

Variational method for estimating the rate of convergence of Markov-chain Monte Carlo algorithms

Fergal P. Casey*

Complex and Adaptive Systems Laboratory, University College Dublin, Dublin 4, Ireland

Joshua J. Waterfall

Department of Molecular Biology and Genetics, Cornell University, Ithaca, New York 14853, USA

Ryan N. Gutenkunst

Department of Biological Statistics and Computational Biology, Cornell University, Ithaca, New York 14853, USA

Christopher R. Myers

Computational Biology Service Unit, Life Sciences Core Laboratories Center, Cornell University, Ithaca, New York 14853, USA

James P. Sethna

Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, New York 14853, USA

(Received 17 July 2008; published 20 October 2008)

We demonstrate the use of a variational method to determine a quantitative lower bound on the rate of convergence of Markov chain Monte Carlo (MCMC) algorithms as a function of the target density and proposal density. The bound relies on approximating the second largest eigenvalue in the spectrum of the MCMC operator using a variational principle and the approach is applicable to problems with continuous state spaces. We apply the method to one dimensional examples with Gaussian and quartic target densities, and we contrast the performance of the random walk Metropolis-Hastings algorithm with a “smart” variant that incorporates gradient information into the trial moves, a generalization of the Metropolis adjusted Langevin algorithm. We find that the variational method agrees quite closely with numerical simulations. We also see that the smart MCMC algorithm often fails to converge geometrically in the tails of the target density except in the simplest case we examine, and even then care must be taken to choose the appropriate scaling of the deterministic and random parts of the proposed moves. Again, this calls into question the utility of smart MCMC in more complex problems. Finally, we apply the same method to approximate the rate of convergence in multidimensional Gaussian problems with and without importance sampling. There we demonstrate the necessity of importance sampling for target densities which depend on variables with a wide range of scales.

DOI: [10.1103/PhysRevE.78.046704](https://doi.org/10.1103/PhysRevE.78.046704)

PACS number(s): 05.10.Ln, 02.70.Tt, 02.50.Ng, 02.70.Rr

I. INTRODUCTION

Markov-chain Monte Carlo (MCMC) methods are important tools in parametric modeling [1,2] where the goal is to determine a posterior distribution of parameters given a particular dataset. Since these algorithms tend to be computationally intensive, the challenge is to produce algorithms that have better convergence rates and are therefore more efficient [3,4]. Of particular concern are situations where there is a large range of scales associated with the target density, which we find are widespread in models from many different fields [5–9].

In this paper we quantify the convergence of the MCMC method by the second largest eigenvalue in absolute value for the associated operator in L^2 . This is not the only numerical quantity that can be used to describe the convergence properties. Other authors quantify convergence with different metrics: computing the constant of geometric convergence

with respect to the total variation norm [10], monitoring sample averages [11], evaluating mixing of parallel chains [12], or looking at the integrated autocorrelation time of functions of the sample [13,14]. The connection between the second eigenvalue and total variation norm is discussed in [15]. To connect the second eigenvalue estimates to metrics based on autocorrelation, we would argue informally that the second eigenvalue determines the autocorrelation time of the slowest mixing function of the sample and as such represents a “worst” case for the length of time you would need to run the chain to reduce the variance of sample averages to a predefined level.

There are a number of techniques to either determine exactly or bound the second eigenvalue or the constant of geometric convergence for MCMC algorithms on discrete state spaces [16–19], but the methods for finding quantitative bounds for continuous state spaces require a more technical formulation. Where work has been done in that area, upper bounds on the convergence rate can be derived using purely analytical [10,20,21] or semianalytical techniques [22], but may not always be very useful for selecting parameters optimally. Therefore, in this work, we show that a conceptually straightforward variational method can provide convergence

*Also at UCD Conway Institute of Biomolecular and Biomedical Research, University College Dublin, Dublin 4, Ireland. fergal.casey@ucd.ie

rate estimates for continuous state space applications. In contrast to earlier closely related work [13,15], we move away from mathematical formalities, focussing from the start on specific examples and step through the calculations that provide the second eigenvalue bounds. In Roberts and Rosenthal [13], rules of thumb are provided for determining the optimal acceptance rate and step lengths for both the random walk Metropolis-Hastings algorithm and the Metropolis adjusted Langevin algorithm, in the asymptotic limit of infinite dimensions where it can be proved that those methods are approximated by diffusion processes. The rules are widely used as they are independent of the specific form of the target density, appear from numerical simulations to be appropriate far from the infinite dimensional asymptotic limit, and are easily implemented. In contrast, the approach proposed here is to establish convergence properties for particular MCMC algorithms based on their performance on simple target distributions without the need to set up a diffusion approximation in an infinite dimensional limit. Poor performance or lack of convergence on these simple distributions then indicates that further application with more complex target densities will also suffer from convergence problems. Conversely, an identification of a range of parameters which provide good convergence properties for simple target distributions may be used as a starting point for further applications. Even though we provide only lower bounds on the second eigenvalue we show these bounds can be remarkably tight due to careful choice of test functions, and computing the approximate convergence rate as a function of algorithm parameters allows us to optimally tune those parameters.

We have been able to obtain explicit formulas for one dimensional example problems but the method may be more generally applicable, when applied in an approximate way, as we demonstrate for a multidimensional problem.

II. MARKOV-CHAIN MONTE CARLO

Typically, one wishes to obtain a sample x_1, x_2, \dots from a probability distribution $\pi(x)$ which is sometimes called the *target* distribution. An MCMC algorithm works by creating a Markov chain that has $\pi(x)$ as its unique stationary distribution, i.e., after many steps of the chain any initial distribution converges to $\pi(x)$. A sufficient condition to establish $\pi(x)$ as the stationary distribution is that the chain be ergodic and that the *transition density*, $t(x, y)$, of the chain satisfy a detailed balance:

$$\pi(x)t(x, y) = \pi(y)t(y, x).$$

Given a *proposal density* $q(x, y)$ for generating moves, one way to construct the required transition density [23–25] is to define $t(x, y) = \alpha(x, y)q(x, y)$ where

$$\alpha(x, y) = \min\left(\frac{q(y, x)\pi(y)}{q(x, y)\pi(x)}, 1\right) \tag{1}$$

is the *acceptance probability* of the step $x \rightarrow y$. Obtaining the sample from the stationary distribution then involves letting the chain run past the transient (*burn-in*) time and recording iterates from the late time trajectory at time intervals exceed-

ing the correlation time. How long it takes to reach the stationary distribution determines the efficiency of the algorithm and for a given target distribution, clearly it depends on the choice of the proposal density. We can write down the one-step evolution of a probability density $p(x)$ as a linear operator:

$$\begin{aligned} (\mathcal{L}p)(y) &= \int t(x, y)p(x)dx + \left(1 - \int t(y, x)dx\right)p(y) \\ &= \int [t(x, y)p(x) - t(y, x)p(y)]dx + p(y), \end{aligned}$$

where $dx = dx_1 \cdots dx_n$, $dy = dy_1 \cdots dy_n$, n is the dimension of the state space, and all integrals are from $-\infty$ to ∞ here and elsewhere in this paper. The second form makes it explicit that $p(y) = \pi(y)$ is the stationary distribution by the detailed balance relation.

Now, if the linear operator has a discrete set of eigenfunctions and eigenvalues, it holds that the asymptotic convergence rate is determined by the second largest eigenvalue in absolute value (the largest being 1) [26,27]. We will write this eigenvalue as λ^* , and will refer to it as the *second eigenvalue* meaning the second largest in *absolute* value. Geometric convergence of the chain is ensured when $\lambda^* < 1$ [15], and then the discrepancy between the density at the m th iterate of the chain and the target density decreases as $(\lambda^*)^m$ for large m . Many previous authors have taken this second eigenvalue approach, in both the finite and continuous state space settings [17–19,22,28], as it provides a useful quantifier for the convergence rate. Ideally we would like algorithm parameters to be adjusted such that λ^* is as small as possible.

The variational calculation allows us to obtain an estimate for λ^* , but before we can do this we need to convert our operator into a self-adjoint form which ensures that the eigenfunctions are orthogonal. This is easily accomplished by a standard technique [16] of modifying the transition density by $s(x, y) = t(x, y)\sqrt{\pi(x)/\pi(y)}$ and our self-adjoint operator is then given by

$$(\mathcal{S}p)(y) = \int s(x, y)p(x)dx + \left(1 - \int t(y, x)dx\right)p(y) \tag{2}$$

$$= \int [s(x, y)p(x) - t(y, x)p(y)]dx + p(y), \tag{3}$$

where the “diagonal” part of the old operator [multiplying $p(y)$] need not be transformed using $s(x, y)$. It is easy to show that defined as above, \mathcal{S} is self-adjoint using the standard inner product in L^2 with respect to Lebesgue measure. Note that if $u(x)$ is an eigenfunction of the operator \mathcal{S} , then $\sqrt{\pi(x)}u(x)$ is an eigenfunction of the original operator \mathcal{L} with the *same* eigenvalue.

A. Metropolis-Hastings and smart Monte Carlo

We consider two MCMC algorithms which essentially differ only in the choice of proposal density and acceptance probability that is used in selecting steps. We will refer to the

random walk Metropolis-Hastings (RWMH) algorithm as that which uses a symmetric proposal density to determine the next move; for example, a Gaussian centered at the current point: $q(x,y)=\sqrt{|L|}/(2\pi)\exp[-(y-x)^T L(y-x)/2]$ where L is an inverse covariance matrix that needs to be chosen appropriately for the given problem (*importance sampling*). In other words, the proposed move from x to y is given by $y=x+R$ where $R\sim\mathbf{N}(0,L^{-1})$ is a normal random variable, the mean is 0, and the covariance is L^{-1} . Thus the update on the current state has no deterministic component. We will see that when the target density is not spherically symmetric, a naive implementation of the Metropolis-Hastings algorithm where the step scales are all chosen to be equal leads to very poor performance of the algorithm. As would be expected the convergence deteriorates as a function of the ratio of the true scales of the target density to the scale chosen for the proposal density.

One variant used to accelerate the standard algorithm is a *smart* Monte Carlo (SMC) method [29] that uses the gradient of the negative of the log target density at every step, $G(x)=-\nabla\ln[\pi(x)]$, to give

$$\begin{aligned}
 q(x,y) &= \frac{\sqrt{|L|}}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\{y-[x-H^{-1}G(x)]\}^T \right. \\
 &\quad \left. \times L\{y-[x-H^{-1}G(x)]\}\right) \quad (4)
 \end{aligned}$$

and H can be considered either as a constant scaling of the gradient part of the step or, if it is the Hessian of $-\ln[\pi(x)]$, as producing a Newton-like optimization step [30]. The move to y is generated as $y=x-H^{-1}G(x)+R$, so now we have a *random* component $R\sim\mathbf{N}(0,L^{-1})$ and a *deterministic* component $-H^{-1}G(x)$. Viewed like this, moves can be considered to be steps in an optimization algorithm (moving to maximize the probability of the target density) with random noise added. We will see that with an optimal choice of H and for Gaussian target densities, the smart Monte Carlo method can converge in one step to the stationary distribution. We will also see that for a one dimensional non-Gaussian distribution it actually fails to converge, independent of the values of the scaling parameters.

B. Variational method

Once we have the self-adjoint operator for the chain \mathcal{S} from Eq. (3), and we know the eigenfunction with eigenvalue $\lambda_1=1$, $\sqrt{\pi(x)}$, we can look for a candidate second eigenfunction in the function space orthogonal to the first eigenfunction where the inner product is defined by $(p_1,p_2)=\int p_1(x)p_2(x)dx$. Given a family of normalized candidate functions in this space, $v_a(x)$, with variational parameter a , the variational principle [26,31] states

$$\max_a |(v_a, \mathcal{S}v_a)| \leq \lambda^* \leq 1 \quad (5)$$

and depending on how accurately our family of candidate functions captures the true second eigenfunction, this can give quite a close approximation to the second dominant eigenvalue. In the problems we examine in the following sections the target densities have an even symmetry which

makes it straightforward to select a variational trial function: any function with odd symmetry will naturally lie in the orthogonal space. For more complicated problems with known symmetries this general principle may be useful in selecting variational families for the purposes of algorithm comparison. Another approach to constructing the variational family is shown in the section on multidimensional target densities: choose the test function as a linear combination of two functions, one with the properties that are required (i.e., slow convergence to the target distribution) and then the additional term is used merely to preserve orthogonality.

This variational method for providing a lower bound to the second eigenvalue of the MCMC algorithm was foreshadowed by a similar approach of Lawler and Sokal [26]. These authors considered the flow of probability out of a subset A of the state space; in our language, their test functions were confined to the family $v_A(x)=[\pi(A^c)\chi_A - \pi(A)\chi_{A^c}]/[\pi(A)\pi(A^c)]$, where χ_A is the indicator function on the set A . By allowing for more general test functions, we establish not only rigorous but also relatively tight bounds on convergence rates, providing guidance for parameter optimization and algorithm comparisons.

Writing out explicitly for \mathcal{S} in $(v_a, \mathcal{S}v_a)$ we have

$$\begin{aligned}
 (v_a, \mathcal{S}v_a) &= \int \int v_a(x)s(x,y)v_a(y)dx dy - \int \int t(y,x) \\
 &\quad \times [v_a(y)]^2 dx dy + 1. \quad (6)
 \end{aligned}$$

As we will see in the following section, the lower bound in Eq. (5) can be arbitrarily close to 1 and therefore equality holds. In these situations we can also show that the chain does not converge geometrically, based on the total variation norm definition of geometric convergence [10]. However, whether the type of convergence changes or not, we still refer to the magnitude of the second eigenvalue estimate in determining efficiency of the algorithm. The rationale is that the second eigenvalue determines the longest possible autocorrelation time of a function of the MCMC sample; the worst case autocorrelation time will be of the order $1.0/\ln(\lambda^*)$ which could be extremely long. We will also see that there can be eigenvalues in the spectrum that are close to -1 which determine the asymptotic convergence rate, i.e., $\lambda^*=|\lambda_n|$ where $\lambda_n<0$. Interestingly, for this situation there is oscillatory behavior of the Markov chain.

III. EXAMPLES

A. Gaussian target density

Consider the simplest case of a one dimensional Gaussian target density $\pi(x)=\sqrt{k/(2\pi)}\exp(-kx^2/2)$ with variance $1/k$. Under the RWMH algorithm, the proposal density is

$$q(x,y) = \sqrt{\frac{l}{2\pi}} \exp\left(-\frac{1}{2}l(y-x)^2\right). \quad (7)$$

The issue is to determine l optimally; a first guess would be that $l=k$ is the best choice. The rationale behind this is that since the target and proposal densities have the same form, if they also have the same scales, then the convergence rate




FIG. 1. Variational estimate on the second eigenvalue for the one dimensional Gaussian problem using the RWMH method, with $k=1.0$. The variational estimate is the solid line and the empirically determined values are marked with stars. Some of the empirical values seem to be less than the lower bound, but this is due to inaccuracies in their estimation. The optimum occurs at $l \approx 0.165$.

might be expected to be optimal. We will see that this is not actually correct.

To begin, define a variational function $v_a(x) \propto \exp(-ax^2/2)$, orthogonal to the target density and normalized such that $\int v_a^2 dx = 1$. We can motivate this choice by recognizing that any initial distribution that is asymmetric will most likely have a component of this test function, and a convergence rate estimate based on it roughly corresponds to how fast probability “equilibrates” between the tails. More commonly, variational calculations will use linear combinations of many basis functions with the coefficients as variational parameters. We find here that including higher order terms in the test function is unnecessary as we obtain tight enough bounds just retaining the lowest order term.

We proceed by evaluating Eq. (6) noting that because of the form of the acceptance probability, Eq. (1), there are two functional forms for the kernels $t(x,y)$ and $s(x,y)$ depending on the sign of $y^2 - x^2$, i.e., whether the “energy” change, $\Delta E(x,y) = -\ln[\pi(y)] + \ln[\pi(x)] = k(y^2 - x^2)/2$, is positive or negative. It is then convenient to use the coordinate transformation $y = rx$, $x = x$ or $x = ry$, $y = y$ where $-1 \leq r \leq 1$ and $-\infty \leq x, y \leq \infty$ to evaluate the integrals. An explicit expression for $(v_a, S v_a)$ can be obtained for this case of a Gaussian target density.

Next, we use a numerical optimization method to maximize the bound defined by Eq. (5) with respect to a . The result of this analysis is shown in Fig. 1 along with an empirically determined convergence rate for comparison. (To obtain the rate empirically, we run the MCMC algorithm for many iterates on a random initial distribution and observe the pointwise differences from the distribution of the m th iterate and the target distribution for large m . These differences are either fit using Hermite polynomial functions or by looking for the multiplicative factor which describes the geometric decay of the m th difference from one iterate to the next.) The

variational bound tightly matches the empirical obtained eigenvalue estimates in this case, and an optimum step size l can be ascertained. Clearly our $l=1$ initial guess for the best scaling is far from optimal.

It is also worth comparing the optimal step scale with those obtained from different methods. In [13,14], a derivation of the optimal step size and acceptance rate is proposed based on minimizing the integrated autocorrelation time of an arbitrary function of the chain’s states in stationarity. By approximating the chain as an infinite dimensional diffusion process, formulas are derived for the optimal scaling of steps. For our one dimensional Gaussian target density, the proposal density’s optimal variance is suggested to be $(1/2.38)^2 = 0.176$ which is surprisingly close to the estimate we have obtained using the variational method from Fig. 1, $l \approx 0.165$. However, given the infinite dimensional limit in which the former approximation is made, and the different convergence criterion based on autocorrelation time rather than second eigenvalue, the agreement may be merely coincidental.

Moving to the one dimensional smart Monte Carlo, we have a Gaussian proposal density of the form

$$q(x,y) = \sqrt{\frac{l}{2\pi}} \exp\left\{-\frac{1}{2}l\left[y - \left(x - \frac{k}{h}x\right)\right]^2\right\}, \quad (8)$$

where $1/l$ is the variance of the random part of the step and $1/h$ is the scale of the deterministic part. [Letting $h \rightarrow \infty$ we recover the RWMH algorithm of Eq. (7).]

Taking $h=k$ corresponds to performing a Newton step at every iterate of the algorithm. Thus, since the log of the target density is purely quadratic, the current point will always be returned to the extremum at 0 by the deterministic component of the smart Monte Carlo step and the random component will give a combined move drawn from $q(x,y) = q(y) = \sqrt{l/(2\pi)} \exp(-ly^2/2)$, which has the form of an *independence sampler* [23]. If we then also choose $l=k$, we see immediately that we are generating moves from the target distribution from the beginning, i.e., we have convergence in one step starting from *any* initial distribution.

In real problems, however, $-\ln[\pi(x)]$ will not be quadratic. We may obtain an estimate for l and h by considering its quadratic approximation or curvature but in many cases those estimates will have to be adjusted. If the curvature is very small (or in multidimensional problems if the quadratic approximations are close to singular), the parameters will have to be increased to provide a step size control to prevent wildly unconstrained moves (analogous to the application of a trust region in optimization methods [30]). If the curvature is large but we believe that the target density is multimodal, we need to decrease the parameters to allow larger steps to escape the local extrema. Therefore we examine in the following the dependence of the convergence rate as we vary both of the parameters l and h .

The acceptance probability Eq. (1) has two functional forms separated by a boundary in the (x,y) plane given by

$$\left[k + l\frac{k}{h}\left(-2 + \frac{k}{h}\right)\right](y^2 - x^2) = b(k,h,l)(y^2 - x^2) = 0. \quad (9)$$

In particular, the acceptance probability is




FIG. 2. Regions in the xy plane where acceptance probability $\alpha(x,y) < 1$ or $\alpha(x,y) = 1$, when (a) $b(1,h,l) \geq 0$ and (b) $b(1,h,l) < 0$. The equation for the boundary is shown in (c); see Eq. (9) with $k=1.0$. [The RWMH algorithm will only have regions described by (a).]

$$\alpha(x,y) = \min \left[\exp \left(-\frac{1}{2} b(k,h,l)(y^2 - x^2) \right), 1 \right]. \quad (10)$$

Now we have a complication over the RWMH method because depending on the sign of the coefficient function $b(k,h,l)$ in Eq. (9), we find that either $\alpha(x,y) < 1$ on $|y| \geq |x|$, $\alpha(x,y) = 1$ on $|y| < |x|$ or vice versa. This is shown in Fig. 2.

As before, for a given value of h and l , we need to break up the double integrals of the scalar product $(v_a, S v_a)$, Eq. (6), into the appropriate regions. Our choice of variational function is the same as before (since the target density is the same) and we again can get an explicit (but complicated) expression for Eq. (6) which we maximize with respect to a . The results of this analysis are shown in Fig. 3(a), where we fix $k=1.0$ and vary h, l . We have confirmed that these lower bounds are quite accurate as shown in Fig. 3(b).

The remarkable feature of these results is that even for this simple Gaussian problem, the selection of step scale parameters h, l is critical to achieve convergence. As already mentioned, there is a trivial choice of optimum with $h=l=k=1$ that gives one step convergence from any initial dis-

tribution (and therefore $\lambda^*=0$). However, if we change parameters infinitesimally such that $l=1+\epsilon, h=1$ ($\epsilon > 0$) we go through a discontinuous transition where we see no convergence from *any* initial distribution. This can be understood by recognizing that after one step we will have a proposal density (before accept or reject) $\propto \exp[-(1+\epsilon)x^2/2]$ which has a factor $\exp(-\epsilon x^2/2)$ less probability in its tails than the target density. Suppose there is an initial distribution or point mass concentrated at $x = \sqrt{2M}/\sqrt{\epsilon}$, $M \gg 1$. The proposed step of the smart Monte Carlo algorithm, starting at x , will revisit x too infrequently by a factor $\exp(-M)$. Thus detailed balance will force the transition $x \rightarrow 0$ to be accepted with a probability of only $\exp(-M)$, and thus the initial distribution will take an arbitrarily long time to converge to the target density.

More formally, we can compute the probability of rejection, $r(x) = 1 - \int t(x,y) dy$, when h, l are as above and we find




FIG. 3. Estimate of second eigenvalue for the symmetrized smart Monte Carlo operator. (a) $k=1$ is fixed and h, l are allowed to vary. $h=1.0, l=1.0$ is the optimal scaling for deterministic and random parts of the step. The solid diagonal line is the parameter restrictions that yield the Metropolis adjusted Langevin algorithm (MALA) algorithm (see text). (b) We take a slice through this surface at $l=1.5$ and empirically determine the second eigenvalue at points along this curve (stars). The error bars are too small to be seen. Dashed lines are discontinuities.




FIG. 4. Forty iterates of the smart Monte Carlo algorithm (solid lines), Eq. (8), when the initial distribution is normal with standard deviation five times the Gaussian target density (dashed line). Parameters are chosen to be in the region of no convergence ($h=2.0$, $l=1.5$); see Fig. 3(a). We see that the tails of the initial distribution are essentially unchanging after many iterates and have failed to converge to the target density

$$r(x) = \Phi\left(\sqrt{\frac{1+\epsilon}{2}}x\right) - \exp\left(\frac{1}{2}b(k,h,l)x^2\right) \times \frac{1+\epsilon}{1+\epsilon+b(k,h,l)}\Phi\left(\frac{1+b(k,h,l)+\epsilon}{2}x\right),$$

where $b(k,h,l) < 0$ and Φ is the cumulative normal (0, 1) distribution function. We note that $\text{ess sup } r(x) = 1$ by continuity of $r(x)$, and then we use Proposition 5.1 from [32] to conclude that the Markov chain is no longer geometrically convergent for these values of h and l .

In fact this is only one of the two disconnected regions where no convergence is observed in Fig. 3. The largest of the two (with $h > 1/2$) is defined exactly by the equation $b(1,h,l) < 0$ [compare Fig. 2(c) with Fig. 3(a)]. In this region the bound on the second eigenvalue approaches 1 as the variational parameter, $a \rightarrow 0$. This corresponds to a perturbation on the target density of $x\sqrt{\pi(x)}$ for the unsymmetrized MCMC operator \mathcal{L} . In other words, we have a test distribution that has exponentially more probability in its tails than the target density. For initial states x arbitrarily far away from the origin, the acceptance probability $\alpha(x,y)$ in the region $|y| < |x|$ is arbitrarily small. To see this, note that Eq. (10) is an exponentially decaying function of $y^2 - x^2$ in this region, and given the form of the proposal density Eq. (8), we see that the expected value of $y^2 - x^2$ is arbitrarily large and negative. Thus states far out will never be “allowed back” and the fat tails of $\sqrt{\pi(x)}$ will never shrink back down those of $\pi(x)$. Furthermore, moves $x \rightarrow y$ where $|y| \geq |x|$ are always accepted [because $\alpha(x,y) = 1$ on $|y| \geq |x|$] which simultaneously prevents convergence. The situation is analogous to that described for $l = 1 + \epsilon$ and $h = k = 1$, except now there is a cutoff both on the deterministic step and the random step. A typical example of this is shown in Fig. 4. Once we cross to the $b(1,h,l) \geq 0$ region, moves $x \rightarrow y$ where $|y| < |x|$ are always accepted by Eq. (10) [Fig. 2(a)]. Therefore

excess probability in the tails is allowed to flow back into the central part of the distribution and the convergence is not blocked.

In the second region where no convergence is observed ($h < 1/2$ in Fig. 3), we have a situation where the deterministic step alone (taking $l \rightarrow \infty$) leads to the proposed moves being generated by an unstable mapping, from the $(n-1)$ th to n th iterate: $x^{(n)} = x^{(n-1)} - \beta x^{(n-1)}$ where $\beta > 2$. The trial variational function for this situation also maximizes the bound as $a \rightarrow 0$, again implying that the tails are not decaying to the stationary distribution. The reason is that, even when $l < \infty$, we have a situation in which the *expected* or mean position of a state x after one step is y where $|y| \geq |x|$. Thus excessive probability in the tails cannot be shifted inward to match the target density.

The lack of convergence in this region was already noted for the Metropolis adjusted Langevin algorithm (MALA), a special case of the SMC algorithm where $h = 2l$. As shown in [33], if π is bounded, a sufficient condition for MALA to fail to be geometrically ergodic is

$$\liminf_{|x| \rightarrow \infty} \frac{|\nabla \ln \pi(x)|}{|x|} > \frac{4}{s},$$

where s is the single stepsize control parameter for that algorithm. The equivalence to the SMC method is established by setting $l = 1/s$. Thus, for the Gaussian target density π , the condition is $l < 1/4$. Referring to the solid white line in Fig. 3(a), the nonconvergent parameter regime for MALA lies along the line segment $h = 2l$ with $l < 1/4$ which matches exactly with the boundary we have determined using the variational method.

The $h = 1/2$ “trough” is a special case where we have oscillatory behavior. That is, the second eigenvalue is negative but greater than -1 and in fact convergence does occur. Interestingly setting $h = k/2$ means that $b(k,h,l) = k$ and the acceptance probability of Eq. (10) looks again like that of the RWMH algorithm, but the convergence is actually faster. In a sense, given that the deterministic part of the step moves $x \rightarrow -x$ and the target distribution is symmetric, the oscillatory behavior allows the chain to sample the distribution twice as fast.

B. Quartic target density

In scientific or statistical applications where MCMC is used, the log of the target density will ordinarily have higher order terms beyond the quadratic order we studied in the previous section. For example, in a Bayesian inference problem the posterior distribution will rarely have a simple Gaussian form. Both finding the maximum *a posteriori* parameter estimates and sampling from the posterior are made more difficult in the presence of these higher order terms.

Therefore we wish to extend the previous example by studying a target distribution of the form $\pi(x) = [2^{(3/4)}k^{(1/4)}/\Gamma(1/4)]\exp(-kx^4/2)$. Here, the log of the target density is quartic and the proposal density (Gaussian) no longer has the same form as the target density. We would like to understand the performance of the Monte Carlo algorithms in this circumstance. [The test distribution is taken to




FIG. 5. Second eigenvalue estimate from the variational method (solid line) and empirical estimates (stars), for the quartic target density ($k=1$) using the RWMH method, Eq. (7). The numerical values for λ^* are now estimated by taking the ratio of the discrepancy from the target density in subsequent iterates and finding a single multiplicative factor which describes the decay. This is done rather than using functional forms analogous to Hermite polynomials to fit the decay, because it appears that there may be more significant contributions from higher order terms. This also explains why the lower bound shown differs more than in Figs. 1 and 3(b). The data point shown at $l=1/\sqrt{2} \approx 0.71$ (see text) does not appear to be optimal.

be $\propto x \exp(-ax^4/2)$, i.e., in the orthogonal space to the stationary distribution.]

The goal is to estimate the optimal value of l , as before. We can argue approximately that the step scale should be such that $kx^4/2 \approx 1$ for a typical move x , i.e., the change in energy is about 1 and the acceptance probability is therefore $\exp(-1)$. This gives a typical value for $x^2 = \sqrt{2}/\sqrt{k}$. Since the proposal density is Gaussian with variance $1/l$, we therefore would naively predict $l = \sqrt{k}/\sqrt{2}$. Applying the variational method, we were unable to find a closed form solution to Eq. (6) so we had to resort to numerical integrals in determining the bound in Eq. (5). The results are shown in Fig. 5 for the RWMH method; it suggests an optimal choice for the step size parameter l , which is an improvement over our initial guess of $1/\sqrt{2} \approx 0.71$ (when $k=1$). Using the formulas for the optimal step scale from [13] coincidentally yields about 0.71, also a little off from our variational estimate.

Turning to the smart Monte Carlo algorithm, if we wish to make the deterministic part of the proposed move a Newton step using the Hessian of $-\ln[\pi(x)]$ at $x=0$ we are left with a singular Hessian and an infinite deterministic step, reinforcing the need for the step length control parameter, h .

Surprisingly, we find that, independent of the value of h and l (k fixed at 1), the scalar product $(v_a, Sv_a) \rightarrow 1$ as $a \rightarrow 0$. Thus there are no choices of scaling parameters which will lead to convergence. This is borne out by numerical simulation; see Fig. 6 for the changes in an initial density under many iterates of the algorithm with an arbitrary choice for s, h .




FIG. 6. Forty iterates of the smart Monte Carlo algorithm (solid lines), Eq. (8), when the target density is quartic (dashed line). The initial distribution of points is normal with standard deviation about five times that of target density (dashed line). Parameters are arbitrarily chosen as ($h=1.0, l=1.0$), and we see that the tails of the initial distribution are unchanged for every iterate of the algorithm. Other parameter sets tested lead to the same behavior.

The failure of the smart Monte Carlo method for the quartic problem is clearly due to nonconvergence of the tails of the distribution, and can be seen by analyzing the integrals defining the operator, Eq. (6), and noting that they all tend to zero as the variational parameter tends to zero, independent of the choice for k, h , and l .

As a partial check on this result, we again apply the condition derived in [33] for the MALA algorithm, which states that geometric convergence is not possible when $\liminf_{|x| \rightarrow \infty} |\nabla \ln \pi(x)|/|x| > 4/s$ where s is the step scale parameter. Now, for the quartic density π , the quantity on the left of the inequality $\rightarrow \infty$, so *no* value of s can give geometric convergence. MALA is a special case of the SMC algorithm, but we have shown here that the latter also has convergence problems indicated by $\lambda^* \rightarrow 1$, for *all* values of its scaling parameters, h and l .

The Gaussian and quartic problems are representative examples of target densities on which we have tested the smart Monte Carlo method. As we have seen there are severe convergence problems on these distributions. We would expect that for real applications, where the log of the target density would contain components of these and higher order nonlinearities, similar convergence difficulties for the smart Monte Carlo method would occur. It may well be that in applications where the method is extensively used (e.g., [34–36]) the convergence criteria are less precise than ours. For example, it may be acceptable to merely monitor the variance of some function of the state space variables and conclude that convergence has been achieved when it ceases to change appreciably, or as in [13], define efficiency by the integrated autocorrelation time.

IV. MULTIDIMENSIONAL TARGET DENSITIES

For multidimensional problems, it is quite common to find a large range of scales associated with the target density

[5,9,8]. That is, the curvature of the probability density along some directions in the parameter space is much larger than in other directions. Clearly, if an MCMC method is not designed to take these different scales into account through importance sampling, the algorithm will perform very poorly. If the curvature is very high in a particular direction and we try to take a moderately sized step, it will almost certainly be rejected but if we take small steps in directions that are essentially flat the MCMC algorithm will be very slow to equilibrate. We would like to show explicitly here what happens to the convergence rate when the scale of the problem has been underestimated or overestimated.

The variational calculations for the one dimensional examples of the previous section either yielded explicit formulas or gave integrals that were relatively fast to compute numerically. However, as we go to multiple dimensions neither of these features are present, in general. Typically the integrals describing (v_a, Sv_a) will *not* factor into one dimensional integrals. For Gaussian target densities the full space is broken into regions analogous to those in Fig. 2, described by an equation like $y^tAy \geq x^tAx$ where A is a symmetric n by n matrix which is not necessarily positive definite. For the RWMH algorithm applied to a multivariate Gaussian target density with inverse covariance matrix K , we have $A=K$, and therefore all the dimensions are coupled through the energy change, $\Delta E=y^tKy-x^tKx$. We would still like to be able to get a lower bound on λ^* , and to this end note that *any* test function orthogonal to the target density will work in Eq. (5); the bound does not explicitly require a variational parameter, however without it the estimate will be less accurate. It is still necessary to make choices for the test functions that are both tractable in computing (v, Sv) and are “difficult” functions for the given algorithm to converge from, i.e., have a significant component along the true second eigenfunction.

As an example, take the multivariate Gaussian distribution of the form

$$\pi(x) = \frac{\sqrt{|K|}}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2}x^tKx\right) \quad (11)$$

with $x=(x_1, \dots, x_n)$, and consider using the MH algorithm with importance sampling, i.e.,

$$q(x,y) = \frac{\sqrt{|L|}}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2}(y-x)^tL(y-x)\right),$$

where again L is the inverse covariance matrix or step size control term and to simplify we assume that both K and L are diagonal. Without any analysis we might guess that the optimum choice for L is K .

First we construct a test function that will provide a useful bound when the proposed steps are too large for the natural scale of the problem. For simplicity, consider putting a delta function distribution at the origin. If we take large steps the acceptance probability should be low and there will be a large overlap between the initial state and the final state. In the limit that the proposed steps have infinite length, the initial state will not be changed at all and the bound on the second eigenvalue in absolute value will approach 1. To do this more carefully we define a test function which is a

Gaussian whose variance will ultimately be taken to zero to represent the delta function. However, we also need to add another term to ensure the test function is orthogonal to the target density, in order to apply the variational bound. Therefore for the unsymmetric operator we write the test function as $u_\sigma(x)=-A\pi(x)+Bw_\sigma(x)$ where $w_\sigma(x)$ is the probability density for a multivariate Gaussian with covariance matrix σ^2I and A and B are constants. For the symmetrized operator the trial function is transformed to $v_\sigma(x)=-A\sqrt{\pi(x)}+Bw_\sigma(x)/\sqrt{\pi(x)}$. A and B are constrained to satisfy the orthogonality relation $(v_\sigma, \pi)=0$ and a normalization $(v_\sigma, v_\sigma)=1$. These lead to the conditions

$$A = B \quad \text{and} \quad B^2 \int \left(\frac{w_\sigma(x)}{\sqrt{\pi(x)}}\right)^2 dx = 1 + B^2.$$

Then it can be seen that

$$(v_\sigma, Sv_\sigma) = -B^2 + B^2 \left(\mathcal{S} \frac{w_\sigma(x)}{\sqrt{\pi(x)}}, \frac{w_\sigma(x)}{\sqrt{\pi(x)}} \right),$$

where we have used the orthogonality condition, the fact that $w_\sigma(x)$ integrates to 1 and that \mathcal{S} is self-adjoint. Writing out the operator \mathcal{S} explicitly we get

$$\begin{aligned} \left(\mathcal{S} \frac{w_\sigma(x)}{\sqrt{\pi(x)}}, \frac{w_\sigma(x)}{\sqrt{\pi(x)}} \right) &= \int \int \frac{w_\sigma(x)}{\sqrt{\pi(x)}} s(x,y) \frac{w_\sigma(y)}{\sqrt{\pi(y)}} dx dy \\ &\quad - \int \int t(x,y) \left(\frac{w_\sigma(x)}{\sqrt{\pi(x)}} \right)^2 dx dy \\ &\quad + \int \left(\frac{w_\sigma(x)}{\sqrt{\pi(x)}} \right)^2 dx. \end{aligned}$$

The last term on the right hand side is $(1+B^2)/B^2$, making use of the normalization condition, so we are left with

$$\begin{aligned} (v_\sigma, Sv_\sigma) &= B^2 \int \int \frac{w_\sigma(x)}{\sqrt{\pi(x)}} s(x,y) \frac{w_\sigma(y)}{\sqrt{\pi(y)}} dx dy - B^2 \int \int t(x,y) \\ &\quad \times \left(\frac{w_\sigma(x)}{\sqrt{\pi(x)}} \right)^2 dx dy + 1. \end{aligned}$$

Since we are ultimately taking a limit as $\sigma \rightarrow 0$ ($w_\sigma \rightarrow$ a delta function) we can make approximations to these integrals as follows:

$$\begin{aligned} \int \int \frac{w_\sigma(x)}{\sqrt{\pi(x)}} s(x,y) \frac{w_\sigma(y)}{\sqrt{\pi(y)}} dx dy \\ \approx s(0,0) \int \int \frac{w_\sigma(x)}{\sqrt{\pi(x)}} \frac{w_\sigma(y)}{\sqrt{\pi(y)}} dx dy \end{aligned}$$

and

$$\int \int t(x,y) \left(\frac{w_\sigma(x)}{\sqrt{\pi(x)}} \right)^2 dx dy \approx \int t(0,y) dy \int \left(\frac{w_\sigma(x)}{\sqrt{\pi(x)}} \right)^2 dx.$$

Finally by taking $\sigma \rightarrow 0$ we have the expression

$$(Sv_0, v_0) = 1 - \int t(0, y) dy.$$

As already mentioned, for the multidimensional problem we expect different functional forms for the kernels $s(x, y)$ and $t(x, y)$ depending on the initial and final state (x, y) and this is what makes decoupling the integrals difficult. However, for this choice of test function the equation for the boundary (with $x=0$) is given by $y^t Ky=0$ and since K is positive semidefinite we always stay on one side of the boundary (the energy never decreases from the initial distribution placed at $x=0$). Then

$$(Sv_0, v_0) = 1 - \frac{\sqrt{|L|}}{(2\pi)^{n/2}} \int \exp\left(-\frac{1}{2}y^t(K+L)y\right) dy \quad (12)$$

$$= 1 - \prod_{i=1}^n \sqrt{\frac{l_i}{l_i + k_i}}, \quad (13)$$

where l_i and k_i are the diagonal elements of the diagonal matrices L and K , respectively. With no importance sampling we would have $L=kI$ where k would be chosen to make sufficiently large steps to enable it to sample $\pi(x)$. A rough argument as follows can give some insight into the form of Eq. (13): $1/\sqrt{l_i}$ is a measure of the scale in the i th coordinate direction of the proposal density, $1/\sqrt{k_i}$ is the scale in the i th coordinate direction of the target density. Suppose that $l_i \ll k_i$ for each i , that is the scales of the proposal density are too large in all directions. Then the ratio of the mean volume of moves generated by $q(0, y)$ to the volume occupied by $\pi(y)$ is exactly $\prod_{i=1}^n \sqrt{l_i}/\sqrt{k_i}$. Intuitively, this ratio is proportional to the acceptance probability, and in the regime $l_i \ll k_i$ the acceptance probability determines the convergence properties.

We want to use Eq. (13) to show how choosing step sizes *too large* even in one direction will result in a very inefficient algorithm. Suppose that for all but one of the directions we make $l_i=k_i$, $i=1, \dots, n-1$ which would be roughly the correct scaling in those directions. Then the bound on the second eigenvalue is

$$(Sv_0, v_0) = 1 - \sqrt{\left(\frac{1}{2}\right)^{n-1}} \sqrt{\frac{1}{1 + k_n/l_n}}. \quad (14)$$

From this we can see that as we go to larger and larger step sizes relative to the scale in the last direction ($k_n/l_n \rightarrow \infty$), the bound on λ^* increases to 1. Conversely we can argue that if one of the directions of the target density has a scale that is considerably smaller than the step scales being used in the proposal density, we will get very few acceptances and the convergence rate will be close to 0. Hence we see explicitly the need for importance sampling to accelerate convergence.

We would also like to address what happens in the other limit as the step size becomes excessively small compared to the natural scale of the problem. [In fact, Eq. (13) gives a lower bound of zero in that case which is not surprising as it is based essentially on the term in the operator equation which gives the probability of *staying* at the current state. If we take infinitesimally small steps, the acceptance probab-

ity will be one and we will never stay at the current state.] When the step scales are infinitesimally small we expect intuitively that the bound on the second eigenvalue will also approach 1; even though the acceptance ratio is close to 1, very small steps will never be able to “explore” the target distribution sufficiently. To compute this limit, we propose a test function which has components of the target density in all directions except the last, where it has an antisymmetric form to make sure it is orthogonal to the target density. With respect to the symmetrized operator \mathcal{S} this means

$$v(x) \propto x_n \prod_{i=1}^n \sqrt{\pi_i(x_i)}. \quad (15)$$

Here $\sqrt{\pi_i(x_i)}$ is the one dimensional Gaussian density which is the i th factor in a diagonalized multivariate Gaussian density. [Recall that since $\pi(x)$ is an eigenfunction of \mathcal{L} , then $\sqrt{\pi(x)}$ is an eigenfunction of \mathcal{S} .] We still have the problem of decoupling the n -dimensional multivariate problem into n one dimensional problems. To manage this we use a device to re-express the operator equation, Eq. (6), explicitly in terms of the change $\frac{1}{2}(y^t Ky - x^t Kx)$. [i.e., $-\ln \frac{\pi(y)}{\pi(x)}$], which we denote by ΔE . That is

$$(v, Sv) = \int \int v(x)s(x, y)v(y) dx dy - \int \int t(x, y)[v(x)]^2 dx dy$$

$$+ 1 = \int \int x_n \pi(x) q(x, y) \left[\int \min(e^{-\Delta E}, 1) \right. \\ \left. \times \delta\left(\Delta E - \frac{1}{2} \sum_{i=1}^n k_i (y_i^2 - x_i^2)\right) d\Delta E \right] dx dy$$

$$- \int \int x_n^2 \pi(x) q(x, y) \left[\int \min(e^{-\Delta E}, 1) \right. \\ \left. \times \delta\left(\Delta E - \frac{1}{2} \sum_{i=1}^n k_i (y_i^2 - x_i^2)\right) d\Delta E \right] dx dy.$$

Then we use the integral representation of the delta function $\delta(x) = \frac{1}{2\pi} \int \exp(-iwx) dw$, factor $q(x, y) = \prod_{i=1}^n q_i(x_i, y_i)$, and rearrange the order of integration to give

$$(v, Sv) = \frac{1}{2\pi} \int \min(\exp(-\Delta E), 1) \\ \times \left(\int A(w) \exp(-iw\Delta E) dw \right) d\Delta E, \quad (16)$$

where $A(w)$ contains the integration over the now decoupled (x, y) coordinates:




FIG. 7. Contour used to evaluate Eq. (19) when $\Delta E < 0$. r_1 is a branch point and $r_{1,0}$ is a pole of order $(n-1)/2$. The contour is the same for $\Delta E > 0$ except restricted to the negative imaginary plane.

$$A(w) = \left(\prod_{i=1}^{n-1} \int \int \pi_i(x_i) q_i(x_i, y_i) \exp\left(\frac{1}{2} i w k_i (y_i^2 - x_i^2)\right) dx_i dy_i \right) \quad (17)$$

$$\times \int \int (x_n y_n - x_n^2) \pi_n(x_n) q_n(x_n, y_n) \times \exp\left(\frac{1}{2} i w k_i (y_n^2 - x_n^2)\right) dx_n dy_n \quad (18)$$

$$= \prod_{i=1}^{n-1} \frac{1}{[1 + (k_i/l_i)w(-i+w)]^{1/2}} \frac{i(k_n/l_n)w}{\left[1 + \frac{k_n}{l_n}w(-i+w)\right]^{3/2}}. \quad (19)$$

Note that the complex integral with respect to dw has a branch point at the roots of $[1 + \frac{k_n}{l_n}w(-i+w)]^{3/2}$ which lie on the imaginary axis at r_1 and r_2 . It simplifies the analysis to consider the situation $k_i=l_i$ for $i=1, \dots, (n-1)$ and assume that $n-1$ is even. This way, the roots of $[1 + w(-i+w)]^{(n-1)/2}$, $r_{1,0}$ and $r_{2,0}$, are $(n-1)/2$ order poles and not branch points, also on the imaginary axis. If we now also assume that $k_n < s_n$, then we can take a contour as shown in Fig. 7 when $\Delta E < 0$ and a similar one in the lower imaginary plane when $\Delta E > 0$. Thus Eq. (19) is reduced to a residue term and a real integral which needs to be evaluated numerically. The result is plotted for $n=11$ in Fig. 8 along with the bound that came from Eq. (14). Thus we see the trade off between taking large steps that potentially can explore the space quickly but have a higher chance of being rejected and taking small steps which will have a high acceptance probability but will be unable to sample the space quickly. As we saw when doing the full variational calculation for the one dimensional problems, the best step scale to use is not what we may have guessed; the natural choice $l_n=k_n=1$ here does not appear to minimize the second eigenvalue. We believe this kind of ‘‘approximate’’ variational approach may be a useful way to deal with problems which are difficult to analyze otherwise.

V. CONCLUSION

By applying a variational method, it is possible to obtain an accurate (lower bound) estimate for the second eigenvalue




FIG. 8. Lower bound on second eigenvalue for the multivariate Gaussian problem, Eq. (11), with $n=11$. Step scale $= \sqrt{1/l_n}$. $k_n=1$ sets the scale of the target density in the last direction. The test function is chosen as the negative of the target density perturbed by a delta function (solid line) or as the target density itself in all directions but the last (dashed line). The estimate for the lower bound is a maximum of the two curves.

of an MCMC operator and thus bound the asymptotic convergence rate of the chain to the target distribution. Given such an estimate we can optimally tune the parameters in the proposal distribution to improve the performance of the algorithm. The procedure has a role to play between the various numerical algorithms that perform convergence diagnostics before the full simulations are run, to allow the user to manually tune parameters, and the adaptive schemes [3,37] that require no preliminary exploration. The simulations we performed to confirm our variational bounds in the case of a one dimensional target density and varying one step scale parameter, Figs. 1 and 3(b), would be infeasible to do as we move to higher dimensions and as we vary additional algorithm parameters. It is in those situations that the variational method can more efficiently identify regions of optimality.

In addition, the variational method allows us to discover weaknesses in variants of the random walk Metropolis-Hastings algorithm which on the surface appear to be reasonable prescriptions for sampling the target density. This is most dramatically seen in the smart Monte Carlo method discussed above which apparently has serious flaws for even the simplest of one dimensional target densities. Although the smart MC method has been widely used in molecular dynamics applications [34–36] the scales are often chosen by physical considerations (for example, to not exceed significantly the step sizes needed to accurately describe the dynamical evolution of the system) and furthermore, the diagnostics of convergence are not as rigorous as ours; typically a physical quantity is monitored until it appears to reach an equilibrium value, the rare events which correspond to the tails of the target distribution are possibly of lesser importance in those studies. Therefore the convergence problems we have discussed here specifically in relation to the smart Monte Carlo method, to our knowledge, have not been previously examined. Presumably the convergence problems can be corrected by a more careful discretization of the underlying diffusion equations, as was shown for the related Langevin-type methods [38].

The smart Monte Carlo and related Langevin-type methods use local information of the target density to try and improve the efficiency in making Monte Carlo moves. It would be interesting to examine convergence issues for the more broadly used gradient based hybrid MC algorithms [39] and other nonadaptive accelerated methods (e.g., parallel tempering [40]) where the alternative techniques for determining convergence via diffusion approximations may be harder to apply. Other MCMC methods, such as biased Metropolis algorithms (BMAs), attempt to capture an approximation of the entire target distribution for use in the proposed moves, by discretizing the target's cumulative distribution function (CDF) ([41,42]). These approaches are not suitable for large-dimensional state spaces where the target distribution varies rapidly in a complex landscape and is not separable (as in Monte Carlo simulations of large clusters of atoms). In low-dimensional applications where BMA is applicable, it can guarantee high acceptance rates and quick convergence in bounded state space applications. However, care would have to be taken in discretizing the target CDF when the state space is unbounded, since the BMA method is an independence sampler (the proposed state is independent of the current state). We saw that in the regime where the

smart Monte is also an independence sampler ($h=1$), the value of the remaining free parameter is critical in ensuring convergence of the tails. For a wrong parameter choice, the ratio of probabilities in the tails of the proposal and the target becomes exponentially small the further we move out and therefore steps will be accepted from those states only at an exponentially small rate. For BMA, the appropriate discretization of the CDF could in fact be determined using the variational method presented here to ensure an analogous situation does not arise.

More generally, the variational analysis could be a useful tool in making comparisons between the convergence properties of the latest MCMC algorithms without extensive numerical simulation.

ACKNOWLEDGMENTS

The authors wish to thank Cyrus Umrigar for discussions and the USDA-ARS plant pathogen systems biology group at Cornell University for computing resources. C.R.M. acknowledges support from USDA-ARS Project No. 1907-21000-017-05. We also acknowledge support from NSF Grant No. DMR 0705167.

-
- [1] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, *Markov Chain Monte Carlo in Practice*, 1st ed. (Chapman and Hall, London, 1996).
 - [2] K. Mosegaard and A. Tarantola, *J. Geophys. Res.* **100**, 12431 (1995).
 - [3] Y. F. Atchade, *Methodol. Comput. Appl. Probab.* **8**, 235 (2006).
 - [4] M. Bedard, technical report, University of Toronto (2006), URL <http://probability.ca/jeff/ftpd/ir/mylene2.pdf>.
 - [5] K. S. Brown and J. P. Sethna, *Phys. Rev. E* **68**, 021904 (2003).
 - [6] K. S. Brown, C. C. Hill, G. A. Calero, C. R. Myers, K. H. Lee, J. P. Sethna, and R. A. Cerione, *Phys. Biol.* **1**, 184 (2004).
 - [7] S. L. Frederiksen, K. W. Jacobsen, K. S. Brown, and J. P. Sethna, *Phys. Rev. Lett.* **93**, 165501 (2004).
 - [8] J. J. Waterfall, F. P. Casey, R. N. Gutenkunst, K. S. Brown, C. R. Myers, P. W. Brouwer, V. Elser, and J. P. Sethna, *Phys. Rev. Lett.* **97**, 150601 (2006).
 - [9] R. N. Gutenkunst, J. J. Waterfall, F. P. Casey, K. S. Brown, C. R. Myers, and J. P. Sethna, *PLOS Comput. Biol.* **3**, e189 (2007).
 - [10] S. P. Meyn and R. L. Tweedie, *Ann. Appl. Probab.* **4**, 981 (1994).
 - [11] S. P. Brooks and G. Roberts, *Stat. Comput.* **8**, 319 (1998).
 - [12] A. Gelman and D. B. Rubin, *Stat. Sci.* **7**, 457 (1992).
 - [13] G. O. Roberts and J. S. Rosenthal, *Stat. Sci.* **16**, 351 (2001).
 - [14] G. O. Roberts and W. R. Gilks, *Ann. Appl. Probab.* **7**, 110 (1997).
 - [15] S. F. Jarner and W. K. Yuen, *Adv. Appl. Probab.* **36**, 243 (2004).
 - [16] E. Behrends, *Introduction to Markov Chains with Special Emphasis on Rapid Mixing*, 1st ed. (Vieweg, Wiesbaden, 2000).
 - [17] A. Frigessi, P. di Stefano, C. R. Hwang, and S. J. Sheu, *J. R. Stat. Soc. Ser. B (Methodol.)* **55**, 205 (1993).
 - [18] A. Sinclair and M. Jerrum, *Inf. Comput.* **82**, 93 (1989).
 - [19] P. Diaconis and D. Stroock, *Ann. Appl. Probab.* **1**, 36 (1991).
 - [20] J. S. Rosenthal, *J. Am. Stat. Assoc.* **90**, 558 (1995).
 - [21] G. L. Jones and J. P. Hobert, *Stat. Sci.* **16**, 312 (2001).
 - [22] S. T. Garren and R. L. Smith, *Bernoulli* **6**, 215 (2000).
 - [23] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*, 1st ed. (Springer-Verlag, New York, 1999).
 - [24] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
 - [25] W. K. Hastings, *Biometrika* **57**, 97 (1970).
 - [26] G. F. Lawler and A. D. Sokal, *Trans. Am. Math. Soc.* **309**, 557 (1988).
 - [27] G. O. Roberts, in *Markov Chain Monte Carlo in Practice*, 1st ed., edited by W. R. Gilks, S. Richardson, and D. J. Spiegelhalter (Chapman and Hall, London, 1996), pp. 45–57.
 - [28] J. S. Rosenthal, *Ann. Appl. Probab.* **3**, 819 (1993).
 - [29] P. J. Rossky, J. D. Doll, and H. L. Friedman, *J. Chem. Phys.* **69**, 4628 (1978).
 - [30] J. E. Dennis and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations* (Prentice-Hall, Englewood Cliffs, NJ, 1983).
 - [31] C. Eckart, *Phys. Rev.* **36**, 878 (1930).
 - [32] G. Roberts and R. Tweedie, *Biometrika* **83**, 95 (1996).
 - [33] G. O. Roberts and R. L. Tweedie, *Bernoulli* **2**, 341 (1996).
 - [34] J. Hu, A. Ma, and A. R. Dinner, *J. Comput. Chem.* **27**, 203 (2006).
 - [35] P. V. Kumar, J. S. Raut, and S. J. Warakomski, *J. Chem. Phys.* **105**, 686 (1996).
 - [36] M. Jardat, O. Bernard, P. Turq, and G. R. Kneller, *J. Chem. Phys.* **110**, 7993 (1999).
 - [37] R. Gilks, G. O. Roberts, and S. K. Sahu, *J. Am. Stat. Assoc.*

- 93**, 1045 (1998).
- [38] O. Stramer and R. Tweedie, *Methodol. Comput. Appl. Probab.* **1**, 283 (1999).
- [39] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, *Phys. Lett. B* **195**, 216 (1987).
- [40] D. J. Earl and M. W. Deem, *Phys. Chem. Chem. Phys.* **7**, 3910 (2005).
- [41] A. Bazavov and B. A. Berg, *Phys. Rev. D* **71**, 114506 (2005).
- [42] B. A. Berg and H. X. Zhou, *Phys. Rev. E* **72**, 016712 (2005).