THE UNIVERSITY OF CHICAGO

COMPUTATIONAL AND APPLIED MATHEMATICS COLLOQUIUM

DAVID BINDEL
Department of Computer Science
Cornell University

Stochastic Linear Algebra for Scalable Gaussian Processes

THURSDAY, January 10, 2018, at 4:00 PM
Jones 226, 5747 South Ellis Avenue

ABSTRACT

Gaussian processes (GPs) define a distribution over functions that generalizes the multivariate normal distribution over vector spaces. Long used as a tool for spatio-temporal statistical modeling, GPs are also a key part of the modern arsenal in machine learning. Unfortunately, Gaussian process regression and kernel hyper-parameter estimation with N training examples involve manipulating a dense N-by-N kernel matrix, and standard factorization-based approaches to the underlying linear algebra problems have $O(N^3)$ scaling. For regression with a fixed covariance kernel, more scalable iterative methods based on fast matrix-vector multiplication with the kernel matrices are available. However, maximum likelihood estimation of kernel hyper-parameters and computation of conditional variances involve operations such as computing log derivatives and their derivatives or extracting the diagonal part of a Schur complement. New tools are needed to address these problems in a scalable manner. In this talk, we discuss our recent work on one such set of tools, based on a combination of Krylov subspace methods for matrix solves

Joint work with Kun Dong, David Eriksson, and Andrew Wilson.

Organizers:
Daniel Sanz-Alonso, Department of Statistics, sanzalonso@uchicago.edu
CAM Colloquium URL: https://cam.uchicago.edu/seminars/colloq/index.shtml.

For further information and inquiries about building access for persons with disabilities, please contact Zellencia Harris, zellenciah@uchicago.edu. If you wish to subscribe to our email list, please visit the following website: https://lists.uchicago.edu/web/subscribe/cam_colloquium/.