

Department of Mathematics and Statistics


Preprint MPS-2016-07

27 April 2016

An ABC interpretation of the multiple auxiliary variable method

by

Dennis Prangle and **Richard G. Everitt**



An ABC interpretation of the multiple auxiliary variable method

Dennis Prangle* and Richard G. Everitt†

Abstract

We show that the auxiliary variable method (Møller et al., 2006; Murray et al., 2006) for inference of Markov random fields can be viewed as an approximate Bayesian computation method for likelihood estimation.

Keywords: ABC, Markov random field, annealed importance sampling, multiple auxiliary variable method

1 Introduction

Markov random fields (MRFs) have densities of the form

$$f(y|\theta) = \gamma(y|\theta)/Z(\theta), \quad (1)$$

where $\gamma(y|\theta)$ can be evaluated numerically but $Z(\theta)$ cannot in a reasonable time. This makes it challenging to perform inference.

This note considers two approaches which both use simulation from $f(y|\theta)$. The single auxiliary variable (SAV) method (Møller et al., 2006) and the multiple auxiliary variable (MAV) method (Murray et al., 2006) provide unbiased likelihood estimates. Approximate Bayesian computation (Marin et al., 2012) finds parameters which produce simulations similar to the observed data. We will demonstrate that these two methods are in fact closely linked.

An additional challenge for inference of MRFs is that exact sampling from $f(y|\theta)$ is difficult. It is possible to implement Markov chain Monte Carlo (MCMC) algorithms which sample from a close approximation to this distribution. These MCMC algorithms have been used for inference through their use as a replacement for an exact sampler in SAV and MAV (Caimo and Friel, 2011; Everitt, 2012) as well as ABC (Grelaud et al., 2009). We will use this approach and discuss it further below.

The remainder of the paper is as follows. Section 2 reviews ABC and MAV methods, and Section 3 derives our result. Throughout the paper y refers to an observed dataset, and x variables refer to simulated datasets used in inference.

*Newcastle University. Email dennis.prangle@newcastle.ac.uk

†University of Reading

2 Background

2.1 Auxiliary variable methods

The SAV method makes use of an unbiased estimate of $f(y|\theta)$, given by using the following *importance sampling* (IS) estimate of $1/Z(\theta)$

$$\frac{\hat{1}}{Z} = \frac{q_x(x|y, \theta)}{\gamma(x|\theta)},$$

where q_x is some arbitrary (normalised) density and $x \sim f(\cdot|\theta)$. MAV extends this idea by instead using *annealed IS* (AIS) (Neal, 2001) for this estimate

$$\frac{\hat{1}}{Z} = \prod_{i=2}^a \frac{\gamma_{i-1}(x_i|\theta, y)}{\gamma_i(x_i|\theta, y)}, \quad (2)$$

where $f_i(\cdot|\theta, y) \propto \gamma_i(\cdot|\theta, y)$ are bridging densities between $f_a(\cdot|\theta, y) = f(\cdot|\theta)$ and $f_1(\cdot|\theta, y) = \gamma_1(\cdot|\theta, y) = q_x(\cdot|\theta, y)$, $x_a \sim f(\cdot|\theta)$ and for $2 \leq i < a$, $x_i \sim K_i(\cdot|x_{i+1})$ where K_i is a reversible Markov kernel with invariant distribution f_i . In this description we have imposed that γ_1 is normalised in order to obtain an estimate of $1/Z$. However we note that a common choice for $q_x(\cdot|\theta, y)$ is $f(\cdot|\tilde{\theta})$ for some estimate $\tilde{\theta}$, in which case the normalising constant $Z(\tilde{\theta})$ is not available. In this case we obtain an estimate of $Z(\tilde{\theta})/Z(\theta)$ from Equation (2).

The SAV and MAV estimates are usually used as constituent parts of other Monte Carlo algorithms for parameter inference: in MCMC (Møller et al., 2006) or IS (Everitt et al., 2016). The estimates of $f(y|\theta)$ just described may be used here since only an unbiased estimate of the posterior up to proportionality is required (Andrieu and Roberts, 2009).

As noted in the introduction, the requirement of being able to draw x_a exactly from $f(\cdot|\theta)$ is potentially problematic. Caimo and Friel (2011) and Everitt (2012) explore the possibility of replacing this exact sampler with a long run (of b iterations) of an MCMC sampler targeting $f(\cdot|\theta)$, and taking the final point. Such an approach results in biased estimates of $1/Z$, although as $b \rightarrow \infty$ this bias goes to zero. Everitt et al. (2016) observes empirically that a similar argument appears to hold when $b = 0$ but a is large.

2.2 Approximate Bayesian computation

ABC refers to a family of inference algorithms (described in Marin et al., 2012) which perform an approximation to Bayesian inference when numerical evaluation of the likelihood function is intractable. They instead use simulation from the model of interest. The core of these algorithms is producing estimates of the likelihood $f(y|\theta)$ using some version of the following method. Simulate a dataset x from $f(\cdot|\theta)$ and return the *ABC likelihood estimate*:

$$L_{\text{ABC}} = \mathbb{1}(\|y - x\| \leq \epsilon).$$

Here $\mathbb{1}$ represents an indicator function, $\|\cdot\|$ is some distance norm, and the acceptance threshold ϵ is a tuning parameter. The expectation of the random variable L_{ABC} is

$$\int f(x|\theta) \mathbb{1}(\|y - x\| \leq \epsilon) dx.$$

This is often referred to as the *ABC likelihood*. It is proportional to a convolution of the likelihood with a uniform density, evaluated at y . For $\epsilon > 0$ this is generally an inexact approximation to the likelihood. For discrete data it is possible to use $\epsilon = 0$ in which case the ABC likelihood equals the exact likelihood, and so L_{ABC} is unbiased.

For MRFs empirically it is observed that, compared with competitors such as the exchange algorithm (Murray et al., 2006), ABC requires a relatively large number of simulations to yield an efficient algorithm (Friel, 2013).

3 Derivation

3.1 ABC for MRF models

Suppose that the model $f(y|\theta)$ has an intractable likelihood but can be targeted by a MCMC chain $\mathbf{x} = (x_1, x_2, \dots, x_n)$. Let π represent densities relating to this chain. Then $\pi_n(y|\theta) := \pi(x_n = y|\theta)$ is an approximation of $f(y|\theta)$ which can be estimated by ABC. For now suppose that y is discrete and consider the ABC likelihood estimate requiring an exact match: simulate from $\pi(\mathbf{x}|\theta)$ and return $\mathbb{1}(x_n = y)$. We will consider an IS variation on this: simulate from $g(\mathbf{x}|\theta)$ and return $\mathbb{1}(x_n = y)\pi(\mathbf{x}|\theta)/g(\mathbf{x}|\theta)$. Under the mild assumption that $g(\mathbf{x}|\theta)$ has the same support as $\pi(\mathbf{x}|\theta)$ (typically true unless n is small), both estimates have the expectation $\Pr(x_n = y|\theta)$.

This can be generalised to cover continuous data using the identity

$$\pi_n(y|\theta) = \int_{x_n=y} \pi(\mathbf{x}|\theta) dx_{1:n-1},$$

where $x_{i:j}$ represents $(x_i, x_{i+1}, \dots, x_j)$. An importance sampling estimate of this integral is

$$w = \frac{\pi(\mathbf{x}|\theta)}{g(x_{1:n-1}|\theta)} \quad (3)$$

where $\mathbf{x}$ is sampled from $g(x_{1:n-1}|\theta)\delta(x_n = y)$, with δ representing a Dirac delta measure. Then, under mild conditions on the support of g , w is an unbiased estimate of $\pi_n(y|\theta)$.

The ideal choice of $g(x_{1:n-1}|\theta)$ is $\pi(x_{1:n-1}|x_n, \theta)$, as then $w = \pi(x_n = y|\theta)$ exactly. This represents sampling from the Markov chain conditional on its final state being y .

3.2 Equivalence to MAV

We now show that natural choices of $\pi(\mathbf{x}|\theta)$ and $g(x_{1:n-1}|\theta)$ in the ABC method just outlined results in the MAV estimator (2). Our choices are

$$g(x_{1:n-1}|\theta) = \prod_{i=1}^{n-1} K_i(x_i|x_{i+1})$$

$$\pi(\mathbf{x}|\theta) = f_1(x_1|\theta, y) \prod_{i=1}^{n-1} K_i(x_{i+1}|x_i).$$

Here $\pi(\mathbf{x}|\theta)$ defines a MCMC chain with transitions $K_i(x_{i+1}|x_i)$. Suppose K_i is as in Section 2.1 for $i \leq a$, and for $i > a$ it is a reversible Markov kernel with invariant distribution $f(\cdot|\theta)$. Also assume $b := n - a \rightarrow \infty$. Then the MCMC chain ends in a long sequence of steps targeting $f(\cdot|\theta)$ so that $\lim_{n \rightarrow \infty} \pi_n(\cdot|\theta) = f(\cdot|\theta)$. Thus the likelihood being estimated converges on the true likelihood for large n . Note this is the case even for fixed a .

The importance density $g(x_{1:n-1}|\theta)$ specifies a reverse time MCMC chain starting from $x_n = y$ with transitions $K_i(x_i|x_{i+1})$. Simulating $\mathbf{x}$ is straightforward by sampling x_{n-1} , then x_{n-2} and so on. This importance density is an approximation to the ideal choice stated at the end of Section 3.1.

The resulting likelihood estimator is

$$w = f_1(x_1|\theta, y) \prod_{i=1}^{n-1} \frac{K_i(x_{i+1}|x_i)}{K_i(x_i|x_{i+1})}.$$

Using detailed balance gives

$$\frac{K_i(x_{i+1}|x_i)}{K_i(x_i|x_{i+1})} = \frac{f_i(x_{i+1}|\theta, y)}{f_i(x_i|\theta, y)} = \frac{\gamma_i(x_{i+1}|\theta, y)}{\gamma_i(x_i|\theta, y)},$$

so that

$$w = f_1(x_1|\theta, y) \prod_{i=1}^{n-1} \frac{\gamma_i(x_{i+1}|\theta, y)}{\gamma_i(x_i|\theta, y)} = \gamma(y|\theta) \prod_{i=2}^n \frac{\gamma_{i-1}(x_i|\theta, y)}{\gamma_i(x_i|\theta, y)}.$$

This is an unbiased estimator of $\pi_n(y|\theta)$. Hence

$$v = \prod_{i=2}^n \frac{\gamma_{i-1}(x_i|\theta, y)}{\gamma_i(x_i|\theta, y)} = \prod_{i=2}^a \frac{\gamma_{i-1}(x_i|\theta, y)}{\gamma_i(x_i|\theta, y)}.$$

is an unbiased estimator of $\pi_n(y|\theta)/\gamma(y|\theta) \rightarrow 1/Z(\theta)$. In the above we have assumed, as in Section 3.1, that γ_1 is normalised. When this is not the case then we instead get an estimator of $Z(\tilde{\theta})/Z(\theta)$, as for MAV methods. Also note that in either case a valid estimator is produced for any choice of y .

The ABC estimate can be viewed by a two stage procedure. First run a MCMC chain of length b with any starting value, targeting $f(\cdot|\theta)$. Let its final value be x_a . Secondly run a MCMC chain $x_a, x_{a-1}, \dots$ using kernels $K_{a-1}, K_{a-2}, \dots$ and evaluate the estimator v . This is unbiased in the limit $b \rightarrow \infty$, so the first stage could be replaced by perfect sampling methods where these exist.

The resulting procedure is thus equivalent to that for MAV.

4 Conclusion

We have demonstrated that the MAV method can be interpreted as an ABC algorithm. We hope this insight will be useful for the development of novel methods for MRFs.

References

- Andrieu, C. and Roberts, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, pages 697–725.
- Caimo, A. and Friel, N. (2011). Bayesian inference for exponential random graph models. *Social Networks*, 33:41–55.
- Everitt, R. G. (2012). Bayesian parameter estimation for latent markov random fields and social networks. *Journal of Computational and Graphical Statistics*, 21:940–960.
- Everitt, R. G., Johansen, A. M., Rowing, E., and Evdemon-Hogan, M. (2016). Bayesian model comparison with un-normalised likelihoods. *Statistics and Computing: early online version*.
- Friel, N. (2013). Evidence and Bayes factor estimation for Gibbs random fields. *Journal of Computational and Graphical Statistics*, 22:518–532.
- Grelaud, A., Robert, C. P., Marin, J.-M., Rodolphe, F., and Taly, J. F. (2009). ABC likelihood-free methods for model choice in Gibbs random fields. *Bayesian Analysis*, 4(2):317–336.
- Marin, J.-M., Pudlo, P., Robert, C. P., and Ryder, R. J. (2012). Approximate Bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180.
- Møller, J., Pettitt, A. N., Reeves, R. W., and Berthelsen, K. K. (2006). An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. *Biometrika*, 93:451–458.
- Murray, I., Ghahramani, Z., and MacKay, D. J. C. (2006). MCMC for doubly-intractable distributions. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 359–366.
- Neal, R. M. (2001). Annealed importance sampling. *Statistics and Computing*, 11:125–139.