

Chapter 5

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

Christopher C. Drovandi, James M. McGree, and Anthony N. Pettitt

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.

5.1 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.

C.C. Drovandi (✉) • J.M. McGree • A.N. Pettitt
Queensland University of Technology, GPO Box 2434, Brisbane 4001, Australia
e-mail: c.drovandi@qut.edu.au; james.mcgree@qut.edu.au; a.pettitt@qut.edu.au

In SMC, new data can be accommodated via a simple re-weighting 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 recompute posterior distributions (see [5]).

From our experience we have found the approach to be successful on several diverse applications, including models for both discrete (see [4]) and continuous (see [7]) data. The purpose of this paper is to collate [4, 7] into a single source describing the SMC mutual information for model discrimination calculation for applications involving a set of discrete or continuous models. Section 5.2 develops the notation, Sect. 5.3 details SMC under model uncertainty and Sect. 5.4 describes the mutual information calculation. Section 5.5 describes the examples our approach has been tested on while Sect. 5.6 concludes the paper.

5.2 Notation

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

5.3 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 SMC algorithms 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 m th model obtained by SMC as $\{W_{m,t}^i, \theta_{m,t}^i\}_{i=1}^N$, 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. The reader is referred to [4] for more details on the algorithm.

5.4 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

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

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

$$\hat{U}(d|\mathbf{y}_t, \mathbf{d}_t) = \sum_{m=1}^K \hat{\pi}(m|\mathbf{y}_t, \mathbf{d}_t) \sum_{z \in \mathcal{S}} \hat{f}(z|m, \mathbf{y}_t, \mathbf{d}_t, d) \log \hat{\pi}(m|\mathbf{y}_t, \mathbf{d}_t, z, d), \quad (2)$$

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

$$\hat{U}(d|\mathbf{y}_t, \mathbf{d}_t) = \sum_{m=1}^K \hat{\pi}(m|\mathbf{y}_t, \mathbf{d}_t) \sum_{i=1}^N W_{m,t}^i \log \hat{\pi}(m|\mathbf{y}_t, \mathbf{d}_t, z_{m,t}^i, d), \quad (3)$$

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

5.5 Examples

The SMC algorithm for designing in the presence of model uncertainty together with the use of the mutual information utility function has been tested on a variety of discrete and continuous model examples spanning several application areas. The SMC algorithm facilitated faster assessment of different utility functions for model discrimination purposes. Drovandi et al. [4] considered binary and count data examples. The applications included memory retention models, dose-response relationships in the context of clinical trials and models for neuronal degeneration. In all cases the mutual information utility led to a more rapid identification of the correct model compared to a random design. McGree et al. [7] applied the algorithm to continuous model examples. The methodology was illustrated on competing models for an asthma dose-finding study, a chemical engineering application and

a pharmacokinetics example. The mutual information utility was compared to a random design and the total separation criterion (see, e.g., [6]), which is another model discrimination utility. We found that the mutual information utility led to designs that were more robust for detecting the correct model across applications.

5.6 Conclusion

Here we have brought together the findings of [4, 7] into a single source for performing adaptive Bayesian model discrimination under discrete or continuous model uncertainty. The methodology relies on SMC, which has already proven to be useful in sequential designs [5] and furthermore provides a convenient estimate of the mutual information utility we advocate for model discrimination. The combination of the SMC algorithm and mutual information utility has been successfully tested on a wide range of applications.

References

1. Box GEP, Hill WJ (1967) Discrimination among mechanistic models. *Technometrics* 9:57–71
2. Chopin N (2002) A sequential particle filter method for static models. *Biometrika* 89:539–551
3. Del Moral P, Doucet A, Jasra A (2006) Sequential Monte Carlo samplers. *J Roy Stat Soc Ser B Stat Methodol* 68:411–436
4. Drovandi CC, McGree JM, Pettitt AN (2012) A sequential Monte Carlo algorithm to incorporate model uncertainty in Bayesian sequential design. *J Comput Graph Stat*. doi:10.1080/10618600.2012.730083
5. Drovandi CC, McGree JM, Pettitt AN (2013) Sequential Monte Carlo for Bayesian sequentially designed experiments for discrete data. *Comput Stat Data Anal* 57:320–335
6. Masoumi S, Duever TA, Reilly PM (2013) Sequential Markov chain Monte Carlo (MCMC) model discrimination. *Cand J Chem Eng* 91:862–869
7. McGree JM, Drovandi CC, Pettitt AN (2013) A sequential Monte Carlo approach to the sequential design for discriminating between rival continuous data models. <http://eprints.qut.edu.au/53813/>.