
BIMS: for Bayesian inference of model structure

Nicos Angelopoulos

Department of Surgery and Cancer, Division of Cancer, Imperial College London,
Hammersmith Hospital Campus, London, UK.

MCMC methods in probabilistic logic programming settings are gaining popularity and a number of different approaches have been proposed recently. We discuss theoretical results and experiences with applications of one of the first approaches in the field. The knowledge representation capabilities of the underlying language, which are less well documented in the literature, are discussed, as well as the machine learning applications of the general framework, which have been presented in a number of publications. We focus on how to express Bayesian prior knowledge in this formalism, and show how it can be used to define generative priors over statistical model spaces: Bayesian networks and classification and regression trees. Finally, we discuss a Metropolis-Hastings algorithm that can take advantage of the defined priors and its application to real-world machine learning tasks. Details of the associated publicly available software are also discussed.