Mach Learn (2012) 89:279–297
DOI 10.1007/s10994-012-5305-8

Online Bayesian inference for the parameters of PRISM programs

James Cussens

Received: 5 February 2012 / Revised: 4 April 2012 / Accepted: 8 June 2012 / Published online: 27 July 2012
© The Author(s) 2012

Abstract This paper presents a method for approximating posterior distributions over the parameters of a given PRISM program. A sequential approach is taken where the distribution is updated one datapoint at a time. This makes it applicable to online learning situations where data arrives over time. The method is applicable whenever the prior is a mixture of products of Dirichlet distributions. In this case the true posterior will be a mixture of very many such products. An approximation is effected by merging products of Dirichlet distributions. An analysis of the quality of the approximation is presented. Due to the heavy computational burden of this approach, the method has been implemented in the Mercury logic programming language. Initial results using a hidden Markov model and a probabilistic graph are presented.

Keywords Inductive logic programming · Bayesian statistics · Statistical relational learning · PRISM · Mixture models · Missing data

1 Introduction

In the Bayesian approach to 'parameter estimation' the goal is to return the joint posterior distribution over all parameters, rather than return the single 'best estimate' of the parameters. The motivation for attempting this complex task is that the posterior captures the combined information given by observed data and prior knowledge, and so provides a much fuller picture of the state of our knowledge about the parameters than can a point estimate.

Unfortunately, many posterior distributions are hard even to represent let alone compute efficiently. This is certainly generally the case for posterior distributions over the parameters of PRISM programs. PRISM programs define distributions over finite or countably infinite sample spaces using potentially complex generative processes. Generally the steps

Editors: S. Muggleton and J. Chen.

J. Cussens (✉)
Dept. of Computer Science & York Centre for Complex Systems Analysis, University of York, Deramore Lane, York, YO10 5GE, UK
e-mail: james.cussens@york.ac.uk