Fast algorithms for Gaussian processes

Michael O'Neil
New York University, USA
Email: oneil@cims.nyu.edu

Gaussian processes are one class of stochastic process that provides a straightforward generalization of random variables to random functions. These random functions are specified via a pointwise multivariate normal distribution whose variance structure is described using a *covariance kernel*. Although their use in mathematical statistics, machine learning, and data analysis for regression and classification problems has been ubiquitous, their computational cost has been dominated by dense linear algebra operations (inverses, applications, and determinants) which scale as $\mathcal{O}(n^2)$ or $\mathcal{O}(n^3)$, where n is the number of pointwise evaluations (or observed evaluations) of the random function. However, by drawing on many ideas central to fast multipole methods, it is possible to design schemes which scale as $\mathcal{O}(n \log n)$ or $\mathcal{O}(n)$ for many widely-used covariance kernels. In this talk, we will give a brief overview of Gaussian processes, their many applications, and demonstrate the performance of these fast schemes via several numerical examples.