Adaptive, delayed-acceptance MCMC for targets with expensive likelihoods

Chris Sherlock[1]*, Andrew Golightly[2] and Daniel A. Henderson[2]

[1]Department of Mathematics and Statistics, Lancaster University, UK
[2]School of Mathematics & Statistics, Newcastle University, UK

Abstract

When conducting Bayesian inference, delayed acceptance (DA) Metropolis-Hastings (MH) algorithms and DA pseudo-marginal MH algorithms can be applied when it is computationally expensive to calculate the true posterior or an unbiased estimate thereof, but a computationally cheap approximation is available. A first accept-reject stage is applied, with the cheap approximation substituted for the true posterior in the MH acceptance ratio. Only for those proposals which pass through the first stage is the computationally expensive true posterior (or unbiased estimate thereof) evaluated, with a second accept-reject stage ensuring that detailed balance is satisfied with respect to the intended true posterior. In some scenarios there is no obvious computationally cheap approximation. A weighted average of previous evaluations of the computationally expensive posterior provides a generic approximation to the posterior. If only the k-nearest neighbours have non-zero weights then evaluation of the approximate posterior can be made computationally cheap provided that the points at which the posterior has been evaluated are stored in a multi-dimensional binary tree, known as a KD-tree. The contents of the KD-tree are potentially updated after every computationally intensive evaluation. The resulting adaptive, delayed-acceptance [pseudo-marginal] Metropolis-Hastings algorithm is justified both theoretically and empirically. Guidance on tuning parameters is provided and the methodology is applied to a discretely observed Markov jump process characterising predator-prey interactions and an ODE system describing the dynamics of an autoregulatory gene network.

Keywords: Delayed-acceptance; surrogate; adaptive MCMC; pseudo-marginal MCMC; KD-tree.

1   Introduction

A major challenge for Bayesian inference in complex statistical models is that evaluation of the likelihood or, for pseudo-marginal MCMC (Andrieu and Roberts, 2009), obtaining a

---

*c.sherlock@lancaster.ac.uk

realisation from an unbiased estimator of the likelihood, can be computationally expensive. The use of a *surrogate* model with a computationally inexpensive likelihood in such circumstances has a long history; see, for example, Sacks et al. (1989), Kennedy and O'Hagan (2001), Rasmussen (2003), Bliznyuk et al. (2008), Fielding et al. (2011), Joseph (2012, 2013), Overstall and Woods (2013) and Conrad et al. (2014). The use of inexpensive surrogates has also been explored in the related context of likelihood free inference or approximate Bayesian computation (ABC) (Wilkinson, 2014; Meeds and Welling, 2014).

In this paper we propose to use a relatively generic surrogate for models with expensive likelihoods and we justify its use in adaptive, delayed-acceptance (pseudo-marginal) MCMC schemes. The delayed acceptance MCMC algorithm of Christen and Fox (2005) is a two-stage Metropolis-Hastings algorithm in which, typically, proposed parameter values are accepted or rejected at the first stage based on a computationally cheap surrogate for the likelihood. Detailed balance with respect to the true posterior is ensured by a second accept-reject step, based on the computationally expensive likelihood, for those parameter values which are accepted in the first stage. Delayed acceptance algorithms thus provide draws from the posterior distribution of interest whilst potentially limiting the number of evaluations of the expensive likelihood. Recent examples of the use of surrogates in delayed acceptance algorithms can be found in Cui et al. (2011), Higdon et al. (2011), Golightly et al. (2015) and Sherlock et al. (2015a), amongst others. Delayed acceptance algorithms which use data subsampling and partitioning for tackling large datasets have also been proposed; see Payne and Mallick (2014), Quiroz (2015) and Banterle et al. (2015).

For some models there may be an obvious cheap surrogate. For example, Golightly et al. (2015) use both the diffusion approximation and the linear noise approximation as surrogates for a Markov jump process in the context of analysing stochastic kinetic models. For many models, however, there are no obvious model-based candidates for the surrogate. It is natural in such scenarios to use regression-based methods which utilise previous evaluations of the computationally expensive likelihood to approximate the likelihood or the unnormalised posterior density at new parameter values. For example, Bliznyuk et al. (2008) use radial basis functions whereas Rasmussen (2003) and Fielding et al. (2011) use Gaussian processes (GPs). In this paper, we focus on a generic surrogate based upon likelihood values at the k-nearest neighbours (e.g. Hastie et al., 2009). It is easy to implement, computationally

cheap, and it adapts as new evaluations of the computationally-expensive likelihood become available. In Section 5.1 we consider the merits and disadvantages of an alternative, GP-based solution.

We approximate the likelihood at proposed parameter values by an inverse-distance-weighted average of the likelihoods of its k-nearest neighbours in the training data; the training data here consist of pairs of parameter vectors and their corresponding likelihoods. This k-NN approach has the advantage of being simple, local, flexible to the shape of the likelihood, and trivially adaptive as new training data become available. Our focus on adapting a local approximation is thus similar to that in Conrad et al. (2014) and we compare and contrast the two approaches in Section 5.1.

Although it is trivial to update our simple k-NN approximation as new training data become available, a naive implementation of adaptive MCMC algorithms may not converge to the intended target (e.g. Roberts and Rosenthal, 2007; Andrieu and Thoms, 2008). Careful control of the adaptation is therefore essential and we prove that, subject to conditions, the strategy we propose is theoretically valid. As with a number of previous adaptive MCMC algorithms (e.g. Roberts and Rosenthal, 2009; Sherlock et al., 2010), our MCMC kernel is a mixture of a fixed kernel and an adaptive kernel. Unlike previous such algorithms, however, the adaptive kernel can be most efficient when the fraction of applications that involve a computationally expensive evaluation is very low; this impacts on the rate of adaptation and on the rate of convergence. We therefore also provide theoretically-justified guidance on choosing both the mixture probability and the total number of iterations of the algorithm.

The main computational expense of the k-NN approximation is searching for the nearest neighbours. In a naive implementation this search takes $O(n)$ operations, where n is the number of training datapoints, so the computational expense of an adaptive, nearest-neighbour-based approach would grow linearly with the length of the MCMC run. Fortunately, there are more efficient algorithms. We use an approach based on storing the training data in a multi-dimensional binary tree, known as a KD-tree (Bentley, 1975; Friedman et al., 1977). The 'K' in 'KD-tree' indicates the dimension of the space, but to avoid confusion with the number of nearest neighbours k, we denote the dimension of the space by d. Because our tree grows on-line as new training data become available we make two major changes to the standard KD-tree algorithm, described at the start of the next section. Our adapted

KD-tree algorithm allows us to efficiently search for the nearest neighbours in approximately $O(d \log n)$ operations and add an additional value to the training set also in $O(d \log n)$ operations. This yields a highly efficient, adaptive surrogate.

The remainder of the paper is structured as follows. Our KD-tree k-NN algorithm is described in Section 2 and its use in adaptive, delayed-acceptance (pseudo-marginal) MCMC is discussed in Section 3. This section also provides the theoretically-justified guidance on the choice of kernel mixture probability and number of iterations. Section 4 applies the methodology to a discretely observed Markov jump process characterising predator-prey interactions and an ODE system describing the dynamics of an autoregulatory gene network. The paper concludes in Section 5 with a discussion.

2   The KD-tree k-nearest neighbour algorithm

Suppose the true parameter is $\theta \in \mathbb{R}^d$. We wish to find a cheap approximation $\hat{\pi}_c(\theta)$ to $\pi(\theta)$ using a weighted average of the expensive values that have already been calculated. These expensive values might be of the true posterior $\pi(\theta_1), \ldots, \pi(\theta_m)$ (involving, for example, the numerical solution of a number of differential equations as in Section 4.2) or the expensive values might be $\hat{\pi}_s(\theta_1), \ldots, \hat{\pi}_s(\theta_m)$, unbiased stochastic estimates obtained as part of a pseudo-marginal MCMC algorithm (Section 4.1).

We will average the k-nearest neighbours. Naive use of a vector of n θ values and associated log-likelihoods, v, is expensive. While adding to the list takes $O(1)$ operations, searching the list for the k nearest neighbours to a particular point, θ^*, takes $O(n)$ operations.

For our applications, typically the dimension, d, of the problem is moderate: between 3 and 12. For very low dimensional problems an analogue of the quadtree (Finkel and Bentley, 1974) would be the most efficient approach, but the number of pointers from each node grows exponentially with dimension.

We therefore use a variation on the KD-tree (Bentley, 1975; Friedman et al., 1977). Creation of the tree from n values takes $O(dn \log n)$ operations and requires storage of $O(n)$. For a balanced (see Section 2.7) tree the computation required to add an additional value is $O(d \log n)$, and to search for a nearest neighbour is also $O(d \log n)$.

The standard KD-tree has a single item of data (θ value and associated information) at

each node; all items further down the tree from this node have been split at this node as described in detail in Section 2.4. The standard structure assumes that all of the θ values to be used are available before the tree is constructed, whereas our tree potentially grows with each new position at which the log-likelihood (or an unbiased estimate thereof) is evaluated. The most efficient 'splitting point' of any node of a tree is the median of all relevant values. Use of the median leads to a balanced tree where all end nodes would be at approximately the same depth. However we do not know the true median. We therefore separate our tree in to *leaf nodes* and *branch nodes*. Branch nodes provide the splitting information and leaf nodes, which occur at the base of the tree, store multiple data values. When a leaf node becomes full, it splits according to the median of the relevant data values it contains, rather than a true median, and becomes a branch node, creating two leaf nodes beneath it. The maximum size of a leaf node is defined so that the leaf median is unlikely to be 'too far' from the true median; the choice of this tuning parameter is investigated in Section 2.7 and in the simulation study in Section 4.1.

2.1   Preliminary run

To estimate the likelihood at some new point, we will take a weighted average of likelihoods, or estimated likelihoods at the k nearest points. If the likelihood varies more quickly with distance along some axes than along others then 'nearest' should be according to some alternative metric such as a Mahalanobis distance. However we also divide the KD-tree along hyperplanes which are perpendicular to one of the Cartesian axes. This leads to gross inefficiencies when there is a strong correlation between components of θ. For example if θ has a bivariate Gaussian distribution with marginal variances of 1 and correlation 0.999 then the hyperplane splits in a tree of depth 5, say, effectively partition the first principal component but fail to partition the second, so that two θ values which are reasonably close according to the Mahalanobis distance can be in different portions of the tree. This problem is exacerbated in higher dimensions.

The algorithm should, therefore, be more efficient if the the parameters are relatively uncorrelated and if the length scales in the direction of each axis are similar. A preliminary run of the MCMC algorithm allows us to find an approximate center $\hat{\mu}$, and variance matrix $\hat{\Sigma}$. For all θ we then define $\psi := \sqrt{\hat{\Sigma}}^{-1}(\theta - \hat{\mu})$ to normalise. Since this transformation gives

a one-to-one mapping from θ to ψ, in what follows, for notational simplicity, we refer to the parameter values to be stored in the KD-tree as θ and implicitly assume that, in practice, the transformation has already been applied. The preliminary run should be of sufficient length that the sample approximately represents the gross relationships in the main posterior and hence that Euclidean distance is a reasonable metric for the transformed parameters; the preliminary chain does not need to have mixed thoroughly.

2.2 Tree preliminaries

The KD-tree stores a large number of vector parameter values, $\theta \in \Theta$, each with an associated vector of interest, $v_\theta \in \mathbb{V}$, in such a way that the time taken to either update the tree or to retrieve the information we require is logarithmic in the number of (θ, v_θ) pairs that are stored in the tree. In our case $v_\theta = (l_\theta, n_\theta) \in \mathbb{R} \times \mathbb{N}$ is the logarithm of the average of all estimates of the posterior at parameter values close to θ and the number of such estimates.

Associated with the vector of interest is a **merge** function $M : \Theta \times \mathbb{V} \times \Theta \times \mathbb{V} \rightarrow \mathbb{V}$, which combines the current vector of interest with the vector at a new value, θ^*.

2.3 Tree structure

The tree consists of **branch nodes** and **leaf nodes**. Each branch node has two children, each of which may either be a branch node or a leaf node.

Leaf nodes. Each leaf node stores up to $2b - 1$ (θ, v_θ) pairs, and the dimension, $d_{split} \in \{1, \ldots, d\}$, on which the node will split.

Branch nodes. Each branch node stores a component, $d_{split} \in \{1, \ldots, d\}$ on which it was split, and a corresponding scalar value θ_{split}, the split point. The left-hand child of the node contains (if it is a leaf node) or points to (if it is a branch) all (θ, v) pairs that have passed through this branch node and that have $\theta_{d_{split}} < \theta_{split}$. The right-hand child contains or points to all (θ, v) pairs with $\theta_{d_{split}} > \theta_{split}$. If $\theta_{d_{split}} = \theta_{split}$ it may be contained by the right hand node or the left hand node.

Root node. The node at the very top of the tree is called the root node. By default the root node has $d_{split} = 1$. If there are fewer than $2b$ leaves in the tree then the root node is a leaf node, otherwise it is a branch node.

2.4   Adding data to a tree

When a leaf has $2b$ entries it immediately spawns two leaf node children by splitting all $2b$ entries along the component, d_{split}. The splitting point, θ_{split}, is the median of the $2b$ values for θ_{dsplit}. The parent node then becomes a branch, and the leaf nodes inherit $d_{split} = d_{split}^{(p)} \oplus 1$, where $d_{split}^{(p)}$ is the component on which the parent has just been split, and $\oplus$ represents regular addition except that $d \oplus 1 = 1$.

Some initial number, n_0 of data points are used to create a balanced tree using the standard recursive procedure given in the supplementary material (Appendix A.1). After this, a single new entry (θ, l_θ) is added to the tree by descending from the root node and, at each branch, comparing component $\theta_{d_{split}}$ with θ_{split} to choose the relevant child (if $\theta_{d_{split}} = \theta_{split}$ then the left child is chosen with probability 0.5). The new entry is added to the leaf node that is reached. If this leaf node still has fewer than $2b$ entries then the algorithm stops, otherwise the leaf-node becomes a branch and spawns two child leaf-nodes as described above.

2.5   Searching the tree

We apply a standard two-stage recursive algorithm to find the r nearest neighbours of a given point, θ^*, together with their distances from θ^*. We provide an overview of the algorithm below; the procedure is detailed in full in the supplementary material (Appendix A.2).

The algorithm first descends the tree to find the leaf node to which θ^* would belong, as described in Section 2.4. It then gradually ascends the tree from this node to the root node; after each ascent from a node to its parent a test is conducted to see whether the other child, that is the child of the current node from which the algorithm has *not* just ascended, or any of its offspring might hold a closer neighbour than the current k nearest. If this is the case then, before any further ascent can take place, the search descends down the tree via this other child.

2.6   Restricting the growth of the tree

The tree will be used to obtain computationally cheap estimates of the posterior density. We would like to ensure that the accuracy of the approximations increases with the amount

of information stored in the tree, whether the information arises from exact evaluations of the posterior or from stochastic estimates. We would also prefer the cost of obtaining this information to increase slowly, if at all, with the amount of information.

To reduce the total number of leaves we set a minimum distance between leaves, ϵ. For any new information, (θ^*, v_{θ^*}), to be added to the tree we first ascertain whether or not any pairs (θ, v_θ) exist with $||\theta^* - \theta|| < \epsilon$. If one or more such pair exists then the merge function, M, described in Section 2.2 is used to combine (θ^*, v_{θ^*}) with the nearest pair, providing a replacement value for v_θ; otherwise (θ^*, v_{θ^*}) is added to the tree as described in Section 2.4.

When each v_{θ^*} contains an exact evaluation of the posterior then the merge simply ignores the new information. However, when v_{θ^*} contains a stochastic estimate of the posterior some weighted average of the new estimate and of the current average will be more appropriate since, by the continuity of the posterior, for sufficiently small ϵ, $\pi(\theta^*) \approx \pi(\theta)$ for all θ^* such that $||\theta^* - \theta|| < \epsilon$. In particular therefore, for pseudo-marginal algorithms, we define

$$M_{PM}(\theta, [l_\theta, n_\theta], \theta^*, [l_{\theta^*}, 1]) := \left[\log \left[n_\theta e^{l_\theta} + e^{l_{\theta^*}} \right] - \log(n_\theta + 1), \ n_\theta + 1 \right].$$

This vector replaces the previous $[l_\theta, n_\theta]$ vector and so is associated with the position θ.

Choice of merge distance

Consider a tree with n existing points, $\theta_1, \ldots \theta_n$ and to which it is proposed that a new point θ^*, chosen at random, will be added. We relate the merge distance, ϵ, to the probability, p_{keep}, that none of the existing points is within the ϵ ball of θ^*, so that θ^* will be added to the tree. This then provides a guide to setting ϵ itself.

Let B_ϵ^* be the ϵ ball around θ^*, define $N_{n,\epsilon}$ to be the number of the n existing points that are inside B_ϵ^* and consider $E_{n,\epsilon} = \mathbb{E}[N_{n,\epsilon}]$. The following is proved in the supplementary material (Appendix B).

Proposition 1. *If* $0 < E_{n,\epsilon} < 1$ *then* $1 - E_{n,\epsilon} < p_{keep} < e^{-E_{n,\epsilon}}$.

To use Proposition 1 we require an expression for $E_{n,\epsilon}$, and this depends on the distribution of $(\theta_1, \ldots, \theta_n, \theta^*)$. For tractability, and because it will often hold approximately with reasonably sized data sets, we suppose that the target is Gaussian, so that $\theta_i \sim N(\mu, \Sigma)$, marginally. This is then normalised (see Section 2.1) so that the following result (see again Appendix B for a proof) can be applied.

Proposition 2. *Let jointly distributed $\theta_1 \sim N(0, I_d), \ldots, \theta_n \sim N(0, I_d)$, be independent of $\theta^* \sim N(0, I_d)$. Then $E_{n,\epsilon} = nF_{\chi_d^2}(\epsilon^2/2)$, where $F_{\chi_d^2}$ is the cumulative distribution function of a χ_d^2 random variable.*

In the simulation studies of Section 4 we choose ϵ such that $E_{n,\epsilon} = 0.5$, giving $0.5 < p_{keep} < 0.61$ and $\epsilon \approx \sqrt{2q_{\chi_d^2}(1/2n)}$, where $q_{\chi_d^2}$ is the quantile function of a χ_d^2 random variable.

2.7 Ensuring that the tree remains balanced

We consider two mechanisms through which a KD-tree that is constructed on-line using an MCMC algorithm may become unbalanced, and for each problem we provide a solution.

When $d = 1$, there are two binary tree structures that allow for online rebalancing: the red-black tree and the AVL tree (e.g. Storer, 2002). Unfortunately there are no known algorithms for rebalancing a KD-tree online and so we consider an alternative which ensures that the KD-tree remains approximately balanced as at grows.

The splitting hyper-planes define a partition of Θ. Let the *node box* corresponding to a particular node be the (possibly unbounded) subset of Θ defined through the constraint at each splitting hyper-plane on the journey from the root node to the node in question, as described in Section 2.4.

Consider, informally, for any node box, an 'effective width' along a particular co-ordinate axis to be some representative width such as the standard deviation of the posterior restricted to the node box. Suppose that the MCMC chain is currently in a node box where an effective width along its splitting co-ordinate is δ. Now, imagine that at each iteration the chain barely moves compared to δ, and each jump proposal - which can give a new evaluation of the log-likelihood even if the chain does not move - is also small compared to δ. In this case all of the new samples for a large number of iterations will descend to this one particular leaf node, which will fill up and then split. This split will not, however, be representative of the marginal median (in terms of the posterior) for the node box along the splitting axis. Hence, over the rest of the MCMC run, once the chain has moved on, there will be one side of the split which takes most of the future sample points and one side which takes very few of them, leading to an unbalanced tree.

2b	[0.4,0.6]	[0.3,0.7]	[0.2,0.8]	$d = 3$	$d = 10$
10	0.49	0.15	0.02	18.3, 14-23 (11-25)	18.4, 14-23 (11-28)
20	0.35	0.05	0.002	17.7, 15-21 (13-23)	17.7, 15-21 (12-23)
30	0.26	0.02	0.0002	17.5, 15-20 (14-22)	17.5, 15-20 (14-21)
40	0.19	0.007	0.00001	17.4, 15-19 (14-21)	17.4, 15-19 (13-20)

Table 1: Left: the probability that an estimated median from a sample of size $2b \in \{10, 20, 30, 40\}$ will be at a quantile outside of the range $[0.4, 0.6]$, $[0.3, 0.7]$ or $[0.2, 0.8]$. Right: tree depths when $2b \times 10^5$ independent entries are added sequentially to a tree: mean over all leaf nodes, range of the central 99% of leaf nodes and the maximum and minimum.

Now suppose that a preliminary run of n_0 iterations has been carried out and that over this run the chain has been seen to mix reasonably across the posterior. Let us now construct a *balanced* tree from this run. Each leaf node will have b^* (or $b^* + 1$) entries, for some $b^* \in \{b, b+1, \ldots, 2b-2\}$. Consider the node box of any specific leaf node in this tree. Since the chain has mixed reasonably, it should represent the posterior within this node box, and in particular (1) in the component over which the leaf node will split, the median should be reasonably approximated, and (2) in the main MCMC run, the chain should also cover this box in approximately b^* iterations or fewer; hence it will also cover any sub-divisions of the box in approximately b^* (or fewer) iterations.

The second reason a tree can become unbalanced is Monte Carlo error. A leaf splits after it has $2b$ entries by finding the median of $\theta_{d_{split}}$ over the $2b$ entries, but this sample median will not be the true median over the node-box. Table 1 provides, for 4 different values of $2b$, the probability that the estimated median will be at a true quantile which is outside the range shown. A tree with $2b \times 10^5$ iid entries was simulated for each value of $2b$ and each of two dimensions. The table also shows the mean depth of the leaf nodes and the range of depths of 99% of the leaf nodes and, in brackets, of all leaf nodes. Both aspects of Table 1 suggest that $2b = 20$ or $2b = 30$ should lead to a reasonably balanced tree with few leaf nodes requiring much more effort to reach than the majority of the leaf nodes and with little effect on the overall mean amount of effort required to reach a leaf node.

3 Adaptive MCMC algorithm

We briefly review the delayed-acceptance algorithm before describing our adaptive version. This is further extended to an adaptive pseudo-marginal version in Section 3.4.

3.1 Delayed-acceptance algorithms

Given a current parameter value, $\theta \in \Theta$, the Metropolis-Hastings (MH) algorithm proposes a new value, θ^* from some density $q(\theta^*|\theta)$ and then accepts or rejects according to

$$\alpha_{MH}(\theta, \theta^*) := 1 \wedge \frac{\pi(\theta^*)q(\theta|\theta^*)}{\pi(\theta)q(\theta^*|\theta)}. \tag{1}$$

The delayed acceptance Metropolis-Hastings (daMH) algorithm utilises a cheap (deterministic or stochastic) approximation $\hat{\pi}_c$ in two stages. At Stage One, $\hat{\pi}_c$ is substituted for π in the standard MH acceptance formula:

$$\tilde{\alpha}_1(\theta, \theta^*) := 1 \wedge \frac{\hat{\pi}_c(\theta^*) \; q(\theta|\theta^*)}{\hat{\pi}_c(\theta) \; q(\theta^*|\theta)}, \tag{2}$$

A second accept/reject stage is applied to any proposals that pass Stage One and a proposal is only accepted if it passes both stages. The Stage Two acceptance probability is:

$$\tilde{\alpha}_2(\theta, \theta^*) := 1 \wedge \frac{\pi(\theta^*)\hat{\pi}_c(\theta)}{\pi(\theta)\hat{\pi}_c(\theta^*)}. \tag{3}$$

The overall acceptance probability, $\tilde{\alpha}_1(\theta, \theta^*)\tilde{\alpha}_2(\theta, \theta^*)$ ensures that detailed balance is satisfied with respect to π; however if a rejection occurs at Stage One then the expensive evaluation of $\pi(\theta)$ at Stage Two is unnecessary.

3.2 Adaptive, delayed-acceptance algorithm

As in Roberts and Rosenthal (2007, 2009) our adaptive kernel consists of a mixture of a fixed kernel and an evolving kernel. At each iteration, with a user-defined probability $\beta \in (0,1)$, the fixed kernel is selected, otherwise the adaptive kernel is used. Our fixed kernel is a standard Metropolis-Hasting kernel and our evolving kernel uses delayed-acceptance with an acceptance rate derived from a cheap approximation $\hat{\pi}_c(\theta^*)$ which is an inverse-distance-weighted average of the expensive evaluations of the true posterior at the k nearest neighbours to θ^* in the KD-tree at iteration n. After iteration $n-1$ let there have been i_{n-1} evaluations

of the true posterior, π. Let these evaluations be at values $\theta^*_{i_1}, \ldots, \theta^*_{i_{n-1}}$. Our adaptive algorithm requires a sequence of probabilities, $\{p_i\}_{i \in \mathbb{N}}$, with

$$\lim_{i \to \infty} p_i = 0. \tag{4}$$

In practice, the algorithm proceeds until some n_{tot} iterations have been performed.

Algorithm 1: *adaptive-KD-tree, delayed-acceptance Metropolis-Hastings.*

1. With probability β go to Step 2 (MH) else go to Step 3 (da-MH).

2. **MH**: Propose θ^* from $q(\theta^*|\theta)$. Evaluate the expensive posterior, $\pi(\theta^*)$, and accept the proposal $(\theta \leftarrow \theta^*)$ with probability given by (1); otherwise reject the proposal $(\theta \leftarrow \theta)$. Go to Step 4.

3. **da-MH**: Propose θ^* from $q'(\theta^*|\theta)$.

    (a) **Stage 1**: Evaluate $\hat{\pi}_a(\theta^*)$ using the current KD-tree; with probability $\tilde{\alpha}_1(\theta, \theta^*)$ as defined in (2) proceed to Step 3b (Stage 2); otherwise reject the proposal $(\theta \leftarrow \theta)$, set $i_n = i_{n-1}$, go to next iteration.

    (b) **Stage 2**: Evaluate $\pi(\theta^*)$; accept the proposal $(\theta \leftarrow \theta^*)$ with probability $\tilde{\alpha}_{2,s}(\theta, \theta^*)$ as defined in (3); otherwise reject the proposal $(\theta \leftarrow \theta)$. Go to Step 4.

4. Set $i_n = i_{n-1} + 1$; add $(\theta^*, \pi(\theta^*))$ to a list of recently-evaluated parameter/posterior pairs; with probability p_{i_n} transfer all pairs from this list to the KD-tree; go to next iteration.

3.3 Delayed acceptance random walk Metropolis

It remains to choose the proposal mechanisms, q and q'. The Random Walk Metropolis (RWM) is a MH algorithm where $q(\theta^*|\theta) = q(||\theta^* - \theta||)$ for some suitable norm, and hence $q(\theta^*|\theta)$ and $q(\theta|\theta^*)$ cancel in the acceptance ratios (1) and (2). We consider the standard choice of

$$q(\theta^*|\theta) = N(\theta^*; \theta, V) \quad \text{and} \quad q'(\theta^*|\theta) = q^\xi(\theta^*|\theta) := N(\theta^*; \theta, \xi^2 V), \tag{5}$$

where $N(\cdot; \theta, V)$ denotes a multivariate Gaussian density with mean θ and variance V and where V has been chosen so as to approximately optimise the efficiency of the standard RWM algorithm.

As we shall discover, the cheap approximation, $\hat{\pi}_c$ is reasonably accurate. As ξ increases from 1 the overall acceptance rate and, in particular, the Stage One acceptance rate, α_1, can decrease quite substantially. If all computational expense is negligible except for the evaluation of the true posterior then for a given amount of computational effort, the number of evaluations of the expensive posterior remains approximately constant, although the total number of iterations of the algorithm increases in proportion to the reciprocal of α_1. However, as each proposed jump is larger, moves which are accepted at Stage Two are typically larger. Thus, provided the Stage Two acceptance rate does not decrease too drastically, the mixing of the algorithm (in terms of movement per CPU second) can, and often does, actually increase for intermediate values of ξ. This heuristic has been noted before (e.g. Christen and Fox, 2005; Banterle et al., 2015) and also applies for the delayed-acceptance pseudo-marginal RWM. A rigorous analysis of the behaviour of these algorithms as a function of the scaling is provided in Sherlock et al. (2015a).

3.4 Adaptive, delayed-acceptance, pseudo-marginal algorithm

We first overview pseudo-marginal Metropolis-Hastings algorithms and then describe the adjustments to the set-up required for the pseudo-marginal version of our algorithm. The algorithm itself is provided in the supplementary material (Appendix C).

The pseudo-marginal algorithm uses a non-negative stochastic estimator $\hat{\pi}_s(\theta; Z)$ of the posterior $\pi(\theta)$, where Z is a collection of random variables whose distribution may, and usually does, depend on θ. Crucially, we require $\mathbb{E}\left[\hat{\pi}_s(\theta; Z)\right] = c\pi(\theta)$, where c is fixed and non-negative. We may therefore rewrite $\hat{\pi}_s(\theta)$ as $\pi(\theta)W$ with $W \in \mathcal{W} \subseteq [0, \infty)$ sampled from some density $q_\theta(w)$, and

$$\mathbb{E}\left[W\right] = \int_0^\infty w q_\theta(w) \, \mathrm{d}w = c. \tag{6}$$

The pseudo-marginal MH (PsMMH) algorithm is simply a Metropolis-Hastings Markov chain acting on the extended statespace $\Theta \times \mathcal{W}$ with a target of

$$\tilde{\pi}(\theta, w) = \frac{1}{c}\pi(\theta)q_\theta(w)w. \tag{7}$$

This has the required marginal for θ by (6). Detailed balance is ensured with respect to this

target by setting the probability for (θ^*, w^*) being accepted to:

$$\alpha_{PM}([\theta, W], [\theta^*, W^*]) := 1 \wedge \frac{\hat{\pi}_s(\theta^*)q(\theta|\theta^*)}{\hat{\pi}_s(\theta)q(\theta^*|\theta)} = 1 \wedge \frac{\pi(\theta^*)q(\theta|\theta^*)W^*}{\pi(\theta)q(\theta^*|\theta)W}. \tag{8}$$

When delayed acceptance is implemented with a pseudo-marginal framework, the Stage One acceptance probability is exactly as in (2). In Stage Two the true posterior in (3) is replaced with the realisation from the unbiased estimator:

$$\tilde{\alpha}_{2,PM}(\theta, \theta^*) := 1 \wedge \frac{\hat{\pi}_s(\theta^*)\hat{\pi}_c(\theta)}{\hat{\pi}_s(\theta)\hat{\pi}_c(\theta^*)}. \tag{9}$$

The kernel $\tilde{P}$ is now a fixed PsMMH kernel on $\Theta \times \mathcal{W}$ and $\{\tilde{P}_\gamma\}_{\gamma \in \mathcal{G}}$ is now a set of pseudo-marginal kernels on $\Theta \times \mathcal{W}$. The common stationary density of all kernels is now given in (7), so, very importantly, all kernels use the same mechanism for generating the estimate $\hat{\pi}_s(\theta^*)$ of the posterior at the proposed value for θ.

The algorithm proceeds as for the non-pseudo-marginal version except that $\hat{\pi}_s$ is substituted for π in (1) (fixed kernel), and (3) is replaced with (9) (evolving, DA kernel). Naturally, instead of storing evaluations of π the KD-tree now stores realisations of the unbiased approximation, $\hat{\pi}_s$.

3.5   Theory and guidance

We show that, subject to conditions, our algorithms (with general proposals, q) are ergodic. We also provide guidance on choosing the probability of using the fixed kernel, β, and on the number of iterations for which the algorithm should be run.

Define $\alpha_{MH}(\theta, \theta^*)$ as follows. For the adaptive KD-tree daMH algorithm of Section 3.2 $\alpha_{MH}(\theta, \theta^*)$ is the acceptance probability for the fixed kernel as given in (1). For the pseudo-marginal version of the algorithm in Section 3.4 it is the acceptance probability for an hypothetical, idealised version of the fixed, pseudo-marginal kernel where the posterior is known exactly, up to a fixed multiplicative constant: $\alpha_{MH}(\theta, \theta^*) = \alpha_{PM}([\theta, 1], [\theta^*, 1])$, where α_{PM} is defined in (8). We require a minorisation condition and, for the daPsMMH algorithm, and additional assumption of uniformly bounded weights. These assumptions are discussed in in the supplementary material (Appendix D), where their main consequence, Theorem 1 is proved.

Assumption 1. There is a density $\nu(\theta)$ and $\delta > 0$ such that $q(\theta^*|\theta)\alpha_{MH}(\theta, \theta^*) \geq \delta\nu(\theta^*)$ for all $\theta \in \Theta$.

Assumption 2. The support for W is uniformly (in θ) bounded above by some $\overline{w} < \infty$.

Theorem 1. *Subject to Assumption 1 the adaptive KD-tree daMH algorithm of Section 3.2 is ergodic. The adaptive KD-tree daPsMMH algorithm of Section 3.4 is ergodic subject to Assumptions 1 and 2.*

Now consider the specific, scaled, proposal q^ξ defined in (5), where increasing ξ decreases the Stage One acceptance rate, α_1. The algorithm may only accept a proposal after a computationally-expensive evaluation (of π for daMH, or $\hat{\pi}_s$ for daPsMMH). Thus, if the probability, β, that the non-DA kernel will be chosen is unaltered, then as $\alpha_1 \to 0$ nearly all of the expensive evaluations will be by the fixed, non-DA kernel, and the relative contribution from the DA kernels will unintentionally dwindle to zero.

Decreasing β in proportion to α_1 would fix the fraction of all expensive evaluations that are by the DA kernel; however with a smaller β the chain can no longer be guaranteed to be as close to π after the same, fixed number of iterations. Theorem 2, which is stated and proved in the supplementary material (Appendix D), shows that, with $\beta \propto \alpha_1$, as α_1 decreases the total number of iterations, n, of the algorithm should be increased so as to maintain the expected total number of expensive evaluations of the posterior, $\mathbb{E}\left[I_n^\xi\right]$, and that $\mathbb{E}\left[I_n^\xi\right]$ can be set so as to maintain any given upper bound on the total variation distance between the chain and π whatever the value of α_1.

Fixing $\mathbb{E}\left[I_n^\xi\right]$ approximately fixes the expected overall CPU cost. Furthermore, adaptation can only occur when the expensive posterior is evaluated, and it occurs with a fixed set of probabilities that depend on i and not on the iteration number. So, fixing $\mathbb{E}\left[I_n^\xi\right]$ also approximately fixes the expected number of adaptation occurrences.

4 Simulation Studies

In this section we evaluate the empirical performance of the proposed da-PsMMH and da-MH algorithms by considering two examples based upon Markov jump processes (MJPs).

The first example (Section 4.1) arises from the Lotka-Volterra system of predator-prey interactions (e.g. Boys et al., 2008). Since the marginal likelihood is intractable, we apply the adaptive da-PsMMH scheme and compare its performance over a range of tuning parameter choices with that of an optimised PsMMH scheme. The second example (Section 4.2) arises from the autoregulatory network proposed by Golightly and Wilkinson (2005). Given the size and complexity of this system, a linear noise approximation (LNA) (van Kampen, 2001) (see also Appendix F of the supplementary material), of the corresponding Markov jump process is taken to be the inferential model of interest. Following the algorithm of Fearnhead et al. (2014) (see Appendix F.1), the marginal likelihood under this model is tractable, but involves the solution of a system of 14 coupled ordinary differential equations (ODEs), which can be time consuming. We therefore apply the da-MH algorithm and compare its performance to a simple MH scheme without delayed acceptance.

Both MJPs are described through a set of r reactions between p different species, $\mathcal{X}_1, \ldots, \mathcal{X}_p$. The hazard rate of each reaction depends on the current species numbers, $X_1, \ldots, X_p$ via an assumption of mass-action kinetics with unknown reaction rate constants $\nu_1, \ldots, \nu_r$; for further details regarding the construction of MJP representations of reaction networks we refer the reader to Wilkinson (2012). Tables E.1 and E.2 in Appendix E.1 list the reactions and associated hazards for each example.

For the Lotka-Volterra model, the Gillespie algorithm (Gillespie (1977)) was applied, using parameter values taken from Wilkinson (2012), to generate a skeleton path comprising 51 values of X_t at integer times in the interval $[0, 50]$. For the autoregulatory system, the LNA itself was used, with parameter values taken from Golightly and Wilkinson (2011), to generate two skeleton paths containing, respectively, 101 and 201 values of X_t at evenly-spaced times covering the intervals $[0, 100]$ and $[0, 1000]$. All skeletons were then corrupted with Gaussian noise to form the data sets on which inference was performed:

$$Y_t | X_t = x_t \ \sim N(x_t, D), \tag{10}$$

where D is a diagonal matrix with diagonal entries $\sigma_1^2, \ldots, \sigma_p^2$. In the autoregulatory example we refer to the data sets with 101 and 201 observations as $\mathcal{D}_1$ and $\mathcal{D}_2$ respectively. Appendix E.1 provides details of the initial conditions and parameter values for each simulation as well as for the variances of the corrupting Gaussian noises.

For inference, in both cases, for simplicity, the initial state of the system is fixed at its (known) true value. As all of the parameters are strictly positive we therefore consider a logarithmic transformation so that the parameter vector of interest is $\theta = (\log(\nu_1), \ldots, \log(\nu_r), \log(\sigma_1), \ldots, \log(\sigma_p))$. In both examples, individual components of θ for which inference is performed are given independent Uniform $U(-8, 8)$ priors. For the Lotka-Volterra model, $r = 2$ and $p = 2$ so that $\dim(\theta) = 4$. For the autoregulatory system, $p = 4$ and, as discussed in Appendix E.1, we perform inference on $r = 6$ rate constants, giving $\dim(\theta) = 10$.

Random-walk proposals of the form (5) are used for the fixed and the adaptive kernels in both examples. The fixed kernels used a proposal variance of $V_{fixed} = \lambda\hat{\Sigma}$, where $\hat{\Sigma}$ is the sample variance from an initial, pilot run using a fixed, non-delayed-acceptance kernel, and λ is chosen to optimise the efficiency of the fixed kernel. The corresponding adaptive kernels use a proposal variance of $\xi^2 V_{fixed}$, with $\xi > 0$ a tuning parameter.

In each example we take the probability of adding to the tree at iteration n to be

$$p_{i_n} = (1 + ci_n)^{-1}. \tag{11}$$

For each example, the computational cost of an evaluation of the expensive stochastic approximation was estimated from the initial training run, and this provided an estimate of the number of expensive evaluations, $\hat{i}$, that would fit within the computational budget. Using the guidelines in Section 2.6, the parameter ϵ was chosen with a desire that points still be as likely as not to be added to the tree after $\hat{i}/2$ evaluations had already been added. For both systems this had the effect of limiting the overall tree size to around four times the size of the initial training set.

4.1 Discretely observed Markov jump process

To implement the adaptive da-PsMMH scheme described in Algorithm 1, we used a bootstrap particle filter with m particles (Andrieu et al., 2010) to obtain each value of $\hat{\pi}_s(\cdot)$. Both m and V_{fixed} were chosen so as to optimise the efficiency of the fixed kernel; see the supplementary material (Appendix E.2). We initialised the KD-tree using the first 10^4 evaluations of the expensive posterior in the pilot run. For all experiments, we assumed a fixed computational budget of 10^4 seconds, which equated to approximately $\hat{i} = 40000$ evaluations of the expensive posterior and, thus, a maximum tree size of five times the initial size.

To assess the effect of scaling ξ on the overall and relative (to PsMMH) efficiency of da-PSMMH we fixed the number of leaf nodes in the KD-tree upon splitting to be $2b = 20$, the number of nearest neighbours to be $k = b = 10$ and took the parameter controlling the rate of adaptation to be $c = 0.001$ to give around a 5% chance of adaptation after half the computational budget. We then ran the algorithm for values of $\xi \in [1, 4]$, following the practical advice of Section 3.5 by choosing $\beta \propto \hat{\alpha}_1$ with $\beta = 0.05$ for $\xi = 1$. Firstly, the sampled posterior values are consistent with the ground truth parameter values that produced the data (see Figure G.1 in the supplementary material for the marginal posterior distributions for a typical run). Figure 1 shows the effect of scaling on minimum effective sample size (mESS) over each parameter chain relative to that obtained under an optimally tuned PsMMH scheme (with an acceptance rate of 9.9% and an mESS of 528) and the effect of scaling on the Stage 1 acceptance probability. The (scaled) values of β used for each run are also shown. Figure 1 suggests that for the values of ξ considered, $\xi = 3$ is optimal in terms of mESS and gives an improvement on overall efficiency over PsMMH of a factor of 6.8. Even simply taking the same scaling as PsMMH still gives a 3-fold increase in efficiency of da-PsMMH over PsMMH.

To assess the effect of adaptation on the performance of the algorithm we fixed $\xi = 3$, $k = b = 10$ and performed runs with $c \in := \{0.0001, 0.001, 0.01, \infty\}$ with $c = \infty$ representing no adaptation. Table 2 summarises our findings. At $\hat{i}/2$ iterations, these values of c correspond to an expected number of expensive evaluations before adaptation occurs of approximately $\{3, 20, 200, \infty\}$, respectively. The larger the pause between adaptations, the less accurate the tree is between adaptations, and while 200 new evaluations is small compared with 10000 or more existing evaluations, it must be remembered that the most recent evaluations will be from a similar part of the state space to the current position and so will be among the most relevant. The reduction in accuracy especially in new, low-density regions, increases the Stage 1 acceptance rate and decreases the Stage 2 acceptance rate giving an overall reduction in statistical efficiency. Moreover, the increase in the Stage 1 rate results in a larger number of expensive posterior evaluations.

With $c = \infty$ the algorithm runs with no adaptation and just uses the initial training set of $10,000$ posterior evaluations. In this case (last row of Table 2) performance is better than for the simple RWM algorithm but worse than for all of the cases that allow adaptation,

providing clear evidence of the importance of adaptation for our algorithm. Further, the more slowly $p_i \downarrow \infty$, the more efficient the algorithm.

We also explore the sensitivity of our method to the choice of the number of leaf nodes in the KD-tree upon splitting ($2b$) and the number of nearest neighbours (k). We fixed $\xi = 3$, $c = 0.001$ and took $2b \in \{4, 10, 20, 30\}$ and $k \in \{2, 5, 10, 15\}$. Table G.3 in Appendix G shows empirical performance for each (k, b) combination considered. Consistent with our findings in Section 2.7, increasing b increases the efficiency until $2b = 20$, but there is little difference when moving from $2b = 20$ to $2b = 30$. Fixing $2b$ and varying k suggests that $k = 5$ is optimal in terms of mESS. Further discussion can be found in the supplementary material (Appendix G). Finally we examine the gain in overall efficiency by using a KD-tree as a storage and look-up method over simply storing posterior evaluations in a list. Running the da-PsMMH scheme with the optimal values ($\xi = 3$, $c = 0.001$ and $k = 5$) gave an mESS of 2009. Thus, in this example, using a KD-tree increases overall efficiency over a naive approach by a factor of 1.9. Naturally, increasing the computational budget (and therefore the number of posterior evaluations to be stored) will increase the advantage of the KD-tree.

c	Tree Size	Mean depth	depth range	$\hat{\alpha}_1$	$\hat{\alpha}_2$	mESS	Rel. mESS
0.0001	41078	11.82	10-14	0.00772	0.339	3845	7.28
0.001	40256	11.79	10-14	0.00915	0.276	3591	6.80
0.01	43248	12.04	10-15	0.0121	0.204	2464	4.67
∞	10000	9.69	9-10	0.0175	0.136	1829	3.46

Table 2: Effect of rate of adaptation c. Final tree size, mean leaf node depth, depth range, empirical stage 1 and 2 acceptance rate, minimum effective sample size (mESS) and relative mESS.

4.2 Discretely observed ODE system

The LNA gives a Gaussian model for X_t which when coupled with the Gaussian observation model above, permits a tractable form for the marginal likelihood which we denote by $\pi(y_{1:n}|\theta)$. An algorithm for evaluating the marginal likelihood, and therefore the posterior (up to proportionality) under the LNA, can be found in the supplementary material

Figure 1: Left panel. Minimum effective sample size (mESS) relative to optimised PsMMH, against scaling. Right panel. Empirical stage 1 acceptance probability $\hat{\alpha}_1$ against scaling. The points represent $\hat{\alpha}_1(\xi = 1)\beta(\xi)/\beta(\xi = 1)$ and show that $\beta(\xi)$ was scaled in proportion to $\alpha_1(\xi)$.

(Appendix F.1). Executing one iteration of the algorithm requires calculation of a full numerical solution of the ODE system (F.1) over $[0, 100]$ or $[0, 1000]$. Our implementation uses standard routines from the GNU scientific library, specifically the explicit embedded Runge-Kutta-Fehlberg $(4, 5)$ method. We limit the computational cost of these calculations by applying the da-MH scheme.

The pilot run for each dataset was of 3×10^4 iterations. We initialised the KD-tree with all 3×10^4 evaluations of the expensive posterior obtained from the initial pilot run. For all experiments, we assumed a fixed computational budget of 5×10^3 seconds, which equated to approximately $\hat{i} = 120000$ evaluations of the expensive posterior and, thus, a maximum tree size of five times the initial size. Following the findings of Section 4.1 we initially set the adaptation rate to $c = 0.001$, and set $2b = 20$ and $k = 5$.

Firstly, Figure G.3 in Appendix G shows that the sampled parameter values are consis-

Algorithm	ξ	c	$\hat{\alpha}_1$	$\hat{\alpha}_2$	mESS	Rel. mESS
			$\mathcal{D}_1$ (101 obs. on $[0, 100]$)			
MMH	1.0	–	0.225	1.000	2383	1.00
da-MMH	1.5	0.0001	0.141	0.482	7699	3.23
	1.5	0.001	0.131	0.480	7638	3.21
	1.5	∞	0.171	0.366	4530	1.90
			$\mathcal{D}_2$ (201 obs. on $[0, 1000]$)			
MMH	1.0	–	0.2291	1.000	632	1.00
da-MMH	2.0	0.0001	0.0455	0.394	2996	4.74
	2.0	0.001	0.0390	0.338	2779	4.40
	2.0	∞	0.0978	0.165	1485	2.35

Table 3: Algorithm, optimal scaling (ξ), adaptation rate parameter (c), empirical stage 1 ($\hat{\alpha}_1$) and 2 ($\hat{\alpha}_2$) acceptance rate, minimum effective sample size (mESS) and relative mESS.

tent with the ground truth, with a decrease in uncertainty when using more observations. With dataset $\mathcal{D}_2$ the LNA equations must be solved over a longer time period than for $\mathcal{D}_1$, and the steps of the algorithm for calculating the marginal likelihood (in Appendix F.1) must be executed twice as many times. Consequently Figure 2 (left panel) shows that the minimum effective sample size (mESS) obtained under da-MH is smaller when using dataset $\mathcal{D}_2$. However, mESS relative to the same quantity under MH is increased when using $\mathcal{D}_2$, since the cost of evaluating the KD-tree is unchanged (for a fixed tree size). The optimal scaling ξ (for the values considered) for each scheme is reported in Table 3. We also report output of additional runs with $c \in := \{0.0001, 0.001, \infty\}$. An optimally tuned da-MH scheme (with $c = 0.001$) gives an increase in overall efficiency of a factor of 3.2 when using $\mathcal{D}_1$ and 4.4 when using $\mathcal{D}_2$. When $c = \infty$ (representing no adaptation), we see an increase in Stage 1 acceptance rate and a decrease at Stage 2. The resulting decrease in empirical performance provides further evidence of the importance of adaptation. Finally, we again note that the algorithm performs best with a very low value of c whilst still providing posterior output consistent with $c = 0.001$ (results not shown).

Figure 2: Left panel. Minimum effective sample size (mESS) from da-MH against scaling. Right panel. Minimum effective sample size from da-MH relative to the same quantity from optimised MH, against scaling. For each panel, the solid and dashed lines indicate output using datasets $\mathcal{D}_1$ and $\mathcal{D}_2$ respectively.

5 Discussion

We have presented standard and pseudo-marginal versions of an adaptive, delayed-acceptance random walk Metropolis algorithm. The delayed-acceptance (DA) step is generic, estimating the posterior using an inverse-distance-weighted average of the k-nearest previous evaluations of the posterior, and the search for these neighbours is made fast by storing a subset of previous evaluations in a customised version of a KD-tree. The kernel is a mixture of a fixed (non-adaptive, non-DA) kernel and an adaptive DA kernel. Easy-to-use C code for creating a KD-tree and storing and retrieving values from it is provided alongside this article.

We have shown that our algorithm is ergodic, subject to conditions. Furthermore, as the scaling of the RWM proposal in the DA kernel is increased, the probability of choosing the fixed kernel should be decreased in proportion to the Stage One acceptance rate of the DA kernel and the total number of iterations should be scaled so that the expected number of

evaluations of the expensive posterior remains constant or, equivalently, so that the total computational budget remains fixed.

Pseudo-marginal and non-pseudo-marginal versions of the methodology were applied, respectively, to synthetic data generated from a discretely observed Lotka-Volterra system and an ODE model of species dynamics in an autoregulatory gene network. In these two examples, our proposed scheme outperforms the standard scheme by factors of approximately 7 and 4 for, respectively.

In both of our examples the algorithm was more efficient the more slowly the adaptation probability approached zero, suggesting that it might be best to simply add each new evaluation of the expensive posterior directly to the KD-tree directly rather than storing them in a queue. Forcing the limit of the adaptation probability to be 0 ensures the diminishing adaptation condition (see, e.g., Theorem 5) is satisfied; diminishing adaptation is itself one of the key conditions required for ergodicity and the concern would be that by removing this direct constraint the algorithm would no longer be ergodic. Whilst this seems likely to be the case in general, Theorem 3 in the supplementary material (Appendix D.5.1) shows that for a variation of Algorithm 1 on a compact state space it is possible to adapt after (at worst) every other expensive iteration and still remain ergodic.

We have focussed on making the cheap approximation to the posterior adaptive by intermittently updating the KD-tree. The covariance matrix of our random walk proposal could have been made adaptive by intermittently updating the covariance matrix of the random walk proposal so that it uses all entries in the chain to date (e.g. Roberts and Rosenthal, 2009; Sherlock et al., 2010). This could even be made 'local', using only the k-nearest neighbours in the KD-tree. It might also be possible to adaptively update the scaling of the proposal; however the mechanism to use is less obvious since, unlike in (e.g. Andrieu and Thoms, 2008; Sherlock et al., 2010; Vihola, 2012), there is no single optimal acceptance rate for our algorithm. Such adaptations would be a distraction to our main innovation and have not been implemented.

As we were revising this article, Kostov and Whiteley (2016) proposed an algorithm for estimating the variance of the estimator of the likelihood that comes from the particle filter. Although we do not pursue it here, this opens up the possibility of weighting each estimate of the likelihood according to the reciprocal of its variance as well as its distance from the

proposed θ^*.

5.1 Alternatives to k-nearest neighbours

As mentioned in Section 1, there are similarities between the non-pseudo-marginal version of
our algorithm and that of Conrad et al. (2014) (henceforth denoted CMPS). Here we discuss
the approach of CMPS, highlighting both similarities to and differences from our algorithm.
We also consider a general GP alternative to our k-NN implementation.

CMPS fit local linear-, quadratic- and Gaussian-process (GP)-based models in neigh-
bourhoods of candidate values, and at any given point in time their algorithm targets the
approximate posterior rather than the true posterior. However, the accuracy of the approxi-
mation is continually assessed and improved by carefully choosing further local points as the
algorithm proceeds so that, asymptotically, the algorithm targets the true posterior distri-
bution. Between adaptations our algorithm targets the true posterior distribution, but it is
perturbed every time an adaptation occurs. Thus our algorithm also asymptotically targets
the true posterior distribution. Both algorithms also use an increasing set of evaluations
of the true posterior. CMPS chooses the next evaluation of the true posterior by design
whereas our algorithm is "opportunistic" and potentially adds the value at each new point
evaluated; however it is an intelligent opportunist in the detail of how it deals with new
points which are very close to existing points.

Our algorithm could be extended to fitting local planes, quadratics or GPs to the k-
nearest neighbours in a similar manner to CMPS, however the weighted average is simpler,
quicker, and for it to be a sensible approach to take it only requires the posterior to be
bounded and continuous, rather than needing additional constraints. When it is the exact
log-likelihood that we are approximating these more complex local models might lead to an
improvement in the accuracy of the surrogate, at the small expense of possibly having to
use more nearest neighbours to fit. In the pseudo-marginal case, however, the efficiency will
depend on the variance in the log-likelihood estimates and the true changes in log-likelihood.
If the former outweighs the latter then there is little to be gained as the extra computation
may not lead to additional accuracy. The cost of re-estimating the GP hyper-parameters
as new training data become available is discussed in the final paragraph of this section.
One final difference between the two approaches is that CMPS focusses on the non-pseudo-

marginal case whereas our algorithm is applied in both pseudo-marginal and non-pseudo marginal settings.

Our likelihood estimate at a proposed point, θ', is a weighted average of the likelihoods at the k-nearest neighbours. An alternative approach would be to use a Gaussian process (GP) fitted to the set of $(\theta, \log p(y|\theta))$ pairs currently in the kd-tree. GPs have been used to approximate the log-likelihood previously. For example in Rasmussen (2003) and Fielding et al. (2011) a GP provides a cheap surrogate for Hamiltonian Monte Carlo calculations. Alterntively, in the context of ABC-MCMC, Wilkinson (2014) uses a Gaussian process to model the logarithm of the ABC approximate likelihood function, while Meeds and Welling (2014) approximate the joint synthetic likelihood at the current and proposed point using independent GPs for each summary statistic.

As with our approach, the final point estimate from a GP is a weighted average of existing values. However, by estimating the parameters of the GP one may represent the scales of variability more accurately and so obtain more accurate point estimates than with our inverse-distance weighting approach. The GP model also supplies an estimate of the uncertainty in the point estimate of the log-likelihood.

As in Conrad et al. (2014), the estimate of uncertainty allows for a choice of new training points to minimise the variance in some region of interest, with the potential of a large reduction in overhead. Algorithms such as those in Wilkinson (2014) and Meeds and Welling (2014) which use this approach target an approximation to the true posterior and an accurate GP approximation is essential for the algorithm to be useful. By design, our algorithm overlays a standard MH or PMMH algorithm and it automatically targets the true posterior; an inaccurate approximation reduces the mixing efficiency but does not invalidate the algorithm. Again, since our algorithm overlays a standard MH algorithm the true likelihood (or an unbiased estimate) must be evaluated at every accepted point and it is not immediately obvious how estimates of uncertainty might be used to circuvent this requirement while maintaining the true posterior as the target.

Finally, estimating the hyper-parameters of a GP is computationally very costly, and this estimation should be repeated as the training data set grows. For this reason and because of the difficulty in identifying $d(d+1)/2$ kernel range parameters, GP methods often use a diagonal covariance structure for the kernel, limiting the flexibility. Our approach of using

the covariance matrix of the sample from an initial training run to provide a map to a new parameter space where Euclidean distance is appropriate has a similar flavour to the pragmatic approach of using an initial training sample to fit the GP hyper-parameters and then keeping them fixed; alternatively, see Shen et al. (2006) for a partial solution.

Acknowledgements

References

Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle Markov chain Monte Carlo methods (with discussion). *J. R. Statist. Soc. B.*, 72(3):1–269.

Andrieu, C. and Roberts, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37:697–725.

Andrieu, C. and Thoms, J. (2008). A tutorial on adaptive MCMC. *Stat. Comput.*, 18(4):343–373.

Banterle, M., Grazian, C., Lee, A., and Robert, C. P. (2015). Accelerating Metropolis-Hastings algorithms by delayed acceptance. `http://arxiv.org/abs/1406.2660`.

Bentley, J. L. (1975). Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517.

Bliznyuk, N., Ruppert, D., Shoemaker, C., Regis, R., Wild, S., and Mugunthan, P. (2008). Bayesian calibration and uncertainty analysis for computationally expensive models using optimization and radial basis function approximation. *Journal of Computational and Graphical Statistics*, 17:270–294.

Boys, R. J., Wilkinson, D. J., and Kirkwood, T. B. L. (2008). Bayesian inference for a discretely observed stochastic-kinetic model. *Stat. Comput.*, 18:125–135.

Christen, J. A. and Fox, C. (2005). Markov chain Monte Carlo using an approximation. *Journal of Computational and Graphical Statistics*, 14:795–810.

Conrad, P. R., Marzouk, Y. M., Pillai, N. S., and Smith, A. (2014). Accelerating asymptotically exact MCMC for computationally intensive models via local approximations. `http://arxiv.org/abs/1402.1694`.

Craiu, R. V., Gray, L., Latuszy'nski, K., Madras, N., Roberts, G. O., and Rosenthal, J. S. (2015). Stability of adversarial Markov chains, with an application to adaptive MCMC algorithms. *Ann. Appl. Probab.*, 25(6):3592–3623.

Cui, T., Fox, C., and O'Sullivan, M. J. (2011). Bayesian calibration of a large-scale geothermal reservoir model by a new adaptive delayed acceptance Metropolis Hastings algorithm. *Water Resources Research*, 47:W10521.

Fearnhead, P., Giagos, V., and Sherlock, C. (2014). Inference for reaction networks using the Linear Noise Approximation. *Biometrics*, 70:457–466.

Fielding, M., Nott, D. J., and Liong, S.-Y. (2011). Efficient MCMC schemes for computationally expensive posterior distributions. *Technometrics*, 53:16–28.

Finkel, R. A. and Bentley, J. L. (1974). Quad trees - a data structure for retrieval on composite keys. *Acta Informatica*, 4:1–9.

Friedman, J. H., Bentley, J. L., and Finkel, R. A. (1977). An algorithm for finding best matches in logarithmic expected time. *ACM Trans. Math. Softw.*, 3(3):209–226.

Gillespie, D. T. (1977). Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81:2340–2361.

Golightly, A., Henderson, D. A., and Sherlock, C. (2015). Delayed acceptance particle MCMC for exact inference in stochastic kinetic models. *Statistics and Computing*, 25:1039–1055.

Golightly, A. and Wilkinson, D. J. (2005). Bayesian inference for stochastic kinetic models using a diffusion approximation. *Biometrics*, 61(3):781–788.

Golightly, A. and Wilkinson, D. J. (2011). Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus*, 1(6):807–820.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* Springer, New York, second edition.

Higdon, D., Reese, C. S., Moulton, J. D., Vrught, J. A., and Fox, C. (2011). Posterior exploration for computationally intensive forward models. In Brooks, S., Gelman, A., Jones, G. L., and Meng, X.-L., editors, *Handbook of Markov Chain Monte Carlo*, chapter 16, pages 401–418. Chapman & Hall/CRC, Boca Raton, FL.

Joseph, V. R. (2012). Bayesian computation using Design of experiments-based Interpolation technique. *Technometrics*, 54:209–225.

Joseph, V. R. (2013). A note on nonnegative DoIt approximation. *Technometrics*, 55:103–107.

Kennedy, M. C. and O'Hagan, A. (2001). Bayesian calibration of computer models (with discussion). *JRSSB*, 63:425–464.

Koblents, E. and Miguez, J. (2015). A population Monte Carlo scheme with transformed weights and its application to stochastic kinetic models. *Statistics and Computing*, 25(2):407–425.

Kostov, S. and Whiteley, N. (2016). An algorithm for approximating the second moment of the normalizing constant estimate from a particle filter. http://arxiv.org/abs/1602.02279.

Meeds, E. and Welling, M. (2014). GPS-ABC: Gaussian process surrogate approximate Bayesian computation. In *Thirtieth Conference on Uncertainty in Artificial Intelligence (UAI)*.

Overstall, A. M. and Woods, D. C. (2013). A strategy for Bayesian inference for computationally expensive models with application to the estimation of stem cell properties. *Biometrics*, 69:458–468.

Owen, J., Wilkinson, D. J., and Gillespie, C. S. (2015). Scalable inference for Markov processes with intractable likelihoods. *Statistics and Computing*, 25(1):145–156.

Payne, R. D. and Mallick, B. K. (2014). Bayesian big data classification: a review with complements. `http://arxiv.org/abs/1411.5653`.

Picchini, U. (2014). Inference for SDE models via Approximate Bayesian Computation. *Journal of Computational and Graphical Statistics*, 23(4):1080–1100.

Quiroz, M. (2015). Speeding up MCMC by delayed acceptance and data subsampling. `http://arxiv.org/abs/1507.06110`.

Rasmussen, C. E. (2003). Gaussian processes to speed up hybrid Monte Carlo for expensive Bayesian integrals. In Bernardo, J. M., Bayarri, M. J., Berger, J. O., Dawid, A. P., Heckerman, D., Smith, A. F. M., and West, M., editors, *Bayesian Statistics 7*, pages 651–659, Oxford. Oxford University Press.

Roberts, G. O. and Rosenthal, J. (2001). Optimal scaling for various Metropolis-Hastings algo- rithms. *Statistical Science*, 16:351–367.

Roberts, G. O. and Rosenthal, J. S. (2004). General state space Markov chains and MCMC algorithms. *Probab. Surv.*, 1:20–71.

Roberts, G. O. and Rosenthal, J. S. (2007). Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms. *J. Appl. Probab.*, 44(2):458–475.

Roberts, G. O. and Rosenthal, J. S. (2009). Examples of adaptive MCMC. *J. Comput. Graph. Statist.*, 18(2):349–367.

Roberts, G. O. and Tweedie, R. L. (1996). Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika*, 83(1):95–110.

Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. (1989). Design and analysis of computer experiments. *Statistical Science*, 4:409–435.

Shen, Y., Ng, A. Y., and Seeger, M. (2006). Fast Gaussian process regression using kd-trees. In *In Advances in Neural Information Processing Systems 18*. MIT Press.

Sherlock, C., Fearnhead, P., and Roberts, G. O. (2010). The random walk Metropolis: linking theory and practice through a case study. *Statist. Sci.*, 25(2):172–190.

Sherlock, C., Thiery, A., and Golightly, A. (2015a). Efficiency of delayed acceptance random walk Metropolis algorithms. *In prepartion.* `http://arxiv.org/abs/1506.08155`.

Sherlock, C., Thiery, A., Roberts, G. O., and Rosenthal, J. S. (2015b). On the efficiency of pseudo-marginal random walk Metropolis algorithms. *The Annals of Statistics*, 43(1):238–275.

Storer, J. (2002). *An introduction to data structures and algorithms.* Springer-Verlag, New York.

van Kampen, N. G. (2001). *Stochastic Processes in Physics and Chemistry.* North-Holland.

Vihola, M. (2012). Robust adaptive Metropolis algorithm with coerced acceptance rate. *Stat. Comput.*, 22(5):997–1008.

Wilkinson, D. J. (2012). *Stochastic Modelling for Systems Biology.* Chapman and Hall/CRC Press, London, 2nd edition.

Wilkinson, R. D. (2014). Accelerating ABC methods using Gaussian processes. In *JMLR Workshop and Conference Proceedings Volume 33: Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics.*

Supplementary material for Adaptive, delayed-acceptance MCMC for targets with expensive likelihoods

A   Operations on the KD-tree

In this appendix we describe the standard operations of creating a balanced KD-tree from a dataset and of searching through the KD-tree for the K-nearest neighbours.

A.1   Creating a balanced KD-tree from an initial dataset

Here we suppose the availability of n_0 pairs of a co-ordinate vector θ, and a vector of interest, v. The algorithm for creating the tree proceeds recursively until the number of leaves in the node is less than $2b$.

The recursive function commences at the root node with $d_{split} = 1$, and a data set $\mathcal{D}$, which is all of the data, so that $n = |\mathcal{D}| = n_0$.

1. If $n < 2b$ then finish this recursion; this node is a leaf node and contains data $\mathcal{D}$ and splitting component d_{split}.

2. Otherwise $n \geq 2b$, so find the median $\mu_{d_{split}}$ of the $|\mathcal{D}|$ scalar values for $\theta_{d_{split}}$.

3. Place all points for which $\theta_{d_{split}} < \mu_{d_{split}}$ into a set $\mathcal{L}$.

4. Place all points for which $\theta_{d_{split}} > \mu_{d_{split}}$ into a set $\mathcal{R}$.

5. If any points have $\theta_{d_{split}} = \mu_{d_{split}}$ then use independent Bernoulli trials with probability 0.5 to place each in $\mathcal{L}$ or $\mathcal{R}$.

6. Create a left and a right child for this node.

7. Recursively call this function for the left child using $\mathcal{D} = \mathcal{L}$ and $d_{split} = d_{split} \oplus 1$.

8. Recursively call this function for the right child using $\mathcal{D} = \mathcal{R}$ and $d_{split} = d_{split} \oplus 1$.

A.2 Finding the k nearest neighbours of a point

The algorithm below finds the k nearest neighbours to a point θ^*. For simplicity it, assumes that $k \leq b$; it is straightforward (but messy) to alter it to allow any k.

Stage One

1. Descend from the root node to find the leaf node to which θ^* would be added if we wished to add a new entry (θ^*, l_{θ^*}).

2. Find and store an ordered list of the k points on this leaf that are nearest to θ^*. Keep track of r_{max}, the distance of the k^{th} furthest point from θ^*.

Stage Two is only applied if the leaf note reached by Stage One has a parent; this is the case in all but the trivial tree.

Stage Two: This part of the algorithm is initialised with a list of k nearest neighbours and a value r_{max}, both obtained from Stage One. It uses an [`ascending/descending`] flag and commences at the parent of the leaf node reached by Stage One with the flag set to `ascending`. To aid the linguistic flow we use the verbs `ascend` and `descend` as shorthand for calling Stage Two with the flag set to `ascending` and `descending` respectively. If the current node is a branch, let θ_{split} and d_{split} refer to the current node.

1. If `ascending`:

    (i) If $\left| \theta_{split} - \theta^*_{d_{split}} \right| < r_{max}$ then set $\mathbf{u} = \theta_{split} \mathbf{e}_{d_{split}}$, where $\mathbf{e}_i$ is the k-vector with all elements set to zero except for the ith, which is set to one, then `descend` to the child (of this node) from which the algorithm has *not* just ascended.

    (ii) If $\left| \theta_{split} - \theta^*_{d_{split}} \right| \geq r_{max}$ and if this node has a parent then `ascend` to this parent.

2. If `descending` and the current node is a leaf then calculate the distance from every point on the leaf to $\boldsymbol{\theta}^*$ and sort the leaves according to this distance. If any are within r_{max} then update the list of the k nearest neighbours and update r_{max}.

3. If `descending` and the current node is a branch then if $\theta^*_{d_{split}} \leq \theta_{split}$ define the *near child* to be the left child, otherwise define it to be the right child. Similarly, define the *far child* to be whichever child is not the near child.

(i) `Descend` to the near child with $\mathbf{u}$ unchanged.

(ii) Update $u_{d_{split}}$ to θ_{split}, keeping the other elements of $\mathbf{u}$ unchanged. If $||\mathbf{u}|| < r_{max}$ then `descend` to the far child.

B   Proofs of Propositions 1 and 2

Proof of Proposition 1

Let B_ϵ^* and B_ϵ^i be the ϵ balls around θ^* and θ_i respectively and denote the empty set by ϕ. Firstly, since θ_i $(i = 1, \ldots, n)$ are identically distributed,

$$1 - \mathbb{P}\left(\theta_1 \notin B_\epsilon^*, \ldots, \theta_n \notin B_\epsilon^*\right) = \mathbb{P}\left(\theta_1 \in B_\epsilon^* \text{ or} \ldots \text{or } \theta_n \in B_\epsilon\right) \leq n\mathbb{P}\left(\theta_1 \in B_\epsilon\right) = E_{n,\epsilon}.$$

Hence $p_{keep} \geq 1 - E_{n,\epsilon}$.

Secondly, requiring that one point be outside of the ϵ ball of another point is equivalent to requiring that the $\epsilon/2$ balls of both points do not intersect. So, whatever the relationship between $\theta_1, \ldots, \theta_{i-1}$,

$$\mathbb{P}\left(\theta_i \notin B_\epsilon^* | \theta_1 \notin B_\epsilon^*, \ldots, \theta_{i-1} \notin B_\epsilon^*\right) = \mathbb{P}\left(B_{\epsilon/2}^i \cap B_{\epsilon/2}^* = \phi | B_{\epsilon/2}^1 \cap B_{\epsilon/2}^* = \phi, \ldots, B_{\epsilon/2}^{i-1} \cap B_{\epsilon/2}^* = \phi\right)$$
$$\leq \mathbb{P}\left(B_{\epsilon/2}^i \cap B_{\epsilon/2}^* = \phi\right)$$
$$= \mathbb{P}\left(\theta_i \notin B_\epsilon^*\right)$$

since the conditioned event only rules out portions of $\mathbb{R}^d \backslash B_{\epsilon/2}^*$ for B_ϵ^i. Hence

$$p_{keep} = \mathbb{P}\left(\theta_1 \notin B_\epsilon^*\right) \prod_{i=2}^{n} \mathbb{P}\left(\theta_i \notin B_\epsilon^* | \theta_1 \notin B_\epsilon^*, \ldots, \theta_{i-1} \notin B_\epsilon^*\right)$$
$$\leq \mathbb{P}\left(\theta_1 \notin B_\epsilon^*\right)^n = \left(1 - E_{n,\epsilon}/n\right)^n \leq e^{-E_{n,\epsilon}}.$$

Proof of Proposition 2

In this case the distance, D between any two points chosen independently satisfies $D^2/2 \sim \chi_d^2$. The θ_i, $(i = 1, \ldots, d)$ are not mutually independent since none can be within ϵ of any other, but this is irrelevant to the calculation that follows. Let $N_{n,\epsilon} := \sum_{i=1}^{n} \mathbf{1}_{\{\theta_i \in B(\theta^*, \epsilon)\}}$, then

$$E_{n,\epsilon} = \sum_{i=1}^{n} \mathbb{P}\left(\theta_i \in B(\theta_*, \epsilon)\right) = nF_{\chi_d^2}(\epsilon^2/2),$$

as required.

C Adaptive-KD-tree, delayed-acceptance pseudo-marginal algorithm

Algorithm 2: *adaptive-KD-tree, delayed-acceptance, pseudo-marginal Metropolis-Hastings.*

At the start of iteration n, let the current parameter value be θ and let the cheap approximation to the posterior be $\hat{\pi}_a(\theta)$ with the expensive unbiased estimate of the posterior denoted by $\hat{\pi}_s(\theta)$.

1. With probability β go to Step 2 (PsMMH) else go to Step 3 (da-PsMMH).

2. **PsMMH** Propose θ^* from $q(\theta^*|\theta)$. Evaluate $\hat{\pi}_s(\theta^*)$ by effectively proposing W^* from $\tilde{q}(W^*|\theta^*)$. Accept the proposal $(\theta \leftarrow \theta^*)$ with probability given by (9); otherwise reject the proposal $(\theta \leftarrow \theta)$. Go to Step 4.

3. **da-PsMMH** Propose θ^* from $q'(\theta^*|\theta)$.

   (a) **Stage 1**: Evaluate $\hat{\pi}_a(\theta^*)$ using the current KD-tree; with probability $\tilde{\alpha}_1(\theta, \theta^*)$ as defined in (2) (but with q' instead of q) proceed to Step 3b (Stage 2); otherwise reject the proposal $(\theta \leftarrow \theta)$, set $i_n = i_{n-1}$ and go to next iteration.

   (b) **Stage 2**: Evaluate $\hat{\pi}_s(\theta^*)$ by, effectively, proposing W^* from $\tilde{q}(W^*|\theta^*)$. Accept the proposal $(\theta \leftarrow \theta^*)$ with probability $\tilde{\alpha}_{2,s}(\theta, \theta^*)$ as defined in (10); otherwise reject the proposal $(\theta \leftarrow \theta)$. Go to Step 4.

4. Set $i_n = i_{n-1} + 1$; add $(\theta^*, \hat{\pi}_s(\theta^*))$ to a list of recently evaluated parameter/posterior pairs; with probability p_{i_n} add all pairs from this list to the KD-tree and then remove all pairs from the list; go to next iteration.

D Theorems 1, 2 and 3: discussion, statement and proofs

We discuss Assumptions 1 and 2, then place our algorithms in the more general framework in terms of which our proof of ergodicity is phrased. Next, we state and discuss Theorem 2 and then show, as Theorem 3, that it is not always necessary to force the adaptation

probabilities, p_i to tend to zero. Finally we prove Theorems 1, 2 and 3. For simplicity, for Theorems 1 and 2, we consider only the daPsMMH algorithm since the daMH is a special case of daPsMMH with $W = W^* = 1$.

D.1 Discussion of Assumptions 1 and 2

Roberts and Rosenthal (2004) show that any Metropolis-Hastings kernel on a compact state-space satisfies Assumption 1 provided $\pi(\theta)$ is continuous and $q(\theta^*|\theta)$ is continuous and positive. In the applications of interest to us all of the parameters are positive and it is common practice (e.g. Golightly and Wilkinson, 2005; Picchini, 2014; Koblents and Miguez, 2015; Owen et al., 2015) to place a vague but proper uniform prior on the logarithm of each parameter; furthermore, π is continuous. We may therefore choose q so that Assumption 1 is satisfied. If Θ is not compact then Assumption 1 still holds, for example, if the proposal density $q(\theta^*|\theta) = q(\theta^*)$ is an independence proposal which satisfies $q(\theta^*) > \delta\pi(\theta^*)$ since then $q(\theta^*|\theta)\alpha(\theta,\theta^*) > q(\theta^*) \wedge \delta\pi(\theta^*) = \delta\pi(\theta^*)$.

Whether or not Assumption 2 is satisfied depends upon the exact pseudo-marginal method used. Suppose, as in our examples, that an unbiased estimate of the likelihood, $\hat{P}(y|\theta)$, (and hence, up to a constant, of the posterior) is obtained using a bootstrap particle filter. This provides $\hat{P}(y|\theta)$ as a finite product of Monte Carlo averages of likelihood terms, where each of these terms arises from an assumed distribution for the error in the observation of some stochastic process Z_t. Suppose, for example, that $Y_t|z_t \sim N(z_t, \theta_1^2)$ and that the state space for Z_t is bounded, so that $0 \le a_t \le (y_t - z_t)^2 \le b_t < \infty$. Then $\log W \le (b_t - a_t)/(2\theta_1^2)$ and Assumption 2 will therefore hold provided the support for θ_1 is bounded away from zero.

D.2 More general set up

We set up the algorithm in more general terms since, in addition to the specific algorithms we have described, our ergodicity result, Theorem 1, is applicable to a general class of adaptive pseudo-marginal algorithms and, to our knowledge, is the first such result.

Let $\tilde{P}$ be a non-adaptive MH kernel on Θ with proposal $q(\theta^*|\theta)$, and let $\{\tilde{P}_\gamma\}_{\gamma \in \mathcal{G}}$ be a (usually infinite) set of kernels on Θ, where P_γ has proposal $q_\gamma(\theta^*|\theta)$. All kernels are assumed to have the same stationary density, $\pi(\theta)$. At iteration n the adaptive kernel that will be used

with probability $1 - \beta$ is $\tilde{P}_{\gamma_n}$ and, since the Markov chain is adaptive, let γ_n be a realisation of the random variable Γ_n, which depends on the history of the chain. At iteration n, our kernel is

$$\tilde{P}^*_{\gamma_n} = \beta\tilde{P} + (1-\beta)\tilde{P}_{\gamma_n} \quad \text{for some } \beta \in (0,1). \tag{D.1}$$

In our case, each $\tilde{P}_\gamma, (\gamma \in \mathcal{G})$ uses the same initial proposal $q'(\theta^*|\theta)$, and this proposal will typically differ from the proposal, $q(\theta^*|\theta)$ used in the fixed kernel. However, our proof of ergodicity applies in the more general set up of (D.1).

D.3   Theorem 2: choice of n and β as functions of α_1

Consider a collection of kernels, indexed by $\xi \in \Xi$, each a mixture of an adaptive kernel, indexed by ξ, and a common fixed kernel:

$$\tilde{P}^{*\xi}_{\gamma_n} = \beta\tilde{P} + (1-\beta)\tilde{P}^\xi_{\gamma_n} \quad \text{for some } \beta \in (0,1) \text{ and } \gamma_n \in G^\xi. \tag{D.2}$$

Let $\tilde{P}$ use a proposal q, such as that in (6). For a given value of ξ, all $\tilde{P}^\xi_\gamma$, $\gamma \in G^\xi$ use the same proposal, q^ξ; this could be the RWM proposal defined in (6), for instance.

Theorem 2. *Consider a set of adaptive daPsMRWM Algorithms indexed by ξ, as described in and around (D.2). Let the common fixed kernel, $\tilde{P}$ satisfy Assumption 1 and let all of the adaptive kernels have the same stationary density as $\tilde{P}$, given in (8); in particular, all use the same mechanism for generating W which satisfies Assumption 2. In addition to (5), let $\{p_i\}_{i\in\mathbb{N}}$ be a strictly decreasing sequence. For each ξ, after n iterations let I^ξ_n be the number of evaluations of the expensive posterior and let $\overline{\alpha}^\xi_{1,n}$ be the average Stage One acceptance rate. Assume that given $\epsilon > 0 \; \exists \; i_{crude}$ such that for all n with $\mathbb{E}\left[I^\xi_n\right] > i_{crude}$ and all ξ,*

$$\mathbb{P}\left(\overline{\alpha}^\xi_{1,n} < 2\overline{\alpha}^\xi\right) > 1 - \epsilon.$$

For some fixed $\kappa > 0$, we then set

$$\beta^\xi = \kappa\overline{\alpha}^\xi \tag{D.3}$$

Subject to the above $\mathbb{E}\left[I^\xi_n\right] \to \infty$ as $n \to \infty$. Further, for any $\epsilon > 0$ and all $\xi \in \Xi$ it is possible to choose a single $E > 0$ so that if $\mathbb{E}\left[I^\xi_n\right] > E$, for some n, then the TVD between the Markov chain and π after n iterations is bounded by ϵ.

36

Thus, provided that after a certain number of expensive iterations the acceptance rate is typically no more than a constant multiple of the long term average, $\overline{\alpha}^\xi$, it is safe to set the probability of choosing the fixed kernel proportional to $\overline{\alpha}^\xi$.

As the chain progresses and the number of expensive evaluations of the posterior increases, the representation of the posterior by the KD-tree improves and $\overline{\alpha}_{1,n}$ initially increases (with a typical relative change of around 5-10%) before settling down. In practice, we find the condition on the convergence of the Stage 1 acceptance probabilities to be a reasonable assumption.

D.4   Diminishing probabilities of adaptation

After every expensive iteration, Algorithms 1 and 2 store the newly-evaluated expensive posterior; they ensure that the diminishing-adaptation condition (see Theorem 5) is satisfied by, after the ith expensive iteration, adding all stored expensive posteriors to the KD-tree with probability $p_i \to 0$. Tables 2 and 3 suggest that the algorithm is more efficient the more slowly $p_i \downarrow 0$. The intuition behind this, and the possible consequences for higher-dimensional systems if p_i were to remain bounded away from 0 are discussed in Section 5; the concern is that such an algorithm might not be ergodic. In Appendix D.5.3 we prove the following.

Theorem 3. *Let $\pi(\theta)$ be a continuous density with respect to Lebesgue measure on a hyper-rectangular state space, Θ, and consider an adaptive, delayed-acceptance MH algorithm the same as Algorithm 1 but with the following alterations: $\forall i \in \mathbb{N}$, $p_{2i} = 1$ and $p_{2i+1} = 0$, and we remove the restriction on the growth of the tree described in Section 2.6. If the kernel satisfies the minorisation condition, Assumption 1, then the diminishing adaptation condition is automatically satisfied.*

D.5   Proofs of Theorems 1, 2, and 3

Throughout this section the total variation distance (TVD) between any two probability measures, $\nu(\cdot)$ and $\pi(\cdot)$, is denoted $||\nu - \pi||$. We require the following definition and two results.

Definition 1. Consider a Markov kernel P on a statespace $\mathcal{X}$. A subset $C \subseteq \mathcal{X}$ is *small* if there exists a positive integer n_0, $\epsilon > 0$, and a probability measure $\nu(\cdot)$ on $\mathcal{X}$ such that the following *minorisation* condition is satisfied

$$P^{n_0}(x, \cdot) \geq \epsilon \nu(\cdot). \tag{D.4}$$

Theorem 4. *(Roberts and Rosenthal, 2004) Consider a Markov chain with invariant probability distribution $\pi(\cdot)$. Suppose that the entire statespace is small (i.e. (D.4) is satisfied with $C = \mathcal{X}$). Then the chain is uniformly ergodic, and in fact $||P^n(x, \cdot) - \pi(\cdot)|| \leq (1 - \delta)^{[n/n_0]}$*

Theorem 5. *(Roberts and Rosenthal, 2007) Consider an adaptive MCMC algorithm on a statespace $\mathcal{X}$ with adaptive kernels P_γ, $\gamma \in \mathcal{G}$, and with $\pi(\cdot)$ stationary for each P_γ; $\gamma \in \mathcal{G}$. Under the following conditions the adaptive algorithm is ergodic.*

1. *(Simultaneous uniform ergodicity) For all $\epsilon > 0$, there is $N = N(\epsilon) \in \mathbb{N}$ such that $\left|\left|P_\gamma^N(x, \cdot) - \pi(\cdot)\right|\right|_{TV} \leq \epsilon$ for all $x \in \mathcal{X}$ and $\gamma \in \mathcal{G}$.*

2. *(Diminishing adaptation) For any $X_0 = x_0$, $\Gamma_0 = \gamma_0$,*

$$\sup_{x \in \mathbb{X}} \left|\left|P_{\Gamma_{n+1}}(x, \cdot) - P_{\Gamma_n}(x, \cdot)\right|\right|_{TV} \xrightarrow{p} 0,$$

   where the convergence in probability is with respect to the distribution of Γ_n and Γ_{n+1} given x_0 and γ_0.

D.5.1 Proof of Theorem 1

We now show that a class of adaptive Metropolis-Hastings algorithms, which includes our algorithms, is ergodic subject to Assumptions 1 and 2. In line with our KD-tree approximation, we assume that the target, π, and the proposals, q, q' and q_θ ($\theta \in \Theta$), are all densities with respect to Lebesgue measure.

Any delayed-acceptance algorithm is simply an accept/reject Markov chain with a non-standard acceptance probability. The key point is that detailed balance is preserved since (for the pseudo-marginal version)

$$\pi(\theta)q_\theta(w)w \ q(\theta^*|\theta)q_{\theta^*}(w^*) \ \tilde{\alpha}_1(\theta, \theta^*)\tilde{\alpha}_{2,PM}([\theta, w], [\theta^*, w^*])$$

$$= \hat{\pi}_c(\theta)q(\theta^*|\theta)\tilde{\alpha}_1(\theta, \theta^*) \times q_\theta(w)q_{\theta^*}(w^*) \times \frac{\pi(\theta)w}{\hat{\pi}_c(\theta)}\tilde{\alpha}_{2,PM}([\theta, w], [\theta^*, w^*]),$$

and each of the three terms in the product is invariant to $(\theta, w) \leftrightarrow (\theta^*, w^*)$. We therefore prove ergodicity (subject to conditions) for any adaptive pseudo-marginal algorithm of the form given in and above (D.1).

We require the following components. Here A and $\tilde{A}$ denote any (Lebesgue) measurable subsets of Θ and $\Theta \times \mathcal{W}$ respectively, and δ represents the Dirac delta function.

1. A fixed pseudo-marginal kernel on $\Theta \times \mathcal{W}$ with stationary density $\pi(\theta)q_\theta(w)w$:

$$\tilde{P}\left([\theta, w], \tilde{A}\right) = (1 - \overline{\alpha}_{PM}([\theta, w]))\, \delta_A([\theta, w])$$
$$+ \int_{\tilde{A}} \mathrm{d}\theta^*\ \mathrm{d}w^*\ q\left(\theta^*|\theta\right) q_{\theta^*}\left(w^*\right) \alpha_{PM}\left([\theta, w], [\theta^*, w^*]\right).$$

Here $\overline{\alpha}_{PM}([\theta, w])$ is the acceptance probability from the current value:

$$\overline{\alpha}_{PM}([\theta, w]) = \int_{\Theta \times \mathcal{W}} \mathrm{d}\theta^*\ \mathrm{d}w^*\ q(\theta^*|\theta)q_{\theta^*}(w^*)\alpha_{PM}([\theta, w], [\theta^*, w^*]).$$

2. The corresponding fixed 'ideal' kernel on Θ with stationary density $\pi(\theta)$,

$$P(\theta, A) = (1 - \overline{\alpha}_{MH}(\theta))\delta_A(\theta) + \int_A \mathrm{d}\theta^*\ q(\theta^*|\theta)\alpha_{MH}(\theta, \theta^*),$$

from which we are unable to sample because $\pi(\theta)$ and $\pi(\theta^*)$ are needed in order to evaluate α_{MH}. Here $\overline{\alpha}_{MH}(\theta) = \int_\Theta \mathrm{d}\theta^*\ q(\theta^*|\theta)\alpha_{MH}(\theta, \theta^*)$.

3. A set of additional (pseudo-marginal) kernels on $\Theta \times \mathcal{W}$: $\{\tilde{P}_\gamma([\theta, w], \cdot)\}_{\gamma \in \mathcal{G}}$, as described in and above (D.1), and with the same stationary density as $\tilde{P}$.

4. A sequence of probabilities p_n satisfying (5).

The generic algorithm is then:

Algorithm 2b: *generic, adaptive, pseudo-marginal, propose and accept/reject algorithm.*

Iteration n commences with current value $[\theta^{(n)}, w^{(n)}]$ and kernel index γ_n and involves the following two steps.

1. Sample $[\theta^{(n+1)}, w^{(n+1)}]$ from $\tilde{P}^*_{\gamma_n}$ as defined in (D.1).

2. With probability p_n update γ_n (i.e. update the adaptive kernel) by including all relevant information obtained since the kernel was last updated to create a new kernel.

Theorem 6. *Subject to Assumptions 1 and 2, Algorithm 2b is ergodic.*

Proof The condition (5) ensures that the diminishing adaptation condition of Theorem 5 is satisfied.

We next show that subject to Assumption 1 , $\tilde{P}$ satisfies a similar condition to Assumption 1 and hence so does each of the kernels in (D.1). This is then shown to ensure simultaneous uniform ergodicity.

We first define

$$\tilde{\nu}(\theta, w) := \nu(\theta)\frac{1}{c}q_\theta(w)w,$$

which is a density by (7); we refer to the corresponding measure as $\tilde{\nu}(\cdot)$. From (D.6),

$$q(\theta^*|\theta)q_{\theta^*}(w^*)\alpha_{PM}\left([\theta, w], [\theta^*, w^*]\right) > q(\theta^*|\theta)q_{\theta^*}(w^*)\frac{w^*}{\overline{w}}\alpha_{MH}(\theta, \theta^*) = \frac{c\delta}{\overline{w}}\tilde{\nu}(\theta^*, w^*).$$

This implies that for any γ and any measurable set $\tilde{A} \in \Theta \times [0, \overline{w}]$,

$$P_\gamma^*([\theta, w], \cdot) \geq \frac{c\beta\delta}{\overline{w}}\tilde{\nu}(\cdot).$$

Hence, the entire-statespace $(\Theta \times [0, \overline{w}])$ is small with $n_0 = 1$ and

$$\epsilon = \tilde{\delta} := c\beta\delta/\overline{w}; \tag{D.5}$$

$\tilde{\delta} < 1$ since $c/\overline{w} \leq 1$. Each kernel therefore individually satisfies the condition of Theorem 4 and hence $N = [\log \epsilon/(1 - \tilde{\delta})] + 1$ ensures that the collection of all kernels P_γ^* satisfy Condition 1 of Theorem 5. ∎

Although the kernel in Algorithm 2b is more general than that in Algorithm 2, in one particular sense Algorithm 2 is not a special case of Algorithm 2b, since the former potentially updates the kernel only after each expensive evaluation rather than after each iteration. Adaptation times enter the proof of Theorem 6 through the *diminishing adaptation* condition (see Theorem 5 of this article), for which it suffices that the probability of a change in the kernel at any given iteration, n, tends to zero as $n \to \infty$. For Algorithm 2 this is guaranteed through Condition (5); however it also holds for Algorithm 1 since $\lim_{n\to\infty} i_n = \infty$ almost surely, as we now demonstrate. Subject to Assumption 2, the acceptance probability for the fixed kernel between $[\theta, w]$ and $[\theta^*, w^*]$ is

$$\alpha_{PM}\left([\theta, w], [\theta^*, w^*]\right) \geq \alpha_{MH}\left(\theta, \theta^*\right)\left(1 \wedge \frac{w^*}{w}\right) \geq \frac{w^*}{\overline{w}}\alpha_{MH}(\theta, \theta^*). \tag{D.6}$$

By (7), the average acceptance rate from $[\theta, w]$ for the fixed kernel is

$$\overline{\alpha}_{PM}\left([\theta, w]\right) \geq \frac{c}{\overline{w}}\mathbb{E}_{\theta^*}\left[\alpha_{MH}(\theta, \theta^*)\right] = \frac{c\overline{\alpha}_{MH}(\theta)}{\overline{w}}.$$

The idealised Metropolis-Hastings kernel $P(\theta, \cdot)$ is uniformly ergodic by Assumption 1 (Roberts and Rosenthl (2004)) so $\overline{\alpha}_{MH}(\theta)$ is bounded below by some $\alpha_0 > 0$ (Roberts and Tweedie (1996) Proposition 5.1). Hence the overall acceptance rate is bounded below by $\beta\alpha_0 c/\overline{w}$, and by the strong law of large numbers, $\lim_{n\to\infty} i_n = \infty$.

D.5.2 Proof of Theorem 2

The proof of ergodicity of adaptive MCMC algorithms in Theorem 5 of Roberts and Rosenthal (2007) relies on a hypothetical Markov chain (for us, $x' := \{[\theta', w']_i\}_{i\in\mathbb{N}}$), which is identical to the real chain up until some iteration n_0 and then continues in parallel with the real chain using the kernel at n_0 without any further adaptation. The kernels for this chain are $\tilde{P}_{\gamma'_n}$, where

$$\gamma'_n = \begin{cases} \gamma_n & n \leq n_0 \\ \gamma_{n_0} & n > n_0. \end{cases}$$

After iteration n_0, the hypothetical chain clearly has π as its stationary distribution. Theorem 4 then informs us that after a further n_1 iterations

$$\left\|\tilde{P}^{n_1}_{\gamma_{n_0}}(x, \cdot) - \pi(\cdot)\right\| \leq (1 - \tilde{\delta})^{n_1}, \tag{D.7}$$

for any x, including the $[\theta, w]$ value of the chain after n_0 iterations; the hypothetical chain is close to the target. Diminishing adaptation is then used to show that if n_0 is large enough then after these further n_1 iterations the real chain is close to the hypothetical chain.

We use the same approach as in Roberts and Rosenthal (2007) and show that for a given required TVD between the true chain and π, the required run length can be specified in terms of the expected number of expensive iterations. Since Stage Two acceptance probabilities are irrelevant to our argument α_n^ξ and $\overline{\alpha}_n^\xi$ henceforth denote, respectively, the Stage One acceptance probability and its average after n iterations. Throughout this proof, for simplicity of notation, we will suppress the superscript ξ from the terms α_n^ξ, $\overline{\alpha}_n^\xi$, $\overline{\alpha}^\xi$, n_0^ξ, n_1^ξ, I^ξ, J^ξ. The point is that the rate of the convergence in (D.10) and the bound in (D.13) only depend on ξ through the expected number of expensive evaluations, I^ξ.

Let I_n and J_n be, respectively, the number of expensive evaluations of the true algorithm between iterations 1 and n and between iterations $n_0 + 1$ and $n_0 + n$. First we define the following events:

$$\mathcal{A}_n := \{\overline{\alpha}_n < 2\overline{\alpha}\} \quad \text{and} \quad \mathcal{B}_{n_0,n_1} := \{1 - (1 - p_{I_{n_0}})^{J_{n_1}} < \epsilon\}. \tag{D.8}$$

The probability that an iteration involves an evaluation of the expensive posterior is

$$\rho_n := \beta + (1 - \beta)\alpha_n = \overline{\alpha}\left(\kappa + (1 - \kappa\overline{\alpha})\frac{\alpha_n}{\overline{\alpha}}\right).$$

Thus

$$\mathbb{E}\left[I_n\right] = \sum_{j=1}^{n} \rho_j = n\overline{\alpha}\left(\kappa + (1 - \kappa\overline{\alpha})\frac{\overline{\alpha}_n}{\overline{\alpha}}\right). \tag{D.9}$$

So $\mathbb{E}\left[I_n\right] > n\overline{\alpha}\kappa$ and $\mathbb{E}\left[I_n\right] \to \infty$. Now, $\mathrm{Var}[I_n - I_{n-1}] = \rho_n(1 - \rho_n) \leq \rho_n$ and whether or not each iteration is expensive is a sequence of independent Bernoulli trials, so $\mathrm{Var}[I_n] \leq \mathbb{E}\left[I_n\right]$. Combined with Chebyshev's inequality shows that as $\mathbb{E}\left[I_n\right] \to \infty$, $I_n \to \infty$ in probability and hence that the diminishing adaptation probabilities satisfy

$$\mathbb{E}\left[p_{I_n}\right] \to 0. \tag{D.10}$$

By assumption, for $\mathbb{E}\left[J_n\right] > i_{crude}$, $\mathbb{P}\left(\mathcal{A}_n\right) > 1 - \epsilon$. Further, from (D.9), conditional on $\mathcal{A}_n$, $\mathbb{E}\left[J_n\right] < n\overline{\alpha}\left(\kappa + 2\right)$. Thus

$$\text{Conditional on } \mathcal{A}_n, \ n\overline{\alpha} > \frac{\mathbb{E}\left[J_n\right]}{\kappa + 2}. \tag{D.11}$$

The TVD between the true adaptive chain and the hypothetical chain described in the preliminaries for this proof is bounded above by the probability that they are not coupled (e.g Roberts and Rosenthal, 2004), and since adaptation can only occur after the true posterior (or an unbiased estimate thereof) is evaluated, the TVD between the real chain and the hypothetical chain after $n_0 + n_1$ iterations is less than $1 - (1 - p_{I_{n_0}})^{J_{n_1}}$. Given (D.7), the triangle inequality and the monotonicity of p_i, the following two conditions, therefore, guarantee that the TVD between the true chain and π is less than 2ϵ: $\mathcal{B}_{n_0,n_1}$ and

$$(1 - \tilde{\delta})^{n_1} < \epsilon. \tag{D.12}$$

By (D.3) and (D.5), $\tilde{\delta} = \overline{\alpha}c\kappa\delta/\overline{w} < 1$ since $c/\overline{w} \leq 1$. Also $\log(1 - \tilde{\delta}) < -\tilde{\delta}$. Thus,

$$n_1 \log(1 - \tilde{\delta}) < -n_1\tilde{\delta} = -\frac{n_1\overline{\alpha}c\kappa\delta}{\overline{w}} < -\mathbb{E}\left[J_{n_1}\right]\frac{c\kappa\delta}{\overline{w}(\kappa + 2)},$$

by (D.11) applied to the iterations from $n_0 + 1$. Hence, conditional on $\mathcal{A}_{n_1}$, (D.12) can be guaranteed by fixing $\mathbb{E}[J_{n_1}]$ (i.e. choosing n_1)

$$\mathbb{E}\left[J_{n_1}\right] > -\frac{\overline{w}(\kappa + 2)\log \epsilon}{c\kappa\delta}. \tag{D.13}$$

To deal with $\mathcal{B}_{n_0,n_1}$, for this fixed $\mathbb{E}\left[J_{n_1}\right]$, apply Jensen's inequality twice and then (D.10):

$$\mathbb{E}\left[(1 - p_{I_{n_0}})^{J_{n_1}}\right] \geq \mathbb{E}\left[(1 - p_{I_{n_0}})^{\mathbb{E}\left[J_{n_1}\right]}\right] \geq \mathbb{E}\left[1 - p_{I_{n_0}}\right]^{\mathbb{E}\left[J_{n_1}\right]} \to 1$$

as $\mathbb{E}\left[I_{n_0}\right] \to \infty$. Hence, by Markov's inequality, for sufficiently large $\mathbb{E}\left[I_{n_0}\right]$ we can ensure $\mathbb{P}\left(\mathcal{B}_{n_0,n_1}\right) > 1 - \epsilon$.

With these choices of $\mathbb{E}\left[I_{n_0}\right]$ and $\mathbb{E}\left[J_{n_1}\right]$, and hence of $E = \mathbb{E}\left[I_{n_0}\right] + \mathbb{E}\left[J_{n_1}\right]$, $\mathcal{A}_{n_1}$ and $\mathcal{B}_{n_0,n_1}$ each holds with probability $1 - \epsilon$, so the TVD between the true chain and π is less than 4ϵ.

D.5.3   Proof of Theorem 3

We denote the k-nearest neighbour approximation to the posterior at θ^* after n iterations by $\hat{\pi}_n(\theta^*)$ and a ball of radius r centred at θ by $B_r(\theta)$. We denote the Stage One and Stage Two acceptance probabilities at iteration n by $\alpha_{1,n}(\theta, \theta^*)$ and $\alpha_{2,n}(\theta, \theta^*)$.

For $\theta \in \Theta$ let

$$A_n(\theta, r) := \{\text{after } n \text{ iterations all } k \text{ nearest neighbours to } \theta \text{ lie within } B_r(\theta)\},$$

and note that $A_{n-1}(\theta, r) \Rightarrow A_n(\theta, r)$. If $A_n(\theta, r)$ holds then then any change (adaptation) in the cheap estimate $\hat{\pi}_n(\theta)$ must occur through one or more new points being added to the tree inside $B_r(\theta)$. Since $\log \pi$ is continuous, for any $\epsilon > 0 \; \exists \; \epsilon_\circ$ such that if $A_{n-1}(\theta, \epsilon_\circ)$ and $A_{n-1}(\theta^*, \epsilon_\circ)$ hold, then

$$1 - \epsilon < \frac{\hat{\pi}_n(\theta)}{\hat{\pi}_{n-1}(\theta)} < 1 + \epsilon \quad \text{and} \quad 1 - \epsilon < \frac{\hat{\pi}_{n+1}(\theta^*)}{\hat{\pi}_n(\theta^*)} < 1 + \epsilon.$$

P_{γ_n} and $P_{\gamma_{n+1}}$ differ only in their acceptance probabilities, with P_{γ_n} using the ratio of $\hat{\pi}_n(\theta^*)$ and $\hat{\pi}_{n-1}(\theta)$ in both $\alpha_{1,n}$ and $\alpha_{2,n}$. Yet, subject to $A_{n-1}(\theta, \epsilon_\circ)$ and $A_{n-1}(\theta^*, \epsilon_\circ)$,

$$(1 - \epsilon)^4 < \frac{\alpha_{1,n+1}(\theta, \theta^*)\alpha_{2,n+1}(\theta, \theta^*)}{\alpha_{1,n}(\theta, \theta^*)\alpha_{2,n}(\theta, \theta^*)} < (1 + \epsilon)^4.$$

Since Θ is compact, $\log \pi$ is uniformly continuous. Hence, for small enough ϵ, and subject to $\cap_{\theta \in \Theta} A_{n-1}(\theta, \epsilon_\circ)$,

$$||P_{\gamma_{n+1}}(\theta, \cdot) - P_{\gamma_n}(\theta, \cdot)|| < 5\epsilon.$$

We show, given any $\epsilon > 0$, $\exists n_\epsilon$ such that $\mathbb{P}\left(\cap_{\theta^* \in \Theta} A_n(\theta^*, \epsilon_\circ)\right) > 1 - 2\epsilon$, for all $n \geq n_\epsilon$. So for $n > n_\epsilon$, $\sup_\theta ||P_{\gamma_{n+1}}(\theta, \cdot) - P_{\gamma_n}(\theta, \cdot)|| < 7\epsilon$.

First, partition the (hyperrectangular) state space, Θ, into $n_\square$ hypercubes of size $\epsilon_\square$, such that any ball of radius $\epsilon_\circ$ must contain at least one hypercube. We now place n points uniformly at random in Θ (i.e. according to a homogeneous Poisson process, U). Denote the number of points that fall in the ith hypercube by $\square_i$. For any fixed $k \in \mathbb{N}$, $\mathbb{P}\left(\exists\, i \in \{1, \ldots, n_\square\} \text{ such that } \square_i < k\right) \leq n_\square \mathbb{P}\left(\square_1 < k\right) \to 0$ as $n \to \infty$. Hence there is an $n_\bullet$ such that for all $n \geq n_\bullet$, $\mathbb{P}\left(\exists\, i \text{ such that } \square_i < k | n \text{ points from } U\right) < \epsilon$.

The minorisation condition holds over the whole state space, and each kernel P_γ is reversible; moreover adaptation only occurs on even-numbered expensive iterations so that each kernel is used (at least) twice before adaptation. Lemma 26 of Craiu et al. (2015) then implies that for each pair of iterations there is a probability of at least $\beta^2 \delta^2 / 4$ of sampling from π. Since $\log \pi$ is continuous and Θ is compact,

$$\rho := \frac{\min_{\theta \in \Theta} \pi(\theta)}{\max_{\theta \in \Theta} \pi(\theta)} > 0$$

So whenever a sample from π is obtained, a sample from the homogeneous Poisson process, U, may be obtained with a probability of at least ρ. Define the event

$$C_n := \{\exists \text{ at least one sample from } U \text{ in } n \text{ iterations}\}.$$

Then $\mathbb{P}\left(C_n^c\right) \leq (1 - \rho \beta^2 \delta^2 / 4)^{n/2}$ and, given $\epsilon > 0$ and $n_\bullet$ there is an n_{once} such that for all $n \geq n_{once}$, $\mathbb{P}\left(C_n^c\right) < \epsilon/(n_\bullet)$. Hence for all $n \geq n_\epsilon := n_{once} n_\bullet$,

$$\mathbb{P}\left(\text{after } n \text{ iterations } \exists\, i \in \{1, \ldots, n_\square\} \text{ such that } \square_i < k\right) < 2\epsilon.$$

If there are at least k entries in each hypercube then, since all balls of radius r contain at least one hypercube, the k nearest neighbours to all $\theta^* \in \Theta$ must be within $B_r(\theta^*)$.

Label	Reaction	Hazard	Description
R_1	$\mathcal{X}_1 \xrightarrow{\nu_1} 2\mathcal{X}_1$	$\nu_1 X_1$	Prey reproduction
R_2	$\mathcal{X}_1 + \mathcal{X}_2 \xrightarrow{\nu_2} 2\mathcal{X}_2$	$\nu_2 X_1 X_2$	Prey death, predator reproduction
R_3	$\mathcal{X}_2 \xrightarrow{\nu_3} \emptyset$	$\nu_3 X_2$	Predator death

Table E.1: Reaction list and hazards for the Lotka-Volterra system.

Label	Reaction	Hazard	Description
R_1	$\mathsf{DNA} + \mathsf{P}_2 \xrightarrow{\nu_1} \mathsf{DNA} \cdot \mathsf{P}_2$	$\nu_1 X_1 X_4$	Dimer binding
R_2	$\mathsf{DNA} \cdot \mathsf{P}_2 \xrightarrow{\nu_2} \mathsf{DNA} + \mathsf{P}_2$	$\nu_2 (k - X_1)$	Dimer unbinding
R_3	$\mathsf{DNA} \xrightarrow{\nu_3} \mathsf{DNA} + \mathsf{RNA}$	$\nu_3 X_1$	Transcription
R_4	$\mathsf{RNA} \xrightarrow{\nu_4} \mathsf{RNA} + \mathsf{P}$	$\nu_4 X_2$	Translation
R_5	$2\mathsf{P} \xrightarrow{\nu_5} \mathsf{P}_2$	$\nu_5 X_3 (X_3 - 1)/2$	Forward dimerisation
R_6	$\mathsf{P}_2 \xrightarrow{\nu_6} 2\mathsf{P}$	$\nu_6 X_4$	Reverse dimerisation
R_7	$\mathsf{RNA} \xrightarrow{\nu_7} \emptyset$	$\nu_7 X_2$	RNA degradation
R_8	$\mathsf{P} \xrightarrow{\nu_8} \emptyset$	$\nu_8 X_3$	Protein degradation

Table E.2: Reaction list and hazards for the auto-regulatory system.

E   Simulation study: model details and further results

E.1   Model and inference details and data simulation

Tables E.1 and E.2 list the reactions and associated hazards for the Lotka Volterra and autoregulatory examples, respectively.

A single data set was simulated from the Lotka-Volterra MJP using an initial value of $X_0 = (71, 79)$ and parameter values taken from Wilkinson (2012), that is $\nu = (1.0, 0.005, 0.6)$. Each X_t, $(t = 1, \ldots, 50)$ was corrupted as in equation (11) with $\sigma_1 = \sigma_2 = 8$.

For the autoregulatory system we generated 2 synthetic datasets (labelled as $\mathcal{D}_1$ and $\mathcal{D}_2$) by taking $X_0 = (5, 8, 8, 8)$, $K = 10$ and using parameter values taken from Golightly and Wilkinson (2011), that is $\nu = (0.1, 0.7, 0.35, 0.2, 0.1, 0.9, 0.3, 0.1)$. Simulated values were corrupted with noise as in (11) with $\sigma_1 = \sigma_2 = 0.5$ and $\sigma_3 = \sigma_4 = 1$. Dataset $\mathcal{D}_1$ consists of 101 observations on the time interval $[0, 100]$ and $\mathcal{D}_2$ consists of 201 observations on the time interval $[0, 1000]$.

For the autoregulatory system, the total number K of $\mathsf{DNA} \cdot \mathsf{P}_2$ and DNA is fixed through-

out the evolution of the system, for in our inferences it is assumed to be known, so that the model comprises of 4 species. We denote the number of molecules of DNA, RNA, P and P_2 as X_1, X_2, X_3 and X_4 respectively. As noted by Golightly and Wilkinson (2011), the rate constants in the reversible reactions can be difficult to infer and we therefore fix ν_1 and ν_5 at the ground truth. We consider inference for the remaining rate constants and the observation error standard deviations, 10 parameters in all.

E.2   Tuning parameters for the fixed kernels

For inference on the Lotka Volterra system, we follow the practical advice of Sherlock et al. (2015b) by choosing the number of particles, m, so that the variance (τ^2) in the log-posterior at the median (estimated from the pilot run) and 4 additional sampled parameter values is less than 3. We set $m = 200$ since this gave $\tau^2 \in [2.06, 3.07]$. Under certain assumptions regarding target and the variance in the log-posterior, in the case that the target is approximately Gaussian, Sherlock et al. (2015b) found that the scaling of the proposal variance should be approximately $\lambda = (2.56^2/d)$ to optimise efficiency. We found further scaling this quantity by 1.1 appeared to give optimal performance in terms of effective sample size (ESS) per second giving $V_{fixed} = 1.1 \times (2.56^2/5) \times \hat{\Sigma}$.

For the autoregulatory example, following Roberts and Rosenthal (2001), we found that scaling variance estimated from the tuning run by $\lambda = (2.38^2/d)$ with $d = 10$ gave optimal performance of the fixed kernel (ie without delayed acceptance).

The choices of $\epsilon = 0.3065$ for the Lotka Volterra and $\epsilon = 0.982$ for the autoregulatory system followed from considerations described immediately after (12).

F   The linear noise approximation

The linear noise approximation (LNA) of the Markov jump process defined by the reaction hazards (e.g. Fearnhead et al., 2014) ignores discreteness (but not stochasticity) and gives the state X_t as a Gaussian: $X_t \sim N(z_t + m_t, V_t)$, where z_t, m_t and V_t satisfy a coupled

ODE system

$$
\begin{cases}
\dot{z}_t = S\,h(z_t, \nu) \\
\dot{m}_t = F_t m_t \\
\dot{V}_t = V_t F_t^T + S\mathrm{diag}\left\{h(z_t, \nu)\right\} S^T + F_t V_t
\end{cases}
\tag{F.1}
$$

Here, $h(x_t, \nu)$ is the length-8 column vector containing the reaction hazards, F_t is a 4×4 matrix whose (i, j)th entry is given by the first partial derivative of the ith component of $S\,h(z_t, \nu)$ with respect to the jth component of z_t and S is the 4×8 stoichiometry matrix whose (i, j)th element gives the effect of reaction j on species i.

F.1   Marginal likelihood under the linear noise approximation

For simplicity of exposition we assume an observation regime of the form

$$
Y_t = X_t + \epsilon_t, \qquad \epsilon_t \sim \mathrm{N}\,(0, \Sigma)
$$

where ϵ_t is a length-d_x Gaussian random vector and $t = 0, 1, \ldots, n$. Suppose that X_1 is fixed at some value x_1. The marginal likelihood $\pi(y_{1:n}|\theta)$ (and hence the posterior up to proportionality) under the LNA can be obtained as follows.

1. Initialisation. Compute

$$
\pi(y_1|\theta) = \phi\,(y_1\,;\,x_1\,,\,\Sigma)
$$

where $\phi\,(y_1\,;\,x_1\,,\,\Sigma)$ denotes the Gaussian density with mean vector x_1 and variance matrix Σ. Set $a_1 = x_1$ and C to be the $d_x \times d_x$ matrix of zeros.

2. For times $t = 1, 2, \ldots, n - 1$,

    (a) Prior at $t + 1$. Initialise the LNA with $z_t = a_t$, $m_t = 0$ and $V_t = C_t$. Note that $m_s = 0$ for all $s > t$. Integrate the ODE system (F.1) forward to $t + 1$ to obtain z_{t+1} and V_{t+1}. Hence

$$
X_{t+1}|y_{1:t}, \theta \sim N(z_{t+1}, V_{t+1})\,.
$$

    (b) One step forecast. Using the observation equation, we have that

$$
Y_{t+1}|y_{1:t}, \theta \sim N\,(z_{t+1}, V_{t+1} + \Sigma)\,.
$$

Compute

$$\pi(y_{1:t+1}|\theta) = \pi(y_{1:t}|\theta)\,\phi\left(y_{t+1}\,;\,z_{t+1}\,,\,V_{t+1}+\Sigma\right).$$

(c) Posterior at $t+1$. Combining the distributions in (a) and (b) gives $X_{t+1}|y_{1:t+1}, \theta \sim N(a_{t+1}, C_{t+1})$ where

$$a_{t+1} = z_{t+1} + V_{t+1}\left(V_{t+1}+\Sigma\right)^{-1}\left(y_{t+1}-z_{t+1}\right)$$

$$C_{t+1} = V_{t+1} - V_{t+1}\left(V_{t+1}+\Sigma\right)^{-1}V_{t+1}\,.$$

# G   Additional graphics and discussion from the simulation study

Figure G.1: Marginal posterior densities of θ_i $(i = 1, \ldots, 5)$ based on the (thinned) output of da-PsMMH.

With regard to Table G.3, we impose the restriction that $k \leq b$, since choosing $k > b$ automatically implies that the k nearest neighbours to a particular parameter value will, at some point, be split over more than one branch node. We found that using $k = 2$ reduces the computational cost of searching the tree but also reduces the accuracy of the KD-tree approximation, resulting in an overall decrease in mESS. Similarly, for $k > 5$ the increased accuracy is offset by increased computational cost. Using $k = 5$ and $2b = 20$ gave a 7-fold improvement in overall efficiency over PsMMH.

Figure G.2: Log-posterior estimates under the MJP ($\log(\hat{\pi}_s(\theta))$) against the corresponding log-posterior estimate given by the KD-tree ($\log(\hat{\pi}_c(\theta))$) based on the training data only (left panel) and the final adapted tree obtained after running da-PsMMH for 10^5 seconds with $\xi = 3$ and $k = b = 10$ (right panel). Both plots are obtained using $5,000$ values of θ sampled from the posterior $\pi(\theta)$.

		2b		
k	4	10	20	30
2	2807 (5.3)	2715 (5.1)	3435 (6.5)	3423 (6.5)
5		3017 (5.7)	3871 (7.3)	3848 (7.3)
10			3591 (6.8)	3377 (6.4)
15				3330 (6.3)

Table G.3: Minimum effective sample size (mESS) and relative mESS in parentheses for the Lotka Volterra model.

Figure G.3: Marginal posterior densities of θ_i $(i = 1, \ldots, 10)$ based on the (thinned) output of da-MMH using dataset $\mathcal{D}_1$ (solid) and $\mathcal{D}_2$ (dashed).