Bayesian model comparison with intractable likelihoods

Richard G. Everitt, Adam M. Johansen, Ellen Rowing and Melina Evdemon-Hogan

1st April 2015

Abstract

Markov random eld models are used widely in computer science, statistical physics and spatial statistics and network analysis. However, Bayesian analysis of these models using standard Monte Carlo methods is not possible due to their intractable likelihood functions. Several methods have been developed that permit exact, or close to exact, simulation from the posterior distribution. However, estimating the evidence and Bayes' factors (BFs) for these models remains challenging in general. This paper describes new random weight importance sampling and sequential Monte Carlo methods for estimating BFs that use simulation to circumvent the evaluation of the intractable likelihood, and compares them to existing methods. In some cases we observe an advantage in the use difiased weight estimates; an initial investigation into the theoretical and empirical properties of this class of methods is presented.

1 Introduction

There has been much recent interest in performing Bayesian inference in models where the posterior is intractable. Speci cally, we have the situation where the posterior distribution (jy) / p()f(yj)

cannot be evaluated due to the presence of the INC. We rst review exact methods for simulating from such a target in sections 1.1.1, 1.1.2 and 1.1.3, before looking at simulation-based methods in sections 1.1.4 and 1.1.5. The methods described here in the context of MCMC form the basis for the methods for evidence estimation we develop in the rest of the paper.

1.1.1 Single and multiple auxiliary variable methods

Møller et al. (2006) avoid the evaluation of the INC by augmenting the target distribution with an extra variable u that lies on the same space as, and use an MH algorithm with target distribution

$$(; u jy) / q_u(uj; y) f(yj) p();$$
 (2)

where q_u is some (normalised) arbitrary distribution. As the MH proposal in (; u)-space they use

$$(; u) f(uj)q(j);$$
 (3)

giving an acceptance probability of

min 1;
$$\frac{q(j)}{q(j)} \frac{p()}{p()} \frac{(yj)}{(yj)} \frac{q_u(uj;y)}{(uj)} \frac{(uj)}{q_u(uj;y)}$$
 : (4)

Note that, by viewing $q_u(u j ; y) = (u j)$ as an unbiased IS estimator of 1=Z(), this algorithm can be seen as an instance of the exact-approximations described in Beaumont (2003) and Andrieu and Roberts (2009), where it is established that if an unbiased estimator of a target density is used appropriately in an MH algorithm, the -marginal of the invariant distribution of this chain is the target distribution of interest. This automatically suggests extensions to the single auxiliary variable (SAV) method described above, where M importance points are used to instead give the estimate

$$\frac{d_1}{Z()} = \frac{1}{M} \frac{X^{V}}{m=1} \frac{q_u(u^{(m)}j;y)}{(u^{(m)}j)}$$
(5)

Andrieu and Vihola (2012) show that the reduced variance of this estimator leads to a reduced asymptotic variance of estimators from the resultant Markov chain. The variance of the IS estimator is strongly dependent on an appropriate choice of IS target_u(j;y), which should ideally have lighter tails than f(j). Møller et al. (2006) suggest that a reasonable choice may bop₂(j;y) = f(j^b), where ^b is the maximum likelihood estimator of . However, in practice q_u(j;y) can be di cult to choose well, particularly when y lies on a high dimensional space. Motivated by this, annealed importance sampling (AIS) (Neal, 2001) can be used as an alterne79091 Tf 5.424 0 Td [d17874.25matheteric]

between f (j) and $q_u(j;y)$. After the initial draw $u_{K+1} = f(j)$, the auxiliary point is taken through a sequence of MCMC moves which successively have target_k(j; ^b,y) for k = K : 1. The resultant IS estimator is given by

$$\frac{d_{1}}{Z()} = \frac{1}{M} \frac{M}{m=1} \sum_{k=0}^{m=1} \frac{W^{k}}{k+1} \frac{(u_{k}^{(m)})}{u_{k}}$$

introduced by this approximation tends to zero as the run length of the internal MCMC increases: the same proof holds for the use of an MCMC chain for the simulation within an ABC-MCMC or SL-MCMC algorithm, as described in sections 1.1.4 and 1.1.5. Although the approach of Girolami et al. (2013) is exact, they comment that it is signi cantly more computationally expensive than this approximate approach. For this reason, we do not pursue Russian Roulette approaches further in this paper.

When a rejection sampler is available for simulating from f(j), Rao et al. (2013) introduce an alternative exact algorithm that has some favourable properties compared to the exchange algorithm. Since a rejection sampler is not available in many cases, we do not pursue this approach further.

1.1.4 Approximate Bayesian computation

ABC (Tavaré et al., 1997) refers to methods that aim to approximate an intractable likelihood f (yj) through the integral

$$f^{e}(S(y)j) = (S(u)jS(y)) f(uj) \frac{@\$}{@\mu}(u) du;$$
(8)

where S() gives a vector of summary statistics, $j \circledast \mathbf{s} = \mathfrak{e}\mu(u)j$ denotes the Jacobian determinant arising from the change of variable, and (jS(y)) is a density centred at S(y) with bandwidth . As ! 0, this distribution becomes more concentrated around S(y), so that in the case where S() gives su cient statistics for estimating , as ! 0 the approximate posterior becomes closer to the true posterior. This approximation is used within standard Monte Carlo methods for simulating from the posterior. For example, it may be used within an MCMC algorithm (known as ABC-MCMC (Marjoram et al., 2003)), where using an exact-approximation argument it can be seen that it is su cient in the calculation of the acceptance probability to use the Monte Carlo approximation

$$f^{(0)}(S(y)j) = \frac{1}{M} \frac{X^{(1)}}{m=1} \qquad S \quad u^{(m)} \quad jS(y)$$
 (9)

for the likelihood at at each iteration, where $u^{(m)}g_{m=1}^{M}$ f (j). Whilst the exact-approximation argument means that there is no additional bias due to this Monte Carlo approximation, the approximation introduced through using a tolerance > 0 or insu cient summary statistics may be large. For this reason it might be considered a last resort to use ABC on likelihoods with an INC, but previous success on these models (e.g Grelaetdal. (2009) and Everitt (2012)) lead us to consider them further in this paper.

1.1.5 Synthetic likelihood

ABC is essentially using, based on simulations from , a nonparameteric estimator off $_{S}$ (Sj), the distribution of the summary statistics of the data given . In some situations, a parametric model might be more appropriate. For example, it might be that the statistic consists of the sum of

independent random variables, in which case a Central Limit Theorem (CLT) might imply that it would be appropriate to assume that $_{S}(Sj)$ is multivariate Gaussian.

The SL approach (Wood, 2010) proceeds by making exactly this Gaussian assumption and uses this approximate likelihood within an MCMC algorithm (SL-MCMC), The parameters (the mean and variance) of this approximating distribution for a given are estimated based on the summary statistics of simulations f $u^{(m)}g_{m=1}^{M}$ f (j). Concretely, the estimate of the likelihood is

their estimation as triply intractable when f has an INC. To our knowledge the only published approach to estimating the evidence for such models is in Friel (2013), with this paper also giving one of the only approaches to estimating BFs in this setting. For estimating BFs, ABC provides a viable alternative (Grelaud et al., 2009), as long as the approximations in this approach described in section 1.1.4 are not too large.

The methods in Friel (2013) are based on Chib's approximation,

$$\mathbf{p}(\mathbf{y}) = \frac{f(\mathbf{y}\mathbf{j}^{\mathbf{e}})\mathbf{p}(^{\mathbf{e}})}{\mathbf{b}(^{\mathbf{e}}\mathbf{j}\mathbf{y})};$$
(13)

where ^e can be an arbitrary value of

with an estimate of the evidence given by

$$\mathbf{p}(y) = \frac{1}{P} \sum_{p=1}^{N} \mathbf{w}^{(p)}:$$
 (15)

To estimate a BF we simply take the ratio of estimates of the evidence for the two models under consideration. However, the presence of the INC in the weight expression in equation 14 means that importance samplers cannot be directly implemented for these models. To circumvent this problem we will investigate the use of the techniques described in section 1.1 in importance sampling. We begin by looking at exact-approximation based methods in section 2.1. We then examine the use to approximate likelihoods based on simulation, including ABC and SL in section 2.2, before looking at the performance of all of these methods on a toy example in section 2.3. Finally, in sections 2.4 and 2.5 we examine applications to exponential random graph models (ERGMs) and Ising models, the latter of which leads us to consider the use of inexact-approximations in IS.

2.1 Auxiliary variable IS

To avoid the evaluation of the INC in equation 14, we propose the use of the auxiliary variable method used in the MCMC context in section 1.1.1. Speci cally, consider IS using the SAV target

p(;ujy) / q_u(uj;y)f(yj)p();

noting that it has the same evidence asp(jy), with proposal

This results in weights

$$ve^{(p)} = q_u(uj p)(p)$$

rather than simple IS, for estimating $1=Z(^{(p)})$ as in equation 7 (giving an algorithm that we refer to as multiple auxiliary variable IS (MAVIS), in common with the terminology in Murray et al. (2006)). Using $q_u(j; y) = f(j^b)$, as described in section 1.1.1, and k as in equation 6, we obtain

$$\frac{d_{1}}{Z()} = \frac{1}{Z(b)} \frac{1}{M} \frac{M}{m_{m=1 \ k=0}} \frac{M}{k=0} \frac{u_{k+1}(u_{k}^{(m)}j; ;y)}{u_{k}(u_{k}^{(m)}j; ;y)}$$
(16)

In this case the (A)IS methods are being used as unbiased estimators of the rat $\overline{z}(b)=Z()$.

2.2 Simulation based methods

Didelot et al. (2011) investigate the use of the ABC approximation when using IS for estimating marginal likelihoods. In this case the weight equation becomes

where $x_r^{(p)} \circ_{r=1}^{0} f(j^{(p)})$, and using the notation from section 1.1.4. However, using these weights within equation 15 gives an estimate forp(S(y)) rather than, as desired, an estimate of the evidence p(y). The only way to obtain an estimate of the evidence from ABC is to use the full data, rather than taking summary statistics.

Fortunately, there are cases in which ABC may be used to estimate BFs. Didelot al. (2011) establishes that, for the BF for two exponential family models: if $S_1(y)$ is su cient for the parameters in model 1 and $S_2(y)$ is su cient for the parameters in model 2, then using $S(y) = (S_1(y); S_2(y))$ gives

$$\frac{p(yjM_1)}{p(yjM_2)} = \frac{p(S(y)jM_1)}{p(S(y)jM_2)}:$$

Outside the exponential family, making an appropriate choice of summary statistics can be more involved (Robert et al., 2011; Prangleet al., 2013; Marin et al., 2013).

Just as in the parameter estimation case, the use of a tolerance> 0 results in estimating an approximation to the true BF. An alternative approximation, not previously used in model comparison, is to use SL (as described in section 1.1.5). In this case the weight equation becomes

2.3 Toy example

In this section we have introduced three alternative methods for estimating BFs: MAVIS, ABC and SL. To further understand their properties we now investigate the performance of each method on a toy example.

Consider i.i.d. observationsy = $f y_i g_{i=1}^{n=100}$ of a discrete random variable that takes values in N. For such a dataset, we will not the BF for the models

$$f_{1}(fy_{i}g_{i=1}^{n}j) = \frac{Y_{i}}{\sum_{i=1}^{y_{i}} \frac{y_{i}exp()}{y_{i}!}}$$
$$= \frac{1}{\exp(n)} \frac{Y_{i}}{\sum_{i=1}^{y_{i}} \frac{y_{i}}{y_{i}!}}$$

2. yj Geometric(), Unif(0;1)

$$\begin{array}{rcl} f_{2}\left(f\,y_{i}\,g_{i=1}^{n}\,\,j\right) &=& \displaystyle \stackrel{Y^{n}}{\underset{i=1}{\overset{j=1}{p}} p(1 \quad p)^{y_{i}} \\ &=& \displaystyle \frac{1}{p^{-n}} \, \stackrel{Y^{n}}{\underset{i=1}{\overset{j=1}{p}} (1 \quad p)^{y_{i}} \, ; \end{array}$$



(a) A box plot of the $\,$ log of the estimated BF divided by the true BF.



(b) The log of the BF estimated by ABC-IS against the log of the true BF.



(c) The log of the BF estimated by SL-IS against the log of the true BF.



(d) The log of the BF estimated by MAVIS against the log of the true BF.

Figure 1: Bayes' factors for the Poisson and geometric models.

by the bias. For this reason it might be more relevant in this example to consider the deviations from the shallow slope, which are likely due to the Monte Carlo variance in the estimator (which becomes more pronounced asis reduced). We see that the choice of essentially governs a bias-variance trade-o, and that the di culty in using the approach more generally is that it is not easy to evaluate whether a choice of



Figure 2: The Gamaneg data.

	ABC ($= 0:1$)	ABC (= 0:05)	SL	MAVIS
$\frac{\hat{p}(yjM_1)}{\hat{p}(yjM_2)}$	4	20	40	41

Table 1: Model comparison results for Gamaneg data. Note that the ABC (= 0:05) estimate was based upon just 5 sample points of non-zero weight. MAVIS also provides estimates of the individual evidence (log [$\mathbf{p}(yjM_1)$] = 69:6, log [$\mathbf{p}(yjM_2)$] = 73:3).

sums of random variables. However, we note that this is not usually the case for ERGMs, particularly when modelling large networks, and that SL is a much more appropriate method for inference in the ERGMs with local dependence (Schweinberger and Handcock, 2015). We might expect a more sophisticated ABC approach to exhibit improved performance, possibly outperforming SL. However, the appeal of SL is in its simplicity, and we nd it to be a useful method for obtaining good results with minimal tuning.

2.5 Application to Ising models

The implementation of MAVIS in the previous section is not an exact-approximate method for two reasons:

- 1. An internal MCMC chain was used in place of an exact sampler;
- The 1=Z(^b) term in equation 16 was estimated before running this algorithm (by using a standard SMC method to estimateZ(^b), and taking the reciprocal) with this xed estimate being used throughout.

However, in practice, we tend to nd that such inexact-approximations do not introduce large errors into Bayes' factor estimates, particularly when compared to standard implementations of ABC (as seen in the previous section). We investigate some of the theoretical aspects of such approximations in section 4. In the current section we investigate this type of approach further empirically, using data simulated from Ising models. In particular we reanalyse the data from Friel (2013), which consists of 20 realisations from a rst-order10 10 Ising model and 20 realisations from a second-order10 10



Figure 3: Box plots of the results of population exchange, SAVIS, bridged SAVIS and MAVIS on the Ising data.

We observe that AIS o ers an improvement over IS, but also that the bridged IS method performs better than IS despite the additional bias in estimating Z() = Z()

here are a natural alternative to the MCMC methods described in section 1.1. and inherently use a population of Monte Carlo points (shown to be bene cial on these models by Caimo and Friel (2011)). In section 3.1 we describe these algorithms, before examining an application to estimating the precision matrix of a Gaussian distribution in high dimensions in section 3.2.

3.1 SMC samplers in the presence of an INC

This section introduces two alternative SMC samplers for use on doubly intractable target distributions. The rst, marginal SMC, directly follows from the IS methods in the previous section. The second, SMC-MCMC, requires a slightly di erent approach, but is more computationally e cient. Finally we brie y discuss simulation-based SMC samplers in section 3.1.4.

3.1.1 Marginal SMC

The rst method we describe results from the use of an approximation to the optimal backward kernel (Peters, 2005; Klaaset al., 2005). In this case the weight update is

$$\mathbf{w}_{t}^{(p)} = \mathbf{P} \frac{p({t}^{(p)}_{t}) \mathbf{f}_{t}(yj{t}^{(p)}_{t})}{\Pr_{r=1}^{P} w_{t}^{(r)} \mathbf{K}_{t}({t}^{(p)}_{t}j{t}^{(r)}_{t})}$$
(20)

$$= \frac{p({(t)}^{(p)}) t(yj{(t)}^{(p)})}{Z_{t}({(t)}^{(p)}) P_{r=1}^{P} w_{t-1}^{(r)} K_{t}({(t)}^{(p)}j{(t)})}$$
(21)

for an arbitrary forward kernel K_t. This method is quite widely used, but is a little non-standard in terms of the framework of Del Moral et al. (2006), since it uses a Monte Carlo estimate of an importance weight de ned on the marginal -space at targett, compared to the usual weight on the entire past history of each particle. This results in a computational complexity ofO(P²) (although as noted by Klaaset al. (2005) that this may often be reduced toO(P log(P)) at the cost of negligible bias), compared toO(P) for a standard SMC method, but we include it here since we notice that equation 21 contains the term1=Z() just as does the corresponding expression (equation 14) for standard IS. This leads us to consider the same approach for avoiding the calculation $\vec{a}f()$ as in section 2.1. Namely, we employ the SAV target and proposal within the SMC algorithm.

We still have an intractable normalising constant in the denominator. Now let us try the SAVM posterior, where in target t we use the distribution q_u for the auxiliary variable u_t , and the SAVM proposal, where $u_t^{(p)} = f_t(j_t^{(p)})$. In this case, the weight update is

$$\mathbf{w}_{t}^{(p)} = \mathbf{P} \frac{\mathbf{q}_{u}(\mathbf{u}_{t}^{(p)}\mathbf{j}_{t}^{(p)};\mathbf{y})\mathbf{p}(\frac{(p)}{t})\mathbf{f}_{t}(\mathbf{y}\mathbf{j}_{t}^{(p)})}{\mathbf{P}_{r=1} \mathbf{K}_{t}(\frac{(p)}{t}\mathbf{j}_{t}^{(r)}\mathbf{j})\mathbf{f}_{t}(\mathbf{u}_{t}^{(p)}\mathbf{j}_{t}^{(p)})\mathbf{w}_{t}^{(r)}\mathbf{1}} \\ = \frac{\mathbf{q}_{u}(\mathbf{u}_{t}^{(p)}\mathbf{j}_{t}^{(p)};\mathbf{y})\mathbf{p}(\frac{(p)}{t})\mathbf{f}_{t}(\mathbf{y}\mathbf{j}_{t}^{(p)})}{\mathbf{t}(\mathbf{u}_{t}^{(p)}\mathbf{j}_{t}^{(p)})\mathbf{P} \mathbf{P}_{r=1} \mathbf{K}_{t}(\frac{(p)}{t}\mathbf{j}_{t}^{(r)}\mathbf{j}\mathbf{w}_{t}^{(r)}\mathbf{1})\mathbf{w}_{t}^{(r)}\mathbf{1}}$$

We see that no normalising constant appears in this weight update, but the algorithm is $O(P^2)$.

3.1.2 SMC with an MCMC kernel

Suppose we were able to use a reversible MCMC kernkl_t with invariant distribution $_t(jy) / p()f_t(yj)$, and choose theL_{t 1} kernel to be the time reversal ofK_t with respect to its invariant

 $Z_{t.\ t}$

multivariate Gaussian distribution with zero mean and p() is a Wishart distribution W(;V) with parameters d and V 2 R^d d. Suppose we observe i.i.d. observations $y = f y_i g_{i=1}^n$, where $y_i 2 R^d$. The true evidence can be calculated analytically, and is given by

$$p(y) = \frac{1}{nd=2} \frac{d(\frac{+n}{2})}{d(\frac{1}{2})} \frac{V^{-1} + \frac{P_{i=1}}{j + 1} y_{i} y_{i}^{T_{i}} + \frac{1 + \frac{-1}{2}}{j + 1}}{j + 1 + \frac{1}{2}};$$
(29)

where $_d$ denotes thed-dimensional gamma function. For ease of implementation, we parametrise the precision using a Cholesky decomposition $^1 = LL^0$ with L a lower triangular matrix whose (i; j)'th element is denoted a_{ij} .

As in section 2.3, we writef (yj) as (yj)=Z() as follows

$$f \quad f \quad y_{i} \quad g_{i=1}^{n} \quad j \quad 1 \quad = \quad \stackrel{Y^{n}}{\overset{i=1}{\overset{i=1}{\overset{i=1}{\overset{i=1}{\overset{j}{\atop{}}}}}} (2) \quad {}^{d=2} \quad j \quad {}^{1=2} \quad \exp \quad \frac{1}{2} y_{i}^{0} \quad {}^{1} y_{i} \qquad ! \\ = \quad (2)^{d} \quad j \quad j \quad {}^{n=2} \quad \exp \quad \frac{1}{2} \sum_{i=1}^{N} y_{i}^{0} \quad {}^{1} y_{i} \quad ; \qquad (30)$$

where in some of the experiments that follow $Z() = \begin{pmatrix} h \\ (2)^d j \end{pmatrix}^{d} j^{n=2}$ is treated as if it is an INC. In the Wishart prior, we take = 10 + d and $V = I_d$.

3.2.2 Data, algorithm description and results

Taking d = 10, n = 30 points were simulated usingy_i MVN $(0_d; 0:1 \ I_d)$. The parameter space is thus 55-dimensional, motivating the use of an SMC sampler in place of IS or the population

approximation of this sequence can then be understood as a simple mean eld approximation and its convergence has been well studied, see for example Del Moral (2004).

In order to do this, we make a number of identi cations in order to allow the consideration of the approximation in an abstract manner. We allow \mathbf{G}_t to denote the incremental weight function at time t, and \mathbf{G}_t to denote the exact weight function which it approximates (any auxiliary random variables needed in order to obtain this approximation are simply added to the state space and their sampling distribution to the transition kernel). The transition kernel M_t combines the proposal distribution of the SMC algorithm together with the sampling distribution of any needed auxiliary variables. We allow x to denote the full collection of variables sampled during an iteration of the sampler, which is assumed to exist on the same space during each iteration of the sampler.

We employ the following assumptions (we assume an in nite sequence of algorithm steps and associated target distributions, proposals and importance weights; naturally, in practice only a nite number would be employed but this formalism allows for a straightforward statement of the result):

A1 (Bounded Relative Approximation Error) There exists < 1 such that:

$$\sup_{t \ge N} \sup_{x} \frac{jG_t(x) \quad \mathfrak{G}_t(x)j}{\mathfrak{G}_t(x)}$$

:

A2 (Strong Mixing; slightly stronger than a global Doeblin condition) There exists (M) > 0 such that:

$$\sup_{t \ge N} \inf_{x;y} \frac{dM_t(x;)}{dM_t(y;)} \qquad (M):$$

A3 (Control of Potential) There exists (G) > 0 such that:

$$\sup_{t \ge N} \inf_{x;y} \frac{G_t(x)}{G_t(y)} \qquad (G):$$

The rst of these assumptions controls the error introduced by employing an inexact weighting

This result is not intended to do any more than demonstrate that, qualitatively, such forgetting can prevent the accumulation of error even in systems with biased importance weighting potentials. In practice, one would wish to make use of more sophisticated ergodicity results such as those of Whiteley (2013), within this framework to obtain results which are somewhat more broadly applicable: assumptions A2 and A3 are very strong, and are used only because they allow stability to be established simply. Although this result is, in isolation, too weak to justify the use of the approximation schemes introduced here in practice it seems su cient, together with the empirical results presented below, to suggest that further investigation of such approximations might be warranted.

4.2.2 Empirical results

We use the precision example introduced in section 3.2.1 to investigate the e ect of using biased weights in SMC samplers. Speci cally we taked = 1 and use a simulated datasety where n = 500 points were simulated usingy_i N (0;0:1). In this case there is only a single parameter to estimate, $_{1}$

listed in section 1. Our results suggest that improved mixing can help combat the accumulation of bias, which may implies that there may be situations where it is useful to perform many iterations of a kernel at a particular target, rather than the more standard approach of using many intermediate targets at each of which a single iteration of a kernel is used. Other variations are also possible, such as the calculation of fast cheap biased weights at each target in order only to adaptively decide when to resample, with more accurate weight estimates (to ensure accurate resampling and accurate estimates based on the particles) only calculated when the method chooses to resample.

5 Conclusions

This paper describes several IS and SMC approaches for estimating the evidence in models with INCs that outperform previously described approaches. These methods may also prove to be useful alternatives to MCMC for parameter estimation. Several of the ideas in the paper are also applicable more generally, in particular the use of synthetic likelihood in the IS context and the notion of using biased weight estimates in IS and SMC. We advise caution in the use of biased weights in SMC due to the potential for bias to accumulate, but also note that this accumulated bias may be small compared to errors resulting from commonly accepted approximate techniques such as ABC.

References

- Del Moral P (2004) Feynman-Kac formulae: genealogical and interacting particle systems with applications. Probability and Its Applications, Springer, New York
- Del Moral P, Doucet A, Jasra A (2006) Sequential Monte Carlo samplers. Journal of the Royal Statistical Society: Series B 68(3):411 436
- Del Moral P, Doucet A, Jasra A (2007) Sequential Monte Carlo for Bayesian computation. Bayesian Statistics 8:115 148
- Didelot X, Everitt RG, Johansen AM, Lawson DJ (2011) Likelihood-free estimation of model evidence. Bayesian Analysis 6(1):49 76
- Drovandi CC, Pettitt AN, Lee A (2014) Bayesian indirect inference using a parametric auxiliary model. Statistical Science
- Everitt RG (2012) Bayesian Parameter Estimation for Latent Markov Random Fields and Social Networks. Journal of Computational and Graphical Statistics 21(4):940 960
- Fearnhead P, Papaspiliopoulos O, Roberts GO, Stuart AM (2010) Random-weight particle Itering of continuous time processes. Journal of the Royal Statistical Society Series B 72(4):497 512
- Friel N (2013) Evidence and Bayes factor estimation for Gibbs random elds. Journal of Computational and Graphical Statistics 22(3):518 532
- Friel N, Rue H (2007) Recursive computing and simulation-free inference for general factorizable models. Biometrika 94(3):661 672
- Girolami MA, Lyne AM, Strathmann H, Simpson D, Atchade Y (2013) Playing Russian Roulette with Intractable Likelihoods. arXiv (1306.4032)
- Grelaud A, Robert CP, Marin JM (2009) ABC likelihood-free methods for model choice in Gibbs random elds. Bayesian Analysis 4(2):317 336
- Klaas M, de Freitas N, Doucet A (2005) Toward practical N² Monte Carlo: The marginal particle Iter. In: Proceedings of the 20th International Conference on Uncertainty in Arti cial Intelligence
- Kong A, Liu JS, Wong WH (1994) Sequential imputations and Bayesian missing data problems. JASA 89(425):278 288
- Marin JM, Pillai NS, Robert CP, Rousseau J (2014) Relevant statistics for Bayesian model choice. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 76(5):833 859
- Marjoram P, Molitor J, Plagnol V, Tavare S (2003) Markov chain Monte Carlo without likelihoods. Proceedings of the National Academy of Sciences of the United States of America 100(26):15,324 15,328

- Whiteley N (2013) Stability properties of some particle Iters. Annals of Applied Probability 23(6):2500 2537
- Wilkinson RD (2013) Approximate Bayesian computation (ABC) gives exact results under the assumption of model error. Statistical Applications in Genetics and Molecular Biology 12(2):129 141
- Wood SN (2010) Statistical inference for noisy nonlinear ecological dynamic systems. Nature 466(August):1102 1104
- Zhou Y, Johansen AM, Aston JAD (2013) Towards automatic model comparison: An adaptive sequential Monte Carlo approach. arXiv (1303.3123)

A Using SAV and exchange MCMC within SMC

A.1 Weight update when using SAV-MCMC

Let us consider the SAVM posterior, with K being the MCMC move used in SAVM. In this case the weight update is

$$\begin{split} \mathbf{w}_{k}^{(p)} &= \frac{p(\frac{t}{t}^{(p)})f_{t}(yj\frac{t}{t}^{(p)})q_{u}(u_{t}^{(p)}j\frac{t}{t}^{(p)};y)}{p(\frac{t}{t}^{(p)})f_{t-1}(yj\frac{t}{t}^{(p)})q_{u}(u_{t}^{(p)}j\frac{t}{t}^{(p)};y)} \frac{L_{t-1}((\frac{t}{t}^{(p)};u_{t}^{(p)});(\frac{t}{t}^{(p)};u_{t}^{(p)}))}{K_{t}((\frac{t}{t}^{(p)})f_{t-1}(yj\frac{t}{t}^{(p)})q_{u}(u_{t}^{(p)}j\frac{t}{t}^{(p)};y)} \frac{p(\frac{t}{t}^{(p)})f_{t}(yj\frac{t}{t}^{(p)})(q_{u}(u_{t}^{(p)}j\frac{t}{t}^{(p)};y)}{p(\frac{t}{t}^{(p)})f_{t-1}(yj\frac{t}{t}^{(p)})q_{u}(u_{t}^{(p)}j\frac{t}{t}^{(p)};y)} \frac{p(\frac{t}{t}^{(p)})f_{t}(yj\frac{t}{t}^{(p)})q_{u}(u_{t}^{(p)}j\frac{t}{t}^{(p)};y)}{p(\frac{t}{t}^{(p)})f_{t}(yj\frac{t}{t}^{(p)})q_{u}(u_{t}^{(p)}j\frac{t}{t}^{(p)};y)} \\ &= \frac{t(yj\frac{t}{t}^{(p)})}{t-1(yj\frac{t}{t}^{(p)})} \frac{Z_{t-1}(\frac{t}{t}^{(p)})}{Z_{t}(\frac{t}{t}^{(p)})}; \end{split}$$

which is the same update as if we could use MCMC directly.

A.2 Weight update when using the exchange algorithm

Nicholls et al. (2012) show the exchange algorithm, when set up to $target_t(jy) / p()f_t(yj)$ in the manner described in section 1.1.2, simulates a transition kernel that is in detailed balance with t(jy). This follows from showing that it satis es a very detailed balance condition, which takes account of the auxiliary variable u. The result is that the derivation of the weight update follows exactly that of equations 22-24.

B An extended space construction for the random weight SMC method in section 3.1.2

The following extended space construction justi es the use of the approximate weights in equation 25 via an explicit sequential importance (re)sampling argument along the lines of Del Moraet al.

(2006), albeit with a slightly di erent sequence of target distributions.

Consider an actual sequence of target distributions $t_t g_{t-0}$. Assume we seek to approximate a normalising constant during every iteration by introducing additional variables $u_t = (u_t^1; \ldots; u_t^M)$ during iteration t > 0.

De ne the sequence of target distributions:

extended space.

C Proof of SMC Sampler Error Bound

A little notation is required. We allow

If A1. holds, then 8 2 P(E) and any t 2 N:

$$jj_{\mathfrak{G}_t}() = G_t()jj_{\mathsf{TV}} 2$$
:

Proof. Let $t := \mathbf{G}_t$ G_t and consider a generic 2 C_b(E):

$$(_{\mathfrak{G}_{t}}() _{G_{t}}())(') = \frac{(\mathfrak{G}_{t}')}{(\mathfrak{G}_{t})} \frac{(G_{t}')}{(G_{t})}$$

$$= \frac{(G_{t}) (\mathfrak{G}_{t}') (\mathfrak{G}_{t}) (G_{t}')}{(\mathfrak{G}_{t}) (G_{t})}$$

$$= \frac{(G_{t}) ((G_{t} + _{t})') ((G_{t} + _{t})) (G_{t}')}{(\mathfrak{G}_{t}) (G_{t})}$$

$$= \frac{(G_{t}) (_{t}') (_{t}) (G_{t})}{(\mathfrak{G}_{t}) (G_{t})}$$

Considering the absolute value of this discrepancy, making using of the triangle inequality:

$$(_{\mathfrak{G}_{t}}() _{G_{t}}())(') = \frac{(G_{t}) (_{t}') (_{t}) (G_{t}')}{(\mathfrak{G}_{t}) (G_{t})}$$
$$\frac{(_{t}')}{(\mathfrak{G}_{t})} + \frac{(_{t}) }{(\mathfrak{G}_{t})} \frac{(G_{t}')}{(G_{t})}$$

Noting that G_t is strictly positive, we can bound $j(G_t') = (G_t)$ with $(G_t j') = (G_t)$ and thus with k' k₁ and apply a similar strategy to the rst term:

$$({}_{\mathfrak{G}_{t}}() {}_{G_{t}}())(') = \frac{(j {}_{t}j)k' k_{1}}{(\mathfrak{G}_{t})} + \frac{({}_{t})}{(\mathfrak{G}_{t})} \frac{(G_{t}j' j)}{(G_{t})}$$
$$k' k_{1} + k' k_{1} = 2 k' k_{1} :$$

(noting that $(j t_j) = (\mathbf{G}_t) < \mathbf{G}_t$

We thus establish (noting that $e_0 = 0$):

$$t \quad e_{t} = \sum_{s=1}^{X^{t}} s_{;t}(s(e_{s-1})) \quad s_{;t}(e_{s}(e_{s-1})):$$
(32)

Turning our attention to an individual term in this expansion, noting that:

$$_{s}()(') = _{G_{s-1}}()M_{s}(')$$
 $e_{s}()(') = _{\mathfrak{G}_{s-1}}()M_{s}(')$

we have, by application of a standard Dobrushin contraction argument and Lemma 1

$$({}_{s}(e_{s-1}) {}^{e}{}_{s}(e_{s-1}))(') = {}_{G_{s-1}}(e_{s-1})M_{s}(') {}_{\mathfrak{G}_{s-1}}(e_{s-1})M_{s}(')$$

$${}_{s}(e_{s-1}) {}^{e}{}_{s}(e_{s-1}) {}_{TV} (1 (M)) {}_{G_{s-1}}(e_{s-1}) {}_{\mathfrak{G}_{s-1}}(e_{s-1}) {}_{TV}$$

$${}_{2} (1 (M)) (33)$$

which controls the error introduced instantaneously during each step.

We now turn our attention to controlling the accumulation of error. We make use of (Del Moral, 2004, Proposition 4.3.6) which, under assumptions A2 and A3, allows us to deduce that for any probability measures ; :

$$k_{s;s+k}() = k_{s;s+k}(k_{TV} = k_{s;s+k})k = k_{TV}$$

where

$$(_{s;s+k}) = \frac{2}{(M)(G)} (1 - {}^{2}(M))^{k}$$

Returning to decomposition (32), applying the triangle inequality and this result, before nally inserting (33) we arrive at:

$$\begin{array}{ccc} X^t & & \\ k_t & e_t k_{TV} & & \\ & s=1 & \\ \end{array} \\ \begin{array}{c} X^t & \\ s;t(s(e_{s-1})) & s;t(e_s(e_{s-1})) & \\ s;t(e_s(e_{s-1})) & s;t(e_s(e_{s-1})) & \\ \end{array} \\ \begin{array}{c} X^t & \\ y \in t \\ s \in t$$

nave, by application of a standard Dobrushin contraction argu/F47 10.9091 Tf m 46.1-6.35 -3.272 Td .9091 1 Tf 4.7320.909n5

result:

$$\frac{4 (1 (M))}{(M) (G)} \frac{\chi}{s=0} (1 - 2(M))^{s} = \frac{4 (1 (M))}{(M) (G)} \frac{1}{1 (1 - 2(M))}$$
$$= \frac{4 (1 (M))}{3(M) (G)}:$$