvectors*, i.e., solutions to $K\mathbf{v} = \mathbf{0}$. Every nonzero element $\mathbf{0} \neq \mathbf{v} \in \ker K$ of its kernel is a null eigenvector.

Example 8.36. The symmetric matrix $K = \begin{pmatrix} 8 & 0 & 1 \\ 0 & 8 & 1 \\ 1 & 1 & 7 \end{pmatrix}$ has characteristic equation

$$\det(K - \lambda I) = -\lambda^3 + 23\lambda^2 - 174\lambda + 432 = -(\lambda - 9)(\lambda - 8)(\lambda - 6),$$

and so its eigenvalues are $9, 8$, and 6. Since they are all positive, K is a positive definite matrix. The associated eigenvectors are

$$\lambda_1 = 9, \qquad \mathbf{v}_1 = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \qquad \lambda_2 = 8, \qquad \mathbf{v}_2 = \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}, \qquad \lambda_3 = 6, \qquad \mathbf{v}_3 = \begin{pmatrix} -1 \\ -1 \\ 2 \end{pmatrix}.$$

Note that the eigenvectors form an orthogonal basis of $\mathbb{R}^3$, as guaranteed by Theorem 8.32. As usual, we can construct an corresponding orthonormal eigenvector basis

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

by dividing each eigenvector by its norm.

Proof of Theorem 8.32: First recall that (see Exercise 3.6.38) if $A = A^T$ is real, symmetric, then

$$(A\mathbf{v}) \cdot \mathbf{w} = \mathbf{v} \cdot (A\mathbf{w}) \qquad \text{for all} \qquad \mathbf{v}, \mathbf{w} \in \mathbb{C}^n, \tag{8.33}$$

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 (8.33),

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

Equating these two expressions, we deduce

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

Since $\mathbf{v}$ is an eigenvector, it must be nonzero. Thus, we deduce that $\overline{\lambda} = \lambda$, proving that the eigenvalue λ must be real.

To prove (b), suppose $A\,\mathbf{v} = \lambda\,\mathbf{v}$, $A\,\mathbf{w} = \mu\,\mathbf{w}$, where $\lambda \neq \mu$ are distinct real eigenvalues. Then, again by (8.33),

$$\lambda\,\mathbf{v}\cdot\mathbf{w} = (A\,\mathbf{v})\cdot\mathbf{w} = \mathbf{v}\cdot(A\,\mathbf{w}) = \mathbf{v}\cdot(\mu\,\mathbf{w}) = \mu\,\mathbf{v}\cdot\mathbf{w}, \quad \text{and hence} \quad (\lambda-\mu)\,\mathbf{v}\cdot\mathbf{w} = 0.$$

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

Finally, the proof of (c) is easy if all the eigenvalues of A are distinct. Theorem 8.21 implies that the eigenvectors form a basis of $\mathbb{R}^n$, and part (b) proves they are orthogonal. (An alternative proof starts with orthogonality, and then applies Proposition 4.4 to prove that the eigenvectors 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 4.2.

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 an associated eigenvector. Let $V^{\perp} = \{\,\mathbf{w}\in\mathbb{R}^n \mid \mathbf{v}_1\cdot\mathbf{w} = 0\,\}$ denote its orthogonal complement — the subspace of all vectors orthogonal to the first eigenvector. Proposition 4.41 implies that $\dim V^{\perp} = n-1$, and so we can choose an orthonormal basis $\mathbf{y}_1,\ldots,\mathbf{y}_{n-1}$. Moreover, by Exercise 8.4.12, $V^{\perp}$ is an invariant subspace of A, and hence A defines a linear transformation on $V^{\perp}$ that is represented by an $(n-1)\times(n-1)$ matrix, say $B = (b_{ij})$, $i,j = 1,\ldots,n-1$, with respect to the chosen orthonormal basis $\mathbf{y}_1,\ldots,\mathbf{y}_{n-1}$, so that $A\,\mathbf{y}_i = \sum_{j=1}^{n-1} b_{ij}\,\mathbf{y}_j$. Moreover, by orthonormality and (8.33),

$$b_{ij} = \mathbf{y}_i\cdot A\,\mathbf{y}_j = (A\,\mathbf{y}_i)\cdot\mathbf{y}_j = b_{ji},$$

and hence $B = B^T$ is symmetric. Our induction hypothesis then implies that there is an orthonormal basis of $V^{\perp}\subset\mathbb{R}^n$ consisting of eigenvectors $\mathbf{u}_2,\ldots,\mathbf{u}_n$ of B, and hence also of A, each of which is orthogonal to $\mathbf{v}_1$. 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.*

Proposition 8.37. Let $A = A^T$ be an $n\times n$ symmetric matrix. Let $\mathbf{v}_1,\ldots,\mathbf{v}_n$ be an orthogonal eigenvector basis such that $\mathbf{v}_1,\ldots,\mathbf{v}_r$ correspond to nonzero eigenvalues, while $\mathbf{v}_{r+1},\ldots,\mathbf{v}_n$ are null eigenvectors corresponding to the zero eigenvalue (if any). Then $r = \operatorname{rank} A$; the non-null eigenvectors $\mathbf{v}_1,\ldots,\mathbf{v}_r$ form an orthogonal basis for $\operatorname{img} A = \operatorname{coimg} A$, while the null eigenvectors $\mathbf{v}_{r+1},\ldots,\mathbf{v}_n$ form an orthogonal basis for $\ker A = \operatorname{coker} A$.

Proof: The zero eigenspace coincides with the kernel, $V_0 = \ker A$. Thus, the linearly independent null eigenvectors form a basis for $\ker A$, which has dimension $n-r$ where $r = \operatorname{rank} A$. Moreover, the remaining r non-null eigenvectors are orthogonal to the null eigenvectors. Therefore, they must form a basis for the kernel's orthogonal complement, namely $\operatorname{coimg} A = \operatorname{img} A$. *Q.E.D.*

Exercises

8.5.1. Find the eigenvalues and an orthonormal eigenvector basis for the following symmetric matrices:

$$(a)\ \begin{pmatrix} 2 & 6 \\ 6 & -7 \end{pmatrix} \quad (b)\ \begin{pmatrix} 5 & -2 \\ -2 & 5 \end{pmatrix}, \quad (c)\ \begin{pmatrix} 2 & -1 \\ -1 & 5 \end{pmatrix}, \quad (d)\ \begin{pmatrix} 1 & 0 & 4 \\ 0 & 1 & 3 \\ 4 & 3 & 1 \end{pmatrix}, \quad (e)\ \begin{pmatrix} 6 & -4 & 1 \\ -4 & 6 & -1 \\ 1 & -1 & 11 \end{pmatrix}.$$

8.5.2. Determine whether the following symmetric matrices are positive definite by computing their eigenvalues. Validate your conclusions by using the methods from Chapter 5.

$$(a)\ \begin{pmatrix} 2 & -2 \\ -2 & 3 \end{pmatrix} \quad (b)\ \begin{pmatrix} -2 & 3 \\ 3 & 6 \end{pmatrix}, \quad (c)\ \begin{pmatrix} 1 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 1 \end{pmatrix}, \quad (d)\ \begin{pmatrix} 4 & -1 & -2 \\ -1 & 4 & -1 \\ -2 & -1 & 4 \end{pmatrix}.$$

8.5.3. Prove that a symmetric matrix is negative definite if and only if all its eigenvalues are negative.

8.5.4. How many orthonormal eigenvector bases does a symmetric $n \times n$ matrix have?

8.5.5. Let $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$. (*a*) Write down necessary and sufficient conditions on the entries a, b, c, d that ensures that A has only real eigenvalues.
(*b*) Verify that all symmetric 2×2 matrices satisfy your conditions.
(*c*) Write down a non-symmetric matrix that satisfies your conditions.

♡ 8.5.6. Let $A^T = -A$ be a real, skew-symmetric $n \times n$ matrix. (*a*) Prove that the only possible real eigenvalue of A is $\lambda = 0$. (*b*) More generally, prove that all eigenvalues λ of A are purely imaginary, i.e., $\operatorname{Re} \lambda = 0$. (*c*) Explain why 0 is an eigenvalue of A whenever n is odd. (*d*) Explain why, if $n = 3$, the eigenvalues of $A \neq \mathrm{O}$ are $0, \mathrm{i}\,\omega, -\mathrm{i}\,\omega$, for some real $\omega \neq 0$. (*e*) Verify these facts for the particular matrices

$$(i)\ \begin{pmatrix} 0 & -2 \\ 2 & 0 \end{pmatrix}, \quad (ii)\ \begin{pmatrix} 0 & 3 & 0 \\ -3 & 0 & -4 \\ 0 & 4 & 0 \end{pmatrix}, \quad (iii)\ \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & -1 \\ 1 & 1 & 0 \end{pmatrix}, \quad (iv)\ \begin{pmatrix} 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & -3 \\ -2 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \end{pmatrix}.$$

♡ 8.5.7. (*a*) Prove that every eigenvalue of a Hermitian matrix A, satisfying $A^T = \overline{A}$ as in Exercise 3.6.45, is real. (*b*) Show that the eigenvectors corresponding to distinct eigenvalues are orthogonal under the Hermitian dot product on $\mathbb{C}^n$. (*c*) Find the eigenvalues and eigenvectors of the following Hermitian matrices, and verify orthogonality:

$$(i)\ \begin{pmatrix} 2 & \mathrm{i} \\ -\mathrm{i} & -2 \end{pmatrix}, \quad (ii)\ \begin{pmatrix} 3 & 2-\mathrm{i} \\ 2+\mathrm{i} & -1 \end{pmatrix}, \quad (iii)\ \begin{pmatrix} 0 & \mathrm{i} & 0 \\ -\mathrm{i} & 0 & \mathrm{i} \\ 0 & -\mathrm{i} & 0 \end{pmatrix}.$$

♡ 8.5.8. Let K, M be $n \times n$ matrices, with $M > 0$ positive definite. A nonzero vector $\mathbf{v} \neq \mathbf{0}$ is called a *generalized eigenvector* of matrix pair K, M if it satisfies the *generalized eigenvalue equation*

$$K\,\mathbf{v} = \lambda\, M\,\mathbf{v}, \qquad \mathbf{v} \neq \mathbf{0}, \tag{8.34}$$

where the scalar λ is the corresponding *generalized eigenvalue*. Note that ordinary eigenvalue/eigenvectors of K are when $M = \mathrm{I}$.(*a*) Prove that λ is a generalized eigenvalue if and only if it satisfies the *generalized characteristic equation* $\det(K - \lambda M) = 0$.
(*b*) Prove that λ is a generalized eigenvalue of the matrix pair K, M if and only if it is an ordinary eigenvalue of the matrix $M^{-1}K$. How are the eigenvectors related? (*c*) Now suppose K is a symmetric matrix. Prove that its generalized eigenvalues are all real.
Hint: First explain why this does *not* follow from part (*a*). Instead mimic the proof of part (*a*) of Theorem 8.32, using the weighted Hermitian inner product $\langle\, \mathbf{v}\,, \mathbf{w}\,\rangle = \mathbf{v}^T M\, \overline{\mathbf{w}}$ in place of the dot product. (*d*) Show that if $K > 0$, then its generalized eigenvalues are all positive: $\lambda > 0$. (*e*) Prove that the eigenvectors corresponding to different generalized eigenvalues are orthogonal under the weighted inner product $\langle\, \mathbf{v}\,, \mathbf{w}\,\rangle = \mathbf{v}^T M\, \mathbf{w}$. (*f*) Show

that, if the matrix pair K, M has n distinct generalized eigenvalues, then the eigenvectors form an orthogonal basis for $\mathbb{R}^n$. **Remark.** One can, by mimicking the proof of part (c) of Theorem 8.32, show that this holds even when there are repeated generalized eigenvalues.

8.5.9. Compute the generalized eigenvalues and eigenvectors, as in (8.34), for the following matrix pairs. Verify orthogonality of the eigenvectors under the appropriate inner product.

(a) $K = \begin{pmatrix} 3 & -1 \\ -1 & 2 \end{pmatrix}$, $M = \begin{pmatrix} 2 & 0 \\ 0 & 3 \end{pmatrix}$, (b) $K = \begin{pmatrix} 3 & 1 \\ 1 & 1 \end{pmatrix}$, $M = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}$,

(c) $K = \begin{pmatrix} 2 & -1 \\ -1 & 4 \end{pmatrix}$, $M = \begin{pmatrix} 2 & -1 \\ -1 & 1 \end{pmatrix}$, (d) $K = \begin{pmatrix} 6 & -8 & 3 \\ -8 & 24 & -6 \\ 3 & -6 & 99 \end{pmatrix}$, $M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 9 \end{pmatrix}$,

(e) $K = \begin{pmatrix} 1 & 2 & 0 \\ 2 & 8 & 2 \\ 0 & 2 & 1 \end{pmatrix}$, $M = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 3 & 1 \\ 0 & 1 & 1 \end{pmatrix}$, (f) $K = \begin{pmatrix} 5 & 3 & -5 \\ 3 & 3 & -1 \\ -5 & -1 & 9 \end{pmatrix}$, $M = \begin{pmatrix} 3 & 2 & -3 \\ 2 & 2 & -1 \\ -3 & -1 & 5 \end{pmatrix}$.

♢ 8.5.10. Let $L = L^* : \mathbb{R}^n \to \mathbb{R}^n$ be a self-adjoint linear transformation with respect to the inner product $\langle \cdot, \cdot \rangle$. Prove that all its eigenvalues are real and the eigenvectors are orthogonal. *Hint*: Mimic the proof of Theorem 8.32, replacing the dot product by the given inner product.

♡ 8.5.11. The *difference map* $\Delta : \mathbb{C}^n \to \mathbb{C}^n$ is defined as $\Delta = S - \mathrm{I}$, where S is the shift map of Exercise 8.2.13. (a) Write down the matrix D corresponding to Δ. (b) Prove that the sampled exponential vectors $\boldsymbol{\omega}_0, \ldots, \boldsymbol{\omega}_{n-1}$ from (5.102) form an eigenvector basis of D. What are the eigenvalues? (c) Prove that $K = D^T D$ has the same eigenvectors as D. What are its eigenvalues? (d) Is K positive definite? (e) According to Theorem 8.32 the eigenvectors of a symmetric matrix are real and orthogonal. Use this to explain the orthogonality of the sampled exponential vectors. But, why aren't they real?

♡ 8.5.12. An $n \times n$ *circulant matrix* has the form $C = \begin{pmatrix} c_0 & c_1 & c_2 & c_3 & \cdots & c_{n-1} \\ c_{n-1} & c_0 & c_1 & c_2 & \cdots & c_{n-2} \\ c_{n-2} & c_{n-1} & c_0 & c_1 & \cdots & c_{n-3} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ c_1 & c_2 & c_3 & c_4 & \cdots & c_0 \end{pmatrix}$,

in which the entries of each succeeding row are obtained by moving all the previous row's entries one slot to the right, the last entry moving to the front. (a) Check that the shift matrix A of Exercise 8.2.13, the difference matrix D, and its symmetric product $K = D^T D$ of Exercise 8.5.11 are all circulant matrices. (b) Prove that the sampled exponential vectors $\boldsymbol{\omega}_0, \ldots, \boldsymbol{\omega}_{n-1}$, cf. (5.102), are eigenvectors of C. Thus, all circulant matrices have *the same eigenvectors*! What are the eigenvalues? (c) Prove that $F_n^{-1} C F_n = \Lambda$, where F_n is the Fourier matrix in Exercise 5.6.9 and Λ is the diagonal matrix with the eigenvalues of C along the diagonal. (d) Find the eigenvalues and eigenvectors of the following circulant matrices:

(i) $\begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$, (ii) $\begin{pmatrix} 1 & 2 & 3 \\ 3 & 1 & 2 \\ 2 & 3 & 1 \end{pmatrix}$, (iii) $\begin{pmatrix} 1 & -1 & -1 & 1 \\ 1 & 1 & -1 & -1 \\ -1 & 1 & 1 & -1 \\ -1 & -1 & 1 & 1 \end{pmatrix}$, (iv) $\begin{pmatrix} 2 & -1 & 0 & -1 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ -1 & 0 & -1 & 2 \end{pmatrix}$.

(e) Find the eigenvalues of the tricirculant matrices in Exercise 1.7.13. Can you find a general formula for the $n \times n$ version? Explain why the eigenvalues must be real and positive. Does your formula reflect this fact? (f) Which of the preceding matrices are invertible? Write down a general criterion for checking the invertibility of circulant matrices.

The Spectral Theorem

Every real, symmetric matrix admits an eigenvector basis, and hence is diagonalizable. Moreover, since we can choose eigenvectors that form an orthonormal basis, the diagonalizing matrix takes a particularly simple form. Recall that an $n \times n$ matrix Q is *orthogonal* if and only if its columns form an orthonormal basis of $\mathbb{R}^n$. Alternatively, one characterizes orthogonal matrices by the condition $Q^{-1} = Q^T$, as in Definition 4.18.

Using the orthonormal eigenvector basis in the diagonalization formula (8.30) results in what is known as the *spectral factorization* of a symmetric matrix.

Theorem 8.38. Let A be a real, symmetric matrix. Then there exists an orthogonal matrix Q such that

$$A = Q\,\Lambda\,Q^{-1} = Q\,\Lambda\,Q^T, \tag{8.35}$$

where Λ 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.

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, [**54**]. The *Spectral Theorem* 8.38 is the finite-dimensional version of the decomposition of quantum mechanical linear operators into their spectral eigenstates.

Warning. Although both involve diagonal matrices, the spectral factorization $A = Q\,\Lambda\,Q^T$ and the Gaussian factorization $A = L\,D\,L^T$ of a regular symmetric matrix, cf. (1.58), are completely different. In particular, the eigenvalues are *not* the pivots, so $\Lambda \neq D$.

The spectral factorization (8.35) provides us with an alternative means of diagonalizing the associated quadratic form $q(\mathbf{x}) = \mathbf{x}^T A\,\mathbf{x}$, i.e., of completing the square. We write

$$q(\mathbf{x}) = \mathbf{x}^T A\,\mathbf{x} = \mathbf{x}^T\,Q\,\Lambda\,Q^T\,\mathbf{x} = \mathbf{y}^T \Lambda\,\mathbf{y} = \sum_{i=1}^{n} \lambda_i\,y_i^2, \tag{8.36}$$

where the entries of $\mathbf{y} = Q^T\mathbf{x} = Q^{-1}\mathbf{x}$ are the coordinates of $\mathbf{x}$ with respect to the orthonormal eigenvector basis of A. In particular, $q(\mathbf{x}) > 0$ for all $\mathbf{x} \neq \mathbf{0}$ and so A is positive definite if and only if each eigenvalue λ_i is strictly positive, reconfirming Theorem 8.35.

Example 8.39. For the 2×2 matrix $A = \begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix}$ considered in Example 8.33, the orthonormal eigenvectors produce the diagonalizing orthogonal matrix $Q = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}$. The reader can validate the resulting spectral factorization:

$$\begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix} = A = Q\,\Lambda\,Q^T = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}.$$

According to formula (8.36), the associated quadratic form is diagonalized as

$$q(\mathbf{x}) = 3\,x_1^2 + 2\,x_1\,x_2 + 3\,x_2^2 = 4\,y_1^2 + 2\,y_2^2,$$

where $\mathbf{y} = Q^T\mathbf{x}$, i.e., $y_1 = \dfrac{x_1 + x_2}{\sqrt{2}}$, $y_2 = \dfrac{-x_1 + x_2}{\sqrt{2}}$.

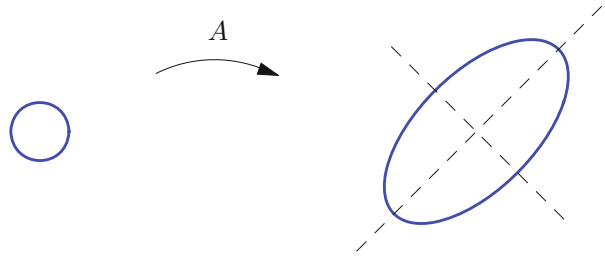

Figure 8.3. Stretching a Circle into an Ellipse.

You can always choose Q to be a proper orthogonal matrix, so $\det Q = 1$, since an improper orthogonal matrix can be made proper by multiplying one of its columns by -1, which does not affect its status as an eigenvector matrix. Since a proper orthogonal matrix Q represents a rigid rotation of $\mathbb{R}^n$, the diagonalization of a symmetric matrix can be interpreted as a rotation of the coordinate system that makes the orthogonal eigenvectors line up along the coordinate axes. Therefore, a linear transformation $L[\mathbf{x}] = A\mathbf{x}$ represented by a positive definite matrix A can be regarded as a combination of stretches in n mutually orthogonal directions. A good way to visualize this is to consider the effect of the linear transformation on the unit (Euclidean) sphere $S_1 = \{ \|\mathbf{x}\| = 1 \}$. Stretching the sphere in mutually orthogonal directions will map it to an ellipsoid $E = L[S_1] = \{ A\mathbf{x} \mid \|\mathbf{x}\| = 1 \}$ whose principal axes are aligned with the directions of stretch; see Figure 8.3 for the two-dimensional case. For instance, in elasticity, the stress tensor of a deformed body is represented by a positive definite matrix. Its eigenvalues are known as the *principal stretches* and its eigenvectors the *principal directions* of the elastic deformation.

Exercises

8.5.13. Write out the spectral factorization of the following matrices:

$$(a)\ \begin{pmatrix} -3 & 4 \\ 4 & 3 \end{pmatrix}, \quad (b)\ \begin{pmatrix} 2 & -1 \\ -1 & 4 \end{pmatrix}, \quad (c)\ \begin{pmatrix} 1 & 1 & 0 \\ 1 & 2 & 1 \\ 0 & 1 & 1 \end{pmatrix}, \quad (d)\ \begin{pmatrix} 3 & -1 & -1 \\ -1 & 2 & 0 \\ -1 & 0 & 2 \end{pmatrix}.$$

8.5.14. Write out the spectral factorization of the matrices listed in Exercise 8.5.1.

8.5.15. Construct a symmetric matrix with the following eigenvectors and eigenvalues, or explain why none exists: (a) $\lambda_1 = 1$, $\mathbf{v}_1 = \left(\frac{3}{5}, \frac{4}{5}\right)^T$, $\lambda_2 = 3$, $\mathbf{v}_2 = \left(-\frac{4}{5}, \frac{3}{5}\right)^T$, (b) $\lambda_1 = -2$, $\mathbf{v}_1 = (1, -1)^T$, $\lambda_2 = 1$, $\mathbf{v}_2 = (1, 1)^T$, (c) $\lambda_1 = 3$, $\mathbf{v}_1 = (2, -1)^T$, $\lambda_2 = -1$, $\mathbf{v}_2 = (-1, 2)^T$, (d) $\lambda_1 = 2$, $\mathbf{v}_1 = (2, 1)^T$, $\lambda_2 = 2$, $\mathbf{v}_2 = (1, 2)^T$.

♡ 8.5.16. (a) Find the eigenvalues and eigenvectors of the matrix $A = \begin{pmatrix} 2 & 1 & -1 \\ 1 & 2 & 1 \\ -1 & 1 & 2 \end{pmatrix}$.

(b) Use the eigenvalues to compute the determinant of A. (c) Is A positive definite? Why or why not? (d) Find an orthonormal eigenvector basis of $\mathbb{R}^3$ determined by A or explain why none exists. (e) Write out the spectral factorization of A if possible. (f) Use orthogonality to write the vector $(1, 0, 0)^T$ as a linear combination of eigenvectors of A.

8.5.17. Use the spectral factorization to diagonalize the following quadratic forms:
(a) $x^2 - 3xy + 5y^2$, (b) $3x^2 + 4xy + 6y^2$, (c) $x^2 + 8xz + y^2 + 6yz + z^2$,
(d) $\frac{3}{2}x^2 - xy - xz + y^2 + z^2$, (e) $6x^2 - 8xy + 2xz + 6y^2 - 2yz + 11z^2$.

♢ 8.5.18. Let $\mathbf{u}_1, \ldots, \mathbf{u}_n$ be an orthonormal basis of $\mathbb{R}^n$. Prove that it forms an eigenvector basis for some symmetric $n \times n$ matrix A. Can you characterize all such matrices?

8.5.19. *True or false*: A matrix with a real orthonormal eigenvector basis is symmetric.

8.5.20. Prove that every quadratic form can be written as $\mathbf{x}^T A\,\mathbf{x} = \|\mathbf{x}\|^2 \left(\sum_{i=1}^{n} \lambda_i \cos^2 \theta_i \right)$, where λ_i are the eigenvalues of A and $\theta_i = \sphericalangle(\mathbf{x}, \mathbf{v}_i)$ denotes the angle between $\mathbf{x}$ and the ith eigenvector.

8.5.21. An elastic body has stress tensor $T = \begin{pmatrix} 3 & 1 & 2 \\ 1 & 3 & 1 \\ 2 & 1 & 3 \end{pmatrix}$. Find the principal stretches and principal directions of stretch.

♢ 8.5.22. Given a solid body spinning around its center of mass, the eigenvectors of its positive definite *inertia tensor* prescribe three mutually orthogonal *principal directions* of rotation, while the corresponding eigenvalues are the *moments of inertia*. Given the inertia tensor $T = \begin{pmatrix} 2 & 1 & 0 \\ 1 & 3 & 1 \\ 0 & 1 & 2 \end{pmatrix}$, find the principal directions and moments of inertia.

♢ 8.5.23. Let K be a positive definite 2×2 matrix. (*a*) Explain why the quadratic equation $\mathbf{x}^T K\,\mathbf{x} = 1$ defines an ellipse. Prove that its principal axes are the eigenvectors of K, and the semi-axes are the reciprocals of the square roots of the eigenvalues.
(*b*) Graph and describe the following curves:

$$(i)\ x^2 + 4y^2 = 1, \qquad (ii)\ x^2 + xy + y^2 = 1, \qquad (iii)\ 3x^2 + 2xy + y^2 = 1.$$

(*c*) What sort of curve(s) does $\mathbf{x}^T K\,\mathbf{x} = 1$ describe if K is not positive definite?

♢ 8.5.24. Let K be a positive definite 3×3 matrix. (*a*) Prove that the quadratic equation $\mathbf{x}^T K\,\mathbf{x} = 1$ defines an ellipsoid in $\mathbb{R}^3$. What are its principal axes and semi-axes?
(*b*) Describe the surface defined by the equation $11x^2 - 8xy + 20y^2 - 10xz + 8yz + 11z^2 = 1$.

8.5.25. Prove that $A = A^T$ has a repeated eigenvalue if and only if it commutes, $AJ = JA$, with a nonzero skew-symmetric matrix: $J^T = -J \neq \mathrm{O}$.
Hint: First prove this when A is a diagonal matrix.

8.5.26. Find all positive definite orthogonal matrices.

♢ 8.5.27. (*a*) Prove that every positive definite matrix K has a unique positive definite *square root*, i.e., a matrix $B > 0$ satisfying $B^2 = K$.
(*b*) Find the positive definite square roots of the following matrices:

$$(i)\ \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}, \quad (ii)\ \begin{pmatrix} 3 & -1 \\ -1 & 1 \end{pmatrix}, \quad (iii)\ \begin{pmatrix} 2 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 9 \end{pmatrix}, \quad (iv)\ \begin{pmatrix} 6 & -4 & 1 \\ -4 & 6 & -1 \\ 1 & -1 & 11 \end{pmatrix}.$$

♢ 8.5.28. *The Polar Decomposition*: Prove that every invertible matrix A has a *polar decomposition*, written $A = QB$, into the product of an orthogonal matrix Q and a positive definite matrix $B > 0$. Show that if $\det A > 0$, then Q is a proper orthogonal matrix. *Hint*: Look at the Gram matrix $K = A^T A$ and use Exercise 8.5.27.

Remark. In mechanics, if A represents the deformation of a body, then Q represents a rotation, while B represents a stretching along the orthogonal eigendirections of K. Thus, every linear deformation of an elastic body can be decomposed into a pure stretching transformation followed by a rotation.

8.5.29. Find the polar decompositions $A = QB$, as defined in Exercise 8.5.28, of the following matrices:

$$(a)\ \begin{pmatrix} 0 & 1 \\ 2 & 0 \end{pmatrix}, \quad (b)\ \begin{pmatrix} 2 & -3 \\ 1 & 6 \end{pmatrix}, \quad (c)\ \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}, \quad (d)\ \begin{pmatrix} 0 & -3 & 8 \\ 1 & 0 & 0 \\ 0 & 4 & 6 \end{pmatrix}, \quad (e)\ \begin{pmatrix} 1 & 0 & 1 \\ 1 & -2 & 0 \\ 1 & 1 & 0 \end{pmatrix}.$$

♢ 8.5.30. *The Spectral Theorem for Hermitian Matrices*. Prove that a complex Hermitian matrix can be factored as $H = U \Lambda U^{\dagger}$ where U is a unitary matrix and Λ is a real diagonal matrix. *Hint*: See Exercises 4.3.25, 8.5.7.

8.5.31. Find the spectral factorization, as in Exercise 8.5.30, of the following Hermitian matrices:

$$(a)\ \begin{pmatrix} 3 & 2\,\mathrm{i} \\ -2\,\mathrm{i} & 6 \end{pmatrix}, \quad (b)\ \begin{pmatrix} 6 & 1-2\,\mathrm{i} \\ 1+2\,\mathrm{i} & 2 \end{pmatrix}, \quad (c)\ \begin{pmatrix} -1 & 5\,\mathrm{i} & -4 \\ -5\,\mathrm{i} & -1 & 4\,\mathrm{i} \\ -4 & -4\,\mathrm{i} & 8 \end{pmatrix}.$$

♡ 8.5.32. *The Spectral Decomposition*: Let A be a symmetric matrix with distinct eigenvalues $\lambda_1, \ldots, \lambda_k$. Let $V_j = \ker(A - \lambda_j \mathrm{I})$ denote the eigenspace corresponding to λ_j, and let P_j be the orthogonal projection matrix onto V_j, as defined in Exercise 4.4.9. (*i*) Prove that the spectral factorization (8.35) can be rewritten as

$$A = \lambda_1 P_1 + \lambda_2 P_2 + \cdots + \lambda_k P_k, \tag{8.37}$$

expressing A as a linear combination of projection matrices. (*ii*) Write out the *spectral decomposition* (8.37) for the matrices in Exercise 8.5.13. (*iii*) Show that

$$\mathrm{I} = P_1 + P_2 + \cdots + P_k, \qquad \text{while} \qquad P_i^2 = P_i, \qquad P_i P_j = \mathrm{O} \quad \text{for} \quad i \neq j.$$

(*iv*) Show that if $p(t)$ is any polynomial, then

$$p(A) = p(\lambda_1) P_1 + p(\lambda_2) P_2 + \cdots + p(\lambda_k) P_k. \tag{8.38}$$

Remark. Replacing $p(t)$ by any function $f(t)$ allows one to define $f(A)$ for any symmetric matrix A.

Optimization Principles for Eigenvalues of Symmetric Matrices

As we learned in Chapter 5, the solution to a linear system with positive definite coefficient matrix can be characterized by a minimization principle. Thus, it should come as no surprise that eigenvalues of positive definite matrices, and even more general symmetric matrices, can also be characterized by some sort of optimization procedure. A number of basic numerical algorithms for computing eigenvalues of matrices are based on such optimization principles.

First, consider the relatively simple case of a real diagonal matrix $\Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_n)$. We assume that the diagonal entries, which are the *same* as the eigenvalues, appear in decreasing order,

$$\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n, \tag{8.39}$$

so λ_1 is the largest eigenvalue, while λ_n is the smallest. The effect of Λ on a vector $\mathbf{y} = (y_1, y_2, \ldots, y_n)^T \in \mathbb{R}^n$ is to multiply its entries by the diagonal eigenvalues: $\Lambda \mathbf{y} = (\lambda_1 y_1, \lambda_2 y_2, \ldots, \lambda_n y_n)^T$. In other words, the linear transformation represented by the coefficient matrix Λ has the effect of stretching† the ith coordinate direction by the factor λ_i. In particular, the maximal stretch occurs in the $\mathbf{e}_1$ direction, with factor λ_1, while the minimal (or largest negative) stretch occurs in the $\mathbf{e}_n$ direction, with factor λ_n. The germ of the optimization principles for characterizing the extreme eigenvalues is contained in this geometrical observation.

Let us turn our attention to the associated quadratic form

$$q(\mathbf{y}) = \mathbf{y}^T \Lambda \mathbf{y} = \lambda_1 y_1^2 + \lambda_2 y_2^2 + \cdots + \lambda_n y_n^2. \tag{8.40}$$

† If $\lambda_i < 0$, then the effect is to stretch and reflect.

Note that $q(t\,\mathbf{e}_1) = \lambda_1 t^2$, and hence if $\lambda_1 > 0$, then $q(\mathbf{y})$ has no maximum; on the other hand, if $\lambda_1 \leq 0$, so all eigenvalues are non-positive, then $q(\mathbf{y}) \leq 0$ for all $\mathbf{y}$, and its maximal value is $q(\mathbf{0}) = 0$. Thus, in either case, a strict maximization of $q(\mathbf{y})$ does not tell us anything of importance.

Suppose, however, that we continue in our quest to maximize $q(\mathbf{y})$, but now restrict $\mathbf{y}$ to be a unit vector (in the Euclidean norm), so

$$\|\mathbf{y}\|^2 = y_1^2 + \cdots + y_n^2 = 1.$$

In view of (8.39),

$$q(\mathbf{y}) = \lambda_1 y_1^2 + \lambda_2 y_2^2 + \cdots + \lambda_n y_n^2 \leq \lambda_1 y_1^2 + \lambda_1 y_2^2 + \cdots + \lambda_1 y_n^2 = \lambda_1 \left(y_1^2 + \cdots + y_n^2 \right) = \lambda_1.$$

Moreover, $q(\mathbf{e}_1) = \lambda_1$. We conclude that the maximal value of $q(\mathbf{y})$ over all unit vectors *is* the largest eigenvalue of Λ:

$$\lambda_1 = \max \{ \, q(\mathbf{y}) \mid \|\mathbf{y}\| = 1 \, \}.$$

By the same reasoning, its minimal value equals the smallest eigenvalue:

$$\lambda_n = \min \{ \, q(\mathbf{y}) \mid \|\mathbf{y}\| = 1 \, \}.$$

Thus, we can characterize the two extreme eigenvalues by optimization principles, albeit of a slightly different character from what we dealt with in Chapter 5.

Now suppose A is any symmetric matrix. We use its spectral factorization (8.35) to diagonalize the associated quadratic form

$$q(\mathbf{x}) = \mathbf{x}^T A\,\mathbf{x} = \mathbf{x}^T Q\,\Lambda\,Q^T\,\mathbf{x} = \mathbf{y}^T \Lambda\,\mathbf{y}, \qquad \text{where} \qquad \mathbf{y} = Q^T\mathbf{x} = Q^{-1}\mathbf{x},$$

as in (8.36). According to the preceding discussion, the maximum of $\mathbf{y}^T\Lambda\,\mathbf{y}$ over all unit vectors $\|\mathbf{y}\| = 1$ is the largest eigenvalue λ_1 of Λ, which is the *same* as the largest eigenvalue of A. Moreover, since Q is an orthogonal matrix, Proposition 7.24 tell us that it maps unit vectors to unit vectors:

$$1 = \|\mathbf{y}\| = \|Q^T\mathbf{x}\| = \|\mathbf{x}\|,$$

and so the maximum of $q(\mathbf{x})$ over all unit vectors $\|\mathbf{x}\| = 1$ is the same maximum eigenvalue λ_1. Similar reasoning applies to the smallest eigenvalue λ_n. In this fashion, we have established the basic optimization principles for the extreme eigenvalues of a symmetric matrix.

Theorem 8.40. Let A be a symmetric matrix with real eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$. Then

$$\lambda_1 = \max \left\{ \, \mathbf{x}^T A\,\mathbf{x} \;\middle|\; \|\mathbf{x}\| = 1 \, \right\}, \qquad \lambda_n = \min \left\{ \, \mathbf{x}^T A\,\mathbf{x} \;\middle|\; \|\mathbf{x}\| = 1 \, \right\}, \tag{8.41}$$

are, respectively its largest and smallest eigenvalues. The maximal value is achieved when $\mathbf{x} = \pm\,\mathbf{u}_1$ is one of the unit eigenvectors associated with the largest eigenvalue λ_1; similarly, the minimal value occurs at $\mathbf{x} = \pm\,\mathbf{u}_n$, a unit eigenvector for the smallest eigenvalue λ_n.

Remark. In multivariable calculus, the eigenvalue λ plays the role of a *Lagrange multiplier* for the constrained optimization problem, [**2**, **78**, **79**].

Example 8.41. The problem is to maximize the value of the quadratic form

$$q(x, y) = 3\,x^2 + 2\,x\,y + 3\,y^2$$

for all x, y lying on the unit circle $x^2 + y^2 = 1$. This maximization problem is precisely of the form (8.41). The symmetric coefficient matrix for the quadratic form is $A = \begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix}$, whose eigenvalues are, according to Example 8.5, $\lambda_1 = 4$ and $\lambda_2 = 2$. Theorem 8.40 implies that the maximum is the largest eigenvalue, and hence equal to 4, while its minimum is the smallest eigenvalue, and hence equal to 2. Thus, evaluating $q(x, y)$ on the unit eigenvectors, we conclude that

$$q\left(-\tfrac{1}{\sqrt{2}}, \tfrac{1}{\sqrt{2}}\right) = 2 \le q(x, y) \le 4 = q\left(-\tfrac{1}{\sqrt{2}}, \tfrac{1}{\sqrt{2}}\right) \qquad \text{for all} \qquad x^2 + y^2 = 1.$$

In practical applications, the restriction of the quadratic form to unit vectors may not be particularly convenient. We can, however, rephrase the eigenvalue optimization principles in a form that utilizes general nonzero vectors. If $\mathbf{v} \neq \mathbf{0}$, then $\mathbf{x} = \mathbf{v}/\|\mathbf{v}\|$ is a unit vector. Substituting this expression for $\mathbf{x}$ in the quadratic form $q(\mathbf{x}) = \mathbf{x}^T A\,\mathbf{x}$ leads to the following optimization principles for the extreme eigenvalues of a symmetric matrix:

$$\lambda_1 = \max\left\{ \frac{\mathbf{v}^T A\,\mathbf{v}}{\|\mathbf{v}\|^2} \;\middle|\; \mathbf{v} \neq \mathbf{0} \right\}, \qquad \lambda_n = \min\left\{ \frac{\mathbf{v}^T A\,\mathbf{v}}{\|\mathbf{v}\|^2} \;\middle|\; \mathbf{v} \neq \mathbf{0} \right\}. \tag{8.42}$$

Thus, we replace optimization of a quadratic polynomial over the unit sphere by optimization of a rational function over all of $\mathbb{R}^n \setminus \{\mathbf{0}\}$. The rational function being optimized is called a *Rayleigh quotient*, after Lord Rayleigh, a prominent nineteenth-century British scientist. For instance, referring back to Example 8.41, the maximum value of

$$r(x, y) = \frac{3x^2 + 2xy + 3y^2}{x^2 + y^2} \qquad \text{for all} \qquad \begin{pmatrix} x \\ y \end{pmatrix} \neq \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

is equal to 4, the same maximal eigenvalue of the corresponding coefficient matrix.

What about characterizing one of the intermediate eigenvalues? Then we need to be a little more sophisticated in designing the optimization principle. To motivate the construction, look first at the diagonal case. If we restrict the quadratic form (8.40) to vectors $\widetilde{\mathbf{y}} = (\,0, y_2, \ldots, y_n\,)^T$ whose first component is zero, we obtain

$$q(\,\widetilde{\mathbf{y}}\,) = q(0, y_2, \ldots, y_n) = \lambda_2 y_2^2 + \;\cdots\; + \lambda_n y_n^2.$$

The maximum value of $q(\,\widetilde{\mathbf{y}}\,)$ over all such $\widetilde{\mathbf{y}}$ of norm 1 is, by the same reasoning, the second largest eigenvalue λ_2. Moreover, $\widetilde{\mathbf{y}} \cdot \mathbf{e}_1 = 0$, and so $\widetilde{\mathbf{y}}$ can be characterized as being orthogonal to the first standard basis vector, which also happens to be the eigenvector of Λ corresponding to the maximal eigenvalue λ_1. Thus, to find the second eigenvalue, we maximize the quadratic form over all unit vectors that are orthogonal to the first eigenvector. Similarly, if we want to find the jth largest eigenvalue λ_j, we maximize $q(\,\widehat{\mathbf{y}}\,)$ over all unit vectors $\widehat{\mathbf{y}}$ whose first $j-1$ components vanish, $y_1 = \cdots = y_{j-1} = 0$, or, stated geometrically, over all vectors $\widehat{\mathbf{y}}$ such that $\|\widehat{\mathbf{y}}\| = 1$ and $\widehat{\mathbf{y}} \cdot \mathbf{e}_1 = \cdots = \widehat{\mathbf{y}} \cdot \mathbf{e}_{j-1} = 0$, that is, over all unit vectors orthogonal to the first $j-1$ eigenvectors of Λ.

A similar reasoning based on the Spectral Theorem 8.38 and the orthogonality of eigenvectors of symmetric matrices leads to the general result.

Theorem 8.42. Let A be a symmetric matrix with eigenvalues $\lambda_1 \ge \lambda_2 \ge \cdots \ge \lambda_n$ and corresponding orthogonal eigenvectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$. Then the maximal value of the quadratic form $q(\mathbf{x}) = \mathbf{x}^T A\,\mathbf{x}$ over all unit vectors that are orthogonal to the first $j-1$ eigenvectors is its jth eigenvalue:

$$\lambda_j = \max\left\{\, \mathbf{x}^T A\,\mathbf{x} \;\middle|\; \|\mathbf{x}\| = 1, \;\; \mathbf{x} \cdot \mathbf{v}_1 = \cdots = \mathbf{x} \cdot \mathbf{v}_{j-1} = 0 \,\right\}. \tag{8.43}$$

Thus, at least in principle, one can compute the eigenvalues and eigenvectors of a symmetric matrix by the following recursive procedure. First, find the largest eigenvalue λ_1 by the basic maximization principle (8.41) and its associated eigenvector $\mathbf{v}_1$ by solving the eigenvector system (8.13). The next largest eigenvalue λ_2 is then characterized by the constrained maximization principle (8.43), and so on. Although of theoretical interest, this algorithm is of somewhat limited value in numerical computations. Practical numerical methods for computing eigenvalues will be developed in Sections 9.5 and 9.6.

Exercises

8.5.33. Find the minimum and maximum values of the quadratic form $5x^2 + 4xy + 5y^2$ where x, y are subject to the constraint $x^2 + y^2 = 1$.

8.5.34. Find the minimum and maximum values of the quadratic form
$$2x^2 + xy + 2xz + 2y^2 + 2z^2 \text{ where } x, y, z \text{ are required to satisfy } x^2 + y^2 + z^2 = 1.$$

8.5.35. What is the minimum and maximum values of the following rational functions:
$$(a)\ \frac{3x^2 - 2y^2}{x^2 + y^2},\ (b)\ \frac{x^2 - 3xy + y^2}{x^2 + y^2},\ (c)\ \frac{3x^2 + xy + 5y^2}{x^2 + y^2},\ (d)\ \frac{2x^2 + xy + 3xz + 2y^2 + 2z^2}{x^2 + y^2 + z^2}.$$

8.5.36. Find the minimum and maximum values of $q(\mathbf{x}) = \sum_{i=1}^{n-1} x_i x_{i+1}$ for $\|\mathbf{x}\|^2 = 1$.
Hint: See Exercise 8.2.47.

8.5.37. Write down and solve an optimization principle characterizing the largest and smallest eigenvalues of the following positive definite matrices:
$$(a)\ \begin{pmatrix} 2 & -1 \\ -1 & 3 \end{pmatrix},\quad (b)\ \begin{pmatrix} 4 & 1 \\ 1 & 4 \end{pmatrix},\quad (c)\ \begin{pmatrix} 6 & -4 & 1 \\ -4 & 6 & -1 \\ 1 & -1 & 11 \end{pmatrix},\quad (d)\ \begin{pmatrix} 4 & -1 & -2 \\ -1 & 4 & -1 \\ -2 & -1 & 4 \end{pmatrix}.$$

8.5.38. Write down a maximization principle that characterizes the middle eigenvalue of the matrices in parts (c) and (d) of Exercise 8.5.37.

8.5.39. Given a 3×3 symmetric matrix, formulate two distinct ways of characterizing its middle eigenvalue λ_2.

8.5.40. Suppose $K > 0$. What is the maximum value of $q(\mathbf{x}) = \mathbf{x}^T K \mathbf{x}$ when $\mathbf{x}$ is constrained to a sphere of radius $\|\mathbf{x}\| = r$?

8.5.41. Let $K > 0$. Prove the product formula
$$\max\left\{ \mathbf{x}^T K \mathbf{x} \;\middle|\; \|\mathbf{x}\| = 1 \right\} \min\left\{ \mathbf{x}^T K^{-1} \mathbf{x} \;\middle|\; \|\mathbf{x}\| = 1 \right\} = 1.$$

♢ 8.5.42. Write out the details in the proof of Theorem 8.42.

♢ 8.5.43. Reformulate Theorem 8.42 as a minimum principle for intermediate eigenvalues.

8.5.44. Under the set-up of Theorem 8.42, explain why
$$\lambda_j = \max\left\{ \frac{\mathbf{v}^T K \mathbf{v}}{\|\mathbf{v}\|^2} \;\middle|\; \mathbf{v} \neq \mathbf{0}, \quad \mathbf{v} \cdot \mathbf{v}_1 = \cdots = \mathbf{v} \cdot \mathbf{v}_{j-1} = 0 \right\}.$$

♡ 8.5.45. (a) Let K, M be positive definite $n \times n$ matrices and $\lambda_1 \geq \cdots \geq \lambda_n$ be their generalized eigenvalues, as in Exercise 8.5.9. Prove that that the largest generalized eigenvalue can be characterized by the maximum principle $\lambda_1 = \max\{ \mathbf{x}^T K \mathbf{x} \mid \mathbf{x}^T M \mathbf{x} = 1 \}$. *Hint*: Use Exercise 8.5.27. (b) Prove the alternative maximum principle $\lambda_1 = \max\left\{ \dfrac{\mathbf{x}^T K \mathbf{x}}{\mathbf{x}^T M \mathbf{x}} \;\middle|\; \mathbf{x} \neq \mathbf{0} \right\}$. (c) How would you characterize the smallest generalized eigenvalue? (d) An intermediate generalized eigenvalue?

8.5.46. Use Exercise 8.5.45 to find the minimum and maximum of the rational functions

(a) $\dfrac{3x^2+2y^2}{4x^2+5y^2}$, (b) $\dfrac{x^2-xy+2y^2}{2x^2-xy+y^2}$, (c) $\dfrac{2x^2+3y^2+z^2}{x^2+3y^2+2z^2}$, (d) $\dfrac{2x^2+6xy+11y^2+6yz+2z^2}{x^2+2xy+3y^2+2yz+z^2}$.

8.5.47. Let A be a complete square matrix, not necessarily symmetric, with all positive eigenvalues. Is the associated quadratic form $q(\mathbf{x}) = \mathbf{x}^T A\,\mathbf{x} > 0$ for all $\mathbf{x} \neq \mathbf{0}$?

8.6 Incomplete Matrices

Unfortunately, not all square matrices are complete. Matrices that do not have enough (complex) eigenvectors to form a basis are considerably less pleasant to work with. However, since they occasionally appear in applications, it is worth learning how to handle them. There are two approaches: the first, named after the twentieth-century Russian/German mathematician Issai Schur, is a generalization of the spectral theorem, and converts an arbitrary square matrix into a similar, upper triangular matrix with the eigenvalues along the diagonal. Thus, although not every matrix can be diagonalized, they can all be "triangularized". Applications of the Schur decomposition, including the numerical computation of eigenvalues, can be found in [**21**].

The second approach, due to the nineteenth-century French mathematician Camille Jordan,† shows how to supplement the eigenvectors of an incomplete matrix in order to obtain a basis in which the matrix assumes a simple, but now non-diagonal, canonical (meaning distinguished) form. Applications of the Jordan canonical form will appear in our study of linear systems of ordinary differential equations. We remark that the two subsections are completely independent of one another.

The Schur Decomposition

As noted above, the Schur decomposition is used to convert a square matrix into similar upper triangular matrix. The similarity transformation can be chosen to be represented by a unitary matrix — a complex generalization of an orthogonal matrix.

Definition 8.43. A complex, square matrix U is called *unitary* if it satisfies

$$U^\dagger U = \mathrm{I}, \qquad \text{where} \qquad U^\dagger = \overline{U^T} \tag{8.44}$$

denotes the *Hermitian transpose*, in which one first transposes and then takes complex conjugates of all entries.

Thus, U is unitary if and only if its inverse equals its Hermitian transpose: $U^{-1} = U^\dagger$. For example, $U = \begin{pmatrix} \frac{1}{\sqrt{2}}\,\mathrm{i} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}}\,\mathrm{i} \end{pmatrix}$ is unitary, since $U^{-1} = U^\dagger = \begin{pmatrix} -\frac{1}{\sqrt{2}}\,\mathrm{i} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}}\,\mathrm{i} \end{pmatrix}$.

The (i,j) entry of the defining equation (8.44) is the Hermitian dot product between the ith and j^{th} columns of U, and hence U is an $n \times n$ unitary matrix if and only if its columns form an orthonormal basis of $\mathbb{C}^n$. In particular, a real matrix is unitary if and only if it is orthogonal. The next result is proved in the same fashion as Proposition 4.23.

† No relation to Wilhelm Jordan of Gauss–Jordan fame.

Proposition 8.44. If U_1 and U_2 are $n \times n$ unitary matrices, so is their product $U_1 U_2$.

The *Schur decomposition* states that every square matrix is unitarily similar to an upper triangular matrix. The method of proof provides a recursive algorithm for constructing the decomposition.

Theorem 8.45. Let A be an $n \times n$ matrix, either real or complex. Then there exists a unitary matrix U and an upper triangular matrix Δ such that

$$A = U \Delta U^{\dagger} = U \Delta U^{-1}. \tag{8.45}$$

The diagonal entries of Δ are the eigenvalues of A.

In particular, if all eigenvalues of A are real, then $U = Q$ can be chosen to be a (real) orthogonal matrix, and Δ is also a real matrix. This follows from the construction outlined in the proof.

Warning. The Schur decomposition (8.45) is not unique. As the method of proof makes clear, there are many inequivalent choices for the matrices U and Δ.

Proof of Theorem 8.45: The proof proceeds by induction on the size of A. According to Theorem 8.11, A has at least one, possibly complex, eigenvalue λ_1. Let $\mathbf{u}_1 \in \mathbb{C}^n$ be a corresponding unit eigenvector, so its Hermitian norm is $\| \mathbf{u}_1 \| = 1$. Let U_1 be an $n \times n$ unitary matrix whose first column is the unit eigenvector $\mathbf{u}_1$. In practice, U_1 can be constructed by applying the Gram–Schmidt process to any basis of $\mathbb{C}^n$ whose first element is the eigenvector $\mathbf{u}_1$. The eigenvector equation $A \mathbf{u}_1 = \lambda_1 \mathbf{u}_1$ forms the first column of the matrix product equation

$$A U_1 = U_1 B, \qquad \text{where} \qquad B = U_1^{\dagger} A U_1 = \begin{pmatrix} \lambda_1 & \mathbf{r} \\ \mathbf{0} & C \end{pmatrix},$$

with C an $(n-1) \times (n-1)$ matrix and $\mathbf{r}$ a row vector. By our induction hypothesis, there is an $(n-1) \times (n-1)$ unitary matrix V such that $V^{\dagger} C V = \Gamma$ is an upper triangular $(n-1) \times (n-1)$ matrix. Set $U_2 = \begin{pmatrix} 1 & \mathbf{0} \\ \mathbf{0} & V \end{pmatrix}$. It is easily checked that U_2 is also unitary, and, moreover,

$$U_2^{\dagger} B U_2 = \begin{pmatrix} \lambda_1 & \mathbf{s} \\ \mathbf{0} & \Gamma \end{pmatrix} = \Delta$$

is upper triangular. The unitary product matrix $U = U_1 U_2$ yields the desired result:

$$U^{\dagger} A U = (U_1 U_2)^{\dagger} A U_1 U_2 = U_2^{\dagger} U_1^{\dagger} A U_1 U_2 = U_2^{\dagger} B U_2 = \Delta,$$

which establishes the Schur decomposition formula (8.45). Finally, since A and Δ are similar matrices, they have the same eigenvalues, which justifies the final statement of the theorem. *Q.E.D.*

Example 8.46. The matrix $A = \begin{pmatrix} 6 & 4 & -3 \\ -4 & -2 & 2 \\ 4 & 4 & -2 \end{pmatrix}$ has a simple eigenvalue $\lambda_1 = 2$, with eigenvector $\mathbf{v}_1 = (\,-1, 1, 0\,)^T$, and a double eigenvalue $\lambda_2 = 0$, with only one independent eigenvector $\mathbf{v}_2 = (\,1, 0, 2\,)^T$. Thus A is incomplete, and so not diagonalizable. To construct a Schur decomposition, we begin with the first eigenvector $\mathbf{v}_1$ and apply the Gram–Schmidt

process to the basis $\mathbf{v}_1, \mathbf{e}_2, \mathbf{e}_3$ to obtain the orthogonal matrix $U_1 = \begin{pmatrix} -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 \\ 0 & 0 & 1 \end{pmatrix}$. The resulting similar matrix† $B = U_1^T A U_1 = \begin{pmatrix} 2 & -8 & \frac{\sqrt{5}}{2} \\ 0 & 2 & -\frac{1}{\sqrt{2}} \\ 0 & 4\sqrt{2} & -2 \end{pmatrix}$ has its first column in upper triangular form. To continue the procedure, we extract the lower 2×2 submatrix $C = \begin{pmatrix} 2 & -\frac{1}{\sqrt{2}} \\ 4\sqrt{2} & -2 \end{pmatrix}$, and find that it has a single (incomplete) eigenvalue 0, with unit eigenvector $\begin{pmatrix} \frac{1}{3} \\ \frac{2\sqrt{2}}{3} \end{pmatrix}$. The corresponding orthogonal matrix $V = \begin{pmatrix} \frac{1}{3} & \frac{2\sqrt{2}}{3} \\ \frac{2\sqrt{2}}{3} & -\frac{1}{3} \end{pmatrix}$ will convert C to upper triangular form $V^T C V = \begin{pmatrix} 0 & \frac{9}{\sqrt{2}} \\ 0 & 0 \end{pmatrix}$. Therefore, $U_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{2\sqrt{2}}{3} \\ 0 & \frac{2\sqrt{2}}{3} & -\frac{1}{3} \end{pmatrix}$ will complete the conversion of the original matrix into upper triangular form

$$\Delta = U_2^T B U_2 = U^T A U = \begin{pmatrix} 2 & \frac{2}{3} & -\frac{37}{3\sqrt{2}} \\ 0 & 0 & \frac{9}{\sqrt{2}} \\ 0 & 0 & 0 \end{pmatrix}, \quad \text{where} \quad U = U_1 U_2 = \begin{pmatrix} -\frac{1}{\sqrt{2}} & \frac{1}{3\sqrt{2}} & \frac{2}{3} \\ \frac{1}{\sqrt{2}} & \frac{1}{3\sqrt{2}} & \frac{2}{3} \\ 0 & \frac{2\sqrt{2}}{3} & -\frac{1}{3} \end{pmatrix}$$

is the desired orthogonal (unitary) matrix. Use of a computer to carry out the detailed calculations is essential in most examples.

Exercises

8.6.1. Establish a Schur Decomposition for the matrices (*a*) $\begin{pmatrix} 1 & -1 \\ 1 & 3 \end{pmatrix}$, (*b*) $\begin{pmatrix} 1 & -2 \\ -2 & 1 \end{pmatrix}$, (*c*) $\begin{pmatrix} 8 & 9 \\ -6 & -7 \end{pmatrix}$, (*d*) $\begin{pmatrix} 1 & 5 \\ -2 & -1 \end{pmatrix}$, (*e*) $\begin{pmatrix} 2 & -1 & 2 \\ -2 & 3 & -1 \\ -6 & 6 & -5 \end{pmatrix}$, (*f*) $\begin{pmatrix} 0 & 2 & -1 \\ -1 & -1 & 1 \\ -1 & 0 & 0 \end{pmatrix}$.

8.6.2. Show that a real unitary matrix is an orthogonal matrix.

♢ 8.6.3. Prove Proposition 8.44.

♢ 8.6.4. Write out a new proof of the Spectral Theorem 8.38 based on the Schur Decomposition.

♡ 8.6.5. A complex matrix A is called *normal* if it commutes with its Hermitian transpose: $A^\dagger A = A A^\dagger$. (*a*) Show that every real symmetric matrix is normal. (*b*) Show that every unitary matrix is normal. (*c*) Show that every real orthogonal matrix is normal. (*d*) Show that an upper triangular matrix is normal if and only if it is diagonal. (*e*) Show that the eigenvectors of a normal matrix form an orthogonal basis of $\mathbb{C}^n$ under the Hermitian dot product. (*f*) Show that the converse is true: a matrix has an orthogonal eigenvector basis of $\mathbb{C}^n$ if and only if it is normal. *Hint*: Use the Schur Decomposition. (*g*) How can you tell when a real matrix has a real orthonormal eigenvector basis?

† Since all matrices are real in this example, the Hermitian transpose † reduces to the ordinary transpose ^{T}.

The Jordan Canonical Form

We now turn to the more sophisticated Jordan canonical form. Throughout this section, A will be an $n \times n$ matrix, with either real or complex entries. We let $\lambda_1, \ldots, \lambda_k$ denote the *distinct* eigenvalues of A, so $\lambda_i \neq \lambda_j$ for $i \neq j$. We recall that Theorem 8.11 guarantees that every matrix has at least one (complex) eigenvalue, so $k \geq 1$. Moreover, we are assuming that $k < n$, since otherwise A would be complete.

Definition 8.47. Let A be a square matrix. A *Jordan chain* of length j for A is a sequence of non-zero vectors $\mathbf{w}_1, \ldots, \mathbf{w}_j \in \mathbb{C}^n$ that satisfies

$$A\mathbf{w}_1 = \lambda \mathbf{w}_1, \qquad A\mathbf{w}_i = \lambda \mathbf{w}_i + \mathbf{w}_{i-1}, \qquad i = 2, \ldots, j, \tag{8.46}$$

where λ is an eigenvalue of A. A Jordan chain associated with a zero eigenvalue, which requires that A be singular, is called a *null Jordan chain*, and satisfies

$$A\mathbf{w}_1 = \mathbf{0}, \qquad A\mathbf{w}_i = \mathbf{w}_{i-1}, \qquad i = 2, \ldots, j. \tag{8.47}$$

Note that the initial vector $\mathbf{w}_1$ in a Jordan chain is a genuine eigenvector, and so Jordan chains exist only when λ is an eigenvalue. The rest, $\mathbf{w}_2, \ldots, \mathbf{w}_j$, are *generalized eigenvectors*, in accordance with the following definition.

Definition 8.48. A nonzero vector $\mathbf{w} \neq \mathbf{0}$ such that

$$(A - \lambda\, \mathrm{I})^k \mathbf{w} = \mathbf{0} \tag{8.48}$$

for some $k > 0$ and $\lambda \in \mathbb{C}$ is called a *generalized eigenvector*[†] of the matrix A.

Note that every ordinary eigenvector is automatically a generalized eigenvector, since we can just take $k = 1$ in (8.48); but the converse is not necessarily valid. We shall call the minimal value of k for which (8.48) holds the *index* of the generalized eigenvector. Thus, an ordinary eigenvector is a generalized eigenvector of index 1. Since $A - \lambda\, \mathrm{I}$ is nonsingular whenever λ is not an eigenvalue of A, its kth power $(A - \lambda\, \mathrm{I})^k$ is also nonsingular. Therefore, generalized eigenvectors can exist only when λ is an ordinary eigenvalue of A — there are no additional "generalized eigenvalues".

Lemma 8.49. The ith vector $\mathbf{w}_i$ in a Jordan chain (8.46) is a generalized eigenvector of index i.

Proof: By definition, $(A - \lambda\, \mathrm{I})\mathbf{w}_1 = \mathbf{0}$, and so $\mathbf{w}_1$ is an eigenvector. Next, we have $(A - \lambda\, \mathrm{I})\mathbf{w}_2 = \mathbf{w}_1 \neq \mathbf{0}$, while $(A - \lambda\, \mathrm{I})^2 \mathbf{w}_2 = (A - \lambda\, \mathrm{I})\mathbf{w}_1 = \mathbf{0}$. Thus, $\mathbf{w}_2$ is a generalized eigenvector of index 2. A simple induction proves that $(A - \lambda\, \mathrm{I})^{i-1}\mathbf{w}_i = \mathbf{w}_1 \neq \mathbf{0}$ while $(A - \lambda\, \mathrm{I})^i \mathbf{w}_i = \mathbf{0}$. *Q.E.D.*

Example 8.50. Consider the 3×3 matrix $A = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 1 \\ 0 & 0 & 2 \end{pmatrix}$. The only eigenvalue is $\lambda = 2$, and $A - 2\,\mathrm{I} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}$. We claim that the standard basis vectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$

---

† Despite the common terminology, this is *not* the same concept as developed in Exercise 8.5.8.

form a Jordan chain. Indeed, $A\mathbf{e}_1 = 2\mathbf{e}_1$, and hence $\mathbf{e}_1 \in \ker(A - 2\,\mathrm{I})$ is a genuine eigenvector. Furthermore, $A\mathbf{e}_2 = 2\mathbf{e}_2 + \mathbf{e}_1$, and $A\mathbf{e}_3 = 2\mathbf{e}_3 + \mathbf{e}_2$, as you can easily check. Thus, $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ satisfy the Jordan chain equations (8.46) for the eigenvalue $\lambda = 2$. Note that $\mathbf{e}_2$ lies in the kernel of $(A - 2\,\mathrm{I})^2 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$, and so is a generalized eigenvector of index 2. Indeed, every vector of the form $\mathbf{w} = a\mathbf{e}_1 + b\mathbf{e}_2$ with $b \neq 0$ is a generalized eigenvector of index 2. (When $b = 0$, $a \neq 0$, the vector $\mathbf{w} = a\mathbf{e}_1$ is an ordinary eigenvector of index 1.) Finally, $(A - 2\,\mathrm{I})^3 = \mathrm{O}$, and so every nonzero vector $\mathbf{0} \neq \mathbf{v} \in \mathbb{R}^3$, including $\mathbf{e}_3$, is a generalized eigenvector of index 3 or less.

A basis of $\mathbb{R}^n$ or $\mathbb{C}^n$ is called a *Jordan basis* for the matrix A if it consists of one or more Jordan chains that have no elements in common. Thus, for the matrix in Example 8.50, the standard basis $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ is, in fact, a Jordan basis. An eigenvector basis, if such exists, qualifies as a Jordan basis, since each eigenvector belongs to a Jordan chain of length 1. Jordan bases are the desired extension of eigenvector bases, and every square matrix has one. The proof of the following *Jordan Basis Theorem* will appear below.

Theorem 8.51. Every $n \times n$ matrix admits a Jordan basis of $\mathbb{C}^n$. The first elements of the Jordan chains form a maximal set of linearly independent eigenvectors. Moreover, the number of generalized eigenvectors in the Jordan basis that belong to the Jordan chains associated with the eigenvalue λ is the same as the eigenvalue's multiplicity.

Example 8.52. Consider the 5×5 matrix $A = \begin{pmatrix} -1 & 0 & 1 & 0 & 0 \\ -2 & 2 & -4 & 1 & 1 \\ -1 & 0 & -3 & 0 & 0 \\ -4 & -1 & 3 & 1 & 0 \\ 4 & 0 & 2 & -1 & 0 \end{pmatrix}$. Its characteristic equation is found to be

$$p_A(\lambda) = \det(A - \lambda\,\mathrm{I}) = \lambda^5 + \lambda^4 - 5\lambda^3 - \lambda^2 + 8\lambda - 4 = (\lambda - 1)^3(\lambda + 2)^2 = 0,$$

and hence A has two eigenvalues: $\lambda_1 = 1$, which is a triple eigenvalue, and $\lambda_2 = -2$, which is double. Solving the associated homogeneous systems $(A - \lambda_j\,\mathrm{I})\mathbf{v} = \mathbf{0}$, we find that, up to constant multiple, there are only two eigenvectors: $\mathbf{v}_1 = (0, 0, 0, -1, 1)^T$ for $\lambda_1 = 1$ and, anticipating our final numbering, $\mathbf{v}_4 = \big(-1, 1, 1, -2, 0\big)^T$ for $\lambda_2 = -2$. Thus, A is far from complete.

To construct a Jordan basis, we first note that since A has 2 linearly independent eigenvectors, the Jordan basis will contain two Jordan chains: the one associated with the triple eigenvalue $\lambda_1 = 1$ will have length 3, while $\lambda_2 = -2$ leads to a Jordan chain of length 2. To construct the former, we need to first solve the system $(A - \mathrm{I})\mathbf{w} = \mathbf{v}_1$. Note that the coefficient matrix is singular — it must be, since 1 is an eigenvalue — and the general solution is $\mathbf{w} = \mathbf{v}_2 + t\,\mathbf{v}_1$, where $\mathbf{v}_2 = (0, 1, 0, 0, -1)^T$, and t is the free variable. The appearance of an arbitrary multiple of the eigenvector $\mathbf{v}_1$ in the solution is not unexpected; indeed, the kernel of $A - \mathrm{I}$ is the eigenspace for $\lambda_1 = 1$. We can choose any solution, e.g., $\mathbf{v}_2$ as the second element in the Jordan chain. We then solve $(A - \mathrm{I})\mathbf{w} = \mathbf{v}_2$ for $\mathbf{w} = \mathbf{v}_3 + t\,\mathbf{v}_1$, where $\mathbf{v}_3 = (0, 0, 0, 1, 0)^T$ can be used as the final element of this Jordan chain. Similarly, to construct the Jordan chain for the second eigenvalue, we solve $(A + 2\,\mathrm{I})\mathbf{w} = \mathbf{v}_4$ for $\mathbf{w} = \mathbf{v}_5 + t\,\mathbf{v}_4$, where $\mathbf{v}_5 = (-1, 0, 0, -2, 1)^T$.

Thus, the desired Jordan basis is

$$\mathbf{v}_1 = \begin{pmatrix} 0 \\ 0 \\ 0 \\ -1 \\ 1 \end{pmatrix}, \qquad \mathbf{v}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ -1 \end{pmatrix}, \qquad \mathbf{v}_3 = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}, \qquad \mathbf{v}_4 = \begin{pmatrix} -1 \\ 1 \\ 1 \\ -2 \\ 0 \end{pmatrix}, \qquad \mathbf{v}_5 = \begin{pmatrix} -1 \\ 0 \\ 0 \\ -2 \\ 1 \end{pmatrix},$$

with $A\mathbf{v}_1 = \mathbf{v}_1$, $A\mathbf{v}_2 = \mathbf{v}_2 + \mathbf{v}_1$, $A\mathbf{v}_3 = \mathbf{v}_3 + \mathbf{v}_2$, $A\mathbf{v}_4 = -2\mathbf{v}_4$, $A\mathbf{v}_5 = -2\mathbf{v}_5 + \mathbf{v}_4$.

Just as an eigenvector basis diagonalizes a complete matrix, a Jordan basis provides a particularly simple form for an incomplete matrix, known as the *Jordan canonical form*.

Definition 8.53. An $n \times n$ matrix of the form†

$$J_{\lambda,n} = \begin{pmatrix} \lambda & 1 & & & & \\ & \lambda & 1 & & & \\ & & \lambda & 1 & & \\ & & & \ddots & \ddots & \\ & & & & \lambda & 1 \\ & & & & & \lambda \end{pmatrix}, \tag{8.49}$$

in which λ is a real or complex number, is known as a *Jordan block*.

In particular, a 1×1 Jordan block is merely a scalar $J_{\lambda,1} = \lambda$. Since every matrix has at least one (complex) eigenvector — see Theorem 8.11 — the Jordan block matrices have the least possible number of independent eigenvectors. The following result is immediate.

Lemma 8.54. The $n \times n$ Jordan block matrix $J_{\lambda,n}$ has a single eigenvalue, λ, and a single independent eigenvector, $\mathbf{e}_1$. The standard basis vectors $\mathbf{e}_1, \ldots, \mathbf{e}_n$ form a Jordan chain for $J_{\lambda,n}$.

Definition 8.55. A *Jordan matrix* is a square matrix of block diagonal form

$$J = \operatorname{diag}(J_{\lambda_1,n_1}, J_{\lambda_2,n_2}, \ldots, J_{\lambda_k,n_k}) = \begin{pmatrix} J_{\lambda_1,n_1} & & & \\ & J_{\lambda_2,n_2} & & \\ & & \ddots & \\ & & & J_{\lambda_k,n_k} \end{pmatrix}, \tag{8.50}$$

in which one or more Jordan blocks, not necessarily of the same size, lie along the diagonal, while the blank off-diagonal blocks are all zero.

Note that the only possibly non-zero entries in a Jordan matrix are those on the diagonal, which can have any complex value, including 0, and those on the superdiagonal, which are either 1 or 0. The positions of the superdiagonal 1's uniquely prescribes the Jordan blocks.

† All non-displayed entries are zero.

For example, the 6×6 matrices

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \quad \begin{pmatrix} -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix}, \quad \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 2 \end{pmatrix},$$

are all Jordan matrices: the first is a diagonal matrix, consisting of 6 distinct 1×1 Jordan blocks; the second has a 4×4 Jordan block followed by a 2×2 block that happen to have the same diagonal entries; the last has three 2×2 Jordan blocks with respective diagonal entries $0, 1, 2$.

As a direct corollary of Lemma 8.54 combined with the matrix's block structure, cf. Exercise 8.2.50, we obtain a complete classification of the eigenvectors and eigenvalues of a Jordan matrix.

Lemma 8.56. The Jordan matrix (8.50) has eigenvalues $\lambda_1, \ldots, \lambda_k$. The standard basis vectors $\mathbf{e}_1, \ldots, \mathbf{e}_n$ form a Jordan basis, which is partitioned into nonoverlapping Jordan chains labeled by the Jordan blocks.

Thus, in the preceding examples of Jordan matrices, the first has three double eigenvalues, 1 and 2 and 3, and corresponding linearly independent eigenvectors $\mathbf{e}_1, \mathbf{e}_6$, and $\mathbf{e}_2, \mathbf{e}_5$, and $\mathbf{e}_3, \mathbf{e}_4$, each of which belongs to a Jordan chain of length 1. The second matrix has only one eigenvalue, -1, but two independent eigenvectors $\mathbf{e}_1, \mathbf{e}_5$, and hence two Jordan chains, namely $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4$, and $\mathbf{e}_5, \mathbf{e}_6$. The last has three eigenvalues $0, 1, 2$, three eigenvectors $\mathbf{e}_1, \mathbf{e}_3, \mathbf{e}_5$, and three Jordan chains of length 2: $\mathbf{e}_1, \mathbf{e}_2$, and $\mathbf{e}_3, \mathbf{e}_4$, and $\mathbf{e}_5, \mathbf{e}_6$. In particular, the only complete Jordan matrices are the diagonal matrices, all of whose Jordan blocks are of size 1×1.

The Jordan canonical form follows from the Jordan Basis Theorem 8.51.

Theorem 8.57. Let A be an $n \times n$ real or complex matrix. Let $S = (\mathbf{w}_1 \; \mathbf{w}_2 \; \ldots \; \mathbf{w}_n)$ be a matrix whose columns form a Jordan basis of A. Then S places A into the *Jordan canonical form*

$$S^{-1} A S = J = \operatorname{diag}(J_{\lambda_1, n_1}, J_{\lambda_2, n_2}, \ldots, J_{\lambda_k, n_k}), \quad \text{or, equivalently,} \quad A = S J S^{-1}. \tag{8.51}$$

The diagonal entries of the similar Jordan matrix J are the eigenvalues of A. In particular, A is complete (diagonalizable) if and only if every Jordan block is of size 1×1 or, equivalently, all Jordan chains are of length 1. The Jordan canonical form of A is uniquely determined up to a permutation of the diagonal Jordan blocks.

For instance, the matrix $A = \begin{pmatrix} -1 & 0 & 1 & 0 & 0 \\ -2 & 2 & -4 & 1 & 1 \\ -1 & 0 & -3 & 0 & 0 \\ -4 & -1 & 3 & 1 & 0 \\ 4 & 0 & 2 & -1 & 0 \end{pmatrix}$ considered in Example 8.52

has the following Jordan basis matrix and Jordan canonical form

$$S = \begin{pmatrix} 0 & 0 & 0 & -1 & -1 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ -1 & 0 & 1 & -2 & -2 \\ 1 & -1 & 0 & 0 & 1 \end{pmatrix}, \qquad J = S^{-1} A S = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -2 & 1 \\ 0 & 0 & 0 & 0 & -2 \end{pmatrix}.$$

Finally, let us outline a proof of the Jordan Basis Theorem 8.51.

Lemma 8.58. If $\mathbf{v}_1, \ldots, \mathbf{v}_n$ forms a Jordan basis for the matrix A, it also forms a Jordan basis for $B = A - c\,\mathrm{I}$, for any scalar c.

Proof: We note that the eigenvalues of B are all of the form $\lambda - c$, where λ is an eigenvalue of A. Moreover, given a Jordan chain $\mathbf{w}_1, \ldots, \mathbf{w}_j$ of A, we have

$$B\,\mathbf{w}_1 = (\lambda - c)\,\mathbf{w}_1, \qquad B\,\mathbf{w}_i = (\lambda - c)\,\mathbf{w}_i + \mathbf{w}_{i-1}, \qquad i = 2, \ldots, j,$$

so $\mathbf{w}_1, \ldots, \mathbf{w}_j$ is also a Jordan chain for B corresponding to the eigenvalue $\lambda - c$. *Q.E.D.*

The proof of Theorem 8.51 will be done by induction on the size n of the matrix. The case $n = 1$ is trivial, since every nonzero element of $\mathbb{C}$ is a Jordan basis for a 1×1 matrix. To perform the induction step, we assume that the result is valid for all matrices of size $\leq n-1$. Let A be an $n \times n$ matrix. According to Theorem 8.11, A has at least one complex eigenvalue λ_1. Let $B = A - \lambda_1\,\mathrm{I}$. Since λ_1 is an eigenvalue of A, we know that 0 is an eigenvalue of B. This means that $\ker B \neq \{\mathbf{0}\}$, and so $r = \operatorname{rank} B < n$. Moreover, by Lemma 8.58, a Jordan basis of B is also a Jordan basis for A, and so we can concentrate all our attention on the singular matrix B from now on.

Recall that $W = \operatorname{img} B \subset \mathbb{C}^n$ is an invariant subspace for B, i.e., $B\,\mathbf{w} \in W$ whenever $\mathbf{w} \in W$. Moreover, since B is singular, $\dim W = r = \operatorname{rank} B < n$. Thus, by fixing a basis of W, we can realize the restriction $B\colon W \to W$ as multiplication by an $r \times r$ matrix. The fact that $r < n$ allows us to invoke the induction hypothesis and deduce the existence of a Jordan basis $\mathbf{w}_1, \ldots, \mathbf{w}_r \in W \subset \mathbb{C}^n$ for the action of B on the subspace W. Our goal is to complete this collection to a full Jordan basis on $\mathbb{C}^n$.

To this end, we append two additional kinds of vectors. Suppose that the Jordan basis of W contains k null Jordan chains associated with its zero eigenvalue. Each null Jordan chain consists of vectors $\mathbf{w}_1, \ldots, \mathbf{w}_j \in W$ satisfying $B\,\mathbf{w}_1 = \mathbf{0}$, $B\,\mathbf{w}_2 = \mathbf{w}_1$, $\ldots$, $B\,\mathbf{w}_j = \mathbf{w}_{j-1}$, cf. (8.47). The number k of null Jordan chains is equal to the number of linearly independent null eigenvectors of B that belong to $W = \operatorname{img} B$, that is $k = \dim(\ker B \cap \operatorname{img} B)$. To each such null Jordan chain, we append a vector $\mathbf{w}_{j+1} \in \mathbb{C}^n$ such that $B\,\mathbf{w}_{j+1} = \mathbf{w}_j$, noting that $\mathbf{w}_{j+1}$ exists because $\mathbf{w}_j \in \operatorname{img} B$. We deduce that $\mathbf{w}_1, \ldots, \mathbf{w}_{j+1} \in \mathbb{C}^n$ forms a null Jordan chain, of length $j+1$. Having extended all the null Jordan chains in W, the resulting collection consists of $r + k$ vectors in $\mathbb{C}^n$ arranged in nonoverlapping Jordan chains. To complete to a basis, we append $n - r - k$ additional linearly independent null vectors $\mathbf{z}_1, \ldots, \mathbf{z}_{n-r-k} \in \ker B \setminus \operatorname{img} B$ that lie outside its image. Since $B\,\mathbf{z}_j = \mathbf{0}$, each $\mathbf{z}_j$ forms a null Jordan chain of length 1. We claim that the complete collection consisting of the r non-null Jordan chains in W, the k extended null chains, and the additional null vectors $\mathbf{z}_1, \ldots, \mathbf{z}_{n-r-k}$, forms the desired Jordan basis. By construction, it consists of nonoverlapping Jordan chains. The only remaining issue is the proof that the resulting collection of vectors is linearly independent; this is left as a challenge for the reader in Exercise 8.6.25. *Q.E.D.*

With the Jordan canonical form in hand, the general result characterizing complex invariant subspaces of a general square matrix can now be stated.

Theorem 8.59. Every complex invariant subspace of a square matrix A is spanned by a finite number of Jordan chains.

Proof: The proof proceeds in the same manner as that of Theorem 8.30. We assume that $\mathbf{v}_k$ is the *last* vector in its Jordan chain. We deduce that $A\mathbf{w} - \lambda_k \mathbf{w}$ is a linear combination of $\mathbf{v}_1, \ldots, \mathbf{v}_{k-1}$ with non-vanishing coefficients, and so the induction step remains valid in this case. The details are left to the reader. *Q.E.D.*

Similarly, a complex conjugate pair of Jordan chains of order k produces two distinct k-dimensional complex invariant subspaces, and hence a real invariant subspace of dimension $2k$ obtained from their real and imaginary parts. A detailed statement of the nature of the most general real invariant subspace of a matrix is also left to the reader.

Exercises

8.6.6. For each of the following Jordan matrices, identify the Jordan blocks. Write down the eigenvalues, the eigenvectors, and the Jordan basis. Clearly identify the Jordan chains.

$$(a)\ \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix},\quad (b)\ \begin{pmatrix} -3 & 0 \\ 0 & 6 \end{pmatrix},\quad (c)\ \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix},\quad (d)\ \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix},\quad (e)\ \begin{pmatrix} 4 & 0 & 0 & 0 \\ 0 & 3 & 1 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix}.$$

8.6.7. Find a Jordan basis and Jordan canonical form for each of the following matrices:

$$(a)\ \begin{pmatrix} 2 & 3 \\ 0 & 2 \end{pmatrix},\quad (b)\ \begin{pmatrix} -1 & -1 \\ 4 & -5 \end{pmatrix},\quad (c)\ \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix},$$

$$(d)\ \begin{pmatrix} -3 & 1 & 0 \\ 1 & -3 & -1 \\ 0 & 1 & -3 \end{pmatrix},\quad (e)\ \begin{pmatrix} -1 & 1 & 1 \\ -2 & -2 & -2 \\ 1 & -1 & -1 \end{pmatrix},\quad (f)\ \begin{pmatrix} 2 & -1 & 1 & 2 \\ 0 & 2 & 0 & 1 \\ 0 & 0 & 2 & -1 \\ 0 & 0 & 0 & 2 \end{pmatrix}.$$

8.6.8. Write down all possible 4×4 Jordan matrices that have only $\lambda = 2$ as an eigenvalue.

8.6.9. Write down all 3×3 Jordan matrices that have eigenvalues 2 and 5 (and no others).

8.6.10. Write down a formula for the inverse of a Jordan block matrix.
Hint: Try some small examples first to help in figuring out the pattern.

8.6.11. *True or false*: If A is complete, every generalized eigenvector is an ordinary eigenvector.

8.6.12. *True or false*: Every generalized eigenvector belongs to a Jordan chain.

8.6.13. *True or false*: If $\mathbf{w}$ is a generalized eigenvector of A, then $\mathbf{w}$ is a generalized eigenvector of every power A^j, for $j \in \mathbb{N}$, thereof.

8.6.14. *True or false*: If $\mathbf{w}$ is a generalized eigenvector of index k of A, then $\mathbf{w}$ is an ordinary eigenvector of A^k.

♢ 8.6.15. Suppose you know all eigenvalues of a matrix as well as their algebraic and geometric multiplicities. Can you determine the matrix's Jordan canonical form?

8.6.16. *True or false*: If $\mathbf{w}_1, \ldots, \mathbf{w}_j$ is a Jordan chain for a matrix A, so are the scalar multiples $c\mathbf{w}_1, \ldots, c\mathbf{w}_j$ for all $c \neq 0$.

8.6.17. Let A and B be $n \times n$ matrices. According to Exercise 8.2.23, the matrix products AB and BA have the same eigenvalues. Do they have the same Jordan form?

8.6.18. *True or false*: If the Jordan canonical form of A is J, then that of A^2 is J^2.

♢ 8.6.19. (*a*) Give an example of a matrix A such that A^2 has an eigenvector that is *not* an eigenvector of A. (*b*) Show that, in general, every eigenvalue of A^2 is the square of an eigenvalue of A.

♢ 8.6.20. (*a*) Prove that a Jordan block matrix $J_{0,n}$ with zero diagonal entries is nilpotent, as in Exercise 1.3.12. (*b*) Prove that a Jordan matrix is nilpotent if and only if all its diagonal entries are zero. (*c*) Prove that a matrix is nilpotent if and only if its Jordan canonical form is nilpotent. (*d*) Explain why a matrix is nilpotent if and only if its only eigenvalue is 0.

8.6.21. Let J be a Jordan matrix. (*a*) Prove that J^k is a complete matrix for some $k \geq 1$ if and only if either J is diagonal, or J is nilpotent with $J^k = \mathrm{O}$. (*b*) Suppose that A is an incomplete matrix such that A^k is complete for some $k \geq 2$. Prove that $A^k = \mathrm{O}$, and hence A is nilpotent. (A simpler version of this problem appears in Exercise 8.3.8.)

♡ 8.6.22. *Cayley–Hamilton Theorem*: Let $p_A(\lambda) = \det(A - \lambda\,\mathrm{I})$ be the characteristic polynomial of A. (a) Prove that if D is a diagonal matrix, then† $p_D(D) = \mathrm{O}$.
Hint: Leave $p_D(\lambda)$ in factored form. (*b*) Prove that if A is complete, then $p_A(A) = \mathrm{O}$. (*c*) Prove that if J is a Jordan block, then $p_J(J) = \mathrm{O}$. (*d*) Prove that this also holds if J is a Jordan matrix. (*e*) Prove the Cayley–Hamilton Theorem: a square matrix A satisfies its own characteristic equation: $p_A(A) = \mathrm{O}$.

♡ 8.6.23. *Minimal polynomial*: Let A be an $n \times n$ matrix. By definition, the *minimal polynomial* of A is the monic polynomial $m_A(t) = t^k + c_{k-1}t^{k-1} + \cdots + c_1 t + c_0$ of *minimal degree* k that annihilates A, so $m_A(A) = A^k + c_{k-1}A^{k-1} + \cdots + c_1 A + c_0\,\mathrm{I} = \mathrm{O}$. (*a*) Prove that the monic minimal polynomial m_A is unique. (*b*) (*c*) Prove that if $r(t)$ is any other polynomial such that $r(A) = \mathrm{O}$, then $r(t) = q(t)\,m_A(t)$ for some polynomial $q(t)$. (*d*) Prove that the matrix's minimal polynomial is a factor of its characteristic polynomial, so $p_A(t) = q_A(t)\,m_A(t)$ for some polynomial $q_A(t)$. *Hint*: Use the Cayley–Hamilton Theorem in Exercise 8.6.22. (*e*) Prove that if A has all distinct eigenvalues, then $p_A = m_A$. (*f*) Prove that $p_A = m_A$ if and only if no two Jordan blocks have the same eigenvalue.

♢ 8.6.24. Prove Lemma 8.54.

♢ 8.6.25. Prove that the n vectors constructed in the proof of Theorem 8.51 are linearly independent and hence form a Jordan basis.
Hint: Suppose that some linear combination vanishes. Apply B to the equation, and then use the fact that we started with a Jordan basis for $W = \operatorname{img} B$.

8.6.26. Find all invariant subspaces of the following matrices: (*a*) $\begin{pmatrix} 4 & 0 \\ -3 & 4 \end{pmatrix}$, (*b*) $\begin{pmatrix} -1 & -1 \\ 4 & -5 \end{pmatrix}$,

(*c*) $\begin{pmatrix} 0 & -1 & 1 \\ 1 & 1 & -1 \\ 3 & 3 & -4 \end{pmatrix}$, (*d*) $\begin{pmatrix} 3 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 3 \end{pmatrix}$, (*e*) $\begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}$, (*f*) $\begin{pmatrix} -1 & 0 & 1 & 2 \\ 0 & 1 & 0 & 1 \\ -1 & -4 & 1 & -2 \\ 0 & 1 & 0 & 1 \end{pmatrix}$.

♡ 8.6.27. The method for constructing a Jordan basis in Example 8.52 is simplified due to the fact that each eigenvalue admits only one Jordan block. On the other hand, the method used in the proof of Theorem 8.51 is rather impractical Devise a general method for constructing a Jordan basis for an arbitrary matrix, paying careful attention to how the Jordan chains having the same eigenvalue are found.

---

† See Exercise 1.2.35 for the basics of matrix polynomials.

8.7 Singular Values

We have already expounded on the central role played by the eigenvalues and eigenvectors of a square matrix in both theory and applications. Further evidence will appear in the subsequent chapters. Alas, rectangular matrices do not have eigenvalues (why?), and so, at first glance, do not appear to possess any quantities of comparable significance. But you no doubt recall that our earlier treatment of least squares minimization problems, as well as the equilibrium equations for structures and circuits, made essential use of the symmetric, positive semi-definite *square* Gram matrix $K = A^T A$ — which can be naturally formed even when A is not square. Perhaps the eigenvalues of K might play a comparably important role for general matrices. Since they are not easily related to the eigenvalues of A — which, in the non-square case, don't even exist — we shall endow them with a new name. They were first studied by the German mathematician Erhard Schmidt in early days of the twentieth century, although intimations can be found in Gauss's work on rigid body dynamics.

Definition 8.60. The *singular values* $\sigma_1, \ldots, \sigma_r$ of an $m \times n$ matrix A are the positive square roots, $\sigma_i = \sqrt{\lambda_i} > 0$, of the nonzero eigenvalues of the associated Gram matrix $K = A^T A$. The corresponding eigenvectors of K are known as the *singular vectors* of A.

Since K is necessarily positive semi-definite, its eigenvalues are always non-negative, $\lambda_i \geq 0$, which justifies the positivity of the singular values of A — independently of whether A itself has positive, negative, or even complex eigenvalues, or is rectangular and has no eigenvalues at all. We will follow the standard convention, and always label the singular values in decreasing order, so that $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r > 0$. Thus, σ_1 will always denote the largest, or *dominant*, singular value. If $K = A^T A$ has repeated eigenvalues, the singular values of A are repeated with the same multiplicities. As we will see, the number r of singular values is equal to the rank of the matrices A and K.

Warning. Some texts include the zero eigenvalues of K as singular values of A. We find this to be less convenient for our development, but you should be aware of the differences between the two conventions.

Example 8.61. Let $A = \begin{pmatrix} 3 & 5 \\ 4 & 0 \end{pmatrix}$. The associated Gram matrix

$$K = A^T A = \begin{pmatrix} 3 & 4 \\ 5 & 0 \end{pmatrix} \begin{pmatrix} 3 & 5 \\ 4 & 0 \end{pmatrix} = \begin{pmatrix} 25 & 15 \\ 15 & 25 \end{pmatrix}$$

has eigenvalues $\lambda_1 = 40$, $\lambda_2 = 10$, and corresponding eigenvectors $\mathbf{v}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$, $\mathbf{v}_2 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$.

Thus, the singular values of A are $\sigma_1 = \sqrt{40} \approx 6.3246$ and $\sigma_2 = \sqrt{10} \approx 3.1623$, with $\mathbf{v}_1, \mathbf{v}_2$ being the singular vectors. Note that the singular values are *not* its eigenvalues, which are $\lambda_1 = \frac{1}{2}(3 + \sqrt{89}) \simeq 6.2170$ and $\lambda_2 = \frac{1}{2}(3 - \sqrt{89}) \simeq -3.2170$, nor are the singular vectors eigenvectors of A.

Only in the special case of symmetric matrices is there a direct connection between their singular values and their (necessarily real) eigenvalues.

Proposition 8.62. If $A = A^T$ is a symmetric matrix, its singular values are the absolute values of its nonzero eigenvalues: $\sigma_i = |\lambda_i| > 0$; its singular vectors coincide with its non-null eigenvectors.

Proof: When A is symmetric, $K = A^T A = A^2$. So, if

$$A\mathbf{v} = \lambda\mathbf{v}, \qquad \text{then} \qquad K\mathbf{v} = A^2\mathbf{v} = A(\lambda\mathbf{v}) = \lambda A\mathbf{v} = \lambda^2\mathbf{v}.$$

Thus, every eigenvector $\mathbf{v}$ of A is also an eigenvector of K with eigenvalue λ^2. Therefore, the eigenvector basis of A guaranteed by Theorem 8.32 is also an eigenvector basis for K, and hence forms a complete system of singular vectors for A. *Q.E.D.*

The generalization of the spectral factorization (8.35) to non-symmetric matrices is known as the *singular value decomposition*, commonly abbreviated SVD. Unlike the former, which applies only to square matrices, every nonzero matrix possesses a singular value decomposition.

Theorem 8.63. A nonzero real $m \times n$ matrix A of rank $r > 0$ can be factored,

$$A = P\,\Sigma\,Q^T, \tag{8.52}$$

into the product of an $m \times r$ matrix P with orthonormal† columns, so $P^T P = \mathrm{I}$, the $r \times r$ diagonal matrix $\Sigma = \operatorname{diag}(\sigma_1, \ldots, \sigma_r)$ that has the singular values of A as its diagonal entries, and an $r \times n$ matrix Q^T with orthonormal rows, so $Q^T Q = \mathrm{I}$.

Proof: Let's begin by rewriting the desired factorization (8.52) as $AQ = P\,\Sigma$. The individual columns of this matrix equation are the vector equations

$$A\mathbf{q}_i = \sigma_i\,\mathbf{p}_i, \qquad i = 1, \ldots, r, \tag{8.53}$$

relating the orthonormal columns of $Q = (\,\mathbf{q}_1\ \mathbf{q}_2\ \ldots\ \mathbf{q}_r\,)$ to the orthonormal columns of $P = (\,\mathbf{p}_1\ \mathbf{p}_2\ \ldots\ \mathbf{p}_r\,)$. Thus, our goal is to find vectors $\mathbf{p}_1, \ldots, \mathbf{p}_r$ and $\mathbf{q}_1, \ldots, \mathbf{q}_r$ that satisfy (8.53). To this end, we let $\mathbf{q}_1, \ldots, \mathbf{q}_r$ be orthonormal eigenvectors of the Gram matrix $K = A^T A$ corresponding to the non-zero eigenvalues, which, according to Proposition 8.37 form a basis for $\operatorname{img} K = \operatorname{coimg} A$, of dimension $r = \operatorname{rank} K = \operatorname{rank} A$. Thus, by the definition of the singular values,

$$A^T A\mathbf{q}_i = K\,\mathbf{q}_i = \sigma_i^2\,\mathbf{q}_i, \qquad i = 1, \ldots, r. \tag{8.54}$$

We claim that the image vectors $\mathbf{w}_i = A\mathbf{q}_i$ are automatically orthogonal. Indeed, in view of the orthonormality of the $\mathbf{q}_i$ combined with (8.54),

$$\mathbf{w}_i \cdot \mathbf{w}_j = \mathbf{w}_i^T \mathbf{w}_j = (A\mathbf{q}_i)^T A\mathbf{q}_j = \mathbf{q}_i^T A^T A\mathbf{q}_j = \sigma_j^2\,\mathbf{q}_i^T\,\mathbf{q}_j = \sigma_j^2\,\mathbf{q}_i \cdot \mathbf{q}_j = \begin{cases} 0, & i \neq j, \\ \sigma_i^2, & i = j. \end{cases}$$

Consequently, $\mathbf{w}_1, \ldots, \mathbf{w}_r$ form an orthogonal system of vectors having respective norms

$$\|\,\mathbf{w}_i\,\| = \sqrt{\mathbf{w}_i \cdot \mathbf{w}_i} = \sigma_i.$$

We conclude that the associated unit vectors

$$\mathbf{p}_i = \frac{\mathbf{w}_i}{\sigma_i} = \frac{A\mathbf{q}_i}{\sigma_i}, \qquad i = 1, \ldots, r, \tag{8.55}$$

form an orthonormal set of vectors satisfying the required equations (8.53). *Q.E.D.*

---

† Throughout this section, we exclusively use the Euclidean dot product and norm.

Remark. If A has distinct singular values, its singular value decomposition (8.52) is almost unique, modulo simultaneously changing the signs of one or more corresponding columns of Q and P. Matrices with repeated singular values have more freedom in their singular value decomposition, since one can choose singular vectors using different orthonormal bases of each eigenspace of the Gram matrix $A^T A$.

Observe that, taking the transpose of (8.52) and noting that $\Sigma^T = \Sigma$ is diagonal, we obtain

$$A^T = Q\,\Sigma\,P, \tag{8.56}$$

which is a singular value decomposition of the transposed matrix A^T. In particular, we obtain the following result:

Corollary 8.64. A matrix A and its transpose A^T have the same singular values.

Note that their singular vectors are not the same; indeed, those of A are the orthonormal columns of Q, whereas those of A^T are the orthonormal columns of P, which are related by (8.53). Thus,

$$A^T \mathbf{p}_i = \sigma_i \mathbf{q}_i, \qquad i = 1, \ldots, r, \tag{8.57}$$

which is also a consequence of (8.54).

Furthermore, the singular value decomposition (8.52) serves to diagonalize the Gram matrix; indeed, since $P^T P = \mathrm{I}$, we have

$$Q^T K\,Q = Q^T A^T A\,Q = Q^T A^T P^T P\,A\,Q = (P\,A\,Q)^T (P\,A\,Q) = \Sigma^T \Sigma = \Sigma^2, \tag{8.58}$$

because Σ is diagonal. If A has maximal rank n, then Q is an $n \times n$ orthogonal matrix, and so (8.58) implies that the linear transformation defined by the Gram matrix K is diagonalized when expressed in terms of the orthonormal basis formed by the singular vectors. If $r = \operatorname{rank} A < n$, then one can supplement the r singular vectors $\mathbf{u}_1, \ldots, \mathbf{u}_r$ with $n - r$ unit vectors $\mathbf{u}_{r+1}, \ldots, \mathbf{u}_n \in \ker K = \ker A$ so as to form an orthonormal basis of $\mathbb{R}^n$. In terms of this basis, the Gram matrix assumes diagonal form with the r nonzero squared singular values (nonzero eigenvalues of K) as its first r diagonal entries, while the remaining diagonal entries are all 0, in accordance with the Spectral Theorem 8.38.

Example 8.65. For the matrix $A = \begin{pmatrix} 3 & 5 \\ 4 & 0 \end{pmatrix}$ in Example 8.61, the orthonormal eigenvector basis of $K = A^T A = \begin{pmatrix} 25 & 15 \\ 15 & 25 \end{pmatrix}$ is given by the unit singular vectors $\mathbf{q}_1 = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}$ and $\mathbf{q}_2 = \begin{pmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}$. Thus, $Q = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}$. Next, according to (8.55),

$$\mathbf{p}_1 = \frac{A\,\mathbf{q}_1}{\sigma_1} = \frac{1}{\sqrt{40}} \begin{pmatrix} 4\sqrt{2} \\ 2\sqrt{2} \end{pmatrix} = \begin{pmatrix} \frac{2}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} \end{pmatrix}, \qquad \mathbf{p}_2 = \frac{A\,\mathbf{q}_2}{\sigma_2} = \frac{1}{\sqrt{10}} \begin{pmatrix} \sqrt{2} \\ -2\sqrt{2} \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{5}} \\ -\frac{2}{\sqrt{5}} \end{pmatrix},$$

and thus $P = \begin{pmatrix} \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \end{pmatrix}$. You may wish to validate the resulting singular value factorization

$$A = \begin{pmatrix} 3 & 5 \\ 4 & 0 \end{pmatrix} = \begin{pmatrix} \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \end{pmatrix} \begin{pmatrix} \sqrt{40} & 0 \\ 0 & \sqrt{10} \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} = P\,\Sigma\,Q^T.$$

The singular value decomposition is revealing some interesting new geometrical information concerning the action of matrices, further supplementing the discussion begun in Section 2.5 and continued in Section 4.4.

Proposition 8.66. Given the singular value decomposition $A = P\,\Sigma\,Q^T$, the columns $\mathbf{q}_1, \ldots, \mathbf{q}_r$ of Q form an orthonormal basis for $\operatorname{coimg} A$, while the columns $\mathbf{p}_1, \ldots, \mathbf{p}_r$ of P form an orthonormal basis for $\operatorname{img} A$.

Proof: The first part of the proposition was proved during the course of the proof of Theorem 8.63. Moreover, the vectors $\mathbf{p}_i = A(\mathbf{q}_i/\sigma_i)$ for $i = 1, \ldots, r$ are mutually orthogonal, of unit length, and belong to $\operatorname{img} A$, which has dimension $r = \operatorname{rank} A$. They therefore form an orthonormal basis for the image. *Q.E.D.*

If A is a nonsingular $n \times n$ matrix, then the matrices P, Σ, Q appearing in its singular value decomposition (8.52) are all of size $n \times n$. If we interpret them as linear transformations of $\mathbb{R}^n$, the two orthogonal matrices represent rigid rotations/reflections, while the diagonal matrix Σ represents a combination of simple stretches, by an amount given by the singular values, in the orthogonal coordinate directions. Thus, every invertible linear transformation on $\mathbb{R}^n$ can be decomposed into a rotation/reflection Q^T, followed by the stretching transformation along the coordinate axes represented by Σ, followed by another rotation/reflection P. See also Exercise 8.5.28.

In the more general rectangular case, the matrix Q^T represents an orthogonal projection from $\mathbb{R}^n$ to $\operatorname{coimg} A$, the matrix Σ continues to represent a stretching transformation within this r-dimensional subspace, while P maps the result to $\operatorname{img} A \subset \mathbb{R}^m$. We already noted in Section 4.4 that the linear transformation $L\colon \mathbb{R}^n \to \mathbb{R}^m$ defined by matrix multiplication, $L[\mathbf{x}] = A\,\mathbf{x}$, can be interpreted as a projection from $\mathbb{R}^n$ to $\operatorname{coimg} A$ followed by an invertible map from $\operatorname{coimg} A$ to $\operatorname{img} A$. The singular value decomposition tells us that not only is the latter map invertible, it is simply a combination of stretches in the r mutually orthogonal singular directions $\mathbf{q}_1, \ldots, \mathbf{q}_r$, whose magnitudes equal the nonzero singular values. In this way, we have at last reached a complete understanding of the subtle geometry underlying the simple operation of multiplying a vector by a matrix!

Finally, we note that practical numerical algorithms for computing singular values and the singular value decomposition can be found in Chapter 9 and in [**32**, **66**, **90**].

The Pseudoinverse

The singular value decomposition enables us to substantially generalize the concept of a matrix inverse. The pseudoinverse was first defined by the American mathematician Eliakim Moore in the 1920's and rediscovered by the British mathematical physicist Sir Roger Penrose in the 1950's, and often has their names attached.

Definition 8.67. The *pseudoinverse* of a nonzero $m \times n$ matrix with singular value decomposition $A = P\,\Sigma\,Q^T$ is the $n \times m$ matrix $A^+ = Q\,\Sigma^{-1}P^T$.

Note that the latter equation is the singular value decomposition of the pseudoinverse A^+, and hence its nonzero singular values are the reciprocals of the nonzero singular values of A. The only matrix without a pseudoinverse is the zero matrix O. If A is a non-singular square matrix, then its pseudoinverse agrees with its ordinary inverse. Indeed, since in this case both P and Q are square orthogonal matrices, it follows that

$$A^{-1} = (P\,\Sigma\,Q^T)^{-1} = (Q^{-1})^T\,\Sigma^{-1}\,P^{-1} = Q\,\Sigma^{-1}\,P^T = A^+,$$

where we used the fact that the inverse of an orthogonal matrix is equal to its transpose. More generally, if A has linearly independent columns, or, equivalently, $\ker A = \{\mathbf{0}\}$, then we can bypass the singular value decomposition to compute its pseudoinverse.

Lemma 8.68. Let A be an $m \times n$ matrix of rank n. Then

$$A^+ = (A^T A)^{-1} A^T. \tag{8.59}$$

Proof: Replacing A by its singular value decomposition (8.52), we find

$$A^T A = (P \Sigma Q^T)^T (P \Sigma Q^T) = Q \Sigma P^T P \Sigma Q^T = Q \Sigma^2 Q^T, \tag{8.60}$$

since $\Sigma = \Sigma^T$ is a diagonal matrix, while $P^T P = \mathrm{I}$, since the columns of P are orthonormal. This is merely the spectral factorization (8.35) of the Gram matrix $A^T A$ — which we in fact already knew from the original definition of the singular values and vectors. Now if A has rank n, then Q is an $n \times n$ orthogonal matrix, and so $Q^{-1} = Q^T$. Therefore,

$$(A^T A)^{-1} A^T = (Q \Sigma^{-2} Q^T)^{-1} (P \Sigma Q^T)^T = (Q \Sigma^{-2} Q^T)(Q \Sigma P^T) = Q \Sigma^{-1} P^T = A^+. \quad Q.E.D.$$

If A is square and nonsingular, then, as we know, the solution to the linear system $A\mathbf{x} = \mathbf{b}$ is given by $\mathbf{x}^\star = A^{-1}\mathbf{b}$. For a general coefficient matrix, the vector $\mathbf{x}^\star = A^+\mathbf{b}$ obtained by applying the pseudoinverse to the right-hand side plays a distinguished role — it is, in fact, the *least squares solution* to the system under the Euclidean norm.

Theorem 8.69. Consider the linear system $A\mathbf{x} = \mathbf{b}$. Let $\mathbf{x}^\star = A^+\mathbf{b}$, where A^+ is the pseudoinverse of A. If $\ker A = \{\mathbf{0}\}$, then $\mathbf{x}^\star$ is the (Euclidean) least squares solution to the linear system. If, more generally, $\ker A \neq \{\mathbf{0}\}$, then $\mathbf{x}^\star = A^+\mathbf{b} \in \operatorname{coimg} A$ is the least squares solution that has the minimal Euclidean norm among all vectors that minimize the least squares error $\| A\mathbf{x} - \mathbf{b} \|^2$.

Proof: To show that $\mathbf{x}^\star = A^+\mathbf{b}$ is the least squares solution to the system, we must, according to Theorem 5.11 check that it satisfies the normal equations $A^T A \mathbf{x}^\star = A^T \mathbf{b}$. If $\operatorname{rank} A = n$, so $A^T A$ is nonsingular, this follows immediately from (8.59). More generally, combining (8.60), the definition of the pseudoinverse, and the fact that Q has orthonormal columns, so $Q^T Q = \mathrm{I}$, yields

$$A^T A \mathbf{x}^\star = A^T A A^+ \mathbf{b} = (Q \Sigma^2 Q^T)(Q \Sigma^{-1} P^T)\mathbf{b} = Q \Sigma P^T \mathbf{b} = A^T \mathbf{b}.$$

This proves that $\mathbf{x}^\star$ solves the normal equations, and hence, by Theorem 5.11 and Exercise 5.4.11, minimizes the least squares error. Moreover,

$$\begin{aligned} \mathbf{x}^\star = A^+\mathbf{b} &= Q \Sigma^{-1} P^T \mathbf{b} \\ &= Q\mathbf{c} = c_1 \mathbf{q}_1 + \cdots + c_r \mathbf{q}_r, \end{aligned} \qquad \text{where} \qquad \mathbf{c} = (c_1, \ldots c_r)^T = \Sigma^{-1} P^T \mathbf{b}.$$

Thus, $\mathbf{x}^\star$ is a linear combination of the singular vectors, and hence, by Proposition 8.66 and Theorem 4.50, $\mathbf{x}^\star \in \operatorname{coimg} A$ is the solution with minimal norm; the most general least squares solution has the form $\mathbf{x} = \mathbf{x}^\star + \mathbf{z}$ for arbitrary $\mathbf{z} \in \ker A$. *Q.E.D.*

Example 8.70. Let us use the pseudoinverse to solve the linear system $A\mathbf{x} = \mathbf{b}$, with

$$A = \begin{pmatrix} 1 & 2 & -1 \\ 3 & -4 & 1 \\ -1 & 3 & -1 \\ 2 & -1 & 0 \end{pmatrix}, \qquad \mathbf{b} = \begin{pmatrix} 1 \\ 0 \\ -1 \\ 2 \end{pmatrix}.$$

In this case, $\ker A \neq \{\mathbf{0}\}$, and so we are not able use the simpler formula (8.59); thus, we begin by establishing the singular value decomposition of A. The corresponding Gram matrix

$$K = A^T A = \begin{pmatrix} 15 & -15 & 3 \\ -15 & 30 & -9 \\ 3 & -9 & 3 \end{pmatrix}$$

has eigenvalues and eigenvectors

$$\lambda_1 = 24 + 3\sqrt{34} \simeq 41.4929, \qquad \lambda_2 = 24 - 3\sqrt{34} \simeq 6.5071, \qquad \lambda_3 = 0,$$

$$\mathbf{v}_1 \simeq \begin{pmatrix} 2.1324 \\ -3.5662 \\ 1. \end{pmatrix}, \qquad \mathbf{v}_2 \simeq \begin{pmatrix} -2.5324 \\ -1.2338 \\ 1. \end{pmatrix}, \qquad \mathbf{v}_3 = \begin{pmatrix} 1 \\ 2 \\ 5 \end{pmatrix}.$$

The singular values are the square roots of the positive eigenvalues, and so

$$\sigma_1 = \sqrt{\lambda_1} \simeq 6.4415, \qquad \sigma_2 = \sqrt{\lambda_2} \simeq 2.5509,$$

which are used to construct the diagonal singular value matrix $\Sigma \simeq \begin{pmatrix} 6.4415 & 0 \\ 0 & 2.5509 \end{pmatrix}$. Note that A has rank 2, because it has just two singular values. The first two eigenvectors of K are the singular vectors of A, and we use the normalized (unit) singular vectors to form the columns of $Q = (\mathbf{q}_1 \ \mathbf{q}_2) \simeq \begin{pmatrix} .4990 & -.8472 \\ -.8344 & -.4128 \\ .2340 & .3345 \end{pmatrix}$. Next, we apply A to the singular vectors and divide by the corresponding singular value as in (8.55); the resulting vectors

$$\mathbf{p}_1 = \frac{A\mathbf{q}_1}{\sigma_1} \simeq \begin{pmatrix} -.2180 \\ .7869 \\ -.5024 \\ .2845 \end{pmatrix}, \qquad \mathbf{p}_2 = \frac{A\mathbf{q}_2}{\sigma_2} \simeq \begin{pmatrix} -.7869 \\ -.2180 \\ -.2845 \\ -.5024 \end{pmatrix},$$

will form the orthonormal columns of $P = (\mathbf{p}_1, \mathbf{p}_2)$, and, as you can verify,

$$A = P\Sigma Q^T \simeq \begin{pmatrix} -.2180 & -.7869 \\ .7869 & -.2180 \\ -.5024 & -.2845 \\ .2845 & -.5024 \end{pmatrix} \begin{pmatrix} 6.4415 & 0 \\ 0 & 2.5509 \end{pmatrix} \begin{pmatrix} .4990 & -.8344 & .2340 \\ -.8472 & -.4128 & .3345 \end{pmatrix}$$

is the singular value decomposition of our coefficient matrix. Its pseudoinverse is immediately computed:

$$A^+ = Q\,\Sigma^{-1} P^T \simeq \begin{pmatrix} .2444 & .1333 & .0556 & .1889 \\ .1556 & -.0667 & .1111 & .0444 \\ -.1111 & 0 & -.0556 & -.0556 \end{pmatrix}.$$

Finally, according to Theorem 8.69, the least squares solution to the original linear system of minimal Euclidean norm is

$$\mathbf{x}^\star = A^+\mathbf{b} \simeq \begin{pmatrix} .5667 \\ .1333 \\ -.1667 \end{pmatrix}.$$

The Euclidean Matrix Norm

Singular values allow us to finally write down a formula for the natural matrix norm induced by the Euclidean norm (or 2 norm) on $\mathbb{R}^n$, as defined in Theorem 3.20.

Theorem 8.71. Let $\|\cdot\|_2$ denote the Euclidean norm on $\mathbb{R}^n$. Let A be a nonzero matrix with singular values $\sigma_1 \geq \cdots \geq \sigma_r$. Then the Euclidean matrix norm of A equals its dominant (largest) singular value:

$$\|A\|_2 = \max\{\, \|A\mathbf{u}\|_2 \mid \|\mathbf{u}\|_2 = 1 \,\} = \sigma_1, \qquad \text{while} \qquad \|\mathrm{O}\|_2 = 0. \tag{8.61}$$

Proof: Let $\mathbf{q}_1, \ldots, \mathbf{q}_n$ be an orthonormal basis of $\mathbb{R}^n$ consisting of the singular vectors $\mathbf{q}_1, \ldots, \mathbf{q}_r$ along with an orthonormal basis $\mathbf{q}_{r+1}, \ldots, \mathbf{q}_n$ of $\ker A$. Thus, by (8.57),

$$A\mathbf{q}_i = \begin{cases} \sigma_i \mathbf{p}_i, & i = 1, \ldots, r, \\ 0, & i = r+1, \ldots, n, \end{cases}$$

where $\mathbf{p}_1, \ldots, \mathbf{p}_r$ form an orthonormal basis for $\operatorname{img} A$. Suppose $\mathbf{u}$ is any unit vector, so

$$\mathbf{u} = c_1 \mathbf{q}_1 + \cdots + c_n \mathbf{q}_n, \qquad \text{where} \qquad \|\mathbf{u}\| = \sqrt{c_1^2 + \cdots + c_n^2} = 1,$$

thanks to the orthonormality of the basis vectors $\mathbf{q}_1, \ldots, \mathbf{q}_n$ and the general Pythagorean formula (4.5). Then

$$A\mathbf{u} = c_1 \sigma_1 \mathbf{p}_1 + \cdots + c_r \sigma_r \mathbf{p}_r, \qquad \text{and hence} \qquad \|A\mathbf{u}\| = \sqrt{c_1^2 \sigma_1^2 + \cdots + c_r^2 \sigma_r^2}\,,$$

since $\mathbf{p}_1, \ldots, \mathbf{p}_n$ are also orthonormal. Now, since $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r$, we have

$$\begin{aligned} \|A\mathbf{u}\|_2 = \sqrt{c_1^2 \sigma_1^2 + \cdots + c_r^2 \sigma_r^2} &\leq \sqrt{c_1^2 \sigma_1^2 + \cdots + c_r^2 \sigma_1^2} \\ &= \sigma_1 \sqrt{c_1^2 + \cdots + c_r^2} \leq \sigma_1 \sqrt{c_1^2 + \cdots + c_n^2} = \sigma_1. \end{aligned}$$

Moreover, if $c_1 = 1, c_2 = \cdots = c_n = 0$, then $\mathbf{u} = \mathbf{q}_1$, and hence $\|A\mathbf{u}\|_2 = \|A\mathbf{q}_1\|_2 = \|\sigma_1 \mathbf{p}_1\|_2 = \sigma_1$. This implies the desired formula (8.61). *Q.E.D.*

Example 8.72. Consider the matrix $A = \begin{pmatrix} 0 & -\frac{1}{3} & \frac{1}{3} \\ \frac{1}{4} & 0 & \frac{1}{2} \\ \frac{2}{5} & \frac{1}{5} & 0 \end{pmatrix}$. The corresponding Gram matrix

$$A^T A \simeq \begin{pmatrix} .2225 & .0800 & .1250 \\ .0800 & .1511 & -.1111 \\ .1250 & -.1111 & .3611 \end{pmatrix},$$

has eigenvalues $\lambda_1 \simeq .4472$, $\lambda_2 \simeq .2665$, $\lambda_3 \simeq .0210$, and hence the singular values of A are their square roots: $\sigma_1 \simeq .6687$, $\sigma_2 \simeq .5163$, $\sigma_3 \simeq .1448$. The Euclidean matrix norm of A is the largest singular value, and so $\|A\|_2 \simeq .6687$.

Condition Number and Rank

Not only do the singular values provide a compelling geometric interpretation of the action of a matrix on vectors, they also play a key role in modern computational algorithms. The relative magnitudes of the singular values can be used to distinguish well-behaved linear systems from ill-conditioned systems, which are more challenging to solve accurately. This information is quantified by the condition number of the matrix.

Definition 8.73. The *condition number* of a nonsingular $n \times n$ matrix is the ratio between its largest and smallest singular values: $\kappa(A) = \sigma_1/\sigma_n$.

Since the number of singular values equals the matrix's rank, an $n \times n$ matrix with fewer than n singular values is singular, and is said to have condition number ∞. A matrix with a very large condition number is close to singular, and designated as *ill-conditioned*; in practical terms, this occurs when the condition number is larger than the reciprocal of the machine's precision, e.g., 10^7 for typical single-precision arithmetic. As the name implies, it is much harder to solve a linear system $A\mathbf{x} = \mathbf{b}$ when its coefficient matrix is ill-conditioned and hence close to singular.

Determining the rank of a large (square or rectangular) matrix can be a numerical challenge. Small numerical inaccuracies can have an unpredictable effect. For example, the rank 1 matrix $A = \begin{pmatrix} 1 & 1 & -1 \\ 2 & 2 & -2 \\ 3 & 3 & -3 \end{pmatrix}$ is very close to $\widetilde{A} = \begin{pmatrix} 1.00001 & 1. & -1. \\ 2. & 2.00001 & -2. \\ 3. & 3. & -3.00001 \end{pmatrix}$, which has rank 3 and so is nonsingular. On the other hand, the latter matrix is very close to singular, and this is highlighted by its singular values, which are $\sigma_1 \approx 6.48075$ while $\sigma_2 \approx \sigma_3 \approx .000001$. The fact that the second and third singular values are very small indicates that $\widetilde{A}$ is close to a matrix of rank 1 and should be viewed as a numerical (or experimental) perturbation of such a matrix. Thus, an effective practical method for computing the rank of a matrix is to first assign a threshold, e.g., 10^{-5}, to distinguish small singular values, and then treat any singular value lying below the threshold as if it were zero.

This idea is justified by the following theorem, which gives a mechanism for constructing the closest low rank approximations to a given matrix A, as measured in the Euclidean matrix norm.

Theorem 8.74. Let the $m \times n$ matrix A have rank r and singular value decomposition $A = P\,\Sigma\,Q^T$. Given $1 \le k \le r$, let Σ_k denote the upper left $k \times k$ diagonal submatrix of Σ containing the largest k singular values on its diagonal. Let Q_k denote the $n \times k$ matrix formed from the first k columns of Q, which are the first k orthonormal singular vectors of A, and let P_k be the $m \times k$ matrix formed from the first k columns of P. Then the $m \times n$ matrix $A_k = P_k \Sigma_k Q_k^T$ has rank k. Moreover, A_k is the closest rank k matrix to A in the sense that, among all $m \times n$ matrices B of rank k, the Euclidean matrix norm $\| A - B \|$ is minimized when $B = A_k$.

Proof: The fact that A_k has rank k is clear, since, by construction, its singular values are $\sigma_1, \ldots, \sigma_k$. Let $\widetilde{\Sigma}_k$ denote the $r \times r$ diagonal matrix whose first k diagonal entries are $\sigma_1, \ldots, \sigma_k$ and whose last $r - k$ diagonal entries are all 0. Clearly $A_k = P\,\widetilde{\Sigma}_k Q^T$ since the additional zero entries have no effect on the product. Moreover, $\Sigma - \widetilde{\Sigma}_k$ is a diagonal matrix whose first k diagonal entries are all 0 and whose last $r - k$ diagonal entries are $\sigma_{k+1}, \ldots, \sigma_r$. Thus, the difference $A - A_k = P\,(\Sigma - \widetilde{\Sigma}_k)\,Q^T$ has singular values $\sigma_{k+1}, \ldots, \sigma_r$. Since σ_{k+1} is the largest of these, Theorem 8.71 implies that

$$\| A - A_k \| = \sigma_{k+1}.$$

We now prove that this is the smallest possible among all $m \times n$ matrices B of rank k. For such a matrix, according to the Fundamental Theorem 2.49, $\dim \ker B = n - k$. Let $V_{k+1} \subset \mathbb{R}^n$ denote the $(k+1)$-dimensional subspace spanned by the first $k+1$ singular vectors $\mathbf{q}_1, \ldots, \mathbf{q}_{k+1}$ of A. Since the dimensions of the subspaces V_{k+1} and $\ker B$ sum up to $k + 1 + n - k = n + 1 > n$, their intersection is a nontrivial subspace, and hence we can

find a non-zero unit vector

$$\mathbf{u}^\star = c_1\mathbf{q}_1 + \cdots + c_{k+1}\mathbf{q}_{k+1} \in V_{k+1} \cap \ker B.$$

Thus, since $\mathbf{q}_1, \ldots, \mathbf{q}_{k+1}$ are orthonormal,

$$\|\mathbf{u}^\star\| = c_1^2 + \cdots + c_{k+1}^2 = 1, \qquad \text{and, moreover,} \qquad B\mathbf{u}^\star = \mathbf{0},$$

which implies

$$(A - B)\mathbf{u}^\star = A\mathbf{u}^\star = c_1 A\mathbf{q}_1 + \cdots + c_{k+1}A\mathbf{q}_{k+1} = c_1\sigma_1\mathbf{p}_1 + \cdots + c_{k+1}\sigma_{k+1}\mathbf{p}_{k+1}.$$

Since $\mathbf{p}_1, \ldots, \mathbf{p}_{k+1}$ are also orthonormal,

$$\|(A - B)\mathbf{u}^\star\|^2 = c_1^2\sigma_1^2 + \cdots + c_{k+1}^2\sigma_{k+1}^2 \geq (c_1^2 + \cdots + c_{k+1}^2)\sigma_{k+1}^2 = \sigma_{k+1}^2.$$

Thus, using the definition (3.39) of the Euclidean matrix norm

$$\|A - B\| = \max\{\, \|(A - B)\mathbf{u}\| \mid \|\mathbf{u}\| = 1 \,\} \geq \|(A - B)\mathbf{u}^\star\| \geq \sigma_{k+1}.$$

This proves that σ_{k+1} minimizes $\|A - B\|$ among all such matrices B. *Q.E.D.*

Remark. One cannot do any better with a matrix of lower rank, i.e., $\|A - B\|$ is also minimized when $B = A_k$ among all matrices with $\operatorname{rank} B \leq k$. Justifying this statement is left to Exercise 8.7.19.

Observe that the closest rank k approximating matrix A_k is unique unless the $(k+1)^{\text{st}}$ singular value equals the kth one: $\sigma_k = \sigma_{k+1}$, in which case one can replace the last columns of P_k and Q_k with those coming from the $(k+1)^{\text{st}}$ singular vectors. To compute the rank k approximating matrix A_k, we need only compute the largest k singular values $\sigma_1, \ldots, \sigma_k$ of A, which form the diagonal entries of Σ_k, and corresponding singular unit vectors $\mathbf{q}_1, \ldots, \mathbf{q}_k$, which form the columns of Q_k. The columns of P_k are formed by their images $\mathbf{p}_1 = A\mathbf{q}_1/\sigma_1, \ldots, \mathbf{p}_k = A\mathbf{q}_k/\sigma_k$.

Consequently, when solving an ill-conditioned linear system $A\mathbf{x} = \mathbf{b}$, a common and effective regularization strategy is to eliminate all "insignificant" singular values below a specified cut-off, replacing A by its rank k approximation A_k specified by Theorem 8.74, where k denotes the number of significant singular values. Applying the corresponding approximating pseudoinverse $A_k^+ = Q_k\Sigma_k^{-1}P_k^T$ to solve for $\mathbf{x}^\star = A_k^+\mathbf{b}$ will, in favorable situations, effectively circumvent the effects of ill-conditioning.

Another common application of low rank approximations is in data compression, in which one replaces a very large data matrix, e.g., one obtained from high-resolution digital images, by a suitable low rank approximation that captures the essential features of the data set while reducing overall storage requirements and thereby accelerating subsequent analysis thereof.

Spectral Graph Theory

Spectral graph theory, [**14**, **76**], refers to the study of the properties of graphs that are captured by the spectrum, meaning the eigenvalues and singular values, of certain naturally associated matrices. The *graph Laplacian matrix*, which we encountered in Section 6.2, is of particular importance. Recall that it is defined as the Gram matrix, $K = A^T A$, constructed from the incidence matrix A of any underlying digraph, noting that the directions assigned to the edges do not affect the ultimate form of K. The eigenvalues of the graph Laplacian matrix are, by definition, the squares of the singular values of the incidence matrix.

As we know — see Exercise 2.6.12 — the dimension of the kernel of the incidence matrix, and hence also thatof its graph Laplacian matrix, equals the number of connected components of the graph. In particular, a connected graph has a one-dimensional kernel spanned by the vector $(1, 1, \ldots, 1)^T$. The magnitude of its final, meaning smallest, singular value, σ_r, can be interpreted as a measure of how close the graph is to being disconnected, since if it were zero (and thus technically not a singular value), the graph would have (at least) two connected components. This is borne out by numerical experiments, which demonstrate that a graph with a small final singular value σ_r can be disconnected by deleting a relatively small number of its edges.

Example 8.75. Consider the graph sketched in Figure 8.4. Using the indicated vertex labels, we can construct its graph Laplacian directly using the recipe found in Section 6.2:

$$K = \begin{pmatrix} 3 & -1 & -1 & -1 & 0 & 0 & 0 & 0 \\ -1 & 3 & -1 & -1 & 0 & 0 & 0 & 0 \\ -1 & -1 & 3 & -1 & 0 & 0 & 0 & 0 \\ -1 & -1 & -1 & 4 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 3 & -1 & -1 & 0 \\ 0 & 0 & 0 & 0 & -1 & 3 & -1 & -1 \\ 0 & 0 & 0 & 0 & -1 & -1 & 3 & -1 \\ 0 & 0 & 0 & 0 & 0 & -1 & -1 & 2 \end{pmatrix}. \tag{8.62}$$

To four decimal places, the eigenvalues are 5.3234, 4., 4., 4., 4., 2.3579, .3187, 0., with corresponding singular values 2.3073, 2., 2., 2., 2., 1.5356, .5645, 0. The relatively small value of $\sigma_7 = .5645$ indicates the graph is not especially well connected. Indeed, we can disconnect it by erasing just the one edge from vertex 4 to vertex 5. The resulting disconnected graph Laplacian is the block diagonal matrix

$$\widetilde{K} = \begin{pmatrix} 3 & -1 & -1 & -1 & 0 & 0 & 0 & 0 \\ -1 & 3 & -1 & -1 & 0 & 0 & 0 & 0 \\ -1 & -1 & 3 & -1 & 0 & 0 & 0 & 0 \\ -1 & -1 & -1 & 3 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2 & -1 & -1 & 0 \\ 0 & 0 & 0 & 0 & -1 & 3 & -1 & -1 \\ 0 & 0 & 0 & 0 & -1 & -1 & 3 & -1 \\ 0 & 0 & 0 & 0 & 0 & -1 & -1 & 2 \end{pmatrix}$$

whose spectrum is the union of the spectra of the two constituent subgraphs, the left one having a triple eigenvalue of 4 and a zero eigenvalue, the right one having eigenvalues $4, 2, 2, 0$. The singular values are their square roots. Note that these are fairly close to those of the original connected graph. Such observations are even more striking when one is dealing with much larger graphs.

Spectral graph theory plays an increasingly important role in theoretical computer science and data analysis. Applications include partitioning and coloring graphs, random graphs and random walks on graphs, routing and networks, and spectral clustering. The PageRank algorithm that underlies Google's search engine is based on representing the web pages on the internet as a gigantic digraph,† which is then viewed as a probabilistic

† According to our conventions, the internet digraph is not simple, since vertices can have two directed edges connecting them, one if the first web page links to the second and another if the second links to the first.

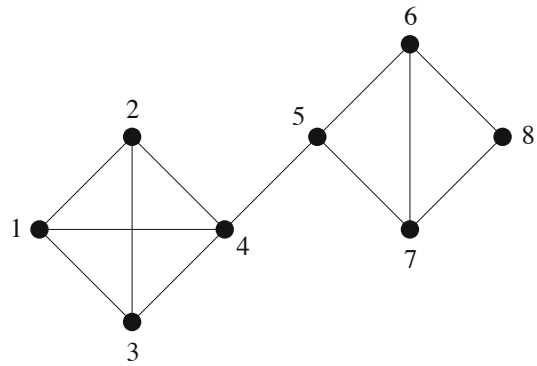

Figure 8.4. An Almost Disconnected Graph.

Markov process; see Section 9.3 for further details.

A basic example is the *complete graph* G_n on n vertices, which has one edge joining every distinct pair of vertices, and hence is *the* most connected simple graph; see Exercise 2.6.10. Its graph Laplacian is easily constructed, and is the $n \times n$ matrix $K_n = n\,\mathrm{I} - E$, where E is the $n \times n$ matrix with every entry equal to 1. Since $\dim \ker E = n - 1$ (why?), we see that K_n has a single nonzero eigenvalue, namely $\lambda_1 = n$, of multiplicity $n - 1$ along with its zero eigenvalue. Thus, the complete graph on n vertices has $n - 1$ identical singular values: $\sigma_1 = \cdots = \sigma_{n-1} = \sqrt{n}$. Motivated by this observation, graphs whose nonzero singular values are close together are, in a certain sense, very highly connected, and are known as *expander graphs*. Expander graphs have many remarkable properties, which underlie their many applications, including communication networks, error-correcting codes, fault-tolerant circuits, pseudo-random number generators, Markov processes, statistical physics, as well as more theoretical disciplines such as group theory and geometry, [**45**].

Exercises

8.7.1. Find the singular values of the following matrices: (*a*) $\begin{pmatrix} 1 & 1 \\ 0 & 2 \end{pmatrix}$, (*b*) $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$, (*c*) $\begin{pmatrix} 1 & -2 \\ -3 & 6 \end{pmatrix}$, (*d*) $\begin{pmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \end{pmatrix}$, (*e*) $\begin{pmatrix} 2 & 1 & 0 & -1 \\ 0 & -1 & 1 & 1 \end{pmatrix}$, (*f*) $\begin{pmatrix} 1 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 1 \end{pmatrix}$.

8.7.2. Write out the singular value decomposition (8.52) of the matrices in Exercise 8.7.1.

8.7.3. (*a*) Construct the singular value decomposition of the shear matrix $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$.

(*b*) Explain how a shear can be realized as a combination of a rotation, and a stretch, followed by a second rotation.

8.7.4. Find the condition number of the following matrices. Which would you characterize as ill-conditioned? (*a*) $\begin{pmatrix} 2 & -1 \\ -3 & 1 \end{pmatrix}$, (*b*) $\begin{pmatrix} -.999 & .341 \\ -1.001 & .388 \end{pmatrix}$, (*c*) $\begin{pmatrix} 1 & 2 \\ 1.001 & 1.9997 \end{pmatrix}$, (*d*) $\begin{pmatrix} -1 & 3 & 4 \\ 2 & 10 & 6 \\ 1 & 2 & -3 \end{pmatrix}$, (*e*) $\begin{pmatrix} 72 & 96 & 103 \\ 42 & 55 & 59 \\ 67 & 95 & 102 \end{pmatrix}$, (*f*) $\begin{pmatrix} 5 & 7 & 6 & 5 \\ 7 & 10 & 8 & 7 \\ 6 & 8 & 10 & 9 \\ 5 & 7 & 9 & 10 \end{pmatrix}$.

8.7.5. Find the closest rank 1 and rank 2 matrices to the matrices in Exercise 8.7.4.

♠ 8.7.6. Solve the following systems of equations using Gaussian Elimination with three-digit rounding arithmetic. Is your answer a reasonable approximation to the exact solution? Compare the accuracy of your answers with the condition number of the coefficient matrix, and discuss the implications of ill-conditioning.

$$(a)\quad \begin{aligned} 1000x + 999y &= 1, \\ 554x + 555y &= -1, \end{aligned} \qquad (b)\quad \begin{aligned} 97x + 175y + 83z &= 1, \\ 44x + 78y + 37z &= 1, \\ 52x + 97y + 46z &= 1. \end{aligned} \qquad (c)\quad \begin{aligned} 3.001x + 2.999y + 5z &= 1, \\ -x + 1.002y - 2.999z &= 2 \\ 2.002x + 4y + 2z &= 1.002. \end{aligned}$$

♠ 8.7.7. (*a*) Compute the singular values and condition numbers of the 2×2, 3×3, and 4×4 Hilbert matrices. (*b*) What is the smallest Hilbert matrix with condition number larger than 10^6?

8.7.8. (*a*) What are the singular values of a $1 \times n$ matrix? (*b*) Write down its singular value decomposition. (*c*) Write down its pseudoinverse.

8.7.9. Answer Exercise 8.7.8 for an $m \times 1$ matrix.

8.7.10. *True or false*: Every matrix has at least one singular value.

8.7.11. Explain why the singular values of A are the same as the nonzero eigenvalues of the positive definite square root matrix $S = \sqrt{A^T A}$, defined in Exercise 8.5.27.

♢ 8.7.12. Prove that if the square matrix A is nonsingular, then the singular values of A^{-1} are the reciprocals of the singular values of A. How are their condition numbers related?

8.7.13. *True or false*: The singular values of A^T are the same as the singular values of A.

♢ 8.7.14. (*a*) Let A be a nonsingular matrix. Prove that the product of the singular values of A equals the absolute value of its determinant: $\sigma_1 \sigma_2 \cdots \sigma_n = |\det A|$.
(*b*) Does their sum equal the absolute value of the trace: $\sigma_1 + \cdots + \sigma_n = |\operatorname{tr} A|$?
(*c*) Show that if $|\det A| < 10^{-k}$, then its minimal singular value satisfies $\sigma_n < 10^{-k/n}$.
(*d*) *True or false*: A matrix whose determinant is very small is ill-conditioned.
(*e*) Construct an ill-conditioned matrix with $\det A = 1$.

8.7.15. *True or false*: If A is a symmetric matrix, then its singular values are the same as its eigenvalues.

8.7.16. *True or false*: If U is an upper triangular matrix whose diagonal entries are all positive, then its singular values are the same as its diagonal entries.

8.7.17. *True or false*: The singular values of A^2 are the squares σ_i^2 of the singular values of A.

8.7.18. *True or false*: If $B = S^{-1} A S$ are similar matrices, then A and B have the same singular values.

♢ 8.7.19. Under the assumptions of Theorem 8.74, show that $\| A - B \|_2$ is also minimized when $B = A_k$ among all matrices with $\operatorname{rank} B \leq k$.

♢ 8.7.20. Suppose A is an $m \times n$ matrix of rank $r < n$. Prove that there exist arbitrarily close matrices of maximal rank, that is, for every $\varepsilon > 0$ there exists an $m \times n$ matrix B with $\operatorname{rank} B = n$such that the Euclidean matrix norm $\| A - B \| < \varepsilon$.

8.7.21. *True or false*: If $\det A > 1$, then A is not ill-conditioned.

8.7.22. Let $A = \begin{pmatrix} 6 & -4 & 1 \\ -4 & 6 & -1 \\ 1 & -1 & 11 \end{pmatrix}$, and let $E = \{ \mathbf{y} = A\mathbf{x} \mid \|\mathbf{x}\| = 1 \}$ be the image of the unit Euclidean sphere under the linear map induced by A. (*a*) Explain why E is an ellipsoid and write down its equation. (*b*) What are its principal axes and their lengths — the semi-axes of the ellipsoid? (*c*) What is the volume of the solid ellipsoidal domain enclosed by E?

♡ 8.7.23. Let A be a nonsingular 2×2 matrix with singular value decomposition $A = P \Sigma Q^T$ and singular values $\sigma_1 \geq \sigma_2 > 0$. (*a*) Prove that the image of the unit (Euclidean) circle under the linear transformation defined by A is an ellipse, $E = \{ A\mathbf{x} \mid \|\mathbf{x}\| = 1 \}$, whose principal axes are the columns $\mathbf{p}_1, \mathbf{p}_2$ of P, and whose corresponding semi-axes are the singular values σ_1, σ_2. (*b*) Show that if A is symmetric, then the ellipse's principal axes are the eigenvectors of A and the semi-axes are the absolute values of its eigenvalues. (*c*) Prove that the area of E equals $\pi \,|\det A\,|$. (*d*) Find the principal axes, semi-axes, and area of the ellipses defined by (*i*) $\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$, (*ii*) $\begin{pmatrix} 2 & 1 \\ -1 & 2 \end{pmatrix}$, (*iii*) $\begin{pmatrix} 5 & -4 \\ 0 & -3 \end{pmatrix}$. (*e*) What happens if A is singular?

♢ 8.7.24. *Optimization Principles for Singular Values*: Let A be any nonzero $m \times n$ matrix. Prove that (*a*) $\sigma_1 = \max \{ \|A\mathbf{u}\| \mid \|\mathbf{u}\| = 1 \}$. (*b*) Is the minimum the smallest singular value? (*c*) Can you design optimization principles for the intermediate singular values?

♢ 8.7.25. Let A be a square matrix. Prove that its maximal eigenvalue is smaller than its maximal singular value: $\max |\lambda_i| \leq \max \sigma_i$. *Hint*: Use Exercise 8.7.24.

8.7.26. Compute the Euclidean matrix norm of the following matrices.

(*a*) $\begin{pmatrix} \frac{1}{2} & \frac{1}{4} \\ \frac{1}{3} & \frac{1}{6} \end{pmatrix}$, (*b*) $\begin{pmatrix} \frac{5}{3} & \frac{4}{3} \\ -\frac{7}{6} & -\frac{5}{6} \end{pmatrix}$, (*c*) $\begin{pmatrix} \frac{2}{7} & -\frac{2}{7} \\ -\frac{2}{7} & \frac{6}{7} \end{pmatrix}$, (*d*) $\begin{pmatrix} \frac{1}{4} & \frac{3}{2} \\ -\frac{1}{2} & \frac{5}{4} \end{pmatrix}$,

(*e*) $\begin{pmatrix} \frac{2}{7} & \frac{2}{7} & -\frac{4}{7} \\ 0 & \frac{2}{7} & \frac{6}{7} \\ \frac{2}{7} & \frac{4}{7} & \frac{2}{7} \end{pmatrix}$, (*f*) $\begin{pmatrix} 0 & .1 & .8 \\ -.1 & 0 & .1 \\ -.8 & -.1 & 0 \end{pmatrix}$, (*g*) $\begin{pmatrix} 1 & -\frac{2}{3} & -\frac{2}{3} \\ 1 & -\frac{1}{3} & -1 \\ \frac{1}{3} & -\frac{2}{3} & 0 \end{pmatrix}$, (*h*) $\begin{pmatrix} \frac{1}{3} & 0 & 0 \\ -\frac{1}{3} & 0 & \frac{1}{3} \\ 0 & \frac{2}{3} & \frac{1}{3} \end{pmatrix}$.

8.7.27. Find a matrix A whose Euclidean matrix norm satisfies $\|A^2\| \neq \|A\|^2$.

♢ 8.7.28. Let $K > 0$ be a positive definite matrix. Characterize the matrix norm induced by the inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^T K \mathbf{y}$. *Hint*: Use Exercise 8.5.45.

8.7.29. Let $A = \begin{pmatrix} 1 & 1 \\ 1 & -2 \end{pmatrix}$. Compute the matrix norm $\|A\|$ using the following norms in $\mathbb{R}^2$: (*a*) the weighted ∞ norm $\|\mathbf{v}\| = \max\{2|v_1|, 3|v_2|\}$; (*b*) the weighted 1 norm $\|\mathbf{v}\| = 2|v_1| + 3|v_2|$; (*c*) the weighted inner product norm $\|\mathbf{v}\| = \sqrt{2v_1^2 + 3v_2^2}$; (*d*) the norm associated with the positive definite matrix $K = \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}$.

♢ 8.7.30. Let A be an $n \times n$ matrix with singular value vector $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_r)$. Prove that (*a*) $\|\boldsymbol{\sigma}\|_\infty = \|A\|_2$; (*b*) $\|\boldsymbol{\sigma}\|_2 = \|A\|_F$, the Frobenius norm of Exercise 3.3.51.

Remark. The 1 norm of the singular value vector $\|\boldsymbol{\sigma}\|_1$ also defines a useful matrix norm, known as the *Ky Fan norm*.

♢ 8.7.31. Let A be an $m \times n$ matrix with singular values $\sigma_1, \ldots, \sigma_r$. Prove that

$$\sum_{i=1}^{r} \sigma_i^2 = \sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij}^2. \tag{8.63}$$

8.7.32. Let A be a nonsingular square matrix. Prove the following formulas for its condition number:

$$(a)\ \kappa(A) = \frac{\max\{ \|A\mathbf{u}\| \mid \|\mathbf{u}\| = 1 \}}{\min\{ \|A\mathbf{u}\| \mid \|\mathbf{u}\| = 1 \}}, \qquad (b)\ \kappa(A) = \|A\|_2 \, \|A^{-1}\|_2.$$

8.7.33. Find the pseudoinverse of the following matrices: (*a*) $\begin{pmatrix} 1 & -1 \\ -3 & 3 \end{pmatrix}$, (*b*) $\begin{pmatrix} 1 & -2 \\ 2 & 1 \end{pmatrix}$,

(*c*) $\begin{pmatrix} 2 & 0 \\ 0 & -1 \\ 0 & 0 \end{pmatrix}$, (*d*) $\begin{pmatrix} 0 & 0 & 1 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}$, (*e*) $\begin{pmatrix} 1 & -1 & 1 \\ -2 & 2 & -2 \end{pmatrix}$, (*f*) $\begin{pmatrix} 1 & 3 \\ 2 & 6 \\ 3 & 9 \end{pmatrix}$, (*g*) $\begin{pmatrix} 1 & 2 & 0 \\ 0 & 1 & 1 \\ 1 & 1 & -1 \end{pmatrix}$.

8.7.34. Use the pseudoinverse to find the least squares solution of minimal norm to the following linear systems: (a) $\begin{aligned} x + y &= 1, \\ 3x + 3y &= -2; \end{aligned}$ (b) $\begin{aligned} x + y + z &= 5, \\ 2x - y + z &= 2; \end{aligned}$ (c) $\begin{aligned} x - 3y &= 2, \\ 2x + y &= -1, \\ x + y &= 0. \end{aligned}$

♡ 8.7.35. Prove that the pseudoinverse satisfies the following identities: (a) $(A^+)^+ = A$, (b) $A A^+ A = A$, (c) $A^+ A A^+ = A^+$, (d) $(A A^+)^T = A A^+$, (e) $(A^+ A)^T = A^+ A$.

8.7.36. Suppose $\mathbf{b} \in \operatorname{img} A$ and $\ker A = \{\mathbf{0}\}$. Prove that $\mathbf{x}^\star = A^+ \mathbf{b}$ is the unique solution to the linear system $A\mathbf{x} = \mathbf{b}$. What if $\ker A \neq \{\mathbf{0}\}$?

8.7.37. Choose a direction for each of the edges and write down the incidence matrix A for the graph sketched in Figure 8.4. Verify that its graph Laplacian (8.62) equals $K = A^T A$.

8.7.38. Determine the spectrum for the graphs in Exercise 2.6.3.

8.7.39. Determine the spectrum of a graph given by the edges of (*i*) a triangle; (*ii*) a square; (*iii*) a pentagon. Can you determine the formula for the spectrum of the graph given by an n sided polygon? *Hint*: See Exercise 8.2.49.

8.7.40. Determine the spectrum for the trees in Exercise 2.6.9. Can you make any conjectures about the nature of the spectrum of a graph that is a tree?

8.8 Principal Component Analysis

Singular values and vectors also underlie contemporary statistical data analysis. In particular, the method of *Principal Component Analysis* has assumed an increasingly essential role in a wide range of applications, including data mining, machine learning, image processing, speech recognition, semantics, face recognition, and health informatics; see [**47**] and the references therein. Given a large data matrix — containing many data points belonging to a high-dimensional vector space — the singular vectors associated with the larger singular values indicate the *principal components* of the data, while small singular values indicate relatively unimportant components. Projection onto the low-dimensional subspaces spanned by the dominant singular vectors can expose structure in their otherwise inscrutable large data sets. The earliest descriptions of the method of Principal Component Analysis are to be found in the first half of the twentieth century in the work of the British statistician Karl Pearson, [**65**], and American statistician Harold Hotelling, [**46**].

Variance and Covariance

We begin with a brief description of basic statistical concepts. Suppose that $x_1, \ldots, x_m \in \mathbb{R}$ represent a collection of m measurements of a single physical quantity, e.g., the distance to a star as measured by various physical apparatuses, the speed of a car at a given instant measured by a collection of instruments, a person's IQ as measured by a series of tests, etc. Experimental error, statistical fluctuations, quantum mechanical effects, numerical approximations, and the like imply that the individual measurements will almost certainly not precisely agree. Nevertheless, one wants to know the most likely value of the measured quantity and the degree of confidence that one has in the proposed value. A variety of statistical tests have been devised to resolve these issues, and we refer the reader to, for example, [**20**, **43**, **87**].

The most basic collective quantity of such a data set is its *mean*, or *average*, denoted by

$$\overline{x} = \frac{x_1 + \cdots + x_m}{m}. \tag{8.64}$$

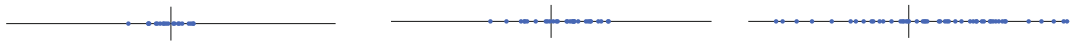

Figure 8.5. Variance.

Barring some inherent statistical or experimental bias, the mean can be viewed as the most likely value, known as the *expected value*, of the quantity being measured, and thus the best bet for its actual value. (More generally, if the measurements are sampled from a known probability distribution, then one works with a suitably weighted average. To keep the formulas relatively simple, we will assume a uniform distribution throughout, and leave generalizations for the reader to pursue by consulting the statistical literature.) Once this has been computed, it will be helpful to *normalize* the measurements to have mean zero, which is done by subtracting off their mean, letting

$$a_i = x_i - \overline{x}, \qquad i = 1, \ldots, m, \qquad \text{with} \qquad \overline{a} = \frac{a_1 + \cdots + a_m}{m} = 0, \tag{8.65}$$

represent the *deviations* of the measurements from their overall mean. It will also help to assemble these quantities into column vectors:

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{pmatrix}, \qquad \mathbf{a} = \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_m \end{pmatrix} = \mathbf{x} - \overline{x}\,\mathbf{e}, \qquad \text{where} \qquad \mathbf{e} = \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} \in \mathbb{R}^m, \qquad \overline{x} = \frac{\mathbf{e} \cdot \mathbf{x}}{m}.$$

Thus, we can write the normalized measurement vector as

$$\mathbf{a} = \left(\mathrm{I} - \frac{1}{m} E \right) \mathbf{x}, \tag{8.66}$$

where I is the $m \times m$ identity matrix and $E = \mathbf{e}\,\mathbf{e}^T$ is the $m \times m$ matrix all of whose entries equal 1.

The measurement *variance* tells us how widely the data points are "scattered" about their mean. As in least squares analysis, this is quantified by summing the squares of their deviations, and denoted by

$$\sigma_x^2 = \nu \big[(x_1 - \overline{x})^2 + \cdots + (x_m - \overline{x})^2 \big] = \nu (a_1^2 + \cdots + a_m^2) = \nu \, \| \mathbf{a} \|^2 = \sigma_a^2, \tag{8.67}$$

where we continue to use the usual Euclidean norm† throughout, and $\nu > 0$ is a certain specified prefactor, which could also be viewed as an overall weight to the Euclidean norm. The square root of the variance is known as the *standard deviation*, and denoted by

$$\sigma = \sigma_x = \sigma_a = \sqrt{\nu} \, \| \mathbf{a} \|. \tag{8.68}$$

The prefactor ν can assume different values depending upon one's statistical objectives; common examples are (*a*) $\nu = 1/m$ for the "naïve" variance; (*b*) $\nu = 1/(m-1)$ (assuming $m > 1$, i.e., there are at least 2 measurements) for an unbiased version; (*c*) $\nu = 1/(m+1)$ for the minimal mean squared estimation of variance; and (*d*) more exotic

† More general probability distributions rely on suitably weighted norms, which can be straightforwardly incorporated into the mathematical framework.

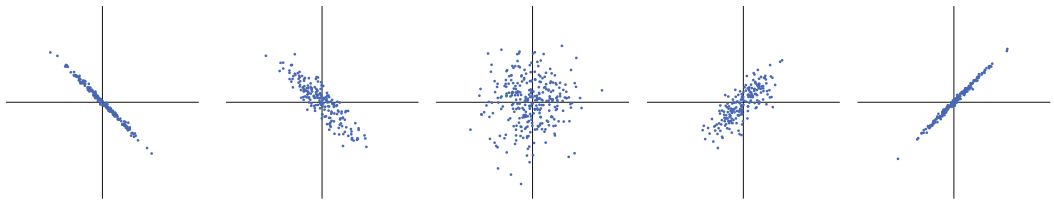

$\rho_{xy} = -.95$ $\rho_{xy} = -.7$ $\rho_{xy} = 0$ $\rho_{xy} = .7$ $\rho_{xy} = .95$

Figure 8.6. Correlation.

choices, e.g., if one desires an unbiased estimation of standard deviation instead of variance, cf. [**43**; p. 349]. Fortunately, apart from the resulting numerical values, the underlying analysis is independent of the prefactor.

The smaller the variance or standard deviation, the less spread out the measurements, and hence the more accurately the mean $\overline{x}$ is expected to approximate the true value of the physical quantity. Figure 8.5 contains several *scatter plots*, in which each real-valued measurement is indicated by a dot and their mean is represented by a small vertical bar. The left plot shows data with relatively small variance, since the measurements are closely clustered about their mean, whereas on the right plot, the variance is large because the data is fairly spread out.

Now suppose we make measurements of several different physical quantities. The individual variances in themselves fail to capture many important features of the resulting data set. For example, Figure 8.6 shows the scatter plots of data sets each representing simultaneous measurements of two quantities, as specified by their horizontal and vertical coordinates. All have the same variances, both individual and cumulative, but clearly represent different interrelationships between the two quantities. In the central plot, they are completely uncorrelated, while on either side they are progressively more correlated (or anti-correlated), meaning that the value of the first measurement is a strong indicator of the value of the second.

This motivates introducing what is known as the *covariance* between a pair of measured quantities $\mathbf{x} = (x_1, x_2, \ldots, x_m)^T$ and $\mathbf{y} = (y_1, y_2, \ldots, y_m)^T$ as the expected value of the product of the deviations from their respective means $\overline{x}, \overline{y}$. We set

$$\sigma_{xy} = \nu \sum_{k=1}^{m} (x_k - \overline{x})(y_k - \overline{y}) = \nu \sum_{k=1}^{m} a_k b_k = \nu\, \mathbf{a} \cdot \mathbf{b}, \tag{8.69}$$

where the normalized vector $\mathbf{a} = (a_1, a_2, \ldots, a_m)^T$ has components $a_k = x_k - \overline{x}$, while $\mathbf{b} = (b_1, b_2, \ldots, b_m)^T$ has components $b_k = y_k - \overline{y}$. Note that, in view of (8.67), the covariance of a set of measurements with itself is its variance: $\sigma_{xx} = \sigma_x^2$. The *correlation* between the two measurement sets is then defined as

$$\rho_{xy} = \frac{\sigma_{xy}}{\sigma_x \sigma_y}, \tag{8.70}$$

and is independent of the prefactor ν. There is an overall bound on the correlation, since the Cauchy–Schwarz inequality (3.18) implies that

$$|\sigma_{xy}| \le \sigma_x \sigma_y, \qquad \text{and hence} \qquad -1 \le \rho_{xy} \le 1. \tag{8.71}$$

The closer ρ_{xy} is to $+1$, the more the variables are correlated; the closer to -1, the more they are anti-correlated, while $\rho_{xy} = 0$ when the variables are uncorrelated. In Figure 8.6,

each scatter plot is labelled by its correlation. Statistically independent variables are automatically uncorrelated, but the converse is not necessarily true, since correlation measures only linear dependencies, and it is possible for nonlinearly dependent variables to nevertheless have zero correlation.

More generally, suppose we are given m measurements of n distinct quantities $\mathbf{x}_1, \ldots, \mathbf{x}_n$. Let X be the $m \times n$ *data matrix* whose entry x_{ij} represents the ith value or measurement of the jth quantity. The column $\mathbf{x}_j = (x_{1j}, \ldots, x_{mj})^T$ of X contains the measurements of the jth quantity, while the row $\boldsymbol{\xi}_i = (\, x_{i1}, \ldots, x_{in}\,)$ is the ith data point. For example, if the data comes from sampling a set $S \subset \mathbb{R}^n$, each row of the data matrix represents a different sample point $\boldsymbol{\xi}_i \in S$. Similarly, in image analysis, each row of the data matrix represents an individual image, whose components are, say, gray scale data for the individual pixels, or color components of pixels — in this case 3 or 4 components per pixel for the RGB or CMKY color scales — or Fourier or wavelet coefficients representing the image, etc.

The (row) vector containing the various measurement means is

$$\overline{\mathbf{x}} = (\overline{x}_1, \ldots, \overline{x}_n) = \frac{1}{m}\, \mathbf{e}^T X. \tag{8.72}$$

We let $\mathbf{a}_i$, with entries $a_{ij} = x_{ij} - \overline{x}_j$ representing the deviations from the mean, denote the corresponding normalized (mean zero) measurement vectors, which form the columns of the *normalized data matrix*

$$A = (\mathbf{a}_1, \ldots, \mathbf{a}_n) = X - \mathbf{e}\,\overline{\mathbf{x}} = \left(\mathrm{I} - \frac{1}{m}\, E \right) X; \tag{8.73}$$

cf. (8.66). The fact that the columns of A all have mean zero is equivalent to the statement that $\mathbf{e} \in \operatorname{coker} A$. We will call the rows $\boldsymbol{\alpha}_i = (a_{i1}, \ldots, a_{in})$ of A the *normalized data points*.

We next define the $n \times n$ *covariance matrix* K of the data set, whose entries equal to the pairwise covariances of the individual measurements:

$$k_{ij} = \sigma_{x_i x_j} = \nu\, \mathbf{a}_i \cdot \mathbf{a}_j = \nu \sum_{k=1}^{m} a_{ki} a_{kj} = \nu \sum_{k=1}^{m} (x_{ki} - \overline{x}_i)(x_{kj} - \overline{x}_j), \quad i, j = 1, \ldots, n. \tag{8.74}$$

The diagonal entries of K are the individual variances: $k_{ii} = \sigma_{x_i x_i} = \sigma_{x_i}^2$. Observe that the covariance matrix is, up to a factor, the symmetric Gram matrix for the normalized measurement vectors $\mathbf{a}_1, \ldots, \mathbf{a}_n$, and hence

$$K = \nu\, A^T A. \tag{8.75}$$

Theorem 3.34 tells us that the covariance matrix is always positive semi-definite, $K \geq 0$, and is positive definite, $K > 0$, unless the columns of A are linearly dependent, or, equivalently, there is a nontrivial exact linear relationship among the normalized measurement vectors: $c_1 \mathbf{a}_1 + \cdots + c_n \mathbf{a}_n = \mathbf{0}$ with $c_1, \ldots, c_n$ not all zero — an unlikely event given the presence of measurement errors and statistical fluctuations.

Remark. Principal Component Analysis requires that *all* n variables be measured the *same* number, m, of times, producing a rectangular $m \times n$ normalized data matrix X, all of whose entries are specified. Extending the analysis to missing or unavailable data is a very active area of contemporary research, which we unfortunately do not have space to examine here. We refer the interested reader to [**24**, **28**] and the references therein.

The Principal Components

The covariance matrix of a data set (8.75) encodes the information concerning the possible linear dependencies and interrelationships among the data. However, due to its potentially large size, it is often not easy to extract the important components and implications. Nor is visualization a good option, since the scatter plots lie in a high-dimensional space. Standard or random projections of high-dimensional data onto two- or three-dimensional subspaces give some limited insight, but the results are highly dependent on the direction of projection and tend to obscure any underlying structure. For example, projecting the data sets in Figure 8.6 onto the x- and y-axes produces more or less the same results, thereby hiding the variety of two-dimensional correlations. A more systematic approach is to locate the so-called principal components of the data, and this leads us us back to the singular value decomposition of the data matrix.

The basic idea behind *Principal Component Analysis*, often abbreviated PCA, is to focus on directions in which the variance of the data is especially large. Given the $m \times n$ normalized data matrix A, we define the first principal direction as that in which the data experiences the most variance. By "direction", we mean a line through the origin in $\mathbb{R}^n$, and the variance is computed from the orthogonal projection of the data measurements onto the line. Each line is spanned by a unit vector $\mathbf{u} = (u_1, u_2, \ldots, u_n)^T$ with $\|\mathbf{u}\| = 1$. (Actually, there are two unit vectors, $\pm\mathbf{u}$, in each line, but, as we will see, it doesn't matter which one we choose.) The orthogonal projection of the ith normalized data point $\boldsymbol{\alpha}_i = (a_{i1}, \ldots, a_{in})$ onto the line spanned by $\mathbf{u}$ is given by the projection formula (4.41), namely $p_i = \boldsymbol{\alpha}_i \mathbf{u}$, $i = 1, \ldots, m$. The result is the projected measurement vector

$$\mathbf{p} = \begin{pmatrix} p_1 \\ p_2 \\ \vdots \\ p_m \end{pmatrix} = A\,\mathbf{u} = u_1 \mathbf{a}_1 + \cdots + u_n \mathbf{a}_n. \tag{8.76}$$

Our goal is to maximize its variance

$$\sigma_p^2 = \nu\,\|\mathbf{p}\|^2 = \nu\,\|A\,\mathbf{u}\|^2 = \nu\,(A\,\mathbf{u})^T A\,\mathbf{u} = \nu\,(\mathbf{u}^T A^T A\,\mathbf{u}) = \nu\,(\mathbf{u}^T K\,\mathbf{u}). \tag{8.77}$$

over all possible choices of unit vector $\mathbf{u}$. But, ignoring the irrelevant factor $\nu > 0$, this is precisely the maximization problem that was solved by Theorem 8.40. We thus immediately deduce that the first principal direction is given by the dominant unit eigenvector $\mathbf{u} = \mathbf{q}_1$ of the covariance matrix $K = \nu\,A^T A$, or, equivalently, the dominant unit singular vector of the normalized data matrix A. The maximum variance is, up to a factor, the dominant, or largest, eigenvalue, or, equivalently, the square of the dominant singular value, namely, $\max_{\mathbf{u}} \sigma_p^2 = \nu\,\lambda_1 = \nu\,\sigma_1^2$, while the dominant singular value is, again up to a factor, the maximal standard deviation of the projected measurements: $\max_{\mathbf{u}} \sigma_p = \sqrt{\nu}\,\sigma_1$.

The second principal direction is assumed to be orthogonal to the first, so as to avoid contaminating it with the already noted direction of maximal variance, and is to be chosen so that the variance of its projected measurements is maximized among all such orthogonal

directions. Thus, the second principal direction will maximize σ_p^2, as given by (8.77), over all unit vectors $\mathbf{u}$ satisfying $\mathbf{u}\cdot\mathbf{q}_1 = 0$. More generally, given the first $j-1$ principal directions $\mathbf{q}_1, \ldots, \mathbf{q}_{j-1}$, the jth principal component is in the direction $\mathbf{u} = \mathbf{q}_j$ that maximizes the variance

$$\sigma_p^2 = \nu\,(\mathbf{u}^T K\,\mathbf{u}) \quad \text{over all vectors } \mathbf{u} \text{ satisfying} \quad \|\,\mathbf{u}\,\| = 1, \quad \mathbf{u}\cdot\mathbf{q}_1 = \;\cdots\; = \mathbf{u}\cdot\mathbf{q}_{j-1} = 0.$$

Theorem 8.42 immediately implies that $\mathbf{q}_j$ is a unit eigenvector of K associated with its jth largest eigenvalue $\lambda_j = \sigma_j^2$, which therefore is, up to a factor, the jth principal variance. Summarizing, we have proved the Fundamental Theorem of Principal Component Analysis.

Theorem 8.76. The jth principal direction of a normalized data matrix A is its jth unit singular vector $\mathbf{q}_j$. The corresponding principal standard deviation $\sqrt{\nu}\,\sigma_j$ is proportional to its jth singular value σ_j.

In applications, one designates a certain number, say k, of the dominant (largest) variances $\nu\,\sigma_1^2 \geq \nu\,\sigma_2^2 \geq \;\cdots\; \geq \nu\,\sigma_k^2$, as "principal" and the corresponding unit singular vectors $\mathbf{q}_1, \ldots, \mathbf{q}_k$ as the *principal directions*. The value of k depends on the user and on the application. For example, in visualization, $k = 2$ or 3 in order to plot these components of the data in the plane or in space. More generally, one could specify k based on some overall size threshold, or where there is a perceived gap in the magnitudes of the variances. Another choice is such that the principal variances are those that make up some large fraction, e.g., $\mu = 95\%$, of the total variance:

$$\nu\left(\sigma_1^2 + \;\cdots\; + \nu\,\sigma_k^2\right) = \mu\left(\nu \sum_{i=1}^{r} \sigma_i^2\right) = \mu\,\nu \sum_{i,j=1}^{n} a_{ij}^2 = \mu\,\nu \sum_{i=1}^{n} \sigma_{x_i}^2, \tag{8.78}$$

where the next-to-last equality follows from (8.63).

Theorem 8.74 says that, in the latter cases, the number of principal components, i.e., the number of significant singular values, will determine the approximate rank k of the covariance matrix K and hence the data matrix A, also. Thus, the normalized data (approximately) lies in a k-dimensional subspace. Further, the variance in any direction orthogonal to principal directions is relatively small and hence relatively unimportant. As a consequence, dimensional reduction by orthogonally projecting the data vectors onto the k-dimensional subspace spanned by the principal directions (singular vectors) $\mathbf{q}_1, \ldots, \mathbf{q}_k$, serves to eliminate significant redundancies.

The coordinates of the data in the jth principal direction are provided by orthogonal projection, namely the entries of the image vectors $A\,\mathbf{q}_j = \sigma_j\,\mathbf{p}_j$. Thus, approximating the data by its k principal components coincides with the closest rank k approximation $A_k = P_k \Sigma_k Q_k^T$ to the data matrix A given by Theorem 8.74. Moreover, rewriting the data in terms of the *principal coordinates*, meaning those supplied by the principal directions, serves to diagonalize the principal covariance matrix $\nu\,K_k = \nu\,A_k^T A_k$, since, as in (8.58),

$$Q_k^T(\nu\,A_k^T A_k)\,Q_k = \nu\,\Sigma_k^2,$$

the result being a diagonal matrix containing the k principal variances along the diagonal. This has the important consequence that the covariance between any two principal components of the data is zero, and thus the principal components are all uncorrelated! In geometric terms, the original data tends to form an ellipsoid in the high-dimensional data space, and the principal directions are aligned with its principal semi-axis, thereby conforming to and exposing the intrinsic geometry of the data set.

Thus, to perform Principal Component Analysis on a complete data set, consisting of m measurements of n quantities, one forms the $m \times n$ normalized data matrix A whose $(i,j)^{\text{th}}$ entry equals the ith measurement of the jth variable minus the mean of the jth variable. The principal directions are the first k singular vectors of A, meaning the eigenvectors of the positive (semi-)definite covariance matrix $K = \nu A^T A$, while the principal variances are, up to the overall factor ν, the corresponding eigenvalues or equivalently the squares of the singular values. The principal coordinates are given by the entries of the resulting matrix $P_k \Sigma_k$ whose columns are the image vectors $\sigma_j \mathbf{p}_j = A \mathbf{q}_j$ specifying the measurements in the principal directions.

Exercises

Note: For simplicity, take the prefactor $\nu = 1$ unless otherwise indicated.

8.8.1. Find the mean, the variance, and the standard deviation of the following data sets: (*a*) $1.1, 1.3, 1.5, 1.55, 1.6, 1.9, 2, 2.1$; (*b*) $2., .9, .7, 1.5, 2.6, .3, .8, 1.4$; (*c*) $-2.9, -.5, .1, -1.5, -3.6, 1.3, .4, -.7$; (*d*) $1.1, .2, .1, .6, 1.3, -.4, -.1, .4$; (*e*) $.9, -.4, -.8, ., 1., -1.6, -1.2, -.7$.

8.8.2. Find the mean, the variance, and the standard deviation of the data sets $\{ f(x) \,|\, x = i/10,\ i = -10, \ldots, 10 \}$ associated with the following functions $f(x)$:
(*a*) $3x + 1$, (*b*) x^2, (*c*) $x^3 - 2x$, (*d*) e^{-x}, (*e*) $\tan^{-1} x$.

8.8.3. Determine the variance and standard deviation of the *normally distributed* data points $\{ e^{-x^2/\sigma} \,|\, x = i/10,\ i = -10, \ldots, 10 \}$ for $\sigma = 1, 2$, and 10.

8.8.4. Prove that $\sigma_{xy} = \overline{xy} - \overline{x}\,\overline{y}$, where $\overline{x}$ and $\overline{y}$ are the means of $\{ x_i \}$ and $\{ y_i \}$, respectively, while $\overline{xy}$ denotes the mean of the product variable $\{ x_i y_i \}$.

♢ 8.8.5. Show that one can compute the variance of a set of measurements without reference to the mean by the following formula

$$\sigma_x^2 = \frac{\nu}{2m} \sum_{i=1}^{m} \sum_{j=1}^{m} (x_i - x_j)^2 = \frac{\nu}{m} \sum_{i<j} (x_i - x_j)^2 .$$

8.8.6. Let A be an $m \times n$ matrix that is normalized, meaning that each of its column sums is zero. Show that AB, where B is any $n \times k$ matrix, is also normalized.

8.8.7. Given a singular value decomposition $A = P \Sigma Q^T$, prove that if the columns of A have zero mean, then so do the columns of $P \Sigma$.

♣ 8.8.8. Construct the 5×5 covariance matrix for the 5 data sets in Exercise 8.8.1 and find its principal variances and principal directions. What do you think is the dimension of the subspace the data lies in?

♣ 8.8.9. For each of the following subsets $S \subset \mathbb{R}^3$, (*i*) Compute a fairly dense sample of data points $z_i \in S$; (*ii*) find the principal components of your data set, using $\mu = .95$ in the criterion in (8.78); (*iii*) using your principal components, estimate the dimension of the set S. Does your estimate coincide with the actual dimension? If not, explain any discrepancies.

(*a*) The line segment $S = \{ (t+1, 3t-1, -2t)^T \mid -1 \le t \le 1 \}$;
(*b*) the set of points z on the three coordinate axes with Euclidean norm $\| z \| \le 1$;
(*c*) the set of "probability vectors" $S = \{ (x, y, z)^T \mid 0 \le x, y, z \le 1,\ x + y + z = 1 \}$;
(*d*) the unit ball $S = \{ \| z \| \le 1 \}$ for the Euclidean norm;
(*e*) the unit sphere $S = \{ \| z \| = 1 \}$ for the Euclidean norm;
(*f*) the unit ball $S = \{ \| z \|_\infty \le 1 \}$ for the ∞ norm;
(*g*) the unit sphere $S = \{ \| z \|_\infty = 1 \}$ for the ∞ norm.

♣ 8.8.10. Using the Euclidean norm, compute a fairly dense sample of points on the unit sphere $S = \{ \mathbf{x} \in \mathbb{R}^3 \mid \|\mathbf{x}\| = 1 \}$. (*a*) Set $\mu = .95$ in (8.78), and then find the principal components of your data set. Do they indicate the two-dimensional nature of the sphere? If not, why not? (*b*) Now look at the subset of your data that is within a distance $r > 0$ of the north pole, i.e., $\| \mathbf{x} - (0, 0, 1)^T \| \leq r$, and compute its principal components. How small does r need to be to reveal the actual dimension of S? Interpret your calculations.

♢ 8.8.11. Show that the first principal direction $\mathbf{q}_1$ can be characterized as the direction of the line that minimizes the sums of the squares of its distances to the data points.

♡ 8.8.12. Let $\mathbf{x}_i = (x_i, y_i)$, $i = 1, \ldots, m$, be a set of data points in the plane. Suppose $L^* \subset \mathbb{R}^2$ is the line that minimizes the sums of the squares of the distances from the data points to it, i.e., $\operatorname{dist}(\mathbf{x}, L) = \sum_{i=1}^{m} \operatorname{dist}(\mathbf{x}_i, L)$, among all lines $L \subset \mathbb{R}^2$. (*a*) Prove that $\overline{\mathbf{x}} = (\overline{x}, \overline{y}) \in L^*$. (*b*) Use Exercise 8.8.11 to find L^*. (*c*) Apply your result to the data points in Example 5.14 and compare the resulting line L^* with the least squares line that was found there.