Abstract

Whether modeling population dynamics across countries, tracking growth patterns of children, or analyzing daily temperature curves in climate research, data measured continuously over time or space are becoming increasingly central to modern statistics. Functional data analysis (FDA) provides the tools to treat such data as realizations of continuous stochastic processes. Many methods in this field rely on a suitable notion of dissimilarity between functional observations – particularly when analyzing structural variation, detecting anomalies, or grouping similar behaviors.

In multivariate statistics, the Mahalanobis distance is a widely used measure of dissimilarity. By incorporating the covariance structure of the data, it offers a correlation-sensitive way to quantify similarity. Extending this concept to the functional domain introduces unique challenges due to the infinite-dimensional nature of functional data and the need for appropriate regularization. A meaningful extension must also be flexible enough to adapt to different types of functional data and analysis goals.

This thesis provides a unifying framework based on an existing functional regularized Mahalanobis distance, by extending and adapting it to several important settings in FDA. A key contribution is the development of a robust covariance estimator, used for outlier detection of univariate functional data. Next, this distance is extended to the Bayes space, enabling a new approach to robust functional principal component analysis (FPCA) of relative functional data. Finally, the metric is generalized to multivariate functional processes with a separable covariance structure and tested in the framework of clustering. Together, these contributions demonstrate diverse applicability of the functional Mahalanobis distance, highlighted by its strong performance during simulation studies as well as real-world examples.