

Importance tempering

Robert Gramacy · Richard Samworth · Ruth King

Received: 27 September 2007 / Accepted: 3 November 2008 / Published online: 4 December 2008
© Springer Science+Business Media, LLC 2008

Abstract Simulated tempering (ST) is an established Markov chain Monte Carlo (MCMC) method for sampling from a multimodal density $\pi(\theta)$. Typically, ST involves introducing an auxiliary variable k taking values in a finite subset of $[0, 1]$ and indexing a set of tempered distributions, say $\pi_k(\theta) \propto \pi(\theta)^k$. In this case, small values of k encourage better mixing, but samples from π are only obtained when the joint chain for (θ, k) reaches $k = 1$. However, the entire chain can be used to estimate expectations under π of functions of interest, provided that importance sampling (IS) weights are calculated. Unfortunately this method, which we call importance tempering (IT), can disappoint. This is partly because the most immediately obvious implementation is naïve and can lead to high variance estimators. We derive a new optimal method for combining multiple IS estimators and prove that the resulting estimator has a highly desirable property related to the notion of effective sample size. We briefly report on the success of the optimal combination in two modelling scenarios requiring reversible-jump MCMC, where the naïve approach fails.

Keywords Simulated tempering · Importance sampling · Markov chain Monte Carlo (MCMC) · Metropolis–coupled MCMC

R. Gramacy (✉) · R. Samworth
Statistical Laboratory, University of Cambridge, Cambridge, UK
e-mail: bobby@statslab.cam.ac.uk

R. Samworth
e-mail: rjs57@statslab.cam.ac.uk

R. King
CREEM, University of St. Andrews, St. Andrews, Scotland
e-mail: ruth@mcs.st-and.ac.uk

1 Introduction

Markov chain Monte Carlo (MCMC) algorithms, in particular Metropolis–Hastings (MH) and Gibbs Sampling (GS), are by now the most widely used methods for simulation-based inference in Bayesian statistics. The beauty of MCMC is its simplicity. Very little user input or expertise is required in order to establish a Markov chain whose stationary distribution is proportional to $\pi(\theta)$, for $\theta \in \Theta \subseteq \mathbb{R}^d$. As long as the chain is irreducible, the theory of Markov chains guarantees that sample averages computed from this realisation will converge in an appropriate sense to their expectations under π . However, difficulties can arise when π has isolated modes, between which the Markov chain moves only rarely. In such cases convergence is slow, meaning that often infeasibly large sample sizes are needed to obtain accurate estimates.

New MCMC algorithms have been proposed to improve mixing. Two related algorithms are Metropolis–coupled MCMC (MC³) (Geyer 1991; Hukushima and Nemoto 1996) and simulated tempering (ST) (Marinari and Parisi 1992; Geyer and Thompson 1995). Both are closely related to the optimisation technique of simulated annealing (SA) (Kirkpatrick et al. 1983). SA works with a set of *tempered* distributions $\pi_k(\theta)$ indexed by an inverse-temperature parameter $k \in [0, \infty)$. One popular form of tempering is called “powering up”, where $\pi_k(\theta) \propto \pi(\theta)^k$. Small values of k have the effect of flattening/widening the peaks and raising troughs in π_k relative to π .

In MC³ and ST we define a *temperature ladder* $1 = k_1 > k_2 > \dots > k_m \geq 0$, and call the k_i its *rungs*. Both MC³ and ST involve simulating from the set of m tempered densities $\pi_{k_1}, \dots, \pi_{k_m}$. MC³ runs m parallel MCMC chains, one at each temperature, and regularly proposes swaps of states at adjacent rungs k_i and k_{i+1} . Usually, samples are

only saved from the “cold distribution” π_{k_1} . In contrast, ST works with a “pseudo-prior” $p(k_i)$ and uses a single chain to sample from the joint distribution, which is proportional to $\pi_k(\theta)p(k)$. Again, it is only at iterations t for which $k^{(t)} = 1$ that the corresponding realisation of $\theta^{(t)}$ is retained. ST has an advantage over MC³ in that only one copy of the process $\{\theta^{(t)} : t = 1, \dots, T\}$ is needed—rather than m —so the chain uses less storage and also has better mixing (Geyer 1991). The disadvantage is that it needs a good choice of pseudo-prior. For further comparison and review, see Jasra et al. (2007a) and Iba (2001).

Both MC³ and ST suffer from inefficiency because they discard all samples from π_k for $k \neq 1$. The discarded samples could be used to estimate expectations under π if they were given appropriate importance sampling (IS) weights. For an inclusive review of IS and related methods see Liu (2001, Chap. 2). Moreover, it may be the case that an IS estimator constructed with samples from a tempered distribution has smaller variance than one based on a sample of the same size from π . As a simple motivating example, let $\pi(\theta) = N(\theta|\mu, \sigma^2)$, and consider estimating $\mu = \mathbb{E}_\pi(\theta)$ by IS from a tempered distribution $\pi_k(\theta) \propto \pi(\theta)^k$. A straightforward calculation shows that the value of k which minimises the variance of the IS estimator is

$$k^* = \begin{cases} 1/2 & \text{if } \mu = 0 \\ \frac{3}{2} + \left(\frac{\sigma}{\mu}\right)^2 - \frac{1}{2}\{1 + 8\left(\frac{\sigma}{\mu}\right)^2 + 4\left(\frac{\sigma}{\mu}\right)^4\}^{1/2} & \text{otherwise.} \end{cases}$$

Note that $k^* \in (1/2, 1)$ for all μ and σ^2 . Moreover, one can compute (numerically) $k^- = k^-(\sigma/\mu) < k^*$ such that for all $k \in (k^-, 1)$, the variance of the IS estimator $\hat{\mu}_k$ based on samples from π_k is smaller than that of one based on a sample of the same size from π . However, $\text{Var}(\hat{\mu}_k) \rightarrow \infty$ as $k \rightarrow 0$ for all μ and σ^2 . Table 1 gives k^* and k^- for various values of σ/μ .

Therefore, there is a trade-off in the choice of tempered IS proposals. On the one hand, low inverse-temperatures k in ST can guard against missing modes of π with large support by encouraging better mixing *between* modes, but can yield very inefficient (IS) estimators overall. On the other hand, “lukewarm” temperatures k , especially $k \in (1/2, 1)$, can yield more efficient estimators *within* modes than those obtained from samples at $k = 1$.

Jennison (1993) was the first to suggest using a single tempered distribution as a proposal in IS, and Neal

(1996, 2001, 2005) has since written several papers combining IS and tempering. Indeed, in the discussion of the 1996 paper on *tempered transitions*, Neal writes “simulated tempering allows data associated with p_i other than p_0 [the cold distribution] to be used to calculate expectations with respect to $\dots p_0$ (using an importance sampling estimator)”.¹ It is this natural extension that we call *importance tempering* (IT), with IMC³ defined similarly. Given the work of the above-mentioned authors, and the fact that calculating importance weights is relatively trivial, it may be surprising that successful IT and IMC³ applications have yet to be published. Liu (2001) comes close in proposing to augment ST with dynamic weighting (Wong and Liang 1997) and in applying the Wang–Landau algorithm (Atchadé and Liu 2007) to ST.

This paper addresses why the straightforward methodology described above has tended not to work well in practice, primarily due to a lack of a principled way of combining the importance weights collected at each temperature to obtain an overall estimator. If we are interested in estimating $\mathbb{E}_\pi\{h(\theta)\}$, one way to do this is with

$$\hat{h} = W^{-1} \sum_{t=1}^T w(\theta^{(t)}, k^{(t)})h(\theta^{(t)}),$$

$$\text{where } W = \sum_{t=1}^T w(\theta^{(t)}, k^{(t)}), \tag{1}$$

and $w(\theta, k) = \pi(\theta)/\pi(\theta)^k = \pi(\theta)^{1-k}$. Observe that this estimator is of the form $\hat{h} = \sum_{i=1}^m \lambda_i \hat{h}_i$, where $0 \leq \lambda_i \leq \sum_{i=1}^m \lambda_i = 1$, with $\lambda_i = W^{-1} \sum_{t=1}^T w(\theta^{(t)}, k^{(t)}) \mathbb{I}_{\{k^{(t)}=k_i\}}$, and where each \hat{h}_i is an IS estimator of $\mathbb{E}_\pi\{h(\theta)\}$ constructed using only the observations at the inverse-temperature k_i . We show how to improve this estimator by choosing $\lambda_1, \dots, \lambda_m$ to maximise the *effective sample size* (see next paragraph), which approximately corresponds to minimising the variance of \hat{h} (Liu 2001, Sect. 2.5.3). For the applications that we have in mind, it is important that our estimator can be constructed without knowledge of the normalising constants of $\pi_{k_1}, \dots, \pi_{k_m}$. It is for this reason that methods motivated by the *balance heuristic* (Veach and Guibas 1995; Owen and Zhou 2000; Madras and Picconi 1999) cannot be applied.

The notion of *effective sample size* plays an important role in the study of IS estimators. Suppose we are interested in estimating $\mathbb{E}_\pi\{h(\theta)\}$ using a vector of observations $\theta = (\theta^{(1)}, \dots, \theta^{(T)})$ from a density π' . Define the vector of importance weights $\mathbf{w} \equiv \mathbf{w}(\theta) = (w(\theta^{(1)}), \dots, w(\theta^{(T)}))$, where $w(\theta) = \pi(\theta)/\pi'(\theta)$. Following Liu (2001, Sect. 2.5.3)

Table 1 Values of k^* and k^- for various values of σ/μ

σ/μ	1/16	1/4	1	4	16
k^*	1.00	0.95	0.70	0.52	0.50
k^-	0.99	0.89	0.42	0.18	0.16

¹A similar note is made in the 2001 paper with regard to *annealed importance sampling*.

we define the *effective sample size* by

$$ESS(\mathbf{w}(\theta)) \equiv ESS(\mathbf{w}) = \frac{T}{1 + cv^2(\mathbf{w})}, \tag{2}$$

where $cv^2(\mathbf{w})$ is the *coefficient of variation* of the weights, given by

$$cv^2(\mathbf{w}) = \frac{\sum_{t=1}^T (w(\theta^{(t)}) - \bar{w})^2}{(T - 1)\bar{w}^2},$$

where $\bar{w} = T^{-1} \sum_{t=1}^T w(\theta^{(t)})$.

This should not be confused with the concept of *effective sample size due to autocorrelation* (Kass et al. 1998) (due to serially correlated samples from a Markov chain). This latter notion is discussed briefly in Sect. 4.

Observe that the swap operations in MC³ require that the state space Θ be common for all m tempered distributions. This is not a requirement for ST, as the state stays fixed when changes in temperature are proposed. Thus applying MC³ is less straightforward in (Bayesian) model selection/averaging problems which typically involve trans-dimensional Markov chains as in reversible-jump MCMC (RJMCMC) (Green 1995), though it is possible (Jasra et al. 2007b). Since RJMCMC algorithms are particularly prone to slow mixing, and hence are an excellent source of applications of our idea (as illustrated in Sect. 3), the rest of the paper will focus on IT. Most of our results apply equally to IMC³ by ignoring the pseudo-prior.

The outline of the paper is as follows. In Sect. 2 we derive the optimal convex combination of multiple IS estimators, and show how this estimator has a particularly attractive property with regard to its effective sample size. In Sect. 3 we briefly report on the effectiveness of optimal IT, and the poor performance of the naïve approach, on several real and synthetic examples. Section 4 concludes with a discussion.

2 Importance tempering

The *simulated tempering* (ST) (Geyer and Thompson 1995) algorithm is an application of MH on the product space of parameters and inverse-temperatures. That is, samples are obtained from the joint chain $\pi(\theta, k) \propto \pi(\theta)^k p(k)$. This is only possible if $\pi(\theta)^k$ is integrable, but Hölder’s inequality may be used to show that this is indeed the case provided that $\mathbb{E}_\pi(\|\theta\|^{\frac{1-k}{k}+\delta}) < \infty$ for some $\delta > 0$, where $\|\cdot\|$ denotes the Euclidean norm. The success of ST depends crucially on the ability of the Markov chain frequently to: (a) visit high temperatures (low k) where the probability of escaping local modes is high; (b) visit $k = 1$ to obtain samples from π . The algorithm can be tuned by: (i) adjusting the number and

location of the rungs of the temperature ladder; or (ii) adjusting the pseudo-prior $p(k)$. Geyer and Thompson (1995) give some automated ways of adjusting the spacing of the rungs of the ladder. Iba (2001) reviews similar techniques from the physics literature. A recent alternative—and very promising—approach involves the Wang–Landau algorithm (Atchadé and Liu 2007). However, many authors prefer to rely on defaults, e.g.,

$$k_i = \begin{cases} (1 + \Delta_k)^{1-i} & \text{geometric spacing} \\ \{1 + \Delta_k(i - 1)\}^{-1} & \text{harmonic spacing,} \end{cases} \tag{3}$$

for $i = 1, \dots, m$. The rate parameter $\Delta_k > 0$ can be problem specific. Motivation for such default spacings is outlined by Liu (2001, Chap. 10, pp. 213, 233). Geometric spacing, or uniform spacing of $\log(k_i)$, is also advocated by Neal (1996, 2001).

Once a suitable ladder has been chosen, the goal is typically to choose the pseudo-prior so that the posterior over temperatures is uniform. The best way to accomplish this is to set $p(k_i) = 1/Z_i$, where $Z_i = \int_\Theta \pi(\theta)^{k_i} d\theta$ is the normalising constant in $\pi_{k_i} = \pi^{k_i}/Z_i$, which is generally unknown. So while normalising constants are not a prerequisite for ST, it can certainly be useful to know them. We follow the suggestions of Geyer and Thompson (1995) in setting the pseudo-prior by a method that roughly approximates the Z_i in two-stages: first by stochastic approximation (Kushner and Lin 1997), and then by observation counts accumulated through pilot runs. To some extent, a non-uniform posterior on the temperatures is less troublesome in the context of IT than ST. So long as the chain still visits the heated temperatures often enough to get good mixing in Θ , and if the ESS of the IS estimators at some temperature(s) is not too low, useful samples can be obtained without ever visiting the cold distribution.

2.1 A new optimal way to combine IS estimators

ST provides us with $\{(\theta^{(t)}, k^{(t)}) : t = 1, \dots, T\}$, where $\theta^{(t)}$ is an sample from $\pi_{k^{(t)}}$. Write $T_i = \{t : k^{(t)} = k_i\}$ for the index set of observations at the i^{th} temperature, and let $T_i = |T_i|$. Let the vector of observations at the i^{th} temperature collect in $\theta_i = (\theta_{i1}, \dots, \theta_{iT_i})$, so that $\{\theta_{ij}\}_{j=1}^{T_i} \sim \pi_{k_i}$. Similarly, the vector of IS weights at the i^{th} temperature is $\mathbf{w}_i = \mathbf{w}_i(\theta_i) = (w_i(\theta_{i1}), \dots, w_i(\theta_{iT_i}))$, where $w_i(\theta) = \pi(\theta)/\pi_{k_i}(\theta)$.

Each vector θ_i can be used to construct an IS estimator of $\mathbb{E}_\pi\{h(\theta)\}$ by setting

$$\hat{h}_i = \frac{\sum_{j=1}^{T_i} w_i(\theta_{ij})h(\theta_{ij})}{\sum_{j=1}^{T_i} w_i(\theta_{ij})} \equiv \frac{\sum_{j=1}^{T_i} w_{ij}h(\theta_{ij})}{W_i}.$$

It is natural to consider an overall estimator of $\mathbb{E}_\pi\{h(\theta)\}$ defined by a convex combination:

$$\hat{h}_\lambda = \sum_{i=1}^m \lambda_i \hat{h}_i, \quad \text{where } 0 \leq \lambda_i \leq \sum_{i=1}^m \lambda_i = 1. \tag{4}$$

Unfortunately, if $\lambda_1, \dots, \lambda_m$ are not chosen carefully, $\text{Var}(\hat{h}_\lambda)$, can be nearly as large as the largest $\text{Var}(\hat{h}_i)$ (Owen and Zhou 2000). Notice that ST is recovered as a special case when $\lambda_1 = 1$ and $\lambda_2 = \dots = \lambda_m = 0$. It may be tempting to choose $\lambda_i = W_i/W$, where $W = \sum_{i=1}^m W_i$, recovering the estimator in (1). This can lead to a very poor estimator, even compared to ST, which is demonstrated empirically in Sect. 3.

Observe that we can write

$$\hat{h}_\lambda = \sum_{i=1}^m \sum_{j=1}^{T_i} w_{ij}^\lambda h(\theta_{ij}), \quad \text{where } w_{ij}^\lambda = \frac{\lambda_i w_{ij}}{W_i}. \tag{5}$$

Let $\mathbf{w}^\lambda = (w_{11}^\lambda, \dots, w_{1T_1}^\lambda, w_{21}^\lambda, \dots, w_{2T_2}^\lambda, \dots, w_{m1}^\lambda, \dots, w_{mT_m}^\lambda)$. Attempting to choose $\lambda_1, \dots, \lambda_m$ to minimise $\text{Var}(\hat{h}_\lambda)$ directly can be difficult. In the balance heuristic, Veach and Guibas (1995) explore combinations of IS estimators of the form (5), where $w_i(\theta) = \pi(\theta)/g_i(\theta)$ for a family of proposal densities g_i , with

$$\lambda_{ij} = \frac{c_i g_i(\theta_{ij})}{\sum_{r=1}^m c_r g_r(\theta_{ij})}, \tag{6}$$

and where $0 \leq c_i \leq \sum_{i=1}^m c_i = 1$ is the proportion of samples taken from g_i . It turns out that this is equivalent to IS with the mixture proposal $\tilde{\pi}(\theta) = \sum_{r=1}^m c_r g_r(\theta)$:

$$\hat{h}_{\text{bal}} \equiv \frac{1}{T} \sum_{t=1}^T w(\theta_t) h(\theta_t),$$

$$\text{where } w(\theta) = \frac{\pi(\theta)}{\sum_{r=1}^m c_r g_r(\theta)}. \tag{7}$$

The balance heuristic has since been generalised by Owen and Zhou (2000); it was reinvented by Madras and Picconi (1999, Sect. 4) in the context of applied probability.

Note that due to the denominator in the definition of $w(\theta)$ in (7), the g_i must be normalised densities. This precludes us from using the balance heuristic with $g_i \propto \pi_{k_i}$. When MCMC is necessary to sample from π , the normalisation constant of π , and therefore π_{k_i} , is generally unknown. The method also requires evaluations of $\pi_{k_i}(\theta^{(t)})$, $i = 1, \dots, m$, at all T rounds, an $O(mT)$ operation that trivialises any computational advantage ST has over MC³. Instead, we consider maximising the ESS of \hat{h}_λ in (4).

Proposition 2.1 Among estimators of the form (4), $\text{ESS}(\mathbf{w}^\lambda)$ is maximised by $\lambda = \lambda^*$, where, for $i = 1, \dots, m$,

$$\lambda_i^* = \frac{\ell_i}{\sum_{i=1}^m \ell_i}, \quad \text{and} \quad \ell_i = \frac{W_i^2}{\sum_{j=1}^{T_i} w_{ij}^2}.$$

Proof Since $\sum_{i=1}^m \sum_{j=1}^{T_i} w_{ij}^\lambda = 1$, the problem of maximising the effective sample size is the same as

$$\min_{\lambda_1, \dots, \lambda_m} \sum_{i=1}^m \sum_{j=1}^{T_i} \left(\lambda_i \frac{w_{ij}}{W_i} - \frac{1}{T} \right)^2,$$

subject to $0 \leq \lambda_i \leq \sum_{i=1}^m \lambda_i = 1$.

The result then follows by a straightforward Lagrange multiplier argument. \square

In the following discussion and in Remark 2.2 below, we assume that for $i = 1, \dots, m$, $T_i \geq 2$. The efficiency of each IS estimator \hat{h}_i can be measured through $\text{ESS}(\mathbf{w}_i)$. Intuitively, we hope that with a good choice of λ , the ESS of \hat{h}_λ , given by

$$\text{ESS}(\mathbf{w}^\lambda) = \frac{T(T-1)}{T^2 \sum_{i=1}^m \lambda_i^2 / \ell_i - 1},$$

would be close to the sum over i of the effective sample sizes of \hat{h}_i , namely

$$\text{ESS}(\mathbf{w}_i) = \frac{T_i(T_i-1)\ell_i}{T_i^2 - \ell_i}. \tag{8}$$

The remark below shows that this is indeed the case for \hat{h}_{λ^*} .

Remark 2.2 We have

$$\text{ESS}(\mathbf{w}^{\lambda^*}) \geq \sum_{i=1}^m \text{ESS}(\mathbf{w}_i) - \frac{1}{4} - \frac{1}{T}.$$

Proof Since $\text{ESS}(\mathbf{w}_i) \leq T_i$, it follows from (8) that $\ell_i \leq T_i$. Thus

$$\begin{aligned} \text{ESS}(\mathbf{w}^{\lambda^*}) &= \frac{(1-T^{-1}) \sum_{i=1}^m \ell_i}{1 - \sum_{i=1}^m \frac{\ell_i}{T^2}} \\ &\geq \left(1 - \frac{1}{T}\right) \left(1 + \frac{1}{T^2} \sum_{i=1}^m \ell_i\right) \sum_{i=1}^m \ell_i \\ &= \sum_{i=1}^m \ell_i - \frac{\sum_{i=1}^m \ell_i}{T} \left(1 - \frac{\sum_{i=1}^m \ell_i}{T}\right) - \frac{(\sum_{i=1}^m \ell_i)^2}{T^3} \\ &\geq \sum_{i=1}^m \ell_i - \frac{1}{4} - \frac{1}{T}, \end{aligned}$$

since $x(1 - x)$ attains its maximum of $1/4$ at $x = 1/2$ and $\sum \ell_i \leq \sum T_i = T$. \square

In practice we have found that this bound is slightly conservative and that often it is the case that $ESS(\mathbf{w}^{\lambda^*}) \geq \sum_{i=1}^m ESS(\mathbf{w}_i)$. Thus our optimally-combined IS estimator has a highly desirable and intuitive property in terms of its effective sample size.

3 Empirical results

Here we briefly report on the success of optimal IT, relative to the naïve approach and ST, on one simple example and two involving RJMCMC.

3.1 A simple mixture of normals

Consider the following toy density π , a mixture of two normals:

$$\pi(\theta) = 0.6N(\theta|\mu_1 = -8, \sigma_1^2 = 0.5^2) + 0.4N(\theta|\mu_2 = 8, \sigma_2^2 = 0.9^2). \tag{9}$$

Table 2 summarises Kolmogorov–Smirnov distances obtained under three IT estimators: ST ($\lambda_1 = 1$), naïve IT ($\lambda_i = W_i/W$) and the optimally-combined IT estimator (\hat{h}_{λ^*}). Observe that the optimally-combined IT estimator has both the largest ESS and the smallest variance of the three estimators, and that $ESS(\mathbf{w}^{\lambda^*}) > \sum_i ESS(\mathbf{w}_i)$. Naïve IT improves upon ST in this example, but has higher variance than \hat{h}_{λ^*} .

3.2 Bayesian treed Gaussian process models

Bayesian treed models extend classification and regression tree (CART) models (Breiman et al. 1984), by putting a prior on the tree structure. We focus on the implementation of Gramacy and Lee (2008) who fit Gaussian Process (GP) models at the leaves of the tree, specify the tree prior through a process that limits its depth, and then define the

Table 2 Summary of K–S distances to the true mixture of normals (9) for ST ($\lambda_1 = 1$), naïve IT ($\lambda_i = W_i/W$), the optimally-combined IT estimator (\hat{h}_{λ^*}). We used 100 repeated samples of size 10^5 , with tempered RWM proposals

Method	K–S distance		
	$ESS(\mathbf{w}^{\lambda^*})$	Mean	Var
ST	2535	0.0938	8.5×10^{-4}
Naïve IT	17779	0.0849	1.4×10^{-4}
\hat{h}_{λ^*}	22913	0.0836	5.2×10^{-5}
$\sum_i ESS(\mathbf{w}_i)$	22910		

tree operations *grow*, *prune*, *change*, and *swap*, to allow inference to proceed by RJMCMC. The RJMCMC chain usually identifies the correct *maximum a posteriori* (MAP) tree, but consistently and significantly over estimates the posterior probability of deep trees.

To guard against the transdimensional chain getting stuck in local modes of the posterior, Gramacy and Lee (2008) resorted regularly restarting the chain from the null tree. ST provides an alternative by increasing the rate of accepted tree operations in higher temperatures. In particular, we find that ST can increase the rate of accepted *prune* operations by an order of magnitude, thus enabling the chain to escape the local modes of deep trees. To demonstrate IT we fit a treed GP model with ST using a geometric ladder with $m = 40$ and $k_m = 0.1$ to two datasets first explored by Gramacy and Lee (2008): the 1-d motorcycle accident data and 2-d exponential data. We refer to that paper for details about the data and models.

For the motorcycle accident data the ST chain was run for $T = 1.5 \times 10^5$ iterations, where a total of $T_1 = 3732$ ($\approx T/m = 3750$) samples were obtained from the cold distribution. That $ESS(\mathbf{w}^{\lambda^*}) = 9338 \approx 2.5T_1$ shows the considerable improvement of IT over ST. Moreover, we have $ESS(\mathbf{w}^{\lambda^*}) > \sum_i ESS(\mathbf{w}_i) = 9334$. The naïve combination $\lambda_i = \frac{W_i}{W}$ in (1) yields $ESS(\mathbf{w}^{\lambda}) = 285 < \frac{1}{10}T_1$, undermining the very motivation of IT. For the exponential data the ST chain was run for a total of $T = 5 \times 10^5$ iterations. A total of $T_1 = 12436$ ($\approx T/m = 12500$) samples were obtained from the cold distribution. We found that $ESS(\mathbf{w}^{\lambda^*}) = 21778 \approx 1.75T_1$, illustrating how IT improves on ST. Moreover, we have $ESS(\mathbf{w}^{\lambda^*}) > \sum_i ESS(\mathbf{w}_i) = 21776$. The naïve combination $\lambda_i = \frac{W_i}{W}$ in (1) yields $ESS(\mathbf{w}^{\lambda}) = 654 \approx \frac{1}{18}T_1$ —worse than ST.

3.3 Mark-recapture-recovery data

We now consider a Bayesian model selection problem with data relating to the mark-recapture and recovery of shags on the Isle of May (King and Brooks 2002). The three demographic parameters of interest are: survival rates, recapture rates and recovery rates. The models considered for each of the demographic parameters allowed a possible age- and/or time-dependence, where the time dependence was conditional on the age structure of the parameters. Typically, movement between the different possible models—by adding/removing time dependence for a given age group, or updating the age structure of the parameters—is slow, with small acceptance probabilities. For further details of the data, model structure, and RJMCMC algorithm see King and Brooks (2002).

Using the same ST setup as above, we ran $T = 10^7$ iterations and discarded the first 10% as burn-in. As with

the treed examples, higher temperatures yielded higher acceptance rates and an order of magnitude better exploration of model space compared to (untempered) RJMCMC. A total of $T_1 = 248158$ ($\approx T/m = 225000$) realisations were obtained from the cold distribution. By comparison, for optimal IT we have $\text{ESS}(\mathbf{w}^{\lambda^*}) = 612026 \approx 2.5T_1$ and $\text{ESS}(\mathbf{w}^{\lambda^*}) > \sum_i \text{ESS}(\mathbf{w}_i) = 612020$. The corresponding naïve IT approach (using $\lambda_i = \frac{W_i}{W}$) performed exceptionally poorly, with $\text{ESS}(\mathbf{w}^{\lambda})$ of only 5.43, due to a few large weights obtained at hot temperatures.

4 Discussion

This paper has addressed the inefficiencies and wastefulness of simulated tempering (ST), and related algorithms that are designed to improve mixing in the Markov chain using tempered distributions. We argued that importance sampling (IS) from tempered distributions can produce estimators that are more efficient than ones based on independent sampling, provided that the temperature is chosen carefully. This motivated augmenting the ST algorithm by calculating importance weights to salvage discarded samples—a technique which we have called *importance tempering* (IT). This idea has been suggested before, but to our knowledge little exploration has been carried out for real, complex, applications. We have derived optimal combination weights for the resulting collection of IS estimators, which can be calculated even when the normalisation constants of the tempered distributions are unknown. The weights are essentially proportional to the effective sample size (ESS) of the individual estimators, and we found that the resulting combined ESS in this case would be approximately equal to their sum.

We note that the overall success of the optimal IT estimator depends crucially on a successful implementation of ST, i.e., having a good temperature ladder and pseudo-prior. However, it is also important to recognise that the optimal combination, as a resource-efficient post-processing step, is equally applicable in other contexts, i.e., within MC³, or even outside of the domain of tempered MCMC to combine any collection IS estimators. Sequential Monte Carlo samplers (Del Moral et al. 2006) may facilitate a natural extension. We have illustrated IT on several examples which benefit from the improved mixing ST provides. For example, the optimal IT methodology can increase the resulting ESS compared to retaining samples only from the cold distribution by roughly a factor of two.

Since IT involves sampling from a Markov chain, ideally one would take into account the serial correlation in the objective criteria for combining the individual estimators. The *effective sample size due to autocorrelation* is defined (Kass

et al. 1998) by

$$\text{ESS}_{\rho}(\boldsymbol{\theta}) = \frac{T}{1 + 2 \sum_{\ell=1}^{T-1} \hat{\rho}(\ell, \boldsymbol{\theta})}, \quad (10)$$

where $\hat{\rho}(\ell, \boldsymbol{\theta})$ is the sample autocorrelation in $\boldsymbol{\theta}$ at lag ℓ ; thus for scalar θ we have that $\hat{\rho}(\ell, \boldsymbol{\theta}) = \hat{\gamma}(\ell, \boldsymbol{\theta})/\hat{\gamma}(0, \boldsymbol{\theta})$, where $\hat{\gamma}(\ell, \boldsymbol{\theta}) = (T - \ell)^{-1} \sum_{t=1}^{T-\ell} (\theta^{(t)} - \bar{\theta})(\theta^{(t+\ell)} - \bar{\theta})$, and $\bar{\theta} = T^{-1} \sum_{t=1}^T \theta^{(t)}$. The results from the previous section suggest that, when the temperature ladder is fixed, a sensible heuristic might be to consider combining the individual estimators with weights λ_i^* proportional to product of $T_i^{-1} \text{ESS}_{\rho}(\boldsymbol{\theta}_i)$ and $\text{ESS}(\mathbf{w}_i)$, say. However, when considering modifications to the number (m) and spacing of inverse temperatures $\mathbf{k} = \{k_1, \dots, k_m\}$, there is clearly a conflict of interest between the two measures of effective sample size. Adding more inverse temperatures near one may increase $\text{ESS}(\mathbf{w}^{\lambda^*})$, but may also increase autocorrelation in the marginal chain for k . Therefore it may be sensible to factor $\text{ESS}_{\rho}(\mathbf{k})$ into the objective as well. Searching for temperature ladders that maximise a hybrid of ESS and ESS_{ρ} would represent a natural extension of this work.

References

- Atchadé, Y., Liu, J.: The Wang–Landau algorithm in general state spaces: applications and convergence analysis. Technical Report, University of Harvard (2007)
- Breiman, L., Friedman, J.H., Olshen, R., Stone, C.: Classification and Regression Trees. Wadsworth, Belmont (1984)
- Del Moral, P., Doucet, A., Jasra, A.: Sequential Monte Carlo samplers. J. R. Stat. Soc. Ser. B **68**, 411–436 (2006)
- Geyer, C.: Markov chain Monte Carlo maximum likelihood. In: Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface, pp. 156–163 (1991)
- Geyer, C., Thompson, E.: Annealing Markov chain Monte Carlo with applications to ancestral inference. J. Am. Stat. Assoc. **90**, 909–920 (1995)
- Gramacy, R.B., Lee, H.K.H.: Bayesian treed Gaussian process models with an application to computer modeling. J. Am. Stat. Assoc. **103**(143), 1119–1130 (2008)
- Green, P.: Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. Biometrika **82**, 711–732 (1995)
- Hukushima, K., Nemoto, K.: Exchange Monte Carlo method and application to spin glass simulations. J. Phys. Soc. Jpn. **65**(4), 1604–1608 (1996)
- Iba, Y.: Extended ensemble Monte Carlo. Int. J. Mod. Phys. **12**(5), 623–656 (2001)
- Jasra, A., Stephens, D., Holmes, C.: On population-based simulation for static inference. Stat. Comput. **17**(3), 263–279 (2007a)
- Jasra, A., Stephens, D., Holmes, C.: Population-based reversible jump Markov chain Monte Carlo. Biometrika **94**(4), 787–807 (2007b)
- Jennison, C.: Discussion on the meeting on the Gibbs sampler and other Markov chain Monte Carlo methods. J. R. Stat. Soc. Ser. B **55**, 54–56 (1993)
- Kass, R.E., Carlin, B.P., Gelman, A., Neal, R.M.: Markov chain Monte Carlo in practice: a roundtable discussion. Am. Stat. **52**(2), 93–100 (1998)

- King, R., Brooks, S.: Model selection for integrated recovery/recapture data. *Biometrics* **58**, 841–851 (2002)
- Kirkpatrick, S., Gelatt, C., Vecchi, M.: Optimization by simulated annealing. *Science* **220**, 671–680 (1983)
- Kushner, H., Lin, G.: *Stochastic Approximation Algorithms and Applications*. Springer, New York (1997)
- Liu, J.S.: *Monte Carlo Strategies in Scientific Computing*. Springer, New York (2001)
- Madras, N., Picconi, M.: Importance sampling for families of distributions. *Ann. Appl. Probab.* **9**, 1202–1225 (1999)
- Marinari, E., Parisi, G.: Simulated tempering: a new Monte Carlo scheme. *Europhys. Lett.* **19**, 451–458 (1992)
- Neal, R.M.: Sampling from multimodal distributions using tempered transition. *Stat. Comput.* **6**, 353–366 (1996)
- Neal, R.M.: Annealed importance sampling. *Stat. Comput.* **11**, 125–129 (2001)
- Neal, R.M.: Estimating ratios of normalizing constants using Linked Importance Sampling. Technical Report 0511, Department of Statistics, University of Toronto (2005), 37 pages
- Owen, A., Zhou, Y.: Safe and effective importance sampling. *J. Am. Stat. Assoc.* **95**(449), 135–143 (2000)
- Veach, E., Guibas, L.J.: Optimally combining sampling techniques for Monte Carlo rendering. In: *SIGGRAPH '95 Conference Proceedings*, pp. 419–428. Reading, Addison-Wesley (1995)
- Wong, W., Liang, F.: (1997). Dynamic weighting in Monte Carlo and optimization. In: *Proceedings of the National Academy of Sciences of USA*, vol. 94(26), pp. 14220–14224