Exploring the generalizability of the optimal 0.234 acceptance rate in random-walk Metropolis and parallel tempering algorithms

Aidan Li^{1*}, Liyan Wang¹, Tianye Dou¹, Jeffrey S. Rosenthal¹

^{1*}Department of Statistical Sciences, University of Toronto.

*Corresponding author. E-mail: aidan.mr.li@gmail.com; Contributing authors: liyan.wang@mail.utoronto.ca; tianye.dou@mail.utoronto.ca; jeff@math.toronto.edu;

Abstract

For random-walk Metropolis (RWM) and parallel tempering (PT) algorithms, an asymptotic acceptance rate of around 0.234 is known to be optimal in certain high-dimensional limits. However, its practical relevance is uncertain due to restrictive derivation conditions. We synthesise previous theoretical advances in extending the 0.234 acceptance rate to more general settings, and demonstrate its applicability with a comprehensive empirical simulation study on examples examining how acceptance rates affect Expected Squared Jumping Distance (ESJD). Our experiments show the optimality of the 0.234 acceptance rate for RWM is surprisingly robust even in lower dimensions across various proposal, multimodal distributions that may not have an i.i.d. product density, and curved Rosenbrock target distributions with nonlinear correlation structure. Parallel tempering experiments also show that the idealized 0.234 spacing of inverse temperatures may be approximately optimal for low dimensions and non i.i.d. product target densities, and that constructing an inverse temperature ladder with spacings given by a swap acceptance of 0.234 is a viable strategy.

Keywords: Markov chain Monte Carlo, Metropolis algorithm, optimal scaling, acceptance rate, parallel tempering

1 Introduction

Markov chain Monte Carlo (MCMC) methods, such as the random-walk Metropolis algorithm, are used to draw samples from complex high-dimensional target probability distributions. They enjoy strong theoretical asymptotic guarantees of accuracy, converging to the target distribution in stationarity (Brooks et al., 2011; Robert and Casella, 2004). However, they may also be inefficient and prohibitively slow to provide a good approximation of the target distribution they are sampling (Rosenthal, 1995). Therefore, the analysis of their running time is an important practical issue to consider.

Roberts et al. (1997) proved an important result for independently and identically distributed (i.i.d.) product target densities: in the state space high-dimensional limit $d \to \infty$, the asymptotic acceptance rate of a Metropolis algorithm is central to its efficiency. Under their assumptions, an asymptotic acceptance rate of approximately 0.234 maximises the efficiency of a random-walk Metropolis (RWM) algorithm with a Gaussian increment proposal distribution starting at stationarity. However, their assumptions on the target and proposals are very restrictive. Many practical scenarios often involve finite and relatively small dimensions. Additionally, many probability distributions do not have densities of the i.i.d. product form, and for those that do, sampling a one-dimensional target is much simpler than using RWM. A burgeoning active research area is dedicated to showing that this figure still applies beyond these assumptions, and much of this focuses on the i.i.d. product assumption.

Many papers have shown that the optimal average acceptance rate of 0.234 for RWM holds greater generality than the assumptions used in Roberts et al. (1997). For more general theoretical results, Sherlock (2013) replaces the restrictive i.i.d. product target and Gaussian proposal setting of Roberts et al. (1997) with a set of "shell", relative variability, and eccentricity conditions on the target, eccentricity condition on the jump distance matrix, and shell condition on the proposal. This framework does not require a product form and is based on the geometric properties of the (log) target distribution in high dimensions. Yang et al. (2020) makes the 0.234 result applicable to targets with sparse dependency structures, such as those arising from Bayesian graphical models, using a set of sufficient conditions based on local dependencies and bounded asymptotic behaviour of derivatives. We expand on these results in Section 2. The 0.234 acceptance rate in optimal scaling has also been shown to apply to the spacing and swapping between inverse temperatures in the parallel tempering MCMC method (Atchadé et al., 2011; Roberts and Rosenthal, 2014). That being said, there are some cases in which 0.234 is not optimal, such as in the case of discontinuous targets (Neal et al., 2012) where the optimal acceptance rate is approximately 0.1353. Additionally, these results of optimality are still theoretical ideals, with required assumptions that could be further relaxed, and in the limit of $d \to \infty$ where the Markov chain converges to a diffusion process.

Other relaxations of the 0.234 assumptions include: non-Gaussian proposals (Neal and Roberts, 2011), discrete hypercube target distributions (Roberts, 1998), Gibbs random fields (Breyer and Roberts, 2000), perturbations from product targets using pre-conditioned proposals (Beskos et al., 2009), independent target components with inhomogeneous scaling (Bédard, 2007; Bédard and Rosenthal, 2008; Roberts and

Rosenthal, 2001), partial updates where not all components are updated at once (Neal and Roberts, 2006), elliptical symmetric unimodal target densities (Sherlock and Roberts, 2009), and infinite-dimensional target distributions with non-trivial dependence structures (Mattingly et al., 2012).

These results are of great interest to practitioners who would like to understand how to tune their MCMC algorithms to maximise their efficiency. Previous simulation studies have been carried out on more general target distributions like products of independent, non-identical components (Bédard, 2008; Roberts and Rosenthal, 2001), and for the Markov modulated Poisson process (Sherlock et al., 2010), but there is still a lack of thorough experimentation and guidance for how handy the 0.234 rule is on realistic target distributions that may have lower or higher dimensionality and/or multimodality. Our research examines how necessary these assumptions are for the theoretical result to remain relevant: where can we relax some assumptions in lower dimensions and still have the optimal acceptance rate of approximately 0.234, and where can we not do this? And in the end, how applicable is the 0.234 figure for the practitioner, who may want to use MCMC for complicated target distributions beyond the necessary assumptions in finite dimensions? We thoroughly dissect various aspects of the Metropolis algorithm and experiment with them, and in doing so, we show empirically where the theoretical ideal value can still align with more realistic scenarios.

We begin our paper with a description of the RWM and its optimal scaling framework, as well as formally stating sufficient theoretical conditions for the 0.234 rule that are most relevant to our work in Section 2. In Section 3, we describe RWM experiments in low dimensions that show where an acceptance rate of 0.234 can be optimal in the sense that it maximises the Expected Squared Jumping Distance (ESJD) for various proposal and target densities which are not necessarily all i.i.d. product forms. Next, we describe the parallel tempering method and its own optimal scaling framework in Section 4, and show using experiments in Section 5 that the 0.234 swap acceptance rate in parallel tempering may also be optimal in lower dimensions for the multivariate Gaussian target density and a Gaussian mixture target density which is not of an i.i.d. product form. Lastly, we discuss the implications of our findings in Section 6.

2 Random-Walk Metropolis Background

We first introduce the Metropolis algorithm and the Optimal Scaling framework. We then discuss measures to evaluate the efficiency of MCMC algorithms, such as the expected squared jumping distance, and explain how the asymptotic acceptance rate of a MCMC algorithm is tied to the algorithm's efficiency. Finally, we discuss some of the conditions for the 0.234 acceptance rate to be optimal that are most relevant to our work.

2.1 Metropolis Algorithm

A Metropolis algorithm (Metropolis et al., 1953) constructs a Markov chain $\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \ldots$ If the chain is constructed to be irreducible and aperiodic, it is guaranteed to have the target distribution π as its unique stationary distribution. This means

that as the number of steps t becomes large, the distribution of $\mathbf{x}^{(t)}$ converges to π . Therefore, after an initial burn-in period, the subsequent states of this single, long chain can be used as a sample from the target distribution π . When taking each step in the Markov chain, we generate a new state $\mathbf{y} = \mathbf{x}^{(t)} + \epsilon$ where ϵ is sampled from an *increment* "proposal" density Q having $Q(\epsilon) = Q(-\epsilon)$ for all ϵ in its domain and accept this new state with probability $\alpha(\mathbf{x}, \mathbf{y})$ given by $\alpha = \min\left\{1, \frac{\pi(\mathbf{y})}{\pi(\mathbf{x})}\right\}$, where, by a common abuse of notation, we let $\pi(\cdot)$ also denote the probability density function of the target distribution π . If the proposed value is accepted, we set the next state $\mathbf{x}^{(t+1)} = \mathbf{y}$. Otherwise, we set $\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)}$.

2.2 Optimal Scaling Framework

Optimal Scaling (Roberts et al., 1997; Roberts and Rosenthal, 2001) is one of the most successful frameworks for performing asymptotic analysis of high-dimensional MCMC methods, and provides mathematically-grounded guidance on how to best optimise MCMC performance by tuning the "scaling" parameter(s) of the proposal distribution. Typically, a multivariate Gaussian distribution with mean 0 is used as the increment proposal distribution. The scaling of the Gaussian proposal is then given by a scaling factor σ , and this is often used as a variance σ^2 that is applied to the $d \times d$ identity matrix I_d . So, $\epsilon \sim \mathcal{N}(0, \sigma^2 I_d)$. If the proposal's variance is too small, many proposed steps are accepted by the algorithm, but each step does not explore the state space much. Vice versa, if the proposal's variance is too large, too few proposed steps are accepted, leading to slow exploration of the state space since the algorithm stays at a state for too long. Therefore, the proposal's variance is crucial to the performance of the algorithm; the optimal scaling should balance the exploration of new areas in the state space with exploiting high-density areas of the target distribution. Different Metropolis algorithms have shown to have different optimal choices of the proposal distribution variance, σ^2 .

We next discuss how to evaluate the efficiency of an MCMC algorithm for determining optimal scaling. A very popular measure of efficiency both theoretically (e.g. Sherlock (2013); Roberts and Rosenthal (2014); Yang et al. (2020)) and in practice (e.g. Pasarica and Gelman (2010); Sherlock et al. (2015)) is the Euclidean *Expected Squared Jumping Distance* (ESJD) metric which we use for our experiments. The expected squared jumping distance measures how far, in expectation, the MCMC chain moves in a single iteration. For the standard random-walk Metropolis algorithm, we define this as

$$\mathbb{E}\left[\left\|\mathbf{x}^{(t+1)} - \mathbf{x}^{(t)}\right\|^{2}\right] \approx \frac{1}{n-1} \sum_{i=1}^{n-1} \left\|\mathbf{x}^{(t+1)} - \mathbf{x}^{(t)}\right\|^{2}$$
(1)

where n is the total number of iterations of the algorithm. Maximising the ESJD aligns with minimizing the first-order auto-correlation of the Markov chain and subsequently maximises efficiency if the higher-order auto-correlations are monotonically increasing relative to the first-order auto-correlation (Pasarica and Gelman, 2010; Yang et al., 2020). It is worth mentioning that ESJD primarily measures local convergence performance. Because heavy-tailed targets disrupt the monotonic relationship between first-order and higher-order autocorrelations, achieving global convergence

for heavy-tailed targets requires maximising an adapted version of ESJD (Kamatani, 2020). However, we do not consider heavy-tailed targets in this work, so the working definition of ESJD is sufficient.

There are other notions of efficiency of a Markov chain, but in the high-dimensional limit $d \to \infty$, if the moment conditions in Assumption 2.1 holds, the chain converges to a diffusion process, and all efficiency measures are effectively equivalent (Roberts et al., 1997; Roberts and Rosenthal, 2001). A key result of these referenced papers is that, given a Metropolis algorithm with a Gaussian increment proposal distribution $Q = \mathcal{N}(0, \frac{\ell^2}{d}I_d)$ where $\ell > 0$ is a fixed scaling constant and I_d is the identity matrix, maximising the diffusion's speed measure $h(\ell)$, which is a function of the scaling constant ℓ , yields the most efficient asymptotic diffusion. Furthermore, the speed measure has a clear relation to a much simpler quantity to tune for: the asymptotic acceptance rate of the proposed new states (moves) of the algorithm, defined as

$$a = \lim_{n \to \infty} \frac{\# \text{ accepted moves}}{n}$$

Both the speed measure $h(\ell)$ and asymptotic acceptance rate $a(\ell)$ are functions of the scaling constant ℓ . Given the restrictive conditions mentioned in Section 2.3, the scaling constant ℓ that maximises the speed measure $h(\ell)$ corresponds to an asymptotic acceptance rate $a(\ell)$ of approximately 0.234. Hence, an asymptotic acceptance rate of approximately 0.234 should necessarily optimise a measure of efficiency such as the ESJD of the algorithm.

2.3 Theoretical Assumptions for Optimality of a 0.234 Acceptance Rate

Certain work has shown that under restrictive assumptions, a 0.234 acceptance rate is asymptotically optimal for random-walk Metropolis algorithms in certain highdimensional limits. We briefly list the theoretical assumptions of these works that are most relevant to our simulations. We first describe the seminal result of Roberts et al. (1997) alluded to in the introduction in full.

Assumption Set 2.1. (Roberts et al. (1997)) Let \mathbf{x} be a random variable from the target density π . The target distribution has an i.i.d. product density $\pi(\mathbf{x}) = \prod_{i=1}^{d} f(x_i)$. The component density $f : \mathbb{R} \to \mathbb{R}^+$ satisfies the following regularity conditions:

- 1. f is positive and twice continuously differentiable on its support.
- 2. The derivative of its logarithm, f'/f, is Lipschitz continuous.
- 3. The following moment conditions hold:

$$\mathbb{E}_f\left[\left(\frac{f'(X)}{f(X)}\right)^8\right] < \infty \quad and \quad \mathbb{E}_f\left[\left(\frac{f''(X)}{f(X)}\right)^4\right] < \infty \tag{2}$$

Under the conditions in Assumption 2.1, Roberts et al. (1997) showed that for an RWM algorithm with a Gaussian proposal distribution starting in stationarity, the sped-up process converges to a limiting Langevin diffusion. Maximizing the speed of

this limiting process, in turn, corresponds to tuning the algorithm to an asymptotic acceptance rate of approximately 0.234.

The Roberts et al. (1997) result works because an i.i.d. product target looks locally the same no matter which coordinate you observe. Then, the analysis of Sherlock (2013) shows that a similar homogeneity emerges for a much wider class of targets once the dimension is large.

Assumption Set 2.2. (Sherlock (2013)) Let \mathbf{x} be a random variable from the target density π , $M(\mathbf{x}) = ||\nabla \log \pi(\mathbf{x})||$ be the norm of the gradient of the log target, and $H(\mathbf{x}) = -\nabla^2 \log \pi(\mathbf{x})$ be the negative Hessian matrix of the log target. The primary conditions are:

1. Target weak shell. There exist sequences $\{M^{(d)}\}\$ and $\{H^{(d)}\}$, as $d \to \infty$:

$$\frac{M^{(d)}(\mathbf{x}^{(d)})}{M^{(d)}} \xrightarrow{p} 1 \quad and \quad \frac{H^{(d)}(\mathbf{x}^{(d)})}{H^{(d)}} \xrightarrow{p} 1 \tag{3}$$

2. Hessian relative variability. The Hessian is locally stable, meaning it does not vary significantly on the scale of a typical proposal jump. Formally, for any fixed $\mu > 0$ and $\delta > 0$, and for a standard Gaussian random vector $\mathbf{z}^{(d)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ independent of the target state $\mathbf{x}^{(d)}$, we require that as $d \to \infty$,

$$\mathbb{P}_{\mathbf{x},\mathbf{z}}\left(\frac{1}{H^{(d)}}\left|\left(\mathbf{z}^{(d)}\right)^{T}\Delta\left(\mathbf{x}^{(d)},t\mu\frac{M^{(d)}}{H^{(d)}}\mathbf{z}^{(d)}\right)\mathbf{z}^{(d)}\right| < \delta \quad \text{for all } t \in [0,1]\right) \to 1.$$
(4)

3. Target eccentricity. The curvature is not disproportionately concentrated in any single direction. Denoting the eigenvalues of $\mathbf{H}(\mathbf{x})$ by $\beta_i(\mathbf{x})$:

$$\frac{\max_i |\beta_i(\mathbf{x}^{(d)})|}{\sum_{j=1}^d \beta_j(\mathbf{x}^{(d)})} \xrightarrow{p} 0.$$
(5)

4. **Proposal shell.** There exists a sequence $\{k_u^{(d)}\}$ such that the sequence of ddimensional spherically symmetric proposals $\{Q^{(d)}\}$ satisfies this as $d \to \infty$:

$$\frac{||Q^{(d)}||}{k_u^{(d)}} \xrightarrow{m.s.} 1. \tag{6}$$

5. Jump distance matrix eccentricity. If the ESJD is defined with respect to a positive definite symmetric matrix $\mathbf{T}^{(d)}$ as $\mathbb{E}[(\mathbf{Y}^{(d)})^T \mathbf{T}^{(d)} \mathbf{Y}^{(d)}]$, then the eigenvalues of this jump distance matrix must be similarly well-behaved. Let $\tau_i(d)$ be the eigenvalues of $\mathbf{T}^{(d)}$. When $d \to \infty$:

$$\frac{\max_i \tau_i(d)}{\sum_{j=1}^d \tau_j(d)} \to 0.$$
(7)

Put together, these geometric regularities recover the same weak-convergence proof as Roberts et al. (1997) but with weaker assumptions on the target.

3 Random-Walk Metropolis Simulations

In this section, we explore the applicability of the 0.234 acceptance rate figure for RWM algorithms using several practical RWM examples in lower dimensions. We begin with simple examples where the target densities meet all the key conditions in Roberts et al. (1997) (provided as Assumption 2.1) required for the 0.234 acceptance rate to be optimal. Next, we investigate examples that do not use Gaussian proposal distributions, and then we explore other target densities that may lack these assumptions, such as smoothness or the i.i.d. product form. We estimate the Expected Squared Jumping Distance (ESJD) of the algorithm as a metric of the algorithm's efficiency as mentioned in Section 2.2 and aim to see under what conditions an acceptance rate of approximately 0.234 maximises the ESJD for various target and proposal distributions in lower dimensions.

Our experimental procedure is as follows: we run RWM simulations using (except for Section 3.2) a Gaussian proposal distribution

$$Q(\mathbf{y} \mid \mathbf{x}) = \mathcal{N}(\mathbf{y} \mid \mathbf{x}, \sigma^2 I_d) \tag{8}$$

where σ is a proposal scaling factor, or σ^2 can be thought of as a variance, and I_d is the *d*-dimensional identity matrix. We run 40 different RWM simulations where each simulation has a different scaling factor, and each simulation has a burn-in of at least 1,000 steps before taking at least 200,000 steps. It is important to note that functions of acceptance rates are somewhat flat around the optimal values and that the values and figures reported are still prone to random error inherent in a Monte Carlo simulation. Therefore, we repeat this procedure over at least 20 computer seeds and average the results to reduce random error. This procedure is applied in each subsection and specific simulation details may be mentioned in a subsection if the experiment used more than these minimum values.

3.1 Simpler Examples: Gamma and Beta i.i.d. Targets

We begin with simpler examples demonstrating that the 0.234 acceptance rate is still roughly optimal for lower dimensions with any i.i.d. product target distribution and a Gaussian proposal. As mentioned in Section 2, Roberts et al. (1997) proved the 0.234 optimal acceptance rate for i.i.d. product target distributions with a standard Gaussian proposal distribution in the infinite-dimensional limit, and Roberts and Rosenthal (2001) and Bédard (2008) extended this with MCMC simulations showing that with a i.i.d. product target distribution and Gaussian proposal, the 0.234 acceptance rate seems approximately optimal for dimensions as low as 10. Our experiment provides further evidence for this being the case, and we extend this with simulations for d = 2, 5. We present two i.i.d. target densities π_1, π_2 with different single-dimension component densities f_1, f_2 :

$$\pi_1(\mathbf{x}) = \prod_{i=1}^n f_1(x_i), \quad f_1(x) = \text{Gamma}(x \mid 3, 2), \tag{9}$$

$$\pi_2(\mathbf{x}) = \prod_{i=1}^n f_2(x_i), \quad f_2(x) = \text{Beta}(x \mid 3, 2)$$
(10)

The first example sets f_1 with shape parameter 3 and scale parameter 2, and the second example sets f_2 with shape parameters 3 and 2. We use a multivariate Gaussian with a mean of 0 as the increment proposal distribution in both cases. We run this experiment in dimensions $d \in \{2, 5, 10, 30, 50, 100\}$.

Table 1: Empirical optimal RWM acceptance rates (i.e., maximizing ESJD) for i.i.d. product targets with a Gaussian proposal.

Target Density	d=2	d = 5	d = 10	d = 30	d = 50	d = 100
$ \pi_1: \text{ i.i.d. Gamma}(3, 2) \pi_2: \text{ i.i.d. Beta}(3, 2) $	$0.3036 \\ 0.3903$	$\begin{array}{c} 0.2378 \\ 0.2937 \end{array}$	$\begin{array}{c} 0.2199 \\ 0.2561 \end{array}$	$\begin{array}{c} 0.2101 \\ 0.2319 \end{array}$	$0.2141 \\ 0.2248$	$0.2140 \\ 0.2159$

The resulting ESJD-maximising acceptance rates are summarised in Table 1 and the i.i.d. Beta target has its trends of ESJD versus acceptance rate with different dimensions displayed in Figure 1. For the Gamma target, the optimal acceptance rate drops from 0.304 at d = 2 to around 0.21 once $d \ge 10$; for the Beta target they fall from 0.390 to around 0.22 over the same range. Figure 1 confirms that the ESJD curves have a moving plateau towards a centre slightly below the classical 0.234 value as the dimension grows. Our simulation provides further evidence that the optimal acceptance rate for maximizing ESJD is approximately 0.234 even in low dimensions.



Fig. 1: ESJD vs. acceptance rate for the i.i.d. Beta(3, 2) target distribution π_2 under RWM with a Gaussian proposal in dimensions $d \in \{2, 5, 10, 30, 50, 100\}$ from top-left to bottom-right. Red dotted line indicates an acceptance rate of 0.234.

Since most ESJD versus acceptance rate plots look fairly similar and have a similar progression as dimension changes, the full set of curves for ESJD versus acceptance rate plots for other subsections are in Appendix A, including the results for the i.i.d. Gamma in Figure 9. As an aside, because the plateau widens appreciably in higher dimensions, the coarse 40-point grid pinpoints the optimum to approximately ± 0.01 .

8

3.2 Non-Gaussian Proposal Densities

Next, we investigate the efficacy of different proposal distributions for the Metropolis algorithm. One of the conditions assumed by the 0.234 theorem of Roberts et al. (1997) is the use of a multivariate Gaussian increment proposal distribution $Q = \mathcal{N}(0, \sigma^2 I_d)$. The additional results in Roberts and Rosenthal (2001) depend on the assumptions of light-tailed proposals. In contrast, Jarner and Roberts (2002) extend this analysis by examining the behaviour of the algorithms with heavy-tailed proposals. They find that, unlike the diffusion behaviour seen with light-tailed proposals, heavy-tailed proposals lead to a different dynamic characterized by abrupt movements followed by periods of inactivity. Furthermore, Neal and Roberts (2011) investigated optimal scaling results for heavy-tailed proposal distributions such as the Cauchy distribution, and spherical proposal distributions of fixed radius. In their RWM simulation study with three different continuous i.i.d. product target densities, both the Gaussian and Cauchy proposals had an estimated asymptotically optimal acceptance rate of approximately 0.234 in d = 100, and a spherical proposal with a fixed radius of 2.38 achieved an optimal acceptance rate of about 0.234 in d = 2 and $d \ge 10$. Finally, Sherlock (2013) proved that proposal distributions following a geometric "shell" condition from Assumption 2.2 (6) should yield an optimal acceptance rate of approximately 0.234 when sampling from a suitable target. As both the Laplace and Uniform proposals satisfy this condition, the 0.234 rule should hold for high dimensions, and our experiments serve to verify this theoretical prediction in lower dimensions.

We experiment with two different proposal distributions in $d \in \{2, 5, 10, 20, 50, 100\}$ and set the target density π_3 to a standard multivariate Gaussian:

$$\pi_{3}(\mathbf{x}) = \mathcal{N}(\mathbf{x} \mid \mathbf{0}, I_{d})$$
$$Q_{L}(\mathbf{y} \mid \mathbf{x}) \propto \prod_{i=1}^{d} \exp\left(-\frac{|x_{i}|}{\sigma}\right), \quad Q_{U}(\mathbf{y} \mid \mathbf{x}) = \prod_{i=1}^{d} \frac{\mathbb{I}\left[|y_{i} - x_{i}| \leq \frac{b}{2}\right]}{b}$$
(11)

The first proposal distribution we consider is the multivariate Laplace distribution, which may also be referred to as the double exponential distribution. We use location parameter $\mu = 0$ and covariance parameter $\Sigma = \sigma^2 I_d$, with σ to be chosen. We also consider an example with the uniform proposal distribution within the closed interval $\left[x_i - \frac{b}{2}, x_i + \frac{b}{2}\right]$ for each component *i*, where b > 0 is the interval length.

Table 2: Empirical optimal RWM acceptance rates across dimensions for different proposal distributions targeting a standard multivariate Gaussian.

Proposal	d=2	d = 5	d = 10	d = 20	d = 50	d = 100
Q_L : Laplace Q_U : Uniform	$\begin{array}{c} 0.3780 \\ 0.3194 \end{array}$	$\begin{array}{c} 0.3036 \\ 0.2516 \end{array}$	$\begin{array}{c} 0.2841 \\ 0.2391 \end{array}$	$\begin{array}{c} 0.2570 \\ 0.2368 \end{array}$	$\begin{array}{c} 0.2429 \\ 0.2365 \end{array}$	$0.2377 \\ 0.2316$

As observed in Table 2 and the corresponding Figures 10, 11 in Appendix A, the optimal acceptance rates for both the Laplace and Uniform proposals exhibit a clear trend. In low dimensions, the optimal rate is significantly higher than 0.234, but converges towards 0.234 from above as the dimension increases. For the Uniform

proposal, the optimal rate is already approximately 0.239 at d = 10, and for both proposals, the rates are close to the high-dimensional limit for $d \ge 20$.

These empirical findings provide strong validation for the generalized theory of optimal scaling. While the foundational result of Roberts et al. (1997) was derived for Gaussian proposals, the more general criteria from Sherlock (2013), summarized in Assumption 2.2, extend this theory. The key requirement on the proposal is the shell condition (6) that requires the norm of the proposal vector to concentrate around a constant in high dimensions. Both the multivariate Laplace and Uniform distributions used here are spherically symmetric and satisfy this condition. The theory therefore predicts that the 0.234 acceptance rate should indeed be asymptotically optimal when using these proposals on a suitable target. Our results confirm this prediction, and characterize the finite-dimensional behaviour, demonstrating that convergence to the 0.234 optimum occurs from above for these proposals.

3.3 Multimodal Target Densities

In this subsection, we investigate target distributions with multimodal densities. Many real-world probability and data distributions are multimodal, and standard RWM algorithms may struggle to explore all the modes of a multimodal distribution in a reasonable amount of time particularly when the modes are far apart. Additionally, multimodal densities may not necessarily fall under the i.i.d. product umbrella of target densities, and even if they do, it is debatable how relevant the 0.234 rule is in these cases. For example, a RWM algorithm may achieve a 0.234 acceptance rate in its run over many iterations and yet completely fail to escape a mode to explore other modes. We illustrate this with Figure 2, where we attempt to draw 10 million samples from a multimodal distribution with three modes but fail to escape the central mode.



Fig. 2: Histogram and traceplot of 10 million samples drawn from a multimodal singledimensional target density. The true density values are traced with the dotted lines.

When we have widely-separated modes, a RWM with the Gaussian proposal distribution struggles to reach the other modes in finite time regardless of the proposal scaling and the dimension of the target distribution. Therefore, for these multimodal target density experiments, we deliberately set the modes to be close enough to each other so that a RWM still mixes between the modes reasonably often. We list out the various target densities before detailing the experimental results.

3.3.1 The "Rough Carpet" Target Density

First, we examine a simpler i.i.d. product density. Our first multimodal distribution has the form

$$\pi_4(\mathbf{x}) = \prod_{i=1}^d f_4(x_i), \quad f_4(x) = 0.5 \,\mathcal{N}(x|\mu_1, 1) + 0.3 \,\mathcal{N}(x|\mu_2, 1) + 0.2 \,\mathcal{N}(x|\mu_3, 1) \quad (12)$$

where the single component density f_3 is a one-dimensional density with three modes at $(\mu_1, \mu_2, \mu_3) = (-5, 0, 5)$ and $\mathcal{N}(x|\mu_i, 1)$ indicates the density of the univariate Gaussian distribution at point x with mean μ_i and variance 1. This forms a "rough carpet"-like appearance for the target distribution when we take the product of this single-component density over many dimensions (see the visualization of the first two components in Figure 3 for an example). Since we have 3^d modes, the density values at each mode may not be very high. Yet, since this is a i.i.d. product form, the theoretical result of 0.234 being the optimal acceptance rate should also apply.



Fig. 3: Density of the first two dimensions of the "rough carpet" target distribution with MCMC samples from one RWM simulation plotted in red. One out of every 200 MCMC samples drawn are plotted.

3.3.2 The "Rough Carpet" Target Density: Inhomogeneously Scaled Components

Next, we examine a multimodal distribution with a density of the form

$$\pi_5(\mathbf{x}) = \prod_{i=1}^d C_i f_5(C_i x_i), \ f_5(x) = 0.5 \,\mathcal{N}(x|\mu_1, 1) + 0.3 \,\mathcal{N}(x|\mu_2, 1) + 0.2 \,\mathcal{N}(x|\mu_3, 1).$$
(13)

Note that $f_5 = f_4$, so this density π_5 is a more general form of π_4 by including inhomogeneous component-wise scaling factors C_i , where $C_i > 0$ and $\mathbb{E}(C_i) = 1$ and

 $\operatorname{Var}(C_i) < \infty$. This inhomogeneously-scaled i.i.d. product target density was discussed by Roberts and Rosenthal (2001), which proved the extension of the 0.234 asymptotically optimal acceptance rate to this case. We independently sample the component scaling factors $C_i \sim \operatorname{Uniform}[0.02, 1.98]$.

3.3.3 A Tale of Three Mixtures

The above "rough carpet" distributions are multimodal, but notice that they still follow the general product form $\pi(\mathbf{x}) = \prod_{i=1}^{d} f(x_i)$. Here, we examine a distribution that does not have this product assumption at all; instead, we have a weighted sum of densities as our target density. We examine a mixture of three Gaussians; a multimodal distribution with just three modes regardless of the number of dimensions. The density is

$$\pi_6(\mathbf{x}) = w_1 \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) + w_2 \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2) + w_3 \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_3, \boldsymbol{\Sigma}_3)$$
(14)

where we have non-negative scalar weights $w_i \ge 0$ such that $w_1 + w_2 + w_3 = 1$, and \mathcal{N} is the density function of a multivariate Gaussian with mean $\boldsymbol{\mu}_i \in \mathbb{R}^d$ and covariance matrix $\Sigma_i \in \mathbb{R}^{d \times d}$.



Fig. 4: A histogram and traceplot of the first component from one of the simulations in the three mixture target distribution experiment after 100,000 iterations.

For these experiments, we set the weights $w_1 = w_2 = w_3 = \frac{1}{3}$, the covariances $\Sigma_1 = \Sigma_2 = \Sigma_3 = I_d$, and set the means $\boldsymbol{\mu}_1 = (\epsilon, 0, 0, \dots, 0), \boldsymbol{\mu}_2 = (0, 0, \dots, 0), \boldsymbol{\mu}_3 = (-\epsilon, 0, 0, \dots, 0)$ for some constant $\epsilon \in \mathbb{R}$, i.e. just varying the first component of the means and setting the other components of the means to 0. As mentioned (refer to Figure 2), we choose the constant $\epsilon = 5.0$ in the first dimension after observing via traceplots that jumping between modes still happens frequently enough. We refer the reader to Figure 4 for an example histogram and traceplot to illustrate how the Markov chain explores the three modes.

3.3.4 A Tale of Three Mixtures: Inhomogenous Scaling Factors

We repeat our experiment with inhomogeneous scaling factors applied to each Gaussian mode. Concretely, we instantiate our target density as

$$\pi_7(\mathbf{x}) = \sum_{k=1}^3 w_k \cdot (\prod_{j=1}^d C_j) \cdot \mathcal{N}(\mathbf{x} \odot \mathbf{C} | \mu_k, I_d)$$
(15)

where $C_i \sim \text{Uniform}[0.2, 1.8]$ and $\mathbf{x} \odot \mathbf{C}$ denotes the Hadamard (element-wise) product $(x_1C_1, ..., x_dC_d)$. This allows each Gaussian mode to follow the functional form defined in Roberts and Rosenthal (2001) Section 6; each mode has its own density of the functional form $\pi(\mathbf{x}) = \prod_{i=1}^{d} C_i f(C_i x_i)$, where $\{C_i\}$ are i.i.d. with $\mathbb{E}(C_i) = 1$ and $\mathbb{E}(C_i^2) < \infty$.

3.3.5 Multimodal Experiments and Results

We conduct our experiments in $d \in \{20, 30, 50\}$ using the default experimental procedure described. We summarize our results in Table 3, and refer the reader to Appendix A Figures 12, 13, 14, and 15 for plots of the trend of ESJD with acceptance rate.

Table 3: Empirical optimal RWM acceptance rates formultimodal target distributions.

Target	d = 20	d = 30	d = 50
π_4 : Rough Carpet (homogeneous) π_5 : Rough Carpet (inhomogeneous) π_6 : 3-Mixture (homogeneous) π_7 : 3-Mixture (inhomogeneous)	$\begin{array}{c} 0.2133 \\ 0.2009 \\ 0.2489 \\ 0.2360 \end{array}$	$\begin{array}{c} 0.2202 \\ 0.2174 \\ 0.2454 \\ 0.2334 \end{array}$	$\begin{array}{c} 0.2291 \\ 0.2351 \\ 0.2394 \\ 0.2371 \end{array}$

The results for our multimodal experiments provide strong empirical support for the generalizability of the 0.234 acceptance rate heuristic. We observe that for all four target distributions, the optimal acceptance rate converges towards 0.234 as the dimension increases, even for those that lack an i.i.d. product structure. Their optimal acceptance rates are already very close to 0.234 even in low dimensions (d = 10). Theoretically, the 0.234 value is likely to be justified by the general criteria of Sherlock (2013) that are stated in Assumption 2.2. The three mixture targets (π_6, π_7), while not product densities, are constructed from isotropic Gaussians. In high dimensions, they form a well-behaved geometric shell with low eccentricity, thus satisfying the target conditions for the 0.234 rule to apply. The rough carpet targets π_4, π_5 are i.i.d. products of a more complex multimodal density. The component-wise multimodality creates a more intricate target landscape, which may require higher dimensions before the geometric properties average out to satisfy the shell and stability conditions as effectively; this may explain why their optimal rates approach 0.234 from below.

There is an interesting note that for both families of targets, introducing inhomogeneous scaling consistently pushes the optimal acceptance rate lower in smaller dimensions. This could be explained by the Target Eccentricity condition of Sherlock (2013) (mentioned in Assumption 2.2 (5)). The inhomogeneous scaling factors introduce anisotropy into the target's curvature, creating directions with much higher curvature than others. With a standard isotropic proposal such as the multivariate Gaussian, the optimal proposal scale must shrink to avoid constant rejections in these high-curvature directions. This makes the sampler inefficient in the less-curved directions, degrading the algorithm's overall performance. This less efficient dynamic alters the trade-off between proposal size and acceptance probability, shifting the peak of the ESJD curve to the left. Our results show that for such low-dimensional inhomogeneously scaled targets, the optimal balance is achieved at a lower acceptance rate than the 0.234 ideal. Then, as d grows, the law of large numbers mitigates the effect of

any single scaling factor, the target becomes effectively less eccentric, and the optimal acceptance rate recovers to 0.234.

3.4 Rosenbrock Target Densities

The multimodal experiments in Section 3.3 demonstrated the effectiveness of the 0.234 acceptance rate for multimodal distributions, where the main challenge is mode isolation. In practice, however, many posterior landscapes are not only multi-peaked but also have nonlinear correlation structure, providing challenges for local proposals even when there is only a single global mode. We examine this setting by sampling three different extensions of the two-dimensional Rosenbrock target density to *n*-dimensions (Goodman and Weare, 2010; Dixon and Mills, 1994; Pagani et al., 2022), each of which have different curvature and therefore different difficulty from a MCMC perspective.

Rosenbrock (1960) originally introduced the Rosenbrock function as an optimization problem: a non-convex function which has a global minimum inside a long, narrow parabolic canyon. The Rosenbrock function can then be converted into a probability density that preserves these properties. Because the local correlation structure changes rapidly along the ridge, step sizes that work near the mode are either rejected or become inefficient elsewhere. Consequently, Rosenbrock targets have become common benchmarks to test MCMC algorithms (Haario et al., 2001; Huijser et al., 2022; Hunt-Smith et al., 2024; Militzer, 2025). Including these targets in our empirical study extends our findings to targets with curved correlation structure.

We first describe the unnormalized Rosenbrock target densities, then describe our experimental procedure and results. We use the Rosenbrock kernels described by Pagani et al. (2022).

3.4.1 Full Rosenbrock

The *d*-dimensional Full Rosenbrock distribution (Goodman and Weare, 2010) has the following kernel for $\mathbf{x} \in \mathbb{R}^d$:

$$\pi_8(\mathbf{x}) \propto \exp\left\{-\sum_{i=1}^{d-1} \left[100(x_{i+1} - x_i^2)^2 + (1 - x_i)^2\right]/20\right\}$$
(16)

The original 2-D Rosenbrock resembles a long banana-shaped valley whose floor follows the parabola $x_2 = x_1^2$. As stated in Pagani et al. (2022), as *d* increases, the long, narrow ridge that makes sampling challenging becomes more concentrated around the mode, decreasing the problem's difficulty. However, the variance of x_d increases steeply as each new random variable is directly dependent on the squared value of the previous variable.

3.4.2 Even Rosenbrock

The *d*-dimensional Even Rosenbrock distribution extended from the optimization function in (Dixon and Mills, 1994) has the following kernel for $\mathbf{x}, \boldsymbol{\mu} \in \mathbb{R}^d$, where *d* must be even:

$$\pi_9(\mathbf{x}) \propto \exp\left\{-\sum_{i=1}^{d/2} \left[(x_{2i-1} - \mu_{2i-1})^2 + 100(x_{2i} - x_{2i-1}^2)^2 \right] / 20 \right\}$$
(17)

The Even Rosenbrock partitions the state into $\frac{d}{2}$ independent 2-D Rosenbrock blocks. Because the blocks do not interact, the global target is the product of smaller bananas; the joint density still exhibits the same curved valleys within each block but lacks the long-range, serial dependence of the Full Rosenbrock. Consequently, the optimal proposal scale can adapt locally to each block, and we expect higher ESJD (and hence slightly lower optimal acceptance rates) than in the fully coupled case.

3.4.3 Hybrid Rosenbrock

The *n*-dimensional Hybrid Rosenbrock distribution introduced by Pagani et al. (2022) has the following kernel for $\mathbf{x} \in \mathbb{R}^d$, where $\mu, x_{j,i} \in \mathbb{R}; a, b_{j,i} \in \mathbb{R}^+(\forall j, i)$:

$$\pi_{10}(\mathbf{x}) \propto \exp\left\{-a(x_1-\mu)^2 - \sum_{j=1}^{n_2} \sum_{i=2}^{n_1} b_{j,i}(x_{j,i}-x_{j,i-1}^2)^2\right\}$$
(18)

The integers n_1 and n_2 describe a grid of blocks: there are n_2 vertical stacks, each containing $n_1 - 1$ Rosenbrock-type links and a shared root coordinate x_1 . The total dimension is therefore $d = (n_1 - 1)n_2 + 1$. Within every column j, the curvature coefficients $b_{j,i}$ can vary with depth i, producing steeper or flatter valleys as we move away from the root. Horizontally, however, columns are independent, so the Hybrid form interpolates between the serial dependence of the Full Rosenbrock (one long chain: $n_2 = 1$) and the independent blocks of the Even Rosenbrock ($n_1 = 3, n_2 = d/2$). This allows the difficulty of the distribution to be modified by setting n_1, n_2 ; larger n_1 increases non-linearity within a column, while larger n_2 raises overall dimension without adding extra curvature.

3.4.4 Simulations and Results

For our Rosenbrock simulations, we found that our initial settings of 200,000 Markov chain steps and 20 seeds were insufficient to produce meaningful trends, so we ran more simulations for more steps with the Rosenbrock targets. In this section, we continue using a Gaussian proposal with 40 different variance values. However, we run each simulation for 1,000,000 steps and we repeated each experiment with 30 different seeds instead of the usual 20. For the Full Rosenbrock, we select $d \in \{5, 10, 20, 30\}$. For the Even Rosenbrock, we select $d \in \{4, 10, 20, 30\}$. For the Hybrid Rosenbrock, we select $(n_1, n_2) \in \{(3, 2), (5, 2), (7, 3), (8, 4)\}$ which gives total dimension $d \in \{5, 9, 19, 29\}$ respectively. We also provide a plot of the ESJD versus acceptance rate trend for the 2-dimensional Rosenbrock. All other plots of ESJD versus acceptance rate trends for these distributions are in Appendix A Figures 16, 17, and 18 for plots of the trend of ESJD with acceptance rate.

Our experiments with the Rosenbrock family of targets, presented in Table 4, demonstrate that the 0.234 heuristic remains effective even for distributions with

Table 4: Empirical optimal RWM acceptance rates for three differentd-dimensional Rosenbrock target distributions.

Target	$d=4\mathrm{or}5$	$d=9\mathrm{or}10$	$d=19{\rm or}20$	$d=29\mathrm{or}30$
π_8 : Full Rosenbrock	0.2047	0.2359	0.2384	0.2337
π_9 : Even Rosenbrock	0.0646	0.1430	0.2003	0.2107
π_{10} : Hybrid Rosenbrock	0.1394	0.2272	0.2349	0.2340

non-linear dependencies and curvature. For the Full Rosenbrock (π_8) and Hybrid Rosenbrock (π_{10}) targets, the optimal acceptance rate robustly converges to approximately 0.234 as the dimension increases. Although these targets are not i.i.d. and possess a highly anisotropic banana-shaped geometry, they should satisfy the general criteria from Sherlock (2013) (stated as Assumption 2.2) in the high-dimensional limit. As *d* grows, the complex local correlations become less influential relative to the global structure, allowing the target's geometric properties to stabilize in the manner required by Assumption 2.2. On the other hand, the Even Rosenbrock (π_9) has a significantly lower optimal acceptance rate that increases slowly with dimension. This seems reasonable, as π_9 is a product of independent 2-dimensional Rosenbrock blocks. The proposal scale can then be fit to the local geometry of each 2-dimensional block where the high-dimensional scaling theory is less applicable.



Fig. 5: ESJD vs. acceptance rate for the 2-dimensional Rosenbrock target distribution π_8 under RWM with a Gaussian proposal. Red dotted line indicates an acceptance rate of 0.234.

The result for the 2-dimensional Rosenbrock target, shown in Figure 5, would add support for this hypothesis. Unlike the higher-dimensional cases, the ESJD does not exhibit a clear peak around a non-zero acceptance rate. Instead, it appears to increase monotonically as the acceptance rate approaches zero. This behaviour indicates that for this low-dimensional, highly anisotropic target, the highest efficiency is achieved by proposing extremely large jumps. Such jumps are almost always rejected, but the rare accepted jump is so large that it dominates the ESJD calculation.

3.5 Non-smooth Target Densities: Continuous Hypercube

Now, consider a RWM using a Gaussian proposal and a target density still of the product form $\pi(\mathbf{x}) = \prod_{i=1}^{d} f(x_i)$, but where the individual component density f is

Uniform [a, b].

$$\pi_{11}(\mathbf{x}) = \prod_{i=1}^{d} f_{11}(x_i), \quad f_{11}(x) = \frac{\mathbb{I}[a \le x \le b]}{b-a}$$
(19)

Unlike the previous examples, the component density f and its logarithm are not continuously differentiable on \mathbb{R} due to the discontinuities at the boundaries $x_i = a$ and $x_i = b$. Since this is one of the assumptions in Assumption 2.1, the asymptotic results where the optimal acceptance rate is 0.234 may not apply. This discontinuous target density case has been explored by Neal et al. (2012), which found that for product densities with a discontinuity at the boundary, such as the unit hypercube $[0, 1]^d$, the efficiency of the RWM is maximised with an asymptotic optimal acceptance rate of 0.1353 as $d \to \infty$. We provide simulation studies in various dimensions to investigate the applicability of this particular 0.1353 acceptance rate figure in lower dimensions on the hypercube $[0, 1]^d$. We use 300,000 iterations instead of the typical 200,000 for this target to make the plots sufficiently smooth.

Table 5: Empirical optimal RWM acceptance rates for the continuous hypercubetarget.

Target	d=2	d = 5	d = 10	d = 20	d = 30	d = 50	d = 100
$\pi_{11} {:}$ Hypercube $[0,1]^{\rm d}$	0.4316	0.2767	0.2027	0.1670	0.1525	0.1439	0.1423

Compared to the previous examples (see Appendix A for their ESJD versus acceptance rate curves, which tend to be smooth and gentle at the plateau), which seem to be fairly consistently optimal with an acceptance rate of approximately 0.234, we find (refer to Table 5 and Figure 6) that the optimal acceptance rate for the hypercube target changes significantly as the dimension increases, and notice that the plateau of the curve tends to get quite sharp as the dimension increases. This is because discontinuities introduce additional complexity in the target distribution, affecting the behaviour and performance of the RWM algorithm.

As the dimension d increases, the algorithm's performance converges to the theoretical 0.1353 value. In this specific case, the shrinking optimal acceptance rate might be because of the unique zero or non-zero form of the target density beyond the hypercube boundary. As discussed by Neal et al. (2012), whether a proposed move is accepted depends completely on whether the proposed state is inside the hypercube. Subsequently, when using a Gaussian proposal, the probability that at least one proposed component will lie outside the boundary increases, leading to a higher probability of instant rejection.

The trends in Figure 6 offer practical guidance for tuning. In high dimensions (e.g., d = 100), the ESJD curve peaks sharply near the theoretical optimum of 0.1353, and tuning to 0.234 would result in a significant loss of efficiency. However, in lower dimensions (e.g., d = 2, 5, 10), the optimal acceptance rate is much closer to 0.234 than to 0.1353. The ESJD curve is also flatter, meaning the penalty for slight mistuning is less severe. This suggests a mini-max strategy for practitioners unsure of the dimensionality effects: tuning to an acceptance rate of approximately 0.234 is a robust choice across all dimensions. While it could be suboptimal for very high dimensions, it is nearly empirically optimal for low dimensions, whereas tuning to 0.1353 in these



Fig. 6: ESJD vs. acceptance rate for the i.i.d. Hypercube $[0, 1]^d$ target distribution π_{11} under RWM with a Gaussian proposal in dimensions $d \in \{2, 5, 10, 30, 50, 100\}$ from top-left to bottom-right. Red dotted line indicates an acceptance rate of 0.234. Purple dotted line indicates an acceptance rate of 0.135.

instances would be highly inefficient even though theory supports it over the 0.234 rate for this specific target distribution.

3.6 Neal's Funnel Target Density

To probe the limits of the 0.234 heuristic, we now examine a target distribution that is canonical for its challenging geometry: the funnel distribution introduced by Neal (2003). Unlike the Rosenbrock target, which has localized curvature, Neal's funnel possesses a global, non-linear dependency structure that presents a severe challenge to standard MCMC algorithms. In this distribution, a single variable controls the variance of all other variables, creating an extreme, state-dependent anisotropy. This structure should violate the geometric regularity conditions outlined in Assumption 2.2, particularly the target eccentricity and Hessian stability.

The *d*-dimensional Neal's funnel target density, which we denote π_{12} , is defined for $\mathbf{x} = (x_1, x_2, \ldots, x_d)$. The density is given by the product of a marginal density for the first component x_1 and conditional densities for the other components:

$$\pi_{12}(\mathbf{x}) = \mathcal{N}(x_1 \mid 0, 3^2) \prod_{i=2}^d \mathcal{N}(x_i \mid 0, \exp(x_1))$$
(20)

The key feature is that the variance of each x_i after x_1 is $\exp(x_1)$. This creates a funnel shape in the state space: when x_1 is large and negative, the funnel is a very narrow neck where the x_i 's are tightly constrained; when x_1 is positive, the funnel opens into a wide mouth where the x_i 's can have enormous variance.

For our experiments, we follow the standard parameterization from Neal (2003). We test this target in dimensions $d \in \{5, 10, 20, 30\}$. Due to the difficulty of sampling

this target, we increased our simulation budget significantly, using 1,000,000 iterations per simulation (across 40 seeds, 40 million total Metropolis steps), a burn-in of 50,000 steps, and averaging the results over 30 random seeds for each data point. We continue using the Gaussian proposal with variance as described in (8).



Fig. 7: ESJD vs. acceptance rate for the Neal's funnel target π_{12} under RWM with a Gaussian proposal in dimensions $d \in \{5, 10, 20, 30\}$ from top-left to bottom-right. Red dotted line indicates an acceptance rate of 0.234.

Similarly to the 2-dimensional Rosenbrock, we find that the ESJD increases monotonically as the acceptance rate reduces in lower dimensions. In higher dimensions (d = 20, 30), we observe a similar trend, but at the extreme ends of the curve, we observe a complete failure of the RWM algorithm to accept any reasonable number of steps as the average acceptance rate decreases towards zero, leading to sudden random drops in ESJD. Notice that the curvature of the graph gets steeper with increasing dimension: it is nearly a straight line at d = 5 and is much more curved by d = 30.

Overall, this behaviour indicates an extreme sensitivity to the proposal scaling in higher dimensions. The practical implication is that even minor mistuning of the proposal variance will cause the sampler's efficiency to collapse to zero. This is a direct consequence of the target's geometry violating the conditions of Assumption 2.2; the target has extreme eccentricity, and the Hessian is not locally stable. Consequently, the 0.234 acceptance rate is not a relevant or useful heuristic for this target. The more important conclusion, which Neal (2003) himself meant to show with this example, is that for targets with strong non-linear dependencies, no single proposal scale can be efficient in both the narrow neck and the wide mouth of the funnel, leading to a sampler that inevitably fails as dimensionality grows. This experiment effectively demonstrates a hard boundary for the class of problems on which the 0.234 RWM heuristic can be successfully applied.

4 Parallel Tempering Background

This section offers a brief introduction to the parallel tempering method and its own optimal scaling framework that we conduct experiments on in Section 5. Many of the terms in Section 2 are re-used but have slightly different definitions for the parallel tempering context.

4.1 Parallel Tempering Method

The parallel tempering method (Geyer, 1991), also known as the Metropolis-coupled Markov chain Monte Carlo algorithm or the replica exchange method, is a version of the Metropolis algorithm that specialises in sampling from multimodal distributions. In this case, we sample from a target distribution with density $\pi(\mathbf{x})$ on a state space \mathcal{X} . We define a sequence of tempered target densities $\pi^{\beta_j}(\mathbf{x})$ where $0 \leq \beta_n < \beta_{n-1} < \cdots < \beta_1 < \beta_0 = 1$ are the inverse temperature values. We require that $\pi^{\beta_0}(\mathbf{x}) = \pi(\mathbf{x})$ (that is, that $\beta_0 = 1$). Typically, the tempered target densities are simply powers of the original density, and we assume this for all parallel tempering experiments. Then, we run a Markov chain for each of the n+1 values of β , each with its own tempered target density. Eventually, the $\beta_0 = 1$ chain, also called the "cold" chain, with its stationary density $\pi^{\beta_0}(\mathbf{x}_0)$ should correspond to the original target density of interest $\pi(\mathbf{x})$, and the idea is that using information from other chains with higher temperatures (lower β) will speed up the mixing in the cold chain.

In each iteration of the algorithm, we alternate between within temperature and temperature swap moves. The former indicates that each chain takes a standard RWM update within its respective π^{β_j} . The latter indicates that for two chains with β_j, β_k inverse temperature values that are adjacent in the sequence of β values, switch the chain values of \mathbf{x}_j and \mathbf{x}_k with probability

$$\min\left\{1, \frac{\pi^{\beta_j}(\mathbf{x}_k)\pi^{\beta_k}(\mathbf{x}_j)}{\pi^{\beta_j}(\mathbf{x}_j)\pi^{\beta_k}(\mathbf{x}_k)}\right\}.$$
(21)

4.1.1 Optimal Scaling for Parallel Tempering

Since the hot chains can explore the state space more quickly, we would like to maximise how frequently we can swap values from the hottest chain to the coldest chain so that the cold chain can mix faster and escape local modes. To do this, we maximise the effective speed with which the chain values move along in the inverse temperature domain. The spacings of the inverse temperatures β_i 's are crucial to this efficiency. If β_j and β_k are too far apart, we usually reject these swaps described by (21), but if they are too close, the swaps will not improve mixing. We would like to swap just the right amount to maximise mixing in the temperature domain.

The expected squared jumping distance (ESJD) for parallel tempering thus refers to the expected squared jump in inverse temperatures. Formalising this, when we attempt to swap the chain values between the inverse temperatures β and $\gamma := \beta + \epsilon$ where $\beta, \epsilon > 0$ and $\beta, \gamma \leq 1$, the swap is either accepted, in which case the values move a squared distance of $(\gamma - \beta)^2 = \epsilon^2$, or the swap is rejected, in which case the

distance moved is 0. This leads to a very natural definition for ESJD, which is

$$ESJD = \mathbb{E}[(\gamma - \beta)^2] = \epsilon^2 \times \mathbb{E}[\Pr(\text{swap accepted})].$$
(22)

Maximising the asymptotic ESJD effectively maximises the efficiency of the attempted swap moves in providing mixing in the temperature space, or in other words, maximises the speed with which the chain values move in the inverse temperature space.

There is also a 0.234 swap acceptance rate theoretical result for parallel tempering. The required assumptions for the result are the following:

Assumption Set 4.1. (Atchadé et al. (2011)) Let \mathbf{x} be a random variable from the target density π .

- 1. The target density is an i.i.d. product $\pi(\mathbf{x}) = \prod_{i=1}^{d} f(x_i)$ where f is a singledimensional component density.
- The tempered distributions are powers of the original density, so π^(β)(**x**) = Π^d_{i=1} f^(β)(x_i) = Π^d_{i=1}(f(x_i))^β.
 The inverse-temperature spread ε = d^{-1/2}ℓ for some positive constant ℓ.

Under these assumptions, choosing ℓ such that the inverse temperature spacing yields a swap acceptance probability of approximately 0.234 is also optimal in the sense that it maximises the ESJD of a parallel tempering algorithm.

5 Parallel Tempering Simulations

In this section, we provide parallel tempering experiments that investigate how well the 0.234 optimal swap acceptance rate holds for the various multimodal target distributions in Section 3.3 in finite dimensions. The parallel tempering method is highly effective for exploring multimodal distributions (Atchadé et al., 2011; Hastings, 1970) and thus we aim our experiments at these distributions to learn more about their applicability in practice.

5.1 Target Distributions

5.1.1 Multivariate Gaussian Distribution

We first give an example of the standard multivariate Gaussian distribution to verify the correctness of our implementation. Note that the standard multivariate Gaussian as a target distribution satisfies the necessary conditions of Atchadé et al. (2011) that are stated in Assumption 4.1.

5.1.2 Rough Carpet Distribution

In this section, we use the same target distribution π_4 from (12) as described in Section 3.3.1. Here, we examine the simpler case with homogeneous scaling factors $f_d(\mathbf{x}) = \prod_{i=1}^n f(x_i)$. In the individual component density f(x) definition, we set $(\mu_1, \mu_2, \mu_3) = (-15, 0, 15)$. We provide an example traceplot of the cold chain from one of the experiment simulations in Appendix A Figure 19 to demonstrate what mixing looks like in the first three components (all components are i.i.d.). Unlike the standard

random-walk Metropolis which cannot escape from the central mode in this setting of the target distribution, the values in the cold chain swap between modes considerably often due to the swapping of values between chains.

5.1.3 Three Mixture Distribution

In this section, we use the same target distribution as described in Section 3.3.3 (14). We set weights $w_1 = w_2 = w_3 = \frac{1}{3}$, means $\boldsymbol{\mu}_1 = (-15, 0, 0, \dots, 0)$, $\boldsymbol{\mu}_2 = (0, 0, \dots, 0)$, $\boldsymbol{\mu}_3 = (15, 0, 0, \dots, 0)$ and covariances $\Sigma_1 = \Sigma_2 = \Sigma_3 = I_d$. We provide an example traceplot of the cold chain in Appendix A Figure 20 to demonstrate what mixing looks like in the first three components.

5.2 Methodology

We conduct experiments in dimension $d \in \{10, 20, 30\}$ that examine the trend of ESJD with swap acceptance rate on these distributions. For each distribution, we run many simulations where each simulation is a parallel tempering algorithm with 30 different average swap acceptance rate values. We explain how we set up these algorithms to match these swap rates (through constructing a specific inverse temperature ladder) in Section 5.2.1. Each parallel tempering algorithm runs for 500,000 iterations, so every individual chain in the algorithm takes 500,000 steps. We attempt a temperature swap every 100 steps. Furthermore, we ran each algorithm instance over 20 different seeds and average the results to reduce the effects of randomness caused by a particular seed.

For each experiment, we set the tempered target distribution to simply be the original distribution raised to the power of β . We set the modes to be far away enough such that the other modes were not reached by a standard random-walk Metropolis algorithm after 100,000 iterations, but still reasonably close so that the parallel tempering method would show the cold chain values would swap between modes somewhat frequently. The within-chain proposal distribution for each chain is given by the multivariate Gaussian (8) with $\sigma^2 = 2.38^2/d$, which is suggested by Roberts et al. (1997) to be optimal for high dimensions. We calculate the optimal swap acceptance rate by finding each simulation's specified swap acceptance rate used to construct the iterative temperature ladder that maximises the ESJD, and take the average of the maximising acceptance rates over 20 seeds.

5.2.1 Constructing an Inverse Temperature Ladder Iteratively

The most common method to construct an inverse temperature ladder selects the inverse temperatures using a geometric series spacing. However, since we are examining the 0.234 swap acceptance rule, we need to have a way of constructing the inverse temperature spacings such that the probability of a swap between adjacent chains is approximately 0.234 (or any other value). To construct an inverse temperature ladder with a desired swap acceptance rate s, we use an iterative procedure adapted from (Atchadé et al., 2011, Section 2.2). Starting with $\beta_0 = 1$ and a minimum value $\beta_{min} = 0.01$, we iteratively add β 's to the ladder. Let β_{curr} denote the most recent β in the ladder. We initialize $\rho_n = 0.5$ (n = 1 initially) and set $\beta^* = \beta_{curr}(1 + e^{\rho_n})^{-1}$.

To determine if β^* should be added to the ladder, we draw N samples (our GPU implementation uses N = 1,000,000) from the target distribution tempered by β^* and β_{curr} , and calculate the average swap probability a. If a is within $s \pm 0.0001$, we add β^* to the ladder and set $\beta_{curr} = \beta^*$. Otherwise, we update $\rho_{n+1} = \rho_n + n^{-0.25}(a-s)$ and recalculate β^* . This process continues until $\beta^* \leq \beta_{min}$, at which point we add β_{min} to the ladder and terminate.

5.3 Results

The empirical optimal swap acceptance rates for the three target distributions are summarized in Table 6, and sample ESJD curves for d = 30 shown in Figure 8. The ESJD curves for PT are visibly more jagged than their RWM counterparts. This might be because the estimation of ESJD in this context is more susceptible to Monte Carlo error even with a large number of simulation steps, since the efficiency or acceptance rate of a swap move depends on the joint stationary distribution of two adjacent chains, and this is repeated for each temperature construction. Nonetheless, clear trends emerge from the results.

Table 6: Empirical swap acceptance rates that maximize ESJD for various PT target densities. Based on Fig 8, the jagged pattern of the ESJD curves suggests that the true optimal acceptance rates are in a neighbourhood of these values. See Section 5.2 for calculation of optimal swap acceptance rates.

Target	d = 10	d = 20	d = 30
π_2 : Standard Multivariate Gaussian π_4 : Rough Carpet π_6 : 3-Mixture	$\begin{array}{c} 0.2618 \\ 0.2660 \\ 0.2609 \end{array}$	$0.2685 \\ 0.2803 \\ 0.2660$	$\begin{array}{c} 0.2474 \\ 0.2761 \\ 0.2372 \end{array}$

For both the standard Gaussian (π_2) and the non-product Three-Mixture (π_6) targets, the optimal swap acceptance rate is consistently in the neighbourhood of 0.234. For the Three-Mixture target, the rate converges cleanly towards the theoretical value, from 0.261 in d = 10 to 0.237 in d = 30. The Gaussian target shows a similar pattern, with an optimal rate of 0.247 at d=30.

The standard Gaussian result provides a successful validation of our experimental setup against the theory of Atchadé et al. (2011), as this target satisfies the i.i.d. product form assumption. More importantly, the result for the Three-Mixture target suggests that the 0.234 heuristic for PT swaps is robust to violations of the i.i.d. product assumption, much like its RWM counterpart. Although the target is a sum of densities, its underlying geometric regularity (a shell composed of three isotropic modes) seems sufficient for the optimal scaling 0.234 result to hold. For a practitioner, this suggests that constructing a temperature ladder with a swap acceptance probability of 0.234 is a viable and likely near-optimal strategy, even for complex, non-product multimodal targets, provided they are not pathologically structured.

Notably, the Rough Carpet target (π_4) shows a significant deviation. The optimal swap acceptance rate is consistently higher than 0.234 and increases away from the theoretical value with increasing dimension. This is surprising because the Rough



Fig. 8: ESJD vs. swap acceptance rate for the parallel tempering algorithm with the π_2, π_4, π_6 target distributions in dimension d = 30. Red dotted line indicates a swap acceptance rate of 0.234.

Carpet target does satisfy the i.i.d. product form and other conditions of Assumption 4.1 by Atchadé et al. (2011). The failure of the 0.234 rule here could be because the extreme multimodality of the target could reduce the practical utility of the "hot" chains. The target possesses 3^d modes, and its probability mass is spread thinly across this vast landscape. For a hot chain where the inverse temperature β is close to 0, the tempered density becomes extremely flat and loses the slight peak that a unimodal target would have. Consequently, it fails to efficiently sample from the specific modal regions that are of interest to the colder chains; the information flow from hot to cold chains is broken. To swap successfully more often, the temperature steps must be made smaller, which forces the swap acceptance rate to be higher.

6 Discussion

We have presented an extensive empirical investigation into the generalizability of the 0.234 acceptance rate heuristic for both Random-Walk Metropolis (RWM) and Parallel Tempering (PT) algorithms, using a variety of target and proposal distributions. Our results demonstrate that the 0.234 rule's utility extends beyond the classical i.i.d. product setting of Roberts et al. (1997). The heuristic remains approximately optimal for: (i) non-Gaussian proposals that satisfy a spherical symmetry (shell) condition, such as the Laplace and Uniform distributions; (ii) multimodal targets such as the "rough carpet" i.i.d product with 3^d modes, and non-product targets constructed as a sum of isotropic components, such as our Three-Mixture model; and (iii) even for highly correlated, non-linear targets like the Full and Hybrid Rosenbrock in sufficiently high dimensions ($d \geq 20$). These findings provide empirical validation for the generalized theory of Sherlock (2013), which predicts that as long as the target's geometry is sufficiently regular and isotropic in the high-dimensional limit, the 0.234 rule should hold.

That being said, there are cases where 0.234 does not seem as relevant. The main case is in very low dimensions; in all our experiments with varying target densities, the optimal acceptance rate strays from 0.234 below a certain threshold which varies for each target density. In particular, targets composed of independent, low-dimensional anisotropic blocks (Even Rosenbrock) or those that are globally anisotropic in low dimensions (2D Rosenbrock), have an optimal acceptance rate significantly below

0.234. In these cases, we hypothesize the high-dimensional averaging effect that stabilizes the target's geometry has not occurred, and the sampler's performance is dictated by the challenging local geometry. As the dimension increases, the law of large numbers mitigates the effect of any single scaling factor or local correlation, so the target becomes effectively less eccentric, complying with the theory of Sherlock (2013), and the optimal acceptance rate recovers towards 0.234. On the other hand, the continuous hypercube and Neal's funnel completely fail the 0.234 heuristic. Although the hypercube is continuously differentiable everywhere except on its boundary, which has zero measure, an intuitive argument is that as d increases, the probability that the proposal has all the values inside decreases. As for the funnel, no single proposal scaling can be globally efficient, which is just a limitation of the RWM algorithm.

Next, we discuss the optimal swap acceptance rate figure for parallel tempering and compare it with the optimal acceptance rate for random-walk Metropolis. On one hand, the results for the multivariate Gaussian and three-mixture density examples show that the 0.234 swap acceptance rate figure may be optimal even in lower dimensions. Although costly, if the practitioner is willing and able to construct the inverse temperature ladder iteratively with a spacing given by a swap acceptance probability, the spacing dictated by a 0.234 swap acceptance probability may be optimal for the practitioner. However, the theoretical 0.234 figure may not always be applicable: despite satisfying the i.i.d. conditions of the PT optimal scaling theory (Atchadé et al., 2011), the extreme multimodality (3^d modes) of the "rough carpet" product density appears to render the hot chains ineffective. We hypothesize that because the probability mass is spread so thinly, the hot chains fail to sample from meaningful modal regions, breaking the flow of information to the cold chain. This forces the algorithm to use smaller temperature steps (and thus higher swap acceptance rates) to achieve mixing, representing a practical failure of the 0.234 PT rule.

Not only that, constructing a temperature ladder iteratively with the "optimal spacing" is challenging and time-consuming. It requires the ability to draw samples from the target distribution, which is itself the main goal of the parallel tempering algorithm. Without a direct way to sample from the target, a practitioner could run an MCMC algorithm with the β value being proposed for the new addition to the ladder, but this may be time-consuming with many attempts required until a suitable β is found for each step of the ladder.

Since this optimal spacing may not yield the optimal ESJD in practice for finite, lower-dimensional targets and with a finite number of runs, an alternative approach is to adaptively adjust the spacings of the temperature ladder during the algorithm's runs based on its recent ESJD performance over a recent window as guidance to practitioners. This is known as adaptive parallel tempering (Miasojedow et al., 2013). Previous literature shows that adaptive parallel tempering algorithms can achieve substantial efficiency gains over a standard geometric temperature spacing (Vousden et al., 2015) while still converging to the target distribution (Miasojedow et al., 2013) under the assumptions that the target distribution is sufficiently regular and has tails decaying faster than exponentially. Adaptive parallel tempering may suffice to be "good enough" for most practitioners whilst avoiding the costly temperature construction algorithm.

Lastly, while we have established that the 0.234 acceptance rate for RWM is still fairly good even in lower dimensions and on some target distributions that may not have the i.i.d. product form, we provide some precautions to the practitioner. The 0.234 acceptance rate is a good heuristic to tune the algorithm to, yet there are still many cases where it is not optimal in terms of efficiency, such as the continuous hypercube example in Section 3.5. Even if the 0.234 acceptance rate is in fact optimal, the acceptance rate alone does not necessarily guarantee good samples, as illustrated by Figure 2. What literature (Roberts and Rosenthal, 2001; Sherlock et al., 2010) indicates is that another very important thing in practice beyond just the acceptance rate is defining a good proposal covariance matrix Σ_d that is approximately proportional to the target covariance matrix Σ_{π} ; doing so can significantly increase the asymptotic relative efficiency. Realistically, it may be impossible to know Σ_{π} in advance or even provide a good estimate for it. In fact, understanding how we can extend this result to correlated targets is an ongoing challenge (Yang et al., 2020) and immediately relevant to the practitioner since real-world data tends to have correlations between variables. Practitioners may consider using the adaptive Metropolis algorithm (Haario et al., 2001) to dynamically update the algorithm's proposal covariance matrix. This adaptive algorithm is very useful and yields significant speed benefits in low and high dimensions (Craiu et al., 2009; Roberts and Rosenthal, 2009), but an important potential pitfall to note is that adaptive methods are not guaranteed to converge to the correct target density; in many cases, they may fail to converge or converge to something completely different (Roberts and Rosenthal, 2007, 2009).

Declarations

The authors did not receive any funding for this submitted work and have no relevant financial or non-financial competing interests. The code implementing the experiments and generating the figures used in this work are publicly available on GitHub at https://github.com/aidanmrli/rwm-pt-pytorch.

During the preparation of this revision, the authors used generative AI models (Gemini 2.5 Pro Preview 05-06, Claude 4.0 Sonnet) to assist with coding and manuscript revision. Generative AI was utilized to help generate, debug, and optimize code for the simulations. The new generated code based on PyTorch allows for the use of Graphics Processing Units (GPUs) when running RWM and PT algorithms to significantly speed up simulations and run more simulations at once. For manuscript revision, Generative AI was used to improve phrasing and clarity on text written by the authors. The authors thoroughly reviewed and edited all AI-generated output and take full responsibility for the scientific integrity and final content of this publication.

References

Atchadé, Y.F., Roberts, G.O., Rosenthal, J.S.: Towards optimal scaling of Metropoliscoupled Markov chain Monte Carlo. Statistics and Computing 21(4), 555–568 (2011) https://doi.org/10.1007/s11222-010-9192-1

- Brooks, S., Gelman, A., Jones, G., Meng, X.-L. (eds.): Handbook of Markov Chain Monte Carlo. Chapman and Hall/CRC, New York (2011). https://doi.org/10.1201/ b10905
- Breyer, L.A., Roberts, G.O.: From Metropolis to diffusions: Gibbs states and optimal scaling. Stochastic Processes and their Applications **90**(2), 181–206 (2000) https://doi.org/10.1016/S0304-4149(00)00041-7
- Bédard, M., Rosenthal, J.S.: Optimal scaling of Metropolis algorithms: Heading toward general target distributions. Canadian Journal of Statistics 36(4), 483–503 (2008) https://doi.org/10.1002/cjs.5550360401
- Beskos, A., Roberts, G., Stuart, A.: Optimal Scalings for Local Metropolis-Hastings Chains on Nonproduct Targets in High Dimensions. The Annals of Applied Probability **19**(3), 863–898 (2009)
- Bédard, M.: Weak convergence of Metropolis algorithms for non-i.i.d. target distributions. The Annals of Applied Probability **17**(4), 1222–1244 (2007) https: //doi.org/10.1214/10505160700000096
- Bédard, M.: Efficient Sampling Using Metropolis Algorithms: Applications of Optimal Scaling Results. Journal of Computational and Graphical Statistics 17(2), 312–332 (2008)
- Craiu, R.V., Rosenthal, J., Yang, C.: Learn From Thy Neighbor: Parallel-Chain and Regional Adaptive MCMC. Journal of the American Statistical Association 104(488), 1454–1466 (2009) https://doi.org/10.1198/jasa.2009.tm08393
- Dixon, L.C.W., Mills, D.J.: Effect of rounding errors on the variable metric method. Journal of Optimization Theory and Applications 80(1), 175–179 (1994) https:// doi.org/10.1007/BF02196600
- Geyer, C.J.: Markov chain Monte Carlo maximum likelihood. Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface, 156–163 (1991)
- Goodman, J., Weare, J.: Ensemble samplers with affine invariance. Communications in Applied Mathematics and Computational Science 5(1), 65–80 (2010) https://doi.org/10.2140/camcos.2010.5.65
- Hastings, W.K.: Monte Carlo Sampling Methods Using Markov Chains and Their Applications. Biometrika 57(1), 97–109 (1970) https://doi.org/10.2307/2334940
- Huijser, D., Goodman, J., Brewer, B.J.: Properties of the affine-invariant ensemble sampler's 'stretch move' in high dimensions. Australian & New Zealand Journal of Statistics 64(1), 1–26 (2022) https://doi.org/10.1111/anzs.12358

Hunt-Smith, N.T., Melnitchouk, W., Ringer, F., Sato, N., Thomas, A.W., White, M.J.:

Accelerating Markov Chain Monte Carlo sampling with diffusion models. Computer Physics Communications **296**, 109059 (2024) https://doi.org/10.1016/j.cpc.2023. 109059

- Haario, H., Saksman, E., Tamminen, J.: An adaptive Metropolis algorithm. Bernoulli, 223–242 (2001)
- Jarner, S., Roberts, G.: Polynomial convergence rates of Markov chains. The Annals of Applied Probability 12, 224–247 (2002)
- Kamatani, K.: Random walk metropolis algorithm in high dimension with nongaussian target distributions. Stochastic Processes and their Applications 130(1), 297–327 (2020) https://doi.org/10.1016/j.spa.2019.03.002
- Militzer, B.: Ensemble Monte Carlo calculations with five novel moves. Computer Physics Communications **307**, 109424 (2025) https://doi.org/10.1016/j.cpc.2024. 109424
- Miasojedow, B., Moulines, E., Vihola, M.: An adaptive parallel tempering algorithm. Journal of Computational and Graphical Statistics **22**(3), 649–664 (2013)
- Mattingly, J.C., Pillai, N.S., Stuart, A.M.: Diffusion limits of the random walk Metropolis algorithm in high dimensions. The Annals of Applied Probability 22(3), 881–930 (2012) https://doi.org/10.1214/10-AAP754
- Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E.: Equation of State Calculations by Fast Computing Machines. The Journal of Chemical Physics 21(6), 1087–1092 (1953) https://doi.org/10.1063/1.1699114
- Neal, R.M.: Slice Sampling. The Annals of Statistics **31**(3), 705–741 (2003)
- Neal, P., Roberts, G.: Optimal scaling for partially updating MCMC algorithms. The Annals of Applied Probability 16(2), 475–515 (2006) https://doi.org/10.1214/ 105051605000000791
- Neal, P., Roberts, G.: Optimal Scaling of Random Walk Metropolis Algorithms with Non-Gaussian Proposals. Methodology and Computing in Applied Probability 13(3), 583–601 (2011) https://doi.org/10.1007/s11009-010-9176-9
- Neal, P., Roberts, G., Yuen, W.K.: Optimal scaling of random walk Metropolis algorithms with discontinuous target densities. The Annals of Applied Probability 22(5) (2012) https://doi.org/10.1214/11-AAP817
- Pasarica, C., Gelman, A.: Adaptively Scaling the Metropolis Algorithm Using Expected Squared Jumped Distance. Statistica Sinica 20(1), 343–364 (2010)
- Pagani, F., Wiegand, M., Nadarajah, S.: An n-dimensional Rosenbrock distribution for Markov chain Monte Carlo testing. Scandinavian Journal of Statistics **49**(2),

657–680 (2022) https://doi.org/10.1111/sjos.12532

- Robert, C.P., Casella, G.: Monte Carlo Statistical Methods. Springer Texts in Statistics. Springer, New York, NY (2004). https://doi.org/10.1007/978-1-4757-4145-2 . http://link.springer.com/10.1007/978-1-4757-4145-2
- Roberts, G.O., Gelman, A., Gilks, W.R.: Weak convergence and optimal scaling of random walk Metropolis algorithms. The Annals of Applied Probability 7(1), 110–120 (1997) https://doi.org/10.1214/aoap/1034625254
- Roberts, G.O.: Optimal metropolis algorithms for product measures on the vertices of a hypercube. Stochastics and Stochastic Reports **62**(3-4), 275–283 (1998) https://doi.org/10.1080/17442509808834136
- Rosenbrock, H.H.: An Automatic Method for Finding the Greatest or Least Value of a Function. The Computer Journal **3**(3), 175–184 (1960) https://doi.org/10.1093/ comjnl/3.3.175
- Rosenthal, J.S.: Minorization Conditions and Convergence Rates for Markov Chain Monte Carlo. Journal of the American Statistical Association 90(430), 558–566 (1995) https://doi.org/10.2307/2291067
- Roberts, G.O., Rosenthal, J.S.: Optimal scaling for various Metropolis-Hastings algorithms. Statistical Science 16(4), 351–367 (2001) https://doi.org/10.1214/ss/1015346320
- Roberts, G.O., Rosenthal, J.S.: Coupling and Ergodicity of Adaptive Markov Chain Monte Carlo Algorithms. Journal of Applied Probability 44(2), 458–475 (2007)
- Roberts, G.O., Rosenthal, J.S.: Examples of Adaptive MCMC. Journal of Computational and Graphical Statistics 18(2), 349–367 (2009) https://doi.org/10.1198/jcgs. 2009.06134
- Roberts, G., Rosenthal, J.: Minimising MCMC variance via diffusion limits, with an application to simulated tempering. The Annals of Applied Probability **24** (2014) https://doi.org/10.1214/12-AAP918
- Sherlock, C., Fearnhead, P., Roberts, G.O.: The Random Walk Metropolis: Linking Theory and Practice Through a Case Study. Statistical Science 25(2), 172–190 (2010) https://doi.org/10.1214/10-STS327
- Sherlock, C.: Optimal Scaling of the Random Walk Metropolis: General Criteria for the 0.234 Acceptance Rule. Journal of Applied Probability 50(1), 1–15 (2013)
- Sherlock, C., Roberts, G.: Optimal scaling of the random walk Metropolis on elliptically symmetric unimodal targets. Bernoulli 15(3), 774–798 (2009) https://doi.org/ 10.3150/08-BEJ176

- Sherlock, C., Thiery, A.H., Roberts, G.O., Rosenthal, J.S.: On the efficiency of pseudomarginal random walk Metropolis algorithms. The Annals of Statistics 43(1), 238– 275 (2015) https://doi.org/10.1214/14-AOS1278
- Vousden, W.D., Farr, W.M., Mandel, I.: Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations. Monthly Notices of the Royal Astronomical Society 455(2), 1919–1937 (2015) https://doi.org/10.1093/mnras/ stv2422
- Yang, J., Roberts, G.O., Rosenthal, J.S.: Optimal scaling of random-walk metropolis algorithms on general target distributions. Stochastic Processes and their Applications 130(10), 6094–6132 (2020) https://doi.org/10.1016/j.spa.2020.05.004

A Plots for RWM simulations



Fig. 9: ESJD vs. acceptance rate for the i.i.d. Gamma(3, 2) target distribution π_2 (Eq 10) under RWM with a Gaussian proposal in dimensions $d \in \{2, 5, 10, 30, 50, 100\}$ from top-left to bottom-right. Red dotted line indicates an acceptance rate of 0.234.



Fig. 10: ESJD vs. acceptance rate for the standard multivariate Gaussian target distribution π_3 under RWM with a multivariate Laplace proposal Q_L (Eq 11) in dimensions $d \in \{2, 5, 10, 30, 50, 100\}$ from top-left to bottom-right. Red dotted line indicates an acceptance rate of 0.234.



Fig. 11: ESJD vs. acceptance rate for the standard multivariate Gaussian target distribution π_3 under RWM with a multivariate Uniform proposal Q_U (Eq 11) in dimensions $d \in \{2, 5, 10, 30, 50, 100\}$ from top-left to bottom-right. Red dotted line indicates an acceptance rate of 0.234.



Fig. 12: ESJD vs. acceptance rate for the i.i.d. rough carpet target distribution π_4 (Eq 12) under RWM with a Gaussian proposal in dimensions $d \in \{20, 30, 50\}$ from left to right. Red dotted line indicates an acceptance rate of 0.234.



Fig. 13: ESJD vs. acceptance rate for the inhomogeneously scaled i.i.d. rough carpet target distribution π_5 (Eq 13) under RWM with a Gaussian proposal in dimensions $d \in \{20, 30, 50\}$ from left to right. Red dotted line indicates an acceptance rate of 0.234.



Fig. 14: ESJD vs. acceptance rate for the three mixture target distribution π_6 (Eq 14) under RWM with a Gaussian proposal in dimensions $d \in \{20, 30, 50\}$ from left to right. Red dotted line indicates an acceptance rate of 0.234.



Fig. 15: ESJD vs. acceptance rate for the inhomogeneously scaled three mixture target distribution π_7 (Eq 14) under RWM with a Gaussian proposal in dimensions $d \in \{20, 30, 50\}$ from left to right. Red dotted line indicates an acceptance rate of 0.234.



Fig. 16: ESJD vs. acceptance rate for the Full Rosenbrock target π_8 (Eq 16) under RWM with a Gaussian proposal in dimensions $d \in \{5, 10, 20, 30\}$ from top-left to bottom-right. Red dotted line indicates an acceptance rate of 0.234.



Fig. 17: ESJD vs. acceptance rate for the Even Rosenbrock target π_9 (Eq 17) under RWM with a Gaussian proposal in dimensions $d \in \{5, 10, 20, 30\}$ from top-left to bottom-right. Red dotted line indicates an acceptance rate of 0.234.



Fig. 18: ESJD vs. acceptance rate for the Hybrid Rosenbrock target π_{10} (Eq 18) under RWM with a Gaussian proposal in dimensions $d \in \{5, 9, 19, 29\}$ from top-left to bottom-right. Red dotted line indicates an acceptance rate of 0.234.



Fig. 19: Traceplot of the first three dimensions of the cold chain from one of the parallel tempering simulations on the rough carpet target distribution π_4 (Eq 12).



Fig. 20: Traceplot of the first three dimensions of the cold chain from one of the parallel tempering simulations on the three-mixture target distribution π_6 (Eq 14).