CRUNCH Seminars at Brown, Division of Applied Mathematics

Friday - February 5, 2021

Deep evidential classification/regression

Apostolos Psaros

Deterministic neural networks (NNs) are increasingly being deployed in safety-critical domains, where calibrated, robust, and efficient measures of uncertainty are crucial. Most techniques that provide uncertainty estimates, including Bayesian NNs, rely on sampling multiple NN parameter sets, making multiple predictions, and subsequently computing the statistics of these predictions (such as mean and variance). In other words, prediction uncertainty is inferred through NN parameter uncertainty. Orthogonally to these approaches, recently published works on evidential classification and regression explicitly model uncertainty by placing priors over the categorical (classification) and the Gaussian (regression) likelihood functions. As a result, by training a deterministic NN that outputs the hyperparameters of these priors (instead of the parameters of the likelihood function as usual) a predictor is constructed that outputs both mean response and uncertainty with a single forward pass.