The University of Chicago
Department of Statistics

Seminar Series

SEAN MEYN

Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign

Average-cost Temporal Difference Learning and Adaptive Control Variates

MONDAY, October 6, 2008 at 4:00 PM
133 Eckhart Hall, 5734 S. University Avenue

Refreshments following the seminar in Eckhart 110.

ABSTRACT

Value functions arise in applications to decision and control theory, and also in simulation of Markov models. A barrier to application is complexity. Temporal difference learning in its standard TD(λ) formulation can almost compute the minimum mean-square optimal value function approximation within a prescribed finite dimensional class. The fit improves as the parameter λ approaches unity, but unfortunately the variance of the estimates approaches infinity simultaneously. We re-interpret TD(λ) learning as a simulation based solution to an adjoint equation. Based on this interpretation we obtain a new algorithm that exactly solves the minimum mean-square approximation problem, and various refinements. Each of these algorithms is convergent, and a finite bound on its variance can be obtained based on a Foster-Lyapunov drift condition. Applications to stochastic networks and to simulation are also presented.

Based on Chapter 11 of
S. P. Meyn. Control techniques for complex networks. Cambridge University Press, Cambridge, 2008 http://black.csl.uiuc.edu/ meyn/pages/CTCN/CTCN.html.

Please send email to Mathias Drton (drton@galton.uchicago.edu) for further information. Information about building access for persons with disabilities may be obtained in advance by calling Kathryn Kraynik or by email (kraynik@galton.uchicago.edu).