
Bayesian Inference in Complex Generative Models*

Max Welling

Abstract

In a time when deep learning and big data claim the center stage in machine learning, generative models and Bayesian inference have moved somewhat out of the limelight. I argue that generative modeling and Bayesian learning will remain key for many exciting applications. I present three advances in this direction developed in my group: 1) a new variational learning algorithm for the Helmholtz machine (the variational auto-encoder) applied to semi-supervised learning, 2) a new large-scale distributed posterior MCMC sampling procedure applied to Matrix Factorization and 3) an efficient posterior MCMC sampling algorithms for complex, likelihood free simulator models. I will conclude with a reflection on what is still missing to build truly large scale, distributed and privacy preserving learning systems.

Joint work with Ted Meeds, Durk Kingma and Sungjin Ahn.

*Machine Learning External Seminar, Gatsby Unit, April 15, 2015.