Bayesian Model Averaging

May 6, 2004

Abstract

Statistical inference is based on a specific model. Ignoring all other possible models is equivalent to ignoring the uncertainty about the model itself and can have serious implications on inference. Bayesian Model Averaging (BMA) provides an elegant solution to this problem. The BMA framework is built on a set of models to which one assigns prior probabilities. Standard practice is to assign equal probabilities, but unequal weights can be assigned when expert opinions are available. The BMA strategy is to successively visit models, obtain model specific estimates and use the rules of probability to **average** over all these inferences. The number of times a model is visited is proportional to its posterior probability. BMA usually has the effect of reducing the significance of parameters and improving prediction.