

Hierarchical Nonparametric Bayes

Michael Jordan

University of California, Berkeley

Much statistical inference is concerned with controlling some form of tradeoff between flexibility and variability. In Bayesian inference, such control is often exerted via hierarchies—stochastic relationships among prior distributions. Nonparametric Bayesian statisticians work with priors that are general stochastic processes (e.g., distributions on spaces of continuous functions, spaces of monotone functions, or general measure spaces). Thus flexibility is emphasized and it is of particular importance to exert hierarchical control. In this talk I discuss Bayesian hierarchical modeling in the setting of two particularly interesting stochastic processes: the Dirichlet process and the beta process. These processes are discrete with probability one, and have interesting relationships to various random combinatorial objects. They yield models with open-ended numbers of "clusters" and models with open-ended numbers of "features," respectively. I discuss Bayesian modeling based on hierarchical Dirichlet process priors and hierarchical beta process priors, and present applications of these models to problems in bioinformatics and computational vision.

[Joint work with Yee Whye Teh and Romain Thibaux.]