Averaging on Riemannian Manifolds: Theoretical and Computational Perspectives

Dr. Bijan Afsari
(Johns Hopkins University)

December 20, 2013

In many data analysis applications one needs to do statistical analysis for data points that take values in a non-Euclidean space, i.e., where addition and multiplication by scalars are not well defined operations. Quite often the space can be given the structure of a Riemannian manifold. Some widely encountered examples include the unit sphere in $\mathbb{R}^n$, the rotation group in $\mathbb{R}^3$, and the Grassmann manifold of k-dimensional subspaces of $\mathbb{R}^n$. In this talk, after some historical remarks, we will look at the most basic statistical tool in the toolbox of nonparametric statistics, namely mean or average. On Riemannian manifolds there are several possible (not necessarily equivalent) definitions. The most standard one is based on minimizing the sum of the square of *geodesic distances*, which has been referred to by names such as Riemannian (or Fréchet) mean, average, or center of mass. As a preparation, we briefly review the convexity and differentiability properties of the Riemannian distance function and see how these properties are related to the topology and curvature of the manifold. We then study some basic properties (e.g., existence, uniqueness, convexity, etc.) of Riemannian averages and examine the role of topology and curvature. We also study some computational aspects of finding the Riemannian average of a set of data points and investigate the challenges involved.