norm

Authors

Alexander Jung[*], Department of Computer Science, Aalto University, Espoo, Finland
[*]Corresponding author: alex.jung@aalto.fi

Abstract

A norm on a vector space is a function that assigns each vector a nonnegative number, measuring its length. A norm must be definite and homogeneous, and it must satisfy the triangle inequality. Every norm defines a metric via the norm of the difference of two vectors, and every inner product induces a norm. A prominent family of norms in machine learning (ML) is the ℓ_p-norms, including the ℓ_1-norm, the Euclidean norm, and the ℓ_∞-norm. Norms are used, for instance, to define a loss function or a regularizer.

Definition

A norm is a function that maps each (vector) element of a vector space to a nonnegative real number. Formally, a norm $\|\cdot\|$ on a vector space $\mathcal{V}$ over a field $\mathbb{F}$ is a function $\|\cdot\| : \mathcal{V} \to \mathbb{R}_+$ that satisfies the following conditions (Horn and Johnson, 2013) for all $\mathbf{u}, \mathbf{v} \in \mathcal{V}$ and $\beta \in \mathbb{F}$:

- Definiteness: $\|\mathbf{u}\| = 0 \iff \mathbf{u} = \mathbf{0}$;
- Homogeneity: $\|\beta\mathbf{u}\| = |\beta| \cdot \|\mathbf{u}\|$;
- Triangle inequality: $\|\mathbf{u} + \mathbf{v}\| \leq \|\mathbf{u}\| + \|\mathbf{v}\|$.

Intuitively, norms measure the length or size of a vector. Norms can also be used to measure distances between vectors, as they define a metric by $d(\mathbf{u}, \mathbf{v}) := \|\mathbf{u} - \mathbf{v}\|$, making $(\mathcal{V}, d(\cdot, \cdot))$ a metric space. On an inner product space, the inner product induces a norm via $\|\mathbf{u}\| := \sqrt{\langle \mathbf{u}, \mathbf{u} \rangle}$; this is the norm of a Hilbert space (see inner product, Hilbert space). A prominent family of norms used in machine learning (ML) is the ℓ_p-norms for $p \geq 1$, defined for a vector $\mathbf{x} \in \mathbb{R}^d$ as $\|\mathbf{x}\|_p = \left(\sum_{j=1}^{d} |x_j|^p \right)^{1/p}$. Important instances include the ℓ_1-norm, the ℓ_2-norm (or Euclidean norm), and the ℓ_∞-norm $\|\mathbf{x}\|_\infty = \max_{j=1,\ldots,d} |x_j|$, which are illustrated by the geometry of their respective unit spheres in Fig. 1. In ML, norms are fundamental, for instance for defining a loss function or a regularizer.

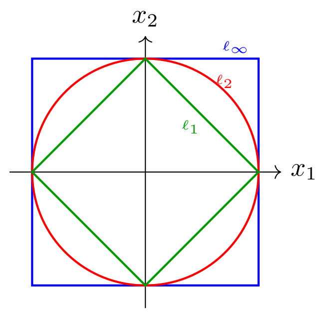

Figure 1: Unit spheres $\mathbb{S}^{(1)}$ with respect to different ℓ_p-norms for $p = 1,\ 2,\ \infty$

Cross References

vector space, metric, metric space, Euclidean space, Hilbert space, inner product, linear regression, Lasso

References

1. Horn and Johnson (2013). Matrix Analysis. 2nd ed., Cambridge Univ. Press, New York, NY, USA.