



Byron Boots
Friday, March 28, 2014
10:30 am
Lindley Hall, Rm. 102

Machine Learning For Modeling Real-World Dynamical Systems

Abstract: A major challenge in machine learning is to reliably and automatically discover hidden structure in high-dimensional data. This is an especially formidable problem for sequential data: revealing the dynamical system that governs a complex time series is often not just difficult, but provably intractable. Popular maximum likelihood strategies for learning dynamical system models are slow in practice and often get stuck at poor local optima, problems that greatly limit the utility of these techniques when learning from real-world data. Although these drawbacks were long thought to be unavoidable, recent work has shown that progress can be made by shifting the focus of learning to realistic instances that rule out the intractable cases.

In this talk, I will present a new family of computational approaches for learning dynamical system models. The key insight is that low-order moments of observed data often possess structure that can be revealed by powerful spectral decomposition methods, and, from this structure, model parameters can be directly recovered. Based on this insight, we design highly effective algorithms for learning parametric models like Kalman Filters and Hidden Markov Models, as well as an expressive new class of nonparametric models via reproducing kernels. Unlike maximum likelihood-based approaches, these new learning algorithms are statistically consistent, computationally efficient, and easy to implement using established matrix-algebra techniques. The result is a powerful framework for learning dynamical system models with state-of-the-art performance on video, robotics, and biological modeling problems.

Biography: Byron Boots is a postdoctoral researcher working with Dieter Fox in the Robotics and State Estimation Lab at the University of Washington. He received his Ph.D. in Machine Learning from Carnegie Mellon University in 2012 where he was advised by Geoffrey Gordon. Byron's work on learning models of dynamical systems received the 2010 Best Paper award at the International Conference on Machine Learning (ICML-2010). His research focuses on modeling and control problems at the intersection of statistical machine learning, artificial intelligence, and robotics.

