

MATH 235: Linear Algebra 2 for Honours Mathematics

Fall 2023 Edition

© Faculty of Mathematics, University of Waterloo

We would like to acknowledge:
Authors: Faisal al-Faisal and Ty Ghaswala
Contributors: Anton Mosunov and Peter Wood
Thanks also to Ian Payne for many helpful comments.

Contents

1	Abstract Vector Spaces	6
1.1	The Definition of a Vector Space	6
1.2	Subspaces	13
1.3	Bases and Dimension	17
2	Linear Transformations	31
2.1	Linear Transformations Between Abstract Vectors	31
2.2	Rank and Nullity	34
2.3	Linear Maps as Matrices	37
2.4	Change of Coordinates	44
2.5	Isomorphisms of Vector Spaces	48
3	Diagonalizability	60
3.1	Eigenvectors and Diagonalization	60
3.2	Diagonalization	68
3.3	Applications of Diagonalization	78
4	Inner Product Spaces	83
4.1	Inner Products	83
4.2	Orthogonality and Norm	90
4.3	Orthonormal Bases	95
4.4	Projections	100
4.5	The Gram–Schmidt Orthogonalization Procedure	103
4.6	Projection onto a Subspace and Orthogonal Complements	107
4.7	Application: Method of Least Squares	115
5	Unitary Diagonalization	122
5.1	The Adjoint	122
5.2	Orthogonal and Unitary Matrices	125
5.3	Schur’s Triangularization Theorem	127
5.4	Orthogonal and Unitary Diagonalization of Matrices	133
5.5	Unitary Diagonalization of Operators	141
5.6	Application: Classifying Quadratic Forms	143
5.7	Application: Inner Products and Gram Matrices	149
6	The Singular Value Decomposition	156
6.1	Singular Values and Singular Vectors	156
6.2	Singular Value Decomposition of Matrices	160
6.3	Singular Value Decomposition of Linear Maps	166
6.4	Applications of the Singular Value Decomposition	168

To the reader

These notes are in the process of being updated in Fall 2023. If you spot any typos/errors or have any feedback that you would like to share, please do so by posting on [Piazza](#).

In order to ensure that you have the most up-to-date version, be sure to download the latest one from LEARN.

Chapter 1

Abstract Vector Spaces

1.1 The Definition of a Vector Space

Linear algebra is the study of vector spaces. Before we formally define a vector space, let's introduce some familiar examples of vector spaces. As you go through each example, pay close attention to the similarities between each.

Example 1.1.1

The vector space \mathbb{R}^n is given by

$$\mathbb{R}^n = \left\{ \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} : a_i \in \mathbb{R} \text{ for all } i \right\}.$$

Addition and scalar multiplication are given by

$$\begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} + \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix} = \begin{bmatrix} a_1 + b_1 \\ \vdots \\ a_n + b_n \end{bmatrix} \quad \text{and} \quad \alpha \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} = \begin{bmatrix} \alpha a_1 \\ \vdots \\ \alpha a_n \end{bmatrix} \quad (\text{for } \alpha \in \mathbb{R}).$$

The intuitive picture that is helpful to have in mind are the cases of \mathbb{R}^2 and \mathbb{R}^3 that you are familiar with from previous courses. You can picture \mathbb{R}^2 as the Cartesian plane, and \mathbb{R}^3 as 3-dimensional space. In both of these vector spaces, you know how vector addition and scalar multiplication work, and intuitively, it's the same for \mathbb{R}^n . Although \mathbb{R}^n is an n -dimensional vector space, it is usually helpful to use the visual imagery of \mathbb{R}^2 and \mathbb{R}^3 .

In this course, we will study both real and complex vector spaces. The previous example is the most basic real vector space. There is an analogous complex vector space.

Example 1.1.2

The vector space \mathbb{C}^n is given by

$$\mathbb{C}^n = \left\{ \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} : a_i \in \mathbb{C} \text{ for all } i \right\}.$$

Addition and scalar multiplication are given by

$$\begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} + \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix} = \begin{bmatrix} a_1 + b_1 \\ \vdots \\ a_n + b_n \end{bmatrix} \quad \text{and} \quad \alpha \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} = \begin{bmatrix} \alpha a_1 \\ \vdots \\ \alpha a_n \end{bmatrix} \quad (\text{for } \alpha \in \mathbb{C})$$

just like in the previous example. Notice however that we now allow ourselves to use *complex scalars* $\alpha \in \mathbb{C}$. The vector space \mathbb{C}^n might appear to be even more challenging to visualize than \mathbb{R}^n . However, the algebra works just the same. In this course we will learn how our intuition from \mathbb{R}^2 and \mathbb{R}^3 will allow us to discover properties of \mathbb{C}^n . This is one of the strengths of linear algebra.

We will use the short-hand notation \mathbb{F} (for “field”) to denote either \mathbb{R} or \mathbb{C} when we do not wish to distinguish between them. This is convenient because many (though not at all!) of our results work equally well over both \mathbb{R} and \mathbb{C} .

Example 1.1.3

The vector space \mathbb{F}^n is given by

$$\mathbb{F}^n = \left\{ \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} : a_i \in \mathbb{F} \text{ for all } i \right\}.$$

Addition and scalar multiplication are given by

$$\begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} + \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix} = \begin{bmatrix} a_1 + b_1 \\ \vdots \\ a_n + b_n \end{bmatrix} \quad \text{and} \quad \alpha \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} = \begin{bmatrix} \alpha a_1 \\ \vdots \\ \alpha a_n \end{bmatrix} \quad (\text{for } \alpha \in \mathbb{F}).$$

If $\mathbb{F} = \mathbb{R}$, then \mathbb{F}^n is the real vector space \mathbb{R}^n from the first example; and if $\mathbb{F} = \mathbb{C}$, then $\mathbb{F}^n = \mathbb{C}^n$. Be mindful that the scalars α are always chosen to be in the appropriate \mathbb{F} . So when we are working with $\mathbb{F}^n = \mathbb{R}^n$, we only use real scalars $\alpha \in \mathbb{R}$.

Example 1.1.4

The vector space $\mathcal{P}_n(\mathbb{F})$ is the set of polynomials of degree at most n with coefficients in \mathbb{F} . That is

$$\mathcal{P}_n(\mathbb{F}) = \{a_0 + a_1x + \cdots + a_nx^n : a_i \in \mathbb{F} \text{ for all } i\}$$

with addition and scalar multiplication defined by

$$(a_0 + a_1x + \cdots + a_nx^n) + (b_0 + b_1x + \cdots + b_nx^n) = (a_0 + b_0) + (a_1 + b_1)x + \cdots + (a_n + b_n)x^n$$

and

$$\alpha(a_0 + a_1x + \cdots + a_nx^n) = (\alpha a_0) + (\alpha a_1)x + \cdots + (\alpha a_n)x^n \quad (\alpha \in \mathbb{F})$$

respectively.

If a coefficient is 0, we usually omit it. So instead of writing $2 + 3x + 0x^2 + 4x^3 + 0x^5$, we'll simply write $2 + 3x + 4x^3$.

For example, $1 + 2x - 3x^2 \in \mathcal{P}_2(\mathbb{R})$ and $1 + (2+i)x - x^5 \in \mathcal{P}_5(\mathbb{C})$. Both of these polynomials are also in $\mathcal{P}_{10}(\mathbb{C})$, if we pretend they are missing some 0 coefficients.

Here is an example of addition in $\mathcal{P}_2(\mathbb{R})$:

$$(4 + 7x) + (1 + x^2) = 5 + 7x + x^2.$$

Here is an example of scalar multiplication in $\mathcal{P}_3(\mathbb{C})$:

$$25i(1 + 2ix^3) = 25i - 50ix^3.$$

Warning: You may be used to thinking of polynomials as functions. In the context of this course, don't! Although it is sometimes useful to evaluate a polynomial at a certain number, in this course, polynomials are not functions. They are simply objects that you can add together and multiply by scalars.

Example 1.1.5

The vector space of m by n matrices with entries in \mathbb{F} is given by

$$M_{m \times n}(\mathbb{F}) = \left\{ \begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{m1} & \cdots & a_{mn} \end{bmatrix} : a_{ij} \in \mathbb{F} \text{ for all } i, j \right\}.$$

Addition and scalar multiplication are given by matrix addition and scalar multiplication of matrices as usual. For instance, in $M_{2 \times 2}(\mathbb{R})$,

$$\begin{bmatrix} 2 & 5 \\ 7 & \pi \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 3 & 5 \\ 8 & \pi + 1 \end{bmatrix} \quad \text{and} \quad \sqrt{2} \begin{bmatrix} 2 & 5 \\ 7 & \pi \end{bmatrix} = \begin{bmatrix} 2\sqrt{2} & 5\sqrt{2} \\ 7\sqrt{2} & \pi\sqrt{2} \end{bmatrix}.$$

Example 1.1.6

The vector space of real-valued continuous functions on the interval $[0, 1]$ is denoted by

$$\mathcal{C}([0, 1]) = \{f: [0, 1] \rightarrow \mathbb{R} : f \text{ is continuous on } [0, 1]\}.$$

Addition and scalar multiplication are defined by

$$(f + g)(x) = f(x) + g(x) \quad \text{and} \quad (\alpha f)(x) = \alpha(f(x)).$$

This example is a little trickier. Here, the ‘vectors’ are continuous functions from $[0, 1]$ to \mathbb{R} . When we add two functions, we get another function, and when we multiply a function by a scalar, we get another function. For example,

$$(\sin(x)) + (\cos(x)) = \sin x + \cos x \quad \text{and} \quad 3(x^2 + 1) = 3x^2 + 3.$$

Example 1.1.7

Here's a slightly more interesting example. Let V be the set of all lines in \mathbb{R}^2 with slope 1. Each such line has equation $y = x + d$ for some $d \in \mathbb{R}$. Addition and scalar multiplication in V are defined by

$$(y = x + d_1) + (y = x + d_2) = (y = x + (d_1 + d_2)) \quad \text{and} \quad \alpha(y = x + d) = (y = x + \alpha d).$$

Now that we have seen a few examples of vector spaces, you might have a perfectly reasonable question in mind: *What is a vector space?* Before we give the formal definition, let's take a look at the similarities between all of these examples.

They each consist of a set of objects called 'vectors' (even though sometimes these vectors can look a little unusual, like a straight line in \mathbb{R}^2 of slope 1) and some set of scalars (\mathbb{R} or \mathbb{C}). Furthermore, there is a way to 'add' two vectors to get another vector, and to 'multiply' a vector by a scalar to get another vector. There is some other structure lurking in the background which is perhaps a little harder to notice just from these examples. Indeed, in each vector space there is a special vector (call it $\vec{0}$, the zero vector) with the property that $\vec{0} + \vec{v} = \vec{v}$ for all vectors \vec{v} in the vector space.

Formally, we define a vector space as follows:

Definition 1.1.8

Vector Space Over \mathbb{F} , Vector Addition, Scalar Multiplication, Vector Space Axioms, Zero Vector, Additive Inverse

A **vector space over \mathbb{F}** is a set V together with an operation $+: V \times V \rightarrow V$ (**vector addition**) so that

$$\text{for all } \vec{x}, \vec{y} \in V, \vec{x} + \vec{y} \in V,$$

and an operation $\cdot: \mathbb{F} \times V \rightarrow V$ (**scalar multiplication**) so that

$$\text{for all } s \in \mathbb{F} \text{ and } \vec{x} \in V, s \cdot \vec{x} \in V.$$

These operations must satisfy the following properties:

1. For all $\vec{x}, \vec{y}, \vec{z} \in V$, $(\vec{x} + \vec{y}) + \vec{z} = \vec{x} + (\vec{y} + \vec{z})$;
2. There exists a vector $\vec{0} \in V$ such that, for all $\vec{x} \in V$, $\vec{0} + \vec{x} = \vec{x} + \vec{0} = \vec{x}$;
3. For all $\vec{x} \in V$, there exists a vector $-\vec{x} \in V$ such that $\vec{x} + (-\vec{x}) = (-\vec{x}) + \vec{x} = \vec{0}$;
4. For all $\vec{x}, \vec{y} \in V$, $\vec{x} + \vec{y} = \vec{y} + \vec{x}$;
5. For all $\vec{x} \in V$ and $s, t \in \mathbb{F}$, $s \cdot (t \cdot \vec{x}) = (st) \cdot \vec{x}$;
6. For all $\vec{x} \in V$ and $s, t \in \mathbb{F}$, $(s + t) \cdot \vec{x} = s \cdot \vec{x} + t \cdot \vec{x}$;
7. For all $\vec{x}, \vec{y} \in V$ and $s \in \mathbb{F}$, $s \cdot (\vec{x} + \vec{y}) = s \cdot \vec{x} + s \cdot \vec{y}$; and
8. For all $\vec{x} \in V$, $1 \cdot \vec{x} = \vec{x}$.

The above properties are called the **vector space axioms**.

The vector $\vec{0}$ in axiom 2 is called **the zero vector** of V . The vector $-\vec{x}$ in axiom 3 is called **the additive inverse** of \vec{x} . (See remarks 3 and 4 below.)

REMARKS

1. It is important to understand that the operations $+$ and \cdot above must be supplied as part of the definition of a vector space. You have the freedom to define any two functions $+: V \times V \rightarrow V$ and $\cdot: \mathbb{F} \times V \rightarrow V$ on any set V . However, you will only get a vector space if your functions satisfy the vector space axioms.

2. Be careful to not confuse the abstract scalar multiplication \cdot and the usual scalar multiplication in \mathbb{F} . For instance, in Axiom 5, on the left side we perform scalar multiplication by t and then by s ; while on the right side we perform scalar multiplication by the scalar st . In general, there is no reason why these should produce the same result. However, our intuition for a vector space suggests that they *should*, and this is why we formally require this property as an axiom.
3. In this section, we will always use \cdot for our abstract scalar multiplication. However, this gets cumbersome very quickly, and so in later sections we will omit it and simply write expressions like $a\vec{x}$ instead of $a \cdot \vec{x}$.
4. We will show below that, in a vector space V , there can be exactly one vector that satisfies axiom 2. It is therefore acceptable to call this vector *the* zero vector of V . Likewise, for each $\vec{x} \in V$, there will be exactly one vector $-\vec{x}$ that satisfies axiom 3. See Proposition 1.1.13.
5. The notation $-\vec{x}$ chosen for the additive inverse of \vec{x} is very suggestive. It resembles the scalar multiplication $(-1) \cdot \vec{x}$ of $-1 \in \mathbb{F}$ and $\vec{x} \in V$. However, since the definition above is very abstract and general, it is conceivable that $-\vec{x}$ and $(-1) \cdot \vec{x}$ might be different vectors. Happily, that is not the case! It turns out that the vector space axioms actually imply that $-\vec{x} = (-1) \cdot \vec{x}$. See Proposition 1.1.14.
6. The vector $\vec{x} + (-\vec{y})$ will usually be written as $\vec{x} - \vec{y}$.

Example 1.1.9

Let's check that \mathbb{R}^2 with the usual definitions of addition and scalar multiplication is a vector space over \mathbb{R} .

To check axiom 1, let $\vec{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$, $\vec{y} = \begin{bmatrix} y_1 \\ y_2 \end{bmatrix}$, and $\vec{z} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix}$ be three arbitrary vectors in \mathbb{R}^2 . Then

$$(\vec{x} + \vec{y}) + \vec{z} = \begin{bmatrix} x_1 + y_1 \\ x_2 + y_2 \end{bmatrix} + \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} x_1 + y_1 + z_1 \\ x_2 + y_2 + z_2 \end{bmatrix} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} y_1 + z_1 \\ y_2 + z_2 \end{bmatrix} = \vec{x} + (\vec{y} + \vec{z}),$$

so axiom 1 holds.

The vector $\vec{0} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ satisfies the properties of axiom 2. If $\vec{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$, then $-\vec{x} = \begin{bmatrix} -x_1 \\ -x_2 \end{bmatrix}$ satisfies $\vec{x} + (-\vec{x}) = (-\vec{x}) + \vec{x} = \vec{0}$, so 3 holds. For 7, let $s \in \mathbb{R}$ be an arbitrary scalar. Then

$$\begin{aligned} s \cdot (\vec{x} + \vec{y}) &= s \cdot \begin{bmatrix} x_1 + y_1 \\ x_2 + y_2 \end{bmatrix} \\ &= \begin{bmatrix} s(x_1 + y_1) \\ s(x_2 + y_2) \end{bmatrix} \\ &= \begin{bmatrix} sx_1 + sy_1 \\ sx_2 + sy_2 \end{bmatrix} \\ &= \begin{bmatrix} sx_1 \\ sx_2 \end{bmatrix} + \begin{bmatrix} sy_1 \\ sy_2 \end{bmatrix} \end{aligned}$$

$$= s \cdot \vec{x} + s \cdot \vec{y}.$$

We leave the rest of the axioms as an exercise for you to check.

Exercise 1 Prove that axioms 4, 5, 6, and 8 hold for \mathbb{R}^2 , implying that \mathbb{R}^2 is indeed a vector space (over \mathbb{R}).

Note: Solutions to these exercises can be found in Appendix A.

Exercise 2 Go through each of Examples 1.1.1–1.1.7 and convince yourself that each of them does in fact give a vector space (over the appropriate \mathbb{F}). In particular, what is the zero vector in Example 1.1.7?

Here is a not very interesting—but important!—example of a vector space.

Example 1.1.10 The **zero vector space** is the vector space $V = \{\vec{0}\}$ consisting of precisely one vector. The definitions of addition and scalar multiplication are:

$$\vec{0} + \vec{0} = \vec{0} \quad \text{and} \quad \alpha \cdot \vec{0} = \vec{0} \quad (\alpha \in \mathbb{F}).$$

It's very easy to check that these definitions satisfy axioms 1 through 8, with $\vec{0}$ (of course) being the zero vector.

Here are examples of things that are *not* a vector space.

Example 1.1.11 Let V be the set of polynomials with coefficients in \mathbb{C} of degree at most 1, with addition of vectors and scalar multiplication (by scalars in \mathbb{C}) given by

$$(a_0 + a_1x) + (b_0 + b_1x) = (a_0 + b_0) + (a_1 + b_1)x \quad \text{and} \quad \alpha \cdot (a_0 + a_1x) = (\alpha a_1) + (\alpha a_0)x$$

respectively. Notice that the addition is standard, but scalar multiplication has the coefficients swapped. Then V is *not* a vector space since $1 \cdot (3 + ix) = i + 3x$, so axiom 8 fails.

Example 1.1.12 Let V be the set of polynomials with coefficients in \mathbb{R} of degree at most 1, with addition of vectors and scalar multiplication (by scalars in \mathbb{R}) given by

$$(a_0 + a_1x) + (b_0 + b_1x) = (a_0 + a_1) + (b_0 + b_1)x \quad \text{and} \quad \alpha \cdot (a_0 + a_1x) = (\alpha a_0) + (\alpha a_1)x$$

respectively. This time scalar multiplication is standard, whereas the addition is not. Notice that V is *not* a vector space since

$$\begin{aligned} (3 + 2x) + (5 + 4x) &= (3 + 2) + (5 + 4)x \\ &\neq (5 + 4) + (3 + 2)x \\ &= (5 + 4x) + (3 + 2x), \end{aligned}$$

so axiom 4 fails.

Proposition 1.1.13

Let V be a vector space over \mathbb{F} . Then

- (a) The zero vector in V is *unique*. That is, if $\vec{0}_1 \in V$ satisfies the property that $\vec{x} + \vec{0}_1 = \vec{x}$ for all $\vec{x} \in V$, and if $\vec{0}_2 \in V$ satisfies the property that $\vec{x} + \vec{0}_2 = \vec{x}$ for all $\vec{x} \in V$, then $\vec{0}_1 = \vec{0}_2$.
- (b) Let $\vec{x} \in V$. The additive inverse of \vec{x} is *uniquely determined by \vec{x}* . That is, if \vec{y} satisfies the property that $\vec{x} + \vec{y} = \vec{y} + \vec{x} = \vec{0}$, then $\vec{y} = -\vec{x}$.

Proof: (a) If $\vec{0}_1$ is a zero vector (i.e. it satisfies axiom 2), then

$$\vec{0}_1 + \vec{0}_2 = \vec{0}_2.$$

Similarly, if $\vec{0}_2$ is a zero vector we have

$$\vec{0}_1 + \vec{0}_2 = \vec{0}_1.$$

The left sides of the preceding equations are the same, so the right sides must be too. Thus, $\vec{0}_1 = \vec{0}_2$, as desired.

- (b) Let $\vec{x}, \vec{y} \in V$ be such that $\vec{x} + \vec{y} = \vec{y} + \vec{x} = \vec{0}$. Since $\vec{x} \in V$, it follows from axiom 3 that there exists a vector $-\vec{x} \in V$ such that $\vec{x} + (-\vec{x}) = \vec{0}$. Now, the equations

$$\vec{x} + (-\vec{x}) = \vec{0} \quad \text{and} \quad \vec{x} + \vec{y} = \vec{0}$$

imply that

$$\vec{x} + (-\vec{x}) = \vec{x} + \vec{y}.$$

Adding \vec{y} to both sides,

$$\vec{y} + (\vec{x} + (-\vec{x})) = \vec{y} + (\vec{x} + \vec{y}).$$

By axiom 1, this is the same as

$$(\vec{y} + \vec{x}) + (-\vec{x}) = (\vec{y} + \vec{x}) + \vec{y}.$$

From the property given in (b), we know that $\vec{y} + \vec{x} = \vec{0}$. So the above equation is simply

$$\vec{0} + (-\vec{x}) = \vec{0} + \vec{y}.$$

Hence $-\vec{x} = \vec{y}$, by axiom 2. □

Part (a) of Proposition 1.1.13 allows us to unambiguously use the symbol $\vec{0}$ to denote the zero vector of a vector space. Here are some further basic properties of the zero vector and the additive inverse.

Proposition 1.1.14 Let V be a vector space over \mathbb{F} . Then

- (a) $0 \cdot \vec{x} = \vec{0}$ for all $\vec{x} \in V$,
- (b) $(-1) \cdot \vec{x} = -\vec{x}$ for all $\vec{x} \in V$, and
- (c) $t \cdot \vec{0} = \vec{0}$ for all $t \in \mathbb{F}$.

Exercise 3 Prove Proposition 1.1.14.

Notice that part (a) of the previous Proposition gives us a quick way of determining the zero vector of a given vector space: we simply scalar multiply any vector by the scalar zero. For instance, if V is the vector space from Example 1.1.7, and if, say, we consider the vector $y = x + 1$, then using the definition of scalar multiplication in V we find that

$$\vec{0} = 0 \cdot (y = x + 1) = (y = x).$$

Similarly, we can use part (b) to quickly determine the additive inverse of any given vector.

Of course, for this to be a useful strategy, we need to be sure that V is in fact a vector space. There are situations where we can be sure of this without having to check each of the 8 axioms.

1.2 Subspaces

We know that if we consider just the plane $\{[x \ y \ 0]^T : x, y \in \mathbb{R}\}$ consisting of only the xy -coordinates of points in \mathbb{R}^3 , then we can think of this as “a copy” of \mathbb{R}^2 living inside \mathbb{R}^3 . This is an example of a subspace of \mathbb{R}^3 . To make this idea precise, we first formally define a subspace.

Definition 1.2.1
Subspace Let V be a vector space over \mathbb{F} and $U \subseteq V$ a subset. We call U a **subspace** of V if U , endowed with the addition and scalar multiplication from V , is itself a vector space over \mathbb{F} .

Example 1.2.2 Every vector space V has two obvious subspaces: V itself and the subspace $\{\vec{0}\}$ consisting of the zero vector of V .

Example 1.2.3 Consider the subset $U \subseteq \mathcal{P}_2(\mathbb{F})$ given by $U = \{p \in \mathcal{P}_2(\mathbb{F}) : p(2) = 0\}$. First to get a feel for U , note that $p_1(x) = x^2 + x - 6$ is in U , but $p_2(x) = x^2$ is not in U , because

$$p_1(2) = 2^2 + 2 - 6 = 0 \quad \text{and} \quad p_2(2) = 2^2 = 4 \neq 0.$$

We claim that U is a subspace of $\mathcal{P}_2(\mathbb{F})$. Let's check some of the axioms to convince ourselves.

First we have to check that the addition and scalar multiplication from $\mathcal{P}_2(\mathbb{F})$ make sense as addition and scalar multiplication in U . That is, we have to make sure that if we take two vectors in U and add them together, we get a vector in U , and that every scalar multiple of a vector in U is in U .

Suppose $p, q \in U$ and $\alpha \in \mathbb{F}$. Then $(p+q)(2) = p(2) + q(2) = 0$ so $p+q \in U$. Furthermore, $(\alpha p)(2) = \alpha p(2) = 0$ so $\alpha p \in U$. Thus, addition and scalar multiplication make sense on U .

Since the addition and scalar multiplication on U is simply that from $\mathcal{P}_2(\mathbb{F})$, and $\mathcal{P}_2(\mathbb{F})$ is a vector space, axioms 1, 4, 5, 6, 7, and 8 obviously hold for U . Since that the zero vector $\vec{0} = 0x^2 + 0x + 0$ of $\mathcal{P}_2(\mathbb{F})$ is in U , we deduce that axiom 2 is satisfied. Finally, by part (b) of Proposition 1.1.14, $-p = (-1)p \in U$, so axiom 3 is satisfied. We may finally conclude that U is a vector space.

Checking that addition and scalar multiplication make sense on U and checking all 8 axioms is a little cumbersome. However, if you carefully examine the previous example, a lot of things came for free from the fact that $\mathcal{P}_2(\mathbb{F})$ was already a vector space. The next theorem allows us never to have to do that much work again, and simply check three things to check whether or not a subset of a vector space is a subspace or not.

Theorem 1.2.4 (The Subspace Test)

Let V be vector space over \mathbb{F} and let U be a subset of V . Then U is a subspace of V if and only if the following three conditions hold.

- (a) U is non-empty.
- (b) For all $\vec{u}_1, \vec{u}_2 \in U$, $\vec{u}_1 + \vec{u}_2 \in U$. (We say that U is **closed under addition**.)
- (c) For all $\alpha \in \mathbb{F}$ and for all $\vec{u} \in U$, $\alpha\vec{u} \in U$. (We say that U is **closed under scalar multiplication**.)

Proof: If U is a subspace, then (b) and (c) hold as part of being a definition of a subspace, and since all vector spaces have a zero vector, U must be non-empty.

Conversely, suppose (a), (b) and (c) hold for a subset U of V . Properties (b) and (c) imply that the addition and scalar multiplication from V restrict to addition and scalar multiplication on U . Axioms 1,4,5,6,7, and 8 hold since V is a vector space. For axiom 2, since U is non-empty, choose a vector $\vec{u} \in U$ and then note by Proposition 1.1.14, $0\vec{u} = \vec{0}$. Property (c) then implies that $\vec{0}$ is in U . Similarly, for axiom 3 let $\vec{u} \in U$. Then by Proposition 1.1.14 and property (c), $-\vec{u} = (-1)\vec{u} \in U$, completing the proof. \square

Example 1.2.5 Prove that $U = \{p \in \mathcal{P}_2(\mathbb{F}) : p(2) = 0\}$ is a subspace of $\mathcal{P}_2(\mathbb{F})$.

Proof: By the Subspace Test, we only need to check three things:

- (a) Since $\vec{0} = 0x^2 + 0x + 0 \in U$, U is non-empty.

- (b) Let $p, q \in U$. Then $(p + q)(2) = p(2) + q(2) = 0$, so $p + q \in U$. Thus, U is closed under addition.
- (c) Let $p \in U$ and $\alpha \in \mathbb{R}$. Then $(\alpha p)(2) = \alpha p(2) = 0$ so $\alpha p \in U$. Thus, U is closed under scalar multiplication.

Therefore, by the Subspace Test, U is a subspace of $\mathcal{P}_2(\mathbb{F})$. \square

It is natural to ask now what kind of things aren't subspaces. If you study the proof of The Subspace Test, you will see that a subspace of a vector space V must contain the zero vector of V . So any subset of V that doesn't contain the zero vector of V cannot possibly be a subspace of V . We state this result formally in the corollary below, where we also address a related subtlety: a subspace U of a vector space V is itself a vector space, so it has a zero vector by definition. Could it be that this zero vector in U is different from the zero vector in V ?

Corollary 1.2.6

Let V be a vector space over \mathbb{F} and suppose that U is a subspace of V . Let $\vec{0}_V$ and $\vec{0}_U$ denote the zero vectors in V and U , respectively. Then $\vec{0}_U = \vec{0}_V$. In particular, the zero vector in V is in U : $\vec{0}_V \in U$.

Exercise 4

Prove Corollary 1.2.6.

This corollary allows us to write $\vec{0}$ unambiguously if we're working inside a fixed vector space V .

Example 1.2.7

Let $S = \{p(x) \in \mathcal{P}_2(\mathbb{F}) : p(2) = 1\}$. Then S is not a subspace of $\mathcal{P}_2(\mathbb{F})$ because it does not contain $\vec{0} = 0 + 0x + 0x^2$, the zero vector of $\mathcal{P}_2(\mathbb{F})$.

Of course a subset of a vector space may contain $\vec{0}$ yet fail to be a subspace.

Example 1.2.8

Let \mathbb{Z} denote the set of all integers, and consider the subset

$$L = \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in M_{2 \times 2}(\mathbb{F}) : a, b, c, d \in \mathbb{Z} \right\}$$

of $M_{2 \times 2}(\mathbb{F})$. This is not a subspace of $M_{2 \times 2}(\mathbb{F})$ since $\frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \notin L$ whereas $\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \in L$. That is, L is not closed under scalar multiplication.

Now that we have studied some examples, an interesting question to think about is how subspaces can be created. One way is to take a subset of vectors in your vector space, and then throw in everything else that needs to be there to make that subset a subspace! This is the same process that you have seen in \mathbb{F}^n . The following definition uses familiar terminology.

Definition 1.2.9**Span, Linear
Combination**

Let $S = \{\vec{v}_1, \dots, \vec{v}_k\}$ be a subset of a vector space V over \mathbb{F} . Define the **span** of S by

$$\text{Span}(S) = \{t_1 \vec{v}_1 + \dots + t_k \vec{v}_k : t_1, \dots, t_k \in \mathbb{F}\}.$$

A vector of the form $t_1 \vec{v}_1 + \dots + t_k \vec{v}_k$ is called a **linear combination** of the vectors $\vec{v}_1, \dots, \vec{v}_k$.

By convention, we define the span of the empty set to be the set consisting of the zero vector: $\text{Span } \emptyset = \{\vec{0}\}$.

REMARK

Notice that the previous definition only applies to finite subsets of V . If S is an infinite subset of V , then the span of S is defined to be the union of the spans of all finite subsets of S . Equivalently, the span of S is the set of all linear combinations of all finite collections of vectors in S .

In this course, we will not be making use of this more general definition. However, in more advanced treatments of linear algebra, this generalization plays an important role.

Example 1.2.10

In $M_{2 \times 2}(\mathbb{F})$, let $S = \left\{ \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix} \right\}$. Then $\begin{bmatrix} 2 & 2 \\ -3 & -3 \end{bmatrix} \in \text{Span}(S)$, since

$$\begin{bmatrix} 2 & 2 \\ -3 & -3 \end{bmatrix} = 2 \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} + (-3) \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix}.$$

On the other hand, $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \notin \text{Span}(S)$, since there are no $a, b \in \mathbb{F}$ such that

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

as you can check.

Example 1.2.11

In $\mathcal{P}_1(\mathbb{F})$, let $S = \{1 + x, 1 - x\}$. Then

$$\text{Span}(S) = \{a(1 + x) + b(1 - x) : a, b \in \mathbb{F}\}.$$

There are many different descriptions of $\text{Span}(S)$. For instance, we claim that

$$\text{Span}(S) = \text{Span}(B), \quad \text{where } B = \{1, x\}.$$

Indeed, since

$$a(1 + x) + b(1 - x) = (a + b) \cdot 1 + (a - b) \cdot x,$$

we see that every element $a(1 + x) + b(1 - x)$ of $\text{Span}(S)$ is also an element of $\text{Span}(B)$, i.e., $\text{Span}(S) \subseteq \text{Span}(B)$. Conversely, notice that

$$c + dx = \frac{c + d}{2}(1 - x) + \frac{c - d}{2}(1 + x),$$

(Where did this seemingly magical expression come from? You may want to review how to solve systems of linear equations! See the Exercise following this example.) This shows that every element $c + dx$ of $\text{Span}(B)$ is also an element of $\text{Span}(S)$, i.e., $\text{Span}(B) \subseteq \text{Span}(S)$. This proves that the sets $\text{Span}(S)$ and $\text{Span}(B)$ are equal.

Exercise 5

In the example above, behind the scenes we ended up solving two systems of linear equations. When proving $\text{Span}(S) \subseteq \text{Span}(B)$, we took an arbitrary element $a(1 + x) + b(1 - x)$ in $\text{Span}(S)$ and solved the equation

$$a(1 + x) + b(1 - x) = c \cdot 1 + d \cdot x$$

for c and d . When proving $\text{Span}(B) \subseteq \text{Span}(S)$, we took an arbitrary element $c + dx$ in $\text{Span}(B)$ and solved the equation

$$c + dx = a(1 + x) + b(1 - x)$$

for a and b . Can you write down the associated systems of linear equations in matrix form and then solve them?

Let's prove now that taking the span of some vectors does actually result in a subspace. The proof is exactly as it was for \mathbb{F}^n . This is a common theme of the subject. Any result that we can prove for \mathbb{F}^n using only the vector space structure of \mathbb{F}^n can usually be carried over word for word to the setting of an abstract vector space over \mathbb{F} . Indeed, this is one motivation for even defining abstract vector spaces! We can now prove one result that simultaneously applies to a wide variety of different-looking spaces, such as \mathbb{F}^n , $M_{m \times n}(\mathbb{F})$, $\mathcal{P}_n(\mathbb{F})$, etc.

Proposition 1.2.12

Let $S = \{\vec{v}_1, \dots, \vec{v}_k\}$ be a subset of a vector space V . Then $\text{Span}(S)$ is a subspace of V .

Proof: Since $\vec{0} = 0\vec{v}_1 + \dots + 0\vec{v}_k$, $\vec{0} \in \text{Span}(S)$ so $\text{Span}(S)$ is non-empty. Suppose $\vec{x}, \vec{y} \in \text{Span}(S)$, and let $\vec{x} = t_1\vec{v}_1 + \dots + t_k\vec{v}_k$ and $\vec{y} = s_1\vec{v}_1 + \dots + s_k\vec{v}_k$ for elements $t_1, \dots, t_k, s_1, \dots, s_k \in \mathbb{F}$. Then

$$\vec{x} + \vec{y} = (t_1 + s_1)\vec{v}_1 + \dots + (t_k + s_k)\vec{v}_k$$

so $\vec{x} + \vec{y} \in \text{Span}(S)$ and hence $\text{Span}(S)$ is closed under addition. Finally, let $\vec{x} \in \text{Span}(S)$ be as above, and let $\alpha \in \mathbb{F}$. Then $\alpha\vec{x} = (\alpha t_1)\vec{v}_1 + \dots + (\alpha t_k)\vec{v}_k$ and since $\alpha t_i \in \mathbb{F}$ for all i , $\alpha\vec{x} \in \text{Span}(S)$ and hence $\text{Span}(S)$ is closed under scalar multiplication. Therefore, by the Subspace Test, $\text{Span}(S)$ is a subspace of V . \square

1.3 Bases and Dimension

We now shift our focus to formalizing the notion of dimension. Intuitively we know that \mathbb{R}^2 is a 2-dimensional space, because there are 2 different directions one can travel in, and no more. We may also have an idea that \mathbb{R}^2 is 2-dimensional since every vector is determined by 2 pieces of information (the x and y coordinates). Similarly, we may guess that \mathbb{R}^n would be

an n -dimensional vector space, and we would be correct! However, this geometric intuition fails us when thinking about other vector spaces. For example, what is the dimension of \mathbb{C}^2 , or $\mathcal{P}_3(\mathbb{R})$, or $\mathcal{C}([0, 1])$?

As you've learned in a previous course, the key to defining a useful notion of dimension is to first define "basis." The definition of basis for \mathbb{F}^n carries over without change to the setting of an abstract vector space. Recall that a basis for \mathbb{F}^n is a linearly independent spanning set for \mathbb{F}^n . Thus, we must begin by defining these concepts in this new abstract setting.

1.3.1 Spanning Sets, Linear Independence, and Bases

Definition 1.3.1

Spanning Set,
Spans

A set of vectors $S = \{\vec{v}_1, \dots, \vec{v}_k\}$ in a vector space V is a **spanning set** for V if

$$\text{Span}(S) = V.$$

We also say that S **spans** V .

Intuitively, a set of vectors span a vector space if every vector in that vector space can be obtained from those vectors. More precisely, every vector in the vector space can be expressed as linear combination of those from the spanning set.

Example 1.3.2

The set $B = \{1, x\}$ is a spanning set for $\mathcal{P}_1(\mathbb{F})$. Indeed, by definition

$$\mathcal{P}_1(\mathbb{F}) = \{a + bx : a, b \in \mathbb{F}\} = \text{Span}(B).$$

In Example 1.2.11, we effectively showed that $S = \{1 - x, 1 + x\}$ is also a spanning set for $\mathcal{P}_1(\mathbb{F})$.

It is easy to see that $T = \{3 + 2x\}$ is not a spanning set for $\mathcal{P}_1(\mathbb{F})$. For instance, there is no way of writing the polynomial x as a multiple of $3 + 2x$.

A spanning set can sometimes have redundant information. For example, the sets

$$\left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\} \quad \text{and} \quad \left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}$$

are both spanning sets for \mathbb{R}^2 , but the vector $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ in the first set is redundant since it is a linear combination of the other two vectors. To formalize this, we introduce the notion of linear independence.

Definition 1.3.3

Linearly
Independent,
Linearly
Dependent

A set of vectors $\{\vec{v}_1, \dots, \vec{v}_k\}$ in a vector space V is **linearly independent** if the only solution to the equation

$$t_1 \vec{v}_1 + \dots + t_k \vec{v}_k = \vec{0}$$

is $t_1 = \dots = t_k = 0$. The set is **linearly dependent** otherwise.

By convention, the empty set \emptyset is linearly independent.

REMARK

We can extend the definition of linear independence to infinite subsets of V by defining such a set to be linearly independent if all of its finite subsets are linearly independent. Just as for spanning sets, this more general definition will not play a role in our course.

Although this is the formal definition we are to work with, the intuition is that a linearly independent set is a set of vectors that all “point in different directions” so that no vector in the set is “redundant.” The next result elaborates on this idea.

Proposition 1.3.4

A subset $S = \{\vec{v}_1, \dots, \vec{v}_k\}$ of a vector space V is linearly dependent if and only if at least one vector in S is a linear combination of other vectors in S .

Proof: If S is linearly dependent, then there is a solution to

$$t_1 \vec{v}_1 + \dots + t_k \vec{v}_k = \vec{0}$$

with some $t_i \neq 0$. But then by moving the i th term on the left to the right-side of the above equation, we get

$$\sum_{j \neq i} t_j \vec{v}_j = (-t_i) \vec{v}_i$$

and then by dividing both sides by $-t_i$ (which is non-zero!), we’ve expressed \vec{v}_i as a linear combination of the other vectors in S .

Conversely, assume that some vector in S , say \vec{v}_i is a linear combination of some other vectors, say

$$\vec{v}_i = \sum_{j \neq i} c_j \vec{v}_j$$

(we can assume that *all* of the other vectors of S appear on the right-side by letting $c_j = 0$ if needed). Then by moving \vec{v}_i to the right, we get

$$c_1 \vec{v}_1 + \dots + (-1) \vec{v}_i + \dots + c_k \vec{v}_k = 0.$$

Thus we’ve found a solution to

$$t_1 \vec{v}_1 + \dots + t_k \vec{v}_k = \vec{0}$$

with $t_i = -1 \neq 0$, which shows that S is linearly dependent. □

Example 1.3.5

The set $\{1 + x, 1\}$ is linearly independent in $\mathcal{P}_1(\mathbb{C})$. To see this, set

$$0 = t_1(1 + x) + t_2(1) = (t_1 + t_2) + t_1x.$$

Then equating the x coefficient gives us $t_1 = 0$, which then implies $t_2 = 0$. Therefore the only solution is $t_1 = t_2 = 0$, so the set is linearly independent.

Example 1.3.6

Since in \mathbb{R}^2 ,

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

the set $\left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \end{bmatrix} \right\}$ is linearly dependent.

Sometimes it's not so easy to stare at a set of vectors and decide whether or not they are linearly independent. However, we do have some tools for solving simultaneous equations from a previous course to help us along the way!

Example 1.3.7

Is $\{x + x^2 - 2x^3, 2x - x^2 + x^3, x + 5x^2 + 3x^3\}$ linearly independent in $\mathcal{P}_3(\mathbb{R})$?

Solution:

To check, we want to solve the equation

$$a(x + x^2 - 2x^3) + b(2x - x^2 + x^3) + c(x + 5x^2 + 3x^3) = 0$$

for a, b, c . Equating coefficients gives us the system of simultaneous equations

$$\begin{aligned} a + 2b + c &= 0 \\ a - b + 5c &= 0 \\ -2a + b + 3c &= 0. \end{aligned}$$

To solve such a system of equations, we plug the coefficients into an augmented matrix and row reduce! We get

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

Therefore the system of equations has exactly one solution, and that solution is $a = b = c = 0$. Therefore the set is linearly independent.

If we have a spanning set that is linearly independent, then in some sense our spanning set is not redundant. Such sets are very special and deserve a name.

Definition 1.3.8

Basis

A **basis** for a vector space V is a linearly independent subset that spans V .

Theorem 1.3.9

Every vector space V has a basis.

This result is mostly of theoretical interest. In practice, almost all of the vector spaces in this course will have bases that we can easily discover and write down explicitly (see Section 1.3.3). The proof of Theorem 1.3.9 in full generality is actually quite intricate and requires ideas beyond the scope of our course. On page 25 (see Item 2), we give a proof of this result in the special case where V is known to have a finite spanning set, which will be the case for the overwhelming majority of our examples.

Example 1.3.10

Here are some examples of bases:

- $\left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \dots, \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix} \right\}$ is the **standard basis** for \mathbb{F}^n .
- $\{1, x, x^2, \dots, x^n\}$ is the **standard basis** for $\mathcal{P}_n(\mathbb{F})$.
- $\{E_{11}, E_{12}, \dots, E_{ij}, \dots, E_{nm}\}$, where E_{ij} is the $m \times n$ matrix with an entry of 1 in the (i, j) th position and 0s elsewhere, is the **standard basis** for $M_{m \times n}(\mathbb{F})$.
- The empty set \emptyset is a basis for the zero vector space $\{\vec{0}\}$.

Example 1.3.11

Is $S = \left\{ \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \right\}$ a basis for $M_{2 \times 2}(\mathbb{R})$?

Solution:

To check, we want to verify whether S is a spanning set for $M_{2 \times 2}(\mathbb{R})$ and whether S is linearly independent.

First, we check whether $\text{Span}(S) = M_{2 \times 2}(\mathbb{R})$. Certainly, $\text{Span}(S) \subseteq M_{2 \times 2}(\mathbb{R})$, so we only need to demonstrate the reverse inclusion $M_{2 \times 2}(\mathbb{R}) \subseteq \text{Span}(S)$. Let $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ be an arbitrary matrix in $M_{2 \times 2}(\mathbb{R})$, and consider the equation

$$a_1 \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + a_2 \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} + a_3 \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} + a_4 \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix}.$$

By adding up the matrices on the left-side and then equating entries with the right-side, we obtain the system of simultaneous equations

$$\begin{aligned} a_1 + a_3 &= a \\ a_2 - a_4 &= b \\ a_2 + a_4 &= c \\ a_1 - a_3 &= d. \end{aligned}$$

To solve this system of equations, we plug the coefficients into an augmented matrix and row reduce. We get

$$\left[\begin{array}{cccc|c} 1 & 0 & 1 & 0 & a \\ 0 & 1 & 0 & -1 & b \\ 0 & 1 & 0 & 1 & c \\ 1 & 0 & -1 & 0 & d \end{array} \right] \longrightarrow \left[\begin{array}{cccc|c} 1 & 0 & 0 & 0 & \frac{a+d}{2} \\ 0 & 1 & 0 & 0 & \frac{b+c}{2} \\ 0 & 0 & 1 & 0 & \frac{a-d}{2} \\ 0 & 0 & 0 & 1 & \frac{c-b}{2} \end{array} \right].$$

We conclude that

$$\frac{a+d}{2} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \frac{b+c}{2} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} + \frac{a-d}{2} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} + \frac{c-b}{2} \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix},$$

which means that $M_{2 \times 2}(\mathbb{R}) \subseteq \text{Span}(S)$. Hence S is a spanning set for $M_{2 \times 2}(\mathbb{R})$.

To check that S is linearly independent, we consider the equation

$$a_1 \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + a_2 \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} + a_3 \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} + a_4 \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$$

which, as can be seen from the RREF above, has exactly one solution, namely $a_1 = a_2 = a_3 = a_4 = 0$. Hence S linearly independent, and so it must be a basis.

In the following subsection we will learn about a more efficient method of checking whether a particular subset S of a vector space V is a basis (provided we know the dimension of V beforehand); see Example 1.3.17.

Exercise 6

Verify that the sets below are bases for the indicated vector spaces.

(a) $\left\{ \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right\}$ is a basis for \mathbb{R}^3 .

(b) $\left\{ \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right\}$ is a basis for \mathbb{C}^3 .

(c) $\left\{ \begin{bmatrix} 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 3i \end{bmatrix} \right\}$ is a basis for \mathbb{C}^2 .

(d) $\{1 - x, 1 + x\}$ is a basis for $\mathcal{P}_1(\mathbb{R})$.

1.3.2 Dimension

Just as for \mathbb{F}^n , we will define the dimension of a vector space V to be the number of vectors in a basis for V . For this to make sense, we must first prove that all bases have the same size. As you examine our proof below, you should compare it to the proof of the same fact for \mathbb{F}^n that you may have seen in a previous course.

Lemma 1.3.12

Let V be a vector space over \mathbb{F} and suppose that $V = \text{Span}(\{\vec{v}_1, \dots, \vec{v}_n\})$. If $\{\vec{u}_1, \dots, \vec{u}_k\}$ is a linearly independent set in V , then $k \leq n$.

Proof: Since $\text{Span}(\{\vec{v}_1, \dots, \vec{v}_n\}) = V$, we have

$$\begin{aligned} \vec{u}_1 &= a_{11}\vec{v}_1 + \cdots + a_{1n}\vec{v}_n \\ &\vdots \\ \vec{u}_k &= a_{k1}\vec{v}_1 + \cdots + a_{kn}\vec{v}_n \end{aligned}$$

where $a_{ij} \in \mathbb{F}$ for all i and j . We will now aim to show that if $k > n$, then there is a solution to $t_1 \vec{u}_1 + \cdots + t_k \vec{u}_k = \vec{0}$ where not all the t_i are 0. We have

$$\begin{aligned} t_1 \vec{u}_1 + \cdots + t_k \vec{u}_k &= t_1(a_{11} \vec{v}_1 + \cdots + a_{1n} \vec{v}_n) + \cdots + t_k(a_{k1} \vec{v}_1 + \cdots + a_{kn} \vec{v}_k) \\ &= (a_{11}t_1 + a_{21}t_2 + \cdots + a_{k1}t_k) \vec{v}_1 + \cdots + (a_{1n}t_1 + \cdots + a_{kn}t_k) \vec{v}_n. \end{aligned}$$

Now, if $k > n$ the system of linear equations

$$\begin{aligned} a_{11}t_1 + \cdots + a_{k1}t_k &= 0 \\ &\vdots \\ a_{1n}t_1 + \cdots + a_{kn}t_k &= 0 \end{aligned}$$

has a solution where not all the t_i are 0. Consider such a solution. We then have

$$\begin{aligned} \vec{0} &= 0\vec{v}_1 + \cdots + 0\vec{v}_n \\ &= (a_{11}t_1 + \cdots + a_{k1}t_k) \vec{v}_1 + \cdots + (a_{1n}t_1 + \cdots + a_{kn}t_k) \vec{v}_n \\ &= t_1 \vec{u}_1 + \cdots + t_k \vec{u}_k \end{aligned}$$

contradicting the assumption that $\{\vec{u}_1, \dots, \vec{u}_k\}$ is linearly independent. So $k \leq n$. \square

Theorem 1.3.13

Suppose $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$ and $\mathcal{C} = \{\vec{u}_1, \dots, \vec{u}_k\}$ are both bases of a vector space V . Then $k = n$.

Proof: Since \mathcal{B} spans V and \mathcal{C} is linearly independent, $k \leq n$. However, since \mathcal{C} spans V and \mathcal{B} is linearly independent, $n \leq k$. Thus, $k = n$. \square

Definition 1.3.14

**Dimension,
Finite-dimensional,
Infinite-dimensional**

We say that a vector space V is **finite-dimensional** if it has a finite basis. Otherwise, we say that V is **infinite-dimensional**.

If V is finite-dimensional, we define its **dimension** to be the size of any finite basis of V . We denote this by $\dim(V)$.

If V is infinite-dimensional, we write $\dim(V) = \infty$.

Note that Theorem 1.3.13 shows that $\dim(V)$ is well-defined if V is finite-dimensional: it doesn't matter what finite basis of V you pick—they all contain the same number of elements.

With the definition of dimension at our disposal, we can now talk about dimension with conviction! Here are four important examples:

- $\dim(\{\vec{0}\}) = 0$ since by convention \emptyset is a basis for $\{\vec{0}\}$.
- $\dim(\mathbb{F}^n) = n$ since the standard basis has size n .
- $\dim(\mathcal{P}_n(\mathbb{F})) = n + 1$ since the standard basis has size $n + 1$.
- $\dim(M_{m \times n}(\mathbb{F})) = mn$ since the standard basis has size mn .

The vector space $\mathcal{C}([0, 1])$ is infinite-dimensional, and we would write $\dim(\mathcal{C}([0, 1])) = \infty$. To convince yourself of this, notice that $\mathcal{C}([0, 1])$ contains all polynomials of arbitrarily high degree. Since no finite set of polynomials can generate all polynomials of arbitrary degree, there can be no finite basis.

Example 1.3.15

Let $U = \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in M_{2 \times 2}(\mathbb{F}) : a + b + c + d = 0 \right\}$. It is an exercise for you to check that U is a subspace of $M_{2 \times 2}(\mathbb{F})$. We will now compute the dimension of U by finding a basis for U .

Note that every matrix in U is of the form $\begin{bmatrix} a & b \\ c & -a - b - c \end{bmatrix}$, so we can write every matrix in U as

$$a \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} + b \begin{bmatrix} 0 & 1 \\ 0 & -1 \end{bmatrix} + c \begin{bmatrix} 0 & 0 \\ 1 & -1 \end{bmatrix}.$$

Thus, $\left\{ \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 0 & -1 \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 1 & -1 \end{bmatrix} \right\}$ is a spanning set for U .

(Notice that this actually proves that U is a subspace of $M_{2 \times 2}(\mathbb{F})$! Indeed, we have just shown that U is equal to a span of some vectors in $M_{2 \times 2}(\mathbb{F})$. Now apply Proposition 1.2.12.)

We now check to see whether the above spanning set is linearly independent. Consider

$$t_1 \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} + t_2 \begin{bmatrix} 0 & 1 \\ 0 & -1 \end{bmatrix} + t_3 \begin{bmatrix} 0 & 0 \\ 1 & -1 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

By comparing the top-left entries, we get $t_1 = 0$; the top-right entries give $t_2 = 0$; and the bottom-left give $t_3 = 0$. Therefore $\left\{ \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 0 & -1 \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 1 & -1 \end{bmatrix} \right\}$ is linearly independent and hence is a basis for U .

Thus, $\dim(U) = 3$.

The next theorem is extremely useful in thinking about dimension. It formally proves things you already know in your heart to be true. Things like “You cannot have 4 linearly independent vectors in \mathbb{R}^3 , there’s just not enough space!” and “You can’t span $M_{2 \times 2}(\mathbb{C})$ with only 3 vectors, that’s not enough because $\dim(M_{2 \times 2}(\mathbb{C})) = 4!$ ” A sketch of the proof is provided, and you should fill in the details as an exercise. (If you get stuck, you can look at the proof in the case of $V = \mathbb{F}^n$ for inspiration.)

Theorem 1.3.16

Let V be an n -dimensional vector space over \mathbb{F} . Then

- (a) A set of more than n vectors in V must be linearly dependent.
- (b) A set of fewer than n vectors in V cannot span V .
- (c) A set with exactly n vectors in V is a spanning set for V if and only if it is linearly independent.

Proof: Parts (a) and (b) are restatements of Lemma 1.3.12. Statement (c) follows from the two paragraphs in Section 1.3.3. \square

In view of this result we can now introduce a much more efficient procedure for verifying whether a given set is a basis of a vector space whose dimension is known.

Example 1.3.17

Let's revisit Example 1.3.11 and demonstrate a different approach for showing that

$$S = \left\{ \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \right\}$$

is a basis for $M_{2 \times 2}(\mathbb{R})$.

Since $M_{2 \times 2}(\mathbb{R})$ is 4-dimensional, every basis for it must contain exactly 4 vectors. Since this is the case for S , it is a potential candidate for being a basis. In view of Theorem 1.3.16 (c), S will be a spanning set for $M_{2 \times 2}(\mathbb{R})$ if and only if it is linearly independent. Thus, it suffices to check *either* that $\text{Span}(S) = M_{2 \times 2}(\mathbb{R})$ *or* that S is linearly independent; once we check one, we get the other one for free.

It's usually easier to check for linear independence. We've already shown in Example 1.3.11 that S is linearly independent. It then follows, by what we've said in the preceding paragraph, that S must be a basis for $M_{2 \times 2}(\mathbb{R})$.

Exercise 7

Let $S = \{p_0(x), p_1(x), p_2(x)\}$ be a subset of $P_2(\mathbb{F})$. Show that if $\deg(p_i(x)) = i$, then S is a basis for $P_2(\mathbb{F})$.

Finally, we state an unsurprising result relating the dimension of a subspace to the dimension of the vector space that contains it.

Theorem 1.3.18

Let V be a finite-dimensional vector space over \mathbb{F} and let W be a subspace of V . Then W is finite-dimensional. Moreover, $\dim(W) \leq \dim(V)$ with equality if and only if $W = V$.

Proof: Since any linearly independent set in W is linearly independent in V as well, it follows from Lemma 1.3.12 (applied to the case where $\{\vec{v}_1, \dots, \vec{v}_n\}$ is a basis for V) that there can be no more than $n = \dim(V)$ distinct linearly independent vectors in W . In particular, the size of any basis for W cannot exceed $\dim(V)$. This simultaneously shows that W is finite-dimensional and that $\dim(W) \leq \dim(V)$.

Suppose now that $\dim(W) = \dim(V)$. Then according to Theorem 1.3.16(c), a basis \mathcal{B} for W will automatically be a basis for V , since it is a linearly independent set of size $\dim(V)$. It follows that $V = \text{Span}(\mathcal{B}) = W$. Conversely, if $W = V$ then of course $\dim(W) = \dim(V)$. \square

1.3.3 Obtaining Bases

There are many ways you could find a basis for a finite-dimensional vector space. Here are a couple of important ways.

1. **Extending a linearly independent subset.** Suppose you have a linearly independent subset $\{\vec{v}_1, \dots, \vec{v}_k\}$ in a finite dimensional vector space V . If it is a spanning set, then you have a basis. If not, choose a vector \vec{v}_{k+1} not in the span of $\{\vec{v}_1, \dots, \vec{v}_k\}$. Then $\{\vec{v}_1, \dots, \vec{v}_{k+1}\}$ must be linearly independent (by Proposition 1.3.4). If this new set spans, then it's a basis. If not, then repeat. This process must eventually stop since our vector space is finite-dimensional, and you will be left with a basis containing $\{\vec{v}_1, \dots, \vec{v}_k\}$.
2. **Reducing an arbitrary finite spanning set.** Suppose you have a finite spanning set $\{\vec{v}_1, \dots, \vec{v}_k\}$ for your vector space, and let's assume that it doesn't contain $\vec{0}$. If it is linearly independent, it is a basis! If not, you can write one of the vectors in the set, say v_i , as a linear combination of the others (by Proposition 1.3.4). Now $\text{Span}(\{\vec{v}_1, \dots, \vec{v}_k\}) = \text{Span}(\{\vec{v}_1, \dots, \vec{v}_{i-1}, \vec{v}_{i+1}, \dots, \vec{v}_k\})$, so $\{\vec{v}_1, \dots, \vec{v}_{i-1}, \vec{v}_{i+1}, \dots, \vec{v}_k\}$ spans our vector space. If this new set is linearly independent, then it is a basis! If not, repeat to remove another vector. This process must eventually stop since we started with finitely many vectors in our spanning set. The final product will be a basis made up entirely out of vectors from our original spanning set.

Here's an example of how method 1 above can be applied.

Example 1.3.19

Suppose we wish to find a basis for $\mathcal{P}_3(\mathbb{R})$ containing the polynomials $1+x$ and $1+x^2$. Notice first of all that $S = \{1+x, 1+x^2\}$ is linearly independent, so we are actually able to accomplish this.

Since $\dim(\mathcal{P}_3(\mathbb{R})) = 4$, must need to add two more polynomials to S in such a way that the resulting set is still linearly independent. Where can we find such polynomials? We have the standard basis vectors $1, x, x^2$ and x^3 , so we can try adding them in one at a time and checking linear independence.

If we add in 1 , the resulting set $T = \{1+x, 1+x^2, 1\}$ will be linearly independent, as you can check (see also the Exercise after Example 1.3.17). So we just need to add one more vector to T .

If we add in x , we obtain a linearly dependent set $\{1+x, 1+x^2, 1, x\}$, since $x = (1+x) + (-1)1$ is a linear combination of two other vectors in the set. Similarly, x^2 is no good. However, if we add in x^3 , then the resulting set $B = \{1+x, 1+x^2, 1, x^3\}$ is linearly independent, and hence (because it contains $4 = \dim(\mathcal{P}_3(\mathbb{R}))$ vectors) is a basis for $\mathcal{P}_3(\mathbb{R})$.

In practice, most of our vector spaces will be given to us as subspaces of a familiar vector space V , like in Examples 1.2.5 and 1.3.15, where we were given a subspace U of V defined by some conditions. In such situations we can attempt to express the defining conditions for U in terms of a known basis for V (such as the "standard basis" in case V is one of \mathbb{F}^n , $\mathcal{P}_n(\mathbb{F})$ or $M_{m \times n}(\mathbb{F})$). This will allow us to determine a spanning set for U , after which we can apply method 2 above (if this spanning set isn't already a basis).

Example 1.3.20

To give an illustration of how this works, let's try to find a basis for the subspace

$$U = \{p \in \mathcal{P}_2(\mathbb{F}) : p(2) = 0\}$$

of $\mathcal{P}_2(\mathbb{F})$ from Example 1.2.5. First we must express the condition $p(2) = 0$ using the standard basis $\{1, x, x^2\}$ of $\mathcal{P}_2(\mathbb{F})$. We can write p as $p(x) = a + bx + cx^2$, in which case the condition $p(2) = 0$ becomes

$$a + 2b + 4c = 0.$$

This is a linear equation in three variables, and we can employ our usual method for solving it in terms of basic variables and free parameters. We can do this quickly: simply notice that $a = -2b - 4c$! Consequently,

$$\begin{aligned} p(x) &= a + bx + cx^2 \\ &= (-2b - 4c) + bx + cx^2 \\ &= b(-2 + x) + c(-4 + x^2). \end{aligned}$$

This shows that $\mathcal{B} = \{-2 + x, -4 + x^2\}$ is a spanning set for U . We'll leave it as an exercise for you to check that \mathcal{B} is linearly independent, and is therefore a basis for U .

Exercise 8

Take another look at Example 1.3.15 and convince yourself that what was done there uses the same process performed in the above example.

1.3.4 Coordinates with Respect to a Basis

Recall that in \mathbb{R}^3 you may have seen that the vector $\begin{bmatrix} 3 \\ 2 \\ 4 \end{bmatrix}$ can be written as $3\vec{e}_1 + 2\vec{e}_2 + 4\vec{e}_3$,

where $\{\vec{e}_1, \vec{e}_2, \vec{e}_3\}$ is the standard basis for \mathbb{R}^3 . You have seen this to mean that the vector

$\begin{bmatrix} 3 \\ 2 \\ 4 \end{bmatrix}$ can be found 3-units in the x -direction, 2 in the y , and 4 in the z .

In fact, once we have a basis for a vector space, we can think of this as a choice of axes, and we can write every vector as a coordinate vector in much the same way as we think about vectors in \mathbb{R}^3 .

Example 1.3.21

Consider the vector $\vec{v} = 3 + 5x - 2x^2$ in $\mathcal{P}_2(\mathbb{R})$, and the bases $\mathcal{B} = \{1, x, x^2\}$ and $\mathcal{C} = \{1, 1 + x, 1 + x + x^2\}$ (as an exercise, prove that \mathcal{C} is a basis).

Then $\vec{v} = 3(1) + 5(x) + (-2)(x^2)$ so we think of \vec{v} as living at the coordinate $(3, 5, -2)$ with respect to the axes defined by \mathcal{B} .

We also have $\vec{v} = -2(1) + 7(1 + x) + (-2)(1 + x + x^2)$ so, with respect to the axes determined by \mathcal{C} , we can think of \vec{v} as living at the point $(-2, 7, -2)$.

More formally, we can write the coordinate vectors of \vec{v} with respect to \mathcal{B} and \mathcal{C} as

$$[\vec{v}]_{\mathcal{B}} = \begin{bmatrix} 3 \\ 5 \\ -2 \end{bmatrix} \quad \text{and} \quad [\vec{v}]_{\mathcal{C}} = \begin{bmatrix} -2 \\ 7 \\ -2 \end{bmatrix},$$

respectively. This gives us two different ways of looking at the same vector.

A natural question to ask is: does it even make sense to talk about coordinate vectors like this? Is it possible that the same vector has two different coordinate vectors with respect to the same basis? The answer is “no.”

Lemma 1.3.22

Let V be a vector space, let $S = \{\vec{v}_1, \dots, \vec{v}_k\}$ be a subset of V , and let $U = \text{Span}(S)$. Then every vector in U can be expressed in a unique way as a linear combination of the vectors in S if and only if S is linearly independent.

Proof: Suppose every vector in U is expressed uniquely as a linear combination of the vectors in S . Then there is only one way to write

$$\vec{0} = t_1 \vec{v}_1 + \dots + t_k \vec{v}_k,$$

which is $t_1 = \dots = t_k = 0$, so S is linearly independent. Conversely, suppose S is linearly independent and

$$t_1 \vec{v}_1 + \dots + t_k \vec{v}_k = s_1 \vec{v}_1 + \dots + s_k \vec{v}_k.$$

Rearranging we have $(t_1 - s_1) \vec{v}_1 + \dots + (t_k - s_k) \vec{v}_k = \vec{0}$. Since S is linearly independent, this can only be true if $t_i = s_i$ for all i , completing the proof. \square

If we apply this lemma to a basis of a vector space, we immediately get the following useful theorem.

Theorem 1.3.23 (Unique Representation Theorem)

Let V be a vector space and let $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$ be a basis of V . Then for all $\vec{v} \in V$ there exist unique scalars $x_1, \dots, x_n \in \mathbb{F}$ such that

$$\vec{v} = x_1 \vec{v}_1 + \dots + x_n \vec{v}_n.$$

We can thus unambiguously define the set of coordinates of a vector with respect to a given basis. But if we want to use these coordinates to form a *coordinate vector*, there is a small subtlety that must be addressed. The next example illustrates the issue.

Example 1.3.24

Let $\vec{v} = 2 - i + 4x - ix^2 \in \mathcal{P}_2(\mathbb{C})$. If $\mathcal{B} = \{1, x, x^2\}$ and $\mathcal{C} = \{1, x^2, x\}$ are bases for $\mathcal{P}_2(\mathbb{C})$, then the coordinates of \vec{v} with respect to these bases are $\begin{bmatrix} 2 - i \\ 4 \\ -i \end{bmatrix}$ and $\begin{bmatrix} 2 - i \\ -i \\ 4 \end{bmatrix}$, respectively.

That is, the order of the basis vectors matters!

Definition 1.3.25

Ordered Basis

Let V be a finite-dimensional vector space over \mathbb{F} . An **ordered basis for V** is a basis $\mathcal{B} = \{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n\}$ for V together with a fixed ordering.

REMARK

When we refer to the set $\{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n\}$ as being *ordered*, we are indicating that \vec{v}_1 is the first element in the ordering, that \vec{v}_2 is the second, and so on.

Thus even though $\{\vec{v}_1, \vec{v}_2\}$ and $\{\vec{v}_2, \vec{v}_1\}$ are the same *set*, they are different from the point of view of orderings.

A basis $\{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n\}$ gives rise to $n!$ ordered bases, one for each possible ordering (permutation) of the vectors in the basis.

Definition 1.3.26**Coordinate Vector**

Let $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$ be an *ordered* basis for a vector space V . If $\vec{x} \in V$ is written as

$$\vec{x} = x_1 \vec{v}_1 + \dots + x_n \vec{v}_n$$

then the **coordinate vector of \vec{x} with respect to \mathcal{B}** is

$$[\vec{x}]_{\mathcal{B}} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}.$$

Example 1.3.27

Consider the ordered basis

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

of $M_{2 \times 2}(\mathbb{R})$. Let $\vec{x} = \begin{bmatrix} 1 & -1 \\ 0 & 3 \end{bmatrix}$. We wish to find $[\vec{x}]_{\mathcal{B}}$. Consider the equation

$$a \begin{bmatrix} 3 & 2 \\ 2 & 2 \end{bmatrix} + b \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} + c \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} + d \begin{bmatrix} 1 & 4 \\ 0 & 3 \end{bmatrix} = \begin{bmatrix} 1 & -1 \\ 0 & 3 \end{bmatrix}.$$

To get the coordinate vector of \vec{x} with respect to \mathcal{B} , we need to solve for a, b, c, d . Equating the entries of the matrices on the left and right hand side of the equals sign gives us the system of equations

$$\begin{aligned} 3a + b + c + d &= 1 \\ 2a + c + 4d &= -1 \\ 2a + b + c &= 0 \\ 2a + b + 3d &= 3. \end{aligned}$$

To solve this equation we create an augmented matrix and row reduce, giving

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

Therefore

$$\vec{x} = 1 \begin{bmatrix} 3 & 2 \\ 2 & 2 \end{bmatrix} + 1 \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} - 3 \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} + 0 \begin{bmatrix} 1 & 4 \\ 0 & 1 \end{bmatrix}$$

and

$$[\vec{x}]_{\mathcal{B}} = \begin{bmatrix} 1 \\ 1 \\ -3 \\ 0 \end{bmatrix}.$$

Example 1.3.28

Earlier you may have noticed that there is some kind of similarity between \mathbb{R}^3 and $\mathcal{P}_2(\mathbb{R})$, and we can somehow identify the vectors

$$\vec{v} = \begin{bmatrix} a \\ b \\ c \end{bmatrix} \quad \text{and} \quad \vec{w} = a + bx + cx^2.$$

Now we can get a glimpse as to how these two vectors may indeed be viewed as the same after picking bases for the two vector spaces. Consider the standard (ordered) bases

$$\mathcal{B} = \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\} \quad \text{and} \quad \mathcal{C} = \{1, x, x^2\}$$

for \mathbb{R}^3 and $\mathcal{P}_2(\mathbb{R})$ respectively. Then we see that

$$[\vec{v}]_{\mathcal{B}} = [\vec{w}]_{\mathcal{C}} = \begin{bmatrix} a \\ b \\ c \end{bmatrix}.$$

Once we have chosen a basis for a vector space V , every vector can now be represented as a column vector. Column vectors, as we know, come with their own addition and scalar multiplication. A natural question to ask is whether or not the column vector addition and scalar multiplication agree with the addition and scalar multiplication on V . Since everything so far in this course has worked out so beautifully, it would be a huge surprise if this wasn't true! Indeed, it is true.

Theorem 1.3.29

Let V be a finite-dimensional vector space over \mathbb{F} with ordered basis \mathcal{B} . Then

$$[\vec{x} + \vec{y}]_{\mathcal{B}} = [\vec{x}]_{\mathcal{B}} + [\vec{y}]_{\mathcal{B}} \quad \text{and} \quad [t\vec{x}]_{\mathcal{B}} = t[\vec{x}]_{\mathcal{B}}$$

for all $\vec{x}, \vec{y} \in V$ and all $t \in \mathbb{F}$.

Proof: This is just a matter of using the definition to determine $[\vec{x}]_{\mathcal{B}}, [\vec{y}]_{\mathcal{B}}, [\vec{x} + \vec{y}]_{\mathcal{B}}$ and $[t\vec{x}]_{\mathcal{B}}$. We'll leave the details as an exercise. \square

Exercise 9

Prove Theorem 1.3.29.

In this next Chapter, we'll see how this theorem effectively allows us to replace computations in an n -dimensional vector space V with computations in the more familiar space \mathbb{F}^n .

Chapter 2

Linear Transformations

2.1 Linear Transformations Between Abstract Vectors

So far in the course we have studied vector spaces in isolation. That is, we've started with a single vector space and studied it, without looking at how it compares to, or interacts with, other vector spaces. However, we have seen glimpses that there is something to be said about comparing vector spaces. For example, \mathbb{R}^3 and $\mathcal{P}_2(\mathbb{R})$ appear to be the same vector space in some sense, just wrapped up in a different package.

In mathematics in general, when we want to study how objects interact, we usually think about functions between them. However, when studying functions between two vector spaces, we don't want to just take any old function. We'd like to take into account that we're working with vector spaces which come with the additional structure of vector addition and scalar multiplication. Ideally, our functions should respect this structure. This leads us to the following (hopefully familiar) definition.

Definition 2.1.1

Linear
Transformation,
Linear Map,
Linearity

If V and W are vector spaces over \mathbb{F} , a function $L: V \rightarrow W$ is called a **linear transformation** (or **linear map**) if it satisfies the **linearity** properties:

1. $L(\vec{x} + \vec{y}) = L(\vec{x}) + L(\vec{y})$, and
2. $L(t\vec{x}) = tL(\vec{x})$

for all $\vec{x}, \vec{y} \in V$, $t \in \mathbb{F}$.

Said another way, it doesn't matter if you add two vectors before or after applying the linear map, and the same with scalar multiplication.

Example 2.1.2

A simple but important linear map is the **identity map** (or **identity transformation**) $\text{id}: V \rightarrow V$, which sends each vector to itself: $\text{id}(\vec{x}) = \vec{x}$.

Example 2.1.3

Consider the map $L: \mathcal{P}_3(\mathbb{R}) \rightarrow \mathbb{R}$ defined by $L(p) = p(2)$. This is a linear map. In general, let $t \in \mathbb{F}$. Define the **evaluation map**

$$\text{ev}_t: \mathcal{P}_n(\mathbb{F}) \rightarrow \mathbb{F}$$

by $\text{ev}_t(p(x)) = p(t)$. This is a linear map, and the proof of this claim is left as an exercise.

Exercise 10

Let $t \in \mathbb{F}$. Prove that the evaluation map $\text{ev}_t: \mathcal{P}_n(\mathbb{F}) \rightarrow \mathbb{F}$ is a linear map.

Example 2.1.4

Let $\text{tr}: M_{n \times n}(\mathbb{F}) \rightarrow \mathbb{F}$ be the map defined by taking the trace of a matrix. Recall, if

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

then $\text{tr}(A) = a_{11} + a_{22} + \cdots + a_{nn}$. We will prove that tr is a linear map.

Let $A, B \in M_{n \times n}(\mathbb{F})$ with $A = \begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{bmatrix}$ and $B = \begin{bmatrix} b_{11} & \cdots & b_{1n} \\ \vdots & \ddots & \vdots \\ b_{n1} & \cdots & b_{nn} \end{bmatrix}$. Then

$$\begin{aligned} \text{tr}(A + B) &= a_{11} + b_{11} + a_{22} + b_{22} + \cdots + a_{nn} + b_{nn} \\ &= a_{11} + a_{22} + \cdots + a_{nn} + b_{11} + b_{22} + \cdots + b_{nn} \\ &= \text{tr}(A) + \text{tr}(B). \end{aligned}$$

If $t \in \mathbb{F}$, then

$$\begin{aligned} \text{tr}(tA) &= ta_{11} + ta_{22} + \cdots + ta_{nn} \\ &= t(a_{11} + a_{22} + \cdots + a_{nn}) \\ &= t(\text{tr}(A)) \end{aligned}$$

so tr is a linear map.

Example 2.1.5

By contrast, the determinant function $\det: M_{n \times n}(\mathbb{F}) \rightarrow \mathbb{F}$ is not linear when $n > 1$. Indeed, $\det(2A) = 2^n \det(A)$, so if $\det(A) \neq 0$, then $\det(2A) \neq 2 \det(A)$.

An important property of linear maps is that not only do they respect addition and scalar multiplication, they also respect the zero vector.

Proposition 2.1.6

Let $L: V \rightarrow W$ be a linear map, and let $\vec{0}_V$ and $\vec{0}_W$ denote the zero vectors of V and W , respectively. Then

$$L(\vec{0}_V) = \vec{0}_W.$$

REMARK

Since the notation $\vec{0}_V$ is somewhat cumbersome, going forwards we will only write $\vec{0}$ when it should be clear from context what the relevant vector space is.

Exercise 11 Prove Proposition 2.1.6. (**Hint:** Consider $L(t\vec{0})$ with $t = 0$.)

Example 2.1.7 The map $L: \mathcal{P}_n(\mathbb{F}) \rightarrow \mathcal{P}_n(\mathbb{F})$ defined by

$$L(p(x)) = p(x) + 1$$

is not linear, because it maps the zero polynomial to the polynomial 1.

The converse to Proposition 2.1.6 is false, as the next two Exercises show.

Exercise 12 Show that the following maps are not linear:

(a) $L: \mathbb{R}^3 \rightarrow \mathcal{P}_2(\mathbb{R})$ defined by $L\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) = a^2 + b^2x + c^2x^2$.

(b) $L: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathbb{R}$ defined by $L(a + bx + cx^2) = abc$.

Several other natural operations you are familiar with are linear maps. For example, integration and differentiation of polynomials are both linear maps.

Example 2.1.8 The differentiation map

$$D: \mathcal{P}_3(\mathbb{R}) \rightarrow \mathcal{P}_2(\mathbb{R})$$

given by $D(p(x)) = p'(x)$, or more explicitly by

$$D(a_0 + a_1x + a_2x^2 + a_3x^3) = a_1 + 2a_2x + 3a_3x^2,$$

can be shown to be a linear map.

Similarly, the integration map

$$I: \mathcal{P}_3(\mathbb{R}) \rightarrow \mathcal{P}_4(\mathbb{R})$$

given by $I(p) = \int_0^x p(t) dt$, or more explicitly by

$$I(a_0 + a_1x + a_2x^2 + a_3x^3) = a_0x + \frac{1}{2}a_1x^2 + \frac{1}{3}a_2x^3 + \frac{1}{4}a_3x^4,$$

is also a linear map.

Exercise 13 Show that the maps D and I in the previous example are linear maps.

2.2 Rank and Nullity

As we did for linear maps from \mathbb{F}^n to \mathbb{F}^m , we associate two important subspaces to each linear map $L: V \rightarrow W$.

Definition 2.2.1
Kernel, Nullspace,
Range

Let $L: V \rightarrow W$ be a linear map. The **kernel** (or **nullspace**) of L is

$$\text{Ker}(L) = \{\vec{x} \in V : L(\vec{x}) = \vec{0}\}.$$

The **range** of L is

$$\text{Range}(L) = \{L(\vec{x}) \in W : \vec{x} \in V\}.$$

The kernel of a linear map $L: V \rightarrow W$ is the set of all the vectors in V that are mapped to $\vec{0} \in W$. The range of L is all the vectors in W that are outputs of L .

Theorem 2.2.2

Let V and W be vector spaces over \mathbb{F} , and let $L: V \rightarrow W$ be a linear map. Then

- (a) $\text{Ker}(L)$ is a subspace of V , and
- (b) $\text{Range}(L)$ is a subspace of W .

Proof: The same proof that you've seen for linear maps from \mathbb{F}^n to \mathbb{F}^m works here. So we will only give a proof of part (a) and leave part (b) to you.

Let's use the Subspace Test. From Proposition 2.1.6, we see that $\vec{0} \in \text{Ker}(L)$, so $\text{Ker}(L)$ is non-empty. Suppose next that $\vec{v}, \vec{w} \in \text{Ker}(L)$. Then $L(\vec{v} + \vec{w}) = L(\vec{v}) + L(\vec{w}) = \vec{0} + \vec{0} = \vec{0}$ so $\vec{v} + \vec{w} \in \text{Ker}(L)$, and therefore $\text{Ker}(L)$ is closed under addition. Finally, let $t \in \mathbb{F}$. Then $L(t\vec{v}) = tL(\vec{v}) = \vec{0}$, so $\text{Ker}(L)$ is closed under scalar multiplication. Thus $\text{Ker}(L)$ is a subspace of V . \square

Exercise 14 Prove part (b) of Theorem 2.2.2.

Example 2.2.3

The linear map $L: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ defined by

$$L\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) = \begin{bmatrix} a \\ b \end{bmatrix}.$$

$$\text{has } \text{Ker}(L) = \left\{ \begin{bmatrix} 0 \\ 0 \\ c \end{bmatrix} \in \mathbb{R}^3 : c \in \mathbb{R} \right\} \text{ and } \text{Range}(L) = \mathbb{R}^2.$$

Example 2.2.4 Let $L: M_{2 \times 2}(\mathbb{R}) \rightarrow \mathcal{P}_2(\mathbb{R})$ be defined by

$$L\left(\begin{bmatrix} a & b \\ c & d \end{bmatrix}\right) = b + c + (c - d)x^2.$$

We leave it to you to check that L is linear. We have

$$\text{Ker}(L) = \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in M_{2 \times 2}(\mathbb{R}) : b + c = c - d = 0 \right\} = \left\{ \begin{bmatrix} a & -c \\ c & c \end{bmatrix} : a, c \in \mathbb{R} \right\}.$$

It is clear that $\text{Range}(L) \subseteq \text{Span}(\{1, x^2\})$. Since $L\left(\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}\right) = 1$ and $L\left(\begin{bmatrix} 0 & 0 \\ 0 & -1 \end{bmatrix}\right) = x^2$, we see $\text{Range}(L) \supseteq \text{Span}(\{1, x^2\})$. Therefore $\text{Range}(L) = \text{Span}(\{1, x^2\})$.

If you examine these examples, you'll notice something interesting about the dimensions of the vector spaces involved. In the first example, $\dim(\mathbb{R}^3) = 3$, $\dim(\text{Range}(L)) = 2$ and $\dim(\text{Ker}(L)) = 1$. In the second we have $\dim(M_{2 \times 2}(\mathbb{R})) = 4$, $\dim(\text{Ker}(L)) = 2$ and $\dim(\text{Range}(L)) = 2$. The dimension of the domain of L in both cases seems to be equal to the sum of the dimensions of $\text{Range}(L)$ and $\text{Ker}(L)$. This is not a coincidence! Let's give these dimensions some names.

Definition 2.2.5
Rank, Nullity

Let V and W be vector spaces over \mathbb{F} . The **rank** of a linear map $L: V \rightarrow W$ is the dimension of the range of L . The **nullity** of L is the dimension of the kernel (nullspace) of L . That is,

$$\text{rank}(L) = \dim(\text{Range}(L)) \quad \text{and} \quad \text{nullity}(L) = \dim(\text{Ker}(L)).$$

The key result about the rank and nullity of a linear map is the following theorem.

Theorem 2.2.6 (Rank–Nullity Theorem)

Let V and W be vector spaces over \mathbb{F} with V finite-dimensional and $\dim(V) = n$. Let $L: V \rightarrow W$ be a linear map. Then $\text{rank}(L) + \text{nullity}(L) = n$.

The idea of the proof is as follows. We will start with a basis of $\text{Ker}(L)$ of size k and we will extend this to a basis of V by adding m vectors to it (so $\dim(V) = k + m$). Then we prove that the image of the m new vectors under L give a basis for $\text{Range}(L)$, which will complete the proof.

Proof: Let $\{\vec{v}_1, \dots, \vec{v}_k\}$ be a basis for $\text{Ker}(L)$ so $\text{nullity}(L) = k$. Extend this to a basis $\{\vec{v}_1, \dots, \vec{v}_k, \vec{w}_1, \dots, \vec{w}_m\}$ for V so $\dim(V) = k + m$. It suffices to show $\mathcal{B} = \{L(\vec{w}_1), \dots, L(\vec{w}_m)\}$ is a basis for $\text{Range}(L)$. We first show $\text{Span}(\mathcal{B}) = \text{Range}(L)$. Clearly $\text{Span}(\mathcal{B}) \subseteq \text{Range}(L)$, so we must prove the reverse containment $\text{Span}(\mathcal{B}) \supseteq \text{Range}(L)$. Let $\vec{w} \in \text{Range}(L)$. Then $\vec{w} = L(\vec{v})$ for some $\vec{v} \in V$, and we may write \vec{v} as

$$\vec{v} = t_1 \vec{v}_1 + \dots + t_k \vec{v}_k + s_1 \vec{w}_1 + \dots + s_m \vec{w}_m.$$

Then

$$\begin{aligned}\vec{w} &= L(\vec{v}) = L(t_1\vec{v}_1 + \cdots + t_k\vec{v}_k + s_1\vec{w}_1 + \cdots + s_m\vec{w}_m) \\ &= t_1L(\vec{v}_1) + \cdots + t_kL(\vec{v}_k) + s_1L(\vec{w}_1) + \cdots + s_mL(\vec{w}_m) \\ &= s_1L(\vec{w}_1) + \cdots + s_mL(\vec{w}_m).\end{aligned}$$

So \mathcal{B} is a spanning set for $\text{Range}(L)$. For linear independence, suppose that

$$s_1L(\vec{w}_1) + \cdots + s_mL(\vec{w}_m) = \vec{0}.$$

Since L is linear, this implies $s_1\vec{w}_1 + \cdots + s_m\vec{w}_m \in \text{Ker}(L)$. Therefore

$$s_1\vec{w}_1 + \cdots + s_m\vec{w}_m = t_1\vec{v}_1 + \cdots + t_k\vec{v}_k$$

for some $t_1, \dots, t_k \in \mathbb{F}$. However, $\{\vec{v}_1, \dots, \vec{v}_k, \vec{w}_1, \dots, \vec{w}_m\}$ is linearly independent, so we must conclude $s_1 = \cdots = s_m = t_1 = \cdots = t_k = 0$. Therefore \mathcal{B} is a basis for $\text{Range}(L)$ and so $\text{rank}(L) = m$. Since $\text{nullity}(L) = k$ and $\dim(V) = m + k$, the proof is complete. \square

REMARK

You might be familiar with the Rank–Nullity theorem for matrices, which states that if $A \in M_{m \times n}(\mathbb{F})$ then

$$n = \text{rank}(A) + \text{nullity}(A),$$

where $\text{rank}(A)$ is the dimension of the column space of A , and $\text{nullity}(A)$ is the dimension of the nullspace of A . We'll show below that this theorem is consequence of the Rank–Nullity theorem for linear maps stated above. See Corollary 2.5.21.

Here are some examples of the kinds of things you can conclude with the Rank–Nullity theorem in your back pocket.

Example 2.2.7

Let $L: \mathcal{P}_3(\mathbb{R}) \rightarrow \mathbb{R}^3$ be a linear map. Since $\dim(\mathbb{R}^3) = 3$, it must be that $\text{rank}(L) \leq 3$. Since $\dim(\mathcal{P}_3(\mathbb{R})) = 4$, the Rank–Nullity theorem implies $\text{nullity}(L) \geq 1$. Therefore without knowing anything about the linear map, we can conclude that there is at least one non-zero vector $\vec{v} \in \mathcal{P}_3(\mathbb{R})$ such that $L(\vec{v}) = \vec{0}$.

Example 2.2.8

Let $L: \mathbb{C}^4 \rightarrow M_{2 \times 2}(\mathbb{C})$ be a linear map. Then $\text{Ker}(L) = \{\vec{0}\}$ if and only if $\text{Range}(L) = M_{2 \times 2}(\mathbb{C})$.

Proof: First note $\dim(\mathbb{C}^4) = \dim(M_{2 \times 2}(\mathbb{C})) = 4$. If $\text{Ker}(L) = \{\vec{0}\}$ then $\text{nullity}(L) = 0$ so the Rank–Nullity theorem says $\text{rank}(L) = 4$. Therefore $\text{Range}(L)$ is a 4-dimensional subspace of $M_{2 \times 2}(\mathbb{C})$ so it must be that $\text{Range}(L) = M_{2 \times 2}(\mathbb{C})$ (why?). Conversely, if $\text{Range}(L) = M_{2 \times 2}(\mathbb{C})$, then $\text{rank}(L) = 4$. Therefore $\text{nullity}(L) = 0$ so $\text{Ker}(L) = \{\vec{0}\}$. \square

2.3 Linear Maps as Matrices

Recall that given an $m \times n$ matrix $A \in M_{m \times n}(\mathbb{F})$, we can define a linear map $L: \mathbb{F}^n \rightarrow \mathbb{F}^m$ by $L(\vec{x}) = A\vec{x}$. Conversely, you learned that every linear map $L: \mathbb{F}^n \rightarrow \mathbb{F}^m$ can be realized in this form for an appropriate matrix A . So, in some sense, linear maps between \mathbb{F}^n and \mathbb{F}^m and matrices in $M_{m \times n}(\mathbb{F})$ are two sides of the same coin.

Example 2.3.1

Consider the linear map $L: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by $L\left(\begin{bmatrix} a \\ b \end{bmatrix}\right) = \begin{bmatrix} a + 2b \\ a - 2b \end{bmatrix}$. We can find a matrix that performs this transformation. Indeed,

$$\begin{bmatrix} 1 & 2 \\ 1 & -2 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a + 2b \\ a - 2b \end{bmatrix}.$$

The fact that a matrix even existed in the previous example was plausible because it's not much of a stretch of the imagination to view a vector in \mathbb{R}^2 as a column vector. But we can do this in *every vector space*—once we fix a basis! Recall that if we fix an ordered basis for a vector space, then every vector can be written as a column vector by simply taking its coordinate vector.

So, now that we have this, it's reasonable to ask whether or not every linear map can be viewed as a matrix transformation on coordinate vectors. Let's take a look at an example.

Example 2.3.2

Let $L: \mathcal{P}_2(\mathbb{F}) \rightarrow M_{2 \times 2}(\mathbb{F})$ be the linear map defined by $L(a + bx + cx^2) = \begin{bmatrix} a - 2b & 4c \\ a + b + c & b - c \end{bmatrix}$. Fix the standard (ordered) bases

$$\mathcal{B} = \{1, x, x^2\} \quad \text{and} \quad \mathcal{C} = \left\{ \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \right\}$$

for $\mathcal{P}_2(\mathbb{F})$ and $M_{2 \times 2}(\mathbb{F})$, respectively. The coordinate vectors of $\vec{v} = a + bx + cx^2$ and $L(\vec{v}) = \begin{bmatrix} a - 2b & 4c \\ a + b + c & b - c \end{bmatrix}$ are

$$[\vec{v}]_{\mathcal{B}} = \begin{bmatrix} a \\ b \\ c \end{bmatrix} \quad \text{and} \quad [L(\vec{v})]_{\mathcal{C}} = \begin{bmatrix} a - 2b \\ 4c \\ a + b + c \\ b - c \end{bmatrix}.$$

So, if there is a matrix A which performs the linear map for us (by matrix multiplication of course), it must be such that

$$A \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} a - 2b \\ 4c \\ a + b + c \\ b - c \end{bmatrix}.$$

We first note that if A is to exist, it must be a 4×3 matrix. With that in mind, if we stare at this for a while (we'll explain how to do this more systematically later on) we can see that we can take

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

For some foreshadowing of notation, we let ${}_C[L]_{\mathcal{B}} = A$.

The previous example is not just a fortunate coincidence. We are always able to find a matrix that performs any given linear transformation, once we fix bases and work with the resulting coordinate vectors. This is the content of our next theorem.

Before we state and prove it though, it is worth addressing why we'd care to do this. Matrices come equipped with machinery to compute many things. It will turn out that once we turn our linear map into a matrix, we can use this machinery to learn about our linear map.

Theorem 2.3.3

Let V be an n -dimensional vector space with ordered basis \mathcal{B} . Let W be an m -dimensional vector space with ordered basis \mathcal{C} . Then, for every linear map $L: V \rightarrow W$, there exists an $m \times n$ matrix A such that $[L(\vec{v})]_{\mathcal{C}} = A[\vec{v}]_{\mathcal{B}}$ for all $\vec{v} \in V$. Conversely, every $m \times n$ matrix A defines a linear map $L: V \rightarrow W$ by $[L(\vec{v})]_{\mathcal{C}} = A[\vec{v}]_{\mathcal{B}}$.

Proof: Since matrix multiplication satisfies $A(B + C) = AB + AC$ and $t(AB) = A(tB)$ for all matrices A, B, C and all scalars $t \in \mathbb{F}$, A defines a linear map $L: V \rightarrow W$ by $A[\vec{v}]_{\mathcal{B}} = [L(\vec{v})]_{\mathcal{C}}$.

For the forward direction, let $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$ and $\mathcal{C} = \{\vec{w}_1, \dots, \vec{w}_m\}$. Let $\vec{v} \in V$, then $\vec{v} = t_1 \vec{v}_1 + \dots + t_n \vec{v}_n$ and $L(\vec{v}) = s_1 \vec{w}_1 + \dots + s_m \vec{w}_m$. Since L is linear we have

$$L(\vec{v}) = t_1 L(\vec{v}_1) + \dots + t_n L(\vec{v}_n) = s_1 \vec{w}_1 + \dots + s_m \vec{w}_m.$$

For each $i \in \{1, \dots, n\}$, let $L(\vec{v}_i) = a_{1i} \vec{w}_1 + \dots + a_{mi} \vec{w}_m$. Then

$$\begin{aligned} L(\vec{v}) &= s_1 \vec{w}_1 + \dots + s_m \vec{w}_m \\ &= t_1 (a_{11} \vec{w}_1 + \dots + a_{m1} \vec{w}_m) + \dots + t_n (a_{1n} \vec{w}_1 + \dots + a_{mn} \vec{w}_m) \\ &= (a_{11} t_1 + a_{12} t_2 + \dots + a_{1n} t_n) \vec{w}_1 + \dots + (a_{m1} t_1 + \dots + a_{mn} t_n) \vec{w}_m. \end{aligned}$$

Therefore we have $s_i = a_{i1} t_1 + \dots + a_{in} t_n$ for all $i \in \{1, \dots, m\}$. This is of course how matrix multiplication works, and we see

$$\begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{m1} & \cdots & a_{mn} \end{bmatrix} \begin{bmatrix} t_1 \\ \vdots \\ t_n \end{bmatrix} = \begin{bmatrix} s_1 \\ \vdots \\ s_m \end{bmatrix}.$$

Since $[\vec{v}]_{\mathcal{B}} = \begin{bmatrix} t_1 \\ \vdots \\ t_n \end{bmatrix}$ and $[L(\vec{v})]_{\mathcal{C}} = \begin{bmatrix} s_1 \\ \vdots \\ s_m \end{bmatrix}$, the proof is completed. \square

Hidden in the proof is the fact that if \vec{v}_i is the i th basis vector of \mathcal{B} , then $[L(\vec{v}_i)]_{\mathcal{C}}$ is simply the i th column of the desired matrix A . This gives us the following corollary.

Corollary 2.3.4

Let V be a vector space with ordered basis $\mathcal{B} = \{\vec{b}_1, \dots, \vec{b}_n\}$. Let W be a vector space with ordered basis $\mathcal{C} = \{\vec{c}_1, \dots, \vec{c}_m\}$. Let $L: V \rightarrow W$ be a linear map. Then the $m \times n$ matrix A such that $[L(\vec{v})]_{\mathcal{C}} = A[\vec{v}]_{\mathcal{B}}$ for all $\vec{v} \in V$, which we denote ${}_c[L]_{\mathcal{B}}$, is given by

$${}_c[L]_{\mathcal{B}} = \left[[L(\vec{b}_1)]_{\mathcal{C}} \cdots [L(\vec{b}_n)]_{\mathcal{C}} \right].$$

The fact that the matrix contains all the information of L , and is determined by the images of the basis vectors, tells us something very interesting about linear maps: They are entirely determined by where they send a basis.

The matrix A for a linear map L is determined once you pick ordered bases for each vector space. We will give this matrix a name.

Definition 2.3.5

Matrix of a Linear Map

We call the matrix ${}_c[L]_{\mathcal{B}}$ the **matrix of the linear map** $L: V \rightarrow W$ with respect to the ordered bases \mathcal{B} and \mathcal{C} of V and W , respectively.

If $V = W$, so that $L: V \rightarrow V$, and we choose the same basis \mathcal{B} for both the domain and codomain of L , then we will write $[L]_{\mathcal{B}}$ instead of ${}_c[L]_{\mathcal{B}}$.

From Corollary 2.3.4 we know that the matrix ${}_c[L]_{\mathcal{B}}$ of the linear map $L: V \rightarrow W$ satisfies a very important identity:

$$[L(\vec{v})]_{\mathcal{C}} = {}_c[L]_{\mathcal{B}}[\vec{v}]_{\mathcal{B}} \quad \text{for all } \vec{v} \in V.$$

Also, the size of ${}_c[L]_{\mathcal{B}}$ is $m \times n$, where $m = \dim(W)$ and $n = \dim(V)$.

Let's see some examples.

Example 2.3.6

Consider the differentiation map $D: \mathcal{P}_3(\mathbb{R}) \rightarrow \mathcal{P}_2(\mathbb{R})$, and let both vector spaces be endowed with the standard ordered bases \mathcal{B} and \mathcal{C} respectively. Then $D(1) = 0$, $D(x) = 1$, $D(x^2) = 2x$, and $D(x^3) = 3x^2$. Therefore

$$\begin{aligned} {}_c[D]_{\mathcal{B}} &= \left[[D(1)]_{\mathcal{C}} \ [D(x)]_{\mathcal{C}} \ [D(x^2)]_{\mathcal{C}} \ [D(x^3)]_{\mathcal{C}} \right] \\ &= \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{bmatrix}. \end{aligned}$$

Let's confirm that this works with a specific example. Let $\vec{v} = 4 + 2x + (-2)x^2 + x^3$. Then $D(\vec{v}) = 2 - 4x + 3x^2$ so $[\vec{v}]_{\mathcal{B}} = \begin{bmatrix} 4 \\ 2 \\ -2 \\ 1 \end{bmatrix}$ and $[D(\vec{v})]_{\mathcal{C}} = \begin{bmatrix} 2 \\ -4 \\ 3 \end{bmatrix}$. Indeed, we can check that the identity $[D(\vec{v})]_{\mathcal{C}} = {}_c[D]_{\mathcal{B}}[\vec{v}]_{\mathcal{B}}$ holds:

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

Example 2.3.7

Consider the trace map $\text{tr}: M_{2 \times 2}(\mathbb{F}) \rightarrow \mathbb{F}$ defined by $\text{tr} \left(\begin{bmatrix} a & b \\ c & d \end{bmatrix} \right) = a + d$. If we give $M_{2 \times 2}(\mathbb{F})$ and \mathbb{F} their standard bases \mathcal{B} and \mathcal{C} , respectively, then ${}_c[\text{tr}]_{\mathcal{B}}$ will be the 1×4 matrix whose columns are the traces of the standard basis matrices in \mathcal{B} :

$${}_c[\text{tr}]_{\mathcal{B}} = \left[\text{tr} \left(\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \right) \text{tr} \left(\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \right) \text{tr} \left(\begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} \right) \text{tr} \left(\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \right) \right] = [1 \ 0 \ 0 \ 1].$$

As a check, if we take $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$, we find that

$${}_c[\text{tr}]_{\mathcal{B}} [A]_{\mathcal{B}} = [1 \ 0 \ 0 \ 1] \begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \end{bmatrix} = 1 + 4.$$

Thus, ${}_c[\text{tr}]_{\mathcal{B}} [A]_{\mathcal{B}} = [\text{tr}(A)]_{\mathcal{C}}$, as expected.

Exercise 15

Let $L: \mathbb{R}^2 \rightarrow \mathcal{P}_2(\mathbb{R})$ be the linear map given by $L \left(\begin{bmatrix} a \\ b \end{bmatrix} \right) = a + (a + b)x + (a + 2b)x^2$.

- Find ${}_c[L]_{\mathcal{B}}$, where \mathcal{B} and \mathcal{C} are the standard bases for \mathbb{R}^2 and $\mathcal{P}_2(\mathbb{R})$, respectively.
- Pick your favorite non-zero vector $\vec{v} \in \mathbb{R}^2$ and confirm that $[L(\vec{v})]_{\mathcal{C}} = {}_c[L]_{\mathcal{B}}[\vec{v}]_{\mathcal{B}}$ by computing both sides of the equation separately.

If you dwell on Theorem 2.3.3, it becomes apparent that the theorem only works because of the way matrix multiplication is defined. In fact, the reason matrix multiplication is defined the way it is is essentially so that Theorem 2.3.3 is true! Even better, the next fact is also true, although we will not prove it here. If you feel like a moderately difficult challenge, you should prove it! You definitely have the tools to do so at this point in the course. The proof is more a matter of careful bookkeeping.

Proposition 2.3.8

Let V , U , and W be vector spaces over \mathbb{F} with bases \mathcal{B} , \mathcal{C} , and \mathcal{D} respectively. Let $L: V \rightarrow U$ and $M: U \rightarrow W$ be linear maps. Then ${}_{\mathcal{D}}[M \circ L]_{\mathcal{B}} = {}_{\mathcal{D}}[M]_{\mathcal{C}} {}_{\mathcal{C}}[L]_{\mathcal{B}}$.

Exercise 16

Prove Proposition 2.3.8.

Now, let's see some of the computational power of matrices in action. First we recall the notions of a column space and nullspace for a matrix, and see how they relate to the range and nullspace of a linear map.

Definition 2.3.9

**Column Space,
Rank, Nullspace,
Nullity of a Matrix**

Let $A \in M_{m \times n}(\mathbb{F})$.

The **column space** of A , denoted by $\text{Col}(A)$, is the span of the columns of A . The **rank** of A is the dimension of its column space:

$$\text{rank}(A) = \dim(\text{Col}(A)).$$

The **nullspace** of A , denoted by $\text{Null}(A)$, is the set of all $\vec{v} \in \mathbb{F}^n$ such that $A\vec{v} = \vec{0}$. The **nullity** of A is the dimension of its nullspace:

$$\text{nullity}(A) = \dim(\text{Null}(A)).$$

REMARK

Recall that if we let $L: \mathbb{F}^n \rightarrow \mathbb{F}^m$ be the linear transformation determined by $A \in M_{m \times n}(\mathbb{F})$, namely the one defined by $L(\vec{x}) = A\vec{x}$, then

$$\text{Col}(A) = \text{Range}(L) \quad \text{and} \quad \text{Null}(A) = \text{Ker}(L).$$

This shows, in particular, that $\text{Col}(A)$ and $\text{Null}(A)$ are subspaces of \mathbb{F}^m and \mathbb{F}^n , respectively. We will see below that a similar type of result is true for linear mappings between general vector spaces.

Let's quickly review how we can find bases for the column space and nullspace of a given matrix. For the column space, one row-reduces the matrix and chooses the original columns corresponding to the columns with leading ones in them. For the nullspace, one solves the system of equations given by augmenting the matrix with a column of 0 and then taking the basic solutions. Let's see this in an example.

Example 2.3.10

Find bases for $\text{Col}(A)$ and $\text{Null}(A)$ where $A = \begin{bmatrix} 1 & 2 & 5 & -3 & -8 \\ -2 & -4 & -11 & 2 & 4 \\ -1 & -2 & -6 & -1 & -4 \\ 1 & 2 & 5 & -2 & -5 \end{bmatrix}$.

Solution:

First, we put the matrix A into row reduced echelon form, which is given by

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

So, immediately, we have that a basis for $\text{Col}(A)$ is

$$\left\{ \begin{bmatrix} 1 \\ -2 \\ -1 \\ 1 \end{bmatrix}, \begin{bmatrix} 5 \\ -11 \\ -6 \\ 5 \end{bmatrix}, \begin{bmatrix} -3 \\ 2 \\ -1 \\ -2 \end{bmatrix} \right\}.$$

Finding a basis for $\text{Null}(A)$ is a little more involved. Finding a vector $\vec{v} = \begin{bmatrix} x_1 \\ \vdots \\ x_5 \end{bmatrix}$ such that

$A\vec{v} = \vec{0}$ is the same as solving the system of equations given by the augmented matrix $[A \mid \mathbf{0}]$. The row-reduced augmented matrix is given by

$$\left[\begin{array}{ccccc|c} 1 & 2 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 3 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right].$$

If we let the variables for this system be x_1, \dots, x_5 , we can write down an entire set of solutions as follows. For every column not corresponding to a leading 1, we let that variable be a free variable, and solve for the rest of them. In this example, the free variables are x_2 and x_5 , so let $x_2 = s$ and $x_5 = t$. Then

$$\begin{aligned} x_1 &= -t - 2s \\ x_2 &= s \\ x_3 &= 0 \\ x_4 &= -3t \\ x_5 &= t \end{aligned}$$

so every vector in $\text{Null}(A)$ is of the form

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = t \begin{bmatrix} -1 \\ 0 \\ 0 \\ -3 \\ 1 \end{bmatrix} + s \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

Finally, we write down a basis for $\text{Null}(A)$ as

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

The next proposition allows us to harness the computational power of matrices to learn about the range and kernel of a linear map.

Proposition 2.3.11

Let $L: V \rightarrow W$ be a linear map between finite-dimensional vector spaces, and let $A = {}_C[L]_{\mathcal{B}}$, where \mathcal{B} and \mathcal{C} are ordered bases for V and W respectively.

- (a) $\vec{v} \in \text{Ker}(L)$ if and only if $[\vec{v}]_{\mathcal{B}} \in \text{Null}(A)$.
 (b) $\vec{w} \in \text{Range}(L)$ if and only if $[\vec{w}]_{\mathcal{C}} \in \text{Col}(A)$.

Proof: (a) Observe that $\vec{v} \in V$ will be in $\text{Ker}(L)$ if and only if $L(\vec{v}) = \vec{0}$ which is the case if and only if $[L(\vec{v})]_{\mathcal{C}} = [\vec{0}]_{\mathcal{C}}$. According to the definition of $A = {}_C[L]_{\mathcal{B}}$, this last condition is equivalent to $A[\vec{v}]_{\mathcal{B}} = \vec{0}$. Thus $\vec{v} \in \text{Ker}(L)$ is equivalent to $[\vec{v}]_{\mathcal{B}} \in \text{Null}(A)$.

(b) Exercise. □

Exercise 17

Prove part (b) of Proposition 2.3.11.

This proposition tells us that if we want to find a basis for the kernel and range of a linear map, we just need to pick some bases, find the matrix associated to the linear map and find bases for the nullspace and column space of the matrix.

Example 2.3.12

Consider the linear map $L: \mathcal{P}_2(\mathbb{R}) \rightarrow M_{2 \times 2}(\mathbb{R})$ given by

$$L(a + bx + cx^2) = \begin{bmatrix} a + b + c & a - b + 3c \\ 3a + b + 5c & 0 \end{bmatrix}.$$

Find bases for $\text{Ker}(L)$ and $\text{Range}(L)$.

Solution:

Let \mathcal{B} and \mathcal{C} be the standard bases for $\mathcal{P}_2(\mathbb{R})$ and $M_{2 \times 2}(\mathbb{R})$ respectively. Since $L(1) = \begin{bmatrix} 1 & 1 \\ 3 & 0 \end{bmatrix}$,

$L(x) = \begin{bmatrix} 1 & -1 \\ 1 & 0 \end{bmatrix}$, and $L(x^2) = \begin{bmatrix} 1 & 3 \\ 5 & 0 \end{bmatrix}$ we have

$${}_C[L]_{\mathcal{B}} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & -1 & 3 \\ 3 & 1 & 5 \\ 0 & 0 & 0 \end{bmatrix}.$$

Call this matrix A . We will now find bases for $\text{Col}(A)$ and $\text{Null}(A)$, and then convert this information back to find bases for $\text{Range}(L)$ and $\text{Ker}(L)$. Row reducing A gives

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & -1 & 3 \\ 3 & 1 & 5 \\ 0 & 0 & 0 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

With a little work we compute bases for $\text{Col}(A)$ and $\text{Null}(A)$ to be

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

respectively. If we convert coordinate vectors back to vectors, we obtain the sets

$$\mathcal{D} = \left\{ \begin{bmatrix} 1 & 1 \\ 3 & 0 \end{bmatrix}, \begin{bmatrix} 1 & -1 \\ 1 & 0 \end{bmatrix} \right\} \quad \text{and} \quad \mathcal{E} = \{-2 + x + x^2\}$$

which are in $\text{Range}(L)$ and $\text{Ker}(L)$, respectively, by Proposition 2.3.11. Since the corresponding sets of coordinate vectors were bases for $\text{Col}(A)$ and $\text{Null}(A)$, we suspect that \mathcal{D} and \mathcal{E} are bases for $\text{Range}(L)$ and $\text{Ker}(L)$. In fact, this is true. We will prove a general result later that will give us this for free (see Proposition 2.5.14). For now, let's prove directly that \mathcal{E} is a basis for $\text{Ker}(L)$. We'll leave the proof that \mathcal{D} is a basis for $\text{Range}(L)$ as an exercise.

If $p \in \text{Ker}(L)$, then $[p]_{\mathcal{B}} \in \text{Null}(A)$, so

$$[p]_{\mathcal{B}} = a \begin{bmatrix} -2 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} -2a \\ a \\ a \end{bmatrix}$$

for some $a \in \mathbb{R}$. Thus, $p = -2a + ax + ax^2 = a(-2 + x + x^2)$, and so \mathcal{E} spans $\text{Ker}(L)$. Since \mathcal{E} contains exactly one non-zero vector, \mathcal{E} is linearly independent and hence is a basis for $\text{Ker}(L)$.

Exercise 18

Show that \mathcal{D} given at the end of Example 2.3.12 is a basis for $\text{Range}(L)$. [Give a proof similar to the one we gave for \mathcal{E} and $\text{Ker}(L)$, by utilizing the coordinates map $[\]_{\mathcal{C}}$.]

2.4 Change of Coordinates

We may be faced with a situation where we want to switch bases for the same vector space, because a particular problem is computationally easier to solve in one of the bases. We do this all the time in physics when we choose a set of coordinates that is natural with respect to the problem at hand.

If we are given bases \mathcal{B} and \mathcal{C} of a vector space V , it would be handy if we could find a matrix that takes a coordinate vector with respect to \mathcal{B} and transforms it into the coordinate vector with respect to \mathcal{C} . This can be achieved by simply finding the matrix of the identity map $\text{id}: V \rightarrow V$.

Example 2.4.1

Let $\mathcal{S} = \{1, x, x^2\}$ be the standard basis for $\mathcal{P}_2(\mathbb{R})$ and $\mathcal{B} = \{1, 1 + x, 1 + x + x^2\}$ another basis. We would like a matrix A such that $A[\vec{v}]_{\mathcal{B}} = [\vec{v}]_{\mathcal{S}}$. To find A , consider the linear map $\text{id}: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathcal{P}_2(\mathbb{R})$ given by $\text{id}(\vec{v}) = \vec{v}$ for all $\vec{v} \in \mathcal{P}_2(\mathbb{R})$. We will find $_{\mathcal{S}}[\text{id}]_{\mathcal{B}}$. This

should be our desired matrix since ${}_S[\text{id}]_{\mathcal{B}}[\vec{v}]_{\mathcal{B}} = [\text{id}(\vec{v})]_{\mathcal{S}} = [\vec{v}]_{\mathcal{S}}$. We will call this matrix ${}_S\mathcal{I}_{\mathcal{B}}$. We have

$$[1]_{\mathcal{S}} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad [1+x]_{\mathcal{S}} = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \quad \text{and} \quad [1+x+x^2]_{\mathcal{S}} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}.$$

Therefore

$${}_S\mathcal{I}_{\mathcal{B}} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}.$$

If we let $\vec{v} = 1 + x - x^2$, we find that ${}_S\mathcal{I}_{\mathcal{B}}[\vec{v}]_{\mathcal{B}} = [\vec{v}]_{\mathcal{S}}$:

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

So ${}_S\mathcal{I}_{\mathcal{B}}$ behaves as expected.

Example 2.4.2

Continuing from the previous example, let's try to find the matrix that changes \mathcal{S} -coordinates to \mathcal{B} -coordinates. By the same reasoning as above, this matrix should be ${}_{\mathcal{B}}[\text{id}]_{\mathcal{S}}$, and we'll denote it by ${}_{\mathcal{B}}\mathcal{I}_{\mathcal{S}}$.

Converting from standard coordinates to non-standard coordinates may take a bit more time. In our example it is not too difficult to notice that

$$[\text{id}(1)]_{\mathcal{B}} = [1]_{\mathcal{B}} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}.$$

Finding the \mathcal{B} -coordinates of x amounts to finding scalars c_1 , c_2 and c_3 such that

$$\begin{aligned} c_1(1) + c_2(1+x) + c_3(1+x+x^2) &= x, \\ (c_1 + c_2 + c_3) + (c_2 + c_3)x + c_3x^2 &= x. \end{aligned}$$

Equating the coefficients of polynomials on both sides of the above equality, we obtain the system

$$\begin{aligned} c_1 + c_2 + c_3 &= 0 \\ c_2 + c_3 &= 1 \\ c_3 &= 0 \end{aligned}$$

The only solution to this system is $c_1 = -1$, $c_2 = 1$ and $c_3 = 0$, meaning that

$$[\text{id}(x)]_{\mathcal{B}} = [x]_{\mathcal{B}} = \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix}.$$

We can find the \mathcal{B} -coordinates of x^2 in a similar way:

$$[\text{id}(x^2)]_{\mathcal{B}} = [x^2]_{\mathcal{B}} = \begin{bmatrix} 0 \\ -1 \\ 1 \end{bmatrix}.$$

Therefore

$${}_{\mathcal{B}}\mathcal{I}_{\mathcal{S}} = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 0 & 0 & 1 \end{bmatrix}.$$

If this matrix does what it should, then we should be able to use it to switch coordinates from \mathcal{S} to \mathcal{B} . Let's check. As in the preceding example, take $\vec{v} = 0(1)+2(1+x)+(-1)(1+x+x^2) = 1+x-x^2$. Then

$$[\vec{v}]_{\mathcal{S}} = \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix} \quad \text{and} \quad [\vec{v}]_{\mathcal{B}} = \begin{bmatrix} 0 \\ 2 \\ -1 \end{bmatrix}.$$

And indeed,

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

which shows that ${}_{\mathcal{B}}\mathcal{I}_{\mathcal{S}}[\vec{v}]_{\mathcal{S}} = [\vec{v}]_{\mathcal{B}}$.

Definition 2.4.3

Change of
Coordinates
Matrix

Let V be a finite dimensional vector space, and let \mathcal{B} and \mathcal{C} be two bases for V . The **change of coordinates matrix** ${}_{\mathcal{C}}\mathcal{I}_{\mathcal{B}}$ is the matrix ${}_{\mathcal{C}}[\text{id}]_{\mathcal{B}}$, where $\text{id}: V \rightarrow V$ is the identity map.

This name makes sense since $[\vec{v}]_{\mathcal{C}} = [\text{id}(\vec{v})]_{\mathcal{C}} = {}_{\mathcal{C}}\mathcal{I}_{\mathcal{B}} [\vec{v}]_{\mathcal{B}}$ for all $\vec{v} \in V$.

Let us address the following natural question: What is the relationship between ${}_{\mathcal{C}}\mathcal{I}_{\mathcal{B}}$ and ${}_{\mathcal{B}}\mathcal{I}_{\mathcal{C}}$? Notice that

$${}_{\mathcal{C}}\mathcal{I}_{\mathcal{B}} {}_{\mathcal{B}}\mathcal{I}_{\mathcal{C}} [\vec{v}]_{\mathcal{C}} = [\vec{v}]_{\mathcal{C}} \quad \text{and} \quad {}_{\mathcal{B}}\mathcal{I}_{\mathcal{C}} {}_{\mathcal{C}}\mathcal{I}_{\mathcal{B}} [\vec{v}]_{\mathcal{B}} = [\vec{v}]_{\mathcal{B}}$$

for all $\vec{v} \in V$. With this in mind you are able to write up a proof of the next proposition.

Proposition 2.4.4

Let V be a finite dimensional vector space with bases \mathcal{B} and \mathcal{C} . Then ${}_{\mathcal{C}}\mathcal{I}_{\mathcal{B}} = ({}_{\mathcal{B}}\mathcal{I}_{\mathcal{C}})^{-1}$.

Exercise 19

Prove Proposition 2.4.4.

Example 2.4.5

Referring to Examples 2.4.2 and 2.4.1, where we found

$${}_{\mathcal{B}}\mathcal{I}_{\mathcal{S}} = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad {}_{\mathcal{S}}\mathcal{I}_{\mathcal{B}} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix},$$

we see that

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

which shows that ${}_{\mathcal{S}}\mathcal{I}_{\mathcal{B}}$ and ${}_{\mathcal{B}}\mathcal{I}_{\mathcal{S}}$ are inverses.

To close off this section, let's consider the following scenario. Suppose we have a linear map $L: V \rightarrow W$ and suppose we choose bases \mathcal{B} and \mathcal{C} for V and W , respectively. This allows us to create the matrix ${}_c[L]_{\mathcal{B}}$. If we choose different bases \mathcal{B}' and \mathcal{C}' for V and W , then we are able to create the matrix ${}_{c'}[L]_{\mathcal{B}'}$. How are ${}_c[L]_{\mathcal{B}}$ and ${}_{c'}[L]_{\mathcal{B}'}$ related? We will show that

$${}_{c'}[L]_{\mathcal{B}'} = {}_{c'}\mathcal{I}_{\mathcal{C}} {}_c[L]_{\mathcal{B}} {}_{\mathcal{B}}\mathcal{I}_{\mathcal{B}'},$$

since to apply the transformation ${}_{c'}[L]_{\mathcal{B}'}$ we can begin by changing coordinates on the domain from \mathcal{B}' to \mathcal{B} , apply ${}_c[L]_{\mathcal{B}}$ and then change coordinates from \mathcal{C} to \mathcal{C}' .

Proposition 2.4.6

Let $L: V \rightarrow W$ be a linear map between two finite-dimensional vector spaces V and W . Suppose that \mathcal{B} and \mathcal{B}' are ordered bases for V and that \mathcal{C} and \mathcal{C}' are ordered bases for W . Then

$${}_{c'}[L]_{\mathcal{B}'} = {}_{c'}\mathcal{I}_{\mathcal{C}} {}_c[L]_{\mathcal{B}} {}_{\mathcal{B}}\mathcal{I}_{\mathcal{B}'}$$

Proof: This is just a matter of unpacking the definitions. We have

$$\begin{aligned} {}_{c'}\mathcal{I}_{\mathcal{C}} {}_c[L]_{\mathcal{B}} {}_{\mathcal{B}}\mathcal{I}_{\mathcal{B}'} [\vec{v}]_{\mathcal{B}'} &= {}_{c'}\mathcal{I}_{\mathcal{C}} {}_c[L]_{\mathcal{B}} [\vec{v}]_{\mathcal{B}} \\ &= {}_{c'}\mathcal{I}_{\mathcal{C}} [L(\vec{v})]_{\mathcal{C}} \\ &= [L(\vec{v})]_{\mathcal{C}'} \\ &= {}_{c'}[L]_{\mathcal{B}'} [\vec{v}]_{\mathcal{B}'} \end{aligned}$$

for all $\vec{v} \in V$. This completes the proof (why?). \square

Example 2.4.7

Let \mathcal{S} and \mathcal{B} be the bases for $\mathcal{P}_2(\mathbb{R})$ from Examples 2.4.2 and 2.4.1, where we had determined ${}_{\mathcal{B}}\mathcal{I}_{\mathcal{S}}$ and ${}_{\mathcal{S}}\mathcal{I}_{\mathcal{B}}$.

Consider the linear map $D: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathcal{P}_2(\mathbb{R})$ given by differentiation. We'll leave it as an exercise for you to check that $[D]_{\mathcal{S}} = {}_{\mathcal{S}}[D]_{\mathcal{S}} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix}$. Let's find $[D]_{\mathcal{B}} = {}_{\mathcal{B}}[D]_{\mathcal{B}}$. We can do so either directly from the definition or by using Proposition 2.4.6.

For the direct approach, we simply note that $[D(1)]_{\mathcal{B}} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$, $[D(1+x)]_{\mathcal{B}} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$, and

$[D(1+x+x^2)]_{\mathcal{B}} = \begin{bmatrix} -1 \\ 2 \\ 0 \end{bmatrix}$. Therefore

$$[D]_{\mathcal{B}} = \begin{bmatrix} 0 & 1 & -1 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix}.$$

On the other hand, using Proposition 2.4.6, we get

$$[D]_{\mathcal{B}} = {}_{\mathcal{B}}\mathcal{I}_{\mathcal{S}} [D]_{\mathcal{S}} {}_{\mathcal{S}}\mathcal{I}_{\mathcal{B}} = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 1 & -1 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix},$$

as expected!

2.5 Isomorphisms of Vector Spaces

Let's return to an observation that has come up a few times so far: \mathbb{R}^3 and $\mathcal{P}_2(\mathbb{R})$ are essentially the same! We could just rename the element $[a \ b \ c]^T$ to $a + bx + cx^2$ and everything would work exactly the same. Somehow it feels like these two elements are the same thing called by different names. We will soon see that the vector spaces \mathbb{R}^3 and $\mathcal{P}_2(\mathbb{R})$, while technically different objects, have exactly the same structure. More precisely, we will see that \mathbb{R}^3 and $\mathcal{P}_2(\mathbb{R})$ are *isomorphic*.

An isomorphism between vector spaces should be thought of like a translator. It's a linear map that preserves information perfectly. No information is lost, and no information is missed. Let's look at a couple of examples to gain a little more intuition.

Example 2.5.1

Consider the linear map $L: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathbb{R}^2$ given by $L(p) = \begin{bmatrix} p(0) \\ p(0) \end{bmatrix}$. This linear map is not an isomorphism because somehow it loses information. For example, $L(x + 2) = L(x^2 + 2) = L(2) = \begin{bmatrix} 2 \\ 2 \end{bmatrix}$ so just by looking at the output of L , we can't tell the difference between $x + 2$ and 2 for example. Furthermore, L somehow misses information. For example, nothing maps to the vector $\begin{bmatrix} 1 \\ 3 \end{bmatrix}$.

Example 2.5.2

On the other hand, consider the linear map $L: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathbb{R}^3$ given by

$$L(p) = \begin{bmatrix} p(-1) \\ p(0) \\ p(1) \end{bmatrix}.$$

There are two very interesting things about this map. First, it turns out that if you know the value of a polynomial in $\mathcal{P}_2(\mathbb{R})$ evaluated at three distinct points, you are able to recover the polynomial. That is, if $L(p) = L(q)$ then $p = q$. Furthermore, for any three numbers $a, b, c \in \mathbb{R}$, there is a polynomial $p \in \mathcal{P}_2(\mathbb{R})$ such that $p(-1) = a$, $p(0) = b$, and $p(1) = c$. Therefore, $\text{Range}(L) = \mathbb{R}^3$. With these two pieces of information, we can see that L is a perfect dictionary between $\mathcal{P}_2(\mathbb{R})$ and \mathbb{R}^3 and both vector spaces contain the same information, just wrapped up in a different package.

Roughly, an isomorphism of vector spaces will be a linear map $L: V \rightarrow W$ which is a perfect dictionary between V and W , that is, no information is lost, and no information is missed, after applying L . More formally, it will be a linear map that is injective and surjective, which we will now define.

Definition 2.5.3

Injective
(One-to-One),
Surjective (Onto)

Let $L: V \rightarrow W$ be a linear map between vector spaces and let $\vec{v}_1, \vec{v}_2 \in V$.

We say L is **injective** (or **one-to-one**) if $L(\vec{v}_1) = L(\vec{v}_2)$ implies $\vec{v}_1 = \vec{v}_2$.

We say L is **surjective** (or **onto**) if $\text{Range}(L) = W$.

Here is a little result that will make checking injectivity that much easier.

Lemma 2.5.4 A linear map $L: V \rightarrow W$ is injective if and only if $\text{Ker}(L) = \{\vec{0}\}$.

Proof: Suppose $L: V \rightarrow W$ is injective and let $\vec{v} \in \text{Ker}(L)$. Then $L(\vec{v}) = L(\vec{0}) = \vec{0}$ so $\vec{v} = \vec{0}$. Therefore $\text{Ker}(L) \subseteq \{\vec{0}\}$, and since the reverse containment is obvious, it follows that $\text{Ker}(L) = \{\vec{0}\}$. Conversely, suppose $\text{Ker}(L) = \{\vec{0}\}$ and let $L(\vec{v}) = L(\vec{w})$. Then $\vec{0} = L(\vec{v}) - L(\vec{w}) = L(\vec{v} - \vec{w})$ so $\vec{v} - \vec{w} \in \text{Ker}(L)$. Since the only vector in $\text{Ker}(L)$ is $\vec{0}$, we have $\vec{v} - \vec{w} = \vec{0}$ so $\vec{v} = \vec{w}$, completing the proof. \square

We now have all the pieces lined up for the following useful result.

Proposition 2.5.5 Let $L: V \rightarrow W$ be a linear map between finite-dimensional vector spaces. Then:

- (a) L is injective if and only if $\text{nullity}(L) = 0$.
- (b) L is surjective if and only if $\text{rank}(L) = \dim W$.

Exercise 20 Prove Proposition 2.5.5.

Example 2.5.6 Consider the linear map $L: \mathcal{P}_3(\mathbb{R}) \rightarrow \mathbb{R}^4$ given by

$$L(p(x)) = \begin{bmatrix} p(0) \\ p'(0) \\ p''(0) \\ p'''(0) \end{bmatrix},$$

where $p'(0)$, $p''(0)$, and $p'''(0)$ are the first, second, and third derivatives of p , evaluated at 0, respectively.

Now suppose $p(x) = a + bx + cx^2 + dx^3 \in \text{Ker}(L)$. Then $p(0) = p'(0) = p''(0) = p'''(0) = 0$ which yield the equations $a = 0$, $b = 0$, $2c = 0$, and $6d = 0$. We must conclude that $p(x) = 0$. Therefore the only polynomial in $\text{Ker}(L)$ is the zero polynomial, so $\text{nullity}(L) = 0$. By Proposition 2.5.5, we can immediately conclude that L is injective.

Since $\mathcal{P}_3(\mathbb{R})$ is a 4-dimensional vector space, the Rank–Nullity theorem tells us $\text{rank}(L) = 4$. Since $\dim(\mathbb{R}^4) = 4$, we can again exploit Proposition 2.5.5 to conclude L is surjective.

The injectivity and surjectivity of L in this example is telling us something very interesting about polynomials of degree at most 3. The injectivity says that such a polynomial p is entirely determined by the 4 numbers $p(0)$, $p'(0)$, $p''(0)$, and $p'''(0)$. The surjectivity of L says that given any 4 real numbers, you can find a polynomial p of degree at most 3 so that $p(0)$, $p'(0)$, $p''(0)$, and $p'''(0)$ are the desired 4 real numbers.

If V and W are finite-dimensional, then we can check the injectivity and surjectivity of any given linear map $L: V \rightarrow W$ rather easily using the matrices and coordinates, as follows.

Begin by choosing ordered bases \mathcal{B} and \mathcal{C} for V and W , respectively, and let $A = {}_c[L]_{\mathcal{B}}$. We've shown in Proposition 2.3.11 that for all $\vec{v} \in V$,

$$\vec{v} \in \text{Ker}(L) \quad \text{if and only if} \quad [\vec{v}]_{\mathcal{B}} \in \text{Null}(A).$$

Likewise, for all $\vec{w} \in W$,

$$\vec{w} \in \text{Range}(L) \quad \text{if and only if} \quad [\vec{w}]_{\mathcal{C}} \in \text{Col}(A).$$

Since we know how to compute the nullspace and column space of a matrix, it follows that we can easily determine $\text{Ker}(L)$ and $\text{Range}(L)$, and therefore whether L is injective or surjective.

Example 2.5.7

Let's re-do Example 2.5.6 and check the injectivity and surjectivity of L .

First, we need bases for $\mathcal{P}_3(\mathbb{R})$ and \mathbb{R}^4 . Let's use the standard bases \mathcal{B} and \mathcal{C} , respectively. Then the matrix of L with respect to these bases is

$$A = {}_c[L]_{\mathcal{B}} = \begin{bmatrix} [L(1)]_{\mathcal{C}} & [L(x)]_{\mathcal{C}} & [L(x^2)]_{\mathcal{C}} & [L(x^3)]_{\mathcal{C}} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix}.$$

Using our usual methods, we can quickly determine that $\text{Null}(A) = \{\vec{0}\}$ and $\text{Col}(A) = \mathbb{R}^4$. From this, we can conclude that $p(x) \in \mathcal{P}_3(\mathbb{R})$ will be in the kernel of L if and only if $[p(x)]_{\mathcal{B}} \in \text{Null}(A) = \{\vec{0}\}$. This can only be the case if $p(x) = 0$ is the zero polynomial. Thus, $\text{Ker}(L) = \{\vec{0}\}$ and so L is injective.

Similarly, $\vec{w} \in \mathbb{R}^4$ will be in the range of L if and only if $[\vec{w}]_{\mathcal{C}} \in \text{Col}(A) = \mathbb{R}^4$, which is always true. Thus, $\text{Range}(L) = \mathbb{R}^4$ and so L is surjective.

Exercise 21

Let $L: M_{2 \times 2}(\mathbb{R}) \rightarrow \mathcal{P}_2(\mathbb{R})$ be the linear map defined by

$$L\left(\begin{bmatrix} a & b \\ c & d \end{bmatrix}\right) = (a+b) + (b+c)x + (c+d)x^2.$$

- Determine $A = {}_c[L]_{\mathcal{B}}$, where \mathcal{B} and \mathcal{C} are the standard ordered bases for $M_{2 \times 2}(\mathbb{R})$ and $\mathcal{P}_2(\mathbb{R})$, respectively.
- Determine $\text{Null}(A)$ and $\text{Col}(A)$.
- Conclude whether L is injective or surjective (or neither).

If you study the previous discussion carefully, you'll come away feeling that $\text{Ker}(L)$ and $\text{Null}({}_c[L]_{\mathcal{B}})$ are *almost* the same thing; similarly, $\text{Range}(L)$ and $\text{Col}({}_c[L]_{\mathcal{B}})$ seem to contain the same information. Of course, these vector spaces are not literally equal—since they can consist of different type of vectors like in the above exercise, where $\text{Ker}(L)$ consists of polynomials in $M_{2 \times 2}(\mathbb{R})$ while $\text{Null}({}_c[L]_{\mathcal{B}})$ consists of column vectors in \mathbb{R}^4 .

This type of relationship between vector spaces is similar to the example we gave in the beginning of this section, where we indicated that $\mathcal{P}_2(\mathbb{R})$ and \mathbb{R}^3 appear to have the same

structure, since $a + bx + cx^2$ and $[a \ b \ c]^T$ feel like two different ways of writing down the same object.

What we appear to be doing is associating vectors of different spaces together. For example, $v \leftrightarrow [v]_{\mathcal{B}}$ is the association between $\text{Ker}(L)$ and $\text{Null}(c[L]_{\mathcal{B}})$, and $a + bx + cx^2 \leftrightarrow \begin{bmatrix} a \\ b \\ c \end{bmatrix}$ is the association between $\mathcal{P}_2(\mathbb{R})$ and \mathbb{R}^3 . What we are *actually* doing, though, is writing down special linear maps between the various spaces!

Definition 2.5.8

**Isomorphism,
Isomorphic to**

Let $L: V \rightarrow W$ be a linear map. If L is injective and surjective, we say L is an **isomorphism**.

If there is an isomorphism $L: V \rightarrow W$, we say that V is **isomorphic to** W , and write $V \cong W$.

In Proposition 2.5.23 below, we will prove that if V is isomorphic to W , then W will be isomorphic to V . As such, we can simply say that V and W are *isomorphic*. The use of the symbol \cong is then apt, because

$$V \cong W \quad \text{if and only if} \quad W \cong V.$$

You should also check for yourself that $V \cong V$; and if $V \cong W$ and $W \cong U$, then $V \cong U$ (we say that \cong is *transitive*). So isomorphism behaves a lot like equality—but be careful: they're not the same thing!

Example 2.5.9

Consider $L: M_{2 \times 2}(\mathbb{C}) \rightarrow \mathbb{C}^3$ given by $L\left(\begin{bmatrix} a & b \\ c & d \end{bmatrix}\right) = \begin{bmatrix} a + b \\ b - 2c \\ a + b + d \end{bmatrix}$. Then $\begin{bmatrix} 2 & -2 \\ -1 & 0 \end{bmatrix} \in \text{Ker}(L)$ so L is not injective, and is therefore not an isomorphism.

Example 2.5.10

Let $L: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathbb{R}^3$ be given by $L(a + bx + cx^2) = \begin{bmatrix} a \\ b \\ c \end{bmatrix}$.

You should check that $\text{Ker}(L) = \{\vec{0}\}$. We can also check directly that $\text{Range}(L) = \mathbb{R}^3$, but here is a clever indirect argument: Since $\text{nullity}(L) = 0$, the Rank–Nullity theorem implies that $\text{rank}(L) = 3$. So $\text{Range}(L)$ is a 3-dimensional subspace of \mathbb{R}^3 , and therefore it must be equal to \mathbb{R}^3 . Thus, L is an isomorphism and $\mathcal{P}_2(\mathbb{R})$ is isomorphic to \mathbb{R}^3 (so we can write $\mathcal{P}_2(\mathbb{R}) \cong \mathbb{R}^3$).

Exercise 22

For L as in the previous example, show using the definitions of kernel and range that $\text{Ker}(L) = \{\vec{0}\}$ and $\text{Range}(L) = \mathbb{R}^3$.

There can be more than one isomorphism between isomorphic vector spaces.

Example 2.5.11 The linear map $L: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathbb{R}^3$ defined by $L(a + bx + cx^2) = \begin{bmatrix} c \\ b \\ a \end{bmatrix}$ is also an isomorphism between $\mathcal{P}_2(\mathbb{R})$ and \mathbb{R}^3 . The proof is similar to the one in Example 2.5.10.

Exercise 23

- (a) Show that linear map $L: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathbb{R}^3$ defined by $L(a + bx + cx^2) = \begin{bmatrix} a \\ a + b \\ a + b + c \end{bmatrix}$ is yet another isomorphism between $\mathcal{P}_2(\mathbb{R})$ and \mathbb{R}^3 .
- (b) Come up with another isomorphism $L: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathbb{R}^3$.

Be careful to note that just because $V \cong W$, it does **not** follow that every linear map $L: V \rightarrow W$ is an isomorphism! Consider, for example, the linear map $L: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathbb{R}^3$ given by $L(p) = \vec{0}$ for all $p \in \mathcal{P}_2(\mathbb{R})$. Then $\text{nullity}(L) = 3$ so L is not injective, even though $\mathcal{P}_2(\mathbb{R}) \cong \mathbb{R}^3$. Isomorphisms are rather special linear maps.

Our next two results give very important examples of isomorphisms. The first is a reformulation of Proposition 2.3.11.

Proposition 2.5.12

Let $L: V \rightarrow W$ be a linear map between finite-dimensional vector spaces, and let $A = {}_C[L]_{\mathcal{B}}$, where \mathcal{B} and \mathcal{C} are ordered bases for V and W respectively.

- (a) $\text{Ker}(L)$ is isomorphic to $\text{Null}(A)$.
- (b) $\text{Range}(L)$ is isomorphic to $\text{Col}(A)$.

Proof: The coordinate map $[\]_{\mathcal{B}}$ is linear (Theorem 1.3.29) and injective (Theorem 1.3.23), and Proposition 2.3.11 says it maps $\text{Ker}(L)$ surjectively onto $\text{Null}(A)$. So it gives an isomorphism $\text{Ker}(L) \cong \text{Null}(A)$. Similarly, $[\]_{\mathcal{C}}$ is injective and maps $\text{Range}(L)$ surjectively onto $\text{Col}(A)$ hence gives an isomorphism $\text{Range}(L) \cong \text{Col}(A)$. \square

The second result is a generalization of Example 2.5.10.

Proposition 2.5.13

Let V be an n -dimensional space over \mathbb{F} . If \mathcal{B} is an ordered basis for V , then the coordinate map $[\]_{\mathcal{B}}: V \rightarrow \mathbb{F}^n$ defined by sending \vec{v} to $[\vec{v}]_{\mathcal{B}}$ is an isomorphism.

Proof sketch: The fact that $[\]_{\mathcal{B}}$ is injective follows from the Unique Representation Theorem (Theorem 1.3.23). To show that $[\]_{\mathcal{B}}$ is surjective, we can mimic the Rank–Nullity argument in Example 2.5.10. \square

Exercise 24

Fill in the details and complete the proof of Proposition 2.5.13.

This proposition says that once we pick a basis, we can view any n -dimensional vector space V as essentially being the same as \mathbb{F}^n . Before showcasing the power of this result and the previous one, let's prove some general facts that will demonstrate how isomorphisms preserve structure perfectly. For instance, the next proposition shows that isomorphisms preserve linear independence, spanning sets and bases.

Proposition 2.5.14

Let $L: V \rightarrow W$ be an isomorphism. Then:

- (a) If $\{\vec{v}_1, \dots, \vec{v}_n\}$ is a linearly independent set of vectors in V , then $\{L(\vec{v}_1), \dots, L(\vec{v}_n)\}$ is a linearly independent set of vectors in W .
- (b) If $\{\vec{v}_1, \dots, \vec{v}_n\}$ is a spanning set for V (meaning, $V = \text{Span}\{\vec{v}_1, \dots, \vec{v}_n\}$), then $\{L(\vec{v}_1), \dots, L(\vec{v}_n)\}$ is a spanning set for W (meaning, $W = \text{Span}\{L(\vec{v}_1), \dots, L(\vec{v}_n)\}$).
- (c) If $\{\vec{v}_1, \dots, \vec{v}_n\}$ is a basis for V , then $\{L(\vec{v}_1), \dots, L(\vec{v}_n)\}$ is a basis for W .

Proof: To prove (a), consider the equation

$$t_1 L(\vec{v}_1) + \dots + t_n L(\vec{v}_n) = \vec{0}.$$

By linearity, we can re-write this as

$$L(t_1 \vec{v}_1 + \dots + t_n \vec{v}_n) = \vec{0}.$$

This shows that $t_1 \vec{v}_1 + \dots + t_n \vec{v}_n \in \text{Ker}(L)$. However, $\text{Ker}(L) = \{\vec{0}\}$ since L is an isomorphism. Thus,

$$t_1 \vec{v}_1 + \dots + t_n \vec{v}_n = \vec{0}$$

and therefore $t_1 = \dots = t_n = 0$ since $\{\vec{v}_1, \dots, \vec{v}_n\}$ is linearly independent.

Next, to prove (b), notice that it suffices to show that $W \subseteq \text{Span}\{L(\vec{v}_1), \dots, L(\vec{v}_n)\}$, since the containment \supseteq is obvious. Thus, let $\vec{w} \in W$. Since L is an isomorphism, hence surjective, we have that $\vec{w} \in \text{Range}(L)$. So there exists a $\vec{v} \in V$ such that $L(\vec{v}) = \vec{w}$. Since $\{\vec{v}_1, \dots, \vec{v}_n\}$ is a spanning set for V , we can write

$$\vec{v} = a_1 \vec{v}_1 + \dots + a_n \vec{v}_n$$

for some $a_i \in \mathbb{F}$. But then

$$\vec{w} = L(\vec{v}) = L(a_1 \vec{v}_1 + \dots + a_n \vec{v}_n) = a_1 L(\vec{v}_1) + \dots + a_n L(\vec{v}_n),$$

which shows that $\vec{w} \in \text{Span}\{L(\vec{v}_1), \dots, L(\vec{v}_n)\}$ and completes the proof of (b).

Finally, (c) follows from (a) and (b). □

We can apply this proposition to the coordinate map $[\]_{\mathcal{B}}$.

Example 2.5.15

In Example 1.3.7, we proved that $\{x + x^2 - 2x^3, 2x - x^2 + x^3, x + 5x^2 + 3x^3\}$ is linearly independent in $\mathcal{P}_3(\mathbb{R})$. We can give another proof of this fact by showing that $\{[x + x^2 - 2x^3]_{\mathcal{B}}, [2x - x^2 + x^3]_{\mathcal{B}}, [x + 5x^2 + 3x^3]_{\mathcal{B}}\}$ is linearly independent in \mathbb{R}^4 , where \mathcal{B} is any ordered basis for $\mathcal{P}_3(\mathbb{R})$.

For example, if \mathcal{B} is the standard ordered basis of $\mathcal{P}_3(\mathbb{R})$, then

$$\{[x + x^2 - 2x^3]_{\mathcal{B}}, [2x - x^2 + x^3]_{\mathcal{B}}, [x + 5x^2 + 3x^3]_{\mathcal{B}}\} = \left\{ \begin{bmatrix} 0 \\ 1 \\ 1 \\ -2 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ -1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 5 \\ 3 \end{bmatrix} \right\}.$$

To check for linear independence in \mathbb{R}^4 , we can simply put these vectors in a matrix and row-reduce:

$$\begin{bmatrix} 0 & 0 & 0 \\ 1 & 2 & 1 \\ 1 & -1 & 5 \\ -2 & 1 & 3 \end{bmatrix} \longrightarrow \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}.$$

We conclude that the rank of this matrix is 3, and hence its columns must be linearly independent, which is exactly what we wanted to show.

Exercise 25

Show, using Proposition 2.5.14, that

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

is a basis for $M_{2 \times 2}(\mathbb{R})$. (Compare with what we did in Examples 1.3.11 and 1.3.17.)

Our next two results examine the relationship between dimension and isomorphism.

If we are to think of an isomorphism as simply being a renaming of vectors, which we should, then we should expect two isomorphic vector spaces to have the same structure. At the very least, it wouldn't be unreasonable to expect two isomorphic finite-dimensional vector spaces to have the same dimension. In fact, this follows immediately from Proposition 2.5.14(c). What's perhaps surprising is that the converse is also true: if V and W are finite-dimensional spaces, and if $\dim(V) = \dim(W)$, then V and W must be isomorphic! Before proving this, let's record the following related result which can be quite useful.

Proposition 2.5.16

Let $L: V \rightarrow W$ be a linear map between finite-dimensional vector spaces.

- (a) If $\dim(V) < \dim(W)$, then L cannot be surjective.
- (b) If $\dim(V) > \dim(W)$, then L cannot be injective.
- (c) If $\dim(V) = \dim(W)$, then L is injective if and only if L is surjective.

Proof: If $\dim(V) < \dim(W)$, then the Rank–Nullity theorem implies that $\text{Range}(L)$ cannot be all of W , so L cannot be surjective. If $\dim(V) > \dim(W)$, the Rank–Nullity theorem

implies that $\text{nullity}(L) \geq 1$, so L cannot be injective. Finally, if $\dim(V) = \dim(W)$, then the Rank–Nullity theorem says that $\dim(V) = \text{rank}(L) + \text{nullity}(L)$ and so our desired result follows from Proposition 2.5.5:

$$\begin{aligned} L \text{ is injective} &\iff \text{nullity}(L) = 0 \\ &\iff \dim(V) = \text{rank}(L) \\ &\iff \dim(W) = \text{rank}(L) \\ &\iff L \text{ is surjective.} \end{aligned}$$

This completes the proof. \square

Exercise 26

Give a careful proof of parts (a) and (b) Proposition 2.5.16 by filling in the details in the above proof.

Theorem 2.5.17

Suppose V and W are finite dimensional vector spaces over the same field \mathbb{F} . Then V and W are isomorphic if and only if $\dim(V) = \dim(W)$.

Proof: Suppose $V \cong W$ via an isomorphism $L: V \rightarrow W$. Then by Proposition 2.5.14(c), L takes a basis for V to a basis for W . So if V has a basis consisting of n elements, then so does W . Thus, $\dim(V) = \dim(W)$.

Conversely, let $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$ be a basis for V and $\mathcal{C} = \{\vec{w}_1, \dots, \vec{w}_n\}$ a basis for W . We want to define an isomorphism $L: V \rightarrow W$. The idea is simple: We should send \vec{v}_i to \vec{w}_i . By linearity, there is only one way to define a linear map that does this. Namely, define $L: V \rightarrow W$ by

$$L(t_1 \vec{v}_1 + \dots + t_n \vec{v}_n) = t_1 \vec{w}_1 + \dots + t_n \vec{w}_n.$$

Note that L is linear since

$$\begin{aligned} L\left(\sum_{i=1}^n t_i \vec{v}_i + \sum_{i=1}^n s_i \vec{v}_i\right) &= L\left(\sum_{i=1}^n (t_i + s_i) \vec{v}_i\right) \\ &= \sum_{i=1}^n (t_i + s_i) \vec{w}_i \\ &= \sum_{i=1}^n t_i \vec{w}_i + \sum_{i=1}^n s_i \vec{w}_i \\ &= L\left(\sum_{i=1}^n t_i \vec{v}_i\right) + L\left(\sum_{i=1}^n s_i \vec{v}_i\right) \end{aligned}$$

and for all $\alpha \in \mathbb{F}$,

$$\begin{aligned} L\left(\alpha \sum_{i=1}^n t_i \vec{v}_i\right) &= L\left(\sum_{i=1}^n \alpha t_i \vec{v}_i\right) \\ &= \sum_{i=1}^n \alpha t_i \vec{w}_i \\ &= \alpha \sum_{i=1}^n t_i \vec{w}_i \\ &= \alpha L\left(\sum_{i=1}^n t_i \vec{v}_i\right). \end{aligned}$$

To see that L is injective, suppose that $L(t_1\vec{v}_1 + \cdots + t_n\vec{v}_n) = t_1\vec{w}_1 + \cdots + t_n\vec{w}_n = \vec{0}$. Then since $\{\vec{w}_1, \dots, \vec{w}_n\}$ is linearly independent, we must have $t_1 = \cdots = t_n = 0$ so $t_1\vec{v}_1 + \cdots + t_n\vec{v}_n = \vec{0}$ and so $\text{Ker}(L) = \{\vec{0}\}$. Finally, the Rank–Nullity theorem implies $\text{rank}(L) = \dim(V) = \dim(W)$ so L is surjective and is therefore an isomorphism. \square

REMARK

This is an incredibly powerful theorem. It immediately tells us, for example, that any two 7-dimensional vector spaces over \mathbb{C} are isomorphic. Furthermore, to find an isomorphism, we simply have to choose bases for both vector spaces and the map that appears in the proof of the theorem will be an isomorphism.

Corollary 2.5.18

Let V be an n -dimensional vector space over \mathbb{F} . Then $V \cong \mathbb{F}^n$.

In fact, an isomorphism is given by the coordinate map $[\]_{\mathcal{B}}: V \rightarrow \mathbb{F}^n$, where \mathcal{B} is any ordered basis for V , as we've already seen.

Example 2.5.19

The vector spaces $M_{2 \times 2}(\mathbb{R})$, \mathbb{R}^4 and $\mathcal{P}_3(\mathbb{R})$ are all isomorphic to each other, since they are each 4-dimensional.

On the other hand, the vector spaces $M_{2 \times 3}(\mathbb{R})$ and $\mathcal{P}_4(\mathbb{R})$ are not isomorphic since the former is 6-dimensional while the latter is 5-dimensional.

Exercise 27

Give an isomorphism $L: M_{2 \times 2}(\mathbb{R}) \rightarrow \mathcal{P}_3(\mathbb{R})$ by following the recipe in the proof of Theorem 2.5.17.

If we apply Theorem 2.5.17 to Proposition 2.5.12, we immediately arrive at the following handy result.

Proposition 2.5.20

Let $L: V \rightarrow W$ be a linear map between finite-dimensional vector spaces, and let $A = {}_C[L]_{\mathcal{B}}$, where \mathcal{B} and \mathcal{C} are ordered bases for V and W respectively. Then:

- (a) $\text{nullity}(L) = \text{nullity}(A)$.
- (b) $\text{rank}(L) = \text{rank}(A)$.

Using this proposition, we obtain the Rank–Nullity theorem for matrices as a consequence of our Rank–Nullity theorem for linear maps (Theorem 2.2.6).

Corollary 2.5.21

(Rank–Nullity Theorem for Matrices)

Let $A \in M_{m \times n}(\mathbb{F})$. Then

$$n = \text{rank}(A) + \text{nullity}(A).$$

Proof: Let $L: \mathbb{F}^n \rightarrow \mathbb{F}^m$ be the linear map corresponding to A , that is, the one defined by $L(\vec{x}) = A\vec{x}$ for $\vec{x} \in \mathbb{F}^n$. Then with respect to the standard bases \mathcal{B} and \mathcal{C} on \mathbb{F}^n and \mathbb{F}^m , respectively, we have $A = {}_{\mathcal{C}}[L]_{\mathcal{B}}$. So, by combining the Rank–Nullity theorem for linear maps and Proposition 2.5.20, we get

$$n = \dim(\mathbb{F}^n) = \text{rank}(L) + \text{nullity}(L) = \text{rank}(A) + \text{nullity}(A),$$

as required. \square

Example 2.5.22

Consider the linear map $L: \mathcal{P}_3(\mathbb{R}) \rightarrow M_{2 \times 2}(\mathbb{R})$ defined by

$$L(a + bx + cx^2 + dx^3) = \begin{bmatrix} a - b & c - d \\ b + d & a + 2c \end{bmatrix}.$$

Is L injective? Is L surjective? Is L an isomorphism?

Solution:

We will use Proposition 2.5.20 to compute $\text{nullity}(L)$ and $\text{rank}(L)$. We begin by determining $A = {}_{\mathcal{C}}[L]_{\mathcal{B}}$, where \mathcal{B} and \mathcal{C} are the standard ordered bases of $\mathcal{P}_3(\mathbb{R})$ and $M_{2 \times 2}(\mathbb{R})$:

$$A = [{}_{\mathcal{C}}[L(1)]_{\mathcal{C}} \quad {}_{\mathcal{C}}[L(x)]_{\mathcal{C}} \quad {}_{\mathcal{C}}[L(x^2)]_{\mathcal{C}} \quad {}_{\mathcal{C}}[L(x^3)]_{\mathcal{C}}] = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 2 & 0 \end{bmatrix}.$$

If we row-reduce A , we get

$$\begin{bmatrix} 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 2 & 0 \end{bmatrix} \longrightarrow \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

From this we can immediately conclude that $\text{nullity}(A) = 0$ and $\text{rank}(A) = 4$. Therefore, $\text{nullity}(L) = 0$ (so L is injective) and $\text{rank}(L) = 4$ (so L is surjective). In particular, L is an isomorphism.

Finally, if the intuition that an isomorphism is a kind of translator is correct, then there should be a way to do an isomorphism in reverse, just like you should be able to translate a word back into English, if you had already translated it into Spanish. This is indeed true, and the next proposition makes it precise.

Proposition 2.5.23

A linear map $L: V \rightarrow W$ is an isomorphism if and only if there exists a linear map $L^{-1}: W \rightarrow V$ such that

$$L \circ L^{-1}(\vec{w}) = \vec{w} \text{ for all } \vec{w} \in W \quad \text{and} \quad L^{-1} \circ L(\vec{v}) = \vec{v} \text{ for all } \vec{v} \in V.$$

Such a map L^{-1} , if it exists, is uniquely determined by L . We call L^{-1} the **inverse** of L . It is also an isomorphism.

Proof sketch: Given an isomorphism $L: V \rightarrow W$, define $L^{-1}: W \rightarrow V$ by $L^{-1}(\vec{w}) = \vec{v}$ where $\vec{v} \in V$ is the *unique* vector such that $L(\vec{v}) = \vec{w}$. Such a vector exists because L is surjective; it is unique because L is injective. It is left to you to prove that L^{-1} is a linear map satisfying the desired properties. For the converse direction, you should check that if such an inverse map exists, then L must necessarily be surjective and injective.

For uniqueness, suppose that $T: W \rightarrow V$ also satisfies the properties $L \circ T(\vec{w}) = \vec{w}$ for all $\vec{w} \in W$ and $T \circ L(\vec{v}) = \vec{v}$ for all $\vec{v} \in V$. Then $T(\vec{w}) = T(L \circ L^{-1}(\vec{w})) = (T \circ L)(L^{-1}(\vec{w})) = L^{-1}(\vec{w})$ for all $\vec{w} \in W$. So $T = L^{-1}$.

We'll leave it to you to check that L^{-1} , if it exists, is an isomorphism. \square

Exercise 28

Fill in the details and complete the proof of Proposition 2.5.23.

Given an isomorphism, it is sometimes very easy to write down the inverse linear map, and sometimes not. For example, return to the isomorphism $L: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathbb{R}^3$ given by

$$L(a + bx + cx^2) = \begin{bmatrix} a \\ b \\ c \end{bmatrix}. \text{ Then } L^{-1}: \mathbb{R}^3 \rightarrow \mathcal{P}_2(\mathbb{R}) \text{ is given by } L^{-1}\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) = a + bx + cx^2.$$

Let's check this is indeed the inverse. We have

$$L \circ L^{-1}\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) = L(a + bx + cx^2) = \begin{bmatrix} a \\ b \\ c \end{bmatrix}$$

and

$$L^{-1} \circ L(a + bx + cx^2) = L^{-1}\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) = a + bx + cx^2$$

so this is the inverse.

On the other hand, what is the inverse to the isomorphism $L: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathbb{R}^3$ given by

$$L(p) = \begin{bmatrix} p(-1) \\ p(0) \\ p(1) \end{bmatrix}?$$

The next proposition, which is a consequence of Proposition 2.5.23, gives us a way to find inverses to isomorphisms. It is yet another demonstration of the power of working with matrices.

Proposition 2.5.24

Let $L: V \rightarrow W$ be an isomorphism. Let \mathcal{B} be a basis for V , and \mathcal{C} a basis for W . Then $c[L]_{\mathcal{B}}$ is an invertible matrix and $(c[L]_{\mathcal{B}})^{-1} = \mathcal{B}[L^{-1}]_{\mathcal{C}}$.

Proof: We have $L \circ L^{-1} = \text{id}$, where $\text{id}: W \rightarrow W$ is the identity map on W . Therefore,

$$c[L \circ L^{-1}]_{\mathcal{C}} = c[\text{id}]_{\mathcal{C}} = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \cdots & 0 & 1 \end{bmatrix}.$$

On the other hand, $c[L \circ L^{-1}]_{\mathcal{C}} = c[L]_{\mathcal{B}} \mathcal{B}[L^{-1}]_{\mathcal{C}}$ by Proposition 2.3.8. So we see that $c[L]_{\mathcal{B}} \mathcal{B}[L^{-1}]_{\mathcal{C}}$ is equal to the identity matrix, hence $c[L]_{\mathcal{B}}^{-1} = \mathcal{B}[L^{-1}]_{\mathcal{C}}$. \square

Let's see this in action!

Example 2.5.25

Let $L: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathbb{R}^3$ be the isomorphism given by $L(p) = \begin{bmatrix} p(-1) \\ p(0) \\ p(1) \end{bmatrix}$. We have already seen that

$${}_C[L]_{\mathcal{B}} = \begin{bmatrix} 1 & -1 & 1 \\ 1 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix}.$$

Use your favourite method of computing the inverse of a matrix to show that

$${}_{\mathcal{B}}[L^{-1}]_C = \begin{bmatrix} 0 & 1 & 0 \\ -\frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{2} & -1 & \frac{1}{2} \end{bmatrix}.$$

Now, since

$$\begin{bmatrix} 0 & 1 & 0 \\ -\frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{2} & -1 & \frac{1}{2} \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} b \\ -\frac{1}{2}a + \frac{1}{2}c \\ \frac{1}{2}a - b + \frac{1}{2}c \end{bmatrix}$$

we have

$$L^{-1} \left(\begin{bmatrix} a \\ b \\ c \end{bmatrix} \right) = b + \left(-\frac{1}{2}a + \frac{1}{2}c \right) x + \left(\frac{1}{2}a - b + \frac{1}{2}c \right) x^2.$$

Chapter 3

Diagonalizability

3.1 Eigenvectors and Diagonalization

As hinted to before, sometimes the standard basis is not the best basis with which to study a particular problem, or a linear map. For example, consider the linear map $L: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ given by

$$L \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \right) = \begin{bmatrix} x \\ y \\ z \end{bmatrix} - \frac{2(x+y+z)}{3} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}.$$

If \mathcal{S} is the standard basis for \mathbb{R}^3 , then you can check that

$$[L]_{\mathcal{S}} = \begin{bmatrix} \frac{1}{3} & -\frac{2}{3} & -\frac{2}{3} \\ -\frac{2}{3} & \frac{1}{3} & -\frac{2}{3} \\ -\frac{2}{3} & -\frac{2}{3} & \frac{1}{3} \end{bmatrix}.$$

Looking at this matrix, it's not clear what this linear map is doing, geometrically or otherwise. However, if we look at the basis

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

then it can be checked that

$$[L]_{\mathcal{B}} = \begin{bmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

The effect of this matrix is easier to understand. It is a reflection! It negates the $[1 \ 1 \ 1]^T$ direction and keeps the 2-dimensional subspace spanned by $[-1 \ 0 \ 1]^T$ and $[-1 \ 1 \ 0]^T$ unchanged.

This example shows us that sometimes looking at a particular problem with the right set of coordinates can prove enlightening. So, with this in mind, the following natural question arises: Given a linear map from a vector space to itself, how can we find an “enlightening” basis with which to view the linear map?

In this chapter we will restrict our attention to linear maps from a vector space to itself—and not between two different vector spaces. This is a fairly natural starting point, and we'll see that it leads to a reasonably nice theory.

Definition 3.1.1**Linear Operator**

A linear map $T: V \rightarrow W$ is called a **linear operator** if $V = W$.

It would be nice to find vectors that are not rotated, but simply scaled when the linear map is applied to them. That is, we'd like to find vectors \vec{v} such that $L(\vec{v}) = \lambda\vec{v}$ for some $\lambda \in \mathbb{F}$. If we can find a basis $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$ of V such that $L(\vec{v}_i) = \lambda_i\vec{v}_i$ for every i , then with respect to \mathcal{B} we would have

$$[L]_{\mathcal{B}} = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_n \end{bmatrix}.$$

Unfortunately, as we will see, we cannot always find such a basis. Let's try anyway!

What we're attempting to do should ring some bells: we are trying to *diagonalize* the linear operator L . The next definitions should come as no surprise. In what follows, we will assume that V is a finite-dimensional vector space over \mathbb{F} .

Definition 3.1.2**Eigenvector,
Eigenvalue**

Let $L: V \rightarrow V$ be a linear operator. A non-zero vector $\vec{v} \in V$ such that $L(\vec{v}) = \lambda\vec{v}$ for some $\lambda \in \mathbb{F}$ is called an **eigenvector** of L . The number λ is called an **eigenvalue** of L .

Definition 3.1.3**Eigenspace, $E_{\lambda}(L)$**

Let $L: V \rightarrow V$ be a linear operator, and let $\lambda \in \mathbb{F}$ be an eigenvalue of L . The **eigenspace of L corresponding to λ** is

$$E_{\lambda}(L) = \{\vec{v} \in V : L(\vec{v}) = \lambda\vec{v}\}.$$

These are the same definitions you've seen for square matrices, except now \vec{v} is an element of a general vector space V over \mathbb{F} , and are not necessarily in \mathbb{F}^n .

Since square matrices are essentially linear operators, as we've learned in Section 2.3, we will be able to transport all our results concerning eigenvalues, eigenvectors and the problem of diagonalizability from the setting of matrices to the setting of linear operators. For instance, we have the following.

Proposition 3.1.4

Let $L: V \rightarrow V$ be a linear operator, and let $\lambda \in \mathbb{F}$ be an eigenvalue of L . The eigenspace of L corresponding to λ is a subspace of V .

Exercise 29

Prove Proposition 3.1.4.

The proof of this result is completely identical to the analogous proof about eigenspaces of matrices. We are going to spend the remainder of this chapter reviewing some of the theory of diagonalization of matrices, but rephrased in the language of linear operators. The proofs will be almost word-for-word identical to the matrix proofs, so we will leave them as exercises for the interested and particularly motivated reader!

To help make this new language a bit more familiar, let's look at a few examples.

Example 3.1.5

Let $D: \mathcal{P}_4(\mathbb{R}) \rightarrow \mathcal{P}_4(\mathbb{R})$ be the differentiation map. (Note that although we can view differentiation as a map from $\mathcal{P}_4(\mathbb{R})$ to $\mathcal{P}_3(\mathbb{R})$, we've chosen the codomain $\mathcal{P}_4(\mathbb{R})$ here to ensure that D is a linear *operator*.) Then 0 is an eigenvalue of D since $D(3) = 0 = 0(3)$, and 3 is not the zero vector in $\mathcal{P}_4(\mathbb{R})$. Furthermore, 0 is the only eigenvalue of D . You can see this by noticing that λp and p have the same degree if and only if $\lambda \neq 0$. So, since $D(p)$ and p never have the same degree (unless $p = 0$ of course), then the only way $D(p) = \lambda p$ can be true is if $\lambda = 0$.


Now let's work out what the eigenspace $E_0(D)$ corresponding to 0 looks like. By definition,

$$E_0(D) = \{p \in \mathcal{P}_4(\mathbb{R}) : D(p) = 0\}$$

and therefore $E_0(D) = \{p \in \mathcal{P}_4(\mathbb{R}) : p = k \text{ for some constant } k \in \mathbb{R}\}$ is the subspace of constant polynomials.

Example 3.1.6

Let W be a plane through the origin in \mathbb{R}^3 and let $P: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be the projection onto W . That is, $P(\vec{v})$ is the vector projection of $\vec{v} \in \mathbb{R}^3$ onto W . We will study projections in more detail later in the course, so if they are not familiar to you, you can just use your intuition for now. In particular, convince yourself that P is linear.



What are the eigenvalues of P ? There's an obvious one: $\lambda = 1$. Indeed, if $\vec{v} \in \mathbb{R}^3$ is in W , then $P(\vec{v}) = \vec{v} = 1\vec{v}$. This shows that all the vectors in W are eigenvectors with eigenvalue 1. In fact, the eigenspace corresponding to 1 is precisely W itself. That is, $E_1(P) = W$.

There is another geometrically obvious eigenvalue, namely $\lambda = 0$. In the next exercise you are asked to determine its eigenvectors.

Exercise 30

Let $P: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be as in the previous example.

- Complete the proof that $E_1(P) = W$ by showing that $E_1(P) \subseteq W$.
- Determine $E_0(P)$.

It's all well and good to make definitions like the above, and work out examples where it's easy to find eigenvalues and eigenspaces by inspection, but how can we actually find eigenvalues and eigenspaces in a systematic way? As is becoming a pattern, we will pick a basis \mathcal{B} of V , turn our linear operator into the matrix $[L]_{\mathcal{B}}$, and harness the computational power of matrices!

Once we've picked a basis, we can think of these definitions purely as definitions for matrices. In this case, we can think of a square matrix as a linear map from \mathbb{F}^n to itself, and our vectors are column vectors in \mathbb{F}^n .

3.1.1 Finding Eigenvectors and Eigenvalues

To find eigenvectors and eigenvalues for a linear operator $L: V \rightarrow V$ on a finite-dimensional vector space V over \mathbb{F} , first pick an ordered basis \mathcal{B} for V so you have an $n \times n$ matrix $A = [L]_{\mathcal{B}}$. Now the problem becomes finding eigenvalues and eigenvectors for A . Let's recall the definitions.

Definition 3.1.7

**Eigenvector,
Eigenvalue,
Eigenspace of a
Matrix, $E_{\lambda}(A)$**

Let $A \in M_{n \times n}(\mathbb{F})$ be an $n \times n$ matrix. A non-zero vector $\vec{v} \in \mathbb{F}^n$ such that $A\vec{v} = \lambda\vec{v}$ for some $\lambda \in \mathbb{F}$ is called an **eigenvector** of A . The number λ is called an **eigenvalue** of A .

The **eigenspace of A corresponding to λ** is

$$E_{\lambda}(A) = \{\vec{v} \in \mathbb{F}^n : A\vec{v} = \lambda\vec{v}\}.$$

Exercise 31

Prove that $\vec{v} \in V$ is an eigenvector of L with eigenvalue $\lambda \in \mathbb{F}$ if and only if $[\vec{v}]_{\mathcal{B}} \in \mathbb{F}^n$ is an eigenvector of $A = [L]_{\mathcal{B}}$ with eigenvalue λ .

To find an eigenvector for A , we're looking for a vector $\vec{v} \neq \vec{0}$ such that

$$A\vec{v} = \lambda\vec{v}$$

for some $\lambda \in \mathbb{F}$. If we rearrange this equation we get

$$A\vec{v} - \lambda\vec{v} = \vec{0}.$$

It would be tempting now to factor out the \vec{v} , which we will do, but we cannot as written. If we did, we would be left with a term $A - \lambda$, which makes no sense since A is a square matrix and λ is an element of \mathbb{F} . To get around this, we observe that $\lambda\vec{v} = \lambda I\vec{v}$ where I is the identity matrix of the appropriate size. Now our equation takes the form

$$(A - \lambda I)\vec{v} = \vec{0}.$$

If the matrix $A - \lambda I$ were invertible, then we could multiply both sides on the left by the inverse and get $\vec{v} = \vec{0}$. Since we're looking for non-zero vectors \vec{v} , this means we are looking for values of λ that make the matrix $A - \lambda I$ not invertible. Equivalently, we want values of λ such that $\det(A - \lambda I) = 0$. Furthermore, once we've found such a λ , a corresponding eigenvector is any non-zero vector such that $(A - \lambda I)\vec{v} = \vec{0}$, which must exist because $A - \lambda I$ is non-invertible. These are precisely the non-zero vectors in $\text{Null}(A - \lambda I)$. Let's summarize.

Proposition 3.1.8

Let $L: V \rightarrow V$ be a linear operator on a finite-dimensional vector space V over \mathbb{F} , and let $\lambda \in \mathbb{F}$ be an eigenvalue of L (provided that it exists). If \mathcal{B} is an ordered basis for V and if $A = [L]_{\mathcal{B}}$, then the eigenspace of A corresponding to λ is $\text{Null}(A - \lambda I)$:

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

Once we determine the eigenspace $E_{\lambda}(A)$, we can obtain $E_{\lambda}(L)$ by converting back from \mathcal{B} -coordinate vectors in \mathbb{F}^n to vectors in V .

Example 3.1.9

Let $D: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathcal{P}_2(\mathbb{R})$ be the differentiation map. If $\mathcal{S} = \{1, x, x^2\}$ is the standard basis for $\mathcal{P}_2(\mathbb{R})$ then as we've noted before

$$[D]_{\mathcal{S}} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix}.$$

Call this matrix A . The argument in Example 3.1.5 shows that 0 is an eigenvalue of D . To find the corresponding eigenspace, we must compute

$$E_0(A) = \text{Null}(A - 0I) = \text{Null} \left(\begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix} \right).$$

To find the nullspace of a matrix, we simply apply the Gauss–Jordan algorithm and row reduce. Skipping the easy details, we find that

$$E_0(A) = \text{Null} \left(\begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix} \right) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right\}.$$

Translating back to polynomials, the vector $[1 \ 0 \ 0]^T$ is the polynomial $1 + 0x + 0x^2 = 1$. So we conclude that $E_0(D)$ is $\text{Span}\{1\}$, the subspace of constant polynomials, perhaps as expected (and as we saw in Example 3.1.5 for the differentiation operator on \mathcal{P}_4).

Exercise 32

Determine $E_0(D)$ for the differentiation operator $D: \mathcal{P}_n(\mathbb{R}) \rightarrow \mathcal{P}_n(\mathbb{R})$.

Thus, we have a computational strategy for finding eigenvectors. What about eigenvalues? Well, they are the numbers $\lambda \in \mathbb{F}$ that satisfy the equation $\det(A - \lambda I) = 0$. Here is another (hopefully familiar) definition.

Definition 3.1.10**Characteristic
Polynomial of a
Matrix**

Let $A \in M_{n \times n}(\mathbb{F})$. The **characteristic polynomial** of A is the polynomial in λ given by $C_A(\lambda) = \det(A - \lambda I)$.

As you may recall, $C_A(\lambda)$ really is a polynomial (see Theorem 3.1.14). Our preceding discussion proves the next proposition.

Proposition 3.1.11

Let $A \in M_{n \times n}(\mathbb{F})$. The eigenvalues of A are the values of $\lambda \in \mathbb{F}$ that are solutions to the equation $\det(A - \lambda I) = 0$. That is, they are the roots of the characteristic polynomial of A that lie in \mathbb{F} .

Example 3.1.12

Going back to Example 3.1.9, with the differentiation operator $D: \mathcal{P}_2(\mathbb{R}) \rightarrow \mathcal{P}_2(\mathbb{R})$ and its standard matrix

$$A = [D]_{\mathcal{S}} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix},$$

we can compute the characteristic polynomial of A to be

$$C_A(\lambda) = \det(A - \lambda I) = \det \left(\begin{bmatrix} -\lambda & 1 & 0 \\ 0 & -\lambda & 2 \\ 0 & 0 & -\lambda \end{bmatrix} \right) = -\lambda^3.$$

The only root is $\lambda = 0$, meaning: the only eigenvalue of D is $\lambda = 0$ as we'd already seen.

Example 3.1.13

Let $L: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the linear operator defined by $L(\vec{x}) = A\vec{x}$, where $A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$. Then $A = [L]_{\mathcal{S}}$ is the standard matrix of L . The characteristic polynomial of A is

$$C_A(\lambda) = \det(A - \lambda I) = \det \left(\begin{bmatrix} -\lambda & 1 \\ -1 & -\lambda \end{bmatrix} \right) = \lambda^2 + 1.$$

This has no roots in \mathbb{R} , and so the operator L has no eigenvalues.

On the other hand, consider the operator $T: \mathbb{C}^2 \rightarrow \mathbb{C}^2$ defined by $T(\vec{x}) = A\vec{x}$, where A is the same matrix above. Then $A = [T]_{\mathcal{S}}$ is again the standard matrix of T , but this time its characteristic polynomial has two roots in \mathbb{C} , namely $\lambda = \pm i$. So T has two eigenvalues $\lambda = \pm i$.

REMARK

The previous example shows that the field of definition plays an important role here. This can get confusing because we can always view a real matrix as being a complex matrix that has real entries.

To avoid any potential ambiguity, we will adopt the following convention in the context of eigenvalues: When we write " $A \in M_{n \times n}(\mathbb{R})$ ", we are specifically viewing A as representing the real operator $L: \mathbb{R}^n \rightarrow \mathbb{R}^n$ defined by $L(\vec{x}) = A\vec{x}$. Thus, its eigenvalues will be required to belong to \mathbb{R} . If we wish to view A as belonging to $M_{n \times n}(\mathbb{C})$, and therefore possibly having non-real eigenvalues, then we will explicitly state this.

To close this section, let's recall a few facts about the characteristic polynomial. The important fact here is that this really is a polynomial.

Theorem 3.1.14 If A is an $n \times n$ matrix with entries in \mathbb{F} , then the characteristic polynomial of A is a polynomial of degree n with coefficients in \mathbb{F} .

This proof is a bit tricky, requiring some finesse with the cofactor expansion definition of the determinant. We shall not write it down here. If you are motivated to try it yourself, first try proving the statement when $n = 2$ and $n = 3$.

Example 3.1.15 In Example 3.1.12, we saw that the characteristic polynomial of the matrix $A = [D]_{\mathcal{S}}$ in $M_{3 \times 3}(\mathbb{R})$ is the degree 3 polynomial $C_A(\lambda) = -\lambda^3$.

In Example 3.1.13, we saw that the characteristic polynomial of the matrix $A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$ in $M_{2 \times 2}(\mathbb{R})$ is the degree 2 polynomial $C_A(\lambda) = \lambda^2 + 1$. Note if we view A as belonging to $M_{2 \times 2}(\mathbb{C})$ then this still conforms with the above theorem, since real coefficients can be viewed as being complex coefficients.

Using the fundamental theorem of algebra, i.e., the fact that every degree n polynomial with coefficients in \mathbb{C} has n (possibly repeated) roots in \mathbb{C} , we obtain the next result.

Corollary 3.1.16 Let $A \in M_{n \times n}(\mathbb{C})$. Then A has n (possibly repeated) eigenvalues in \mathbb{C} .

REMARK

Pay careful attention to the field of definition here. As we saw in Example 3.1.13, if $A \in M_{n \times n}(\mathbb{R})$, then A need not have any eigenvalues in \mathbb{R} at all! But viewed as a matrix in $M_{n \times n}(\mathbb{C})$, it will always have eigenvalues in \mathbb{C} .

Example 3.1.17 If our characteristic polynomial is $C_A(\lambda) = -\lambda^3 = -(\lambda - 0)^3$, then the root 0 is *repeated with multiplicity 3*.

If instead we had $C_A(\lambda) = (\lambda - 1)(\lambda - (1 + i))^2(\lambda + 5)^4$, then we would say that the root 1 is not repeated (or *repeated with multiplicity 1*), while the roots $1 + i$ and -5 are *repeated with multiplicities 2 and 4*, respectively.

By thinking carefully about polynomials and how the roots relate to the coefficients, you can prove the following corollary. (We will be able to give a very easy proof in Chapter 5. See Example 5.3.3.)

Corollary 3.1.18 Let $A \in M_{n \times n}(\mathbb{C})$. Then

- The determinant of A is the product of eigenvalues of A , where each eigenvalue is repeated according to its multiplicity.
- The trace of A is the sum of eigenvalues of A , where each eigenvalue is repeated according to its multiplicity.

REMARK

It's worth pointing out that the previous result applies to a real matrix A if we view it as lying in $M_{n \times n}(\mathbb{C})$. As we've seen, such a matrix can have non-real eigenvalues in \mathbb{C} . However, that $\det(A)$ and $\operatorname{tr}(A)$ will be real if all of the entries of A are real. So we arrive at the interesting observation that the sum and product of all of the complex eigenvalues of A must both be real numbers! This is illustrated in Examples 3.1.19 and 3.1.21 below.

Example 3.1.19 If $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \in M_{2 \times 2}(\mathbb{C})$, then we saw in Example 3.1.13 that the characteristic polynomial of A is $C_A(\lambda) = \lambda^2 + 1 = (\lambda + i)(\lambda - i)$. So the eigenvalues of A are i and $-i$ (non-repeated), and indeed we have

$$\det(A) = i(-i) \quad \text{and} \quad \operatorname{tr}(A) = i + (-i)$$

since $\det(A) = 1$ and $\operatorname{tr}(A) = 0$. Notice that these two quantities are real, even though the eigenvalues appearing in the above equations are not!

Example 3.1.20

For

$$A = [D]_{\mathcal{S}} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix},$$

where $C_A(\lambda) = -\lambda^3$, we can indeed verify that

$$\det(A) = 0 \cdot 0 \cdot 0$$

and

$$\operatorname{tr}(A) = 0 + 0 + 0.$$

This example is particularly trivial because the matrix A is upper-triangular, which if you recall means the diagonal entries of A are precisely the eigenvalues of A (repeated according to multiplicity).

It is perhaps more interesting to consider the matrix of D with respect to another basis. For instance, if we take $\mathcal{B} = \{1 + x^2, x, 1 + x - x^2\}$, then we'd find that

$$B = [D]_{\mathcal{B}} = \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ 2 & -\frac{1}{2} & -\frac{5}{2} \\ 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}.$$

We'll leave it to you to check that $C_B(\lambda) = -\lambda^3$, so that the only eigenvalue of B is 0 repeated with multiplicity 3, and that

$$\det(B) = 0 = 0 \cdot 0 \cdot 0 \quad \text{and} \quad \operatorname{tr}(B) = 0 + \left(-\frac{1}{2}\right) + \frac{1}{2} = 0 + 0 + 0.$$

Notice that the matrices A and B have the same characteristic polynomial (and therefore the same determinant and trace). This is not a coincidence! We will explore this in the next section.

Exercise 33 Verify the claims about the matrix B in the preceding example.

Example 3.1.21 Consider the matrix

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

As a matrix in $M_{2 \times 2}(\mathbb{R})$ it has only one eigenvalue in \mathbb{R} , namely $\lambda_1 = 1$. However, as a matrix in $M_{2 \times 2}(\mathbb{C})$, it has three eigenvalues in \mathbb{C} , namely $\lambda_1 = 1$, $\lambda_2 = i$ and $\lambda_3 = -i$. The determinant of this matrix is $\det(A) = 1$ and we can verify that it is equal to the product of the eigenvalues in \mathbb{C} :

$$\det(A) = 1 = 1 \cdot i \cdot (-i).$$

We can also verify that the trace of A is equal to the sum of the eigenvalues in \mathbb{C} :

$$\operatorname{tr}(A) = 1 + 0 + 0 = 1 + i + (-i).$$

3.2 Diagonalization

The fundamental definition in this section is the following.

Definition 3.2.1

**Diagonalizable
Operator,
Diagonalizes**

Let V be a finite-dimensional vector space over \mathbb{F} . A linear operator $L: V \rightarrow V$ is **diagonalizable** if there exists an ordered basis \mathcal{D} for V such that $[L]_{\mathcal{D}}$ is a diagonal matrix. We say that the basis \mathcal{D} **diagonalizes** L .

The point being: an operator L is diagonalizable if there is some coordinate system in which the action of L is easy to interpret.

Example 3.2.2

As noted in the opening to this chapter, the linear operator $L: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ given by

$$L \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \right) = \begin{bmatrix} x \\ y \\ z \end{bmatrix} - \frac{2(x+y+z)}{3} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}.$$

is diagonalizable. Indeed, if we take the ordered basis

$$\mathcal{D} = \left\{ \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix} \right\}$$

for \mathbb{R}^3 , we find that

$$[L]_{\mathcal{D}} = \begin{bmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

We can use this to interpret L as a reflection in the coordinate system provided by the basis \mathcal{D} for \mathbb{R}^3 .

How do we determine if a linear operator $L: V \rightarrow V$ is diagonalizable? And if it is, how do we diagonalize it? That is, how do we determine a basis \mathcal{D} for V such that $[L]_{\mathcal{D}}$ is diagonal?

To answer these questions we will begin by considering $[L]_{\mathcal{B}}$ for an arbitrary basis \mathcal{B} . We would like to understand how these matrices, for the various bases of V , are related.

Proposition 3.2.3

Let $L: V \rightarrow V$ be a linear operator, and let \mathcal{B} and \mathcal{C} be ordered bases for V . Then

$$[L]_{\mathcal{B}} = ({}_C\mathcal{I}_{\mathcal{B}})^{-1} [L]_{\mathcal{C}} {}_C\mathcal{I}_{\mathcal{B}}.$$

Proof: Since $({}_C\mathcal{I}_{\mathcal{B}})^{-1} = {}_{\mathcal{B}}\mathcal{I}_{\mathcal{C}}$ (by Proposition 2.4.4), and since $[L]_{\mathcal{C}} = {}_C[L]_{\mathcal{C}}$ and $[L]_{\mathcal{B}} = {}_{\mathcal{B}}[L]_{\mathcal{B}}$, the result follows from Proposition 2.4.6. To spell it out,

$$\begin{aligned} ({}_C\mathcal{I}_{\mathcal{B}})^{-1} [L]_{\mathcal{C}} {}_C\mathcal{I}_{\mathcal{B}} &= ({}_C\mathcal{I}_{\mathcal{B}})^{-1} {}_C[L]_{\mathcal{C}} {}_C\mathcal{I}_{\mathcal{B}} \\ &= ({}_C\mathcal{I}_{\mathcal{B}})^{-1} {}_C[L]_{\mathcal{B}} \\ &= {}_{\mathcal{B}}\mathcal{I}_{\mathcal{C}} {}_C[L]_{\mathcal{B}} \\ &= {}_{\mathcal{B}}[L]_{\mathcal{B}} \\ &= [L]_{\mathcal{B}}. \end{aligned}$$

□

This proposition motivates our next definition.

Definition 3.2.4

Similar

If B and C are $n \times n$ matrices such that $B = P^{-1}CP$ for some invertible matrix P in $M_{n \times n}(\mathbb{F})$, then we say B is **similar to** C over \mathbb{F} .

Proposition 3.2.3 tells us that any two matrix representations $[L]_{\mathcal{B}}$ and $[L]_{\mathcal{C}}$ of a fixed operator are similar, with P being the change of basis matrix ${}_C\mathcal{I}_{\mathcal{B}}$. Conversely, two similar $n \times n$ matrices can be always be viewed as representing the same operator $L: \mathbb{F}^n \rightarrow \mathbb{F}^n$. Indeed, if $B = P^{-1}CP$, then the columns of P form a basis for \mathbb{F}^n , since P is invertible. If we call this basis \mathcal{C} , and if we let \mathcal{B} be the standard basis, then we have $P = {}_{\mathcal{B}}\mathcal{I}_{\mathcal{C}}$. Now let $L: \mathbb{F}^n \rightarrow \mathbb{F}^n$ be the operator defined by $L(\vec{x}) = Bx$. Then $[L]_{\mathcal{B}} = B$ and so, by changing coordinates to \mathcal{C} , $[L]_{\mathcal{C}} = P^{-1}[L]_{\mathcal{B}}P$ since $P = {}_{\mathcal{B}}\mathcal{I}_{\mathcal{C}}$ and $P^{-1} = {}_C\mathcal{I}_{\mathcal{B}}$.

That was probably confusing. The takeaway is this: two $n \times n$ matrices are similar over \mathbb{F} if and only if they represent the same operator on \mathbb{F}^n . This justifies the usage of the adjective “similar” for such matrices. The next set of exercises and the proposition that follows give further justification.

Exercise 34

Let $A, B, C \in M_{n \times n}(\mathbb{F})$.

- Show that A is similar to A .
- Show that if A is similar to B then B is similar to A .
- Show that if A is similar to B and if B is similar to C then A is similar to C .

The next result shows that similar matrices share a lot of common features, which is to be expected—since they represent the same operator!

Proposition 3.2.5

Let $A, B \in M_{n \times n}(\mathbb{F})$. If A is similar to B over \mathbb{F} , then A and B have the same

- (a) characteristic polynomial,
- (b) eigenvalues,
- (c) determinant,
- (d) trace,
- (e) rank, and
- (f) nullity.

Proof: Let's begin with part (a). If the $n \times n$ matrices A and B are similar, then there exists an invertible matrix $P \in M_{n \times n}(\mathbb{F})$ such that $A = P^{-1}BP$, and therefore

$$\begin{aligned}
 C_A(\lambda) &= \det(A - \lambda I) \\
 &= \det(P^{-1}BP - P^{-1}(\lambda I)P) \\
 &= \det(P^{-1}(B - \lambda I)P) \\
 &= \det(P^{-1}) \det(B - \lambda I) \det(P) && \text{(since det is multiplicative)} \\
 &= \det(P)^{-1} \det(B - \lambda I) \det(P) \\
 &= \det(B - \lambda I) \\
 &= C_B(\lambda).
 \end{aligned}$$

This completes the proof of part (a). Parts (b), (c) and (d) follow immediately from this. (How?)

Finally, for parts (e) and (f), we can assume (by the discussion following Definition 3.2.4) that A and B represent the same operator $L: \mathbb{F}^n \rightarrow \mathbb{F}^n$. That is, we can assume there are ordered bases \mathcal{A} and \mathcal{B} for \mathbb{F}^n such that $[L]_{\mathcal{A}} = A$ and $[L]_{\mathcal{B}} = B$. Then, by Proposition 2.5.20, $\text{rank}(A) = \text{rank}(L)$ and $\text{rank}(B) = \text{rank}(L)$. Thus, $\text{rank}(A) = \text{rank}(B)$. Similarly, $\text{nullity}(A) = \text{nullity}(L) = \text{nullity}(B)$. This completes the proof. \square

Now let's return to the problem of diagonalizing a given operator $L: V \rightarrow V$, where we ask if it's possible to find a matrix representation $[L]_{\mathcal{D}}$ of L that is diagonal. By what we've just learned, this is the same as asking if *any* matrix representation of A is *similar* to a diagonal matrix.

Example 3.2.6

If $L: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is in Example 3.2.2, then with respect to the standard basis \mathcal{S} of \mathbb{R}^3 we have

$$[L]_{\mathcal{S}} = \begin{bmatrix} \frac{1}{3} & -\frac{2}{3} & -\frac{2}{3} \\ -\frac{2}{3} & \frac{1}{3} & -\frac{2}{3} \\ -\frac{2}{3} & -\frac{2}{3} & \frac{1}{3} \end{bmatrix}.$$

If \mathcal{D} is as in Example 3.2.2 then the change of basis matrix from \mathcal{D} to \mathcal{S} is given by

$${}_S\mathcal{I}_{\mathcal{D}} = \begin{bmatrix} 1 & -1 & -1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}$$

and you can check that

$$({}_S\mathcal{I}_{\mathcal{D}})^{-1} [L]_{\mathcal{S}} {}_S\mathcal{I}_{\mathcal{D}} = \begin{bmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

So $[L]_{\mathcal{S}}$ is similar to a diagonal matrix. The above matrix is, in fact, equal to $[L]_{\mathcal{D}}$ by our change-of-coordinates formula (Proposition 2.4.6).

With this new perspective, let's re-frame our diagonalization problem in terms of matrices.

Definition 3.2.7

**Diagonalizable
Matrix,
Diagonalizes**

A matrix $A \in M_{n \times n}(\mathbb{F})$ is **diagonalizable (over \mathbb{F})** if there exists an invertible matrix $P \in M_{n \times n}(\mathbb{F})$ such that $P^{-1}AP = D$ where D is a diagonal matrix. We say that the matrix P **diagonalizes A** .

REMARK

You will recall that the field \mathbb{F} plays an important role here. For instance, a matrix A in $M_{n \times n}(\mathbb{R})$ might not be diagonalizable if we insist on having $P \in M_{n \times n}(\mathbb{R})$. However, we might be able to find a suitable P in $M_{n \times n}(\mathbb{C})$. The standard example of this is a 2×2 rotation matrix, which, if the angle of rotation isn't a multiple of π , is not diagonalizable over \mathbb{R} but is diagonalizable over \mathbb{C} . (Do you remember why?)

So the question now becomes: When is an $n \times n$ matrix A diagonalizable? This is a problem that you have studied in a previous course. We will quickly review—without proof—the solution to this problem.

If we think about how the matrix of a linear map works, then we wish to find a basis of eigenvectors.

Theorem 3.2.8

An $n \times n$ matrix $A \in M_{n \times n}(\mathbb{F})$ is diagonalizable over \mathbb{F} if and only if there exists a basis $\mathcal{D} = \{\vec{v}_1, \dots, \vec{v}_n\}$ for \mathbb{F}^n such that each \vec{v}_i is an eigenvector for A .

If such a basis \mathcal{D} exists, and if we let $P = [\vec{v}_1 \ \cdots \ \vec{v}_n]$ be the matrix whose columns are the vectors in \mathcal{D} , then

$$P^{-1}AP = D = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_n \end{bmatrix}$$

where λ_i is the eigenvalue corresponding to the eigenvector \vec{v}_i .

So, if we are to diagonalize a matrix, we need to find a basis for \mathbb{F}^n consisting entirely of eigenvectors. Let's take a look at some examples.

Example 3.2.9

Let $L: \mathcal{P}_2(\mathbb{F}) \rightarrow \mathcal{P}_2(\mathbb{F})$ be the differentiation operator defined by $L(p(x)) = p'(x)$. If $\mathcal{S} = \{1, x, x^2\}$ is the standard basis of $\mathcal{P}_2(\mathbb{F})$ then we have seen in Examples 3.1.9 and 3.1.12 that

$$A = [L]_{\mathcal{S}} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix}$$

and that $C_A(\lambda) = -\lambda^3$. So the only eigenvalue of A is $\lambda = 0$ and we saw that the corresponding eigenspace is

$$E_0(A) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right\}.$$

So there can be no basis of \mathbb{F}^3 consisting entirely of eigenvectors of A , since there is *at most* one linearly independent eigenvector of A in \mathbb{F}^3 . Thus A (and hence L) is **not** diagonalizable.

Example 3.2.10

Let $L: \mathcal{P}_1(\mathbb{F}) \rightarrow \mathcal{P}_1(\mathbb{F})$ be defined by $L(a + bx) = (a + 2b) + (2a + b)x$. Then, using the standard basis $\mathcal{S} = \{1, x\}$, we have that

$$A = [L]_{\mathcal{S}} = \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix}.$$

The characteristic polynomial of A is $C_A(\lambda) = (1 - \lambda)^2 - 4 = (\lambda - 3)(\lambda + 1)$ and so the eigenvalues of A are $\lambda_1 = 3$ and $\lambda_2 = -1$.

Using our usual method for finding nullspaces, we obtain

$$E_{\lambda_1}(A) = \text{Null}(A - 3I) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$$

and

$$E_{\lambda_2}(A) = \text{Null}(A - (-I)) = \text{Span} \left\{ \begin{bmatrix} 1 \\ -1 \end{bmatrix} \right\}.$$

Since the two eigenvectors $\vec{v}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\vec{v}_2 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$ are linearly independent, we can take $\mathcal{D} = \{\vec{v}_1, \vec{v}_2\}$ as a basis for \mathbb{F}^2 . If we let $P = [\vec{v}_1 \ \vec{v}_2] = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$ then

$$P^{-1}AP = \begin{bmatrix} 3 & 0 \\ 0 & -1 \end{bmatrix}.$$

So we have diagonalized the matrix A . We can convert back from coordinate vectors to polynomials to obtain a basis for $\mathcal{P}_2(\mathbb{F})$. Recall that we are using the standard basis $\mathcal{S} = \{1, x\}$, so $\vec{v}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\vec{v}_2 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$ correspond to the polynomials $1 + x$ and $1 - x$, respectively. Our work above amounts to the fact that if we let $\mathcal{D} = \{1 + x, 1 - x\}$ then $[L]_{\mathcal{D}} = \begin{bmatrix} 3 & 0 \\ 0 & -1 \end{bmatrix}$.

In the previous example we could check by inspection that the eigenvectors \vec{v}_1 and \vec{v}_2 corresponding to the two distinct eigenvalues λ_1 and λ_2 are linearly independent. In fact, this is a special case of:

Proposition 3.2.11

Suppose $\lambda_1, \dots, \lambda_k \in \mathbb{F}$ are *distinct* eigenvalues of a square matrix $A \in M_{n \times n}(\mathbb{F})$ with corresponding eigenvectors $\vec{v}_1, \dots, \vec{v}_k$. Then $\{\vec{v}_1, \dots, \vec{v}_k\}$ is linearly independent.

Combining this with Theorem 3.2.8, we obtain the following useful proposition.

Proposition 3.2.12

If $A \in M_{n \times n}(\mathbb{F})$ has n distinct eigenvalues in \mathbb{F} , then A is diagonalizable over \mathbb{F} .

Warning: The converse of Proposition 3.2.12 is false. There are plenty of diagonalizable matrices that do not have distinct eigenvalues. For instance, the $n \times n$ identity matrix when $n > 1$.

Example 3.2.13

Let $L: M_{1 \times 2}(\mathbb{F}) \rightarrow M_{1 \times 2}(\mathbb{F})$ be defined by $L\left(\begin{bmatrix} x & y \end{bmatrix}\right) = \begin{bmatrix} y & -2x - 3y \end{bmatrix}$. Then with respect to the standard basis \mathcal{S} of $M_{1 \times 2}(\mathbb{F})$ we have

$$A = [L]_{\mathcal{S}} = \begin{bmatrix} 0 & 1 \\ -2 & -3 \end{bmatrix}.$$

The characteristic polynomial of A is $C_A(\lambda) = \lambda^2 + 3\lambda + 2 = (\lambda + 1)(\lambda + 2)$. Thus, A has distinct eigenvalues, and therefore must be diagonalizable.

Going further, we can show that $\left\{\begin{bmatrix} -1 \\ 1 \end{bmatrix}\right\}$ and $\left\{\begin{bmatrix} -1 \\ 2 \end{bmatrix}\right\}$ are bases for the eigenspaces $E_{-1}(A)$ and $E_{-2}(A)$, respectively, and therefore $\left\{\begin{bmatrix} -1 \\ 1 \end{bmatrix}, \begin{bmatrix} -1 \\ 2 \end{bmatrix}\right\}$ is a basis for \mathbb{F}^2 consisting entirely of eigenvectors of A . So, by Theorem 3.2.8, we must have $P^{-1}AP = D$ where

$$P = \begin{bmatrix} -1 & -1 \\ 1 & 2 \end{bmatrix} \quad \text{and} \quad D = \begin{bmatrix} -1 & 0 \\ 0 & -2 \end{bmatrix}.$$

Let's check this. We have

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

and

$$PD = \begin{bmatrix} -1 & -1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} -1 & 0 \\ 0 & -2 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ -1 & -4 \end{bmatrix}$$

so $P^{-1}AP = D$.

Translating back to $M_{1 \times 2}(\mathbb{F})$, we can conclude that the basis $\mathcal{D} = \left\{\begin{bmatrix} -1 & 1 \end{bmatrix}, \begin{bmatrix} -1 & 2 \end{bmatrix}\right\}$ diagonalizes L .

If the matrix A has repeated eigenvalues then, as you might recall, their multiplicities play an important role in the problem of diagonalizability.

Definition 3.2.14

Algebraic
Multiplicity,
Geometric
Multiplicity

Let $A \in M_{n \times n}(\mathbb{F})$ and let $\lambda \in \mathbb{F}$ be an eigenvalue of A . The **algebraic multiplicity** of λ is the multiplicity of λ as a root of the characteristic polynomial of A . The **geometric multiplicity** of λ is defined to be the dimension of the eigenspace $E_\lambda(A) = \text{Null}(A - \lambda I)$.

The geometric multiplicity of an eigenvalue is really what we are interested in. It tells us how many linearly independent eigenvectors an eigenvalue can have. After all, we want to find a basis consisting entirely of eigenvectors, so we want to be able to find sufficiently many linearly independent eigenvectors. The obvious thing to do is to find bases for each of the eigenspaces and then combine them together (i.e. take their union) to form a set \mathcal{D} . Two potential issues arise:

1. Is the resulting set \mathcal{D} linearly independent?
2. Are there enough vectors in \mathcal{D} to span all of \mathbb{F}^n ?

Amazingly, the answer to question 1 is always yes! (Compare to Proposition 3.2.11.)

Proposition 3.2.15

Suppose $\lambda_1, \dots, \lambda_k \in \mathbb{F}$ are distinct eigenvalues of a matrix $A \in M_{n \times n}(\mathbb{F})$, and let $\{\vec{v}_{i,1}, \vec{v}_{i,2}, \dots, \vec{v}_{i,m_i}\}$ be a basis for the eigenspace corresponding to λ_i (so the dimension of the eigenspace corresponding to λ_i is m_i). Then

$$\{\vec{v}_{1,1}, \vec{v}_{1,2}, \dots, \vec{v}_{1,m_1}, \vec{v}_{2,1}, \dots, \vec{v}_{2,m_2}, \dots, \vec{v}_{k,1}, \dots, \vec{v}_{k,m_k}\}$$

is a linearly independent subset of \mathbb{F}^n .

The other key fact here is that the algebraic multiplicity acts as an upper bound on the geometric multiplicity.

Proposition 3.2.16

Let $A \in M_{n \times n}(\mathbb{F})$ and let $\lambda \in \mathbb{F}$ be an eigenvalue of A . Then

$$1 \leq \text{geometric multiplicity of } \lambda \leq \text{algebraic multiplicity of } \lambda \leq n.$$

Now, if $\mathbb{F} = \mathbb{R}$ and the sum of algebraic multiplicities is strictly less than n , then the characteristic polynomial $C_A(\lambda)$ does not have n roots in \mathbb{F} , and so we can immediately conclude that the matrix is not diagonalizable. Otherwise, when the sum of the algebraic multiplicities is equal to $\deg C_A(\lambda) = n = \dim \mathbb{F}^n$, we need to check if geometric multiplicity is equal to algebraic multiplicity for each eigenvalue. If the geometric multiplicity is ever strictly less than the algebraic multiplicity for any eigenvalue, we immediately know that we cannot find enough linearly independent eigenvectors to diagonalize the matrix. Conversely, if the sum of algebraic multiplicities is equal to n , and if the geometric multiplicity is equal to the algebraic multiplicity for each eigenvalue, then we will be able to find enough linearly independent eigenvectors.

Theorem 3.2.17 (Diagonalizability Test)

Let $A \in M_{n \times n}(\mathbb{F})$. Suppose that the characteristic polynomial of A factors over \mathbb{F} as

$$C_A(\lambda) = (\lambda - \lambda_1)^{a_{\lambda_1}} \cdots (\lambda - \lambda_k)^{a_{\lambda_k}} h(\lambda),$$

where $\lambda_1, \dots, \lambda_k$ are all of the distinct eigenvalues of A in \mathbb{F} with corresponding algebraic multiplicities $a_{\lambda_1}, \dots, a_{\lambda_k}$ and $h(\lambda)$ is a polynomial in λ that has no roots in \mathbb{F} . Then A is diagonalizable over \mathbb{F} if and only if $h(\lambda)$ is a constant polynomial and for all $i = 1, \dots, k$,

$$\text{algebraic multiplicity of } \lambda_i = \text{geometric multiplicity of } \lambda_i.$$

Exercise 35

Supply proofs for as many of the unproved propositions and theorems above as you can. For those you don't know how to prove, look up their proofs in the notes from your previous linear algebra course.

The diagonalizability test completely answers the question of whether or not an $n \times n$ matrix is diagonalizable over \mathbb{F} . Let us summarize all our previous results.

ALGORITHM (Diagonalization of an Operator)

To diagonalize a linear operator $L: V \rightarrow V$:

1. Pick any basis \mathcal{B} for V and determine the matrix $A = [L]_{\mathcal{B}}$.
2. Compute and factor the characteristic polynomial $C_A(\lambda)$ to find all the distinct eigenvalues $\lambda_1, \dots, \lambda_k \in \mathbb{F}$ of A . Let a_i denote the algebraic multiplicity of λ_i . If $a_{\lambda_1} + \cdots + a_{\lambda_k} \neq n$, conclude that A is not diagonalizable over \mathbb{F} .
3. Otherwise, if $a_{\lambda_1} + \cdots + a_{\lambda_k} = n$, determine a basis \mathcal{B}_i for the eigenspace $E_{\lambda_i}(A)$, for each $i = 1, \dots, k$. Let $g_{\lambda_i} = \dim E_{\lambda_i}$ denote the geometric multiplicity of λ_i .
4. A (hence L) is diagonalizable if and only if $a_{\lambda_i} = g_{\lambda_i}$ for all $i = 1, \dots, k$.
5. If A is diagonalizable, then $\mathcal{D} = \mathcal{B}_1 \cup \cdots \cup \mathcal{B}_k$ is a basis for \mathbb{F}^n (where $n = \dim V$) consisting of eigenvectors of A . If P is the matrix whose columns are the vectors in \mathcal{D} , then $D = P^{-1}AP$ is a diagonal matrix. The diagonal entries of D are λ_1 (listed a_1 times), \dots , λ_k (listed a_k times). The order of eigenvalues matches the order in which their corresponding eigenvectors occur as columns in P .
6. To determine a basis for V that diagonalizes L , take each of the vectors in \mathcal{D} , view it as a coordinate vector in \mathbb{F}^n with respect to the basis \mathcal{B} from Step 1, and thereby convert it into vector V . The set of all these vectors is then the desired basis for V .

Example 3.2.18

Let $L: M_{2 \times 2}(\mathbb{F}) \rightarrow M_{2 \times 2}(\mathbb{F})$ be defined by

$$L \left(\begin{bmatrix} a & b \\ c & d \end{bmatrix} \right) = \begin{bmatrix} 5a + 2b + d & -2a + b - d \\ 4a + 4b + 3c + 2d & 16a - 8c - 5d \end{bmatrix}.$$

Using the standard basis \mathcal{S} of $M_{2 \times 2}(\mathbb{F})$, we have

$$A = [L]_{\mathcal{S}} = \begin{bmatrix} 5 & 2 & 0 & 1 \\ -2 & 1 & 0 & -1 \\ 4 & 4 & 3 & 2 \\ 16 & 0 & -8 & -5 \end{bmatrix}.$$

Then $C_A(\lambda) = (\lambda - 3)^3(\lambda + 5)$. So the eigenvalues of A are $\lambda_1 = 3$ and $\lambda_2 = -5$ with algebraic multiplicities $a_{\lambda_1} = 3$ and $a_{\lambda_2} = 1$, respectively.

We must now determine the geometric multiplicities. The geometric multiplicity g_{λ_2} is easy: since $1 \leq g_{\lambda_2} \leq a_{\lambda_2} = 1$, we must have that $g_{\lambda_2} = 1$.

For g_{λ_1} , we need to find the dimension of the eigenspace corresponding to $\lambda_1 = 3$. We have

$$E_{\lambda_1}(A) = \text{Null}(A - 3I) = \text{Null} \left(\begin{bmatrix} 2 & 2 & 0 & 1 \\ -2 & -2 & 0 & -1 \\ 4 & 4 & 0 & 2 \\ 16 & 0 & -8 & -8 \end{bmatrix} \right).$$

Row reduction leads to

$$\begin{bmatrix} 2 & 2 & 0 & 1 \\ -2 & -2 & 0 & -1 \\ 4 & 4 & 0 & 2 \\ 16 & 0 & -8 & -8 \end{bmatrix} \rightarrow \begin{bmatrix} 2 & 2 & 0 & 1 \\ 0 & 16 & 8 & 16 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

The nullity of these matrices is 2. Hence, $g_{\lambda_1} = \dim E_{\lambda_1}(A) = \text{nullity}(A - 3I) = 2$. However, $a_{\lambda_1} = 3$. So, since $g_{\lambda_1} \neq a_{\lambda_1}$, we conclude that A (and hence L) is not diagonalizable.

Example 3.2.19

Let $L: \mathcal{P}_3(\mathbb{F}) \rightarrow \mathcal{P}_3(\mathbb{F})$ be defined by

$$L(a + bx + cx^2 + dx^3) = (-2b + 2c) + (-2a + 2d)x + (2a - 2d)x^2 + (2b - 2c)x^3.$$

Using the standard basis \mathcal{S} of $\mathcal{P}_3(\mathbb{F})$, we have

$$A = [L]_{\mathcal{S}} = \begin{bmatrix} 0 & -2 & 2 & 0 \\ -2 & 0 & 0 & 2 \\ 2 & 0 & 0 & -2 \\ 0 & 2 & -2 & 0 \end{bmatrix}.$$

Then $C_A(\lambda) = \lambda^2(\lambda + 4)(\lambda - 4)$. So the eigenvalues of A are $\lambda_1 = 0$, $\lambda_2 = -4$ and $\lambda_3 = 4$, with algebraic multiplicities $a_{\lambda_1} = 2$, $a_{\lambda_2} = a_{\lambda_3} = 1$. Just as in the previous example, we can immediately conclude that $g_{\lambda_2} = g_{\lambda_3} = 1$.

So it remains to determine $g_{\lambda_1} = \dim E_{\lambda_1}(A) = \text{nullity}(A - 0I) = \text{nullity}(A)$. A quick row reduction of A leads to

$$\begin{bmatrix} 0 & -2 & 2 & 0 \\ -2 & 0 & 0 & 2 \\ 2 & 0 & 0 & -2 \\ 0 & 2 & -2 & 0 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & 0 & 0 & -1 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

Thus, $\text{nullity}(A) = 2$. So $g_{\lambda_1} = a_{\lambda_1} = 2$ and therefore A (hence L) is diagonalizable.

From this alone we know that A must be similar to the diagonal matrix $D = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -4 & 0 \\ 0 & 0 & 0 & 4 \end{bmatrix}$.

But let's actually find a diagonalizing basis. For this, we need to find bases for $E_0(A)$, $E_{-4}(A)$ and $E_4(A)$.

Going through our usual row reduction process (steps omitted), we find that

$$E_0(A) = \text{Null}(A - 0I) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix} \right\},$$

$$E_4(A) = \text{Null}(A - 4I) = \text{Span} \left\{ \begin{bmatrix} 1 \\ -1 \\ 1 \\ -1 \end{bmatrix} \right\},$$

and

$$E_{-4}(A) = \text{Null}(A + 4I) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 1 \\ -1 \\ -1 \end{bmatrix} \right\}.$$

Therefore,

$$\left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \\ 1 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ -1 \\ -1 \end{bmatrix} \right\}$$

is a basis for \mathbb{F}^4 that diagonalizes A . If we view these vectors as being coordinate vectors with respect to the standard basis $\mathcal{S} = \{1, x, x^2, x^3\}$, we can convert them to obtain the basis $\mathcal{D} = \{1 + x^3, x + x^2, 1 - x + x^2 - x^3, 1 + x - x^2 - x^3\}$ that diagonalizes L . Let's check directly that

$$[L]_{\mathcal{D}} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -4 & 0 \\ 0 & 0 & 0 & 4 \end{bmatrix}.$$

We have

$$L(1 + x^3) = 0, \quad \text{so } [L(1 + x^3)]_{\mathcal{D}} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$$L(x + x^2) = 0, \quad \text{so } [L(x + x^2)]_{\mathcal{D}} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$$L(1 - x + x^2 - x^3) = 4(1 - x + x^2 - x^3), \quad \text{so } [L(1 - x + x^2 - x^3)]_{\mathcal{D}} = \begin{bmatrix} 0 \\ 0 \\ 4 \\ 0 \end{bmatrix}$$

$$L(1 + x - x^2 - x^3) = -4(1 + x - x^2 - x^3), \quad \text{so } [L(1 + x - x^2 - x^3)]_{\mathcal{D}} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ -4 \end{bmatrix}.$$

Example 3.2.20 Let $L: \mathcal{P}_1(\mathbb{F}) \rightarrow \mathcal{P}_1(\mathbb{F})$ be defined by

$$L(a + bx) = (a + 2b) + (b - 2a)x.$$

Using the standard basis \mathcal{S} of $\mathcal{P}_1(\mathbb{F})$, we have

$$A = [L]_{\mathcal{S}} = \begin{bmatrix} 1 & 2 \\ -2 & 1 \end{bmatrix}.$$

Then $C_A(\lambda) = \lambda^2 - 2\lambda + 5$. If $\mathbb{F} = \mathbb{R}$ then there are no real eigenvalues, and so no hope for diagonalization over \mathbb{R} !

If $\mathbb{F} = \mathbb{C}$ then, using the quadratic formula, we find that the roots of $C_A(\lambda)$ are $\lambda_1 = 1 + 2i$ and $\lambda_2 = 1 - 2i$. Thus, A is diagonalizable over \mathbb{C} . We leave it to you to find a diagonalizing basis for L .

Exercise 36 Complete the previous example and find a basis \mathcal{D} for $\mathcal{P}_1(\mathbb{C})$ such that $[L]_{\mathcal{D}}$ is diagonal.

3.3 Applications of Diagonalization

We now know how to figure out whether or not a linear operator $L: V \rightarrow V$ is diagonalizable, and even better, how to find a basis that diagonalizes L . We have learned that this is equivalent to the problem of determining whether or not a square matrix is diagonalizable.

There are several practical applications of diagonalization. You might be already familiar with one: taking powers of matrices. If A is diagonalizable, say with $A = PDP^{-1}$, then we can quickly compute A^k as

$$\begin{aligned} A^k &= (PDP^{-1})(PDP^{-1}) \cdots (PDP^{-1}) \\ &= PD(P^{-1}P)D(P^{-1}P) \cdots (P^{-1}P)DP^{-1} \\ &= PD^kP^{-1}. \end{aligned}$$

This is useful because if

$$D = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_n \end{bmatrix}$$

then

$$D^k = \begin{bmatrix} \lambda_1^k & 0 & \cdots & 0 \\ 0 & \lambda_2^k & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_n^k \end{bmatrix}.$$

So, for instance, to compute A^{1000} we don't need to multiply A with itself 1000 times. We can simply perform *two* matrix multiplications to compute $PD^{1000}P^{-1}$ instead.

Example 3.3.1

Suppose $A = \begin{bmatrix} -1 & 3 & -1 \\ -3 & 5 & -1 \\ -3 & 3 & 1 \end{bmatrix}$. Then $A = PDP^{-1}$ where

$$P = \begin{bmatrix} 1 & 1 & -1 \\ 1 & 1 & 0 \\ 1 & 0 & 3 \end{bmatrix} \quad \text{and} \quad D = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix}.$$

Therefore

$$\begin{aligned} A^{100} &= PD^{100}P^{-1} \\ &= \begin{bmatrix} 1 & 1 & -1 \\ 1 & 1 & 0 \\ 1 & 0 & 3 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2^{100} & 0 \\ 0 & 0 & 2^{100} \end{bmatrix} \begin{bmatrix} 3 & -3 & 1 \\ -3 & 4 & -1 \\ -1 & 1 & 0 \end{bmatrix} \\ &= \begin{bmatrix} 3 - 2 \cdot 2^{100} & -3 + 3 \cdot 2^{100} & 1 - 2^{100} \\ 3 - 3 \cdot 2^{100} & -3 + 4 \cdot 2^{100} & 1 - 2^{100} \\ 3 + 3 \cdot 2^{100} & -3 - 3 \cdot 2^{100} & 1 \end{bmatrix} \end{aligned}$$

Taking powers of matrices like this arises when studying dynamical systems or Markov chains in probability theory, to name just a couple. In particular, the Google PageRank algorithm involves these ideas. If you're interested to learn how, check out the video [How does Google Google?](#) on mathtube.org by Margot Gerritsen.

To close this section, we'll take a look at another application of diagonalization. We'll see some other applications later in this course.

Decoupling Differential Equations

It often comes up when modelling nature that you have a few differential equations where the variables from each equation appear in every other one. These can be extremely difficult to solve, but occasionally we can separate out the variables in a process called *decoupling*. It turns out that if your differential equations take on a very precise form, we can diagonalize a matrix and as a result decouple the equations. Let's see this in an example.

Example 3.3.2

Xavier and Yvonne are in a zombie apocalypse, and they are continuously killing zombies, and the following things are true about the way they kill zombies.

- Both get better with practice. (A reasonable assumption.)

- Both slow down as the other kills zombies. (They stop to congratulate the other person, and to give them a high-five if they are within slapping distance.)
- For every zombie Xavier kills, his kill rate goes up by a factor of 5 and Yvonne's goes down by a factor of 3.
- For every zombie Yvonne kills, her kill rate increases by a factor of 2 and Xavier's decreases by a factor of 6.

So, if we let x be the number of zombies killed by Xavier, y the number killed by Yvonne and t the time since the apocalypse started, we can set up the following system of differential equations that models the situation at hand.

$$\begin{aligned}\frac{dx}{dt} &= 5x - 3y \\ \frac{dy}{dt} &= -6x + 2y.\end{aligned}$$

Seemingly out of nowhere, let's make the substitutions

$$u = -\frac{2}{3}x + \frac{1}{3}y \quad \text{and} \quad w = \frac{1}{3}x + \frac{1}{3}y.$$

We now have

$$\begin{aligned}\frac{du}{dt} &= -\frac{2}{3}\frac{dx}{dt} + \frac{1}{3}\frac{dy}{dt} \\ &= -\frac{2}{3}(5x - 3y) + \frac{1}{3}(-6x + 2y) \\ &= -\frac{16}{3}x + \frac{8}{3}y \\ &= 8u\end{aligned}$$

and

$$\begin{aligned}\frac{dw}{dt} &= \frac{1}{3}(5x - 3y) + \frac{1}{3}(-6x + 2y) \\ &= -\frac{1}{3}x - \frac{1}{3}y \\ &= -w.\end{aligned}$$

These differential equations are much easier to deal with, and can be easily solved, and then converted back to our x and y variables. That's not the important thing here. The question you should have burning in your mind is, "*how did we choose u and w ?*"

To answer this, we do what we do best: bring matrices into the picture!

Let

$$\vec{x} = \begin{bmatrix} x \\ y \end{bmatrix}, \quad \frac{d\vec{x}}{dt} = \begin{bmatrix} \frac{dx}{dt} \\ \frac{dy}{dt} \end{bmatrix}, \quad \text{and} \quad A = \begin{bmatrix} 5 & -3 \\ -6 & 2 \end{bmatrix}.$$

Then

$$\frac{d\vec{x}}{dt} = A\vec{x}.$$

Now, it turns out that A is diagonalizable. In fact, $P^{-1}AP = D$ where

$$P = \begin{bmatrix} -1 & 1 \\ 1 & 2 \end{bmatrix} \quad \text{and} \quad D = \begin{bmatrix} 8 & 0 \\ 0 & -1 \end{bmatrix}$$

so

$$\frac{d\vec{x}}{dt} = PDP^{-1}\vec{x}.$$

If we let

$$\vec{v} = \begin{bmatrix} u \\ w \end{bmatrix} = P^{-1}\vec{x} = \begin{bmatrix} -\frac{2}{3}x + \frac{1}{3}y \\ \frac{1}{3}x + \frac{1}{3}y \end{bmatrix}$$

we have $\vec{x} = P\vec{v}$. Putting this back into the differential equation gives

$$\frac{d(P\vec{v})}{dt} = P\frac{d\vec{v}}{dt} = PD\vec{v}.$$

Multiplying on the left by P^{-1} gives

$$\frac{d\vec{v}}{dt} = D\vec{v}$$

or

$$\begin{bmatrix} \frac{du}{dt} \\ \frac{dw}{dt} \end{bmatrix} = \begin{bmatrix} 8 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} u \\ w \end{bmatrix} = \begin{bmatrix} 8u \\ -w \end{bmatrix}$$

giving us our decoupled differential equations.

The take home message here is that the matrix P told us what substitution to make. This shouldn't be too surprising because P can always be thought of as a change of coordinate matrix—one that changes coordinates from the ones we started with to a more natural set of coordinates depending on the problem at hand.

In general, suppose you have a system of differential equations of the form

$$\begin{aligned} \frac{dx}{dt} &= ax + by \\ \frac{dy}{dt} &= cx + dy. \end{aligned}$$

Here the variables x and y depend on each other, but it would be great if they didn't, since then you could solve two separate differential equations. If the matrix $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ is diagonalizable, then by a simple change of coordinates, we can decouple the differential equation.

Let $\vec{x} = \begin{bmatrix} x \\ y \end{bmatrix}$ and $\frac{d\vec{x}}{dt} = \begin{bmatrix} \frac{dx}{dt} \\ \frac{dy}{dt} \end{bmatrix}$. Then the set of equations above can be written as $\frac{d\vec{x}}{dt} = A\vec{x}$.

Suppose $P^{-1}AP = D$ for some diagonal matrix $D = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$. Let $\vec{x} = \begin{bmatrix} x \\ y \end{bmatrix}$ and let

$\vec{v} = \begin{bmatrix} u \\ w \end{bmatrix} = P^{-1}\vec{x}$. Then

$$\frac{d\vec{x}}{dt} = A\vec{x} \iff P\frac{d\vec{v}}{dt} = PDP^{-1}\vec{x} \iff \frac{d\vec{v}}{dt} = D\vec{v}.$$

Rewriting this last equation as a pair of differential equations we get

$$\begin{aligned}\frac{du}{dt} &= \lambda_1 u \\ \frac{dw}{dt} &= \lambda_2 w.\end{aligned}$$

Now we have two decoupled equations, each of which can be solved independently of the other. This process of decoupling is easily generalized to n variables and n equations.

Chapter 4

Inner Product Spaces

4.1 Inner Products

We have seen that if you have a finite-dimensional real vector space, once you pick a basis, you may as well think of the vector space as \mathbb{R}^n (and indeed, any n -dimensional real vector space is isomorphic to \mathbb{R}^n). This provides us with plenty of geometric intuition. A very useful feature of \mathbb{R}^2 and \mathbb{R}^3 is that they come with well-defined notions of length and angle.


In \mathbb{R}^2 , we know that the length of a vector $\vec{v} = [v_1 \ v_2]^T$ is given by

$$\|\vec{v}\| = \sqrt{v_1^2 + v_2^2}.$$

This formula is given to us by the Pythagorean theorem.

To compute the angle θ between two non-zero vectors $\vec{v} = [v_1 \ v_2]^T$ and $\vec{w} = [w_1 \ w_2]^T$, we invoke the cosine rule:

$$\begin{aligned} \cos \theta &= \frac{\|\vec{v}\|^2 + \|\vec{w}\|^2 - \|\vec{v} - \vec{w}\|^2}{2 \|\vec{v}\| \|\vec{w}\|} \\ &= \frac{v_1^2 + v_2^2 + w_1^2 + w_2^2 - (v_1 - w_1)^2 - (v_2 - w_2)^2}{2 \|\vec{v}\| \|\vec{w}\|} \\ &= \frac{v_1 w_1 + v_2 w_2}{\|\vec{v}\| \|\vec{w}\|}. \end{aligned}$$



In \mathbb{R}^3 , a similar thing occurs. Let $\vec{v} = [v_1 \ v_2 \ v_3]^T$ and $\vec{w} = [w_1 \ w_2 \ w_3]^T$ be two vectors in \mathbb{R}^3 . Then the length of \vec{v} is

$$\|\vec{v}\| = \sqrt{v_1^2 + v_2^2 + v_3^2}$$

(see the exercise below) and the angle θ between \vec{v} and \vec{w} (if they are non-zero) is given by


$$\cos \theta = \frac{v_1 w_1 + v_2 w_2 + v_3 w_3}{\|\vec{v}\| \|\vec{w}\|}.$$

By examining the diagram on the right and applying the Pythagorean theorem, show that

$$\|\vec{v}\| = \sqrt{v_1^2 + v_2^2 + v_3^2}.$$

Exercise 37

Can you similarly derive the expression for $\cos \theta$ given above?



Just like electricity and magnetism are actually two sides of the same coin, angles and lengths in \mathbb{R}^2 and \mathbb{R}^3 are just two sides of the same coin, and that coin is the dot product.

Recall that the dot product on \mathbb{R}^n is defined as

$$\begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \cdot \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} = v_1 w_1 + \cdots + v_n w_n.$$

So it is a function that takes as input two vectors in \mathbb{R}^n and produces a real number as output. Our expressions above for length and angle can be reformulated as

$$\|\vec{v}\| = \sqrt{\vec{v} \cdot \vec{v}} \quad \text{and} \quad \cos \theta = \frac{\vec{v} \cdot \vec{w}}{\|\vec{v}\| \|\vec{w}\|}$$

where θ is the angle between \vec{v} and \vec{w} . Since the dot product is defined on \mathbb{R}^n (and not just \mathbb{R}^2 and \mathbb{R}^3), we can use this to *define* length and angle in \mathbb{R}^n .

If we are to generalize the dot product to other vector spaces, we would want it to satisfy certain properties. Whatever our generalization is, it should be a function

$$\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{F}$$

that takes as input two vectors $\vec{v}, \vec{w} \in V$ and produces a real number $\langle \vec{v}, \vec{w} \rangle$ as output. We would then like to use this to define the length of a vector \vec{v} as $\|\vec{v}\| = \sqrt{\langle \vec{v}, \vec{v} \rangle}$.

This expression should behave the same way the length does in \mathbb{R}^2 and \mathbb{R}^3 . For example, we would like the length of a vector to be a positive real number, that is $\|\vec{v}\| \geq 0$. We would also want the length of a vector to be zero if and only if the vector itself is the zero vector.

We would then like to define the angle θ between two non-zero vectors $\vec{v}, \vec{w} \in V$ via

$$\cos \theta = \frac{\langle \vec{v}, \vec{w} \rangle}{\|\vec{v}\| \|\vec{w}\|}$$

just as we did above. For this to make sense, we need to know that the number on the right-side lies in the interval $[-1, 1]$. That is, we would like to have

$$\frac{|\langle \vec{v}, \vec{w} \rangle|}{\|\vec{v}\| \|\vec{w}\|} \leq 1.$$

Furthermore, in \mathbb{R}^2 and \mathbb{R}^3 we know the shortest path between two points is a straight line, which is reflected in the triangle inequality. That is, we would like it to be true if

$$\|\vec{v} + \vec{w}\| \leq \|\vec{v}\| + \|\vec{w}\|.$$

Finally, the length of a vector ought to behave nicely under scalar multiplication. That is, we would like to have $\|\alpha\vec{v}\| = |\alpha| \|\vec{v}\|$.

With these desires in mind, let's define an inner product on an arbitrary vector space. Our definition will produce a notion of length and angle that will meet *all* of our requirements above!

Definition 4.1.1

**Inner Product,
Conjugate
Symmetry,
Linearity in First
Argument,
Positive-Definite**

Let V be a vector space over \mathbb{F} . An **inner product** on V is a function

$$\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{F}$$

such that for all $\vec{u}, \vec{v}, \vec{w} \in V$ and $\alpha \in \mathbb{F}$,

1. $\langle \vec{v}, \vec{w} \rangle = \overline{\langle \vec{w}, \vec{v} \rangle}$.
2. $\langle \alpha\vec{v}, \vec{w} \rangle = \alpha \langle \vec{v}, \vec{w} \rangle$.
3. $\langle \vec{u} + \vec{v}, \vec{w} \rangle = \langle \vec{u}, \vec{w} \rangle + \langle \vec{v}, \vec{w} \rangle$.
4. (a) $\langle \vec{v}, \vec{v} \rangle \geq 0$.
(b) If $\langle \vec{v}, \vec{v} \rangle = 0$ then $\vec{v} = \vec{0}$.

A vector space V equipped with an inner product $\langle \cdot, \cdot \rangle$ is called an **inner product space**.

Property 1 is called **conjugate symmetry**. Properties 2 and 3 together are called **linearity in the first argument**. Property 4 is called **positive definiteness**, and we say that $\langle \cdot, \cdot \rangle$ is **positive definite**.

Note that if $\mathbb{F} = \mathbb{R}$, then $\bar{a} = a$ for all $a \in \mathbb{R}$ so the first property becomes $\langle \vec{v}, \vec{w} \rangle = \langle \vec{w}, \vec{v} \rangle$.

Let's see some examples.

Example 4.1.2

Of course, the dot product on \mathbb{R}^n satisfies properties 1–4, and so it defines an inner product on \mathbb{R}^n . The dot product on \mathbb{C}^n , however, is **not** an inner product.

Exercise 38

- (a) Carefully check that the dot product is indeed an inner product on \mathbb{R}^n .
- (b) Show that the dot product on \mathbb{C}^2 is not an inner product by finding a vector $\vec{z} \in \mathbb{C}^2$ such that $\vec{z} \cdot \vec{z}$ is a negative real number.

Example 4.1.3 (Standard inner product on \mathbb{C}^n)

Let $\vec{v}, \vec{w} \in \mathbb{C}^n$ where

$$\vec{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \quad \text{and} \quad \vec{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix}.$$

We define the standard inner product on \mathbb{C}^n by

$$\langle \vec{v}, \vec{w} \rangle = v_1 \bar{w}_1 + \cdots + v_n \bar{w}_n.$$

Let's prove that the standard inner product is indeed an inner product.

Proof: Let $\vec{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}$, $\vec{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix}$, $\vec{u} = \begin{bmatrix} u_1 \\ \vdots \\ u_n \end{bmatrix} \in V$ and $\alpha \in \mathbb{C}$. Then

$$\begin{aligned} \langle \vec{v}, \vec{w} \rangle &= v_1 \bar{w}_1 + \cdots + v_n \bar{w}_n \\ &= \bar{w}_1 v_1 + \cdots + \bar{w}_n v_n \\ &= \overline{w_1 \bar{v}_1 + \cdots + w_n \bar{v}_n} \\ &= \overline{\langle \vec{w}, \vec{v} \rangle} \end{aligned}$$

so property 1 holds. We have

$$\begin{aligned} \langle \alpha \vec{v}, \vec{w} \rangle &= \alpha v_1 \bar{w}_1 + \cdots + \alpha v_n \bar{w}_n \\ &= \alpha (v_1 \bar{w}_1 + \cdots + v_n \bar{w}_n) \\ &= \alpha \langle \vec{v}, \vec{w} \rangle \end{aligned}$$

and

$$\begin{aligned} \langle \vec{u} + \vec{v}, \vec{w} \rangle &= (u_1 + v_1) \bar{w}_1 + \cdots + (u_n + v_n) \bar{w}_n \\ &= u_1 \bar{w}_1 + v_1 \bar{w}_1 + \cdots + u_n \bar{w}_n + v_n \bar{w}_n \\ &= (u_1 \bar{w}_1 + \cdots + u_n \bar{w}_n) + (v_1 \bar{w}_1 + \cdots + v_n \bar{w}_n) \\ &= \langle \vec{u}, \vec{w} \rangle + \langle \vec{v}, \vec{w} \rangle \end{aligned}$$

so properties 2 and 3 hold. For 4(a) we have

$$\langle \vec{v}, \vec{v} \rangle = v_1 \bar{v}_1 + \cdots + v_n \bar{v}_n = |v_1|^2 + \cdots + |v_n|^2 \geq 0.$$

Finally, suppose $\langle \vec{v}, \vec{v} \rangle = |v_1|^2 + \cdots + |v_n|^2 = 0$. Since each $|v_i|^2$ is a positive real number, the only way this can be true is if $|v_1| = \cdots = |v_n| = 0$. This in turn implies $v_1 = \cdots = v_n = 0$ so $\vec{v} = \vec{0}$.

Thus, the standard inner product on \mathbb{C}^n is indeed an inner product. \square

In summary, for $\vec{v}, \vec{w} \in \mathbb{F}^n$, where

$$\vec{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \quad \text{and} \quad \vec{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix},$$

we define the standard inner product on \mathbb{F}^n by

$$\langle \vec{v}, \vec{w} \rangle = v_1 \bar{w}_1 + \cdots + v_n \bar{w}_n.$$

If $\mathbb{F} = \mathbb{R}$, this gives the dot product, and if $\mathbb{F} = \mathbb{C}$ this gives the standard inner product on \mathbb{C}^n defined in the previous example.

Before we move on to other examples of inner products, let's state a few of their important properties.

Proposition 4.1.4

Let V be an inner product space. For all $\vec{v}, \vec{u}, \vec{w} \in V$ and $\alpha \in \mathbb{F}$ the following properties are true.

- (a) $\langle \vec{0}, \vec{v} \rangle = \langle \vec{v}, \vec{0} \rangle = 0$.
- (b) $\langle \vec{v}, \alpha \vec{w} \rangle = \bar{\alpha} \langle \vec{v}, \vec{w} \rangle$.
- (c) $\langle \vec{v}, \vec{u} + \vec{w} \rangle = \langle \vec{v}, \vec{u} \rangle + \langle \vec{v}, \vec{w} \rangle$.

Proof: We will prove part (a). There are a couple of ways to do this. Perhaps the slickest is to note that $\vec{0} = 0 \cdot \vec{0}$. So linearity in the first argument gives

$$\langle \vec{0}, \vec{v} \rangle = \langle 0 \cdot \vec{0}, \vec{v} \rangle = 0 \langle \vec{0}, \vec{v} \rangle = 0.$$

This proves half of part (a). For the second half, we use conjugate-symmetry:

$$\langle \vec{v}, \vec{0} \rangle = \overline{\langle \vec{0}, \vec{v} \rangle} = \bar{0} = 0.$$

Parts (b) and (c) can be proved quickly by appealing to conjugate symmetry. We'll leave the details as an exercise. \square

Exercise 39

Prove parts (b) and (c) of Proposition 4.1.4.

Now let's explore some examples of inner products on other vector spaces.

Example 4.1.5

For $p, q \in \mathcal{P}_n(\mathbb{R})$, define

$$\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx.$$

This is an inner product on $\mathcal{P}_n(\mathbb{R})$. We will check property 4 and leave the rest as an exercise.

If $p \in \mathcal{P}_n(\mathbb{R})$ then

$$\langle p, p \rangle = \int_{-1}^1 (p(x))^2 dx$$

is non-negative because it is the integral of a non-negative function. Moreover, as we know from calculus, the integral of a non-negative continuous function (such as a polynomial) over an interval is zero if and only if that function is itself zero on the interval. This shows that $\langle p, p \rangle = 0$ if and only if $p(x) = 0$ for all $x \in [-1, 1]$. Any polynomial which evaluates to 0 at infinitely many distinct points is the zero polynomial (why?), so we can conclude $p = 0$.

Exercise 40

Check that $\langle \cdot, \cdot \rangle$ defined in the previous example satisfies properties 1–3 and therefore is an inner product on $\mathcal{P}_3(\mathbb{R})$.

Example 4.1.6

(Frobenius inner product on $M_{m \times n}(\mathbb{R})$)

For $A, B \in M_{m \times n}(\mathbb{R})$, define

$$\langle A, B \rangle = \text{tr}(B^T A).$$

This is an inner product on $M_{m \times n}(\mathbb{R})$. It is called the **Frobenius inner product**.

If $A = [a_{ij}]$ and $B = [b_{ij}]$, then the above expression can be expanded as

$$\langle A, B \rangle = \sum_{i=1}^m \sum_{j=1}^n a_{ij} b_{ij}.$$

This makes it easy to compute the Frobenius inner product: simply multiply the (i, j) th entries of A and B together and then add up all these products. For instance,

$$\left\langle \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}, \begin{bmatrix} 1 & -1 \\ 0 & 5 \end{bmatrix} \right\rangle = 1(1) + 2(-1) + 3(0) + 4(5) = 19.$$

Exercise 41

- Verify that $\langle \cdot, \cdot \rangle$ defined in the preceding example is an inner product on $M_{m \times n}(\mathbb{R})$.
- What does this inner product give if $n = 1$ (so that $M_{m \times n}(\mathbb{R}) = \mathbb{R}^m$)?

Exercise 42

(Frobenius inner product on $M_{m \times n}(\mathbb{C})$)

For a matrix $A \in M_{m \times n}(\mathbb{C})$, we define its **adjoint** (or **conjugate-transpose**) to be the matrix $A^* = \overline{A^T}$ in $M_{n \times m}(\mathbb{C})$. We will return to this notion in Chapter 5.

- Show that $\langle A, B \rangle = \text{tr}(B^* A)$ defines an inner product, called the **Frobenius inner product**, on $M_{m \times n}(\mathbb{C})$.
- Show that $\langle A, B \rangle = \text{tr}(A^* B)$ does **not** define an inner product on $M_{m \times n}(\mathbb{C})$.

In summary, the Frobenius inner product on $M_{m \times n}(\mathbb{F})$ can be defined as $\langle A, B \rangle = \text{tr}(B^* A)$. While this may seem like a rather strange way to define an inner product, in reality it is nothing but the inner product in $\mathbb{C}^{m \times n}$ between two coordinate vectors $[A]_{\mathcal{S}}$ and $[B]_{\mathcal{S}}$ with respect to the standard basis \mathcal{S} of $M_{m \times n}(\mathbb{F})$. For example, if

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix},$$

then

$$\langle A, B \rangle = \left\langle \begin{bmatrix} a_{11} \\ a_{12} \\ a_{21} \\ a_{22} \end{bmatrix}, \begin{bmatrix} a_{11} \\ a_{12} \\ a_{21} \\ a_{22} \end{bmatrix} \right\rangle = a_{11}\overline{b_{11}} + a_{12}\overline{b_{12}} + a_{21}\overline{b_{21}} + a_{22}\overline{b_{22}}.$$

It may be natural to ask at this point whether or not every vector space can be turned into an inner product space. Let's answer that now.

Proposition 4.1.7 Every finite-dimensional vector space admits an inner product.

Proof: Let V be a vector space with basis $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$. Let $\vec{v} = t_1\vec{v}_1 + \dots + t_n\vec{v}_n$ and $\vec{w} = s_1\vec{v}_1 + \dots + s_n\vec{v}_n$. Then it is left as an exercise to check

$$\langle \vec{v}, \vec{w} \rangle = t_1\overline{s_1} + \dots + t_n\overline{s_n}$$

is an inner product on V . □

Exercise 43 Complete the proof of Proposition 4.1.7.

It is true that every infinite-dimensional vector space also admits an inner product, but we won't be proving that here.

Since the inner product of a vector with itself is always a positive real number, it now makes sense to define the length of a vector.

Definition 4.1.8 Let \vec{v} be a vector in an inner product space V . The **norm** (or **length**) of \vec{v} is defined by

$$\|\vec{v}\| = \sqrt{\langle \vec{v}, \vec{v} \rangle}.$$

This definition of “length” satisfies many of the properties one would expect length to satisfy. See the following exercise and Proposition 4.2.7.

Exercise 44 Prove that, in any inner product space, $\|\vec{0}\| = 0$.

Example 4.1.9

On $\mathcal{P}_2(\mathbb{C})$ consider

$$\langle p, q \rangle = p(i)\overline{q(i)} + p(-i)\overline{q(-i)} + p(1)\overline{q(1)}.$$

This is an inner product. Indeed, it is left as an exercise to check properties 1, 2, and 3. For 4a,

$$\langle p, p \rangle = |p(i)|^2 + |p(-i)|^2 + |p(1)|^2 \geq 0.$$

For 4b, suppose $\langle p, p \rangle = 0$. Then

$$|p(i)|^2 + |p(-i)|^2 + |p(1)|^2 = 0$$

so $p(i) = p(-i) = p(1) = 0$. Since p is a polynomial of degree at most 2 that evaluates to 0 at three distinct points, we must have $p = 0$. Alternatively, suppose $p = ax^2 + bx + c$. Then $p(i) = p(-i) = p(1) = 0$ gives the system of equations

$$-a + bi + c = 0$$

$$-a - bi + c = 0$$

$$a + b + c = 0.$$

Plugging this into an augmented matrix and row-reducing, we get

$$\left[\begin{array}{ccc|c} -1 & i & 1 & 0 \\ -1 & -i & 1 & 0 \\ 1 & 1 & 1 & 0 \end{array} \right] \rightarrow \left[\begin{array}{ccc|c} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{array} \right],$$

so $a = b = c = 0$ and therefore $p = 0$.

So this is an inner product. With respect to this inner product, let's compute the norm of some vectors. We have

$$\|1\| = \sqrt{\langle 1, 1 \rangle} = \sqrt{3}, \quad \|x\| = \sqrt{\langle x, x \rangle} = \sqrt{|i|^2 + |-i|^2 + |1|^2} = \sqrt{3}$$

and

$$\|1 + x\| = \sqrt{\langle 1 + x, 1 + x \rangle} = \sqrt{|1 + i|^2 + |1 - i|^2 + |1 + 1|^2} = \sqrt{8}.$$

Notice that $\|1 + x\| \neq \|1\| + \|x\|$.

In the next section, we'll explore the properties of the norm more extensively.

4.2 Orthogonality and Norm

Recall in \mathbb{R}^n that the angle θ between two non-zero vectors \vec{v} and \vec{w} is given by

$$\cos \theta = \frac{\vec{v} \cdot \vec{w}}{\|\vec{v}\| \|\vec{w}\|}.$$

Therefore if the dot product of two vectors is 0, we know they are perpendicular. The notion of being perpendicular (or *orthogonal* as we will call it) turns out to be an extremely useful notion in inner product spaces.

Definition 4.2.1**Orthogonal, \perp**

Let V be an inner product space. We say \vec{v} is **orthogonal** to \vec{w} , and write $\vec{v} \perp \vec{w}$, if $\langle \vec{v}, \vec{w} \rangle = 0$.

Notice that $\langle \vec{v}, \vec{w} \rangle = 0$ if and only if $\langle \vec{w}, \vec{v} \rangle = 0$ (why?) so the definition is symmetric in \vec{v} and \vec{w} . That is, we're safe to say that \vec{v} and \vec{w} are orthogonal, instead of \vec{v} is orthogonal to \vec{w} or vice versa.

Example 4.2.2

Consider $\mathcal{P}_2(\mathbb{R})$ with the inner product

$$\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx.$$

Then

$$\langle 1, x \rangle = \int_{-1}^1 x dx = 0$$

so 1 and x are orthogonal. However,

$$\langle 1, x^2 \rangle = \int_{-1}^1 x^2 dx = \frac{2}{3}$$

so 1 and x^2 are not orthogonal.

One thing we know from geometry is that if we have a right-angled triangle, then Pythagoras' theorem holds. So, if we are to believe that being orthogonal really means that two vectors are at right angles to each other, we should expect the Pythagorean theorem to hold. Indeed it does!

Proposition 4.2.3**(Pythagorean Theorem)**

Let V be an inner product space. Suppose $\vec{v} \perp \vec{w}$. Then $\|\vec{v} + \vec{w}\|^2 = \|\vec{v}\|^2 + \|\vec{w}\|^2$.

Proof: Since $\vec{v} \perp \vec{w}$, we have $\langle \vec{v}, \vec{w} \rangle = 0$ and $\langle \vec{w}, \vec{v} \rangle = 0$. Consequently,

$$\begin{aligned} \|\vec{v} + \vec{w}\|^2 &= \langle \vec{v} + \vec{w}, \vec{v} + \vec{w} \rangle \\ &= \langle \vec{v}, \vec{v} + \vec{w} \rangle + \langle \vec{w}, \vec{v} + \vec{w} \rangle \\ &= \langle \vec{v}, \vec{v} \rangle + \langle \vec{v}, \vec{w} \rangle + \langle \vec{w}, \vec{v} \rangle + \langle \vec{w}, \vec{w} \rangle \\ &= \langle \vec{v}, \vec{v} \rangle + 0 + 0 + \langle \vec{w}, \vec{w} \rangle \\ &= \|\vec{v}\|^2 + \|\vec{w}\|^2. \end{aligned}$$

This completes the proof. □

Our short term goal is to prove the Cauchy–Schwarz inequality (Theorem 4.2.5 below), which will allow us to define the angle between two vectors in an arbitrary inner product space. To do that we first need the following technical lemma.

Lemma 4.2.4

Let V be an inner product space and let $\vec{v}, \vec{w} \in V$ such that $\vec{w} \neq \vec{0}$. Then \vec{w} is orthogonal to $\vec{v} - \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \vec{w}$.

Proof: We simply need to take the inner product between these two vectors and show it is zero. We have

$$\begin{aligned} \left\langle \vec{v} - \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \vec{w}, \vec{w} \right\rangle &= \langle \vec{v}, \vec{w} \rangle - \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \langle \vec{w}, \vec{w} \rangle \\ &= \langle \vec{v}, \vec{w} \rangle - \langle \vec{v}, \vec{w} \rangle \\ &= 0, \end{aligned}$$

as desired. □

You may have come across the orthogonal projection of a vector \vec{v} onto \vec{w} in \mathbb{R}^n before, and the vector is given by $\frac{\vec{v} \cdot \vec{w}}{\vec{w} \cdot \vec{w}} \vec{w}$. This expression appeared in the previous proof, and we will revisit it later (see Section 4.4). Phrased in this language, Lemma 4.2.4 proves that the perpendicular component of \vec{v} when projected onto \vec{w} is indeed perpendicular to \vec{w} .

Theorem 4.2.5 (Cauchy–Schwarz Inequality)

Let V be an inner product space. Then

$$|\langle \vec{v}, \vec{w} \rangle| \leq \|\vec{v}\| \|\vec{w}\| \quad \text{for all } \vec{v}, \vec{w} \in V,$$

with equality if and only if \vec{v} and \vec{w} are scalar multiples of each other.

Proof: If $\vec{w} = \vec{0}$ we have $|\langle \vec{v}, \vec{w} \rangle| = \|\vec{v}\| \|\vec{w}\| = 0$ so the statement is true. Assume now that $\vec{w} \neq \vec{0}$. We have

$$\begin{aligned} \|\vec{v}\|^2 &= \left\| \vec{v} - \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \vec{w} + \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \vec{w} \right\|^2 \\ &= \left\| \vec{v} - \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \vec{w} \right\|^2 + \left\| \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \vec{w} \right\|^2 \quad (\text{by Proposition 4.2.3 and Lemma 4.2.4}) \\ &\geq 0 + \left\| \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \vec{w} \right\|^2 \\ &= \frac{|\langle \vec{v}, \vec{w} \rangle|^2}{\|\vec{w}\|^4} \|\vec{w}\|^2 \\ &= \frac{|\langle \vec{v}, \vec{w} \rangle|^2}{\|\vec{w}\|^2}. \end{aligned}$$

Since $\|\vec{w}\|^2$ is positive, this implies

$$\|\vec{v}\|^2 \|\vec{w}\|^2 \geq |\langle \vec{v}, \vec{w} \rangle|^2.$$

Since norms are always positive real numbers we can take square roots to obtain

$$\|\vec{v}\| \|\vec{w}\| \geq |\langle \vec{v}, \vec{w} \rangle|,$$

giving us the desired inequality. Tracing back our steps, we see that the inequality will be an *equality* if and only if

$$\left\| \vec{v} - \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \vec{w} \right\|^2 = \left\langle \vec{v} - \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \vec{w}, \vec{v} - \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \vec{w} \right\rangle = 0$$

which, by definition of the inner product, happens if and only if

$$\vec{v} - \frac{\langle \vec{v}, \vec{w} \rangle}{\langle \vec{w}, \vec{w} \rangle} \vec{w} = \vec{0}.$$

This completes the proof. \square

There is actually a geometric interpretation of the Cauchy–Schwarz inequality, at least in \mathbb{R}^2 . It turns out that the inequality is a rephrasing of the following fact from geometry: If you have a parallelogram with side lengths x and y , then the area of that parallelogram is maximized exactly when the parallelogram is a rectangle. It’s a fun exercise to try and see how this fact relates to the Cauchy–Schwarz inequality for the dot product on \mathbb{R}^2 !

Because of the Cauchy–Schwarz inequality, we can now sensibly define the angle between two non-zero vectors, at least when our vector space is over the field \mathbb{R} .

Definition 4.2.6

Angle

Let V be a real inner product space. The **angle** θ between two non-zero vectors \vec{v} and \vec{w} in V is defined by

$$\cos(\theta) = \frac{\langle \vec{v}, \vec{w} \rangle}{\|\vec{v}\| \|\vec{w}\|},$$

that is, θ is the unique real number in the interval $[0, \pi]$ given by

$$\theta = \cos^{-1} \left(\frac{\langle \vec{v}, \vec{w} \rangle}{\|\vec{v}\| \|\vec{w}\|} \right).$$

We do not define the angle between \vec{v} and \vec{w} if one of them is the zero vector.

There are various ways to define the angle between vectors in a complex inner product space, each serving a different purpose. We won’t be talking about the angle between vectors in a complex vector space in this course, except for the case when vectors are orthogonal.

We will finish this section by returning to one of our motivations for defining an inner product: a sensible notion of length. The next proposition shows us that our definition of norm provides such a notion.

Proposition 4.2.7

(Properties of Norm)

Let V be an inner product space. For all $\vec{v}, \vec{w} \in V$ and $\alpha \in \mathbb{F}$, the following properties are true.

- (a) $\|\alpha \vec{v}\| = |\alpha| \|\vec{v}\|$.
- (b) $\|\vec{v} + \vec{w}\| \leq \|\vec{v}\| + \|\vec{w}\|$ (this is called the **Triangle Inequality**).
- (c) $\|\vec{v}\| \geq 0$, and $\|\vec{v}\| = 0$ if and only if $\vec{v} = \vec{0}$.

Proof: Properties (a) and (c) are left as an exercise. For the triangle inequality we will make use of the Cauchy–Schwarz inequality. We have

$$\begin{aligned}
 \|\vec{v} + \vec{w}\|^2 &= \langle \vec{v} + \vec{w}, \vec{v} + \vec{w} \rangle \\
 &= \|\vec{v}\|^2 + \|\vec{w}\|^2 + \langle \vec{v}, \vec{w} \rangle + \overline{\langle \vec{v}, \vec{w} \rangle} \\
 &= \|\vec{v}\|^2 + \|\vec{w}\|^2 + 2 \operatorname{Re}(\langle \vec{v}, \vec{w} \rangle) \\
 &\leq \|\vec{v}\|^2 + \|\vec{w}\|^2 + 2 |\operatorname{Re}(\langle \vec{v}, \vec{w} \rangle)| \\
 &\leq \|\vec{v}\|^2 + \|\vec{w}\|^2 + 2 |\langle \vec{v}, \vec{w} \rangle| \\
 &\leq \|\vec{v}\|^2 + \|\vec{w}\|^2 + 2 \|\vec{v}\| \|\vec{w}\| \\
 &= (\|\vec{v}\| + \|\vec{w}\|)^2.
 \end{aligned}$$

Since both sides are positive we have $\|\vec{v} + \vec{w}\| \leq \|\vec{v}\| + \|\vec{w}\|$ completing the proof. \square

Exercise 45

- (a) Prove properties (a) and (c) in Proposition 4.2.7.
 (b) Determine when the Triangle Inequality is in fact an *equality*.

There is an abstract notion of a *norm* on an arbitrary vector space V that is not necessarily an inner product space. It's a function $\|\cdot\| : V \rightarrow \mathbb{F}$ that satisfies the three properties of Proposition 4.2.7. So what we have proved is that what we are calling the norm on an inner product space V is in fact an example of an abstract norm. It is true that there are norms that do not arise this way (i.e. norms that are not built from an inner product), but we will not be concerned with such norms in this course.

Finally, once we have a notion of length, we can speak about the distance between vectors.

Definition 4.2.8

Distance


Let V be an inner product space, and let $\vec{v}, \vec{w} \in V$. The **distance between \vec{v} and \vec{w}** is defined as

$$\operatorname{dist}(\vec{v}, \vec{w}) = \|\vec{v} - \vec{w}\|.$$

This definition agrees with our usual notion of distance in \mathbb{R}^2 and \mathbb{R}^3 .

Notice that the norm of a vector \vec{v} can now be interpreted to be the distance between \vec{v} and $\vec{0}$:

$$\|\vec{v}\| = \|\vec{v} - \vec{0}\| = \operatorname{dist}(\vec{v}, \vec{0}).$$



Example 4.2.9

In $\mathcal{P}_2(\mathbb{R})$ with $\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx$, the distance between $1 - x$ and $1 + x^2$ is given by

$$\begin{aligned} \text{dist}(1 - x, 1 + x^2) &= \|(1 - x) - (1 + x^2)\| \\ &= \|-x - x^2\| \\ &= |-1| \|x + x^2\| \\ &= \sqrt{\langle x + x^2, x + x^2 \rangle} \\ &= \sqrt{\int_{-1}^1 (x + x^2)^2 dx} \\ &= \sqrt{\frac{16}{15}}. \end{aligned}$$

Exercise 46 (Properties of Distance)

Let V be an inner product space. Show that the following properties are true for all $\vec{x}, \vec{y}, \vec{z} \in V$.

- (a) $\text{dist}(\vec{x}, \vec{y}) \geq 0$, and $\text{dist}(\vec{x}, \vec{y}) = 0$ if and only if $\vec{x} = \vec{y}$.
- (b) $\text{dist}(\vec{x}, \vec{y}) = \text{dist}(\vec{y}, \vec{x})$.
- (c) $\text{dist}(\vec{x}, \vec{z}) \leq \text{dist}(\vec{x}, \vec{y}) + \text{dist}(\vec{y}, \vec{z})$ (Triangle Inequality).

4.3 Orthonormal Bases

Consider the standard basis in \mathbb{R}^n equipped with the dot product. Each vector in this basis has length 1, and even better, any two vectors in the basis are orthogonal. This will be our gold standard to head towards.

Definition 4.3.1
Orthogonal Set

A set $\{\vec{v}_1, \dots, \vec{v}_k\}$ in an inner product space is called **orthogonal** if $\langle \vec{v}_i, \vec{v}_j \rangle = 0$ whenever $i \neq j$.

Definition 4.3.2
Unit Vector

A vector \vec{v} in an inner product space is a **unit vector** if $\|\vec{v}\| = 1$ (or, equivalently, if $\langle \vec{v}, \vec{v} \rangle = 1$).

Definition 4.3.3
Orthonormal Set

A set $\{\vec{v}_1, \dots, \vec{v}_k\}$ in an inner product space is an **orthonormal set** if it is an orthogonal set and if each vector \vec{v}_i in the set is a unit vector.

Example 4.3.4 Consider $M_{2 \times 2}(\mathbb{R})$ with inner product $\langle A, B \rangle = \text{tr}(B^T A)$ from Example 4.1.6.

Define the matrices

$$A = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}, \quad B = \begin{bmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{bmatrix}, \quad \text{and} \quad C = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

Then

$$\begin{aligned} \langle A, B \rangle &= \text{tr}(B^T A) = \text{tr} \left(\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \right) = 0 \\ \langle A, A \rangle &= \text{tr}(A^T A) = \text{tr} \left(\begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix} \right) = 1 \\ \langle B, B \rangle &= \text{tr}(B^T B) = \text{tr} \left(\begin{bmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{bmatrix} \right) = 1. \end{aligned}$$

Since $\langle A, C \rangle = \langle B, C \rangle = 0$, it follows that $\{A, B, C\}$ is an orthogonal set, but it is not orthonormal (since $\|C\| = 0$). However, $\{A, B\}$ is an orthonormal set.

Example 4.3.5 The set $\left\{ \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}, \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix} \right\}$ is orthonormal in \mathbb{R}^2 with respect to the dot product.

Example 4.3.6 (**Legendre polynomials in $\mathcal{P}_3(\mathbb{R})$**)

There are special polynomials called **Legendre polynomials** which arise in physics (specifically when solving Laplace's equation in spherical coordinates), and also in some special trigonometric identities!

The first four Legendre polynomials are the polynomials $\{1, x, \frac{3}{2}x^2 - \frac{1}{2}, \frac{5}{2}x^3 - \frac{3}{2}x\}$ in $\mathcal{P}_3(\mathbb{R})$. You can check that this is an orthogonal set in $\mathcal{P}_3(\mathbb{R})$ with respect to the inner product

$$\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx.$$

However, it is not an orthonormal set because

$$\|1\| = \sqrt{\int_{-1}^1 1 dx} = \sqrt{2}$$

$$\|x\| = \sqrt{\int_{-1}^1 x^2 dx} = \sqrt{\frac{2}{3}}$$

$$\left\| \frac{3}{2}x^2 - \frac{1}{2} \right\| = \sqrt{\frac{2}{5}}$$

$$\left\| \frac{5}{2}x^3 - \frac{3}{2}x \right\| = \sqrt{\frac{2}{7}}.$$

So none of these Legendre polynomials are unit vectors. If we divide each by its norm, the resulting set

$$\left\{ \frac{1}{\sqrt{2}}, \sqrt{\frac{3}{2}}x, \sqrt{\frac{5}{2}} \left(\frac{3}{2}x^2 - \frac{1}{2} \right), \sqrt{\frac{7}{2}} \left(\frac{5}{2}x^3 - \frac{3}{2}x \right) \right\}$$

would then be orthonormal. See the next exercise.

Given a non-zero vector \vec{v} , it is always possible to produce a vector of \hat{v} that points in the same direction of \vec{v} and has norm equal to one. We refer to this process as *normalization*.

Definition 4.3.7

Normalization

Let \vec{v} be a non-zero vector in an inner product space. The **normalization** of \vec{v} is the vector $\hat{v} = \frac{\vec{v}}{\|\vec{v}\|}$.

Exercise 47

- Prove that \hat{v} is a unit vector.
- Suppose that $S = \{\vec{v}_1, \dots, \vec{v}_k\}$ is an orthogonal set. Prove that $T = \{\hat{v}_1, \dots, \hat{v}_k\}$ is an orthonormal set.

If our model for an orthogonal set is the standard basis in \mathbb{R}^n , we should expect an orthogonal set to be linearly independent. Indeed that is the case—at least if none of the vectors is the zero vector.

Proposition 4.3.8

Suppose $\{\vec{v}_1, \dots, \vec{v}_k\}$ is orthogonal and $\vec{v}_i \neq \vec{0}$ for all i . Then $\{\vec{v}_1, \dots, \vec{v}_k\}$ is linearly independent.

Proof: Suppose $t_1 \vec{v}_1 + \dots + t_k \vec{v}_k = \vec{0}$. Fix $i \in \{1, \dots, k\}$. Then

$$\begin{aligned} 0 &= \langle t_1 \vec{v}_1 + \dots + t_k \vec{v}_k, \vec{v}_i \rangle \\ &= t_1 \langle \vec{v}_1, \vec{v}_i \rangle + \dots + t_i \langle \vec{v}_i, \vec{v}_i \rangle + \dots + t_k \langle \vec{v}_k, \vec{v}_i \rangle \\ &= t_i \|\vec{v}_i\|^2. \end{aligned}$$

Since $\vec{v}_i \neq \vec{0}$, we must have $t_i = 0$. Since this is true for all i , we conclude $\{\vec{v}_1, \dots, \vec{v}_k\}$ is linearly independent. \square

We now make a special definition for the case that we have an orthonormal set which forms a basis for an inner product space.

Definition 4.3.9

Orthogonal Basis, Orthonormal Basis

A set $\{\vec{v}_1, \dots, \vec{v}_n\}$ in an inner product space V is an **orthogonal basis** (resp. **orthonormal basis**) if it is a basis for V and it is an orthogonal set (resp. an orthonormal set).

Example 4.3.10

The following are orthonormal bases, as you can easily check.

- The standard basis of \mathbb{R}^n with respect to the dot product.
- The standard basis of \mathbb{C}^n with respect to the standard inner product on \mathbb{C}^n .

3. The standard basis of $M_{m \times n}(\mathbb{R})$ with respect to the Frobenius inner product $\langle A, B \rangle = \text{tr}(B^T A)$.
4. The standard basis of $M_{m \times n}(\mathbb{C})$ with respect to the Frobenius inner product $\langle A, B \rangle = \text{tr}(B^* A)$.

Example 4.3.11

The standard basis of $\mathcal{P}_n(\mathbb{R})$ with respect to the inner product

$$\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx$$

is **not** an orthonormal basis. In fact, it is not even an orthogonal set if $n > 1$. (See Example 4.2.2.)

In Example 4.3.6 we showed that the set

$$\left\{ \frac{1}{\sqrt{2}}, \frac{\sqrt{3}}{\sqrt{2}}x, \sqrt{\frac{5}{2}} \left(\frac{3}{2}x^2 - \frac{1}{2} \right), \sqrt{\frac{7}{2}} \left(\frac{5}{2}x^3 - \frac{3}{2}x \right) \right\}$$

is an orthonormal basis for $\mathcal{P}_3(\mathbb{R})$ with respect to the above inner product. In Section 4.5 we'll see how to create an orthonormal basis for $\mathcal{P}_n(\mathbb{R})$ for all $n \geq 1$.

Example 4.3.12

In \mathbb{C}^2 , with the inner product

$$\left\langle \begin{bmatrix} a \\ b \end{bmatrix}, \begin{bmatrix} c \\ d \end{bmatrix} \right\rangle = 2a\bar{c} + b\bar{d},$$

the set $\left\{ \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}$ is an orthonormal basis, as you can check.

Exercise 48

Check that the sets given in Examples 4.3.10 and 4.3.12 are orthonormal bases.

We close this section by giving a result that might explain why we sometimes prefer to work with an orthonormal (or even orthogonal) basis: it makes it very easy to find the coordinates of any given vector.

Proposition 4.3.13**(Coordinates Relative to an Orthogonal Basis)**

Let V be an inner product space and let $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$ be a basis for V . If $\vec{x} \in V$ is given by $\vec{x} = x_1 \vec{v}_1 + \dots + x_n \vec{v}_n$, then:

- (a) If \mathcal{B} is an orthogonal basis, then $x_i = \frac{\langle \vec{x}, \vec{v}_i \rangle}{\langle \vec{v}_i, \vec{v}_i \rangle}$ for $1 \leq i \leq n$.
- (b) If \mathcal{B} is an orthonormal basis, then $x_i = \langle \vec{x}, \vec{v}_i \rangle$ for $1 \leq i \leq n$.

Proof: If \mathcal{B} is orthogonal, then by taking the inner product of \vec{x} and \vec{v}_i we obtain

$$\begin{aligned}\langle \vec{x}, \vec{v}_i \rangle &= \langle x_1 \vec{v}_1 + \cdots + x_n \vec{v}_n, \vec{v}_i \rangle \\ &= x_1 \langle \vec{v}_1, \vec{v}_i \rangle + \cdots + x_i \langle \vec{v}_i, \vec{v}_i \rangle + \cdots + x_n \langle \vec{v}_n, \vec{v}_i \rangle \\ &= x_1 0 + \cdots + x_{i-1} 0 + x_i \langle \vec{v}_i, \vec{v}_i \rangle + x_{i+1} 0 + \cdots + x_n 0 \\ &= x_i \langle \vec{v}_i, \vec{v}_i \rangle.\end{aligned}$$

Now since $\langle \vec{v}_i, \vec{v}_i \rangle \neq 0$ (why?), we can divide through by $\langle \vec{v}_i, \vec{v}_i \rangle$ to obtain

$$x_i = \frac{\langle \vec{x}, \vec{v}_i \rangle}{\langle \vec{v}_i, \vec{v}_i \rangle}.$$

This proves (a). Part (b) follows immediately since if \mathcal{B} is orthonormal then $\langle \vec{v}_i, \vec{v}_i \rangle = 1$ for all i . \square

Example 4.3.14

In \mathbb{R}^n with the standard basis $\mathcal{S} = \{\vec{e}_1, \dots, \vec{e}_n\}$ (which is orthonormal with respect to the dot product), the above proposition says that we can obtain the i^{th} component of $\vec{x} = [x_1 \cdots x_n]^T$ as $\vec{x} \cdot \vec{e}_i$. And indeed we can:

$$\vec{x} \cdot \vec{e}_i = \begin{bmatrix} x_1 \\ \vdots \\ x_i \\ \vdots \\ x_n \end{bmatrix} \cdot \begin{bmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{bmatrix} = x_i.$$

Example 4.3.15

In $M_{2 \times 2}(\mathbb{R})$, let $\mathcal{B} = \left\{ \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ -1 & 0 \end{bmatrix}, \begin{bmatrix} 2 & -1 \\ 2 & -3 \end{bmatrix} \right\}$ and $W = \text{Span}(\mathcal{B})$. You can check that \mathcal{B} is an orthogonal basis for W with respect to the inner product $\langle A, B \rangle = \text{tr}(B^T A)$. Supposing we know that $A = \begin{bmatrix} 3 & -2 \\ -1 & -4 \end{bmatrix}$ is in W , we can find its coordinate vector $[A]_{\mathcal{B}} = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix}$ quickly as follows. We have

$$\begin{aligned}a_1 &= \frac{\left\langle A, \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \right\rangle}{\left\| \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \right\|^2} = \frac{3(1) + (-2)(1) + (-1)(1) + (-4)(1)}{1^2 + 1^2 + 1^2 + 1^2} = -1 \\ a_2 &= \frac{\left\langle A, \begin{bmatrix} 1 & 0 \\ -1 & 0 \end{bmatrix} \right\rangle}{\left\| \begin{bmatrix} 1 & 0 \\ -1 & 0 \end{bmatrix} \right\|^2} = \frac{3(1) + (-2)(0) + (-1)(-1) + (-4)(0)}{1^2 + (-1)^2 + 0^2 + 0^2} = 2 \\ a_3 &= \frac{\left\langle A, \begin{bmatrix} 2 & -1 \\ 2 & -3 \end{bmatrix} \right\rangle}{\left\| \begin{bmatrix} 2 & -1 \\ 2 & -3 \end{bmatrix} \right\|^2} = \frac{3(2) + (-2)(-1) + (-1)(2) + (-4)(-3)}{2^2 + (-1)^2 + 2^2 + (-3)^2} = 1.\end{aligned}$$

$$\text{Thus, } [A]_{\mathcal{B}} = \begin{bmatrix} -1 \\ 2 \\ 1 \end{bmatrix}.$$

Our goal now is to show that *every* finite-dimensional inner product space admits an orthonormal basis. As a fun exercise, try to see if you can prove the following special case.


Exercise 49

Let V be an inner product space and let $W = \text{Span}(\{\vec{v}, \vec{w}\})$, where $\{\vec{v}, \vec{w}\}$ is a linearly independent subset of V . Obtain an orthonormal basis for W . [**Hint:** Start by trying to find a vector $a\vec{v} + b\vec{w}$ in W that is orthogonal to \vec{v} .]

4.4 Projections

You may have seen the projection of a vector onto another vector in a previous course, and we have seen hints of it in the proof of the Cauchy–Schwarz inequality. We now shift our attention to fleshing out the details in full.

Let's think about what projection looks like in \mathbb{R}^2 . Suppose \vec{v} and \vec{w} are two non-zero vectors in \mathbb{R}^2 , and we wish to project \vec{v} onto \vec{w} . We can think of this as shining a light perpendicular to \vec{w} , and drawing a vector $\text{proj}_{\vec{w}}(\vec{v})$ representing the shadow of \vec{v} .



Suppose θ is the angle between \vec{v} and \vec{w} . Then by drawing out a right triangle, we see the length of the projection must be $\|\vec{v}\| \cos \theta$. The direction we wish the vector to go in is the direction \vec{w} is pointing, so we can obtain the projection by scalar multiplying the unit vector in the direction of \vec{w} by $\|\vec{v}\| \cos \theta$. Since $\cos \theta = \frac{\vec{v} \cdot \vec{w}}{\|\vec{v}\| \|\vec{w}\|}$, this gives the projection as

$$\text{proj}_{\vec{w}}(\vec{v}) = \|\vec{v}\| \cos \theta \frac{1}{\|\vec{w}\|} \vec{w} = \frac{\vec{v} \cdot \vec{w}}{\|\vec{w}\|^2} \vec{w}.$$

We will use this as motivation for defining projections in inner product spaces in general.

Definition 4.4.1

Projection onto a Vector, $\text{proj}_{\vec{w}}$, Perpendicular Vector with respect to a Vector, $\text{perp}_{\vec{w}}$

Let V be an inner product space, and let $\vec{w}, \vec{v} \in V$ with $\vec{w} \neq \vec{0}$. The **projection of \vec{v} onto \vec{w}** is defined to be the vector

$$\text{proj}_{\vec{w}}(\vec{v}) = \frac{\langle \vec{v}, \vec{w} \rangle}{\|\vec{w}\|^2} \vec{w}.$$

We also define the **perpendicular vector of \vec{v} with respect to \vec{w}** by

$$\text{perp}_{\vec{w}}(\vec{v}) = \vec{v} - \frac{\langle \vec{v}, \vec{w} \rangle}{\|\vec{w}\|^2} \vec{w}.$$

We already saw in Lemma 4.2.4 that $\text{perp}_{\vec{w}}(\vec{v})$ is orthogonal to \vec{w} , which is what we'd expect to be true if these definitions imitate the situation in \mathbb{R}^2 . Furthermore, notice that according to the definition

$$\vec{v} = \text{proj}_{\vec{w}}(\vec{v}) + \text{perp}_{\vec{w}}(\vec{v})$$

for all $\vec{v}, \vec{w} \in V$ with $\vec{w} \neq \vec{0}$.

Example 4.4.2

Let $\vec{v} = \begin{bmatrix} 4 \\ 1+i \\ 2 \end{bmatrix} \in \mathbb{C}^3$. Let $\vec{w}_1 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$ and $\vec{w}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$. With respect to the standard inner product we have

$$\begin{aligned} \text{proj}_{\vec{w}_1}(\vec{v}) &= \frac{2}{1} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 2 \end{bmatrix} \\ \text{proj}_{\vec{w}_2}(\vec{v}) &= \frac{i+1}{1} \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 1+i \\ 0 \end{bmatrix} \end{aligned}$$


These computations are what we'd expect if we are to project onto the z and y axes of \mathbb{C}^3 .

Another way to think about a projection is as finding the closest vector on a subspace to a given vector. More specifically, the projection of \vec{v} onto \vec{w} is the closest vector in the subspace spanned by $\{\vec{w}\}$ to \vec{v} . (See Proposition 4.6.10 below.) Remember—since we're in an inner product space, we have a notion of length and distance, so asking for the *closest* vector makes sense.

With this in mind, what we're really doing when we're projecting onto a vector is we're projecting onto the one-dimensional subspace spanned by that vector. It's natural to now ask how we can project onto a general subspace. Let's look at an example.

Example 4.4.3

Consider \mathbb{R}^3 with the dot product. Let $\vec{v} = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$, $\vec{e}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$, and $\vec{e}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$. The subspace $W = \text{Span}(\{\vec{e}_1, \vec{e}_2\})$ is the xy -plane in \mathbb{R}^3 . So the projection of \vec{v} onto W , let's call it $\text{proj}_W(\vec{v})$, should be the vector $\begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}$.



If we compute the projections of \vec{v} onto \vec{e}_1 and \vec{e}_2 , we find that

$$\text{proj}_{\vec{e}_1}(\vec{v}) = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \quad \text{and} \quad \text{proj}_{\vec{e}_2}(\vec{v}) = \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix}.$$

The sum of these two vector projections is what we believe $\text{proj}_W(\vec{v})$ should be. That is, it appears that

$$\text{proj}_W(\vec{v}) = \text{proj}_{\vec{e}_1}(\vec{v}) + \text{proj}_{\vec{e}_2}(\vec{v}).$$

From this example, it may be tempting to guess the following: The projection of a vector \vec{v} onto a subspace W is simply obtained by choosing a basis for W , projecting \vec{v} onto each basis vector, and summing up the resulting vectors.

Unfortunately this does not work all the time, but we will see that it does work when we have an orthogonal basis for W . So it may be tempting to refine our initial guess and define the projection of \vec{v} onto W as follows. Let $\{\vec{w}_1, \dots, \vec{w}_k\}$ be an orthogonal basis for W . Then

$$\text{proj}_W(\vec{v}) = \text{proj}_{\vec{w}_1}(\vec{v}) + \dots + \text{proj}_{\vec{w}_k}(\vec{v}).$$

For this to be a usable definition, there are two issues that must be addressed.

1. Does W even have an orthogonal basis? If so, how do we find one?
2. Does the above definition of $\text{proj}_W(\vec{v})$ depend on the chosen orthogonal basis for W ? That is, if $\{\vec{u}_1, \dots, \vec{u}_k\}$ is another orthogonal basis for W , how can we be sure that

$$\text{proj}_{\vec{w}_1}(\vec{v}) + \dots + \text{proj}_{\vec{w}_k}(\vec{v}) = \text{proj}_{\vec{u}_1}(\vec{v}) + \dots + \text{proj}_{\vec{u}_k}(\vec{v})$$

so that our definition of $\text{proj}_W(\vec{v})$ is well-defined?

We will address both these issues in the next two sections.

4.5 The Gram–Schmidt Orthogonalization Procedure

In this section we will show that every finite-dimensional inner product space has an orthonormal basis. In fact, we will describe a procedure that allows us to take an arbitrary basis and create an orthonormal basis from it. Let's see how this works in an example.

Example 4.5.1

In \mathbb{R}^3 equipped with the dot product, suppose we have

$$\vec{v}_1 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \quad \vec{v}_2 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \quad \text{and} \quad \vec{v}_3 = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}.$$

It turns out that $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ is a basis for \mathbb{R}^3 . However, $\vec{v}_1 \cdot \vec{v}_2 = 1$ so it is not an orthogonal set (or an orthonormal one for that matter). We will now create an orthogonal basis starting from this one, and then scale the vectors to obtain an orthonormal basis.

Let's first create an orthogonal set $\{\vec{w}_1, \vec{w}_2, \vec{w}_3\}$. We may as well start with

$$\vec{w}_1 = \vec{v}_1.$$

Now, whatever we make \vec{w}_2 , it needs to be orthogonal to \vec{w}_1 . We have proved earlier that $\text{perp}_{\vec{w}_1}(\vec{v}_2)$ is orthogonal to \vec{w}_1 , so let's use that. We have

$$\begin{aligned} \text{perp}_{\vec{w}_1}(\vec{v}_2) &= \vec{v}_2 - \frac{\langle \vec{v}_2, \vec{w}_1 \rangle}{\|\vec{w}_1\|^2} \vec{w}_1 \\ &= \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} \frac{1}{2} \\ -\frac{1}{2} \\ 1 \end{bmatrix}. \end{aligned}$$

So that we don't have to deal with fractions, set

$$\vec{w}_2 = \begin{bmatrix} 1 \\ -1 \\ 2 \end{bmatrix}.$$

So far we've obtained two orthogonal vectors \vec{w}_1 and \vec{w}_2 . To create \vec{w}_3 , we do the same thing except we will want to take

$$\text{perp}_W(\vec{v}_3) = \vec{v}_3 - \text{proj}_W(\vec{v}_3), \quad \text{where } W = \text{Span}(\{\vec{w}_1, \vec{w}_2\}),$$

which will give us a vector orthogonal to both \vec{w}_1 and \vec{w}_2 . Of course, the problem here is that we haven't formally defined $\text{perp}_W(\vec{v}_3)$ or $\text{proj}_W(\vec{v}_3)$, but based on Example 4.4.3 and the discussion that followed it, we suspect that

$$\text{perp}_W(\vec{v}_3) = \vec{v}_3 - \text{proj}_{\vec{w}_1} \vec{v}_3 - \text{proj}_{\vec{w}_2} \vec{v}_3.$$

Let's just take \vec{w}_3 to be this vector! That is, let

$$\vec{w}_3 = \vec{v}_3 - \text{proj}_{\vec{w}_1} \vec{v}_3 - \text{proj}_{\vec{w}_2} \vec{v}_3$$

$$\begin{aligned}
&= \vec{v}_3 - \frac{\langle \vec{v}_3, \vec{w}_1 \rangle}{\|\vec{w}_1\|^2} \vec{w}_1 - \frac{\langle \vec{v}_3, \vec{w}_2 \rangle}{\|\vec{w}_2\|^2} \vec{w}_2 \\
&= \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} - \frac{1}{6} \begin{bmatrix} 1 \\ -1 \\ 2 \end{bmatrix} \\
&= \begin{bmatrix} -\frac{2}{3} \\ \frac{2}{3} \\ \frac{2}{3} \end{bmatrix}.
\end{aligned}$$

Again, to make our lives easier, set

$$\vec{w}_3 = \begin{bmatrix} -2 \\ 2 \\ 2 \end{bmatrix}.$$

We have to confirm that \vec{w}_3 behaves as expected and is in fact orthogonal to both \vec{w}_1 and \vec{w}_2 . Indeed, this is the case! We have

$$\vec{w}_1 \cdot \vec{w}_3 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} -2 \\ 2 \\ 2 \end{bmatrix} = -2 + 2 = 0$$

and

$$\vec{w}_2 \cdot \vec{w}_3 = \begin{bmatrix} 1 \\ -1 \\ 2 \end{bmatrix} \cdot \begin{bmatrix} -2 \\ 2 \\ 2 \end{bmatrix} = -2 - 2 + 4 = 0.$$

Now we have three vectors that are orthogonal to each other, so $\mathcal{B} = \{\vec{w}_1, \vec{w}_2, \vec{w}_3\}$ is an orthogonal set. Since $\vec{0} \notin \mathcal{B}$, it follows that \mathcal{B} linearly independent and thus a basis for \mathbb{R}^3 .

To create an orthonormal basis for \mathbb{R}^3 we simply normalize the vectors in S . Therefore

$$\mathcal{C} = \left\{ \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \frac{1}{\sqrt{6}} \begin{bmatrix} 1 \\ -1 \\ 2 \end{bmatrix}, \frac{1}{\sqrt{3}} \begin{bmatrix} -1 \\ 1 \\ 1 \end{bmatrix} \right\}$$

is an orthonormal basis of \mathbb{R}^3 with respect to the dot product.

The method we illustrated in this example works in general, and it is called the **Gram–Schmidt orthogonalization procedure**. Here it is in detail.

ALGORITHM (Gram–Schmidt Orthogonalization Procedure)

Let V be an inner product space with basis $\{\vec{v}_1, \dots, \vec{v}_n\}$. To obtain an orthogonal basis for V , define $\vec{w}_1, \dots, \vec{w}_n$ as follows:

$$\begin{aligned}
\vec{w}_1 &= \vec{v}_1 \\
\vec{w}_2 &= \vec{v}_2 - \frac{\langle \vec{v}_2, \vec{w}_1 \rangle}{\|\vec{w}_1\|^2} \vec{w}_1
\end{aligned}$$

$$\begin{aligned}\vec{w}_3 &= \vec{v}_3 - \frac{\langle \vec{v}_3, \vec{w}_1 \rangle}{\|\vec{w}_1\|^2} \vec{w}_1 - \frac{\langle \vec{v}_3, \vec{w}_2 \rangle}{\|\vec{w}_2\|^2} \vec{w}_2 \\ &\vdots \\ \vec{w}_n &= \vec{v}_n - \frac{\langle \vec{v}_n, \vec{w}_1 \rangle}{\|\vec{w}_1\|^2} \vec{w}_1 - \dots - \frac{\langle \vec{v}_n, \vec{w}_{n-1} \rangle}{\|\vec{w}_{n-1}\|^2} \vec{w}_{n-1}.\end{aligned}$$

Then $\mathcal{B} = \{\vec{w}_1, \dots, \vec{w}_n\}$ is an orthogonal basis for V .

To obtain an orthonormal basis, set $\vec{u}_i = \widehat{w}_i = \frac{1}{\|\vec{w}_i\|} \vec{w}_i$ for all $i = 1, \dots, n$, and take $\mathcal{C} = \{\vec{u}_1, \dots, \vec{u}_n\}$.

To prove that this procedure actually works, we need to ensure that at each step we actually get a non-zero vector. This amounts to proving the following statement.

Exercise 50

Let $\{\vec{v}_1, \dots, \vec{v}_n\}$ be a basis for an inner product space V . Let $\{\vec{w}_1, \dots, \vec{w}_i\}$ be the first i vectors obtained from the Gram–Schmidt orthogonalization procedure. Prove $\text{Span}(\{\vec{v}_1, \dots, \vec{v}_i\}) = \text{Span}(\{\vec{w}_1, \dots, \vec{w}_i\})$. Use this to prove that the resulting basis $\{\vec{w}_1, \dots, \vec{w}_n\}$ is an orthogonal basis for V .

REMARK

You may have noticed that the algorithm described above does not exactly match the procedure that we followed in Example 4.5.1. Indeed, notice that at each stage of the Gram–Schmidt procedure we can pick $c\vec{w}_i$ instead of \vec{w}_i for some non-zero scalar c , which is exactly what we did in the example. This little modification of the algorithm can make calculations a bit nicer, and is especially useful when doing calculations by hand. We recommend that you take time to reflect as to why it works.

As a consequence of the Gram–Schmidt procedure, we get the following corollary.

Corollary 4.5.2

Every finite-dimensional inner product space has an orthonormal basis.

If $V = \{\vec{0}\}$ is the zero vector space, then we agree to consider its basis (the empty set) as being an orthogonal basis. This makes sense because the condition for the empty set to be an orthogonal set is vacuously true.

Example 4.5.3

Let's apply the Gram–Schmidt process to find an orthogonal basis for $P_3(\mathbb{R})$ with respect to the inner product $\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx$. Starting from the standard basis $\{1, x, x^2, x^3\}$, we take $\vec{w}_1 = 1$, and then

$$\vec{w}_2 = x - \frac{\langle x, 1 \rangle}{\|1\|^2} 1 = x - \frac{0}{2} 1 = x.$$

(Recall that we'd already seen in Example 4.2.2 that $1 \perp x$, so this is a promising start.) Next,

$$\vec{w}_3 = x^2 - \frac{\langle x^2, 1 \rangle}{\|1\|^2} 1 - \frac{\langle x^2, x \rangle}{\|x\|^2} x = x^2 - \frac{(2/3)}{2} 1 - \frac{0}{(2/3)} x = x^2 - \frac{1}{3}.$$

To avoid having to deal with fractions, let's multiply by 3 and take $\vec{w}_3 = 3x^2 - 1$ instead. Finally,

$$\vec{w}_4 = x^3 - \frac{\langle x^3, 1 \rangle}{\|1\|^2} 1 - \frac{\langle x^3, x \rangle}{\|x\|^2} x - \frac{\langle x^3, 3x^2 - 1 \rangle}{\|3x^2 - 1\|^2} (3x^2 - 1) = x^3 - \frac{3}{5}x.$$

Let's multiply by 5 and take $\vec{w}_4 = 5x^3 - 3x$ instead.

So now we have an orthogonal basis $\{1, x, 3x^2 - 1, 5x^3 - 3x\}$ for $\mathcal{P}_3(\mathbb{R})$ with respect to $\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx$. If we normalize this, we obtain the orthonormal basis

$$\mathcal{C} = \left\{ \frac{1}{\sqrt{2}}, \sqrt{\frac{3}{2}}x, \sqrt{\frac{5}{8}}(3x^2 - 1), \sqrt{\frac{7}{8}}(5x^3 - 3x) \right\}$$

which is exactly the same as the basis we'd seen in Example 4.3.6.

Exercise 51

Find an orthonormal basis for $\mathcal{P}_4(\mathbb{R})$ with respect to $\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx$.

Example 4.5.4

Let's find an orthogonal basis for the subspace

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

of \mathbb{R}^3 with respect to the dot product.

Our first step should be to find a basis for W , which means we must eliminate any redundant vectors in the given spanning set. (We know for sure that at least one of them must be redundant, since W cannot be a 4-dimensional subspace of \mathbb{R}^3 !) But let's *not* do this, and let's see what happens if we apply the Gram–Schmidt procedure to the given spanning set of W .

So take $\vec{w}_1 = [1 \ 1 \ 1]^T$, and then

$$\vec{w}_2 = \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} - \frac{\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}}{\left\| \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right\|^2} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}.$$

Next, we are supposed to take

$$\vec{w}_3 = \begin{bmatrix} 2 \\ 0 \\ 1 \end{bmatrix} - \frac{\begin{bmatrix} 2 \\ 0 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}}{\left\| \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right\|^2} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} - \frac{\begin{bmatrix} 2 \\ 0 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}}{\left\| \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} \right\|^2} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix},$$

but we cannot have the zero vector in a basis! Did the Gram–Schmidt procedure fail? No, of course not. In fact, the procedure was able to detect the linear dependence present in the spanning set of W . Indeed, the above shows that the vector $[2 \ 0 \ 1]^T$ is a linear combination of $[1 \ 1 \ 1]^T$ and $[1 \ -1 \ 0]^T$. So that means we toss it aside and move on to the next vector in the spanning set. Hence, we now take

$$\vec{w}_3 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} - \frac{\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}}{\left\| \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right\|^2} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} - \frac{\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}}{\left\| \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} \right\|^2} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} = \frac{1}{6} \begin{bmatrix} -1 \\ -1 \\ 2 \end{bmatrix}.$$

We can scale this vector by 6 and then arrive at the following orthogonal basis for W :

$$\left\{ \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \\ 2 \end{bmatrix} \right\}.$$

What occurred in the previous example will always occur. The Gram–Schmidt process will automatically locate linear dependency in a given spanning set. This is formalized in the next exercise.

Exercise 52

Let $S = \{\vec{v}_1, \dots, \vec{v}_k\}$ be linearly independent vectors in an inner product space, and let $\vec{w}_1, \dots, \vec{w}_k$ be the vectors produced by the Gram–Schmidt procedure applied to S . Show that if $\vec{v}_{k+1} \in \text{Span}(S)$, and if we take

$$\vec{w}_{k+1} = \vec{v}_{k+1} - \frac{\langle \vec{v}_{k+1}, \vec{w}_1 \rangle}{\|\vec{w}_1\|^2} \vec{w}_1 - \dots - \frac{\langle \vec{v}_{k+1}, \vec{w}_k \rangle}{\|\vec{w}_k\|^2} \vec{w}_k$$

then necessarily $\vec{w}_{k+1} = 0$.

4.6 Projection onto a Subspace and Orthogonal Complements


Let's now return to the problem of determining the projection of a vector onto a subspace of an inner product space, which we discussed at the end of Section 4.4. This is a problem that has many important practical and theoretical implications. For instance, if the vectors

in our inner product space somehow encode the data of an MRI image, one might have a subspace of images that are known to be associated with a certain condition or disease. Then it would be highly desirable to determine how close a given image is to this subspace, and this is something that can be easily done with projections. We will see another, very widely-used, application in Section 4.7.

Here is the basic problem. Let W be a subspace of an inner product space V and let $\vec{v} \in V$. We would like to define vectors $\text{proj}_W(\vec{v})$ and $\text{perp}_W(\vec{v})$ so that (i) $\text{proj}_W(\vec{v})$ is in W , (ii) $\text{perp}_W(\vec{v})$ is orthogonal to W , and (iii) \vec{v} can be expressed as

$$\vec{v} = \text{proj}_W(\vec{v}) + \text{perp}_W(\vec{v}).$$

In a sense, we would like $\text{proj}_W(\vec{v})$ to be the vector in W that is closest to \vec{v} . See the diagram below.



We saw in Example 4.4.3 how to do this if W was the xy -plane in \mathbb{R}^3 . The case of a general plane W in \mathbb{R}^3 can be made to look exactly the same: simply choose a new set of *orthogonal* axes for \mathbb{R}^3 , two of which lie in W , and the third of which is orthogonal to W . Then W will effectively look like “the horizontal plane” in this coordinate system, and so we can easily project onto it.

The case of an arbitrary subspace of an inner product space can be handled in the same fashion. It all boils down to being able to pick out the correct orthogonal basis. Here is the key result.

Proposition 4.6.1

Let W be a subspace of a finite-dimensional inner product space V . Then we can find an orthogonal basis $\{\vec{w}_1, \dots, \vec{w}_k, \vec{n}_1, \dots, \vec{n}_l\}$ for V such that $\{\vec{w}_1, \dots, \vec{w}_k\}$ is an orthogonal basis for W .

Proof: Let $\{\vec{u}_1, \dots, \vec{u}_k\}$ be a basis for W . Apply the Gram–Schmidt procedure to it to obtain an orthogonal basis $\{\vec{w}_1, \dots, \vec{w}_k\}$ for W and then extend this to a basis $\{\vec{w}_1, \dots, \vec{w}_k, \vec{v}_1, \dots, \vec{v}_l\}$ for V . Apply the Gram–Schmidt procedure to this basis. The first k vectors $\vec{w}_1, \dots, \vec{w}_k$ will be unchanged (why?), and $\vec{v}_1, \dots, \vec{v}_l$ will be refined to give us our desired $\vec{n}_1, \dots, \vec{n}_l$. \square

Notice that the proof of Proposition 4.6.1 shows that we can extend any orthogonal basis for the subspace W to an orthogonal basis of the whole space V . In fact, more is true: we can also extend any *orthonormal* basis for the subspace W to an *orthonormal* basis of the whole space V . We strongly encourage you to think how this can be done!

Example 4.6.2

Let W be the plane in \mathbb{R}^3 with scalar equation $x + y + 2z = 0$. (Let's also note in passing that any vector orthogonal to this plane will have to be a scalar multiple of $[1 \ 1 \ 2]^T$.) A basis for W will consist of any two linear independent vectors lying in this plane, e.g.,

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

If we apply the Gram–Schmidt procedure to $\{\vec{u}_1, \vec{u}_2\}$ we obtain $\vec{w}_1 = \vec{u}_1$ and

$$\vec{w}_2 = \begin{bmatrix} 2 \\ 0 \\ -1 \end{bmatrix} - \frac{\begin{bmatrix} 2 \\ 0 \\ -1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}}{\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix}.$$

This gives us the orthogonal basis $\{\vec{w}_1, \vec{w}_2\}$ for W . To extend this to a basis for \mathbb{R}^3 , we need one more vector. We can either find this vector by inspection, or (and this is what we'll do) we can apply the Gram–Schmidt procedure to $\{\vec{w}_1, \vec{w}_2, \vec{e}_1, \vec{e}_2, \vec{e}_3\}$. Indeed, the procedure will return the first two vectors unchanged, and then will eliminate any linear dependencies presented by adding in $\vec{e}_1, \vec{e}_2, \vec{e}_3$, while simultaneously producing for us an orthogonal basis for \mathbb{R}^3 . We obtain the vector

$$\vec{n} = \vec{e}_1 - \frac{\vec{e}_1 \cdot \vec{w}_1}{\vec{w}_1 \cdot \vec{w}_1} \vec{w}_1 - \frac{\vec{e}_1 \cdot \vec{w}_2}{\vec{w}_2 \cdot \vec{w}_2} \vec{w}_2 = \frac{1}{6} \begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix}.$$

And indeed, this is a normal vector for the plane W .

Now that we have an orthogonal basis $\{\vec{w}_1, \vec{w}_2, \vec{n}\}$ for \mathbb{R}^3 , Proposition 4.3.13 tells us that every vector $\vec{v} \in \mathbb{R}^3$ can be expressed as

$$\vec{v} = \text{proj}_{\vec{w}_1}(\vec{v}) + \text{proj}_{\vec{w}_2}(\vec{v}) + \text{proj}_{\vec{n}}(\vec{v}).$$

The sum of the first two projections above gives us the component of \vec{v} that lies in W while the third gives us the component of \vec{v} that is orthogonal to W , i.e., these should be $\text{proj}_W(\vec{v})$ and $\text{perp}_W(\vec{v})$, respectively.

We will now generalize what we did in the previous example. First, let's identify the significance of the vectors $\vec{n}_1, \dots, \vec{n}_l$ in Proposition 4.6.1: they are each orthogonal to all the vectors in W . In general, every subspace W of an inner product space V determines a “complementary” subspace that consists of all the vectors in V that are orthogonal to every vector in W . This is analogous to how a plane in \mathbb{R}^3 determines a normal line.

Definition 4.6.3**Orthogonal
Complement**

Let V be an inner product space and let $W \subseteq V$ be a subspace. The **orthogonal complement** of W is the set

$$W^\perp = \{\vec{v} \in V : \langle \vec{v}, \vec{w} \rangle = 0 \text{ for all } \vec{w} \in W\}.$$

Here are two basic properties of W^\perp whose proofs are left as an exercise.

Proposition 4.6.4

Let V be an inner product space and $W \subseteq V$ a subspace. Then:

- (a) W^\perp is a subspace of V .
- (b) $W \cap W^\perp = \{\vec{0}\}$.

Exercise 53

Prove Proposition 4.6.4.

Example 4.6.5

Consider \mathbb{R}^3 with the dot product and let $W = \text{Span} \left(\left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \right\} \right)$. We would expect $U = \text{Span} \left(\left\{ \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\} \right)$ to be equal to W^\perp . Let's prove this.

Proof: We want to show $U = W^\perp$. Let $\vec{u} \in U$. Then $\vec{u} = \begin{bmatrix} 0 \\ 0 \\ t \end{bmatrix}$ for some $t \in \mathbb{R}$. Let \vec{w} be an arbitrary vector in W , so $\vec{w} = \begin{bmatrix} a \\ b \\ 0 \end{bmatrix}$ for some $a, b \in \mathbb{R}$. Then $\vec{u} \cdot \vec{w} = 0$ so $\vec{u} \in W^\perp$ and therefore $U \subseteq W^\perp$. Conversely, suppose $\vec{u} = \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} \in W^\perp$. Then since $\vec{u} \cdot \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = 0$, we must have $u_1 = 0$. Similarly we must have $u_2 = 0$. Therefore $\vec{u} \in U$ and $W^\perp \subseteq U$. We can now conclude $W^\perp = U$, completing the proof. \square

Example 4.6.6

We saw earlier that

$$\left\{ 1, x, \frac{3}{2}x^2 - \frac{1}{2}, \frac{5}{2}x^3 - \frac{3}{2}x \right\}$$

is an orthogonal basis for $\mathcal{P}_3(\mathbb{R})$ with respect to the inner product $\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx$. Therefore if $W = \text{Span} \left(\left\{ 1, \frac{3}{2}x^2 - \frac{1}{2} \right\} \right)$ then $W^\perp = \text{Span} \left(\left\{ x, \frac{5}{2}x^3 - \frac{3}{2}x \right\} \right)$. We can prove this as we did in the previous example.

However, if we stop and think for a moment, we might suspect that there is a general result along the lines of: If $\{\vec{v}_1, \dots, \vec{v}_n\}$ is an orthogonal basis for V , and if $W = \text{Span}\{\vec{v}_1, \dots, \vec{v}_k\}$, then $W^\perp = \text{Span}\{\vec{v}_{k+1}, \dots, \vec{v}_n\}$. This is indeed true and the proof conceptually is identical to what we did in the previous example.

Exercise 54 Prove the claim in the final paragraph of Example 4.6.6.

There is a sort-of converse to this exercise. We record it together with a few other basic properties of the orthogonal complement in the next proposition. Part (b) of this proposition can be viewed as a more refined version of Proposition 4.6.1.

Proposition 4.6.7 Let V be a finite-dimensional inner product space and $W \subseteq V$ a subspace. Then:

- (a) If $\{\vec{w}_1, \dots, \vec{w}_k\}$ is a spanning set for W then $\vec{v} \in W^\perp$ if and only if $\langle \vec{v}, \vec{w}_i \rangle = 0$ for all $i = 1, \dots, k$.
- (b) If $\mathcal{B} = \{\vec{w}_1, \dots, \vec{w}_k\}$ is an orthogonal basis for W , then there exists an orthogonal basis $\mathcal{C} = \{\vec{n}_1, \dots, \vec{n}_l\}$ for W^\perp such that $\mathcal{B} \cup \mathcal{C} = \{\vec{w}_1, \dots, \vec{w}_k, \vec{n}_1, \dots, \vec{n}_l\}$ is an orthogonal basis for V .
- (c) $\dim(V) = \dim(W) + \dim(W^\perp)$.
- (d) $(W^\perp)^\perp = W$.

Proof: (a) If $\vec{v} \in W^\perp$, then $\langle \vec{v}, \vec{w} \rangle = 0$ for all $\vec{w} \in W$, so in particular for all \vec{w}_i . Conversely, assume that $\langle \vec{v}, \vec{w}_i \rangle = 0$ for all \vec{w}_i . To show that $\vec{v} \in W^\perp$, we must prove that $\langle \vec{v}, \vec{w} \rangle = 0$ for all $\vec{w} \in W$. Given $\vec{w} \in W$, we can write it as $\vec{w} = \sum_{i=1}^k a_i \vec{w}_i$. Then

$$\langle \vec{v}, \vec{w} \rangle = \left\langle \vec{v}, \sum_{i=1}^k a_i \vec{w}_i \right\rangle = \sum_{i=1}^k \overline{a_i} \langle \vec{v}, \vec{w}_i \rangle = \sum_{i=1}^k \overline{a_i} 0 = 0,$$

as desired.

- (b) The proof of Proposition 4.6.1 shows that we can find an orthogonal set of vectors $\mathcal{C} = \{\vec{n}_1, \dots, \vec{n}_l\}$ such that $\mathcal{B} \cup \mathcal{C}$ is a orthogonal basis for V . Since $\vec{n}_1, \dots, \vec{n}_l$ are each orthogonal to all the vectors in \mathcal{B} , they must all lie in W^\perp , by part (a). It remains to show that $\mathcal{C} = \{\vec{v}_1, \dots, \vec{v}_m\}$ is a basis for W^\perp . Since \mathcal{C} is linearly independent (being a subset of a basis), it suffices to show that \mathcal{C} spans W^\perp .

So suppose $\vec{v} \in W^\perp$. Then since $\mathcal{B} \cup \mathcal{C}$ is a basis for V , we can write \vec{v} as

$$\vec{v} = \sum_{i=1}^k a_i \vec{w}_i + \sum_{j=1}^m b_j \vec{n}_j.$$

Taking inner product of both sides with \vec{w}_l , we obtain

$$\langle \vec{v}, \vec{w}_l \rangle = \sum_{i=1}^k a_i \langle \vec{w}_i, \vec{w}_l \rangle + \sum_{j=1}^m b_j \langle \vec{n}_j, \vec{w}_l \rangle$$

$$0 = a_l \langle \vec{w}_l, \vec{w}_l \rangle + \sum_{j=1}^m b_j 0$$

since $\vec{v}, \vec{n}_j \in W^\perp$ for all j and since $\vec{w}_i \perp \vec{w}_l$ for $i \neq l$ (because \mathcal{B} is orthogonal). This leaves us with the equation $a_l \langle \vec{w}_l, \vec{w}_l \rangle = 0$ which shows that $a_l = 0$ for all l (since $\vec{w}_l \neq \vec{0}$), hence

$$\vec{v} = \sum_{j=1}^m b_j \vec{n}_j.$$

This proves that \mathcal{C} is a spanning set of W^\perp , as desired.

(c) This follows from part (b).

(d) If we apply part (c) to the subspace W^\perp (instead of W), we obtain

$$\dim(V) = \dim(W^\perp) + \dim((W^\perp)^\perp).$$

On the other hand, by part (c) applied to W , we have

$$\dim(V) = \dim(W) + \dim(W^\perp).$$

Equating both expressions, we obtain $\dim(W) = \dim((W^\perp)^\perp)$. However, W is a subspace of $(W^\perp)^\perp$, since all the vectors in W are orthogonal to all the vectors in W^\perp . This forces $W = (W^\perp)^\perp$ by Theorem 1.3.18.

□

REMARK

Part (a) of the previous proposition remains true if V were infinite-dimensional, but the remaining parts fail. While it is always true that $W \subseteq (W^\perp)^\perp$, there are examples where $W \neq (W^\perp)^\perp$ in an infinite-dimensional inner product space.

We're going to show that every vector $\vec{v} \in V$ can be decomposed as the sum of a vector in W and a vector in W^\perp . The idea here is that we can write \vec{v} as the sum of $\text{proj}_W(\vec{v})$ and $\text{perp}_W(\vec{v})$, but we don't want to use this terminology just yet (since we're going to use this result to *define* proj_W and perp_W).

Theorem 4.6.8 (Orthogonal Decomposition)

Let W be a subspace of a finite-dimensional inner product space V . Then every $\vec{v} \in V$ can be written as $\vec{v} = \vec{p} + \vec{r}$ where $\vec{p} \in W$ and $\vec{r} \in W^\perp$ are uniquely determined by \vec{v} .

Moreover, if $\mathcal{B} = \{\vec{w}_1, \dots, \vec{w}_k\}$ is an orthogonal basis for W , then \vec{p} is given by

$$\vec{p} = \sum_{i=1}^k \text{proj}_{\vec{w}_i}(\vec{v}).$$

Proof: Let $\{\vec{w}_1, \dots, \vec{w}_k, \vec{n}_1, \dots, \vec{n}_l\}$ be a basis for V as in Proposition 4.6.1. Then we can write \vec{v} as

$$\vec{v} = \sum_{i=1}^k a_i \vec{w}_i + \sum_{j=1}^l b_j \vec{n}_j.$$

We can take $\vec{p} = \sum_{i=1}^k a_i \vec{w}_i$ and $\vec{r} = \sum_{j=1}^l b_j \vec{n}_j$.

To prove uniqueness, suppose that we have $\vec{v} = \vec{p}_1 + \vec{r}_1$ with $\vec{p}_1 \in W$ and $\vec{r}_1 \in W^\perp$. Then

$$\vec{0} = \vec{v} - \vec{v} = (\vec{p} + \vec{r}) - (\vec{p}_1 + \vec{r}_1) = (\vec{p} - \vec{p}_1) + (\vec{r} - \vec{r}_1)$$

hence $\vec{p} - \vec{p}_1 = -(\vec{r} - \vec{r}_1)$. Now, $\vec{p} - \vec{p}_1 \in W$ and $-(\vec{r} - \vec{r}_1) \in W^\perp$ since W and W^\perp are subspaces. So these identical vectors are in $W \cap W^\perp$ hence they are both $\vec{0}$ by Proposition 4.6.4(b). Thus, $\vec{p} = \vec{p}_1$ and $\vec{r} = \vec{r}_1$.

If we apply Proposition 4.3.13 to \vec{p} , we find that $a_i = \frac{\langle \vec{v}, \vec{w}_i \rangle}{\|\vec{w}_i\|^2}$, and so

$$\vec{p} = \sum_{i=1}^k a_i \vec{w}_i = \sum_{i=1}^k \frac{\langle \vec{v}, \vec{w}_i \rangle}{\|\vec{w}_i\|^2} \vec{w}_i = \sum_{i=1}^k \text{proj}_{\vec{w}_i}(\vec{v}),$$

completing the proof. \square

With this result, we are now finally able to properly define the projection onto a subspace W of an inner product space V .

Definition 4.6.9

Projection onto a Subspace, proj_W , Perpendicular Vector with Respect to a Subspace, perp_W

Let V be an inner product space and let $W \subseteq V$ be a subspace. Let $\{\vec{w}_1, \dots, \vec{w}_k\}$ be an orthogonal basis for W . Let $\vec{v} \in V$. The **projection of \vec{v} onto W** and the **perpendicular vector of \vec{v} with respect to W** are defined to be

$$\text{proj}_W(\vec{v}) = \text{proj}_{\vec{w}_1}(\vec{v}) + \dots + \text{proj}_{\vec{w}_k}(\vec{v}) \quad \text{and} \quad \text{perp}_W(\vec{v}) = \vec{v} - \text{proj}_W(\vec{v}),$$

respectively.

The Orthogonal Decomposition Theorem shows that $\text{proj}_W(\vec{v})$ and hence $\text{perp}_W(\vec{v})$ are well-defined—that is, they do not depend on the choice of orthogonal basis for W ; they are, respectively, the unique vectors \vec{p} and \vec{r} associated to \vec{v} by the theorem.

The next result provides us with the alternative characterization of $\text{proj}_W(\vec{v})$ as being the unique vector in W that is *closest* to \vec{v} .


Proposition 4.6.10

Let V be a finite-dimensional inner product space, W a subspace of V and $\vec{v} \in V$. Let $\vec{p} = \text{proj}_W(\vec{v})$.

- For all $\vec{w} \in W$, $\|\vec{v} - \vec{p}\| \leq \|\vec{v} - \vec{w}\|$. That is, among all vectors in W , the vector $\vec{p} = \text{proj}_W(\vec{v})$ is the closest to \vec{v} .
- If $\|\vec{v} - \vec{p}\| = \|\vec{v} - \vec{w}\|$ for some $\vec{w} \in W$, then $\vec{w} = \vec{p}$. That is, among all vectors in W , the vector that is closest to \vec{v} is unique.

Proof: (a) Notice that $\vec{v} - \vec{w} = (\vec{v} - \vec{p}) + (\vec{p} - \vec{w})$. Now, $\vec{r} = \vec{v} - \vec{p}$ is in W^\perp and $\vec{p} - \vec{w}$ is in W since W is a subspace. So $\vec{v} - \vec{p} \perp \vec{p} - \vec{w}$, and thus the Pythagorean theorem yields

$$\|\vec{v} - \vec{w}\|^2 = \|\vec{v} - \vec{p}\|^2 + \|\vec{p} - \vec{w}\|^2 \geq \|\vec{v} - \vec{p}\|^2.$$



(b) The only way we have equality in (a) is if $\|\vec{p} - \vec{w}\|^2 = 0$, or equivalently, if $\vec{w} = \vec{p}$, as required. □

It is important to keep in mind that our definition of $\text{proj}_W(\vec{v})$ requires an *orthogonal* basis for W .

Exercise 55

Let $\{\vec{z}_1, \dots, \vec{z}_k\}$ be a basis for W that is *not* orthogonal. Show that there exists a vector $\vec{v} \in V$ such that $\text{proj}_{\vec{z}_1}(\vec{v}) + \dots + \text{proj}_{\vec{z}_k}(\vec{v}) \neq \text{proj}_W(\vec{v})$.

Example 4.6.11

Consider the vector space $\mathcal{P}_3(\mathbb{R})$ with the inner product $\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx$. We know from an earlier example that $1 \perp x$, so $\{1, x\}$ is an orthogonal basis for $W = \text{Span}(\{1, x\})$. Let's find the projection of x^2 onto W .

We have

$$\text{proj}_W(x^2) = \frac{\langle x^2, 1 \rangle}{\|1\|^2} 1 + \frac{\langle x^2, x \rangle}{\|x\|^2} x.$$

It can be checked that $\langle x^2, 1 \rangle = \frac{2}{3}$, $\langle x^2, x \rangle = 0$, and $\|1\|^2 = 2$. Therefore

$$\text{proj}_W(x^2) = \frac{1}{2} \frac{2}{3} 1 = \frac{1}{3}$$

and

$$\text{perp}_W(x^2) = x^2 - \frac{1}{3}.$$

This tells us that the closest vector in W to x^2 is the vector $\frac{1}{3}$. Go figure!

We close this section with a neat example of projection.

Example 4.6.12 (Fourier expansion)

Consider the vector space $V = \mathcal{C}([-\pi, \pi])$ of continuous functions $f: [-\pi, \pi] \rightarrow \mathbb{R}$ equipped with the inner product $\langle f, g \rangle = \int_{-\pi}^{\pi} f(x)g(x) dx$. You can check that $S = \{1, \sin x, \cos x\}$ is an orthogonal subset of V .

Let $W = \text{Span}(S)$. Let's find the projection of the function $f(x) = x$ onto W . We have

$$\text{proj}_W(f) = \text{proj}_1(x) + \text{proj}_{\sin x}(x) + \text{proj}_{\cos x}(x).$$

We leave it to you to show that

$$\text{proj}_1(x) = 0, \quad \text{proj}_{\sin x}(x) = 2 \quad \text{and} \quad \text{proj}_{\cos x}(x) = 0.$$

Hence $\text{proj}_W(x) = 2 \sin x$. This is the beginning of the so-called *Fourier expansion* of $f(x) = x$ on the interval $[-\pi, \pi]$, which is a way of approximating f with sinusoidal functions.

We can project onto the subspace W_n spanned by the orthogonal set

$$\{1, \cos(x), \sin(x), \cos(2x), \sin(2x), \dots, \cos(nx), \sin(nx)\}$$

to obtain more terms of the Fourier expansion of f . This is an entry point to a very rich area of mathematics that has broad applications to physics, computer science, engineering, finance, and so on. Matters become much more interesting once we allow ourselves to dabble in infinite-dimensional vector spaces, which will grant us access to infinite Fourier expansions.

Exercise 56

Let V be the inner product space from the previous example.

(a) Show that, for all $n \geq 1$, the set

$$S_n = \{1, \cos(x), \sin(x), \cos(2x), \sin(2x), \dots, \cos(nx), \sin(nx)\}$$

is an orthogonal set in V .

(b) Let $W_n = \text{Span}(S_n)$, where S_n is as in part (a). Find $\text{proj}_{W_n}(f)$, where $f \in V$ is the function $f(x) = |x|$.

(c) On the same set of axes, plot the graphs of $f(x) = |x|$ and $\text{proj}_{W_n}(f)$ for $n = 1, 3, 9$. What do you notice?

4.7 Application: Method of Least Squares

In this section we will work over $\mathbb{F} = \mathbb{R}$.

Suppose we have a system of equations expressed in matrix form as $A\vec{x} = \vec{b}$. We know that this system has a solution if and only if $\vec{b} \in \text{Col}(A)$. If $\vec{b} \notin \text{Col}(A)$, then there is a way to obtain an *approximate* solution to this system. We first define $\vec{p} = \text{proj}_{\text{Col}(A)}(\vec{b})$ to be the vector in $\text{Col}(A)$ closest to \vec{b} . Then we know that the system $A\vec{x} = \vec{p}$ has a solution.

Definition 4.7.1**Least Squares Solution**

Let $A \in M_{m \times n}(\mathbb{R})$ and $\vec{b} \in \mathbb{R}^m$. The vector $\vec{s} \in \mathbb{R}^n$ is called a **least squares solution** to $A\vec{x} = \vec{b}$ if it is a solution to the system $A\vec{x} = \vec{p}$, where $\vec{p} = \text{proj}_{\text{Col}(A)}(\vec{b})$.

We think of least squares solutions as being approximate solutions to the system $A\vec{x} = \vec{b}$. The reason for the name stems from the fact that a least squares solution $\vec{x} = \vec{s}$ minimizes the quantity $\|A\vec{x} - \vec{b}\|^2$, which is a sum of squares. Indeed, since $A\vec{x}$ represents an arbitrary element of $\text{Col}(A)$, the distance between $A\vec{x}$ and \vec{b} is minimized precisely when $A\vec{x} = \text{proj}_{\text{Col}(A)}(\vec{b})$, by Proposition 4.6.10(b) and (c).

Example 4.7.2

Let $A = \begin{bmatrix} 1 & 2 \\ 1 & 2 \end{bmatrix}$ and observe that $\text{Col}(A) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$. Let $\vec{b} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$. Then $\vec{b} \notin \text{Col}(A)$ so the equation $A\vec{x} = \vec{b}$ does not have a solution.

Let's find $\vec{p} = \text{proj}_{\text{Col}(A)}(\vec{b})$. This is simply equal to

$$\vec{p} = \text{proj}_{\begin{bmatrix} 1 \\ 1 \end{bmatrix}} \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) = \frac{\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \end{bmatrix}}{\begin{bmatrix} 1 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \end{bmatrix}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Now the equation $A\vec{x} = \vec{p}$ has a solution—in fact, it has infinitely many. The solution set is given by

$$S = \left\{ \begin{bmatrix} \frac{1}{2} \\ 0 \end{bmatrix} + t \begin{bmatrix} -2 \\ 1 \end{bmatrix} : t \in \mathbb{R} \right\}$$

as you can check. Any $\vec{s} \in S$ is a least squares solution to the system $A\vec{x} = \vec{b}$.

Let's give an alternative, and more practical, characterization of least squares solutions.

Proposition 4.7.3

Let $A \in M_{m \times n}(\mathbb{R})$ and $\vec{b} \in \mathbb{R}^m$. The vector $\vec{s} \in \mathbb{R}^n$ is a least squares solution to $A\vec{x} = \vec{b}$ if and only if it is a solution to $A^T A \vec{x} = A^T \vec{b}$.

Proof: Suppose that \vec{s} is a least squares solution to $A\vec{x} = \vec{b}$. Then $A\vec{s} = \vec{p}$ where $\vec{p} = \text{proj}_{\text{Col}(A)}(\vec{b})$. Now, $\vec{b} - \vec{p} = \text{perp}_{\text{Col}(A)}(\vec{b})$ is in $(\text{Col}(A))^\perp$ hence is orthogonal to the columns of A . So, letting \vec{a}_i denote the i^{th} column of A , we have

$$\vec{a}_i \cdot (\vec{b} - \vec{p}) = 0 \iff \vec{a}_i^T (\vec{b} - \vec{p}) = 0$$

for all i . That is, $A^T(\vec{b} - \vec{p}) = \vec{0}$, or equivalently $A^T(\vec{b} - A\vec{s}) = \vec{0}$. Hence $A^T \vec{b} = A^T A \vec{s}$, as required.

Conversely, if $A^T A \vec{s} = A^T \vec{b}$, then $A^T(A\vec{s} - \vec{b}) = \vec{0}$ and by the same reasoning as above, we see that $A\vec{s} - \vec{b}$ is orthogonal to the columns of A and hence must be in $\text{Col}(A)^\perp$ by Proposition 4.6.7(a). Since $A\vec{s}$ is in $\text{Col}(A)$, it follows that $A\vec{s}$ must be the projection of \vec{b} onto $\text{Col}(A)$ (see exercise below). That is, $A\vec{s} = \vec{p}$, and so \vec{s} is a least squares solution. This completes the proof. \square

We used the following observation to complete the proof of the preceding proposition.

Exercise 57

Let V be an inner product space and W a subspace. Let $\vec{v} \in V$ and $\vec{w} \in W$. Prove that $\vec{w} = \text{proj}_W(\vec{v})$ if and only if $\vec{v} - \vec{w} \in W^\perp$.

Example 4.7.4

Let $A = \begin{bmatrix} 1 & 2 \\ 1 & 2 \end{bmatrix}$ and $\vec{b} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, as in the previous example. Then

$$A^T A = \begin{bmatrix} 2 & 4 \\ 4 & 8 \end{bmatrix} \quad \text{and} \quad A^T \vec{b} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}.$$

The system

$$\begin{bmatrix} 2 & 4 \\ 4 & 8 \end{bmatrix} \vec{x} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$$


has solution set

$$S = \left\{ \begin{bmatrix} \frac{1}{2} \\ 0 \end{bmatrix} + t \begin{bmatrix} -2 \\ 1 \end{bmatrix} : t \in \mathbb{R} \right\}$$

which coincides with what we found in Example 4.7.2 above, confirming Proposition 4.7.3.


4.7.1 Least Squares Curve Fitting

Suppose we're doing a super-serious study, and we've gathered a collection of data which is looking for some kind of relationship between "self-perceived karaoke ability" and "alcohol consumed." The data we've collected looks like this when plotted:



Our goal is to model this data by some quadratic equation $y = a + bx + cx^2$ where y is the perceived karaoke ability and x is the alcohol consumed. After all, we would expect this to occur in reality: a person while sober thinks they're quite good, after a couple of drinks is aware they will be slurring a little, but after drinking more will begin to think they are god's gift to vocal performance!

So, we would like to find a quadratic that looks something like the blue curve:



Furthermore, we would like such a quadratic to make the lengths of the vertical green lines as small as possible, since the vertical green lines represent the error between our model and the experimental data.

So let's say we had the data points $(x_1, y_1), \dots, (x_n, y_n)$ which we want to approximate by $y = a + bx + cx^2$. We want this curve to best fit the points. But what do we mean by "best fit"? We want to choose the curve that minimizes the sum of the squares of the lengths of the vertical green lines:

$$(y_1 - (a + bx_1 + cx_1^2))^2 + \dots + (y_n - (a + bx_n + cx_n^2))^2.$$

(Hence the name: *least squares*!) This looks an awful lot like a norm in \mathbb{R}^n with respect to the dot product. In fact, if we let

$$\vec{y} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}, \quad \vec{1} = \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix}, \quad \vec{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}, \quad \vec{x}^2 = \begin{bmatrix} x_1^2 \\ \vdots \\ x_n^2 \end{bmatrix}$$

be vectors in \mathbb{R}^n , then minimizing the sum of the squares of the errors (the vertical green bars) is the same as minimizing

$$\left\| \vec{y} - (a\vec{1} + b\vec{x} + c\vec{x}^2) \right\|^2$$

with respect to the dot product. In other words, to find a , b , and c , we need to find the vector on the subspace $W = \text{Span}(\{\vec{1}, \vec{x}, \vec{x}^2\})$ closest to the vector \vec{y} . We know how to do this! Putting all of these observations together, we find a , b , and c by setting

$$a\vec{1} + b\vec{x} + c\vec{x}^2 = \text{proj}_W(\vec{y}).$$

If we let

$$X = [\vec{1} \ \vec{x} \ \vec{x}^2] \quad \text{and} \quad \vec{s} = \begin{bmatrix} a \\ b \\ c \end{bmatrix},$$

then the above equation becomes

$$X\vec{s} = \text{proj}_W(\vec{y}).$$

That is, we are looking for a least squares solution \vec{s} to the equation $X\vec{x} = \vec{y}$! By Proposition 4.7.3, we know that this is equivalent to finding a solution to

$$X^T X \vec{s} = X^T \vec{y}.$$

Now, *if* the 3×3 matrix $X^T X$ were invertible, then we'd be able to get our solution \vec{s} as

$$\vec{s} = (X^T X)^{-1} X^T \vec{y}.$$

Let's see this in action.

Example 4.7.5

Suppose we have the following data:

$$\begin{array}{c|cccc} x & -1 & 0 & 1 & 2 \\ \hline y & 4 & 1 & 1 & -1 \end{array}$$

Let's first try to approximate this data set by a linear equation $y = a + bx$. So we let

$$\vec{s} = \begin{bmatrix} a \\ b \end{bmatrix}, \quad X = \begin{bmatrix} 1 & -1 \\ 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{bmatrix}, \quad \text{and} \quad \vec{y} = \begin{bmatrix} 4 \\ 1 \\ 1 \\ -1 \end{bmatrix}.$$


Then

$$X^T X = \begin{bmatrix} 1 & 1 & 1 & 1 \\ -1 & 0 & 1 & 2 \end{bmatrix} \begin{bmatrix} 1 & -1 \\ 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{bmatrix} = \begin{bmatrix} 4 & 2 \\ 2 & 6 \end{bmatrix},$$

which is an invertible matrix! Thus,

$$\vec{s} = (X^T X)^{-1} X^T \vec{y} = \left(\begin{bmatrix} 4 & 2 \\ 2 & 6 \end{bmatrix} \right)^{-1} \begin{bmatrix} 5 \\ -5 \end{bmatrix} = \frac{1}{20} \begin{bmatrix} 6 & -2 \\ -2 & 4 \end{bmatrix} \begin{bmatrix} 5 \\ -5 \end{bmatrix} = \begin{bmatrix} 2 \\ -\frac{3}{2} \end{bmatrix}.$$

Therefore $y = 2 - \frac{3}{2}x$ is the line of best fit to the given data. Let's see what this line looks like.




While this is good, maybe it's not as good as we'd like! Let's see if we can do better approximating the data by a quadratic equation $y = a + bx + cx^2$. This time we have

$$\vec{s} = \begin{bmatrix} a \\ b \\ c \end{bmatrix}, \quad X = \begin{bmatrix} 1 & -1 & 1 \\ 1 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 2 & 4 \end{bmatrix}, \quad \text{and} \quad \vec{y} = \begin{bmatrix} 4 \\ 1 \\ 1 \\ -1 \end{bmatrix}.$$

Again, $X^T X$ turns out to be invertible (we'll let you check this), and so we get

$$\vec{s} = (X^T X)^{-1} X^T \vec{y} = \begin{bmatrix} \frac{7}{4} \\ -\frac{7}{4} \\ \frac{1}{4} \end{bmatrix}.$$

Therefore the quadratic curve of best fit is $y = \frac{7}{4} - \frac{7}{4}x + \frac{1}{4}x^2$. Here's a plot:



That's a little better!

In general, suppose we have some data points

$$\begin{array}{c|ccc} x & x_1 & \cdots & x_n \\ \hline y & y_1 & \cdots & y_n \end{array}$$

and we want to find the equation $y = a_0 + a_1x + \cdots + a_kx^k$ of best fit to this data. Let

$$\vec{1} = \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix}, \vec{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}, \vec{x}^2 = \begin{bmatrix} x_1^2 \\ \vdots \\ x_n^2 \end{bmatrix}, \dots, \vec{x}^k = \begin{bmatrix} x_1^k \\ \vdots \\ x_n^k \end{bmatrix}, \vec{y} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}, \quad \text{and} \quad \vec{s} = \begin{bmatrix} a_0 \\ \vdots \\ a_k \end{bmatrix}.$$

Let $X = [\vec{1} \ \vec{x} \ \cdots \ \vec{x}^k]$. If $X^T X$ is invertible, then $\vec{s} = (X^T X)^{-1} X^T \vec{y}$ gives the coefficients of the equation of best fit.

A natural question, then, is: When is the matrix $X^T X$ invertible?

Proposition 4.7.6

Let $X \in M_{m \times n}(\mathbb{R})$. Then $X^T X \in M_{n \times n}(\mathbb{R})$ is invertible if and only if the columns of X are linearly independent.

Proof: The matrix $X^T X$ will be invertible if and only if the columns of $X^T X$ are linearly independent, which will be the case if and only if $\text{Null}(X^T X) = \{\vec{0}\}$. We will show that $\text{Null}(X^T X) = \text{Null}(X)$. This will show that $X^T X$ is invertible if and only if the columns of X are linearly independent, as desired.

If $\vec{x} \in \text{Null}(X)$, then $X^T X \vec{x} = X^T \vec{0} = \vec{0}$, so $\vec{x} \in \text{Null}(X^T X)$ and $\text{Null}(X) \subseteq \text{Null}(X^T X)$. Conversely, if $\vec{x} \in \text{Null}(X^T X)$ then $\vec{x}^T X^T X \vec{x} = \vec{x}^T \vec{0} = 0$ hence $(X \vec{x})^T (X \vec{x}) = 0$, or equivalently, $(X \vec{x}) \cdot (X \vec{x}) = 0$. Thus $X \vec{x} = \vec{0}$, so $\vec{x} \in \text{Null}(X)$ and $\text{Null}(X^T X) \subseteq \text{Null}(X)$, which completes the proof. \square

The next exercise addresses the issue in the case of the best quadratic fit.

Exercise 58

$$\text{Let } X = \begin{bmatrix} 1 & x_1 & x_1^2 \\ 1 & x_2 & x_2^2 \\ \vdots & \vdots & \vdots \\ 1 & x_n & x_n^2 \end{bmatrix} \in M_{n \times 3}(\mathbb{R}).$$

- (a) Assume that $n \geq 3$. Show that $X^T X$ is invertible if and only if at least three of the x_i are distinct.
- (b) Show that $X^T X$ is never invertible if $n < 3$.

The assertions in this exercise should seem plausible. For instance, part (b) says that there is no *unique* best fitting quadratic curve (parabola) through two or fewer points.

Chapter 5

Unitary Diagonalization

5.1 The Adjoint

When working with a linear operator $L: V \rightarrow V$ on a finite-dimensional vector space, we have seen in Chapter 3 how convenient it can be to have a basis \mathcal{D} for V consisting of eigenvectors of L . If V also happens to be an inner product space, then we've learned in Chapter 4 how useful it can be to have an orthonormal basis \mathcal{B} for V . For a typical operator, its eigenvectors will not be orthogonal, and so we will have to choose to either work with eigenvectors or with orthogonal vectors.

So a natural question, then, is: Given a linear operator $L: V \rightarrow V$ on a finite-dimensional inner product space over \mathbb{F} , can we find a basis for V consisting of orthogonal eigenvectors of L ? If $\mathbb{F} = \mathbb{R}$, this will be possible if and only if L is *self-adjoint*. If $\mathbb{F} = \mathbb{C}$, this will be possible if and only if L is *normal*. We'll explain what this means in due course.

The first important definition we need to make is that of the adjoint of a matrix. To motivate it, let's take a slightly different look at the standard inner product in \mathbb{C}^n . Recall that

$$\left\langle \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}, \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} \right\rangle = v_1 \overline{w_1} + v_2 \overline{w_2} + \cdots + v_n \overline{w_n}.$$

With a bit of squinting and turning our head, we can recast this as a matrix multiplication. Here it is:

$$\begin{bmatrix} \overline{w_1} & \cdots & \overline{w_n} \end{bmatrix} \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} = \left[v_1 \overline{w_1} + v_2 \overline{w_2} + \cdots + v_n \overline{w_n} \right].$$

The important trick here, which will come up again and again, is to take a matrix, conjugate every entry, and take its transpose. This is called taking the adjoint of a matrix.

Definition 5.1.1
Adjoint, A^* , A^\dagger

If $A \in M_{m \times n}(\mathbb{F})$, the **adjoint** of A is the matrix $\overline{A^T} \in M_{n \times m}(\mathbb{F})$. It is denoted by A^* (read *A star*). (In some texts the notation A^\dagger , read *A dagger*, is also used.)

The bar denotes complex conjugation. So $A^* = \overline{A^T}$ is the matrix whose entries are the complex conjugates of the entries of A^T . For example,

$$\begin{bmatrix} 2 & 4 \\ -i & 2+i \end{bmatrix}^* = \begin{bmatrix} 2 & i \\ 4 & 2-i \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 2 & 1-i \\ 1+i & 3 \end{bmatrix}^* = \begin{bmatrix} 2 & 1-i \\ 1+i & 3 \end{bmatrix}.$$

If $\mathbb{F} = \mathbb{R}$ then $A^* = A^T$ is just the transpose of A .

Let's return to the standard inner product on \mathbb{C}^n . With

$$\vec{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \quad \text{and} \quad \vec{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix}$$

we have

$$\vec{w}^* \vec{v} = [\langle \vec{v}, \vec{w} \rangle].$$

Notice that the matrix multiplication doesn't quite spit out the inner product, but instead spits out a 1×1 matrix with the inner product as its entry.

This all works wonderfully for the standard inner product on \mathbb{C}^n , but we know there are other inner product spaces! It would be nice if we could do something similar for all of those as well—and, of course, we can!

Recall that for an arbitrary linear map $L: V \rightarrow W$ between finite-dimensional vector spaces, after choosing ordered bases for V and W , we can find a matrix that performs the mapping for us by matrix multiplication. Similarly, for an arbitrary finite-dimensional inner product space V , after choosing an orthonormal basis (which we know exists by the Gram–Schmidt procedure; see Corollary 4.5.2), we can simply take the adjoint of one of the coordinate vectors, and matrix multiplication then computes the inner product for us!

Proposition 5.1.2

Let V be a finite-dimensional inner product space, and let \mathcal{B} an orthonormal basis for V . Then for all $\vec{x}, \vec{y} \in V$,

$$[\langle \vec{x}, \vec{y} \rangle] = [\vec{y}]_{\mathcal{B}}^* [\vec{x}]_{\mathcal{B}}.$$

Proof: Let $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$ and let $\vec{x} = x_1 \vec{v}_1 + \dots + x_n \vec{v}_n$ and $\vec{y} = y_1 \vec{v}_1 + \dots + y_n \vec{v}_n$. Then $\langle \vec{x}, \vec{y} \rangle = x_1 \bar{y}_1 + \dots + x_n \bar{y}_n$. Also,

$$[\vec{x}]_{\mathcal{B}} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \quad \text{and} \quad [\vec{y}]_{\mathcal{B}} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}.$$

It is now easily checked that $[\vec{y}]_{\mathcal{B}}^* [\vec{x}]_{\mathcal{B}} = [\langle \vec{x}, \vec{y} \rangle]$. □

One way of interpreting this Proposition is that once you choose an orthonormal basis for an inner product space, it looks like \mathbb{F}^n with the standard inner product! This is reminiscent of the fact that once you choose a basis for a vector space, it looks like \mathbb{F}^n .

It is important that an orthonormal basis is chosen in Proposition 5.1.2, and in fact the result is false if you do not choose an orthonormal basis.

Exercise 59 Show that Proposition 5.1.2 is false if the basis \mathcal{B} is not orthonormal.

Taking the adjoint of a matrix may seem like a strange thing to do, which just happens to be useful when computing inner products. It turns out that taking the adjoint of a matrix gives us much more than a tool to compute inner products. In this chapter, the adjoint plays a fundamental role in the main theorems. With this promise for future payoff in your back pocket, let's investigate some basic properties of the adjoint.

Proposition 5.1.3 (Properties of the Adjoint)

Let $A, B \in M_{n \times m}(\mathbb{F})$ and $C \in M_{m \times k}(\mathbb{F})$. Then:

- (a) $(A + B)^* = A^* + B^*$.
- (b) $(A^*)^* = A$.
- (c) $(\alpha A)^* = \bar{\alpha}A^*$ for all $\alpha \in \mathbb{F}$.
- (d) $(AC)^* = C^*A^*$.

Exercise 60 Prove Proposition 5.1.3.

We will close out the first section of this chapter by investigating the relationship between adjoints and linear maps. We have emphasised over and over again that after choosing bases, a matrix is a linear map, and a linear map is a matrix. “What linear map does the adjoint of a matrix give us?” I hear you ask. This is absolutely the right question.

Let's restrict our attention to the case where $L: V \rightarrow V$ is an operator on an inner product space V , since this is the main interest of this chapter. The more general case of a linear map $L: V \rightarrow W$ between inner product spaces is dealt with in an exercise at the end of this section.

To make our lives easier, let's choose an orthonormal basis \mathcal{B} of V . Let $M: V \rightarrow V$ be the operator obtained by simply taking the matrix $[L]_{\mathcal{B}}$, and taking its adjoint. That is, $[M]_{\mathcal{B}} = [L]_{\mathcal{B}}^*$. Let's see what we can say about the operator M .

Since \mathcal{B} is an orthonormal basis for V , by Proposition 5.1.2 we know the inner product is given to us by matrix multiplication (after taking an adjoint of course). So, using the properties of the adjoint (Proposition 5.1.3) we have that for all $\vec{v}, \vec{w} \in V$,

$$\begin{aligned} \langle \vec{v}, M(\vec{w}) \rangle &= [M(\vec{w})]_{\mathcal{B}}^* [\vec{v}]_{\mathcal{B}} \\ &= ([M]_{\mathcal{B}} [\vec{w}]_{\mathcal{B}})^* [\vec{v}]_{\mathcal{B}} \\ &= [\vec{w}]_{\mathcal{B}}^* [M]_{\mathcal{B}}^* [\vec{v}]_{\mathcal{B}} \\ &= [\vec{w}]_{\mathcal{B}}^* [L]_{\mathcal{B}} [\vec{v}]_{\mathcal{B}} \\ &= [\vec{w}]_{\mathcal{B}}^* [L(\vec{v})]_{\mathcal{B}} \\ &= \langle L(\vec{v}), \vec{w} \rangle. \end{aligned}$$

Thus M and L are related by the fact that for all $\vec{v}, \vec{w} \in V$, $\langle L(\vec{v}), \vec{w} \rangle = \langle \vec{v}, M(\vec{w}) \rangle$.

In the special case when we are in \mathbb{F}^n with the standard inner product, this discussion leads to the following important property of adjoints.

Proposition 5.1.4 (The Fundamental Property of the Adjoint of a Matrix)

Let $A \in M_{n \times n}(\mathbb{F})$. Equip \mathbb{F}^n with the standard inner product. Then for all $\vec{v}, \vec{w} \in \mathbb{F}^n$,

$$\langle A\vec{v}, \vec{w} \rangle = \langle \vec{v}, A^*\vec{w} \rangle.$$

Proof: We have,

$$\langle A\vec{v}, \vec{w} \rangle = \vec{w}^* A\vec{v} = (\vec{w}^* (A^*)^*) \vec{v} = (A^*\vec{w})^* \vec{v} = \langle \vec{v}, A^*\vec{w} \rangle$$

completing the proof. □

Adjoint allows us to jump across an inner product, a fact that will be incredibly useful to us as we progress through this chapter.

Exercise 61

Let V and W be inner product spaces, and $L : V \rightarrow W$ a linear map. Prove that there exists a unique linear map $M : W \rightarrow V$ with the property that for all $\vec{v} \in V$ and all $\vec{w} \in W$, $\langle L(\vec{v}), \vec{w} \rangle = \langle \vec{v}, M(\vec{w}) \rangle$.

5.2 Orthogonal and Unitary Matrices

Our goal is to address the question of when an operator $L : V \rightarrow V$ has the property that there exists an orthonormal basis of V that are eigenvectors of L . We can approach this question by first choosing an arbitrary orthonormal basis of V . As revealed in Section 5.1, this puts us firmly in the setting of \mathbb{F}^n with the standard inner product. It is here that we shall remain for the remainder of this chapter.

Now recall that a matrix $A \in M_{n \times n}(\mathbb{F})$ admits a basis of eigenvectors for \mathbb{F}^n precisely when A is diagonalizable (Theorem 3.2.8), that is, if and only if there is an invertible matrix P such that $P^{-1}AP$ is diagonal. When this is the case, the columns of P will form a basis of \mathbb{F}^n consisting of eigenvectors of A . Thus being able to find a basis of eigenvectors for \mathbb{F}^n is effectively equivalent to being able to construct this matrix P . In our setting, we wish to impose the additional assumption that the columns of P form an *orthonormal* basis of eigenvectors with respect to the standard inner product on \mathbb{F}^n . What can we then say about P ?

Proposition 5.2.1

Let $P \in M_{n \times n}(\mathbb{F})$. Equip \mathbb{F}^n with the standard inner product. Then the following properties are equivalent:

- (a) The columns of P form an orthonormal basis for \mathbb{F}^n .
- (b) $P^* = P^{-1}$.
- (c) The rows of P form an orthonormal basis for \mathbb{F}^n .

Proof: We will prove the equivalence of (a) and (b). Let $P = [\vec{v}_1 \cdots \vec{v}_n]$. Then

$$P^*P = \begin{bmatrix} \vec{v}_1^* \\ \vdots \\ \vec{v}_n^* \end{bmatrix} \begin{bmatrix} \vec{v}_1 & \cdots & \vec{v}_n \end{bmatrix} = \begin{bmatrix} \langle \vec{v}_1, \vec{v}_1 \rangle & \cdots & \langle \vec{v}_1, \vec{v}_n \rangle \\ \vdots & \ddots & \vdots \\ \langle \vec{v}_n, \vec{v}_1 \rangle & \cdots & \langle \vec{v}_n, \vec{v}_n \rangle \end{bmatrix}.$$

From this we see that $P^*P = I_n$ if and only if

$$\langle \vec{v}_i, \vec{v}_j \rangle = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$$

equivalently, if and only if $\{\vec{v}_1, \dots, \vec{v}_n\}$ is an orthonormal set, which is precisely what we wanted to prove.

The proof of the equivalence of (c) with (a) and (b) is left as an exercise. \square

Exercise 62

Show that properties (b) and (c) of Proposition 5.2.1 are equivalent. (Hint: The rows of P are the columns of P^T .) Observe that since (a) and (b) are equivalent, this also proves the equivalence of (c) and (a).

Let's give a name to matrices that satisfy any of the equivalent properties in Proposition 5.2.1. It's customary to separate the case where $\mathbb{F} = \mathbb{R}$.

Definition 5.2.2

**Unitary Matrix,
Orthogonal Matrix**

A matrix $U \in M_{n \times n}(\mathbb{F})$ is called a **unitary matrix** if $U^* = U^{-1}$.

A matrix $Q \in M_{n \times n}(\mathbb{R})$ is called an **orthogonal matrix** if $Q^T = Q^{-1}$.

Note that an orthogonal matrix is by definition a real matrix. Of course, an orthogonal matrix is by definition also a unitary matrix. However, we generally (but not always) reserve the adjective “unitary” for when we are working with complex matrices. You might argue that an orthogonal matrix should be called an *orthonormal* matrix, since its columns are in fact orthonormal and not just orthogonal. You would have a point. Alas, the definition “orthogonal” is deeply entrenched in the literature.

Example 5.2.3

The $n \times n$ identity matrix is unitary. (In fact, orthogonal.)

Example 5.2.4

Let $U = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}}i & \frac{1}{\sqrt{2}}i \end{bmatrix}$. Then U is unitary since

$$UU^* = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ -i & i \end{bmatrix} \begin{bmatrix} 1 & i \\ 1 & -i \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

which proves that $U^* = U^{-1}$.

Example 5.2.5 Let $A = \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$. Notice that the columns of A are orthogonal. However, A is **not** an orthogonal matrix since its columns are not orthonormal!

Example 5.2.6 For $\theta \in \mathbb{R}$, let $R_\theta = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$ (the rotation-by- θ -counterclockwise matrix in \mathbb{R}^2). Then R_θ is orthogonal. We leave it to you to check that

$$R_\theta R_\theta^T = I.$$

Unitary matrices turn out to be interesting for a variety of reasons. When viewed as linear maps with respect to the standard basis in \mathbb{F}^n , they do not affect the standard inner product. That is, unitary matrices preserve length and angle!

Proposition 5.2.7 Let $U \in M_{n \times n}(\mathbb{F})$ be a unitary matrix and consider \mathbb{F}^n with the standard inner product. Then:

- (a) $\langle U\vec{v}, U\vec{w} \rangle = \langle \vec{v}, \vec{w} \rangle$ for all $\vec{v}, \vec{w} \in \mathbb{F}^n$.
- (b) $\|U\vec{v}\| = \|\vec{v}\|$ for all $\vec{v} \in \mathbb{F}^n$.

Proof: (a) Since the standard inner product on \mathbb{F}^n is given by $\langle \vec{v}, \vec{w} \rangle = [\vec{w}^* \vec{v}]$, we have

$$[\langle U\vec{v}, U\vec{w} \rangle] = (U\vec{w})^* U\vec{v} = \vec{w}^* U^* U\vec{v} = \vec{w}^* \vec{v} = [\langle \vec{v}, \vec{w} \rangle]$$

completing the proof.

(b) Using part (a), we have

$$\|U\vec{v}\| = \sqrt{\langle U\vec{v}, U\vec{v} \rangle} = \sqrt{\langle \vec{v}, \vec{v} \rangle} = \|\vec{v}\|.$$

□

Exercise 63 Let $Q \in M_{n \times n}(\mathbb{R})$ be an orthogonal matrix and let $\vec{v}, \vec{u} \in \mathbb{R}^n$ be non-zero vectors. Prove that the angle between \vec{v} and \vec{u} is equal to the angle between $Q\vec{v}$ and $Q\vec{u}$.

5.3 Schur's Triangularization Theorem

We have just seen that a unitary matrix can be viewed as a special kind of change of basis matrix, one that preserves length and angles. While it is not true that every matrix is diagonalizable, we will now see that every matrix is upper-triangularizable (at least over \mathbb{C}), that is, for every matrix we can find a basis for \mathbb{C}^n with respect to which the matrix is upper-triangular. Even better, we can choose this basis to be orthonormal (with respect to the standard inner product)—meaning, the change of basis matrix will be unitary. We can sometimes do this over \mathbb{R} , too.

Theorem 5.3.1 (Schur's Triangularization Theorem)

Let $A \in M_{n \times n}(\mathbb{C})$. There is a unitary matrix $U \in M_{n \times n}(\mathbb{C})$ and an upper-triangular matrix $T \in M_{n \times n}(\mathbb{C})$ such that $U^*AU = T$:

$$U^*AU = T = \begin{bmatrix} \lambda_1 & * & * \\ 0 & \ddots & * \\ 0 & 0 & \lambda_n \end{bmatrix}.$$

The diagonal entries λ_i of T are the complex eigenvalues of A (repeated according to multiplicity).

Furthermore, if A has real coefficients, and if all the eigenvalues λ_i of A are in \mathbb{R} , then U may be chosen to be real and orthogonal and T will be in $M_{n \times n}(\mathbb{R})$.

NOTE: The $*$ entries in the matrix just indicate that any element of \mathbb{F} can appear in those entries of the matrix.

REMARK

In stating this theorem, we're using the fact that the characteristic polynomial of a matrix $A \in M_{n \times n}(\mathbb{C})$ will be a degree n polynomial with complex coefficients. Hence, by the fundamental theorem of algebra, A will have n eigenvalues in \mathbb{C} (possibly repeated).

This applies in particular to the case where $A \in M_{n \times n}(\mathbb{R})$. Such a matrix will typically have non-real roots. Schur's theorem shows that A can be triangularized using a unitary matrix $U \in M_{n \times n}(\mathbb{C})$. In the special case where the n eigenvalues of A are all real, the theorem asserts that we can arrange for U and T to be in $M_{n \times n}(\mathbb{R})$.

Proof of Theorem 5.3.1: We will proceed by induction on n . For $n = 1$, this is clearly true since every 1×1 matrix is upper-triangular.

Now suppose A is an $n \times n$ matrix, and assume that the theorem is true for all $(n - 1) \times (n - 1)$ matrices. Let \vec{v}_1 be a unit eigenvector of A with eigenvalue λ . Extend $\{\vec{v}_1\}$ to a basis for \mathbb{C}^n and perform the Gram-Schmidt procedure to obtain an orthonormal basis $\{\vec{v}_1, \vec{w}_2, \dots, \vec{w}_n\}$.

Let $V_1 = [\vec{v}_1 \ \vec{w}_2 \ \cdots \ \vec{w}_n]$. Since the columns are orthonormal, V_1 is a unitary matrix. We then have

$$\begin{aligned} V_1^*AV_1 &= \begin{bmatrix} \vec{v}_1^* \\ \vec{w}_2^* \\ \vdots \\ \vec{w}_n^* \end{bmatrix} A [\vec{v}_1 \ \vec{w}_2 \ \cdots \ \vec{w}_n] \\ &= \begin{bmatrix} \vec{v}_1^* \\ \vec{w}_2^* \\ \vdots \\ \vec{w}_n^* \end{bmatrix} [A\vec{v}_1 \ A\vec{w}_2 \ \cdots \ A\vec{w}_n] \end{aligned}$$

$$\begin{aligned}
&= \begin{bmatrix} \vec{v}_1^* \\ \vec{w}_2^* \\ \vdots \\ \vec{w}_n^* \end{bmatrix} [\lambda \vec{v}_1 \quad A\vec{w}_2 \quad \cdots \quad A\vec{w}_n] \\
&= \begin{bmatrix} \langle \lambda \vec{v}_1, \vec{v}_1 \rangle & \langle A\vec{w}_2, \vec{v}_1 \rangle & \cdots & \langle A\vec{w}_n, \vec{v}_1 \rangle \\ \langle \lambda \vec{v}_1, \vec{w}_2 \rangle & \langle A\vec{w}_2, \vec{w}_2 \rangle & \cdots & \langle A\vec{w}_n, \vec{w}_2 \rangle \\ \vdots & \vdots & \ddots & \vdots \\ \langle \lambda \vec{v}_1, \vec{w}_n \rangle & \langle A\vec{w}_2, \vec{w}_n \rangle & \cdots & \langle A\vec{w}_n, \vec{w}_n \rangle \end{bmatrix} \\
&= \left[\begin{array}{c|ccc} \lambda & * & * & * \\ \hline 0 & & & \\ \vdots & & B & \\ 0 & & & \end{array} \right].
\end{aligned}$$

Now B is an $(n-1) \times (n-1)$ matrix, so by the inductive hypothesis there is a unitary matrix V_2 such that $V_2^* B V_2 = T_2$ where T_2 is upper-triangular. Let

$$U = V_1 \left[\begin{array}{c|c} 1 & 0 \\ \hline 0 & V_2 \end{array} \right].$$

Then $U^* U = I_n$ (see the Remark below) so U is unitary, and we have

$$\begin{aligned}
U^* A U &= \left[\begin{array}{c|c} 1 & 0 \\ \hline 0 & V_2 \end{array} \right]^* V_1^* A V_1 \left[\begin{array}{c|c} 1 & 0 \\ \hline 0 & V_2 \end{array} \right] \\
&= \left[\begin{array}{c|c} 1 & 0 \\ \hline 0 & V_2^* \end{array} \right] \left[\begin{array}{c|ccc} \lambda & * & * & * \\ \hline 0 & & B & \end{array} \right] \left[\begin{array}{c|c} 1 & 0 \\ \hline 0 & V_2 \end{array} \right] \\
&= \left[\begin{array}{c|ccc} \lambda & & * & \\ \hline 0 & & V_2^* B V_2 & \end{array} \right] \\
&= \left[\begin{array}{c|ccc} \lambda & & * & \\ \hline 0 & & T_2 & \end{array} \right] = T.
\end{aligned}$$

Since this last matrix T is upper-triangular, the first part of the theorem has been proved by the principle of mathematical induction.

Now observe that since U is unitary, $U^* = U^{-1}$ so we see that A is similar to T . Consequently, the eigenvalues of A and T are the same. However, since T is upper-triangular, its eigenvalues are its diagonal entries. This proves the second part of the theorem.

Finally, if A has real coefficients and all of its eigenvalues are real, then in the inductive step above, \vec{v}_1 will be in \mathbb{R}^n , λ will be in \mathbb{R} , and all the consequent steps can be carried out in \mathbb{R}^n with real arithmetic. We'll leave the careful verification to you. \square

REMARK

In the proof of Schur's Triangularization Theorem, a trick was used, commonly referred to as *block multiplication* of matrices. Here's how it works in general. Suppose you have two square $(n+m) \times (n+m)$ matrices, and you imagine drawing imaginary lines, splitting up the two matrices into sub-matrices as follows:

$$\left[\begin{array}{c|c} A_1 & B_1 \\ \hline C_1 & D_1 \end{array} \right] \quad \text{and} \quad \left[\begin{array}{c|c} A_2 & B_2 \\ \hline C_2 & D_2 \end{array} \right]$$

where A_1 and A_2 are $n \times n$ matrices, B_1 and B_2 are $n \times m$ matrices, C_1 and C_2 are $m \times n$ matrices and D_1 and D_2 are $m \times m$ matrices. Then the product of the two matrices is given by

$$\left[\begin{array}{c|c} A_1 & B_1 \\ \hline C_1 & D_1 \end{array} \right] \left[\begin{array}{c|c} A_2 & B_2 \\ \hline C_2 & D_2 \end{array} \right] = \left[\begin{array}{c|c} A_1A_2 + B_1C_2 & A_1B_2 + B_1D_2 \\ \hline C_1A_2 + D_1C_2 & C_1B_2 + D_1D_2 \end{array} \right].$$

One way to think about this is that you are doing regular matrix multiplication, but instead of the entries of the matrix being numbers, the entries are smaller matrices! Let's look at the top left $n \times n$ block of the resulting matrix, just to see whether or not this is plausible. Since A_1 and A_2 are both $n \times n$ matrices, the product A_1A_2 exists and is an $n \times n$ matrix. The block B_1 is an $n \times m$ matrix and C_2 is an $m \times n$ matrix. Therefore B_1C_2 exists and is an $n \times n$ matrix. Alas, the sum $A_1A_2 + B_1C_2$ is a sum of two $n \times n$ matrices, so the sum makes sense and the result is an $n \times n$ matrix. You can do a similar analysis for the other three blocks to see that the sizes of the matrices allow for the products and sums written there to be defined.

Verify this yourself for two 5×5 matrices, where $n = 2$ and $m = 3$.

The take home message is that every matrix is similar (over \mathbb{C} !) to an upper-triangular one. (But not to a unique upper-triangular matrix. See the next example.)

Example 5.3.2

Let $T_1 = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 2 & 2 \\ 0 & 0 & 2 \end{bmatrix}$ and $T_2 = \begin{bmatrix} 2 & -1 & \sqrt{2} \\ 0 & 1 & -\sqrt{2} \\ 0 & 0 & 2 \end{bmatrix}$. Then you can check that $U^*T_1U = T_2$ with the unitary matrix

$$U = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

So unlike with diagonalization, where the resulting diagonal matrix is unique up to a re-ordering of diagonal entries, the situation with triangularization is a bit more complicated.

Example 5.3.3

To illustrate the power of Schur's theorem, let's give a quick proof of the fact that, for any matrix $A \in M_{n \times n}(\mathbb{C})$, its determinant is the product of its eigenvalues and its trace is the sum of its eigenvalues. (This was given without proof in Corollary 3.1.18.)

By Schur's theorem we know there is a unitary matrix U and an upper-triangular matrix

$$T = \begin{bmatrix} \lambda_1 & * & * \\ 0 & \ddots & * \\ 0 & 0 & \lambda_n \end{bmatrix}$$

such that $U^*AU = T$. The characteristic polynomial of T is

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

so the λ_i are the eigenvalues of T . Furthermore we have

$$\text{tr}(T) = \lambda_1 + \cdots + \lambda_n \quad \text{and} \quad \det(T) = \lambda_1 \cdots \lambda_n.$$

Since A and T are similar, they have the same eigenvalues, determinant and trace, completing the proof.

At this point we would be remiss not to mention one of the more intriguing consequences of Schur's theorem: the Cayley–Hamilton theorem, which says that if you plug a matrix into its own characteristic polynomial you get the zero matrix. That is, A is a “root” of its own characteristic polynomial!

For example, if $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$, then its characteristic polynomial is $C_A(\lambda) = \lambda^2 - 5\lambda - 2$. If we plug $\lambda = A$ into the expression $\lambda^2 - 5\lambda - 2I$, we get:

$$A^2 - 5A - 2I = \begin{bmatrix} 7 & 10 \\ 15 & 22 \end{bmatrix} - \begin{bmatrix} 5 & 10 \\ 15 & 20 \end{bmatrix} - \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

This is a general phenomenon. Let's agree that if $p(x) = a_0 + a_1x + \cdots + a_nx^n$ is a polynomial and if A is a square matrix, then $p(A) = a_0I + a_1A + \cdots + a_nA^n$.

Theorem 5.3.4 (Cayley–Hamilton Theorem)

Let $A \in M_{n \times n}(\mathbb{C})$. Then $C_A(A) = 0_{n \times n}$.

Proof: By Schur's theorem, we can find a unitary matrix U and an upper triangular matrix

$$T = \begin{bmatrix} \lambda_1 & * & * \\ 0 & \ddots & * \\ 0 & 0 & \lambda_n \end{bmatrix}$$

in $M_{n \times n}(\mathbb{C})$ such that $A = UTU^*$. The diagonal entries of T are the eigenvalues of A , and therefore

$$C_A(\lambda) = (\lambda_1 - \lambda) \cdots (\lambda_n - \lambda).$$

As a result,

$$\begin{aligned} C_A(A) &= (\lambda_1 I - A) \cdots (\lambda_n I - A) \\ &= (\lambda_1 I - UTU^*) \cdots (\lambda_n I - UTU^*) \\ &= (\lambda_1 UU^* - UTU^*) \cdots (\lambda_n UU^* - UTU^*) \\ &= U(\lambda_1 I - T)U^* \cdots U(\lambda_n I - T)U^* \\ &= UC_A(T)U^*. \end{aligned}$$

So the proof will be complete if we can show that $C_A(T) = (\lambda_1 I - T) \cdots (\lambda_n I - T)$ is the zero matrix. We will do so by proving that, for every k with $1 \leq k \leq n$, the first k columns of the matrix $(\lambda_1 I - T)(\lambda_2 I - T) \cdots (\lambda_k I - T)$ only contain zeros.

We proceed by induction. When $k = 1$, the first column of the matrix $\lambda_1 I - T$ contains only zeros, since the first column of T is $[\lambda_1 \ 0 \ \cdots \ 0]^T$. Next, assume that the first k columns of $(\lambda_1 I - T)(\lambda_2 I - T) \cdots (\lambda_k I - T)$ contain only zeros. Convince yourself that, since each of the first $k + 1$ columns of $\lambda_{k+1} I - T$ has $n - k$ zeros at the bottom, the product $(\lambda_1 I - T)(\lambda_2 I - T) \cdots (\lambda_k I - T)(\lambda_{k+1} I - T)$ will have zeros everywhere in the first $k + 1$ columns. Thus, $C_A(T) = (\lambda_1 I - T)(\lambda_2 I - T) \cdots (\lambda_n I - T)$ is the $n \times n$ zero matrix, completing the proof. \square

Exercise 64

What is wrong with the following “proof” of the Cayley–Hamilton theorem?

“Since $C_A(\lambda) = \det(A - \lambda I)$, if we let $\lambda = A$ we get $C_A(A) = \det(A - AI) = \det(0) = 0$.”

We close this section by giving two fun applications of the Cayley–Hamilton theorem. The first one shows that we can express the inverse of an invertible matrix A as a linear combination of powers A ; the second one illustrates a new approach to computing powers A^k of an arbitrary $n \times n$ matrix A . Recall that we had seen one approach to computing powers of diagonalizable matrices in Section 3.3. Now we have something that works for *any* matrix, diagonalizable or not.

Example 5.3.5 (A^{-1} via Cayley–Hamilton)

Let $A \in M_{n \times n}(\mathbb{F})$. If the characteristic polynomial of A is expressed as

$$C_A(\lambda) = c_0 + c_1\lambda + \cdots + c_n\lambda^n$$

then the Cayley–Hamilton theorem tells us that

$$c_0I + c_1A + \cdots + c_nA^n = 0.$$

If A is invertible, then $c_0 \neq 0$ (since c_0 equals the determinant of A). Then we can multiply both sides of the above equation by $\frac{1}{c_0}A^{-1}$ to get

$$A^{-1} + \frac{c_1}{c_0}I + \frac{c_2}{c_0}A + \cdots + \frac{c_n}{c_0}A^{n-1} = 0.$$

Thus,

$$A^{-1} = -\frac{c_1}{c_0}I - \frac{c_2}{c_0}A - \cdots - \frac{c_n}{c_0}A^{n-1}.$$

For instance, if $A = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 1 & -1 \\ 2 & 1 & 1 \end{bmatrix}$ then the characteristic polynomial of A is

$$C_A(\lambda) = -\lambda^3 + 3\lambda^2 - 2\lambda - 4.$$

So

$$-A^3 + 3A^2 - 2A - 4I = 0.$$

Therefore, after multiplying by through by A^{-1} and re-arranging, we get

$$\begin{aligned} A^{-1} &= -\frac{1}{4}(A^2 - 3A + 2I) \\ &= -\frac{1}{4} \left(\begin{bmatrix} 3 & 5 & 0 \\ -2 & 0 & -2 \\ 4 & 6 & 2 \end{bmatrix} - \begin{bmatrix} 3 & 6 & 3 \\ 0 & 3 & -3 \\ 6 & 3 & 3 \end{bmatrix} + \begin{bmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix} \right) \\ &= \frac{1}{4} \begin{bmatrix} -2 & 1 & 3 \\ 2 & 1 & -1 \\ 2 & -3 & -1 \end{bmatrix}. \end{aligned}$$

Example 5.3.6 (A^k via Cayley–Hamilton)

Let's illustrate how the Cayley–Hamilton theorem can be used to compute powers of the 3×3 matrix $A = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \\ 0 & 0 & 2 \end{bmatrix}$ (which, incidentally, is **not** diagonalizable). The characteristic polynomial of A is $C_A(\lambda) = -\lambda^3 + 4\lambda^2 - 5\lambda + 2$. So

$$-A^3 + 4A^2 - 5A + 2I = 0.$$

Thus,

$$A^3 = 4A^2 - 5A + 2I.$$

Multiplying through by A and then re-using this expression for A^3 , we get

$$A^4 = 4A^3 - 5A^2 + 2A = 4(4A^2 - 5A + 2I) - 5A^2 + 2A = 11A^2 - 18A + 8I.$$

Proceeding this way, we are able to express A^k as a linear combination of I , A and A^2 .

At first sight this seems too tedious, but if you really think about it, you'll realize that we're effectively carrying out the Euclidean algorithm for polynomial division! More specifically, if we divide the polynomial λ^k by $C_A(\lambda)$, then we get

$$\lambda^k = q_k(\lambda)C_A(\lambda) + r_k(\lambda)$$

for some polynomials q_k and r_k (quotient and remainder). Plugging A into this, and using the Cayley–Hamilton theorem, we find that

$$A^k = r_k(A).$$

So all we have to do is find r_k . Fortunately, there are some fairly efficient software implementations of the Euclidean algorithm. Using one of these, we can find that r_{10} (the remainder of λ^{10} divided by $C_A(\lambda)$) is

$$r_{10}(\lambda) = 1013\lambda^2 - 2016\lambda + 1004.$$

Thus,

$$A^{10} = 1013A^2 - 2016A + 1004I = \begin{bmatrix} 1 & 10 & 3049 \\ 0 & 1 & 2046 \\ 0 & 0 & 1024 \end{bmatrix}.$$

5.4 Orthogonal and Unitary Diagonalization of Matrices

Let's return to our problem of trying to find a basis of orthogonal eigenvectors for a given $n \times n$ matrix. We begin by introducing some handy terminology. First, recall that a matrix $A \in M_{n \times n}(\mathbb{F})$ is diagonalizable over \mathbb{F} if there is an invertible matrix $P \in M_{n \times n}(\mathbb{F})$ such that $P^{-1}AP$ is diagonal. This prompts the following.

Definition 5.4.1

**Orthogonally
Diagonalizable,
Unitarily
Diagonalizable**

A matrix $A \in M_{n \times n}(\mathbb{R})$ is said to be **orthogonally diagonalizable** if there is an orthogonal matrix $Q \in M_{n \times n}(\mathbb{R})$ such that $Q^T A Q$ is diagonal.

A matrix $A \in M_{n \times n}(\mathbb{C})$ is said to be **unitarily diagonalizable** if there is a unitary matrix $U \in M_{n \times n}(\mathbb{C})$ such that $U^* A U$ is diagonal.

In Section 5.2, we proved:

Proposition 5.4.2**(Criterion for Orthogonal and Unitary Diagonalizability)**

(a) A matrix $A \in M_{n \times n}(\mathbb{R})$ is orthogonally diagonalizable if and only if there is a basis for \mathbb{R}^n consisting of orthonormal eigenvectors of A (orthonormal with respect to the standard inner product on \mathbb{R}^n , i.e., the dot product).

(b) A matrix $A \in M_{n \times n}(\mathbb{C})$ is unitarily diagonalizable if and only if there is a basis for \mathbb{C}^n consisting of orthonormal eigenvectors for A (orthonormal with respect to the standard inner product on \mathbb{C}^n).

Now the burning question is: when is a matrix in $M_{n \times n}(\mathbb{R})$ orthogonally diagonalizable and when is a matrix in $M_{n \times n}(\mathbb{C})$ unitarily diagonalizable? The answers are surprisingly simple: precisely when the matrix is *symmetric* or *normal*, respectively.

The formal definitions are given below, but here is the key observation. If $A = U D U^*$ is unitarily diagonalizable then $A^* = (U D U^*)^* = U D^* U^*$. If D is diagonal, then so is D^* ; if D is *real* too, then $D^* = D$, and so $A = A^*$. Thus, *if* a real matrix can be orthogonally diagonalized, it must be the case that $A = A^T$. On the other hand, if D has non-real entries, then we cannot say that $A = A^*$. However, notice that

$$A A^* = (U D U^*)(U D^* U^*) = U D D^* U^* = U D^* D U^* = A^* A.$$

Thus *if* a complex matrix can be unitarily diagonalized, it must be the case that $A A^* = A^* A$. This leads us to single out the following classes of matrices.

Definition 5.4.3

**Normal,
Self-adjoint,
Symmetric**

A matrix $A \in M_{n \times n}(\mathbb{F})$ is said to be **normal** if $A A^* = A^* A$.

A matrix $A \in M_{n \times n}(\mathbb{F})$ is said to be **self-adjoint** if $A = A^*$. Note that if $A \in M_{n \times n}(\mathbb{R})$ and self-adjoint, then in fact we have $A = A^T$ and we say that A is **symmetric**.

So, normal matrices are the matrices that commute with their adjoints. So, in particular, a self-adjoint matrix is normal. According to the definition, a real symmetric matrix is self-adjoint, but we more often use the term “symmetric” in the real case.

In some textbooks, if $A \in M_{n \times n}(\mathbb{C})$ is a complex self-adjoint matrix, then we say that A is **self-adjoint**.

Example 5.4.4

The matrices $A = \begin{bmatrix} 1 & 2-i \\ 2+i & 3 \end{bmatrix}$ and $B = \begin{bmatrix} 2 & 4 \\ 4 & 1 \end{bmatrix}$ are self-adjoint and B is symmetric.

The matrix $C = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$ is normal but not self-adjoint. Indeed, $C^* = -C$ and therefore $C C^* = C^* C = -C^2$.

Exercise 65 Prove that if $A \in M_{n \times n}(\mathbb{F})$ is self-adjoint, unitary or diagonal then A is normal.

With Schur's triangularization theorem (Theorem 5.3.1) in hand, we can easily prove the following result.

Theorem 5.4.5 (Spectral Theorem for Self-adjoint Matrices)

A square matrix in $M_{n \times n}(\mathbb{C})$ is self-adjoint if and only if it is unitarily diagonalizable and its eigenvalues are all real.

Proof: Suppose $A \in M_{n \times n}(\mathbb{C})$ is unitarily diagonalizable, say $A = UDU^*$ with U unitary and D diagonal. If the eigenvalues of A are real, then $D \in M_n(\mathbb{R})$. So $D^* = D^T = D$ since D is diagonal and real, and therefore

$$A^* = (UDU^*)^* = UDU^* = A.$$

This proves that A is self-adjoint.

Conversely, suppose that A is self-adjoint. By Schur's theorem, we know that there is a unitary matrix U such that $U^*AU = T$ where T is upper-triangular. Notice that

$$T^* = (U^*AU)^* = U^*A^*U = U^*AU = T$$

so T is self-adjoint. Since T is also upper-triangular, it must be the case that T is diagonal. So A is unitarily diagonalizable. Furthermore, the entries on the diagonal of any self-adjoint matrix are real (why?), so all of the eigenvalues of A are real, since they are the diagonal entries of T . \square

Exercise 66 Prove the following fact that was used in the preceding proof. If $A \in M_{n \times n}(\mathbb{C})$ is self-adjoint, then the diagonal entries of A are real.

Example 5.4.6 Let $A = \begin{bmatrix} 1 & 1-i \\ 1+i & 0 \end{bmatrix}$ and notice that $A^* = A$, so A is self-adjoint. Now, the Spectral Theorem for self-adjoint matrices guarantees us that A is unitarily diagonalizable, so let us find a unitary matrix U that unitarily diagonalizes A .

We find that the characteristic polynomial of A is $C_A(\lambda) = (\lambda + 1)(\lambda - 2)$, so A has the distinct eigenvalues $\lambda_1 = -1$ and $\lambda_2 = 2$. (As an aside, observe that these eigenvalues are real, as predicted by the spectral theorem for self-adjoint matrices.)

Since a 2×2 matrix A has 2 distinct eigenvalues, it follows from Proposition 3.2.12 that, at the very least, it is diagonalizable. Let's find eigenvectors. For λ_1 , we row reduce

$$A - \lambda_1 I = \begin{bmatrix} 2 & 1-i \\ 1+i & 1 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & \frac{1}{2}(1-i) \\ 0 & 0 \end{bmatrix}.$$

(for this computation, use $(1+i)^{-1} = \frac{1}{2}(1-i)$). Therefore a basis for the eigenspace corresponding to λ_1 is $\left\{ \begin{bmatrix} -\frac{1}{2} + \frac{1}{2}i \\ 1 \end{bmatrix} \right\}$, or more simply $\left\{ \begin{bmatrix} -1+i \\ 2 \end{bmatrix} \right\}$.

Next, for λ_2 , we row reduce

$$A - \lambda_2 i = \begin{bmatrix} -1 & 1-i \\ 1+i & -2 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & -1+i \\ 0 & 0 \end{bmatrix}.$$

So a basis for the corresponding eigenspace is $\left\{ \begin{bmatrix} 1-i \\ 1 \end{bmatrix} \right\}$.

Putting these two bases together, we get the basis $\left\{ \begin{bmatrix} -1+i \\ 2 \end{bmatrix}, \begin{bmatrix} 1-i \\ 1 \end{bmatrix} \right\}$ for \mathbb{C}^2 . Notice that this basis is orthogonal, since

$$\left\langle \begin{bmatrix} -1+i \\ 2 \end{bmatrix}, \begin{bmatrix} 1-i \\ 1 \end{bmatrix} \right\rangle = (-1+i)\overline{(1-i)} + 2\overline{1} = -2 + 2 = 0.$$

This is not a coincidence: in Proposition 5.4.10 we will show that any two eigenvectors of a normal matrix that correspond to distinct eigenvalues must be orthogonal. To get an *orthonormal* basis, we simply normalize, obtaining

$$\left\{ \frac{1}{\sqrt{6}} \begin{bmatrix} -1+i \\ 2 \end{bmatrix}, \frac{1}{\sqrt{3}} \begin{bmatrix} 1-i \\ 1 \end{bmatrix} \right\}.$$

We now set

$$U = \begin{bmatrix} \frac{1}{\sqrt{6}}(-1+i) & \frac{1}{\sqrt{3}}(1-i) \\ \frac{2}{\sqrt{6}} & \frac{1}{\sqrt{3}} \end{bmatrix}$$

to obtain our unitary matrix. Convince yourself that the identity $U^*AU = \text{diag}(-1, 2)$ holds.

If we specialize the Spectral Theorem for self-adjoint matrices to *real* self-adjoint matrices (otherwise known as real symmetric matrices), we arrive at the following result.

Theorem 5.4.7 (Spectral Theorem for Symmetric Matrices)

A square matrix in $M_{n \times n}(\mathbb{R})$ is symmetric if and only if it is orthogonally diagonalizable.

Proof: A matrix A in $M_{n \times n}(\mathbb{R})$ can be regarded as a matrix in $M_{n \times n}(\mathbb{C})$. In this light, we know from the preceding theorem that if $A \in M_{n \times n}(\mathbb{R})$ is symmetric (hence self-adjoint) then there is a unitary matrix $U \in M_{n \times n}(\mathbb{C})$ and a real diagonal matrix $D \in M_{n \times n}(\mathbb{R})$ such that $A = UDU^*$. Note in particular that this means that all of the eigenvalues of A are real, since they are the diagonal entries of D . The matrix U was provided by Schur's theorem, which when A is real and has all real eigenvalues, we know we can choose to be orthogonal. That is, we can find a real orthogonal matrix U such that $A = UDU^* = UDU^T$. This proves that A is orthogonally diagonalizable.

We leave the proof of the converse as an easy exercise. □

Exercise 67

Complete the proof of Theorem 5.4.7 by showing that if $A \in M_{n \times n}(\mathbb{R})$ is orthogonally diagonalizable then A is symmetric.

Here is the definitive result that says when a matrix can be unitarily diagonalized.

Theorem 5.4.8 (Spectral Theorem for Normal Matrices)

A square matrix in $M_{n \times n}(\mathbb{C})$ is normal if and only if it is unitarily diagonalizable.

Proof: Suppose that $A \in M_{n \times n}(\mathbb{C})$ is normal. By Schur's triangularization theorem, we can find a unitary matrix U and an upper-triangular matrix T such that $U^*AU = T$. Then

$$TT^* = (U^*AU)(U^*A^*U) = U^*AA^*U = U^*A^*AU = T^*T.$$

Let t_{ij} denote the (i, j) th entry of T . Then the $(1, 1)$ entry of TT^* is

$$t_{11}\overline{t_{11}} + t_{12}\overline{t_{12}} + \cdots + t_{1n}\overline{t_{1n}} = |t_{11}|^2 + |t_{12}|^2 + \cdots + |t_{1n}|^2.$$

On the other hand, the $(1, 1)$ entry of T^*T is $|t_{11}|^2$. It follows that $t_{12} = \cdots = t_{1n} = 0$. Next, comparing the $(2, 2)$ entries in a similar manner, we find that $t_{21} = t_{23} = t_{24} = \cdots = t_{2n} = 0$. Continuing in this way, we see that T must be diagonal.

Thus, if A is normal, it is unitarily diagonalizable. The converse is left as an exercise. \square

Exercise 68

Complete the proof of Theorem 5.4.8 by showing that if $A \in M_{n \times n}(\mathbb{F})$ is unitarily diagonalizable then A is normal.

Example 5.4.9

Let $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$. We find the eigenvalues of A are i and $-i$ with bases for the eigenspaces given by

$$\left\{ \begin{bmatrix} i \\ 1 \end{bmatrix} \right\} \quad \text{and} \quad \left\{ \begin{bmatrix} -i \\ 1 \end{bmatrix} \right\},$$

respectively. Notice that with the standard inner product on \mathbb{C}^2 we have

$$\left\langle \begin{bmatrix} i \\ 1 \end{bmatrix}, \begin{bmatrix} -i \\ 1 \end{bmatrix} \right\rangle = i^2 + 1 = 0.$$

So the eigenvectors are orthogonal. Normalizing them and putting them in a matrix gives us the unitary matrix

$$U = \begin{bmatrix} \frac{1}{\sqrt{2}}i & -\frac{1}{\sqrt{2}}i \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix}.$$

Then, since both the columns of U are eigenvectors of A , we have $U^*AU = \begin{bmatrix} i & 0 \\ 0 & -i \end{bmatrix}$.

So we see here that A is unitarily diagonalizable, but the eigenvalues are not real, and A is not self-adjoint. (However, notice that A is normal!)

To summarize: we now know that if we have a normal (e.g., self-adjoint) matrix A , then it is diagonalizable (over \mathbb{C}) with a basis of orthonormal eigenvectors; further, if A is real and symmetric, then this can be done over \mathbb{R} . Unfortunately, the results so far don't really give any indication as to how we can find this basis of orthogonal eigenvectors—or equivalently,

how we can find the unitary U or real orthogonal Q that diagonalizes A . It is tempting to simply find a basis of eigenvectors as usual (that is, diagonalize as usual) and perform Gram–Schmidt on that basis to obtain an orthonormal basis. There’s only one problem with this idea: there’s a chance that while performing Gram–Schmidt, you are no longer left with eigenvectors!

Example 5.4.9 gives us a hint as to how to get around this. Notice that the two eigenvectors we found were already orthogonal. This was not a coincidence! For normal matrices (so in particular for self-adjoint and real symmetric matrices), eigenvectors corresponding to different eigenvalues are orthogonal. We prove this, along with a few other useful facts about normal matrices, in the following proposition.

Proposition 5.4.10 (Properties of Normal Matrices)

Let $A \in M_{n \times n}(\mathbb{F})$ be a normal matrix. Equip \mathbb{F}^n with the standard inner product. Then:

- (a) For all $\vec{x} \in \mathbb{F}^n$, $\|A\vec{x}\| = \|A^*\vec{x}\|$.
- (b) If $\vec{x} \in \mathbb{F}^n$ is an eigenvector for A with eigenvalue λ , then \vec{x} is an eigenvector for A^* with eigenvalue $\bar{\lambda}$.
- (c) If \vec{x} and \vec{y} in \mathbb{F}^n are eigenvectors of A with *distinct* eigenvalues λ and μ , then \vec{x} is orthogonal to \vec{y} .

Proof: (a) We have

$$\|A\vec{x}\|^2 = \langle A\vec{x}, A\vec{x} \rangle = \langle \vec{x}, A^*A\vec{x} \rangle = \langle \vec{x}, AA^*\vec{x} \rangle = \langle A^*\vec{x}, A^*\vec{x} \rangle = \|A^*\vec{x}\|^2.$$

- (b) Suppose that $A\vec{x} = \lambda\vec{x}$. We want to prove that $A^*\vec{x} = \bar{\lambda}\vec{x}$. It will suffice to show that $\|A^*\vec{x} - \bar{\lambda}\vec{x}\| = 0$. Now, by part (a), we have

$$\|A^*\vec{x} - \bar{\lambda}\vec{x}\| = \|(A^* - \bar{\lambda}I)\vec{x}\| = \|(A^* - \bar{\lambda}I)^*\vec{x}\| = \|(A - \lambda I)\vec{x}\|.$$

Since $A\vec{x} = \lambda\vec{x}$, the last term above is 0, completing the proof.

- (c) We want to show that $\langle \vec{x}, \vec{y} \rangle = 0$. The trick is to consider $\lambda \langle \vec{x}, \vec{y} \rangle$:

$$\begin{aligned} \lambda \langle \vec{x}, \vec{y} \rangle &= \langle \lambda\vec{x}, \vec{y} \rangle \\ &= \langle A\vec{x}, \vec{y} \rangle \\ &= \langle \vec{x}, A^*\vec{y} \rangle \\ &= \langle \vec{x}, \mu\vec{y} \rangle \quad (\text{by part (b)}) \\ &= \mu \langle \vec{x}, \vec{y} \rangle. \end{aligned}$$

Thus $(\lambda - \mu) \langle \vec{x}, \vec{y} \rangle = 0$. Since $\lambda \neq \mu$, it follows that $\langle \vec{x}, \vec{y} \rangle = 0$, as required. \square

Proposition 5.4.10(c) tells us that if A is normal, then eigenvectors corresponding to distinct eigenvalues must be orthogonal. This result should remind you of Lemma 3.2.11, which says that eigenvectors corresponding to distinct eigenvalues are linearly independent. It’s important to note that eigenvectors of a normal matrix that correspond to the same eigenvalue need not be orthogonal. (Consider then eigenvector \vec{x} and its scalar multiple

$2\vec{x}$ for example.) Now that we have Proposition 5.4.10(c), we now know that we don't have to perform Gram–Schmidt on the *entire* basis of eigenvectors, just on the basis for each eigenspace! This leads to the following algorithm to unitarily diagonalize a normal matrix (so, in particular, a self-adjoint or real symmetric matrix).

ALGORITHM (Unitary Diagonalization of a Normal Matrix)

To unitarily diagonalize a normal matrix $A \in M_{n \times n}(\mathbb{F})$:

1. Diagonalize A as usual (see the Algorithm “Diagonalization of an Operator” in Section 3.2), obtaining $D = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix}$ and a basis of eigenvectors for \mathbb{F}^n .
2. Perform the Gram–Schmidt procedure on the bases of each of the eigenspaces E_{λ_i} of A , obtaining orthonormal bases for the eigenspaces. (This step is to be carried out using the standard inner product on \mathbb{F}^n .)
3. Take the union of the orthonormal bases from Step 2 to obtain $\mathcal{D} = \{\vec{w}_1, \dots, \vec{w}_n\}$, which is an orthonormal basis of eigenvectors for \mathbb{F}^n . Order the basis \mathcal{D} so that the orthonormal basis for E_{λ_1} is followed by the orthonormal basis for E_{λ_2} , etc.
4. Let $U = [\vec{w}_1 \ \cdots \ \vec{w}_n]$. Then U is unitary and $U^*AU = D$, with D as in Step 1.

If the above process is carried out on a real symmetric matrix $A \in M_{n \times n}(\mathbb{R})$, then the resulting matrix U will be real and orthogonal, and the diagonal matrix D will have real entries. Thus, we will have $U^T A U = D$ with the columns of U forming an orthonormal basis for \mathbb{R}^n .

In particular, suppose an $n \times n$ self-adjoint (or real symmetric) matrix A has n distinct eigenvalues. Since A has distinct eigenvalues, we know it's diagonalizable, and moreover, any basis of eigenvectors for \mathbb{F}^n will consist of one eigenvector for each eigenvalue. Since A is self-adjoint, hence normal, this basis must be orthogonal, since it consists of eigenvectors corresponding to distinct eigenvalues. In this case, to obtain an orthonormal basis (and therefore a unitary U such that U^*AU is diagonal) we simply need to normalize the eigenvectors in this basis. This special case was demonstrated in Example 5.4.6.

And now here is an example of the algorithm in all its glory.

Example 5.4.11

Let's unitarily diagonalize the matrix

$$A = \begin{bmatrix} 5 & -4 & -2 \\ -4 & 5 & -2 \\ -2 & -2 & 8 \end{bmatrix}.$$

First note that A is symmetric. A quick computation gives $C_A(\lambda) = \lambda(9 - \lambda)^2$. We'll now find bases for the eigenspaces corresponding to $\lambda = 0$ and $\lambda = 9$. Since A is self-adjoint, we know it's diagonalizable so we should have the geometric multiplicities $g_0 = 1$ and $g_9 = 2$.

For $\lambda = 0$, we row reduce

$$A - 0I = \begin{bmatrix} 5 & -4 & -2 \\ -4 & 5 & -2 \\ -2 & -2 & 8 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & 0 & -2 \\ 0 & 1 & -2 \\ 0 & 0 & 0 \end{bmatrix}.$$

Therefore a basis for the eigenspace corresponding to 0 is $\left\{ \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix} \right\}$. We wish to have an orthonormal basis for this eigenspace, so we normalize and choose instead the basis $\left\{ \frac{1}{3} \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix} \right\}$. To find the eigenspace corresponding to $\lambda = 9$ we row reduce

$$A - 9I = \begin{bmatrix} -4 & -4 & -2 \\ -4 & -4 & -2 \\ -2 & -2 & -1 \end{bmatrix} \rightarrow \begin{bmatrix} 2 & 2 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

giving a basis $\{\vec{v}_1, \vec{v}_2\} = \left\{ \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 2 \end{bmatrix} \right\}$. Note that \vec{v}_1 and \vec{v}_2 are not orthogonal, so we must perform Gram-Schmidt to obtain an orthogonal basis for this eigenspace.

Doing so, we let $\vec{w}_1 = \vec{v}_1$ and

$$\vec{w}_2 = \vec{v}_2 - \frac{\langle \vec{v}_2, \vec{w}_1 \rangle}{\|\vec{w}_1\|^2} \vec{w}_1 = \begin{bmatrix} -1 \\ 0 \\ 2 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} -1 \\ -1 \\ 4 \end{bmatrix}.$$

Now $\{\vec{w}_1, \vec{w}_2\}$ is an orthogonal basis for the eigenspace corresponding to $\lambda = 9$. To obtain an orthonormal basis $\{\vec{u}_1, \vec{u}_2\}$ we set

$$\vec{u}_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix} \quad \text{and} \quad \vec{u}_2 = \frac{1}{\sqrt{18}} \begin{bmatrix} -1 \\ -1 \\ 4 \end{bmatrix}.$$

Finally, $U^*AU = D$ where

$$U = \begin{bmatrix} \frac{2}{3} & -\frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{18}} \\ \frac{2}{3} & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{18}} \\ \frac{1}{3} & 0 & \frac{4}{\sqrt{18}} \end{bmatrix} \quad \text{and} \quad D = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 9 & 0 \\ 0 & 0 & 9 \end{bmatrix}.$$

As usual, the order of diagonal entries in D matches the order of eigenvectors in our change-of-basis matrix U .

Exercise 69

A matrix $A \in M_{n \times n}(\mathbb{C})$ is said to be **skew-self-adjoint** if $A^* = -A$.

(a) Prove that a skew-self-adjoint matrix is unitarily diagonalizable.

(b) Let $A = \begin{bmatrix} 0 & i \\ i & 0 \end{bmatrix}$. Find a unitary matrix U and a diagonal matrix D such that $U^*AU = D$.

5.5 Unitary Diagonalization of Operators

The spectral theorems from Section 5.4 are all stated in terms of matrices, which we viewed as operators on \mathbb{F}^n with the standard inner product. However, we know from Section 5.1 that given an arbitrary inner product space V and an orthonormal basis \mathcal{B} , the inner product applied to the coordinate vectors acts like the standard inner product on \mathbb{F}^n . More precisely, for all $\vec{v}, \vec{w} \in V$, $[\langle \vec{v}, \vec{w} \rangle] = [\vec{w}]_{\mathcal{B}}^* [\vec{v}]_{\mathcal{B}}$.

So, suppose $L : V \rightarrow V$ is an arbitrary operator on an inner product space, and we want to know whether or not there exists an *orthonormal* basis for V consisting of eigenvectors of L . The upshot is that we can simply choose an orthonormal basis \mathcal{B} of V , and apply our work from Section 5.4 to $[L]_{\mathcal{B}}$:

Theorem 5.5.1 (Spectral Theorem for Operators)

Let $L : V \rightarrow V$ be a linear operator on a finite-dimensional inner product space over \mathbb{F} . Let \mathcal{B} be an *orthonormal* basis for V and let $A = [L]_{\mathcal{B}}$. Then:

- (a) If $\mathbb{F} = \mathbb{C}$, there is an orthonormal basis for V consisting of eigenvectors of L if and only if A is normal.
- (b) If $\mathbb{F} = \mathbb{R}$, there is an orthonormal basis for V consisting of eigenvectors of L if and only if A is symmetric.

Proof: This follows from our preceding discussion and the corresponding spectral theorems for matrices (Theorems 5.4.8 and 5.4.7). \square

We emphasize once more that the previous theorem only works if \mathcal{B} is an orthonormal basis for V . See the following example.

Example 5.5.2 Endow $V = \mathcal{P}_1(\mathbb{R})$ with the inner product

$$\langle p, q \rangle = p(0)q(0) + p(1)q(1).$$

(Check that this is indeed an inner product using what we learned in the previous section!) Consider the linear operator $L : V \rightarrow V$ defined by

$$L(a + bx) = (a + b) + (a + b)x.$$

If we let $\mathcal{S} = \{1, x\}$ be the standard basis for $\mathcal{P}_1(\mathbb{R})$, then the corresponding matrix of L is

$$A = [L]_{\mathcal{S}} = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}.$$

Notice that A is symmetric. We'll leave it to you to check that the eigenvalues of A are 0 and 2, with eigenspaces

$$E_0 = \text{Span} \left\{ \begin{bmatrix} -1 \\ 1 \end{bmatrix} \right\} \quad \text{and} \quad E_2 = \text{Span} \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}.$$

Converting back to polynomials, we see that $-1 + x$ and $1 + x$ span the two corresponding eigenspaces of L . So there is no chance of finding orthogonal eigenvectors for L , since $-1 + x$ and $1 + x$ are not orthogonal! Indeed,

$$\langle -1 + x, 1 + x \rangle = (-1)(1) + (-1 + 1)(1 + 1) = -1.$$

This does **not** contradict the spectral theorem. Although the matrix A is symmetric, the problem here is that the standard basis \mathcal{S} is not orthonormal with respect to $\langle \cdot, \cdot \rangle$. Let's find an orthonormal basis and see what happens. Applying the Gram–Schmidt procedure to \mathcal{S} , we arrive at the orthonormal basis $\mathcal{B} = \left\{ \frac{1}{\sqrt{2}}, \sqrt{2} \left(x - \frac{1}{2} \right) \right\}$. The matrix of L with respect to \mathcal{B} is

$$[L]_{\mathcal{B}} = \frac{1}{2} \begin{bmatrix} 3 & 3 \\ 1 & 1 \end{bmatrix}$$

which is **not** symmetric, so L cannot be orthogonally diagonalized, according to the spectral theorem, and as we discovered above.

We close this chapter with a remark that gives a different perspective on what we did in this section.

REMARK

Our approach to the problem of finding an orthonormal basis for an operator on an arbitrary inner product space is not the only one. We decided to tackle it using matrices to keep everything as concrete as possible. It is possible to attack the problem head on, without choosing any bases, and without reducing to the case of matrices.

The key result is that every operator $L: V \rightarrow V$ on a finite-dimensional inner product space gives rise to another operator $L^*: V \rightarrow V$, called its *adjoint*. These two operators are linked by the following fundamental property:

$$\langle L(\vec{v}), \vec{u} \rangle = \langle \vec{v}, L^*(\vec{u}) \rangle \quad \text{for all } \vec{v}, \vec{u} \in V.$$

(Compare this with Proposition 5.1.4 and the exercise immediately after the proposition.)

Starting from this innocent seeming identity, it is possible to develop the theory of the adjoint operator all the way to proving the spectral theorem for operators. From this we can then deduce the spectral theorems for matrices. This is the reverse of what we've done!

The connection between the two approaches is the following. If we choose an orthonormal basis \mathcal{B} for V , then the \mathcal{B} -matrices $[L]_{\mathcal{B}}$ and $[L^*]_{\mathcal{B}}$ turn out to be adjoints of one another. That is, $[L]_{\mathcal{B}}^* = [L^*]_{\mathcal{B}}$. (This is not true if \mathcal{B} is not orthonormal.) So our two notions of “adjoint” are thus linked.

A natural question is: why do this? Why be unsatisfied with working with matrices as we had done? Here are two reasons.

1. Our approach to the spectral theorem for operators involved *choosing* an orthonormal basis for the underlying inner product space only to later discard this basis in favour of a diagonalizing one. This seems odd. More-so because there is no obvious choice to be made at the outset, besides picking a random basis and applying Gram–Schmidt to it.
2. The matrix approach is restricted to finite-dimensional inner product spaces. The theory of orthonormal diagonalization has important consequences for operators on infinite-dimensional inner product spaces (e.g. it is of great use in quantum mechanics). As such, an approach that is not tied down to matrices is very desirable.

5.6 Application: Classifying Quadratic Forms

There are many situations where you might find yourself interested in maximizing or minimizing a certain quantity. A physicist would want to determine when a mass moving down a hill reaches a stable equilibrium, which will be the case when the potential energy due to gravity is minimized. A statistician carrying out an experiment will want to minimize the error between real-world observations and the predictions of their model. A student of linear algebra will want to solve a system of equations $A\vec{x} = \vec{b}$, which will amount to minimizing $\|A\vec{x} - \vec{b}\|$ or, equivalently, minimizing $\|A\vec{x} - \vec{b}\|^2$. (We've already considered this last problem when we discussed the method of least squares in Section 4.7.)

In the simplest of these situations, you will be faced with the task of minimizing some kind of *quadratic* function. For instance, to solve the system

$$\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \vec{x} = \begin{bmatrix} 5 \\ 11 \end{bmatrix}$$

by means of solving a minimization problem as described above, you will want to minimize the quantity

$$\begin{aligned} f(\vec{x}) &= \left\| \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \vec{x} - \begin{bmatrix} 5 \\ 11 \end{bmatrix} \right\|^2 \\ &= (x_1 + 2x_2 - 5)^2 + (3x_1 + 4x_2 - 11)^2 \\ &= 10x_1^2 + 28x_1x_2 + 20x_2^2 - 76x_1 - 108x_2 + 146. \end{aligned}$$

Calculus teaches us that the local maximum and minimum values of a sufficiently differentiable function $f(\vec{x})$ (where $\vec{x} \in \mathbb{R}^n$) occur at the *critical points* of f , i.e. the points where all the partial derivatives of f vanish:

$$\frac{\partial f}{\partial x_i}(\vec{x}) = 0 \quad \text{for all } i = 1, \dots, n.$$

To determine whether a critical point gives a local maximum or minimum value of f (or neither), one employs a type of *second derivative test*. The idea is that near a critical point \vec{a} , f can be approximated by its second degree Taylor polynomial

$$f(\vec{x}) \approx f(\vec{a}) + \underbrace{\sum_{i=1}^n \frac{\partial f}{\partial x_i}(\vec{a})}_{=0} (x_i - a_i) + \frac{1}{2} \sum_{i,j=1}^n \frac{\partial^2 f}{\partial x_i \partial x_j}(\vec{a}) (x_i - a_i)(x_j - a_j). \quad (5.1)$$

The difference $f(\vec{x}) - f(\vec{a})$ will therefore be approximated by the quadratic quantity

$$\frac{1}{2} \sum_{i,j=1}^n \frac{\partial^2 f}{\partial x_i \partial x_j}(\vec{a})(x_i - a_i)(x_j - a_j). \quad (5.2)$$

For instance, if (5.2) is positive for all $\vec{x} \approx \vec{a}$, then we will have $f(\vec{x}) > f(\vec{a})$ for all $\vec{x} \approx \vec{a}$, and so $f(\vec{a})$ will be a local minimum value. Similarly, if (5.2) is negative for all $\vec{x} \approx \vec{a}$, $f(\vec{a})$ will be a local maximum. Thus we find ourselves interested in determining the sign of (5.2).

To make the previous analysis rigorous, we need to carry out a careful examination of the approximation (5.1). However, if $f(\vec{x})$ is itself already a *quadratic* polynomial, we can side-step this issue, since f will be *equal* to its degree-2 Taylor polynomial:

$$f(\vec{x}) = f(\vec{a}) + \sum_{i=1}^n \frac{\partial f}{\partial x_i}(\vec{a})(x_i - a_i) + \frac{1}{2} \sum_{i,j=1}^n \frac{\partial^2 f}{\partial x_i \partial x_j}(\vec{a})(x_i - a_i)(x_j - a_j). \quad (5.3)$$

In fact, what is going on here is a change of basis in the space of degree ≤ 2 polynomials in x_1, \dots, x_n from

$$\{1, x_1^2, \dots, x_n^2, x_1 x_2, \dots, x_{n-1} x_n\}$$

to

$$\{1, (x_1 - a_1)^2, \dots, (x_n - a_n)^2, (x_1 - a_1)(x_2 - a_2), \dots, (x_{n-1} - a_{n-1})(x_n - a_n)\}.$$

Example 5.6.1

If $f(x_1, x_2) = 10x_1^2 + 28x_1x_2 + 20x_2^2 - 76x_1 - 108x_2 + 146$ then f has the unique critical point $\vec{a} = (1, 2)$, and if we try to re-write f in terms of powers of $(x_1 - 1)$ and $(x_2 - 2)$, we would find that

$$f(x_1, x_2) = 10(x_1 - 1)^2 + 28(x_1 - 1)(x_2 - 2) + 20(x_2 - 2)^2.$$

We'll leave it to you to check that this is true, and to verify that the above expression is identical to

$$f(x_1, x_2) = f(1, 2) + \sum_{i=1}^2 \frac{\partial f}{\partial x_i}(1, 2)(x_i - a_i) + \frac{1}{2} \sum_{i,j=1}^2 \frac{\partial^2 f}{\partial x_i \partial x_j}(1, 2)(x_i - a_i)(x_j - a_j).$$

Returning to (5.2), if we let

$$u_i = x_i - a_i \quad \text{and} \quad a_{ij} = \frac{1}{2} \frac{\partial^2 f}{\partial x_i \partial x_j}(\vec{a})$$

then the expression takes the simpler form

$$\sum_{i,j=1}^n a_{ij} u_i u_j.$$

Definition 5.6.2 Quadratic Form

A (real) **quadratic form** in the variables $\vec{u} = (u_1, \dots, u_n)$ is a polynomial of the form

$$Q(\vec{u}) = \sum_{i,j=1}^n a_{ij} u_i u_j, \quad \text{where } a_{ij} \in \mathbb{R}.$$

Thus, a quadratic form is effectively a quadratic polynomial that does not have any linear or constant terms. We are interested in determining the *sign* of a quadratic form. Observe that $Q(\vec{0}) = 0$.

Definition 5.6.3

Positive Definite,
Negative Definite,
Semi-Definite,
Indefinite
Quadratic Form

A quadratic form $Q(\vec{u})$ is said to be

- **positive definite** if $Q(\vec{u}) > 0$ for all non-zero $\vec{u} \in \mathbb{R}^n$;
- **positive semi-definite** if $Q(\vec{u}) \geq 0$ for all $\vec{u} \in \mathbb{R}^n$;
- **negative definite** if $Q(\vec{u}) < 0$ for all non-zero $\vec{u} \in \mathbb{R}^n$;
- **negative semi-definite** if $Q(\vec{u}) \leq 0$ for all $\vec{u} \in \mathbb{R}^n$;
- **indefinite** if there exist $\vec{u}, \vec{v} \in \mathbb{R}^n$ such that $Q(\vec{u}) > 0$ and $Q(\vec{v}) < 0$.

Notice that every positive (resp. negative) definite quadratic form is also positive (negative) semi-definite.

Example 5.6.4

The quadratic form $Q_1(u_1, u_2) = u_1^2 + u_2^2$ is positive definite.

The quadratic form $Q_2(u_1, u_2) = u_1^2$ is positive semi-definite but not positive definite.

The quadratic form $Q_3(u_1, u_2) = -3u_1^2 - 4u_2^2$ is negative definite.

The quadratic form $Q_4(u_1, u_2) = -4u_1^2$ is negative semi-definite.

The quadratic form $Q_5(u_1, u_2) = 3u_1^2 - 2u_2^2$ is indefinite.

The verification of the above is entirely trivial. But how do we classify something more complicated like $Q_6(u_1, u_2) = 10u_1^2 + 28u_1u_2 + 20u_2^2$? (This is the quadratic form that appeared in the previous example.)

Our main tool in being able to classify a quadratic form into one of the above classes is its associated matrix. Consider an arbitrary two-variable quadratic form

$$Q(u_1, u_2) = \sum_{i,j=1}^2 a_{ij}u_iu_j = a_{11}u_1^2 + a_{12}u_1u_2 + a_{21}u_2u_1 + a_{22}u_2^2.$$

We can re-write this as

$$Q(u_1, u_2) = \begin{bmatrix} u_1 & u_2 \end{bmatrix} \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}$$

or more simply as

$$Q(\vec{u}) = \vec{u}^T \begin{bmatrix} a_{11} & \frac{a_{12}+a_{21}}{2} \\ \frac{a_{12}+a_{21}}{2} & a_{22} \end{bmatrix} \vec{u}.$$

(This type of expression should ring some bells. Compare it to our construction of the Gram matrix of an inner product in Section 5.7.) The same kind of construction is possible for n -variable quadratic forms:

Proposition 5.6.5

Let $Q(\vec{u}) = \sum_{i,j=1}^n a_{ij}u_iu_j$ be a quadratic form. If we let A be the $n \times n$ matrix whose (i, j) th entry is $\frac{a_{ij} + a_{ji}}{2}$, then

$$Q(\vec{u}) = \vec{u}^T A \vec{u}.$$

Exercise 70

Prove Proposition 5.6.5.

Definition 5.6.6

Matrix Associated
to a Quadratic
Form

The matrix $A \in M_{n \times n}(\mathbb{R})$ constructed in Proposition 5.6.5 is called the **matrix associated to the quadratic form** $Q(\vec{u})$.

Example 5.6.7

The matrix associated to the quadratic form $Q(u_1, u_2) = 10u_1^2 + 28u_1u_2 + 20u_2^2$ is

$$A = \begin{bmatrix} 10 & 14 \\ 14 & 20 \end{bmatrix}.$$

Notice that the diagonal entries are the coefficients of the pure square terms u_1^2 and u_2^2 while the off-diagonal entries are *one-half* of the coefficient of the mixed term u_1u_2 .

Example 5.6.8

The matrix associated to the quadratic form $Q(u_1, u_2, u_3) = 3u_1^2 + u_1u_2 - 2u_2u_3 + 3u_3^2$ is

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

Again, notice that the diagonal entries are the coefficients of the pure square terms, while the off-diagonal entries are one-half the coefficients of the mixed terms.

The key thing to glean from the previous two examples is that the matrices we got were symmetric. This is true in general.

Proposition 5.6.9

Let $A \in M_{n \times n}(\mathbb{R})$ be the matrix associated to the quadratic form $Q(\vec{u})$. Then A is symmetric.

Exercise 71

Prove Proposition 5.6.9.

REMARK

We can reverse the sequence of ideas presented above. Starting from a symmetric matrix $A \in M_{n \times n}(\mathbb{R})$, we can create a quadratic form $Q(\vec{u}) = \vec{u}^T A \vec{u}$. Propositions 5.6.5 and 5.6.9 guarantee that in this way we are able to produce all quadratic forms.

Thus, quadratic forms in n -variables and $n \times n$ symmetric matrices are essentially one and the same. Can you formulate this as some kind of isomorphism between two vector spaces?

Given this, we can now appeal to the spectral theorem for symmetric matrices (Theorem 5.4.7). The upshot is the following result, which says that the sign of $Q(\vec{u})$ is determined by the signs of the eigenvalues of A .

Theorem 5.6.10 (Classification of Quadratic Forms)

Let $Q(\vec{u})$ be a quadratic form with associated matrix A . Let $\lambda_1, \dots, \lambda_n$ be the eigenvalues of A . Then:

1. $Q(\vec{u})$ is positive definite if and only if $\lambda_i > 0$ for all $i \in \{1, \dots, n\}$.
2. $Q(\vec{u})$ is positive semi-definite if and only if $\lambda_i \geq 0$ for all $i \in \{1, \dots, n\}$.
3. $Q(\vec{u})$ is negative definite if and only if $\lambda_i < 0$ for all $i \in \{1, \dots, n\}$.
4. $Q(\vec{u})$ is negative semi-definite if and only if $\lambda_i \leq 0$ for all $i \in \{1, \dots, n\}$.
5. $Q(\vec{u})$ is indefinite if and only if $\lambda_i > 0$ and $\lambda_j < 0$ for some $i, j \in \{1, \dots, n\}$.

Proof: By the spectral theorem for symmetric matrices (Theorem 5.4.7), we can orthogonally diagonalize A . That is, there exists an orthogonal matrix $P \in M_{n \times n}(\mathbb{R})$ such that $A = PDP^T$, where $D = \text{diag}(\lambda_1, \dots, \lambda_n)$. Without loss of generality, let's label the eigenvalues by descending order according to size: $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. (Note that the eigenvalues are all real, by the spectral theorem, so this makes sense.) Then

$$Q(\vec{u}) = \vec{u}^T A \vec{u} = \vec{u}^T P D P^T \vec{u} = (P^T \vec{u})^T D (P^T \vec{u}).$$

If we let $\vec{y} = P^T \vec{u} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}$ (equivalently, $\vec{u} = P \vec{y}$), then we can write

$$Q(\vec{u}) = \vec{y}^T D \vec{y} = \vec{y}^T \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} \vec{y} = \lambda_1 y_1^2 + \dots + \lambda_n y_n^2.$$

The sign of Q is now very easy to determine. For instance, if all the λ_i are positive, then $Q(\vec{u}) > 0$ for all $\vec{u} \neq \vec{0}$. This proves (a). The proofs of (b), (c) and (d) are similar. For (e), consider $\vec{u} = P \vec{y}$, where $\vec{y} = \vec{e}_i$ and $\vec{y} = \vec{e}_j$ (where \vec{e}_k is the k th standard basis vector for \mathbb{R}^n). Then $Q(\vec{u}) = \lambda_i > 0$ and $Q(\vec{u}) = \lambda_j < 0$, respectively, proving that Q is indefinite. \square

The take-away from the above proof is that when we diagonalize A , the quadratic form becomes simple: all the mixed terms disappear, and we are left only with pure terms. (Essentially, we *completed the squares*.) This is generally what diagonalization does. It removes unnecessary complications that are present because we are, in a sense, working with a less than optimal point of view.

Example 5.6.11

The matrix associated to the quadratic form $Q(u_1, u_2) = 10u_1^2 + 28u_1u_2 + 20u_2^2$ is

$$A = \begin{bmatrix} 10 & 14 \\ 14 & 20 \end{bmatrix}.$$

Its eigenvalues are

$$\lambda_1 = 15 + \sqrt{221} \quad \text{and} \quad \lambda_2 = 15 - \sqrt{221}.$$

Since these are positive, we conclude that $Q(\vec{u})$ is positive definite.

Example 5.6.12

The matrix associated to the quadratic form $Q(u_1, u_2, u_3) = 3u_1^2 + u_1u_2 - 2u_2u_3 + 3u_3^2$ is

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

Its eigenvalues are

$$\lambda_1 = \frac{3 + \sqrt{14}}{2}, \quad \lambda_2 = 3 \quad \text{and} \quad \lambda_3 = \frac{3 - \sqrt{14}}{2}.$$

Since $\lambda_1 > 0$ and $\lambda_3 < 0$, it follows that $Q(\vec{u})$ is indefinite.

Let's return now to our motivating problem of optimizing a quadratic function $f(\vec{x})$. We had reduced the issue to determining the sign of the quadratic form

$$Q(\vec{u}) = \frac{1}{2} \sum_{i,j=1}^n \frac{\partial^2 f}{\partial x_i \partial x_j}(\vec{a}) u_i u_j,$$

where \vec{a} is a critical point of the function. The matrix associated to Q is $A = \frac{1}{2}H$, with

$$H = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2}(\vec{a}) & \frac{\partial^2 f}{\partial x_1 \partial x_2}(\vec{a}) & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n}(\vec{a}) \\ \frac{\partial^2 f}{\partial x_2 \partial x_1}(\vec{a}) & \frac{\partial^2 f}{\partial x_2^2}(\vec{a}) & \cdots & \frac{\partial^2 f}{\partial x_2 \partial x_n}(\vec{a}) \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1}(\vec{a}) & \frac{\partial^2 f}{\partial x_n \partial x_2}(\vec{a}) & \cdots & \frac{\partial^2 f}{\partial x_n^2}(\vec{a}) \end{bmatrix},$$

where we've implicitly used the fact that for a sufficiently differentiable function f , its mixed second-order partial derivatives are equal:

$$\frac{\partial^2 f}{\partial x_i \partial x_j} = \frac{\partial^2 f}{\partial x_j \partial x_i}.$$

The matrix H above is called the **Hessian matrix** of f . It plays a role analogous to the second derivative of a single-variable function. Our discussion above tells us that the nature of the critical point \vec{a} is tied to the sign of Q hence to the eigenvalues of H . In particular, $f(\vec{a})$ will be a local maximum value of f if Q is negative definite, and will be a local minimum if Q is positive definite. This is part of the so-called *second derivative test for multivariable functions*. To learn more, take a course in multivariable calculus!

Example 5.6.13 Let's revisit our function

$$f(x_1, x_2) = \left\| \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \vec{x} - \begin{bmatrix} 5 \\ 11 \end{bmatrix} \right\|^2 = 10x_1^2 + 28x_1x_2 + 20x_2^2 - 76x_1 - 108x_2 + 146.$$

Our work in Example 5.6.1 shows that the Hessian matrix of f at the critical point $\vec{a} = (1, 2)$ is

$$H = 2A = \begin{bmatrix} 20 & 28 \\ 28 & 40 \end{bmatrix},$$

where A is the matrix associated to the quadratic form $Q(u_1, u_2) = 20u_1^2 + 28u_1u_2 + 40u_2^2$. In Example 5.6.11, we showed that Q is positive-definite. Thus, $f(1, 2) = 0$ is a local minimum value of f . In fact, since $f(x_1, x_2) \geq 0$ for all (x_1, x_2) , it must be the case that this is a *global* minimum value.

Of course, this is to be expected, since $\vec{x} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$ is a solution to the system

$$\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \vec{x} = \begin{bmatrix} 5 \\ 11 \end{bmatrix},$$

hence definitely minimizes $f(x_1, x_2)$! In fact, we know that $(x_1, x_2) = (1, 2)$ gives us the unique global minimum value of f . But it's good to know that our methods above allow us to reach this same conclusion. Their true power emerges once we consider more complicated (in particular, non-quadratic) functions f .

5.7 Application: Inner Products and Gram Matrices

The goal of this section is to describe what all inner products on \mathbb{F}^n (and indeed on any finite-dimensional inner product space) look like in terms of matrices.

Suppose we have an inner product $\langle \cdot, \cdot \rangle$ on \mathbb{C}^2 and suppose we know what it does to the standard basis $\{\vec{e}_1, \vec{e}_2\}$:

$$\langle \vec{e}_1, \vec{e}_1 \rangle = 2, \quad \langle \vec{e}_2, \vec{e}_2 \rangle = 3, \quad \text{and} \quad \langle \vec{e}_1, \vec{e}_2 \rangle = 1 + i.$$

Is it possible to recover the entire inner product just from this information? The answer is yes! Given $\vec{v}, \vec{u} \in \mathbb{C}^2$, we can express them uniquely in terms of the standard basis as

$$\vec{v} = a\vec{e}_1 + b\vec{e}_2 \quad \text{and} \quad \vec{u} = c\vec{e}_1 + d\vec{e}_2 \quad (a, b, c, d \in \mathbb{C}).$$

Then

$$\begin{aligned} \langle \vec{v}, \vec{u} \rangle &= \langle a\vec{e}_1 + b\vec{e}_2, c\vec{e}_1 + d\vec{e}_2 \rangle \\ &= \langle a\vec{e}_1, c\vec{e}_1 \rangle + \langle b\vec{e}_2, c\vec{e}_1 \rangle + \langle a\vec{e}_1, d\vec{e}_2 \rangle + \langle b\vec{e}_2, d\vec{e}_2 \rangle \\ &= a\bar{c}\langle \vec{e}_1, \vec{e}_1 \rangle + b\bar{c}\langle \vec{e}_2, \vec{e}_1 \rangle + a\bar{d}\langle \vec{e}_1, \vec{e}_2 \rangle + b\bar{d}\langle \vec{e}_2, \vec{e}_2 \rangle \\ &= 2a\bar{c} + (1-i)b\bar{c} + (1+i)a\bar{d} + 3b\bar{d}. \end{aligned}$$

This completely defines our inner product. So, just like linear maps are determined by what they do to basis vectors, it appears, at least in this example, that inner products may be determined by what they do to basis vectors.

Recall that the fact that linear maps are determined by the images of the basis vectors is intimately tied to the fact that every linear map can be represented by a matrix (once we've chosen bases for our vector spaces of course). So a natural question arises: can we represent an inner product by a matrix, and does every matrix represent an inner product?

Continuing the example above, let \mathcal{S} be the standard basis for \mathbb{C}^2 . Then $[\vec{v}]_{\mathcal{S}} = \begin{bmatrix} a \\ b \end{bmatrix}$, $[\vec{u}]_{\mathcal{S}} = \begin{bmatrix} c \\ d \end{bmatrix}$ and (forgive us for pulling this identity out of thin air)

$$[\bar{c} \ \bar{d}] \begin{bmatrix} 2 & 1-i \\ 1+i & 3 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = [2a\bar{c} + (1-i)b\bar{c} + (1+i)a\bar{d} + 3b\bar{d}].$$

Hmm! If we identify a 1×1 matrix with an element of \mathbb{C} , the above equation can be written as

$$[\vec{u}]_{\mathcal{S}}^T A [\vec{v}]_{\mathcal{S}} = \langle \vec{v}, \vec{u} \rangle$$

for the 2×2 matrix $A = \begin{bmatrix} 2 & 1-i \\ 1+i & 3 \end{bmatrix}$.

Does this type of expression seem familiar? You might recall that the standard inner product on \mathbb{C}^2 is given by

$$\langle \vec{v}, \vec{u} \rangle = \overline{\vec{u}}^T \vec{v} = \overline{\vec{u}}^T I \vec{v} = [\vec{u}]_{\mathcal{S}}^T I [\vec{v}]_{\mathcal{S}}.$$

Comparing both equations, the difference appears to be that we changed the definition of the inner product on the standard basis, which resulted in the identity matrix changing to this peculiar matrix A .

Let's phrase our motivating question a little better. Let V be an n -dimensional vector space with basis \mathcal{B} .

1. Let $\langle \cdot, \cdot \rangle$ be an inner product on V . Is there an $n \times n$ matrix A such that

$$\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}} \quad \text{for all } \vec{v}, \vec{w} \in V?$$

2. Let A be an $n \times n$ matrix. Does

$$\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}}$$

define an inner product on V ?

To try and get some headway into the problem, let's assume that $\dim(V) = 2$ say with basis $\mathcal{B} = \{\vec{g}_1, \vec{g}_2\}$. Let $\vec{v} = a\vec{g}_1 + b\vec{g}_2$ and $\vec{w} = c\vec{g}_1 + d\vec{g}_2$ be arbitrary vectors in V , so that $[\vec{v}]_{\mathcal{B}} = \begin{bmatrix} a \\ b \end{bmatrix}$ and $[\vec{w}]_{\mathcal{B}} = \begin{bmatrix} c \\ d \end{bmatrix}$.

Suppose there is a 2×2 matrix $A = \begin{bmatrix} w & x \\ y & z \end{bmatrix}$ such that $\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}}$. Then

$$\langle \vec{v}, \vec{w} \rangle = [\bar{c} \ \bar{d}] \begin{bmatrix} w & x \\ y & z \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = [a\bar{c}w + a\bar{d}y + b\bar{c}x + b\bar{d}z].$$

Since

$$\langle \vec{v}, \vec{w} \rangle = a\bar{c}\langle \vec{g}_1, \vec{g}_1 \rangle + a\bar{d}\langle \vec{g}_1, \vec{g}_2 \rangle + b\bar{c}\langle \vec{g}_2, \vec{g}_1 \rangle + b\bar{d}\langle \vec{g}_2, \vec{g}_2 \rangle$$

we must have

$$a\bar{c}w + \bar{a}dy + \bar{b}cx + \bar{b}dz = a\bar{c}\langle \vec{g}_1, \vec{g}_1 \rangle + \bar{a}d\langle \vec{g}_1, \vec{g}_2 \rangle + \bar{b}c\langle \vec{g}_2, \vec{g}_1 \rangle + \bar{b}d\langle \vec{g}_2, \vec{g}_2 \rangle$$

for all $a, b, c, d \in \mathbb{F}$. Thus,

$$\begin{bmatrix} w & x \\ y & z \end{bmatrix} = \begin{bmatrix} \langle \vec{g}_1, \vec{g}_1 \rangle & \langle \vec{g}_2, \vec{g}_1 \rangle \\ \langle \vec{g}_1, \vec{g}_2 \rangle & \langle \vec{g}_2, \vec{g}_2 \rangle \end{bmatrix}.$$

(For instance, if we let $b = d = 0$ and $a = c = 1$, then we find that $w = \langle \vec{g}_1, \vec{g}_1 \rangle$. Similar choices allow us to determine x, y and z .)

For the general case of an n -dimensional inner product space with basis $\mathcal{B} = \{\vec{g}_1, \dots, \vec{g}_n\}$, if there is to be an $n \times n$ matrix $A = [A_{ij}]$ such that $\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}}$ then we must have $A_{ij} = \langle \vec{g}_j, \vec{g}_i \rangle$. Furthermore, since $\langle \vec{v}, \vec{w} \rangle = \overline{\langle \vec{w}, \vec{v} \rangle}$ it follows that $A_{ij} = \overline{A_{ji}}$ (take $\vec{v} = \vec{g}_i$ and $\vec{w} = \vec{g}_j$). That is, it must be the case that $A = A^*$ is self-adjoint. What other features must A have?

Theorem 5.7.1 (Characterization of Inner Products in Terms of Matrices)

Let V be a vector space over \mathbb{F} with basis $\mathcal{B} = \{\vec{g}_1, \dots, \vec{g}_n\}$, and let $A \in M_{n \times n}(\mathbb{F})$. Then

$$\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}} \quad (\vec{v}, \vec{w} \in V)$$

defines an inner product on V if and only if $A = A^*$ and all the eigenvalues of A are positive.

Furthermore, if $\langle \cdot, \cdot \rangle$ is an inner product of V , then there is a self-adjoint matrix A such that

$$\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}} \quad \text{for all } \vec{v}, \vec{w} \in V.$$

Explicitly, this matrix is given by

$$A = \begin{bmatrix} \langle \vec{g}_1, \vec{g}_1 \rangle & \langle \vec{g}_2, \vec{g}_1 \rangle & \cdots & \langle \vec{g}_n, \vec{g}_1 \rangle \\ \vdots & \vdots & \ddots & \vdots \\ \langle \vec{g}_1, \vec{g}_n \rangle & \langle \vec{g}_2, \vec{g}_n \rangle & \cdots & \langle \vec{g}_n, \vec{g}_n \rangle \end{bmatrix}.$$

Proof: First, assume that $\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}}$ defines an inner product on V . Then it must satisfy the inner product axioms:

1. $\langle \vec{v}, \vec{w} \rangle = \overline{\langle \vec{w}, \vec{v} \rangle}$.
2. $\langle \alpha \vec{v}, \vec{w} \rangle = \alpha \langle \vec{v}, \vec{w} \rangle$.
3. $\langle \vec{u} + \vec{v}, \vec{w} \rangle = \langle \vec{u}, \vec{w} \rangle + \langle \vec{v}, \vec{w} \rangle$.
4. (a) $\langle \vec{v}, \vec{v} \rangle \geq 0$.
(b) If $\langle \vec{v}, \vec{v} \rangle = 0$ then $\vec{v} = \vec{0}$.

The discussion preceding the statement of the theorem shows that we must have $A = A^*$. For the claim about the eigenvalues, suppose that $A[\vec{x}]_{\mathcal{B}} = \lambda[\vec{x}]_{\mathcal{B}}$ with $[\vec{x}]_{\mathcal{B}} \neq \vec{0}$ an eigenvector of A . Then by multiplying both sides by $[\vec{x}]_{\mathcal{B}}^*$ we get

$$[\vec{x}]_{\mathcal{B}}^* A [\vec{x}]_{\mathcal{B}} = \lambda [\vec{x}]_{\mathcal{B}}^* [\vec{x}]_{\mathcal{B}}.$$

The left-side is $\langle \vec{x}, \vec{x} \rangle$ which by axiom 4 is positive since $\vec{x} \neq \vec{0}$. The number $[\vec{x}]_{\mathcal{B}}^* [\vec{x}]_{\mathcal{B}}$ on the right-side is the standard inner product of $[\vec{x}]_{\mathcal{B}}$ with itself, hence must also be positive. Thus

$$\lambda = \frac{\langle \vec{x}, \vec{x} \rangle}{[\vec{x}]_{\mathcal{B}}^* [\vec{x}]_{\mathcal{B}}}$$

is positive.

Conversely, assume that $A = A^*$ and that the eigenvalues of A are positive. We want to prove that $\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}}$ satisfies the inner product axioms above. To show that axiom 1 is satisfied, we will “un-abuse” notation and revert back to viewing $[\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}}$ as the 1×1 matrix $[\langle \vec{v}, \vec{w} \rangle]$. Then

$$\begin{aligned} [\langle \vec{v}, \vec{w} \rangle]^* &= ([\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}})^* \\ &= [\vec{v}]_{\mathcal{B}}^* A^* [\vec{w}]_{\mathcal{B}} \\ &= [\vec{v}]_{\mathcal{B}}^* A [\vec{w}]_{\mathcal{B}} \quad (\text{since } A = A^*) \\ &= [\langle \vec{w}, \vec{v} \rangle] \end{aligned}$$

which, since the transpose of a 1×1 matrix is just the matrix itself, implies $\overline{\langle \vec{v}, \vec{w} \rangle} = \langle \vec{w}, \vec{v} \rangle$.

For the other axioms we resume our abuse of notation and identify the 1×1 matrix $[\langle \vec{v}, \vec{w} \rangle]$ with the number $\langle \vec{v}, \vec{w} \rangle$. We have

$$\begin{aligned} \langle \vec{v} + \vec{u}, \vec{w} \rangle &= [\vec{w}]_{\mathcal{B}}^* A [\vec{v} + \vec{u}]_{\mathcal{B}} \\ &= [\vec{w}]_{\mathcal{B}}^* A ([\vec{v}]_{\mathcal{B}} + [\vec{u}]_{\mathcal{B}}) \\ &= [\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}} + [\vec{w}]_{\mathcal{B}}^* A [\vec{u}]_{\mathcal{B}} \\ &= \langle \vec{v}, \vec{w} \rangle + \langle \vec{u}, \vec{w} \rangle. \end{aligned}$$

Similarly,

$$\begin{aligned} \langle \alpha \vec{v}, \vec{w} \rangle &= [\vec{w}]_{\mathcal{B}}^* A [\alpha \vec{v}]_{\mathcal{B}} \\ &= \alpha [\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}} \\ &= \alpha \langle \vec{v}, \vec{w} \rangle. \end{aligned}$$

As usual, the interesting situation is axiom 4. It’s not clear what conditions we need on A so that $[\vec{v}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}} \geq 0$.

Since $A^* = A$, i.e., A is self-adjoint, the spectral theorem (Theorem 5.4.5) tells us there is a unitary matrix $U \in M_{n \times n}(\mathbb{C})$ and a diagonal matrix $D \in M_{n \times n}(\mathbb{R})$ so that $U^* A U = D$. Then

$$\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{B}}^* U D U^* [\vec{v}]_{\mathcal{B}} = (U^* [\vec{w}]_{\mathcal{B}})^* D (U^* [\vec{v}]_{\mathcal{B}}).$$

Since U^* is invertible, we can think of it as a change of basis matrix. In fact, we have $U^* [\vec{v}]_{\mathcal{B}} = [\vec{v}]_{\mathcal{C}}$ for all $\vec{v} \in V$, where \mathcal{C} is some other basis (it’s not important what it is). So we now have

$$\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{C}}^* D [\vec{v}]_{\mathcal{C}}.$$

Remember that D is a diagonal matrix with entries along the diagonal equal to the eigenvalues of A , and these eigenvalues are real and positive (by assumption). So if

$$[\vec{w}]_{\mathcal{C}} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix}, \quad [\vec{v}]_{\mathcal{C}} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}, \quad \text{and} \quad D = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix}.$$

Then

$$\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{C}}^* D [\vec{v}]_{\mathcal{C}} = \lambda_1 v_1 \bar{w}_1 + \cdots + \lambda_n v_n \bar{w}_n.$$

In particular,

$$\langle \vec{v}, \vec{v} \rangle = \lambda_1 |v_1|^2 + \cdots + \lambda_n |v_n|^2.$$

Therefore $\langle \vec{v}, \vec{v} \rangle \geq 0$ since $\lambda_i > 0$ for all i . Furthermore, if $\langle \vec{v}, \vec{v} \rangle = 0$ then necessarily $|v_i|^2 = 0$ for all i , and consequently $\vec{v} = \vec{0}$. This completes the proof that $\langle \vec{v}, \vec{w} \rangle$ is an inner product.

The final assertion in the theorem was proved in the discussion immediately preceding the statement of the theorem. \square

Let's give the matrix A in the previous theorem a name.

Definition 5.7.2

Gram Matrix

Let V be a finite-dimensional inner product space over \mathbb{F} with inner product $\langle \cdot, \cdot \rangle$ and basis $\mathcal{B} = \{\vec{g}_1, \dots, \vec{g}_n\}$. The **Gram matrix of $\langle \cdot, \cdot \rangle$ with respect to \mathcal{B}** is the matrix whose (i, j) th entry is $A_{ij} = \langle \vec{g}_j, \vec{g}_i \rangle$. That is,

$$A = \begin{bmatrix} \langle \vec{g}_1, \vec{g}_1 \rangle & \cdots & \langle \vec{g}_n, \vec{g}_1 \rangle \\ \vdots & \ddots & \vdots \\ \langle \vec{g}_1, \vec{g}_n \rangle & \cdots & \langle \vec{g}_n, \vec{g}_n \rangle \end{bmatrix}.$$

Theorem 5.7.1 shows that the Gram matrix can be used to compute the inner product via the formula

$$\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{B}}^* A [\vec{v}]_{\mathcal{B}}.$$

Example 5.7.3

In \mathbb{F}^n with the standard inner product, if we let $\mathcal{B} = \{\vec{e}_1, \dots, \vec{e}_n\}$ be the standard basis, then the corresponding Gram matrix is

$$\begin{bmatrix} \langle \vec{e}_1, \vec{e}_1 \rangle & \cdots & \langle \vec{e}_n, \vec{e}_1 \rangle \\ \vdots & \ddots & \vdots \\ \langle \vec{e}_1, \vec{e}_n \rangle & \cdots & \langle \vec{e}_n, \vec{e}_n \rangle \end{bmatrix} = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & 1 \end{bmatrix},$$

i.e. it's the $n \times n$ identity matrix! This re-establishes our familiar formula

$$\langle \vec{v}, \vec{w} \rangle = \vec{w}^* I_n \vec{v} = \vec{w}^* \vec{v}$$

for the standard inner product on \mathbb{F}^n .

More generally, if V is any n -dimensional inner product space, and if \mathcal{B} is an *orthonormal* basis for V , then the corresponding Gram matrix is the $n \times n$ identity matrix. This means that we can compute the inner product on V as though it were the standard inner product on \mathbb{F}^n (once we convert to \mathcal{B} -coordinates):

$$\langle \vec{v}, \vec{w} \rangle = [\vec{w}]_{\mathcal{B}}^* [\vec{v}]_{\mathcal{B}}.$$

Another reason why orthonormal bases are nice!

Exercise 72

Let V be an n -dimensional inner product space and let \mathcal{B} be an *orthonormal* basis for V .

- (a) Show that the Gram matrix with respect to \mathcal{B} is I_n .
- (b) What does the Gram matrix with respect to an orthogonal basis look like?

As you may have experienced, sometimes it is difficult to prove or disprove that a potential inner product satisfies axiom 4 (positive-definiteness) of Definition 4.1.1. Theorem 5.7.1 gives us an algorithmic way to determine if a potential inner product is actually an inner product. We simply choose a basis, compute the corresponding potential Gram matrix A , check that A is self-adjoint, verify that A actually performs the potential inner product for us, and then determine if the eigenvalues of A are all positive.

Example 5.7.4

Is

$$\langle a + bx, c + dx \rangle = 2a\bar{c} + (1 + i)b\bar{c} + (1 - i)a\bar{d} + 3b\bar{d}$$

an inner product on $\mathcal{P}_1(\mathbb{C})$? Let's find out.

Let \mathcal{B} be the standard basis for $\mathcal{P}_1(\mathbb{C})$. The corresponding Gram matrix is given by

$$A = \begin{bmatrix} \langle 1, 1 \rangle & \langle x, 1 \rangle \\ \langle 1, x \rangle & \langle x, x \rangle \end{bmatrix} = \begin{bmatrix} 2 & 1 + i \\ 1 - i & 3 \end{bmatrix}.$$

First thing we need to check is that this matrix is self-adjoint, which it is. Now we need to check that it actually performs the inner product for us. That is, we need to check if $\langle p, q \rangle = [q]_{\mathcal{B}}^* A [p]_{\mathcal{B}}$. We have

$$[c + dx]_{\mathcal{B}}^* A [a + bx]_{\mathcal{B}} = [\bar{c} \ \bar{d}] \begin{bmatrix} 2 & 1 + i \\ 1 - i & 3 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = [2a\bar{c} + (1 + i)b\bar{c} + (1 - i)a\bar{d} + 3b\bar{d}].$$

Therefore this matrix does the trick! So, to check whether or not it's an inner product, we need to compute the eigenvalues and make sure they're all positive. We have

$$\begin{aligned} |A - \lambda I| &= \begin{vmatrix} 2 - \lambda & 1 + i \\ 1 - i & 3 - \lambda \end{vmatrix} \\ &= (2 - \lambda)(3 - \lambda) - (1 + i)(1 - i) \\ &= \lambda^2 - 5\lambda + 4 \\ &= (\lambda - 4)(\lambda - 1). \end{aligned}$$

Since the eigenvalues are 1 and 4, which are both positive, Theorem 5.7.1 allows us to conclude that this is indeed an inner product.

Example 5.7.5

Let's determine whether or not

$$\left\langle \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix}, \begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix} \right\rangle = 2v_1w_1 + 2v_2w_2 + 2v_3w_3 - v_1w_2 - v_2w_1 - v_2w_3 - v_3w_2$$

is an inner product on \mathbb{R}^3 .

Let $\{\vec{e}_1, \vec{e}_2, \vec{e}_3\}$ be the standard basis for \mathbb{R}^3 . Our candidate matrix for this potential inner product is

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

This matrix is certainly self-adjoint, which is a good start. Furthermore we have

$$[w_1 \ w_2 \ w_3] \begin{bmatrix} 2 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = [2v_1w_1 + 2v_2w_2 + 2v_3w_3 - v_1w_2 - v_2w_1 - v_2w_3 - v_3w_2]$$

so the matrix does actually perform the inner product for us. Finally, you can compute its eigenvalues and get them to be $\{2, 2 + \sqrt{2}, 2 - \sqrt{2}\}$, all of which are positive! Therefore this is an inner product.

Exercise 73

Determine whether

$$\langle \vec{z}, \vec{w} \rangle = z_1 \overline{w_1} - iz_1 \overline{w_2} + iz_2 \overline{w_1}$$

is an inner product on \mathbb{C}^2 .

Chapter 6

The Singular Value Decomposition

6.1 Singular Values and Singular Vectors

In Chapter 3 we studied the problem of finding a simple matrix representation of a given linear operator $L: V \rightarrow V$ on a finite-dimensional vector space V . We learned that, if certain conditions are satisfied by L , we're able to find a basis \mathcal{D} for V consisting of eigenvectors of L such that $[L]_{\mathcal{D}}$ is a diagonal matrix. Unfortunately, since the needed conditions are not always met, many linear operators are left without simple matrix representations.

In this chapter we aim to rectify this situation. What made diagonalization difficult to attain is our insistence on using the same basis \mathcal{D} for both the domain and codomain of $L: V \rightarrow V$. Let's instead ask for two (possibly different) bases \mathcal{B} and \mathcal{C} of V so that ${}_c[L]_{\mathcal{B}}$ is diagonal. Are we always able to find these? The answer is *yes!* Actually, this is quite easy to do—but what is more interesting is that, if V is an inner product space, we can arrange for \mathcal{B} and \mathcal{C} to be orthonormal bases. Additionally, all this can be done not just for operators, but for linear maps $L: V \rightarrow W$ as well! (Of course, in this case ${}_c[L]_{\mathcal{B}}$ won't be a square matrix, so we'll need to redefine what we mean by “diagonal matrix”—but you can probably guess what the definition is going to be.)

Although this seems like it would be only of theoretical interest, this is very far from the truth. The results of this chapter in fact have an abundance of practical, real-world applications—see Section 6.4.

We will begin by seeing how all this works for a matrix $A \in M_{m \times n}(\mathbb{F})$, which we think of as giving a linear map $\mathbb{F}^n \rightarrow \mathbb{F}^m$. We give \mathbb{F}^n and \mathbb{F}^m their standard inner product. The key to our approach to diagonalization of square matrices (and linear operators) was the idea of eigenvectors and eigenvalues. Our starting point here will be to find a suitable analogue for non-square matrices; these are the so-called *singular vectors* and *singular values* of A , to be introduced momentarily. First we need a preliminary lemma.

Lemma 6.1.1

Let $A \in M_{m \times n}(\mathbb{F})$. Then A^*A is an $n \times n$ self-adjoint matrix and its eigenvalues are non-negative real numbers.

Proof: It's clear that A^*A is $n \times n$ and self-adjoint. Let λ be an eigenvalue of A^*A , say with eigenvector \vec{x} . Then $A^*A\vec{x} = \lambda\vec{x}$, and therefore

$$\lambda\|\vec{x}\|^2 = \lambda\langle\vec{x}, \vec{x}\rangle = \langle\lambda\vec{x}, \vec{x}\rangle = \langle A^*A\vec{x}, \vec{x}\rangle = \langle A\vec{x}, A\vec{x}\rangle = \|A\vec{x}\|^2.$$

Since \vec{x} is an eigenvector, $\|\vec{x}\| \neq 0$, and consequently $\lambda = \frac{\|A\vec{x}\|^2}{\|\vec{x}\|^2} \geq 0$. \square

Definition 6.1.2**Singular Values,
Singular Vectors**

Let $A \in M_{m \times n}(\mathbb{F})$. The **singular values** of A are the non-negative square-roots $\sigma_i = \sqrt{\lambda_i}$ of the eigenvalues λ_i of A^*A .

The corresponding eigenvectors of A^*A are called the **singular vectors** of A .

REMARK

The name *singular value* originates in the theory of integral equations, and was coined by Emile Picard for a value that is of special (or *singular*) interest. It has nothing to do with our modern mathematical usage of the word “singular.”

By convention, the singular values are always ordered in descending order:

$$\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n \geq 0.$$

At this stage it's not at all clear what singular values and vectors have to do with our motivating problem. The connection is explained in the discussion preceding Theorem 6.2.1 in the next section. For now, however, let's see some examples.

Example 6.1.3

Let $A = \begin{bmatrix} 1 & -4 \\ -2 & 2 \\ 2 & 4 \end{bmatrix}$. Then

$$A^*A = A^T A = \begin{bmatrix} 1 & -2 & 2 \\ -4 & 2 & 4 \end{bmatrix} \begin{bmatrix} 1 & -4 \\ -2 & 2 \\ 2 & 4 \end{bmatrix} = \begin{bmatrix} 9 & 0 \\ 0 & 36 \end{bmatrix}.$$

(Notice that this matrix is symmetric, as expected.)

The eigenvalues of A^*A are $\lambda_1 = 36$ and $\lambda_2 = 9$. (Notice that they are non-negative, as expected.) Thus the singular values of A are $\sigma_1 = \sqrt{36} = 6$ and $\sigma_2 = \sqrt{9} = 3$.

The vectors $\vec{v}_1 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ and $\vec{v}_2 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ are corresponding singular vectors, since they are eigenvectors for A^*A with eigenvalues $\lambda_1 = 36$ and $\lambda_2 = 9$, respectively.

Example 6.1.4

Let $A = \begin{bmatrix} 1 & 1 & 0 \\ -1 & 1 & 1 \end{bmatrix}$. Then

$$A^*A = A^T A = \begin{bmatrix} 1 & -1 \\ 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & 0 \\ -1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 2 & 0 & -1 \\ 0 & 2 & 1 \\ -1 & 1 & 1 \end{bmatrix}.$$

You can check that the eigenvalues of A^*A are $\lambda_1 = 3$, $\lambda_2 = 2$ and $\lambda_3 = 0$. Thus the singular values of A are $\sigma_1 = \sqrt{3}$, $\sigma_2 = \sqrt{2}$ and $\sigma_3 = 0$. Corresponding singular vectors are $\vec{v}_1 = \begin{bmatrix} -1 \\ 1 \\ 1 \end{bmatrix}$, $\vec{v}_2 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$ and $\vec{v}_3 = \begin{bmatrix} 1 \\ -1 \\ 2 \end{bmatrix}$, respectively.

Example 6.1.5

Let $A = \begin{bmatrix} 0 & 2 \\ -3 & 0 \end{bmatrix}$. Then

$$A^*A = A^T A = \begin{bmatrix} 0 & -3 \\ 2 & 0 \end{bmatrix} \begin{bmatrix} 0 & 2 \\ -3 & 0 \end{bmatrix} = \begin{bmatrix} 9 & 0 \\ 0 & 4 \end{bmatrix}.$$

The eigenvalues of A^*A are $\lambda_1 = 9$ and $\lambda_2 = 4$. Thus the singular values of A are $\sigma_1 = 3$ and $\sigma_2 = 2$, and $\vec{v}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $\vec{v}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ are corresponding singular vectors.

Example 6.1.6

Let $A = \begin{bmatrix} 1 & 2 & -1 \\ 2 & 4 & -2 \\ -1 & -2 & 1 \\ i & 2i & -i \end{bmatrix}$. Then

$$A^*A = \begin{bmatrix} 1 & 2 & -1 & -i \\ 2 & 4 & -2 & -2i \\ -1 & -2 & 1 & i \\ i & 2i & -i \end{bmatrix} \begin{bmatrix} 1 & 2 & -1 \\ 2 & 4 & -2 \\ -1 & -2 & 1 \\ i & 2i & -i \end{bmatrix} = \begin{bmatrix} 7 & 14 & -7 \\ 14 & 28 & -14 \\ -7 & -14 & 7 \end{bmatrix}.$$

The eigenvalues of A^*A are $\lambda_1 = 42$ and $\lambda_2 = \lambda_3 = 0$. For eigenvectors, we can take $\vec{v}_1 = \begin{bmatrix} -1 \\ -2 \\ 1 \end{bmatrix}$, $\vec{v}_2 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$ and $\vec{v}_3 = \begin{bmatrix} -1 \\ 1 \\ 1 \end{bmatrix}$, respectively. Here we decided to pick two orthogonal eigenvectors \vec{v}_2 and \vec{v}_3 for the repeated eigenvalue 0. We are able to do this because A^*A is symmetric, hence orthogonally diagonalizable by the spectral theorem.

The singular values of A are $\sigma_1 = \sqrt{42}$ and $\sigma_2 = \sigma_3 = 0$, with corresponding singular vectors \vec{v}_1 , \vec{v}_2 and \vec{v}_3 as above.

Exercise 74

Let $A \in M_{n \times n}(\mathbb{F})$ be a self-adjoint matrix. What are the singular values of A ? What can you say about the corresponding singular vectors?

There are two things we can notice from the previous examples:

1. The singular vectors of A corresponding to different singular values are orthogonal.
2. The rank of A is equal to the number of non-zero singular values of A .

These resemble two facts that we know about eigenvalues and eigenvectors of a square matrix A . The first one is that eigenvectors corresponding to distinct eigenvalues are linearly independent (Proposition 3.2.11). The second one is that the nullity of A is equal to the geometric multiplicity of 0 as an eigenvalue. Consequently, if A is diagonalizable, $\text{nullity}(A)$ is equal to the algebraic multiplicity of 0, and so $\text{rank}(A)$ is the sum of the algebraic multiplicities of the non-zero eigenvalues which we can interpret as being the number of non-zero eigenvalues, each counted according to multiplicity.

Let's now prove that our two observations above are in fact always true.

Proposition 6.1.7

Let $A \in M_{m \times n}(\mathbb{F})$, and let $\vec{x}, \vec{y} \in \mathbb{F}^n$ be singular vectors of A corresponding to the singular values σ_1 and σ_2 . If $\sigma_1 \neq \sigma_2$, then \vec{x} and \vec{y} are orthogonal.

Proof: Let $B = A^*A$ and let $\lambda_i = \sigma_i^2$. Notice that $\lambda_1 \neq \lambda_2$ since otherwise we'd have $\sigma_1 = \pm\sigma_2$ hence $\sigma_1 = \sigma_2$ since both are non-negative. Thus, \vec{x} and \vec{y} are eigenvectors of B with distinct eigenvalues λ_1 and λ_2 , respectively. Since B is self-adjoint, hence normal, the orthogonality of \vec{x} and \vec{y} follows from Proposition 5.4.10(c). \square

To prove our second observation, we will need the following lemma.

Lemma 6.1.8

Let $A \in M_{m \times n}(\mathbb{F})$. Then $\text{Null}(A^*A) = \text{Null}(A)$.

Proof: Suppose that $\vec{x} \in \text{Null}(A)$, so that $A\vec{x} = \vec{0}$. Then $A^*A\vec{x} = A^*\vec{0} = \vec{0}$, so $\vec{x} \in \text{Null}(A^*A)$. Thus, $\text{Null}(A) \subseteq \text{Null}(A^*A)$. Conversely, suppose that $\vec{x} \in \text{Null}(A^*A)$. Then

$$\|A\vec{x}\|^2 = \langle A\vec{x}, A\vec{x} \rangle = \langle \vec{x}, A^*A\vec{x} \rangle = \langle \vec{x}, \vec{0} \rangle = 0$$

so that $A\vec{x} = \vec{0}$, proving that $\vec{x} \in \text{Null}(A)$ and hence that $\text{Null}(A^*A) \subseteq \text{Null}(A)$. This completes the proof. \square

Proposition 6.1.9

Let $A \in M_{m \times n}(\mathbb{F})$. The number of non-zero singular values of A is equal to $\text{rank}(A)$, where each repeated singular value is counted according to its multiplicity. (The *multiplicity* of a singular value σ is the algebraic multiplicity of σ^2 as an eigenvalue of A^*A .)

Proof: Since A^*A is self-adjoint, hence diagonalizable by the spectral theorem, the argument preceding Proposition 6.1.7 shows that $\text{rank}(A^*A)$ is equal to the number of non-zero eigenvalues of A^*A , which in turn is equal to the number of singular values of A , each counted according to multiplicity. (Here we're again using the fact that distinct singular values of A come from distinct eigenvalues of A^*A , and vice versa, as observed in the proof of Proposition 6.1.7.)

So the proof will be complete if we can show that $\text{rank}(A) = \text{rank}(A^*A)$. Since A is $m \times n$ and A^*A is $n \times n$, the rank-nullity theorem and Lemma 6.1.8 give

$$\text{rank}(A) = n - \text{nullity}(A) = n - \text{nullity}(A^*A) = \text{rank}(A^*A),$$

as desired. \square

In particular, an $m \times n$ matrix will have at most $\min\{m, n\}$ non-zero singular values.

6.2 Singular Value Decomposition of Matrices

In this section we will show that we can always diagonalize a matrix $A \in M_{m \times n}(\mathbb{F})$ by choosing appropriate orthonormal bases for \mathbb{F}^n and \mathbb{F}^m . If we let U and V be the matrices whose columns are these basis vectors, then this amounts to a factorization of A of the form

$$A = U\Sigma V^*,$$

where U is an $m \times m$ unitary matrix, V is an $n \times n$ unitary matrix, and where Σ is an $m \times n$ matrix whose (i, j) th entry is 0 for $i \neq j$.

The idea behind the proof is really simple: we reverse engineer what we want. That is, suppose we know that $A = U\Sigma V^*$. Then we'd also have that $A^* = V\Sigma^*U^*$ and therefore

$$A^*A = V\Sigma^*U^*U\Sigma V^* = V\Sigma^*\Sigma V^*.$$

The matrix $\Sigma^*\Sigma$ is a square diagonal matrix whose diagonal entries are non-negative real numbers. So, we conclude that **if** $A = U\Sigma V^*$, then we'd be able to unitarily diagonalize A^*A with non-negative real eigenvalues. But since A^*A is self-adjoint, we *know* that this is possible thanks to the spectral theorem (Theorem 5.4.5) and Lemma 6.1.1. So this suggests what to take for the diagonal entries of Σ and for V : namely, the positive square roots of the eigenvalues of A^*A (i.e. the singular values of A), and the same V that unitarily diagonalizes A^*A (i.e. the corresponding singular vectors)!

By applying the same analysis to $AA^* = U\Sigma\Sigma^*U^*$, we know what we must take for U . Now all that remains is to show that if we do take these U , Σ and V , then we get our desired decomposition $A = U\Sigma V^*$. The proof below will basically do this, except it will construct U directly because this makes the verification easier.

Theorem 6.2.1 (Singular Value Decomposition of Matrices)

Let $A \in M_{m \times n}(\mathbb{F})$ be a matrix of rank r with non-zero singular values $\sigma_1 \geq \dots \geq \sigma_r > 0$. Then there exist unitary matrices $U \in M_{m \times m}(\mathbb{F})$ and $V \in M_{n \times n}(\mathbb{F})$ such that

$$A = U\Sigma V^*,$$

where Σ is the $m \times n$ matrix whose entries are

$$\Sigma_{ij} = \begin{cases} \sigma_i & \text{if } i = j \leq r \\ 0 & \text{otherwise.} \end{cases}$$

If $A \in M_{m \times n}(\mathbb{R})$ is real, then U and V can be chosen to be orthogonal matrices.

Proof: The $n \times n$ self-adjoint matrix A^*A is unitarily diagonalizable, by the spectral theorem. So we can find an orthonormal basis $\{\vec{v}_1, \dots, \vec{v}_n\}$ for \mathbb{F}^n consisting of eigenvectors of A^*A with corresponding eigenvalues $\lambda_1 \geq \dots \geq \lambda_n$ (which are all real, by the spectral theorem, so it is possible to order them like this). Then $V = [\vec{v}_1 \cdots \vec{v}_n]$ is a unitary matrix and

$$AV = [A\vec{v}_1 \cdots A\vec{v}_n].$$

According to Proposition 6.1.9, $\lambda_i = \sigma_i^2 = 0$ for all $i > r$. Thus the vectors \vec{v}_i for $i > r$ are all in $\text{Null}(A^*A)$, hence in $\text{Null}(A)$ by Lemma 6.1.8. So

$$AV = [A\vec{v}_1 \cdots A\vec{v}_r \ \vec{0} \cdots \vec{0}]$$

Let $\vec{u}_i = \frac{1}{\sigma_i} A \vec{v}_i$ for $i \leq r$ and notice that, since σ_i is real and positive,

$$\langle \vec{u}_i, \vec{u}_j \rangle = \frac{1}{\sigma_i \sigma_j} \langle A \vec{v}_i, A \vec{v}_j \rangle = \frac{1}{\sigma_i \sigma_j} \langle \vec{v}_i, A^* A \vec{v}_j \rangle = \frac{1}{\sigma_i \sigma_j} \langle \vec{v}_i, \sigma_j^2 \vec{v}_j \rangle = \frac{\sigma_j}{\sigma_i} \langle \vec{v}_i, \vec{v}_j \rangle.$$

Thus, $\langle \vec{u}_i, \vec{u}_j \rangle$ is 0 for $i \neq j$ and is 1 otherwise, since $\{\vec{v}_1, \dots, \vec{v}_n\}$ is an orthonormal set. It follows that $\{\vec{u}_1, \dots, \vec{u}_r\}$ is an orthonormal set in \mathbb{F}^m . Extend it to an orthonormal basis $\{\vec{u}_1, \dots, \vec{u}_m\}$. Then $U = [\vec{u}_1 \cdots \vec{u}_m]$ is unitary and we have

$$\begin{aligned} AV &= [A \vec{v}_1 \cdots A \vec{v}_r \ \vec{0} \cdots \vec{0}] \\ &= [\sigma_1 \vec{u}_1 \cdots \sigma_r \vec{u}_r \ \vec{0} \cdots \vec{0}] \\ &= [\vec{u}_1 \cdots \vec{u}_r \ \vec{u}_{r+1} \cdots \vec{u}_m] \begin{bmatrix} \text{diag}(\sigma_1, \dots, \sigma_r) & 0_{r \times (n-r)} \\ 0_{(m-r) \times r} & 0_{(n-r) \times (m-r)} \end{bmatrix} \\ &= U \Sigma. \end{aligned}$$

Multiplying both sides by V^* on the right, we get our desired decomposition $A = U \Sigma V^*$. \square

Exercise 75

By examining the proof of Theorem 6.2.1, show that if $A \in M_{m \times n}(\mathbb{R})$ is real, then we can choose U and V to be orthogonal.

Definition 6.2.2

Singular Value
Decomposition,
SVD

A decomposition $A = U \Sigma V^*$ of the type occurring in Theorem 6.2.1 is called a **singular value decomposition (SVD)** of A .

Although the entries of Σ are uniquely determined by A (they are its singular values), there is generally quite a bit of freedom in choosing U and V . For instance, if $A = I_n$ is the $n \times n$ identity matrix, then $A = U I_n U^*$ will be an SVD of A for *any* unitary $n \times n$ unitary matrix U . Thus, an SVD is not unique.

If we let $D = \text{diag}(\sigma_1, \dots, \sigma_r) \in M_{r \times r}(\mathbb{F})$, we see that the matrix Σ in an SVD of A takes one of the following shapes, depending on $r = \text{rank}(A)$, m and n :

- If $r < \min\{m, n\}$, then $\Sigma = \begin{bmatrix} D & 0_{r \times (n-r)} \\ 0_{(m-r) \times r} & 0_{(n-r) \times (m-r)} \end{bmatrix}$.
- If $r = m < n$, then $\Sigma = [D \ 0_{r \times (n-r)}]$.
- If $r = n < m$, then $\Sigma = \begin{bmatrix} D \\ 0_{(m-r) \times r} \end{bmatrix}$.
- If $r = m = n$, then $\Sigma = D$.

We will illustrate each of these scenarios below.

First, let's note that the proof of Theorem 6.2.1 describes an algorithm for constructing an SVD for a given $m \times n$ matrix A .

ALGORITHM (Finding an SVD for a Matrix)

Let $A \in M_{m \times n}(\mathbb{F})$ be a matrix of rank r . To find an SVD for A :

1. Find the eigenvalues $\lambda_1 \geq \cdots \geq \lambda_n$ and a corresponding set of **orthonormal** eigenvectors $\{\vec{v}_1, \dots, \vec{v}_n\}$ for A^*A . (See the algorithm “Unitary Diagonalization of a Normal Matrix” in Section 5.4.)
2. Set $\sigma_i = \sqrt{\lambda_i}$ for $i \leq r$.
3. Set $\vec{u}_i = \frac{1}{\sigma_i} A \vec{v}_i$ for $i \leq r$. If $r < m$, extend $\{\vec{u}_1, \dots, \vec{u}_r\}$ to an orthonormal basis $\{\vec{u}_1, \dots, \vec{u}_r, \vec{u}_{r+1}, \dots, \vec{u}_m\}$ of \mathbb{F}^m .
4. Set $V = [\vec{v}_1 \cdots \vec{v}_n]$ and $U = [\vec{u}_1 \cdots \vec{u}_m]$, and let Σ be the $m \times n$ matrix whose entries are

$$\Sigma_{ij} = \begin{cases} \sigma_i & \text{if } i = j \leq r \\ 0 & \text{otherwise.} \end{cases}$$

Then U and V are unitary square matrices (orthogonal matrices if A is real) and

$$A = U\Sigma V^*.$$

Notice that in the above algorithm we will have $A\vec{v}_i = \sigma_i\vec{u}_i$ for all $i \leq r$ (and even for $r < i \leq n$, since in that case $\sigma_i = 0$ and $A\vec{v}_i = \vec{0}$ by Lemma 6.1.8), which should remind you of eigenvectors and eigenvalues. See Section 6.2.1 for a geometric interpretation.

To facilitate Step 3 of the algorithm, let’s observe that the vectors $\vec{u}_1, \dots, \vec{u}_r$ form a set of $r = \text{rank}(A)$ linearly independent vectors in $\text{Col}(A)$, thus they must be an orthonormal basis for $\text{Col}(A)$. It follows that $\{\vec{u}_{r+1}, \dots, \vec{u}_m\}$ must be an orthonormal basis for $\text{Col}(A)^\perp$. Here is a handy result, which tells us that one way of finding the vectors $\{\vec{u}_{r+1}, \dots, \vec{u}_m\}$ is by taking them to be an orthonormal basis for $\text{Null}(A^*)$.

Proposition 6.2.3

Let $A \in M_{m \times n}(\mathbb{F})$. Then $\text{Col}(A)^\perp = \text{Null}(A^*)$.

Proof: The vectors in $\text{Col}(A)$ are those of the form $A\vec{x}$, where $\vec{x} \in \mathbb{F}^n$ is arbitrary. So a vector $\vec{y} \in \mathbb{F}^m$ will be in $\text{Col}(A)^\perp$ if and only if

$$0 = \langle A\vec{x}, \vec{y} \rangle = \langle \vec{x}, A^*\vec{y} \rangle \quad \text{for all } \vec{x} \in \mathbb{F}^n,$$

which is the case if and only if $A^*\vec{y} \in (\mathbb{F}^n)^\perp = \{\vec{0}\}$, i.e., if and only if $\vec{y} \in \text{Null}(A^*)$, proving the proposition. \square

Let’s now compute SVDs for the matrices in Examples 6.1.3—6.1.6.

Example 6.2.4

Let $A = \begin{bmatrix} 1 & -4 \\ -2 & 2 \\ 2 & 4 \end{bmatrix}$. The singular values of A were found to be $\sigma_1 = 6$ and $\sigma_2 = 3$, with corresponding singular vectors $\vec{v}_1 = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}$ and $\vec{v}_2 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$ (which are unit vectors). Now let

$$\vec{u}_1 = \frac{1}{\sigma_1} A \vec{v}_1 = \frac{1}{6} \begin{bmatrix} -4 \\ 2 \\ 4 \end{bmatrix} = \frac{1}{3} \begin{bmatrix} -2 \\ 1 \\ 2 \end{bmatrix} \quad \text{and} \quad \vec{u}_2 = \frac{1}{\sigma_2} A \vec{v}_2 = \frac{1}{3} \begin{bmatrix} 1 \\ -2 \\ 2 \end{bmatrix}.$$

Notice that $\{\vec{u}_1, \vec{u}_2\}$ is an orthonormal set in \mathbb{F}^3 (as guaranteed by the proof of Theorem 6.2.1). We must extend it to an orthonormal basis for \mathbb{F}^3 . This can be achieved in a variety of ways. For instance, we can take the cross product of \vec{u}_1 and \vec{u}_2 or we can find an orthonormal basis for $\text{Null}(A^*)$. In any case, we obtain $\vec{u}_3 = \frac{1}{3} \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}$.

Finally, set

$$V = [\vec{v}_1 \ \vec{v}_2] = \begin{bmatrix} 0 & 1 \\ 1 & 0 \\ 1 & 0 \end{bmatrix} \quad \text{and} \quad U = [\vec{u}_1 \ \vec{u}_2 \ \vec{u}_3] = \frac{1}{3} \begin{bmatrix} -2 & 1 & 2 \\ 1 & -2 & 2 \\ 2 & 2 & 1 \end{bmatrix}$$

and notice that these matrices are orthogonal. The resulting SVD of A is:

$$A = U \begin{bmatrix} 6 & 0 \\ 0 & 3 \\ 0 & 0 \end{bmatrix} V^T.$$

Example 6.2.5

Let $A = \begin{bmatrix} 1 & 1 & 0 \\ -1 & 1 & 1 \end{bmatrix}$. The singular values of A were found to be $\sigma_1 = \sqrt{3}$, $\sigma_2 = \sqrt{2}$ and $\sigma_3 = 0$, with corresponding singular vectors $\begin{bmatrix} -1 \\ 1 \\ 1 \end{bmatrix}$, $\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 1 \\ -1 \\ 2 \end{bmatrix}$, respectively.

Let $\vec{v}_1 = \frac{1}{\sqrt{3}} \begin{bmatrix} -1 \\ 1 \\ 1 \end{bmatrix}$, $\vec{v}_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$ and $\vec{v}_3 = \frac{1}{\sqrt{6}} \begin{bmatrix} 1 \\ -1 \\ 2 \end{bmatrix}$, so that $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ is orthonormal. Notice that $r = \text{rank}(A) = 2$, since there are only two non-zero singular values. Next, let

$$\vec{u}_1 = \frac{1}{\sigma_1} A \vec{v}_1 = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \quad \text{and} \quad \vec{u}_2 = \frac{1}{\sigma_2} A \vec{v}_2 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}.$$

Then $\{\vec{u}_1, \vec{u}_2\}$ is already a basis for \mathbb{F}^2 and we can construct the orthogonal matrices

$$V = \frac{1}{\sqrt{6}} \begin{bmatrix} -\sqrt{2} & \sqrt{3} & 1 \\ \sqrt{2} & \sqrt{3} & -1 \\ \sqrt{2} & 0 & 2 \end{bmatrix} \quad \text{and} \quad U = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

The resulting SVD of A is

$$A = U \begin{bmatrix} \sqrt{3} & 0 & 0 \\ 0 & \sqrt{2} & 0 \end{bmatrix} V^T.$$

Example 6.2.6

Let $A = \begin{bmatrix} 0 & 2 \\ -3 & 0 \end{bmatrix}$. The singular values of A were found to be $\sigma_1 = 3$ and $\sigma_2 = 2$, with corresponding singular vectors $\vec{v}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $\vec{v}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$, respectively. These are unit vectors, so we may set

$$\vec{u}_1 = \frac{1}{\sigma_1} A \vec{v}_1 = \begin{bmatrix} 0 \\ -1 \end{bmatrix} \quad \text{and} \quad \vec{u}_2 = \frac{1}{\sigma_2} A \vec{v}_2 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

Thus, in this case our orthogonal matrices are

$$V = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad U = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

and the resulting SVD of A is

$$A = U \begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix} V^T.$$

Interestingly, while A is **not** diagonalizable over \mathbb{R} , it is diagonalizable over \mathbb{C} (with distinct eigenvalues $\pm\sqrt{-6}$). However, A is not unitarily diagonalizable since it is not normal, as you can check.

Example 6.2.7

Let $A = \begin{bmatrix} 1 & 2 & -1 \\ 2 & 4 & -2 \\ -1 & -2 & 1 \\ i & 2i & -i \end{bmatrix}$. The singular values of A were found to be $\sigma_1 = \sqrt{42}$ and $\sigma_2 =$

$\sigma_3 = 0$ with corresponding orthonormal eigenvectors $\vec{v}_1 = \frac{1}{\sqrt{6}} \begin{bmatrix} -1 \\ -2 \\ 1 \end{bmatrix}$, $\vec{v}_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$

and $\vec{v}_3 = \frac{1}{\sqrt{3}} \begin{bmatrix} -1 \\ 1 \\ 1 \end{bmatrix}$, respectively. Here there is only one non-zero singular value, and correspondingly we set

$$\vec{u}_1 = \frac{1}{\sigma_1} A \vec{v}_1 = \frac{1}{\sqrt{7}} \begin{bmatrix} -1 \\ -2 \\ 1 \\ -i \end{bmatrix}.$$

Now we must extend $\{\vec{u}_1\}$ to an orthonormal basis for \mathbb{C}^4 . We can do so by applying the Gram–Schmidt process to $\{\vec{u}_1, \vec{e}_1, \vec{e}_2, \vec{e}_3, \vec{e}_4\}$, where the \vec{e}_i are the standard basis vectors for \mathbb{C}^4 , or by finding an orthonormal basis for $\text{Null}(A^*)$. In either case it's going to

be a tedious computation! We'll spare you the boring details and just give the end result:

$$\vec{u}_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} i \\ 0 \\ 0 \\ 1 \end{bmatrix}, \quad \vec{u}_3 = \frac{1}{\sqrt{6}} \begin{bmatrix} 1 \\ 0 \\ 2 \\ i \end{bmatrix} \quad \text{and} \quad \vec{u}_4 = \frac{1}{\sqrt{21}} \begin{bmatrix} -2 \\ 3 \\ 2 \\ -2i \end{bmatrix}.$$

We thus have our unitary matrices $V = [\vec{v}_1 \ \vec{v}_2 \ \vec{v}_3]$ and $U = [\vec{u}_1 \ \vec{u}_2 \ \vec{u}_3 \ \vec{u}_4]$ with corresponding SVD

$$A = U \begin{bmatrix} \sqrt{42} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} V^*.$$

Exercise 76

Let $A \in M_{n \times n}(\mathbb{F})$ be a self-adjoint matrix. How does an SVD $A = U\Sigma V^*$ compare to a unitary diagonalization $A = WDW^*$?


6.2.1 The Geometry of the Singular Value Decomposition

In this section we will work over $\mathbb{F} = \mathbb{R}$.

If we diagonalize a square matrix $A \in M_{n \times n}(\mathbb{R})$ to get $A = PDP^{-1}$, then we can interpret this factorization as follows. The matrix A sends the vector $\vec{x} \in \mathbb{R}^n$ to the vector $A\vec{x} = PDP^{-1}\vec{x}$. The matrix P^{-1} changes coordinates from the standard basis for \mathbb{R}^n to our basis of eigenvectors, then the matrix D scales the coordinate of each eigenvector by the corresponding eigenvalue, and finally P converts the result back to the standard basis.

If we can *orthogonally* diagonalize A , then P and $P^{-1} = P^T$ are still change of basis matrices, but they change bases from the one orthonormal basis to another. In a sense, they are built up of rotations and reflections, but they don't perform any "stretching."

Now if $A \in M_{m \times n}(\mathbb{R})$ has an SVD given by $A = U\Sigma V^T$, then we can interpret the orthogonal matrices V^T and U as each performing rotations and/or reflections, while the diagonal matrix Σ performs a scaling (possibly by zero in some directions). So the singular value decomposition now paints the following picture of any linear map from \mathbb{R}^n to \mathbb{R}^m : it is built up of rotations and/or reflections, followed by scalings, then followed by additional rotations and/or reflections! For instance, the effect of $A \in M_{2 \times 2}(\mathbb{R})$ on the unit circle can be pictured as follows:



6.3 Singular Value Decomposition of Linear Maps

We now come to the result that we wished to establish in the introduction to this chapter. We're going to prove that a linear map between finite-dimensional inner product spaces can always be "diagonalized" by picking appropriate (and possibly distinct) orthonormal bases for its domain and codomain.

Theorem 6.3.1 (Singular Value Decomposition of Linear Maps)

Let $L: W_1 \rightarrow W_2$ be a linear map between finite-dimensional inner product spaces of dimensions n and m , respectively. If $r = \text{rank}(A)$, then there exist orthonormal basis \mathcal{B} and \mathcal{C} for W_1 and W_2 and an $r \times r$ diagonal matrix D such that

$$c[L]_{\mathcal{B}} = \begin{bmatrix} D & 0_{r \times (n-r)} \\ 0_{(m-r) \times r} & 0_{(m-r) \times (n-r)} \end{bmatrix}.$$

Proof: Choose orthonormal bases \mathcal{B}' and \mathcal{C}' for W_1 and W_2 , let $A = c'[L]_{\mathcal{B}'}$, and let $A = U\Sigma V^*$ be an SVD of A . Since \mathcal{B}' is an orthonormal basis, the discussion preceding Theorem 5.5.1 shows that the inner product of two vectors in W_1 is the same as the inner product of their \mathcal{B}' -coordinate vectors; and the analogous statement is true for W_2 and \mathcal{C}' . Thus if we let \mathcal{B} be the basis for W_1 consisting of the vectors whose \mathcal{B}' -coordinates are the columns of V , and if we let \mathcal{C} be the basis for W_2 consisting of the vectors whose \mathcal{C}' -coordinates are the columns of U , then \mathcal{B} and \mathcal{C} are orthonormal bases, and we have $U = c'\mathcal{I}_{\mathcal{C}}$ and $V = \mathcal{B}'\mathcal{I}_{\mathcal{B}}$ so that

$$c[L]_{\mathcal{B}} = c\mathcal{I}_{\mathcal{C}'} c'[L]_{\mathcal{B}'} \mathcal{B}'\mathcal{I}_{\mathcal{B}} = U^{-1}AV = \Sigma,$$

as required. □

Let's see the proof of this theorem in action.

Example 6.3.2

Consider the differentiation map $D: \mathcal{P}_2(\mathbb{F}) \rightarrow \mathcal{P}_1(\mathbb{F})$ given by $D(p(x)) = p'(x)$, and suppose that both polynomial spaces are endowed with the inner product $\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx$.

The first order of business is to find orthonormal bases \mathcal{B}' and \mathcal{C}' for $\mathcal{P}_2(\mathbb{R})$ and $\mathcal{P}_1(\mathbb{R})$, respectively. According to Example 4.3.6, we can take

$$\mathcal{B}' = \left\{ \frac{1}{\sqrt{2}}, \sqrt{\frac{3}{2}}x, \sqrt{\frac{5}{2}} \left(\frac{3}{2}x^2 - \frac{1}{2} \right) \right\} \quad \text{and} \quad \mathcal{C}' = \left\{ \frac{1}{\sqrt{2}}, \sqrt{\frac{3}{2}}x \right\}.$$

Then let

$$A = c'[D]_{\mathcal{B}'} = \begin{bmatrix} 0 & \sqrt{3} & 0 \\ 0 & 0 & \sqrt{15} \end{bmatrix}.$$

We must now find an SVD of A . We have

$$A^T A = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 15 \end{bmatrix}.$$

So the eigenvalues of A^*A are $\lambda_1 = 15$, $\lambda_2 = 3$ and $\lambda_3 = 0$ with corresponding eigenvectors $\vec{v}_1 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$, $\vec{v}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$ and $\vec{v}_3 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$.

Thus the singular values of A are $\sigma_1 = \sqrt{15}$, $\sigma_2 = \sqrt{3}$ and $\sigma_3 = 0$, with corresponding singular vectors \vec{v}_1 , \vec{v}_2 and \vec{v}_3 as above.

Now let

$$\vec{u}_1 = \frac{1}{\sigma_1}A\vec{v}_1 = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad \text{and} \quad \vec{u}_2 = \frac{1}{\sigma_2}A\vec{v}_2 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

Then $\{\vec{u}_1, \vec{u}_2\}$ is an orthonormal basis for \mathbb{R}^2 . So if we let

$$V = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} \quad \text{and} \quad U = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

then the corresponding SVD of A is

$$A = U \begin{bmatrix} \sqrt{15} & 0 & 0 \\ 0 & \sqrt{3} & 0 \end{bmatrix} V^T.$$

Let's convert everything back to polynomials. Our desired basis \mathcal{B} for $\mathcal{P}_2(\mathbb{R})$ comes from the columns of V (thought of as being \mathcal{B}' -coordinate vectors):

$$\mathcal{B} = \left\{ \sqrt{\frac{5}{2}} \left(\frac{3}{2}x^2 - \frac{1}{2} \right), \sqrt{\frac{3}{2}}x, \frac{1}{\sqrt{2}} \right\}.$$

Similarly, our desired basis \mathcal{C} for $\mathcal{P}_1(\mathbb{R})$ comes from the columns of U (thought of as being \mathcal{C}' -coordinate vectors):

$$\mathcal{C} = \left\{ \sqrt{\frac{3}{2}}x, \frac{1}{\sqrt{2}} \right\}.$$

We'll leave it to you to check that

$$c[D]_{\mathcal{B}} = \begin{bmatrix} \sqrt{15} & 0 & 0 \\ 0 & \sqrt{3} & 0 \end{bmatrix},$$

as claimed by Theorem 6.3.1 (bear in mind that $r = \text{rank}(D) = 2$).

6.4 Applications of the Singular Value Decomposition

In this final section we will highlight a select few applications of the singular value decomposition. We will barely scratch the surface of what is possible. Indeed, the applications of the singular value decomposition are too numerous and too broad, ranging from computational mathematics, statistics, text analysis, image and data compression, machine learning, quantum physics, to even psychology! Essentially, if you are able to put your data into a matrix somehow, then it's very likely that the SVD will say something interesting about this matrix (and hence your data).

Here's a simple, but relatable, example. You've probably used software to compute the rank or nullspace of a matrix. Have you wondered how the software does these computations? You might be surprised to learn that it very likely used SVD at some point. For instance, Mathematica (hence WolframAlpha), Maple and MATLAB all use SVD in their rank and nullspace routines. One reason for this is that SVD offers certain numerical advantages that makes it more stable than, say, the more efficient Gaussian elimination algorithm.

6.4.1 Low-rank Approximations

Many real-life applications of linear algebra involve large matrices built out of data. For instance, the results of a demographic survey of the population of Ontario can be recorded in a matrix each of whose rows represents a resident of Ontario, and whose columns represent values such as age, gender, income and city of residence. Movies on a streaming platform could be represented as rows in a matrix whose columns contain information such as language, running time and genre. A gray-scale image could be represented as a matrix whose entries are the color intensities of each pixel.

In practice, these large data matrices tend to contain certain dominant features due to the inherent correlation in the data (e.g. nearby pixels in an image tend to have the same shade of colour), and it's highly desirable that these dominant features be identified. The SVD provides us with a way to do this. It turns out that the singular vectors associated to the larger singular values contain most of the information about the matrix, in a certain sense. Therefore by “forgetting” about the parts of the matrix coming from the smaller singular values, we're able to somehow compress our matrix down into something that is simpler to analyze but that is still fairly representative. We will explain how this works in this section, and at the end we will illustrate by showing how these ideas can help with image compression.

As a first step, we will introduce a more compact version of SVD that removes unnecessary rows and columns from U , Σ and V^T . Let's look back at the SVDs from Examples 6.2.4–6.2.6 to illustrate the idea. We had obtained the decompositions

$$\begin{bmatrix} 1 & -4 \\ -2 & 2 \\ 2 & 4 \end{bmatrix} = [\vec{u}_1 \ \vec{u}_2 \ \vec{u}_3] \begin{bmatrix} 6 & 0 \\ 0 & 3 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \vec{v}_1^T \\ \vec{v}_2^T \end{bmatrix} \quad (6.1)$$

$$\begin{bmatrix} 1 & 1 & 0 \\ -1 & 1 & 1 \end{bmatrix} = [\vec{u}_1 \ \vec{u}_2] \begin{bmatrix} \sqrt{3} & 0 & 0 \\ 0 & \sqrt{2} & 0 \end{bmatrix} \begin{bmatrix} \vec{v}_1^T \\ \vec{v}_2^T \\ \vec{v}_3^T \end{bmatrix} \quad (6.2)$$

$$\begin{bmatrix} 0 & 2 \\ -3 & 0 \end{bmatrix} = [\vec{u}_1 \ \vec{u}_2] \begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} \vec{v}_1^T \\ \vec{v}_2^T \end{bmatrix}. \quad (6.3)$$

In (6.1), the row of 0s in Σ indicates that the vector \vec{u}_3 is rather unnecessary since it will get multiplied by 0. So if we delete \vec{u}_3 and the row of 0s from Σ , we would obtain the simpler decomposition

$$\begin{bmatrix} 1 & -4 \\ -2 & 2 \\ 2 & 4 \end{bmatrix} = [\vec{u}_1 \ \vec{u}_2] \begin{bmatrix} 6 & 0 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} \vec{v}_1^T \\ \vec{v}_2^T \end{bmatrix}$$

Likewise, the column of 0s in (6.2) indicates that we could've done without \vec{v}_3^T . By deleting both, we obtain the simpler decomposition

$$\begin{bmatrix} 1 & 1 & 0 \\ -1 & 1 & 1 \end{bmatrix} = [\vec{u}_1 \ \vec{u}_2] \begin{bmatrix} \sqrt{3} & 0 \\ 0 & \sqrt{2} \end{bmatrix} \begin{bmatrix} \vec{v}_1^T \\ \vec{v}_2^T \end{bmatrix}.$$

On the other hand, there are no redundancies in (6.3).

In general, suppose that A is an $m \times n$ matrix of rank r with SVD $A = U\Sigma V^*$. The $m \times n$ matrix Σ will have r non-zero entries on its diagonal. By deleting all zero rows and columns from Σ , we are left with an $r \times r$ diagonal matrix $\Sigma_r = \text{diag}(\sigma_1, \dots, \sigma_r)$. Let U_r and V_r be the matrices formed from the first r columns of U and V . Thus, U_r is an $m \times r$ matrix and V_r^* is an $r \times n$ matrix.

Definition 6.4.1
Compact SVD

Let $A \in M_{m \times n}(\mathbb{F})$ be a rank r matrix with singular value decomposition $A = U\Sigma V^*$. Let Σ_r , U_r and V_r be as described in the preceding paragraph. The decomposition

$$A = U_r \Sigma_r V_r^*$$

is called a **compact** singular value decomposition of A .

Example 6.4.2

In Example 6.2.7 we obtained the singular value decomposition

$$A = \begin{bmatrix} 1 & 2 & -1 \\ 2 & 4 & -2 \\ -1 & -2 & 1 \\ i & 2i & -i \end{bmatrix} = U \begin{bmatrix} \sqrt{42} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} V^*,$$

where

$$U = \begin{bmatrix} -\frac{1}{\sqrt{7}} & \frac{i}{\sqrt{2}} & \frac{1}{\sqrt{6}} & -\frac{2}{\sqrt{21}} \\ \frac{-2}{\sqrt{7}} & 0 & 0 & \frac{3}{\sqrt{21}} \\ \frac{1}{\sqrt{7}} & 0 & \frac{2}{\sqrt{6}} & \frac{\sqrt{21}}{2} \\ \frac{1}{\sqrt{7}} & 0 & \frac{2}{\sqrt{6}} & \frac{\sqrt{21}}{2} \\ -\frac{i}{\sqrt{7}} & \frac{1}{\sqrt{2}} & \frac{i}{\sqrt{6}} & -\frac{2i}{\sqrt{21}} \end{bmatrix} \quad \text{and} \quad V = \begin{bmatrix} -\frac{1}{\sqrt{6}} & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{3}} \\ -\frac{2}{\sqrt{6}} & 0 & \frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{6}} & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{3}} \end{bmatrix}.$$

The matrix A has rank $r = 1$ (since it has only one non-zero singular value). Let's delete the last three zero rows from Σ (hence the last three columns from U), and then let's delete the last two columns from Σ (hence the last two columns from V). We are left with

$$\Sigma_r = [\sqrt{42}], \quad U_r = \frac{1}{\sqrt{7}} \begin{bmatrix} -1 \\ -2 \\ 1 \\ -i \end{bmatrix} \quad \text{and} \quad V_r = \frac{1}{\sqrt{6}} \begin{bmatrix} -1 \\ -2 \\ 1 \end{bmatrix}.$$

The corresponding compact SVD is

$$\begin{bmatrix} 1 & 2 & -1 \\ 2 & 4 & -2 \\ -1 & -2 & 1 \\ i & 2i & -i \end{bmatrix} = \begin{bmatrix} -\frac{1}{\sqrt{7}} \\ \frac{2}{\sqrt{7}} \\ \frac{1}{\sqrt{7}} \\ -\frac{i}{\sqrt{7}} \end{bmatrix} [\sqrt{42}] \begin{bmatrix} -\frac{1}{\sqrt{6}} & -\frac{2}{\sqrt{6}} & \frac{1}{\sqrt{6}} \end{bmatrix}.$$

You should multiply out the right side to confirm that it does indeed give the left side!

Notice that we've replaced our original decomposition $A = U\Sigma V^*$ with the much simpler $A = \sigma_1 \vec{u}_1 \vec{v}_1^*$.

The compact SVD in the example above took on the particularly simple form $A = \sigma_1 \vec{u}_1 \vec{v}_1^*$ owing to the fact that $r = \text{rank}(A)$ was equal to 1. In the general situation, we'll be able to express A as a sum of r matrices of this form.

Proposition 6.4.3

Let $A \in M_{m \times n}(\mathbb{F})$ have rank r and compact SVD $A = U_r \Sigma_r V_r^*$, where $\Sigma_r = \begin{bmatrix} \sigma_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \sigma_r \end{bmatrix}$.

Let $\vec{u}_1, \dots, \vec{u}_r$ and $\vec{v}_1, \dots, \vec{v}_r$ be the columns of U_r and V_r , respectively. Then:

$$A = \sigma_1 \vec{u}_1 \vec{v}_1^* + \dots + \sigma_r \vec{u}_r \vec{v}_r^*.$$

Proof: We have

$$\begin{aligned} A &= U_r \Sigma_r V_r^* \\ &= [\vec{u}_1 \ \dots \ \vec{u}_r] \begin{bmatrix} \sigma_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \sigma_r \end{bmatrix} \begin{bmatrix} \vec{v}_1^* \\ \vdots \\ \vec{v}_r^* \end{bmatrix} \\ &= [\sigma_1 \vec{u}_1 \ \dots \ \sigma_r \vec{u}_r] \begin{bmatrix} \vec{v}_1^* \\ \vdots \\ \vec{v}_r^* \end{bmatrix} \\ &= \sigma_1 \vec{u}_1 \vec{v}_1^* + \dots + \sigma_r \vec{u}_r \vec{v}_r^*, \end{aligned}$$

using the definition of matrix multiplication. □

Example 6.4.4

From Example 6.2.4, we have the SVD

$$\begin{bmatrix} 1 & -4 \\ -2 & 2 \\ 2 & 4 \end{bmatrix} = \begin{bmatrix} -\frac{2}{3} & \frac{1}{3} & \frac{2}{3} \\ \frac{1}{3} & -\frac{2}{3} & \frac{2}{3} \\ \frac{2}{3} & \frac{2}{3} & \frac{1}{3} \end{bmatrix} \begin{bmatrix} 6 & 0 \\ 0 & 3 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

hence compact SVD

$$\begin{bmatrix} 1 & -4 \\ -2 & 2 \\ 2 & 4 \end{bmatrix} = \begin{bmatrix} -\frac{2}{3} & \frac{1}{3} \\ \frac{1}{3} & -\frac{2}{3} \\ \frac{2}{3} & \frac{2}{3} \end{bmatrix} \begin{bmatrix} 6 & 0 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Thus,

$$\begin{bmatrix} 1 & -4 \\ -2 & 2 \\ 2 & 4 \end{bmatrix} = 6 \begin{bmatrix} -\frac{2}{3} \\ \frac{1}{3} \\ \frac{2}{3} \end{bmatrix} [0 \ 1] + 3 \begin{bmatrix} \frac{1}{3} \\ -\frac{2}{3} \\ \frac{2}{3} \end{bmatrix} [1 \ 0].$$

Exercise 77

Let $\vec{u}, \vec{v} \in \mathbb{F}^n$ be non-zero vectors. Show that the $n \times n$ matrix $\vec{u}\vec{v}^*$ has rank 1. This is sometimes called the **outer product** of \vec{u} and \vec{v} . (Recall that the standard *inner* product of \vec{u} and \vec{v} is given by $\vec{v}^*\vec{u}$.)

In view of the previous exercise and Proposition 6.4.3, we're now able to use the SVD to express a given rank r matrix A as the sum of r matrices of rank 1:

$$A = \sigma_1 \vec{u}_1 \vec{v}_1^* + \cdots + \sigma_r \vec{u}_r \vec{v}_r^*.$$

Since the singular values of A are ordered in descending order $\sigma_1 \geq \cdots \geq \sigma_r$, the tail-end of this representation will in some sense be “less significant” if the smaller singular values are very small. In many real-life scenarios, the small singular values will be numerically much smaller than the first few large singular values, and so if we discard them then not much is lost. This motivates the following definition.

Definition 6.4.5

Rank- k Truncation

Let $A \in M_{m \times n}(\mathbb{F})$ be a rank r matrix with singular values $\sigma_1 \geq \cdots \geq \sigma_r > 0$ and compact singular value decomposition $A = U_r \Sigma_r V_r^*$, where $U = [\vec{u}_1 \cdots \vec{u}_r]$ and $V = [\vec{v}_1 \cdots \vec{v}_r]$. Let $k \leq r$ be a positive integer. The **rank- k truncation** of A is

$$A_k = \sigma_1 \vec{u}_1 \vec{v}_1^* + \cdots + \sigma_k \vec{u}_k \vec{v}_k^*.$$

That is, the rank- k truncation only keeps the contributions of the largest k singular values and corresponding singular vectors and discards the rest. Here are some key properties.

Proposition 6.4.6

Let $A \in M_{m \times n}(\mathbb{F})$ be a matrix of rank r with rank- k truncation

$$A_k = \sigma_1 \vec{u}_1 \vec{v}_1^* + \cdots + \sigma_k \vec{u}_k \vec{v}_k^*.$$

Then:

- $A = A_r$.
- $\text{rank}(A_k) = k$.
- $\|A - A_k\| \leq \sum_{i=k+1}^r \sigma_i$, where $\|X\| = \sqrt{\text{tr}(X^*X)}$ is the norm induced from the Frobenius inner product $\langle A, B \rangle = \text{tr}(B^*A)$ on $M_{m \times n}(\mathbb{F})$. (See Example 4.1.6 and the two exercises that follow it.)

Proof: (a) This is just Proposition 6.4.3.

- (b) Let's view A_k as giving the linear map $L_k: \mathbb{F}^n \rightarrow \mathbb{F}^m$ defined by $L_k(\vec{x}) = A_k \vec{x}$. We claim that the range of L_k is $\text{Span}\{\vec{u}_1, \dots, \vec{u}_k\}$. Since the vectors $\{\vec{u}_1, \dots, \vec{u}_k\}$ are linearly independent, this will show that the rank of L_k , hence of A_k , is equal to k , as required.

To see why the claim about the range of L_k is true, consider first any $\vec{x} \in \mathbb{F}^n$. Then

$$L_k(\vec{x}) = A_k \vec{x} = \sum_{i=1}^k \sigma_i \vec{u}_i \vec{v}_i^* \vec{x} = \sum_{i=1}^k \sigma_i \vec{u}_i \langle \vec{x}, \vec{v}_i \rangle$$

is a linear combination of $\{\vec{u}_1, \dots, \vec{u}_k\}$, proving that $\text{Range}(L_k) \subseteq \text{Span}\{\vec{u}_1, \dots, \vec{u}_k\}$. Conversely, given any $j \in \{1, \dots, k\}$, consider

$$L_k(\vec{v}_j) = \sum_{i=1}^k \sigma_i \vec{u}_i \langle \vec{v}_j, \vec{v}_i \rangle.$$

Since $\{\vec{v}_1, \dots, \vec{v}_j\}$ is an orthonormal set, the above reduces to $L_k(\vec{v}_j) = \sigma_j \vec{u}_j$. And since $\sigma_j \neq 0$ by the definition of the compact SVD, it follows that $\vec{u}_j = L_k(\frac{1}{\sigma_j} \vec{v}_j)$ is in the range of L_k . This completes the proof that $\text{Range}(L_k) = \text{Span}\{\vec{u}_1, \dots, \vec{u}_k\}$, as claimed.

- (c) We have

$$\begin{aligned} \|A - A_k\| &= \left\| \sum_{i=k+1}^r \sigma_i \vec{u}_i \vec{v}_i^* \right\| \\ &\leq \sum_{i=k+1}^r |\sigma_i| \|\vec{u}_i \vec{v}_i^*\| \\ &= \sum_{i=k+1}^r \sigma_i \text{tr}((\vec{u}_i \vec{v}_i^*)^* \vec{u}_i \vec{v}_i^*) \\ &= \sum_{i=k+1}^r \sigma_i \text{tr}(\vec{v}_i \vec{u}_i^* \vec{u}_i \vec{v}_i^*) \\ &= \sum_{i=k+1}^r \sigma_i \text{tr}(\vec{v}_i \vec{v}_i^*), \end{aligned}$$

where the last equality follows since $\vec{u}_i^* \vec{u}_i = \langle \vec{u}_i, \vec{u}_i \rangle = 1$. Now, using the fact that $\text{tr}(XY) = \text{tr}(YX)$ for matrices X and Y of compatible sizes, we see that $\text{tr}(\vec{v}_i \vec{v}_i^*) = \text{tr}(\vec{v}_i^* \vec{v}_i) = \langle \vec{v}_i, \vec{v}_i \rangle = 1$, completing the proof. \square

Exercise 78

We can do better than the inequality in part (c) of Proposition 6.4.6. Show that, in fact,

$$\|A - A_k\| = \sqrt{\sum_{i=k+1}^r \sigma_i^2}.$$

Hint: Show that the matrices $\vec{u}_i \vec{v}_i^*$ and $\vec{u}_j \vec{v}_j^*$ are orthogonal with respect to the Frobenius inner product.

Part (b) of Proposition 6.4.6 justifies that the name “rank- k truncation” for A_k , while part (c) shows that if the singular values σ_i for $i > k$ are small, then $\|A - A_k\|$ will be small too, and so A_k will in this sense be a good approximation to A . Of course, this is only a heuristic justification, since even if the σ_i 's are small, their sum might not be small. However, there is a sense in which A_k is *the best* rank k approximation to A . This is the content of the next theorem, which we state without proof. (The proof is not difficult, but it requires a few extra ideas that will veer us off course.)

Theorem 6.4.7 (Eckart–Young Theorem)

Let $A \in M_{m \times n}(\mathbb{F})$, and let A_k be the rank- k truncation of A . Let $B \in M_{m \times n}(\mathbb{F})$ be an arbitrary rank k matrix. Then

$$\|A - B\| \geq \|A - A_k\|.$$

Thus, A_k is closer to A than any other rank k matrix B .

Example 6.4.8

Consider the matrix $A = \begin{bmatrix} 1 & 0 & -2 \\ 1 & 1 & 3 \\ -1 & 0 & 1 \\ -4 & 0 & 2 \end{bmatrix}$. Using software, we find $A = U\Sigma V^*$ with

$$U \approx \begin{bmatrix} -0.408 & -0.207 & -0.833 & -0.312 \\ 0.272 & 0.894 & -0.356 & 0 \\ 0.274 & 0.003 & 0.217 & -0.937 \\ 0.827 & -0.397 & -0.365 & 0.156 \end{bmatrix}, \quad V \approx \begin{bmatrix} -0.721 & 0.677 & 0.148 \\ 0.053 & 0.266 & -0.962 \\ 0.691 & 0.686 & 0.228 \end{bmatrix}$$

and

$$\Sigma \approx \begin{bmatrix} 5.156 & 0 & 0 \\ 0 & 3.358 & 0 \\ 0 & 0 & 0.370 \\ 0 & 0 & 0 \end{bmatrix}.$$

We can delete the bottom row from Σ and the fourth column from U to obtain a compact SVD. Here are the rank truncations of A :

$$\begin{aligned} A_1 &\approx 5.156 \begin{bmatrix} -0.408 \\ 0.272 \\ 0.274 \\ 0.827 \end{bmatrix} \begin{bmatrix} -0.721 & 0.053 & 0.691 \end{bmatrix} = \begin{bmatrix} 1.517 & -0.111 & -1.452 \\ -1.011 & 0.074 & 0.968 \\ -1.019 & 0.075 & 0.975 \\ -3.074 & 0.226 & 2.943 \end{bmatrix} \\ A_2 &\approx A_1 + 3.358 \begin{bmatrix} -0.207 \\ 0.894 \\ 0.003 \\ -0.397 \end{bmatrix} \begin{bmatrix} 0.677 & 0.266 & 0.686 \end{bmatrix} = \begin{bmatrix} 1.046 & -0.295 & -1.930 \\ 1.021 & 0.872 & 3.028 \\ -1.012 & 0.077 & 0.983 \\ -3.977 & -0.131 & 2.032 \end{bmatrix} \\ A_3 &= A. \end{aligned}$$

We have $\|A - A_1\| \approx 3.379 \leq \sigma_2 + \sigma_3$ and $\|A - A_2\| \approx 0.36983 \leq \sigma_3$, as predicted by Proposition 6.4.6(c).

Let's now illustrate how all this can be applied to perform image compression. The basic idea is to encode an image into a matrix, say by converting to gray-scale and then recording the (i, j) th pixel intensity as the (i, j) th entry of a matrix A . This will usually create a rather large matrix, with large rank. However, in practice, most of the singular values of A will be extremely small. Thus we will be able to use the rank- k approximation for $k \ll \text{rank}(A)$ to approximate the original image. This results in us having to use less memory to store the image, since we only need to use the truncated SVD data to reconstruct it.

Here is a simple example of three rank- k truncations of an image.




Figure 6.4.1: Low-rank image compression. Ritmüller, E. (n.d.). Portrait of Carl Friedrich Gauss [Lithograph]. Smithsonian Libraries. <https://library.si.edu/image-gallery/73624>. Adapted using Python code.

We should mention in closing that this is not a very sophisticated image compression technique (in particular, it is *lossy*), and there are many better compression algorithms available. Nonetheless, it is a very compelling application of SVD! You can visit [Tim Baumann's SVD-Demo webpage](#) to try it yourself.

6.4.2 The Pseudoinverse and Least Squares Revisited

Let's consider once more the basic problem of solving a system of equations given in matrix form as $A\vec{x} = \vec{b}$, where $A \in M_{m \times n}(\mathbb{F})$ and $\vec{b} \in \mathbb{F}^m$. As you well know, this system has a solution if and only if $\vec{b} \in \text{Col}(A)$. If A is square and invertible, then this condition is always satisfied, and in fact we know that the system will have a *unique* solution, which is given by $\vec{x} = A^{-1}\vec{b}$.

If A is an arbitrary $m \times n$ matrix, we can use the SVD $A = U\Sigma V^*$ to re-write the equation $A\vec{x} = \vec{b}$ as

$$U\Sigma V^* \vec{x} = \vec{b}.$$

We now wish to “invert” U , Σ and V . With U and V , this is no problem, since they are square and unitary. However, Σ is an $m \times n$ matrix, so it doesn't make sense to invert it. But let's pretend we can: the non-zero diagonal entries of Σ are the non-zero singular values $\sigma_1, \dots, \sigma_r$, so let's define Σ^\dagger to be the $n \times m$ matrix (not the $m \times n$ matrix!) whose diagonal entries are $\frac{1}{\sigma_1}, \dots, \frac{1}{\sigma_r}$ and all of whose other entries are 0.

Example 6.4.9 If $\Sigma = \begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix}$ then $\Sigma^\dagger = \begin{bmatrix} \frac{1}{3} & 0 \\ 0 & \frac{1}{2} \end{bmatrix}$. Notice that in this case Σ^\dagger is equal to Σ^{-1} . On the other hand,

$$\begin{bmatrix} 6 & 0 \\ 0 & 3 \\ 0 & 0 \end{bmatrix}^\dagger = \begin{bmatrix} \frac{1}{6} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} \sqrt{3} & 0 & 0 \\ 0 & \sqrt{2} & 0 \end{bmatrix}^\dagger = \begin{bmatrix} \frac{1}{\sqrt{3}} & 0 \\ 0 & \frac{1}{\sqrt{2}} \\ 0 & 0 \end{bmatrix}.$$

Now, given the equation

$$U\Sigma V^* \vec{x} = \vec{b}$$

we can “invert” each matrix on the left-side, in order, to get the “solution”

$$\vec{x}_0 = V\Sigma^\dagger U^* \vec{b}.$$

So, in effect, the matrix $V\Sigma^\dagger U^*$ is acting like some kind of inverse for $A = U\Sigma V^*$. We’ll explain the meaning of this \vec{x}_0 below, but first let’s introduce some terminology.

Definition 6.4.10

Pseudoinverse

Let $A \in M_{m \times n}(\mathbb{F})$ be a matrix of rank r with SVD $A = U\Sigma V^*$. The **pseudoinverse** of A is the $n \times m$ matrix

$$A^\dagger = V\Sigma^\dagger U^*,$$

where Σ^\dagger is the $n \times m$ matrix whose (i, j) th entry is $\frac{1}{\sigma_i}$ for $i = j \leq r$ and 0 otherwise.

Notice that Σ^\dagger is in fact the pseudoinverse of Σ , so the notation is consistent. The pseudoinverse of A is uniquely determined by A , and does not depend on the choice of singular value decomposition (i.e. on a choice of U and V).

Example 6.4.11

If A is an invertible $n \times n$ matrix, then all n of its singular values are nonzero thanks to Proposition 6.1.9. Thus in this case $\Sigma^\dagger = \Sigma^{-1}$ and it follows that the pseudoinverse of A is actually the inverse of A :

$$A^\dagger = A^{-1}.$$

Exercise 79

Supply the remaining details to prove that if A is invertible then $A^\dagger = A^{-1}$.

Example 6.4.12

For $A = \begin{bmatrix} 1 & -4 \\ -2 & 2 \\ 2 & 4 \end{bmatrix}$ as in Example 6.2.4, we found that $A = U\Sigma V^*$ with

$$U = \frac{1}{3} \begin{bmatrix} -2 & 1 & 2 \\ 1 & -2 & 2 \\ 2 & 2 & 1 \end{bmatrix}, \quad \Sigma = \begin{bmatrix} 6 & 0 \\ 0 & 3 \\ 0 & 0 \end{bmatrix} \quad \text{and} \quad V = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Thus,

$$A^\dagger = V\Sigma^\dagger U^*$$

$$\begin{aligned}
&= \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{6} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \end{bmatrix} \frac{1}{3} \begin{bmatrix} -2 & 1 & 2 \\ 1 & -2 & 2 \\ 2 & 2 & 1 \end{bmatrix} \\
&= \frac{1}{18} \begin{bmatrix} 2 & -4 & 4 \\ -2 & 1 & 2 \end{bmatrix}.
\end{aligned}$$

Let's return now to our system of linear equations

$$A\vec{x} = \vec{b}, \quad (6.4)$$

where $A \in M_{m \times n}(\mathbb{F})$ and $\vec{b} \in \mathbb{F}^m$. Using the pseudoinverse of A , we create the vector $\vec{x}_0 = A^\dagger \vec{b} \in \mathbb{F}^n$. If A is an invertible square matrix, then $\vec{x}_0 = A^{-1} \vec{b}$ is the unique solution to the system (6.4). In the general case, \vec{x}_0 is still interesting, as the next proposition explains.

Proposition 6.4.13 (Minimal Norm Solutions)

Consider the system of linear equations $A\vec{x} = \vec{b}$, where $A \in M_{m \times n}(\mathbb{F})$ and $\vec{b} \in \mathbb{F}^m$. Let $\vec{x}_0 = A^\dagger \vec{b}$. Then:

- (a) If $A\vec{x} = \vec{b}$ is consistent, then \vec{x}_0 is a solution to the system. Moreover, it is the solution of minimal norm, i.e., if \vec{x} is any solution to the system, then $\|\vec{x}\| \geq \|\vec{x}_0\|$ with equality if and only if $\vec{x} = \vec{x}_0$.
- (b) If $A\vec{x} = \vec{b}$ is inconsistent, then \vec{x}_0 is a least squares solution (Definition 4.7.1). Moreover, it is the least squares solution of minimal norm, i.e., if \vec{s} is any least squares solution to the system, then $\|\vec{s}\| \geq \|\vec{x}_0\|$ with equality if and only if $\vec{s} = \vec{x}_0$.

REMARK

Although in our initial discussion of least squares (Section 4.7) we only worked over $\mathbb{F} = \mathbb{R}$, everything works just as well over $\mathbb{F} = \mathbb{C}$. The one difference is: instead of using transposes, we should be using conjugate-transposes, as you probably would have guessed.

Proof of Proposition 6.4.13: Let's show that \vec{x}_0 is always a least squares solution. In the case where the system is consistent, this will automatically imply that \vec{x}_0 is an actual solution (why?). By Proposition 4.7.3, we must show that $A^* A \vec{x}_0 = A^* \vec{b}$. We have

$$\begin{aligned}
A^* A \vec{x}_0 &= A^* A A^\dagger \vec{b} \\
&= (V \Sigma^* U^*) (U \Sigma V^*) (V \Sigma^\dagger U^*) \vec{b} \\
&= V \Sigma^* \Sigma \Sigma^\dagger U^* \vec{b}.
\end{aligned} \quad (*)$$

Now, the matrix $\Sigma\Sigma^\dagger$ will take the form $\begin{bmatrix} I_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$, with enough 0s to fill out an $m \times m$ matrix.

On the other hand, Σ^* will take the form $\begin{bmatrix} D^* & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$, where $D = \text{diag}(\sigma_1, \dots, \sigma_r)$ is $r \times r$ and with enough 0s to fill out an $n \times m$ matrix. It follows that

$$\Sigma^*\Sigma\Sigma^\dagger = \begin{bmatrix} D^* & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} I_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} = \begin{bmatrix} D^* & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} = \Sigma^*.$$

So from (*) we obtain

$$A^*A\vec{x}_0 = V\Sigma^*U^*\vec{b} = A^*\vec{b},$$

proving that \vec{x}_0 is a least squares solution to $A\vec{x} = \vec{b}$, as desired. All that remains now is proving that \vec{x}_0 has minimal norm.

If $A\vec{x} = \vec{b}$ is consistent, then an arbitrary solution will take the form $\vec{x}_0 + \vec{z}$ where $\vec{z} \in \text{Null}(A)$, since \vec{x}_0 is a solution. Now we make two observations. First, $\vec{z} \in \text{Col}(A^*)^\perp$, by Proposition 6.2.3. Next, $\vec{x}_0 = A^\dagger\vec{b} = V(\Sigma^\dagger U^*\vec{b})$ is a linear combination of the first r columns of V , which are eigenvectors of A^*A corresponding to nonzero eigenvalues; so each of these eigenvectors is in $\text{Col}(A^*)$ (since if $A^*A\vec{v} = \lambda\vec{v}$ with $\lambda \neq 0$ then $\vec{v} = A^*(\frac{1}{\lambda}A\vec{v})$), and therefore $\vec{x}_0 \in \text{Col}(A^*)$. Combining both observations, we deduce that $\vec{x}_0 \perp \vec{z}$. Hence, by the Pythagorean theorem,

$$\|\vec{x}_0 + \vec{z}\|^2 = \|\vec{x}_0\|^2 + \|\vec{z}\|^2 \geq \|\vec{x}_0\|^2,$$

completing the proof of the minimality of $\|\vec{x}_0\|$ amongst norms of solutions. Notice also that the above inequality is an equality if and only if $\vec{z} = \vec{0}$.

Finally, if $A\vec{x} = \vec{b}$ is inconsistent, then an arbitrary least squares solution will take the form $\vec{x}_0 + \vec{z}$ where $\vec{z} \in \text{Null}(A)$. The same proof of minimality above, which never used the fact that \vec{x}_0 was an actual solution to $A\vec{x} = \vec{b}$, applies once more to give us the desired minimality result in this case. This completes the proof of the proposition. \square

Example 6.4.14

Let $A = \begin{bmatrix} 1 & 2 \\ 1 & 2 \end{bmatrix}$ and consider the system $A\vec{x} = \vec{b}$ where $\vec{b} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$. In Example 4.7.2 we noted that this system is inconsistent, and we determined the set of least squares solutions to be

$$S = \left\{ \begin{bmatrix} \frac{1}{2} \\ 0 \end{bmatrix} + t \begin{bmatrix} -2 \\ 1 \end{bmatrix} : t \in \mathbb{R} \right\}.$$

The norm of an arbitrary vector in S is given by

$$\left\| \begin{bmatrix} \frac{1}{2} \\ 0 \end{bmatrix} + t \begin{bmatrix} -2 \\ 1 \end{bmatrix} \right\| = \sqrt{\left(\frac{1}{2} - 2t\right)^2 + t^2}.$$

Using calculus, it's easy to show that this norm is minimized precisely when $t = \frac{1}{5}$. The corresponding least squares solution is $\frac{1}{10} \begin{bmatrix} 2 \\ 1 \end{bmatrix}$.

Let's confirm that this is in line with Proposition 6.4.13. An SVD for A is $A = U\Sigma V^*$, where

$$U = \begin{bmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \sqrt{10} & 0 \\ 0 & 0 \end{bmatrix} \quad \text{and} \quad V = \begin{bmatrix} \frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \end{bmatrix}.$$

Hence the pseudoinverse of A is

$$A^\dagger = V\Sigma^\dagger U^* = \begin{bmatrix} \frac{1}{\sqrt{5}} & -\frac{-2}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{10}} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} = \frac{1}{10} \begin{bmatrix} 1 & 1 \\ 2 & 2 \end{bmatrix}.$$

Thus, the minimal norm least squares solution should be

$$A^\dagger = A^\dagger \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \frac{1}{10} \begin{bmatrix} 1 \\ 2 \end{bmatrix},$$

just as we've computed above!

Appendix A

Solutions to In-Chapter Exercises

Chapter 1: Abstract Vector Spaces

Exercise 1: These all follow from the properties of real numbers. For example, for axiom 4: if $\vec{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$ and $\vec{y} = \begin{bmatrix} y_1 \\ y_2 \end{bmatrix}$, then

$$\vec{x} + \vec{y} = \begin{bmatrix} x_1 + y_1 \\ x_2 + y_2 \end{bmatrix} \quad \text{and} \quad \vec{y} + \vec{x} = \begin{bmatrix} y_1 + x_1 \\ y_2 + x_2 \end{bmatrix}.$$

However, for real numbers addition is commutative, meaning $x_i + y_i = y_i + x_i$ for $i = 1, 2$. So $\vec{x} + \vec{y} = \vec{y} + \vec{x}$.

Exercise 2: The zero vector is the line $y = x + 0$. Indeed, if $y = x + d$ is any line in V , then

$$(y = x + d) + (y = x + 0) = (y = x + (d + 0)) = (y = x + d).$$

So $y = x + 0$ satisfies axiom 2.

Exercise 3: For part (a), notice that $0 = 0 + 0$, and so we have

$$0 \cdot \vec{x} = (0 + 0) \cdot \vec{x} = 0 \cdot \vec{x} + 0 \cdot \vec{x},$$

where in the last step we used axiom 6. Next, add $-(0 \cdot \vec{x})$ to both sides to get

$$\begin{aligned} \vec{0} &= -(0 \cdot \vec{x}) + (0 \cdot \vec{x} + 0 \cdot \vec{x}) && \text{(axiom 3 used on LHS)} \\ &= (-(0 \cdot \vec{x}) + 0 \cdot \vec{x}) + 0 \cdot \vec{x} && \text{(axiom 1)} \\ &= \vec{0} + 0 \cdot \vec{x} && \text{(axiom 3)} \\ &= 0 \cdot \vec{x}. && \text{(axiom 2)} \end{aligned}$$

This completes the proof of part (a).

For part (b), add \vec{x} and $(-1) \cdot \vec{x}$ together:

$$\vec{x} + (-1) \cdot \vec{x} = (1) \cdot \vec{x} + (-1) \cdot \vec{x} = (1 + (-1)) \cdot \vec{x} = 0 \cdot \vec{x}$$

and then use part (a).

Finally, for part (c), note that $\vec{0} = \vec{0} + \vec{0}$ so that

$$t \cdot (\vec{0} + \vec{0}) = t \cdot \vec{0} + t \cdot \vec{0}.$$

Now proceed as in part (a).

Exercise 4: Since U is a subspace, it must be non-empty, so there is a $\vec{u} \in U$. Then, by closure under scalar multiplication, $0 \cdot \vec{u}$ will be in U too. Now we can appeal to Proposition 1.1.14(a): If we view this scalar multiplication as occurring in U , then $0 \cdot \vec{u} = \vec{0}_U$. On the other hand, we can also view this scalar multiplication as occurring in V , since $U \subseteq V$. Thus, $0 \cdot \vec{u} = \vec{0}_V$. It follows then that $\vec{0}_V = 0 \cdot \vec{u} = \vec{0}_U$.

In particular, $\vec{0}_V$ is in U (since $\vec{0}_U$ is).

Exercise 5: By equating coefficients of 1 and x on both sides, we get the system of equations

$$\begin{aligned} a + b &= c \\ a - b &= d. \end{aligned}$$

In matrix form, this is

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

We can quickly solve this equation by inverting the matrix on the left-side (which is indeed invertible!):

$$\begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}^{-1} \begin{bmatrix} c \\ d \end{bmatrix} = \frac{1}{-2} \begin{bmatrix} -1 & -1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} c \\ d \end{bmatrix} = \frac{1}{2} \begin{bmatrix} c + d \\ c - d \end{bmatrix}.$$

Thus, $a = \frac{c+d}{2}$ and $b = \frac{c-d}{2}$, as given in the Example. The second half of the exercise is similar.

Exercise 6: Parts (a) and (b) are essentially identical, and we can solve both at once if we work over \mathbb{F} . To check for linear independence, we have to solve the equation

$$c_1 \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} + c_2 \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + c_3 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

This equation is equivalent to

$$\begin{bmatrix} c_1 + c_2 + c_3 \\ c_1 + c_2 \\ c_1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

From the bottom entry, we get $c_1 = 0$, and then from the middle entry we get $c_2 = 0$, and finally from the first entry, we get $c_3 = 0$. Thus, $c_1 = c_2 = c_3 = 0$ and our set is linearly independent. To check for spanning, we wish to show that given any $[a \ b \ c]^T \in \mathbb{F}^3$, we can find $c_1, c_2, c_3 \in \mathbb{F}$ such that

$$c_1 \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} + c_2 \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + c_3 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} a \\ b \\ c \end{bmatrix}.$$

This leads us to the system of equations

$$\begin{aligned}c_1 + c_2 + c_3 &= a \\c_1 + c_2 &= b \\c_1 &= c.\end{aligned}$$

We can solve this to find: $c_1 = c$, $c_2 = b - c$ and $c_3 = a - b - c$. This proves that our set spans \mathbb{F}^3 and hence is a basis for \mathbb{F}^3 .

For part (c), we begin by noting that the set is clearly linearly independent, since it contains two vectors that are not scalar multiples of each other. To check for spanning, we want to show that for all $[a \ b]^T \in \mathbb{C}^2$, the equation

$$c_1 \begin{bmatrix} 2 \\ 0 \end{bmatrix} + c_2 \begin{bmatrix} 0 \\ 3i \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix}$$

can be solved for $c_1, c_2 \in \mathbb{C}$. This is indeed possible, since we can take $c_1 = \frac{a}{2}$ and $c_2 = \frac{b}{3i}$. Thus, the set in (c) is a basis for \mathbb{C}^2 .

Finally, for part (d), again we note that $\{1+x, 1-x\}$ is linearly independent since it contains two vectors that are not scalar multiples. We proved that $\{1+x, 1-x\}$ spans $\mathcal{P}_1(\mathbb{R})$ in Example 1.2.11. Thus, it is a basis for $\mathcal{P}_1(\mathbb{R})$.

Exercise 7: Since $\dim(\mathcal{P}_2(\mathbb{F})) = 3$ and since S contains exactly 3 vectors, it suffices to prove that S is linearly independent. So consider the equation

$$c_0 p_0(x) + c_1 p_1(x) + c_2 p_2(x) = 0 + 0x + 0x^2.$$

The only term of degree 2 on the left-side occurs in $p_2(x)$. If $p_2(x) = ax^2 + \dots$ (with $a \neq 0$), then by equating coefficients of x^2 on both sides, we get $c_2 a = 0$. Hence $c_2 = 0$ since $a \neq 0$. So our equation is reduced to

$$c_0 p_0(x) + c_1 p_1(x) = 0 + 0x + 0x^2.$$

Now repeat the same argument with coefficients of x to get that $c_1 = 0$. Then finally, and in the same way, $c_0 = 0$. Thus $c_0 = c_1 = c_2 = 0$, and so S is linearly independent, as required.

Exercise 8: *Omitted.*

Exercise 9: Suppose $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$. Given $\vec{x}, \vec{y} \in V$, we can write them in the form

$$\vec{x} = \sum_{i=1}^n a_i \vec{v}_i \quad \text{and} \quad \vec{y} = \sum_{i=1}^n b_i \vec{v}_i,$$

for some $a_i, b_i \in \mathbb{F}$. This means:

$$[\vec{x}]_{\mathcal{B}} = \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} \quad \text{and} \quad [\vec{y}]_{\mathcal{B}} = \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix}.$$

Now,

$$\vec{x} + \vec{y} = \sum_{i=1}^n a_i \vec{v}_i + \sum_{i=1}^n b_i \vec{v}_i = \sum_{i=1}^n (a_i + b_i) \vec{v}_i$$

and therefore

$$[\vec{x} + \vec{y}]_{\mathcal{B}} = \begin{bmatrix} a_1 + b_1 \\ \vdots \\ a_n + b_n \end{bmatrix}.$$

The above is clearly equal to $[\vec{x}]_{\mathcal{B}} + [\vec{y}]_{\mathcal{B}}$.

Similarly, if $t \in \mathbb{F}$, then

$$t\vec{x} = t \sum_{i=1}^n a_i \vec{v}_i = \sum_{i=1}^n (ta_i) \vec{v}_i.$$

So

$$[t\vec{x}]_{\mathcal{B}} = \begin{bmatrix} ta_1 \\ \vdots \\ ta_n \end{bmatrix} = t \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} = t[\vec{x}]_{\mathcal{B}}.$$

Chapter 2: Linear Transformations

Exercise 10: For $p, q \in \mathcal{P}_n(\mathbb{F})$, we have $(p + q)(x) = p(x) + q(x)$, and so

$$\text{ev}_t(p + q) = (p + q)(t) = p(t) + q(t) = \text{ev}_t(p) + \text{ev}_t(q).$$

Similarly, if $c \in \mathbb{F}$, then

$$\text{ev}_t(cp) = (cp)(t) = c(p(t)) = c\text{ev}_t(p).$$

Thus, ev_t is linear.

Exercise 11: We have

$$L(\vec{0}_V) = L(0 \cdot \vec{0}_V) = 0 \cdot L(\vec{0}_V) = \vec{0}_W,$$

where we have used Proposition 1.1.14 (once in V and once in W).

Exercise 12: For (a), note that $L([2 \ 0 \ 0]^T) \neq 2L([1 \ 0 \ 0]^T)$. For (b), note that $L(2 + 2x + 2x^2) \neq 2L(1 + x + x^2)$.

Exercise 13: This follows from the familiar rules of differentiation and integration:

$$(f + g)' = f' + g', \quad (cf)' = cf', \quad \int f + g = \int f + \int g, \quad \text{and} \quad \int cf = c \int f.$$

But you can also check this directly using the given formulas for D and I .

Exercise 14: We have $\vec{0}_W \in \text{Range}(L)$ because $L(\vec{0}_V) = \vec{0}_W$. For closure under addition, suppose that $\vec{w}, \vec{w}' \in \text{Range}(L)$. Then $\vec{w} = L(\vec{v})$ and $\vec{w}' = L(\vec{v}')$ for some $\vec{v}, \vec{v}' \in V$ and therefore

$$\vec{w} + \vec{w}' = L(\vec{v}) + L(\vec{v}') = L(\vec{v} + \vec{v}')$$

so $\vec{w} + \vec{w}' \in \text{Range}(L)$. Similarly,

$$c\vec{w} = cL(\vec{v}) = L(c\vec{v})$$

so $c\vec{w} \in \text{Range}(L)$ and $\text{Range}(L)$ is closed under scalar multiplication. So by the Subspace Test, $\text{Range}(L)$ is a subspace of W .

Exercise 15: For part (a), we have

$$\begin{aligned} c[L]_{\mathcal{B}} &= \left[\left[L \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \right]_c \quad \left[L \left(\begin{bmatrix} 0 \\ 1 \end{bmatrix} \right) \right]_c \right] \\ &= \left[[1 + x + x^2]_c \quad [x + 2x^2]_c \right] \\ &= \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{bmatrix}. \end{aligned}$$

For part (b), let's take $\vec{v} = [1 \ 1]^T$ for example. Then $L(\vec{v}) = 1 + 2x + 3x^2$ so

$$[L(\vec{v})]_c = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}.$$

On the other hand,

$$c[L]_{\mathcal{B}} [\vec{v}]_{\mathcal{B}} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix},$$

which is indeed equal to $[L(\vec{v})]_c$!

Exercise 16: Let $\mathcal{B} = \{\vec{b}_1, \dots, \vec{b}_n\}$. Then the i th column of $_{\mathcal{D}}[M \circ L]_{\mathcal{B}}$ is

$$[M \circ L(\vec{b}_i)]_{\mathcal{D}} = [M(L(\vec{b}_i))]_{\mathcal{D}}.$$

Now, the key fact about $_{\mathcal{D}}[M]_c$ is that we have

$$[M(\vec{u})]_{\mathcal{D}} = {}_{\mathcal{D}}[M]_c [\vec{u}]_c \text{ for all } \vec{u} \in U.$$

Applying this to the above equation with $\vec{u} = L(\vec{b}_i)$, we obtain

$$[M \circ L(\vec{b}_i)]_{\mathcal{D}} = {}_{\mathcal{D}}[M]_c [L(\vec{b}_i)]_c.$$

Now, using the same key property but applied to ${}_c[L]_{\mathcal{B}}$, we have

$$[L(\vec{b}_i)]_c = {}_c[L]_{\mathcal{B}} [\vec{b}_i]_{\mathcal{B}}.$$

Here, however, we also have that $[\vec{b}_i]_{\mathcal{B}} = \vec{e}_i$ is the i th standard basis vector in \mathbb{F}^n . So ${}_c[L]_{\mathcal{B}} [\vec{b}_i]_{\mathcal{B}} = {}_c[L]_{\mathcal{B}} \vec{e}_i$ is the i th column of ${}_c[L]_{\mathcal{B}}$.

So, putting all of this together, we arrive at the following:

$$\begin{aligned} i\text{th column of } {}_{\mathcal{D}}[M \circ L]_{\mathcal{B}} &= [M \circ L(\vec{b}_i)]_{\mathcal{D}} \\ &= {}_{\mathcal{D}}[M]_c (i\text{th column of } {}_c[L]_{\mathcal{B}}) \\ &= i\text{th column of } ({}_{\mathcal{D}}[M]_c {}_c[L]_{\mathcal{B}}), \end{aligned}$$

where in the last step we used the definition of matrix multiplication. It follows from this that the i th columns of ${}_{\mathcal{D}}[M \circ L]_{\mathcal{B}}$ and ${}_{\mathcal{D}}[M]_c {}_c[L]_{\mathcal{B}}$ are equal for all i , so the two matrices themselves are equal.

Exercise 17: Let $\vec{w} \in W$. Then

$$\begin{aligned} \vec{w} \in \text{Range}(L) &\iff \vec{w} = L(\vec{v}) \text{ for some } \vec{v} \in V \\ &\iff [\vec{w}]_{\mathcal{C}} = [L(\vec{v})]_{\mathcal{C}} \quad (\text{by the Unique Representation Theorem 1.3.23}) \\ &\iff [\vec{w}]_{\mathcal{C}} = {}_c[L]_{\mathcal{B}}[\vec{v}]_{\mathcal{B}} \\ &\iff [\vec{w}]_{\mathcal{C}} \in \text{Col}({}_c[L]_{\mathcal{B}}). \end{aligned}$$

Exercise 18: We must show that \mathcal{D} is linearly independent and spans $\text{Range}(L)$. For linear independence, consider the equation

$$a \begin{bmatrix} 1 & 1 \\ 3 & 0 \end{bmatrix} + b \begin{bmatrix} 1 & -1 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

Applying the coordinate map $[\]_{\mathcal{C}}$ (which is linear), the above becomes

$$a \begin{bmatrix} 1 \\ 1 \\ 3 \\ 0 \end{bmatrix} + b \begin{bmatrix} 1 \\ -1 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

Since we know that $\{[1 \ 1 \ 3 \ 0]^T, [1 \ -1 \ 1 \ 0]^T\}$ is linearly independent (as it's a basis for $\text{Col}(A)$), it follows that $a = b = 0$. So \mathcal{D} is linearly independent.

For spanning, let $\vec{w} \in \text{Range}(L)$. Then $[\vec{w}]_{\mathcal{C}} \in \text{Col}(A)$, so

$$[\vec{w}]_{\mathcal{C}} = a \begin{bmatrix} 1 \\ 1 \\ 3 \\ 0 \end{bmatrix} + b \begin{bmatrix} 1 \\ -1 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} a+b \\ a-b \\ 3a+b \\ 0 \end{bmatrix}$$

for some $a, b \in \mathbb{R}$. Converting from \mathcal{C} -coordinates back to a vector in $M_{2 \times 2}(\mathbb{R})$, we get

$$\vec{w} = \begin{bmatrix} a+b & a-b \\ 3a+b & 0 \end{bmatrix} = a \begin{bmatrix} 1 & 1 \\ 3 & 0 \end{bmatrix} + b \begin{bmatrix} 1 & -1 \\ 1 & 0 \end{bmatrix}.$$

Thus, $\vec{w} \in \text{Span}(\mathcal{D})$, as required.

Exercise 19: To show that ${}_c\mathcal{I}_{\mathcal{B}} = ({}_{\mathcal{B}}\mathcal{I}_{\mathcal{C}})^{-1}$, it suffices to show that ${}_c\mathcal{I}_{\mathcal{B}} {}_{\mathcal{B}}\mathcal{I}_{\mathcal{C}}$ is the identity matrix. But observe that

$$\begin{aligned} {}_c\mathcal{I}_{\mathcal{B}} {}_{\mathcal{B}}\mathcal{I}_{\mathcal{C}} &= {}_c[\text{id}]_{\mathcal{B}} {}_{\mathcal{B}}[\text{id}]_{\mathcal{C}} \\ &= {}_c[\text{id}]_{\mathcal{C}} \end{aligned} \quad (\text{by Proposition 2.3.8})$$

and ${}_c[\text{id}]_{\mathcal{C}}$ is the $m \times m$ identity matrix, where $m = \dim(W)$.

Exercise 20: For part (a), note that L is injective if and only if $\text{Ker}(L) = \{\vec{0}\}$ if and only if $\dim(\text{Ker}(L)) = 0$ (since a non-zero subspace has non-zero dimension) and so the result follows since $\text{nullity}(L) = \dim(\text{Ker}(L))$.

For part (b), note that L is surjective if and only if $\text{Range}(L) = W$. Since $\text{Range}(L)$ is a subspace of W , $\dim(\text{Range}(L)) \leq \dim(W)$ with equality if and only if $\text{Range}(L) = W$. Putting the previous two sentences together, we conclude that L is surjective if and only

if $\dim(\text{Range}(L)) = \dim(W)$. Since $\text{rank}(L) = \dim(\text{Range}(L))$, this is exactly what we wanted to prove.

Exercise 21: We have

$$A = \begin{bmatrix} [1]_C & [1+x]_C & [x+x^2]_C & [x^2]_C \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}.$$

Row reducing, we find that

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

From this, we see that $\text{rank}(A) = 3$, $\text{nullity}(A) = 1$,

$$\text{Col}(A) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right\}$$

and

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

Therefore $\text{Col}(A) = \mathbb{R}^3$, so $\text{Range}(L) = \mathcal{P}_2(\mathbb{R})$ and L is surjective. On the other hand, $\text{Ker}(L) \neq \{\vec{0}\}$ (it is spanned by $\begin{bmatrix} -1 & 1 \\ -1 & 1 \end{bmatrix}$), so L is not injective.

Exercise 22: We have

$$a + bx + cx^2 \in \text{Ker}(L) \iff L(a + bx + cx^2) = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \iff \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

So the only vector in $\text{Ker}(L)$ is $0 + 0x + 0x^2$. Thus, $\text{Ker}(L) = \{\vec{0}\}$.

On the other hand, it's clear that $\text{Range}(L) = \mathbb{R}^3$, since given any $\vec{v} = [a \ b \ c]^T \in \mathbb{R}^3$, we have $L(a + bx + cx^2) = \vec{v}$.

Exercise 23: For part (a), note that

$$a + bx + cx^2 \in \text{Ker}(L) \iff L(a + bx + cx^2) = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \iff \begin{bmatrix} a \\ a + b \\ a + b + c \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

Equating components and solving the resulting equations, we find that $a = b = c = 0$. So $\text{Ker}(L) = \{\vec{0}\}$ and thus L is injective.

Next, to show that L is surjective, let $\vec{v} = [v_1 \ v_2 \ v_3] \in \mathbb{R}^3$ be an arbitrary vector. We wish to find a polynomial $a + bx + cx^2 \in \mathcal{P}_2(\mathbb{R})$ such that $L(a + bx + cx^2) = \vec{v}$. This amounts to solving the system of equations

$$\begin{aligned} a &= v_1 \\ a + b &= v_2 \end{aligned}$$

$$a + b + c = v_3$$

for $a, b, c \in \mathbb{R}$. Doing so, we find the solution $a = v_1$, $b = v_2 - v_1$ and $c = v_3 - v_2 - v_1$. Thus L is indeed surjective and therefore an isomorphism.

For part (b), there are many examples. An easy one is: $L(a + bx + cx^2) = \begin{bmatrix} a \\ c \\ b \end{bmatrix}$.

Exercise 24: Let $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$. If $[\vec{x}]_{\mathcal{B}} = [0 \ \dots \ 0]^T$, then $\vec{x} = \sum_{i=1}^n 0\vec{v}_i = \vec{0}$. So $\text{Ker}([\]_{\mathcal{B}}) = \{\vec{0}\}$ and $[\]_{\mathcal{B}}$ is injective. We can also invoke the Unique Representation Theorem and note that the *only* vector that has coordinates $[0 \ \dots \ 0]^T$ is the zero vector.

Since $[\]_{\mathcal{B}}$ is injective, the Rank–Nullity Theorem now tells us that

$$n = \dim(V) = \text{nullity}([\]_{\mathcal{B}}) + \text{rank}([\]_{\mathcal{B}}) = 0 + \text{rank}([\]_{\mathcal{B}}).$$

That is, $\dim(\text{Range}([\]_{\mathcal{B}})) = n$. But $\text{Range}([\]_{\mathcal{B}})$ is a subspace of \mathbb{R}^n , so if its dimension is equal to n then it must itself be all of \mathbb{R}^n . So $\text{Range}([\]_{\mathcal{B}}) = \mathbb{R}^n$ and $[\]_{\mathcal{B}}$ is surjective.

We can also give a very easy direct proof of surjectivity: if $\vec{v} = [a_1 \ \dots \ a_n]$ is an arbitrary vector in \mathbb{R}^n , then let $\vec{x} = \sum_{i=1}^n a_i \vec{v}_i$. This is a vector in V with $[\vec{x}]_{\mathcal{B}} = \vec{v}$. So $\vec{v} \in \text{Range}([\]_{\mathcal{B}})$.

Exercise 25: Since the coordinate map is an isomorphism, it suffices to prove that the standard coordinate vectors of the vectors in \mathcal{B} form a basis for \mathbb{R}^4 . That is, we wish to prove that

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

is a basis for \mathbb{R}^4 . Since \mathcal{B}' contains exactly $4 = \dim(\mathbb{R}^4)$ vectors, it suffices to prove that \mathcal{B}' is linear independent. This is a straightforward exercise and so the details are omitted.

Exercise 26: The Rank–Nullity theorem says that $\dim(V) = \text{rank}(L) + \text{nullity}(L)$. If $\dim(V) < \dim(W)$, then $\text{rank}(L) \leq \dim(V) < \dim(W)$ too. So $\text{rank}(L) \neq \dim(W)$ and thus L cannot be surjective.

If $\dim(V) > \dim(W)$ then $\text{nullity}(L) > \dim(W) - \text{rank}(L)$. Note however that $\text{rank}(L) \leq \dim(W)$ since $\text{rank}(L) = \dim(\text{Range}(L))$ and $\text{Range}(L)$ is a subspace of W . Consequently, $\dim(W) - \text{rank}(L) \geq 0$. So we conclude that $\text{nullity}(L) > 0$, and thus L cannot be injective.

Exercise 27: The recipe tells us to pick bases for $M_{2 \times 2}(\mathbb{R})$ and $\mathcal{P}_3(\mathbb{R})$ and then map one to the other. More precisely, let $\mathcal{B} = \{E_{11}, E_{12}, E_{21}, E_{22}\}$ and $\mathcal{C} = \{1, x, x^2, x^3\}$ be the standard bases for $M_{2 \times 2}(\mathbb{R})$ and $\mathcal{P}_3(\mathbb{R})$ and then map E_{11} to 1, E_{12} to x , etc.

To spell this out, the desired isomorphism $L: M_{2 \times 2}(\mathbb{R}) \rightarrow \mathcal{P}_3(\mathbb{R})$ is defined by

$$L\left(\begin{bmatrix} a & b \\ c & d \end{bmatrix}\right) = L(aE_{11} + bE_{12} + cE_{21} + dE_{22}) = a1 + bx + cx^2 + dx^3.$$

Exercise 28: First we must check that L^{-1} , defined as in the proof, is linear. So let $\vec{w}, \vec{u} \in W$ and $c \in \mathbb{F}$. Then there are $\vec{x}, \vec{y} \in V$ such that $L(\vec{x}) = \vec{w}$ and $L(\vec{y}) = \vec{u}$, and by definition $L^{-1}(\vec{w}) = \vec{x}$ and $L^{-1}(\vec{u}) = \vec{y}$. Now,

$$L(c\vec{x} + \vec{y}) = cL(\vec{x}) + L(\vec{y}) = c\vec{w} + \vec{u}.$$

Therefore, $L^{-1}(c\vec{w} + \vec{u}) = c\vec{x} + \vec{y} = cL^{-1}(\vec{w}) + L^{-1}(\vec{u})$, which shows that L^{-1} is linear.

Next, suppose we are given that there exists a map $L^{-1}: W \rightarrow V$ such that $L \circ T(\vec{w}) = \vec{w}$ for all $\vec{w} \in W$ and $T \circ L(\vec{v}) = \vec{v}$ for all $\vec{v} \in V$. We must prove that this implies that L is injective and surjective. For injectivity, note that

$$L(\vec{v}) = \vec{0} \implies L^{-1}(L(\vec{v})) = \vec{0} \implies \vec{v} = \vec{0}.$$

Thus $\text{Ker}(L) = \{\vec{0}\}$ and L is injective.

For surjectivity, let $\vec{w} \in W$, and note that

$$\vec{w} = L(L^{-1}(\vec{w}))$$

so $\vec{w} \in \text{Range}(L)$, showing that $\text{Range}(L) = W$.

Finally, we must prove that if L^{-1} exists, then it is itself an isomorphism. In fact this follows from what we've just proved. Reversing the roles of L and L^{-1} , our proof above shows that L^{-1} is injective and surjective.

Chapter 3: Diagonalizability

Exercise 29: Since $L(\vec{0}) = \vec{0} = \lambda\vec{0}$, $\vec{0} \in E_\lambda(L)$. Let $\vec{v}, \vec{w} \in W$ and $t \in \mathbb{F}$. Then

$$L(\vec{v} + \vec{w}) = L(\vec{v}) + L(\vec{w}) = \lambda\vec{v} + \lambda\vec{w} = \lambda(\vec{v} + \vec{w})$$

so $\vec{v} + \vec{w} \in E_\lambda(L)$. Finally,

$$L(t\vec{v}) = tL(\vec{v}) = t\lambda\vec{v} = \lambda(t\vec{v})$$

so $t\vec{v} \in E_\lambda(L)$. Since $E_\lambda(L)$ is non-empty, closed under addition, and closed under scalar multiplication, $E_\lambda(L)$ is a subspace of V by the Subspace Test.

Exercise 30: If $\vec{v} \in E_1(P)$, then $P(\vec{v}) = \vec{v}$, meaning P projects \vec{v} onto itself—so \vec{v} must've been on W to begin with. That is, $\vec{v} \in W$. This completes part (a).

For part (b), the vectors in $E_0(P)$ are precisely the ones that get projected onto $\vec{0}$. Thus they are all perpendicular to W . So if \vec{n} is a normal vector to W , then $E_0(P) = \text{Span}\{\vec{n}\}$.

Exercise 31: Simply observe that

$$L(\vec{v}) = \lambda\vec{v} \iff [L(\vec{v})]_{\mathcal{B}} = [\lambda\vec{v}]_{\mathcal{B}} \iff [L]_{\mathcal{B}}[\vec{v}]_{\mathcal{B}} = \lambda[\vec{v}]_{\mathcal{B}}.$$

Also keep in mind that \vec{v} is non-zero if and only if $[\vec{v}]_{\mathcal{B}}$ is non-zero, by the Unique Representation Theorem.

Exercise 32: Proceeding just as in Example 3.1.9, we find that $E_0(D) = \text{Span}\{1\}$ is the space of constant polynomials in $P_n(\mathbb{R})$.

Exercise 33: Direct calculations. *Details omitted.*

Exercise 34: For part (a), note that $A = I_n^{-1}AI_n$, where I_n is the $n \times n$ identity matrix. For part (b), if $A = P^{-1}BP$, then $B = PAP^{-1} = (P^{-1})^{-1}A(P^{-1})$. For part (c), if $A = P^{-1}BP$ and $B = Q^{-1}CQ$, then $A = P^{-1}(Q^{-1}CQ)P = (QP)^{-1}C(QP)$.

Exercise 35: *Omitted.*

Exercise 36: We need eigenvectors for A for the eigenvalues λ_1 and λ_2 . We can find some by computing $\text{Null}(A - \lambda_1 I)$ and $\text{Null}(A - \lambda_2 I)$. Doing so, we obtain the eigenvectors $\vec{v}_1 = [-i \ 1]^T$ and $\vec{v}_2 = [i \ 1]^T$, resp. Converting back to polynomials, we have the eigenvectors $p_1 = -i + x$ and $p_2 = i + x$ for L with eigenvalues $\lambda_1 = 1 + 2i$ and $\lambda_2 = 2 - i$, resp. So if we take $\mathcal{D} = \{-i + x, i + x\}$ then $[L]_{\mathcal{D}} = \begin{bmatrix} 1 + 2i & 0 \\ 0 & 1 - 2i \end{bmatrix}$.

Chapter 4: Inner Product Spaces

Exercise 37: In the given diagram, let l denote the length of the hypotenuse of the right-triangle in the xy -plane with side-lengths v_1 and v_2 . Then, by the Pythagorean theorem,

$$l^2 = v_1^2 + v_2^2.$$

On the other hand, we can view l as being the side length of the vertical right triangle with hypotenuse \vec{v} . So again, by the Pythagorean theorem,

$$\|\vec{v}\|^2 = l^2 + v_3^2.$$

By combining both of these equations and taking square roots, we arrive at

$$\|\vec{v}\| = \sqrt{v_1^2 + v_2^2 + v_3^2},$$

as required.

The expression for $\cos \theta$ again comes from the cosine rule:

$$\cos \theta = \frac{\|\vec{v}\|^2 + \|\vec{w}\|^2 - \|\vec{v} - \vec{w}\|^2}{2 \|\vec{v}\| \|\vec{w}\|}.$$

If we expand the right-side and simplify, we will arrive at the desired expression. [A shortcut is possible here if we use the dot product:

$$\|\vec{v} - \vec{w}\|^2 = (\vec{v} - \vec{w}) \cdot (\vec{v} - \vec{w}) = \|\vec{v}\|^2 + \|\vec{w}\|^2 - 2(\vec{v} \cdot \vec{w}).]$$

Exercise 38: Part (a) is a straightforward computation and is therefore omitted. For part (b), take $\vec{z} = [i \ 0]^T$. Then $\vec{z} \cdot \vec{z} = i^2 + 0^2 = -1$.

Exercise 39: For part (b),

$$\begin{aligned} \langle \vec{v}, \alpha \vec{w} \rangle &= \overline{\langle \alpha \vec{w}, \vec{v} \rangle} && \text{(conjugate symmetry)} \\ &= \overline{\alpha \langle \vec{w}, \vec{v} \rangle} && \text{(linearity in 1st argument)} \\ &= \overline{\alpha} \overline{\langle \vec{w}, \vec{v} \rangle} \\ &= \overline{\alpha} \langle \vec{v}, \vec{w} \rangle. && \text{(conjugate symmetry)} \end{aligned}$$

For part (c),

$$\begin{aligned} \langle \vec{v}, \vec{u} + \vec{w} \rangle &= \overline{\langle \vec{u} + \vec{w}, \vec{v} \rangle} && \text{(conjugate symmetry)} \\ &= \overline{\langle \vec{u}, \vec{v} \rangle + \langle \vec{w}, \vec{v} \rangle} && \text{(linearity in 1st argument)} \end{aligned}$$

$$\begin{aligned}
&= \overline{\langle \vec{u}, \vec{v} \rangle} + \overline{\langle \vec{w}, \vec{v} \rangle} \\
&= \langle \vec{v}, \vec{u} \rangle + \langle \vec{v}, \vec{w} \rangle. \qquad \text{(conjugate symmetry)}
\end{aligned}$$

Exercise 40: We have $\langle p, q \rangle = \int_{-1}^1 p(x)q(x) dx = \int_{-1}^1 q(x)p(x) dx = \langle q, p \rangle$, and this is equal to $\overline{\langle q, p \rangle}$ since everything here is real. This proves axiom 1. Axioms 2 and 3 follow from the linearity of the integral.

Exercise 41: For axiom 1, notice that $\langle A, B \rangle = \text{tr}(B^T A) = \text{tr}((B^T A)^T) = \text{tr}(A^T B) = \langle B, A \rangle$, where we've used the fact that the trace of a matrix is equal to the trace of the transpose of that matrix. Axioms 2 and 3 follow from the linearity of trace and matrix multiplication. Finally, for axiom 3, note that $\langle A, A \rangle = \sum_{i=1}^n \sum_{j=1}^n a_{ij}a_{ij} = \sum_{i,j} a_{ij}^2$ is the sum of the squares of the entries of A . So $\langle A, A \rangle \geq 0$ and it's equal to zero if and only if all of the entries of A are equal to zero, that is, if and only if $A = 0$. This completes the proof that the Frobenius inner product is an inner product.

Finally, if $n = 1$, then the Frobenius inner product is just the dot product on \mathbb{R}^m !

Exercise 42: The proof that $\langle A, B \rangle$ defines an inner product is very similar to the proof in the previous exercise. For part (b), the given definition is not linear in the first variable, since

$$\langle cA, B \rangle = \text{tr}((cA)^* B) = \text{tr}(\bar{c}A^* B) = \bar{c} \text{tr}(A^* B) = \bar{c} \langle A, B \rangle.$$

In particular, if c is not real (e.g. if $c = i$), then $\langle cA, B \rangle \neq c \langle A, B \rangle$.

Exercise 43: *Omitted.* This is effectively the same as checking that the standard inner product on \mathbb{F}^n is an inner product.

Exercise 44: We have $\|\vec{0}\| = \sqrt{\langle \vec{0}, \vec{0} \rangle} = \sqrt{0} = 0$.

Exercise 45: We have

$$\|\alpha \vec{v}\| = \sqrt{\langle \alpha \vec{v}, \alpha \vec{v} \rangle} = \sqrt{\alpha \bar{\alpha} \langle \vec{v}, \vec{v} \rangle} = \sqrt{|\alpha|^2 \langle \vec{v}, \vec{v} \rangle} = |\alpha| \|\vec{v}\|.$$

Next, $\|\vec{v}\| = \sqrt{\langle \vec{v}, \vec{v} \rangle}$ is certainly non-negative; it will be equal to zero if and only if $\langle \vec{v}, \vec{v} \rangle = 0$ which is the case if and only if $\vec{v} = \vec{0}$, by positive definiteness of the inner product. This completes part (a).

For part (b), by examining the proof of the Triangle Inequality, we see that the inequality will be an equality if and only if $|\langle \vec{v}, \vec{w} \rangle| = \|\vec{v}\| \|\vec{w}\|$. This is precisely the equality condition for the Cauchy–Schwarz inequality! So we conclude that the Triangle Inequality will be an equality if and only if \vec{v} and \vec{w} are scalar multiples of each other.

Exercise 46: Part (a) follows from the fact that $\|\vec{x} - \vec{y}\| \geq 0$ with equality if and only if $\vec{x} - \vec{y} = \vec{0}$. Part (b) follows from $\|\vec{x} - \vec{y}\| = \|(-1)(\vec{x} - \vec{y})\| = \|\vec{y} - \vec{x}\|$. And part (c) follows from the Triangle Inequality for the norm:

$$\begin{aligned}
\text{dist}(\vec{x}, \vec{z}) &= \|\vec{x} - \vec{z}\| \\
&= \|\vec{x} - \vec{y} + \vec{y} - \vec{z}\| \\
&\leq \|\vec{x} - \vec{y}\| + \|\vec{y} - \vec{z}\| \\
&= \text{dist}(\vec{x}, \vec{y}) + \text{dist}(\vec{y}, \vec{z}).
\end{aligned}$$

Exercise 47: For part (a), notice that

$$\|\hat{v}\| = \left\| \frac{\vec{v}}{\|\vec{v}\|} \right\| = \frac{1}{\|\vec{v}\|} \|\vec{v}\| = 1.$$

Part (b) follows from part (a) together with the observation that

$$\langle \hat{v}_i, \hat{v}_j \rangle = \left\langle \frac{\vec{v}_i}{\|\vec{v}_i\|}, \frac{\vec{v}_j}{\|\vec{v}_j\|} \right\rangle = \frac{1}{\|\vec{v}_i\| \|\vec{v}_j\|} \langle \vec{v}_i, \vec{v}_j \rangle = 0$$

for $i \neq j$.

Exercise 48: These are all routine computations, so we will only deal with the set given in Example 4.3.12. Since $\dim(\mathbb{C}^2) = 2$, all we have to do is check that the set is orthonormal, as this will automatically imply linear independence and hence that the set is a basis.

The two given vectors are orthogonal, since

$$\left\langle \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\rangle = 2 \frac{1}{\sqrt{2}} \bar{0} + (0) \bar{1} = 0,$$

and they each have unit norm:

$$\left\| \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \end{bmatrix} \right\| = \sqrt{\left\langle \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \end{bmatrix}, \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \end{bmatrix} \right\rangle} = \sqrt{2 \left(\frac{1}{\sqrt{2}} \right)^2 + 0^2} = 1,$$

$$\left\| \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\| = \sqrt{\left\langle \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\rangle} = \sqrt{2(0)^2 + 1^2} = 1.$$

This proves that $\left\{ \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}$ is an orthonormal basis for \mathbb{C}^2 , as required.

Exercise 49: Let's begin by finding an orthogonal basis for W . To make our lives easier, let's put \vec{v} in this basis. So all we have to do is find a non-zero vector \vec{w} that is orthogonal to \vec{v} . Then $\{\vec{v}, \vec{w}\}$ will be an orthogonal basis for W (it will be a basis because $\dim(W) = 2$).

Now, any vector in W will be of the form $\vec{u} = a\vec{v} + b\vec{w}$ for some $a, b \in \mathbb{F}$. If we want \vec{u} to be orthogonal to \vec{v} , then we must have

$$0 = \langle a\vec{v} + b\vec{w}, \vec{v} \rangle = a \langle \vec{v}, \vec{v} \rangle + b \langle \vec{w}, \vec{v} \rangle.$$

We want to find *some* $a, b \in \mathbb{F}$ for which this is true. Let's take $a = 1$ and try to find an appropriate b . Solving for b , we find that $b = -\frac{\langle \vec{v}, \vec{v} \rangle}{\langle \vec{w}, \vec{v} \rangle}$, at least provided $\langle \vec{w}, \vec{v} \rangle \neq 0$. However if $\langle \vec{w}, \vec{v} \rangle = 0$ then $\vec{w} \perp \vec{v}$ and we're done: $\{\vec{v}, \vec{w}\}$ is already an orthogonal basis. So we may assume that $\langle \vec{w}, \vec{v} \rangle \neq 0$. Then our computation above shows that $\{\vec{v}, \vec{v} - \frac{\langle \vec{v}, \vec{v} \rangle}{\langle \vec{w}, \vec{v} \rangle} \vec{w}\}$ is an orthogonal basis for W .

To obtain an orthonormal basis, we can simply normalize the above orthogonal basis.

Exercise 50: We can proceed by induction on i . If $i = 1$, then $\vec{w}_1 = \vec{v}_1$ and so obviously $\text{Span}(\{\vec{v}_1\}) = \text{Span}(\{\vec{w}_1\})$.

So assume the result holds for some $i = k \geq 1$, and consider the case $i = k + 1$. So, in particular, we have that $\text{Span}(\{\vec{v}_1, \dots, \vec{v}_k\}) = \text{Span}(\{\vec{w}_1, \dots, \vec{w}_k\})$. Now, according to the algorithm,

$$\vec{w}_{k+1} = \vec{v}_{k+1} - \frac{\langle \vec{v}_{k+1}, \vec{w}_1 \rangle}{\|\vec{w}_1\|^2} \vec{w}_1 - \dots - \frac{\langle \vec{v}_{k+1}, \vec{w}_k \rangle}{\|\vec{w}_k\|^2} \vec{w}_k.$$

So \vec{w}_{k+1} is a linear combination of \vec{v}_{k+1} and $\vec{w}_1, \dots, \vec{w}_k$, hence is a linear combination of \vec{v}_{k+1} and $\vec{v}_1, \dots, \vec{v}_k$, by the induction hypothesis. Thus, $\vec{w}_k \in \text{Span}(\{\vec{v}_1, \dots, \vec{v}_{k+1}\})$ and therefore

$$\text{Span}(\{\vec{w}_1, \dots, \vec{w}_{k+1}\}) \subseteq \text{Span}(\{\vec{v}_1, \dots, \vec{v}_{k+1}\}).$$

Conversely, from the above equation, we see that

$$\vec{v}_{k+1} = \vec{w}_{k+1} + \frac{\langle \vec{v}_{k+1}, \vec{w}_1 \rangle}{\|\vec{w}_1\|^2} \vec{w}_1 - \dots - \frac{\langle \vec{v}_{k+1}, \vec{w}_k \rangle}{\|\vec{w}_k\|^2} \vec{w}_k$$

so $\vec{v}_{k+1} \in \text{Span}(\{\vec{w}_1, \dots, \vec{w}_{k+1}\})$, hence

$$\text{Span}(\{\vec{v}_1, \dots, \vec{v}_{k+1}\}) \subseteq \text{Span}(\{\vec{w}_1, \dots, \vec{w}_{k+1}\}).$$

It follows that

$$\text{Span}(\{\vec{v}_1, \dots, \vec{v}_{k+1}\}) = \text{Span}(\{\vec{w}_1, \dots, \vec{w}_{k+1}\}).$$

This completes the induction.

For the second part in the exercise, all we have to do is observe that $\{\vec{w}_1, \dots, \vec{w}_n\}$ is an orthogonal set of n vectors that spans V (thanks to what we just proved above). So it must be a basis too (since $\dim(V) = n$).

Exercise 51: $\left\{ \frac{1}{\sqrt{2}}, \sqrt{\frac{3}{2}}x, \sqrt{\frac{5}{8}}(3x^2 - 1), \sqrt{\frac{7}{8}}(5x^3 - 3x), \sqrt{\frac{9}{128}}(35x^4 - 30x^2 + 3) \right\}$. *Details omitted.*

Exercise 52: From [Exercise 50](#), we know that $\{\vec{w}_1, \dots, \vec{w}_k\}$ is an orthogonal basis for $\text{Span}(S)$. Hence, if we apply [Proposition 4.3.13](#) to $\vec{v}_{k+1} \in \text{Span}(S)$, we get

$$\vec{v}_{k+1} = \frac{\langle \vec{v}_{k+1}, \vec{w}_1 \rangle}{\|\vec{w}_1\|^2} \vec{w}_1 + \dots + \frac{\langle \vec{v}_{k+1}, \vec{w}_k \rangle}{\|\vec{w}_k\|^2} \vec{w}_k.$$

So

$$\vec{v}_{k+1} - \frac{\langle \vec{v}_{k+1}, \vec{w}_1 \rangle}{\|\vec{w}_1\|^2} \vec{w}_1 - \dots - \frac{\langle \vec{v}_{k+1}, \vec{w}_k \rangle}{\|\vec{w}_k\|^2} \vec{w}_k = \vec{0}.$$

Exercise 53: For part (a), note that $\vec{0} \in W^\perp$ since $\vec{0} \perp \vec{w}$ for all $\vec{w} \in W$. Next, if $\vec{x}, \vec{y} \in W^\perp$ and $c \in \mathbb{F}$, then for all $\vec{w} \in W$ we have

$$\langle c\vec{x} + \vec{y}, \vec{w} \rangle = c\langle \vec{x}, \vec{w} \rangle + \langle \vec{y}, \vec{w} \rangle = c(0) + 0 = 0.$$

So $c\vec{x} + \vec{y} \in W^\perp$. So, by the Subspace Test, W^\perp is a subspace of V .

For part (b), suppose that $\vec{x} \in W \cap W^\perp$. Then $\vec{x} \in W$ and $\vec{x} \in W^\perp$, so $\vec{x} \perp \vec{x}$, meaning $\langle \vec{x}, \vec{x} \rangle = 0$. So, by the positive definiteness of the inner product, it must be the case that $\vec{x} = \vec{0}$.

Exercise 54: Since $\{\vec{v}_1, \dots, \vec{v}_n\}$ is an orthogonal set, it's clear that $\{\vec{v}_{k+1}, \dots, \vec{v}_n\} \subseteq W^\perp$, and hence that $\text{Span}\{\vec{v}_{k+1}, \dots, \vec{v}_n\} \subseteq W^\perp$ since W^\perp is a subspace.

For the reverse containment, take $\vec{v} \in W^\perp$. Then $\vec{v} \in V$ and since $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$ is a basis for V we can write

$$\vec{v} = \sum_{i=1}^n a_i \vec{v}_i$$

where, since \mathcal{B} is an *orthogonal* basis, we have

$$a_i = \frac{\langle \vec{v}, \vec{v}_i \rangle}{\|\vec{v}_i\|^2}.$$

Now since $\vec{v} \in W^\perp$ and $\vec{v}_i \in W$ for $i = 1, \dots, k$, it must be the case that $a_i = 0$ for $i = 1, \dots, k$. Hence $\vec{v} = \sum_{i=k+1}^n a_i \vec{v}_i$ is in $\text{Span}\{\vec{v}_{k+1}, \dots, \vec{v}_n\}$. This proves that $W^\perp \subseteq \text{Span}\{\vec{v}_{k+1}, \dots, \vec{v}_n\}$, as desired.

Exercise 55: Suppose without loss of generality that \vec{z}_1 and \vec{z}_2 are not orthogonal, so that $\langle \vec{z}_1, \vec{z}_2 \rangle \neq 0$. Take $\vec{v} = \vec{z}_1$, so that $\text{proj}_W(\vec{v}) = \vec{v} = \vec{z}_1$, since \vec{v} is in W . We also have $\text{proj}_{\vec{z}_1} \vec{v} = \vec{z}_1$ and $\text{proj}_{\vec{z}_2} \vec{v} = c \vec{z}_2$ for some non-zero constant c (namely $c = \frac{\langle \vec{z}_1, \vec{z}_2 \rangle}{\langle \vec{z}_2, \vec{z}_2 \rangle}$). Then $\text{proj}_W(\vec{v}) \neq \text{proj}_{\vec{z}_1}(\vec{v}) + \dots + \text{proj}_{\vec{z}_k}(\vec{v})$ since otherwise we would have

$$\vec{z}_1 = \vec{z}_1 + c \vec{z}_2 + \dots + \frac{\langle \vec{z}_1, \vec{z}_k \rangle}{\langle \vec{z}_k, \vec{z}_k \rangle} \vec{z}_k$$

which contradicts the linear independence of $\{\vec{z}_1, \dots, \vec{z}_k\}$ since $c \neq 0$.

Exercise 56: Part (a) is really a calculus exercise. The orthogonality of S_n follows from the following trigonometric integrals:

$$\int_{-\pi}^{\pi} \cos(kx) dx = 0 \quad (\text{if } k \neq 0)$$

$$\int_{-\pi}^{\pi} \cos(kx) \sin(lx) dx = 0$$

where $k, l \in \mathbb{Z}$. (Note, in particular, that $\cos(kx) = 1$ when $k = 0$.)

For part (b), we have

$$\text{proj}_{W_n}(f) = \sum_{k=0}^n a_k \cos(kx) + \sum_{l=1}^n b_l \sin(lx),$$


where

$$a_k = \frac{\langle |x|, \cos(kx) \rangle}{\langle \cos(kx), \cos kx \rangle} = \frac{1}{\pi} \int_{-\pi}^{\pi} |x| \cos(kx) dx = \begin{cases} \pi & \text{if } k = 0 \\ 0 & \text{if } k > 0 \text{ is even} \\ -\frac{4}{\pi k^2} & \text{if } k > 0 \text{ is odd} \end{cases}$$

and

$$b_l = \frac{\langle |x|, \sin(lx) \rangle}{\langle \sin(lx), \sin lx \rangle} = \frac{1}{\pi} \int_{-\pi}^{\pi} |x| \sin(lx) dx = 0.$$

Below is a plot of both $f(x) = |x|$ and the projections $\text{proj}_{W_n}(f)$ for $n = 1, 3, 9$.



Exercise 57: If $\vec{w} = \text{proj}_W(\vec{v})$, then $\vec{v} - \vec{w} = \text{proj}_{W^\perp}(\vec{v})$ by the Orthogonal Decomposition Theorem, so $\vec{v} - \vec{w} \in W^\perp$. Conversely, if $\vec{v} - \vec{w} \in W^\perp$, then $\text{proj}_W(\vec{v} - \vec{w}) = \vec{0}$ hence $\text{proj}_W(\vec{v}) - \text{proj}_W(\vec{w}) = \vec{0}$ by linearity of proj_W . But $\text{proj}_W(\vec{w}) = \vec{w}$ since $\vec{w} \in W$, so we get $\text{proj}_W(\vec{v}) - \vec{w} = \vec{0}$ and consequently $\vec{w} = \text{proj}_W(\vec{v})$ as required.

Exercise 58: By Proposition 4.7.6, we only need to determine whether the columns of X are linearly independent. If $n < 3$, then X being an $n \times 3$ matrix means its rank is less than 3, so its columns cannot be linearly independent. Hence $X^T X$ cannot be invertible.

On the other hand, if $n \geq 3$ then the columns of X will be linearly independent if and only if $\text{rank}(X) = 3$, which will be the case if and only if three of the rows of X are linearly independent. So now we have to determine when three rows of the form $[1 \ x_i \ x_i^2]$, $[1 \ x_j \ x_j^2]$ and $[1 \ x_k \ x_k^2]$ are linearly independent. This will be the case if and only if the matrix

$$A = \begin{bmatrix} 1 & x_i & x_i^2 \\ 1 & x_j & x_j^2 \\ 1 & x_k & x_k^2 \end{bmatrix}$$

is invertible. If we compute the determinant of A , we find that

$$\det(A) = (x_j - x_i)(x_k - x_i)(x_k - x_j).$$

(This is an example of a **Vandermonde determinant**.) Thus, A is invertible if and only if $\det(A) \neq 0$, which is the case if and only if x_i, x_j, x_k are distinct. This completes the proof.

Chapter 5: Unitary Diagonalization

Exercise 59: Consider $V = \mathcal{P}_1(\mathbb{R})$ with inner product $\langle p, q \rangle = \int_0^1 pq$. Let $\mathcal{B} = \{1, x\}$ be the standard basis. Note that $\langle 1, x \rangle = \frac{1}{2}$ while $[x]_{\mathcal{B}}^* [1]_{\mathcal{B}} = [0 \ 1] \begin{bmatrix} 1 \\ 0 \end{bmatrix} = 0$.

Exercise 60: These follow at once from the properties of the transpose, plus the following facts about complex conjugation:

$$\overline{z + w} = \bar{z} + \bar{w}, \quad \overline{zw} = \bar{z}\bar{w}, \quad \text{and} \quad \overline{\bar{z}} = z.$$

Exercise 61: *No solution provided.*

Exercise 62: Since the rows of P are the columns of P^T , the equivalence of parts (a) and (b) in Proposition 5.2.1 allows us to conclude that the rows of P form an orthonormal bases for \mathbb{F}^n if and only if $(P^T)^* = (P^T)^{-1}$. This in turn is equivalent to $(P^*)^T = (P^{-1})^T$ and hence to $P^* = P^{-1}$.

Exercise 63: The angle between \vec{v} and \vec{u} is given by

$$\theta = \cos^{-1} \left(\frac{\langle \vec{v}, \vec{u} \rangle}{\|\vec{v}\| \|\vec{u}\|} \right).$$

The angle between $Q\vec{v}$ and $Q\vec{u}$ is given by

$$\theta' = \cos^{-1} \left(\frac{\langle Q\vec{v}, Q\vec{u} \rangle}{\|Q\vec{v}\| \|Q\vec{u}\|} \right).$$

Since Q is orthogonal, Proposition 5.2.7 tells us that $\langle Q\vec{v}, Q\vec{u} \rangle = \langle \vec{v}, \vec{u} \rangle$, $\|Q\vec{v}\| = \|\vec{v}\|$, $\|Q\vec{u}\| = \|\vec{u}\|$, so $\theta' = \theta$.

Exercise 64: In $\det(A - \lambda I)$, λ is in \mathbb{F} , so we cannot substitute in the matrix A .

Exercise 65: If A is self-adjoint, then AA^* and A^*A are both equal to A^2 .

If A is unitary, then $AA^* = A^*A = I$.

Finally, if $A = \text{diag}(\lambda_1, \dots, \lambda_n)$, then $AA^* = A^*A = \text{diag}(\lambda_1 \bar{\lambda}_1, \dots, \lambda_n \bar{\lambda}_n)$.

Exercise 66: The diagonal entries of A are a_{ii} . The diagonal entries of A^* are \bar{a}_{ii} . Since $A = A^*$, it follows that $a_{ii} = \bar{a}_{ii}$, and hence that $a_{ii} \in \mathbb{R}$.

Exercise 67: If $A \in M_{n \times n}(\mathbb{R})$ is orthogonally diagonalizable, then we can write $A = QDQ^T$ with $Q, D \in M_{n \times n}(\mathbb{R})$ that are unitary and diagonal, respectively. Then $A^T = (QDQ^T)^T = QD^TQ^T$. Further, since D is diagonal, $D^T = D$, so $A^T = A$.

Exercise 68: If $A \in M_{n \times n}(\mathbb{C})$ is unitarily diagonalizable, then we can write $A = UDU^*$ with $U, D \in M_{n \times n}(\mathbb{C})$ that are unitary and diagonal, respectively. Then $A^* = (UDU^*)^* = UD^*U$. Hence $A^*A = UD^*DU$ and $AA^* = UDD^*U$. Note that $D^* = \bar{D}$ is diagonal, so $DD^* = D^*D$ since diagonal matrices commute. Thus, $A^*A = AA^*$.

Exercise 69: For part (a), note that if A is skew-self-adjoint then A is normal, because AA^* and A^*A are both equal to $-A^2$. Hence, by the spectral theorem for normal operators, A is unitarily diagonalizable.

For part (b), we follow our unitary diagonalization algorithm and find $U = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$ and

$$D = \begin{bmatrix} i & 0 \\ 0 & -i \end{bmatrix}.$$

Exercise 70: This is just a matter of multiplying out $\vec{u}^T A \vec{u}$.

Exercise 71: The (i, j) th entry of A is $\frac{a_{ij} + a_{ji}}{2}$. This is clearly equal to the (j, i) th entry. So A is symmetric.

Exercise 72: Let A be the Gram matrix with respect to \mathcal{B} , and let $\mathcal{B} = \{\vec{g}_1, \dots, \vec{g}_n\}$. Then the (i, j) th entry of A is $A_{ij} = \langle \vec{g}_j, \vec{g}_i \rangle$. So if \mathcal{B} is orthonormal, we see that $A_{ij} = 1$ if $i = j$ and $A_{ij} = 0$ if $i \neq j$ —meaning, $A = I_n$.

If \mathcal{B} is orthogonal, then we still have that $A_{ij} = 0$ for all $i \neq j$, but $A_{ii} = \langle \vec{v}_i, \vec{v}_i \rangle = \|\vec{v}_i\|^2$. So in this case A is a diagonal matrix whose diagonal entries are the norms-squared of the vectors in \mathcal{B} .

Exercise 73: The Gram matrix of $\langle \cdot, \cdot \rangle$ with respect to the standard basis \mathcal{B} of \mathbb{C}^2 is

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

The eigenvalues of A can be computed to be $\frac{1}{2}(1 \pm \sqrt{5})$. One of these is negative. So $\langle \cdot, \cdot \rangle$ is not an inner product (specifically, it is not positive-definite).

Chapter 6: The Singular Value Decomposition

Exercise 74: If A is self-adjoint then $A^*A = A^2$.

We claim that the eigenvalues of A are squares of the eigenvalues of A .

This follows from the spectral theorem (which can be applied because A is self-adjoint). If $A = UDU^*$ then $A^2 = UD^2U^*$. The eigenvalues of A are the diagonal entries of D and the eigenvalues of A^2 are the diagonal entries of D^2 . Since the entries of D^2 are the squares of the entries of D , this proves our claim. Thus, the eigenvalues of A^*A are of the form λ^2 where λ is an eigenvalue of A .

From this we deduce that if σ is a singular value of A , then $\sigma = \sqrt{\lambda^2} = |\lambda|$ where λ is an eigenvalue of A . That is, *the singular values of A are the absolute values of the eigenvalues of A* .

In the above unitary diagonalization of $A^*A = A^2$, we found that the unitary matrix U is the same one that unitarily diagonalizes A . Thus, *the singular vectors of A are eigenvectors of A* .

Exercise 75: The matrix V is obtained from applying the spectral theorem to the self-adjoint matrix A^*A . If A is real, then $A^*A = A^T A$ is real symmetric, so we can guarantee V to be real, by the spectral theorem for symmetric matrices. The matrix U is built out of two types of vectors. First, we use vectors of the form $\frac{1}{\sigma} A \vec{v}$ where σ is a singular value and \vec{v} is a column of V —both of which are real. The second type of vectors are obtained by applying the Gram–Schmidt process in \mathbb{R}^n to the first type of vectors, which also results in real vectors.

Exercise 76: Refer to [Exercise 74](#). An SVD $A = U\Sigma V^*$ and unitary diagonalization $A = WDW^*$ may share the unitary matrices (i.e. we may have $U = V = W$). The key difference lies in D and Σ : the diagonal entries of Σ are the *absolute values* of the diagonal entries of D . So if we have a unitarily diagonalization of A we can easily obtain an SVD, but not conversely (not unless we know what the eigenvalues of A are).

Exercise 77: If you multiply out $\vec{u} \vec{v}^*$ you will obtain an $n \times n$ matrix whose rows are all multiples of \vec{v}^* . Since \vec{v}^* is non-zero, it thus forms a basis for the row-space of $\vec{u} \vec{v}^*$, and so $\text{rank}(\vec{u} \vec{v}^*) = 1$.

Exercise 78: Note that, with respect to the Frobenius inner product,

$$\langle \vec{u}_i \vec{v}_i^*, \vec{u}_j \vec{v}_j^* \rangle = \text{tr}((\vec{u}_j \vec{v}_j^*)^* \vec{u}_i \vec{v}_i^*) = \text{tr}(\vec{v}_j \vec{u}_j^* \vec{u}_i \vec{v}_i^*).$$

But $\vec{u}_j^* \vec{u}_i = 0$ for $i \neq j$, since this is the standard inner product of \vec{u}_j and \vec{u}_i —and these vectors are orthogonal since they are distinct columns of a unitary matrix. Similarly, notice that if $i = j$ in the above, then

$$\vec{v}_i \vec{u}_i^* \vec{u}_i \vec{v}_i^* = \vec{v}_i (1) \vec{v}_i^* = \vec{v}_i \vec{v}_i^*$$

(why?) so

$$\langle \vec{u}_i \vec{v}_i^*, \vec{u}_i \vec{v}_i^* \rangle = \text{tr}(\vec{v}_i \vec{v}_i^*) = 1$$

just as we observed at the end of the proof of part (c) of Proposition 6.4.6.

Thus, $\{\vec{u}_1 \vec{v}_1^*, \dots, \vec{u}_r \vec{v}_r^*\}$ is an orthonormal set with respect to the Frobenius inner product. So, by the Pythagorean theorem, we have

$$\begin{aligned} \|A - A_k\|^2 &= \left\| \sum_{i=k+1}^r \sigma_i \vec{u}_i \vec{v}_i^* \right\|^2 \\ &= \sum_{i=k+1}^r |\sigma_i|^2 \|\vec{u}_i \vec{v}_i^*\|^2 \\ &= \sum_{i=k+1}^r |\sigma_i|^2 \\ &= \sum_{i=k+1}^r \sigma_i^2 \end{aligned} \quad (\text{since } \sigma_i \geq 0)$$

from which the desired result follows.

Exercise 79: The key point is that if $A = U\Sigma V^*$ is an SVD of A , then $A^{-1} = V\Sigma^{-1}U^*$ is an SVD of A^{-1} . Since Example 6.4.11 shows that $\Sigma^\dagger = \Sigma^{-1}$, it follows that $A^\dagger = A^{-1}$.

Index

- $E_\lambda(L)$, 61
- additive inverse, 9
- adjoint, 122, 124
- algorithm
 - diagonalization, 75
 - Gram–Schmidt procedure, 104
 - singular value decomposition, 162
 - SVD, 162
 - unitary diagonalization, 139
- angle, 93
- basis, 20
 - ordered, 28
 - orthogonal, 97
 - orthonormal, 97
- Cauchy–Schwarz Inequality, 92
- Cayley–Hamilton Theorem, 131
- characteristic polynomial, 64
- column space, 41
- compact singular value decomposition, 169
- conjugate symmetry, 85
- coordinate
 - change of, 46
 - relative to an orthogonal basis, 98
 - vector, 29
- cosine, 83, 90, 93, 100
- decoupling, 79
- diagonalizability
 - orthogonal, 134
 - unitary, 134
- diagonalizable, 68, 71, 75
 - orthogonally, 134
 - unitarily, 134
- diagonalization
 - orthogonal, 139
 - unitary, 139
- diagonalizes, 68, 71
- dimension, 23
- distance, 94
 - properties, 95
- Eckart–Young Theorem, 173
- eigenspace, 61, 63
- eigenvalue, 61, 63, 64
- eigenvector, 61, 63
- evaluation map, 32
- \mathbb{F} , 7
- field, 7
- Fourier expansion, 115
- Gram matrix, 153
- Gram–Schmidt procedure, 104
- identity
 - map, 31
 - transformation, 31
- image compression, 174
- indefinite, 145
- injective, 48
- inner product, 85
 - dot product, 85
 - space, 85
 - standard, 86
- inverse, 57, 58
- isomorphic to, 51
- isomorphism, 51
- kernel, 34
- least squares, 116, 174
 - curve fitting, 117
 - minimal norm, 176
 - solution, 116
- Legendre polynomials, 96
- length, 89
- linear

- combination, 16
 - map, 31
 - operator, 61
 - transformation, 31
- linearity, 31
 - in first argument, 85
- linearly independent, 18
- matrix
 - adjoint, 122
 - orthogonal, 126
 - unitary, 126
 - associated to a quadratic form, 146
 - change of coordinates, 46
 - Gram, 153
 - Hessian, 148
 - normal, 134, 138
 - of a linear map, 39
 - of an inner product, 149
 - self-adjoint, 134
 - skew-self-adjoint, 140
 - symmetric, 134
- minimal norm, 176
- $M_{m \times n}(\mathbb{F})$, 8
- multiplicity, 66, 159
 - algebraic, 74
 - geometric, 74
- negative
 - definite, 145
 - semi-definite, 145
- norm, 89
 - properties, 93
- normalization, 97
- nullity, 35, 41
- nullspace, 34, 41
- one-to-one, 48
- onto, 48
- ordered basis, 28
- orthogonal, 91
 - complement, 110
 - decomposition, 112
 - set, 95
- orthogonal complement, 110
- orthonormal
 - set, 95
- outer product, 171
- perpendicular
 - with respect to a subspace, 113
 - with respect to a vector, 101
- perp_W , 113
- $\text{perp}_{\vec{w}}$, 101
- $\mathcal{P}_n(\mathbb{F})$, 7
- positive
 - definite, 85, 145
 - semi-definite, 145
- projection, 62, 100, 101
 - onto a plane, 101
 - onto a subspace, 113
 - onto a vector, 101
- proj_W , 113
- $\text{proj}_{\vec{w}}$, 101
- pseudoinverse, 175
- Pythagorean Theorem, 91
- quadratic form, 144
 - classification, 147
- range, 34
- rank, 35, 41
- Rank–Nullity Theorem, 35
- scalar multiplication, 9
- Schur’s Triangularization Theorem, 128
- similar, 69
- singular
 - value, 157
 - value decomposition, 160, 161, 166
 - value decomposition, compact, 169
 - vector, 157
- span, 16
- spanning set, 18
- spans, 18
- Spectral Theorem, 135–137
 - for normal matrices, 137
 - for operators, 141
 - for Self-Adjoint matrices, 135
 - for symmetric matrices, 136
- subspace, 13
 - Subspace Test, 14
- surjective, 48
- SVD, 160, 161, 166
 - compact, 169
 - geometry, 165
 - image compression, 174
 - low-rank approximation, 168
 - rank- k truncation, 171
- Unique Representation Theorem, 28
- vector

coordinate vector, 29
addition, 9
unit, 95
zero, 9
vector space, 9

dimension, 23
finite-dimensional, 23
infinite-dimensional, 23
subspace, 13
axioms, 9