Bayesian computation and machine learning

Nicolas Chopin (ENSAE, IPP)

PAC-Bayes

Uses as an estimator the expectation of pseudo-posterior:

$$p(x|y) \propto p(x) \exp\{-\gamma R(x, y)\}$$

where $R(x, y)$ is the empirical risk, for parameter x and data y.

How to compute this expectation?

1. Fast variational approximation: but can you we obtain the same non-asymptotic bounds? See Alquier, Ridgway and C. (2016, JMLR).
2. Monte Carlo methods: isn't that slow? not if you do it right, e.g. Sequential Monte Carlo (Ridgway et al, NIPS, 2014).

Other applications of Bayesian computation

1. Probabilistic machine learning.
2. Sequential learning: use Sequential Monte Carlo?
3. Non-convex optimisation