Variance Bounding of Delayed-Acceptance Kernels

Chris Sherlock¹ · Anthony Lee²

Received: 21 May 2021 / Revised: 10 November 2021 / Accepted: 11 November 2021 /
Published online: 22 November 2021
© The Author(s) 2021

Abstract

A delayed-acceptance version of a Metropolis–Hastings algorithm can be useful for Bayesian inference when it is computationally expensive to calculate the true posterior, but a computationally cheap approximation is available; the delayed-acceptance kernel targets the same posterior as its associated “parent” Metropolis-Hastings kernel. Although the asymptotic variance of the ergodic average of any functional of the delayed-acceptance chain cannot be less than that obtained using its parent, the average computational time per iteration can be much smaller and so for a given computational budget the delayed-acceptance kernel can be more efficient. When the asymptotic variance of the ergodic averages of all L^2 functionals of the chain are finite, the kernel is said to be variance bounding. It has recently been noted that a delayed-acceptance kernel need not be variance bounding even when its parent is. We provide sufficient conditions for inheritance: for non-local algorithms, such as the independence sampler, the discrepancy between the log density of the approximation and that of the truth should be bounded; for local algorithms, two alternative sets of conditions are provided. As a by-product of our initial, general result we also supply sufficient conditions on any pair of proposals such that, for any shared target distribution, if a Metropolis-Hastings kernel using one of the proposals is variance bounding then so is the Metropolis-Hastings kernel using the other proposal.

Keywords Metropolis-Hastings · Delayed-acceptance · Variance bounding · Conductance

AMS Classification 60J10 · 65C40 · 47A10

1 Introduction

The Metropolis-Hastings (MH) algorithm is widely used to approximately compute expectations with respect to complicated high-dimensional posterior distributions (e.g. Gilks et al. (1996); Geyer (2011)). The algorithm requires that it be possible to evaluate

✉ Chris Sherlock
c.sherlock@lancaster.ac.uk
<https://chrissherlock.github.io/>

¹ Department of Mathematics and Statistics, Lancaster University, Lancaster LA1 4YF, UK

² School of Mathematics, Fry Buiding, University of Bristol, Bristol BS8 1UG, UK

point-wise the density of the distribution of interest (throughout this article, all densities are with respect to Lebesgue measure) up to an arbitrary constant of proportionality.

In many problems the target posterior density is computationally expensive to evaluate. When a computationally-cheap approximation, or surrogate, is available, the delayed-acceptance Metropolis-Hastings (DAMH) algorithm also known as the two-stage algorithm, and a special case of the surrogate-transition method Liu (2001); Christen and Fox (2005); Higdon et al. (2011) leverages the surrogate to produce a new Markov chain that still targets the original distribution of interest. A first ‘screening’ stage substitutes the surrogate density for the true density in the standard formula for the MH acceptance probability; proposals which fail at this stage are discarded. Only proposals that pass the first stage are considered in the second ‘correction’ stage, where it is necessary to evaluate the true posterior density at the proposed value.

Delayed acceptance (DA) algorithms have been applied in a variety of settings with the approximate density obtained in a variety of different ways, for example: a coarsening of a numerical grid in Bayesian inverse problems Christen and Fox (2005); Moulton et al. (2008); Cui et al. (2011), subsampling from big-data Payne and Mallick (2014); Banterle et al. (2019); Quiroz et al. (2018), a tractable approximation to a stochastic process Smith (2011); Golightly et al. (2015), or a direct, nearest-neighbour approximation to the truth using previous values Sherlock et al. (2017).

For a Markov kernel, P , with a stationary distribution of π , and an associated chain $\{X_t\}_{t=1}^\infty$, the *asymptotic variance* of any functional, h , is defined to be

$$\text{var}(h, P) := \lim_{n \rightarrow \infty} n \text{Var} \left[\frac{1}{n} \sum_{i=1}^n h(X_i) \right],$$

where $X_1 \sim \pi$. A lower asymptotic variance is thus associated, in practice, with a greater accuracy in estimating $\mathbb{E}_\pi[h(X)]$ using a realisation of length $n \gg 1$ from the distribution of the chain. In terms of the asymptotic variance of any functional of the chain, the DAMH kernel cannot be more efficient than the parent MH kernel; however the computational cost per iteration is, typically, reduced considerably. The almost-negligible computational cost of the screening stage also, typically, facilitates proposals that have a larger chance of being rejected than the MH proposal, but where the pay-off on acceptance is so much larger that the expected overall movement per unit of time increases. When efficiency is measured in terms of effective samples per second, gains of over an order of magnitude have been reported (e.g. Golightly et al. (2015)).

A Markov kernel P with a stationary distribution of π is termed *variance bounding* if $\text{var}(h, P) < \infty$ for all $h \in L^2(\pi)$, the Hilbert space of functions that are square-integrable with respect to π . Equivalently there exists $K < \infty$ such that $\text{var}(h, P) \leq K \text{Var}_\pi[h]$ for all such h . This property was named and studied in Roberts and Rosenthal (2008), where it was shown to be equivalent to the existence of a ‘usual’ central limit theorem (CLT); that is, a CLT where the limiting variance is the asymptotic variance.

Intuitively, the variance-bounding property embodies desirable behaviour for a chain started at equilibrium. In practice, the chain is not started at equilibrium, but asymptotically the bias that results from this is negligible compared with the variance. An alternative natural requirement is that the chain converge to equilibrium geometrically quickly (rather than, say, polynomially quickly). A Markov chain kernel, P , with stationary distribution π is *geometrically ergodic* (e.g. Roberts and Rosenthal (1997,

2004); Meyn and Tweedie (1993) Chapter 15) if there exist $\rho > 0$ and $M : \mathcal{X} \rightarrow [0, \infty)$ that is finite π -almost everywhere, such that

$$|P^n(x, \mathcal{A}) - \pi(\mathcal{A})| \leq M(x)\rho^n$$

for all $\mathcal{A} \in \mathcal{F}$, $x \in \mathcal{X}$ and $n \in \mathbb{N}$, where P^n denotes the n -step transition kernel.

Although the motivations behind the definitions of variance bounding and geometric ergodicity, mixing at equilibrium and convergence to equilibrium, are quite different, for a large class of algorithms, including those studied in this article, these two properties are very closely linked as we will describe in Sect. 2.2. Indeed, for delayed-acceptance algorithms, under weak conditions the two properties are equivalent (see Proposition 1).

Theoretical properties of the efficiency of delayed-acceptance algorithms have been studied in Banterle et al. (2019), Sherlock et al. (2021) and Franks and Vihola (2020). The first contribution from Banterle et al. (2019) is an example delayed-acceptance algorithm which fails to inherit geometric ergodicity from its parent Metropolis-Hastings algorithm (see Example 1 in Sect. 2.3 of this article); a simple sufficient condition for inheritance of geometric ergodicity, uniformly good behaviour of the ratios r_1 and r_2 that we define in Eq. 3, is also supplied. Finally, an idealised setting where the cheap approximation is perfectly accurate is explored to obtain tuning guidelines for λ in the delayed-acceptance random walk Metropolis algorithm. Sherlock et al. (2021) examines this tuning issue further, proving a limiting diffusion for the first component of the delayed-acceptance Markov chain, and providing robust tuning guidelines that account for the error in the cheap approximation; the article then extends these guidelines to the pseudo-marginal version of the algorithm. Finally, Franks and Vihola (2020) compares the asymptotic variance of a general pseudo-marginal delayed-acceptance algorithm with the variance of an algorithm that applies importance-sampling to the output of an MCMC algorithm targeting the cheap approximation directly.

Using our Proposition 1, the lack of inheritance of geometric ergodicity in the example in Banterle et al. (2019) is equivalent to a lack of inheritance of the variance bounding property: even though the asymptotic variance using the parent MH kernel is finite for all $h \in L^2(\pi)$, there exist $h \in L^2(\pi)$ for which the asymptotic variance using the DA kernel is infinite. For such h , estimated quantities such as effective sample size (e.g. Hoff (2009) are invalid, and consequent, standard CLT-based intuitions about the sizes of typical errors in estimates of $\mathbb{E}_\pi[h]$ from the chain *do not hold*.

We investigate the conditions under which a DAMH kernel inherits variance bounding from its MH parent and, as a by product, discover conditions under which two different proposals produce MH kernels that are equivalent in terms of whether or not they are variance bounding. Section 2 provides the background and two motivating examples, while Sect. 3 provides some key definitions, a general inheritance result applicable to all propose-accept-reject kernels, and sufficient conditions for variance-bounding equivalence between two Metropolis-Hastings proposals. Section 4 contains our results for standard DA algorithms with further illustrative examples, and includes parent MH algorithms where the proposal depends upon the form of the density, so that the proposal for a computationally cheap DA kernel would naturally depend on the surrogate. Numerical experiments are performed in Sect. 5 and the article concludes with a discussion. All proofs are deferred to Appendix 1.

2 Background, Notation and Motivation

Throughout this article all Markov chains are assumed to be on a statespace $(\mathcal{X}, \mathcal{F})$, with $\mathcal{X} \subseteq \mathbb{R}^d$ Lebesgue measurable, and $\mathcal{F}$ the σ -algebra of all Lebesgue-measurable sets in $\mathcal{X}$. The target and surrogate distributions are denoted by π and $\hat{\pi}$, respectively, and they are assumed to have densities of $\pi(x)$ and $\hat{\pi}(x)$ with respect to Lebesgue measure.

2.1 Metropolis-Hastings and Delayed-Acceptance Kernels

The Metropolis-Hastings kernel has a proposal density $q(x, y)$ and an acceptance probability $\alpha(x, y) = 1 \wedge r(x, y)$ where

$$r(x, y) := \frac{\pi(y)q(y, x)}{\pi(x)q(x, y)}. \quad (1)$$

With $\bar{\alpha}(x) := \int \alpha(x, y)q(x, y)dy$, the Metropolis-Hastings (MH) kernel is then

$$P(x, dy) := q(x, y)dy \alpha(x, y) + [1 - \bar{\alpha}(x)]\delta_x(dy). \quad (2)$$

An iteration of the corresponding MH algorithm proceeds from a current value, x , to the next value, y , as follows. A value x' is sampled from the distribution with a density of $q(x, x')$. With a probability of $\alpha(x, x')$, $y \leftarrow x'$, else $y \leftarrow x$.

Now, suppose that we have an approximation, $\hat{\pi}(x)$, to $\pi(x)$. The standard delayed-acceptance kernel uses the same proposal, $q(x, y)$, but has an acceptance probability of $\tilde{\alpha}(x, y) = [1 \wedge r_1(x, y)][1 \wedge r_2(x, y)]$, where

$$r_1(x, y) := \frac{\hat{\pi}(y)q(y, x)}{\hat{\pi}(x)q(x, y)} \quad \text{and} \quad r_2(x, y) := \frac{\pi(y)/\hat{\pi}(y)}{\pi(x)/\hat{\pi}(x)}. \quad (3)$$

With $\tilde{\alpha}(x) := \int \tilde{\alpha}(x, y)q(x, y)dy$, the delayed-acceptance (DA) kernel is

$$\tilde{P}(x, dy) := q(x, y)dy \tilde{\alpha}(x, y) + [1 - \tilde{\alpha}(x)]\delta_x(dy). \quad (4)$$

An iteration of the corresponding DA algorithm proceeds from a current value, x , to a next value, y , as follows.

Stage One: A value x' is sampled from the distribution with a density of $q(x, x')$. With a probability of $1 \wedge r_1(x, x')$ the algorithm proceeds to Stage Two, else $y \leftarrow x$.

Stage Two: With a probability of $1 \wedge r_2(x, x')$, $y \leftarrow x'$, else $y \leftarrow x$.

Now, $\tilde{\alpha}(x, y) \leq \alpha(x, y)$, and so $\text{var}(h, \tilde{P}) \geq \text{var}(h, P)$ for each $h \in L^2(\pi)$ Peskun (1973); Tierney (1998). At first glance this might suggest that the DA algorithm is never worthwhile; however for any proposal that is rejected at Stage One there is no need to complete the expensive calculation of $\pi(x')$ that is required at every iteration of the MH algorithm and in Stage Two of the DA algorithm. As mentioned in the Introduction, for a fixed computational time, the decreased average computational cost per iteration, and alterations of any tuning parameters to take advantage of this, can lead to a DA algorithm where the variance of an estimator can be over an order of magnitude smaller than that of the MH algorithm.

Since $\text{var}(h, \tilde{P}) \geq \text{var}(h, P)$, if $\tilde{P}$ is variance bounding then so is P ; however it is feasible that P may be variance bounding while $\tilde{P}$ is not.

2.2 Key Terminology, Equivalences and Implications

The MH and DA kernels are both *reversible* with respect to the target. A kernel P is reversible with respect to a distribution π iff for all $\mathcal{A} \in \mathcal{F}$ and $\mathcal{B} \in \mathcal{F}$, $\int_{\mathcal{A}} \pi(\mathrm{d}x)P(x, \mathcal{B}) = \int_{\mathcal{B}} \pi(\mathrm{d}x)P(x, \mathcal{A})$. This article utilises a number of existing results for reversible Markov chains on the relationship between variance bounding, conductance, spectral gaps and geometric ergodicity. Here we define conductance and spectral gaps and summarise the relationships between the four properties.

Define the Hilbert space $L_0^2(\pi) = \{f : \mathcal{X} \rightarrow \mathbb{R}; \mathbb{E}_{\pi}[f(X)] = 0, \mathbb{E}_{\pi}[f^2(X)] < \infty\}$ with the inner product $\langle f, g \rangle = \int_{\mathcal{X}} f(x)g(x)\pi(\mathrm{d}x)$, and consider P as an operator acting on $L_0^2(\pi)$ according to $(Pf)(x) = \int_{\mathcal{X}} P(x, \mathrm{d}y)f(y)$. If P is reversible then it is a self adjoint operator on $L_0^2(\pi)$ and by the spectral theorem for bounded self-adjoint operators, for each $f \in L_0^2(\pi)$, $\langle f, P^n f \rangle = \int_{-1}^1 \lambda^n H_f(\mathrm{d}\lambda)$ for some positive measure H_f on $[-1, 1]$. Let

$$r_1 := \inf_{f \in L_0^2(\pi)} \frac{\langle f, Pf \rangle}{\langle f, f \rangle} \geq -1 \quad \text{and} \quad r_2 := \sup_{f \in L_0^2(\pi)} \frac{\langle f, Pf \rangle}{\langle f, f \rangle} \leq 1, \quad (5)$$

or, equivalently (e.g. Yosida (1980), p320, Theorem 2), that the smallest closed interval containing the support of H_f for all $f \in L_0^2(\pi)$ is $[r_1, r_2]$. The *spectral gap* of P is $1 - \max(|r_1|, |r_2|)$ (e.g. Geyer (1992); Roberts and Rosenthal (1997)), the *right spectral gap* is $1 - r_2$ and the *left spectral gap* is $1 + r_1$. P is said to have a spectral gap (or a left or right spectral gap) if its spectral gap (or left or right spectral gap) is non-zero.

For any set $\mathcal{A} \in \mathcal{F}$ with $\pi(\mathcal{A}) > 0$ consider the probability of leaving $\mathcal{A}$ at the next iteration given that the stationary chain is currently in $\mathcal{A}$:

$$\kappa(\mathcal{A}) := \frac{1}{\pi(\mathcal{A})} \int_{\mathcal{A}} P(x, \mathcal{A}^c) \pi(\mathrm{d}x).$$

The *conductance*, κ for a Markov kernel P with invariant measure π is then (e.g. Lawler and Sokal (1988)) (see also Jerrum and Sinclair (1988))

$$\kappa := \inf_{\mathcal{A}: 0 < \pi(\mathcal{A}) \leq 1/2} \kappa(\mathcal{A}).$$

For any reversible Markov chain we have the following relationships:

P is variance bounding	$\Leftrightarrow P$ has a right spectral gap	(Thrm 14 of Roberts and Rosentdal (2008))
	$\Leftrightarrow P$ has a positive conductance	(Thrm 2.1 of Lawler and Sokal (1988))
	$\Leftrightarrow P$ has a spectral gap	
	$\Leftrightarrow P$ is geometrically ergodic	(Thrm 2.1 of Roberts and Rosenthal (1997)).

These relationships will be used repeatedly in the sequel without further reference.

2.3 Example Algorithms

To exemplify our theoretical results we will consider four specific, frequently-used MH algorithms.

1. The Metropolis-Hastings independence sampler (MHIS): $q(x, y) = q(y, x)$.

2. The random walk Metropolis (RWM): $q(x, y) = q(y, x) = q(x - y)$; e.g. $q(x, y) = N(y; x, \lambda^2 I)$.
3. The Metropolis-adjusted Langevin algorithm (MALA): $q(x, y) = N(y; x + \frac{1}{2} \lambda^2 \nabla \log \pi, \lambda^2 I)$.
4. The truncated MALA:

$$q(x, y) = N\left(y; x + \frac{1}{2} \lambda^2 R(x), \lambda^2 I\right), \text{ where } R(x) = \frac{D \nabla \log \pi}{D \vee \|\nabla \log \pi\|}, \quad (6)$$

for some $D > 0$.

In Proposals of type 2, 3 and 4, λ is often referred to as the scale parameter of the proposal. The MHIS and RWM have been used since the early days of MCMC (e.g. Tierney (1994)); conditions under which they are geometrically ergodic (and, hence, variance bounding) have been well studied; see, for example, Liu (1996) and Mengersen and Tweedie (1996) for the MHIS and Mengersen and Tweedie (1996), Roberts and Tweedie (1996b) and Jarner and Hansen (2000) for the RWM. Essentially, for the MHIS the proposal, q , must not have lighter tails than the target, and for the RWM the target must have sufficiently smooth and exponentially decreasing tails. The MALA was introduced in Besag (1994) and was analysed in Roberts and Tweedie (1996a), in which the truncated MALA was also introduced. The MALA can be much more efficient than the RWM in moderate to high dimensions. As with the RWM, for geometric ergodicity the MALA requires exponentially decreasing tails, but if the tails decrease too quickly, $\|\nabla \log \pi\|$ grows too quickly and the MALA can fail to be geometrically ergodic. The truncated MALA circumvents this problem.

In Banterle et al. (2019) it is shown that the geometric ergodicity of an RWM algorithm need not be inherited by the resulting DA algorithm.

Example 1 Banterle et al. (2019) Let $\mathcal{X} = \mathbb{R}$ with $\pi(x) \propto e^{-x^2/2}$ and $q(x, y) \propto e^{-(y-x)^2/(2\lambda^2)}$. If $\hat{\pi}(x) \propto e^{-x^2/(2\sigma^2)}$, with $\sigma^2 < 1$ then $\tilde{P}$ is not geometrically ergodic.

The following conditional equivalence (proved in Sect. A.2) is used throughout the sequel. If the parent kernel is geometrically ergodic then the DA kernel must have a left spectral gap, and with this constraint geometric ergodicity and variance bounding are equivalent.

Proposition 1 Let P be a MH kernel targeting π as specified in Eq. 2. Let $\tilde{P}$ be the DA kernel derived from this through the approximation $\hat{\pi}$ as in Eq. 4. If P is geometrically ergodic then

$$\tilde{P} \text{ is geometrically ergodic} \iff \tilde{P} \text{ is variance bounding.}$$

The original random walk Metropolis algorithm on $\pi(x)$ is geometrically ergodic Mengersen and Tweedie (1996), and hence variance bounding, so the DA kernel in Example 1 has not inherited its parent's desirable properties. As a direct corollary of our Theorem 3 (see Sect. 4.2) we find that $\sigma^2 \geq 1$ is exactly the right condition in this case:

Example 2 Let $\mathcal{X} = \mathbb{R}$ with $\pi(x) \propto e^{-x^2/2}$ and $q(x, y) \propto e^{-(y-x)^2/(2\lambda^2)}$. If $\hat{\pi}(x) \propto e^{-x^2/(2\sigma^2)}$, with $\sigma^2 \geq 1$ then $\tilde{P}$ is variance bounding and geometrically ergodic.

Examples 1 and 2 suggest an intuition that problems may arise when $\hat{\pi}(x)$ has lighter tails than $\pi(x)$. As we shall see, this is a part of the story; however, in general, heavier

tails are not sufficient to guarantee inheritance of the variance bounding property, and for a class of algorithms where heavy tails are sufficient, lighter tails can also be sufficient provided they are not too much lighter, in a sense we make precise.

3 Variance Bounding: Inheritance and Equivalence

Throughout this section we use the following generic formulation for two Markov kernels.

Definition 1 Let $P_A(x, dy)$ and $P_B(x, dy)$ be propose-accept-reject Markov kernels both targeting a distribution π , and using, respectively, proposal densities of $q_A(x, y)$ and $q_B(x, y)$ and acceptance probabilities of $\alpha_A(x, y)$ and $\alpha_B(x, y)$.

Theorem 1, below, follows from Lemma 1, which is proved in Sect. A.1.1. It generalises Corollary 12 of Roberts and Rosenthal (2008) to allow for different acceptance probabilities and, more importantly, removes the need for a fixed, uniform minorisation condition. The minorisation needs only hold in a region $\mathcal{D}(x)$ such that under P_A there is “unlikely” to be an accepted proposal in $\mathcal{D}(x)^c$.

Lemma 1 Let $P_A, P_B, q_A, q_B, \alpha_A$ and α_B be as in Definition 1, and let the conductances of P_A and P_B be κ_A and κ_B respectively. If $\kappa_A > 0$ and there is an $\epsilon < \kappa_A$ and a $\delta > 0$ such that for π -almost all $x \in \mathcal{X}$, there is a region $\mathcal{D}(x) \in \mathcal{F}$ such that

$$\int_{\mathcal{D}(x)^c} q_A(x, y) \alpha_A(x, y) dy \leq \epsilon, \quad (7)$$

and

$$y \in \mathcal{D}(x) \Rightarrow q_B(x, y) \alpha_B(x, y) \geq \delta q_A(x, y) \alpha_A(x, y), \quad (8)$$

then $\kappa_B \geq (1 - \epsilon/\kappa_A)\delta\kappa_A$.

If P_A is variance bounding, $\kappa_A > 0$; choose an $\epsilon \in (0, \kappa_A)$ and for each x a corresponding $\mathcal{D}(x)$ so as to satisfy Eqs. 7 and 8 to obtain:

Theorem 1 Let $P_A, P_B, q_A, q_B, \alpha_A$ and α_B be as in Definition 1. If P_A is variance bounding and for any $\epsilon > 0$ there is a $\delta > 0$ such that for π -almost all $x \in \mathcal{X}$ there is a region $\mathcal{D}(x) \in \mathcal{F}$ such that Eqs. 7 and 8 hold, then P_B is also variance bounding.

The relationship between conductance and right spectral gap has recently Lee and Latuszyński (2014); Rudolph and Sprungk (2016) been used in other contexts to bound the behaviour of one Markov kernel in terms of that of another. Lemma 1 itself shows that condition Eq. 7 need only hold for a single $\epsilon < \kappa_A$; however, since in practice κ_A is unlikely to be known, the conditions of Theorem 1 are more practically useful.

From Sect. 4 we apply Theorem 1 to provide sufficient conditions for a delayed-acceptance kernel to inherit variance bounding from its Metropolis-Hastings parent. However, if a DA kernel is variance bounding then so is its parent MH kernel. Thus, the sufficient conditions in Sect. 4 imply an *equivalence* between the two kernels with respect to the variance bounding property. In this section, after two key definitions, we

return, briefly, to this equivalence with regard to the variance bounding property and provide sufficient conditions for equivalence (over potential targets) between Metropolis–Hastings kernels arising from two different proposal densities.

The most natural special case of Eq. 7 in practice is where the kernel is *uniformly local*, which we define as follows:

Definition 2 (Uniformly Local) A proposal is *uniformly local* if, given any $\epsilon > 0$,

$$\exists r < \infty \text{ such that for all } x \in \mathcal{X}, \int_{B(x,r)^c} q(x,y) dy < \epsilon. \quad (9)$$

A propose–accept–reject kernel is defined to be uniformly local when its proposal is uniformly local.

Here and throughout this article, $B(x, r) := \{y \in \mathcal{X} : \|y - x\| < r\}$ is the open ball of radius r centred on x . In our examples, $\|x\|$ indicates the Euclidean norm, although the results are equally valid for other norms such as the Mahalanobis norm.

Control of the ratio $q(y, x)/q(x, y)$ will also be important and so we define the following.

Definition 3 For any proposal density $q(x, y)$,

$$\Delta(x, y; q) := \log q(y, x) - \log q(x, y).$$

Clearly, the RWM is a uniformly local kernel; moreover $\Delta(x, y; q_{\text{RWM}}) = 0$. In contrast, on any target with unbounded support, the MHIS cannot be uniformly local; as we shall see, the behaviour of Δ is then irrelevant. For the MALA and the truncated MALA we have:

Proposition 2

(A) Let $q(x, y)$ be the proposal for the truncated MALA in Eq. 6 or for the MALA on a target where $\text{ess sup}_x \|\nabla \log \pi(x)\| = D < \infty$. Then

- (i) For all x , $\mathbb{P}_q(\|Y - x\| > r) \rightarrow 0$ uniformly in x as $r \rightarrow \infty$, so q is uniformly local, as defined in Eq. 9.
- (ii) $|\Delta(x, y; q)| \leq h(\|y - x\|) := D\|y - x\| + \lambda^2 D^2/8$.

(B) The proposal, $q(x, y)$, for the MALA on a target where $\text{ess sup} \|\nabla \log \pi(x)\| = \infty$ is not uniformly local.

The applicability of Lemma 1 and Theorem 1 ranges beyond delayed-acceptance kernels. Here we supply sufficient conditions for an *equivalence* between Metropolis–Hastings proposals.

Theorem 2 Let $P_A, P_B, q_A, q_B, \alpha_A$ and α_B be as in Definition 1 except that $q_A(x, y)$ and $q_B(x, y)$ are uniformly local proposal kernels, with $\log q_A(x, y) - \log q_B(x, y)$ a continuous function from $\mathbb{R}^{2d}$ to $\mathbb{R}$. If, for π -almost all x and for some function $h : [0, \infty) \rightarrow [0, \infty)$ with $h(r) < \infty$ for all $r < \infty$,

$$|\Delta(x, y; q_B) - \Delta(x, y; q_A)| \leq h(\|y - x\|), \quad (10)$$

then P_A is variance bounding if and only if P_B is variance bounding.

Thus, for example, any two random-walk Metropolis algorithms with Gaussian jumps are equivalent, in that if, on a particular target, one is variance bounding then so is the other. When restricted to targets with a continuous gradient this equivalence extends to truncated MALA algorithms. The continuity requirement on $\log q_A - \log q_B$ rules out, for example, an equivalence between a Gaussian random walk and a random walk where the proposal has bounded support; indeed, the latter may not even be ergodic if the target has gaps in its support.

4 Application to Delayed-Acceptance Kernels

4.1 Key Definitions and Properties

For uniformly local kernels we will describe two general sets of sufficient conditions for Eq. 8 to hold. The first is based upon the fact that the acceptance probability for $\tilde{P}$ can be written as

$$\begin{aligned} \tilde{\alpha}(x, y) &= [1 \wedge r_1(x, y)] \left[1 \wedge \frac{r(x, y)}{r_1(x, y)} \right] \geq [1 \wedge r_1(x, y)] \left[1 \wedge \frac{1}{r_1(x, y)} \right] [1 \wedge r(x, y)], \\ \text{or } \tilde{\alpha}(x, y) &= \left[1 \wedge \frac{r(x, y)}{r_2(x, y)} \right] [1 \wedge r_2(x, y)] \geq [1 \wedge r(x, y)] \left[1 \wedge \frac{1}{r_2(x, y)} \right] [1 \wedge r_2(x, y)], \end{aligned}$$

where r_1 and r_2 are as defined in Eq. 3. So, if $|\log r_1(x, y)| \leq m$ or $|\log r_2(x, y)| \leq m$ then $\tilde{\alpha}(x, y) \geq e^{-m} \alpha(x, y)$. The quantity $|\log r_2(x, y)| = |[\log \hat{\pi}(y) - \log \pi(y)] - [\log \hat{\pi}(x) - \log \pi(x)]|$ measures the discrepancy between the error in the approximation at the proposed value and the error in the approximation at the current value. We name this intuitive quantity, the *log-error discrepancy*. The quantity $\log r_1$ is less natural since it relates $\hat{\pi}(x)$, $\hat{\pi}(y)$ and $q(x, y)$.

The second set of conditions is based upon the fact that if either $r_1(x, y) \leq 1$ and $r_2(x, y) \leq 1$ or if $r_1(x, y) \geq 1$ and $r_2(x, y) \geq 1$ then $\tilde{\alpha}(x, y) = \alpha(x, y)$, whatever the log-error discrepancy.

These considerations lead to the natural definitions of a ‘potential problem’ set, $\mathcal{M}_m(x)$, and a ‘no problem’ set $\mathcal{C}(x)$, as follows:

$$\mathcal{M}_m(x) := \{y \in \mathcal{X} : \min |\log r_1(x, y)|, |\log r_2(x, y)| > m\}, \quad (11)$$

$$\mathcal{C}(x) := \{y \in \mathcal{X} : \text{sign}[\log r_1(x, y)] \text{sign}[\log r_2(x, y)] \geq 0\}. \quad (12)$$

Theorem 1 then leads directly to the following.

Corollary 1 *Let P be the Metropolis-Hastings kernel given in Eq. 2 and let $\tilde{P}$ be the corresponding delayed-acceptance kernel given in Eq. 4. Suppose that for all $\epsilon > 0$ there is an $m < \infty$ such that for π -almost all x there exists a set $\mathcal{D}(x) \subseteq \mathcal{X}$ such that*

$$\mathcal{M}_m(x) \cap \mathcal{D}(x) \subseteq \mathcal{C}(x), \quad (13)$$

and

$$\int_{\mathcal{D}(x)^c} q(x, y) dy \leq \epsilon. \quad (14)$$

Subject to these conditions, if P is variance bounding then so is $\tilde{P}$.

When $\hat{\pi}$ has heavier tails than π then for large x , the set $\mathcal{C}(x)$ can play an important role in the inheritance of the variance bounding property. In a dimension $d > 1$, there are numerous possible definitions of ‘heavier tails’. The following is precisely that required for our purposes:

Definition 4 (heavy tails) An approximate density $\hat{\pi}$ is said to have *heavy tails* with respect to a density π if

$$\exists r_* > 0 \text{ such that if } \|x\| > r_* \text{ and } \|y\| > r_* \text{ then } \hat{\pi}(x) \leq \hat{\pi}(y) \Rightarrow \frac{\hat{\pi}(x)}{\pi(x)} \geq \frac{\hat{\pi}(y)}{\pi(y)}. \quad (15)$$

Intuitively, the left hand side is true when x is ‘further from the centre’ (according to $\hat{\pi}$) than y , and the implication is that the further out a point, the larger $\hat{\pi}$ is compared with π .

For uniformly local kernels we show (Corollary 3) that it is sufficient that either the log error discrepancy should satisfy a growth condition that is uniform in $\|y - x\|$, or (Theorem 3) that the tails of the approximation should be heavier than those of the target and that $|\Delta(x, y; q)|$ should satisfy a growth condition that is uniform in $\|y - x\|$.

For all kernels, boundedness of the error $\hat{\pi}(x)/\pi(x)$ away from 0 and ∞ will ensure the required inheritance (Corollary 2). This is a very strong condition, but we exhibit MHIS and MALA algorithms where the weaker conditions, that are sufficient for a uniformly local kernel, are satisfied, but the DA kernel is not variance bounding even though the MH kernel is.

4.2 DA Kernels with the Same Proposal Distribution as the Parent

Suppose that for all $x \in \mathcal{X}$, $\gamma_{lo} \leq \hat{\pi}(x)/\pi(x) \leq \gamma_{hi}$, then $|\log r_2(x, y)| \leq \log(\gamma_{hi}/\gamma_{lo})$, so applying Corollary 1 with $\mathcal{D}(x) = \mathcal{X}$ and $m = \log \gamma_{hi} - \log \gamma_{lo}$ leads to:

Corollary 2 Let P and $\tilde{P}$ be as described in Corollary 1. If there exist $\gamma_{lo} > 0$ and $\gamma_{hi} < \infty$ such that $\gamma_{lo} \leq \hat{\pi}(x)/\pi(x) \leq \gamma_{hi}$, and if P is variance bounding then so is $\tilde{P}$.

A more direct proof of Corollary 2 is possible using Dirichlet forms. However, Corollary 1 comes into its own when the error discrepancy is unbounded.

We first provide a cautionary example which shows that once the errors are unbounded the delayed-acceptance kernel need not inherit the variance bounding property from the Metropolis-Hastings kernel even if the growth of the log error discrepancy is uniformly bounded or if $\hat{\pi}$ has heavier tails than π .

Example 3 Let $\mathcal{X} = \mathbb{R}$, let P be an MHIS with $q(x, y) = q(y) = \pi(y) = e^{-y} \mathbb{1}(y > 0)$, and let $\tilde{P}$ be the corresponding delayed-acceptance kernel Eq. 4, with $\hat{\pi}(y) = ke^{-ky} \mathbb{1}(y > 0)$ with

$k > 0$ and $k \neq 1$. P is geometrically ergodic, but $\tilde{P}$ is neither geometrically ergodic nor variance bounding.

The problem with the algorithm in Example 3 is that for some x values the proposal, y , is very likely to be a long way from x and yet $y \notin \mathcal{C}(x)$. Our definition of a *uniformly local* proposal, Eq. 9, provides uniform control on the probability that $\|y - x\|$ is large. Since this is only strictly necessary for $y \notin \mathcal{C}(x)$, Eq. 9 is stronger than necessary, but it is much easier to check.

Our first sufficient condition for uniformly local kernels insists on uniformly bounded growth in the log-error discrepancy except when $\tilde{\alpha}(x, y) = \alpha(x, y)$. For π -almost all x and for some function $h : [0, \infty) \rightarrow [0, \infty)$ with $h(r) < \infty$ for all $r < \infty$,

$$\{y \in \mathcal{X} : |\log r_2(x, y)| > h(\|y - x\|)\} \subseteq \mathcal{C}(x). \quad (16)$$

If a proposal is uniformly local, given $\epsilon > 0$ find $r(\epsilon)$ according to Eq. 9. Then Eq. 16 implies that for $y \in B(x, r)$, $\mathcal{M}_{h(r)} \subseteq \mathcal{C}(x)$. Applying Corollary 1 with $\mathcal{D}(x) = B(x, r)$ leads to the following.

Corollary 3 *Let P and $\tilde{P}$ be as described in Corollary 1. In addition let $q(x, y)$ be a uniformly local proposal as in Eq. 9, and let the error discrepancy satisfy Eq. 16. If P is variance bounding then so is $\tilde{P}$.*

Because most of the mass from the proposal, y , is not too far away from the current value, x , the discrepancy between the error at x and the error at y remains manageable provided the discrepancy grows in a manner that is controlled uniformly across the state-space. Since the random walk Metropolis on an exponential target density is geometrically ergodic Mengersen and Tweedie (1996) we may apply Corollary 3 with $h(r) = |k - 1|r$, and then Proposition 1, to obtain the following contrast to Example 3, and showing that the variance bounding property can be inherited even when the approximation has *lighter* tails than the target.

Example 4 Let $\mathcal{X} = \mathbb{R}$ and let P be a RWM algorithm on $\pi(x) = e^{-x} \mathbb{1}(x > 0)$ using $q(x, y) \propto e^{-(y-x)^2/(2\lambda^2)}$. For any $k > 0$, let $\tilde{P}$ be the corresponding delayed-acceptance RWM algorithm using a surrogate of $\hat{\pi}(x) = ke^{-kx} \mathbb{1}(x > 0)$. $\tilde{P}$ is variance bounding and geometrically ergodic.

As yet, the set $\mathcal{C}(x)$ has not played a part in any of our examples. It is precisely this set that allows a delayed-acceptance random walk Metropolis kernel to inherit the variance bounding property from its parent even when the error discrepancy is not controlled uniformly, provided $\hat{\pi}$ has tails that are heavier than those of π . For general MH algorithms an additional control on the behaviour of q is enough to guarantee inheritance of the variance bounding property.

Theorem 3 *Let P be the Metropolis-Hastings kernel given in Eq. 2 and let $\tilde{P}$ be the corresponding delayed-acceptance kernel given in Eq. 4. Further, let $q(x, y)$ be a uniformly local proposal in the sense of Eq. 9, let π and $\hat{\pi}$ be continuous, and let $\hat{\pi}$ have heavier tails than π in the sense of Eq. 15. Suppose that, in addition, for any $\mathcal{D}(x)$ required by Eqs. 13 and 14 there exists a function $h : [0, \infty) \rightarrow [0, \infty)$ with $h(r) < \infty$ for all $r < \infty$, such that for π -almost all x*

$$\{y \in \mathcal{D}(x) : |\Delta(x, y; q)| > h(|y - x|)\} \subseteq \mathcal{C}(x). \quad (17)$$

Subject to these conditions, if P is variance bounding then so is $\tilde{P}$.

We now consider the delayed-acceptance versions of the random walk Metropolis, the truncated MALA, and the MALA. Before doing this we provide the details of a property that was anticipated in Roberts and Tweedie (1996a).

Proposition 3 Let P_{RWM} be a random walk Metropolis kernel using $q(x, y) \propto e^{-\frac{1}{2\lambda^2}\|y-x\|^2}$ and targeting a density $\pi(x)$. Let P be a Metropolis-Hastings kernel on π of the form $q(x, y) \propto e^{-\frac{1}{2}\lambda^2\|y-x-v(x)\|^2}$, where $\pi - \text{ess sup}_x \|v(x)\| < \infty$. P_{RWM} is variance bounding if and only if P is variance bounding.

Proposition 3 clearly applies to a truncated MALA kernel on $\pi(x)$ using q as in Eq. 6. It, together with each of our subsequent results for the truncated MALA, also applies to a MALA kernel on a target where $\pi - \text{ess sup}_x \|\nabla \log \pi(x)\| = D < \infty$; in practice, however, the useful set of such kernels is limited to targets with exponentially decaying tails, since MALA is not geometrically ergodic on targets with heavier tails Roberts and Tweedie (1996a).

Given Proposition 2 and its prelude, a direct application of Theorem 3 then leads to the following.

Example 5 Let P_{RWM} and P_{TMALA} be, respectively, a random walk Metropolis kernel and a truncated MALA kernel on the differentiable density, $\pi(x)$. Let $\tilde{P}_{\text{RWM}}$ and $\tilde{P}_{\text{TMALA}}$ be the corresponding delayed-acceptance kernels, created as in Eq. 4 through the continuous density, $\hat{\pi}(x)$. Suppose also that $\hat{\pi}$ has heavier tails than π in the sense of Eq. 15. Subject to these conditions, if P_{RWM} is variance bounding then so is $\tilde{P}_{\text{RWM}}$, and if P_{TMALA} is variance bounding then so is $\tilde{P}_{\text{TMALA}}$.

The MALA is geometrically ergodic when applied to one-dimensional targets of the form $\pi(x) \propto e^{-|x|^\beta}$ for $\beta \in [1, 2)$ Roberts and Tweedie (1996a); when $\beta = 2$ geometric ergodicity occurs provided λ is sufficiently small, and for $\beta > 2$ the MALA is not geometrically ergodic. Even when $\beta > 1$, however, Theorem 3 does not apply because the proposal is not uniformly local.

Example 6 Let $\mathcal{X} = \mathbb{R}$ and let P be a MALA algorithm on $\pi(x) \propto e^{-x^\beta} 1(x > 0)$ with $1 \leq \beta < 2$. Let $\hat{\pi}(x) \propto e^{-x^\gamma} 1(x > 0)$ and let $\tilde{P}$ be the corresponding delayed-acceptance MALA kernel Eq. 4 (i.e. using a proposal of $Y = x + \frac{1}{2}\lambda^2 \nabla \log \pi(x) + \lambda Z$, where $Z \sim N(0, 1)$). $\tilde{P}$ is neither geometrically ergodic nor variance bounding, except when $\gamma = \beta$.

The contrast between the truncated MALA and the MALA in Examples 5 and 6 highlights the importance of a uniformly local proposal. In practice, however, if $\pi(x)$ is computationally expensive to evaluate then, typically, $\nabla \log \pi(x)$ will also be expensive to evaluate and it might seem more reasonable to base the proposal for delayed-acceptance MALA and delayed-acceptance truncated MALA on $\nabla \log \hat{\pi}(x)$.

4.3 Kernels Where the Proposal is Based Upon $\hat{\pi}$

On some occasions, the proposal $q(x, y)$ is a function of the posterior, $\pi(x)$, and on such occasions it may be expedient for the delayed-acceptance algorithm to use a proposal $\hat{q}(x, y)$, which is based upon $\hat{\pi}(x)$. The acceptance rate is $\tilde{\alpha}_b(x, y) = [1 \wedge r_{1b}(x, y)][1 \wedge r_2(x, y)]$, where

$$r_{1b}(x, y) := \frac{\hat{\pi}(y)\hat{q}(y, x)}{\hat{\pi}(x)\hat{q}(x, y)}.$$

With $\tilde{\alpha}_b(x) := \mathbb{E}_q[\tilde{\alpha}_b(x, Y)]$, the corresponding delayed acceptance kernel is

$$\tilde{P}_b(x, dy) := \hat{q}(x, y)dy \tilde{\alpha}_b(x, y) + [1 - \tilde{\alpha}_b(x)]\delta_x(dy). \quad (18)$$

Let $r_{\text{hyp}}(x, y) := \pi(y)\hat{q}(y, x)/[\pi(x)\hat{q}(x, y)]$, $\alpha_{\text{hyp}}(x, y) = 1 \wedge r_{\text{hyp}}(x, y)$, and, with $\bar{\alpha}_{\text{hyp}}(x) = \mathbb{E}_{\hat{q}}[\alpha_{\text{hyp}}(x, Y)]$, consider the hypothetical Metropolis-Hastings kernel:

$$P_{\text{hyp}}(x, dy) := \hat{q}(x, y)dy \alpha_{\text{hyp}}(x, y) + [1 - \bar{\alpha}_{\text{hyp}}(x)]\delta_x(dy). \quad (19)$$

Now, $\tilde{\alpha}_b(x, y) \leq \alpha_{\text{hyp}}(x, y)$, so if P_{hyp} is not variance bounding then $\tilde{P}_b$ is not variance bounding either. There is an exact correspondence between P from the previous section, and P_{hyp} , and it is natural to consider inheritance of geometric ergodicity from P_{hyp} exactly as in the previous section we considered inheritance from P . The theoretical results are analogous and will not be restated; moreover, the theoretical properties of kernels of the form P_{hyp} are less well investigated. Instead we illustrate inheritance of variance bounding (or its lack) through two examples.

Example 7 Let P_{TMALA} be a truncated MALA kernel on the differentiable density, $\pi(x)$. Let $\tilde{P}_{\text{TMALab}}$ be the corresponding delayed-acceptance kernel, created as in Eq. 18 through the differentiable density $\hat{\pi}(x)$. $\tilde{P}_{\text{TMALab}}$ inherits the variance bounding property from P_{TMALA} if either of the following conditions holds. (i) There is uniformly bounded growth in the log error discrepancy, in the sense of Eq. 16, or (ii) $\hat{\pi}$ has heavier tails than π in the sense of Eq. 15.

Our penultimate example suggests that a delayed-acceptance MALA based upon an approximation that has heavier (though not too much heavier) tails is a reasonable choice.

Example 8 Let $\mathcal{X} = \mathbb{R}$ and let P be a MALA algorithm on $\pi(x) \propto e^{-x^\beta} \mathbb{1}(x > 0)$ with $1 \leq \beta < 2$. Let $\hat{\pi}(x) \propto e^{-x^\gamma} \mathbb{1}(x > 0)$ and let $\tilde{P}$ be the corresponding delayed-acceptance MALA kernel created as in Eq. 18 through the differentiable density $\hat{\pi}(x)$. $\tilde{P}$ is variance bounding $\iff \tilde{P}$ is geometrically ergodic $\iff 1 \leq \gamma \leq \beta$.

We summarise the consequences of Examples 5 to 8 for $\pi(x) \propto e^{-x^\beta} \mathbb{1}(x > 0)$ and $\hat{\pi}(x) \propto e^{-x^\gamma} \mathbb{1}(x > 0)$ in Table 1, filling in the two blanks with Example 9 below. The table displays the results in terms of variance bounding, which is equivalent to geometric ergodicity in all these cases by Proposition 1.

Example 9 Let $\mathcal{X} = \mathbb{R}$ and let P be a RWM or truncated MALA algorithm on $\pi(x) \propto e^{-x^\beta} \mathbb{1}(x > 0)$ with $1 \leq \beta < 2$. Let $\hat{\pi}(x) \propto e^{-x^\gamma} \mathbb{1}(x > 0)$ and let $\tilde{P}$ be the corresponding delayed-acceptance RWM or truncated MALA kernel created either from Eq. 4 or Eq. 18

Table 1 Whether or not the DA algorithm for $\pi(x) \propto e^{-x^\beta}$ ($x > 0$) using $\hat{\pi}(x) \propto e^{-x^\gamma}$ ($x > 0$) is variance bounding as a function of γ and β and the specific DA algorithm. The final two columns indicate that $\hat{\pi}$ rather than π is used to create the proposal

Algorithm	MALA	RWM/TMALA	MALA ($\nabla \log \hat{\pi}$)	TMALA ($\nabla \log \hat{\pi}$)
$1 \leq \gamma < \beta \leq 2$	× Ex 6	✓ Ex 5	✓ Ex 8	✓ Ex 7
$1 \leq \beta < \gamma$	× Ex 6	× Ex 9	× Ex 8	× Ex 9

through the differentiable density $\hat{\pi}(x)$. If $1 < \beta < \gamma$, $\tilde{P}$ is neither geometrically ergodic nor variance bounding.

5 Numerical Demonstrations

The theoretical results from Sect. 4 were made more concrete through Examples 1 to 9. In this section we investigate the numerical performance of delayed acceptance algorithms in examples similar to those used in earlier sections. The specific targets in the earlier Examples were chosen to demonstrate particular points as simply as possible; here we deliberately investigate a broader class of targets, the exponential family class (e.g. Roberts and Tweedie (1996a); Livingstone et al. (2019)):

$$\pi(x) \propto \exp(-||x||^\beta) \quad \text{and} \quad \hat{\pi}(x) \propto \exp(-||x||^\gamma / \kappa^\gamma). \quad (20)$$

The parameters β and γ in (19) govern the lightness of the tails in the target and the approximation to it respectively, and allow us to vary these separately.

A lack of variance bounding can be seen in terms of the chain struggling to leave a certain region, which typically has a low probability under π . In practice, this lack of variance bounding (or a lack of geometric ergodicity) can manifest in two ways.

1. When a sensible starting value is not known, a starting value with poor properties may be chosen unwittingly and the algorithm may struggle to move from this initial point or region of the space.
2. Even when started from a reasonable value, over the course of a sufficiently long run the algorithm will visit this “danger region” and then struggle to leave.

For the target Eq. 20, the “danger region” corresponds to the tails of π .

Our experiments deliberately start the algorithm in the tails of π and then measure the number of iterations to reach the centre of the distribution. To make “reaching the centre” concrete, we find the number of iterations until $||x||$ is less than its median value under π . To decide where in the tails we start, we set $||x||$ to its $1 - p_0$ quantile under π , for $p_0 \in \{10^{-1}, 10^{-2}, \dots, 10^{-6}\}$ in Scenarios (i) and (ii), and $p_0 \in \{10^{-4}, 10^{-8}, \dots, 10^{-24}\}$ in Scenarios (iii) and (iv); we start the algorithm from a uniformly random point on the surface of that hypersphere. In practical MCMC, many runs are of $\mathcal{O}(10^6)$ iterations, so it is not unreasonable that issues which are detected for $p_0 \geq 10^{-8}$ might occur in practice even when the algorithm is started from a sensible value. We work in dimension $d = 5$ and repeat each experiment 20 times, except for scenario (iii) where we repeat 10 times to avoid excessive clutter.

We consider four specific scenarios, and so as to bound the amount of computing time, in each scenario we set a maximum number of iterations for which the algorithm should be run. In all scenarios the time until convergence increases with the starting quantile, whether or not the algorithm is variance bounding, for the most part simply because the algorithm is starting further from the main mass of the target. However, when the disparity between algorithms grows towards an order of magnitude, this suggests danger.

For the DARWM and DAMALA, the scaling parameter, λ , was chosen so that for the RWM or MALA itself, the acceptance rate was a little larger than the theoretical optimum values of approximately 23% and 57% respectively. DATMALA used the same scaling as DAMALA and a truncation value such that when TMALA explored the true posterior, fewer than 4% of the gradients were truncated.

Scenario i ($\beta = \gamma = 2$, $\kappa = 1/2$ and $\kappa = 2$). The results appear in the top-left of Figure 1 and demonstrate the undesirable behaviour when the target and the approximation are both Gaussian but the approximation has lighter tails than the target (see Example 1), and the reasonable behaviour when the approximation's tails are less tight than the target's (Example 2).

Scenario ii ($\beta = \gamma = 1$, $\kappa = 1/2$ and $\kappa = 2$). The results, in the top-right of Figure 1, demonstrate that, in alignment with Example 3, the worst behaviour by some margin is exhibited by the only non-variance bounding algorithm: the independence sampler where $\hat{\pi}$ uses a smaller scaling than π has. In particular, aligning with Example 4, the DARWM that uses the same $\hat{\pi}$ as the poor independence sampler performs only marginally worse than the DARWM which uses the notionally 'safer' $\hat{\pi}$.

Scenario iii ($\beta = 1.5$, $\kappa = 1$, $\gamma = 1.2$ and $\gamma = 1.8$). This corresponds to Examples 5, 6 and 9 and is consistent with the DARWM and DATMALA, but not DAMALA, being variance bounding when $1 < \gamma < \beta$, and none being variance bounding when $\gamma > \beta$.

Scenario iv ($\beta = 1.5$, $\kappa = 1$, $\gamma = 1.2$ and $\gamma = 1.8$, proposal uses $\nabla \log \hat{\pi}$) and suggests that as with Examples 7 and 8, DAMALA and DATMALA are both variance bounding when $\gamma < \beta$, and following Example 9, neither is variance bounding when $\gamma > \beta$.

In scenarios (i), (iii) and (iv) the target itself has lighter-than-exponential tails, so even though the x-axis is linear in $\log p$ it is sublinear in the magnitude of the initial value, $\|x_0\|$. Hence, issues with the algorithms might be expected to appear more slowly as $-\log p_0$ increases than they do with scenario (ii). Whilst exceptionally poor behaviour is unlikely to be seen, therefore, during a typical run that has been started from the main posterior mass, it could easily occur as a result of a poor starting value.

6 Discussion

Delayed acceptance Metropolis-Hastings algorithms are popular when the posterior is computationally intensive to evaluate yet a cheap approximation is available. Approximations can arise through many mechanisms, including the coarsening of a numerical-integration grid, subsampling from big data, Gaussian process approximation and nearest neighbour averaging. To date, with the exception of Franks and Vihola (2020) and a note in Banterle et al. (2019), little consideration has been given to the properties of the resulting algorithm and, in particular as to whether the delayed-acceptance algorithm might inherit good properties, such as variance bounding, from its parent Metropolis-Hastings algorithm. From the MCMC output, one might reasonably hope to be able to estimate any quantity with a finite variance under π and be confident that the Monte Carlo error would reduce in inverse

Fig. 1 Plots of $\log_{10}$ convergence time against (jittered) $-\log_{10} p_0$ for scenarios (i) top left, (ii) top right, (iii) bottom left and (iv) bottom right. In plots (i)–(iii), a dark blue + corresponds to DARWM where the target has a ‘good’ parameter value (scaling in (i) and (ii), and power in (iii)) and a red x to DARWM with a ‘poor’ parameter value. In plot (ii) a light blue o corresponds to DAIS (DA independence sampler) with $\kappa = 2$ and a magenta $\triangle$ to DAIS with $\kappa = 1/2$. In plots (iii) and (iv), the light blue o corresponds to truncated DATMALA with $\gamma < \beta$, and magenta $\triangle$ to DATMALA with $\gamma > \beta$, whilst a green • corresponds to DAMALA with $\gamma < \beta$, and black $\blacktriangle$ to DAMALA with $\gamma > \beta$

proportion to the square-root of the run length; however, if the algorithm is not variance bounding then this may not be the case.

We have investigated the inheritance of the variance bounding property and provided sufficient conditions for it to occur. A general rule of thumb for algorithms with *uniformly local* (see Definition 2) proposals, such as the random walk Metropolis and the truncated MALA, is that the approximation should have *heavier tails* (see Definition 4) than the target; however, this is not always necessary (see Example 4). The MALA algorithm does not enjoy the same good properties as the truncated MALA and, in particular, does not necessarily inherit variance bounding even when the approximation does have heavier tails than the target (see Example 6).

A note of caution is also in order: variance bounding (and/or geometric ergodicity) are helpful properties as, in particular, they guarantee the existence of a usual central limit

theorem for ergodic averages. However, whilst non-zero, the conductance of a kernel could be exceedingly small (or the geometric rate of convergence exceptionally close to one) so that the algorithm might not be useful in practice. Thus, whilst we recommend following the advice in this article when choosing the approximation so as to reduce the chance of false confidence in the resulting Monte Carlo estimates, one should also continue to check other diagnostics, such as trace plots, and to vary any tuning parameters to optimise performance.

Proofs of Results

Proofs of Results in Sect. 3

Proof of Lemma 1

Since $\epsilon < \kappa_A$ we may define $\beta \in (0, 1)$ such that $\epsilon = (1 - \beta)\kappa_A$. For any $\mathcal{A} \in \mathcal{F}$,

$$\begin{aligned} \int_{\mathcal{A}^c} q_B(x, y) \alpha_B(x, y) dy &\geq \int_{\mathcal{A}^c \cap \mathcal{D}(x)} q_B(x, y) \alpha_B(x, y) dy \\ &\geq \delta \int_{\mathcal{A}^c \cap \mathcal{D}(x)} q_A(x, y) \alpha_A(x, y) dy \text{ by Eq. 8.} \\ &= \delta \left[\int_{\mathcal{A}^c} q_A(x, y) \alpha_A(x, y) dy - \int_{\mathcal{A}^c \cap \mathcal{D}(x)^c} q_A(x, y) \alpha_A(x, y) dy \right] \\ &\geq \delta \left[\int_{\mathcal{A}^c} q_A(x, y) \alpha_A(x, y) dy - \int_{\mathcal{D}(x)^c} q_A(x, y) \alpha_A(x, y) dy \right] \\ &\geq \delta \left[\int_{\mathcal{A}^c} q_A(x, y) \alpha_A(x, y) dy - (1 - \beta)\kappa_A \right] \text{ by Eq. 7.} \end{aligned}$$

Integrating both sides over $x \in \mathcal{A}$ with respect to π gives

$$\begin{aligned} \pi(\mathcal{A})\kappa_B(\mathcal{A}) &= \int_{x \in \mathcal{A}} \int_{y \in \mathcal{A}^c} \pi(dx) q_B(x, y) \alpha_B(x, y) dy \\ &\geq \delta \left[\int_{x \in \mathcal{A}} \int_{y \in \mathcal{A}^c} \pi(dx) q_A(x, y) \alpha_A(x, y) dy - (1 - \beta)\kappa_A \int_{x \in \mathcal{A}} \pi(dx) \right] \\ &= \delta [\pi(\mathcal{A})\kappa_A(\mathcal{A}) - (1 - \beta)\pi(\mathcal{A})\kappa_A] \\ &\geq \beta\delta\pi(\mathcal{A})\kappa_A. \end{aligned}$$

The result follows since only sets with $\pi(\mathcal{A}) > 0$ are relevant. $\square$

Proof of Proposition 2

Let $v(x) := \frac{1}{2}\lambda^2 \nabla \log \pi(x)$.

- (A) (i) $\|Y - x\| = \|v(x) + \lambda Z\| \leq \frac{1}{2}\lambda^2 D + \lambda\|Z\|$, where Z is a vector with iid $N(0, 1)$ components. So $\mathbb{P}(\|Y - x\| > r) \leq \mathbb{P}\left(\lambda\|Z\| + \frac{1}{2}\lambda^2 D > r\right) \rightarrow 0$ as $r \rightarrow \infty$.
(ii) Algebra shows that

$$\Delta(x, y; q) = \frac{1}{2\lambda^2} \{2(x - y) \cdot [v(x) + v(y)] + \|v(x)\|^2 - \|v(y)\|^2\}. \quad (21)$$

Since $\|v(x)\| \leq \lambda^2 D/2$, $|\Delta(x, y; q)| \leq h(|y - x|) := D|y - x| + \lambda^2 D^2/8$, as required.

(B) Let $Z_* = Z \cdot \nabla \log \pi(x) / \|\nabla \log \pi\| \sim N(0, 1)$. For any $D > 0$ we may find $\mathcal{A}_D \in \mathcal{F}$ with $\pi(\mathcal{A}_D) > 0$ and $\|\nabla \log \pi(x)\| \geq D$ for all $x \in \mathcal{A}_D$. Hence, for $x \in \mathcal{A}_D$,

$$\|Y - x\| = \left\| \frac{1}{2} \lambda^2 \nabla \log \pi(x) + \lambda Z \right\| \geq \frac{1}{2} \lambda^2 \|\nabla \log \pi(x)\| - \lambda |Z_*| \geq \frac{1}{2} \lambda^2 D - \lambda |Z_*|.$$

So for any $r > 0$, $\mathbb{P}(\|Y - x\| \geq r) \geq \mathbb{P}(|Z_*| \leq \frac{1}{2} \lambda D - r/\lambda)$, which can be made as close to 1 as desired by taking D to be sufficiently large. $\square$

Proof of Theorem 2

Since $\kappa_A > 0$ and q_A is uniformly local, we may take $\mathcal{D}(x) = \overline{B}(x, r)$, the closure of $B(x, r)$, where r is chosen so as to satisfy Eq. 9 for π -almost all x , and with $\epsilon = \kappa_A/2$.

Next, let $t = 0 \wedge (c + b) - 0 \wedge (c + a)$. Since $0 \wedge (c + a)$ is upper bounded by both 0 and $c + a$, if $0 < c + b$ then $t \geq 0$, and if $c + b < 0$ then $t \geq b - a$. Thus $t \geq 0 \wedge (b - a)$. Hence for $y \in \mathcal{D}(x)$,

$$\begin{aligned} \log \frac{q_B(x, y) \alpha_B(x, y)}{q_A(x, y) \alpha_A(x, y)} &= \log q_B(x, y) - \log q_A(x, y) \\ &\quad + 0 \wedge [\log(\pi(y)/\pi(x)) + \Delta(x, y; q_B)] - 0 \wedge [\log(\pi(y)/\pi(x)) + \Delta(x, y; q_A)] \\ &\geq \log q_B(x, y) - \log q_A(x, y) + 0 \wedge [\log \Delta(x, y; q_B) - \log \Delta(x, y; q_A)] \\ &\geq \log q_B(x, y) - \log q_A(x, y) - h(r). \end{aligned}$$

The first term is bounded on $\mathcal{D}(x)$ since $\log q_B(x, y) - \log q_A(x, y)$ is continuous, and so Eq. 8 holds and we may apply Lemma 1. Repeat with $A \leftrightarrow B$.

Proofs of Results in Sect. 2

Proof of Proposition 1

Since P is geometrically ergodic, it must have a left spectral gap. From Eq. 5,

$$\text{Gap}_L(P) = 1 + \inf_{f \in L_0^2(\pi), \langle f, f \rangle = 1} \int \pi(dx) P(x, dy) f(x) f(y) = 2 - \sup_{f \in L_0^2(\pi), \langle f, f \rangle = 1} \mathcal{E}_P(f),$$

where the Dirichlet form for the functional f of the Markov chain is

$$\mathcal{E}_P(f) = \frac{1}{2} \int \pi(dx) P(x, dy) [f(x) - f(y)]^2 = \frac{1}{2} \int \pi(dx) q(x, y) \alpha(x, y) [f(x) - f(y)]^2 dy,$$

for a propose-accept-reject chain.

Since $\tilde{\alpha}(x, y) \leq \alpha(x, y)$, $\mathcal{E}_{\tilde{P}}(f) \leq \mathcal{E}_P(f)$. So if P is geometrically ergodic, then $\text{Gap}_L(\tilde{P}) \geq \text{Gap}_L(P) > 0$. The result follows as all geometrically ergodic kernels are variance bounding and the only kernels which are variance bounding but not geometrically ergodic have

no left-spectral gap, yet we have just shown that $\tilde{P}$ must have a left spectral gap because P is geometrically ergodic. $\square$

Proof of Example 3

Since $\alpha(x, y) = 1$, the MH algorithm produces iid samples from π and so it is geometrically ergodic (with a spectral gap of 1) and, hence, variance bounding. For the DA algorithm,

$$\tilde{\alpha}(x, y) = [1 \wedge e^{(k-1)(x-y)}] [1 \wedge e^{(k-1)(y-x)}].$$

For any $r > 0$ let

$$\mathcal{A}_r := [2r, \infty), \mathcal{B}_r := (0, r), \text{ and } \mathcal{C}_r := [r, 2r).$$

For $(x, y) \in \mathcal{A}_r \times \mathcal{B}_r$, $\tilde{\alpha}(x, y) \leq e^{-|k-1|r}$, whilst for $(x, y) \in \mathcal{A}_r \times \mathcal{C}_r$, $\tilde{\alpha}(x, y) \leq 1$. Also, $\int_{\mathcal{B}_r} q(x, y) dy \leq 1$, whilst $\int_{\mathcal{C}_r} q(x, y) dy = e^{-kr} - e^{-2kr} \leq e^{-kr}$. Therefore, for $r > \log 2$ (so $\pi(\mathcal{A}_r) < 1/2$) the flow out of $\mathcal{A}_r$ satisfies

$$\begin{aligned} \pi(\mathcal{A}_r) \kappa_{DA}(\mathcal{A}_r) &= \int_{x \in \mathcal{A}_r, y \in \mathcal{B}_r} \pi(x) q(x, y) \tilde{\alpha}(x, y) dy dx + \int_{x \in \mathcal{A}_r, y \in \mathcal{C}_r} \pi(x) q(x, y) \tilde{\alpha}(x, y) dy dx \\ &\leq e^{-|k-1|r} \int_{x \in \mathcal{A}_r} \pi(x) dx + e^{-kr} \int_{x \in \mathcal{A}_r} \pi(x) dx = (e^{-|k-1|r} + e^{-kr}) \pi(\mathcal{A}_r). \end{aligned}$$

So, for any $\epsilon > 0 \exists r$ such that $\kappa_{DA}(\mathcal{A}_r) < \epsilon$, and the conductance of the chain is therefore 0; the chain is not variance bounding. The lack of geometric ergodicity follows from Proposition 1. $\square$

Proofs of Results in Sect. 4

Shorthand for Delayed-Acceptance Kernels

The following short-hand is used through the remainder of this section.

$$b_1(x, y) := \log r_1 = [\log \hat{\pi}(y) - \log \hat{\pi}(x)] + \Delta(x, y; q), \quad (22)$$

$$b_2(x, y) := \log r_2 = [\log \pi(y) - \log \pi(x)] - [\log \hat{\pi}(y) - \log \hat{\pi}(x)]. \quad (23)$$

Proof of Theorem 3

For any $\epsilon > 0$, by Eq. 9, choose r_ϵ such that $\int_{B(x, r_\epsilon)^c} q(x, y) dy \leq \epsilon$; set $\mathcal{D}(x) := B(x, r_\epsilon)$ so Eq. 14 holds.

The ‘heavier-tail’ condition Eq. 15 is equivalent to $b_1(x, y) + \Delta(x, y; q) \geq 0 \Rightarrow b_2(x, y) \geq 0$. Applying the identity $[b_1(y, x), b_2(y, x), \Delta(y, x; q)] = -[b_1(x, y), b_2(x, y), \Delta(x, y; q)]$ and relabelling x and y then gives $b_1(x, y) + \Delta(x, y; q) \leq 0 \Rightarrow b_2(x, y) \leq 0$.

Next, suppose that $\|x\| > r_*$, $\|y\| > r_*$ and $y \in \mathcal{D}(x)$ but $y \notin \mathcal{C}(x)$, so that $b_1(x, y)$ and $b_2(x, y)$ have opposite signs. By Eq. 17 $|\Delta(x, y; q)| \leq h(r_\epsilon)$, and the implications derived in the previous paragraph then imply that both $|b_1(x, y)| \leq h(r_\epsilon)$ and $|b_2(x, y)| \leq h(r_\epsilon)$. Thus $\mathcal{D}(x) \cap \mathcal{M}_{h(r_\epsilon)}(x) \subseteq \mathcal{C}(x)$.

Finally, let $\bar{B}(x, r)$ be the closure of $B(x, r)$, let $\mathcal{D} := \bar{B}(0, r_\epsilon + r_*) \times \bar{B}(0, 2r_\epsilon + r_*)$ and let $m_* := \sup_{(x,y) \in \mathcal{D}} |b_2(x, y)|$; since $b_2(x, y)$ is continuous and $\mathcal{D}$ is compact, $m_* < \infty$. For $x \in \bar{B}(0, r_\epsilon + r_*)$ and $y \in \mathcal{D}(x)$, $(x, y) \in \mathcal{D}$ and so $|b_2(x, y)| \leq m_*$. For $x \in \bar{B}(0, r_\epsilon + r_*)^c$ and $y \in \mathcal{D}(x)$, $\|x\| > r_*$ and $\|y\| > r_*$ and, from the previous paragraph, $\mathcal{D}(x) \cap \mathcal{M}_{h(r_\epsilon)}(x) \subseteq \mathcal{C}(x)$. Hence Eq. 13 holds with $m = \max(h(r_\epsilon), m_*)$, and the result follows from the proof of Corollary 1. $\square$

Proof of Proposition 3

Let $q(x, y)$ and $q_{\text{RWM}}(x, y)$ be the proposal densities for the Metropolis-Hastings and RWM algorithms, respectively, let $\alpha(x, y)$ and $\alpha_{\text{RWM}}(x, y)$ be the corresponding acceptance probabilities, and let $v_* := \text{ess sup } \|v(x)\|$.

Firstly, if $\int_{B(x, r)^c} q_{\text{RWM}}(x, y) dy < \epsilon$ then $\int_{B(x, r+v_*)^c} q(x, y) dy < \epsilon$; since q_{RWM} is uniformly local, so, therefore, is q . Next, algebra shows that

$$\log q(x, y) - \log q_{\text{RWM}}(x, y) = \frac{1}{2\lambda^2} [2(y - x) \cdot v(x) + \|v(x)\|^2].$$

Now consider $y \in B(x, r)$ and apply the triangle inequality to obtain,

$$|\log q(x, y) - \log q_{\text{RWM}}(x, y)| \leq \frac{1}{2\lambda^2} [2rv_* + v_*^2] =: m(r).$$

Thus $q(x, y) \geq e^{-m(r)} q_{\text{RWM}}(x, y)$ and $q_{\text{RWM}}(x, y) \geq e^{-m(r)} q(x, y)$. Also $q(x, y)\alpha(x, y) = q(x, y) \wedge [q(y, x)\pi(y)/\pi(x)] \geq e^{-m(r)} q_{\text{RWM}}(x, y)\alpha_{\text{RWM}}(x, y)$ and, similarly, $q_{\text{RWM}}(x, y)\alpha_{\text{RWM}}(x, y) \geq e^{-m(r)} q(x, y)\alpha(x, y)$. Both implications then follow from Theorem 1 with $\mathcal{D}(x) = B(x, r)$ and with r chosen so that both $\int_{B(x, r)^c} q_{\text{RWM}}(x, y) dy \leq \epsilon$ and $\int_{B(x, r)^c} q(x, y) dy \leq \epsilon$. $\square$

Proof of Example 7

Let P_{RWM} be the RWM kernel using a Gaussian proposal and targeting π , as in Proposition 3. Applying Proposition 3 twice shows that P_{TMALA} is variance bounding if and only if P_{RWM} is variance bounding, which occurs if and only if P_{hyp} is variance bounding. The sufficiency of (i) then arises directly from Corollary 3. For (ii), by Proposition 2 A(ii), applied to the proposal $\hat{q}$, we may use Theorem 3. Geometric ergodicity then follows from Proposition 1. $\square$

Proofs of Examples 6 and 8

Let the proposal be $Y = x - \frac{1}{2}\lambda^2\xi x^{\xi-1} + \lambda Z$, where $Z \sim N(0, 1)$. Example 6 uses $\xi = \beta$ and Example 8 uses $\xi = \gamma$. Now $v(x) = -\xi\lambda^2 x^{\xi-1}/2$, so, from Eq. 21,

$$\Delta^{(\xi)}(x, y; q) = -\frac{\xi}{2}(x - y)(x^{\xi-1} + y^{\xi-1}) + \frac{\lambda^2\xi^2}{8}(x^{2\xi-2} - y^{2\xi-2}). \quad (24)$$

Also $Y^k = x^k \left(1 - \frac{1}{2}\lambda^2\xi x^{\xi-2} + \lambda Z/x\right)^k$. Hence, for $k > 0$,

$$\xi < 2 \Rightarrow U_x := Y^{k-1}/x^{\xi-1} \xrightarrow{p} 1,$$

$$\xi \geq 1 \Rightarrow V_{x,k} := \frac{x^k - Y^k}{x^{\xi+k-2}} \xrightarrow{p} \frac{1}{2}k\xi\lambda^2 - kZ\mathbb{1}_{(\xi=1)},$$

as $x \rightarrow \infty$, and where here, and throughout this proof $\xrightarrow{p}$ indicates convergence in probability. Now

$$\Delta^{(\xi)} = -\frac{1}{2}\xi x^{\xi-1} \left(\frac{1}{2}\lambda^2 \xi x^{\xi-1} - \lambda Z \right) (1 + U_x) + \frac{1}{8}\lambda^2 \xi^2 x^{3\xi-4} V_{x,2\xi-2}.$$

However, if $\xi < 2$, $x^{3\xi-4}/x^{2\xi-2} \rightarrow 0$ as $x \rightarrow \infty$, so, for $1 \leq \xi < 2$,

$$T_{x,\xi} := \frac{\Delta^{(\xi)}}{x^{2\xi-2}} \xrightarrow{p} -\frac{1}{2}\xi^2 \lambda^2 + \xi \lambda Z \mathbb{1}_{(\xi=1)},$$

as $x \rightarrow \infty$. Finally,

$$b_1(x, Y) = x^{\xi+\gamma-2} V_{x,\gamma} + x^{2\xi-2} T_{x,\xi} \quad \text{and} \quad b_2(x, Y) = x^{\xi+\beta-2} V_{x,\beta} - x^{\xi+\gamma-2} V_{x,\gamma}.$$

Example 6 ($\xi = \beta$). Consider the behaviour of $b_1(x, Y)$ and $b_2(x, Y)$ as $x \rightarrow \infty$. If $1 \leq \gamma < \beta$, b_1 is dominated by $x^{2\beta-2} T_{x,\beta}$ and so $b_1 \xrightarrow{p} -\infty$. If $1 \leq \beta < \gamma$, b_1 is dominated by $x^{\beta+\gamma-2} V_{x,\gamma}$ and b_2 is dominated by $-x^{\beta+\gamma-2} V_{x,\gamma}$; thus, when $\beta > 1$, $b_2 \rightarrow -\infty$, and when $\beta = 1$, either $b_1 \rightarrow -\infty$ or $b_2 \rightarrow -\infty$, depending on the value of Z . In either case, $\tilde{\alpha}(x, Y) \xrightarrow{p} 0$ as $x \rightarrow \infty$. Given $\epsilon > 0$, we choose x_* such that for all $x > x_*$ $\mathbb{P}(\tilde{\alpha}(x_*, Y) > \epsilon) < \epsilon$ and set $\mathcal{A}_x := [x, \infty)$, so that $\kappa(\mathcal{A}_{x_*}) < 2\epsilon$. But ϵ can be made as small as desired, so $\kappa_{DA} = 0$.

Example 8 ($\xi = \gamma$). When $1 \leq \beta < \gamma$, as $x \rightarrow \infty$, b_2 is dominated by $-x^{2\gamma-2} V_{x,\gamma} \xrightarrow{p} -\infty$ so $\kappa_{DA} = 0$, by an analogous argument to that for Example 6.

If $1 \leq \gamma < \beta$ then we note that the proof of geometric ergodicity of the MALA algorithm in Theorem 4.1 of Roberts and Tweedie (1996a) applies to any algorithm with a proposal of the form $q(x, y) \propto \exp[-|y - [x + v(x)]|^2/2\lambda^2]$. For an irreducible and aperiodic kernel with a continuous proposal density, such as the one under consideration, geometric ergodicity is therefore guaranteed provided the following two conditions are satisfied:

$$\eta := \liminf_{|x| \rightarrow \infty} (|x| - |x + v(x)|) > 0 \quad \text{and} \quad \lim_{|x| \rightarrow \infty} \int_{\mathbf{R}(x) \cap \mathbf{I}(x)} q(x, y) dy = 0,$$

where $\mathbf{I}(x) := \{y : |y| \leq |x|\}$ is the interior and $\mathbf{R}(x) := \{y : \alpha(x, y) < 1\}$ is the region where a rejection is possible. Now, $x + v(x) = x - \gamma x^{\gamma-1}$ so $\gamma \geq 1 \Rightarrow \eta > 0$. We will show that if $\gamma \in [1, 2)$, for $x > 1$ and $y \in \mathbf{I}(x)$, both $b_1(x, y) \geq 0$ and $b_2(x, y) \geq 0$, so that $\tilde{\alpha}_b(x, y) = 1$ and hence $\mathbf{R}(x) \cap \mathbf{I}(x)$ is empty.

Now $b_2(x, y) = x^\gamma(x^{\beta-\gamma} - 1) + y^\gamma(y^{\beta-\gamma} - 1)$, so if $x \geq y$ and $x \geq 1$ then $b_2(x, y) \geq 0$. Further, from Eq. 24,

$$b_1(x, y) = x^\gamma - y^\gamma - \frac{\gamma}{2}(x - y)(x^{\gamma-1} + y^{\gamma-1}) + \frac{\lambda^2 \gamma^2}{8}(x^{2\gamma-2} - y^{2\gamma-2}).$$

The final term is non-negative when $x \geq y$. Directly from the concavity of $f(t) = \gamma t^{\gamma-1}$, we obtain

$$x^\gamma - y^\gamma = \int_y^x f(t) dt \geq (x - y) \frac{f(x) + f(y)}{2} = \frac{\gamma}{2}(x^{\gamma-1} + y^{\gamma-1}),$$

so the sum of the first two terms is also non-negative when $x \geq y$. Hence, for $x \geq 1$, $\mathbf{I}(x) \cap \mathbf{R}(x)$ is empty, as claimed. $\square$

Proof of Example 9

Consider any proposal of the form $Y = x + \lambda^2 v(x) + \lambda Z$, where $Z \sim \mathbf{N}(0, 1)$ and $|v(x)| \leq v_*$ for all x . Firstly,

$$\log \hat{\pi}(Y) - \log \hat{\pi}(x) = \frac{1}{\gamma} (x^\gamma - Y^\gamma).$$

Secondly, as $x \rightarrow \infty$,

$$\mathbb{P}(Y > x/2) = \mathbb{P}(\lambda Z > -x/2 - \lambda^2 v(x)) > \mathbb{P}(\lambda Z > -x/2 + \lambda^2 v_*) \rightarrow 1. \quad (25)$$

Also, for any $\delta > 0$,

$$\mathbb{P}(|Y - x| \geq \lambda\delta) = \mathbb{P}(|\lambda^2 v(x) + \lambda Z| \geq \delta) \geq 1 - 2\delta/\sqrt{2\pi}. \quad (26)$$

The intermediate value theorem supplies: $\log \hat{\pi}(Y) - \log \hat{\pi}(x) = \eta^{\gamma-1}(Y - x)$ for some $\eta \geq \min(x, Y)$. Given any $\epsilon > 0$, set $\delta = \sqrt{2\pi\epsilon}/4$ and choose x_* such that $\mathbb{P}(Y > x/2) > 1 - \epsilon/2$ for all $x > x_*$. Then with probability at most ϵ , $|\log \hat{\pi}(Y) - \log \hat{\pi}(x)| > (x/2)^{\gamma-1} \{|Y - x| \vee \delta\}$. Next,

$$\begin{aligned} \Delta(x, Y) &= \log q(Y, x) - \log q(x, Y) = \frac{1}{2\lambda^2} [\lambda^4 v(x)^2 - \lambda^4 v(y)^2 - 2\lambda(Y - x)\{\lambda^2 v(x) + \lambda^2 v(y)\}] \\ &\leq \frac{1}{2\lambda^2} [\lambda^4 v_*^2 - 4\lambda^3(Y - x)v_*]. \end{aligned}$$

Hence, $|\log \hat{\pi}(Y) - \log \hat{\pi}(x)| \rightarrow \infty$ and dominates $\Delta(x, Y)$ in probability as $x \rightarrow \infty$; i.e., $\log r_1(x, Y) \sim \log \hat{\pi}(Y) - \log \hat{\pi}(x)$ and $|\log r_1(x, Y)| \rightarrow \infty$ in probability as $x \rightarrow \infty$.

After some algebra,

$$\log r_2(x, Y) = \frac{1}{\gamma} Y^\gamma \left(1 - \frac{\gamma}{\beta Y^{\gamma-\beta}} \right) - \frac{1}{\gamma} x^\gamma \left(1 - \frac{\gamma}{\beta x^{\gamma-\beta}} \right).$$

Since $\gamma > \beta$, as x (and hence, Y), becomes large, we have

$$\log r_2(x, Y) \sim -\{\log \hat{\pi}(Y) - \log \hat{\pi}(x)\}$$

in probability as $x \rightarrow \infty$.

Combining these two ideas, $0 \wedge \log r_1(x, Y) + 0 \wedge \log r_2(x, Y) \rightarrow -\infty$ in probability. Thus $\text{ess sup}_x P(x, \{x\}) = 1$ and the algorithm cannot be geometrically ergodic by Theorem 5.1 of Roberts and Tweedie (1996b); by Proposition 1 it also cannot be variance bounding.

Acknowledgements This paper was motivated by initial conversations with Alexandre Thiery about the preservation, and lack of preservation, of geometric ergodicity for delayed-acceptance kernels.

Data Availability Data sharing is not applicable to this article as no new data were created or analysed in this study.

Code Availability R code to produce the plots in Section 5 is available from <https://chrisgsherlock.github.io/Research/publications.html>.

Declarations

Conflict of Interests None

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Banterle M, Grazian C, Lee A, Robert CP (2019) Accelerating metropolis-hastings algorithms by delayed acceptance. *Foundations of Data Science* 1(2639-8001-2019-2-103):103
- Besag J (1994) In discussion of 'Representations of knowledge in complex systems' by U. Grenander and M. Miller. *J Roy Stat Soc Ser B* 56:591–592
- Christen JA, Fox C (2005) Markov chain Monte Carlo using an approximation. *J Comp Graph Stat* 14(4):795–810
- Cui T, Fox C, O'Sullivan M (2011) Bayesian calibration of a large-scale geothermal reservoir model by a new adaptive delayed acceptance metropolis hastings algorithm. *Water Resour Res* 47(10)
- Franks J, Vihola M (2020) Importance sampling correction versus standard averages of reversible MCMCs in terms of the asymptotic variance. *Stochastic Processes and their Applications*, Early availability online
- Geyer CJ (1992) Practical Markov Chain Monte Carlo. *Stat Sci* 7(4):473–483
- Geyer CJ (2011) Introduction to Markov chain Monte Carlo. In: Brooks S, Gelman A, Jones GL, Meng X-L (eds) *Handbook of Markov chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press, Boca Raton, FL, pp 3–48
- Gilks WR, Richardson S, Spiegelhalter DJ (1996) *Markov Chain Monte Carlo in practice*. Chapman and Hall, London, UK
- Golightly A, Henderson DA, Sherlock C (2015) Delayed acceptance particle MCMC for exact inference in stochastic kinetic models. *Stat Comput* 25(5):1039–1055
- Higdon D, Reese CS, Moulton JD, Vrugt JA, Fox C (2011) Posterior exploration for computationally intensive forward models. In: Brooks S, Gelman A, Jones GL, Meng X-L (eds) *Handbook of Markov chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press, Boca Raton, FL, pp 401–418
- Hoff PD (2009) *A first course in Bayesian statistical methods*. Springer, Dordrecht, Heidelberg, London, New York
- Jarner SF, Hansen E (2000) Geometric ergodicity of Metropolis algorithms. *Stochastic Process Appl* 85(2):341–361
- Jerrum M, Sinclair A (1988) Conductance and the rapid mixing property for markov chains: The approximation of permanent resolved. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing, STOC '88*, page 235–244, New York, NY, USA. Association for Computing Machinery
- Lawler GF, Sokal AD (1988) Bounds on the L^2 spectrum for Markov chains and Markov processes: a generalization of Cheeger's inequality. *Trans Amer Math Soc* 309(2):557–580
- Lee A, Łatuszyński K (2014) Variance bounding and geometric ergodicity of Markov chain Monte Carlo kernels for approximate Bayesian computation. *Biometrika* 101(3):655–671
- Liu JS (1996) Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Stat Comp* 113–119
- Liu JS (2001) *Monte Carlo Strategies In Scientific Computing*. Springer
- Livingstone S, Betancourt M, Byrne S, Girolami M (2019) On the geometric ergodicity of Hamiltonian Monte Carlo. *Bernoulli* 25(4A):3109–3138
- Mengersen KL, Tweedie RL (1996) Rates of convergence of the Hastings and Metropolis algorithms. *Ann Statist* 24(1):101–121
- Meyn S, Tweedie R (1993) *Markov chains and stochastic stability*. Springer-Verlag, London

- Moulton JD, Fox C, Svyatskiy D (2008) Multilevel approximations in sample-based inversion from the Dirichlet-to-Neumann map. *J Phys: Conf Ser* 124(1)
- Payne RD, Mallick BK (2014) Bayesian Big Data Classification: A Review with Complements. *ArXiv e-prints*
- Peskun PH (1973) Optimum Monte-Carlo sampling using Markov chains. *Biometrika* 60:607–612
- Quiroz M, Tran M-N, Villani M, Kohn R (2018) Speeding up MCMC by delayed acceptance and data sub-sampling. *J Comput Graph Stat* 27(1):12–22
- Roberts G, Rosenthal J (1997) Geometric Ergodicity and Hybrid Markov Chains. *Electron Commun Probab* 2:13–25
- Roberts GO, Rosenthal JS (2004) General state space Markov chains and MCMC algorithms. *Probab Surv* 1:20–71
- Roberts GO, Rosenthal JS (2008) Variance bounding Markov chains. *Ann Appl Probab* 18(3):1201–1214
- Roberts GO, Tweedie RL (1996) Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* 2(4):341–363
- Roberts GO, Tweedie RL (1996) Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika* 83(1):95–110
- Rudolph A, Sprungk B (2016) On a generalization of the preconditioned Crank–Nicolson algorithm. *Found Comput Math*
- Sherlock C, Henderson DA (2017) Adaptive, delayed-acceptance MCMC for targets with expensive likelihoods. *J Comput Graph Stat* 26(2):434–444
- Sherlock C, Thiery A, Golightly A (2021) Efficiency of delayed-acceptance random walk Metropolis algorithms. *The Annals of Statistics*. <https://projecteuclid.org/journals/annals-of-statistics/volume-49/issue-5/Efficiency-of-delayed-acceptance-random-walk-Metropolis-algorithms/10.1214/21-AOS2068.full>
- Smith ME (2011) Estimating nonlinear economic models using surrogate transitions. Available from https://files.nyu.edu/mes473/public/Smith/_Surrogate.pdf
- Tierney L (1994) Markov chains for exploring posterior distributions. *Ann Statist* 22(4):1701–1762. With discussion and a rejoinder by the author
- Tierney L (1998) A note on Metropolis-Hastings kernels for general state spaces. *Ann Appl Probab* 8(1):1–9
- Yosida K (1980) *Functional analysis*. Springer, Berlin, 6 edition

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.