

CRUNCH Seminars at Brown, Division of Applied Mathematics

Friday - October 9, 2020

**Uncertainty in Neural Networks: Approximately Bayesian
Ensembling**

Liu Yang

Understanding the uncertainty of a neural network's (NN) predictions is essential for many purposes. The Bayesian framework provides a principled approach to this, however applying it to NNs is challenging due to large numbers of parameters and data. Ensembling NNs provides an easily implementable, scalable method for uncertainty quantification, however, it has been criticised for not being Bayesian. This work proposes one modification to the usual process that we argue does result in approximate Bayesian inference; regularising parameters about values drawn from a distribution which can be set equal to the prior. A theoretical analysis of the procedure in a simplified setting suggests the recovered posterior is centred correctly but tends to have underestimated marginal variance, and overestimated correlation. However, two conditions can lead to exact recovery. We argue that these conditions are partially present in NNs. Empirical evaluations demonstrate it has an advantage over standard ensembling, and is competitive with variational methods.