



Statistical Sciences
UNIVERSITY OF TORONTO

SEMINAR

October 6, 2016 at 3:30pm

Refreshments will be provided at 3:15pm

Sidney Smith Hall, Room 2108

Speaker: Roger Grosse, University of Toronto

Host: Stanislav Volgushev

Sandwiching the marginal likelihood using bidirectional Monte Carlo

Computing the marginal likelihood (ML) of a model requires marginalizing out all of the parameters and latent variables, a difficult high-dimensional summation or integration problem. To make matters worse, it is often hard to measure the accuracy of one's ML estimates. I will present bidirectional Monte Carlo, a technique for obtaining accurate log-ML estimates on data simulated from a model.

This method obtains stochastic lower bounds on the log-ML using annealed importance sampling or sequential Monte Carlo, and obtains stochastic upper bounds by running these same algorithms in reverse starting from an exact posterior sample. The true value can be

sandwiched between these two stochastic bounds with high probability.

Using the ground truth log-ML estimates obtained from our method, we quantitatively evaluate a wide variety of existing ML estimators on several latent variable models: clustering, a low rank approximation, and a binary attributes model. These experiments yield insights into how to accurately estimate marginal likelihoods.