

# MULTILEVEL BAYESIAN DEEP NEURAL NETWORKS

NEIL CHADA

In this talk we consider Bayesian inference associated to deep neural networks (DNNs) and in particular, trace-class neural network (TNN) priors which were proposed by Sell et al. Such priors were developed as more robust alternatives to classical architectures in the context of inference problems. For this work we develop multilevel Monte Carlo (MLMC) methods for such models. MLMC is a popular variance reduction technique, with particular applications in Bayesian statistics and uncertainty quantification. We show how a particular advanced MLMC method that was introduced in [6] can be applied to Bayesian inference from DNNs, and used to compute expectations associated to the posterior predictive. We establish that the cost to achieve a given mean square error can be reduced by several orders, in comparison to more conventional techniques. To verify such results we provide numerous numerical experiments on several model problems arising in machine learning, including regression, classification, and reinforcement learning.