

MIT Open Access Articles

Rate-optimal refinement strategies for local approximation MCMC

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Statistics and Computing. 2022 Aug 09;32(4):60

Published Version: <https://doi.org/10.1007/s11222-022-10123-0>

Publisher: Springer US

Permanent Link: <https://hdl.handle.net/1721.1/144274>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: <http://creativecommons.org/licenses/by-nc-sa/4.0/>



Rate-optimal refinement strategies for local approximation MCMC

Cite this Accepted Manuscript (AM) as Accepted Manuscript (AM) version of Andrew D. Davis, Youssef Marzouk, Aaron Smith, Natesh Pillai, Rate-optimal refinement strategies for local approximation MCMC, *Statistics and Computing* <https://doi.org/10.1007/s11222-022-10123-0>

This AM is a PDF file of the manuscript accepted for publication after peer review, when applicable, but does not reflect post-acceptance improvements, or any corrections. Use of this AM is subject to the publisher's embargo period and AM terms of use. Under no circumstances may this AM be shared or distributed under a Creative Commons or other form of open access license, nor may it be reformatted or enhanced, whether by the Author or third parties. See here for Springer Nature's terms of use for AM versions of subscription articles: <https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>

The Version of Record of this article, as published and maintained by the publisher, is available online at: <https://doi.org/10.1007/s11222-022-10123-0>. The Version of Record is the version of the article after copy-editing and typesetting, and connected to open research data, open protocols, and open code where available. Any supplementary information can be found on the journal website, connected to the Version of Record.

Accepted manuscript

Noname manuscript No.
(will be inserted by the editor)

Rate-optimal refinement strategies for local approximation MCMC

Andrew D. Davis · Youssef Marzouk · Aaron Smith ·
Natesh Pillai

the date of receipt and acceptance should be inserted later

Abstract Many Bayesian inference problems involve target distributions whose density functions are computationally expensive to evaluate. Replacing the target density with a local approximation based on a small number of carefully chosen density evaluations can significantly reduce the computational expense of Markov chain Monte Carlo (MCMC) sampling. Moreover, continual refinement of the local approximation can guarantee asymptotically exact sampling. We devise a new strategy for balancing the decay rate of the bias due to the approximation with that of the MCMC variance. We prove that the error of the resulting local approximation MCMC (LA-MCMC) algorithm decays at roughly the expected $1/\sqrt{T}$ rate, and we demonstrate this rate numerically. We also introduce an algorithmic parameter that guarantees convergence given very weak tail bounds, significantly strengthening previous convergence results. Finally, we apply LA-MCMC to a computationally intensive Bayesian inverse problem arising in groundwater hydrology.

Keywords Markov chain Monte Carlo · local regression · Bayesian inference · surrogate models · sampling methods

PACS 02.50.r · 02.50.Ng · 02.50.Tt · 02.70.Uu

1 Introduction

Markov chain Monte Carlo (MCMC) is a sampling algorithm that in principle can simulate any probability distribution. In practice, many MCMC algorithms are infeasible when target density evaluations are computationally intensive. In this paper, we present an algorithm that mitigates this issue, achieving two desirable properties: (i) vastly reducing the number of target density evaluations per MCMC step, while (ii) *provably* retaining a convergence rate roughly equal to that of the original expensive algorithm.

Challenges associated with computationally intractable target densities are extremely well known in the Monte Carlo literature. We now give a short survey of approaches to the problem, emphasizing the contributions of the current paper. Replacing the intractable target density with an approximation or “surrogate model” is the simplest way to significantly reduce the computational cost of MCMC. Replacing the target density with a *fixed* surrogate, however, introduces a non-vanishing bias (Marzouk and Xiu, 2009; Cotter et al., 2010; Bliznyuk et al., 2012; Li and Marzouk, 2014; Cui et al., 2016; Stuart and Teckentrup, 2018). The obvious solution to the

AS was supported by NSERC. AD and YM were supported by the SciDAC program of the DOE Office of Advanced Scientific Computing Research.

A. D. Davis
Courant Institute of Mathematical Sciences, New York, NY USA E-mail: davisad@alum.mit.edu

Y. M. Marzouk
Massachusetts Institute of Technology, Cambridge, MA USA E-mail: ymarz@mit.edu

A. M. Smith
University of Ottawa, Ottawa, Canada E-mail: asmi28@uottawa.ca

N. Pillai
Harvard University, Cambridge, MA USA E-mail: pillai@fas.harvard.edu

problem of non-vanishing bias is to iteratively refine the surrogate during sampling. This strategy introduces some non-trivial technical difficulties: continual surrogate refinement within MCMC results in a process that is not Markovian. Such non-Markovian processes can have surprising and terrible convergence properties even when the surrogate has very low pointwise error (Łatuszyński and Rosenthal, 2014), sometimes failing to converge at all. In our previous work (Conrad et al., 2016, 2018), we presented an MCMC algorithm with a continually refined surrogate based on *local approximations*. We showed that this algorithm avoided the worst of these convergence problems and produced asymptotically exact results under certain strong assumptions.

These strong assumptions are symptoms of a practical problem with how local approximation and MCMC interact in the “tails” of the target distribution. Conrad et al. (2016, 2018) built a piecewise polynomial approximation by solving a local regression problem, using a small number of nearby exact target density evaluations. Refinements triggered by a cross-validation heuristic on the acceptance ratio added new points to the set of target density evaluations. Ideally, this approach should avoid expensive density evaluations in regions of low probability—particularly the tails of the distribution. In practice, there is an important tension: the cross-validation heuristic requires many *more* tail evaluations than a good global design would suggest, but gives many *fewer* tail evaluations than are sometimes required for global convergence and stability of the MCMC algorithm.

In Conrad et al. (2016, 2018) we showed that this tension could be avoided under certain strong tail conditions, with the additional requirement of randomly triggered surrogate refinements to ensure convergence.

In more general settings, the underlying tension remains and leads to two serious issues which we address in the present paper. First is the behavior of the cross-validation heuristic noted above: it demands more refinements in low probability regions than the goal of “good” approximation of the target density (in a natural sense that we will make precise later) would dictate. Second, and even worse, local polynomial approximations typically have very bad tail behavior, even failing to be integrable unless nearly every point in the tail has a density evaluation.

To resolve the first issue, we introduce here a new surrogate refinement strategy that relaxes the acceptable error threshold in the tails of the distribution. The refinement strategy also *balances* the rate of decay of the surrogate bias with that of Monte Carlo variance, using a local error indicator to characterize the bias. To resolve the second issue, we introduce a new correction term for the Metropolis-Hastings acceptance ratio that allows the algorithm to retain good convergence properties even under very weak assumptions about tail behavior, and without requiring integrability of the surrogate. The result is a local approximation MCMC (LA-MCMC) algorithm that tends to require far fewer density evaluations (resulting in faster runs) while delivering robust performance for a much broader class of target distributions. From a practical point of view, work such as Conrad et al. (2016, 2018); Angelikopoulos et al. (2015) showed that the number of expensive target density evaluations can be reduced by orders of magnitude with minimal impact on the accuracy of target expectation estimates. Our new algorithm further reduces the number of density evaluations to nearly the rate-optimal number, without sacrificing MCMC stability.

From a theoretical point of view, we *guarantee* that LA-MCMC converges quickly by providing useful bounds on the convergence of the algorithm after a *finite* number of MCMC steps. Our main result is that the mean-square error (MSE) of our LA-MCMC estimates decays at roughly the expected $1/T$ rate, where T is the number of MCMC steps. We also show that LA-MCMC converges under very weak tail bounds that hold for many statistical examples. Comparatively, Conrad et al. (2016) only established a law of large numbers that applied for special tail shapes.

Using our optimal surrogate refinement strategy, the refinement rate (i.e., the rate at which new density evaluations are demanded) decreases with T . This implies that the decay of the MSE as the number of density evaluations n increases is not only faster than the $1/n$ rate expected for standard geometrically ergodic MCMC, but also that this convergence rate (in n) may accelerate as the number of MCMC steps T increases.

Other related work. There have been many attempts to accelerate MCMC for models that are expensive to compute. We have already discussed work based on the construction of surrogate models/approximations for the target density. This discussion was not exhaustive, and we point to the survey paper Llorente et al. (2021) for a broader overview of various other techniques and when they are most useful. We also briefly describe several approaches that are quite different from ours. The multi-level MCMC approach of Dodwell et al. (2015), and related multi-index extensions (Jasra et al., 2018), rely on predefined hierarchies of models—e.g., corresponding to different mesh refinements of an underlying differential equation model—that induce trade-offs between computational cost and accuracy. These approaches can drastically reduce the number of expensive model evaluations and total wallclock time of an MCMC run, but require a careful understanding of numerical

approximation errors to achieve optimal convergence rates. Another part of the literature, including Kaipio and Somersalo (2007); Chkrebtii et al. (2016) and other related work, creates statistical models of numerical discretization error in ODE or PDE models. The delayed-acceptance MCMC method (e.g., Christen and Fox (2005); Cui et al. (2011)) can also exploit a hierarchy of models, but instead “screens” MCMC proposals through the approximate model(s) and evaluates the expensive target density at least once for each accepted sample, thus reducing cost by a constant factor.

The MCMC method of Sherlock et al. (2017) makes central use of delayed acceptance, but also builds surrogates through local approximation of the target density—and in that sense has many interesting links to the present approach. We summarize some of the similarities and differences as follows. First, rather than using a local polynomial approximation as we do here, Sherlock et al. (2017) approximates the target density (or likelihood) with a simpler weighted average of its values at the k nearest previous evaluation points. In principle, however, both algorithms can be run with a wide variety of local approximations. A more substantive distinction is that Sherlock et al. (2017) fundamentally relies on the delayed-acceptance construction, and additionally on the presence of a standard non-adaptive Metropolis transition kernel employing the expensive target, to ensure ergodicity. The ratio of model evaluations to MCMC steps in the approximate MCMC algorithm is thus bounded below by a nonzero constant. Moreover, to ensure diminishing adaptation, not all of the expensive target density evaluations are used to inform the surrogate, which reduces efficiency. In contrast, both Conrad et al. (2016, 2018) and the new construction we propose here do not employ delayed acceptance, and instead carefully control the convergence of the local approximation to ensure ergodicity for the exact target. All target density evaluations are used to build the surrogate, and we use a controlled, rather than opportunistic, design strategy to choose new evaluation points.

A key improvement of the present work over both Sherlock et al. (2017) and Conrad et al. (2016, 2018) is that our new algorithm has much stronger theoretical convergence guarantees, and is applicable to a broader class of target distributions. For instance, earlier local approximation methods assume that the target distribution has tails that are very light or otherwise very special (Sherlock et al. (2017) assume a uniform minorization condition; Conrad et al. (2016) assumes either compact support or something very similar to a uniform minorization condition). Under such conditions, these efforts showed that their algorithms are ergodic (that is, Monte Carlo averages converge at some unspecified rate). Our current effort substantially improves on all of these: we allow the ratio of model evaluations to MCMC steps to decay to zero quite quickly, we allow the target distribution to have quite general tails, and we give a nearly-optimal bound on the rate of convergence. We emphasize that these improvements are not merely theoretical; they represent real improvements in algorithm performance. For example, it is shockingly easy to write down an algorithm that looks quite a bit like the one in our present paper that *does not* converge at the correct rate or *does* require the number of model evaluations to grow roughly linearly in the number of MCMC steps (see, e.g., Section 4.2).

Scope. This paper focuses on the core problems of surrogate construction and refinement, tail correction, and convergence analysis, all in the context of a Metropolis–Hastings-type scheme. Many variations and extensions are possible, and indeed it is often fruitful to *merge* the most important techniques presented in various papers. For instance, one could consider using gradient information from the surrogate within the Metropolis proposal, as in the MALA scheme of Conrad et al. (2018). Sherlock et al. (2017) has an extensive discussion of kd -trees and their use in facilitating fast nearest neighbor searching; we use kd -trees in our implementation as well (see Section 2), but without some of the online point addition and re-balancing heuristics proposed in that paper. Sherlock et al. (2017) also consider a pseudo-marginal version of their approximate MCMC algorithm, where only noisy unbiased evaluations of the expensive target density are available. An exhaustive evaluation and comparison of these methods is beyond the scope of the present work. A complete implementation of our algorithm is available within the software package MUQ (<https://muq.mit.edu>), so that users can combine our LA-MCMC scheme with more advanced transition kernels or other computational tools that improve performance.

We also note that our emphasis here is on inference problems with expensive models (for instance, numerical discretizations of partial differential equations), but in *moderate* parameter dimensions (for instance, $d = 9$ in our last example, or $d = 12$ in Conrad et al. (2018)). Posterior sampling in higher dimensional settings presents an additional set of challenges, outside of the present scope. Moreover, function approximation (and hence surrogate modeling) in high dimensions is generically subject to the curse of dimensionality, unless one can exploit some special structure. For example, Cui et al. (2016) combines parameter-space dimension reduction with intrusive model reduction methods for expensive likelihoods. It should be feasible to combine our LA-MCMC approach

with dimension reduction methods for Bayesian inverse problems (Cui et al., 2014; Constantine et al., 2016; Zahm et al., 2022) in a similar way.

Organization of the paper. In Section 2 we present our new algorithm, including the refinement scheme, the bias-variance tradeoff, and the tail correction. Section 3 describes our main theoretical results. Section 4 provides a range of numerical examples, beginning with simple configurations intended to illustrate specific features and variations of the algorithm, and culminating in the inference of a spatially distributed coefficient in a set of nonlinear PDEs.

2 Local approximation MCMC

We assume that our target distribution has a density $\pi(x) = \exp \mathcal{L}(x)$ on \mathbb{R}^d . Our goal is to construct an algorithm that exploits regularity in π to reduce the computational cost of simulating from the target distribution. As discussed in the Introduction, replacing target density evaluations with a continually and infinitely refined surrogate model, within MCMC, can asymptotically generate samples from the exact target distribution (Conrad et al., 2016). Given a finite number of MCMC samples, however, replacing the target density with such an approximation introduces a *surrogate bias*. In this section, we present a LA-MCMC algorithm that extends Conrad et al. (2016) in two significant ways: (i) we use the trade-off between surrogate bias and Monte Carlo variance to develop a refinement strategy with near-optimal convergence guarantees; and (ii) we introduce a parameter that allows LA-MCMC to easily characterize distributions with heavier tails given little *a priori* knowledge, while enabling additional speedups when the user has substantial *a priori* knowledge.

Our approach will build a *local polynomial* approximation (Conn et al., 2009; Kohler, 2002; Stone, 1977) of a function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, where g is chosen such that evaluating $\pi(x)$ is trivial given $g(x)$. For instance, g might be the logarithm of the target density, \mathcal{L} ; or it might be the log-likelihood function in a Bayesian setting, if the prior density is relatively simple to evaluate. Let the “evaluated set” $\mathcal{S}_n = \{x_1, \dots, x_n\}$ comprise the set of parameter values $x_i \in \mathbb{R}^d$ at which we have evaluated $g(x_i)$. For now, we take this set as given; later, we discuss how to construct it. With an appropriate construction of \mathcal{S} , local approximations allow refinement to focus on regions where the target distribution has greater mass. Polynomial approximations are appealing in this setting because they can be built easily, are cheap to evaluate, and have known analytic derivatives.

2.1 Local polynomial approximations

We construct local polynomial approximations using weighted regression (Kohler, 2002). Let \mathcal{P} be a polynomial space over \mathbb{R}^d . The approximation at each $x \in \mathbb{R}^d$ is $\hat{g}(x, \mathcal{S}_n)$ such that

$$\hat{g}(x, \mathcal{S}_n) = \arg \min_{m \in \mathcal{P}} \sum_{i=1}^n (m(x_i) - g(x_i))^2 W(x, x_i), \quad (1)$$

where $W(x, x')$ is a locally supported kernel. Typically, we choose the k -nearest neighbor kernel

$$W(x, x') = \begin{cases} 1 & \text{if } x' \in \mathcal{B}_k(x) \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

where $\mathcal{B}_k(x)$ is the smallest ball centered at x containing k elements of \mathcal{S}_n . Define a basis $\Phi = \{\phi_j\}_{j=1}^q$ for \mathcal{P} , let $\phi(x) = [\phi_1(x), \dots, \phi_q(x)]^T$, and define $\mathbf{x}_k(x, \mathcal{S}_n) = \{\tilde{x} \in \mathcal{S}_n : W(x, \tilde{x}) > 0\}$. The elements of $\mathbf{x}_k(x, \mathcal{S}_n)$ are the k nearest neighbors to x in \mathcal{S}_n . The optimal polynomial kernel estimate is

$$\hat{g}(x, \mathcal{S}_n) = \phi(x)^T \mathbf{a}(x, \mathbf{x}_k(x, \mathcal{S}_n)) \quad (3)$$

such that

$$\mathbf{a}(x, \mathbf{x}_k) = \arg \min_{\alpha \in \mathbb{R}^q} \|\mathbf{V}(\mathbf{x}_k) \alpha - \mathbf{g}(\mathbf{x}_k)\|_{\mathbf{W}(x, \mathbf{x}_k)}^2, \quad (4)$$

where $\|\mathbf{x}\|_{\mathbf{A}}^2 = \mathbf{x}^T \mathbf{A} \mathbf{x}$,

$$\mathbf{W}(x, \mathbf{x}_k) = \text{diag}[W(x, x_1), \dots, W(x, x_k)], \quad (5a)$$

$\mathbf{V}(\mathbf{x}_k)$ is the Vandermonde matrix

$$\mathbf{V}(\mathbf{x}_k) = \begin{bmatrix} \phi(x_1)^T \\ \vdots \\ \phi(x_k)^T \end{bmatrix}, \quad (5b)$$

and

$$\mathbf{g}(\mathbf{x}_k) = \begin{bmatrix} g(x_1) \\ \vdots \\ g(x_k) \end{bmatrix}. \quad (5c)$$

Assuming $\mathbf{V}(\mathbf{x}_k)$ has full column rank and $k \geq q$, the local polynomial approximation exists and is unique (Kohler, 2002; Stone, 1977). The solution to (4) is

$$\mathbf{a}(x, \mathbf{x}_k(x, \mathcal{S}_n)) = (\mathbf{V}^T \mathbf{W} \mathbf{V})^{-1} \mathbf{V}^T \mathbf{g}. \quad (6)$$

2.1.1 Error analysis

Now we derive local error bounds that, for a fixed evaluated set \mathcal{S}_n , depend on the number of nearest neighbors, the size of the ball containing them, and the local behavior of the surrogate. We assume that the kernel $W(x, x')$ is the hat kernel defined in (2) and that \mathcal{P} is the space of polynomials of degree less than or equal to p ; we write $q = \dim(\mathcal{P})$.

We first investigate the local behavior of the polynomials in the ball $\mathcal{B}_k(x)$. Let

$$\Delta(x) = \max_{x_i \in \mathbf{x}_k(x, \mathcal{S}_n)} \|x - x_i\|$$

be the radius of $\mathcal{B}_k(x)$ —the ball that contains k points in the evaluated set \mathcal{S}_n . We note that the radius of $\mathcal{B}_k(x)$ is not uniform in $x \in \mathbb{R}^d$. Analogous to (4), define k Lagrange polynomials

$$\lambda_j(x) = \arg \min_{m \in \mathcal{P}} \sum_{i=1}^k (m(x_i) - \delta_{ij})^2, \quad (7)$$

where δ_{ij} is the Dirac delta. Here, the summation is over the k nearest neighbors to x , i.e., $x_i \in \mathbf{x}_k(x)$. Let $\boldsymbol{\lambda}(x) = [\lambda_1(x), \dots, \lambda_k(x)]^T$. The nearest neighbors $\mathbf{x}_k(x)$ are called Λ -poised in $\mathcal{B}_k(x)$ if

$$\max_{x' \in \mathcal{B}_k(x)} \|\boldsymbol{\lambda}(x')\|_2 = \Lambda_2(x) \leq \Lambda. \quad (8)$$

This definition of Λ -poisedness is slightly different than in Conn et al. (2009), which uses the infinity norm. Since the number of nearest neighbors k is fixed and finite, however, these definitions are equivalent. Λ -poisedness measures how well distributed the points are within the ball. For example, consider points i and j in a unit ball and quadratic polynomials. The i^{th} Lagrange polynomial is close to one at point i and close to zero at point j . (If the number of nearest neighbors k is exactly the number required to interpolate ($k = q$) then the Lagrange polynomial will be exactly one or zero at these points.) The quadratic Lagrange polynomial is narrower if the points are close together and wider if they are far apart. The wider parabola has a smaller poisedness constant than the narrower one.

Assuming Λ -poisedness and that $g(x)$ is at least $(p+1)$ times differentiable, we have the error bound

$$|\hat{g}(x, \mathcal{S}_n) - g(x)| \leq \frac{k}{(p+1)!} \Lambda \Delta(x)^{p+1} \sup_{x' \in \mathcal{B}_k(x)} g^{(p+1)}(x'), \quad (9)$$

where p is the polynomial degree and $g^{(p+1)}(x)$ is the $(p+1)^{\text{th}}$ derivative of $g(x)$ (Conn et al., 2009). The radius $\Delta(x)$ decreases as points are added to \mathcal{S}_n and thus the error bound decreases, assuming that points are chosen in a way that maintains Λ -poisedness.

2.1.2 Local refinements

Given a surrogate $\hat{g}(x, \mathcal{S}_n)$, we can improve its accuracy in a neighborhood around x by performing a *local refinement*. A local refinement adds a new point x^* to the evaluated set—i.e., $\mathcal{S}_{n+1} = \mathcal{S}_n \cup \{x^*\}$ —and thus allows the local error bound (9) to decrease. Randomly choosing a nearby point tends to form clusters (Rote and Tichy, 1996), however, and therefore refining using a random point inside $\mathcal{B}_k(x)$ fails to maintain Λ -poisedness.

We instead choose the refinement location based on the poisedness constant $\Lambda_2(x)$ defined in (8). Computing $\Lambda_2(x)$ by solving the optimization problem defined in (8) also defines the point

$$x_\lambda(x) = \arg \min_{x' \in \mathcal{B}_k(x)} \|\lambda(x')\|_2. \quad (10)$$

If $\mathbf{x}_k(x)$ is poorly poised in $\mathcal{B}_k(x)$ ($\Lambda_2(x) \gg 1$) then the new point $x_\lambda(x)$ will be relatively far from any clusters in $\mathbf{x}_k(x)$ and will tend to improve the poisedness of the updated set. Conversely, if $\mathbf{x}_k(x)$ is well poised in $\mathcal{B}_k(x)$ ($\Lambda_2(x) \approx 1$), then the updated set remains well poised. We therefore refine the surrogate by setting $\mathcal{S}_{n+1} = \mathcal{S}_n \cup \{x_\lambda(x)\}$.

Having discussed how to refine, we must also discuss *when*, i.e., under what conditions, to refine. Our primary criterion will be derived from the local error bound (9); this process is described in Section 2.2.1. In the meantime, we discuss an additional but natural secondary criterion, which is to trigger a refinement if the poisedness constant exceeds a user-prescribed threshold $\bar{\Lambda}$ (i.e., if $\Lambda_2(x) > \bar{\Lambda}$ then set $\mathcal{S}_{n+1} = \mathcal{S}_n \cup \{x_\lambda(x)\}$). In practice, we find that adding $x_\lambda(x)$ to the evaluated set whenever a refinement is triggered tends to maintain poisedness automatically, and that the secondary criterion thus rarely triggers refinements. Computing the poisedness constant by solving (8) can also be computationally burdensome (although still significantly cheaper than an expensive density evaluation). We therefore often “turn off” this secondary refinement criterion by setting $\bar{\Lambda} = \infty$, and only compute the poisedness constant when refinements are triggered by the primary criterion.

2.2 Sampling methods using local approximations

Each step of a typical MCMC algorithm consists of three stages: (i) propose a new state, (ii) compute the acceptance probability, and (iii) accept or reject the proposed state. LA-MCMC adds a fourth stage—possibly refine the surrogate model—and replaces the target density evaluations in stage (ii) exclusively with cheaper surrogate evaluations; see Algorithm 1. The refinement frequency ensures that the error incurred by using a surrogate model balances the Monte Carlo variance.

2.2.1 Bias-variance trade-off

We now present a heuristic for balancing bias and variance in LA-MCMC, deferring a rigorous discussion to Section 3.

Fix a function f and Markov chain $\{X_t\}_{t \geq 0}$ with stationary measure π . For $T \in \mathbb{N}$, denote by $\hat{\pi}_T(f) = \frac{1}{T} \sum_{t=1}^T f(X_t)$ the usual Monte Carlo estimate of $\pi(f) = \mathbb{E}[f(X)]$ —the expectation with respect to π . Under modest conditions (see, e.g., Meyn and Tweedie (2012)), the asymptotic bias exists and satisfies

$$\lim_{T \rightarrow \infty} T |\mathbb{E}[\hat{\pi}_T(f)] - \pi(f)| = 0 \quad (11)$$

while the asymptotic variance exists and satisfies

$$\lim_{T \rightarrow \infty} T \text{Var}(\hat{\pi}_T(f)) = C_v \in (0, \infty). \quad (12)$$

When this happens, there is (asymptotically) no trade-off: the bias quickly becomes a negligible source of error.

In the context of LA-MCMC, more care is required because of an additional bias from the surrogate model. We assume that π has density $\pi(x) = \exp \mathcal{L}(x)$, and view LA-MCMC as an “approximate” MCMC algorithm that tries to target¹ $\hat{\pi}(x) = \exp \hat{\mathcal{L}}(x)$, where $\hat{\mathcal{L}}(x) \approx \mathcal{L}(x)$. In particular, we assume that the surrogate $\hat{\mathcal{L}}$ adheres to the error bound in (9), which is a pointwise condition of the form:

$$|\hat{\mathcal{L}}(x) - \mathcal{L}(x)| \leq C_b \Delta(x)^{p+1} \quad (13)$$

¹ We note that our $\hat{\pi}(x) = \exp \hat{\mathcal{L}}(x)$ often fails to be a probability density. We resolve this technical problem in Section 3, and somewhat surprisingly our approach means that this problem has very little impact on the following heuristic calculations.

for $C_b > 0$. We will show in Section 3 that the bias incurred from using this surrogate model inherits a similar error bound,

$$|\hat{\pi}(f) - \pi(f)| \leq C'_b \bar{\Delta}^{p+1}, \quad (14)$$

where $C'_b > 0$ and, informally, we can think of $\bar{\Delta}^{p+1}$ as the maximum of $\Delta(x)^{p+1}/V(x)$, where $V(x) > 0$ is a penalty function to be described below. We will thus use the local radius $\Delta(x)^{p+1}$ as an *error indicator*, and sequentially refine our approximation so that $\bar{\Delta} = \bar{\Delta}(T)$ decays with time T .

We would like to balance the additional error introduced by the surrogate model against the usual ‘‘Monte Carlo’’ error of the original Markov chain. We choose our parameters so that the upper bound on the surrogate error decays at the same rate as the Monte Carlo error. To ensure that the pointwise surrogate error is small enough, we partition the chain into consecutive intervals called *levels*, $\ell = 0, 1, 2, \dots$, and prescribe a piecewise constant error threshold for each level ℓ ,

$$\gamma_\ell(x) = \gamma_0 \ell^{-\gamma_1} V(x), \quad (15)$$

with $\gamma_0, \gamma_1 > 0$. Here, $V(x) > 0$ is a π -integrable penalty function that allows the error threshold to be larger in low-probability regions (e.g., in the tails of the distribution). We assume that $V(x)$ satisfies the Lyapunov inequality (21); this requirement guides how to choose the function. As a reasonable default, we set $V(x) = \exp(\nu_0 \|x - \bar{x}\|^{\nu_1})$, where $\nu_0 > 0$, $0 < \nu_1 \leq 1$, and \bar{x} is a user-prescribed estimate of the centroid of the distribution (e.g., the mean or mode). During MCMC, we locally refine the surrogate model if the error indicator $\Delta(x)^{p+1}$ at the current state exceeds the current threshold $\gamma_\ell(x)$.

Level ℓ ends at step T_ℓ , when the MCMC variance is of the same size as the squared surrogate bias:

$$\left(\gamma_0 \ell^{-\gamma_1}\right)^2 \sim C_v T_\ell^{-1}. \quad (16)$$

The number of steps in the ℓ^{th} level is, therefore,

$$T_\ell = \tau_0 \ell^{2\gamma_1}, \quad (17)$$

where $\tau_0 = C_v/\gamma_0^2$ is the length of the first level. In practice, we set the level as a function of the MCMC step t ,

$$\ell(t) = \lfloor (t/\tau_0)^{1/(2\gamma_1)} \rfloor. \quad (18)$$

The length of each level, $\tau_\ell = T_\ell - T_{\ell-1}$, will then decrease with ℓ if $\gamma_1 < 0.5$. This is undesirable because then, beyond some critical step $t^* < \infty$, the level length will become less than one. If this occurs, we will need to increment the level more than once per MCMC step. Rather than addressing this complication, we simply require that lengths of the levels strictly increase with ℓ by setting $\gamma_1 > 0.5$.

The constant τ_0 is rarely known, and we treat it as an algorithmic parameter. Given the initial error threshold $\gamma_0 > 0$, error decay rate $\gamma_1 > 0.5$, initial level length $\tau_0 \geq 1$, and the Lyapunov function $V(x)$, we locally refine whenever the error indicator at an accepted state in the chain exceeds the decaying error threshold. See Algorithm 1 for a complete summary.

2.2.2 Tail correction

There is a major technical difficulty in making the heuristic from Section 2.2.1 precise: it is possible that our surrogate function has small pointwise error as in (13), but fails to be globally integrable (and thus cannot possibly be an unnormalized probability density). This results in large practical and theoretical difficulties: the associated stochastic process may wander off to infinity, and the pointwise bound in (13) will not give any bound on the Monte Carlo error. In previous work (e.g., Conrad et al. (2016)), we avoided this problem by working on compact state spaces or obtaining much stronger pointwise approximations than (13). Here, we obtain much stronger results by slightly tweaking our algorithm’s acceptance probability.

Our tweak uses the Lyapunov function of the MCMC algorithm with exact evaluations, $V(x)$, to change the acceptance probability. Rather than computing the typical Metropolis-Hastings acceptance probability targeting $\hat{\pi}$, we slightly *increase* the chances of moves that would decrease the Lyapunov function and *decrease* the chances of moves that would increase it. Roberts and Tweedie (1996) give very general conditions under which we can easily define the Lyapunov function.

More precisely, given the current state x_t and proposed state x'_t , we tweak the usual Metropolis-Hastings acceptance probability with the estimate

$$\log \pi(x'_t) \approx \tilde{\mathcal{L}}(x'_t) \equiv \hat{\mathcal{L}}(x'_t) + Q_V(x_t, x'_t), \quad (19)$$

where

$$Q_V(x_t, x'_t) = \begin{cases} \eta(\gamma(x'_t) + \gamma(x_t)) & \text{if } V(x'_t) < V(x_t) \\ -\eta(\gamma(x'_t) + \gamma(x_t)) & \text{if } V(x'_t) \geq V(x_t), \end{cases} \quad (20)$$

and $\gamma(\cdot) = \gamma_{\ell(t)}(\cdot)$ is defined by (15) and (18). Here, $\widehat{\mathcal{L}}(x'_t)$ is an approximation of the log-density using the local polynomial approximation. The factor $\gamma(x'_t) + \gamma(x_t)$ causes the correction to become less apparent as the error threshold decays, $\gamma_{\ell(t)}(\cdot) \rightarrow 0$, and $\eta \geq 0$ is a user-defined parameter to control the correction.

Algorithm 1 Pseudocode for the LA-MCMC algorithm. LA-MCMC requires eight parameters: (i) the initial error threshold $\gamma_0 > 0$, (ii) the error threshold decay rate $\gamma_1 > 0.5$, (iii) the maximum poisedness constant \bar{A} , (iv) the length of the first level $\tau_0 \geq 1$, (v) the tail-correction parameter $\eta \geq 0$, (vi) the number of nearest neighbors k used to construct the local polynomial surrogate, (vii) the degree of the local polynomial surrogate p , and (viii) a guessed Lyapunov function V .

Set initial state X_0 and surrogate model $\widehat{\mathcal{L}}_0$ for the log-target density

for $t \leftarrow 1$ to ∞ **do**

Propose $X' \sim q_t(\cdot|X_t)$

Possibly refine at X_t : $(\gamma(X_t), \widehat{\mathcal{L}}_t) = \text{CHECKANDREFINE}(X_t, \widehat{\mathcal{L}}_{t-1}, t)$

Compute error threshold at the proposed point

$$\gamma(X') = \gamma_0(\lfloor (t/\tau_0)^{1/(2\gamma_1)} \rfloor)^{-\gamma_1} V(X')$$

Compute $\widetilde{\mathcal{L}}_t(X')$ via (19) and evaluate the acceptance probability

$$\alpha(X_t, X') = \min \left(1, \exp(\widetilde{\mathcal{L}}_t(X') - \widehat{\mathcal{L}}_t(X_t)) \frac{q_t(X_t|X')}{q_t(X'|X_t)} \right)$$

Accept/reject step:

$$X_{t+1} = \begin{cases} X' & \text{with probability } \alpha(X_t, X') \\ X_t & \text{else} \end{cases}$$

procedure CHECKANDREFINE($x, \widehat{\mathcal{L}}, t$)

Compute the local error threshold

$$\gamma(x) = \gamma_0(\lfloor (t/\tau_0)^{1/(2\gamma_1)} \rfloor)^{-\gamma_1} V(x)$$

and (optionally) compute the poisedness constant $A_2(x)$

if $\Delta(x)^{p+1} > \gamma(x)$ or (optionally) $A_2(x) > \bar{A}$ **then**

$\widehat{\mathcal{L}}^* = \text{REFINESURROGATE}(x, \widehat{\mathcal{L}})$

else $\widehat{\mathcal{L}}^* = \widehat{\mathcal{L}}$

return The error threshold $\gamma(x)$ and $\widehat{\mathcal{L}}^*$

procedure REFINESURROGATE($x, \widehat{\mathcal{L}}$)

Compute poisedness-based refinement location $x_\lambda(x)$ defined in (10)

if $x_\lambda(x) \notin \mathcal{S}_n$ **then** Set $\mathcal{S}_n \leftarrow \mathcal{S}_n \cup \{x_\lambda(x)\}$ and $n \leftarrow n + 1$

if $x_\lambda(x) \in \mathcal{S}_n$ **then** Randomly choose $X^* \in \mathcal{B}_k(x)$ and set $\mathcal{S}_n \leftarrow \mathcal{S}_n \cup \{X^*\}$ and $n \leftarrow n + 1$

return The updated surrogate model $\widehat{\mathcal{L}}^*$ using \mathcal{S}_n

2.2.3 Algorithm parameters

We briefly describe some heuristics for choosing the input parameters to Algorithm 1. Ideally, we choose the initial threshold γ_0 so that the required ball size Δ on the first error level is similar to the standard deviation of the posterior. Although the latter quantity is unknown, we can often estimate it using derivative information around the posterior mode. We have found that the decay rate $\gamma_1 > 0.5$ does not have much impact on the performance of the algorithm (see Section 4.1). We normally set $\tau_0 = 1$. Typically we use quadratic surrogates, i.e., $p = 2$; we further discuss this choice in Sections 4.1 and 4.3. We usually prescribe the number of nearest neighbors to be $k = 2q$, where $q = \dim(\mathcal{P})$ is the dimension of the polynomial space. (In principle we only need $k \geq q$, but the extra regression points add stability.)

There is usually a great deal of flexibility in the choice of Lyapunov function, as it should only have a large impact on the tails of the target distribution. In order to prevent the stochastic process from wandering to infinity, it must satisfy $\pi(V) < \infty$. If one has knowledge of the tail behavior of the target density π , choosing $V(x) \propto 1/\sqrt{\pi(x)}$ often works well. If we merely have bounds on the tails, we denote by Π a distribution with heavier tails than π and try $V(x) \propto 1/\sqrt{\Pi(x)}$. We choose the tail correction parameter η to be as small as possible without letting the chain wander out to infinity. Typically, setting $\eta = 0$ and slowly increasing the parameter until the chains stop wandering is sufficient. Finally, the bound on the poisedness constant, $\bar{\Lambda}$, is mostly included for reasons related to our proof technique. As discussed in Section 2.1.2, we often set $\bar{\Lambda} = \infty$, and even the theory allows for very large values.

Efficiently finding $x_k(x, \mathcal{S}_n)$, i.e., the k nearest neighbors to a point x , is a non-trivial but fortunately well studied problem. We store points of the evaluated set \mathcal{S}_n in kd -trees, which enable efficient nearest neighbor searching. Specifically, we use the kd -tree implementation in the library `nanoflann` (Blanco and Rai, 2014), which is quite efficient; as reported in Blanco and Rai (2014), it can build kd -trees from point sets of size $n = 10^6$ in microseconds. (This benchmark is for $d = 3$, but tree construction cost scales only linearly with d .) For the examples we have studied, which typically have $n < 10^4$ and $d \leq 12$, kd -tree construction cost is negligible. Our current implementation rebuilds the tree structure at each refinement step; this process ensures that the time needed to search the tree, a task invoked much more frequently, remains small. For problems where tree construction time is non-negligible, online point addition and re-balancing approaches, as discussed in Sherlock et al. (2017), could be beneficial to incorporate into the present workflow.

3 Theoretical results

The main theoretical results of this paper are that the LA-MCMC algorithm described in Section 2 inherits approximately the same $1/T$ convergence rate for the mean squared error as MCMC with exact evaluations and that our slight tweak to the acceptance probably allows LA-MCMC to converge even for heavier tailed distributions. We summarize the results and their implications here and refer to Appendix A for detailed discussion.

Denote by K the transition kernel of a discrete-time Markov chain on \mathbb{R}^d with unique stationary distribution π_X that has density $\pi(x)$. We wish to study the stochastic processes that approximate K , despite not being Markov chains. Note that most of the calculations in this section are similar to those in the ‘‘approximate’’ Markov chain literature (see, e.g., Johndrow et al. (2015); Medina-Aguayo et al. (2018); Pillai and Smith (2014); Rudolf et al. (2018)). The biggest differences are:

1. Our processes are not quite Markov chains. This does not substantially change any calculations, but does require us to be slightly more careful in a few steps.
2. We make slightly atypical assumptions about our Lyapunov functions. These are justified in Section A.2.

The purpose of the ‘‘tweak’’ in (19) was to force our main algorithm to satisfy the same Lyapunov condition as the baseline MCMC algorithm that it is approximating. To be more precise, we assume that the Metropolis-Hastings transition kernel K associated with proposal kernel q_t and target distribution π_X satisfies:

Assumption 1 (*Lyapunov inequality*). *There exists $V : \mathbb{R}^d \rightarrow [1, \infty)$ and constants $0 < \alpha \leq 1$ and $0 \leq \beta < \infty$ so that*

$$(KV)(x, \cdot) \leq (1 - \alpha)V(x) + \beta \quad (21)$$

for all $x \in \mathbb{R}^d$.

Assumption 2 (*Geometric ergodicity*). *Let Assumption 1 (or Assumption 4; see Appendix A) hold. There exist $0 < R < \infty$ and $0 \leq \gamma < 1$ so that*

$$\sup_{x : V(x) \leq 4\beta/\alpha} \|K^s(x, \cdot) - \pi(\cdot)\|_{TV} \leq R\gamma^s \quad (22)$$

for all $s \geq 0$.

Note that we denote by $K^s(x, \cdot)$ the s -step transition density for a chain sampled from K with starting point x , so that $K^1(x, \cdot) = K(x, \cdot)$.

We now link to Algorithm 1 with the same proposal q_t and target π_X . Let $\{X_t\}_{t \geq 0}$ be the result of a run of this algorithm, and let $\{\hat{\mathcal{L}}_t\}_{t \geq 0}$ and $\{X'_t\}_{t \geq 0}$ be the sequence of surrogates and proposals computed during the run. We make the following assumption about the ‘‘goodness’’ of our approximation:

Assumption 3 (*Approximation goodness*). The parameters of Algorithm 1 are such that the approximations $\widehat{\mathcal{L}}_t$ to the log-density satisfy

$$\left| \widehat{\mathcal{L}}_t(X_t) - \mathcal{L}(X_t) \right| + \left| \widehat{\mathcal{L}}_t(X'_t) - \mathcal{L}(X'_t) \right| \leq |Q_V(X_t, X'_t)| \quad (23)$$

deterministically, where $Q_V(X_t, X'_t)$ is as defined in (20).

Fix a function $f : \mathbb{R}^d \rightarrow [-1, 1]$ with $\pi(f) = 0$. Our main theoretical result is:

Theorem 1 Let Assumptions 1, 2, and 3 hold and let $\{X_t\}_{t=1}^T$ be generated by Algorithm 1 with starting point x . Then there exists a constant $C > 0$ so that

$$\left| \mathbb{E} \left[\left(\frac{1}{T} \sum_{t=1}^T f(X_t) \right)^2 \right] \right| \leq C \frac{\log(T)^3}{T} \quad (24)$$

for all $T > T_0(x)$ sufficiently large.

To give Theorem 1 some context, the “usual” MCMC estimate satisfies $\mathbb{E}[(\widehat{\pi}_T(f))^2] \leq C/T$ for all $T > T_0$ sufficiently large under the same Lyapunov and minorization assumptions (Assumptions 1 and 2—see Meyn and Tweedie (2012)). Furthermore, the usual bound is sharp. Our bound is nearly identical, giving up only logarithmic terms. Informally, this suggests that using our algorithm has nearly the same finite-time guarantees as typical MCMC algorithms and, depending on the specific problem, may be substantially cheaper to run.

4 Numerical examples

We present three numerical experiments. The first focuses on understanding the convergence of the LA-MCMC algorithm and the impact of various algorithmic parameters controlling the approximation. The second experiment illustrates the tail correction approach described in Section 2.2.2. The third then demonstrates the practical performance of LA-MCMC in a computationally challenging large-scale application: an inverse problem arising in groundwater hydrology.

4.1 One-dimensional toy example

We first use the one-dimensional density

$$\log \pi(x) \propto -0.5x^2 + \sin(4\pi x), \quad x \in \mathbb{R} \quad (25)$$

to demonstrate how changing various algorithmic parameters affects the performance of LA-MCMC. Figure 1 shows binned MCMC samples computed with a random-walk Metropolis algorithm that uses exact target evaluations, compared with binned samples from an LA-MCMC algorithm that uses the same proposal. The sample histograms match very closely. Figure 2 shows the error indicator $\Delta(x)^{p+1}$ and error threshold (15) computed in a single run of LA-MCMC. Both the error indicator and threshold depend on the current state X_t : the indicator depends on the local ball size, and the threshold depends on the Lyapunov function $V(X_t) = \exp(\|X_t\|)$. Intuitively, the Lyapunov function relaxes the refinement threshold in the tails of the distribution and thus prevents excessive refinement in low probability regions, where the ball size $\Delta(x)$ tends to be large.

LA-MCMC algorithmic parameters for the runs described here and below are $\gamma_0 = 0.1$, $\bar{A} = \infty$, $\tau_0 = 1$, $\eta = 0$, $k = 2(p+1)$, and $V(x) = \exp(\|x\|)$ (see Algorithm 1). We use local polynomial degree $p = 2$ and $\gamma_1 = 1$ unless otherwise indicated.

Let σ_t^2 and $\check{\sigma}_t^2$ be running t -sample estimates of the target variance computed using sequences of samples $\{X_t\}_{t>0}$ generated by MCMC with exact evaluations (which we refer to as “exact MCMC” for shorthand) and LA-MCMC, respectively. As a baseline for comparison, we also compute a ‘high fidelity’ approximation of the variance from 50 realizations of exact MCMC: $\bar{\sigma}^2 = \sum_{i=1}^{50} \sigma_{T,i}^2$, with $T = 10^6$. Then we evaluate the error in variance estimates produced by exact MCMC and LA-MCMC,

$$e_t = |\bar{\sigma}^2 - \sigma_t^2| \quad \text{and} \quad \check{e}_t = |\bar{\sigma}^2 - \check{\sigma}_t^2|. \quad (26)$$

We compute the expectations of these errors by averaging e_t and \check{e}_t over multiple independent realizations of each chain.

The bias-variance trade-off used to construct our algorithm (see Section 2.2.1) ensures that the error in an expectation computed with LA-MCMC decays at essentially the same rate as the error in an expectation computed with exact MCMC. However, we need to tune the initial error threshold γ_0 and initial level length τ_0 to ensure that the expected errors are of the same magnitude. In general, we set $\tau_0 = 1$. The initial error threshold also determines the initial local radius $\Delta(x) = \gamma_0^{1/(p+1)}V(x)$. As a heuristic, we choose γ_0 so that the initial radius is smaller than the radius of a ball containing the non-trivial support of the target density.




Fig. 1 Binned MCMC samples computed with exact evaluations (grey line) and using local approximations (red line).




Fig. 2 The local error indicator $\Delta(x)^{p+1}$ (blue line) and error threshold $\gamma_\ell(x)$ (Equation (15)—red line) used to trigger refinement. The sharp lower bound for the error threshold (bottom border of the red region) is the error threshold with $V(x) = 1$. We, however, allow $V(x) = \exp(\|x\|)$ to be larger in the tails, which relaxes the error threshold in low probability regions. This example run triggered 472 refinements. Many more (unnecessary) refinements would be required if the Lyapunov function did not relax the allowable error in the tails.

The bias-variance trade-off renders LA-MCMC insensitive to the error decay rate γ_1 , as long as $\gamma_1 \geq 0.5$. This is borne out in Figure 3(a), which shows that the error in the variance estimate decays at the same $(1/\sqrt{t})$ rate as in the exact evaluation case for all values of γ_1 except $\gamma_1 = 0.25$. Recall that we impose a piecewise constant error threshold $\gamma_0 \ell^{-\gamma_1}$, fixed for each level $\ell(t)$. If $\gamma_1 < 0.5$, then the level lengths τ_ℓ decrease as $t \rightarrow \infty$; see Section 2.2.1. In our practical implementation, we can increase ℓ at most once per MCMC step and, therefore, when the level length is less than one step, the error threshold cannot decay quickly enough. In this case, the surrogate bias dominates the error and, as we see in Figure 3(a), the error decays more slowly as a function of MCMC steps.

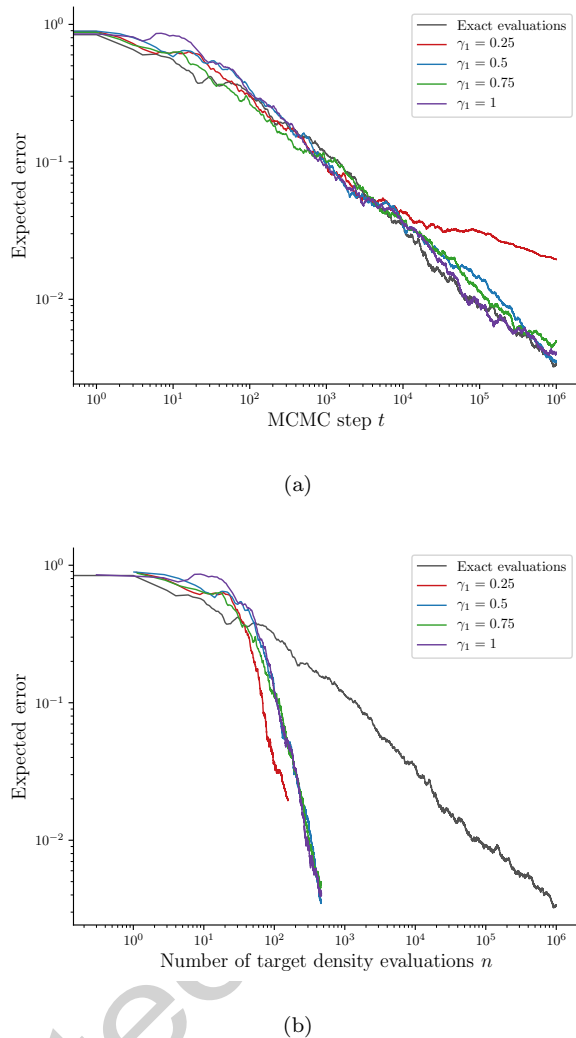


Fig. 3 The expected errors in variance (e_t and \tilde{e}_t) averaged over 50 MCMC chains as a function of (a) MCMC steps t and (b) the number of target density evaluations n . The LA-MCMC construction ensures that, if $\gamma_1 \geq 0.5$, the expected error decays at essentially the same $1/\sqrt{t}$ rate as in the exact evaluation case (Theorem 1). Since we do not need to evaluate the target density at every MCMC step, and in fact evaluate the target much less frequently over time (see Figure 5), the error decays much more quickly as a function of the number of target density evaluations.

Figure 3(b) shows that, since we do not evaluate the target density every MCMC step, the convergence of LA-MCMC as a function of the number of *density evaluations* n is in general much *faster* than in the exact evaluation case.

Approximating the target density with polynomials of higher degree p increases the efficiency of LA-MCMC; we explore this in Figure 4. Figure 4(a) shows that controlling the bias-variance trade-off ensures the error decay rate—as a function of the number of MCMC steps t —is the same regardless of p . Since the local error indicator is $\Delta(x)^{p+1}$, larger values of p achieve the same error threshold with larger radius $\Delta(x) < 1$. Using higher-degree polynomials therefore requires *fewer* target density evaluations, as shown in Figure 4(b). Yet higher-degree polynomials also require more evaluated points inside the local ball $\mathcal{B}_k(x)$. In this one-dimensional example, the local polynomial requires $p + 1$ points to interpolate, and here we choose $k = 2(p + 1)$ nearest neighbors to solve the regression problem (1). We thus see diminishing returns as p increases: higher order polynomials achieve the same accuracy with larger $\Delta(x)$ but require more target density evaluations within each ball $\mathcal{B}_k(x)$.

As the number of MCMC steps $t \rightarrow \infty$, the rate at which LA-MCMC requires new target density evaluations—i.e., the refinement rate—slows significantly. Figure 5 illustrates this pattern: the number of target density evaluations increases much more slowly than t . While the bias-variance trade-off ensures that the error decay rate remains $1/\sqrt{t}$, viewed in terms of the number of target density evaluations n , the picture is different. If

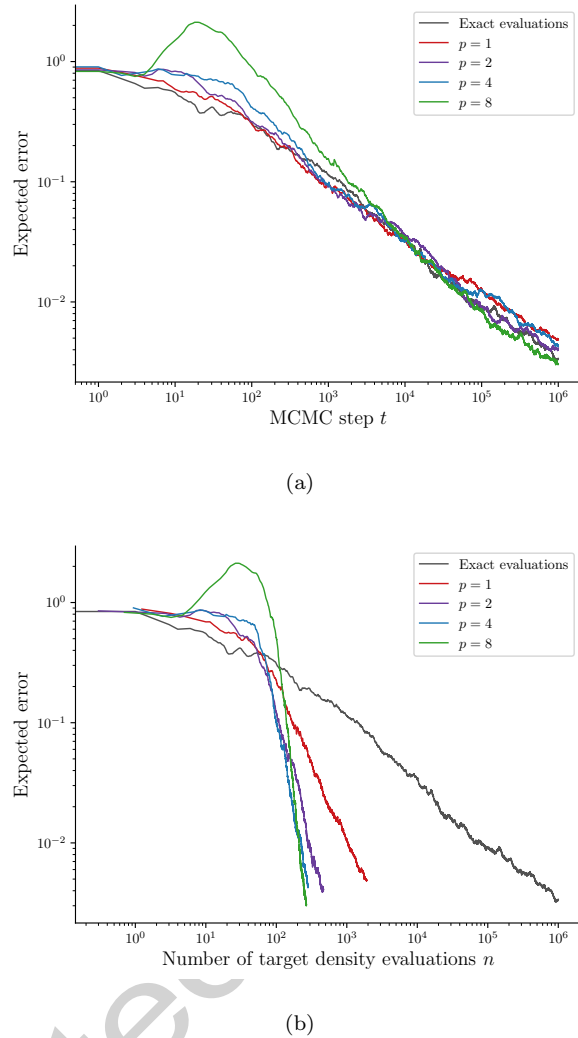


Fig. 4 The expected errors e_t and \tilde{e}_t averaged over 50 MCMC chains as a function of (a) MCMC steps t , and (b) the number of target density evaluations n . Controlling the bias-variance trade-off ensures that error decays at essentially the same $1/\sqrt{t}$ rate as in the exact evaluation case. As we increase the order of the local polynomial approximation, however, LA-MCMC requires fewer target density evaluations n to achieve the same error.

target density evaluations are the dominant computational expense—as is typical in many applications—then LA-MCMC generates samples more efficiently as $t \rightarrow \infty$.

4.2 Controlling tail behavior

The algorithmic parameter η (see (20)) controls how quickly LA-MCMC explores the tails of the target distribution. In the previous example, the tails of the log-target density decayed quadratically—like a Gaussian—and therefore we set $\eta = 0$. Now we consider the “banana-shaped” density

$$\log \pi(x) \propto -x_1^2 - (x_2 - 5x_1^2)^2 \quad (27)$$

for $x \in \mathbb{R}^2$. Figure 6 illustrates this target distribution; note the long tails in the x_2 direction. The two traces in Figure 7 show the mixing of an adaptive Metropolis (Haario et al., 2001) chain that uses exact density evaluations. We see that the chain does explore the tails in x_2 , but always returns to high probability regions; this is the expected and desired behavior of an MCMC algorithm for this target.

We can control how quickly LA-MCMC explores the tails of the distribution by varying η . Here, the local polynomial approximation of $\mathcal{L} = \log \pi$ might not (at any finite time) correctly capture the tail behavior; the




Fig. 5 The expected number of refinements $n(t)$, computed from 50 independent LA-MCMC chains, given different local polynomial degrees p . The refinement rate decreases as $t \rightarrow \infty$, making MCMC more efficient as t increases.




Fig. 6 The two dimensional density defined in (27).

surrogate model is therefore, in general, not an unnormalized probability density. Using LA-MCMC with no corrections ($\eta = 0$) thus allows the chain to wander into the tail without returning to the high-probability region, as shown in Figure 8(c). Increasing the tail correction parameter η biases the acceptance probability so that proposed points close to the centroid are more likely to be accepted and those that are farther are more likely to be rejected. As the number of MCMC steps $t \rightarrow \infty$, this biasing diminishes. The traces in Figure 8(a) show that setting $\eta > 0$ (here $\eta = 0.01$) prevents the chain from wandering too far into the tails, and Figure 9(a) shows that the resulting samples correctly characterize the target distribution.

If η is too large, however, then the correction will reduce the efficiency with which the chain explores the distribution's tail. The trace plot in Figure 8(b) shows that for $\eta = 5$, the chain appears to be mixing well. When we compare this chain to Figure 7 or Figure 8(a), though, we see that the chain is not spending any time in the tail of the distribution. Indeed, the density estimate in Figure 9(b) shows that the tails of the distribution are missing. Asymptotically, the tail correction decays and the algorithm *will* correctly characterize the target distribution for any $\eta > 0$. But too large an η can have an impact at finite time. In general, we choose the smallest η that prevents the chain from wandering into the distribution's tail.

4.3 Inferring aquifer transmissivity

Now we demonstrate the usefulness of LA-MCMC in a more computationally demanding example: inferring the spatially heterogeneous transmissivity of an unconfined aquifer. Although this example is physically motivated,




Fig. 7 Trace plots of an MCMC chain targeting (27), using exact evaluations.

our main goal is to highlight important aspects of the LA-MCMC algorithm and its performance. We choose this example because it is related to examples used in previous work (Conrad et al., 2018) and because similar models are used in the groundwater literature—see, for example, Matott (2012); Janetti et al. (2010); Al-Murad et al. (2018); Jardani et al. (2012); Willmann et al. (2007); Pool et al. (2015); None (2003). Though our model is idealized, it is not unreasonably different from many models that are used in practice. Moreover, explaining state-of-the-art groundwater models is well beyond the scope of this paper. We refer to existing work (e.g., Janetti et al. (2010)) for a detailed discussion of inferring transmissivity fields in hydrological applications, and here we focus on the computational demonstration of LA-MCMC. In this problem, the likelihood contains a set of coupled partial differential equations that model transport of a nonreactive tracer through the aquifer; the tracer concentration is then observed, with noise, at selected locations in the domain. Each likelihood evaluation is thus computationally intensive.

More specifically, the tracer is advected and diffused through the unconfined aquifer (a groundwater resource whose top boundary is not capped by an impermeable layer of rock/soil) by a steady state velocity field. We model the aquifer’s log-transmissivity, which determines the permeability of the soil, as a random field on the unit square $\mathcal{D} = [0, 1]^2$, parameterized as

$$\log \kappa(z) = \sum_{i=1}^d \kappa_i \sqrt{\lambda_i} e_i(z), \quad (28)$$

where κ_i are scalar coefficients, $z \in \mathcal{D}$, and $\{(\lambda_i, e_i(z))\}_{i=1}^d$ are the $d = 9$ leading eigenvalues/eigenfunctions of the integral operator associated with the squared exponential kernel $k(z, z') = \exp(-\|z - z'\|^2/2L^2)$, with $L = 0.1$. In other words, we have

$$\int_{\mathcal{D}} k(z, z') e_i(z') dz' = \lambda_i e_i(z). \quad (29)$$

Figure 10 shows the “true” log-transmissivity, where the true κ_i are drawn from a standard normal distribution. An “injection” well is located at $(x_w, y_w) = (0.45, 0.6)$. Given a value of the field $\kappa(z)$, we first compute the steady state hydraulic head h by solving

$$-\nabla \cdot (\kappa h \nabla h) = f_h \quad (30a)$$

on the domain $\mathcal{D} \ni z \equiv (x, y)$, with Dirichlet boundary conditions $h(x, y) = 1$ for $x = 0$, $x = 1$, $y = 0$, or $y = 1$. We set the source term above to:

$$f_h(x, y) = 250 \exp(-((x - x_w)^2 + (y - y_w)^2)/0.005),$$

which models hydraulic forcing due to well pumping. The velocity is then

$$u = -\kappa h \nabla h. \quad (31)$$

The hydraulic head and velocity resulting from the “true” parameter values (Figure 10) are illustrated in Figure 11. Given the steady-state velocity field u and an initial tracer concentration $c(x, y, 0) = 0$, advection



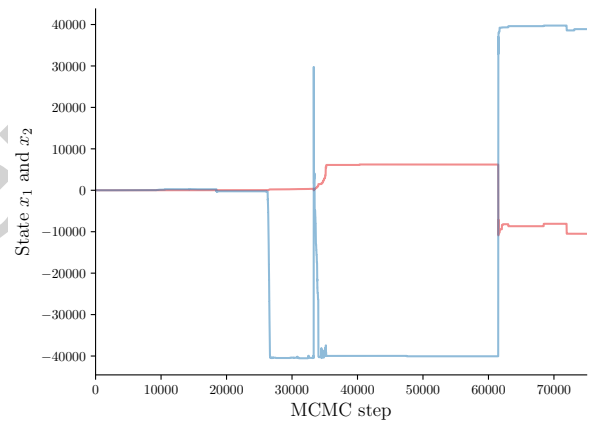
(a) $\eta = 0.01$ (b) $\eta = 5$ (c) $\eta = 0$

Fig. 8 Trace plots of LA-MCMC chains targeting (27) using different values of the tail correction parameter η . Other algorithmic parameters are $\gamma_0 = 2$, $\gamma_1 = 1$, $\bar{\Lambda} = \infty$, $\tau_0 = 1$, $k = 15$, $p = 2$, and $V(x) = \exp(0.25\|x\|^{0.75})$.




Fig. 9 Estimates of the target density in (27) constructed from 10^6 LA-MCMC samples. We vary $\eta \in \{0.01, 5\}$ and set $\gamma_0 = 2$, $\gamma_1 = 1$, $\bar{A} = \infty$, $\tau_0 = 1$, $k = 15$, $p = 2$, and $V(x) = \exp(0.25\|x\|^{0.75})$.

and diffusion of the tracer throughout the domain is modeled by a time-dependent concentration field $c(x, y, t)$ that obeys the following transport equation:

$$\frac{\partial c}{\partial t} + \nabla \cdot ((d_m \mathbf{I} + d_\ell u u^T) \nabla c) - u^T \nabla c = -f_t \quad (32)$$

with $d_m = 0.05$, $d_\ell = 0.001$, and

$$f_t(x, y) = \exp(-((x - x_w)^2 + (y - y_w)^2)/0.005). \quad (33)$$

The tracer forcing $f_t(x, y)$ models tracer leakage into the domain. The concentration field at time $t = 1$ is shown in Figure 12.

Our goal is to infer the log-transmissivity parameters $(\kappa_i)_{i=1}^d$ given observations of the tracer concentration. We endow each κ_i with an independent standard Gaussian prior distribution, $\kappa_i \sim N(0, 1)$. We could also infer other model parameters (e.g., d_m , d_ℓ , or the forcing functions f_h and f_t) by including them in the set of inferred parameters. For the sake of this example, however, we assume that these parameters are known and therefore fixed. Data are collected on a 10×10 array of points evenly spaced in $x \in [0, 1]$ and $y \in [0, 1]$. These locations are marked in Figure 12. We make observations at 10 evenly spaced times $t \in [0, 1]$. This is a relatively large amount of data—typical observation arrays only partially observe the aquifer, although there are often still hundreds of observations (see, e.g., Al-Murad et al. (2018); Janetti et al. (2010)). The large data set here is




Fig. 10 The “true” log-transmissivity for an unconfined aquifer model. The coefficients in (28) are sampled from a standard Gaussian distribution. The diamond denotes a well location.




Fig. 11 The steady-state hydraulic head and velocity field computed by solving (30) and (31) given the “true” parameter values are sampled from a standard Gaussian distribution. The white diamond represents the location of a well in the aquifer.

relatively informative, such that the posterior distribution is quite concentrated relative to the prior and the inference problem is thus more computationally challenging. The observations are modeled as

$$y = f(\kappa) + \varepsilon, \quad (34)$$

where $f(\kappa)$ is the *forward model* induced by the partial differential equations above, mapping the log-transmissivity parameters $\kappa \equiv (\kappa_i)_{i=1}^d$ to a prediction of the tracer concentration at the chosen sensor locations/times. The sensor noise is $\varepsilon \sim N(0, 10^{-4}I)$, yielding the conditional distribution

$$y|\kappa \sim N(f(\kappa), 10^{-4}I). \quad (35)$$

The posterior density then follows from Bayes’ rule:

$$\pi(\kappa|y) \propto \pi(y|\kappa)\pi(\kappa). \quad (36)$$

To avoid a so-called “inverse crime,” where data are generated from the same numerical model used to perform inference (Kaipio and Somersalo, 2006), we generate the data y by solving the forward model with a well-refined 65×65 numerical discretization using the “true” parameters, and then perform the inference using a 25×25 numerical discretization of the forward model.

Evaluating the forward model at every MCMC step is computationally prohibitive. The run time of the model is $\mathcal{O}(1)$ seconds on an Intel Core i7-7700 CPU at 3.60GHz, with some variability depending on the




Fig. 12 Tracer concentration computed at time $t = 1$ by solving (32) given the “true” parameter values in Figure 10 and the steady state velocity field shown in Figure 11. The white diamond represents the location of a well in the aquifer, and the circles are sensor locations used in the Bayesian inference problem.

value of log-transmissivity parameters κ . Generating 10^6 samples thus takes 5–10 days of computation time. LA-MCMC is essential to making this Bayesian inference problem computationally feasible. We note that high performance and parallel computing resources could certainly reduce the computational cost of each model evaluation. Parallel computing can even enable the shared construction of surrogate models using concurrent chains, as described in our previous work (Conrad et al., 2018). Since our goal here is to demonstrate the impact of new local approximation strategies, however, we focus on a serial implementation. A useful hardware- and implementation-independent measure of computational cost is the number of refinements (i.e., expensive model/target density evaluations) in a given run, which we discuss below.

Recall that the local polynomial surrogate (1) requires computing $q = \dim(\mathcal{P})$ coefficients, and in the total-degree setting of our theory, we have $q = \binom{d+p}{p}$. This quantity grows rapidly with the dimension d of the parameters for $p > 1$. Retaining a total-degree construction, we could set $p = 0$ and thus include only a constant term in \mathcal{P} . This results in a constant surrogate model within each ball; now the number of terms in the local polynomial approximation is always one, independent of the parameter dimension. Similarly, setting $p = 1$ lets q grow only linearly with d . Yet in the one-dimensional example of Section 4.1, we showed that including higher-degree terms in the approximation can reduce the number of expensive target density evaluations required to achieve a given accuracy (see Figure 4).

We implement these total-degree local polynomial approximations but also compare them empirically with other choices of polynomial approximation space—in particular, truncations that retain higher-order terms more selectively. A common practice in high-dimensional approximation is to employ *sparse truncations* of the relevant index set (Blatman and Sudret, 2011). Let $\psi_s(\kappa_i)$ be a univariate polynomial of degree s in κ_i . (Typically, we choose an orthogonal polynomial family, but this choice is immaterial.) Each d -variate polynomial basis function $\phi_{\alpha}(\kappa) := \prod_{i=1}^d \psi_{\alpha_i}(\kappa_i)$ is defined by a multi-index $\alpha \in \mathbb{N}_0^d$, where the integers $\alpha_i \geq 0$ are components of α . Now define the ℓ^ν norm for $\nu > 0$ (a quasi-norm for $0 < \nu < 1$):

$$\|\alpha\|_\nu = \left(\sum_{i=1}^d \alpha_i^\nu \right)^{1/\nu}. \quad (37)$$

Given a maximum degree p , we can build our *local* polynomial approximations in the polynomial space $\mathcal{P}_\nu^{p,d}$ spanned by basis functions with multi-indices $\|\alpha\|_\nu \leq p$. When $\nu = 1$, this recovers a total-degree expansion, which is consistent with the error bounds in Section 2.1.1. However, the dimensionality of $\mathcal{P}_\nu^{p,d}$ decreases as $\nu \rightarrow 0$. We define the special case of $\nu = 0$ to be a polynomial space with no cross terms: each multi-index has one nonzero element, and that element is at most p .

We now apply the LA-MCMC algorithm, as well as an exact MCMC algorithm for comparison, to the posterior distribution of the tracer transport problem. All MCMC chains use the same (fixed) random-walk proposal $q_t = \mathcal{N}(\kappa_t, 0.005I)$, where κ_t is the current state of the chain. We tested different configurations of




Fig. 13 The diagonal and lower diagonal subplots show the one- and two-dimensional marginal distributions of the posterior distribution (36) estimated using LA-MCMC with $\nu = 0.75$, $\gamma_0 = 2$, $\gamma_1 = 0.5$, $\bar{\Lambda} = \infty$, $\tau_0 = 1$, $\eta = 0$, $k = 75$, $p = 3$, and $V(x) = \exp(\|\kappa - \kappa_{\text{MAP}}\|)$. Here, $\kappa_{\text{MAP}} = \arg \max_{\kappa} \pi(\kappa|y)$ is the posterior mode. The subplot on the upper right shows the two-dimensional prior marginal for any pair of parameters.

LA-MCMC, parameters of which are given in Table 1. The Lyapunov function V used in the tail correction is defined using the posterior mode, $\bar{\kappa} = \arg \max_{\kappa} \pi(\kappa|y)$, which we obtain by maximizing the log-posterior density using optimization algorithms in `nlopt` (Johnson, 2014). The rows in Table 1 with $\nu = 1$ correspond to using total-degree linear and quadratic local polynomial approximations (consistent with our theory). Setting $\nu < 1$ defines a sparse approximation. The required number of nearest neighbors k is chosen to be slightly more than the number of points $q = \dim(\mathcal{P}_\nu^{p,d})$ required to interpolate. Reducing ν reduces q and k . Chains produced by all these algorithmic configurations yield essentially identical posterior estimates. A trace plot of selected states from the LA-MCMC chain corresponding to the $p = 3$ configuration is shown in Figure 14, and compared to a trace plot of the exact-model chain. One- and two-dimensional marginals of the posterior distribution are shown in Figure 13. While this figure is generated with $\nu = 0.75$ and $p = 3$, the results are essentially identical for the other algorithmic configurations. An exact chain of 10^6 steps took about 6 days to complete, whereas the LA-MCMC chains of the same length with quadratic and cubic surrogate models took roughly 6 to 10 hours, and with linear models took just over one day. These timings are somewhat imprecise, however, as other computational tasks were competing for resources on the same workstation (a Intel Core i7-7700 CPU at 3.60GHz) used for our runs.

degree p	ν	k	# to interpolate (q)	γ_0
Exact	—	—	—	—
1	1	20	10	1.9
2	1	65	55	1
3	0.75	75	64	1

Table 1 Parameter configurations for Algorithm 1, used to generate samples from the posterior distribution (36) of the PDE/tracer transport problem. In all cases $\gamma_1 = 0.5$, $\eta = 0$, $\log V(\kappa) = \|\kappa - \kappa_{\text{MAP}}\|$, and $\bar{\Lambda} = \infty$. Here, $\kappa_{\text{MAP}} = \arg \max_{\kappa} \pi(\kappa|y)$ is the posterior mode, which we obtain by maximizing the log-posterior density using `nlopt` (Johnson, 2014).

Though all these chains are successful in characterizing the posterior distribution, the required number of (expensive) likelihood evaluations—and hence the overall computational cost—depend on a non-trivial relationship between the accuracy of the local approximations and the number of nearest neighbors required to define them. As noted earlier, though, the number of such likelihood evaluations is a robust and translatable measure of computational cost, independent of implementation and hardware, and also most meaningful in the setting where model evaluations dominate all other costs of the MCMC machinery. Figure 15 shows the number of likelihood evaluations n as a function of the number of MCMC steps t for each experiment in Table 1. We see that the quadratic and cubic local approximations (purple and blue lines) require far fewer expensive likelihood evaluations. In all cases, we see a reduction relative to the number of likelihood evaluations that would be required by exact MCMC. However, this improvement is drastically increased, by orders of magnitude, when choosing $p > 1$.

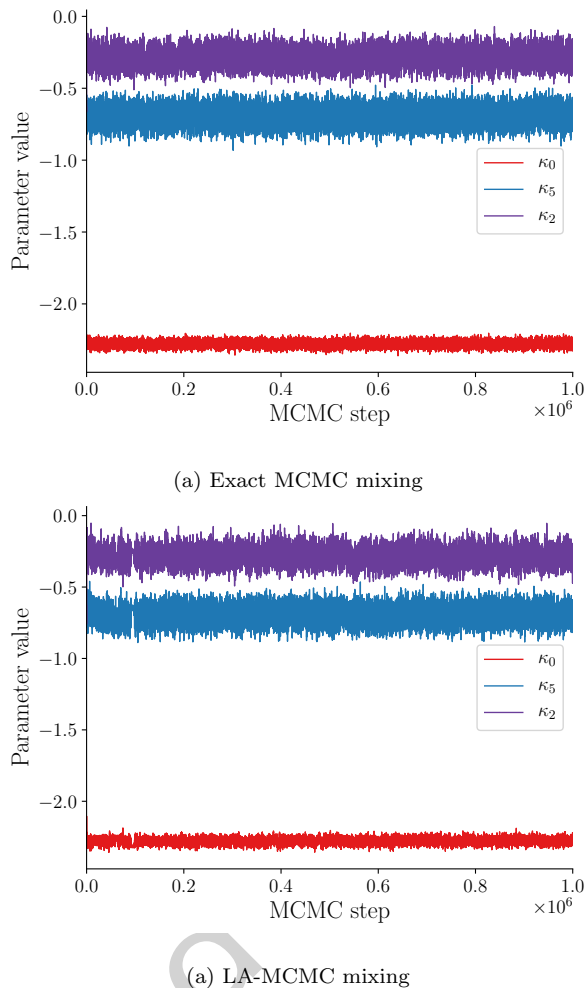
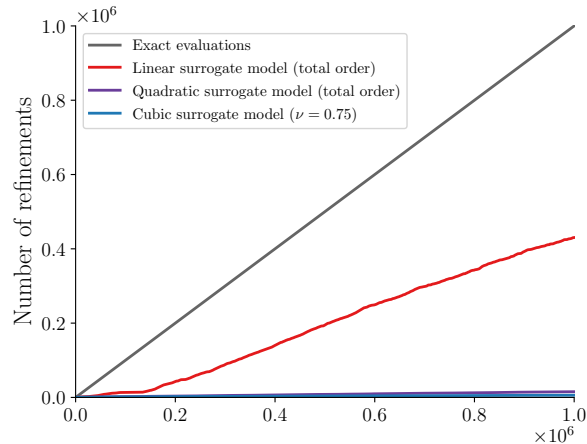


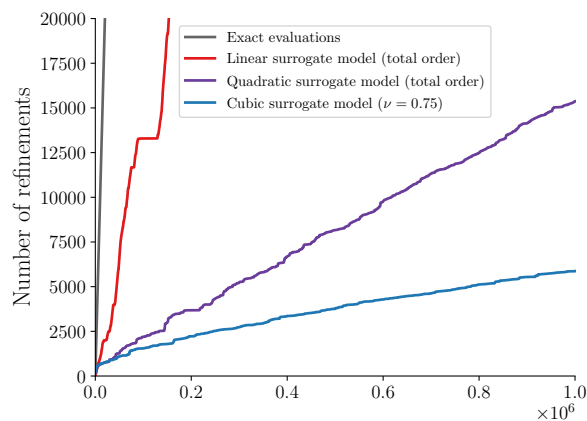
Fig. 14 PDE/tracer transport problem: trace plot of three parameters of (a) an exact MCMC chain and (b) an LA-MCMC chain using a locally cubic ($p = 3$) surrogate model with $\nu = 0.75$. Additional parameters are defined on the last row of Table 1.

It is important to note that our approach to controlling the bias-variance tradeoff, and thus the error thresholds for triggering refinement via (9) and (15), are derived for expansions of total degree p and thus in principle applicable only for $\nu = 1$. We use the same rules here for $\nu < 1$, but this approach cannot guarantee that bias and variance decay at the same rate. Indeed, a full analysis of the local sparse approximations may depend on understanding the magnitudes of mixed derivatives of the target function g ; this can become quite problem-specific, and we defer such an investigation to future work. The empirical results for $\nu < 1$ here are intended to be practical and exploratory.

The choice of local approximation not only affects the overall computational cost of each chain, but also its mixing. Figure 16(a) shows the number of effectively independent samples (i.e., the effective sample size (ESS) (Wolff et al., 2004)) produced by each configuration, as a function of the number of MCMC steps t . The chain using exact density evaluations generates effectively independent samples fastest *as a function of* t . Mixing is slightly faster with cubic local approximations than with local quadratic or linear approximations. Figures 16(b)–(c), however, show a non-trivial relationship between complexity of the surrogate model and how quickly it generates ESS *as a function of the number of model evaluations* n . Chains employing local linear approximations ($p = 1$) (red line) do not mix as efficiently as the sparse local cubic approximation or the total-degree quadratic approximation (blue and purple lines). Because the number of model evaluations in the $p = 1$ case is not so drastically reduced over the exact case (see Figure 15(a)), poorer mixing leads to worse performance in the overall metric of ESS per model evaluation. On the other hand, the $p = 3$ case achieves the same ESS with



(a) Refinements per MCMC step



(b) Refinements per MCMC step (rescaled)


Fig. 15 PDE/tracer transport problem: number of likelihood evaluations (i.e., refinements) as a function of MCMC steps, for LA-MCMC chains with parameters defined in Table 1. Panel (a) shows the total number of refinements for all chains, and panel (b) rescales the vertical axis to emphasize how many fewer model evaluations LA-MCMC requires than the exact case. For local surrogate models with $p \geq 2$, we see significant reductions (of three to four orders of magnitude) over the number of likelihood evaluations required by exact MCMC.

5000 model evaluations that the exact chain achieves with 5×10^5 evaluations, an improvement of two orders of magnitude.


Recall that Figure 13 shows strong posterior correlations for this problem. We speculate that log-likelihood approximations that include cross terms (i.e., $p > 1$) can more easily approximate the true posterior density and, therefore, significantly outperform local approximations that do not include these factors. In theory, even a locally constant model will yield convergence as $\Delta \rightarrow 0$. In practice, we find that the choice of basis for the local approximation may have a substantial impact. While $p = 2$ or $p = 3$ are good default choices, a method for adaptively selecting polynomial basis functions as we learn about correlations in the posterior could possibly improve the algorithm. We leave this to future work.

5 Conclusions


This paper investigated the design and convergence properties of local approximation MCMC (LA-MCMC) algorithms, which exploit regularity of the target density to build and refine local polynomial surrogate models during sampling. When density evaluations are expensive, these algorithms can reduce the computational cost of MCMC by orders of magnitude.



(a) ESS per MCMC step



(b) ESS per model evaluation



(c) ESS per model evaluation (rescaled)

Fig. 16 PDE/tracer transport problem, computational efficiency for LA-MCMC chains with parameters defined in Table 1. (a) effective sample size (ESS) as a function of the number of MCMC steps; (b-c) ESS as a function of the number of likelihood evaluations (refinements). Panel (c) rescales the horizontal axis to more clearly show how quickly LA-MCMC with cubic and quadratic surrogate models generates independent samples as a function of the number of refinements.

We introduce a new LA-MCMC algorithm with two key features: (1) a rate-optimal refinement strategy that balances the decay of surrogate-induced bias with the decay of Monte Carlo variance; and (2) a modification to the Metropolis acceptance step of the algorithm, which helps ensure stability of the resulting chain. Together these features enable broader applicability and stronger theoretical guarantees than previous efforts. Earlier versions of LA-MCMC (Conrad et al., 2016) employed a cross-validation heuristic in conjunction with random refinements of the surrogate model; the latter made possible a guarantee of ergodicity on compact state spaces or under strong tail conditions. In the present effort, we dispense with both cross-validation and random refinement, using instead the more principled bias-variance control strategy mentioned above. Not only is the resulting sampling method asymptotically exact; here we also show that the error in LA-MCMC estimates decays at approximately the expected $1/\sqrt{T}$ rate, where T is the number of MCMC steps. Moreover, this rate holds without the restrictive tail conditions that were required, both in theory and in practice, for stability and convergence of previous versions of the algorithm. Using our refinement strategy, we also observe that the rate at which new target density evaluations are demanded decays with T . This means that as a function of the number of target density evaluations, convergence is *faster* than that of standard geometrically ergodic MCMC, accelerating as T increases, in accordance with the underlying function approximation machinery.

Our numerical examples demonstrate both the predicted convergence rates and stability of the sampler for target distributions with heavier tails. We also show how LA-MCMC can be employed in a computationally intensive inference problem, where evaluations of the likelihood require solving a coupled set of nonlinear partial differential equations; in this example, motivated by inverse problems in groundwater hydrology, LA-MCMC reduces the number of expensive likelihood evaluations by roughly two orders of magnitude.

While the present algorithm is well suited to expensive target densities on parameter spaces of moderate dimension (e.g., $d = 9$ in the current PDE example), it becomes more challenging to apply in higher dimensions, even $d = 15$ in our experience. The fundamental reason is that generic isotropic local polynomial approximation is subject to the curse of dimensionality, with at least $\binom{d+p}{p}$ nearest neighbors required to evaluate the surrogate at a given point. Hence a larger number of target density evaluations are required as d increases, and the benefits of function approximation over sampling begin to diminish. As with any function approximation problem, the key to handling higher dimensions is to exploit some form of structure, e.g., anisotropy or sparsity. For instance, future work could investigate the use of sparse local polynomial approximations as way of reducing the complexity of the local surrogate in higher dimensions. The preliminary empirical study in this paper shows that sparse local polynomial approximation can be helpful, but further analysis—perhaps building on the analysis of *global* sparse polynomial approximations of likelihood functions in Bayesian inverse problems (Schillings and Schwab, 2014) and exploiting anisotropy—is needed to understand the convergence properties of such approximations within LA-MCMC. Other methods of explicit dimension reduction (Zahm et al., 2022) may also be useful in this setting; the idea would be to find a low-dimensional subspace in which the posterior departs most strongly from the prior and to apply our LA-MCMC machinery to the log-likelihood function only within that subspace. This approach has antecedents in Cui et al. (2016). In the present paper, however, we have focused on the essential approximation and sampling machinery of LA-MCMC and on the analysis of its convergence. Combinations with other techniques, while practically important, are left to future work.

An open source implementation of the LA-MCMC algorithm introduced in this paper is available as part of the MUQ library (<http://muq.mit.edu>).

A Theoretical results

We include all theoretical results from the paper. Throughout this section, we only deal with the case that the underlying proposal distribution q_t does not change with t . To deal with typical small adaptations, we believe that the following framework can be combined with, e.g., the approach of Roberts and Rosenthal (2007), but this would result in a significantly longer paper and these adaptations are not central to our approach.

A.1 General bounds on non-Markovian approximate MCMC algorithms

Proceeding more formally, let $\{\hat{X}_t, \hat{K}_t, \mathcal{F}_t\}_{t \geq 0}$ be a triple satisfying:

1. $\{\hat{X}_t\}_{t \geq 0}$ is a sequence of *random variables* on \mathbb{R}^d ;
2. $\{\hat{K}_t\}_{t \geq 0}$ is a (typically random) sequence of *transition kernels* on \mathbb{R}^d ;
3. $\{\mathcal{F}_t\}_{t \geq 0}$ is a *filtration*, and $\{\hat{X}_t, \hat{K}_t\}_{t \geq 0}$ is *adapted* to this filtration;

4. the three agree in the sense that

$$\mathbb{P}[X_{s+1} \in A | \mathcal{F}_s] = \hat{K}_s(\hat{X}_s, A) \quad (38)$$

for all $s \geq 0$ and all measurable A . Note that, in particular, both left- and right-hand sides are \mathcal{F}_s -measurable random variables in $[0, 1]$.

In practice, \mathcal{F}_s is generated by our sequence of approximations to the true log-target. We use the following quantitative assumptions:

Assumption 4 (*Lyapunov inequality*). There exists $V : \mathbb{R}^d \rightarrow [1, \infty)$ and constants $0 < \alpha \leq 1$ and $0 \leq \beta < \infty$ so that

$$(\hat{K}_s V)(\hat{X}_s) \leq (1 - \alpha)V(\hat{X}_s) + \beta \quad (39)$$

and

$$(KV)(x) \leq (1 - \alpha)V(x) + \beta \quad (40)$$

for all $s \geq 0$. The second inequality should hold deterministically; note that this is an \mathcal{F}_s -measurable event.

Assumption 5 (*Good approximation*). Let Assumption 1 or 4 hold. There exists a monotonically decreasing function $\delta : [0, \infty) \rightarrow [0, 0.5]$ so that

$$\|K(\hat{X}_s, \cdot) - \hat{K}_s(\hat{X}_s, \cdot)\|_{TV} \leq \delta(s)V(x) \quad (41)$$

for all $s \geq 0$ and $x \in \mathbb{R}^d$. Again, this inequality should hold deterministically, which is an \mathcal{F}_s -measurable event. For notational convenience, we define $\delta(s) = \delta(0)$ for all $s < 0$.

A.1.1 Initial coupling bounds

The following is our main technical lemma. It is *not* monotone in the time s ; this will be remedied in applications.

Lemma 1 Let Assumptions 2, 4, and 5 hold. There exists a constant $0 < C < \infty$ depending only on α , β , R , and γ so that for $x \in \mathbb{R}^d$ and triple $\{\hat{X}_t, \hat{K}_t, \mathcal{F}_t\}_{t \geq 0}$ started at $\hat{X}_0 = x$, we have

$$\|\mathbb{P}[\hat{X}_s \in \cdot] - \pi(\cdot)\|_{TV} \leq \begin{cases} 1 & \text{if } s \leq C_0 \\ C\delta(0)s & \text{if } s > C_0, \end{cases}$$

where $C_0 = C \log(\delta(0)^{-1}V(x))$.

Proof Define the “small set”

$$\mathcal{C} = \left\{ y : V(y) \leq \frac{4\beta}{\alpha} \right\} \quad (42)$$

and the associated hitting time

$$\tau_{\mathcal{C}} = \min \{ t : \hat{X}_t \in \mathcal{C} \}. \quad (43)$$

Denote by $T_b \geq 0$ a “burn-in” time whose value will be fixed toward the end of the proof.

By the triangle inequality, for all measurable $A \subset \mathbb{R}^d$

$$\begin{aligned} |\mathbb{P}[\hat{X}_s \in A] - \pi(A)| &\leq |\mathbb{P}[\hat{X}_s \in A, \tau_{\mathcal{C}} \leq T_b] - \pi(A)\mathbb{P}[\tau_{\mathcal{C}} \leq T_b]| + |\mathbb{P}[\hat{X}_s \in A, \tau_{\mathcal{C}} > T_b] - \pi(A)\mathbb{P}[\tau_{\mathcal{C}} > T_b]| \\ &\leq |\mathbb{P}[\hat{X}_s \in A, \tau_{\mathcal{C}} \leq T_b] - \pi(A)\mathbb{P}[\tau_{\mathcal{C}} \leq T_b]| + \mathbb{P}[\tau_{\mathcal{C}} > T_b] \end{aligned}$$

To bound the first term, note that

$$\begin{aligned} |\mathbb{P}[\hat{X}_s \in A, \tau_{\mathcal{C}} \leq T_b] - \pi(A)\mathbb{P}[\tau_{\mathcal{C}} \leq T_b]| &\leq \sup_{y \in \mathcal{C}, 0 \leq u \leq T_b} |\mathbb{P}[\hat{X}_s \in A | \hat{X}_u = y, \tau_{\mathcal{C}} = u] - \pi(A)| \\ &\leq \sup_{y \in \mathcal{C}, 0 \leq u \leq T_b} |\mathbb{P}[\hat{X}_s \in A | \hat{X}_u = y, \tau_{\mathcal{C}} = u] - K^{s-u-1}(y, A)| \\ &\quad + \sup_{y \in \mathcal{C}, 0 \leq u \leq T_b} |K^{s-u-1}(y, A) - \pi(A)|. \end{aligned} \quad (44)$$

Assumption 2 gives

$$\sup_{y \in \mathcal{C}, 0 \leq u \leq T_b} |K^{s-u-1}(y, A) - \pi(A)| \leq R\gamma^{s-T_b-1}$$

and Assumption 5 gives

$$\sup_{y \in \mathcal{C}, 0 \leq u \leq T_b} |\mathbb{P}[\hat{X}_s \in A | \hat{X}_u = y, \tau_{\mathcal{C}} = u] - K^{s-u-1}(y, A)| \leq \sup_{y \in \mathcal{C}, 0 \leq u \leq T_b} \sum_{t=u}^{s-1} \delta(t)(K^{s-t-1}V)(y).$$

Furthermore,

$$\sup_{y \in \mathcal{C}, 0 \leq u \leq T_b} \sum_{t=u}^{s-1} \delta(t)(K^{s-t-1}V)(y) + R\gamma^{s-T_b-1} \leq \sup_{y \in \mathcal{C}} \sum_{t=0}^{s-1} \delta(t)(K^{s-t-1}V)(y) + R\gamma^{s-T_b-1}.$$

Substituting this back into (44) and by Assumption 4,

$$|\mathbb{P}[\hat{X}_s \in A, \tau_C \leq T_b] - \pi(A)\mathbb{P}[\tau_C \leq T_b]| \leq \sup_{y \in C} \sum_{t=0}^{s-1} \delta(t)((1-\alpha)^{s-t-1}V(y) + \beta/\alpha) + R\gamma^{s-T_b-1} \quad (45a)$$

$$\leq \sum_{t=0}^{s-1} \delta(t)((1-\alpha)^{s-t-1}4\beta/\alpha + \beta/\alpha) + R\gamma^{s-T_b-1} \quad (45b)$$

$$\leq \delta(0)5s\beta/\alpha + R\gamma^{s-T_b-1}. \quad (45c)$$

To bound the second term in (44), recall from Assumption 4 that

$$\begin{aligned} \mathbb{E}[V(\hat{X}_{t+1})\mathbf{1}_{\tau_C > t} | \mathcal{F}_t] &\leq ((1-\alpha)V(\hat{X}_t) + \beta)\mathbf{1}_{V(\hat{X}_t) > 4\beta/\alpha} \\ &\leq \left(1 - \frac{3\alpha}{4}\right)V(\hat{X}_t) + \left(\beta - \frac{\alpha}{4}V(\hat{X}_t)\right)\mathbf{1}_{V(\hat{X}_t) > 4\beta/\alpha} \\ &\leq \left(1 - \frac{3\alpha}{4}\right)V(\hat{X}_t) + \left(\beta - \frac{\alpha}{4}\frac{4\beta}{\alpha}\right)\mathbf{1}_{V(\hat{X}_t) > 4\beta/\alpha} \\ &\leq \left(1 - \frac{3\alpha}{4}\right)V(\hat{X}_t) \end{aligned}$$

for all $t \geq 0$. Iterating, we find by induction on t that

$$\mathbb{E}[V(\hat{X}_t)\mathbf{1}_{\tau_C > t}] \leq \left(1 - \frac{3\alpha}{4}\right)^t V(\hat{X}_0) = \left(1 - \frac{3\alpha}{4}\right)^t V(x).$$

Thus, by Markov's inequality,

$$\mathbb{P}[\tau_C > T_b] = \mathbb{P}\left[V(\hat{X}_{\tau_C})\mathbf{1}_{\tau_C > T_b} > \frac{4\beta}{\alpha}\right] \leq \frac{\alpha}{4\beta} \left(1 - \frac{3\alpha}{4}\right)^{T_b} V(x). \quad (46)$$

Combining (45) and (46), we have shown that

$$|\mathbb{P}[\hat{X}_s \in A] - \pi(A)| \leq \delta(0)\frac{5s\beta}{\alpha} + R\gamma^{s-T_b-1} + \frac{\alpha}{4\beta} \left(1 - \frac{3\alpha}{4}\right)^{T_b} V(x). \quad (47)$$

Finally, we can choose T_b . Set

$$T(s) = \max\left\{t : R\gamma^{s-t-1}, \frac{\alpha}{4\beta} \left(1 - \frac{3\alpha}{4}\right)^t V(x) \leq \frac{1}{2}\delta(0)\right\}$$

and define $T_b = \lfloor T(s) \rfloor$ when $0 < T(s) < \infty$ and $T_b = 0$ otherwise. Define $S = \min\{s : T(s) \in (0, \infty)\}$. Noting that $S = \Theta(\log(\delta(0)^{-1}) + \log(V(x)))$ for fixed α, β, R , and γ completes the proof. \square

We strengthen Lemma 1 by first observing that, if \hat{X}_0 satisfies $V(\hat{X}_0) \leq 4\beta/\alpha$, then

$$\mathbb{E}[V(\hat{X}_1)] \leq (1-\alpha)\mathbb{E}[V(\hat{X}_0)] + \beta \leq (1-\alpha)\frac{4\beta}{\alpha} + \beta \leq \frac{4\beta}{\alpha}. \quad (48)$$

Thus, by induction, if $\mathbb{E}[V(\hat{X}_0)] \leq 4\beta/\alpha$, then $\mathbb{E}[V(\hat{X}_s)] \leq 4\beta/\alpha$ for all time $s \geq 0$. Using this (and possibly relabelling the starting time to the quantity denoted by $T_0(s)$), Lemma 1 has the immediate slight strengthening:

Lemma 2 *Let Assumptions 2, 4, and 5 hold. There exists a constant $0 < C < \infty$ depending only on α, β, R , and γ so that for all starting distributions μ on \mathbb{R}^d with $\mu(V) \leq 4\beta/\alpha$ and triple $\{\hat{X}_t, \hat{K}_t, \mathcal{F}_t\}_{t \geq 0}$ started at $\hat{X}_0 \sim \mu$, we have*

$$\|\mathbb{P}[\hat{X}_s \in \cdot] - \pi(\cdot)\|_{TV} \leq \begin{cases} 1, & s \leq C_0 \\ C\delta(T_0(s)) \log(\delta(T_0(s))^{-1}), & s > C_0, \end{cases}$$

where $C_0 = C\delta(0) \log(\delta(0)^{-1})$ and

$$T_0(s) = s - C\delta(0) \log(\delta(0)^{-1}). \quad (49)$$

A.1.2 Application to bounds on mean-squared error

Recall that $f : \mathbb{R}^d \rightarrow [-1, 1]$ with $\pi(f) = 0$. We apply Lemma 2 to obtain the following bound on the Monte Carlo bias:

Lemma 3 (*Bias estimate*). *Let Assumptions 2, 4, and 5 hold. There exists a constant $0 < C < \infty$ depending only on α, β, R , and γ so that for all starting points $x \in \mathbb{R}^d$ with $V(x) \leq 4\beta/\alpha$ and triple $\{\hat{X}_t, \hat{K}_t, \mathcal{F}_t\}_{t \geq 0}$ started at $\hat{X}_0 = x$ we have*

$$\left| \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T f(\hat{X}_t) \right] \right| \leq \begin{cases} 1, & T \leq C_0 \\ \frac{C_0}{T} + \frac{C}{T} \sum_{s=C_0}^T \delta(T_0(s)) \log(\delta(T_0(s))^{-1}) & T > C_0, \end{cases}$$

where $C_0 = C\delta(0) \log(\delta(0)^{-1})$.

Proof In the notation of Lemma 2, we have for $T > C_0$ sufficiently large

$$\begin{aligned} \left| \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T f(\hat{X}_t) \right] \right| &\leq \frac{1}{T} \sum_{t=1}^{C_0} |\mathbb{E}[f(\hat{X}_t)]| + \frac{1}{T} \sum_{t=C_0}^T |\mathbb{E}[f(\hat{X}_t)]| \\ &\leq \frac{C_0}{T} + \frac{C}{T} \sum_{s=C_0}^T \delta(T_0(s)) \log(\delta(T_0(s))^{-1}). \end{aligned}$$

□

We have a similar bound for the Monte Carlo variance:

Lemma 4 (*Covariance estimate*). *Let Assumptions 2, 4, and 5 hold. There exists a constant $0 < C < \infty$ depending only on α, β, R , and γ so that for all starting points $x \in \mathbb{R}^d$ with $V(x) \leq 4\beta/\alpha$ and triple $\{\hat{X}_t, \hat{K}_t, \mathcal{F}_t\}_{t \geq 0}$ started at $\hat{X}_0 = x$ we have*

$$\sqrt{|\mathbb{E}[f(\hat{X}_s)f(\hat{X}_t)] - \mathbb{E}[f(\hat{X}_s)]\mathbb{E}[f(\hat{X}_t)]|} \leq \begin{cases} 1, & m(s, t) \leq C_0 \\ C\delta(T_0) \log(\delta(T_0)^{-1}), & m(s, t) > C_0, \end{cases}$$

where $m(s, t) = \min(s, t, |t - s|)$, $T_0 = T_0(m(s, t))$ is as in (49), and $C_0 = C\delta(0) \log(\delta(0)^{-1})$ as before.

Proof By the triangle inequality

$$|\mathbb{E}[f(\hat{X}_s)f(\hat{X}_t)] - \mathbb{E}[f(\hat{X}_s)]\mathbb{E}[f(\hat{X}_t)]| \leq |\mathbb{E}[f(\hat{X}_s)f(\hat{X}_t)]| + |\mathbb{E}[f(\hat{X}_s)]\mathbb{E}[f(\hat{X}_t)]|.$$

As above, applying Lemma 2 completes the proof.

□

The above bias and variance estimates immediate imply our main theorem on the total error of the Monte Carlo estimator:

Theorem 2 *Let Assumptions 2, 4, and 5 hold. There exists a constant $0 < C < \infty$ depending only on α, β, R , and γ so that for all starting points $x \in \mathbb{R}^d$ with $V(x) \leq 4\beta/\alpha$ and triple $\{\hat{X}_t, \hat{K}_t, \mathcal{F}_t\}_{t \geq 0}$ started at $\hat{X}_0 = x$ we have*

$$\mathbb{E} \left[\left(\frac{1}{T} \sum_{t=1}^T f(\hat{X}_t) \right)^2 \right] \leq \frac{2C_0}{T^2} \sum_{s=1}^T C(s) + \frac{3}{T} \sum_{s=1}^T C(s)^2,$$

where $C(s) = C\delta(T_0(s)) \log(\delta(T_0(s))^{-1})$ and we write $C_0 = C_0(0)$.

Proof We calculate

$$\begin{aligned} \mathbb{E} \left[\left(\frac{1}{T} \sum_{t=1}^T f(\hat{X}_t) \right)^2 \right] &= T^{-2} \left[\sum_{t=1}^T \mathbb{E}[f(\hat{X}_t)^2] + \sum_{s, t : m(s, t) < C_0} \mathbb{E}[f(\hat{X}_s)f(\hat{X}_t)] + \sum_{s, t : m(s, t) \geq C_0} \mathbb{E}[f(\hat{X}_s)f(\hat{X}_t)] \right] \\ &\leq T^{-2} \left[\sum_{s=1}^T C_0(s) + C_0 \sum_{s=1}^T C_0(s) + 3T \sum_{s=1}^T C_0(s)^2 \right] \\ &\leq \frac{2C_0}{T^2} \sum_{s=1}^T C_0(s) + \frac{3}{T} \sum_{s=1}^T C_0(s)^2. \end{aligned}$$

□

A.2 Inheriting Lyapunov conditions

Observe that each step of the main “for” loop in Algorithm 1 determines an entire transition kernel from *any* starting point; denote the kernel in step t by \mathcal{K}_t . Finally, let \mathcal{F}_t be the associated filtration.

Lemma 5 *Let Assumptions 1 and 3 hold. Then in fact Assumption 4 holds as well.*

Proof Under Assumption 3, all proposals that *decrease* V are *more* likely to be accepted under \hat{K}_t than under K , while all proposals that *increase* V are *less* likely to be accepted under \hat{K}_t than under K . Thus, for all x and all t ,

$$(\hat{K}_t V)(\hat{X}_t) \leq (KV)(\hat{X}_t) \leq (1 - \alpha)V(\hat{X}_t) + \beta, \quad (50)$$

which completes the proof. \square

A.3 Final estimates

We combine the theoretical results in the previous sections to obtain a final estimate on the error of our algorithm. Continuing the notation as above, we have our main theoretical result: Theorem 1, whose proof we give here.

Proof We set some notation. Define

$$E(T) \equiv \mathbb{E} \left[\left(\frac{1}{T} \sum_{t=1}^T f(\hat{X}_t) \right)^2 \right]. \quad (51)$$

Also define the “burn-in” time $T_b = \log(T)^2$, and define the hitting time τ_C as in Equation (43).

By Lemma 5, Assumption 4 in fact holds. Note that our assumptions also immediately give Assumption 5 with

$$\delta(t) \leq 2\gamma_0 \sqrt{\frac{\tau_0}{t}}.$$

Thus, applying Theorem 2 in line 3 and then Assumption 4 and Markov’s inequality in line 4, we have (in the notation of that theorem and assumption):

$$\begin{aligned} E(T) &\leq \mathbb{E} \left[\left(\frac{1}{T} \sum_{t=1}^{T_b} f(\hat{X}_t) \right)^2 \right] + \mathbb{E} \left[\left(\frac{1}{T} \sum_{t=T_b+1}^T f(\hat{X}_t) \right)^2 \right] + 2\mathbb{E} \left[\frac{1}{T^2} \left(\sum_{t=1}^{T_b} f(\hat{X}_t) \right) \left(\sum_{t=T_b+1}^T f(\hat{X}_t) \right) \right] \\ &\leq \mathbb{E} \left[\left(\frac{1}{T} \sum_{t=1}^{T_b} f(\hat{X}_t) \right)^2 \right] + \frac{T_b^2}{T^2} + \frac{2T_b}{T} \\ &\leq \frac{2C_0}{T^2} \sum_{s=1}^T C_0(s) + \frac{3}{T} \sum_{s=1}^T C_0(s)^2 + \mathbb{P}[\tau_C > T_b] + \frac{T_b^2}{T^2} + \frac{2T_b}{T} \\ &\leq \frac{2C_0}{T^2} \sum_{s=1}^T C_0(s) + \frac{3}{T} \sum_{s=1}^T C_0(s)^2 + \frac{\alpha}{4\beta} (1 - \alpha)^{T_b} V(x) + \frac{T_b^2}{T^2} + \frac{2T_b}{T} \\ &= O \left(\frac{1}{T^2} \sum_{s=1}^T \sqrt{\frac{\tau_0}{s}} \log \left(\sqrt{\frac{\tau_0}{s}} \right) + \frac{1}{T} \sum_{s=1}^T \frac{\tau_0}{s} \log \left(\sqrt{\frac{\tau_0}{s}} \right)^2 + \frac{\log(T)^2}{T} \right) \\ &= O \left(\frac{\log(T)^3}{T} \right). \end{aligned}$$

\square

References

- Al-Murad M, Zubari WK, Uddin S (2018) Geostatistical characterization of the transmissivity: An example of kuwait aquifers. *Water* 10(7):828
- Angelikopoulos P, Papadimitriou C, Koumoutsakos P (2015) X-TMCMC: adaptive kriging for Bayesian inverse modeling. *Computer Methods in Applied Mechanics and Engineering* 289:409–428
- Blanco JL, Rai PK (2014) nanoflann: a C++ header-only fork of FLANN, a library for nearest neighbor (NN) with kd-trees. <https://github.com/jlblancoc/nanoflann>
- Blatman G, Sudret B (2011) Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics* 230(6):2345–2367

- Bliznyuk N, Ruppert D, Shoemaker CA (2012) Local derivative-free approximation of computationally expensive posterior densities. *Journal of Computational and Graphical Statistics* 21(2):476–495
- Chkrebtii OA, Campbell DA, Calderhead B, Girolami MA, et al. (2016) Bayesian solution uncertainty quantification for differential equations. *Bayesian Analysis* 11(4):1239–1267
- Christen JA, Fox C (2005) Markov chain Monte Carlo using an approximation. *Journal of Computational and Graphical Statistics* 14(4):795–810
- Conn AR, Scheinberg K, Vicente LN (2009) *Introduction to derivative-free optimization*, vol 8. Siam
- Conrad PR, Marzouk YM, Pillai NS, Smith A (2016) Accelerating asymptotically exact MCMC for computationally intensive models via local approximations. *Journal of the American Statistical Association* 111(516):1591–1607
- Conrad PR, Davis AD, Marzouk YM, Pillai NS, Smith A (2018) Parallel local approximation MCMC for expensive models. *SIAM/ASA Journal on Uncertainty Quantification* 6(1):339–373
- Constantine PG, Kent C, Bui-Thanh T (2016) Accelerating Markov chain Monte Carlo with active subspaces. *SIAM Journal on Scientific Computing* 38(5):A2779–A2805
- Cotter SL, Dashti M, Stuart AM (2010) Approximation of Bayesian inverse problems for PDEs. *SIAM Journal on Numerical Analysis* 48(1):322–345
- 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 Resources Research* 47(10)
- Cui T, Martin J, Marzouk YM, Solonen A, Spantini A (2014) Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems* 30(11):114015
- Cui T, Marzouk Y, Willcox K (2016) Scalable posterior approximations for large-scale Bayesian inverse problems via likelihood-informed parameter and state reduction. *Journal of Computational Physics* 315:363–387
- Dodwell TJ, Ketelsen C, Scheichl R, Teckentrup AL (2015) A hierarchical multilevel Markov chain Monte Carlo algorithm with applications to uncertainty quantification in subsurface flow. *SIAM/ASA Journal on Uncertainty Quantification* 3(1):1075–1108
- Haario H, Saksman E, Tamminen J (2001) An adaptive Metropolis algorithm. *Bernoulli* 7(2):223–242
- Janetti EB, Riva M, Straface S, Guadagnini A (2010) Stochastic characterization of the montalto uffugo research site (italy) by geostatistical inversion of moment equations of groundwater flow. *Journal of hydrology* 381(1-2):42–51
- Jardani A, Dupont JP, Revil A, Massei N, Fournier M, Laignel B (2012) Geostatistical inverse modeling of the transmissivity field of a heterogeneous alluvial aquifer under tidal influence. *Journal of Hydrology* 472:287–300
- Jasra A, Kamatani K, Law KJ, Zhou Y (2018) A multi-index Markov chain Monte Carlo method. *International Journal for Uncertainty Quantification* 8(1)
- Johndrow JE, Mattingly JC, Mukherjee S, Dunson D (2015) Optimal approximating Markov chains for Bayesian inference. arXiv preprint arXiv:150803387
- Johnson SG (2014) The nlopt nonlinear-optimization package
- Kaipio J, Somersalo E (2006) *Statistical and computational inverse problems*, vol 160. Springer Science & Business Media
- Kaipio J, Somersalo E (2007) Statistical inverse problems: discretization, model reduction and inverse crimes. *Journal of Computational and Applied Mathematics* 198(2):493–504
- Kohler M (2002) Universal consistency of local polynomial kernel regression estimates. *Annals of the Institute of Statistical Mathematics* 54(4):879–899
- Latuszyński K, Rosenthal JS (2014) The containment condition and AdapFail algorithms. *Journal of Applied Probability* 51(4):1189–1195
- Li J, Marzouk YM (2014) Adaptive construction of surrogates for the Bayesian solution of inverse problems. *SIAM Journal on Scientific Computing* 36(3):A1163–A1186
- Llorente F, Martino L, Read J, Delgado D (2021) A survey of Monte Carlo methods for noisy and costly densities with application to reinforcement learning. arXiv preprint arXiv:210800490
- Marzouk Y, Xiu D (2009) A stochastic collocation approach to Bayesian inference in inverse problems. *Communications in Computational Physics* 6(4):826–847
- Matott LS (2012) Screening-level sensitivity analysis for the design of pump-and-treat systems. *Groundwater Monitoring & Remediation* 32(2):66–80
- Medina-Aguayo F, Rudolf D, Schweizer N (2018) Perturbation bounds for Monte Carlo within Metropolis via restricted approximations. arXiv preprint arXiv:180909547
- Meyn SP, Tweedie RL (2012) *Markov chains and stochastic stability*. Springer Science & Business Media
- None N (2003) 40 cfr part 191 subparts b and c and 40 cfr 194 monitoring implementation plan (rev. 3). Tech. rep., Waste Isolation Pilot Plant (WIPP), Carlsbad, NM (United States); Washington . . .
- Pillai NS, Smith A (2014) Ergodicity of approximate MCMC chains with applications to large data sets. arXiv preprint arXiv:14050182
- Pool M, Carrera J, Alcolea A, Bocanegra E (2015) A comparison of deterministic and stochastic approaches for regional scale inverse modeling on the Mar del Plata aquifer. *Journal of Hydrology* 531:214–229
- Roberts GO, Rosenthal JS (2007) Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms. *Journal of applied probability* 44(2):458–475
- Roberts GO, Tweedie RL (1996) Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika* 83(1):95–110
- Rote G, Tichy RF (1996) Quasi-Monte Carlo methods and the dispersion of point sequences. *Mathematical and computer modelling* 23(8-9):9–23
- Rudolf D, Schweizer N, et al. (2018) Perturbation theory for Markov chains via Wasserstein distance. *Bernoulli* 24(4A):2610–2639
- Schillings C, Schwab C (2014) Sparsity in Bayesian inversion of parametric operator equations. *Inverse Problems* 30(6):065007
- Sherlock C, Golightly A, Henderson DA (2017) Adaptive, delayed-acceptance mcmc for targets with expensive likelihoods. *Journal of Computational and Graphical Statistics* 26(2):434–444, DOI 10.1080/10618600.2016.1231064
- Stone CJ (1977) Consistent nonparametric regression. *The Annals of Statistics* pp 595–620

- Stuart A, Teckentrup A (2018) Posterior consistency for Gaussian process approximations of Bayesian posterior distributions. *Mathematics of Computation* 87(310):721–753
- Willmann M, Carrera J, Sánchez-Vila X, Vázquez-Suñé E (2007) On the meaning of the transmissivity values obtained from recovery tests. *Hydrogeology Journal* 15(5):833–842
- Wolff U, Collaboration A, et al. (2004) Monte carlo errors with less errors. *Computer Physics Communications* 156(2):143–153
- Zahm O, Cui T, Law K, Spantini A, Marzouk Y (2022) Certified dimension reduction in nonlinear Bayesian inverse problems. *Mathematics of Computation*

Accepted manuscript