Linear Algebra *in Situ*

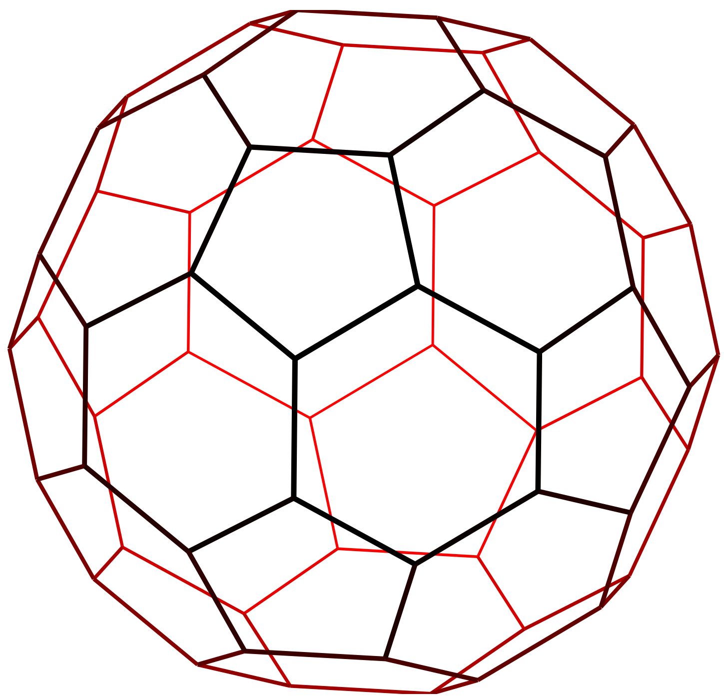

Steven J. Cox

```
Poet, oracle and wit
Like unsuccessful anglers by
The ponds of apperception sit,
Baiting with the wrong request
The vectors of their interest;
At nightfall tell the angler's lie.

With time in tempest everywhere,
To rafts of frail assumption cling
The saintly and the insincere;
Enraged phenomena bear down
In overwhelming waves to drown
Both sufferer and suffering.

The waters long to hear our question put
Which would release their longed for answer, but.

                        W.H. Auden
```

Preface

This is a text where concrete physical problems are posed and the ensuing mathematical theory is developed, tested, applied and associated with existing theory. The problems I pose spring from questions of equilibria, dynamics, optimization and inference of large electrical, mechanical and chemical networks. Following Gil Strang, I demonstrate throughout that Linear Algebra is both a tool for expressing these questions and for achieving, computing and representing their solutions.

The theory needed to resolve the questions of network equilibria, optimization and inference is now well enshrined in the Fundamental Theorem of Linear Algebra, and it appears difficult to improve on this approach. Regarding dynamics however there are two distinct paths to the spectral theorem; one via zeros of the characteristic polynomial, $\det(zI - A)$, the other via poles of the resolvent, $(zI - A)^{-1}$. The first is common among introductory texts while the latter, to my knowledge, has yet to succeed at that level – although, since the treatise of Kato, it is well known to be considerably cleaner and more flexible. I feel strongly that students new to linear algebra can grasp the resolvent more readily than the determinant. For, with eigenvalues defined as those z for which $(zI-A)$ does not have an inverse, the **direct approach** is to simply construct $(zI-A)^{-1}$ and observe the offending z. The construction of $(zI - A)^{-1}$, say via Gauss–Jordan, is straightforward though tedious. Once they understand the process however they may turn the tedium over to one of a number of "symbolic algebra" routines. I make systematic use of the symbolic toolbox in MATLAB . By contrast, the **indirect approach** ignores the inverse and relies on the determinant as a mere numerical test of invertibility. The approach via the resolvent comes however at the cost of presuming familiarity with the residue theorem of complex integration. I see this rather as a win–win situation, for the residue theorem is also key to making proper sense of the Inverse Laplace and Fourier Transforms. Hence, our two brief chapters on complex variables pay multiple dividends.

The reader will find here an introductory course, an advanced course, an array of intermediate courses, and a reference for self–study and/or use in advanced courses across Science, Engineering and Mathematics. The general audience introductory course, assuming only one year of calculus, that I have taught to sophomores at Rice University for more than 20 years, is composed of the following sections from the first 13 chapters:

Introductory Course

1. Orientation, §§1–3
2. Electrical Networks, §§1–2
3. Mechanical Networks, §§1–3
4. The Column and Null Spaces, §§1–3
5. The Fundamental Theorem and Beyond, §§1–3
6. Least Squares, §§1–4
7. Metabolic Networks, §§1–3
8. Dynamical Systems, §§1–4
9. Complex Numbers, Vectors and Functions, §§1–3
10. Complex Integration, §§1–3
11. The Eigenvalue Problem, §§1–2
12. The Hermitian Eigenvalue Problem, §§1–2
13. The Singular Value Decomposition, §§1–2.

This course stresses applications, methods and computation over theory and algorithms. As the audience has been predominantly students of engineering and science I have used application chapters to motivate theory chapters and then used this theory to both revisit old applications and to embark on new ones. For example, the pseudo–inverse is invoked in Chapter 3 in order to ignore the rigid body motion of a mechanical network. This provokes discussion of null and column spaces but does not get resolved until the spectral representation and singular value decomposition in Chapters 11–13. Similary, the resolvent and eigenvalues arise naturally in our consideration, in Chapter 8, of dynamical systems but do not get resolved until the spectral representation is reached. As such the material, including the exercises, in the early sections of the first 13 Chapters (with the exclusion of Chapter 7 on Metabolic Networks) is highly integrated.

For audiences with either prior exposure to linear algebra or motivating applications one can skim Chapter 1 and the early sections of Chapters of 2, 3 and 7 and use the time saved to delve more deeply into the latter, more challenging, sections of Chapters 2–13 or perhaps into the more advanced material of Chapters 14–16. These last three chapters, presuming a solid foundation in Linear Algebra, develop the Group, Representation and Graph Theory that underly the exact solution to three exciting problems concerning large networks. In particular: I provide a detailed derivation of the exact formulas, due to Chung and Sternberg, for the 60 eigenvalues that govern the electronic structure of the Buckyball, and I provide detailed proofs that concrete constructions of Margulis achieve large girth in one case and establish a family of expander graphs in the other.

Steve Cox

Contents

1 **Orientation** **1**
1.1 **Objects** . . . 1
1.2 **Computations** . . . 6
1.3 **Proofs** . . . 10
1.4 **Notes and Exercises** . . . 13

2 **Electrical Networks** **19**
2.1 **Neurons and the Strang Quartet** . . . 19
2.2 **Resistor Nets with Current Sources and Batteries** . . . 20
2.3 **Operational Amplifiers** . . . 23
2.4 **Notes and Exercises** . . . 26

3 **Mechanical Networks** **31**
3.1 **Elastic Fibers and the Strang Quartet** . . . 31
3.2 **Gaussian Elimination and LU Decomposition** . . . 33
3.3 **Planar Network Examples** . . . 38
3.4 **Equilibrium and Energy Minimization** . . . 43
3.5 **Notes and Exercises** . . . 45

4 **The Column and Null Spaces** **50**
4.1 **The Column Space** . . . 50
4.2 **The Null Space** . . . 52
4.3 **On the Stability of Mechanical Networks** . . . 53
4.4 **The Structure of Nilpotent Matrices** . . . 57
4.5 **Notes and Exercises** . . . 61

5 **The Fundamental Theorem and Beyond** **64**
5.1 **The Row Space** . . . 64
5.2 **The Left Null Space** . . . 66
5.3 **Synthesis** . . . 68
5.4 **Vector Spaces and Linear Transformations** . . . 69
5.5 **Convex Sets** . . . 72
5.6 **Exercises** . . . 78

6 **Least Squares** **80**
6.1 **The Normal Equations** 80
6.2 **Application to a Biaxial Test Problem** 82
6.3 **Projections** 84
6.4 **The QR Decomposition** 85
6.5 **Orthogonal Polynomials** 88
6.6 **Detecting Integer Relations** 93
6.7 **Probabilistic and Statistical Interpretations** 96
6.8 **Autoregressive Models and Levinson's Algorithm** 99
6.9 **Notes and Exercises** 103

7 **Metabolic Networks** **108**
7.1 **Flux Balance and Optimal Yield** 108
7.2 **Linear Programming** 109
7.3 **The Simplex Method** 110
7.4 **The Geometric Point of View** 111
7.5 **Succinate Production** 113
7.6 **Elementary Flux Modes and Extremal Rays** 115
7.7 **Notes and Exercises** 120

8 **Dynamical Systems** **122**
8.1 **Neurons and the Dynamic Strang Quartet** 122
8.2 **The Laplace Transform** 124
8.3 **The Backward–Euler Method** 127
8.4 **Dynamics of Mechanical Systems** 129
8.5 **Exercises** 131

9 **Complex Numbers, Functions and Derivatives** **134**
9.1 **Complex Numbers** 134
9.2 **Complex Functions** 136
9.3 **Complex Differentiation and the First Residue Theorem** 139
9.4 **Fourier Series and Transforms** 142
9.5 **The Power Spectra of Stationary Processes** 145
9.6 **Notes and Exercises** 148

10 **Complex Integration** **151**

10.1 **Cauchy's Theorem** 151
10.2 **The Second Residue Theorem** 155
10.3 **The Inverse Laplace Transform and Return to Dynamics** 158
10.4 **The Inverse Fourier Transform and the Causal Wiener Filter** 160
10.5 **Further Applications of the Second Residue Theorem** 162
10.6 **Notes and Exercises** 164

11 The Eigenvalue Problem 166

11.1 **The Resolvent** 166
11.2 **The Spectral Representation** 171
11.3 **Diagonalization of a Semisimple Matrix** 173
11.4 **The Schur Form and the QR Algorithm** 176
11.5 **The Jordan Canonical Form** 178
11.6 **Positive Matrices and the PageRank Algorithm** 184
11.7 **Notes and Exercises** 188

12 The Hermitian Eigenvalue Problem 192

12.1 **The Spectral Representation** 192
12.2 **Orthonormal Diagonalization of Hermitian Matrices** 194
12.3 **Perturbation Theory** 196
12.4 **Rayleigh's Principle and the Power Method** 199
12.5 **Hückel's Molecular Orbital Theory** 201
12.6 **Optimal Damping of Mechanical Networks** 204
12.7 **Notes and Exercises** 210

13 The Singular Value Decomposition 214

13.1 **The Decomposition** 214
13.2 **The SVD in Image Compression** 217
13.3 **Low Rank Approximation** 219
13.4 **Principal Component Analysis** 220
13.5 **Notes and Exercises** 220

14 Matrix Groups 222

14.1 **Orthogonal Groups** 222
14.2 **Symmetry Groups** 224
14.3 **Permutation Groups** 230

14.4 **Linear, Free, and Quotient Groups** 235
14.5 **Group Action and Counting Theory** 242
14.6 **Notes and Exercises** 247

15 **Group Representation Theory** **249**
15.1 **Representations** 249
15.2 **Characters** 253
15.3 **The Electronic Structure of the Buckyball** 258
15.4 **Fourier Analysis on Abelian Groups** 264
15.5 **Notes and Exercises** 268

16 **Graph Theory** **271**
16.1 **Graphs, Matrices and Groups** 271
16.2 **Trees and Molecules** 271
16.3 **Spanning Trees and Electrical Networks** 275
16.4 **Cycles and Girth** 279
16.5 **The Isoperimetric Constant and Expanders** 283
16.6 **Notes and Exercises** 287

17 **References** **290**

1. Orientation

You have likely encountered vectors, and perhaps matrices in your introductory calculus and/or physics courses. My goal in this chapter is to strengthen these encounters and so prepare you for the applications, computations and theory to come. I begin in §1.1 with a careful presentation of the basic objects – and the laws that govern their arithmetic combinations. I then introduce MATLAB in §1.2 as a means to visually explore the sense in which matrices transform vectors. I complete our orientation in §1.3 with an introduction to the principle methods of proof used in Linear Algebra. Throughout the chapter I introduce and reinforce concepts through examples and stress that you gain confidence and expertise by generating examples of your own. The exercises at the end of the chapter should help toward that end.

1.1. Objects

A **vector** is a column of real numbers, and is written, e.g.,

$$x = \begin{pmatrix} 2 \\ -4 \\ 1 \end{pmatrix}. \tag{1.1}$$

The vector has 3 elements and so lies in the class of all 3–element vectors, denoted, $\mathbb{R}^3$, where $\mathbb{R}$ stands for "real". We denote "is a member of" by the symbol $\in$. So, e.g., $x \in \mathbb{R}^3$. We denote the first element of x by x_1, its second element by x_2 and so on. For example, $x_2 = -4$ in (1.1).

We will typically use the positive integer n to denote the ambient dimension of our problem, and so will be working in $\mathbb{R}^n$. The sum of two vectors, x and y, in $\mathbb{R}^n$ is defined elementwise by

$$z = x + y, \quad \text{where} \quad z_j = x_j + y_j, \quad j = 1, \ldots, n.$$

The multiplication of a vector, $x \in \mathbb{R}^n$, by a scalar $s \in \mathbb{R}$ is defined elementwise by

$$z = sx, \quad \text{where} \quad z_j = sx_j, \quad j = 1, \ldots, n.$$

For example,

$$\begin{pmatrix} 2 \\ 5 \end{pmatrix} + \begin{pmatrix} 1 \\ -3 \end{pmatrix} = \begin{pmatrix} 3 \\ 2 \end{pmatrix} \quad \text{and} \quad 6\begin{pmatrix} 4 \\ 2 \end{pmatrix} = \begin{pmatrix} 24 \\ 12 \end{pmatrix}.$$

The most common product of two vectors, x and y, in $\mathbb{R}^n$ is the **inner product**,

$$x^T y \equiv \begin{pmatrix} x_1 & x_2 & \cdots & x_n \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} = x_1 y_1 + x_2 y_2 + \cdots + x_n y_n = \sum_{j=1}^{n} x_j y_j. \tag{1.2}$$

As $x_j y_j = y_j x_j$ for each j it follows that $x^T y = y^T x$. For example,

$$\begin{pmatrix} 10 & 1 & 3 \end{pmatrix} \begin{pmatrix} 8 \\ 2 \\ -4 \end{pmatrix} = 10 \cdot 8 + 1 \cdot 2 + 3 \cdot (-4) = 70. \tag{1.3}$$

So, the inner product of two vectors is a number. The superscript T on the x on the far left of Eq. (1.2) stands for **transpose** and, when applied to a column yields a **row**. Columns are vertical and rows are horizontal and so we see, in Eq. (1.2), that x^T is x laid on its side. We follow Euclid and measure the magnitude, or more commonly the **norm**, of a vector by the square root of the sum of the squares of its elements. In symbols,

$$\|x\| \equiv \sqrt{x^T x} = \sqrt{\sum_{j=1}^{n} x_j^2}. \tag{1.4}$$

For example, the norm of the vector in Eq. (1.1) is $\sqrt{21}$. As Eq. (1.4) is a direct generalization of the Euclidean distance of high school planar geometry we may expect that $\mathbb{R}^n$ has much the same "look." To be precise, let us consider the situation of Figure 1.1.

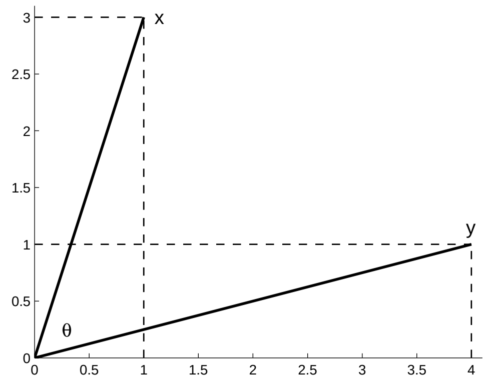

Figure 1.1. A guide to interpreting the inner product.

We have x and y in $\mathbb{R}^2$ and

$$x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 1 \\ 3 \end{pmatrix} \quad \text{and} \quad y = \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} 4 \\ 1 \end{pmatrix}$$

and we recognize that both x and y define right triangles with hypotenuses $\|x\|$ and $\|y\|$ respectively. We have denoted by θ the angle that x makes with y. If θ_x and θ_y denotes the angles that x and y respectively make with the positive horizontal axis then $\theta = \theta_x - \theta_y$ and the Pythagorean Theorem permits us to note that

$$x_1 = \|x\|\cos(\theta_x), \quad x_2 = \|x\|\sin(\theta_x), \quad \text{and} \quad y_1 = \|y\|\cos(\theta_y), \quad y_2 = \|y\|\sin(\theta_y),$$

and these in turn permit us to express the inner product of x and y as

$$\begin{aligned} x^T y &= x_1 y_1 + x_2 y_2 \\ &= \|x\|\|y\|(\cos(\theta_x)\cos(\theta_y) + \sin(\theta_x)\sin(\theta_y)) \\ &= \|x\|\|y\|\cos(\theta_x - \theta_y) \\ &= \|x\|\|y\|\cos(\theta). \end{aligned} \tag{1.5}$$

We interpret this by saying that the inner product of two vectors is proportional to the cosine of the angle between them. Now given two vectors in say $\mathbb{R}^8$ we don't panic, rather we orient ourselves by observing that they together lie in a particular plane and that this plane, and the angle they make with one another is in no way different from the situation illustrated in Figure 1.1. And for this reason we say that x and y are perpendicular, or **orthogonal**, to one another whenever $x^T y = 0$.

In addition to the geometric interpretation of the inner product it is often important to be able to estimate it in terms of the products of the norms. Here is an argument that works for x and y in $\mathbb{R}^n$. Once we know where to start, we simply expand

$$\begin{aligned}\|(y^Ty)x-(x^Ty)y\|^2 &= ((y^Ty)x-(x^Ty)y)^T((y^Ty)x-(x^Ty)y)\\ &= \|y\|^4\|x\|^2-2\|y\|^2(x^Ty)^2+(x^Ty)^2\|y\|^2\\ &= \|y\|^2(\|x\|^2\|y\|^2-(x^Ty)^2)\end{aligned} \tag{1.6}$$

and then note that as the initial expression is nonnegative, the final expression requires (after taking square roots) that

$$\boxed{|x^Ty| \le \|x\|\|y\|.} \tag{1.7}$$

This is known as the **Cauchy–Schwarz inequality**.

As a vector is simply a column of numbers, a matrix is simply a row of columns, or a column of rows. This necessarily requires two numbers, the row and column indices, to specify each matrix element. For example

$$A = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix} = \begin{pmatrix} 5 & 0 & 1 \\ 2 & 3 & 4 \end{pmatrix} \tag{1.8}$$

is a 2-by-3 matrix. The first dimension is the number of rows and the second is the number of columns and this ordering is also used to address individual elements. For example, the element in row 1 column 3 is $a_{13} = 1$. We will consistently use upper–case letters to denote matrices.

The addition of two matrices (of the same size) and the multiplication of a matrix by a scalar proceed exactly as in the vector case. In particular,

$$(A+B)_{ij} = a_{ij} + b_{ij}, \quad \text{e.g.,} \quad \begin{pmatrix} 5 & 0 & 1 \\ 2 & 3 & 4 \end{pmatrix} + \begin{pmatrix} 2 & 4 & 6 \\ 1 & -3 & 4 \end{pmatrix} = \begin{pmatrix} 7 & 4 & 7 \\ 3 & 0 & 8 \end{pmatrix},$$

and

$$(cA)_{ij} = ca_{ij}, \quad \text{e.g.,} \quad 3\begin{pmatrix} 5 & 0 & 1 \\ 2 & 3 & 4 \end{pmatrix} = \begin{pmatrix} 15 & 0 & 3 \\ 6 & 9 & 12 \end{pmatrix}.$$

The product of two commensurate matrices proceeds through a long sequence of inner products. In particular if $C = AB$ then the ij element of C is the product of the ith row of A and the jth column of B. Hence, for two A and B to be commensurate it follows that each row of A must have the same number of elements as each column of B. In other words, the number of columns of A must match the number of rows of B. Hence, if A is m-by-n and B is n-by-p then the ij element of their product $C = AB$ is

$$c_{ij} = \sum_{k=1}^{n} a_{ik}b_{kj} = A(i,:)B(:,k), \tag{1.9}$$

where $A(i,:)$ denotes row i of A and $B(:,k)$ denotes column k of B. For example,

$$\begin{pmatrix} 5 & 0 & 1 \\ 2 & 3 & 4 \end{pmatrix}\begin{pmatrix} 2 & 4 \\ 6 & 1 \\ -3 & 4 \end{pmatrix} = \begin{pmatrix} 5\cdot 2+0\cdot 6+1\cdot(-3) & 5\cdot 4+0\cdot 1+1\cdot 4 \\ 2\cdot 2+3\cdot 6+4\cdot(-3) & 2\cdot 4+3\cdot 1+4\cdot(-4) \end{pmatrix} = \begin{pmatrix} 7 & 24 \\ 10 & -5 \end{pmatrix}.$$

In this case, the product BA is not even defined. If A is m-by-n and B is n-by-m then both AB and BA are defined, but unless $m = n$ they are of distinct dimensions and so not comparable. If

$m = n$ so A and B are square then we may ask if $AB = BA$? and learn that the answer is typically no. For example,

$$\begin{pmatrix} 5 & 0 \\ 2 & 3 \end{pmatrix} \begin{pmatrix} 2 & 4 \\ 6 & 1 \end{pmatrix} = \begin{pmatrix} 10 & 20 \\ 22 & 11 \end{pmatrix} \neq \begin{pmatrix} 2 & 4 \\ 6 & 1 \end{pmatrix} \begin{pmatrix} 5 & 0 \\ 2 & 3 \end{pmatrix} = \begin{pmatrix} 18 & 12 \\ 32 & 3 \end{pmatrix}. \tag{1.10}$$

We will often abbreviate the awkward phrase "A is m-by-n" with the declaration $A \in \mathbb{R}^{m\times n}$. The matrix algebra of multiplication, though tedious, is easy enough to follow. It stemmed from a row-centric point of view. It will help to consider the columns. If $A \in \mathbb{R}^{m\times n}$ and the jth column of A is $A(:,j)$ and $x \in \mathbb{R}^n$ then we recognize the product

$$Ax = [A(:,1)\ A(:,2)\ \cdots\ A(:,n)] \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} = x_1 A(:,1) + x_2 A(:,2) + \cdots + x_n A(:,n), \tag{1.11}$$

as a weighted sum of the columns of A. For example

$$\begin{pmatrix} 2 & 3 \\ 1 & 4 \end{pmatrix} \begin{pmatrix} 2 \\ 3 \end{pmatrix} = 2 \begin{pmatrix} 2 \\ 1 \end{pmatrix} + 3 \begin{pmatrix} 3 \\ 4 \end{pmatrix} = \begin{pmatrix} 13 \\ 14 \end{pmatrix}. \tag{1.12}$$

We illustrate this in Figure 1.2(A) and then proceed to illustrate in the second panel the transformation by this A of a representative collection of unit vectors.

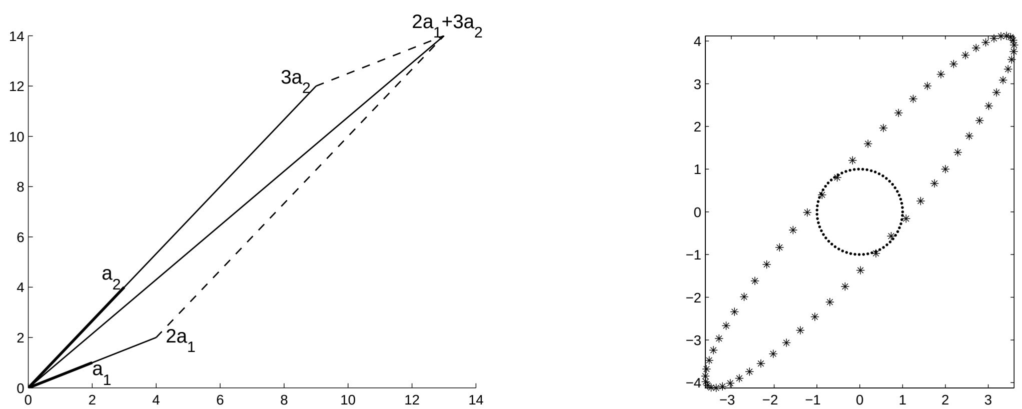

Figure 1.2. (A) An illustration of the matrix vector multiplication conducted in Eq. (1.12). Both $A(:,1)$ and $A(:,2)$ are plotted heavy for emphasis. We see that their multiples, by 2 and 3, simply extend them, while their weighted sum simply completes the natural parallelogram. (B) For a given x on the unit circle (denoted by a dot) we plot its transformation by the A matrix of Eq. (1.12) (denoted by an asterisk). `mymult.m`

A common goal of matrix analysis is to describe m-by-n matrices by many fewer than mn numbers. The simplest such descriptor is the sum of the matrice's diagonal elements. We call this the **trace** and abbreviate it by

$$\mathrm{tr}(A) \equiv \sum_{i=1}^{n} a_{ii}. \tag{1.13}$$

Looking for matrices to trace you scan Eq. (1.10) and note that $10 + 11 = 18 + 3$ and you ask, knowing that $AB \neq BA$, whether

$$\boxed{\mathrm{tr}(AB) = \mathrm{tr}(BA)} \tag{1.14}$$

might possibly be true in general. For arbitrary A and B in $\mathbb{R}^{n\times n}$ we therefore construct $\mathrm{tr}(AB)$

$$(AB)_{ii} = \sum_{k=1}^{n} a_{ik}b_{ki} \quad \text{so} \quad \mathrm{tr}(AB) = \sum_{i=1}^{n}\sum_{k=1}^{n} a_{ik}b_{ki},$$

and $\mathrm{tr}(BA)$

$$(BA)_{ii} = \sum_{k=1}^{n} b_{ik}a_{ki} \quad \text{so} \quad \mathrm{tr}(BA) = \sum_{i=1}^{n}\sum_{k=1}^{n} b_{ik}a_{ki}.$$

These sums indeed coincide, for both are simply the sum of the product of each element of A and the reflected (interchange i and k) element of B.

In general, if A is m-by-n then the matrix that results on exchanging its rows for its columns is called the **transpose** of A, denoted A^T. It follows that A^T is n-by-m and

$$(A^T)_{ij} = a_{ji}.$$

For example,

$$\begin{pmatrix} 5 & 0 & 1 \\ 2 & 3 & 4 \end{pmatrix}^T = \begin{pmatrix} 5 & 2 \\ 0 & 3 \\ 1 & 4 \end{pmatrix}.$$

We will have frequent need to transpose a product, so let us contrast

$$((AB)^T)_{ij} = \sum_{k=1}^{n} a_{jk}b_{ki}$$

with

$$(B^TA^T)_{ij} = \sum_{k=1}^{n} a_{jk}b_{ki} \tag{1.15}$$

and so conclude that

$$\boxed{(AB)^T = B^TA^T,} \tag{1.16}$$

i.e., that the transpose of a product is the product of the transposes in reverse order.

Regarding the norm of a matrix it seems natural, on recalling our definition of the norm of a vector, to simply define it as the square root of the sum of the squares of each element. This definition, where $A \in \mathbb{R}^{m\times n}$ is viewed as a collection of vectors, is associated with the name Frobenius and hence the subscript in the definition of the **Frobenius norm** of A,

$$\|A\|_F \equiv \left(\sum_{i=1}^{m}\sum_{j=1}^{n} a_{ij}^2\right)^{1/2}. \tag{1.17}$$

As scientific progress and mathematical insight most often come from seeing things from multiple angles we pause to note Eq. (1.17) may be seen as the trace of a product. In particular, with $B = A^T$ and $j = i$ in the general formula Eq. (1.15) we arrive immediately at

$$(AA^T)_{ii} = \sum_{k=1}^{n} a_{ik}^2.$$

As the sum over i is precisely the trace of AA^T we have established the equivalent definition

$$\|A\|_F = (\mathrm{tr}(AA^T))^{1/2}. \tag{1.18}$$

For example, the Frobenius norm of the A in Eq. (1.8) is $\sqrt{55}$. Just as the vector norm can help us bound (recall Eq. (1.7)) the inner product of two vectors, this matrix norm can help us bound the product of a matrix and vector. More precisely, lets prove that

$$\|Ax\| \le \|A\|_F \|x\|, \tag{1.19}$$

for arbitrary A and x. To see this we complement Eq. (1.11) with a row representation

$$Ax = \begin{pmatrix} A(1,:)x \\ A(2,:)x \\ \vdots \\ A(m,:)x \end{pmatrix}$$

and so

$$\begin{aligned} \|Ax\| &= \sqrt{(A(1,:)x)^2 + (A(2,:)x)^2 + \cdots + (A(m,:)x)^2} \\ &\le \sqrt{\|A(1,:)\|^2\|x\|^2 + \|A(2,:)\|^2\|x\|^2 + \cdots + \|A(:,n)\|^2\|x\|^2} \\ &= \|A\|_F\|x\|, \end{aligned}$$

where we have used Eq. (1.7) to conclude that each $|A(j,:)x| \le \|A(j,:)\|\|x\|$. The simple rearrangement of Eq. (1.19),

$$\frac{\|Ax\|}{\|x\|} \le \|A\|_F \qquad \forall\ x, \tag{1.20}$$

has the nice geometric interpretation: "The matrix A can stretch no vector by more than $\|A\|_F$." We can reinforce this interpretation by returning to Figure 1.2 and noting that no vector in the ellipse is longer than $\|A\|_F = \sqrt{30}$.

1.2. Computations

The objects of the previous section turn stale and are easily forgotten unless handled. We are fortunate to work in a time in which both the tedium of their manipulation and the task of illustrating our "findings" have been automated – leaving one's imagination the only obstacle to discovery.

To prepare you to "handle" your own objects we now present a brief introduction to Matlab via experiments on the innocent looking

$$A = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}. \tag{1.21}$$

It is inert until it acts. Its action is spelled out in

$$Ax = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} x_1 + 2x_2 \\ x_2 \end{pmatrix} \tag{1.22}$$

but perhaps these symbols do not yet speak to you. To illustrate or animate this action we might turn to devices like Figure 1.2 where we plot its deformation of the unit circle. Though this gives a general sense of its influence it neglects to track the transformation of individual unit vectors. We correct for this and display our findings in Figure 1.3, by marking 12 unit vectors in black and their 12 deformations, under A, in red.

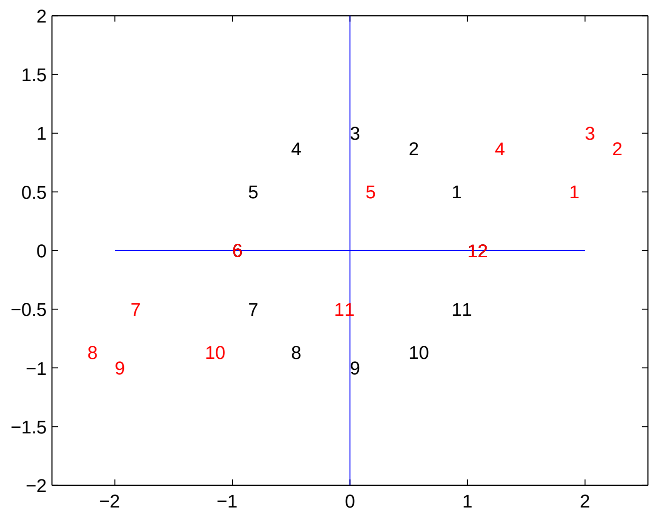

Figure 1.3. Illustration of the action, Ax in red, specified in Eq. (1.22) for the twelve x vectors (black). That is, A takes the black 1 to the red 1, the black 2 to the red 2 and so on. Yes, both the black 6 and black 12 remain unmoved by A.

Now we are really on to something – for this figure suggests so many new questions! But before getting carried away lets take a careful look at the MATLAB script, `Morb.m`, that generated Figure 1.3. For ease of reference we have numbered each line in our program.

```
1       A = [1 2; 0 1];                         % the matrix
2       plot([-2 2],[0 0])                      % plot the horizontal axis
3       hold on                                 % plot future info in same figure
4       plot([0 0],[-2 2])                      % plot the vertical axis
5       for j=1:12,                             % do what follows 12 times
6           ang = j*2*pi/12;                    % angle
7           x = [cos(ang); sin(ang)];           % a point on the unit circle
8           y = A*x;                            % transformed by A
9           text(x(1),x(2),num2str(j))          % place the counter value at x
10          s = text(y(1),y(2),num2str(j));     % place the counter value at y
11          set(s,'color','r')                  % paint that last value red
12      end
13      hold off                                % let go of the picture
14      axis equal                              % fiddle with the axes
```

Our actor, `A`, gets line 1 billing. We specify matrices, and columns, between square brackets and terminate each row (except the last) with a semicolon. Note that line 1 is not an equation but rather an *assignment*. MATLAB assigns what it finds to the right of `=` to the symbol it finds at the left.

In line 2 we instruct MATLAB to plot a line in the plane from $(-2, 0)$ to $(2, 0)$ using the default color, blue. In line 4 we instruct MATLAB to plot a blue line from $(0, -2)$ to $(0, 2)$.

In line 5 we enter a loop that terminates at line 12 when the counter, `j`, reaches its terminal value. The colon is a powerful synonym for 'to,' in the sense that we read line 5 as "for j equal 1 to 12 execute lines 6 through 11." You see that `ang` will then take on multiples of $\pi/6$ and that `x` will be the associated unit vector and `y` its transformation under `A`. In line 9 through 10 we take the important step of actually *marking* our tracks by turning the counter value to a text string that is then placed at `(x(1),x(2))` in the default (black) color and then again at `(y(1),y(2))`, but this time in red.

This script now belongs to your list of objects and as such invites experimentation. For example, What must change to up the action from 12 to 24 players? Once you've learned this script we can return to pondering Figure 1.3. Do you see that it `shears` the circle in the sense that it drags the top half to the right and bottom half to the left while the equator remains unmoved? Does this suggest that we could learn more be deforming shape other than circles? Though many shapes come to mind we might miss something if we stick to regular objects. One of the key advantages of computational experimentation is the ability to simultaneously observe the action upon many random players. One difficulty with `many` is that it becomes more difficult to mark our tracks. To get round this we will restrict our players to one half of the plane and paint each black while painting red their action by A. So how should we divide the plane. The simple guess of top, $x_2 > 0$, and bottom, $x_2 < 0$ does not seem to expose any new patterns and so one might instead tilt this guess to say align with diagonals and so divide the plane into the two bow-ties

$$E = \{x \in \mathbb{R}^2 : |x_2| > |x_1|\} \quad \text{and} \quad F = \{x \in \mathbb{R}^2 : |x_1| > |x_2|\}. \tag{1.23}$$

We illustrate our remarkable findings in Figure 1.4.

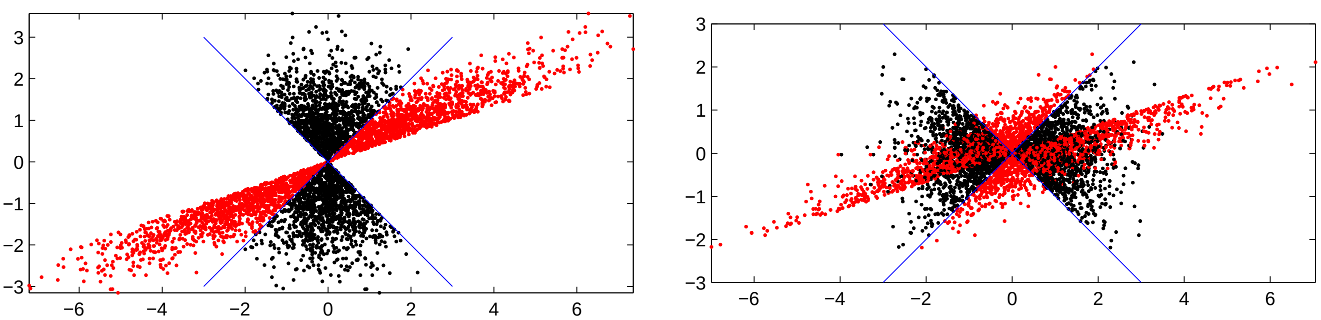

Figure 1.4. (A) The deformation (red) by A of 2500 random vectors (black) from E. We surmise that A takes E to F. (B) The deformation (red) by A of 2500 random vectors (black) from F.

The difference in clarity between between panels (A) and (B) is striking – for these are drawn from the *same* matrix. Panel (A) leads immediately, via Eq. (1.22), to the conjecture: if $|x_2| > |x_1|$ then $|x_1 + 2x_2| > |x_2|$. We leave its proof (and more) to Exer. 1.3 in order that we may explicate the script that generated Figure 1.4.

```
A = [1 2; 0 1];                         % the matrix
for n=1:2500,                           % do the following 2500 times
    x = randn(2,1);                     % generate a random point
    [sax,ord] = sort(abs(x),1,'ascend'); % sort their magnitudes
    x = x(ord);                         % reorder the elements
    y = A*x;                            % transform via A
    plot(x(1),x(2),'k.')                % mark the original point black
```

```
        hold on                                  % save this picture
        plot(y(1),y(2),'r.')                     % mark the transformed point red
    end
    plot([-3 3],[3 -3])                          % plot the NW-SE diagonal
    plot([-3 3],[-3 3])                          % plot the SW-NE diagonal
    axis equal                                   % fiddle with axes
    hold off                                     % let go of the picture
```

There are two key differences with the previous script. Our x vectors are now generated (and reordered) at random and we are plotting points rather than texting strings. The `x = randn(2,1)` places two random samples of the normal (or Gaussian, or bell–curve) distribution into the 2–by–1 vector `x`. In order to ensure that this `x` lies in E we sort its absolute values via `sort` in an ascending fashion. The `sort` function returns two objects: `sax`, the sorted values and `ord`, the order in which they appeared. More precisely if `abs(`x_1`)<abs(`x_2`)` then `ord=[1 2]` and `x=x(ord)` changes nothing while if instead `abs(`x_1`)>abs(`x_2`)` then `ord=[2 1]` and `x=x(ord)` corrects their order. If instead we wish to restrict `x` to F, to generate panel (B), we switch `ascend` to `descend`.

Now that we understand how matrices like $A = [1\ 2; 0\ 1]$ act on objects like circles and bowties we may inspect their action on much more complicated objects. MATLAB has a large library of stock images that we may manipulate. We present such a before an after in Figure 1.5.

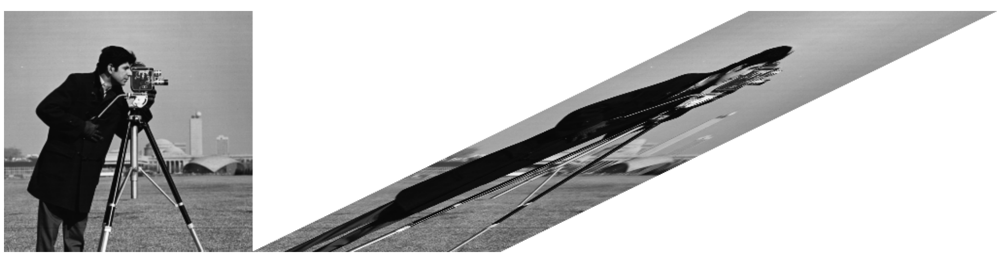

Figure 1.5. An image of a camerman, normal and sheared by the A matrix in (1.22).

The code that achieves this transformation is

```
        P = imread('cameraman.tif');         % read the image
        [m,n] = size(P);                     % record its size
1       SP = 256*ones(m,2*m+n,'uint8');      % create a white canvas
        for i=1:m                            % inspect every pixel
        for j=1:n,                           %    of the original image
2           SP(i,2*m+j-2*i) = P(i,j);        % and shear it with the matrix A
        end
        end
        imshow([P SP])                       % display both images
```

We have numbered the "interesting lines." Regarding line 1, Why does 256 designate white? and Why have we added $2m$ columns? Regarding line 2, where exactly is A? You can discover the answers by observing the result of small changes to these lines.

1.3. Proofs

Regarding the **proofs** in the text, and more importantly in the exercises and exams, many will be of the type that brought us Eq. (1.14) and Eq. (1.16). These are what one might call **confirmations**. They require a clear head and may require a bit of rearrangement but as they follow directly from definitions they do not require magic, clairvoyance or even ingenuity. As further examples of confirmations let us prove (confirm) that

$$\mathrm{tr}(A) = \mathrm{tr}(A^T). \tag{1.24}$$

It would be acceptable to say that "As A^T is the reflection of A across its diagonal both A and A^T agree on the diagonal. As the trace of matrix is simply the sum of its diagonal terms we have confirmed Eq. (1.24)." It would also be acceptable to proceed in symbols and say "from $(A^T)_{ii} = a_{ii}$ for each i it follows that

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

It would not be acceptable to confirm Eq. (1.24) on a particular numerical matrix, nor even on a class of matrices of a particular size.

As a second example, lets confirm that

$$\text{if} \quad \|x\| = 0 \quad \text{then} \quad x = 0. \tag{1.25}$$

It would be acceptable to say that "As the sum of the squares of each element of x is zero then in fact each element of x must vanish." Or, in symbols, as

$$\sum_{i=1}^{n} x_i^2 = 0$$

we conclude that each $x_i = 0$.

Our third example is a slight variation on the second.

$$\text{if } x \in \mathbb{R}^n \quad \text{and} \quad x^T y = 0 \quad \text{for all } y \in \mathbb{R}^n \quad \text{then} \quad x = 0. \tag{1.26}$$

This says that the only vector that is orthogonal to every vector in the space is the zero vector. The most straightforward proof is probably the one that reduces this to the previous Proposition, Eq. (1.25). Namely, since $x^T y = 0$ for each y we can simply use $y = x$ and discern that $x^T x = 0$ and conclude from Eq. (1.25) that indeed $x = 0$. As this section is meant to be an introduction to proving let us apply instead a different strategy, one that replaces a proposition with its equivalent contra–positive. More precisely, if your proposition reads "if c then d" then its contrapositive reads "if not d then not c." Do you see that a proposition is true if and only its contrapositive is true? Why bother? Sometimes the contrapositive is "easier" to prove, sometimes it throws new light on the original proposition, and it always expands our understanding of the landscape. So let us construct the contra–positive of Eq. (1.26). As clause d is simply $x = 0$, not d is simply $x \neq 0$. Clause c is a bit more difficult, for it includes the clause "for all," that is often called the **universal quantifier** and abbreviated by $\forall$. So clause c states $x^T y = 0 \; \forall \, y$. The negation of "some thing happens for every y" is that "there exists a y for which that thing does not happen." This "there exists" is called the **existential quantifier** and is often abbreviated $\exists$. Hence, the contra–positive of Eq. (1.26) is

$$\text{if } x \in \mathbb{R}^n \quad \text{and} \quad x \neq 0 \quad \text{then} \quad \exists \, y \in \mathbb{R}^n \quad \text{such that} \quad x^T y \neq 0. \tag{1.27}$$

It is a matter of taste, guided by experience, that causes one to favor (or not) the contra–positive over the original. At first sight the student new to proofs and unsure of "where to start" may feel that the two are equally opaque. Mathematics however is that field that is, on first sight, opaque to everyone, but that on second (or third) thought begins to clarify, suggest pathways, and offer insight and rewards. The key for the beginner is not to despair but rather to generate as many starting paths as possible, in the hope that one of them will indeed lead to a fruitful second step, and on to a deeper understanding of what you are attempting to prove. So, investigating the contra–positive fits into our bigger strategy of generating multiple starting points and, even when a dead-end, is a great piece of guilt–free procrastination.

Back to the problem at hand I'd like to point out two avenues "suggested" by Eq. (1.27). The first is the old avenue – "take $y = x$" for then $x \neq 0$ surely implies that $x^T x \neq 0$. The second I feel is more concrete, more pedestrian, less clever and therefore hopefully contradicts the belief that one either "gets the proof or not." The concreteness I speak of is generated by the $\exists$ for it says we only have to find one – and I typically find that easier to do than finding many or all. To be precise, if $x \neq 0$ then a particular element $x_i \neq 0$. From here we can custom build a y, namely choose y to be 0 at each element except for the ith in which you set $y_i = 1$. Now $x^T y = x_i$ which, by not c, is presumed nonzero.

As a final example lets prove that

$$\text{if } A \in \mathbb{R}^{n\times n} \quad \text{and} \quad Ax = 0 \;\; \forall\, x \in \mathbb{R}^n \quad \text{then} \quad A = 0. \tag{1.28}$$

In fact, lets offer three proofs.

The first is a "row proof." We denote row j of A by $A(j,:)$ and notes that $Ax = 0$ implies that the inner product $A(j,:)x = 0$ for every x. By our proof of Eq. (1.26) it follows that the jth row vanishes, i.e., $A(j,:) = 0$. As this holds for each j it follows that the entire matrix is 0.

Our second is a "column proof." We interpret $Ax = 0$, $\forall\, x$, in light of Eq. (1.11), to say that every weighted sum of the columns of A must vanish. So lets get concrete and choose an x that is zero in every element except the jth, for which we set $x_j = 1$. Now Eq. (1.11) and the if clause in Eq. (1.28) reveal that $A(:,j) = 0$, i.e., the jth column vanishes. As j was arbitrary it follows that every column vanishes ans so the entire matrix is zero.

Our third proof will address the contrapositive,

$$\text{if } A \neq 0 \in \mathbb{R}^{n\times n} \quad \text{then} \quad \exists\; x \in \mathbb{R}^n \quad \text{such that} \quad Ax \neq 0. \tag{1.29}$$

We now move concretely and infer from $A \neq 0$ that for some particular i and j that $a_{ij} \neq 0$. We then construct (yet again) an x of zeros except we set $x_j = 1$. It follows (from either the row or column interpretation of Ax) that the ith element of Ax is a_{ij}. As this is not zero we have proven that $Ax \neq 0$.

We next move on to a class of propositions that involve infinity in a substantial way. If there are in fact an infinite number of claims we may use the Principle of Mathematical Induction, if it is a claim about equality of infinite sets then we may use the method of reciprocal inclusion, while if it is a claim about convergence of infinite sequences of vectors we may use the ordering of the reals.

The **Principle of Mathematical Induction** states that the truth of the infinite sequence of statements $\{P(n) : n = 1, 2, \ldots\}$ follows from establishing that
(PMI1) $P(1)$ is true.
(PMI2) if $P(n)$ is true then $P(n+1)$ is true, for arbitrary n.

For example, let us prove by induction that

$$\begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}^n = \begin{pmatrix}1 & n\\ 0 & 1\end{pmatrix} \qquad n = 1, 2, \ldots. \tag{1.30}$$

We first check the base case, here Eq. (1.30) holds by inspection when $n = 1$. We now suppose it holds for some n then deduce its validity for $n + 1$. Namely

$$\begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}^{n+1} = \begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}\begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}^n = \begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}\begin{pmatrix}1 & n\\ 0 & 1\end{pmatrix} = \begin{pmatrix}1 & n+1\\ 0 & 1\end{pmatrix}.$$

Regarding infinite sets, the **Principle of Mutual Inclusion** states that two sets coincide if each is a subset of the other. For example, given an $x \in \mathbb{R}^n$ lets consider the outer product matrix $xx^T \in \mathbb{R}^{n\times n}$ and let us prove that the two sets

$$N_1 \equiv \{y : x^T y = 0\} \quad \text{and} \quad N_2 \equiv \{z : xx^T z = 0\}$$

coincide. If $x = 0$ both sides are simply $\mathbb{R}^n$. So lets assume $x \neq 0$ and check the reciprocal inclusions, $N_1 \subset N_2$ and $N_2 \subset N_1$. The former here looks to be the "easy" direction. For if $x^T y = 0$ then surely $xx^T y = 0$. Next, if $xx^T z = 0$ then $x^T xx^T z = 0$, i.e., $\|x\|^2 x^T z = 0$ which, as $x \neq 0$ implies that $x^T z = 0$.

Regarding **Infinite Sequences** $\{x_n\}_{n=1}^{\infty} \subset \mathbb{R}$ we note that although the elements may change erratically with n we may always extract a well ordered subsequence. For example, from the oscillatory $x_n = (-1)^n/n$ we may extract the decreasing $x_{n_k} \equiv x_{2k} = 1/(2k)$. More generally we call a sequence **monotone** if either $x_n \leq x_{n+1}$ for all n or $x_n \geq x_{n+1}$ for all n. We state and prove the general case:

Proposition 1.1. Given $\{x_n\}_{n=1}^{\infty} \subset \mathbb{R}$ there exists a monotone subsequence $\{x_{n_k}\}_{k=1}^{\infty} \subset \mathbb{R}$.

Proof: Call x_n a peak if $x_n > x_m$ for all $m < n$. If our sequence has no peaks then it is already monotone. If our sequence has an infinite number of peaks (as in our example above) at $n_1 < n_2 < \cdots$ then $x_{n_1} \geq x_{n_2} \geq \cdots$ is a monotone subsequence. It remains to study sequences with at least one but at most finitely many peaks. In this case, if x_N is the peak with the biggest index then x_{n_1} where $n_1 = N + 1$ is not a peak and so $\exists$ and $n_2 > n_1$ such that $x_{n_2} \geq x_{n_1}$. In the same fashion, as n_2 is not a peak $\exists$ and $n_3 > n_2$ such that $x_{n_3} \geq x_{n_2}$. On repetition this procedure generates an infinite monotone subsequence. End of Proof.

The great attraction of (bounded) monotone sequences is that they must converge to their smallest or largest value. To make this precise we call u an upper bound for $\{x_n\}$ if $x_n \leq u$ for all n and we denote by x^u the **least upper bound**. For example, 1 is the least upper bound of $\{1 - 1/n\}_n$.

Proposition 1.2. If $\{x_n\}_n$ is monotonically nondecreasing and x^u is its least upper bound then

$$\lim_{n\to\infty} x_n = x^u.$$

That is, given any $\varepsilon > 0$ $\exists$ $N > 0$ such that $|x_n - x^u| \leq \varepsilon$ $\forall$ $n > N$. We often abbreviate this as $x_n \to x^u$.

Proof: Given $\varepsilon > 0$ if there exists an $N > 0$ such that $x_n \le x^u - \varepsilon$ for $n > N$ then $x^u - \varepsilon/2$ is an upper bound less than x^u, contrary to its definition. End of Proof.

In a similar fashion we call ℓ a lower bound for $\{x_n\}$ if $x_n \ge \ell$ for all n and we denote by x^ℓ the **greatest lower bound**. For example, 0 is the greatest lower bound of $\{1/n\}_n$. If $\{x_n\}$ is nonincreasing then $x_n \to x^\ell$. Combining these last two propositions we find that every bounded sequence of real numbers has a convergent subsequence. Our argument in fact translates nicely to vectors.

Proposition 1.3. If $\{x_j\}_j \subset \mathbb{R}^n$ and there exists a finite M for which $\|x_j\| \le M$ for all j then there exists a subsequence $\{x_{j_k}\} \subset \{x_j\}_j$ and an $x \in \mathbb{R}^n$ such that $x_{j_k} \to x$. That is given any $\varepsilon > 0 \; \exists \; N > 0$ such that $\|x_{j_k} - x\| \le \varepsilon \; \forall \; j_k > N$.

Proof: We note the elements of x_j by $x_j(1)$ through $x_j(n)$. As $\{x_j(1)\}_j$ is a bounded sequence in $\mathbb{R}$ it has a subsequence, $\{x_{j_k}(1)\}_j$, that converges to a number that we label $x(1)$. As $\{x_{j_k}(2)\}_j$ is a bounded sequence in $\mathbb{R}$ it has a subsequence, $\{x_{j_{kl}}(2)\}_l$, that converges to a number that we label $x(2)$. Moreover, this new subsequence does not affect the convergence of the first element. In particular, $x_{j_{kl}}(1) \to x(1)$ as $l \to \infty$. We now continue to extract a subsequence from the previous sequence until we have exhausted all n dimensions. End of Proof.

Our first application of this is to an alternate notion of matrix norm. We observed in Eq. (1.20) that the Frobenius norm is larger than the biggest stretch. The word "biggest" suggest that we are looking for the least upper bound. This three word phrase is a bit awkward and so is often rephrased as *supremum* which itself it abbreviated to sup. All this suggests that

$$\|A\| \equiv \sup_{\|x\|=1} \|Ax\| \tag{1.31}$$

is worthy of study. By definition there exists a sequence $\{x_j\}_j$ of unit vectors for which $\|Ax_j\| \to \|A\|$. By Prop. 1.3 there exists a convergent subsequence, $x_{j_k} \to \tilde{x}$. It follows that $\|x_{j_k}\| \to \|\tilde{x}\|$ and so $\|\tilde{x}\| = 1$. In addition,

$$\|Ax_{j_k} - A\tilde{x}\| = \|A(x_{j_k} - \tilde{x})\| \le \|A\|_F \|x_{j_k} - \tilde{x}\|$$

permits us to conclude that $Ax_{j_k} \to A\tilde{x}$ and so $\|Ax_{j_k}\| \to \|A\tilde{x}\|$ and recalling $\|Ax_{j_k}\| \to \|A\|$ we conclude that $\|A\tilde{x}\| = \|A\|$. The upshot is that the *supremum* in Eq. (1.31) is actually attained. We distinguish this fact by writing

$$\boxed{\|A\| \equiv \max_{\|x\|=1} \|Ax\|.} \tag{1.32}$$

By definition we know that $\|A\| \le \|A\|_F$ for every matrix. A simple example that shows up the disparity involves I_n, the identity matrix on $\mathbb{R}^n$. Please confirm that $\|I_n\| = 1$ while $\|I_n\|_F = \sqrt{n}$.

1.4. Notes and Exercises

For thousands more worked examples I recommend Lipschutz (1989). Higham and Higham (2005) is an excellent guide to Matlab . For a more thorough guide to proofs please see Velleman (2006).

1. Consider the matrix

$$A = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \tag{1.33}$$

Evaluate the product Ax for several choices of x. Sketch both x and Ax in the plane for several carefully marked x and explain why A is called a "rotation." Argue, on strictly geometric grounds, why $A^5 = A$.

2. Consider the matrix

$$A = \begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}. \tag{1.34}$$

Evaluate the product Ax for several choices of x. Sketch both x and Ax in the plane for several carefully marked x and explain why A is called a "reflection." Argue, on strictly geometric grounds, why $A^3 = A$.

3. We will consider the action of

$$A = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}, \tag{1.35}$$

on the bow-ties, E and F, of Eq. (1.23).

(a) Show that if $x \in E$ then $Ax \in F$,

(b) Show that if $x \in F$ then $Bx \in E$.

(c) Prove by induction that

$$A^n = \begin{pmatrix} 1 & 2n \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad B^n = \begin{pmatrix} 1 & 0 \\ 2n & 1 \end{pmatrix},$$

for positive integer n.

(d) Use (c) to generalize (a) and (b). That is, show that if $x \in E$ then $A^n x \in F$ while if $x \in F$ then $B^n x \in E$ for all positive integer n.

4. We will make frequent use of the **identity matrix**, $I \in \mathbb{R}^{n\times n}$, comprised of zeros off the diagonal and ones on the diagonal. In symbols, $I_{ij} = 0$ if $i \neq j$, while $I_{ii} = 1$. Prove the two propositions, if $A \in \mathbb{R}^{n\times n}$ then $AI = IA = A$. The identity also gives us a means to define the **inverse** of a matrix. One (square) matrix is the inverse of another (square) matrix if their product is the identity matrix. Please show that

$$A^{-1} = \begin{pmatrix} 1 & -2 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad B^{-1} = \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix}, \tag{1.36}$$

are the inverses of the A and B matrices of Eq. (1.36).

5. Write a Matlab program to investigate the shear of the integer diamond by the A and B matrices, Eq. (1.35), and their inverses, Eq. (1.36). More precisely, write a program that generates Figure 1.6.

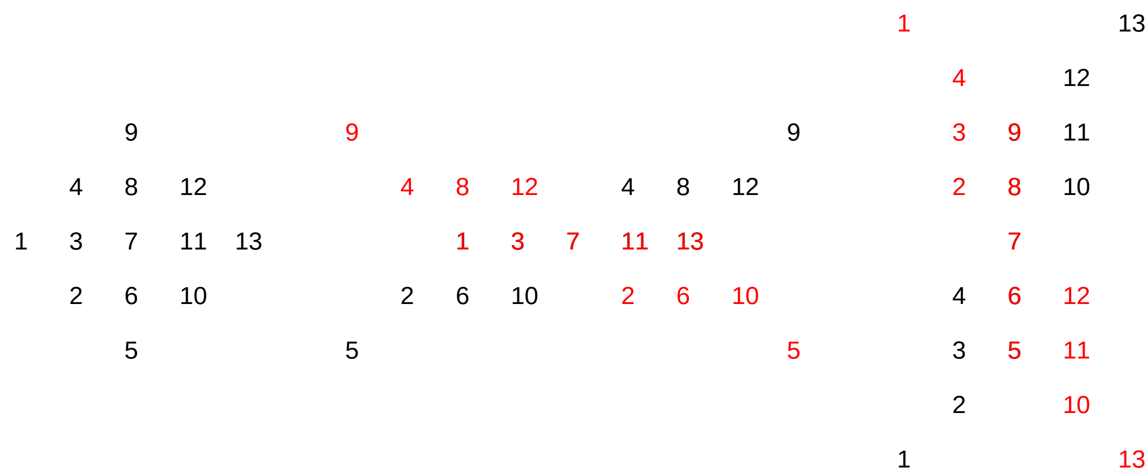

Figure 1.6. Shearing the integral diamond. (Left) The labels are at integral points, $1 = (-2,0)$, $2 = (-1,-1)$, $3 = (-10)$, $4 = (-1,1)$ and so on. (Center) Transformation by A (black) and A^{-1} (red) of the points in panel (Left). (Right) Transformation by B (black) and B^{-1} (red) of the points in panel (Left).

6. We can view, see Figure 1.7(A), vector sums as parallelogram generators. Please show that the area of this parallelogram is $ad - bc$. Show all of your work.

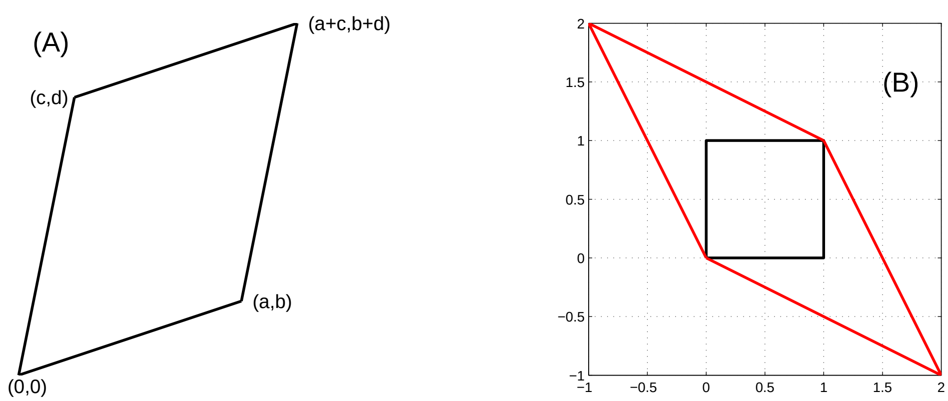

Figure 1.7. (A) The vectors (a,b) and (c,d) drawn from the origin, $(0,0)$, sum to the fourth vertex of a parallelogram. (B) A black square and its deformation (red diamond) by the matrix in (1.37)

7. Show that

$$A = \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix} \tag{1.37}$$

takes the black square to the red diamond in Figure 1.7(B). Use the previous exercise to compute the area of the red diamond.

8. Each of the following chapters will demonstrate the fundamental role that matrices play in modeling the world. Perhaps one of the simplest contexts is in the field of information retrieval. Here one has m "terms" and n "documents" and builds a so–called term-by-document matrix A where a_{ij} is the number of times that term i appears in document j. In Figure 1.8(A) below we depict such a matrix where the documents are the 81 chapters of the *Tao Te Ching* and our

10 terms are *heaven, virtue, nature, life, knowledge, understand, fear, death, good*, and *right*. This matrix is then used to process new queries. For example, if the disciple is looking for the chapters most expressive of *virtue* and *good* then, as these are the second and ninth of our our terms we build the query vector

$$q = (0\ 1\ 0\ 0\ 0\ 0\ 0\ 0\ 1\ 0) \tag{1.38}$$

and search for means to compare this to the columns of A. The standard approach is to exploit the geometric interpretation (recall Eq. (1.5)) of the inner product and to so rank the chapters by the cosine of the angle they make with the query. More precisely, for the jth document we compute

$$\cos(\theta_j) = \frac{qa_j}{\|q\|\|a_j\|}. \tag{1.39}$$

and present these scores in Figure 1.8(B). As small angles correspond to values of cosine near 1 our analysis would direct the disciple to chapter 49 of the *Tao Te Ching*. Typically a threshold is chosen, e.g., 0.8, and a rank ordered list of all documents that exceed that threshold is returned.

Please change `tao.m` to find the chapter most expressive of *heaven, nature* and *knowledge*.

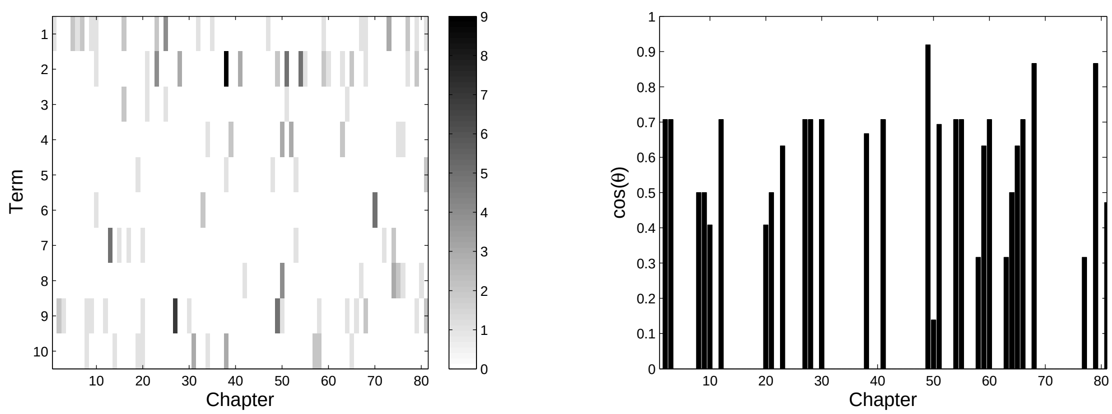

Figure 1.8. Query matching. (A) The 10×81 term-by-document matrix for the *Tao Te Ching*, illustrated with the help of the MATLAB command `imagesc`. (B) The cosine scores associated with the query in Eq. (1.38) as expressed in Eq. (1.39). `tao.m`

9. Prove that matrix multiplication is associative, i.e., that $(AB)C = A(BC)$.

10. Prove that if x and y lie in $\mathbb{R}^n$ and $A \in \mathbb{R}^{n\times n}$ then

$$x^T Ay = y^T A^T x.$$

Hint: The left side is a number. Now argue as we did in achieving Eq. (1.16).

11. Suppose that $A \in \mathbb{R}^{n\times n}$ and $x^T Ax = 0\ \forall\ x \in \mathbb{R}^n$. Does this imply that $A = 0$? If so, prove it. If not, offer a counterexample.

12. Prove that $\text{tr}(A + B) = \text{tr}(A) + \text{tr}(B)$.

13. Use Eq. (1.14) to prove that the fundamental commutator relation of Quantum Mechanics,

$$AB - BA = I,$$

can *not* hold for matrices.

14. Construct a nonzero $A \in \mathbb{R}^{2\times 2}$ for which $A^2 = 0$.

15. A matrix that equals its transpose is called **symmetric**. Suppose $S = A^T G A$ where $A \in \mathbb{R}^{m\times n}$ and $G \in \mathbb{R}^{m\times m}$. Prove that if $G = G^T$ then $S = S^T$.

16. Establish the **triangle inequality**

$$\|x+y\| \le \|x\| + \|y\| \qquad \forall\; x, y \in \mathbb{R}^n. \tag{1.40}$$

First draw this for two concrete planar x and y and discuss the aptness of the name. Then, for the general case expand $\|x+y\|^2$, invoke the Cauchy–Schwarz inequality, Eq. (1.7), and finish with a square root.

17. The other natural vector product is the **outer product**. Note that if $x \in \mathbb{R}^n$ then the outer product of x with itself, xx^T, lies in $\mathbb{R}^{n\times n}$. Please prove that $\|xx^T\|_F = \|x\|^2$.

18. The outer product is also a useful ingredient in the **Householder Reflection**

$$H = I - 2xx^T, \tag{1.41}$$

associated with the unit vector x.

(a) How does H transform vectors that are multiples of x?

(b) How does H transform vectors that are orthogonal to x?

(c) How does H transform vectors that are neither colinear with nor orthogonal to x? Illustrate your answers to (a-c) with a careful drawing.

(d) Confirm that $H^T = H$ and that $H^2 = I$.

19. There is a third way of computing the product of two vectors in $\mathbb{R}^3$, perhaps familiar from vector calculus. The **cross product** of u and v is written $u \times v$ and defined as the matrix vector product

$$u \times v \equiv \mathtt{X}(u)v = \begin{pmatrix} 0 & -u_3 & u_2 \\ u_3 & 0 & -u_1 \\ -u_2 & u_1 & 0 \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} = \begin{pmatrix} -u_3v_2 + u_2v_3 \\ u_3v_1 - u_1v_3 \\ -u_2v_1 + u_1v_2 \end{pmatrix}$$

(a) How does $\mathtt{X}(u)$ transform vectors that are multiples of u?

(b) How does $\mathtt{X}(u)$ transform vectors that are orthogonal to u?

(c) How does $\mathtt{X}(u)$ transform vectors that are neither colinear with nor orthogonal to u? Illustrate your answers to (a-c) with a careful drawing. You may wish to use the MATLAB function `cross`.

(d) Confirm that $\mathtt{X}(u)^T = -\mathtt{X}(u)$ and that $\mathtt{X}(u)^2 = uu^T - \|u\|^2 I$.

(e) Use (d) to derive

$$\|u \times v\|^2 = \|u\|^2\|v\|^2 - (u^Tv)^2.$$

(f) If θ is the angle between u and v use (e) and (1.5) to show that

$$\|u \times v\| = \|u\|\|v\||\sin\theta|.$$

(g) Use (f) and Figure 1.9(A) to conclude that $\|u \times v\|$ is the area (base times height) of the parallelogram with sides u and v.

(h) Use (g) and Figure 1.9(B) to conclude that $|w^T(u \times v)|$ is the volume (area of base times height) of the parallelepiped with sides u, v and w. Hint: Let u and v define the base. Then $u \times v$ is parallel to the height vector obtained by drawing a perpendicular from w to the base.

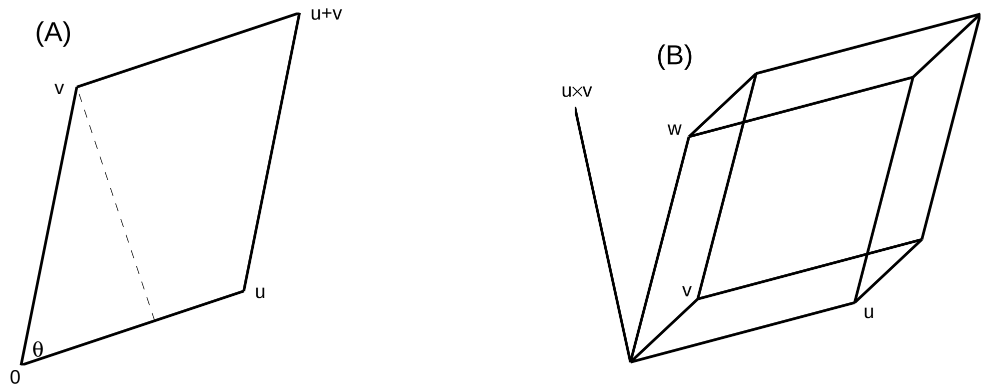

Figure 1.9. (A) Parallelogram. (B) Parallelepiped.

20. Show that if $A \in \mathbb{R}^{m\times n}$ and $B \in \mathbb{R}^{n\times p}$ then $\|AB\|_F \le \|A\|_F\|B\|_F$. Hint: Adapt the proof of Eq. (1.19).

21. Via experimentation with small n arrive (show your work) at a formula for f_n in

$$\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}^n = \begin{pmatrix} 1 & n & f_n \\ 0 & 1 & n \\ 0 & 0 & 1 \end{pmatrix}$$

and prove, via induction, that your formula holds true for all n.

22. Suppose that $\{a_j : j = 0, \pm 1, \pm 2, \ldots\}$ is a doubly infinite sequence. Prove, via induction, that

$$\sum_{j=0}^{n}\sum_{k=0}^{n} a_{j-k} = \sum_{m=-n}^{n} (n+1-|m|)a_m. \tag{1.42}$$

23. For the matrix of (1.37) compute, by hand and showing all work, that $\|A\| = 3$ and $\|A\|_F = \sqrt{10}$. Hint: For the former, choose $x = (\cos(\theta), \sin(\theta))^T$ and show that $\|Ax\|^2 = 5 - 8\cos(\theta)\sin(\theta)$. Now take a derivative in order to find the θ that gives the largest $\|Ax\|$.