

Quick Review of Matrix and Real Linear Algebra

KC Border

Subject to constant revision

Last major revision: December 12, 2016

v. 2016.12.30:16.10

Contents

1	Scalar fields	1
2	Vector spaces	1
2.1	The vector space \mathbf{R}^n	2
2.2	Other examples of vector spaces	3
3	Some elementary consequences	4
3.1	Linear combinations and subspaces	4
3.2	Linear independence	5
3.3	Bases and dimension	5
3.4	Linear transformations	6
3.5	The coordinate mapping	7
4	Inner product	7
4.1	The Parallelogram Law	9
4.2	Orthogonality	10
4.3	Orthogonal projection and orthogonal complements	11
4.4	Orthogonality and alternatives	13
4.5	The geometry of the Euclidean inner product	14
5	The dual of an inner product space	14
6*	Linear functionals on infinite dimensional spaces	17
7	Linear transformations	18
7.1	Inverse of a linear transformation	19
7.2	Adjoint of a transformation	20
7.3	Orthogonal transformations	22
7.4	Symmetric transformations	23
8	Eigenvalues and eigenvectors	24
9	Matrices	27
9.1	Matrix operations	27
9.2	Systems of linear equations	28
9.3	Matrix representation of a linear transformation	30
9.4	Gershgorin's Theorem	32
9.5	Matrix representation of a composition	34

9.6 Change of basis 34

9.7 The Principal Axis Theorem 36

9.8 Simultaneous diagonalization 36

9.9 Trace 38

9.10 Matrices and orthogonal projection 40

10 Quadratic forms 41

10.1 Diagonalization of quadratic forms 42

11 Determinants 43

11.1 Determinants as multilinear forms 44

11.2 Some simple consequences 47

11.3 Minors and cofactors 48

11.4 Characteristic polynomials 52

11.5 The determinant as an “oriented volume” 53

11.6 Computing inverses and determinants by Gauss’ method 54

Foreword

These notes have accreted piecemeal from courses in econometrics and microeconomics I have taught over the last thirty-something years, so the notation, hyphenation, and terminology may vary from section to section. They were originally compiled to be a centralized personal reference for finite-dimensional real vector spaces, although I occasionally mention complex or infinite-dimensional spaces. If you too find them useful, so much the better. There is a sketchy index, which I think is better than none.

For a thorough course on linear algebra I now recommend Axler [7]. My favorite reference for infinite-dimensional vector spaces is the *Hitchhiker’s Guide* [2], but it needn’t be your favorite.

1 Scalar fields

A **field** is a collection of mathematical entities that we shall call **scalars**. There are two binary operations defined on scalars, **addition** and **multiplication**. We denote the **sum** of α and β by $\alpha + \beta$, and the product by $\alpha \cdot \beta$ or more simply $\alpha\beta$. These operations satisfy the following familiar properties:

F.1 (Commutativity of Addition and Multiplication) $\alpha + \beta = \beta + \alpha$, $\alpha\beta = \beta\alpha$.

F.2 (Associativity of Addition and Multiplication) $(\alpha + \beta) + \gamma = \alpha + (\beta + \gamma)$, $(\alpha\beta)\gamma = \alpha(\beta\gamma)$.

F.3 (Distributive Law) $\alpha(\beta + \gamma) = (\alpha\beta) + (\alpha\gamma)$.

F.4 (Existence of Additive and Multiplicative Identities) There are two distinct scalars, denoted 0 and 1, such that for every scalar α , we have $\alpha + 0 = \alpha$ and $1\alpha = \alpha$.

F.5 (Existence of Additive Inverse) For each scalar α there exists a scalar β such that $\alpha + \beta = 0$. We often write $-\alpha$ for such a β .

F.5 (Existence of Multiplicative Inverse) For each *nonzero* scalar α there exists a scalar β such that $\alpha\beta = 1$. We often write α^{-1} for such a β .

There are many elementary theorems that point out many things that you probably take for granted. For instance,

- $-\alpha$ and α^{-1} are unique. This justifies the notation by the way.
- $-\alpha = (-1) \cdot \alpha$. (Here -1 is the scalar β that satisfies $1 + \beta = 0$.)
- $\alpha + \gamma = \beta + \gamma \iff \alpha = \beta$.
- $0 \cdot \alpha = 0$.
- $\alpha\beta = 0 \implies [\alpha = 0 \text{ or } \beta = 0]$.

There are many more. See, for instance, Apostol [5, p. 18].

The most important scalar fields are the **real numbers** \mathbf{R} , the **rational numbers** \mathbb{Q} , and the field of **complex numbers** \mathbb{C} . Computer scientists, e.g., Klein [16], are fond of the $\{0, 1\}$ field sometimes known as GF(2) or the **Boolean field**.

These notes are mostly concerned with the field of real numbers. This is not because the other fields are unimportant—it's because I myself have limited use for the others, which is probably a personal shortcoming.

2 Vector spaces

Let K be a field of **scalars**—usually either the real numbers \mathbf{R} or the complex numbers \mathbb{C} , or occasionally the rationals \mathbb{Q} .

1 Definition A **vector space** over K is a nonempty set V of **vectors** equipped with two operations, vector addition $(x, y) \mapsto x + y$, and scalar multiplication $(\alpha, x) \mapsto \alpha x$, where $x, y \in V$ and $\alpha \in K$. The operations satisfy:

V.1 (Commutativity of Vector Addition) $x + y = y + x$

V.2 (Associativity of Vector Addition) $(x + y) + z = x + (y + z)$

V.3 (Existence of Zero Vector) There is a vector 0_V , often denoted simply 0 , satisfying $x + 0 = x$ for every vector x .

V.4 (Additive Inverse) $x + (-1)x = 0$

V.5 (Associativity of Scalar Multiplication) $\alpha(\beta x) = (\alpha\beta)x$

V.6 (Scalar Multiplication by 1) $1x = x$

V.7 (Distributive Law) $\alpha(x + y) = (\alpha x) + (\alpha y)$

V.8 (Distributive Law) $(\alpha + \beta)x = (\alpha x) + (\beta x)$

The term **real vector space** refers to a vector space over the field of real numbers, and a **complex vector space** is a vector space over the field of complex numbers. The term **linear space** is a synonym for vector space.

I shall try to adhere to the following notational conventions that are used by Gale [13]. Vectors are typically denoted by lower case Latin letters, and scalars are typically denoted by lower case Greek letters. When we have occasion to refer to the components of a vector in \mathbf{R}^m , I will try to use the Greek letter corresponding to the Latin letter of the vector. For instance, $x = (\xi_1, \dots, \xi_m)$, $y = (\eta_1, \dots, \eta_m)$, $z = (\zeta_1, \dots, \zeta_m)$. Upper case Latin letters tend to denote matrices, and their entries will be denoted by the corresponding lower case Greek letters.^a

That said, these notes come from many different courses, and I may not have standardized everything. I'm sorry.

^aThe correspondence between Greek and Latin letters is somewhat arbitrary. For instance, one could make the case that Latin y corresponds to Greek v , and not to Greek η . I should write out a guide.

2.1 The vector space \mathbf{R}^n

By far the most important example of a vector space, and indeed the mother of all vector spaces, is the space \mathbf{R}^n of ordered lists of n real numbers. Given ordered lists $x = (\xi_1, \dots, \xi_n)$ and $y = (\eta_1, \dots, \eta_n)$ the vector sum $x + y$ is the ordered list $(\xi_1 + \eta_1, \dots, \xi_n + \eta_n)$. The scalar product is given by the ordered list $\alpha x = (\alpha\xi_1, \dots, \alpha\xi_n)$. The zero of this vector space is the ordered list $0 = (0, \dots, 0)$. It is a trivial matter to verify that the axioms above are satisfied by these operations. As a special case, the set of real numbers \mathbf{R} is a real vector space.

In \mathbf{R}^n we identify several special vectors, the **unit coordinate vectors**. These are the ordered lists e_i that consist of zeroes except for a single 1 in the i^{th} position. We use the notation e_i to denote this vector regardless of the length n of the list.

2.2 Other examples of vector spaces

The following are also vector spaces. The definition of the vector operations is usually obvious.

- $\{0\}$ is a vector space, called the **trivial** vector space. A nontrivial vector space contains at least one nonzero vector.
- The field K is a vector space over itself.
- The set $L_1(I)$ of integrable real-valued functions on an interval I of the real line,

$$\{f: I \rightarrow \mathbf{R} : \int_I |f(t)| dt < \infty\}$$

is a real vector space under the **pointwise operations**: $(f + g)(t) = f(t) + g(t)$ and $(\alpha f)(t) = \alpha f(t)$ for all $t \in I$.¹

- The set $L_2(P)$ of square-integrable random variables,

$$E X^2 < \infty$$

on the probability space (S, \mathcal{E}, P) is a real vector space.

- The set of solutions to a homogeneous linear differential equation, e.g., $f'' + \alpha f' + \beta f = 0$, is a real vector space.
- The sequence spaces ℓ_p ,

$$\ell_p = \left\{ x = (\xi_1, \xi_2, \dots) \in \mathbf{R}^{\mathbf{N}} : \sum_{n=1}^{\infty} |\xi_n|^p < \infty \right\}, \quad 0 < p < \infty$$

$$\ell_{\infty} = \left\{ x = (\xi_1, \xi_2, \dots) : \sup_n |\xi_n| < \infty \right\}$$

are all real vector spaces.

- The set $M(m, n)$ of $m \times n$ real matrices is a real vector space.
- The set of linear transformations from one vector space into another is a linear space. See Proposition 12 below.
- We can even consider the set \mathbf{R} of real numbers as an (infinite-dimensional, see below) vector space over the field \mathbb{Q} of rational numbers. This leads to some very interesting (counter-)examples.

¹See how quickly I forgot the convention that scalars are denoted by Greek letters. I used t here instead of τ .

3 Some elementary consequences

Here are some simple consequences of the axioms that we shall use repeatedly without further comment.

1. 0 is unique.

$$0_2 \underbrace{=}_{\text{by V.3}} 0_2 + 0_1 \underbrace{=}_{\text{by V.1}} 0_1 + 0_2 \underbrace{=}_{\text{by V.3}} 0_1$$

2. $-x = (-1)x$ is the unique vector z satisfying $x + z = 0$. Suppose

$$\begin{aligned} x + y &= 0 \\ \underbrace{-x + x} + y &= -x \\ 0 + y &= -x \\ y &= -x. \end{aligned}$$

3. $0_K x = 0_V$:

$$0_V = 0_K x + (-1)(0_K x) = (0_K - 0_K)x = 0_K x.$$

4. $\underbrace{x + \cdots + x}_n = nx$:

$$\begin{aligned} x + x &= 1x + 1x = (1 + 1)x = 2x \\ x + x + x &= (x + x) + x = 2x + x = 2x + 1x = (2 + 1)x = 3x \end{aligned}$$

etc.

3.1 Linear combinations and subspaces

A **linear combination** of the vectors x_1, \dots, x_m is any sum of scalar multiples of vectors of the form $\alpha_1 x_1 + \cdots + \alpha_m x_m$, $\alpha_i \in K$, $x_i \in V$. A **linear subspace** M of V is a nonempty subset of V that is closed under linear combinations. A linear subspace of a vector space is a vector space in its own right. A linear subspace may also be called a **vector subspace**.

Let $E \subset V$. The **span** of E , denoted $\text{span } E$, is the set of all linear combinations from E . That is,

$$\text{span } E = \left\{ \sum_{i=1}^m \alpha_i x_i : \alpha_i \in K, x_i \in E, m \in \mathbb{N} \right\}.$$

2 Exercise Prove the following.

1. $\{0\}$ is a linear subspace.
2. If M is a linear subspace, then $0 \in M$.
3. The intersection of a family of linear subspaces is a linear subspace.
4. The set $\text{span } E$ is the smallest (with respect to inclusion) linear subspace that includes E .

□

3.2 Linear independence

3 Definition A set E of vectors is **linearly dependent** if there are distinct vectors x_1, \dots, x_m belonging to E , and nonzero scalars $\alpha_1, \dots, \alpha_m$, such that $\alpha_1 x_1 + \dots + \alpha_m x_m = 0$. A set of vectors is **linearly independent** if it is not dependent. That is, E is independent if for every set x_1, \dots, x_m of distinct vectors in E , $\sum_{i=1}^m \alpha_i x_i = 0$ implies $\alpha_1 = \dots = \alpha_m = 0$. We also say that the vectors x_1, \dots, x_m are independent instead of saying that the set $\{x_1, \dots, x_m\}$ is independent.

4 Proposition (Uniqueness of linear combinations) If E is a linearly independent set of vectors and z belongs to $\text{span } E$, then z is a unique linear combination of elements of E .

Proof: If z is zero, the conclusion follows by definition of independence. If z is nonzero, suppose

$$z = \sum_{i=1}^m \alpha_i x_i = \sum_{j=1}^n \beta_j y_j,$$

where the x_i s are distinct elements of E , the y_j s are distinct elements of E (but may overlap with the x_i s), and $\alpha_i, \beta_j > 0$ for $i = 1, \dots, m$ and $j = 1, \dots, n$. Enumerate $A = \{x_i : i = 1, \dots, m\} \cup \{y_j : j = 1, \dots, n\}$ as $A = \{z_k : k = 1, \dots, p\}$. (If $x_i = y_j$ for some i and j , then p is strictly less than $m + n$.) Then we can rewrite $z = \sum_{k=1}^p \hat{\alpha}_k z_k = \sum_{k=1}^p \hat{\beta}_k z_k$, where

$$\hat{\alpha}_k = \begin{cases} \alpha_i & \text{if } z_k = x_i \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \hat{\beta}_k = \begin{cases} \beta_j & \text{if } z_k = y_j \\ 0 & \text{otherwise.} \end{cases}$$

Then

$$0 = z - z = \sum_{k=1}^p (\hat{\alpha}_k - \hat{\beta}_k) z_k \implies \hat{\alpha}_k - \hat{\beta}_k = 0, \quad k = 1, \dots, p$$

since E is independent. Therefore $\hat{\alpha}_k = \hat{\beta}_k$, $k = 1, \dots, p$, which in turn implies $m = n = p$ and $\{x_i : i = 1, \dots, m\} = \{y_j : j = 1, \dots, n\}$, and the proposition is proved. ■

The coefficients of this linear combination are called the **coordinates of x with respect to the set E** .

5 Exercise Prove the following.

1. The empty set is independent.
2. If $0 \in E$, then E is dependent.

□

3.3 Bases and dimension

6 Definition A **Hamel basis**, or more succinctly, a **basis** for the linear space V is a linearly independent set B such that $\text{span } B = V$. The plural of basis is bases.

The next result is immediate from Proposition 4.

7 Proposition Every element of the vector space V is a unique linear combination of basis vectors.

8 Example (The standard basis for \mathbf{R}^m) The set of unit coordinate vectors e_1, \dots, e_m in \mathbf{R}^m is a basis for \mathbf{R}^m , called the **standard basis**.

Observe that the vector $x = (\xi_1, \dots, \xi_m)$ can be written uniquely as $\sum_{j=1}^m \xi_j e_j$. \square

The fundamental facts about bases are these. For a proof of the first assertion, see the *Hitchhiker's Guide* [2, Theorem 1.8, p. 15].

9 Fact Every nontrivial vector space has a basis. Any two bases have the same cardinality, called the **dimension** of V .

Mostly these notes deal with finite-dimensional vector spaces. The next result summarizes Theorems 1.5, 1.6, and 1.7 in Apostol [6, pp. 10–12]

10 Theorem In an n -dimensional space, every set of more than n vectors is dependent. Consequently, any independent set of n vectors is a basis.

Prove the second half for finite-dimensional case.

3.4 Linear transformations

11 Definition Let V, W be vector spaces. A function $T: V \rightarrow W$ is a **linear transformation** or **linear operator** or **homomorphism** if

$$T(\alpha x + \beta y) = \alpha T(x) + \beta T(y).$$

The set of linear transformations from the vector space V into the vector space W is denoted

$$L(V, W).$$

If T is a linear transformation from a vector space into the reals \mathbf{R} , then it is customary to call T a **linear functional**.

The set $L(V, \mathbf{R})$ of linear functionals on V is called the **dual space** of V , and is denoted V' . When there is a notion of continuity, we shall let V^* denote the vector space of continuous linear functionals on V .

The set V' is often called the **algebraic dual** of V and V^* is the **topological dual** of V .

Note that if T is linear, then $T0 = 0$. We shall use this fact without any special mention. It is traditional to write a linear transformation without parentheses, that is, to write Tx rather than $T(x)$.

12 Proposition $L(V, W)$ is itself a vector space under the usual pointwise addition and scalar multiplication of functions.

3.5 The coordinate mapping

In an n -dimensional space V , if we fix a basis in some particular order, we have an **ordered basis** or **frame**. Given this ordered basis the coordinates of a vector comprise an ordered list, and so correspond to an element in \mathbf{R}^n . This mapping from vectors to their coordinates is called the **coordinate mapping** for the ordered basis. The coordinate mapping preserves all the vector operations. The mathematicians' term for this is isomorphism. If two vector spaces are isomorphic, then for many purposes we can think of them as being the same space, the only difference being that the names of the points have been changed.

13 Definition An **isomorphism** between vector spaces is a bijective linear transformation. That is, a function $\varphi: V \rightarrow W$ between vector spaces is an isomorphism if φ is one-to-one and onto, and

$$\varphi(x + y) = \varphi(x) + \varphi(y) \quad \text{and} \quad \varphi(\alpha x) = \alpha\varphi(x).$$

In this case we say that V and W are **isomorphic**.

We have just argued the following.

14 Proposition Given an ordered basis x_1, \dots, x_n for an n -dimensional vector space V , the coordinate mapping is an isomorphism from V to \mathbf{R}^n .

If any n -dimensional vector space is isomorphic to \mathbf{R}^n , is there any reason to consider n -dimensional vector spaces other than \mathbf{R}^n ? The answer, I believe, is yes. Sometimes there is no "natural" ordered basis for a vector space, so there is no "natural" way to treat it as \mathbf{R}^n . For instance, let V be the set of vectors $x = (\xi_1, \xi_2, \xi_3)$ in \mathbf{R}^3 , such that $\xi_1 + \xi_2 + \xi_3 = 0$. This is a 2-dimensional vector space, and there are infinitely many isomorphisms onto \mathbf{R}^2 , but I assert that there is no unique obvious, natural way to identify points in V with points in \mathbf{R}^2 .

4 Inner product

15 Definition A real linear space V has an **inner product** if for each pair of vectors x and y there is a real number, traditionally denoted (x, y) , satisfying the following properties.

IP.1 $(x, y) = (y, x)$.

IP.2 $(x, y + z) = (x, y) + (x, z)$.

IP.3 $\alpha(x, y) = (\alpha x, y) = (x, \alpha y)$.

IP.4 $(x, x) > 0$ if $x \neq 0$.

A vector space V equipped with an inner product is called an **inner product space**.

For a complex vector space, the inner product is complex-valued, and property (1) is replaced by $(x, y) = \overline{(y, x)}$, where the bar denotes complex conjugation, and (3) is replaced by $\alpha(x, y) = (\alpha x, y) = (x, \overline{\alpha}y)$.

16 Proposition In a real inner product space, the inner product is **bilinear**. That is,

$$(\alpha x + \beta y, z) = \alpha(x, z) + \beta(y, z) \quad \text{and} \quad (z, \alpha x + \beta y) = \alpha(z, x) + \beta(z, y)$$

Another simple consequence for real inner product spaces is the following identity that we shall use often.

$$(x + y, x + y) = (x, x) + 2(x, y) + (y, y)$$

For complex inner product spaces, this becomes

$$(x + y, x + y) = (x, x) + (x, y) + (y, x) + (y, y) = (x, x) + (x, y) + \overline{(x, y)} + (y, y)$$

17 Example The **dot product** of two vectors $x = (\xi_1, \dots, \xi_n)$ and $y = (\eta_1, \dots, \eta_n)$ in \mathbf{R}^n is defined by

$$x \cdot y = \sum_{i=1}^n \xi_i \eta_i.$$

The dot product is an inner product, and **Euclidean n -space** is \mathbf{R}^n equipped with this inner product. The dot product of two vectors is zero if they meet at a right angle. \square

18 Lemma If $(x, z) = (y, z)$ for all z , then $x = y$.

Proof: This implies $((x - y), z) = 0$ for all z , so setting $z = x - y$ yields $((x - y), (x - y)) = 0$, which by IP.4 implies $x - y = 0$. \blacksquare

19 Cauchy–Schwartz Inequality In a real inner product space,

$$(x, y)^2 \leq (x, x)(y, y) \tag{1}$$

with $=$ if and only if x and y are dependent.

Proof: (Cf. MacLane and Birkhoff [18, p. 353]) If x or y is zero, then we have equality, so assume x, y are nonzero. Define the quadratic polynomial $Q: \mathbf{R} \rightarrow \mathbf{R}$ by

$$Q(\lambda) = (\lambda x + y, \lambda x + y) = (x, x)\lambda^2 + 2(x, y)\lambda + (y, y).$$

By Property IP.4 of inner products (Definition 15), $Q(\lambda) \geq 0$ for each $\lambda \in \mathbf{R}$. Therefore the discriminant of Q is nonpositive,² that is, $4(x, y)^2 - 4(x, x)(y, y) \leq 0$, or $(x, y)^2 \leq (x, x)(y, y)$. Equality in (1) can occur only if the discriminant is zero, in which case Q has a real root. That is, there is some λ for which $Q(\lambda) = (\lambda x + y, \lambda x + y) = 0$. But this implies that $\lambda x + y = 0$, which means the vectors x and y are linearly dependent. \blacksquare

20 Definition A **norm** $\|x\|$ is a real function on a vector space that satisfies:

N.1 $\|0\| = 0$

N.2 $\|x\| > 0$ if $x \neq 0$

N.3 $\|\alpha x\| = |\alpha| \|x\|$.

²In case you have forgotten how you derived the quadratic formula in Algebra I, rewrite the polynomial as

$$f(z) = \alpha z^2 + \beta z + \gamma = \frac{1}{\alpha} \left(\alpha z + \frac{\beta}{2} \right)^2 - (\beta^2 - 4\alpha\gamma)/4\alpha,$$

and note that the only way to guarantee that $f(z) \geq 0$ for all z is to have $\alpha > 0$ and $\beta^2 - 4\alpha\gamma \leq 0$.

N.4 $\|x + y\| \leq \|x\| + \|y\|$ with equality if and only if $x = 0$ or $y = 0$ or $y = \alpha x$, $\alpha > 0$.

21 Proposition If (\cdot, \cdot) is an inner product, then $\|x\| = (x, x)^{\frac{1}{2}}$ is a norm.

Proof: The only nontrivial part is N.4. Now $\|x + y\| \leq \|x\| + \|y\|$ if and only if $\|x + y\|^2 \leq (\|x\| + \|y\|)^2$, which is equivalent to

$$\underbrace{(x + y, x + y)}_{=(x,x)+2(x,y)+(y,y)} \leq (x, x) + 2\sqrt{(x, x)(y, y)} + (y, y),$$

or


$$(x, y) \leq \sqrt{(x, x)(y, y)}.$$

But this is just the square root of the Cauchy–Schwartz Inequality. ■

If the norm induced by the inner product gives rise to a complete metric space,³ then the inner product space is called a **Hilbert space** or **complete inner product space**.

4.1 The Parallelogram Law

The next result asserts that in an inner product space, the sum of the squares of the lengths of the diagonals of a parallelogram is equal to the sum of the squares of the lengths of the sides. Consider the parallelogram with vertices $0, x, y, x + y$. Its diagonals are the segments $[0, x + y]$ and $[x, y]$, and their lengths are $\|x + y\|$ and $\|x - y\|$. It has two sides of length $\|x\|$ and two of length $\|y\|$. So the claim is:



22 The parallelogram law In an inner product space,

$$\|x + y\|^2 + \|x - y\|^2 = 2\|x\|^2 + 2\|y\|^2.$$

Proof: Note that

$$\begin{aligned} ((x + y), (x + y)) &= (x, x) + 2(x, y) + (y, y) \\ ((x - y), (x - y)) &= (x, x) - 2(x, y) + (y, y). \end{aligned}$$

Add these two to get

$$((x + y), (x + y)) + ((x - y), (x - y)) = 2(x, x) + 2(y, y),$$

and the desired result is restated in terms of norms. ■

On a related note, we have the following.

23 Proposition In an inner product space.

$$\|x + y\|^2 - \|x - y\|^2 = 4(x, y).$$

³A metric space is **complete** if every Cauchy sequence has a limit point in the space. A sequence x_1, x_2, \dots is a Cauchy sequence if $\lim \|x_m - x_n\| = 0$ as $n, m \rightarrow \infty$.

Proof: In the proof above, instead of adding the two equations, subtract them.

$$\begin{aligned}\|x + y\|^2 - \|x - y\|^2 &= ((x + y), (x + y)) - ((x - y), (x - y)) \\ &= (x, x) + 2(x, y) + (y, y) - ((x, x) - 2(x, y) + (y, y)) \\ &= 4(x, y).\end{aligned}$$

■

As an aside, a norm on a vector space is induced by an inner product if and only if it satisfies the parallelogram law; see for instance [3, Problem 32.10, p. 303].

4.2 Orthogonality

24 Definition Vectors x and y are **orthogonal** if $(x, y) = 0$, written

$$x \perp y.$$

A set of vectors $E \subset V$ is **orthogonal** if it is pairwise orthogonal. That is, for all $x, y \in E$ with $x \neq y$ we have $(x, y) = 0$. A set E is **orthonormal** if E is orthogonal and $(x, x) = 1$ for all $x \in E$.

25 Lemma If a set of nonzero vectors is pairwise orthogonal, then the set is independent.

Proof: Suppose $\sum_{i=1}^m \alpha_i x_i = 0$, where the x_i s are pairwise orthogonal. Then for each k ,

$$0 = (x_k, 0) = \left(x_k, \sum_{i=1}^m \alpha_i x_i\right) = \sum_{i=1}^m \alpha_i (x_k, x_i) = \alpha_k (x_k, x_k).$$

Since each $x_k \neq 0$, we have each $\alpha_k = 0$. Thus the vectors x_i are linearly independent. ■

26 Definition The **orthogonal complement** of M in V is the set

$$\{x \in V : (\forall y \in M) [x \perp y]\},$$

denoted M_\perp .

27 Lemma If $x \perp y$ and $x + y = 0$, then $x = y = 0$.

Proof: Now $(x + y, x + y) = (x, x) + 2(x, y) + (y, y)$, so $(x, y) = 0$ implies $(x, x) + (y, y) = 0$. Since $(x, x) \geq 0$ and $(y, y) \geq 0$ we have $x = y = 0$. ■

The next lemma is left as an exercise.

28 Lemma For any set M , the set M_\perp is a linear subspace. If M is a subspace, then

$$M \cap M_\perp = \{0\}.$$

The fact that M_\perp is a subspace is not by itself very useful. It might be the trivial subspace $\{0\}$. We shall see below in Corollary 33 that for finite dimensional spaces, the dimensions of M and M_\perp add up to the dimension of the entire space.

4.3 Orthogonal projection and orthogonal complements

The next result may be found, for instance, in Apostol [6, Theorem 1.14, p. 24], and the proof is known as the **Gram–Schmidt procedure**.

29 Proposition *Let x_1, x_2, \dots be a sequence (finite or infinite) of vectors in an inner product space. Then there exists a sequence y_1, y_2, \dots such that for each m , the span of y_1, \dots, y_m is the same as the span of x_1, \dots, x_m , and the y_i s are orthogonal.*

Proof: Set $y_1 = x_1$. For $m > 1$ recursively define

$$y_m = x_m - \frac{(x_m, y_{m-1})}{(y_{m-1}, y_{m-1})}y_{m-1} - \frac{(x_m, y_{m-2})}{(y_{m-2}, y_{m-2})}y_{m-2} - \dots - \frac{(x_m, y_1)}{(y_1, y_1)}y_1.$$

Use induction on m to prove that the vectors y_1, \dots, y_m are orthogonal and span the same space as the x_i s. Observe that y_2 is orthogonal to $y_1 = x_1$:

$$(y_2, y_1) = (y_2, x_1) = (x_2, x_1) - \frac{(x_2, x_1)}{(x_1, x_1)}(x_2, x_1) = 0.$$

Furthermore any linear combination of x_1 and x_2 can be replicated with y_1 and y_2 .

For $m > 2$, suppose that y_1, \dots, y_{m-1} are orthogonal and span the same space as x_1, \dots, x_{m-1} . Now compute (y_m, y_k) for $k \leq m$:

$$\begin{aligned} (y_m, y_k) &= (x_m, y_k) - \sum_{i=1}^{m-1} \left(\frac{(x_m, y_i)}{(y_i, y_i)} y_i, y_k \right) \\ &= (x_m, y_k) - \sum_{i=1}^{m-1} \frac{(x_m, y_i)}{(y_i, y_i)} (y_i, y_k) \\ &= (x_m, y_k) - \frac{(x_m, y_k)}{(y_k, y_k)} (y_k, y_k) \\ &= 0, \end{aligned}$$

so y_m is orthogonal to every y_1, \dots, y_{m-1} . As an exercise verify that the span of y_1, \dots, y_m is the same as the span of x_1, \dots, x_m . ■

30 Corollary *Every nontrivial finite-dimensional subspace of an inner product space has an orthonormal basis.*

Proof: Apply the Gram–Schmidt procedure to a basis. To obtain an orthonormal basis, simply normalize each y_k by dividing by its norm $(y_k, y_k)^{\frac{1}{2}}$. ■

31 Orthogonal Projection Theorem *Let M be a linear subspace of the finite-dimensional real inner product space V . For each $x \in V$ we can write x in a unique way as $x = x_M + x_\perp$, where $x_M \in M$ and $x_\perp \in M^\perp$.*

Proof: Let y_1, \dots, y_m be an orthonormal basis for M . Put $z_i = (x, y_i)y_i$ for $i = 1, \dots, m$. Put $x_M = \sum_{i=1}^m z_i$, and $x_\perp = x - x_M$. Let $y \in M$, $y = \sum_{i=1}^m \alpha_i y_i$. Then

$$\begin{aligned} (y, x_\perp) &= \left(\sum_{i=1}^m \alpha_i y_i, x - \sum_{j=1}^m (x, y_j)y_j \right) \\ &= \sum_{i=1}^m \alpha_i (y_i, x - \sum_{j=1}^m (x, y_j)y_j) \\ &= \sum_{i=1}^m \alpha_i \left\{ (y_i, x) - (y_i, \sum_{j=1}^m (x, y_j)y_j) \right\} \\ &= \sum_{i=1}^m \alpha_i \left\{ (y_i, x) - \sum_{j=1}^m (x, y_j)(y_i, y_j) \right\} \\ &= \sum_{i=1}^m \alpha_i \{ (y_i, x) - (x, y_i) \} \\ &= 0. \end{aligned}$$

Uniqueness: Let $x = x_M + x_\perp = z_M + z_\perp$. Then $0 = \underbrace{(x_M - z_M)}_{\in M} + \underbrace{(x_\perp - z_\perp)}_{\in M_\perp}$. Rewriting this yields $x_M - z_M = 0 - (x_\perp - z_\perp)$, so $x_M - z_M$ also lies in M_\perp . Similarly $x_\perp - z_\perp$ also lies in M . Thus $x_M - z_M \in M \cap M_\perp = \{0\}$ and $x_\perp - z_\perp \in M \cap M_\perp = \{0\}$. ■

32 Definition The vector x_M given by the theorem above is called the **orthogonal projection** of x on M .

33 Corollary For a finite-dimensional space V and any subspace M of V , $\dim M + \dim M_\perp = \dim V$.

There is another important characterization of orthogonal projection.

34 Proposition (Orthogonal projection minimizes the norm of the “residual”) Let M be a linear subspace of the real inner product space V . Let $y \in V$. Then

$$\|y - y_M\| \leq \|y - x\| \quad \text{for all } x \in M.$$

Proof: Let $x \in M$ and note

$$\begin{aligned} \|y - x\|^2 &= \|(y - y_M) + (y_M - x)\|^2 \\ &= \left((y - y_M) + (y_M - x), (y - y_M) + (y_M - x) \right) \\ &= (y - y_M, y - y_M) + 2(y - y_M, y_M - x) + (y_M - x, y_M - x) \end{aligned}$$

but $y - y_M = y_\perp \perp M$, and $y_M - x \in M$, so $(y - y_M, y_M - x) = 0$, so

$$\begin{aligned} &= \|y - y_M\|^2 + \|y_M - x\|^2 \\ &\geq \|y - y_M\|^2. \end{aligned}$$

That is, $x = y_M$ minimizes $\|y - x\|$ over M . ■

35 Proposition (Linearity of Projection) *The orthogonal projection satisfies*

$$(x + z)_M = x_M + z_M \quad \text{and} \quad (\alpha x)_M = \alpha x_M.$$

Proof: Use $x_M = \sum_{j=1}^k (x, y_j) y_j$ and $z_M = \sum_{j=1}^k (z, y_j) y_j$. Then

$$(x + z)_M = \sum_{j=1}^k (x + z, y_j) y_j.$$

Use linearity of (\cdot, \cdot) . ■

4.4 Orthogonality and alternatives

We now present a lemma about linear functions that is true in quite general linear spaces, see the *Hitchhiker's Guide* [2, Theorem 5.91, p. 212], but we will prove it using some of the special properties of inner products.

36 Lemma *Let V be an inner product space. Then y is a linear combination of x_1, \dots, x_m if and only if*

$$\bigcap_{i=1}^m \{x_i\}_\perp \subset \{y\}_\perp. \tag{2}$$

Proof: If y is a linear combination of x_1, \dots, x_m , say $y = \sum_{i=1}^m \alpha_i x_i$, then

$$(z, y) = \sum_{i=1}^m \alpha_i (z, x_i),$$

so (2) holds.

For the converse, suppose (2) holds. Let $M = \text{span}\{x_1, \dots, x_m\}$ and orthogonally project y on M to get $y = y_M + y_\perp$, where $y_M \in M$ and $y_\perp \perp M$. In particular, $(x_i, y_\perp) = 0$, $i = 1, \dots, m$. Consequently, by hypothesis, $(y, y_\perp) = 0$ too. But

$$0 = (y, y_\perp) = (y_M, y_\perp) + (y_\perp, y_\perp) = 0 + (y_\perp, y_\perp).$$

Thus $y_\perp = 0$, so $y = y_M \in M$. That is, y is a linear combination of x_1, \dots, x_m . ■

We can rephrase the above result as an alternative.

37 Corollary (Fredholm alternative) *Either there exist real numbers $\alpha_1, \dots, \alpha_m$ such that*

$$y = \sum_{i=1}^m \alpha_i x_i$$

or else there exists a vector z satisfying

$$(z, x_i) = 0, \quad i = 1, \dots, m \quad \text{and} \quad (z, y) = 1.$$

4.5 The geometry of the Euclidean inner product

For nonzero vectors x and y in a Euclidean space,

$$x \cdot y = \|x\| \|y\| \cos \theta,$$

where θ is the angle between x and y .

To see this, orthogonally project y on the space spanned by x . That is, write $y = \alpha x + z$ where $z \cdot x = 0$. Thus

$$z \cdot x = (y - \alpha x) \cdot x = y \cdot x - \alpha x \cdot x = 0 \implies \alpha = x \cdot y / x \cdot x.$$

Referring to Figure 1 we see that

$$\cos \theta = \alpha \|x\| / \|y\| = x \cdot y / (\|x\| \|y\|).$$

Thus in an inner product space it makes sense to talk about the angle between nonzero vectors.




Figure 1. Dot product and angles: $\cos \theta = \alpha \|x\| / \|y\| = x \cdot y / (\|x\| \|y\|)$.

38 Definition In a real inner product space, the **angle** $\angle xy$ between two nonzero vectors x and y is defined to be

$$\angle xy = \arccos \frac{(x, y)}{\sqrt{(x, x)(y, y)}}.$$

5 The dual of an inner product space

Recall that the topological dual V^* of a vector space V is the vector subspace of $L(V, \mathbf{R}) = V'$ consisting of continuous real linear functionals on V . When V is an inner product space, V^* has a particularly nice structure. It is clear from the bilinearity of the inner product (Proposition 16) that for every y in the real inner product space V , the function ℓ on V defined by

$$\ell(x) = (y, x)$$

is linear. Moreover, it is continuous with respect to the norm induced by the inner product.

39 Proposition The inner product is jointly norm-continuous. That is, if $\|y_n - y\| \rightarrow 0$ and $\|x_n - x\| \rightarrow 0$, then $(y_n, x_n) \rightarrow (y, x)$.

Proof: By bilinearity and the Cauchy–Schwartz Inequality 19,

$$\begin{aligned}
 |(y_n, x_n) - (y, x)| &= |(y_n \underbrace{-y + y}, x_n \underbrace{-x + x}) - (y, x)| \\
 &= |(y_n - y, x_n - x) + (y_n - y, x) + (y, x_n - x) + (y, x) - (y, x)| \\
 &\leq \sqrt{\|y_n - y\| \|x_n - x\|} + \sqrt{\|y_n - y\| \|x\|} + \sqrt{\|y\| \|x_n - x\|} \\
 &\xrightarrow{n \rightarrow \infty} 0.
 \end{aligned}$$

■

The converse is true if the inner product space is complete as a metric space, that is, if it is a Hilbert space. Every continuous linear functional on a real Hilbert space has the form of an inner product with some vector. The question is, which vector? If $\ell(x) = (y, x)$, then Section 4.5 suggests that ℓ is maximized on the unit sphere when the angle between x and y maximizes the cosine. The maximum value of the cosine is one, which occurs when the angle is zero, that is, when x is a positive multiple of y . Thus to find y given ℓ , we need to find the maximizer of ℓ on the unit ball. But first we need to know that such a maximizer exists.

40 Theorem *Let V be a Hilbert space, and let U be its unit ball:*

$$U = \{x \in V : (x, x) \leq 1\}.$$

Let ℓ be a continuous linear functional on V . Then ℓ has a maximizer in U . If ℓ is nonzero, the maximizer is unique and has norm 1.

Note that if V is finite dimensional, then U is compact, and the result follows from the Weierstrass Theorem. When V is infinite dimensional, the unit ball is not compact, so another technique is needed.

Proof: (Cf. Murray [20, pp. 12–13]) The case where ℓ is the zero functional is obvious, so restrict attention to the case where ℓ is not identically zero.

First observe that if ℓ is continuous, then it is bounded on U . To see this, consider the standard ε - δ definition of continuity at 0, where $\varepsilon = 1$. Then there is some $\delta > 0$ such that if $\|x\| = \|x - 0\| < \delta$, then $|\ell(x)| = |\ell(x) - \ell(0)| < 1$. Thus ℓ is bounded by $|1/\delta|$ on U .

So let $\mu = \sup_{x \in U} \ell(x)$, and note that $0 < \mu < \infty$. Since $x \in U$ implies $-x \in U$ we also have that $-\mu = \inf_{x \in U} \ell(x)$. The scalar μ is called the **operator norm** of ℓ . It has the property that for any $x \in V$,

$$|\ell(x)| \leq \mu \|x\|. \quad (3)$$

To see this, observe that $x/\|x\| \in U$ so $|\ell(x/\|x\|)| \leq \mu$, and the result follows by multiplying by both sides by $\|x\|$.

Pick a sequence x_n in U with $\ell(x_n)$ approaching the supremum μ that satisfies

$$\ell(x_n) \geq \frac{n-1}{n} \mu. \quad (4)$$

We shall show that the sequence x_n is a Cauchy sequence, and so has a limit. This limit is the maximizer. To see that we have a Cauchy sequence, we use the Parallelogram Law 22 to write

$$\|x_n - x_m\|^2 = 2\|x_n\|^2 + 2\|x_m\|^2 - \|x_n + x_m\|^2. \tag{5}$$

Now observe that by (3) and (4),

$$\mu\|x_n\| \geq |\ell(x_n)| = \ell(x_n) \geq \frac{n-1}{n}\mu.$$

Dividing by $\mu > 0$ and recalling that $1 \geq \|x_n\|$ we see that $\|x_n\| \rightarrow 1$ as $n \rightarrow \infty$. (This makes sense. If a nonzero linear functional is going to achieve a maximum over the unit ball, it should happen on the boundary.)

Similarly we have

$$\mu\|x_n + x_m\| \geq |\ell(x_n + x_m)| = \ell(x_n) + \ell(x_m) \geq \left(\frac{n-1}{n} + \frac{m-1}{m}\right)\mu.$$

Dividing by $\mu > 0$ gives

$$\|x_n + x_m\| \geq \frac{n-1}{n} + \frac{m-1}{m}.$$

Square both sides and substitute in (5) to get

$$\|x_n - x_m\|^2 \leq 2\|x_n\|^2 + 2\|x_m\|^2 - \left(\frac{n-1}{n} + \frac{m-1}{m}\right)^2.$$

Since $\|x_n\|, \|x_m\| \rightarrow 1$, the right-hand side tends to $2 + 2 - (1 + 1)^2 = 0$ as $n, m \rightarrow \infty$, so the sequence is indeed a Cauchy sequence.

Thus there is a limit point $x \in U$, with $\|x\| = 1$, and which by continuity satisfies

$$\ell(x) = \mu = \max_{x \in U} \ell(x).$$

To see that the maximizer is unique, suppose $\ell(x) = \ell(y) = \mu > 0$. (Note that this rules out $y = -x$.) Then by the Cauchy–Schwartz Inequality $\|(x + y)/2\| < 1$ (unless y and x are dependent, which in this case means $y = x$), so $\ell((x + y)/\|(x + y)\|) > \mu$, a contradiction. ■

41 Theorem *Let V be a Hilbert space. For every continuous linear functional ℓ on V , there is a vector y in V such that for every $x \in V$,*

$$\ell(x) = (y, x).$$

The correspondence $\ell \leftrightarrow y$ is a homomorphism between V^ and V .*

Proof: (Cf. Murray [20, p. 13]) If ℓ is the zero functional, let $y = 0$. Otherwise, let \hat{y} be the unique maximizer of ℓ over the unit ball (so $\|\hat{y}\| = 1$), and let $\mu = \ell(\hat{y})$. Set

$$y = \mu\hat{y}.$$

Then

$$(y, \hat{y}) = (\mu\hat{y}, \hat{y}) = \mu(\hat{y}, \hat{y}) = \mu = \ell(\hat{y}),$$

so we are off to a good start. We need to show that for every x , we have $(y, x) = \ell(x)$.

We start by showing that if $\ell(x) = 0$, then $(y, x) = 0$. So assume that $\ell(x) = 0$. Then $\ell(\hat{y} \pm \lambda x) = \ell(\hat{y}) = \mu$ for every λ . But by (3) above,

$$\mu = \ell(\hat{y} \pm \lambda x) \leq \mu \|\hat{y} \pm \lambda x\|,$$

so

$$\|\hat{y} \pm \lambda x\| \geq 1 = \|\hat{y}\|.$$

Squaring both sides gives

$$\|\hat{y}\|^2 \pm 2\lambda(\hat{y}, x) + \lambda^2\|x\|^2 \geq \|\hat{y}\|^2$$

or

$$\lambda\|x\|^2 \geq \mp 2(\hat{y}, x)$$

for all $\lambda > 0$. Letting $\lambda \rightarrow 0$ implies $(\hat{y}, x) = 0$, so $(y, x) = 0$.

For an arbitrary x , let

$$x' = x - \ell(x)(\hat{y}/\mu).$$

Then

$$\ell(x') = \ell(x) - \ell(x)\ell(\hat{y}/\mu) = 0.$$

So $(y, x') = 0$. Now observe that

$$\begin{aligned} (y, x) &= (y, x' + \ell(x)(\hat{y}/\mu)) \\ &= (y, x') + \ell(x)(y, \hat{y}/\mu) \\ &= 0 + \ell(x)(\mu\hat{y}, \hat{y}/\mu) \\ &= \ell(x). \end{aligned}$$

This completes the proof. ■

6★ Linear functionals on infinite dimensional spaces

If ℓ is a linear functional on the finite dimensional vector space \mathbf{R}^n , then by Theorem 41, there is a vector $y = (\eta_1, \dots, \eta_n)$ such that for any $x = (\xi_1, \dots, \xi_n)$ we have

$$\ell(x) = \sum_{i=1}^n \eta_i \xi_i,$$

so that ℓ is a continuous function. However, for infinite dimensional normed spaces there are always linear functionals that are discontinuous! Lots of them in fact.

42 Proposition *Every infinite dimensional normed space has a discontinuous linear functional.*

Proof: If X is infinite dimensional normed space, then it has an infinite Hamel basis B . We may normalize each basis vector to have norm one. Let $S = \{x_1, x_2, \dots\}$ be a countable subset of the basis B . Define the function ℓ on the basis B by $\ell(x_n) = n$ for $x_n \in S$, and $\ell(v) = 0$ for $v \in B \setminus S$. Every $x \in X$ has a unique representation as

$$x = \sum_{v \in B} \alpha_v v,$$

where only finitely many α_v are nonzero. Extend ℓ from B to X by

$$\ell(x) = \sum_{v \in B} \alpha_v \ell(v).$$

Then ℓ is a linear functional, but it is not bounded on the unit ball (as $\ell(x_n) = n$). The same argument used in the proof of Theorem 40 shows that in order for ℓ to be continuous, it must be bounded on the unit ball, so it is not continuous. ■

7 Linear transformations

Recall that a linear transformation is a function $T: V \rightarrow W$ between vector spaces satisfying $T(\alpha x + \beta y) = \alpha Tx + \beta Ty$.

There are three important classes of linear transformations of a space into itself. The first is just a rescaling along various dimensions. The identity function is of this class, where the scale factor is one in each dimension. Another important example is orthogonal projection onto a linear subspace. And another important class is rotation about an axis. These do not cover all the linear transformation, since for instance the composition of two of these kinds of transformations does not belong to any of these groups.

43 Definition For a linear transformation $T: V \rightarrow W$, the **null space** or **kernel** of T is the inverse image of 0, that is, $\{x \in V : Tx = 0\}$.

44 Proposition The null space of T , $\{x : Tx = 0\}$, is a linear subspace of V , and the range of T , $\{Tx : x \in V\}$, is a linear subspace of W .

45 Proposition A linear transformation T is one-to-one if and only if its null space is $\{0\}$. In other words, T is one-to-one if and only if $Tx = 0$ implies $x = 0$.

The dimension of the range of T is called the **rank** of T . The dimension of the null space of T is called the **nullity** of T . The next result may be found in [6, Theorem 2.3, p. 34]. You can prove it using Corollary 33 applied to the null space of T .

46 Nullity Plus Rank Theorem Let $T: V \rightarrow W$, where V and W are finite-dimensional inner product spaces. Then

$$\text{rank } T + \text{nullity } T = \dim V.$$

Proof: Let N denote the null space of T . For x in V , decompose it orthogonally as $x = x_N + x_\perp$, where $x_N \in N$ and $x_\perp \in N_\perp$. Then $Tx = Tx_\perp$, so the range of T is just $T(N_\perp)$. Now let x_1, \dots, x_k be a basis for N_\perp . I claim that Tx_1, \dots, Tx_k are independent and therefore constitute a basis for the range of T . So suppose some linear combination $\sum_{i=1}^k \alpha_i Tx_i$ is zero. By linearity of T we have $T\left(\sum_{i=1}^k \alpha_i x_i\right) = 0$. Which implies $\sum_{i=1}^k \alpha_i x_i$ belongs to the null space N . But this combination also lies in N_\perp , so it must be zero. But since the x_i s are independent, it follows that $\alpha_1 = \dots = \alpha_k = 0$. ■

47 Corollary Let $T: V \rightarrow W$ be a linear transformation between m -dimensional spaces. Then T is one-to-one if and only if T is onto.

Proof: (\implies) Assume T is one-to-one. If $Tx = 0 = T0$, $x = 0$ by one-to-oneness. In other words, the null space of T is just $\{0\}$. Then by the Nullity Plus Rank Theorem 46 the rank of T is m , but this mean the image under T of V is all of W .

(\impliedby) Assume T maps V onto W . Then the rank of T is m , so by the Nullity Plus Rank Theorem 46 its null space of T contains only 0. Suppose $Tx = Ty$. Then $T(x - y) = 0$, so $x - y = 0$, which implies T is one-to-one. ■

7.1 Inverse of a linear transformation

Let $T: V \rightarrow W$ be a linear transformation between two vector spaces. A **left inverse** for T is a function $L: \text{range } T \rightarrow V$ such that for all $x \in V$, we have $LTx = x$. A **right inverse** for T is a function $R: \text{range } T \rightarrow V$ such that for all $y \in \text{range } T$, we have $TRy = y$.

48 Lemma (Left vs. right inverses) *If T has a left inverse, then it is unique, and also a right inverse.*

Proof: (Cf. Apostol [6, Theorem 2.8, pp. 39–40].) Let L and M be left inverses for T . Let $y \in \text{range } T$, say $y = Tx$. Then

$$Ly = LTx = x = MTx = My.$$

That is $L = M$ on $\text{range } T$.

Now we show that L is also a right inverse. Let

$$y = Tx$$

belong to $\text{range } T$. Then

$$\begin{aligned} TLy &= TLTx \\ &= T(LT)x = Tx = y. \end{aligned}$$

That is, L is a right inverse. ■

Theorem 117 below describes how to compute the inverse of a matrix.

49 Proposition *Let $T: V \rightarrow W$ be linear, one-to-one, and onto, where V and W are finite-dimensional inner product spaces. Then $\dim V = \dim W$, T has full rank, the inverse T^{-1} exists and is linear.*

Proof: First we show that T^{-1} is linear: Let $x_i = T^{-1}(u_i)$, $i = 1, 2$. Now

$$T(\alpha x_1 + \beta x_2) = \alpha T x_1 + \beta T x_2 = \alpha u_1 + \beta u_2.$$

So taking T^{-1} of both sides gives

$$\alpha \underbrace{T^{-1}(u_1)}_{x_1} + \beta \underbrace{T^{-1}(u_2)}_{x_2} = T^{-1}(\alpha u_1 + \beta u_2)$$

Next we show that $\dim W \leq \dim V$: Let y_1, \dots, y_n be a basis for V . Then Ty_1, \dots, Ty_n span W since T is linear and onto. Therefore $\dim W \leq n = \dim V$.

On the other hand, Ty_1, \dots, Ty_n are linearly independent, so $\dim W \geq \dim V$. To see this, suppose

$$0 = \sum_{i=1}^n \xi_i T e_i = T \left(\sum_{i=1}^n \xi_i e_i \right)$$

Since T is one-to-one and $T0 = 0$, we see that $\sum_{i=1}^n \xi_i y_i = 0$. But this implies $\xi_1 = \dots = \xi_n = 0$.

This shows that $\dim V = \dim W$, and since T is onto it has rank $\dim W$. ■

7.2 Adjoint of a transformation

50 Definition Let

$$T: V \rightarrow W$$

be linear where V and W are Hilbert spaces. The **transpose**, or **adjoint**,⁴ of T , denoted T' , is the linear transformation

$$T': W \rightarrow V$$


such that for every x in V and every y in W .

$$(T'y, x) = (y, Tx). \tag{6}$$

A natural question is, does such a linear transformation T' exist?

51 Proposition The adjoint of T exists, and is uniquely determined by (6).

Proof: By bilinearity of the inner product (Proposition 16) on W , for each $y \in W$, the mapping $\ell_y: z \mapsto (y, z)$ from W to \mathbf{R} is linear, so the composition $\ell_y \circ T: V \rightarrow \mathbf{R}$ is linear. By Theorem 41 on the representation of linear functionals on Hilbert spaces, there is some vector, call it $T'y$, in V so that $(y, Tx) = (\ell_y \circ T)(x) = (T'y, x)$, holds for each $x \in V$. The correspondence $y \mapsto T'y$ defines the transformation $T': W \rightarrow V$.



Let y and z belong to W . Then again by bilinearity,

$$\ell_{\alpha y + \beta z} \circ T = \alpha \ell_y \circ T + \beta \ell_z \circ T.$$

That is, the mapping T' is linear.

Finally, Lemma 18 implies that $T'y$ is uniquely determined by (6). ■

Elaborate this.

52 Proposition Let $S, T: V \rightarrow W$ and $U: W \rightarrow X$. Then

$$(UT)' = T'U'$$

$$(\alpha S + \beta T)' = \alpha S' + \beta T'$$

$$(T')' = T$$

⁴When dealing with complex vector spaces, the definition of the adjoint is modified to $(y, Tx) = \overline{(T'y, x)}$.

53 Theorem Let $T: V \rightarrow W$, where V and W are inner product spaces, so $T': W \rightarrow V$. Then $T'y = 0 \iff y \perp \text{range } T$. In other words,

$$\text{null space } T' = (\text{range } T)_{\perp}.$$

Proof: (\implies) If $T'y = 0$, then $(T'y, x) = (0, x) = 0$ for all x in V . But $(T'y, x) = (y, Tx)$, so $(y, Tx) = 0$ for all x in V . That is, $y \perp \text{range } T$.

(\impliedby) If $y \perp \text{range } T$, then $(y, Tx) = 0$ for all x in V . Therefore $(T'y, x) = 0$ for all x in V , so $T'y = 0$. ■

54 Corollary Let $T: V \rightarrow W$, where V and W are inner product spaces. Then $\text{range } T'T = \text{range } T'$. Consequently, $\text{rank } T' = \text{rank } T'T$.

Proof: Clearly $\text{range } T'T \subset \text{range } T'$.

Let x belong to the range of T' , so $x = T'y$ for some y in W . Let M denote the range of T and consider the orthogonal decomposition $y = y_M + y_{\perp}$. Then $T'y = T'y_M + T'y_{\perp}$, but $T'y_{\perp} = 0$ by Theorem 53. Now $y_M = Tz$ some $z \in W$. Then $x = T'Tz$, so x belongs to the range of $T'T$. ■

55 Corollary Let $T: V \rightarrow W$, where V and W are finite-dimensional inner product spaces. Then $\text{rank } T' = \text{rank } T$

Proof: By Theorem 53, $\text{null space } T' = (\text{range } T)_{\perp} \subset W$. Therefore

$$\dim W - \text{rank } T' = \text{nullity } T' = \dim(\text{range } T)_{\perp} = \dim W - \text{rank } T,$$

where the first equality follows from the Nullity Plus Rank Theorem 46, the second from Theorem 53, and the third from Corollary 33. Therefore, $\text{rank } T = \text{rank } T'$. ■

56 Corollary Let $T: \mathbf{R}^n \rightarrow \mathbf{R}^m$. Then $\text{null space } T = \text{null space } T'T$.

Proof: Clearly $\text{null space } T \subset \text{null space } T'T$, since $Tx = 0 \implies T'Tx = T'0 = 0$.

Now suppose $T'Tx = 0$. Then $(x, T'Tx) = (x, 0) = 0$. But

$$(x, T'Tx) = (T'Tx, x) = (Tx, Tx), \text{IP}$$

where the first equality is IP.1 and the second is the definition of T' , so $Tx = 0$. ■

We can even get a version of the Fredholm Alternative 37 as a corollary. I leave it to you to unravel why I think this a version of the Fredholm Alternative.

57 Corollary (Fredholm Alternative II) Let $T: V \rightarrow W$, where V and W are inner product spaces, and let z belong to W . Then either

$$\text{there exists } x \in V \text{ with } z = Tx,$$

or else

$$\text{there exists } y \in W \text{ with } (y, z) \neq 0 \text{ \& } T'y = 0.$$

Proof: The first alternative states that $z \in \text{range } T$, or equivalently (by Theorem 53), that $z \in (\text{null space } T')_{\perp}$. If this is not the case, that is, if $z \notin (\text{null space } T')_{\perp}$, then there must be some y in null space T' that is not orthogonal to z . Such a y satisfies $T'y = 0$ and $(y, z) \neq 0$, which is the second alternative.

To see that the two alternatives are inconsistent, suppose that x and y satisfy the alternatives. Then $0 \neq (y, z) = (y, Tx) = (T'y, x) = (0, x) = 0$, a contradiction. (The middle equality is just the definition of the transpose.) ■

58 Proposition (Summary) For a linear transformation T between finite-dimensional spaces, $\text{range } T'T = \text{range } T'$ and $\text{range } TT' = \text{range } T$, so

$$\text{rank } T = \text{rank } T' = \text{rank } T'T = \text{rank } TT'.$$

7.3 Orthogonal transformations

59 Definition Let V be a real inner product space, and let $T: V \rightarrow V$ be a linear transformation of V into itself. We say that T is an **orthogonal transformation** if its adjoint is its inverse,

$$T' = T^{-1}.$$

60 Proposition For a linear transformation $T: V \rightarrow V$ on an inner product space, the following are equivalent.

1. T is orthogonal. That is, $T' = T^{-1}$.
2. T preserves norms. That is, for all x ,

$$(Tx, Tx) = (x, x). \tag{7}$$

3. T preserves inner products, that is, for every $x, y \in V$,

$$(Tx, Ty) = (x, y).$$

Proof: (1) \implies (2) Assume T is orthogonal. Fix x and let $y = Tx$. By the definition of T' we have

$$(T'y, x) = (y, Tx),$$

so

$$(x, x) = (T'Tx, x) = (T'y, x) = (y, Tx) = (Tx, Tx).$$

- (2) \implies (3) Assume T preserves norms. By Proposition 23,

$$\begin{aligned} (Tx, Ty) &= \frac{\|Tx + Ty\|^2 - \|Tx - Ty\|^2}{4} \\ &= \frac{\|T(x + y)\|^2 - \|T(x - y)\|^2}{4} \\ &= \frac{\|x + y\|^2 - \|x - y\|^2}{4} \\ &= (x, y). \end{aligned}$$

(3) \implies (1) Assume T preserves inner products. By the definition of T' , for all x, y ,

$$(T'y, x) = (y, Tx).$$

Taking $y = Tz$, we have

$$(T'Tz, x) = (T'y, x) = (y, Tx) = (Tz, Tx) = (z, x),$$

so by Lemma 18, $T'Tz - z = 0$ for every z , which is equivalent to $T' = T^{-1}$. ■

A norm preserving mapping is also called an **isometry**. Since the composition of norm-preserving mappings preserves norms we have the following.

61 Corollary *The composition of orthogonal transformations is an orthogonal transformation.*

An orthogonal transformation preserves angles (since it preserves inner products) and distances between vectors. Reflection and rotation are the basic orthogonal transformations in a finite-dimensional Euclidean space.

7.4 Symmetric transformations

62 Definition A transformation $T: \mathbf{R}^m \rightarrow \mathbf{R}^m$ is **symmetric** or **self-adjoint** if $T' = T$. A transformation T is **skew-symmetric** if $T' = -T$.

63 Proposition Let $\pi_M: \mathbf{R}^m \rightarrow \mathbf{R}^m$ be the orthogonal projection on M . Then π_M is symmetric.

Proof: To show $(x, \pi_M z) = (\pi_M x, z)$. Observe that

$$(x, z_M) = (x_M + x_\perp, z_M) = (x_M, z_M) + (x_\perp, z_M) = (x_M, z_M) + 0.$$

Similarly,

$$(x_M, z) = (x_M, z_M + z_\perp) = (x_M, z_M) + (x_M, z_\perp) = (x_M, z_M) + 0.$$

Therefore $(x, z_M) = (x_M, z) = (x_M, z_M)$. ■

In light of the following exercise, a linear transformation is an orthogonal projection if and only if it is symmetric and idempotent. A transformation T is **idempotent** if $T^2x = TTx = Tx$ for all x .

64 Exercise Let $P: \mathbf{R}^m \rightarrow \mathbf{R}^m$ be a linear transformation satisfying

$$P' = P \quad \text{and} \quad P^2x = Px \quad \text{for all } x \in \mathbf{R}^m.$$

That is, P is idempotent and symmetric. Set $M = I - P$ (where I is the identity on \mathbf{R}^m). Prove the following.

1. For any x , $x = Mx + Px$.
2. $M^2x = Mx$ for all x and $M' = M$.
3. null space $P = (\text{range } P)_\perp = \text{range } M$ and null space $M = (\text{range } M)_\perp = \text{range } P$
4. P is the orthogonal projection onto its range. Likewise for M .

□

8 Eigenvalues and eigenvectors

65 Definition Let V be a real vector space, and let T be a linear transformation of V into itself, $T: V \rightarrow V$. A real number λ is an **eigenvalue** of T if there is a nonzero vector x in V such that $Tx = \lambda x$. The vector x is called an **eigenvector** of T associated with λ . Note that the vector 0 is by definition not an eigenvector of T .⁵

If T has an eigenvalue λ with eigenvector x , the transformation “stretches” the space by a factor λ in the direction x .

While the vector 0 is never an eigenvector, the scalar 0 may be an eigenvalue. Indeed 0 is the eigenvalue associated with any nonzero vector in the null space of T .

There are linear transformations with no (real) eigenvalues. For instance, consider the rotation of \mathbf{R}^2 by ninety degrees. This is given by the transformation $(x, y) \mapsto (-y, x)$. In order to satisfy $Tx = \lambda x$ we must have $\lambda x = -y$ and $\lambda y = x$. This cannot happen for any nonzero real vector (x, y) and real λ .

On the other hand, the identity transformation has an eigenvalue 1 , associated with every nonzero vector.

Observe that there is a unique eigenvalue associated with each eigenvector: If $Tx = \lambda x$ and $Tx = \alpha x$, then $\alpha x = \lambda x$, so $\alpha = \lambda$, since by definition x is nonzero.

On the other hand, one eigenvalue must be associated with many eigenvectors, for if x is an eigenvector associated with λ , so is any nonzero scalar multiple of x . More generally, a linear combination of eigenvectors corresponding to an eigenvalue is also an eigenvector corresponding to the same eigenvalue (provided the linear combination does not equal the zero vector). The span of the set of eigenvectors associated with the eigenvalue λ is called the **eigenspace** of T corresponding to λ . Every nonzero vector in the eigenspace is an eigenvector associated with λ . The dimension of the eigenspace is called the **multiplicity** of λ .

66 Proposition If $T: V \rightarrow V$ is idempotent, then each of its eigenvalues is either 0 or 1 .

Proof: Suppose $Tx = \lambda x$ with $x \neq 0$. Since T is idempotent, we have $\lambda x = Tx = T^2x = \lambda^2x$. Since $x \neq 0$, this implies $\lambda = \lambda^2$, so $\lambda = 0$ or $\lambda = 1$. ■

For distinct eigenvalues we have the following, taken from [6, Theorem 4.1, p. 100].

67 Theorem Let x_1, \dots, x_n be eigenvectors associated with distinct eigenvalues $\lambda_1, \dots, \lambda_n$. Then the vectors x_1, \dots, x_n are independent.

Proof: The proof is by induction n . The case $n = 1$ is trivial, since by definition eigenvectors are nonzero. Now consider $n > 1$ and suppose that the result is true for $n - 1$. Now let

$$\sum_{i=1}^n \alpha_i x_i = 0. \quad (8)$$

⁵For a linear transformation of a complex vector space, eigenvalues may be complex, but I shall only deal with real vector spaces here.

Applying the transformation T to both sides gives

$$\sum_{i=1}^n \alpha_i \lambda_i x_i = 0. \tag{9}$$

Let us eliminate x_n from (9) by multiplying (8) by λ_n and subtracting to get

$$\sum_{i=1}^n \alpha_i (\lambda_i - \lambda_n) x_i = \sum_{i=1}^{n-1} \alpha_i (\lambda_i - \lambda_n) x_i = 0.$$

But x^1, \dots, x^{n-1} are independent so by the induction hypothesis $\alpha_i (\lambda_i - \lambda_n) = 0$ for $i = 1, \dots, n - 1$. Since the eigenvalues are distinct, this implies each $\alpha_i = 0$ for $i = 1, \dots, n - 1$. Thus (8) reduces to $\alpha_n x_n = 0$, which implies $\alpha_n = 0$. This shows that x_1, \dots, x_n are independent. ■

68 Corollary *A linear transformation on an n -dimensional space has at most n distinct eigenvalues. If it has n , then the space has a basis made up of eigenvectors.*

When T is a symmetric transformation of an inner product space into itself, not only are eigenvectors associated with distinct eigenvalues independent, they are orthogonal.

69 Proposition *Let V be a real inner product space, and let T be a symmetric linear transformation of V into itself. Let x and y be eigenvectors of T corresponding to eigenvalues α and λ with $\alpha \neq \lambda$. Then $x \perp y$.*

Proof: We are given $Tx = \alpha x$ and $Ty = \lambda y$. Thus $(Tx, y) = (\lambda x, y) = \lambda(x, y)$ and $(x, Ty) = (x, \alpha y) = \alpha(x, y)$. Since T is symmetric, $(Tx, y) = (x, Ty)$, so $\alpha(x, y) = \lambda(x, y)$. Since $\lambda \neq \alpha$ we must then have $(x, y) = 0$. ■

Also if T is symmetric, we are guaranteed that it has plenty of eigenvectors.

70 Theorem *Let $T: V \rightarrow V$ be symmetric, where V is an n -dimensional inner product space. Then V has an orthonormal basis consisting of eigenvectors of T .*



Proof: This proof uses some well known results from topology and calculus that are beyond the scope of these notes. Cf. Anderson [4, pp. 273–275], Carathéodory [10, § 195], Franklin [12, Section 6.2, pp. 141–145], Rao [21, 1f.2.iii, p. 62].

Let $S = \{x \in V : (x, x) = 1\}$ denote the unit sphere in V . Set $M_0 = \{0\}$ and define $S_0 = S \cap M_{0\perp}$. Define the quadratic form $Q: V \rightarrow \mathbf{R}$ by

$$Q(x) = (x, Tx).$$

It is easy to see that Q is continuous, so it has a maximizer on S_0 , which is compact. (This maximizer cannot be unique, since $Q(-x) = Q(x)$, and indeed if T is the identity, then Q is constant on S .) Fix a maximizer x_1 of Q over S_0 .

Proceed recursively for $k = 1, \dots, n - 1$. Let M_k denote the span of x_1, \dots, x_k , and set $S_k = S \cap M_{k\perp}$. Let x_{k+1} maximize Q over S_k . By construction, $x_{k+1} \in M_{k\perp}$, so the x_k 's are orthogonal, indeed orthonormal.

Since x_1 maximizes Q on $S = S_0$, it maximizes Q subject to the constraint $1 - (x, x) = 0$. Now $Q(x) = (x, Tx)$ is continuously differentiable and $Q'(x) = 2Tx$, and the gradient of the constraint function is $-2x$, which is clearly nonzero (hence linearly independent) on S . It is a nuisance to have these 2s popping up, so let us agree to maximize $\frac{1}{2}(x, Tx)$ subject $\frac{1}{2}(1 - (x, x)) = 0$ instead. Therefore by the well known Lagrange Multiplier Theorem, there exists λ_1 satisfying

$$Tx_1 - \lambda_1 x_1 = 0.$$

This obviously implies that the Lagrange multiplier λ_1 is an eigenvalue of T and x_1 is a corresponding eigenvector. Further, it is the value of the maximum:

$$Q(x_1) = (x_1, Tx_1) = (x_1, \lambda_1 x_1) = \lambda_1,$$

since $(x_1, x_1) = 1$.

Let x_1, \dots, x_n be defined as above and assume that for $i = 1, \dots, k < n$, each x_i is an eigenvector of T and that $\lambda_i = Q(x_i)$ is its corresponding eigenvalue. We wish to show that x_{k+1} is an eigenvector of T and $\lambda_{k+1} = Q(x_{k+1})$ is its corresponding eigenvalue.

By construction, x_{k+1} maximizes $\frac{1}{2}Q(x)$ subject to the $k + 1$ constraints

$$\frac{1}{2}(1 - (x, x)) = 0, \quad (x, x_1) = 0, \quad \dots \quad (x, x_k) = 0.$$

The gradients of these constraint functions are $-x$ and x_1, \dots, x_k respectively. By construction, x_1, \dots, x_{k+1} are orthonormal, so at $x = x_{k+1}$ the gradients of the constraint are linearly independent. Therefore there exist Lagrange multipliers λ_{k+1} and μ_1, \dots, μ_k satisfying

$$Tx_{k+1} - \lambda_{k+1}x_{k+1} + \mu_1x_1 + \dots + \mu_kx_k = 0. \tag{10}$$

Therefore

$$Q(x_{k+1}) = (x_{k+1}, Tx_{k+1}) = \lambda_{k+1}(x_{k+1}, x_{k+1}) - \mu_1(x_{k+1}, x_1) - \dots - \mu_k(x_{k+1}, x_k) = \lambda_{k+1},$$

since x_1, \dots, x_{k+1} are orthonormal. That is, the multiplier λ_{k+1} is the maximum value of Q over S_k .

By hypothesis, $Tx_i = \lambda_i x_i$ for $i = 1, \dots, k$. Then since T is symmetric,

$$(x_i, Tx_{k+1}) = (x_{k+1}, Tx_i) = (x_{k+1}, \lambda_i x_i) = 0, \quad i = 1, \dots, k.$$

That is, $Tx_{k+1} \in M_{k\perp}$, so $Tx_{k+1} - \lambda_{k+1}x_{k+1} \in M_{k\perp}$, so by Lemma 27 equation (10) implies

$$\underbrace{Tx_{k+1} - \lambda_{k+1}x_{k+1}}_{\in M_{k\perp}} = 0 \quad \text{and} \quad \underbrace{\mu_1x_1 + \dots + \mu_kx_k}_{\in M_k} = 0.$$

We conclude therefore that $Tx_{k+1} = \lambda_{k+1}x_{k+1}$, so that x_{k+1} is an eigenvector of T and λ_{k+1} is the corresponding eigenvalue.

Since V is n -dimensional, x_1, \dots, x_n is an orthonormal basis of eigenvectors, as we desire. ■

9 Matrices

A **matrix** is merely a rectangular array of numbers, or equivalently, a doubly indexed ordered list of real numbers. An $m \times n$ matrix has m rows and n columns. The entries in a matrix are doubly indexed, with the first index denoting its row and the second its column. Here is a generic matrix:

$$A = \begin{bmatrix} \alpha_{1,1} & \cdots & \alpha_{1,n} \\ \vdots & & \vdots \\ \alpha_{m,1} & \cdots & \alpha_{m,n} \end{bmatrix}$$

The generally accepted plural of *matrix* is *matrices*, although *matrixes* is okay too. Matrices are of interest for two separate but hopelessly intertwined reasons. One is their relation to systems of linear equations and inequalities, and the other is their connection to linear transformations between finite-dimensional vector spaces.

The set of $m \times n$ matrices is denoted

$$M(m, n).$$

Given a matrix A , let A_i denote the i^{th} row of A and let A^j denote the j^{th} column. The i^{th} row and j^{th} column entry is generally denoted by a lower case Greek letter, e.g., $\alpha_{i,j}$. We can identify the rows or columns of a matrix with singly indexed lists of real numbers, that is, elements of some \mathbf{R}^k . If A is $m \times n$, the **column space** of A is the subset of \mathbf{R}^m spanned by the n columns of A . Its **row space** is the subspace of \mathbf{R}^n spanned by its m rows.

9.1 Matrix operations

If both A and B are $m \times n$ matrices, the **sum** $A + B$ is the $m \times n$ matrix C defined by

$$\gamma_{i,j} = \alpha_{i,j} + \beta_{i,j}.$$

The **scalar multiple** αA of a matrix A by a scalar α is the matrix M defined by $m_{i,j} = \alpha \alpha_{i,j}$. Under these operations the set of $m \times n$ matrices becomes an mn -dimensional vector space.

71 Proposition *The set $M(m, n)$ is a vector space under the operations of matrix addition and scalar multiplication. It has dimension mn .*

If A is $m \times p$ and B is $p \times n$, the **product** of A and B is the $m \times n$ matrix whose i^{th} row, j^{th} column entry is the dot product $A_i \cdot B^j$ of the i^{th} row of A with the j^{th} column of B .

$$(AB)_{i,j} = A_i \cdot B^j$$

The reason for this peculiar definition is explained in the next section.

Vectors in \mathbf{R}^n may also be considered matrices as either a column or a row vector. Let x be a row m -vector (a $1 \times m$ matrix), y be a column n -vector (an $n \times 1$ matrix), and let A be an $m \times n$ matrix. Then the matrix product xA considered as a vector in \mathbf{R}^m belongs to the row space of A , and Ay as a vector in \mathbf{R}^m belongs to the column space of A .

$$xA = \sum_{i=1}^m \xi_i A_i \quad \text{and} \quad Ay = \sum_{j=1}^n \eta_j A^j$$

Note that the i^{th} row of AB is given by

$$(AB)_i = (A_i)B$$

and the j^{th} column of AB is given by

$$(AB)^j = A(B^j).$$

The **main diagonal** of a square matrix $A = [\alpha_{i,j}]$ is the set of $\alpha_{i,j}$ with $i = j$. A matrix is called a **diagonal matrix** if it is square and all its nonzero entries are on the main diagonal. A square matrix is **upper triangular** if the only nonzero elements are on or above the main diagonal, that is, if $i > j$ implies $\alpha_{i,j} = 0$. A square matrix is **lower triangular** if $i < j$ implies $\alpha_{i,j} = 0$.

The $n \times n$ diagonal matrix I whose diagonal entries are all 1 and off-diagonal entries are all 0 has the property that

$$AI = IA = A$$

for any $n \times n$ matrix A . The matrix I is called the **$n \times n$ identity matrix**. The **zero matrix** has all its entries zero, and satisfies $A + 0 = A$.

72 Fact (Summary) *Direct computation reveals the following facts.*

$$\begin{aligned} (AB)C &= A(BC) \\ AB &\neq BA \quad (\text{in general}) \\ A(B+C) &= (AB) + (AC) \\ (A+B)C &= AC + BC \end{aligned}$$

73 Exercise Verify the following.

1. The product of upper triangular matrices is upper triangular.
2. The product of lower triangular matrices is lower triangular.
3. The product of diagonal matrices is diagonal.
4. The inverse of an upper triangular matrix is upper triangular (if it exists).
5. The inverse of a lower triangular matrix is lower triangular (if it exists).
6. The inverse of a diagonal matrix is diagonal (if it exists).

□

9.2 Systems of linear equations

One of the main uses of matrices is the simplification of representing a system of linear equations. For instance, consider the following system of equations.

$$\begin{aligned} 3x_1 + 2x_2 &= 8 \\ 2x_1 + 3x_2 &= 7 \end{aligned}$$

It has a unique solution $x_1 = 2$ and $x_2 = 1$. One way to solve this is to take the second equation and solve for $x_1 = \frac{7}{2} - \frac{3}{2}x_2$ and substitute this into the first equation to get $3(\frac{7}{2} - \frac{3}{2}x_2) + 2x_2 = 8$, so $x_2 = 1$ and $x_1 = \frac{7}{2} - \frac{3}{2} = 2$. However Gauss formalized a computationally efficient way to attack these problems using **elementary row operations**. The first step is to write down the so-called augmented coefficient matrix of the system, which is the 2×3 matrix of just the numbers above:

$$\begin{bmatrix} 3 & 2 & 8 \\ 2 & 3 & 7 \end{bmatrix}.$$

There are three elementary row operations, and they correspond to steps used to solve a system of equations. One of these operations is to interchange two rows. We won't use that here. Another is to multiply a row by a nonzero scalar. This does not change the solution. The third operation is to add one row to another. These last two operations can be combined, and we can think of adding a scalar multiple of one row to another as an elementary row operation. We apply these operations until we get a matrix of the form

$$\begin{bmatrix} 1 & 0 & a \\ 0 & 1 & b \end{bmatrix}$$

which is the augmented matrix of the system

$$\begin{aligned} x_1 &= a \\ x_2 &= b \end{aligned}$$

and the system is solved. There is a simple algorithm for deciding which elementary row operations to apply, namely, the **Gaussian elimination algorithm**. In a section below, we shall go into this algorithm in detail, but let me just give you a hint here. First we multiply the first row by $\frac{1}{3}$, to get a leading 1:

$$\begin{bmatrix} 1 & \frac{2}{3} & \frac{8}{3} \\ 2 & 3 & 7 \end{bmatrix}$$

We want to eliminate x_1 from the second equation, so we add an appropriate multiple of the first row to the second. In this case the multiple is -2 , the result is:

$$\begin{bmatrix} 1 & \frac{2}{3} & \frac{8}{3} \\ 2 - 2 \cdot 1 & 3 - 2 \cdot \frac{2}{3} & 7 - 2 \cdot \frac{8}{3} \end{bmatrix} = \begin{bmatrix} 1 & \frac{2}{3} & \frac{8}{3} \\ 0 & \frac{5}{3} & \frac{5}{3} \end{bmatrix}.$$

Now multiply the second row by $\frac{3}{5}$ to get

$$\begin{bmatrix} 1 & \frac{2}{3} & \frac{8}{3} \\ 0 & 1 & 1 \end{bmatrix}.$$

Finally to eliminate x_2 from the first row we add $-\frac{2}{3}$ times the second row to the first and get

$$\begin{bmatrix} 1 - \frac{2}{3} \cdot 0 & \frac{2}{3} - \frac{2}{3} \cdot 1 & \frac{8}{3} - \frac{2}{3} \cdot 1 \\ 0 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 1 \end{bmatrix},$$

which accords with our earlier result.

9.3 Matrix representation of a linear transformation

Let T be a linear transformation from the n -dimensional space V into the m -dimensional space W . Let x_1, \dots, x_n be an ordered basis for V and y_1, \dots, y_m be an ordered basis for W . Let

$$\begin{aligned} Tx_1 &= \sum_{i=1}^m \alpha_{i,1} y_i \\ Tx_2 &= \sum_{i=1}^m \alpha_{i,2} y_i \\ &\vdots \\ Tx_n &= \sum_{i=1}^m \alpha_{i,n} y_i. \end{aligned}$$

The $m \times n$ array of numbers

$$M(T) = \begin{bmatrix} \alpha_{1,1} & \dots & \alpha_{1,n} \\ \vdots & & \vdots \\ \alpha_{m,1} & \dots & \alpha_{m,n} \end{bmatrix}$$

is the **matrix representation of T with respect to the ordered bases $(x_j), (y_i)$** . Note that the j^{th} column of this matrix is the coordinate vector of Tx_j with respect to the ordered basis y_1, \dots, y_m .

This representation provides an isomorphism between matrices and linear transformations. The proof is left as an exercise.

74 Proposition *Let V be an n -dimensional vector space, and let W be an m -dimensional vector space. Fix an ordered basis for each, and let $M(T)$ be the matrix representation a linear transformation $T: V \rightarrow W$. Then the mapping*

$$T \rightarrow M(T)$$

is a linear isomorphism from $L(V, W)$ to $M(m, n)$.

Given the coordinate vector of an element x of V , we can use matrix multiplication to compute the coordinates of Tx . Let $x = \sum_{j=1}^n \xi_j x_j$. Then


$$\begin{aligned} Tx &= \sum_{j=1}^n \xi_j Tx_j \\ &= \sum_{j=1}^n \xi_j \sum_{i=1}^m \alpha_{i,j} y_i \\ &= \sum_{i=1}^m \left(\sum_{j=1}^n \alpha_{i,j} \xi_j \right) y_i \end{aligned}$$

The coordinates of Tx with respect to (y_i) are given by

$$(Tx)_i = \sum_{j=1}^n \alpha_{i,j} \xi_j.$$

That is, the coordinate vector of Tx is $M(T)$ times the column vector of coordinates of x .

The matrix representation can be thought of in the following terms. Let X denote the coordinate mapping of V onto \mathbf{R}^n with respect to the ordered basis x_1, \dots, x_n . Similarly Y denotes the coordinate mapping from W onto \mathbf{R}^m . In order to compute the coordinates of Tz , we first find the coordinate vector of Xz , and then multiply by the matrix $M(T)$, as the following “commutative diagram” shows.



75 Example (The matrix of the coordinate mapping) Consider the case $V = \mathbf{R}^m$. That is, elements of V are already thought of as ordered lists of real numbers. Let x_1, \dots, x_m be an ordered basis, and let X denote the coordinate mapping from \mathbf{R}^m with this basis to \mathbf{R}^m with the standard basis. Then $M(X)$ is simply the matrix whose j^{th} column is x_j . \square

76 Example (Matrices as linear transformations) Let $A = [\alpha_{i,j}]$ be an $m \times n$ matrix and

$$x = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$$

a $n \times 1$ matrix. The matrix product Ax is an $m \times 1$ matrix whose i^{th} row is

$$\sum_{j=1}^n \alpha_{i,j} x_j \quad i = 1, \dots, m.$$

Then $T: x \mapsto Ax$ defines a linear transformation from \mathbf{R}^n to \mathbf{R}^m . The matrix $M(T)$ of this transformation with respect to the standard ordered bases (the bases of unit coordinate vectors) is just A . \square

77 Definition *The rank of a matrix is the largest number of linearly independent columns.*

It follows from Proposition 49 that:

78 Proposition *An $m \times m$ matrix has an inverse if and only if it has rank m .*

79 Example (The identity matrix) What is the matrix representation for the identity mapping $I: \mathbf{R}^m \rightarrow \mathbf{R}^m$?

$$M(I) = \begin{bmatrix} 1 & & 0 \\ & \ddots & \\ 0 & & 1 \end{bmatrix}$$

is the $m \times m$ identity matrix I . If the transformation T is invertible, so that $TT^{-1} = I$, then $M(T)M(T^{-1}) = I$. The matrix $M(T^{-1})$ is naturally referred to as the **inverse of the matrix $M(T)$** . In general, if A defines an invertible transformation from \mathbf{R}^m onto \mathbf{R}^m , the matrix A^{-1} satisfies

$$AA^{-1} = A^{-1}A = I.$$

□

80 Example (The zero matrix) The matrix for the zero transformation $0: \mathbf{R}^n \rightarrow \mathbf{R}^m$, is

$$\begin{bmatrix} 0 & \cdots & 0 \\ \vdots & & \vdots \\ 0 & \cdots & 0 \end{bmatrix},$$

the $m \times n$ zero matrix.

□

The **transpose of a matrix** is the matrix formed by interchanging rows and columns. This definition is justified by the following lemma.

81 Lemma $M(T)_{i,j} = M(T')_{j,i}$.

Proof:

$$M(T)_{i,j} = (e_i, Te_j) = (T'e_i, e_j) = (e_j, T'e_i) = M(T')_{j,i}$$

■

9.4 Gershgorin's Theorem

For a diagonal matrix, the diagonal elements of the matrix are the eigenvalues of the corresponding linear transformation. Even if the matrix is not diagonal, its eigenvalues are “close” to the diagonal elements.

82 Gershgorin's Theorem Let $A = [\alpha_{i,j}]$ be an $m \times m$ real matrix. If λ is an eigenvalue of A , that is, if $Ax = \lambda x$ for some nonzero x , then for some i ,

$$|\lambda - \alpha_{i,i}| \leq \sum_{j:j \neq i} |\alpha_{i,j}|.$$

Proof: (Cf. Franklin [12, p. 162].) Let x be an eigenvector of A belonging to the eigenvalue λ . Choose i so that

$$|x_i| = \max\{|x_1|, \dots, |x_m|\},$$

and note that since $x \neq 0$, we have $|x_i| > 0$. Then by the definition of eigenvalue and eigenvector,

$$(\lambda I - A)x = 0.$$

Now the i^{th} component of the vector $(\lambda I - A)x$ is just $(\lambda - \alpha_{i,i})x_i - \sum_{j:j \neq i} \alpha_{i,j}x_j = 0$, so

$$(\lambda - \alpha_{i,i})x_i = \sum_{j:j \neq i} \alpha_{i,j}x_j$$

so taking absolute values,

$$\begin{aligned} |\lambda - \alpha_{i,i}| |x_i| &= \left| \sum_{j:j \neq i} \alpha_{i,j}x_j \right| \\ &\leq \sum_{j:j \neq i} |\alpha_{i,j}| |x_j| \end{aligned}$$

so dividing by $|x_i| > 0$,

$$\begin{aligned} |\lambda - \alpha_{i,i}| &\leq \sum_{j:j \neq i} |\alpha_{i,j}| \frac{|x_j|}{|x_i|} \\ &\leq \sum_{j:j \neq i} |\alpha_{i,j}|. \end{aligned}$$

■

A square matrix A has a **dominant diagonal** if for each i ,

$$|a_{i,i}| > \sum_{j:j \neq i} |a_{i,j}|.$$

Note that any diagonal matrix with nonzero diagonal elements has a dominant diagonal. This leads to the following corollary of Gershgorin's Theorem.

83 Corollary *If A is a nonnegative square matrix with a dominant diagonal, then every eigenvalue of A is strictly positive.*

Proof: Since A is nonnegative, if λ is an eigenvalue, then Gershgorin's theorem implies that for some i ,

$$\alpha_{i,i} - \lambda \leq |\lambda - \alpha_{i,i}| \leq \sum_{j:j \neq i} \alpha_{i,j},$$

so

$$0 < \alpha_{i,i} - \sum_{j:j \neq i} \alpha_{i,j} \leq \lambda,$$

where the first inequality is the dominant diagonal property. ■

For application of this corollary and related results, see McKenzie [19].

9.5 Matrix representation of a composition

Let S take $\mathbf{R}^p \rightarrow \mathbf{R}^n$ be linear with matrix $M(S) = [\beta_{j,k}]_{j=1,\dots,n}^{k=1,\dots,p}$. Let T take $\mathbf{R}^n \rightarrow \mathbf{R}^m$ be linear with matrix $M(T) = [\alpha_{i,j}]_{i=1,\dots,m}^{j=1,\dots,n}$. Then $T \circ S: \mathbf{R}^p \rightarrow \mathbf{R}^m$ is linear. What is $M(TS)$?

Let $x = \begin{bmatrix} cx_1 \\ \vdots \\ x_p \end{bmatrix}$. Then

$$Sx = \sum_{j=1}^n \left(\sum_{k=1}^p \beta_{j,k} x_k \right) e_j.$$

$$T(Sx) = \sum_{i=1}^m \left(\sum_{j=1}^n \alpha_{i,j} \left(\sum_{k=1}^p \beta_{j,k} x_k \right) \right) e_i.$$

Set

$$\gamma_{i,k} = \sum_{j=1}^n \alpha_{i,j} \beta_{j,k}.$$

Then

$$T(Sx) = \sum_{i=1}^m \left(\sum_{k=1}^p \gamma_{i,k} x_k \right) e_i.$$

Thus $M(TS) = [\gamma_{i,k}]_{i=1,\dots,m}^{k=1,\dots,p}$. This proves the following theorem.

84 Theorem $M(TS) = M(T)M(S)$.

Thus multiplication of matrices corresponds to composition of linear transformations.

Let $S, T: \mathbf{R}^n \rightarrow \mathbf{R}^m$ be linear transformations with matrices A, B . Then $S + T$ is linear. The matrix for $S + T$ is $A + B$ where $(A + B)_{i,j} = \alpha_{i,j} + \beta_{i,j}$.

9.6 Change of basis

A linear transformation may have different matrix representations for different bases. Is there some way to tell if two matrices represent the same transformation?


We shall answer this question for the important case where T maps an m -dimensional vector space V into itself, and we use the same basis for V as both the domain and the range. Let $A = [\alpha_{i,j}]$ represent T with respect to the ordered basis x_1, \dots, x_m , and let $B = [\beta_{i,j}]$ represent T with respect to the ordered basis y_1, \dots, y_m .

That is,

$$Tx_i = \sum_{k=1}^m \alpha_{k,i} x_k \quad \text{and} \quad Ty_i = \sum_{k=1}^m \beta_{k,i} y_k.$$

Let X be the coordinate map from V into \mathbf{R}^m with respect to the ordered basis x_1, \dots, x_m , and let Y be the coordinate map for the ordered basis y_1, \dots, y_m . Then X and Y have full rank and so have inverses.

Consider the following commutative diagram.



The mapping XY^{-1} from \mathbf{R}^m onto \mathbf{R}^m followed by A from \mathbf{R}^m into \mathbf{R}^m , which maps the upper left \mathbf{R}^m into the lower right \mathbf{R}^m , is the same as B followed by XY^{-1} . Let C be the matrix representation of XY^{-1} with respect to the standard ordered basis of unit coordinate vectors. Then

$$A = CBC^{-1} \quad \text{and} \quad B = C^{-1}AC. \tag{*}$$

85 Definition Two square matrices A and B are called **similar** if there is some nonsingular C such that $(*)$ holds.

So we have already proved half of the following. The second half is left for you. (See Apostol [6, Theorem 4.7, p. 110] if you get stuck.)

86 Theorem Two matrices are similar if and only if they represent the same linear transformation.

The following are corollaries, but have simple direct proofs.

87 Proposition If A, B are similar with $A = CBC^{-1}$, then λ is an eigenvalue of A if it is an eigenvalue of B . If x is an eigenvector of A , $C^{-1}x$ is an eigenvector of B .

Proof: Suppose x is an eigenvector of A , $Ax = \lambda x$. Let $y = C^{-1}x$. Since $A = CBC^{-1}$,

$$\lambda x = Ax = CBC^{-1}x = CB y.$$

Premultiplying by C^{-1} ,

$$\lambda y = \lambda C^{-1}x = C^{-1}CB y = B y.$$

■

88 Proposition If A and B are similar, then $\text{rank } A = \text{rank } B$.

Proof: We prove $\text{rank } B \geq \text{rank } A$. Symmetry completes the argument. Let z_1, \dots, z_k be a basis for $\text{range } A$, $z_i = Ay_i$. Put $w_i = C^{-1}y_i$. Then the Bw_i s are independent. To see this suppose

$$0 = \sum_{i=1}^k \alpha_i (Bw_i) = \sum_{i=1}^k \alpha_i C^{-1} A C C^{-1} y_i = \sum_{i=1}^k \alpha_i C^{-1} z_i = C^{-1} \left(\sum_{i=1}^k \alpha_i z_i \right).$$

Since C^{-1} is nonsingular, this implies $\sum_{i=1}^k \alpha_i z_i = 0$, which in turn implies $\alpha_i = 0$, $i = 1, \dots, k$. ■

9.7 The Principal Axis Theorem

The next result describes the diagonalization of a symmetric matrix.

89 Definition A square matrix X is **orthogonal** if $X'X = I$, or equivalently $X' = X^{-1}$.

90 Principal Axis Theorem Let $A: \mathbf{R}^m \rightarrow \mathbf{R}^m$ be a symmetric matrix. Let x_1, \dots, x_m be an orthonormal basis for \mathbf{R}^m made up of eigenvectors of A , with corresponding eigenvalues $\lambda_1, \dots, \lambda_m$. Set

$$\Lambda = \begin{bmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_m \end{bmatrix},$$

and set $X = [x_1, \dots, x_m]$.

Then

$$A = X \Lambda X^{-1},$$

$$\Lambda = X^{-1} A X,$$

and X is orthogonal, that is,

$$X^{-1} = X'.$$

Proof: Now $X'X = I$ by orthonormality, so $X^{-1} = X'$. Pick any z and set $y = X^{-1}z$, so $z = Xy = \sum_{j=1}^m \eta_j x_j$. Then

$$\begin{aligned} Az &= \sum_{j=1}^m \eta_j A x_j = \sum_{j=1}^m \eta_j (\lambda_j x_j) \\ &= X \Lambda y \\ &= X \Lambda X^{-1} z. \end{aligned}$$

Since z is arbitrary $A = X \Lambda X^{-1}$. ■

This result is called the principal axis theorem because in the case where A is positive definite (see Definition 101 below), the columns of X are the principal axes of the ellipsoid $\{x : x \cdot Ax = 1\}$. See Franklin [12, § 4.6, pp. 80–83].

9.8 Simultaneous diagonalization

91 Theorem (Simultaneous Diagonalization) *Let A, B be symmetric $m \times m$ matrices.*

Then $AB = BA$ if and only if there exists an orthonormal basis consisting of vectors that are eigenvectors of both A and B . Then letting X be the orthogonal matrix whose columns are the basis we have

$$\begin{aligned} A &= X\Lambda_A X^{-1} \\ B &= X\Lambda_B X^{-1}, \end{aligned}$$

where Λ_A and Λ_B are diagonal matrices of eigenvalues of A and B respectively.

Partial proof: (\Leftarrow)

$$AB = X\Lambda_A X^{-1} X\Lambda_B X^{-1} = X\Lambda_A \Lambda_B X^{-1} = X\Lambda_B \Lambda_A X^{-1} = X\Lambda_B X^{-1} X\Lambda_A X = BA,$$

since diagonal matrices commute.

(\Rightarrow) We shall prove the result for the special case where A has distinct eigenvalues. In this case, the eigenvectors associated with any eigenvalue are distinct up to scalar multiplication.

Let x be an eigenvector of A corresponding to eigenvalue λ . Suppose A and B commute. Then

$$A(Bx) = BAx = B(\lambda x) = \lambda(Bx).$$

This means that Bx too is an eigenvector of A corresponding to λ , provided $Bx \neq 0$. But as remarked above, this implies that Bx is a scalar multiple of x , so x is an eigenvector of B too. So let X be a matrix whose columns are a basis of orthonormal eigenvectors for both A and B . Then it follows that $A = X\Lambda_A X^{-1}$, where Λ_A is the diagonal matrix of eigenvalues, and similarly $B = X\Lambda_B X^{-1}$.

We now present a sketch of the proof for the general case. The crux of the proof is that in general, the eigenspace M associated with λ may be more than one-dimensional, so it is harder to conclude that the eigenvectors of A and B are the same. To get around this problem observe that $B^2A = BBA = BAB = ABB = AB^2$, and in general, $B^nA = AB^n$, so that if $Ax = \lambda x$, then $A(B^n x) = B^n Ax = B^n \lambda x = \lambda(B^n x)$. That is, for every n , the vector $B^n x$ is also an eigenvector of A corresponding to λ . Since the eigenspace M associated with λ is finite-dimensional, for some minimal k , the vectors $x, Bx, \dots, B^k x$ are dependent. That is, there are $\alpha_0, \dots, \alpha_k$, not all zero, with

$$\alpha_0 x + \alpha_1 Bx + \alpha_2 B^2 x + \dots + \alpha_k B^k x = 0.$$

Let μ_1, \dots, μ_k be the roots of the polynomial $\alpha_0 + \alpha_1 y + \alpha_2 y^2 + \dots + \alpha_k y^k$. Then

$$[(B - \mu_1 I)(B - \mu_2 I) \dots (B - \mu_k I)]x = 0.$$

$$\text{Set } z = [(B - \mu_2 I)(B - \mu_3 I) \dots (B - \mu_k I)]x$$

If k is minimal, then $z \neq 0$. (Even if coefficients are complex. Independence over real field implies independence over complex field. Just look at real and imaginary parts.) Therefore $(B - \mu_1 I)z = 0$

Claim: μ_1 and z are real.

Proof of claim: Let $\mu_1 = \alpha + i\beta$ and $z = x + iy$

$$B(x + iy) = (\alpha + i\beta)(x + iy)$$

Therefore $Bx = \alpha x - \beta y$ and $By = \beta x + \alpha y$ (equate real and imaginary parts). Thus $y'Bx = \alpha y'x - \beta y'y$ and by symmetry $y'Bx = x'By = \beta x'x + \alpha x'y$.

Now,

$$\alpha(y'x) - \beta(y'y) = \beta(x'x) + \alpha(x'y),$$

so $\beta = 0$ or $z = 0$, i.e., $x = 0$ and $y = 0$. But $z \neq 0$, so $\beta = 0$ and μ_1 is real. Similarly each μ_i is real. □

Therefore z is a real linear combination of $B^n x$ (all eigenvectors of A) satisfying $(B - \mu_1 I)z = 0$. In other words, $Bz = \mu_1 z$, so z is an eigenvector of both B and A !

We now consider the orthogonal complement of z in the eigenspace M to recursively construct a basis for M composed of eigenvectors for both A and B .

We must do this for each eigenvalue λ of A . More details are in Rao [21, Result (iii), pp. 41–42]. ■

9.9 Trace

92 Definition Let A be an $m \times m$ matrix. The **trace** of A , denoted $\text{tr } A$, is defined by

$$\text{tr } A = \sum_{i=1}^m \alpha_{ii}.$$

The trace is a linear functional on the set of $m \times m$ matrices.

93 Lemma Let A and B be $m \times m$ matrices. Then

$$\text{tr}(\alpha A + \beta B) = \alpha \text{tr } A + \beta \text{tr } B \tag{11}$$

$$\text{tr}(AB) = \text{tr}(BA). \tag{12}$$

Proof: The proof of linearity (11) is straightforward. For the proof of (12), observe that

$$\text{tr } AB = \sum_{i=1}^m \left(\sum_{j=1}^m \alpha_{i,j} \beta_{j,i} \right) = \sum_{j=1}^m \left(\sum_{i=1}^m \beta_{j,i} \alpha_{i,j} \right) = \text{tr } BA.$$

■

Equation (11) says that the trace is a linear functional on the vector space $M(m, m)$ of $m \times m$ matrices. The next result says that up to a scale multiple it is the only linear functional to satisfy (12).

94 Proposition If ℓ is a linear functional on $M(m, m)$ satisfying

$$\ell(AB) = \ell(BA) \quad \text{for all } A, B \in M(m, m),$$

then there is a constant α such that for all $A \in M(m, m)$,

$$\ell(A) = \alpha \text{tr } A.$$

Proof: To show uniqueness, let $\ell: M(m, m) \rightarrow \mathbf{R}$ be a linear functional satisfying

$$\ell(AB) = \ell(BA) \quad \text{for all } A, B \in M(m, m).$$

Now ℓ belongs to the space $L(M(m, m), \mathbf{R})$. The ordered basis on V induces a matrix representation for ℓ as an $mm \times 1$ matrix, call it L so that

$$\ell(A) = \sum_{i=1}^m \sum_{j=1}^m L_{ij} A_{ij}.$$

We also know that $\ell(AB - BA) = 0$ for every $A, B \in M(m, m)$. That is,

$$\sum_{i=1}^m \sum_{j=1}^m L_{ij} (AB - BA)_{ij} = \sum_{i=1}^m \sum_{j=1}^m \sum_{k=1}^m L_{ij} \{A_{ik} B_{kj} - B_{ik} A_{kj}\} = 0. \tag{13}$$

Now consider the matrix A with all its entries zero except for the i^{th} row, which is all ones, and B , which has all its entries zero, except for the j^{th} column, which is all ones. If $i \neq j$, then the i, j entry of AB is m and the rest are zero, whereas all the entries of BA are zero. In this case, (13) implies $L_{ij} = 0$.

Next consider the matrices A and B where the nonzero entries A are rows i and j , which consist of ones; and the nonzero entries of B are column i , which consists of ones, and column j , which is column of minus ones. Then $BA = 0$ and AB is zero except for $(AB)_{ii} = (AB)_{ji} = m$ and $(AB)_{ij} = (AB)_{jj} = -m$. Since $L_{ij} = 0$ whenever $i \neq j$, equation (13) reduces to $L_{ii}m + L_{jj}(-m) = 0$, which implies $L_{ii} = L_{jj}$.

Let α be the common value of the elements of the diagonal matrix L . We have just shown that $\ell(A) = \alpha \operatorname{tr} A$. ■

The trace of a matrix depends only on the linear transformation of \mathbf{R}^m into \mathbf{R}^m that it represents.

95 Proposition *The trace of a matrix depends only on the linear transformation of \mathbf{R}^m into \mathbf{R}^m that it represents. In other words, if A and B are similar matrices, that is, if $B = C^{-1}AC$, then $\operatorname{tr} B = \operatorname{tr} A$.*

Proof: Theorem 86 asserts that two matrices represent the same transformation if and only if they are similar. By Lemma 93, equation (12),

$$\operatorname{tr} B = \operatorname{tr} C^{-1}AC = \operatorname{tr} ACC^{-1} = \operatorname{tr} A. \quad \blacksquare$$

96 Corollary *Let V be an m -dimensional inner product space. There is a unique linear functional tr on $L(V, V)$ satisfying*

$$\operatorname{tr} ST = \operatorname{tr} TS \quad \text{for all } S, T \in L(V, V),$$

and

$$\operatorname{tr} I = m.$$

Add: The trace is also equal to the sum of the characteristic roots (eigenvalues).

97 Theorem *If A is symmetric and idempotent, then $\text{tr } A = \text{rank } A$.*

Proof: Since A is symmetric, $A = XBX^{-1}$ where $X = [x_1, \dots, x_m]$ is an orthogonal matrix whose columns are eigenvectors of A , and B is a diagonal matrix whose diagonal elements are the eigenvalues of A , which are either 0 or 1.

Thus $\text{tr } B$ is the number of nonzero eigenvalues of A . Also $\text{rank } B$ is the number of nonzero diagonal elements. Thus $\text{tr } B = \text{rank } B$, but since A and B are similar, $\text{tr } A = \text{tr } B = \text{rank } B = \text{rank } A$. ■

It also follows that on the space of symmetric matrices, the trace can be used to define an inner product.

98 Proposition *Let X denote the linear space of symmetric $m \times m$ real matrices. Then*

$$(A, B) = \text{tr } AB$$

is an inner product.

Proof: Lemma 93, equation (12) shows that $(A, B) = (B, A)$ so **IP.1** is satisfied. Moreover equation (11) implies

$$\begin{aligned} (A, B + C) &= \text{tr } A(B + C) = \text{tr}(AB + AC) = \text{tr } AB + \text{tr } AC = (A, B) + (A, C), \\ (\alpha A, B) &= \text{tr } \alpha AB = \alpha \text{tr } AB = \alpha(A, B) \\ (A, \alpha B) &= \text{tr } A\alpha B = \alpha \text{tr } AB = \alpha(A, B) \end{aligned}$$

so **IP.2** and **IP.3** are satisfied.

To see that **IP.4** is satisfied, observe that if A is a symmetric matrix, then $(AA)_{ii}$ is just the inner product of the i^{th} row of A with itself, which is ≥ 0 and equals zero only if the row is zero. Thus $(A, A) = \text{tr } AA = \sum_i (AA)_{ii} \geq 0$ and $= 0$ only if every row of A is zero, that is if A itself is zero. ■

99 Exercise What does it mean for symmetric matrices A and B to be orthogonal under this inner product? □

9.10 Matrices and orthogonal projection

Section 4.3 discussed orthogonal projection as a linear transformation. In this section we discuss matrix representations for orthogonal projection.

Let M be a k -dimensional subspace of \mathbf{R}^m , and let b_1, \dots, b_k be a basis for M . Given a vector y , the orthogonal projection y_M of y is the vector in M that minimizes the distance to y (Proposition 34). The difference $y_{\perp} = y - y_M$ is orthogonal to M (Theorem 31).

The next result is crucial to the statistical analysis of linear regression models.

100 Least squares regression *Let B the $m \times k$ matrix with columns b_1, \dots, b_k , where $\{b_1, \dots, b_k\}$ is a basis for the k -dimensional subspace M . Given a vector y in \mathbf{R}^m , the orthogonal projection y_M of y onto M satisfies*

$$y_M = B(B'B)^{-1}B'y.$$

Proof: The first thing to note is that since $\{b_1, \dots, b_k\}$ is a basis for M , the matrix B has rank k , so by Corollary 54, the $k \times k$ matrix $B'B$ has rank k , so by Proposition 78 it is invertible.

Next note that the $m \times 1$ column vector $B(B'B)^{-1}B'y$ belongs to M . In fact, setting

$$a = (B'B)^{-1}B'y$$

(a $k \times 1$ column matrix), we see that it is the linear combination

$$Ba = \sum_{j=1}^k \alpha_j b_j$$

of basis vectors. Thus by the Orthogonal Projection Theorem 31, to show that $y_M = Ba$, it suffices to show that $y - Ba$ is orthogonal to M . This in turn is equivalent to $y - Ba$ being orthogonal to each basis vector b_j . Now for any x , the $k \times 1$ column matrix $B'x$ has as its j^{th} (row) entry the dot product $b_j \cdot x$. Thus all we need do is show that $B'(y - Ba) = 0$. To this end, compute

$$B'(y - Ba) = B'y - B'B(B'B)^{-1}B'y = B'y - B'y = 0.$$

■

A perhaps more familiar way to restate this result is that the vector a of coefficients that minimizes the sum of squared residuals, $(y - Ba) \cdot (y - Ba)$, is given by $a = (B'B)^{-1}B'y$.

10 Quadratic forms

We introduced quadratic forms in the proof of Theorem 70. We go a little deeper here. If you want to know even more, I recommend my on-line notes [9].

Let A be an $n \times n$ symmetric matrix, and let x be an n -vector. Then $x \cdot Ax$ is a scalar, and

$$x \cdot Ax = \sum_{i=1}^n \sum_{j=1}^n \alpha_{ij} \xi_i \xi_j. \tag{14}$$

(We may also write this as $x'Ax$ in matrix notation.)

The mapping $Q: x \mapsto x \cdot Ax$ is the **quadratic form** defined by A .⁶

101 Definition A symmetric matrix A (or its associated quadratic form) is called

- **positive definite** if $x \cdot Ax > 0$ for all nonzero x .

⁶For decades I was baffled by the term *form*. I once asked Tom Apostol at a faculty cocktail party what it meant. He professed not to know (it was a cocktail party, so that is excusable), but suggested that I should ask John Todd. He hypothesized that mathematicians don't know the difference between form and function, a clever reference to modern architectural philosophy. I was too intimidated by Todd to ask, but I subsequently learned (where, I can't recall) that *form* refers to a polynomial function in several variables where each term in the polynomial has the same degree. (The *degree* of the term is the sum of the exponents. For example, in the expression $xyz + x^2y + xz + z$, the first two terms have degree three, the third term has degree two and the last one has degree one. It is thus not a form.) This is most often encountered in the phrases *linear form* (each term has degree one) or *quadratic form* (each term has degree two).

- **negative definite** if $x \cdot Ax < 0$ for all nonzero x .
- **positive semidefinite** if $x \cdot Ax \geq 0$ for all x .
- **negative semidefinite** if $x \cdot Ax \leq 0$ for all x .

We want all our (semi)definite matrices to be symmetric so that their eigenvectors generate an orthonormal basis for \mathbf{R}^n . If A is not symmetric, then $\frac{A+A'}{2}$ is symmetric and $x \cdot Ax = x \cdot (\frac{A+A'}{2})x$ for any x , so we do not lose much applicability by this assumption. Some authors use the term **quasi (semi)definite** when they do not wish to impose symmetry.

10.1 Diagonalization of quadratic forms

By the Principal Axis Theorem 90 we may write

$$A = X\Lambda X',$$

where X is an orthogonal matrix with columns that are eigenvectors of A , and Λ is a diagonal matrix of eigenvalues of A . Then the quadratic form can be written in terms of the diagonal matrix Λ :

$$x \cdot Ax = x'Ax = x'X\Lambda X'x = y'\Lambda y = \sum_{i=1}^n \lambda_i y_i^2,$$

where

$$y = X'x.$$

102 Proposition (Eigenvalues and definiteness) *The symmetric matrix A is*

1. *positive definite if and only if all its eigenvalues are strictly positive.*
2. *negative definite if and only if all its eigenvalues are strictly negative.*
3. *positive semidefinite if and only if all its eigenvalues are nonnegative.*
4. *negative semidefinite if and only if all its eigenvalues are nonpositive.*

Proof: As above, write

$$x'Ax = \sum_{i=1}^n \lambda_i y_i^2.$$

where the λ_i s are the eigenvalues of A . All the statements above follow from this equation and the fact that $y_i^2 \geq 0$ for all k . ■

103 Proposition (Definiteness of the inverse) *If A is positive definite (negative definite), then A^{-1} exists and is also positive definite (negative definite).*

Proof: First off, how do we know the inverse of A exists? Suppose $Ax = 0$. Then $x \cdot Ax = x \cdot 0 = 0$. Since A is positive definite, we see that $x = 0$. Therefore A is invertible. Here are two proofs of the proposition.

First proof. Since $(Ax = \lambda x) \implies (x = \lambda A^{-1}x) \implies (A^{-1}x = \frac{1}{\lambda}x)$, the eigenvalues of A and A^{-1} are reciprocals, so they must have the same sign. Apply Proposition 102.

Second proof.

$$x \cdot A^{-1}x = y \cdot Ay \quad \text{where} \quad y = A^{-1}x.$$

■

11 Determinants

The main quick references here are Apostol [6, Chapter 3] and Dieudonné [11, Appendix A.6]. The main things to remember are:

- The determinant assigns a number to each square matrix A , denoted either $\det A$ or $|A|$. A matrix is **singular** if its determinant is zero, otherwise it is **nonsingular**. For an $n \times n$ identity matrix, $\det I = 1$.
- $\det(AB) = \det A \cdot \det B$.
- Thus a matrix has an inverse if and only if its determinant is nonzero.
- Multiplying a row or a column by a scalar multiplies the determinant by the same amount. Consequently for an $n \times n$ matrix A ,

$$\det(-A) = (-1)^n \det A.$$

- Also consequently, the determinant of a diagonal matrix is the product of its diagonal elements.
- Adding a multiple of one row to another does not change the determinant.
- Consequently, the determinant of an upper (or lower) triangular matrix is the product of its diagonal.
- Moreover if a square matrix A is block upper triangular, that is, of the form

$$A = \begin{bmatrix} B & C \\ 0 & D \end{bmatrix},$$

where B and D are square, then $\det A = \det B \cdot \det D$. Likewise for block lower triangular matrices.

- $\det A = \det A'$.
- The determinant can be defined recursively in terms of minors (determinants of submatrices).

- The inverse of a matrix can be computed in terms of these minors. The inverse of A is the transpose of its cofactor matrix divided by $\det A$.
- Cramer's Rule: If $Ax = b$ for a nonsingular matrix A , then

$$x_i = \frac{\det(A^1, \dots, A^{i-1}, b, A^{i+1}, \dots, A^n)}{\det A}.$$

- The determinant is the oriented volume of the n -dimensional "cube" formed by its columns.
- The determinant $\det(\lambda I - A)$, where A and I are $n \times n$ is an n^{th} degree polynomial in λ , called the **characteristic polynomial** of A .
- A root (real or complex) of the characteristic polynomial of A is called a **characteristic root** of A . Characteristic roots that are real are also eigenvalues. If nonzero x belongs to the null space of $\lambda I - A$, then it is an eigenvector corresponding to the eigenvalue λ .
- The determinant of A is the product of its characteristic roots.
- If A is symmetric, then $\det A$ is the product of its eigenvalues.
- If A has rank k then every minor of size greater than k has zero determinant and there is at least one minor of order k with nonzero determinant.
- The determinant of an orthogonal matrix is ± 1 .

11.1 Determinants as multilinear forms

There are several ways to think about determinants. Perhaps the most useful is as an alternating multilinear n -form:

A function $\varphi: \underbrace{\mathbf{R}^n \times \dots \times \mathbf{R}^n}_{n \text{ copies}} \rightarrow \mathbf{R}$ is **multilinear** if it is linear in each variable separately.

That is, for each $i = 1, \dots, n$,

$$\begin{aligned} \varphi(x_1, \dots, x_{i-1}, \alpha x_i + \beta y_i, x_{i+1}, \dots, x_n) \\ = \alpha \varphi(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n) + \beta \varphi(x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_n). \end{aligned}$$

A consequence of this is that if any x_i is zero, then so is φ . That is,

$$\varphi(x_1, \dots, x_{i-1}, 0, x_{i+1}, \dots, x_n) = 0.$$

The multilinear function φ is **alternating** if $x_i = x_j = z$ for distinct i and j , then

$$\varphi(x_1, \dots, z, \dots, z, \dots, x_n) = 0.$$

The reason for this terminology is the following lemma.

104 Lemma *The multilinear function φ is alternating if and only if interchanging x_i and x_j changes the sign of φ , that is,*

$$\varphi(x_1, \dots, x_i, \dots, x_j, \dots, x_n) = -\varphi(x_1, \dots, x_j, \dots, x_i, \dots, x_n).$$

Proof: Suppose first that φ is alternating. Then

$$\begin{aligned} 0 &= \varphi(x_1, \dots, x_i + x_j, \dots, x_j + x_i, \dots, x_n) \\ &= \varphi(x_1, \dots, x_i, \dots, x_j + x_i, \dots, x_n) + \varphi(x_1, \dots, x_j, \dots, x_j + x_i, \dots, x_n) \\ &= \varphi(x_1, \dots, x_i, \dots, x_j, \dots, x_n) + \underbrace{\varphi(x_1, \dots, x_i, \dots, x_i, \dots, x_n)}_{=0} \\ &\quad + \underbrace{\varphi(x_1, \dots, x_j, \dots, x_j, \dots, x_n)}_{=0} + \varphi(x_1, \dots, x_j, \dots, x_i, \dots, x_n) \\ &= \varphi(x_1, \dots, x_i, \dots, x_j, \dots, x_n) + \varphi(x_1, \dots, x_j, \dots, x_i, \dots, x_n). \end{aligned}$$

So $\varphi(x_1, \dots, x_i, \dots, x_j, \dots, x_n) = -\varphi(x_1, \dots, x_j, \dots, x_i, \dots, x_n)$.

Now suppose interchanging x_i and x_j changes the sign of φ . Then if $x_i = x_j = z$,

$$\varphi(x_1, \dots, z, \dots, z, \dots, x_n) = -\varphi(x_1, \dots, z, \dots, z, \dots, x_n),$$

which implies $\varphi(x_1, \dots, z, \dots, z, \dots, x_n) = 0$, so φ is alternating. ■

There is an obvious identification of $n \times n$ square matrices with the elements of $\mathbf{R}^n \times \dots \times \mathbf{R}^n$. Namely $A \leftrightarrow (A^1, \dots, A^n)$, where you will recall A^j denotes the j^{th} column of A interpreted as a vector in \mathbf{R}^n . (We could have used rows just as well.) Henceforth, for a multilinear form φ and $n \times n$ matrix A , we write $\varphi(A)$ for $\varphi(A^1, \dots, A^n)$. The next fact is rather remarkable, so pay close attention.

105 Proposition For every $n \times n$ matrix A , there is a number $\det A$ with the property that for any alternating multilinear n -form φ ,

$$\varphi(A) = \det A \cdot \varphi(I). \tag{*}$$

Proof: Before we demonstrate the proposition in general, let us start with a special case, $n = 2$. Let

$$A = \begin{bmatrix} \alpha_{1,1} & \alpha_{1,2} \\ \alpha_{2,1} & \alpha_{2,2} \end{bmatrix}.$$

Then

$$\begin{aligned} \varphi(A) &= \varphi(A^1, A^2) \\ &= \varphi(\alpha_{1,1}e_1 + \alpha_{2,1}e_2, \alpha_{1,2}e_1 + \alpha_{2,2}e_2) \\ &= \alpha_{1,1}\varphi(e_1, \alpha_{1,2}e_1 + \alpha_{2,2}e_2) + \alpha_{2,1}\varphi(e_2, \alpha_{1,2}e_1 + \alpha_{2,2}e_2) \\ &= \alpha_{1,1}\alpha_{1,2}\varphi(e_1, e_1) + \alpha_{1,1}\alpha_{2,2}\varphi(e_1, e_2) + \alpha_{2,1}\alpha_{1,2}\varphi(e_2, e_1) + \alpha_{2,1}\alpha_{2,2}\varphi(e_2, e_2) \\ &= 0 + \alpha_{1,1}\alpha_{2,2}\varphi(e_1, e_2) + \alpha_{2,1}\alpha_{1,2}\varphi(e_2, e_1) + 0 \\ &= \alpha_{1,1}\alpha_{2,2}\varphi(e_1, e_2) - \alpha_{2,1}\alpha_{1,2}\varphi(e_1, e_2) \\ &= (\alpha_{1,1}\alpha_{2,2} - \alpha_{2,1}\alpha_{1,2})\varphi(I), \end{aligned}$$

so we see that $\det A = \alpha_{1,1}\alpha_{2,2} - \alpha_{2,1}\alpha_{1,2}$. Thus the whole of φ is determined by the single number $\varphi(I)$.

This is true more generally. Write

$$\varphi(A) = \varphi \left(\sum_{i_1=1}^n \alpha_{i_1,1} e_{i_1}, \sum_{i_2=1}^n \alpha_{i_2,2} e_{i_2}, \dots, \sum_{i_j=1}^n \alpha_{i_j,j} e_{i_j}, \dots, \sum_{i_n=1}^n \alpha_{i_n,n} e_{i_n} \right).$$

Now expand this using linearity in the first component:

$$\varphi(A) = \sum_{i_1=1}^n \alpha_{i_1,1} \varphi \left(e_{i_1}, \sum_{i_2=1}^n \alpha_{i_2,2} e_{i_2}, \dots, \sum_{i_j=1}^n \alpha_{i_j,j} e_{i_j}, \dots, \sum_{i_n=1}^n \alpha_{i_n,n} e_{i_n} \right).$$

Repeating this for the other components leads to

$$\varphi(A) = \sum_{i_1=1}^n \sum_{i_2=1}^n \dots \sum_{i_n=1}^n \alpha_{i_1,1} \alpha_{i_2,2} \dots \alpha_{i_n,n} \varphi(e_{i_1}, e_{i_2}, \dots, e_{i_n}).$$

Now consider $\varphi(e_{i_1}, e_{i_2}, \dots, e_{i_n})$. Since φ is alternating, this term is zero unless i_1, i_2, \dots, i_n are distinct. When these are distinct, then by switching pairs we get to $\pm \varphi(e_1, e_2, \dots, e_n)$ where the sign on whether we need an odd or an even number of switches. It now pays to introduce some new terminology and notation. A **permutation** \mathbf{i} is an ordered list $\mathbf{i} = (i_1, \dots, i_n)$ of the numbers $1, 2, \dots, n$. The **signature** $\text{sgn}(\mathbf{i})$ of \mathbf{i} is 1 if \mathbf{i} can be put in numerical order by switching terms an even number of times and is -1 if it requires an odd number. It follows then that defining

How do we know that $\text{sgn}(\mathbf{i})$ is well defined?

$$\det A = \sum_{\mathbf{i}} \text{sgn}(\mathbf{i}) \cdot \alpha_{i_1,1} \alpha_{i_2,2} \dots \alpha_{i_n,n}, \tag{15}$$

where the sum runs over all permutations \mathbf{i} , satisfies the conclusion of the proposition. ■

This result still leaves the following question: *Are there any alternating multilinear n -forms at all?* The reason the result above does not settle this is that it would be vacuously true if there were none. Fortunately, it is not hard to verify that \det itself is such an n -form.

106 Proposition *The function $A \mapsto \det A$ as defined in (15) is an alternating multilinear n -form.*

Proof: Observe that in each product $\text{sgn}(\mathbf{i}) \cdot \alpha_{i_1,1} \alpha_{i_2,2} \dots \alpha_{i_n,n}$ in the sum in (15) there is exactly one element from each row and each column of A . This makes it obvious that $\det(A^1, \dots, \alpha A^j, \dots, A^n) = \alpha \det A$, and it is straightforward to verify that

$$\det(A^1, \dots, x + y, \dots, A^n) = \det(A^1, \dots, x, \dots, A^n) + \det(A^1, \dots, y, \dots, A^n).$$

To see that \det is alternating, suppose $A^i = A^j$ with $i \neq j$. Then for any permutation \mathbf{i} , $i_j \neq i_k$, and there is exactly one permutation \mathbf{i}' satisfying $i_p = i'_p$ for $p \notin \{i, j\}$ and $i'_j = i_k$ and $i'_k = i_j$. Now observe that $\text{sgn}(\mathbf{i}) = -\text{sgn}(\mathbf{i}')$ as it requires an odd number of interchanges to swap two elements in a list (why?). Thus we can rewrite (15) as

$$\det A = \sum_{\mathbf{i}:\text{sgn}(\mathbf{i})=1} \alpha_{i_1,1} \alpha_{i_2,2} \dots \alpha_{i_n,n} - \alpha_{i'_1,1} \alpha_{i'_2,2} \dots \alpha_{i'_n,n},$$

but each $\alpha_{i_1,1} \alpha_{i_2,2} \dots \alpha_{i_n,n} - \alpha_{i'_1,1} \alpha_{i'_2,2} \dots \alpha_{i'_n,n} = 0$. (Why?) Therefore $\det A = 0$, so \det is alternating. ■

To sum things up we have:

107 Corollary *An alternating multilinear n -form is identically zero if and only if $\varphi(I) = 0$. The determinant is the unique alternating multilinear n -form φ that satisfies $\varphi(I) = 1$. Any other alternating multilinear n -form φ is of the form $\varphi = \varphi(I) \cdot \det$.*

11.2 Some simple consequences

108 Proposition *If A' is the transpose of A , then $\det A = \det A'$.*

Write out a proof.

109 Proposition *Adding a scalar multiple of one column of A to a different column leaves the determinant unchanged. Likewise for rows.*

Proof: By multilinearity, $\det(A^1, \dots, A^j + \alpha A^k, \dots, A^k, \dots, A^n) = \det(A^1, \dots, A^j, \dots, A^k, \dots, A^n) + \alpha \det(A^1, \dots, A^k, \dots, A^k, \dots, A^n)$, but $\det(A^1, \dots, A^k, \dots, A^k, \dots, A^n) = 0$ since \det is alternating. The conclusion for rows follows from that for columns and Proposition 108 on transposes. ■

110 Proposition *The determinant of an upper triangular matrix is the product of the diagonal entries.*

Proof: Recall that an upper triangular matrix is one for which $i > j$ implies $\alpha_{i,j} = 0$. (Diagonal matrices are also upper triangular.) Now examine equation (15). The only summand that is nonzero comes from the permutation $(1, 2, 3, \dots, n)$, since for any other permutation there is some j satisfying $i_j > j$. (Why?) ■

By the way, the result also holds for lower triangular matrices.

111 Lemma *If A is $n \times n$ and φ is an alternating multilinear n -form, so is the form φ_A defined by*

$$\varphi_A(x_1, \dots, x_n) = \varphi(Ax_1, \dots, Ax_n).$$

Furthermore

$$\varphi_A(I) = \varphi(A) = \det A \cdot \varphi(I). \tag{16}$$

Proof: That φ_A is an alternating multilinear n -form is straightforward. Therefore φ_A is proportional to φ . To see that the coefficient of proportionality is $\det A$, consider $\varphi_A(I)$. Direct computation shows that $\varphi_A(I) = \varphi(A)$. ■

As an aside, I mention that we could have defined the determinant directly for linear transformations as follows. For a linear transformation $T: \mathbf{R}^n \rightarrow \mathbf{R}^n$ it follows that φ_T defined by $\varphi_T(x_1, \dots, x_n) = \varphi(Tx_1, \dots, Tx_n)$ is an alternating n -form whenever φ is. It also follows that we could use (16) to define $\det T$ to be the scalar satisfying $\varphi_T(x_1, \dots, x_n) = \det T \cdot \varphi(x_1, \dots, x_n)$. This is precisely Dieudonné's approach. It has the drawback that minors and cofactors are awkward to describe in his framework.

112 Theorem *Let A and B be $n \times n$ matrices. Then*

$$\det AB = \det A \cdot \det B.$$

Proof: Let φ be an alternating multilinear n -form. Applying (16) and (\star) , we see

$$\varphi_{AB}(I) = \varphi(AB) = \det AB \cdot \varphi(I).$$

On the other hand

$$\varphi_{AB}(I) = \varphi_A(B) = \det B \cdot \varphi_A(I) = \det B \cdot \det A \cdot \varphi(I).$$

Therefore $\det AB = \det A \cdot \det B$. ■

113 Corollary *If $\det A = 0$, then A has no inverse.*

Proof: Observe that if A has an inverse, then

$$1 = \det I = \det A \cdot \det A^{-1}$$

so $\det A \neq 0$. ■

114 Corollary *The determinant of an orthogonal matrix is ± 1 .*

Proof: Recall that A is orthogonal if $A'A = I$ (Definition 89). By Theorem 112, $\det(A') \det(A) = 1$, but by Proposition 108, $\det(A') = \det A$, so $(\det A)^2 = 1$. ■

11.3 Minors and cofactors

Different authors assign different meanings to the term **minor**. Given a square $n \times n$ matrix

$$A = \begin{bmatrix} \alpha_{1,1} & \cdots & \alpha_{1,n} \\ \vdots & & \vdots \\ \alpha_{n,1} & \cdots & \alpha_{n,n} \end{bmatrix}$$

Apostol [6, p. 87] defines the i, j **minor** of A to be the $(n-1) \times (n-1)$ *submatrix* obtained from A by deleting the i^{th} row and j^{th} column, and denotes it $A_{i,j}$. Gantmacher [14, p. 2] defines a **minor** of a (not necessarily square) matrix A to be the *determinant* of a square submatrix of A , and uses the following notation for minors in terms of the remaining rows and columns:

$$A \left(\begin{matrix} i_1, \dots, i_p \\ j_1, \dots, j_p \end{matrix} \right) = \begin{vmatrix} \alpha_{i_1, j_1} & \cdots & \alpha_{i_1, j_p} \\ \vdots & & \vdots \\ \alpha_{i_p, j_1} & \cdots & \alpha_{i_p, j_p} \end{vmatrix}.$$

Here we require that $1 \leq i_1 < i_2 < \cdots < i_p \leq n$ and $1 \leq j_1 < j_2 < \cdots < j_p \leq n$. If the deleted (and hence remaining) rows and columns are the same, that is, if $i_1 = j_1, i_2 = j_2, \dots, i_p = j_p$, then $A \left(\begin{matrix} i_1, \dots, i_p \\ j_1, \dots, j_p \end{matrix} \right)$ is called a **principal minor of order p** . I think the following hybrid terminology is useful: a **minor submatrix** is any square submatrix of A (regardless of whether A is square) and a **minor** of A is the determinant of a minor submatrix (same as

Gantmacher). To be on the safe side, I may use the redundant term **minor determinant** to mean minor.

Given a square $n \times n$ matrix

$$A = \begin{bmatrix} \alpha_{1,1} & \cdots & \alpha_{1,n} \\ \vdots & & \vdots \\ \alpha_{n,1} & \cdots & \alpha_{n,n} \end{bmatrix}$$

the **cofactor** $\text{cof } \alpha_{i,j}$ of $\alpha_{i,j}$ is the determinant obtained by replacing the j^{th} column of A with the i^{th} unit coordinate vector e_i . That is,

$$\text{cof } \alpha_{i,j} = \det(A^1, \dots, A^{j-1}, e_i, A^{j+1}, \dots, A^n).$$

By multilinearity we have for any column j ,

$$\det A = \sum_{i=1}^n \alpha_{i,j} \text{cof } \alpha_{i,j}.$$

115 Lemma (Cofactors and minors) For any square matrix A ,

$$\text{cof } \alpha_{i,j} = (-1)^{i+j} \det A_{i,j},$$

where $A_{i,j}$ is the minor submatrix obtained by deleting the i^{th} row and j^{th} column from A . Consequently

$$\det A = \sum_{i=1}^n (-1)^{i+j} \alpha_{i,j} \det A_{i,j}.$$

Similarly (interchanging the roles of rows and columns), for any row i ,

$$\det A = \sum_{j=1}^n (-1)^{i+j} \alpha_{i,j} \det A_{i,j}.$$

Proof: Cf. Apostol [6, Theorem 3.9, p. 87]. By definition

$$\text{cof } \alpha_{i,j} = \begin{vmatrix} \alpha_{1,1} & \cdots & \alpha_{1,j-1} & 0 & \alpha_{1,j+1} & \cdots & \alpha_{1,n} \\ \vdots & & \vdots & \vdots & \vdots & & \vdots \\ \alpha_{i-1,1} & \cdots & \alpha_{i-1,j-1} & 0 & \alpha_{i-1,j+1} & \cdots & \alpha_{i-1,n} \\ \alpha_{i,1} & \cdots & \alpha_{i,j-1} & 1 & \alpha_{i,j+1} & \cdots & \alpha_{i,n} \\ \alpha_{i+1,1} & \cdots & \alpha_{i+1,j-1} & 0 & \alpha_{i+1,j+1} & \cdots & \alpha_{i+1,n} \\ \vdots & & \vdots & \vdots & \vdots & & \vdots \\ \alpha_{n,1} & \cdots & \alpha_{n,j-1} & 0 & \alpha_{n,j+1} & \cdots & \alpha_{n,n} \end{vmatrix}.$$

Adding $-\alpha_{i,k}e_i$ to column k does not change the determinant. Doing this for all $k \neq j$ yields

$$\text{cof } \alpha_{i,j} = \begin{vmatrix} \alpha_{1,1} & \dots & \alpha_{1,j-1} & 0 & \alpha_{1,j+1} & \dots & \alpha_{1,n} \\ \vdots & & \vdots & \vdots & \vdots & & \vdots \\ \alpha_{i-1,1} & \dots & \alpha_{i-1,j-1} & 0 & \alpha_{i-1,j+1} & \dots & \alpha_{i-1,n} \\ 0 & \dots & 0 & 1 & 0 & \dots & 0 \\ \alpha_{i+1,1} & \dots & \alpha_{i+1,j-1} & 0 & \alpha_{i+1,j+1} & \dots & \alpha_{i+1,n} \\ \vdots & & \vdots & \vdots & \vdots & & \vdots \\ \alpha_{n,1} & \dots & \alpha_{n,j-1} & 0 & \alpha_{n,j+1} & \dots & \alpha_{n,n} \end{vmatrix}.$$

Now by repeatedly interchanging columns a total of $j - 1$ times we obtain

$$\text{cof } \alpha_{i,j} = (-1)^{j-1} \begin{vmatrix} 0 & \alpha_{1,1} & \dots & \alpha_{1,j-1} & \alpha_{1,j+1} & \dots & \alpha_{1,n} \\ \vdots & & & \vdots & & & \vdots \\ 0 & \alpha_{i-1,1} & \dots & \alpha_{i-1,j-1} & \alpha_{i-1,j+1} & \dots & \alpha_{i-1,n} \\ 1 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & \alpha_{i+1,1} & \dots & \alpha_{i+1,j-1} & \alpha_{i+1,j+1} & \dots & \alpha_{i+1,n} \\ \vdots & & & \vdots & & & \vdots \\ 0 & \alpha_{n,1} & \dots & \alpha_{n,j-1} & \alpha_{n,j+1} & \dots & \alpha_{n,n} \end{vmatrix}.$$

Interchanging rows $i - 1$ times yields

$$\text{cof } \alpha_{i,j} = (-1)^{i+j} \begin{vmatrix} 1 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & \alpha_{1,1} & \dots & \alpha_{1,j-1} & \alpha_{1,j+1} & \dots & \alpha_{1,n} \\ \vdots & & & \vdots & & & \vdots \\ 0 & \alpha_{i-1,1} & \dots & \alpha_{i-1,j-1} & \alpha_{i-1,j+1} & \dots & \alpha_{i-1,n} \\ 0 & \alpha_{i+1,1} & \dots & \alpha_{i+1,j-1} & \alpha_{i+1,j+1} & \dots & \alpha_{i+1,n} \\ \vdots & & & \vdots & & & \vdots \\ 0 & \alpha_{n,1} & \dots & \alpha_{n,j-1} & \alpha_{n,j+1} & \dots & \alpha_{n,n} \end{vmatrix}.$$

(Recall that $(-1)^{j-1+i-1} = (-1)^{i+j}$.) This last determinant is block diagonal, so we see that it is just $|A_{i,j}|$, which completes the proof.

The conclusion for rows follows from that for columns and Proposition 108 on transposes. ■

By repeatedly applying this result, we can express $\det A$ in terms of 1×1 determinants.

If we take the cofactors from an *alien* column or row, we have:

116 Lemma (Expansion by alien cofactors) *Let A be a square $n \times n$ matrix. For any column j and any $k \neq j$,*

$$\sum_{i=1}^n \alpha_{i,j} \text{cof } \alpha_{i,k} = 0.$$

Likewise for any row i and any $k \neq i$,

$$\sum_{j=1}^n \alpha_{i,j} \text{cof } \alpha_{k,j} = 0.$$

Proof: Consider the matrix $A_1 = [\tilde{\alpha}_{i,j}]$ obtained from A by replacing the k^{th} column with another copy of column j . Then the i, k cofactors, $i = 1, \dots, n$, of A and A_1 are the same. (The cofactors don't depend on the column they belong to, since it is replaced by a unit coordinate vector.) So by Lemma 115 $|A_1| = \sum_{i=1}^n \tilde{\alpha}_{i,k} \text{cof } \alpha_{i,k} = \sum_{i=1}^n \alpha_{i,j} \text{cof } \alpha_{i,k}$. But $|A_1| = 0$ since it has two identical columns.

The conclusion for rows follows from that for columns and Proposition 108 on transposes. ■

The transpose of the cofactor matrix is also called the **adjugate** matrix of A . Combining the previous two lemmas yields the following on the adjugate.

117 Theorem (Cofactors and the inverse matrix) *For a square matrix*

$$A = \begin{bmatrix} \alpha_{1,1} & \dots & \alpha_{1,n} \\ \vdots & & \vdots \\ \alpha_{n,1} & \dots & \alpha_{n,n} \end{bmatrix},$$

we have

$$\begin{bmatrix} \alpha_{1,1} & \dots & \alpha_{1,n} \\ \vdots & & \vdots \\ \alpha_{n,1} & \dots & \alpha_{n,n} \end{bmatrix} \begin{bmatrix} \text{cof } \alpha_{1,1} & \dots & \text{cof } \alpha_{1,n} \\ \vdots & & \vdots \\ \text{cof } \alpha_{1,n} & \dots & \text{cof } \alpha_{n,n} \end{bmatrix} = \begin{bmatrix} |A| & & 0 \\ & \ddots & \\ 0 & & |A| \end{bmatrix}.$$

That is, the product of A and its adjugate is $(\det A)I_n$. Consequently, if $\det A \neq 0$, then

$$A^{-1} = \frac{1}{\det A} \begin{bmatrix} \text{cof } \alpha_{1,1} & \dots & \text{cof } \alpha_{1,n} \\ \vdots & & \vdots \\ \text{cof } \alpha_{1,n} & \dots & \text{cof } \alpha_{n,n} \end{bmatrix}.$$

Combining this with Corollary 113 yields the following.

118 Corollary (Determinants and invertibility) *A square matrix is invertible if and only if its determinant is nonzero.*

Similar reasoning leads to the following theorem due to Jacobi. It expresses a p^{th} order minor of the adjugate in terms of the corresponding complementary minor of A . The **complement** of the p^{th} -order minor $A \begin{pmatrix} i_1, \dots, i_p \\ j_1, \dots, j_p \end{pmatrix}$ is the $n - p^{\text{th}}$ -order minor obtained by deleting rows i_1, \dots, i_p and columns j_1, \dots, j_p from A .

119 Theorem (Jacobi) *For a square matrix*

$$A = \begin{bmatrix} \alpha_{1,1} & \dots & \alpha_{1,n} \\ \vdots & & \vdots \\ \alpha_{n,1} & \dots & \alpha_{n,n} \end{bmatrix},$$

we have for any $1 \leq p \leq n$,

$$|A| \cdot \begin{vmatrix} \text{cof } \alpha_{1,1} & \dots & \text{cof } \alpha_{p,1} \\ \vdots & & \vdots \\ \text{cof } \alpha_{1,p} & \dots & \text{cof } \alpha_{p,p} \end{vmatrix} = |A|^p \cdot \begin{vmatrix} \alpha_{p+1,p+1} & \dots & \alpha_{p+1,n} \\ \vdots & & \vdots \\ \alpha_{n,p+1} & \dots & \alpha_{n,n} \end{vmatrix}.$$

Proof: Observe, recalling Theorem 116 on alien cofactors, that

$$\begin{aligned} & \left[\begin{array}{ccc|ccc} \alpha_{1,1} & \dots & \alpha_{1,p} & \alpha_{1,p+1} & \dots & \alpha_{1,n} \\ \vdots & & \vdots & \vdots & & \vdots \\ \alpha_{p,1} & \dots & \alpha_{p,p} & \alpha_{p,p+1} & \dots & \alpha_{p,n} \\ \hline \alpha_{p+1,1} & \dots & \alpha_{p+1,p} & \alpha_{p+1,p+1} & \dots & \alpha_{p+1,n} \\ \vdots & & \vdots & \vdots & & \vdots \\ \alpha_{n,1} & \dots & \alpha_{n,p} & \alpha_{n,p+1} & \dots & \alpha_{n,n} \end{array} \right] \left[\begin{array}{ccc|ccc} \text{cof } \alpha_{1,1} & \dots & \text{cof } \alpha_{p,1} & 0 & \dots & 0 \\ \vdots & & \vdots & \vdots & & \vdots \\ \text{cof } \alpha_{1,p} & \dots & \text{cof } \alpha_{p,p} & 0 & \dots & 0 \\ \hline \text{cof } \alpha_{1,p+1} & \dots & \text{cof } \alpha_{p,p+1} & 1 & & 0 \\ \vdots & & \vdots & & \ddots & \\ \text{cof } \alpha_{1,n} & \dots & \text{cof } \alpha_{n,n} & 0 & & 1 \end{array} \right] \\ & = \left[\begin{array}{cc|ccc} |A| & 0 & \alpha_{1,p+1} & \dots & \alpha_{1,n} \\ & \ddots & \vdots & & \vdots \\ 0 & |A| & \alpha_{p,p+1} & \dots & \alpha_{p,n} \\ \hline 0 & \dots & 0 & \alpha_{p+1,p+1} & \dots & \alpha_{p+1,n} \\ \vdots & & \vdots & \vdots & & \vdots \\ 0 & \dots & 0 & \alpha_{n,p+1} & \dots & \alpha_{n,n} \end{array} \right] \end{aligned}$$

and take determinants on both sides. ■

11.4 Characteristic polynomials

The **characteristic polynomial** f of a square matrix A is defined by $f(\lambda) = \det(\lambda I - A)$. Roots of this polynomial are called **characteristic roots** of A .

120 Lemma *Every eigenvalue of a matrix is a characteristic root, and every real characteristic root is an eigenvalue.*

Proof: To see this note that if λ is an eigenvalue with eigenvector $x \neq 0$, then $(\lambda I - A)x = \lambda x - Ax = 0$, so $(\lambda I - A)$ is singular, so $\det(\lambda I - A) = 0$. That is, λ is a characteristic root of A .

Conversely, if $\det(\lambda I - A) = 0$, then there is some nonzero x with $(\lambda I - A)x = 0$, or $Ax = \lambda x$. ■

121 Lemma *The determinant of a square matrix is the product of its characteristic roots.*

Proof: (Cf. [6, p. 106]) Let A be an $n \times n$ square matrix and let f be its characteristic polynomial. Then $f(0) = \det(-A) = (-1)^n \det A$. On the other hand, we can factor f as

$$f(\lambda) = (\lambda - \lambda_1) \cdots (\lambda - \lambda_n)$$


where $\lambda_1, \dots, \lambda_n$ are its characteristic roots. Thus $f(0) = (-1)^n \lambda_1 \cdots \lambda_n$. ■

122 Corollary *The determinant of a symmetric matrix is the product of its eigenvalues.*

11.5 The determinant as an “oriented volume”

There is another interpretation of the determinant $\det(x_1, \dots, x_n)$. Namely it is the n -dimensional volume of the parallelotope defined by the vectors x_1, \dots, x_n , perhaps multiplied by -1 depending on the “orientation” of the vectors. I don’t want to go into detail on the notion of orientation or even to prove this assertion in any generality, but I shall present enough so that you can get a glimmer of what the assertion is.

First, what is a parallelotope? It is a polytope with parallel faces, for example, a cube. The parallelotope generated by the vectors x_1, \dots, x_n is the convex hull of zero and the vectors of the form $x_{i_1} + x_{i_2} + \dots + x_{i_k}$, where $1 \leq k \leq n$ and i_1, \dots, i_n are distinct. For instance, in \mathbf{R}^2 the parallelotope generated by x and y is the plane parallelogram with vertexes $0, x, y,$ and $x + y$.



The notion of n -dimensional volume is straightforward. For $n = 1$, it is length, for $n = 2$, it is area, etc. Oriented volume is more complicated. It distinguishes parallelotopes based on the order that the vectors are presented. For instance, in the figure above the angle swept out from x to y counterclockwise is positive, while from y to x is clockwise. The oriented volume is positive in the first case and negative in the second.

Let’s start simply in \mathbf{R}^2 , with vectors $a = (a_x, a_y)$ and $b = (b_x, b_y)$. See Figure 2. The area




Figure 2. The area of the parallelotope determined by a and b .

(two-dimensional volume, if you will) of the parallelotope P generated by the origin, $a, b,$ and $a + b$ is the area of the large square determined by the origin and $a + b$ less the area of the six regions labeled R_1, \dots, R_3 . Thus

$$\begin{aligned} \text{area } P &= (a_x + b_x)(a_y + b_y) - 2\text{area } R_1 - 2\text{area } R_2 - 2\text{area } R_3 = \\ &= (a_x + b_x)(a_y + b_y) - a_x a_y - b_x b_y - 2a_y b_x \\ &= a_x b_y - a_y b_x. \end{aligned}$$

This is just the determinant $\det(a, b)$ or the matrix determinant

$$\det A = \det \begin{bmatrix} a_x & b_x \\ a_y & b_y \end{bmatrix}.$$

I’ll leave it to you to verify that if we interchange a and b , the “oriented” area changes sign. What determines the orientation? The matrix A defines a linear transformation from \mathbf{R}^2 into

itself. It maps the first standard basis element e_1 into the first column of A , and the second basis element e_2 into the second column of A . The basis elements define a “cube” of volume 1. The image of this cube under A is the parallelotope defined by the columns of A . The vector e_2 is the vector e_1 rotated $1/4$ turn counterclockwise. If the second column of A is a counterclockwise rotation of the first column, then the columns of A have the same “orientation” as the basis vectors, and the determinant is positive.

Perhaps simplest way to convince yourself that the determinant is the oriented volume of the parallelotope, is to realize that oriented volume is an alternating multilinear form. It is clear that multiplying one of the defining vectors by a positive scalar multiplies the volume by the same factor. It is also clear that if two of the vectors are the same, then the parallelotope is degenerate and has zero volume. The additivity in the defining vectors can be appreciated, at least in two dimensions, by considering Figure 3, .




Figure 3. The area of the parallelogram defined by x and y is additive in y .

For more details see, for instance, Franklin [12, § 6.6, pp. 153–157].

11.6 Computing inverses and determinants by Gauss’ method

This section describes how to use the method of Gaussian elimination to find the inverse of a matrix and to compute its determinant. Let A be an $n \times n$ invertible matrix, so that

$$AA^{-1} = I.$$

To find the j^{th} column $x = A^{-1}I^j$ of A^{-1} , recall from the definition of matrix multiplication that x satisfies $Ax = I^j$, where I^j is the j^{th} column of the identity matrix. Indeed if A is invertible, then x is the unique vector with this property. Thus we can form the $n \times (n + 1)$ augmented coefficient matrix $(A|I^j)$ and use elementary row operations of Gaussian elimination to transform it to $(I|x) = (I|A^{-1}I^j)$. The same sequence of row operations is employed to transform A into I regardless of which column I^j we use to augment with. So we can solve for *all* the columns simultaneously by augmenting with all the columns of I . That is, form the $n \times 2n$ augmented matrix $(A|I^1, I^2, \dots, I^n) = (A|I)$. If we can transform this by elementary row operations into $(I|X)$, it must be that $X = A^{-1}$. Furthermore, if we cannot make the transformation, then A is not invertible.

You may ask, how can we tell if we cannot transform A into I ? Perhaps it is possible, but we are not clever enough. To attack this question we must write down a specific algorithm. So here is one version of an algorithm for inverting a matrix by Gaussian elimination, or else showing that it is not invertible.

Let $A^0 = A$. At each stage t we apply an elementary operation to transform A^{t-1} into A^t in such a way that

$$\det A^t \neq 0 \iff \det A \neq 0.$$

The rules for selecting the elementary row operation are described below. By stage $t = n$, either $A^t = I$ or else we shall have shown that $\det A = 0$, so A is not invertible.

Stage t : At stage t , assume that all the columns $j = 1, \dots, t-1$ of A^{t-1} have been transformed into the corresponding columns of I , and that $\det A \neq 0 \iff \det A^{t-1} \neq 0$.

Step 1: Normalize the diagonal to 1.

Case 1: $a_{t,t}^{t-1}$ is nonzero. Divide row t by $a_{t,t}^{t-1}$ to set $a_{t,t}^t = 1$. This has the side effect of setting $\det A^t = a_{t,t}^{t-1} \det A^{t-1}$.

Case 2: $a_{t,t}^{t-1} = 0$, but there is row i with $i > t$ for which $a_{i,t}^{t-1} \neq 0$. Divide row i by $a_{i,t}^{t-1}$ and add it to row t . This sets $a_{t,t}^t = 1$, and leaves $\det A^t = \det A^{t-1}$.

Case 3: If $a_{t,t}^{t-1} = 0$, but there is no row i with $i > t$ for which $a_{i,t}^{t-1} \neq 0$. In this case, the first t columns of A^{t-1} must be dependent. This is because there are t column vectors whose only nonzero components are in the first $t-1$ rows. This implies $\det A^{t-1} = 0$, and hence $\det A = 0$. This shows that A is not invertible.

If Case 3 occurs stop. We already know that A is not invertible. In cases 1 and 2 proceed to:

Step 2: Eliminate the off diagonal elements. Since $a_{t,t}^t = 1$, for $i \neq t$ multiply row t by $a_{i,t}^{t-1}$ and subtract it from row i . This sets $a_{i,t}^t = 0$, for $i \neq t$, and does not change $\det A^t$.

This completes the construction of A^t from A^{t-1} . Proceed to stage $t+1$, and note that all the columns $j = 1, \dots, t$ of A^t have been transformed into the corresponding columns of I , and that $\det A \neq 0 \iff \det A^t \neq 0$.

Now observe how this also can be used to calculate the determinant. Suppose the process runs to completion so that $A^n = I$. Every time a row was divided by its diagonal element $a_{t,t}^{t-1}$ to normalize it, we had $\det A^t = \frac{1}{a_{t,t}^{t-1}} \det A^{t-1}$. Thus we have

$$1 = \det I = \det A^n = \prod_{t: a_{t,t}^{t-1} \neq 0} \frac{1}{a_{t,t}^{t-1}} \det A$$

or

$$\det A = \prod_{t: a_{t,t}^{t-1} \neq 0} a_{t,t}^{t-1}.$$

One of the virtues of this approach is that it is extremely easy to program, and reasonably efficient. Each elementary row operation has at most $2n$ multiplications and $2n$ additions, and at most n^2 elementary row operations are required, so the number of steps grows no faster than $4n^3$.

Here are two illustrative examples, which, by the way, were produced (text and all) by a 475-line program written in `perl` of all things.

123 Example (Matrix inversion by Gaussian elimination) Invert the following 3×3 matrix using Gaussian elimination and compute its determinant as a byproduct.

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

We use only two elementary row operations: dividing a row by a scalar, which also divides the determinant, and adding a multiple of one row to a different row, which leaves the determinant unchanged. Thus the determinant of A is just the product of all the scalars used to divide the rows of A to normalize its diagonal elements to ones. We shall keep track of this product in the variable μ as we go along. At any given stage, μ times the determinant of the left hand block is equal to the determinant of A . Before each step, we put a box around the target entry to be transformed.

$\alpha_{1,1}$ is zero,

$$\mu = 1 \quad \left(\begin{array}{ccc|ccc} \boxed{0} & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 \end{array} \right)$$

so add row 3 to row 1. To eliminate $\alpha_{3,1} = 1$,

$$\mu = 1 \quad \left(\begin{array}{ccc|ccc} 1 & 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ \boxed{1} & 0 & 0 & 0 & 0 & 1 \end{array} \right)$$

subtract row 1 from row 3. To normalize $\alpha_{3,3} = -1$,

$$\mu = 1 \quad \left(\begin{array}{ccc|ccc} 1 & 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & \boxed{-1} & -1 & 0 & 0 \end{array} \right)$$

divide row 3 (and multiply μ) by -1 . To eliminate $\alpha_{1,3} = 1$,

$$\mu = -1 \quad \left(\begin{array}{ccc|ccc} 1 & 0 & \boxed{1} & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \end{array} \right)$$

subtract row 3 from row 1. This gives us

$$\mu = -1 \quad \left(\begin{array}{ccc|ccc} 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \end{array} \right)$$

To summarize:

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

And the determinant of A is -1 . We could have gotten this result faster by interchanging rows 1 and 3, but it is hard (for me) to program an algorithm to recognize when to do this. \square

124 Example (Gaussian elimination on a singular matrix) Consider the following 3×3 matrix.

$$A = \begin{pmatrix} 1 & 1 & 2 \\ 2 & 4 & 4 \\ 3 & 6 & 6 \end{pmatrix}$$

We shall use the method of Gaussian elimination to attempt transform the augmented block matrix $(A|I)$ into $(I|A^{-1})$. To eliminate $\alpha_{2,1} = 2$,

$$\mu = 1 \quad \left(\begin{array}{ccc|ccc} 1 & 1 & 2 & 1 & 0 & 0 \\ \boxed{2} & 4 & 4 & 0 & 1 & 0 \\ 3 & 6 & 6 & 0 & 0 & 1 \end{array} \right)$$

multiply row 1 by 2 and subtract it from row 2. To eliminate $\alpha_{3,1} = 3$,

$$\mu = 1 \quad \left(\begin{array}{ccc|ccc} 1 & 1 & 2 & 1 & 0 & 0 \\ 0 & 2 & 0 & -2 & 1 & 0 \\ \boxed{3} & 6 & 6 & 0 & 0 & 1 \end{array} \right)$$

multiply row 1 by 3 and subtract it from row 3. To normalize $\alpha_{2,2} = 2$,

$$\mu = 1 \quad \left(\begin{array}{ccc|ccc} 1 & 1 & 2 & 1 & 0 & 0 \\ 0 & \boxed{2} & 0 & -2 & 1 & 0 \\ 0 & 3 & 0 & -3 & 0 & 1 \end{array} \right)$$

divide row 2 (and multiply μ) by 2. To eliminate $\alpha_{1,2} = 1$,

$$\mu = 2 \quad \left(\begin{array}{ccc|ccc} 1 & \boxed{1} & 2 & 1 & 0 & 0 \\ 0 & 1 & 0 & -1 & \frac{1}{2} & 0 \\ 0 & 3 & 0 & -3 & 0 & 1 \end{array} \right)$$

subtract row 2 from row 1. To eliminate $a_{3,2} = 3$,

$$\mu = 2 \quad \left(\begin{array}{ccc|ccc} 1 & 0 & 2 & 2 & -\frac{1}{2} & 0 \\ 0 & 1 & 0 & -1 & \frac{1}{2} & 0 \\ 0 & \boxed{3} & 0 & -3 & 0 & 1 \end{array} \right)$$

multiply row 2 by 3 and subtract it from row 3. Now notice that the first 3 columns of A are dependent, as each column has at most its first 2 entries nonzero, and any 3 vectors in a 2-dimensional space are dependent.

$$\mu = 2 \quad \left(\begin{array}{ccc|ccc} \boxed{1} & 0 & 2 & 2 & -\frac{1}{2} & 0 \\ 0 & 1 & 0 & -1 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 & -1\frac{1}{2} & 1 \end{array} \right)$$

Therefore A has no inverse, and the determinant of A is zero. (If we had interchanged columns 1 and 3, we would have arrived at this result sooner, but hindsight is hard to program.) \square

125 Exercise Using the language of your choice, write a program to accept a square matrix and invert it by Gaussian elimination, or stop when you find it is not invertible. \square

References

- [1] A. C. Aitken. 1954. *Determinants and matrices*, 8th. ed. New York: Interscience.
- [2] C. D. Aliprantis and K. C. Border. 2006. *Infinite dimensional analysis: A hitchhiker's guide*, 3d. ed. Berlin: Springer-Verlag.
- [3] C. D. Aliprantis and O. Burkinshaw. 1999. *Problems in real analysis*, 2d. ed. San Diego: Academic Press. Date of publication, 1998. Copyright 1999.
- [4] T. W. Anderson. 1958. *An introduction to multivariate statistical analysis*. A Wiley Publication in Mathematical Statistics. New York: Wiley.
- [5] T. M. Apostol. 1967. *Calculus*, 2d. ed., volume 1. Waltham, Massachusetts: Blaisdell.
- [6] ———. 1969. *Calculus*, 2d. ed., volume 2. Waltham, Massachusetts: Blaisdell.
- [7] S. Axler. 1997. *Linear algebra done right*, 2d. ed. Undergraduate Texts in Mathematics. New York: Springer.
- [8] A. Berman and A. Ben-Israel. 1971. More on linear inequalities with applications to matrix theory. *Journal of Mathematical Analysis and Applications* 33(3):482–496.
DOI: [http://dx.doi.org/10.1016/0022-247X\(71\)90072-2](http://dx.doi.org/10.1016/0022-247X(71)90072-2)
- [9] K. C. Border. 2001. More than you wanted to know about quadratic forms. On-line note. <http://www.hss.caltech.edu/~kcb/Notes/QuadraticForms.pdf>
- [10] C. Carathéodory. 1982. *Calculus of variations*, 2d. ed. New York: Chelsea. This was originally published in 1935 in two volumes by B. G. Teubner in Berlin as *Variationsrechnung und Partielle Differentialgleichungen erster Ordnung*. In 1956 the first volume was edited and updated by E. Hölder. The revised work was translated by Robert B. Dean and Julius J. Brandstatter and published in two volumes as *Calculus of variations and partial differential equations of the first order* by Holden-Day in 1965–66. The Chelsea second edition combines and revises the 1967 edition.
- [11] J. Dieudonné. 1969. *Foundations of modern analysis*. Number 10-I in Pure and Applied Mathematics. New York: Academic Press. Volume 1 of Treatise on Analysis.
- [12] J. Franklin. 2000. *Matrix theory*. Mineola, New York: Dover. Reprint of 1968 edition published by Prentice-Hall, Englewood Cliffs, N.J.
- [13] D. Gale. 1989. *Theory of linear economic models*. Chicago: University of Chicago Press. Reprint of the 1960 edition published by McGraw-Hill.
- [14] F. R. Gantmacher. 1959. *Matrix theory*, volume 1. New York: Chelsea.
- [15] P. R. Halmos. 1974. *Finite dimensional vector spaces*. New York: Springer-Verlag. Reprint of the edition published by Van Nostrand, 1958.
- [16] P. N. Klein. 2013. *Coding the matrix: Linear algebra through computer science applications*. Newton, Massachusetts: Newtonian Press.

- [17] L. H. Loomis and S. Sternberg. 1968. *Advanced calculus*. Reading, Massachusetts: Addison–Wesley.
- [18] S. MacLane and G. Birkhoff. 1993. *Algebra*, 3d. ed. New York: Chelsea.
- [19] L. W. McKenzie. 1960. Matrices with dominant diagonals and economic theory. In K. J. Arrow, S. Karlin, and P. Suppes, eds., *Mathematical Methods in the Social Sciences, 1959*, number 4 in Stanford Mathematical Studies in the Social Sciences, chapter 4, pages 47–62. Stanford, California: Stanford University Press.
- [20] F. J. Murray. 1941. *An introduction to linear transformations in Hilbert space*. Number 4 in Annals of Mathematics Studies. Princeton, New Jersey: Princeton University Press. Reprinted 1965 by Kraus Reprint Corporation, New York.
- [21] C. R. Rao. 1973. *Linear statistical inference and its applications*, 2d. ed. Wiley Series in Probability and Mathematical Statistics. New York: Wiley.
- [22] H. Theil. 1971. *Principles of econometrics*. New York: Wiley.

Index

- $\angle xy$, 14
- addition
 - of scalars, 1
- adjoint, 20
 - existence, 20
- adjugate, 50
- algebraic dual, 6
- alien cofactor, 50
- angle, 14
- Apostol, Tom, 41n
- basis, 5
 - Hamel, 5
 - ordered, 7
 - standard, 6
- bilinear, 7
- Boolean field, 1
- Cauchy sequence, 9n, 16
- characteristic polynomial, 43, 52
- characteristic root, 44, 52
 - v. eigenvalue, 52
- cofactor, 48
- column space, 27
- complete inner product space, 9
- complex numbers, 1
- complex vector space, 2
- coordinate mapping, 7
- coordinates, 5
- dependence, linear, 5
- determinant, 45, 46
 - = product of characteristic roots, 52
 - and invertibility, 48, 51
 - and row operations, 47
 - as oriented volume, 53
 - expansion by alien cofactors, 50
 - Jacobi's Theorem, 51
 - of a linear transformation, 47
 - of a product, 47
 - of a transpose, 47
 - of a triangular matrix, 47
 - of an orthogonal matrix, 48
 - recursive computation, 49
- diagonal matrix, 28
- dimension, 6
- dominant diagonal matrix, 33
- dot product, 8
- dual space, 6
 - algebraic, 6
 - topological, 6, 14
- eigenspace, 24
- eigenvalue, 24
 - multiplicity of, 24
 - v. characteristic root, 52
- eigenvector, 24
- elementary row operations, 29
- Euclidean n -space, 8
- field, 1
 - scalar, 1
- form, 41
- frame, 7
- Gaussian elimination algorithm, 29
 - for computing determinant, 55–58
 - for matrix inversion, 55–58
- Gershgorin's Theorem, 32
- GF(2), 1
- Gram–Schmidt procedure, 11
- Hilbert space, 9, 15, 16, 20
- homomorphism, 6
- identity matrix, 28
- independence, linear, 5
- inner product, 7, 40
- inner product space
 - complete = Hilbert, 9
- inverse, 19
- isometry, 23
- isomorphism, 7
- Jacobi's theorem on determinants, 51

- kernel, 18
- $L(V, W)$, 6
- ℓ_p , 3
- left inverse, 19
- linear combination, 4
- linear dependence, 5
- linear functional, 6
 - discontinuous, 17
- linear independence, 5
- linear operator, 6
- linear space = vector space, 2
- linear subspace, 4
- linear transformation, 6
 - idempotent, 23
 - matrix representation of, 30
- matrix, 27
 - characteristic polynomial, 43
 - determinant of, 45
 - diagonal, 28
 - identity, 28
 - inverse, 32
 - main diagonal, 28
 - nonsingular, 43
 - orthogonal, 36
 - product, 27
 - representation of a linear transformation, 30
 - scalar multiple of, 27
 - similarity of, 35, 39
 - singular, 43
 - sum, 27
 - trace, 38
 - transpose, 32
 - triangular, 28
 - zero, 28
- minor, 48
- minor determinant, 48
- minor submatrix, 48
- $M(m, n)$, 3, 27
- multilinear form, 44
 - alternating, 44
- multiplication
 - of scalars, 1
- norm, 8
- null space, 18
- nullity, 18
- operator norm, 15
- ordered basis, 7
- orthogonal complement, 10
- orthogonal matrix, 36
 - determinant of, 48
- orthogonal projection, 12
- orthogonal transformation, 22
 - is norm-preserving, 22
 - preserves inner products, 22
- orthogonality, 10
- orthonormality, 10
- parallelogram law, 9
- permutation, 46
 - signature, 46
- Principal Axis Theorem, 36
- principal minor of order p , 48
- quadratic form, 41
 - and eigenvalues, 42
 - definite, semidefinite, 41
- \mathbf{R}^n , 2
- rank, 18
- rational numbers, 1
- real numbers, 1
- real vector space, 2
- right inverse, 19
- row space, 27
- scalar, 1
- scalar addition, 1
- scalars, 1
- self-adjoint transformation, 23
- signature of a permutation, 46
- similar matrices, 35, 39
- skew-symmetric transformation, 23
- span, 4
- standard basis, 6
- sum, 1
- symmetric transformation, 23
- Todd, John, 41n

topological dual, 6

trace, 38

transpose, 20

unit coordinate vectors, 2

V' , 6

V^* , 6

vector, 2

vector space, 2

 trivial, 3

vector subspace, 4