

Inner Products

Definition 14.1 (Inner Product): The inner product between two vectors $\vec{x}, \vec{y} \in \mathbb{R}^n$ is defined as:

$$\langle \vec{x}, \vec{y} \rangle \equiv \vec{x}^T \vec{y} \quad (1)$$

$$= [x_1 \quad x_2 \quad \cdots \quad x_n] \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} \quad (2)$$

$$= x_1 y_1 + x_2 y_2 + \cdots + x_n y_n \quad (3)$$

$$= \sum_{i=1}^n x_i y_i \quad (4)$$

In physics, inner products are often called dot products (denoted by $\vec{x} \cdot \vec{y}$) and can be used to calculate values such as work. In this class we will use the notation $\langle \vec{x}, \vec{y} \rangle$ to denote inner product.

Example 14.1 (Inner product of two vectors): Compute the inner product of the two vectors $[-1 \quad 3.5 \quad 0]^T$ and $[1 \quad 0 \quad 2]^T$.

$$\left\langle \begin{bmatrix} -1 \\ 3.5 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 2 \end{bmatrix} \right\rangle = [-1 \quad 3.5 \quad 0] \begin{bmatrix} 1 \\ 0 \\ 2 \end{bmatrix} \quad (5)$$

$$= -1 \times 1 + 3.5 \times 0 + 0 \times 2 \quad (6)$$

$$= -1 + 0 + 0 = -1. \quad (7)$$

Properties of Inner Products

Symmetry We can prove algebraically that the inner product is a symmetric operation, that is, it remains the same even if we switch its arguments: for any $\vec{x}, \vec{y} \in \mathbb{R}^n$, $\langle \vec{x}, \vec{y} \rangle = \langle \vec{y}, \vec{x} \rangle$,

$$\begin{aligned} \langle \vec{x}, \vec{y} \rangle &= x_1 y_1 + \cdots + x_n y_n \\ &= y_1 x_1 + \cdots + y_n x_n \\ &= \langle \vec{y}, \vec{x} \rangle \end{aligned} \quad (8)$$

Homogeneity We shall see how scaling affects inner products. If c is a real number, algebraically:

$$\begin{aligned}\langle c\vec{x}, \vec{y} \rangle &= (cx_1)y_1 + \cdots + (cx_n)y_n \\ &= c(x_1y_1) + \cdots + c(x_ny_n) \\ &= c \langle \vec{x}, \vec{y} \rangle\end{aligned}\tag{9}$$

By a similar argument (exercise: prove it yourself),

$$\langle \vec{x}, c\vec{y} \rangle = c \langle \vec{x}, \vec{y} \rangle\tag{10}$$

This means the scaling any vector in the inner product by any real constant c will scale the inner product by the same constant.

Additivity What happens when we take the inner product between a sum of two vectors and another vector?

We can write:

$$\begin{aligned}\langle \vec{x} + \vec{y}, \vec{z} \rangle &= (x_1 + y_1)z_1 + \cdots + (x_n + y_n)z_n \\ &= x_1z_1 + y_1z_1 + \cdots + x_nz_n + y_nz_n \\ &= x_1z_1 + \cdots + x_nz_n + y_1z_1 + \cdots + y_nz_n \\ &= \langle \vec{x}, \vec{z} \rangle + \langle \vec{y}, \vec{z} \rangle\end{aligned}\tag{11}$$

Similarly:

$$\langle \vec{x}, \vec{y} + \vec{z} \rangle = \langle \vec{x}, \vec{y} \rangle + \langle \vec{x}, \vec{z} \rangle\tag{12}$$

Homogeneity together with additivity gives rise to a property called *Bilinearity*. This means that the inner product is a function that is linear in each argument. These properties are also very useful in proving properties about inner products.

Orthogonal Vectors

Two vectors \vec{x}, \vec{y} are said to be **orthogonal** if their inner product is equal to zero, i.e., $\langle \vec{x}, \vec{y} \rangle = 0$. In 2-D and 3-D coordinate spaces, perpendicular and orthogonal mean the same thing; in higher-dimension spaces, it is harder to visualize vectors being perpendicular, so the term orthogonal comes in handy to abstract away the need to visualize.

Example 14.2 (Orthogonal vectors): We have the two vectors $\vec{x} = \begin{bmatrix} 1 \\ 1 \\ 3 \end{bmatrix}$ and $\vec{y} = \begin{bmatrix} 4 \\ -1 \\ -1 \end{bmatrix}$ in \mathbb{R}^3 . Are they orthogonal to each other? We compute the inner product of the two vectors: $\langle \vec{x}, \vec{y} \rangle = (1)(4) + (1)(-1) + (3)(-1) = 0$. Since their inner product is equal to zero, by definition, \vec{x} and \vec{y} are orthogonal.

A very important fact is that the standard unit vectors are always orthogonal to each other. For example, in \mathbb{R}^2 , the standard unit vectors are $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$.

Special Vector Operations Using Inner Products

The inner product is a basic building block for many operations. In programming terms, if you have a black box function `INNERPRODUCT`, here are some useful operations you can form by controlling what vectors you input into the function.

Sum of Components

$$\left\langle [1 \ 1 \ \dots \ 1]^T, [x_0 \ x_1 \ \dots \ x_n]^T \right\rangle = x_0 + x_1 + \dots + x_n \quad (13)$$

Average

$$\left\langle \left[\frac{1}{n} \ \frac{1}{n} \ \dots \ \frac{1}{n} \right]^T, [x_0 \ x_1 \ \dots \ x_n]^T \right\rangle = \frac{x_0 + x_1 + \dots + x_n}{n} \quad (14)$$

Sum of Squares

$$\left\langle [x_0 \ x_1 \ \dots \ x_n]^T, [x_0 \ x_1 \ \dots \ x_n]^T \right\rangle = x_0^2 + x_1^2 + \dots + x_n^2 \quad (15)$$

Selective Sum Here, the values of the first vector has a set of 0's and 1's, which correspond to which values of the variable matrix (the vertical one) are to be used, and which ones are to be thrown away for the sum. This, for example, can take the place of a for loop checking every element of a vector.

$$\left\langle [0 \ 0 \ 1 \ 0 \ 1 \ \dots \ 0 \ 1]^T, [x_0 \ x_1 \ \dots \ x_n]^T \right\rangle = x_2 + x_4 + \dots + x_n \quad (16)$$

These things are often important in computer programming contexts because computers (and programming languages) are often optimized to be able to do vector operations like inner products very fast. So, seeing an inner-product way of representing something can often speed up calculations considerably.

Norms

The **Euclidean Norm** of a vector is given by:

$$\|\vec{x}\|_2 = \sqrt{x_1^2 + x_2^2 + \dots + x_n^2} = \sqrt{\langle \vec{x}, \vec{x} \rangle} \quad (17)$$

This corresponds to the usual notion of distance in \mathbb{R}^2 or \mathbb{R}^3 . It is interesting to note that the set of points with equal Euclidean norm is a circle in \mathbb{R}^2 , or a sphere in \mathbb{R}^3 .

You may have noticed that the subscript 2 in the definition of the norm given above. The subscript differentiates the Euclidean norm (or 2-norm) from other useful norms. In general, the p-norm is defined as:

$$\|\vec{x}\|_p = (x_1^p + x_2^p + \dots + x_n^p)^{1/p} \quad (18)$$

These other norms might feel esoteric but turn out to be spectacularly useful in many engineering settings. Follow-on courses like EE127 will show you how these can be useful in applications like machine learning.

Why is the norm important? The 2-norm of a vector is also the magnitude of the vector (or length of the vector).

Properties of Norms

Absolute Homogeneity For any scalar $\alpha \in \mathbb{R}$ and $\vec{x} \in \mathbb{R}^n$,

$$\|\alpha\vec{x}\| = |\alpha| \|\vec{x}\|. \quad (19)$$

Triangle Inequality For any two vectors $\vec{x}, \vec{y} \in \mathbb{R}^n$,

$$\|\vec{x} + \vec{y}\| \leq \|\vec{x}\| + \|\vec{y}\|. \quad (20)$$

Non-negativity For any vector $\vec{x} \in \mathbb{R}^n$,

$$\|\vec{x}\| \geq 0. \quad (21)$$

In addition, $\|\vec{x}\| = 0$ if and only if $\vec{x} = \vec{0}$.


Interpretations of Inner Products

Now that we have defined inner products, it's time to get an intuition about what an inner product is.

First we take the unit vector $\vec{e}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and a general unit vector in \mathbb{R}^2 , $\vec{x} = \begin{bmatrix} \cos \alpha \\ \sin \alpha \end{bmatrix}$. \vec{x} is a unit vector because $\|\vec{x}\| = \sqrt{\cos^2 \alpha + \sin^2 \alpha} = 1$.

When we draw the vectors, the angle between them is clearly α . We can calculate $\langle \vec{e}_1, \vec{x} \rangle = 1 \times \cos \alpha + 0 \times \sin \alpha = \cos \alpha$. With this we might guess that the inner product between two unit vectors is the cosine of the angle between them.

Let's try to prove this for two general unit vectors in \mathbb{R}^2 , $\vec{x} = \begin{bmatrix} \cos \alpha \\ \sin \alpha \end{bmatrix}$ and $\vec{y} = \begin{bmatrix} \cos \beta \\ \sin \beta \end{bmatrix}$.



$$\|\vec{x}\| = 1, \quad \|\vec{y}\| = 1$$

Using the trigonometric identity $\cos \alpha \cos \beta + \sin \alpha \sin \beta = \cos(\alpha - \beta)$, we find:

$$\begin{aligned}\langle \vec{x}, \vec{y} \rangle &= \cos \alpha \cos \beta + \sin \alpha \sin \beta \\ &= \cos(\alpha - \beta) \\ &= \cos \theta\end{aligned}\tag{22}$$

Again, that is the cosine of the angle between these two vectors!

What about any two general vectors \vec{x} and \vec{y} ? We can first convert them into the unit vectors $\vec{x}/\|\vec{x}\|$ and $\vec{y}/\|\vec{y}\|$ by dividing each vector by its norm. Then we can use (??) to get:

$$\left\langle \frac{\vec{x}}{\|\vec{x}\|}, \frac{\vec{y}}{\|\vec{y}\|} \right\rangle = \cos \theta\tag{23}$$

We can then use the homogeneity property (??) (since both norms are scalars):

$$\begin{aligned}\frac{\langle \vec{x}, \vec{y} \rangle}{\|\vec{x}\| \|\vec{y}\|} &= \cos \theta \\ \langle \vec{x}, \vec{y} \rangle &= \|\vec{x}\| \|\vec{y}\| \cos \theta\end{aligned}\tag{24}$$

Now we have a geometric interpretation of the inner product: the inner product of vectors \vec{x} and \vec{y} is their lengths multiplied by the cosine of the angle between them. One remarkable observation is that the inner product *does not depend on the coordinate system the vectors are in*, it only depends on the relative angle between these vectors and their length. This is the reason it is very useful in physics, where the physical laws do not depend on the coordinate system used to measure them. This is also the reason why this property holds in higher dimensions as well. For any two vectors we can look at the plane passing through them and the angle between them is the angle θ measured in the plane.


The Cauchy-Schwarz Inequality

Now we are ready to prove the Cauchy-Schwarz inequality, which relates the inner product of two vectors to their length: for any two vectors $\vec{x}, \vec{y} \in \mathbb{R}^n$, $|\langle \vec{x}, \vec{y} \rangle| \leq \|\vec{x}\| \|\vec{y}\|$. To prove this, we will use the fact that $|\cos \theta| \leq 1$ for all real θ ,

$$\begin{aligned}|\langle \vec{x}, \vec{y} \rangle| &= \|\vec{x}\| \|\vec{y}\| |\cos \theta| \\ &= \|\vec{x}\| \|\vec{y}\| \cdot |\cos \theta| \\ &\leq \|\vec{x}\| \|\vec{y}\|\end{aligned}\tag{25}$$

Projections

Knowing that the inner products of two vectors in \mathbb{R}^n is their lengths multiplied by the angle between them, we can write the *projection* of a vector onto another using the inner product.



The projection of \vec{y} onto \vec{x} is the component of \vec{y} lying in the direction of \vec{x} . First let's find the length of this component. By simple geometry, this is $\|\vec{y}\| \cos \theta$, where θ is the angle between \vec{y} and \vec{x} . We can write it as $\langle \vec{y}, \vec{u} \rangle$, where $\vec{u} = \vec{x} / \|\vec{x}\|$ is the unit vector in the direction of \vec{x} . Now we know the length of the projection, all we need is a direction. This is just \vec{u} , so the projection of \vec{y} onto \vec{x} is given by:

$$\langle \vec{y}, \vec{u} \rangle \vec{u} = \left\langle \vec{y}, \frac{\vec{x}}{\|\vec{x}\|} \right\rangle \frac{\vec{x}}{\|\vec{x}\|} \quad (26)$$