

MIT Open Access Articles

MALA-within-Gibbs Samplers for High-Dimensional Distributions with Sparse Conditional Structure

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

As Published: 10.1137/19M1284014

Publisher: Society for Industrial & Applied Mathematics (SIAM)

Persistent URL: <https://hdl.handle.net/1721.1/135471>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



MALA-WITHIN-GIBBS SAMPLERS FOR HIGH-DIMENSIONAL DISTRIBUTIONS WITH SPARSE CONDITIONAL STRUCTURE*

X. T. TONG[†], M. MORZFELD[‡], AND Y. M. MARZOUK[§]

Abstract. Markov chain Monte Carlo (MCMC) samplers are numerical methods for drawing samples from a given target probability distribution. We discuss one particular MCMC sampler, the MALA-within-Gibbs sampler, from the theoretical and practical perspectives. We first show that the acceptance ratio and step size of this sampler are independent of the overall problem dimension when (i) the target distribution has sparse conditional structure, and (ii) this structure is reflected in the partial updating strategy of MALA-within-Gibbs. If, in addition, the target density is blockwise log-concave, then the sampler's convergence rate is independent of dimension. From a practical perspective, we expect that MALA-within-Gibbs is useful for solving high-dimensional Bayesian inference problems where the posterior exhibits sparse conditional structure at least approximately. In this context, a partitioning of the state that correctly reflects the sparse conditional structure must be found, and we illustrate this process in two numerical examples. We also discuss trade-offs between the block size used for partial updating and computational requirements that may increase with the number of blocks.

Key words. Bayesian computation, high-dimensional distributions, Markov chain Monte Carlo

AMS subject classifications. 62F15, 65C05, 68Q25

DOI. 10.1137/19M1284014

1. Introduction. Markov chain Monte Carlo (MCMC) samplers are numerical methods for drawing samples from an arbitrary target probability distribution whose density is known up to a normalizing constant. Generically, a Metropolis–Hastings MCMC sampler proposes a move by drawing from a proposal distribution and accepts or rejects the move with a probability that ensures that the stationary distribution of the Markov chain is the target distribution.

To design or choose a sampler for a given distribution, one typically considers the following three criteria. First, the type of proposal distribution is chosen based on how much information about the target distribution is available. For example, the Metropolis adjusted Langevin algorithm (MALA) requires derivatives of the target, while the random walk Metropolis (RWM) algorithm does not. Second, the step size, which controls how far the proposed sample strays from the current MCMC state, needs to be tuned. Put simply, too large a step size leads to poor mixing because the acceptance probability is too low; too small a step size leads to a large acceptance probability, but the mixing of the chain is poor because a large number of steps is required to produce an effectively independent sample. Step size tuning must find a

*Submitted to the journal's Methods and Algorithms for Scientific Computing section August 27, 2019; accepted for publication (in revised form) March 17, 2020; published electronically June 16, 2020.

<https://doi.org/10.1137/19M1284014>

Funding: The work of the first author was supported by Singapore MOE AcRF Tier 1 grant R-146-000-292-114. The work of the second author was supported by the Office of Naval Research grant number N00173-17-2-C003. The work of the third author was supported by the DOE Office of Advanced Scientific Computing Research, AEOLUS project through grant number DE-SC0019303, and by the NSF through grant number DMS-1723011.

[†]Department of Mathematics, National University of Singapore, Singapore, 119076 (mattxin@nus.edu.sg).

[‡]Institute of Geophysics and Planetary Physics, Scripps Institution of Oceanography, University of California, San Diego, La Jolla, CA 92093 (matti@ucsd.edu).

[§]Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139 (ymarz@mit.edu).

practical solution to this trade-off and is problem-dependent. Optimal or practical choices of the step size may depend, among other things, on the choice of proposal distribution, the computational resources available, the (apparent or effective) dimension of the problem, and the overall desired accuracy of the MCMC computation. Lastly, in an n -dimensional problem, one can propose an n -dimensional update via an n -dimensional proposal, or one can propose, at each step in the chain, an update for an n/m -dimensional block of variables. Such samplers are called within-Gibbs samplers, partially updating MCMC, componentwise MCMC, or partial resampling algorithms; see, e.g., [37, 29, 5].

The distributions one wishes to sample by MCMC are often high-dimensional. Yet the convergence of MCMC samplers often slows in high dimensions, to the extent that calculations become practically infeasible. Our main motivation for this work is that, while it is certainly difficult to sample generic high-dimensional distributions, distributions that exhibit certain special structure can be feasible to sample, independent of their dimension, provided that the sampler exploits this structure. Examples of samplers in the current literature that leverage various special problem structures are given in section 4. In this paper, we focus on high-dimensional sampling via the MALA-within-Gibbs sampler in the presence of sparse conditional structure.

We define sparse conditional structure in section 3 via the Hessian of the logarithm of the target density. In the special case of a Gaussian target distribution, sparse conditional structure is equivalent to the precision matrix being sparse. More generally, sparse conditional structure is equivalent to the existence of many conditional independence relationships, or the distribution being Markov with respect to a sparse graph [31]. We prove in section 3 that the partial updating strategy of MALA-within-Gibbs, with carefully defined updates that make use of the sparse conditional structure, leads to acceptance ratios that depend on the dimension of the block-update but are independent of the overall dimension. We further show that MALA-within-Gibbs converges with a rate independent of the dimension if the target distribution is blockwise log-concave.

We then discuss MALA-within-Gibbs from a practical perspective in section 5. In this context it is important to realize that the sparse conditional structure may become apparent only after a suitable change of coordinates. For MALA-within-Gibbs to be an effective sampler, we thus need a means of discovering these coordinates or, equivalently, identifying sparse conditional structure. We expect that many Bayesian inference problems (see, e.g., [15, 41, 2]) are naturally formulated in coordinates that exhibit sparse conditional structure, but we also consider an example where a coordinate transformation is required to reveal sparse conditional structure. We further discuss the overall computational efficiency of the MALA-within-Gibbs approach and raise the issue that the partial updating of MALA-within-Gibbs requires m simulations of the numerical model per sample, m being the number of blocks. This implies that there may be an optimal, problem-dependent choice for the dimensionality of the update that can lead to significant computational savings. We explore all of the above issues numerically by applying MALA-within-Gibbs to two well-known test problems: a log-Gaussian Cox point process [24, 33] and an elliptic PDE inverse problem [35, 36, 38, 4, 20, 48]. We further compare the computational efficiency of MALA-within-Gibbs to the efficiencies of other samplers including MALA, pCN [17], and manifold MALA (MMALA) [24].

2. Notation, assumptions, and background. We consider probability distributions with density functions $c\pi(\mathbf{x})$, where c is an unknown normalization constant

and π is a known function. We partition the n -dimensional vector \mathbf{x} into m blocks, $\mathbf{x}_1, \dots, \mathbf{x}_m$, where the subscripts are called block indices. Note that the blocks \mathbf{x}_j are not necessarily consecutive elements of the vector \mathbf{x} .

2.1. Notation. MALA will require gradients of the logarithm of the target density π , which we write as $\mathbf{v}(\mathbf{x}) = \nabla_{\mathbf{x}} \log \pi(\mathbf{x})$. Similarly, we sometimes write derivatives of π with respect to the blocks as $\mathbf{v}_j(\mathbf{x}) = \nabla_{\mathbf{x}_j} \log \pi(\mathbf{x})$. We write $\nabla_{\mathbf{x}_i, \mathbf{x}_j}^2$ to denote second derivatives with respect to blocks i and j ; i.e., $\nabla_{\mathbf{x}_i, \mathbf{x}_j}^2 \log \pi(\mathbf{x})$ is a matrix of size $\dim(\mathbf{x}_i) \times \dim(\mathbf{x}_j)$. Throughout this paper, we use the Euclidean norm for vectors, i.e., $\|\mathbf{x}\| = \sqrt{\mathbf{x}^T \mathbf{x}}$, where superscript T denotes a transpose, and the l_2 -operator norm, $\|\mathbf{A}\|$, for matrices. We write $\mathbf{A} \preceq \mathbf{B}$, where \mathbf{A} and \mathbf{B} are two $n \times n$ matrices, when the matrix $\mathbf{A} - \mathbf{B}$ is negative semidefinite. We write $\lambda_{\min}(\mathbf{A})$ for the smallest eigenvalue of the matrix \mathbf{A} .

We write conditional densities of one block, $\pi(\mathbf{x}_j | \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{j-1}, \mathbf{x}_{j+1}, \dots, \mathbf{x}_m)$, as $\pi(\mathbf{x}_j | \mathbf{x}_{\setminus j})$, i.e., the block index set $\setminus j = \{1, 2, \dots, j-1, j+1, \dots, m\}$. More generally, we write \mathcal{I} for a subset of block indices; i.e., \mathcal{I} is a subset of $\{1, 2, \dots, m\}$. The cardinality of \mathcal{I} will be denoted as $|\mathcal{I}|$. We denote the complement of \mathcal{I} by \mathcal{I}^c ; i.e., \mathcal{I}^c is the subset of $\{1, 2, \dots, m\}$ which excludes the block indices in \mathcal{I} .

An important concept we will use repeatedly is conditional independence. Conditional independence means that conditioning block i on all but a few other blocks is irrelevant, which we write as

$$\mathbf{x}_j \perp\!\!\!\perp \mathbf{x}_{\mathcal{I}_j^c} \mid \mathbf{x}_{\mathcal{I}_j \setminus \{j\}},$$

where the index set \mathcal{I}_j depends on j and includes the block index j and where $\mathcal{I}_j \setminus \{j\}$ is the index set \mathcal{I}_j with index j removed. In terms of probability distributions, conditional independence means that

$$\pi(\mathbf{x}_j | \mathbf{x}_{\setminus j}) = \pi(\mathbf{x}_j | \mathbf{x}_{\mathcal{I}_j \setminus \{j\}}).$$

We assume throughout that \mathcal{I}_j has at most $S \ll m$ elements.

2.2. Assumptions. We assume throughout this paper that $\pi(\mathbf{x})$ has continuous second derivatives and that

- (i) the dimension, n , of \mathbf{x} and the number of blocks, m , are large;
- (ii) any block \mathbf{x}_j is conditionally independent of most other blocks.

We refer to assumption (ii) as sparse conditional structure. This terminology is inspired by linear algebra and Gaussian $\pi(\mathbf{x})$ —a Gaussian with sparse conditional structure is characterized by a sparse precision matrix. We make assumption (ii) mathematically more precise in section 3.1. For simplicity, we assume that n/m (dimension divided by the number of blocks) is an integer.

2.3. Background: MALA and MALA-within-Gibbs. The MALA sampler with n -dimensional updates and step size τ generates a sequence of iterates \mathbf{x}^k by repeating the following two steps, starting from a given \mathbf{x}^0 :

1. Draw a sample $\tilde{\mathbf{x}}^k$ from the MALA proposal by

$$\tilde{\mathbf{x}}^k = \mathbf{x}^k + \tau \mathbf{v}(\mathbf{x}^k) + \sqrt{2\tau} \boldsymbol{\xi}^k,$$

where $\boldsymbol{\xi}^k$ is an independent sample from $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$.

2. Accept this proposal with probability

$$\alpha(\mathbf{x}^k, \tilde{\mathbf{x}}^k) = \min \left\{ 1, \frac{\pi(\tilde{\mathbf{x}}^k) \exp(-\frac{1}{4\tau} \|\mathbf{x}^k - \tilde{\mathbf{x}}^k - \tau \mathbf{v}(\tilde{\mathbf{x}}^k)\|^2)}{\pi(\mathbf{x}^k) \exp(-\frac{1}{4\tau} \|\tilde{\mathbf{x}}^k - \mathbf{x}^k - \tau \mathbf{v}(\mathbf{x}^k)\|^2)} \right\};$$

i.e., let $\mathbf{x}^{k+1} = \tilde{\mathbf{x}}^k$ with probability $\alpha(\mathbf{x}^k, \tilde{\mathbf{x}}^k)$, and $\mathbf{x}^{k+1} = \mathbf{x}^k$ with probability $1 - \alpha(\mathbf{x}^k, \tilde{\mathbf{x}}^k)$.

It is straightforward to show that $c\pi(\mathbf{x})$ is the invariant distribution of the Markov chain \mathbf{x}^k . Therefore, when $k \rightarrow \infty$, \mathbf{x}^k can be viewed as a sample from $c\pi(\mathbf{x})$.

MALA-within-Gibbs is a variation of MALA that uses n/m -dimensional updates. We use superscripts to index “time” in the Markov chain (see above) and subscripts to index the blocks. Thus, starting with a vector \mathbf{x}^0 and step size τ , MALA-within-Gibbs iterates the following steps:

1. Set $\mathbf{x}^k = \mathbf{x}^{k-1}$. Repeat steps (a) and (b) below for $j = 1, \dots, m$ to update all m blocks of $\mathbf{x}^k = [\mathbf{x}_1^k, \dots, \mathbf{x}_m^k]$.

- (a) Sample a standard Gaussian $\boldsymbol{\xi}_j^k$ of the same dimension as \mathbf{x}_j and use the MALA proposal for the current block \mathbf{x}_j :

$$(2.1) \quad \tilde{\mathbf{x}}_j^k = \mathbf{x}_j^k + \tau \mathbf{v}_j(\mathbf{x}^k) + \sqrt{2\tau} \boldsymbol{\xi}_j^k.$$

- (b) Define $\tilde{\mathbf{x}}^k = [\mathbf{x}_1^k, \dots, \mathbf{x}_{j-1}^k, \tilde{\mathbf{x}}_j^k, \mathbf{x}_{j+1}^k, \dots, \mathbf{x}_m^k]$, i.e., $\tilde{\mathbf{x}}^k$ is equal to \mathbf{x}^k , except at its j th block. Compute the block acceptance ratio

$$\alpha_j(\mathbf{x}^k, \tilde{\mathbf{x}}^k) = \min \left\{ 1, \frac{\pi(\tilde{\mathbf{x}}^k) \exp(-\frac{1}{4\tau} \|\mathbf{x}^k - \tilde{\mathbf{x}}^k - \tau \mathbf{v}_j(\tilde{\mathbf{x}}^k)\|^2)}{\pi(\mathbf{x}^k) \exp(-\frac{1}{4\tau} \|\tilde{\mathbf{x}}^k - \mathbf{x}^k - \tau \mathbf{v}_j(\mathbf{x}^k)\|^2)} \right\}.$$

Set \mathbf{x}^k to be $\tilde{\mathbf{x}}^k$ with probability $\alpha_j(\mathbf{x}^k, \tilde{\mathbf{x}}^k)$; else \mathbf{x}^k maintains its value.

2. Increase the time index from k to $k + 1$ and go to 1.

As before, it is straightforward to verify that the target $c\pi(\mathbf{x})$ is the invariant distribution of the MALA-within-Gibbs iterates. The partial updating can be derived from applying MALA within a Gibbs iteration (hence the name), i.e., with target distributions $\pi(\mathbf{x}_j | \mathbf{x}_{\setminus j})$.

3. Dimension-independent acceptance and convergence rate. Both the MALA and MALA-within-Gibbs samplers can, in principle, be used for arbitrary target distributions, but in generic high-dimensional problems we expect that convergence is slow. In high-dimensional problems with sparse conditional structure, however, MALA-within-Gibbs can be effective if the partitioning of \mathbf{x} , defining the partial updates, is chosen in accordance with the sparse conditional structure. With a suitable partial updating strategy, we show that the step size and the acceptance ratio of MALA-within-Gibbs (within each block) can be made independent of the overall dimension. We then show, under additional assumptions of blockwise log-concavity, that MALA-within-Gibbs converges to the target distribution at a dimension-independent rate. The proofs of the propositions and the theorem can be found in the Supplementary Material.

3.1. Dimension-independent acceptance rate under sparse conditional structure. To simplify the proofs, the conditional independence assumption is formulated in terms of the gradient $\mathbf{v}(\mathbf{x})$.

ASSUMPTION 3.1 (sparse conditional structure). *For $\pi(\mathbf{x})$ and the partition $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_m)$, there are constants S and q independent of n , so that the following hold:*

- (i) *The dimension of each block \mathbf{x}_j is bounded by q .*
- (ii) *For each block index $j \in 1, \dots, m$, there is an \mathbf{x} -independent block index set $\mathcal{I}_j \subset \{1, \dots, m\}$ with $j \in \mathcal{I}_j$ and cardinality $|\mathcal{I}_j| \leq S$ so that*

$$\nabla_{\mathbf{x}_k, \mathbf{x}_j}^2 \log \pi(\mathbf{x}) = \nabla_{\mathbf{x}_k} \mathbf{v}_j(\mathbf{x}) = \mathbf{0} \quad \text{if } k \notin \mathcal{I}_j.$$

Note that (i) is trivial because we deal with finite-dimensional problems and note that (ii), by Lemma 2 of [46], is equivalent to $\mathbf{x}_j \perp\!\!\!\perp \mathbf{x}_{\mathcal{I}_j^c} \mid \mathbf{x}_{\mathcal{I}_j \setminus \{j\}}$ if the density is strictly positive and smooth. In other words, Assumption 3.1 is equivalent to the assumption of sparse conditional structure, as described earlier, but the formulation in terms of gradients is easier to use in our proofs.

Whether or not Assumption 3.1 is satisfied for a given target distribution depends, to a large extent, on how the blocks of \mathbf{x} are defined. Using physical insight into the problem, it is often possible to group components of \mathbf{x} such that Assumption 3.1 is satisfied or approximately satisfied. We discuss this issue more in section 5 below, but it is important to understand that Assumption 3.1 essentially requires a good understanding of the target distribution and that the results we derive under this assumption make use of the fact that one understands and leverages conditional independencies among the components of \mathbf{x} .

We also assume that the gradient $\mathbf{v}(\mathbf{x})$ and its derivatives are bounded.

ASSUMPTION 3.2 (bounded vector fields). *The vector field $\mathbf{v}_j(\mathbf{x}) = \nabla_{\mathbf{x}_j} \log \pi(\mathbf{x})$, for $j = 1, \dots, m$, and its first derivatives are bounded, i.e., there exist constants M_v and H_v , independent of the overall dimension n , such that*

$$\|\mathbf{v}_j(\mathbf{x})\| \leq M_v, \quad \|\nabla_{\mathbf{x}_i} \mathbf{v}_j(\mathbf{x})\| \leq H_v, \quad \|\nabla_{\mathbf{x}_j} \mathbf{v}_j(\mathbf{x}) - \nabla_{\mathbf{z}_j} \mathbf{v}_j(\mathbf{z})\| \leq H_v \|\mathbf{x} - \mathbf{z}\|.$$

By Assumption 3.1, $\mathbf{v}_j(\mathbf{x})$ has no dependence on $\mathbf{x}_{\mathcal{I}_j^c}$, so one can write it as $\mathbf{v}_j(\mathbf{x}_{\mathcal{I}_j})$. If the support of $\pi(\mathbf{x})$ is bounded, then $\mathbf{v}_j(\mathbf{x})$ having no dependence on $\mathbf{x}_{\mathcal{I}_j^c}$, along with the fact that the dimension of each block \mathbf{x}_i is at most q , often yields Assumption 3.2. Unbounded support is more complicated. A Gaussian, for example, violates Assumption 3.2 because the norm of the gradient is not bounded. This boundedness assumption, however, is made for simplicity and may not be required in practice. More sophisticated constructions may be used in the future to relax this assumption and to derive more general results.

Under Assumptions 3.1 and 3.2, the following proposition shows that the step size and the acceptance ratio of MALA-within-Gibbs are independent of the overall dimension.

PROPOSITION 3.1 (block acceptance). *Suppose $c\pi(\mathbf{x})$ is the density of a distribution with sparse conditional structure (Assumption 3.1) and that, in addition, Assumption 3.2 holds. There is a constant M , independent of the number of blocks m , so that, for any given state $\mathbf{x} \in \mathbb{R}^n$, the block acceptance ratio $\alpha_j(\mathbf{x}^k, \tilde{\mathbf{x}}^k)$ is bounded below by*

$$\mathbb{E}[\alpha_j(\mathbf{x}^k, \tilde{\mathbf{x}}^k)] \geq 1 - M\sqrt{\tau}$$

for all blocks $j \in \{1, \dots, m\}$.

Proposition 3.1 follows directly from Lemma SM1.2, which we prove in section SM1. The dimension-independent block acceptance ratio is intuitive. Partial updating implies that the proposed updates are, by design, low-dimensional: their dimension depends on the block size, q , but is independent of the number of blocks, m , or the overall dimension $n = m \cdot q$. Thus, only the dimension of the update controls the block acceptance ratio. The overall number of low-dimensional updates, which defines the overall dimension, is irrelevant.

It might seem that Proposition 3.1 contradicts earlier results on optimal scalings of MALA step sizes, where the optimal step size decreases with dimension at a well-understood rate [42, 43, 8, 7]. These earlier scaling results, however, do not assume sparse conditional structure. Thus, in general, the optimal step size of MALA and

MALA-within-Gibbs should decrease with dimension, but if the target distribution has sparse conditional structure and if, in addition, this structure is used in the block-updates, then the step size (and acceptance ratio) can be independent of the overall dimension.

Assumption 3.1 ensures that the sparse conditional structure of the target is exploited by the MALA-within-Gibbs sampler. We thus assume away any difficulties of discovering sparse conditional structure, but we discuss practical aspects of this assumption in section 5, including a brief discussion of what happens when the assumptions are only nearly met. We also emphasize that we have no claims at optimal step sizes of MALA-within-Gibbs—we merely show that the step size need not decrease with dimension to ensure a constant average block acceptance ratio. Moreover, local tunings as discussed in [5] may further improve efficiency, but we do not pursue such ideas here.

Partial updating of MALA has also been considered in [37], where the conclusion is that the updates in MALA-within-Gibbs should be high-dimensional. Again, this is true in general, but if the target has sparse conditional structure, the dimensionality of the updates may depend on this structure. We revisit this issue in section 5, where we also bring up a trade-off between the block size and computational requirements that may increase with the number of blocks. One can also perform random partial updates, i.e., choosing at random which components of \mathbf{x} are next updated. Asymptotically, MALA-within-Gibbs with a random partial updating strategy converges to the target distribution, but we expect that the convergence will be slow for problems with sparse conditional structure because this structure is not used by random partial updates.

3.2. Dimension-independent convergence rate. We have shown above that the step size and acceptance ratio of MALA-within-Gibbs can be independent of dimension if the sparse conditional structure of the target distribution is known and used via a suitable partition of the variables during the within-Gibbs moves. This is not enough to guarantee fast convergence of MALA-within-Gibbs. To study the convergence rate of MALA-within-Gibbs, we require, as an additional assumption, that the target distribution be unimodal and blockwise log-concave (see below for a definition). The reason is that difficulties with MCMC that arise from high dimensionality or multimodality are independent of each other: if the target distribution has multiple modes, a large number of samples may be required even if the dimension is small. We focus on aspects of high-dimensional problems with a single mode.

The additional assumption we need in our proof (see section SM1) is blockwise log-concavity. To define blockwise log-concavity, we first construct an $m \times m$ matrix $\mathbf{H}(\mathbf{x})$, where m is the number of blocks, with the following properties.

DEFINITION 3.3. *A symmetric $m \times m$ matrix function $\mathbf{H}(\mathbf{x})$ with entries $H_{j,i}(x)$ is uniformly bounded and negative if there are strictly positive constants H_v and λ_H such that for all j, i and all \mathbf{x} ,*

$$|H_{j,i}(\mathbf{x})| \leq H_v, \quad \lambda_{\max}(\mathbf{H}(\mathbf{x})) \leq -\lambda_H < 0.$$

As a simple example, a constant symmetric negative definite matrix is uniformly bounded and negative. Blockwise log-concavity can now be formulated as follows.

ASSUMPTION 3.4. *A probability density $c\pi(\mathbf{x})$ is blockwise log-concave with block size m if there exists an $m \times m$ uniformly bounded and negative matrix $\mathbf{H}(\mathbf{x})$ such that*

TABLE 1

Blockwise log-concavity with \mathbf{H} as defined in the text for different block sizes q and different correlation length scales l . Positive numbers, highlighted in bold, indicate blockwise log-concavity.

Setting	$q = 1$	$q = 2$	$q = 4$	$q = 8$	$q = 16$	$q = 32$
$l = 2$	-0.71	-1.28	-1.44	-1.28	-0.71	0.24
$l = 1$	0.04	-0.21	-0.29	-0.21	0.04	0.46
$l = 0.5$	0.62	0.54	0.52	0.54	0.62	0.76

- (i) $\nabla_{\mathbf{x}_j, \mathbf{x}_j}^2 \log \pi(\mathbf{x}) \preceq H_{j,j}(\mathbf{x}) \mathbf{I}$, where \mathbf{I} is the identity matrix of size $\dim(\mathbf{x}_j) \times \dim(\mathbf{x}_j)$;
- (ii) the off-diagonal elements bound the conditional dependence between blocks, that is, $\|\nabla_{\mathbf{x}_j, \mathbf{x}_i}^2 \log \pi(\mathbf{x})\| \leq H_{j,i}(\mathbf{x})$ for all $i \neq j$.

Note that if the dimension of the blocks is $q = 1$, so that $m = n$, and if the Hessian of $\log \pi(\mathbf{x})$ is diagonally dominant, then $\mathbf{H}(\mathbf{x})$ can be taken as the Hessian of π , with all off-diagonal entries replaced by their absolute value. Further note that blockwise log-concavity is a stronger assumption than log-concavity—the function $\pi(\mathbf{x})$ can be log-concave but not blockwise log-concave (see example below). On the other hand, a distribution that is blockwise log-concave, for any block size, is also log-concave.

As an illustration, we consider a Gaussian distribution for $\mathbf{x} = [x_1, \dots, x_{64}]$ with mean zero and covariance matrix \mathbf{C} with elements

$$[C]_{i,j} = \exp\left(-\frac{1}{2l}|i-j|\right), \quad i, j = 1, \dots, 64.$$

Interpreting this Gaussian as a discretization of a 1D random field with exponential covariance kernel (and discretization $\Delta x = 1$), the quantity l is a correlation length scale. If l is small, only those components of \mathbf{x} that are near each other in the 1D domain are significantly correlated. This suggests partitioning \mathbf{x} based on neighborhoods in the 1D domain which correspond to consecutive elements of \mathbf{x} . For example, the block size $q = 4$ results in $m = 16$ blocks

$$\mathbf{x}_1 = [x_1, \dots, x_4], \quad \mathbf{x}_2 = [x_5, \dots, x_8], \quad \dots, \quad \mathbf{x}_{16} = [x_{61}, \dots, x_{64}].$$

Recall that, for Gaussian distributions, the precision matrix \mathbf{P} is equal to $-2\nabla^2 \log \pi$, which suggests constructing the matrix $\mathbf{H}(\mathbf{x}) \in \mathbb{R}^{m \times m}$ by

$$H_{i,i}(\mathbf{x}) \equiv -\lambda_{\min}(\mathbf{P}_{i,i}), \quad H_{i,j}(\mathbf{x}) \equiv \|\mathbf{P}_{i,j}\|.$$

Here, $\mathbf{P}_{i,j}$ is the i, j th $q \times q$ subblock of \mathbf{P} with indices corresponding to the blocks \mathbf{x}_i and \mathbf{x}_j . For example, with $q = 4$, $\mathbf{P}_{1,2}$ is a subblock of \mathbf{P} consisting of rows 1–4 and columns 5–8. Assumption 3.4 is then equivalent to assuming that \mathbf{H} is negative definite, i.e., $\lambda_{\min}(-\mathbf{H}) > 0$. We can numerically check this condition by computing eigenvalues of \mathbf{H} . Table 1 lists values of $\lambda_{\min}(-\mathbf{H})$ for varying correlation length scales l and block sizes q .

We note that while the Gaussian is log-concave for any l , blockwise log-concavity depends on the length scale l and the size of the blocks q . If l is large, only large blocks lead to blockwise log-concavity (with $q = 64$ guaranteeing log-concavity and blockwise log-concavity). If the correlation is (essentially) confined to small neighborhoods, i.e., if l is small, then small block sizes q also lead to blockwise log-concavity.

With the definition of blockwise log-concavity, we can now state a theorem about the dimension-independent convergence rate of MALA-within-Gibbs. The proof is given in section SM1.

THEOREM 3.2. *Under Assumptions 3.1, 3.2, and 3.4, for any $\delta > 0$, there exists a $\tau_0 > 0$ independent of the number of blocks m , so that when the step size $\tau < \tau_0$, we can couple two MALA-within-Gibbs samples \mathbf{x}^k and \mathbf{z}^k , such that*

$$\sum_{i=1}^m (\mathbb{E}\|\mathbf{x}_i^k - \mathbf{z}_i^k\|)^2 \leq (1 - (1 - \delta)\lambda_H\tau)^{2k} \sum_{i=1}^m (\mathbb{E}\|\mathbf{x}_i^0 - \mathbf{z}_i^0\|)^2.$$

In particular, one can let $\mathbf{z}^0 \sim \pi$. It follows that $\mathbf{z}^k \sim \pi$, which in turn shows that \mathbf{x}^k converges to π geometrically fast.

Theorem 3.2 indicates that MALA-within-Gibbs can be a fast sampler for high-dimensional problems if (i) the target distribution has sparse conditional structure and this structure is used by the MALA-within-Gibbs sampler; and (ii) the target distribution is blockwise log-concave.

Blockwise log-concavity implies that the target has only one mode. Assuming blockwise log-concavity thus allows us to study computational barriers due to high dimensionality without requiring that we simultaneously consider challenges due to multimodality. Notably, blockwise log-concavity is more restrictive than log-concavity, which also implies that the target is unimodal. We use blockwise log-concavity here to gain stronger control over the coupling between blocks in the analysis. Ultimately, a less restrictive assumption (e.g., log-concavity) may be preferable, and one may view our results as a first step towards a full understanding of how MALA-within-Gibbs can operate effectively in high-dimensional problems.

Theorem 3.2 also has connections to previous work on Gibbs samplers for Gaussian distributions with sparse conditional structure [34]. In particular, Theorems 3.1 and 3.2 of [34] show dimension-independent convergence of a Gibbs sampler for Gaussian distributions. By interpreting the blockwise log-concavity assumption as a generalization of the Gaussian assumptions in Theorems 3.1 and 3.2 of [34], one can understand Theorem 3.2 as a generalization of this result to a within-Gibbs sampler for non-Gaussian distributions.

4. Discussion of efficient samplers in high dimensions. We suggested earlier that sampling generic high-dimensional distributions is difficult, but if the target distribution has a special structure, then efficient samplers can be constructed. One example is Gaussian distributions. Gaussians can be sampled efficiently even if their dimension is large, either by direct samplers (using techniques from numerical linear algebra for computing matrix square roots), or by MCMC, using analogies between Gibbs samplers and linear solvers to construct matrix splittings for accelerated sampling [23, 22]. Connections between linear solvers and MCMC are also discussed in [25].

There are also several routes to making MCMC samplers effective for high-dimensional distributions that are not Gaussian, and we discuss some of them here in relation to our results. Recall that the MALA-within-Gibbs sampler relies on sampling blocks of variables; this idea is also used in multigrid Monte Carlo (MGMC) for lattice systems [26, 27, 21]. The basic ideas in MGMC, however, are different. MGMC proposes the same move simultaneously for all variables within one block. The MALA-within-Gibbs sampler proposes moves for each variable within a block, but handles the blocks (nearly) independently. In terms of analogues to linear solvers, MALA-within-Gibbs is perhaps more akin to domain decomposition than to multigrid.

Metropolis–Hastings (MH) samplers offer a general tool for sampling non-Gaussian

target distributions. The step size τ of many MH samplers needs to be tuned, which often requires that τ decrease with dimension n . In fact, optimal scalings of τ with dimension for various MCMC samplers have been derived for RWM $\tau_{\text{opt}} = O(n^{-1})$, for MALA $\tau_{\text{opt}} = O(n^{-1/3})$, and for HMC $\tau_{\text{opt}} = O(n^{-1/4})$. Fixing the acceptance ratio with small τ , however, comes with a price: the acceptance ratio may be large, but all accepted steps are small (on average on the order of τ), so that the sampler moves often, but slowly. As a rule of thumb, it takes about $O(1/\tau)$ iterations to move through the support of the target distribution (putting aside issues of reaching stationarity [16]). For very large dimensions, many MCMC samplers are thus slow to converge. These results hold for general target distributions. Even if the analyses that lead to the optimal scalings rely on certain assumptions, e.g., that the target measure is of product type, this problem structure is not directly used by the samplers. Our results on dimension-independent step size and convergence rates hold only for target distributions with sparse conditional structure and when this structure is explicitly used by the MALA-within-Gibbs sampler (via Assumption 3.1).

Another strategy for effective sampling of high-dimensional distributions applies if the parameters \mathbf{x} can be decomposed as $\mathbf{x} = (\mathbf{y}, \mathbf{z})$, where \mathbf{z} is low-dimensional and, conditioned on \mathbf{z} , there are fast (perhaps direct) samplers for \mathbf{y} ; see, e.g., [10, 11, 13]. In Bayesian inverse problems, this structure can often be identified by a suitable choice of basis or reparameterization, as in [19, 47, 50]: typically \mathbf{z} represents directions where the posterior departs significantly from the prior, while \mathbf{y} represents prior-dominated directions of the parameter space, which may even be conditionally Gaussian and/or (approximately) independent of \mathbf{z} . Note that the MALA-within-Gibbs sampler does not require partially or conditionally Gaussian target distributions, but our analysis of its efficiency requires sparse conditional structure and blockwise log-concavity.

The theory of function-space MCMC also has led to effective MCMC methods for another class of high-dimensional Bayesian inverse problems; see, e.g., [17, 28, 48, 39]. The basic idea is to design MCMC samplers that are well defined on function spaces: consider, for example, the inference of a spatially distributed parameter, and send the spatial discretization parameter h to zero while keeping the number of observations fixed. Since the spatial discretization controls the dimension of the problem, the limit $h \rightarrow 0$ corresponds to an infinite-dimensional problem. The proposal distributions of function-space MCMC are chosen such that, for a fixed step size, the acceptance ratio remains constant, independent of h or, equivalently, the dimension of the problem. There are many variations of such discretization invariant MCMC samplers [19, 44, 14, 6], with applications discussed in, e.g., [9, 40]. Some (e.g., [18, 3, 30]) combine function-space MCMC with the decomposition of the parameter space into (finite-dimensional) likelihood-informed directions and (infinite-dimensional) prior-dominated directions, as described above.

The dimension-independence of the MALA-within-Gibbs sampler discussed here is markedly different from the discretization invariance of function-space MCMC. We do not consider the infinite-dimensional limit resulting from the discretization of a function on a given domain. Rather, we show that the dimension of the inference problem may not affect the acceptance and convergence rates of the sampler if the problem has sparse conditional structure that is suitably exploited. This allows us, for example, to consider a sequence of inference problems posed on increasingly large domains and with an increasing number of observations, while keeping the spatial discretization fixed (see section 5.3 and Figure 5.1). Provided that each of the problems within the sequence has suitable sparse conditional structure, the acceptance and convergence rates of MALA-within-Gibbs are the same for all problems within the se-

quence. The convergence rate of function-space MCMC, e.g., the pCN method, is not constant for all problems within such a sequence (see section 5.3 for a specific example). This does not come as a surprise, because the dimension independence of pCN holds only in the case in which the dimension increases due to grid refinement, with the number of observations and the domain size held constant. A sequence of problems with increasing domain size and increasing numbers of observations violates the assumptions underpinning the dimension independence of pCN. We refer the reader to [34] for a more thorough discussion of the various notions of high dimensionality that may occur when solving Bayesian inverse problems.

5. Practical considerations and numerical experiments. Our results on dimension-independent step size, acceptance probabilities, and convergence rates hold under precise mathematical assumptions of sparse conditional structure and log-concavity (see section 3). We now focus on posterior distributions that arise in Bayesian inference problems, because of their practical importance and because we anticipate that the assumption of sparse conditional structure is often satisfied in such problems. We also demonstrate how to use MALA-within-Gibbs in two numerical examples and discuss and compare the computational costs of MALA-within-Gibbs and other MCMC samplers.

5.1. Posterior distributions with sparse conditional structure. Below we provide the Bayesian problem setup. Let \mathbf{x} be an n -dimensional vector endowed with a prior probability density $\pi_0(\mathbf{x})$. In many problems, \mathbf{x} arises from a discretization of a spatially distributed quantity (i.e., a field) and, for that reason, is high-dimensional. The prior reflects assumptions about the smoothness of the field and is often assumed to be Gaussian with a known mean and covariance. A computational model, $\mathcal{M}(\mathbf{x})$, maps \mathbf{x} to observations \mathbf{y} . Typically, the model is nonlinear and the number of observations is less than the dimension of \mathbf{x} . Any model errors are represented by a random variable $\boldsymbol{\varepsilon}$, and, often, model errors are additive, i.e.,

$$(5.1) \quad \mathbf{y} = \mathcal{M}(\mathbf{x}) + \boldsymbol{\varepsilon}.$$

The distribution of $\boldsymbol{\varepsilon}$ is assumed to be known (often Gaussian with mean zero and diagonal covariance matrix). Equation (5.1) defines a likelihood $\pi_l(\mathbf{y}|\mathbf{x})$, and the likelihood and prior jointly define the posterior distribution

$$\pi(\mathbf{x}|\mathbf{y}) \propto \pi_0(\mathbf{x})\pi_l(\mathbf{y}|\mathbf{x}).$$

Sparse conditional structure arises naturally in Bayesian posterior distributions when (i) the parameters \mathbf{x} are high-dimensional, but not all components of \mathbf{x} have significant statistical interactions; and (ii) each observation is informative for only a small subset of the components of \mathbf{x} (see also [34]). Put differently, we assume that the prior has sparse conditional structure and that the observations do not significantly densify the conditional structure of the prior. This happens in many geophysical applications, e.g., in numerical weather prediction (NWP), where the posterior distribution is defined jointly by a global atmospheric model (with dimension $O(10^8)$) and observations of the atmospheric state (typically $O(10^7)$ observations). In a global atmospheric model, each model component stores information about the atmospheric state at a specific location at a given time, and each component has significant statistical interactions with nearby components, but not with components that are far away. A discussion of the mathematical mechanisms that lead to this property can be found in [12].

5.2. Implementation of MALA-within-Gibbs. The partitioning of \mathbf{x} into blocks is important for effective sampling of the posterior by MALA-within-Gibbs, because only a suitable partition will indeed put the sparse conditional structure to use. We do not have a general strategy to find a suitable partition, but we expect that a workable partition is often intuitive. For example, if \mathbf{x} is defined over a spatial domain (1D, 2D, or 3D) and if correlations are limited to small neighborhoods, then the partitioning should be based on these neighborhoods and the block size should take the correlation lengths' scales into account. We demonstrate this process in a numerical example in section 5.3. Our second example in section 5.4 demonstrates how to choose a partition for the partial updating based on prior covariances in the absence of a spatial scale.

The partial updating of MALA-within-Gibbs requires m likelihood evaluations per sample. In a typical Bayesian inverse problem, each likelihood evaluation will require a full forward solve with the numerical model \mathcal{M} , even if only one block of the model's components is updated. Using the common effective sample size

$$(5.2) \quad N_{\text{eff}} = N_e / \text{IACT},$$

where N_e is the number of MCMC samples (length of the Markov chain) and IACT is the integrated auto correlation time (see, e.g., [49, 45]), we can estimate the cost per effective sample by

$$(5.3) \quad \text{cost per effective sample} = \text{IACT} \times (\# \text{ of blocks}) \times \text{cost of likelihood evaluation}.$$

It is now clear that the computational cost of MALA-within-Gibbs grows with the number of blocks, even if IACT is independent of dimension. The computational requirements of MALA-within-Gibbs are, therefore, not independent of the dimension of the problem, since a higher-dimensional problem requires a larger number of blocks (keeping other parameters that define the model unchanged; see examples below). This also points to a trade-off for sampling in high dimensions that may not be easy to resolve: to keep the efficiency high (small IACT), one may want to use a large number of small blocks (with a lower bound on the block size depending on the correlation structure), but on the other hand, one may want to use a small number of large blocks to keep the number of model evaluations per sample small.

5.3. Numerical illustration 1: Log-Gaussian Cox point processes. We consider inference in a log-Gaussian Cox point process similar to the numerical experiments in [24]. A uniform $N_u \times N_v$ grid, with spacing $\Delta u = \Delta v = 1$, covers the 2D (spatial) domain, $[1, L] \times [1, L]$. The parameter to be inferred is defined over the domain, $\{X_{i,j}, i, j = 1, \dots, L/\Delta u\}$. Its prior is $\mathcal{N}(\mu \mathbf{1}, \mathbf{B})$, where \mathbf{B} is a discretization of the exponential covariance kernel, i.e.,

$$\text{cov}(X_{s_1, t_1}, X_{s_2, t_2}) = \sigma_s^2 \sigma_t^2 \exp\left(-\frac{1}{2} \frac{|s_1 - s_2|}{l_s} - \frac{1}{2} \frac{|t_1 - t_2|}{l_t}\right),$$

where $\sigma_s^2 = \sigma_t^2 = 2$, $\mu = 4$, $l_s = 2$, $l_t = 4$.

Observations are made at each grid point, denoted by $Y_{i,j}$. The observations are conditionally independent and Poisson distributed with means $\exp(X_{i,j})$. Our goal is to estimate $X_{i,j}$ from $Y_{i,j}$. The prior and likelihood define the posterior distribution

$$\pi(\mathbf{x}|\mathbf{y}) \propto \exp\left(-\frac{1}{2} \|\mathbf{B}^{-1/2}(\mathbf{x} - \mu \mathbf{1})\|^2\right) \prod_{i,j} \exp(Y_{i,j} X_{i,j} - \exp(X_{i,j})),$$

where \mathbf{x} is the column stack of $X_{i,j}$, i.e., an $n = L^2$ dimensional vector.




FIG. 5.1. True values of X for three problems with increasing domain size L (drawn to scale).


5.3.1. Problem setup. We consider three problems with increasing domain size $L = 16$, $L = 32$, and $L = 64$, leading to sampling problems of dimensions 256, 1024, and 4096. The true values of $X_{i,j}$ for the three domains are shown in Figure 5.1. Note that the (apparent) dimension of the problem ($n = L^2$) and the number of observations (L^2) increase with increasing domain size, but the prior length scales are fixed and short compared to all three domain sizes. Moreover, each observation $Y_{i,j}$ carries information about only one grid point, $X_{i,j}$.

We note that our setup is different from the problems usually considered in function-space MCMC, where the increasing dimension is caused by refining the spatial discretization, while keeping the size of the domain and the number of observations constant. In fact, we consider the opposite scenario: the number of observations and the size of the domain increase, but the spatial discretization remains unchanged. Our setup is also slightly different from that considered in [24], where the means of the Poisson distributions are $\exp([X]_{i,j})/L^2$ (in our notation) and where a different covariance kernel is used to define the prior. The latter is minor. We do not scale the mean values with domain size, L , because we want to study MALA-within-Gibbs on problems with increasing dimension while leaving all other parameters that define the problem unchanged.

The prior precision matrix is sparse, as illustrated in Figure 5.2 for the problem of size 16×16 . The prior precision matrices of the larger problems (32×32 and 64×64) have similar sparsity patterns. We now investigate whether Assumptions 3.1 (sparse conditional structure) and 3.4 (blockwise log-concavity) are satisfied in this problem. Sparse conditional structure can be verified by inspection: the chosen Gaussian prior is a Markov random field where each pixel has only four neighbors, and the likelihood is purely local, introducing no new dependencies. We can verify this structure more carefully as follows, partitioning the state \mathbf{x} based on 2D neighborhoods. Recall that the log posterior density is

$$(5.4) \quad \log \pi(\mathbf{x}|\mathbf{y}) = C - \frac{1}{2} \|\mathbf{B}^{-1/2}(\mathbf{x} - \mu\mathbf{1})\|^2 + \sum_{i,j} (Y_{i,j}X_{i,j} - \exp(X_{i,j})),$$

where C is a constant whose value is irrelevant. Fixing the block size at $q = n/m$, we

FIG. 5.2. Prior precision matrix of the 16×16 problem.

find that

$$(5.5) \quad \nabla_{\mathbf{x}_i, \mathbf{x}_j} \log \pi(\mathbf{x}|\mathbf{y}) = -[\mathbf{B}^{-1}]_{\mathbf{x}_i, \mathbf{x}_j} - \mathbf{1}_{i=j} \mathbf{D}_i,$$

where the $q \times q$ matrices $[\mathbf{B}^{-1}]_{\mathbf{x}_i, \mathbf{x}_j}$ are constructed from \mathbf{B}^{-1} based on the blocks $\mathbf{x}_i, \mathbf{x}_j$, and \mathbf{D}_i is a diagonal $q \times q$ matrix with entries being $\exp(X_{k,l})$ for each $X_{k,l}$ in \mathbf{x}_i . Due to the sparse structure of the prior precision \mathbf{B}^{-1} (see Figure 5.2), $[\mathbf{B}^{-1}]_{\mathbf{x}_i, \mathbf{x}_j}$ is zero if the blocks i and j are far from each other in the 2D domain. In this case, $\mathbf{1}_{i=j} \mathbf{D}_i = 0$, so that by (5.5), $\nabla_{\mathbf{x}_i, \mathbf{x}_j}^2 \log \pi(\mathbf{x}|\mathbf{y})$ is also zero. The problem is thus indeed characterized by sparse conditional structure.

The blockwise log-concavity, Assumption 3.4, may not be satisfied in this example. By (5.5), the Hessian is bounded above by $-\mathbf{B}^{-1}$, which suggests using

$$H_{i,i}(x) \equiv -\lambda_{\min}([\mathbf{B}^{-1}]_{\mathbf{x}_i, \mathbf{x}_i}), \quad H_{i,j}(x) \equiv \|\mathbf{B}^{-1}\|_{\mathbf{x}_i, \mathbf{x}_j}, \quad i, j = 1, \dots, m,$$

where $[\mathbf{B}^{-1}]_{\mathbf{x}_i, \mathbf{x}_i}$ and $[\mathbf{B}^{-1}]_{\mathbf{x}_i, \mathbf{x}_j}$ are constructed from \mathbf{B}^{-1} , based on the blocks with indices i and j . With this choice, Assumption 3.4 requires that

$$(5.6) \quad c := \lambda_{\min}(-\mathbf{H}) > 0.$$

With this choice of \mathbf{H} and with the length scales $l_s = 2$, $l_t = 4$, the condition in (5.6) is not satisfied, suggesting that the problem is not blockwise log-concave.

5.3.2. MCMC samplers. We apply simplified manifold MALA (MMALA) [24] and MMALA-within-Gibbs to draw samples from the posterior distributions of the 16×16 , 32×32 , and 64×64 problems. We also apply pCN, as an example of a function-space MCMC scheme, to illustrate that function-space MCMC is not dimension-independent when some of its underlying assumptions are not met.

The MMALA proposal is

$$\tilde{\mathbf{x}}^k = \mathbf{x}^k + \tau \mathbf{M} \nabla \log p(\mathbf{x}^k | \mathbf{y}) + \sqrt{2\tau} \mathbf{M}^{1/2} \xi^{k+1}, \quad \xi^{k+1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}),$$

where the choice $\mathbf{M} = \mathbf{A} + \mathbf{B}^{-1}$ turns MALA into simplified manifold MALA. The matrix \mathbf{A} is diagonal and the i th diagonal element is $[\mathbf{A}]_{i,i} = \exp(\mu + [\mathbf{B}]_{i,i})$; see [24]. We implement MMALA-within-Gibbs using blocks of size $q = d \times d$ and consider $d = 8, 16, 32, 64$. We emphasize that MMALA-within-Gibbs with a single block, covering the entire domain, is equivalent to the MMALA sampler. For example, if $L = 64$ and $d = 64$, the sampler does not use partial updating and we recover the usual MMALA; with $L = 64$ and $d = 16$, we divide the domain into 16 blocks, each




FIG. 5.3. Illustration of results obtained by 10^4 samples of a MMALA-within-Gibbs sampler with $d = 8$ and step size $\tau = 0.5$. Top row: Posterior mean (left) and observations $Y_{i,j}$ (right). Bottom row: Posterior variance at each grid point (left) and observations corresponding to posterior mean (right).

of size 16×16 . The blocks define a neighborhood of components $X_{i,j}$ of size $d \times d$ on the 2D domain and are ordered left-to-right and top-to-bottom.

All samplers are initialized at the maximum a posteriori point (MAP), which we find by solving the optimization problem

$$\min_{\mathbf{x}} -\log \pi(\mathbf{x}|\mathbf{y}),$$

using a Gauss–Newton method. We consider various step sizes τ , and, for each one, we run pCN to generate 10^5 samples and MMALA or MMALA-within-Gibbs to generate 10^4 samples. We then compute the integrated auto correlation time (IACT) of each pixel using the techniques described in [49]. Note that we use all samples (no burn-in) to compute the average acceptance ratios and IACT. We inspected some of the chains and could not identify an apparent transient phase, likely because our initialization point makes the transients negligible.

5.3.3. MCMC results. Results of an MMALA-within-Gibbs sampler with $d = 8$ and step size $\tau = 0.5$ are shown in Figure 5.3. The panels in the top row show the posterior mean (average of all MCMC samples) and the observations $Y_{i,j}$ (on a log-scale). The panels in the bottom row show the posterior variance at each grid point and the observations (on a log-scale) corresponding to the posterior mean. We note a good agreement between the posterior mean and the true field (see Figure 5.1), as well as a good agreement between the observations and the reconstructed observations.

Our tuning of the step size is illustrated in Figure 5.4, where the average acceptance ratio of MMALA and MMALA-within-Gibbs is plotted as a function of the step sizes we tried for the problem with $L = 64$. The results are qualitatively similar for




FIG. 5.4. *Left: Average acceptance ratio of MMALA and MMALA-within-Gibbs as a function of the step size for the 64×64 problem. Right: IACT of MMALA and MMALA-within-Gibbs as a function of the step size for the 64×64 problem.*

the problems with $L = 16$ and $L = 32$. We see that the acceptance ratio decreases with step size, but for any fixed step size τ , the acceptance ratios of the within-Gibbs samplers increase when the block sizes are decreased. The reason is that the partial updating of the MMALA-within-Gibbs sampler results in large acceptance ratios for large step sizes independently of the dimension of the problem. Figure 5.4 also shows IACT as a function of the step size. We note that, for a fixed step size, IACT increases with the block size and that the step size that minimizes IACT increases as the block size decreases. Again, the reason is that the partial updating strategy of MMALA-within-Gibbs allows larger steps for smaller blocks, which decreases IACT and accelerates the mixing of the Markov chain.

IACT, averaged over all grid points, and the average acceptance probabilities (averages taken over the MCMC moves) are listed in Table 2. Here, we list results where we fixed the step size for each considered block size to make the resulting average acceptance probabilities comparable. An even better agreement between the acceptance probabilities at each block would require a more careful tuning of the step size, but the step size tuning we carried out is sufficient to make our points and to illustrate the relevant characteristics of the samplers.

TABLE 2

IACT/average acceptance probability/ τ of pCN, MMALA, and MMALA-within-Gibbs for three problems with increasing domain size L (and thus increasing dimension). Note that MMALA-within-Gibbs with block size equal to the domain size corresponds to MMALA.

L	pCN ($N_e = 10^5$)	MMALA-within-Gibbs ($N_e = 10^4$)			
		64×64 blocks	32×32 blocks	16×16 blocks	8×8 blocks
16	4626/0.18/0.002	-	-	342/0.95/0.2	204/0.93/0.5
32	5363/0.26/0.002	-	437/0.75/0.1	330/0.73/0.2	203/0.78/0.5
64	6884/0.21/0.001	627/0.48/0.05	529/0.61/0.1	394/0.75/0.2	249/0.80/0.5

For a fixed block size, MMALA-within-Gibbs yields the same IACT independently of the domain size (dimension). For example, with blocks of size $d = 16$, IACT of the 4096-dimensional problem is similar to the IACT of the 256- or 1024-dimensional problem. Moreover, the step size and corresponding acceptance ratios seem to be independent of the overall problem dimension. The numerical experiments thus corroborate our theoretical results on dimension-independent convergence of MALA-within-Gibbs, even when the assumption of blockwise log-concavity, required for our proofs, is not satisfied with our choices of block size.

Comparing MMALA-within-Gibbs with MMALA and pCN, we note that the IACTs of pCN and of MMALA with n -dimensional updates increase with dimension. In the case of pCN, this behavior is not surprising because the way that dimension increases in this sequence of problems breaks some of the assumptions that are required for the dimension independence of pCN. For MMALA, this behavior is also to be expected since MMALA in itself is not dimension-independent. Using the within-Gibbs framework turns MMALA into a dimension-independent sampling algorithm.

The dimension-independent convergence of MMALA-within-Gibbs, however, does not necessarily imply that MMALA-within-Gibbs is the most efficient sampler for this problem. Using the cost-per-effective-sample in (5.3), it is evident that MMALA is more efficient than MMALA-within-Gibbs (at the block sizes we consider). The cost estimate (5.3), however, assumes that a full likelihood evaluation is required for each proposed sample of MMALA-within-Gibbs, which is a conservative estimate. One can easily envision making use of the problem structure during likelihood evaluations in each block. For example, evaluation of the prior term in (5.4) in each block does not require computing the full matrix-vector product $\mathbf{B}^{-1/2}(\mathbf{x} - \mu\mathbf{1})$. One can speed up the computations by only updating the relevant components that are modified in the current block. We did not pursue such ideas because this problem is relatively simple and because our main goal is to demonstrate that MMALA-within-Gibbs can exhibit dimension independence.

5.4. Numerical illustration 2: Inverse problems with an elliptic PDE.

We consider the PDE

$$-\nabla \cdot (\kappa \nabla u) = g$$

on a square domain $(s, t) \in [0, 1]^2$ with Dirichlet boundary conditions; here u represents a pressure field, and g is a given source term, which consists of four delta functions (sources) at four locations in the domain. Details on the boundary conditions and source term are given in [35]. The quantity $\kappa > 0$ represents the permeability of the medium; we use a log-normal prior for the permeability to enforce the non-negativity constraint. Thus, $K = \log \kappa$ is a Gaussian random field. We set its mean to be zero and employ the covariance kernel

$$k(s_1, t_1; s_2, t_2) = \exp\left(-\frac{(s_1 - s_2)^2}{2l_s^2} - \frac{(t_1 - t_2)^2}{2l_t^2}\right),$$

where (s_1, t_1) and (s_2, t_2) are two points in the square domain and l_s and l_t are correlation length scales. Our goal is to estimate the permeability given 128 noisy observations of the pressure u in the center of the domain. This problem setup is also described in [35]. The inverse problems we consider here differ from those in [35] only in the correlation lengths of the prior, which do not affect the numerics of the PDE solve, the gradient computations, or the observation and forcing network. We thus refer the reader to [35] for the details of the numerical solution of the PDE, and in particular to Figure 2 of [35] for descriptions of the locations of the forcing terms.

5.4.1. Discretization and problem setups. For computations, we discretize the PDE using a standard finite element method with a uniform grid of 16×16 points (see [35] for details of the discretization). The discretization leads to the algebraic equation

$$(5.7) \quad \mathbf{A}(\hat{\kappa})\hat{\mathbf{u}} = \hat{\mathbf{g}},$$

where the hat over variables denotes discretized quantities, i.e., $\hat{\mathbf{k}}$, $\hat{\mathbf{u}}$, and $\hat{\mathbf{g}}$ are vectors of size $N_u = 256$, and \mathbf{A} is a 256×256 matrix that depends on the permeability $\hat{\mathbf{k}}$. We will be computing with the discretized PDE from now on, and, for that reason, we drop the hats above all variables. The pressure observations are modeled by the equation

$$(5.8) \quad \mathbf{y} = \mathbf{H}\mathbf{u} + \boldsymbol{\eta}, \quad \boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}),$$

where \mathbf{H} is an $N_y \times N_u$ matrix that has exactly one 1 in each row and picks out every other component of \mathbf{u} . The observation noise covariance is set to be $\mathbf{R} = 0.1^2 \mathbf{I}$.

After discretization, the log-permeability \mathbf{k} is finite-dimensional and its prior distribution is the finite-dimensional Gaussian $\mathcal{N}(\mathbf{0}, \mathbf{B})$. Due to the squared exponential covariance model, \mathbf{B} can be well approximated by a low-rank matrix, i.e.,

$$(5.9) \quad \mathbf{B} \approx \mathbf{U}_\theta \mathbf{L}_\theta \mathbf{U}_\theta^T,$$

where \mathbf{L}_θ is an $N_\theta \times N_\theta$ diagonal matrix whose diagonal elements are the $N_\theta < N_u$ largest eigenvalues of \mathbf{B} (see [35] for details).

The Gaussian prior for the log-permeability and the likelihood in (5.8) define the posterior distribution $\pi(\mathbf{k}|\mathbf{y}) \propto \pi_0(\mathbf{k})\pi_l(\mathbf{y}|\mathbf{k})$ for the log-permeability:

$$\pi(\mathbf{k}|\mathbf{y}) \propto \exp\left(-\frac{1}{2}\|\mathbf{B}^{-1/2}\mathbf{k}\|^2 - \frac{1}{2}\|\mathbf{R}^{-1/2}(\mathcal{M}(\mathbf{k}) - \mathbf{y})\|^2\right);$$

here \mathcal{M} maps the log-permeability to the pressure at observation locations, i.e., $\mathcal{M}(\mathbf{k}) = \mathbf{H}\mathbf{u}(\exp(\mathbf{k}))$, with the \mathbf{u} being the solution to the discretized PDE (5.7).

Since symmetric positive semidefinite matrices can always be diagonalized by a coordinate transformation, we consider the change of variables

$$(5.10) \quad \boldsymbol{\theta} = \mathbf{L}_\theta^{-1/2} \mathbf{U}_\theta^T \mathbf{k} \approx \mathbf{B}^{-1/2} \mathbf{k},$$

which leads to the posterior distribution

$$(5.11) \quad \pi(\boldsymbol{\theta}|\mathbf{y}) \propto \exp\left(-\frac{1}{2}\|\boldsymbol{\theta}\|^2 - \frac{1}{2}\|\mathbf{R}^{-1/2}(\mathcal{M}(\mathbf{k}(\boldsymbol{\theta})) - \mathbf{y})\|^2\right).$$

Below, we use MCMC samplers to draw samples from the posterior distribution of $\boldsymbol{\theta}$. The corresponding (log-)permeabilities are computed from posterior samples of $\boldsymbol{\theta}$ via the inverse of the transformation (5.10).

We consider two problem setups, which differ in the correlation lengths of the log-normal prior. The correlation lengths define the dimension of $\boldsymbol{\theta}$ in that the latter is chosen to retain 95% of the integrated prior variance. Specifically, if the correlation lengths are short compared to the $[0, 1] \times [0, 1]$ domain, then the dimension of $\boldsymbol{\theta}$ is large; if the correlation lengths are large, the dimension of $\boldsymbol{\theta}$ is small. The correlation length scales and implied dimensions of $\boldsymbol{\theta}$ of Setups 1 and 2 are summarized in Table 3.

TABLE 3
Correlation lengths and reduced dimensions for Setups 1 and 2.

	l_s	l_t	N_θ
Setup 1	0.4	0.8	30
Setup 2	0.2	0.1	136




FIG. 5.5. Left: Eigenvalues of the prior covariance matrix for Setup 1 (blue) and Setup 2 (red). Center: True log-permeability of Setup 1. Right: True log-permeability of Setup 2. (Color available online.)

We illustrate the decay of the prior covariance eigenvalues and the true log-permeabilities of Setups 1 and 2 in Figure 5.5. Specifically, we note that the eigenvalues decay more quickly for Setup 1 than for Setup 2 because Setup 1 is characterized by larger correlation length scales than Setup 2. The true log-permeabilities of Setups 1 and 2 are random draws from the prior and are shown in the right panels of Figure 5.5. There is more small-scale structure in the log-permeability of Setup 2 than in Setup 1, again due to the shorter prior correlation length scales.

Setup 2 is intended to have a higher dimension than Setup 1, not just in the apparent dimension of θ but also in the sense of the prior-to-posterior update and hence the influence of the data (i.e., the effective dimension, as defined in [1]). We achieve this by keeping the domain size fixed while decreasing the correlation lengths, which effectively increases the number of degrees of freedom in the unknown. In the previous log-Cox example, we imposed a similar growth by keeping the correlation lengths fixed but increasing the domain size. Note that, however, in the previous example we also increased the number of observations with the dimension (size of the domain), while in this example, we keep the number of observations fixed. This is a minor issue because in Setup 1, due to the large prior correlation lengths, many of the observations are strongly dependent (i.e., in the prior predictive $\pi(\mathbf{y})$). When the correlation lengths decrease in Setup 2, the number of effectively independent observations increases, and thus the relative influence of the likelihood, and hence the effective dimension, increase as well.

We emphasize that the way dimension increases in these two elliptic PDE inverse problem setups is different from what is usually considered in function-space MCMC. Dimension independence of function-space MCMC requires that the dimension increase due to refinement of a discretization, while keeping all other aspects of the problem setup (e.g., the prior and forward operator, assuming a consistent discretization scheme) fixed. As in the previous example, we in fact do the opposite and keep the discretization fixed but decrease the prior correlation length scales. For this reason, one cannot expect that pCN, or other function-space MCMC techniques, will exhibit dimension independence in the problem setups we consider.

5.4.2. Sparse conditional structure and blockwise log-concavity. The theory we created for the dimension-independent convergence of the MALA-within-Gibbs sampler relies on assumptions of sparse conditional structure and blockwise log-concavity. With our choice of prior, the problem does not have sparse conditional structure in (s, t) -coordinates. Yet the coordinate transformation (5.10) produces a sparse conditional structure—indeed complete independence—in the prior for the

$\boldsymbol{\theta}$ -coordinates, which correspond to discretized Karhunen–Loève (KL) modes. Conditioning on the observations, however, can introduce dependence among the smoother KL modes because an observation at a given (s, t) -location is in principle influenced by all of the modes—due to the nature of the elliptic operator and the KL modes’ global support. Conversely, changes in one KL mode can affect the solution everywhere in (s, t) -coordinates. Nonetheless, our experiments, along with various other experiments with this problem found in the literature, suggest that this dependence is weak and that the problem thus has an approximate sparse conditional structure in the KL modes. The assumption of blockwise log-concavity is difficult to verify in this example, in either the $\boldsymbol{\theta}$ - or (s, t) -coordinates. The reason is that the discretization of the PDE, e.g., in (5.11), makes computations difficult because we do not have second-order adjoints to compute the required Hessian.

5.4.3. MCMC samplers. We use pCN, MALA, and MALA-within-Gibbs to draw samples from the posterior distribution (5.11). Again, we emphasize that MALA-within-Gibbs with a sufficiently large block size is the same as the usual MALA without partial updating. The pCN proposal is

$$\tilde{\boldsymbol{\theta}}^{k+1} = \sqrt{1 - \beta^2} \boldsymbol{\theta}^k + \beta \boldsymbol{\xi}^{k+1},$$

where $\boldsymbol{\theta}^k$ is the current state of the MCMC and where $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{N_\theta})$, \mathbf{I}_{N_θ} being the identity matrix of order N_θ . The proposed $\tilde{\boldsymbol{\theta}}^{k+1}$ is accepted with probability

$$\alpha_{\text{pCN}} = 1 \wedge \exp\left(\frac{1}{2} \|\mathbf{R}^{-1/2}(\mathcal{M}(\mathbf{k}(\boldsymbol{\theta}^k)) - \mathbf{y})\|^2 - \frac{1}{2} \|\mathbf{R}^{-1/2}(\mathcal{M}(\mathbf{k}(\tilde{\boldsymbol{\theta}}^{k+1})) - \mathbf{y})\|^2\right),$$

where $1 \wedge x$ denotes $\min\{1, x\}$. We initialize the pCN chain at the MAP, which we find by quasi-Newton optimization (MATLAB `fminunc`) of the cost function

$$(5.12) \quad F(\boldsymbol{\theta}) = \log(\pi(\boldsymbol{\theta}|\mathbf{y})) = -\frac{1}{2} \|\boldsymbol{\theta}\|^2 - \frac{1}{2} \|\mathbf{R}^{-1/2}(\mathcal{M}(\mathbf{k}(\boldsymbol{\theta})) - \mathbf{y})\|^2 + C,$$

where C is a constant that is irrelevant. We tune the parameter β to obtain minimal IACT. As above, IACT is computed using the techniques and definitions of [49].

The MALA proposal for this problem is

$$\tilde{\boldsymbol{\theta}}^{k+1} = \boldsymbol{\theta}^k - \tau \mathbf{J}^{-1} \nabla_{\boldsymbol{\theta}} F(\boldsymbol{\theta}^k) + \sqrt{2\tau} \mathbf{J}^{-1/2} \boldsymbol{\xi}^{k+1},$$

where $F(\boldsymbol{\theta})$ is as in (5.12) and where \mathbf{J} is the Hessian of F at the MAP. As with pCN, we initialize MALA at the MAP and tune the step size τ of MALA to find a minimal IACT. As in the previous example, we use all samples for our computations (no burn-in).

MALA-within-Gibbs requires that we partition $\boldsymbol{\theta}$ into blocks. Above, we argued that this problem has an approximate sparse conditional structure in the $\boldsymbol{\theta}$ coordinates. For this reason, we use partitions of $\boldsymbol{\theta}$ that group consecutive elements of $\boldsymbol{\theta}$ together. Below, we consider several block sizes, and for each one, we initialize MALA-within-Gibbs at the MAP (as before) and tune the step size to achieve a minimal IACT.

5.4.4. MCMC results. Typical results one can obtain via MCMC are shown in Figure 5.6, where we plot an approximation of the posterior mean of the log-permeability and the approximate posterior standard deviations (on the grid) computed via MALA-within-Gibbs. We obtain an approximate posterior mean of the

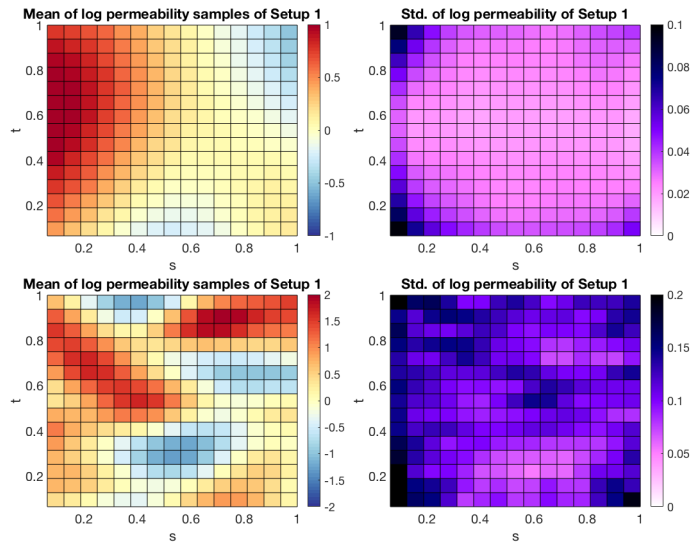


FIG. 5.6. Top row: Approximate posterior mean (left) and approximate standard deviation (right) computed from $N_e = 10^4$ MALA-within-Gibbs samples with block size $q = 1$ for Setup 1. Bottom row: Approximate posterior mean (left) and approximate standard deviation (right) computed from $N_e = 10^3$ MALA-within-Gibbs samples with block size $q = 1$ for Setup 2.

TABLE 4
Summary of simulation results of Setups 1 and 2.

	Method	Length of chain	IACT	Acc. ratio	Step
Setup 1	MALA-within-Gibbs, $q = 1$	10^4	25	0.43	0.5
	MALA-within-Gibbs, $q = 15$	10^4	141	0.22	0.05
	MALA/MALA-within-Gibbs, $q = 30$	10^5	246	0.44	0.01
	pCN	10^6	6,102	0.24	0.01
Setup 2	MALA-within-Gibbs, $q = 1$	10^3	20	0.38	0.500
	MALA-within-Gibbs, $q = 68$	10^4	367	0.23	0.010
	MALA/MALA-within-Gibbs, $q = 136$	10^5	923	0.23	0.005
	pCN	10^6	28,015	0.45	0.010

log-permeability, \mathbf{k} , from the posterior mean of $\boldsymbol{\theta}$, by mapping $\boldsymbol{\theta}$ to \mathbf{k} via the inverse of (5.10). The approximate posterior mean of \mathbf{k} should be compared to the true log-permeability in Figure 5.5.

A summary of the numerical experiments we performed is provided in Table 4. The table lists IACT, step sizes, and average acceptance ratios for the various MCMC samplers. The numbers shown are tuned, in the sense that we only show results for the step size that leads to minimal IACT (over all step sizes we tried).

We note that the IACT of MALA-within-Gibbs with block size one is nearly identical for the two problem setups, indicating that the dimension independence results we obtained under more restrictive assumptions may indeed hold in practice. As in the previous example, we also note that the step size τ that leads to minimal IACT decreases as we increase the size of the blocks of MALA-within-Gibbs. This is further illustrated in Figure 5.7, where we plot the average acceptance ratio as a function of the step size for MALA-within-Gibbs (several block sizes) and MALA. As in the previous example, we note that for a given fixed step size, the average acceptance ratio increases as we decrease the block size. The figure also shows IACT




FIG. 5.7. *Left: Average acceptance ratio of MALA-within-Gibbs and MALA, as a function of step size, for Setup 1. Right: Average IACT of MALA-within-Gibbs and MALA, as a function of step size, for Setup 1.*

as a function of the step size for the various samplers, with optimal step sizes clearly visible. We note, as before, that the partial updating of MALA-within-Gibbs pushes the step size that minimizes IACT towards larger values.

We also note that the IACT of pCN is larger than that of MALA, and that the IACTs of pCN and MALA increase with the dimension of θ . The reason that the IACT of pCN is larger than that of MALA may be that MALA makes use of gradients of the target, while the pCN proposal does not directly exploit gradient or likelihood information. Increase of the IACT of pCN with dimension is due to the way in which dimension increases when going from Setup 1 to Setup 2. As in the previous example (and as explained above), some of the fundamental assumptions that are needed for dimension independence of pCN are not satisfied when transitioning from Setup 1 to Setup 2. For that reason, one cannot expect pCN to be dimension-independent in the scenario we consider here.

Finally, we note that the acceptance rate of pCN that minimizes IACT (over the step sizes we tried) in Setup 2 is substantially larger than optimal acceptance rates of RWM (45% rather than about 20%). This is in line with recent numerical experiments and analyses which suggest that for Gaussian targets, a good acceptance rate may be near 50%. We make no claim, however, that our tuning of pCN is perfect. We considered a wide range of step sizes and ran pCN chains of length 10^6 for each choice. One could possibly achieve a slightly better IACT with further tuning, but nonetheless the IACT of pCN can be expected to be significantly larger than that of MALA or MALA-within-Gibbs. Moreover, given the overall chain length of only 10^6 , the estimated IACT of 28015 for pCN may not be entirely precise, but all of our numerical experiments indicate that it is in any case very large.

Recall that a small and dimension-independent IACT of MALA-within-Gibbs does not necessarily imply that the algorithm is a computationally efficient sampler; generating one sample requires several likelihood evaluations due to the partial updating strategy, as described above. Estimating the cost per effective sample by (5.3), we see that MALA-within-Gibbs is not an efficient sampler for Setup 1, but MALA-within-Gibbs with $q = 68 \times 68$ (leading to two blocks and two likelihood evaluations per sample) is indeed the most effective sampler for Setup 2. This further illustrates that there is a trade-off between the need to reduce IACT by using partial updating and the need to keep the cost-per-sample, which we take to be proportional to the number of blocks, reasonable. In the future, such issues may be addressed by incor-

porating the partial updating into local likelihood evaluations which do not require solving the full PDE, similar to the localization of forward dynamics discussed in [32], but such issues are beyond the scope of this paper.

We also note that our numerical experiments are limited in the sense that we only considered pCN, MALA, and MALA-within-Gibbs. Other samplers may turn out to be more practical than MALA-within-Gibbs. Specifically, note that the pressure field is relatively well observed, which implies that the posterior differs strongly from the prior in many directions (high effective dimension). This explains, at least in part, why we observe such large IACT for pCN. Other anisotropic samplers that are modifications of pCN, e.g., DILI [18], pCNL [17], and generalized pCN [44], might be effective in this problem. Our goal, however, is not to find the most appropriate sampler for this Bayesian inverse problem but rather to use this example to demonstrate some of the practical and theoretical aspects of the MALA-within-Gibbs sampler.

6. Conclusion. Markov chain Monte Carlo (MCMC) samplers are used to draw samples from a given target probability distribution in a wide array of applications. We have discussed the numerical efficiency of a particular sampler, the MALA-within-Gibbs sampler, when the target distribution exhibits a particular sparse conditional structure. In simple terms, the latter is just a structured conditional independence relationship, or block conditional independence relationship, among the variables of interest. For Gaussians, sparse conditional structure is equivalent to a (block-)sparse precision matrix. MALA-within-Gibbs samplers are natural tools to make effective use of sparse conditional structure for numerical efficiency via a suitable partial updating. We have shown that the acceptance ratio and step size of MALA-within-Gibbs are independent of the overall dimension of the problem if the partial updating is chosen to be in line with the sparse conditional structure of the target distribution. Under additional assumptions of blockwise log-concavity, we could prove that the convergence rate of MALA-within-Gibbs is independent of dimension. This suggests that MALA-within-Gibbs can be an effective sampler for high-dimensional problems.

We have investigated the applicability of MALA-within-Gibbs in the context of Bayesian inverse problems, where we expect to encounter sparse conditional structure. In many Bayesian inverse problems, we expect that sparse conditional structure can be anticipated based on the prior distribution and the locality of the likelihood, in appropriate coordinates. Numerical experiments on two well-known test problems suggest that our theoretical results are indeed indicative of what to expect in practice, where the required assumptions may only hold approximately. For example, in both numerical examples, we could show that measures of performance of the MALA-within-Gibbs sampler, e.g., integrated autocorrelation time (IACT), step size, and acceptance ratio, are indeed independent of the overall dimension of the problem. Nonetheless, the actual computational cost of MALA-within-Gibbs is dependent on the problem dimension because the partial updating requires repeated likelihood evaluations (which are costly) per sample. Our numerical experiments suggest that there is a trade-off between additional computational costs due to the partial updating and the increase in computational cost due to larger IACT or decreasing step size, without partial updating. To keep, for example, IACT small, a large number of partial updates should be used, but this in turn requires several likelihood evaluations per sample. This trade-off may not always be easy to resolve in practice. We have provided examples in which MALA-within-Gibbs leads to significant gains in the computational cost per (effective) sample, but we have also encountered examples in which a global update is, ultimately, the right choice.

REFERENCES

- [1] S. AGAPIOU, O. PAPANILOPOULOS, D. SANZ-ALONSO, AND A. STUART, *Importance sampling: Computational complexity and intrinsic dimension*, *Stat. Sci.*, 32 (2017), pp. 405–431.
- [2] M. ASCH, M. BOCQUET, AND M. NODET, *Data Assimilation: Methods, Algorithms and Applications*, *Fund. Algorithms* 11, SIAM, Philadelphia, 2016, <https://doi.org/10.1137/1.9781611974546>.
- [3] J. BARDSLEY, T. CUI, Y. MARZOUK, AND Z. WANG, *Scalable Optimization-Based Sampling on Function Space*, preprint, <https://arxiv.org/abs/1903.00870>, 2019.
- [4] J. BEAR, *Modeling Groundwater Flow and Pollution*, Kluwer, Dordrecht, The Netherlands, 1990.
- [5] M. BÉDARD, *Hierarchical models: Local proposal variances for RWM-within-Gibbs and MALA-within-Gibbs*, *Comput. Stat. Data Anal.*, 109 (2017), pp. 231–246.
- [6] A. BESKOS, *A stable manifold MCMC method for high dimensions*, *Stat. Probab. Lett.*, 90 (2014), pp. 46–52.
- [7] A. BESKOS, N. PILLAI, G. O. ROBERTS, J. M. SANZ-SERNA, AND A. M. STUART, *Optimal tuning of the hybrid Monte Carlo algorithm*, *Bernoulli*, 19 (2013), pp. 1501–1534.
- [8] A. BESKOS, G. O. ROBERTS, AND A. M. STUART, *Optimal scalings for local Metropolis-Hastings chains on nonproduct targets in high dimensions*, *Ann. Appl. Probab.*, 19 (2009), pp. 863–898.
- [9] T. BUI-THANH, O. GHATTAS, J. MARTIN, AND G. STADLER, *A computational framework for infinite-dimensional Bayesian inverse problems. Part I: The linearized case, with application to global seismic inversion*, *SIAM J. Sci. Comput.*, 35 (2013), pp. A2494–A2523, <https://doi.org/10.1137/12089586X>.
- [10] N. CHEN AND A. J. MAJDA, *Filtering nonlinear turbulent dynamical systems through conditional Gaussian statistics*, *Mon. Weather Rev.*, 144 (2016), pp. 4885–4917.
- [11] N. CHEN AND A. J. MAJDA, *Conditional Gaussian systems for multiscale nonlinear stochastic systems: Prediction, state estimation and uncertainty quantification*, *Entropy*, 20 (2018), 509.
- [12] N. CHEN, A. J. MAJDA, AND X. T. TONG, *Spatial Localization for Nonlinear Dynamical Stochastic Models for Excitable Media*, preprint, <https://arxiv.org/abs/1901.07318>, 2019.
- [13] N. CHEN, A. J. MAJDA, AND X. T. TONG, *Rigorous analysis for efficient statistically accurate algorithms for solving Fokker-Planck equations in large dimensions*, *SIAM/ASA J. Uncertain. Quantif.*, 6 (2018), pp. 1198–1223, <https://doi.org/10.1137/17M1142004>.
- [14] Y. CHEN, D. KEYES, K. J. H. LAW, AND H. LTAIEF, *Accelerated dimension-independent adaptive Metropolis*, *SIAM J. Sci. Comput.*, 38 (2016), pp. S539–S565, <https://doi.org/10.1137/15M1026432>.
- [15] A. CHORIN AND O. HALD, *Stochastic Tools in Mathematics and Science*, 3rd ed., Springer, New York, 2013.
- [16] O. F. CHRISTENSEN, G. O. ROBERTS, AND J. S. ROSENTHAL, *Scaling limits for the transient phase of local Metropolis-Hastings algorithms*, *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67 (2005), pp. 253–268.
- [17] S. L. COTTER, G. O. ROBERTS, A. M. STUART, AND D. WHITE, *MCMC methods for functions: Modifying old algorithms to make them faster*, *Stat. Sci.*, 28 (2013), pp. 424–446.
- [18] T. CUI, K. J. H. LAW, AND Y. M. MARZOUK, *Dimension-independent likelihood-informed MCMC*, *J. Comput. Phys.*, 304 (2016), pp. 109–137.
- [19] T. CUI, Y. M. MARZOUK, AND K. WILLCOX, *Scalable posterior approximations for large-scale Bayesian inverse problems via likelihood-informed parameter and state reduction*, *J. Comput. Phys.*, 315 (2016), pp. 363–387.
- [20] M. DASHTI AND A. M. STUART, *Uncertainty quantification and weak approximation of an elliptic inverse problem*, *SIAM J. Numer. Anal.*, 49 (2011), pp. 2524–2542, <https://doi.org/10.1137/100814664>.
- [21] R. G. EDWARDS, J. GOODMAN, AND A. D. SOKAL, *Multi-grid Monte Carlo (II). Two-dimensional xy model*, *Nuclear Phys. B*, 354 (1991), pp. 289–327.
- [22] C. FOX AND R. A. NORTON, *Fast sampling in a linear-Gaussian inverse problem*, *SIAM/ASA J. Uncertain. Quantif.*, 4 (2016), pp. 1191–1218, <https://doi.org/10.1137/15M1029527>.
- [23] C. FOX AND A. PARKER, *Accelerated Gibbs sampling of normal distributions using matrix splittings and polynomials*, *Bernoulli*, 23 (2017), pp. 3711–3743.
- [24] M. GIROLAMI AND B. CALDERHEAD, *Riemann manifold Langevin and Hamiltonian Monte Carlo methods*, *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 73 (2011), pp. 123–214.
- [25] J. GOODMAN AND N. MADRAS, *Random-walk interpretations of classical iteration methods*, *Linear Algebra Appl.*, 216 (1995), pp. 61–79.

- [26] J. GOODMAN AND A. D. SOKAL, *Multigrid Monte Carlo method for lattice field theories*, Phys. Rev. Lett., 56 (1986), pp. 1015–1018.
- [27] J. GOODMAN AND A. D. SOKAL, *Multigrid Monte Carlo method. conceptual foundations*, Phys. Rev. Lett., 40 (1989), pp. 2035–2071.
- [28] M. HAIRER, A. STUART, AND S. VOLLMER, *Spectral gaps for a Metropolis–Hastings algorithm in infinite dimensions*, Ann. Appl. Probab., 24 (2014), pp. 2455–2490.
- [29] A. A. JOHNSON, G. L. JONES, AND R. C. NEATH, *Component-wise Markov chain Monte Carlo: Uniform and geometric ergodicity under mixing and composition*, Stat. Sci., 28 (2013), pp. 360–375.
- [30] S. LAN, *Adaptive dimension reduction to accelerate infinite-dimensional geometric Markov chain Monte Carlo*, J. Comput. Phys., 392 (2019), pp. 71–95.
- [31] S. L. LAURITZEN, *Graphical Models*, Oxford Stat. Sci. Ser. 17, The Clarendon Press, Oxford University Press, New York, 1996.
- [32] Q. LIU AND X. T. TONG, *Accelerating Metropolis-within-Gibbs sampler with localized computations of differential equations*, Stat. Comput., 30 (2020), pp. 1037–1056.
- [33] J. MÖLLER, A. R. SYVERSVEEN, AND R. P. WAAGEPETERSEN, *Log Gaussian Cox processes*, Scand. J. Statist., 25 (1998), pp. 451–482.
- [34] M. MORZFELD, X. T. TONG, AND Y. M. MARZOUK, *Localization for MCMC: Sampling high-dimensional posterior distributions with local structure*, J. Comput. Phys., 310 (2019), pp. 1–28.
- [35] M. MORZFELD, X. TU, J. WILKENING, AND A. CHORIN, *Parameter estimation by implicit sampling*, Commun. Appl. Math. Comput. Sci., 10 (2015), pp. 205–225.
- [36] T. MOSELHY AND Y. MARZOUK, *Bayesian inference with optimal maps*, J. Comput. Phys., 231 (2012), pp. 7815–7850.
- [37] P. NEAL AND G. O. ROBERTS, *Optimal scaling for partially updating MCMC algorithms*, Ann. Appl. Probab., 16 (2006), pp. 475–515.
- [38] D. S. OLIVER, A. C. REYNOLDS, AND N. LIU, *Inverse Theory for Petroleum Reservoir Characterization and History Matching*, Cambridge University Press, Cambridge, UK, 2008.
- [39] M. OTTOBRE, N. S. PILLAI, F. J. PINSKI, AND A. M. STUART, *A function space HMC algorithm with second order Langevin diffusion limit*, Bernoulli, 22 (2016), pp. 60–106.
- [40] N. PETRA, J. MARTIN, G. STADLER, AND O. GHATTAS, *A computational framework for infinite-dimensional Bayesian inverse problems. Part II: Stochastic Newton MCMC with application to ice sheet flow inverse problems*, SIAM J. Sci. Comput., 36 (2014), pp. A1525–1555, <https://doi.org/10.1137/130934805>.
- [41] S. REICH AND C. COTTER, *Probabilistic Forecasting and Bayesian Data Assimilation*, Cambridge University Press, Cambridge, UK, 2015.
- [42] G. O. ROBERTS, A. GELMAN, AND W. R. GILKS, *Weak convergence and optimal scaling of random walk Metropolis algorithms*, Ann. Appl. Probab., 7 (1997), pp. 110–120.
- [43] G. O. ROBERTS AND J. S. ROSENTHAL, *Optimal scaling of discrete approximations to Langevin diffusions*, J. R. Stat. Soc. Ser. B Stat. Methodol., 60 (1998), pp. 255–268.
- [44] D. RUDOLF AND B. SPRUNGK, *On a generalization of the preconditioned Crank–Nicolson metropolis algorithm*, Found. Comput. Math., 18 (2018), pp. 309–343.
- [45] A. D. SOKAL, *Monte Carlo methods in statistical mechanics: Foundations and new algorithms*, in Functional Integration (Cargèse, 1996), Plenum, New York, 1997, pp. 131–192.
- [46] A. SPANTINI, D. BIGONI, AND Y. M. MARZOUK, *Inference via low-dimensional couplings*, J. Mach. Learn. Res., 19 (2018), pp. 1–71.
- [47] A. SPANTINI, A. SOLONEN, T. CUI, J. MARTIN, L. TENORIO, AND Y. M. MARZOUK, *Optimal low-rank approximations of Bayesian linear inverse problems*, SIAM J. Sci. Comput., 37 (2015), pp. A2451–A2487, <https://doi.org/10.1137/140977308>.
- [48] S. J. VOLLMER, *Dimension-independent MCMC sampling for inverse problems with non-Gaussian priors*, SIAM/ASA J. Uncertain. Quantif., 3 (2015), pp. 535–561, <https://doi.org/10.1137/130929904>.
- [49] U. WOLFF, *Monte Carlo errors with less errors*, Comput. Phys. Commun., 156 (2004), pp. 143–153.
- [50] O. ZAHM, T. CUI, K. LAW, A. SPANTINI, AND Y. MARZOUK, *Certified Dimension Reduction in Nonlinear Bayesian Inverse Problems*, preprint, <https://arxiv.org/abs/1807.03712>, 2018.