Title:
Capturing Heteroscedasticity and Long-Range Dependencies with Gaussian Processes

Abstract:
Many inferential tasks, such as analyzing the functional connectivity of the brain via coactivation patterns or capturing the changing correlations amongst a set of geographically-distinct flu estimates, rely on modeling a covariance matrix whose elements evolve as a function of time or some other covariate. The limited set of existing multivariate heteroscedastic methods struggle to cope with three key challenges common to many datasets: (i) high dimensionality, (ii) unequally spaced observations, and (iii) missing values in the observation vector. Another key challenge in modeling such datasets is capturing potentially long-range, non-Markovian dependencies in the individual time series, especially in the presence of abrupt changes.

We first focus on developing a class of nonparametric covariance regression models, which allow an unknown p x p covariance matrix to change flexibly with predictors (e.g., time, space, categories, etc.). To cope with the dimensionality of the data, the framework harnesses a latent factor model representation with predictor-dependent factor loadings modeled as a sparse combination of Gaussian process random functions. Our prior specification leads to a highly-flexible, but computationally tractable formulation with simple conjugate posterior updates that can readily handle missing data and unequally spaced observations.

We then turn to the problem of modeling long-range dependencies and abrupt changes in the mean regression. We propose a multiresolution GP that hierarchically couples a collection of smooth GPs, each defined over an element of a random nested partition. Long-range dependencies are captured by the top-level GP while the partition points define the abrupt changes in the time series. The inherent conjugacy of the GPs allows for efficient inference of the hierarchical partition, for which we employ graph-theoretic techniques.

Theoretical properties are discussed and the methods are illustrated through an application to the Google Flu Trends dataset and the task of word classification based on single-trial Magnetoencephalography (MEG) recordings of brain activity.

Joint work with David Dunson, Alona Fyshe, and Tom Mitchell.