A Sequential Monte Carlo Framework for Adaptive Bayesian Model Discrimination Designs using Mutual Information

Christopher C. Drovandi[1], James M. McGree[1], Anthony N. Pettitt[1]

Bayesian Young Statisticians Meeting (BAYSM), Milan June, 5-6, 2013
Paper no. 12

[1]Queensland University of Technology, Brisbane, Australia
c.drovandi@qut.edu.au

Abstract

In this paper we present a unified sequential Monte Carlo (SMC) framework for performing sequential experimental design for discriminating between a set of models. The model discrimination utility that we advocate is fully Bayesian and based upon the mutual information. SMC provides a convenient way to estimate the mutual information. Our experience suggests that the approach works well on either a set of discrete or continuous models and outperforms other model discrimination approaches.

Keywords: Model Discrimination; Mutual Information; Sequential design; Sequential Monte Carlo.

1 Introduction

The problem of model choice within a Bayesian framework has received an abundance of attention in the literature. Therefore, when a set of competing models is proposed a priori, it is important to determine the optimal selection of the controllable aspects (when available) of the experiment for discriminating between the models. A sequential experimental design allows experiments to be performed in batches, so that adaptive decisions can be made for each new batch.

In this paper we adopt a unified sequential Monte Carlo (SMC) framework for performing model discrimination in sequential experiments. We consider as a utility the mutual information between the model indicator and the next observation(s) [1]. SMC allows for convenient estimation of posterior model probabilities [3] as well as the mutual information utility, both of which are generally difficult to calculate. In SMC, new data can be accommodated via a simple re-weight step. Thus, the simulation properties of various utilities can be discovered in a timely manner with SMC compared with approaches that use Markov chain Monte Carlo to re-compute posterior distributions (see [4]).

From our experience we have found the approach to be successful on several diverse applications, including models for both discrete (see [5]) and continuous (see [6]) data. The purpose of this paper is to collate [5] and [6] into a single source describing the SMC mutual information for model discrimination calculation for applications involving a set of discrete or continuous models. Section 2 develops the notation, Section 3 details SMC under model uncertainty and Section 4 describes the mutual information calculation.

2 Notation

We use the following notation. We consider a finite number, K, models, described by the random variable $M \in \{1, \ldots, K\}$. We assume one of the K models is responsible for data generation. Each model contains a parameter, $\theta_{\mathbf{m}}$, with a likelihood function, $f(\mathbf{y_t}|m, \theta_{\mathbf{m}}, \mathbf{d_t})$, where $\mathbf{y_t}$ represents the data collected up to current t based on the selected design points, $\mathbf{d_t}$. We place a prior distribution over $\theta_{\mathbf{m}}$ for each model, denoted by $\pi(\theta_{\mathbf{m}}|m)$. $\pi(m)$ and $\pi(m|\mathbf{y_t}, \mathbf{d_t})$ is the prior and posterior probability of model m, respectively.

3 Sequential Monte Carlo Incorporating Model Uncertainty

SMC consists of a series of re-weighting, re-sampling and mutation steps. For a single model, we use the algorithm of [2]. For sequential designs involving model uncertainty, we run an SMC algorithm in parallel for each model and combine them after introducing each observation to compute posterior model probabilities and the mutual information utility. We denote the particle set at target t for the mth model obtained by SMC as $\{W_{m,t}^{i}, \theta_{\mathbf{m,t}}^{\mathbf{i}}\}_{i=1}^{N}$, where N is the number of particles. It is well known that SMC provides a simple way to estimate the evidence for a particular model based on importance weights, which can be converted to estimates of the posterior model probabilities.

4 Mutual Information for Model Discrimination

For model discrimination, we advocate the use of the mutual information utility between the model indicator and the next observation, first proposed in [1]. This utility provides us with the expected gain in information about the model indicator introduced by the next observation. In general it is difficult to calculate, however SMC allows efficient calculation. One can show that the utility for the design d to apply for the next observation z is given by

$$U(d|\mathbf{y_t}, \mathbf{d_t}) = \sum_{m=1}^{K} \pi(m|\mathbf{y_t}, \mathbf{d_t}) \int_{z\in\mathcal{S}} f(z|m, \mathbf{y_t}, \mathbf{d_t}, d) \log \pi(m|\mathbf{y_t}, \mathbf{d_t}, z, d) dz,$$

where $\mathcal{S}$ is the sample space of the response z. Below, we denote SMC estimates of predictive distributions and posterior model probabilities with a hat. If z is discrete, a summation replaces the integral

$$\hat{U}(d|\mathbf{y_t}, \mathbf{d_t}) = \sum_{m=1}^{K} \hat{\pi}(m|\mathbf{y_t}, \mathbf{d_t}) \sum_{z\in\mathcal{S}} \hat{f}(z|m, \mathbf{y_t}, \mathbf{d_t}, d) \log \hat{\pi}(m|\mathbf{y_t}, \mathbf{d_t}, z, d),$$

[5]. When z is continuous, the integral can be approximated using the SMC particle population for each model

$$\hat{U}(d|\mathbf{y_t}, \mathbf{d_t}) = \sum_{m=1}^{K} \hat{\pi}(m|\mathbf{y_t}, \mathbf{d_t}) \sum_{i=1}^{N} W_{m,t}^{i} \log \hat{\pi}(m|\mathbf{y_t}, \mathbf{d_t}, z_{m,t}^{i}, d),$$

[6] where $z_{m,t}^{i} \sim f(z|m, \theta_{\mathbf{m,t}}^{\mathbf{i}}, d)$ if the observations are independent.

5 Conclusion

Here we have brought together the research of [5] and [6] into a single source for performing adaptive Bayesian model discrimination under discrete or continuous model uncertainty. The methodology relies on sequential Monte Carlo, which has already proven to be useful in sequential designs [4], and furthermore provides a convenient estimate of the mutual information utility we advocate for model discrimination.

References

[1] G. E. P. Box, W. J. Hill. **Discrimination among mechanistic models**. *Technometrics*; 1967; 9(1); pp. 57-71.

[2] N. Chopin. **A sequential particle filter method for static models**. *Biometrika*; 2002; 89(3); 539-551.

[3] P. Del Moral, A. Doucet, A. Jasra. **Sequential Monte Carlo samplers**. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*; 2006; 68(3); 411-436.

[4] C. C. Drovandi, J. M. McGree, A. N. Pettitt. **Sequential Monte Carlo for Bayesian sequentially designed experiments for discrete data**. *Computational Statistics & Data Analysis*; 2013; 57(1); 320-335.

[5] C. C. Drovandi, J. M. McGree, A. N. Pettitt. **A sequential Monte Carlo algorithm to incorporate model uncertainty in Bayesian sequential design**. *To appear in Journal of Computational and Graphical Statistics*; 2013.

[6] J. M. McGree, C. C. Drovandi, A. N. Pettitt. **A sequential Monte Carlo approach to the sequential design for discriminating between rival continuous data models.**. *http://eprints.qut.edu.au/53813/*; 2013.