Bayesian nonparametric learning through randomized loss functions

Chris Holmes (University of Oxford, UK)

We discuss Bayesian nonparametric learning whereby Bayesian nonparametric models are used to train Bayesian parametric models by way of suitably randomized objective (loss) functions. The resulting posterior models exhibit provably better properties than their conventional Bayesian counterparts when the sampling distribution (likelihood function) is misspecified. For additive log-likelihoods inference is achieved through posterior sampling obtained by independent optimizations of randomly re-weighted loss-functions, as opposed to Markov chain Monte Carlo sampling. This avoids issues with MCMC such as burn-in and chain dependence, and is highly scalable on modern computer architectures allowing for samples to be drawn in parallel for the price of a single model optimization. We demonstrate the approach on a number of examples including nonparametric learning for Bayesian logistic regression, variational Bayes (VB), mixture models, and random forests. The work has its foundations in the weighted-likelihood bootstrap of Newton and Raftery (1994).