

MAT 2114 Intro to Linear Algebra

Joe Wells
Virginia Tech

Spring 2020

Contents

Preface	4
1.1 The Geometry and Algebra of Vectors	5
1.1.1 Vector operations	6
1.1.2 Geometry of vector operations	7
1.1.3 Linear combinations	7
1.2 Length and Angle: The Dot Product	10
1.2.1 Length	10
1.2.2 Distances	12
1.2.3 Angles	13
1.2.4 Projections	14
2.1 Introduction to Linear Systems	16
2.2 Direct Methods for Solving Linear Systems	19
2.2.1 Row Operations	19
2.2.2 (Reduced) Row Echelon Form	20
2.2.3 Gaussian Elimination and Gauss–Jordan Elimination	22
2.2.4 Rank and Number of Solutions	24
2.3 Spanning Sets and Linear Independence	26
2.4 Span and Spanning Sets	26
2.4.1 Linear (In)dependence	27
2.4.2 Using Matrices to Determine Linear (In)dependence	29
3.1 Matrix Operations	31

3.1.1	Matrix Basics	31
3.1.2	Matrix Operations	32
3.1.3	Matrix Powers	34
3.1.4	Transpose	34
3.2	Matrix Algebra	36
3.2.1	Properties of Matrix Multiplication	36
3.3	The Inverse of a Matrix	39
3.5	Subspaces, Basis, Dimension, and Rank	44
3.5.1	Subspaces Associated with Matrices	45
3.5.2	Dimension and Rank	49
3.5.3	Coordinates	51
6.3	Change of Basis	53
3.6	Introduction to Linear Transformations	55
3.6.1	Types of Linear Transformations of \mathbb{R}^2	56
3.7	Applications	59
3.7.1	Markov Chains	59
4.1	Introduction to Eigenvalues and Eigenvectors	61
Appendix C: Complex Numbers		64
Matrices with complex entries		64
4.2	Determinants	65
4.2.1	Determinant of a 2×2 Matrix	65
4.2.2	Determinant of a $n \times n$ Matrix	67
4.2.3	Properties of Determinants	68
4.3	Eigenvalues and Eigenvectors of $n \times n$ Matrices	70
4.3.1	Relationship to Matrix Operations	71
4.4	Similarity and Diagonalization	74
4.4.1	Properties of similarity and similar matrices	74
4.4.2	Diagonalization	75
4.4.3	Computational power of diagonal matrices	78
5.1	Orthogonality in \mathbb{R}^n	79
5.1.1	Orthonormality	81

5.2	Orthogonal Complements and Orthogonal Projections	83
5.2.1	Orthogonal Complement	83
5.2.2	Orthogonal Projections	84
5.3	The Gram–Schmidt Process and the QR Factorization	87
5.3.1	Gram–Schmidt	87
5.3.2	QR Factorization	88
7.3	Least Squares	90
7.3.1	Least Squares Approximation	92

Preface

There are many different approaches to linear algebra, and everyone has their preference. This document is compiled from the course I taught in the Spring of 2020 at Virginia Tech, where both the book (*Linear Algebra: A Modern Introduction* 4th Ed. by David Poole) and order of topics covers were suggested to me by some others in the department. Although not formally stated anywhere, this class was largely geared towards math-adjacent students (engineering, physics, computer science, etc.) and so these notes and the presentation are at a lower level of abstraction (and occasionally rigor) than what one might experience in another introductory linear algebra course. In hindsight, I probably would have picked both a different text and order in which to introduce the topics – it seems perverse to leave the phrase “vector space” until the 6th chapter! Nevertheless, I did my best to gently introduce concepts as needed in order to more smoothly segue the topics, and many of the homework exercises are designed to bridge certain theoretical gaps in the material (which is precisely where I’ve hidden the definition of a vector space).


I would like to thank the many students who inadvertently served as my copy editors for the semester.

1.1 The Geometry and Algebra of Vectors

Let $A = (a_1, \dots, a_n)$ and $B = (b_1, \dots, b_n)$ be points in space.

Definition. The *vector* \mathbf{AB} is the arrow from A (the *initial point* or *tail*) to B (the *terminal point* or *head*). In coordinates, use square brackets and write $\mathbf{AB} = [b_1 - a_1, \dots, b_n - a_n]$, and the individual coordinates of this vector are called *components*. A vector \mathbf{AB} is in *standard position* if A is at the origin. The *zero vector*, denoted $\mathbf{0}$, is the vector of all zeros.

Example 1.1.1. Consider the points $A(-2, 4)$ and $B(1, 2)$ in the plane. Draw the vector \mathbf{AB} emanating from A . Draw the vector \mathbf{AB} in standard position.



Remark. Unless otherwise specified, given an arbitrary vector \mathbf{v} , we will always draw it in standard position.

This leads us to the following fact:

Fact. Two vectors are *equal* if and only if their components are the same.

We will use the following notation repeatedly throughout the semester:

- \mathbb{R} - This is the set of real numbers.
- \mathbb{R}^n - This is the set of vectors with n real components.

There are different ways to represent the vector \mathbf{v} , either as a *row vector*,

$$\mathbf{v} = [v_1, \dots, v_n],$$

or as a *column vector*,

$$\mathbf{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}.$$

As will become apparent in future sections, we'll eventually come to prefer column vectors, but for now we won't make a distinction.

Remark. Since writing a column vector takes up a lot of vertical space on the page, we will sometimes write $\mathbf{v} = [v_1, \dots, v_n]^T$ to denote the column form of the vector \mathbf{v} (and we will also formally describe the superscript T notation in a later section).

1.1.1 Vector operations

Because the real numbers have proven to be so nice to us over the years, we would like to try to come up with an analogous algebraic structure that works for higher dimensions. The structure we're building is formally known as a *(real) vector space*, but we'll avoid that terminology and generality in favor of trying to motivate concrete understanding.

Definition. There are two operations on \mathbb{R}^n :

- *Vector addition* – For vectors $\mathbf{u} = [u_1, \dots, u_n]^T$ and $\mathbf{v} = [v_1, \dots, v_n]^T$ in \mathbb{R}^n

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

- *Scalar multiplication.* For a vector $\mathbf{u} = [u_1, \dots, u_n]^T$ in \mathbb{R}^n and a *scalar* a in \mathbb{R} ,

$$a\mathbf{u} = \begin{bmatrix} au_1 \\ \vdots \\ au_n \end{bmatrix}$$

It's worthwhile to note that scalar multiplication is slightly weaker than what we would want – ideally we could multiply two vectors together to get a vector, but instead we can only multiply a vector and a scalar (at least the output is still a vector). The reasons for this are somewhat subtle and deep, but I hope in a later section to provide some motivation for why trying to define other types of multiplication is ultimately more problematic.

Theorem 1.1.2 (*Poole Theorem 1.1 - Algebraic Properties of Vectors in \mathbb{R}^n*). *Let $\mathbf{u}, \mathbf{v}, \mathbf{w}$ be vectors in \mathbb{R}^n and let a, b be scalars in \mathbb{R} . With vector addition and scalar multiplication on \mathbb{R}^n as defined above, the following are true:*

1. $\mathbf{u} + \mathbf{v}$ is in \mathbb{R}^n .
2. $\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$.
3. $(\mathbf{u} + \mathbf{v}) + \mathbf{w} = \mathbf{u} + (\mathbf{v} + \mathbf{w})$.
4. The zero vector $\mathbf{0}$ satisfies $\mathbf{u} + \mathbf{0} = \mathbf{u}$ for all vectors \mathbf{u} .
5. For each \mathbf{u} in \mathbb{R}^n , there is some vector $-\mathbf{u}$ for which $\mathbf{u} + (-\mathbf{u}) = \mathbf{0}$.
6. $a\mathbf{u}$ is in \mathbb{R}^n .
7. $a(\mathbf{u} + \mathbf{v}) = a\mathbf{u} + a\mathbf{v}$.
8. $(a + b)\mathbf{u} = a\mathbf{u} + b\mathbf{u}$.
9. $(ab)\mathbf{u} = a(b\mathbf{u})$.
10. $1\mathbf{u} = \mathbf{u}$.

Proof. All parts of the above theorem are straightforward to check, so they are left as an exercise to the reader. \square

The above theorem effectively says that \mathbb{R}^n , with our two operations, satisfies all of the same algebraic properties as \mathbb{R} .

1.1.2 Geometry of vector operations

Given two vectors \mathbf{u} and \mathbf{v} , their sum is the vector formed from adjoining the tail of \mathbf{v} to the head of \mathbf{u} . In \mathbb{R}^2 , this is equivalent to the diagonal to the parallelogram formed from sides \mathbf{u} and \mathbf{v} . Scalar multiplication merely *scales* the vector and leaves its direction unchanged. To see this, consider the following example:



Example 1.1.3. Let $\mathbf{u} = [-1, -2]^T$, $\mathbf{v} = [4, 1]^T$, and $a = 2$. Then

$$\mathbf{u} + \mathbf{v} = \begin{bmatrix} -1 + 4 \\ -2 + 1 \end{bmatrix} = \begin{bmatrix} 3 \\ -1 \end{bmatrix}$$

and

$$a\mathbf{u} = \begin{bmatrix} 2(-1) \\ 2(-2) \end{bmatrix} = \begin{bmatrix} -2 \\ -4 \end{bmatrix}.$$

Drawing them out



1.1.3 Linear combinations

Definition. A vector \mathbf{u} is a *linear combination* of the vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ if there are scalars a_1, \dots, a_n so that

$$\mathbf{u} = a_1\mathbf{v}_1 + \dots + a_n\mathbf{v}_n.$$

You can think of a linear combination as some sort of recipe - the \mathbf{v}_i 's are the ingredients, the a_i 's are the quantities of those ingredients, and \mathbf{u} is the finished product.

Definition. In \mathbb{R}^n , there are n vectors

$$\mathbf{e}_1 = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix}, \quad \dots \quad \mathbf{e}_n = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}.$$

which we call the *standard basis vectors for \mathbb{R}^n* .

For now, ignore the word *basis* above; we will give technical meaning to that later. The reason these are standard is because, when looking to decompose a vector \mathbf{u} into a linear combination of vectors, then simply picking apart the components is probably the most natural thing to try first.


Example 1.1.4. The vector $\mathbf{u} = \begin{bmatrix} 5 \\ 6 \\ 7 \end{bmatrix}$ is a linear combination of the standard basis vectors in the following way:

$$\mathbf{u} = \begin{bmatrix} 5 \\ 6 \\ 7 \end{bmatrix} = \begin{bmatrix} 5 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 6 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 7 \end{bmatrix} = 5 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + 6 \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} + 7 \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = 5\mathbf{e}_1 + 6\mathbf{e}_2 + 7\mathbf{e}_3$$

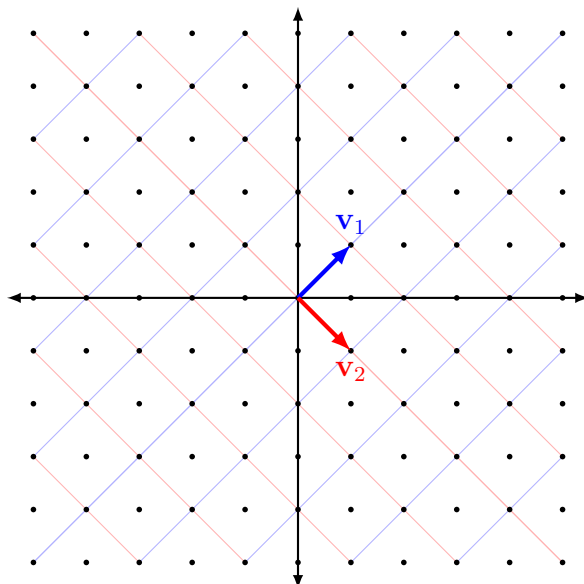
Definition. A *coordinate grid* is the grid formed by translates of a set of vectors.

I don't think the above definition is a very good way to explain the idea, which is actually fairly intuitive. We want to form a grid where (1) the lines parallel to the given vectors and (2) the intersections of these grids correspond to integer-linear combinations of vectors. It may be more enlightening to just see a few examples.

Example 1.1.5. The coordinate grid for \mathbb{R}^2 formed from the standard basis vectors \mathbf{e}_1 and \mathbf{e}_2 is the usual Cartesian grid.





Example 1.1.6. The coordinate grid for \mathbb{R}^2 formed from the vectors $\mathbf{v}_1 = [1, 1]^T$ and $\mathbf{v}_2 = [1, -1]^T$ is below.



Combined with the geometric intuition about vector addition and scalar multiplication, these coordinate grids provide us with a way to visually identify the linear combination.

Example 1.1.7. Write the vector $\mathbf{u} = [2, 4]^T$ as a linear combination of the vectors...

1. $\mathbf{e}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, $\mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$
2. $\mathbf{v}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$, $\mathbf{v}_2 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$



We can verify this algebraically:

$$\begin{aligned}
 \mathbf{u} &= 2\mathbf{e}_1 + 4\mathbf{e}_2 \\
 &= 2[1, 0] + 4[0, 1] \\
 &= [2, 0] + [0, 4] \\
 &= [2, 4]
 \end{aligned}$$

We can verify this algebraically:

$$\begin{aligned}
 \mathbf{u} &= 3\mathbf{v}_1 + (-1)\mathbf{v}_2 \\
 &= 3[1, 1] + (-1)[1, -1] \\
 &= [3, 3] + [-1, 1] \\
 &= [2, 4]
 \end{aligned}$$

1.2 Length and Angle: The Dot Product

Definition. For vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$, the *dot product* of \mathbf{u} and \mathbf{v} , denoted $\mathbf{u} \cdot \mathbf{v}$ is

$$\mathbf{u} \cdot \mathbf{v} = u_1v_1 + \cdots + u_nv_n.$$

Remark. Note that the dot product of two vectors *is a scalar*.

The dot product has the following nice properties.

Theorem 1.2.1 (Poole Theorem 1.2). *Let $\mathbf{u}, \mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ and let k be some scalar. Then*

1. $\mathbf{u} \cdot \mathbf{v} = \mathbf{v} \cdot \mathbf{u}$
2. $\mathbf{u} \cdot (\mathbf{v} + \mathbf{w}) = (\mathbf{u} \cdot \mathbf{v}) + (\mathbf{u} \cdot \mathbf{w})$
3. $(\mathbf{v} + \mathbf{w}) \cdot \mathbf{u} = (\mathbf{v} \cdot \mathbf{u}) + (\mathbf{w} \cdot \mathbf{u})$
4. $(k\mathbf{u}) \cdot \mathbf{v} = \mathbf{u} \cdot (k\mathbf{v}) = k(\mathbf{u} \cdot \mathbf{v})$
5. *For every \mathbf{u} we have that $\mathbf{u} \cdot \mathbf{u} \geq 0$, with equality if and only if $\mathbf{u} = \mathbf{0}$.*


Proof. The proof is entirely straightforward and left as an exercise to the reader. □

1.2.1 Length

Notice that for a vector $\mathbf{v} = [x, y] \in \mathbb{R}^2$,

$$\mathbf{v} \cdot \mathbf{v} = x^2 + y^2,$$

which, from the Pythagorean theorem, is precisely the square of the length of \mathbf{v} .



Definition. The *length* (or *norm*) of a vector $\mathbf{v} \in \mathbb{R}^n$ is the scalar defined by

$$\|\mathbf{v}\| = \sqrt{\mathbf{v} \cdot \mathbf{v}} = \sqrt{v_1^2 + v_2^2 + \cdots + v_n^2}$$

The following are immediate consequences of the properties of the dot product in Theorem 1.2.1

Theorem 1.2.2 (Poole Theorem 1.3). *For $\mathbf{v} \in \mathbb{R}^n$ and a scalar k ,*


1. $\|\mathbf{v}\| = 0$ if and only if $\mathbf{v} = \mathbf{0}$.

$$2. \|k\mathbf{v}\| = |k|\|\mathbf{v}\|.$$

The following follows from the classical geometry result of the same name.

Theorem 1.2.3 (Triangle Inequality). For $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$,

$$\|\mathbf{u} + \mathbf{v}\| \leq \|\mathbf{u}\| + \|\mathbf{v}\|.$$




Definition. A vector \mathbf{v} is called a *unit vector* if $\|\mathbf{v}\| = 1$.

Remark. Every unit vector in \mathbb{R}^2 corresponds to a point on the unit circle. Every unit vector in \mathbb{R}^3 corresponds to a point on the unit sphere. Generally, every unit vector in \mathbb{R}^n corresponds to a point on the unit $(n - 1)$ -sphere.

Let \mathbf{v} be any nonzero vector and let $\ell = \|\mathbf{v}\|$ be its length. Then the vector $\frac{\mathbf{v}}{\ell}$ is a unit vector because

$$\left\| \frac{\mathbf{v}}{\ell} \right\| = \frac{\|\mathbf{v}\|}{\ell} = \frac{\ell}{\ell} = 1$$

Definition. The process above is called *normalization*, and it always produces a vector in the same direction as \mathbf{v} but with unit length.



Remark. If $\|\mathbf{v}\| > 1$, then normalization corresponds to shrinking \mathbf{v} (pictured above), but if $\|\mathbf{v}\| < 1$, then normalization stretches \mathbf{v} .

Remark. Despite the similarities in name, “normalization” is unrelated to the concept of a “normal vector.” What you’ll find is that “normal” is probably the most over-used word in mathematics. Because there aren’t any around me as I type this, I’m going to go ahead and blame the physicists for the abuse of language.

1.2.2 Distances

Recall that, for two points $P(x_1, y_1)$ and $Q(x_2, y_2)$ in the plane, we have that the distance between them is given by


$$d(P, Q) = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$$

If we identify the point $P(x_1, y_1)$ with the vector $\mathbf{u} = [x_1, y_1]$ and the point $Q(x_2, y_2)$ with the vector $\mathbf{v} = [x_2, y_2]$, then the right-hand side of the equation is just $\|\mathbf{u} - \mathbf{v}\|$. As such, we can define distances between vectors using the obvious analog.

Definition. Given two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$, the *distance* between \mathbf{u} and \mathbf{v} is


$$d(\mathbf{u}, \mathbf{v}) = \|\mathbf{u} - \mathbf{v}\|.$$

Remark. Visualizing vectors as arrows emanating from the origin, distance, as above, is actually measuring the distance between the heads of the arrows.



1.2.3 Angles

Consider a triangle $\triangle ABC$ and the angle $\theta = \angle ABC$ (pictured below)



Recall that the law of cosines says

$$b^2 = a^2 + c^2 - 2ac \cos(\theta)$$

Replacing the triangle $\triangle ABC$ with the triangle formed from vectors \mathbf{u} , \mathbf{v} , $\mathbf{u} - \mathbf{v}$ (as in the picture above on the right), we have

$$\|\mathbf{u} - \mathbf{v}\|^2 = \|\mathbf{v}\|^2 + \|\mathbf{u}\|^2 - 2\|\mathbf{u}\|\|\mathbf{v}\| \cos \theta$$

Expanding out the left-hand side of the above equation in terms of dot products, we get

$$\|\mathbf{u}\|^2 + \|\mathbf{v}\|^2 - 2\mathbf{u} \cdot \mathbf{v} = \|\mathbf{v}\|^2 + \|\mathbf{u}\|^2 - 2\|\mathbf{u}\|\|\mathbf{v}\| \cos \theta$$

Canceling appropriately and rearranging the equation yields

Definition. For nonzero vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$, the *angle* θ between \mathbf{u} and \mathbf{v} satisfies

$$\cos \theta = \frac{\mathbf{u} \cdot \mathbf{v}}{\|\mathbf{u}\|\|\mathbf{v}\|}$$

Example 1.2.4. Compute the angle between the vectors $\mathbf{u} = [0, 3, 3]^T$ and $\mathbf{v} = [-1, 2, 1]^T$.

From the above, we get that

$$\cos \theta = \frac{\mathbf{u} \cdot \mathbf{v}}{\|\mathbf{u}\|\|\mathbf{v}\|} = \frac{9}{(3\sqrt{2})(\sqrt{6})} = \frac{\sqrt{3}}{2}$$

and thus

$$\theta = \arccos\left(\frac{\sqrt{3}}{2}\right) = \frac{\pi}{6}.$$


The following follows immediately from the definition.

Corollary 1.2.5. $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$ are perpendicular if and only if $\mathbf{u} \cdot \mathbf{v} = 0$.

[NOTE: Move coverage of this next part to Section 5.1]

1.2.4 Projections

Up until now we haven't really given any geometric description of what the dot product actually measures, so let's consider $\mathbf{e}_1 = [1, 0]^T$ and $\mathbf{v} = [x, y]^T$. A straightforward computation shows that $\mathbf{e}_1 \cdot \mathbf{v} = x$. Looking at the picture below



We see that x is precisely the length of the “shadow” cast by \mathbf{v} on the line spanned by \mathbf{e}_1 . More generally, for two vectors \mathbf{u} and \mathbf{v} , the dot product $\mathbf{u} \cdot \mathbf{v}$ is the length of “shadow” of \mathbf{v} cast onto \mathbf{u} , but then scaled by the length of \mathbf{v} (by symmetry of the dot product, this description also works *mutatis mutandis* for the “shadow” cast by \mathbf{u} onto \mathbf{v}).

Let's formalize this notion of a “shadow” cast by \mathbf{v} onto \mathbf{u} keeping in mind two things:


1. \mathbf{u} should be normalized to be a unit vector so as to preserve the length of \mathbf{v} 's shadow, and
2. the shadow of a vector should also be a vector in the direction of \mathbf{u} (i.e. a scalar multiple of \mathbf{u}).

Definition. Given two vectors \mathbf{u} and \mathbf{v} , the *projection* (or *orthogonal projection*) of \mathbf{v} onto \mathbf{u} is the vector $\text{proj}_{\mathbf{u}}(\mathbf{v})$ given by

$$\text{proj}_{\mathbf{u}}(\mathbf{v}) = \left(\frac{\mathbf{u} \cdot \mathbf{v}}{\|\mathbf{u}\|^2} \right) \mathbf{u} = \left(\frac{\mathbf{u} \cdot \mathbf{v}}{\mathbf{u} \cdot \mathbf{u}} \right) \mathbf{u}.$$

Example 1.2.6. Let $\mathbf{u} = [4, -1]^T$ and $\mathbf{v} = [3, 5]^T$. Using the formula from the definition we have

$$\text{proj}_{\mathbf{u}}(\mathbf{v}) = \left(\frac{\mathbf{u} \cdot \mathbf{v}}{\mathbf{u} \cdot \mathbf{u}} \right) \mathbf{u} = \left(\frac{7}{17} \right) \mathbf{u} = \left[\frac{28}{17}, \frac{-7}{17} \right]^T.$$



Notice that the vector $\mathbf{v} - \text{proj}_{\mathbf{u}}(\mathbf{v})$ (represented by a dashed line) is perpendicular to \mathbf{u} . Perpendicular vectors are called *orthogonal* (hence the name “orthogonal projection”). What’s more

$$\begin{aligned}
 (\mathbf{v} - \text{proj}_{\mathbf{u}}(\mathbf{v})) \cdot \mathbf{u} &= \mathbf{v} \cdot \mathbf{u} - \left(\frac{\mathbf{u} \cdot \mathbf{v}}{\mathbf{u} \cdot \mathbf{u}} \right) \mathbf{u} \cdot \mathbf{u} \\
 &= \mathbf{v} \cdot \mathbf{u} - \mathbf{u} \cdot \mathbf{v} \\
 &= \mathbf{u} \cdot \mathbf{v} - \mathbf{u} \cdot \mathbf{v} \\
 &= 0.
 \end{aligned}$$

So we deduce from this the following fact

Theorem 1.2.7. *Two vectors \mathbf{u}, \mathbf{v} in \mathbb{R}^n are orthogonal if and only if $\mathbf{u} \cdot \mathbf{v} = 0$.*

2.1 Introduction to Linear Systems

Definition. A *linear equation* in the variables x_1, \dots, x_n is an equation that can be written in the form

$$a_1x_1 + \dots + a_nx_n = b$$

where a_1, \dots, a_n are real numbers called *coefficients* and b is a real number called the *constant term*. A *solution* of the equation is a vector $[v_1, \dots, v_n]^T$ so that

$$a_1v_1 + \dots + a_nv_n = b.$$

Example 2.1.1. $4x - y = 2$ is an example of a linear equation. And notice we can rearrange it as $y = 4x - 2$, which is the equation of a line (hence why we call these “linear”). The vector $[1, 2]^T$ is a solution because

$$4(1) - (2) = 2.$$

In fact, for any real number t , the vector $[t, 4t - 2]^T$ is a solution because

$$4(t) - (4t - 2) = 2.$$

This means there are infinitely many possible solutions.

Definition. The collection of all solutions to a linear equation is called the

Definition. solution set of that equation.

Noticing that

$$\begin{bmatrix} t \\ 4t - 2 \end{bmatrix} = \begin{bmatrix} 0 \\ -2 \end{bmatrix} + \begin{bmatrix} t \\ 4t \end{bmatrix} = \begin{bmatrix} 0 \\ -2 \end{bmatrix} + t \begin{bmatrix} 1 \\ 4 \end{bmatrix}$$

we can write the solution set to the previous example as

$$\left\{ \begin{bmatrix} 0 \\ -2 \end{bmatrix} + t \begin{bmatrix} 1 \\ 4 \end{bmatrix} \quad \text{where } t \in \mathbb{R} \right\}$$

Definition. The *parametric form* of the solution set is when it is written as

$$\{ \mathbf{v}_0 + t_1\mathbf{v}_1 + \dots + t_n\mathbf{v}_n \quad \text{where } t_i \in \mathbb{R} \}$$

for some vectors \mathbf{v}_i .

Example 2.1.2. $\sin\left(\frac{\pi}{82364423}\right)x + \sqrt{540.6464}y + z = e^{71}$ is a linear equation, because complicated as they are, $\sin\left(\frac{\pi}{82364423}\right)$, $\sqrt{540.6464}$, e^{71} are just real numbers.

Example 2.1.3. $x + xy + y = 7$ is *not* a linear equation because of that xy term.

Example 2.1.4. $x^2 + 3^y = 8$ is *not* a linear equation because of both the x^2 and 3^y terms.

Definition. A *system of linear equations* is a finite set of linear equations, each with the same variables (and probably different coefficients). A *solution* of a system of linear equations is a vector that is simultaneously a solution for each linear equation in the system. A *solution set* is the collection of all possible solutions to the system.

Example 2.1.5. The system

$$\begin{cases} 2x - y = 3 \\ x + 3y = 5 \end{cases}$$

has the vector $[2, 1]^T$ as a solution; in fact, this is the only solution.

Definition. A system of linear equations is called *consistent* if it has at least one solution, and *inconsistent* if it has no solutions.

Fact. A system of linear equations with real coefficients has either

- (a) a unique solution (consistent)
- (b) infinitely many solutions (consistent)
- (c) no solutions (inconsistent)

You can convince yourself of the above trichotomy by considering how it works for systems of equations with two variables (whose solution sets are graphically lines in the Cartesian plane). Two lines can either intersect in a single point (if they are transverse), intersect in infinitely many points (if they coincide), or no points (if they are parallel).

Example 2.1.6. The system in Example 2.1.5 is consistent and the solution is unique.

Example 2.1.7. The system

$$\begin{cases} x - y = 0 \\ 2x - 2y = 0 \end{cases}$$

is consistent. It has the solution $[x, y]^T = [1, 1]^T$, but this is not the only solution. For any real number t , the vector $[t, t]^T$ is a solution, so there are infinitely many.

Example 2.1.8. The system

$$\begin{cases} x + y = 0 \\ x + y = 2 \end{cases}$$

has no solutions.

Definition. Two systems of linear equations are called *equivalent* if they have the same solution set.

Notice how easy the next system of equations is to solve by *back-substitution*.

Example 2.1.9. Consider the system

$$\begin{cases} x + 3y + 5z = 7 \\ \quad 2y - 4z = 6 \\ \quad \quad 8z = 16 \end{cases}$$

Because of this kind of “triangular structure,” we quickly deduce $z = 2$, and then $2y - 4(2) = 6$ implies that $y = 7$, and then $x + 3(7) + 5(2) = 7$ implies that $x = -24$.

2.2 Direct Methods for Solving Linear Systems

2.2.1 Row Operations

Example 2.2.1. In Example 2.1.5 we saw that the system

$$\begin{cases} 2x - y = 3 \\ x + 3y = 5 \end{cases}$$

was consistent and had the unique solution $[x, y]^T = [2, 1]^T$. Consider now the following three systems of equations. The first system is obtained by merely swapping the equations. The second system is obtained by scaling the second equation. The third system is obtained by replacing the second equation with the sum of the first and second equations.

$$\begin{cases} x + 3y = 5 \\ 2x - y = 3 \end{cases} \quad \begin{cases} 2x - y = 3 \\ 100x + 300y = 500 \end{cases} \quad \begin{cases} 2x - y = 3 \\ 3x + 2y = 8 \end{cases}$$

All of these have the (unique) solution $[x, y]^T = [2, 1]^T$ (this is left as an exercise for the reader), and so they are all equivalent.

It turns out that this fact isn't specific to this system, but is generally true of any linear system: these three operations do not change the solution set of the system! The "elimination method" (which you may be familiar with from a previous algebra/precalculus class) uses this fact to solve systems of linear equations. If we think about what this is doing to the corresponding augmented matrices, we get what we call the *elementary row operations*.

Definition. The *elementary row operations* of a given matrix are the following operations:

1. Swapping Row i and Row j (denoted $R_i \leftrightarrow R_j$).
2. Multiplying Row i by a nonzero constant (denoted $kR_i \mapsto R_i$).
3. Adding (a multiple of) Row j to Row i (denoted $R_i + kR_j \mapsto R_i$).

Remark. These operations are not specific to augmented matrices, but are true of any matrices. In fact, unless explicitly stated otherwise, you should probably not ever assume that a matrix is augmented.

Given two (augmented) matrices, the above operations do not change the solution set for the corresponding linear system. So since two linear systems are equivalent if they have the same solution set, the following is a natural definition

Definition. Two matrices A and B are *row equivalent* if there is a sequence of elementary row operations transforming A into B .

Example 2.2.2. Using the systems in Example 2.2.1, we will show that the corresponding augmented matrices are row equivalent:

$$\begin{array}{ccc} & \begin{bmatrix} 2 & -1 & | & 3 \\ 1 & 3 & | & 5 \end{bmatrix} & \\ & \swarrow^{R_1 \leftrightarrow R_2} & \downarrow^{100R_2 \mapsto R_2} \\ \begin{bmatrix} 1 & 3 & | & 5 \\ 2 & -1 & | & 3 \end{bmatrix} & \begin{bmatrix} 2 & -1 & | & 3 \\ 100 & 300 & | & 500 \end{bmatrix} & \searrow^{R_1 + R_2 \mapsto R_2} \\ & & \begin{bmatrix} 2 & -1 & | & 3 \\ 3 & 2 & | & 8 \end{bmatrix} \end{array}$$

2.2.2 (Reduced) Row Echelon Form

The following systems are equivalent (it's again an exercise to the reader to verify this):

$$\begin{cases} x - y - z = 2 \\ 3x - 3y + 2z = 16 \\ 2x - y + z = 9 \end{cases} \quad \begin{cases} x - y - z = 2 \\ y + 3z = 5 \\ 5z = 10 \end{cases} \quad \begin{cases} x = 3 \\ y = -1 \\ z = 2 \end{cases}$$

and thus they correspond to the following row equivalent augmented matrices

$$\left[\begin{array}{ccc|c} 1 & -1 & -1 & 2 \\ 3 & -3 & 2 & 16 \\ 2 & -1 & 1 & 9 \end{array} \right] \quad \left[\begin{array}{ccc|c} 1 & -1 & -1 & 2 \\ 0 & 1 & 3 & 5 \\ 0 & 0 & 5 & 10 \end{array} \right] \quad \left[\begin{array}{ccc|c} 1 & 0 & 0 & 3 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & 2 \end{array} \right]$$

The second and third systems are much more useful for actually *solving* the system because they have the nice triangular structure that allows us to back-substitute (or in the case of the third one, simply reading off the solution). Let's give names to this triangular structure that we like so much.

Definition. A matrix is in *row echelon form* (REF) if it satisfies the following properties:

- Any rows consisting entirely of zeros are at the bottom
- In each nonzero row, the first nonzero entry (the *leading entry*) is in a column to the left of any leading entries below it. The column containing the leading entry is sometimes called the *pivot column*.

Example 2.2.3. The following matrices are in row echelon form.

$$\begin{bmatrix} 1 & 2 & 3 \\ 0 & 4 & 5 \\ 0 & 0 & 0 \end{bmatrix} \quad \begin{bmatrix} 1 & 2 & 3 \\ 0 & 4 & 5 \\ 0 & 0 & 6 \end{bmatrix} \quad \begin{bmatrix} 0 & 1 & 2 & 3 & 4 & 5 \\ 0 & 0 & 6 & 7 & 8 & 9 \\ 0 & 0 & 0 & 0 & 10 & 11 \\ 0 & 0 & 0 & 0 & 0 & 5 \end{bmatrix}$$

Example 2.2.4. The following matrices are not in row echelon form. (Why?)

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

Definition. The *reduced row echelon form* (RREF) of a matrix is essentially the same as the row echelon form with the following additional requirements:

- Each leading entry is 1.
- Any entries above a *leading 1* are also 0.

Example 2.2.5. The following matrices are in reduced row echelon form.

$$\begin{bmatrix} 1 & 0 & 3 \\ 0 & 1 & 5 \\ 0 & 0 & 0 \end{bmatrix} \quad \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \begin{bmatrix} 0 & 1 & 0 & 3 & 0 & 0 \\ 0 & 0 & 1 & 7 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Example 2.2.6. The following matrices are not in reduced row echelon form. (Why?)

need examples here

Theorem 2.2.7. *Every matrix is equivalent to a matrix in (reduced) row echelon form.*

The proof of this is actually procedural, so let's see it done in the context of an example.

Example 2.2.8.

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

1. Working left to right, find the first nonzero column in the matrix.

The first column is nonzero

2. Among all of the rows with nonzero entries in this column, choose one and move it to Row 1.

We'll just keep the first row where it is

3. Use elementary row operations to clear all other nonzero entries in this column (below Row 1).

$$\left[\begin{array}{ccc|c} 1 & -1 & -1 & 2 \\ 3 & -3 & 2 & 16 \\ 2 & -1 & 1 & 9 \end{array} \right] \xrightarrow{R_2 - 3R_1 \mapsto R_2} \left[\begin{array}{ccc|c} 1 & -1 & -1 & 2 \\ 0 & 0 & 5 & 10 \\ 2 & -1 & 1 & 9 \end{array} \right] \quad (2.2.1)$$

$$\xrightarrow{R_3 - 2R_1 \mapsto R_3} \left[\begin{array}{ccc|c} 1 & -1 & -1 & 2 \\ 0 & 0 & 5 & 10 \\ 0 & 1 & 3 & 5 \end{array} \right] \quad (2.2.2)$$

$$(2.2.3)$$

4. Ignoring Row 1, find the next nonzero column in this matrix.

Ignoring Row 1, the second column is now the next nonzero column.

5. Among all of the rows below Row 1 with nonzero entries in this column, choose one and move it to Row 2.

$$\xrightarrow{R_2 \leftrightarrow R_3} \left[\begin{array}{ccc|c} 1 & -1 & -1 & 2 \\ 0 & 1 & 3 & 5 \\ 0 & 0 & 5 & 10 \end{array} \right] \quad (2.2.4)$$

6. Use elementary row operations to clear all other nonzero entries in this column (below Row 2).

Already done.

7. Repeat this process until the matrix is in row echelon form.

Huzzah, the matrix in Equation 2.2.5 is in row echelon form!

8. Now scale every row so that the leading term is a 1. The result will be in reduced row echelon form.

$$\xrightarrow{\frac{1}{5}R_3 \mapsto R_3} \left[\begin{array}{ccc|c} 1 & -1 & -1 & 2 \\ 0 & 1 & 3 & 5 \\ 0 & 0 & 1 & 2 \end{array} \right] \quad (2.2.5)$$

9. Working from left to right, use elementary row operations to clear all nonzero entries above each leading 1.

$$\xrightarrow{R_1+R_2\rightarrow R_1} \left[\begin{array}{ccc|c} 1 & 0 & 2 & 7 \\ 0 & 1 & 3 & 5 \\ 0 & 0 & 1 & 2 \end{array} \right] \quad (2.2.6)$$

$$\xrightarrow{R_1-2R_3\rightarrow R_1} \left[\begin{array}{ccc|c} 1 & 0 & 0 & 3 \\ 0 & 1 & 3 & 5 \\ 0 & 0 & 1 & 2 \end{array} \right] \quad (2.2.7)$$

$$\xrightarrow{R_2-3R_3\rightarrow R_2} \left[\begin{array}{ccc|c} 1 & 0 & 0 & 3 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & 2 \end{array} \right] \quad (2.2.8)$$

Remark. The row echelon form of a given matrix is not unique.

Remark. The reduced row echelon form of a matrix is unique.

Definition. The process described in the example above is called *row reduction*.

Theorem 2.2.9 (Poole Theorem 2.1). *Matrices A and B are row equivalent if and only if they can be row reduced to the same echelon form.*

2.2.3 Gaussian Elimination and Gauss–Jordan Elimination

Definition. Given a linear system with augmented matrix $[A|\mathbf{b}]$ in (reduced) row echelon form, the pivot columns correspond to *leading variables* in the system, and the other nonzero columns correspond to *free variables* in the system.

Definition. *Gaussian elimination* is the following process:

1. Write a linear system as an augmented matrix.
2. Put the matrix into row echelon form.
3. Reinterpret as a linear system and use back-substitution to solve the system for the leading variables.

Definition. *Gauss–Jordan Elimination* is the following process:

1. Write a linear system as an augmented matrix.
2. Put the matrix into reduced row echelon form.
3. Reinterpret as a linear system and solve the system.

Both processes take about the same amount of time by hand. But since the reduced row echelon form is unique and most matrix algebra software has an RREF feature, Gauss–Jordan is more efficient in practice.

Example 2.2.10. Use Gauss–Jordan elimination to find the solution set for the given system

$$\begin{cases} x_1 - x_2 + x_3 + 4x_4 = 0 \\ 2x_1 + x_2 - x_3 + 2x_4 = 9 \\ 3x_1 - 3x_2 + 3x_3 + 12x_4 = 0 \end{cases}$$

We set up the augmented matrix and row-reduce

$$\begin{array}{c}
 \left[\begin{array}{cccc|c} 1 & -1 & 1 & 4 & 0 \\ 2 & 1 & -1 & 2 & 9 \\ 3 & -3 & 3 & 12 & 0 \end{array} \right] \xrightarrow{R_2 - 2R_1 \mapsto R_2} \left[\begin{array}{cccc|c} 1 & -1 & 1 & 4 & 4 \\ 0 & 3 & -3 & -6 & -6 \\ 3 & -3 & 3 & 12 & 0 \end{array} \right] \\
 \\
 \xrightarrow{R_3 - 3R_1 \mapsto R_3} \left[\begin{array}{cccc|c} 1 & -1 & 1 & 4 & 4 \\ 0 & 3 & -3 & -6 & -6 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right] \\
 \\
 \xrightarrow{\frac{1}{3}R_2 \mapsto R_2} \left[\begin{array}{cccc|c} 1 & -1 & 1 & 4 & 4 \\ 0 & 1 & -1 & -2 & -2 \\ 3 & 0 & 0 & 0 & 0 \end{array} \right] \\
 \\
 \xrightarrow{R_1 + R_2 \mapsto R_1} \left[\begin{array}{cccc|c} 1 & 0 & 0 & 2 & 2 \\ 3 & 0 & 0 & 0 & 0 \\ 0 & 1 & -1 & -2 & -2 \\ 3 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right]
 \end{array}$$

The corresponding system is

$$\begin{cases} x_1 & & & + 2x_4 & = & 3 \\ & x_2 & - & x_3 & - & x_4 & = & 3 \end{cases}$$

Solving for the leading variables, we get

$$\begin{cases} x_1 = 3 - 2x_4 \\ x_2 = 3 + x_3 + 2x_4 \end{cases}$$

and hence any solution is of the form

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

Replacing our free variables x_3 and x_4 with parameters s and t (respectively), our solution set is

$$\left\{ \begin{bmatrix} 3 \\ 3 \\ 0 \\ 0 \end{bmatrix} + s \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix} + t \begin{bmatrix} -2 \\ 2 \\ 0 \\ 1 \end{bmatrix} \text{ where } s, t \in \mathbb{R} \right\}$$

What we have seen is that both row echelon form and reduced row echelon form are useful in the same way, but both have pros and cons. Row echelon form isn't unique and, in the case of augmented matrices, it takes a little bit more work to solve the system at the end. Reduced row echelon form is unique and makes the solution at the end easier, but requires more steps initially.

2.2.4 Rank and Number of Solutions

Example 2.2.11. In Example 2.1.8, we stated that the system

$$\begin{cases} x + y = 0 \\ x + y = 2 \end{cases}$$

was inconsistent. Look at what happens when we set up the augmented matrix and row-reduce:

$$\left[\begin{array}{cc|c} 1 & 1 & 0 \\ 1 & 1 & 2 \end{array} \right] \xrightarrow{R_2 - R_1 \rightarrow R_2} \left[\begin{array}{cc|c} 1 & 1 & 0 \\ 0 & 0 & 2 \end{array} \right]$$

That last row corresponds to the linear equation $0 = 2$, which is patently false. This means there can't possibly be a solution to the system, i.e., it is inconsistent. We state this observation as a proposition.

Proposition 2.2.12. *Let $A = [a_{ij}]$ be the coefficient matrix for a system and $\mathbf{b} = [b_1, \dots, b_m]^T$ the vector of constant terms. If the row echelon form of the augmented matrix $[A|\mathbf{b}]$ contains a row where $a_{i1} = a_{i2} = \dots = a_{im} = 0$ and $b_m \neq 0$, then the system is inconsistent.*

One might ask if we can say anything about a consistent system from its (reduced) row echelon form. To answer this, we first introduce the following definition.

Definition. The *rank* of a matrix A is the number of nonzero rows in its (reduced) row echelon form, and is denoted $\text{Rank}(A)$.

Example 2.2.13. The rank of the coefficient matrix in Example 2.1.8 is 1, and the rank of the coefficient matrix in Example ?? is 2.

Theorem 2.2.14 (Poole Theorem 2.2 - The Rank Theorem). *If A is the coefficient matrix of a consistent system of linear equations with n variables, then*

$$n = \text{Rank}(A) + \text{number of free variables.}$$

Remark. It turns out this theorem is actually just a special interpretation of a much more powerful theorem called the "Rank-Nullity Theorem," but that discussion will have to wait for a later section.

Definition. A system of linear equations, $A\mathbf{x} = \mathbf{b}$ is *homogeneous* if $\mathbf{b} = \mathbf{0}$.

Remark. Homogeneous systems are nice because they ALWAYS have at least one solution (namely the *trivial solution* $\mathbf{x} = \mathbf{0}$).

Theorem 2.2.15. *If $A\mathbf{x} = \mathbf{0}$ is a homogeneous system of m linear equations and n variables, where $m < n$, then the system has infinitely many solutions.*

Proof. Since the system is homogeneous, it has at least one solution. Since $\text{Rank}(A) \leq m$, then by the Rank Theorem

$$\text{number of free variables} = n - \text{Rank}(A) \geq n - m > 0$$

and a nonzero number of free variables implies that there are infinitely-many solutions. \square

Example 2.2.16. Use Gauss-Jordan elimination to find the solution set for the given system

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

Creating the augmented matrix and doing the corresponding row operations, we have

$$\begin{aligned} \left[\begin{array}{cccc|c} 1 & -1 & 3 & 4 & 0 \\ 1 & 1 & -1 & -2 & 0 \end{array} \right] &\xrightarrow{R_2 - R_1 \mapsto R_2} \left[\begin{array}{cccc|c} 1 & -1 & 3 & 4 & 0 \\ 0 & 2 & -4 & -6 & 0 \end{array} \right] \\ &\xrightarrow{\frac{1}{2}R_2 \mapsto R_2} \left[\begin{array}{cccc|c} 1 & -1 & 3 & 4 & 0 \\ 0 & 1 & -2 & -3 & 0 \end{array} \right] \\ &\xrightarrow{R_1 + R_2 \mapsto R_1} \left[\begin{array}{cccc|c} 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & -2 & -3 & 0 \end{array} \right] \end{aligned}$$

From here, we can see that x_3 and x_4 are free variables, so letting $x_3 = s$ and $x_4 = t$, we get that the solution is

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

The way this last example differs from Example 2.2.10 is that we have exactly as many free variables as we have vectors in the linear combination (instead of also having the extra constant vector added on). This is more ideal because, with the usual vector operations, the collection of all of these solutions is actually a vector space! We will explore this idea a bit further in the next section.

2.3 Spanning Sets and Linear Independence

2.4 Span and Spanning Sets

Notice that we can rewrite the linear system

$$\begin{cases} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n = b_1 \\ \vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n = b_m \end{cases}$$

as an equation of vectors

$$x_1 \begin{bmatrix} a_{11} \\ \vdots \\ a_{m1} \end{bmatrix} + x_2 \begin{bmatrix} a_{12} \\ \vdots \\ a_{m2} \end{bmatrix} + \cdots + x_n \begin{bmatrix} a_{1n} \\ \vdots \\ a_{mn} \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_m \end{bmatrix}$$

In this way a solution to the system corresponds to a linear combination.

Theorem 2.4.1 (Poole Theorem 2.4). *A system of linear equations $[A \mid \mathbf{b}]$ is consistent if and only if \mathbf{b} is a linear combination of the columns of A .*

The number of solutions to the system also tells us how many ways we can make such a linear combination. If there is a unique solution, then there is exactly one way. If there are infinitely-many solutions, there are infinitely-many ways to make the linear combination, so it may be reasonable to ask more qualitative questions about the set of all possible linear combinations and study the space of linear combinations instead.

Definition. Given a set of vectors $S = \{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ in a vector space V , we define the *span* of $\mathbf{v}_1, \dots, \mathbf{v}_k$ to be the set of all linear combinations of these vectors, and we write $\text{Span}(\mathbf{v}_1, \dots, \mathbf{v}_k)$ or $\text{Span}(S)$. If $V = \text{Span}(S)$, then we call S a *spanning set* for V .

With this definition, we can restate Theorem 2.4.1 as follows:

Theorem 2.4.2. *A system of linear equations $A\mathbf{x} = \mathbf{b}$ is consistent if and only if \mathbf{b} is in $\text{Span}(\mathbf{a}_1, \dots, \mathbf{a}_n)$ (where \mathbf{a}_i is the i^{th} column of A).*


Exercise 2.4.1. If V is a (real) vector space and $\mathbf{v}_1, \dots, \mathbf{v}_k$ is some collection of vectors in V , then the set $\text{Span}(\mathbf{v}_1, \dots, \mathbf{v}_k)$ is also (real) vector space.

Example 2.4.3. Let $\mathbf{e}_1 = [1, 0]^T$, $\mathbf{e}_2 = [0, 1]^T$ be the standard basis vectors for \mathbb{R}^2 . By definition, an arbitrary vector in $\text{Span}(\mathbf{e}_1, \mathbf{e}_2)$ is of the form

$$x\mathbf{e}_1 + y\mathbf{e}_2 = x \begin{bmatrix} 1 \\ 0 \end{bmatrix} + y \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \end{bmatrix}$$

and so $\text{Span}(\mathbf{e}_1, \mathbf{e}_2) = \mathbb{R}^2$.

Example 2.4.4. Let $\mathbf{e}_1 = [1, 0, 0]^T$ and $\mathbf{e}_2 = [0, 1, 0]^T$ be standard basis vectors in \mathbb{R}^3 . $\text{Span}(\mathbf{e}_1, \mathbf{e}_2)$ is the collection of all vectors in \mathbb{R}^3 of the form $[x, y, 0]$, which is just the xy -plane. This set does not span \mathbb{R}^3 , however, because it is missing all vectors with a nonzero 3rd coordinate (i.e. the z -direction).



Example 2.4.5. For any two vectors \mathbf{u}, \mathbf{v} in \mathbb{R}^3 (with \mathbf{u} not a scalar multiple of \mathbf{v}), then $\text{Span}(\mathbf{u}, \mathbf{v})$ is a plane through the origin in \mathbb{R}^3 .

2.4.1 Linear (In)dependence

Suppose we can write one vector \mathbf{w} as a linear combination of \mathbf{u} and \mathbf{v} , say $\mathbf{w} = a\mathbf{u} + b\mathbf{v}$. Then \mathbf{w} “depends” on \mathbf{u} and \mathbf{v} . Clearly we can rewrite this as

$$\mathbf{w} = a\mathbf{u} + b\mathbf{v} \quad \implies \quad a\mathbf{u} + b\mathbf{v} - \mathbf{w} = \mathbf{0},$$

and so we introduce the following definition.

Definition. A set of vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$ in a vector space is *linearly dependent* if there are scalars a_1, \dots, a_k (not all zero) such that

$$a_1\mathbf{v}_1 + \dots + a_k\mathbf{v}_k = \mathbf{0}.$$

A set of vectors that is not linearly dependent is called *linearly independent*.

Remark. It is always true that the equation above holds if $a_1 = \dots = a_k = 0$, so linearly dependence says that there is some other collection of a_i 's for which the equation is also true. In this way, linear independence can be thought of as saying that the only way the above equation is true is if $a_1 = \dots = a_k = 0$.

Example 2.4.6. The set $\mathbf{v}_1 = [1, 1]$, $\mathbf{v}_2 = [2, 1]$ and $\mathbf{v}_3 = [1, 2]$ is linearly dependent because $\mathbf{v}_2 + \mathbf{v}_3 - 3\mathbf{v}_1 = \mathbf{0}$.

It should be clear from the way we defined linear dependence that the idea is to capture when one vector can be written as a linear combination of the others. In the above example we can easily write $\mathbf{v}_3 = \frac{1}{3}\mathbf{v}_2 + \frac{1}{3}\mathbf{v}_2$. In fact, this is an equivalent characterization of linear dependence.

Theorem 2.4.7 (Poole Theorem 2.5). *The set of vectors $S = \{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ in a vector space is linearly dependent if and only if at least one of the vectors can be written as a linear combination of the others.*

The proof of this fact is essentially exactly what happens in the example, so we provide it fully below. The hard part is that we have to prove two separate things (because the statement of the theorem is a “biconditional statement”).

Proof. If S is linearly dependent, then we can find scalars a_1, \dots, a_k , not all zero, so that

$$a_1 \mathbf{v}_1 + \dots + a_k \mathbf{v}_k = \mathbf{0}.$$

Since one of the coefficients $a_i \neq 0$, then we can rearrange this as

$$\mathbf{v}_i = -\frac{a_1}{a_i} \mathbf{v}_1 - \dots - \frac{a_{i-1}}{a_i} \mathbf{v}_{i-1} - \frac{a_{i+1}}{a_i} \mathbf{v}_{i+1} - \dots - \frac{a_k}{a_i} \mathbf{v}_k.$$

Conversely, suppose \mathbf{v}_1 is nonzero and is a linear combination of the remaining vectors in S . Then there are constants a_2, \dots, a_k , not all zero, for which

$$\mathbf{v}_1 = a_2 \mathbf{v}_2 + \dots + a_k \mathbf{v}_k$$

which rearranges to

$$-\mathbf{v}_1 + a_2 \mathbf{v}_2 + \dots + a_k \mathbf{v}_k = \mathbf{0}.$$

hence S is linearly dependent. □

Remark. The above theorem is actually a bit subtle. It doesn't say that *every* vector can be written as a linear combination of the others, just that there's at least one that can be written this way.

Why is linear independence important?

It all comes down to uniqueness. If $\mathbf{v}_1, \dots, \mathbf{v}_n$ are linearly independent and $\mathbf{b} \in \text{Span}(\mathbf{v}_1, \dots, \mathbf{v}_n)$, then linear independence of $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ tells us that the linear combination

$$\mathbf{b} = c_1 \mathbf{v}_1 + \dots + c_n \mathbf{v}_n$$

is actually the unique one representing \mathbf{b} . This is great because it means that if two people can agree on a particular (ordered) linearly independent set of vectors, then we can make a vector of the coefficients $[c_1, \dots, c_n]^T$ that unambiguously represents the vector \mathbf{b} . Contrast this with the following example:

Example 2.4.8. Consider the vectors $\mathbf{v}_1 = [1, 0]^T$, $\mathbf{v}_2 = [0, 1]^T$, $\mathbf{v}_3 = [1, 1]^T$. Clearly, the set $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ is linearly dependent. So we can write $\mathbf{b} = [3, 1]^T$ in multiple ways, say,

$$\mathbf{b} = 3\mathbf{v}_1 + 1\mathbf{v}_2 + 0\mathbf{v}_3 = 2\mathbf{v}_1 + 0\mathbf{v}_2 + 1\mathbf{v}_3 = 0\mathbf{v}_1 - 2\mathbf{v}_2 + 3\mathbf{v}_3$$

and if we tried to refer to \mathbf{b} as just a vector of its coefficients, we could have that $[3, 1, 0]^T$, $[0, 2, 1]^T$, and $[0, -2, 3]^T$ all represent the same vector.

This particular idea will become especially important when we discuss bases for a vector space, but there are also several other ideas that we will see are reliant upon the notion of linear independence.

2.4.2 Using Matrices to Determine Linear (In)dependence

Given a collection of vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, determining linear (in)dependence comes down to finding whether (or not) there exist nonzero real numbers x_1, \dots, x_n such that

$$x_1\mathbf{v}_1 + \dots + x_n\mathbf{v}_n = \mathbf{0}$$

and this equivalent to checking whether or not the following system has any nontrivial solutions (i.e. solutions other than the zero vector).

$$A\mathbf{x} = \begin{bmatrix} \mathbf{v}_1 & \cdots & \mathbf{v}_n \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} = x_1\mathbf{v}_1 + \dots + x_n\mathbf{v}_n = \mathbf{0}.$$

Theorem 2.4.9 (Poole Theorem 2.6). *Let $\mathbf{v}_1, \dots, \mathbf{v}_m$ be vectors in \mathbb{R}^n and let A be the $n \times m$ matrix with these vectors as its columns, $A = [\mathbf{v}_1, \dots, \mathbf{v}_m]$. Then that collection of vectors is linearly dependent if and only if the homogeneous system $[A|\mathbf{0}]$ has at least one nontrivial solution.*

The following theorem is logically equivalent to the above, but is stated to make the connection between this system and linear independence completely clear.

Theorem 2.4.10. *Let $\mathbf{v}_1, \dots, \mathbf{v}_m$ be vectors in \mathbb{R}^n and let A be the $n \times m$ matrix with these vectors as its columns, $A = [\mathbf{v}_1, \dots, \mathbf{v}_m]$. Then that collection of vectors is linearly independent if and only if the homogeneous system $[A|\mathbf{0}]$ has no nontrivial solutions.*

Example 2.4.11. Consider the column vectors $\mathbf{v}_1 = [0, 1, 2]^T$, $\mathbf{v}_2 = [2, 1, 3]^T$, $\mathbf{v}_3 = [2, 0, 2]^T$. We can check for linear (in)dependence by row reducing $[\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3|\mathbf{0}]$ and checking for consistency.

$$\begin{aligned} [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3|\mathbf{0}] &= \begin{bmatrix} 0 & 2 & 2 & | & 0 \\ 1 & 1 & 0 & | & 0 \\ 2 & 3 & 2 & | & 0 \end{bmatrix} \longrightarrow \begin{array}{c} R_1 \leftrightarrow R_2 \end{array} \begin{bmatrix} 1 & 1 & 0 & | & 0 \\ 0 & 2 & 2 & | & 0 \\ 2 & 3 & 2 & | & 0 \end{bmatrix} \\ &\longrightarrow \begin{array}{c} \frac{1}{2}R_2 \mapsto R_2 \\ R_3 - 2R_1 \mapsto R_3 \end{array} \begin{bmatrix} 1 & 1 & 0 & | & 0 \\ 0 & 1 & 1 & | & 0 \\ 0 & 1 & 2 & | & 0 \end{bmatrix} \\ &\longrightarrow \begin{array}{c} R_1 - R_2 \mapsto R_1 \\ R_3 - R_2 \mapsto R_3 \end{array} \begin{bmatrix} 1 & 0 & -1 & | & 0 \\ 0 & 1 & 1 & | & 0 \\ 0 & 0 & 1 & | & 0 \end{bmatrix} \\ &\longrightarrow \begin{array}{c} R_1 + R_3 \mapsto R_1 \\ R_2 - R_3 \mapsto R_2 \end{array} \begin{bmatrix} 1 & 0 & 0 & | & 0 \\ 0 & 1 & 0 & | & 0 \\ 0 & 0 & 1 & | & 0 \end{bmatrix} \end{aligned}$$

The system thus has no nontrivial solutions (because there are no free variables). By Theorem 2.4.10, the set $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ is linearly independent.

Example 2.4.12. Consider the column vectors $\mathbf{v}_1 = [0, 1, 2]^T$, $\mathbf{v}_2 = [2, 1, 3]^T$, $\mathbf{v}_3 = [2, 0, 1]^T$. We can check linear dependence by using Gaussian (or Gauss-Jordan) Elimination to see if $[\mathbf{v}_1 \mathbf{v}_2 \mathbf{v}_3 | \mathbf{0}]$ is consistent.

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

It's an exercise to the reader to check that the reduced row echelon form is correct. The system thus has no nontrivial solutions. By Theorem 2.4.9 the set $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ is linearly dependent. Of course, this isn't hard to see, as $\mathbf{v}_2 = \mathbf{v}_1 + \mathbf{v}_3$.

Example 2.4.13. Let $\mathbf{v}_1 = [1, 1]^T$, $\mathbf{v}_2 = [1, -1]^T$, $\mathbf{v}_3 = [x, y]^T$ be column vectors in \mathbb{R}^2 . Using the same method from Theorem 2.4.9 we can check for linear dependence via Gaussian (or Gauss-Jordan) elimination:

$$[\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3 \mid \mathbf{0}] = \left[\begin{array}{ccc|c} 1 & 1 & x & 0 \\ 1 & -1 & y & 0 \end{array} \right] \xrightarrow{\text{RREF}} \left[\begin{array}{ccc|c} 1 & 0 & \frac{x-y}{2} & 0 \\ 0 & 1 & \frac{x+y}{2} & 0 \end{array} \right]$$

It's an exercise to the reader to check that the reduced row echelon form is correct. The system has nontrivial solutions for every nonzero vector $[x, y]^T$, so by Theorem 2.4.9 the set $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ is linearly dependent.

Of course, three nonzero vectors in \mathbb{R}^2 being linearly dependent doesn't sound too unreasonable - the corresponding linear system must have rank at most 2 and thus at least one free variable. We deduce from this the following result:

Theorem 2.4.14 (*Poole Theorem 2.8*). *Any set of m vectors in \mathbb{R}^n is linearly dependent if $m > n$.*

3.1 Matrix Operations

3.1.1 Matrix Basics

Definition. A *matrix* is an array of numbers (called *entries*) and has *size* $m \times n$ if it has m rows and n columns.

$$A = [a_{ij}] = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \vdots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix}$$

The subscripts on the entries a_{ij} tell us that we're looking at the entry in the i^{th} row and the j^{th} column.

Fact. Two matrices are equal if and only if both (1) their sizes are equal and (2) their corresponding entries are all equal.

Example 3.1.1. Consider the matrices

$$A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 2 \\ 4 & 5 \end{bmatrix}, \quad \text{and } C = \begin{bmatrix} 7 & 8 \\ 9 & 10 \end{bmatrix}.$$

$A \neq B$ because their sizes are different, and $A \neq C$ because their corresponding entries are not equal.

Definition. A matrix is *square* if it has size $n \times n$ and the *diagonal of the matrix* $A = [a_{ij}]$ are the entries where $i = j$.

$$A = \begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{bmatrix}$$

A square matrix is *diagonal* if the only nonzero entries are along the diagonal. You may see this written as $A = \text{diag}(a_{11}, \dots, a_{nn})$.

A bit of a subtlety – the definition of diagonal just says that nonzero entries must occur along the diagonal, but the diagonal entries *do not necessarily need to be nonzero*.

Example 3.1.2. The matrices

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 7 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

are both diagonal because every entry off of the diagonal is 0.

Definition. A matrix is *scalar* if it is a diagonal matrix and the diagonal entries are all equal, i.e., the matrix $A = \text{diag}(r, r, \dots, r)$ for some $r \in \mathbb{R}$.

Definition. The *zero matrix*, often denoted O is the matrix for which all entries are 0. Its size should be clear from context, but we may write $O_{m \times n}$ if we need to specify.

Definition. The *Kroenecker delta*, denoted δ_{ij} , δ_i^j , or δ^{ij} , is the following:

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

Definition. The *identity matrix* I_n is the diagonal $n \times n$ matrix with all 1's along the diagonal. You may sometimes see this written as $I_n = [\delta_{ij}]$.

Example 3.1.3. It may be useful to see exactly how the Kroenecker delta leads to the identity matrix.

$$I_3 = [\delta_{ij}] = \begin{bmatrix} \delta_{11} & \delta_{12} & \delta_{13} \\ \delta_{21} & \delta_{22} & \delta_{23} \\ \delta_{31} & \delta_{32} & \delta_{33} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

3.1.2 Matrix Operations

Definition. Given two $m \times n$ matrices $A = [a_{ij}]$ and $B = [b_{ij}]$, the *matrix sum* is

$$A + B = [a_{ij} + b_{ij}] = \begin{bmatrix} a_{11} + b_{11} & \cdots & a_{1n} + b_{1n} \\ \vdots & \ddots & \vdots \\ a_{m1} + b_{m1} & \cdots & a_{mn} + b_{mn} \end{bmatrix}$$

Example 3.1.4. For $A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$ and $B = \begin{bmatrix} 7 & 8 & 9 \\ 10 & 11 & 12 \end{bmatrix}$,

$$A + B = \begin{bmatrix} 1+7 & 2+8 & 3+9 \\ 4+10 & 5+11 & 6+12 \end{bmatrix} = \begin{bmatrix} 8 & 10 & 12 \\ 14 & 16 & 18 \end{bmatrix}.$$

Remark. When adding two matrices, they must have the same size.

Definition. For a matrix $A = [a_{ij}]$ and a scalar r , the *scalar multiple of A* is the matrix

$$rA = [ra_{ij}]$$

Example 3.1.5. Given the matrix $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$ and the scalar $r = 5$,

$$rA = \begin{bmatrix} 5(1) & 5(2) \\ 5(3) & 5(4) \end{bmatrix} = \begin{bmatrix} 5 & 10 \\ 15 & 20 \end{bmatrix}$$

Definition. Subtraction of matrices is then defined in the obvious way: $A - B = A + (-1)B$.

Remark. The zero matrix satisfies the properties you want it to: $A + \mathcal{O} = A$ and $A - A = \mathcal{O}$

Definition. Given an $m \times n$ matrix $A = [a_{ij}]$ and an $n \times p$ matrix $B = [b_{ij}]$, the *matrix product* of A and B is

$$AB = [c_{ij}]$$

where

$$c_{ij} = a_{i1}b_{1j} + a_{i2}b_{2j} + \cdots + a_{in}b_{nj}$$

in other words, the $(i, j)^{\text{th}}$ entry of AB is the dot product of the i^{th} row of A (thought of as a row vector) and the j^{th} of column B (thought of as a column vector).

Example 3.1.6. Let A and B be the following matrices

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix} = \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \mathbf{a}_3 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 7 & 8 \\ 9 & 10 \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 & \mathbf{b}_2 \end{bmatrix}$$

where the \mathbf{a}_i 's are vectors forming the rows in A and the \mathbf{b}_i 's are the vectors forming the columns of B . We have that

$$\begin{aligned} c_{11} &= \mathbf{a}_1 \cdot \mathbf{b}_1 = 1(7) + 2(9) = 23, & c_{12} &= \mathbf{a}_1 \cdot \mathbf{b}_2 = 1(8) + 2(10) = 28, \\ c_{21} &= \mathbf{a}_2 \cdot \mathbf{b}_1 = 3(7) + 4(9) = 57, & c_{22} &= \mathbf{a}_2 \cdot \mathbf{b}_2 = 3(8) + 4(10) = 64, \\ c_{31} &= \mathbf{a}_3 \cdot \mathbf{b}_1 = 5(7) + 6(9) = 89, & c_{32} &= \mathbf{a}_3 \cdot \mathbf{b}_2 = 5(8) + 6(10) = 100, \end{aligned}$$

and so the product matrix is

$$AB = [c_{ij}] = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \\ c_{31} & c_{32} \end{bmatrix} = \begin{bmatrix} 23 & 28 \\ 57 & 64 \\ 89 & 100 \end{bmatrix}.$$

Thinking of \mathbf{b}_1 and \mathbf{b}_2 as matrices of size 3×1 , we may notice that the first column of AB is precisely $A\mathbf{b}_1$ and the second column is $A\mathbf{b}_2$:

$$\begin{aligned} \begin{bmatrix} c_{11} \\ c_{21} \\ c_{31} \end{bmatrix} &= \begin{bmatrix} 1(7) + 2(9) \\ 3(7) + 4(9) \\ 5(7) + 6(9) \end{bmatrix} = 7 \begin{bmatrix} 1 \\ 3 \\ 5 \end{bmatrix} + 9 \begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix} = A\mathbf{b}_1. \\ \begin{bmatrix} c_{12} \\ c_{22} \\ c_{32} \end{bmatrix} &= \begin{bmatrix} 1(8) + 2(10) \\ 3(8) + 4(10) \\ 5(8) + 6(10) \end{bmatrix} = 8 \begin{bmatrix} 1 \\ 3 \\ 5 \end{bmatrix} + 10 \begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix} = A\mathbf{b}_2. \end{aligned}$$

As such, we have

$$AB = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \\ c_{31} & c_{32} \end{bmatrix} = \begin{bmatrix} A\mathbf{b}_1 & A\mathbf{b}_2 \end{bmatrix}$$

The matrix on the right is sometimes called the *matrix-column representation* of the product AB .

Remark. Two matrices A, B can be multiplied even if their sizes are different. As long as the number of columns of A is equal to the number of rows of B , then the product AB exists. Moreover

$$\begin{matrix} A & B & = & AB \\ m \times n & n \times p & & m \times p \end{matrix}$$

Fact. If A is an $m \times n$ matrix, then we have that

$$I_m A = A \quad \text{and} \quad A I_n = A$$

and it is for this reason that we call I_n the identity matrix.

Give linear system motivation

For now you'll have to take it on faith that our definition matrix multiplication, although odd and tedious, is actually the correct notion. The parallel story of matrices is that they correspond to certain functions on vector spaces, and matrix multiplication corresponds to function composition (which, as you can imagine, can be quite confusing). We'll fully make this connection later this semester.

3.1.3 Matrix Powers

If A is a square matrix, then for positive integers k , then we can define the power of a matrix in the intuitive way,

$$A^k = \underbrace{AA \cdots A}_{k \text{ factors}}$$

and the usual rules for exponents hold:

- $A^r A^s = A^{r+s}$
- $(A^r)^s = A^{rs}$

Example 3.1.7. If $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$, then

$$A^2 = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} = \begin{bmatrix} 7 & 10 \\ 15 & 22 \end{bmatrix}$$

and

$$A^3 = A^2 A = \begin{bmatrix} 7 & 10 \\ 15 & 22 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} = \begin{bmatrix} 37 & 54 \\ 81 & 118 \end{bmatrix}$$

3.1.4 Transpose

Definition. If $A = [a_{ij}]$ is an $m \times n$ matrix, then its *transpose*, $A^T = [a_{ji}]$ is the $n \times m$ matrix formed by swapping the $(i, j)^{\text{th}}$ and $(j, i)^{\text{th}}$ entries of A .

Visually, the transpose amounts to flipping the matrix across the red line below

$$A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} \xrightarrow{\text{flip}} \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix} = A^T$$

Definition. A matrix A is *symmetric* if $A = A^T$.

Example 3.1.8. $A = \begin{bmatrix} 1 & 2 \\ 2 & 3 \end{bmatrix}$ is symmetric, while $B = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$ is not.

Example 3.1.9. Let $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix}$. Then

$$\begin{aligned} A^T A &= \begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix} \\ &= \begin{bmatrix} 1^2 + 3^2 + 5^2 & 1(2) + 3(4) + 5(6) \\ 2(1) + 4(3) + 6(5) & 2^2 + 4^2 + 6^2 \end{bmatrix} \\ &= \begin{bmatrix} 35 & 44 \\ 44 & 56 \end{bmatrix} \end{aligned}$$

What's interesting to notice is that, while A is not symmetric (and not even square), $A^T A$ is symmetric (and hence also a square matrix).

3.2 Matrix Algebra

We'll introduce another piece of notation (that I probably should have used in the previous section):

$\mathbb{R}^{m \times n}$ represents the collection of $m \times n$ matrices with real number entries.

Theorem 3.2.1 (Poole Theorem 3.2 - Algebraic Properties of Matrix Addition and Scalar Multiplication). *Let $A, B, C \in \mathbb{R}^{m \times n}$ and let $c, d \in \mathbb{R}$. The following are true:*

(a) $A + B = B + A$

(b) $(A + B) + C = A + (B + C)$

(c) $A + O_{m \times n} = A$

(d) $A + (-A) = O_{m \times n}$

(e) $c(A + B) = cA + cB$

(f) $(c + d)A = cA + dA$

(g) $c(dA) = (cd)A$

(h) $1A = A$

Remark. In short, Theorem 3.2.1 above says that $\mathbb{R}^{m \times n}$ is a real vector space (see page ??).

3.2.1 Properties of Matrix Multiplication

Matrix multiplication is *not commutative in general*, and it is often the case that $AB \neq BA$. This fact is clear if $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{n \times m}$ where $m \neq n$ (just compare the sizes of AB and BA), but is possibly less obvious in the case where A, B are both square matrices. It is an exercise to find an example of this in the case of 2×2 matrices.

So what properties does matrix multiplication have?

Theorem 3.2.2 (Poole Theorem 3.3 - Properties of Matrix Multiplication). *Let A, B, C be matrices (whose sizes are such that the following exist) and $k \in \mathbb{R}$ a scalar. Then*

(a) $A(BC) = (AB)C$

(b) $A(B + C) = AB + AC$

(c) $(A + B)C = AC + BC$

(d) $k(AB) = (kA)B = A(kB)$

(e) $I_m A = A = A I_n$ (if A is $m \times n$)

Remark. This theorem implies that $\mathbb{R}^{n \times n}$ is a fancy object called a (*non-commutative*) *algebra*. Informally, this is a vector space with an additional operation that lets us multiply two vectors together (which, if you look closely, isn't a feature of vector spaces normally). This is outside the scope of the course, but it may be interesting to you to know that such things exist and that these properties are not unique to $\mathbb{R}^{n \times n}$.

The proof of this theorem will require the properties of the dot product (recall Proposition ??).

Proof. For simplicity, we'll introduce some notation. For a matrix M

- (this one is standard notation) M_{ij} denotes the $(i, j)^{\text{th}}$ entry of M ,
- $\text{row}_i(M)$ denotes the i^{th} row of M , and
- $\text{col}_j(M)$ denotes the j^{th} column of M .

- (a) This is clunky to do fully. Instead we'll note that for $A \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{n \times p}$ and $C \in \mathbb{R}^{p \times r}$, both products $(AB)C$ and $A(BC)$ have the same size $m \times r$, so at least there's hope these will be equal as matrices. Instead of proving the rest of this part, we'll just do an example to demonstrate that it's true. [Add proof later maybe?]
- (b) Let $A \in \mathbb{R}^{m \times n}$ and $B, C \in \mathbb{R}^{n \times p}$. Notice that $A(B + C)$ and $AB + AC$ have the same size, hence they are equal if they have the same corresponding elements.

$$\begin{aligned} (A(B + C))_{ij} &= \text{row}_i(A) \cdot \text{col}_j(B + C) \\ &= \text{row}_i(A) \cdot (\text{col}_j(B) + \text{col}_j(C)) \\ &= \text{row}_i(A) \cdot \text{col}_j(B) + \text{row}_i(A) \cdot \text{col}_j(C) = (AC)_{ij} + (BC)_{ij}. \end{aligned}$$

- (c) Let $A, B \in \mathbb{R}^{m \times n}$ and $C \in \mathbb{R}^{n \times p}$. Notice that $(A + B)C$ and $AC + BC$ have the same size, hence they are equal if they have the same corresponding elements.

$$\begin{aligned} ((A + B)C)_{ij} &= \text{row}_i(A + B) \cdot \text{col}_j(C) \\ &= (\text{row}_i(A) + \text{row}_i(B)) \cdot \text{col}_j(C) \\ &= \text{row}_i(A) \cdot \text{col}_j(C) + \text{row}_i(B) \cdot \text{col}_j(C) = (AC)_{ij} + (BC)_{ij}. \end{aligned}$$

- (d) Let $A \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{n \times p}$, and $k \in \mathbb{R}$. Notice that $k(AB)$, $(kA)B$ and $A(kB)$ all have the same size $m \times p$, hence they are equal if they have the same corresponding elements.

$$\begin{aligned} (k(AB))_{ij} &= k(\text{row}_i(A) \cdot \text{col}_j(B)) \\ &= \text{row}_i(kA) \cdot \text{col}_j(B) = ((kA)B)_{ij} \\ &= \text{row}_i(A) \cdot \text{col}_j(kB) = (A(kB))_{ij} \end{aligned}$$

- (e) Let $A \in \mathbb{R}^{m \times n}$. Writing the $m \times m$ identity matrix $I_m = [\delta_{ij}]$ using the Kroenecker delta (c.f. page 32), we note that $I_m A$ and A have the same size, hence they are equal if they have the same corresponding elements.

$$\begin{aligned} (I_m A)_{ij} &= \text{row}_i(I_m) \cdot \text{col}_j(A) \\ &= \delta_{i1}A_{1j} + \delta_{i2}A_{2j} + \cdots + \delta_{im}A_{mj} \\ &= \delta_{ii}A_{ij} && \text{(the only nonzero term)} \\ &= A_{ij} \end{aligned}$$

Similarly, for the $n \times n$ identity matrix I_n ,

$$\begin{aligned} (AI_n)_{ij} &= \text{row}_i(A) \cdot \text{col}_j(I_n) \\ &= A_{i1}\delta_{1j} + A_{i2}\delta_{2j} + \cdots + A_{in}\delta_{nj} \\ &= A_{ij}\delta_{jj} && \text{(the only nonzero term)} \\ &= A_{ij} \end{aligned}$$

□

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

$$\begin{aligned} A(BC) &= \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} \left(\begin{bmatrix} 3 & 2 \\ 1 & 0 \\ -1 & -2 \end{bmatrix} \begin{bmatrix} 1 & -2 \\ 3 & -4 \end{bmatrix} \right) \\ &= \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} \left(\begin{bmatrix} 9 & -14 \\ 1 & -2 \\ -7 & 10 \end{bmatrix} \right) \\ &= \begin{bmatrix} -10 & 12 \\ -1 & -6 \end{bmatrix} \end{aligned}$$

and

$$\begin{aligned} (AB)C &= \left(\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} \begin{bmatrix} 3 & 2 \\ 1 & 0 \\ -1 & -2 \end{bmatrix} \right) \begin{bmatrix} 1 & -2 \\ 3 & -4 \end{bmatrix} \\ &= \left(\begin{bmatrix} 2 & -4 \\ 11 & -4 \end{bmatrix} \right) \begin{bmatrix} 1 & -2 \\ 3 & -4 \end{bmatrix} \\ &= \begin{bmatrix} -10 & 12 \\ -1 & -6 \end{bmatrix} \end{aligned}$$

Theorem 3.2.4 (Poole Theorem 3.4 - Properties of the Transpose). *Let A and B be matrices (whose sizes are such that the indicated operations can be performed) and let k be a scalar. Then*

- (a) $(A^T)^T = A$
- (b) $(kA)^T = k(A^T)$
- (c) $(A^r)^T = (A^T)^r$ for all nonnegative integers r .
- (d) $(A + B)^T = A^T + B^T$
- (e) $(AB)^T = B^T A^T$

Proof. We use the same notation as in the proof of Theorem 3.2.2.

- (a) [Add proof later?]
- (b) [Add proof later?]
- (c) [Add proof later?]
- (d) [Add proof later?]
- (e) Let $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{n \times p}$. Note that $(AB)^T$ and $B^T A^T$ both have the same size, hence they are equal if their corresponding entries are equal.

$$\begin{aligned} ((AB)^T)_{ij} &= (AB)_{ji} \\ &= \text{row}_j(A) \cdot \text{col}_i(B) \\ &= \text{col}_j(A^T) \cdot \text{row}_i(B^T) \\ &= \text{row}_i(B^T) \cdot \text{col}_j(A^T) = B^T A^T. \end{aligned}$$

□

3.3 The Inverse of a Matrix

Motivation: If $a, b \in \mathbb{R}$ and x is some unknown and we wanted to solve for x in the equation

$$ax = b,$$

we would do so multiplying both sides by $a^{-1} = \frac{1}{a}$ to get that $x = a^{-1}b$. We'd like to be able to do this same thing for the system of linear equations

$$A\mathbf{x} = \mathbf{b}$$

where $A \in \mathbb{R}^{n \times n}$ and $\mathbf{b} \in \mathbb{R}^n$. But alas, we don't have a notion of division of matrices.

Let's think – if $a \in \mathbb{R}$ is some nonzero number, then a^{-1} is just some other real number for which $aa^{-1} = a^{-1}a = 1$. Since the $n \times n$ identity matrix I_n plays the role of 1, multiplicatively, then the natural way to define the matrix we desire is

Definition. For an nonzero $n \times n$ matrix A , the *inverse of A* , denoted A^{-1} , is the $n \times n$ matrix satisfying

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

If the inverse exists, we say that A is *invertible*.

Fact. Not every nonzero matrix is invertible, and we'll devote the latter half of this section to exploring when a matrix is invertible.

Remark. We only define inverses for square matrices. You will explore what happens in your homework for non-square matrices.

Theorem 3.3.1 (Poole Theorem 3.6). *The inverse is unique.*

Proof. Suppose X, Y are both inverses of A . Then

$$X = X(I_n) = X(AY) = (XA)Y = (I_n)Y = Y.$$

□

Theorem 3.3.2 (Poole Theorem 3.9). *If $A, B \in \mathbb{R}^{n \times n}$ are invertible and $c \in \mathbb{R}$ is some nonzero scalar, then*

- a. A^{-1} is invertible and $(A^{-1})^{-1} = A$.
- b. $(cA)^{-1} = \frac{1}{c}A^{-1}$
- c. $(AB)^{-1} = B^{-1}A^{-1}$
- d. A^T is invertible and $(A^T)^{-1} = (A^{-1})^T$
- e. A^n is invertible for all positive integers n and $(A^k)^{-1} = (A^{-1})^k$

Proof. Since inverses are unique, and matrices are invertible if their inverses exist, then each of these is essentially proven by merely checking that the multiplication checks out.

- a. $(A^{-1})(A) = I_n$
- b. $(cA)(\frac{1}{c}A^{-1}) = \frac{c}{c}AA^{-1} = I_n$.

- c. $(AB)(B^{-1}A^{-1}) = AI_nA^{-1} = AA^{-1} = I_n$.
- d. $A^T(A^{-1})^T = (A^{-1}A)^T = I_n^T = I_n$
- e. $(A^k)(A^{-1})^k = \underbrace{A \cdots A}_k \underbrace{A^{-1} \cdots A^{-1}}_k = \underbrace{A \cdots A}_{k-1} I_n \underbrace{A^{-1} \cdots A^{-1}}_{k-1} = \cdots = I_n$

□

Remark. Because of the above theorem, some will use the notation A^{-n} (for n a positive integer) and A^{-T} (the transpose) to mean the obvious things:

$$A^{-n} = (A^{-1})^n = (A^n)^{-1}$$

$$A^{-T} = (A^T)^{-1} = (A^{-1})^T$$

Suppose A is invertible. How do we find A^{-1} ?

Well, you secretly already did this on your last homework assignment *cue nefarious laughter*. If $A = [a_{ij}]$ is given and $X = [x_{ij}]$ is a matrix of indeterminates, The matrix equation $AX = I_n$ yields a linear system (with n^2 equations and n^2 unknowns)

$$\begin{cases} \sum_{k=1}^n a_{1k}x_{k1} & = 1 \\ \sum_{k=1}^n a_{1k}x_{k2} & = 0 \\ & \vdots \\ \sum_{k=1}^n a_{ik}x_{kj} & = \delta_{ij} \quad (\text{the Kroenecker delta}) \\ & \vdots \\ \sum_{k=1}^n a_{nk}x_{kn} & = 1 \end{cases}$$

and you can use standard techniques to solve this system.

Example 3.3.3. Find A^{-1} given $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$.

Let $X = \begin{bmatrix} x_1 & x_2 \\ x_3 & x_4 \end{bmatrix}$ be a matrix of indeterminates. Then the matrix equation

$$AX = I_2$$

$$\begin{bmatrix} x_1 + 2x_3 & x_2 + 2x_4 \\ 3x_1 + 4x_3 & 3x_2 + 4x_4 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

yields the system

$$\begin{cases} x_1 & + & 2x_3 & & = & 1 \\ & x_2 & & + & 2x_4 & = & 0 \\ 3x_1 & & + & 4x_3 & & = & 0 \\ & 3x_2 & & + & 4x_4 & = & 1 \end{cases}$$

And solving this in the usual way, we get that there is a unique solution $x_1 = -2$, $x_2 = 1$, $x_3 = \frac{3}{2}$, $x_4 = -\frac{1}{2}$. So

$$A^{-1} = \begin{bmatrix} -2 & 1 \\ \frac{3}{2} & -\frac{1}{2} \end{bmatrix}.$$

While this technique does work, it seems to get unwieldy pretty quickly because the system size is huge. The other way to solve for inverses (which was also on your homework), was via elementary matrices.

Definition. An *elementary matrix* is a matrix obtained by performing an elementary row operation on the identity matrix.

The other way to think about it is that, given a matrix $A \in \mathbb{R}^{m \times n}$, an elementary matrix $E \in \mathbb{R}^{m \times m}$ is one for which the product EA has the same effect as doing an elementary row operation on A .

One can then get the inverse as the product of all of the elementary matrices.

Example 3.3.4. Given $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$. Use elementary matrices to compute A^{-1} .

In the following string of equalities, we'll denote the row reduction on the left-hand side and the corresponding product by elementary matrices on the right-hand side.

$$\begin{aligned} \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} &= A \\ R_2 - 3R_1 \mapsto R_2 \quad \begin{bmatrix} 1 & 2 \\ 0 & -2 \end{bmatrix} &= \begin{bmatrix} 1 & 0 \\ -3 & 1 \end{bmatrix} A \\ \frac{-1}{2} R_2 \mapsto R_2 \quad \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix} &= \begin{bmatrix} 1 & 0 \\ 0 & -\frac{1}{2} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -3 & 1 \end{bmatrix} A \\ R_1 - 2R_2 \mapsto R_1 \quad \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} &= \begin{bmatrix} 1 & -2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & -\frac{1}{2} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -3 & 1 \end{bmatrix} A \\ I_2 &= \begin{bmatrix} -2 & 1 \\ \frac{3}{2} & -\frac{1}{2} \end{bmatrix} A \end{aligned}$$

so again we get that $A^{-1} = \begin{bmatrix} -2 & 1 \\ \frac{3}{2} & -\frac{1}{2} \end{bmatrix}$.

If we write the elementary matrices above as E_1, E_2, E_3 , respectively, (so that $A^{-1} = E_3E_2E_1$) then applying them to the matrix equation $AX = I$ looks something like the following

$$\begin{aligned} AX &= I \\ E_1AX &= E_1 \\ E_2E_1AX &= E_2E_1 \\ E_3E_2E_1AX &= E_3E_2E_1 \\ IX &= E_3E_2E_1 \end{aligned}$$

What this tells us is that computing the inverse can just be done with row operations on an augmented matrix $[A|I]$ until we get $[I|A^{-1}]$ (your book calls this the *Gauss-Jordan Method*).

Example 3.3.5. With A as before

$$\begin{aligned} & \left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 3 & 4 & 0 & 1 \end{array} \right] = [A \mid I] \\ & \xrightarrow{R_2 - 3R_1 \rightarrow R_2} \left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 0 & -2 & -3 & 1 \end{array} \right] \\ & \xrightarrow{\frac{-1}{2}R_2 \rightarrow R_2} \left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 0 & 1 & \frac{3}{2} & -\frac{1}{2} \end{array} \right] \\ & \xrightarrow{R_1 - 2R_2 \rightarrow R_1} \left[\begin{array}{cc|cc} 1 & 0 & -2 & 1 \\ 0 & 1 & \frac{3}{2} & -\frac{1}{2} \end{array} \right] = [I \mid A^{-1}] \end{aligned}$$

One can, of course, do the above process for generic 2×2 matrices, which yields the following result.

Theorem 3.3.6 (Poole thm 3.8). *If $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, then the inverse is given by $A^{-1} = \frac{1}{ad-bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$, provided $ad - bc \neq 0$.*

Proof. Using the Gauss-Jordan method above,

$$\begin{aligned} & \left[\begin{array}{cc|cc} a & b & 1 & 0 \\ c & d & 0 & 1 \end{array} \right] \xrightarrow{\frac{1}{a}R_1 \rightarrow R_1} \left[\begin{array}{cc|cc} 1 & \frac{b}{a} & \frac{1}{a} & 0 \\ c & d & 0 & 1 \end{array} \right] \\ & \xrightarrow{R_2 - cR_1 \rightarrow R_2} \left[\begin{array}{cc|cc} 1 & \frac{b}{a} & \frac{1}{a} & 0 \\ 0 & \frac{ad-bc}{a} & -\frac{c}{a} & 1 \end{array} \right] \\ & \xrightarrow{\frac{a}{ad-bc}R_2 \rightarrow R_2} \left[\begin{array}{cc|cc} 1 & \frac{b}{a} & \frac{1}{a} & 0 \\ 0 & 1 & -\frac{c}{ad-bc} & \frac{a}{ad-bc} \end{array} \right] \\ & \xrightarrow{R_1 - \frac{b}{a}R_2 \rightarrow R_1} \left[\begin{array}{cc|cc} 1 & 0 & \frac{d}{ad-bc} & -\frac{b}{ad-bc} \\ 0 & 1 & -\frac{c}{ad-bc} & \frac{a}{ad-bc} \end{array} \right] \end{aligned}$$

□

You could do this same system-solving process for larger matrices, but the formulas are significantly worse.

How does this help us solve our system?

If A is invertible and $A\mathbf{x} = \mathbf{b}$ then, by design, it should have the solution $\mathbf{x} = A^{-1}\mathbf{b}$. Moreover, since A^{-1} is unique, we expect this solution to be unique.

Theorem 3.3.7 (Poole Theorem 3.7). *If A is an invertible $n \times n$ matrix, then for every $\mathbf{b} \in \mathbb{R}^n$, the linear system $A\mathbf{x} = \mathbf{b}$ is consistent and has the unique solution $\mathbf{x} = A^{-1}\mathbf{b}$.*

Example 3.3.8. Solve the system $A\mathbf{x} = \mathbf{b}$ given $A = \begin{bmatrix} 1 & 4 \\ 3 & 13 \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} 1 \\ 7 \end{bmatrix}$.

We could solve this the old way, or we can try our nifty new method. We quickly deduce that A^{-1} is given by

$$A^{-1} = \begin{bmatrix} 13 & -4 \\ -3 & 1 \end{bmatrix}$$

(which can be seen either by appealing to Theorem 3.3.6 or using the Gauss-Jordan Method). Hence the solution is

$$\mathbf{x} = A^{-1}\mathbf{b} = \begin{bmatrix} 13 & -4 \\ -3 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 7 \end{bmatrix} = \begin{bmatrix} -15 \\ 4 \end{bmatrix}.$$

Example 3.3.9. Solve the system

$$\begin{cases} x + y = 1 \\ -x - y = -1 \end{cases}$$

Notice that this system is equivalent to the system $\{x + y = 1\}$, which has infinitely-many solutions. Notice also that the coefficient matrix for this system is

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

which isn't invertible (because otherwise, attempting to apply Theorem 3.3.6 we would be dividing by 0).

Okay, so tell me, when is A invertible?

Putting it all together, we can wrap it up into the following theorem

Theorem 3.3.10 (Poole Theorem 3.12 - The Fundamental Theorem of Invertible Matrices: Pt I). *Let $A \in \mathbb{R}^{n \times n}$. The following are equivalent:*

- a. A is invertible.*
- b. A is row equivalent to I_n (i.e. its reduced row echelon form is I_n).*
- c. A is the product of elementary matrices.*
- d. $A\mathbf{x} = \mathbf{b}$ has a unique solution for every $\mathbf{b} \in \mathbb{R}^n$.*
- e. $A\mathbf{x} = \mathbf{0}$ has only the trivial solution.*
- f. The columns of A are linearly independent.*

For part (c) above, we saw that it worked in one particular example. In fact, it's true for all \mathbf{b} because, if A is row equivalent to the identity, then it will never have a row of all 0's in its row-reduced form (which is precisely what happened in the cases that $A\mathbf{x} = \mathbf{b}$ had no solutions or infinitely many solutions).

3.5 Subspaces, Basis, Dimension, and Rank

We've thought about solution sets as spans of vectors and also, alternatively, as lines and planes in 3-dimensional space. Now we'll formalize these ideas so that we can talk about these things in more generality. Recall the definition of a real vector space:

Definition. A (*real*) *vector space* V is a set of objects (called *vectors*) with two operations *vector addition* (denoted $+$) and *scalar multiplication* (no symbol) satisfying the following properties: For all vectors $\mathbf{u}, \mathbf{v}, \mathbf{w}$ and real numbers a, b (called *scalars*),

- (a) $\mathbf{u} + \mathbf{v}$ is in V [closure of addition]
- (b) $\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$ [commutativity of addition]
- (c) $(\mathbf{u} + \mathbf{v}) + \mathbf{w} = \mathbf{u} + (\mathbf{v} + \mathbf{w})$ [associativity of $+$]
- (d) There is some vector $\mathbf{0}$, called the *zero vector*, so that $\mathbf{u} + \mathbf{0} = \mathbf{u}$ for all vectors \mathbf{u} . [additive identity]
- (e) For each \mathbf{u} in V , there is some vector $-\mathbf{u}$ for which $\mathbf{u} + (-\mathbf{u}) = \mathbf{0}$. [additive inverse]
- (f) $a\mathbf{u}$ is in V [closure of scalar mult.]
- (g) $a(\mathbf{u} + \mathbf{v}) = a\mathbf{u} + a\mathbf{v}$ [distributivity]
- (h) $(a + b)\mathbf{u} = a\mathbf{u} + b\mathbf{u}$ [distributivity]
- (i) $(ab)\mathbf{u} = a(b\mathbf{u})$ [associativity of scalar mult.]
- (j) $1\mathbf{u} = \mathbf{u}$ [scalar mult. identity]

With this in mind, we introduce the following definition:

Definition. Let V be a vector space and let W be a subset of vectors in V . We say that W is a *subspace* of V if it is also a vector space (with the same vector addition/scalar multiplication operations).

In order to check that a set of vectors is a subspace, one would have to check all of the axioms of the vector space definition – eww. Instead, here is an equivalent characterization of a subspace (note: this is typically a theorem in most textbooks, but your book presents it as the definition).

Definition. Let V be a vector space and let W be a subset of vectors in V . W is a *subspace* of V if it has the following properties:

1. $\mathbf{0}$ is in W (where $\mathbf{0}$ is the same zero vector in V).
2. If $\mathbf{u}, \mathbf{v} \in W$, then $\mathbf{u} + \mathbf{v} \in W$. [closure of addition]
3. If $\mathbf{u} \in W$ and $k \in \mathbb{R}$ is a scalar, then $k\mathbf{u} \in W$. [closure of scalar multiplication]

(where vector addition and scalar multiplication in W are the same operations for V).

Example 3.5.1. Every vector space V is a subspace of itself.

Example 3.5.2. For any vector space V , the set $\{\mathbf{0}\}$ is a subspace of V (sometimes called the *trivial subspace*).

Example 3.5.3. Let W be the set of all vectors in \mathbb{R}^3 of the form $[x, y, 0]$. Then W is a subspace of \mathbb{R}^3 because

1. $[0, 0, 0] \in W$
2. $[x_1, y_1, 0] + [x_2, y_2, 0] = [x_1 + x_2, y_1 + y_2, 0] \in W$
3. $k[x, y, 0] = [kx, ky, 0] \in W$

Example 3.5.4. Let W be the set of all vectors in \mathbb{R}^3 of the form $[x, y, 1]$. Then W is not a subspace of \mathbb{R}^3 because it fails property 1 in the definition: the zero vector $[0, 0, 0]$ does not have the form $[x, y, 1]$. (It actually fails the other two properties as well, but this is left as an exercise to the reader).

Theorem 3.5.5 (Poole Theorem 3.19). *Let $\mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{R}^n$. Then $\text{Span}(\mathbf{v}_1, \dots, \mathbf{v}_k)$ is a subspace of \mathbb{R}^n .*

Proof. For simplicity, let $W = \text{Span}(\mathbf{v}_1, \dots, \mathbf{v}_k)$.

1. Since $\mathbf{0} = 0\mathbf{v}_1 + \dots + 0\mathbf{v}_k$, then $\mathbf{0} \in W$.
2. Let $\mathbf{u} = c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k$ and $\mathbf{w} = d_1\mathbf{v}_1 + \dots + d_k\mathbf{v}_k$ be vectors in W . Then

$$\mathbf{u} + \mathbf{w} = (c_1 + d_1)\mathbf{v}_1 + \dots + (c_k + d_k)\mathbf{v}_k.$$

Since $\mathbf{u} + \mathbf{w}$ is a linear combination of the vectors \mathbf{v}_i , then $\mathbf{u} + \mathbf{w} \in W$.

3. Let \mathbf{u} be as above and $c \in \mathbb{R}$ be some scalar. Then

$$c\mathbf{u} = (cc_1)\mathbf{v}_1 + \dots + (cc_k)\mathbf{v}_k.$$

Since $c\mathbf{u}$ is a linear combination of the vectors \mathbf{v}_i , then $c\mathbf{u} \in W$.

□

3.5.1 Subspaces Associated with Matrices

With the notion of a "subspace" in mind, let's try to revisit some ideas involving matrices. First, a new definition

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

1. The *column space* of A is a subspace \mathbb{R}^n spanned by the columns of A . We denote it as $\text{Col}(A)$.
2. The *row space* of A is a subspace of \mathbb{R}^m spanned by the rows of A . We denote it as $\text{Row}(A)$

Remark. Since we will prefer to think about column vectors whenever possible, it may be more useful to define $\text{Row}(A) := \text{Col}(A^T)$.

Theorem 3.5.6 (Poole Theorem 3.21). *Let A be an $m \times n$ matrix and let N be the set of solutions to the homogeneous system $A\mathbf{x} = \mathbf{0}$. Then N is a subspace of \mathbb{R}^n .*

Proof. 1. Clear $\mathbf{x} = \mathbf{0}$ is a solution to the homogeneous system.

2. Let $\mathbf{x}_1, \mathbf{x}_2$ be in N . Then

$$A(\mathbf{x}_1 + \mathbf{x}_2) = A\mathbf{x}_1 + A\mathbf{x}_2 = \mathbf{0} + \mathbf{0} = \mathbf{0}.$$

so N is closed under addition.

3. Let \mathbf{x} be in N and k be a scalar. Then

$$A(k\mathbf{x}) = k(A\mathbf{x}) = k\mathbf{0} = \mathbf{0}$$

so N is closed under scalar multiplication. □

Definition. N , as above, is called the *null space* of A , and is denoted $\text{Null}(A)$. (In some texts, it is called the *kernel* of A and is denoted $\ker(A)$.)

Example 3.5.7. Compute $\text{Col}(A)$, $\text{Row}(A)$, and $\text{Null}(A)$ for $A = \begin{bmatrix} 1 & 2 & 4 \\ 1 & 2 & 4 \\ 1 & 2 & 4 \end{bmatrix}$.

Letting \mathbf{a}_i denote the i^{th} column of A , we see that $\mathbf{a}_3 = 2\mathbf{a}_2 = 4\mathbf{a}_1$, hence

$$\text{Col}(A) = \text{Span} \left(\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right).$$

Similarly, letting \mathbf{A}_i denote the i^{th} row of A , we see that $\mathbf{A}_3 = \mathbf{A}_2 = \mathbf{A}_1$, hence

$$\text{Row}(A) = \text{Span}([1, 2, 4]).$$

Examining the homogeneous system $A\mathbf{x} = \mathbf{0}$,

$$[A \mid \mathbf{0}] = \left[\begin{array}{ccc|c} 1 & 2 & 4 & 0 \\ 1 & 2 & 4 & 0 \\ 1 & 2 & 4 & 0 \end{array} \right] \xrightarrow[\substack{R_2 - R_1 \rightarrow R_2 \\ R_3 - R_1 \rightarrow R_3}]{R_3 - R_1 \rightarrow R_3} \left[\begin{array}{ccc|c} 1 & 2 & 4 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{array} \right]$$

we see that A has rank 1 (hey wait, that's how many vectors span both $\text{Row}(A)$ and $\text{Col}(A)$... weird), hence there are two free variables in this system, $x_2 = s$ and $x_3 = t$. We thus get that the solution set is of the form

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} -2s - 4t \\ s \\ t \end{bmatrix} = s \begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix} + t \begin{bmatrix} -4 \\ 0 \\ 1 \end{bmatrix}$$

hence

$$\text{Null}(A) = \text{Span} \left(\begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -4 \\ 0 \\ 1 \end{bmatrix} \right).$$

Definition. Let W be a subspace of a vector space and $\mathcal{B} = \{w_1, \dots, w_k\}$ a set of vectors in W . \mathcal{B} is a *basis* for W if

1. $W = \text{Span}(\mathcal{B})$ and
2. \mathcal{B} is a linearly independent set.

Remark. Since every vector space is a subspace of itself, this definition is valid for all vector spaces. We've merely stated it in terms of subspaces to make it clear that the basis vectors must each be contained in that subspace.

Example 3.5.8. The standard basis vectors \mathbf{e}_i in \mathbb{R}^n form a basis for \mathbb{R}^n . We refer to $\mathcal{E} = \{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ as the *standard basis* for \mathbb{R}^n .

Example 3.5.9. Find basis for the column space $A = \begin{bmatrix} 1 & 1 & 3 & 1 & 6 \\ 2 & -1 & 0 & 1 & -1 \\ -3 & 2 & 1 & -2 & 1 \\ 4 & 1 & 6 & 1 & 3 \end{bmatrix}$.

Notice that

$$\text{RREF}(A) = \begin{bmatrix} 1 & 0 & 1 & 0 & -1 \\ 0 & 1 & 2 & 0 & 3 \\ 0 & 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Let \mathbf{a}_i represent the i^{th} column vector for $A = [\mathbf{a}_1 \ \dots \ \mathbf{a}_5]$. What the reduced row echelon form suggests is that

$$\mathbf{a}_3 = 1\mathbf{a}_1 + 2\mathbf{a}_2 \quad \text{and} \quad \mathbf{a}_5 = -1\mathbf{a}_1 + 3\mathbf{a}_2 + 4\mathbf{a}_4$$

(and you can check that this is true). Hence $\text{Col}(A) = \text{Span}(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_4)$. Moreover,

$$\text{RREF}([\mathbf{a}_1 \ \mathbf{a}_2 \ \mathbf{a}_4 \mid \mathbf{0}]) = \left[\begin{array}{ccc|c} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{array} \right]$$

and so $\{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_4\}$ is a linearly independent set. Hence

$$\{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_4\} = \left\{ \begin{bmatrix} 1 \\ 2 \\ -3 \\ 4 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \\ 2 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ -2 \\ 1 \end{bmatrix} \right\}$$

is a basis for $\text{Col}(A)$.

Strategy for finding a basis for the column space of a matrix:

1. Row reduce the matrix (just row-echelon form is fine)
2. Take as a basis every column (in the original matrix) which contains a leading 1/pivot.

Example 3.5.10. Find basis for the null space of $A = \begin{bmatrix} 1 & 1 & 3 & 1 & 6 \\ 2 & -1 & 0 & 1 & -1 \\ -3 & 2 & 1 & -2 & 1 \\ 4 & 1 & 6 & 1 & 3 \end{bmatrix}$.

Notice that

$$\text{RREF}(A) = \begin{bmatrix} 1 & 0 & 1 & 0 & -1 \\ 0 & 1 & 2 & 0 & 3 \\ 0 & 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

When we go to solve the system

$$A\mathbf{x} = A \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = \mathbf{0}$$

we see that x_3 and x_5 are free variables and

$$\begin{aligned} x_1 &= -x_3 + x_5, \\ x_2 &= -2x_3 - 3x_5, \\ x_4 &= -4x_5. \end{aligned}$$

By setting $x_3 = s$ and $x_5 = t$, we can parameterize the solution set as

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = \begin{bmatrix} -s + t \\ -2s - 3t \\ s \\ -4t \\ t \end{bmatrix} = s \begin{bmatrix} -1 \\ -2 \\ 1 \\ 0 \\ 0 \end{bmatrix} + t \begin{bmatrix} 1 \\ -3 \\ 0 \\ -4 \\ 1 \end{bmatrix} = s\mathbf{v}_1 + t\mathbf{v}_2.$$

Clearly $\text{Null}(A) = \text{Span}(\mathbf{v}_1, \mathbf{v}_2)$ and it is straightforward to check that $\{\mathbf{v}_1, \mathbf{v}_2\}$ are linearly independent, hence

$$\{\mathbf{v}_1, \mathbf{v}_2\} = \left\{ \begin{bmatrix} -1 \\ -2 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ -3 \\ 0 \\ -4 \\ 1 \end{bmatrix} \right\}$$

is a basis for $\text{Null}(A)$.

Strategy for finding a basis for the null space of a matrix A :

1. Row reduce the matrix.
2. Solve the system $A\mathbf{x} = \mathbf{0}$.
3. Solve for all columns with leading 1's/pivots in terms of free variables.
4. Write the solution set as a parameterized linear combination of k vectors (where k is the number of free variables).
5. Take these k vectors as a basis for $\text{Null}(A)$.

How many bases can a vector space have?

Infinitely-many:

Exercise 3.5.1. Show that, for any nonzero real numbers m, n , the set $\left\{ \begin{bmatrix} m \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ n \end{bmatrix} \right\}$ is a basis for \mathbb{R}^2 .

Maybe the better question is

How many vectors must a basis have? Can two different bases have different numbers of vectors?

Example 3.5.11. Suppose V is a subspace of \mathbb{R}^3 and has two different bases, $\mathcal{B}_1 = \{\mathbf{u}_1, \mathbf{u}_2\}$ and $\mathcal{B}_2 = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$.

Since \mathcal{B}_1 is a basis, we can write each of the \mathcal{B}_2 -basis vectors as a linear combination of the \mathcal{B}_1 -basis vectors

$$\mathbf{v}_1 = a_{11}\mathbf{u}_1 + a_{21}\mathbf{u}_2$$

$$\mathbf{v}_2 = a_{12}\mathbf{u}_1 + a_{22}\mathbf{u}_2$$

$$\mathbf{v}_3 = a_{13}\mathbf{u}_1 + a_{23}\mathbf{u}_2$$

Now if we consider the vector equation

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

it *should be* that the only solution is when each of the $x_i = 0$ (since the we claim the \mathbf{v}_i 's are linearly independent). Notice, however, that

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

$$\mathbf{0} = x_1(a_{11}\mathbf{u}_1 + a_{21}\mathbf{u}_2) + x_2(a_{12}\mathbf{u}_1 + a_{22}\mathbf{u}_2) + x_3(a_{13}\mathbf{u}_1 + a_{23}\mathbf{u}_2)$$

$$\mathbf{0} = (a_{11}x_1 + a_{12}x_2 + a_{13}x_3)\mathbf{u}_1 + (a_{21}x_1 + a_{22}x_2 + a_{23}x_3)\mathbf{u}_2$$

Since the \mathbf{u}_i 's are linearly independent, we must have that

$$\begin{cases} a_{11}x_1 + a_{12}x_2 + a_{13}x_3 = 0 \\ a_{21}x_1 + a_{22}x_2 + a_{23}x_3 = 0 \end{cases}$$

which is a homogeneous linear system in 3 variables (and only 2 equations), hence there are infinitely many solutions, and hence \mathcal{B}_2 is a linearly dependent set and not a basis.

3.5.2 Dimension and Rank

Theorem 3.5.12 (Poole Theorem 3.23 - The Basis Theorem). *Let V be a vector space with two different bases \mathcal{B}_1 and \mathcal{B}_2 . Then \mathcal{B}_1 and \mathcal{B}_2 have the same number of vectors.*

Because the number of vectors in the basis is invariant of the choice of basis, we can define the following term.

Definition. The *dimension* of a vector space V is the number of vectors in a basis for V . We denote this $\dim(V)$.

Remark. The trivial vector space $\{\mathbf{0}\}$ is defined to have dimension 0.

Example 3.5.13. $\dim(\mathbb{R}^n) = n$.

Example 3.5.14. Let $V = \text{Span} \left(\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right)$ be a subspace of \mathbb{R}^3 . Clearly the first vector is a linear combination of the other two (which are linearly independent), so $\dim(V) = 2$.

Theorem 3.5.15 (Poole Theorem 3.24). *For a matrix A , $\dim(\text{Row}(A)) = \dim(\text{Col}(A))$.*

Definition. The *rank* of a matrix A is the dimension of its column space (denoted $\text{Rank}(A)$). If A has size $n \times n$ and $\text{Rank}(A) = n$, then sometimes we say that A has *full rank*.

Remark. This new notion of rank still agrees with our old notion, because the number of linearly independent rows in A is the same as the number of nonzero rows in $RREF(A)$ and $\dim(\text{Row}(A)) = \dim(\text{Col}(A))$.

Theorem 3.5.16 (Poole Theorem 3.25). For any matrix A , $\text{Rank}(A) = \text{Rank}(A^T)$

Definition. The *nullity* of a matrix A is the dimension of its null space. We denote it by $\text{nullity}(A)$.

Theorem 3.5.17 (Poole Theorem 3.26 - Rank-Nullity). If A is an $m \times n$ matrix, then $\text{Rank}(A) + \text{nullity}(A) = n$.

Example 3.5.18. Show that $\mathcal{B} = \left\{ \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 2 \\ 1 \\ -2 \end{bmatrix}, \begin{bmatrix} 3 \\ 1 \\ -3 \end{bmatrix} \right\}$ is a basis for \mathbb{R}^3

We consider the matrix $A = \begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 1 \\ 1 & -2 & -3 \end{bmatrix}$ and compute its rank.

$$\begin{aligned} \begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 1 \\ 1 & -2 & -3 \end{bmatrix} &\xrightarrow{R_3 - R_1 \rightarrow R_3} \begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 1 \\ 0 & -4 & -6 \end{bmatrix} \\ &\xrightarrow{R_3 + 4R_2 \rightarrow R_3} \begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 1 \\ 0 & 0 & -2 \end{bmatrix} \end{aligned}$$

Thus $\text{Rank}(A) = 3$. Hence the columns of A span a 3-dimensional subspace of \mathbb{R}^3 , i.e. $\text{Col}(A) = \mathbb{R}^3$. Since the columns of A are linearly independent and span \mathbb{R}^3 , \mathcal{B} is a basis for \mathbb{R}^3 .

Example 3.5.19. Show that $\mathcal{B} = \left\{ \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}, \begin{bmatrix} 2 \\ 0 \\ 2 \end{bmatrix}, \begin{bmatrix} 3 \\ -7 \\ 3 \end{bmatrix} \right\}$ is not a basis for \mathbb{R}^3 .

Employ the same technique as in the previous example, see that the rank is only 2.

Theorem 3.5.20 (Fundamental Theorem of Invertible Matrices, Pt II). Suppose A is an $n \times n$ matrix. The following are equivalent:

- a. A is invertible.
- ⋮
- g. The column vectors of A span \mathbb{R}^n .
- h. The column vectors of A form a basis for \mathbb{R}^n .
- i. The row vectors of A are linearly independent.
- j. The row vectors of A span \mathbb{R}^n .
- k. The row vectors of A form a basis for \mathbb{R}^n .
- l. $\text{Rank}(A) = n$
- m. $\text{nullity}(A) = 0$

3.5.3 Coordinates

Theorem 3.5.21. Let V be a vector space with an ordered basis $\mathcal{B} = \{\mathbf{v}_1, \dots, \mathbf{v}_k\}$. For every vector $\mathbf{u} \in V$, there is a unique linear combination of \mathcal{B} -basis vectors such that

$$\mathbf{u} = c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k.$$

Definition. The c_i in the previous theorem are called the *coordinates of \mathbf{u} with respect to \mathcal{B}* and the column vector

$$[\mathbf{u}]_{\mathcal{B}} = \begin{bmatrix} c_1 \\ \vdots \\ c_k \end{bmatrix}$$

is called the *coordinate vector of \mathbf{u} with respect to \mathcal{B}* .

Example 3.5.22. Let $P(0,0)$ and $Q(3,1)$ be points in the plane and consider the vector $\mathbf{v} = \overrightarrow{PQ}$. Given the standard basis $\mathcal{E} = \{\mathbf{e}_1, \mathbf{e}_2\}$ for \mathbb{R}^2 , we can write

$$\mathbf{v} = 3\mathbf{e}_1 + 1\mathbf{e}_2$$

hence

$$[\mathbf{v}]_{\mathcal{E}} = \begin{bmatrix} 3 \\ 1 \end{bmatrix}.$$


Example 3.5.23. With P, Q, \mathbf{v} as above, we consider now the basis $\mathcal{B} = \{\mathbf{b}_1, \mathbf{b}_2\} = \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \end{bmatrix} \right\}$ for \mathbb{R}^2 . Since

$$\mathbf{v} = 2\mathbf{b}_1 + 1\mathbf{b}_2$$


then

$$[\mathbf{v}]_{\mathcal{B}} = \begin{bmatrix} 2 \\ 1 \end{bmatrix}.$$

Visually, the previous two examples just give a bit of formality to the hand-wavy "coordinate grid" discussion from Section 1.1



\mathbf{v} in the \mathcal{E} -basis.



\mathbf{v} in the \mathcal{B} -basis.

Remark. We typically don't write the subscript \mathcal{E} for vectors when they are written in the standard basis.

6.3 Change of Basis

Let $\mathcal{E} = \left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}$ be the standard basis for \mathbb{R}^2 and let $\mathcal{B} = \left\{ \mathbf{b}_1 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \mathbf{b}_2 = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$ another basis for \mathbb{R}^2 . Any vector $\mathbf{v} \in \mathbb{R}^2$ can be represented in coordinates as $[\mathbf{v}]_{\mathcal{E}}$ or $[\mathbf{v}]_{\mathcal{B}}$ by just solving for a linear combination. Rather than solve for two different linear combinations, however, it would be useful to solve only one linear combination and then find a way to convert between the bases. Notice that

$$\mathbf{b}_1 = 1\mathbf{e}_1 - 1\mathbf{e}_2 \quad \text{and} \quad \mathbf{b}_2 = 1\mathbf{e}_1 + 1\mathbf{e}_2. \quad (6.3.1)$$

Now, given $[\mathbf{v}]_{\mathcal{B}} = [x, y]^T$, we have that

$$\begin{aligned} \begin{bmatrix} \mathbf{b}_1 & \mathbf{b}_2 \end{bmatrix} [\mathbf{v}]_{\mathcal{B}} &= \begin{bmatrix} \mathbf{b}_1 & \mathbf{b}_2 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \\ &= x\mathbf{b}_1 + y\mathbf{b}_2 \\ &= x(\mathbf{e}_1 - \mathbf{e}_2) + y(\mathbf{e}_1 + \mathbf{e}_2) \\ &= (x + y)\mathbf{e}_1 + (-x + y)\mathbf{e}_2 \\ &= \begin{bmatrix} \mathbf{e}_1 & \mathbf{e}_2 \end{bmatrix} \begin{bmatrix} x + y \\ -x + y \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{e}_1 & \mathbf{e}_2 \end{bmatrix} [\mathbf{v}]_{\mathcal{E}} \end{aligned}$$

It should be clear that

$$[\mathbf{e}_1]_{\mathcal{E}} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad \text{and} \quad [\mathbf{e}_2]_{\mathcal{E}} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

so if we write out our matrix columns in the \mathcal{E} -basis coordinates, we get

$$\begin{bmatrix} [\mathbf{b}_1]_{\mathcal{E}} & [\mathbf{b}_2]_{\mathcal{E}} \end{bmatrix} [\mathbf{v}]_{\mathcal{B}} = \begin{bmatrix} [\mathbf{e}_1]_{\mathcal{E}} & [\mathbf{e}_2]_{\mathcal{E}} \end{bmatrix} [\mathbf{v}]_{\mathcal{E}} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} [\mathbf{v}]_{\mathcal{E}} = [\mathbf{v}]_{\mathcal{E}}$$

and so the matrix

$$P_{\mathcal{E} \leftarrow \mathcal{B}} := \begin{bmatrix} [\mathbf{b}_1]_{\mathcal{E}} & [\mathbf{b}_2]_{\mathcal{E}} \end{bmatrix}$$

has the feature that it converts vectors from the \mathcal{B} -basis to the \mathcal{E} -basis.

Example 6.3.1. With \mathcal{E} and \mathcal{B} above, find the coordinate representation of $\mathbf{v} = [3, 1]^T$ in both bases, and verify that $P_{\mathcal{E} \leftarrow \mathcal{B}}$ converts these representations accordingly.

One can easily verify that

$$[\mathbf{v}]_{\mathcal{E}} = \begin{bmatrix} 3 \\ 1 \end{bmatrix} \quad \text{and} \quad [\mathbf{v}]_{\mathcal{B}} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$$

and from Equation 6.3.1, one can also readily see that

$$[\mathbf{b}_1]_{\mathcal{E}} = \begin{bmatrix} 1 \\ -1 \end{bmatrix} \quad \text{and} \quad [\mathbf{b}_2]_{\mathcal{E}} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Thus

$$\begin{bmatrix} [\mathbf{b}_1]_{\mathcal{E}} & [\mathbf{b}_2]_{\mathcal{E}} \end{bmatrix} [\mathbf{v}]_{\mathcal{B}} = \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \end{bmatrix} = \begin{bmatrix} 3 \\ 1 \end{bmatrix} = [\mathbf{v}]_{\mathcal{E}}.$$

Our choice to write the matrix columns in the \mathcal{E} -basis was just to give us the identity matrix. If we had two non-standard bases, we could also do the same thing (and one of the matrices would always be the identity), hence we take the following definition

Definition. Let $\mathcal{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ and $\mathcal{C} = \{\mathbf{c}_1, \dots, \mathbf{c}_n\}$ be two ordered bases for \mathbb{R}^n . The $n \times n$ matrix

$$P_{\mathcal{C} \leftarrow \mathcal{B}} = \begin{bmatrix} [\mathbf{b}_1]_{\mathcal{C}} & \cdots & [\mathbf{b}_n]_{\mathcal{C}} \end{bmatrix}.$$

is called the *change-of-basis matrix* from \mathcal{B} to \mathcal{C} . It has the effect

$$P_{\mathcal{C} \leftarrow \mathcal{B}} [\mathbf{v}]_{\mathcal{B}} = [\mathbf{v}]_{\mathcal{C}}.$$

Example 6.3.2. Let $\mathcal{B} = \left\{ \mathbf{b}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \mathbf{b}_2 = \begin{bmatrix} 1 \\ -1 \end{bmatrix} \right\}$ and $\mathcal{C} = \left\{ \mathbf{c}_1 = \begin{bmatrix} 1 \\ -2 \end{bmatrix}, \mathbf{c}_2 = \begin{bmatrix} -3 \\ 3 \end{bmatrix} \right\}$ be bases for \mathbb{R}^2 . Compute the change of basis matrix $P_{\mathcal{C} \leftarrow \mathcal{B}}$.

We need to find $[\mathbf{b}_1]_{\mathcal{C}}$ and $[\mathbf{b}_2]_{\mathcal{C}}$.

$$\begin{aligned} \mathbf{b}_1 = -2\mathbf{c}_2 - \mathbf{c}_1 &\implies [\mathbf{b}_1]_{\mathcal{C}} = \begin{bmatrix} -2 \\ -1 \end{bmatrix} \\ \mathbf{b}_2 = 0\mathbf{c}_2 - \frac{1}{3}\mathbf{c}_2 &\implies [\mathbf{b}_2]_{\mathcal{C}} = \begin{bmatrix} 0 \\ -\frac{1}{3} \end{bmatrix} \end{aligned}$$

so

$$P_{\mathcal{C} \leftarrow \mathcal{B}} = \begin{bmatrix} [\mathbf{b}_1]_{\mathcal{C}} & [\mathbf{b}_2]_{\mathcal{C}} \end{bmatrix} = \begin{bmatrix} -2 & 0 \\ -1 & -\frac{1}{3} \end{bmatrix}.$$

Proposition 6.3.3. Given two different bases for \mathbb{R}^n , \mathcal{B} and \mathcal{C} , the following are true

- $\left(P_{\mathcal{C} \leftarrow \mathcal{B}} \right)^{-1} = P_{\mathcal{B} \leftarrow \mathcal{C}}$, and
- $\left(P_{\mathcal{C} \leftarrow \mathcal{E}} \right) \left(P_{\mathcal{E} \leftarrow \mathcal{B}} \right) = P_{\mathcal{C} \leftarrow \mathcal{B}}$,

where \mathcal{E} is the standard basis.

Remark. It is extremely fast to find change-of-basis matrices to the standard basis, so these two facts above make it very quick to find a change-of-basis matrix between arbitrary bases.

Exercise 6.3.1. With the same setup as in the previous example, compute $P_{\mathcal{E} \leftarrow \mathcal{B}}$ and $P_{\mathcal{E} \leftarrow \mathcal{C}}$, then verify that $\left(P_{\mathcal{E} \leftarrow \mathcal{C}} \right)^{-1} \left(P_{\mathcal{E} \leftarrow \mathcal{B}} \right) = P_{\mathcal{C} \leftarrow \mathcal{B}}$.

3.6 Introduction to Linear Transformations

Definition. A *transformation* (aka *function* or *map*) is a function T with domain \mathbb{R}^n and codomain \mathbb{R}^m , written

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

T is a *linear transformation* if

1. $T(\mathbf{u} + \mathbf{v}) = T(\mathbf{u}) + T(\mathbf{v})$ for all $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$, and
2. $T(k\mathbf{v}) = kT(\mathbf{v})$ for all scalars $k \in \mathbb{R}$ and vectors $\mathbf{v} \in \mathbb{R}^m$.

Example 3.6.1 (Identity transformation). $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ given by $T(\mathbf{v}) = \mathbf{v}$ is a linear transformation.

Example 3.6.2 (Trivial transformation). $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ given by $T(\mathbf{v}) = \mathbf{0}$ is a linear transformation.

Example 3.6.3. Consider $T : \mathbb{R} \rightarrow \mathbb{R}$ given by $T(x) = mx + b$, where b is some constant. Then

1. $T(x_1 + x_2) = mx_1 + mx_2 + b = (mx_1 + b) + (mx_2 + b) = T(x_1) + T(x_2)$ (precisely when $b = 0$), and
2. $T(kx) = mkx + b = k(mx + b) = kT(x)$ (precisely when $b = 0$)

so T is a linear transformation only when $b = 0$.

Example 3.6.4. Consider the map

$$\begin{aligned} T_A : \mathbb{R}^2 &\rightarrow \mathbb{R}^3 \\ T_A(\mathbf{x}) &= A\mathbf{x} \end{aligned}$$

where $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix}$. It follows from matrix arithmetic rules that This is a linear transformation because one can easily check that

1. $T_A(\mathbf{x}_1 + \mathbf{x}_2) = A(\mathbf{x}_1 + \mathbf{x}_2) = A\mathbf{x}_1 + A\mathbf{x}_2 = T_A(\mathbf{x}_1) + T_A(\mathbf{x}_2)$
2. $T_A(k\mathbf{x}) = A(k\mathbf{x}) = kA\mathbf{x} = kT_A(\mathbf{x})$

The above argument actually shows the following more general result

Theorem 3.6.5. If $A \in \mathbb{R}^{m \times n}$, then the transformation

$$\begin{aligned} T_A : \mathbb{R}^n &\rightarrow \mathbb{R}^m \\ T_A(\mathbf{x}) &= A\mathbf{x} \end{aligned}$$

is a linear transformation.

Notice that, explicitly, we could write

$$T_A\left(\begin{bmatrix} x \\ y \end{bmatrix}\right) = \begin{bmatrix} x + 2y \\ 3x + 4y \\ 5x + 6y \end{bmatrix} = x \begin{bmatrix} 1 \\ 3 \\ 5 \end{bmatrix} + y \begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix}.$$

The range of T_A then is just

$$\begin{aligned} \text{Range}(T_A) &= \left\{ T_A \left(\begin{bmatrix} x \\ y \end{bmatrix} \right) : \text{where } x, y \in \mathbb{R} \right\} \\ &= \left\{ x \begin{bmatrix} 1 \\ 3 \\ 5 \end{bmatrix} + y \begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix} : \text{where } x, y \in \mathbb{R} \right\} = \text{Col}(A) \end{aligned}$$

As it turns out, we can write every linear transformation as multiplication by a matrix.


Theorem 3.6.6. *Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear transformation. Then we can write $T(\mathbf{x}) = A\mathbf{x}$ where A is the $m \times n$ matrix whose i^{th} column is the column vector $T(\mathbf{e}_i)$, i.e.*

$$A = \left[\begin{array}{c|c|c} T(\mathbf{e}_1) & \cdots & T(\mathbf{e}_m) \end{array} \right].$$

Definition. The matrix in the above theorem is called the *standard matrix of T* . We may sometimes write $[T]$ to denote the standard matrix of T .

3.6.1 Types of Linear Transformations of \mathbb{R}^2

Example 3.6.7. A **Reflection about the x -axis**, R_x , is a linear transformation of \mathbb{R}^2 .



Explicitly, it sends a points (x, y) to $(x, -y)$, hence the transformation is given by


$$R_x \left(\begin{bmatrix} x \\ y \end{bmatrix} \right) = \begin{bmatrix} x \\ -y \end{bmatrix}$$

The standard matrix for R_x is thusly given by

$$[R_x] = \left[R_x \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \quad R_x \left(\begin{bmatrix} 0 \\ 1 \end{bmatrix} \right) \right] = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

Exercise 3.6.1. Show that reflection about the y -axis is also a linear transformation and find the standard matrix for this.

Example 3.6.8. Rotation by an angle θ about the origin, R_θ , is a linear transformation of \mathbb{R}^2 .




For this one, we'll first find the standard transformation matrix. Note that $\mathbf{e}_1 = [\cos(0), \sin(0)]^T$ and $\mathbf{e}_2 = [\cos(\frac{\pi}{2}), \sin(\frac{\pi}{2})]^T$. So rotation by an angle θ should add θ to the angle arguments of sine and cosine, i.e.

$$\begin{aligned} [R_\theta] &= \left[R_\theta \left(\begin{bmatrix} \cos(0) \\ \sin(0) \end{bmatrix} \right) \quad R_\theta \left(\begin{bmatrix} \cos(\frac{\pi}{2}) \\ \sin(\frac{\pi}{2}) \end{bmatrix} \right) \right] \\ &= \begin{bmatrix} \cos(0 + \theta) & \cos(\frac{\pi}{2} + \theta) \\ \sin(0 + \theta) & \sin(\frac{\pi}{2} + \theta) \end{bmatrix} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \end{aligned}$$

hence the linear transformation is given by

$$R_\theta \left(\begin{bmatrix} x \\ y \end{bmatrix} \right) = \begin{bmatrix} x \cos(\theta) - y \sin(\theta) \\ x \sin(\theta) + y \cos(\theta) \end{bmatrix}$$

Example 3.6.9. A **dilation** (with **dilation factor** k) is a transformation D_k that expands out from the origin by a factor of k .




Explicitly, it sends a point (x, y) to a point (kx, ky) so for vectors,

$$D_k \left(\begin{bmatrix} x \\ y \end{bmatrix} \right) = \begin{bmatrix} kx \\ ky \end{bmatrix}$$

The standard matrix for D_k is given by

$$[D_k] = \left[D_k \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \quad D_k \left(\begin{bmatrix} 0 \\ 1 \end{bmatrix} \right) \right] = \begin{bmatrix} k & 0 \\ 0 & k \end{bmatrix}.$$

Example 3.6.10. A **horizontal shear** (with **shear factor** m), S_m , is a transformation that slides the top edge of the unit square m units to the right (making a parallelogram).



In particular, it sends (x, y) to the point $(x + my, y)$,

$$S_m \left(\begin{bmatrix} x \\ y \end{bmatrix} \right) = \begin{bmatrix} x + my \\ y \end{bmatrix}$$

The standard matrix for S_m is given by

$$[S_m] = \left[S_m \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \quad S_m \left(\begin{bmatrix} 0 \\ 1 \end{bmatrix} \right) \right] = \begin{bmatrix} 1 & m \\ 0 & 1 \end{bmatrix}$$

Theorem 3.6.11. If $T_A : \mathbb{R}^m \rightarrow \mathbb{R}^n$ and $T_B : \mathbb{R}^n \rightarrow \mathbb{R}^p$ are two linear transformations (with standard matrices A and B), then

$$T_B \circ T_A = T_{BA}.$$

Definition. A transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is *one-to-one* (or *injective*) if, for every vector $\mathbf{v} \in \mathbb{R}^n$, there is a unique $\mathbf{w} \in \mathbb{R}^m$ for which $T(\mathbf{v}) = \mathbf{w}$. T is *onto* (or *surjective*) if, for every $\mathbf{w} \in \mathbb{R}^m$, there is at least one $\mathbf{v} \in \mathbb{R}^n$ for which $T(\mathbf{v}) = \mathbf{w}$. The *kernel* of T is the collection of vectors \mathbf{x} in \mathbb{R}^n for which $T(\mathbf{x}) = \mathbf{0}$.

Remark. Maybe the right way to think about it is this: “one-to-one” means that the range of T is a copy of \mathbb{R}^n living inside of \mathbb{R}^m , and “onto” means that the range of T is all of \mathbb{R}^m .

Theorem 3.6.12. Let $T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$.

1. T_A is one-to-one if and only if $\text{Rank}(A) = n$.
2. T_A is onto if and only if $\text{Rank}(A) = m$.
3. $\ker(T_A) = \text{Null}(A)$.

Definition. Two transformations $S, T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are *inverses* if, for every $\mathbf{v} \in \mathbb{R}^n$,

$$S \circ T(\mathbf{v}) = T \circ S(\mathbf{v}) = \mathbf{v}$$

Theorem 3.6.13. If $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ has an inverse, then the standard matrix $[T]$ is invertible and $[T^{-1}] = [T]^{-1}$


Theorem 3.6.14. A is invertible if and only if $T_A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is both one-to-one and onto.

3.7 Applications

3.7.1 Markov Chains

A Markov chain is just a process by which one models probabilistic scenarios

Example 3.7.1. Researchers have found that Democrat (D) voters are 70% likely to continue voting Democrat in the next election, 10% likely to vote Republican in the next election, and 20% likely to vote Independent in the next election. Similar data was compiled for Republican (R) and Independent (I) voters, and can be modeled in the following graph:



If there are D_0 Democrat voters, R_0 Republican voters, and I_0 Independent voters in this current election cycle, how many of each will there be for the next election cycle? We can write

$$\begin{aligned}D_1 &= 0.7(D_0) + 0.1(R_0) + 0.3(I_0) \\R_1 &= 0.1(D_0) + 0.8(R_0) + 0.2(I_0) \\I_1 &= 0.2(D_0) + 0.1(R_0) + 0.5(I_0)\end{aligned}$$

or, as a matrix/vectors

$$\mathbf{x}_1 = \begin{bmatrix} D_1 \\ R_1 \\ I_1 \end{bmatrix} = \begin{bmatrix} 0.7 & 0.1 & 0.3 \\ 0.1 & 0.8 & 0.2 \\ 0.2 & 0.1 & 0.5 \end{bmatrix} \begin{bmatrix} D_0 \\ R_0 \\ I_0 \end{bmatrix} = A\mathbf{x}_0$$

It follows that, after k elections cycles, $\mathbf{x}_k = A^k\mathbf{x}_0$.

Definition. The vector \mathbf{x}_0 is known as the *initial state vector*, A is known as the *transition matrix*, and the entire process above is called a *Markov chain* (with 3 states). If the transition matrix has columns with positive entries that all sum to 1, it is called a *stochastic matrix*, the columns are called *probability vectors*, and the (i, j) -entry represents the probability that object in state i (e.g. voting preference) switches to state j .

Below are a few computations for different k -values, assuming $D_0 = 1000$, $R_0 = 800$ and $I_0 = 300$:

$$\mathbf{x}_1 = \begin{bmatrix} 870 \\ 800 \\ 430 \end{bmatrix}, \quad \mathbf{x}_{20} \approx \begin{bmatrix} 764 \\ 859 \\ 477 \end{bmatrix}, \quad \mathbf{x}_{100} \approx \begin{bmatrix} 764 \\ 859 \\ 477 \end{bmatrix}, \quad \mathbf{x}_{1000} \approx \begin{bmatrix} 764 \\ 859 \\ 477 \end{bmatrix}$$

What we notice is that the vector $\mathbf{x}_k = A^k \mathbf{x}_0$ seems to stop changing as $k \rightarrow \infty$.

If you're interested in playing around with this yourself, say with different initial conditions or a different number of steps in the Markov chain process, you can use the MATLAB code below:

```
transMat = [ 0.7 0.3 0.1 ; 0.1 0.2 0.8 ; 0.2 0.5 0.1 ]; %transition matrix
x0 = [ 1000 ; 800 ; 300 ]; %initial state vector [D0,R0,I0]
maxLoop = 100; %number of iterations in Markov chain

for k=1:maxLoop
    transpose(transMat^k*x0) %outputs xk = [Dk,Rk,Ik], kth step in Markov chain process
end
```

What's fascinating is that this limit doesn't change as long as $D_0 + R_0 + I_0 = 2100$ (really, try it).

Definition. In a Markov chain process with transition matrix A , a vector \mathbf{x} for which $A\mathbf{x} = \mathbf{x}$ is called a *steady state vector*.

Remark. Every Markov chain has a steady state vector, but unless all of the entries of the transition matrix are positive, it may not be the limit $\lim_{k \rightarrow \infty} A^k \mathbf{x}_0$ (such a limit may not even exist).

Even more interesting - in the last example, not only did the limit not change under mild assumptions about the initial state vector, it turns out that no matter what nonzero vector we choose for \mathbf{x}_0 , it will always be the case that

$$\lim_{k \rightarrow \infty} A^k \mathbf{x}_0 \approx \lambda \begin{bmatrix} 764 \\ 859 \\ 477 \end{bmatrix}$$

(for some real number λ).

4.1 Introduction to Eigenvalues and Eigenvectors

Definition. Let A be an $n \times n$ matrix. A scalar λ is an *eigenvalue* of A if there is a nonzero vector $\mathbf{v} \in \mathbb{R}^n$ so that $A\mathbf{v} = \lambda\mathbf{v}$. Such a vector is called an *eigenvector* of A corresponding to λ .

Remark. The prefix *eigen-* is not a name, but is derived from German and means “special” or “characteristic.”

Example 4.1.1. In the Markov chain example, the steady state vector was precisely the vector \mathbf{x} satisfying $A\mathbf{x} = \mathbf{x}$, so \mathbf{x} was an eigenvector of A corresponding to the eigenvalue 1.

Example 4.1.2. $A = \begin{bmatrix} -6 & 3 \\ 4 & 5 \end{bmatrix}$ has eigenvector $\begin{bmatrix} -3 \\ 1 \end{bmatrix}$. Find the corresponding eigenvalue.

$$\begin{bmatrix} -6 & 3 \\ 4 & 5 \end{bmatrix} \begin{bmatrix} -3 \\ 1 \end{bmatrix} = \begin{bmatrix} 21 \\ -7 \end{bmatrix} = \lambda \begin{bmatrix} -3 \\ 1 \end{bmatrix}$$

hence the corresponding eigenvalue is $\lambda = -7$.

Example 4.1.3. Show that 6 is another eigenvalue of $A = \begin{bmatrix} -6 & 3 \\ 4 & 5 \end{bmatrix}$ and find its corresponding eigenvector.

We need to find a vector \mathbf{v} such that

$$A\mathbf{v} = 6\mathbf{v} \quad \implies \quad A\mathbf{v} - 6\mathbf{v} = \mathbf{0} \quad \implies \quad (A - 6I)\mathbf{v} = \mathbf{0}$$

so really we need to compute $\text{Null}(A - 6I)$.

$$[A - 6I \mid \mathbf{0}] = \left[\begin{array}{cc|c} -12 & 3 & 0 \\ 4 & -1 & 0 \end{array} \right] \xrightarrow{\text{RREF}} \left[\begin{array}{cc|c} 4 & -1 & 0 \\ 0 & 0 & 0 \end{array} \right]$$

So a vector \mathbf{v} is in $\text{Null}(A - 6I)$ if it has the form $\begin{bmatrix} t \\ 4t \end{bmatrix}$. As such, any nonzero vector of this form is an eigenvector of A corresponding to 6.

Definition. Let A be an $n \times n$ matrix and λ an eigenvalue with corresponding eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_k$. The *eigenspace* corresponding to λ is

$$E_\lambda := \text{Span}(\mathbf{v}_1, \dots, \mathbf{v}_k).$$


Remark. It may at first seem surprising that any linear combination of the above \mathbf{v}_i 's is still an eigenvector for λ , but it is a straightforward computation to see that it is true:

$$\begin{aligned} A(c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k) &= c_1A\mathbf{v}_1 + \dots + c_kA\mathbf{v}_k \\ &= c_1\lambda\mathbf{v}_1 + \dots + c_k\lambda\mathbf{v}_k \\ &= \lambda(c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k). \end{aligned}$$


Example 4.1.4. In the previous example, $E_6 = \text{Span}\left(\begin{bmatrix} 1 \\ 4 \end{bmatrix}\right)$

Geometrically, eigenvectors tell us about lines that are preserved by the linear transformation (since all that happens is a vector in that line is scaled).

Example 4.1.5. Using A as before, notice that the vectors $[-3, 1]^T$ and $[1, 4]^T$ do not change direction after a transformation, but are merely scaled.



Before applying transformation A .



After applying transformation A .

How do I compute eigenvalues for a given matrix A ?

Definition. If $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, the *determinant* of A is

$$\det(A) = ad - bc$$

We'll define the determinants for all $n \times n$ matrices in the next section.

An eigenvector for the eigenvalue λ is a vector in $\text{Null}(A - \lambda I)$. By the Fundamental Theorem of Invertible matrices tells us that $\text{nullity}(A - \lambda I) \neq 0$ if and only if $A - \lambda I$ is *not* invertible. Moreover, from a previous theorem, a 2×2 matrix is *not* invertible if and only if $\det(A - \lambda I) = 0$. Since $\det(A - \lambda I)$ is a polynomial in λ , this means that eigenvalues are precisely the zeroes of this polynomial.

Example 4.1.6. Find the eigenvalues for $A = \begin{bmatrix} -6 & 3 \\ 4 & 5 \end{bmatrix}$.

$$\begin{aligned} 0 &= \det(A - \lambda I) \\ &= \det\left(\begin{bmatrix} -6 - \lambda & 3 \\ 4 & 5 - \lambda \end{bmatrix}\right) \\ &= (-6 - \lambda)(5 - \lambda) - 12 \\ &= \lambda^2 + \lambda - 42. \end{aligned}$$

Using your favorite method of solving for the zeroes of this polynomial, we exactly see that its zeroes are $\lambda = -6, 7$, which are precisely the eigenvalues we expected to get from the previous examples.

Example 4.1.7. Find the eigenvalues and corresponding eigenspaces for the scalar matrix $A = \begin{bmatrix} 7 & 0 \\ 0 & 7 \end{bmatrix}$.

$$0 = \det(A - \lambda I) = \det\left(\begin{bmatrix} 7 - \lambda & 0 \\ 0 & 7 - \lambda \end{bmatrix}\right) = (7 - \lambda)(7 - \lambda).$$

This matrix has a single eigenvalue $\lambda = 7$ and

$$E_7 = \text{Null}(A - 7I) = \text{Null}\left(\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}\right) = \text{Span}\left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}\right).$$

Exercise 4.1.1. Find the eigenvalues and corresponding eigenspaces for the horizontal shear matrix $A = \begin{bmatrix} 1 & 7 \\ 0 & 1 \end{bmatrix}$.

APPENDIX C Complex Numbers

Some matrices may have complex eigenvalues (what does that mean geometrically? hmmm...), so below is an list of important properties of complex numbers.

Let $i = \sqrt{-1}$, the so-called *imaginary unity*. A *complex number* is a number $z = a + bi$ where a and b are real numbers. a is called the *real part* and b is called the *imaginary part* of z . The *conjugate of z* is the complex number $\bar{z} = a - bi$.

Addition of two complex numbers z_1 and z_2 is done via by adding the real and imaginary parts separately:

$$z_1 + z_2 = (a_1 + b_1i) + (a_2 + b_2i) = (a_1 + a_2) + i(b_1 + b_2).$$

Multiplication of complex numbers follows the usual distributive law:

$$\begin{aligned} z_1 z_2 &= (a_1 + b_1i)(a_2 + b_2i) \\ &= a_1 a_2 + a_1 b_2 i + a_2 b_1 i + b_1 b_2 i^2 \\ &= (a_1 a_2 - b_1 b_2) + (a_1 b_2 + a_2 b_1) i. \end{aligned}$$

Noting that $z\bar{z}$ is a real number for any z , *division* of complex numbers is done by multiplying by the conjugate and scaling by $1/z\bar{z}$:

$$\frac{z_1}{z_2} = \frac{z_1 \bar{z}_2}{z_2 \bar{z}_2} = \frac{1}{z_2 \bar{z}_2} (z_1 \bar{z}_2) = \frac{1}{a_2^2 + b_2^2} ((a_1 a_2 + b_1 b_2) + (-a_1 b_2 + a_2 b_1) i)$$

Exercise. To each complex number $z = a + bi$ we can associate the matrix $[z] = \begin{bmatrix} a & -b \\ b & a \end{bmatrix}$. Show that the product $z_1 z_2$ can be derived from multiplying the associated matrices $[z_1][z_2]$.

Matrices with complex entries

Although we won't see them in this course, matrices with complex entries appear all the time (they may even be the central objects in your instructor's dissertation) and have uses in physics and electrical engineering.

If $A = [a_{ij}]$ is a matrix with complex entries, the *conjugate* is

$$\bar{A} = \begin{bmatrix} \overline{a_{11}} & \cdots & \overline{a_{1n}} \\ \vdots & & \vdots \\ \overline{a_{m1}} & \cdots & \overline{a_{mn}} \end{bmatrix}.$$

and the *conjugate transpose* (sometimes called the *adjoint* or *Hermitian transpose*) is obtained by first conjugating the matrix and then transposing it

$$\bar{A}^T = \begin{bmatrix} \overline{a_{11}} & \cdots & \overline{a_{m1}} \\ \vdots & & \vdots \\ \overline{a_{1n}} & \cdots & \overline{a_{nn}} \end{bmatrix}.$$

The conjugate-transpose is also sometimes denoted A^\dagger or A^* , depending on whether or not a physicist wrote the paper.

4.2 Determinants

Definition. The *determinant* (denoted \det) is a function whose inputs are all of the columns of an $n \times n$ real matrix A , whose output is a single real number, and which satisfies the following properties:

1. $\det I_n = 1$.
2. \det is a *multilinear* function (it is a linear transformation in each input).
3. \det is an *alternating* function (its sign changes whenever two columns are swapped).


Remark.

- The determinant is actually the unique such function with these properties.
- Some write $|A|$ to mean $\det A$.


How do we compute the determinant of a given matrix?

4.2.1 Determinant of a 2×2 Matrix

The first property of the determinant is that $\det(I_2) = 1$, so we'd like to find a natural place where the value 1 naturally occurs relative to the columns of the identity matrix. Plotting the vectors \mathbf{e}_1 and \mathbf{e}_2 , we see that they naturally form a parallelogram (in particular, a square) of area 1, so maybe the natural choice is that the determinant should be the (signed) area of this parallelogram.



More generally, let $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$. For purposes of the picture, we will assume that $a > c > 0$ and $d > b > 0$.



The (signed) area of the parallelogram is

$$\text{area} = (a+c)(b+d) - 2bc - cd - ab = ad - bc.$$

Proposition 4.2.1. $\det \begin{bmatrix} a & b \\ c & d \end{bmatrix} = ad - bc.$

Proof. By construction we know that this satisfies the property that $\det I_2 = 1$, so we only need to check the multilinearity and alternating conditions.

Multilinearity: it is a straightforward computation to show that

$$\det \begin{bmatrix} (a_1 + a_2) & b \\ (c_1 + c_2) & d \end{bmatrix} = \det \begin{bmatrix} a_1 & b \\ c_1 & d \end{bmatrix} + \det \begin{bmatrix} a_2 & b \\ c_2 & d \end{bmatrix}$$

and

$$\det \begin{bmatrix} a & (b_1 + b_2) \\ c & (d_1 + d_2) \end{bmatrix} = \det \begin{bmatrix} a & b_1 \\ c & d_1 \end{bmatrix} + \det \begin{bmatrix} a & b_2 \\ c & d_2 \end{bmatrix}$$

and that for any scalar k ,

$$\det \begin{bmatrix} ka & b \\ kc & d \end{bmatrix} = k \det \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \det \begin{bmatrix} a & kb \\ c & kd \end{bmatrix}.$$

Alternating: it is a straightforward computation to show that

$$\det \begin{bmatrix} a & b \\ c & d \end{bmatrix} = -\det \begin{bmatrix} b & a \\ d & c \end{bmatrix}.$$

Since the determinant is the unique function with these properties, this must mean that our definition of determinant is correct. \square

Theorem 4.2.2. For any $n \times n$ matrix A , $\det A$ is the signed volume of the parallelepiped formed by the columns of A .

4.2.2 Determinant of a $n \times n$ Matrix

We already learned the determinant of a 2×2 matrix. So how about for $n \times n$ matrices where $n > 2$? Computing the parallelepiped volume is a bit of a chore, but thankfully we can do it iteratively.

Definition. For an $n \times n$ matrix A , the (i, j) -minor of A , denoted $M_{i,j}$, is the determinant of the *submatrix* formed by removing row i and column j from A .

Example 4.2.3. For $A = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}$, the $(2, 1)$ -minor is

$$M_{2,1} = \det \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} = \det \begin{bmatrix} a_{12} & a_{13} \\ a_{32} & a_{33} \end{bmatrix} = a_{12}a_{33} - a_{32}a_{13}.$$

Remark. Your book uses the notation $\det A_{ij}$ to denote the (i, j) -minor, but I think it just makes things more confusing.

Theorem 4.2.4. The determinant of an $n \times n$ matrix A can be computed along the i^{th} row as the sum

$$\det A = \sum_{j=1}^n (-1)^{i+j} a_{ij} M_{i,j}.$$

or along the j^{th} column as the sum

$$\det A = \sum_{i=1}^n (-1)^{i+j} a_{ij} M_{i,j}.$$

It is an exercise to the reader to verify that the above sums satisfy the definition of the determinant and that they agree no matter along which row or column they are computed.

Remark. The quantity $(-1)^{i+j} M_{i,j}$ is sometimes called the (i, j) -cofactor and the above sums are cofactor expansions.

Example 4.2.5. Compute $\det A$ given $A = \begin{bmatrix} 3 & 1 & 4 \\ -1 & 5 & -9 \\ 2 & 6 & -5 \end{bmatrix}$ along the first row.

$$\begin{aligned} \det A &= a_{11}M_{1,1} - a_{12}M_{1,2} + a_{13}M_{1,3} \\ &= 3 \det \begin{bmatrix} 3 & 1 & 4 \\ -1 & 5 & -9 \\ 2 & 6 & -5 \end{bmatrix} - 1 \det \begin{bmatrix} 3 & 1 & 4 \\ -1 & 5 & -9 \\ 2 & 6 & -5 \end{bmatrix} + 4 \det \begin{bmatrix} 3 & 1 & 4 \\ -1 & 5 & -9 \\ 2 & 6 & -5 \end{bmatrix} \\ &= 3 \det \begin{bmatrix} 5 & -9 \\ 6 & -5 \end{bmatrix} - \det \begin{bmatrix} -1 & -9 \\ 2 & -5 \end{bmatrix} + 4 \det \begin{bmatrix} -1 & 5 \\ 2 & 6 \end{bmatrix} \\ &= 3(29) - (23) + 4(-16) = 0. \end{aligned}$$

Having a zero determinant tells us that the corresponding parallelepiped has no volume, which means all three column vectors must live in the same plane (or even same line). So a determinant of zero is enough to tell us that the columns are linearly dependent!

Proposition 4.2.6. *For an $n \times n$ matrix, $\det(A) = 0$ if and only if the columns of A are linearly dependent.*

This means we can add onto the fundamental theorem of invertible matrices:

Theorem 4.2.7. *A is invertible if and only if $\det(A) \neq 0$.*

4.2.3 Properties of Determinants

Definition. A square matrix $A = [a_{ij}]$ is *upper triangular* if $a_{ij} = 0$ whenever $i > j$, and is *lower triangular* if $a_{ij} = 0$ whenever $i < j$.

Example 4.2.8. Consider the following matrices

$$A = \begin{bmatrix} 1 & 2 & 3 \\ 0 & 4 & 5 \\ 0 & 0 & 6 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 3 & 0 \\ 4 & 5 & 6 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{bmatrix}.$$

A is upper triangular, B is lower triangular, and C (a diagonal matrix) is both upper and lower triangular.

Theorem 4.2.9 (Poole Theorem 4.2). *If A is an $n \times n$ triangular matrix, then $\det(A) = a_{11}a_{22} \cdots a_{nn}$, the product of the numbers along the diagonal.*

Proof. Perform the cofactor expansion along the first column (in the upper triangular case) or along the first row (in the lower triangular case). \square

Example 4.2.10. Let's look at the matrices in Example 4.2.8. By computing the determinant of A along the first column, we have

$$\begin{aligned} \det A &= 1M_{1,1} - 0M_{2,1} + 0M_{3,1} \\ &= 1 \det \begin{bmatrix} 4 & 5 \\ 0 & 6 \end{bmatrix} = 1(4)(6). \end{aligned}$$

Now computing the determinant of B along the first row, we have

$$\begin{aligned} \det B &= 1M_{1,1} - 0M_{1,2} + 0M_{1,3} \\ &= 1 \det \begin{bmatrix} 3 & 0 \\ 5 & 6 \end{bmatrix} = 1(3)(6). \end{aligned}$$

Finally, computing the determinant of C along the first row (although the first column is perfectly fine as well),

$$\begin{aligned} \det C &= 1M_{1,1} - 0M_{1,2} + 0M_{1,3} \\ &= 1 \det \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix} = 1(2)(3). \end{aligned}$$

Theorem 4.2.11 (Poole Theorems 4.7 - 4.10). *If A and B are $n \times n$ matrices and k is a scalar, then*

1. $\det(AB) = (\det A)(\det B)$,
2. $\det(kA) = k^n(\det A)$,
3. $\det A = \det A^T$,
4. *and if A is invertible, $\det A^{-1} = \frac{1}{\det A}$.*

Example 4.2.12. Verify each part of the theorem above with the following:

$$A = \begin{bmatrix} 3 & -1 \\ 8 & -2 \end{bmatrix}, \quad B = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}, \quad \text{and} \quad k = 5.$$

1. $\det A = -6 + 8 = 2$, $\det B = 4 - 1 = 3$, and

$$\det(AB) = \det \begin{bmatrix} 5 & 1 \\ 14 & 4 \end{bmatrix} = 6 = (2)(3) = (\det A)(\det B).$$

2. $\det kA = \det \begin{bmatrix} 15 & -5 \\ 40 & -10 \end{bmatrix} = 50 = 25(2) = k^2 \det A.$

3. $\det A^T = \det \begin{bmatrix} 3 & 8 \\ -1 & -2 \end{bmatrix} = -6 + 8 = 2 = \det A$

4. $\det A^{-1} = \det \begin{bmatrix} -1 & \frac{1}{2} \\ -4 & \frac{3}{2} \end{bmatrix} = -\frac{3}{2} + 2 = \frac{1}{2} = \frac{1}{\det A}$

4.3 Eigenvalues and Eigenvectors of $n \times n$ Matrices

We've already seen eigenvalues and eigenvectors for 2×2 matrices, but now that we have defined determinants for $n \times n$ matrices, we'll extend these definitions accordingly.

Definition. If A is a square matrix, then $\det(A - \lambda I)$ is a polynomial with indeterminate λ and is called the *characteristic polynomial* of A . The *eigenvalues* of A are precisely the roots of the characteristic polynomial. For each eigenvalue λ , the corresponding *eigenspace* is $E_\lambda = \text{Null}(A - \lambda I)$ and the nonzero vectors in E_λ are *eigenvectors*.

Remark. Non-square matrices do not have eigenvalues/eigenvectors, because if $v \in \mathbb{R}^n$ is an eigenvector for A , then $Av = \lambda v$ implies that $v \in \mathbb{R}^n$ as well, hence A is $n \times n$. Non-square matrices have something called *singular values* which, in some sense, play the role of eigenvalues, but this is outside of the scope of this course.

Example 4.3.1. Find the eigenvalues and eigenvectors for $A = \begin{bmatrix} 2 & 12 & 10 \\ 0 & -4 & -4 \\ 1 & 2 & 1 \end{bmatrix}$.

We first compute $\det(A - \lambda I)$ via cofactor expansion along the first column.

$$\begin{aligned} p(\lambda) = \det(A - \lambda I) &= \det \begin{bmatrix} 2 - \lambda & 12 & 10 \\ 0 & -4 - \lambda & -4 \\ 1 & 2 & 1 - \lambda \end{bmatrix} \\ &= (2 - \lambda) \det \begin{bmatrix} -4 - \lambda & -4 \\ 2 & 1 - \lambda \end{bmatrix} + 1 \det \begin{bmatrix} 12 & 10 \\ -4 - \lambda & -4 \end{bmatrix} \\ &= (2 - \lambda) ((-4 - \lambda)(1 - \lambda) + 8) + 1(-48 - 10(-4 - \lambda)) \\ &= -(\lambda^3 + \lambda^2 - 12\lambda) \\ &= -\lambda(\lambda + 4)(\lambda - 3) \end{aligned}$$

The characteristic polynomial factors nicely and the eigenvalues are $-4, 0, 3$. The corresponding eigenspaces are

$$\begin{aligned} E_{-4} &= \text{Null}(A + 4I) = \text{Null} \left(\begin{bmatrix} 6 & 12 & 10 \\ 0 & 0 & -4 \\ 1 & 2 & 5 \end{bmatrix} \right) = \text{Span} \left(\begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix} \right), \\ E_0 &= \text{Null}(A - 0I) = \text{Null} \left(\begin{bmatrix} 2 & 12 & 10 \\ 0 & -4 & -4 \\ 1 & 2 & 1 \end{bmatrix} \right) = \text{Span} \left(\begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix} \right), \\ E_3 &= \text{Null}(A - 3I) = \text{Null} \left(\begin{bmatrix} -1 & 12 & 10 \\ 0 & -7 & -4 \\ 1 & 2 & -2 \end{bmatrix} \right) = \text{Span} \left(\begin{bmatrix} 22 \\ -4 \\ 7 \end{bmatrix} \right). \end{aligned}$$

Example 4.3.2. Consider the matrix $A = \begin{bmatrix} 3 & -2 & -4 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{bmatrix}$. The characteristic polynomial for A is

$$p(\lambda) = (3 - \lambda)(1 - \lambda)^2$$

hence the eigenvalues are 1, 3. The corresponding eigenspaces are

$$E_3 = \text{Span} \left(\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right) \quad \text{and} \quad E_1 = \text{Span} \left(\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \right).$$

Notice that the above 3×3 matrix had only two distinct eigenvalues, but if we counted multiplicity as well ($(1 - \lambda)$ appears with multiplicity 2 in the characteristic polynomial), then in fact we have exactly 3 eigenvalues. It may also be interesting to notice that $\dim(E_1) = \dim(E_3)$.

Definition. The *algebraic multiplicity* of an eigenvalue λ is the multiplicity as a root of the characteristic polynomial, and the *geometric multiplicity* is the dimension of the eigenspace E_λ .

These two different notions of multiplicity will be important in the next section. We'll note that if A is an $n \times n$ matrix, then the sum of all of the algebraic multiplicities will always be n .

Theorem 4.3.3 (Poole - Theorem 4.15). *If A is a triangular (or diagonal) matrix, then the eigenvalues are precisely the entries appearing along the diagonal.*

Proof. If A is triangular, then by Theorem 4.2.9

$$p(\lambda) = \det(A - \lambda I) = (a_{11} - \lambda)(a_{22} - \lambda) \cdots (a_{nn} - \lambda).$$

□

Theorem 4.3.4. *If A is an $n \times n$ square matrix with eigenvalues $\lambda_1, \dots, \lambda_n$ (not necessarily all distinct), then $\det A = \lambda_1 \lambda_2 \cdots \lambda_n$, the product of all of the eigenvalues.*

Proof. The above theorem is easy to see in the case that A is triangular - the product of the roots of a polynomial are precisely the constant term, and the constant term of the characteristic polynomial is exactly $a_{11} a_{22} \cdots a_{nn}$. □

From this it follows that we have yet another test for invertibility:

Theorem 4.3.5. *A square matrix A is invertible if and only if 0 is not an eigenvalue of A .*

4.3.1 Relationship to Matrix Operations

It is natural to ask about the interplay between eigenvalues/eigenvectors and matrix operations like inversion and exponentiation.

Example 4.3.6. Consider $A = \begin{bmatrix} 1 & -2 \\ 1 & 4 \end{bmatrix}$. The characteristic polynomial is

$$p_A(\lambda) = (2 - \lambda)(3 - \lambda)$$

hence the eigenvalues are 2, 3 and the corresponding eigenspaces are

$$E_2 = \text{Span} \left(\begin{bmatrix} -2 \\ 1 \end{bmatrix} \right) \quad \text{and} \quad E_3 = \text{Span} \left(\begin{bmatrix} -1 \\ 1 \end{bmatrix} \right).$$

With A as above, we have that $A^{-1} = \frac{1}{6} \begin{bmatrix} 4 & 2 \\ -1 & 1 \end{bmatrix}$ and the eigenvalues are $\frac{1}{2}, \frac{1}{3}$ – reciprocals of A 's eigenvalues. What's more, notice that

$$A^{-1} \begin{bmatrix} -2 \\ 1 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} -2 \\ 1 \end{bmatrix} \quad \text{and} \quad A^{-1} \begin{bmatrix} -1 \\ 1 \end{bmatrix} = \frac{1}{3} \begin{bmatrix} -1 \\ 1 \end{bmatrix}$$

so the reciprocal eigenvalues of A^{-1} have the same eigenvectors as the eigenvalues of A ! With A as before, we have that $A^2 = \begin{bmatrix} -1 & -10 \\ 5 & 14 \end{bmatrix}$ and the eigenvalues are 4, 9 – squares of A 's eigenvalues. What's more, notice that


$$A^2 \begin{bmatrix} -2 \\ 1 \end{bmatrix} = 2^2 \begin{bmatrix} -2 \\ 1 \end{bmatrix} \quad \text{and} \quad A^2 \begin{bmatrix} -1 \\ 1 \end{bmatrix} = 3^2 \begin{bmatrix} -1 \\ 1 \end{bmatrix}$$

so the squared eigenvalues of A^2 have the same eigenvectors as the eigenvalues of A !


Theorem 4.3.7. *Let A be a square matrix with eigenvalue λ and corresponding eigenvector \mathbf{v} .*

1. *For any positive integer n , λ^n is an eigenvalue of A^n with corresponding eigenvector \mathbf{v} .*
2. *If A is invertible, then $\frac{1}{\lambda}$ is an eigenvalue of A^{-1} with corresponding eigenvector \mathbf{v} .*

The above theorem also makes sense geometrically. Each application of the transformation A stretches its eigenvector \mathbf{v} by a factor of λ :



Similarly, each application of A^{-1} “undoes” the stretching of its eigenvector \mathbf{v} by a factor of λ (i.e., stretches instead by a factor of $\frac{1}{\lambda}$): The above theorem also makes sense geometrically. Each application of A stretches its eigenvector by a factor of λ



4.4 Similarity and Diagonalization

As we saw in the previous section, triangular and diagonal matrices were very nice from a computational standpoint, so it would be nice to convert a matrix into triangular form in a meaningful way. We already know that we can do this with row reduction, but this process does not preserve eigenvalues (any invertible matrix row reduces to the identity, for example), so in this section we will look at another process that does retain the useful eigen-information.

Definition. Two $n \times n$ matrices A and B are called *similar* if there is an invertible $n \times n$ matrix P for which $P^{-1}AP = B$. We sometimes write “ $A \sim B$ ” to mean “ A is similar to B .” We also sometimes refer to the product $P^{-1}AP$ as “conjugation of A by P .”

Remark. Such a P is not unique. For example, $P^{-1}IP = I$ is true for every invertible matrix P .

Example 4.4.1. $A = \begin{bmatrix} 2 & 3 \\ 1 & 4 \end{bmatrix}$ and $B = \begin{bmatrix} 1 & 0 \\ 0 & 5 \end{bmatrix}$ are similar. With $P = \begin{bmatrix} -3 & 1 \\ 1 & 1 \end{bmatrix}$, we have

$$\begin{aligned} P^{-1}AP &= \begin{bmatrix} -3 & 1 \\ 1 & 1 \end{bmatrix}^{-1} \begin{bmatrix} 2 & 3 \\ 1 & 4 \end{bmatrix} \begin{bmatrix} -3 & 1 \\ 1 & 1 \end{bmatrix} \\ &= \frac{1}{-4} \begin{bmatrix} 1 & -1 \\ -1 & -3 \end{bmatrix} \begin{bmatrix} 2 & 3 \\ 1 & 4 \end{bmatrix} \begin{bmatrix} -3 & 1 \\ 1 & 1 \end{bmatrix} \\ &= \begin{bmatrix} 1 & 0 \\ 0 & 5 \end{bmatrix} = B \end{aligned}$$

Example 4.4.2. The matrices $A = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$ and $B = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ are not similar.

If they were, we could find a matrix $P = \begin{bmatrix} x & y \\ z & w \end{bmatrix}$ for which $B = P^{-1}AP$. In this case, we would have

$$\begin{aligned} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} &= \frac{1}{xw - yz} \begin{bmatrix} w & -y \\ -z & x \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} x & y \\ z & w \end{bmatrix} \\ &= \frac{1}{xw - yz} \begin{bmatrix} wx + yz & 2wy \\ -2xz & -wx - yz \end{bmatrix} \end{aligned}$$

and it's impossible that both $(wx + yz) = 1$ and $(-wx - yz) = 1$ (otherwise $wx + yz = -(wx + yz)$, hence $wx + yz = 0$).

4.4.1 Properties of similarity and similar matrices

Theorem 4.4.3 (Poole Theorem 4.21). *Let A, B, C be $n \times n$ matrices.*

- a. reflexive: $A \sim A$.
- b. symmetric: *If $A \sim B$ then $B \sim A$.*
- c. transitive: *If $A \sim B$ and $B \sim C$ then $A \sim C$.*

Each of the following properties are easily verified, say with the matrices from Example 4.4.1.

Theorem 4.4.4 (Poole Theorem 4.22). *Let A and B be similar $n \times n$ matrices. Then*

- a. $\det A = \det B$
- b. A is invertible if and only if B is invertible.
- c. A and B have the same rank.
- d. A and B have the same characteristic polynomial.
- e. A and B have the same eigenvalues.
- f. $A^m \sim B^m$ for any positive integer m .

Partial proof sketch. Let P be a matrix for which $B = P^{-1}AP$.

- a. $\det B = \det(P^{-1}AP) = (\det P^{-1})(\det A)(\det P) = \left(\frac{1}{\det P}\right)(\det A)(\det P) = \det A$
- b. $B^{-1} = (P^{-1}AP)^{-1} = P^{-1}A^{-1}(P^{-1})^{-1} = P^{-1}A^{-1}P$
- c. This follows from the fact that for any invertible matrix P , $\text{Rank}(A) = \text{Rank}(PA) = \text{Rank}(AP)$.
- d. $\det(B - \lambda I) = \det(P^{-1}AP - \lambda P^{-1}IP) = \det(P^{-1}(A - \lambda I)P) = \det(A - \lambda I)$
- e. A and B have the same characteristic polynomials
- f. $B^m = \underbrace{(P^{-1}AP) \cdots (P^{-1}AP)}_m = P^{-1} \underbrace{A \cdots A}_m P = P^{-1}A^mP$

□

To check whether two given matrices A and B are similar requires finding the matrix P satisfying $P^{-1}AP = B$, which as we saw from 4.4.2, could be quite laborious. The above theorem is actually most useful for showing that two matrices are not similar (in fact, no single part of the theorem is enough to deduce that two matrices are similar).

Example 4.4.5. The matrices A and B in Example 4.4.2 are not similar because $\det A = -1$ and $\det B = 1$.

Example 4.4.6. Although matrices $A = \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix}$ and $B = \begin{bmatrix} 1 & 0 \\ 0 & 6 \end{bmatrix}$ have the same rank ($\text{Rank } A = \text{Rank } B = 2$) and determinant ($\det A = \det B = 6$), they are not similar because their eigenvalues and characteristic polynomials are different.

4.4.2 Diagonalization

Definition. A matrix A is *diagonalizable* if it is similar to a diagonal matrix D , i.e. if there is some invertible matrix P so that $P^{-1}AP = D$.

Example 4.4.7. From Example 4.4.1, $A = \begin{bmatrix} 2 & 3 \\ 1 & 4 \end{bmatrix}$ is diagonalizable since it is similar to the diagonal matrix $B = \begin{bmatrix} 1 & 0 \\ 0 & 5 \end{bmatrix}$.

Notice that the characteristic polynomial of A above is

$$\det(A - \lambda I) = (2 - \lambda)(4 - \lambda) - 3 = \lambda^2 - 6\lambda + 5 = (\lambda - 1)(\lambda - 5)$$

and thus B contains A 's eigenvalues along the diagonal. This gives us a clue as to how one can go about finding the matrix P used to conjugate A into a diagonal matrix (if possible).

Theorem 4.4.8 (Poole Theorem 4.23). *Let A be an $n \times n$ matrix. Then A is diagonalizable if and only if A has n linearly independent eigenvectors.*

More precisely, $D = P^{-1}AP$ if and only if the columns of P are the eigenvectors of A and if the (i, i) entry of D is the eigenvalue corresponding to the i^{th} column of P .

I won't sketch the proof, but the core observation is the following:

If $P^{-1}AP = D$, then this rearranges to $AP = PD$. So if \mathbf{p}_i is the i^{th} column of P and λ_i is the (i, i) entry in D , then

$$\begin{aligned}
 AP &= PD \\
 A \begin{bmatrix} \mathbf{p}_1 & \cdots & \mathbf{p}_n \end{bmatrix} &= \begin{bmatrix} \mathbf{p}_1 & \cdots & \mathbf{p}_n \end{bmatrix} \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} \\
 \begin{bmatrix} A\mathbf{p}_1 & & A\mathbf{p}_n \end{bmatrix} &= \begin{bmatrix} \lambda_1\mathbf{p}_1 & & \lambda_n\mathbf{p}_n \end{bmatrix}
 \end{aligned}$$

and so $A\mathbf{p}_i = \lambda_i\mathbf{p}_i$, hence the λ_i are eigenvalues for A with corresponding eigenvectors \mathbf{p}_i .

Example 4.4.9. Let $A = \begin{bmatrix} 2 & 3 \\ 0 & 2 \end{bmatrix}$. The characteristic polynomial for A is

$$\det(A - \lambda I) = \det \begin{bmatrix} 2 - \lambda & 3 \\ 0 & 2 - \lambda \end{bmatrix} = (2 - \lambda)^2$$

so A has a single eigenvalue of 2 with algebraic multiplicity 2. The corresponding eigenspace is

$$E_2 = \text{Null}(A - 2I) = \text{Span} \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right).$$

and so the eigenvalue 2 has geometric multiplicity 1. This means that there are not enough linearly independent eigenvectors to form our invertible matrix P (the one for which $P^{-1}AP$ is a diagonal matrix), hence A is not diagonalizable.

Example 4.4.10. Consider the 3×3 matrix $A = \begin{bmatrix} 4 & -3 & -3 \\ 3 & -2 & -3 \\ -1 & 1 & 2 \end{bmatrix}$. The characteristic polynomial is

$$\det(A - \lambda I) = \det \begin{bmatrix} 4 - \lambda & -3 & -3 \\ 3 & -2 - \lambda & -3 \\ -1 & 1 & 2 - \lambda \end{bmatrix} = -(\lambda - 1)^2(\lambda - 2)$$

and the eigenvalues are 1 and 2 (with algebraic multiplicities 2 and 1, respectively). The corresponding eigenspaces are

$$E_1 = \text{Null}(A - I) = \text{Span} \left(\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \right) \quad E_2 = \text{Null}(A - 2I) = \text{Span} \left(\begin{bmatrix} -3 \\ -3 \\ 1 \end{bmatrix} \right)$$

and so the eigenvalues 1 and 2 have geometric multiplicities 2 and 1 (respectively). It is readily seen that the vectors we used to define E_1 are linearly independent, so the following matrix is invertible:

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

We then diagonalize A :

$$P^{-1}AP = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{bmatrix}.$$

What these examples have highlighted is that a matrix may only fail to be diagonalizable if it has repeated eigenvalues.

Theorem 4.4.11 (Poole Theorem 4.24). *Let A be an $n \times n$ matrix with distinct eigenvalues $\lambda_1, \dots, \lambda_k$. Let \mathcal{B}_i be the basis for E_{λ_i} . The union of the \mathcal{B}_i 's (i.e. the collection of all basis vectors in the \mathcal{B}_i 's) is a linearly independent set.*

Corollary 4.4.12 (Poole Theorem 4.25). *If A is an $n \times n$ matrix with n distinct eigenvalues, then A is diagonalizable.*

What is it about the repeated eigenvalues that causes the failure of diagonalizability of a matrix $A \in \mathbb{R}^{n \times n}$? Well, we need there to be n linearly independent eigenvectors, so we need the geometric multiplicity for each eigenvalue to be as large as possible.

Lemma 4.4.13 (Poole Lemma 4.26). *The geometric multiplicity of an eigenvalue λ is less than or equal to its algebraic multiplicity.*

All of this culminates in the following result:

Theorem 4.4.14 (Diagonalization Theorem). *Let A be an $n \times n$ matrix with distinct eigenvalues $\lambda_1, \dots, \lambda_k$. The following are equivalent:*

a. *A is diagonalizable.*

b. *The union of the basis vectors from each E_{λ_i} is a set of n vectors. In other words, $n = \sum_{i=1}^k \dim(E_{\lambda_i})$.*

c. *For each i , the algebraic multiplicity of λ_i is equal to the geometric multiplicity of λ_i .*

Example 4.4.15. The matrix $A = \begin{bmatrix} 3 & -1 & 2 \\ 3 & -1 & 6 \\ -2 & 2 & -2 \end{bmatrix}$ has characteristic polynomial

$$\det(A - \lambda I) = (2 - \lambda)^2(4 - \lambda).$$

The eigenvalue 2 has geometric multiplicity 2 and the eigenvalue 4 has geometric multiplicity 1. By the Diagonalization Theorem, A is diagonalizable – you can verify that an appropriate conjugating matrix is

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

Example 4.4.16. The matrix $A = \begin{bmatrix} 1 & 0 & 2 \\ 0 & 3 & 4 \\ 0 & 0 & 3 \end{bmatrix}$ has characteristic polynomial

$$\det(A - \lambda I) = (1 - \lambda)(3 - \lambda)^2.$$

Both eigenvalues 1 and 3 have geometric multiplicity 1, so by the Diagonalization Theorem, A is not diagonalizable.

4.4.3 Computational power of diagonal matrices

Notice that for a diagonal matrix $D = \text{diag}(d_1, \dots, d_n)$ and any positive integer k ,

$$D^k = \begin{bmatrix} d_1^k & & \\ & \ddots & \\ & & d_n^k \end{bmatrix}.$$

Moreover, if D is invertible, then

$$D^{-k} = \begin{bmatrix} d_1^{-k} & & \\ & \ddots & \\ & & d_n^{-k} \end{bmatrix} = \begin{bmatrix} \frac{1}{d_1^k} & & \\ & \ddots & \\ & & \frac{1}{d_n^k} \end{bmatrix}.$$

This is instantaneous. For a general $n \times n$ matrix A , computing A^k in the usual way is *extremely* computationally expensive. However, if A is diagonalizable, we can write $P^{-1}AP = D$, hence

$$D^k = (P^{-1}AP)^k = P^{-1}A^kP \quad \implies \quad A^k = PD^kP^{-1}.$$

In this way, computing the k^{th} power of A is only as computationally difficult as diagonalizing A .

Example 4.4.17. Let $A = \begin{bmatrix} 11 & -6 \\ 15 & -8 \end{bmatrix}$. One can readily check that A has eigenvalues 1, 2, hence is diagonalizable (and since the eigenvalues are all nonzero, A is invertible). Through the usual methods, we can obtain

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

whence, for any integer k ,

$$A^k = \begin{bmatrix} 3 & 2 \\ 5 & 3 \end{bmatrix} \begin{bmatrix} 1^k & 0 \\ 0 & 2^k \end{bmatrix} \begin{bmatrix} 3 & 2 \\ 5 & 3 \end{bmatrix}^{-1} = \begin{bmatrix} 3 & 2 \\ 5 & 3 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 2^k \end{bmatrix} \begin{bmatrix} -3 & 2 \\ 5 & -3 \end{bmatrix} = \begin{bmatrix} -9 + 10(2^k) & 6 - 6(2^k) \\ -15 + 15(2^k) & 10 - 9(2^k) \end{bmatrix}.$$

Similarly,

$$A^{-2} = \begin{bmatrix} -9 + 10(2^{-2}) & 6 - 6(2^{-2}) \\ -15 + 15(2^{-2}) & 10 - 9(2^{-2}) \end{bmatrix} = \begin{bmatrix} -13/2 & 9/2 \\ -45/4 & 31/4 \end{bmatrix}$$

5.1 Orthogonality in \mathbb{R}^n

At the end of the last section, we saw that orthogonality could be detected with the dot product.

Definition. A set of vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ is an *orthogonal set* of vectors in \mathbb{R}^n if $\mathbf{v}_i \cdot \mathbf{v}_j = 0$ whenever $i \neq j$ for $i, j = 1, \dots, k$.

Theorem 5.1.1. *If $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ is an orthogonal set of nonzero vectors in \mathbb{R}^n , then it is linearly independent.*

Proof. Let a_i be scalars for which

$$\mathbf{0} = a_1 \mathbf{v}_1 + \cdots + a_k \mathbf{v}_k$$

Then for each $i = 1, \dots, k$,

$$\begin{aligned} 0 &= \mathbf{0} \cdot \mathbf{v}_i \\ &= (a_1 \mathbf{v}_1 + \cdots + a_k \mathbf{v}_k) \cdot \mathbf{v}_i \\ &= a_1 (\mathbf{v}_1 \cdot \mathbf{v}_i) + \cdots + a_k (\mathbf{v}_k \cdot \mathbf{v}_i) \\ &= a_i (\mathbf{v}_i \cdot \mathbf{v}_i) \end{aligned}$$

and since $\mathbf{v}_i \cdot \mathbf{v}_i \neq 0$, then it must be that $a_i = 0$. □

Definition. A basis \mathcal{B} for \mathbb{R}^n is an *orthogonal basis* if it is also an orthogonal set.

Example 5.1.2. The standard basis \mathcal{E} for \mathbb{R}^n is orthogonal.

Example 5.1.3. $\mathcal{B} = \left\{ \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$ is an orthogonal basis for \mathbb{R}^2 (this is straightforward to check).

Notice that for a vector $\mathbf{v} = [x, y] \in \mathbb{R}^2$, we have that $\mathbf{v} = x\mathbf{e}_1 + y\mathbf{e}_2$

Theorem 5.1.4 (Poole Theorem 5.2). *Suppose $\mathcal{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_k\}$ is an orthogonal basis for the subspace W of \mathbb{R}^n and let \mathbf{w} be any vector in W . Then the coefficients c_i of the linear combination*

$$\mathbf{w} = c_1 \mathbf{b}_1 + \cdots + c_k \mathbf{b}_k$$

are obtained by

$$c_i = \frac{\mathbf{w} \cdot \mathbf{b}_i}{\mathbf{b}_i \cdot \mathbf{b}_i}.$$


In other words,

$$\mathbf{w} = \text{proj}_{\mathbf{b}_1}(\mathbf{w}) + \cdots + \text{proj}_{\mathbf{b}_k}(\mathbf{w}).$$

Example 5.1.5. Let $\mathbf{v} = [4, 2]^T \in \mathbb{R}^2$ and consider the standard basis \mathcal{E} . A straightforward computation shows that

$$\text{proj}_{\mathbf{e}_1}(\mathbf{v}) = 4\mathbf{e}_1 \quad \text{and} \quad \text{proj}_{\mathbf{e}_2}(\mathbf{v}) = 2\mathbf{e}_2$$

and clearly $\mathbf{v} = 4\mathbf{e}_1 + 2\mathbf{e}_2$.




Example 5.1.6. Let \mathbf{v} be as in the previous example and consider the orthogonal basis $\mathcal{B} = \left\{ \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$. Computing the orthogonal projection of \mathbf{v} onto the basis vectors we have that

$$\text{proj}_{\mathbf{b}_1}(\mathbf{v}) = \frac{4-2}{2}\mathbf{b}_1 = \mathbf{b}_1 \quad \text{and} \quad \text{proj}_{\mathbf{b}_2}(\mathbf{v}) = \frac{4+2}{2}\mathbf{b}_2 = 3\mathbf{b}_2$$

and certainly

$$\mathbf{v} = \mathbf{b}_1 + 3\mathbf{b}_2. \tag{5.1.1}$$




Example 5.1.7. To see why orthogonality of the basis is important, let \mathbf{v} be as in the previous two examples and consider the basis $\mathcal{B} = \left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$, which is not orthogonal. Computing the orthogonal projection of \mathbf{v} onto the basis vectors, we have that

$$\text{proj}_{\mathbf{b}_1}(\mathbf{v}) = 4\mathbf{b}_1 \quad \text{and} \quad \text{proj}_{\mathbf{b}_2}(\mathbf{v}) = \frac{4+2}{2}3\mathbf{b}_2$$

but

$$\mathbf{v} = \begin{bmatrix} 4 \\ 2 \end{bmatrix} \neq \begin{bmatrix} 7 \\ 3 \end{bmatrix} = 4\mathbf{b}_1 + 3\mathbf{b}_2$$

(in fact, by inspection we see that $\mathbf{v} = 2\mathbf{b}_1 + 2\mathbf{b}_2$)



5.1.1 Orthonormality

Definition. An orthogonal set of vectors is *orthonormal* if each vector is also a unit vector. An *orthogonal basis* for \mathbb{R}^n is an orthonormal set that is also a basis for \mathbb{R}^n .

Proposition 5.1.8. Suppose $\mathcal{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ is a basis for \mathbb{R}^n . Then \mathcal{B} is orthonormal if and only if $\mathbf{b}_i \cdot \mathbf{b}_j = \delta_{ij}$ for all $i, j = 1, \dots, n$ (where δ_{ij} is the Kronecker delta).

Example 5.1.9. The standard basis \mathcal{E} for \mathbb{R}^n is orthonormal.

Example 5.1.10. The basis \mathcal{B} from Example 5.1.3 is not orthonormal because $\left\| \begin{bmatrix} 1 \\ -1 \end{bmatrix} \right\| = \left\| \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\| = \sqrt{2}$.


Obviously we like the standard basis and its orthonormality, so it would be nice if we could supe up an orthogonal basis to be a normal basis. Indeed, this is easily achieved by merely normalizing each vector.

Proposition 5.1.11. If $\mathcal{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ is an orthogonal basis for \mathbb{R}^n , then $\mathcal{B}' = \left\{ \frac{\mathbf{b}_1}{\|\mathbf{b}_1\|}, \dots, \frac{\mathbf{b}_n}{\|\mathbf{b}_n\|} \right\}$ is an orthonormal basis for \mathbb{R}^n

Example 5.1.12. Normalizing the basis vectors in Example 5.1.3, we get that

$$\mathcal{B}' = \left\{ \begin{bmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{bmatrix}, \begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix} \right\}$$

is an orthonormal basis for \mathbb{R}^2 .



What you may notice is that the orthonormal basis \mathcal{B}' in the previous example is just a $-\frac{\pi}{4}$ rotation of the standard basis. Amazingly, this same feature is very nearly true of all orthonormal bases.

Theorem 5.1.13. *Every orthonormal basis of \mathbb{R}^n can be obtained by a rotation of the standard basis (possibly followed by a permutation of the basis vectors).*

It follows that, given two orthonormal bases \mathcal{B} and \mathcal{C} of \mathbb{R}^2 , the change of basis matrix is either a rotation matrix

$$P_{\mathcal{C} \leftarrow \mathcal{B}} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

or a rotation matrix times a matrix that swaps the basis vectors

$$P_{\mathcal{C} \leftarrow \mathcal{B}} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} = \begin{bmatrix} \sin \theta & -\cos \theta \\ \cos \theta & \sin \theta \end{bmatrix}$$

(although, since the particular angle θ is freely chosen, we would often make the substitution $\theta = \varphi + \frac{\pi}{2}$ to get

$$\begin{bmatrix} \cos \varphi & \sin \varphi \\ \sin \varphi & -\cos \varphi \end{bmatrix}$$

as this provides some nice symmetries in the forms of the matrix).

5.2 Orthogonal Complements and Orthogonal Projections

5.2.1 Orthogonal Complement

Definition. Let W be a subspace of \mathbb{R}^n . A vector $\mathbf{v} \in \mathbb{R}^n$ is *orthogonal to W* if it is orthogonal to every vector $\mathbf{w} \in W$. The collection of all vectors orthogonal to W is a subspace called the *orthogonal complement to W* and is denoted W^\perp .

Example 5.2.1. Suppose W is the xy -plane in \mathbb{R}^3 (i.e. the set of vectors $[x, y, 0]^T$). Then W^\perp is the z -axis (i.e. the set of vectors $[0, 0, z]^T$). Geometrically, W^\perp is the line through the origin that is perpendicular to W .

Theorem 5.2.2 (Poole Theorem 5.9). *Let W be a subspace of \mathbb{R}^n .*

1. W^\perp is also a subspace of \mathbb{R}^n .
2. $(W^\perp)^\perp = W$
3. The only vector common to both W and W^\perp is $\mathbf{0}$ (we say that W and W^\perp have trivial intersection).
4. If $W = \text{Span}(\mathbf{w}_1, \dots, \mathbf{w}_k)$, then W^\perp is the set of vectors perpendicular to each \mathbf{w}_i .

The following result encodes what you may have noticed from a previous homework problem.

Theorem 5.2.3 (Poole Theorem 5.10). *Let A be an $m \times n$ matrix. Then $\text{Null}(A) = (\text{row } A)^\perp$ and $\text{Null}(A^T) = (\text{col } A)^\perp$.*

Proof. If A is an $m \times n$ matrix with rows $\mathbf{A}_1, \dots, \mathbf{A}_m$ and $\mathbf{x} \in \mathbb{R}^n$, then

$$A\mathbf{x} = \begin{bmatrix} \mathbf{A}_1 \cdot \mathbf{x} \\ \vdots \\ \mathbf{A}_m \cdot \mathbf{x} \end{bmatrix}$$

so $A\mathbf{x} = \mathbf{0}$ precisely when $\mathbf{A}_i \cdot \mathbf{x} = 0$ for each $i = 1, \dots, m$. □

As such, solving for the orthogonal complement can be done by explicitly solving for the null space of the appropriate matrix of vectors.

Example 5.2.4. Let $W = \text{Span}\left(\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}, \begin{bmatrix} 2 \\ -1 \\ 0 \end{bmatrix}\right)$ be a plane in \mathbb{R}^3 . By viewing the above vectors as row vectors of a matrix and applying the above theorem, we have that

$$W^\perp = \text{Null}\left(\begin{bmatrix} 1 & 2 & 3 \\ 2 & -1 & 0 \end{bmatrix}\right).$$

Since

$$\text{RREF}\left(\begin{bmatrix} 1 & 2 & 3 \\ 2 & -1 & 0 \end{bmatrix}\right) = \begin{bmatrix} 1 & 0 & 3/5 \\ 0 & 1 & 6/5 \end{bmatrix}$$

we deduce that $W^\perp = \text{Span}\left(\begin{bmatrix} 3/5 \\ 6/5 \\ -1 \end{bmatrix}\right)$.

5.2.2 Orthogonal Projections

Recall that, for two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^2$, $\text{proj}_{\mathbf{u}}(\mathbf{v})$ is a vector contained in $\text{Span}(\mathbf{u})$. Moreover, $\mathbf{v} - \text{proj}_{\mathbf{u}}(\mathbf{v})$ is a vector perpendicular to \mathbf{u} . In this way, we can “decompose” \mathbf{v} into a sum of two vectors. If $W = \text{Span}(\mathbf{u})$ then we have

$$\mathbf{v} = \underbrace{\text{proj}_{\mathbf{u}}(\mathbf{v})}_{\text{in } W} + \underbrace{(\mathbf{v} - \text{proj}_{\mathbf{u}}(\mathbf{v}))}_{\text{in } W^\perp}$$

Remark. Your book uses the notation $\text{perp}_{\mathbf{u}}(\mathbf{v})$ to mean $\mathbf{v} - \text{proj}_{\mathbf{u}}(\mathbf{v})$. This is reasonable, but I can’t say its particularly common.

The following definition extends the idea of orthogonal projection onto an entire subspace.

Definition. Let W be a subspace of \mathbb{R}^n and $\{\mathbf{w}_1, \dots, \mathbf{w}_k\}$ an orthogonal basis for W . For a vector $\mathbf{v} \in \mathbb{R}^n$, the *orthogonal projection of \mathbf{v} onto W* is

$$\text{proj}_W(\mathbf{v}) = \text{proj}_{\mathbf{w}_1}(\mathbf{v}) + \dots + \text{proj}_{\mathbf{w}_k}(\mathbf{v})$$

and the *component of \mathbf{v} orthogonal to W* is

$$\text{perp}_W(\mathbf{v}) = \mathbf{v} - \text{proj}_W(\mathbf{v}).$$


Remark. Once again, for any subspace W of \mathbb{R}^n and $\mathbf{v} \in \mathbb{R}^n$, we have the *orthogonal decomposition* of \mathbf{v} :

$$\mathbf{v} = \underbrace{\text{proj}_{\mathbf{u}}(\mathbf{v})}_{\text{in } W} + \underbrace{(\mathbf{v} - \text{proj}_{\mathbf{u}}(\mathbf{v}))}_{\text{in } W^\perp}$$

Example 5.2.5. Suppose W is the xy -plane in \mathbb{R}^3 and let $\mathbf{v} = [3, 4, 5]^T$. Then

$$\text{proj}_W(\mathbf{v}) = \text{proj}_{\mathbf{e}_1}(\mathbf{v}) + \text{proj}_{\mathbf{e}_2}(\mathbf{v}) = 3\mathbf{e}_1 + 4\mathbf{e}_2 = [3, 4, 0]^T.$$


Visually,



Example 5.2.6. Suppose $\mathbf{w}_1 = \begin{bmatrix} 2 \\ -2 \\ 2 \end{bmatrix}$, $\mathbf{w}_2 = \begin{bmatrix} -2 \\ 0 \\ 2 \end{bmatrix}$, and $\mathbf{v} = \begin{bmatrix} 3 \\ 4 \\ 5 \end{bmatrix}$ are vectors in \mathbb{R}^3 and $W = \text{Span}(\mathbf{w}_1, \mathbf{w}_2)$. We first check that $\mathbf{w}_1 \cdot \mathbf{w}_2 = 0$, whence $\{\mathbf{w}_1, \mathbf{w}_2\}$ is an orthogonal basis for W . To compute the projection of \mathbf{v} onto W

$$\begin{aligned} \text{proj}_W(\mathbf{v}) &= \text{proj}_{\mathbf{w}_1}(\mathbf{v}) + \text{proj}_{\mathbf{w}_2}(\mathbf{v}) \\ &= \left(\frac{\mathbf{v} \cdot \mathbf{w}_1}{\mathbf{w}_1 \cdot \mathbf{w}_1} \right) \mathbf{w}_1 + \left(\frac{\mathbf{v} \cdot \mathbf{w}_2}{\mathbf{w}_2 \cdot \mathbf{w}_2} \right) \mathbf{w}_2 \\ &= \left(\frac{6 - 8 + 10}{12} \right) \mathbf{w}_1 + \left(\frac{-6 + 10}{8} \right) \mathbf{w}_2 \\ &= \frac{2}{3} \mathbf{w}_1 + \frac{3}{4} \mathbf{w}_2 = \frac{1}{6}[-1, -8, 17]^T \end{aligned}$$

Visually,



What's amazing is that the orthogonal decomposition of a vector is actually unique.

Theorem 5.2.7 (Orthogonal Decomposition Theorem). *Let W be a subspace of \mathbb{R}^n and $\mathbf{v} \in \mathbb{R}^n$. Then there are unique vectors $\mathbf{w} \in W$ and $\mathbf{w}^\perp \in W^\perp$ for which $\mathbf{v} = \mathbf{w} + \mathbf{w}^\perp$.*

In particular, $\mathbf{w} = \text{proj}_W(\mathbf{v})$ and $\mathbf{w}^\perp = \mathbf{v} - \text{proj}_W(\mathbf{v})$.

Corollary 5.2.8. $(W^\perp)^\perp = W$

Theorem 5.2.9. *If W is a subspace of \mathbb{R}^n , then $\dim W + \dim W^\perp = n$*

The above isn't particularly surprising; it's exactly what we saw happen in the first example of this section. More generally, given a basis of column vectors $\{\mathbf{w}_1, \dots, \mathbf{w}_k\}$ for W , we find W^\perp by using Theorem 5.2.3:

$$W^\perp = \text{Null} \left(\begin{bmatrix} -\mathbf{w}_1^T & - \\ \vdots & \\ -\mathbf{w}_k^T & - \end{bmatrix} \right).$$

The matrix above has size $k \times n$ and rank k , and so the result follows by Rank-Nullity.

Of course, the proof given in the book doesn't rely on Rank-Nullity at all. In fact, Rank-Nullity actually a corollary of the above theorem (and the proof is given by the same observation as above).

5.3 The Gram–Schmidt Process and the QR Factorization

5.3.1 Gram–Schmidt

It sure would be nice to be able to find an orthogonal basis for every subspace, huh?

Example 5.3.1. Let $\mathbf{b}_1 = \begin{bmatrix} 3 \\ 4 \\ 0 \end{bmatrix}$, $\mathbf{b}_2 = \begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix}$ be vectors in \mathbb{R}^3 and let $W = \text{Span}(\mathbf{b}_1, \mathbf{b}_2)$. Recall that

$$\mathbf{b}_2 - \text{proj}_{\mathbf{b}_1}(\mathbf{b}_2) = \begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix} - \begin{bmatrix} 21/25 \\ 28/25 \\ 0 \end{bmatrix} = \begin{bmatrix} 4/25 \\ -3/25 \\ 2 \end{bmatrix}$$

is perpendicular to \mathbf{b}_1 and is still contained within W . This means that $\{\mathbf{b}_1, \mathbf{b}_2 - \text{proj}_{\mathbf{b}_1}(\mathbf{b}_2)\}$ is an orthogonal basis for W !

As it turns out, the above example can be extended into any dimension, and this iterative process is known as the *Gram–Schmidt orthogonalization*.

Theorem 5.3.2 (Gram–Schmidt). *Let $\mathcal{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_k\}$ be a basis for W , a subspace of \mathbb{R}^n .*

1. Let $\mathbf{x}_1 = \mathbf{b}_1$, and let $W_1 = \text{Span}(\mathbf{x}_1)$.
2. For each $i = 2, \dots, k$, let $\mathbf{x}_i = \mathbf{b}_i - \text{proj}_{W_{i-1}}(\mathbf{b}_i)$ and set $W_i = \text{Span}(\mathbf{x}_1, \dots, \mathbf{x}_i)$

For each i , $\{x_1, \dots, x_i\}$ is an orthogonal basis for W_i and $W_k = W$.

Remark. The basis produced by the Gram–Schmidt process is not unique, as every step required a choice of basis vector \mathbf{b}_i . In \mathbb{R}^2 for example, running the procedure on the basis $\left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$ will

produce either the standard basis or the basis $\left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1/2 \\ -1/2 \end{bmatrix} \right\}$.

Remark. One can always scale the basis elements to have norm 1, further producing an orthonormal basis.

Example 5.3.3. Find an orthonormal basis for $W = \text{Span} \left(\begin{bmatrix} 1 \\ 2 \\ 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 2 \\ 2 \end{bmatrix}, \begin{bmatrix} 2 \\ 0 \\ 1 \\ 2 \end{bmatrix} \right)$.

We choose

$$\mathbf{x}_1 = [1, 2, 2, 0]^T$$

and we set $W_1 = \text{Span}(\mathbf{x}_1)$. Then

$$\begin{aligned} \mathbf{x}_2 &= [0, 1, 2, 2]^T - \text{proj}_{W_1}([0, 1, 2, 2]^T) \\ &= [0, 1, 2, 2]^T - \text{proj}_{\mathbf{x}_1}([0, 1, 2, 2]^T) \\ &= [0, 1, 2, 2]^T - \left[\frac{2}{3}, \frac{4}{3}, \frac{4}{3}, 0 \right]^T \\ &= \left[-\frac{2}{3}, -\frac{1}{3}, \frac{2}{3}, 2 \right]^T \end{aligned}$$

and we set $W_2 = \text{Span}(\mathbf{x}_1, \mathbf{x}_2)$. Then

$$\begin{aligned} \mathbf{x}_3 &= [2, 0, 1, 2]^T - \text{proj}_{W_2}([2, 0, 1, 2]^T) \\ &= [2, 0, 1, 2]^T - \text{proj}_{\mathbf{x}_1}([2, 0, 1, 2]^T) - \text{proj}_{\mathbf{x}_2}([2, 0, 1, 2]^T) \\ &= [2, 0, 1, 2]^T - \left[\frac{4}{9}, \frac{8}{9}, \frac{8}{9}, 0 \right]^T - \left[-\frac{4}{9}, -\frac{2}{9}, \frac{4}{9}, \frac{4}{9} \right]^T \\ &= \left[2, -\frac{2}{3}, -\frac{1}{3}, \frac{2}{3} \right]^T. \end{aligned}$$

and $\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$ is an orthogonal basis for W . To form an orthonormal basis, we normalize each of these vectors, hence an orthonormal basis for W is

$$\left\{ \begin{bmatrix} 1/3 \\ 1/3 \\ 2/3 \\ 0 \end{bmatrix}, \begin{bmatrix} -2/3\sqrt{5} \\ -1/3\sqrt{5} \\ 2/3\sqrt{5} \\ 2/\sqrt{5} \end{bmatrix}, \begin{bmatrix} 2/\sqrt{5} \\ -2/3\sqrt{5} \\ -1/3\sqrt{5} \\ 2/3\sqrt{5} \end{bmatrix} \right\}.$$

5.3.2 QR Factorization

If A is an $m \times n$ matrix with linearly independent columns (implying $m \geq n$), then applying the Gram–Schmidt process to the columns yields a useful factorization of A .

Theorem 5.3.4 (QR Factorization). *Let A be an $m \times n$ matrix with linearly independent columns. Then there exists a matrix Q with orthonormal columns and an invertible upper triangular matrix R for which $A = QR$.*

The proof is constructive. Let $\{\mathbf{A}_1, \dots, \mathbf{A}_n\}$ be the columns of A and let $\{\mathbf{Q}_1, \dots, \mathbf{Q}_n\}$ be the orthonormal basis produced from applying Gram-Schmidt to the \mathbf{A}_i 's. Notice that in the Gram-Schmidt process, we have

$$\begin{aligned} \mathbf{Q}_1 &= c_1 \mathbf{A}_1 \\ \mathbf{Q}_2 &= c_2 \left(\mathbf{A}_2 - \left(\frac{\mathbf{Q}_1 \cdot \mathbf{A}_2}{\mathbf{Q}_1 \cdot \mathbf{Q}_1} \right) \mathbf{Q}_1 \right) \\ \mathbf{Q}_3 &= c_3 \left(\mathbf{A}_3 - \left(\frac{\mathbf{Q}_1 \cdot \mathbf{A}_3}{\mathbf{Q}_1 \cdot \mathbf{Q}_1} \right) \mathbf{Q}_1 - \left(\frac{\mathbf{Q}_2 \cdot \mathbf{A}_3}{\mathbf{Q}_2 \cdot \mathbf{Q}_2} \right) \mathbf{Q}_2 \right) \\ &\vdots \end{aligned}$$

where the c_i 's are all the scalars normalizing the vectors.

Since all of the dot products are just scalars, we can write

$$r_{ij} = \begin{cases} 1/c_j & \text{if } i = j \\ \left(\frac{\mathbf{Q}_i \cdot \mathbf{A}_j}{\mathbf{Q}_i \cdot \mathbf{Q}_i} \right) & \text{if } i \neq j \end{cases}$$

and rearrange the above equations to be

$$\begin{aligned}\mathbf{A}_1 &= r_{11}\mathbf{Q}_1 \\ \mathbf{A}_2 &= r_{12}\mathbf{Q}_1 + r_{22}\mathbf{Q}_2 \\ \mathbf{A}_3 &= r_{13}\mathbf{Q}_1 + r_{23}\mathbf{Q}_2 + r_{33}\mathbf{Q}_3 \\ &\vdots\end{aligned}$$

The above system can be represented as the following matrix product:

$$A = \begin{bmatrix} | & & | \\ \mathbf{A}_1 & \cdots & \mathbf{A}_n \\ | & & | \end{bmatrix} = \begin{bmatrix} | & & | \\ \mathbf{Q}_1 & \cdots & \mathbf{Q}_n \\ | & & | \end{bmatrix} \begin{bmatrix} r_{11} & r_{12} & \cdots & r_{1n} \\ & r_{22} & \cdots & r_{2n} \\ & & \ddots & \vdots \\ & & & r_{nn} \end{bmatrix} = QR$$

Remark. We can always take the diagonal entries r_{ii} to be positive: if $r_{ii} < 0$, then simply replace \mathbf{Q}_i with $-\mathbf{Q}_i$.

Remark. Since Q is $m \times n$ with orthonormal columns, then $Q^T Q = I_n$, so in fact $R = Q^T A$, saving us some time in computing R .

Example 5.3.5. Compute the QR factorization of $A = \begin{bmatrix} 12 & -51 & -4 \\ 6 & 167 & 68 \\ -4 & 24 & 41 \end{bmatrix}$

We first apply the Gram-Schmidt process to the columns. Let \mathbf{A}_i denote the i^{th} column of A . We take $\mathbf{x}_1 = \mathbf{A}_1$. Letting $W_1 = \text{Span}(\mathbf{x}_1)$,

$$\mathbf{x}_2 = \mathbf{A}_2 - \text{proj}_{W_1}(\mathbf{A}_2) = [-69, 158, 30]^T.$$

Letting $W_2 = \text{Span}(\mathbf{x}_1, \mathbf{x}_2)$,

$$\mathbf{x}_3 = \mathbf{A}_3 - \text{proj}_{W_2}(\mathbf{A}_3) = \left[\frac{58}{5}, -\frac{6}{5}, 33 \right]^T.$$

Now $\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$ is an orthogonal basis for \mathbb{R}^3 . Letting $\mathbf{Q}_i = \frac{\mathbf{x}_i}{\|\mathbf{x}_i\|}$, we form the orthogonal matrix

$$Q = \begin{bmatrix} | & | & | \\ \mathbf{Q}_1 & \mathbf{Q}_2 & \mathbf{Q}_3 \\ | & | & | \end{bmatrix} = \begin{bmatrix} \frac{6}{7} & -\frac{69}{175} & \frac{58}{175} \\ \frac{3}{7} & \frac{158}{175} & -\frac{6}{175} \\ -\frac{2}{7} & \frac{6}{35} & \frac{33}{35} \end{bmatrix}$$

and

$$R = Q^T A = \begin{bmatrix} 14 & 21 & 14 \\ 0 & 175 & 70 \\ 0 & 0 & 35 \end{bmatrix}.$$


7.3 Least Squares

Example 7.3.1. Let $A = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & -1 \\ 1 & -1 & -1 \\ 1 & -1 & 1 \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} 2 \\ 0 \\ 2 \\ 0 \end{bmatrix}$. We consider the system $A\mathbf{x} = \mathbf{b}$. Through row reducing the augmented matrix $[A \mid \mathbf{b}]$, we see that the system is inconsistent:

$$[A \mid \mathbf{b}] = \left[\begin{array}{ccc|c} 1 & 1 & 1 & 2 \\ 1 & 1 & -1 & 0 \\ 1 & -1 & -1 & 2 \\ 1 & -1 & 1 & 0 \end{array} \right] \xrightarrow{\text{RREF}} \left[\begin{array}{ccc|c} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{array} \right]$$

From this it follows that \mathbf{b} is not in $\text{col}(A)$. We may be interested to know which vector in $\text{col}(A)$ is closest to \mathbf{b} .

Certainly we could use some calculus techniques to find this, but we could also appeal to a fact of Euclidean geometry – the distance between a point p and a subspace W of \mathbb{R}^n is the length of the line segment ℓ which is perpendicular to W and has endpoint p .



As such, $\text{proj}_{\text{Col}(A)}(\mathbf{b})$ is the unique vector in $\text{Col}(A)$ that is closest to \mathbf{b} .

Definition. Given an $m \times n$ matrix A and a vector $\mathbf{b} \in \mathbb{R}^m$, and solution to the linear system

$$A\mathbf{x} = \text{proj}_{\text{Col}(A)}(\mathbf{b})$$

is called a *least squares solution* of the linear system $A\mathbf{x} = \mathbf{b}$.

Why might we care about this? As it happens quite frequently in life, we sample real-world data and use it to form a vector space (essentially, a linear approximation representing the data). If this vector space approximation doesn't exactly match the data, then solving linear systems outright becomes largely impossible. Least squares is one tool that allows us to continue doing computations - if our vector space approximates the data really well, then the least squares solution approximates the actual solution really well too.

In these notes I won't go through any real-world examples of this. Sorry about that. For more, you might look more into the CMDA program. For now, we'll look at some more of the theoretical/computational aspects of least squares approximations.

Remark. For any linear system, $A\mathbf{x} = \mathbf{b}$, there always exists a least squares solution (because $\text{proj}_{\text{Col}(A)}(\mathbf{b})$ is a vector in $\text{Col}(A)$). This least squares solution is unique if $\text{nullity}(A) = 0$.

Remark. If $\mathbf{b} \in \text{Col}(A)$, then $\mathbf{b} = \text{proj}_{\text{Col}(A)}(\mathbf{b})$, hence the least squares solution set is exactly the same as the actual solution set.

Example 7.3.2. Find the least squares solution set for $A\mathbf{x} = \mathbf{b}$ where $A = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$.

The size of the matrix shows that $\text{nullity}(A) > 0$ (and in fact, a quick computation shows that $\text{nullity}(A) = 2$), so the remark above indicates that there are infinitely many least squares solutions.

To get the least squares solutions, we first compute $\text{proj}_{\text{Col}(A)}(\mathbf{b})$. We first note that

$$\text{Col}(A) = \text{Span} \left(\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \right)$$

and the basis for $\text{Col}(A)$ is orthogonal. Hence

$$\text{proj}_{\text{Col}(A)}(\mathbf{b}) = \text{proj}_{[1,0,1]^T}(\mathbf{b}) + \text{proj}_{[0,1,0]}(\mathbf{b}) = \begin{bmatrix} 2 \\ 0 \\ 2 \end{bmatrix} + \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix} = \begin{bmatrix} 2 \\ 2 \\ 2 \end{bmatrix}.$$

The least squares solution set is now the solution set to the equation

$$A\mathbf{x} = \text{proj}_{\text{Col}(A)}(\mathbf{b}) \quad \Rightarrow \quad \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} 2 \\ 2 \\ 2 \end{bmatrix}$$

hence

$$\left[\begin{array}{cccc|c} 1 & 0 & 0 & 1 & 2 \\ 0 & 1 & 1 & 0 & 2 \\ 1 & 0 & 0 & 1 & 2 \end{array} \right] \xrightarrow{RREF} \left[\begin{array}{cccc|c} 1 & 0 & 0 & 1 & 2 \\ 0 & 1 & 1 & 0 & 2 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right]$$

hence $x_1 = 2 - x_4$ and $x_2 = 2 - x_3$, so the least squares solution set is

$$\left\{ \begin{bmatrix} 2 \\ 2 \\ 0 \\ 0 \end{bmatrix} + s \begin{bmatrix} 0 \\ -1 \\ 1 \\ 0 \end{bmatrix} + t \begin{bmatrix} -1 \\ 0 \\ 0 \\ 1 \end{bmatrix} : \text{ where } s, t \in \mathbb{R} \right\}$$


Theorem 7.3.3. Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. The vector $\mathbf{x} \in \mathbb{R}^n$ is a least squares solution to $A\mathbf{x} = \mathbf{b}$ if and only if \mathbf{x} is a solution to the normal equations

$$A^T A\mathbf{x} = A^T \mathbf{b}.$$

Exercise 7.3.1. Show that the least squares solutions in Example 7.3.2 also satisfy the normal equations.

7.3.1 Least Squares Approximation

Suppose that we have points P_1, P_2, P_3 in the plane and we approximate these three points with the line $y = mx + b$. Let $\varepsilon_1, \varepsilon_2, \varepsilon_3$ be the vertical distance between these points and the line. The *error vector* is $\boldsymbol{\varepsilon} = [\varepsilon_1, \varepsilon_2, \varepsilon_3]$ and the *least squares error* is $\|\boldsymbol{\varepsilon}\|$.



Example 7.3.4. The image above shows the case where $P_1(1, 3)$, $P_2(3, 3)$, $P_3(-2, -4)$, $y = x + 1$, and $\varepsilon_1 = 1$, $\varepsilon_2 = 1$ and $\varepsilon_3 = 3$. Thus the least squares error is

$$\|\boldsymbol{\varepsilon}\| = \sqrt{\varepsilon_1^2 + \varepsilon_2^2 + \varepsilon_3^2} = \sqrt{11}$$

Often times we're mostly interested with finding the line that minimizes this error, and this is called the *least squares error approximating line*.

If we have a bunch of points (x_i, y_i) and we approximate them with the line $y = mx + b$, then we have that

$$\varepsilon_i = y_i - y = y_i - mx_i + b.$$

Letting

$$A = \begin{bmatrix} 1 & x_1 \\ \vdots & \vdots \\ 1 & x_n \end{bmatrix}, \quad \mathbf{x} = \begin{bmatrix} b \\ m \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}$$

then we have that

$$\boldsymbol{\varepsilon} = \mathbf{b} - A\mathbf{x}.$$

The following theorem tells us that a least-squares solution is actually the one that minimizes the least squares error

Theorem 7.3.5. Given $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^n$, a vector $\hat{\mathbf{x}}$ is a least squares solutions to $A\mathbf{x} = \mathbf{b}$ if, for every vector $\mathbf{x} \in \mathbb{R}^n$,

$$\|\mathbf{b} - A\hat{\mathbf{x}}\| \leq \|\mathbf{b} - A\mathbf{x}\|$$

Example 7.3.6. We want to find a least squares approximating line $y = mx + b$ for the points in Example 7.3.4. Using the above discussion, this is equivalent to finding a least squares solution to $A\mathbf{x} = \mathbf{b}$ where

$$A = \begin{bmatrix} 1 & 1 \\ 1 & 3 \\ 1 & -2 \end{bmatrix}, \quad \mathbf{x} = \begin{bmatrix} b \\ m \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 3 \\ 3 \\ -4 \end{bmatrix}.$$

We thus look for the solution set to

$$A^T A \mathbf{x} = A^T \mathbf{b}$$

$$\begin{bmatrix} 3 & 2 \\ 2 & 14 \end{bmatrix} \begin{bmatrix} b \\ m \end{bmatrix} = \begin{bmatrix} 2 \\ 20 \end{bmatrix}.$$

By row reducing the augmented system $[A^T A \mid A^T \mathbf{b}]$, we get

$$[A^T A \mid A^T \mathbf{b}] = \left[\begin{array}{cc|c} 3 & 2 & 2 \\ 2 & 14 & 20 \end{array} \right] \xrightarrow{\text{RREF}} \left[\begin{array}{cc|c} 1 & 0 & -6/19 \\ 0 & 1 & 28/29 \end{array} \right]$$


whence the least squares approximating line is

$$y = \frac{28}{19}x - \frac{6}{19}.$$

The least squares error vector for this line is

$$\boldsymbol{\varepsilon} = \mathbf{b} - A\mathbf{x} = \begin{bmatrix} 35/19 \\ -21/19 \\ -14/19 \end{bmatrix}$$

and the least squares error is $\|\boldsymbol{\varepsilon}\| \approx 2.2711$. This is better than the previous line we had, as that had an error of $\sqrt{11} \approx 3.3166$.



The above strategy can be employed for approximating planes, etc, making the obvious changes to A , \mathbf{x} , and \mathbf{b} .