

6

Optimizing and Adapting the Metropolis Algorithm

Jeffrey S. Rosenthal

University of Toronto, Toronto, ON

6.1 Introduction

Many modern scientific questions involve high-dimensional data and complicated statistical models. For example, data on weather consist of huge numbers of measurements across spatial grids, over a period of time. Even in simpler settings, data can be complex: for example, Bartolucci et al. (2007) consider recurrence rates for melanoma (skin cancer) patients after surgery. The probability of recurrence for an individual may depend on physical or biological characteristics of their cancerous lesion, as well as other factors.

Typically, a statistical model for high-dimensional data not only involves a large number of variables but also a correspondingly large number of parameters, which are often represented by a vector θ of some dimension d . To assess the relevance of specific variables for disease recurrence, and to build models that give a risk of recurrence for any given individual, researchers often use Bayesian analysis; see, e.g., Box and Tiao (1973), Gelman et al. (2003) and Carlin and Louis (2009). In this framework, the parameter vector is assumed to follow some probability distribution (of dimension d), and the challenge is to combine a “prior” distribution for θ (typically based on background information about the scientific area) with data that are collected, so as to produce a “posterior” distribution for θ . This probability distribution, call it $\pi(\theta)$, can then be used to answer important scientific questions (e.g., is the size of a cancerous lesion related to the risk of recurrence after surgery?) and to calculate specific probabilities (e.g., this person has a 20% probability of a recurrence within the next five years).

One challenge for Bayesian analysis in situations where the data and parameter vectors are high dimensional is that it is difficult or impossible to compute probabilities based on the posterior distribution. If there is some outcome A of interest (e.g., the outcome that a specific individual’s cancer

will recur), then its probability is given by an integral over A , written

$$\Pi(A) = \int_A \pi(x) dx. \quad (6.1)$$

Also, the expected (average) value of any particular quantity h (e.g., the size of the lesion if the cancer does recur) is also given by an integral (now over the state space \mathcal{X} of all possible vectors), namely

$$E_\pi(h) = \int_{\mathcal{X}} h(x) \pi(x) dx. \quad (6.2)$$

So, to draw conclusions from a Bayesian statistical model, “all” we have to do is compute integrals like (6.1) and (6.2).

Unfortunately, integrals like (6.1) and (6.2) are sometimes very difficult to compute. For example, one commonly-used posterior density (corresponding to a “variance components model” in which individuals are divided into groups) is given by the formula

$$\begin{aligned} \pi(V, W, \mu, \theta_1, \dots, \theta_K) = & C e^{-b_1/V} V^{-a_1-1} \\ & \times e^{-b_2/W} W^{-a_2-1} e^{-(\mu-a_3)^2/2b_3} V^{-K/2} W^{-\frac{1}{2} \sum_{i=1}^K J_i} \\ & \times \exp \left\{ - \sum_{i=1}^K (\theta_i - \mu)^2 / 2V - \sum_{i=1}^K \sum_{j=1}^{J_i} (Y_{ij} - \theta_i)^2 / 2W \right\}. \end{aligned}$$

Here K is the number of different groups (i.e., the number of different values θ_i), V , W and μ are additional unknown parameters of the model, while the a_i, b_i and Y_{ij} are known constants. A typical application might have, say, $K = 19$, so that π is a 22-dimensional function. In such cases, direct computation of integrals like (6.1) and (6.2), using calculus tricks or numerical integration or anything else, appears to be impossible, and alternative methods must be sought (Evans and Swartz, 2000). What can be done to compute estimates of quantities like $E_\pi(h)$ in such complicated cases?

6.2 Monte Carlo Algorithms

One answer to this question is provided by Monte Carlo algorithms. These algorithms, named after the famous casino in Monaco, use randomness to estimate quantities like (6.1) and (6.2). Surprisingly, this turns out to be very helpful!

The most basic (“classical”) form of Monaco requires us to sample from π , i.e., generate a sequence of independent d -dimensional random vectors (variables) X_1, \dots, X_M which each follow the density π , i.e., each have probabilities

given by $\Pr(X_i \in A) = \int_A \pi(x) dx$ for all (measurable) $A \subset \mathcal{X}$. We can then use this random sample to estimate quantities like (6.2), by

$$\mathbb{E}_\pi(h) \approx \frac{1}{M} \sum_{i=1}^M h(X_i). \quad (6.3)$$

If the Monte Carlo sample size M is sufficiently large, then by the Law of Large Numbers the estimate (6.3) will be close to the true expected value (6.2).

Classical Monte Carlo can be very effective, and is widely used for lots of different applications. However, in many cases (such as the above variance components example), it is infeasible to directly sample from π in this sense, i.e., there is no known way to run a computer program which will produce the required sequence of vectors X_i . This state of affairs presented a serious limitation to the use of Monte Carlo algorithms in Bayesian statistical inference problems, until it was largely solved with the introduction of Markov chain Monte Carlo algorithms, as we now discuss.

6.3 Markov Chain Monte Carlo (MCMC)

Markov chain Monte Carlo (MCMC) algorithms were first developed in the statistical physics community by Metropolis et al. (1953), and later expanded by statisticians such as Hastings (1970) before being introduced into the wider statistical community by Gelfand and Smith (1990). They do not use or require a sequence of independent vectors X_i as above. Instead, they define a dependent sequence, more precisely a Markov chain, with each new vector X_{i+1} constructed using the previously-constructed vector X_i .

Suppose we can define a sequence of vectors X_0, X_1, \dots , where each X_{i+1} is constructed using the previous X_i , such that for large n , X_n is approximately a sample (i.e., observation) from π . That is, for large n ,

$$\Pr(X_n \in A) \approx \int_A \pi(x) dx$$

for all $A \subset \mathcal{X}$. This situation is worse than for classical Monte Carlo as in (6.3), since it is just approximate, and furthermore the random vectors X_n are no longer independent but rather are each constructed sequentially using the previous vector. Nevertheless, it is sometimes true that if M is sufficiently large, then we can still approximate $\mathbb{E}_\pi(h)$ by

$$\mathbb{E}_\pi(h) \approx \frac{1}{M - B} \sum_{i=B+1}^M h(X_i). \quad (6.4)$$

That is, we still average many values together, similar to (6.3), even though they are no longer independent. Also, by convention we drop the first B “burn-in” observations since they might be too heavily influenced by our initial values X_0 and hence bias our estimate.

So when is an approximation like (6.4) valid? It turns out that this approximation holds as $M \rightarrow \infty$ provided that the Markov chain is ergodic for π , i.e., that the chain’s probabilities converge to those of π in the sense that, for all $A \subset \mathcal{X}$,

$$\lim_{n \rightarrow \infty} \Pr(X_n \in A) = \int_A \pi(x) dx.$$

This raises the question of how to define a computationally simple Markov chain (X_n) which guarantees that, as $n \rightarrow \infty$,

$$\Pr(X_n \in A) \rightarrow \int_A \pi(x) dx.$$

By standard Markov chain theory, this will hold if the updating rules for (X_n) are irreducible (i.e., the chain can eventually get to anywhere in \mathcal{X}), and aperiodic (i.e., the chain doesn’t have any forced cyclical behavior), and leave the density π stationary (i.e., if it starts distributed according to π , then it will remain distributed according to π at all future times too). How can we ensure that those conditions hold?

6.4 Metropolis Algorithm

An answer to this question was developed in the physics community over fifty years ago (Metropolis et al., 1953). Specifically, given a (possibly important and complicated and high-dimensional) target density π on some state space $\mathcal{X} \subset \mathbb{R}^d$ (with $\pi(x) = 0$ for $x \in \mathbb{R}^d \setminus \mathcal{X}$), the original Metropolis algorithm proceeds as follows.

First, choose a symmetric d -dimensional increment distribution Q ; the most common choice is $Q = \mathcal{N}(0, \Sigma)$ for some fixed covariance matrix Σ like $\Sigma = cI_d$. Here $\mathcal{N}(0, \Sigma)$ is a d -dimensional Normal (Gaussian) distribution, and I_d is the $d \times d$ identity matrix.

Also, choose some initial vector X_0 . Then, iteratively for $n = 1, 2, \dots$, compute X_n from X_{n-1} as follows:

1. Let $Y_n = X_{n-1} + Z_n$, where $(Z_n) \sim Q$ are independent and identically distributed (“proposal”).
2. Let $\alpha = \pi(Y_n)/\pi(X_{n-1})$ if this ratio is smaller than 1; otherwise, set $\alpha = 1$ (“acceptance probability”).

- 3a. With probability α , let $X_n = Y_n$ (“accept”).
- 3b. Otherwise, with probability $1 - \alpha$, let $X_n = X_{n-1}$ (“reject”).

Steps 2, 3a, and 3b can all be accomplished by drawing a random variable U_n uniformly between 0 and 1, and then setting

$$X_n = \begin{cases} Y_n & \text{if } U_n \leq \pi(Y_n)/\pi(X_{n-1}), \\ X_{n-1} & \text{if } U_n > \pi(Y_n)/\pi(X_{n-1}). \end{cases}$$

Intuitively, the above acceptance probabilities α are useful because they encourage the algorithm to accept more moves toward larger values of π . More precisely, the formula for α turns out to be exactly the right one to ensure that the Markov chain (X_n) leaves the density π stationary, a key property for convergence (as discussed at the end of the previous section). Furthermore, the irreducibility property will almost always hold, e.g., it is guaranteed if Q has an everywhere-positive density like $\mathcal{N}(0, \Sigma)$. And, the aperiodicity property is essentially never a problem since the algorithm will eventually reject and thus avoid cyclic behavior. So, this simple algorithm has all the right properties to guarantee that, as $n \rightarrow \infty$, $\Pr(X_n \in A) \rightarrow \pi(A)$. It follows that we can use this algorithm to estimate $E_\pi(h)$ as in (6.4). Good!

The only problem is that sometimes the Metropolis algorithm will be too inefficient, i.e., it will take far too long (i.e., require too many iterations) to provide a decent approximation to π , which is a very important consideration; see, e.g., Rosenthal (1995). In some cases, even running the algorithm for months on the world’s fastest computers would not provide a remotely reasonable approximation to π . Overcoming such problems has often necessitated new and more complicated MCMC algorithms; see, e.g., Bélisle et al. (1993), Neal (2003), Jain and Neal (2004), and Hamze and de Freitas (2012). In a different direction, detecting convergence of MCMC to π is so challenging that some authors have developed perfect sampling algorithms which guarantee complete convergence at the expense of a more complicated algorithm; see, e.g., Propp and Wilson (1996), Murdoch and Green (1998) or Fill et al. (2000). However, such perfect sampling algorithms are often infeasible to run, so we do not discuss them further here.

All of this raises the question of how to improve or optimize the speed of convergence of the Metropolis algorithm, for example by modifying the increment distribution Q , as we discuss next.

6.5 Goldilocks Principle

To illustrate the Metropolis optimization issues, consider the very simple case where $\pi = \mathcal{N}(0, 1)$, i.e., where the target density is just the standard Normal

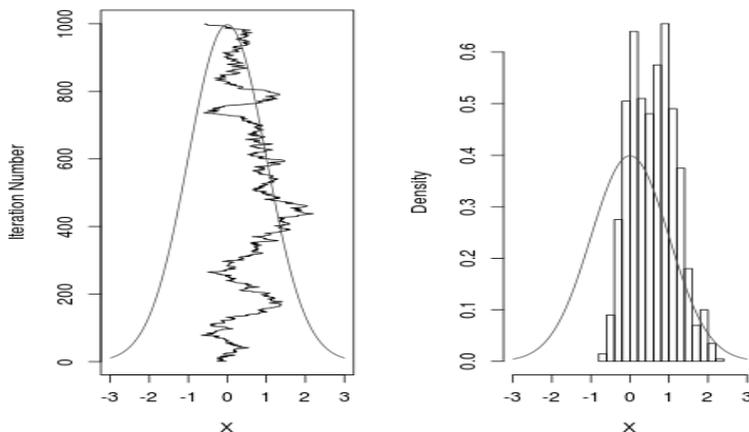


FIGURE 6.1: The trace plot (left) and histogram (right), along with the mound-shaped target density for a one-dimensional Metropolis algorithm with too small a proposal scaling σ , showing slow mixing and poor convergence.

distribution. (Of course we wouldn't actually need to use MCMC in such a simple case.) Assume that the proposal distribution is given by $Q = \mathcal{N}(0, \sigma^2)$. Our question of interest is, how should we choose σ ?

As a first try, let's choose a small value of σ , say $\sigma = .1$, and run the Metropolis algorithm for 1000 iterations with that σ . The corresponding trace plot, graphing the values of the Markov chain (horizontal axis) at each iteration n (vertical axis), is shown in Figure 6.1 (left panel). Looking at this trace plot, we can see that the chain moves very slowly. It starts at the value zero, and takes many hundreds of iterations before it moves appreciably away from zero. In particular, it does not do a very good job of exploring the superposed mound-shaped target density. This is also illustrated by the histogram of those 1000 iterations, in Figure 6.1 (right panel), which does not match up very well with the target density.

As a second try, let's choose a large value of σ , say $\sigma = 25$, and again run the algorithm for 1000 iterations. The trace plot in this case is shown in Figure 6.2 (left panel). In this case, when the chain finally accepts a move, it jumps quite far, which is good. However, since it proposes such large moves, it hardly ever accepts them. (Indeed, it accepted just 5.4% of the proposed moves, compared to 97.7% when $\sigma = .1$.) So, this chain doesn't perform very well either, as illustrated by the histogram in Figure 6.2 (right panel), which again does not match up very well with the target density.

As a third try, let's choose a compromise value of σ , say $\sigma = 2.38$, and again run the algorithm for 1000 iterations. In this case, the chain performs very well. It accepts a medium fraction of its proposals (44.5%), and moves reasonably

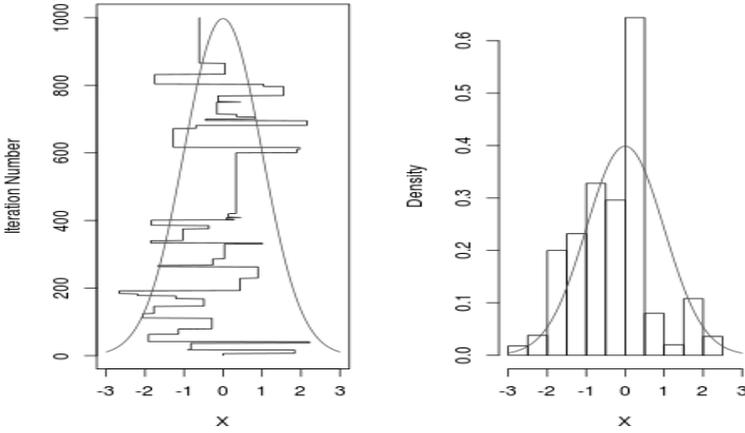


FIGURE 6.2: The trace plot (left) and histogram (right), along with the mound-shaped target density for a one-dimensional Metropolis algorithm with too large a proposal scaling σ , again showing slow mixing and poor convergence.

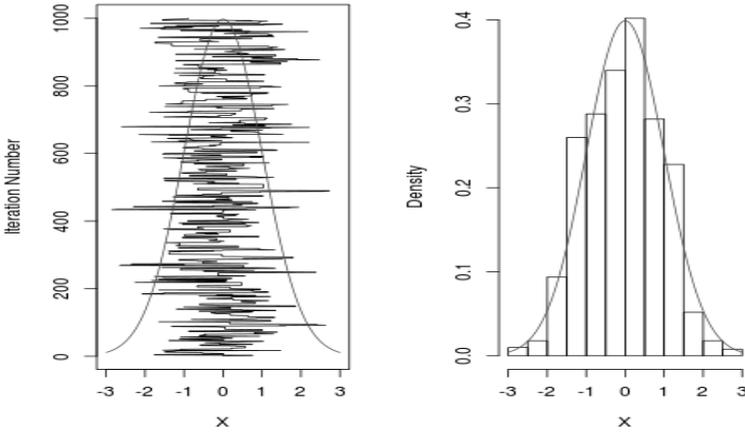


FIGURE 6.3: The trace plot (left) and histogram (right), along with the mound-shaped target density for a one-dimensional Metropolis algorithm with a good choice of proposal scaling σ , showing much better mixing and convergence properties.

far when it does accept. It thus explores the target density efficiently and well, as illustrated by the trace plot in Figure 6.3 (left panel). Furthermore it now provides fairly good samples from the target distribution, as illustrated by the histogram in Figure 6.3 (right panel).

We learn from this that it is best to choose values of the proposal increment scaling σ which are between the two extremes, i.e., not too small and not too big, but rather “just right” (as the little girl Goldilocks says in the classic children’s fairy tale *The Three Bears*; <http://w8r.com/the-colorful-story-book/the-three-bears>). Correspondingly, the acceptance rate (i.e., the percentage of proposed moves which are accepted) should be far from 0% but also far from 100%.

6.6 Optimal Scaling

The above intuition was made more precise in a pioneering paper by Roberts et al. (1997). They decided to consider a Metropolis algorithm (X_n) in dimension d , with increment distribution

$$Q = \mathcal{N}\left(0, \frac{\ell^2}{d} I_d\right)$$

for some fixed scaling constant $\ell > 0$, and take the limit as $d \rightarrow \infty$. The beauty of their approach was that they could compute the speed of the algorithm in this limit, as an explicit (but messy and uncomputable) function $h(\ell)$ of the scaling constant ℓ . They then argued that the best choice of ℓ is the one which maximizes the limiting speed $h(\ell)$. This provided a clear standard for how to optimize the Metropolis algorithm.

The story gets even better. Roberts and his collaborators also considered the asymptotic acceptance rate, i.e., the fraction of Metropolis proposals that would be accepted in the limit as $d \rightarrow \infty$, and they computed an explicit function $A(\ell)$ for this as well. They then showed that the limiting speed $h(\ell)$ has a simple relation to the asymptotic acceptance rate $A(\ell)$. This in turn allowed them to compute that if ℓ_{opt} is the value of ℓ which maximizes the speed $h(\ell)$, then $A(\ell_{\text{opt}}) \approx .234 = 23.4\%$, a specific number that does not depend on any unknown quantities about π or anything else. This means that, at least under their (strong) assumptions, the optimal acceptance probability is 23.4%, which leads to the fastest limiting speed regardless of the target density π .

This provides a clear, simple rule for tuning Metropolis algorithms: adjust the proposal scaling ℓ so that the acceptance rate is roughly 23.4%. This rule appears to be quite robust, i.e., .234 is often a nearly optimal acceptance rate even if the theorem’s formal assumptions are not satisfied. It has been

implemented in numerous applied papers and software, including the hugely popular WinBUGS computer package (Lunn et al., 2000).

A number of authors have attempted to weaken and generalize the original strong assumptions; see e.g., Bédard (2007, 2008), Bédard and Rosenthal (2008), Beskos et al. (2009), and Sherlock and Roberts (2009). Corresponding results have been developed for Langevin MCMC algorithms (Roberts and Rosenthal, 1998), and for simulated tempering algorithms (Atchadé et al., 2011; Roberts and Rosenthal, 2013). It is also known that any acceptance rate between about 15% and 50% is still reasonably efficient (though 23.4% is best); see, e.g., Figure 3 in Roberts and Rosenthal (2001). Over all, this school of research has been very influential in guiding both applied usage and theoretical investigations for MCMC.

6.7 Proposal Shape

Despite the .234 rule's great success, an algorithm's acceptance rate is just a single scalar quantity which does not completely control the algorithm's efficiency. To illustrate this, consider the 20-dimensional target density $\pi = \mathcal{N}(0, \Sigma_*)$, where Σ_* is a 20-dimensional covariance matrix generated randomly as $\Sigma_* = M^\top M$ with M being a 20×20 matrix consisting of independent and identically distributed $\mathcal{N}(0, 1)$ entries, which shall remain fixed throughout the remainder of this chapter, and M^\top is the transpose of M . We shall try running Metropolis algorithms on this density.

Based on the discussion in Section 6.6, we use an increment distribution of the form $Q = \mathcal{N}(0, \sigma^2 I_{20})$, where we try to adjust σ so that the resulting acceptance rate is approximately .234. After some experimenting, we take $\sigma = .5$, leading to an acceptance rate of .228 (close enough). Figure 6.4 shows the trace plot (left panel) and histogram (right panel) of the first coordinate of this run. The mixing and convergence aren't so bad, given the relatively large dimension, but they aren't great either. Indeed, the performance appears to be similar to our earlier one-dimensional example's first attempt with $\sigma = .1$: it only slowly explores the support of π .

Next, we instead try the increment distribution

$$Q = \mathcal{N}\left(0, \frac{(2.38)^2}{20} \Sigma_*\right),$$

which we shall justify later. This leads to an acceptance rate of .252 (again, close enough). Figure 6.5 shows the trace plot (left panel) and histogram (right panel) of the first coordinate of this run. Direct inspection (as well as more precise measurements like squared jumping distance and functional variance, not discussed here) indicate that this choice of Q , despite leading to a very similar acceptance rate (approximately .234), actually performs much

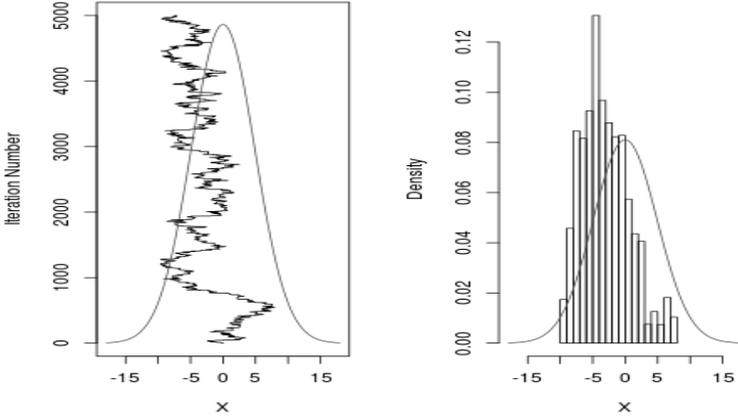


FIGURE 6.4: The trace plot (left) and histogram (right), along with the mound-shaped target density of the first coordinate of a 20-dimensional Metropolis algorithm with proposal scaling proportional to the identity matrix, showing slow mixing and poor convergence.

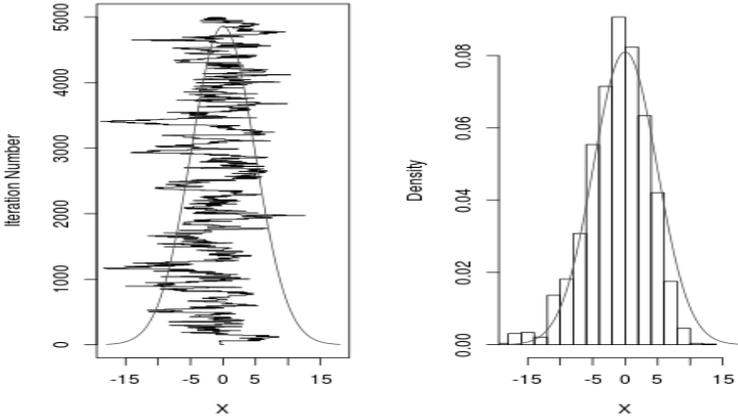


FIGURE 6.5: The trace plot (left) and histogram (right), along with the mound-shaped target density of the first coordinate of a 20-dimensional Metropolis algorithm with proposal scaling proportional to the target covariance matrix Σ_* , showing much faster mixing and much better convergence properties.

better, exploring the support of π and converging to the correct probabilities much more efficiently. This confirms that there is more to the story than just acceptance rate, and indeed that the shape of the proposal distribution (determined by the proposal covariance matrix, in this case $(2.38)^2 \Sigma_*/20$) is also very important.

This concept was formalized by Roberts and Rosenthal (2001). They proved that under strong assumptions (similar to before), the optimal Gaussian proposal distribution is given (to three significant figures) by

$$Q = \mathcal{N}\left(0, \frac{(2.38)^2}{d} \Sigma_\pi\right), \quad (6.5)$$

where Σ_π is the $d \times d$ covariance matrix of the target density π , rather than $Q = \mathcal{N}(0, \sigma^2 I_d)$ for some σ^2 . Furthermore, with this choice, the asymptotic acceptance rate will again be approximately .234. And, as before, this result appears to be robust in the sense of being nearly optimal even when the strong assumptions do not hold.

6.8 Adaptive MCMC

The optimization result of the previous section requires us to know and use the target covariance matrix Σ_π . Now, in most realistic situations, Σ_π would not be known in advance, and indeed would be at least as difficult to estimate as the quantity $E_\pi(h)$ of ultimate interest. How can we optimize the Metropolis algorithm in such situations?

In a pioneering paper, Haario et al. (2001) proposed to optimize the algorithm adaptively. That is, even if we don't know the optimal algorithm at the beginning, the computer can learn it during the run, and update the algorithm "on the fly" as it proceeds.

In its simplest form, this algorithm finds an approximately optimal increment distribution by replacing the unknown target $d \times d$ covariance matrix Σ_π by the sample $d \times d$ covariance $\Sigma_n = \text{cov}(X_0, \dots, X_{n-1})$ of the vectors visited so far during the run. If those X_i are indeed good approximate samples from π , then Σ_n will be a good approximation to Σ_π , and hence

$$Q = \mathcal{N}\left(0, \frac{(2.38)^2}{d} \Sigma_n\right)$$

can be used (after an initial phase, e.g., for $n \geq 40$ only) as a good approximation to the optimal proposal (6.5). If not, then the algorithm will not work well initially, but it will hopefully improve as it goes.

Now, Σ_n is easily computed, so this algorithm is quite feasible to run in practice. Running it on the above 20-dimensional target density $\pi = \mathcal{N}(0, \Sigma_*)$,

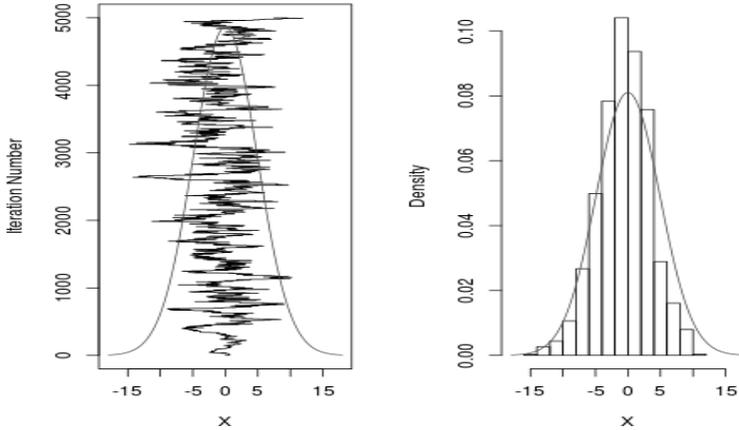


FIGURE 6.6: The trace plot (left) and histogram (right), along with the mound-shaped target density of the first coordinate of a 20-dimensional adaptive Metropolis algorithm with proposal scaling computed from previous iterations, showing fairly rapid mixing and good convergence properties, especially in later iterations (top of left plot).

the resulting trace plot and histogram of the first coordinate are shown in Figure 6.6. Direct inspection, and precise measurements, both indicate that this algorithm performs poorly during the initial phase (bottom of left plot), but then performs very well later on (top of left plot), exploring the support of π nearly as efficiently as the optimal algorithm presented above, even though it does not require any prior knowledge about Σ_π .

Adaptive MCMC algorithms have recently been used in a number of different statistical applications, and often lead to significant speed-ups, even in hundreds of dimensions; see, e.g., Roberts and Rosenthal (2009), Craiu et al. (2009), Giordani and Kohn (2010) and Richardson et al. (2011). On the other hand, adaptive MCMC algorithms use previous iterations to determine their future transitions, so they violate the Markov property which provides the justification for conventional MCMC. This raises the question of whether adaptive MCMC algorithms are valid, i.e., whether they converge (asymptotically, at least) to the target density π .

The answer to this question is “no” in general. However it is “yes” under various conditions; see, e.g., Andrieu and Moulines (2006), Haario et al. (2001), Atchadé and Rosenthal (2005), Roberts and Rosenthal (2007), Saksman and Vihola (2010), Fort et al. (2012) and Łatuszyński et al. (2013). In particular, Roberts and Rosenthal (2007) show that adaptive algorithms will still converge to the target density π provided they satisfy two fairly mild conditions: “Diminishing Adaptation” (the algorithm adapts by less and less

as time goes on), and “Containment” (the chain never gets too lost, in the sense that it remains bounded in probability). Conditions such as these have been used to formally justify adaptive algorithms in many examples; see, e.g., Roberts and Rosenthal (2009) and Richardson et al. (2011). Adaptive MCMC appears to hold great promise for improving statistical computation in many application areas in the years ahead.

6.9 Summary

We summarize the points made in this chapter as follows.

- a) The Metropolis algorithm is very important in many applications.
- b) This algorithm sometimes runs so slowly that computations are infeasible; thus, optimization of the algorithm can be crucial.
- c) The simplest optimization result is the Goldilocks Principle that the acceptance rate should be far from 0, but also far from 1.
- d) A more detailed result says that the optimal acceptance rate is .234 under certain strong assumptions (though the conclusion appears to be fairly robust even when the assumptions are not satisfied).
- e) Another theorem says that the optimal increment distribution is $\mathcal{N}(0, (2.38)^2 \Sigma_\pi / d)$, again under strong assumptions (though again with a fairly robust conclusion).
- f) When certain key information is unknown (e.g., Σ_π), it may still be possible to adapt toward the optimal algorithm. Such adaptation is not valid in general, but is valid under various conditions such as “Diminishing Adaptation” and “Containment.” It can eventually lead to tremendous speed-ups, even in very high dimensions.
- g) In short, to improve statistical computation in important applied areas, optimization and adaption may well be worth the trouble!

Acknowledgments

This work was financed in part by the Natural Sciences and Engineering Research Council of Canada. The author thanks the editors and reviewers for many detailed comments.

About the Author

Jeffrey S. Rosenthal is a professor of statistics at the University of Toronto. He received his BSc from the University of Toronto and his PhD in Mathematics from Harvard University. His research interests include Markov Chain Monte Carlo algorithms and interdisciplinary applications. He has also published two textbooks, and a general-interest book about probability. He has received the 2006 CRM–SSC Prize, the 2007 COPSS Award, and the 2013 SSC Gold Medal for research. He is a fellow of the Institute of Mathematical Statistics and was elected a fellow of the Royal Society of Canada in 2012.

Bibliography

- Andrieu, C. and Moulines, É. (2006). On the ergodicity properties of some adaptive Markov Chain Monte Carlo algorithms. *The Annals of Applied Probability*, 16:1462–1505.
- Atchadé, Y. F., Roberts, G. O., and Rosenthal, J. S. (2011). Towards optimal scaling of Metropolis-Coupled Markov Chain Monte Carlo. *Statistics and Computing*, 21:555–568.
- Atchadé, Y. F. and Rosenthal, J. S. (2005). On adaptive Markov Chain Monte Carlo algorithms. *Bernoulli*, 11:815–828.
- Bartolucci, A. A., Singh, K. P., and Bae, S. (2007). Analyzing clinical trial data via the Bayesian multiple logistic random effects model. In *MODSIM 2007 conference proceedings, Modelling and Simulation Society of Australia and New Zealand*, pp. 2861–2866.
- Bédard, M. (2007). Weak convergence of Metropolis algorithms for non-i.i.d. target distributions. *The Annals of Applied Probability*, 17:1222–1244.
- Bédard, M. (2008). Optimal acceptance rates for Metropolis algorithms: Moving beyond 0.234. *Stochastic Processes and their Applications*, 118:2198–2222.
- Bédard, M. and Rosenthal, J. S. (2008). Optimal scaling of Metropolis algorithms: Heading toward general target distributions. *The Canadian Journal of Statistics*, 36:483–503.
- Bélisle, C., Romeijn, H. E., and Smith, R. L. (1993). Hit-and-run algorithms for generating multivariate distributions. *Mathematics of Operations Research*, 18:255–266.
- Beskos, A., Roberts, G., and Stuart, A. (2009). Optimal scalings for local Metropolis–Hastings chains on nonproduct targets in high dimensions. *The Annals of Applied Probability*, 19:863–898.

- Box, G. E. P. and Tiao, G. C. (1973). *Bayesian Inference in Statistical Analysis*. Addison-Wesley, Reading, MA.
- Carlin, B. P. and Louis, T. A. (2009). *Bayesian Methods for Data Analysis*, Third Edition. CRC Press, Boca Raton, FL.
- Craiu, R. V., Rosenthal, J. S., and Yang, C. (2009). Learn from thy neighbor: Parallel-chain and regional adaptive MCMC. *Journal of the American Statistical Association*, 104:1454–1466.
- Evans, M. J. and Swartz, T. B. (2000). *Approximating Integrals via Monte Carlo and Deterministic Methods*. Oxford University Press, Oxford.
- Fill, J. A., Machida, M., Murdoch, D. J., and Rosenthal, J. S. (2000). Extension of Fill’s perfect rejection sampling algorithm to general chains. In *Proceedings of the Ninth International Conference “Random Structures and Algorithms” (Poznan, 1999)*, volume 17, pp. 290–316.
- Fort, G., Moulines, É., and Priouret, P. (2012). Convergence of adaptive and interacting Markov Chain Monte Carlo algorithms. *The Annals of Statistics*, 39:3262–3289.
- Gelfand, A. E. and Smith, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85:398–409.
- Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2003). *Bayesian Data Analysis*, Second Edition. Chapman & Hall, London.
- Giordani, P. and Kohn, R. (2010). Adaptive independent Metropolis–Hastings by fast estimation of mixtures of normals. *Journal of Computational and Graphical Statistics*, 19:243–259.
- Haario, H., Saksman, E., and Tamminen, J. (2001). An adaptive Metropolis algorithm. *Bernoulli*, 7:223–242.
- Hamze, F. and de Freitas, N. (2012). Intracluster moves for constrained discrete-space MCMC. *CoRR*, abs/1203.3484.
- Hastings, W. (1970). Monte Carlo sampling methods using Markov Chains and their applications. *Biometrika*, 57:97–109.
- Jain, S. and Neal, R. M. (2004). A split-merge Markov Chain Monte Carlo procedure for the Dirichlet process mixture model. *Journal of Computational and Graphical Statistics*, 13:158–182.
- Latuszyński, K., Roberts, G. O., and Rosenthal, J. S. (2013). Adaptive Gibbs samplers and related MCMC methods. *The Annals of Applied Probability*, 23:66–98.
- Lunn, D., Thomas, A., Best, N., and Spiegelhalter, D. (2000). WinBUGS: A Bayesian modelling framework: Concepts, structure, and extensibility. *Statistics and Computing*, 10:325–337.

- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equations of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21:1087–1092.
- Murdoch, D. J. and Green, P. J. (1998). Exact sampling from a continuous state space. *Scandinavian Journal of Statistics*, 25:483–502.
- Neal, R. M. (2003). Slice sampling (with discussion). *The Annals of Statistics*, 31:705–767.
- Propp, J. G. and Wilson, D. B. (1996). Exact sampling with coupled Markov Chains and applications to statistical mechanics. In *Proceedings of the Seventh International Conference on Random Structures and Algorithms (Atlanta, GA, 1995)*, 9:223–252.
- Richardson, S., Bottolo, L., and Rosenthal, J. S. (2011). Bayesian models for sparse regression analysis of high dimensional data. In *Bayesian Statistics 9*, pp. 539–560. Oxford University Press.
- Roberts, G. O., Gelman, A., and Gilks, W. R. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability*, 7:110–120.
- Roberts, G. O. and Rosenthal, J. S. (1998). Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society, Series B*, 60:255–268.
- Roberts, G. O. and Rosenthal, J. S. (2001). Optimal scaling for various Metropolis–Hastings algorithms. *Statistical Science*, 16:351–367.
- Roberts, G. O. and Rosenthal, J. S. (2007). Coupling and ergodicity of adaptive MCMC. *Journal of Applied Probability*, 44:458–475.
- Roberts, G. O. and Rosenthal, J. S. (2009). Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics*, 18:349–367.
- Roberts, G. O. and Rosenthal, J. S. (2013). Minimising MCMC variance via diffusion limits, with an application to simulated tempering. *The Annals of Applied Probability*, in press.
- Rosenthal, J. S. (1995). Minorization conditions and convergence rates for Markov Chain Monte Carlo. *Journal of the American Statistical Association*, 90:558–566. [Correction: p. 1136].
- Saksman, E. and Vihola, M. (2010). On the ergodicity of the adaptive Metropolis algorithm on unbounded domains. *The Annals of Applied Probability*, 20:2178–2203.
- Sherlock, C. and Roberts, G. O. (2009). Optimal scaling of the random walk Metropolis on elliptically symmetric unimodal targets. *Bernoulli*, 15:774–798.