Chapter I

Background in Matrix Theory and Linear Algebra

This chapter reviews basic matrix theory and introduces some of the elementary notation used throughout the book. Matrices are objects that represent linear mappings between vector spaces. The notions that will be predominantly used in this book are very intimately related to these linear mappings and it is possible to discuss eigenvalues of linear operators without ever mentioning their matrix representations. However, to the numerical analyst, or the engineer, any theory that would be developed in this manner would be insufficient in that it will not be of much help in developing or understanding computational algorithms. The abstraction of linear mappings on vector spaces does however provide very concise definitions and some important theorems.

1. Matrices

When dealing with eigenvalues it is more convenient, if not more relevant, to manipulate complex matrices rather than real matrices. A complex $n \times m$ matrix A is an $n \times m$ array of complex numbers

$$a_{ij}, \quad i = 1, \ldots, n, \; j = 1, \ldots, m.$$

The set of all $n \times m$ matrices is a complex vector space denoted by $\mathbb{C}^{n \times m}$. The main operations with matrices are the following:

- Addition: $C = A + B$, where A, B and C are matrices of size $n \times m$ and
$$c_{ij} = a_{ij} + b_{ij} \ ,$$
$i = 1, 2, \ldots n$, $j = 1, 2, \ldots m$.

- Multiplication by a scalar: $C = \alpha A$, where $c_{ij} = \alpha \ a_{ij}$.

- Multiplication by another matrix:
$$C = AB,$$
where $A \in \mathbb{C}^{n \times m}, B \in \mathbb{C}^{m \times p}, C \in \mathbb{C}^{n \times p}$, and
$$c_{ij} = \sum_{k=1}^{m} a_{ik} b_{kj}.$$

A notation that is often used is that of column vectors and row vectors. The column vector $a_{.j}$ is the vector consisting of the j-th column of A, i.e., $a_{.j} = (a_{ij})_{i=1,\ldots,n}$. Similarly we will use the notation $a_{i.}$ to denote the i-th row of the matrix A. For example, we may write that

$$A = (a_{.1}, a_{.2}, \ldots, a_{.m}) \ .$$

or

$$A = \begin{pmatrix} a_{1.} \\ a_{2.} \\ . \\ . \\ a_{n.} \end{pmatrix}$$

The *transpose* of a matrix A in $\mathbb{C}^{n\times m}$ is a matrix C in $\mathbb{C}^{m\times n}$ whose elements are defined by $c_{ij} = a_{ji}, i = 1, \ldots, m,\ j = 1, \ldots, n$. The transpose of a matrix A is denoted by A^T. It is more relevant in eigenvalue problems to use the *transpose conjugate* matrix denoted by A^H and defined by

$$A^H = \bar{A}^T = \overline{A^T}$$

in which the bar denotes the (element-wise) complex conjugation.

Finally, we should recall that matrices are strongly related to linear mappings between vector spaces of finite dimension. They are in fact representations of these transformations with respect to two given bases; one for the initial vector space and the other for the image vector space.

2. Square Matrices and Eigenvalues

A matrix belonging to $\mathbb{C}^{n\times n}$ is said to be square. Some notions are only defined for square matrices. A square matrix which is very important is the identity matrix

$$I = \{\delta_{ij}\}_{i,j=1,\ldots,n}$$

where δ_{ij} is the Kronecker symbol. The identity matrix satisfies the equality $AI = IA = A$ for every matrix A of size n. The inverse of a matrix, when it exists, is a matrix C such that $CA = AC = I$. The inverse of A is denoted by A^{-1}.

The determinant of a matrix may be defined in several ways. For simplicity we adopt here the following recursive definition. The determinant of a 1×1 matrix (a) is defined as the scalar a. Then the determinant of an $n \times n$ matrix is given by

$$\det(A) = \sum_{j=1}^{n} (-1)^{j+1} a_{1j} \det(A_{1j})$$

where A_{1j} is an $(n-1) \times (n-1)$ matrix obtained by deleting the 1-st row and the $j-th$ column of A. The determinant of a matrix determines whether or not a matrix is singular since A is singular if and only if its determinant is zero. We have the following simple properties:

- $\det(AB) = \det(BA)$,
- $\det(A^T) = \det(A)$,
- $\det(\alpha A) = \alpha^n \det(A)$,
- $\det(\bar{A}) = \overline{\det(A)}$,
- $\det(I) = 1$.

From the above definition of the determinant it can be shown by induction that the function that maps a given complex value λ to the value $p_A(\lambda) = \det(A - \lambda I)$ is a polynomial of degree n (Problem P-1.6). This is referred to as the *characteristic polynomial* of the matrix A.

Definition 1.1 *A complex scalar λ is called an eigenvalue of the square matrix A if there exists a nonzero vector u of $\mathbb{C}^n$ such that $Au = \lambda u$. The vector u is called an eigenvector of A associated with λ. The set of all the eigenvalues of A is referred to as the spectrum of A and is denoted by $\sigma(A)$.*

An eigenvalue of A is a root of the characteristic polynomial. Indeed λ is an eigenvalue of A iff $\det(A - \lambda I) \equiv p_A(\lambda) = 0$. So there are at most n distinct eigenvalues. The maximum modulus of the eigenvalues is called *spectral radius* and is denoted by $\rho(A)$:

$$\rho(A) = \max_{\lambda \in \sigma(A)} |\lambda|.$$

The *trace* of a matrix is equal to the sum of all its diagonal elements,

$$\mathrm{tr}(A) = \sum_{i=1}^{n} a_{ii}.$$

It can be easily shown that this is also equal to the sum of its eigenvalues counted with their multiplicities as roots of the characteristic polynomial.

Proposition 1.1 *If λ is an eigenvalue of A then $\bar{\lambda}$ is an eigenvalue of A^H. An eigenvector v of A^H associated with the eigenvalue $\bar{\lambda}$ is called left eigenvector of A.*

When a distinction is necessary, an eigenvector of A is often called a right eigenvector. Thus the eigenvalue λ and the right and left eigenvectors, u and v, satisfy the relations

$$Au = \lambda u \ , \quad v^H A = \lambda v^H$$

or, equivalently,

$$u^H A^H = \bar{\lambda} u^H \ , \quad A^H v = \bar{\lambda} v \ .$$

3. Types of Matrices

The properties of eigenvalues and eigenvectors of square matrices will sometimes depend on special properties of the matrix A. For example, the eigenvalues or eigenvectors of the following types of matrices will all have some special properties.

- *Symmetric matrices:* $A^T = A$;
- *Hermitian matrices:* $A^H = A$;
- *Skew-symmetric matrices:* $A^T = -A$;
- *Skew-Hermitian matrices:* $A^H = -A$;
- *Normal matrices:* $A^H A = A A^H$;
- *Nonnegative matrices:* $a_{ij} \geq 0, \ i, j = 1, \ldots, n$ (similar definition for nonpositive, positive, and negative matrices);

- *Unitary matrices:* $Q^H Q = I$.

Often, a matrix Q such that $Q^H Q$ is diagonal is called orthogonal. It is worth noting that a unitary matrix Q is a matrix whose inverse is its transpose conjugate Q^H.

Some matrices have particular structures that are often convenient for computational purposes and play important roles in numerical analysis. The following list though incomplete, gives an idea of the most important special matrices arising in applications and algorithms.

- *Diagonal matrices:* $a_{ij} = 0$ for $j \neq i$. Notation:

$$A = \text{diag}\ (a_{11}, a_{22}, \ldots, a_{nn}).$$

- *Upper triangular matrices:* $a_{ij} = 0$ for $i > j$.
- *Lower triangular matrices:* $a_{ij} = 0$ for $i < j$.
- *Upper bidiagonal matrices:* $a_{ij} = 0$ for $j \neq i$ or $j \neq i + 1$.
- *Lower bidiagonal matrices:* $a_{ij} = 0$ for $j \neq i$ or $j \neq i - 1$.
- *Tridiagonal matrices:* $a_{ij} = 0$ for any pair i, j such that $|j - i| > 1$. Notation:

$$A = \text{tridiag}\ (a_{i,i-1}, a_{ii}, a_{i,i+1}).$$

- *Banded matrices:* there exist two integers m_l and m_u such that $a_{ij} \neq 0$ only if $i - m_l \leq j \leq i + m_u$. The number $m_l + m_u + 1$ is called the bandwidth of A.
- *Upper Hessenberg matrices:* $a_{ij} = 0$ for any pair i, j such that $i > j + 1$. One can define lower Hessenberg matrices similarly.
- *Outer product matrices:* $A = uv^H$, where both u and v are vectors.

- *Permutation matrices:* the columns of A are a permutation of the columns of the identity matrix.
- *Block diagonal matrices:* generalizes the diagonal matrix by replacing each diagonal entry by a matrix. Notation:

$$A = \text{diag}\ (A_{11}, A_{22}, \ldots, A_{mm})\,.$$

- *Block tri-diagonal matrices:* generalizes the tri-diagonal matrix by replacing each nonzero entry by a square matrix. Notation:

$$A = \text{tridiag}\ (A_{i,i-1}, A_{ii}, A_{i,i+1})\,.$$

The above properties emphasize structure, i.e., positions of the nonzero elements with respect to the zeros, and assume that there are many zero elements or that the matrix is of low rank. No such assumption is made for, say, orthogonal or symmetric matrices.

4. Vector Inner Products and Norms

We define the Hermitian inner product of the two vectors $x = (x_i)_{i=1,\ldots,n}$ and $y = (y_i)_{i=1,\ldots,n}$ of $\mathbb{C}^n$ as the complex number

$$(x, y) = \sum_{i=1}^{n} x_i \bar{y}_i, \tag{1.1}$$

which is often rewritten in matrix notation as

$$(x, y) = y^H x.$$

A vector norm on $\mathbb{C}^n$ is a real-valued function on $\mathbb{C}^n$, which satisfies the following three conditions,

$$\begin{aligned} &\|x\| \geq 0 \quad \forall\ x, \quad \text{and} \quad \|x\| = 0 \text{ iff } x = 0; \\ &\|\alpha x\| = |\alpha| \|x\|, \quad \forall\ x \in \mathbb{C}^n, \quad \forall \alpha \in \mathbb{C}\ ; \\ &\|x + y\| \leq \|x\| + \|y\|, \quad \forall x, y \in \mathbb{C}^n\ . \end{aligned}$$

Associated with the inner product (1.1) is the Euclidean norm of a complex vector defined by

$$\|x\|_2 = (x, x)^{1/2} .$$

A fundamental additional property in matrix computations is the simple relation

$$(Ax, y) = (x, A^H y) \quad \forall x, y \in \mathbb{C}^n \tag{1.2}$$

the proof of which is straightforward. The following proposition is a consequence of the above equality.

Proposition 1.2 *Unitary matrices preserve the Hermitian inner product, i.e.,* $(Qx, Qy) = (x, y)$ *for any unitary matrix* Q*.*

Proof. Indeed $(Qx, Qy) = (x, Q^H Qy) = (x, y)$. ∎

In particular a unitary matrix preserves the 2-norm metric, i.e., it is isometric with respect to the 2-norm.

The most commonly used vector norms in numerical linear algebra are special cases of the Hölder norms

$$\|x\|_p = \left(\sum_{i=1}^{n} |x_i|^p \right)^{1/p} . \tag{1.3}$$

Note that the limit of $\|x\|_p$ when p tends to infinity exists and is equal to the maximum modulus of the x_i's. This defines a norm denoted by $\|.\|_\infty$. The cases $p = 1$, $p = 2$, and $p = \infty$ lead to the most important norms in practice,

$$\|x\|_1 = |x_1| + |x_2| + \cdots + |x_n|$$
$$\|x\|_2 = \left[|x_1|^2 + |x_2|^2 + \cdots + |x_n|^2 \right]^{1/2}$$
$$\|x\|_\infty = \max_{i=1,..,n} |x_i| .$$

A useful relation concerning the 2-norm is the so-called Cauchy-Schwartz inequality:

$$|(x, y)| \le \|x\|_2 \|y\|_2 .$$

5. Matrix Norms

For a general matrix A in $\mathbb{C}^{n\times m}$ we define a special set of norms of matrices as follows

$$\|A\|_{pq} = \max_{x\in\mathbb{C}^m,\ x\neq 0} \frac{\|Ax\|_p}{\|x\|_q} . \tag{1.4}$$

We say that the norms $\|.\|_{pq}$ are induced by the two norms $\|.\|_p$ and $\|.\|_q$. These norms satisfy the usual properties of norms, i.e.,

$$\|A\| \geq 0 \quad \forall A \ \in \mathbb{C}^{n\times m} \quad \text{and} \quad \|A\| = 0 \quad \text{iff} \quad A = 0 ;$$
$$\|\alpha A\| = |\alpha| \|A\|, \forall A \ \in \mathbb{C}^{n\times m}, \quad \forall \alpha \in \mathbb{C} ;$$
$$\|A + B\| \leq \|A\| + \|B\|, \quad \forall A, B \ \in \mathbb{C}^{n\times m} .$$

Again the most important cases are the ones associated with the cases $p, q = 1, 2, \infty$. The case $q = p$ is of particular interest and the associated norm $\|.\|_{pq}$ is simply denoted by $\|.\|_p$.

A fundamental property of these norms is that

$$\|AB\|_p \leq \|A\|_p \|B\|_p,$$

which is an immediate consequence of the definition (1.4). Matrix norms that satisfy the above property are sometimes called *consistent*. As a result of the above inequality, for example, we have that for any square matrix A,

$$\|A^n\|_p \leq \|A\|_p^n ,$$

which implies in particular that the matrix A^n converges to zero if *any* of its p-norms is less than 1.

The Frobenius norm of a matrix is defined by

$$\|A\|_F = \left(\sum_{j=1}^{m} \sum_{i=1}^{n} |a_{ij}|^2 \right)^{1/2} . \tag{1.5}$$

This can be viewed as the 2-norm of the column (or row) vector in $\mathbb{C}^{n^2}$ consisting of all the columns (resp. rows) of A listed from

1 to m (resp. 1 to n). It can easily be shown that this norm is also consistent, in spite of the fact that is not induced by a pair of vector norms, i.e., it is not derived from a formula of the form (1.4), see Problem P-1.3. However, it does not satisfy some of the other properties of the p-norms. For example, the Frobenius norm of the identity matrix is not unity. To avoid these difficulties, *we will only use the term matrix norm for a norm that is induced by two norms as in the definition (1.4).* Thus, we will not consider the Frobenius norm to be a proper matrix norm, according to our conventions, even though it is consistent.

It can be shown that the norms of matrices defined above satisfy the following equalities that lead to alternative definitions that are often easier to work with.

$$\|A\|_1 = \max_{j=1,..,m} \sum_{i=1}^{n} |a_{ij}| \; ; \tag{1.6}$$

$$\|A\|_\infty = \max_{i=1,..,n} \sum_{j=1}^{m} |a_{ij}| \; ; \tag{1.7}$$

$$\|A\|_2 = \left[\rho(A^H A)\right]^{1/2} = \left[\rho(AA^H)\right]^{1/2} \; ; \tag{1.8}$$

$$\|A\|_F = \left[\mathrm{tr}(A^H A)\right]^{1/2} = \left[\mathrm{tr}(AA^H)\right]^{1/2} \; . \tag{1.9}$$

As will be shown in Section 5, the eigenvalues of $A^H A$ are nonnegative. Their square roots are called singular values of A and are denoted by $\sigma_i, i = 1, \ldots, m$. Thus, the relation (1.8) shows that $\|A\|_2$ is equal to σ_1, the largest singular value of A.

Example 1.1 From the above properties, it is clear that the spectral radius $\rho(A)$ is equal to the 2-norm of a matrix when the matrix is Hermitian. However, it is not a matrix norm in general. For example, the first property of norms is not satisfied, since for

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

we have $\rho(A) = 0$ while $A \neq 0$. The triangle inequality is also not satisfied for the pair A, and $B = A^T$ where A is defined above. Indeed,

$$\rho(A + B) = 1 \quad \text{while} \quad \rho(A) + \rho(B) = 0.$$

6. Subspaces

A subspace of $\mathbb{C}^n$ is a subset of $\mathbb{C}^n$ that is also a complex vector space. The set of all linear combinations of a set of vectors $G = \{a_1, a_2, ..., a_q\}$ of $\mathbb{C}^n$ is a vector subspace called the linear span of G,

$$\begin{aligned} \text{span}\{G\} &= \text{span}\{a_1, a_2, \ldots, a_q\} \\ &= \left\{ z \ \in \mathbb{C}^n \mid z = \sum_{i=1}^{q} \alpha_i a_i \ ; \ \{\alpha\}_{i=1,\ldots,q} \in \mathbb{C}^q \right\} . \end{aligned}$$

If the a_i's are linearly independent, then each vector of span$\{G\}$ admits a unique expression as a linear combination of the a_i's. The set G is then called a basis of the subspace *span*$\{G\}$.

Given two vector subspaces S_1 and S_2, their sum S is a subspace defined as the set of all vectors that are equal to the sum of a vector of S_1 and a vector of S_2. The intersection of two subspaces is also a subspace. If the intersection of S_1 and S_2 is reduced to $\{0\}$ then the sum of S_1 and S_2 is called their direct sum and is denoted by $S = S_1 \bigoplus S_2$. When S is equal to $\mathbb{C}^n$ then every vector x of $\mathbb{C}^n$ can be decomposed in a unique way as the sum of an element x_1 of S_1 and an element x_2 of S_2. The transformation P that maps x into x_1 is a linear transformation that is idempotent ($P^2 = P$). It is called a *projector*, onto S_1 along S_2.

Two subspaces of importance that are associated with a matrix A of $\mathbb{C}^{n\times m}$ are its *range* defined by

$$\text{Ran}(A) = \{Ax \mid x \in \mathbb{C}^m\} \tag{1.10}$$

and its *kernel* or null space

$$\text{Ker}(A) = \{x \in \mathbb{C}^m \mid Ax = 0 \ \} \ .$$

The range of A is clearly equal to the linear *span* of its columns. The *rank* of a matrix is equal to the dimension of the range of A.

A subspace S is said to be *invariant* under a (square) matrix A whenever $AS \subset S$. In particular for any eigenvalue λ of A

the subspace $\text{Ker}(A - \lambda I)$ is invariant under A. The subspace $\text{Ker}(A - \lambda I)$ is called the eigenspace associated with λ and consists of all the eigenvectors of A associated with λ and the vector 0.

7. Orthogonal Vectors and Subspaces

A set of vectors $G = \{a_1, a_2, \ldots, a_r\}$ is said to be *orthogonal* if

$$(a_i, a_j) = 0 \quad \text{when} \quad i \neq j$$

It is *orthonormal* if in addition every vector of G has a 2-norm equal to unity. Every subspace admits an orthonormal basis which is obtained by taking any basis and "orthonormalizing" it. The orthonormalization can be achieved by an algorithm referred to as the Gram-Schmidt process which we now describe. Given a set of linearly independent vectors $\{x_1, x_2, \ldots, x_r\}$, we first normalize the vector x_1, i.e., we divide it by its 2-norm, to obtain the scaled vector q_1. Then x_2 is orthogonalized against the vector q_1 by subtracting from x_2 a multiple of q_1 to make the resulting vector orthogonal to q_1, i.e.,

$$x_2 \leftarrow x_2 - (x_2, q_1)q_1.$$

The resulting vector is again normalized to yield the second vector q_2. The i-th step of the Gram-Schmidt process consists of orthogonalizing the vector x_i against all previous vectors q_j.

ALGORITHM **1.1 Gram-Schmidt**

1. **Start**: *Compute* $r_{11} := \|x_1\|_2$. *If* $r_{11} = 0$ *stop, else* $q_1 := x_1/r_{11}$.
2. **Loop**: *For* $j = 2, \ldots, r$ *do:*
   (a) *Compute* $r_{ij} := (x_j, q_i)$ *for* $i = 1, 2, \ldots, j-1$,
   (b) $\hat{q} := x_j - \sum_{i=1}^{j-1} r_{ij} q_i$,

(c) $r_{jj} := \|\hat{q}\|_2$,

(d) *If* $r_{jj} = 0$ *then stop, else* $q_j := \hat{q}/r_{jj}$.

It is easy to prove that the above algorithm will not break down, i.e., all r steps will be completed, if and only if the family of vectors $x_1, x_2, \ldots, x_r$ is linearly independent. From 2-(b) and 2-(c) it is clear that at every step of the algorithm the following relation holds:

$$x_j = \sum_{i=1}^{j} r_{ij} q_i \ .$$

If we let $X = [x_1, x_2, \ldots, x_r]$, $Q = [q_1, q_2, \ldots, q_r]$, and if R denotes the $r \times r$ upper triangular matrix whose nonzero elements are the r_{ij} defined in the algorithm, then the above relation can be written as

$$X = QR \ . \tag{1.11}$$

This is called the QR decomposition of the $n \times r$ matrix X. Thus, from what was said above the QR decomposition of a matrix exists whenever the column vectors of X form a linearly independent set of vectors.

The above algorithm is the standard Gram-Schmidt process. There are other formulations of the same algorithm which are mathematically equivalent but have better numerical properties. The Modified Gram-Schmidt algorithm (MGSA) is one such alternative.

ALGORITHM **1.2 Modified Gram-Schmidt**

1. *Start: define* $r_{11} := \|x_1\|_2$. *If* $r_{11} = 0$ *stop, else* $q_1 := x_1/r_{11}$.

2. *Loop: For* $j = 2, \ldots, r$ *do:*

   (a) *Define* $\hat{q} := x_j$,

   (b) *For* $i = 1, \ldots, j-1$, *do* $\left\{ \begin{array}{l} r_{ij} := (\hat{q}, q_i) \\ \hat{q} := \hat{q} - r_{ij} q_i \end{array} \right.$

   (c) *Compute* $r_{jj} := \|\hat{q}\|_2$,

(d) *If* $r_{jj} = 0$ *then stop, else* $q_j := \hat{q}/r_{jj}$.

A vector that is orthogonal to all the vectors of a subspace S is said to be orthogonal to that subspace. The set of all the vectors that are orthogonal to S is a vector subspace called the *orthogonal complement* of S and denoted by $S^{\perp}$. The space $\mathbb{C}^n$ is the direct sum of S and its orthogonal complement. The projector onto S along its orthogonal complement is called an *orthogonal projector* onto S. If $V = [v_1, v_2, \ldots, v_r]$ is an orthonormal matrix then $V^H V = I$, i.e., V is orthogonal. However, VV^H is not the identity matrix but represents the orthogonal projector onto span$\{V\}$, see Section 1 of Chapter III for details.

8. Canonical Forms of Matrices

In this section we will be concerned with the reduction of square matrices into matrices that have simpler forms, such as diagonal or bidiagonal, or triangular. By reduction we mean a transformation that preserves the eigenvalues of a matrix.

Definition 1.2 *Two matrices A and B are said to be similar if there is a nonsingular matrix X such that*

$$A = XBX^{-1}$$

The mapping $B \rightarrow A$ is called a similarity transformation.

It is clear that *similarity* is an equivalence relation. Similarity transformations preserve the eigenvalues of matrix. An eigenvector u_B of B is transformed into the eigenvector $u_A = Xu_B$ of A. In effect, a similarity transformation amounts to representing the matrix B in a different basis.

We now need to define some terminology.

1. An eigenvalue λ of A is said to have *algebraic multiplicity* μ if it is a root of multiplicity μ of the characteristic polynomial.

2. If an eigenvalue is of algebraic multiplicity one it is said to be *simple*. A nonsimple eigenvalue is said to be *multiple*.

3. An eigenvalue λ of A is said to have *geometric multiplicity* γ if the maximum number of independent eigenvectors associated with it is γ. In other words the geometric multiplicity γ is the dimension of the eigenspace Ker $(A - \lambda I)$.

4. A matrix is said to be *derogatory* if the geometric multiplicity of at least one of its eigenvalues is larger than one.

5. An eigenvalue is said to be *semi-simple* if its algebraic multiplicity is equal to its geometric multiplicity. An eigenvalue that is not semi-simple is called *defective* .

We will often denote by $\lambda_1, \lambda_2, \ldots, \lambda_p$, $(p \leq n)$, all the *distinct* eigenvalues of A. It is a simple exercise to show that the characteristic polynomials of two similar matrices are identical, see Exercise P-1.7. Therefore, the eigenvalues of two similar matrices are equal and so are their algebraic multiplicities. Moreover if v is an eigenvector of B then Xv is an eigenvector of A and, conversely, if y is an eigenvector of A then $X^{-1}y$ is an eigenvector of B. As a result the number of independent eigenvectors associated with a given eigenvalue is the same for two similar matrices, i.e., their geometric multiplicity is also the same.

The possible desired forms are numerous but they all have the common goal of attempting to simplify the original eigenvalue problem. Here are some possibilities with comments as to their usefulness.

- *Diagonal:* the simplest and certainly most desirable choice but it is not always achievable.

- *Jordan:* this is an upper bidiagonal matrix with ones or zeroes on the super diagonal. Always possible but not numerically trustworthy.

- *Upper triangular:* in practice this is the most reasonable compromise as the similarity from the original matrix to a triangular form can be chosen to be isometric and therefore the transformation can be achieved via a sequence of elementary unitary transformations which are numerically stable.

8.1. Reduction to the Diagonal Form.

The simplest form in which a matrix can be reduced is undoubtedly the diagonal form but this reduction is, unfortunately, not always possible. A matrix that can be reduced to the diagonal form is called diagonalizable. The following theorem characterizes such matrices.

Theorem 1.1 *A matrix of dimension n is diagonalizable if and only if it has n linearly independent eigenvectors.*

Proof. A matrix A is diagonalizable if and only if there exists a nonsingular matrix X and a diagonal matrix D such that $A = XDX^{-1}$ or equivalently $AX = XD$, where D is a diagonal matrix. This is equivalent to saying that there exist n linearly independent vectors – the n column-vectors of X – such that $Ax_i = d_i x_i$, i.e., each of these column-vectors is an eigenvector of A. ■

A matrix that is diagonalizable has only semi-simple eigenvalues. Conversely, if all the eigenvalues of a matrix are semi-simple then there exist n eigenvectors of the matrix A. It can be easily shown that these eigenvectors are linearly independent, see Exercise P-1.1. As a result we have the following proposition.

Proposition 1.3 *A matrix is diagonalizable if and only if all its eigenvalues are semi-simple.*

Since every simple eigenvalue is semi-simple, an immediate corollary of the above result is that when A has n distinct eigenvalues then it is diagonalizable.

8.2. The Jordan Canonical Form

From the theoretical viewpoint, one of the most important canonical forms of matrices is the well-known Jordan form. In what follows, the main constructive steps that lead to the Jordan canonical decomposition are outlined. For details, the reader is referred to a standard book on matrix theory or linear algebra.

• For every integer l and each eigenvalue λ_i it is true that

$$\mathrm{Ker}(A - \lambda_i I)^{l+1} \supset \mathrm{Ker}(A - \lambda_i I)^{l} \ .$$

• Because we are in a finite dimensional space the above property implies that there is a first integer l_i such that

$$\mathrm{Ker}(A - \lambda_i I)^{l_i+1} = \mathrm{Ker}(A - \lambda_i I)^{l_i},$$

and in fact $\mathrm{Ker}(A - \lambda_i I)^{l} = \mathrm{Ker}(A - \lambda_i I)^{l_i}$ for all $l \geq l_i$. The integer l_i is called the index of λ_i.

• The subspace $M_i = \mathrm{Ker}(A - \lambda_i I)^{l_i}$ is invariant under A. Moreover, the space $\mathbb{C}^n$ is the direct sum of the subspaces M_i's, for $i = 1, 2, \ldots, p$. Let $m_i = \dim(M_i)$.

• In each invariant subspace M_i there are γ_i independent eigenvectors, i.e., elements of $\mathrm{Ker}(A - \lambda_i I)$, with $\gamma_i \leq m_i$. It turns out that this set of vectors can be completed to form a basis of M_i by adding to it elements of $\mathrm{Ker}(A - \lambda_i I)^2$, then elements of $\mathrm{Ker}(A - \lambda_i I)^3$, and so on. These elements are generated by starting separately from each eigenvector u, i.e., an element of $\mathrm{Ker}(A - \lambda_i I)$, and then seeking an element that satisfies $(A - \lambda_i I)z_1 = u$. Then, more generally we construct z_{i+1} by solving the equation $(A - \lambda_i I)z_{i+1} = z_i$ when possible. The vector z_i belongs to $\mathrm{Ker}(A - \lambda_i I)^{i+1}$ and is called a principal vector (sometimes generalized eigenvector). The process is continued until no more principal vectors are found. There are at most l_i principal vectors for each of the γ_i eigenvectors.

• The final step is to represent the original matrix A with respect to the basis made up of the p bases of the invariant subspaces M_i defined in the previous step.

The matrix representation J of A in the new basis described above has the block diagonal structure,

$$X^{-1}AX = J = \begin{pmatrix} J_1 & & & & & \\ & J_2 & & & & \\ & & \ddots & & & \\ & & & J_i & & \\ & & & & \ddots & \\ & & & & & J_p \end{pmatrix}$$

where each J_i corresponds to the subspace M_i associated with the eigenvalue λ_i. It is of size m_i and it has itself the following structure,

$$J_i = \begin{pmatrix} J_{i1} & & & \\ & J_{i2} & & \\ & & \ddots & \\ & & & J_{i\gamma_i} \end{pmatrix} \text{ with } J_{ik} = \begin{pmatrix} \lambda_i & 1 & & \\ & \ddots & \ddots & \\ & & \lambda_i & 1 \\ & & & \lambda_i \end{pmatrix}.$$

Each of the blocks J_{ik} corresponds to a different eigenvector associated with the eigenvalue λ_i. Its size is equal to the number of principal vectors found for the eigenvector to which the block is associated and does not exceed l_i.

Theorem 1.2 *Any matrix A can be reduced to a block diagonal matrix consisting of p diagonal blocks, each associated with a distinct eigenvalue. Each diagonal block number i has itself a block diagonal structure consisting of γ_i subblocks, where γ_i is the geometric multiplicity of the eigenvalue λ_i. Each of the subblocks, referred to as a Jordan block, is an upper bidiagonal matrix of size not exceeding l_i, with the constant λ_i on the diagonal and the constant one on the super diagonal.*

We refer to the i-th diagonal block, $i = 1, \ldots, p$ as the i-th Jordan submatrix (sometimes "Jordan Box"). The Jordan submatrix number i starts in column $j_i \equiv m_1 + m_2 + \cdots + m_{i-1} + 1$. From the above form it is not difficult to see that $M_i = \text{Ker}(A - \lambda_i I)^{l_i}$ is merely the span of the columns $j_i, j_i + 1, \ldots, j_{i+1} - 1$ of the matrix X. These vectors are all the eigenvectors and the principal vectors associated with the eigenvalue λ_i.

Since A and J are similar matrices their characteristic polynomials are identical. Hence, it is clear that the algebraic multiplicity of an eigenvalue λ_i is equal to the dimension of M_i:

$$\mu_i = m_i \equiv \dim\,(M_i)\ .$$

As a result,

$$\mu_i \geq \gamma_i.$$

Because $\mathbb{C}^n$ is the direct sum of the subspaces $M_i, i = 1, \ldots, p$ each vector x can be written in a unique way as

$$x = x_1 + x_2 + \cdots + x_i + \cdots + x_p,$$

where x_i is a member of the subspace M_i. The linear transformation defined by

$$P_i : x \rightarrow x_i$$

is a projector onto M_i along the direct sum of the subspaces $M_j, j \neq i$. The family of projectors $P_i, i = 1, \ldots, p$ satisfies the following properties,

$$\text{Ran}(P_i) = M_i \tag{1.12}$$

$$P_i P_j = P_j P_i = 0, \text{ if } i \neq j \tag{1.13}$$

$$\sum_{i=1}^{p} P_i = I \tag{1.14}$$

In fact it is easy to see that the above three properties define a decomposition of $\mathbb{C}^n$ into a direct sum of the images of the projectors P_i in a unique way. More precisely, any family of projectors that satisfies the above three properties is uniquely determined and is associated with the decomposition of $\mathbb{C}^n$ into the direct sum of the images of the P_i 's.

It is helpful for the understanding of the Jordan canonical form to determine the matrix representation of the projectors P_i. Consider the matrix $\hat{J}_i$ which is obtained from the Jordan matrix by replacing all the diagonal submatrices by zero blocks except the i^{th} submatrix which is replaced by the identity matrix.

$$\hat{J}_i = \begin{pmatrix} 0 & & & & \\ & 0 & & & \\ & & I & & \\ & & & 0 & \\ & & & & 0 \end{pmatrix}$$

In other words if each i-th Jordan submatrix starts at the column number j_i, then the columns of $\hat{J}_i$ will be zero columns except columns $j_i, \ldots, j_{i+1} - 1$ which are the corresponding columns of the identity matrix. Let $\hat{P}_i = X\hat{J}_i X^{-1}$. Then it is not difficult to verify that $\hat{P}_i$ is a projector and that,

1. The range of $\hat{P}_i$ is the span of columns $j_i, \ldots, j_{i+1} - 1$ of the matrix X. This is the same subspace as M_i.
2. $\hat{P}_i \hat{P}_j = \hat{P}_j \hat{P}_i = 0$ whenever $i \neq j$
3. $\hat{P}_1 + \hat{P}_2 + \cdots + \hat{P}_p = I$

According to our observation concerning the uniqueness of a family of projectors that satisfy (1.12) - (1.14) this implies that

$$\hat{P}_i = P_i \quad , \quad i = 1, \ldots, p$$

Example 1.2 Let us assume that the eigenvalue λ_i is simple. Then,

$$P_i = X e_i e_i^H X^{-1} \equiv u_i w_i^H,$$

in which we have defined $u_i = Xe_i$ and $w_i = X^{-H} e_i$. It is easy to show that u_i and w_i are right and left eigenvectors, respectively, associated with λ_i and normalized so that $w_i^H u_i = 1$.

Consider now the matrix $\hat{D}_i$ obtained from the Jordan form of A by replacing each Jordan submatrix by a zero matrix except the i-th submatrix which is obtained by zeroing its diagonal elements, i.e.,

$$\hat{D}_i = \begin{pmatrix} 0 & & & & & \\ & 0 & & & & \\ & & \ddots & & & \\ & & & J_i - \lambda_i I & & \\ & & & & \ddots & \\ & & & & & 0 \end{pmatrix}$$

Define $D_i = X\hat{D}_i X^{-1}$. Then it is a simple exercise to show by means of the explicit expression for $\hat{P}_i$, that

$$D_i = (A - \lambda_i I)P_i. \tag{1.15}$$

Moreover, $D_i^{l_i} = 0$, i.e., D_i is a *nilpotent matrix* of index l_i. We are now ready to state the following important theorem which can be viewed as an alternative mathematical formulation of Theorem 1.2 on Jordan forms.

Theorem 1.3 *Every square matrix A admits the decomposition*

$$A = \sum_{i=1}^{p} (\lambda_i P_i + D_i) \tag{1.16}$$

where the family of projectors $\{P_i\}_{i=1,\ldots,p}$ satisfies the conditions (1.12), (1.13), and (1.14), and where $D_i = (A - \lambda_i I)P_i$ is a nilpotent operator of index l_i.

Proof. From (1.15), we have

$$AP_i = \lambda_i P_i + D_i \quad i = 1, 2, \ldots, p$$

Summing up the above equalities for $i = 1, 2, \ldots, p$ we get

$$A \sum_{i=1}^{p} P_i = \sum_{i=1}^{p} (\lambda_i P_i + D_i)$$

The proof follows by substituting (1.14) into the left-hand-side. ■

The projector P_i is called the *spectral projector* associated with the eigenvalue λ_i. The linear operator D_i is called the *nilpotent* associated with λ_i. The decomposition (1.16) is referred to as the spectral decomposition of A. Additional properties that are easy to prove from the various expressions of P_i and D_i are the following

$$P_i D_j = D_j P_i = \delta_{ij} P_i \tag{1.17}$$
$$AP_i = P_i A = P_i A P_i = \lambda_i P_i + D_i \tag{1.18}$$
$$A^k P_i = P_i A^k = P_i A^k P_i = P_i(\lambda_i I + D_i)^k = (\lambda_i I + D_i)^k P_i \tag{1.19}$$
$$AP_i = [x_{j_i}, \ldots, x_{j_{i+1}-1}] B_i [y_{j_i}, \ldots, y_{j_{i+1}-1}]^H \tag{1.20}$$

where B_i is the i-th Jordan submatrix and where the columns y_j are the columns of the matrix X^{-H}.

Corollary 1.1 *For any matrix norm $\|.\|$, the following relation holds*

$$\lim_{k\to\infty} \|A^k\|^{1/k} = \rho(A) . \tag{1.21}$$

Proof. The proof of this corollary is the subject of exercise P-1.8. ■

Another way of stating the above corollary is that there is a sequence ϵ_k such that

$$\|A^k\| = (\rho(A) + \epsilon_k)^k$$

where $\lim_{k\to\infty} \epsilon_k = 0$.

8.3. The Schur Canonical Form

We will now show that any matrix is unitarily similar to an upper triangular matrix. The only result needed to prove the following theorem is that any vector of 2-norm one can be completed by $n-1$ additional vectors to form an orthonormal basis of $\mathbb{C}^n$.

Theorem 1.4 *For any given matrix A there exists a unitary matrix Q such that $Q^H AQ = R$ is upper triangular.*

Proof. The proof is by induction over the dimension n. The result is trivial for $n = 1$. Let us assume that it is true for $n-1$ and consider any matrix A of size n. The matrix admits at least one eigenvector u that is associated with an eigenvalue λ. We assume without loss of generality that $\|u\|_2 = 1$. We can complete the vector u into an orthonormal set, i.e., we can find an $n \times (n-1)$ matrix V such that the $n \times n$ matrix $U = [u, V]$ is unitary. Then we have $AU = [\lambda u, AV]$ and hence,

$$U^H AU = \begin{bmatrix} u^H \\ V^H \end{bmatrix} \quad [\lambda u, AV] = \begin{pmatrix} \lambda & u^H AV \\ 0 & V^H AV \end{pmatrix} \tag{1.22}$$

We now use our induction hypothesis for the $(n-1) \times (n-1)$ matrix $B = V^H AV$: there exists an $(n-1) \times (n-1)$ unitary matrix Q_1 such that $Q_1^H BQ_1 = R_1$ is upper triangular. Let us define the $n \times n$ matrix

$$\hat{Q}_1 = \begin{pmatrix} 1 & 0 \\ 0 & Q_1 \end{pmatrix}$$

and multiply both members of (1.22) by $\hat{Q}_1^H$ from the left and $\hat{Q}_1$ from the right. The resulting matrix is clearly upper triangular and this shows that the result is true for A, with $Q = \hat{Q}_1 U$ which is a unitary $n \times n$ matrix. ∎

A simpler proof that uses the Jordan canonical form and the QR decomposition is the subject of Exercise P-1.5. Since the matrix

R is triangular and similar to A, its diagonal elements are equal to the eigenvalues of A ordered in a certain manner. In fact it is easy to extend the proof of the theorem to show that we can obtain this factorization with *any order* we want for the eigenvalues. One might ask the question as to which order might be best numerically but the answer to the question goes beyond the scope of this book. Despite its simplicity, the above theorem has far reaching consequences some of which will be examined in the next section.

It is important to note that for any $k \leq n$ the subspace spanned by the first k columns of Q is invariant under A. This is because from the Schur decomposition we have, for $1 \leq j \leq k$,

$$Aq_j = \sum_{i=1}^{i=j} r_{ij} q_i \ .$$

In fact, letting $Q_k = [q_1, q_2, \ldots, q_k]$ and R_k be the principal leading submatrix of dimension k of R, the above relation can be rewritten as

$$AQ_k = Q_k R_k$$

which we refer to as the partial Schur decomposition of A. The simplest case of this decomposition is when $k = 1$, in which case q_1 is an eigenvector. The vectors q_i are usually referred to as Schur vectors. Note that the Schur vectors are not unique and in fact they depend on the order chosen for the eigenvalues.

A slight variation on the Schur canonical form is the quasi Schur form, also referred to as the real Schur form. Here, diagonal blocks of size 2 x 2 are allowed in the upper triangular matrix R. The reason for this is to avoid complex arithmetic when the original matrix is real. A 2×2 block is associated with each complex conjugate pair of eigenvalues of the matrix.

Example 1.3 Consider the 3×3 matrix

$$A = \begin{pmatrix} 1 & 10 & 0 \\ -1 & 3 & 1 \\ -1 & 0 & 1 \end{pmatrix}$$

The matrix A has the pair of complex conjugate eigenvalues

$$2.4069.. \pm i \times 3.2110..$$

and the real eigenvalue 0.1863... The standard (complex) Schur form is given by the pair of matrices

$$V = \begin{pmatrix} 0.3381 - 0.8462i & 0.3572 - 0.1071i & 0.1749 \\ 0.3193 - 0.0105i & -0.2263 - 0.6786i & -0.6214 \\ 0.1824 + 0.1852i & -0.2659 - 0.5277i & 0.7637 \end{pmatrix}$$

and

$$S = \begin{pmatrix} 2.4069 + 3.2110i & 4.6073 - 4.7030i & -2.3418 - 5.2330i \\ 0 & 2.4069 - 3.2110i & -2.0251 - 1.2016i \\ 0 & 0 & 0.1863 \end{pmatrix}.$$

It is possible to avoid complex arithmetic by using the quasi-Schur form which consists of the pair of matrices

$$U = \begin{pmatrix} -0.9768 & 0.1236 & 0.1749 \\ -0.0121 & 0.7834 & -0.6214 \\ 0.2138 & 0.6091 & 0.7637 \end{pmatrix}$$

and

$$R = \begin{pmatrix} 1.3129 & -7.7033 & 6.0407 \\ 1.4938 & 3.5008 & -1.3870 \\ 0 & 0 & 0.1863 \end{pmatrix}.$$

We would like to conclude this section by pointing out that the Schur and the quasi Schur forms of a given matrix are in no way unique. In addition to the dependence on the ordering of the eigenvalues, any column of Q can be multiplied by a complex sign $e^{i\theta}$ and a new corresponding R can be found. For the quasi Schur form there are infinitely many ways of selecting the 2×2 blocks, corresponding to applying arbitrary rotations to the columns of Q associated with these blocks.

9. Normal and Hermitian Matrices

In this section we look at the specific properties of normal matrices and Hermitian matrices regarding among other things their spectra and some important optimality properties of their eigenvalues. The most common normal matrices that arise in practice are Hermitian or skew-Hermitian. In fact, symmetric real matrices form a large part of the matrices that arise in practical eigenvalue problems.

9.1. Normal Matrices

By definition a matrix is said to be normal if it satisfies the relation

$$A^H A = AA^H. \tag{1.23}$$

An immediate property of normal matrices is stated in the following proposition.

Proposition 1.4 *If a normal matrix is triangular then it is necessarily a diagonal matrix.*

Proof. Assume for example that A is upper triangular and normal and let us compare the first diagonal element of the left hand side matrix of (1.23) with the corresponding element of the matrix on the right hand side. We obtain that

$$|a_{11}|^2 = \sum_{j=1}^{n} |a_{1j}|^2,$$

which shows that the elements of the first row are zeros except for the diagonal one. The same argument can now be used for the second row, the third row, and so on to the last row, to show that $a_{ij} = 0$ for $i \neq j$. ∎

As a consequence of this we have the following important result.

Theorem 1.5 *A matrix is normal if and only if it is unitarily similar to a diagonal matrix.*

Proof. It is straightforward to verify that a matrix which is unitarily similar to a diagonal matrix is normal. Let us now show that any normal matrix A is unitarily similar to a diagonal matrix. Let $A = QRQ^H$ be the Schur canonical form of A where we recall that Q is unitary and R is upper triangular. By the normality of A we have

$$QR^H Q^H QRQ^H = QRQ^H QR^H Q^H$$

or,

$$QR^H RQ^H = QRR^H Q^H$$

Upon multiplication by Q^H on the left and Q on the right this leads to the equality $R^H R = RR^H$ which means that R is normal, and according to the previous proposition this is only possible if R is diagonal. ■

Thus, any normal matrix is diagonalizable and admits an orthonormal basis of eigenvectors, namely the column vectors of Q.

Clearly, Hermitian matrices are just a particular case of normal matrices. Since a normal matrix satisfies the relation $A = QDQ^H$, with D diagonal and Q unitary, the eigenvalues of A are the diagonal entries of D. Therefore, if these entries are real it is clear that we will have $A^H = A$. This is restated in the following corollary.

Corollary 1.2 *A normal matrix whose eigenvalues are real is Hermitian.*

As will be seen shortly the converse is also true, in that a Hermitian matrix has real eigenvalues.

An eigenvalue λ of any matrix satisfies the relation

$$\lambda = \frac{(Au, u)}{(u, u)}$$

where u is an associated eigenvector. More generally one might consider the complex scalars,

$$\mu(x) = \frac{(Ax, x)}{(x, x)} \tag{1.24}$$

defined for any nonzero vector in $\mathbb{C}^n$. These ratios are referred to as *Rayleigh quotients* and are important both from theoretical and practical purposes. The set of all possible Rayleigh quotients as x runs over $\mathbb{C}^n$ is called the *field of values* of A. This set is clearly bounded since each $|\mu(x)|$ is bounded by the the 2-norm of A, i.e., $|\mu(x)| \leq \|A\|_2$ for all x.

If a matrix is normal then any vector x in $\mathbb{C}^n$ can be expressed as

$$\sum_{i=1}^{n} \xi_i q_i$$

where the vectors q_i form an orthogonal basis of eigenvectors, and the expression for $\mu(x)$ becomes,

$$\mu(x) = \frac{(Ax, x)}{(x, x)} = \frac{\sum_{k=1}^{n} \lambda_k |\xi_k|^2}{\sum_{k=1}^{n} |\xi_k|^2} \equiv \sum_{k=1}^{n} \beta_k \lambda_k \tag{1.25}$$

where

$$0 \leq \beta_i = \frac{|\xi_i|^2}{\sum_{k=1}^{n} |\xi_k|^2} \leq 1 \ , \quad \text{and} \quad \sum_{i=1}^{n} \beta_i = 1$$

From a well-known characterization of convex hulls due to Hausdorff, (Hausdorff's convex hull theorem) this means that the set of all possible Rayleigh quotients as x runs over all of $\mathbb{C}^n$ is equal to the convex hull of the λ_i's. This leads to the following theorem.

Theorem 1.6 *The field of values of a normal matrix is equal to the convex hull of its spectrum.*

The question that arises next is whether or not this is also true for non-normal matrices and the answer is no, i.e., the convex hull of the eigenvalues and the field of values of a non-normal matrix are different in general, see Exercise P-1.10 for an example. As a generic example, one can take any nonsymmetric real matrix that has real eigenvalues only; its field of values will contain imaginary values. It has been shown (Hausdorff) that the field of values of a matrix is a convex set. Since the eigenvalues are members of the field of values, their convex hull is contained in the field of values. This is summarized in the following proposition.

Proposition 1.5 *The field of values of an arbitrary matrix is a convex set which contains the convex hull of its spectrum. It is equal to the convex hull of the spectrum when the matrix in normal.*

9.2. Hermitian Matrices

A first and important result on Hermitian matrices is the following.

Theorem 1.7 *The eigenvalues of a Hermitian matrix are real, i.e.,* $\sigma(A) \subset \mathbb{R}$.

Proof. Let λ be an eigenvalue of A and u an associated eigenvector or 2-norm unity. Then

$$\lambda = (Au, u) = (u, Au) = \overline{(Au, u)} = \overline{\lambda}$$

∎

Moreover, it is not difficult to see that if, in addition, the matrix is real then the eigenvectors can be chosen to be real, see Exercise P-1.16. Since a Hermitian matrix is normal an immediate consequence of Theorem 1.5 is the following result.

Theorem 1.8 *Any Hermitian matrix is unitarily similar to a real diagonal matrix.*

In particular a Hermitian matrix admits a set of orthonormal eigenvectors that form a basis of $\mathbb{C}^n$.

In the proof of Theorem 1.6 we used the fact that the inner products (Au, u) are real. More generally it is clear that any Hermitian matrix is such that (Ax, x) is real for any vector $x \in \mathbb{C}^n$. It turns out that the converse is also true, i.e., it can be shown that if (Az, z) is real for all vectors z in $\mathbb{C}^n$ then the matrix A is Hermitian, see Problem P-1.14.

Eigenvalues of Hermitian matrices can be characterized by optimality properties of the Rayleigh quotients (1.24). The best known of these is the Min-Max principle. Let us order all the eigenvalues of A in descending order:

$$\lambda_1 \ \geq \ \lambda_2 \ \ldots \ \geq \ \lambda_n .$$

Here the eigenvalues are not necessarily distinct and they are repeated, each according to its multiplicity. In what follows, we denote by S a generic subspace of $\mathbb{C}^n$. Then we have the following theorem.

Theorem 1.9 (Min-Max theorem) *The eigenvalues of a Hermitian matrix A are characterized by the relation*

$$\lambda_k = \min_{S, \ \dim\ (S)=n-k+1} \quad \max_{x \in S, x \neq 0} \frac{(Ax, x)}{(x, x)} \tag{1.26}$$

Proof. Let $\{q_i\}_{i=1,\ldots,n}$ be an orthonormal basis of $\mathbb{C}^n$ consisting of eigenvectors of A associated with $\lambda_1, \ldots, \lambda_n$ respectively. Let S_k be the subspace spanned by the first k of these vectors and denote by $\mu(S)$ the maximum of $(Ax, x)/(x, x)$ over all nonzero vectors of a subspace S. Since the dimension of S_k is k, a well-known theorem of linear algebra shows that its intersection with any subspace S of dimension $n - k + 1$ is not reduced to $\{0\}$, i.e., there is vector x in $S \bigcap S_k$. For this $x = \sum_{i=1}^{k} \xi_i q_i$ we have

$$\frac{(Ax, x)}{(x, x)} \ = \ \frac{\sum_{i=1}^{k} \lambda_i |\xi_i|^2}{\sum_{i=1}^{k} |\xi_i|^2} \quad \geq \quad \lambda_k$$

so that $\mu(S) \; \geq \; \lambda_k$.

Consider on the other hand the particular subspace S_0 of dimension $n-k+1$ which is spanned by $q_k, \ldots, q_n$. For each vector x in this subspace we have

$$\frac{(Ax,x)}{(x,x)} = \frac{\sum_{i=k}^{n} \lambda_i |\xi_i|^2}{\sum_{i=k}^{n} |\xi_i|^2} \quad \leq \quad \lambda_k$$

so that $\mu(S_0) \leq \lambda_k$. In other words, as S runs over all $n-k+1$-dimensional subspaces $\mu(S)$ is always $\geq \lambda_k$ and there is at least one subspace S_0 for which $\mu(S_0) \leq \lambda_k$ which shows the result. ■

This result is attributed to Courant and Fisher, and to Poincaré and Weyl. It is often referred to as Courant-Fisher min-max principle or theorem. As a particular case, the largest eigenvalue of A satisfies

$$\lambda_1 = \max_{x \neq 0} \frac{(Ax,x)}{(x,x)} \; . \tag{1.27}$$

Actually, there are four different ways of rewriting the above characterization. The second formulation is

$$\lambda_k = \max_{S, \; \dim \; (S)=k} \; \min_{x \in S, x \neq 0} \; , \; \frac{(Ax,x)}{(x,x)} \tag{1.28}$$

and the two other ones can be obtained from the above two formulations by simply relabeling the eigenvalues increasingly instead of decreasingly. Thus, with our labeling of the eigenvalues in descending order, (1.28) tells us that the smallest eigenvalue satisfies,

$$\lambda_n = \min_{x \neq 0} \frac{(Ax,x)}{(x,x)} \; .$$

with λ_n replaced by λ_1 if the eigenvalues are relabeled increasingly.

In order for all the eigenvalues of a Hermitian matrix to be positive it is necessary and sufficient that

$$(Ax,x) > 0, \quad \forall \; x \in \mathbb{C}^n, \quad x \neq 0.$$

Such a matrix is called *positive definite*. A matrix that satisfies $(Ax, x) \geq 0$ for any x is said to be *positive semi-definite*. In particular the matrix $A^H A$ is semi-positive definite for any rectangular matrix, since

$$(A^H Ax, x) = (Ax, Ax) \ \geq \ 0 \quad \forall \ x.$$

Similarly, AA^H is also a Hermitian semi-positive definite matrix. The square roots of the eigenvalues of $A^H A$ for a general rectangular matrix A are called the *singular values* of A and are denoted by σ_i. In Section 1.5 we have stated without proof that the 2-norm of any matrix A is equal to the largest singular value σ_1 of A. This is now an obvious fact, because

$$||A||_2^2 = \max_{x \neq 0} \frac{||Ax||_2^2}{||x||_2^2} \ = \ \max_{x \neq 0} \frac{(Ax, Ax)}{(x, x)} \ = \ \max_{x \neq 0} \frac{(A^H Ax, x)}{(x, x)} \ = \ \sigma_1^2$$

which results from (1.27).

Another characterization of eigenvalues, known as the Courant characterization, is stated in the next theorem. In contrast with the min-max theorem this property is recursive in nature.

Theorem 1.10 *The eigenvalue λ_i and the corresponding eigenvector q_i of a Hermitian matrix are such that*

$$\lambda_1 = \frac{(Aq_1, q_1)}{(q_1, q_1)} \ = \ \max_{x \in \mathbb{C}^n, x \neq 0} \frac{(Ax, x)}{(x, x)}$$

and for $k > 1$:

$$\lambda_k = \frac{(Aq_k, q_k)}{(q_k, q_k)} \ = \ \max_{x \neq 0, q_1^H x = \ldots = q_{k-1}^H x = 0} \frac{(Ax, x)}{(x, x)} \ . \tag{1.29}$$

In other words, the maximum of the Rayleigh quotient over a subspace that is orthogonal to the first $k-1$ eigenvectors is equal to λ_k and is achieved for the eigenvector q_k associated with λ_k. The proof follows easily from the expansion (1.25) of the Rayleigh quotient.

10. Nonnegative Matrices

A nonnegative matrix is a matrix whose entries are nonnegative,

$$a_{ij} \geq 0 \ .$$

Nonnegative matrices arise in many applications and play a crucial role in the theory of matrices. They play for example a key role in the analysis of convergence of iterative methods for partial differential equations. They also arise in economics, queuing theory, chemical engineering, etc..

A matrix is said to be reducible if, there is a permutation matrix P such that PAP^T is block upper triangular. An important result concerning nonnegative matrices is the following theorem known as the Perron-Frobenius theorem.

Theorem 1.11 *Let A be a real $n \times n$ nonnegative irreducible matrix. Then $\lambda \equiv \rho(A)$, the spectral radius of A, is a simple eigenvalue of A. Moreover, there exists an eigenvector u with positive elements associated with this eigenvalue.*

PROBLEMS

P-1.1 Show that two eigenvectors associated with two distinct eigenvalues are linearly independent. More generally show that a family of eigenvectors associated with distinct eigenvalues forms a linearly independent family.

P-1.2 Show that if λ is any eigenvalue of the matrix AB then it is also an eigenvalue of the matrix BA. Start with the particular case where A and B are square and B is nonsingular then consider the more general case where A, B may be singular or even rectangular (but such that AB and BA are square).

P-1.3 Show that the Frobenius norm is consistent. Can this norm be associated to two vector norms via (1.4)? What is the Frobenius norm of a diagonal matrix? What is the p-norm of a diagonal matrix (for any p)?

P-1.4 Find the Jordan canonical form of the matrix:

$$A = \begin{pmatrix} 1 & 2 & -4 \\ 0 & 1 & 2 \\ 0 & 0 & 2 \end{pmatrix}.$$

Same question for the matrix obtained by replacing the element a_{33} by 1.

P-1.5 Give an alternative proof of Theorem 1.4 on the Schur form by starting from the Jordan canonical form. [Hint: write $A = XJX^{-1}$ and use the QR decomposition of X.]

P-1.6 Show from the definition of determinants used in Section (1.2) that the characteristic polynomial is a polynomial of degree n for an $n \times n$ matrix.

P-1.7 Show that the characteristic polynomials of two similar matrices are equal.

P-1.8 Show that

$$\lim_{k \to \infty} \|A^k\|^{1/k} = \rho(A),$$

for any matrix norm. [Hint: use the Jordan canonical form or Theorem 1.3]

P-1.9 Let X be a nonsingular matrix and, for any matrix norm $\|.\|$, define $\|A\|_X = \|AX\|$. Show that this is indeed a matrix norm. Is this matrix norm consistent? Similar questions for $\|XA\|$ and $\|YAX\|$ where Y is also a nonsingular matrix. These norms are not, in general, associated with any vector norms, i.e., they can't be defined by a formula of the form (1.4). Why? What about the particular case $\|A\|' = \|XAX^{-1}\|$?

P-1.10 Find the field of values of the matrix

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

and verify that it is not equal to the convex hull of its eigenvalues.

P-1.11 Show that any matrix can be written as the sum of a Hermitian and a skew-Hermitian matrix (or the sum of a symmetric and a skew-symmetric matrix).

P-1.12 Show that for a skew-Hermitian matrix S, we have

$$\Re e(Sx, x) = 0 \quad \text{for any } x \in \mathbb{C}^n.$$

P-1.13 Given an arbitrary matrix S, show that if $(Sx, x) = 0$ for all x in $\mathbb{C}^n$ then we must have

$$(Sy, z) + (Sz, y) = 0 \quad \forall\, y\ ,\ z\ \in\ \mathbb{C}^n.$$

[Hint: expand $(S(y + z), y + z)$].

P-1.14 Using the result of the previous two problems, show that if (Ax, x) is real for all x in $\mathbb{C}^n$, then A must be Hermitian. Would this result be true if we were to replace the assumption by: (Ax, x) *is real for all real* x? Explain.

P-1.15 The definition of a positive definite matrix is that (Ax, x) be real and positive for all real vectors x. Show that this is equivalent to requiring that the Hermitian part of A, namely $\frac{1}{2}(A + A^H)$, be (Hermitian) positive definite.

P-1.16 Let A be a real symmetric matrix and λ an eigenvalue of A. Show that if u is an eigenvector associated with λ then so is $\bar{u}$. As a result, prove that for any eigenvalue of a real symmetric matrix, there is an associated eigenvector which is real.

P-1.17 Show that a Hessenberg matrix H such that $h_{j+1,j} \neq 0, j = 1, 2, \ldots, n-1$ cannot be derogatory.

NOTES AND REFERENCES. For additional reading on the material presented in this Chapter, see Golub and Van Loan [63] and Stewart [167]. More details on matrix eigenvalue problems can be found in Gantmacher's book [54] and Wilkinson [183]. Stewart and Sun's recent book [172] devotes a separate chapter to matrix norms and contains a wealth of information. Some of the terminology we used is borrowed from Chatelin [14, 15] and Kato [85]. For a good overview of the linear algebra aspects of matrix theory and a complete proof of Jordan's canonical form we recommend Halmos' book [69]. ♠