Max Welling
University of Amsterdam, The Netherlands

Lecture 1: Approximate Bayesian Posterior Inference for Big Data.

In this lecture I will explain the three main inference algorithms for probabilistic graphical models: Markov Chain Monte Carlo (MCMC), Expectation Propagation (EP) and Variational Bayesian Inference (VB). Besides explaining the algorithms, I will emphasize how to scale these algorithms up to very large datasets (Big Data) through stochastic gradient updates and parallelization.

Lecture 2: A Unifying Framework for Deep Learning, Graphical Models and Bayesian Estimation

In this lecture I will explain how deep learning and graphical models can be unified into a single powerful modeling framework, which we call Variational Auto-Encoders (VAE). In VAEs we have a generative model (e.g. a Bayesian networks or a probabilistic program) and a discriminative network (a deep mixture density network). Both models are trained jointly using variational expectation maximization. I will also discuss fully Bayesian versions of the VAE where also the parameters are treated as random variables.