Lecture 18: Inner Product Spaces

March 19, 2024

Lecture 18: Inner Product Spaces

Inner Products

In $\mathbb{R}^2$, the scalar (or dot) product of the vectors $x = \binom{x_1}{x_2}$ and $y = \binom{y_1}{y_2}$ is given by

$$x \cdot y = x_1 y_1 + x_2 y_2 = x^T y = y^T x = y \cdot x.$$

We can interpret the *length* $||x||$ of the vector x as the length of the directed line segment from the origin to (x_1, x_2), which by the Theorem of Pythagoras is $\sqrt{x_1^2 + x_2^2}$ or $\sqrt{x \cdot x}$.

Once we have a concept of length of a vector, we can define the *distance* $d(x, y)$ between two vectors x and y as the length of their difference: $d(x, y) = ||x - y||$.

Inner Products

In $\mathbb{R}^2$, the scalar (or dot) product of the vectors $x = \binom{x_1}{x_2}$ and $y = \binom{y_1}{y_2}$ is given by

$$x \cdot y = x_1 y_1 + x_2 y_2 = x^T y = y^T x = y \cdot x.$$

Similarly, from the Cosine Rule we can observe that $x \cdot y = ||x||\,||y|| \cos\theta$, where θ is the angle between the directed line segments representing x and y. In particular, x is orthogonal to y (or $x \perp y$) if and only if $x \cdot y = 0$.

So the scalar product encodes geometric information in $\mathbb{R}^2$, and it also provides a mechanism for defining concepts of length, distance and orthogonality on real vector spaces that do not necessarily have an obvious geometric structure.
The scalar product is an example of an inner product.

Real Inner Products

An inner product on a vector space V is a function from $V \times V$ to $\mathbb{R}$ that assigns an element of $\mathbb{R}$ to every ordered pair of elements of V, and has the following properties.

1. Symmetry: $\langle x, y\rangle = \langle y, x\rangle$ for all $x, y \in V$
2. Linearity in both slots (bilinearity): For all $x, y, z \in V$ and all $a, b \in \mathbb{R}$, we have $\langle ax + by, z\rangle = a\langle x, z\rangle + b\langle y, z\rangle$ and $\langle x, ay + bz\rangle = a\langle x, y\rangle + b\langle x, z\rangle$.
3. Non-negativity: $\langle x, x\rangle \geq 0$ for all $x \in V$, and $\langle x, x\rangle = 0$ only if $x = 0_V$.

The ordinary scalar product on $\mathbb{R}^n$ is the best known example of an inner product.

Examples of inner products

1. The ordinary scalar product on $\mathbb{R}^n$.
2. Let C be the vector space of all continuous real-valued functions on the interval $[0, 1]$. The analogue of the ordinary scalar product on C is the inner product given by
$$\langle f, g\rangle = \int_0^1 f(x)g(x)\,dx, \text{ for } f, g \in C.$$
3. On the space $M_{m\times n}(\mathbb{R})$, the *Frobenius inner product* or *trace inner product* is defined by $\langle A, B\rangle = \text{trace}(A^T B)$. Note that trace$A^T B$ is the sum over all positions (i, j) of the products $A_{ij}B_{ij}$. So this is closely related to the ordinary scalar product, if the matrices A and B were regarded as vectors with mn entries over $\mathbb{R}$.

It is possible for a single vector space to have many different inner products defined on it, and if there is any risk of ambiguity we need to specify which one we are considering.

Length, Distance and Orthogonality

Given a real vector space and equipped with an inner product $\langle \cdot, \cdot, \rangle$, we make the following two definitions.

Definition We define the length or norm of any vector v by

$$||v|| = \sqrt{\langle v, v \rangle},$$

and we define the distance between the vectors u and v by

$$d(u, v) = ||u - v||.$$

Definition We say that vectors u and v are orthogonal (with respect to $\langle \cdot, \cdot, \rangle$) if $\langle u, v \rangle = 0$.

These definitions are consistent with "typical" geometrically motivated concepts of distance and orthogonality.

Unit Vectors and Scaling

An element v of V is referred to as a unit vector if $||v|| = 1$.
The norm of elements of V has the property that $||cv|| = |c|\,||v||$ for any vector v and real scalar c. To see this we can note that

$$||cv|| = \sqrt{\langle cv, cv\rangle} = \sqrt{c^2\langle v, v\rangle} = |c|\,||v||.$$

So we can adjust the norm of any element of V, while preserving its direction, by multiplying it by a positive scalar.

Definition If v is a non-zero vector in an inner product space V, then

$$\hat{v} := \frac{1}{||v||}v$$

is a unit vector in the same direction as v, referred to as the *normalization* of v.