

An Introduction to Advanced Topics in Linear Algebra

Sam Roven and Rekha Thomas

Contents

0.1	Introduction	4
1	Systems of Linear Equations	5
1.1	Solving Linear Equations	5
1.2	Linear Systems and Matrices	8
2	Euclidean Space	12
2.1	Vectors	12
2.2	Span	15
2.2.1	Some modern math techniques	17
2.3	Linear Independence	22
3	Linear Transformations	26
3.1	The Basics of Linear Maps	26
3.2	Matrix Algebra	32
3.3	Inverses	36
4	Basis and Subspaces	39
4.1	Subspaces	39
4.2	Basis and Dimension	42
4.3	Row Space, Column Space, and Rank	46
4.4	Change of Basis	49
5	Determinants	52
5.1	The Determinant Function	52
5.2	Properties of the Determinant	53
6	Eigenvalues and Diagonalization	56
6.1	Eigenvalues and Eigenvectors	56
6.2	Diagonalization	60
7	Review of Eigenvalues and Diagonalization	65
7.1	Eigenvalues and Eigenvectors	65
7.1.1	Complex Eigenvalues	68
7.2	Diagonalization	68
7.3	Similar Matrices	70
8	Polynomial Vector Spaces	71
8.1	$\mathbb{R}[x]$: The polynomial vector space	71
8.1.1	Linear Maps between Polynomial Vector Spaces	72
8.2	Some Helpful Logic Facts	74
8.3	Problem Set 1	75

9	Positive Matrices and much much more	79
9.1	Difference Equations	79
9.2	Positive Matrices	82
9.3	Applications: Adjacency Matrices and Google Page Rank	86
9.3.1	Google Page Rank	87
9.4	Problem Set 2	89
10	Orthogonality and Projections	94
10.1	Orthogonal Subspaces	94
10.2	Orthogonal complements	98
10.3	Projections	100
10.4	Applications: Least Squares Regression	102
10.5	Problem Set 3	106
11	Symmetric Matrices: The Most Important Matrices You'll Ever See	112
11.1	Spectral Theorem	112
12	Change of Basis Review and Why Orthogonal Diagonalization is So Great	116
12.1	Change of Basis	116
12.2	Linear Maps in Different Bases	117
12.3	Cramer's Rule	119
12.4	Problem Set 4	120
13	Positive Definite/Positive Semi-Definite Matrices	124
13.1	Motivation and Definitions	124
13.2	Laplacian of a Graph and a Peek into Spectral Gap Theory	129
13.3	Application: Distance Realization	133
13.4	Problem Set 5	135
14	The Singular Value Decomposition	142
14.1	Motivation	142
14.2	The Mechanics of the Singular Value Decomposition	143
14.3	The Geometry of the Singular Value Decomposition	146
14.4	The Reduced Singular Value Decomposition	151
14.5	Matrix Norms	153
14.6	Rank One Decompositions	154
14.7	Best Fit k -Planes	157
14.8	Application: Principal Component Analysis	161
14.9	Problem Set 6	162
15	Quotients of Vector Spaces and Error Detecting Codes	171
15.1	Quotient Spaces	171
15.2	Quotients of Polynomial Vector Spaces	177
15.3	Application: Solving a Univariate Polynomial	180
15.4	Problem Set 7	183
16	Error Correcting Codes	185
16.1	Vector spaces over finite fields	185
16.2	Error Correcting Codes	186
16.3	Problem Set 8	189

17 Vector Spaces over \mathbb{C}	193
17.1 Complex Numbers	193
17.2 Complex Vectors and the Vector Space \mathbb{C}^n	194
17.3 $\mathbb{C}^{m \times n}$ and the Complex Spectral Theorem	196
17.4 Problem Set 9	198

0.1 Introduction

More to come on this page. For now just keep reading!

Chapter 1

Systems of Linear Equations

1.1 Solving Linear Equations

We have all seen a linear system of equations at some point in gradeschool, and we first learned how to attack these systems using the methods of substitution and elimination. We begin with a refreshing example of a linear system with **three** equations and **three** unknowns.

Example 1.1.1.
$$\begin{cases} x_1 + x_2 - x_3 = 7 \\ 2x_1 + 3x_3 = 5 \\ -5x_2 = -10 \end{cases}$$

We define a solution of this system as an ordered triple of real numbers, $\underbrace{(x_1, x_2, x_3)}_{\text{also called a 3-tuple}}$, which simultaneously satisfies **all** equations.

One solution of this system is $(4, 2, -1)$ because

$$4 + 2 - (-1) = 7$$

$$2(4) + 3(-1) = 5$$

and

$$-5(2) = -10$$

It is also worth noting that this solution is the same thing as an ordinary point in 3-dimensional Euclidean space, and we are immediately able to talk about geometry (much more to come).

Next, we must get our hands around the vocabulary of linear systems, the first of which is distinguishing one type of variable from another.

Definition 1.1.2. A variable that appears as the first (left-most) term of at least one equation is a **leading variable**. In the above example, x_1 and x_2 are leading variables.

Definition 1.1.3. If a linear system has no solutions, then it is **inconsistent**. If a linear system has at least one solution, then it is **consistent**.

Example 1.1.4.
$$\begin{cases} 2x_1 - 3x_2 + x_3 = 8 \\ 2x_2 = 5 \end{cases}$$

This is an example of a linear system with infinitely many solutions. In fact, for any real number t , the tuple

$$\left(\frac{31}{4} - \frac{1}{2}t, \frac{5}{2}, t\right)$$

represents a solution and we can verify directly that it is a solution by plugging in $\frac{31}{4} - \frac{1}{2}t$ for x_1 , $\frac{5}{2}$ for x_2 , and t for x_3 and checking that all the t 's cancel to give equality.

Here, t is called a **free variable** or **free parameter**.

Now that we have some language to work with, we will need to investigate the possible forms a system can have. In particular, there are two.

1. Triangular form: An example of a linear system in triangular form is

$$\begin{cases} 4x_1 - 2x_2 + 3x_3 + x_4 = 17 \\ x_2 - 2x_3 - x_4 = 0 \\ 5x_3 + 2x_4 = 20 \\ 3x_4 = 15 \end{cases}$$

We can solve a system like this using **back substitution** (using the last equation first). In doing this we see that

$$3x_4 = 15 \implies x_4 = 5$$

We then apply this to the third equation and get

$$5x_3 + 2x_4 = 5x_3 + 2(5) = 20 \implies 5x_3 = 10 \implies x_3 = 2$$

Applying the same procedure to the second and first equation we find that $x_2 = 9$ and $x_1 = 6$ (you should verify this for yourself!). The final solution is then given by

$$(x_1, x_2, x_3, x_4) = (6, 9, 2, 5)$$

In general, triangular forms have three main properties:

- There are the same number of equations as variables.
- Every variable is the leading variable of exactly one equation.
- A triangular system has **exactly one solution**. We refer to this as a unique solution.

2. Echelon Form

This is the more general form that a linear system can have and we can characterize it according to two (or three) main properties:

- Every variable is the leading variable of at most one equation.
- The system is organized in a descending stair-step pattern.

If a linear system satisfies both of these properties then we say the system is in **echelon form**. The last property of a system in echelon form is

- There are either no solutions, exactly one solution, or infinitely many solutions.

To build off of the last point, we can actually say something more general.

Theorem 1.1.5. Any system of linear equations has either

- No solutions (this is known as an **inconsistent** linear system).
- Exactly one solution.
or
- Infinitely many solutions.

The latter two cases define what we call a **consistent** linear system.

Example 1.1.6. The following linear system **is** in echelon form.

$$\begin{cases} 3x_1 & - x_3 = 7 \\ & x_2 & = 10 \end{cases}$$

Example 1.1.7. The following linear system **is not** in echelon form because the linear equations do not form a stair-step pattern.

$$\begin{cases} 3x_1 + x_2 - x_3 = 7 \\ & x_3 = 5 \\ & x_2 + 9x_3 = 11 \end{cases}$$

Example 1.1.8. The following linear system **is not** in echelon form because x_1 is the leading variable of more than one equation.

$$\begin{cases} 3x_1 & - x_3 = 7 \\ x_1 + x_2 & = 9 \end{cases}$$

Definition 1.1.9. For a system in echelon form, any variable that does not appear as a leading variable is called a **free variable**, hence all variables in a system are either leading or free.

Here are some nice facts to remember about systems in echelon form.

1. If an echelon system has no free variables, it must be triangular and therefore has exactly one solution.
2. If an echelon system has at least one free variable, then it has infinitely many solutions.

Now that we have much of the needed vocabulary, let's end the section with a fully worked example.

Example 1.1.10. Consider the linear system

$$\begin{cases} 2x_1 - x_2 + 5x_3 - x_4 = -30 \\ & x_3 + x_4 = -6 \end{cases}$$

This is a system in echelon form with x_1, x_3 as leading variables and x_2, x_4 as free variables. We solve the system in two steps.

Step 1: Denote free variables. Let $x_2 = t_1$ and $x_4 = t_2$ and remember these can be any real number!

Step 2: Plug the free variables into the system and solve for leading variables. Starting with the second equation we have

$$x_3 + t_2 = -6 \implies x_3 = -t_2 - 6$$

Plugging this into the first equation we have

$$2x_1 - t_1 + 5(-t_2 - 6) - t_2 = -30 \implies 2x_1 - t_1 - 5t_2 - t_2 = 0 \implies 2x_1 - t_1 - 6t_2 = 0$$

Using this to solve for x_1 we get

$$2x_1 = t_1 + 6t_2 \implies x_1 = \frac{1}{2}t_1 + 3t_2$$

The (infinitely many) solutions of this linear system have the form

$$(x_1, x_2, x_3, x_4) = \left(\frac{1}{2}t_1 + 3t_2, t_1, -t_2 - 6, t_2\right)$$

with t_1, t_2 as free variables.

1.2 Linear Systems and Matrices

In this section we dive deeper into the procedures for solving linear systems and in the process, encounter matrices for the first time. These procedures will transform any linear system into one in echelon form and produce a new linear system with the exact same solution set.

Definition 1.2.1. Two linear systems are **equivalent** if they have the same solution set. The notion of being equivalent is denoted with the symbol “ \sim ”.

The way in which we get from an arbitrary linear system to an echelon one is by applying **elementary row operations**. These consist of three possible “moves” that transform a system into an equivalent one:

1. Interchange two equations.
2. Replace one equation with a non-zero multiple of itself.
3. Add one equation to a multiple of another.

Example 1.2.2.
$$\begin{cases} -4x_1 + 5x_2 = 20 \\ x_1 - 2x_2 = 14 \end{cases}$$

$$\sim \begin{cases} x_1 - 2x_2 = 14 \\ -4x_1 + 5x_2 = 20 \end{cases} \text{ (interchange equations)}$$

$$\sim \begin{cases} 4x_1 - 8x_2 = 56 \\ -4x_1 + 5x_2 = 20 \end{cases} \text{ (multiply equation 1 by 4)}$$

$$\sim \begin{cases} 4x_1 - 8x_2 = 56 \\ -3x_2 = 76 \end{cases} \text{ (add equation 1 to equation 2)}$$

Notice the last (equivalent) system is in echelon form!

This example illustrates the general procedure, but the main tool that we use to streamline the procedure is that of augmented matrices. When we solve a linear system, we are only working with the coefficients of the linear equations, so we place the coefficients in an array called an **augmented matrix**.

Example 1.2.3. The linear system

$$\begin{cases} x_1 - 2x_2 + 3x_3 = 9 \\ -x_1 + 3x_3 = -4 \\ 2x_1 - 5x_2 + 5x_3 = 17 \end{cases}$$

has associated augmented matrix given by

$$\left[\begin{array}{ccc|c} 1 & -2 & 3 & 9 \\ -1 & 0 & 3 & -4 \\ 2 & -5 & 5 & 17 \end{array} \right]$$

We can now translate vocabulary from linear systems into that for matrices. Once we do this we will never look back. Similar to that of linear systems, there are two special types of matrices.

1. Echelon Form.

- Every leading term (the first nonzero number in a row) is in a column to the **left** of the leading term of the row below it.
- Any zero rows (rows of all zeroes) are at the bottom.

In general, we call any leading term of a non-zero row a **pivot**.

Example 1.2.4.

$$\begin{bmatrix} 3 & 0 & 4 & 5 \\ 0 & 1 & 3 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

is a matrix in echelon form with pivots being the entries 3 and 1.

The matrices

$$\begin{bmatrix} 0 & 1 & 0 & 3 \\ 4 & 5 & 6 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 2 \end{bmatrix}, \begin{bmatrix} 1 & 2 & 3 & 4 \\ 2 & 3 & 4 & 5 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

are all not in echelon form. Can you see why?

2. Reduced Echelon Form.

- It is in echelon form.
- All pivot positions contain a 1.
- All other entries in a pivot column (a column that contains a pivot) are 0.

Example 1.2.5. $\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 2 \\ 0 & 0 & 0 & 0 \end{bmatrix}$ is a matrix that is in reduced echelon form.

$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 2 & 1 \end{bmatrix}$ is neither in echelon nor reduced echelon form.

$\begin{bmatrix} 2 & 0 & 1 & 3 \\ 0 & -1 & 1 & 4 \\ 0 & 0 & 0 & 1 \end{bmatrix}$ is in echelon form but **not** in reduced echelon form.

$\begin{bmatrix} 1 & 0 & 2 & 1 \\ 0 & 1 & 3 & 4 \\ 0 & 0 & 1 & 0 \end{bmatrix}$ is also in echelon form but **not** in reduced echelon form.

When working through a solution to a linear system, we can easily follow our own steps by adopting the following notation for row operations.

1. Interchange row i and row j is denoted

$$R_i \leftrightarrow R_j$$

2. Replacing row i with a non-zero multiple (c) times row j is denoted

$$cR_i \rightarrow R_i$$

3. Adding a non-zero multiple of row i to row j and applying the change to row j is denoted

$$cR_i + R_j \rightarrow R_j$$

In practice, we will use these row operations to transform augmented matrices into systems that are in echelon or reduced echelon form, at which point we will be able to solve them by back substitution.

This whole process will be most easily learned via examples so let's jump right in with a continuation of Example 1.2.3.

Example 1.2.6.

$$\begin{aligned} & \left[\begin{array}{ccc|c} 1 & -2 & 3 & 9 \\ -1 & 0 & 3 & -4 \\ 2 & -5 & 5 & 17 \end{array} \right] \\ (R_1 + R_2 \rightarrow R_2) & \implies \left[\begin{array}{ccc|c} 1 & -2 & 3 & 9 \\ 0 & -2 & 6 & 5 \\ 2 & -5 & 5 & 17 \end{array} \right] \\ (-2R_1 + R_3 \rightarrow R_3) & \implies \left[\begin{array}{ccc|c} 1 & -2 & 3 & 9 \\ 0 & -2 & 6 & 5 \\ 0 & -1 & -1 & -1 \end{array} \right] \\ (-\frac{1}{2}R_2 \rightarrow R_2) & \implies \left[\begin{array}{ccc|c} 1 & -2 & 3 & 9 \\ 0 & 1 & -3 & -5/2 \\ 0 & -1 & -1 & -1 \end{array} \right] \\ (R_2 + R_3 \rightarrow R_3) & \implies \underbrace{\left[\begin{array}{ccc|c} 1 & -2 & 3 & 9 \\ 0 & -2 & 6 & 5 \\ 0 & 0 & -4 & -7/2 \end{array} \right]}_{\text{echelon form!}} \end{aligned}$$

This matrix represents the (triangular) linear system

$$\begin{cases} x_1 - 2x_2 + 3x_3 = 9 \\ x_2 + 6x_3 = -5/2 \\ -4x_3 = -7/2 \end{cases}$$

hence we can use back substitution to obtain the (unique) solution

$$(x_1, x_2, x_3) = (-113/8, -41/4, 7/8)$$

Definition 1.2.7. The process of using row operations (like above) to transform a matrix into echelon form is called **Gaussian Elimination**.

We can take this one step further, if we prefer, by reducing the given matrix to **reduced** echelon form. This is known as **Gauss-Jordan Elimination**.

Example 1.2.8. Use Gauss-Jordan elimination to solve the linear system

$$\begin{cases} x_1 - 3x_3 = -2 \\ 3x_1 + x_2 - 2x_3 = 5 \\ 2x_1 + 2x_2 + x_3 = 4 \end{cases}$$

We begin with the augmented matrix for this linear system and write a string of equivalent matrices, ending with the reduced echelon form. We leave the row operations to be determined by the reader.

$$\left[\begin{array}{ccc|c} 1 & 0 & -3 & -2 \\ 3 & 1 & -2 & 5 \\ 2 & 2 & 1 & 4 \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 0 & -3 & -2 \\ 0 & 1 & 7 & 11 \\ 0 & 2 & 7 & 8 \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 0 & -3 & -2 \\ 0 & 1 & 7 & 11 \\ 0 & 0 & -7 & -14 \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 0 & -3 & -2 \\ 0 & 1 & 7 & 11 \\ 0 & 0 & 1 & 2 \end{array} \right]$$

We note here that at this point we could stop and use back substitution, we have performed Gaussian elimination and have arrived at the echelon form. Continuing onward we have

$$\sim \left[\begin{array}{ccc|c} 1 & 0 & -3 & -2 \\ 0 & 1 & 0 & -3 \\ 0 & 0 & 1 & 2 \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 0 & 0 & 4 \\ 0 & 1 & 0 & -3 \\ 0 & 0 & 1 & 2 \end{array} \right]$$

Translating back to the linear system, we have the (unique) solution

$$(x_1, x_2, x_3) = (4, -3, 2)$$

Before ending the chapter, we note that there is a methodical way to clear out entries of augmented matrices, starting in the upper left corner, moving down column 1, then to the entry in the second column and second row, then down the entire second column, etc. Having a methodical approach to row reductions will reduce errors and make row reductions much easier with a little practice.

Chapter 2

Euclidean Space

We now translate from the algebraic nature of linear systems to their underlying geometry. We begin with a quick refresher on vectors and Euclidean space, then spend the majority of the chapter introducing the all important notions of span and linear independence.

2.1 Vectors

Vectors are the fundamental object of linear algebra and we will use them frequently.

Definition 2.1.1. A **vector** is an ordered list of real numbers that can be expressed in two ways:

- Column vector

$$\mathbf{u} = \begin{bmatrix} u_1 \\ u_2 \\ \vdots \\ u_n \end{bmatrix}$$

- Row vector

$$\mathbf{u} = (u_1, u_2, \dots, u_n)$$

We will use column vectors most of the time, but it is good to know that both notations can mean the same thing.

Just like with real numbers, we can perform arithmetic with vectors. Let $\mathbf{u} = (u_1, u_2, \dots, u_n)$ and $\mathbf{v} = (v_1, v_2, \dots, v_n)$ with c a real number (also known as a scalar).

We can multiply vectors by scalars as follows

$$c\mathbf{u} = \begin{bmatrix} cu_1 \\ \vdots \\ cu_n \end{bmatrix}$$

We can also add two vectors, as long as they have the same number of coordinates.

$$\mathbf{u} + \mathbf{v} = \begin{bmatrix} u_1 + v_1 \\ u_2 + v_2 \\ \vdots \\ u_n + v_n \end{bmatrix}$$

Lastly, $\mathbf{u} = \mathbf{v}$ if and only if $u_1 = v_1, u_2 = v_2, \dots, u_n = v_n$.

Definition 2.1.2. The set of all vectors with n entries (components), together with the above operations of scalar multiplication and vector addition, form what is known as n -dimensional Euclidean space. We denote this space by \mathbb{R}^n . For the vectors \mathbf{u} and \mathbf{v} defined above, we use the symbol “ \in ” to denote that the vector **lives in** \mathbb{R}^n . Similarly, since the scalar c is a real number, it **lives in** the set of real numbers, which we denote by writing $c \in \mathbb{R}$. We will use this notation **frequently** from now on.

In \mathbb{R}^2 and \mathbb{R}^3 we usually represent vectors with “arrows”. The previous three vector properties can also be expressed geometrically.

- Two vectors are equal if and only if they have the same **length** and point in the same **direction**.
- Given a vector \mathbf{u} , the vector $c\mathbf{u}$ (for $c \neq 0$, and $c \in \mathbb{R}$) is parallel to \mathbf{u} , with length equal to $|c|$ times the length of \mathbf{u} . Multiplying a vector by a negative scalar switches the direction that it points in.
- Given $\mathbf{u}, \mathbf{v} \in \mathbb{R}$, the vector $\mathbf{u} + \mathbf{v}$ can be found by using the usual parallelogram law (or tip-to-tail rule) from calculus 3.

Now that we have the fundamentals refreshed, we can move onto one of the central topics of the course.

Definition 2.1.3. If $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m \in \mathbb{R}^n$ and $c_1, c_2, \dots, c_m \in \mathbb{R}$ then the vector

$$c_1\mathbf{u}_1 + c_2\mathbf{u}_2 + \dots + c_m\mathbf{u}_m$$

is called a **linear combination** of $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m$.

Example 2.1.4. Given the vectors $\mathbf{u}_1 = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$ and $\mathbf{u}_2 = \begin{bmatrix} 5 \\ -3 \end{bmatrix}$, three different linear combinations of \mathbf{u}_1 and \mathbf{u}_2 are

$$\mathbf{u}_1 + \mathbf{u}_2 = \begin{bmatrix} 6 \\ -1 \end{bmatrix}, \mathbf{u}_1 - \mathbf{u}_2 = \begin{bmatrix} -4 \\ 5 \end{bmatrix}, 2\mathbf{u}_1 + 30\mathbf{u}_2 = \begin{bmatrix} 152 \\ -86 \end{bmatrix}$$

A very important idea tied to linear combinations is finding when a given vector is a linear combination of a fixed set of vectors.

Example 2.1.5. Let $\mathbf{v}_1 = \begin{bmatrix} 1 \\ 0 \\ -2 \end{bmatrix}$, $\mathbf{v}_2 = \begin{bmatrix} 2 \\ 1 \\ 1 \end{bmatrix}$, and $\mathbf{v}_3 = \begin{bmatrix} 5 \\ 2 \\ -1 \end{bmatrix}$ and determine if $\mathbf{b} = \begin{bmatrix} 19 \\ 7 \\ -9 \end{bmatrix}$ is a linear combination of $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$.

When approaching a question like this one, starting the problem is often the hardest part. How in the world can we figure this out? We figure it out by assuming it is true and following our nose until we arrive at two possible outcomes. Either we find a solution and we are done or the system is inconsistent and we see that there is no such linear combination. The starting point of this problem is **the most important thing we will learn this far**.

If \mathbf{b} is a linear combination of $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ then **there exist scalars** $c_1, c_2, c_3 \in \mathbb{R}$ such that

$$c_1 \begin{bmatrix} 1 \\ 0 \\ -2 \end{bmatrix} + c_2 \begin{bmatrix} 2 \\ 1 \\ 1 \end{bmatrix} + c_3 \begin{bmatrix} 5 \\ 2 \\ -1 \end{bmatrix} = \begin{bmatrix} 19 \\ 7 \\ -9 \end{bmatrix}$$

This is how we always approach these problems. We find values for the c_i or realize that they cannot exist. The way in which we find the c_i is by unpacking what it means for two vectors to be equal. Using vector addition on the left hand side of the equation we get that

$$\begin{bmatrix} c_1 + 2c_2 + 5c_3 \\ c_2 + 2c_3 \\ -2c_1 + c_2 - c_3 \end{bmatrix} = \begin{bmatrix} 19 \\ 7 \\ -9 \end{bmatrix}$$

which translates to the linear system
$$\begin{cases} c_1 + 2c_2 + 5c_3 = 19 \\ c_2 + 2c_3 = 7 \\ -2c_1 + c_2 - c_3 = -9 \end{cases}$$

We solve this linear system by solving the corresponding augmented matrix

$$\left[\begin{array}{ccc|c} 1 & 2 & 5 & 19 \\ 0 & 1 & 2 & 7 \\ -2 & 1 & -1 & -9 \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 0 & 0 & 2 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & 4 \end{array} \right]$$

Note that all row reducing from now on will not be explicitly worked out. You are an expert row reducer and you can work it out yourself!

Gauss-Jordan elimination here tells us that

$$(c_1, c_2, c_3) = (2, -1, 4)$$

hence

$$\mathbf{b} = 2\mathbf{v}_1 - \mathbf{v}_2 + 4\mathbf{v}_3$$

and we are done!

This example illustrated the best case scenario, that is, we wonder if a fixed vector is a linear combination of some others, and we directly find the coefficients that give us the desired linear combination. If such a linear combination does not exist, we unravel a different conclusion.

Example 2.1.6. Let $\mathbf{v}_1 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$, $\mathbf{v}_2 = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}$, $\mathbf{v}_3 = \begin{bmatrix} 2 \\ 1 \\ -1 \end{bmatrix}$, and $\mathbf{b} = \begin{bmatrix} 1 \\ 3 \\ -1 \end{bmatrix}$. Is \mathbf{b} a linear combination of $\mathbf{v}_1, \mathbf{v}_2$, and \mathbf{v}_3 ?

Just like we did in the previous example, we set up the corresponding linear system as if there did exist such a linear combination. We then proceed by attempting to solve the linear system. In this case we get

$$\left[\begin{array}{ccc|c} 1 & 0 & 2 & 1 \\ 1 & 1 & 1 & 3 \\ 0 & 1 & -1 & 1 \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 0 & 2 & 1 \\ 0 & 1 & -1 & 2 \\ 0 & 0 & 0 & -3 \end{array} \right]$$

The equivalent matrix we have found represents an **inconsistent** linear system, therefore \mathbf{b} is **not** a linear combination of $\mathbf{v}_1, \mathbf{v}_2$, and \mathbf{v}_3

Before ending this section we make one last use of vector notation by expressing solution sets in terms of linear combinations.

Example 2.1.7. Suppose we have the following linear system
$$\begin{cases} 4x_1 - 2x_2 + x_3 - x_4 = -5 \\ x_3 + x_4 = 1 \end{cases}$$

This linear system will have infinitely many solutions because there are two free variables. We can express all such solutions in a compact way.

We first label the free variables, namely, $x_2 = t_1$ and $x_4 = t_2$. Then, using the second equation we get that

$$x_3 = 1 - t_2$$

Plugging all of this back into the first equation we see that

$$x_1 = \frac{-6 + 2t_1 + 2t_2}{4} = -\frac{3}{2} + \frac{1}{2}t_1 + \frac{1}{2}t_2$$

We then express this general solution in vector form by grouping together free variables, that is

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} \frac{3}{2} + \frac{1}{2}t_1 + \frac{1}{2}t_2 \\ t_1 \\ 1 - t_2 \\ t_2 \end{bmatrix} = \begin{bmatrix} -\frac{3}{2} \\ 0 \\ 1 \\ 0 \end{bmatrix} + t_1 \begin{bmatrix} \frac{1}{2} \\ 1 \\ 0 \\ 0 \end{bmatrix} + t_2 \begin{bmatrix} \frac{1}{2} \\ 0 \\ -1 \\ 1 \end{bmatrix}$$

The expression of a solution set in terms of a linear combination of vectors is known as the **general solution in vector form**.

2.2 Span

In this section we dig deeper into the question “Can we express one vector as a linear combination of others?” Geometrically, this is the same as asking if we can travel to a point in space, by moving along fixed directions. For example, suppose we were a little dot in \mathbb{R}^2 , located at the origin, and we wanted to find a path to the point (a, b) **but** we could only move along a line with slope 1 or slope zero, i.e. we can only move parallel to the line $y = x$ or horizontally. By translating this into the language of linear algebra, we are asking if the point (a, b) can be expressed as a linear combination of $\mathbf{u}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\mathbf{u}_2 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$.

Can we get to the point $(5, 3)$? Yes!

$$3\mathbf{u}_1 + 2\mathbf{u}_2 = \begin{bmatrix} 5 \\ 3 \end{bmatrix}$$

hence the vector $\begin{bmatrix} 5 \\ 3 \end{bmatrix}$ is a linear combination of \mathbf{u}_1 and \mathbf{u}_2 .

Example 2.2.1. We can extend the question to an arbitrary point in \mathbb{R}^2 . That is, can we express any vector $\begin{bmatrix} a \\ b \end{bmatrix}$ (for $a, b \in \mathbb{R}$) as a linear combination of $\mathbf{u}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\mathbf{u}_2 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$? If we could, then there would exist scalars $x_1, x_2 \in \mathbb{R}$ such that

$$\begin{bmatrix} a \\ b \end{bmatrix} = x_1 \begin{bmatrix} 1 \\ 1 \end{bmatrix} + x_2 \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

Finding such values of the x_i is equivalent to solving the linear system with augmented matrix

$$\left[\begin{array}{cc|c} 1 & 1 & a \\ 1 & 0 & b \end{array} \right]$$

By performing Gauss-Jordan elimination we see that

$$\left[\begin{array}{cc|c} 1 & 1 & a \\ 1 & 0 & b \end{array} \right] \sim \left[\begin{array}{cc|c} 1 & 1 & a \\ 0 & -1 & b-a \end{array} \right] \sim \left[\begin{array}{cc|c} 1 & 1 & a \\ 0 & 1 & a-b \end{array} \right] \sim \left[\begin{array}{cc|c} 1 & 0 & b \\ 0 & 1 & a-b \end{array} \right]$$

hence $x_1 = b$ and $x_2 = a - b$. In other words

$$\begin{bmatrix} a \\ b \end{bmatrix} = b \begin{bmatrix} 1 \\ 1 \end{bmatrix} + (a - b) \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

and we can write **any** vector in \mathbb{R}^2 as a linear combination of $\mathbf{u}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\mathbf{u}_2 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$. This example lends itself to the central object of this section.

Definition 2.2.2. Suppose $\mathbf{u}_1, \dots, \mathbf{u}_m \in \mathbb{R}^n$. The **span** of the vectors $\mathbf{u}_1, \dots, \mathbf{u}_m$ denoted $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$, is the set of all linear combinations of $\mathbf{u}_1, \dots, \mathbf{u}_m$. In other words, $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ consists of all vectors of the form

$$\mathbf{v} = x_1\mathbf{u}_1 + x_2\mathbf{u}_2 + \dots + x_m\mathbf{u}_m$$

for some scalars $x_1, x_2, \dots, x_n \in \mathbb{R}$.

If $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\} = \mathbb{R}^n$ we say that $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ **spans** \mathbb{R}^n .

Note that in the above example we showed that $\text{Span}\left\{\begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}\right\} = \mathbb{R}^2$ so $\left\{\begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}\right\}$ spans \mathbb{R}^2 . Now lets look at some more examples.

Example 2.2.3. Let $\mathbf{u}_1 = \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}$, $\mathbf{u}_2 = \begin{bmatrix} 2 \\ -1 \\ 1 \end{bmatrix}$ and $\mathbf{u}_3 = \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix}$. Show that $\{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3\}$ spans \mathbb{R}^3 .

Let $\mathbf{v} = \begin{bmatrix} a \\ b \\ c \end{bmatrix}$ denote an arbitrary vector in \mathbb{R}^3 . We need to show that there always exist scalars $x_1, x_2, x_3 \in \mathbb{R}$ such that

$$x_1 \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix} + x_2 \begin{bmatrix} 2 \\ -1 \\ 1 \end{bmatrix} + x_3 \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} a \\ b \\ c \end{bmatrix}$$

i.e. \mathbf{v} is a linear combination $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$.

In trying to solve the system corresponding to the vector equation above we see that

$$\begin{aligned} \left[\begin{array}{ccc|c} 1 & 2 & 4 & a \\ 2 & -1 & 0 & b \\ 0 & 1 & 1 & c \end{array} \right] &\sim \left[\begin{array}{ccc|c} 1 & 2 & 4 & a \\ 0 & -5 & -8 & b-2a \\ 0 & 1 & 1 & c \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 2 & 4 & a \\ 0 & 1 & 1 & c \\ 0 & -5 & -8 & b-2a \end{array} \right] \\ &\sim \left[\begin{array}{ccc|c} 1 & 2 & 4 & a \\ 0 & 1 & 1 & c \\ 0 & 0 & -3 & b-2a+5c \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 2 & 4 & a \\ 0 & 1 & 1 & c \\ 0 & 0 & 1 & \frac{b-2a+5c}{-3} \end{array} \right] \end{aligned}$$

From here we can use back substitution and solve for x_1, x_2 , and x_3 which means that $\begin{bmatrix} a \\ b \\ c \end{bmatrix} \in \text{Span}\{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3\}$

for **every** vector $\begin{bmatrix} a \\ b \\ c \end{bmatrix} \in \mathbb{R}^3$. This precisely means that $\text{Span}\{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3\} = \mathbb{R}^3$.

Example 2.2.4. Let $\mathbf{u}_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$, $\mathbf{u}_2 = \begin{bmatrix} 2 \\ 4 \\ -3 \end{bmatrix}$ and $\mathbf{v} = \begin{bmatrix} 2 \\ 2 \\ 5 \end{bmatrix}$. Is $\mathbf{v} \in \text{Span}\{\mathbf{u}_1, \mathbf{u}_2\}$?

We try to solve the vector equation $x_1\mathbf{u}_1 + x_2\mathbf{u}_2 = \mathbf{v}$ by looking at the augmented matrix $[\mathbf{u}_1 \quad \mathbf{u}_2 \mid \mathbf{v}]$.

$$\left[\begin{array}{cc|c} 1 & 2 & 2 \\ 1 & 4 & 2 \\ 1 & -3 & 5 \end{array} \right] \sim \left[\begin{array}{cc|c} 1 & 2 & 2 \\ 0 & 2 & 0 \\ 0 & -5 & 3 \end{array} \right] \sim \left[\begin{array}{cc|c} 1 & 2 & 2 \\ 0 & 1 & 0 \\ 0 & -5 & 3 \end{array} \right] \sim \left[\begin{array}{cc|c} 1 & 2 & 2 \\ 0 & 1 & 0 \\ 0 & 0 & 3 \end{array} \right]$$

The third line of the last equivalent matrix translates to the equation $0 = 3$ hence the linear system is inconsistent! This means there are **no** scalars $x_1, x_2 \in \mathbb{R}$ such that $x_1\mathbf{u}_1 + x_2\mathbf{u}_2 = \mathbf{v}$, hence $\mathbf{v} \notin \text{Span}\{\mathbf{u}_1, \mathbf{u}_2\}$.

To recap, we have seen three vectors that spanned \mathbb{R}^3 and two vectors that did not span \mathbb{R}^3 . It turns out that no two vectors in \mathbb{R}^3 will ever be able to span \mathbb{R}^3 , we will actually need at least 3. Will any three vectors span \mathbb{R}^3 or do we need to choose them more carefully? The next example tells us that we must choose them more carefully.

Example 2.2.5. Let $\mathbf{u}_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$, $\mathbf{u}_2 = \begin{bmatrix} 2 \\ 4 \\ -3 \end{bmatrix}$ and $\mathbf{u}_3 = \begin{bmatrix} 9 \\ 13 \\ -1 \end{bmatrix}$. Is $\mathbf{v} = \begin{bmatrix} 4 \\ -2 \\ 3 \end{bmatrix} \in \text{Span}\{\mathbf{u}_1, \mathbf{u}_2\}$?

Performing row operations reveals that

$$\left[\begin{array}{ccc|c} 1 & 2 & 9 & 4 \\ 1 & 4 & 13 & -2 \\ 1 & -3 & -1 & 3 \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 2 & 9 & 4 \\ 0 & 1 & 2 & -3 \\ 0 & 0 & 0 & -16 \end{array} \right]$$

This means that $v = \begin{bmatrix} 4 \\ -2 \\ 3 \end{bmatrix} \notin \text{Span}\{\mathbf{u}_1, \mathbf{u}_2\}$ hence any random set of three vectors will not always span \mathbb{R}^3 .

We can drill down the needed specifications a bit more in the following proposition.

Proposition 2.2.6. *Suppose $\mathbf{u}_1, \dots, \mathbf{u}_m \in \mathbb{R}^n$.*

- *If $m < n$, then $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ does not span \mathbb{R}^n .*
- *If $m \geq n$, then $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ **may or may not** span \mathbb{R}^n (we have seen that both cases are possible when $m = n$).*

This proposition prompts further investigation on how two spans are related. We will begin this investigation by proving another proposition, and before we do, we lay out some foundational ideas surrounding proofs.

2.2.1 Some modern math techniques

We begin by recalling the definition of a span of a set of vectors. Given vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$ the span of these vectors, written as $\text{span}(\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n)$ is the set of all linear combinations of the vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$.

Tying this into what we mentioned above, we can see that the span of a set of vectors is a set! What does it mean for something to be an element of this set? For this (and all other sets we encounter), being an element of a given set means the element in question satisfies the definition of what it means to be in that set. Stated in the context of span, a vector \mathbf{v} is in the span of $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$, written as

$$\mathbf{v} \in \text{span}(\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n)$$

if \mathbf{v} is a linear combination of the \mathbf{u}_i for $i = 1, 2, \dots, n$. Digging a little further, we can apply the definition and write

If $\mathbf{v} \in \text{span}(\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n)$, then there exist c_1, \dots, c_n such that

$$c_1\mathbf{u}_1 + c_2\mathbf{u}_2 + \dots + c_n\mathbf{u}_n = \mathbf{v}$$

In proving things about spans, we will constantly come back to this definition, and in general, you should remember that being an element of a set generally involves looking at the definition of what it means to be in that set. This is a very common starting point for many proofs. It should help you get your mind moving and prevent you from getting too stuck.

Some remarks on proofs

Although I don't plan to discuss proofs very much in this course, there are several basic techniques that you will be required to know. They are:

- 1) Knowing how to show that two sets are equal (in particular we will apply this to spans)

2) The implications of what an “if and only if” statement means.

1) By definition, two sets, A and B , are equal if any element of A is also an element of B , and similarly, every element of B is an element of A . If only one of these conditions holds, say every element of A is an element of B , but not every element of B is an element of A , then we say A is a *subset* of B and write $A \subset B$. Since a span of a set of vectors is a set, we will be interested in showing that two spans are equal.

The key idea is to take an arbitrary element of one set, and show it belongs to the other, then repeat the process in the other direction. Using the notation above we can write out this process in a series of steps.

- i) Pick an arbitrary element $a \in A$, and show that $a \in B$. This means that $A \subset B$.
- ii) Pick an arbitrary element $b \in B$ and show that $b \in A$. This shows that $B \subset A$.

To summarize, we have that $A = B$ if and only if $A \subset B$ and $B \subset A$. Now we explain a short bit about if and only if statements, then illustrate the above proof method with an example.

2) For if and only if statements there is not much to know. The one take away is that you have 4 useful statements that come out of it. If P and Q are two facts, say P is the fact that all cats are black and Q is the fact that all dogs are brown, then P if and only if Q (also written as $P \Leftrightarrow Q$, or P iff Q) means that all cats are black if and only if all dogs are brown. The 4 statements that we can get out of this come from breaking down the statement into parts.

If we have that P if and only if Q , then this means that

- i) If P is true then Q is true (also written as $P \implies Q$).
- ii) If Q is true then P is true (also written as $Q \implies P$).
- iii) If P is false, then Q is false.
- iv) If Q is false, then P is false.

Note that the last two statements are the negation of the first two (if this confuses you then just ignore it).

One last thing worth mentioning is what it means if we have a series of statements A, B, C and there is a theorem saying

The following are equivalent:

- i) A
- ii) B
- iii) C

What does this mean? Well the statement “the following are equivalent” means that the statements that follow can all be stated with if and only iff statements between them. The above example then reads as A if and only if B if and only if C . We can pick apart these however we please, i.e. since A if and only if B , then in particular, B implies A .

Taking an if and only if statement in the context of linear algebra, we can see how the four statements can give us different results. Recall the following theorem:

Proposition 2.2.7. *Let $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n$ be vectors in \mathbb{R}^n (we could also write $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n \in \mathbb{R}^n$). Then the following statements are equivalent:*

- i) \mathbf{b} is in $\text{span}\{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n\}$
- ii) The vector equation $x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \dots + x_n\mathbf{a}_n$ has at least one solution.

Unpacking all of this we have that:

$$\mathbf{b} \in \text{span}\{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n\} \Leftrightarrow x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \dots + x_n\mathbf{a}_n \text{ has at least one solution}$$

From this we get the four statements

i) If $\mathbf{b} \in \text{span}\{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n\}$ then $x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \dots + x_n\mathbf{a}_n$ has at least one solution.

ii) If $x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \dots + x_n\mathbf{a}_n$ has at least one solution, then $\mathbf{b} \in \text{span}\{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n\}$.

iii) If $\mathbf{b} \notin \text{span}\{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n\}$ (i.e. if the vector \mathbf{b} is NOT in the span, then $x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \dots + x_n\mathbf{a}_n$ has NO solutions.

iv) If $x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \dots + x_n\mathbf{a}_n$ has at NO solutions, then $\mathbf{b} \notin \text{span}\{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n\}$.

Now that we're a bit more familiar with if and only if statements, let's finish off with a concrete example of a proof that the spans of two different sets of vectors are equal. Remember that spans of vectors are still sets! This means that showing equality of spanning sets is done in the same way that we show equality of sets.

Example:

Prove that

$$\text{span}\left\{\begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix}\right\} = \text{span}\left\{\begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}\right\}$$

To avoid writing the above vectors as much we let

$$\mathbf{v}_1 = \begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix} \quad \mathbf{v}_2 = \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix} \quad \mathbf{v}_3 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$$

We begin by showing that

$$\text{span}\{\mathbf{v}_1, \mathbf{v}_2\} \subset \text{span}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$$

Let $\mathbf{x} \in \text{span}\{\mathbf{v}_1, \mathbf{v}_2\}$. This means that there exist scalars a_1, a_2 such that

$$\mathbf{x} = a_1\mathbf{v}_1 + a_2\mathbf{v}_2 = a_1 \begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix} + a_2 \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix}$$

This is the “unraveling the definition part”.

Now it will be super useful to remember that when we say linear combination, we can include 0 as a scalar! This will prove to be a handy trick and in this context means that

$$\mathbf{x} = a_1 \begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix} + a_2 \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix} + 0 \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} = a_1\mathbf{v}_1 + a_2\mathbf{v}_2 + 0\mathbf{v}_3$$

So we just wrote \mathbf{x} as a linear combination of $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$! Thus, $\mathbf{x} \in \text{span}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ hence we have shown that $\text{span}\{\mathbf{v}_1, \mathbf{v}_2\} \subset \text{span}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$.

Now what remains to show is the other direction, namely that $\text{span}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\} \subset \text{span}\{\mathbf{v}_1, \mathbf{v}_2\}$ and we apply the same procedure. Letting $\mathbf{x} \in \text{span}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ this means that there exist scalars b_1, b_2, b_3 such that

$$\mathbf{x} = b_1\mathbf{v}_1 + b_2\mathbf{v}_2 + b_3\mathbf{v}_3 = b_1 \begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix} + b_2 \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix} + b_3 \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$$

Now, we need to show that $\mathbf{x} \in \text{span}\{\mathbf{v}_1, \mathbf{v}_2\}$ so how can we do this? Well, showing that \mathbf{v}_3 is a linear combination of the other two will allow us to write \mathbf{x} as a linear combo ONLY in \mathbf{v}_1 and \mathbf{v}_2 . So lets try and do

that. (You may see this method and think, “how in the world was I supposed to think of that?!”, but while seeing it now may seem foreign, you will be doing this trick several times and it will seem less crazy each time).

We want to write \mathbf{v}_3 as a linear combo of \mathbf{v}_1 and \mathbf{v}_2 , so lets take a look at what that linear combination would look like. It would give us some scalars a_1, a_2 such that

$$\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} = a_1 \begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix} + a_2 \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix}$$

so all we need to do is FIND the scalars. We do this by looking at the components of each vector and deducing what the scalars MUST be in order for the above equation to hold. Let’s zoom in on the first components. For the above equality to hold, this must give us the equation

$$1 = a_1 \cdot 1 + a_2 \cdot 0 = a_1$$

hence we need to have $a_1 = 1$. Now lets look at the second components, assuming we’ve found a_1 this reduces to the equation

$$0 = 1 \cdot -2 + a_2$$

hence $a_2 = 2$. We can look at the third component and verify that indeed $a_1 = 1, a_2 = 2$ give us

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

So we have the desired linear combo. Plugging this into the original linear combination that we started with, we see that

$$\begin{aligned} \mathbf{x} &= b_1 \mathbf{v}_1 + b_2 \mathbf{v}_2 + b_3 \mathbf{v}_3 = b_1 \begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix} + b_2 \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix} + b_3 \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \\ &= b_1 \begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix} + b_2 \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix} + b_3 \left(\begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix} + 2 \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix} \right) = (b_1 + b_3) \begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix} + (b_2 + 2) \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix} \end{aligned}$$

which we can now see is a linear combination of $\mathbf{v}_1, \mathbf{v}_2$! Thus $\mathbf{x} \in \text{span}\{\mathbf{v}_1, \mathbf{v}_2\}$ which now implies that

$$\text{span} \left\{ \begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix} \right\} = \text{span} \left\{ \begin{bmatrix} 1 \\ -2 \\ 3 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \right\}$$

We now return to the main investigation concerning how two spanning sets can be related.

Proposition 2.2.8. *If $\mathbf{u} \in \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ then $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}\} = \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$.*

Proof. Recall what was mentioned about if-then statements and showing two sets are equal. Our goal will be to show that the two sets $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}\}$ and $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ are equal. Our hypothesis is that $\mathbf{u} \in \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ and we must use this somewhere along the way.

Let’s first assume that $\mathbf{u} \in \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$. This means that there exist scalars $x_1, x_2, \dots, x_m \in \mathbb{R}$ such that

$$\mathbf{u} = x_1 \mathbf{u}_1 + \dots + x_m \mathbf{u}_m$$

Since we want to ultimately show that $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}\} = \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ we pick an arbitrary $\mathbf{v} \in \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}\}$ and show that it is also in $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$. This will show that $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}\} \subseteq \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$.

$\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$. For this \mathbf{v} that we have chosen, there must exist some other scalars $y_0, y_1, \dots, y_m \in \mathbb{R}$ such that

$$\mathbf{v} = y_0\mathbf{u} + y_1\mathbf{u}_1 + \dots + y_m\mathbf{u}_m$$

Now, we use our assumption that $\mathbf{u} \in \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ and substitute $\mathbf{u} = x_1\mathbf{u}_1 + \dots + x_m\mathbf{u}_m$ into the equation for \mathbf{v} . This tells us that

$$\mathbf{v} = y_0\mathbf{u} + y_1\mathbf{u}_1 + \dots + y_m\mathbf{u}_m = y_0(x_1\mathbf{u}_1 + \dots + x_m\mathbf{u}_m) + y_1\mathbf{u}_1 + \dots + y_m\mathbf{u}_m = (y_0x_1 + y_1)\mathbf{u}_1 + \dots + (y_0x_m + y_m)\mathbf{u}_m$$

which is an element of $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$! This means that $\mathbf{v} \in \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ and

$$\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}\} \subseteq \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$$

This shows the first part. It remains to show that $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\} \subseteq \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}\}$ so we pick a vector $\mathbf{w} \in \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ and conclude that it is also in $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}\}$. The assumption on our vector \mathbf{w} implies that there exist scalars z_1, z_2, \dots, z_m such that

$$\mathbf{w} = z_1\mathbf{u}_1 + \dots + z_m\mathbf{u}_m$$

Now, observe that 0 is a scalar that we can always use when constructing linear combinations, hence we can write \mathbf{w} as a linear combination of $\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}$ by writing

$$\mathbf{w} = 0\mathbf{u} + z_1\mathbf{u}_1 + \dots + z_m\mathbf{u}_m$$

hence $\mathbf{w} \in \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}\}$ and we can conclude that

$$\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\} \subseteq \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}\}$$

This now means that

$$\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{u}\} = \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$$

which completes the proof. □

Before ending this section, we exhibit one more bit of compact (and very useful!) notation, namely that of representing a linear system via matrix notation.

Let A be a matrix with columns $\mathbf{a}_1 = \begin{bmatrix} 3 \\ 0 \\ 1 \end{bmatrix}$, $\mathbf{a}_2 = \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix}$, and $\mathbf{a}_3 = \begin{bmatrix} -4 \\ 1 \\ -5 \end{bmatrix}$. We can write the matrix A as

$$A = [\mathbf{a}_1 \quad \mathbf{a}_2 \quad \mathbf{a}_3] = \begin{bmatrix} 3 & 2 & -4 \\ 0 & 1 & 1 \\ 1 & 0 & -5 \end{bmatrix}$$

Let $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$, then we can define the following product.

Definition 2.2.9. The **product** $A\mathbf{x}$ is given by

$$A\mathbf{x} = \begin{bmatrix} 3 & 2 & -4 \\ 0 & 1 & 1 \\ 1 & 0 & -5 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = x_1 \begin{bmatrix} 3 \\ 0 \\ 1 \end{bmatrix} + x_2 \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix} + x_3 \begin{bmatrix} -4 \\ 1 \\ -5 \end{bmatrix} = x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + x_3\mathbf{a}_3$$

This allows us to succinctly write out linear systems in terms of matrices as follows.

Example 2.2.10. The linear system
$$\begin{cases} 3x_1 + 2x_2 - 4x_3 = 1 \\ x_2 + x_3 = 0 \\ x_1 - 5x_3 = 2 \end{cases}$$

has augmented matrix

$$\left[\begin{array}{ccc|c} 3 & 2 & -4 & 1 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & -5 & 2 \end{array} \right]$$

Using the product definition above we can see that for a vector $\mathbf{b} = \begin{bmatrix} 1 \\ 0 \\ 2 \end{bmatrix}$ the above system can be written

as $A\mathbf{x} = \mathbf{b}$ for $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$.

We now end the section with (what we will soon see is) a very useful theorem. We note that any time one says “the following are equivalent”, it means that “if and only if” statements should be placed between every item in the list. That is to say, if one sentence in the list is true, all others are true, and likewise, if one sentence is false then all others are false.

Theorem 2.2.11. Let $\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{b} \in \mathbb{R}^n$. The following statements are equivalent:

1. $\mathbf{b} \in \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$.
2. The vector equation $x_1\mathbf{u}_1 + \dots + x_m\mathbf{u}_m = \mathbf{b}$ has at least one solution.
3. The linear system with augmented matrix $[\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_m \mid \mathbf{b}]$ is consistent.
4. The equation $A\mathbf{x} = \mathbf{b}$ with $A = [\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_m]$ has at least one solution for every choice of $\mathbf{b} \in \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$.

2.3 Linear Independence

The topic of linear independence will be precisely what we need to understand when a set of vectors spans some euclidean space. In order to wrap our heads around it, we need one new definition.

Definition 2.3.1. A linear system is **homogeneous** if it has the form

$$x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \dots + x_n\mathbf{a}_n = \mathbf{0}$$

In other words, every linear equation is set equal to zero (the $\mathbf{0}$ denotes the zero vector, all of whose entries are 0).

The beauty of a homogeneous linear system is that it is **always** consistent since we can always find the solution $x_1 = x_2 = \dots = x_n = 0$. We call this solution the **trivial** solution, any other solutions are referred to as **non-trivial** solutions. It is this notion that allows us to define linear independence.

Definition 2.3.2. Suppose $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m \in \mathbb{R}^n$. If the trivial solution to the linear system

$$x_1\mathbf{u}_1 + x_2\mathbf{u}_2 + \dots + x_m\mathbf{u}_m = \mathbf{0}$$

is the trivial solution, then we say $\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m\}$ is a **linearly independent** set of vectors, or that the vectors are linearly independent. If the vector equation above has non-trivial solutions then $\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m\}$ is a **linearly dependent** set of vectors. Recalling that a linear system either has 0, 1, or infinitely many solutions, we can say that a set of vectors is linearly dependent if the associated homogeneous linear system involving those vectors has at least one free variable. If this confuses you then feel free to ignore it.

Example 2.3.3. Let $\mathbf{u}_1 = \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}$, $\mathbf{u}_2 = \begin{bmatrix} 0 \\ 2 \\ -2 \end{bmatrix}$, $\mathbf{u}_3 = \begin{bmatrix} 1 \\ 3 \\ 3 \end{bmatrix} \in \mathbb{R}^3$. Is $\{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3\}$ a linearly independent set of vectors?

Considering the homogeneous linear system

$$x_1\mathbf{u}_1 + x_2\mathbf{u}_2 + x_3\mathbf{u}_3 = \mathbf{0}$$

we can row reduce the corresponding augmented matrix

$$\left[\begin{array}{ccc|c} 1 & 0 & 1 & 0 \\ -1 & 2 & 3 & 0 \\ 1 & -2 & 3 & 0 \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 0 & 1 & 0 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 6 & 0 \end{array} \right]$$

Using back substitution we see that

$$x_3 = 0 \implies x_2 = 0 \implies x_1 = 0$$

hence the only solution is the trivial one. This means that

$$\begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ 3 \\ 3 \end{bmatrix}$$

are linearly independent.

Example 2.3.4. Let $\mathbf{u}_1 = \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}$, $\mathbf{u}_2 = \begin{bmatrix} 0 \\ 2 \\ -2 \end{bmatrix}$, $\mathbf{u}_3 = \begin{bmatrix} 4 \\ 2 \\ -2 \end{bmatrix} \in \mathbb{R}^3$. Are $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$ linearly independent?

Considering the augmented matrix for the homogeneous linear system associated to the three vectors above we have

$$\left[\begin{array}{ccc|c} 1 & 0 & 4 & 0 \\ -1 & 2 & 2 & 0 \\ 1 & -2 & -2 & 0 \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 0 & 4 & 0 \\ 0 & 1 & 3 & 0 \\ 0 & 0 & 0 & 0 \end{array} \right]$$

This equivalent linear system has x_3 as a free variable, and from this we obtain the **non-trivial** solution

$$x_3 = t, x_2 = -t, x_1 = -4t$$

where $t \in \mathbb{R}$. The existence of a non-trivial solution implies that $\begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}$, $\begin{bmatrix} 0 \\ 2 \\ -2 \end{bmatrix}$, and $\begin{bmatrix} 4 \\ 2 \\ -2 \end{bmatrix}$ are linearly dependent.

Now that we have a little bit of a feel for linear independence, let's dig into some important propositions that we may want to use in the future.

Proposition 2.3.5. *If $\mathbf{u}_1, \dots, \mathbf{u}_m \in \mathbb{R}^n$ then $\{\mathbf{0}, \mathbf{u}_1, \dots, \mathbf{u}_m\}$ is always linearly dependent.*

Proof. Given $\mathbf{u}_1, \dots, \mathbf{u}_m \in \mathbb{R}^n$ the equation $x_0\mathbf{0} + x_1\mathbf{u}_1 + \dots + x_m\mathbf{u}_m = \mathbf{0}$ always has the nontrivial solution $x_0 = 1, x_1 = 0, \dots, x_m = 0$. \square

We can actually say much more about when certain vectors are linearly dependent.

Proposition 2.3.6. *If $\mathbf{u}_1, \dots, \mathbf{u}_m \in \mathbb{R}^n$ and $m > n$ then $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ is linearly dependent.*

Proof. We begin by observing that the vector equation

$$x_1 \mathbf{u}_1 + \cdots + x_m \mathbf{u}_m = \mathbf{0}$$

always has at least one solution (the trivial one). This means that if we set up the usual augmented matrix and row reduce to a matrix B in echelon form, i.e.

$$[\mathbf{u}_1 \quad \mathbf{u}_2 \quad \cdots \quad \mathbf{u}_m \mid \mathbf{0}] \sim B$$

then the matrix B does not have any rows of the form $[0 \quad 0 \quad \cdots \quad 0 \mid c]$, where $c \neq 0$. Now, observing that the number of components of each vector is n (this is what it means to say that $\mathbf{u}_i \in \mathbb{R}^n$ and that $m > n$), we can conclude that there are more vectors than there are components of each vector. This means that the corresponding augmented matrix has more columns than rows, hence there **must** be at least one free variable, hence infinitely many (non-trivial) solutions, which completes the proof. \square

This Proposition will be a very important one moving forward so we will want to keep it in our toolbox. Next, we get after a bigger question. How are the ideas of span and linear independence related? The answer as we will soon see, is quite nice, especially when phrased in terms of pivots. Recall that a pivot position in a matrix is a coefficient that sits in front of what would be a leading variable, in the corresponding linear system. We now give three relationships between these two ideas, and prove the third statement in detail.

Proposition 2.3.7. *Let $\mathbf{u}_1, \dots, \mathbf{u}_m \in \mathbb{R}^n$ and suppose $A = [\mathbf{u}_1 \quad \cdots \quad \mathbf{u}_m] \sim B$ where B is a matrix in echelon form.*

1. $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\} = \mathbb{R}^n$ exactly when B has a pivot in every **row**.
2. $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ is linearly independent exactly when B has a pivot in every **column**.

This proposition is a personal favorite of many. It essentially gives an algorithm for determining when a given set of vectors span \mathbb{R}^n and/or are linearly independent. The question of spanning is a question about pivots of rows and the question of independence is a question about pivots of columns. All one needs to do before checking rows and/or columns, is put the given vectors as the columns of a matrix and row reduce to echelon form.

The last relationship is the following theorem.

Theorem 2.3.8. *Let $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ be a set of vector in \mathbb{R}^n . This set is linearly dependent if and only if one of the vectors in the set is in the span of the others.*

Proof. As with any “if and only if” proof, we must show both directions of the statement, We begin by assuming that the given vectors are linearly dependent, then deduce that one of the vectors is in the span of the others. This is the forward direction of the proof and is indicated with “ \rightarrow ”. After proving this direction, we tackle the reverse direction, denoted by “ \leftarrow ”, where we assume that one of the vectors is in the span of the others, and conclude linear dependence.

\rightarrow

Suppose $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ is linearly dependent. Then the vector equation $x_1 \mathbf{u}_1 + \cdots + x_m \mathbf{u}_m = \mathbf{0}$ has a non-trivial solution, which we call (x_1, \dots, x_m) . Note that this solution being non-trivial means that **at least one** of the x_i is non-zero (so we can divide by it!). Without loss of generality, let's assume that $x_1 \neq 0$. Using the vector equation above, we can then solve for \mathbf{u}_1

$$x_1 \mathbf{u}_1 + \cdots + x_m \mathbf{u}_m = \mathbf{0} \implies x_1 \mathbf{u}_1 = -(x_2 \mathbf{u}_2 + \cdots + x_m \mathbf{u}_m) \implies \mathbf{u}_1 = \frac{-(x_2 \mathbf{u}_2 + \cdots + x_m \mathbf{u}_m)}{x_1}$$

hence $\mathbf{u}_1 \in \text{Span}\{\mathbf{u}_2, \dots, \mathbf{u}_m\}$.

←

Now assume that one of the vectors (say \mathbf{u}_1) is a linear combination of the others. Then there exist scalars c_2, \dots, c_m such that

$$\mathbf{u}_1 = c_2\mathbf{u}_2 + \dots + c_m\mathbf{u}_m \quad \text{hence} \quad \mathbf{u}_1 - c_2\mathbf{u}_2 - \dots - c_m\mathbf{u}_m = \mathbf{0}$$

so we have a non-trivial solution to the equation $x_1\mathbf{u}_1 + \dots + x_m\mathbf{u}_m = \mathbf{0}$, which is exactly what it means for the vectors $\mathbf{u}_1, \dots, \mathbf{u}_m$ to be linearly dependent. \square

Example 2.3.9. One can show that the set $\left\{ \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \begin{bmatrix} 10 \\ 9 \end{bmatrix}, \begin{bmatrix} -4 \\ 17 \end{bmatrix} \right\}$ is linearly dependent (this is a good exercise), hence the above theorem implies that one of them is a linear combination of the others. In fact, we have

$$\begin{bmatrix} -4 \\ 17 \end{bmatrix} = \frac{-206}{19} \begin{bmatrix} 1 \\ -1 \end{bmatrix} + \frac{13}{19} \begin{bmatrix} 10 \\ 9 \end{bmatrix}$$

Warning: This does not mean that **every** vector is a linear combination of the others. An easy example of this is $\left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}$. Can you see which vectors are not in the span of the others?

We end the chapter with what is arguably the most important theorem of linear algebra which we refer to it as the big theorem. It is given as a list of equivalent statements and we will add to the list throughout the course. The key thing to note about the big theorem is that its statements are only true if we have \mathbf{n} vectors in \mathbb{R}^n . In most of the statements of propositions we have m vectors in \mathbb{R}^n and we do not assume that m and n are the same. This is something you should always be aware of if you try to use the big theorem to solve a problem.

Theorem 2.3.10. Let $S = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ be a set of vectors in \mathbb{R}^n and let $A = [\mathbf{u}_1 \ \dots \ \mathbf{u}_n]$. The following statements are equivalent:

1. S spans \mathbb{R}^n .
2. S is linearly independent.
3. The system $A\mathbf{x} = \mathbf{b}$ has a solution for every $\mathbf{b} \in \mathbb{R}^n$.

We end the chapter with an example question that would be impossible to solve without the big theorem.

Example 2.3.11. Let $A = [\mathbf{u}_1 \ \mathbf{u}_2 \ \mathbf{u}_3]$ for $\mathbf{u}_1 = \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}$, $\mathbf{u}_2 = \begin{bmatrix} 0 \\ 2 \\ -2 \end{bmatrix}$, $\mathbf{u}_3 = \begin{bmatrix} 1 \\ 3 \\ 3 \end{bmatrix}$, and show that $A\mathbf{x} = \mathbf{b}$ has a solution for every $\mathbf{b} \in \mathbb{R}^3$.

By the big theorem, $A\mathbf{x} = \mathbf{b}$ has a solution for every $\mathbf{b} \in \mathbb{R}^3$ if and only if the three vectors are linearly independent. This was shown in Example 2.3.3, hence the question is true by the big theorem.

Chapter 3

Linear Transformations

Up to this point, we have done a tremendous amount of algebra with vectors and matrices, but we have not examined the geometry underlying linear systems. As we will soon see, the notion of a linear transformation allows us to translate our algebraic notions into geometric ones. Often times in practice, we aim to answer hard geometric questions and the methods we use involve translating the geometry into an algebraic question involving matrices, then using the matrices to answer the question, and translating the answer back to the underlying geometric picture.

3.1 The Basics of Linear Maps


We begin by outlining the basic vocabulary of linear transformations, otherwise known as linear maps. A priori, a linear map is just a function that takes vectors as input and outputs vectors (of possibly different size than the input). The notation

$$T : \mathbb{R}^m \rightarrow \mathbb{R}^n$$

reads as “ T is a function from \mathbb{R}^m to \mathbb{R}^n ”.

- The set \mathbb{R}^m is the **domain** of T (and T must be defined for every element of \mathbb{R}^m).
- The set \mathbb{R}^n is the **codomain** of T . It is the set where all the output vectors live.
- The subset of \mathbb{R}^n consisting of all output vectors, that is, all vectors of the form $\mathbf{w} = T(\mathbf{x})$ for some $\mathbf{x} \in \mathbb{R}^m$ is known as the **Range of T** , denoted $\text{Range}(T)$. It is also often called the image of T .

The following picture can serve as a visual summary of these definitions



Before defining what a linear map is, let's look at an example of a vector valued function.

Example 3.1.1. Let $T: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ be defined by

$$T\left(\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}\right) = \begin{bmatrix} x_1x_2 \\ x_2 - x_3 \end{bmatrix}$$

The domain of this map is \mathbb{R}^3 and the codomain is \mathbb{R}^2 . The vector $\begin{bmatrix} 2 \\ 1 \end{bmatrix}$ is in $\text{Range}(T)$ because $T\left(\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}\right) = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$

Definition 3.1.2. A function $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a **linear transformation** or **linear map** if, for every $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$ and every scalar $r \in \mathbb{R}$, we have:

- $T(\mathbf{u} + \mathbf{v}) = T(\mathbf{u}) + T(\mathbf{v})$.
- $T(r\mathbf{u}) = rT(\mathbf{u})$

Some people like to combine the two conditions of linearity by saying that T is a linear transformation if

$$T(r\mathbf{u} + s\mathbf{v}) = rT(\mathbf{u}) + sT(\mathbf{v})$$

for all vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$ and all scalars $r, s \in \mathbb{R}$.

Example 3.1.3. Let's show that the map $T: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ defined by

$$T\left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}\right) = \begin{bmatrix} -x_2 \\ x_1 + x_2 \\ 4x_1 \end{bmatrix}$$

is a linear map.

Let $\mathbf{u} = \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}$ and $\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}$ be arbitrary vectors in the domain (\mathbb{R}^2). Then $\mathbf{u} + \mathbf{v} = \begin{bmatrix} u_1 + v_1 \\ u_2 + v_2 \end{bmatrix}$ hence

$$T(\mathbf{u} + \mathbf{v}) = \begin{bmatrix} -(u_1 + v_1) \\ (u_1 + v_1) + (u_2 + v_2) \\ 4(u_1 + v_1) \end{bmatrix} = \begin{bmatrix} -u_2 \\ u_1 + u_2 \\ 4u_1 \end{bmatrix} + \begin{bmatrix} -v_2 \\ v_1 + v_2 \\ 4v_1 \end{bmatrix} = T(\mathbf{u}) + T(\mathbf{v})$$

Moreover, if $r \in \mathbb{R}$ then $r\mathbf{u} = \begin{bmatrix} ru_1 \\ ru_2 \end{bmatrix}$ and

$$T(r\mathbf{u}) = \begin{bmatrix} -ru_2 \\ ru_1 + ru_2 \\ 4ru_1 \end{bmatrix} = r \begin{bmatrix} -u_2 \\ u_1 + u_2 \\ 4u_1 \end{bmatrix} = rT(\mathbf{u})$$

hence T is indeed a linear transformation.

Example 3.1.4. Let $T: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ be the map defined earlier by

$$T\left(\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}\right) = \begin{bmatrix} x_1x_2 \\ x_2 - x_3 \end{bmatrix}$$

This is **not** a linear map since, for example, if $r = 2$ then

$$T(2\mathbf{x}) = \begin{bmatrix} 4x_1x_2 \\ 2(x_2 - x_3) \end{bmatrix} \neq 2T(\mathbf{x}) = \begin{bmatrix} 2x_1x_2 \\ 2(x_2 - x_3) \end{bmatrix}$$

One way to see why this is not a linear map is that the first coordinate of an arbitrary output vector is a quadratic function in the input variables. In general, linear maps have coordinate functions that are linear.

One of the most amazing things about linear maps is that they are intimately tied to matrices.

Definition 3.1.5. A matrix with n rows and m columns has **dimensions** $n \times m$ and is referred to as an $n \times m$ matrix. An $n \times n$ matrix is often called a **square** matrix.

Now, by recalling Definition 2.2.9 (the product $A\mathbf{x}$) we can see the connection with matrices and linear maps. If A is an $n \times m$ matrix and $\mathbf{x} \in \mathbb{R}^m$ then the product $A\mathbf{x}$ is always a vector in \mathbb{R}^n (you should verify this for yourself). In other words, an $n \times m$ matrix, when multiplied by a vector in $\mathbf{x} \in \mathbb{R}^m$, takes \mathbf{x} to a vector $A\mathbf{x}$, in \mathbb{R}^n .

Theorem 3.1.6. Let A be an $n \times m$ matrix and define $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ via

$$T(\mathbf{x}) = A\mathbf{x}$$

then T is a linear transformation.

The above theorem is powerful and can be used to easily show that a given map is linear, without verifying the two properties of the original definition. That is, to show that a function $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a linear map, it suffices to find a matrix A such that $T(\mathbf{x}) = A\mathbf{x}$.

Example 3.1.7. Consider the linear map from Example 3.1.3, which we now know is indeed linear. Using Definition 2.2.9 we have

$$T(\mathbf{x}) = \begin{bmatrix} 0x_1 - x_2 \\ x_1 + x_2 \\ 4x_1 + 0x_2 \end{bmatrix} = x_1 \begin{bmatrix} 0 \\ 1 \\ 4 \end{bmatrix} + x_2 \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 & -1 \\ 1 & 1 \\ 4 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

so $T(\mathbf{x}) = A\mathbf{x}$ for $A = \begin{bmatrix} 0 & -1 \\ 1 & 1 \\ 4 & 0 \end{bmatrix}$ and by the above theorem, T is a linear map.

Continuing with this example, let $\mathbf{w} = \begin{bmatrix} 10 \\ 5 \\ 2 \end{bmatrix}$. Is $\mathbf{w} \in \text{Range}(T)$? That is, does there exist a vector $\mathbf{x} \in \mathbb{R}^2$ such that $T(\mathbf{x}) = \mathbf{w}$. Since $T(\mathbf{x}) = A\mathbf{x}$, the existence of such a vector \mathbf{x} would imply that $A\mathbf{x} = \mathbf{w}$ so to find the vector \mathbf{x} we need to solve the system $A\mathbf{x} = \mathbf{w}$ which has augmented matrix

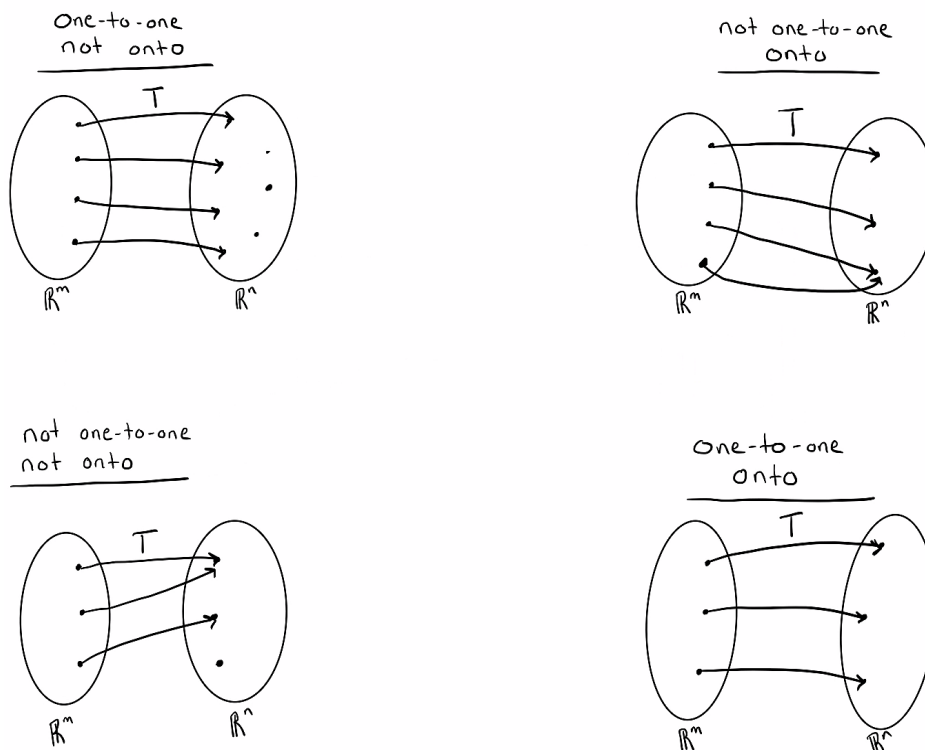
$$\left[\begin{array}{cc|c} 0 & -1 & 10 \\ 1 & 1 & 5 \\ 4 & 0 & 2 \end{array} \right] \sim \left[\begin{array}{cc|c} 1 & 1 & 5 \\ 0 & 1 & -10 \\ 0 & 0 & -58 \end{array} \right]$$

hence $\mathbf{w} \notin \text{Range}(T)$ because there does not exist a vector \mathbf{x} with $T(\mathbf{x}) = \mathbf{w}$. This is an example of a linear map that is **not onto**, which leads us to our next set of definitions.

Definition 3.1.8. Let $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear transformation.

1. T is **one-to-one** if for each $w \in \mathbb{R}^n$, there is at most one vector $\mathbf{x} \in \mathbb{R}^m$ such that $T(\mathbf{x}) = \mathbf{w}$. In other words, every domain vector \mathbf{x} goes to exactly one vector in the codomain. It is not possible for one-to-one maps to send two different vectors to the same one. This would be “two-to-one”.
2. T is **onto** if for every $\mathbf{w} \in \mathbb{R}^n$, there is at least one vector $\mathbf{x} \in \mathbb{R}^m$ such that $T(\mathbf{x}) = \mathbf{w}$. In other words, T is onto if every vector in the codomain is mapped to by some vector in the domain.

Less formally, T is one-to-one if nothing in the codomain gets “hit” more than once, and T is onto if everything in the codomain gets “hit”.



All possibilities involving these definitions are most easily understood through these helpful pictures. Now let's get familiar with these concepts through examples.

Example 3.1.9. Let $A = \begin{bmatrix} 1 & -2 \\ -2 & 4 \end{bmatrix}$ with $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by $T(\mathbf{x}) = A\mathbf{x}$. We have

$$T\left(\begin{bmatrix} 1 \\ -2 \end{bmatrix}\right) = \begin{bmatrix} 1 & -2 \\ -2 & 4 \end{bmatrix} \begin{bmatrix} 1 \\ -2 \end{bmatrix} = \begin{bmatrix} 5 \\ -10 \end{bmatrix} \quad \text{and} \quad T\left(\begin{bmatrix} -3 \\ -4 \end{bmatrix}\right) = \begin{bmatrix} 1 & -2 \\ -2 & 4 \end{bmatrix} \begin{bmatrix} -3 \\ -4 \end{bmatrix} = \begin{bmatrix} 5 \\ -10 \end{bmatrix}$$

which means that T is not one-to-one. Moreover, (exercise) T is not onto since there is no $\mathbf{x} \in \mathbb{R}^2$ with $T(\mathbf{x}) = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$.

Example 3.1.10. Let $A = \begin{bmatrix} 2 & 0 \\ 1 & -1 \\ 0 & 0 \end{bmatrix}$ with $T: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ given by $T(\mathbf{x}) = A\mathbf{x}$. Is T onto?

If it were, then for any vector $\mathbf{w} = \begin{bmatrix} a \\ b \\ c \end{bmatrix} \in \mathbb{R}^3$ we could always find a vector $\mathbf{x} \in \mathbb{R}^2$ such that $T(\mathbf{x}) = A\mathbf{x} = \mathbf{w}$. Solving the associated linear system in the usual way we get that

$$\left[\begin{array}{cc|c} 2 & 0 & a \\ 1 & -1 & b \\ 0 & 0 & c \end{array} \right] \sim \left[\begin{array}{cc|c} 1 & -1 & b \\ 0 & 2 & a-2b \\ 0 & 0 & c \end{array} \right]$$

which corresponds to a linear system whose third equation is $0 = c$. Now, if \mathbf{w} was a vector with non-zero

third coordinate, then $\mathbf{w} \notin \text{Range}(T)$ by what we have stated above. For example, $\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \notin \text{Range}(T)$ whereas

$$\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \in \text{Range}(T).$$

Example 3.1.11. Let $A = \begin{bmatrix} 2 & 0 \\ 0 & 4 \end{bmatrix}$ with $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by $T(\mathbf{x}) = A\mathbf{x}$. In coordinates, we have $T\left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}\right) = \begin{bmatrix} 2x_1 \\ 4x_2 \end{bmatrix}$. If $\mathbf{w} = \begin{bmatrix} a \\ b \end{bmatrix}$ then $T(\mathbf{x}) = \mathbf{w}$ has **exactly** one solution, namely $\mathbf{x} = \begin{bmatrix} a/2 \\ b/4 \end{bmatrix}$. There are no other vectors that get mapped to $\mathbf{w} = \begin{bmatrix} a \\ b \end{bmatrix}$. This means that every vector gets “hit” **and** there is exactly one \mathbf{x} such that $T(\mathbf{x}) = \mathbf{w}$ for any \mathbf{w} , hence T is both one-to-one and onto.

There is an alternative definition of one-to-one that some may find useful.

Definition 3.1.12. T is one-to-one if

$$T(\mathbf{u}) = T(\mathbf{v}) \text{ implies } \mathbf{u} = \mathbf{v}$$

Continuing along with the notion of a one-to-one map, we have one essential property of a linear map, that closely ties into being one-to-one.

If $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a linear transformation, then $T(\mathbf{0}) = \mathbf{0}$.

Theorem 3.1.13. Let $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear map. T is one-to-one if and only if the equation $T(\mathbf{x}) = \mathbf{0}$ has only the trivial solution $\mathbf{x} = \mathbf{0}$. That is, if $T(\mathbf{x}) = \mathbf{0}$, we must have $\mathbf{x} = \mathbf{0}$.

Proof. If $T(\mathbf{x}) = \mathbf{0}$ has only the trivial solution, then by the alternative definition of one-to-one, we can conclude that if $T(\mathbf{u}) = T(\mathbf{v})$ then $T(\mathbf{u}) - T(\mathbf{v}) = \mathbf{0}$. Moreover, since T is linear, we have $T(\mathbf{u}) - T(\mathbf{v}) = T(\mathbf{u} - \mathbf{v})$ hence $T(\mathbf{u} - \mathbf{v}) = \mathbf{0}$. The trivial solution here implies that $\mathbf{u} - \mathbf{v} = \mathbf{0}$, thus we must have $\mathbf{u} = \mathbf{v}$, implying that T is indeed one-to-one. \square

In practice, when checking if a linear map is one-to-one, this theorem is the easiest method to use. In general, the following theorem outlines some other useful methods of checking when linear maps are one-to-one or onto.

Proposition 3.1.14. Let A be an $n \times m$ matrix and define $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ via $T(\mathbf{x}) = A\mathbf{x}$. Then

1. T is one-to-one if and only if the columns of A are linearly independent.
2. If $m > n$, then T is never one-to-one.
3. T is onto if and only if the columns of A span \mathbb{R}^n (the codomain).
4. If $m < n$, then T is never onto.

In practice, you should always try to use statements 2 and 4 from the proposition, they are super useful!

Next, let's illustrate the last proposition with an example.

Example 3.1.15. Let $A = \begin{bmatrix} 2 & 0 \\ 0 & 1 \\ 3 & -3 \end{bmatrix}$ with $T: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ given by $T(\mathbf{x}) = A\mathbf{x}$. Since $m < n$ we know immediately that T is not onto, but it could be one-to-one. To find out if it is, we look at the equation

$T(\mathbf{x}) = A\mathbf{x} = \mathbf{0}$. This is a linear system which reduces via

$$\left[\begin{array}{cc|c} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 3 & -3 & 0 \end{array} \right] \sim \left[\begin{array}{cc|c} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{array} \right]$$

hence only the trivial solution exists and we must have $\mathbf{x} = \mathbf{0} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$. We can then conclude, by the theorem we just proved, that T is one-to-one.

Now we have seen quite a few linear maps in action. Every one that we have seen was given by some matrix so it is natural to ask if **all** linear maps are given by matrices. The emphatic answer is yes!

Theorem 3.1.16. *If $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a linear transformation, then there exists an $n \times m$ matrix A such that $T(\mathbf{x}) = A\mathbf{x}$.*

This means that we can **always** find the matrix associated to a linear map (and we should!). Working with linear maps is always easier when working with their associated matrices and because of this, we move interchangeably between linear maps and matrices from here on out. When you think of a matrix you should always be thinking about what it does as a linear transformation.

The beauty of this theorem extends further. In fact, given any linear map, we can **always** find its associated matrix fairly easily.

Example 3.1.17. Let $T: \mathbb{R}^3 \rightarrow \mathbb{R}^5$ be given by

$$T\left(\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}\right) = \begin{bmatrix} x_3 - x_1 \\ x_2 + x_3 \\ 4x_1 + 3x_2 \\ x_1 - 5x_3 \\ 9x_2 \end{bmatrix}$$

Lets find the matrix for T .

Let

$$\mathbf{e}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \mathbf{e}_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

We call these the standard basis vectors of the domain. If the domain is \mathbb{R}^m then there are m of these vectors. The general rule is that the formula for A is

$$A = [T(\mathbf{e}_1) \quad T(\mathbf{e}_2) \quad T(\mathbf{e}_3)]$$

so by using the coordinate definition of T we have that $T(\mathbf{e}_1) = \begin{bmatrix} -1 \\ 0 \\ 4 \\ 1 \\ 0 \end{bmatrix}$, $T(\mathbf{e}_2) = \begin{bmatrix} 0 \\ 1 \\ 3 \\ 0 \\ 9 \end{bmatrix}$, $T(\mathbf{e}_3) = \begin{bmatrix} 1 \\ 1 \\ 0 \\ -5 \\ 0 \end{bmatrix}$ hence

$$A = \begin{bmatrix} -1 & 0 & 1 \\ 0 & 1 & 1 \\ 4 & 3 & 0 \\ 1 & 0 & -5 \\ 0 & 9 & 0 \end{bmatrix}$$

This formula is our new best friend is is **extremely useful**. Let's summarize how this works in general.

Proposition 3.1.18. *Let $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear transformation. Then $T(\mathbf{x}) = A\mathbf{x}$ with A an $n \times m$ matrix given by*

$$A = [T(\mathbf{e}_1) \quad T(\mathbf{e}_2) \quad \cdots \quad T(\mathbf{e}_m)]$$

Last but not least, we use results from this section to add to the big theorem.

Theorem 3.1.19. *Let $S = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ be a set of vectors in \mathbb{R}^n and let $A = [\mathbf{u}_1 \quad \cdots \quad \mathbf{u}_n]$ with associated linear transformation given by $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$. The following statements are equivalent:*

1. S spans \mathbb{R}^n .
2. S is linearly independent.
3. The system $A\mathbf{x} = \mathbf{b}$ has a solution for every $\mathbf{b} \in \mathbb{R}^n$.
4. T is onto.
5. T is one-to-one.

An interesting consequence of this, in stark contrast to Proposition 3.1.14, is that a linear map from \mathbb{R}^n to itself is either one-to-one **and** onto, or neither.

3.2 Matrix Algebra

In continuing with our geometric theme, the tools of matrix algebra provide the algebraic notions that we will use to answer geometric questions concerning multiple linear transformations.

The first notion we need is matrix addition. This is done component-wise, in a way that is similar to vectors.

Example 3.2.1. Let $A = \begin{bmatrix} 4 & 0 & -1 \\ 2 & 2 & 5 \end{bmatrix}$ and $B = \begin{bmatrix} 9 & 10 & 6 \\ -1 & 0 & -1 \end{bmatrix}$. We define $A + B$ to be the 2×3 matrix

$$A + B = \begin{bmatrix} 13 & 10 & 5 \\ 1 & 2 & 4 \end{bmatrix}$$

Note that we obtained this matrix just by adding matching coordinates of each matrix. For a given scalar $r \in \mathbb{R}$ we define rA to be

$$rA = \begin{bmatrix} 4r & 0 & -r \\ 2r & 2r & 5r \end{bmatrix}$$

In general, there are just several things to note about matrix addition.

1. One can only add matrices of the same size. That is, if C is a 2×3 matrix and D is a 3×4 matrix then $C + D$ is undefined.
2. We denote the zero matrix with n rows and m columns by 0_{nm} , or simply write 0 when the context is clear. The zero matrix satisfies the property that $0 + A = A$ for any matrix A where the addition is defined.
3. Matrix addition is commutative, that is, $A + B = B + A$.

We now move onto the slightly more complicated (but also more important) notion of matrix multiplication. This can be thought of as a generalization of multiplying a matrix by a vector.

Definition 3.2.2. If A is an $n \times k$ matrix and B is a $k \times m$ matrix, written column-wise as $B = [\mathbf{b}_1 \ \mathbf{b}_2 \ \cdots \ \mathbf{b}_m]$ then the **product** AB is the $n \times m$ matrix given by

$$AB = [A\mathbf{b}_1 \ A\mathbf{b}_2 \ \cdots \ A\mathbf{b}_m]$$

where each column $A\mathbf{b}_i$ is computed using Definition 2.2.9.

Example 3.2.3. Let $A = \begin{bmatrix} 4 & 0 & -1 \\ 2 & 2 & 5 \end{bmatrix}$ and $B = \begin{bmatrix} -2 & 1 & 2 & 0 \\ 6 & 0 & -3 & -1 \\ 7 & -1 & 4 & 1 \end{bmatrix}$. Then

$$\begin{aligned} AB &= \begin{bmatrix} 4 & 0 & -1 \\ 2 & 2 & 5 \end{bmatrix} \begin{bmatrix} -2 & 1 & 2 & 0 \\ 6 & 0 & -3 & -1 \\ 7 & -1 & 4 & 1 \end{bmatrix} \\ &= A \begin{bmatrix} -2 \\ 6 \\ 7 \end{bmatrix} + A \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix} + A \begin{bmatrix} 2 \\ -3 \\ 4 \end{bmatrix} + A \begin{bmatrix} 0 \\ -1 \\ -1 \end{bmatrix} = \begin{bmatrix} -15 & 5 & 4 & -1 \\ 43 & -3 & 18 & 3 \end{bmatrix} \end{aligned}$$

This can be taken to be the original definition of matrix multiplication, but in practice, there is a much easier way of computing it.

Given an $n \times k$ matrix A and a $k \times m$ matrix B , the product AB is the $n \times m$ matrix, whose ij -entry (the entry in row i column j) is the dot product of the i^{th} row of A with the j^{th} column of B . It would be a great exercise to run back through the example above using this method. In doing so, you should also see why the product is not defined when the number of columns of A does not equal the number of rows of B .

Warning: In general, the order in which one multiplies matrices matters. With the example above, even though AB is defined, BA is not. Always exercise care with the order in which you multiply matrices.

It is now a good time to introduce some special types of matrices that we will encounter more frequently as well as some useful ideas that come from our new perspective of matrices. We begin with two definitions, then lay out some special classes of matrices.

Definition 3.2.4. The **transpose** of an $m \times n$ matrix A , denoted A^\top , is the $m \times n$ matrix obtained by interchanging the rows and columns of A . For example, if $A = \begin{bmatrix} 3 & 0 & 1 \\ 4 & 1 & -2 \end{bmatrix}$ then $A^\top = \begin{bmatrix} 3 & 4 \\ 0 & 1 \\ 1 & -1 \end{bmatrix}$. The main properties of transposing matrices is that the transpose of a product is the product of transposes (with order swapped), that is

$$(AB)^\top = B^\top A^\top$$

Definition 3.2.5. Given a square matrix A , we define the k^{th} power of A to be the matrix A^k . That is, the matrix obtained by multiplying A by itself k times. For example, $A^2 = AA$ and $A^3 = A(A^2) = AAA$.

- **The $n \times n$ identity Matrix, I_n :** Considering (again) the standard basis vectors for \mathbb{R}^n ,

$$\mathbf{e}_1 = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix}, \dots, \mathbf{e}_n = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}, \text{ the } n \times n \text{ identity matrix is given by}$$

$$I_n = [\mathbf{e}_1 \ \mathbf{e}_2 \ \cdots \ \mathbf{e}_n]$$

I_n is the unique matrix for which $AI_n = I_nA = A$ for any $n \times n$ matrix A .

To see some small examples, we have

$$I_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad I_3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

- **Diagonal Matrices:** If the only non-zero entries of a square matrix A lie on the main diagonal, then we call A a **diagonal matrix**. For example,

$$A = \begin{bmatrix} 2 & 0 & 0 & 0 \\ 0 & 4 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix}$$

and we can check that

$$A^2 = \begin{bmatrix} 4 & 0 & 0 & 0 \\ 0 & 16 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 36 \end{bmatrix} \quad \text{and} \quad A^3 = \begin{bmatrix} 8 & 0 & 0 & 0 \\ 0 & 64 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 216 \end{bmatrix}$$

In general, if

$$A = \begin{bmatrix} a_{11} & 0 & \dots & 0 \\ 0 & a_{22} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & a_{nn} \end{bmatrix} \quad \text{then} \quad A^k = \begin{bmatrix} a_{11}^k & 0 & \dots & 0 \\ 0 & a_{22}^k & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & a_{nn}^k \end{bmatrix}$$

- **Triangular Matrix:** If A is a square matrix with zeroes in each entry below the main diagonal, then A is an **upper triangular matrix**. We can similarly define a lower triangular matrix to have zeroes below the main diagonal. If A is either upper or lower triangular, then we say that A is a **triangular matrix**. For example, given the two matrices

$$A = \begin{bmatrix} 1 & 2 & 2 \\ 0 & 4 & 5 \\ 0 & 0 & 6 \end{bmatrix} \quad B = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 5 & 0 \\ 4 & 4 & 4 \end{bmatrix}$$

we have that A is upper triangular and B is lower triangular.

Expanding on what we said for diagonal matrices, we have the (great) fact.

Proposition 3.2.6. *If A is a triangular matrix, then A^k is triangular.*

Note that what was said about powers of diagonal matrices follows from this Proposition because all diagonal matrices are triangular.

We now outline some things that we need to be very careful about, when it comes to matrices and products of them. These are things that you will want to always keep in mind when computing matrix products.

1. In general, if AB is defined, the product BA is not defined. This is always the case if A or B are not square matrices and can be seen in the previous example.
2. The commutative property does not hold for matrix multiplication. For example,

$$\begin{bmatrix} 1 & 2 \\ 3 & -4 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ -3 & 4 \end{bmatrix} = \begin{bmatrix} -1 & 8 \\ 15 & -10 \end{bmatrix} \neq \begin{bmatrix} 1 & 2 \\ -3 & 4 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 3 & -4 \end{bmatrix}$$

3. Unlike with numbers, it is possible to multiply two non-zero matrices together and obtain the zero matrix. For example,

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

The takeaway from this is that if $AB = 0$ we cannot conclude that either $A = 0$ or $B = 0$.

4. Its possible that $AC = BC$ but $A \neq B$ and $C \neq 0$. For example,

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

The takeaway from this is that if $AC = BC$ and $C \neq 0$ then we cannot conclude that $A = B$.

We now end the section with what is arguably the most important aspect of matrix multiplication.

Proposition 3.2.7. *Let $T_1: \mathbb{R}^m \rightarrow \mathbb{R}^k$ be given by $T_1(\mathbf{x}) = A_1\mathbf{x}$ and let $T_2: \mathbb{R}^k \rightarrow \mathbb{R}^n$ be given by $T_2(\mathbf{x}) = A_2\mathbf{x}$. The matrix associated to the composition $T_2 \circ T_1(\mathbf{x}) = T_2(T_1(\mathbf{x}))$ is A_2A_1 , that is, matrix multiplication corresponds to composition of associated linear maps.*

Proof. We can quickly verify that

$$T_2 \circ T_1(\mathbf{x}) = T_2(T_1(\mathbf{x})) = T_2(A_1\mathbf{x}) = A_2A_1\mathbf{x}$$

□

Example 3.2.8. Let $T_1, T_2: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be given by $T_1(\mathbf{x}) = A_1\mathbf{x}$ and $T_2(\mathbf{x}) = A_2\mathbf{x}$ with

$$A_1 = \begin{bmatrix} 1 & 2 \\ -3 & 4 \end{bmatrix} \quad \text{and} \quad A_2 = \begin{bmatrix} 2 & 1 \\ 3 & -2 \end{bmatrix}$$

If $\mathbf{x} = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$ then we can compute $T_1(T_2(\mathbf{x}))$. First observe that

$$T_2(\mathbf{x}) = \begin{bmatrix} 2 & 1 \\ 3 & -2 \end{bmatrix} \begin{bmatrix} 1 \\ -1 \end{bmatrix} = \begin{bmatrix} 1 \\ 7 \end{bmatrix}$$

Now we have

$$T_1(T_2(\mathbf{x})) = \begin{bmatrix} 1 & 2 \\ -3 & 4 \end{bmatrix} \begin{bmatrix} 1 \\ 7 \end{bmatrix} = \begin{bmatrix} 15 \\ 25 \end{bmatrix}$$

We can also verify that

$$A_1A_2 = \begin{bmatrix} 8 & -7 \\ 6 & -19 \end{bmatrix}$$

which means that $T_1(T_2(\mathbf{x})) = \begin{bmatrix} 8 & -7 \\ 6 & -19 \end{bmatrix} \begin{bmatrix} 1 \\ -1 \end{bmatrix}$. A direct computation indeed yields the desired result.

As a last note, we emphasize that the order of matrix multiplication is an essential component of computing a composition of linear maps correctly. In practice, always make sure that the order in which you multiply is correct.

3.3 Inverses

We now come to the last topic of this chapter, the all important idea of an inverse. We will see that the notion of an inverse will correspond to the same notion of an inverse function. They will also make solving certain linear systems much easier.

Definition 3.3.1. If A is an $n \times n$ matrix and there exists another $n \times n$ matrix, A^{-1} (pronounced A inverse), satisfying

$$A^{-1}A = AA^{-1} = I_n$$

then A is **invertible** and we say A^{-1} is the inverse of A .

Example 3.3.2. Let $A = \begin{bmatrix} 1 & -1 \\ 3 & 2 \end{bmatrix}$. We can see that A is invertible and $A^{-1} = \begin{bmatrix} 2/5 & 1/5 \\ -3/5 & 1/5 \end{bmatrix}$ since an easy computation shows that $A^{-1}A = AA^{-1} = I_2$.

Going along the lines of the example, we actually have a nice closed formula for the inverse of a 2×2 matrix (larger matrices do not have such nice formulas). Given $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, if $ad - bc \neq 0$ then A is invertible and

$$A^{-1} = \frac{1}{ad - bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$$

Our main task in this section will be to compute inverses for $n \times n$ matrices where $n > 2$. The process is as follows:

Suppose we are given the $n \times n$ matrix $A = [\mathbf{a}_1 \ \mathbf{a}_2 \ \cdots \ \mathbf{a}_n]$. To find A^{-1} (if it exists) we

1. Augment A with the $n \times n$ identity matrix $I_n = [\mathbf{e}_1 \ \mathbf{e}_2 \ \cdots \ \mathbf{e}_n]$ to get

$$[\mathbf{a}_1 \ \mathbf{a}_2 \ \cdots \ \mathbf{a}_n \mid \mathbf{e}_1 \ \mathbf{e}_2 \ \cdots \ \mathbf{e}_n]$$

2. Reduce the left hand side (the matrix A) to reduced echelon form and apply the same row operations to I_n .
3. If this algorithm can be completed, the right hand side of the augmented matrix will be A^{-1} . That is

$$[A \mid I_n] \sim [I_n \mid A^{-1}]$$

Example 3.3.3. Find A^{-1} if $A = \begin{bmatrix} -1 & 4 & 1 \\ 1 & 0 & 1 \\ 2 & 0 & 1 \end{bmatrix}$. By applying all the necessary row operations we get

$$\begin{aligned} & \left[\begin{array}{ccc|ccc} -1 & 4 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 2 & 0 & 1 & 0 & 0 & 1 \end{array} \right] \sim \left[\begin{array}{ccc|ccc} 1 & -4 & -1 & -1 & 0 & 0 \\ 0 & 4 & 2 & 1 & 1 & 0 \\ 0 & 8 & 3 & 2 & 0 & 1 \end{array} \right] \sim \left[\begin{array}{ccc|ccc} 1 & -4 & -1 & -1 & 0 & 0 \\ 0 & 1 & 1/2 & 1/4 & 1/4 & 0 \\ 0 & 8 & 3 & 2 & 0 & 1 \end{array} \right] \\ & \sim \left[\begin{array}{ccc|ccc} 1 & -4 & -1 & -1 & 0 & 0 \\ 0 & 1 & 1/2 & 1/4 & 1/4 & 0 \\ 0 & 0 & 1 & 0 & 2 & -1 \end{array} \right] \sim \left[\begin{array}{ccc|ccc} 1 & -4 & 0 & -1 & 2 & -1 \\ 0 & 1 & 0 & 1/4 & -3/4 & 1/2 \\ 0 & 0 & 1 & 0 & 2 & -1 \end{array} \right] \sim \left[\begin{array}{ccc|ccc} 1 & 0 & 0 & 0 & -1 & 1 \\ 0 & 1 & 0 & 1/4 & -3/4 & 1/2 \\ 0 & 0 & 1 & 0 & 2 & -1 \end{array} \right] \end{aligned}$$

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

Definition 3.3.4. An $n \times n$ matrix A is **non-singular** if it has an inverse, otherwise we say it is **singular**. It is also important to note that if A^{-1} exists, it is unique.

Inverses also relate nicely to linear transformations.

Definition 3.3.5. If $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a linear transformation that is one-to-one and onto then T is **invertible**. Its inverse is the function $T^{-1}: \mathbb{R}^n \rightarrow \mathbb{R}^n$ with the property that for each $\mathbf{x} \in \mathbb{R}^n$ we have $T^{-1}(T(\mathbf{x})) = \mathbf{x}$. In fact, if T is given by $T(\mathbf{x}) = A\mathbf{x}$, then if T is invertible, we always have $T^{-1}(\mathbf{x}) = A^{-1}\mathbf{x}$.

Example 3.3.6. Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be given by $T\left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}\right) = \begin{bmatrix} 4x_1 + 3x_2 \\ -6x_1 + 5x_2 \end{bmatrix}$ so that $T(\mathbf{x}) = A\mathbf{x}$ with $A = \begin{bmatrix} 4 & 3 \\ -6 & 5 \end{bmatrix}$. Using the formula for the inverse of a 2×2 matrix, we have that

$$A^{-1} = \frac{1}{38} \begin{bmatrix} 5 & -3 \\ 6 & 4 \end{bmatrix}$$

We can then verify that

$$T^{-1}\left(T\left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}\right)\right) = T^{-1}\left(\begin{bmatrix} 4 & 3 \\ -6 & 5 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}\right) = \left(\frac{1}{38} \begin{bmatrix} 5 & -3 \\ 6 & 4 \end{bmatrix} \begin{bmatrix} 4 & 3 \\ -6 & 5 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}\right) = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Now that we have a bit of a handle of inverses of matrices, we can return to the algebraic mishaps of last section, and see that invertibility was indeed the solution we needed to make sense of when matrix multiplication behaves like regular multiplication of numbers.

Proposition 3.3.7. *Suppose A and B are non-singular $n \times n$ matrices and C and D are $n \times m$ matrices. Then*

1. A^{-1} is invertible with inverse $(A^{-1})^{-1} = A$.
2. AB is invertible and $(AB)^{-1} = B^{-1}A^{-1}$. This is known as the shoes and socks lemma. If B represents putting on your socks and A represents putting on your shoes, then undoing this process translates to **first** taking off your shoes (A^{-1}), then taking off your socks (B^{-1}).
3. If $AC = AD$, then $C = D$. We can obtain this logically by taking the first equation and multiplying **on the left** by A^{-1} on both sides.
4. If $AC = O_{nm}$ then $C = O_{nm}$. This can similarly be obtained by multiplying both sides by A^{-1} on the left.

It is essential to note here that invertibility of A is precisely what gives us the ability to draw all the conclusions we have made. Without invertibility of A , we cannot deduce any of the four statements.

We can now add some more results to the big theorem (which some refer to as the invertible matrix theorem).

Theorem 3.3.8. *Let $S = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ be a set of vectors in \mathbb{R}^n and let $A = [\mathbf{u}_1 \ \dots \ \mathbf{u}_n]$ with associated linear transformation given by $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$. The following statements are equivalent:*

1. S spans \mathbb{R}^n .
2. S is linearly independent.
3. The system $A\mathbf{x} = \mathbf{b}$ has a solution for every $\mathbf{b} \in \mathbb{R}^n$.
4. T is onto.
5. T is one-to-one.
6. A is invertible.

We now end the section with one illustration of why we love invertible matrices.

Example 3.3.9. Consider the linear system

$$\begin{cases} 4x_1 + 3x_2 = 5 \\ -2x_1 - x_2 = 7 \end{cases}$$

This system is the same as $A\mathbf{x} = \mathbf{b}$ for $A = \begin{bmatrix} 4 & 3 \\ -2 & -1 \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} 5 \\ 7 \end{bmatrix}$. Moreover, A is invertible with inverse given by $A^{-1} = \begin{bmatrix} -1/2 & -3/2 \\ 1 & 2 \end{bmatrix}$. Note that we found this matrix by using the formula for 2×2 matrices. Looking at the matrix equation $A\mathbf{x} = \mathbf{b}$, we can see that isolating \mathbf{x} is equivalent to multiplying both sides by A^{-1} **on the left**, hence

$$A^{-1}A\mathbf{x} = A^{-1}\mathbf{b} \implies \mathbf{x} = A^{-1}\mathbf{b}$$

and

$$A^{-1}\mathbf{b} = \begin{bmatrix} -1/2 & -3/2 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} 5 \\ 7 \end{bmatrix} = \begin{bmatrix} -13 \\ 19 \end{bmatrix} = \mathbf{x}$$

which uniquely solves the system.

Chapter 4

Basis and Subspaces

We now enter the second half of the topics list for this course, the first of which is subspaces. The language of subspaces gives us precise notions that allow one to describe things like planes and lines in \mathbb{R}^3 in greater generality. Once we have the basics of subspaces, we will define the all important notion of a basis, which will also lead us to the definition of dimension. We then take an in depth look at some of the most important subspaces related to a matrix, namely the column space and null space. We then finish the chapter with a description of change of basis, a central theme in all of linear algebra.

4.1 Subspaces

Definition 4.1.1. A subset S of \mathbb{R}^n is a subspace of \mathbb{R}^n if vectors in S satisfy the three following conditions:

1. $\mathbf{0} \in S$.
2. If $\mathbf{u}, \mathbf{v} \in S$, then $\mathbf{u} + \mathbf{v} \in S$. This is known as closure under addition.
3. If $r \in \mathbb{R}$ and $\mathbf{u} \in S$ then $r\mathbf{u} \in S$. This is known as closure under scaling.

It is worth noting that the first condition introduces the necessary condition that no two parallel subspaces can ever exist.

Example 4.1.2. Let $\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{R}^n$ and $S = \text{Span}\{\mathbf{u}_1, \mathbf{u}_2\}$. Is S a subspace of \mathbb{R}^n ?

We must verify each of the three conditions for S to be a subspace.

1. $\mathbf{0} \in S$ because $\mathbf{0} = 0\mathbf{u}_1 + 0\mathbf{u}_2$.
2. Elements of S are closed under addition. Let $\mathbf{u} = c_1\mathbf{u}_1 + c_2\mathbf{u}_2$ and $\mathbf{v} = d_1\mathbf{u}_1 + d_2\mathbf{u}_2$. Clearly both of these are arbitrary elements of S . By adding them together we see that

$$\mathbf{u} + \mathbf{v} = (c_1 + d_1)\mathbf{u}_1 + (c_2 + d_2)\mathbf{u}_2 \in S$$

which shows that the second condition is met.

3. If $r \in \mathbb{R}$ and $\mathbf{u} = c_1\mathbf{u}_1 + c_2\mathbf{u}_2 \in S$, then

$$r\mathbf{u} = rc_1\mathbf{u}_1 + rc_2\mathbf{u}_2 \in S$$

which shows that the third condition is met. We can now conclude that S is a subspace.

In general, the span of any set of vectors in \mathbb{R}^n is always a subspace of \mathbb{R}^n . This can easily be seen by reworking the above example with arbitrarily many vectors. This is such a fundamental fact that we state it as a theorem, which may be freely used from here on out.

Theorem 4.1.3. If $\mathbf{u}_1, \dots, \mathbf{u}_m \in \mathbb{R}^n$, then $\text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ is a subspace of \mathbb{R}^n .

Example 4.1.4. Let S be the set of solutions of the linear system

$$\begin{cases} 5x_1 + 5x_2 = 10 \\ x_1 + x_2 = 5 \end{cases}$$

Is S a subspace of \mathbb{R}^2 ?

NO! The easiest way to see this is by verifying that the first subspace condition is broken, That is, $\mathbf{0} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \notin S$ because it is not a solution to the *non-homogeneous* set of equations which define S .

Example 4.1.5. Let S be a subset of vectors in \mathbb{R}^3 consisting of the vectors $\begin{bmatrix} a \\ b \\ c \end{bmatrix}$ such that $ab = 0$. It turns

out that S is not a subspace of \mathbb{R}^3 because S is not closed under addition. For example, the vectors $\mathbf{u} = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$

and $\mathbf{v} = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}$ are both in S (condition that $ab = 0$ is satisfied) but

$$\mathbf{u} + \mathbf{v} = \begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix}$$

and $a = b = 1$ so $ab \neq 0$, which means that $\mathbf{u}, \mathbf{v} \in S$ but $\mathbf{u} + \mathbf{v} \notin S$.

We now introduce one of the fundamental subspaces associated to a matrix.

Theorem 4.1.6. Let A be an $n \times m$ matrix. If S is the set of solutions of the homogeneous linear system $A\mathbf{x} = \mathbf{0}$, then S is a subspace of \mathbb{R}^m .

Proof. First, we can see that $A\mathbf{0} = \mathbf{0}$ for any matrix A , hence $\mathbf{0} \in S$. Moreover, if $A\mathbf{u} = \mathbf{0}$ and $A\mathbf{v} = \mathbf{0}$ (meaning $\mathbf{u}, \mathbf{v} \in S$), then $\mathbf{u} + \mathbf{v} \in S$ because

$$A(\mathbf{u} + \mathbf{v}) = A\mathbf{u} + A\mathbf{v} = \mathbf{0} + \mathbf{0} = \mathbf{0}$$

Finally, we can also see that $r\mathbf{u} \in S$ for any $\mathbf{u} \in S$. By assuming that $A\mathbf{u} = \mathbf{0}$ (since $\mathbf{u} \in S$) we have that for any scalar $r \in \mathbb{R}$

$$A(r\mathbf{u}) = rA\mathbf{u} = r(\mathbf{0}) = \mathbf{0}$$

This shows that S is a subspace. □

Definition 4.1.7. If A is an $n \times m$ matrix, then the set of all solutions to the homogeneous linear system $A\mathbf{x} = \mathbf{0}$ is called **the null space of A** . It is denoted $\text{Null}(A)$ and is a subspace of \mathbb{R}^m . In other words

$$\text{Null}(A) = \left\{ \mathbf{x} \in \mathbb{R}^m : A\mathbf{x} = \mathbf{0} \right\}$$

Example 4.1.8. Find the null space of $A = \begin{bmatrix} 1 & -1 & 0 \\ 2 & 4 & 3 \end{bmatrix}$.

The procedure for finding the null space of a matrix is always the same. We begin by augmenting with the zero vector and row reducing.

$$\left[\begin{array}{ccc|c} 1 & -1 & 0 & 0 \\ 2 & 4 & 3 & 0 \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & -1 & 0 & 0 \\ 0 & 1 & 1/2 & 0 \end{array} \right]$$

Looking at the echelon matrix, we can see that x_3 is a free variable so we set $x_3 = t$. Using back substitution from here get the general solution

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} -\frac{1}{2}t \\ -\frac{1}{2}t \\ t \end{bmatrix} = t \begin{bmatrix} -1/2 \\ -1/2 \\ 1 \end{bmatrix}$$

which means that

$$\text{Null}(A) = \text{Span} \left\{ \begin{bmatrix} -1/2 \\ -1/2 \\ 1 \end{bmatrix} \right\}$$

In general, one may encounter a situation where they have to determine if a given set is a subspace. Here are some helpful tips to carry out this task successfully:

1. Check if $\mathbf{0} \in S$. If not, then S is not a subspace.
2. If you can find **specific** vectors whose span is precisely equal to S , then you can leverage Theorem 4.1.3 to argue that S is a subspace.
3. Recognize that S can in fact be expressed as the null space of some matrix and leverage Theorem 4.1.6 to show that S is a subspace (this method is powerful if you can get good at using it). As

an example, consider the set of vectors of the form $\begin{bmatrix} a \\ b \\ c \end{bmatrix}$ such that $a - b = -c$. The condition that

$a - b = -c$ is equivalent to $a - b + c = 0$. This set of vectors is then the solution set of the linear system $x_1 - x_2 + x_3 = 0$ which can be expressed as $\text{Null}(A)$ where $A = \begin{bmatrix} 1 & -1 & 1 \end{bmatrix}$. Note that in general, if you can algebraically manipulate something to obtain a zero somewhere, you are probably looking at a null space in disguise.

4. If all else fails, show closure under addition and scaling directly. If you encounter a road block in trying to prove this, it may mean that S is not a subspace. If you suspect this is the case, you should then seek out a counterexample. Either two vectors in S whose sum is not in S , or a fixed vector and fixed scalar which break closure under scaling.

We end this section by investigating how this relates to linear maps.

Definition 4.1.9. Let $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear transformation. The set of all vectors $x \in \mathbb{R}^m$ such that $T(\mathbf{x}) = \mathbf{0}$ is called **the kernel of T** and is denoted $\ker(T)$.

Theorem 4.1.10. If $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a linear transformation, then $\ker(T)$ is a subspace of \mathbb{R}^m and $\text{Range}(T)$ is a subspace of \mathbb{R}^n (recall that $\text{Range}(T) = \{\mathbf{y} \in \mathbb{R}^n : T(\mathbf{x}) = \mathbf{y} \text{ for some } \mathbf{x} \in \mathbb{R}^m\}$).

Proof. The proof of this is very instructive and will be useful for the remainder of the course.

Since T is a linear transformation, we know that there exists a matrix A such that $T(\mathbf{x}) = A\mathbf{x}$. This means that if $T(\mathbf{x}) = \mathbf{0}$ then $A\mathbf{x} = \mathbf{0}$ hence $\ker(T)$ and $\text{Null}(A)$ are the same! By Theorem 4.1.6 we can conclude that $\ker(T)$ is a subspace. Similarly by recalling the formula of Definition 2.2.9, we can see that

$$\text{Range}(T) = \text{Span}\{\mathbf{a}_1, \dots, \mathbf{a}_m\}$$

where $A = [\mathbf{a}_1 \ \dots \ \mathbf{a}_m]$ (Remember this fact!). It then follows from Theorem 4.1.3 that $\text{Range}(T)$ is a subspace. \square

Tracing back to the results of the previous chapter, we now have a nice new fact.

Proposition 4.1.11. *Let $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear transformation. T is one-to-one if and only if $\ker(T) = \{\mathbf{0}\}$.*

We now end the section by adding to the big theorem.

Theorem 4.1.12. *Let $S = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ be a set of vectors in \mathbb{R}^n and let $A = [\mathbf{u}_1 \ \cdots \ \mathbf{u}_n]$ with associated linear transformation given by $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$. The following statements are equivalent:*

1. S spans \mathbb{R}^n .
2. S is linearly independent.
3. The system $A\mathbf{x} = \mathbf{b}$ has a solution for every $\mathbf{b} \in \mathbb{R}^n$.
4. T is onto.
5. T is one-to-one.
6. A is invertible.
7. $\ker(T) = \{\mathbf{0}\}$.

4.2 Basis and Dimension

We saw in the previous section that spans of any number of vectors always forms a subspace. From this fact, we can ask the question, is every subspace the span of some set of vectors? The answer to this is yes! Moreover, we can go one step further and ask whether or not we can find the smallest set of vectors that span a given subspace. It is the notion of a basis that stems from this idea.

Definition 4.2.1. Let S be a subspace of \mathbb{R}^n . A **basis** for S is a set of vectors $\mathcal{B}_S = \{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ that spans S and is linearly independent.

Example 4.2.2. Let $S = \text{Span} \left\{ \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \begin{bmatrix} -3 \\ -6 \end{bmatrix}, \begin{bmatrix} 10 \\ 20 \end{bmatrix} \right\}$. We can observe that $\begin{bmatrix} -3 \\ -6 \end{bmatrix} = -3 \begin{bmatrix} 1 \\ 2 \end{bmatrix}$ and $\begin{bmatrix} 10 \\ 20 \end{bmatrix} = 10 \begin{bmatrix} 1 \\ 2 \end{bmatrix}$, hence, we can see that the spanning vectors for S are linearly dependent. Based on the definition for a basis, this means that the given vectors are not a basis. Moreover, along the lines of Proposition 2.2.8, we have that

$$\text{Span} \left\{ \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \begin{bmatrix} -3 \\ -6 \end{bmatrix}, \begin{bmatrix} 10 \\ 20 \end{bmatrix} \right\} = \text{Span} \left\{ \begin{bmatrix} 1 \\ 2 \end{bmatrix} \right\}$$

Since $\left\{ \begin{bmatrix} 1 \\ 2 \end{bmatrix} \right\}$ spans S and is linearly independent, we have that $\mathcal{B}_S = \left\{ \begin{bmatrix} 1 \\ 2 \end{bmatrix} \right\}$ forms a basis for S .

Example 4.2.3. Let

$$\mathbf{u}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \mathbf{u}_2 = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}, \mathbf{u}_3 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \text{ and } \mathbf{u}_4 = \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$$

and let $S = \text{Span}\{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \mathbf{u}_4\}$. We can observe that $\mathbf{u}_3 = \mathbf{u}_1 + \mathbf{u}_2$ and $\mathbf{u}_4 = \mathbf{u}_1 + 2\mathbf{u}_2$ hence $S = \text{Span}\{\mathbf{u}_1, \mathbf{u}_2\}$. Since $\{\mathbf{u}_1, \mathbf{u}_2\}$ is linearly independent, we can conclude that $\{\mathbf{u}_1, \mathbf{u}_2\}$ is a basis for S .

Example 4.2.4. Consider the zero vector $\mathbf{0} \in \mathbb{R}^n$. The **zero subspace** $S = \{\mathbf{0}\}$ is the only subspace of \mathbb{R}^n that has no basis. It consists of the origin and nothing else.

A task that will arise again and again is that of finding a basis for a given subspace. There are two ways of doing this and we break down each "recipe". Both have their advantages depending on the context in which you want to find a basis. In both cases, assume $S = \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_n\}$.

1. **Recipe 1:**

- Form a matrix A whose **ROWS** are the vectors $\mathbf{u}_1, \dots, \mathbf{u}_n$.
- Use row reductions to transform A to an echelon matrix B .
- The **non-zero** rows of B form a basis for S .

Example 4.2.5. Let $\mathbf{u}_1 = \begin{bmatrix} 1 \\ -2 \\ 3 \\ -2 \end{bmatrix}$, $\mathbf{u}_2 = \begin{bmatrix} 1 \\ 1 \\ -1 \\ 0 \end{bmatrix}$, and $\mathbf{u}_3 = \begin{bmatrix} 3 \\ -3 \\ 5 \\ 4 \end{bmatrix}$ and suppose $S = \text{Span}\{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3\}$.

Then

$$A = \begin{bmatrix} 1 & -2 & 3 & -2 \\ 1 & 1 & -1 & 0 \\ 3 & -3 & 5 & -4 \end{bmatrix} \sim \begin{bmatrix} 1 & -2 & 3 & -2 \\ 0 & 3 & -4 & 2 \\ 0 & 0 & 0 & 0 \end{bmatrix} = B$$

The non-zero rows of B form a basis for S hence $\mathcal{B}_S = \left\{ \begin{bmatrix} -1 \\ -2 \\ 3 \\ -2 \end{bmatrix}, \begin{bmatrix} 0 \\ 3 \\ -4 \\ 2 \end{bmatrix} \right\}$ is a basis for S .

2. **Recipe 2:**

- Form a matrix A whose **COLUMNS** are the vectors $\mathbf{u}_1, \dots, \mathbf{u}_n$.
- Use row reductions to transform A to an echelon matrix B .
- The columns of A that correspond to the pivot columns of B form a basis for S .

Example 4.2.6. Let $\mathbf{u}_1 = \begin{bmatrix} 1 \\ -2 \\ 3 \\ -2 \end{bmatrix}$, $\mathbf{u}_2 = \begin{bmatrix} 1 \\ 1 \\ -1 \\ 0 \end{bmatrix}$, and $\mathbf{u}_3 = \begin{bmatrix} 3 \\ -3 \\ 5 \\ 4 \end{bmatrix}$ and suppose $S = \text{Span}\{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3\}$.

Then

$$A = \begin{bmatrix} 1 & 1 & 3 \\ -2 & 1 & -3 \\ 3 & -1 & 5 \\ -2 & 0 & -4 \end{bmatrix} \sim \begin{bmatrix} 1 & 1 & 3 \\ 0 & 3 & 3 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} = B$$

The pivot columns of B are columns 1 and 2 hence our basis for S is

$$\mathcal{B}_S = \left\{ \begin{bmatrix} 1 \\ -2 \\ 3 \\ -2 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ -1 \\ 0 \end{bmatrix} \right\}$$

As a general rule, one should always remember that both recipes always work but recipe 1 tends to give "simpler" basis vectors (with more zeroes) whereas recipe 2 always gives basis vectors that are a subset of the vectors you started with. It is very common to want to reduce a spanning set to a basis (known as **reducing to a basis**), and this makes recipe 2 especially useful in many scenarios. We can now use the notion of a basis to define dimension. The first fundamental fact that we need is the following.

Proposition 4.2.7. *If S is a subspace of \mathbb{R}^n , then every basis of S has the same number of vectors in it.*

Definition 4.2.8. The **dimension** of a subspace S , denoted $\dim(S)$, is the number of vectors in any basis for S . Note that in the previous example, we had $\dim(S) = 2$. In general, we always have $\dim(\{\mathbf{0}\}) = 0$.

Example 4.2.9. If S is a subspace of \mathbb{R}^3 , what are the possible values of $\dim(S)$?

- S could be the zero subspace, in which case $\dim(S) = 0$.
- S could be a line through the origin in which case it has the form $S = \text{Span}\{\mathbf{u}_1\}$ for $\mathbf{u}_1 \neq \mathbf{0}$ and $\dim(S) = 1$.
- S could be a plane through the origin in which case it has the form $S = \text{Span}\{\mathbf{u}_1, \mathbf{u}_2\}$ for $\mathbf{u}_1, \mathbf{u}_2$ linearly independent. In this case we have $\dim(S) = 2$.
- S could be all of \mathbb{R}^3 , which we could write as

$$\mathbb{R}^3 = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\}$$

We call $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ the **standard basis of \mathbb{R}^3** . In this case, $\dim(S) = 3$ and in general, this is the only 3-dimensional subspace of \mathbb{R}^3 .

This completes our list because any subspace $S = \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ where $m > 3$ can never be m dimensional. This follows from the fact that any set of $m > 3$ vectors in \mathbb{R}^3 is never linearly independent, hence we can never have a basis containing more than 3 vectors.

Let's illustrate all of these ideas on some more complex examples.

Example 4.2.10. Find a basis for \mathbb{R}^4 containing the vectors

$$\mathbf{u}_1 = \begin{bmatrix} 1 \\ 0 \\ 2 \\ 0 \end{bmatrix} \quad \text{and} \quad \mathbf{u}_2 = \begin{bmatrix} 2 \\ 2 \\ 1 \\ 1 \end{bmatrix}$$

We know that $\mathcal{B}_{\mathbb{R}^4} = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4\}$ is a basis for \mathbb{R}^4 and since $\mathbf{u}_1, \mathbf{u}_2 \in \text{Span}\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4\}$ we know (by Proposition 2.2.8) that

$$\mathbb{R}^4 = \text{Span}\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4\} = \text{Span}\{\mathbf{u}_1, \mathbf{u}_2, \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4\}$$

We can then apply recipe 2, placing \mathbf{u}_1 and \mathbf{u}_2 as the left-most vectors. Upon row reducing we get that

$$A = \begin{bmatrix} 1 & 2 & 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 1 & 0 & 0 \\ 2 & 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \end{bmatrix} \sim \begin{bmatrix} 1 & 2 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & -2 & 0 & 1 & 3 \\ 0 & 0 & 0 & 1 & 0 & -2 \end{bmatrix} = B$$

The pivots columns of B are columns 1, 2, 3, and 4 hence

$$\mathcal{B}_{\mathbb{R}^4} = \left\{ \begin{bmatrix} 1 \\ 0 \\ 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 2 \\ 2 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \right\}$$

is a basis for \mathbb{R}^4 containing the prescribed vectors.

Next, let's up the difficulty a little bit and find a basis for a new but increasingly familiar subspace.

Example 4.2.11. Let $A = \begin{bmatrix} 1 & -1 & 1 & 0 \\ 0 & 1 & -2 & 3 \\ 2 & -1 & 0 & 3 \end{bmatrix}$ and compute $\dim(\text{Null}(A))$.

We first need to find $\text{Null}(A)$ which involves solving the linear system $A\mathbf{x} = \mathbf{0}$. We see that

$$\left[\begin{array}{cccc|c} 1 & -1 & 1 & 0 & 0 \\ 0 & 1 & -2 & 3 & 0 \\ 2 & -1 & 0 & 3 & 0 \end{array} \right] \sim \left[\begin{array}{cccc|c} 1 & -1 & 1 & 0 & 0 \\ 0 & 1 & -2 & 3 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right]$$

We have 2 free variables so we set $x_3 = t$ and $x_4 = s$. Then, by back substitution, we get $x_2 = 2t - 3s$ and $x_1 = t - 2s$ hence

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = t \begin{bmatrix} 1 \\ 2 \\ 1 \\ 0 \end{bmatrix} + s \begin{bmatrix} -3 \\ -3 \\ 0 \\ 1 \end{bmatrix}$$

It then follows immediately that

$$\mathcal{B}_{\text{Null}(A)} = \left\{ \begin{bmatrix} 1 \\ 2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -3 \\ -3 \\ 0 \\ 1 \end{bmatrix} \right\}$$

is a basis for $\text{Null}(A)$ and we can conclude that $\dim(\text{Null}(A)) = 2$. This procedure for finding a basis **always** works because the free variables will always contribute a 1 to one entry of a basis vector and a 0 to the corresponding entries of all other vectors. The offset 0's and 1's always ensure linear independence of the spanning vectors that we find, hence a basis is obtained automatically.

This number is so important that it has its own name.

Definition 4.2.12. The **nullity** of a matrix A , denoted $\text{nullity}(A)$, is the number $\dim(\text{Null}(A))$.

We will have much more to say about this numerical invariant in the next section, but before ending this section, we add to the big theorem once more.

Theorem 4.2.13. Let $S = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ be a set of vectors in \mathbb{R}^n and let $A = [\mathbf{u}_1 \ \dots \ \mathbf{u}_n]$ with associated linear transformation given by $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$. The following statements are equivalent:

1. S spans \mathbb{R}^n .
2. S is linearly independent.
3. The system $A\mathbf{x} = \mathbf{b}$ has a solution for every $\mathbf{b} \in \mathbb{R}^n$.
4. T is onto.
5. T is one-to-one.
6. A is invertible.
7. $\ker(T) = \{\mathbf{0}\}$.
8. S is a basis for \mathbb{R}^n .

4.3 Row Space, Column Space, and Rank

We now introduce several more fundamental subspaces associated to a matrix. Once we have these additional definitions, we state the all important Rank-Nullity theorem, sometimes known as the fundamental theorem of linear algebra. This theorem allows us to “decompose” \mathbb{R}^n into disjoint subspaces.

Definition 4.3.1. Let A be an $n \times m$ matrix.

- The **row space** of A is the subspace of \mathbb{R}^m spanned by the row vectors of A . It is denoted $\text{row}(A)$.
- The **column space** of A is the subspace of \mathbb{R}^n spanned by the columns of A . It is denoted $\text{col}(A)$ and is the set of all outputs of the form $A\mathbf{x}$ or alternatively, just the span of the columns of A .

Combining these definition with our “recipes” from the last section we can deduce that given any matrix $A \sim B$ with B in echelon form

- The non-zero rows of B form a basis for $\text{row}(A)$.
- The columns of A corresponding to the pivot columns of B form a basis for $\text{col}(A)$.

Example 4.3.2.

$$A = \begin{bmatrix} 1 & 2 & 1 & -1 \\ 0 & 1 & 1 & 0 \\ -1 & 5 & 3 & 0 \end{bmatrix} \sim \begin{bmatrix} 1 & 2 & 1 & -2 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} = B$$

Using the recipes, we can see that

$$\mathcal{B}_{\text{row}(A)} = \left\{ \begin{bmatrix} 1 \\ 2 \\ 1 \\ -1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} \quad \text{and} \quad \mathcal{B}_{\text{col}(A)} = \left\{ \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}, \begin{bmatrix} 2 \\ 1 \\ 5 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix} \right\}$$

You may notice in this example that the row space and column space have the same dimension, even though one of them is a subspace of \mathbb{R}^4 and the other one is a subspace of \mathbb{R}^3 . It turns out this phenomenon is always true.

Theorem 4.3.3. *Given any matrix A we have*

$$\dim(\text{Row}(A)) = \dim(\text{Col}(A))$$

Proof. Let B be a matrix in echelon form that is row equivalent to A . Every non-zero row of B contains a pivot and similarly, the pivot in each pivot column must lie in one of these non-zero rows. This means that the number of non-zero rows of B must equal the number of pivot columns of B . The number of pivot rows (resp. columns) is precisely what we use to find bases of these subspaces, hence these numbers always being equal implies that $\text{row}(A)$ and $\text{col}(A)$ must always have the same dimension. \square

This new numerical invariant also has its own name.

Definition 4.3.4. The **rank** of a matrix A , denoted $\text{rank}(A)$, is the dimension of the row, or column, space of A . In the above example we have $\text{rank}(A) = 3$ and we say that the matrix A has rank 3.

We now have everything we need to state what is, without question, the most amazing and useful theorem in this course, known most commonly as the rank-nullity theorem.

Theorem 4.3.5. *If A is an $n \times m$ matrix then*

$$\text{rank}(A) + \text{Nullity}(A) = m$$

Proof. If A is an $n \times m$ matrix, and $A \sim B$, then the number of non-zero rows of B is the rank of A by definition. This is also equal to the number of pivot columns of B . Each non-pivot column will correspond to a free variable and each free variable corresponds to a basis vector for $\text{Null}(A)$ (think about how we did this in Example 4.2.11). Putting this all together we have that

$\text{rank}(A) =$ the number of pivot columns of B

and

$\text{Nullity}(A) =$ the number of non-pivot columns of B .

The total number of columns of B , which is equal to m , is then the sum of the $\text{rank}(A)$ and $\text{Nullity}(A)$. \square

The power of this theorem pops up again and again but at this stage, we can already find it useful in doing routine computations. In particular, if you want to find the rank or nullity of a matrix, you only need to find one and you get the other for free. This allows you to apply the “find a basis for the null space” procedure of Example 4.2.11 to find the nullity, or, apply your favorite recipe to find the rank, then the other numerical invariant follows immediately from rank-nullity.

Example 4.3.6. Consider the following matrix, and an equivalent echelon form

$$A = \begin{bmatrix} 1 & 2 & 1 & -1 \\ 0 & 1 & 1 & 0 \\ -1 & 5 & 3 & 0 \end{bmatrix} \sim \begin{bmatrix} 1 & 2 & 1 & -2 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} = B$$

We saw in the previous example that $\text{rank}(A) = 3$ so we immediately know that this matrix has nullity equal

to 1. You should verify this for yourself and in doing so will see, that $\left\{ \begin{bmatrix} 5/3 \\ 1/3 \\ -1/3 \\ 1 \end{bmatrix} \right\}$ is a basis for $\text{Null}(A)$.

We can also relate this theorem to linear transformations. Recall that given a linear transformation $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ with associated matrix A , we deduced that the span of the columns of A was equal to the range. If this does not ring a bell, take a look at Definition 2.2.9. With our new terminology, this means that $\text{Col}(A) = \text{Range}(T)$. Moreover, the solution set of $A\mathbf{x} = \mathbf{0}$ consisted of the vectors \mathbf{x} such that $T(\mathbf{x}) = \mathbf{0}$. In other words, we had $\ker(T) = \text{Null}(A)$. This means that $\text{rank}(A) = \dim(\text{Range}(T))$ and $\text{Nullity}(A) = \dim(\ker(T))$. This is a dense paragraph but is worth spending the time to understand every sentence.

Combining these geometric notions with the rank-nullity theorem we can see that

$$m = \dim(\text{Range}(T)) + \dim(\ker(T))$$

It is worth noting that the dimension of the row space being equal to $\text{rank}(A)$ and the dimension of the null space being equal to $\text{Nullity}(A)$ says something significant about \mathbb{R}^m . Both the row space and null space of A are subspaces of \mathbb{R}^m , whose dimensions add up to m . It is rank-nullity that allows us to conclude that \mathbb{R}^m “decomposes” into the row space and the null space of the given matrix. This would not be possible to understand without our notions of linear maps and the rank-nullity theorem.

We now finish the section with one more example, followed by one more addition to the big theorem.

Example 4.3.7. Let $T: \mathbb{R}^{11} \rightarrow \mathbb{R}^9$ be given by $T(\mathbf{x}) = A\mathbf{x}$ and further assume that T is onto. How many dimensions of \mathbb{R}^{11} are occupied by $\ker(T)$?

Since T is onto, we know that its range is the entire codomain, that is, $\text{Range}(T) = \mathbb{R}^9$. This means that $\dim(\text{Range}(T)) = \dim(\text{Col}(A)) = \text{rank}(A) = 9$. Rank-nullity then implies that

$$11 = 9 + \dim(\ker(T))$$

hence $\dim(\ker(T)) = 2$.

Theorem 4.3.8. Let $S = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ be a set of vectors in \mathbb{R}^n and let $A = [\mathbf{u}_1 \ \cdots \ \mathbf{u}_n]$ with associated linear transformation given by $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$. The following statements are equivalent:

1. S spans \mathbb{R}^n .
2. S is linearly independent.
3. The system $A\mathbf{x} = \mathbf{b}$ has a solution for every $\mathbf{b} \in \mathbb{R}^n$.
4. T is onto.
5. T is one-to-one.
6. A is invertible.
7. $\ker(T) = \{\mathbf{0}\}$.
8. S is a basis for \mathbb{R}^n .
9. $\text{rank}(A) = n$.
10. $\text{Nullity}(A) = 0$.

This is quite a bit of information so we briefly summarize the main ideas in the following list:

- $\text{Null}(A) = \{\mathbf{x} \in \mathbb{R}^m : A\mathbf{x} = \mathbf{0}\} = \{\mathbf{x} \in \mathbb{R}^m : T(\mathbf{x}) = \mathbf{0}\} = \ker(T)$.
- $\text{Col}(A) = \text{span of columns of } A = \text{Range}(T)$.
- $\dim(\text{Col}(A)) = \dim(\text{Row}(A)) = \text{rank}(A)$.
- $\dim(\text{Null}(A)) = \text{Nullity}(A)$.
- If $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ then $\text{Null}(A) \subset \mathbb{R}^m$, $\text{Row}(A) \subset \mathbb{R}^m$, and $\text{Col}(A) \subset \mathbb{R}^n$.

4.4 Change of Basis

We now encounter the all important idea surrounding changing a basis. This can be one of the trickiest concepts to understand, but the hard work will pay off. Reading this section several times over may be helpful in gaining a full understanding and when in doubt, do more examples!

Let's first address notation. Let $\mathbf{x} = \begin{bmatrix} 3 \\ -2 \end{bmatrix} \in \mathbb{R}^2$ be written in the standard basis. The coordinates of \mathbf{x} are expressing its geometric location in the plane. That is, to arrive at the tip of the vector \mathbf{x} , you move 3 units to the right of the origin (3 units along \mathbf{e}_1) and -2 units down from there (-2 units along \mathbf{e}_2). This is because

$$\mathbf{x} = 3\mathbf{e}_1 - 2\mathbf{e}_2$$

The coefficients of \mathbf{x} in this expression involving the standard basis are what determine its coordinates. This is the general idea behind change of basis.

Example 4.4.1. Let $\mathcal{B} = \left\{ \begin{bmatrix} 2 \\ 7 \end{bmatrix}, \begin{bmatrix} 1 \\ 4 \end{bmatrix} \right\}$ be a (non-standard) basis of \mathbb{R}^2 . In this basis we can express the same vector \mathbf{x} as

$$\mathbf{x} = 14 \begin{bmatrix} 2 \\ 7 \end{bmatrix} - 25 \begin{bmatrix} 1 \\ 4 \end{bmatrix}$$

and we express this notationally as

$$[\mathbf{x}]_{\mathcal{B}} = \begin{bmatrix} 14 \\ -25 \end{bmatrix}$$

With this idea in mind, we can now define this notion in greater generality.

Definition 4.4.2. Let $\mathcal{B} = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n\}$ be a basis for \mathbb{R}^n and let

$$\mathbf{y} = a_1\mathbf{u}_1 + a_2\mathbf{u}_2 + \dots + a_n\mathbf{u}_n$$

then **the coordinate vector of \mathbf{y} with respect to the basis \mathcal{B}** is

$$[\mathbf{y}]_{\mathcal{B}} = \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{bmatrix}$$

Let $U = [\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_n] \in \mathbb{R}^{n \times n}$. We call U the **change of basis matrix for the basis \mathcal{B}** (note that it has the basis vectors as its columns). If \mathbf{y} is taken to be a vector written in the standard basis, then

$$U[\mathbf{y}]_{\mathcal{B}} = [\mathbf{u}_1 \ \dots \ \mathbf{u}_n] \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} = a_1\mathbf{u}_1 + a_2\mathbf{u}_2 + \dots + a_n\mathbf{u}_n = \mathbf{y}$$

Example 4.4.3. Let

$$\mathcal{B} = \left\{ \begin{bmatrix} 1 \\ 3 \\ -2 \end{bmatrix}, \begin{bmatrix} 2 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 4 \\ 5 \\ 1 \end{bmatrix} \right\}$$

be a basis for \mathbb{R}^3 and let $[\mathbf{x}]_{\mathcal{B}} = \begin{bmatrix} -2 \\ 3 \\ 1 \end{bmatrix}$. Find \mathbf{x} with respect to the standard basis for \mathbb{R}^3 .

Given the basis \mathcal{B} , our change of basis matrix is

$$U = \begin{bmatrix} 1 & 2 & 4 \\ 3 & 0 & 5 \\ -2 & 1 & 1 \end{bmatrix}$$

so we can find \mathbf{x} via

$$\mathbf{x} = U[\mathbf{x}]_{\mathcal{B}} = \begin{bmatrix} 1 & 2 & 4 \\ 3 & 0 & 5 \\ -2 & 1 & 1 \end{bmatrix} \begin{bmatrix} -2 \\ 3 \\ 1 \end{bmatrix} = \begin{bmatrix} 8 \\ -1 \\ 8 \end{bmatrix}$$

Note that U took a vector **from** the non-standard basis **to** the standard basis.

A natural question one can ask is, how can we go the other direction? That is, if we are given a vector, written in the standard basis, how can we find its representation in some other non-standard basis?

The key is to look at the equation we get from the change of basis matrix, namely

$$\mathbf{x} = U[\mathbf{x}]_{\mathcal{B}}$$

We can see that U is **always** invertible (by the big theorem) because its columns form a basis, hence we can take the above equation and multiply both sides by U^{-1} on the left to obtain

$$U^{-1}\mathbf{x} = [\mathbf{x}]_{\mathcal{B}}$$

We summarize in the following proposition.

Proposition 4.4.4. *Let \mathbf{x} be expressed in the standard basis with $\mathcal{B} = \{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ a non-standard basis for \mathbb{R}^n . If $U = [\mathbf{u}_1 \ \dots \ \mathbf{u}_n]$ is the change of basis matrix for the basis \mathcal{B} then*

$$U[\mathbf{x}]_{\mathcal{B}} = \mathbf{x} \quad \text{and} \quad [\mathbf{x}]_{\mathcal{B}} = U^{-1}\mathbf{x}$$

Example 4.4.5. Continuing from example 4.4.1, we have $\mathcal{B} = \left\{ \begin{bmatrix} 2 \\ 7 \end{bmatrix}, \begin{bmatrix} 1 \\ 4 \end{bmatrix} \right\}$ and $\mathbf{x} = \begin{bmatrix} 3 \\ -2 \end{bmatrix}$. Going from the standard basis to this one we see that

$$[\mathbf{x}]_{\mathcal{B}} = \begin{bmatrix} 14 \\ -25 \end{bmatrix} = \begin{bmatrix} 2 & 1 \\ 7 & 4 \end{bmatrix}^{-1} \begin{bmatrix} 3 \\ -2 \end{bmatrix}$$

In words, the proposition is saying that U takes a vector **from** the non-standard basis to the standard basis, and its inverse does the opposite.

What remains is to find a fluid way to go from one non-standard basis to another. The short solution is to “go through the standard basis” but this requires some explanation.

Let $\mathcal{B}_1 = \{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ and $\mathcal{B}_2 = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ be non-standard bases for \mathbb{R}^n . We aim to find a matrix that takes $[\mathbf{x}]_{\mathcal{B}_1}$ as input, and outputs $[\mathbf{x}]_{\mathcal{B}_2}$. Let \mathcal{B}_{st} denote the standard basis for \mathbb{R}^n . We carry out the task in two steps

1. Go from $[\mathbf{x}]_{\mathcal{B}_1}$ to $[\mathbf{x}]_{\mathcal{B}_{st}}$.
2. Go from $[\mathbf{x}]_{\mathcal{B}_{st}}$ to $[\mathbf{x}]_{\mathcal{B}_2}$.

We use matrix multiplication to combine the steps.

Theorem 4.4.6. Let $\mathcal{B}_1 = \{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ and $\mathcal{B}_2 = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ be non-standard bases for \mathbb{R}^n with change of basis matrices given by $U = [\mathbf{u}_1 \ \dots \ \mathbf{u}_m]$ and $V = [\mathbf{v}_1 \ \dots \ \mathbf{v}_n]$ respectively. Then

$$[\mathbf{x}]_{\mathcal{B}_2} = V^{-1}U[\mathbf{x}]_{\mathcal{B}_1}$$

and

$$[\mathbf{x}]_{\mathcal{B}_1} = U^{-1}V[\mathbf{x}]_{\mathcal{B}_2}$$

Proof. We know that U and V are change of basis matrices, hence by Proposition 4.4.4 we know that


$$\mathbf{x} = U[\mathbf{x}]_{\mathcal{B}_1} \quad \text{and} \quad [\mathbf{x}]_{\mathcal{B}_2} = V^{-1}\mathbf{x}$$

Note here that we are writing \mathbf{x} to mean $[\mathbf{x}]_{\mathcal{B}_{st}}$ (this is standard convention). Combining these two equations we see that

$$[\mathbf{x}]_{\mathcal{B}_2} = V^{-1}\mathbf{x} = V^{-1}(U[\mathbf{x}]_{\mathcal{B}_1}) = V^{-1}U[\mathbf{x}]_{\mathcal{B}_1}$$

This means that the change of basis matrix from \mathcal{B}_1 to \mathcal{B}_2 is $V^{-1}U$. By taking inverses and using the shoes and socks lemma, we get the second result. \square

We end this chapter with an illustration of this entire idea. The whole of change of basis can be summarized in the following picture.



Chapter 5

Determinants

The determinant can be thought of as a useful number that we can associate with a fixed matrix. In particular, viewing it as a function, it takes an $n \times n$ matrix as input and outputs a real number. In this chapter we will begin by discussing ways to compute the determinant of a matrix, and once we have the basics down, we will see how it can be used.

5.1 The Determinant Function

We can compute the determinant of $n \times n$ matrices, for small n , quite easily.

1. $n = 1$: If $A = [a_{11}]$ then $\det(A) = a_{11}$.
2. $n = 2$: If $A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$ then $\det(A) = a_{11}a_{22} - a_{12}a_{21}$.
3. $n = 3$: If $A = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}$ then

$$\det(A) = a_{11} \det \left(\begin{bmatrix} a_{22} & a_{23} \\ a_{32} & a_{33} \end{bmatrix} \right) - a_{12} \det \left(\begin{bmatrix} a_{21} & a_{23} \\ a_{31} & a_{33} \end{bmatrix} \right) + a_{13} \det \left(\begin{bmatrix} a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} \right)$$

The following definition is very formal and can be quite complicated to understand. The best way to get a grasp on it is to do LOTS of examples! Starting with a 3×3 matrix is the best place to begin. For practice (after reading the definition), compute the determinant of $A = \begin{bmatrix} 1 & 2 & 3 \\ 1 & -1 & 4 \\ 5 & 6 & 2 \end{bmatrix}$ and verify that the final answer is 43.

Definition 5.1.1. Let A be an $n \times n$ matrix given by

$$(a_{ij}) = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ \vdots & \vdots & \cdots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix}$$

Note that writing a matrix as (a_{ij}) is common compact matrix notation, which denotes that the entry in row i and column j is the real number a_{ij} . For $n = 2, \dots, n$, let M_{ij} be that $(n-1) \times (n-1)$ matrix obtained by removing the i^{th} row and j^{th} column

of A . The minor of a_{ij} is the real number $\det(M_{ij})$. The **cofactor** of a_{ij} is $C_{ij} = (-1)^{i+j} \det(M_{ij})$. The **determinant of A** is then the scalar

$$\begin{aligned} \det(A) &= a_{i1}C_{i1} + a_{i2}C_{i2} + \cdots + a_{in}C_{in} \\ &= a_{1j}C_{1j} + a_{2j}C_{2j} + \cdots + a_{nj}C_{nj} \end{aligned}$$

where i and j can be any **fixed** values from 1 up to n . Note that the first equation is known as the **cofactor expansion along the i^{th} row**, and the second equation is known as the **cofactor expansion along the j^{th} column**.

The profound fact concerning computation of determinants is the following.

Theorem 5.1.2. *Given any $n \times n$ matrix A , the value of $\det(A)$ obtained by performing a cofactor expansion along the i^{th} row or j^{th} column is always the same.*

The main consequence of this theorem is that, when computing the determinant of a matrix, we can seek out the least labor intensive method possible. In practice, this involves finding the row or column of the given matrix that has the most zeroes, and computing a cofactor expansion along that row or column. This is always the least computationally expensive method.

Example 5.1.3. Compute the determinant of

$$A = \begin{bmatrix} 1 & 2 & 2 & 4 \\ 1 & 0 & 0 & 3 \\ 5 & 6 & 0 & 7 \\ 3 & 1 & 0 & 8 \end{bmatrix}$$

Observe that the third column of A has the most zeroes. Moreover, if we compute the cofactor expansion along the 3rd column, we will only need to compute one minor explicitly (as opposed to as many as 4!).

$$\det(A) = 2 \det \left(\begin{bmatrix} 1 & 0 & 3 \\ 5 & 6 & 7 \\ 3 & 1 & 8 \end{bmatrix} \right) = 2 \left(1 \cdot \det \left(\begin{bmatrix} 6 & 7 \\ 1 & 8 \end{bmatrix} \right) + 3 \cdot \det \left(\begin{bmatrix} 5 & 6 \\ 3 & 1 \end{bmatrix} \right) \right) = 4$$

5.2 Properties of the Determinant

One other way in which we can compute a determinant is to row reduce the given matrix and track how the determinant changes at each step. We do so according to the following proposition.

Proposition 5.2.1. *Suppose B is an $n \times n$ matrix obtained by performing one of the following row operations on A . The determinant of B and A are related as follows:*

1. *Switch two rows of A to get $B \implies \det(B) = -\det(A)$.*
2. *Multiply a row of A by a non-zero constant c to get $B \implies \det(B) = c \det(A)$.*
3. *Add a multiple of one row to another to get $B \implies \det(B) = \det(A)$.*

This Proposition implies the following shortcuts:

- If A has a row or column of zeroes then $\det(A) = 0$.
- If A has a two identical rows then $\det(A) = 0$.

We also add one more useful trick in computing determinants of triangular matrices.

Proposition 5.2.2. *If A is a triangular matrix, then $\det(A)$ is the product of the diagonal entries of A .*

Proof. You can do this one yourself! Try drawing an arbitrary 3×3 upper triangular matrix (with entries labeled a_{ij}), then compute the determinant by doing a cofactor expansion along the first column or third row. \square

A nice (but sort of obvious) corollary of this is the following.

Corollary 5.2.3. $\det(I_n) = 1$

We now end this chapter with arguably the most useful and important theorems concerning determinants.

Theorem 5.2.4. *Let A be an $n \times n$ matrix. Then A is invertible if and only if $\det(A) \neq 0$.*

Proof. We know that there exists a sequence of row operations taking A to B , where B is in reduced echelon form. This means that every row of B contains a pivot, or the main diagonal has at least one 0 entry. Since B is triangular and $A \sim B$, we know that $\det(A) = c \det(B)$ for some non-zero scalar c . Based on both possibilities for the diagonal entries, we can conclude that if A was invertible, then every column of B is a pivot column so the product of the diagonal entries must be non-zero. If there is a non-pivot column, then there must be a zero entry on the diagonal, hence $\det(B) = 0$. \square

The second useful fact concerns the determinant of a product.

Proposition 5.2.5. *If A and B are $n \times n$ matrices, then*

$$\det(AB) = \det(A) \det(B)$$

The third useful fact, is that when A is invertible, we have a nice explicit form for the determinant of A^{-1} .

Proposition 5.2.6. *If A is invertible, then*

$$\det(A^{-1}) = \frac{1}{\det(A)}$$

Proof. Invertibility of A implies that A^{-1} exists and satisfies the equation $AA^{-1} = I_n$. Taking determinants of both sides and using properties of determinants (which ones?) we conclude that

$$\det(AA^{-1}) = \det(I_n) \implies \det(A) \det(A^{-1}) = 1 \implies \det(A^{-1}) = \frac{1}{\det(A)}$$

\square

Before ending the chapter with an addition to the big theorem, we add several interesting notes on how the determinant relates to geometry and area.

Proposition 5.2.7. *Let S denote the unit square in \mathbb{R}^2 and let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear map with associated matrix A . If $P = T(S)$ denotes the image of the unit square under T , then we have $\text{Area}(P) = |\det(A)|$.*

This means that if $\det(A) = 1$, the associated linear transformation preserves area. An example of this is rotation. Building on this, we have a similar result in higher dimensions.

Proposition 5.2.8. *Let D be a region of finite volume in \mathbb{R}^n and suppose $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a linear map with associated matrix A . If $T(D)$ denotes the image of D under T , then $\text{Volume}(T(D)) = |\det(A)| \cdot \text{Volume}(D)$.*

We now end with an updated (and very powerful!) big theorem.

Theorem 5.2.9. Let $S = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ be a set of vectors in \mathbb{R}^n and let $A = [\mathbf{u}_1 \ \cdots \ \mathbf{u}_n]$ with associated linear transformation given by $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$. The following statements are equivalent:

1. S spans \mathbb{R}^n .
2. S is linearly independent.
3. The system $A\mathbf{x} = \mathbf{b}$ has a solution for every $\mathbf{b} \in \mathbb{R}^n$.
4. T is onto.
5. T is one-to-one.
6. A is invertible.
7. $\ker(T) = \{\mathbf{0}\}$.
8. S is a basis for \mathbb{R}^n .
9. $\text{rank}(A) = n$.
10. $\text{Nullity}(A) = 0$.
11. $\det(A) \neq 0$.

Chapter 6

Eigenvalues and Diagonalization

All of our hard work thus far will finally pay off in this chapter. Much of linear algebra past this point is centered around the idea of eigenvalues and eigenvectors and it is certainly something you will want to remember for future classes in any stem field..

6.1 Eigenvalues and Eigenvectors

Let's quickly recall the basics of the geometry of linear transformations. Given a linear map $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by $T(\mathbf{x}) = A\mathbf{x}$, for some 2×2 matrix A , we can plug in any vector $\mathbf{x} \in \mathbb{R}^2$ and T will output a new vector $A\mathbf{x}$ with a (potentially different) direction and length. The idea of eigenvalues and eigenvectors investigates when the direction and/or length of the output vector is related to the input vector.

Definition 6.1.1. Let A be an $n \times n$ matrix. If \mathbf{u} is a non-zero vector and $\lambda \in \mathbb{R}$ is a scalar such that $A\mathbf{u} = \lambda\mathbf{u}$, then λ is an **eigenvalue** of A and \mathbf{u} is an **eigenvector** of A associated with eigenvalue λ .

There are a few fundamental facts concerning eigenvectors that will allow us to gain extra structure on the set of all eigenvectors associated to some fixed eigenvalue. The first is that the sum of two eigenvectors associated to the same eigenvalue is another (different!) eigenvector associated to the same eigenvalue (you should verify this for yourself). We also have a related result.

Proposition 6.1.2. Suppose A is a square matrix and λ is an eigenvalue of A with associated eigenvector \mathbf{u} , that is, $A\mathbf{u} = \lambda\mathbf{u}$. Then for any non-zero scalar c , we have that $c\mathbf{u}$ is an eigenvector of A associated to λ .

Proof. If $A\mathbf{u} = \lambda\mathbf{u}$ then A being linear implies that for any $c \in \mathbb{R}$

$$A(c\mathbf{u}) = cA\mathbf{u} = c\lambda\mathbf{u} = \lambda(c\mathbf{u})$$

hence $c\mathbf{u}$ is an eigenvector of A associated to eigenvalue λ . □

Combining the last two facts, we obtain the notion of eigenspaces.

Definition 6.1.3. Let A be an $n \times n$ matrix with eigenvalue λ . The set \mathcal{S} consisting of the zero vector and all eigenvectors of A associated with λ forms a subspace of \mathbb{R}^n known as the **eigenspace associated to eigenvalue** λ , often denoted by E_λ .

Example 6.1.4. Let $A = \begin{bmatrix} 6 & -2 \\ 5 & -1 \end{bmatrix}$. One can check that if $\mathbf{u} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\mathbf{v} = \begin{bmatrix} 2 \\ 5 \end{bmatrix}$ then $A\mathbf{u} = 4\mathbf{u}$ and $A\mathbf{v} = \mathbf{v}$. This means that \mathbf{u} is an eigenvector of A of eigenvalue 4 and \mathbf{v} is an eigenvector of A with eigenvalue 1. It follows (by reasons we will soon see) that the eigenspace of eigenvalue 4 is

$$E_4 = \text{Span} \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$$

and the eigenspace of eigenvalue 1 is

$$E_1 = \text{Span} \left\{ \begin{bmatrix} 2 \\ 5 \end{bmatrix} \right\}$$

What we need moving forward is a streamlined way to find eigenvalues and a basis for each associated eigenspace, when given an arbitrary matrix A . What follows is the reasoning behind how we find eigenvalues.

If we have an eigenvalue/eigenvector pair so that $A\mathbf{u} = \lambda\mathbf{u}$ for some vector \mathbf{u} and scalar λ , then we can obtain the closely related equation

$$A\mathbf{u} - \lambda\mathbf{u} = \mathbf{0}$$

By rewriting \mathbf{u} as $I\mathbf{u}$, where I is the $n \times n$ identity matrix, the above equation can be more compactly written as

$$(A - \lambda I)\mathbf{u} = \mathbf{0}$$

Note that $\lambda I = \begin{bmatrix} \lambda & & \\ & \ddots & \\ & & \lambda \end{bmatrix}$ and if $\mathbf{u} = \begin{bmatrix} u_1 \\ \vdots \\ u_n \end{bmatrix}$ then

$$\lambda I\mathbf{u} = \begin{bmatrix} \lambda & & \\ & \ddots & \\ & & \lambda \end{bmatrix} \begin{bmatrix} u_1 \\ \vdots \\ u_n \end{bmatrix} = \begin{bmatrix} \lambda u_1 \\ \vdots \\ \lambda u_n \end{bmatrix} = \lambda\mathbf{u}$$

so the expression $A - \lambda I$ does indeed make sense. With this equation being understood, we can now classify how one finds eigenvalues of a matrix.

Proposition 6.1.5. *Let A be an $n \times n$ matrix. A scalar $\lambda \in R$ is an eigenvalue of A if and only if $\det(A - \lambda I) = 0$.*

Proof. Summarizing what was said above, we have that λ is an eigenvalue of A if and only if $A\mathbf{u} = \lambda\mathbf{u}$ for some vector $\mathbf{u} \neq \mathbf{0}$ if and only if $A\mathbf{u} - \lambda I\mathbf{u} = \mathbf{0}$ if and only if $(A - \lambda I)\mathbf{u} = \mathbf{0}$. This means that λ is an eigenvalue of A if and only if the homogeneous equation $(A - \lambda I)\mathbf{u} = \mathbf{0}$ has a non-trivial solution, and this is true if and only if $A - \lambda I$ is **not** invertible (by the big theorem). It follows that $A - \lambda I$ is **not** invertible if and only if $\det(A - \lambda I) = 0$ which completes the proof. \square

The heart of our method lies in this proof. We will soon see that $\det(A - \lambda I)$ is a polynomial in the variable λ (note that λ is merely a placeholder at first and the values of λ that satisfy $\det(A - \lambda I) = 0$ are the eigenvalues of A). Looking more closely at the polynomial $\det(A - \lambda I) = 0$, we will see that the eigenvalues of A are the roots of this polynomial. The above proposition then takes the task of finding eigenvalues to the task of finding roots of a polynomial. In general we call $\det(A - \lambda I)$ the **characteristic polynomial of A** .

Example 6.1.6. Let $A = \begin{bmatrix} 1 & 1 \\ 2 & 0 \end{bmatrix}$. Find the eigenvalues and a basis for each eigenspace.

We first see that

$$A - \lambda I = \begin{bmatrix} 1 & 1 \\ 2 & 0 \end{bmatrix} - \begin{bmatrix} \lambda & 0 \\ 0 & \lambda \end{bmatrix} = \begin{bmatrix} 1 - \lambda & 1 \\ 2 & -\lambda \end{bmatrix}$$

and

$$\det(A - \lambda I) = (1 - \lambda)(-\lambda) - 2 = \lambda^2 - \lambda - 2 = (\lambda - 2)(\lambda + 1)$$

We need the solutions to $\det(A - \lambda I) = 0$ and these are the solutions to $(\lambda - 2)(\lambda + 1) = 0$ hence $\lambda = 2$ and $\lambda = -1$ are the eigenvalues of A . It is worth noting that no other scalars are eigenvalues of A , these two are the only ones. To find bases for the eigenspaces, we then only need to find the vectors \mathbf{u} and \mathbf{v} respectively,

that satisfy $A\mathbf{u} = 2\mathbf{u}$ and $A\mathbf{v} = -\mathbf{v}$.

If $A\mathbf{x} = \lambda\mathbf{x}$ for some eigenvalue λ , then \mathbf{x} satisfies the equation $(A - \lambda I)\mathbf{x} = \mathbf{0}$. In other words, all the eigenvectors with eigenvalue λ are precisely the vectors in $\text{Null}(A - \lambda I)$. This means that the eigenspace for eigenvalue λ is the same thing as $\text{Null}(A - \lambda I)$. That is

$$E_\lambda = \text{Null}(A - \lambda I)$$

We can now find a basis for E_{-1} . We need to find a basis for $\text{Null}(A - \lambda I)$ with $\lambda = -1$ so we plug in $\lambda = -1$ to $A - \lambda I$ and we get

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

We see that all vectors in the null space of this matrix are of the form $t \begin{bmatrix} -1 \\ 2 \end{bmatrix}$ for some free variable t , hence the basis for E_{-1} is $\left\{ \begin{bmatrix} -1 \\ 2 \end{bmatrix} \right\}$. We leave the computation of a basis for E_2 to the reader as practice. The answer you should get is

$$\mathcal{B}_{E_2} = \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$$

Next, we outline so shortcuts that can be used in finding eigenvalues of simple types of matrices.

Example 6.1.7. Find the eigenvalues of the triangular matrix

$$A = \begin{bmatrix} 1 & 2 & 0 & 1 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \end{bmatrix}$$

By computing the characteristic polynomial of A we see that

$$\text{Det}(A - \lambda I) = \det \left(\begin{bmatrix} 1 - \lambda & 2 & 0 & 1 \\ 0 & 1 - \lambda & 0 & 1 \\ 0 & 0 & 2 - \lambda & 0 \\ 0 & 0 & 0 & 2 - \lambda \end{bmatrix} \right)$$

Recalling that the determinant of a triangular matrix is the product of the diagonal entries, it follows that

$$\det(A - \lambda I) = (1 - \lambda)^2(2 - \lambda)^2$$

Looking back at the matrix A , we can see that the eigenvalues of A are exactly the diagonal entries. This is in fact true for eigenvalues of all triangular matrices.

Next, let's find bases for the eigenspaces E_1 and E_2 , using some shortcuts along the way.

To compute a basis for E_1 , we need to find a basis for $\text{Null}(A - I)$ which is

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

We can see that this matrix has three pivot columns hence $\text{rank}(A - I) = 3$. By rank-nullity this means its null space is 1 dimensional, hence is of the form $\text{Span}\{\mathbf{x}\}$ for some non-zero vector $\mathbf{x} \in \mathbb{R}^4$. Since the first

column of $A - I$ is the zero vector, this means that $A - I$ sends \mathbf{e}_1 to $\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$ hence

$$\mathcal{B}_{E_1} = \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right\}$$

As an exercise, we leave the computation of a basis for E_2 to the reader, but to check your work you should get

$$\mathcal{B}_2 = \left\{ \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 3 \\ 1 \\ 0 \\ 1 \end{bmatrix} \right\}$$

No tricks to doing this one, just compute the correct null space in the usual way.

Looking back at the example above, we can see that one of the eigenspaces was one-dimensional while the other one was two dimensional. This is a phenomenon that is subtle and requires a bit more discussion.

The first thing to note is that if A is an $n \times n$ matrix, then the characteristic polynomial $\det(A - \lambda I)$ is always a degree n polynomial. The concept we will need to understand further is that of (algebraic) multiplicity of a root of a polynomial, which we now define.

Definition 6.1.8. Let $P(x)$ denote a polynomial of degree n , in the variable x , and suppose $P(x)$ is the characteristic polynomial of some matrix A . If we can factor this polynomial as

$$P(x) = (x - \alpha)^m Q(x)$$

where $Q(\alpha) \neq 0$, then we say $x = \alpha$ is an eigenvalue of A with **multiplicity m**. In other words, the exponent attached to the linear term of a polynomial is the multiplicity we associate to the root of that polynomial that comes from the given linear term.

This notion of multiplicity is precisely what we need to say more about dimensions of eigenspaces.

Theorem 6.1.9. Let λ be an eigenvalue of a matrix A and let $m(\lambda)$ denote the multiplicity of the eigenvalue λ . Then we always have

$$\dim E_\lambda \leq m(\lambda)$$

That is, the dimension of the eigenspace for eigenvalue λ never exceeds the multiplicity of that eigenvalue.

Looking back at the previous example, we can see that both eigenvalues 1 and 2 have multiplicity 2, yet $\dim E_1 = 1$ and $\dim E_2 = 2$. The inequality holds in both cases but we only obtained equality in one. There is lots more that one can say about multiplicities of eigenvalues but we leave it at this for now, and say a bit more in the next section. We now end this section with one more important fact, which will be our last addition to the big theorem.

Proposition 6.1.10. $\lambda = 0$ is not an eigenvalue of A if and only if $\det(A) \neq 0$.

Proof. We show that $\lambda = 0$ is an eigenvalue of A if and only if $\det(A) = 0$. We can see that $\lambda = 0$ is an eigenvalue of A if and only if $\det(A - \lambda I) = \det(A - 0I) = \det(A) = 0$, which is all we needed to show. \square

Since we can only discuss eigenvalues for square matrices, this proposition can extend our list of results coming from the big theorem.

Theorem 6.1.11. *Let $S = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ be a set of vectors in \mathbb{R}^n and let $A = [\mathbf{u}_1 \ \cdots \ \mathbf{u}_n]$ with associated linear transformation given by $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$. The following statements are equivalent:*

1. S spans \mathbb{R}^n .
2. S is linearly independent.
3. The system $A\mathbf{x} = \mathbf{b}$ has a solution for every $\mathbf{b} \in \mathbb{R}^n$.
4. T is onto.
5. T is one-to-one.
6. A is invertible.
7. $\ker(T) = \{\mathbf{0}\}$.
8. S is a basis for \mathbb{R}^n .
9. $\text{rank}(A) = n$.
10. $\text{Nullity}(A) = 0$.
11. $\det(A) \neq 0$.
12. $\lambda = 0$ is not an eigenvalue of A .

6.2 Diagonalization

Let's jump right in.

Definition 6.2.1. An $n \times n$ matrix is **diagonalizable** if there exists $n \times n$ matrices Λ and X such that

- Λ is diagonal.
- X is invertible.

and

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

Note that Λ is the greek capital letter for λ . This is intentional, and we will soon see that the diagonal entries of Λ are precisely the eigenvalues of A .

Example 6.2.2. If $X = \begin{bmatrix} 1 & 0 \\ 2 & 1 \end{bmatrix}$ and $\Lambda = \begin{bmatrix} 4 & 0 \\ 0 & -3 \end{bmatrix}$, then $X^{-1} = \begin{bmatrix} 1 & 0 \\ -2 & 1 \end{bmatrix}$ and $X\Lambda X^{-1} = \begin{bmatrix} 4 & 0 \\ 14 & -3 \end{bmatrix}$. If $A = \begin{bmatrix} 4 & 0 \\ 14 & -3 \end{bmatrix}$ then we say A is diagonalizable.

This example doesn't help much. In general, we need to find a way to construct the matrices X and Λ that diagonalize A , and in doing so, we will see when a given matrix is not diagonalizable. Before embarking on this adventure, it is worth noting one of the many reasons why diagonalization is useful. In many applied fields, systems can be modeled by matrix multiplication and iterates of our system can be taken via computing powers of a matrix. If the given matrix is diagonalizable, computing powers can be very easy.

Assume that A is diagonalizable so that we can write $A = X\Lambda X^{-1}$, then

$$A^2 = (X\Lambda X^{-1})(X\Lambda X^{-1}) = X\Lambda^2 X^{-1}$$

and

$$A^3 = A^2 A = (X\Lambda^2 X^{-1})(X\Lambda X^{-1}) = X\Lambda^3 X^{-1}$$

Continuing this process we can see that

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

and since Λ is a diagonal matrix, computing powers of it is excessively easy. We now begin the investigation of when A is diagonalizable by stating the main result and digging into the details.

Theorem 6.2.3. *An $n \times n$ matrix A is diagonalizable if and only if A has eigenvectors that form a basis for \mathbb{R}^n .*

There are a few important things to point out regarding what we mean when we say eigenvectors.

Proposition 6.2.4. *If λ_1 and λ_2 are eigenvalues of a matrix A and $\lambda_1 \neq \lambda_2$, then if $\mathbf{x} \in E_{\lambda_1}$ and $\mathbf{y} \in E_{\lambda_2}$ it is always true that $\{\mathbf{x}, \mathbf{y}\}$ form a linearly independent set. That is, eigenvectors corresponding to different eigenvalues are always linearly independent.*

This means that when given a square matrix A we can

1. Find all the eigenvalues of A (by finding roots of $\det(A - \lambda I) = 0$).
2. Find bases for all eigenspaces.
3. Put all basis vectors from different eigenspaces in a set and see if this set forms a basis for \mathbb{R}^n . By way of the above proposition, we know that the eigenvectors will form a basis (called an eigenbasis) if there are n of them.

We now prove Theorem 6.2.3.

Proof. Let $\mathbf{u}_1, \dots, \mathbf{u}_n$ be n (linearly independent) eigenvectors for a matrix A , with eigenvalues labeled $\lambda_1, \dots, \lambda_n$ (note here that we are assuming there are n distinct eigenvalues for simplicity of the proof but this is not always the case), so that $\mathcal{B} = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ forms an (eigen)basis for \mathbb{R}^n . Let $X = [\mathbf{u}_1 \ \dots \ \mathbf{u}_n]$ and

$$\Lambda = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix}$$

be the diagonal matrix with eigenvalues on the diagonal. Since \mathcal{B} is a basis for \mathbb{R}^n we know that X is invertible, by the big theorem. Looking at the matrix multiplication, we see that

$$AX = A[\mathbf{u}_1 \ \dots \ \mathbf{u}_n] = [A\mathbf{u}_1 \ \dots \ A\mathbf{u}_n] = [\lambda_1\mathbf{u}_1 \ \dots \ \lambda_n\mathbf{u}_n] = [\mathbf{u}_1 \ \dots \ \mathbf{u}_n] \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} = X\Lambda$$

Since $AX = X\Lambda$ we can conclude that $A = X\Lambda X^{-1}$ and A is diagonalizable. □

Example 6.2.5. Let $A = \begin{bmatrix} 4 & -2 \\ 4 & -2 \end{bmatrix}$ and show that A is diagonalizable by finding matrices X and Λ such that $A = X\Lambda X^{-1}$.

We begin by finding the eigenvalues of A as well as bases for the eigenspaces. By computing the characteristic polynomial for A we see that

$$\det(A - \lambda I) = -(4 - \lambda)(2 + \lambda) + 8 = -(8 + 2\lambda + \lambda^2) + 8 = \lambda^2 - 2\lambda = \lambda(\lambda - 2) = 0$$

This means that $\lambda = 0, 2$ are the eigenvalues of A . In computing bases for both eigenspaces, we just need to find bases for $\text{Null}(A)$ and $\text{Null}(A - 2I)$ respectively. We get that

$$\text{Null}(A) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 2 \end{bmatrix} \right\} \quad \text{and} \quad \text{Null}(A - 2I) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$$

Following the proof of Theorem 6.2.3, we set $X = \begin{bmatrix} 1 & 1 \\ 2 & 1 \end{bmatrix}$ and $\Lambda = \begin{bmatrix} 0 & 0 \\ 0 & 2 \end{bmatrix}$ and these are precisely the matrices that diagonalize A . Note the the order in which we place the vectors is very important, if we swapped the order of 0 and 2, on the diagonal of Λ , while leaving the columns of X unchanged, the resulting matrix product would not equal A . You must always have the columns of X correspond, in the same order, with the eigenvalues for those column vectors. Moreover, if you have an eigenvalue of multiplicity k , then there will be exactly k diagonal entries of Λ that are equal to that given eigenvalue.

Example 6.2.6. Construct a 3×3 matrix A with the following eigenvalues and eigenvectors.

$$\lambda_1 = 2, \mathbf{u}_1 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \quad \lambda_2 = -1, \mathbf{u}_2 = \begin{bmatrix} -1 \\ 0 \\ 2 \end{bmatrix} \quad \lambda_3 = 5, \mathbf{u}_3 = \begin{bmatrix} 4 \\ 4 \\ 2 \end{bmatrix}$$

This can quickly be done by working backwards through the mechanics of the proof of Theorem 6.2.3. Let

$$X = \begin{bmatrix} 0 & -1 & 4 \\ 1 & 0 & 4 \\ 0 & 2 & 2 \end{bmatrix} \quad \text{and} \quad \Lambda = \begin{bmatrix} 2 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 5 \end{bmatrix}$$

By computing X^{-1} (which we know exists), the resulting matrix $X\Lambda X^{-1}$ will have the prescribed eigenvalues and eigenvectors.

We now use the full strength of Theorem 6.2.3.

Example 6.2.7. Is $A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ diagonalizable?

It suffices to see if there exists a basis of eigenvectors of A . We have

$$\det(A - \lambda I) = \det \left(\begin{bmatrix} 1 - \lambda & 1 \\ 0 & 1 - \lambda \end{bmatrix} \right) = (1 - \lambda)^2$$

hence 1 is the only eigenvalue of A , with multiplicity 2. Since we need a basis of eigenvectors in order to diagonalize A , we must have the dimension of E_1 be equal to 2. If it is not, then there is no way for our eigenvectors to form a basis for \mathbb{R}^2 , since we will not have enough of them. In computing the eigenspace we see that

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

If we wanted, we could stop right here since the null space of this matrix can never be 2 dimensional, because it has rank 1 (rank-nullity is being used here). If we want to be more explicit, we can directly compute $E_1 = \text{Null}(A - I)$ and find that

$$\text{Null}(A - I) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right\}$$

Regardless of the argument we prefer, we can now see that A is not diagonalizable.

Shortcuts like this one can be very helpful in practice. There is one shortcut in particular that can really come in handy.

Proposition 6.2.8. *Let A be an $n \times n$ matrix and assume that $\{\lambda_1, \dots, \lambda_n\}$ are **distinct** eigenvalues (the word *distinct* here means that there are exactly n eigenvalues, no two of which are equal), then A is always diagonalizable.*

Proof. If A has n distinct eigenvalues, then the characteristic polynomial of A has exactly n distinct roots. That is,

$$\det(A - \lambda I) = (\lambda - \lambda_1)(\lambda - \lambda_2) \cdots (\lambda - \lambda_n)$$

This means that the multiplicity of each eigenvalue is 1. Now, recalling that $\dim(E_\lambda) \leq m(\lambda)$ for all eigenvalues, we have that $\dim(E_{\lambda_i}) \leq m(\lambda_i) = 1$ for all $i = 1, \dots, n$, hence we must have $\dim(E_{\lambda_i}) = 1$ for all i , because eigenspaces of actual eigenvalues of a matrix are never 0 dimensional (in fact, the only instance when $E_\lambda = \{\mathbf{0}\}$ for some matrix A is when λ is **not** an eigenvalue of A). This means that we get exactly one (linearly independent) eigenvector coming from each eigenspace, of which there are n in total. Putting them all together in one set, we obtain a set of n linearly independent vectors in \mathbb{R}^n , which (by the theorem) forms our desired eigenbasis. This implies that A is diagonalizable. \square

Example 6.2.9. Is $A = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 2 & 1 \\ 0 & 0 & 3 \end{bmatrix}$ diagonalizable?

Recalling our nice little trick for triangular matrices, we can see that the eigenvalues are 1, 2, and 3 respectively, which are distinct! This means that A is diagonalizable, by the above proposition.

Although this proposition is great, we still need to treat it with care. In particular, not all diagonalizable matrices have distinct eigenvalues.

For an easy example, one should note that the zero matrix is diagonalizable. Moreover, if we consider the identity matrix $I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, then we can write I as

$$I = III^{-1}$$

which (in a silly way) satisfies the definition of being diagonalizable. Moreover, we could compute the null space of $I - I$, and see that it admits a basis of eigenvectors (in particular it admits the standard basis).

Before ending this chapter, we provide one last alternative way to check if a matrix is diagonalizable. One can think of this as a generalization of the proposition on distinct eigenvalues.

Proposition 6.2.10. *Suppose A is an $n \times n$ matrix with **only real eigenvalues** (so none of them are complex numbers). A is diagonalizable if and only if the dimension of each eigenspace is equal to the multiplicity of the corresponding eigenvalue.*

In general, one calls the dimension of E_λ the *geometric* multiplicity of λ whereas the usual multiplicity $m(\lambda)$ is known as the *algebraic* multiplicity. This proposition is saying that the geometric multiplicity is always less than or equal to the algebraic multiplicity, and when they are equal, the given matrix is diagonalizable.

Proof. An $n \times n$ matrix A is diagonalizable if and only if it admits n linearly independent eigenvectors (by Theorem 6.2.3). Moreover, each eigenspace has dimension no greater than the multiplicity of the associated eigenvalue, i.e. $\dim(E_\lambda) \leq m(\lambda)$. Since A is an $n \times n$ matrix, the sum of the multiplicities of the eigenvalues must equal n , because $\det(A - \lambda I)$ is a degree n polynomial. Lastly, since the eigenvectors coming from different eigenspaces are always linearly independent, we can conclude that the sum of the dimensions of all eigenspaces equals n . Counting up one basis vector for each dimension, we end up with exactly n eigenvectors, hence a basis for \mathbb{R}^n , completing the proof. \square

This marks the end of the introductory material. We resume with a "review" of eigenvalues, introducing a new topic in the next chapter.

Chapter 7

Review of Eigenvalues and Diagonalization

To begin our review, let's lay out some possibly new terminology and notation. When we say the word “vector space” V , we mean a set of “vectors”, such that if $\mathbf{u}, \mathbf{v} \in V$, then $\mathbf{u} + \mathbf{v} \in V$, and if $r \in \mathbb{R}$ and $\mathbf{u} \in V$ then $r\mathbf{u} \in V$. We also always have the zero vector in V , which is also something that should be shown in practice. When we hear the word vector space we should just think to ourselves, “subspace” just like we learned in our first linear algebra course.

We will see that what we take as our “vectors” can vary greatly. We have gotten quite used to one vector space, namely \mathbb{R}^n , but we are about to see another.

7.1 Eigenvalues and Eigenvectors

Definition 7.1.1. Define $\mathbb{R}^{m \times n}$ to be the vector space of $m \times n$ matrices with real number entries. That is, using compact matrix notation and writing $A = (a_{ij}) = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ \vdots & \vdots & \dots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix}$ we have

$$\mathbb{R}^{m \times n} = \left\{ A = (a_{ij}) : a_{ij} \in \mathbb{R} \quad \forall i = 1, 2, \dots, m, \quad \forall j = 1, 2, \dots, n \right\}$$

Recalling that $m \times n$ matrices represent linear transformations from \mathbb{R}^n to \mathbb{R}^m , given some $A \in \mathbb{R}^{m \times n}$ we have $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ given by $\mathbf{x} \mapsto A\mathbf{x}$.

Now onto the definition of eigenvalues, eigenvectors, and eigenspaces. Note that the definition only makes sense for square matrices so for the remainder of this section, any matrix A is assumed to be in the vector space $\mathbb{R}^{n \times n}$.

Definition 7.1.2. Let $A \in \mathbb{R}^{n \times n}$. If there exists a non-zero vector \mathbf{x} and some scalar $\lambda \in \mathbb{R}$ such that $A\mathbf{x} = \lambda\mathbf{x}$ we say that \mathbf{x} is an **eigenvector of A with eigenvalue λ** .

Along with eigenvalues come subspaces which leads us to the notion of an eigenspace.

Definition 7.1.3. The **eigenspace of eigenvalue λ** is defined to be $E_\lambda = \text{Null}(A - \lambda I)$. It is the vector space (subspace) consisting of all eigenvectors with eigenvalue λ .

From this definition we may recall that eigenvalues are the roots of a polynomial that relates to the matrix A . This is called the characteristic polynomial and is given by $\det(A - \lambda I)$. This gives the following crucial proposition which is our main tool for finding eigenvalues in practice.

Proposition 7.1.4. λ is an eigenvalue of A if and only if $\det(A - \lambda I) = 0$

Proof. See online notes for a proof □

Example 7.1.5. Let $A \in \mathbb{R}^{2 \times 2}$ be the linear map that reflects the plane about the line $y = x$. To find the matrix of this we need to recall what is arguably the most useful fact from our first course in linear algebra.

Matrix of a linear transformation

Given a linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$, the matrix for T can be written as

$$[T(\mathbf{e}_1) \quad T(\mathbf{e}_2) \quad \cdots \quad T(\mathbf{e}_n)]$$

This means any linear transformation is completely determined by where it takes a basis.

Using this fact, we find that the matrix for the linear transformation mentioned above is $A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$.

We find the eigenvalues and bases for eigenspaces in three steps.

Step 1: $\det(A - \lambda I) = \lambda^2 - 1$

Step 2: Solving $\det(A - \lambda I) = \lambda^2 - 1 = 0$ we find that the eigenvalues are $\lambda = \pm 1$

Step 3: For each λ , we find a basis of $E_\lambda = \text{Null}(A - \lambda I)$

$$\lambda = 1 \implies A - I = \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix} \sim \begin{bmatrix} -1 & 1 \\ 0 & 0 \end{bmatrix} \text{ which gives the linear system } -x_1 + x_2 = 0 \implies x_1 = x_2$$

This implies that $E_1 = \text{Span} \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$.

$$\lambda = -1 \implies A + I = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \sim \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} \text{ which gives the linear system } x_1 + x_2 = 0 \implies x_1 = -x_2$$

This implies that $E_{-1} = \text{Span} \left\{ \begin{bmatrix} -1 \\ 1 \end{bmatrix} \right\}$.

Along with the notion of an eigenvalue comes its multiplicity. When we first saw multiplicities we defined it to be an exponent that was tied to the characteristic polynomial. This was in fact just one type of multiplicity, and there is one more, which we will soon see does not always agree with our original definition. The original definition we had is what was known as the *algebraic multiplicity*. Before seeing the definition, we recall that for any $A \in \mathbb{R}^{n \times n}$, the characteristic polynomial $\det(A - \lambda I)$ is a degree n polynomial. This will good to keep in the back of our minds.

Definition 7.1.6. Denote the characteristic polynomial by $p(\lambda)$. We can factor p as

$$\det(A - \lambda I) = p(\lambda) = (\lambda - \lambda_1)^{\alpha_1} (\lambda - \lambda_2)^{\alpha_2} \cdots (\lambda - \lambda_k)^{\alpha_k}$$

where $\alpha_1 + \alpha_2 + \cdots + \alpha_k = n$ and $\lambda_1, \dots, \lambda_k$ are the eigenvalues. We say that the **algebraic multiplicity** of the eigenvalue λ_i is the exponent α_i . We denote this value by $\text{AM}(\lambda_i)$.

Note that from the fact that the degree of $\det(A - \lambda I)$ equals n , we can conclude that the sum of the algebraic multiplicities equals n .

We now come across our first new definition of the course.

Definition 7.1.7. The **geometric multiplicity** of the eigenvalue λ_i , denoted $\text{GM}(\lambda_i)$, is the dimension of its corresponding eigenspace. That is

$$\text{GM}(\lambda_i) = \dim(E_{\lambda_i})$$

One may be inclined to ask when these two quantities are equal.

Example 7.1.8. Let $A = \begin{bmatrix} 8 & -9 \\ 4 & -4 \end{bmatrix}$. A quick computation of the characteristic polynomial shows that $p(\lambda) = (\lambda - 2)^2$ so we can conclude that $AM(2) = 2$. To find the geometric multiplicity, we do a little more matrix algebra.

$$A - \lambda I = \begin{bmatrix} 8 - \lambda & -9 \\ 4 & -4 - \lambda \end{bmatrix} \implies A - 2I = \begin{bmatrix} 6 & -9 \\ 4 & -6 \end{bmatrix}$$

We want to find $\dim(E_2) = \dim(\text{Null}(A - 2I))$. The columns of this matrix are linearly dependent so we know that it has nonzero nullity, but the only 2×2 matrix with nullity equal to 2 is the zero matrix. From this we can conclude that $\dim(E_2) = 1$ hence $GM(2) < AM(2)$.

This example gives us some sense of how these two values are related. The following result will be referred to as the AM-GM inequality.

Proposition 7.1.9. AM-GM inequality

Given $A \in \mathbb{R}^{n \times n}$ with eigenvalues λ_i for $i = 1, 2, \dots, k$ we always have $GM(\lambda_i) \leq AM(\lambda_i) \quad \forall i$.

Question 7.1.10. For what matrices are these values equal?

Great question! The answer is that the values are equal precisely when A is **diagonalizable**. This answer is precisely the content of the next section, but first, we state several more general facts about eigenvalues.

1. If λ is an eigenvalue of A then λ^k is an eigenvalue of A^k .

Why?

$$A\mathbf{x} = \lambda\mathbf{x} \implies A^2\mathbf{x} = A(A\mathbf{x}) = A\lambda\mathbf{x} = \lambda A\mathbf{x} = \lambda^2\mathbf{x} \implies A^3\mathbf{x} = A(A^2\mathbf{x}) = A\lambda^2\mathbf{x} = \lambda^2 A\mathbf{x} = \lambda^3\mathbf{x}$$

Continuing with this process we can see that

$$A^k\mathbf{x} = A(A^{k-1}\mathbf{x}) = A\lambda^{k-1}\mathbf{x} = \lambda^{k-1}A\mathbf{x} = \lambda^k\mathbf{x}$$

2. If λ is an eigenvalue of A and A is invertible, then $\lambda^{-1} = \frac{1}{\lambda}$ is an eigenvalue of A^{-1} .

Why? If $A\mathbf{x} = \lambda\mathbf{x}$, we can multiply both sides of this equation on the left by A^{-1} (which we know exists!). We then get $\mathbf{x} = A^{-1}\lambda\mathbf{x}$ which implies that $\frac{1}{\lambda}\mathbf{x} = A^{-1}\mathbf{x}$. Note that the eigenvalues change but the eigenvectors don't.

3. $A \in \mathbb{R}^{n \times n}$ has n eigenvalues (counting (algebraic) multiplicities).

Why? To see why we need a “fundamental” theorem that we should never forget.

Fundamental Theorem of Algebra

Let $p(x)$ be a degree n polynomial. Then n has (counting multiplicities), precisely n complex roots. In the language of characteristic polynomials, this means we can factor the characteristic polynomial of any square matrix as

$$\det(A - \lambda I) = (\lambda - \lambda_1)(\lambda - \lambda_2) \cdots (\lambda - \lambda_n)$$

Some example of this include

$$x^2 - 1 = 0 \implies x = \pm 1 \implies x^2 - 1 = (x - 1)(x + 1) = 0$$

or in the case of complex roots

$$x^2 + 1 = 0 \implies x = \pm i \implies x^2 + 1 = (x - i)(x + i) = 0$$

Notice that the coefficients of this polynomial were real numbers, not complex numbers. This hints at the fact that matrices with real number entries can have complex eigenvalues.

7.1.1 Complex Eigenvalues

It will be important to keep in mind that our set of possible eigenvalues can leave the set of real numbers and be complex. We have no easy way to get a geometric picture from this, but nonetheless it is a case we must consider. Rotation matrices are some of most classical examples of real matrices with complex eigenvalues.

Example 7.1.11. Let A denote the linear transformation that rotates any vector \mathbb{R}^2 by $\pi/2$. Recall that rotation matrices in \mathbb{R}^2 have a generic form

$$R_\theta = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

From this we can quickly see that

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

We compute its characteristic polynomial and see that

$$\det \begin{bmatrix} -\lambda & -1 \\ 1 & -\lambda \end{bmatrix} = \lambda^2 + 1 = 0 \implies \lambda = \pm i$$

If we were to go further and compute the eigenvectors we would see that an eigenvector of eigenvalue i is $\begin{bmatrix} i \\ 1 \end{bmatrix}$ and an eigenvector of eigenvalue $-i$ is $\begin{bmatrix} 1 \\ i \end{bmatrix}$. As mentioned, geometry is hopeless here, but after we develop some more machinery we will be able to make some geometric sense of this.

Example 7.1.12. If B denoted rotation by π rather than $\pi/2$, we could obtain B from A by recalling that the product of matrices corresponds to the composition of linear transformations. Since rotation by π is the same as rotation by $\pi/2$ twice, we would have that $B = A^2$. Geometrically, every vector will flip its sign, which is the same as saying that it has eigenvalue -1 . Algebraically, a quick computation shows that

$$B = A^2 = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}$$

hence

$$\det \begin{bmatrix} -1 - \lambda & 0 \\ 0 & -1 - \lambda \end{bmatrix} = (-1 - \lambda)^2 = (\lambda + 1)^2 \implies \lambda = -1 \text{ with } \text{AM}(-1) = 2$$

. This is one of the few rotation matrices that has real eigenvalues. To quickly tie this into the notion of geometric multiplicity, we observe that

$$E_{-1} = \text{Null}(A^2 + I) = \text{Null}(-I + I) = \text{Null}(O_{22}) = \mathbb{R}^2$$

so here $\text{GM}(-1) = \text{AM}(-1) = 2$.

7.2 Diagonalization

Definition 7.2.1. A matrix $A \in \mathbb{R}^{n \times n}$ is **diagonalizable** if there exists an invertible matrix X and a diagonal matrix Λ (this is a capital lambda) such that

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

. Recall that $\Lambda = \text{diag}((\lambda_1, \lambda_2, \dots, \lambda_n))$, the diagonal matrix with the eigenvalues of A on the diagonal. Moreover, the columns of X are the eigenbasis vectors, in order.

The following theorem is an important one, and we can take it as an equivalent definition of diagonalizability.

Theorem 7.2.2. *If $A \in \mathbb{R}^{n \times n}$ has n linearly independent eigenvectors, then A is diagonalizable.*

Proof. [see online notes for a proof](#) □

As an exercise, you should try to convince yourself of the following fact, which will prove (depending on your preference) to be the most useful equivalent definition of diagonalizability.

Theorem 7.2.3. *$A \in \mathbb{R}^{n \times n}$ is diagonalizable if and only if $\text{AM}(\lambda) = \text{GM}(\lambda)$ for all eigenvalues λ .*

Proof. [see online notes for a proof](#) □

We have already seen some examples of diagonalizable matrices, namely Example 7.1.12. Here is one more:

Example 7.2.4. $A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. The eigenvalues of this matrix are $\lambda = \pm 1$ with respective eigenspaces given by $E_1 = \text{Span} \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$ and $E_{-1} = \text{Span} \left\{ \begin{bmatrix} -1 \\ 1 \end{bmatrix} \right\}$. Using the basis vectors for our eigenspaces we obtain the eigenbasis of \mathbb{R}^2 given by $\left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} -1 \\ 1 \end{bmatrix} \right\}$. Since we have two linearly independent eigenvectors, we know A is diagonalizable with diagonalization given by

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

In the first linear algebra course, we saw that diagonalization made computing powers of a matrix quite easy. If $A = X\Lambda X^{-1}$ then

$$A^k = \underbrace{(X\Lambda X^{-1})(X\Lambda X^{-1}) \cdots (X\Lambda X^{-1})}_{k \text{ times}} = X\Lambda^k X^{-1}$$

Recalling basic facts about products of diagonal matrices we can also conclude that $\Lambda = \text{diag}(\lambda_1^k, \lambda_2^k, \dots, \lambda_n^k)$.

Continuing with our example above, we can see that

$$A^{53} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}^{53} = \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}^{53} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

In general, it may take some work to see if a matrix is diagonalizable or not but in certain cases, we can deduce the result fairly quickly.

Theorem 7.2.5. *If $A \in \mathbb{R}^{n \times n}$ has n **distinct** eigenvalues, then A is diagonalizable.*

Proof. First note that if λ_i is an eigenvalue of A it **must** have some non-zero eigenvector \mathbf{x} such that $A\mathbf{x} = \lambda_i\mathbf{x}$. Moreover, for any $\mathbf{x} \in E_{\lambda_i}$, we know that $c\mathbf{x} \in E_{\lambda_i}$ for any $c \in \mathbb{R}$, thus if $\mathbf{x} \in E_{\lambda_i}$ then $\text{Span}\{\mathbf{x}\} \subseteq E_{\lambda_i}$. This means that for each eigenvalue λ_i , we have $\text{GM}(\lambda_i) > 0$.

If A has n distinct eigenvalues, then we can list them out $\lambda_1, \lambda_2, \dots, \lambda_n$ and are guaranteed that $\lambda_i \neq \lambda_j$ for all $i \neq j$. This means that the characteristic polynomial looks like

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

hence $\text{AM}(\lambda_i) = 1 \forall i$. From proposition 7.1.9 we can conclude that

$$0 < \text{GM}(\lambda_i) \leq \text{AM}(\lambda_i) = 1$$

hence $\text{AM}(\lambda_i) = \text{GM}(\lambda_i) \forall i$. The result now follows from 7.2.3. □

Warning: Not all matrices are diagonalizable! You *must* have n linearly independent eigenvectors so that the matrix X^{-1} exists.

Example 7.2.6. Let $A = \begin{bmatrix} 6 & -1 \\ 1 & 4 \end{bmatrix}$. In finding a basis of eigenvectors we see that

$$\det \begin{bmatrix} 6 - \lambda & -1 \\ 1 & 4 - \lambda \end{bmatrix} = \lambda^2 - 10\lambda + 25 = (\lambda - 5)^2 = 0 \implies \lambda = 5 \text{ with } \text{AM}(5) = 2$$

This tells us that the eigenvalues are not distinct so this may not be diagonalizable. Computing the eigenspace we see that

$$E_5 = \text{Null} \begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix} = \text{Null} \begin{bmatrix} 1 & -1 \\ 0 & 0 \end{bmatrix} = \text{Span} \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$$

so $\text{GM}(5) = 1$ and $\text{AM}(5) \neq \text{GM}(5)$. This means A is not diagonalizable because there's not enough eigenvectors.

We end the section with a quick notational refresher, which we will use **a lot**.

Definition 7.2.7. Given $A = (a_{ij}) \in \mathbb{R}^{m \times n}$, A **transpose**, denoted A^T is the matrix obtained by swapping rows and columns of A . That is $A^T = (a_{ji}) \in \mathbb{R}^{n \times m}$.

Example 7.2.8. If $A = \begin{bmatrix} 1 & -1 \\ 3 & 1 \end{bmatrix}$ then $A^T = \begin{bmatrix} 1 & 3 \\ -1 & 1 \end{bmatrix}$.

We finish off the chapter with a completely new notion.

7.3 Similar Matrices

Definition 7.3.1. Two matrices A and C are **similar**, sometimes denoted $A \sim C$, if there exists an invertible matrix B such that

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

There is one main example that probably comes to mind.

Example 7.3.2. If A is diagonalizable, then $A = X\Lambda X^{-1}$ for some λ and some X , hence $A \sim \Lambda$.

Theorem 7.3.3. If $A \sim C$ then A and C have the same eigenvalues.

Proof. If $A \sim C$ then \exists an invertible matrix B such that $A = BCB^{-1}$. Our goal is to show that A and C have the same eigenvalues. Suppose $C\mathbf{x} = \lambda\mathbf{x}$, then since $AB = BC$ we know that $AB\mathbf{x} = BC\mathbf{x}$, hence

$$A(B\mathbf{x}) = B(C\mathbf{x}) = B\lambda\mathbf{x} = \lambda B\mathbf{x}$$

This implies that λ is an eigenvalue of A (with eigenvector $B\mathbf{x}$), hence all eigenvalues of C are also eigenvalues of A . \square

Chapter 8

Polynomial Vector Spaces

Now that we have some concrete things to think when we hear the word “vector space”, we can introduce some new and less obvious vector spaces.

8.1 $\mathbb{R}[x]$: The polynomial vector space

Before we define $\mathbb{R}[x]$, we define smaller pieces of it.

Definition 8.1.1.

$$\mathbb{R}[x]_n = \{a_n x^n + a_{n-1} x^{n-1} + \cdots + a_1 x + a_0 : a_i \in \mathbb{R}\} = \{\text{polynomials of degree at most } n \text{ with coefficients in } \mathbb{R}\}$$

You should convince yourself that this is indeed a vector space with polynomials of degree at most n as its vectors. With this new definition, we can ask ourselves some natural questions.

Question 8.1.2. What is a basis for $\mathbb{R}[x]_n$?

When we think of a basis of a vector space we think of the fundamental building blocks of that vector space. That is, what are the vectors that **every** vector is a linear combination of? Looking at the definition above, its not too hard to see that every degree n polynomial has a general form, and any given polynomial differs from another by its coefficients. Looking at a polynomial as a linear combination of **monomials** (elements of the form x^k), we can see that every degree at most n polynomial is a linear combination of the monomials $1, x, x^2, \dots, x^n$, hence a basis for $\mathbb{R}[x]_n$ is given by $\{1, x, x^2, \dots, x^n\}$. We call this basis the **monomial basis**.

Question 8.1.3. What is $\dim(\mathbb{R}[x]_n)$?

This question follows pretty quickly from the previous one and we get that $\dim(\mathbb{R}[x]_n) = n + 1$.

We can now define a *larger* vector space.

Definition 8.1.4.

$$\mathbb{R}[x] = \{\cdots + a_n x^n + a_{n-1} x^{n-1} + \cdots + a_1 x + a_0 : a_i \in \mathbb{R}\} = \{\text{All polynomials with coefficients in } \mathbb{R}\}$$

For every positive integer n , this vector space has $\mathbb{R}[x]_n$ as a subspace of it. This vector space is a little harder to understand, as can be seen by asking the same questions.

Question 8.1.5. What is a basis for $\mathbb{R}[x]$?

We can take polynomials of any degree so writing all of these as a linear combination of monomials, we get a monomial basis given by $\{1, x, x^2, \dots\}$

Question 8.1.6. What is $\dim(\mathbb{R}[x])$?

Since dimension is defined to be the size of a basis, we see that $\mathbb{R}[x]$ has infinite dimension! This is likely our first encounter with infinite-dimensional vector spaces. The first thing to notice about them is that we can use nifty theorems like Rank-Nullity, and we have to be a little more tactical in the way that we approach problems.

8.1.1 Linear Maps between Polynomial Vector Spaces

The best way to get a grasp on this is through examples.

Example 8.1.7. (Multiplication by x^2)

$$T : \mathbb{R}[x]_n \rightarrow \mathbb{R}[x]_{n+2} \quad p(x) \mapsto x^2 p(x)$$

For example $T(x^n + x - 5) = x^{n+2} + x^3 - 5x^2$. You should check for yourself that T indeed defines a linear map (satisfies the two main properties) between the indicated domain and codomain vector spaces.

Example 8.1.8. (Integration)

$$T : \mathbb{R}[x] \mapsto \mathbb{R} \quad p(x) \mapsto \int_0^1 p(x) dx$$

For example,

$$T(x^2 - 1) = \int_0^1 x^2 - 1 dx = \left. \frac{x^3}{3} - x \right|_0^1 = \frac{1}{3} - 1 = -\frac{2}{3}$$

ou should check for yourself that T indeed defines a linear map (satisfies the two main properties) between the indicated domain and codomain vector spaces, i.e. integration is linear!

We can use linear maps like this to find interesting subspaces of $\mathbb{R}[x]_n$.

Example 8.1.9. Recall that the kernel of a linear map is always a subspace of the domain. Using the map of Example 8.1.8 we have that

$$S = \{p(x) \in \mathbb{R}[x] : \int_0^1 p(x) = 0\}$$

is a subspace of $\mathbb{R}[x]$ because it is the kernel of the integration map.

Example 8.1.10. (Differentiation)

$$D : \mathbb{R}[x]_n \rightarrow \mathbb{R}[x]_n \quad p(x) \mapsto p'(x)$$

where $p'(x) = \frac{d}{dx}(p(x))$, usual differentiation. For example $D(x^n + x - 5) = nx^{n-1} + 1$. Let's verify that it is linear.

1) $D(p(x) + q(x)) = (p(x) + q(x))' = p'(x) + q'(x) = D(p(x)) + D(q(x))$

2) For $r \in \mathbb{R}$ we have $D(rp(x)) = (rp(x))' = rp'(x) = rD(p(x))$

Now, since every linear map is given by a matrix, lets find the matrix of the linear map $D : \mathbb{R}[x]_4 \rightarrow \mathbb{R}[x]_4$.

Step 1: Recall from 9.1 that we have a formula for the matrix of any linear transformation. This matrix need to be written in terms of specified bases of the domain and codomain. That is, we need a set of domain basis vectors to plug in to D , and when we look at the outputs, we must write those outputs in terms of a set of specified basis vectors of the codomain. Since the domain and codomain of this linear map are the same, we will write this map in terms of the usual monomial bases of the domain and codomain. Unless otherwise

specified, we always write the matrix of a linear map in terms of standard bases of the domain and codomain.

Step 2: We associate degree 4 polynomials with vectors in \mathbb{R}^5 (recall that $\mathbb{R}[x]_4$ is 5-dimensional) via

$$a_4x^4 + a_3x^3 + a_2x^2 + a_1x + a_0 \leftrightarrow \begin{bmatrix} a_4 \\ a_3 \\ a_2 \\ a_1 \\ a_0 \end{bmatrix}$$

so for example, $3x^4 + 7x^2 - 3x - 12$ corresponds to $\begin{bmatrix} 3 \\ 0 \\ 7 \\ -3 \\ -12 \end{bmatrix}$

Step 3: Determine the matrix for D . We need to find $D(\mathbf{e}_1), D(\mathbf{e}_2), D(\mathbf{e}_3), D(\mathbf{e}_4), D(\mathbf{e}_5)$.

$$\mathbf{e}_1: \mathbf{e}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \leftrightarrow x^4 \text{ and } D(x^4) = 4x^3 \leftrightarrow \begin{bmatrix} 0 \\ 4 \\ 0 \\ 0 \\ 0 \end{bmatrix} \text{ so } D(\mathbf{e}_1) = 4\mathbf{e}_2$$

$$\mathbf{e}_2: \mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \leftrightarrow x^3 \text{ and } D(x^3) = 3x^2 \leftrightarrow \begin{bmatrix} 0 \\ 0 \\ 3 \\ 0 \\ 0 \end{bmatrix} \text{ so } D(\mathbf{e}_2) = 3\mathbf{e}_3$$

Continuing with these computations we get that $D(\mathbf{e}_3) = 2\mathbf{e}_4, D(\mathbf{e}_4) = \mathbf{e}_5, D(\mathbf{e}_5) = \mathbf{0}$ thus the matrix is

$$D = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 4 & 0 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

We can check that this actually works. $(3x^4 + 2x^2 + 9x)' = 12x^3 + 4x + 9$ and applying the associated vector to our newfound matrix we see that

$$D \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 4 & 0 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 3 \\ 0 \\ 2 \\ 9 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 12 \\ 0 \\ 4 \\ 9 \end{bmatrix} \leftrightarrow 12x^3 + 4x + 9$$

We finish off with some questions concerning other properties of D .

Question 8.1.11. Is D one-to-one?

Recall that any linear map is one-to-one if and only if the columns of its associated matrix are linearly independent. Alternatively, $T: V \rightarrow W$ is one-to-one if and only if $\ker(T) = \{\mathbf{x} \in V: T(\mathbf{x}) = \mathbf{0}\} = \{\mathbf{0}\}$. The columns of D are $\{4\mathbf{e}_2, 3\mathbf{e}_3, 2\mathbf{e}_4, \mathbf{e}_5, \mathbf{0}\}$, which is a linearly dependent set. This means that $\ker(D) \neq \{\mathbf{0}\}$.

An example of a vector in the kernel is $\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ k \end{bmatrix}$, i.e. constants.

Question 8.1.12. Is D onto?

Recall that any linear map is onto if and only if the columns of its associated matrix span the codomain. Alternatively, $T : V \rightarrow W$ is onto if and only if $\text{Range}(T) = \{\mathbf{y} \in W : \mathbf{y} = T(\mathbf{x}) \text{ for some } \mathbf{x} \in V\} = W$. Looking at the matrix D , we can see that $\text{Span}\{\mathbf{e}_1\} \subsetneq \text{Col}(D)$, so D is not onto. We can in fact find a polynomial in the codomain that does not get mapped to. Namely any multiple of \mathbf{e}_1 (any polynomial with an x^4 term).

We will see in some problems, that when we adjust the domain and codomain of some linear maps, the behavior of the maps themselves change.

8.2 Some Helpful Logic Facts

Before taking our first look at the problem set, we take a quick peek at several ways to logically argue that something is true. The problems in these notes ask for true facts to be *argued* and this word is used to emphasize that a formal proof is not needed. No knowledge of proofs is required to get through this material but there are three logical equivalences that one can use to argue that something is true, and you may find them helpful from time to time.

Direct: The direct method is the most commonly used one. Given an *if-then* statement of the form $P \implies Q$ (pronounced P implies Q or if P then Q), one assumes the *if* statement and works to conclude the statement that comes after *then*.

Example 8.2.1. Fix a matrix $A \in \mathbb{R}^{m \times n}$ and let $S = \{\mathbf{x} \in \mathbb{R}^n : A^2\mathbf{x} = A\mathbf{x}\}$. Argue that if $\mathbf{x}, \mathbf{y} \in S$, then $\mathbf{x} + \mathbf{y} \in S$.

Doing this directly, we *assume* that $\mathbf{x}, \mathbf{y} \in S$ and want to conclude that $\mathbf{x} + \mathbf{y} \in S$. Since $\mathbf{x}, \mathbf{y} \in S$ we know that $A^2\mathbf{x} = A\mathbf{x}$ and $A^2\mathbf{y} = A\mathbf{y}$. Adding them up, we check that they are still in S and see that

$$A^2(\mathbf{x} + \mathbf{y}) = A^2\mathbf{x} + A^2\mathbf{y} = A\mathbf{x} + A\mathbf{y} = A(\mathbf{x} + \mathbf{y})$$

Note that the second to last equality is where we used our assumption.

The next method is often a good idea when arguing something directly seems too hard. In many cases like these, this other method proves much easier.

Contrapositive: Given the statement, “if P then Q ”, its contrapositive is obtained by reversing the direction that you read it, and *negating* both statements. That is, the contrapositive of $P \implies Q$ is, not Q implies not P , i.e. “if Q is false then P is false”. This is logically equivalent to the statement if P then Q and can be a useful method to show that $P \implies Q$ is a true statement.

Example 8.2.2. Consider the statement, “if $x^2 - 6x + 5$ is even, then x is odd”. In doing this directly, we would assume that there is some number a such that $x^2 - 6x + 5 = 2a$ but then we would need to show that $x = 2b + 1$ for some number b , and this feels hard. It turns out the contrapositive makes this much more tractable.

The contrapositive statement is that “if x is not odd, then $x^2 - 6x + 5$ is not even”, in other words, “if x is even, then $x^2 - 6x + 5$ is odd”. Assuming that $x = 2a$ for some number a , we plug it into the equation and get that

$$(2a)^2 - 6(2a) + 5 = 4a^2 - 12a + 4 + 1 = 2(2a^2 - 6a + 2) + 1$$

Therefore, $x = 2b + 1$ where $b = 2a^2 - 6a + 2$ and consequently, $x^2 - 6x + 5$ is odd.

The last one has a slightly different flavor but can also be very effective if you're stuck.

Contradiction: Given the statement “if P then Q ”, one can show it is true by assuming that P is true and Q is false, then hunting down a statement that is obviously false. This obviously false statement is what we call the contradiction, and it implies that our original assumption that P is true but Q is false, was a false assumption all together. Therefore, it must have been true that if P was true then Q must also be true.

Example 8.2.3. Consider the statement, “If $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ form a basis for \mathbb{R}^n , then $m = n$ ”. The assumption here that we begin with is that $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ form a basis for \mathbb{R}^n and $n \neq m$. The latter assumption tells us that either $n < m$ or $n > m$,

If $n < m$ then we have more than n vectors in \mathbb{R}^n and no set of more than n vectors can be linearly independent. That is, $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ form a basis consisting of linearly dependent vectors. This is super duper false because bases must be linearly independent by definition, hence a contradiction.

If $n > m$, then we have a basis for \mathbb{R}^n consisting of fewer than n vectors, but no set of fewer than n vectors can ever span \mathbb{R}^n , giving us another contradiction. Therefore, our assumption must have been false, hence if $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ form a basis for \mathbb{R}^n it must always be true that $m = n$.

Before ending the section we mention two other quick things.

When are two sets equal: By definition, two sets, A and B , are equal if any element of A is also an element of B , and similarly, every element of B is an element of A . If only one of these conditions holds, say every element of A is an element of B , but not every element of B is an element of A , then we say A is a *subset* of B and write $A \subset B$.

The key idea is to take an arbitrary element of one set, and show it belongs to the other, then repeat the process in the other direction. Using the notation above we can write out this process in a series of steps.

- i) Pick an arbitrary element $a \in A$, and show that $a \in B$. This means that $A \subset B$.
- ii) Pick an arbitrary element $b \in B$ and show that $b \in A$. This shows that $B \subset A$.

To summarize, we have that $A = B$ if and only if $A \subset B$ and $B \subset A$. This will prove to be useful a number of times throughout the course. Remember, at its core, **many** of the things we look at are sets! For example, $\text{Col}(A)$, $\text{Row}(A)$, $\text{Null}(A)$, $\text{Range}(T)$, and $\ker(T)$ are all sets, so if we want to argue that any of them are equal, we use methods similar to what we've just described.

Uniqueness: This is a more subtle concept but it will come up twice. We mainly use it when thinking about diagonalizations of matrices. The main statement is that given a diagonalizable matrix A , **its diagonalization** $A = X\Lambda X^{-1}$ **is unique** (up to reordering of the columns). This means that there is only one diagonalization of A and if you are looking at two, they **must** be the same. For example, if $A = X\Lambda X^{-1} = YD Y^{-1}$, up to reordering of the columns of Y and X , we must have $X = Y$ and $\Lambda = D$. The points in the course when this could prove useful should be obvious when they show up. Now, onto the problems!

8.3 Problem Set 1

1. Subtleties with similarity

We saw in Theorem 1.3.3 of the notes that similar matrices have the same eigenvalues. The converse is not always true.

- (a) Find examples of two matrices with the same eigenvalues (counting multiplicities) that are similar

- (b) Find examples of two other matrices with the same eigenvalues (counting multiplicities) that are not similar.
- (c) Argue that if A and B have the same **distinct** eigenvalues, then A and B are similar.

2. Finding Eigenspaces

Compute all eigenvalues of the following matrix and a basis for each eigenspace:

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

Answer the following questions using your computations:

- (a) What are the eigenvalues and eigenvectors of A^3 and $A + 10I$?
- (b) What are the eigenvalues and eigenvectors of A^T ? In general, how are the eigenvalues of A related to the eigenvalues of A^T ? **explain why**
- (c) Is A diagonalizable? If yes, write its diagonalization and compute A^3 .
- (d) From your eigenvalue computation, decide if A is invertible. If yes, what are the eigenvalues of A^{-1} ?
- (e) Can you tell definitively from the eigenvalues of a square matrix A whether A is invertible? If yes, say how. If not, give an example or reason to justify your answer.

3. Invertibility vs. Diagonalizability

In each of the following cases, find an example of a matrix that satisfies the given conditions or say why there can be no such matrix. You must explicitly show the diagonalization of the matrix you chose or explain why your matrix cannot be diagonalized by computing eigenvalues and eigenvectors. Small matrices will work in all cases — 2×2 or 3×3 .

- (a) a matrix that is invertible and diagonalizable
- (b) a matrix that is invertible but not diagonalizable
- (c) a matrix that is singular but diagonalizable
- (d) a matrix that is singular and not diagonalizable

What can you conclude about the relationship between invertibility and diagonalizability of a matrix?

4. Products and Sums of eigenvalues

Let A be an $n \times n$ matrix with eigenvalues $\lambda_1, \dots, \lambda_n$ (maybe not all distinct).

- (a) (6.1 #16) Show that the determinant of A is the product of eigenvalues. i.e., $\det(A) = \lambda_1 \lambda_2 \cdots \lambda_n$.
Hint: Start with the polynomial $\det(A - \lambda I)$ factored as follows:

$$\det(A - \lambda I) = (\lambda - \lambda_1)(\lambda - \lambda_2) \cdots (\lambda - \lambda_n).$$

- (b) (6.1 #17) Show that the sum of the eigenvalues of a 2×2 matrix A is the trace of A by which we mean the sum of the diagonal entries of A .

Hint: If $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, then $\text{trace}(A) = a + d$ and $\det(A - \lambda I) = \lambda^2 - (a + d)\lambda + ad - bc$.

** (You do not need to write up an answer to this part) The trace of any $n \times n$ matrix A is the sum of its eigenvalues. Test this on the 3×3 matrix from (1) and think about why it's true there. Why is this true in general for all $n \times n$ matrices?

5. Rotation Matrices

- (a) (6.1 #14) Argue that counterclockwise rotation by θ degrees in \mathbb{R}^2 is modeled by the linear transformation with matrix:

$$Q = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}.$$

Compute the diagonalization of Q . (Recall that $i^2 = -1$.)

- (b) (6.2 #34):

- i. Argue that the matrix for rotation by $n\theta$ is given by Q^n .
- ii. Compute Q^n using the diagonalization of Q and show that

$$Q^n = \begin{bmatrix} \cos n\theta & -\sin n\theta \\ \sin n\theta & \cos n\theta \end{bmatrix}.$$

The following formulae may help:

$$\cos \theta = \frac{e^{i\theta} + e^{-i\theta}}{2}, \quad \sin \theta = \frac{e^{i\theta} - e^{-i\theta}}{2i}$$

6. Permutation Matrices

(6.1 #34) A *permutation* of $1, 2, 3, \dots, n$ is a reordering of the n numbers. For example (1324) is a permutation of $1, 2, 3, 4$ in which 1 goes to position 3, 3 goes to position 2, 2 goes to position 4 and 4 goes to position 1. The notation is $1 \mapsto 3, 2 \mapsto 4, 3 \mapsto 2, 4 \mapsto 1$.

There are $2 = 2!$ permutations of 1 and 2: (1, 2) and (2, 1). There are $6 = 3!$ permutations of 1, 2, 3: (123), (132), (12)(3), (13)(2), (23)(1), (1)(2)(3). There are $24 = 4!$ permutations of 1, 2, 3, 4 etc

Consider the following *permutation matrix*:

$$P = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Let's examine where the name of P comes from and the eigenvalues and eigenvectors of P .

- (a) What is the effect of multiplying the vector $x = (1, 2, 3, 4)$ by the matrix P , i.e., what is Px ? Do you see why P is called a permutation matrix?
- (b) Write down the permutation matrix that sends $1 \mapsto 1, 2 \mapsto 3, 3 \mapsto 2, 4 \mapsto 4$.

- (c) Compute all eigenvalues of P .
- (d) For each eigenvalue use geometry and your understanding of permutations to find a corresponding eigenvector. You shouldn't have to do tedious computations.

7. * **Fun with polynomials and linear maps.** The following problems are unrelated unless it is mentioned.

- (a) Give an example of a linear map $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that $\text{Range}(T) = \text{Ker}(T)$ or explain why no such example exists.
(Hint: Pick your favorite subspace of \mathbb{R}^2 that will work and construct the matrix to have kernel and range equal to that subspace.)
- (b) Give an example of a linear map $T : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ such that $\text{Range}(T) = \text{Ker}(T)$ or explain why no such example exists.
- (c) Consider the differentiation map $D : \mathbb{R}[x]_3 \rightarrow \mathbb{R}[x]_2$ given by $D(p(x)) = p'(x)$. Find a basis of $\mathbb{R}[x]_3$ and $\mathbb{R}[x]_2$ such that the matrix of D with respect to these bases is

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

and be sure to explain why the sets you have chosen are bases.

- (d) Recall that $\mathbb{R}[x]$ is the (infinite-dimensional) vector space consisting of all polynomials of all degrees. It contains each $\mathbb{R}[x]_n$ as a subspace. Given any $p \in \mathbb{R}[x]$, use **linear algebra** to argue that there exists a polynomial $q \in \mathbb{R}[x]$ such that

$$5q'' + 3q' = p$$

An answer that uses integration **does not count!**

Hint: This question is most easily answered by showing that a certain linear transformation $T : \mathbb{R}[x] \rightarrow \mathbb{R}[x]$ is onto. To take this approach, you need to make sense of what an onto linear map between infinite dimensional vector spaces is.

Chapter 9

Positive Matrices and much much more

In this chapter we will see a new class of matrices and encounter a theorem that is extremely useful in applied settings. We then look at applications of this theorem, ending with an introduction to adjacency matrices of graphs and the Google page rank algorithm.

9.1 Difference Equations

We can define the Fibonacci sequence recursively as follows. Let F_i denote the i^{th} Fibonacci number, and define

$$F_0 = 0, F_1 = 1, F_2 = 1, F_{k+2} = F_{k+1} + F_k$$

This recursive definition let's us enumerate any Fibonacci number we want

$$F_0 = 0, F_1 = 1, F_2 = 1, F_3 = 2, F_4 = 3, F_5 = 5, F_6 = 8, F_7 = 13, F_8 = 21, F_9 = 33, F_{10} = 54, \dots$$

We can see the numbers get relatively large relatively quickly. Now suppose, we wanted to easily find F_{100} . This would take a bit of time by hand, but we will soon see that linear algebra makes this incredibly easy. Let's introduce the close cousin of the Fibonacci sequence first, namely, the golden ratio.

Definition 9.1.1. The **golden ratio**, denoted by the greek letter ϕ , is defined to be $\phi = \frac{1+\sqrt{5}}{2}$. It is also a root of the quadratic polynomial $x^2 - x - 1 = 0$.

The golden ratio came about from the idea of trying to draw the “perfect” rectangle, that is, a rectangle that was the most pleasing to the human eye. It appears everywhere in nature and is intimately related to the Fibonacci sequence.

Example 9.1.2. Let's stop doing linear algebra for a moment and suppose we were a plant. One that grows up and has leaves coming off of its main stem. How would we grow more and more leaves and ensure that the leaf spacing maximized sunlight on the surface of our leaves? We could start with leaf 1, and then rotate halfway around the stem to let leaf 2 grow there. That is, leaf 2 is a rotation of $\frac{1}{2}$ units around the stem from leaf 1.

Next, we would want leaf 3 to be positioned so it does not block too much sunlight from leaves 1 and 2. If we drew a picture and thought about it for a bit, we would end up rotating $\frac{3}{5}$ units away from leaf 2. How about leaf 4? We would rotate this one $\frac{5}{8}$ units from leaf 3. Assuming the plant was immortal, we would continue this process indefinitely, with a new spacing for each leaf.

Question 9.1.3. What would this angle eventually tend to?

If we write it out we see that we obtain a sequence of fractions $\frac{1}{2}, \frac{3}{5}, \frac{5}{8}, \dots, \frac{F_k}{F_{k+1}}$ and the limit of this sequence as k tends to infinity is

$$\lim_{k \rightarrow \infty} \frac{F_k}{F_{k+1}} = \phi$$

That is, the limit of successive ratios of Fibonacci numbers tends to the golden ratio! For those who are amazed and wondering why, linear algebra can be used to compute this limit and is left as an optional exercise.

Now let's try to find F_{100} . Let $\mathbf{u}_k = \begin{bmatrix} F_{k+1} \\ F_k \end{bmatrix}$ and $\mathbf{u}_{k+1} = \begin{bmatrix} F_{k+2} \\ F_{k+1} \end{bmatrix}$. Can we find a matrix that eats two successive Fibonacci numbers and outputs the next one? That is, can we find a 2×2 matrix, A such that $A\mathbf{u}_k = \mathbf{u}_{k+1}$? Sure we can! Using the recursive definition of Fibonacci numbers we find that our matrix is $\begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$ and we can check that

$$\begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \mathbf{u}_k = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} F_{k+1} \\ F_k \end{bmatrix} = \begin{bmatrix} F_{k+1} + F_k \\ F_{k+1} \end{bmatrix} = \begin{bmatrix} F_{k+2} \\ F_{k+1} \end{bmatrix} = \mathbf{u}_{k+1}$$

Now let's start this process with $\mathbf{u}_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$. We know that $A\mathbf{u}_0 = \mathbf{u}_1$ and applying A to both sides of this equation we see that

$$\mathbf{u}_2 = A\mathbf{u}_1 = A(A\mathbf{u}_0) = A^2\mathbf{u}_0$$

Applying the same idea for arbitrary powers we see that $\mathbf{u}_k = A^k\mathbf{u}_0$ hence $\mathbf{u}_{100} = A^{100}\mathbf{u}_0$, so it remains to compute A^{100} . We do this easily if A is diagonalizable.

We find the eigenvalues of our Fibonacci matrix by computing

$$\det(A - \lambda I) = \begin{vmatrix} 1-\lambda & 1 \\ 1 & -\lambda \end{vmatrix} = -\lambda(1-\lambda) - 1 = \lambda^2 - \lambda - 1 = 0$$

Do we remember what the roots of this are?!? We have eigenvalues being $\lambda_1 = \frac{1+\sqrt{5}}{2} = \phi$ and $\lambda_2 = \frac{1-\sqrt{5}}{2}$. If we weren't yet convinced of the relationship between the Fibonacci numbers and the golden ratio, we should be sufficiently amazed now.

We find the eigenvectors by finding bases for our respective null spaces (eigenspaces) and get that the eigenvector corresponding to eigenvalue λ_1 is $\mathbf{x}_1 = \begin{bmatrix} \lambda_1 \\ 1 \end{bmatrix}$ and likewise for λ_2 we get $\mathbf{x}_2 = \begin{bmatrix} \lambda_2 \\ 1 \end{bmatrix}$. This information allows us to diagonalize A and write $A = X\Lambda X^{-1}$. This means that

$$\mathbf{u}_k = A^k\mathbf{u}_0 = X\Lambda^k X^{-1}\mathbf{u}_0$$

At this stage, computing F_{100} seems reasonable, but still requires some matrix computation, in addition to taking large powers of Λ , so let's try and do better and get a closed form for the k^{th} Fibonacci number, purely in terms of these eigenvalues. We do so by carefully looking at what happens in a more general setting.

General Setting: The general phenomenon that is occurring is that we have a matrix $A \in \mathbb{R}^{n \times n}$ which maps our initial state vector \mathbf{u}_0 to the k^{th} state vector, \mathbf{u}_k via $A^k\mathbf{u}_0 = \mathbf{u}_k$. Assuming that A is diagonalizable, we have that $\mathbf{u}_k = A^k\mathbf{u}_0 = X\Lambda^k X^{-1}\mathbf{u}_0$ and express \mathbf{u}_k in terms of our eigenbasis in the following way.

Set $\mathbf{c} = X^{-1}\mathbf{u}_0$ where $X = [\mathbf{x}_1 \ \mathbf{x}_2 \ \cdots \ \mathbf{x}_n]$ and $\mathbf{c} = \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{bmatrix}$. Note that you find the vector \mathbf{c} by solving

the system $[X|\mathbf{u}_0]$.

Next, for this vector \mathbf{c} , we have

$$\begin{aligned} \mathbf{u}_k &= X\Lambda^k\mathbf{c} = [\mathbf{x}_1 \ \mathbf{x}_2 \ \cdots \ \mathbf{x}_n] \begin{bmatrix} \lambda_1^k & 0 & \cdots & 0 \\ 0 & \lambda_2^k & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_n^k \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{bmatrix} \\ &= [\mathbf{x}_1 \ \mathbf{x}_2 \ \cdots \ \mathbf{x}_n] \begin{bmatrix} \lambda_1^k c_1 \\ \lambda_2^k c_2 \\ \vdots \\ \lambda_n^k c_n \end{bmatrix} = c_1 \lambda_1^k \mathbf{x}_1 + c_2 \lambda_2^k \mathbf{x}_2 + \cdots + c_n \lambda_n^k \mathbf{x}_n \end{aligned}$$

It is worth noting that two important things just happened. First, we found a nice uniform way to express the k^{th} state vector as a linear combination of our eigenbasis vectors \mathbf{x}_i . Second, we used a fundamental fact regarding matrix multiplication in terms of its columns.

Matrix multiplication in terms of matrix columns

Let $A = [\mathbf{a}_1 \ \mathbf{a}_2 \ \cdots \ \mathbf{a}_n] \in \mathbb{R}^{n \times n}$ and $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$. Then $A\mathbf{x} = x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \cdots + x_n\mathbf{a}_n$. **This allows us to easily see that $\text{Range}(T) = \text{Col}(A)$, and furthermore, that every element of $\text{Col}(A)$ is of the form $A\mathbf{x}$ for some \mathbf{x} .** It is a worthwhile use of your time to convince yourself of these facts.

We now summarize the whole section thus far in a proposition.

Proposition 9.1.4. *Let \mathbf{u}_0 denote an initial state vector and let $A \in \mathbb{R}^{n \times n}$ satisfy $\mathbf{u}_k = A^k\mathbf{u}_0$. If A is diagonalizable, with eigenvalues λ_i and eigenbasis given by $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ then*

$$\mathbf{u}_k = \sum_{i=1}^n c_i \lambda_i^k \mathbf{x}_i = c_1 \lambda_1^k \mathbf{x}_1 + c_2 \lambda_2^k \mathbf{x}_2 + \cdots + c_n \lambda_n^k \mathbf{x}_n$$

where $\mathbf{c} = \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{bmatrix}$ satisfies the linear system $X^{-1}\mathbf{u}_0 = \mathbf{c}$ with $X = [\mathbf{x}_1 \ \mathbf{x}_2 \ \cdots \ \mathbf{x}_n]$.

Now let's return to the main task of computing F_{100} . If we were to solve the linear system $X^{-1}\mathbf{u}_0 = \mathbf{c}$, we would see that

$$\mathbf{c} = \begin{bmatrix} 1 \\ \frac{\lambda_1 - \lambda_2}{\lambda_1 - \lambda_2} \\ \lambda_1 - \lambda_2 \end{bmatrix}$$

hence

$$\mathbf{u}_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \underbrace{\frac{1}{\lambda_1 - \lambda_2}}_{c_1} \underbrace{\begin{bmatrix} \lambda_1 \\ 1 \end{bmatrix}}_{\mathbf{x}_1} - \underbrace{\frac{1}{\lambda_1 - \lambda_2}}_{c_2} \underbrace{\begin{bmatrix} \lambda_2 \\ 1 \end{bmatrix}}_{\mathbf{x}_2}$$

therefore

$$\mathbf{u}_k = \frac{1}{\lambda_1 - \lambda_2} \lambda_1^k \begin{bmatrix} \lambda_1 \\ 1 \end{bmatrix} - \frac{1}{\lambda_1 - \lambda_2} \lambda_2^k \begin{bmatrix} \lambda_2 \\ 1 \end{bmatrix}$$

To obtain F_k , we take the second coordinate of $\mathbf{u}_k = \begin{bmatrix} F_{k+1} \\ F_k \end{bmatrix}$, which we can see from the above equation is

$$F_k = \frac{\lambda_1^k - \lambda_2^k}{\lambda_1 - \lambda_2}$$

What happens as k tend to infinity? We would need to compute $\lim_{k \rightarrow \infty} \mathbf{u}_k = \lim_{k \rightarrow \infty} c_1 \lambda_1^k \mathbf{x}_1 + c_2 \lambda_2^k \mathbf{x}_2$. Notice that $\lambda_2 = -0.618$, so $-1 < \lambda_2 < 0$ and $\lim_{k \rightarrow \infty} \lambda_2^k = 0$. This means that

$$\lim_{k \rightarrow \infty} c_1 \lambda_1^k \mathbf{x}_1 + c_2 \lambda_2^k \mathbf{x}_2 = \lim_{k \rightarrow \infty} c_1 \lambda_1^k \mathbf{x}_1$$

and we can conclude that $F_k \sim \frac{\lambda_1^k}{\lambda_1 - \lambda_2} = \frac{1}{\sqrt{5}} \left(\frac{1 + \sqrt{5}}{2} \right)^k$.

While this is a sort of phenomenon in itself, the key to determining the limiting behavior was that the second eigenvalue tended to 0 as k got large. This is the central idea surrounding positive matrices and the main focus of our attention in the next section.

9.2 Positive Matrices

We got a hint of the phenomenon that is occurring but we need some more examples before really seeing how it works.

Example 9.2.1. Suppose we own a rental car agency in Seattle. The agency has cars in Seattle as well as cars that have been rented and returned to other locations outside of Seattle. The current stats are

- Fraction of rental cars in Seattle at the start: 0.02
- Fraction of rental cars outside Seattle at the start: 0.98
- Every month 20% of Seattle cars leave and 5% of outside cars come in.

Question 9.2.2. What fraction of rental cars are in Seattle in the long run?

Step 1: Begin with an initial state vector whose first coordinate is the fraction of cars in Seattle and second coordinate is the fraction of cars outside Seattle, that is $\mathbf{u}_0 = \begin{bmatrix} 0.02 \\ 0.98 \end{bmatrix}$.

Step 2: Determine \mathbf{u}_1 . Since 20% of cars leave every month and 5% come back in, we know that after one month the fraction of cars in Seattle is

$$(.8)(0.02) + (0.05)(.98)$$

and the fraction of cars outside Seattle is

$$(.2)(.02) + (.95)(.98)$$

Step 3: Write \mathbf{u}_1 in terms of this information.

$$\mathbf{u}_1 = \begin{bmatrix} (.8)(0.02) + (0.05)(.98) \\ (.2)(.02) + (.95)(.98) \end{bmatrix} = \underbrace{\begin{bmatrix} .8 & .05 \\ .2 & .95 \end{bmatrix}}_A \underbrace{\begin{bmatrix} .02 \\ .98 \end{bmatrix}}_{\mathbf{u}_0}$$

Step 4: Using the same reasoning from the Fibonacci example, we can conclude that at the end of month k , our k^{th} state vector is $\mathbf{u}_k = A^k \mathbf{u}_0$. In computing arbitrary powers of A , we use diagonalization as our tool and find that the eigenvalues and eigenvectors are

$$\lambda_1 = 1, \mathbf{x}_1 = \begin{bmatrix} .2 \\ .8 \end{bmatrix} \quad \lambda_2 = 0.75, \mathbf{x}_2 = \begin{bmatrix} -1 \\ 1 \end{bmatrix}$$

Applying proposition 9.1.4, we get

$$\mathbf{u}_0 = 1 \begin{bmatrix} .2 \\ .8 \end{bmatrix} + .18 \begin{bmatrix} -1 \\ 1 \end{bmatrix}$$

hence

$$\mathbf{u}_k = (1)(1^k) \begin{bmatrix} .2 \\ .8 \end{bmatrix} + \underbrace{(.18)(.75)^k}_{\rightarrow 0 \text{ as } k \rightarrow \infty} \begin{bmatrix} -1 \\ 1 \end{bmatrix}$$

We then conclude that the limiting behavior is $\lim_{k \rightarrow \infty} \mathbf{u}_k = \begin{bmatrix} .2 \\ .8 \end{bmatrix}$, that is, the limiting behavior tended towards the eigenvector associated to the *largest* eigenvalue. We conclude that in the long run, 20% of cars end up in Seattle.

Is this a general phenomenon? Does the biggest eigenvalue always determine the limiting behavior? Is the biggest eigenvalue always 1, with all other eigenvalues being smaller in absolute value? The emphatic answer is YES!

Definition 9.2.3. Let $A = (a_{ij}) \in \mathbb{R}^{m \times n}$.

- A is a **positive** matrix, denoted $A > 0$, if $a_{ij} > 0$ for all i, j .
- A is a **non-negative matrix**, denoted $A \geq 0$ if $a_{ij} \geq 0$ for all i, j .
- A is **Markov** (also known as stochastic) if the entries in each column add up to 1. That is, if

$$\sum_{i=1}^m a_{ij} = 1 \quad \forall j$$

- A is a **positive Markov** matrix if $a_{ij} > 0 \quad \forall i, j$ and $\sum_{i=1}^m a_{ij} = 1 \quad \forall j$.

Note that the matrix $\begin{bmatrix} .8 & .05 \\ .2 & .95 \end{bmatrix}$ of the previous example is positive Markov. The following theorem explains the phenomenon that we have now seen several times. The ideal hypothesis is that our matrix is positive Markov, but dropping the Markov assumption still yields powerful results.

Theorem 9.2.4. (Perron-Frobenius for positive matrices) *If $A > 0$ then A has a dominant eigenvalue λ_A with the following properties:*

1. $\lambda_A > 0$ and its associated eigenvector $\mathbf{x}_A > 0$ (has all positive entries).
2. $\text{AM}(\lambda_A) = 1$
3. If μ is another eigenvalue of A , then $|\mu| < \lambda_A$.
4. A has no other eigenvectors with non-negative entries

The proof of this theorem involves the notion of compactness, which is not a part of this course. For the interested reader, one can find a proof in the book *Numerical Linear Algebra* by Lloyd Trefethen and David Bau. The results that we will be using the most will be for positive Markov matrices. To prove these results we will assume all results of theorem 9.2.4 without proof. Adding the Markov assumption to the hypothesis of the previous theorem greatly strengthens our result.

Theorem 9.2.5. (Perron-Frobenius for positive Markov matrices) *If A is positive Markov then A has a dominant eigenvalue λ_A with the following properties:*

1. $\lambda_A = 1$.
2. If μ is another eigenvalue of A , then $|\mu| < 1$.
3. If $\mathbf{u}_0 \geq 0$, then $\lim_{k \rightarrow \infty} A^k \mathbf{u}_0 = c \mathbf{x}_A$ where $c \geq 0$.

Proof. Recall (from the previous problem set) that A and A^\top have the same eigenvalues. Since the column

entries of A sum to 1, the row entries of A^\top sum to 1. Let $\mathbf{1} = \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix}$. Since the row entries of A^\top sum to 1,

this implies that

$$A^\top \mathbf{1} = A^\top \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix}$$

hence 1 is an eigenvalue of A^\top with eigenvector $\mathbf{1}$. This implies that the dominant eigenvalue of A^\top , λ_{A^\top} , is 1 by part 4 of theorem 9.2.4. This in turn, tells us that $\lambda_A = 1$. Note that we are implicitly using the fact that A^\top is positive.

The second statement of the theorem follows immediately from theorem 9.2.4.

To prove the third statement, we assume for simplicity that A is diagonalizable. In general this is not always the case but the extent of cases that we see will only involve this case. If $A = X \Lambda X^{-1}$ then by proposition 9.1.4 we know that

$$\mathbf{u}_k = \sum_{i=1}^n c_i \lambda_i^k \mathbf{x}_i = c_1 (1)^k \mathbf{x}_A + c_2 \lambda_2^k \mathbf{x}_2 + \cdots + c_n \lambda_n^k \mathbf{x}_n$$

The previous statement implies that $|\lambda_i| < 1$ for all $i = 2, 3, \dots, n$ hence $\lim_{k \rightarrow \infty} \lambda_i^k = 0$ for all non-dominant eigenvalues. From this we can conclude that

$$\lim_{k \rightarrow \infty} \mathbf{u}_k = \lim_{k \rightarrow \infty} c_1 (1)^k \mathbf{x}_A + \underbrace{c_2 \lambda_2^k \mathbf{x}_2 + \cdots + c_n \lambda_n^k \mathbf{x}_n}_{\rightarrow 0 \text{ as } k \rightarrow \infty} = c_1 \mathbf{x}_A$$

□

Warning: If A is non-negative, the Perron-Frobenius theorem, as written, is false. Witness the counterexample $A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. We have $\lambda_1 = 1$ but $\lambda_2 = -1$.

The proof of the first statement contains a method that we shouldn't forget. In general, when trying to argue that a given matrix is Markov, we almost never want to look at the individual matrix entries and work up from there. Instead, we'll want to use something a little more slick.

Alternative Definition of Markov

Let $\mathbf{1}$ denote the vector of all 1's. If we transpose both sides of the equation $A^\top \mathbf{1} = \mathbf{1}$ using the fact that $(AB)^\top = B^\top A^\top$ we get that $\mathbf{1}^\top A = \mathbf{1}^\top$. Note that multiplying a matrix *on the left* by a row vector is defined. You should check this for yourself. We then have the alternative definition of Markov given by

$$A \text{ is Markov} \Leftrightarrow \mathbf{1}^\top A = \mathbf{1}^\top$$

Even though Perron-Frobenius fails for non-negative Markov matrices, not all hope is lost.

Proposition 9.2.6. *If $A \geq 0$ but A^k is positive Markov, then 1 is still the unique dominant eigenvalue of A with $|\mu| < 1$ for all other eigenvalues μ .*

Proof. If μ is an eigenvalue of A , then μ^k is an eigenvalue of A^k . A^k is a positive Markov matrix, hence by theorem 9.2.5 1 is its dominant eigenvalue. This means that $1 = \lambda^k$ for some eigenvalue λ of A , and $|\mu^k| < 1$ for all other eigenvalues μ , thus $\lambda^k = 1$ and $|\mu| < 1$ for all other eigenvalues μ . \square

This last proposition can be useful in determining limiting behavior of certain systems.

Example 9.2.7. Suppose we have three groups with populations p_1, p_2, p_3 respectively and further assume that after each week, each group splits in half and joins the others. This behavior can be modeled by a

non-negative Markov matrix. Let $\mathbf{u}_0 = \begin{bmatrix} p_1 \\ p_2 \\ p_3 \end{bmatrix}$, then after one month we have

$$\mathbf{u}_1 = A\mathbf{u}_0 = \underbrace{\begin{bmatrix} 0 & 1/2 & 1/2 \\ 1/2 & 0 & 1/2 \\ 1/2 & 1/2 & 0 \end{bmatrix}}_{\text{Markov and non-negative}} \begin{bmatrix} p_1 \\ p_2 \\ p_3 \end{bmatrix}$$

Moreover,


$$\mathbf{u}_2 = A^2\mathbf{u}_0 = \underbrace{\begin{bmatrix} 1/2 & 1/4 & 1/4 \\ 1/4 & 1/2 & 1/4 \\ 1/4 & 1/4 & 1/2 \end{bmatrix}}_{\text{positive Markov}} \begin{bmatrix} p_1 \\ p_2 \\ p_3 \end{bmatrix}$$

Computing eigenvalues and eigenvectors we see that $\lambda_A = 1$ with $\mathbf{x}_A = \begin{bmatrix} 1/2 \\ 1/3 \\ 1/3 \end{bmatrix}$ and $\lambda_2 = \lambda_3 = -1/2$. By

Perron-Frobenius, the limiting behavior will approach \mathbf{x}_A and in the long run, all populations will have the same size.

9.3 Applications: Adjacency Matrices and Google Page Rank

A **graph** G is a collection of **nodes** or **vertices**, labeled $1, 2, \dots, n$ and **edges** which are represented by a pair of nodes $\{i, j\}$ namely the node that the given edge connects. An example of a graph looks like



Graphs are sometimes written as $G = (E, V)$ where E is the set of edges and V is the set of vertices. Our example graph above would be written as $G = (\{1, 2\}, \{2, 3\}, 1, 2, 3)$. To any graph G we can associate a matrix A_G , known as the adjacency matrix of G .

Definition 9.3.1. The **adjacency matrix** A_G of a graph G with n nodes is a $n \times n$ matrix defined as follows. The rows and columns of A are indexed by the node labels $1, \dots, n$ and the (i, j) -entry of A is 1 if the pair $\{i, j\}$ is an edge in G and 0 otherwise. That is $A_G = (a_{ij})$ where

$$a_{ij} = \begin{cases} 1 & \text{there exists an edge from } i \text{ to } j \\ 0 & \text{otherwise} \end{cases}$$

The adjacency matrix for the graph G from above would be

$$A_G = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

Notice that this matrix is non-negative. We can say a few things about it thanks to yet another version of Perron-Frobenius, this one only attributed to Frobenius (1912).

Proposition 9.3.2. *If $A \geq 0$ then A has an eigenvalue λ_A such that*

1. $\lambda_A \geq 0$ with eigenvector(s) $\mathbf{x}_A \geq 0$.
2. If μ is another eigenvalue of A then $|\mu| \leq \lambda_A$.

Note that if λ_A is not positive, then $\text{AM}(\lambda_A)$ is not necessarily 1. This is witnessed by the counterexample $A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$ which has a $\lambda_A = 0$ and $\text{AM}(\lambda_A) = 2$.

Computing the eigenvalues and eigenvectors of A_G , we get

$$\lambda_1 = \sqrt{2} \quad \mathbf{x}_1 = \begin{bmatrix} \sqrt{2} \\ 2 \\ \sqrt{2} \end{bmatrix}, \quad \lambda_2 = 0 \quad \mathbf{x}_2 = \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}, \quad \lambda_3 = -\sqrt{2} \quad \mathbf{x}_3 = \begin{bmatrix} \sqrt{2} \\ -2 \\ \sqrt{2} \end{bmatrix}$$

Note that here, in the non-negative case, we have another eigenvalue μ such that $|\mu| = \lambda_{A_G}$. You may also have notice that in addition to being non-negative, A_G is also Markov. We can combine part of the proof of theorem 9.2.5 to obtain the following corollary.


Corollary 9.3.3. *If A is non-negative and Markov, then $\lambda_A = 1$ and $\mathbf{x}_A \geq 0$.*

Over the coming weeks we will be looking at lots of graphs and adjacency matrices but we leave further investigation to problem set 2 and later lectures. We finish this section with a powerful application of what we have learned in this chapter.

9.3.1 Google Page Rank

The Google page rank algorithm was first discovered by Larry Page and Sergey Brin in 1998. The central idea is to decide the “rank” or “importance” of a webpage by the importance of pages that link to it. When we submit a query, we want to see the most relevant pages at the top of the list, so a useful measurement of importance is needed.

This is where a graph comes into the picture. The world wide web can be thought of as a *directed* graph or network, with nodes indicating webpages and edges indicating links from one page to another. An example network with 6 webpages and links could look like



From this network we construct a *Page rank matrix* by modeling movement between pages via a 6×6 matrix of probabilities. We compute it according to the following rules:

When your browser is at page i , you have two choices

1. Teleportation: Go directly to a different page by typing the webpage directly into the browser
2. Follow a link.

If you teleport, assume all pages are equally likely to move to.

If you follow a link, assume all links are equally likely.

The ij entry of the page rank matrix is then the probability of going **from** j **to** i .

Let's roll a die that predetermines whether we link or teleport, this will come into play when computing our matrix entries.

- If we roll 1, 2, 3, 4, or 5, then we link with equal probability.
- If we roll a 6, then we teleport

Lets focus on column 1 of this matrix. The entries are a_{i1} for $i = 1, 2, 3, 4, 5, 6$. Remember, to compute a_{ij} we only focus on going **from** j **to** i . We use the example network at each step.

- a_{11} : We can only teleport from 1 to 1 with no linking possible. We have a $1/6$ chance of rolling a 6 and a $1/6$ chance of teleporting to 1 out of the 6 possible webpages. The total probability is then $(1/6)(1/6) = 1/36$.

- a_{21} : Similarly to the previous case, we have a $1/6$ chance of teleporting from 1 to 2. Selecting 2 out of 6 possibilities means that the total probability that we teleport from 1 to 2 is $(1/6)(1/6) = 1/36$. Note that while we have a $5/6$ chance of linking (i.e. rolling anything but a 6), there is no link from 1 to 2 in our network. At the beginning we could either link **or** teleport, so adding these probabilities we get that $a_{21} = (1/6)(1/6) + (5/6)(0) = 1/36$.
- a_{31} : This is the same as above
- a_{41} : This is the only computation that is different because there is a link from 1 to 4. We have a $1/6$ chance of rolling a 6 so the teleportation probability is still $1/36$ as usual. We have a $5/6$ chance of rolling anything other than a 6, and a $1/1$ chance of linking from 1 to 4. Note that if 1 had k outgoing edges, then the chance of linking would then be $1/k$. This gives the total linking probability as $a_{41} = (1/6)(1/6) + (5/6)(1) = 31/36$.
- a_{51} : This is the same as a_{21} .
- a_{61} : This is the same as a_{21} .

Summarizing, the first column of our page rank matrix is $\begin{bmatrix} 1/36 \\ 1/36 \\ 1/36 \\ 31/36 \\ 1/36 \\ 1/36 \end{bmatrix}$. As practice, you should fill out the

rest of the google page rank matrix and check your answer with the one below. Before giving the full matrix, we make several key observations about this procedure and the resulting matrix that we obtain. This is where Perron-Frobenius comes to the rescue!

Important Facts

- Teleportation **guarantees** that our matrix will be positive since every entry will have at least $1/36$.
- Since the entries are probabilities, we will be certain that the resulting matrix will be Markov. **Hint:** This is how you will compute column 6. It also provides many shortcuts in computing the columns of page rank matrices.
- We assume all pages are equally likely at the beginning, so our initial state vector, whose entries rank

the importance of each webpage at time 0, is $\mathbf{u}_0 = \begin{bmatrix} 1/6 \\ 1/6 \\ 1/6 \\ 1/6 \\ 1/6 \\ 1/6 \end{bmatrix}$.

- The importance transfers to $\mathbf{u}_1 = P\mathbf{u}_0$ and $\mathbf{u}_k = P^k\mathbf{u}_0$ where P is the page rank matrix. Therefore, by theorem 9.2.5, we have $\lim_{k \rightarrow \infty} \mathbf{u}_k = c_1(1^k)\mathbf{x}_P$. That is, the importance of webpages, in the long run, is determined by the dominant eigenvector. Since this system limits to a multiple of it, we must normalize \mathbf{x}_P for it to accurately represent a probability vector.

Definition 9.3.4. We call \mathbf{x}_P the **page rank vector**. It is the dominant eigenvector of the page rank matrix and is always taken to be normalized in the sense that the entries sum to 1.

In finishing out the computation, we obtain

$$P = \begin{bmatrix} 1/36 & 31/36 & 31/36 & 1/36 & 1/36 & 6/36 \\ 1/36 & 1/36 & 1/36 & 11/36 & 1/36 & 6/36 \\ 1/36 & 1/36 & 1/36 & 11/36 & 16/36 & 6/36 \\ 31/36 & 1/36 & 1/36 & 1/36 & 1/36 & 6/36 \\ 1/36 & 1/36 & 1/36 & 11/36 & 1/36 & 6/36 \\ 1/36 & 1/36 & 1/36 & 1/36 & 16/36 & 6/36 \end{bmatrix}$$

We use Julia to compute the page rank vector of the example network and we obtain

$$\mathbf{x}_P = \begin{bmatrix} 0.599066 \\ 0.253028 \\ 0.358456 \\ .588717 \\ 0.253028 \\ 0.194924 \end{bmatrix}$$

Notice that this is not yet normalized, but ignoring the scaling factor, we can see that the first entry has the highest “rank” and the fourth entry has second highest “rank”. We then conclude that the most important webpage in our network is webpage 1, and the second most important is webpage 4.

9.4 Problem Set 2

1. (6.2 #9) Suppose a sequence $\{G_k\}$ is defined as $G_{k+2} = \frac{1}{2}G_{k+1} + \frac{1}{2}G_k$.

(a) Find the matrix A such that

$$\begin{bmatrix} G_{k+2} \\ G_{k+1} \end{bmatrix} = A \begin{bmatrix} G_{k+1} \\ G_k \end{bmatrix}.$$

(b) Find the eigenvalues and eigenvectors of A .

(c) Find the limit as $n \rightarrow \infty$ of the matrices A^n .

(d) If $G_0 = 0$ and $G_1 = 1$ show that G_n approaches $\frac{2}{3}$ as $n \rightarrow \infty$.

2. (6.2 #15,#16) Consider the matrices

$$A_1 = \begin{bmatrix} .6 & .9 \\ .4 & .1 \end{bmatrix} \text{ and } A_2 = \begin{bmatrix} .6 & .9 \\ .1 & .6 \end{bmatrix}.$$

Which of these matrices have the property that A^k approaches the zero matrix as $k \rightarrow \infty$? Why does this happen?

The following questions are all about A_1 .


(a) Diagonalize A_1 as $A_1 = X\Lambda X^{-1}$.

(b) What is the limit of Λ^k as $k \rightarrow \infty$?

(c) What is the limit of A_1^k as $k \rightarrow \infty$? What do you see in the columns of this limiting matrix?

3. (6.2 #20, #21)

- (a) Suppose $A \in \mathbb{R}^n$ is diagonalizable and show that $\det(A) = \lambda_1 \lambda_2 \cdots \lambda_n$. **Note:** You already showed this last week but this is here for you to realize that the diagonalizability assumption makes the problem considerably easier to show.
- (b) Show that $\text{trace}(PQ) = \text{trace}(QP)$ for any two $n \times n$ matrices P and Q . You might ramp up by first checking it for two 2×2 matrices, then 3×3 and so on.
- (c) Whether you can do (b) or not, use the result to show that if A is diagonalizable, then the trace of A is the sum of the eigenvalues of A .
4. (6.2 #25) Recall that the column space of a $n \times n$ matrix A , denoted $\text{Col}(A)$, is the span of the columns of A . Suppose A is a $n \times n$ non-zero matrix such that $A^2 = A$.
- (a) Show that $\lambda = 1$ is an eigenvalue of A with eigenspace equal to $\text{Col}(A)$.
- (b) What is the eigenspace of $\lambda = 1$ if A is invertible? Is A diagonalizable in this case? If yes, write its diagonalization.
- (c) If A is not invertible, what are its eigenvalues and their eigenspaces? Is A diagonalizable in this case? If yes, write its diagonalization. **Note:** If you can write its diagonalization, it will not be with explicit vectors, but rather a general diagonalization in terms of vectors coming from the various eigenspaces that you have found.
5. (6.2 #24, #29)
- (a) Consider the set of all 4×4 matrices that are diagonalized by the same eigenvector matrix X :
- $$S_X = \{X\Lambda X^{-1} \in \mathbb{R}^{4 \times 4} : \Lambda \text{ is a diagonal matrix} \}.$$
- Show that S is a subspace of $\mathbb{R}^{4 \times 4}$. (Check the properties of a subspace.)
- (b) What is S_I where I is the identity matrix?
- (c) Suppose the same X diagonalizes A and B , i.e., $A = X\Lambda_1 X^{-1}$ and $B = X\Lambda_2 X^{-1}$. Argue that $AB = BA$. (Recall that in general matrix multiplication is not commutative, i.e., $AB \neq BA$.)
6. (6.2 #38) Recall that a matrix A is *similar* to a matrix C if there is an invertible matrix B such that $A = BCB^{-1}$. Also, we saw in class that similar matrices have the same eigenvalues. Suppose Λ is the diagonal matrix with the eigenvalues of A on its diagonal. Are A and Λ always similar? If yes, say why. If not, provide an example in which they are not similar and explain what happened. Under what conditions are A and Λ similar?
7. * **How to find a triangle in a graph.** A *graph* G is a collection of *nodes* labeled $1, 2, \dots, n$ and *edges* which are pairs of nodes. Shown below is a graph with 5 nodes labeled $1, \dots, 5$ and edges $\{1, 2\}, \{1, 3\}, \{2, 4\}, \{3, 4\}, \{3, 5\}, \{4, 5\}$. A triangle in a graph G is a triple of nodes i, j, k such that all edges $\{i, j\}, \{i, k\}, \{j, k\}$ are present in G . For example $3, 4, 5$ forms a triangle in the graph below. Two nodes i and j are *neighbors* in G if $\{i, j\}$ is an edge in G .



If the graph G is very large it becomes hard to decide if there is a triangle in it by simply looking at the graph. In this problem we will see that linear algebra can be used to decide if a graph contains a triangle.

- (a) The *adjacency matrix* A of a graph G with n nodes is a $n \times n$ matrix defined as follows. The rows and columns of A are indexed by the node labels $1, \dots, n$ and the (i, j) -entry of A is 1 if the pair $\{i, j\}$ is an edge in G and 0 otherwise. An example can be seen in Problem 2.4A on page 76 in Strang's book. Write down the adjacency matrix of the graph G shown above.
 - (b) Let $B = A^2$ where A is the adjacency matrix of G . The entry b_{ij} in B is the dot product of two vectors sitting in A . Which vectors are they? In our example, how is b_{23} formed?
 - (c) For three nodes i, j, k from G , argue that $a_{ik}a_{kj}$ is 1 exactly when k is a common neighbor of i and j .
 - (d) Using the above, what does b_{ij} count?
 - (e) The nodes i, j, k form a triangle if and only if i and j are neighbors and k is a common neighbor of i and j . If i, j, k is a triangle what property must a_{ij} and b_{ij} have?
 - (f) Putting all this together can you construct an algorithm that takes as input two indices i, j , and determines whether or not there exists a triangle with the edge between i and j as one of the sides. Justify your algorithm. **Your algorithm MUST use both A and B . An algorithm that only uses A is not efficient!** You can write an informal algorithm if you're not familiar with programming. Just writing the steps out is fine.
 - (g) Use your algorithm to find the triangles in the example graph shown above.
 - (h) (optional) If you know about running times of algorithms, do you see how fast this algorithm runs? Is it faster than checking all triples of nodes in G ?
8. (10.3 #5) Every year 2% of young people become old and 3% of old people die. There are no births. Without doing any math, what do you think happens in the long run to these people? Check that the difference equation for this population is

$$\begin{bmatrix} \text{young} \\ \text{old} \\ \text{dead} \end{bmatrix}_{k+1} = \begin{bmatrix} 0.98 & 0 & 0 \\ 0.02 & 0.97 & 0 \\ 0 & 0.03 & 1 \end{bmatrix} \begin{bmatrix} \text{young} \\ \text{old} \\ \text{dead} \end{bmatrix}_k.$$

Confirm your suspicion by computing the limit of u_k as k goes to infinity.

9. (10.3 #11) Complete the following to a Markov matrix


$$\begin{bmatrix} .7 & .1 & .2 \\ .1 & .6 & .3 \\ - & - & - \end{bmatrix}$$

so that $(1, 1, 1)$ is the dominant eigenvector. What is the dominant eigenvalue? Write down the general principle that you are using and explain why it is true.

10. (10.3 #6 and #9)

- (a) Suppose A is a 2×2 Markov matrix and $u \in \mathbb{R}^2$ is a nonnegative vector ($u \neq 0$) whose entries add to α . Argue that Au is a nonnegative vector whose entries add to α .
Hint: Let $\mathbf{1}$ be the column vector with all entries equal to 1. Think of $\mathbf{1}$ as a $n \times 1$ matrix. What are $\mathbf{1}^\top A$, $\mathbf{1}^\top u$ and $\mathbf{1}^\top Au$?
- (b) Using the above argue that all powers of A are Markov. Check A^2 first.
Hint: What should be the value of α that can help you here?
- (c) Suppose $Au = \lambda u$ where $\lambda \neq 1$. Then what must α be? Illustrate on a 2×2 Markov matrix of your choice.
- (d) Do you expect these results to hold if we replace $A \in \mathbb{R}^{2 \times 2}$ with $A \in \mathbb{R}^{n \times n}$ and $u \in \mathbb{R}^2$ with $u \in \mathbb{R}^n$? Explain your answer.

11. Compute the page rank of the 5 webpages you see in the following network. An arrow from i to j indicates that page i contains a link to page j . Use the rule that when you are at a webpage you will teleport with probability $\frac{1}{2}$ and follow a link with probability $\frac{1}{2}$. Assume that all links from a page have equal probability of being followed. You can also assume that you will teleport to any of the vertices in the network with equal probability as we assumed in class. (You will want to use a software package for this.)



12. (4.1, #4, #30) Recall from class that for any vector x , Ax is a linear combination of the columns of A and for any vector y , $y^\top A$ is a linear combination of the rows of A . Check this if you are not convinced.
- (a) Let A and B be two matrices such that AB is defined. What is the relationship between
- column space of AB and column space of A ?
 - row space of AB and row space of B ?
 - Using (i) and (ii) argue that $\text{rank}(AB) \leq \min\{\text{rank}(A), \text{rank}(B)\}$.
 - How can you use (iii) to see that if the columns of a matrix $A \in \mathbb{R}^{n \times k}$ are not linearly independent, then $A^\top A$ cannot be inverted? In class we showed the converse statement, namely, *if the columns of A are linearly independent then $A^\top A$ is invertible.*
- (b) Now suppose $AB = 0$. What is the relationship between
- $\text{Null}(A)$ and $\text{Col}(B)$?
 - $(\text{Col}(B))^\perp$ and $\text{Row}(A)$?
- (c) If $AB = 0$, can A and B be 3×3 matrices of rank 2?
- (d) Suppose $A \in \mathbb{R}^{3 \times 4}$ and $B \in \mathbb{R}^{4 \times 5}$ and $AB = 0$. Argue that $\text{rank}(A) + \text{rank}(B) \leq 4$.

13. ***How to limit clubs in your community.** Suppose there are n people in your community and the community leaders are feeling overwhelmed by the number of different clubs that are getting formed (for which they need to fund cookies and drinks each week). The leaders devise the following rules to limit the number of clubs that can be formed:

- (a) *Each club has to have an odd number of members.*
- (b) *Every two clubs must have an even number of members in common.*

Let's use linear algebra to argue that no more than n clubs can be formed under these rules. **Make a small example that you can keep using as you do the various parts of this question.** Always do small examples! Examples are very enlightening.

- (a) Let's call the members of the community $1, 2, \dots, n$ and the clubs C_1, \dots, C_m . Form the $m \times n$ matrix with rows indexed by clubs and columns by people as follows:

$$a_{ij} = \begin{cases} 1 & \text{if person } j \text{ is in club } C_i \\ 0 & \text{otherwise} \end{cases}$$

Write an inequality that relates $\text{rank}(A)$ and n .

- (b) Now compute AA^T which is an $m \times m$ matrix. Argue that the (i, k) entry of AA^T counts the number of people common to both club C_i and club C_k . In particular, the (i, i) entry counts the number of people in C_i .
- (c) Next we replace all the odd numbers you see in AA^T with 1 and all the even numbers with 0. (In mathematical language we are working in the field F_2 with two elements 0 and 1 where $1 + 1 = 0$. Or equivalently, we are computing mod 2. Don't worry if haven't seen this before.) After you have made these replacements, what is AA^T given the rules on clubs? What is $\text{rank}(AA^T)$? (Note that this is meant to be the rank of AA^T as a matrix with real number entries, not entries in the field F_2 . If this note confuses you then you can safely ignore it)
- (d) Using the result of problem 5 (a) iii), and the previous step, argue that $m \leq n$. In other words your community of n people cannot form more than n clubs under the rules.

Chapter 10

Orthogonality and Projections

This chapter marks the beginning of our march towards the singular value decomposition of a matrix. We will need all the linear algebra we have to understand how it works, and that begins with the notion of orthogonality. The reader should consider everything from here onward as something they will use the rest of the class.

10.1 Orthogonal Subspaces

We begin by introducing some notation.

Let $\mathbf{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}$, $\mathbf{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} \in \mathbb{R}^n$. We define the **dot product** of \mathbf{v} and \mathbf{w} to be

$$\mathbf{v}^\top \mathbf{w} = v_1 w_1 + \cdots + v_n w_n = \sum_{i=1}^n v_i w_i$$

We use this notion to further define vector norms as follows. The **norm** of \mathbf{v} is given by

$$\|\mathbf{v}\| = \sqrt{v_1^2 + \cdots + v_n^2} = \sqrt{\sum_{i=1}^n v_i^2} = \sqrt{\mathbf{v}^\top \mathbf{v}}$$

and the **square norm** or **norm squared** as

$$\|\mathbf{v}\|^2 = \mathbf{v}^\top \mathbf{v}$$

Definition 10.1.1. For any $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ we say that \mathbf{v} and \mathbf{w} are **orthogonal** if $\mathbf{v}^\top \mathbf{w} = 0$ or equivalently, if $\mathbf{w}^\top \mathbf{v} = 0$.

Using our new notation, we have a fun restatement of the Pythagorean theorem for \mathbb{R}^n .

Theorem 10.1.2. If $\mathbf{v}^\top \mathbf{w} = 0$ then

$$\|\mathbf{v}\|^2 + \|\mathbf{w}\|^2 = \|\mathbf{v} - \mathbf{w}\|^2 = \|\mathbf{v} + \mathbf{w}\|^2$$


Now that we have a notion of orthogonality amongst vectors, we can extrapolate and define the central concept of this chapter, namely orthogonal subspaces.

Definition 10.1.3. Let V and W be subspaces of \mathbb{R}^n . V and W are **orthogonal**, denoted $V \perp W$, if every $\mathbf{v} \in V$ is orthogonal to every $\mathbf{w} \in W$.

We can visualize some orthogonal subspaces, and we can't visualize some others. It will be important to have something to think when we encounter the orthogonal subspaces in greater generality.

Examples you can see


Example 10.1.4. Let $V = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right\}$ and $W = \text{Span} \left\{ \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}$.



We can see that these are orthogonal because every vector in V looks like $\begin{bmatrix} c \\ 0 \end{bmatrix}$ for some $c \in \mathbb{R}$ and any vector in W looks like $\begin{bmatrix} d \\ 0 \end{bmatrix}$ for some $d \in \mathbb{R}$. Taking the dot product of these arbitrary vectors is always zero, hence V and W are orthogonal subspaces. This is also clear from the picture and this is the intuition that we should always keep in mind.

In the previous example we orthogonal subspaces of equal dimension but this is not always the case.

Example 10.1.5. Let $V = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \right\}$ and $W = \text{Span} \left\{ \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\}$.



We can see that these are orthogonal because every vector in V looks like $\mathbf{v} = \begin{bmatrix} a \\ b \\ 0 \end{bmatrix}$ for some $a, b \in \mathbb{R}$ and

any vector in W looks like $\mathbf{w} = \begin{bmatrix} 0 \\ 0 \\ c \end{bmatrix}$ for some $c \in \mathbb{R}$. We can easily check that $\mathbf{w}^\top \mathbf{v} = 0$ so V and W are orthogonal subspaces.

Examples you can't see

Example 10.1.6. Let $V = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\}$ and $W = \text{Span} \left\{ \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\}$. You should check, by similar methods of the above example, that these are orthogonal subspaces.

For the last example, we illustrate orthogonal subspaces of a vector space that is not \mathbb{R}^n .

Example 10.1.7. Let $V = \{a_3x^3 + a_1x : a_1, a_3 \in \mathbb{R}\}$ and $W = \{a_2x^2 + a_0 : a_0, a_2 \in \mathbb{R}\}$. Note that we have $V, W \in \mathbb{R}[x]_3$. Using the association we have between vectors and polynomials, we can see that V and W are the same subspaces of the previous example, hence are orthogonal by the same reason. We certainly have no hope of obtaining a picture from either of these examples but we can still have a notion of orthogonality

Now that we have a feel for what orthogonal subspaces can look like, we can encounter the first (amazing) proposition concerning orthogonality and matrices.

Proposition 10.1.8. $\text{Row}(A) \perp \text{Null}(A)$ for any $A \in \mathbb{R}^{m \times n}$.

Proof. We begin by writing A in terms of its rows.

$$A = \begin{bmatrix} \mathbf{a}_1^\top \\ \vdots \\ \mathbf{a}_m^\top \end{bmatrix}$$

If the new notation is unfamiliar, please reference the notation section at the beginning of the notes. Recall that $\text{Row}(A) = \text{Span}\{\mathbf{a}_1^\top, \dots, \mathbf{a}_m^\top\}$ and $\text{Null}(A) = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{0}\}$.

Now, assume that $\mathbf{x} \in \text{Null}(A)$. We aim to show that $\mathbf{a}_i^\top \mathbf{x} = 0$ for all i . By looking at the entries of $A\mathbf{x}$ using the rows of A we see that

$$A\mathbf{x} = \begin{bmatrix} \mathbf{a}_1^\top \\ \vdots \\ \mathbf{a}_m^\top \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{a}_1^\top \mathbf{x} \\ \vdots \\ \mathbf{a}_m^\top \mathbf{x} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$$

Since two vectors are equal if and only if their entries are equal, we obtain the desired result. In other words, every vector in $\text{Row}(A)$ is orthogonal to every vector in $\text{Null}(A)$. □

There is another way of seeing why this proposition is true, and it involves another incredibly useful fact which we can think of as the analog of the fact that $\text{Col}(A) = \{A\mathbf{x} : \mathbf{x} \in \mathbb{R}^n\}$.

Alternative Definition of Row(A)

$$\text{Row}(A) = \mathbf{y}^\top A: \mathbf{y} \in \mathbb{R}^m\}$$

To see why this is true, think of $\text{Row}(A)$ as all linear combinations of the rows of A . Write $A = [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n] = (a_{ij})$ and $\mathbf{y}^\top = [y_1 \ \cdots \ y_n]$, then

$$\begin{aligned} \mathbf{y}^\top A &= \mathbf{y}^\top [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n] = [\mathbf{y}^\top \mathbf{a}_1 \ \cdots \ \mathbf{y}^\top \mathbf{a}_n] = \left[\mathbf{y}^\top \begin{bmatrix} a_{11} \\ a_{12} \\ \vdots \\ a_{1m} \end{bmatrix} \ \cdots \ \mathbf{y}^\top \begin{bmatrix} a_{n1} \\ a_{n2} \\ \vdots \\ a_{nm} \end{bmatrix} \right] \\ &= y_1 [a_{11} \ a_{21} \ \cdots \ a_{n1}] + y_2 [a_{12} \ a_{22} \ \cdots \ a_{n2}] + \cdots + y_n [a_{1n} \ a_{2n} \ \cdots \ a_{nn}] \end{aligned}$$

This last object is a linear combination of the rows of A .

Now, we can see the alternative proof of proposition 10.1.8 by using this fact. If $\mathbf{x} \in \text{Null}(A)$, we just need to compute the dot product of \mathbf{x} with an *arbitrary* element of $\text{Row}(A)$, and show that it is 0. Note that being an element of the row space means **you are a row vector already**, so there is no need to transpose. Let $\mathbf{r} \in \text{Row}(A)$. Since every element of $\text{Row}(A)$ looks like $\mathbf{y}^\top A$ for some vector $\mathbf{y} \in \mathbb{R}^m$, we have $\mathbf{r} = \mathbf{y}^\top A$ for some \mathbf{y} . Then

$$(\mathbf{r})\mathbf{x} = (\mathbf{y}^\top A)\mathbf{x} = \mathbf{y}^\top (A\mathbf{x}) = \mathbf{y}^\top (\mathbf{0}) = \mathbf{0}$$

Now you may be wondering if there is a nice analog of this proposition for $\text{Col}(A)$ and sure enough there is! The key to seeing this fact is to notice that $\text{Col}(A) = \text{Row}(A^\top)$ for any matrix A .

Proposition 10.1.9. $\text{Col}(A) \perp \text{Null}(A^\top)$ for any matrix $A \in \mathbb{R}^{n \times m}$.

Proof. Let $\mathbf{x} \in \text{Null}(A^\top)$ and write $A = [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n]$. Our goal is to show that $\mathbf{x}^\top \mathbf{a}_i = 0$ for all columns \mathbf{a}_i of A . By assumption, we have that $A^\top \mathbf{x} = \mathbf{0}$. Transposing both sides of this equation yields $\mathbf{x}^\top A = \mathbf{0}^\top$, but

$$\mathbf{x}^\top A = \mathbf{x}^\top [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n] = [\mathbf{x}^\top \mathbf{a}_1 \ \cdots \ \mathbf{x}^\top \mathbf{a}_n] = [0 \ \cdots \ 0]$$

This means that if $\mathbf{x} \in \text{Null}(A^\top)$ and $\mathbf{a} \in \text{Col}(A)$ then $\mathbf{x}^\top \mathbf{a} = 0$. □

In practice, there are nice ways of checking that two subspaces are orthogonal. As with many reductions in linear algebra, it is enough to verify that the basis vectors for the respective nullspaces are orthogonal. We can easily illustrate why with an almost generalized example.

Example 10.1.10. Let V and W be subspaces with respective bases given by $\mathcal{B}_V = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ and $\mathcal{B}_W = \{\mathbf{w}\}$. By the definition of basis, we know that any $\mathbf{v} \in V$ looks like

$$\mathbf{v} = a_1 \mathbf{v}_1 + \cdots + a_n \mathbf{v}_n$$

and similarly, any vector in W look like $a\mathbf{w}$ for some $a \in \mathbb{R}$. Computing this dot product we get

$$\mathbf{w}^\top \mathbf{v} = a_1 \mathbf{w}^\top \mathbf{v}_1 + a_2 \mathbf{w}^\top \mathbf{v}_2 + \cdots + a_n \mathbf{w}^\top \mathbf{v}_n = 0 \Leftrightarrow \mathbf{w}^\top \mathbf{v}_i = 0$$

for all $i = 1, \dots, n$. This means that any vectors in V and W are orthogonal if the bases for the respective subspaces are orthogonal.

In some texts, these four subspaces associated to any matrix are called the *four fundamental subspaces*. We are familiar with three of them and we call $\text{Null}(A^\top)$ the **left nullspace** of A .

Example 10.1.11. Let $A = \begin{bmatrix} 1 & 2 & 3 \\ -1 & 0 & 2 \end{bmatrix}$. We have $\text{Row}(A) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 2 \end{bmatrix} \right\}$. Before computing the Null space, we know by rank-nullity that it must be one-dimensional. Carrying out the usual computation we get that $\text{Null}(A) = \text{Span} \left\{ \begin{bmatrix} 2 \\ -5/2 \\ 1 \end{bmatrix} \right\}$. Computing the dot product of $\begin{bmatrix} 2 \\ -5/2 \\ 1 \end{bmatrix}$ with $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$ and $\begin{bmatrix} -1 \\ 0 \\ 2 \end{bmatrix}$ respectively, we see that we get 0 in both cases, verifying that $\text{Row}(A) \perp \text{Null}(A)$.

10.2 Orthogonal complements



The first two examples of this section provided scenarios where we had orthogonal subspaces of equal and different dimensions. In addition to having these properties, we also had that the sum of their respective dimensions was equal to the dimension of the ambient space, \mathbb{R}^2 and \mathbb{R}^3 respectively. This was a nice coincidence but we could have also used the following example.

Example 10.2.1. Let $V = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right\}$ and $W = \text{Span} \left\{ \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\}$.

In this example, the dimensions did not add up to 3. These provide an example of orthogonal subspaces that are not *complementary*. In contrast, complementary orthogonal subspaces allow us to decompose our entire space into orthogonal chunks, and as a result, will be of central interest to us when discussing orthogonality.

Definition 10.2.2. Let V be a subspace of \mathbb{R}^n . The **orthogonal complement** of V , denoted V^\perp is the subspace of all vectors orthogonal to V . That is

$$V^\perp = \{ \mathbf{y} \in \mathbb{R}^n : \mathbf{x}^\top \mathbf{y} = 0 \ \forall \mathbf{x} \in V \}$$



We have one giant example that encompasses this whole idea. The statement is short but packed with information so taking some times to digest it will be time well spent.

Example 10.2.3.

$$\text{Row}(A)^\perp = \text{Null}(A)$$

$$\text{Col}(A)^\perp = \text{Null}(A^\top)$$

Looking back at example 10.1.11, we see that in addition to the subspaces being orthogonal, they are also complementary. We dig deeper into this with a series of fundamental propositions and their proofs. In each statement, we assume that V is a subspace of \mathbb{R}^n

Proposition 10.2.4. $V \cap V^\perp = \{\mathbf{0}\}$. That is, the only vector in common to any subspace and its orthogonal complement is the zero subspace.

Proof. If $\mathbf{x} \in V \cap V^\perp$ and $\mathbf{x} \neq \mathbf{0}$ then $\mathbf{x}^\top \mathbf{x} = 0$ which implies $\|\mathbf{x}\| = 0$ hence $\mathbf{x} = \mathbf{0}$. □

Proposition 10.2.5. We can always find a matrix $A \in \mathbb{R}^{n \times n}$ such that $V^\perp = \text{Null}(A)$, or equivalently, that $V = \text{Row}(A)$.

Proof. Let $V = \text{Span}\{\mathbf{a}_1, \dots, \mathbf{a}_k\}$ and consider the matrix $A = \begin{bmatrix} \mathbf{a}_1^\top \\ \vdots \\ \mathbf{a}_k^\top \end{bmatrix}$. It is immediate that $V = \text{Row}(A)$

hence $V^\perp = \text{Null}(A)$. □

This proposition actually tells us that any pair of subspaces (V, V^\perp) is of the form $(\text{Row}(A), \text{Null}(A))$ for some matrix A .

Proposition 10.2.6. $\dim(V) + \dim(V^\perp) = n$

Proof. Suppose $\dim(V) = r$ and let $\mathcal{B}_V = \{\mathbf{a}_1, \dots, \mathbf{a}_r\}$ be a basis for V . By the previous proposition we know that $V = \text{Row}(A)$ where $A = \begin{bmatrix} \mathbf{a}_1^\top \\ \vdots \\ \mathbf{a}_r^\top \end{bmatrix}$, $\text{rank}(A) = r$, and $V^\perp = \text{Null}(A)$. It then follows by rank-nullity that $\dim(V^\perp) = \dim(\text{Null}(A)) = n - r$ thus $\dim(V) + \dim(V^\perp) = r + n - r = n$. □

Proposition 10.2.7. If $\mathcal{B}_V = \{\mathbf{a}_1, \dots, \mathbf{a}_r\}$ is a basis for V and $\mathcal{B}_{V^\perp} = \{\mathbf{a}'_1, \dots, \mathbf{a}'_{n-r}\}$ is a basis for V^\perp then

$$\mathcal{B} = \mathcal{B}_V \cup \mathcal{B}_{V^\perp} = \{\mathbf{a}_1, \dots, \mathbf{a}_r, \mathbf{a}'_1, \dots, \mathbf{a}'_{n-r}\}$$

is a basis for \mathbb{R}^n .

Proof. We know that \mathcal{B} has n elements, so we only need to show that they are linearly independent. Each set separately is linearly independent because we assumed that \mathcal{B}_V and \mathcal{B}_{V^\perp} were bases. We will now show that none of the \mathbf{a}'_i are in $V = \text{Span}\{\mathbf{a}_1, \dots, \mathbf{a}_r\}$ by assuming that they are and arriving at a contradiction.

Given an arbitrary $\mathbf{a}' \in V^\perp$, assume that $\mathbf{a}' \in V$. Then there exist scalars $c_1, \dots, c_r \in \mathbb{R}$ such that

$$\mathbf{a}' = c_1 \mathbf{a}_1 + \dots + c_r \mathbf{a}_r$$

Note that $\mathbf{a}'^\top \mathbf{a}_i = 0$ for all $i = 1, \dots, r$ by the assumption that $\mathbf{a}' \in V^\perp$. It follows that

$$\|\mathbf{a}'\|^2 = \mathbf{a}'^\top \mathbf{a}' = c_1 \underbrace{\mathbf{a}'^\top \mathbf{a}_1}_{=0} + c_2 \underbrace{\mathbf{a}'^\top \mathbf{a}_2}_{=0} + \dots + c_r \underbrace{\mathbf{a}'^\top \mathbf{a}_r}_{=0} = 0$$

thus $\mathbf{a}' = \mathbf{0}$. □

Proposition 10.2.8. Every $\mathbf{a} \in \mathbb{R}^n$ can be written uniquely as

$$\mathbf{a} = \underbrace{c_1 \mathbf{a}_1 + \dots + c_r \mathbf{a}_r}_{\mathbf{a}_V} + \underbrace{c'_1 \mathbf{a}'_1 + \dots + c'_{n-r} \mathbf{a}'_{n-r}}_{\mathbf{a}_{V^\perp}}$$

That is, every vector has the form $\mathbf{a} = \mathbf{a}_V + \mathbf{a}_{V^\perp}$ where $\mathbf{a}_V \in V$ and $\mathbf{a}_{V^\perp} \in V^\perp$.

Proof. This follows immediately from the previous proposition. □

The notion of writing any vector as a sum of vectors from V and V^\perp is an extremely useful tool for a number of reasons. It is also useful to think about the picture that we can associate to this idea.

Warning: We may be inclined to think proposition 10.2.7 implies that any $\mathbf{x} \in \mathbb{R}^n$ lives in either V or V^\perp but this is not the case. There is a fundamental difference in saying that $\mathbf{x} = \mathbf{a}_V + \mathbf{a}_{V^\perp}$ versus saying that \mathbf{x} in V or V^\perp . The following example will help us never forget this.

Example 10.2.9. Let $\mathbf{x} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$ with $V = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \right\}$ and $V^\perp = \text{Span} \left\{ \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\}$. We can write


$$\mathbf{x} = \underbrace{\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}}_{\mathbf{a}_V} + \underbrace{\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}}_{\mathbf{a}_{V^\perp}} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

but $\mathbf{x} \notin V$ and $\mathbf{x} \notin V^\perp$.

10.3 Projections

The notion of projecting onto a subspace is one of the most widely used concepts in many branches of mathematics. In the realm of linear algebra, it is centered around the following question.

Question 10.3.1. Given a subspace $V \subset \mathbb{R}^n$ and any vector $\mathbf{a} \in \mathbb{R}^n$, can one always find a linear map that projects \mathbf{a} onto V ? The following picture illustrates the idea




The purpose of this section is to describe the way in which we obtain this projection matrix given a subspace V . We then end with an example. We could approach the situation naively and use the idea from proposition 10.2.7. Given a fixed subspace V and an arbitrary vector $\mathbf{a} \in \mathbb{R}^n$, the algorithm would carry out as follows:

1. Compute a basis \mathcal{B}_V of V .
2. Compute a basis \mathcal{B}_{V^\perp} of V^\perp .
3. Write $\mathbf{a} = \mathbf{a}_V + \mathbf{a}_{V^\perp}$
4. The projection of \mathbf{a} into V is the vector \mathbf{a}_V , written as $\text{proj}_V \mathbf{a}$.

This method works but finding bases is time consuming and inefficient. Moreover, this method does not give us an explicit matrix that projects **any** vector onto V . Let's work methodically.

Suppose $V = \text{Span}\{\mathbf{a}_1, \dots, \mathbf{a}_k\} \subset \mathbb{R}^n$ and write $A = [\mathbf{a}_1 \ \dots \ \mathbf{a}_k] \in \mathbb{R}^{n \times k}$. Since $A\mathbf{x}$ is a linear combination of the columns of A for any $\mathbf{x} \in \mathbb{R}^k$, we know that $A\mathbf{x} \in V$ for all $\mathbf{x} \in \mathbb{R}^k$. In particular, there exists some $\hat{\mathbf{x}}$ such that $\mathbf{a}_V = A\hat{\mathbf{x}}$ and furthermore

$$\mathbf{a} - \mathbf{a}_V \perp V$$



From the fact that $\mathbf{a} - \mathbf{a}_V \perp V$ we can conclude the following.

$$\mathbf{a}_1^\top (\mathbf{a} - \mathbf{a}_V) = \mathbf{a}_2^\top (\mathbf{a} - \mathbf{a}_V) = \dots = \mathbf{a}_k^\top (\mathbf{a} - \mathbf{a}_V) = 0 \implies A^\top (\mathbf{a} - \mathbf{a}_V) = \begin{bmatrix} \mathbf{a}_1^\top \\ \vdots \\ \mathbf{a}_k^\top \end{bmatrix} (\mathbf{a} - \mathbf{a}_V) = \mathbf{0}$$

Since $\mathbf{a}_V = A\hat{\mathbf{x}}$ we can conclude that

$$A^\top (\mathbf{a} - A\hat{\mathbf{x}}) = \mathbf{0} \implies A^\top \mathbf{a} = A^\top A\hat{\mathbf{x}}$$

The next step is to investigate $A^\top A$ a little more. It is worth noting here that it is **always** a square matrix.

Proposition 10.3.2. *If $\mathbf{a}_1, \dots, \mathbf{a}_k$ are linearly independent, then $A^\top A$ is invertible.*

Proof. Since $A = [\mathbf{a}_1 \ \dots \ \mathbf{a}_k]$ has linearly independent columns, we know that $\text{Null}(A) = \{\mathbf{0}\}$, hence if $A\mathbf{x} = \mathbf{0}$ then $A^\top A\mathbf{x} = \mathbf{0}$ and $\mathbf{x} \in \text{Null}(A^\top A)$. Using this, we need to show that if $A^\top A\mathbf{y} = \mathbf{0}$ then $\mathbf{y} = \mathbf{0}$. This will imply (by the big theorem) that $A^\top A$ is invertible. Assuming that $A^\top A\mathbf{y} = \mathbf{0}$, we can multiply both sides (on the left) by \mathbf{y}^\top and we obtain the following string of implications

$$\mathbf{y}^\top A^\top A\mathbf{y} = \mathbf{0} \implies (A\mathbf{y})^\top = A\mathbf{y} = \mathbf{0} \implies \|A\mathbf{y}\|^2 = 0 \implies \|A\mathbf{y}\| = 0 \implies A\mathbf{y} = \mathbf{0}$$

This means that $\mathbf{y} \in \text{Null}(A)$ but $\text{Null}(A) = \{\mathbf{0}\}$ hence $\mathbf{y} = \mathbf{0}$. □

Now let's characterize projections in a nice way. We have obtained the equation $A^\top \mathbf{a} = A^\top A\hat{\mathbf{x}}$ with $\mathbf{a}_V = A\hat{\mathbf{x}}$ and we want to find a matrix M such that $M\mathbf{a} = A\hat{\mathbf{x}} = \mathbf{a}_V$, i.e. the projection matrix. Looking a little bit closer at $A^\top \mathbf{a} = A^\top A\hat{\mathbf{x}}$ we can see that the goal would be to first isolate the vector $\hat{\mathbf{x}}$, then multiply it on the left by A to obtain $A\hat{\mathbf{x}} = \mathbf{a}_V$. Carrying out the first step requires inverting $A^\top A$, which, by the previous proposition, requires A to have linearly independent columns. We then **reduce the columns of A to a basis for V** to ensure this is the case. Do not forget this.

Now that we can invert, we have

$$A^T \mathbf{a} = A^T A \hat{\mathbf{x}} \implies (A^T A)^{-1} A^T \mathbf{a} = \hat{\mathbf{x}}$$

and left multiplying by A we get

$$A(A^T A)^{-1} A^T \mathbf{a} = A \hat{\mathbf{x}} = \mathbf{a}_V = \text{proj}_V \mathbf{a}$$

Proposition 10.3.3. Let V be a subspace of \mathbb{R}^n with basis given by $\{\mathbf{a}_1, \dots, \mathbf{a}_k\}$ and let $A = [\mathbf{a}_1 \ \dots \ \mathbf{a}_k]$. Then projection onto V is given by the projection matrix

$$P = A(A^T A)^{-1} A^T$$

It is worth noting that any (symmetric) matrix that satisfies $P^2 = P$ is a projection onto some subspace and it is projection onto a proper subspace of 0 is an eigenvalue. Problem 4 of the problem set from chapter 2 further tells us that **all projection matrices are diagonalizable** which is a very nice fact.

We end the section with a much needed example.

Example 10.3.4. Lets project $\mathbf{a} = \begin{bmatrix} 3 \\ 4 \\ 4 \end{bmatrix}$ onto the line spanned by $\begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}$.

Let $V = \text{Span} \left\{ \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix} \right\}$. Going along with the procedure just described, we obtain the matrix $A = \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}$

and then compute $A^T A = \begin{bmatrix} 2 & 2 & 1 \end{bmatrix} \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix} = 9$. This is clearly an invertible (1×1) matrix. The last step is to plug into the formula for a projection matrix and multiply by \mathbf{a} to get $\text{proj}_V \mathbf{a}$.

$$\mathbf{a}_V = A(A^T A)^{-1} A^T \mathbf{a} = \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix} \left(\frac{1}{9}\right) \begin{bmatrix} 2 & 2 & 1 \end{bmatrix} \begin{bmatrix} 3 \\ 4 \\ 4 \end{bmatrix} = \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix} \left(\frac{1}{9}\right)(18) = \begin{bmatrix} 4 \\ 4 \\ 2 \end{bmatrix}$$

10.4 Applications: Least Squares Regression

Projections come up in many different places in linear algebra but one of the most widely used applications is that of least squares regression, otherwise known as linear regression. Using notions of projecting onto a subspace we can find the best fit line to a given set of data points.

We begin with the same setup and notation as we did in the previous section. That is, $V = \text{Span}\{\mathbf{a}_1, \dots, \mathbf{a}_k\} \subset \mathbb{R}^n$ with \mathbf{a}_i linearly independent and $A = [\mathbf{a}_1 \ \dots \ \mathbf{a}_k] \in \mathbb{R}^{n \times k}$. Given a vector $\mathbf{b} \in \mathbb{R}^n$ we denote its projection onto V with \mathbf{b}_V and we let the “error” be given by $\mathbf{e} = \mathbf{b} - \mathbf{b}_V$ (the terminology used here will make more sense in a bit). Note that $\mathbf{e} \perp V$.


Lastly, we have what is know as the *normal equations* for finding \mathbf{b}_V which are

$$A^T A \hat{\mathbf{x}} = A^T \mathbf{b} \quad \text{and} \quad \hat{\mathbf{x}} = (A^T A)^{-1} A^T \mathbf{b}$$

and we combine them to get the general equation

$$A \hat{\mathbf{x}} = \mathbf{b}_V = A(A^T A)^{-1} A^T \mathbf{b}$$

The central idea of the whole setup is the following picture, which should help you understand what is to follow.



Question 10.4.1. Suppose $\mathbf{b} \notin \text{Col}(A)$. What is the closest point in $\text{Col}(A)$ to \mathbf{b} ? (This is a minimization problem)

Since any point in $\text{Col}(A)$ looks like $A\mathbf{x}$ for some \mathbf{x} , the question is asking us to find $\mathbf{x} \in \mathbb{R}^k$ such that $\|A\mathbf{x} - \mathbf{b}\|^2$ is minimized.

General Idea: Since $\mathbf{e} = \mathbf{b} - \mathbf{b}_V$ we have $\mathbf{b} = \mathbf{b}_V + \mathbf{e}$ with $\mathbf{b}_V \in \text{Col}(A)$. Using this fact we can look more closely at the quantity that we want to minimize. We see that

$$\|A\mathbf{x} - \mathbf{b}\|^2 = \|A\mathbf{x} - \mathbf{b}_V - \mathbf{e}\|^2 = \|A\mathbf{x} - \mathbf{b}_V\|^2 + \|\mathbf{e}\|^2$$

The last equality follows from the Pythagorean theorem coupled with the fact that $A\mathbf{x} - \mathbf{b}_V \in \text{Col}(A)$ and $\mathbf{e} \perp \text{Col}(A)$.


Putting this all together, we can conclude that in order to minimize $\|A\mathbf{x} - \mathbf{b}\|^2$ we want to minimize $\|A\mathbf{x} - \mathbf{b}_V\|^2 + \|\mathbf{e}\|^2$. We can further simplify this problem by making two observations:

1. $\|\mathbf{e}\|$ cannot be changed since it is the perpendicular distance from \mathbf{b} to $\text{Col}(A)$
2. By setting $\mathbf{x} = \hat{\mathbf{x}}$ where $A\hat{\mathbf{x}} = \mathbf{b}_V$, we see that $\hat{\mathbf{x}}$ is the choice of \mathbf{x} that minimizes $\|A\mathbf{x} - \mathbf{b}\|^2$ because $A\hat{\mathbf{x}} = \mathbf{b}_V$ is the closest point to \mathbf{b} in $\text{Col}(A)$. In other words

$$\|A\hat{\mathbf{x}} - \mathbf{b}_V\|^2 = 0$$

These two observations are the key. Lets now see least squares in action before ending with a description of the general algorithm.

Example 10.4.2. Find the closest line to the points $(0, 6)$, $(1, 0)$, and $(2, 0)$



From the picture we can see that no line passes through all 3 points so we will need to find the best fit line. Lets call each point on this best fit line (t, b) so that b records the vertical height of a point.

The general line in the $t - b$ plane has the form $C + Dt$ where C and D are constants. If the line passed through the 3 points then

$$6 = C + D \cdot 0$$

$$0 = C + D \cdot 1$$

and

$$0 = C + D \cdot 2$$

would be a system of linear equations with a solution. In other words, the linear system

$$\underbrace{\begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{bmatrix}}_A \underbrace{\begin{bmatrix} C \\ D \end{bmatrix}}_b = \underbrace{\begin{bmatrix} 6 \\ 0 \\ 0 \end{bmatrix}}_b$$

would have a solution.

Since there is no such line, we know this system has no solution, i.e. $\mathbf{b} \notin \text{Col}(A)$ so we project \mathbf{b} onto $\text{Col}(A)$ via the projection formula. When we do this (using the normal equations) we get $\mathbf{b}_V = A\hat{\mathbf{x}}$ with $\hat{\mathbf{x}} = \begin{bmatrix} 5 \\ -3 \end{bmatrix}$ hence $\mathbf{b}_V = A\hat{\mathbf{x}} = \begin{bmatrix} 5 \\ 2 \\ -1 \end{bmatrix}$.

This means that the shortest distance $\|A\mathbf{x} - \mathbf{b}\|^2$ was obtained via the vector $\hat{\mathbf{x}} = \begin{bmatrix} 5 \\ -3 \end{bmatrix}$ thus we take the line to have equation $b = 5 - 3t$ by setting $\begin{bmatrix} C \\ D \end{bmatrix} = \begin{bmatrix} 5 \\ -3 \end{bmatrix}$. Therefore the points on the best fit line are of the form $(t, 5 - 3t)$.


Using the equation of the best fit line with the values of t we had initially ($t_1 = 0, t_2 = 1, t_3 = 2$), combined with the projection of \mathbf{b} onto the column space of A we have

$$\begin{bmatrix} 5 - 3t_1 \\ 5 - 3t_2 \\ 5 - 3t_3 \end{bmatrix} = \begin{bmatrix} 5 \\ 2 \\ -1 \end{bmatrix}$$

A few important things to observe here are

- The points on the best fit line are represented by (t_i, b_i) where b_i records a vertical distance.
- The vertical distance between a given data point and a **red point** on the best fit line above or below it is the value e_i . That is, the i^{th} entry of the error vector \mathbf{e} . This is because $\mathbf{b} - \mathbf{b}_V = \mathbf{e}$. In this example this translates to

$$\underbrace{\begin{bmatrix} 6 \\ 0 \\ 0 \end{bmatrix}}_{\mathbf{b}} - \underbrace{\begin{bmatrix} 5 \\ 2 \\ -1 \end{bmatrix}}_{\mathbf{b}_V} = \underbrace{\begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}}_{\mathbf{e}} = \begin{bmatrix} e_1 \\ e_2 \\ e_3 \end{bmatrix}$$



This line minimizes $\|\mathbf{e}\|^2 = e_1^2 + e_2^2 + e_3^2$.

Note that since $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \in \text{Col}(A)$ (this is always the case!) and $\mathbf{e} \perp \text{Col}(A)$, we **always** have $\mathbf{e}^T \mathbf{1} = e_1 + e_2 + e_3 = 0$.

We end the section with an outline of the general algorithm outlined in the example.

General method of least squares

Input: Data points $(t_1, b_1), (t_2, b_2), \dots, (t_m, b_m)$.

Output:

1. Set $\mathbf{b} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix}$

2. If there was a line $b = C + Dt$ through all points, then $\mathbf{b} \in \text{Col}(A)$ where

$$\underbrace{\begin{bmatrix} 1 & t_1 \\ 1 & t_2 \\ \vdots & \vdots \\ 1 & t_n \end{bmatrix}}_A \begin{bmatrix} C \\ D \end{bmatrix} = \underbrace{\begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix}}_{\mathbf{b}}$$

(a) If the line exists, solve for C and D to obtain the equation of the best fit line.

(b) If not, project \mathbf{b} onto $\text{Col}(A)$ to get $\mathbf{b}_V = A\hat{\mathbf{x}}$. Then set $\begin{bmatrix} C \\ D \end{bmatrix} = \hat{\mathbf{x}}$ to get a line L .

- The points on L are of the form $(t, C + Dt)$
- The points on the line above or below the data points are of the form $(t_i, C + Dt_i)$
- The values $e_i = b_i - (C + Dt_i)$ are the differences in height between the i^{th} data point and the point on L above or below it.
- L minimizes $\|\mathbf{e}\|^2 = e_1^2 + \cdots + e_m^2$.
- $e_1 + \cdots + e_m = 0$ because $\mathbf{1} \in \text{Col}(A)$ and $\mathbf{e} \perp \text{Col}(A)$.

10.5 Problem Set 3

1. (4.3 #1) Find the line that is closest to the following four points in the sense of Section 4.3:

$$(0, 0), (1, 8), (3, 8), (4, 20).$$

Show all steps of your work and plot the points and the line.

If $b = (0, 8, 8, 20)$ records the second coordinates of the given four points, and $p = (p_1, p_2, p_3, p_4)$ records the second coordinates of the points on the line you found, at times 0, 1, 3, 4, and $e = b - p$. What are the following?

- (a) p
- (b) $e_1^2 + e_2^2 + e_3^2 + e_4^2$
- (c) $e_1 + e_2 + e_3 + e_4$

2. (4.1 #3) Construct a matrix with the required property or say why you cannot:

(a) Column space contains $\begin{pmatrix} 1 \\ 2 \\ -3 \end{pmatrix}$ and $\begin{pmatrix} 2 \\ -3 \\ 5 \end{pmatrix}$, nullspace contains $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$.

(b) Row space contains $\begin{pmatrix} 1 \\ 2 \\ -3 \end{pmatrix}$ and $\begin{pmatrix} 2 \\ -3 \\ 5 \end{pmatrix}$, nullspace contains $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$.

(c) $Ax = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$ has a solution and $A^T \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$.

(d) Every row is orthogonal to every column but the matrix is not the zero matrix.

(e) Columns add to a column of zeros, rows add to a row of ones.

3. The following three parts are unrelated.

(a) Let

$$S = \text{Span} \left\{ \begin{pmatrix} 1 \\ 2 \\ 2 \\ 3 \end{pmatrix}, \begin{pmatrix} 1 \\ 3 \\ 3 \\ 2 \end{pmatrix} \right\}$$

Find a basis for S^\perp

(b) Let P be the following plane in \mathbb{R}^4

$$P = \left\{ \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} : x_1 + x_2 + x_3 + x_4 = 0 \right\}$$

Find a basis for P^\perp

(c) Let $P \in \mathbb{R}^{n \times n}$ be a matrix satisfying

- $P^2 = P$
- Every vector in $\text{Null}(P)$ is orthogonal to every vector in $\text{Range}(P)$.

P is a projection matrix. What subspace of \mathbb{R}^n does P project onto? Explain why your answer is true.

4. (4.1, #17)

(a) If S is a subspace of \mathbb{R}^3 containing only the origin, what is S^\perp ?

(b) If S is spanned by $(1, 1, 1)$, what is S^\perp ?

(c) Project $b = \begin{pmatrix} 2 \\ 3 \\ 4 \end{pmatrix}$ onto the S from (b).

(d) Project $b = \begin{pmatrix} 2 \\ 3 \\ 4 \end{pmatrix}$ onto the S^\perp from (b).

5. (4.2 #30)

- (a) Find the projection matrix P_C onto the column space of

$$A = \begin{bmatrix} 3 & 6 & 6 \\ 4 & 8 & 8 \end{bmatrix}.$$

Look carefully at this matrix before you start. Remember that $A^T A$ is invertible if and only if the columns of A are linearly independent.

- (b) Find the projection matrix P_R onto the row space of A .
(c) Compute $B = P_C A P_R$. Can you explain why B is the way it is?
6. (4.1, #6, #7) The following system $Ax = b$ has no solution. You can check if you like, or just believe me.

$$\begin{aligned} x + 2y + 2z &= 5 \\ 2x + 2y + 3z &= 5 \\ 3x + 4y + 5z &= 9 \end{aligned}$$

- (a) Find a vector $y = (y_1, y_2, y_3)$ such that $y^T A = 0$ and $y^T b \neq 0$.
(b) If you could produce such a y can you convince your boss that $Ax = b$ has no solution without solving the system? How would you do that?
(c) Which subspace associated to A did y come from?
(d) Is it true that whenever $Ax = b$ has no solution there will be a y such that $y^T A = 0$ and $y^T b \neq 0$? Explain.

Hint: If $Ax = b$ has no solution how does the echelon form of the augmented matrix $[A | b]$ look? Call this echelon form $[B | b']$. If $Bx = b'$ has no solution, can you easily produce a y such that $y^T B = 0$ and $y^T b' \neq 0$? What is the relationship between the solutions of $Ax = b$ and $Bx = b'$?

7. A shop with no small change**

Suppose you own an art supply shop that sells n different items (n is very large), and suppose m children have placed orders for the start of school (m is much smaller than n). Suddenly all coins of value less than \$1 go out of circulation, and you now need to round the prices of your items up or down. How can you round the prices so that the total price of each order is not affected too much?

We will show that the following is possible using linear algebra.

Theorem 10.5.1. *Suppose at most t items of each type has been ordered in total, and no order asks for more than one item of each type. Then it is possible to round the prices so that the total price of each order changes by no more than t dollars.*

Mathematical formulation:

- Suppose the price of item j is c_j . Note that we can assume $0 < c_j < 1$ for all j since only the rounding matters.
- Since each order contains only one of each item, we can represent an order S_i as a subset of $\{1, 2, 3, \dots, n\}$. For example, if $S_1 = \{2, 6, 9\}$ then child 1 has ordered one of item 2, item 6 and item 9.

- We are also told that for each j , item j is in no more than t sets among S_1, \dots, S_m . (*Note the role of t in the theorem. Don't forget what t is!*)
- The theorem says that we can find numbers $z_1, \dots, z_n \in \{0, 1\}$ (the rounded down prices for the n items) so that each order changes in price by at most t dollars. The change in price of order S_i is $|\sum_{j \in S_i} (c_j - z_j)|$. So we'll get

$$\left| \sum_{j \in S_i} (c_j - z_j) \right| \leq t \quad \text{for all } i = 1, 2, \dots, m \quad (10.5.1)$$

Running example: Suppose you carry $n = 7$ items in your shop, $m = 3$ children place orders, and no more than $t = 2$ items of each type are ordered in total. The three orders could be:

$$S_1 = \{1, 2, 3, 5, 7\}, S_2 = \{1, 2, 6, 7\}, S_3 = \{3, 4, 5, 6\}.$$

Check n, m, t on this example. Further suppose the costs of the 7 items are

$$c_1 = \frac{1}{2}, \quad c_2 = \frac{1}{4}, \quad c_3 = \frac{1}{4}, \quad c_4 = \frac{3}{4}, \quad c_5 = \frac{1}{2}, \quad c_6 = \frac{1}{4}, \quad c_7 = \frac{3}{4}.$$

By the theorem, we will be able to round (up or down) the 7 prices to z_1, \dots, z_7 each of which will be 0 or 1, so that each order changes by at most \$2 ($t = 2$).

- (a) Suppose each order contained at most s items. In our example, $s = 5$. Argue that we can easily round prices so that each order changes by at most s dollars.

So the interesting part about this theorem is that you can do much better if s is large, i.e., each order has lots of items, but t is small, i.e., the total number of erasers (or easels, or whatever) ordered is small. In our running example, it is easy to round prices so that no order changes by more than \$5 in price, but the cool thing is that we can get the change in price to be at most \$2. Note that the theorem doesn't care what the original prices c_j are.

The algorithm: The way to round prices is via the following *iterative algorithm*. This means we will repeat a procedure over and over again until we get what we want. Comments are in italics.

Initialize: For each item j , let x_j be a **floating variable** that is initially set to c_j .

*The following iterative method will move each floating x_j to 0 or 1 which then becomes the value of z_j (the rounded price). Once this happens, x_j is permanently set to z_j and we say x_j becomes **fixed**. In each step of our procedure, at least one floating variable will become fixed.*

Call S_i **dangerous** if it has more than t indices j for which x_j is still floating; the others sets are **safe**. In our running example, all sets are dangerous at the start and all variables are floating. We will always maintain the following equality:

$$\sum_{j \in S_i} x_j = \sum_{j \in S_i} c_j \quad \text{for all dangerous sets } S_i \quad (10.5.2)$$

At the start, all variables are floating and the above equation is true since $x_j = c_j$ for all j .

- Write down the equations (10.5.2) for all the currently dangerous sets. Think of this as a system of linear equations with the floating variables as unknowns and the fixed variables as constants. Find a solution of this system where at least one of the floating variables becomes 0 or 1. *We will argue later that this is always possible.*
- If x_j gets set to 0 or 1 then set z_j to the value of x_j . Declare x_j fixed. *Several x_j 's can get fixed at the same time.*

iii. If there are no more dangerous sets, then stop. Otherwise, write down the new system of linear equations (10.5.2) with all the fixed x_j 's turned into z_j 's and thought of as constants. Do not remove anything from the old system; simply replace x_j by its fixed value z_j in each equation. A fixed x_j is no longer a variable – it has a value — and the number of x_j variables have decreased from the old system to the new system. Go back to step (i).

(b) Run the above algorithm on our example and check that the theorem is true.

In the rest of this problem we argue that the algorithm always produces prices as stated in the theorem.

(c) Argue that at the start, the linear system in (i) is feasible.

You can say a bit more. Suppose F is the set of indices of the floating variables in the system and $|F|$ denotes the cardinality of F , i.e., the number of elements in F . In our example, at the start, $F = \{1, 2, 3, 4, 5, 6, 7\}$ and so $|F| = 7$. Argue that the system has a solution strictly inside the unit cube $[0, 1]^{|F|}$. **Hint:** this is a one line answer, just think about what the values of x_j are at the start.

By $[0, 1]^k$ we mean a cube with side lengths 1 and with opposite corners $(0, 0, \dots, 0)$ and $(1, 1, \dots, 1)$ in \mathbb{R}^k . See Figure 10.1 to see examples and what it means to be strictly inside the cube $[0, 1]^{|F|}$.




Figure 10.1: On the left you see the square $[0, 1]^2$ which is the unit cube in dimension 2. Next you see the boundary of the square in thick lines and a point on the boundary. Then you see the square without its boundary and a point strictly inside the square (i.e., not on the boundary). At the end is the unit cube $[0, 1]^3$, which is the usual cube. Its boundary consists of the 6 squares that form the outside of the cube. A point is strictly inside the cube if it does not lie on any of these 6 squares.

(d) Suppose now you are in some iteration of the algorithm.

i. Argue that there are fewer dangerous sets than floating variables in step (i). In our example, at the start, we have 3 dangerous sets and 7 floating variables at the start since all variables are floating at the start, and indeed $3 < 7$.

This is the trickiest part of the question, so let's break it down. Suppose there are f floating variables and d dangerous sets.

- A. Each floating variable can be in at most t dangerous sets. So argue that if we sum up the number of floating variables in each dangerous set, we get at most ft .
- B. Each dangerous set contains at least $t + 1$ floating variables. So if we again sum up the number of floating variables in each dangerous set, we get at least $d(t + 1)$.
- C. Write the inequality that relates ft and $d(t + 1)$ and conclude that $d < f$.

ii. Argue that the linear system obtained by taking all the equations (10.5.2) as you vary over dangerous sets S_i has a solution space of dimension at least one, i.e., contains a line l . **Hint:** Use the previous part and think about when linear systems of equations always have a solution.

- iii. Now argue that there is a solution to the linear system that lies on the boundary of the cube $[0, 1]^{|F|}$. **Hint:** use the line l to find such a point \mathbf{y} .
 - iv. Use the coordinates of \mathbf{y} as the values of the floating variables. Argue that at least one floating variable will get set to 0 or 1.
- (e) We keep going through iterations of the algorithm, each time setting up equations (10.5.2) and fixing some floating variables. Argue that after finitely many iterations, there will be no more dangerous sets and the algorithm will stop. How many iterations will be needed at worst?
- (f) To finish we need to argue that when there are no more dangerous sets, we will have the inequality (10.5.1).
- i. Consider an S_i . Argue that it satisfies (10.5.2) and has at most t floating variables.
 - ii. If there are at most t floating variables in each set S_i , can you conclude that we have (10.5.1)? **Hint:** Do you see a way to round the values of the remaining floating variables to 0 or 1 (at most t of them) so that $|\sum_{j \in S_i} (c_j - z_j)| \leq t$?

Chapter 11

Symmetric Matrices: The Most Important Matrices You'll Ever See

In this chapter we will define symmetric matrices and prove the all powerful Spectral theorem for real matrices. Once we are equipped with this theorem, we will dive deeper into the class of symmetric matrices, looking at positive definite and positive semi-definite matrices. We will finish by seeing a variety of applications.

11.1 Spectral Theorem

Definition 11.1.1. A matrix $A \in \mathbb{R}^{n \times n}$ is **symmetric** if $A = A^T$

With this definition in hand, we will prove a sequence of propositions that combine to give the statement of the Spectral theorem for real matrices. Note that there is also a spectral theorem for complex matrices that we will see at the end of the course. Before beginning our proof we recall that the **conjugate** of a complex number $z = a + bi$, denoted \bar{z} , is obtained by changing the sign of the imaginary part of z . That is

$$\bar{z} = \overline{a + bi} = a - bi$$

The complex conjugate satisfies a number of nice properties, the first of which is that

$$\overline{(a + bi)(c + di)} = \overline{(a + bi)}\overline{(c + di)}$$

We also note that the conjugate of a matrix, denoted \bar{A} , is obtained by conjugating all entries of A . These will be useful notions for us since eigenvalues of real matrices can be complex numbers, but a complex number z is in fact a real number if and only if $\bar{z} = z$. Similarly, if $A \in \mathbb{R}^{n \times n}$ then $\bar{A} = A$. These two facts will be needed to prove the first proposition.

Proposition 11.1.2. *If A is symmetric then all eigenvalues of A are real.*

Proof. Suppose that $A\mathbf{x} = \lambda\mathbf{x}$ where λ may be some complex number and \mathbf{x} may have some complex entries. Taking conjugates of both sides of our eigenvalue equation we get

$$\overline{A\mathbf{x}} = \overline{\lambda\mathbf{x}} \implies \bar{A}\bar{\mathbf{x}} = \bar{\lambda}\bar{\mathbf{x}} \implies A\bar{\mathbf{x}} = \bar{\lambda}\bar{\mathbf{x}}$$

We now transpose both sides of this equation. This is where we use the symmetric hypothesis!

$$(A\bar{\mathbf{x}})^T = (\bar{\lambda}\bar{\mathbf{x}})^T \implies \bar{\mathbf{x}}^T A^T = \bar{\mathbf{x}}^T \bar{\lambda} \implies \bar{\mathbf{x}}^T A = \bar{\mathbf{x}}^T \bar{\lambda}$$

Next, we take the two equations $A\mathbf{x} = \lambda\mathbf{x}$ and $\bar{\mathbf{x}}^T A = \bar{\mathbf{x}}^T \bar{\lambda}$ and multiply by $\bar{\mathbf{x}}^T$ on the left and \mathbf{x} on the right, respectively. In doing this we obtain

$$\bar{\mathbf{x}}^T A\mathbf{x} = \lambda\bar{\mathbf{x}}^T \mathbf{x}$$

from the first equation and

$$\bar{\mathbf{x}}^\top A\mathbf{x} = \bar{\lambda}\bar{\mathbf{x}}^\top \mathbf{x}$$

from the second. Combining these two equations we get that

$$\lambda\bar{\mathbf{x}}^\top \mathbf{x} = \bar{\lambda}\bar{\mathbf{x}}^\top \mathbf{x} \quad \text{which implies that} \quad \lambda\|\mathbf{x}\| = \bar{\lambda}\|\mathbf{x}\|$$

hence $\lambda = \bar{\lambda}$ and $\lambda \in \mathbb{R}$. □

We can also easily see that the eigenvectors are in fact in \mathbb{R}^n

Proposition 11.1.3. *If $A\mathbf{x} = \lambda\mathbf{x}$, where $\lambda \in \mathbb{R}$, then $\mathbf{x} \in \mathbb{R}^n$.*

Proof. If $A\mathbf{x} = \lambda\mathbf{x}$ then $(A - \lambda I)\mathbf{x} = \mathbf{0}$. Since $A - \lambda I$ is a real matrix, it has a real vector in its kernel, hence $\mathbf{x} \in \mathbb{R}^n$. □

We saw an example of these propositions in action on the first homework set.

Example 11.1.4. The matrix $A = \begin{bmatrix} 1 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 1 \end{bmatrix}$ is symmetric. It's eigenvalues are $\lambda = 0, 1, 3$ with respective eigenvectors $\mathbf{x}_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$, $\mathbf{x}_2 = \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}$, and $\mathbf{x}_3 = \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}$. Not only are the eigenvectors and eigenvalues real, but we can also see that

$$\mathbf{x}_1^\top \mathbf{x}_2 = \mathbf{x}_1^\top \mathbf{x}_3 = \mathbf{x}_2^\top \mathbf{x}_3 = 0$$

That is, the eigenvectors are mutually orthogonal. It turns out that this is always true.

Proposition 11.1.5. *All eigenvectors coming from different eigenspaces are mutually orthogonal. That is, if $\mathbf{x}_1 \in E_{\lambda_1}$ and $\mathbf{x}_2 \in E_{\lambda_2}$ with $\lambda_1 \neq \lambda_2$, then $\mathbf{x}_1^\top \mathbf{x}_2 = 0$.*

Proof. Suppose $A\mathbf{x} = \lambda_1\mathbf{x}$ and $A\mathbf{y} = \lambda_2\mathbf{y}$ with $\lambda_1 \neq \lambda_2$.

$$A\mathbf{x} = \lambda_1\mathbf{x} \implies (\lambda_1\mathbf{x})^\top = (A\mathbf{x})^\top \implies \mathbf{x}^\top \lambda_1 = \mathbf{x}^\top A^\top = \mathbf{x}^\top A$$

Multiplying both sides of this equation on the right by \mathbf{y} and using the fact that \mathbf{y} is an eigenvector we get

$$\mathbf{x}^\top \lambda_1 \mathbf{y} = \mathbf{x}^\top A \mathbf{y} = \mathbf{x}^\top \lambda_2 \mathbf{y} \implies \lambda_1 \mathbf{x}^\top \mathbf{y} = \lambda_2 \mathbf{x}^\top \mathbf{y} \implies (\lambda_1 - \lambda_2) \mathbf{x}^\top \mathbf{y} = 0$$

Since $\lambda_1 \neq \lambda_2$ we must have $\mathbf{x}^\top \mathbf{y} = 0$. □

Proposition 11.1.6. *For each eigenvalue λ we have $\text{AM}(\lambda) = \text{GM}(\lambda)$.*

Proof. We will prove this one later but can assume it for now. □

Before stating and proving the last proposition we will need some new terminology. The previous proposition tells us that any symmetric matrix admits n linearly independent eigenvectors, i.e. it is diagonalizable. If we combine this with proposition 11.1.5 we see that we can choose a special set of eigenvectors.

Definition 11.1.7. Let $S = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n\}$ with $\mathbf{u}_i \in \mathbb{R}^n$. S is a set of **orthonormal** vectors if

1. $\mathbf{u}_i^\top \mathbf{u}_j = 0 \forall i \neq j$
2. $\|\mathbf{u}_i\| = 1 \forall i$

Example 11.1.8. The set

$$S = \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\}$$

forms an orthonormal set because all vectors have length 1 and are mutually orthogonal.

In some linear algebra courses, extra time is taken to describe a nice algorithm known as the Gram-Schmidt algorithm. We will assume existence of this algorithm but will not go into the details of explaining it.

Gram-Schmidt Algorithm

Given any vector space V and a basis \mathcal{B} for V , a procedure, known as the Gram-Schmidt algorithm, can be applied to this basis. It takes the basis \mathcal{B} as input and outputs an orthonormal basis for V . We will use this algorithm in that if we are given a basis for a subspace, we can always obtain an orthonormal one for the same subspace. Now onto the proposition.

Proposition 11.1.9. *If A is symmetric, then A has n orthonormal eigenvectors, i.e. A admits an orthonormal basis for \mathbb{R}^n .*

Proof. By proposition 11.1.6, we know that each eigenspace E_λ , has maximum possible dimension, equal to $AM(\lambda)$. By applying the Gram-Schmidt procedure, we can obtain an orthonormal basis for each E_λ , lets call it

$$\mathcal{B}_{E_{\lambda_i}} = \{\mathbf{x}_{i_1}, \mathbf{x}_{i_2}, \dots, \mathbf{x}_{i_{AM(\lambda_i)}}\}$$

From proposition 11.1.5, we have that the basis vectors in $\mathcal{B}_{E_{\lambda_i}}$ are all orthogonal to the basis vectors in $\mathcal{B}_{E_{\lambda_j}}$ for all eigenvalues $\lambda_i \neq \lambda_j$. By putting all of these basis vectors together, we obtain n linearly independent vectors, all of whom are orthogonal and have unit length. That is, we obtain an orthonormal basis of eigenvectors for \mathbb{R}^n . \square

We can now state the Spectral theorem for real matrices.

Theorem 11.1.10. *If $A \in \mathbb{R}^{n \times n}$ is symmetric, then*

$$A = Q\Lambda Q^{-1} = Q\Lambda Q^\top$$

All eigenvalues of Λ are real and the columns of Q form an orthonormal basis of eigenvectors for \mathbb{R}^n .

Proof. The proof of this is a combination of the above propositions. Proposition 11.1.2 implies that the eigenvalues of A are all real. Proposition 11.1.6 and proposition 11.1.9 imply that A is diagonalizable by a matrix whose columns form an orthonormal basis. \square

A few important things must be noted here:

1. The diagonalization of a symmetric matrix is known as an orthogonal diagonalization and the matrix Q is known as an orthogonal matrix. It satisfies the property that $Q^\top = Q^{-1}$.
2. Orthogonal diagonalizations of symmetric matrices are unique up to reordering of the columns. This is a subtle but crucially important fact. It means that if A is a symmetric matrix and $A = Q_1\Lambda_1Q_1^\top = Q_2\Lambda_2Q_2^\top$, then by reordering the columns of Q_i and diagonal entries of Λ_i , we must have $Q_1 = Q_2$ and $\Lambda_1 = \Lambda_2$. There will be moments where we will encounter several diagonalizations of a symmetric matrix, and uniqueness will come to the rescue, allowing us to conclude that many of the matrices in question are in fact equal.
3. In practice, we will not have to apply Gram-Schmidt to obtain our desired orthonormal basis. The eigenvectors we get from our usual computation will already be orthogonal by proposition 11.1.5, so we will only need to normalize them in order to obtain an orthonormal basis.

For the sake of completeness, we define orthogonal matrices here, but we will not use them explicitly for another few sections.

Definition 11.1.11. A matrix $Q \in \mathbb{R}^{n \times n}$ is **orthogonal** if $Q^\top Q = QQ^\top = I$. That is, $Q^\top = Q^{-1}$.

We can alternatively define a matrix to be orthogonal if its columns form an orthonormal basis of \mathbb{R}^n . We can check that mutually orthogonal columns of unit length will imply that $Q^T Q = I$ as follows.

Let $Q = [\mathbf{u}_1 \ \cdots \ \mathbf{u}_n]$ so that

$$Q^T Q = \begin{bmatrix} \mathbf{u}_1^T \\ \vdots \\ \mathbf{u}_n^T \end{bmatrix} [\mathbf{u}_1 \ \cdots \ \mathbf{u}_n] = \begin{bmatrix} \|\mathbf{u}_1\|^2 & & \mathbf{u}_j^T \mathbf{u}_i \\ & \ddots & \\ \mathbf{u}_i^T \mathbf{u}_j & & \|\mathbf{u}_n\|^2 \end{bmatrix} = \begin{bmatrix} 1 & & 0 \\ & \ddots & \\ 0 & & 1 \end{bmatrix} = I_n$$

The definition of orthogonal matrix differs by author. Our definition includes the assumption that the columns are all of unit length, in addition to being orthogonal. Other texts sometimes refer to an orthogonal matrix as a matrix with orthogonal columns, not necessarily of unit length. Any future reference to the word *orthogonal matrix* is taken to mean the former, but if you read texts elsewhere it is good to be aware of both definitions.

Now lets look at the Spectral theorem in action.

Example 11.1.12. Let $A = \begin{bmatrix} 1 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 1 \end{bmatrix}$. The eigenvalues of this matrix are $\lambda_1 = 0, \lambda_2 = 1$, and

$\lambda_3 = 3$ with respective eigenvalues given by $\mathbf{x}_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$, $\mathbf{x}_2 = \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}$, and $\mathbf{x}_3 = \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}$. After normalizing

these vectors we get $\mathbf{x}_1 = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$, $\mathbf{x}_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}$, and $\mathbf{x}_3 = \frac{1}{\sqrt{6}} \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}$. This gives rise to the orthogonal diagonalization

$$A = Q \Lambda Q^T = \begin{bmatrix} 1/\sqrt{3} & 1/\sqrt{2} & 1/\sqrt{6} \\ 1/\sqrt{3} & 0 & -2/\sqrt{6} \\ 1/\sqrt{3} & -1/\sqrt{2} & 1/\sqrt{6} \end{bmatrix} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 3 \end{bmatrix} \begin{bmatrix} 1/\sqrt{3} & 1/\sqrt{3} & 1/\sqrt{3} \\ 1/\sqrt{2} & 0 & -1/\sqrt{2} \\ 1/\sqrt{6} & -2/\sqrt{6} & 1/\sqrt{6} \end{bmatrix}$$

Chapter 12

Change of Basis Review and Why Orthogonal Diagonalization is So Great

We all learned about change of basis in math 308 but some of it may have gone over our head. Even if it made perfect sense the first time, there's no harm in discussing it again. The main idea behind a change of basis is that it allows us to represent the same transformation geometrically but in a different, and sometimes more advantageous, coordinate system. We will begin with highlighting a few of the main facts of change of basis, but for a more thorough treatment of the subject, take a look at the notes for lecture 12 here

<http://www.samroven.com/linear>

12.1 Change of Basis

Let's first address notation. Let $\mathbf{x} = \begin{bmatrix} 3 \\ -2 \end{bmatrix} \in \mathbb{R}^2$ be written in the standard basis. The coordinates of \mathbf{x} are expressing its geometric location in the plane. That is, to arrive at the tip of the vector \mathbf{x} , you move 3 units to the right of the origin (3 units along \mathbf{e}_1) and -2 units down from there (-2 units along \mathbf{e}_2). This is because

$$\mathbf{x} = 3\mathbf{e}_1 - 2\mathbf{e}_2$$

The coefficients of \mathbf{x} in this expression involving the standard basis are what determine its coordinates. This is the general idea behind change of basis.

Let $\mathcal{B} = \left\{ \begin{bmatrix} 2 \\ 7 \end{bmatrix}, \begin{bmatrix} 1 \\ 4 \end{bmatrix} \right\}$ be a (non-standard) basis of \mathbb{R}^2 , In this basis we can express the same vector \mathbf{x} as

$$\mathbf{x} = 14 \begin{bmatrix} 2 \\ 7 \end{bmatrix} - 25 \begin{bmatrix} 1 \\ 4 \end{bmatrix}$$

and we express this notationally as

$$[\mathbf{x}]_{\mathcal{B}} = \begin{bmatrix} 14 \\ -25 \end{bmatrix}$$

With this idea in mind, we can now define this notion in greater generality.

Definition 12.1.1. Let $\mathcal{B} = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n\}$ be a basis for \mathbb{R}^n and let

$$\mathbf{y} = a_1\mathbf{u}_1 + a_2\mathbf{u}_2 + \dots + a_n\mathbf{u}_n$$

then the coordinate vector of \mathbf{y} with respect to the basis \mathcal{B} is

$$[\mathbf{y}]_{\mathcal{B}} = \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{bmatrix}$$

Let $U = [\mathbf{u}_1 \ \mathbf{u}_2 \ \cdots \ \mathbf{u}_n] \in \mathbb{R}^{n \times n}$. We call U the **change of basis matrix for the basis \mathcal{B}** . If \mathbf{y} is taken to be a vector written in the standard basis, then

$$U[\mathbf{y}]_{\mathcal{B}} = [\mathbf{u}_1 \ \cdots \ \mathbf{u}_n] \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} = a_1 \mathbf{u}_1 + a_2 \mathbf{u}_2 + \cdots + a_n \mathbf{u}_n = \mathbf{y}$$

This definition was the upshot when we learned this in the first linear algebra course and we could go both ways since the change of basis matrix is always invertible.

Proposition 12.1.2. *Let \mathbf{y} be expressed in the standard basis with \mathcal{B} a non-standard basis for \mathbb{R}^n . If U is the change of basis matrix for the basis \mathcal{B} then*

$$U[\mathbf{y}]_{\mathcal{B}} = \mathbf{y} \quad \text{and} \quad [\mathbf{y}]_{\mathcal{B}} = U^{-1}\mathbf{y}$$

Example 12.1.3. Continuing from example 11.1.12, we have $\mathcal{B} = \left\{ \begin{bmatrix} 2 \\ 7 \end{bmatrix}, \begin{bmatrix} 1 \\ 4 \end{bmatrix} \right\}$ and $\mathbf{x} = \begin{bmatrix} 3 \\ -2 \end{bmatrix}$. Going from the standard basis to this one we see that

$$[\mathbf{x}]_{\mathcal{B}} = \begin{bmatrix} 14 \\ -25 \end{bmatrix} = \begin{bmatrix} 2 & 1 \\ 7 & 4 \end{bmatrix}^{-1} \begin{bmatrix} 3 \\ -2 \end{bmatrix}$$

12.2 Linear Maps in Different Bases

Now that we have seen how to write vectors in different bases, we need to dig into how to write a linear map in a different basis. Recall that any linear transformation is, at first, always given in the standard basis of the domain and codomain. This is reflected in the formula

$$A = [T(\mathbf{e}_1) \ \cdots \ T(\mathbf{e}_n)]$$

Geometrically, this linear map may look horrible. It may be turning \mathbb{R}^n around on itself, making it very hard to see what it does geometrically. We can ask the natural question:

Question 12.2.1. Does there exist a basis in which this transformation looks as nice as possible?

When we say “as nice as possible” we mean to say that the linear map is just scaling along a set of coordinate axes. Any linear map that looks like this geometrically is given by some sort of diagonal matrix (you should convince yourself of this).

The answer to this question is yes **if** the matrix is diagonalizable. We will soon see that this is in some sense always true, but what we can see now is that this is true in the best way possible for symmetric matrices. We will dig into this shortly but let’s see why diagonalizable matrices always admit these “ideal” bases.

Suppose $A = Q\Lambda Q^{-1}$ with $\mathcal{B} = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ and $Q = [\mathbf{u}_1 \ \cdots \ \mathbf{u}_n]$. Note that \mathcal{B} is our eigenbasis for \mathbb{R}^n . Next, let’s take a vector $\mathbf{x} \in \mathbb{R}^n$ with coordinates $\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$ in the standard basis. Let $\mathbf{y} = [\mathbf{x}]_{\mathcal{B}} = Q^{-1}\mathbf{x}$.

In the basis \mathcal{B} , we have

$$\mathbf{x} = Q\mathbf{y}$$

Now, the fact that A is diagonalizable implies the following

$$A\mathbf{x} = Q\Lambda Q^{-1} \implies A\mathbf{x} = Q\Lambda(Q^{-1}\mathbf{x}) = Q\Lambda\mathbf{y}$$

We would like to look at the coordinates of the image vector $A\mathbf{x}$ in the basis \mathcal{B} . Since $A\mathbf{x} = Q\Lambda\mathbf{y}$ we apply Q^{-1} to both sides of this because it is this matrix that takes us from the standard basis to \mathcal{B} . Carrying out the computation we see that

$$Q^{-1}(A\mathbf{x}) = \Lambda\mathbf{y}$$

Let's summarize.

If we change basis from the standard basis to the eigenbasis \mathcal{B} , the effect of applying the linear transformation A is to send

$$[\mathbf{x}]_{\mathcal{B}} = \mathbf{y} \mapsto \Lambda\mathbf{y}$$

This is ideal since linear maps given by diagonal matrices are the easiest ones to understand. If

$$[\mathbf{x}]_{\mathcal{B}} = \mathbf{y} = \begin{bmatrix} a_1 \\ \cdots \\ a_n \end{bmatrix} \quad \text{and} \quad \Lambda = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix}$$

then the linear map sends \mathbf{y} to $\Lambda\mathbf{y} = \begin{bmatrix} \lambda_1 a_1 \\ \lambda_2 a_2 \\ \vdots \\ \lambda_n a_n \end{bmatrix}$. This just scales along new axes given by the eigenbasis.

This is where the beauty of symmetric matrices come into play. If we picture a generic diagonalizable matrix, and what its associated linear map looks like in the eigenbasis, we may encounter the fact that the eigenbasis vectors may not be orthogonal. While the matrix still looks nice, it could be nicer! This best case scenario would be if we scaled along **orthogonal** axes, and from the Spectral theorem, we know that this is always the case for symmetric matrices.

Let's see this idea in action on a low dimensional example.

Example 12.2.2. Consider the following (symmetric) matrix, along with its orthogonal diagonalization

$$A = \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix} = \underbrace{\begin{bmatrix} 2 & 1 \\ -1 & 2 \end{bmatrix}}_Q \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & 5 \end{bmatrix}}_{\Lambda} \underbrace{\left(\frac{1}{5}\right) \begin{bmatrix} 2 & -1 \\ 1 & 2 \end{bmatrix}}_{Q^{-1}}$$

Here, our eigenbasis is $\mathcal{B} = \left\{ \begin{bmatrix} 2 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ 2 \end{bmatrix} \right\}$ and we note that these basis vectors have not been normalized for the sake of having nicer entries.

Let $\mathbf{x} = \begin{bmatrix} 5 \\ 5 \end{bmatrix}$, then $A\mathbf{x} = \begin{bmatrix} 15 \\ 30 \end{bmatrix}$. Lets change the coordinates of \mathbf{x} and $A\mathbf{x}$ with respect to the eigenbasis \mathcal{B} . We get

$$[\mathbf{x}]_{\mathcal{B}} = \mathbf{y} = Q^{-1}\mathbf{x} = \frac{1}{5} \begin{bmatrix} 2 & -1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} 5 \\ 5 \end{bmatrix} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$$

and

$$[A\mathbf{x}]_{\mathcal{B}} = Q^{-1}A\mathbf{x} = \frac{1}{5} \begin{bmatrix} 2 & -1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} 15 \\ 20 \end{bmatrix} = \begin{bmatrix} 0 \\ 15 \end{bmatrix} = \Lambda\mathbf{y} = \begin{bmatrix} 0 & 0 \\ 0 & 5 \end{bmatrix} \begin{bmatrix} 1 \\ 3 \end{bmatrix}$$

Sure enough, in the eigenbasis, A is given by stretching along **new** coordinate axes, i.e. crushing the first coordinate of any vector to the origin and scaling the second by a factor of 5.

12.3 Cramer's Rule

While Cramer's rule does not have anything to do with change of basis, no second linear algebra course can be taught without a mention of it. It is a fun and interesting method of solving linear systems of equations that can be used to interpret solutions to linear systems in different ways.

Question 12.3.1. How did we use to solve $A\mathbf{x} = \mathbf{b}$?

We would set up an augmented matrix $[A|\mathbf{b}]$ and row reduce to find our solution. We saw at a later time, that if A was invertible, we could compute A^{-1} directly and use it to solve via

$$\mathbf{x} = A^{-1}\mathbf{b}$$

Most people (if not all) don't usually enjoy computing inverses of matrices, even small ones. If the given matrix of coefficients is invertible, Cramer's rule comes to the rescue and gives an explicit solution to the given system.

Theorem 12.3.2. (*Cramer's Rule*)

If $A \in \mathbb{R}^{n \times n}$ has non-zero determinant, then $A\mathbf{x} = \mathbf{b}$ is solved by computing determinants.

Let $\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$ and let B_j be the matrix obtained from A by replacing the j^{th} column of A with the vector $\mathbf{b} = \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix}$. The solution to the system is then given by

$$x_1 = \frac{\det B_1}{\det A}, x_2 = \frac{\det B_2}{\det A}, \dots, x_n = \frac{\det B_n}{\det A}$$

Example 12.3.3. Solve the linear system

$$3x_1 + 4x_2 = 2$$

$$5x_1 + 6x_2 = 4$$

This is represented via $A\mathbf{x} = \mathbf{b}$ where $A = \begin{bmatrix} 3 & 4 \\ 5 & 6 \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} 2 \\ 4 \end{bmatrix}$. Computing determinants of the three matrices needed to solve the system we get

$$\det(A) = \det \left(\begin{bmatrix} 3 & 4 \\ 5 & 6 \end{bmatrix} \right) = -2, \quad \det(B_1) = \det \left(\begin{bmatrix} 2 & 4 \\ 4 & 6 \end{bmatrix} \right) = -4, \quad \det(B_2) = \det \left(\begin{bmatrix} 3 & 2 \\ 5 & 4 \end{bmatrix} \right) = 2$$

so

$$x_1 = \frac{-4}{-2} = 2, \quad x_2 = \frac{2}{-2} = -1$$

Plugging these back in we see that it is the unique solution to the system.

Before ending the section we leave a few parting remarks about Cramer's Rule.

- If the give system is purely symbolic, then Cramer's rule allows us to write the solution explicitly in terms of the inputs.
- Since the solution is given in terms of the input parameters, we can see how solutions change as parameters change. This is essential in some branches of economics that analyze supply and demand.

12.4 Problem Set 4

1. (6.4 # 7)

(a) Diagonalize the following matrix with an *orthogonal* eigenvector matrix Q :

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

(b) Compute the coordinates y of the point $x = (1, 1, 1)$ in the eigenbasis given by the columns of Q .

(c) Compute the coordinates z of Ax in the eigenbasis given by the columns of Q .

(d) Check that $z = \Lambda y$.

2. (5.3 #4) **Cramer's Rule** Let $A = [\vec{a}_1 \ \vec{a}_2 \ \cdots \ \vec{a}_n] \in \mathbb{R}^{n \times n}$ be an invertible matrix. We will solve the general equation

$$\vec{b} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix} = A\vec{x} = x_1\vec{a}_1 + x_2\vec{a}_2 + \cdots + x_n\vec{a}_n \quad (12.4.1)$$

(a) Use properties of determinants to show that $x_1 = \frac{\det(B_1)}{\det(A)}$, where $B_1 = [\vec{b} \ \vec{a}_2 \ \cdots \ \vec{a}_n]$. We proceed in steps:

i. Recall the following facts about determinants, letting B be a matrix obtained from A via some row operation.

- If B is obtained by swapping rows of A , then $\det(B) = -\det(A)$.
- If B is obtained by multiplying a row of A by a constant c , then $\det(B) = c\det(A)$
- If B is obtained by adding a multiple of one row of A to another, then $\det(B) = \det(A)$.

Argue that the same properties hold true for column operations. (this is one of the rare instances where column operations actually don't mess things up)

ii. Find a relationship between $\det(B_1)$ and $\det(A)$ by making an appropriate substitution for $\vec{b} = x_1\vec{a}_1 + x_2\vec{a}_2 + \cdots + x_n\vec{a}_n$ in equation (1) and using part (i).

iii. Conclude that $x_1 = \frac{\det(B_1)}{\det(A)}$

(b) Use part (a) to argue that $x_i = \frac{\det(B_i)}{\det(A)}$ for all $i = 2, 3, \dots, n$, where $B_i = [\vec{a}_1 \ \vec{a}_2 \ \cdots \ \vec{a}_{i-1} \ \vec{b} \ \vec{a}_{i+1} \ \cdots \ \vec{a}_n]$.

3. Suppose $\vec{u} \in \mathbb{R}^n$ is a unit vector. This problem is about the matrix $H = I - 2\vec{u}\vec{u}^\top$

(a) Compute H^2 . Is H symmetric? Is H orthogonal? Explain your answers.

(b) Show that \vec{u} is an eigenvector of H and find its corresponding eigenvalue $\lambda_{\vec{u}}$.

(c) Let $\vec{v} \in \mathbb{R}^n$ be orthogonal to \vec{u} . Argue that \vec{v} is an eigenvector of H and find its corresponding eigenvalue $\lambda_{\vec{v}}$.

(d) What is the multiplicity of $\lambda_{\vec{v}}$? Is H diagonalizable? Why or why not?

(e) Let h_{ii} denote the diagonal entries of H . What is $\sum_{i=1}^n h_{ii}$? How does this compare to the sum of the eigenvalues?

- (f) Explain what the transformation H is doing geometrically and explain why? (This shouldn't be longer than one or two sentences)
4. (6.4 #27) Find all 2×2 matrices that are both symmetric and orthogonal. In each case what are the eigenvalues and what must the determinant of such matrices be?
5. (6.4 #8) Find all *orthogonal* matrices that diagonalize $\begin{bmatrix} 9 & 12 \\ 12 & 16 \end{bmatrix}$.
6. (6.4 #10)
- If $A^3 = 0$ then what are the eigenvalues of A ?
 - Find a matrix A that is not the zero matrix for which $A^3 = 0$.
 - Is there a symmetric matrix A that is not the zero matrix for which $A^3 = 0$?
7. We call a **real** matrix $A \in \mathbb{R}^{n \times n}$ *anti-symmetric* if $A^\top = -A$. Explain the following facts about A :
- Give an example of a 3×3 anti-symmetric matrix.
 - Argue that the diagonal elements of any anti-symmetric matrix must be 0.
 - Argue that $\vec{x}^\top A \vec{x} = 0$ for all $\vec{x} \in \mathbb{R}^n$ (**Hint**: Use the fact that $\vec{x}^\top A \vec{x}$ is a scalar so that $(\vec{x}^\top A \vec{x})^\top = \vec{x}^\top A \vec{x}$).
 - Recall that a complex number $z \in \mathbb{C}$ is of the form $z = a + ib$ where $a, b \in \mathbb{R}$, where $i = \sqrt{-1}$ satisfies the equations $i^2 = -1, i^3 = -i$, and $i^4 = 1$ (if you'd like, you may think of \mathbb{C} as a two-dimensional real vector space with basis given by $\{(1, 0), (0, i)\}$). The *complex conjugate* of a complex number is given by $\bar{z} = \overline{a + ib} = a - ib$ and satisfies the property that $\overline{z_1 z_2} = \bar{z}_1 \bar{z}_2$. We say that a complex number $z = a + ib$ is purely imaginary if $a = 0$, that is, z is of the form bi for some $b \in \mathbb{R}$. Using the notion of a complex conjugate, **we can say that a non-zero complex number z is purely imaginary if $z = -\bar{z}$** , this is the optimal definition for this question. Given a vector \mathbf{x} with possibly complex entries, we let $\bar{\mathbf{x}}$ denote the vector whose entries are conjugates of the entries of \mathbf{x} . For example if $\mathbf{x} = \begin{bmatrix} i \\ -i \\ 1 + 2i \end{bmatrix}$, then $\bar{\mathbf{x}} = \begin{bmatrix} -i \\ i \\ 1 - 2i \end{bmatrix}$. Just like we did for vectors with real entries, we denote the square norm of a complex vector by $\bar{\mathbf{x}}^\top \mathbf{x} = \|\mathbf{x}\|^2$. Checking this with our example we get

$$\bar{\mathbf{x}}^\top \mathbf{x} = \begin{bmatrix} -i & i & 1 - 2i \end{bmatrix} \begin{bmatrix} i \\ -i \\ 1 + 2i \end{bmatrix} = -i^2 - i^2 + (1 + 2i)(1 - 2i) = 7 = \|\mathbf{x}\|^2$$

We also have the exact same notions for matrices and vectors, that is $\overline{A\mathbf{x}} = \bar{A}\bar{\mathbf{x}}$. Now, onto the question:

Argue that the eigenvalues of a real anti-symmetric matrix are purely imaginary.

(**Hint:** You may want to look at the proof in the notes that real symmetric matrices have real eigenvalues for inspiration. Begin by considering the equation $A\mathbf{x} = \lambda\mathbf{x}$ where \mathbf{x} can have complex entries and λ is a complex number. First multiply both sides by $\bar{\mathbf{x}}^\top$ and use the fact that $\bar{\mathbf{x}}^\top \mathbf{y} = \bar{\mathbf{y}}^\top \mathbf{x}$ to get a statement involving the square norm of the eigenvector \mathbf{x} . Then, conjugate both sides of $A\mathbf{x} = \lambda\mathbf{x}$ to simplify the previous equation further. You will want to conclude that $\lambda = -\bar{\lambda}$.)



- (e) Use (d) to argue that if n is odd, then $\det(A) = 0$ and if n is even then $\det(A) \geq 0$. Give an example where A is invertible.

8. * The Petersen Puzzle

Before you attempt this problem, you should read the recent article about Ringel's conjecture in Quanta Magazine. It will help understand the problem statement and also show you the current day status of problems related to this one! The animations there illustrate the notion of tiling needed in this problem. <https://www.quantamagazine.org/mathematicians-prove-ringels-graph-theory-conjecture-20200219/>

Below you see two important graphs:

- On the left is the *Petersen graph* with 10 nodes and 15 edges.
- On the right is the *complete graph* on 10 vertices called K_{10} with 10 nodes and 45 edges (i.e., all possible edges among 10 nodes).



Each node in K_{10} has 9 edges incident to it while each node in the Petersen graph has 3 edges incident to it. So it is plausible that K_{10} can be covered perfectly (the technical word is *tiling*) by 3 Petersen graphs. This means that you can lay down three Petersens on K_{10} so that vertices go to vertices and each edge of K_{10} lies under an edge of exactly one of the three Petersens. In the exercise below we will use several things we have learned so far to argue that it is NOT possible to cover K_{10} with 3 Petersens.

Fact: The adjacency matrix of the Petersen graph has eigenvalue 1 with multiplicity 5. It does not have -3 as an eigenvalue.

- (a) If J_{10} is the 10×10 matrix with all entries equal to 1 and I_{10} is the 10×10 identity matrix, then argue that the adjacency matrix of K_{10} is $J_{10} - I_{10}$.
- (b) If there were three Petersen graphs called P, Q, R that cover K_{10} , and their adjacency matrices are A_P, A_Q, A_R , then argue that

$$A_P + A_R + A_Q = J_{10} - I_{10}.$$

- (c) Argue that the matrix $A_P - I_{10}$ has a 5-dimensional nullspace. **Hint:** look at the given multiplicity of 1 as an eigenvalue of A_P . You will need to use the Spectral theorem here!
- (d) Argue that the nullspace of $A_P - I_{10}$ is in the orthogonal complement of $\mathbf{1} = (1, 1, \dots, 1) \in \mathbb{R}^{10}$. **Hint:** Where must $\mathbf{1}$ lie for this to be true?
- (e) What is the dimension of the orthogonal complement of $\mathbf{1}$?
- (f) The above results are also true for $A_Q - I_{10}$ since Q is also a Petersen graph. Therefore, argue using dimensions of the subspaces you have looked at, that
- i. there is a non-zero vector \mathbf{w} in the intersection of $\text{nullspace}(A_P - I_{10})$ and $\text{nullspace}(A_Q - I_{10})$, and
 - ii. $\mathbf{1}^\top \mathbf{w} = 0$.
- Hint:** Can there be two subspaces in some n dimensional vector space of dimensions a and b that only intersect at the origin if $a + b > n$? Try some small examples in \mathbb{R}^2 and \mathbb{R}^3 to gain some intuition.
- (g) Now compute $A_R \mathbf{w}$ using the expression $A_R = (J_{10} - I_{10} - A_P - A_Q)$ from (b) and observe that -3 is an eigenvalue of A_R .
- (h) What can you conclude?

Chapter 13

Positive Definite/Positive Semi-Definite Matrices

We now examine a new class of matrices, which live inside of the space of all symmetric matrices. The positive definite and positive semi-definite matrices have a number of surprising applications to many fields of math, the main one of interest to us being Laplacians of graphs. The first observation of this chapter is surprising in itself.

13.1 Motivation and Definitions

The first nice observation is that to every symmetric matrix $A \in \mathbb{R}^{n \times n}$, we can associate a quadratic function in n variables. When $n = 2$ it can be easily seen

$$\begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} a & b \\ b & c \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} ax + by & bx + cy \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = ax^2 + bxy + bxy + cy^2 = ax^2 + 2bxy + cy^2$$

In general, we can build up any quadratic function this way. Given a symmetric matrix $A \in \mathbb{R}^{n \times n}$ and any

vector $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$ we obtain a quadratic function

$$q(x_1, x_2, \dots, x_n) = \mathbf{x}^\top A \mathbf{x}$$

in n variables. With a little thought one can see that given any quadratic function, we can find its associated symmetric matrix.

Example 13.1.1. Let $q(x, y) = 2x^2 - 15xy + 3y^2$. Then if $A = \begin{bmatrix} 2 & -15/2 \\ -15/2 & 3 \end{bmatrix}$ we have


$$\begin{bmatrix} x & y \end{bmatrix} A \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 2 & -15/2 \\ -15/2 & 3 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = q(x, y)$$

The behavior of this quadratic function depends on the symmetric matrix we started with and this dependence leads to a natural definition of positive definite (abbreviated PD) and positive semidefinite (abbreviated PSD) matrices.

Example 13.1.2. Let $A = \begin{bmatrix} 1 & -5 \\ -5 & 1 \end{bmatrix}$. The associated quadratic is

$$q(x, y) = \begin{bmatrix} x & y \end{bmatrix} A \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 1 & -5 \\ -5 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = x^2 - 10xy + y^2$$

The eigenvalues of this matrix are $\lambda = -4, 6$ and $\det(A) = -24$. A graph of the associated quadratic looks as follows




Note that the quadratic contains points with negative values and is not convex. Moreover, it has a negative eigenvalue and negative determinant.

Example 13.1.3. Let $A = \begin{bmatrix} 10 & -5 \\ -5 & 2 \end{bmatrix}$. The associated quadratic is

$$q(x, y) = \begin{bmatrix} x & y \end{bmatrix} A \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 10 & -5 \\ -5 & 2 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = 10x^2 - 10xy + 2y^2$$

The eigenvalues of this matrix are $\lambda = -0.403, 12$ and $\det(A) = -5$. A graph of the associated quadratic looks as follows




Note that the quadratic contains points with negative values and is also not convex. Moreover, it has a negative eigenvalue and negative determinant.

Example 13.1.4. Let $A = \begin{bmatrix} 100 & -5 \\ -5 & 20 \end{bmatrix}$. The associated quadratic is

$$q(x, y) = \begin{bmatrix} x & y \end{bmatrix} A \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 100 & -5 \\ -5 & 20 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = 100x^2 - 10xy + 20y^2$$

The eigenvalues of this matrix are $\lambda = 19.88, 100.31$ and $\det(A) = 1975$. A graph of the associated quadratic looks as follows




Note that the quadratic contains **no** points with negative values and **is** convex.. Moreover, it has all positive eigenvalues and a positive (and large) determinant.

Example 13.1.5. Let $A = \begin{bmatrix} 500 & -5 \\ -5 & 500 \end{bmatrix}$. The associated quadratic is

$$q(x, y) = \begin{bmatrix} x & y \end{bmatrix} A \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 500 & -5 \\ -5 & 500 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = 500x^2 - 10xy + 500y^2$$

The eigenvalues of this matrix are $\lambda = 495, 505$ and $\det(A) = 249975$. A graph of the associated quadratic looks as follows



Note that the quadratic contains **no** points with negative values and **is** convex. In some sense it is more positive than the previous example. Moreover, it has all positive eigenvalues and a positive (and huge) determinant.

The last two examples were special in the sense that $q(x, y) \geq 0$ for all x, y in the domain of the function. The distinguishing characteristics of the behavior of these functions allows us to break symmetric matrices up into smaller pieces, giving rise to the main notions of this chapter. Before defining positive (semi)definite matrices, we need a new definition.

Definition 13.1.6. Given a matrix $A \in \mathbb{R}^{n \times n}$, a **principal minor** of A is the determinant of the submatrix obtained by deleting any number of the same rows and columns. Along the same lines, a **leading principal minor** of A is a determinant of a sub matrix obtained by deleting the last $n - k$ rows and columns of A . This is denoted by D_k .

This definition involves a bit of unpacking so lets look at an example.

Example 13.1.7. Let $A = \begin{bmatrix} 1 & 4 & 6 \\ 4 & 2 & 1 \\ 6 & 1 & 6 \end{bmatrix}$. Given any $n \times n$ matrix, there are exactly n leading principal minors, one for each value of $1 \leq k \leq n$. Note that when we delete columns we can think of deleting them in order, going from the bottom right corner to the top left. The leading principal minors are

$$D_3 = \det \left(\begin{bmatrix} 1 & 4 & 6 \\ 4 & 2 & 1 \\ 6 & 1 & 6 \end{bmatrix} \right), \quad D_2 = \det \left(\begin{bmatrix} 1 & 4 \\ 4 & 2 \end{bmatrix} \right), \quad D_1 = 1$$

Computing the principal minors is more involved since there are many more of them. The leading principal minors are also principal minors but we exclude those here. We can categorize principal minors by the size of the submatrix we are taking the determinant of. There is only one 3×3 principal minor, namely the leading one that we already computed. There are several 2×2 principal minors. The first is obtained by deleting the first row and column of A and the second is obtained by deleting the second row and second column. They are

$$\det \left(\begin{bmatrix} 2 & 1 \\ 1 & 6 \end{bmatrix} \right) = 11, \quad \det \left(\begin{bmatrix} 1 & 7 \\ 6 & 6 \end{bmatrix} \right) = -36$$

There are two 1×1 principal minors, both of which are obtained by deleting two rows and columns. By deleting rows and columns 1 and 3 from A we obtain 2, and by deleting rows and columns 1 and 2 from A we obtain 6.

We can now define the namesake matrices of the chapter.

Definition 13.1.8. $A \in \mathbb{R}^{n \times n}$ is **positive semidefinite**, denoted $A \succeq 0$, if any of the following equivalent conditions are true:

- $\mathbf{x}^\top A \mathbf{x} \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$. Note that this is saying the associated quadratic has non-negative output for all values in the domain.
- All eigenvalues of A are non-negative.
- There exists some matrix B such that $A = B^\top B$. This definition is important and such a factorization of A is called a **Cholesky factorization**. Note that we could equivalently write $A = B B^\top$ as a Cholesky factorization so the two forms are equivalent.
- All principal minors are non-negative.

Definition 13.1.9. $A \in \mathbb{R}^{n \times n}$ is **positive definite**, denoted $A \succ 0$, if any of the following equivalent conditions are true:

- $\mathbf{x}^\top A \mathbf{x} > 0$ for all nonzero $\mathbf{x} \in \mathbb{R}^n$. Note that this is saying the associated quadratic has positive output for all values in the domain.

- All eigenvalues of A are positive.
- There exists some matrix B such that $A = B^T B$ with B having linearly independent columns. We note that Cholesky factorizations are not unique and there are many different sized matrices B that can admit such a factorization.
- All leading principal minors are positive.

A nice result of these definitions is that all positive (semi)definite matrices must be symmetric since the Cholesky factorization of a given matrix implies symmetry right away due to the fact that $(B^T B)^T = B^T B$

Example 13.1.10. Let $A = \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix}$. This is a psd matrix which we will verify with each equivalent definition.

•

$$q(x, y) = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = x^2 + 4xy + 4y^2$$

It is not immediately clear but a graph of this function will show that $q(x, y) \geq 0$ for all $\begin{bmatrix} x \\ y \end{bmatrix} \in \mathbb{R}^2$. This means it is a positive semidefinite matrix.

- The eigenvalues are 0 and 5 which are both non-negative.

•

$$\begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 \\ 2 \end{bmatrix}}_B \underbrace{\begin{bmatrix} 1 & 2 \end{bmatrix}}_{B^T}$$

Cholesky factorizations can be computed by inspection many times once you choose a rank of your matrix B . Picking B to be the simplest kind of rank 1 matrix here ends up yielding a factorization.

- The 1×1 principal minors are 1 and 4. The 2×2 principal minors are 0. All of these values are non-negative as we expected.

Example 13.1.11. Let $A = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix}$. This is a positive definite matrix which we will verify with each equivalent definition.

•

$$q(x, y, z) = \begin{bmatrix} x & y & z \end{bmatrix} \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = 2x^2 + 2y^2 + 2z^2 + 2xy + 2xz + 2yz$$

It is again not immediately clear but a graph of this function will show that $q(x, y, z) > 0$ for all $\begin{bmatrix} x \\ y \\ z \end{bmatrix} \in \mathbb{R}^3$. This means it is a positive definite matrix. In practice, checking this condition graphically is a bad idea and should not be done. We will see that this equivalent definition acts more as a useful tool for making argument involving positive (semi)definite matrices.

- The eigenvalues are 1, 1, and 4 which are all positive.

- One can show that

$$A = \underbrace{\begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{bmatrix}}_B \underbrace{\begin{bmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}}_{B^\top}$$

Note that B is invertible and has linearly independent columns.

- The leading principal minors are 2, 3, and 4 respectively, which are all positive, as expected.

Now that we have a feel for what these definitions entail, let's see why some of them are equivalent. We do it for the positive semidefinite case and leave out a proof of the equivalent definition involving minors.

Proposition 13.1.12. *A is positive semidefinite if and only if any of the following conditions are true*

1. $\mathbf{x}^\top A \mathbf{x} \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$.
2. All eigenvalues of A are non-negative.
3. There exists some matrix B such that $A = B^\top B$.

Proof. We first show that $\mathbf{x}^\top A \mathbf{x} \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$ if and only if all eigenvalues of A are non-negative. If $\mathbf{x}^\top A \mathbf{x} \geq 0 \forall \mathbf{x}$ and $A \mathbf{x} = \lambda \mathbf{x}$ then

$$\mathbf{x}^\top A \mathbf{x} = \mathbf{x}^\top \lambda \mathbf{x} = \lambda \mathbf{x}^\top \mathbf{x} = \lambda \|\mathbf{x}\|^2$$

We can then conclude that

$$\mathbf{x}^\top A \mathbf{x} \geq 0 \forall \mathbf{x} \in \mathbb{R}^n \Leftrightarrow \lambda \mathbf{x}^\top \mathbf{x} \geq 0 \forall \mathbf{x} \in \mathbb{R}^n \Leftrightarrow \lambda \geq 0$$

Next, we show that if $A = B^\top B$ then $\mathbf{x}^\top A \mathbf{x} \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$. Using the Cholesky factorization for A in our quadratic function we get that

$$\mathbf{x}^\top A \mathbf{x} = \mathbf{x}^\top B^\top B \mathbf{x} = (B \mathbf{x})^\top (B \mathbf{x}) = \|B \mathbf{x}\|^2 \geq 0 \quad \forall \mathbf{x}$$

Lastly, we assume that the quadratic function is non-negative and obtain a Cholesky factorization. Remember, anytime we are considering a psd matrix, it must be symmetric to begin with! This means that A is orthogonally diagonalizable by the spectral theorem hence we can write $A = Q \Lambda Q^\top$. Moreover, we

know that $\Lambda \geq 0$ by the previous equivalence that we just proved. Let $\sqrt{\Lambda} = \begin{bmatrix} \sqrt{\lambda_1} & & \\ & \ddots & \\ & & \sqrt{\lambda_n} \end{bmatrix}$ and set

$B = \sqrt{\Lambda} Q^\top$. We then see that


$$B^\top B = (\sqrt{\Lambda} Q^\top)^\top (\sqrt{\Lambda} Q^\top) = Q \sqrt{\Lambda} \sqrt{\Lambda} Q^\top = Q \Lambda Q^\top = A$$

□

Now that we are equipped with the machinery of these matrices, we can dig into some applications.

13.2 Laplacian of a Graph and a Peek into Spectral Gap Theory

Consider the following graph G



Definition 13.2.1. Let G be a graph with vertices indexed by i . The **degree of vertex i** , denoted d_i , is the number of edges incident (next to) vertex i .

With the notion of degree, we can construct a diagonal matrix D_G whose i^{th} diagonal entry is the degree d_i . And for the graph G above we have

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

Recall that the adjacency matrix for a graph G is given by $A_G = (a_{ij})$ where

$$a_{ij} = \begin{cases} 1 & \text{there exists an edge from } i \text{ to } j \\ 0 & \text{otherwise} \end{cases}$$

We can now define the Laplacian.

Definition 13.2.2. Given a graph G with diagonal matrix D_G as above and adjacency matrix A_G . The **Laplacian** of G is defined to be


$$L_G = D_G - A_G$$

With the graph given above we have

$$A_G = \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} \quad \text{and} \quad L_G = D_G - A_G = \begin{bmatrix} 2 & -1 & 0 & -1 \\ -1 & 3 & -1 & -1 \\ 0 & -1 & 1 & 0 \\ -1 & -1 & 0 & 2 \end{bmatrix}$$

L_G holds a tremendous amount of information about G . Before seeing why, we state some nice facts concerning L_G .

- $\lambda = 0$ is always an eigenvalue of L_G because the rows (by definition) always sum to 0. This means not only that 0 is an eigenvalue but it also means that $\mathbf{1}$ is always an eigenvector of eigenvalue 0.
- The more amazing fact is that L_G is always positive semidefinite. This is because we can always factor the Laplacian as $L_G = B_G B_G^T$ where B_G is the “**directed**” **node-edge incidence matrix** of G . B_G is computed using the following rules:
 - Label the rows of B_G according to the nodes of the graph G
 - Label the columns according to the edges of the graph with a direction. We label every edge as $\{ij\}$ where we ensure that $i < j$. Then we turn G into a directed graph by drawing arrows along our edges according to the rule that all arrows go **from i to j** for $i < j$. Turning G into a directed graph according to this rule looks like



- Let $B_G = (b_{k,ij})$ where the rows are indexed by the nodes, k , and the columns are indexed by the edges, ij , we then have

$$b_{k,ij} = \begin{cases} 1 & k = i \\ -1 & k = j \\ 0 & \text{otherwise} \end{cases}$$

With our graph G , the corresponding directed node-edge incidence matrix is

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

One of the beautiful facts about this matrix is that it always satisfies the equation

$$L_G = B_G B_G^\top$$

We may safely assume this fact without proof.

The last fact we state as a proposition.

Proposition 13.2.3. *Let E be the set of edges of G . Then if $\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$ we always have*

$$\mathbf{x}^\top L_G \mathbf{x} = \sum_{\{ij\} \in E} (x_i - x_j)^2$$

Proof. We prove this just with our example but from this one case we will be able to see how a general argument could follow.

The first observation is to use the Cholesky factorization of L_G to express the quadratic function as a sum of squares. That is

$$\mathbf{x}^\top L_G \mathbf{x} = \mathbf{x}^\top B_G B_G^\top \mathbf{x}$$

By multiplying out $\mathbf{x}^\top B_G$ and $B_G^\top \mathbf{x}$ we can see that each edge (column of B_G) gives rise to exactly one entry of 1 and one entry of -1 . When we multiply $\mathbf{x}^\top B_G$ the resulting row vector contains an entry of the form $(x_i - x_j)$ for each instance of ± 1 in the column corresponding to the edge $\{ij\}$. We get the same thing for $B_G^\top \mathbf{x}$. Looking at the example we can see this in action

$$\begin{aligned} (\mathbf{x}^\top B_G)(B_G^\top \mathbf{x}) &= [x_1 \quad x_2 \quad x_3 \quad x_4] \underbrace{\begin{bmatrix} 1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & -1 \end{bmatrix}}_{B_G} \underbrace{\begin{bmatrix} 1 & -1 & 0 & 0 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{bmatrix}}_{B_G^\top} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} \\ &= [(x_1 - x_2) \quad (x_1 - x_4) \quad (x_2 - x_3) \quad (x_2 - x_4)] \begin{bmatrix} (x_1 - x_2) \\ (x_1 - x_4) \\ (x_2 - x_3) \\ (x_2 - x_4) \end{bmatrix} \\ &= (x_1 - x_2)^2 + (x_1 - x_4)^2 + (x_2 - x_3)^2 + (x_2 - x_4)^2 = \sum_{\{ij\} \in E} (x_i - x_j)^2 \end{aligned}$$

□


We can now summarize all of these facts with a theorem.

Theorem 13.2.4. *Let G be a graph with n vertices and let L_G denote the Laplacian of G . Let $0 = \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \dots \leq \lambda_n$ denote the eigenvalues of L_G .*

1. L_G is positive semi-definite.
2. $\lambda_1 = 0$ is an eigenvalue with $\mathbf{1}$ as an eigenvector.
3. $\lambda_2 > 0$ if and only if G is connected (you can get from one vertex to any other via a sequence of edges)

Proof. 1 follows from the Cholesky factorization of L_G and 2 follows from the fact that the rows of L_G sum to 0. We illustrate the third fact on a different example, from which the general case should be clear. The following is an example of a disconnected graph and we show that the graph is disconnected if and only if the second eigenvalue is 0.

Consider the following graph G



Computing the Laplacian of this graph, we see that it breaks up into blocks, with zero blocks coming from the lack of connectivity between connected components of the graph.

$$L_G = \left[\begin{array}{ccc|ccc} 2 & -1 & -1 & & & \\ -1 & 2 & -1 & & & \\ -1 & -1 & 2 & & & \\ \hline & & & \mathbf{0} & & \\ & & & & 2 & -1 & -1 \\ & & & & -1 & 2 & -1 \\ & & & & -1 & -1 & 2 \end{array} \right]$$

From here, we can explicitly find two linearly independent eigenvectors by leveraging the fact that the individual Laplacians of the connected subgraphs of G have $\mathbf{1}$ as an eigenvector of eigenvalue zero. From this we get eigenvectors of eigenvalue zero being

$$\begin{bmatrix} 1 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 1 \end{bmatrix}$$

This tells us that $\text{AM}(0) \geq 2$, hence $\lambda_2 = 0$. This means that if G is disconnected, then $\lambda_2 = 0$.

If G is connected, we need to show that $\text{AM}(0) = 1$. This will imply that $\lambda_2 > 0$ from the fact that L_G is positive semidefinite. Let \mathbf{x} be an eigenvector of eigenvalue 0. The fact that $L_G \mathbf{x} = \mathbf{0}$ implies that

$$0 = \mathbf{x}^\top L_G \mathbf{x} = \sum_{\{ij\} \in E} (x_i - x_j)^2$$

Since each summand, $(x_i - x_j)^2$ is non-negative, it must be true that $x_i = x_j$. In other words, $x_i = x_j$ for

all $\{ij\} \in E$. Since G is connected, every instance of $1 \leq i, j \leq n$ occurs, so with $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$ this means that

$$\mathbf{x} = c \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} = c \mathbf{1}$$

This means that the only eigenvectors of eigenvalue 0 are multiples of $\mathbf{1}$, hence $\text{GM}(0) = \text{AM}(0) = 1$ and $\lambda_2 > 0$. \square

λ_2 is known as the **spectral gap**. It measures the connectivity of the graph G . The larger λ_2 is, the more connected G is, and when λ_2 is “as big as possible”, G is complete. The notion of the spectral gap is extremely useful in clustering algorithms which we will soon see in homework. For a very nice article on clustering take a look at

<https://towardsdatascience.com/spectral-clustering-aba2640c0d5b>

13.3 Application: Distance Realization

We end the chapter with one last application of positive semidefinite matrices.


Question 13.3.1. Are there three points $p, q, r \in \mathbb{R}^2$ such that

$$\|p - q\| = \|q - r\| = 1 \quad \text{and} \quad \|p - r\| = 3?$$

The answer is no! The reason is because of the triangle inequality which states

$$\|p - r\| \leq \|p - q\| + \|q - r\| \implies 3 \leq 2$$

We could investigate the same sort of question in higher dimensions. That is, are there four points $p, q, r, s \in \mathbb{R}^3$ with the following distances?



The answer is again no! However, the triangle inequality presents no problems this time. The following theorem, known as Schoenberg’s theorem (1935), is the main tool for answering these kinds of questions.

Theorem 13.3.2. Given “distances” $d_{ij} \geq 0$ for $i, j = 0, 1, \dots, n$ and $d_{ii} = 0$ for all i , there exist points $p_0, p_1, \dots, p_n \in \mathbb{R}^n$ with $\|p_i - p_j\| = d_{ij}$ if and only if the matrix $M = (M_{ij})$, where $M_{ij} = \frac{1}{2}(d_{i0}^2 + d_{0j}^2 - d_{ij}^2)$, is positive semidefinite. If such points exist we say the points are **realizable**

Before proving the theorem, we illustrate its use on the example with four points above. The trick is computing the matrix M . First, we must construct a matrix whose entries are the prescribed distances. The rows and columns of this matrix are **indexed starting from 0 to n** so the d_{0i} and d_{j0} entry make sense. For the example above, the distance matrix is

$$D = \begin{bmatrix} 0 & 2 & 3 & 2 \\ 2 & 0 & 2 & 3 \\ 3 & 2 & 0 & 2 \\ 2 & 3 & 2 & 0 \end{bmatrix}$$

From here, we label the entries of the matrix $D = (d_{ij})$ and compute the entries of the matrix $M = (M_{ij})$ via the formula

$$M_{ij} = \frac{1}{2}(d_{i0}^2 + d_{0j}^2 - d_{ij}^2)$$

The matrix M in this example is

$$M = \begin{bmatrix} 8 & 9 & -1 \\ 9 & 18 & 9 \\ -1 & 9 & 8 \end{bmatrix}$$

A quick computation shows that $\det(M) < 0$ hence M is not psd and the points are not realizable.

Now let's prove the theorem. We will need to use the cosine theorem as a lemma.

Lemma 13.3.3. *Given $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$*

$$\|\mathbf{x} - \mathbf{y}\|^2 = \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - 2\mathbf{x}^\top \mathbf{y}$$

always holds. Note that this can alternatively be written as

$$\mathbf{x}^\top \mathbf{y} = \frac{1}{2}(\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - \|\mathbf{x} - \mathbf{y}\|^2)$$

We omit the proof of the cosine theorem but use the result to now prove Shoenbergs's theorem.

Proof. We first assume that the distances are realizable and aim to show that $M \succeq 0$. If the points are realizable then there exist points $p_0, p_2, \dots, p_n \in \mathbb{R}^n$ such that

$$\|p_i - p_j\| = d_{ij}$$

We can shift everything to the origin and assume that $p_0 = \mathbf{0}$. Next, set $\mathbf{x}_i = p_i - p_0$ for $i = 1, \dots, n$. This implies that $\|\mathbf{x}_i\| = d_{i0} = d_{0i}$, $\|\mathbf{x}_j\| = d_{0j} = d_{j0}$, and $d_{ij} = \|p_i - p_j\| = \|\mathbf{x}_i - \mathbf{x}_j\|$. By the cosine theorem, we have

$$\mathbf{x}_i^\top \mathbf{x}_j = \frac{1}{2}(\|\mathbf{x}_i\|^2 + \|\mathbf{x}_j\|^2 - \|\mathbf{x}_i - \mathbf{x}_j\|^2) = \frac{1}{2}(d_{i0}^2 + d_{0j}^2 - d_{ij}^2) = M_{ij}$$

Upon closer inspection we can see that

$$M = (M_{ij}) = (\mathbf{x}_i^\top \mathbf{x}_j) = \begin{bmatrix} \mathbf{x}_1^\top \\ \mathbf{x}_2^\top \\ \vdots \\ \mathbf{x}_n^\top \end{bmatrix} [\mathbf{x}_1 \quad \mathbf{x}_2 \quad \cdots \quad \mathbf{x}_n] = B^\top B$$

hence M is positive semidefinite.

The other direction follows similarly. If M is positive semi definite, then $M = B^\top B$ for some matrix $B \in \mathbb{R}^{n \times k}$. Set $p_0 = \mathbf{0}$ and set $\mathbf{p}_i =$ the i^{th} column of B for $i = 1, \dots, n$. Then we have $M_{ij} = \mathbf{p}_i^\top \mathbf{p}_j$. Working backwards from the first half of the proof we get that the cosine theorem holds for all i, j and can conclude that the diastances are realizable. \square

Next we are off to the SVD!

13.4 Problem Set 5

Many of the questions below are based on the four equivalent definitions of a matrix being positive semidefinite (psd) or positive definite (pd). In each case, there is usually one definition that will be the most efficient for what you need to do.

1. (All parts are unrelated)

(a) Is the following matrix positive semidefinite?

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

(b) (6.5 #12) Find all values of d for which the following matrix is positive definite.

$$\begin{bmatrix} 1 & 2 & 3 \\ 2 & d & 4 \\ 3 & 4 & 5 \end{bmatrix}$$

(c) Argue that the only orthogonal, symmetric, and positive definite matrix is the identity.

2. (6.5 # 14)

(a) If S is positive definite, is S invertible? If S is invertible, is S^{-1} also positive definite?

(b) If S is a positive semidefinite matrix is S always invertible? If not provide an example.

3. Use Shoenberg's theorem to decide if the following distances are realizable. By this I mean, are there four points p_0, p_1, p_2, p_3 in some dimension such that the (i, j) entry in the following matrix is the distance between p_i and p_j ?

$$\begin{bmatrix} 0 & 1 & 3 & 2 \\ 1 & 0 & \sqrt{10} & \sqrt{5} \\ 3 & \sqrt{10} & 0 & \sqrt{13} \\ 2 & \sqrt{5} & \sqrt{13} & 0 \end{bmatrix}$$

If the distances are realizable, find four points that realize them (Look at the proof of Shoenberg's theorem if you need to find the points).

4. (a) Let $q(x, y) = x^2 + 4xy + 9y^2$.

i. Write down the symmetric matrix Q such that $q(x, y) = \begin{pmatrix} x & y \end{pmatrix} Q \begin{pmatrix} x \\ y \end{pmatrix}$.

ii. Why is $q(x, y) \geq 0$ for all $(x, y) \in \mathbb{R}^2$?

iii. Show that $q(x, y) \geq 0$ for all $(x, y) \in \mathbb{R}^2$ by writing it as a sum of squares by factoring Q as BB^\top .

- (b) Find a quadratic function in x, y that is going to be negative for some (x, y) . What is your strategy?
5. (a) Argue that the set of all $n \times n$ symmetric matrices forms a subspace of $\mathbb{R}^{n \times n}$.
- (b) Recall that psd matrices are symmetric. Show that the set of all $n \times n$ psd matrices DOES NOT form a subspace in the space of all $n \times n$ symmetric matrices.
- (c) However psd matrices do have structure. Argue that for $n \times n$ matrices,
- i. if $A \succeq 0$ and $B \succeq 0$, then $A + B \succeq 0$,
 - ii. if $A \succeq 0$ and $\lambda \geq 0$ then $\lambda A \succeq 0$.

6. Given any psd matrix A , we can associate it with its determinant. Since a psd matrix must have all non-negative minors, it follows that psd matrices must have non-negative determinant. We can therefore associate any psd matrix to a non-negative real number $r \in [0, \infty)$. We have known for years that any positive real number r has two square roots, namely $\pm\sqrt{r}$, and \sqrt{r} is the unique positive number whose square is r . In this problem, we will see that the corresponding notion of “positive number” for matrices, is **positive semi-definite**. Given a matrix $A \in \mathbb{R}^{n \times n}$, we say a matrix B is a *square root* of A if $B^2 = A$. First, some examples:

- (a) Consider the linear transformations $T, S : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ given by

$$T\left(\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}\right) = \begin{bmatrix} x_3 \\ 0 \\ 0 \end{bmatrix} \quad \text{and} \quad S\left(\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}\right) = \begin{bmatrix} x_2 \\ x_3 \\ 0 \end{bmatrix}$$

Show that S is a square root of T .


- (b) Find the square roots of the matrix

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

How many did you find? **Hint:** Compute an orthogonal diagonalization of A and look back in the psd notes to see how one shows that the existence of a Cholesky factorization is an equivalent definition of psd.

- (c) Let $A \in \mathbb{R}^{n \times n}$ be an arbitrary positive semi-definite matrix. Argue that A has a square root. **Hint:** Mirror what you did in part (c) in a general setting. Are the square roots of a psd matrix always psd? Why or why not? (**Note:** A matrix can have many square roots and the last question is asking if **all** square roots of a given psd matrix are psd)
- (d) Find infinitely many square roots of I_2 . Can you find them all? (You aren't required to find them all. Just finding infinitely many is sufficient)
7. We always order the eigenvalues of the Laplacian of a graph on n vertices as $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. This is possible since all eigenvalues are real.

Consider the following graph H :




- (a) Write down the Laplacian L_H of the graph H .
- (b) Factor L_H as BB^T .
- (c) Compute the quadratic function $x^T L_H x$ and write it as a sum of squares.
- (d) Find a basis for the eigenspace of L_H corresponding to the eigenvalue 0.
- (e) Based on the above which is the first eigenvalue of L_H that is going to be positive?
- (f) Do you see a connection between the arithmetic multiplicity of the eigenvalue 0 and the number of connected components in a graph?

8. Spectral Clustering - Part 1

The second smallest eigenvalue λ_2 of the Laplacian L_G of a graph G is called the *Fiedler value* of G , and its eigenvector \mathbf{w} is called the *Fiedler vector* of G . We saw that G is connected if and only if $\lambda_2 \neq 0$. In this exercise we will see an application of the Fiedler vector \mathbf{w} .

An important task in data science is to find *clusters* in a graph. By a cluster we mean a group of vertices that are relatively well connected amongst themselves but not so well connected to the rest of the vertices. For example, suppose G is a social network graph with people as vertices and an edge between two people who know each other. Then some natural clusters might be all people who belong to the same church, or soccer club, or do Tai Chi etc. Someone from church may also know someone who does Tai Chi, but perhaps there are only a few such pairs. Knowing clusters in graphs allows one to understand how information or infection might spread in that network. Advertisers use cluster information to target similar ads to people in a given cluster. If you bought a particular knee support for soccer, then chances are that your soccer friends might also buy it, where as your church friends may not. In this exercise we will use linear algebra to find clusters in a graph.

Running example: The graph below is a reproduction of the example from <https://towardsdatascience.com/spectral-clustering> with vertices relabeled as $1, \dots, 10$.



This graph has two obvious clusters in it, the cluster of vertices $\{4, 5, 6, 7, 8\}$ and the cluster of vertices $\{1, 2, 3, 9, 10\}$. There is only one edge between these two groups while vertices in a group have more connections among them.

How do we find clusters in large complicated graphs? To talk about clusters, we use the math terminology of *cuts* in graphs.

Definition 13.4.1. Let G be a graph with vertex set $V = \{1, 2, 3, \dots, n\}$ and edge set E consisting of pairs of vertices $\{i, j\}$. If $A \subseteq V$ is a subset of vertices, we use the notation $V \setminus A$ for the vertices in V that are not in A . Also, $|A|$ denotes the cardinality of the set A , which is the number of elements in A .

In our example, $V = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$ and $E = \{\{1, 2\}, \{1, 3\}, \{2, 3\}, \dots\}$. If $A = \{4, 5, 6, 7, 8\}$, then $|A| = 5$, $V \setminus A = \{1, 2, 3, 9, 10\}$ and $|V \setminus A| = 5$.

- A **cut** in G is a partition of V into two sets A and $V \setminus A$ for some subset of vertices $A \subset V$. This is sometimes called the *cut induced by A* .
- Let $E(A, V \setminus A)$ be the edges that go between the vertices in A and $V \setminus A$. These are the edges holding A and $V \setminus A$ together in G . In our example, the cut induced by $A = \{4, 5, 6, 7, 8\}$ has $E(A, V \setminus A) = \{\{1, 6\}\}$.


Then the **density** of the cut induced by A is

$$\phi(A, V \setminus A) = n \cdot \frac{|E(A, V \setminus A)|}{|A| \cdot |V \setminus A|}$$

For $A = \{1, 3, 4, 5, 6, 7\}$ (different from the A we considered previously), $E(A, V \setminus A) = \{\{1, 2\}, \{2, 3\}, \{1, 9\}, \{1, 10\}\}$, and $\phi(A, V \setminus A) = 10 \cdot \frac{6}{6 \cdot 4} = \frac{5}{2}$. Note that there are 6 edges between A and $V \setminus A$ but there could have been $6 \cdot 4 = 24$ between A and $V \setminus A$ (which would have been the case if G was the complete graph K_{10}). So $\frac{6}{6 \cdot 4}$ is the ratio of the number of edges connecting A and $V \setminus A$ in this graph and in the complete graph. The cut density is the product of this ratio and the total number of vertices, 10.

- (c) Let ϕ_G denote the smallest possible density of a cut in G . We call a cut with density ϕ_G , the *sparsest cut* in G . What is the density of the cut induced by $A = \{4, 5, 6, 7, 8\}$? Would you guess that this is the sparsest cut in our example graph?

The questions below will be based on the following graph H :



Q1 Calculate the density of the cut in H induced by $A = \{1, 2, 3, 4\}$.

Q2 Below is an algorithm that uses the Fiedler vector \mathbf{w} to break a graph into its two main clusters. Since it uses eigenvalues and vectors, the method is called *spectral clustering*.

- i. Sort the components of \mathbf{w} in descending order so that $w_{i_1} \geq w_{i_2} \geq \dots \geq w_{i_n}$. In our running example, Julia says that $\lambda_2 = 0.2984$ and the Fiedler vector $\mathbf{w} = (0.23, 0.33, 0.33, -0.33, -0.33, -0.23, -0.33, -0.33, -0.23, -0.33)$. One way to sort the components of \mathbf{w} is as

$$w_2 \geq w_3 \geq w_9 \geq w_{10} \geq w_1 \geq w_6 \geq w_4 \geq w_5 \geq w_7 \geq w_8.$$

When two components tie in value, you can break the tie as you wish.

- ii. Let $A_k := \{i_1, \dots, i_k\}$ for $k = 1, \dots, n - 1$. Among the cuts $(A_k, V \setminus A_k)$, output the one with the smallest density. In the first example, $A_1 = \{2\}$, $A_2 = \{2, 3\}$, $A_3 = \{2, 3, 9\}$, $A_4 = \{2, 3, 9, 10\}$ etc. The cut $(A_4, V \setminus A_4)$ has density $10 \frac{4}{4 \cdot 6} = \frac{5}{3}$. In our running example, what is A_5 ?

Calculate the Fiedler value λ_2 and its eigenvector \mathbf{w} for H using Julia (I would suggest just finding the Laplacian of the graph H and typing it directly into Julia. Coding graphs into Julia is more involved). Then sort the components of \mathbf{w} and find all the sets A_1, \dots, A_8 . Pick one A_i from your list such that the density of the cut it induces has a chance to be sparser than the cut induced by $\{1, 2, 3, 4\}$. Compute the density of this cut.

The following theorem says how well our algorithm can do.

Theorem 13.4.2. *The following hold for G :*

- (a) $\phi_G \geq \lambda_2$.
- (b) *The above algorithm always finds a cut of density at most $4\sqrt{d_G \lambda_2}$ where d_G is the largest degree of a vertex in G .*

Q3 Argue that this theorem is saying that $\lambda_2 \leq \phi_G \leq 4\sqrt{d_G\lambda_2}$. (**Caution:** The sparsest cut output by the algorithm may not be the overall sparsest cut in G . How does the smallest cut density obtained from the algorithm compare to ϕ_G ?)

What are these bounds for ϕ_G in H ? What are these bounds for our running example?

Call a vector \mathbf{x} *non-constant* if it is not a multiple of $\mathbf{1} = (1, 1, \dots, 1)$. For $\mathbf{x} = (x_1, \dots, x_n)$, define the function

$$Q(\mathbf{x}) = n \cdot \frac{\sum_{\{i,j\} \in E(G)} (x_i - x_j)^2}{\sum_{1 \leq i < j \leq n} (x_i - x_j)^2}$$

Note that the numerator is $\mathbf{x}^\top L_G \mathbf{x}$ and the denominator is the same quadratic function for the complete graph on n vertices, i.e., $\mathbf{x}^\top L_{K_n} \mathbf{x}$.

Q4 For a subset $A \subseteq V$, let \mathbf{c}_A be the vector in \mathbb{R}^n with i th coordinate equal to 1 if $i \in A$ and 0 otherwise. This is called the *characteristic vector* of A . In our running example, $\mathbf{c}_{\{4,5,6,7,8\}} = (0, 0, 0, 1, 1, 1, 1, 1, 0, 0)$, and $Q((0, 0, 0, 1, 1, 1, 1, 1, 0, 0)) = \frac{2}{5} = \phi(\{4, 5, 6, 7, 8\}, \{0, 1, 2, 3, 9, 10\})$.

Compute the characteristic vector \mathbf{c}_A for $A = \{1, 2, 3, 4\}$ in H , and check that for this A , $Q(\mathbf{c}_A)$ is exactly the density of the cut induced by A .

Q5 Argue that ϕ_G is the minimum of $Q(\mathbf{x})$ as \mathbf{x} varies over all \mathbf{c}_A . If you were to use this method to find ϕ_G how many \mathbf{c}_A 's would you need to compute in the graph G shown above? Hopefully you see that this is not a very practical way to find ϕ_G .

9. **Spectral Clustering - Part 2*** In the remaining exercises, we will prove the first part of the theorem which says that the smallest density of a cut in G is at least as big as the Fiedler value.

Instead of finding the minimum of $Q(\mathbf{x})$ over all \mathbf{c}_A 's, let's be less ambitious and *relax* the problem and minimize $Q(\mathbf{x})$ over all non-constant vectors \mathbf{x} . **Call the minimum value μ .** This procedure is called a relaxation since now we are enlarging the allowable \mathbf{x} we can plug into Q from characteristic vectors \mathbf{c}_A to all non-constant vectors. This means the region we are minimizing over has become larger, more relaxed.

Q6 Argue that $\mu \leq \phi_G$. **Hint:** if we minimize over a larger set than what we are supposed to, what happens to the minimum value of Q ? Does it become smaller or bigger?

We'll now show that $\mu = \lambda_2$ completing the first part of the theorem.

Q7 Show that $Q(\mathbf{x}) = Q(\mathbf{x} + t\mathbf{1})$ for all $t \in \mathbb{R}$. This means that if we start at \mathbf{x} and move in direction $\mathbf{1}$ or $-\mathbf{1}$, the value of Q does not change.

Q8 Therefore, argue that

$$\mu = \min\{Q(\mathbf{x}) : \mathbf{x} \in \mathbb{R}^n \setminus \{\mathbf{0}\}, \mathbf{1}^\top \mathbf{x} = 0\}.$$

This is saying that μ is the minimum value of $Q(\mathbf{x})$ among the non-zero vectors \mathbf{x} in the orthogonal complement of $\mathbf{1}$. (**Hint:** Any $\mathbf{x} \in \mathbb{R}^n$ can be written as the sum of an element of $\text{span}(\mathbf{1})$ and $\mathbf{1}^\perp$. Apply Q to this sum and use the previous result.)

Q9 We are now going to think about the minimization problem in **Q8**.

- i. Argue that the Laplacian of the complete graph K_n is $nI_n - J_n$ where J_n is the $n \times n$ matrix filled with 1.
- ii. Use the fact that $\mathbf{1}^\top \mathbf{x} = 0$ to argue that $J_n \mathbf{x} = \mathbf{0}$.
- iii. Now show that the denominator of $Q(\mathbf{x})$ is $n\|\mathbf{x}\|^2$.

- iv. Therefore, what is the relationship between $Q(\mathbf{x})$ and $Q(\alpha\mathbf{x})$ for any non-zero $\alpha \in \mathbb{R}$?
 v. Using the above, argue that

$$\mu = \min\{\mathbf{x}^\top L_G \mathbf{x} : \|\mathbf{x}\| = 1, \mathbf{1}^\top \mathbf{x} = 0\}$$

Q10 To finish we'll argue that $\mu = \lambda_2$. Note that this is giving you a new interpretation of λ_2 .

Let $\mathbf{v}_1, \dots, \mathbf{v}_n$ be a set of orthonormal eigenvectors of L_G corresponding to the eigenvalues $\lambda_1 = 0, \lambda_2, \dots, \lambda_n$ in increasing order. Recall that \mathbf{v}_1 is a multiple of $\mathbf{1}$.

- i. Show that if $\mathbf{1}^\top \mathbf{x} = 0$, then \mathbf{x} is a linear combination of $\mathbf{v}_2, \dots, \mathbf{v}_n$.
 ii. Since $\|\mathbf{x}\| = 1$, argue that $\mathbf{x} = \sum_{i=2}^n \alpha_i \mathbf{v}_i$ with $\sum_{i=2}^n \alpha_i^2 = 1$.
 iii. Using $\mathbf{x} = \sum_{i=2}^n \alpha_i \mathbf{v}_i$, expand $\mathbf{x}^\top L_G \mathbf{x}$ and get that

$$\mathbf{x}^\top L_G \mathbf{x} = \alpha_2^2 \lambda_2 + \alpha_3^2 \lambda_3 + \dots + \alpha_n^2 \lambda_n.$$

iv. Now argue that $\alpha_2^2 \lambda_2 + \alpha_3^2 \lambda_3 + \dots + \alpha_n^2 \lambda_n \geq (\sum_{i=2}^n \alpha_i^2) \lambda_2$, and from this conclude that $\lambda_2 \leq \mu$.

v. Now show that $\mu \leq \lambda_2$.

Hint: how can you choose $\alpha_2, \alpha_3, \dots, \alpha_n$ that will give you an \mathbf{x} for which $\mathbf{x}^\top L_G \mathbf{x} = \lambda_2$?

vi. Conclude that $\lambda_2 = \mu$.

Chapter 14

The Singular Value Decomposition

We now spend a considerable amount of our efforts understanding the singular value decomposition of a matrix. In doing so, we will see that everything we have learned up to this point will be necessary. We will begin with some motivation, followed by the mechanics and inner workings of the definition. After that, we move onto the geometry of the singular value decomposition before finishing off with an introduction to matrix norms and rank one approximations to a matrix.

14.1 Motivation

We motivate the singular value decomposition with one geometric fact.


Given any $A \in \mathbb{R}^{m \times n}$ the image of the unit n -sphere under A is a hyperellipse

Let's try to understand this statement. We let

$$\mathbb{S}^{n-1} = \{\mathbf{x} \in \mathbb{R}^n : x_1^2 + x_2^2 + \cdots + x_n^2\}$$

denote the unit sphere in \mathbb{R}^n , otherwise known as the unit n sphere. When we say hyperellipse, we mean an m -dimensional generalization of an ellipse. It is defined as the surface obtained by stretching the unit sphere in \mathbb{R}^m by some factors $\sigma_1, \sigma_2, \dots, \sigma_m$ (some possibly zero) in orthogonal directions $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m \in \mathbb{R}^m$.

We can see what this looks like in low dimensional cases. The unit sphere in \mathbb{R}^2 (this is the unit 2-sphere which we denote by \mathbb{S}^1) is just the unit circle and a hyperellipse in \mathbb{R}^2 is just a usual ellipse. If $A \in \mathbb{R}^{2 \times 2}$ has full rank, the picture associated to A is



Note that

$$\Sigma^T \Sigma = \begin{bmatrix} \sigma_1^2 & 0 & \dots & 0 \\ 0 & \sigma_2^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix} \in \mathbb{R}^{n \times n}$$

Since $A^T A$ is symmetric, we must have $V \Sigma^T \Sigma V^T$ as its orthogonal diagonalization. This allows us to conclude two things (from uniqueness of orthogonal diagonalizations).

1. $\Sigma^T \Sigma$ is the diagonal matrix of eigenvalues of $A^T A$.
2. V is the orthogonal matrix whose columns are the eigenvectors of $A^T A$.

Combining these two facts tells us that

$$A^T A \mathbf{v}_i = \sigma_i^2 \mathbf{v}_i$$

We can now define the singular values.

Definition 14.2.2. Given $A \in \mathbb{R}^{m \times n}$, the **singular values** of A are the square roots of (non-zero) eigenvalues of $A^T A$. We will soon see that these numbers are also the square roots of (non-zero) eigenvalues of AA^T .

Warning: When one says the word “singular values” the context should always be taken into account. Singular values are most often meant to be non-zero numbers but in some settings it makes sense to say the word singular value and have it pertain to square roots of **all** eigenvalues of $A^T A$, non-zero or not. You will likely run into this issue several times and when you do, always take the context into account and think about whether or not allowing zero singular values will mess up the context of the setting in question.

All of this information allows us to compute the σ_i and the \mathbf{v}_i , so it remains to find the \mathbf{u}_i . This is where the equation

$$A \mathbf{v}_i = \sigma_i \mathbf{u}_i$$

comes to the rescue. It tells us that

$$\mathbf{u}_i = \frac{A \mathbf{v}_i}{\sigma_i} \quad \text{for all } i = 1, 2, \dots, r$$

Note that this allows us to find r of the \mathbf{u}_i vectors. The remaining ones must be found via direct null space computations.

We can use this fact to verify that $\mathbf{u}_i \perp \mathbf{u}_j$ for $i \neq j$ as follows.

$$\mathbf{u}_i^T \mathbf{u}_j = \left(\frac{A \mathbf{v}_i}{\sigma_i} \right)^T \left(\frac{A \mathbf{v}_j}{\sigma_j} \right) = \frac{\mathbf{v}_i^T (A^T A \mathbf{v}_j)}{\sigma_i \sigma_j} = \frac{\mathbf{v}_i^T (\sigma_j^2 \mathbf{v}_j)}{\sigma_i \sigma_j} = \frac{\sigma_j \mathbf{v}_i^T \mathbf{v}_j}{\sigma_i} = 0$$

Note that the last equality follows from the fact that the columns of V form an orthonormal basis.

This summarizes the main points of singular value decomposition computations so we now illustrate the details with an example.

Example 14.2.3. Consider the following rank 2 matrix

$$A = \begin{bmatrix} 3 & 0 \\ 4 & 5 \end{bmatrix}$$

Finding the σ_i 's and the \mathbf{v} 's first we see that

$$A^T A = \begin{bmatrix} 25 & 20 \\ 20 & 25 \end{bmatrix} \quad \text{with } \lambda_1 = 45 = \sigma_1^2, \lambda_2 = 5 = \sigma_2^2 \quad \text{and } \mathbf{v}_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \mathbf{v}_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 1 \end{bmatrix}$$

This implies that

$$V = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix} \quad \text{and} \quad \Sigma = \begin{bmatrix} \sqrt{45} & 0 \\ 0 & \sqrt{5} \end{bmatrix}$$

We now use the equation $A\mathbf{v}_i = \sigma_i\mathbf{u}_i$ to find the \mathbf{u}_i . We have

$$A\mathbf{v}_1 = \begin{bmatrix} 3 & 0 \\ 4 & 5 \end{bmatrix} \begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix} = \begin{bmatrix} 3/\sqrt{2} \\ 9/\sqrt{2} \end{bmatrix} = \frac{3}{\sqrt{2}} \begin{bmatrix} 1 \\ 3 \end{bmatrix} = \sqrt{45} \frac{1}{\sqrt{10}} \begin{bmatrix} 1 \\ 3 \end{bmatrix}$$

so $\mathbf{u}_1 = \frac{1}{\sqrt{10}} \begin{bmatrix} 1 \\ 3 \end{bmatrix}$.

$$A\mathbf{v}_2 = \begin{bmatrix} 3 & 0 \\ 4 & 5 \end{bmatrix} \begin{bmatrix} -1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix} = \begin{bmatrix} -3/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix} = \sqrt{5} \frac{1}{\sqrt{10}} \begin{bmatrix} -3 \\ 1 \end{bmatrix}$$

so $\mathbf{u}_2 = \frac{1}{\sqrt{10}} \begin{bmatrix} -3 \\ 1 \end{bmatrix}$. It is important to note that since $\text{rank}(A) = 2$ we did not have to compute any vectors from $\text{Null}(A)$ or $\text{Null}(A^\top)$. In more general scenarios, bases for these subspaces will have to be computed and normalized in order to obtain the matrices V and U .

We now can conclude that

$$U = \frac{1}{\sqrt{10}} \begin{bmatrix} 1 & -3 \\ 3 & 1 \end{bmatrix}$$

hence the singular value decomposition is

$$A = \begin{bmatrix} 3 & 0 \\ 4 & 5 \end{bmatrix} = \frac{1}{\sqrt{10}} \begin{bmatrix} 1 & -3 \\ 3 & 1 \end{bmatrix} \begin{bmatrix} \sqrt{45} & 0 \\ 0 & \sqrt{5} \end{bmatrix} \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{bmatrix}$$

Before ending this section, we point out a different way to understand the matrix U .

Earlier in this chapter, we computed $A^\top A$ and uncovered the inner workings of the matrices V and Σ in the process. Now lets try and analyze AA^\top . Using the fact that $A\mathbf{v}_i = \sigma_i\mathbf{u}_i$ we have that

$$AA^\top\mathbf{u}_i = AA^\top\left(\frac{A\mathbf{v}_i}{\sigma_i}\right) = \frac{A\left(A^\top A\mathbf{v}_i\right)}{\sigma_i} = \frac{A\sigma_i^2\mathbf{v}_i}{\sigma_i} = \sigma_i A\mathbf{v}_i = \sigma_i^2\mathbf{u}_i$$

hence the \mathbf{u}_i are eigenvectors of AA^\top with eigenvalues σ_i^2 .

Lastly, we can also see that the \mathbf{u}_i are indeed unit vectors.

$$\|\mathbf{u}_i\|^2 = \mathbf{u}_i^\top \mathbf{u}_i = \left(\frac{A\mathbf{v}_i}{\sigma_i}\right)^\top \left(\frac{A\mathbf{v}_i}{\sigma_i}\right) = \frac{\mathbf{v}_i^\top A^\top A\mathbf{v}_i}{\sigma_i^2} = \mathbf{v}_i^\top \mathbf{v}_i = \|\mathbf{v}_i\|^2 = 1$$

hence $\|\mathbf{u}_i\| = 1$.

We finish this section by pointing out one remaining subtlety. We have now seen that the singular values are square roots of eigenvalues of both $A^\top A$ and AA^\top . If A is an $m \times n$ matrix, with $m \neq n$, then $A^\top A$ and AA^\top have different sizes, hence different numbers of eigenvalues. This poses potential confusion, and to avoid it, we note that the number of **non-zero** eigenvalues of both are the same, since this number is equal to the rank. There will potentially be a different multiplicity of 0 eigenvalues of $A^\top A$ and AA^\top but this does not pose any problems to the mechanics of what is happening.

14.3 The Geometry of the Singular Value Decomposition

Let's start off with a quick refresher of the main geometric notions.

Let the unit n sphere be given by

$$\mathbb{S}^{n-1} = \{\mathbf{x} \in \mathbb{R}^n : x_1^2 + x_2^2 + \cdots + x_n^2 = 1\}$$

and the unit n ball by

$$\mathbb{B}^n = \{\mathbf{x} \in \mathbb{R}^n : x_1^2 + x_2^2 + \cdots + x_n^2 \leq 1\}$$


The fundamental fact that motivated the singular value decomposition was the following.

The image of \mathbb{S}^n under $A \in \mathbb{R}^{m \times n}$ is a hyperellipse

Ellipses and hyperellipses are defined by specific equations. A hyperellipse in \mathbb{R}^2 is just a normal ellipse and is given by the equation

$$\frac{x_1^2}{a} + \frac{x_2^2}{b} = 1$$

Taking $a = 4$ and $b = 2$ we obtain the example




We call the values \sqrt{a} and \sqrt{b} the length of the semi-axes, and the axes of the ellipse themselves are known as the semi-axes.

A (general) hyperellipse in \mathbb{R}^n is given by

$$\frac{y_1^2}{\sigma_1^2} + \frac{y_2^2}{\sigma_2^2} + \dots + \frac{y_n^2}{\sigma_n^2} = 1$$

An example of a hyperellipse in \mathbb{R}^3 would look like



Example 14.3.1. Consider the singular value decomposition of the following matrix

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

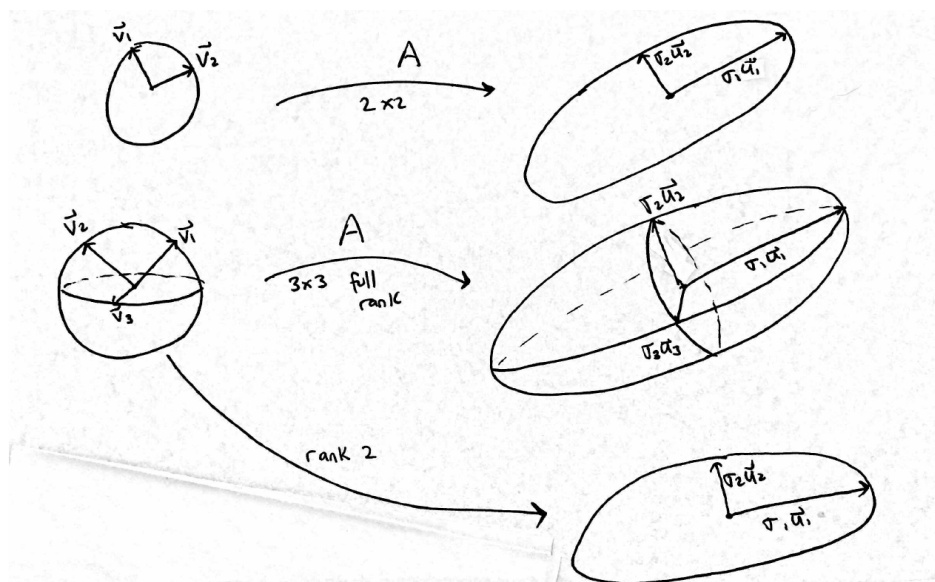
Thinking of A as a linear map $A: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ sending $\mathbf{x} \mapsto A\mathbf{x}$ we can apply A to the unit sphere in the domain (this is \mathbb{S}^1) with domain basis $\mathbf{v}_1 = \mathbf{e}_2$ and $\mathbf{v}_2 = \mathbf{e}_1$ we see that it maps \mathbb{S}^1 to the hyperellipse in \mathbb{R}^2 with semi-axes given by $\sigma_1 \mathbf{u}_1 = 3\mathbf{e}_2$ and $\sigma_2 \mathbf{u}_2 = 2\mathbf{e}_1$. The vectors \mathbf{u}_1 and \mathbf{u}_2 form an orthonormal basis for \mathbb{R}^2 and the lengths of the semi-axes are given by the singular values of A .

This example illustrates the general phenomenon that happens with any matrix. We can summarize everything up to this point with a sequence of facts.

Let $A \in \mathbb{R}^{m \times n}$ and assume $\text{rank}(A) = r$. Then

1. $A(\mathbb{S}^{n-1})$ is a hyperellipse in \mathbb{R}^m of dimension r .
2. The semiaxes of the image hyperellipse are $\sigma_1 \mathbf{u}_1, \sigma_2 \mathbf{u}_2, \dots, \sigma_r \mathbf{u}_r$.
3. The singular values are the lengths of the semiaxes of the image hyperellipse.
4. The vectors $\mathbf{u}_1, \dots, \mathbf{u}_r$ are the unit vectors along the semiaxes of the image hyperellipse.
5. The vectors $\mathbf{v}_1, \dots, \mathbf{v}_r$ map to $\sigma_1 \mathbf{u}_1, \sigma_2 \mathbf{u}_2, \dots, \sigma_r \mathbf{u}_r$ under A .

The rank of the given matrix determines the dimension of the image hyperellipse as we can see in the following pictures.



We can also see that the geometric information agrees with the algebraic picture. In particular, we have

- $\mathbf{u}_1, \dots, \mathbf{u}_r$ are orthonormal vectors in $\text{Range}(A)$.
- $\sigma_1, \dots, \sigma_r > 0$ because they are lengths of semiaxes.
- $\mathbf{v}_1, \dots, \mathbf{v}_n$ are orthonormal vectors in the domain (hence live on the unit sphere \mathbb{S}^{n-1}).

Our next job will be to understand the singular value decomposition of A as a linear map. Since it is given as a factorization ($A = U\Sigma V^T$), we can view the linear map associated to A as a composition of 3 linear maps whose net composition has the same effect as A . Since U and V are orthogonal (hence invertible), they are change of basis matrices. As a result of this observation, we will find that the geometry of the singular value decomposition is best described using the language of change of bases. First, some facts concerning orthogonal matrices.

Proposition 14.3.2. *If $Q \in \mathbb{R}^{n \times n}$ is an orthogonal matrix, then Q preserves dot products (hence preserves angles between vectors and lengths of vectors).*

Proof. If $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$, we show that

$$\mathbf{v}^\top \mathbf{w} = (Q\mathbf{v})^\top (Q\mathbf{w})$$

Observe that

$$(Q\mathbf{v})^\top (Q\mathbf{w}) = \mathbf{v}^\top Q^\top Q\mathbf{w} = \mathbf{v}^\top I\mathbf{w} = \mathbf{v}^\top \mathbf{w}$$

This also allows us to see that Q preserves lengths since

$$\|Q\mathbf{v}\|^2 = (Q\mathbf{v})^\top (Q\mathbf{v}) = \mathbf{v}^\top Q^\top Q\mathbf{v} = \mathbf{v}^\top \mathbf{v} = \|\mathbf{v}\|^2$$

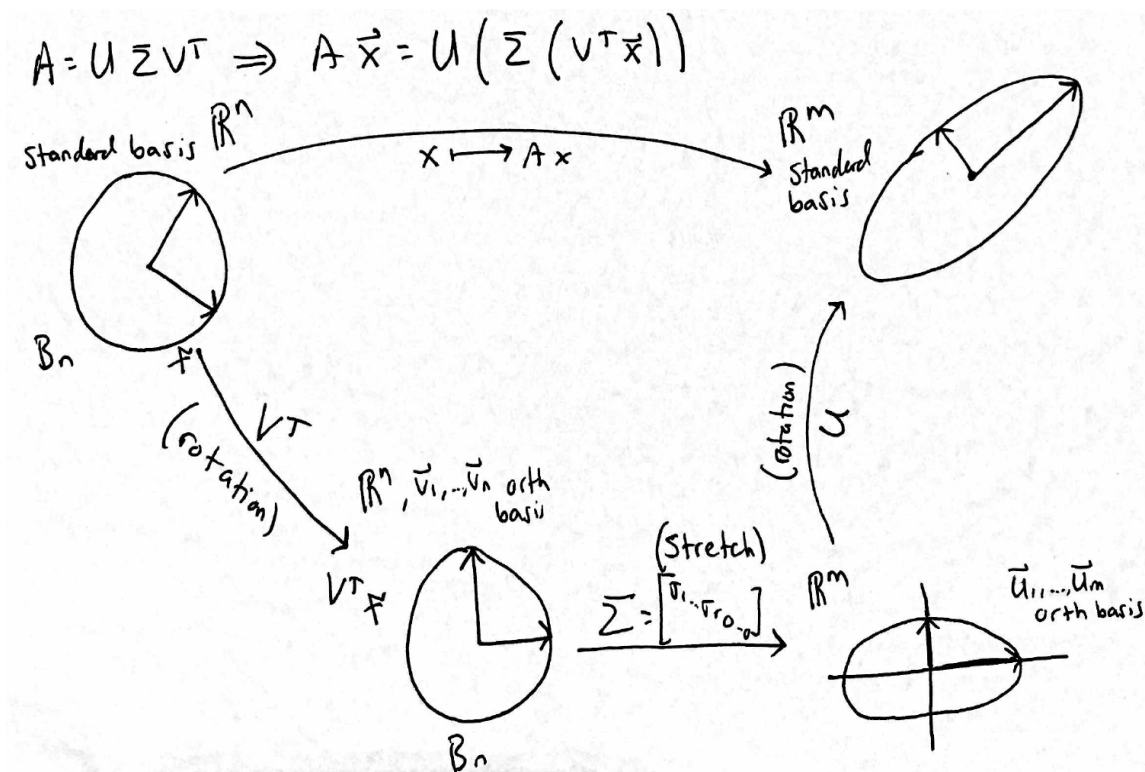
□

We can conclude from this that orthogonal matrices act as “rotations”. We write this in quotation marks because they are not identically rotations as we know them, but for the sake of understanding the geometry, we may think of them simply as rotating space.

Let’s now look back at the singular value decomposition, keeping in mind that U and V are orthogonal. By viewing the factorization as a composition of linear maps, two of which are “rotations”, we can determine where a vector goes by seeing what each individual matrix does to that vector. That is

$$A = U\Sigma V^\top \implies A\mathbf{x} = U(\Sigma(V^\top \mathbf{x}))$$

The following picture summarizes the composition



We now illustrate the details of this with a series of important statements, the first of which does not need to be thought of as a formal proposition but will serve as a guide for our change of basis perspective.

Proposition 14.3.3. Let $A = U\Sigma V^\top \in \mathbb{R}^{m \times n}$. If we take $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ and $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ as orthonormal bases of the domain and codomain of A , respectively, then A behaves like the matrix Σ .

This can be explained by the following picture

$$\underbrace{\mathbf{x}}_{\text{(In standard basis)}} \xrightarrow{V^\top} \underbrace{V^\top \mathbf{x}}_{\text{(In } V \text{ basis, still in } \mathbb{R}^n)} \xrightarrow{\Sigma} \underbrace{\Sigma V^\top \mathbf{x}}_{\text{(In } U \text{ basis, now in } \mathbb{R}^m)} \xrightarrow{U} \underbrace{U \Sigma V^\top \mathbf{x}}_{\text{(Back to standard basis, now in } \mathbb{R}^n)}$$

With this idea in mind, we can now prove the motivating statement of the chapter.

Proposition 14.3.4. *Let $A \in \mathbb{R}^{m \times n}$ and let \mathbb{S}^{n-1} be the unit n sphere in the domain of A . Then $A(\mathbb{S}^{n-1})$ is always a hyperellipse.*

Proof. First, assume without loss of generality that $\text{rank}(A) = n$. Pick any vector $\mathbf{x} \in \mathbb{S}^{n-1}$ and write \mathbf{x} in the V basis (rather than the standard basis), so that

$$\mathbf{x} = x_1 \mathbf{v}_1 + x_2 \mathbf{v}_2 + \cdots + x_n \mathbf{v}_n$$

where

$$x_1^2 + x_2^2 + \cdots + x_n^2 = 1$$

Note that the last condition is true because $\mathbf{x} \in \mathbb{S}^{n-1}$ and will show up again. Next, look at the image of \mathbf{x} under A and label it \mathbf{y} , so $A\mathbf{x} = \mathbf{y}$. We get that

$$\mathbf{y} = A\mathbf{x} = A(x_1 \mathbf{v}_1 + x_2 \mathbf{v}_2 + \cdots + x_n \mathbf{v}_n) = x_1 \sigma_1 \mathbf{u}_1 + \cdots + x_n \sigma_n \mathbf{u}_n$$

since $A\mathbf{v}_i = \sigma_i \mathbf{u}_i$. Next, we relabel our coefficients and write

$$\mathbf{y} = \sum_{i=1}^n x_i \sigma_i \mathbf{u}_i = \sum_{i=1}^n y_i \mathbf{u}_i$$

to ease notation. Now, we write \mathbf{y} in the U basis, and see that

$$[\mathbf{y}]_{\mathcal{B}_U} = \begin{bmatrix} x_1 \sigma_1 \\ \vdots \\ x_n \sigma_n \end{bmatrix} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}$$

Now we have obtained the image of the vector \mathbf{x} in the U basis, so it remains to verify that the coordinates of \mathbf{y} (the y_i) satisfy the equation of a hyperellipse. Since $y_i = x_i \sigma_i$ we have $x_i = \frac{y_i}{\sigma_i}$ hence

$$\left(\frac{y_1}{\sigma_1}\right)^2 + \cdots + \left(\frac{y_n}{\sigma_n}\right)^2 = x_1^2 + \cdots + x_n^2 = 1$$

We can now conclude that $\mathbf{y} = A\mathbf{x}$ is a vector on the hyperellipse with semiaxes of lengths $\sigma_1, \dots, \sigma_n$ in directions $\mathbf{u}_1, \dots, \mathbf{u}_n$. \square

We finish off the section with one more useful fact concerning singular values. We will need a lemma to finish the proof of our useful proposition so we state and prove the lemma here.

Lemma 14.3.5. *Let $A \in \mathbb{R}^{m \times n}$ and let $B \in \mathbb{R}^{n \times n}$ be invertible, then $\text{rank}(A) = \text{rank}(AB)$.*

Proof. Recall that $\text{rank}(AB) \leq \min\{\text{rank}(A), \text{rank}(B)\}$ and suppose that B^{-1} exists. Then $\text{rank}(A) = \text{rank}((AB)B^{-1}) \leq \text{rank}(AB)$ since the rank of B^{-1} is as big as possible. We can then conclude that

$$\text{rank}(A) = \text{rank}((AB)B^{-1}) \leq \text{rank}(AB) \leq \min\{\text{rank}(A), \text{rank}(B)\} = \text{rank}(A)$$

This implies that

$$\text{rank}(A) \leq \text{rank}(AB) \leq \text{rank}(A)$$

hence we must have $\text{rank}(A) = \text{rank}(AB)$. \square

Now we can state and prove the proposition.

Proposition 14.3.6. $\text{rank}(A)$ is equal to the number of (non-zero) singular values of A , counting multiplicities.

Proof. Given $A = U\Sigma V^\top$, begin by observing that U and V^\top have full rank and that $\text{rank}(\Sigma) = r$. Also notice that $U\Sigma$ has r linearly independent rows and columns hence $\text{rank}(U\Sigma) = r$ (you should verify this for yourself). Now, looking at the product $(U\Sigma)V^\top$ and noting that V^\top is invertible, we can apply the above lemma and conclude that

$$\text{rank}(A) = \text{rank}(U\Sigma V^\top) = \text{rank}(U\Sigma) = r$$

□

This proposition gives a nice way of computing the rank of a matrix in practice.

1. Compute the singular value decomposition (in Julia).
2. If some singular values are VERY close to 0, then they are 0 (this happens because of rounding errors in the computer).
3. $\text{rank}(A)$ = the number of nonzero singular values of A .

We finish off this section with a small additional notion that may come up later, namely the **polar decomposition** of a square matrix.

Let $A \in \mathbb{R}^{n \times n}$. Observe that we can use the singular value decomposition to write A as

$$A = U\Sigma V^\top = (UV^\top)(V\Sigma V^\top)$$

Noting that UV^\top is orthogonal and $V\Sigma V^\top$ is positive semidefinite, we can conclude that any square matrix can be written as the product of an orthogonal matrix with a positive semidefinite matrix. This factorization is known as the **polar decomposition of A** . This will come up in some exercises but we will leave it at the definition for now.

14.4 The Reduced Singular Value Decomposition

We now begin covering the preliminary ideas needed to understand how one finds low rank matrices that are “close” to a given matrix. This is one of the central ideas behind principal component analysis in addition to many other notions related to the singular value composition.

Thus far, we have seen that for any matrix A of rank r , we have the singular values of A as the square roots of $\sigma_1^2, \dots, \sigma_r^2$, which are eigenvalues of AA^\top and $A^\top A$ (the remaining eigenvalues are 0). The reduced singular value decomposition is a different factorization of A which captures most of the data that A holds, without adding additional computational hurdles.

Definition 14.4.1. Given

$$A = [\mathbf{u}_1 \ \cdots \ \mathbf{u}_r \ \mathbf{u}_{r+1} \ \cdots \ \mathbf{u}_m] \Sigma [\mathbf{v}_1 \ \cdots \ \mathbf{v}_r \ \mathbf{v}_{r+1} \ \cdots \ \mathbf{v}_n]^\top$$

the **reduced singular value decomposition of A** is the factorization

$$A = [\mathbf{u}_1 \ \cdots \ \mathbf{u}_r] \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix} [\mathbf{v}_1 \ \cdots \ \mathbf{v}_r]^\top = \hat{U} \hat{\Sigma} \hat{V}^\top$$

We call the vectors $\mathbf{u}_1, \dots, \mathbf{u}_r$ the **left singular vectors** of A and similarly, the vectors $\mathbf{v}_1, \dots, \mathbf{v}_r$ are called the **right singular vectors** of A .

Note that in contrast to the (full) singular value decomposition, the reduced version has $\hat{\Sigma}$ as a square diagonal matrix and \hat{U} and \hat{V} as rectangular matrices, if A is not square.

Example 14.4.2. Let A be the following matrix

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

Note that the only eigenvalue of A is 0 with $\text{AM}(0) = 4$ and $\text{GM}(0) = 1$. Moreover, $\text{rank}(A) = 3$ and A is not invertible. Computing the singular value decomposition of A we see that

$$A^T A = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 9 \end{bmatrix}$$

hence $\sigma_1 = 3, \sigma_2 = 2, \sigma_3 = 1$, and $\sigma_4 = 0$. The respective eigenvectors of $A^T A$ are

$$\mathbf{v}_1 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \mathbf{v}_2 = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \mathbf{v}_3 = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \text{ and } \mathbf{v}_4 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

hence

$$V^T = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}$$

Next, we use the fact that $A\mathbf{v}_i = \sigma_i\mathbf{u}_i$ for $i = 1, 2, 3$ to find the first three columns of U . They are

$$\mathbf{u}_1 = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \mathbf{u}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \text{ and } \mathbf{u}_3 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \text{ The last step is to find the one orthonormal basis vector for } \text{Null}(A^T),$$

which by inspection is $\mathbf{u}_4 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}$. This rounds out the computation, giving us the matrix U hence the full singular value decomposition is

$$A = \underbrace{\begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}}_U \underbrace{\begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}}_{\Sigma} \underbrace{\begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}}_{V^T}$$

Computing the reduced singular value decomposition reduces to eliminating 0 singular values and any vectors coming from null spaces. This yields

$$\underbrace{\begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{\hat{U}} \underbrace{\begin{bmatrix} 3 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{\hat{\Sigma}} \underbrace{\begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}}_{\hat{V}^T}$$

Next, we use this to approximate the matrix A . In order to understand approximations to a matrix, we must define matrix norms which give us a notion of “closeness” in $\mathbb{R}^{m \times n}$.

14.5 Matrix Norms

We define matrix norms in terms of vector norms, of which there are many.

Definition 14.5.1. Let $\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$. The **1-norm** of \mathbf{x} is the real number

$$\|\mathbf{x}\|_1 = |x_1| + |x_2| + \cdots + |x_n|$$

The **2-norm** of \mathbf{x} is the real number

$$\|\mathbf{x}\|_2 = \sqrt{x_1^2 + \cdots + x_n^2}$$

You may now realize that the 2-norm of a vector is the usual norm we have known all along. This is the one we care about the most, as will also be the case with matrices. We now define matrix norms in terms of vector norms.

Definition 14.5.2. Let $A = (a_{ij}) \in \mathbb{R}^{m \times n}$. The **∞ -norm** of A is the maximal row sum, also defined as the real number

$$\|A\|_\infty = \max_{i=1, \dots, n} \left\{ \sum_{j=1}^n |a_{ij}| \right\}$$

The **1-norm** of A is the maximal column sum, also defined as the real number

$$\|A\|_1 = \max_{j=1, \dots, m} \left\{ \sum_{i=1}^m |a_{ij}| \right\}$$

The **2-norm** of A is defined to be the real number

$$\|A\|_2 = \max_{\mathbf{x} \neq \mathbf{0}} \left\{ \frac{\|A\mathbf{x}\|_2}{\|\mathbf{x}\|_2} \right\}$$

$\|A\|_2$ is the only matrix norm we care about in this course. The observation that for any real number α we have that

$$\|\alpha \mathbf{x}\| = |\alpha| \|\mathbf{x}\|$$

implies that

$$\frac{\|A\alpha \mathbf{x}\|_2}{\|\alpha \mathbf{x}\|_2} = \frac{|\alpha| \|A\mathbf{x}\|_2}{|\alpha| \|\mathbf{x}\|_2} = \frac{\|A\mathbf{x}\|_2}{\|\mathbf{x}\|_2}$$

hence we can actually define this norm to be

$$\|A\|_2 = \max_{\|\mathbf{x}\|_2=1} \left\{ \|A\mathbf{x}\|_2 \right\}$$

In words, this means that the 2-norm of the matrix A is the maximum stretch that A applies to a unit vector.

This maximum stretch is none other than the largest singular value of A !!!

That is, if

$$A = [\mathbf{u}_1 \quad \cdots \quad \mathbf{u}_r] \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix} [\mathbf{v}_1 \quad \cdots \quad \mathbf{v}_r]^\top = \hat{U} \hat{\Sigma} \hat{V}^\top$$

then $\|A\|_2 = \sigma_1$.

Before digging deeper into how we use this norm to approximate A we note two facts about matrix norms.

$$\|A + B\|_2 \leq \|A\|_2 + \|B\|_2 \quad \text{and} \quad \|AB\|_2 \leq \|A\|_2 \|B\|_2$$

14.6 Rank One Decompositions

Recall that the singular value decomposition of A (with $\text{rank}(A) = r$) allows one to write A as a sum of r rank 1 matrices.

$$[\mathbf{u}_1 \quad \cdots \quad \mathbf{u}_r] \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix} \begin{bmatrix} \mathbf{v}_1^\top \\ \vdots \\ \mathbf{v}_r^\top \end{bmatrix} = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^\top + \cdots + \sigma_r \mathbf{u}_r \mathbf{v}_r^\top$$

In general, there are many ways of writing a matrix as the sum of rank one matrices but the one we have above is special and allows us to define lower rank approximations to A .

Definition 14.6.1. Let $\text{rank}(A) = r$. For $k < r$, define the **rank k approximation to A** as the k^{th} partial sum

$$A_k = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^\top + \cdots + \sigma_k \mathbf{u}_k \mathbf{v}_k^\top = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^\top$$

Note that the case where $k = r$ is what we call the rank one **decomposition** whereas $k < r$ gives rise to the rank k **approximation**.

The word approximation is meant to imply that the object of interest is “close” to A . When we say the word close, we mean close in the 2-norm. This is made precise with the following theorem due to Eckart and Young.

Theorem 14.6.2. *The closest rank k matrix (in the 2-norm) to A is the matrix*

$$A_k = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^\top$$

More formally, if B is any $m \times n$ matrix of rank k , then

$$\|A - B\|_2 \geq \|A - A_k\|_2 = \sigma_{k+1}$$

We do not prove this theorem here but we will show that $\|A - A_k\|_2 = \sigma_{k+1}$. Since the singular values of A are ordered, this tells us that A_k gets closer to A as k gets closer to r .

To see why this is true, we will need a lemma that, as stated, is simple but has immense use in practice.

Lemma 14.6.3.

$$A = B \Leftrightarrow A\mathbf{v} = B\mathbf{v} \quad \forall \mathbf{v} \text{ in the domain} \Leftrightarrow A\mathbf{u}_i = B\mathbf{u}_i \text{ for all basis vectors } \mathbf{u}_i$$

Proof. It is easy to see that if $A = B$ then $A\mathbf{v} = B\mathbf{v}$ for all domain vectors \mathbf{v} , since A and B are exactly the same matrix. On the other hand, if we assume that $A\mathbf{v} = B\mathbf{v}$ for all vectors \mathbf{v} , then in particular, we have that $A\mathbf{e}_i = B\mathbf{e}_i$ for all standard basis vectors \mathbf{e}_i of the domain. Since $A\mathbf{e}_i$ is the i^{th} column of A and $B\mathbf{e}_i$ is the i^{th} column of B , we can conclude that the i^{th} column of A and B are equal for all i , hence A and B are the same matrix. We leave the proof of the last equivalence (the one for basis vectors) to the reader. You may find it useful in homework problems but if you use it you must argue why it is true. \square

Now onto the main idea.

Proposition 14.6.4.

$$\|A - A_k\|_2 = \sigma_{k+1}$$

Proof. Using the result about equality of matrices from the lemma above, we can work backwards from a certain sum of rank one matrices to conclude that

$$A - A_k = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^\top - \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^\top = \sum_{i=k+1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^\top = [\mathbf{u}_{k+1} \quad \cdots \quad \mathbf{u}_r] \begin{bmatrix} \sigma_{k+1} & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix} \begin{bmatrix} \mathbf{v}_{k+1}^\top \\ \vdots \\ \mathbf{v}_r^\top \end{bmatrix}$$

Note that the lemma above was used to obtain the last equality. Looking at the final product of matrices, we are left with none other than the singular value decomposition of the matrix $A - A_k$, thus the value $\|A - A_k\|_2$ is the largest singular value of this matrix, which is σ_{k+1} . \square

Before moving onto the next section, lets look back at example 14.4.2 in the context of rank one approximations.


Example 14.6.5. We computed the reduced singular value decomposition in the previous example and obtained

$$\begin{aligned} A &= \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{\hat{U}} \underbrace{\begin{bmatrix} 3 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}}_{\hat{\Sigma}} \underbrace{\begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}}_{\hat{V}^\top} \\ &= 3 \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} [0 \quad 0 \quad 0 \quad 1] + 2 \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} [0 \quad 0 \quad 1 \quad 0] + \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} [0 \quad 1 \quad 0 \quad 0] \end{aligned}$$

The closest rank 1 matrix to A is

$$A_1 = 3 \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} [0 \quad 0 \quad 0 \quad 1] = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$


This is a linear map from \mathbb{R}^4 to \mathbb{R}^4 of rank 1 so the image of the unit sphere in \mathbb{R}^4 is a one dimensional hyperellipse which looks like




The closest rank 2 matrix to A is

$$A_2 = 3 \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} [0 \ 0 \ 0 \ 1] + 2 \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} [0 \ 0 \ 1 \ 0] = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

This is a linear map from \mathbb{R}^4 to \mathbb{R}^4 of rank 2 so the image of the unit sphere in \mathbb{R}^4 is a two dimensional hyperellipse which looks like



The full matrix A is a rank 3 linear map from \mathbb{R}^4 to \mathbb{R}^4 and the image of the unit sphere in \mathbb{R}^4 is the 3 dimensional hyperellipse which looks like



Geometrically, we can see that the higher rank approximations build up the image hyperellipse one semiaxis at a time.


14.7 Best Fit k -Planes

We begin this section with a general question. Suppose we obtain m data points $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{R}^n$ and we input them into a matrix as the rows

$$A = \begin{bmatrix} \mathbf{a}_1^T \\ \vdots \\ \mathbf{a}_n^T \end{bmatrix} \in \mathbb{R}^{m \times n}$$

Question 14.7.1. What is the “best fit k -dimensional subspace” to the data points $\mathbf{a}_1, \dots, \mathbf{a}_n$?

Let’s think about the best fit line l . This is the case when $k = 1$. The best fit line is the line that minimizes the sum of squared distances of $\mathbf{a}_1, \dots, \mathbf{a}_m$ to the line.



We can see that $d_i^2 = \|\mathbf{a}_i\|^2 - \|\mathbf{p}_i\|^2$, where $\mathbf{p}_i = \text{proj}_l \mathbf{a}_i$ is the (orthogonal) projection of \mathbf{a}_i onto the line l . From this we can conclude that in order to minimize the value $\sum d_i^2$, we want to maximize $\sum \|\mathbf{p}_i\|^2$. This

value is the sum of squares of lengths of the projections of $\mathbf{a}_1, \dots, \mathbf{a}_m$ onto the best fit line, l . This leads us to our second question.

Question 14.7.2. Which line maximizes the value $\sum \|\mathbf{p}_i\|^2$?

Let's begin by considering a unit vector \mathbf{v} on the best fit line l (note that every line has a unit vector lying on it). Suppose that \mathbf{a}_i and \mathbf{v} have θ radians between them and consider the dot product of \mathbf{a}_i and \mathbf{v} . We see that

$$\mathbf{a}_i^\top \mathbf{v} = \|\mathbf{a}_i\| \|\mathbf{v}\| \cos \theta = \|\mathbf{a}_i\| \cos \theta$$

Using our old friend SOHCAHTOA, we can conclude that $\cos \theta = \frac{\|\mathbf{p}_i\|}{\|\mathbf{a}_i\|}$ which implies that

$$|\mathbf{a}_i^\top \mathbf{v}| = \|\mathbf{a}_i\| \cos \theta = \|\mathbf{p}_i\|$$

We can then conclude that $\|\mathbf{p}_i\|^2 = |\mathbf{a}_i^\top \mathbf{v}|^2$. Recalling that the \mathbf{a}_i were the rows of the matrix A , we can

connect the two ideas and see that since $A = \begin{bmatrix} \mathbf{a}_1^\top \\ \vdots \\ \mathbf{a}_n^\top \end{bmatrix}$ we have

$$A\mathbf{v} = \begin{bmatrix} \mathbf{a}_1^\top \\ \vdots \\ \mathbf{a}_n^\top \end{bmatrix} \mathbf{v} = \begin{bmatrix} \mathbf{a}_1^\top \mathbf{v} \\ \vdots \\ \mathbf{a}_n^\top \mathbf{v} \end{bmatrix}$$

hence

$$\sum_{i=1}^m \|\mathbf{p}_i\|^2 = \|A\mathbf{v}\|^2$$

This translates our original question to a new one.

Question 14.7.3. Which unit vector $\mathbf{v} \in \mathbb{R}^n$ maximizes the value $\|A\mathbf{v}\|^2$?

This is now something that is fresh and familiar. Recalling that $\|A\|_2 = \max_{\|\mathbf{x}\|_2=1} \{\|A\mathbf{x}\|_2\}$ we can conclude that the choice of \mathbf{v} which maximizes this value is the first right singular vector \mathbf{v}_1 ! Since $A\mathbf{v}_i = \sigma_i \mathbf{u}_i$, we can see that

$$\|A\|_2 = \max_{\|\mathbf{x}\|_2=1} \{\|A\mathbf{x}\|_2\} = \sigma_1 = \|A\mathbf{v}_1\| = \|\sigma_1 \mathbf{u}_1\| = |\sigma_1| \|\mathbf{u}_1\|$$

We can summarize all of this with a proposition.

Proposition 14.7.4. If $A = \begin{bmatrix} \mathbf{a}_1^\top \\ \vdots \\ \mathbf{a}_n^\top \end{bmatrix} \in \mathbb{R}^{m \times n}$, then the best fit line to the rows of A is $\text{Span}\{\mathbf{v}_1\}$, where \mathbf{v}_1 is the first right singular vector of A .

We leave out the details of the proof but state the all important generalization of this idea.

Theorem 14.7.5. Let $\text{rank}(A) = r$. For $1 \leq k \leq r$, let $V_k = \text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ where the \mathbf{v}_i denote the right singular vectors of A . Then the best fit k -plane to the rows of A is the subspace V_k .

By transposing A , we can obtain a similar theorem concerning the columns of A . Observe that if $A = U\Sigma V^\top$ then $A^\top = V\Sigma^\top U^\top$. Since the rows of A^\top are the columns of A , we can directly apply the previous theorem to conclude the following.

Theorem 14.7.6. Let $U_k = \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_k\}$ where \mathbf{u}_i denote the left singular vectors of A . Then the best fit k -plane to the columns of A is the subspace U_k .

We should pause for a moment and let this sink in. We have seen many reasons why the singular value decomposition of A is important. The singular values alone tell us a tremendous amount of information, but we now have more. Mainly that the left and right singular vectors give the best fit k -dimensional subspaces to the columns (resp. rows) of A . Before returning to example 14.4.2, we give one more important proposition concerning the relationship between the best fit k -planes and the rank one decomposition of A .

Proposition 14.7.7. *Suppose $A = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^\top$ and $A_k = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^\top$. The rows of the matrix A_k are the projections of the rows of A onto the subspace $V_k = \text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$*

Proof. Let \mathbf{a} be a row of A . Recall that the \mathbf{v}_i are mutually orthogonal unit vectors which form a basis for \mathbb{R}^n , hence there exist scalars a_i such that

$$\mathbf{a} = a_1 \mathbf{v}_1 + \dots + a_n \mathbf{v}_n$$

Given $V_k = \text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ we know that $V_k^\perp = \text{Span}\{\mathbf{v}_{k+1}, \dots, \mathbf{v}_n\}$ so that

$$\mathbf{a} = \mathbf{a}_{V_k} + \mathbf{a}_{V_k^\perp}$$

where $\mathbf{a}_{V_k} = a_1 \mathbf{v}_1 + \dots + a_k \mathbf{v}_k$ and $\mathbf{a}_{V_k^\perp} = a_{k+1} \mathbf{v}_{k+1} + \dots + a_n \mathbf{v}_n$. In this way, we can think of the projection of \mathbf{a} onto V_k as the vector \mathbf{a}_{V_k} .

Now, along these lines, we know that the projection of \mathbf{a} onto V_1 is $a_1 \mathbf{v}_1^\top$ (we are writing the vector \mathbf{v}_1 as a row vector because we are projecting a row vector onto V_1). We can extract the coefficient a_1 by computing a dot product. That is

$$\mathbf{a}^\top \mathbf{v}_1 = (a_1 \mathbf{v}_1 + \dots + a_n \mathbf{v}_n)^\top \mathbf{v}_1 = a_1$$

hence

$$\text{proj}_{V_1} \mathbf{a} = (\mathbf{a}^\top \mathbf{v}_1) \mathbf{v}_1^\top = a_1 \mathbf{v}_1^\top$$

Along these same lines, we get that $a_2 = \mathbf{a}^\top \mathbf{v}_2$ so that

$$\text{proj}_{V_2} \mathbf{a} = (\mathbf{a}^\top \mathbf{v}_1) \mathbf{v}_1^\top + (\mathbf{a}^\top \mathbf{v}_2) \mathbf{v}_2^\top$$

and in general we have

$$\text{proj}_{V_k} \mathbf{a} = \sum_{i=1}^k (\mathbf{a}^\top \mathbf{v}_i) \mathbf{v}_i^\top$$

Now let's look at the matrix whose rows are the projections of the rows of A (the \mathbf{a}_i) onto V_k . We have explicit formulas for the projections hence can write this matrix out as

$$\begin{bmatrix} \sum_{i=1}^k (\mathbf{a}_1^\top \mathbf{v}_i) \mathbf{v}_i^\top \\ \sum_{i=1}^k (\mathbf{a}_2^\top \mathbf{v}_i) \mathbf{v}_i^\top \\ \vdots \\ \sum_{i=1}^k (\mathbf{a}_m^\top \mathbf{v}_i) \mathbf{v}_i^\top \end{bmatrix} = \sum_{i=1}^k (A \mathbf{v}_i) \mathbf{v}_i^\top = \sum_{i=1}^k (\sigma_i \mathbf{u}_i) \mathbf{v}_i^\top = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^\top = A_k$$

This is none other than the rank k approximation of A , finishing the proof. □

Now let's return to example 14.4.2 and see these projections in action.

Example 14.7.8. Thus far we have computed


$$A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{\hat{U}} \underbrace{\begin{bmatrix} 3 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{\hat{\Sigma}} \underbrace{\begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}}_{\hat{V}^\top}$$

$$= 3 \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} [0 \ 0 \ 0 \ 1] + 2 \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} [0 \ 0 \ 1 \ 0] + \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} [0 \ 1 \ 0 \ 0]$$

Now let's look at the rows of A . They are (written as columns)

$$\mathbf{a}_1 = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{a}_2 = \begin{bmatrix} 0 \\ 0 \\ 2 \\ 0 \end{bmatrix}, \quad \mathbf{a}_3 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 3 \end{bmatrix}, \quad \mathbf{a}_4 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

Even though the rows of A are vectors in \mathbb{R}^4 , none of the vectors have non-zero fourth coordinate so we can associate a picture with these vectors.



Looking at the rows of

$$A_1 = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T = 3 \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} [0 \ 0 \ 0 \ 1] = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

we see the proposition in action.

\mathbf{a}_1 and \mathbf{a}_2 get projected to $\mathbf{0}$ (see picture) whereas \mathbf{a}_3 gets projected to itself because it is already on the best fit line, $V_1 = \text{Span}\{\mathbf{a}_3\}$.

Looking at the best fit plane $V_2 = \text{Span} \left\{ \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \right\} = zw\text{-plane}$. and the rows of

$$A_2 = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T + \sigma_2 \mathbf{u}_2 \mathbf{v}_2^T = 3 \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} [0 \ 0 \ 0 \ 1] + 2 \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} [0 \ 0 \ 1 \ 0] = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

we see that row 1 of A_2 is $\mathbf{0}$ and \mathbf{a}_1 projects to $\mathbf{0}$. Moreover, \mathbf{a}_2 and \mathbf{a}_3 map to themselves because $\mathbf{a}_2, \mathbf{a}_3 \in V_2$. Summarized more succinctly, we have that the rows of A_2 are the projections of $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, \mathbf{a}_4$ onto V_2 .

We finish off the chapter with a description of principal component analysis.

14.8 Application: Principal Component Analysis

When modeling many real world systems, data points $\mathbf{a}_i \in \mathbb{R}^m$ are obtained and input as the columns (or rows) of a matrix $A = [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n] \in \mathbb{R}^{m \times n}$. Once we have the matrix A , several steps are taken to obtain what is known as the **principal components**.

1. First, the mean of each row is computed and subtracted from each row. This centers the data in the sense that the average along every row is now 0, hence the “center” of the data points is the origin.
2. Then the singular value decomposition is computed, and the basis vectors for the best fit k -planes to the columns, i.e. the vectors $\mathbf{u}_1, \dots, \mathbf{u}_r$, are found (they are actually computed as eigenvectors of $\frac{AA^T}{n-1}$ for statistical reasons). The vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_r$ are called the **principal components** of the data points.

The best fit lines, planes, and k -planes are the subspaces that most closely fit the data. There is much more that can be said about the fitting of these data points, but in short, we can say that in many situations, a large proportion of the variance in a data set is explained by the first several principal components. We illustrate the general idea with an example, but then turn to a nice visual representation of a large data set.

Example 14.8.1. Suppose we have math scores and history scores of six students. Our data points are of the form (x_i, y_i) where x_i denotes the math score of student i and y_i denotes the history score of student i . Inputting some sample data into a matrix we have

$$A = \begin{bmatrix} 3 & -4 & 7 & 1 & -4 & -3 \\ 7 & -6 & 8 & -1 & -1 & -7 \end{bmatrix}$$

Finding the principal components reduces to finding the eigenvectors of $\frac{AA^T}{5}$ and these are

$$\mathbf{u}_1 = \begin{bmatrix} .6 \\ .8 \end{bmatrix} \quad \text{and} \quad \mathbf{u}_2 = \begin{bmatrix} .8 \\ -.6 \end{bmatrix}$$

with respective singular values being $\sigma_1 = \sqrt{57}$ and $\sigma_2 = \sqrt{3}$. The first rank one piece of A is $\sqrt{57}\mathbf{u}_1\mathbf{v}_1^T$ and is much larger than the second piece $\sqrt{3}\mathbf{u}_2\mathbf{v}_2^T$. The leading left singular vector \mathbf{u}_1 gives the direction of the line that most closely fits the data (the best fit line). Since σ_1 is large relative to σ_2 , it means that a large proportion of the variance in our data is explained by the first principal component.

In general, the power of principal component analysis is most easily seen with large data sets. The following website

<https://setosa.io/ev/principal-component-analysis/>

gives a nice illustration of this. If you rotate the second example data set, one can see that the most variance among the data points can be seen along the first principal component, and less variance can be seen when looking along different lines given by other right singular vectors.

We also mention that principal component analysis is central to the Netflix problem. For a great explanation of this see

<https://www.youtube.com/watch?v=8wLKuscyO9I>

14.9 Problem Set 6

1. (7.2 #2, #4, #5) Calculate the singular value decomposition (SVD) of the following matrices.

$$\begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 1 \\ 3 & 3 \end{bmatrix}, \begin{bmatrix} -1 & 2 \\ 2 & -4 \\ 0 & 0 \end{bmatrix}$$

In each case, draw a picture of the unit sphere in the domain and the hyperellipse that is its image in the codomain, and mark the vectors v_i and $\sigma_i u_i$ in the sphere and hyperellipse.

2. (a) (7.2 #13) If A is a square invertible matrix then use the SVD of A to compute the SVD of A^{-1} . What are the singular values of A^{-1} ?
- (b) (7.2 #14) If u_1, \dots, u_n and v_1, \dots, v_n are orthonormal bases for \mathbb{R}^n . Construct the matrix A such that $Av_1 = u_1, Av_2 = u_2, \dots, Av_n = u_n$ using what you know about SVD.
- (c) (7.2 #16) If A has orthogonal columns w_1, \dots, w_n of lengths $\sigma_1, \dots, \sigma_n$. What is the SVD of A ?
3. **True/false and why?** For the following questions, indicate whether they are true or false. If they are true, argue why it is so and if they are false, indicate why or give a counterexample if possible. Note that if an if and only if statement is true you will need to explain why both directions are true.

- (a) Let $A \in \mathbb{R}^{m \times n}$. A and A^\top can have different singular values.
- (b) Let $A \in \mathbb{R}^{n \times n}$. A is invertible if and only if 0 is not a singular value of A .
- (c) Let $A \in \mathbb{R}^{n \times n}$. $\|A\vec{x}\| = \|\vec{x}\|$ for all non-zero $\vec{x} \in \mathbb{R}^n$ if and only if all singular values of A equal 1.
- (d) If $A, B \in \mathbb{R}^{n \times n}$ are similar matrices (recall that they are similar if there exists a matrix C such that $A = CBC^{-1}$) then every singular value of A is also a singular value of B .
- (e) Let $A_1, A_2 \in \mathbb{R}^{n \times n}$ be matrices of equal rank. A_1 and A_2 have the same singular values if and only if there exist orthogonal matrices Q_1, Q_2 such that $A_1 = Q_1 A_2 Q_2$.
4. (a) Argue that for any two matrices A and B , $\text{rank}(A+B) \leq \text{rank}(A) + \text{rank}(B)$.
- (b) Find examples of 2×2 matrices where $\text{rank}(A+B) < \text{rank}(A) + \text{rank}(B)$ and where $\text{rank}(A+B) = \text{rank}(A) + \text{rank}(B)$.

5. (a) Check that

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} e & f \\ g & h \end{bmatrix} = \begin{bmatrix} a \\ c \end{bmatrix} [e \ f] + \begin{bmatrix} b \\ d \end{bmatrix} [g \ h].$$

- (b) Can you always find a decomposition of AB where $A \in \mathbb{R}^{m \times k}$ and $B \in \mathbb{R}^{k \times n}$ as a sum of rank one matrices as above? If yes, what is the formula? **Hint:** To do this correctly, you will need to check that the ij entry agrees on both sides.
- (c) Use the SVD of the first and second matrices in problem (2) to express each matrix as a sum of rank one matrices.
- (d) Use the reduced SVD of an arbitrary matrix $A \in \mathbb{R}^{m \times n}$ of rank r , to show that A can always be written as a sum of exactly r rank 1 matrices.

6. This problem explores the geometry of the SVD interpreted as a rotation + stretch + rotation (as in Section 7.4 of Strang's book). Draw figures at each step.

- Compute the SVD of the matrix $A = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$.
- Compute the image of the square in \mathbb{R}^2 with corners $(1, 1), (1, -1), (-1, 1), (-1, -1)$ under the linear transformation given by A .
- Compute the image of the square under the linear transformation V^\top . What is the geometric relationship between the original square and the one you just got?
- Compute the image of the square from part (c) under the matrix Σ in the SVD.
- Compute the image of the answer in (d) under the linear transformation U . Does your answer agree with what you computed in (b)?
- Track the point $p = (1, 0)$ through the above transformations on the figures you drew: find (i) coordinates of p in the V -basis, call it q . (ii) coordinates of Σq in the U basis and (iii) coordinates of Σq in the standard basis. Do you get Ap ?

7. Consider a rank one matrix $M = uv^\top$ of size 5×5 . (This problem works for any size $n \times n$.)

- Argue that M has at most one nonzero eigenvalue. Call it λ_1 .
- Check that $\lambda_1 = v^\top u$ and u is an eigenvector of M with eigenvalue λ_1 .
- Argue that M has exactly one nonzero singular value. Call it σ_1 .
- Argue that v is a singular vector of M by showing that it is an eigenvector of $M^\top M$.
- Using the above show that $\sigma_1 = \|u\|\|v\|$.
- Do you see a relationship between $|\lambda_1|$ and σ_1 ?

8. Using the command `svd(A)` in Julia you can compute the SVD of

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

```
julia> using LinearAlgebra

julia> A = [1 1 0 0 0; 0 1 1 0 1; 1 2 1 0 1]
3Array{Int64,2}:
 1  1  0  0  0
 0  1  1  0  1
 1  2  1  0  1

julia> svd(A)
SVD{Float64,Float64,Array{Float64,2}}
U factor:
3Array{Float64,2}:
-0.339287  0.742665  -0.57735
-0.473523 -0.665163  -0.57735
-0.81281  0.0775016  0.57735
```

```
singular values:
3-element Array{Float64,1}:
 3.253087102270064
 1.1905563006612325
 1.812986607347358e-16
```

```
Vt factor:
3Array{Float64,2}:
-0.354155 -0.749574 -0.395419  0.0 -0.395419
 0.688894  0.195291 -0.493603  0.0 -0.493603
 0.632456 -0.632456  0.316228  0.0  0.316228
```

What are the singular vectors $u_1, u_2, \dots, v_1, v_2 \dots$ that you see in the above SVD? You can use these labels in answering the following questions.

- What is the rank of A ?
- Find an orthonormal basis for the row space of A and column space of A .
- Which singular vector(s) lie in the nullspace of A^T ?
- Are some of the above singular vectors in the nullspace of A ? If yes, which ones? Do we get a basis for the nullspace of A ?
- Why is the fourth column of V^T filled with zeros?
- Write down the rank one decomposition of A . How many rank one matrices are there in the decomposition?
- What is the dimension of the unit ball and ellipsoid in the domain and codomain of the linear transformation given by A such that the ellipsoid is the image of the ball.
- What are the semiaxes of the ellipsoid?

9. (7.3 #4) Consider the matrix

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

- Compute the best fitting line L and best fitting plane P to the four columns of A . Express L and P as the span of vectors.
 - Compute the rank one decomposition of A , and the rank one approximations A_1 and A_2 .
 - Check that the columns of A_1 are the projections of the columns of A on L .
 - Check that the third column of A_2 is the projection of the third column of A on P .
 - Compute the 2-norm of $A - A_1$.
 - Construct a rank one matrix B of your choice and the same size as A . Check that B is not closer to A than A_1 .
10. (a) If A is a symmetric matrix of size $n \times n$, argue that $\sigma_i = |\lambda_i|$ for all i . Here σ_i is the i th singular value of A and λ_i is the i th eigenvalue of A .
- (b) If A is a psd matrix of size $n \times n$ then what is the relationship between its singular values and eigenvalues? What is the SVD of A ?

(c) Use the SVD to argue that for any square matrix A of size $n \times n$, $|\det(A)| = \sigma_1 \sigma_2 \cdots \sigma_n$.

11. (**Rank one matrices in $\mathbb{R}^{m \times n}$**) We saw in previous homework that matrices of the form uv^\top where u and v are nonzero vectors have rank one. You also saw that if you add two rank one matrices you get a matrix of rank at most two, and sometimes less than two. The following question explores more in this direction.

- (a) Argue using SVD that **all** rank one matrices of size $m \times n$ can be written as uv^\top for a vector $u \in \mathbb{R}^m$ and $v \in \mathbb{R}^n$.
- (b) Consider the rank one decomposition of a matrix $A \in \mathbb{R}^{m \times n}$ and let A_k be the partial sum of the first k rank one matrices in this decomposition. What is the rank of A_2 ? In general what is the rank of A_k ?
- (c) Can a rank k matrix always be written as the sum of k rank one matrices? What would be your algorithm for doing this?

12. (**Rank one psd matrices in $\mathbb{R}^{n \times n}$**) This exercise piggybacks on the previous one.


- (a) Argue that all rank one psd matrices of size $n \times n$ can be written as bb^\top for a vector $b \in \mathbb{R}^n$. (Note that unlike in 2(a) you have just one vector b now.)
- (b) Use SVD to argue that all $n \times n$ psd matrices of rank r can be written as a sum of r rank one psd matrices in which all coefficients are positive. (**Note:** The word sum is intended to mean a linear combination in which all coefficients are positive, or alternatively, something of the form $A_1 + A_2 + \cdots + A_n$ where A_i are rank one matrices. The word sum is used to emphasize that only “+” signs are used and any negatives are absorbed into the summands themselves)
- (c) Compute the SVD of the symmetric matrix

$$B = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 5 \\ 3 & 5 & 6 \end{bmatrix}.$$

- (d) Do you see a connection between the SVD of B and the diagonalization of B ?
- (e) Is it possible to write B as a sum of rank one **psd** matrices? The word sum has the same intended meaning as in part (b). **Hint:** you saw the relationship between σ_i and λ_i above.
- (f) In general, can a symmetric matrix of rank r be written as a sum of r rank one psd matrices? What is the difference (if any) with the result in (b)?
- (g) Describe 3×3 psd matrices of rank 1 that have 0s and 1s on the diagonal. **Hint:** Use part (a) to deduce what any rank one matrix looks like, then find necessary conditions for the diagonals of your psd matrix to be 0 or 1. You do not need to explicitly give all matrices but you should argue how many you could get (double counting is ok) and explain how you get them.

13. Polar decomposition of matrices and polar forms of complex numbers

We have seen that every complex number has the form $a + ib$. This allows us to think of complex numbers as living in the *complex plane* where the number $a + ib$ corresponds to the vector (a, b) .



We have heard (and possibly seen) that any complex number alternatively has the form $re^{i\theta} = \cos \theta + i \sin \theta$ where $r \geq 0$ is the length and $0 \leq \theta < 2\pi$ is the angle. This is more widely known as the *polar form* of a complex number and completely describes the number in terms of its length and angle.

In this problem we will derive the polar form of a complex number and see that the *polar decomposition* of a square matrix relates to the polar form of a complex number. Recall that given $A \in \mathbb{R}^{n \times n}$, the **polar decomposition** of A is a factorization

$$A = RS$$

where R is an orthogonal matrix and S is a psd matrix.

- (a) We can associate any complex number $a + bi$, with $a^2 + b^2 \neq 0$ to the matrix

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

Verify that multiplying and adding complex numbers is the same as multiplying and adding their associated matrices.

- (b) Recall the power series expansions of e^x , $\sin x$, and $\cos x$ given by


$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \frac{x^4}{4!} + \frac{x^5}{5!} + \frac{x^6}{6!} + \dots$$

$$\cos x = 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \frac{x^6}{6!} + \frac{x^8}{8!} - \frac{x^{10}}{10!} + \dots$$

$$\sin x = x - \frac{x^3}{3!} + \frac{x^5}{5!} - \frac{x^7}{7!} + \frac{x^9}{9!} - \frac{x^{11}}{11!} + \dots$$

Use these formulas to show that $e^{i\theta} = \cos \theta + i \sin \theta$. and conclude that any complex number of length r has the form $re^{i\theta}$

- (c) Setting $a_1 + ib_1 = r_1e^{i\theta}$ and $a_2 + ib_2 = r_2e^{i\gamma}$, determine what happens to the lengths and angles of the respective complex numbers when you multiply them together. Finish by arguing that multiplication of one complex number by another complex number is geometrically a stretch and a rotation applied to the first one.



- (d) Use the SVD to compute the polar decomposition of the matrix $A = \begin{bmatrix} a & -b \\ b & a \end{bmatrix}$ and argue that this matrix can be described as a “length” and an “angle”.

14. **Two Distances***

What is the largest number of points in \mathbb{R}^2 such that any two of them have the same distance? Three points are ok, we can put them at the vertices of an equilateral triangle, but there is no set of four points for which all pairwise distances are the same. Do you see why?

What if we allow **two** possible distances? A regular pentagon has two distances among pairs of vertices: all the diagonals have the same length and all the sides have the same length.




Figure 14.1: A regular pentagon

Question: What is the maximum number n of points in \mathbb{R}^d such that all pairwise distances among the points are one of two (positive) numbers?

It seems that n should depend on d , so better to write $n(d)$ instead of n . The above examples are maximal in the sense that $n(2) = 3$ and $n(3) = 5$. In this exercise we will show that $n(d) \leq \frac{1}{2}(d^2 + 5d + 4)$.

- (a) Let the n points in \mathbb{R}^d be $\mathbf{p}_1, \dots, \mathbf{p}_n$, and the two allowed distances be a and b . We have that the square of the distance between \mathbf{p}_i and \mathbf{p}_j is

$$\|\mathbf{p}_i - \mathbf{p}_j\|^2 = (p_{i1} - p_{j1})^2 + (p_{i2} - p_{j2})^2 + \dots + (p_{id} - p_{jd})^2 \in \{a^2, b^2\}$$

Associate to each \mathbf{p}_i the function

$$f_i : \mathbb{R}^d \rightarrow \mathbb{R}, \text{ such that } f_i(\mathbf{x}) = (\|\mathbf{x} - \mathbf{p}_i\|^2 - a^2)(\|\mathbf{x} - \mathbf{p}_i\|^2 - b^2)$$

where $\mathbf{x} = (x_1, \dots, x_d)$. Show that

$$f_i(\mathbf{p}_j) = \begin{cases} 0 & \text{for } i \neq j \\ a^2b^2 & \text{for } i = j \end{cases}$$

Hint: You might choose a few actual points \mathbf{p}_i in \mathbb{R}^2 or \mathbb{R}^3 and write out the function f_i to get a feel for this question.

- (b) Consider the set V of all functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$. Show that V is a vector space but checking all the requirements to be a vector space.

Hint: The sum of two functions, $f_1 + f_2$ is defined as $(f_1 + f_2)(\mathbf{x}) = f_1(\mathbf{x}) + f_2(\mathbf{x})$. If f is a function and $\alpha \in \mathbb{R}$, then αf is the function from $\mathbb{R}^d \rightarrow \mathbb{R}$ defined as $(\alpha f)(\mathbf{x}) = \alpha(f(\mathbf{x}))$.

- (c) Let W be the subspace of V spanned by the functions f_1, \dots, f_n that we defined before. Argue that f_1, \dots, f_n form a basis of W , i.e., they are linearly independent functions in V .
Hint: Suppose they are not, then there is a some linear combination of them $\alpha_1 f_1 + \alpha_2 f_2 + \dots + \alpha_n f_n = 0$ where 0 is the zero function that sends everything to 0. Use the definition of the functions f_i to show that this forces $\alpha_i = 0$ for all i proving what we want. **Hint:** Think about useful vectors that you could plug into the function $\alpha_1 f_1(\mathbf{x}) + \alpha_2 f_2(\mathbf{x}) + \dots + \alpha_n f_n(\mathbf{x}) = 0$ to show that $\alpha_i = 0$. A choice of one point will show that one of the $\alpha_i = 0$ and a choice of different point will show that $\alpha_j = 0$ for some $j \neq i$.
- (d) Remember we are trying to put an upper bound on $n(d)$. Here is a strategy: suppose we can find another set of functions g_1, \dots, g_t such that W lies in their span. Argue that $t \geq n(d)$. This means t is an upper bound on $n(d)$.

In the remaining part of this problem we will see how to find such functions g_1, \dots, g_t . Of course we want as small an upper bound as possible so we want a t that is as small as possible.

- (e) Check that each f_i is a polynomial of degree 4 in x_1, \dots, x_d .

An example of a degree four polynomial in $d = 3$ variables x_1, x_2, x_3 is

$$x_1^3 x_3 + x_2^2 x_3 - 15x_1 x_3 + 100.$$

Each such polynomial is a linear combination of *monomials*. The monomials in the above example are $x_1^3 x_3$, $x_2^2 x_3$, $x_1 x_3$ and 1. A monomial is a polynomial with one term and coefficient 1.

- (f) Argue that the set of all polynomials in d variables x_1, \dots, x_d , of degree at most 4, is a vector space. Call this vector space P .
- (g) Write out all monomials in x_1, x_2 of degree at most 4 and check that all polynomials in x_1, x_2 of degree at most 4 can be written as a linear combination of these monomials.
- (h) Argue that the vector space P is spanned by all monomials of degree at most 4 in x_1, \dots, x_d .

In fact, the monomials of degree at most 4 form a basis of P and there are precisely $\binom{d+4}{4}$ monomials of degree at most 4 in d variables.

- (i) Using the above, argue that

$$\binom{d+4}{4} \geq n(d).$$

This upper bound is a 4th degree polynomial in d . We need to get to a quadratic in d to get the result we want.

- (j) To get a smaller upper bound, we need to look for a smaller set of functions that span W . Maybe we don't need all monomials of degree at most 4 to generate W since the functions f_i have rather special structure. So we should look at it more carefully. Expand f_i and show that it is a linear combination of the following functions:

$$\begin{aligned} & (x_1^2 + \dots + x_d^2)^2 \\ & x_j(x_1^2 + \dots + x_d^2) \quad j = 1, 2, \dots, d \\ & x_j^2 \quad j = 1, 2, \dots, d \\ & x_i x_j \quad 1 \leq i < j \leq d \\ & x_j \quad j = 1, 2, \dots, d \\ & 1 \end{aligned}$$

(k) Show that there are $\frac{1}{2}(d^2 + 5d + 4)$ functions in the above list. Why is this number an upper bound on $n(d)$?

15. * **Equiangular lines in \mathbb{R}^d .** It is important in areas like coding theory to understand the largest number of lines through the origin in \mathbb{R}^d such that the angle between any two of them is the same. A collection of lines through the origin in \mathbb{R}^d such that the angle between any two of them is the same is a set of *equiangular* lines. For example, the maximum number of lines in \mathbb{R}^3 such that the angle between any two of them is 90° is 3 – take for example the 3 coordinate axes. However, if the common angle is different from 90° there can be more lines. The 6 diagonals of a regular icosahedron are equiangular. Google for the regular icosahedron if you haven't seen it before.

In the following exercise we will argue that you cannot have more than 6 equiangular lines in \mathbb{R}^3 no matter what angle you choose.

Suppose we have a collection of n equiangular lines in \mathbb{R}^3 and \vec{v}_i is a unit vector in the direction of the i th line. It does not matter if you choose \vec{v}_i or its negative, but choose one and call it \vec{v}_i .

(a) Argue that the condition of equiangularity means that if $i \neq j$ then $\vec{v}_i^\top \vec{v}_j = \cos \theta$ for a fixed angle θ .

(b) You showed last week that the set of all symmetric 3×3 matrices form a vector space. Argue that the dimension of this vector space is 6.

Hint: What is a basis of the space of all 2×2 symmetric matrices? Then try 3×3 .

Now consider the rank one psd matrices $\vec{v}_i \vec{v}_i^\top$ of size 3×3 .

(c) We will now argue that the rank one psd matrices $\vec{v}_i \vec{v}_i^\top$ of size 3×3 , coming from the equiangular lines, are **linearly independent** (as matrices). This will give us the result since all these rank one psd matrices are in the 6 dimensional vector space of symmetric matrices, there cannot be more than 6 of them. This will imply that there cannot be more than 6 lines.

i. Suppose the **matrices** $\vec{v}_i \vec{v}_i^\top$ are **linearly dependent**. Then there are $a_1, \dots, a_n \in \mathbb{R}$ such that

$$\sum a_i \vec{v}_i \vec{v}_i^\top = 0.$$

Multiply this expression on the left with \vec{v}_j^\top and on the right with \vec{v}_j and get that

$$0 = a_j + \sum_{i \neq j} a_i \cos^2 \theta.$$

ii. Find a matrix M such that you can express the equations from (b) in the form $M\vec{a} = 0$.

iii. Check that $M = (1 - \cos^2 \theta)I_n + \cos^2 \theta J_n$ where J_n is the $n \times n$ matrix of all ones.

iv. Argue that I_n and J_n are psd.

v. Argue that the coefficient $(1 - \cos^2 \theta)$ is positive and hence M is psd.

vi. Is M positive definite? If so, what is \vec{a} if $M\vec{a} = 0$?

vii. Conclude that the psd matrices $\vec{v}_i \vec{v}_i^\top$ are linearly independent.

(d) Can you extend each step above to see that you cannot have more than $\binom{d+1}{2} = \frac{d(d+1)}{2}$ equiangular lines in \mathbb{R}^d .

Chapter 15

Quotients of Vector Spaces and Error Detecting Codes

This chapter covers two completely separate topics. The notion of a quotient of a vector space is a central idea in many areas of pure mathematics. While it is very pure at its heart, it still has nice applications which yield surprising results, as we will soon see. After investigating quotients of various vector spaces we move on to error correcting codes and the ideas behind vector spaces over finite fields.

15.1 Quotient Spaces

When we first came across the idea of a subspace S , we may recall that the first and simplest property was that $\mathbf{0} \in S$. This property ensures that **we can never have parallel subspaces**. That being said, parallel lines to a given line through the origin in \mathbb{R}^2 are still worth thinking about and the idea of a quotient space allows us to still look at these parallel lines which are almost subspaces.

The first idea we meet is the idea of adding a vector to a subspace.


Definition 15.1.1. Let V be a vector space and $S \subset V$ be a subspace. Given $v \in V$ we define $v + S$ to be the subset of V defined by

$$v + S = \{v + s : s \in S\}$$

Example 15.1.2. Let S denote the line in \mathbb{R}^2 with equation $y = 2x$. That is

$$S = \left\{ \begin{bmatrix} x \\ 2x \end{bmatrix} : x \in \mathbb{R} \right\} = \left\{ \begin{bmatrix} x \\ y \end{bmatrix} \in \mathbb{R}^2 : y = 2x \right\}$$


The subset $\begin{bmatrix} 17 \\ 20 \end{bmatrix} + S$ is the line in \mathbb{R}^2 containing the point $(17, 20)$ with slope 2. As a subset, it is the line obtained by adding $\begin{bmatrix} 17 \\ 20 \end{bmatrix}$ to every point on S .



We can extend this idea to a more general one.

Definition 15.1.3. An **affine subset of V** is a subset of the form $v + S$ for some $v \in V$ and some subspace $S \subset V$. The affine subset $v + S$ is said to be parallel to S .

Example 15.1.4. Let $S = \left\{ \begin{bmatrix} x \\ y \\ z \end{bmatrix} \in \mathbb{R}^3 : z = 0 \right\}$ = the xy -plane. All affine subsets of \mathbb{R}^3 parallel to S are of the form $\mathbf{v} + S$ for some $\mathbf{v} \in \mathbb{R}^3$. Geometrically, they are all the planes parallel to the xy -plane



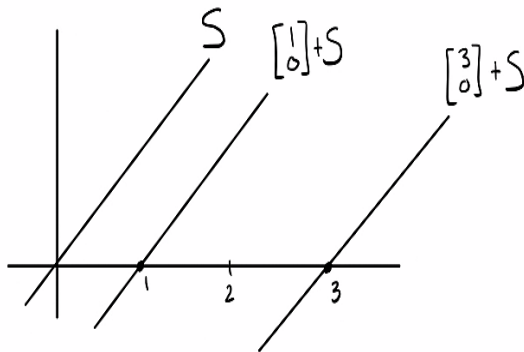
Note that we always have $v \in v + S$. This follows from the fact that $\mathbf{0} \in S$ hence the vector v , as an element of $v + S$ is obtained via $v + \mathbf{0}$. We can now generalize further and introduce the central idea.

Definition 15.1.5. Let S be a subspace of V . The **quotient space** V/S , sometimes pronounced “ V mod S ”, is the set of all affine subsets of V , parallel to S . In other words


$$V/S = \{v + S : v \in V\}$$

We will soon see that this is in fact a vector space in its own right and the “vectors” in this space are the affine subsets! Let’s revisit the same examples in this new context.

Example 15.1.6. If $S = \left\{ \begin{bmatrix} x \\ y \end{bmatrix} \in \mathbb{R}^2 : y = 2x \right\}$ then \mathbb{R}^2/S is the set of all lines in \mathbb{R}^2 of slope 2.



Example 15.1.7. If $S = \left\{ \begin{bmatrix} x \\ 0 \\ z \end{bmatrix} \in \mathbb{R}^3 : x, z \in \mathbb{R} \right\} =$ the xz -plane, then \mathbb{R}^3/S is the set of all planes parallel to the xz -plane. The planes themselves are the elements of this set.



Theorem 15.1.8. Given a real vector space V and a subspace $S \subset V$, the set V/S forms a real vector space with vectors given by the affine subsets of the form $v + S$ for $v \in V$

Proof. Let $v, w \in V$ and $r \in \mathbb{R}$. We define vector addition and scaling as follows:

- $(v + S) + (w + S) = (v + w) + S$.
- $r(v + S) = (rv) + S$.

where $v + w$ and rv are defined with the operations coming from V . □

Now that we have a vector space structure on V/S , we will want to know the answer to the following question.


Question 15.1.9. When are two affine subsets equal in V/S ?

The answer is simple but takes time to digest. We have that

$$v + S = w + S \quad \text{if and only if} \quad v - w \in S$$

Let's see why this is true.

If $v + S = w + S$ then in particular, we know that $v \in w + S$ because $v \in v + S$. Going further, this means that there exists some element $s \in S$ such that $v = w + s$ hence $v - w = s$ and $v - w \in S$.



In this picture, we can see that $\begin{bmatrix} 2 \\ 0 \end{bmatrix} + S = \begin{bmatrix} 3 \\ 2 \end{bmatrix} + S$ where S is the line $y = 2x$. We can also verify that $\begin{bmatrix} 3 \\ 2 \end{bmatrix} - \begin{bmatrix} 2 \\ 0 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \end{bmatrix} \in S$. The two vectors $\begin{bmatrix} 3 \\ 2 \end{bmatrix}$ and $\begin{bmatrix} 2 \\ 0 \end{bmatrix}$ are in some sense, indistinguishable in V/S , which leads us to the next, ultra important, idea.


The notion of a representative of an affine subset is an especially subtle detail that can be easily overlooked when dealing with quotient spaces. We call v a **representative** of the affine subset $v + S$ and if $v + S = w + S$ then both v and w are representatives of the same affine subset. The central component of V/S being a vector space is that **addition and scaling of vectors in V does not depend on the representative that we**

choose. This is one of the single most important facts about quotient spaces and one that makes our life considerably easier. Stated more formally, this means that if $v + S = v' + S$ and $w + S = w' + S$ then

$$(v + w) + S = (v' + w') + S$$

so to correctly add vectors in V/S , we don't need to worry about the representative that we choose to add with. In general, affine subsets have infinitely many choices for representatives and in many cases we want to take the simplest one. In either case, it is very useful to know when two affine subsets are the same for this reason.

Revisiting the example with S being equal to the line in \mathbb{R}^2 with equation $y = 2x$, we know that the vectors in \mathbb{R}^2/S are lines of slope 2 and we can add any two of them by adding any representatives of those lines.



We can see from that picture that

$$\left(\begin{bmatrix} 2 \\ 0 \end{bmatrix} + S \right) + \left(\begin{bmatrix} 5 \\ 0 \end{bmatrix} + S \right) = \begin{bmatrix} 7 \\ 0 \end{bmatrix} + S$$

and

$$\left(\begin{bmatrix} 3 \\ 2 \end{bmatrix} + S \right) + \left(\begin{bmatrix} 6 \\ 2 \end{bmatrix} + S \right) = \begin{bmatrix} 9 \\ 4 \end{bmatrix} + S$$

but $\begin{bmatrix} 9 \\ 4 \end{bmatrix} + S = \begin{bmatrix} 7 \\ 0 \end{bmatrix} + S$ because $\begin{bmatrix} 9 \\ 4 \end{bmatrix} - \begin{bmatrix} 7 \\ 0 \end{bmatrix} = \begin{bmatrix} 2 \\ 4 \end{bmatrix} \in S$

We end the section with some fundamental facts about quotient spaces in general. Since quotient spaces are vector spaces, we can investigate bases for them and in the process, find a nice formula for their dimension.

Definition 15.1.10. Let S be a subspace of V . There always exists a (linear!) map from V into the quotient space and this map is known as the **quotient map**

$$\pi : V \rightarrow V/S$$

We define this map by sending any vector $v \in V$ to the affine subset that it represents. That is

$$\pi(v) = v + S$$

The quotient map is what allows us to determine the dimension of any quotient space. Since it is a linear map we can look at its rank and nullity and this is the key.

Proposition 15.1.11.

$$\dim(V/S) = \dim(V) - \dim(S)$$

The first thing to observe is that the zero vector in the vector space V/S is the “vector” S . That is, the construction taking us from V to V/S creates a new vector space in which the subspace S is now playing the role of the zero vector. This is reflected in the fact that if $v + S \in V/S$, then $(v + S) + S = v + S$. If we want to look at addition with respect to adding representatives, then we can think of adding anything to S as adding its representative to $\mathbf{0} \in S$. Taking $\mathbf{0}$ as the representative of S is always the simplest choice. This is the key idea to the proof of this proposition.

Proof. Since the zero vector of V/S is S itself, we can conclude that $\ker(\pi) = \{v \in V : \pi(v) = S\}$ and

$$\pi(v) = v + S = S \quad \text{if and only if} \quad v \in S$$

therefore $\ker(\pi) = S$. Moreover, we know that every $v \in V$ lies in some affine subset parallel to S which implies that

$$\text{Range}(\pi) = V/S$$

That is, given any affine subset $v + S \in V/S$ we always have $\pi(v) = v + S$. We can now directly imply rank-nullity and deduce that

$$\dim(V) = \dim(\ker \pi) + \dim(\text{Range } \pi) = \dim(S) + \dim(V/S)$$

which means that

$$\dim(V/S) = \dim(V) - \dim(S)$$

□

Example 15.1.12. Let S be the line $y = 2x$ in \mathbb{R}^2 . We know that $\dim(\mathbb{R}^2) = 2$ and that $\dim(S) = 1$ hence $\dim(\mathbb{R}^2/S) = 1$. Since this new vector space is 1-dimensional, we can take any non-zero vector as a basis element. For example, we can say that

$$\mathbb{R}^2/S = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix} + S \right\}$$

We know that $\begin{bmatrix} 1 \\ 0 \end{bmatrix} + S$ is a non-zero vector in V/S because $\begin{bmatrix} 1 \\ 0 \end{bmatrix} \notin S$.

15.2 Quotients of Polynomial Vector Spaces

One of the most useful instances of quotient spaces comes from quotients of polynomial vector spaces like $\mathbb{R}[x]$. Since some fixed subspace S plays the role of the zero vector in V/S , we are able to talk about zeroes of polynomials in spaces like $\mathbb{R}[x]/S$, because subspaces in this case look like sets of polynomials. We investigate functions on hypercubes and remainders of polynomials upon division in this section and end with a nice application relating polynomials to determinants of matrices.


Let's begin by considering the vertices of the line segment $C_1 = [0, 1]$, namely, $V_1 = \{0, 1\} \subset \mathbb{R}$. Suppose $f(x) \in \mathbb{R}[x]$ is the univariate polynomial

$$f(x) = x^5 + 20x^4 + 8x^2 - 10$$

Observe that $f(0) = -10$ and $f(1) = 1 + 20 + 8 - 10 = 19$.

Question 15.2.1. How can we express the values that f takes on V_1 without having to look at a degree 5 polynomial?

The potential answer to this question is to find the linear function L that goes through both points



Surprisingly, we find the function L by looking at a quotient space of $\mathbb{R}[x]$!

First observe that the polynomial $p(x) = x^2 - x$ has both 0 and 1 as a root, that is, $p(0) = p(1) = 0$. This is the same as writing the factorization

$$x^2 - x = x(x - 1) = 0$$

If we divide $f(x)$ by $p(x)$ (remember long division from high school!) we obtain a new expression for $f(x)$, namely

$$\underbrace{x^5 + 20x^4 + 8x^2 - 10}_{f(x)} = \underbrace{(x^2 - x)}_{p(x)} \underbrace{(x^3 + 21x^2 + 21x + 29)}_{\text{multiple of } p \text{ that goes into } f} + \underbrace{(29x - 10)}_{\text{remainder}}$$

We write this compactly as

$$f(x) = p(x)m(x) + r(x)$$

where $m(x)$ is the polynomial that you get upon division and $r(x)$ is the remainder.

We can now check that $r(x)$ (a linear function!) has the same values that $f(x)$ does at the points 0 and 1.

$$f(0) = p(0)m(0) + r(0) = 0m(0) + r(0) = r(0)$$

and

$$f(1) = p(1)m(1) + r(1) = 0m(1) + r(1) = r(1)$$

This implies that the function $r(x) = 29x - 10$ satisfies

$$r(0) = -10 = f(0)$$

and

$$r(1) = 19 = f(1)$$

We can see that $r(x)$ is much simpler than $f(x)$ but conveys the same information about $f(x)$ on the points 0 and 1.

THIS ALWAYS WORKS!

That is, if $f(x)$ is any univariate polynomial and we have an expression

$$f(x) = p(x)m(x) + r(x)$$

, where $r(x)$ is the remainder of $\frac{f(x)}{p(x)}$, then $f(x) = r(x)$ at all the roots of $p(x)$. This is true for any polynomial $p(x)$ whose degree does not exceed that of f . Formalizing this a little bit more, we can say that every polynomial $f(x) \in \mathbb{R}[x]$ has a “representative” function $r(x)$ (for some given $p(x)$) where $\deg(r(x)) < \deg(p(x))$. Sticking with the example of $p(x) = x^2 - x$, we can say that every polynomial $f(x) \in \mathbb{R}[x]$ has a representative function $r(x)$ that is linear (degree 1), since $\deg(r(x)) < \deg(p(x)) = 2$, in this case. Conversely, if $f(x)$ is a linear polynomial, it is its own representative function, as a function on $V_1 = \{0, 1\}$. We can now formalize this completely.

Definition 15.2.2. Let $(x^2 - x)$ denote the subspace of $\mathbb{R}[x]$ containing all polynomial multiples of $x^2 - x$. That is

$$(x^2 - x) = \{g(x)(x^2 - x) : g(x) \in \mathbb{R}[x]\}$$

We say $\mathbb{R}[x]/(x^2 - x)$ is the vector space of all representatives of polynomial functions on $V_1 = \{0, 1\}$ (the roots of $x^2 - x$).

More generally, given $p(x) \in \mathbb{R}[x]$, $\mathbb{R}[x]/(p(x))$ is the vector space containing all representatives of polynomial functions on the roots of $p(x)$. **Each polynomial is represented by its remainder, after division by $p(x)$.**

There are some less confusing and more intuitive ways of thinking about $\mathbb{R}[x]/(p(x))$, all of which are valid and can be used.

1. $\mathbb{R}[x]/(p(x))$ “ = ” {All polynomials of degree less than that of $p(x)$ }. This way of viewing it follows from the fact that the remainder of any polynomial, after division by $p(x)$ must have strictly smaller degree than that of $p(x)$.

$$(a) \mathbb{R}[x]/(x^2 - x) \text{ “ = ” } \{\text{All linear polynomials in } x\}.$$

2. The vector space of all polynomials, under the relation that $p(x) = 0$. We can view it in this way because in V/S , the subspace S becomes the zero vector. In this case, S is multiples of $p(x)$, hence $p(x) = 0$ in $\mathbb{R}[x]/(p(x))$.

- (a) In $\mathbb{R}[x]/(x^2 - x)$, we have $x^2 - x = 0$ hence every instance of x^2 can be replaced with x (since $x^2 - x = 0 \implies x^2 = x$). We can apply this procedure to any polynomial, and we are done when we obtain a polynomial of degree strictly less than 2. This is in many ways, the best way to think about elements of $\mathbb{R}[x]/(p(x))$.

The second point described above is extremely useful. Using the relation introduced by $p(x)$ will always allow us to reduce a given polynomial down to its remainder as follows.

Example 15.2.3. Let $p(x) = x^2 - x$ and consider $\mathbb{R}[x]/(x^2 - x)$. The relation $x^2 - x = 0$ allows us to replace any instance of x^2 with x . We can apply this procedure to $f(x) = x^5 + 20x^4 + 8x^2 - 10$. In the quotient space, we have the following string of equalities

$$\begin{aligned} f(x) &= x^5 + 20x^4 + 8x^2 - 10 = x(x^2)^2 + 20(x^2)^2 + 8x^2 - 10 \\ &= x(x)^2 + 20(x)^2 + 8x^2 - 10 = x^2 + 28x^2 - 10 = x + 28x - 10 = 29x - 10 \end{aligned}$$

Using the method of this example, we can see that every polynomial can be reduced down to a linear one, and it is this linear polynomial that represents any $f(x) \in \mathbb{R}[x]$, in this quotient space. That is, every polynomial in this space is represented by a linear one, hence a basis for $\mathbb{R}[x]/(x^2 - x)$ is $\{1, x\}$.

Before moving on to multivariable polynomials, we make note of one last crucial fact that will prove useful in the next section. Recall that two elements $v, w \in V/S$ are equal if and only if $v - w \in S$. The same definition applies to $\mathbb{R}[x]/(p(x))$ since it is a quotient space, but we can obtain an even more specific definition of equality in this case, which will follow from the definition of the subspace $(p(x))$.

Equality in $\mathbb{R}[x]/(p(x))$

If $p(x) \in \mathbb{R}[x]$, then two polynomials $q(x), h(x) \in \mathbb{R}[x]/(p(x))$ are **equal** if there exists a polynomial $m(x) \in \mathbb{R}[x]$ such that $q(x) - h(x) = m(x)p(x)$. In other words, $q(x) = h(x)$ in $\mathbb{R}[x]/(p(x))$ if and only if $q(x) - h(x) \in (p(x)) =$ the subspace containing all multiples of $p(x)$.

Now, we move away from the single variable case. We mention that by looking at the quotient by a subspace of the form $(p(x))$, we are looking at values that all polynomials take on the roots of $p(x)$. The roots of a single variable polynomial are all just numbers, but what would we do if we wanted to look at values that a function takes on points of the form (x, y) ? To do this, we would need to look at multivariable polynomials. The notion of division still carries over to this setting, as do all the other properties we have already seen.

Consider the vertices of the unit square $V_2 = \{(0, 0), (1, 0), (0, 1), (1, 1)\}$. In terms of polynomials, the set of points V_2 is the set of simultaneous solutions to $x^2 - x = 0$ and $y^2 - y = 0$. In light of our previous work, this is explained by division of a multivariable polynomial by both $x^2 - x$ and $y^2 - y$.

In general, division by multivariable polynomials is much more delicate than the single variable case, however, by using relations that come from $x^2 - x = 0$ and $y^2 - y = 0$, we can obtain the desired remainder in a similar fashion that we done before. Letting $\mathbb{R}[x, y]$ denote the vector space of all polynomials in two variables, with real coefficients, we can replace any instances of x^2 or y^2 with x and y respectively.

Example 15.2.4. Let $(x^2 - x, y^2 - y)$ denote the subspace consisting of (sums of) polynomial multiples of $x^2 - x$ and $y^2 - y$. That is

$$(x^2 - x, y^2 - y) = \{m(x, y)(x^2 - x) + n(x, y)(y^2 - y) : m(x, y), n(x, y) \in \mathbb{R}[x, y]\}$$

We can define the quotient space $\mathbb{R}[x, y]/(x^2 - x, y^2 - y)$ to be obtained by the usual construction, where elements of the subspace $(x^2 - x, y^2 - y)$ are set equal to zero and are used as a relation to reduce polynomials

down to their representatives.

Let $f(x, y) = x^2 + x^2y + y^3 \in \mathbb{R}[x, y]$. In $\mathbb{R}[x, y]/(x^2 - x, y^2 - y)$, we have $x^2 = x$ and $y^2 = y$ hence

$$f(x, y) = x^2 + x^2y + y^3 = x^2 + (x^2)y + (y^2)y = x + (x)y + y(y) = x + xy + y^2 = x + xy + y$$

Using the example as our main model for how to work in a quotient space like this, we can see that a basis for $\mathbb{R}[x, y]/(x^2 - x, y^2 - y)$ is $\{1, x, y, xy\}$.

We finish the section by describing these quotient spaces in complete generality.

Let $V_n = \{(x, y) : x \in \{0, 1\}, y \in \{0, 1\}\} =$ vertices of an n -dimensional hypercube. Then

$$\mathbb{R}[x_1, \dots, x_n]/(x_1^2 - x_1, x_2^2 - x_2, \dots, x_n^2 - x_n)$$

is the set of all polynomial functions on V_n . All polynomials in this vector space can be reduced down to their representatives via the relations $x_i^2 = x_i$. It is a vector space spanned by all square-free monomials in x_1, \dots, x_n .

We finish by emphasizing (one last time) that division by $x_i^2 - x_i$ replaces x_i^2 by x_i , everywhere that it occurs. Given any polynomial $f(x_1, \dots, x_n) \in \mathbb{R}[x_1, \dots, x_n]$, we can replace all instances of x_i^2 with x_i until there are no more multiples of x_i^2 . This yields a representative of $f(x_1, \dots, x_n)$ in the quotient space.

15.3 Application: Solving a Univariate Polynomial

As we have seen, a univariate polynomial is a polynomial in one variable. As an example, consider

$$p(x) = x^5 - 20x^4 + 8x^2 - 10$$

The polynomial $p(x)$ has *degree* 5. The *terms* of $p(x)$ are $x^5, -20x^4, 8x^2$ and -10 . Note that there is no x^3 term and no x term in $p(x)$. The coefficients of $p(x)$ are the coefficients of all terms in $p(x)$ of degree at most 5 including the missing terms. In this example, the coefficients of $p(x)$ are $-10, 0, 8, 0, -20, 1$ written in increasing order of the degree of the term they appear with. The *leading term* of $p(x)$ is x^5 and its *leading coefficient* is 1.

The general univariate polynomial of degree d is

$$p(x) = a_d x^d + a_{d-1} x^{d-1} + a_{d-2} x^{d-2} + \dots + a_1 x + a_0$$

where the coefficients $a_0, a_1, \dots, a_{d-1}, a_d$ are real numbers and $a_d \neq 0$.

Goal: We wish to solve $p(x) = 0$.

All the values of x for which $p(x) = 0$ are called the *roots* of $p(x)$. By the **Fundamental Theorem of Algebra**, if $\text{degree}(p(x)) = d$ then $p(x)$ has d roots, some of which may be complex or might occur more than once.

Example 15.3.1. The general quadratic polynomial is $q(x) = ax^2 + bx + c = 0$ where $a, b, c \in \mathbb{R}$. Its two roots are

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

The quantity $b^2 - 4ac$ is the *discriminant* of the quadratic polynomial $q(x)$. The quadratic formula shown above tells us the following:

1. If $b^2 - 4ac > 0$, then $q(x)$ has two real roots.

2. If $b^2 - 4ac = 0$ then $q(x)$ has a double real root,
3. and if $b^2 - 4ac < 0$, $q(x)$ has two complex roots of the form $\alpha + i\beta$ and $\alpha - i\beta$.

Try making examples of quadratics that have all these possibilities. □

Idea: Suppose we could find a $d \times d$ matrix A_p such that $p(x)$ is the characteristic polynomial of A_p . Then the roots of $p(x)$ would be the eigenvalues of A .

We will see that this is always possible!!

Here is the algorithm:

1. Input: $p(x) = a_d x^d + a_{d-1} x^{d-1} + a_{d-2} x^{d-2} + \cdots + a_1 x + a_0$
2. Set up the following $d \times d$ matrix

$$A_p = \begin{bmatrix} 0 & 0 & 0 & 0 & \cdots & 0 & -a_0/a_d \\ 1 & 0 & 0 & 0 & \cdots & 0 & -a_1/a_d \\ 0 & 1 & 0 & 0 & \cdots & 0 & -a_2/a_d \\ 0 & 0 & 1 & 0 & \cdots & 0 & -a_3/a_d \\ \vdots & & & & \ddots & & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 & -a_{d-1}/a_d \end{bmatrix}$$

3. The characteristic polynomial of A_p , namely $\det(A_p - \lambda I)$, is either $p(\lambda)$ or $-p(\lambda)$. Therefore, the eigenvalues of A_p are the roots of $p(x)$.

Let's see an example of this in action.

Example 15.3.2. Using Julia let's find the roots of the polynomial $p(x) = x^5 - 20x^4 + 8x^2 - 10$. See Julia documentation on polynomials if the following commands are not self evident.

```
julia> using Polynomials

julia> p = Poly([-10,0,8,0,-20,1])
Poly(-10 + 8*x^2 - 20*x^4 + x^5)

julia> roots(p)
5-element Array{Complex{Float64},1}:
-0.6685161487282535 - 0.4961354572272264im
-0.6685161487282535 + 0.4961354572272264im
 0.6785047807938136 - 0.5116507106132901im
 0.6785047807938136 + 0.5116507106132901im
19.980022735868893 + 0.0im
```

The polynomial $p(x)$ has 5 roots as expected. The last one is real. The other 4 are complex and they come in conjugate pairs: $a + ib$ and $a - ib$. This is a general fact: if a polynomial with real coefficients has complex roots then the complex roots come in conjugate pairs.

Now we set up the matrix A_p as follows. I call it A in Julia.

```
julia> A = [0 0 0 0 10; 1 0 0 0 0; 0 1 0 0 -8; 0 0 1 0 0; 0 0 0 1 20]
55 Array{Int64,2}:
 0  0  0  0 10
 1  0  0  0  0
 0  1  0  0 -8
 0  0  0  1 20
```

```

0 0 1 0 0
0 0 0 1 20

```

```
julia> using LinearAlgebra
```

```

julia> eigvals(A)
5-element Array{Complex{Float64},1}:
-0.6685161487282542 - 0.4961354572272264im
-0.6685161487282542 + 0.4961354572272264im
 0.678504780793814 - 0.5116507106132913im
 0.678504780793814 + 0.5116507106132913im
19.98002273586887 + 0.0im

```

Check that the eigenvalues of A are exactly the roots of $p(x)$.

Why does this work?

A mechanical proof: Check that the determinant of $A_p - \lambda I$ is exactly $p(\lambda)$ or $-p(\lambda)$. The best strategy is to expand the determinant along the last column and you will see that this will work out. Try it on some examples first and this might help you warm up for the general proof. You might need to know some proof techniques like induction to do this precisely. For right now, just check it on examples and you have a proof by example!

A more sophisticated proof: We will explain this on the example $p(x) = x^5 - 20x^4 + 8x^2 - 10$, leaving out several steps for you to complete as homework problem.

We are going to look at *polynomials mod $p(x)$* in the proof below.

Consider the set $\mathcal{B} = \{1, x, x^2, x^3, x^4\}$ which is the set of all monomials of degree less than the degree of p . Let $V = \text{Span}(\mathcal{B})$. The elements of V are all linear combinations of $1, x, x^2, x^3, x^4$, i.e., elements of the form $a_0 + a_1x + a_2x^2 + a_3x^3 + a_4x^4$ where $a_0, \dots, a_4 \in \mathbb{R}$, which are precisely all polynomials in x of degree at most 4.

Step 1: Check that V is a vector space. If you add two polynomials of degree at most 4 you get a polynomial of degree at most 4. If you scale a polynomial of degree at most 4 by a real number you again get a polynomial of degree at most 4.

Step 2: Note that any polynomial of degree at most 4 cannot be divided by $p(x)$. Therefore, it is its own remainder when you divide it by $p(x)$. On the other hand, the remainder of any polynomial $q(x)$ after division by $p(x)$ is a polynomial of degree at most 4. So the vector space V is the set of all remainders of all polynomials after division by $p(x)$.

This observation allows us to think of V as really the set of all polynomials in x , where each polynomial is represented by its remainder after division by $p(x)$. Note that many polynomials can have the same remainder when divided by $p(x)$, so one polynomial $r(x)$ in V is the representative of many polynomials. In particular, the 0 in V is the representative of all polynomials that have 0 remainder when divided by $p(x)$, namely all polynomials that are polynomial multiples of $p(x)$.

Step 3: Consider the linear transformation $T : V \rightarrow V$ that sends $r(x) \mapsto xr(x)$. In other words, $T(r(x))$ is the product of $r(x)$ with the variable x . For example

$$T(3x^2 + 5x - 1) = x(3x^2 + 5x - 1) = 3x^3 + 5x^2 - x$$

What happens if $T(r(x))$ is a polynomial of degree larger than 4? For example

$$T(x^4 + 5x - 1) = x(x^4 + 5x - 1) = x^5 + 5x^2 - x$$

Then we need to replace the answer, $x^5 + 5x^2 - x$, by its representative in V , which is the remainder after division by $p(x)$. Doing long division, we see that the remainder is $20x^4 - 3x^2 - x + 10$. Therefore, in this special vector space V of remainders,

$$T(x^4 + 5x - 1) = 20x^4 - 3x^2 - x + 10$$

You can check that T is a linear transformation by checking the following:

$$T(r_1(x) + r_2(x)) = x(r_1(x) + r_2(x)) = xr_1(x) + xr_2(x), \quad T(\alpha r(x)) = x(\alpha r(x)) = \alpha xr(x).$$

Note that if the degree of $x(r(x))$ exceeds 4 then you have to replace it with its remainder on division by $p(x)$.

Step 4: We are almost done. Suppose $\lambda \in \mathbb{R}$ and $f(x) \in V$ is an eigenvalue/eigenvector pair of T , meaning $T(f(x)) = \lambda f(x)$. This means that $xf(x) = \lambda f(x)$ since $T(f(x)) = xf(x)$. In other words, $(x - \lambda)f(x) = 0$ in the vectors space V . Now remember what 0 means: $(x - \lambda)f(x) = 0$ means that $p(x)$ divides $(x - \lambda)f(x)$, equivalently, there is some other polynomial $h(x)$ such that

$$h(x)p(x) = (x - \lambda)f(x).$$

This equality is in the usual sense. Let's look at the degrees on the left and right. On the right we have a polynomial of degree at most 5 since the degree of $f(x)$ is at most 4 as it came from V . Therefore, $h(x)$ must be just a real number and the degrees on both sides better be 5. This means $(x - \lambda)f(x)$ is a factorization of $p(x)$ (up to some scalar multiple perhaps) and one of its factors is $(x - \lambda)$. Therefore, λ is a root of $p(x)$.

Step 5: The matrix A_p is the matrix representing T in the basis $\mathcal{B} = \{1, x, x^2, x^3, x^4\}$. Therefore, its columns should be the coordinates of $T(1), T(x), T(x^2), T(x^3)$ and $T(x^4)$. Check for instance that $T(1) = x$ has coordinates $(0, 1, 0, 0, 0)$ with respect to \mathcal{B} since $x = 0(1) + 1(x) + 0(x^2) + 0(x^3) + 0(x^4)$. So $(0, 1, 0, 0, 0)$ is the first column of A_p . $T(x) = x^2$ has coordinates $(0, 0, 1, 0, 0)$ in the basis \mathcal{B} , and it is the second column of A_p etc. Finally, $T(x^4) = x^5$ which should be replaced by its remainder after division by $p(x) = x^5 - 20x^4 + 8x^2 - 10$. This is precisely $20x^4 - 8x^2 + 10$ whose coordinates in the basis \mathcal{B} is $(10, 0, -8, 0, 20)$. This is the last column of A_p .

In homework, you will do this in generality but the main idea is represented in this last step.

Some comments:

1. The matrix A_p is called the *companion matrix* of the polynomial $p(x)$.
2. This lecture shows you that every univariate polynomial is a determinant. Therefore the polynomials you get as determinants are not rare! It was, and continues to be, an interesting question in mathematics to ask if a polynomial (in many variables) can be expressed as a determinant of special classes of matrices. A positive answer can have good algorithmic consequences.

15.4 Problem Set 7

1. (a) Find the roots of the following polynomial using its companion matrix:

$$p(x) = x^7 - 3x^5 + 100x^4 - 2x - 5$$

and double check your answer by computing all roots of p in Julia.

- (b) Let $p(x) = x^n + a_{n-1}x^{n-1} + \dots + a_1x + a_0$ and recall that $\mathbb{R}[x]/(p)$ is defined to be the set of all (unique) remainders of polynomials in $\mathbb{R}[x]$ after division by $p(x)$. Argue that $\mathbb{R}[x]/(p)$ is a vector space over \mathbb{R} and find a basis for it.

Note: You can show that this is a vector space in a number of ways, depending on how you think of elements in $\mathbb{R}[x]/(p)$. You can think of them as remainders or you can think of them as elements of the form $f(x) + (p(x))$ like we did in the first lecture on quotient spaces.

- (c) Argue that the map

$$T : \mathbb{R}[x]/(p) \rightarrow \mathbb{R}[x]/(p) \quad \text{s.t.} \quad r(x) \mapsto xr(x)$$

is a linear transformation on the vector space $\mathbb{R}[x]/(p)$.

(**Note:** To do this correctly, you must check that the linear map is **well-defined** in addition to satisfying the properties of being a linear transformation. That is, if $q(x)$ and $h(x)$ represent the same element of $\mathbb{R}[x]/(p)$ then $T(q(x)) = T(h(x))$. You'll want to use the fact that $q(x) = h(x)$ in $\mathbb{R}[x]/(p)$ if and only if $q(x) - h(x) \in (p)$, where elements of (p) are all polynomial multiples of p .)

- (d) The companion matrix A_p is the matrix representing the above linear transformation. Justify the formula for A_p using the basis you computed in (b).

(**Hint:** Recall how to compute the matrix of a linear map by looking at what the map does to a basis of the domain vector space.)

2. (Warm-up for next week)

- (a) Let $q(x) = ax^2 + bx + c$ be a quadratic polynomial where a, b, c are real numbers. Suppose it has a complex root $\alpha + i\beta$. Then argue that $q(x)$ will also have $\alpha - i\beta$ as a root. **Hint:** You know that $q(\alpha + i\beta) = 0$. Write this out and see if it follows that $q(\alpha - i\beta) = 0$.
- (b) Argue that for any polynomial in one variable whose coefficients are all real, complex roots come in conjugate pairs.

3. (a) What are all the points in \mathbb{R}^3 that are (simultaneously) the solutions of the three polynomials $x^2 - 1 = 0, y^2 - 1 = 0, z^2 - 1 = 0$? Call this set C and draw a picture of it.

- (b) Consider the polynomial $f = x^3y - 3xy^2 + 4x^2y^2 + 2z^4$. The lowest degree polynomial r for which there exists polynomials g_1, g_2, g_3 such that $f = g_1(x^2 - 1) + g_2(y^2 - 1) + g_3(z^2 - 1) + r$ is the *remainder* of f after division by $x^2 - 1, y^2 - 1, z^2 - 1$. What procedure would you use to find the remainder of any polynomial g in $\mathbb{R}[x, y, z]$ after division by $x^2 - 1, y^2 - 1, z^2 - 1$? **Hint:** Polynomial long division doesn't work the same way with multivariable polynomials. You will need to think about f in a certain quotient space.

- (c) Use your procedure to find the remainder of the given f after division by $x^2 - 1, y^2 - 1, z^2 - 1$.

- (d) As with univariate polynomials, $\mathbb{R}[x, y, z]/(x^2 - 1, y^2 - 1, z^2 - 1)$ denotes the set of all remainders of polynomials $f(x, y, z) \in \mathbb{R}[x, y, z]$ after division by $x^2 - 1, y^2 - 1, z^2 - 1$. This is a vector space.

i. Which polynomials in $\mathbb{R}[x, y, z]$ correspond to 0 in this vector space?

ii. Find a basis for the vector space $\mathbb{R}[x, y, z]/(x^2 - 1, y^2 - 1, z^2 - 1)$.

- (e) In general, what is a basis for the vector space $\mathbb{R}[x_1, \dots, x_n]/(x_1^2 - 1, x_2^2 - 1, \dots, x_n^2 - 1)$?

Chapter 16

Error Correcting Codes

Nobody can get through an advanced linear algebra class without seeing some vector spaces over finite fields, which is partly why this chapter is here. Error correcting codes pop up in many useful areas of math, including secure encryption and ISBN numbers. Throughout this chapter, the word **exercise** is more of an optional suggestion than a requirement. Some exercises are reworded as problems at the end of the chapter, others are sprinkled throughout the text as a way to check your own understanding of the material, but may be skipped.

16.1 Vector spaces over finite fields

If you're not familiar with modular arithmetic, take a quick look at the handout on it.

In this lecture we will focus on $\mathbb{Z}_2 = \{0, 1\}$, the integers mod 2 or *binary numbers*. How do we add and multiply in \mathbb{Z}_2 ? The only thing to remember is that we should do all these operations as we normally would but everytime you get a number different from 0, 1 you should replace it by its representative in \mathbb{Z}_2 .

Add as follows:

$$0 + 0 = 0, 0 + 1 = 1 + 0 = 1, 1 + 1 = 0.$$

The reason why $1 + 1 = 0$ is because normally $1 + 1 = 2$ and the representative of 2 in \mathbb{Z}_2 is 0.

Multiplication: We can only use elements from \mathbb{Z}_2 and we should stay in \mathbb{Z}_2 :

$$0 \cdot 0 = 0, 0 \cdot 1 = 0, 1 \cdot 1 = 1$$

The multiplicative inverse of 1 is 1 since $1 \cdot 1 = 1$.

All other rules of addition and multiplication are the same as in \mathbb{R} : addition and multiplication are both commutative and associative and so on.

A set that has all these properties is called a *field*. I am being a bit vague about "all these properties". The set of real numbers \mathbb{R} , the set of rational numbers \mathbb{Q} and the set of complex numbers \mathbb{C} are all fields. \mathbb{Z}_2 is called the finite field of two elements. In general, for any prime number p , \mathbb{Z}_p is the finite field of p elements. In other words \mathbb{Z}_p has the properties of \mathbb{R} and \mathbb{C} . The word "finite" is used since the field \mathbb{Z}_p has only finitely many (i.e., p) elements.

Just as with \mathbb{R} and \mathbb{C} we can have vector spaces over \mathbb{Z}_2 . By this we mean sets that satisfy all the rules of being a vector space and where linear combinations have coefficients from \mathbb{Z}_2 and scalar multiplication only uses elements of \mathbb{Z}_2 .

Example 16.1.1. What is $(\mathbb{Z}_2)^3$? Remember \mathbb{R}^3 is the set of all triples of real numbers (a, b, c) . So $(\mathbb{Z}_2)^3$ must be all triples of elements from \mathbb{Z}_2 .

$$(\mathbb{Z}_2)^3 = \{(0, 0, 0), (1, 0, 0), (0, 1, 0), (0, 0, 1), (1, 1, 0), (1, 0, 1), (0, 1, 1), (1, 1, 1)\}$$

Note that these are precisely the corners of a special cube of side length 1 (called the *unit cube*) in \mathbb{R}^3 .

Check that $(\mathbb{Z}_2)^3$ is a vector space over \mathbb{Z}_2 just like \mathbb{R}^3 is a vector space over \mathbb{R} :

$$\text{If } a, b \in (\mathbb{Z}_2)^3 \text{ and } \alpha, \beta \in \mathbb{Z}_2 \text{ then } \alpha a + \beta b \in (\mathbb{Z}_2)^3.$$

For example, if $a = (1, 0, 0), b = (1, 0, 1), \alpha = 0, \beta = 1$ then $0(1, 0, 0) + 1(1, 0, 1) = (0, 0, 0) + (1, 0, 1) = (1, 0, 1)$ which is still in $(\mathbb{Z}_2)^3$. Try a few more examples.

$(\mathbb{Z}_2)^4 = \{(0, 0, 0, 0), (1, 0, 0, 0), (0, 1, 0, 0), \dots, (1, 1, 0, 0), \dots, (0, 1, 1, 1), (1, 1, 1, 1)\}$ which are precisely all the corners of the unit cube in \mathbb{R}^4 . It has $2^4 = 16$ elements. This is again a vector space over \mathbb{Z}_2 .

In general $(\mathbb{Z}_2)^t$ is a vector space over \mathbb{Z}_2 with 2^t elements.

Exercise 16.1.2. What is a basis for $(\mathbb{Z}_2)^3$? In general, $(\mathbb{Z}_2)^t$? What are the dimensions of these vector spaces?

Since $(\mathbb{Z}_2)^t$ is a vector space, it can have *subspaces*. These would be subsets of $(\mathbb{Z}_2)^t$ that themselves form a vector space over \mathbb{Z}_2 . We will see subspaces in the next section.

16.2 Error Correcting Codes

The following material is taken from the book *Thirty Three Miniatures* by J. Matoušek.

Suppose we wish to transmit a message as a string \mathbf{v} of 0s and 1s. The transmission channel could introduce errors. For example you might send the string $\mathbf{v} = 1011$ but the string your buddy receives is $\mathbf{w} = 1001$, which has one error. We assume that the probability of many errors is small, say the probability of two errors is very small, but there is a chance of one error. In general, the probability of k errors might be very small, but there is a significant chance of $k - 1$ or less errors. Error correcting codes will pad your original message with extra digits that can help you correct errors. Below we see how this works.

Example 16.2.1. Suppose the probability of two errors is very small but one error is quite possible. Then if we wish to send the string 1011, we could make the rule that we will triple every digit and send 111000111111. Then if your buddy receives 110000111111, they will know that there is an error in the first digit and the message really is 1011. Of course there might be more errors but since the chance of two errors is small, this assessment seems reasonable. The question is, do you really need to triple every digit? Can you be more economical?

One of the best known error correcting codes is the **Hamming code** which was discovered in the 1950s.

Example 16.2.2. Here is what a Hamming code would do with a string $\mathbf{v} = abcd$ where $a, b, c, d \in \{0, 1\}$. It would send $\mathbf{w} = abcdefg$ where

$$e = a + b + c \pmod 2, \quad f = a + b + d \pmod 2, \quad g = a + c + d \pmod 2$$

So if $\mathbf{v} = 1011$, the code would send $\mathbf{w} = 1011001$. We will see that this can correct one error.

Now here is where \mathbb{Z}_2 and $(\mathbb{Z}_2)^n$ come into the picture. Recall that the elements of $(\mathbb{Z}_2)^n$ are precisely all strings of length n from the *alphabet* $\{0, 1\} = \mathbb{Z}_2$. An element of $(\mathbb{Z}_2)^n$ is called a *string* or *word* of length n (or with n -bits) from the alphabet \mathbb{Z}_2 .

Definition 16.2.3. 1. A **code** of length n is any subset of $(\mathbb{Z}_2)^n$.

2. A **linear code** of length n is a subspace of $(\mathbb{Z}_2)^n$.

Example 16.2.4. Consider all 7-bit strings that the Hamming code in the previous example will produce starting with all 4-bit strings:

$$C = \{0000000, 0001011, 0010101, 0011110, 0100110, 0101101, 0110011, 0111000, 1000111, \\ 1001100, 1010010, 1011001, 1100001, 1101010, 1110100, 1111111\}$$

C is a code since it is a subset of $(\mathbb{Z}_2)^7$.

We will now see that the code C is in fact a linear code. To prove this we need to argue that C is a subspace of $(\mathbb{Z}_2)^7$. How do we prove such a thing?

Remember that there are two ways of representing a subspace. We can either find a generating set (or basis) or write the set as the solutions to a finite number of linear equations. All this also works over \mathbb{Z}_2 as long as all calculations are done mod 2. Let's look at these methods for a subspace C of $(\mathbb{Z}_2)^n$.

1. **By basis:** Suppose G is a $k \times n$ matrix whose rows form a basis of C . Then all elements of C are linear combinations of the rows of G and so we can write C as

$$C = \{y^\top G : y \in (\mathbb{Z}_2)^k\}$$

We call G a *generator matrix* of C .

How to use G to encode messages? If we need to send a message $\mathbf{v} \in (\mathbb{Z}_2)^k$ we would send the string $\mathbf{w} = \mathbf{v}^\top G$ which lies in C . If there is no error, then we can recover \mathbf{v} by solving $\mathbf{w} = \mathbf{v}^\top G$ which has a unique solution since the rows of G are linearly independent.

Exercise 16.2.5. It is always possible to choose G so that it looks like $G = [I_k \ A]$.

Hint: Maybe first think about why this is always possible in a subspace of \mathbb{R}^n .

Example 16.2.6. In our example C , we can take

$$G = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 & 1 \end{bmatrix}$$

Check that every one of the 16 elements in C is a linear combination of the four rows of G and there are no more linear combinations. This is one way to see that C is a linear code but this is laborious.

Now check where 1011 would be sent by G . It would go to $(1, 0, 1, 1)G = 1011001$ as we had before. Note that because G has the identity matrix I_4 at the start, the first four bits in $\mathbf{v}^\top G$ is exactly \mathbf{v} .

2. **By linear equations:** There is a trick to finding a system of equations $Px = 0$ that represents C . If $G = [I_k \ A]$, then $P = [-A^\top \ I_{n-k}]$. In our example,

$$P = \begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 & 1 \end{bmatrix}$$

Check that the 16 elements in C satisfy $Px = 0$. Can you see that there are no more solutions? This is a lot easier to check than the checking if all elements of C are combinations of the rows of G .

If you write out the equations, do you see that they say

$$a + b + c \equiv e \pmod{2}, \quad a + b + d \equiv f \pmod{2}, \quad a + c + d \equiv g \pmod{2}$$

The matrix P is called the *parity check matrix* of the code C . Note that all nonzero elements of $(\mathbb{Z}_2)^3$ appears as a column of P .

Exercise 16.2.7. If you have a basis of any subspace stored as the rows of $G = [I_k \ A]$, then show that the subspace is the set of solutions of $[-A^\top \ I_{n-k}]x = 0$.

To finish we need some coding theory terminology which will help us see that the Hamming code above can correct one error.

Definition 16.2.8. 1. The **Hamming distance** of two words \mathbf{u}, \mathbf{v} in $(\mathbb{Z}_2)^n$ is the number of places in which they are different. Mathematically:

$$d(\mathbf{u}, \mathbf{v}) = |\{i : u_i \neq v_i, \text{ for } i = 1, \dots, n\}|$$

For example, $d(1011, 1001) = 1$ since the two words differ only in position 3, while $d(1011, 1000) = 2$.

It is useful to think of the Hamming distance geometrically as the smallest number of edges of the unit cube that you need to walk on to go from \mathbf{u} to \mathbf{v} . The way to “walk” from a corner of the unit cube to another, is to start with the initial corner and move successively to a neighboring corner, which is a word that differs from the given word in exactly one digit. For example, we can “walk” from 1101 to 1000 by successively making the moves $1011 \rightarrow 1001 \rightarrow 1000$. You could have also done $1011 \rightarrow 1010 \rightarrow 1000$. There are other longer routes between 1011 and 1000 through the corners of the cube but the smallest number of steps needed is unique and this number of steps is the Hamming distance between 1101 and 1000, namely $d(1011, 1000) = 2$.

2. A code $C \subseteq (\mathbb{Z}_2)^n$ **corrects t errors** if for every $\mathbf{u} \in (\mathbb{Z}_2)^n$ there is at most one $\mathbf{v} \in C$ such that $d(\mathbf{u}, \mathbf{v}) \leq t$.

For example, our code C corrects one error, if for every $\mathbf{u} \in (\mathbb{Z}_2)^7$ there is at most one string in C at distance 1 or 0 from $\mathbf{u} \in (\mathbb{Z}_2)^7$. Our example is indeed a one-error correcting code. Please check on a few examples. We will prove this shortly.

3. The **minimum distance** of a code C is the smallest distance between any two words in C . Mathematically,

$$d(C) := \min\{d(\mathbf{u}, \mathbf{v}) : \mathbf{u}, \mathbf{v} \in C, \mathbf{u} \neq \mathbf{v}\}$$

In our example C , $d(C) = 3$. Check that any two code words differ in at least 3 bits and there are pairs with Hamming distance exactly 3.

Theorem 16.2.9. A code C corrects t errors if and only if $d(C) \geq 2t + 1$.

We will prove the following special case which should help you see how to prove the general case. Also, everything below is about 1-correcting codes. The arguments generalize.

Theorem 16.2.10. A code C corrects 1 error if and only if $d(C) \geq 3$.

Proof. Suppose $d(C) \leq 2$. Then there are two code words $\mathbf{u}, \mathbf{v} \in C$ such that $d(\mathbf{u}, \mathbf{v}) \leq 2$. This means that either \mathbf{u} and \mathbf{v} are neighboring vertices of the unit cube or there is way to walk along the edges of the unit cube in \mathbb{R}^n from \mathbf{u} to \mathbf{v} via a word \mathbf{w} which is also a corner of the unit cube. In the first case, there is a code word within distance 1 from \mathbf{u} and in the second case, the word \mathbf{w} has two code words within distance 1 from it. Either way, C cannot correct 1 error by definition. This proves that if C corrects 1 error then $d(C) \geq 3$.

For the converse, suppose C is not 1-correcting. Then there is some $\mathbf{w} \in (\mathbb{Z}_2)^n$ such that there are two or more code words within distance 1 of it. Suppose $\mathbf{u}, \mathbf{v} \in C$ are two of these code words. Then by the same argument as above, we can walk from \mathbf{u} to \mathbf{v} via \mathbf{w} in two steps. This in turn means that $d(C) \leq 2$. \square

How can we encode and decode given a 1-correcting linear code $C \subseteq (\mathbb{Z}_2)^n$ with generator matrix G of size $k \times n$ and parity check matrix P of size $(n - k) \times n$?

Given a word $\mathbf{v} \in (\mathbb{Z}_2)^k$ we encode it as $\mathbf{w} = \mathbf{v}^\top G \in C$. If we receive $\mathbf{w}' \in (\mathbb{Z}_2)^n$, then we look for a word $\mathbf{w}'' \in C$ closest to \mathbf{w}' in Hamming distance. Since $\mathbf{w}'' \in C$, there is a unique \mathbf{v}' such that $(\mathbf{v}')^\top G = \mathbf{w}''$. We declare \mathbf{v}' to be the decoding.

Why does this work? Suppose at most 1 error occurred during the transmission. Then \mathbf{w}' is at distance at most 1 from \mathbf{w} which lies in C . On the other hand, since C is 1-error correcting, there is at most one code word within distance 1 of \mathbf{w}' and this one code word must be \mathbf{w} and from that we recover \mathbf{v} .

To finish we will argue that Hamming codes are 1-correcting and so our example code C is 1-correcting.

Definition 16.2.11. Fix a positive integer l . The **generalized Hamming code** (for l) is the linear code in $(\mathbb{Z}_2)^n$ where $n = 2^l - 1$ with parity matrix P whose columns are all the nonzero elements of $(\mathbb{Z}_2)^l$. In particular, generalized Hamming codes are linear codes since they are solutions of $Px = 0$.

Example 16.2.12. In our code C , $l = 3$. The code words are in $(\mathbb{Z}_2)^7$ and $7 = 2^3 - 1$. The parity matrix P has all the nonzero elements of $(\mathbb{Z}_2)^3$ as columns.

Theorem 16.2.13. The generalized Hamming code C has $d(C) = 3$ and thus is a 1-error correcting code.

Proof. We first note that for any linear code C ,

$$d(C) = \min\{d(\mathbf{0}, \mathbf{u}) : \mathbf{u} \in C, \mathbf{u} \neq \mathbf{0}\}.$$

In other words, to compute the smallest distance between two code words, it is enough to compute the smallest distance between $\mathbf{0}$ and any code word. Suppose not. Then there are two nonzero code words \mathbf{u}, \mathbf{w} whose distance is $d(C)$. Now consider $\mathbf{0} = \mathbf{u} - \mathbf{u}$ and $\mathbf{w} = \mathbf{v} - \mathbf{u}$. Since C is a subspace, $\mathbf{0} = \mathbf{u} - \mathbf{u}$ and $\mathbf{w} = \mathbf{v} - \mathbf{u}$ are also in C and distances don't change under subtraction, so $d(C) = d(\mathbf{0}, \mathbf{w})$.

To prove our theorem, we need to show that $d(C) \geq 3$ which by the above is same as showing that $d(\mathbf{0}, \mathbf{w}) \geq 3$ for every nonzero $\mathbf{w} \in C$. This is in turn is same as showing that every nonzero $\mathbf{w} \in C$ has at least 3 nonzero bits. The parity matrix now helps. We can show that no word \mathbf{w} with at most 2 nonzero bits satisfies $P\mathbf{w} = 0$.

If \mathbf{w} had only one nonzero bit then $P\mathbf{w} = 0$ if and only if a column of P is $\mathbf{0}$, but this is not allowed in the definition of P . If \mathbf{w} had two nonzero bits and $P\mathbf{w} = 0$ then two columns of P are the same which is also not true. Therefore we are done. \square

Thus we see that our running example code C is 1-error correcting. The matrix P allows for easy decoding.

Decoding a generalized Hamming code: Suppose we send the code word \mathbf{w} and receive \mathbf{w}' . If at most one error has occurred we have $\mathbf{w}' = \mathbf{w}$ or $\mathbf{w}' = \mathbf{w} + \mathbf{e}_i$ for some $i \in \{1, 2, \dots, n\}$.

If $\mathbf{w}' = \mathbf{w}$ then $P\mathbf{w}' = 0$. If $\mathbf{w}' = \mathbf{w} + \mathbf{e}_i$ then $P\mathbf{w}' = P\mathbf{w} + P\mathbf{e}_i = P\mathbf{e}_i$ which is the i th column of P . Thus if there was at most one error, we can immediately tell if an error occurred and we see which bit was wrong, namely the i th bit was wrong and there is a unique correction.

16.3 Problem Set 8

1. Let C be a subspace of $(\mathbb{Z}_2)^n$, i.e., C is a linear code. (Note that parts (a) and (b) are true for ANY subspace of any vector space. The argument has nothing to do with $(\mathbb{Z}_2)^n$. In general, writing a subspace as the span of the rows of a matrix is hard to work with. Writing the same subspace as the kernel of a matrix is much more efficient and useful in practice. The first two problems investigate how to do this in general.)
 - (a) If G is a $k \times n$ matrix whose rows form a basis of C , then argue that we can always choose G to look like $G = [I_k \ A]$.
 - (b) Recall that every vector space can be written as a span of vectors or as the kernel of a matrix. Suppose C is the kernel of an $(n - k) \times n$ matrix P (i.e., $Px = 0$ is the representation of C by equations), then argue that P can be chosen to be $P = [-A^T \ I_{n-k}]$. (**Hint:** Given G as above, show that $GP^T = 0$ and deduce that $PG^T = 0$. Use this to show that the columns of G^T equal the kernel of P .)
 - (c) Consider the set $C = \{(0, 0, 0), (1, 1, 0), (1, 0, 1), (0, 1, 1)\}$.
 - i. Argue that C is a linear code, i.e., C is a subspace of $(\mathbb{Z}_2)^3$.

- ii. Find a matrix G and P for C (Examples of how we construct these matrices are on page 3 of the handout).
 - iii. Verify that $Px = 0$ for all $x \in C$. Notice that $x \notin C$ if and only if $Px \neq 0$. This a handy way to show that something is not a code word.
2. (a) A code C is said to be k -separated if $d(C) = k$.
- i. Draw the three-dimensional unit cube and locate the elements of C in Problem 4(c) among its corners.
 - ii. Use your picture to verify that the code C in Problem 4(c) is 2-separated. **Hint:** Think about what k -separated means in terms of how many edges of the cube you need to walk along to get from any given code word to another.
- (b) We proved that the Hamming code from Example 6 in the handout, in $(\mathbb{Z}_2)^7$, is 3-separated by using its parity matrix P . For this example, argue that every non-code word has *exactly one* code word within Hamming distance 1 from it.
- Hints: You should work through this problem using the following steps**
- i. Argue that each element of $(\mathbb{Z}_2)^7$ has 7 immediate neighbors (at Hamming distance 1 from it).
 - ii. Use the fact that C is 3-separated to argue that all 7 words of Hamming distance 1 from a given code word are not code words.
 - iii. Think of each code word along with its 7 neighbors as a cloud of 8 vertices of the 7-dimensional cube. Argue that two of these clouds do not intersect.
 - iv. Now argue that the union of all 16 clouds is $(\mathbb{Z}_2)^7$.
 - v. Use this to conclude that every non-code word has exactly one code word as an immediate neighbor on the 7-dimensional cube.
This property allows us to correct a non-code word by replacing it with the unique code word in its cloud – i.e., at Hamming distance 1 from it.
- (c) Suppose you receive the string $\mathbf{w} = 0111101$.
- i. Argue that this is not in the Hamming code. (Use Problem 4(c)(iii).)
 - ii. Find the unique decoding of this string. **Hint:** Think about what went wrong in $P\mathbf{w}$ and which single digit you can change in \mathbf{w} to get the unique code word at Hamming distance 1 from it. Remember, the matrix P from the handout is what we use to check if something is a code word.

3. Kakeya sets*

Let \mathbb{F} denote a finite field with q elements which we can think of as the integers mod q for some prime number q , i.e., $\mathbb{F} = \mathbb{Z}_q$. For example $\mathbb{F} = \{0, 1\}$ is the finite field with $q = 2$ elements which we called \mathbb{Z}_2 before. The vector space \mathbb{F}^n consists of all n -tuples of points (a_1, a_2, \dots, a_n) such that $a_i \in \mathbb{F}$ for all $i = 1, \dots, n$. Therefore, $|\mathbb{F}^n| = q^n$.

- (a) For $\mathbb{F} = \{0, 1\}$ write down \mathbb{F}^3 .
- (b) A *line* in \mathbb{F}^n is defined as follows: fix a point $\mathbf{a} \in \mathbb{F}^n$ and a vector $\mathbf{u} \in \mathbb{F}^n$. Then the line through \mathbf{a} in direction \mathbf{u} is

$$\ell_{\mathbf{a}, \mathbf{u}} = \{\mathbf{a} + t\mathbf{u} : t \in \mathbb{F}\}.$$

For $\mathbb{F} = \{0, 1\}$, compute all the lines in \mathbb{F}^3 that pass through $\mathbf{a} = (1, 0, 0)$. Note that each line consists of 2 points, namely $\{\mathbf{a}, \mathbf{a} + \mathbf{u}\}$. (Yes, this is weird. Vector spaces over finite fields are very strange objects to work with and the fact that a line consists of just two points is definitely strange.)

- (c) A set $K \subset \mathbb{F}^n$ is called a *Keakeya set* if it contains a line in every possible direction.

These lines don't have to go through the same point \mathbf{a} , but for every $\mathbf{u} \in \mathbb{F}^n$ there must be some $\mathbf{a} \in \mathbb{F}^n$ such that the line $\ell_{\mathbf{a}, \mathbf{u}}$ lies in K . Note that in \mathbb{F}^n there are 2^n possible directions, each one given by a vector $\mathbf{u} \in \mathbb{F}^n$.

Use your computation in (b) to argue that there is a Keakeya set in \mathbb{F}^3 with 7 elements.

In what follows we will prove the baby case of the following theorem: **If \mathbb{F} is a finite field with q elements, and $K \subseteq \mathbb{F}^n$ is a Keakeya set, then $|K| \geq \binom{q+n-1}{n}$.**

From now on assume that $\mathbb{F} = \{0, 1\}$, i.e., $q = 2$ (this is the baby case).

- (d) What is the statement of the theorem when $\mathbb{F} = \{0, 1\}$?

The trick to proving the theorem is to use polynomials in $\mathbb{F}[x_1, \dots, x_n]$. A polynomial $p(\mathbf{x})$ of degree at most d in $\mathbb{F}[x_1, \dots, x_n]$ is a linear combination of monomials in x_1, \dots, x_n of degree at most d and coefficients in \mathbb{F} . It looks like

$$p(\mathbf{x}) = \sum c_{\alpha_1, \dots, \alpha_n} x_1^{\alpha_1} x_2^{\alpha_2} \cdots x_n^{\alpha_n} \quad (16.3.1)$$

as $(\alpha_1, \dots, \alpha_n)$ varies over all vectors in \mathbb{N}^n whose sum is at most d , and the coefficients $c_{\alpha_1, \dots, \alpha_n} \in \mathbb{F}$. Here \mathbb{N} denotes the natural numbers which is just the set of all positive (non-zero) integers $\{1, 2, 3, \dots\}$.

- (e) Give two examples of polynomials in $\mathbb{F}[x_1, \dots, x_n]$ for your choices of n and d . Remember $\mathbb{F} = \{0, 1\}$.

Let $\mathbb{F}[x_1, \dots, x_n]_{\leq d}$ be the set of all polynomials in $\mathbb{F}[x_1, \dots, x_n]$ of degree at most d . Recall from last week's star problem that $\mathbb{F}[x_1, \dots, x_n]_{\leq d}$ is a vector space of dimension $\binom{n+d}{d}$.

- (f) Now suppose $\mathbf{a}_1, \dots, \mathbf{a}_N$ are points in \mathbb{F}^n where $N < \binom{n+d}{d}$, and $p(\mathbf{x})$ is a polynomial in $\mathbb{F}[x_1, \dots, x_n]$ of degree at most d , like the one in equation (1) above. Consider the linear system of equations in $c_{\alpha_1, \dots, \alpha_n}$ that you get by setting $p(\mathbf{a}_i) = 0$ for all $i = 1, \dots, N$. Argue that this system has at least one non-zero solution. Equivalently, that there is at least one non-zero polynomial $p(\mathbf{x})$ of degree at most d such that $p(\mathbf{a}_i) = 0$ for all $i = 1, \dots, N$.

(Hint: Do an example with $n = 3$ to see what this system looks like.)

- (g) We are now ready to prove the theorem in the case $q = 2$. Suppose K is a Keakeya set in \mathbb{F}^n and $|K| < \binom{n+1}{n} = n + 1$. Argue that there is a nonzero linear polynomial $\ell(\mathbf{x})$ in $\mathbb{F}[x_1, \dots, x_n]$ that vanishes at all points in K .

(Hint: This is a direct application of the previous result.)

Now that we have this linear polynomial, let's assume that it looks like the following

$$\ell(\mathbf{x}) = c_0 + c_1 x_1 + c_2 x_2 + \cdots + c_n x_n$$

where $c_0, c_1, \dots, c_n \in \mathbb{F}$.

- (h) Take any nonzero $\mathbf{u} \in \mathbb{F}^n$. Since K is a Keakeya set, there is some \mathbf{a} such that the line $\{\mathbf{a} + t\mathbf{u} : t \in \mathbb{F}\}$ lies in K . Assume $\mathbf{a} = (a_1, \dots, a_n)$ and $\mathbf{u} = (u_1, \dots, u_n)$.

- i. Write out the polynomial $f(t) := \ell(\mathbf{a} + t\mathbf{u})$ in the single variable t using the formula for ℓ .
- ii. Check that f is a polynomial of degree at most 1 in t . What is the coefficient of t in f ? Call it β .
- iii. Argue that $f(t) = 0$ for all values of $t \in \mathbb{F}$ using the definition of f and the property of ℓ from part (g).
If $f(t) = 0$ for all values of t then f must be the zero polynomial which means that all its coefficients are 0 and in particular, $\beta = 0$.
- iv. Since \mathbf{u} was any element of \mathbb{F}^n and we just saw that $\beta = 0$, conclude that $h(\mathbf{x}) = c_1x_1 + \cdots + c_nx_n$ vanishes on all $\mathbf{u} \in \mathbb{F}^n$.

Now we'll invoke the following cool result to get our contradiction.

Schwartz-Zippel theorem: Any nonzero polynomial h of degree d can vanish on at most dq^{n-1} many points of \mathbb{F}^n where \mathbb{F} is a finite field with q elements.

- (i) What does this theorem say about the polynomial $h(\mathbf{x}) = c_1x_1 + \cdots + c_nx_n$ from the previous part? (Remember $q = 2$.)
- (j) Conclude that we have a contradiction and that the $|K|$ cannot be smaller than $n + 1$.

History: Kakeya sets over finite fields were inspired by the Kakeya conjecture in \mathbb{R}^n which is still open. A Kakeya set in \mathbb{R}^n is a *compact* set K that contains a line segment of length 1 in every direction. If you had a needle of length 1, you could rotate it continuously in all directions inside a Kakeya set (sort of). See the Wikipedia page:

`{\em https://en.wikipedia.org/wiki/Kakeya_set}`

to see a Kakeya set. Also, google “Numberphile Kakeya Needle Problem” to see a video about this problem. It turns out that Kakeya sets can have arbitrarily small area in \mathbb{R}^2 which is crazy!

The Kakeya conjecture says that a Kakeya set in \mathbb{R}^n cannot be too small – it has *Hausdorff dimension* n , whatever that is. This conjecture is open, but when phrased over finite fields you just showed using linear algebra that indeed Kakeya sets in \mathbb{F}^n cannot be too small. Sometimes it's good to practice over finite fields when you have a difficult problem over \mathbb{R}^n . The Kakeya conjecture has been proved in \mathbb{R}^2 . Even though this problem seems like a silly puzzle it has surprisingly strong connections to many parts of mathematics like number theory, partial differential equations, harmonic analysis etc.

Chapter 17

Vector Spaces over \mathbb{C}

The first key step to understanding vector spaces over \mathbb{C} , is understanding the basics of \mathbb{C} itself. After introducing complex numbers, we move up to complex vectors, and finish with complex matrices and the complex version of the spectral theorem. The problems for this section involve an introduction to Fourier series and the fast Fourier transform.

17.1 Complex Numbers


Any complex number $z \in \mathbb{C}$ has the form $z = a + bi$ where $a, b \in \mathbb{R}$. $i = \sqrt{-1}$ denotes the imaginary unit ($i^2 = -1$). We call a the real part of z , denoted $\text{Re}(z)$ and we call b the imaginary part of z , denoted $\text{Im}(z)$. We can also define addition and multiplication of complex numbers in the usual sense.

$$(a + bi) + (c + di) = (a + c) + (b + d)i \quad \text{and} \quad (a + bi)(c + di) = (ac - bd) + (ad + bc)i$$

One of the most useful notions attached to a complex number is its complex conjugate.

Definition 17.1.1. If $z = a + bi$ then the complex conjugate of z , denoted \bar{z} , is the complex number $\bar{z} = a - bi$.

Note that if $z \in \mathbb{R}$ then $\bar{z} = z$. We can visualize numbers (and their conjugates) in the complex plane with imaginary and real axes representing the lengths a and b as follows:



z and \bar{z} play nicely together and have four fundamental equations that relate the two. Namely

$$z + \bar{z} = (a + bi) + (a - bi) = 2a = 2\operatorname{Re}(z) \in \mathbb{R}$$

$$z\bar{z} = (a + bi)(a - bi) = a^2 + b^2 \in \mathbb{R}$$

$$\overline{zw} = \bar{z}\bar{w} \quad \text{and} \quad \overline{z + w} = \bar{z} + \bar{w}$$

The notion of a complex conjugate also allows us to define the length (or modulus) of a complex number as the real number

$$|z| = \sqrt{z\bar{z}} = \sqrt{a^2 + b^2}$$

Note that

$$\frac{1}{z} = \frac{1}{a + bi} = \frac{a - bi}{a + bi} \frac{a - bi}{a - bi} = \frac{a - bi}{a^2 + b^2} = \frac{\bar{z}}{z\bar{z}} = \frac{\bar{z}}{|z|^2} \in \mathbb{C}$$

and if $a^2 + b^2 = |z| = 1$ then $\frac{1}{z} = \bar{z}$.

When taking powers of complex numbers, as written, it can be computationally difficult. The polar form of a complex number greatly reduces the difficulty in computing this and will be of central importance moving forward.

Given any $z \in \mathbb{C}$, we can write

$$z = |z|(\cos \theta + i \sin \theta) = r(\cos \theta + i \sin \theta) = re^{i\theta}$$

Given any complex number we can easily compute its real and imaginary parts using this formula.

Example 17.1.2. Let $z = 3 - 2i$. We can see that $|z| = \sqrt{9 + 4} = \sqrt{13}$ hence

$$z = \sqrt{13}\left(\frac{3}{\sqrt{13}} - \frac{2}{\sqrt{13}}i\right)$$

This means that $\cos \theta = \frac{3}{\sqrt{13}}$ and $\sin \theta = \frac{-2}{\sqrt{13}}$ hence we can find θ via $\theta = \cos^{-1}\left(\frac{3}{\sqrt{13}}\right) = \sin^{-1}\left(\frac{-2}{\sqrt{13}}\right)$

Recall from last weeks homework problem that $\cos \theta + i \sin \theta = e^{i\theta}$, hence if $z = r(\cos \theta + i \sin \theta)$ then $z = re^{i\theta}$ and we can easily compute powers of z . That is

$$z^n = (re^{i\theta})^n = r^n e^{in\theta} = r^n(\cos n\theta + i \sin n\theta)$$

Moreover, if $|z| = 1$, then $r = 1$ and $z^n = e^{in\theta}$. That is, z^n is just z rotated by θ , n times. The polar form tells us the all important fact that when you multiply two complex numbers you **multiply the lengths and add the angles**.

Now that we have a grasp on the elements of \mathbb{C} , we can look at vector spaces over \mathbb{C} .

17.2 Complex Vectors and the Vector Space \mathbb{C}^n

Define a complex vector to be $\mathbf{z} = \begin{bmatrix} z_1 \\ \vdots \\ z_n \end{bmatrix} \in \mathbb{C}^n$ where $z_j = a_j + ib_j$ and $a_j, b_j \in \mathbb{R}$. One can verify that \mathbb{C}^n is a vector space over \mathbb{C} (meaning the scalars are all complex numbers) with vectors of the form \mathbf{z} as above.

Addition and scaling is defined similarly. If $\mathbf{z} = \begin{bmatrix} z_1 \\ \vdots \\ z_n \end{bmatrix}$ and $\mathbf{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix}$ then

$$\mathbf{z} + \mathbf{w} = \begin{bmatrix} z_1 + w_1 \\ \vdots \\ z_n + w_n \end{bmatrix}$$

and for $a + bi \in \mathbb{C}$ we have

$$(a + bi)\mathbf{z} = \begin{bmatrix} (a + bi)z_1 \\ \vdots \\ (a + bi)z_n \end{bmatrix}$$

Similar to the notion of a conjugate for an element of \mathbb{C} , we have the notion of conjugate transpose for complex vectors (and matrices).

Definition 17.2.1. Given $\mathbf{z} = \begin{bmatrix} z_1 \\ \vdots \\ z_n \end{bmatrix} \in \mathbb{C}^n$, the **conjugate transpose** of \mathbf{z} to be

$$\bar{\mathbf{z}}^\top = [\bar{z}_1 \quad \cdots \quad \bar{z}_n]$$

Given any complex number z_i , recall that $\bar{z}_i z_i \in \mathbb{R}$. We can expand on this idea and use conjugate transposes to define norms of complex vectors. Given $\mathbf{z} \in \mathbb{C}^n$ one can check that $\bar{\mathbf{z}}^\top \mathbf{z} \in \mathbb{R}$. We define the **norm** of \mathbf{z} to be the real number

$$\|\mathbf{z}\| = \sqrt{\bar{\mathbf{z}}^\top \mathbf{z}}$$

To ease notation from here on out, we denote the conjugate transpose by $\mathbf{z}^* = \bar{\mathbf{z}}^\top$.

We make an important notational note here since norms (or lengths) of complex numbers and complex vectors are denoted differently and must be handled with care. Given $z \in \mathbb{C}$ we have $|z| = \sqrt{\bar{z}z}$ whereas the norm of a vector $\mathbf{z} \in \mathbb{C}^n$ is denoted $\|\mathbf{z}\| = \sqrt{\bar{\mathbf{z}}^\top \mathbf{z}}$. This will come up several times and should be clear from context whenever it arises.

Example 17.2.2. Let $\mathbf{z} = \begin{bmatrix} 2 - i \\ 3 + 5i \end{bmatrix}$ and $\mathbf{z}^* = [2 + i \quad 3 - 5i]$. We compute the norm of \mathbf{z} squared via

$$\mathbf{z}^* \mathbf{z} = [2 + i \quad 3 - 5i] \begin{bmatrix} 2 - i \\ 3 + 5i \end{bmatrix} = |2 - i|^2 + |3 + 5i|^2 = 5 + 34$$

We can extend this idea further to define an inner product on \mathbb{C}^n known as the Hermitian inner product. It plays the same role that the dot product does on \mathbb{R}^n .

Definition 17.2.3. Given $\mathbf{u} = \begin{bmatrix} u_1 \\ \vdots \\ u_n \end{bmatrix}, \mathbf{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \in \mathbb{C}^n$ we define their inner product to be

$$\mathbf{v}^* \mathbf{u} = [\bar{v}_1 \quad \cdots \quad \bar{v}_n] \begin{bmatrix} u_1 \\ \vdots \\ u_n \end{bmatrix} = \bar{v}_1 u_1 + \cdots + \bar{v}_n u_n$$

We say that \mathbf{u} and \mathbf{v} are orthogonal if $\mathbf{v}^* \mathbf{u} = 0$, or equivalently, if $\mathbf{u}^* \mathbf{v} = 0$.

Example 17.2.4. Let $\mathbf{u} = \begin{bmatrix} 1 \\ i \end{bmatrix}$ and $\mathbf{v} = \begin{bmatrix} i \\ 1 \end{bmatrix}$, then

$$\mathbf{u}^* \mathbf{v} = [1 \quad -i] \begin{bmatrix} i \\ 1 \end{bmatrix} = 0$$

Note that we can use this notion to define orthonormal sets as well. We say that \mathbf{u} and \mathbf{v} are orthonormal if both $\mathbf{u}^* \mathbf{v} = 0$ and $\|\mathbf{u}\| = \|\mathbf{v}\| = 1$

Before defining complex matrices we note that the notion of conjugate transpose is more natural than the notion of just transposing when dealing with complex vectors. In general we can not always guarantee that $\mathbf{z}^T \mathbf{z}$ is a real number, and the conjugation is needed.

17.3 $\mathbb{C}^{m \times n}$ and the Complex Spectral Theorem

We can now (finally!) define complex matrices. We define a complex matrix to be $A = (z_{ij}) \in \mathbb{C}^{m \times n}$ with $z_{ij} \in \mathbb{C}$ for all i, j .

The notion of a conjugate transpose carries over identically from the previous section. That is, if $A = (z_{ij}) \in \mathbb{C}^{m \times n}$ then its conjugate transpose is given by

$$A^* = (\bar{z}_{ji}) \in \mathbb{C}^{n \times m}$$

Example 17.3.1. Let $A = \begin{bmatrix} 1 & i \\ 0 & 1+i \end{bmatrix}$ then the conjugate transpose is $A^* = \begin{bmatrix} 1 & 0 \\ -i & 1-i \end{bmatrix}$.

We also mention that the usual properties of transpose carry over in the complex setting just like they did before. Namely that

$$(\mathbf{A}\mathbf{u})^* \mathbf{v} = \mathbf{u}^* (\mathbf{A}^* \mathbf{v}) \quad \text{and} \quad (\mathbf{A}\mathbf{B})^* = \mathbf{B}^* \mathbf{A}^*$$

We can now define the most important type of complex matrices.

Definition 17.3.2. A matrix $A \in \mathbb{C}^{n \times n}$ is **Hermitian** if $A = A^*$. A consequence of this definition is that every real symmetric matrix is Hermitian.

Example 17.3.3. If $A = \begin{bmatrix} 2 & 3-3i \\ 3+3i & 5 \end{bmatrix}$ then $A^* = \begin{bmatrix} 2 & 3-3i \\ 3+3i & 5 \end{bmatrix} = A$ hence A is Hermitian.

As we had with symmetric matrices, Hermitian matrices have three important properties that combine to give the main result of this section.

Proposition 17.3.4. *If A is Hermitian and $\mathbf{z} \in \mathbb{C}^n$ then $\mathbf{z}^* \mathbf{A} \mathbf{z} \in \mathbb{R}$.*

Proof. Recall that if $z \in \mathbb{C}$ and $\bar{z} = z$ then $z \in \mathbb{R}$. We know that $\mathbf{z}^* \mathbf{A} \mathbf{z} \in \mathbb{C}$ so taking the conjugate transpose is the same as taking the conjugate. Applying this operation we see that

$$(\mathbf{z}^* \mathbf{A} \mathbf{z})^* = \mathbf{z}^* \mathbf{A}^* (\mathbf{z}^*)^* = \mathbf{z}^* \mathbf{A}^* \mathbf{z} = \mathbf{z}^* \mathbf{A} \mathbf{z}$$

hence $\mathbf{z}^* \mathbf{A} \mathbf{z}$ equals its (conjugate) transpose, and $\mathbf{z}^* \mathbf{A} \mathbf{z} \in \mathbb{R}$. □

Proposition 17.3.5. *Every eigenvalue of a Hermitian matrix is real.*

Proof. Assume that $A^* = A$ and $\mathbf{A} \mathbf{z} = \lambda \mathbf{z}$ with $\lambda \in \mathbb{C}$. Then from proposition 17.3.4 we know that

$$\mathbf{z}^* \mathbf{A} \mathbf{z} = \mathbf{z}^* \lambda \mathbf{z} = \lambda \mathbf{z}^* \mathbf{z} = \lambda \|\mathbf{z}\|^2 \in \mathbb{R}$$

Since $\|\mathbf{z}\|^2 \in \mathbb{R}$, we must have $\lambda \in \mathbb{R}$ as well. □

Example 17.3.6. Continuing the example from above with $A = \begin{bmatrix} 2 & 3-3i \\ 3+3i & 5 \end{bmatrix}$ we have that

$$\det(A - \lambda I) = \det \left(\begin{bmatrix} 2-\lambda & 3-3i \\ 3+3i & 5-\lambda \end{bmatrix} \right) = \lambda^2 - 7\lambda - 8 = (\lambda - 8)(\lambda + 1)$$

Proposition 17.3.7. Let A be Hermitian and assume that $A\mathbf{z} = \lambda\mathbf{z}$, $A\mathbf{y} = \beta\mathbf{y}$ with $\lambda \neq \beta$. We always have $\mathbf{y}^*\mathbf{z} = 0$. That is, eigenvectors of a Hermitian matrix corresponding to different eigenvalues are always orthogonal.

Proof. First, observe that

$$\mathbf{y}^*A\mathbf{z} = \mathbf{y}^*\lambda\mathbf{z} = \lambda\mathbf{y}^*\mathbf{z}$$

Furthermore, since $A\mathbf{y} = \beta\mathbf{y}$ we have that

$$(A\mathbf{y})^* = \beta\mathbf{y}^* \implies \mathbf{y}^*A^* = \beta\mathbf{y}^*$$

Multiplying both sides of this equation by \mathbf{z} on the right (and using the fact that A is Hermitian) we can conclude that

$$\mathbf{y}^*A^*\mathbf{z} = \mathbf{y}^*A\mathbf{z} = \beta\mathbf{y}^*\mathbf{z}$$

Now that we have two equations involving $\mathbf{y}^*A\mathbf{z}$ we combine them and see that

$$\mathbf{y}^*A\mathbf{z} = \lambda\mathbf{y}^*\mathbf{z} = \beta\mathbf{y}^*\mathbf{z} \implies (\lambda - \beta)\mathbf{y}^*\mathbf{z} \implies \mathbf{y}^*\mathbf{z} = 0$$

since $\lambda \neq \beta$ □

Example 17.3.8. Carrying on with the same matrix from the previous two examples we have $A = \begin{bmatrix} 2 & 3-3i \\ 3+3i & 5 \end{bmatrix}$ with eigenvalues $\lambda = 8$ and $\lambda = -1$. The corresponding eigenvectors are $\mathbf{z} = \begin{bmatrix} 1 \\ 1+i \end{bmatrix}$ and $\mathbf{y} = \begin{bmatrix} 1-i \\ -1 \end{bmatrix}$ respectively. Computing their inner product we can see that

$$\mathbf{y}^*\mathbf{z} = [1+i \quad -1] \begin{bmatrix} 1 \\ 1+i \end{bmatrix} = (1+i) - (1+i) = 0$$

Furthermore, we can divide these vectors by their norms to obtain orthonormal vectors

$$\frac{\mathbf{z}}{\|\mathbf{z}\|} = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 \\ 1+i \end{bmatrix} \quad \text{and} \quad \frac{\mathbf{y}}{\|\mathbf{y}\|} = \frac{1}{\sqrt{3}} \begin{bmatrix} 1-i \\ 1 \end{bmatrix}$$

Since these vectors live in \mathbb{C}^2 we can conclude that they are actually an orthonormal basis (with respect to the Hermitian inner product). This phenomenon always happens with Hermitian matrices. We have an orthonormal basis of eigenvectors, therefore, we can diagonalize A by writing

$$A = Q \begin{bmatrix} 8 & 0 \\ 0 & -1 \end{bmatrix} Q^*$$

where

$$Q = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 & 1-i \\ 1+i & -1 \end{bmatrix}$$

This matrix is special in that its conjugate transpose equals its inverse. This is a consequence of the fact that its columns form an orthonormal basis and matrices like this deserve their own name.

Definition 17.3.9. A complex square matrix Q with the property that $Q^*Q = QQ^* = I$ is called a **unitary** matrix. It is the complex analogue of an orthogonal matrix.

We can now combine all these ideas to end the chapter with the all important complex spectral theorem.

Theorem 17.3.10. *If $A \in \mathbb{C}^{n \times n}$ is Hermitian then*

- All eigenvalues of A are real.
- \mathbb{C}^n has an orthonormal basis of eigenvectors of A .
- If Q is the eigenvector matrix for A , then Q is unitary.
- A is **unitarily diagonalizable**, i.e. there exists a diagonal matrix $\Lambda \in \mathbb{R}^{n \times n}$ and a unitary matrix Q such that

$$A = Q\Lambda Q^*$$

Example 17.3.11. Writing out the unitary diagonalization for the matrix $A = \begin{bmatrix} 2 & 3-3i \\ 3+3i & 5 \end{bmatrix}$ we get

$$\begin{bmatrix} 2 & 3-3i \\ 3+3i & 5 \end{bmatrix} = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 & 1-i \\ 1+i & -1 \end{bmatrix} \begin{bmatrix} 8 & 0 \\ 0 & -1 \end{bmatrix} \frac{1}{\sqrt{3}} \begin{bmatrix} 1 & 1-i \\ 1+i & -1 \end{bmatrix}$$

This theorem concludes the chapter but we note that there is **much** more to discover about the world of complex matrices. Many of the nice theorems and properties we have seen with real matrices have their complex analogues, and often times the statements for the complex setting simply involve interchanging the word conjugate, with conjugate transpose, symmetric with Hermitian, and orthogonal with unitary.

17.4 Problem Set 9

1. (9.2 #8, #9)

- (a) Which class of matrices does P belong to: invertible, Hermitian, unitary?

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

- (b) Compute P^2, P^3, P^{100} .

- (c) What are the eigenvalues of P ?

- (d) Find the unit eigenvectors of P and put them into the columns of a unitary matrix U . Check that any two of them are orthogonal. What property of P makes these eigenvectors orthogonal?

- (e) Is PP^* invertible, Hermitian, unitary, psd?

2. **(Inner products and norms)** The notion of an *inner product* on a vector space is a central to many areas of mathematics. In this problem, V stands for a vector space over \mathbb{C} .

Definition 17.4.1. An **inner product** on V is a function that takes an ordered pair (u, v) of elements of V to a number $\langle u, v \rangle \in \mathbb{C}$ and satisfies the following properties:

- **Positivity:** $\langle v, v \rangle > 0 \quad \forall v \in V, v \neq 0$.
- **Conjugate symmetry:** $\langle u, v \rangle = \overline{\langle v, u \rangle} \quad \forall u, v \in V$.
- **Linearity in the first slot:** $\langle \lambda u + \mu v, w \rangle = \lambda \langle u, w \rangle + \mu \langle v, w \rangle \quad \forall u, v, w \in V, \lambda, \mu \in \mathbb{C}$.

The usual dot product of two vectors in \mathbb{R}^n defined as $\langle u, v \rangle = u^\top v$ is an inner product on \mathbb{R}^n . Inner products lead to notions of orthogonality and norm in V :

Definition 17.4.2. • Two vectors $u, v \in V$ are **orthogonal** if $\langle u, v \rangle = 0$.

- The norm of a vector $v \in V$ is defined to be $\|v\| := \sqrt{\langle v, v \rangle}$.

- (a) The Hermitian inner product on \mathbb{C}^n is given by $\langle u, v \rangle = v^*u$.
- Show that $\langle u, v \rangle = \overline{\langle v, u \rangle}$ for any two $u, v \in \mathbb{C}^n$. **Hint:** Write u and v in coordinates and compute both sides explicitly.
 - Check the remaining two properties of positivity and linearity in the first slot.
 - Show that $\langle u, \lambda_1 v + \lambda_2 w \rangle = \overline{\lambda_1} \langle u, v \rangle + \overline{\lambda_2} \langle u, w \rangle$. We say that the Hermitian inner product is **conjugate linear** in the second slot. **Hint:** To do this correctly, you will need to use conjugate symmetry AND linearity in the first slot. Alternatively, you could just do it in coordinates and this second method doesn't require using the other two properties of inner products.
- (b) The **Cauchy-Schwarz inequality** says that for any $u, v \in \mathbb{C}^n$,

$$|\langle u, v \rangle| \leq \|u\| \|v\| \quad (17.4.1)$$

- Show why the Cauchy-Schwarz inequality holds in \mathbb{R}^n with the usual dot product. **Hint:** Recall that $u^\top v = \|u\| \|v\| \cos \theta$.
- The **triangle inequality** says that for all $u, v \in \mathbb{C}^n$:

$$\|u + v\| \leq \|u\| + \|v\|. \quad (17.4.2)$$

- Why is the triangle inequality true in \mathbb{R}^n (with the usual dot product as inner product)? **Hint:** Start with $\|u + v\|^2 = (u + v)^\top (u + v)$.
- Show that for any $z \in \mathbb{C}$, $z + \bar{z} = 2\operatorname{Re}(z)$.
- If $u, v \in \mathbb{C}^n$ check that $\|u + v\|^2 = \|u\|^2 + \|v\|^2 + 2\operatorname{Re}\langle u, v \rangle$ where $\operatorname{Re}(a + ib) = a$, the real part of the complex number $a + ib$. **Hint:** Use part (B).
- Show that the triangle inequality holds in \mathbb{C}^n under the Hermitian inner product. That is, that

$$\|u + v\| \leq \|u\| + \|v\|$$

(**Hint:** Argue that $\operatorname{Re}\langle u, v \rangle \leq \|u\| \|v\|$ by using Cauchy-Schwarz (remember $\langle u, v \rangle$ is a complex number). Then use part (C) to complete the square. Your answer should be a string of inequalities starting with $\|u + v\|^2$ and ending with $(\|u\| + \|v\|)^2$.)

- (c) Let u_1, \dots, u_n be an orthonormal basis for V with respect to an inner product that we denote by $\langle -, - \rangle$. In particular, $\|u_i\| = 1$ for all i and $\langle u_i, u_j \rangle = 0$ for all $i \neq j$. Given any $a \in V$ we know $a \in \operatorname{span}\{u_1, \dots, u_n\}$ i.e there exist coefficients c_i such that $a = \sum_{i=1}^n c_i u_i$. Argue that

$$a = \langle a, u_1 \rangle u_1 + \dots + \langle a, u_n \rangle u_n.$$

That is, show that $c_i = \langle a, u_i \rangle$. **Hint:** Look at a as an arbitrary linear combination of the basis vectors (like above) and take inner products of a with a *reasonable* choice of vectors.

In other words, the unique coefficients of a with respect to the basis u_1, \dots, u_n are the inner products $\langle a, u_i \rangle$. Check for yourself that you secretly use this fact all the time – when we say $x \in \mathbb{R}^n$ has coordinates x_1, \dots, x_n in the standard basis e_1, \dots, e_n , we have that $x_i = x^\top e_i$. This exercise allows an easy way to find the coordinates of a vector in V with respect to an orthonormal basis of V .

3. (**Fourier series**) In this problem let

$$V = \{f : \mathbb{R} \rightarrow \mathbb{C} : f \text{ is piece-wise continuous, } f(x) = f(x + 2\pi) \forall x \in \mathbb{R}, \text{ and } \int_{-\pi}^{\pi} |f(x)|^2 dx < \infty\}$$

In other words, V is the set of all piece-wise continuous functions from $\mathbb{R} \rightarrow \mathbb{C}$ that are periodic with period 2π . The last condition is a necessary one but after this part of the problem this condition will not come up again. Examples of such functions are $f(x) = \sin x$ or $g(x) = e^{ix} = \cos x + i \sin x$. Note that piece-wise continuous means that the function is continuous on intervals.

(a) Define an inner product between functions in V as follows:

$$\langle f, g \rangle = \int_{-\pi}^{\pi} f(x) \overline{g(x)} dx$$

Check that under this inner product the norm of a function $f \in V$ is $\|f\| = (\int_{-\pi}^{\pi} |f(x)|^2 dx)^{\frac{1}{2}}$.

(b) Argue that V is a vector space over \mathbb{R} . You may assume that sums and multiples of piece-wise continuous functions are piece-wise continuous, you just need to check that the latter two conditions are preserved by sums and multiplication by scalars. **Hint:** You may want to use the previous problem to check the last condition.

(c) Calculate the norm of the function $f(x) = e^{ikx}$ where k is a fixed integer.

(d) Compute the inner product of the functions e^{inx} and e^{ikx} for two integers n, k . (You should get two possible values depending on whether $k = n$ or not. You may use the fact that for any integer m , $\int_{-\pi}^{\pi} e^{imx} dx = \begin{cases} 2\pi & \text{if } m = 0 \\ 0 & \text{if } m \neq 0 \end{cases}$ We will see why this is true in part (g).)

(e) Consider the following infinite set of functions in V :

$$\phi_n(x) = \frac{1}{\sqrt{2\pi}} e^{inx}, \quad n = 0, \pm 1, \pm 2, \pm 3, \dots$$

Argue that the set of functions $\{\phi_n\}$ is orthonormal. That is, verify that $\langle \phi_n(x), \phi_m(x) \rangle = 0 \forall n \neq m$ and that $\|\phi_n(x)\| = 1 \forall n$.

(f) It turns out that the orthonormal functions $\{\phi_n(x)\}$ form a basis of V which means two things:

- V is an infinite dimensional vector space, and
- any $f \in V$ can be written uniquely as an infinite series of the form

$$f(x) = \sum_{-\infty}^{\infty} c_n e^{inx}$$

*This series is called the **Fourier series** of f and the coefficients c_n are called the **Fourier coefficients** of f .*

Using the result of 2(c), show that $c_n = \frac{1}{\sqrt{2\pi}} \langle f, \phi_n \rangle$.

(g) Show that the Fourier series of the periodic function $f(x) = x$ when $-\pi \leq x \leq \pi$ is

$$\sum_{n \neq 0} \frac{(-1)^{n+1}}{in} e^{inx}$$

Hints:

1) You'll first need to show that you can integrate e^{ix} in the usual way, treating i as a scalar. To do this, break it down in terms of sin and cos (treating i as a scalar) then conclude from basic integration that $\int e^{ix} = \frac{1}{i}e^{ix}$.

2) You will need to do some integration by parts. If you're not in a calculus mood you can just look up the needed integration formula.) **Get PICS**

Fourier series underlie Fourier Analysis which is the basis of Signal Processing in Electrical Engineering. All your Zoom calls work so well because of sophisticated signal processing — one of the greatest applications of linear algebra.

This question and the next are based on Section 9.3 in Strang.

4. (**The Discrete Fourier Transform**) By the fundamental theorem of algebra, the polynomial equation $x^n = 1$ has n complex roots. These are called the n th roots of unity and they are denoted as $\omega_n^0 = 1, \omega_n, \omega_n^2, \dots, \omega_n^{n-1}$. For example, the 4th roots of unity are $1, i, -1, -i$, the solutions of $z^4 = 1$.

- Check that $\omega_n = e^{\frac{2\pi i}{n}}$. (You need to check that ω_n^k is an n th root of 1 for all $k = 0, \dots, n-1$, i.e., $(\omega_n^k)^n = 1$ for all $k = 0, \dots, n-1$.)
- Write out all the 8th roots of unity and plot them on the unit circle in the complex plane.
- Show that $1 + \omega_n + \omega_n^2 + \dots + \omega_n^{n-1} = 0$ **Hint:** Use the factorization $x^n - 1 = (x-1)(1+x+x^2+\dots+x^{n-1})$.
- The n th Fourier matrix is the following symmetric matrix:

$$F_n = \begin{bmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & \omega_n & \omega_n^2 & \dots & \omega_n^{n-1} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & \omega_n^{n-1} & \omega_n^{2(n-1)} & \dots & \omega_n^{(n-1)^2} \end{bmatrix}.$$

- Write F_4 . Check that F_4 is not Hermitian.
 - Argue that $\frac{1}{2}F_4$ is unitary. (There are several different ways to do this) Note: In general, $\frac{1}{\sqrt{n}}F_n$ is unitary.
 - What is F_4^{-1} ? More generally, F_n^{-1} ?
- Consider $p(x) = a_0 + a_1x + a_2x^2 + a_3x^3$ and $a = (a_0, a_1, a_2, a_3)$ its vector of coefficients (a is a column vector, we just write it as a row vector to save space). Compute $F_4a = \hat{a}$.
 - More generally, suppose $p(x) = a_0 + a_1x + a_2x^2 + \dots + a_{n-1}x^{n-1}$ is a univariate polynomial and $a = (a_0, \dots, a_{n-1})$ is its vector of coefficients (a is a column vector). Argue that the components of $F_na = \hat{a}$ are, respectively, $p(1), p(\omega_n), p(\omega_n^2), \dots, p(\omega_n^{n-1})$, the evaluations of p are the n th roots of unity.

Comments:

- For any vector $a \in \mathbb{C}^n$, the vector $\hat{a} = F_na$ is called the *discrete Fourier transform* of a .
- Computer scientists are interested in how fast you can compute something. It takes about $(2n)^2$ additions and multiplications to compute F_na and they would say that the computation time is $O(n^2)$ steps after suppressing all the constants.

- The Fourier transform is what allows you to zoom in on your phone, in addition to many many other things.
- (g) If we are given the evaluations of a degree $n - 1$ polynomial $p(x)$ at the n th roots of unity, can we find the polynomial? How would you do it using Fourier matrices?

*Finding a polynomial from its values at specified points is called **interpolation** and has a huge number of applications. For example, the temperature in an experiment might be an unknown polynomial $p(t)$ in the time t . If you measure the temperature at various times, and you have enough measurements, you can use interpolation to find the temperature function.*

5. (**Fast Fourier Transform (FFT)**) This problem builds on the previous problem. We are going to show that there is a faster way to compute the discrete Fourier transform than by the straight multiplication $F_n a$. This in turn amounts to a clever recursive way to factorize Fourier matrices. See Section 9.3 in Strang for this factorization point of view. As an application we will see a fast way to multiply two polynomials together.

Assume throughout that n is a power of 2, say $n = 2^m$. You can always pad your polynomial with terms having 0 coefficient until it is a polynomial of degree $2^m - 1$, so there is no harm in this assumption.

- (a) If $n = 2^m$ then what are
- $(\omega_n)^{\frac{n}{4}}$?
 - $(\omega_n)^{\frac{n}{2}}$?
 - $(\omega_n)^{\frac{3n}{4}}$?
 - $(\omega_n)^n$?
 - Show that $\omega_n^{\frac{n}{2}+k} = -\omega_n^k$.
 - Show that $\omega_n^{n+k} = \omega_n^k$.
- (Use the $n = 8$ example from the previous problem as a guide.)
- (b) Given $p(x) = a_0 + a_1x + a_2x^2 + \cdots + a_{n-1}x^{n-1}$. Remember, we still have $n = 2^m$. Define the polynomials

$$p_0(x) = a_0 + a_2x + a_4x^2 + \cdots + a_{n-2}x^{\frac{n}{2}-1} \quad \text{and} \quad p_1(x) = a_1 + a_3x + a_5x^2 + \cdots + a_{n-1}x^{\frac{n}{2}-1}.$$

- Argue that $p(x) = p_0(x^2) + xp_1(x^2)$.
 - Argue that to compute the values of $p(x)$ at $1, \omega_n, \omega_n^2, \dots, \omega_n^{n-1}$ it suffices to compute the values of p_0 and p_1 at $1, (\omega_n)^2, (\omega_n^2)^2, \dots, (\omega_n^{n-1})^2$.
 - Argue that there are only $\frac{n}{2}$ elements in the list $1, (\omega_n)^2, (\omega_n^2)^2, \dots, (\omega_n^{n-1})^2$.
 - Suppose we use the method of discrete Fourier transforms from the previous problem to find the values of p_0 and p_1 at these $\frac{n}{2}$ roots of unity. What is the size of the Fourier matrices you would use and how would they look? **Hint:** Look at problem 4e and think about what part of the fourier matrix they need to do this. **You will not need the entire Fourier matrix to find the values of p_0 , respectively p_1 .**
- (c) We can now devise a recursive algorithm that breaks p_0 and p_1 into two polynomials say p_{00}, p_{01} and p_{10}, p_{11} respectively, using the same rule as above, and then each of these into two further polynomials and so on until we cannot break down any more. These can then be assembled back to produce p . Break $p(x) = a_0 + a_1x + a_2x^2 + \cdots + a_7x^7$ according to this recursive procedure and check that you get back p .

An aside: If this were a computer science class, we would calculate that this recursive divide-and-conquer algorithm takes $O(n \log_2 n)$ multiplications and additions to compute all the values of p at the n th roots of unity. This is faster than $O(n^2)$. This is the method of Fast Fourier Transform or FFT which was considered to be one of the top 10 algorithms of the 20th century.

Now we are going to apply this to multiply two polynomials together.

(d) Suppose you are given the following two polynomials:

$$p(x) = a_0 + a_1x + a_2x^2 + \cdots + a_{n-1}x^{n-1} \quad \text{and} \quad q(x) = b_0 + b_1x + b_2x^2 + \cdots + b_{n-1}x^{n-1}.$$

- i. Argue that to compute $p(x)q(x)$ you need about n^2 operations (additions and multiplications).
- ii. Here is the algorithm that uses FFT to speed up the computation of $p(x)q(x)$. These are the steps, your solution to this question will be using the algorithm on an example:
 - A. Compute the values of $p(x)$ and $q(x)$ at the $2n$ points $1, \omega_{2n}, \omega_{2n}^2, \dots, \omega_{2n}^{2n-1}$ which are the $2n$ th roots of unity.
 - B. Compute the evaluation of pq at the same roots of unity by computing the products

$$\begin{aligned} (pq)(1) &= p(1) \cdot q(1) \\ (pq)(\omega_{2n}) &= p(\omega_{2n}) \cdot q(\omega_{2n}) \\ &\vdots \\ (pq)(\omega_{2n}^{2n-1}) &= p(\omega_{2n}^{2n-1}) \cdot q(\omega_{2n}^{2n-1}) \end{aligned}$$

- C. Now that we have the values of pq at the $2n$ th roots of unity use part (d) of problem 4 to find the polynomial pq .

If you have background in algorithms, you can check that this takes only $O(n \log_2 n)$ steps.

Use the above procedure to multiply the polynomials:

$$p(x) = 1 + 2x + 3x^2 + 4x^3, \quad q(x) = x + 2x^2 + 3x^3$$

and check your answer by multiplying them as usual.

While this was chosen to be doable by hand, you will need to take much larger values of n to see the savings of this method.