

Nested Adaptation of MCMC Algorithms

Dao Nguyen^{*}, Perry de Valpine[†], Yves Atchade[‡], Daniel Turek[§],
Nicholas Michaud^{¶,||}, and Christopher Paciorek^{**}

Abstract. Markov chain Monte Carlo (MCMC) methods are ubiquitous tools for simulation-based inference in many fields but designing and identifying good MCMC samplers is still an open question. This paper introduces a novel MCMC algorithm, namely, Nested Adaptation MCMC. For sampling variables or blocks of variables, we use two levels of adaptation where the inner adaptation optimizes the MCMC performance within each sampler, while the outer adaptation explores the space of valid kernels to find the optimal samplers. We provide a theoretical foundation for our approach. To show the generality and usefulness of the approach, we describe a framework using only standard MCMC samplers as candidate samplers and some adaptation schemes for both inner and outer iterations. In several benchmark problems, we show that our proposed approach substantially outperforms other approaches, including an automatic blocking algorithm, in terms of MCMC efficiency and computational time.

Keywords: adaptive MCMC, MCMC efficiency, integrated autocorrelation time, mixing.

1 Introduction

Markov chain Monte Carlo (MCMC) has become a widely used approach for simulation from an arbitrary distribution of interest, typically a Bayesian posterior distribution, known as the target distribution. MCMC really represents a family of sampling methods. Generally speaking, any new sampler that can be shown to preserve the ergodicity of the Markov chain such that it converges to the target distribution is a member of the family and can be combined with other samplers as part of a valid MCMC kernel. The key to MCMC's success is its simplicity and applicability. In practice, however, it sometimes needs a lot of non-trivial tuning to work well (Haario et al., 2005).

To deal with this problem, many adaptive MCMC algorithms have been proposed (Gilks et al., 1998; Haario et al., 2001; Andrieu and Robert, 2001; Sahu et al., 2003). These allow parameters of the MCMC kernel to be automatically tuned based on previous samples. This breaks the Markovian property of the chain so has required special schemes and proofs that the resulting chain will converge to the target distribution (Andrieu and Moulines, 2003; Atchadé et al., 2005; Andrieu and Atchadé, 2006). Un-

^{*}Departments of Mathematics, University of Mississippi, Oxford, dxnguyen@olemiss.edu

[†]Department of Environmental Science, Policy, and Management, University of California, Berkeley

[‡]Department of Statistics, University of Michigan, Ann Arbor

[§]Department of Mathematics & Statistics, Williams College, Williamstown

[¶]Department of Environmental Science, Policy, and Management, University of California, Berkeley

^{||}Department of Statistics, University of California, Berkeley

^{**}Department of Statistics, University of California, Berkeley

der some weaker and easily verifiable conditions, namely “diminishing adaptation” and “containment”, Roberts and Rosenthal (2007) proved ergodicity of adaptive MCMC and proposed many useful samplers.

It is important to realize, however, that such adaptive MCMC samplers address only a small aspect of a much larger problem. A typical adaptive MCMC sampler will approximately optimize performance given the kind of sampler chosen in the first place, but it will not optimize among the variety of samplers that could have been chosen. For example, an adaptive random walk Metropolis-Hastings sampler will adapt the scale of its proposal distribution, but that adaptation won’t reveal whether an altogether different kind of sampler would be more efficient. In many cases it would, and the exploration of different sampling strategies often remains a human-driven trial-and-error affair.

Here we present a method for a higher level of MCMC adaptation. The adaptation explores a potentially large space of valid MCMC kernels composed of different samplers. One starts with an arbitrary set of candidate samplers for each dimension or block of dimensions in the target distribution. The main idea is to iteratively try different candidates that compose a valid MCMC kernel, run them for a relatively short time, generate the next set of candidates based on the results thus far, and so on. Since relative performance of different samplers is specific to each model and even to each computing environment, it is doubtful whether there is a universally optimal kind of sampler. Hence we view the choice of efficient samplers for a particular problem as well-suited to empirical determination via computation.

The goal of computationally exploring valid sampler combinations in search of an efficient model-specific MCMC kernel raises a number of challenges. First, one must prove that the samples collected as the algorithm proceeds indeed converge to the target distribution, even when some of the candidate samplers are internally adaptive, such as conventional adaptive Metropolis-Hastings samplers. We provide such a proof for a general framework.

Second, one must determine efficient methods for exploring the very large, discrete space of valid sampler combinations. This is complicated by a combinatorial explosion, which is exacerbated by the fact that any multivariate samplers can potentially be used for arbitrary blocks of model dimensions. Here we take a practical approach to this problem, setting as our goal only to show basic schemes that can yield substantial improvements in useful time frames. Future work can aim to develop improvements within the general framework presented here. We also limit ourselves to relatively simple candidate samplers, but the framework can accommodate many more choices.

Third, one must determine how to measure the efficiency of a particular MCMC kernel for each dimension and for the entire model, in order to have a metric to seek to optimize. As a first step, it is vital to realize that there can be a trade-off between good mixing and computational speed. When considering adaptation within one kind of sampler, say adaptive Metropolis-Hastings, one can roughly assume that computational cost does not depend on the proposal scale, and hence mixing measured by integrated autocorrelation time, or the related effective sample size, is a sensible measure of efficiency. But when comparing two samplers with very different computational costs, say

adaptive Metropolis-Hastings and slice samplers, good mixing may or may not be worth its computational cost. Metropolis-Hastings samplers may mix more slowly than slice samplers on a per iteration basis, but they do so at higher computational speed because slice samplers can require many evaluations of model density functions. Thus the greater number of random walk iterations per unit time could outperform the slice sampler. An additional issue is that different dimensions of the model may mix at different rates, and often the slowest-mixing dimensions limit the validity of all results (Turek et al., 2017). In view of these considerations, we define MCMC efficiency as the effective sample size per computation time and use that as a metric of performance per dimension. Performance of an MCMC kernel across all dimensions is defined as the minimum efficiency among all dimensions.

The rest of the paper is organized as follows. Section 2 begins with a general theoretical framework for Nested Adaptation MCMC, then extends these methods to a specific Nested Adaptation algorithm involving block MCMC updating. Section 3 presents an example algorithm that fits within the framework, and provides some explanations on its details. Section 4 then outlines some numerical examples comparing the example algorithm with existing algorithms for a variety of benchmark models. Finally, Section 5 concludes and discusses some future research directions.

2 A general Nested Adaptation MCMC

In this section, we present a general Nested Adaptation MCMC algorithm and give theoretical results establishing its correctness.

Let \mathcal{X} be a state space and π the probability distribution on \mathcal{X} that we wish to sample from. Let \mathcal{I} be a countable set (this set indexes the discrete set of MCMC kernels we wish to choose from). For $\iota \in \mathcal{I}$, let Θ_ι be a parameter space in some subset space \mathbb{R}^m . For $\iota \in \mathcal{I}$ and $\theta \in \Theta_\iota$, let $P_{\iota, \theta}$ denote a Markov kernel on \mathcal{X} with invariant distribution π . We set $\bar{\Theta} = \bigcup_{\iota \in \mathcal{I}} \{\iota\} \times \Theta_\iota$ the adaptive MCMC parameter space. We want to build a stochastic process (an adaptive Markov chain) $\{(X_n, \iota_n, \theta_n), n \geq 0\}$ on $\mathcal{X} \times \bar{\Theta}$ such that as $n \rightarrow \infty$, the distribution of X_n converges to π , and a law of large numbers holds. We call ι the external adaptation parameter and θ the internal adaptation parameter.

We will follow the general adaptive MCMC recipe of Roberts and Rosenthal (2009). Assume that any internal adaptation on Θ_ι is done using a function $H_\iota: \Theta_\iota \times \mathcal{X} \rightarrow \Theta_\iota$, and an ‘‘internal clock’’ sequence $\{\gamma_n, n \geq 0\}$ such that $\lim_{n \rightarrow \infty} \gamma_n = 0$. The function H_ι depends on γ_n and updates parameters for internal adaptation. Also, let $\{p_k, k \geq 1\}$ be a sequence of numbers $p_k \in (0, 1)$ such that $\lim_{k \rightarrow \infty} p_k = 0$. p_k will be the probability of performing external adaptation at external iteration k . During the algorithm we will also keep track of two variables: κ_n , the number of external adaptations performed up to step n ; and τ_n , the number of iterations between n and the last time an external adaptation is performed. These two variables are used to manage the internal clock based on external iterations, which in most situations can simply be the number of adaptation steps. We build the stochastic process $\{(X_n, \iota_n, \theta_n), n \geq 0\}$ on $\mathcal{X} \times \bar{\Theta}$ as follows.

1. We start with $\kappa_0 = \tau_0 = 0$. We start also with some $X_0 \in \mathcal{X}$, $\iota_0 \in \mathcal{I}$, and $\theta_0 \in \Theta_{\iota_0}$.
2. At the n -th iteration, given $\mathcal{F}_n \stackrel{\text{def}}{=} \sigma\{(X_k, \iota_k, \theta_k), k \leq n\}$, and given κ_n, τ_n :
 - (a) Draw $X_{n+1} \sim P_{\iota_n, \theta_n}(X_n, \cdot)$.
 - (b) Independently of \mathcal{F}_n and X_{n+1} , draw $B_{n+1} \sim \text{Bern}(p_{n+1}) \in \{0, 1\}$.
 - i. If $B_{n+1} = 0$, there is no external adaptation: $\iota_{n+1} = \iota_n$. We update κ_n and τ_n :

$$\kappa_{n+1} = \kappa_n, \tau_{n+1} = \tau_n + 1. \quad (2.1)$$

Then we perform an internal adaptation: set $c_{n+1} = \kappa_{n+1} + \tau_{n+1}$, and compute

$$\theta_{n+1} = \theta_n + \gamma_{c_{n+1}} H_{\iota_n}(\theta_n, X_{n+1}). \quad (2.2)$$

Note that the internal adaptation interval could vary between iterations.

- ii. If $B_{n+1} = 1$, then we do an external adaptation: we choose a new ι_{n+1} . And we choose a new value $\theta_{n+1} \in \Theta_{\iota_{n+1}}$ based on \mathcal{F}_n and X_{n+1} . Then we update κ_n and τ_n .

$$\kappa_{n+1} = \kappa_n + 1, \tau_{n+1} = 0. \quad (2.3)$$

For this Nested Adaptation MCMC algorithm to be valid we must show that it satisfies three assumptions:

1. For each $(\iota, \theta) \in \bar{\Theta}$, $P_{\iota, \theta}$ has invariant distribution π .
2. (diminishing adaptation):

$$\Delta_{n+1} \stackrel{\text{def}}{=} \sup_{x \in \mathcal{X}} \|P_{\iota_n, \theta_n}(x, \cdot) - P_{\iota_{n+1}, \theta_{n+1}}(x, \cdot)\|_{\text{TV}}$$

converges in probability to zero, as $n \rightarrow \infty$.

3. (containment): For all $\epsilon > 0$, the sequence $\{M_\epsilon(\iota_n, \theta_n, X_n)\}$ is bounded in probability, where

$$M_\epsilon(\iota, \theta, x) \stackrel{\text{def}}{=} \inf\{n \geq 1 : \|P_{\iota, \theta}^n(x, \cdot) - \pi\|_{\text{TV}} \leq \epsilon\}.$$

Remark 2.1. Here the first assumption holds by construction. We will show that by the design, our Nested Adaptation algorithm satisfies the diminishing adaptation.

For $\iota \in \mathcal{I}$, $\theta, \theta' \in \Theta_\iota$, define

$$D_\iota(\theta, \theta') \stackrel{\text{def}}{=} \sup_{x \in \mathcal{X}} \|P_{\iota, \theta}(x, \cdot) - P_{\iota, \theta'}(x, \cdot)\|_{\text{TV}}.$$

Proposition 2.1. Suppose that \mathcal{I} is finite, and for any $\iota \in \mathcal{I}$, the adaptation function H_ι is bounded, and there exists $C < \infty$ such that

$$D_\iota(\theta, \theta') \leq C\|\theta - \theta'\|.$$

Then the diminishing adaptation holds.

Proof. We have

$$\begin{aligned} \mathbb{E}(\Delta_{n+1}) &= p_{n+1}\mathbb{E}(\Delta_{n+1}|B_{n+1} = 1) + (1 - p_{n+1})\mathbb{E}(\Delta_{n+1}|B_{n+1} = 0), \\ &\leq 2p_{n+1} + \mathbb{E}(\Delta_{n+1}|B_{n+1} = 0), \\ &= 2p_{n+1} + \mathbb{E}[D_{\iota_n}(\theta_n, \theta_{n+1})], \\ &\leq 2p_{n+1} + C\mathbb{E}[\|\theta_{n+1} - \theta_n\|], \\ &\leq 2p_{n+1} + C_1\gamma_{c_{n+1}}, \end{aligned}$$

where $c_{n+1} = \kappa_{n+1} + \tau_{n+1}$. It is easy to see that $c_n \rightarrow \infty$ as $n \rightarrow \infty$. The result follows since $\lim_{n \rightarrow \infty} p_n = \lim_{n \rightarrow \infty} \gamma_n = 0$. \square

2.1 On the containment assumption

In general, the containment condition is more challenging to check than the diminishing adaptation condition, and this technical assumption might not even be necessary sometimes (Roberts and Rosenthal, 2007). Since the containment assumption is an assumption on the mixing time process $\{M_\epsilon(\iota_n, \theta_n, X_n), n \geq 0\}$, adaptive MCMC algorithms for which the Markov kernels $P_{\iota, \theta}$ have similar sufficiently fast mixing behavior typically satisfy the containment assumption. This intuition was rigorously established by Roberts and Rosenthal (2007) which showed that the containment assumption holds when the family $\{P_{\iota, \theta} : (\iota, \theta) \in \Theta\}$ possesses a certain simultaneous ergodicity property (either uniform, geometric, or sub-geometric). This means in a nutshell that all the Markov kernels $P_{\iota, \theta}$ have qualitatively the same convergence rate, be it uniform, geometric or subgeometric. Hence to check the containment assumption, we typically need to study the convergence rate of each kernel $P_{\iota, \theta}$, and verify that all the rates are within a multiplicative constant of each other. Despite recent advances in Markov Chain Monte Carlo theory, there is no general theory that can be easily employed to establish the convergence rate of a given MCMC algorithm. The rate depends as much on the algorithm, as on the target distribution. In this context, it is clear that verifying the containment assumption is a mathematically tedious endeavor that is beyond the scope of this paper. We refer the interested reader to Roberts and Rosenthal (2007); Andrieu and Thoms (2008); Atchade and Fort (2010); Atchadé et al. (2011) for some detailed examples.

3 Example algorithms

We present one specific approach as an example of a Nested Adaptation algorithm. Our approach to “outer adaptation” will be to identify the “worst-mixing dimension” (i.e., some parameter or latent state of the statistical model) and update the kernel by assigning different sampler(s) for that dimension. To explain the method, we will give some terminology for describing our algorithm. In particular, we will define a valid kernel, MCMC efficiency, and worst-mixing dimension. We will define a set of candidate samplers for a given dimension, which could include scalar samplers or block samplers. In either case, a sampler may also have internal adaptation for each parameter or combination. To implement the internal clock of each sampler (c_n of the general algorithm),

we need to formulate all internal adaptation in the framework using 2.2. We use P (without subscripts) in this section to represent $P_{\iota, \theta}$ of the general theory, so the kernel and parameters are implicit.

3.1 Valid kernel

Assume our model of interest is \mathcal{M} , which could be represented as a graphical model where vertices or nodes represent states or data while edges represent dependencies among them. Here we are using “state” as Bayesian statisticians do to mean any dimension of the model to be sampled by MCMC, including model parameters and latent states. We denote the set of all dimensions of the target distribution, as $\mathcal{X} = \{\mathcal{X}_1, \dots, \mathcal{X}_m\}$. Since we will construct a new MCMC kernel as an ordered set of samplers at each outer iteration, it is useful to define requirements for a kernel to be valid. We require that each kernel, if used on its own, would be a valid MCMC to sample from the target distribution $\pi(X)$, $X \in \mathcal{X}$ (typically defined from Bayes Rule as the conditional distribution of states given the data). This is the case if it satisfies the detailed balance equation, $\pi = P\pi$.

In more detail, we need to ensure that a new MCMC kernel does not omit some subspace of \mathcal{X} from mixing. Denote the kernel P as a sequence of (new choice of c_{n+1}) samplers P_i , $i = 1 \dots j$, such that $P = P_j P_{j-1} \dots P_1$. By some abuse of terminology, P is a valid kernel if each sampler P_i operates on a non-empty subset b_i of \mathcal{X} , satisfying $\bigcup_{i=1}^j b_i = \mathcal{X}$.

At iteration n , assume the kernel is $P^{(n)}$ and the samples are $X_n = (X_{n,1}, \dots, X_{n,m})$ where the set of initial values is X_0 . For each dimension \mathcal{X}_k , $k = 1, \dots, m$ let $\mathbf{X}_k = \{X_{0,k}, X_{1,k}, \dots\}$ be the scalar chain of samples of \mathcal{X}_k .

3.2 Worst mixing state and MCMC efficiency

We define MCMC efficiency for state \mathcal{X}_k from a sample of size N from kernel P as effective sample size per computation time

$$\omega_k(N, P) = \frac{N/\tau_k(P)}{t(N, P)},$$

where $t(N, P)$ is the computation time for kernel P to run N iterations (often $t(N, P) \approx Nt(1, P)$) and $\tau_k(P)$ is the integrated autocorrelation time for chain \mathbf{X}_k defined as

$$\tau_k = 1 + 2 \sum_{i=1}^{\infty} \text{cor}(X_{0,k}, X_{i,k})$$

(Straatsma et al., 1986). The ratio N/τ_k is the effective sample size (ESS) for state \mathcal{X}_k (Roberts and Rosenthal, 2001). Note that $t(N, P)$ is computation time for the entire kernel, not just samplers that update \mathcal{X}_k . τ_k can be interpreted as the number of effective samples per actual sample. The worst-mixing state is defined as the state with minimum MCMC efficiency among all states. Let k_{min} be the index of the worst-mixing state,

that is

$$k_{min} = \arg \min_k \tau_k^{-1}.$$

Since the worst mixing dimension will limit the validity of the entire posterior sample (Thompson, 2010), we define the efficiency of a MCMC algorithm as $\omega_{k_{min}}(N, P)$, the efficiency of the worst-mixing state of model \mathcal{M} .

There are several ways to estimate ESS, but we use `effectiveSize` function in the R coda package (Plummer et al., 2006) since this function provides a stable estimation of ESS. This method, which is based on the spectral density at frequency zero, has been proven to have the highest convergence rate, thus giving a more accurate and stable result (Thompson, 2010).

3.3 Candidate samplers

A set of candidate samplers $\{P_j, j \in \mathcal{S}\}$ is a list of all possible samplers that could be used for a parameter of the model \mathcal{M} . These may differ depending on the parameter's characteristics and role in the model (e.g., whether there is a valid Gibbs sampler, or whether it is restricted to $[0, \infty)$). In addition to univariate candidate samplers, nodes can also be sampled by block samplers. Denote $|b|$ the number of elements of b . If $|b_i| > 1$, $P_i^{(n)}$, the sampler applied to block b_i at iteration n , is called a block sampler; otherwise it is a univariate or scalar sampler.

In the examples below we considered up to four univariate candidate samplers and three kinds of block samplers. The univariate samplers included adaptive Metropolis-Hastings (AMH), adaptive Metropolis-Hastings on a log scale (AMHLS) for states taking only positive real values, Gibbs samplers for states with a conjugate prior-posterior pairing, and slice samplers. The block samplers included adaptive Metropolis-Hastings with multivariate normal proposals, automated factor slice sampler (Tibbits et al., 2014) (slice samplers in a set of orthogonal rotated coordinates), and automated factor random walk (univariate random walks in a set of orthogonal rotated coordinates). These choices are by no means exhaustive but serve to illustrate the algorithms here.

Block samplers and how to block

Turek et al. (2017) suggested different ways to block the states efficiently: (a) based on correlation clustering, (b) based on model structure. Here we use the first method.

At each iteration, we use the generated samples to create the empirical posterior correlation matrix. To stabilize the estimation, all of the samples are used to compute a correlation matrix $\rho_{d \times d}$. This in turn is used to make a distance matrix $D_{d \times d}$ where $D_{i,j} = 1 - |\rho_{i,j}|$ for $i \neq j$ and $D_{i,i} = 0$ for every i, j in $1, \dots, d$. To ensure minimum absolute pairwise correlation between clusters, we construct a hierarchical cluster tree from the distance matrix D (Everitt et al. (2011) chapter 4). Given a selected height, we cluster the hierarchical tree into distinct groups of states. Different parts of the tree may have different optimal heights for forming blocks. Instead of using a global height to cut the tree, we only choose a block that contains the worst-mixing state from the

cut and keep the other samplers unchanged. At each outer iteration, if the algorithm is choosing a new block sampler as part of a new kernel, it will cut the tree at a lower correlation threshold, thus creating a cluster of nodes with lower posterior correlations. In our implementation, we use the R function `hclust` to build the hierarchical clustering tree with “complete linkage” from the distance matrix D . By construction, the absolute correlation between states within each group is at least $1 - h$ for h in $[0, 1]$. We then use the R function `cutree` to choose a block that contains the worst-mixing state. This process is justified in the sense that the partitioning adapts according to the model structure through the posterior correlation. The details and validity of the block sampling in our general framework are provided in Supplement S1 (Nguyen et al., 2019).

3.4 How to choose new samplers

To choose new samplers to compose a new kernel, we determine the worst-mixing state and choose randomly from candidate samplers to replace whatever sampler was updating it in the previous kernel while keeping other samplers the same. There are some choices to make when considering a block sampler. If the worst-mixing parameter is x , and the new kernel will use a block sampler for x together with one or more parameters y , we keep the current sampler(s) used for y . Future work can consider other schemes such as changing group of samplers together based on model structure.

3.5 Internal clock variables

In the Algorithm 1, θ represents the internal adaptation parameter of a particular sampler and c represents its internal clock. In general, an internal clock variable is defined as a variable used in a sampler to determine the size of internal adaptation steps such that any internal adaptation would converge in a typical MCMC setting. An example of an internal clock variable is a number of internal iterations that have occurred. To use a sampler in the general framework, we need to establish what are its internal adaptation and clock variables. A few examples of internal adaptation variables of different samplers are summarized as follows:

- For adaptive Metropolis-Hastings: proposal scale is used.
- For block adaptive Metropolis-Hastings: proposal scale and covariance matrix are used.
- For automated factor slice sampler: covariance matrix (or equivalent, i.e. coordinate rotation) is used.
- For automated factor random walk: covariance matrix (ditto) and proposal scales for each rotated coordinate axis are used.

These internal adaption variables are set to default initial values when its sampler is first used. After that, they are retained along with internal clock variables so that whenever we revisit a sampler, we will use the stored values to set up this sampler.

This setting guarantees the diminishing adaption property, which is essential for the convergence of the algorithm. Pseudo-code for Nested Adaptation MCMC is given in Algorithm 1.

Algorithm 1 Nested Adaptation MCMC

Input:

Bayesian model with initial state (including latent variables) X_0
 $\{p_n, n \in \mathbb{N} | p_n \in (0, 1), \lim_n p_n = 0\}$, maximum iteration M
 Candidate samplers $\{P_j, j \in \mathcal{S}\}$
 $P_{i_0, \theta_0} :=$ ordered set of initial samplers $\{P_j^{(0)}\}_{j \in \mathcal{S}}$ from Bayesian model

Output:

An ordered set of samplers $\{P_{i^*}\}_{i^* \in \mathcal{S}}$ with the best MCMC efficiency so far

- 1: Initialize $\text{EFF}, \text{EFF}_{\text{best}}, n, \kappa_0, \tau_0, c_0$ to 0 ▷ Denote MCMC efficiency EFF
- 2: **while** ($\text{EFF} \geq \text{EFF}_{\text{best}}$) or ($n < M$) **do**
- 3: Sample N samples from the current sampler set $\{P_j^{(n)}\}_{j \in \mathcal{S}}$
- 4: Store internal clocks c_n and adaption variables θ_n for each sampler ▷ Section 3.5
- 5: Compute $\text{EFF}_k = \omega_k(N, P) = \frac{N/\tau_k(P)}{t(N, P)}$ ▷ k is an index of parameters
- 6: Identify $k_{\min} = \arg \min_k \tau_k^{-1}$, $\text{EFF} = \text{EFF}_{k_{\min}}$ ▷ See Section 3.2
- 7: **if** ($\text{EFF} \geq \text{EFF}_{\text{best}}$) **then**
- 8: Set $\{P_{i^*}\}_{i^* \in \mathcal{S}} = \{P_i^{(n)}\}_{i \in \mathcal{S}}$
- 9: Set $\text{EFF}_{\text{best}} = \text{EFF}$
- 10: **else**
- 11: Set $\{P_i^{(n)}\}_{i \in \mathcal{S}} = \{P_{i^*}\}_{i^* \in \mathcal{S}}$
- 12: Draw $B_{n+1} \sim \text{Bern}(p_{n+1}) \in \{0, 1\}$
- 13: **if** $B_{n+1} = 0$ **then**
- 14: $\kappa_{n+1} = \kappa_n, \tau_{n+1} = \tau_n + 1, c_{n+1} = \kappa_{n+1} + \tau_{n+1}$
- 15: **else**
- 16: $\kappa_{n+1} = \kappa_n + 1, \tau_{n+1} = 0, c_{n+1} = \kappa_{n+1} + \tau_{n+1}$
- 17: Set $P_i^{(n+1)} = P_i^{(n)}, i \neq k_{\min}$, choose $P_{k_{\min}}^{(n+1)}$ from candidate samplers ▷ Section 3.4
- 18: **if** ($P_{k_{\min}}^{(n+1)}$ has been used before) **then**
- 19: Use c_n, θ_n to set up the sampler $P_{k_{\min}}^{(n+1)}$
- 20: **else**
- 21: Use default internal adaptation value of $P_{k_{\min}}^{(n+1)}$ ▷ Section 3.5
- 22: Set $n = n + 1$

4 Examples

In this section, we evaluate our algorithm on some benchmark examples and compare them to different MCMC algorithms. In particular, we compare our approach to the following MCMC algorithms.

- All Scalar algorithm: Every dimension is sampled using an adaptive scalar normal random walk sampler.
- All Blocked algorithm: All dimensions are sampled in one adaptive multivariate normal random walk sampler.

- Default algorithm: Samplers are assigned as follows. When possible, a conjugate (Gibbs) sampler is used. Otherwise, adaptive random-walk Metropolis-Hastings samplers are used. These will be block samplers for parameters arising from multivariate distributions and scalar samplers otherwise.
- Auto Block algorithm: The Auto Block method (Turek et al., 2017) searches blocking schemes based on hierarchical clustering from posterior correlations to determine a highly efficient (but not necessarily optimal) set of blocks that are sampled with multivariate normal random-walk samplers. Thus, Auto Block uses only either scalar or multivariate adaptive random walk, concentrating more on partitioning the correlation matrix than trying different sampling methods. Note that the initial sampler of both the Auto Block algorithm and our proposed algorithm is the All Scalar algorithm.

All experiments were carried out using the NIMBLE package (de Valpine et al., 2017) for R (R Core Team, 2013) on a cluster using 32 cores of Intel Xeon E5-2680 2.7 Ghz with 256 GB memory. Models are coded using NIMBLE’s version of the BUGS model declaration language (Lunn et al., 2000, 2012). All MCMC algorithms are written in NIMBLE, which provides user-friendly interfaces in R and efficient execution in custom-generated C++, including matrix operations in the C++ Eigen library (Guennebaud et al., 2010).

To measure the performance of an MCMC algorithm, we use MCMC efficiency. MCMC efficiency depends on ESS, estimates of which can have a high variance for a short Markov chain. This presents a tuning-parameter trade-off for the Nested Adaptation method: Is it better to move cautiously (in sampler space) by running long chains for each outer adaptation in order to gain an accurate measure of efficiency, or is it better to move adventurously by running short chains, knowing that some algorithm decisions about samplers will be based on noisy efficiency comparisons? In the latter case, the final samplers may be less optimal, but that may be compensated by the saved computation time. To explore this trade-off, we try our Nested Adaptation algorithm with different sample sizes in each outer adaptation and label results accordingly. For example, Nested Adaptation 10K will refer to the Nested Adaptation method with samples of 10,000 per outer iteration.

We present algorithm comparisons in terms of time spent in an adaptation phase, final MCMC efficiency achieved, and the time required to obtain a fixed effective sample size (e.g., 10,000). Only Auto Block and Nested Adaptation have adaptation phases. An important difference is that Auto Block did not come with a proof of valid adaptive MCMC convergence (it could be modified to work in the current framework, but we compare to the published version). Therefore, samples from its adaptation phase are not normally included in the final samples, while the adaptation samples of Nested Adaptation can be included.

To measure final MCMC efficiency, we conducted a single long run of length N with the final kernel of each method solely for the purpose of obtaining an accurate ESS estimate. One would not normally do such a run in a real application. The calculation of time to obtain a fixed effective sample size incorporates both adaptation time and

efficiency of the final samplers. For both Nested Adaptation and Auto Block, we placed them on a similar playing field by assuming for this calculation that samples are not retained from the adaptation phase, making the results conservative.

For all comparisons, we used 20 independent runs of each method and present the average results from these runs. To show the variation in runs, we present box-plots of efficiency in relation to computation time from the 20 runs of Nested Adaptation. The final (right-most) box-plot in each such figures shows the 20 final efficiency estimates from larger runs. Not surprisingly, these can be lower than obtained by shorter runs. These final estimates are reported in the tables.

A public Github repository containing scripts for reproducing our results may be found at <https://github.com/nxdao2000/AutoAdaptMCMC>. Some additional experiments are also provided there.

4.1 Linear state space model

First, to illustrate the distinction between the Nested Adaptation algorithm and other methods, we will use a linear state space model (Durbin and Koopman, 2012) in which all parameters are fixed. Such a model will allow us to provide a simple assessment of the performances of the methods for a range of situations, including those in which the models were true.

Let x_t be the latent state at time t , y_t be the observed data, and suppose we have the number of time points $T = 50$. Let the initial priors be:

$$\begin{aligned} a &\sim \text{Uniform}(0, 1), \\ b &\sim N(1, 1), \\ \mu_0 &\sim N(0, 1), \end{aligned}$$

and observation density $y_t \sim N(bx_t, 1)$ for $t = 2, \dots, 50$. State transitions are governed by a first order autoregressive (AR) process

$$x_t \sim N(ax_{t-1}, 1) \quad \text{for } t = 2, \dots, 50.$$

x_1 and y_1 were randomly generated from simulating $x_1 \sim N(\mu_0, 1)$, $y_1 \sim N(x_1, 1)$ with $\mu_0 = 0$ while x_t and y_t are subsequently generated using transition density with the parameter values $a = 0.2$ and $b = 2$. These values were arbitrarily chosen to cover the range observed for the explanatory variables; we desire high efficiency regardless of model fit, so the particular choice of this simple model is tangential to our main points.

In this relatively simple example, we try our Nested Adaptation algorithm with large numbers of iterations in each outer adaptation as well as a large number of iterations for the final efficiency. Specifically, we used samples sizes of 5000, 10000 and 20000 per outer adaptation and we used $N = 100000$ for computing final efficiency.

For this example (Table 1) All Scalar sampling produces MCMC efficiency of about 121.8, while the All Blocked algorithm, which consists of a single block sampler of dimension 53, has MCMC efficiency of approximately 5.2. All Blocked samples all 53

Algorithms	Adapt time	Efficiency	Time to 10^4 effective samples
All Blocked	0.00	52.359	1910
Default	0.00	1198.128	83
All Scalar	0.00	1218.060	82
Auto Block	22.92	1277.639	101
Nested Adaptation 5K	0.9	1796.079	56
Nested Adaptation 10K	1.7	1693.929	61
Nested Adaptation 20K	3.3	1616.752	65

Table 1: Summary results of different MCMC algorithms for the linear state space model. Runtime is presented as seconds, and efficiency is in units of effective samples produced per second of algorithm runtime. Time to N effective samples is computed by $N/\text{efficiency}$ for static algorithms and that plus adaptation time for Auto Block and Nested Adaptation algorithms.

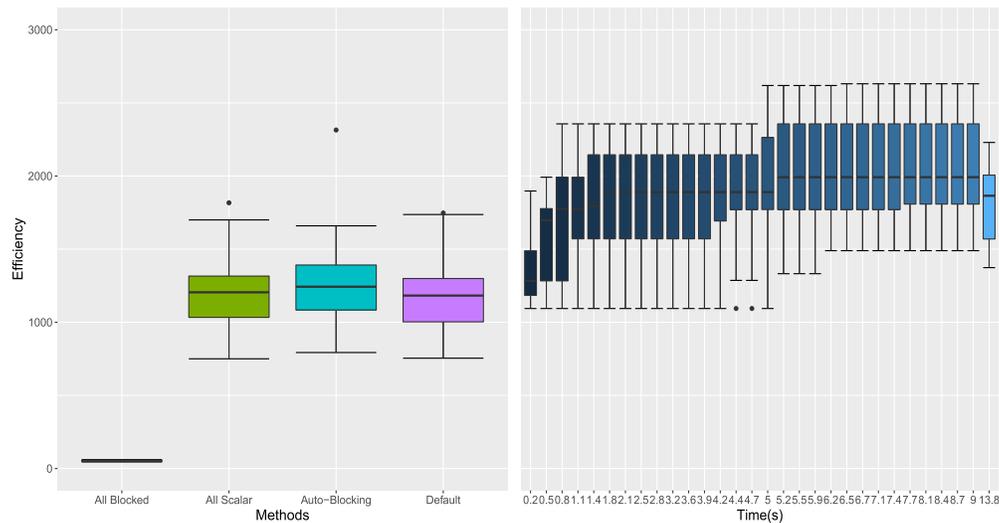


Figure 1: MCMC efficiencies of different methods for linear state space model. The left panel shows the box-plots of MCMC efficiencies of All Blocked, All Scalar, Auto Block and Default algorithms computed from 20 replications. The right panel shows the box-plots of MCMC efficiencies of Nested Adaptation 5K algorithm computed from 20 replications at each outer adaptation. The last (right-most) box-plot is computed from the chain of 100000 samples generated. The time axis shows the average computational time of 20 replications.

dimensions jointly, which requires computation time roughly double that of All Scalar and yields only rather low ESS. The Default algorithm performs similarly to Auto Block but worse than Nested Adaptation algorithm. Since Auto Block converges to All Scalar in this case, the Auto Block algorithm performs no better than standard All Scalar (Figure 1) as would be expected. It is clear that Nested Adaptation method has

dramatic improvements even when we take into account the adaptation time. Amongst these Auto methods, Auto Block performs worse than all Nested Adaptation 5K, Nested Adaptation 10K and Nested Adaptation 20K in both computational time and MCMC efficiencies. Overall, Nested Adaptation 5K appears to be the most efficient method in terms of time to 10000 effective samples. One interpretation is that Nested Adaptation 5K trades off well between adaptation time and MCMC efficiency in this model. The final samplers from Nested Adaptation included an automated factor slice sampler on block (a, b), an automated factor random walk sampler on block (μ_0, x_1) and adaptive random-walk Metropolis-Hastings samplers on other nodes.

Figure 1 illustrates that Nested Adaptation algorithm outperforms the other algorithms. This comes from both the flexibility to trade-off the number of outer adaptations vs. adaptive time to reach a good sampler as well as the larger space of kernels being explored. Since MCMC efficiency is highly dependent upon hierarchical model structures, using scalar and multivariate normal random walks alone, as done by the Auto Block algorithm, can be quite limiting. Nested Adaptation, can overcome this limitation with the flexibility to choose different type of samplers. We will see that more strongly in the next examples, where the models are more complex.

4.2 A random effect model

We consider the “litters” model, which is an original example model provided with the MCMC package WinBUGS. This model is chosen because of its notoriously slow mixing, which is due to the strong correlation between parameter pairs. It is desirable to show how much improvement can be achieved compared to other approaches on this benchmark example. The purpose of using this simple example is to establish the potential utility of the Nested Adaptation approach, while saving more advanced applications for future work. In this case, we show that our algorithm indeed outperforms by a significant margin the other approaches. This model’s specification is given following Deely and Lindley (1981) and Kass and Steffey (1989) as follows.

Suppose we observe the data in i groups. In each group, the data y_{ij} , $j = \{1, \dots, n\}$ are conditionally independent given the parameters p_{ij} , with the observation density

$$y_{ij} \sim \text{Bin}(n_{ij}, p_{ij}).$$

In addition, assume that p_{ij} for fixed i are conditionally independent given the “hyper-parameters” α_i, β_i , with conjugate density

$$p_{ij} \sim \text{Beta}(\alpha_i, \beta_i).$$

Assume that α_j, β_j follow prior densities,

$$\alpha_1 \sim \text{Gamma}(1, 0.001),$$

$$\beta_1 \sim \text{Gamma}(1, 0.001),$$

$$\alpha_2 \sim \text{Uniform}(0, 100),$$

$$\beta_2 \sim \text{Uniform}(0, 50).$$

Algorithms	Adapt time	Efficiency	Time to 10^4 effective samples
All Blocked	0.00	0.5855	17079
Default	0.00	1.8385	5439
All Scalar	0.00	1.6870	5928
Auto Block	21.97	12.1205	847
Nested Adaptation 50K	4.09	16.4532	612
Nested Adaptation 70K	5.95	18.0358	560
Nested Adaptation 100K	8.34	17.6278	576

Table 2: Summary results of different MCMC algorithms for the litters model. Runtime is presented as seconds, and efficiency is in units of effective samples produced per second of algorithm runtime. Time to N effective samples is computed by $N/\text{efficiency}$ for static algorithms and that plus adaptation time for Auto Block and Nested Adaptation algorithms.

Following the setup of Rue and Held (2005) as discussed in Turek et al. (2017), we could jointly sample the top-level parameters and conjugate latent states as the beta-binomial conjugacy relationships allow the use of what Turek et al. (2017) call cross-level sampling, but, for demonstration purposes, we do not include this here.

Since the litters model mixes poorly, we run a large number of iterations (i.e. $N = 300000$) to produce stable estimates of final MCMC efficiency. We start both Auto Block and Auto Adapt algorithms with All Scalar and adaptively explore the space of all given candidate samplers. We use Nested Adaptation with either 50000, 70000 or 100000 iterations per outer adaptation.

Results (Table 2) show that Auto Block generates samples with MCMC efficiency about seven-fold, six-fold and twenty-fold that of the All Scalar, Default and All Blocked methods, respectively. We can also see that as the outer adaptation sample size increases, the performance of Nested Adaptation improves. Final MCMC efficiencies of Nested Adaptation 50K, Nested Adaptation 70K and Nested Adaptation 100K are 135%, 148% and 145% of MCMC efficiency of Auto Block, respectively. In addition, the adaptation time for all cases of Nested Adaptation are much shorter than for Auto Block. Combining adaptation time and final efficiency into the resulting time to 10000 effective samples, we see that in this case, larger samples in each outer iteration are worth their computational cost. In this example, the final samplers from Nested Adaptation included automated factor random walk samplers on blocks (α_1, β_1) and (α_2, β_2) , a slice sampler on $p_{1,16}$ and adaptive Metropolis-Hastings samplers on all other nodes.

Figure 2 shows the box-plots computed from 20 independent runs on the litters model of All Blocked, All Scalar, Auto Block, Default and Nested Adaptation 50K algorithms. The left panel of the figure confirms that MCMC efficiency of Auto Block is well dominated that of other static adaptive algorithms. The right panel of the figure shows the MCMC efficiency of Nested Adaptation 50K gradually improves with time. The right-most box-plot verifies that the MCMC efficiency of selected samplers from Nested Adaptation algorithm (computed from large samples) is better than that of

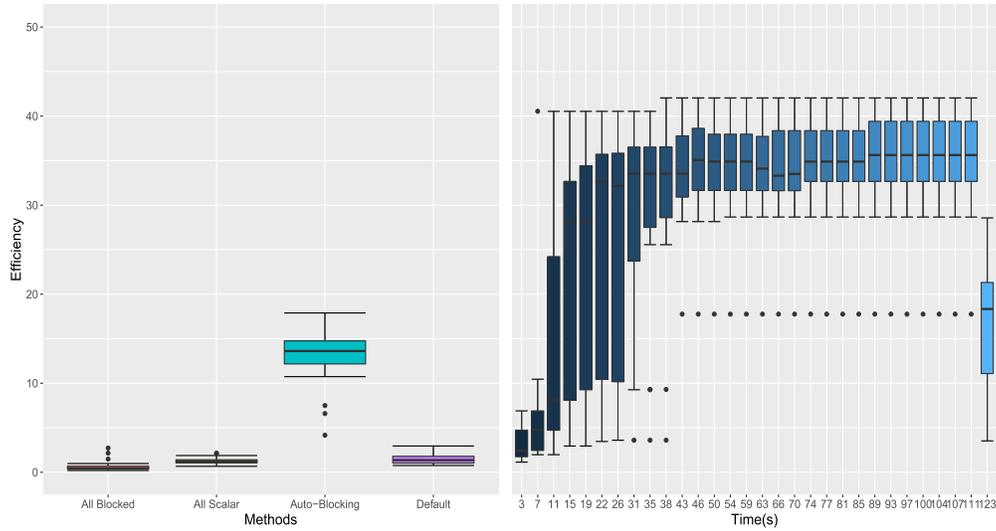


Figure 2: MCMC efficiencies of different methods for litters model. The left panel shows the box-plots of MCMC efficiencies of All Blocked, All Scalar, Auto Block and Default algorithms computed from 20 replications. The right panel shows the box-plots of MCMC efficiencies of Nested Adaptation 50K algorithm computed from 20 replications at each outer adaptation. The last (right-most) box-plot is computed from the chain of 300000 samples generated. The time axis shows the average computational time of 20 replications.

Auto Block algorithm. Last but not least, Nested Adaptation algorithms are much more efficient than Auto Block in the sense that we can keep every sample while Auto Block algorithm throws away most of the samples.

4.3 Spatial model

In this section, we consider a hierarchical spatial model as the final example. We use the classical scallops dataset for this model. This dataset is chosen since we want to compare our approach with other standard approaches in the presence of spatial dependence. This data collects observations of scallop abundance at 148 locations from the New York to New Jersey coastline in 1993. It was surveyed by the Northeast Fisheries Science Center of the National Marine Fisheries Service and made publicly available at <http://www.biostat.umn.edu/~brad/data/myscallops.txt>. It has been analyzed many times, such as Ecker and Heltshe (1994); Ecker and Gelfand (1997); Banerjee et al. (2004) and references therein. Following Banerjee et al., assume the log-abundance $\mathbf{g} = (g_1, \dots, g_N)$ follows a multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, defined by covariances that decay exponentially as a function of distance. Specifically, let y_i be measured scallop abundance at site i , $d_{i,j}$ be the distance between

Algorithms	Adapt time	Efficiency	Time to 10^4 effective samples
All Blocked	0.00	0.0100	1000000
Default	0.00	0.0020	5000000
All Scalar	0.00	0.1150	86956
Auto Block	19094.89	0.3565	47145
Nested Adaptation 5K	2967.55	0.4420	25592
Nested Adaptation 10K	6221.61	0.4565	28127
Nested Adaptation 20K	11278.78	0.4948	31488

Table 3: Summary results of different MCMC algorithms for spatial model. Runtime is presented as seconds, and efficiency is in units of effective samples produced per second of algorithm runtime. Time to N effective samples is computed by $N/\text{efficiency}$ for static algorithms and that plus adaptation time for Auto Block and Nested Adaptation algorithms.

sites i and j , and ρ be a valid correlation. Then

$$\mathbf{g} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma}),$$

where each component $\Sigma_{ij} = \sigma^2 \exp(-d_{i,j}/\rho)$.

We model observations as $y_i \sim \text{Poisson}(\exp(g_i))$. Priors for σ and ρ are Uniform over a large range of interest. The parameters in the posterior distribution are expected to be correlated as the covariance structure induces a trade-off between σ and ρ . This can be sampled well by Auto Block algorithm, and we would like to show that our approach can achieve even higher efficiency with lower computational cost of adaptation.

This spatial model, with 858 parameters, is computationally expensive to estimate. Therefore, we will use Nested Adaptation 5K, Nested Adaptation 10K and Nested Adaptation 20K algorithms for comparison and run $N = 50000$ for estimating final efficiency. In this case, the typical final samplers from Nested Adaptation included an automated factor random walk sampler on block (σ, ρ) , adaptive random walk samplers on blocks $(\mu, g_{34}, g_{66}, g_{148})$ and (g_{39}, g_{52}) , an automated factor slice sampler on block $(g_{11}, g_{12}, g_{33}, g_{48}, g_{92})$, slice samplers on g_{54}, g_{65} , and g_{68} , and adaptive random-walk Metropolis-Hastings samplers on the rest.

As can be seen from Table 3, All Blocked and Default algorithms mix very poorly, resulting in extremely low efficiencies of 0.01 and 0.002, respectively. The All Scalar algorithm, while achieving higher ESS, runs slowly because large matrix calculations are needed for every univariate sampler. The Auto Block algorithm, on the other hand, selects an optimal threshold to cut the entire hierarchical clustering tree into different groups, increasing the ESS about 3 times. With a few small blocks, the computation cost of Auto Block is somewhat cheaper than All Scalar algorithm. As a result, the efficiency mean is about 3.5 times that of All Scalar. Meanwhile, our Nested Adaptation 5K, 10K and 20K algorithms perform best. It should be noted that the Nested Adaptation algorithm can achieve good mixing with adaptation times that are only 15.5%, 32.5% and 59% compared to the adaptation time of Auto Block. In Figure 3, while the left panel shows a distinction between Auto Block and other

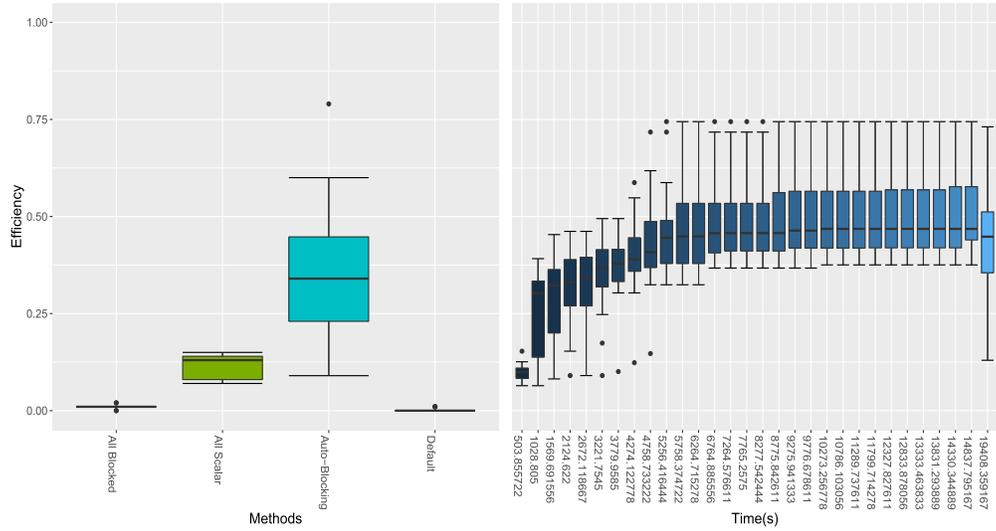


Figure 3: MCMC efficiencies of different methods for spatial model. The left panel shows the box-plots of MCMC efficiencies of All Blocked, All Scalar, Auto Block and Default algorithms computed from 20 replications. The right panel shows the box-plots of MCMC efficiencies of Nested Adaptation 5K algorithm computed from 20 replications at each outer adaptation. The last (right-most) box-plot is computed from the chain of 10000 samples generated. The time axis shows the average computational time of 20 replications.

static algorithms, the right panel shows that Nested Adaptation 20K surpasses Auto Block in just a few outer iterations, indicating substantial improvements in some models.

5 Discussion

We have proposed a general Nested Adaptation MCMC algorithm. Our algorithm traverses a space of valid MCMC kernels to find an efficient algorithm automatically. There is only one previous approach, namely Auto Block sampling, of this kind that we are aware of. We have shown that our approach can substantially outperform Auto Block in some cases, and that both outperform simple static approaches. Using some benchmark models, we observe that our approach can yield improvements, which can be substantial compared to the Default method.

The comparisons presented have deliberately used fairly simple samplers as options for Nested Adaptation in order to avoid comparisons among vastly different computational implementations. A major feature of our framework is that it can incorporate almost any sampler as a candidate and almost any strategy for choosing new kernels from compositions of samplers based on results so far. Samplers to be explored in the

future could include auxiliary variable algorithms such as slice sampling or derivative-based sampling algorithms such as Hamiltonian Monte Carlo (Duane et al., 1987). Now that the basic framework is established and shown to be useful in simple cases, it merits extension to more advanced cases.

The Nested Adaptation method can be viewed as a generalization of the Auto Block method. It is more general in the sense that it can use more kinds of samplers and explore the space of samplers more generally than the correlation-clustering of Auto Block. Thus, our framework can be considered to provide a broad class of automated kernel construction algorithms that use a wide range of sampling algorithm as components.

If block sampling is included in the space of the candidate samplers, choosing optimal blocks is important and can greatly increase the efficiency of the algorithm. For this reason, we extended the cutting of a hierarchical cluster tree to allow different cut heights on different branches (different parts of the model). This differs from Auto Block, which forms all blocks by cutting the entire tree at the same height. We also have different multivariate adaptive sampling other than random walk normal distribution such as automated factor slice sampler and automated factor random walk sampler. With these extensions, the final efficiency achieved by our algorithm specifically among blocking schemes is often substantially better and is found in a shorter time.

Beyond hierarchical clustering, there are other approaches one might consider to find efficient blocking schemes. One such approach would be to use the structure of the graph instead of posterior correlations to form blocks. This would allow conservation of calculations that are shared by some parts of the graph, whether or not they are correlated. Another future direction could be to improve how a new kernel is determined from previous results, essentially to determine an effective strategy for exploring the very high-dimensional kernel space. Finally, the trade-off between computational cost and the accuracy of effective sample size estimates is worth further exploration.

Development of Nested Adaptation MCMC has also raised several theoretical questions for future work. First, similarly to other adaptive MCMC methods, can more general conditions of validity be proven? Second, the estimation of effective sample size from a chain generated using Nested Adaptation is difficult because the properties of the kernel will have changed during the course of sampling. Methods for estimating effective sample size typically don't consider such a situation. Third, what are good choices for the external and internal adaptation schedules (e.g. γ_n and p_k in Section 2)? Fourth, how can we disentangle the contribution of each sampler to mixing achieved by a kernel comprising multiple samplers? Doing so could enable better moves in kernel space during outer adaptation. Finally, what are good strategies for parallelization of Nested Adaptation?

Supplementary Material

Nested adaptation of MCMC algorithms: Appendix A (DOI: [10.1214/19-BA1190SUPP](https://doi.org/10.1214/19-BA1190SUPP); .pdf).

References

- Andrieu, C. and Atchadé, Y. F. (2006). “On the efficiency of adaptive MCMC algorithms.” In *Proceedings of the 1st International Conference on Performance Evaluation Methodologies and Tools*, 43. ACM. MR2350572. doi: <https://doi.org/10.1214/ECP.v12-1320>. 1323
- Andrieu, C. and Moulines, E. (2003). “Ergodicity of some adaptive Markov chain Monte Carlo algorithm.” Technical report. 1323
- Andrieu, C. and Robert, C. P. (2001). “Controlled MCMC for optimal sampling.” *Preprint*. MR0784226. 1323
- Andrieu, C. and Thoms, J. (2008). “A tutorial on adaptive MCMC.” *Statistics and Computing*, 18(4): 343–373. MR2461882. doi: <https://doi.org/10.1007/s11222-008-9110-y>. 1327
- Atchade, Y. and Fort, G. (2010). “Limit theorems for some adaptive MCMC algorithms with sub-geometric kernels.” *Bernoulli*, 16(1): 116–154. MR2648752. doi: <https://doi.org/10.3150/09-BEJ199>. 1327
- Atchadé, Y., Fort, G., Moulines, E., and Priouret, P. (2011). “Adaptive Markov chain Monte Carlo: theory and methods.” In *Bayesian time series models*, 32–51. Cambridge Univ. Press, Cambridge. MR2894232. 1327
- Atchadé, Y. F., Rosenthal, J. S., et al. (2005). “On adaptive Markov chain Monte Carlo algorithms.” *Bernoulli*, 11(5): 815–828. MR2172842. doi: <https://doi.org/10.3150/bj/1130077595>. 1323
- Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2004). *Hierarchical Modeling and Analysis for Spatial Data*. Chapman and Hall/CRC. MR3362184. 1337
- de Valpine, P., Turek, D., Paciorek, C. J., Anderson-Bergman, C., Lang, D. T., and Bodik, R. (2017). “Programming with models: writing statistical algorithms for general model structures with NIMBLE.” *Journal of Computational and Graphical Statistics*, 26(2): 403–413. MR3640196. doi: <https://doi.org/10.1080/10618600.2016.1172487>. 1332
- Deely, J. and Lindley, D. (1981). “Bayes empirical bayes.” *Journal of the American Statistical Association*, 76(376): 833–841. MR0650894. 1335
- Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D. (1987). “Hybrid Monte Carlo.” *Physics Letters B*, 195(2): 216–222. MR3960671. 1340
- Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods*, volume 38. Oxford University Press. MR3014996. doi: <https://doi.org/10.1093/acprof:oso/9780199641178.001.0001>. 1333
- Ecker, M. and Heltshe, J. (1994). “Geostatistical estimates of scallop abundance.” *Case Studies in Biometry*, 107–124. 1337
- Ecker, M. D. and Gelfand, A. E. (1997). “Bayesian variogram modeling for an isotropic

- spatial process.” *Journal of Agricultural, Biological, and Environmental Statistics*, 347–369. MR1817044. doi: <https://doi.org/10.2307/1400508>. 1337
- Everitt, B. S., Landau, S., Leese, M., and Stahl, D. (2011). “Hierarchical clustering.” *Cluster Analysis, 5th Edition*, 71–110. MR3155074. doi: <https://doi.org/10.1002/9780470977811>. 1329
- Gilks, W. R., Roberts, G. O., and Sahu, S. K. (1998). “Adaptive Markov chain Monte Carlo through regeneration.” *Journal of the American Statistical Association*, 93(443): 1045–1054. MR1649199. doi: <https://doi.org/10.2307/2669848>. 1323
- Guennebaud, G., Jacob, B., et al. (2010). “Eigen.” URL: <http://eigen.tuxfamily.org>. 1332
- Haario, H., Saksman, E., and Tamminen, J. (2001). “An adaptive Metropolis algorithm.” *Bernoulli*, 223–242. MR1828504. doi: <https://doi.org/10.2307/3318737>. 1323
- Haario, H., Saksman, E., and Tamminen, J. (2005). “Component-wise adaptation for high dimensional MCMC.” *Computational Statistics*, 20(2): 265–273. MR2323976. doi: <https://doi.org/10.1007/BF02789703>. 1323
- Kass, R. E. and Steffey, D. (1989). “Approximate Bayesian inference in conditionally independent hierarchical models (parametric empirical Bayes models).” *Journal of the American Statistical Association*, 84(407): 717–726. MR1132587. 1335
- Lunn, D., Jackson, C., Best, N., Thomas, A., and Spiegelhalter, D. (2012). *The BUGS Book: A Practical Introduction to Bayesian Analysis*. CRC Press. 1332
- Lunn, D. J., Thomas, A., Best, N., and Spiegelhalter, D. (2000). “WinBUGS—a Bayesian modelling framework: concepts, structure, and extensibility.” *Statistics and Computing*, 10(4): 325–337. 1332
- Nguyen, D., de Valpine, P., Atchade, Y., Turek, D., Michaud, N., and Paciorek, C. (2019). “Nested adaptation of MCMC algorithms: Appendix A.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/19-BA1190SUPP>. 1330
- Plummer, M., Best, N., Cowles, K., and Vines, K. (2006). “CODA: convergence diagnosis and output analysis for MCMC.” *R News*, 6(1): 7–11. 1329
- R Core Team (2013). “R: A Language and Environment for Statistical Computing.” URL: <http://www.R-project.org/>. 1332
- Roberts, G. O. and Rosenthal, J. S. (2001). “Optimal scaling for various Metropolis-Hastings algorithms.” *Statistical Science*, 16(4): 351–367. MR1888450. doi: <https://doi.org/10.1214/ss/1015346320>. 1328
- Roberts, G. O. and Rosenthal, J. S. (2007). “Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms.” *Journal of Applied Probability*, 44(2): 458–475. MR2340211. doi: <https://doi.org/10.1239/jap/1183667414>. 1324, 1327
- Roberts, G. O. and Rosenthal, J. S. (2009). “Examples of adaptive MCMC.” *Journal of*

- Computational and Graphical Statistics*, 18(2): 349–367. MR2749836. doi: <https://doi.org/10.1198/jcgs.2009.06134>. 1325
- Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*. CRC Press. MR2130347. doi: <https://doi.org/10.1201/9780203492024.1336>
- Sahu, S. K., Zhigljavsky, A. A., et al. (2003). “Self-regenerative Markov chain Monte Carlo with adaptation.” *Bernoulli*, 9(3): 395–422. MR1997490. doi: <https://doi.org/10.3150/bj/1065444811>. 1323
- Straatsma, T., Berendsen, H., and Stam, A. (1986). “Estimation of statistical errors in molecular simulation calculations.” *Molecular Physics*, 57(1): 89–95. 1328
- Thompson, M. B. (2010). “Graphical comparison of MCMC performance.” *arXiv preprint arXiv:1011.4457*. 1329
- Tibbits, M. M., Groendyke, C., Haran, M., and Liechty, J. C. (2014). “Automated factor slice sampling.” *Journal of Computational and Graphical Statistics*, 23(2): 543–563. MR3215824. doi: <https://doi.org/10.1080/10618600.2013.791193>. 1329
- Turek, D., de Valpine, P., Paciorek, C. J., and Anderson-Bergman, C. (2017). “Automated parameter blocking for efficient Markov chain Monte Carlo sampling.” *Bayesian Analysis*, 12(2): 465–490. 1325, 1329, 1332, 1336

Acknowledgments

This work was funded in part by a grant from the U.S. National Science Foundation SI2-SSI program (ACI-1550488).