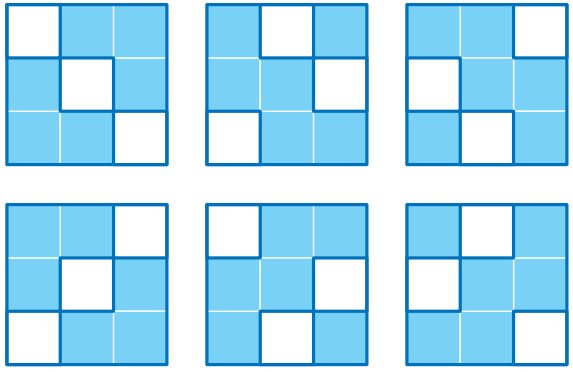

Linear Algebra

Dave Bayer

Projective Press
New York, NY USA

http://projectivepress.com/LinearAlgebra

Dave Bayer
Department of Mathematics
Barnard College
Columbia University
New York, New York

Draft of January 12, 2007

Contents

Preface

This is an early draft of a textbook on linear algebra. There is an accompanying book, *An Atlas of Matrices*, giving examples suitable for working by hand. I am making these texts available in incomplete form for the benefit of my linear algebra students at Barnard College and Columbia University.

Linear algebra is a transitional subject in the mathematics curriculum. It can be one of the last math courses taken by students heading off to various applications, and these students are well served by existing textbooks that focus on applications. However, pulled in many directions by these competing interests, the *algebra* in linear algebra has a tendency to get lost. Mathematics itself, and algebra in particular, is a great application of linear algebra.

Linear algebra is also the first semester of algebra for math majors, leading directly into the *modern algebra* sequence. Students who haven't picked up a fair bit of algebra by the time that they begin modern algebra have a tendency to hit a brick wall. I envision this book as a supplementary text for math majors who have the ambition to learn as much algebra as they can, as soon as they can. For example, we use linear algebra to study the algebra of polynomials, and then apply this theory to better understand functions of a matrix.

This draft is currently in a very fragmented state. For now, there are various stubs for chapters-to-be. I am completing the chapter on eigenvalues and eigenvectors first, so that I will know exactly what material is needed to prepare for this chapter. For this reason, partial drafts of this text will not be self-contained.

This is a draft of January 12, 2007. The most recent drafts of these texts may be downloaded from

http://projectivepress.com/LinearAlgebra

DAVE BAYER
New York, NY
January 2007

Introduction

A *matrix* is a box of numbers arranged in a grid, like

$$A = \begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix}$$

Linear algebra is the study of matrices. Even in its most advanced form, where one studies linear operators on infinite-dimensional spaces with no coordinate system in sight, one relies on intuitions built by getting really good at understanding small matrices that can be manipulated by hand. Such matrices are the focus of this introductory text. We will pile on lots of theory, but theory that is relevant to tackling increasingly difficult problems involving small matrices.

Matrices act by multiplication on *vectors*, elements of n-dimensional space, in the same way that numbers act by multiplication on other numbers in the real line. One can think of matrices as "rich" numbers. Like computer programmers following a modular design, we will alternately look inside the box of numbers to fiddle with the low-level behavior of a matrix, and close up the box to think of the matrix as a single entity. The analogy between matrices and numbers is a powerful one; we can add and multiply matrices, solve systems of equations by multiplying by inverse matrices, substitute matrices for variables in polynomial expressions, make matrices of matrices, and so forth. This is the spirit of algebra, accepting a broad notion as to what constitutes a value that may be manipulated in algebraic expressions. Linear algebra is the algebra of matrices.

There is room in n dimensions for more subtle behavior than in one dimension, so the behavior of matrices can be more subtle than that of numbers. For example, in the real line there is only room to make half-turns, facing alternately in the directions of ∞ and $-\infty$. Multiplication by -1 can be thought of as a half-turn. Already in the plane there is room to rotate by an arbitrary angle, and there are matrices that represent each of these rotations, just as there are complex numbers that represent each of these rotations.

Three objects in a row can be permuted in more than one way, e.g. swap the two objects on the left, or take the object on one end and move it to the other end, slid-

ing the other objects over. This is the simplest example of operations where order matters; carrying out these steps gives different results, depending on which step comes first. Matrices can replicate this behavior, by rearranging the order of the coordinates in 3-dimensional space, and matrix multiplication can give different results, depending on which matrix comes first. We say that matrix multiplication is *noncommutative*.

This again is in keeping with the spirit of algebra. Rather than hard-wiring the rules of real arithmetic into our brains, we learn to flexibly flip switches to reconfigure the rules that we apply, so that we don't automatically make simplifications that aren't allowed. To work with matrices, we learn to keep track of the order of multiplication, which is not an issue in real arithmetic. Why always play chess when there are other games? The appeal of algebra is the appeal of learning new games. For all its complexity, matrix arithmetic is far more interesting than real arithmetic.

The complexity of matrices is the complexity of the real world. It makes a difference whether one cracks an egg before frying it, or fries it in its shell and cracks it later. Such is a matter of taste, but the order matters. More tragically, one can drink a glass of wine and then drop the glass, but not easily the other way around. Order matters in a vast array of applications which can be modeled by matrices, but not by numbers. Not only can we think of matrices as numbers, we can think of matrices as *actions*. For example, we will first learn to solve systems of linear equations by rescaling, adding, and swapping around the equations. Then, we will store these actions as matrices, with the property that multiplying by the matrix carries out the action. This is a common mode of thought in linear algebra. When we say a matrix acts on a space, we mean this quite literally.

Historically, the study of matrices began with the study of solutions to linear systems of equations. Only later was the subject generalized to the study of *linear maps*, functions with a property exemplified by matrices. We approach each of these points of view on an equal footing; the first two chapters can be read independently of each other.

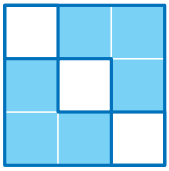

Chapter 1

Systems of equations

Lorem ipsum dolor sit amet, consectetur adipisicing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

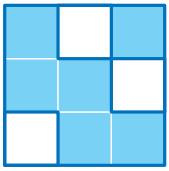

Chapter 2

Linear maps

Lorem ipsum dolor sit amet, consectetur adipisicing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

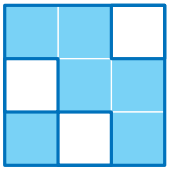

Chapter 3

Vector spaces

Lorem ipsum dolor sit amet, consectetur adipisicing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

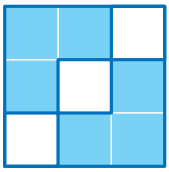

Chapter 4

Determinants

The *determinant* is an expression in the entries of a square matrix, that appears throughout linear algebra. If it didn't already come with hundreds of years of history and a myriad of interpretations, the attentive student would notice the determinant as a pattern cropping up all over the place, and give it a name. For us, the determinant is first and foremost a formula, whose pattern we want to understand. Then we will consider interpretations of this formula.

For an operational definition, the determinant of a matrix is the number that most closely resembles the matrix. Matrices are capable of far more subtle behavior than numbers, such as representing actions where the order of operations matters, so it is too much to ask that there be a number that exactly reflects the properties of a matrix. If this were possible, then we wouldn't need matrices. Nevertheless, it is reasonable to ask that an invertible matrix be represented by an invertible number, for the product of two matrices to be represented by the product of their corresponding numbers, and so forth. The determinant is that number.

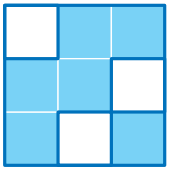

Chapter 5

Coordinates

Lorem ipsum dolor sit amet, consectetur adipisicing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

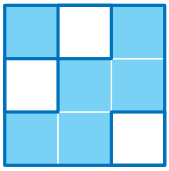

Chapter 6

Polynomials

There are striking parallels between the algebra of matrices and the algebra of polynomials. To prepare for the theory of functions of a matrix, we review some polynomial algebra.

If one steps back and takes a fresh look at polynomials, they now look very much a part of linear algebra. A polynomial in the variable x is after all a *linear combination* of powers of x. This is a useful point of view. If we take all powers of some object x in an arbitrary algebraic setting over a field F, either the powers are linearly independent, spanning an infinite dimensional subspace, or they are linearly dependent, and we can write the first dependent power

$$x^d = c_{d-1}x^{d-1} + \ldots + c_1x_1 + c_0$$

as a linear combination of the preceding powers. This rule can then be used repeatedly to rewrite any higher power of x in terms of $\{x^{d-1}, \ldots, x, 1\}$, showing that this set is a basis for all powers of x. In this case, x behaves as if we are working modulo the polynomial

$$p(x) = x^d - c_{d-1}x^{d-1} - \ldots - c_1x_1 - c_0$$

Square matrices are such a setting. The space of all $n \times n$ matrices over a field F is an n^2-dimensional vector space, so it isn't possible for all powers of an $n \times n$ matrix A to be linearly independent of each other. There simply isn't enough room. It must be the case that for some $d \leq n^2$,

$$\mathrm{A}^d = c_{d-1}\mathrm{A}^{d-1} + \ldots + c_1\mathrm{A}_1 + c_0\mathrm{I}$$

In other words, it is inevitable that square matrices behave as if we are working modulo some polynomial, because their powers live in a finite dimensional vector space.

We are very interested in understanding such polynomials. It turns out that the least such d is smaller than this argument suggests; we can always find such a dependence with $d \leq n$. Viewing A as a linear map $V \to V$, it is the dimension n of V, not the dimension n^2 of possible matrices, that matters. Still, this first argument makes it clear that there must be some such polynomial, hence our interest now in polynomials.

6.1 Modular arithmetic

Integers

Let m be an integer. Two integers a, b are equivalent mod m if their difference $a - b$ is a multiple of m. We write

$$a \equiv b \mod m$$

and say that we are computing in the ring $\mathbb{Z}/m\mathbb{Z}$ of integers mod m. For example, $8 \equiv 2 \mod 3$, because $8 - 2 = 6$ is a multiple of 3. The distinct elements of $\mathbb{Z}/3\mathbb{Z}$ are $\{0, 1, 2\}$, which are the possible remainders under division by 3.

It is a tremendous simplification to be able to work modulo an integer m; our calculations can take place in the finite set $\{0, 1, \ldots, m-1\}$ of remainders under division by m. This can mean the difference between a value fitting in one word of memory on a computer, and a value not fitting on the computer at all.

Example 6.1. The integer 2^{1024} is a 309 digit number, but $2^{1024} \bmod 5$ can be computed by repeated squaring mod 5; after the second squaring, the value becomes and stays 1:

$$\begin{aligned} & & 2^2 & \equiv 4 \mod 5 \\ 2^4 & \equiv (2^2)^2 & \equiv 1 & \mod 5 \\ \ldots\ 2^{1024} & \equiv (2^{512})^2 & \equiv 1 & \mod 5 \end{aligned}$$

Polynomials

One can also work with polynomials modulo a polynomial $p(x)$. Computationally, this is again a tremendous simplification; our calculations can all take place in the finite dimensional vector space of remainders under division by p.

Two polynomials f, g are equivalent mod p if their difference $f - g$ is a multiple of p. We again write

$$f(x) \equiv g(x) \mod p(x)$$

and say that we are computing in the ring $F[x]/p(x)F[x]$ of polynomials mod p. In this notation, $F[x]$ stands for the ring of all polynomials in the variable x with coefficients in our field F.

Example 6.2. Let $p(x) = x^2 + 1$. We have

$$x^2 \equiv -1 \mod x^2 + 1$$

because $x^2 - (-1) = x^2 + 1$. The distinct elements of $\mathbb{R}[x]/(x^2+1)\mathbb{R}[x]$ are now all possible remainders under division by $x^2 + 1$. This is a 2-dimensional vector space over $\mathbb{R}$ with basis $\{1, x\}$.

The complex numbers

The complex numbers $\mathbb{C}$ can be viewed as an example of polynomial modular arithmetic.

Example 6.2 looks familiar. The complex numbers $\mathbb{C}$ are also a 2-dimensional vector space over $\mathbb{R}$, with basis $\{1, \mathrm{i}\}$, and $\mathrm{i}^2 = -1$. The complex numbers look like $\mathbb{R}[x]/(x^2+1)\mathbb{R}[x]$, only with x replaced by i. They are an important special case of working modulo a polynomial, with x getting the special name i.

Viewing $\mathbb{C}$ as the vector space $\mathbb{R}^2$, we can represent multiplication by i as a 2×2 matrix A. Multiplication by i maps the basis $\{1, \mathrm{i}\}$ to $\{\mathrm{i}, -1\}$, so in vector notation we want $\mathrm{A}(1,0) = (0,1)$ and $\mathrm{A}(0,1) = (-1,0)$. This is the matrix

$$\mathrm{A} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$$

which we recognize as a rotation by $\pi/2$. Indeed, multiplication by i rotates the complex plane by $\pi/2$. For all intents and purposes, this matrix A is the imaginary number i. The matrix A goes with the notation $\mathbb{R}^2$, the variable x goes with the notation $\mathbb{R}[x]/(x^2+1)\mathbb{R}[x]$, and i goes with the notation $\mathbb{C}$, but they are all the same thing.

If A and i are the same thing, what about $\mathrm{i}^2 + 1 = 0$? We have

$$\mathrm{A}^2 + \mathrm{I} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}^2 + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = 0$$

A acts just like i, so this is no surprise. We now have three views of the complex numbers: $\mathbb{C}$, and $\mathbb{R}[x]/(x^2+1)\mathbb{R}[x]$, and the ring of all polynomial expressions in A. For example, the complex number $2 - 3\mathrm{i}$ can be viewed as the polynomial $2 - 3x \mod x^2 + 1$, and as the matrix expression $2\mathrm{I} - 3\mathrm{A}$.

For an arbitrary square matrix A, the ring of all polynomial expressions in A is not so different. The governing equation $x^2 + 1 = 0$ changes, but the idea is the same. $\mathbb{C}$ is just a special case, with its own notation.

The dual numbers

A closely related example is the ring of *dual numbers*, modeled after the idea of an infinitesimal in calculus. In calculus, ϵ is taken to be so small that ϵ^2 is effectively zero. In algebra, we simply dictate that $\epsilon^2 = 0$ by working modulo ϵ^2. In the notation of modular arithmetic, we have

$$\epsilon^2 \equiv 0 \mod \epsilon^2$$

We are again working modulo a polynomial, with our variable getting the special name ϵ. The ring of dual numbers can be denoted $\mathbb{R}[\epsilon]/\epsilon^2\mathbb{R}[\epsilon]$.

Not surprisingly given its origins, this ϵ can be used to differentiate algebraic expressions. If $f(x)$ is a polynomial or power series, then

$$f(a+\epsilon) \equiv f(a) + f'(a)\epsilon \mod \epsilon^2 \tag{6.1}$$

This is just a version of the familiar equation

$$f'(a) = \lim_{\epsilon \to 0} \frac{f(a+\epsilon) - f(a)}{\epsilon}$$

In calculus, a limit is required to suppress the effects of ϵ^2. Here, no limit is required because we are working modulo ϵ^2.

When we are firmly ensconced in this setting, we will drop the notation $\equiv$ in favor of $=$ and stop writing "mod ϵ^2" after every equation.

Example 6.3. If $f(x) = x^3$ then working modulo ϵ^2,

$$f(1+\epsilon) = (1+\epsilon)^3 = 1 + 3\epsilon + 3\epsilon^2 + \epsilon^3 = 1 + 3\epsilon$$

This agrees with $f'(1) = 3$.

The dual numbers are a 2-dimensional vector space over $\mathbb{R}$ with basis $\{\epsilon, 1\}$. Viewing the dual numbers as the vector space $\mathbb{R}^2$, we can represent multiplication by ϵ as a 2×2 matrix N. Multiplication by ϵ maps the basis $\{\epsilon, 1\}$ to $\{0, \epsilon\}$, so in vector notation we want $\mathrm{N}(1,0) = (0,0)$ and $\mathrm{N}(0,1) = (1,0)$. This is the matrix

$$\mathrm{N} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

This matrix is the simplest example of a *nilpotent* matrix:

Definition 6.4. A square matrix N is *nilpotent* if some power of N is the zero matrix.

Here we have $\mathrm{N}^2 = 0$, mirroring the fact that $\epsilon^2 = 0$. We can view the dual numbers as the ring of all polynomial expressions in N.

Example 6.5. If the matrix N acts exactly like ϵ, then we should be able to differentiate using N. Indeed, again taking $f(x) = x^3$,

$$\begin{aligned} f(\mathrm{I}+\mathrm{N}) &= \left(\begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix} + \begin{bmatrix}0 & 1\\ 0 & 0\end{bmatrix}\right)^3 \\ &= \begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix} + 3\begin{bmatrix}0 & 1\\ 0 & 0\end{bmatrix} \\ &= f(1)\,\mathrm{I} + f'(1)\,\mathrm{N} \end{aligned}$$

This calculation agrees with that of example 6.3, with ϵ replaced by N.

In general, equation 6.1 takes the matrix form

$$f(a\mathrm{I}+\mathrm{N}) \;=\; f(a)\,\mathrm{I} \;+\; f'(a)\,\mathrm{N} \tag{6.2}$$

where N is any nilpotent matrix that acts like ϵ. In other words, N can be any square matrix with $\mathrm{N}^2 = 0$.

Rational canonical form

In both of the preceding examples, we could view the ring of polynomials modulo $p(x)$ as a 2-dimensional vector space with basis $\{x, 1\}$. In both of these rings, x^2 was a linear combination of these basis elements, allowing us to reduce higher powers of x to an expression in x and 1.

In general, when $p(x)$ is a polynomial of degree n, the distinct elements of $F[x]/p(x)F[x]$ are the possible remainders under division by p. Our ring is now an n-dimensional vector space with basis $\{x^{n-1}, \ldots, x, 1\}$. In this ring, x^n is a linear combination of these basis elements, allowing us to reduce higher powers of x to an expression in $x^{n-1}, \ldots, x, 1$. For some coefficients $c_{n-1}, \ldots, c_1, c_0$ we have

$$x^n \equiv c_{n-1}x^{n-1} \;+\; \ldots \;+\; c_1x_1 \;+\; c_0 \mod p(x)$$

This would not change if we replace p by any nonzero multiple of itself, so we might as well take p to be the monic polynomial

$$p(x) \;=\; x^n \;-\; c_{n-1}\,x^{n-1} \;-\; \ldots \;-\; c_1\,x \;-\; c_0$$

We can represent multiplication by x as an $n \times n$ matrix A. Multiplication by x maps the basis $\{x^{n-1}, \ldots, x, 1\}$ to

$$\{\, c_{n-1}x^{n-1} + \ldots + c_1x + c_0,\; x^{n-1},\; \ldots,\; x \,\}$$

Example 6.6. If

$$p(x) \;=\; x^3 - ax^2 - bx - c$$

then multiplication by x maps the basis $\{x^2, x, 1\}$ to $\{ax^2 + bx + c, x^2, x\}$. Let $\boldsymbol{e}_1 = (1,0,0)$, $\boldsymbol{e}_2 = (0,1,0)$, and $\boldsymbol{e}_3 = (0,0,1)$, as usual. In vector notation we want $\mathrm{A}\boldsymbol{e}_1 = (a,b,c)$, $\mathrm{A}\boldsymbol{e}_2 = \boldsymbol{e}_1$, and $\mathrm{A}\boldsymbol{e}_3 = \boldsymbol{e}_2$. This is the matrix

$$\mathrm{A} \;=\; \begin{bmatrix} a & 1 & 0 \\ b & 0 & 1 \\ c & 0 & 0 \end{bmatrix}$$

This matrix is in *rational canonical form*. An arbitrary square matrix is similar to a block diagonal matrix with blocks of this form.

Definition 6.7. A matrix A is in *rational canonical form*[1] if it has the form

$$\mathrm{A} \;=\; \begin{bmatrix} c_{n-1} & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ c_1 & 0 & \dots & 1 \\ c_0 & 0 & \dots & 0 \end{bmatrix}$$

in which case we say that A is the *companion matrix* of the polynomial

$$p(x) = x^n - c_{n-1}x^{n-1} - \dots - c_1 x - c_0$$

The matrix A constructed in example 6.6 acts like x mod $p(x)$, so it will be the case that $p(\mathrm{A}) = 0$. However, we would like to be able to see directly that for any matrix A which is the companion matrix of a polynomial $p(x)$, we have $p(\mathrm{A}) = 0$. There is a pattern to powers of A that could be used to show this, writing $\mathrm{A} = \mathrm{P} + \mathrm{N}$ for a nilpotent matrix N, but the pattern isn't particularly illuminating.

The simplest way to proceed is to argue that such an A is the same matrix that we would get by following the above construction, so of course $p(\mathrm{A}) = 0$. This is correct reasoning, but it begs the question of whether we really understand the above construction. A lot of math gets built quickly this way, so to learn to sketch new theories one should get the hang of thinking like this. At same time, we always want to be able to extract an explicit argument from such reasoning.

The following argument is essentially this idea, but spelled out with no mention of the ring $F[x]/p(x)F[x]$:

Proposition 6.8. *If a matrix* A *is the companion matrix of the polynomial* $p(x)$*, then* $p(\mathrm{A}) = 0$*.*

[1] There are many versions of rational canonical form in use, varying with the sign convention for the coefficients of p, the order of the basis $\{x^{n-1}, \dots, x, 1\}$, and whether or not one transposes to obscure the meaning of the columns. We have chosen a form that agrees with the prevailing convention for writing nilpotent matrices and Jordan canonical form.

Proof. We give a proof for the 3×3 case; the proof for $n \times n$ matrices is the same. Let A be the companion matrix

$$\mathrm{A} \;=\; \begin{bmatrix} a & 1 & 0 \\ b & 0 & 1 \\ c & 0 & 0 \end{bmatrix}$$

of the polynomial

$$p(x) \;=\; x^3 - ax^2 - bx - c$$

We want to show that $p(\mathrm{A}) = 0$.

We have the progression

$$\begin{array}{ccccccc} \boldsymbol{e}_3 & \overset{\mathrm{A}}{\longmapsto} & \boldsymbol{e}_2 & \overset{\mathrm{A}}{\longmapsto} & \boldsymbol{e}_1 & \overset{\mathrm{A}}{\longmapsto} & a\boldsymbol{e}_1 + b\boldsymbol{e}_2 + c\boldsymbol{e}_3 \\ 1 & & x & & x^2 & & x^3 \end{array}$$

so

$$(\mathrm{A}^3 - a\,\mathrm{A}^2 - b\,\mathrm{A} - c\,\mathrm{I})\,\boldsymbol{e}_3 \;=\; (a\boldsymbol{e}_1 + b\boldsymbol{e}_2 + c\boldsymbol{e}_3) - a\boldsymbol{e}_1 - b\boldsymbol{e}_2 - c\boldsymbol{e}_3 \;=\; 0$$

In other words, the matrix

$$p(A) \;=\; \mathrm{A}^3 \;-\; a\,\mathrm{A}^2 \;-\; b\,\mathrm{A} \;-\; c\,\mathrm{I}$$

maps $\boldsymbol{e}_3$ to zero.

What about $\boldsymbol{e}_2$ and $\boldsymbol{e}_1$, joining the class late? We have

$$\boldsymbol{e}_2 \;=\; \mathrm{A}\,\boldsymbol{e}_3 \quad \text{and} \quad \boldsymbol{e}_1 \;=\; \mathrm{A}^2\,\boldsymbol{e}_3$$

so

$$\begin{array}{rcccl} p(\mathrm{A})\,\boldsymbol{e}_2 & = & p(\mathrm{A})\,\mathrm{A}\,\boldsymbol{e}_3 & = & \mathrm{A}\,p(A)\,\boldsymbol{e}_3 \;=\; 0 \\ p(\mathrm{A})\,\boldsymbol{e}_1 & = & p(\mathrm{A})\,\mathrm{A}^2\,\boldsymbol{e}_3 & = & \mathrm{A}^2\,p(A)\,\boldsymbol{e}_3 \;=\; 0 \end{array}$$

Since $p(A)$ maps the basis $\{\boldsymbol{e}_1, \boldsymbol{e}_2, \boldsymbol{e}_3\}$ to zero, it must be the zero matrix. □

So far, the flow of ideas has been to interpret algebras such as the field $\mathbb{C}$ first as modular polynomial arithmetic, then as all polynomial expressions in a matrix. We have now reversed this flow. Starting with a matrix A in rational canonical form, we have interpreted the ring of polynomial expressions in A as modular polynomial arithmetic. Doing this for any square matrix A is the subject of chapter 7.

A geometric interpretation

When a polynomial p factors into distinct linear factors, there is a close connection between arithmetic mod $p(x)$ and the set of roots of p. Namely, a polynomial is equivalent to zero mod $p(x)$ if and only if it vanishes on the roots of p. It follows that two polynomials are equivalent mod $p(x)$ if and only if they have the same values on the roots of p.

Theorem 6.9. *Let f and p be polynomials in x with coefficients in the field F, and suppose that p factors as*

$$p(x) \;=\; (x - a_1) \cdots (x - a_d)$$

where the roots $a_1, \ldots, a_d$ are distinct. Then $f(x) \equiv 0$ mod $p(x)$ if and only if

$$f(a_1) \;=\; \ldots \;=\; f(a_d) \;=\; 0$$

Proof. If $f(x) \equiv 0 \bmod p(x)$ then f is a multiple of p, so $f(a_i) = 0$ for each root a_i of p.

Conversely, suppose that $f(a_i) = 0$ for each root a_i of p. For any element a of F and for any polynomial $g(x)$ with coefficients in F we have

$$g(a) = 0 \quad \Leftrightarrow \quad (x - a) \mid g(x)$$

This part of the argument works for coefficients in any ring; dividing $x - a$ into $g(x)$ by long division yields a remainder of $g(a)$, so $x - a$ divides $g(x)$ if and only if $g(a) = 0$.

Applying this to $f(a_1) = 0$, write

$$f(x) \;=\; (x - a_1)\, g(x)$$

Substituting $x = a_2$, we have

$$f(a_2) \;=\; (a_2 - a_1)\, g(a_2) \;=\; 0 \tag{6.3}$$

Because $a_1 \neq a_2$, it follows that $g(a_2) = 0$, so $(x - a_2)$ divides g. Continuing by induction, we can write

$$f(x) \;=\; (x - a_1)\,(x - a_2) \cdots (x - a_d)\, h(x)$$

so f is a multiple of p. Therefore $f(x) \equiv 0 \bmod p(x)$. □

Equation 6.3 is the crux of this argument. Our reasoning actually works for coefficients in an *integral domain*, a ring in which the product of any two nonzero elements is nonzero. Soon, we will be tempted to view a_1 and $a_1 + \epsilon$ as distinct values and apply theorem 6.9. We can't; the ring of dual numbers is the poster child

for a ring that isn't an integral domain, because ϵ is nonzero but ϵ^2 is zero. However, the spirit of this idea points us in the right direction. What we can do instead is to consider the coefficient of ϵ in various dual number expressions, analogous to considering the imaginary part of a complex number.

Working mod $p(x)$ is adopting the stance that the domain of our polynomials is the set of roots of p. We're taking p itself to be zero because p is zero on these roots. Two polynomials g and h may disagree on some elements of R other than these roots, but we're adopting the stance that we don't care about any other elements of R. As long as g and h are the same function when restricted to the roots of p, to us they are the same function.

This point of view is the beginning of a subject called *algebraic geometry*. There, one studies systems of polynomial equations in many variables, both in terms of the geometry of the solution sets, and in terms of the algebra modulo the defining equations. Again, this is taking the stance that the solution sets are the domains of the polynomials, as if all other points don't even exist. This turns out to be a powerful idea. Around the same time that painting went abstract, so did algebraic geometry. Somebody flipped a switch, and suddenly all of the ambient spaces were gone, leaving just bare point sets, curves and surfaces to study by themselves. We're flipping this switch by working modulo $p(x)$.

When two roots of p come together, say a_1 and a_2, theorem 6.9 requires modification. The condition $f(a_1) = 0$ insures that $x - a_1$ divides f, but it is not enough to insure that $(x - a_1)^2$ divides f. It takes d conditions to determine whether $f(x) \equiv 0$ mod $p(x)$. When a_1 and a_2 come together, we have lost a condition, which we recover as the condition $f'(a_1) = 0$.

Proposition 6.10. *Let $f(x)$ be a polynomial in x with coefficients in a field F. If a is a root of the polynomial $f(x)$, then a is a repeated root of $f(x)$ if and only if a is also a root of $f'(x)$.*

Proof. Write $f(x) = (x - a)g(x)$. Then a is a repeated root of $f(x)$ if and only a is a root of $g(x)$. By the product rule,

$$f'(x) \;=\; g(x) \;+\; (x - a)\, g'(x)$$

so $f'(a) = g(a)$. □

Example 6.11. If

$$f(x) \;=\; (x-a)^2\,(x-b)\,(x-c)$$

then

$$f'(x) \;=\; 2\,(x-a)\,(x-b)\,(x-c) \;+\; (x-a)^2\,(x-c) \;+\; (x-a)^2\,(x-b)$$

Because a is a repeated root of $f(x)$, differentiating once cannot rid f of the factor $(x - a)$.

This pattern reappears throughout this chapter. Working over the field $\mathbb{R}$, knowing $f(a_1)$ and $f'(a_1)$ is pretty much the same information as knowing $f(a_1)$ and $f(a_2)$ when a_1 and a_2 are close; we can use $f'(a_1)$ to estimate $f(a_2)$. However, this pattern holds for any field, including fields where we cannot reason analytically. We can work with an arbitrary field F by using the ring of dual numbers.

Proposition 6.12. *Let f and p be polynomials in x with coefficients in the field F, and suppose that p factors as*

$$p(x) \;=\; (x-a_1)^2\,(x-a_3)\cdots(x-a_d)$$

where the roots $a_1, a_3, \ldots, a_d$ are distinct. Then $f(x) \equiv 0$ mod $p(x)$ if and only if

$$f(a_1) \;=\; f'(a_1) \;=\; f(a_3) \;=\; \ldots \;=\; f(a_d) \;=\; 0 \tag{6.4}$$

Proof. If $f(x) \equiv 0 \bmod p(x)$, then f is a multiple of p. We work with the values

$$a_1,\; a_1+\epsilon,\; a_3,\; \ldots,\; a_d \tag{6.5}$$

Note that working modulo ϵ^2,

$$p(a_1+\epsilon) \;=\; \epsilon^2\,(a_1+\epsilon-a_3)\cdots(a_1+\epsilon-a_d) \;=\; 0$$

Therefore $p(x)$ vanishes on each of the values 6.5, so we have

$$f(a_1) \;=\; f(a_1+\epsilon) \;=\; f(a_3) \;=\; \ldots \;=\; f(a_d) \;=\; 0$$

Since $f(a_1+\epsilon) = f(a_1) + \epsilon f'(a_1)$, we get the condition $f'(a_1) = 0$.

Conversely, suppose that the conditions 6.4 hold. From $f(a_1) = f'(a_1) = 0$ we get $f(a_1+\epsilon) = 0$. Write

$$f(x) \;=\; (x-a_1)\,g(x)$$

where g is the quotient of f under division by $x-a_1$. Then

$$f(a_1+\epsilon) \;=\; (a_1+\epsilon-a_1)\,g(a_1+\epsilon) \;=\; \epsilon\,(g(a_1)+\epsilon\,g'(a_1)) \;=\; \epsilon\,g(a_1) \;=\; 0$$

This is identically zero as an expression in the ring of dual numbers, so the coefficient $g(a_1)$ of ϵ must be zero. Therefore, $x-a_1$ divides $g(x)$, so $x-a_1$ divides $f(x)$ twice. It follows that f is a multiple of p, so $f(x) \equiv 0 \bmod p(x)$. □

We could have instead used proposition 6.10 here. We applied the ring of dual numbers to illustrate their use; they will appear again.

6.2 Polynomial interpolation

A polynomial of degree $< d$ is determined by its values at d distinct points. There is a steep, direct assault on this statement using the *Vandermonde matrix*, and a clever way to sidestep its complexity using *Lagrange interpolation*.

The Vandermonde matrix

Proposition 6.13. *Let $a_1, \ldots, a_d$ be distinct values in the field F. There is a unique polynomial of degree $< d$*

$$f(x) \;=\; c_{d-1}x^{d-1} \;+\; \ldots \;+\; c_1 x \;+\; c_0$$

determined by the values $f(a_1), \ldots, f(a_d)$.

Proof. The coefficients of $f(x)$ are determined by the system of equations

$$\begin{bmatrix} a_1^{d-1} & a_1^{d-2} & \ldots & a_1 & 1 \\ a_2^{d-1} & a_2^{d-2} & \ldots & a_2 & 1 \\ \vdots & \vdots & & \vdots & \vdots \\ a_{d-1}^{d-1} & a_{d-1}^{d-2} & \ldots & a_{d-1} & 1 \\ a_d^{d-1} & a_d^{d-2} & \ldots & a_d & 1 \end{bmatrix} \begin{bmatrix} c_{d-1} \\ c_{d-2} \\ \vdots \\ c_1 \\ c_0 \end{bmatrix} = \begin{bmatrix} f(a_1) \\ f(a_2) \\ \vdots \\ f(a_{d-1}) \\ f(a_d) \end{bmatrix} \tag{6.6}$$

where the ith row is the equation

$$c_{d-1}\, a_i^{d-1} \;+\; c_{d-2}\, a_i^{d-2} \;+\; \ldots \;+\; c_1\, a_i \;+\; c_0 \;=\; f(a_i)$$

Let A denote the coefficient matrix of this system of equations. A matrix of this form is called a *Vandermonde matrix*. We will show that A has the determinant

$$\det(\mathrm{A}) \;=\; \prod_{i<j} (a_i - a_j) \tag{6.7}$$

It follows that this system of equations has a unique solution whenever each $a_i - a_j \neq 0$.

It is a worthy exercise to deduce equation 6.7 directly by induction. However, there is also an argument using modular arithmetic. Let g be the polynomial given by the right hand side of equation 6.7, where we take $a_1, \ldots, a_d$ to be variables:

$$g(a_1, \ldots, a_d) \;=\; \prod_{i<j} (a_i - a_j)$$

Consider the term of g obtained by taking the product of the first variables in each factor,

$$g(a_1, \ldots, a_d) \;=\; a_1^{d-1}\, a_2^{d-2} \cdots a_{d-1} \;+\; \ldots$$

We recognize this term as the product of the diagonal entries of A, so it is also a term of $\det(\mathrm{A})$, viewed as a polynomial in $a_1, \ldots, a_d$:

$$\det(\mathrm{A}) \;=\; a_1^{d-1}\, a_2^{d-2} \cdots a_{d-1} \;+\; \ldots$$

The polynomials g and $\det(\mathrm{A})$ are each homogeneous of degree

$$\binom{d}{2} \;=\; (d-1) \;+\; (d-2) \;+\; \ldots \;+\; 1$$

so if g divides $\det(\mathrm{A})$, they must be the same polynomial.

We now work modulo $a_i - a_j$. If we substitute $a_i = a_j$ in the matrix A, then the ith and jth rows of A become the same, so

$$\det(\mathrm{A}) \equiv 0 \mod a_i - a_j$$

Therefore $\det(\mathrm{A})$ is a multiple of $a_i - a_j$ for each $i < j$, so g divides $\det(\mathrm{A})$. □

Example 6.14. Let $d = 3$, and let a, b, c be values in F. Then

$$\mathrm{A} \;=\; \begin{bmatrix} a^2 & a & 1 \\ b^2 & b & 1 \\ c^2 & c & 1 \end{bmatrix}$$

and

$$(a-b)(a-c)(b-c) \;=\; a^2b + b^2c + ac^2 - a^2c - ab^2 - bc^2 \;=\; \det(\mathrm{A})$$

Notice that of the eight possible terms in this product, two are $abc - abc$, leaving the six shown.

When two points come together, say a_1 and a_2, the determinant of A vanishes and we can no longer solve for the polynomial f using the system of equations 6.6. In other words, f is no longer determined by its values at $a_1, a_2, \ldots, a_d$ because the value at $a_1 = a_2$ is only useful once; we need another value. However, the polynomial f is determined by these values and the derivative $f'(a_1)$.

To find the polynomial $f(x)$ determined by the values

$$f(a_1),\; f'(a_1),\; f(a_3),\; \ldots,\; f(a_d) \tag{6.8}$$

we want to solve the system of equations

$$\begin{bmatrix} a_1^{d-1} & \ldots & a_1^2 & a_1 & 1 \\ (d-1)a_1^{d-2} & \ldots & 2a_1 & 1 & 0 \\ a_3^{d-1} & \ldots & a_3^2 & a_3 & 1 \\ \vdots & & \vdots & \vdots & \\ a_d^{d-1} & \ldots & a_d^2 & a_d & 1 \end{bmatrix} \begin{bmatrix} c_{d-1} \\ \vdots \\ c_2 \\ c_1 \\ c_0 \end{bmatrix} \;=\; \begin{bmatrix} f(a_1) \\ f'(a_1) \\ f(a_3) \\ \vdots \\ f(a_d) \end{bmatrix} \tag{6.9}$$

where the second row is the equation

$$(d-1)\, c_{d-1}\, a_1^{d-2} \;+\; \ldots \;+\; 2\, c_2\, a_1 \;+\; c_1 \;=\; f'(a_1)$$

We can use dual numbers to understand this matrix. Working modulo ϵ^2, we can transform 6.6 into 6.9 by substituting $a_2 = a_1 + \epsilon$, applying equation 6.1, subtracting the first row from the second, and factoring out an ϵ from the second row.

Making this same substitution into g yields

$$g(a_1, a_1 + \epsilon, a_3, \ldots, a_d) \;=\; \epsilon \tfrac{\partial g}{\partial a_2}(a_1, a_1, a_3, \ldots, a_d)$$

because $g(a_1, a_1, a_3, \ldots, a_d)$ vanishes. Writing this out, we have

$$\begin{aligned}\prod_{i<j}(a_i - a_j) \;&=\; (a_1 - (a_1+\epsilon)) \prod_{j=3}^{d} (a_1 - a_j)((a_1+\epsilon) - a_j) \prod_{3\le i<j\le d} (a_i - a_j) \\ &=\; -\epsilon \prod_{j=3}^{d} (a_1 - a_j)^2 \prod_{3\le i<j\le d} (a_i - a_j)\end{aligned}$$

Making this substitution into $\det(\mathrm{A})$ yields

$$\begin{vmatrix} a_1^{d-1} & \ldots & a_1^2 & a_1 & 1 \\ (a_1+\epsilon)^{d-1} & \ldots & (a_1+\epsilon)^2 & a_1+\epsilon & 1 \\ a_3^{d-1} & \ldots & a_3^2 & a_3 & 1 \\ \vdots & & \vdots & \vdots & \vdots \\ a_d^{d-1} & \ldots & a_d^2 & a_d & 1 \end{vmatrix} \;=\; \epsilon \begin{vmatrix} a_1^{d-1} & \ldots & a_1^2 & a_1 & 1 \\ (d-1)a_1^{d-2} & \ldots & 2a_1 & 1 & 0 \\ a_3^{d-1} & \ldots & a_3^2 & a_3 & 1 \\ \vdots & & \vdots & \vdots & \vdots \\ a_d^{d-1} & \ldots & a_d^2 & a_d & 1 \end{vmatrix}$$

following the same steps that transform 6.6 into 6.9.

Comparing these equivalences, we see that the determinant of the system of equations 6.9 is nonzero whenever $a_1, a_3, \ldots, a_d$ are distinct values, so the polynomial $f(x)$ is uniquely determined by the values 6.8. The general case follows the same pattern.

We could be forgiven for expecting proposition 6.13 to be useless in the situation where the value a_1 repeats. However, the ring of dual numbers provides us with a device for viewing the repeated value a_1 as the two distinct values a_1 and $a_1 + \epsilon$, allowing us to apply proposition 6.13 after all. Intuitively, the two values a_1 and $a_1 + \epsilon$ are right next to each other, differing infinitesimally by ϵ. We have replaced differentiation with algebra[2]. In *algebraic geometry*, one generalizes this idea to model multiplicities of higher dimensional sets, as an algebraic counterpart to multivariable calculus.

We will see a version of this phenomenon in chapter 7. Ordinarily, a function applied to a matrix is determined by its values at the *eigenvalues* of the matrix. However, when two eigenvalues come together, we also need to consider the derivative at the repeated eigenvalue. This is parallel to polynomial interpolation; our decomposition of the matrix will sprout a nilpotent matrix playing the role of ϵ here.

[2]Few mathematicians see this as an even trade, but there are plenty of takers on both sides.

Lagrange interpolation

The complexity of the Vandermonde matrix is due to the fact that while

$$\{x^{d-1},\ x^{d-2},\ \ldots,\ x, 1\}$$

may be the obvious choice of a basis for the polynomials of degree $< d$, it is poorly adapted to the problem of interpolating polynomials from their values. We would be much happier if the Vandermonde matrix were a diagonal matrix, but it isn't.

The idea of *Lagrange interpolation* is to choose a basis for the polynomials of degree $< d$ that diagonalizes the problem of interpolating polynomials from their values. Indeed, the method is so simple that none of these big words are necessary; one can apply this method knowing next to nothing[3]. The virtue of understanding Lagrange interpolation in the language of linear algebra is that it represents a common pattern we want to reapply. Whenever possible, it pays to find a basis that diagonalizes the problem one is studying.

Example 6.15. Let $f(x)$ be the degree two polynomial determined by the values $f(a)$, $f(b)$, and $f(c)$. Then

$$f(x) \;=\; f(a)\,\frac{(x-b)(x-c)}{(a-b)(a-c)} \;+\; f(b)\,\frac{(x-a)(x-c)}{(b-a)(b-c)} \;+\; f(c)\,\frac{(x-a)(x-b)}{(c-a)(c-b)}$$

It is clear by inspection that this expression yields the desired values; substituting $x = a, b$ or c, all but one quotient vanishes, and the remaining quotient cancels out to 1. We have expressed the polynomial $f(x)$ in *Lagrange form*.

Conceptually, we have chosen the basis

$$g_1(x) = \frac{(x-b)(x-c)}{(a-b)(a-c)},\ \ g_2(x) = \frac{(x-a)(x-c)}{(b-a)(b-c)},\ \ g_3(x) = \frac{(x-a)(x-b)}{(c-a)(c-b)}$$

for the space of polynomials of degree < 3. Now, any such polynomial can be expressed as a linear combination of g_1, g_2, and g_3. We want to solve for r_1, r_2, and r_3 in the system of equations

$$\begin{aligned} r_1\,g_1(a) \;+\; r_2\,g_2(a) \;+\; r_3\,g_3(a) \;&=\; f(a) \\ r_1\,g_1(b) \;+\; r_2\,g_2(b) \;+\; r_3\,g_3(b) \;&=\; f(b) \\ r_1\,g_1(c) \;+\; r_2\,g_2(c) \;+\; r_3\,g_3(c) \;&=\; f(c) \end{aligned}$$

[3] I remember thinking of Lagrange interpolation in high school, only to be gently scolded by my teacher that this was not a new idea. While some math students worry about when they will start to do original work, one generally starts to think of new ideas long after the distinction stops mattering. It is simply fun to think of things for oneself.

However, this leads to the matrix equation

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \\ r_3 \end{bmatrix} = \begin{bmatrix} f(a) \\ f(b) \\ f(c) \end{bmatrix}$$

from which we can simply read off the solution.

6.3 Interpolation mod $\mathbf{p}(\mathbf{x})$

Suppose that the polynomial $p(x)$ factors into the linear factors

$$p(x) = (x - a_1) \cdots (x - a_d)$$

for distinct roots $a_1, \ldots, a_d$, and that we are working modulo $p(x)$. Then all of the results of this chapter are applicable. This is the setting that most closely resembles working with a diagonalizable matrix A.

Now, any polynomial $f(x)$ is equivalent mod $p(x)$ to a polynomial of degree $< d$, so any polynomial $f(x)$ can be interpolated from its values $f(a_1), \ldots, f(a_d)$.

Example 6.16. Let

$$p(x) = (x-2)(x-3)(x-4)$$

Then $f(x) = x^k$ can be interpolated mod $p(x)$ from its values at the roots 2, 3, and 4. By Lagrange interpolation,

$$x^k \equiv 2^k \frac{(x-3)(x-4)}{(2-3)(2-4)} + 3^k \frac{(x-2)(x-4)}{(3-2)(3-4)} + 4^k \frac{(x-2)(x-3)}{(4-2)(4-3)} \mod p(x)$$

This polynomial is the unique degree < 3 polynomial determined by the values $f(2) = 2^k$, $f(3) = 3^k$, and $f(4) = 4^k$. On the other hand, f is equivalent mod p to its remainder under division by p, which is a polynomial of degree < 3. Therefore, $f(x)$ is this polynomial.

Recall the geometric interpretation of polynomial modular arithmetic. In the above example, only the values 2, 3, and 4 matter. Working mod $p(x)$, our domain is the set $\{2, 3, 4\}$, as if the rest of our field F isn't even there. As soon as we have found a function that agrees with f on this domain, we have found f.

In fact, the distinction between polynomials and other functions breaks down when our domain is a finite set; any function is equivalent mod $p(x)$ to a polynomial of degree $< d$, which can be found by Lagrange interpolation from its values on the roots of p.

In particular, f can be a function of several variables, where the domain of one of the variables is the set of roots of p.

Example 6.17. Let $f(x,t) = \mathrm{e}^{xt}$ be the exponential function, and again take

$$p(x) \;=\; (x-2)(x-3)(x-4)$$

Working mod $p(x)$, the function f has the form

$$f : \{2,3,4\} \times \mathbb{R} \to \mathbb{R}$$

By Lagrange interpolation, we have

$$\mathrm{e}^{xt} \;=\; \mathrm{e}^{2t}\,\frac{(x-3)(x-4)}{(2-3)(2-4)} \;+\; \mathrm{e}^{3t}\,\frac{(x-2)(x-4)}{(3-2)(3-4)} \;+\; \mathrm{e}^{4t}\,\frac{(x-2)(x-3)}{(4-2)(4-3)} \tag{6.10}$$

In other words, f is equivalent mod $p(x)$ to the polynomial determined by the list of functions e^{2t}, e^{3t}, e^{4t}.

When a polynomial vanishes at all but one value in the domain $\{2,3,4\}$, multiplying that polynomial by x mod $p(x)$ is multiplication by the remaining value:

$$\begin{aligned}
x\,(x-3)(x-4) \;&=\; 2\,(x-3)(x-4) \;+\; (x-2)(x-3)(x-4) \\
&\equiv\; 2\,(x-3)(x-4) \mod p(x) \\
x\,(x-2)(x-4) \;&\equiv\; 3\,(x-2)(x-4) \mod p(x) \\
x\,(x-2)(x-3) \;&\equiv\; 4\,(x-2)(x-3) \mod p(x)
\end{aligned}$$

Therefore,

$$\begin{aligned}
x\mathrm{e}^{xt} \;&=\; x\left(\frac{\mathrm{e}^{2t}}{2}\,(x-3)(x-4) \;-\; \mathrm{e}^{3t}\,(x-2)(x-4) \;+\; \frac{\mathrm{e}^{4t}}{2}\,(x-2)(x-3)\right) \\
&=\; \mathrm{e}^{2t}\,(x-3)(x-4) \;-\; 3\mathrm{e}^{3t}\,(x-2)(x-4) \;+\; 2\mathrm{e}^{4t}\,(x-2)(x-3) \\
&=\; \tfrac{\partial}{\partial t}\,\mathrm{e}^{xt}
\end{aligned}$$

Thus, our expression for e^{xt} satisfies the differential equation $\frac{\partial}{\partial t}\mathrm{e}^{xt} = x\mathrm{e}^{xt}$.

Let B be the matrix

$$\mathrm{B} \;=\; \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 4 \end{bmatrix}$$

Because B is a diagonal matrix, it acts like the three numbers which are its diagonal entries. These numbers do not interact with each other. In particular, any function of B is simply that function applied to the diagonal entries of B. Therefore,

$$f(\mathrm{B},t) \;=\; \mathrm{e}^{\mathrm{B}t} \;=\; \begin{bmatrix} \mathrm{e}^{2t} & 0 & 0 \\ 0 & \mathrm{e}^{3t} & 0 \\ 0 & 0 & \mathrm{e}^{4t} \end{bmatrix}$$

On the other hand, we have

$$p(\mathrm{B}) \;=\; \begin{bmatrix} p(2) & 0 & 0 \\ 0 & p(3) & 0 \\ 0 & 0 & p(4) \end{bmatrix} \;=\; 0$$

so the ring of polynomial expressions in B behaves like polynomials in x mod $p(x)$. Substituting the matrix B for the variable x in equation 6.10, we get

$$\begin{aligned} \mathrm{e}^{\mathrm{B}t} &= \mathrm{e}^{2t}\,\frac{(\mathrm{B}-3\mathrm{I})(\mathrm{B}-4\mathrm{I})}{(2-3)(2-4)} + \mathrm{e}^{3t}\,\frac{(\mathrm{B}-2\mathrm{I})(\mathrm{B}-4\mathrm{I})}{(3-2)(3-4)} + \mathrm{e}^{4t}\,\frac{(\mathrm{B}-2\mathrm{I})(\mathrm{B}-3\mathrm{I})}{(4-2)(4-3)} \\ &= \mathrm{e}^{2t}\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} + \mathrm{e}^{3t}\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} + \mathrm{e}^{4t}\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \mathrm{e}^{2t} & 0 & 0 \\ 0 & \mathrm{e}^{3t} & 0 \\ 0 & 0 & \mathrm{e}^{4t} \end{bmatrix} \end{aligned}$$

which agrees with our first answer.

Now let A be the matrix

$$\underset{\mathrm{A}}{\begin{bmatrix} 3 & -1 & 1 \\ -1 & 3 & 1 \\ -1 & 1 & 3 \end{bmatrix}} = \underset{\mathrm{C}}{\begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix}} \underset{\mathrm{B}}{\begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 4 \end{bmatrix}} \underset{\mathrm{C}^{-1}}{\begin{bmatrix} 0 & 1 & -1 \\ 1 & -1 & 1 \\ -1 & 1 & 0 \end{bmatrix}}$$

It is also the case that $p(\mathrm{A}) = 0$, because A is similar to B:

$$p(\mathrm{A}) \;=\; p(\mathrm{CBC}^{-1}) \;=\; \mathrm{C}\,p(\mathrm{B})\,\mathrm{C}^{-1} \;=\; 0$$

Therefore, the ring of polynomial expressions in A also behaves like polynomials in x mod $p(x)$. Substituting the matrix A for the variable x in equation 6.10, we get

$$\begin{aligned} \mathrm{e}^{\mathrm{A}t} &= \mathrm{e}^{2t}\,\frac{(\mathrm{A}-3\mathrm{I})(\mathrm{A}-4\mathrm{I})}{(2-3)(2-4)} + \mathrm{e}^{3t}\,\frac{(\mathrm{A}-2\mathrm{I})(\mathrm{A}-4\mathrm{I})}{(3-2)(3-4)} + \mathrm{e}^{4t}\,\frac{(\mathrm{A}-2\mathrm{I})(\mathrm{A}-3\mathrm{I})}{(4-2)(4-3)} \\ &= \mathrm{e}^{2t}\begin{bmatrix} 0 & 1 & -1 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{bmatrix} + \mathrm{e}^{3t}\begin{bmatrix} 1 & -1 & 1 \\ 1 & -1 & 1 \\ 1 & -1 & 1 \end{bmatrix} + \mathrm{e}^{4t}\begin{bmatrix} 0 & 0 & 0 \\ -1 & 1 & 0 \\ -1 & 1 & 0 \end{bmatrix} \end{aligned}$$

which can be checked by expanding out $\mathrm{e}^{\mathrm{A}t} = \mathrm{C}\,\mathrm{e}^{\mathrm{B}t}\,\mathrm{C}^{-1}$.

To apply this method for computing matrix exponentials, it is not necessary to find an explicit similarity to a diagonal matrix with distinct entries, as we had available here. To use Lagrange interpolation, all we need is a polynomial $p(x)$ with distinct roots such that $p(\mathrm{A}) = 0$. If such a polynomial exists, then there are a variety of ways to find one, and all are fair game.

Example 6.18. Let A be the matrix

$$\mathrm{A} \;=\; \begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix}$$

Emboldened by the previous example, we go looking for a polynomial $p(x)$ such that $p(\mathrm{A}) = 0$. Taking a few powers of A, we have

$$\mathrm{A}^0 \;=\; \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathrm{A}^1 \;=\; \begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix}, \quad \mathrm{A}^2 \;=\; \begin{bmatrix} 6 & 5 \\ 10 & 11 \end{bmatrix}$$

These powers of A are linearly dependent;

$$\begin{bmatrix} 6 & 5 \\ 10 & 11 \end{bmatrix} \;=\; 5 \begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix} \;-\; 4 \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

picking the coefficient 5 to get the off-diagonal entries to work, then the coefficient -4 to get the diagonal entries to work. Thus

$$\mathrm{A}^2 \;-\; 5\,\mathrm{A} \;+\; 4\,\mathrm{I} \;=\; 0$$

In other words, $p(\mathrm{A}) = 0$ for the polynomial $p(x)$ given by

$$p(x) \;=\; x^2 - 5x + 4 \;=\; (x-1)(x-4)$$

which has distinct roots. By Lagrange interpolation we have

$$f(x) \;\equiv\; f(1)\,\frac{(x-4)}{(1-4)} \;+\; f(4)\,\frac{(x-1)}{(4-1)} \mod p(x)$$

for any function $f(x)$. For the exponential function $f(x,t) = \mathrm{e}^{xt}$ this yields

$$\mathrm{e}^{\mathrm{A}t} \;=\; \mathrm{e}^t\,\frac{(\mathrm{A} - 4\,\mathrm{I})}{(1-4)} \;+\; \mathrm{e}^{4t}\,\frac{(\mathrm{A} - \mathrm{I})}{(4-1)} \;=\; \frac{\mathrm{e}^t}{3}\begin{bmatrix} 2 & -1 \\ -2 & 1 \end{bmatrix} \;+\; \frac{\mathrm{e}^{4t}}{3}\begin{bmatrix} 1 & 1 \\ 2 & 2 \end{bmatrix}$$

We can check our work by confirming that $\mathrm{e}^{\mathrm{A}t}|_{t=0} = \mathrm{I}$ and $\frac{\partial}{\partial t}\mathrm{e}^{\mathrm{A}t}|_{t=0} = \mathrm{A}$.

Example 6.19. Let A_s be the matrix

$$\mathrm{A}_s \;=\; \begin{bmatrix} 2 & 1 \\ 0 & 2+s \end{bmatrix}$$

and let $\mathrm{A} = \mathrm{A}_0$. We can think of A_s is a parametrized family of matrices, with parameter s. Intuitively, A_s is a movie, animating a *deformation* of the matrix A as s moves away from zero. We want to understand the limit for s near zero. Physicists

ponder the first moments of the universe, just after the big bang. As algebraists, we instead ponder life modulo s^2.

We look for a polynomial $p(x)$ such that $p(\mathrm{A}_s) = 0$. We have

$$\mathrm{A}_s^0 \;=\; \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathrm{A}_s^1 \;=\; \begin{bmatrix} 2 & 1 \\ 0 & 2+s \end{bmatrix}, \quad \mathrm{A}_s^2 \;=\; \begin{bmatrix} 4 & 4+s \\ 0 & (2+s)^2 \end{bmatrix}$$

These powers of A_s are linearly dependent;

$$\begin{bmatrix} 4 & 4+s \\ 0 & (2+s)^2 \end{bmatrix} \;=\; (4{+}s) \begin{bmatrix} 2 & 1 \\ 0 & 2+s \end{bmatrix} \;-\; (4{+}2s) \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

picking the coefficient $4+s$ to get the off-diagonal entries to work, then the coefficient $4+2s$ to get the diagonal entries to work. Thus

$$\mathrm{A}_s^2 \;-\; (4{+}s)\,\mathrm{A}_s \;+\; (4{+}2s)\,\mathrm{I} \;=\; 0$$

In other words, $p(\mathrm{A}_s) = 0$ for the polynomial $p(x)$ given by

$$p(x) \;=\; x^2 \;-\; (4{+}s)\,x \;+\; (4{+}2s) \;=\; (x-2)(x-(2{+}s))$$

which has the distinct roots 2 and $2+s$ when $s \neq 0$. At $s = 0$ the root 2 repeats, and

$$p(x)\,|_{s=0} \;=\; (x-2)^2$$

By Lagrange interpolation we have

$$\begin{aligned} f(x) \;&\equiv\; f(2)\,\frac{(x-(2{+}s))}{(2-(2{+}s))} \;+\; f(2{+}s)\,\frac{(x-2)}{((2{+}s)-2)} \\ &\equiv\; -\,\frac{f(2)}{s}(x-(2{+}s)) \;+\; \frac{f(2{+}s)}{s}(x-2) \\ &\equiv\; f(2) \;+\; \frac{f(2{+}s)-f(2)}{s}(x-2) \mod p(x) \end{aligned}$$

for any function $f(x)$. Taking the limit as $s \to 0$ gives us

$$f(x) \;\equiv\; f(2) \;+\; f'(2)(x-2) \mod (x-2)^2 \tag{6.11}$$

For the exponential function $f(x,t) = \mathrm{e}^{xt}$ this yields

$$\mathrm{e}^{xt} \;\equiv\; \mathrm{e}^{2t} \;+\; t\mathrm{e}^{2t}\,(x-2) \mod (x-2)^2$$

For this formula we have

$$\tfrac{\partial}{\partial t}\,\mathrm{e}^{xt} \;=\; 2\,\mathrm{e}^{2t} \;+\; (x-2)\,\mathrm{e}^{2t} \;+\; 2\,t\mathrm{e}^{2t}\,(x-2) \;=\; x\,\mathrm{e}^{2t} \;+\; 2\,t\mathrm{e}^{2t}\,(x-2)$$

and

$$x\,e^{xt} \;=\; x\,e^{2t} \;+\; x\,te^{2t}\,(x-2)$$

Using

$$x(x-2) \;\equiv\; 2(x-2) \mod (x-2)^2$$

we see that e^{xt} satisfies the differential equation $\frac{\partial}{\partial t}e^{xt} = xe^{xt}$. Substituting A for x gives the formula

$$e^{At} \;=\; e^{2t}\begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix} \;+\; te^{2t}\begin{bmatrix}0 & 1\\ 0 & 0\end{bmatrix}$$

The ring of polynomial expressions in A behaves like polynomials in x mod $(x-2)^2$, so we have already checked this formula. We found this formula by deformation, using the parameter s to jiggle the entries of A a bit so that we could apply Lagrange interpolation.

Taking limits is an analytic process available to us for the fields $\mathbb{R}$ and $\mathbb{C}$, but not necessarily for other fields. Again, algebraic geometry provides tools for making continuity arguments over arbitrary fields. The ring of dual numbers is the simplest incarnation of this approach. Rather than taking a limit to find equation 6.11, we can work mod s^2, mimicking the relation $\epsilon^2 = 0$. Clearing denominators in Lagrange interpolation to avoid division, we have

$$\begin{aligned} s\,f(x) \;&=\; -f(2)(x-(2+s)) \;+\; f(2+s)(x-2)\\ &\equiv\; s\,\big(f(2) \;+\; f'(2)(x-2)\big) \mod s^2 \end{aligned}$$

Canceling s recovers equation 6.11 for any field.

Canceling s is the algebraic analog to taking the limit as $s \to 0$; these two points of view meet up in the calculus exercise

$$\lim_{s\to 0}\frac{s}{s} \;=\; 1$$

Notice that all of the information in our equivalence is carried by the coefficents of s, and would be lost if we substituted $s = 0$. The information is visible mod s^2, but not visible mod s. The polynomial s^2 has zero twice as a root, so working mod s^2 we can view zero as two values, and capture the effect of working with distinct values. This is the usual pattern, working with dual numbers.

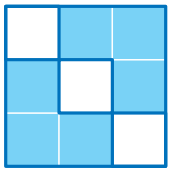

Chapter 7

Functions of matrices

In this chapter, A will always denote a square $n \times n$ matrix representing the linear map $L : V \to V$, where V is an n-dimensional vector space over a field F. For example, one can take F to either be the field of real numbers $\mathbb{R}$, in which case $V = \mathbb{R}^n$, or the field of complex numbers $\mathbb{C}$, in which case $V = \mathbb{C}^n$. In order for polynomials to have full sets of roots, it will sometimes be necessary to move from $\mathbb{R}$ to $\mathbb{C}$.

We consider the problem of computing functions $f(\mathrm{A})$ of the matrix A, such as the matrix exponential $\mathrm{e}^{\mathrm{A}t}$ used to solve systems of linear differential equations. This problem is closely linked to the study of the *eigenvalues* and *eigenvectors* of A.

7.1 Polynomials and power series

The functions that we typically want to apply to matrices will either be polynomials of the form

$$f(x) \;=\; c_d\, x^d \;+\; c_{d-1}\, x^{d-1} \;+\; \ldots \;+\; c_1\, x \;+\; c_0$$

or power series of the form

$$f(x) \;=\; \sum_{k=0}^{\infty} c_k\, x^k$$

with coefficients in our field F.

When $f(x)$ is a polynomial, to compute $f(\mathrm{A})$ for a square matrix A we substitute A for the variable x

$$f(\mathrm{A}) \;=\; c_d\, \mathrm{A}^d \;+\; c_{d-1}\, \mathrm{A}^{d-1} \;+\; \ldots \;+\; c_1\, \mathrm{A} \;+\; c_0$$

yielding a matrix as the result. A formula for powers of the matrix A is helpful but not essential.

When $f(x)$ is a power series, to compute $f(\mathrm{A})$ for a square matrix A we again substitute A for the variable x

$$f(\mathrm{A}) \;=\; \sum_{k=0}^{\infty} c_k\, \mathrm{A}^k$$

yielding a matrix as the result. A formula for powers of the matrix A is now essential, until we develop better methods.

Our favorite power series are those for e^x, $\cos(x)$ and $\sin(x)$, which are about as user-friendly as power series can be, converging for all real numbers. In general, the use of power series poses convergence issues. It turns out that for a square matrix A, the value $f(\mathrm{A})$ is determined by the value of f and possibly of some of its derivatives at the *eigenvalues* of A, to be defined shortly. We ask that our power series and their needed derivatives converge at the eigenvalues of A, and give convergence issues no further thought.

Matrix Exponentials

The exponential function $\mathrm{e}^{at} : \mathbb{R} \to \mathbb{R}$ can be defined by the power series expansion

$$\mathrm{e}^{at} \;=\; \sum_{k=0}^{\infty} \frac{(at)^k}{k!} \tag{7.1}$$

and the other familiar properties of this function, such as

$$\tfrac{\partial}{\partial t}\, \mathrm{e}^{at} \;=\; a\mathrm{e}^{at}$$

follow from this definition, which is used to extend the exponential to a complex function $\mathrm{e}^{at} : \mathbb{C} \to \mathbb{C}$. We use this definition to extend the exponential to a matrix-valued function $\mathrm{e}^{\mathrm{A}t} : \mathbb{R} \to \mathbb{R}^{n^2}$ with the corresponding property

$$\tfrac{\partial}{\partial t}\, \mathrm{e}^{\mathrm{A}t} \;=\; \mathrm{A}\mathrm{e}^{\mathrm{A}t}$$

Example 7.1. Suppose that the matrix similarity $\mathrm{A} = \mathrm{C}\mathrm{B}\mathrm{C}^{-1}$ is given to us, where B is in as simple a form as possible:

$$\underset{\mathrm{A}}{\begin{bmatrix} 1 & 1 & 0 \\ -1 & 3 & 0 \\ 1 & -1 & 3 \end{bmatrix}} \;=\; \underset{\mathrm{C}}{\begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}} \; \underset{\mathrm{B}}{\begin{bmatrix} 2 & 1 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{bmatrix}} \; \underset{\mathrm{C}^{-1}}{\begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 1 & -1 & 1 \end{bmatrix}}$$

We would like to find the matrix exponential $\mathrm{e}^{\mathrm{A}t}$ of A.

By multiplying out a few powers of B we find the formula

$$B^k = \begin{bmatrix} 2^k & k\,2^{k-1} & 0 \\ 0 & 2^k & 0 \\ 0 & 0 & 3^k \end{bmatrix}$$

which can be proved by induction. Substituting B^k for a in equation 7.1 we get

$$\begin{aligned} e^{Bt} = \sum_{k=0}^{\infty} \frac{B^k t^k}{k!} &= \sum_{k=0}^{\infty} \frac{t^k}{k!} \begin{bmatrix} 2^k & k\,2^{k-1} & 0 \\ 0 & 2^k & 0 \\ 0 & 0 & 3^k \end{bmatrix} \\ &= \begin{bmatrix} \sum_{k=0}^{\infty} \frac{2^k t^k}{k!} & \sum_{k=0}^{\infty} \frac{k\,2^{k-1} t^k}{k!} & 0 \\ 0 & \sum_{k=0}^{\infty} \frac{2^k t^k}{k!} & 0 \\ 0 & 0 & \sum_{k=0}^{\infty} \frac{3^k t^k}{k!} \end{bmatrix} \\ &= \begin{bmatrix} e^{2t} & te^{2t} & 0 \\ 0 & e^{2t} & 0 \\ 0 & 0 & e^{3t} \end{bmatrix} \end{aligned}$$

Changing back to our original coordinates, we get

$$\begin{aligned} e^{At} = C\,e^{Bt}\,C^{-1} &= \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} e^{2t} & te^{2t} & 0 \\ 0 & e^{2t} & 0 \\ 0 & 0 & e^{3t} \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 1 & -1 & 1 \end{bmatrix} \\ &= \begin{bmatrix} e^{2t} - te^{2t} & te^{2t} & 0 \\ -te^{2t} & e^{2t} + te^{2t} & 0 \\ e^{3t} - e^{2t} & e^{2t} - e^{3t} & e^{3t} \end{bmatrix} \end{aligned}$$

This answer is both hard on the eyes, and needlessly tedious to multiply out. Its one virtue is that it could be burying a mistake a grader might never catch.

With a little preparation, we can put this answer in a much nicer form. Write B as the linear combination

$$B = 2 \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} + 3 \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

so

$$e^{Bt} = e^{2t} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} + te^{2t} \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} + e^{3t} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Now write

$$\begin{aligned}
\mathrm{A}_1 &= \mathrm{C}\begin{bmatrix}1&0&0\\0&1&0\\0&0&0\end{bmatrix}\mathrm{C}^{-1} = \begin{bmatrix}1&0&0\\0&1&0\\-1&1&0\end{bmatrix}\\
\mathrm{N}_1 &= \mathrm{C}\begin{bmatrix}0&1&0\\0&0&0\\0&0&0\end{bmatrix}\mathrm{C}^{-1} = \begin{bmatrix}-1&1&0\\-1&1&0\\0&0&0\end{bmatrix}\\
\mathrm{A}_2 &= \mathrm{C}\begin{bmatrix}0&0&0\\0&0&0\\0&0&1\end{bmatrix}\mathrm{C}^{-1} = \begin{bmatrix}0&0&0\\0&0&0\\1&-1&1\end{bmatrix}
\end{aligned}$$

The similarities $\mathrm{I} = \mathrm{CIC}^{-1}$, $\mathrm{A} = \mathrm{CBC}^{-1}$, $\mathrm{A}^k = \mathrm{CB}^k\mathrm{C}^{-1}$, and $\mathrm{e}^{\mathrm{A}t} = \mathrm{Ce}^{\mathrm{B}t}\mathrm{C}^{-1}$ can be expanded in terms of A_1, N_1, and A_2. We get

$$\begin{aligned}
\mathrm{I} &= \mathrm{A}_1 + + \mathrm{A}_2\\
\mathrm{A} &= 2\,\mathrm{A}_1 + \mathrm{N}_1 + 3\,\mathrm{A}_2\\
\mathrm{A}^k &= 2^k\,\mathrm{A}_1 + k\,2^{k-1}\,\mathrm{N}_1 + 3^k\,\mathrm{A}_2\\
\mathrm{e}^{\mathrm{A}t} &= \mathrm{e}^{2t}\,\mathrm{A}_1 + t\mathrm{e}^{2t}\,\mathrm{N}_1 + \mathrm{e}^{3t}\,\mathrm{A}_2
\end{aligned} \tag{7.2}$$

This expresses the matrix exponential $\mathrm{e}^{\mathrm{A}t}$ as

$$\mathrm{e}^{\mathrm{A}t} = \mathrm{e}^{2t}\begin{bmatrix}1&0&0\\0&1&0\\-1&1&0\end{bmatrix} + t\mathrm{e}^{2t}\begin{bmatrix}-1&1&0\\-1&1&0\\0&0&0\end{bmatrix} + \mathrm{e}^{3t}\begin{bmatrix}0&0&0\\0&0&0\\1&-1&1\end{bmatrix}$$

Equations 7.2 are instances of the pattern

$$f(\mathrm{A}) = f(2)\,\mathrm{A}_1 + f'(2)\,\mathrm{N}_1 + f(3)\,\mathrm{A}_2$$

for the functions $f(x) = 1$, $f(x) = x$, $f(x) = x^k$, and $f(x) = \mathrm{e}^{xt}$. If we apply this pattern to the function $f(x) = \sqrt{x}$, we get

$$\sqrt{\mathrm{A}} = \sqrt{2}\,A_1 + \frac{1}{2\sqrt{2}}\,N_1 + \sqrt{3}\,A_2$$

This formula can be checked by squaring:

$$\begin{aligned}
\left(\sqrt{\mathrm{A}}\right)^2 &= \left(\sqrt{2}\,A_1 + \frac{1}{2\sqrt{2}}\,N_1 + \sqrt{3}\,A_2\right)^2\\
&= 2\,{\mathrm{A}_1}^2 + \frac{1}{2}\left(\mathrm{A}_1\mathrm{N}_1 + \mathrm{N}_1\mathrm{A}_1\right) + 3\,{\mathrm{A}_2}^2\\
&\quad + \frac{1}{8}\,{\mathrm{N}_1}^2 + \frac{\sqrt{3}}{2\sqrt{2}}\left(\mathrm{A}_2\mathrm{N}_1 + \mathrm{N}_1\mathrm{A}_2\right) + \sqrt{6}\left(\mathrm{A}_1\mathrm{A}_2 + \mathrm{A}_2\mathrm{A}_1\right)\\
&= 2\,\mathrm{A}_1 + \mathrm{N}_1 + 3\,\mathrm{A}_2 = \mathrm{A} \quad \checkmark
\end{aligned}$$

We have noticed and used the relations

$$\begin{array}{ccc} {A_1}^2 = A_1 & {A_2}^2 = A_2 & {N_1}^2 = 0 \\ A_1N_1 = N_1A_1 = N_1 & A_2N_1 = N_1A_2 = 0 & A_1A_2 = A_2A_1 = 0 \end{array}$$

This example exhibits behavior typical of the general case of an $n \times n$ matrix A.

It would appear in this example that the values 2 and 3 and the matrices A_1, N_1, and A_2 play a fundamental role in the theory of functions of the matrix A. We have found these values and matrices by ad hoc means, and that was only after the similarity $A = CBC^{-1}$ was handed to us on a silver platter. This similarity is nontrivial to work out from scratch.

We would like to understand these values and matrices better, and work out faster ways to compute them. That is the goal of this chapter.

Eigenvalues

Eigenvalues are key to finding matrix similarities such as the similarity $A = CBC^{-1}$ used in example 7.1.

Definition 7.2. Let A be a square matrix. If

$$Av = \lambda v$$

for some scalar λ and nonzero vector v, then we say that λ is an *eigenvalue* of A, with *eigenvector* v.

The matrix A acts on v like multiplication by λ. We can think of v as a "stretch direction" for A, with stretching factor λ. *Eigen* is a German word, roughly suggesting that these values belong to the matrix A.

In the simplest case, we can find n distinct eigenvalues $\lambda_1, \ldots, \lambda_n$ for the matrix A, and we can find n matrices $A_1, \ldots, A_n$ such that

$$\begin{array}{rcl} I & = & A_1 + \ldots + A_n \\ A & = & \lambda_1 A_1 + \ldots + \lambda_n A_n \end{array}$$

and we are able to compute the function f of A as

$$f(A) = f(\lambda_1) A_1 + \ldots + f(\lambda_n) A_n$$

This is wonderful. Under function evaluation, the matrix A behaves like the list of eigenvalues $\lambda_1, \ldots, \lambda_n$ with the helper matrices $A_1, \ldots, A_n$ divvying up the effect.

In exceptional cases which do arise in practice, eigenvalues can repeat, leaving us with $m < n$ distinct eigenvalues $\lambda_1, \ldots, \lambda_m$. We can find matrices $A_1, \ldots, A_m$ and N such that

$$\begin{array}{rcl} I & = & A_1 + \ldots + A_m \\ A & = & \lambda_1 A_1 + \ldots + \lambda_m A_m + N \end{array}$$

where N is a nilpotent matrix; see definition 6.4. Here, $\mathrm{N}^\ell = 0$ for some ℓ bounded by the multiplicity of the eigenvalue of A that repeats the most. One could imagine that N is present in our first version of these formulas, but we have $\ell = 1$ when each eigenvalue of A occurs only once, so N is the zero matrix.

Now, the matrix N and some derivatives of f are involved in the computation of $f(A)$. For example, suppose that in the list $\lambda_1, \ldots, \lambda_n$ of eigenvalues of A, the eigenvalues λ_1 and λ_2 are the same. We say that $\lambda_1 = \lambda_2$ has multiplicity two, and we work instead with the shorter list of distinct eigenvalues $\lambda_1, \lambda_3, \ldots, \lambda_n$. Then

$$f(\mathrm{A}) \;=\; f(\lambda_1)\,\mathrm{A}_1 \;+\; f'(\lambda_1)\,\mathrm{N}_1 \;+\; f(\lambda_3)\,\mathrm{A}_3 \;+\; \ldots \;+\; f(\lambda_n)\,\mathrm{A}_n \tag{7.3}$$

where $\mathrm{N}_1 = \mathrm{A}_1\mathrm{N}$. In other words, when λ_1 occurs twice as an eigenvalue, we need to know both $f(\lambda_1)$ and $f'(\lambda_1)$ to compute $f(\mathrm{A})$.

This is a frequently occurring pattern in mathematics, and it makes sense if we imagine λ_1 and λ_2 moving together as real numbers, as we vary the matrix A. Once these eigenvalues get close to each other, knowing $f(\lambda_1)$ and $f'(\lambda_1)$ is pretty much the same information as knowing $f(\lambda_1)$ and $f(\lambda_2)$; we can use $f'(\lambda_1)$ to estimate $f(\lambda_2)$. Once these eigenvalues come together, our original formula can break, but its limit is our new formula, giving the correct answer.

It turns out that the matrix N_1 is nilpotent, and

$$f(\lambda_1)\,\mathrm{A}_1 \;+\; f'(\lambda_1)\,\mathrm{N}_1$$

is an instance of the pattern that we saw in equation 6.2. The matrix A_1 is a piece of the identity matrix I, and N_1 is the corresponding piece of the nilpotent matrix N.

7.2 The characteristic polynomial

The *characteristic polynomial* p_A of a square matrix A is the polynomial in x defined by

$$\boxed{p_\mathrm{A}(x) \;=\; \det(x\mathrm{I} - \mathrm{A})} \tag{7.4}$$

We are particularly interested in the *roots* of p_A, which are those values λ such that $p_\mathrm{A}(\lambda) = 0$. For these λ, we have

$$\det(\lambda\mathrm{I} - \mathrm{A}) \;=\; 0$$

so $\lambda\mathrm{I} - \mathrm{A}$ is a singular matrix, and we can find a nonzero vector $\boldsymbol{v}$ such that $(\lambda\mathrm{I} - \mathrm{A})\,\boldsymbol{v} = 0$. We have

$$(\lambda\mathrm{I} \;-\; \mathrm{A})\,\boldsymbol{v} \;=\; 0 \;\Leftrightarrow\; \mathrm{A}\,\boldsymbol{v} \;=\; \lambda\,\boldsymbol{v}$$

In other words, if λ is a root of the polynomial p_{A}, then λ is an eigenvalue of A, with eigenvector v. If we can find a basis of eigenvectors for A, then A can be expressed as a diagonal matrix in terms of this basis.

In general, the characteristic polynomial p_{A} can have repeated roots, and we may not be able to find a basis of eigenvectors for A. Nevertheless the matrix equation

$$p_{\mathrm{A}}(\mathrm{A}) \;=\; 0$$

always holds. This identity is known as the *Cayley-Hamilton theorem*.

Many problems involving the matrix A can be solved either by diagonalizing A if possible, or by applying the identity $p_{\mathrm{A}}(\mathrm{A}) = 0$. Either way, understanding p_{A} is essential to working with the matrix A.

In this section, we develop formulas for computing the characteristic polynomial p_{A}.

2×2 matrices

For the matrix

$$\mathrm{A} \;=\; \begin{bmatrix} a & c \\ b & d \end{bmatrix}$$

we have

$$p_{\mathrm{A}}(x) \;=\; \begin{vmatrix} x-a & -c \\ -b & x-d \end{vmatrix} \;=\; x^2 \;-\; (a+d)\,x \;+\; (ad-bc)$$

This yields the formula

$$\boxed{p_{\mathrm{A}}(x) \;=\; x^2 \;-\; \mathrm{trace}(\mathrm{A})\,x \;+\; \det(\mathrm{A})} \tag{7.5}$$

where $\mathrm{trace}(\mathrm{A})$ is the sum of the diagonal entries of A.

One systematic way to carry out this computation is to write

$$\begin{bmatrix} x-a \\ -b \end{bmatrix} \;=\; x \begin{bmatrix} 1 \\ 0 \end{bmatrix} \;-\; \begin{bmatrix} a \\ b \end{bmatrix}$$
$$\begin{bmatrix} x-c \\ -d \end{bmatrix} \;=\; x \begin{bmatrix} 0 \\ 1 \end{bmatrix} \;-\; \begin{bmatrix} c \\ d \end{bmatrix}$$

and to expand $p_{\mathrm{A}}(x)$ by linearity in each column of the determinant:

$$\begin{aligned}
p_{\mathrm{A}}(x) &= \begin{vmatrix} x-a & -c \\ -b & x-d \end{vmatrix} \\
&= x \begin{vmatrix} 1 & -c \\ 0 & x-d \end{vmatrix} - \begin{vmatrix} a & -c \\ b & x-d \end{vmatrix} \\
&= x \left(x \begin{vmatrix} 1 & 0 \\ 0 & 1 \end{vmatrix} - \begin{vmatrix} 1 & c \\ 0 & d \end{vmatrix} \right) - \left(x \begin{vmatrix} a & 0 \\ b & 1 \end{vmatrix} - \begin{vmatrix} a & c \\ b & d \end{vmatrix} \right) \\
&= \begin{vmatrix} 1 & 0 \\ 0 & 1 \end{vmatrix} x^2 - \left(\begin{vmatrix} a & 0 \\ b & 1 \end{vmatrix} + \begin{vmatrix} 1 & c \\ 0 & d \end{vmatrix} \right) x + \begin{vmatrix} a & c \\ b & d \end{vmatrix} \\
&= x^2 - (a+d)\, x + \begin{vmatrix} a & c \\ b & d \end{vmatrix}
\end{aligned}$$

3×3 matrices

For the matrix

$$\mathrm{A} = \begin{bmatrix} a & d & g \\ b & e & h \\ c & f & i \end{bmatrix}$$

we have

$$\begin{aligned}
p_{\mathrm{A}}(x) &= \begin{vmatrix} x-a & -d & -g \\ -b & x-e & -h \\ -c & -f & x-i \end{vmatrix} \\
&= x^3 - (a+e+i)\, x^2 \\
&\quad + ((ae-bd) + (ai-cg) + (ei-fh))\, x \\
&\quad - (aei + bfg + cdh - afh - bdi - ceg)
\end{aligned}$$

This yields the formula

$$\boxed{p_{\mathrm{A}}(x) = x^3 - \operatorname{trace}(\mathrm{A})\, x^2 + \operatorname{trace}(\wedge^2 \mathrm{A})\, x - \det(\mathrm{A})} \tag{7.6}$$

where

$$\operatorname{trace}(\wedge^2 \mathrm{A}) = \begin{vmatrix} a & d \\ b & e \end{vmatrix} + \begin{vmatrix} a & g \\ c & i \end{vmatrix} + \begin{vmatrix} e & h \\ f & i \end{vmatrix}$$

is the sum of the diagonal 2×2 minors of A.

Again, one systematic way to carry out this computation is to write

$$\begin{bmatrix} x-a \\ -b \\ -c \end{bmatrix} = x \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} - \begin{bmatrix} a \\ b \\ c \end{bmatrix}$$
$$\begin{bmatrix} -d \\ x-e \\ -f \end{bmatrix} = x \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} - \begin{bmatrix} d \\ e \\ f \end{bmatrix}$$
$$\begin{bmatrix} -g \\ -h \\ x-i \end{bmatrix} = x \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} - \begin{bmatrix} g \\ h \\ i \end{bmatrix}$$

and to expand $p_{\mathrm{A}}(x)$ by linearity in each column of the determinant:

$$\begin{aligned} p_{\mathrm{A}}(x) &= \begin{vmatrix} x-a & -d & -g \\ -b & x-e & -h \\ -c & -f & x-i \end{vmatrix} \\ &= \begin{vmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{vmatrix} x^3 - \left(\begin{vmatrix} a & 0 & 0 \\ b & 1 & 0 \\ c & 0 & 1 \end{vmatrix} + \begin{vmatrix} 1 & d & 0 \\ 0 & e & 0 \\ 0 & f & 1 \end{vmatrix} + \begin{vmatrix} 1 & 0 & g \\ 0 & 1 & h \\ 0 & 0 & i \end{vmatrix} \right) x^2 \\ &\quad + \left(\begin{vmatrix} a & d & 0 \\ b & e & 0 \\ c & f & 1 \end{vmatrix} + \begin{vmatrix} a & 0 & g \\ b & 1 & h \\ c & 0 & i \end{vmatrix} + \begin{vmatrix} 1 & d & g \\ 0 & e & h \\ 0 & f & i \end{vmatrix} \right) x - \begin{vmatrix} a & d & g \\ b & e & h \\ c & f & i \end{vmatrix} \\ &= x^3 - (a+e+i)\, x^2 + \left(\begin{vmatrix} a & d \\ b & e \end{vmatrix} + \begin{vmatrix} a & g \\ c & i \end{vmatrix} + \begin{vmatrix} e & h \\ f & i \end{vmatrix} \right) x - \begin{vmatrix} a & d & g \\ b & e & h \\ c & f & i \end{vmatrix} \end{aligned}$$

The general case

Let $\wedge^i \mathrm{A}$ denote the matrix of all $i \times i$ minors of A. Then $\mathrm{trace}(\wedge^i \mathrm{A})$ is the sum of the diagonal $i \times i$ minors of A, which are those minors defined using the same rows and columns of A. For example, if

$$\mathrm{A} = \begin{bmatrix} a & d & g \\ b & e & h \\ c & f & i \end{bmatrix}$$

then

$$\wedge^2 \mathrm{A} \;=\; \begin{bmatrix} \begin{vmatrix} a & d \\ b & e \end{vmatrix} & \begin{vmatrix} a & g \\ b & h \end{vmatrix} & \begin{vmatrix} d & g \\ e & h \end{vmatrix} \\ \begin{vmatrix} a & d \\ c & f \end{vmatrix} & \begin{vmatrix} a & g \\ c & i \end{vmatrix} & \begin{vmatrix} d & g \\ f & i \end{vmatrix} \\ \begin{vmatrix} b & e \\ c & f \end{vmatrix} & \begin{vmatrix} b & h \\ c & i \end{vmatrix} & \begin{vmatrix} e & h \\ f & i \end{vmatrix} \end{bmatrix}$$

The 1×1 minors of A are the entries of A, so $\wedge^1 \mathrm{A} = \mathrm{A}$. The unique 3×3 minor of A is $\det(\mathrm{A})$, so $\wedge^3 \mathrm{A}$ is a 1×1 matrix whose sole entry is $\det(\mathrm{A})$. Thus

$$\begin{aligned} \operatorname{trace}(\wedge^1 \mathrm{A}) &= \operatorname{trace}(\mathrm{A}) \\ \operatorname{trace}(\wedge^2 \mathrm{A}) &= \begin{vmatrix} a & d \\ b & e \end{vmatrix} + \begin{vmatrix} a & g \\ c & i \end{vmatrix} + \begin{vmatrix} e & h \\ f & i \end{vmatrix} \\ \operatorname{trace}(\wedge^3 \mathrm{A}) &= \det(\mathrm{A}) \end{aligned}$$

We can now rewrite our formula for $p_{\mathrm{A}}(x)$ in the 3×3 case as

$$p_{\mathrm{A}}(x) \;=\; x^3 \;-\; \operatorname{trace}(\wedge^1 \mathrm{A})\, x^2 \;+\; \operatorname{trace}(\wedge^2 \mathrm{A})\, x^1 \;-\; \operatorname{trace}(\wedge^3 \mathrm{A})\, x^0$$

The general formula for the characteristic polynomial of an $n \times n$ matrix A follows this pattern; it is

$$\boxed{p_{\mathrm{A}}(x) \;=\; x^n \;+\; \sum_{i=1}^{n} (-1)^i \;\operatorname{trace}(\wedge^i \mathrm{A})\; x^{n-i}} \tag{7.7}$$

where $\operatorname{trace}(\wedge^i \mathrm{A})$ is the sum of the diagonal $i \times i$ minors of A. This formula can be established by expanding $p_{\mathrm{A}}(x)$ by linearity in each column of the determinant, exactly as we did for $n = 2$ and $n = 3$.

7.3 Rings and fields

Roots of polynomial equations

Working with polynomials, we must contend with the fact that factorizations of polynomials depend on the coefficients that we allow. For example, the polynomial $x^2 - 2$ can't be factored into linear factors over the rational numbers $\mathbb{Q}$, but it can be factored over the real numbers $\mathbb{R}$ as

$$x^2 - 2 \;=\; (x + \sqrt{2})\,(x - \sqrt{2})$$

Similarly, $x^2 + 1$ can't be factored into linear factors over $\mathbb{R}$, but it can be factored over the complex numbers $\mathbb{C}$ as

$$x^2 + 1 \;=\; (x + \mathrm{i})\,(x - \mathrm{i})$$

where i is the imaginary number $\sqrt{-1}$. Equivalently, the equation $x^2 + 1 = 0$ has no solutions in $\mathbb{R}$, but it has the two solutions $-\mathrm{i}$, i in $\mathbb{C}$.

The *fundamental theorem of algebra* states that every polynomial in one variable can be factored into linear factors over $\mathbb{C}$. Equivalently, if $f(x)$ is a degree d polynomial with complex coefficients, then the equation $f(x) = 0$ always has d complex solutions, where we count each solution as many times as the corresponding linear factor appears in $f(x)$. This is called counting with *multiplicity*. For example, the degree four equation

$$x^4 + x^2 \;=\; (x + \mathrm{i})\,(x - \mathrm{i})\,(x - 0)\,(x - 0) \;=\; 0$$

has the four solutions $-\mathrm{i}$, i, 0, 0 over $\mathbb{C}$, where we count the solution 0 with multiplicity 2.

Our matrices typically have entries in either the integers $\mathbb{Z}$, or in $\mathbb{Q}$, $\mathbb{R}$, or $\mathbb{C}$. The integers $\mathbb{Z}$ are an example of a *ring*, a number system in which one can't always divide. When we work with integer matrices, we do divide as necessary, moving to the rational numbers $\mathbb{Q}$ when these divisions create fractions. A similar issue arises in working with matrices with polynomial entries, where divisions can create rational functions.

$\mathbb{Q}$, $\mathbb{R}$, and $\mathbb{C}$ are examples of a *field*, a number system in which one can always divide by any nonzero number. For most of the operations of linear algebra, we are implicitly working over a field, even if our matrix entries start out in a smaller ring contained in that field, and generally any field is as good as any other field.

$\mathbb{C}$ is an example of an *algebraically closed field*, a field over which every polynomial in one variable can be factored into linear factors. In a course on *modern algebra*, one proves that every field can be embedded in an algebraically closed field, generalizing the fact that $\mathbb{Q}$ and $\mathbb{R}$ are contained in $\mathbb{C}$.

To compute functions of matrices, we often end up working with matrices with entries in an algebraically closed field, because we need the characteristic polynomial p_{A} to have a full set of roots. If A is an $n \times n$ matrix, then we want p_{A} to factor into n linear factors, so $p_{\mathrm{A}}(x) = 0$ has n solutions counted with multiplicity. Starting with an integer, rational, or real matrix, it may be necessary to move to the complex numbers to find all these roots.

Euler's formula

We're nearly out of the woods, working with real numbers: Every polynomial in one variable with coefficients in $\mathbb{R}$ can be factored into linear and quadratic factors.

In calculus one studies exponential and trigonometric functions, corresponding to solutions to certain degree one and two differential equations. Ever wonder why there isn't some ornate theory that comes next, studying functions that correspond to solutions to certain degree three differential equations? Because real polynomials factor into linear and quadratic factors, we can reduce to the study of exponential and trigonometric functions.

Over the complex numbers, we have *Euler's formula*

$$\boxed{e^{ix} = \cos(x) + i\,\sin(x)} \tag{7.8}$$

which can easily be established by comparing power series expansions. Because complex polynomials factor into linear factors, we can reduce to the study of exponential functions alone. Euler's formula expresses how to reduce questions involving trigonometric functions to ones involving complex exponentials.

For example, what were the addition laws for sine and cosine again? We can form the complex exponential

$$\begin{aligned}
\cos(a+b) + i\,\sin(a+b) &= e^{i(a+b)} \\
&= e^{ia}\,e^{ib} \\
&= (\cos(a) + i\sin(a))(\cos(b) + i\sin(b)) \\
&= (\cos(a)\cos(b) - \sin(a)\sin(b)) \\
&\quad + i\,(\cos(a)\sin(b) + \cos(b)\sin(a))
\end{aligned}$$

and by taking real and imaginary parts, we get

$$\begin{aligned}
\cos(a+b) &= \cos(a)\cos(b) - \sin(a)\sin(b) \\
\sin(a+b) &= \cos(a)\sin(b) + \cos(b)\sin(a)
\end{aligned}$$

In general, Euler's formula is a radical simplification of the trigonometric identities. One can understand Euler's formula as factoring the Pythagorean trigonometric identity as a difference of squares:

$$\begin{aligned}
\cos^2(x) + \sin^2(x) &= (\cos(x) + i\sin(x))\,(\cos(x) - i\sin(x)) \\
&= e^{ix}e^{-ix} = 1
\end{aligned}$$

This approach really comes into its own for solving integration problems. Computer programs for symbolic integration work over $\mathbb{C}$ in order to avoid the intricacies of trigonometric integration; a broad swath of disparate integration problems can be viewed uniformly as rational functions of matrix exponentials. Many people also make this migration to $\mathbb{C}$, with the same motivation.

We are at a similar juncture in our study of linear algebra. One could develop a complete theory of functions of real matrices without ever involving $\mathbb{C}$, by reducing matrices to 1×1 and 2×2 blocks corresponding to linear and quadratic factors of $p_A(x)$. Instead, we choose to move to $\mathbb{C}$ as necessary, so a linear factorization is always available. This leads to a simpler theory.

Taking this point of view, by far the most important function of a matrix A is the matrix exponential e^{At}, which is used to solve systems of linear differential equations. One can also compute trigonometric functions of matrices, but with the availability of the complex numbers their importance is diminished. We again can use Euler's formula to reduce trigonometry to exponentiation:

$$\begin{bmatrix} 1 & i \\ 1 & -i \end{bmatrix} \begin{bmatrix} \cos(x) \\ \sin(x) \end{bmatrix} = \begin{bmatrix} \cos(x) + i \sin(x) \\ \cos(x) - i \sin(x) \end{bmatrix} = \begin{bmatrix} e^{ix} \\ e^{-ix} \end{bmatrix}$$

so

$$\begin{bmatrix} \cos(x) \\ \sin(x) \end{bmatrix} = \begin{bmatrix} 1 & i \\ 1 & -i \end{bmatrix}^{-1} \begin{bmatrix} e^{ix} \\ e^{-ix} \end{bmatrix}$$

This leads in particular to the matrix identities

$$\begin{aligned} \cos(At) &= (e^{iAt} + e^{-iAt})/2 \\ \sin(At) &= (e^{iAt} - e^{-iAt})/2i \end{aligned}$$

which allow us to compute trigonometric functions of matrices in terms of matrix exponentials.

In practice, many polynomials do factor into linear factors over $\mathbb{R}$. We will generally choose examples where $p_A(x)$ factors over $\mathbb{Z}$, while reserving the option to move to $\mathbb{R}$ or $\mathbb{C}$ as necessary.

7.4 Diagonal and triangular forms

This will be a section on diagonal and triangular forms. Over an algebraically closed field, a matrix whose characteristic polynomial has distinct roots is similar to a diagonal matrix, and any matrix is similar to a triangular matrix.

7.5 The Cayley-Hamilton Theorem

It is a remarkable and useful fact that a square matrix A satisfies its own characteristic polynomial. Namely,

$$p_A(A) = 0$$

In other words, if A is an $n \times n$ matrix and

$$p_A(x) = x^n + c_1 x^{n-1} + \ldots + c_{n-1} x + c_n$$

then substituting the matrix A for the variable x yields the zero matrix:

$$\boxed{A^n + c_1 A^{n-1} + \ldots + c_{n-1} A + c_n I = 0} \tag{7.9}$$

We often like to think of matrices as if they are single values, a kind of generalized number. Bearing in mind the caveats that order of multiplication matters, and that there are many singular matrices in place of the unique number zero, much can be learned about matrices by manipulating algebraic expressions involving matrix values.

The significance of the Cayley-Hamilton theorem is that *all* computations involving a matrix A can be viewed as taking place modulo p_A. Using this result, we will develop methods that make short work of typical eigenvalue problems such as computing matrix exponentials.

A formula for the inverse

Equation 7.9 has many applications; one is to express the inverse of an invertible matrix A as a polynomial in A. We have

$$c_n = p_A(0) = \det(-A) = (-1)^n \det(A)$$

If $c_n \neq 0$, then we can rewrite equation 7.9 as

$$c_n I = -(A^{n-1} + c_1 A^{n-2} + \ldots + c_{n-1} I)\, A$$

so

$$\boxed{A^{-1} = -\frac{1}{c_n}(A^{n-1} + c_1 A^{n-2} + \ldots + c_{n-1} I)} \tag{7.10}$$

Example

Let

$$A = \begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix}$$

By equation 7.5,

$$\begin{aligned} p_A(x) &= x^2 - \operatorname{trace}(A)\, x + \det(A) \\ &= x^2 - (2+3)\, x + (2 \cdot 3 - 2 \cdot 1) \\ &= x^2 - 5x + 4 \\ &= (x-1)(x-4) \end{aligned}$$

so by equation 7.9,

$$\begin{aligned} p_A(A) &= A^2 - 5A + 4I \\ &= \begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix}^2 - 5\begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix} + 4\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} 6 & 5 \\ 10 & 11 \end{bmatrix} - \begin{bmatrix} 10 & 5 \\ 10 & 15 \end{bmatrix} + \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \end{aligned}$$

and

$$\begin{aligned} p_A(A) &= (A - I)(A - 4I) \\ &= \left(\begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix} - \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}\right)\left(\begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix} - 4\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}\right) \\ &= \begin{bmatrix} 1 & 1 \\ 2 & 2 \end{bmatrix}\begin{bmatrix} -2 & 1 \\ 2 & -1 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \end{aligned}$$

and by equation 7.10,

$$\begin{aligned} A^{-1} &= -\frac{1}{4}(A - 5I) \\ &= \frac{1}{4}\left(-\begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix} + \begin{bmatrix} 5 & 0 \\ 0 & 5 \end{bmatrix}\right) \\ &= \begin{bmatrix} 3 & -1 \\ -2 & 2 \end{bmatrix}/4 \end{aligned}$$

One confirms that

$$A\,A^{-1} = \begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix}\begin{bmatrix} 3 & -1 \\ -2 & 2 \end{bmatrix}/4 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = I$$

2 × 2 matrices

One can confirm equation 7.9 by direct computation, for a general 2 × 2 matrix. Let

$$A = \begin{bmatrix} a & c \\ b & d \end{bmatrix}$$

Then

$$p_A\left(\begin{bmatrix} a & c \\ b & d \end{bmatrix}\right) = \begin{bmatrix} a & c \\ b & d \end{bmatrix}^2 - (a+d)\begin{bmatrix} a & c \\ b & d \end{bmatrix} + (ad - bc)\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

$$= \begin{bmatrix} a^2 + bc & ac + cd \\ ab + bd & bc + d^2 \end{bmatrix} - \begin{bmatrix} a^2 + ad & ac + cd \\ ab + bd & ad + d^2 \end{bmatrix} + \begin{bmatrix} ad - bc & 0 \\ 0 & ad - bc \end{bmatrix}$$

$$= \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$$

Diagonal matrices

One can confirm equation 7.9 by direct computation, for a diagonal matrix. We demonstrate using a 3×3 diagonal matrix. Let

$$\mathrm{A} = \begin{bmatrix} a & & \\ & b & \\ & & c \end{bmatrix}$$

Then

$$p_{\mathrm{A}}(x) = \begin{vmatrix} x-a & & \\ & x-b & \\ & & x-c \end{vmatrix} = (x-a)(x-b)(x-c)$$

so

$$\begin{aligned} p_{\mathrm{A}}(\mathrm{A}) &= (\mathrm{A}-a\mathrm{I})(\mathrm{A}-b\mathrm{I})(\mathrm{A}-c\mathrm{I}) \\ &= \begin{bmatrix} a-a & & \\ & b-a & \\ & & c-a \end{bmatrix} \begin{bmatrix} a-b & & \\ & b-b & \\ & & c-b \end{bmatrix} \begin{bmatrix} a-c & & \\ & b-c & \\ & & c-c \end{bmatrix} \\ &= \begin{bmatrix} p_{\mathrm{A}}(a) & & \\ & p_{\mathrm{A}}(b) & \\ & & p_{\mathrm{A}}(c) \end{bmatrix} = \begin{bmatrix} 0 & & \\ & 0 & \\ & & 0 \end{bmatrix} \end{aligned}$$

In general, any polynomial function of a diagonal matrix is the diagonal matrix obtained by applying that function to each diagonal entry. Since $p_{\mathrm{A}}(x)$ vanishes on each diagonal entry of A, we have $p_{\mathrm{A}}(\mathrm{A}) = 0$.

Triangular matrices

One can also confirm equation 7.9 by direct computation, for a triangular matrix. We demonstrate using a 3×3 triangular matrix. Let

$$\mathrm{A} = \begin{bmatrix} a & d & f \\ & b & e \\ & & c \end{bmatrix}$$

Then

$$p_{\mathrm{A}}(x) = \begin{vmatrix} x-a & -d & -f \\ & x-b & -e \\ & & x-c \end{vmatrix} = (x-a)(x-b)(x-c)$$

so

$$\begin{aligned} p_{\mathrm{A}}(\mathrm{A}) &= (\mathrm{A}-a\mathrm{I})(\mathrm{A}-b\mathrm{I})(\mathrm{A}-c\mathrm{I}) \\ &= \begin{bmatrix} 0 & d & f \\ & b-a & e \\ & & c-a \end{bmatrix} \begin{bmatrix} a-b & d & f \\ & 0 & e \\ & & c-b \end{bmatrix} \begin{bmatrix} a-c & d & f \\ & b-c & e \\ & & 0 \end{bmatrix} \end{aligned}$$

The pattern of zeros here forces this product to be zero: Let $*$ denote a matrix entry that is not known to be zero. Then we have the pattern

$$\begin{aligned} p_A(A) &= \begin{bmatrix} 0 & * & * \\ & * & * \\ & & * \end{bmatrix} \begin{bmatrix} * & * & * \\ & 0 & * \\ & & * \end{bmatrix} \begin{bmatrix} * & * & * \\ & * & * \\ & & 0 \end{bmatrix} \\ &= \begin{bmatrix} 0 & * & * \\ & * & * \\ & & * \end{bmatrix} \begin{bmatrix} * & * & * \\ & 0 & 0 \\ & & 0 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 & 0 \\ & 0 & 0 \\ & & 0 \end{bmatrix} \end{aligned}$$

Algebraically closed fields

If our field F is algebraically closed, then any matrix A is similar to a triangular matrix B:

$$A = C\,B\,C^{-1}$$

for some change of basis matrix C. If A and B are similar, then $xI - A$ and $xI - B$ are similar, so they have the same determinant:

$$\begin{aligned} \det(xI - B) &= \det(C(xI - B)C^{-1}) \\ &= \det(xI - CBC^{-1}) \\ &= \det(xI - A) \end{aligned}$$

so $p_A(x) = p_B(x)$. Because B is triangular, $p_B(B) = 0$, so

$$p_A(A) = p_B(A) = p_B(CBC^{-1}) = C\,p_B(B)\,C^{-1} = 0$$

This proves equation 7.9 when the field F is algebraically closed. Any field can be embedded in an algebraically closed field, so this gives one method of proof for an arbitrary matrix A.

Over an algebraically closed field, most matrices A are in fact similar to a diagonal matrix D. Intuitively, one can make any matrix diagonalizable by deforming its entries a bit. More precisely, those $n \times n$ matrices that aren't diagonalizable form a lower dimensional set in the space of all $n \times n$ matrices. By reasoning like the above argument, $p_A(A)$ is a function on this space that vanishes on diagonalizable matrices, so it vanishes by continuity on all matrices. The details are straightforward for the complex numbers, but require the tools of algebraic geometry for other fields. We avoid these technicalities by working instead with triangular matrices.

The general case

We now establish equation 7.9 by a systematic approach that works for any matrix A, over any field F. The Cayley-Hamilton theorem in fact holds for matrices defined over arbitrary commutative rings. We have only developed a theory of bases when F is a field, but the following argument is quite general, and can be adapted to the case where F is a ring. One sees this same argument as the lead-in to Nakayama's lemma, in any commutative algebra textbook.

We demonstrate the method using a 2×2 matrix. Again let

$$\mathrm{A} = \begin{bmatrix} a & c \\ b & d \end{bmatrix}$$

Let $\mathrm{M}_\mathrm{A}(x)$ be the matrix function

$$\mathrm{M}_\mathrm{A}(x) = (x\mathrm{I} - \mathrm{A})^T = \begin{bmatrix} x-a & -b \\ -c & x-d \end{bmatrix}$$

The transpose does not change the determinant, so we have

$$\det(\mathrm{M}_\mathrm{A}(x)) = p_\mathrm{A}(x)$$

Therefore, if we multiply $\mathrm{M}_\mathrm{A}(x)$ by its adjoint, we obtain a diagonal matrix whose entries are $p_\mathrm{A}(x)$:

$$\begin{bmatrix} x-d & b \\ c & x-a \end{bmatrix} \begin{bmatrix} x-a & -b \\ -c & x-d \end{bmatrix} = \begin{bmatrix} p_\mathrm{A}(x) & 0 \\ 0 & p_\mathrm{A}(x) \end{bmatrix}$$

$\mathrm{M}_\mathrm{A}(x)$ creates a 2×2 matrix of elements of the same form as x. If we substitute a matrix for x, then M creates a 2×2 matrix of matrices. In particular,

$$\mathrm{M}_\mathrm{A}(\mathrm{A}) = \begin{bmatrix} \mathrm{A}-a\mathrm{I} & -b\mathrm{I} \\ -c\mathrm{I} & \mathrm{A}-d\mathrm{I} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} a-a & c \\ b & d-a \end{bmatrix} \begin{bmatrix} -b & 0 \\ 0 & -b \end{bmatrix} \\ \begin{bmatrix} -c & 0 \\ 0 & -c \end{bmatrix} \begin{bmatrix} a-d & c \\ b & d-d \end{bmatrix} \end{bmatrix}$$

Note that

$$\begin{bmatrix} \begin{bmatrix} a-a & c \\ b & d-a \end{bmatrix} \begin{bmatrix} -b & 0 \\ 0 & -b \end{bmatrix} \\ \begin{bmatrix} -c & 0 \\ 0 & -c \end{bmatrix} \begin{bmatrix} a-d & c \\ b & d-d \end{bmatrix} \end{bmatrix} \begin{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\ \begin{bmatrix} 0 \\ 1 \end{bmatrix} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} a-a \\ b-b \end{bmatrix} \\ \begin{bmatrix} c-c \\ d-d \end{bmatrix} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} 0 \\ 0 \end{bmatrix} \\ \begin{bmatrix} 0 \\ 0 \end{bmatrix} \end{bmatrix}$$

Using $\boldsymbol{e}_1 = (1, 0)$, $\boldsymbol{e}_2 = (0, 1)$, we can write this more concisely as

$$\begin{bmatrix} \mathrm{A}-a\mathrm{I} & -b\mathrm{I} \\ -c\mathrm{I} & \mathrm{A}-d\mathrm{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{e}_1 \\ \boldsymbol{e}_2 \end{bmatrix} = \begin{bmatrix} A\boldsymbol{e}_1 - a\boldsymbol{e}_1 - b\boldsymbol{e}_2 \\ A\boldsymbol{e}_2 - c\boldsymbol{e}_1 - d\boldsymbol{e}_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

Now, multiply $\mathrm{M_A}(A)$ by its adjoint:

$$\begin{bmatrix} \mathrm{A} - d\mathrm{I} & b\mathrm{I} \\ c\mathrm{I} & \mathrm{A} - a\mathrm{I} \end{bmatrix} \begin{bmatrix} \mathrm{A} - a\mathrm{I} & -b\mathrm{I} \\ -c\mathrm{I} & \mathrm{A} - d\mathrm{I} \end{bmatrix} = \begin{bmatrix} p_\mathrm{A}(\mathrm{A}) & 0 \\ 0 & p_\mathrm{A}(\mathrm{A}) \end{bmatrix}$$

Putting this together, we have

$$\begin{bmatrix} p_\mathrm{A}(\mathrm{A}) & 0 \\ 0 & p_\mathrm{A}(\mathrm{A}) \end{bmatrix} \begin{bmatrix} \boldsymbol{e}_1 \\ \boldsymbol{e}_2 \end{bmatrix} = \begin{bmatrix} p_\mathrm{A}(\mathrm{A})\,\boldsymbol{e}_1 \\ p_\mathrm{A}(\mathrm{A})\,\boldsymbol{e}_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

Since $p_\mathrm{A}(\mathrm{A})$ maps the basis $\boldsymbol{e}_1, \boldsymbol{e}_2$ to 0, it must be the zero matrix.

7.6 Repeated roots

If the characteristic polynomial p_A of a matrix A has repeated roots, then it may not be possible to diagonalize A.

A 2×2 example

Let $V = \mathbb{R}^2$ with the usual basis $S = \{\boldsymbol{e}_1, \boldsymbol{e}_2\}$. Let $L : V \to V$ be the linear map represented in S coordinates by the matrix

$$\mathrm{A} = \overset{L}{\underset{S \leftarrow S}{\begin{bmatrix} 2 & -4 \\ 1 & 6 \end{bmatrix}}}$$

Then

$$p_\mathrm{A}(x) = x^2 - 8x + 16 = (x-4)^2$$

so A has the single eigenvalue $\lambda = 4$, with multiplicity two.

If A could be diagonalized, then it would be similar to the scalar matrix $4\,\mathrm{I}$. However, for any change of basis matrix C we have

$$\mathrm{C} \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix} \mathrm{C}^{-1} = 4\,\mathrm{C} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \mathrm{C}^{-1} = \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix} \neq \begin{bmatrix} 2 & -4 \\ 1 & 6 \end{bmatrix}$$

In other words, scalar matrices aren't similar to any other matrix, and A isn't a scalar matrix, so it can't be diagonalized. Put differently, if L looks like multiplication by 4 in some coordinate system, then it looks like multiplication by 4 in every coordinate system. The matrix A certainly doesn't look like multiplication by 4, so it cannot be diagonalized. This reasoning will apply whenever A has only one eigenvalue.

Let

$$\mathrm{B} \;=\; \mathrm{A} - 4\,\mathrm{I} \;=\; \begin{bmatrix} -2 & -4 \\ 1 & 2 \end{bmatrix}$$

B is indeed singular, as expected because $p_{\mathrm{A}}(4) = 0$. The nullspace of B is the eigenspace of A with eigenvalue $\lambda = 4$. For A to have a basis of eigenvectors v_1, v_2, this nullspace would have to have dimension two, so B would have to have rank zero. However, only the zero matrix has rank zero. B is not the zero matrix, so we again see that A cannot be diagonalized.

This means that we cannot find two linearly independent vectors v_1, v_2 such that

$$v_1 \xrightarrow{\mathrm{B}} 0, \qquad v_2 \xrightarrow{\mathrm{B}} 0$$

However, by the Cayley-Hamilton theorem,

$$p_{\mathrm{A}}(\mathrm{A}) \;=\; (\mathrm{A} - 4\,\mathrm{I})^2 \;=\; \mathrm{B}^2 \;=\; 0$$

so B^2 is the zero matrix, as we would have liked for B itself. We check this. Indeed,

$$\mathrm{B}^2 \;=\; \begin{bmatrix} -2 & -4 \\ 1 & 2 \end{bmatrix}\begin{bmatrix} -2 & -4 \\ 1 & 2 \end{bmatrix} \;=\; \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \quad \checkmark$$

It would be nice if we could simply choose an arbitrary basis for the nullspace of B^2 and be done. However, S is already such a basis, and the appearance of A isn't exactly illuminating. We can do better.

The best we can do is to find a basis of vectors v_1, v_2 forming a chain

$$v_2 \xrightarrow{\mathrm{B}} v_1 \xrightarrow{\mathrm{B}} 0$$

In terms of such a basis $T = \{v_1, v_2\}$ we have

$$\begin{aligned} \mathrm{A}v_1 \;&=\; (4\,\mathrm{I} + \mathrm{B})\,v_1 \;&=\; 4\,v_1 \;+\; 0 \\ \mathrm{A}v_2 \;&=\; (4\,\mathrm{I} + \mathrm{B})\,v_2 \;&=\; 4\,v_2 \;+\; v_1 \end{aligned}$$

allowing us to represent the linear map L in T coordinates by the matrix

$$\mathrm{E} \;=\; \overset{L}{\underset{T \leftarrow T}{\begin{bmatrix} 4 & 1 \\ 0 & 4 \end{bmatrix}}}$$

This matrix E is in *Jordan canonical form*.

Any vector v_2 that is not in the nullspace of B will yield the desired chain. For example, the standard basis vectors e_1, e_2 can't both be in the nullspace of B, so we try both of them:

$$\begin{aligned} e_1 &= (1,0) \xrightarrow{\mathrm{B}} (-2,1) \xrightarrow{\mathrm{B}} 0 \\ e_2 &= (0,1) \xrightarrow{\mathrm{B}} (-4,2) \xrightarrow{\mathrm{B}} 0 \end{aligned}$$

We prefer the first chain; it leads to a simpler basis, with a change of basis matrix having determinant 1. Had we balked at using either of these chains, we could have chosen a nice solution to $\mathrm{B}v_1 = 0$, and then solved $\mathrm{B}v_2 = v_1$.

We have found the basis T given by

$$v_1 = (-2,1), \qquad v_2 = (1,0)$$

Let C be the change of basis matrix with columns v_1, v_2. Then we have the change of basis

$$\underset{S \leftarrow S}{\overset{L}{\begin{bmatrix} 2 & -4 \\ 1 & 6 \end{bmatrix}}} = \underset{S \leftarrow T}{\overset{I}{\begin{bmatrix} -2 & 1 \\ 1 & 0 \end{bmatrix}}} \underset{T \leftarrow T}{\overset{L}{\begin{bmatrix} 4 & 1 \\ 0 & 4 \end{bmatrix}}} \underset{T \leftarrow S}{\overset{I}{\begin{bmatrix} 0 & 1 \\ 1 & 2 \end{bmatrix}}}$$

$$\mathrm{A} \qquad\qquad \mathrm{C} \qquad\quad \mathrm{E} \qquad \mathrm{C}^{-1}$$

We can expand E as the linear combination $\mathrm{E} = 4\mathrm{I} + \mathrm{N}$ of the identity matrix and a nilpotent matrix N

$$\begin{bmatrix} 4 & 1 \\ 0 & 4 \end{bmatrix} = 4 \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

Changing back to S coordinates, we expand

$$\mathrm{A} = \mathrm{CEC}^{-1} = \mathrm{C}(4\mathrm{I}+\mathrm{N})\mathrm{C}^{-1} = 4\mathrm{CIC}^{-1} + \mathrm{CNC}^{-1}$$

This gives us

$$\begin{aligned} \begin{bmatrix} 2 & -4 \\ 1 & 6 \end{bmatrix} &= 4 \underbrace{\begin{bmatrix} -2 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & 2 \end{bmatrix}}_{4 \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}} + \underbrace{\begin{bmatrix} -2 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & 2 \end{bmatrix}}_{\begin{bmatrix} -2 & -4 \\ 1 & 2 \end{bmatrix}} \\ &= 4 \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} -2 & -4 \\ 1 & 2 \end{bmatrix} \end{aligned}$$

expressing our original matrix A as the same linear combination of the identity matrix and the nilpotent matrix B. Seeing that this is where we ended up, it would seem that we could have skipped right to this step.

Indeed, we have

$$\begin{array}{rclcrcl} A &=& C\,E\,C^{-1} &\quad\Rightarrow\quad& A^n &=& C\,E^n\,C^{-1} \\ A &=& \lambda I + B &\quad\Rightarrow\quad& A^n &=& \lambda^n I + n\,\lambda^{n-1} B \end{array}$$

using $B^2 = 0$. The first implication expands into the second, and either can be used to compute functions of A: We have

$$E^n = \begin{bmatrix} 4 & 1 \\ 0 & 4 \end{bmatrix}^n = \begin{bmatrix} 4^n & n\,4^{n-1} \\ 0 & 4^n \end{bmatrix} = 4^n \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + n\,4^{n-1} \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

so

$$A^n = 4^n \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + n\,4^{n-1} \begin{bmatrix} -2 & -4 \\ 1 & 2 \end{bmatrix}$$

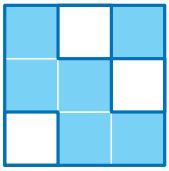

Chapter 8

Inner products

Lorem ipsum dolor sit amet, consectetur adipisicing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

Index