International Workshop on Computer Vision, April 26-30, 2004, IPM, Tehran

Shape Description via Persistent Homology

A. Zomorodian
Department of Computer Science
Stanford University
USA

This talk is about a blending of techniques from geometry and topology for shape description and classification. One can view geometry as the finest level of classification as it focuses on the local properties of shapes. In this sense, geometry has a quantitative nature and can answer low level questions about a shape. But most of our questions about shapes have qualitative feel and take a higher view of a shape. This prompts us to look at topological techniques which classify shapes according to the way they are connected globally – their connectivity. Unfortunately, this view is often too coarse to be useful. For example, the topological invariant homology classifies shapes according to the number of components, tunnels, and enclosed spaces. This classification cannot distinguish between circles and ellipses, between circles and rectangles, or even between Euclidean spaces of different dimensions. Further, it cannot identify sharp features, such as corners, edges, or cone points: their neighborhoods are all homeomorphic or connected the same way. Whereas geometry is too sensitive for shape description, topology seems to be too insensitive.

In this talk, I focus on an approach that combines the differentiating power of geometry with the classification power of topology. The key idea is to apply homology not to the space itself, but to derived spaces with richer geometric content. In this case, we construct spaces using the tangential information of the shape. This method gives us a simple and compact shape descriptor, a combinatorial invariant called a "barcode". We define a metric over the set of barcodes that allows us to classify and compare shapes.

I will spend most of the talk describing the geometric and topological techniques and verifying the effectiveness of the approach using explicit calculations on known families of mathematical shapes. I will then discuss applying the method to finite samples of shapes or point cloud data (PCD). After describing a computational pipeline for computing barcodes, I conclude the talk with some examples of computing barcodes for curves, including scanned letters and noisy data.