Linear Algebra

Sara Pohland

Created: January 3, 2021
Last Modified: December 14, 2023

Contents

Part I

Vectors

Chapter 1

Vector Basics

1.1 Overview of Vectors

1.1.1 Cartesian Product

Given two sets X and Y, the **Cartesian product** of these sets is denoted $X \times Y$ and is defined as the set of all ordered pairs (x, y) such that $x \in X$ and $y \in Y$. Expressed more formally, the Cartesian product of sets X and Y is

$$X \times Y = \{(x, y) : x \in X,\ y \in Y\}.$$

The Euclidean n-space is the set of ordered n-tuples of real numbers, defined as

$$\mathbb{R}^n = \{(x_1, x_2, \ldots, x_n) : x_i \in \mathbb{R},\ i = 1, \ldots, n\}.$$

Similarly, the set of all ordered n-tuples of complex numbers is expressed as

$$\mathbb{C}^n = \{(x_1, x_2, \ldots, x_n) : x_i \in \mathbb{C},\ i = 1, \ldots, n\}.$$

1.1.2 Column and Row Vectors

An n-dimensional vector, $\boldsymbol{x}$, is composed of n scalar values, $x_1, \ldots, x_n \in \mathbb{C}$. If $\boldsymbol{x}$ is a column vector, then it has the following form:

$$\boldsymbol{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \in \mathbb{C}^n.$$

Similarly, if $\boldsymbol{x}$ is an n-dimensional row vector, then it has the following form:

$$\boldsymbol{x} = \begin{bmatrix} x_1 & \ldots & x_n \end{bmatrix} \in \mathbb{C}^n.$$

In general, if a vector is given without any indication of whether it is a row or column vector, it is assumed to be a column vector.

I will use $\mathbf{0}_n$ to denote the n-dimensional zero vector, which is a column vector composed of all zeros. Similarly, I will use $\mathbf{1}_n$ to denote the n-dimensional ones vector, which is a column vector composed of all ones.

1.1.3 Vector Transpose

The **transpose** of a column vector is a row vector with the same elements, and the transpose of a row vector is the corresponding column vector. The transpose of a vector $\boldsymbol{x}$ is typically denoted as $\boldsymbol{x}^T$ and sometimes as $\boldsymbol{x}'$. For example,

$$\boldsymbol{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \iff \boldsymbol{x}^T = \begin{bmatrix} x_1 & \dots & x_n \end{bmatrix}$$

and

$$\boldsymbol{x} = \begin{bmatrix} x_1 & \dots & x_n \end{bmatrix} \iff \boldsymbol{x}^T = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}.$$

1.1.4 Conjugate Transpose

The **conjugate transpose** of a vector is defined as its transpose, where all of its elements have been replaced with their conjugates. The conjugate transpose of a vector $\boldsymbol{x}$ is typically denoted as $\boldsymbol{x}^*$. For example,

$$\boldsymbol{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \iff \boldsymbol{x}^* = \begin{bmatrix} \bar{x}_1 & \dots & \bar{x}_n \end{bmatrix}$$

and

$$\boldsymbol{x} = \begin{bmatrix} x_1 & \dots & x_n \end{bmatrix} \iff \boldsymbol{x}^* = \begin{bmatrix} \bar{x}_1 \\ \vdots \\ \bar{x}_n \end{bmatrix}.$$

Chapter 2

Vector Subspaces

2.1 Definition of Vector Subspace

2.1.1 Fields

A **field** is defined as an object consisting of a set of elements and two binary operations: addition and multiplication. In order for F to be considered a valid field, its elements must satisfy the following set of axioms.

1. The addition operation must...
   (a) be associative (i.e. $(\alpha + \beta) + \gamma = \alpha + (\beta + \gamma),\ \forall \alpha, \beta, \gamma \in F$),
   (b) be commutative (i.e. $\alpha + \beta = \beta + \alpha,\ \forall \alpha, \beta \in F$),
   (c) have an identity element 0 (i.e. $\exists\ 0 \in F$ s.t. $\alpha + 0 = \alpha,\ \forall \alpha \in F$),
   (d) and have an inverse (i.e. $\exists\ -\alpha \in F$ s.t. $\alpha + (-\alpha) = 0,\ \forall \alpha \in F$).
2. The multiplication operation must...
   (a) be associative (i.e. $(\alpha \cdot \beta) \cdot \gamma = \alpha \cdot (\beta \cdot \gamma),\ \forall \alpha, \beta, \gamma \in F$),
   (b) be commutative (i.e. $\alpha \cdot \beta = \beta \cdot \alpha,\ \forall \alpha, \beta \in F$),
   (c) have an identity element 1 (i.e. $\exists\ 1 \in F$ s.t. $\alpha \cdot 1 = \alpha,\ \forall \alpha \in F$),
   (d) and have an inverse (i.e. $\exists\, \alpha^{-1} \in F$ s.t. $\alpha \cdot \alpha^{-1} = 1,\ \forall \alpha \in F :\ \alpha \neq 0$).
3. The operations together must...
   (a) be distributive (i.e. $\alpha \cdot (\beta + \gamma) = \alpha \cdot \beta + \alpha \cdot \gamma,\ \forall \alpha, \beta, \gamma \in F$ and $(\beta + \gamma) \cdot \alpha = \beta \cdot \alpha + \gamma \cdot \alpha,\ \forall \alpha, \beta, \gamma \in F$).

Some examples of fields are the real line, $\mathbb{R}$, the complex plane, $\mathbb{C}$, the set of rational functions in s with coefficients in $\mathbb{R}$, $\mathbb{R}(s)$, and the set of rational functions in s with coefficients in $\mathbb{C}$, $\mathbb{C}(s)$.

Some examples of sets that are not fields are the set of integers, $\mathbb{Z}$, the set of polynomials in s with coefficients in $\mathbb{R}$, $\mathbb{R}[s]$, and strictly proper rational functions in s, $\mathbb{R}_P(s)$. Not all of the elements in each of these sets have inverses that are also contained in that set, so these sets are not considered fields.

2.1.2 Vector Spaces

A **vector space**, which is also referred to as a **linear space**, is defined as a set of vectors, V, and a field of scalars, F, equipped with vector addition and scalar multiplication. In order for (V, F) to be considered a valid vector space, its elements must satisfy the following set of axioms.

1. The vector addition operation must...
   (a) be associative (i.e. $(\boldsymbol{v_1}+\boldsymbol{v_2})+\boldsymbol{v_3} = \boldsymbol{v_1}+(\boldsymbol{v_2}+\boldsymbol{v_3}),\ \forall \boldsymbol{v_1}, \boldsymbol{v_2}, \boldsymbol{v_3} \in V$),
   (b) be commutative (i.e. $\boldsymbol{v_1} + \boldsymbol{v_2} = \boldsymbol{v_2} + \boldsymbol{v_1},\ \forall \boldsymbol{v_1}, \boldsymbol{v_2} \in V$),
   (c) have an identity element $\mathbf{0}$ (i.e. $\exists\ \mathbf{0} \in V$ s.t. $\boldsymbol{v} + \mathbf{0} = \boldsymbol{v},\ \forall \boldsymbol{v} \in V$),
   (d) and have an inverse (i.e. $\exists\ -\boldsymbol{v} \in V$ s.t. $\boldsymbol{v} + (-\boldsymbol{v}) = \mathbf{0},\ \forall \boldsymbol{v} \in V$).
2. The scalar multiplication operation must...
   (a) Be associative – $(\alpha \cdot \beta) \cdot \boldsymbol{v} = \alpha \cdot (\beta \cdot \boldsymbol{v}),\ \forall \alpha, \beta \in F,\ \forall \boldsymbol{v} \in V$
   (b) Have a multiplicative identity 1 – $\exists\ 1 \in F$ s.t. $1 \cdot \boldsymbol{v} = \boldsymbol{v},\ \forall \boldsymbol{v} \in V$
   (c) Have an additive identity 0 – $\exists\ 0 \in F$ s.t. $0 \cdot \boldsymbol{v} = \mathbf{0},\ \forall \boldsymbol{v} \in V$
3. The operations together must...
   (a) Be distributive – $(\alpha + \beta) \cdot \boldsymbol{v} = \alpha \cdot \boldsymbol{v} + \beta \cdot \boldsymbol{v},\ \forall \alpha, \beta \in F,\ \forall \boldsymbol{v} \in V$
   $\alpha \cdot (\boldsymbol{v_1} + \boldsymbol{v_2}) = \alpha \cdot \boldsymbol{v_1} + \alpha \cdot \boldsymbol{v_1},\ \forall \alpha \in F,\ \forall \boldsymbol{v_1}, \boldsymbol{v_2} \in V$

(V, F) is a vector space if and only if the set V is closed under vector addition and scalar multiplication, meaning that if you perform these operations on any elements of the vector space, the resulting vector is also within this vector space. We call this property of vector spaces **closure**. Written more formally, this says

$$\alpha_1 \boldsymbol{v_1} + \alpha_2 \boldsymbol{v_2} \in V,\ \forall \alpha_1, \alpha_2 \in F,\ \forall \boldsymbol{v_1}, \boldsymbol{v_2} \in V.$$

One common example of a vector space is the space of n-tuples in $\mathbb{F}^n$ over the scalar field $\mathbb{F}$, where $\mathbb{F}$ is the field of real numbers $\mathbb{R}$ or complex numbers $\mathbb{C}$. Another example of a vector space is the set of continuous functions on some interval $[t_0, t_1]$ over the field of reals, which can be expressed as $\Big(C\big([t_0, t_1], \mathbb{R}^n\big), \mathbb{R}\Big)$. Similarly, the set of k times differentiable functions on $[t_0, t_1]$ over the field of reals, $\Big(C^k\big([t_0, t_1], \mathbb{R}^n\big), \mathbb{R}\Big)$, is a vector space. Note that $(\mathbb{R}^n, \mathbb{C})$ is not a vector space because it is not closed under scalar multiplication.

2.1.3 Vector Subspaces

Let's assume (V, F) is a vector space. The space (W, F) is considered a **vector subspace** of (V, F) if W is a non-empty subset of V and (W, F) is closed under vector addition and scalar multiplication. Written more formally, (W, F) is a vector subspace if it satisfies the following two properties:

1. $W \subseteq V,\ W \neq \emptyset$
2. $\alpha_1 \boldsymbol{w_1} + \alpha_2 \boldsymbol{w_2} \in W,\ \forall \boldsymbol{w_1}, \boldsymbol{w_2} \in W,\ \forall \alpha_1, \alpha_2 \in F$

To show that a set is a subspace, we need to show that both of these properties hold for some vector space V. To show that a set is not a subspace, we simply need to find one example for which the closure property does not hold.

2.2 Properties of Vector Subspaces

2.2.1 Intersection, Union, and Sum

If W and X are two subspaces of the vector space V, then the **intersection** of these two subspaces is defined as

$$W \cap X = \{v \in V : v \in W,\ v \in X\}.$$

Similarly, the **union** of these two subspaces is defined as

$$W \cup X = \{v \in V : v \in W \text{ or } v \in X\}.$$

If $W \cap X = \emptyset$, meaning that there are no elements shared between the two subspaces, then we can define the **direct sum** as

$$W \oplus X = \{w + x : w \in W,\ x \in X\}.$$

If $W \cap X \neq \emptyset$, the **sum** of these two subspaces is simply defined as

$$W + X = \{w + x : w \in W,\ x \in X\}.$$

2.2.2 Linear Independence

Let's assume (V, F) is a valid vector space. A set of vectors $S = \{\boldsymbol{v_1}, \ldots, \boldsymbol{v_m}\}$, where $\boldsymbol{v_i} \in V$ for $i = 1, \ldots, m$, is **linearly independent** if and only if

$$\left[\sum_{i=1}^{m} \alpha_i \boldsymbol{v_i} = \boldsymbol{0}_n\right] \implies \big[\alpha_i = 0,\ \forall i = 1, \ldots, m\big].$$

This says that S is linearly dependent if and only if there exist scalars $\alpha_i \in F$ that are not all equal to zero such that $\sum_{i=1}^{m} \alpha_i \boldsymbol{v_i} = \boldsymbol{0}$. Note that linear independence is with respect to a field. A set of vectors may be linearly independent over one field but linearly dependent over another field.

2.2.3 Basis

Suppose (V, F) is a vector space. The **span** of the set of vectors $B = \{\boldsymbol{b_1}, \ldots, \boldsymbol{b_m}\}$ is the **linear combination** of the vectors in B. This can be expressed as

$$\text{span}(B) = \left\{ \sum_{i=1}^{m} \alpha_i \boldsymbol{b_i} : \alpha_i \in F,\ i = 1, \ldots, m \right\}.$$

A set $B = \{\boldsymbol{b_1}, \ldots, \boldsymbol{b_m}\}$ is a **basis** for V if B satisfies two requirements: (1) B is a linearly independent set and (2) $\text{span}(B) = V$. If B is a basis for V, any vector in V can be expressed as a linear combination of the vectors in B, i.e.

$$[\mathbf{v} \in V] \Longrightarrow \left[v = \sum_{i=1}^{m} \alpha_i \boldsymbol{b_i},\ \alpha_i \in F,\ i = 1, \ldots, m \right].$$

The **standard basis** for the vector space $(\mathbb{R}^n, \mathbb{R})$ is $B = \{\boldsymbol{e_1}, \boldsymbol{e_2}, \ldots, \boldsymbol{e_n}\}$, where

$$\boldsymbol{e_1} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix},\ \boldsymbol{e_2} = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix},\ \ldots,\ \boldsymbol{e_n} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}.$$

The **dimension** of a vector space is the number of elements in the basis, which is m for the example vector space (V, F) and n for $(\mathbb{R}^n, \mathbb{R})$. A subspace can have infinitely many bases, but the dimension of a given subspace is fixed.

Chapter 3

Vector Inner Products & Norms

3.1 Vector Inner Products

3.1.1 Properties of Inner Products

Consider the vector space (V, F). An **inner product** is a function, denoted $\langle \cdot, \cdot \rangle$, between two vectors, which must satisfy the following properties:

1. $\langle \boldsymbol{x}, \boldsymbol{y} + \boldsymbol{z} \rangle = \langle \boldsymbol{x}, \boldsymbol{y} \rangle + \langle \boldsymbol{x}, \boldsymbol{z} \rangle, \ \forall \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in V$
   $\langle \boldsymbol{x} + \boldsymbol{y}, \boldsymbol{z} \rangle = \langle \boldsymbol{x}, \boldsymbol{z} \rangle + \langle \boldsymbol{y}, \boldsymbol{z} \rangle, \ \forall \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in V$
2. $\langle \alpha \boldsymbol{x}, \boldsymbol{y} \rangle = \alpha \langle \boldsymbol{x}, \boldsymbol{y} \rangle, \ \forall \boldsymbol{x}, \boldsymbol{y} \in V, \ \alpha \in F$
   $\langle \boldsymbol{x}, \alpha \boldsymbol{y} \rangle = \bar{\alpha} \langle \boldsymbol{x}, \boldsymbol{y} \rangle, \ \forall \boldsymbol{x}, \boldsymbol{y} \in V, \ \alpha \in F$
3. $||\boldsymbol{x}||^2 := \langle \boldsymbol{x}, \boldsymbol{x} \rangle \geq 0, \ \forall \boldsymbol{x} \in V$
   $\langle \boldsymbol{x}, \boldsymbol{x} \rangle = 0 \iff \boldsymbol{x} = \boldsymbol{0}$
4. $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \overline{\langle \boldsymbol{y}, \boldsymbol{x} \rangle}$
5. $\langle \alpha \boldsymbol{x} + \beta \boldsymbol{y}, \gamma \boldsymbol{z} + \delta \boldsymbol{w} \rangle = \bar{\alpha}\gamma \langle \boldsymbol{x}, \boldsymbol{z} \rangle + \bar{\beta}\gamma \langle \boldsymbol{y}, \boldsymbol{z} \rangle + \bar{\alpha}\delta \langle \boldsymbol{x}, \boldsymbol{w} \rangle + \bar{\beta}\delta \langle \boldsymbol{y}, \boldsymbol{w} \rangle$
   $\forall \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}, \boldsymbol{w} \in V, \ \alpha, \beta, \gamma, \delta \in F$

3.1.2 Inner Product Spaces

An **inner product space** is a vector space equipped with an inner product. The vector space of n-dimensional real-valued vectors, $(\mathbb{R}^n, \mathbb{R})$, is an inner product space equipped with the standard, or Euclidean, inner product:

$$\langle \boldsymbol{x}, \boldsymbol{y} \rangle_{\mathbb{R}^n} = \boldsymbol{x}^T \boldsymbol{y} = \sum_{i=1}^{n} x_i y_i.$$

The vector space of n-dimensional complex-valued vectors, $(\mathbb{C}^n, \mathbb{C})$, is an inner product space equipped with a similar inner product:

$$\langle \boldsymbol{x}, \boldsymbol{y} \rangle_{\mathbb{C}^n} = \boldsymbol{x}^* \boldsymbol{y} = \sum_{i=1}^{n} \bar{x}_i y_i.$$

3.1.3 Cauchy Schwarz

The **Cauchy-Schwarz inequality** states that for all vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ in an inner product space (V, F) with valid inner product $\langle \cdot, \cdot \rangle$, the following holds:

$$|\langle \boldsymbol{x}, \boldsymbol{y} \rangle|^2 \leq \langle \boldsymbol{x}, \boldsymbol{x} \rangle \langle \boldsymbol{y}, \boldsymbol{y} \rangle.$$

Using the third property of the inner product, we can express this inequality as

$$|\langle \boldsymbol{x}, \boldsymbol{y} \rangle| \leq ||\boldsymbol{x}|| ||\boldsymbol{y}||.$$

In the above inequalities, equality holds if and only if either $\boldsymbol{x}$ or $\boldsymbol{y}$ are linearly dependent, meaning one is a scalar multiple of the other.

Proof: We can prove this inequality using properties of the inner product. If $\boldsymbol{x}$ or $\boldsymbol{y}$ is the zero vector, this inequality hold trivially because both sides of the inequality will be zero. Let's then assume that $\boldsymbol{x}$ and $\boldsymbol{y}$ are non-zero.

$$\left\langle \boldsymbol{x} - \frac{\langle \boldsymbol{x}, \boldsymbol{y} \rangle \boldsymbol{y}}{||\boldsymbol{y}||^2}, \boldsymbol{x} - \frac{\langle \boldsymbol{x}, \boldsymbol{y} \rangle \boldsymbol{y}}{||\boldsymbol{y}||^2} \right\rangle \geq 0$$

$$\langle \boldsymbol{x}, \boldsymbol{x} \rangle - \frac{1}{||\boldsymbol{y}||^2} \langle \boldsymbol{x}, \langle \boldsymbol{x}, \boldsymbol{y} \rangle \boldsymbol{y} \rangle - \frac{1}{||\boldsymbol{y}||^2} \langle \langle \boldsymbol{x}, \boldsymbol{y} \rangle \boldsymbol{y}, \boldsymbol{x} \rangle + \frac{1}{||\boldsymbol{y}||^4} \langle \langle \boldsymbol{x}, \boldsymbol{y} \rangle \boldsymbol{y}, \langle \boldsymbol{x}, \boldsymbol{y} \rangle \boldsymbol{y} \rangle \geq 0$$

$$\langle \boldsymbol{x}, \boldsymbol{x} \rangle - \frac{1}{||\boldsymbol{y}||^2} \overline{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} \langle \boldsymbol{x}, \boldsymbol{y} \rangle - \frac{1}{||\boldsymbol{y}||^2} \langle \boldsymbol{x}, \boldsymbol{y} \rangle \langle \boldsymbol{y}, \boldsymbol{x} \rangle + \frac{1}{||\boldsymbol{y}||^4} \langle \boldsymbol{x}, \boldsymbol{y} \rangle \overline{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} \langle \boldsymbol{y}, \boldsymbol{y} \rangle \geq 0$$

$$\langle \boldsymbol{x}, \boldsymbol{x} \rangle - \frac{1}{||\boldsymbol{y}||^2} \overline{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} \langle \boldsymbol{x}, \boldsymbol{y} \rangle - \frac{1}{||\boldsymbol{y}||^2} \langle \boldsymbol{x}, \boldsymbol{y} \rangle \overline{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} + \frac{1}{||\boldsymbol{y}||^4} \langle \boldsymbol{x}, \boldsymbol{y} \rangle \overline{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} ||\boldsymbol{y}||^2 \geq 0$$

$$\langle \boldsymbol{x}, \boldsymbol{x} \rangle - \frac{1}{||\boldsymbol{y}||^2} |\langle \boldsymbol{x}, \boldsymbol{y} \rangle|^2 - \frac{1}{||\boldsymbol{y}||^2} |\langle \boldsymbol{x}, \boldsymbol{y} \rangle|^2 + \frac{1}{||\boldsymbol{y}||^2} |\langle \boldsymbol{x}, \boldsymbol{y} \rangle|^2 \geq 0$$

$$\langle \boldsymbol{x}, \boldsymbol{x} \rangle - \frac{1}{||\boldsymbol{y}||^2} |\langle \boldsymbol{x}, \boldsymbol{y} \rangle|^2 \geq 0$$

$$\langle \boldsymbol{x}, \boldsymbol{x} \rangle \langle \boldsymbol{y}, \boldsymbol{y} \rangle \geq |\langle \boldsymbol{x}, \boldsymbol{y} \rangle|^2$$

From the third inner product property, this inequality becomes an equality iff

$$\boldsymbol{x} - \frac{\langle \boldsymbol{x}, \boldsymbol{y} \rangle \boldsymbol{y}}{||\boldsymbol{y}||^2} = 0 \iff \boldsymbol{x} = \frac{\langle \boldsymbol{x}, \boldsymbol{y} \rangle}{||\boldsymbol{y}||^2} \boldsymbol{y}$$

This says that equality holds if and only if either $\boldsymbol{x}$ or $\boldsymbol{y}$ can be expressed as scalar multiples of each other, meaning that they are linearly dependent.

3.2 Vector Norms

3.2.1 Properties of Vector Norms

Consider the vector space (V, F). In order for a vector function to be considered a **norm**, it must satisfy the following three properties:

1. $||\boldsymbol{x}|| \geq 0, \ \forall \boldsymbol{x} \in V$
   $||\boldsymbol{x}|| = 0 \iff \boldsymbol{x} = \boldsymbol{0}$
2. $||\boldsymbol{x} + \boldsymbol{y}|| \leq ||\boldsymbol{x}|| + ||\boldsymbol{y}||, \ \forall \boldsymbol{x}, \boldsymbol{y} \in V$
3. $||\alpha \boldsymbol{x}|| = |\alpha| ||\boldsymbol{x}||, \ \forall \alpha \in F, \ \forall \boldsymbol{x} \in V$

The second property is commonly referred to as the **triangle inequality**. Vectors also satisfy a similar property called the **reverse triangle inequality**:

$$||\boldsymbol{x} - \boldsymbol{y}|| \geq \Big| ||\boldsymbol{x}|| - ||\boldsymbol{y}|| \Big|, \ \forall \boldsymbol{x}, \boldsymbol{y} \in V.$$

One other property of vector norms is referred to as the **parallelogram law** and states that for all vectors $\boldsymbol{x}, \boldsymbol{y} \in V$,

$$||\boldsymbol{x} + \boldsymbol{y}||^2 + ||\boldsymbol{x} - \boldsymbol{y}||^2 = 2||\boldsymbol{x}||^2 + 2||\boldsymbol{y}||^2.$$

Proof: Notice that from the properties of inner products we can write:

$$\begin{aligned}
||\boldsymbol{x} + \boldsymbol{y}||^2 = \langle \boldsymbol{x} + \boldsymbol{y}, \boldsymbol{x} + \boldsymbol{y} \rangle &= \langle \boldsymbol{x}, \boldsymbol{x} \rangle + \langle \boldsymbol{x}, \boldsymbol{y} \rangle + \langle \boldsymbol{y}, \boldsymbol{x} \rangle + \langle \boldsymbol{y}, \boldsymbol{y} \rangle \\
&= ||\boldsymbol{x}||^2 + ||\boldsymbol{y}||^2 + \langle \boldsymbol{x}, \boldsymbol{y} \rangle + \langle \boldsymbol{y}, \boldsymbol{x} \rangle \\
||\boldsymbol{x} - \boldsymbol{y}||^2 = \langle \boldsymbol{x} - \boldsymbol{y}, \boldsymbol{x} - \boldsymbol{y} \rangle &= \langle \boldsymbol{x}, \boldsymbol{x} \rangle + \langle \boldsymbol{x}, -\boldsymbol{y} \rangle + \langle -\boldsymbol{y}, \boldsymbol{x} \rangle + \langle -\boldsymbol{y}, -\boldsymbol{y} \rangle \\
&= \langle \boldsymbol{x}, \boldsymbol{x} \rangle - \langle \boldsymbol{x}, \boldsymbol{y} \rangle - \langle \boldsymbol{y}, \boldsymbol{x} \rangle + \langle \boldsymbol{y}, \boldsymbol{y} \rangle \\
&= ||\boldsymbol{x}||^2 + ||\boldsymbol{y}||^2 - \langle \boldsymbol{x}, \boldsymbol{y} \rangle - \langle \boldsymbol{y}, \boldsymbol{x} \rangle
\end{aligned}$$

Summing together these two sums, we get the desired result:

$$||\boldsymbol{x} + \boldsymbol{y}||^2 + ||\boldsymbol{x} - \boldsymbol{y}||^2 = 2||\boldsymbol{x}||^2 + 2||\boldsymbol{y}||^2$$

From the proof of the parallelogram law for an arbitrary vector space, we can see that if $\boldsymbol{x}$ and $\boldsymbol{y}$ are real-valued vectors, then

$$||\boldsymbol{x} + \boldsymbol{y}||^2 = ||\boldsymbol{x}||^2 + ||\boldsymbol{y}||^2 + 2\langle \boldsymbol{x}, \boldsymbol{y} \rangle$$

and

$$||\boldsymbol{x} - \boldsymbol{y}||^2 = ||\boldsymbol{x}||^2 + ||\boldsymbol{y}||^2 - 2\langle \boldsymbol{x}, \boldsymbol{y} \rangle.$$

3.2.2 l_p Norms

Consider the inner product space $(\mathbb{R}^n, \mathbb{R})$. Some common vector norms defined on this vector space are the $\boldsymbol{l_p}$ **norms**, which are defined for $1 \leq p \leq \infty$ as

$$||\boldsymbol{x}||_p = \left(\sum_{i=1}^{n} |x_i|^p \right)^{1/p}.$$

The $\boldsymbol{l_2}$ **norm** is the standard **Euclidean length**, which is defined as

$$||\boldsymbol{x}||_2 = \left(\sum_{i=1}^{n} |x_i|^2 \right)^{1/2}.$$

The $\boldsymbol{l_1}$ **norm** is called the **sum-of-absolute values length** and is defined as

$$||\boldsymbol{x}||_1 = \sum_{i=1}^{n} |x_i|.$$

The $\boldsymbol{l_\infty}$ **norm** is the maximum absolute value of a vector and is sometimes referred to as the **Chebyshev distance**. It is defined as

$$||\boldsymbol{x}||_\infty = \max_{i=1,\ldots,n} |x_i|.$$

To see why the l_∞ norm is defined this way, consider the following:

$$\left(\sum_{i=1}^{n} |x_i|^p \right)^{1/p} \leq \left(\sum_{i=1}^{n} \max_{i=1,\ldots,n} |x_i|^p \right)^{1/p} = \left(n \max_{i=1,\ldots,n} |x_i|^p \right)^{1/p} = n^{1/p} \max_{i=1,\ldots,n} |x_i|$$

$$\lim_{p\to\infty} \left(\sum_{i=1}^{n} |x_i|^p \right)^{1/p} \leq \max_{i=1,\ldots,n} |x_i|$$

$$\left(\sum_{i=1}^{n} |x_i|^p \right)^{1/p} \geq \left(\max_{i=1,\ldots,n} |x_i|^p \right)^{1/p} = \max_{i=1,\ldots,n} |x_i|$$

$$\lim_{p\to\infty} \left(\sum_{i=1}^{n} |x_i|^p \right)^{1/p} \geq \max_{i=1,\ldots,n} |x_i|$$

Therefore, the definition of the l_∞ norm we gave is what we would find if we took the limit of the l_p norm as p approaches infinity:

$$\lim_{p\to\infty} ||\boldsymbol{x}||_p = \max_{i=1,\ldots,n} |x_i| =: ||\boldsymbol{x}||_\infty.$$

3.2.3 l_p Norm Balls

The $\boldsymbol{l_p}$ **norm ball** is denoted B_p and is defined a the set of vectors in the Euclidean space whose l_p norm is less than or equal to one, which is written as

$$B_p = \{\boldsymbol{x} \in \mathbb{R}^n : ||\boldsymbol{x}||_p \leq 1\}.$$

For a two-dimensional vector, we can also define the following more general sets:

$\{\boldsymbol{x} \in \mathbb{R}^2 : ||\boldsymbol{x} - \boldsymbol{x_c}||_1 \leq r\}$ – Diamond with diagonal length $2r$ and center at $\boldsymbol{x_c}$

$\{\boldsymbol{x} \in \mathbb{R}^2 : ||\boldsymbol{x} - \boldsymbol{x_c}||_2 \leq r\}$ – Circle with diameter $2r$ and center at $\boldsymbol{x_c}$

$\{\boldsymbol{x} \in \mathbb{R}^2 : ||\boldsymbol{x} - \boldsymbol{x_c}||_\infty \leq r\}$ – Square with side length $2r$ and center at $\boldsymbol{x_c}$

These sets are visualized in figure 3.1. For a three-dimensional vector, the set defined above using the l_1 norm is an octahedron, the set defined using the l_2 norm is a sphere, and the set defined using the l_∞ norm is a cube.

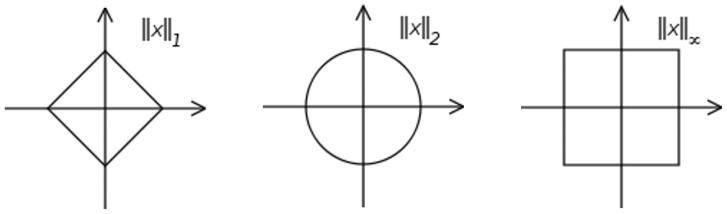

Figure 3.1: For a two-dimensional vector $\boldsymbol{x} \in \mathbb{R}^2$, the l_1, l_2, and l_∞ norm balls are depicted in the figure above.

3.2.4 Holder's Inequality

Holder's inequality says that for any two measurable real-valued or complex-valued functions f and g and any constants $p, q \geq 1$ such that $\frac{1}{p} + \frac{1}{q} = 1$,

$$||fg||_1 \leq ||f||_p ||g||_q.$$

From this general definition of Holder's inequality, we can say that for any vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ and any constants $p, q \geq 1$ such that $\frac{1}{p} + \frac{1}{q} = 1$,

$$|\langle \boldsymbol{x}, \boldsymbol{y} \rangle_{\mathbb{R}^n}| = \left| \sum_{i=1}^{n} x_i y_i \right| \leq \sum_{i=1}^{n} |x_i y_i| \leq ||\boldsymbol{x}||_q ||\boldsymbol{y}||_p$$

The last inequality comes from Holder's inequality because we can express the third term as $||\boldsymbol{x} \odot \boldsymbol{y}||_1$, where $\boldsymbol{x} \odot \boldsymbol{y}$ is used to denote element-wise multiplication.

3.2.5 Equivalent Norms

Two vector norms $||\cdot||_p$ and $||\cdot||_q$ on (V, F) are said to be equivalent if for all vectors $\boldsymbol{x} \in V$, there exist positive scalar values m_l and m_u such that

$$m_l||\boldsymbol{x}||_p \leq ||\boldsymbol{x}||_q \leq m_u||\boldsymbol{x}||_p.$$

We can show that for the inner product space $(\mathbb{F}^n, \mathbb{F})$, where $\mathbb{F}$ is the field of real numbers $\mathbb{R}$ or complex numbers $\mathbb{C}$, the l_1, l_2, and l_∞ norms are all equivalent because they satisfy the following inequalities:

1. $||\boldsymbol{x}||_\infty \leq ||\boldsymbol{x}||_1 \leq n||\boldsymbol{x}||_\infty$
2. $||\boldsymbol{x}||_\infty \leq ||\boldsymbol{x}||_2 \leq \sqrt{n}||\boldsymbol{x}||_\infty$
3. $\frac{1}{\sqrt{n}}||\boldsymbol{x}||_1 \leq ||\boldsymbol{x}||_2 \leq ||\boldsymbol{x}||_1$

Proof: We will first show that the l_1 and l_∞ norms are equivalent:

$$||\boldsymbol{x}||_1 = \sum_{i=1}^{n} |x_i| \geq |x_i|, \ \forall i \in [1, n] \implies ||\boldsymbol{x}||_1 \geq \max_{i=1,\ldots,n} |x_i| = ||\boldsymbol{x}||_\infty$$

$$||\boldsymbol{x}||_1 = \sum_{i=1}^{n} |x_i| \leq \sum_{i=1}^{n} \max_{i=1,\ldots,n} |x_i| = \sum_{i=1}^{n} ||\boldsymbol{x}||_\infty = n||\boldsymbol{x}||_\infty$$

Now we will show that the l_2 and l_∞ norms are equivalent:

$$||\boldsymbol{x}||_2 = \left(\sum_{i=1}^{n} |x_i|^2\right)^{1/2} \geq |x_i|, \ \forall i \in [1, n] \implies ||\boldsymbol{x}||_2 \geq \max_{i=1,\ldots,n} |x_i| = ||\boldsymbol{x}||_\infty$$

$$||\boldsymbol{x}||_2^2 = \sum_{i=1}^{n} |x_i|^2 \leq \sum_{i=1}^{n} \max_{i=1,\ldots,n} |x_i|^2 = \sum_{i=1}^{n} ||\boldsymbol{x}||_\infty^2 = n||\boldsymbol{x}||_\infty^2 \implies ||\boldsymbol{x}||_2 \leq \sqrt{n}||\boldsymbol{x}||_\infty$$

Finally, we will show that the l_1 and l_2 norms are equivalent:

$$||\boldsymbol{x}||_1^2 = \left(\sum_{i=1}^{n} |x_i|\right)^2 = \sum_{i=1}^{n} |x_i|^2 + 2\sum_{i \neq j} |x_i||x_j| \geq \sum_{i=1}^{n} |x_i|^2 = ||\boldsymbol{x}||_2^2 \implies ||\boldsymbol{x}||_1 \geq ||\boldsymbol{x}||_2$$

Let's define a new vector $\tilde{\boldsymbol{x}} := \text{sign}(\boldsymbol{x})$ whose elements are defined such that $\tilde{x}_i = 1$ if $x_i \geq 0$ and $\tilde{x}_i = -1$ if $x_i < 0$. Using the Cauchy Schwartz inequality,

$$|\langle \boldsymbol{x}, \tilde{\boldsymbol{x}} \rangle| \leq ||\boldsymbol{x}||_2 ||\tilde{\boldsymbol{x}}||_2 \implies \left|\sum_{i=1}^{n} x_i \tilde{x}_i\right| \leq ||\boldsymbol{x}||_2 \Big(\sum_{i=1}^{n} |\tilde{x}_i|^2\Big)^{1/2}$$

$$\left|\sum_{i=1}^{n} x_i \text{sign}(x_i)\right| \leq ||\boldsymbol{x}||_2 \Big(\sum_{i=1}^{n} \big(\text{sign}(x_i)\big)^2\Big)^{1/2}$$

$$\left|\sum_{i=1}^{n} |x_i|\right| \leq ||\boldsymbol{x}||_2 \Big(\sum_{i=1}^{n} 1\Big)^{1/2} \implies \sum_{i=1}^{n} |x_i| \leq ||\boldsymbol{x}||_2 \sqrt{n} \implies ||\boldsymbol{x}||_1 \leq \sqrt{n}||\boldsymbol{x}||_2$$

3.2.6 Dual Norms

Let $||\cdot||$ be a norm defined on the vector space $(\mathbb{R}^n, \mathbb{R})$. The dual of the norm $||\cdot||$ is denoted $||\cdot||^*$ and is defined for a vector $\boldsymbol{x} \in \mathbb{R}^n$ as

$$||\boldsymbol{x}||^* = \sup_{\boldsymbol{z}:||\boldsymbol{z}||\leq 1} \boldsymbol{x}^T\boldsymbol{z}.$$

If the constant p and q satisfy $1 \leq p, q \leq \infty$ and $\frac{1}{p} + \frac{1}{q} = 1$, then

$$||\cdot||_q^* = ||\cdot||_p.$$

Proof: Given $1 \leq p, q \leq \infty$ and $\frac{1}{p} + \frac{1}{q} = 1$, Holder's inequality says that

$$\boldsymbol{x}^T\boldsymbol{z} \leq |\boldsymbol{x}^T\boldsymbol{z}| = |\langle \boldsymbol{x}, \boldsymbol{z}\rangle_{\mathbb{R}^n}| \leq ||\boldsymbol{x}||_p||\boldsymbol{z}||_q.$$

Because a valid norm is necessarily non-negative, under the constraint $||\boldsymbol{z}||_q \leq 1$, the inner product $\boldsymbol{x}^T\boldsymbol{z}$ is maximized when $||\boldsymbol{z}||_q = 1$. This implies that

$$||\boldsymbol{x}||_q^* = \sup_{\boldsymbol{z}:||\boldsymbol{z}||_q\leq 1} \boldsymbol{x}^T\boldsymbol{z} = ||\boldsymbol{x}||_p(1) = ||\boldsymbol{x}||_p.$$

This has important implications for the common l_p norms we discussed:

1. $||\boldsymbol{x}||_2 = ||\boldsymbol{x}||_2^* = \sup_{\boldsymbol{z}:||\boldsymbol{z}||_2\leq 1} \boldsymbol{x}^T\boldsymbol{z}$
2. $||\boldsymbol{x}||_\infty = ||\boldsymbol{x}||_1^* = \sup_{\boldsymbol{z}:||\boldsymbol{z}||_1\leq 1} \boldsymbol{x}^T\boldsymbol{z}$
3. $||\boldsymbol{x}||_1 = ||\boldsymbol{x}||_\infty^* = \sup_{\boldsymbol{z}:||\boldsymbol{z}||_\infty\leq 1} \boldsymbol{x}^T\boldsymbol{z}$

3.2.7 Cardinality

The **cardinality** of a vector is often called the $\boldsymbol{l_0}$ **norm** and is denoted $||\boldsymbol{x}||_0$, but it is defined differently than the other l_p norms. The cardinality of a vector is the number of nonzero elements, which can be expressed as

$$\text{card}(\boldsymbol{x}) = \sum_{i=1}^{n} \tilde{x}_i \quad \text{where} \quad \tilde{x}_i = \begin{cases} 1 & \text{if } x_i \neq 0 \\ 0 & \text{if } x_i = 0 \end{cases}.$$

3.3 Orthogonality

3.3.1 Angle Between Vectors

The angle between two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ is denoted θ and satisfies

$$\cos\theta = \frac{\langle \boldsymbol{x}, \boldsymbol{y}\rangle}{||\boldsymbol{x}||||\boldsymbol{y}||}.$$

The angle between two vectors measures the similarity between those vectors. Values of $|\cos\theta|$ close to one indicate high levels of similarity, while values of $|\cos\theta|$ close to zero indicate low levels of similarity. If $\cos\theta = \pm 1$, then $\boldsymbol{x}$ and $\boldsymbol{y}$ are **parallel** or **colinear**. This also indicates that they are linearly dependent. If $\cos\theta = 0$, then $\boldsymbol{x}$ and $\boldsymbol{y}$ are **perpendicular** or **orthogonal**.

3.3.2 Orthogonal Vectors

Now we will look more at the concept of orthogonality. Two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ in an inner product space are **orthogonal** if $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = 0$. We often use $\boldsymbol{x} \perp \boldsymbol{y}$ to indicate that the vectors x and y are orthogonal to each other. The **Pythagoras theorem** says that if two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ are orthogonal, then

$$||\boldsymbol{x} + \boldsymbol{y}||^2 = ||\boldsymbol{x}||^2 + ||\boldsymbol{y}||^2.$$

Proof: Notice that from the properties of inner products we can write:

$$\begin{aligned} ||\boldsymbol{x} + \boldsymbol{y}||^2 = \langle \boldsymbol{x} + \boldsymbol{y}, \boldsymbol{x} + \boldsymbol{y} \rangle &= \langle \boldsymbol{x}, \boldsymbol{x} \rangle + \langle \boldsymbol{x}, \boldsymbol{y} \rangle + \langle \boldsymbol{y}, \boldsymbol{x} \rangle + \langle \boldsymbol{y}, \boldsymbol{y} \rangle \\ &= ||\boldsymbol{x}||^2 + ||\boldsymbol{y}||^2 + \langle \boldsymbol{x}, \boldsymbol{y} \rangle + \langle \boldsymbol{y}, \boldsymbol{x} \rangle \end{aligned}$$

If $\boldsymbol{x}$ and $\boldsymbol{y}$ are orthogonal, then $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \langle \boldsymbol{y}, \boldsymbol{x} \rangle = 0$, which implies that

$$||\boldsymbol{x} + \boldsymbol{y}||^2 = ||\boldsymbol{x}||^2 + ||\boldsymbol{y}||^2.$$

A set of vectors $S = \{\boldsymbol{x_1}, \ldots, \boldsymbol{x}_d\}$ is considered **mutually orthogonal** if $\langle \boldsymbol{x_i}, \boldsymbol{x_j} \rangle = 0$ for all $i \neq j$. Mutually orthogonal vectors are linearly independent because the angle between each pair of vectors is non-zero. This set of vectors is called **orthonormal** if it is mutually orthogonal and each vector has norm one. Therefore, the set S is orthonormal if for all i and j,

$$\langle \boldsymbol{x_i}, \boldsymbol{x_j} \rangle = \begin{cases} 0 & \text{if } i \neq j \\ 1 & \text{if } i = j \end{cases}.$$

3.3.3 Orthogonal Complement

Let's assume (V, F) is an inner product space, $\boldsymbol{x}$ is a vector in V, and the set S is a subset of the set of vectors in the inner product space. The vector $\boldsymbol{x}$ is orthogonal to the subset S if $\boldsymbol{x} \perp \boldsymbol{y}$ for all $\boldsymbol{y}$ in S. Note that if $\boldsymbol{x}$ is orthogonal to the set $\mathbb{R}^n$ or $\mathbb{C}^n$, then $\boldsymbol{x}$ must be the n-dimensional zero vector, $\boldsymbol{0}_n$. The set of vectors in V that are orthogonal to S make up the **orthogonal complement** of S, which we denote as $S^\perp$. We can express this set mathematically as

$$S^\perp = \{\boldsymbol{x} \in V : \langle \boldsymbol{x}, \boldsymbol{y} \rangle = 0, \ \forall \boldsymbol{y} \in S\}.$$

From this definition of the orthogonal complement, we can see that

$$\boldsymbol{x} \perp \boldsymbol{y}, \ \forall \boldsymbol{x} \in S, \ \boldsymbol{y} \in S^\perp.$$

For any subset $S \subseteq V$, we know $S \cap S^\perp = \emptyset$ and $V = S \oplus S^\perp$. This implies that the dimension of the set of vectors V is given by $\dim V = \dim S + \dim S^\perp$.

3.3.4 Gram Schmidt Orthonormalization

Given a collection of basis vectors $B = \{\boldsymbol{b_1}, \ldots, \boldsymbol{b_m}\}$ for an inner product space with the set of vectors $V = \text{span}(B)$, we can construct an **orthonormal basis**. To find the orthonomal basis $\{\boldsymbol{c_1}, \ldots, \boldsymbol{c_m}\}$, we will go through an iterative process called the **Gram-Schmidt orthonormalization** procedure:

1. Initialize $k = 2$ and $\boldsymbol{y_1} = \boldsymbol{b_1}$. Set $\boldsymbol{c_1} = \frac{\boldsymbol{y_1}}{||\boldsymbol{y_1}||}$.
2. Define $\boldsymbol{y_k} = \boldsymbol{b_k} - \sum_{i=1}^{k-1} \langle \boldsymbol{c_i}, \boldsymbol{b_k} \rangle \boldsymbol{c_i}$
3. If $\boldsymbol{y_k} \neq \boldsymbol{0}$, then set $\boldsymbol{c_k} = \frac{\boldsymbol{y_k}}{||\boldsymbol{y_k}||}$. Otherwise, let $\boldsymbol{c_k} = \boldsymbol{0}$.
4. If the set B has not been exhausted, increment k and move to step 2.

Chapter 4

Projections

4.1 Projection onto a Subspace

The goal of **projection** is to find a point in a given set that is closest in norm to a given point that is not in the desired set. Given a vector $\boldsymbol{x}$ in an inner product space $\mathcal{X}$ and a closed set $S \subseteq \mathcal{X}$, the projection of $\boldsymbol{x}$ onto S is

$$\Pi_S(\boldsymbol{x}) = \arg\min_{\boldsymbol{y} \in S} ||\boldsymbol{y} - \boldsymbol{x}||.$$

I will generally assume that $\mathcal{X}$ is the Euclidean space (or some subset of the Euclidean space) and the norm defined in the projection equation is the l_2 norm.

4.1.1 Projection Theorem

For a vector $\boldsymbol{x}$ in a Euclidean set $\mathcal{X}$ and a closed set $S \subseteq \mathcal{X}$, the **projection theorem** says that there always exists a unique solution to the following:

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{y} \in S} ||\boldsymbol{y} - \boldsymbol{x}||_2 = \arg\min_{\boldsymbol{y} \in S} ||\boldsymbol{y} - \boldsymbol{x}||_2^2.$$

Moreover, a necessary and sufficient condition for $\hat{\boldsymbol{x}}$ being optimal is

1. $\hat{\boldsymbol{x}} \in S$, and
2. $(\boldsymbol{x} - \hat{\boldsymbol{x}}) \perp S$.

Proof: To prove the projection theorem, recall that if $S^\perp$ is the orthogonal complement of S, then any $\boldsymbol{x} \in \mathcal{X}$ can be expressed as $\boldsymbol{x} = \boldsymbol{u} + \boldsymbol{z}$, where $\boldsymbol{u} \in S$ and $\boldsymbol{z} \in S^\perp$. We can then express the objective in the projection problem as

$$||\boldsymbol{y} - \boldsymbol{x}||_2^2 = ||\boldsymbol{y} - (\boldsymbol{u} + \boldsymbol{z})||_2^2 = ||(\boldsymbol{y} - \boldsymbol{u}) - \boldsymbol{z}||_2^2 = ||\boldsymbol{y} - \boldsymbol{u}||_2^2 + ||\boldsymbol{z}||_2^2 - 2\langle \boldsymbol{y} - \boldsymbol{u}, \boldsymbol{z} \rangle.$$

Because we assume that $\boldsymbol{y} \in S$ and $\boldsymbol{u} \in S$ are both elements of a vector subspace S, we know that $(\boldsymbol{y} - \boldsymbol{u}) \in S$. We also assumed that $\boldsymbol{z} \in S^{\perp}$, so we can write $\boldsymbol{z} \perp (\boldsymbol{y} - \boldsymbol{u})$, which implies that $\langle \boldsymbol{y} - \boldsymbol{u}, \boldsymbol{z} \rangle = 0$. Now the objective becomes

$$||\boldsymbol{y} - \boldsymbol{x}||_2^2 = ||\boldsymbol{y} - \boldsymbol{u}||_2^2 + ||\boldsymbol{z}||_2^2.$$

With this new objective function, the projection problem becomes

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{y} \in S} ||\boldsymbol{y} - \boldsymbol{u}||_2^2 + ||\boldsymbol{z}||_2^2.$$

The term $||\boldsymbol{z}||_2^2$ is a constant that does not depend on $\boldsymbol{y}$, so

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{y} \in S} ||\boldsymbol{y} - \boldsymbol{u}||_2^2.$$

Now we can see that the optimal solution for this problem is simply $\hat{\boldsymbol{x}} = \boldsymbol{u}$. Recall that we defined $\boldsymbol{x}$ such $\boldsymbol{x} = \boldsymbol{u} + \boldsymbol{z}$, where $\boldsymbol{u} \in S$ and $\boldsymbol{z} \in S^{\perp}$. This allows us to write $\boldsymbol{x} - \hat{\boldsymbol{x}} = \boldsymbol{z}$, so $(\boldsymbol{x} - \hat{\boldsymbol{x}}) \in S^{\perp}$ or equivalently $(\boldsymbol{x} - \hat{\boldsymbol{x}}) \perp S$.

4.2 Common Projections

There are some common projection problems whose solutions are good to know. Among these are projections onto a one-dimensional line, onto a two-dimensional hyperplane, and onto a subspace defined by a span of vectors.

4.2.1 Projection onto a Line

Suppose we are interested in projecting a point $\boldsymbol{x} \in \mathbb{R}^n$ onto a line

$$\mathcal{L} = \{\boldsymbol{y} \in \mathbb{R}^n : \boldsymbol{y} = \boldsymbol{x}_0 + \lambda \boldsymbol{v}, \lambda \in \mathbb{R}\},$$

which passes through the point $\boldsymbol{x}_0 \in \mathbb{R}^n$ in the direction $\boldsymbol{v} \in \mathbb{R}^n$. This problem can be expressed as the following optimization problem:

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{y} \in \mathcal{L}} ||\boldsymbol{y} - \boldsymbol{x}||_2.$$

Using the projection theorem, we know $\hat{\boldsymbol{x}} \in \mathcal{L}$ and $(\boldsymbol{x} - \hat{\boldsymbol{x}}) \perp \mathcal{L}$. Given that $\hat{\boldsymbol{x}} \in \mathcal{L}$, there is some constant $\hat{\lambda} \in \mathbb{R}$ such that $\hat{\boldsymbol{x}} = \boldsymbol{x}_0 + \hat{\lambda}\boldsymbol{v}$. Because $(\boldsymbol{x} - \hat{\boldsymbol{x}}) \perp \mathcal{L}$, we can also write $\boldsymbol{v}^T(\boldsymbol{x} - \hat{\boldsymbol{x}}) = 0$. Combining these facts, we can notice that

$$0 = \boldsymbol{v}^T(\boldsymbol{x} - \hat{\boldsymbol{x}}) = \boldsymbol{v}^T\big(\boldsymbol{x} - (\boldsymbol{x}_0 + \hat{\lambda}\boldsymbol{v})\big) = \boldsymbol{v}^T\boldsymbol{x} - \boldsymbol{v}^T\boldsymbol{x}_0 - \hat{\lambda}\boldsymbol{v}^T\boldsymbol{v} \quad \Longrightarrow \quad \hat{\lambda} = \frac{\boldsymbol{v}^T(\boldsymbol{x} - \boldsymbol{x}_0)}{\boldsymbol{v}^T\boldsymbol{v}}$$

$$\hat{\boldsymbol{x}} = \boldsymbol{x}_0 + \hat{\lambda}\boldsymbol{v} = \boldsymbol{x}_0 + \frac{\boldsymbol{v}^T(\boldsymbol{x} - \boldsymbol{x}_0)}{||\boldsymbol{v}||_2^2}\boldsymbol{v}$$

4.2.2 Projection onto a Hyperplane

Suppose we are interested in projecting a point $\boldsymbol{x} \in \mathbb{R}^n$ onto a hyperplane

$$\mathcal{H} = \{\boldsymbol{y} \in \mathbb{R}^n : \boldsymbol{a}^T\boldsymbol{y} = b\},$$

where $\boldsymbol{a} \in \mathbb{R}^n$ is the normal vector and $b \in \mathbb{R}$ is the offset. This problem can be expressed as the following optimization problem:

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{y} \in \mathcal{H}} ||\boldsymbol{y} - \boldsymbol{x}||_2$$

Using the projection theorem, we know $\hat{\boldsymbol{x}} \in \mathcal{H}$ and $(\boldsymbol{x} - \hat{\boldsymbol{x}}) \perp \mathcal{H}$. Given that $\hat{\boldsymbol{x}} \in \mathcal{H}$, the vector must satisfy $\boldsymbol{a}^T\hat{\boldsymbol{x}} = b$. Because $(\boldsymbol{x} - \hat{\boldsymbol{x}}) \perp \mathcal{H}$ and the vector $\boldsymbol{a}$ is normal to the hyperplane $\mathcal{H}$, there is some constant $\alpha \in \mathbb{R}$ such that $(\boldsymbol{x} - \hat{\boldsymbol{x}}) = \alpha\boldsymbol{a}$. Combining these facts, we can determine the following:

$$\boldsymbol{a}^T(\boldsymbol{x} - \hat{\boldsymbol{x}}) = \boldsymbol{a}^T(\alpha\boldsymbol{a})$$

$$\boldsymbol{a}^T\boldsymbol{x} - b = \alpha\boldsymbol{a}^T\boldsymbol{a}$$

$$\alpha = \frac{\boldsymbol{a}^T\boldsymbol{x} - b}{\boldsymbol{a}^T\boldsymbol{a}} = \frac{\boldsymbol{a}^T\boldsymbol{x} - b}{||\boldsymbol{a}||_2^2}$$

$$\hat{\boldsymbol{x}} = \boldsymbol{x} - \alpha\boldsymbol{a} = \boldsymbol{x} - \frac{\boldsymbol{a}^T\boldsymbol{x} - b}{||\boldsymbol{a}||_2^2}\boldsymbol{a}$$

We can also notice that the distance from the point $\boldsymbol{x}$ to the hyperplane $\mathcal{H}$ is

$$d = ||\boldsymbol{x} - \hat{\boldsymbol{x}}||_2 = ||\alpha\boldsymbol{a}||_2 = |\alpha|||\boldsymbol{a}||_2 = \frac{|\boldsymbol{a}^T\boldsymbol{x} - b|}{||\boldsymbol{a}||_2^2}||\boldsymbol{a}||_2 = \frac{|\boldsymbol{a}^T\boldsymbol{x} - b|}{||\boldsymbol{a}||_2}.$$

4.2.3 Projection onto a Vector Span

Suppose we are interested in projecting a point $\boldsymbol{x} \in \mathcal{X}$ onto a subspace $S \subseteq \mathcal{X}$, which is defined as the span of a set of arbitrary vectors:

$$S = \text{span}(\boldsymbol{v_1}, \ldots, \boldsymbol{v_d}).$$

Using the projection theorem, we know $\hat{\boldsymbol{x}} \in S$ and $(\boldsymbol{x} - \hat{\boldsymbol{x}}) \perp S$. We can equivalently write these two conditions as the following:

$$\hat{\boldsymbol{x}} = \sum_{i=1}^{d} \alpha_i \boldsymbol{v_i} \quad \text{for some} \quad \alpha_i \in \mathbb{R},\ i = 1, \ldots, d, \text{ and}$$

$$\langle \boldsymbol{x} - \hat{\boldsymbol{x}}, \boldsymbol{v_i} \rangle = 0 \quad \text{for} \quad i = 1, \ldots, d.$$

From the second condition, we can notice that

$$\boldsymbol{v_k}^T(\boldsymbol{x} - \hat{\boldsymbol{x}}) = 0, \text{ which implies } \boldsymbol{v_k}^T\boldsymbol{x} = \boldsymbol{v_k}^T\hat{\boldsymbol{x}} \quad \text{for} \quad k = 1, \ldots, d.$$

Combining this with our first condition, we can see that

$$\boldsymbol{v_k}^T\boldsymbol{x} = \boldsymbol{v_k}^T\left(\sum_{i=1}^{d}\alpha_i\boldsymbol{v_i}\right) = \sum_{i=1}^{d}\alpha_i\boldsymbol{v_k}^T\boldsymbol{v_i} \quad \text{for} \quad k = 1,\ldots,d.$$

Now we have a system of d equations and d unknowns: $\alpha_1,\ldots,\alpha_d$. From this system of equations, we can solve for $\alpha_1,\ldots,\alpha_d$, and then express $\hat{\boldsymbol{x}}$ as

$$\hat{\boldsymbol{x}} = \sum_{i=1}^{d}\alpha_i\boldsymbol{v_i}.$$

If we assume $\boldsymbol{v_1},\ldots,\boldsymbol{v_d}$ are orthonormal, we can find a more precise solution. Recall that if $\boldsymbol{v_1},\ldots,\boldsymbol{v_d}$ are orthonormal, their inner products are given by

$$\boldsymbol{v_k}^T\boldsymbol{v_i} = \begin{cases} 1 & \text{if } i = k \\ 0 & \text{otherwise} \end{cases}.$$

Now the system of equations we wrote for an arbitrary set of vectors becomes

$$\boldsymbol{v_k}^T\boldsymbol{x} = \boldsymbol{v_k}^T\left(\sum_{i=1}^{d}\alpha_i\boldsymbol{v_i}\right) = \sum_{i=1}^{d}\alpha_i\boldsymbol{v_k}^T\boldsymbol{v_i} = \alpha_k \quad \text{for} \quad k = 1,\ldots,d.$$

Now we can easily determine the value of each of the constants $\alpha_1,\ldots,\alpha_d$, which allows us to express the optimal solution, $\hat{\boldsymbol{x}}$, as

$$\hat{\boldsymbol{x}} = \sum_{i=1}^{d}\alpha_i\boldsymbol{v_i} = \sum_{i=1}^{d}(\boldsymbol{v_i}^T\boldsymbol{x})\boldsymbol{v_i}.$$

Part II

Matrices

Chapter 5

Matrix Basics

5.1 Overview of Matrices

An $m \times n$ matrix $\boldsymbol{A}$ is composed of m n-dimensional row vectors and n m-dimensional column vectors. I will denote the ith row $\boldsymbol{a_i^T}$, the jth column $\boldsymbol{a_j}$, and the ijth element a_{ij}. In general, an $m \times n$ matrix $\boldsymbol{A}$ has the form

$$\boldsymbol{A} = \begin{bmatrix} — & \boldsymbol{a_1^T} & — \\ & \vdots & \\ — & \boldsymbol{a_m^T} & — \end{bmatrix} = \begin{bmatrix} | & & | \\ \boldsymbol{a_1} & \dots & \boldsymbol{a_n} \\ | & & | \end{bmatrix} = \begin{bmatrix} a_{11} & \dots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{m1} & \dots & a_{mn} \end{bmatrix} \in \mathbb{F}^{m \times n},$$

where $\mathbb{F}$ is the field of real numbers $\mathbb{R}$ or complex numbers $\mathbb{C}$.

5.1.1 Matrix Product

If $\boldsymbol{A} \in \mathbb{F}^{m \times n}$ and $\boldsymbol{B} \in \mathbb{F}^{n \times p}$, we can multiply these two matrices together, resulting in an $m \times p$ matrix. We multiply the matrices in the following way:

$$\boldsymbol{AB} = \begin{bmatrix} — & \boldsymbol{a_1^T} & — \\ & \vdots & \\ — & \boldsymbol{a_m^T} & — \end{bmatrix} \begin{bmatrix} | & & | \\ \boldsymbol{b_1} & \dots & \boldsymbol{b_n} \\ | & & | \end{bmatrix} = \begin{bmatrix} \boldsymbol{a_1^T b_1} & \dots & \boldsymbol{a_1^T b_n} \\ \vdots & \ddots & \vdots \\ \boldsymbol{a_m^T b_1} & \dots & \boldsymbol{a_m^T b_n} \end{bmatrix}.$$

Each element of the resulting matrix $\boldsymbol{AB}$ can be expressed as

$$(\boldsymbol{AB})_{ij} = \sum_{k=1}^{n} a_{ik} b_{kj}.$$

5.1.2 Matrix Transpose

The transpose of a matrix $\boldsymbol{A}$ shown previously can be expressed

$$\boldsymbol{A}^T = \begin{bmatrix} | & & | \\ \boldsymbol{a_1} & \dots & \boldsymbol{a_m} \\ | & & | \end{bmatrix} = \begin{bmatrix} — & \boldsymbol{a_1^T} & — \\ & \vdots & \\ — & \boldsymbol{a_n^T} & — \end{bmatrix} = \begin{bmatrix} a_{11} & \dots & a_{m1} \\ \vdots & \ddots & \vdots \\ a_{1n} & \dots & a_{mn} \end{bmatrix} \in \mathbb{F}^{n \times m}.$$

The matrix transpose has the following properties:

1. $(\boldsymbol{AB})^T = \boldsymbol{B}^T\boldsymbol{A}^T$
2. $(\boldsymbol{A_1 A_2 \dots A_p})^T = \boldsymbol{A_p^T \dots A_2^T A_1^T}$
3. $\begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{C} & \boldsymbol{D} \end{bmatrix}^T = \begin{bmatrix} \boldsymbol{A}^T & \boldsymbol{C}^T \\ \boldsymbol{B}^T & \boldsymbol{D}^T \end{bmatrix}$

The complex conjugate transpose, or Hermitian transpose, of the matrix $\boldsymbol{A}$ is

$$\boldsymbol{A}^* = \begin{bmatrix} — & \boldsymbol{a_1^*} & — \\ & \vdots & \\ — & \boldsymbol{a_n^*} & — \end{bmatrix} = \begin{bmatrix} \bar{a}_{11} & \dots & \bar{a}_{m1} \\ \vdots & \ddots & \vdots \\ \bar{a}_{1n} & \dots & \bar{a}_{mn} \end{bmatrix} \in \mathbb{F}^{n\times m}.$$

5.1.3 Range and Null Space

Consider an $m \times n$ matrix $\boldsymbol{A} \in \mathbb{F}^{m\times n}$. The **range space**, which is also known as the **column space**, **image**, or **span**, of the matrix $\boldsymbol{A}$ is defined as

$$R(\boldsymbol{A}) = \{\boldsymbol{y} \in \mathbb{F}^m : \boldsymbol{y} = \boldsymbol{Ax},\ \boldsymbol{x} \in \mathbb{F}^n\} \subseteq \mathbb{F}^m.$$

The **row space** of the matrix $\boldsymbol{A}$ is defined as

$$\text{Row}(\boldsymbol{A}) = \{\boldsymbol{z} \in \mathbb{F}^n : \boldsymbol{z} = \boldsymbol{A}^T\boldsymbol{x},\ \boldsymbol{x} \in \mathbb{F}^m\} \subseteq \mathbb{F}^n.$$

The **null space**, or **kernel**, of the matrix $\boldsymbol{A}$ is defined as

$$N(\boldsymbol{A}) = \{\boldsymbol{x} \in \mathbb{F}^n : \boldsymbol{Ax} = \boldsymbol{0}_m\} \subseteq \mathbb{F}^n.$$

Note that the range space of $\boldsymbol{A}^T$ is the same as the row space of $\boldsymbol{A}$. Additionally, the orthogonal complement of the null space of $\boldsymbol{A}$ is the row space of $\boldsymbol{A}$.

The **rank** of matrix $\boldsymbol{A}$ is defined as the dimension of the range space, and the **nullity** of $\boldsymbol{A}$ is defined as the dimension of the null space. Because the range space is a subset of $\mathbb{F}^m$, the rank is less than or equal m. Similarly, because the null space is a subset of $\mathbb{F}^n$, the nullity is less than or equal to n.

5.1.4 Elementary Operations

There are three elementary row operations performed on matrices:

1. interchange two rows of the matrix,
2. multiply a row by a non-zero constant, and
3. add one row to another row.

Elementary row operations are equivalent to pre-multiplying a matrix $\boldsymbol{A}$ by a left elementary matrix $\boldsymbol{L}$, which is obtained by performing the desired operations on the identity matrix. Note that the null space of the transformed matrix is the same as the null space of the original matrix (i.e. $N(\boldsymbol{LA}) = N(\boldsymbol{A})$).

Similarly, there are three elementary column operations performed on matrices:

1. interchange two columns of the matrix,
2. multiply a column by a non-zero constant, and
3. add one column to another column.

Elementary column operations are equivalent to post-multiplying a matrix $\boldsymbol{A}$ by a right elementary matrix $\boldsymbol{R}$, which is obtained by performing the desired operations on the identity matrix. The range space of the transformed matrix is the same as the range space of the original matrix (i.e. $R(\boldsymbol{AR}) = R(\boldsymbol{A})$).

5.1.5 Hadamard Product

Given two $m \times n$ matrices $\boldsymbol{A}$ and $\boldsymbol{B}$, their **Hadamard product** is the $m \times n$ matrix resulting from the element-wise product of the entries of the matrices and is denoted $\boldsymbol{A} \circ \boldsymbol{B}$. From this definition, the ijth element of $\boldsymbol{A} \circ \boldsymbol{B}$ is

$$(\boldsymbol{A} \circ \boldsymbol{B})_{ij} = a_{ij} b_{ij}.$$

5.2 Special Matrices

There are a number of different types of matrices with unique properties. Some of the most useful types of special matrices are listed below.

1. **Sparse** – A matrix is sparse if most of the elements in the matrix are 0.
2. **Square** – A matrix is square if it has the same number of rows as columns.
3. **Symmetric** – A square matrix $\boldsymbol{A}$ is symmetric if it is equal to its transpose (i.e. $\boldsymbol{A} = \boldsymbol{A}^T$), meaning $a_{ij} = a_{ji}$ for all i and j.
4. **Skew Symmetric** – A square matrix $\boldsymbol{A}$ is skew symmetric if it is equal to the negative of its transpose (i.e. $\boldsymbol{A} = -\boldsymbol{A}^T$), meaning $a_{ij} = -a_{ji}$ for all i and j. Note that all the diagonal entries, a_{ii}, must be zero.
5. **Hermitian** – A square matrix $\boldsymbol{A}$ is Hermitian if it is equal to its complex conjugate transpose (i.e. $\boldsymbol{A} = \boldsymbol{A}^*$), meaning $a_{ij} = \bar{a}_{ji}$ for all i and j.
6. **Diagonal** – A square matrix $\boldsymbol{A}$ is diagonal if all of its nonzero elements are along its diagonal, meaning $a_{ij} = 0$ if $i \neq j$. For example,

$$\boldsymbol{A} = \text{diag}(\boldsymbol{a}) = \text{diag}(a_1, \ldots, a_n) = \begin{bmatrix} a_1 & & \\ & \ddots & \\ & & a_n \end{bmatrix}.$$

The most popular diagonal matrix is the identity matrix whose diagonal elements are all one. I denote the $n \times n$ identity matrix as $\boldsymbol{I_n}$.

7. **Upper Triangular** – A square matrix $\boldsymbol{A}$ is upper triangular if all of its nonzero elements fall on or above the main diagonal, meaning $a_{ij} = 0$ if $i > j$. We can express a general upper triangular matrix as

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & \dots & a_{1n} \\ & \ddots & \vdots \\ & & a_{nn} \end{bmatrix}.$$

8. **Lower Triangular** – A square matrix $\boldsymbol{A}$ is lower triangular if all of its nonzero elements fall on or below the main diagonal, meaning $a_{ij} = 0$ if $i < j$. We can express a general lower triangular matrix as

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & & \\ \vdots & \ddots & \\ a_{n1} & \dots & a_{nn} \end{bmatrix}.$$

9. **Unitary** – A matrix $\boldsymbol{U}$ is unitary if its columns form an orthonormal basis for $\mathbb{C}^n$. A unitary matrix generally has the form

$$\boldsymbol{U} = \begin{bmatrix} \boldsymbol{u_1} & \dots & \boldsymbol{u_n} \end{bmatrix}, \quad \text{where} \quad \boldsymbol{u_i^* u_j} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}.$$

Unitary matrices must satisfy the property $\boldsymbol{U^*U} = \boldsymbol{UU^*} = \boldsymbol{I_n}$. From this property, we can also notice that multiplying a vector by a unitary matrix preserves the length of the vector. To see this, notice that

$$||\boldsymbol{Ux}||_2^2 = (\boldsymbol{Ux})^*(\boldsymbol{Ux}) = \boldsymbol{x^*U^*Ux} = \boldsymbol{x^*I_n x} = \boldsymbol{x^*x} = ||\boldsymbol{x}||_2^2.$$

10. **Orthogonal** – A matrix $\boldsymbol{U}$ is orthgonal if its columns form an orthonormal basis for $\mathbb{R}^n$. An orthogonal matrix generally has the form

$$\boldsymbol{U} = \begin{bmatrix} \boldsymbol{u_1} & \dots & \boldsymbol{u_n} \end{bmatrix} \quad \text{where} \quad \boldsymbol{u_i^T u_j} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}.$$

Orthogonal matrices must satisfy the property $\boldsymbol{U^TU} = \boldsymbol{UU^T} = \boldsymbol{I_n}$. From this property, we can also notice that multiplying a vector by an orthogonal matrix preserves the length of the vector. To see this, notice that

$$||\boldsymbol{Ux}||_2^2 = (\boldsymbol{Ux})^T(\boldsymbol{Ux}) = \boldsymbol{x^TU^TUx} = \boldsymbol{x^TI_n x} = \boldsymbol{x^Tx} = ||\boldsymbol{x}||_2^2.$$

11. **Dyad** – A matrix is a dyad if it is the outer product of two vectors. An $m \times n$ matrix $\boldsymbol{A}$ is a dyad if it can be expressed as $A = uv^T$, where $u \in \mathbb{F}^m$ and $v \in \mathbb{F}^n$. The rank of a dyad is always equal to one and any rank one matrix can be expressed as the outer product of two nonzero vectors.

12. **Square Dyad** – A square dyad is a square matrix that can be expressed as the outer product of two vectors. An $n \times n$ matrix $\boldsymbol{A}$ is a square dyad if it can be expressed as $A = uv^T$, where $u \in \mathbb{F}^n$ and $v \in \mathbb{F}^n$.

13. **Nilpotent** – A square matrix $\boldsymbol{A}$ is nilpotent if $A^k = 0$ for some positive integer k. The smallest such k is called the **index**, or **degree**, of $\boldsymbol{A}$.

5.2.1 Block Matrices

Block matrices can be broken into submatrices or blocks. There are a number of different types of block matrices with unique properties. Some of the most useful types of block matrices are listed below.

1. **Block Diagonal** – A square matrix is block diagonal if its diagonal blocks are square matrices of any size and the off-diagonal blocks are all zero matrices. A block diagonal matrix generally has the form

$$A = \begin{bmatrix} \boldsymbol{A_1} & & \\ & \ddots & \\ & & \boldsymbol{A_n} \end{bmatrix}.$$

2. **Block Upper Triangular** – A matrix is block upper triangular if all of its nonzero blocks fall on or above its diagonal. A block upper triangular matrix generally has the form

$$A = \begin{bmatrix} \boldsymbol{A_{11}} & \dots & \boldsymbol{A_{1n}} \\ & \ddots & \vdots \\ & & \boldsymbol{A_{nn}} \end{bmatrix}.$$

3. **Block Lower Triangular** – A matrix is block lower triangular if all of its nonzero blocks fall on or below its diagonal. A block lower triangular matrix generally has the form

$$A = \begin{bmatrix} \boldsymbol{A_{11}} & & \\ \vdots & \ddots & \\ \boldsymbol{A_{n1}} & \dots & \boldsymbol{A_{nn}} \end{bmatrix}.$$

Chapter 6

Square Matrices

6.1 Trace

The **trace** is a property that is only defined for square matrices. The trace of an $n \times n$ matrix $\boldsymbol{A}$ is the sum of its diagonal elements, which is expressed as

$$\text{tr}(\boldsymbol{A}) = \sum_{i=1}^{n} a_{ii}.$$

For two $n \times n$ matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ and a scalar α, the trace satisfies the following:

1. $\text{tr}(\boldsymbol{A}) = \text{tr}(\boldsymbol{A}^T)$
2. $\text{tr}(\alpha\boldsymbol{A}) = \alpha\text{tr}(\boldsymbol{A})$
3. $\text{tr}(\boldsymbol{AB}) = \text{tr}(\boldsymbol{BA})$
4. $\text{tr}(\boldsymbol{A} + \boldsymbol{B}) = \text{tr}(\boldsymbol{A}) + \text{tr}(\boldsymbol{B})$

Another important property of the trace is the **cyclic property**, which says that the trace is invariant under cyclic perumuations. For example, for the matrices $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, $\boldsymbol{B} \in \mathbb{R}^{n\times p}$, and $\boldsymbol{C} \in \mathbb{R}^{p\times m}$, the cyclic property says

$$\text{tr}(\boldsymbol{ABC}) = \text{tr}(\boldsymbol{BCA}) = \text{tr}(\boldsymbol{CAB}).$$

Note that arbitrary permutations are generally not allowed. For example, for the matrices $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, $\boldsymbol{B} \in \mathbb{R}^{n\times p}$, and $\boldsymbol{C} \in \mathbb{R}^{p\times m}$, we can notice that

$$\text{tr}(\boldsymbol{ABC}) \neq \text{tr}(\boldsymbol{ACB}).$$

However, if $\boldsymbol{A}$, $\boldsymbol{B}$, and $\boldsymbol{C}$ are all symmetric, then we can see that

$$\text{tr}(\boldsymbol{ABC}) = \text{tr}\Big((\boldsymbol{ABC})^T\Big) = \text{tr}(\boldsymbol{C}^T\boldsymbol{B}^T\boldsymbol{A}^T) = \text{tr}(\boldsymbol{CBA}) = \text{tr}(\boldsymbol{ACB}).$$

6.2 Determinant

Another useful property that is only defined for square matrices is the **determinant**. If we define $\boldsymbol{A_{(i,j)}}$ as the $(n-1)\times(n-1)$ submatrix of $\boldsymbol{A}$ obtained by eliminating row i and column j from $\boldsymbol{A}$, then the determinant of $\boldsymbol{A}$ is

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

For two $n\times n$ matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ and a scalar α, the determinant satisfies:

1. $\det(\boldsymbol{I_n}) = 1$
2. $\det(\boldsymbol{A}) = \det(\boldsymbol{A}^T)$
3. $\det(\boldsymbol{A}^{-1}) = \frac{1}{\det(\boldsymbol{A})}$
4. $\det(\alpha\boldsymbol{A}) = \alpha^n\det(\boldsymbol{A})$
5. $\det(\boldsymbol{AB}) = \det(\boldsymbol{BA}) = \det(\boldsymbol{A})\det(\boldsymbol{B})$

In general, the determinant of a matrix can be difficult to compute, but it is straightforward for 2×2 and 3×3 matrices. If $\boldsymbol{A}$ is a 2×2 matrix given by

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix},$$

then its determinant is simply $\det(\boldsymbol{A}) = a_{11}a_{22} - a_{12}a_{21}$. If $\boldsymbol{A}$ is the 3×3 matrix

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix},$$

then its determinant can be expressed as

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

6.3 Adjugate

Recall that we defined $\boldsymbol{A_{(i,j)}}$ as the $(n-1)\times(n-1)$ submatrix of $\boldsymbol{A}$ obtained by eliminating row i and column j from $\boldsymbol{A}$. The **adjugate** of a square matrix $n\times n$ matrix $\boldsymbol{A}$ is defined component-wise as

$$\big(\mathrm{adj}(\boldsymbol{A})\big)_{ij} = (-1)^{i+j}\det(\boldsymbol{A}_{(i,j)}).$$

6.4 Matrix Inverse

A square matrix $\boldsymbol{A} \in \mathbb{F}^n$ is said to be **invertible**, or **non-singular**, if its determinant is nonzero. This is equivalent to saying that the columns of $\boldsymbol{A}$ form a basis for $\mathbb{F}^n$ or, equivalently, the rows of $\boldsymbol{A}$ form a basis for $\mathbb{F}^n$. The **inverse** of an invertible $n \times n$ matrix $\boldsymbol{A}$ is denoted $\boldsymbol{A}^{-1}$ and is defined such that

$$\boldsymbol{A}\boldsymbol{A}^{-1} = \boldsymbol{A}^{-1}\boldsymbol{A} = \boldsymbol{I_n}.$$

The inverse of $\boldsymbol{A}$ can be expressed in terms of its adjugate and determinant as

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

For two $n \times n$ matrices $\boldsymbol{A}$ and $\boldsymbol{B}$, the matrix inverse satisfies the following:

1. $(\boldsymbol{AB})^{-1} = \boldsymbol{B}^{-1}\boldsymbol{A}^{-1}$
2. $(\boldsymbol{A}^T)^{-1} = (\boldsymbol{A}^{-1})^T$

6.4.1 Inverse of Block Matrices

1. **Block Diagonal** – Consider a block diagonal matrix

$$\boldsymbol{A} = \begin{bmatrix} \boldsymbol{A_1} & & \\ & \ddots & \\ & & \boldsymbol{A_n} \end{bmatrix}.$$

   $\boldsymbol{A}$ is invertible if and only if each of its diagonal blocks are invertible. When all of its blocks are invertible, the inverse of $\boldsymbol{A}$ inverse is

$$\boldsymbol{A}^{-1} = \begin{bmatrix} \boldsymbol{A_1^{-1}} & & \\ & \ddots & \\ & & \boldsymbol{A_n^{-1}} \end{bmatrix}.$$

2. **Block Upper Triangular** – Consider a block upper triangular matrix

$$\boldsymbol{A} = \begin{bmatrix} \boldsymbol{A_{11}} & \boldsymbol{A_{12}} \\ 0 & \boldsymbol{A_{22}} \end{bmatrix}.$$

   $\boldsymbol{A}$ is invertible if and only if its diagonal blocks $\boldsymbol{A_{11}}$ and $\boldsymbol{A_{22}}$ are invertible. If both of its diagaonal blocks are invertible, the the inverse of $\boldsymbol{A}$ is

$$\boldsymbol{A}^{-1} = \begin{bmatrix} \boldsymbol{A_{11}^{-1}} & -\boldsymbol{A_{11}^{-1}A_{12}A_{22}^{-1}} \\ 0 & \boldsymbol{A_{22}^{-1}} \end{bmatrix}.$$

3. **Block Lower Triangular** – Consider a block lower triangular matrix

$$\boldsymbol{A} = \begin{bmatrix} \boldsymbol{A_{11}} & 0 \\ \boldsymbol{A_{21}} & \boldsymbol{A_{22}} \end{bmatrix}.$$

$\boldsymbol{A}$ is invertible if and only if its diagonal blocks $\boldsymbol{A_{11}}$ and $\boldsymbol{A_{22}}$ are invertible. If both of its diagaonal blocks are invertible, the the inverse of $\boldsymbol{A}$ is

$$\boldsymbol{A}^{-1} = \begin{bmatrix} \boldsymbol{A_{11}^{-1}} & 0 \\ -\boldsymbol{A_{22}^{-1}}\boldsymbol{A_{21}}\boldsymbol{A_{11}^{-1}} & \boldsymbol{A_{22}^{-1}}. \end{bmatrix}$$

6.5 Additional Properties

6.5.1 Sylvester's Determinant Theorem

Sylvester's determinant theorem states that for $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ and $\boldsymbol{B} \in \mathbb{R}^{n\times m}$,

$$\det(\boldsymbol{I_m} + \boldsymbol{AB}) = \det(\boldsymbol{I_n} + \boldsymbol{BA}).$$

From this theorem, we can also derive the following consequences:

1. For two column vectors $\boldsymbol{u} \in \mathbb{R}^m$ and $\boldsymbol{v} \in \mathbb{R}^m$,

$$\det(\boldsymbol{I_m} + \boldsymbol{uv}^T) = \det(1 + \boldsymbol{v}^T\boldsymbol{u}) = 1 + \boldsymbol{v}^T\boldsymbol{u}.$$

2. For $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, $\boldsymbol{B} \in \mathbb{R}^{n\times m}$, and an invertible matrix $\boldsymbol{X} \in \mathbb{R}^{m\times m}$,

$$\begin{aligned} \det(\boldsymbol{X} + \boldsymbol{AB}) &= \det\big(\boldsymbol{X}(\boldsymbol{I_m} + \boldsymbol{X}^{-1}\boldsymbol{AB})\big) \\ &= \det(\boldsymbol{X})\det(\boldsymbol{I_m} + \boldsymbol{X}^{-1}\boldsymbol{AB}) \\ &= \det(\boldsymbol{X})\det(\boldsymbol{I_n} + \boldsymbol{BX}^{-1}\boldsymbol{A}). \end{aligned}$$

3. For column vectors $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^m$ and an invertible matrix $\boldsymbol{X} \in \mathbb{R}^{m\times m}$,

$$\begin{aligned} \det(\boldsymbol{X} + \boldsymbol{uv}^T) &= \det\big(\boldsymbol{X}(\boldsymbol{I_m} + \boldsymbol{X}^{-1}\boldsymbol{uv}^T)\big) = \det(\boldsymbol{X})\det(\boldsymbol{I_m} + \boldsymbol{X}^{-1}\boldsymbol{uv}^T) \\ &= \det(\boldsymbol{X})\det(1 + \boldsymbol{v}^T\boldsymbol{X}^{-1}\boldsymbol{u}) = \det(\boldsymbol{X})(1 + \boldsymbol{v}^T\boldsymbol{X}^{-1}\boldsymbol{u}) \\ &= \det(\boldsymbol{X}) + \boldsymbol{v}^T\det(\boldsymbol{X})\boldsymbol{X}^{-1}\boldsymbol{u} = \det(\boldsymbol{X}) + \boldsymbol{v}^T\mathrm{adj}(\boldsymbol{X})\boldsymbol{u}. \end{aligned}$$

6.5.2 Determinant of Block Matrices

There are also special properties of the determinant for block matrices composed of four submatrices: $\boldsymbol{A} \in \mathbb{R}^{n\times n}$, $\boldsymbol{B} \in \mathbb{R}^{n\times m}$, $\boldsymbol{C} \in \mathbb{R}^{m\times n}$, and $\boldsymbol{D} \in \mathbb{R}^{m\times m}$.

1. For an upper triangular, lower triangular, or diagonal matrix,

$$\det\begin{bmatrix} \boldsymbol{A} & 0 \\ \boldsymbol{C} & \boldsymbol{D} \end{bmatrix} = \det\begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ 0 & \boldsymbol{D} \end{bmatrix} = \det\begin{bmatrix} \boldsymbol{A} & 0 \\ 0 & \boldsymbol{D} \end{bmatrix} = \det(\boldsymbol{A})\det(\boldsymbol{D}).$$

2. If $\boldsymbol{A}$ is invertible, then

$$\det \begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{C} & \boldsymbol{D} \end{bmatrix} = \det(\boldsymbol{A}) \det(\boldsymbol{D} - \boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B}).$$

   If $\boldsymbol{D}$ is invertible, then

$$\det \begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{C} & \boldsymbol{D} \end{bmatrix} = \det(\boldsymbol{D}) \det(\boldsymbol{A} - \boldsymbol{B}\boldsymbol{D}^{-1}\boldsymbol{C}).$$

3. If the blocks are square matrices of the same size (i.e. $n = m$) and $\boldsymbol{C}$ and $\boldsymbol{D}$ commute (i.e. $\boldsymbol{C}\boldsymbol{D} = \boldsymbol{D}\boldsymbol{C}$), then

$$\det \begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{C} & \boldsymbol{D} \end{bmatrix} = \det(\boldsymbol{A}\boldsymbol{D} - \boldsymbol{B}\boldsymbol{C}).$$

4. If $\boldsymbol{A} = \boldsymbol{D}$ and $\boldsymbol{B} = \boldsymbol{C}$, then regardless of whether $\boldsymbol{A}$ and $\boldsymbol{B}$ commute,

$$\det \begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{B} & \boldsymbol{A} \end{bmatrix} = \det(\boldsymbol{A} - \boldsymbol{B}) \det(\boldsymbol{A} + \boldsymbol{B}).$$

6.6 Eigenvalues and Eigenvectors

If $\boldsymbol{A}$ is an $n \times n$ matrix, then it has n **eigenvalues**, $\lambda_i \in \mathbb{C}$ for $i = 1, \ldots, n$, and n **(right) eigenvectors**, $\boldsymbol{u_i} \in \mathbb{C}^n$ for $i = 1, \ldots, n$, which satisfy

$$\boldsymbol{A}\boldsymbol{u_i} = \lambda_i \boldsymbol{u_i}, \ i = 1, \ldots, n.$$

This relationship can be equivalently expressed as

$$(\lambda_i \boldsymbol{I_n} - \boldsymbol{A})\boldsymbol{u_i} = \boldsymbol{0}_n, \ i = 1, \ldots, n.$$

$\boldsymbol{A}$ also has n **left eigenvectors**, $\boldsymbol{v_i}^T \in \mathbb{C}^n$ for $i = 1, \ldots, n$, which satisfy

$$\boldsymbol{v_i}^T \boldsymbol{A} = \lambda_i \boldsymbol{v_i}^T, \ i = 1, \ldots, n.$$

This relationship can be equivalently expressed as

$$\boldsymbol{v_i}^T(\lambda_i \boldsymbol{I_n} - \boldsymbol{A}) = \boldsymbol{0}_n, \ i = 1, \ldots, n.$$

6.6.1 Characteristic Polynomial

The **characteristic polynomial** of an $n \times n$ matrix $\boldsymbol{A}$ is denoted χ_A and is a polynomial of degree n in s, which is defined such that

$$\chi_A(s) = \det(s\boldsymbol{I_n} - \boldsymbol{A}).$$

The eigenvalues of $\boldsymbol{A}$ are the roots are the characteristic polynomial, which means that if λ_i is an eigenvalue of $\boldsymbol{A}$, then $\chi_A(\lambda_i) = 0$. The set of eigenvalues of $\boldsymbol{A}$ is called the **spectrum** of $\boldsymbol{A}$ and is often denoted $\lambda(\boldsymbol{A}) = \{\lambda_1, \ldots, \lambda_n\}$.

Any $n \times n$ matrix $\boldsymbol{A}$ has n eigenvalues, $\lambda_1, \ldots, \lambda_n$, but some of these eigenvalues may be repeated. Suppose that the first k eigenvalues of $\boldsymbol{A}$ are distinct, where $k \leq n$. Each of these k eigenvalues has an associated algebraic multiplicity $\mu_i \geq 1$, and all of the multiplicities must sum to the total number of eigenvalues (i.e. $\sum_{i=1}^{k} \mu_i = n$). The characteristic polynomial of $\boldsymbol{A}$ can be expressed as

$$\chi_A(s) = (s - \lambda_1)^{\mu_1}(s - \lambda_2)^{\mu_2} \ldots (s - \lambda_k)^{\mu_k}.$$

6.6.2 Eigenspaces

For each distinct eigenvalue, λ_i, there is a corresponding subspace of eigenvectors called the **eigenspace**, which is denoted ϕ_i and is defined as

$$\phi_i = N(\lambda_i \boldsymbol{I_n} - \boldsymbol{A}) = \{\boldsymbol{x} \in \mathbb{C}^n : (\lambda_i \boldsymbol{I_n} - \boldsymbol{A})\boldsymbol{x} = \boldsymbol{0}_n\}.$$

If an eigenvector $\boldsymbol{u_i}$ is an element of the eigenspace ϕ_i, then $(\lambda_i \boldsymbol{I_n} - \boldsymbol{A})\boldsymbol{u_i} = \boldsymbol{0}_n$. Within a given eigenspace, there are a infinitely many linearly dependent eigenvectors and finitely many linearly independent eigenvectors. If two eigenvectors are in two different eigenspaces, then they are linearly independent.

Each eigenvalue λ_i with multiplicity μ_i has μ_i linearly independent base eigenvectors and generalized eigenvectors. While base eigenvectors are an element of the null space of $(\lambda_i \boldsymbol{I_n} - \boldsymbol{A})$, generalized eigenvectors are an element of the null space of $(\lambda_i \boldsymbol{I_n} - \boldsymbol{A})^j$, where j is an integer greater than one. Generalized eigenvectors are covered in more detail in the Jordan form section.

Note that for any $n \times n$ matrix $\boldsymbol{A}$ with eigenvalue λ_i, we can write

$$N\big((\lambda_i \boldsymbol{I_n} - \boldsymbol{A})^j\big) \subseteq N\big((\lambda_i \boldsymbol{I_n} - \boldsymbol{A})^{j+1}\big)$$

$$\dim N\big((\lambda_i \boldsymbol{I_n} - \boldsymbol{A})^j\big) \leq \dim N\big((\lambda_i \boldsymbol{I_n} - \boldsymbol{A})^{j+1}\big)$$

6.6.3 Minimum Polynomial

The **characteristic equation** of $\boldsymbol{A}$ is given by $\chi_A(s) = 0$, where χ_A is the characteristic polynomial. The **Cayley Hamilton theorem** says that every square matrix satisfies its own characteristic equation. This means that if an $n \times n$ matrix $\boldsymbol{A}$ admits the characteristic polynomial $\chi_A(s)$, then $\chi_A(A) = 0$.

The **minimum polynomial** of $\boldsymbol{A}$, which we denote ψ_A, is the polynomial of least degree degree that satisfies $\psi_A(A) = 0$. We call the equation $\psi_A(s) = 0$ the **minimum equation**. The minimum polynomial has the general form

$$\psi_A(s) = (s - \lambda_1)^{m_1}(s - \lambda_2)^{m_2} \ldots (s - \lambda_k)^{m_k},$$

where $0 < m_i \leq \mu_i$ for $i = 1, \ldots, k$. From the minimum polynomial, we can say the following about the eigenspaces of $\boldsymbol{A}$:

$$\mathbb{C}^n = N\Big((A - \lambda_1 \boldsymbol{I_n})^{m_1}\Big) \oplus \ldots \oplus N\Big((A - \lambda_k \boldsymbol{I_n})^{m_k}\Big)$$

$$\dim N\Big((A - \lambda_i \boldsymbol{I_n})^{m_i}\Big) = \mu_i, \ i = 1, \ldots, k$$

6.6.4 Trace & Determinant

The trace and determinant of a matrix are closely related to the eigenvalues of the matrix. If $\boldsymbol{A}$ is an $n \times n$ matrix whose eigenvalues are $\lambda_i,\ i = 1, \ldots, n$, then

$$\text{tr}(\boldsymbol{A}) = \sum_{i=1}^{n} \lambda_i \quad \text{and} \quad \det(\boldsymbol{A}) = \prod_{i=1}^{n} \lambda_i.$$

This gives us the following relationships between trace and determinant:

$$\det\big(\exp(\boldsymbol{A})\big) = \exp\big(\text{tr}(\boldsymbol{A})\big) \quad \text{and} \quad \text{tr}\big(\log(\boldsymbol{A})\big) = \log\big(\det(\boldsymbol{A})\big).$$

6.6.5 Similar Matrices

Two $n \times n$ matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ are considered **similar matrices** if there exists an invertible $n \times n$ matrix $\boldsymbol{P}$ such that

$$\boldsymbol{A} = \boldsymbol{P}\boldsymbol{B}\boldsymbol{P}^{-1}.$$

The map $\boldsymbol{B} \mapsto \boldsymbol{P}\boldsymbol{B}\boldsymbol{P}^{-1}$ is called a **similarity transformation**. Any similarity transformation can be represented by a series of elementary row and column operations. Similar matrices have the same eigenvalues, which also implies that they have the same rank, determinant, and trace.

6.6.6 Special Matrices

1. **Real-Valued** – If $\boldsymbol{A}$ has all real entries, it can still have complex eigenvalues. However, its complex eigenvalues must form a pair of complex conjugates and the corresponding eigenvectors will also be conjugate pairs.

2. **Symmetric** – If $\boldsymbol{A}$ is an $n \times n$ symmetric matrix with all real entries, then it has n real eigenvalues. Additionally, it has n orthogonal eigenvectors.

3. **Diagonal** – If $\boldsymbol{A}$ is a diagonal matrix, then its eigenvalues are simply its diagonal entries. This also implies that its determinant is the product of its diagonal elements. In addition, the standard basis vectors, $\boldsymbol{e_1}, \ldots, \boldsymbol{e_n}$, form a set of n orthogonal eigenvectors for $\boldsymbol{A}$.

4. **Upper/Lower Triangular** – If $\boldsymbol{A}$ is upper or lower triangular, then its eigenvalues are simply its diagonal entries. This also implies that its determinant is the product of its diagonal elements. In addition, the standard basis vectors form a set of n orthogonal eigenvectors for $\boldsymbol{A}$.

5. **Square Dyad** – If $\boldsymbol{A}$ is a square dyad that can be expressed as $\boldsymbol{A} = \boldsymbol{u}\boldsymbol{v}^T$, it has one non-zero eigenvalue $\lambda = \boldsymbol{v}^T\boldsymbol{u}$ with corresponding eigenvector $\boldsymbol{u}$.

6. **Block Diagonal** – If $\boldsymbol{A}$ is a block diagonal matrix, its eigenvalues are the union of the eigenvalues of each block along its diagonal. This implies that its determinant is the product of the determinant of each of its diagonal blocks, and its trace is the sum of the trace of each of its diagonal blocks.

7. **Block Upper/Lower Triangular** – If $\boldsymbol{A}$ is a block upper or lower triangular matrix, then its eigenvalues are the union of the eigenvalues of each block along its diagonal. This implies that its determinant is the product of the determinant of each of its diagonal blocks, and its trace is the sum of the trace of each of its diagonal blocks.

6.7 Diagonal Form

Suppose $\boldsymbol{A}$ is an $n \times n$ matrix whose eigenvalues are $\lambda_1, \ldots, \lambda_n$ and whose corresponding eigenvectors are $\boldsymbol{u_1}, \ldots, \boldsymbol{u_n}$. $\boldsymbol{A}$ is diagonalizable if and only if $\boldsymbol{A}$ has n linearly independent eigenvectors. Note that if $\boldsymbol{A}$ has n distinct eigenvalues, then it has n linearly independent eigenvectors, but this is not a necessary condition. The eigenvalues and eigenvectors of $\boldsymbol{A}$ must satisfy $A\boldsymbol{u_i} = \lambda_i \boldsymbol{u_i}$ for $i = 1, \ldots, n$. We can combine these equations into a single matrix equation:

$$\boldsymbol{A}\begin{bmatrix} | & & | \\ \boldsymbol{u_1} & \ldots & \boldsymbol{u_n} \\ | & & | \end{bmatrix} = \begin{bmatrix} | & & | \\ \boldsymbol{u_1} & \ldots & \boldsymbol{u_n} \\ | & & | \end{bmatrix} \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix}.$$

Now we can define the following matrices to simplify our expression:

$$\boldsymbol{U} = \begin{bmatrix} | & & | \\ \boldsymbol{u_1} & \ldots & \boldsymbol{u_n} \\ | & & | \end{bmatrix} \in \mathbb{C}^{n \times n} \quad \text{and} \quad \boldsymbol{\Lambda} = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} \in \mathbb{C}^{n \times n}.$$

This allows us to write the previous matrix equation as $\boldsymbol{AU} = \boldsymbol{U\Lambda}$. Because the matrix $\boldsymbol{U}$ is composed of n linearly independent eigenvectors, it is invertible. This allows us to express $\boldsymbol{A}$ in its **diagonal form**:

$$\boldsymbol{A} = \boldsymbol{U\Lambda U}^{-1}.$$

Now let's define $\boldsymbol{U}^{-1}$ as the matrix $\boldsymbol{V}$, which can be expressed as

$$\boldsymbol{V} = \begin{bmatrix} — & \boldsymbol{v_1}^T & — \\ & \vdots & \\ — & \boldsymbol{v_n}^T & — \end{bmatrix} \in \mathbb{C}^{n \times n}, = .$$

Now we have the following matrix equation: $\boldsymbol{A} = \boldsymbol{V}^{-1}\boldsymbol{\Lambda V}$. Because $\boldsymbol{V}^{-1}$ is invertible, we can also express this equation as $\boldsymbol{VA} = \boldsymbol{\Lambda V}$, or equivalently

$$\begin{bmatrix} — & \boldsymbol{v_1}^T & — \\ & \vdots & \\ — & \boldsymbol{v_n}^T & — \end{bmatrix} \boldsymbol{A} = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} \begin{bmatrix} — & \boldsymbol{v_1}^T & — \\ & \vdots & \\ — & \boldsymbol{v_n}^T & — \end{bmatrix}.$$

We can break this matrix equation into n independent equations: $\boldsymbol{v_i}^T\boldsymbol{A} = \lambda_i \boldsymbol{v_i}^T$ for $i = 1, \ldots, n$. Now we can see that $\boldsymbol{v_i}^T$ is a left eigenvector of $\boldsymbol{A}$. Therefore, for

an $n \times n$ diagonalizable matrix $\boldsymbol{A}$ with eigenvalues $\lambda_1, \ldots, \lambda_n$, right eigenvectors $\boldsymbol{u_1}, \ldots, \boldsymbol{u_n}$, and left eigenvectors $\boldsymbol{v_1^T}, \ldots, \boldsymbol{v_n^T}$, we can express $\boldsymbol{A}$ as

$$\boldsymbol{A} = \begin{bmatrix} | & & | \\ \boldsymbol{u_1} & \ldots & \boldsymbol{u_n} \\ | & & | \end{bmatrix} \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} \begin{bmatrix} — & \boldsymbol{v_1^T} & — \\ & \vdots & \\ — & \boldsymbol{v_n^T} & — \end{bmatrix}.$$

If $\boldsymbol{\Lambda}$ is a diagonal matrix whose diagonal entries are the eigenvectors of $\boldsymbol{A}$, $\boldsymbol{U}$ is an orthogonal matrix whose columns are the right eigenvectors of $\boldsymbol{A}$, $\boldsymbol{V}$ is an orthogonal matrix columns are the left eigenvectors of $\boldsymbol{A}$, then

$$\boldsymbol{A} = \boldsymbol{U\Lambda V}.$$

This also allows us to express the matrix $\boldsymbol{A}$ as a sum of dyads:

$$\boldsymbol{A} = \sum_{i=1}^{n} \lambda_i \boldsymbol{u_i v_i^T}.$$

6.7.1 Matrix Inverse & Product

The diagonal form is useful because diagonal matrices have some important properties. As mentioned previously, the eigenvalues of a diagonal matrix are simply the elements along its diagonal. This means that the determinant of a diagonal matrix is the product of its diagonal entries. If a matrix is diagonalizable, we can use its diagonal form to more easily compute matrix inverses.

$$A^{-1} = (\boldsymbol{U\Lambda U}^{-1})^{-1} = \boldsymbol{U\Lambda}^{-1}\boldsymbol{U}^{-1}$$

$$\boldsymbol{\Lambda}^{-1} = \begin{bmatrix} \frac{1}{\lambda_1} & & \\ & \ddots & \\ & & \frac{1}{\lambda_n} \end{bmatrix}$$

In a similar way, if a matrix is diagonalizable, we can use its diagonal form to more easily compute matrix products.

$$A^k = (\boldsymbol{U\Lambda U}^{-1})^k = \boldsymbol{U\Lambda}^k\boldsymbol{U}^{-1}$$

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

6.8 Jordan Form

We cannot always diagonalize a square matrix $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ because $\boldsymbol{A}$ may not have n linearly independent columns. However, given any matrix $\boldsymbol{A} \in \mathbb{R}^{n \times n}$, we can transform it into **Jordan canonical form** by the similarity transform

$\boldsymbol{A} = \boldsymbol{T}\boldsymbol{J}\boldsymbol{T}^{-1}$, where $\boldsymbol{T}$ is an invertible matrix whose columns are the base and generalized eigenvectors of $\boldsymbol{A}$ and $\boldsymbol{J}$ is a block diagonal matrix of the form

$$\boldsymbol{J} = \begin{bmatrix} \boldsymbol{J}_1 & & \\ & \ddots & \\ & & \boldsymbol{J}_l \end{bmatrix}, \quad \text{where} \quad \boldsymbol{J}_i = \begin{bmatrix} \lambda_i & 1 & & \\ & \ddots & \ddots & \\ & & \lambda_i & 1 \\ & & & \lambda_i \end{bmatrix}.$$

6.8.1 Obtaining the Jordan Form

To find the Jordan form of $\boldsymbol{A} \in \mathbb{R}^{n \times n}$, we go through the following procedure:

1. Find all of the eigenvalues of $\boldsymbol{A}$, $\lambda_1, \ldots, \lambda_n$, by computing the roots of the characteristic polynomial $\chi_A(s) = \det(sI - A)$.
2. For each distinct eigenvalue λ_i,
   (a) Find all of the linearly independent base eigenvectors $\boldsymbol{u_1}, \ldots, \boldsymbol{u_{c_i}}$ by solving the equation $(\boldsymbol{A} - \lambda_i \boldsymbol{I_n})\boldsymbol{u_j} = \boldsymbol{0_n}$, where $j = 1, \ldots, c_i$.
   (b) For each base eigenvector $\boldsymbol{u_j}$, compute the generalized eigenvectors:
      i. Check if there is a solution to $(\boldsymbol{A} - \lambda_i \boldsymbol{I_n})\boldsymbol{w_j^1} = \boldsymbol{u_j}$. If there is, find $\boldsymbol{w_j^1}$. If there is not, look at the next eigenvector.
      ii. Continue finding generalized eigenvectors by solving $(\boldsymbol{A} - \lambda_i \boldsymbol{I_n})\boldsymbol{w_j^l} = \boldsymbol{w_j^{l-1}}$ such that $\boldsymbol{w_j^l}$ and $\boldsymbol{w_j^{l-1}}$ are linearly independent.
3. Construct $\boldsymbol{T}$ using the eigenvectors and generalized eigenvectors such that
$$\boldsymbol{T} = \begin{bmatrix} \boldsymbol{u_1} & \boldsymbol{w_1^1} & \ldots & \boldsymbol{u_2} & \boldsymbol{w_2^1} & \ldots \end{bmatrix}.$$
4. Construct the blocks of $\boldsymbol{J}$ with the eigenvalues λ_i along the diagonal, corresponding to each base eigenvector, and ones along the upper off-diagonal, corresponding to each generalized eigenvector.

6.8.2 Jordan Form & Minimum Polynomial

The Jordan form of a matrix is also closely related to its characteristic and minimum polynomial. Consider a square matrix $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ with the Jordan form $\boldsymbol{A} = \boldsymbol{T}\boldsymbol{J}\boldsymbol{T}^{-1}$, whose characteristic and minimum polynomial are given by

$$\chi_A(s) = (s - \lambda_1)^{\mu_1}(s - \lambda_2)^{\mu_2} \ldots (s - \lambda_k)^{\mu_k}$$

and

$$\psi_A(s) = (s - \lambda_1)^{m_1}(s - \lambda_2)^{m_2} \ldots (s - \lambda_k)^{m_k}.$$

These two polynomials tell us the following about the Jordan matrix $\boldsymbol{J}$:

1. μ_i = sum of sizes of the Jordan blocks corresponding to eigenvalue λ_i
2. m_i = size of largest Jordan block corresponding to eigenvalue λ_i
3. number of linearly independent eigenvectors = number of Jordan blocks

6.9 Matrix Functions

6.9.1 Polynomial Matrix Functions

If p is a polynomial function in s defined such that

$$p(s) = a_m s^m + a_{m-1}s^{m-1} + \ldots + a_1 s + a_0,$$

then we can expand this definition for a square matrix $\boldsymbol{A}$ and express $p(\boldsymbol{A})$ as

$$p(\boldsymbol{A}) = a_m \boldsymbol{A}^m + a_{m-1}\boldsymbol{A}^{m-1} + \ldots + a_1 \boldsymbol{A} + a_0.$$

If $(\lambda_i, \boldsymbol{u_i})$ is an eigenvalue-eigenvector pair for the matrix $\boldsymbol{A}$, then $(p(\lambda_i), \boldsymbol{u_i})$ is an eigenvalue-eigenvector pair for the matrix $p(\boldsymbol{A})$. This leads to the **eigenvalue shift rul**e, which says if $\lambda_i(\boldsymbol{A})$ is the ith eigenvalue of the matrix $\boldsymbol{A}$,

$$\lambda_i(\boldsymbol{A} + \alpha \boldsymbol{I_n}) = \lambda_i(\boldsymbol{A}) + \alpha.$$

If the $n \times n$ matrix $\boldsymbol{A}$ is diagonalizable and admits a diagonal factorization of the form $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{-1}$, then we can express $p(\boldsymbol{A})$ as

$$p(\boldsymbol{A}) = \boldsymbol{U}p(\boldsymbol{\Lambda})\boldsymbol{U}^{-1}, \text{ where } p(\boldsymbol{\Lambda}) = \begin{bmatrix} p(\lambda_1) & & \\ & \ddots & \\ & & p(\lambda_n) \end{bmatrix}.$$

6.9.2 Cayley Hamilton Theorem

Recall that the Cayley Hamilton theorem says that every square matrix satisfies its own characteristic equation. This means that if the matrix $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ admits the characteristic polynomial

$$\chi_A(s) = s^n + \alpha_1 s^{n-1} + \ldots + \alpha_{n-1} s + \alpha_n,$$

then the Cayley Hamilton theorem tells us that

$$\chi_A(\boldsymbol{A}) = \boldsymbol{A}^n + \alpha_1 \boldsymbol{A}^{n-1} + \ldots + \alpha_{n-1}\boldsymbol{A} + \alpha_n I = 0.$$

The Cayley Hamilton theorem is useful because it allows us to write any polynomial function of $\boldsymbol{A}$ as a function of order $n-1$ or less. This is because

$$\begin{aligned} \boldsymbol{A}^n &= -\alpha_1 \boldsymbol{A}^{n-1} - \ldots - \alpha_{n-1}\boldsymbol{A} - \alpha_n I \\ \boldsymbol{A}^{n+1} &= \boldsymbol{A}(\boldsymbol{A}^n) = \boldsymbol{A}\big(-\alpha_1 \boldsymbol{A}^{n-1} - \ldots - \alpha_{n-1}\boldsymbol{A} - \alpha_n I\big) \\ &= -\alpha_1 \boldsymbol{A}^n - \ldots - \alpha_{n-1}\boldsymbol{A}^2 - \alpha_n \boldsymbol{A} \\ &= -\alpha_1\big(-\alpha_1 \boldsymbol{A}^{n-1} - \ldots - \alpha_{n-1}\boldsymbol{A} - \alpha_n I\big) - \ldots - \alpha_{n-1}\boldsymbol{A}^2 - \alpha_n \boldsymbol{A} \end{aligned}$$

6.9.3 Analytic Matrix Functions

If f is an analytic function in s, it can be expressed as a convergent power series:

$$f(s) = \sum_{n=0}^{\infty} a_n(s - s_0)^n = a_0 + a_1(s - s_0) + a_2(s - s_0)^2 + \dots$$

Some examples of analytic functions are polynomials, exponential functions, trigonometic functions, logarithms, and power functions. There are other types of analytic functions, but these are the most common. We can expand the definition of analytic functions for a square matrix $\boldsymbol{A}$ and express $f(\boldsymbol{A})$ as

$$f(\boldsymbol{A}) = \sum_{n=0}^{\infty} a_n(\boldsymbol{A} - \boldsymbol{A}_0)^n = a_0 I + a_1(\boldsymbol{A} - \boldsymbol{A}_0) + a_2(\boldsymbol{A} - \boldsymbol{A}_0)^2 + \dots$$

The **spectral mapping theorem** states that if f is an analytic function and $\boldsymbol{A}$ is a square matrix with eigenvalues $\lambda(\boldsymbol{A})$, then the eigenvalues of $f(\boldsymbol{A})$ are $f(\lambda(\boldsymbol{A}))$. Additionally, if $\boldsymbol{A}$ is diagonalizable and admits a diagonal factorization of the form $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{-1}$, then $f(\boldsymbol{A})$ can be expressed as

$$f(\boldsymbol{A}) = \boldsymbol{U} f(\boldsymbol{\Lambda}) \boldsymbol{U}^{-1}, \text{ where } f(\boldsymbol{\Lambda}) = \begin{bmatrix} f(\lambda_1) & & \\ & \ddots & \\ & & f(\lambda_n) \end{bmatrix}.$$

If $\boldsymbol{A}$ is not diagonalizable, we can represent it in its Jordan form as $\boldsymbol{A} = \boldsymbol{T}\boldsymbol{J}\boldsymbol{T}^{-1}$. If we define $f^{(k)}(\lambda_i) = \frac{\partial^k}{\partial s^k} f(s)|_{s=\lambda_i}$, then we can express $f(\boldsymbol{A})$ as

$$f(\boldsymbol{A}) = \boldsymbol{T} f(\boldsymbol{J}) \boldsymbol{T}^{-1}, \text{ where } f(\boldsymbol{J}) = \begin{bmatrix} f(\boldsymbol{J_1}) & & \\ & \ddots & \\ & & f(\boldsymbol{J_l}) \end{bmatrix} \text{ and}$$

$$f(\boldsymbol{J_i}) = \begin{bmatrix} f(\lambda_i) & f'(\lambda_i) & \dots & \frac{1}{(n_i-1)!} f^{(n_i-1)}(\lambda_i) \\ 0 & f(\lambda_i) & \dots & \frac{1}{(n_i-2)!} f^{(n_i-2)}(\lambda_i) \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & f'(\lambda_i) \\ 0 & 0 & \dots & f(\lambda_i) \end{bmatrix}.$$

Note that λ_i is the eigenvalue corresponding to the ith Jordan block in $\boldsymbol{J}$ and n_i is the size of the ith Jordan block in $\boldsymbol{J}$.

6.9.4 Analytic Functions as Polynomials

Let $\boldsymbol{A}$ be an $n \times n$ matrix with eigenvalues $\lambda_1, \dots, \lambda_n$. Assume that the first k eigenvalues of $\boldsymbol{A}$ are distinct, where $k \leq n$. Let $m_1, \dots, m_k$ be the powers in

the minimum polynomial of $\boldsymbol{A}$. If f is an analytic function of s defined on the spectrum of A and p is a polynomial function of s such that

$$f^{(l)}(\lambda_i) = p^{(l)}(\lambda_i),\ l = 0, \ldots, m_i - 1,\ i = 1, \ldots, k,$$

then $f(\boldsymbol{A}) = p(\boldsymbol{A})$. We can choose any form of polynomial to express an analytic function in this way, but we generally choose a polynomial of the form

$$p(s) = a_1 s^{m-1} + \ldots + a_{m-1} s + a_m,$$

where $m = \sum_{i=1}^{k} m_i$. This allows us to set up a system of m equations that satisfy the equalities $f^{(l)}(\lambda_i) = p^{(l)}(\lambda_i)$ for $l = 0, \ldots, m_i - 1,\ i = 1, \ldots, k$, and solve for the coefficients a_i in terms of $f^{(l)}(\lambda_i)$. We can then write $f(\boldsymbol{A})$ as

$$f(\boldsymbol{A}) = a_1 \boldsymbol{A}^{m-1} + \ldots a_{m-1} \boldsymbol{A} + a_m \boldsymbol{I_n}.$$

6.10 Matrix Exponential

One of the most common analytic matrix functions is the matrix exponential. The matrix exponential $e^{\boldsymbol{A}t} \in \mathbb{R}^{n \times n}$ for a given matrix $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ is defined as

$$e^{\boldsymbol{A}t} = \sum_{j=0}^{\infty} \frac{1}{j!} (\boldsymbol{A}t)^j = \boldsymbol{I_n} + \boldsymbol{A}t + \frac{\boldsymbol{A}^2 t^2}{2!} + \ldots$$

The matrix exponential has several useful properties:

1. $e^{\boldsymbol{0_{n,n}}} = \boldsymbol{I_n}$
2. $e^{\boldsymbol{A}(t+s)} = e^{\boldsymbol{A}t} e^{\boldsymbol{A}s}$
3. $e^{(\boldsymbol{A}+\boldsymbol{B})t} = e^{\boldsymbol{A}t} e^{\boldsymbol{B}t}$ iff $\boldsymbol{AB} = \boldsymbol{BA}$
4. $\left(e^{\boldsymbol{A}t}\right)^{-1} = e^{-\boldsymbol{A}t}$
5. $\frac{d}{dt} e^{\boldsymbol{A}t} = \boldsymbol{A} e^{\boldsymbol{A}t} = e^{\boldsymbol{A}t} \boldsymbol{A}$

6.10.1 Computing the Matrix Exponential

There are a few different ways to compute the matrix exponential.

1. **Method 1: Infinite Series**

   We can use the infinite series representation to compute the matrix exponential. This is straightforward if $\boldsymbol{A}$ is a nilpotent matrix, or if the terms of $e^{\boldsymbol{A}t}$ can be recognized as basic transformations of common infinite series. However, this method is generally difficult to use. Below are some of the most common infinite series representations.

   $$e^t = \sum_{j=0}^{\infty} \frac{x^j}{j!} = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \ldots$$

$$\sin t = \sum_{j=0}^{\infty} \frac{(-1)^j}{(2j+1)!} x^{2j+1} = x - \frac{x^3}{3!} + \frac{x^5}{5!} - \dots$$

$$\cos t = \sum_{j=0}^{\infty} \frac{(-1)^j}{(2j)!} x^{2j} = 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \dots$$

$$\frac{1}{1-x} = \sum_{j=0}^{\infty} x^j = 1 + x + x^2 + \dots$$

2. **Method 2: Laplace Transforms**

   In some cases, we may be able to use Laplace transforms to compute

   $$e^{\boldsymbol{A}t} = \mathcal{L}^{-1}\{(sI - \boldsymbol{A})^{-1}\}.$$

   Note that we compute the inverse Laplace transform element-wise. We generally want to use partial fraction decomposition, so that we can look up the inverse Laplace transform in a Laplace transform table.

3. **Method 3: Diagonal and Jordan Form**

   If $\boldsymbol{A}$ is a diagonalizable matrix that admits a diagonalization $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{-1}$,

   $$e^{\boldsymbol{A}t} = \boldsymbol{U}e^{\boldsymbol{\Lambda}t}\boldsymbol{U}^{-1}, \text{ where } e^{\boldsymbol{\Lambda}t} = \begin{bmatrix} e^{\lambda_1 t} & & \\ & \ddots & \\ & & e^{\lambda_n t} \end{bmatrix}.$$

   If $\boldsymbol{A}$ is not diagonalizable, we can still represent it in its Jordan form as $\boldsymbol{A} = \boldsymbol{T}\boldsymbol{J}\boldsymbol{T}^{-1}$. Now we can write the matrix exponential of $\boldsymbol{A}$ as

   $$e^{\boldsymbol{A}t} = \boldsymbol{T}e^{\boldsymbol{J}t}\boldsymbol{T}^{-1}, \text{ where } e^{\boldsymbol{J}t} = \begin{bmatrix} e^{\lambda_i t} & te^{\lambda_i t} & \dots & \frac{1}{(n_i-1)!}t^{(n_i-1)}e^{\lambda_i t} \\ 0 & e^{\lambda_i t} & \dots & \frac{1}{(n_i-2)!}t^{(n_i-2)}e^{\lambda_i t} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & te^{\lambda_i t} \\ 0 & 0 & \dots & e^{\lambda_i t} \end{bmatrix}.$$

6.10.2 Eigenvalues & Eigenvectors

Suppose $(\lambda_i, \boldsymbol{u_i})$ is an eigenvalue-eigenvector pair for a matrix $\boldsymbol{A}$ satisfying $\boldsymbol{A}\boldsymbol{u_i} = \lambda_i \boldsymbol{u_i}$. This implies that $e^{\boldsymbol{A}t}\boldsymbol{u_i} = e^{\lambda_i t}\boldsymbol{u_i}$. To see this, notice that

$$e^{\boldsymbol{A}t}\boldsymbol{u_i} = \sum_{j=0}^{\infty} \frac{t^j}{j!}\boldsymbol{A}^j\boldsymbol{u_i} = \sum_{j=0}^{\infty} \frac{t^j}{j!}\lambda^j\boldsymbol{u_i} = e^{\lambda_i t}\boldsymbol{u_i}.$$

Therefore, if $(\lambda_i, \boldsymbol{u_i})$ is an eigenvalue-eigenvector pair for $\boldsymbol{A}$, then $(e^{\lambda_i t}, \boldsymbol{u_i})$ is an eigenvalue-eigenvector pair for $e^{\boldsymbol{A}t}$. This implies that any eigenvector of $\boldsymbol{A}$

is also an eigenvector of $e^{\boldsymbol{A}t}$. However, in general, an eigenvector of $e^{\boldsymbol{A}t}$ is not necessarily and eigenvector of $\boldsymbol{A}$. Consider the following couterexample:

$$\boldsymbol{A} = \begin{bmatrix} 0 & \beta \\ -\beta & 0 \end{bmatrix} \qquad e^{\boldsymbol{A}t} = \begin{bmatrix} \cos(\beta t) & \sin(\beta t) \\ -\sin(\beta t) & \cos(\beta t) \end{bmatrix}$$

At time $t = \frac{2\pi}{\beta} n$, where n is an integer, the matrix exponential is simply the identity matrix. The standard basis vectors are eigenvectors of the identity matrix, but they are clearly not eigenvectors of the matrix $\boldsymbol{A}$. Therefore, the eigenvectors of $e^{\boldsymbol{A}t}$ are not eigenvectors of $\boldsymbol{A}$ in general.

Chapter 7

Singular Value Decomposition

7.1 Singular Value Decomposition (SVD)

7.1.1 Singular Values

For a square matrix $\boldsymbol{A}$, $\lambda_i \in \mathbb{C}$ is an eigenvalue of $\boldsymbol{A}$ if there exists a non-zero vector $\boldsymbol{u_i}$ such that $\boldsymbol{A u_i} = \lambda_i \boldsymbol{u_i}$. Only square matrices have eigenvalues, but matrices generally are not square. However if $\boldsymbol{A}$ is an $m \times n$, then $\boldsymbol{A A^*}$ is $m \times m$ and $\boldsymbol{A^* A}$ is $n \times n$, so both are square matrices and have eigenvalues.

Let $\lambda_1, \ldots, \lambda_m$ be the real, non-negative eigenvalues of $\boldsymbol{A A^*}$. If the rank of $\boldsymbol{A}$ is r, then r of those eigenvalues will be strictly positive and $m - r$ will be zero. Let's assume $\lambda_1 \geq \ldots \geq \lambda_r > 0$ and $\lambda_{r+1} = \ldots = \lambda_m = 0$. The non-zero singular values of the matrix $\boldsymbol{A}$ are defined as $\sigma_i = \sqrt{\lambda_i}$, $i = 1, \ldots, r$.

We also could have let $\lambda_1, \ldots, \lambda_n$ be the real, non-negative eigenvalues of $\boldsymbol{A^* A}$. If the rank of $\boldsymbol{A}$ is r, then r of those eigenvalues will be strictly positive and $n - r$ will be zero. Let's assume $\lambda_1 \geq \ldots \geq \lambda_r > 0$ and $\lambda_{r+1} = \ldots = \lambda_n = 0$. The non-zero singular values of the matrix $\boldsymbol{A}$ are $\sigma_i = \sqrt{\lambda_i}$, $i = 1, \ldots, r$.

Therefore, if $\text{rank}(\boldsymbol{A}) = r$, then the non-zero **singular values** of $\boldsymbol{A}$ are

$$\sigma_i = \sqrt{\lambda_i(\boldsymbol{A A^*})} = \sqrt{\lambda_i(\boldsymbol{A^* A})}, \; i = 1, \ldots, r.$$

7.1.2 Singular Vectors

Suppose $\boldsymbol{A}$ is an $m \times n$ matrix with rank r. The matrix $\boldsymbol{A A^*}$ is then $m \times m$ and has m eigenvalues $\lambda_1(\boldsymbol{A A^*}), \ldots, \lambda_m(\boldsymbol{A A^*})$ and corresponding eigenvectors $\boldsymbol{u_1}, \ldots, \boldsymbol{u_m}$, which satisfy $\boldsymbol{A A^* u_i} = \lambda_i(\boldsymbol{A A^*})\boldsymbol{u_i}$ for $i = 1, \ldots, m$. The vectors $\boldsymbol{u_1}, \ldots, \boldsymbol{u_m}$ are also referred to as the **left singular vectors** of $\boldsymbol{A}$.

The matrix $\boldsymbol{A^* A}$ is $n \times n$ and has n eigenvalues $\lambda_1(\boldsymbol{A^* A}), \ldots, \lambda_n(\boldsymbol{A^* A})$ and corresponding eigenvectors $\boldsymbol{v_1}, \ldots, \boldsymbol{v_n}$, which satisfy $\boldsymbol{A^* A v_i} = \lambda_i(\boldsymbol{A^* A})\boldsymbol{v_i}$. The vectors $\boldsymbol{v_i}$ that satisfy this equation are the **right singular vectors** of $\boldsymbol{A}$.

7.1.3 Full Form SVD

Suppose $\boldsymbol{A}$ is an $m \times n$ matrix with rank $r \leq \min(m, n)$. Let $\sigma_1, \ldots, \sigma_r$ denote the singular values of $\boldsymbol{A}$, $\boldsymbol{u_1}, \ldots, \boldsymbol{u_m}$ denote the left singular vectors of $\boldsymbol{A}$, and $\boldsymbol{v_1}, \ldots, \boldsymbol{v_n}$ denote the right singular vectors. There exist unitary matrices $\boldsymbol{U} \in \mathbb{C}^{m \times m}$ and $\boldsymbol{V} \in \mathbb{C}^{n \times n}$ and a block diagonal matrix $\boldsymbol{\Sigma} \in \mathbb{C}^{m \times n}$ such that

$$\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^*, \text{ where}$$

$$\boldsymbol{U} = \begin{bmatrix} \boldsymbol{u_1} & \boldsymbol{u_2} & \ldots & \boldsymbol{u_m} \end{bmatrix} \in \mathbb{C}^{m \times m},$$

$$\boldsymbol{V} = \begin{bmatrix} \boldsymbol{v_1} & \boldsymbol{v_2} & \ldots & \boldsymbol{v_n} \end{bmatrix} \in \mathbb{C}^{n \times n},$$

$$\boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma_r} & \boldsymbol{0}^{r \times (n-r)} \\ \boldsymbol{0}^{(m-r) \times r} & \boldsymbol{0}^{(m-r) \times (n-r)} \end{bmatrix} \in \mathbb{C}^{m \times n}.$$

The first unitary matrix is composed of the m left singular vectors of $\boldsymbol{A}$, and the second unitary matrix is composed of the n right singular vectors of $\boldsymbol{A}$. The block diagonal matrix $\boldsymbol{\Sigma}$ is composed of zero matrices and a diagonal matrix whose elements are the r singular values of $\boldsymbol{A}$: $\boldsymbol{\Sigma_r} = \text{diag}(\sigma_1, \ldots, \sigma_r) \in \mathbb{C}^{r \times r}$. This is called the **full form singular value decomposition (SVD)**.

7.1.4 Compact Form SVD

Previously, we wrote the full form singular value decomposition (SVD) for the $m \times n$ matrix $\boldsymbol{A}$ with rank $r \leq \min(m, n)$. We can write the SVD in a more compact form. The first unitary matrix, $\boldsymbol{U}$, can be broken into two submatrices: an $m \times r$ matrix $\boldsymbol{U_r}$ and an $m \times (m - r)$ matrix $\boldsymbol{U_{mr}}$ such that

$$\boldsymbol{U} = \begin{bmatrix} \boldsymbol{U_r} & \boldsymbol{U_{mr}} \end{bmatrix} \in \mathbb{C}^{m \times m}, \text{ where}$$

$$\boldsymbol{U_r} = \begin{bmatrix} \boldsymbol{u_1} & \ldots & \boldsymbol{u_r} \end{bmatrix} \in \mathbb{C}^{m \times r},$$

$$\boldsymbol{U_{mr}} = \begin{bmatrix} \boldsymbol{u_{r+1}} & \ldots & \boldsymbol{u_m} \end{bmatrix} \in \mathbb{C}^{m \times (m-r)}.$$

In a similar way, the second unitary matrix, $\boldsymbol{V}$, can be broken into two submatrices: an $n \times r$ matrix $\boldsymbol{V_r}$ and an $n \times (n - r)$ matrix $\boldsymbol{V_{nr}}$ such that

$$\boldsymbol{V} = \begin{bmatrix} \boldsymbol{V_r} & \boldsymbol{V_{nr}} \end{bmatrix} \in \mathbb{C}^{n \times n}, \text{ where}$$

$$\boldsymbol{V_r} = \begin{bmatrix} \boldsymbol{v_1} & \ldots & \boldsymbol{v_r} \end{bmatrix} \in \mathbb{C}^{n \times r},$$

$$\boldsymbol{V_{nr}} = \begin{bmatrix} \boldsymbol{v_{r+1}} & \ldots & \boldsymbol{v_n} \end{bmatrix} \in \mathbb{C}^{n \times (n-r)}.$$

Now we can decompose $\boldsymbol{A}$ as

$$\boldsymbol{A} = \boldsymbol{U_r}\boldsymbol{\Sigma_r}\boldsymbol{V_r}^* = \sum_{i=1}^{r} \sigma_i \boldsymbol{u_i} \boldsymbol{v_i}^*$$

This is called the **compact form singular value decomposition (SVD)**.

7.1.5 Range and Null Space

Let $\boldsymbol{A}$ be an $m \times n$ matrix with rank r that can be expressed in terms of its compact form SVD as $\boldsymbol{A} = \boldsymbol{U_r \Sigma_r V_r^*}$. This representation allows us to characterize the range space and null space of $\boldsymbol{A}$, as well as its Hermitian transpose.

Range Space of $\boldsymbol{A}$

$R(\boldsymbol{A}) = R(\boldsymbol{U_r}) = \text{span}(\{\boldsymbol{u_1}, \ldots, \boldsymbol{u_r}\})$

$R(\boldsymbol{A}) = \left\{\boldsymbol{y} \in \mathbb{C}^m : \boldsymbol{y} = \boldsymbol{U_r z},\ \boldsymbol{z} \in \mathbb{C}^r\right\}$

$\text{dim}R(\boldsymbol{A}) = \text{dim}R(\boldsymbol{U_r}) = r$

Range Space of $\boldsymbol{A}^*$

$R(\boldsymbol{A}^*) = R(\boldsymbol{V_r}) = \text{span}(\{\boldsymbol{v_1}, \ldots, \boldsymbol{v_r}\})$

$R(\boldsymbol{A}^*) = \left\{\boldsymbol{y} \in \mathbb{C}^n : \boldsymbol{y} = \boldsymbol{V_r z},\ \boldsymbol{z} \in \mathbb{C}^r\right\}$

$\text{dim}R(\boldsymbol{A}^*) = \text{dim}R(\boldsymbol{V_r}) = r$

Null Space of $\boldsymbol{A}$

$N(\boldsymbol{A}) = R(\boldsymbol{V_{nr}}) = \text{span}(\{\boldsymbol{v_{r+1}}, \ldots, \boldsymbol{v_n}\})$

$N(\boldsymbol{A}) = \left\{\boldsymbol{x} \in \mathbb{C}^n : \boldsymbol{x} = \boldsymbol{V_{nr}\alpha},\ \boldsymbol{\alpha} \in \mathbb{C}^{n-r}\right\}$

$\text{dim}N(\boldsymbol{A}) = \text{dim}R(\boldsymbol{V_{nr}}) = n - r$

Null Space of $\boldsymbol{A}^*$

$N(\boldsymbol{A}^*) = R(\boldsymbol{U_{mr}}) = \text{span}(\{\boldsymbol{u_{r+1}}, \ldots, \boldsymbol{u_m}\})$

$N(\boldsymbol{A}^*) = \left\{\boldsymbol{x} \in \mathbb{C}^m : \boldsymbol{x} = \boldsymbol{U_{mr}\alpha},\ \boldsymbol{\alpha} \in \mathbb{C}^{m-r}\right\}$

$\text{dim}N(\boldsymbol{A}^*) = \text{dim}R(\boldsymbol{U_{mr}}) = m - r$

7.2 Matrix Pseudoinverse

For some matrix $\boldsymbol{A} \in \mathbb{C}^{n \times m}$, a **pseudoinverse** of A is a matrix $\boldsymbol{A}^\dagger \in \mathbb{C}^{m \times n}$ that satisfies the following criteria, known as the Moore–Penrose conditions:

1. $\boldsymbol{A A^\dagger A} = \boldsymbol{A}$
2. $\boldsymbol{A^\dagger A A^\dagger} = \boldsymbol{A^\dagger}$
3. $(\boldsymbol{A A^\dagger})^* = \boldsymbol{A A^\dagger}$
4. $(\boldsymbol{A^\dagger A})^* = \boldsymbol{A^\dagger A}$

From conditions one and two, we can see that the pseudoinverse acts like a weak inverse. While $\boldsymbol{A}\boldsymbol{A}^\dagger$ is not necessarily the identity matrix, it does map $\boldsymbol{A}$ to itself when it left multiplies $\boldsymbol{A}$. Similarly, $\boldsymbol{A}\boldsymbol{A}^\dagger$ is not necessarily the identity matrix, but it maps $\boldsymbol{A}^\dagger$ to itself when it right multiplies $\boldsymbol{A}^\dagger$.
Let $\boldsymbol{A}$ be an $m \times n$ matrix with rank r that can be expressed with its full form SVD as $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^*$. The **Moore-Penrose pseudoinverse** of $\boldsymbol{A}$ is defined as

$$\boldsymbol{A}^\dagger = \boldsymbol{V}\boldsymbol{\Sigma}^\dagger\boldsymbol{U}^*, \text{ where}$$

$$\boldsymbol{\Sigma}^\dagger = \begin{bmatrix} \boldsymbol{\Sigma}_{\boldsymbol{r}}^{-1} & \boldsymbol{0}^{r\times(m-r)} \\ \boldsymbol{0}^{(n-r)\times r} & \boldsymbol{0}^{(n-r)\times(m-r)} \end{bmatrix} \in \mathbb{C}^{n\times m},$$

$$\boldsymbol{\Sigma}_{\boldsymbol{r}}^{-1} = \operatorname{diag}\Big(\frac{1}{\sigma_1}, \ldots, \frac{1}{\sigma_r}\Big) \in \mathbb{C}^{r\times r}.$$

We can also write the Moore–Penrose pseudoinverse in compact form as

$$\boldsymbol{A}^\dagger = \boldsymbol{V}_{\boldsymbol{r}}\boldsymbol{\Sigma}_{\boldsymbol{r}}^{-1}\boldsymbol{U}_{\boldsymbol{r}}^* \in \mathbb{C}^{n\times m}.$$

The pseudoinverse takes on recognizable forms given certain conditions on $\boldsymbol{A}$.

1. **Invertible** – If $\boldsymbol{A}$ is invertible (i.e. $r = n = m$), then $\boldsymbol{A}^\dagger = \boldsymbol{A}^{-1}$.
2. **Full Column Rank** – If $\boldsymbol{A}$ has full column rank (i.e. $r = n \leq m$), then $\boldsymbol{A}^\dagger\boldsymbol{A} = \boldsymbol{V}\boldsymbol{V}^* = \boldsymbol{I}_{\boldsymbol{n}}$, so $\boldsymbol{A}^\dagger$ is the left inverse of $\boldsymbol{A}$. For this case,
$$\boldsymbol{A}^\dagger = (\boldsymbol{A}^*\boldsymbol{A})^{-1}\boldsymbol{A}^*.$$
3. **Full Row Rank** – If $\boldsymbol{A}$ has full row rank (i.e. $r = m \leq n$), then $\boldsymbol{A}\boldsymbol{A}^\dagger = \boldsymbol{U}\boldsymbol{U}^* = \boldsymbol{I}_{\boldsymbol{m}}$, so $\boldsymbol{A}^\dagger$ is the right inverse of $\boldsymbol{A}$. For this case,
$$\boldsymbol{A}^\dagger = \boldsymbol{A}^*(\boldsymbol{A}\boldsymbol{A}^*)^{-1}.$$

Chapter 8

Symmetric Matrices

8.1 Symmetric Matrices

8.1.1 Spectral Decomposition

Recall that if an $n \times n$ matrix is symmetric, then it has n real eigenvalues and n orthogonal eigenvectors. Let $\mathbb{S}^n$ denote the set of $n \times n$ symmetric matrices. Any matrix $\boldsymbol{A} \in \mathbb{S}^n$ with real eigenvalues $\lambda_1, \ldots, \lambda_n$ and orthonormal eigenvectors $\boldsymbol{u_1}, \ldots, \boldsymbol{u_n}$ can be factored into its **spectral decomposition** as $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^T$, where $\boldsymbol{U}$ is an orthogonal matrix composed of the eigenvectors of $\boldsymbol{A}$ such that

$$\boldsymbol{U} = \begin{bmatrix} \boldsymbol{u_1} & \ldots & \boldsymbol{u_n} \end{bmatrix} \in \mathbb{C}^{n \times n}$$

and $\boldsymbol{\Lambda}$ is a diagonal matrix whose diagonal entries are the eigenvalues of $\boldsymbol{A}$:

$$\boldsymbol{\Lambda} = \text{diag}(\lambda_1, \ldots, \lambda_n) \in \mathbb{C}^{n \times n}.$$

We can also express the symmetric matrix $\boldsymbol{A}$ as a sum of dyads:

$$\boldsymbol{A} = \sum_{i=1}^{n} \lambda_i \boldsymbol{u_i} \boldsymbol{u_i}^T.$$

8.1.2 Rayleigh Quotient

Let $\boldsymbol{A}$ be an $n \times n$ symmetric matrix with eigenvalues $\lambda_1 \geq \ldots \geq \lambda_n$ and corresponding orthonormal eigenvectors $\boldsymbol{u_1}, \ldots, \boldsymbol{u_n}$. The **Rayleigh quotient** of $\boldsymbol{A}$ for some n-dimensional vector $\boldsymbol{x}$ is the following fraction:

$$R_{\boldsymbol{A}}(\boldsymbol{x}) = \frac{\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x}}{\boldsymbol{x}^T \boldsymbol{x}}.$$

The **Rayleigh quotient theorem** says that if λ_1 denotes the largest eigenvalue of $\boldsymbol{A}$ and λ_n denotes the smallest eigenvalue, then the following inequalities hold:

$$\lambda_n \leq \frac{\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x}}{\boldsymbol{x}^T \boldsymbol{x}} \leq \lambda_1, \ \forall \boldsymbol{x} \neq \boldsymbol{0}_n.$$

Furthermore, the largest and smallest eigenvalues of A can be expressed as

$$\lambda_1 = \max_{\boldsymbol{x}:||\boldsymbol{x}||_2=1} \boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} \quad \text{and} \quad \lambda_n = \min_{\boldsymbol{x}:||\boldsymbol{x}||_2=1} \boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x}.$$

The maximum value of $\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x}$ is attained for $\hat{\boldsymbol{x}}_{max} = \boldsymbol{u_1}$, where $\boldsymbol{u_1}$ is the eigenvector corresponding to λ_1. Similarly, the minimum value of $\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x}$ is attained for $\hat{\boldsymbol{x}}_{min} = \boldsymbol{u_n}$, where $\boldsymbol{u_n}$ is the eigenvector corresponding to the λ_n.

Proof: Because we assume $\boldsymbol{A}$ is symmetric, it admits the spectral decomposition $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^T$, where $\boldsymbol{U}$ is an orthogonal matrix composed of the eigenvectors of $\boldsymbol{A}$ and $\boldsymbol{\Lambda}$ is a diagonal matrix whose diagonal entries are the eigenvalues of $\boldsymbol{A}$. If we define the vector $\boldsymbol{y}$ such that $\boldsymbol{y} = \boldsymbol{U}^T\boldsymbol{x}$, then we can see the following:

$$\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} = \boldsymbol{x}^T\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^T\boldsymbol{x} = \boldsymbol{y}^T\boldsymbol{\Lambda}\boldsymbol{y}$$

$$||\boldsymbol{y}||_2^2 = \boldsymbol{y}^T\boldsymbol{y} = \boldsymbol{x}^T\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{x} = \boldsymbol{x}^T\boldsymbol{I_n}\boldsymbol{x} = \boldsymbol{x}^T\boldsymbol{x} = ||\boldsymbol{x}||_2^2$$

$$\max_{||\boldsymbol{x}||_2=1} \boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} = \max_{||\boldsymbol{y}||_2=1} \boldsymbol{y}^T\boldsymbol{\Lambda}\boldsymbol{y} = \max_{||\boldsymbol{y}||_2=1} \sum_{i=1}^{n} \lambda_i y_i^2$$

$$\min_{||\boldsymbol{x}||_2=1} \boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} = \min_{||\boldsymbol{y}||_2=1} \boldsymbol{y}^T\boldsymbol{\Lambda}\boldsymbol{y} = \min_{||\boldsymbol{y}||_2=1} \sum_{i=1}^{n} \lambda_i y_i^2$$

The constraint on $\boldsymbol{y}$ tells us that all of the squared elements of $\boldsymbol{y}$ must sum to one. Now it is relatively clear that the optimal solution to the maximization problem is $\hat{\boldsymbol{y}}_{max} = \boldsymbol{e_1}$ and the optimal solution to the minimization problem is $\hat{\boldsymbol{y}}_{min} = \boldsymbol{e_n}$, where e_i is the ith standard basis vector. Because $\boldsymbol{U}$ is an orthogonal matrix composed of the eigenvectors of $\boldsymbol{A}$ and we defined $\boldsymbol{y}$ such that $\boldsymbol{y} = \boldsymbol{U}^T\boldsymbol{x}$, we can also write the relationship $\boldsymbol{x} = \boldsymbol{U}\boldsymbol{y}$. Therefore,

$$\hat{\boldsymbol{x}}_{max} = \boldsymbol{U}\hat{\boldsymbol{y}}_{max} = \boldsymbol{U}\boldsymbol{e_1} = \boldsymbol{u_1}$$

$$\hat{\boldsymbol{x}}_{min} = \boldsymbol{U}\hat{\boldsymbol{y}}_{min} = \boldsymbol{U}\boldsymbol{e_n} = \boldsymbol{u_n}$$

Furthermore, the optimal values are given by

$$\max_{||\boldsymbol{x}||_2=1} \boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} = \lambda_1(1)^2 + \sum_{i=2}^{n} \lambda_i(0)^2 = \lambda_1$$

$$\min_{||\boldsymbol{x}||_2=1} \boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} = \sum_{i=1}^{n-1} \lambda_i(0)^2 + \lambda_n(1)^2 = \lambda_n$$

Extension to Singular Values

The Rayleigh quotient also allows us to characterize the singular values of $\boldsymbol{A}$. From our previous definitions, we can write:

$$\lambda_{min}(\boldsymbol{A}^T\boldsymbol{A}) \leq \frac{\boldsymbol{x}^T\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^T\boldsymbol{x}} \leq \lambda_{max}(\boldsymbol{A}^T\boldsymbol{A}), \ \forall \boldsymbol{x} \neq 0$$

We can rearrange these inequalities to obtain the following inequalities:

$$\lambda_{min}(\boldsymbol{A}^T\boldsymbol{A})||\boldsymbol{x}||_2^2 \leq ||\boldsymbol{A}\boldsymbol{x}||_2^2 \leq \lambda_{max}(\boldsymbol{A}^T\boldsymbol{A})||\boldsymbol{x}||_2^2$$

$$\sigma_{min}(\boldsymbol{A})||\boldsymbol{x}||_2 \leq ||\boldsymbol{A}\boldsymbol{x}||_2 \leq \sigma_{max}(\boldsymbol{A})||\boldsymbol{x}||_2$$

These relationships allow us to see how multiplying any n-dimensional vector $\boldsymbol{x}$ by a matrix $\boldsymbol{A}$ changes the length of that vector. From the Rayleigh quotient theorem, we can obtain a couple additional relationships for the singular values.

$$\lambda_{max}(\boldsymbol{A}^T\boldsymbol{A}) = \max_{\boldsymbol{x}:||\boldsymbol{x}||_2=1} \boldsymbol{x}^T\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{x} = \max_{\boldsymbol{x}:||\boldsymbol{x}||_2=1} ||\boldsymbol{A}\boldsymbol{x}||_2^2$$

$$\sigma_{max}(\boldsymbol{A}) = \sqrt{\lambda_{max}(\boldsymbol{A}^T\boldsymbol{A})} = \max_{\boldsymbol{x}:||\boldsymbol{x}||_2=1} ||\boldsymbol{A}\boldsymbol{x}||_2$$

$$\lambda_{min}(\boldsymbol{A}^T\boldsymbol{A}) = \min_{\boldsymbol{x}:||\boldsymbol{x}||_2=1} \boldsymbol{x}^T\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{x} = \min_{\boldsymbol{x}:||\boldsymbol{x}||_2=1} ||\boldsymbol{A}\boldsymbol{x}||_2^2$$

$$\sigma_{min}(\boldsymbol{A}) = \sqrt{\lambda_{min}(\boldsymbol{A}^T\boldsymbol{A})} = \min_{\boldsymbol{x}:||\boldsymbol{x}||_2=1} ||\boldsymbol{A}\boldsymbol{x}||_2$$

8.1.3 Minmax Principle

The **minmax principle** is similar to the Rayleigh quotient theorem but is defined for vector subspaces. Let $\boldsymbol{A}$ be an $n \times n$ symmetric matrix with eigenvalues $\lambda_1 \geq \ldots \geq \lambda_n$ and let V be any k-dimensional subspace of $\mathbb{R}^n$, where $1 \leq k \leq n$. There exist vectors $\boldsymbol{x}, \boldsymbol{y} \in V$ with unit norm $||\boldsymbol{x}||_2 = ||\boldsymbol{y}||_2 = 1$ such that

$$\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} \leq \lambda_k \quad \text{and} \quad \boldsymbol{y}^T\boldsymbol{A}\boldsymbol{y} \geq \lambda_{n-k+1}.$$

Furthermore, the minmax principle says that for $k \in \{1, \ldots, n\}$,

$$\lambda_k = \min_{V:\dim V=k} \max_{\boldsymbol{x}\in V:||\boldsymbol{x}||_2=1} \boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} = \max_{V:\dim V=n-k+1} \min_{\boldsymbol{x}\in V:||\boldsymbol{x}||_2=1} \boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x}.$$

8.2 Positive (Semi)Definite Matrices

8.2.1 Positive and Negative Definiteness

Let $\boldsymbol{A}$ be an $n \times n$ matrix with eigenvalues $\lambda_1, \ldots, \lambda_n$ and components a_{ij} for $i = 1, \ldots, n$ and $j = 1, \ldots, n$. We use $\boldsymbol{A} \succeq 0$ to express that $\boldsymbol{A}$ is positive semidefinite (PSD), $\boldsymbol{A} \succ 0$ to express that $\boldsymbol{A}$ is positive definite (PD), $\boldsymbol{A} \preceq 0$ to express that $\boldsymbol{A}$ is negative semidefinite (NSD), and $\boldsymbol{A} \prec 0$ to express that $\boldsymbol{A}$ is negative definite (ND). We define these characterizations as follows:

1. **Positive Semidefinite (PSD)**
   (a) $\boldsymbol{A} \succeq 0 \iff \boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} \geq 0, \ \forall \boldsymbol{x} \in \mathbb{R}^n$
   (b) $\boldsymbol{A} \succeq 0 \iff \lambda_i \geq 0, \ i = 1, \ldots, n$

(c) $\boldsymbol{A} \succeq 0 \implies a_{ii} \geq 0,\ i = 1, \ldots, n$

(d) $\boldsymbol{A} \succeq 0 \implies \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} \geq 0$

2. **Positive Definite (PD)**

(a) $\boldsymbol{A} \succ 0 \iff \boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} > 0,\ \forall \boldsymbol{x} \in \mathbb{R}^n :\ \boldsymbol{x} \neq \boldsymbol{0}_n$

(b) $\boldsymbol{A} \succ 0 \iff \lambda_i > 0,\ i = 1, \ldots, n$

(c) $\boldsymbol{A} \succ 0 \implies a_{ii} > 0,\ i = 1, \ldots, n$

(d) $\boldsymbol{A} \succ 0 \implies \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} > 0$

3. **Negative Semidefinite (NSD)**

(a) $\boldsymbol{A} \preceq 0 \iff \boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} \leq 0,\ \forall \boldsymbol{x} \in \mathbb{R}^n$

(b) $\boldsymbol{A} \preceq 0 \iff \lambda_i \leq 0,\ i = 1, \ldots, n$

(c) $\boldsymbol{A} \preceq 0 \implies a_{ii} \leq 0,\ i = 1, \ldots, n$

(d) $\boldsymbol{A} \preceq 0 \implies \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} \leq 0$

4. **Negative Definite (ND)**

(a) $\boldsymbol{A} \prec 0 \iff \boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} < 0 :\ \forall \boldsymbol{x} \in \mathbb{R}^n,\ \boldsymbol{x} \neq \boldsymbol{0}_n$

(b) $\boldsymbol{A} \prec 0 \iff \lambda_i < 0,\ i = 1, \ldots, n$

(c) $\boldsymbol{A} \prec 0 \implies a_{ii} < 0,\ i = 1, \ldots, n$

(d) $\boldsymbol{A} \prec 0 \implies \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} < 0$

From these characterizations, we can see the following relationships:

1. $\boldsymbol{A} \preceq 0 \iff -\boldsymbol{A} \succeq 0$,
2. $\boldsymbol{A} \prec 0 \iff -\boldsymbol{A} \succ 0$, and
3. $\boldsymbol{A} \succ 0 \iff \boldsymbol{A} \succeq 0$ and $\boldsymbol{A}$ is invertible.

8.2.2 Symmetric Positive (Semi)definite Matrices

We denote the set of symmetric positive semidefinite (PSD) matrices in $\mathbb{R}^{n \times n}$

$$\mathbb{S}_{+}^{n} = \{\boldsymbol{A} \in \mathbb{S}^n : \boldsymbol{A} \succeq 0\}.$$

We denote the set of symmetric positive definite (PD) matrices in $\mathbb{R}^{n \times n}$

$$\mathbb{S}_{++}^{n} = \{\boldsymbol{A} \in \mathbb{S}^n : \boldsymbol{A} \succ 0\}.$$

8.2.3 Sylvester's Criterion

Sylvester's criterion says that a matrix $\boldsymbol{A} \in \mathbb{S}^n$ is positive semidefinite if and only if all its principal minors are non-negative. Similarly, it says that $\boldsymbol{A} \in \mathbb{S}^n$ is positive definite if and only if its leading principal minors are strictly positive.

Principal minors and **leading principal minors** are easiest to understand for the case of 3×3 matrices. A 3×3 matrix $\boldsymbol{A} \in \mathbb{S}^3$, has seven total principal minors. Because $\boldsymbol{A}$ is symmetric, I will denote its components such that

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{12} & a_{22} & a_{23} \\ a_{13} & a_{23} & a_{33} \end{bmatrix}.$$

The principal minors of the matrix $\boldsymbol{A} \in \mathbb{S}^3$ are given by:

1. Three principal one-minors
   (a) $m_1 = a_{11}$
   (b) $m_2 = a_{22}$
   (c) $m_3 = a_{33}$
2. Three principal two-minors
   (a) $m_{12} = \det \begin{bmatrix} a_{11} & a_{12} \\ a_{12} & a_{22} \end{bmatrix} = a_{11}a_{22} - a_{12}^2$
   (b) $m_{13} = \det \begin{bmatrix} a_{11} & a_{13} \\ a_{13} & a_{33} \end{bmatrix} = a_{11}a_{33} - a_{13}^2$
   (c) $m_{23} = \det \begin{bmatrix} a_{22} & a_{23} \\ a_{23} & a_{33} \end{bmatrix} = a_{22}a_{33} - a_{23}^2$
3. One principal three minor
   (a) $m_{123} = \det(\boldsymbol{A})$

Its leading principal minors are m_1, m_{12}, and m_{123}.

8.2.4 Schur Complements

Let $\boldsymbol{A} \in \mathbb{S}^n$ and $\boldsymbol{B} \in \mathbb{S}^m$ and consider the block diagonal matrix

$$\boldsymbol{M} = \begin{bmatrix} \boldsymbol{A} & \boldsymbol{0}^{n \times m} \\ \boldsymbol{0}^{m \times n} & \boldsymbol{B} \end{bmatrix}.$$

We can write the following two implications for this matrix:

$$\boldsymbol{M} \succeq 0 \iff \boldsymbol{A} \succeq 0,\ \boldsymbol{B} \succeq 0$$

$$\boldsymbol{M} \succ 0 \iff \boldsymbol{A} \succ 0,\ \boldsymbol{B} \succ 0$$

Now let $\boldsymbol{A} \in \mathbb{S}^n$, $\boldsymbol{B} \in \mathbb{S}^m$, and $\boldsymbol{X} \in \mathbb{R}^{n \times m}$ and consider the block matrix

$$\boldsymbol{M} = \begin{bmatrix} \boldsymbol{A} & \boldsymbol{X} \\ \boldsymbol{X}^T & \boldsymbol{B} \end{bmatrix}.$$

For this block matrix, we can write similar relationships among the blocks:

$$\boldsymbol{M} \succeq 0 \implies \boldsymbol{A} \succeq 0, \ \boldsymbol{B} \succeq 0$$

$$\boldsymbol{M} \succ 0 \implies \boldsymbol{A} \succ 0, \ \boldsymbol{B} \succ 0$$

$$\text{If } \boldsymbol{A} = 0, \text{ then } \boldsymbol{M} \succeq 0 \iff \boldsymbol{B} \succeq 0, \ \boldsymbol{X} = 0$$

$$\text{If } \boldsymbol{B} = 0, \text{ then } \boldsymbol{M} \succeq 0 \iff \boldsymbol{A} \succeq 0, \ \boldsymbol{X} = 0$$

If we assume $\boldsymbol{B} \succ 0$, then the **Schur complement** matrix of $\boldsymbol{A}$ in $\boldsymbol{M}$ is

$$S_{\boldsymbol{A}|\boldsymbol{M}} = \boldsymbol{A} - \boldsymbol{X}\boldsymbol{B}^{-1}\boldsymbol{X}^T.$$

Now we can express the following relationships for the symmetric block matrix:

$$\boldsymbol{M} \succeq 0 \iff S_{\boldsymbol{A}|\boldsymbol{M}} \succeq 0$$

$$\boldsymbol{M} \succ 0 \iff S_{\boldsymbol{A}|\boldsymbol{M}} \succ 0$$

If we assume $\boldsymbol{A} \succ 0$, then the **Schur complement** matrix of $\boldsymbol{B}$ in $\boldsymbol{M}$ is

$$S_{\boldsymbol{B}|\boldsymbol{M}} = \boldsymbol{B} - \boldsymbol{X}^T\boldsymbol{A}^{-1}\boldsymbol{X}.$$

Now we can express the following relationships for the symmetric block matrix:

$$\boldsymbol{M} \succeq 0 \iff S_{\boldsymbol{B}|\boldsymbol{M}} \succeq 0$$

$$\boldsymbol{M} \succ 0 \iff S_{\boldsymbol{B}|\boldsymbol{M}} \succ 0$$

If we cannot assume that $\boldsymbol{A}$ or $\boldsymbol{B}$ is positive definite, then we can still say that

$$\boldsymbol{M} \succeq 0 \iff \boldsymbol{A} - \boldsymbol{X}^T\boldsymbol{B}^\dagger\boldsymbol{X}^T \succeq 0, \ R(\boldsymbol{X}^T) \subseteq R(\boldsymbol{B})$$

$$\boldsymbol{M} \succeq 0 \iff \boldsymbol{B} - \boldsymbol{X}^T\boldsymbol{A}^\dagger\boldsymbol{X} \succeq 0, \ R(\boldsymbol{X}) \subseteq R(\boldsymbol{A})$$

8.2.5 Congruence Transformations

Let $\boldsymbol{A} \in \mathbb{S}^n$ and $\boldsymbol{B} \in \mathbb{R}^{m \times n}$. Consider the product $\boldsymbol{C} = \boldsymbol{B}\boldsymbol{A}\boldsymbol{B}^T \in \mathbb{S}^m$. We can say the following about this congruence transformation:

1. If $\boldsymbol{A} \succeq 0$, then $\boldsymbol{C} \succeq 0$.
2. If $\boldsymbol{A} \succ 0$, then $\boldsymbol{C} \succ 0$ if and only if $\text{rank}(\boldsymbol{B}) = m$.
3. If $\boldsymbol{B}$ is square and invertible, then $\boldsymbol{A} \succeq 0$ if and only if $\boldsymbol{C} \succeq 0$.
4. If $\boldsymbol{B}$ is square and invertible, then $\boldsymbol{A} \succ 0$ if and only if $\boldsymbol{C} \succ 0$.

8.2.6 Summations

Let $\boldsymbol{S}$ be the sum of two matrices $\boldsymbol{A} \in \mathbb{S}^n$ and $\boldsymbol{B} \in \mathbb{S}^n$ such that $\boldsymbol{S} = \boldsymbol{A} + \boldsymbol{B}$. We can deduce the following about the matrix $\boldsymbol{S}$:

1. $\boldsymbol{A} \succ 0,\ \boldsymbol{B} \succ 0 \implies \boldsymbol{S} \succ 0$
2. $\boldsymbol{A} \succeq 0,\ \boldsymbol{B} \succeq 0 \implies \boldsymbol{S} \succeq 0$
3. $\boldsymbol{A} \succ 0,\ \boldsymbol{B} \succeq 0 \implies \boldsymbol{S} \succ 0$

Now let $\boldsymbol{A} \in \mathbb{S}^n_{++}$ and $\boldsymbol{B} \in \mathbb{S}^n$. Consider the symmetric sum $\boldsymbol{S} = \boldsymbol{AB} + \boldsymbol{BA}$. Now the matrix $\boldsymbol{S}$ can tell us the following information about $\boldsymbol{B}$:

1. $\boldsymbol{S} \succ 0 \implies \boldsymbol{B} \succ 0$
2. $\boldsymbol{S} \succeq 0 \implies \boldsymbol{B} \succeq 0$

8.2.7 Matrix Product

For any matrix $\boldsymbol{U} \in \mathbb{R}^{m \times n}$, we can say the following about its matrix products:

1. $\boldsymbol{U}^T\boldsymbol{U} \succeq 0$
2. $\boldsymbol{U}\boldsymbol{U}^T \succeq 0$
3. $\boldsymbol{U}^T\boldsymbol{U} \succ 0$ if and only if $\boldsymbol{U}$ is full column rank (i.e. $\text{rank}(\boldsymbol{U}) = n$)
4. $\boldsymbol{U}\boldsymbol{U}^T \succ 0$ if and only if $\boldsymbol{U}$ is full row rank (i.e. $\text{rank}(\boldsymbol{U}) = m$)

Additionally, for any symmetric positive semidefinite matrix $\boldsymbol{A} \in \mathbb{S}^n_+$ with rank r, there exists a matrix $\boldsymbol{U} \in \mathbb{R}^{r \times n}$ such that $\boldsymbol{A} = \boldsymbol{U}^T\boldsymbol{U}$. This also implies that for any symmetric positive definite matrix $\boldsymbol{A} \in \mathbb{S}^n_{++}$, there exists an invertible matrix $\boldsymbol{U} \in \mathbb{R}^{n \times n}$ such that $\boldsymbol{A} = \boldsymbol{U}^T\boldsymbol{U}$.

8.2.8 Matrix Square Root

For any symmetric matrix $\boldsymbol{A} \in \mathbb{S}^n$, we can write the following:

$$\boldsymbol{A} \succeq 0 \iff \exists \boldsymbol{B} \succeq 0 : \boldsymbol{A} = \boldsymbol{B}^2$$

$$\boldsymbol{A} \succ 0 \iff \exists \boldsymbol{B} \succ 0 : \boldsymbol{A} = \boldsymbol{B}^2$$

The matrix $\boldsymbol{B}$ is unique and called the matrix square root of $\boldsymbol{A}$. We sometimes denote $\boldsymbol{B}$ as $\boldsymbol{B} = \boldsymbol{A}^{1/2}$. Recall that the spectral decomposition theorem says that any symmetric matrix $\boldsymbol{A}$ can be expressed as $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^T$, where the columns of $\boldsymbol{U}$ are the eigenvectors of $\boldsymbol{A}$ and $\boldsymbol{\Lambda} = \text{diag}(\lambda_1, \ldots, \lambda_n)$. The matrix square root is given by $\boldsymbol{B} = \boldsymbol{U}\boldsymbol{\Lambda}^{1/2}\boldsymbol{U}^T$, where $\boldsymbol{\Lambda}^{1/2} = \text{diag}(\sqrt{\lambda_1}, \ldots, \sqrt{\lambda_n})$.

8.2.9 Partial Order

If $(\boldsymbol{A} - \boldsymbol{B}) \succeq 0$, then we can write $\boldsymbol{A} \succeq \boldsymbol{B}$. Similarly, if $(\boldsymbol{A} - \boldsymbol{B}) \succ 0$, then we can write $\boldsymbol{A} \succ \boldsymbol{B}$. We call the relationships $\boldsymbol{A} \succeq \boldsymbol{B}$ and $\boldsymbol{A} \succ \boldsymbol{B}$ partial orders. When these relationships hold, we can say the following about $\boldsymbol{A}$ and $\boldsymbol{B}$:

$$\boldsymbol{A} \succeq \boldsymbol{B} \iff \boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{x}^T \boldsymbol{B} \boldsymbol{x}, \ \forall \boldsymbol{x} \in \mathbb{R}^n$$

$$\boldsymbol{A} \succ \boldsymbol{B} \iff \boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} > \boldsymbol{x}^T \boldsymbol{B} x, \ \forall \boldsymbol{x} \neq \boldsymbol{0}_n$$

Based on these partial order relationships, we can also notice that

$$\max_{\boldsymbol{B}:\boldsymbol{B} \preceq \boldsymbol{A}} \boldsymbol{x}^T \boldsymbol{B} \boldsymbol{x} = \boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} \quad \text{and} \quad \min_{\boldsymbol{B}:\boldsymbol{B} \succeq \boldsymbol{A}} \boldsymbol{x}^T \boldsymbol{B} \boldsymbol{x} = \boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x}.$$

From the partial order relationships given, we can also say that if $\boldsymbol{A} \succeq \boldsymbol{B}$ and both $\boldsymbol{A}$ and $\boldsymbol{B}$ are positive definite, then $\boldsymbol{B}^{-1} \succeq \boldsymbol{A}^{-1}$.

8.2.10 Hadamard's Inequality

For a matrix $\boldsymbol{A} \in \mathbb{S}_+^n$ with diagonal elements a_{ii}, Hadamard's inequality says

$$\det(\boldsymbol{A}) \leq \prod_{i=1}^{n} a_{ii}.$$

Chapter 9

Matrix Inner Product & Norms

9.1 Matrix Inner Product

The inner product of two $m \times n$ matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ is defined as

$$\langle \boldsymbol{A}, \boldsymbol{B} \rangle = \text{tr}(\boldsymbol{A}^T \boldsymbol{B}) = \sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij} b_{ij}.$$

From the properties of the trace, we could equivalently write this definition as

$$\langle \boldsymbol{A}, \boldsymbol{B} \rangle = \text{tr}(\boldsymbol{A}^T \boldsymbol{B}) = \text{tr}(\boldsymbol{B}^T \boldsymbol{A}) = \text{tr}(\boldsymbol{A}\boldsymbol{B}^T) = \text{tr}(\boldsymbol{B}\boldsymbol{A}^T).$$

9.2 Matrix Norms

9.2.1 Properties of Matrix Norms

In order for a matrix function to be considered a norm, it must satisfy the following three properties for any two matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ and any scalar α:

1. $||\boldsymbol{A}|| \geq 0 \ \forall \boldsymbol{A} \in \mathbb{F}^{m \times n}$
   $||\boldsymbol{A}|| = 0 \iff \boldsymbol{A} = \boldsymbol{0}^{m \times n}$
2. $||\boldsymbol{A} + \boldsymbol{B}|| \leq ||\boldsymbol{A}|| + ||\boldsymbol{B}||, \ \forall \boldsymbol{A}, \boldsymbol{B} \in \mathbb{F}^{m \times n}$
3. $||\alpha \boldsymbol{A}|| = |\alpha| \ ||\boldsymbol{A}||, \ \forall \alpha \in \mathbb{F}, \ \boldsymbol{A} \in \mathbb{F}^{m \times n}$

Often, matrix norms also satisfy the **sub-multiplicativity** property:

$$||\boldsymbol{A}\boldsymbol{B}|| \leq ||\boldsymbol{A}|| \cdot ||\boldsymbol{B}||$$

However, this is not a required property for a valid matrix norm.

9.2.2 Frobenius Norm

Suppose $\boldsymbol{A}$ is an $m \times n$ matrix whose ijth element is a_{ij} and whose ith row vector is $\boldsymbol{a_i}^T$. The **Frobenius norm** is denoted $||\cdot||_F$ and defined as such that

$$||\boldsymbol{A}||_F = \left(\sum_{i=1}^{m}\sum_{j=1}^{n}|a_{ij}|^2\right)^{1/2} = \left(\sum_{i=1}^{m}||\boldsymbol{a_i}^T||_2^2\right)^{1/2}.$$

If the rank of $\boldsymbol{A}$ is r, we can also express the Frobenius norm as

$$||\boldsymbol{A}||_F = \sqrt{\text{tr}(\boldsymbol{A}\boldsymbol{A}^*)} = \left(\sum_{i=1}^{m}\lambda_i(\boldsymbol{A}\boldsymbol{A}^*)\right)^{1/2} = \left(\sum_{i=1}^{r}\sigma_i(\boldsymbol{A})^2\right)^{1/2}.$$

Equivalenty, we could express the Frobenius norm of $\boldsymbol{A}$ as

$$||\boldsymbol{A}||_F = \sqrt{\text{tr}(\boldsymbol{A}^*\boldsymbol{A})} = \left(\sum_{i=1}^{n}\lambda_i(\boldsymbol{A}^*\boldsymbol{A})\right)^{1/2} = \left(\sum_{i=1}^{r}\sigma_i(\boldsymbol{A})^2\right)^{1/2}.$$

Properties of the Frobenious Norm

The Frobenius norm satisfies the sub-multipicativity property, which means that for any $m \times n$ matrix $\boldsymbol{A}$ and any $n \times p$ matrix $\boldsymbol{B}$,

$$||\boldsymbol{A}\boldsymbol{B}||_F \leq ||\boldsymbol{A}||_F||\boldsymbol{B}||_F.$$

Another interesting property of the Frobenius norm is that for any $m \times n$ matrix $\boldsymbol{A}$ and n-dimensional vector $\boldsymbol{x}$, the following inequality holds:

$$||\boldsymbol{A}\boldsymbol{x}||_2 \leq ||\boldsymbol{A}||_F||\boldsymbol{x}||_2.$$

One final property is that if $\boldsymbol{U}$ and $\boldsymbol{V}$ are two orthogonal matrices, then

$$||\boldsymbol{U}\boldsymbol{A}\boldsymbol{V}||_F = ||\boldsymbol{A}||_F.$$

9.2.3 Induced Norms

For $p \in [1, \infty)$, the **induced matrix norms** have the general form:

$$||\boldsymbol{A}||_p = \sup_{\boldsymbol{x}\neq\boldsymbol{0}_n}\frac{||\boldsymbol{A}\boldsymbol{x}||_p}{|||\boldsymbol{x}||_p} = \sup_{||\boldsymbol{x}||_p=1}||\boldsymbol{A}\boldsymbol{x}||_p.$$

Properties of the Induced Norm

The induced norm satisfies the sub-multipicativity property, which means that for any $m \times n$ matrix $\boldsymbol{A}$ and any $n \times p$ matrix $\boldsymbol{B}$,

$$||\boldsymbol{A}\boldsymbol{B}||_p \leq ||\boldsymbol{A}||_p||\boldsymbol{B}||_p.$$

Based on how we defined the induced norm, we can also deduce that for any $m \times n$ matrix $\boldsymbol{A}$ and any n-dimensional vector $\boldsymbol{x}$,

$$||\boldsymbol{A}\boldsymbol{x}||_p \leq ||\boldsymbol{A}||_p||\boldsymbol{x}||_p.$$

9.2.4 l_2 Induced Norm

The $\boldsymbol{l_2}$ **induced norm** is referred to as the **spectral norm** or the maximum singular value. If $\boldsymbol{A}$ is an $m \times n$ matrix of rank r, its l_2 induced norm is

$$||\boldsymbol{A}||_2 = \max_{||\boldsymbol{x}||_2=1} ||\boldsymbol{A}\boldsymbol{x}||_2 = \sigma_{max}(\boldsymbol{A}) := \max_{j\in[1,r]} \sigma_j(\boldsymbol{A}).$$

Proof: When discussing the extension of the Rayleigh quotient theorem to singular values in Section 8.1.2, we showed that

$$\sigma_{max}(\boldsymbol{A}) = \sqrt{\lambda_{max}(\boldsymbol{A}^T\boldsymbol{A})} = \max_{\boldsymbol{x}:||\boldsymbol{x}||_2=1} ||\boldsymbol{A}\boldsymbol{x}||_2.$$

Therefore, the l_2 induced norm of $\boldsymbol{A}$ is the largest singular value of $\boldsymbol{A}$.

9.2.5 l_1 Induced Norm

The $\boldsymbol{l_1}$ **induced norm** is referred to as the largest absolute column sum. If $\boldsymbol{A}$ is an $m \times n$ matrix whose columns are $a_1, \ldots, a_n$, its l_1 induced norm is

$$||\boldsymbol{A}||_1 = \max_{||\boldsymbol{x}||_1=1} ||\boldsymbol{A}\boldsymbol{x}||_1 = \max_{j\in[1,n]} ||a_j||_1 = \max_{j\in[1,n]} \sum_{i=1}^{m} |a_{ij}|.$$

Proof: To see why this holds, notice from the definition of the l_1 norm,

$$||\boldsymbol{A}\boldsymbol{x}||_1 = \sum_{i=1}^{m} |(\boldsymbol{A}\boldsymbol{x})_i| = \sum_{i=1}^{m} \left| \sum_{j=1}^{n} a_{ij}x_j \right| \leq \sum_{i=1}^{m}\sum_{j=1}^{n} |a_{ij}x_j| \leq \sum_{i=1}^{m}\sum_{j=1}^{n} |a_{ij}||x_j|.$$

Under the constraint $||\boldsymbol{x}||_1 = 1$, the absolute value of all of the elements of $\boldsymbol{x}$ must sum to one. Given this constraint, the maximum value of the upper bound on $||\boldsymbol{A}\boldsymbol{x}||_1$ is attained for the vector $\hat{\boldsymbol{x}} = e_{\hat{j}}$, where $\hat{j}$ is defined as

$$\hat{j} = \arg\max_{j\in[1,n]} \sum_{i=1}^{m} |a_{ij}|.$$

Notice that for this choice of $\hat{\boldsymbol{x}}$, we have

$$||\boldsymbol{A}\hat{\boldsymbol{x}}||_1 = \sum_{i=1}^{m} \left| \sum_{j=1}^{n} a_{ij}x_j \right| = \sum_{i=1}^{m} |a_{i\hat{j}}| = \max_{j\in[1,n]} \sum_{i=1}^{m} |a_{ij}|.$$

Therefore, the maximum value of $||\boldsymbol{A}\boldsymbol{x}||_1$ under the constraint $||\boldsymbol{x}||_1 = 1$ (i.e. the l_1 induced norm of $\boldsymbol{A}$) is the largest absolute row column of $\boldsymbol{A}$.

9.2.6 l_∞ Induced Norm

The $\boldsymbol{l_\infty}$ **induced norm** is referred to as the largest absolute row sum. If $\boldsymbol{A}$ is an $m \times n$ matrix whose rows are $a_1^T, \ldots, a_m^T$, its l_∞ induced norm is

$$||\boldsymbol{A}||_\infty = \max_{||\boldsymbol{x}||_\infty = 1} ||\boldsymbol{A}\boldsymbol{x}||_\infty = \max_{i \in [1,m]} ||a_i^T||_1 = \max_{i \in [1,m]} \sum_{j=1}^{n} |a_{ij}|.$$

Proof: To see why this holds, notice from the definition of the l_∞ norm,

$$\begin{aligned} ||\boldsymbol{A}\boldsymbol{x}||_\infty = \max_{i \in [1,m]} |(\boldsymbol{A}\boldsymbol{x})_i| = \max_{i \in [1,m]} \left| \sum_{j=1}^{n} a_{ij} x_j \right| &\leq \max_{i \in [1,m]} \sum_{j=1}^{n} |a_{ij}||x_j| \\ \leq \max_{i \in [1,m]} \sum_{j=1}^{n} |a_{ij}| \max_{k \in [1,n]} |x_k| &= \max_{i \in [1,m]} \sum_{j=1}^{n} |a_{ij}| ||\boldsymbol{x}||_\infty. \end{aligned}$$

Under the constraint $||\boldsymbol{x}||_\infty = 1$, we can clearly see that the maximum value of the upper bound on $||\boldsymbol{A}\boldsymbol{x}||_\infty$ is the largest absolute row sum of $\boldsymbol{A}$. Now it remains to show that we can find a vector $\boldsymbol{x}$, satisfying the constraint on the l_∞ norm, that attains this upper bound. Consider the vector $\hat{\boldsymbol{x}}$ whose elements are $\hat{x}_j = \text{sign}(a_{\hat{i}j})$, where $\hat{i}$ is defined such that

$$\hat{i} = \arg\max_{i \in [1,m]} \sum_{j=1}^{n} |a_{ij}|.$$

Notice that for this choice of $\hat{\boldsymbol{x}}$, we have

$$||\boldsymbol{A}\hat{\boldsymbol{x}}||_\infty = \max_{i \in [1,m]} \left| \sum_{j=1}^{n} a_{ij} \hat{x}_j \right|, \text{ where}$$

$$\left| \sum_{j=1}^{n} a_{ij} \hat{x}_j \right| = \begin{cases} \left| \sum_{j=1}^{n} a_{\hat{i}j} \text{sign}(a_{\hat{i}j}) \right| = \sum_{j=1}^{n} |a_{\hat{i}j}| & \text{if } i = \hat{i} \\ \left| \sum_{j=1}^{n} a_{ij} \text{sign}(a_{\hat{i}j}) \right| \leq \sum_{j=1}^{n} |a_{\hat{i}j}| & \text{if } i \neq \hat{i} \end{cases}.$$

Now we can see that for this choice of $\hat{\boldsymbol{x}}$, we can conclude that

$$||\boldsymbol{A}\hat{\boldsymbol{x}}||_\infty = \max_{i \in [1,m]} \sum_{j=1}^{n} |a_{ij}|.$$

Therefore, the maximum value of $||\boldsymbol{A}\boldsymbol{x}||_\infty$ under the constraint $||\boldsymbol{x}||_\infty = 1$ (i.e. the l_∞ induced norm of $\boldsymbol{A}$) is the largest absolute row sum of $\boldsymbol{A}$.

9.3 Condition Number

The **condition number** of a matrix $\boldsymbol{A} \in \mathbb{F}^{m \times n}$ is the ratio between the largest and smallest singular values of that matrix. If we assume that $\boldsymbol{A}$ has r singular values $\sigma_1 \geq \ldots \geq \sigma_r > 0$, then we can express its condition number as

$$K(\boldsymbol{A}) = \frac{\sigma_1}{\sigma_r} = \frac{\sup_{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{x} \neq \boldsymbol{0}_n} \frac{||\boldsymbol{A}\boldsymbol{x}||_2}{||\boldsymbol{x}||_2}}{\inf_{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{x} \neq \boldsymbol{0}_n} \frac{||\boldsymbol{A}\boldsymbol{x}||_2}{||\boldsymbol{x}||_2}} = \frac{\sup_{\boldsymbol{x} \in \mathbb{R}^n : ||\boldsymbol{x}||_2 = 1} ||\boldsymbol{A}\boldsymbol{x}||_2}{\inf_{\boldsymbol{x} \in \mathbb{R}^n : ||\boldsymbol{x}||_2 = 1} ||\boldsymbol{A}\boldsymbol{x}||_2}.$$

9.3.1 Condition Number of Invertible Matrices

Now suppose that $\boldsymbol{A}$ is a square $n \times n$ matrix. If $\boldsymbol{A}$ is an invertible matrix with strictly positive singular values $\sigma_1 \geq \ldots \geq \sigma_n > 0$, then we can equivalently express its largest and smallest singular values as

$$\sigma_1 = ||\boldsymbol{A}||_2 \quad \text{and} \quad \sigma_n = \frac{1}{||\boldsymbol{A}^{-1}||_2}.$$

This then allows us to express the condition number as

$$K(\boldsymbol{A}) = \frac{\sigma_1}{\sigma_n} = ||\boldsymbol{A}||_2 ||\boldsymbol{A}^{-1}||_2.$$

For an invertible matrix $\boldsymbol{A}$, the condition number gives a quantitative measure of how close $\boldsymbol{A}$ is to being singular. The larger the value of the condition number $K(\boldsymbol{A})$, the closer $\boldsymbol{A}$ is to being singular.

9.3.2 Subspace Condition Number

Given a subspace $V \subseteq \mathbb{R}^n$, the subspace condition number is defined as

$$K_V(\boldsymbol{A}) = \frac{\sup_{\boldsymbol{x} \in V : \boldsymbol{x} \neq \boldsymbol{0}_n} \frac{||\boldsymbol{A}\boldsymbol{x}||_2}{||\boldsymbol{x}||_2}}{\inf_{\boldsymbol{x} \in V : \boldsymbol{x} \neq \boldsymbol{0}_n} \frac{||\boldsymbol{A}\boldsymbol{x}||_2}{||\boldsymbol{x}||_2}} = \frac{\sup_{\boldsymbol{x} \in V : ||\boldsymbol{x}||_2 = 1} ||\boldsymbol{A}\boldsymbol{x}||_2}{\inf_{\boldsymbol{x} \in V : ||\boldsymbol{x}||_2 = 1} ||\boldsymbol{A}\boldsymbol{x}||_2}.$$

Note that if the singular value decomposition of $\boldsymbol{A}$ is $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^*$ and we define the subspace $W = \{\boldsymbol{V}^*\boldsymbol{x} : \boldsymbol{x} \in V\}$, then $K_V(\boldsymbol{A}) = K_W(\boldsymbol{\Sigma})$.

9.4 Eckart-Young-Mirsky Theorem

Suppose we want to find a rank k matrix that best approximates the matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ with rank $r > k$. The optimal solution $\hat{\boldsymbol{A}}$ is given by

$$\begin{aligned} \hat{\boldsymbol{A}} = \underset{\boldsymbol{B} \in \mathbb{R}^{m \times n}}{\arg\min} \; & ||\boldsymbol{A} - \boldsymbol{B}||_F \\ \text{s.t.} \quad & \text{rank}(\boldsymbol{B}) \leq k \end{aligned}$$

Suppose that $\boldsymbol{A}$ admits the singular value decomposition $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T$, where $\boldsymbol{U} \in \mathbb{R}^{m\times m}$, $\boldsymbol{\Sigma} \in \mathbb{R}^{m\times n}$, and $\boldsymbol{V} \in \mathbb{R}^{n\times n}$. We can express these matrices as

$$\begin{aligned}
\boldsymbol{U} &= \begin{bmatrix} \boldsymbol{U}_1 & \boldsymbol{U}_2 \end{bmatrix} \quad \text{where} \quad \boldsymbol{U}_1 \in \mathbb{R}^{m\times k},\ \boldsymbol{U}_2 \in \mathbb{R}^{m\times(m-k)} \\
\boldsymbol{\Sigma} &= \begin{bmatrix} \boldsymbol{\Sigma_1} & 0 \\ 0 & \boldsymbol{\Sigma_2} \end{bmatrix} \quad \text{where} \quad \boldsymbol{\Sigma_1} \in \mathbb{R}^{k\times k},\ \boldsymbol{\Sigma_2} \in \mathbb{R}^{(m-k)\times(n-k)} \\
\boldsymbol{V} &= \begin{bmatrix} \boldsymbol{V_1} & \boldsymbol{V_2} \end{bmatrix} \quad \text{where} \quad \boldsymbol{V_1} \in \mathbb{R}^{n\times k},\ \boldsymbol{V_2} \in \mathbb{R}^{n\times(n-k)}
\end{aligned}$$

The Eckhart-Young-Mirsky theorem says that the optimal solution $\hat{\boldsymbol{A}}$ satisfies

$$\hat{\boldsymbol{A}} = \boldsymbol{U}_1\boldsymbol{\Sigma_1}\boldsymbol{V_1}^T$$

Moreover, this solution is unique if and only if $\sigma_{k+1} \neq \sigma_k$, assuming $\sigma_1 \geq \ldots \geq \sigma_r > 0$ are the non-zero singular values of $\boldsymbol{A}$.

Proof: `https://en.wikipedia.org/wiki/Low-rank_approximation`

Part III

Functions and Maps

Chapter 10

Functions

10.1 Domain and Range

Given two sets X and Y, there is some **function** f that maps X into Y, which we express as $f : X \to Y$. The set X is the input space, known as the **domain**, and Y is the output space, known as the **codomain**. For all values $x \in X$, a valid function f must assign a unique value $y = f(x) \in Y$. If f maps some $x \in X$ to more than one value in Y, as in Figure 10.1, f is not a valid function.

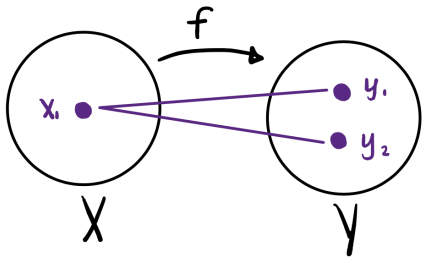

Figure 10.1: The mapping f is not a function because it maps the point x_1 in the domain, X, to both y_1 and y_2 in the codomain, Y.

The **range** of f is the set of all values in the codomain to which the function f can map values in the domain. We express the range as the set $\{f(x) : x \in X\}$. The relationship between the codomain and range is shown in figure 10.2.

10.2 Graphs and Level Sets

The **graph** of a function f with domain X and codomain Y is the set of input-output pairs the function can attain and is expressed as

$$\text{graph}(f) = \{(x, y) \in X \times Y : y = f(x)\}.$$

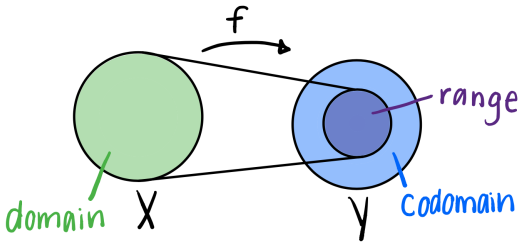

Figure 10.2: The function f maps points in X to points in Y. The green region, X, is the domain of f, the blue region, Y, is the codomain, and the purple region is the range.

The **epigraph** of a function f is the set of input-output pairs the function can attain and anything above these values. The epigraph of f can be expressed as

$$\text{epi}(f) = \{(x, y) \in X \times Y : y \geq f(x)\}.$$

A depiction of the graph and epigraph of a function is shown in figure 10.3.

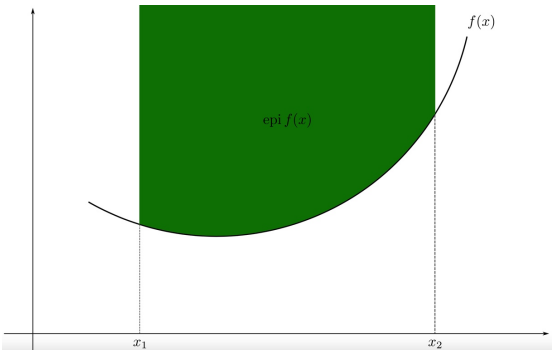

Figure 10.3: The graph of the function, f, is the black line and the epigraph is the green region.

The **level set**, or **contour line**, of a function is the set of points that achieve exactly some value for a function. The c-level set of the function f is given by

$$L_c(f) = \{x \in X : f(x) = c\}.$$

The **sublevel set** of a function is the set of points that achieve at most a certain value of the function. The c-sublevel set of the function f is given by

$$L_c^-(f) = \{x \in X : f(x) \leq c\}.$$

The **superlevel set** of a function is the set of points that achieve above a certain value of the function. The c-superlevel set of the function f is given by

$$L_c^+(f) = \{x \in X : f(x) \geq c\}.$$

10.3 Injectivity and Surjectivity

A few important properties of a function are injectivity, surjectivity, and bijectivity. A function f is **injective**, or **one-to-one**, if and only if

$$x_1 \neq x_2 \implies f(x_1) \neq f(x_2), \ \forall x_1, x_2 \in X.$$

This means that if there exist two distinct values $x_1, x_2 \in X$ that map to the same value in the codomain, then f is not injective. A function f is **surjective**, or **onto**, if and only if its range is equal to the codomain Y, meaning that

$$\exists \ x \in X \text{ s.t. } y = f(x), \ \forall y \in Y.$$

A function f is **bijective** if and only if f is both injective and surjective, meaning

$$\exists! \ x \in X \text{ s.t. } y = f(x), \ \forall y \in Y.$$

Figure 10.4 demonstrates the concept of injectivity, surjectivity, and bijectivity.

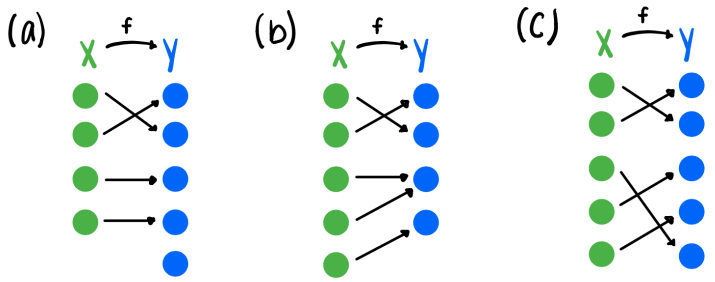

Figure 10.4: (a) f is injective but not surjective. (b) f is sujective but not injective. (c) f is bijective (both injective and surjective).

10.4 Function Inner Product & Norms

10.4.1 Function Inner Product

Assume f and g are two square, integrable, $\mathbb{F}^n$-valued functions defined on $[t_0, t_1]$. The inner product of two functions is commonly defined such that

$$\langle f, g \rangle = \int_{t_0}^{t_1} f(t)^* g(t) dt,$$

where the $f(t)^*$ operation is the complex conjugate transpose of $f(t)$. With this definition of the inner product, the **energy** of f is defined as

$$E_f = \langle f, f \rangle = \int_{t_0}^{t_1} f(t)^* f(t) dt = \int_{t_0}^{t_1} ||f(t)||_2^2 dt.$$

10.4.2 Function Norm

Suppose f is a continuous, $\mathbb{F}^n$-valued function defined over the domain $[t_0, t_1]$. The $\boldsymbol{L_1}$ **norm** of f is defined such that

$$||f||_1 = \int_{t_0}^{t_1} ||f(t)||dt.$$

Similarly, $\boldsymbol{L_2}$ **norm** of f is defined such that

$$||f||_2 = \left(\int_{t_0}^{t_1} ||f(t)||^2 dt \right)^{1/2}.$$

Similarly, $\boldsymbol{L_\infty}$ **norm** of f is defined such that

$$||f||_\infty = \max_{t \in [t_0, t_1]} ||f(t)||.$$

Note that for all of these function norms, we can use any vector norm and simply specify which vector norm we are using to define the function norm.

10.5 Derivative Operators

10.5.1 Gradients

Consider a function $f : \mathbb{F}^n \to \mathbb{F}$ that maps vectors to scalars. For any vector $\boldsymbol{x} \in \mathbb{F}^n$ with elements $x_1, \ldots, x_n$, the **gradient** of f is defined such that

$$\nabla_{\boldsymbol{x}} f(\boldsymbol{x}) = \frac{\partial f}{\partial \boldsymbol{x}} = \begin{bmatrix} \frac{\partial f}{\partial x_1} \\ \vdots \\ \frac{\partial f}{\partial x_n} \end{bmatrix} \in \mathbb{F}^n.$$

Properties of the Gradient

Consider two functions $f : \mathbb{F}^n \to \mathbb{F}$ and $g : \mathbb{F}^n \to \mathbb{F}$ that map vectors to scalars. The gradient operation exhibits the property of linearity, which says that

$$\nabla_{\boldsymbol{x}} \Big(\alpha f(\boldsymbol{x}) + \beta g(\boldsymbol{x}) \Big) = \alpha \nabla_{\boldsymbol{x}} f(\boldsymbol{x}) + \beta \nabla_{\boldsymbol{x}} g(\boldsymbol{x}).$$

The gradient operation also satisfies the product rule, which says that

$$\nabla_{\boldsymbol{x}} \Big(f(\boldsymbol{x}) g(\boldsymbol{x}) \Big) = f(\boldsymbol{x}) \nabla_{\boldsymbol{x}} g(\boldsymbol{x}) + g(\boldsymbol{x}) \nabla_{\boldsymbol{x}} f(\boldsymbol{x}).$$

The gradient operation also satisfies the quotient rule, which says that

$$\nabla_{\boldsymbol{x}} \left(\frac{f(\boldsymbol{x})}{g(\boldsymbol{x})} \right) = \frac{g(\boldsymbol{x}) \nabla_{\boldsymbol{x}} f(\boldsymbol{x}) - f(\boldsymbol{x}) \nabla_{\boldsymbol{x}} g(\boldsymbol{x})}{g^2(\boldsymbol{x})}.$$

If we define a function $f : \mathbb{F}^m \to \mathbb{F}$ and $g : \mathbb{F}^n \to \mathbb{F}^m$, then the composition of these two functions can be expressed as $\phi(x) = f(g(x))$, where $\phi : \mathbb{F}^n \to \mathbb{F}$. Now we can define a final property of the gradient, called the chain rule:

$$\Big(\nabla_{\boldsymbol{x}}\phi(\boldsymbol{x})\Big)_i = \begin{bmatrix} \frac{\partial g_1(\boldsymbol{x})}{\partial x_i} & \dots & \frac{\partial g_m(\boldsymbol{x})}{\partial x_i} \end{bmatrix} \nabla_{\boldsymbol{x}} f(g(\boldsymbol{x}))$$

$$\nabla_{\boldsymbol{x}}\phi(\boldsymbol{x}) = \begin{bmatrix} \Big(\nabla_{\boldsymbol{x}}\phi(\boldsymbol{x})\Big)_1 \\ \vdots \\ \Big(\nabla_{\boldsymbol{x}}\phi(\boldsymbol{x})\Big)_n \end{bmatrix}$$

Common Gradients

Assume $\boldsymbol{x}$ and $\boldsymbol{w}$ are n-dimensional vectors, $\boldsymbol{A}$ is an $n \times n$ matrix, and $\boldsymbol{P}$ is an $n \times n$ symmetric matrix. Below are a few common gradients:

1. $\nabla_{\boldsymbol{x}}\big(\boldsymbol{x}^T\boldsymbol{w}\big) = \nabla_{\boldsymbol{x}}\big(\boldsymbol{w}^T\boldsymbol{x}\big) = \boldsymbol{w}$
2. $\nabla_{\boldsymbol{x}}\big(\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x}\big) = (\boldsymbol{A} + \boldsymbol{A}^T)\boldsymbol{x}$
3. $\nabla_{\boldsymbol{x}}\big(\boldsymbol{x}^T\boldsymbol{P}\boldsymbol{x}\big) = 2\boldsymbol{P}\boldsymbol{x}$

10.5.2 Hessian

Again, consider a function $f : \mathbb{F}^n \to \mathbb{F}$ that maps vectors to scalars. For any vector $\boldsymbol{x} \in \mathbb{F}^n$ with elements $x_1, \dots, x_n$, the **Hessian** of f is defined such that

$$\nabla^2_{\boldsymbol{x}} f(\boldsymbol{x}) = \frac{\partial^2 f}{\partial \boldsymbol{x}^2} = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2} & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1} & \cdots & \frac{\partial^2 f}{\partial x_n^2} \end{bmatrix} \in \mathbb{F}^{n \times n}.$$

Assume $\boldsymbol{x}$ and $\boldsymbol{b}$ are n-dimensional vectors, $\boldsymbol{A}$ is an $n \times n$ matrix, and c is a scalar value. It is useful to know the Hessian of the quadratic function:

$$\nabla^2_{\boldsymbol{x}}(\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}^T\boldsymbol{x} + c) = (\boldsymbol{A} + \boldsymbol{A}^T).$$

Note that if $\boldsymbol{A}$ is a symmetric matrix, then this Hessian is equal to $2\boldsymbol{A}$.

10.5.3 Jacobian

Now consider a function $f : \mathbb{F}^n \to \mathbb{F}^m$ that maps vectors to vectors. For any vector $\boldsymbol{x} \in \mathbb{F}^n$ with elements $x_1, \dots, x_n$, the **Jacobian** of f is defined such that

$$\boldsymbol{D}_{\boldsymbol{x}} f(\boldsymbol{x}) = \frac{\partial f}{\partial \boldsymbol{x}} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \cdots & \frac{\partial f_m}{\partial x_n} \end{bmatrix} \in \mathbb{F}^{m \times n}.$$

Chapter 11

Linear Maps

11.1 Properties of Linear Maps

Let (U, F) and (V, F) be two vector spaces over the same field F and let $\mathcal{A} : U \to V$ be a map from the set of vectors U to vectors V. For any $\boldsymbol{u} \in U$, we can write $\mathcal{A}(\boldsymbol{u}) = \boldsymbol{v}$ for some $\boldsymbol{v} \in V$. $\mathcal{A}$ is considered a **linear map** if and only if it satisfies the **principle of superposition**, which says that

$$\mathcal{A}(\alpha_1 \boldsymbol{u_1} + \alpha_2 \boldsymbol{u_2}) = \alpha_1 \mathcal{A}(\boldsymbol{u_1}) + \alpha_2 \mathcal{A}(\boldsymbol{u_2}),$$

$$\forall \alpha_1, \alpha_2 \in F,\ \forall \boldsymbol{u_1}, \boldsymbol{u_2} \in U.$$

11.2 Matrix Representation

Any linear map between finite dimensional vector spaces can be represented as matrix multiplication on coordinate vectors. Let $\mathcal{A} : U \to V$ be a linear map from (U, F) to (V, F), where the dimension of U is n and the dimension of V is m. Let $\{\boldsymbol{u_j}\}_{j=1}^{n}$ be a basis for U and $\{\boldsymbol{v_i}\}_{i=1}^{m}$ be a basis for V. For each $j = 1, \ldots, n$, there exists a unique set of constants $\{a_{1j}, \ldots, a_{mj}\}$ such that

$$\mathcal{A}(\boldsymbol{u_j}) = \sum_{i=1}^{m} a_{ij} \boldsymbol{v_i}.$$

The matrix representation of the linear map $\mathcal{A}$ is the $m \times n$ matrix $\boldsymbol{A}$ whose ijth element is a_{ij}. To find this matrix representation, we can find the mapping of each of the basis vectors in U: $\mathcal{A}(\boldsymbol{u_1}), \ldots, \mathcal{A}(\boldsymbol{u_m})$. Then for each mapping $\mathcal{A}(\boldsymbol{u_j})$, we can set up the summation equation with the basis vectors in V to determine the elements of the matrix $\boldsymbol{A}$.

Note that the matrix representation of $\mathcal{A}$ with respect to the standard basis is simply $A = \begin{bmatrix} \mathcal{A}(\boldsymbol{u_1}) & \ldots & \mathcal{A}(\boldsymbol{u_m}) \end{bmatrix}$. Additionally, the composition of linear maps (i.e. $\mathcal{B}(\mathcal{A}(x))$) corresponds to matrix multiplication (i.e. BAx).

11.2.1 Change of Basis

Let (U, F) and (V, F) be two linear spaces over the same field F. Suppose $\{\boldsymbol{u_j}\}_{j=1}^{n}$ and $\{\tilde{\boldsymbol{u}}_{\boldsymbol{j}}\}_{j=1}^{n}$ are two bases for U and $\{\boldsymbol{v_i}\}_{i=1}^{m}$ and $\{\tilde{\boldsymbol{v}}_{\boldsymbol{i}}\}_{i=1}^{m}$ are two bases for V. Now let $\boldsymbol{A}$ be the matrix representation of the linear map $\mathcal{A} : U \to V$ with respect to the bases $\{\boldsymbol{u_j}\}_{j=1}^{n}$ and $\{\boldsymbol{v_i}\}_{i=1}^{m}$, and let $\tilde{\boldsymbol{A}}$ be the matrix representation of the linear map $\mathcal{A} : U \to V$ with respect to the bases $\{\tilde{\boldsymbol{u}}_{\boldsymbol{j}}\}_{j=1}^{n}$ and $\{\tilde{\boldsymbol{v}}_{\boldsymbol{i}}\}_{i=1}^{m}$. The matrices $\boldsymbol{A}$ and $\tilde{\boldsymbol{A}}$ are said to be equivalent because they represent the same linear map, and they respect the following relationship:

$$\tilde{\boldsymbol{A}} = \boldsymbol{QAP}, \text{ where}$$

$$\boldsymbol{P} = \begin{bmatrix} \boldsymbol{u_1} & \dots & \boldsymbol{u_n} \end{bmatrix}^{-1} \begin{bmatrix} \tilde{\boldsymbol{u}}_{\boldsymbol{1}} & \dots & \tilde{\boldsymbol{u}}_{\boldsymbol{n}} \end{bmatrix},$$

$$\boldsymbol{Q} = \begin{bmatrix} \tilde{\boldsymbol{v}}_{\boldsymbol{1}} & \dots & \tilde{\boldsymbol{v}}_{\boldsymbol{m}} \end{bmatrix}^{-1} \begin{bmatrix} \boldsymbol{v_1} & \dots & \boldsymbol{v_m} \end{bmatrix}.$$

11.3 Range and Null Space

11.3.1 Definitions

Given two vector spaces (U, F) and (V, F) and a linear map $\mathcal{A} : U \to V$, the **range space**, or **image**, of the map $\mathcal{A}$ is defined as

$$R(\mathcal{A}) = \{\boldsymbol{v} \in V : \boldsymbol{v} = \mathcal{A}(\boldsymbol{u}),\ \boldsymbol{u} \in U\} \subseteq V.$$

If the matrix representation of $\mathcal{A}$ is $\boldsymbol{A}$, we can also express the range space as

$$R(\boldsymbol{A}) = \{\boldsymbol{v} \in V : \boldsymbol{v} = \boldsymbol{Au},\ \boldsymbol{u} \in U\} \subseteq V.$$

Given the same linear map $\mathcal{A}$, the **null space**, or **kernel**, of the map $\mathcal{A}$ is

$$N(\mathcal{A}) = \{\boldsymbol{u} \in U : \mathcal{A}(\boldsymbol{u}) = \boldsymbol{0}_V\} \subseteq U.$$

Equivalently, using the matrix representation, we can define the null space as

$$N(\boldsymbol{A}) = \{\boldsymbol{u} \in U : \boldsymbol{Au} = \boldsymbol{0}_V\} \subseteq U.$$

The **rank** of the linear map $\mathcal{A}$ is equal to the dimension of the range space, and the **nullity** of $\mathcal{A}$ is equal to the dimension of the null space. Because the range space is a subset of the codomain, the rank is less than or equal to the dimension of the codomain. Similarly, because the null space is a subset of the domain, the nullity is less than or equal to the dimension of the domain.

11.3.2 Rank-Nullity Theorem

Consider a linear map $\mathcal{A} : U \to V$ with the matrix representation $\boldsymbol{A} \in \mathbb{F}^{m \times n}$. The **rank-nullity theorem** states that the rank of $\mathcal{A}$ plus the nullity of $\mathcal{A}$ is equal to the dimension of the domain, which we can express as

$$\dim R(\mathcal{A}) + \dim N(\mathcal{A}) = \dim U = n.$$

The rank-nullity theorem is depicted pictorially in figure 11.1.

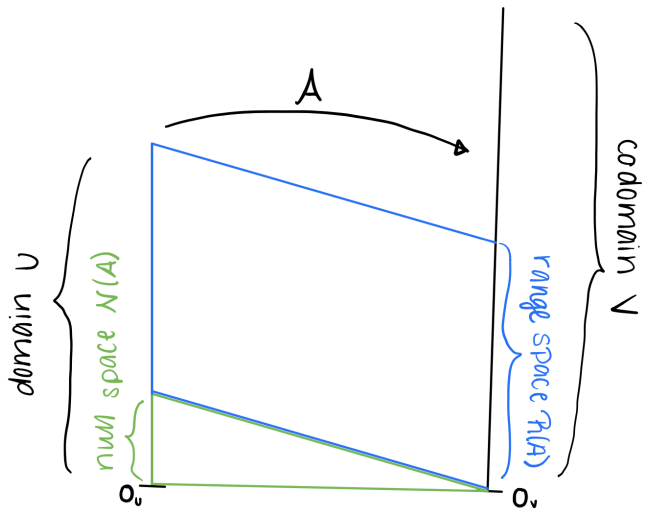

Figure 11.1: $\mathcal{A}$ is a linear map from the domain U to the codomain V. The null space of $\mathcal{A}$ is a subset of the domain, as shown in green. The range space of $\mathcal{A}$ is a subset of the codomain, as shown in blue. Notice that the dimension of the null space is less than the dimension of the domain, and the dimension of the range space is less than the dimension of the codomain. In addition, the dimension of the null space and range space sum to the dimension of the domain.

Proof: To prove the rank-nullity theorem, let's first denote the dimension of the null space as $\dim N(\mathcal{A}) = k$ and the basis of the null space as $\{\boldsymbol{u_1}, \ldots, \boldsymbol{u_k}\}$. Note that the null space $N(\mathcal{A})$ is a subset of the vector space U, so the basis of the null space is a set of linearly independent vectors in U. We can extend this set with $n-k$ different linearly independent vectors to form a full basis for U, which we'll write as $\{\boldsymbol{u_1}, \ldots, \boldsymbol{u_k}, \boldsymbol{u_{k+1}}, \ldots, \boldsymbol{u_n}\}$. Now the range space of $\mathcal{A}$ is

$$R(\mathcal{A}) = \text{span}\{\mathcal{A}(\boldsymbol{u_1}), \ldots, \mathcal{A}(\boldsymbol{u_k}), \mathcal{A}(\boldsymbol{u_{k+1}}), \ldots, \mathcal{A}(\boldsymbol{u_n})\}$$

Recall that we defined $\{\boldsymbol{u_1}, \ldots, \boldsymbol{u_k}\}$ to be a basis for the null space of $\mathcal{A}$, which implies $\mathcal{A}(\boldsymbol{u_1}), \ldots, \mathcal{A}(\boldsymbol{u_k}) = \boldsymbol{0}_V$. Now we can see that the range space of $\mathcal{A}$ is

$$R(\mathcal{A}) = \text{span}\{\boldsymbol{0}_V, \ldots, \boldsymbol{0}_V, \mathcal{A}(\boldsymbol{u_{k+1}}), \ldots, \mathcal{A}(\boldsymbol{u_n})\} = \text{span}\{\mathcal{A}(\boldsymbol{u_{k+1}}), \ldots, \mathcal{A}(\boldsymbol{u_n})\}.$$

Now we want to show that the set $\{\mathcal{A}(\boldsymbol{u_{k+1}}), \ldots, \mathcal{A}(\boldsymbol{u_n})\}$ is actually a basis for $R(\mathcal{A})$. To do so, we need to prove that the set is linearly independent, which can be shown through contradiction. Let's suppose that the set is not linearly independent and there exist scalars $\alpha_{k+1}, \ldots, \alpha_n$ that are not all zero such that

$$\alpha_{k+1}\mathcal{A}(\boldsymbol{u_{k+1}}) + \ldots + \alpha_n\mathcal{A}(\boldsymbol{u_n}) = \boldsymbol{0}_V.$$

Because $\mathcal{A}$ is a linear map, we could write this as

$$\mathcal{A}(\alpha_{k+1}\boldsymbol{u_{k+1}} + \ldots + \alpha_n \boldsymbol{u_n}) = \boldsymbol{0}_V.$$

This would then imply that

$$(\alpha_{k+1}\boldsymbol{u_{k+1}} + \ldots + \alpha_n \boldsymbol{u_n}) \in N(\mathcal{A}).$$

Recall that we denoted $\{\boldsymbol{u_1}, \ldots, \boldsymbol{u_k}\}$ to be a basis for $N(\mathcal{A})$, which means any element of the null space could be represented as a linear combination of these vectors. This implies that there exist scalars $\alpha_1, \ldots, \alpha_k$ such that

$$(\alpha_{k+1}\boldsymbol{u_{k+1}} + \ldots + \alpha_n \boldsymbol{u_n}) = (\alpha_1 \boldsymbol{u_1} + \ldots + \alpha_k \boldsymbol{u_k})$$

$$(\alpha_{k+1}\boldsymbol{u_{k+1}} + \ldots + \alpha_n \boldsymbol{u_n} - \alpha_1 \boldsymbol{u_1} - \ldots - \alpha_k \boldsymbol{u_k}) = 0$$

$$(\beta_1 \boldsymbol{u_1} + \ldots + \beta_n \boldsymbol{u_n}) = 0$$

Now remember that the set $\{\boldsymbol{u_1}, \ldots, \boldsymbol{u_k}, \boldsymbol{u_{k+1}}, \ldots, \boldsymbol{u_n}\}$ is a full basis for U. Because this set is a basis, the vectors must be linearly independent, and we cannot find non-zero scalars $\beta_1, \ldots, \beta_n$ such that $\beta_1 \boldsymbol{u_1} + \ldots + \beta_n \boldsymbol{u_n} = 0$.

Therefore, we have found a contradiction to our previous assumption that $\{\mathcal{A}(\boldsymbol{v_{k+1}}), \ldots, \mathcal{A}(\boldsymbol{v_n})\}$ is not a linearly independent set. This set must be linearly independent and is a basis for the range space of $\mathcal{A}$. Therefore, the dimension of the range space is $\dim R(\mathcal{A}) = n - k$.

To complete our proof of the rank-nullity theorem, let's recall the dimension of the vector space $\dim(U) = n$, the dimension of the null space $\dim N(\mathcal{A}) = k$, and the dimension of the range space $\dim R(\mathcal{A}) = n-k$. Clearly, $(n-k)+k = n$, so we have shown the rank-nullity theorem:

$$\dim R(\mathcal{A}) + \dim N(\mathcal{A}) = \dim U.$$

11.3.3 Sylvester's Inequality

A similar theorem to the rank-nullity theorem that relates the ranks of two linear maps is Sylvester's inequality. Let $\mathcal{A} : U \to V$ be a linear map with the matrix representation $\boldsymbol{A} \in \mathbb{F}^{m \times n}$ and $\mathcal{B} : W \to U$ be a linear map with the matrix representation $\boldsymbol{B} \in \mathbb{F}^{n \times p}$. The composition $\mathcal{A}(\mathcal{B}(\cdot))$ can be represented by the matrix $\boldsymbol{AB} \in \mathbb{F}^{m \times p}$. **Sylvester's inequality** states that

$$\dim R(\boldsymbol{A}) + \dim R(\boldsymbol{B}) - n \leq \dim R(\boldsymbol{AB}) \leq \min\{\dim R(\boldsymbol{A}), \dim R(\boldsymbol{B})\}$$

Figure 11.2 helps to demonstrate the concept of Sylvester's inequality.

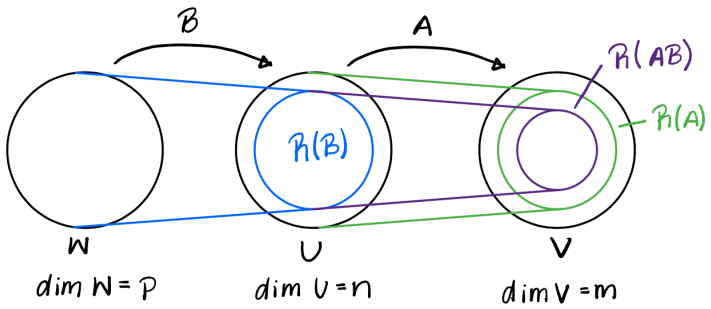

Figure 11.2: $\mathcal{B}$ maps vector space W to vector space U, then $\mathcal{A}$ maps U to V. The range space of $\boldsymbol{B}$ is shown in blue, the range space of $\boldsymbol{A}$ is shown in green, and the range space of $\boldsymbol{AB}$ is shown in purple. Notice than the range space of $\boldsymbol{AB}$ is smaller than the range space of $\boldsymbol{A}$ and the range space of $\boldsymbol{B}$.

Proof: We'll start with the first inequality. The rank-nullity theorem says that

$$\dim R(\boldsymbol{A}) + \dim N(\boldsymbol{A}) = n$$

$$\dim R(\boldsymbol{B}) + \dim N(\boldsymbol{B}) = p$$

$$\dim R(\boldsymbol{AB}) + \dim N(\boldsymbol{AB}) = p$$

Now, from these three inequalities, notice that

$$\dim R(\boldsymbol{A}) + \dim R(\boldsymbol{B}) - n - \dim R(\boldsymbol{AB}) = \dim N(\boldsymbol{AB}) - \dim N(\boldsymbol{A}) - \dim N(\boldsymbol{B})$$

$$\dim R(\boldsymbol{A}) + \dim R(\boldsymbol{B}) - n \leq \dim R(\boldsymbol{AB}) \iff \dim N(\boldsymbol{A}) + \dim N(\boldsymbol{B}) \geq \dim N(\boldsymbol{AB})$$

Therefore, to show Sylvester's inequality is true, it is equivalent to show that the following relationship holds: $\dim N(\boldsymbol{A}) + \dim N(\boldsymbol{B}) \geq \dim N(\boldsymbol{AB})$.

If a vector $\boldsymbol{x}$ is in the null space of $\boldsymbol{B}$, it satisfies $\boldsymbol{Bx} = \boldsymbol{0_n}$. It follows that it must also satisfy $\boldsymbol{ABx} = \boldsymbol{0_m}$, so the vector $\boldsymbol{x}$ is also in the null space of $\boldsymbol{AB}$. Clearly, every element of $N(\boldsymbol{B})$ is also an element of $N(\boldsymbol{AB})$, so $N(\boldsymbol{B}) \subseteq N(\boldsymbol{AB})$. Let $\dim N(\boldsymbol{B}) = r_1$ and $\dim N(\boldsymbol{AB}) = r_2$, where $r_1 \leq r_2$ because $N(\boldsymbol{B}) \subseteq N(\boldsymbol{AB})$. Assume that $\{\boldsymbol{v_1}, \ldots, \boldsymbol{v_{r_1}}\}$ is a basis for $N(\boldsymbol{B})$. Because $N(\boldsymbol{B})$ is a subset of $N(\boldsymbol{AB})$, the basis we defined for $N(\boldsymbol{B})$ is a set of linearly independent vectors in $N(\boldsymbol{AB})$. We can extend this set with $r_2 - r_1$ different linearly independent vectors to form a full basis for $N(\boldsymbol{AB})$, which we'll express as $\{\boldsymbol{v_1}, \ldots, \boldsymbol{v_{r_1}}, \boldsymbol{v_{r_1+1}}, \ldots, \boldsymbol{v_{r_2}}\}$. Based on this definition,

$$\boldsymbol{ABv_i} = 0 \text{ for } 1 \leq i \leq r_2$$

$$\boldsymbol{Bv_i} = 0 \text{ for } 1 \leq i \leq r_1$$

$$\boldsymbol{Bv_i} \neq 0 \text{ for } (r_1 + 1) \leq i \leq r_2$$

Now we can define a new set of vectors $\{\boldsymbol{z_{r_1+1}}, \ldots, \boldsymbol{z_{r_2}}\}$ such that

$$\boldsymbol{z_i} := \boldsymbol{B v_i} \neq 0 \text{ for } (r_1 + 1) \leq i \leq r_2$$

$$\boldsymbol{A z_i} = \boldsymbol{AB} v_i = 0 \text{ for } (r_1 + 1) \leq i \leq r_2$$

Now if $\{\boldsymbol{z_{r_1+1}}, \ldots, \boldsymbol{z_{r_2}}\}$ is a linearly independent set, then it forms a part of the basis for the null space $N(\boldsymbol{A})$. Let's suppose the set is not linearly independent, then for scalars $\alpha_{r_1+1}, \ldots, \alpha_{r_2}$ that are not all zero-valued,

$$\alpha_{r_1+1}\boldsymbol{z_{r_1+1}} + \ldots + \alpha_{r_2}\boldsymbol{z_{r_2}} = 0$$

$$\boldsymbol{B}(\alpha_{r_1+1}\boldsymbol{v_{r_1+1}} + \ldots + \alpha_{r_2}\boldsymbol{v_{r_2}}) = 0$$

This would imply that the vector $\alpha_{r_1+1}\boldsymbol{v_{r_1+1}} + \ldots + \alpha_{r_2}\boldsymbol{v_{r_2}}$ is in the null space of $\boldsymbol{B}$. We previously defined the basis for the null space $N(\boldsymbol{B})$ as $\{\boldsymbol{v_1}, \ldots, \boldsymbol{v_{r_1}}\}$, which would imply that we can use arbitary scalars to write

$$\alpha_{r_1+1}\boldsymbol{v_{r_1+1}} + \ldots + \alpha_{r_2}\boldsymbol{v_{r_2}} = \beta_1\boldsymbol{v_1} + \ldots + \beta_{r_1}\boldsymbol{v_{r_1}}$$

However, recall that $\{\boldsymbol{v_1}, \ldots, \boldsymbol{v_{r_1}}, \boldsymbol{v_{r_1+1}}, \ldots, \boldsymbol{v_{r_2}}\}$ is a linearly independent set of vectors, so we have found a contradiction. The set $\{\boldsymbol{z_{r_1+1}}, \ldots, \boldsymbol{z_{r_2}}\}$ is linearly independent and must form at least a partial basis for $N(\boldsymbol{A})$. This set contains $r_2 - r_1$ elements, so $\dim N(\boldsymbol{A}) \geq r_2 - r_1$. Finally, we can write

$$\dim N(\boldsymbol{A}) + \dim N(\boldsymbol{B}) \geq (r_2 - r_1) + r_1 = r_2$$

$$\dim N(\boldsymbol{A}) + \dim N(\boldsymbol{B}) \geq \dim N(\boldsymbol{AB})$$

This proves the first inequality. Now we can prove the second inequality by showing that $\dim R(\boldsymbol{AB}) \leq \dim R(\boldsymbol{A})$ and $\dim R(\boldsymbol{AB}) \leq \dim R(\boldsymbol{B})$.

If a vector $\boldsymbol{y}$ is an element of $R(\boldsymbol{AB})$, then $\boldsymbol{y} = \boldsymbol{ABx}$ for some vector $\boldsymbol{x} \in \mathbb{R}^p$. We could write this vector as $\boldsymbol{y} = A\boldsymbol{z}$ for some vector $\boldsymbol{z} = \boldsymbol{Bx} \in \mathbb{R}^n$, which implies that the vector $\boldsymbol{y}$ is also an element of $R(\boldsymbol{A})$. Clearly, every element of $R(\boldsymbol{AB})$ is also an element of $R(\boldsymbol{A})$, so $R(\boldsymbol{AB}) \subseteq R(\boldsymbol{A})$. This tells us that

$$\dim R(\boldsymbol{AB}) \leq \dim R(\boldsymbol{A}).$$

Recall that earlier in this problem we showed that $N(\boldsymbol{B}) \subseteq N(\boldsymbol{AB})$, implying that $\dim N(\boldsymbol{B}) \leq \dim N(\boldsymbol{AB})$. From the rank-nullity theorem, we also said that $\dim R(\boldsymbol{B}) + \dim N(\boldsymbol{B}) = p$ and $\dim R(\boldsymbol{AB}) + \dim N(\boldsymbol{AB}) = p$. Now we see that

$$\dim R(\boldsymbol{AB}) \leq \dim R(\boldsymbol{B}).$$

This shows that the second inequality holds, which completes our proof.

11.4 A-Invariant Subspaces

11.4.1 A-Invariance

Consider a vector space (V, F) and a linear map $\mathcal{A} : V \to V$ with matrix representation $\boldsymbol{A}$. A subspace $M \subset V$ is considered **A-invariant** if

$$\boldsymbol{x} \in M \implies \boldsymbol{x} \in M.$$

Some examples of $\mathcal{A}$-invariant subspaces are $N(\boldsymbol{A})$, $R(\boldsymbol{A})$, and $N(f(\boldsymbol{A}))$, where f is an analytic function. If two subspaces M_1 and M_2 are $\mathcal{A}$-invariant, then the intersection of these spaces, $M_1 \cap M_2$, and the sum, $M_1 + M_2$, are also $\mathcal{A}$-invariant. However, the union, $M_1 \cup M_2$, is not necessarily $\mathcal{A}$-invariant.

11.4.2 Second Representation Theorem

Consider a vector space (V, F) and a linear map $\mathcal{A} : V \to V$ with matrix representation $\boldsymbol{A} \in \mathbb{F}^{n\times n}$. Let V be a finite dimensional vector space that can be represented as the direct sum of two subspaces such that $V = M_1 \oplus M_2$. Assume $\dim V = n$, $\dim M_1 = k$, and $\dim M_2 = n - k$. We can express $\boldsymbol{A}$ as

$$\boldsymbol{A} = \begin{bmatrix} \boldsymbol{A_{11}} & \boldsymbol{A_{12}} \\ \boldsymbol{A_{21}} & \boldsymbol{A_{22}} \end{bmatrix},$$

where $\boldsymbol{A_{11}} \in \mathbb{F}^{k\times k}$, $\boldsymbol{A_{12}} \in \mathbb{F}^{k\times(n-k)}$, $\boldsymbol{A_{21}} \in \mathbb{F}^{(n-k)\times k}$ and $\boldsymbol{A_{22}} \in \mathbb{F}^{(n-k)\times(n-k)}$. The matrix $\boldsymbol{A_{11}}$ maps elements in the subspace M_1 to M_1, $\boldsymbol{A_{12}}$ maps elements in M_2 to M_1, $\boldsymbol{A_{21}}$ maps elements in M_1 to M_2, and $\boldsymbol{A_{22}}$ map M_2 to M_2.

If M_1 is $\mathcal{A}$-invariant, elements in M_1 cannot be mapped to M_2, which implies that $\boldsymbol{A_{21}}$ is the zero matrix. Therefore, $\boldsymbol{A}$ has a representation of the form:

$$\boldsymbol{A}_{\boldsymbol{M_1}} = \begin{bmatrix} \boldsymbol{A_{11}} & \boldsymbol{A_{12}} \\ \boldsymbol{0}^{(n-k)\times k} & \boldsymbol{A_{22}} \end{bmatrix}.$$

If M_2 is $\mathcal{A}$-invariant, elements in M_2 cannot be mapped to M_1, which implies that $\boldsymbol{A_{12}}$ is the zero matrix. Therefore, $\boldsymbol{A}$ has a representation of the form:

$$\boldsymbol{A}_{\boldsymbol{M_2}} = \begin{bmatrix} \boldsymbol{A_{11}} & \boldsymbol{0}^{k\times(n-k)} \\ \boldsymbol{A_{21}} & \boldsymbol{A_{22}} \end{bmatrix}.$$

If both M_1 and M_2 are both $\mathcal{A}$-invariant, elements in M_1 cannot be mapped to M_2 and elements in M_2 cannot be mapped to M_1, so both $\boldsymbol{A_{12}}$ and $\boldsymbol{A_{21}}$ are matrices of all zeros. Therefore, $\boldsymbol{A}$ has a representation of the form:

$$\boldsymbol{A}_{\boldsymbol{M_1 M_2}} = \begin{bmatrix} \boldsymbol{A_{11}} & \boldsymbol{0}^{k\times(n-k)} \\ \boldsymbol{0}^{(n-k)\times k} & \boldsymbol{A_{22}} \end{bmatrix}.$$

Justification: To understand this theorem, let $V = \mathbb{F}^n$ and assume M_1 and M_2 can be expressed as $M_1 = \text{span}(\boldsymbol{e_1}, \ldots, \boldsymbol{e_k})$ and $M_2 = \text{span}(\boldsymbol{e_{k+1}}, \ldots, \boldsymbol{e_n})$ respectively, where e_i is the ith standard basis vector in $\mathbb{F}^n$.

Any vector $\boldsymbol{x} \in V$ can be expressed as

$$\boldsymbol{x} = \begin{bmatrix} \boldsymbol{x_1} \\ \boldsymbol{x_2} \end{bmatrix}, \quad \text{where} \quad \boldsymbol{x_1} \in \mathbb{F}^k, \ \boldsymbol{x_2} \in \mathbb{F}^{n-k}.$$

With this representation of $\boldsymbol{x}$, we can write

$$\boldsymbol{A}_{M_1}\boldsymbol{x} = \begin{bmatrix} \boldsymbol{A_{11}x_1} + \boldsymbol{A_{12}x_2} \\ \boldsymbol{A_{22}x_2} \end{bmatrix} \quad \text{and} \quad \boldsymbol{A}_{M_2}\boldsymbol{x} = \begin{bmatrix} \boldsymbol{A_{11}x_1} \\ \boldsymbol{A_{21}x_1} + \boldsymbol{A_{22}x_2} \end{bmatrix}.$$

If $\boldsymbol{x}$ is an element of M_1, then it can be expressed in terms of $\{\boldsymbol{e_1}, \ldots, \boldsymbol{e_k}\}$, so $\boldsymbol{x_2} = \boldsymbol{0}_{n-k}$. Therefore, if $\boldsymbol{x} \in M_1$, then our expression for $\boldsymbol{A}_{M_1}\boldsymbol{x}$ becomes

$$\boldsymbol{A}_{M_1}\boldsymbol{x} = \begin{bmatrix} \boldsymbol{A_{11}x_1} \\ \boldsymbol{0}_{n-k} \end{bmatrix} \in M_1.$$

Now we can see that is $\boldsymbol{x}$ is an element of M_1, then $\boldsymbol{A}_{M_1}\boldsymbol{x}$ is also an element of M_1, which means that M_1 is in fact $\mathcal{A}$-invariant. Therefore, $\boldsymbol{A}_{M_1}$ is a valid representation for $\boldsymbol{A}$ if M_1 is $\mathcal{A}$-invariant.

If $\boldsymbol{x}$ is an element of M_2, then it can be expressed in terms of $\{\boldsymbol{e_{k+1}}, \ldots, \boldsymbol{e_n}\}$, so $\boldsymbol{x_1} = \boldsymbol{0}_k$. Therefore, if $\boldsymbol{x} \in M_2$, then our expression for $\boldsymbol{A}_{M_2}\boldsymbol{x}$ becomes

$$\boldsymbol{A}_{M_2}\boldsymbol{x} = \begin{bmatrix} \boldsymbol{0}_k \\ \boldsymbol{A_{22}x_2} \end{bmatrix} \in M_2.$$

Now we can see that is $\boldsymbol{x}$ is an element of M_2, then $\boldsymbol{A}_{M_2}\boldsymbol{x}$ is also an element of M_2, which means that M_2 is in fact $\mathcal{A}$-invariant. Therefore, $\boldsymbol{A}_{M_2}$ is a valid representation for $\boldsymbol{A}$ if M_2 is $\mathcal{A}$-invariant.

We can use this same argument to show that $\boldsymbol{A}_{M_1 M_2}$ is a valid representation for $\boldsymbol{A}$ if both M_1 and M_2 are $\mathcal{A}$-invariant.

11.5 Adjoints

11.5.1 Definition

Let (U, F) be a vector space with inner product $\langle \cdot, \cdot \rangle_U$ and (V, F) be a vector space with inner product $\langle \cdot, \cdot \rangle_V$. Let $\mathcal{A} : U \to V$ be a continuous linear map. The **adjoint** of $\mathcal{A}$, denoted $\mathcal{A}^*$, is the linear map

$$\mathcal{A}^* : V \to U \text{ s.t. } \langle \boldsymbol{v}, \mathcal{A}(\boldsymbol{u}) \rangle_V = \langle \mathcal{A}^*(\boldsymbol{v}), \boldsymbol{u} \rangle_U, \ \forall \boldsymbol{u} \in U, \boldsymbol{v} \in V.$$

Suppose the linear map $\mathcal{A}$ is represented by the matrix $\boldsymbol{A}$. If $U = \mathbb{F}^n$ and $V = \mathbb{F}^m$, then the matrix representation of the adjoint map $\mathcal{A}^*$ is $\boldsymbol{A}^*$.

11.5.2 Self-Adjoint Maps

Suppose (V, F) is a vector space equipped with the inner product $\langle \cdot, \cdot \rangle_V$ and $\mathcal{A} : V \to V$ is a continuous linear map with adjoint $\mathcal{A}^* : V \to V$. We say that the map $\mathcal{A}$ is **self-adjoint** if and only if $\mathcal{A} = \mathcal{A}^*$, or equivalently,

$$\langle \boldsymbol{x}, \mathcal{A}(\boldsymbol{y}) \rangle_V = \langle \mathcal{A}(\boldsymbol{x}), \boldsymbol{y} \rangle_V, \ \forall \boldsymbol{x}, \boldsymbol{y} \in V.$$

Now suppose $V = \mathbb{F}^n$ and let the linear map $\mathcal{A}$ be represented by a matrix $\boldsymbol{A} \in \mathbb{F}^{n \times n}$. The map $\mathcal{A}$ is self-adjoint if and only if the matrix $\boldsymbol{A}$ is Hermitian.

11.5.3 Fundamental Theorem of Linear Algebra

Let $\mathcal{A} : U \to V$ be a linear map with the adjoint $\mathcal{A}^*$. The domain and codomain of $\mathcal{A}$ can be decomposed into two orthogonal subspaces:

1. Domain – $U = N(\mathcal{A}) \oplus R(\mathcal{A}^*)$
2. Codomain – $V = R(\mathcal{A}) \oplus N(\mathcal{A}^*)$

From these relationships, we can also see that

1. $\dim U = \dim N(\mathcal{A}) + \dim R(\mathcal{A}^*)$
2. $\dim V = \dim R(\mathcal{A}) + \dim N(\mathcal{A}^*)$

Figure 11.3 helps to demonstrate the decomposition of the domain and codomain.

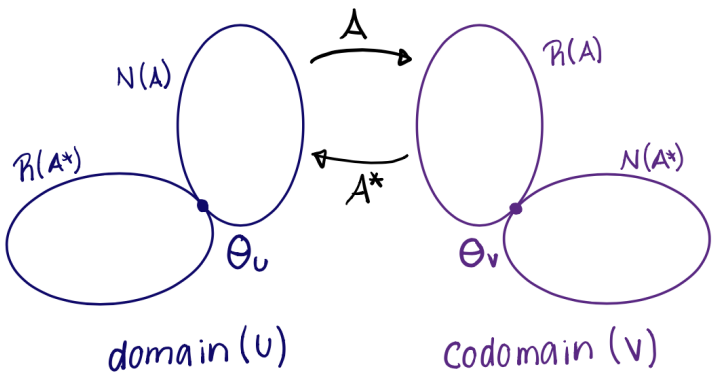

Figure 11.3: The linear map $\mathcal{A}$ maps elements in the domain U to the codomain V, and its adjoint, $\mathcal{A}^*$, maps elements in V to U. The domain is composed of the null space of $\mathcal{A}$ and the range space of $\mathcal{A}^*$, and the codomain is composed of the range space of $\mathcal{A}$ and the null space of $\mathcal{A}^*$.

This decomposition leads us to the **fundamental theorem of linear algebra**:

1. $R(\mathcal{A}) = N(\mathcal{A}^*)^\perp$
2. $N(\mathcal{A}) = R(\mathcal{A}^*)^\perp$
3. $R(\mathcal{A}\mathcal{A}^*) = R(\mathcal{A})$
4. $R(\mathcal{A}^*\mathcal{A}) = R(\mathcal{A}^*)$
5. $N(\mathcal{A}^*\mathcal{A}) = N(\mathcal{A})$
6. $N(\mathcal{A}\mathcal{A}^*) = N(\mathcal{A}^*)$

Proof (1 and 2): To prove part one of the fundamental theorem of linear algebra, we will first assume $\boldsymbol{y} \in N(\mathcal{A}^*)$. If $\boldsymbol{y}$ is in the null space of $\mathcal{A}^*$, then $\mathcal{A}^*(\boldsymbol{y}) = \boldsymbol{0}_U$. By the definition of the adjoint, this implies

$$\langle \boldsymbol{y}, \mathcal{A}(\boldsymbol{x})\rangle_V = \langle \mathcal{A}^*(\boldsymbol{y}), \boldsymbol{x}\rangle_U = \langle \boldsymbol{0}_U, \boldsymbol{x}\rangle_U = 0, \ \forall \boldsymbol{x} \in U.$$

Because $\boldsymbol{y}$ is orthogonal to $\mathcal{A}(\boldsymbol{x})$ for all $\boldsymbol{x} \in U$, we can say that $\boldsymbol{y} \in R(\mathcal{A})^\perp$. We showed $\boldsymbol{y} \in N(\mathcal{A}^*)$ implies $\boldsymbol{y} \in R(\mathcal{A})^\perp$, which tells us that $N(\mathcal{A}^*) \subseteq R(\mathcal{A})^\perp$.

Now we will assume $\boldsymbol{y} \in R(\mathcal{A})^\perp$. If $\boldsymbol{y}$ is in the orthogonal complement of $R(\mathcal{A})$, then $\boldsymbol{y}$ must be perpendicular to every element in $R(\mathcal{A})$, which can be expressed as $\langle \boldsymbol{y}, \mathcal{A}(\boldsymbol{x})\rangle_V = 0$ for every $\boldsymbol{x} \in U$. By the definition of the adjoint, this implies

$$\langle \mathcal{A}^*(\boldsymbol{y}), \boldsymbol{x}\rangle_U = \langle \boldsymbol{y}, \mathcal{A}(\boldsymbol{x})\rangle_V = 0, \ \forall \boldsymbol{x} \in U.$$

Because $\mathcal{A}^*(\boldsymbol{y})$ is orthogonal to $\boldsymbol{x}$ for all $\boldsymbol{x} \in U$, we know that $\mathcal{A}^*(\boldsymbol{y}) = \boldsymbol{0}_U$, so $\boldsymbol{y} \in N(\mathcal{A}^*)$. We showed $\boldsymbol{y} \in R(\mathcal{A})^\perp$ implies $\boldsymbol{y} \in N(\mathcal{A}^*)$, which tells us that $R(\mathcal{A})^\perp \subseteq N(\mathcal{A}^*)$. Now we can conclude that $N(\mathcal{A}^*) = R(\mathcal{A})^\perp$. We can prove the second part of the fundamental theorem in a very similar way.

Proof (3 and 4): To prove part three of the fundamental theorem of linear algebra, we will first assume $\boldsymbol{y} \in R(\mathcal{A})$. If this is true, then $\boldsymbol{y} = \mathcal{A}(\boldsymbol{x})$ for some $\boldsymbol{x} \in U$. From the second part of the fundamental theorem, we know that $N(\mathcal{A}) = R(\mathcal{A}^*)^\perp$, which means that $N(\mathcal{A})$ and $R(\mathcal{A}^*)$ compose the entire domain, U. Therefore, any $\boldsymbol{x} \in U$ can be expressed as $\boldsymbol{x} = \mathcal{A}^*(\bar{\boldsymbol{y}}) + \boldsymbol{z}$, where $\bar{\boldsymbol{y}} \in V$ and $\boldsymbol{z} \in N(\mathcal{A})$. Because the map $\mathcal{A}$ and its adjoint $\mathcal{A}^*$ are linear,

$$\boldsymbol{y} = \mathcal{A}(\boldsymbol{x}) = \mathcal{A}(\mathcal{A}^*(\bar{\boldsymbol{y}}) + \boldsymbol{z}) = \mathcal{A}(\mathcal{A}^*(\bar{\boldsymbol{y}})) + \mathcal{A}(\boldsymbol{z}) = \mathcal{A}(\mathcal{A}^*(\bar{\boldsymbol{y}})) + \boldsymbol{0}_V = \mathcal{A}(\mathcal{A}^*(\bar{\boldsymbol{y}})).$$

Now we can see $\boldsymbol{y} \in R(\mathcal{A})$ implies $\boldsymbol{y} \in R(\mathcal{A}\mathcal{A}^*)$, which tells us $R(\mathcal{A}) \subseteq R(\mathcal{A}\mathcal{A}^*)$.

Now we will assume $\boldsymbol{y} \in R(\mathcal{A}\mathcal{A}^*)$. If this is true, then $\boldsymbol{y} = \mathcal{A}(\mathcal{A}^*(\bar{\boldsymbol{y}}))$ for some $\bar{\boldsymbol{y}} \in V$. If we define $\boldsymbol{x} := \mathcal{A}^*(\bar{\boldsymbol{y}}) \in U$, then we have $\boldsymbol{y} = \mathcal{A}(\boldsymbol{x})$. Now we can see $\boldsymbol{y} \in R(\mathcal{A}\mathcal{A}^*)$ implies $\boldsymbol{y} \in R(\mathcal{A})$, which tells us $R(\mathcal{A}\mathcal{A}^*) \subseteq R(\mathcal{A})$. Now we can conclude that $R(\mathcal{A}) = R(\mathcal{A}\mathcal{A}^*)$. We can prove the fourth part of the fundamental theorem in a very similar way.

Proof (5 and 6): To prove part five of the fundamental theorem of linear algebra, we will first assume $\boldsymbol{x} \in N(\mathcal{A})$. If this is true, then $\mathcal{A}(\boldsymbol{x}) = \mathbf{0}_V$, thus

$$\mathcal{A}^*(\mathcal{A}(\boldsymbol{x})) = \mathcal{A}^*(\mathbf{0}_V) = \mathbf{0}_U.$$

We can see $\boldsymbol{x} \in N(\mathcal{A})$ implies $\boldsymbol{x} \in N(\mathcal{A}^*\mathcal{A})$, which implies $N(\mathcal{A}) \subseteq N(\mathcal{A}^*\mathcal{A})$.

Now we will assume $N(\mathcal{A}^*\mathcal{A})$, which means $\mathcal{A}^*(\mathcal{A}(\boldsymbol{x})) = \mathbf{0}_U$. Combining this observation with the definition of the adjoint, we can notice that

$$||\mathcal{A}(\boldsymbol{x})||^2 = \langle \mathcal{A}(\boldsymbol{x}), \mathcal{A}(\boldsymbol{x}) \rangle_V = \langle \boldsymbol{x}, \mathcal{A}^*(\mathcal{A}(\boldsymbol{x})) \rangle_U = \langle \boldsymbol{x}, \mathbf{0}_U \rangle_U = \mathbf{0}_U$$

From the first property of vector norms, this finding implies that $\mathcal{A}(\boldsymbol{x}) = \mathbf{0}_U$. Now we can see that $\boldsymbol{x} \in N(\mathcal{A}^*\mathcal{A})$ implies $\boldsymbol{x} \in N(\mathcal{A})$, which tells us that $N(\mathcal{A}^*\mathcal{A}) \subseteq N(\mathcal{A})$. Now we can conclude that $N(\mathcal{A}) \subseteq N(\mathcal{A}^*\mathcal{A})$. We can prove the sixth part of the fundamental theorem in a very similar way.

Chapter 12

Linear Equations

12.1 Linear Matrix Equation

12.1.1 Existence & Uniqueness of Solution

Consider a linear map $\mathcal{A} : U \to V$ and its adjoint $\mathcal{A}^* : V \to U$, where $U = \mathbb{F}^n$ and $V = \mathbb{F}^m$. Assume the matrix representation of the map $\mathcal{A}$ is given by $\boldsymbol{A} \in \mathbb{F}^{m \times n}$. Suppose we would like to find a solution $\boldsymbol{x} \in \mathbb{F}^n$ to the linear equation $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}$. The solution set for this equation can be expressed as

$$S = \{\boldsymbol{x} \in \mathbb{F}^n : \boldsymbol{A}\boldsymbol{x} = \boldsymbol{y}\}.$$

A solution to the given linear equation exists if and only if $\boldsymbol{y}$ is in the range space of $\boldsymbol{A}$. If $\boldsymbol{y}$ is not in the range space of $\boldsymbol{A}$, then S is the empty set. Note that $\boldsymbol{y}$ is in the range space of $\boldsymbol{A}$ if $\text{rank}(\boldsymbol{A}) = \text{rank}(\tilde{\boldsymbol{A}})$, where $\tilde{\boldsymbol{A}} = [\boldsymbol{A} \mid \boldsymbol{y}]$.

If $\boldsymbol{x}$ and $\bar{\boldsymbol{x}}$ are two distinct solutions to the linear equation $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{y}$, then $\boldsymbol{A}(\boldsymbol{x} - \bar{\boldsymbol{x}}) = 0$, which implies that $(\boldsymbol{x} - \bar{\boldsymbol{x}}) \in N(\boldsymbol{A})$. If $\bar{\boldsymbol{x}}$ is a known solution to the linear equation (i.e. $\boldsymbol{y} = \boldsymbol{A}\bar{\boldsymbol{x}}$), then the solution set can be given as

$$S_{\bar{\boldsymbol{x}}} = \{\boldsymbol{x} = \bar{\boldsymbol{x}} + \boldsymbol{z} : \boldsymbol{z} \in N(\boldsymbol{A})\}.$$

Therefore, the solution $\bar{\boldsymbol{x}}$ is unique if and only if $N(\boldsymbol{A}) = \{\boldsymbol{0}_n\}$. Note that $N(\boldsymbol{A}) = \{\boldsymbol{0}_n\}$ if and only if the linear map $\mathcal{A}$ is injective (one-to-one). Equivalently, $N(\boldsymbol{A}) = \{\boldsymbol{0}_n\}$ if and only if $\boldsymbol{A}$ has full column rank (i.e. $\text{rank}(\boldsymbol{A}) = n$), which is equivalent to saying that the columns of $\boldsymbol{A}$ are linearly independent

12.1.2 Unique Solution

Suppose the linear equation $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{y}$ has a unique solution, which we know occurs if and only if $N(\boldsymbol{A}) = \{\boldsymbol{0}_n\}$, or equivalently $\dim N(\boldsymbol{A}) = 0$. From the fundamental theorem of linear algebra, $N(\boldsymbol{A}) = R(\boldsymbol{A}^*)^\perp$. This implies that

$$\dim R(\boldsymbol{A}^*) = \dim U - \dim N(\boldsymbol{A}) = n - 0 = n.$$

From the fundamental theorem of linear algebra, $R(\boldsymbol{A}^*\boldsymbol{A}) = R(\boldsymbol{A}^*)$, so

$$\text{dim}R(\boldsymbol{A}^*\boldsymbol{A}) = \text{dim}R(\boldsymbol{A}^*) = n.$$

$\boldsymbol{A}^*\boldsymbol{A}$ is an $n \times n$ matrix, so this finding implies that $\boldsymbol{A}^*\boldsymbol{A}$ has full rank and is invertible. If we left multiply our linear equation by the matrix $\boldsymbol{A}^*$, we get

$$\boldsymbol{A}^*\boldsymbol{y} = \boldsymbol{A}^*\boldsymbol{A}\boldsymbol{x}.$$

We just found that if $N(\boldsymbol{A}) = \{\boldsymbol{0}_n\}$, then the matrix $\boldsymbol{A}^*\boldsymbol{A}$ is invertible. Therefore, the unique solution in this case is

$$\boldsymbol{x} = (\boldsymbol{A}^*\boldsymbol{A})^{-1}\boldsymbol{A}^*\boldsymbol{y}.$$

12.2 Minimum Norm Solution

Assume that $\boldsymbol{y} \in R(\boldsymbol{A})$, so there exists a solution to the linear equation $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}$. Consider the case when the linear map $\mathcal{A}$ is not injective. Because $N(\boldsymbol{A}) \neq \{\boldsymbol{0}_n\}$, the solution set, S, must contain more than one element. Because there are infinitely many elements in $N(\boldsymbol{A})$ if its dimension is greater than zero, S actually contains infinitely many solutions. Generally, we want to select the solution with the smallest Euclidean norm, which we call the minimum norm solution. We can express the minimum norm solution, which I denote $\hat{\boldsymbol{x}}$, as

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}:\boldsymbol{A}\boldsymbol{x}=\boldsymbol{y}}{\arg\min} ||\boldsymbol{x}||_2 = \underset{\boldsymbol{x}\in S_{\hat{\boldsymbol{x}}}}{\arg\min} ||\boldsymbol{x}||_2.$$

12.2.1 Surjective Map

In order to minimize the norm of the solution $\hat{\boldsymbol{x}}$, we should choose $\hat{\boldsymbol{x}}$ to be perpendicular to the null space of $\boldsymbol{A}$. From the fundamental theorem of linear algebra, the orthogonal complement of $N(\mathcal{A})$ is $R(\mathcal{A}^*)$. Therefore, we should choose $\hat{\boldsymbol{x}}$ to be in $R(\boldsymbol{A}^*)$. This means there exists a vector $\boldsymbol{v} \in V$ such that $\hat{\boldsymbol{x}} = \boldsymbol{A}^*\boldsymbol{v}$. The minimum norm solution must satisfy the linear equation, so

$$\boldsymbol{y} = \boldsymbol{A}\hat{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{A}^*\boldsymbol{v}.$$

Let's consider the case when the map $\mathcal{A}$ is surjective (onto). From the definition of surjectivity, we know that, in this case, $R(\boldsymbol{A}) = V$, so $\text{dim}R(\boldsymbol{A}) = \text{dim}V$. From the fundamental theorem of linear algebra $R(\boldsymbol{A}\boldsymbol{A}^*) = R(\boldsymbol{A})$, which implies

$$\text{dim}R(\boldsymbol{A}\boldsymbol{A}^*) = \text{dim}R(\boldsymbol{A}) = \text{dim}V = m$$

$\boldsymbol{A}\boldsymbol{A}^*$ is an $m \times m$ matrix, so it has full rank, indicating that it is invertible. Therefore, assuming $\mathcal{A}$ is surjective, we can see that

$$\boldsymbol{v} = (\boldsymbol{A}\boldsymbol{A}^*)^{-1}\boldsymbol{y}.$$

Plugging this expression for $\boldsymbol{v}$ into our equation for $\hat{\boldsymbol{x}}$, we find that when the map $\mathcal{A}$ is surjective, the minimum norm solution is

$$\hat{\boldsymbol{x}} = \boldsymbol{A}^*\boldsymbol{v} = \boldsymbol{A}^*(\boldsymbol{A}\boldsymbol{A}^*)^{-1}\boldsymbol{y}.$$

This is the minimum norm solution when the linear map $\mathcal{A}$ is surjective. Note that $\mathcal{A}$ is surjective if and only if $\boldsymbol{A}$ has full row rank (i.e. $\text{rank}(\boldsymbol{A}) = m$), which is equivalent to saying that the rows of $\boldsymbol{A}$ are linearly independent

12.2.2 Non-Surjective Map

Now let's consider the case when the map $\mathcal{A}$ is neither injective nor surjective. In this case, we need to use the singular value decomposition of $\boldsymbol{A}$ to compute its pseudoinverse $\boldsymbol{A}^\dagger$. The minimum norm solution is now given by

$$\hat{\boldsymbol{x}} = \boldsymbol{A}^\dagger\boldsymbol{y}.$$

Recall that when $\boldsymbol{A}$ has full column rank (i.e. $r = n \leq m$), $\boldsymbol{A}^\dagger = (\boldsymbol{A}^*\boldsymbol{A})^{-1}\boldsymbol{A}^*$, and when $\boldsymbol{A}$ has full row rank (i.e. $r = m \leq n$), $\boldsymbol{A}^\dagger = \boldsymbol{A}^*(\boldsymbol{A}\boldsymbol{A}^*)^{-1}$. This aligns with our previous findings for the solution to the linear equation. Therefore, in general, we can say that the minimum norm solution to the linear equation is $\hat{\boldsymbol{x}} = \boldsymbol{A}^\dagger\boldsymbol{y}$. This tells us that the set of solutions is given by

$$S = \{\boldsymbol{A}^\dagger\boldsymbol{y} + \boldsymbol{z} : \boldsymbol{z} \in N(\boldsymbol{A})\}.$$

12.3 Least Squares Solution

Previously we said that a solution to the linear equation $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}$ exists if and only if $\boldsymbol{y} \in R(\boldsymbol{A})$. Now we want to consider the case when $\boldsymbol{y} \notin R(\boldsymbol{A})$. Now there is no $\boldsymbol{x} \in U$ such that $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{y}$, so we aim to find the best possible solution, $\hat{\boldsymbol{x}}$, such that the vector $\hat{\boldsymbol{y}} = \boldsymbol{A}\hat{\boldsymbol{x}}$ is the closest possible vector to $\boldsymbol{y}$ within the range space of $\boldsymbol{A}$. More formally, we want to solve the following problem:

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} ||\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}||_2^2.$$

The solution $\hat{\boldsymbol{x}}$ is referred to as the least-squares (LS) solution. To compute this solution, we first compute the gradient of the objective with respect to $\boldsymbol{x}$:

$$\begin{aligned}
\nabla_{\boldsymbol{x}}||\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}||_2^2 &= \nabla_{\boldsymbol{x}}\Big((\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y})^*(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y})\Big) \\
&= \nabla_{\boldsymbol{x}}\Big(\boldsymbol{x}^*\boldsymbol{A}^*\boldsymbol{A}\boldsymbol{x} - \boldsymbol{x}^*\boldsymbol{A}^*\boldsymbol{y} - \boldsymbol{y}^*\boldsymbol{A}\boldsymbol{x} + \boldsymbol{y}^*\boldsymbol{y}\Big) \\
&= 2\boldsymbol{A}^*\boldsymbol{A}\boldsymbol{x} - 2\boldsymbol{A}^*\boldsymbol{y}
\end{aligned}$$

Plugging in the optimal solution $\hat{\boldsymbol{x}}$ and setting it equal to zero, we get

$$2\boldsymbol{A}^*\boldsymbol{A}\hat{\boldsymbol{x}} - 2\boldsymbol{A}^*\boldsymbol{y} = 0 \implies \boldsymbol{A}^*\boldsymbol{A}\hat{\boldsymbol{x}} = \boldsymbol{A}^*\boldsymbol{y}.$$

By definition, the vector $\boldsymbol{A}^*\boldsymbol{y}$ is in the range space of $\boldsymbol{A}^*$. By the fundamental theorem of linear algebra, the range space of $\boldsymbol{A}^*$ is the same as the range space of $\boldsymbol{A}^*\boldsymbol{A}$. Therefore, the vector $\boldsymbol{A}^*\boldsymbol{y}$ is in the range space of $\boldsymbol{A}^*\boldsymbol{A}$. This tells us that the equation $\boldsymbol{A}^*\boldsymbol{A}\boldsymbol{x} = \boldsymbol{A}^*\boldsymbol{y}$ always has at least one solution.

12.3.1 Injective Map

As we showed previously, if the map $\mathcal{A}$ is injective, then $\boldsymbol{A}^*\boldsymbol{A}$ is invertible. In this case, the least-squares solution is simply

$$\hat{\boldsymbol{x}} = (\boldsymbol{A}^*\boldsymbol{A})^{-1}\boldsymbol{A}^*\boldsymbol{y}.$$

This looks very similar to the solution we found when $\boldsymbol{y}$ was in the range of $\boldsymbol{A}$. The only difference here is that $\hat{\boldsymbol{y}} = \boldsymbol{A}\hat{\boldsymbol{x}}$ is not in the range of $\boldsymbol{A}$.

12.3.2 Surjective Map

If the $\mathcal{A}$ is not injective, then we want to find the minimum norm solution, as we did previously. Again, we should choose $\hat{\boldsymbol{x}}$ to be perpendicular to the null space of $\boldsymbol{A}$. From the fundamental theorem of linear algebra, this implies that there exists a vector $\boldsymbol{v} \in V$ such that $\hat{\boldsymbol{x}} = \boldsymbol{A}^*\boldsymbol{v}$. The minimum norm solution must satisfy the optimality equation we found for the least squares problem, so

$$\boldsymbol{A}^*\boldsymbol{A}\boldsymbol{A}^*\boldsymbol{v} = \boldsymbol{A}^*\boldsymbol{y}.$$

As we showed previously, if the map $\mathcal{A}$ is surjective, then $\boldsymbol{A}\boldsymbol{A}^*$ is invertible. This observation allows us to find an expression for v in terms of y.

$$\boldsymbol{A}\boldsymbol{A}^*\boldsymbol{A}\boldsymbol{A}^*\boldsymbol{v} = \boldsymbol{A}\boldsymbol{A}^*\boldsymbol{y}$$

$$\boldsymbol{v} = (\boldsymbol{A}\boldsymbol{A}^*)^{-1}\boldsymbol{y}$$

Now we can see that if $\mathcal{A}$ is surjective, then the least-squares solution is simply

$$\hat{\boldsymbol{x}} = \boldsymbol{A}^*\boldsymbol{v} = \boldsymbol{A}^*(\boldsymbol{A}\boldsymbol{A}^*)^{-1}\boldsymbol{y}.$$

Once again, the only difference between this solution and the one we found for a surjective map when $\boldsymbol{y}$ was in the range of $\boldsymbol{A}$ is that $\hat{\boldsymbol{y}} = \boldsymbol{A}\hat{\boldsymbol{x}} \notin R(\boldsymbol{A})$.

12.3.3 General Map

Finally, as we stated previously, if $\mathcal{A}$ is neither injective nor surjective, then

$$\hat{\boldsymbol{x}} = \boldsymbol{A}^\dagger\boldsymbol{y}.$$

The set of optimal solutions to the least-squares problem is thus

$$\mathcal{X}_{\text{opt}} = \{\boldsymbol{A}^\dagger\boldsymbol{y} + \boldsymbol{z} : \boldsymbol{z} \in N(A)\}.$$

Any solution within this optimal set, solves the least-squares optimization problem. However, $\hat{\boldsymbol{x}} = \boldsymbol{A}^\dagger\boldsymbol{y}$ is the unique minimum norm solution.

12.4 Weighted Least Squares

As shown previously, the ordinary least squares objective is given by

$$f_{LS} = ||\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}||_2^2.$$

The minimum possible value of f_{LS} is zero, which occurs when $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{y}$. The vector $\boldsymbol{r} = \boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}$ represents residual terms. If $\boldsymbol{a_i}^T$ is the ith row of $\boldsymbol{A}$, then $r_i = \boldsymbol{a_i}^T\boldsymbol{x} - \boldsymbol{y}_i$ is the residual corresponding to the ith element of $\boldsymbol{y}$.

We can give the residuals relative importance by introducing weights into the least squares objective. Suppose we want to assign residual r_i a relative weight of w_i for $i = 1, \ldots, m$. To incorporate these weights, we can define the matrix $\boldsymbol{W} = \text{diag}(w_1, \ldots, w_m)$ and express the weighted least squares objective as

$$f_{WLS} = ||\boldsymbol{W}(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y})||_2^2.$$

If we define the matrix $\boldsymbol{A_W} = \boldsymbol{W}\boldsymbol{A}$ and the vector $\boldsymbol{y_W} = \boldsymbol{W}\boldsymbol{y}$, then we can express this weighted least squares objective as

$$f_{WLS} = ||\boldsymbol{A_W}\boldsymbol{x} - \boldsymbol{y_W}||_2^2.$$

Now that we have put the weighted least squares objective in this form, we can see that its solution must satisfy the following set of equations:

$$\boldsymbol{A}_{\boldsymbol{W}}^*\boldsymbol{A_W}\boldsymbol{x} = \boldsymbol{A}_{\boldsymbol{W}}^*\boldsymbol{y_W}$$

$$(\boldsymbol{W}\boldsymbol{A})^*(\boldsymbol{W}\boldsymbol{A})\hat{\boldsymbol{x}} = (\boldsymbol{W}\boldsymbol{A})^*(\boldsymbol{W}\boldsymbol{y})$$

$$\boldsymbol{A}^*\boldsymbol{W}^*\boldsymbol{W}\boldsymbol{A}\hat{\boldsymbol{x}} = \boldsymbol{A}^*\boldsymbol{W}^*\boldsymbol{W}\boldsymbol{y}$$

If we assume the matrix $\boldsymbol{A}^*\boldsymbol{W}^*\boldsymbol{W}\boldsymbol{A}$ is invertible, then our solution becomes

$$\hat{\boldsymbol{x}} = (\boldsymbol{A}^*\boldsymbol{W}^*\boldsymbol{W}\boldsymbol{A})^{-1}\boldsymbol{A}^*\boldsymbol{W}^*\boldsymbol{W}\boldsymbol{y}.$$

12.5 Regularized Least Squares

As shown previously, the ordinary least squares objective is given by

$$f_{LS} = ||\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}||_2^2.$$

We can add a "regularization" or penalty term $\phi(\boldsymbol{x})$ to the ordinary least squares objective and express the regularized least squares objective as

$$f_{RLS} = ||\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}||_2^2 + \phi(\boldsymbol{x}).$$

Most often, $\phi(\boldsymbol{x})$ is chosen to be proportional to the Euclidean norm of $\boldsymbol{x}$. This type of problem is called the l_2-regularized least squares problem and is often referred to as ridge regression. This objective can be written as

$$f_{RR} = ||\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}||_2^2 + \gamma||\boldsymbol{x}||_2^2, \ \gamma \geq 0.$$

We can write this objective in the form of the ordinary least squares objective by defining the matrix $\tilde{\boldsymbol{A}}$ and vector $\tilde{\boldsymbol{y}}$ such that

$$\tilde{\boldsymbol{A}} = \begin{bmatrix} \boldsymbol{A} \\ \sqrt{\gamma}\boldsymbol{I_n} \end{bmatrix} \text{ and } \tilde{\boldsymbol{y}} = \begin{bmatrix} \boldsymbol{y} \\ \boldsymbol{0}_n \end{bmatrix}.$$

This allows us to express the ridge regression objective as

$$f_{RR} = ||\tilde{\boldsymbol{A}}\boldsymbol{x} - \tilde{\boldsymbol{y}}||_2^2.$$

Now that we have put the ridge regression objective in this form, we can see that its solution must satisfy the following set of equations:

$$\tilde{\boldsymbol{A}}^*\tilde{\boldsymbol{A}}\boldsymbol{x} = \tilde{\boldsymbol{A}}^*\tilde{\boldsymbol{y}}$$

$$\begin{bmatrix} \boldsymbol{A} \\ \sqrt{\gamma}\boldsymbol{I_n} \end{bmatrix}^* \begin{bmatrix} \boldsymbol{A} \\ \sqrt{\gamma}\boldsymbol{I_n} \end{bmatrix} \hat{\boldsymbol{x}} = \begin{bmatrix} \boldsymbol{A} \\ \sqrt{\gamma}\boldsymbol{I_n} \end{bmatrix}^* \begin{bmatrix} \boldsymbol{y} \\ 0_n \end{bmatrix}$$

$$(\boldsymbol{A}^*\boldsymbol{A} + \gamma\boldsymbol{I_n})\hat{\boldsymbol{x}} = \boldsymbol{A}^*\boldsymbol{y}$$

From our discussion of positive semidefinite matrices in section 8.2, if the constant γ is strictly positive, then the matrix $(\boldsymbol{A}^*\boldsymbol{A}+\gamma\boldsymbol{I_n})$ is necessarily invertible. In this case the solution to the ridge regression problem is

$$\hat{\boldsymbol{x}} = (\boldsymbol{A}^*\boldsymbol{A} + \gamma\boldsymbol{I_n})^{-1}\boldsymbol{A}^*\boldsymbol{y}.$$

12.6 Tikhonov Regularization

Tikhonov regularization combines weighted least squares with regularized least squares, giving us the following objective:

$$f_{TR} = ||\boldsymbol{W_1}(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y})||_2^2 + ||\boldsymbol{W_2}(\boldsymbol{x} - \boldsymbol{x}_0)||_2^2,$$

where $\boldsymbol{W_1}, \boldsymbol{W_2} \succ 0$ are weighting matrices and $\boldsymbol{x}_0$ is some nominal value of $\boldsymbol{x}$. For simplicity, let's define $\boldsymbol{W_1}$ as the diagonal matrix $\boldsymbol{W_1} = \text{diag}(w_1, \ldots, w_m)$ and $\boldsymbol{W_2}$ as the diagonal matrix $\boldsymbol{W_2} = \text{diag}(\gamma_1, \ldots, \gamma_n)$. We can write this objective in the form of the ordinary least squares objective by defining

$$\tilde{\boldsymbol{A}} = \begin{bmatrix} \boldsymbol{W_1}\boldsymbol{A} \\ \boldsymbol{W_2} \end{bmatrix} \text{ and } \tilde{\boldsymbol{y}} = \begin{bmatrix} \boldsymbol{W_1}\boldsymbol{y} \\ \boldsymbol{W_2}\boldsymbol{x}_0 \end{bmatrix}.$$

This allows us to express the Tikhonov regularization objective as

$$f_{TR} = ||\tilde{\boldsymbol{A}}\boldsymbol{x} - \tilde{\boldsymbol{y}}||_2^2.$$

Now that we have put the weighted least squares objective in this form, we can see that its solution must satisfy the following set of equations:

$$\tilde{\boldsymbol{A}}^*\tilde{\boldsymbol{A}}\boldsymbol{x} = \tilde{\boldsymbol{A}}^*\tilde{\boldsymbol{y}}$$

$$\begin{bmatrix} \boldsymbol{W_1 A} \\ \boldsymbol{W_2} \end{bmatrix}^* \begin{bmatrix} \boldsymbol{W_1 A} \\ \boldsymbol{W_2} \end{bmatrix} \hat{\boldsymbol{x}} = \begin{bmatrix} \boldsymbol{W_1 A} \\ \boldsymbol{W_2} \end{bmatrix}^* \begin{bmatrix} \boldsymbol{W_1 y} \\ \boldsymbol{W_2 x}_0 \end{bmatrix}$$

$$(\boldsymbol{A}^* \boldsymbol{W_1}^* \boldsymbol{W_1 A} + \boldsymbol{W_2}^* \boldsymbol{W_2}) \hat{\boldsymbol{x}} = \boldsymbol{A}^* \boldsymbol{W_1}^* \boldsymbol{W_1 y} + \boldsymbol{W_2}^* \boldsymbol{W_2 x}_0$$

If we assume $(\boldsymbol{A}^* \boldsymbol{W_1}^* \boldsymbol{W_1 A} + \boldsymbol{W_2}^* \boldsymbol{W_2})$ is invertible, the solution is simply

$$\hat{\boldsymbol{x}} = (\boldsymbol{A}^* \boldsymbol{W_1}^* \boldsymbol{W_1 A} + \boldsymbol{W_2}^* \boldsymbol{W_2})^{-1} (\boldsymbol{A}^* \boldsymbol{W_1}^* \boldsymbol{W_1 y} + \boldsymbol{W_2}^* \boldsymbol{W_2 x}_0).$$