

Bayesian Hidden Markov Models and Extensions

Invited Talk

Zoubin Ghahramani

Engineering Department, University of Cambridge, Cambridge, UK

zoubin@eng.cam.ac.uk

Hidden Markov models (HMMs) are one of the cornerstones of time-series modelling. I will review HMMs, motivations for Bayesian approaches to inference in them, and our work on variational Bayesian learning. I will then focus on recent nonparametric extensions to HMMs. Traditionally, HMMs have a known structure with a fixed number of states and are trained using maximum likelihood techniques. The infinite HMM (iHMM) allows a potentially unbounded number of hidden states, letting the model use as many states as it needs for the data. The recent development of 'Beam Sampling' — an efficient inference algorithm for iHMMs based on dynamic programming — makes it possible to apply iHMMs to large problems. I will show some applications of iHMMs to unsupervised POS tagging and experiments with parallel and distributed implementations. I will also describe a factorial generalisation of the iHMM which makes it possible to have an unbounded number of binary state variables, and can be thought of as a time-series generalisation of the Indian buffet process. I will conclude with thoughts on future directions in Bayesian modelling of sequential data.