samsi
NSF·Duke·NCSU·UNC

Deep Learning Program Opening Workshop
August 12-16, 2019

SPEAKER TITLES/ABSTRACTS

Deborshee Sen
Duke University

"Neural Network Density Estimation"

In conducting non-linear dimensionality reduction and feature learning, it is common to suppose that the data lie near a lower-dimensional manifold. However, there are very few model-based approaches for density estimation that can accommodate such structure. One such class includes latent variables in an unknown non-linear regression function; this includes Gaussian process latent variable models (GP-LVMs) and variational auto-encoders (VAEs) as special cases. VAEs are similar to GP-LVMs, but instead of using a GP to model the unknown regression function, one uses neural networks and additionally employs approximations to make the computation tractable. Current implementations of such frameworks lack adequate uncertainty quantification in estimating the unknown density and the lower-dimensional subspace, and can be unstable and lack reproducibility in practice. We attempt to solve this problem by designing Markov chain Monte Carlo (MCMC) sampling algorithms for fully Bayesian inferences in neural network models with latent variables. Most sampling algorithms tend to have very poor mixing â€œoff-the-shelfâ€• for non-linear regression with latent variables, but building on Hamiltonian Monte Carlo (HMC) approaches, we develop efficient algorithms that are shown to produce good performance in a variety of settings. This approach can be used for not only uncertainty quantification in density estimation, but also to conduct inferences on any functional of the density as well as prediction. We also discuss issues of identifiability.