Rapidly Mixing Multiple-try Metropolis Algorithms for Model Selection Problems

Hyunwoong Chang*
Department of Statistics
Texas A&M University
hwchang@stat.tamu.edu

Changwoo J. Lee*
Department of Statistics
Texas A&M University
c.lee@stat.tamu.edu

Zhao Tang Luo
Department of Statistics
Texas A&M University
ztluo@stat.tamu.edu

Huiyan Sang
Department of Statistics
Texas A&M University
huiyan@stat.tamu.edu

Quan Zhou
Department of Statistics
Texas A&M University
quan@stat.tamu.edu

Abstract

The multiple-try Metropolis (MTM) algorithm is an extension of the Metropolis-Hastings (MH) algorithm by selecting the proposed state among multiple trials according to some weight function. Although MTM has gained great popularity owing to its faster empirical convergence and mixing than the standard MH algorithm, its theoretical mixing property is rarely studied in the literature due to its complex proposal scheme. We prove that MTM can achieve a mixing time bound smaller than that of MH by a factor of the number of trials under a general setting applicable to high-dimensional model selection problems with discrete state spaces. Our theoretical results motivate a new class of weight functions called *locally balanced weight functions* and guide the choice of the number of trials, which leads to improved performance over standard MTM algorithms. We support our theoretical results by extensive simulation studies and real data applications with several Bayesian model selection problems.

1 Introduction

The Markov chain Monte Carlo (MCMC) method has become a powerful and standard tool for Bayesian posterior computation. In particular, the Metropolis-Hastings (MH) algorithm is widely employed in many statistical and machine learning models to sample from posterior distributions with intractable normalizing constants. A prominent usage example of MH is the vast class of Bayesian model selection problems with discrete-valued high-dimensional parameters, whose posterior probabilities are infeasible to evaluate because the size of model space potentially grows (super-)exponentially with the number of parameters. Examples of Bayesian model selection problems and their applications are ubiquitous, including Bayesian variable selection (BVS) [45, and references therein], stochastic block model (SBM) [20, 34], Bayesian structure learning [10, 15, 29], change-point detection model [5, 11], spatial clustering models [23, 26, 28], and many others.

Despite the popularity of MCMC algorithms, their convergence rate analysis is often challenging. Mixing time is a key concept of interest when analyzing the convergence of MCMC, which defines the number of iterations needed to converge to a stationary distribution within a small total variation distance [25]. We say a chain is *rapidly mixing* if the mixing time grows at most polynomially in some complexity parameters. Mixing time bound has been studied extensively for problems such as card shuffling [3], random walks on groups [2], graph colorings [21], and the Ising model [14], across

*Equal contribution

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

the fields of probability, physics, and computer science [40]. Although there is a rich literature on discrete-state-space mixing time analysis in many areas, the focus has been primarily on approximate counting and random generation, but only recently has research on the MCMC convergence in the context of statistical model selection begun to emerge. [47] proved rapid mixing of an MH algorithm in a high-dimensional BVS problem, and [38] improved the bound when a "warm start" is available. Rapid mixing results of MH algorithms are established for Bayesian community detection problem with SBM [55] and for learning equivalence class in a high-dimensional Bayesian structure learning [52]. Recently, locally informed MH algorithms have been developed to incorporate local information about the target distribution. [48] proposed a locally balanced proposal and proved their optimality in terms of Peskun ordering [37] for discrete space models without providing mixing time rate analysis. [54] also proposed an informed MCMC and proved a dimension-free mixing time bound but only for a BVS problem.

The multiple-try Metropolis (MTM) algorithm [27] is an extension of the MH algorithm, which uses a weight function to choose the proposed state among randomly sampled trials from a proposal distribution. The variety of weight functions and the number of trials enrich the algorithm class but result in a much more complex transition probability than standard MH algorithms, hampering the theoretical development on the mixing time of MTM. As a result, there is a lack of theoretical guidance for practitioners on the choices of weight functions and number of trials in MTM and also the choice between MTM and other MCMC algorithms. Only very recently, [46] considered the MTM independence sampler (MTM-IS), a special case of MTM where the proposal does not depend on the current state. They conducted the convergence rate analysis of MTM-IS, which shows it is less efficient than the simpler approach of repeated Metropolised independent sampling at the same computational cost, but they did not discuss the choices of weight functions and the number of trials. As acknowledged in [46], MTM is significantly different from MTM-IS, and its convergence rate is more challenging to analyze. In this paper, we establish the mixing time bound of MTM for general model selection problems. Our main contributions are the following:

1) We show that MTM with existing popular choices of weight functions can have mixing issues and propose a new class of weight functions called *locally balanced weight functions*.
2) With locally balanced weight functions, we prove that the mixing time bound of the MTM algorithm is smaller than that of the MH algorithm by a factor of the number of trials N in a general model selection setting, under some regularity conditions including a rate condition on N.
3) We provide theoretical justification for the counterintuitive phenomenon that an increase in the number of trials N may not always lead to a proportionate improvement in mixing, and suggest a theoretically guided algorithm to choose N for practical use.
4) We validate our theoretical findings via extensive simulation studies and real data applications with various model selection problems: Bayesian variable selection, stochastic block models, spatial clustering models, and structure learning models.

2 Preliminary

2.1 Notation

For a positive integer $m \in \mathbb{N}$, let $[m] = \{1, \ldots, m\}$ and $\lfloor m \rfloor$ be the largest integer no greater than m. Let $|\cdot|$ denote the cardinality of a set. Let $\mathcal{X} = \mathcal{X}_p$ denote a finite state space where $|\mathcal{X}|$ grows in a complexity parameter $p \in \mathbb{N}$. We use $\mathbf{P} : \mathcal{X} \times \mathcal{X} \to [0, 1]$ to denote the transition probability matrix of an irreducible, aperiodic and reversible Markov chain. We denote by $\mathbf{I}$ an identity matrix. We call a probability distribution π on $\mathcal{X}$ a *stationary distribution* or a *target distribution* of a chain $\mathbf{P}$ if it satisfies $\sum_{z \in \mathcal{X}} \pi(z)\mathbf{P}(z, x) = \pi(x)$ for all $x \in \mathcal{X}$. In Bayesian inference, π is a posterior distribution. Let $\mathcal{N}\colon \mathcal{X} \to 2^{\mathcal{X}}$ be a set-valued map called a neighborhood relation, which maps a state to a set of states. We say the neighborhood relation $\mathcal{N}$ is symmetric if $x' \in \mathcal{N}(x)$ iff $x \in \mathcal{N}(x')$ for all $x \neq x'$. We define the random walk proposal matrix as $\mathbf{K}_{\mathrm{RW}}(x, y) = \mathbb{1}_{\mathcal{N}(x)}(y)/|\mathcal{N}(x)|$ where $\mathbb{1}_{\mathcal{N}(x)}$ denotes the indicator function of $\mathcal{N}(x)$. We say $\mathbf{K}$ is symmetric if $\mathbf{K}(x, y) = \mathbf{K}(y, x)$. We denote a graph by (V, E) where V is the vertex set and $E \subset V \times V$ is the set of edges. We say $\delta(x, y) = (v_0, \ldots, v_m)$ is an E-path (with length m) from x to y if $e_i = (v_{i-1}, v_i) \in E$ for $i \in [m]$, $v_0 = x, v_m = y$, and all vertices are distinct.

Algorithm 1: Multiple-try Metropolis (MTM) algorithm with proposal $\mathbf{K}_{\mathrm{RW}}$

Input: An initial state x_0, a neighborhood relation $\mathcal{N}$, the number of trials N, a weight function w defined in (1), the number of Markov chain iterations T.

for $t = 0, \ldots, T-1$ **do**

Step 1. Draw $y_1, \ldots, y_N$ uniformly at random from $\mathcal{N}(x_t)$, and compute $w(y_j \,|\, x_t)$ for $j = 1 \ldots, N$.

Step 2. Select $j \in [N]$ with probability proportional to $w(y_j \,|\, x_t)$ and define $y = y_j$.

Step 3. Sample $x_1^\star, \ldots, x_{N-1}^\star$ uniformly at random from $\mathcal{N}(y)$ and define

$$\alpha = \min\left\{1, \frac{w(y \,|\, x_t) + \sum_{l \in [N] \setminus \{j\}} w(y_l \,|\, x_t)}{w(x_t \,|\, y) + \sum_{l \in [N-1]} w(x_l^\star \,|\, y)}\right\}. \tag{2}$$

Step 4. With probability α, accept y and let $x_{t+1} = y$; otherwise, let $x_{t+1} = x_t$.

Output: A set of Markov chain samples $\{x_t\}_{t=1}^T$.

2.2 Multiple-try Metropolis algorithm

We present the MTM algorithm [27] on a discrete state space in Algorithm 1. Here N trial states $y_1, \ldots, y_N$ are first sampled from the proposal $\mathbf{K}_{\mathrm{RW}}(x, \cdot)$, and then the proposed state y is selected among the trials using the weight function

$$w(y \,|\, x) = \pi(y)\mathbf{K}_{\mathrm{RW}}(y, x)\lambda(y, x), \tag{1}$$

where $\lambda(y, x)$ can be any symmetric function that is positive whenever $\mathbf{K}_{\mathrm{RW}}(y, x) > 0$. Note that Algorithm 1 becomes the standard random walk MH algorithm when the number of trials N is one.

There have been efforts to generalize the MTM algorithm [9, 12, 30]. [35] extends the weight function class by proposing a general form of acceptance probability to achieve the reversibility of a chain when the weight function does not belong to (1). While one can easily find a weight function satisfying (1), there is little research on the most appropriate form of weight functions. [27] investigated some particular class of weight functions and claimed the performance is insensitive to the choice of weight functions within this class; $w(y \,|\, x) = \pi(y)$ (when $\mathbf{K}_{\mathrm{RW}}$ is symmetric) and $w(y \,|\, x) = \pi(y)/\mathbf{K}_{\mathrm{RW}}(x, y)$ are popular choices in this class [6, 33].

Another important component of the MTM method is the number of trials N. It has been reported with empirical evidence that increasing N does not necessarily result in a corresponding improvement in mixing under the random walk proposal [32]. They illustrate a counterintuitive phenomenon that the acceptance probability for a proposal move to higher probability regions can be extremely low. As a solution, they suggest randomly selecting N at each iteration. Despite the anecdotal observation that certain choices of N make the algorithm less efficient, the principled way to choose N is rarely discussed in the literature.

2.3 Analysis on mixing time via geometric tools

A theoretical foundation for employing MCMC methods to sample from a posterior distribution on a finite state space is the well-known convergence theorem, which states that an irreducible and aperiodic Markov chain with a stationary distribution converges to the stationary distribution regardless of the initial distribution [25, Theorem 4.9]. However, since it is impossible to simulate a chain infinitely long, our main interest is to analyze how well the chain approximates the target distribution after a certain number of steps. To this end, we introduce ϵ-mixing time $t_{\mathrm{mix}}(\epsilon) = \max_{x \in \mathcal{X}} \min\left\{t \in \mathbb{N} \colon \|\mathbf{P}_{\mathrm{lazy}}^t(x, \cdot) - \pi(\cdot)\|_{\mathrm{TV}} \leq \epsilon\right\}$, where $\|\cdot\|_{\mathrm{TV}}$ is the total variation (TV) distance, ϵ is a small positive constant, and $\mathbf{P}_{\mathrm{lazy}} = (\mathbf{P} + \mathbf{I})/2$ denotes the lazy version of $\mathbf{P}$, which is introduced merely for theoretical convenience since all eigenvalues of $\mathbf{P}_{\mathrm{lazy}}$ are non-negative (the use of lazy Markov chains is standard in mixing time analysis). In words, t_{mix} quantifies the necessary number of steps for the chain to have a small TV distance from the stationary distribution. Among several techniques for characterizing the mixing property [8, 18, 25, 42], a path method is employed to derive our main result [13, 43]. We also refer readers to [25, Section 13.4] and [42, Section 3] for more details about the path method. Intuitively, the mixing time depends on how well the states "communicate", especially among those which have high values of π. In contrast to

samplers defined on a continuous state space with a Euclidean topology, posterior landscape and modality on a finite state space are more difficult to envision since it highly depends on how one defines the neighborhood relation. Given a transition probability matrix $\mathbf{P}$, we define the neighborhood of a state x by $\mathcal{N}(x) = \{x' \in \mathcal{X} \setminus \{x\} : \mathbf{P}(x, x') > 0\}$, the set of reachable states by one-step transition from x, and we say x is a mode if $\pi(x) > \max_{x' \in \mathcal{N}(x)} \pi(x')$. This specification of the neighborhood enables us to define a graph $(\mathcal{X}, E)$ where $E = \{(x, x') \in \mathcal{X} \times \mathcal{X} : x' \in \mathcal{N}(x)\}$ is an (undirected since we have a symmetric neighborhood relation by the reversibility of $\mathbf{P}$) edge set. For each $(x, y) \in \mathcal{X} \times \mathcal{X}$ with $x \neq y$, we choose only one E-path from x to y and denote it as $\delta(x, y)$. A *path ensemble* is defined as a collection of the paths $\Delta = \{\delta(x, y) : x, y \in \mathcal{X}, x \neq y\}$. A path ensemble is a "configuration of path network" that helps quantify the quality of the communication within the state space $\mathcal{X}$. The following proposition, which serves as a base of our main result, gives a mixing time bound by means of a path ensemble.

Proposition 1 ([43])**.** *For any path ensemble Δ,*

$$t_{\mathrm{mix}}(\epsilon) \leq 2\rho(\Delta)\ell(\Delta) \left[\log(1/\epsilon) + \log \left\{ \min_{x \in \mathcal{X}} \pi(x) \right\}^{-1} \right],$$

where $\ell(\Delta) = \max_{x,y} |\delta(x, y)|$ and

$$\rho(\Delta) = \max_{(u,v) \in E} \frac{1}{\pi(u)\mathbf{P}(u, v)} \sum_{x,y : \delta(x,y) \ni (u,v)} \pi(x)\pi(y). \tag{3}$$

For completeness, we provide a proof of Proposition 1 with a detailed discussion in Appendix A.1. The quantity $\rho(\Delta)$ is called the *congestion parameter*, and $\ell(\Delta)$ is the length of the longest path in Δ. To interpret $\rho(\Delta)$ and $\ell(\Delta)$, some terminologies from transportation networks in graph theory are often useful. Consider $\pi(u)\mathbf{P}(u, v)$ as the capacity of an edge (u, v), and $\pi(x)\pi(y)$ as the unit flow of a path $\delta(x, y)$. Then $\rho(\Delta)$ quantifies the maximum loading of an edge under the configuration Δ. If one edge with a small capacity is traversed by a huge number of paths with a high unit flow, it results in a constriction. Consider $\ell(\Delta)$ as the diameter of the configuration Δ. If the diameter is extremely large, it causes inefficient communication, especially among the states lying on the margin of the configuration. If it is possible to design the path system in a way that no edge is overloaded by paths and no two states are too far apart, then the chain can have a good mixing behavior. One of the key steps in the proof of the main result is to identify a good path ensemble Δ^* (see Appendix A.3).

3 Main result

In this section, we prove the rapid mixing of the MTM algorithm for general model selection problems on a finite model space $\mathcal{X}$, where we assume that there is only one state with the highest posterior mass, denoted by x^*. Throughout the section, we may assume that $|\mathcal{X}|$ grows with the complexity parameter p (e.g., the number of variables in BVS), and the posterior distribution π and other related objects are implicitly parameterized by p.

3.1 A mixing time bound with locally balanced weight functions

While extending the possible form of weight functions has been the main focus in the literature [33, 35], the choice of weight function is rarely discussed. We introduce a class of weight functions by using a function $h \colon \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ that satisfies $h(u) = uh(1/u)$ for all $u > 0$; such a function h is called a *balancing function* [31, 48]. The class of balancing functions is very broad: we can choose an arbitrary non-negative function h on $(0, 1]$, and then $h(u)$ on $u \in (1, \infty)$ is defined by the "balancing rule", $h(u) = uh(1/u)$. Typical choices of balancing functions include $\sqrt{u}$, $\min\{1, u\}$, $\max\{1, u\}$, $u/(u + 1)$, and $u + 1$. The following proposition states that Algorithm 1 with the proposed weight function induces a reversible Markov chain with π as its stationary distribution.

Proposition 2 (Locally balanced weight functions)**.** *Suppose that a weight function is defined as*

$$w(y \,|\, x) = h\left(\frac{\pi(y)\mathbf{K}_{\mathrm{RW}}(y, x)}{\pi(x)\mathbf{K}_{\mathrm{RW}}(x, y)} \right), \tag{4}$$

where h is a balancing function. Then, a Markov chain $\mathbf{P}_{\mathrm{MTM}}$ induced by Algorithm 1 has a stationary distribution π.

The proof is deferred to Appendix A.2. We call weight functions of the form (4) *locally balanced weight functions*; though we only consider random walk proposals in this paper, $\mathbf{K}_{\mathrm{RW}}$ in Proposition 2 can be replaced by any other proposal matrix. Compared with the popular choice of weight function in the existing MTM literature such as $w(y \mid x) = \pi(y)$, the locally balanced weight function can be viewed as a transformed and scaled version with respect to the posterior probability at current state x. For example, if $h(u) = \sqrt{u}$ and $\mathbf{K}_{\mathrm{RW}}$ is symmetric, then $w(y \mid x) = \sqrt{\pi(y)/\pi(x)}$.

Before providing a rigorous mixing time bound of MTM under our new choice of $w(y \mid x)$, we give an intuitive explanation for why the existing unscaled weight function may have poor mixing properties. First, it is reasonable to assume that the true data-generating model, x^*, will have the largest posterior probability when the posterior contraction or the model selection consistency holds asymptotically. Indeed, such results have been established for a variety of model selection problems [47, 52, 55]. On the other hand, a model space $\mathcal{X}$ is often endowed with a natural neighborhood relation such that the degree of "model fit" (measured by goodness of fit and penalization on model size) between neighboring models tends to have a small difference. This implies that a small modification of a model may lead to a marginal change in the posterior probability. For example, in a Bayesian variable selection problem with the number of variables p, the model space can be represented as a set of binary vectors $\gamma \in \{0,1\}^p$ where $\gamma_j = 1$ (resp. $\gamma_j = 0$) indicates j-th variable is included (resp. excluded) in the model. Let us assume that γ^* is the true data generating model and has the highest posterior probability, and define a Hamming distance $d_{\mathrm{H}}(\gamma, \gamma') = \sum_{j=1}^{p} \mathbb{1}\{\gamma_j \neq \gamma'_j\}$ for $\gamma, \gamma' \in \{0,1\}^p$. One often defines the neighborhood $\mathcal{N}(\gamma)$ as the set of states with a small Hamming distance to γ [19, 47, 54], and models in the neighborhood tend to have similar posterior probabilities. A state with a large Hamming distance from γ^* usually has a small posterior probability because its structure is largely different from the true data-generating model. Figure 1 depicts such a relationship between posterior probabilities and structural similarities among the models. Given that γ is the current state, Step 2 of Algorithm 1 tends to propose a neighboring state γ' with a higher posterior probability, so γ' tends to be closer to γ^* in Hamming distance, i.e. $d_{\mathrm{H}}(\gamma', \gamma^*) < d_{\mathrm{H}}(\gamma, \gamma^*)$. On average, the Hamming distance between a neighboring state of γ' and γ^* will be smaller than that between a neighboring state of γ with γ^*. Therefore, the expected value of posterior probability of a neighboring state of γ' (in the denominator part in (2), shown in red in Figure 1) will be typically larger than that of a neighboring state of γ (in the numerator part in (2), shown in blue in Figure 1). If we resort to the unscaled weight function $w(\gamma' \mid \gamma) = \pi(\gamma')$, this can cause a significantly low acceptance probability, especially when the number of trials N is large (due to the law of large numbers, the acceptance ratio converges to the ratio of average posterior probabilities of the two neighborhoods); see also Appendix A.4 for a concrete example. A locally balanced weight function (4), however, is able to mitigate this counterintuitive phenomenon by "scaling" the weight function with respect to the current state. We shall demonstrate the effectiveness of locally balanced weight functions in the simulation studies. Although the scope of our study is limited to finite state spaces, employing (4) also improves the algorithm on continuous state spaces. We provide the relevant discussion in Appendix D.2.

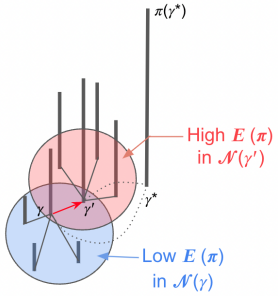

Figure 1: An illustration of the reason that the weight function $w(y \mid x) = \pi(y)$ may fail. Each vertical bar represents the posterior probability on the corresponding state, and thin lines connect neighboring states. The expected value of the posterior probability of the neighborhood of γ' (red) is typically greater than that of γ (blue).

In order for the MTM procedure to efficiently sample states from a high posterior region, it is sensible to assign a larger weight to those states with a higher posterior probability in the proposal, i.e. Step 2 of Algorithm 1. To this end, we only require h to be non-decreasing. Now we introduce the main result, which states that the mixing time bound can be improved linearly in the number of trials N for our locally balanced weighting scheme.

Theorem 1. *Let $\mathbf{P}_{\mathrm{MTM}}$ denote a Markov chain induced by Algorithm 1 with a locally balanced weight function (4) with some non-decreasing balancing function h. Define*

$$\mathcal{S}(x) = \{x' \in \mathcal{N}(x) \colon \pi(x')/\pi(x) > p^{t_1}\}, \tag{5}$$

for some constant t_1 and $s_0 = \max_{x \in \mathcal{X}} |\mathcal{S}(x)|$. Suppose that the following conditions hold with $p \geq 2$ and $t_2, t_3, t_4 \geq 0$ being some constants that satisfy $t_1 < t_2$, $t_3 \leq t_4 < t_2$.

(i) There exists $x' \in \mathcal{N}(x)$ such that $\pi(x')/\pi(x) \geq p^{t_2}$ for every $x \neq x^*$.

(ii) For any $x \in \mathcal{X}$, $p^{t_3} \leq |\mathcal{N}(x)| \leq p^{t_4}$.

(iii) $N = o\left(\min\left\{\frac{h(p^{t_2+t_3-t_4})}{h(p^{t_1+t_4-t_3})}, \frac{p^{t_3}}{s_0}\right\}\right)$.

Then, we have the ϵ-mixing time of $\mathbf{P}_{\mathrm{MTM}}$

$$t_{\mathrm{mix}}(\epsilon) = O\left(\frac{1}{N}(p^{-t_4} - p^{-t_2})^{-1}\ell(\Delta^*)\log\left(\min_{x\in\mathcal{X}}\pi(x)\right)^{-1}\right), \tag{6}$$

where Δ^* is a path ensemble defined in Appendix A.3.

Remark 1. By the recent result of [53, Lemma 3] based on a refined path argument, we can remove the term $\ell(\Delta^*)$ from (6) without changing the order of the bound.

Sketch of the proof. The complete proof is deferred to Appendix A.3. The first key step of our proof is to identify a suitable path ensemble Δ^* by defining a function $g : \mathcal{X} \to \mathcal{X}$ such that $g(x)$ has the highest posterior probability in the neighborhood of the state x. With π being unimodal and $\mathcal{X}$ being finite, repeated composition of g for any x will always lead to the mode x^*, that is $g^m(x) = x^*$ for some $m \in \mathbb{N}$, which we utilize to construct $\delta^*(x, y)$ for any $x, y \in \mathcal{X}, x \neq y$. See Appendix A.5 for an toy example. The other key step is to bound the congestion parameter $\rho(\Delta^*)$ defined in (3). To this end, we prove that the lower bound of the transition probability $\mathbf{P}(x, g(x))$ gets closer to Np^{-t_4} asymptotically, for any state $x \in \mathcal{X}\backslash\{x^*\}$.

Condition (i) implies that the posterior distribution is unimodal with the peak at the model x^*. We emphasize that the unimodality is with respect to the neighborhood relation $\mathcal{N}$ and is usually satisfied with some appropriate choice of a large enough neighborhood. For example, for Bayesian variable selection, it is shown in [47] that their proposed posterior may not be unimodal if a neighborhood is defined as a set of models within 1-Hamming distance, but expanding to the 2-Hamming distance neighborhood satisfies Condition (i) with $t_2 = 2$. Similarly, the community detection problem considered in [55] and the structure learning problem in [52] satisfy (i) with some positive constant t_2 by choosing an appropriate $\mathcal{N}$. Next, Condition (ii) states that the neighborhood size $|\mathcal{N}(x)|$ should neither be too huge nor vary much from state to state. The purpose of this condition is to control the ratio $\mathbf{K}_{\mathrm{RW}}(y, x)/\mathbf{K}_{\mathrm{RW}}(x, y)$ in (4) so that the posterior ratio $\pi(y)/\pi(x)$ dominates the term. Last but not least, Condition (iii) implies that the number of trials N should not be arbitrarily large. This is consistent with the empirical observation of [32] that additional trials do not always result in better mixing. Condition (iii) also implies that $s_0 = o(p^{t_3})$ since $N \geq 1$. This seems to be a strong condition, since the number of neighboring states with a relatively large posterior probability with respect to the current state can be very large if we consider a state with the smallest posterior that is surrounded by states with greater posterior. Yet, it is a common practice to restrict the support of the prior to rule out unrealistic models by imposing some high-dimensional regularity assumptions, such as sparsity. For example, [47] introduced the parameter of the maximum number of important covariates in variable selection, and [52] considered the maximum in-degree and out-degree in structure learning to restrict the model space. In the same context, [55] suggested the use of a feasible set for the initial partition in the community detection problem. In this regard, we may consider $\mathcal{X}$ as the restricted space so that s_0 can be controlled. Most importantly, Condition (iii) provides a key idea on how to select the number of trials N, which will be discussed in Section 3.2.

Remark 2 (Problem-specific applications). Theorem 1 can be used to prove rapid mixing of the MTM algorithm for Bayesian variable selection [47] and structure learning [52] with order N improvement of the mixing bound, given that the choice of N satisfies the condition (iii).

Remark 3 (Computational benefits). Since MTM with N trials requires calculating $2N - 1$ weight functions at each iteration, improving the mixing time bound by the factor of N leads to the same overall computational complexity as the single-try MH algorithm until the convergence of the Markov chain. However, compared to the single-try MH with N times longer sequential iterations, the calculation of N weight functions can be efficiently done in parallel. For example, calculating multiple weight functions can often be converted into a series of matrix multiplication problems, and optimized linear algebra libraries such as BLAS [7] can be used to exploit multiple execution units and pipe-lining. As examples, we outline parallelization strategies for the Bayesian variable selection (Section 4.1) and stochastic block model (Section 4.2) in Appendix D.1. Indeed, Table 1 in our simulation study results shows that MTM with moderate choice of N significantly reduces the wall-clock time until the convergence of the Markov chain.

3.2 Choice of the number of trials

Motivated by Condition (iii) in Theorem 1, we propose an algorithm to choose the number of trials N under a general setting applicable to high-dimensional model selection problems. For conciseness, we shall only focus on illustrating the case when we have balancing function $h(u) = \sqrt{u}$ and symmetric proposals (known $t_3 = t_4$) so that Condition (iii) becomes $N = o\left(\min\{p^{(t_2-t_1)/2}, p^{t_3}/s_0\}\right)$, although the algorithm can be easily generalized to other choices of h and non-symmetric proposals.

Algorithm 2: The choice of the number of trials, when $h(u) = \sqrt{u}$ and known $t_3 = t_4$.

Input: An initial state x_0, neighborhood $\mathcal{N}$, constant $\psi \in (0, 1)$.

Step 1. Calculate $\log_p(\pi(y_j)/\pi(x_0))$ for all $y_j \in \mathcal{N}(x_0)$, $j = 1, \dots, |\mathcal{N}(x_0)|$.

Step 2. Run k-means algorithm with k=2 on log-ratios $\{\log_p(\pi(y_j)/\pi(x_0))\}$ to obtain a partition, $\mathcal{C}_1$ (set of small log-ratios) and $\mathcal{C}_2$ (set of big log-ratios).

Step 3. Let $\hat{t}_2 = \min \mathcal{C}_2$, $\hat{t}_1 = \max \mathcal{C}_1$ and $\hat{s}_0 = |\mathcal{C}_2|$. †

while $p^{(\hat{t}_2-\hat{t}_1)/2} < p^{t_3}/\hat{s}_0$ **do**

  Update $\mathcal{C}_1 \leftarrow \mathcal{C}_1 \backslash \{\hat{t}_1\}$, $\hat{t}_1 \leftarrow \max \mathcal{C}_1$ and $\hat{s}_0 \leftarrow \hat{s}_0 + 1$.

Output: The number of trials $N = \lfloor (p^{t_3}/\hat{s}_0)^{\psi} \rfloor$.

† If $\hat{t}_2 < t_4$, redefine $\mathcal{C}_1$, $\mathcal{C}_2$ such that $\hat{t}_1 = \max \mathcal{C}_1 \leq t_4 < \min \mathcal{C}_2 = \hat{t}_2$. If no such $\mathcal{C}_2$ exists, set $\hat{t}_2 = t_4$, $\hat{s}_0 = 1$ and $\hat{t}_1 = \max \mathcal{C}_1$.

Algorithm 2 is depicted in Figure 2. In Step 1, the log probability ratio is calculated for all neighboring states of the initial state x_0. Step 2 identifies two subsets $\mathcal{C}_1$ and $\mathcal{C}_2$, corresponding to "bad moves" and "good moves" respectively. In Step 3, motivated by Condition (i), we set $\hat{t}_2 = \min \mathcal{C}_2$, a lower bound of log probability ratios among good moves. After t_2 is fixed, we first set $\hat{t}_1 = \max \mathcal{C}_1$, $\hat{s}_0 = |\mathcal{C}_2|$ and gradually decrease $\hat{t}_1$ (so that $\hat{s}_0$ increases) until the "crossing" $p^{(\hat{t}_2-\hat{t}_1)/2} \geq p^{t_3}/\hat{s}_0$ happens. The choice of $N = \lfloor (p^{t_3}/\hat{s}_0)^{\psi} \rfloor$ immediately after the crossing point approximates the worst-case scenario of $\min\{p^{(t_2-t_1)/2}, p^{t_3}/s_0\}$, hence providing a conservative estimate of N. We set $\psi = 0.9$ to match with the asymptotic dominance condition (iii).

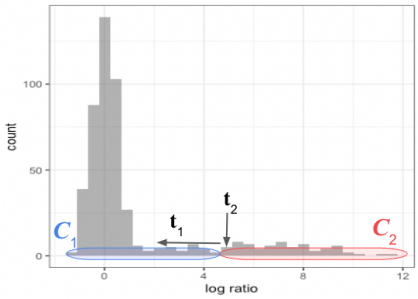

Figure 2: An illustration of Algorithm 2.

We run Algorithm 2 only once for an initial state x_0. It is based on the insight that for model selection problems, log probability ratios can often be partitioned into "bad moves" and "good moves". Although we considered any $x \in \mathcal{X} \backslash \{x^*\}$ in Condition (i), analyzing the initial state x_0 is not only computationally simple but also gives a good guess of t_2 since as chain proceeds to the highest posterior state x^*, the number of good moves and the magnitude of ratios generally decreases, for example in the Bayesian variable selection problem [54]. Another possible way of using Algorithm 2 is to re-evaluate N after a certain number of MCMC iterations. We note that another clustering algorithm can be substituted for the k-means algorithm in Step 2.

4 Simulation studies

In this section, we present the results for two Bayesian model selection problems, BVS and SBM. We defer the result for spatial clustering models to Appendix B.3 due to the page limit. Throughout the simulation studies, we use symmetric proposal schemes. We run single-try MH and MTM with different choices of $N \in \{5, 10, 50, 100, 500, 1000, 2000, 5000\}$ and four different choices of weight functions, one with (unscaled) ordinary weight function $w_{\text{ord}}(y \,|\, x) = \pi(y)$, and three with the proposed locally balanced weight functions $w_{\text{sqrt}}(y \,|\, x) = \sqrt{\pi(y)/\pi(x)}$, $w_{\min}(y \,|\, x) = \min\{1, \pi(y)/\pi(x)\}$, and $w_{\max}(y \,|\, x) = \max\{1, \pi(y)/\pi(x)\}$. We consider the scenarios where the data-generating model x^* receives the highest posterior probability to verify our main theoretical results. Similarly to [54], two performance measures are considered: 1) $H = \min\{t : x_t = x^*\}$ (hitting iteration): the number of MCMC iterations until the Markov chain $(x_0, x_1, \dots)$ hits x^*, 2) T_H (wall-clock hitting time): the wall-clock time taken until the chain hits x^*.

4.1 Bayesian variable selection (BVS)

Consider a high-dimensional linear model with response vector $\boldsymbol{y} \in \mathbb{R}^n$ and design matrix $\boldsymbol{X} \in \mathbb{R}^{n \times p}$, where the number of predictors p is much larger than the sample size n. BVS seeks to find the best subset of predictors, denoted as a binary vector $\gamma \in \{0,1\}^p$ described in Section 3.1, using the sparsity-inducing prior. We adopt the prior introduced by [47]:

$$\begin{aligned} &\text{Linear model:} && \boldsymbol{y} = \boldsymbol{X}_\gamma \boldsymbol{\beta}_\gamma + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathsf{N}\left(\mathbf{0}, \phi^{-1}\mathbf{I}_n\right) \\ &\text{Prior on } \boldsymbol{\beta}_\gamma \text{ and } \phi: && \boldsymbol{\beta}_\gamma \mid \phi, \gamma \sim \mathsf{N}\left(\mathbf{0}, \mathscr{G}\phi^{-1}(\boldsymbol{X}_\gamma^\top \boldsymbol{X}_\gamma)^{-1}\right), \quad \pi(\phi) \propto 1/\phi \\ &\text{Sparsity prior:} && \pi(\gamma) \propto p^{-\kappa|\gamma|}\mathbb{1}\left[|\gamma| \le s_{max}\right], \end{aligned}$$

where $\boldsymbol{X}_\gamma$ is a submatrix of $\boldsymbol{X}$ consisting of all j-th columns with $\gamma_j = 1$, $\boldsymbol{\beta}_\gamma \in \mathbb{R}^{|\gamma|}$ is the subvector of $\boldsymbol{\beta} \in \mathbb{R}^p$ with nonzero coefficients, where $|\gamma| = \sum_{j=1}^p \gamma_j$. The hyperparameters $\mathscr{G}$, κ and $s_{\max}$ control the weight of prior effect [50], sparsity strength and the maximum model size, respectively. We sample γ from the posterior distribution $\pi(\gamma \mid \boldsymbol{y})$ where other parameters $\boldsymbol{\beta}_\gamma$ and ϕ are marginalized out; see Appendix B.1 for details. We consider the following random walk proposal:

$$\mathbf{K}_{\mathrm{RW}}(\gamma, \gamma') = \begin{cases} \frac{1}{p}\mathbb{1}_{\mathcal{N}_1(\gamma)}(\gamma') \text{ (single flip)} & \text{if } |\gamma| < s_{\max}, \\ \frac{1}{2p}\mathbb{1}_{\mathcal{N}_1(\gamma)}(\gamma') + \frac{1}{2|\gamma|(p-|\gamma|)}\mathbb{1}_{\mathcal{N}_2(\gamma)}(\gamma') \text{ (single or double flip)} & \text{if } |\gamma| = s_{\max}, \end{cases}$$

where $\mathcal{N}_1(\gamma) = \{\gamma' : d_{\mathrm{H}}(\gamma, \gamma') = 1\}$ and $\mathcal{N}_2(\gamma) = \{\gamma' = (\gamma \cup \{j\}) \backslash \{\ell\} : j \notin \gamma, \ell \in \gamma\}$. With $n = 1000$ and $p = 5000$, we consider different settings of design matrix where each row $\boldsymbol{x}_i$ is i.i.d. sampled from $\boldsymbol{x}_i \sim \mathsf{N}(\mathbf{0}, \mathbf{I}_p)$ (independent) or $\boldsymbol{x}_i \sim \mathsf{N}(\mathbf{0}, \boldsymbol{\Sigma})$, $\Sigma_{jj'} = \exp(-|j - j'|)$ (dependent). The response vector $\boldsymbol{y}$ is generated from $\boldsymbol{y} \sim \mathsf{N}(\boldsymbol{X}\boldsymbol{\beta}, \mathbf{I}_n)$, with $\boldsymbol{\beta}$ having only 10 nonzero entries at the first 10 coordinates: $\boldsymbol{\beta} = \mathrm{SNR}\sqrt{(\log p)/n}(2, -3, 2, 2, -3, 3, -2, 3, -2, 3)$ where $\mathrm{SNR} \in \{2, 4\}$ is a signal-to-noise ratio. For each of the four settings, we simulate 50 replicate datasets. Per each replicate, algorithms are randomly initialized with state γ_0 such that $\gamma_0 \cap \gamma^* = \emptyset$ and $d_{\mathrm{H}}(\gamma_0, \gamma^*) = 20$ which implies $H = 20$ is the minimum required hitting iteration. All hyperparameter setups, the number of MCMC iterations, and other details are deferred to Appendix B.1.

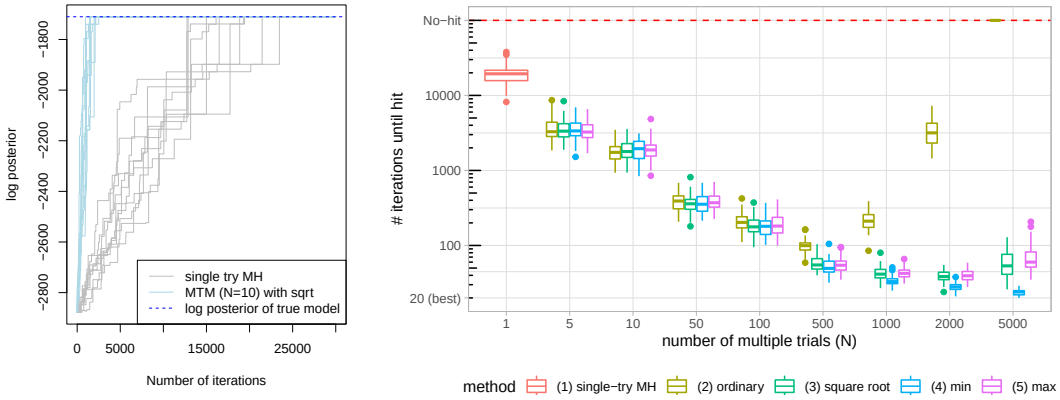

Figure 3: BVS simulation results for independent design and SNR $= 4$. (Left) Trace plots of unnormalized log-posterior probability using single-try MH and MTM ($N = 10$) with square root weighting function. (Right) Boxplot of H, the number of iterations until hit, against different choices of weight function and N based on 50 replicates.

Results under independent design with SNR $= 4$ are summarized in Figure 3. The trace plot shows that single-try MH reaches the true state γ^* at around 20,000 iterations, whereas the MTM with $N = 10$ reaches the true state γ^* at around 2,000 iterations, smaller by a factor of 10. The boxplot confirms our findings that a larger number of trials N is not always desirable, especially when using the ordinary weight function w_{ord}. The proposed locally balanced weight functions using three different choices of h perform much better than w_{ord} for large N. Especially, the performance of $w_{\min}$ seems to be robust to the choice of large N in the high SNR case. This can be explained by the fact that the denominator in (2) is often larger in high SNR data, and the choice of $\min\{1, \pi(y)/\pi(x)\}$ limits its growth; see also Section 3.1. The choice of N suggested from Algorithm 2 with $\psi = 0.9$ is overall reasonable, which has a median of 349 over 50 datasets. The detailed results for other settings are available in Appendix B.1.

4.2 Stochastic block model (SBM)

SBM [20] is a popular generative model of an undirected graph which assumes a block structure of edge connection probabilities $\boldsymbol{Q}$ to describe community structure. We consider Bayesian SBM [34] by assigning prior on $\boldsymbol{Q}$ and partition $\mathbf{z} = (z_1, \dots, z_p) \in \{1, \dots, K\}^p$ with fixed number of blocks K, where z_i indicates membership label of i-th node. The main goal is to find partition $\mathbf{z}$ (up to a label permutation) that best describes the graph, denoted as an adjacency matrix $\boldsymbol{A} \in \{0,1\}^{p \times p}$. We follow the prior of [55]:

$$\begin{aligned} \text{Edge appearance:} \quad & A_{ij} \mid \boldsymbol{Q}, \mathbf{z} \overset{\text{ind}}{\sim} \mathsf{Bernoulli}(Q_{z_i z_j}), 1 \le i < j \le p \\ \text{Prior on blockwise probabilities:} \quad & Q_{uv} \overset{\text{iid}}{\sim} \mathsf{Beta}(\kappa_1, \kappa_2), 1 \le u \le v \le K \\ \text{Prior on partition:} \quad & \pi(\mathbf{z}) \propto \mathbb{1}(\mathbf{z} \in S_\alpha), \end{aligned}$$

with restricted partition space $S_\alpha = \{\mathbf{z} : \sum_{i=1}^p \mathbb{1}(z_i = u) \in \left[\frac{p}{\alpha K}, \frac{\alpha p}{K}\right] \text{ for all } u = 1, \dots, K\}$ for some $\alpha > 0$, which excludes partitions whose block sizes differ too much. We sample $\mathbf{z}$ from the posterior $\pi(\mathbf{z} \mid \boldsymbol{A})$ where $\boldsymbol{Q}$ is marginalized out; see Appendix B.2 for details. We consider the following proposal where $\tilde{d}_{\mathrm{H}}$ is a permutation-invariant Hamming distance [51]:

$$\mathbf{K}_{\mathrm{RW}}(\mathbf{z}, \mathbf{z}') = 1/(p(K-1))\mathbb{1}_{\mathcal{N}(\mathbf{z})}(\mathbf{z}'), \quad \mathcal{N}(\mathbf{z}) = \{\mathbf{z}' : \tilde{d}_{\mathrm{H}}(\mathbf{z}, \mathbf{z}') = 1\} \quad \text{(single flip)}.$$

With the number of nodes $p = 1000$, we consider two different settings of the blocks $K \in \{2, 5\}$ with true partition $\mathbf{z}^* = (1, \dots, 1, \cdots, K, \dots, K)$ being balanced (p/K times each) which ensures $\mathbf{z}^* \in S_\alpha$. In each setting of K, we generate a graph from the homogeneous SBM [16, 55], where $Q_{uv} = a$ if $u = v$ and $Q_{uv} = b$ otherwise. We choose probability pairs (a, b) based on the Chernoff-Hellinger divergence (CH) [1], defined as $\mathrm{CH} := p(\sqrt{a} - \sqrt{b})^2/(K \log p)$ for $a, b \asymp (\log p)/p$, which can be interpreted as a signal-to-noise ratio. Our choices are $\mathrm{CH} \approx 2$ (weak signal) and $\mathrm{CH} \approx 10$ (high signal), both being greater than 1 which is necessary to satisfy posterior consistency [1, 51, 55]. For each of the the four combinations of (K, CH), we simulate 50 replicate datasets. For each replicate, algorithms are randomly initialized with state $\mathbf{z}_0$ such that $\tilde{d}_{\mathrm{H}}(\mathbf{z}_0, \mathbf{z}^*) = 400$ which implies $H = 400$ is the minimum required hitting iteration. All hyperparameter setups, the number of MCMC iterations, and other details are deferred to Appendix B.2.

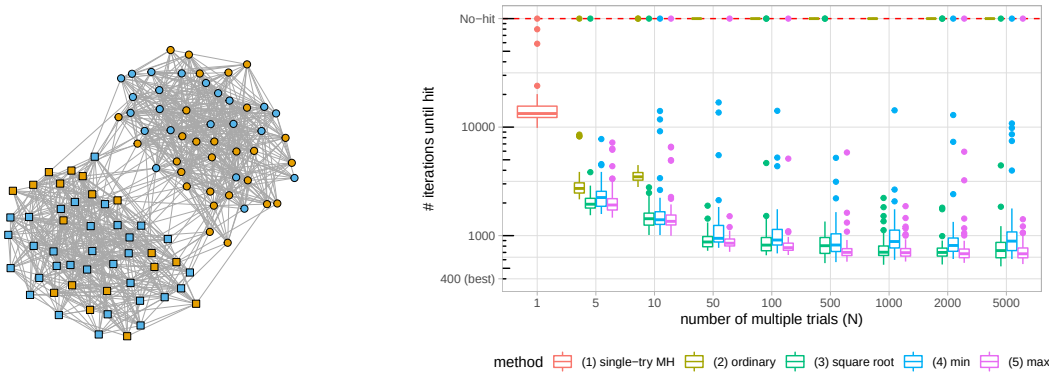

Figure 4: (Left) An example graph generated from SBM with $p = 100$, $K = 2, a = 0.32$, and $b = 0.02$ so that $\mathrm{CH} \approx 2$. Node shape (square, circle) denotes the true partition. Node color (blue, yellow) represents the initial state $\mathbf{z}_0$ satisfying $\tilde{d}_{\mathrm{H}}(\mathbf{z}_0, \mathbf{z}^*) = 40$. (Right) Boxplot of H against different weight functions and N based on 50 replicates when $p = 1000, K = 2$, and $\mathrm{CH} \approx 2$.

Results for $K = 2$ and $\mathrm{CH} \approx 2$ are summarized in Figure 4. The boxplot of H is similar to that of Figure 3: w_{ord} quickly deteriorates as N increases, but the locally balanced weight functions mostly converge with a larger choice of N, although some Markov chains fail to converge due to the low CH. Especially, now the performance of $w_{\max}$ or w_{sqrt} is slightly better than $w_{\min}$ when N is large. This can be explained by the fact that in low signal case ($\mathrm{CH} \approx 2$), there are few neighboring states that have high probability ratios. It is important to catch such states to increase the acceptance ratio (2), and the $w_{\min}$ cannot do this efficiently as it is upper bounded by 1. The choice of N suggested from Algorithm 2 with $\psi = 0.9$ has a median of 15 over 50 datasets. The detailed simulation results for other settings are available at Appendix B.2.

Table 1: Summary of T_H, median wall-clock time (in seconds) until the Markov chain hit the highest posterior state, for the previous two examples using the square root weight function over 50 replicates.

N		1	5	10	50	100	500	1000	2000	5000
T_H	BVS, indep, SNR$=4$	1.30	0.81	0.46	0.11	0.07	**0.04**	0.05	0.09	0.30
	SBM, $K=2$, CH ≈ 2	0.82	0.40	**0.35**	0.48	0.75	1.49	2.15	3.78	9.30

Finally, Table 1 provides the summary of wall-clock hitting time T_H for the previous two examples, which shows the clear computational benefit of MTM over the single-try MH thanks to the parallelism. As N increases, computation becomes more demanding compared to the decreasing rate of H, so the optimal choice of N with respect to T_H is often less than the minimizer of H.

Spatial clustering models and additional information. In Appendix B.3, we present simulation study results with spatial clustering models where the key findings align with those from BVS and SBM. In Appendix B.4, we also analyze the behavior of MTM algorithms on multimodal target distributions. The code is available at `https://github.com/changwoo-lee/rapidMTM`.

5 Real data applications and discussion

Real data applications. We carry out two real data application analyses to corroborate our findings beyond the scope of the simulation settings. The performance measures suggested in Section 4, the hitting iteration H and the corresponding wall-clock hitting time T_H, are no longer available for real data analysis since we cannot identify the true data generating state x^*. Instead, we evaluate the performance based on the acceptance rate and the number of unique states visited. Both real data analyses confirm our theory: the performance deteriorates significantly as the number of trials N grows when using w_{ord}, but it does not when using w_{sqrt}, $w_{\min}$, or $w_{\max}$.

The first real data application is to find a subset of genetic variants from a genome-wide association study (GWAS) dataset that best explains the cup-to-disk ratio averaged over two eyes, which is used to assess the risk of glaucoma. We employ BVS model described in Section 4.1, with the number of samples $n = 5418$ and the number of gene variants $p = 7255$ which are selected after the preliminary screening procedure described in [54, Section 6]. We tabulate the acceptance rate and the number of unique states visited along the different number of trials N, and also report the posterior inclusion probabilities of the top 10 genetic variants in Appendix C.1 and E with discussion.

The other application is to learn the underlying directed acyclic graph (DAG) model for the single-cell RNA dataset on Alzheimer's disease [22]. We preprocess the dataset as [10, Section 6], which yields the sample size $n = 1666$ and 73 genes, and we utilize the DAG model described in [10, Section 2.2], whose model size is equal to $73! \approx 4.5 \times 10^{105}$. We defer the performance result and other details (including the log-posterior trace plots for 4 different weight functions) to Appendix C.2.

Discussion and future work. We prove that, under some assumptions, the mixing time bound of the MTM algorithm is smaller than that of the MH algorithm by a factor of the number of trials. Motivated by the observation that popular choices of weight function can cause mixing problems, we propose locally balanced weight functions for which a mixing time bound is proven. We also suggest a theoretically guided choice of N for practical use. Future research may investigate the mixing time of other generalizations of MTM such as the case with correlated weight functions. It will also be of interest to extend MTM and analyze its mixing time to better handle the potential multimodality in the target distribution, by combining it with techniques such as annealing or tempering [9].

Acknowledgments and Disclosure of Funding

We thank James Cai for providing us with the processed single-cell data set on Alzheimer's disease. We thank the submitters and participants of the two dbGaP studies (phs000308.v1.p1 and phs000238.v1.p1), which were funded by NIH. We thank Florian Maire, Giacomo Zanella, and Philippe Gagnon for the helpful discussion at ISBA 2022. The research of Changwoo Lee, Zhao Tang Luo, and Huiyan Sang was partially supported by NSF DMS-1854655 and DMS-2210456.

References

[1] Emmanuel Abbe. Community detection and stochastic block models: recent developments. *The Journal of Machine Learning Research*, 18(1):6446–6531, 2017.

[2] David Aldous. Random walks on finite groups and rapidly mixing Markov chains. In *Séminaire de Probabilités XVII 1981/82*, pages 243–297. Springer, 1983.

[3] David Aldous and Persi Diaconis. Shuffling cards and stopping times. *The American Mathematical Monthly*, 93(5):333–348, 1986.

[4] Dominique Bakry, Ivan Gentil, Michel Ledoux, et al. *Analysis and geometry of Markov diffusion operators*, volume 103. Springer, 2014.

[5] Daniel Barry and John A Hartigan. A Bayesian analysis for change point problems. *Journal of the American Statistical Association*, 88(421):309–319, 1993.

[6] Mylène Bédard, Randal Douc, and Eric Moulines. Scaling analysis of multiple-try MCMC methods. *Stochastic Processes and their Applications*, 122(3):758–786, 2012.

[7] L Susan Blackford, Antoine Petitet, Roldan Pozo, Karin Remington, R Clint Whaley, James Demmel, Jack Dongarra, Iain Duff, Sven Hammarling, Greg Henry, et al. An updated set of basic linear algebra subprograms (BLAS). *ACM Transactions on Mathematical Software*, 28(2):135–151, 2002.

[8] Russ Bubley and Martin Dyer. Path coupling: A technique for proving rapid mixing in Markov chains. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 223–231. IEEE, 1997.

[9] Roberto Casarin, Radu Craiu, and Fabrizio Leisen. Interacting multiple try algorithms with different proposal distributions. *Statistics and Computing*, 23(2):185–200, 2013.

[10] Hyunwoong Chang, James Cai, and Quan Zhou. Order-based structure learning without score equivalence. *arXiv preprint arXiv:2202.05150*, 2022.

[11] Siddhartha Chib. Estimation and comparison of multiple change-point models. *Journal of econometrics*, 86(2):221–241, 1998.

[12] Radu V Craiu and Christiane Lemieux. Acceleration of the multiple-try Metropolis algorithm using antithetic and stratified sampling. *Statistics and computing*, 17(2):109–120, 2007.

[13] Persi Diaconis and Daniel Stroock. Geometric bounds for eigenvalues of Markov chains. *The annals of applied probability*, pages 36–61, 1991.

[14] Roland Lvovich Dobrushin, Roman Koteckỳ, and Senya Shlosman. *Wulff construction: a global shape from local interaction*, volume 104. American Mathematical Society Providence, 1992.

[15] Nir Friedman and Daphne Koller. Being Bayesian about network structure. A Bayesian approach to structure discovery in Bayesian networks. *Machine learning*, 50(1):95–125, 2003.

[16] Junxian Geng, Anirban Bhattacharya, and Debdeep Pati. Probabilistic community detection with unknown number of communities. *Journal of the American Statistical Association*, 114(526):893–905, 2019.

[17] Edward I George and Robert E McCulloch. Approaches for Bayesian variable selection. *Statistica sinica*, pages 339–373, 1997.

[18] Venkatesan Guruswami. Rapidly mixing Markov chains: a comparison of techniques (a survey). *arXiv preprint arXiv:1603.01512*, 2016.

[19] Chris Hans, Adrian Dobra, and Mike West. Shotgun stochastic search for "large p" regression. *Journal of the American Statistical Association*, 102(478):507–516, 2007.

[20] Paul W Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983.

[21] Mark Jerrum. A very simple algorithm for estimating the number of k-colorings of a low-degree graph. *Random Structures & Algorithms*, 7(2):157–165, 1995.

[22] Jing Jiang, Cankun Wang, Ren Qi, Hongjun Fu, and Qin Ma. scREAD: a single-cell rna-seq database for alzheimer's disease. *Iscience*, 23(11):101769, 2020.

[23] Changwoo Lee, Zhao Tang Luo, and Huiyan Sang. T-LoHo: A Bayesian regularization model for structured sparsity and smoothness on graphs. *Advances in Neural Information Processing Systems*, 34:598–609, 2021.

[24] Sirio Legramanti, Tommaso Rigon, Daniele Durante, and David B Dunson. Extended stochastic block models with application to criminal networks. *The Annals of Applied Statistics*, 16(4):2369–2395, 2022.

[25] David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

[26] Furong Li and Huiyan Sang. Spatial homogeneity pursuit of regression coefficients for large datasets. *Journal of the American Statistical Association*, 114(527):1050–1062, 2019.

[27] Jun S Liu, Faming Liang, and Wing Hung Wong. The multiple-try method and local optimization in Metropolis sampling. *Journal of the American Statistical Association*, 95(449):121–134, 2000.

[28] Zhao Tang Luo, Huiyan Sang, and Bani Mallick. A Bayesian contiguous partitioning method for learning clustered latent variables. *Journal of Machine Learning Research*, 22(37):1–52, 2021.

[29] David Madigan, Jeremy York, and Denis Allard. Bayesian graphical models for discrete data. *International Statistical Review/Revue Internationale de Statistique*, pages 215–232, 1995.

[30] Luca Martino. A review of multiple try MCMC algorithms for signal processing. *Digital Signal Processing*, 75:134–152, 2018.

[31] Luca Martino and Victor Elvira. Metropolis sampling. *Wiley StatsRef: Statistics Reference Online*, pages 1–18, 2017.

[32] Luca Martino and Francisco Louzada. Issues in the multiple try Metropolis mixing. *Computational Statistics*, 32(1):239–252, 2017.

[33] Luca Martino and Jesse Read. On the flexibility of the design of multiple try Metropolis schemes. *Computational Statistics*, 28(6):2797–2823, 2013.

[34] Krzysztof Nowicki and Tom A B Snijders. Estimation and prediction for stochastic blockstructures. *Journal of the American statistical association*, 96(455):1077–1087, 2001.

[35] Silvia Pandolfi, Francesco Bartolucci, and Nial Friel. A generalization of the multiple-try Metropolis algorithm for Bayesian estimation and model selection. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 581–588. JMLR Workshop and Conference Proceedings, 2010.

[36] Yuval Peres and Perla Sousi. Mixing times are hitting times of large sets. *Journal of Theoretical Probability*, 28(2):488–519, 2015.

[37] Peter H Peskun. Optimum Monte-Carlo sampling using Markov chains. *Biometrika*, 60(3):607–612, 1973.

[38] David Pollard and Dana Yang. Rapid mixing of a Markov chain for an exponentially weighted aggregation estimator. *arXiv preprint arXiv:1909.11773*, 2019.

[39] William M Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical association*, 66(336):846–850, 1971.

[40] Dana Randall. Rapidly mixing Markov chains with applications in computer science and physics. *Computing in Science & Engineering*, 8(2):30–41, 2006.

[41] B Ramakrishna Rau and Joseph A Fisher. Instruction-level parallel processing: History, overview, and perspective. *The journal of Supercomputing*, 7(1):9–50, 1993.

[42] Laurent Saloff-Coste. Lectures on finite Markov chains. *Lectures on probability theory and statistics*, pages 301–413, 1997.

[43] Alistair Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, probability and Computing*, 1(4):351–370, 1992.

[44] Michael Smith and Robert Kohn. Nonparametric regression using Bayesian variable selection. *Journal of Econometrics*, 75(2):317–343, 1996.

[45] Mahlet G Tadesse and Marina Vannucci. *Handbook of Bayesian variable selection*. CRC Press, 2021.

[46] Xiaodong Yang and Jun S Liu. Convergence rate of multiple-try Metropolis independent sampler. *arXiv preprint arXiv:2111.15084*, 2021.

[47] Yun Yang, Martin J Wainwright, and Michael I Jordan. On the computational complexity of high-dimensional Bayesian variable selection. *The Annals of Statistics*, 44(6):2497–2532, 2016.

[48] Giacomo Zanella. Informed proposals for local MCMC in discrete spaces. *Journal of the American Statistical Association*, 115(530):852–865, 2020.

[49] Giacomo Zanella and Gareth Roberts. Scalable importance tempering and Bayesian variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(3):489–517, 2019.

[50] Arnold Zellner. On assessing prior distributions and Bayesian regression analysis with g-prior distributions. *Bayesian inference and decision techniques*, 1986.

[51] Anderson Y Zhang and Harrison H Zhou. Minimax rates of community detection in stochastic block models. *The Annals of Statistics*, 44(5):2252–2280, 2016.

[52] Quan Zhou and Hyunwoong Chang. Complexity analysis of Bayesian learning of high-dimensional DAG models and their equivalence classes. *arXiv preprint arXiv:2101.04084*, 2021.

[53] Quan Zhou and Aaron Smith. Rapid convergence of informed importance tempering. In *International Conference on Artificial Intelligence and Statistics*, pages 10939–10965. PMLR, 2022.

[54] Quan Zhou, Jun Yang, Dootika Vats, Gareth O Roberts, and Jeffrey S Rosenthal. Dimension-free mixing for high-dimensional Bayesian variable selection. *Journal of the Royal Statistical Society: Series B, to appear.*, 2022+.

[55] Bumeng Zhuo and Chao Gao. Mixing time of Metropolis-Hastings for Bayesian community detection. *Journal of Machine Learning Research*, 22(10):1–89, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 5 for the future work direction.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A] This work is mainly theoretical and does not present any foreseeable societal consequences.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix A for all proofs.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code is available at `https://github.com/changwoo-lee/rapidMTM`.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Please see Section 4 and Appendix B.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See boxplots associated with simulation results.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix B for the detailed settings of the simulation results.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] The code is available at `https://github.com/changwoo-lee/rapidMTM`.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Appendices

In Appendix A, we provide proofs of Proposition 1, Proposition 2, and Theorem 1 in the main text. In Appendix B, we provide more details of the Bayesian variable selection (BVS) and stochastic block model (SBM) in Section 4 as well as a detailed simulation study on the spatial clustering model (SCM). In addition, we study the performance of multiple-try Metropolis for the case with multimodal target distributions, following the BVS simulation setting of [54]. In Appendix C, we present details of two real data application analyses. In Appendix D, we add a more detailed discussion on parallelization, state space of interest, and the behavior of MTM on continuous state space. Finally, we provide additional tables on the real data analysis results in Appendix E.

A A path method for proving the mixing time bound for multiple-try Metropolis algorithm

A.1 Proof of Proposition 1

This section aims to provide a summary of the existing results on proving mixing time bound via path methods. We refer readers to [13, 25, 42, 43] for more details. Let $\mathbf{P}(x, y)$ denote the transition probability for an irreducible, aperiodic chain on the finite state space $\mathcal{X}$. Assume $\mathbf{P}$ satisfies the detailed balance condition with respect to the probability distribution π, that is, $\pi(x)\mathbf{P}(x, y) = \pi(y)\mathbf{P}(y, x)$ for $x, y \in \mathcal{X}$, which leads to π being stationary for $\mathbf{P}$ [25, Proposition 1.20]. $\mathbf{P}$ may be thought as a $|\mathcal{X}| \times |\mathcal{X}|$ stochastic matrix, which means $\mathbf{P}(x, y) \geq 0$ and $\sum_{z\in\mathcal{X}} \mathbf{P}(x, z) = 1$ for all $x, y \in \mathcal{X}$, and π can be regarded as a $|\mathcal{X}|$-dimensional stochastic vector since $\sum_{x\in\mathcal{X}} \pi(x) = 1$. By the spectral decomposition, we can sort the eigenvalues of $\mathbf{P}$ as

$$1 = \lambda_0 > \lambda_1 \geq \cdots \geq \lambda_{|\mathcal{X}|-1} > -1,$$

due to $\mathbf{P}$ being stochastic, irreducible and aperiodic [25, Lemma 12.1]. Let $\lambda_{\max} = \max\{\lambda_1, |\lambda_{|\mathcal{X}|-1}|\}$. We say $\mathrm{Gap}(\mathbf{P}) = 1 - \lambda_{\max}$ is the *spectral gap* of the chain $\mathbf{P}$. Intuitively, if the spectral gap is close to zero, the chain requires a large number of steps to be close to the stationary distribution in total variation distance. The following lemma draws the connection between the spectral gap and mixing time defined in Section 2.3.

Lemma 1. *Let $t_{\mathrm{mix}}(\epsilon)$ denote ϵ-mixing time defined in Section 2.3, then*

$$t_{\mathrm{mix}}(\epsilon) \leq 2\{\mathrm{Gap}(\mathbf{P})\}^{-1} \left[\log(1/\epsilon) + \log\left\{\min_{x\in\mathcal{X}} \pi(x)\right\}^{-1}\right].$$

Proof. We consider $\mathbf{P}_{\mathrm{lazy}} = (\mathbf{I} + \mathbf{P})/2$ so that all eigenvalues of $\mathbf{P}_{\mathrm{lazy}}$ are positive and the spectral gap becomes $\mathrm{Gap}(\mathbf{P}_{\mathrm{lazy}}) = 1 - \lambda_1$. By (1.10) of [13, Proposition 3],

$$4 \max_{x\in\mathcal{X}} \|\mathbf{P}^t_{\mathrm{lazy}}(x, \cdot) - \pi(\cdot)\|^2_{\mathrm{TV}} \leq \max_{x\in\mathcal{X}} \{(1 - \pi(x))/\pi(x)\} \exp[-2t\mathrm{Gap}(\mathbf{P}_{\mathrm{lazy}})],$$

for $t \in \mathbb{N}$. We set $\max_{x\in\mathcal{X}}\{(1 - \pi(x))/\pi(x)\} \exp[-2t\mathrm{Gap}(\mathbf{P}_{\mathrm{lazy}})] \leq 4\epsilon^2$, and solving the inequality with t gives

$$t \geq \{\mathrm{Gap}(\mathbf{P}_{\mathrm{lazy}})\}^{-1} \left(\frac{1}{2} \max_{x\in\mathcal{X}} \log(\pi(x)^{-1} - 1) - \log 2 + \log(1/\epsilon)\right).$$

Since

$$\begin{aligned}
&\{\mathrm{Gap}(\mathbf{P}_{\mathrm{lazy}})\}^{-1} \left(\frac{1}{2} \max_{x\in\mathcal{X}} \log(\pi(x)^{-1} - 1) - \log 2 + \log(1/\epsilon)\right) \\
&\leq \{\mathrm{Gap}(\mathbf{P}_{\mathrm{lazy}})\}^{-1} \left[\log(1/\epsilon) + \log\left\{\min_{x\in\mathcal{X}} \pi(x)\right\}^{-1}\right] \\
&\leq 2\{\mathrm{Gap}(\mathbf{P})\}^{-1} \left[\log(1/\epsilon) + \log\left\{\min_{x\in\mathcal{X}} \pi(x)\right\}^{-1}\right],
\end{aligned}$$

it follows that to achieve "ϵ-mixing", it suffices to choose

$$t \geq 2\{\mathrm{Gap}(\mathbf{P})\}^{-1}\left[\log(1/\epsilon) + \log\left\{\min_{x\in\mathcal{X}} \pi(x)\right\}^{-1}\right].$$

This concludes the proof. □

Our next interest is to find the lower bound of the spectral gap, which leads to the upper bound of the mixing time. It is often the case that the functional analysis tool is useful to establish such bound. Let $\mathbb{R}^{\mathcal{X}} = \{f\colon \mathcal{X} \to \mathbb{R}\}$ and $\ell^2(\pi) \subset \mathbb{R}^{\mathcal{X}}$ be the vector space equipped with an inner product $\langle\cdot,\cdot\rangle_\pi$, which is defined by $\langle f_1, f_2\rangle_\pi = \sum_{x\in\mathcal{X}} f_1(x)f_2(x)\pi(x)$ for all $f_1, f_2 \in \ell^2(\pi)$. We can regard the transition probability $\mathbf{P}$ as a function operator in $\ell^2(\pi)$, which can be defined as $\mathbf{P}f(x) = \sum_{y\in\mathcal{X}} \mathbf{P}(x,y)f(y)$. We define the *Dirichlet form* associated to the pair $(\mathbf{P}, \pi)$ by

$$\mathcal{E}(f_1, f_2) = \langle(\mathbf{I} - \mathbf{P})f_1, f_2\rangle_\pi \quad \text{for } f_1, f_2 \in \ell^2(\pi).$$

By the reversibility of $\mathbf{P}$, we can easily check that $\mathcal{E}(f) = \mathcal{E}(f,f) = \frac{1}{2}\sum_{x,y\in\mathcal{X}}[f(x) - f(y)]^2\pi(x)\mathbf{P}(x,y)$. The spectral gap can be defined using the Dirichlet form as follows [25, Remark 13.8]:

$$\mathrm{Gap}(\mathbf{P}) = \min_{\substack{f\in\ell^2(\pi)\\ \mathrm{Var}_\pi(f)\neq 0}} \frac{\mathcal{E}(f)}{\mathrm{Var}_\pi(f)},$$

where $\mathrm{Var}_\pi(f) = \sum_{x\in\mathcal{X}}(f(x) - \mathbb{E}_\pi f)^2\pi(x)$. This definition also has a link to the famous *Poincaré inequality* [4], that is, $\mathrm{Var}_\pi(f) \leq C\mathcal{E}(f)$ for all $f \in \ell^2(\pi)$, because the smallest constant C is equal to $\{\mathrm{Gap}(\mathbf{P})\}^{-1}$. The next lemma uses Poincaré inequality and an arbitrary path ensemble Δ defined in the main text.

Lemma 2 (Corollary 6, [43])**.** *For an arbitrary path ensemble Δ,*

$$\mathrm{Gap}(\mathbf{P}) \geq \frac{1}{\rho(\Delta)l(\Delta)}.$$

where $\ell(\Delta) = \max_{x,y}|\delta(x,y)|$ and

$$\rho(\Delta) = \max_{(u,v)\in E} \frac{1}{\pi(u)\mathbf{P}(u,v)} \sum_{x,y:\delta(x,y)\ni(u,v)} \pi(x)\pi(y)$$

Proof. We follow the proof given in [42, Theorem 3.2.1]. For each $(x,y) \in \mathcal{X}\times\mathcal{X}$ and for any function $f \in \mathbb{R}^{\mathcal{X}}$, we can write $f(y) - f(x) = \sum_{(u,v)\in\delta(x,y)} f(v) - f(u)$. By using Cauchy-Schwarz, multiplying $\pi(x)\pi(y)/2$, and summing over x and y,

$$|f(y) - f(x)|^2 \leq |\delta(x,y)| \sum_{(u,v)\in\delta(x,y)} |f(v) - f(u)|^2$$

$$\implies \underbrace{\frac{1}{2}\sum_{x,y}|f(y) - f(x)|^2\pi(x)\pi(y)}_{=\mathrm{Var}_\pi(f)} \leq \frac{1}{2}\sum_{x,y}|\delta(x,y)| \sum_{(u,v)\in\delta(x,y)} |f(v) - f(u)|^2\pi(x)\pi(y),$$

where the right-hand side becomes

$$\frac{1}{2}\sum_{(u,v)\in E}\left\{\frac{1}{\pi(u)\mathbf{P}(u,v)} \sum_{x,y:\delta(x,y)\ni(u,v)} |\delta(x,y)|\pi(x)\pi(y)\right\}|f(v) - f(u)|^2\pi(u)\mathbf{P}(u,v)$$

$$\leq \underbrace{\max_{(u,v)\in E}\left\{\frac{1}{\pi(u)\mathbf{P}(u,v)} \sum_{x,y:\delta(x,y)\ni(u,v)} \pi(x)\pi(y)\right\}}_{\rho(\Delta)} \underbrace{\left(\max_{x,y}|\delta(x,y)|\right)}_{\ell(\Delta)} \underbrace{\left(\frac{1}{2}\sum_{(u,v)\in E}|f(v) - f(u)|^2\pi(u)\mathbf{P}(u,v)\right)}_{\mathcal{E}(f)}.$$

This satisfies the Poincaré inequality, which yields the conclusion. □

By combining the results of Lemma 1 and Lemma 2, we get the conclusion of Proposition 1.

A.2 Proof of Proposition 2

Proof. Recall that the form of the weight function suggested in [27] is given by

$$w(y \mid x) = \pi(y)\mathbf{K}_{\mathrm{RW}}(y,x)\lambda(y,x),$$

where $\lambda(x,y) = \lambda(y,x)$ is a non-negative symmetric function in x and y, and satisfies $\lambda(x,y) > 0$ whenever $\mathbf{K}_{\mathrm{RW}}(x,y) > 0$. If we put

$$\lambda(x,y) = \frac{1}{\pi(y)\mathbf{K}_{\mathrm{RW}}(y,x)} h\left(\frac{\pi(y)\mathbf{K}_{\mathrm{RW}}(y,x)}{\pi(x)\mathbf{K}_{\mathrm{RW}}(x,y)}\right),$$

it is easy to check that the conditions are met. □

A.3 Proof of Theorem 1

In this section, we prove our main result by using Proposition 1. The main step of the proof is identifying the path ensemble Δ^* that makes the mixing time bound tight. To this end, we need to choose exactly one path $\delta^*(x,y)$ for each tuple $(x,y) \in \mathcal{X} \times \mathcal{X}$. (Note that x and y cannot be identical by the definition of path). From the intuition described in Section 2.3, an edge (u,v) has a large capacity if $\pi(u)$ and $\pi(v)$ are large. For example, if an edge contains the highest posterior state x^*, we can let the edge be traversed by a large number of paths. Given an edge with a small capacity, however, we need to ensure that the edge overlies with a small number of paths. Importantly, we do not let a path $\delta^*(x,y)$ pass through an edge with a small capacity if both $\pi(x)$ and $\pi(y)$ are large, so that the edge can maintain a small unit flow size. We may envision the topography of the path ensemble; x^* becomes the hub, while states with low posterior probability are located on the outskirt.

We construct the path ensemble Δ^* according to the description above. The construction of Δ^* is similar to that of [47] and [52]. With a neighborhood relation $\mathcal{N}$ that satisfies the conditions in Theorem 1, $g\colon \mathcal{X} \to \mathcal{X}$,

$$g(x) = \begin{cases} \arg\max_{x' \in \mathcal{N}(x)} \pi(x')^{1} & \text{if } x \neq x^*, \\ x^* & \text{otherwise.} \end{cases} \tag{7}$$

By Condition (i) in Theorem 1, there exists $m \in \mathbb{N}$ such that $g^m(x) = (\overbrace{g \circ \cdots \circ g}^{m \text{ times}})(x) = x^*$ for any $x \in \mathcal{X}$ and x^* is the only fixed point of g. (x^* can be thought as an attractor in dynamic systems.) For all $x, y \in \mathcal{X}$ with $x \neq y$, we have three cases; (i) $g^m(x) = y$ for some $m \in \mathbb{N} \setminus \{0\}$, (ii) $g^m(y) = x$ for some $m \in \mathbb{N} \setminus \{0\}$, or (iii) neither (i) nor (ii). If $(x,y) \in \mathcal{X} \times \mathcal{X}$ belongs to (i), we define $\delta^*(x,y) = (x, g(x), \ldots, g^m(x) = y)$. Similarly, if $(x,y) \in \mathcal{X} \times \mathcal{X}$ belongs to (ii), let $\delta^*(x,y) = (x = g^m(y), \ldots, g(y), y)$. For the case (iii), if $m_1, m_2 \in \mathbb{N} \setminus \{0\}$ are the minimum numbers that satisfy $g^{m_1}(x) = x^*, g^{m_2}(y) = x^*$, respectively, we let $\delta^*(x,y) = (x, g(x), \ldots, g^{m_1}(x) = x^* = g^{m_2}(y), \ldots, g(y), y)$. This yields the path ensemble Δ^*. We provide a toy example on how to construct g and any path $\delta^*(x,y)$ for each tuple $(x,y) \in \mathcal{X} \times \mathcal{X}$ associated with g in Appendix A.5.

Next, we make a bound for the congestion parameter $\rho(\Delta^*)$. We let $\Lambda(u) = \{x \in \mathcal{X} : u = g^k(x), k \in \mathbb{N}\}$ denote the ancestor set of u with respect to g. If $(u,v) \in \delta^*(x,y)$ for some $x, y \in \mathcal{X}$, we can easily verify that $x \in \Lambda(u)$ by the construction of δ^*. This implies $\{(x,y) \in \mathcal{X} \times \mathcal{X} : (u,v) \in \delta^*(x,y)\} \subseteq \Lambda(u) \times \mathcal{X}$. It follows that

$$\begin{aligned}
\rho(\Delta^*) &\leq \max_{(u,v):v=g(u)} \frac{1}{\pi(u)\mathbf{P}(u,v)} \sum_{(x,y)\in\Lambda(u)\times\mathcal{X}} \pi(x)\pi(y) \\
&= \max_{(u,v):v=g(u)} \frac{1}{\pi(u)\mathbf{P}(u,v)} \left(\sum_{x\in\Lambda(u)} \pi(x)\right)\left(\sum_{y\in\mathcal{X}} \pi(y)\right) \\
&= \max_{(u,v):v=g(u)} \frac{\pi(\Lambda(u))}{\pi(u)\mathbf{P}(u,v)} \\
&\leq \max_{(u,v):v=g(u)} \frac{1}{(1-p^{-(t_2-t_4)})\mathbf{P}(u,v)}.
\end{aligned}$$

[1] If multiple states tie, we randomly pick one of them.

The last inequality holds by the following. Let $g^{-k}(u) = \{x \in \mathcal{X} \colon g^k(x) = u, g^{k-1}(x) \neq u\}$ for $k \in \mathbb{N}$, then $\Lambda(u) = \uplus_{k\in\mathbb{N}} g^{-k}(u)$. By Condition (ii) in Theorem 1, we have $|g^{-k}(u)| \leq p^{kt_4}$ which yields,

$$\frac{\pi(\Lambda(u))}{\pi(u)} = \sum_{k\in\mathbb{N}} \frac{\pi(g^{-k}(u))}{\pi(u)} \leq \sum_{k\in\mathbb{N}} p^{-k(t_2-t_4)} = \frac{1}{1-p^{-(t_2-t_4)}},$$

where we use Condition (i) and the definition of g, and $t_2 > t_4$ in Theorem 1.

Finally, we show that $\mathbf{P}(x, g(x)) \geq C' \frac{N}{p^{t_4}}$ for $x \neq x^*$ and for some universal constant $C' > 0$. Recall that (see also [27, Theorem 1])

$$\begin{aligned}
\mathbf{P}(x, g(x)) = N \sum_{y_1,\ldots,y_{N-1}} \sum_{x_1^\star,\ldots,x_{N-1}^\star} \frac{w(g(x)\,|\,x)}{w(g(x)\,|\,x) + \sum_{j=1}^{N-1} w(y_j\,|\,x)} \min\left\{1, \frac{w(g(x)\,|\,x) + \sum_{j=1}^{N-1} w(y_j\,|\,x)}{w(x\,|\,g(x)) + \sum_{j=1}^{N-1} w(x_j^\star\,|\,g(x))}\right\} \times \\
\mathbf{K}_{\mathrm{RW}}(x, g(x))\mathbf{K}_{\mathrm{RW}}(x, y_1)\cdots\mathbf{K}_{\mathrm{RW}}(x, y_{N-1})\mathbf{K}_{\mathrm{RW}}(g(x), x_1^\star)\cdots\mathbf{K}_{\mathrm{RW}}(g(x), x_{N-1}^\star) \\
\geq \left(N \sum_{y_1,\ldots,y_{N-1}} \frac{w(g(x)\,|\,x)}{w(g(x)\,|\,x) + \sum_{j=1}^{N-1} w(y_j\,|\,x)} \mathbf{K}_{\mathrm{RW}}(x, g(x))\mathbf{K}_{\mathrm{RW}}(x, y_1)\cdots\mathbf{K}_{\mathrm{RW}}(x, y_{N-1})\right) \times \\
\left(\sum_{x_1^\star,\ldots,x_{N-1}^\star} \min\left\{1, \frac{w(g(x)\,|\,x)}{w(x\,|\,g(x)) + \sum_{j=1}^{N-1} w(x_j^\star\,|\,g(x))}\right\} \mathbf{K}_{\mathrm{RW}}(g(x), x_1^\star)\cdots\mathbf{K}_{\mathrm{RW}}(g(x), x_{N-1}^\star)\right),
\end{aligned}$$

and we denote the first and the second terms of the right-hand side by

$$\mathbf{K}(x, g(x)) = N \sum_{y_1,\ldots,y_{N-1}} \frac{w(g(x)\,|\,x)}{w(g(x)\,|\,x) + \sum_{j=1}^{N-1} w(y_j\,|\,x)} \mathbf{K}_{\mathrm{RW}}(x, g(x))\mathbf{K}_{\mathrm{RW}}(x, y_1)\cdots\mathbf{K}_{\mathrm{RW}}(x, y_{N-1}), \tag{8}$$

$$\eta(x, g(x)) = \sum_{x_1^\star,\ldots,x_{N-1}^\star} \min\left\{1, \frac{w(g(x)\,|\,x)}{w(x\,|\,g(x)) + \sum_{j=1}^{N-1} w(x_j^\star\,|\,g(x))}\right\} \mathbf{K}_{\mathrm{RW}}(g(x), x_1^\star)\cdots\mathbf{K}_{\mathrm{RW}}(g(x), x_{N-1}^\star). \tag{9}$$

Hence, we have $\mathbf{P}(x, g(x)) \geq \mathbf{K}(x, g(x))\eta(x, g(x))$. We remark that the formulation of $\mathbf{K}(x, g(x))$ in (8) is based on the exchangeability, where the N-th trial state $y_N = g(x)$ is assumed to be selected as a proposal state and its probability is multiplied by N. However, we will not utilize such exchangeability to calculate a lower bound of $\mathbf{K}(x, g(x))$. Instead, we define an event A that selects $g(x)$ as the proposal from Step 1 and Step 2 in Algorithm 1, so that $\mathbb{P}(A) = \mathbf{K}(x, g(x))$. We aim to lower bound the $\mathbb{P}(A)$ by using the law of total probability. To this end, we introduce the event F for Step 1 in Algorithm 1 that we include the state $g(x)$ at least once among the N trials while we don't sample any "high" posterior states of the neighborhood of x for the rest of trials, i.e. they do not belong to the set $\mathcal{S}(x)$. Using conditional probability rule, the probability of the event F is equal to

$$\mathbb{P}(F) = \left(\frac{|\mathcal{N}(x)| - |\mathcal{S}(x)| + 1}{|\mathcal{N}(x)|}\right)^N \left(1 - \left(\frac{|\mathcal{N}(x)| - |\mathcal{S}(x)|}{|\mathcal{N}(x)| - |\mathcal{S}(x)| + 1}\right)^N\right), \tag{10}$$

where we take the probability over the uniform samples in Step 1 in Algorithm 1. Using the inequality $a^n - b^n = (a-b)(a^{n-1} + a^{n-2}b + \cdots + b^{n-1}) \geq (a-b)nb^{n-1}$ for any $a \geq b \geq 0$, we find that the lower bound of $\mathbb{P}(F)$ can be obtained by

$$\begin{aligned}
\mathbb{P}(F) &= \left(\left(1 - \frac{|\mathcal{S}(x)| - 1}{|\mathcal{N}(x)|}\right)^N - \left(1 - \frac{|\mathcal{S}(x)|}{|\mathcal{N}(x)|}\right)^N\right) \\
&\geq \frac{N}{|\mathcal{N}(x)|}\left(1 - \frac{|\mathcal{S}(x)|}{|\mathcal{N}(x)|}\right)^{N-1} \\
&\geq \frac{N}{p^{t_4}}\left(1 - (N-1)\frac{s_0}{p^{t_3}}\right) \\
&\geq \frac{N}{p^{t_4}}(1 + o(1)),
\end{aligned}$$

where Condition (ii) and Bernoulli's inequality are used in the second inequality, and Condition (iii) is used in the last inequality. We further define F_k as the event F with k number of $g(x)$ among N trials. (Note that since we sample trials with replacement in Step 1 of Algorithm 1, we may sample $g(x)$ multiple times.) Observe that $F = \uplus_{k=1}^{N} F_k$. Given the event F_k, the probability to select $g(x)$ in (2) of Algorithm 1 is upper bounded by

$$\begin{aligned}
\mathbb{P}(A \mid F_k) &= \frac{kw(g(x) \mid x)}{kw(g(x) \mid x) + \sum_{i=1}^{N-k} w(y_j \mid x)} \\
&= \left\{ 1 + k^{-1} \sum_{i=1}^{N-k} \frac{w(y_j \mid x)}{w(g(x) \mid x)} \right\}^{-1} \\
&\overset{(\star)}{\geq} \left\{ 1 + N \frac{h(p^{t_1+t_4-t_3})}{h(p^{t_2+t_3-t_4})} \right\}^{-1} = 1 + o(1).
\end{aligned}$$

To see $(\star)$, we have used the fact that h is a non-decreasing function, Condition (iii), and we have

$$\begin{aligned}
\frac{\pi(y_j)}{\pi(x)} \cdot \frac{\mathbf{K}_{\mathrm{RW}}(y_j, x)}{\mathbf{K}_{\mathrm{RW}}(x, y_j)} &\leq p^{t_1} p^{t_4 - t_3}, \\
\frac{\pi(g(x))}{\pi(x)} \cdot \frac{\mathbf{K}_{\mathrm{RW}}(g(x), x)}{\mathbf{K}_{\mathrm{RW}}(x, g(x))} &\geq p^{t_2} p^{t_3 - t_4},
\end{aligned}$$

from Condition (i) and (ii). Note that the right-hand side of the inequality $(\star)$ does not depend on k. Using the law of total probability, we yield the lower bound of $\mathbb{P}(A) = \mathbf{K}(x, g(x))$ by combining the previous results:

$$\begin{aligned}
\mathbf{K}(x, g(x)) &= \mathbb{P}(A|F)\mathbb{P}(F) + \mathbb{P}(A|F^c)\mathbb{P}(F^c) \geq \mathbb{P}(A|F)\mathbb{P}(F) \\
&= \sum_{k=1}^{N} \mathbb{P}(A|F_k)\mathbb{P}(F_k) \geq \sum_{k=1}^{N} (1 + o(1))\mathbb{P}(F_k) \\
&= (1 + o(1))\mathbb{P}(F) \geq \frac{N}{p^{t_4}}(1 + o(1)),
\end{aligned} \tag{11}$$

Similarly, we can calculate the lower bound of $\eta(x, g(x))$. We consider the event G that we don't select any "high" posterior states of the neighborhood of $g(x)$ for $N - 1$ trials, that is, any of them are not in the set $\mathcal{S}(g(x))$. A simple calculation yields

$$\mathbb{P}(G) \geq \left(\frac{|\mathcal{N}(g(x))| - s_0}{|\mathcal{N}(g(x))|} \right)^{N-1}.$$

Under the event G, on the other hand,

$$\begin{aligned}
\frac{w(g(x) \mid x)}{w(x \mid g(x)) + \sum_{j=1}^{N-1} w(x_j^\star \mid g(x))} &= \left(\frac{w(x \mid g(x))}{w(g(x) \mid x)} + \sum_{j=1}^{N-1} \frac{w(x_j^\star \mid g(x))}{w(g(x) \mid x)} \right)^{-1} \\
&= \left(\frac{\pi(x)\mathbf{K}_{\mathrm{RW}}(x, g(x))}{\pi(g(x))\mathbf{K}_{\mathrm{RW}}(g(x), x)} + \sum_{j=1}^{N-1} \frac{w(x_j^\star \mid g(x))}{w(g(x) \mid x)} \right)^{-1} \\
&\geq \left(p^{-t_2-t_3+t_4} + (N-1) \frac{h(p^{t_1-t_3+t_4})}{h(p^{t_2+t_3-t_4})} \right)^{-1},
\end{aligned}$$

where the second equality is due to the property of the balancing function $h(u) = uh(1/u)$ and the last inequality follows from a similar argument as before using Conditions (i), (ii) and

$$\frac{\pi(x_j^\star)}{\pi(g(x))} \cdot \frac{\mathbf{K}_{\mathrm{RW}}(x_j^\star, g(x))}{\mathbf{K}_{\mathrm{RW}}(g(x), x_j^\star)} \leq p^{t_1} p^{t_4 - t_3},$$

by non-decreasing h. Then for $x \neq x^*$, by Conditions (ii), (iii) and we have

$$\begin{aligned}
\eta(x, g(x)) &\geq \min\left\{1, \left(p^{-t_2-t_3+t_4} + (N-1)\frac{h(p^{t_1-t_3+t_4})}{h(p^{t_2+t_3-t_4})}\right)^{-1}\right\}\left(\frac{|\mathcal{N}(g(x))| - s_0}{|\mathcal{N}(g(x))|}\right)^{N-1} \\
&\geq (1+o(1))(1-s_0/p^{t_3})^{N-1} \\
&\geq 1 - \frac{Ns_0}{p^{t_3}} + o(1) \\
&= 1 + o(1). \qquad (12)
\end{aligned}$$

Combining the lower bounds (11) and (12) leads to $\mathbf{P}(x, g(x)) \geq C' \frac{N}{p^{t_4}}$, which concludes the proof of the theorem.

A.4 An example on a different weight function.

The example below shows an undesirable behavior of the MTM algorithm if we use a weight function which is not in the class of (4). Here an undesirable behavior means that the acceptance probability is close to zero even when $\pi(x_{\text{prop}}) \gg \pi(x_{\text{curr}})$, where we denote $x_{\text{curr}}, x_{\text{prop}}$ as the current state and the proposed state from (2) in Algorithm 1, respectively.

Example 1. We consider the weight function $w(y \,|\, x) = \pi(y)$. For the sake of simplicity, let $|\mathcal{N}(x)| = p^{t_3}$ and for $x' \in \mathcal{N}(x)$,

$$\frac{\pi(x')}{\pi(x)} = \begin{cases} p^{t_2}, & \text{for } x' \in \mathcal{S}(x), \\ p^{t_1/2}, & \text{for } x' \notin \mathcal{S}(x), \end{cases}$$

for all $x \neq x^*$. Let $x_j, x_j^\star$ be uniform samples from $\mathcal{N}(x_{\text{curr}})$ and $\mathcal{N}(x_{\text{prop}})$, respectively for $j \in [N-1]$. Assume $x_{\text{prop}} \in \mathcal{S}(x_{\text{curr}})$ to reflect $\pi(x_{\text{prop}}) \gg \pi(x_{\text{curr}})$. Although the conditions on Theorem 1 are met (except for those related to a balancing function h), the acceptance probability from x_{curr} to x_{prop} is upper bounded as

$$\begin{aligned}
\alpha(x_{\text{curr}}, x_{\text{prop}}) &= \min\left\{1, \frac{\sum_{i=1}^{N-1}\pi(x_i) + \pi(x_{\text{prop}})}{\sum_{i=1}^{N-1}\pi(x_i^\star) + \pi(x_{\text{curr}})}\right\} \\
&\leq \frac{Np^{t_2}\pi(x_{\text{curr}})}{(1+(N-1)p^{t_2+t_1/2})\pi(x_{\text{curr}})} \\
&\leq \frac{2}{p^{t_1/2}} = o(1).
\end{aligned}$$

Notice that the acceptance probability $\alpha(x_{\text{curr}}, x_{\text{prop}}) = 1$ in the MH algorithm.

A.5 A toy example on path construction.

We provide a simple example on how to construct the canonical path ensemble Δ^* described in A.3. Define $\mathcal{X} = \{0,1\}^3$ as our state space and let $x^* = (1,1,0)$ be the mode of the distribution on $\mathcal{X}$. Let the target distribution $\pi(x) \propto \exp(-d_{\mathrm{H}}(x, x^*))$, where d_{H} is a Hamming distance. We specify a neighborhood of any $x \in \mathcal{X}$ as $\mathcal{N}(x) = \{y \in \mathcal{X} : d_{\mathrm{H}}(y, x) = 1\}$, where d_{H} is a Hamming distance. In the left panel of Figure 5, neighboring states are linked to black undirected edges. We follow the rule described in (7) to define a transition function $g : \mathcal{X} \to \mathcal{X}$ and black directed edges indicate the defined moves by a function g. The height of each bar indicates $\pi(x)$ associated with the corresponding state x. In the right panel of Figure 5, we provide three examples to illustrate the three possible cases to construct a path from the defined transition g, described in A.3.

- The red directed edges indicate a path from $(0,0,1)$ to $(1,0,0)$, which corresponds to the case (i), since $(1,0,0) = g^2((0,0,1))$,
- The green directed edge indicates a path from $(1,1,0)$ to $(0,1,0)$, which corresponds to the case (ii), since $(1,1,0) = g((0,1,0))$,
- The blue directed edges indicate a path from $(0,0,1)$ to $(1,1,1)$, which corresponds to the case (iii).

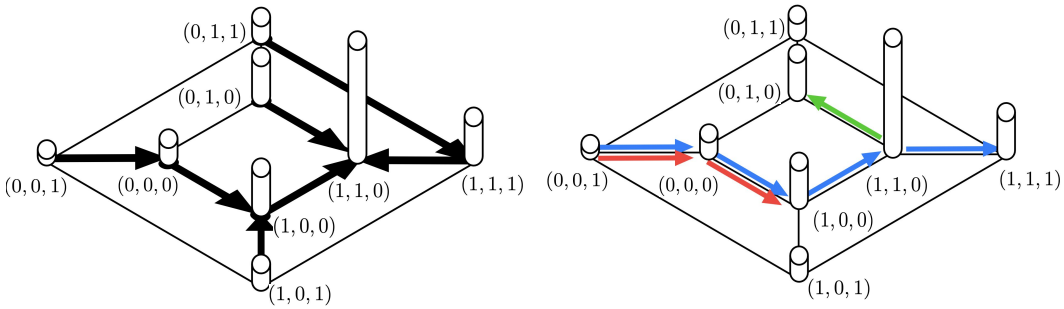

Figure 5: A toy example on path construction. Black undirected edges connect neighboring states and the target distribution π is represented as the heights of the cylinders. (Left) Black directed edges indicate the defined moves by a function g. (Right) The colored paths exemplify the three possible cases of path construction.

B Details of simulation studies

The scope of this paper is to theoretically study the mixing time for the family of MTM algorithms, and hence we mainly focus on experiments to empirically verify our theoretical insights, that the MTM mixing time is smaller by a factor of the number of trials N and that locally balanced weight functions tend to perform better under suitable assumptions. Nevertheless, in some experiments, we compare the MTM algorithm with the locally balanced MH algorithm (denoted as LBMH) [48], as it has been reported to outperform the other state-of-the-art methods. Before describing the details of simulation studies, here we briefly describe the locally balanced MH algorithm.

Specification of LBMH requires balancing function h and uninformed symmetric distribution $\mathbf{K}_{\mathrm{sym}}(x,\cdot)$ supported on $\mathcal{N}(x)$. LBMH chooses a proposal state y from a pointwise informed proposal distribution

$$Q_h(x,y) = (Z_h(x))^{-1} h\left(\pi(y)/\pi(x)\right) \mathbf{K}_{\mathrm{sym}}(x,y), \tag{13}$$

where $Z_h(x) = \sum_{z\in\mathcal{N}(x)} h\left(\pi(z)/\pi(x)\right) \mathbf{K}_{\mathrm{sym}}(x,z)$ is a normalizing constant. Then, y is accepted with probability $\alpha = \min\{1, \frac{\pi(y)Q_h(y,x)}{\pi(x)Q_h(x,y)}\} = \min\{1, \frac{Z_h(x)}{Z_h(y)}\}$, by defintion of balancing function $h(u) = uh(1/u)$ and symmetry of $\mathbf{K}_{\mathrm{sym}}$. Unlike MTM where a subset of $\mathcal{N}(x)$ is selected (with replacement) as a trial and choose a proposal among them, LBMH needs to evaluate $h(\pi(y)/\pi(x))$ for *all* $y \in \mathcal{N}(x)$ to get a proposal state y which can be viewed as an exhaustive search of $\mathcal{N}(x)$. In terms of computation, MTM requires calculating $2N-1$ weight functions at each iteration where N can be chosen at one's disposal, LBMH requires calculating $|\mathcal{N}(y)|$ number of ratios to calculate $Z_h(y)$ at each iteration, where proposal probabilities $\{Q_h(x,y) : y \in \mathcal{N}(x)\}$ and normalizing constant $Z_h(x)$ can be saved and reused from the previous iteration. Since random walk proposals in BVS and SBM examples are both symmetric, we compare MTM with LBMH by letting $\mathbf{K}_{\mathrm{sym}} = \mathbf{K}_{\mathrm{RW}}$ with three different balancing functions: $h(u) = \sqrt{u}$, $h(u) = \min\{1,u\}$ and $h(u) = \max\{1,u\}$ (corresponding to $w_{\mathrm{sqrt}}, w_{\mathrm{min}}, w_{\mathrm{max}}$ respectively).

B.1 Details of Bayesian variable selection (BVS)

After marginalizing out $\boldsymbol{\beta}$ and ϕ, the posterior distribution $\pi(\gamma \,|\, \boldsymbol{y})$ is written as [47, §A.1]

$$\pi(\gamma \,|\, \boldsymbol{y}) = C \cdot \frac{1}{p^{\kappa|\gamma|}(1+\mathscr{G})^{|\gamma|/2}} \mathrm{SSR}(\gamma)^{-n/2} \mathbb{1}(|\gamma| \le s_{\max}), \tag{14}$$

where $\mathrm{SSR}(\gamma) = \boldsymbol{y}^\top \left(\mathbf{I}_n - \frac{\mathscr{G}}{\mathscr{G}+1}\boldsymbol{X}_\gamma \left(\boldsymbol{X}_\gamma^\top \boldsymbol{X}_\gamma\right)^{-1} \mathbf{X}_\gamma^\top\right) \boldsymbol{y}$ is a term having a similar role as a sum of squared residuals and C is a normalizing constant.

MCMC setup. Hyperparameters are specified as $\mathscr{G} = p^3 = 5000^3$, $\kappa = 2$, and $s_{\max} = 100$. For each dataset, we run a chain of 10^5 iteration for single-try MH, 2×10^4 iteration for MTM with $N = 5$, and 10^4 iteration for MTM with $N = 5, 10, 50, 100, 500, 1000, 2000, 5000$ using four different weight functions. Algorithms are randomly initialized with state γ_0 such that $\gamma_0 \cap \gamma^* = \emptyset$ and $d_{\mathrm{H}}(\gamma_0, \gamma^*) = 20$ which implies $H = 20$ is the minimum required hitting iteration. For each simulated dataset, the true data generated model achieves the highest posterior probability ($\gamma^* = x^*$). All simulation studies are performed on a Linux cluster with Intel(R) Xeon(R) Gold 6132 CPU @ 2.60GHz and 96GB memory.

Results from Table 2 show that H decreases roughly by a factor of N until $N = 100$, which confirms our theoretical findings, given that the model setting satisfies that the mixing time is equivalent to the hitting iteration up to constant factors [36]. When N becomes larger, the performance of unscaled weight function w_{ord} deteriorates and never converges when $N = 5000$. In contrast, locally balanced weight functions generally perform well even when N is large. Table 3 suggests that choosing moderate N is beneficial in terms of computational savings. When the design matrix is correlated and SNR$= 2$, the result suggests that the chain often stuck when we choose the weight function as $w_{\max}$. Since the shape of the posterior distribution becomes irregular when the SNR is intermediate [47] and design matrix is correlated, to get a more clear insight we further perform additional simulation study when the posterior distribution exhibits multimodality; see Appendix B.4. Finally, under different settings of the design and SNR, the median N estimated from Algorithm 2 using $\psi = 0.9$ over 50 replicate datasets is $\hat{N} = 349$ (indep, SNR=4), 501 (indep, SNR=2), 328 (dep, SNR=4) and 158 (dep, SNR=2).

Table 2: (BVS) Median of H, the number of iterations until the chain hit γ^* over 50 replicates. Entry with "Fail" indicates that chains never hit γ^* in more than half of the replicated datasets.

	SNR	N	1	5	10	50	100	500	1000	2000	5000	LBMH
ind.	4	w_{ord}	19414	3283	1742	392	203	100	211	3168	Fail	N/A
		w_{sqrt}		3340	1787	360	177	55	42	38	54	Fail
		$w_{\min}$		3365	1948	354	180	50	33	28	24	20
		$w_{\max}$		3246	1876	372	182	54	42	40	60	Fail
	2	w_{ord}	20088	3684	1865	392	213	89	137	1020	Fail	N/A
		w_{sqrt}		3666	1955	398	200	58	40	33	32	137
		$w_{\min}$		3928	2034	366	202	55	34	29	24	20
		$w_{\max}$		3696	2000	418	229	62	44	36	34	Fail
dep.	4	w_{ord}	21292	3898	1989	422	234	91	117	735	Fail	N/A
		w_{sqrt}		3977	2256	394	209	64	44	36	45	Fail
		$w_{\min}$		4196	2065	412	226	51	35	30	24	20
		$w_{\max}$		4360	2137	504	240	70	46	42	40	Fail
	2	w_{ord}	66020	7724	6324	1033	528	150	145	404	Fail	N/A
		w_{sqrt}		9458	4226	1088	660	180	97	59	72	Fail
		$w_{\min}$		8357	4363	1212	484	109	68	54	31	37
		$w_{\max}$		11541	7057	6794	6782	3668	3124	3729	7246	Fail

Table 3: (BVS) Median of T_H, wall-clock time (in seconds) until the chain hit γ^* over 50 replicates. Entry with "Fail" indicates that chains never hit γ^* in more than half of the replicated datasets.

	SNR	N	1	5	10	50	100	500	1000	2000	5000	LBMH
ind.	4	w_{ord}	1.30	0.80	0.42	0.12	0.07	0.07	0.27	6.95	Fail	N/A
		w_{sqrt}		0.81	0.46	0.11	0.07	0.04	0.05	0.09	0.30	Fail
		$w_{\min}$		0.89	0.53	0.12	0.07	0.04	0.04	0.06	0.11	0.07
		$w_{\max}$		0.88	0.53	0.13	0.07	0.04	0.05	0.09	0.33	Fail
	2	w_{ord}	1.25	0.82	0.43	0.11	0.07	0.06	0.15	1.93	Fail	N/A
		w_{sqrt}		0.81	0.44	0.12	0.07	0.04	0.04	0.06	0.13	0.39
		$w_{\min}$		1.01	0.55	0.11	0.08	0.04	0.04	0.06	0.11	0.07
		$w_{\max}$		0.92	0.51	0.13	0.08	0.05	0.05	0.06	0.13	Fail
dep.	4	w_{ord}	1.35	0.85	0.42	0.12	0.08	0.06	0.12	1.37	Fail	N/A
		w_{sqrt}		0.87	0.50	0.11	0.07	0.04	0.05	0.06	0.17	Fail
		$w_{\min}$		1.04	0.55	0.13	0.08	0.04	0.04	0.06	0.10	0.06
		$w_{\max}$		1.07	0.53	0.15	0.08	0.05	0.05	0.07	0.14	Fail
	2	w_{ord}	3.38	1.67	1.29	0.26	0.16	0.09	0.14	0.70	Fail	N/A
		w_{sqrt}		1.91	0.90	0.28	0.19	0.10	0.10	0.10	0.24	Fail
		$w_{\min}$		1.89	1.01	0.34	0.16	0.07	0.07	0.10	0.13	0.09
		$w_{\max}$		2.58	1.66	1.86	1.97	2.10	2.78	5.74	22.36	Fail

In contrast to MTM, LBMH fails to converge to γ^* when $h(u) = \sqrt{u}$ or $h(u) = \max\{1, u\}$. It is easier for MTM to escape from such local modes by randomly searching part of its neighborhood to select the proposal. The exhaustive search nature of LBMH makes it difficult to escape from the local mode since some high values of $\pi(y^*)/\pi(y)$, $y^* \in \mathcal{N}(y)$ involved in the denominator makes the acceptance ratio small. This phenomenon disappears when $h(u) = \min\{1, u\}$ is used. In terms of wall-clock hitting time T_H, LBMH is not as efficient as MTM with a smaller choice of N.

In addition, we also consider the case when SNR $= 0.5$ (very weak SNR) so that the null model $\gamma^* = \mathbf{0}$ receives the highest probability across all simulated datasets. For each replicated dataset, algorithms are randomly initialized with state γ_0 such that $d_{\mathrm{H}}(\gamma_0, \gamma^*) = 10$ which implies $H = 10$ is the minimum required hitting iteration. Table 4 provides the result similar to Tables 2 and 3, since the posterior distribution is unimodal with the peak at the null model $\gamma^* = \mathbf{0}$ due to the sparsity prior. The median N estimated from Algorithm 2 using $\psi = 0.9$ over 50 replicate datasets is $\hat{N} = 171$for independent design and is $\hat{N} = 212$ for dependent design.

Table 4: (BVS, very weak SNR $= 0.5$) Median of H and T_H over 50 replicates. Entry with "Fail" indicates that chains never hit γ^* in more than half of the replicated datasets.

	SNR $= 0.5$	N	1	5	10	50	100	500	1000	2000	5000
H	indep.	w_{ord}	14358	2708	1309	281	142	50	46	106	8968
		w_{sqrt}		2788	1402	285	148	37	20	16	12
		w_{min}		2680	1481	282	149	34	19	16	12
		w_{max}		2476	1276	266	131	38	23	17	12
	dep.	w_{ord}	12596	2924	1432	304	142	43	42	104	Fail
		w_{sqrt}		2776	1532	278	146	35	21	15	12
		w_{min}		2836	1354	260	154	34	22	14	11
		w_{max}		2562	1289	270	140	32	22	15	12
T_H	indep.	w_{ord}	0.89	0.56	0.3	0.07	0.04	0.03	0.05	0.19	39.38
		w_{sqrt}		0.6	0.32	0.07	0.05	0.02	0.02	0.03	0.04
		w_{min}		0.68	0.36	0.09	0.05	0.02	0.02	0.03	0.04
		w_{max}		0.62	0.35	0.08	0.05	0.02	0.02	0.03	0.04
	dep.	w_{ord}	0.67	0.57	0.29	0.07	0.04	0.03	0.04	0.17	Fail
		w_{sqrt}		0.56	0.31	0.07	0.04	0.02	0.02	0.02	0.04
		w_{min}		0.66	0.33	0.07	0.05	0.02	0.02	0.02	0.04
		w_{max}		0.58	0.31	0.08	0.04	0.02	0.02	0.02	0.04

B.2 Details of stochastic block model (SBM)

After marginalizing out $\{Q_{uv}\}_{1\le u\le v\le K}$, the posterior distribution $\pi(\mathbf{z}\,|\,\boldsymbol{A})$ is written as (see [24, §2.1] and [55, §2.2])

$$\pi(\mathbf{z}\,|\,\boldsymbol{A}) = C\cdot \prod_{1\le u\le v\le K} B(\kappa_1 + m_{uv}, \kappa_2 + \overline{m}_{uv})\cdot \mathbb{1}(\mathbf{z}\in S_\alpha), \tag{15}$$

where $B(\kappa_1,\kappa_2) = \Gamma(\kappa_1)\Gamma(\kappa_2)/\Gamma(\kappa_1+\kappa_2)$ is a beta function, C is a normalizing constant,

$$m_{uv} = \begin{cases} \sum_{i,j} A_{ij}\mathbb{1}(z_i = u, z_j = v) & \text{if } u < v, \\ \sum_{i<j} A_{ij}\mathbb{1}(z_i = u, z_j = u) & \text{if } u = v, \end{cases}$$

is the number of edges between blocks u and v, and using the notation $n_u(\mathbf{z}) = \sum_i \mathbb{1}(z_i = u)$,

$$\overline{m}_{uv} = \begin{cases} n_u(\mathbf{z})n_v(\mathbf{z}) - m_{uv} & \text{if } u < v, \\ n_u(\mathbf{z})(n_u(\mathbf{z})-1)/2 - m_{uu} & \text{if } u = v, \end{cases}$$

is the number of non-edges between blocks u and v. We note that $\pi(\mathbf{z}\,|\,\boldsymbol{A})$ is invariant of a label permutation.

Data generation. When $K = 2$, there are two true clusters (blocks) of nodes each with 500 nodes. When $K = 5$, there are five true clusters of nodes each with 200 nodes. We generated a graph from the homogeneous SBM and where within- and cross-community edge connection probabilities are a and b respectively. Specifically, for $K = 2$ we set $(a, b) = (0.222, 0.01)$ and $(a, b) = (0.07, 0.01)$ so that $\mathrm{CH} \approx 10$ and $\mathrm{CH} \approx 2$, and for $K = 5$ we set $(a, b) = (0.473, 0.01)$ and $(a, b) = (0.13, 0.01)$ so that $\mathrm{CH} \approx 10$ and $\mathrm{CH} \approx 2$. For each setting, we simulate 50 datasets.

MCMC setup. Hyperparameters are specified as $\kappa_1 = \kappa_2 = 1$, and $\alpha = 1000$ so that the size of the feasible set S_α is maximized. For each dataset, we run a chain of 10^5 iteration for single-try MH, 5×10^4 iteration for MTM with $N = 5$, and 2×10^4 iteration for MTM with $N = 5, 10, 50, 100, 500, 1000, 2000, 5000$ using four different weight functions. Algorithms are randomly initialized with state $\mathbf{z}_0$ such that $\tilde{d}_{\mathrm{H}}(\mathbf{z}_0, \mathbf{z}^*) = 400$ which implies $H = 400$ is the minimum required hitting iteration. For each simulated dataset, the true data generated model achieves the highest posterior probability ($\mathbf{z}^* = x^*$). All simulations are performed on a Linux cluster with Intel(R) Xeon(R) Gold 6132 CPU @ 2.60GHz and 96GB memory.

Results from Table 5 show that H decreases roughly by a factor of N until $N = 10$ for locally balanced weight functions, but not for unscaled weight function w_{ord}. Even when $N \ge 50$, MTM with w_{ord} never converges to the highest probability model, highlighting the necessity of the use of locally balanced weight function for the general model selection problems. When N is very large, the performance of locally balanced weight functions generally deteriorates, which matches with our theoretical findings regarding the rate condition on N. Table 6 also suggests that a moderate choice of N (in SBM case, around 10) is beneficial in terms of computation savings. Finally, under different settings of (K, CH), the median N estimated from Algorithm 2 using $\psi = 0.9$ over 50 replicate datasets is $\hat{N} = 15$ for (K, CH) = (2, 2), $\hat{N} = 8$ for (2, 10), $\hat{N} = 5$ for (5, 2) and $\hat{N} = 4$ for (5, 10).

Comparison of hitting iteration H with LBMH gives an insight similar to the BVS example. When $K = 5$, LBMH often gets stuck at a local mode and never converges to $\mathbf{z}^*$. However when $K = 2$, LBMH performs similarly to MTM with larger choices of N. We note that when $K = 2$, the shape of posterior distribution can be significantly different from that of $K = 5$, as the minimax rate and posterior contraction rate analysis are often treated separately when $K = 2$ and $K \ge 3$ [51, 55]. The comparison of wall-clock hitting time T_H also suggests MTM with moderate choice of N is much more efficient.

Table 5: (SBM) Median of H, the number of iterations until the chain hit $\mathbf{z}^*$ over 50 replicates. Entry with "Fail" indicates that chains never hit $\mathbf{z}^*$ in more than half of the replicated datasets.

	CH	N	1	5	10	50	100	500	1000	2000	5000	LBMH
$K = 2$	≈ 10	w_{ord}	11572	2495	5542	Fail	Fail	Fail	Fail	Fail	Fail	N/A
		w_{sqrt}		1603	1136	692	644	602	610	637	684	740
		w_{min}		1558	974	544	493	444	434	431	424	418
		w_{max}		1657	1142	762	708	726	830	1224	2842	6364
	≈ 2	w_{ord}	13343	2722	3484	Fail	Fail	Fail	Fail	Fail	Fail	N/A
		w_{sqrt}		1948	1432	874	818	808	703	700	728	682
		w_{min}		2244	1400	944	911	820	882	812	890	866
		w_{max}		1916	1354	851	774	701	696	681	680	709
$K = 5$	≈ 10	w_{ord}	25328	5852	4614	Fail	Fail	Fail	Fail	Fail	Fail	N/A
		w_{sqrt}		5400	3008	1210	987	858	874	992	1564	Fail
		w_{min}		5376	2695	1127	907	719	690	675	660	Fail
		w_{max}		5230	2977	1184	1022	905	960	1286	3895	Fail
	≈ 2	w_{ord}	25883	6388	4422	Fail	Fail	Fail	Fail	Fail	Fail	N/A
		w_{sqrt}		5552	2885	1067	805	628	574	542	506	Fail
		w_{min}		5426	3056	1168	966	802	775	752	740	1517
		w_{max}		5245	2904	1100	882	703	654	630	614	Fail

Table 6: (SBM) Median of T_H, wall-clock time (in seconds) until the chain hit $\mathbf{z}^*$ over 50 replicates. Entry with "Fail" indicates that chains never hit $\mathbf{z}^*$ in more than half of the replicated datasets.

	CH	N	1	5	10	50	100	500	1000	2000	5000	LBMH
$K = 2$	≈ 10	w_{ord}	0.73	0.50	1.32	Fail	Fail	Fail	Fail	Fail	Fail	N/A
		w_{sqrt}		0.33	0.27	0.39	0.62	1.64	2.75	5.04	13.49	1.95
		w_{min}		0.41	0.29	0.36	0.53	1.29	2.03	3.52	8.27	1.09
		w_{max}		0.42	0.35	0.49	0.75	2.05	3.75	9.90	54.05	16.13
	≈ 2	w_{ord}	0.82	0.54	0.83	Fail	Fail	Fail	Fail	Fail	Fail	N/A
		w_{sqrt}		0.4	0.35	0.48	0.75	1.49	2.15	3.78	9.30	1.28
		w_{min}		0.56	0.41	0.57	0.89	1.55	2.62	4.31	11.26	1.60
		w_{max}		0.47	0.39	0.51	0.75	1.29	2.07	3.57	8.35	1.35
$K = 5$	≈ 10	w_{ord}	1.97	1.35	1.39	Fail	Fail	Fail	Fail	Fail	Fail	N/A
		w_{sqrt}		1.25	0.92	0.97	1.38	4.01	7.36	15.71	60.95	Fail
		w_{min}		2.27	1.50	1.31	1.78	3.97	6.69	12.30	29.76	Fail
		w_{max}		2.21	1.52	1.39	2.00	4.94	9.31	23.38	174.89	Fail
	≈ 2	w_{ord}	1.88	1.41	1.24	Fail	Fail	Fail	Fail	Fail	Fail	N/A
		w_{sqrt}		1.22	0.84	0.81	1.05	2.31	3.84	6.89	15.90	Fail
		w_{min}		1.35	0.97	0.91	1.31	2.99	5.17	9.61	23.64	6.52
		w_{max}		1.33	0.89	0.83	1.16	2.61	4.36	7.93	19.18	Fail

B.3 Spatial clustering model (SCM)

We consider a spatial clustering problem for a given set of spatial locations $\mathcal{S} = \{\mathbf{s}_1, \ldots, \mathbf{s}_p\} \subset \mathbb{R}^2$ where the responses $\zeta(\mathbf{s}_i)$'s are observed. The goal of an SCM is to identify a spatially contiguous partition on $\mathcal{S}$, denoted by $\mathcal{P} = \{\mathcal{S}_1, \ldots, \mathcal{S}_K\}$, where $\mathcal{S}_j$'s are disjoint subsets of $\mathcal{S}$ whose union is $\mathcal{S}$, such that the responses within a cluster $\{\zeta(\mathbf{s}) : \mathbf{s} \in \mathcal{S}_j\}$ are identically distributed and have different means across clusters.

We follow [28] to adopt a probabilistic model for $\mathcal{P}$ that utilizes a spanning tree graph $\mathcal{T}$ on $\mathcal{S}$ (with p vertices and $p-1$ edges) as a "spatial order" of $\mathcal{S}$. The spanning tree $\mathcal{T}$ is chosen in a way that two locations connected by an edge are spatially proximate to each other. A partition $\mathcal{P}$ with K clusters can be defined by removing $K-1$ edges from $\mathcal{T}$. Specifically, the SCM we consider can be written as

$$\begin{aligned}
\zeta(\mathbf{s}_i)|\{\mu(\mathbf{s}_i)\}, \mathcal{P}, K, \sigma^2 &\overset{\text{ind}}{\sim} \mathsf{N}(\mu_j, \sigma^2), \quad \text{with } \mu(\mathbf{s}_i) = \mu_j \text{ if } \mathbf{s}_i \in \mathcal{S}_j, \text{ for } i \in [p], \\
\mu_j|\mathcal{P}, K, \sigma^2 &\overset{\text{iid}}{\sim} \mathsf{N}(0, \lambda^{-1}\sigma^2), \quad \text{for } j \in [K], \\
\sigma^2 &\sim \mathsf{InvGamma}(a_0/2, b_0/2), \\
\pi(\mathcal{P}|K) &\propto \mathbb{1}\{\mathcal{P} \text{ can be obtained by removing } K-1 \text{ edges from } \mathcal{T}\}, \\
\pi(K) &\propto (1-c_0)^K, \quad K = 1, \ldots, p,
\end{aligned} \tag{16}$$

where $a_0 > 0$, $b_0 > 0$, $0 \le c_0 < 1$, and $\lambda > 0$ are hyperparameters. See Figure 6a for an example of partition $\mathcal{P}$ obtained by cutting edges from the spanning tree $\mathcal{T}$.

Thanks to the conjugate priors, μ_j's and σ^2 can be analytically marginalized out, and hence the inference problem boils down to drawing samples from the (discrete) posterior distribution $\pi(\mathcal{P} \,|\, \text{data})$. Although [28] considered random spanning trees by assigning a prior distribution on $\mathcal{T}$, we stress that the main focus of this paper is the mixing time analysis on the posterior distribution of $\mathcal{P}$. Thus, following [26], we fix $\mathcal{T}$ as the Euclidean minimum spanning tree in a Delaunay triangulation graph on $\mathcal{S}$, otherwise we can only sample from the conditional distribution $\pi(\mathcal{P} \,|\, \mathcal{T}, \text{data})$ which complicates the mixing time analysis of our target distribution $\pi(\mathcal{P} \,|\, \text{data})$.

We consider the following proposal:

$$\mathbf{K}_{\text{RW}}(\mathcal{P}, \mathcal{P}') = 1/(p-1)\,\mathbb{1}_{\mathcal{N}_b(\mathcal{P}) \cup \mathcal{N}_d(\mathcal{P})}(\mathcal{P}'),$$

where $\mathcal{N}_b(\mathcal{P})$ is the set of all possible partitions obtained by splitting a cluster in $\mathcal{P}$ into two clusters by selecting a cut-edge of $\mathcal{T}$ and $\mathcal{N}_d(\mathcal{P})$ is the set of all possible partitions obtained by merging two neighboring (with respect to $\mathcal{T}$) clusters in $\mathcal{P}$. See [23, 28] for detailed discussion on how to perform an appropriate split or merge on $\mathcal{P}$ given a spanning tree $\mathcal{T}$.

Data generation. We generate $p = 1000$ uniform locations $\mathbf{s}_i \overset{\text{iid}}{\sim} \mathsf{Unif}([0,1]^2), i \in [p]$ and specify the true means $\{\mu(\mathbf{s}_i)\}$ as in Figure 6a. Responses are generated according to (16) with $\sigma = \sqrt{\text{Var}(\mu(\mathbf{s}))}/\text{SNR}$ and we simulate 50 replicate datasets under SNR $\in \{3, 10\}$ respectively.

MCMC setup. Following [28], we initialize the chain using the estimates from the spatially clustered coefficient model of [26]. However, this initialization does not guarantee the same minimum number of iterations required to hit the true partition $\mathcal{P}^{\text{true}}$ for different replicate datasets. For a fair comparison, throughout this subsection, we redefine H as the number of *extra* iterations until hit, which is the iterations until hit minus the minimum number of iterations required to reach $\mathcal{P}^{\text{true}}$. Hyperparameters are specified as $a_0 = b_0 = 1$, $c_0 = 0.5$, and $\lambda = 0.01$. We consider the number of trials $N \in \{5, 10, 100, 500, 1000\}$. For each replicate dataset, we run a chain of $10,000$ iterations for each MTM specification and a chain of $30,000$ for standard single-try MH. All simulation studies are performed on a Linux cluster with Intel(R) Xeon(R) Gold 6132 CPU @ 2.60GHz and 96GB memory.

Table 7 summarizes H and T_H of various weight functions and numbers of trials N. The distributions of H for the setting of SNR = 10 are provided in Figure 6b. The results from single-try MH are also included as a baseline. When SNR = 10, the proposed locally balanced weight functions, especially w_{sqrt} and $w_{\min}$, considerably outperform the ordinary weight function w_{ord} and the single-try MH, in the sense that the proposed ones can reach $\mathcal{P}^{\text{true}}$ by much fewer iterations when $N \in \{100, 500, 1000\}$. In contrast, the performance of w_{ord} deteriorates when $N \ge 100$ and it fails to reach $\mathcal{P}^{\text{true}}$ when $N = 500$ or 1000. For the proposed weight functions, the wall-clock time until hit T_H is minimized when $N = 100$, since the benefit of having fewer iterations until hit is offset by the computational cost of extra trials when N is large.

When SNR = 3, the chains never visit the true partition, possibly because $\mathcal{P}^{\text{true}}$ does not lead to the highest posterior probability. In this case, we redefine H and T_H to be the number of extra iterations and the wall-clock time, respectively, to reach the 0.99 Rand index neighborhood of $\mathcal{P}^{\text{true}}$, defined as

$$\mathcal{N}_{\text{Rand}}(\mathcal{P}^{\text{true}}) := \{\mathcal{P} : \text{Rand}(\mathcal{P}, \mathcal{P}^{\text{true}}) \geq 0.99\},$$

where $\text{Rand}(\cdot, \cdot)$ is the Rand index [39] measuring the proportion of agreements between two partitions. The findings on H and T_H for reaching $\mathcal{N}_{\text{Rand}}(\mathcal{P}^{\text{true}})$ are similar to the ones when SNR = 10.

Finally, the median N estimated from Algorithm 2 using $\psi = 0.9$ over 50 replicate datasets is $\hat{N} = 13$ when SNR = 10 and $\hat{N} = 21$ when SNR = 3.

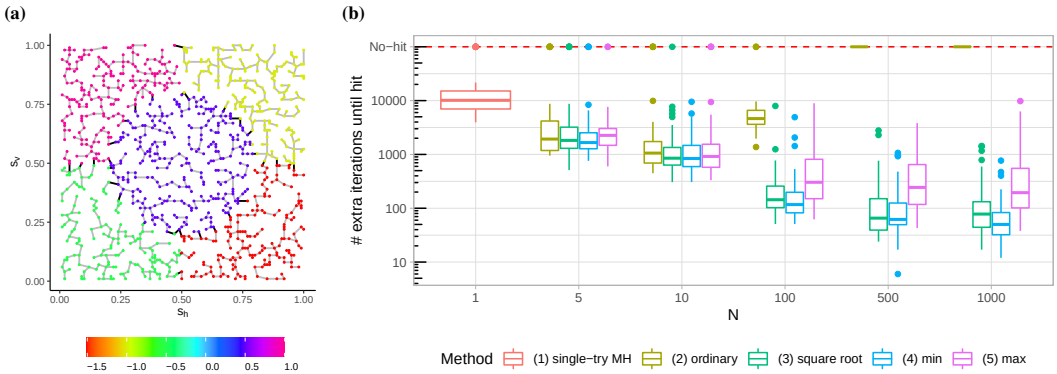

Figure 6: (a) True $\mu(\mathbf{s})$ and the Euclidean minimum spanning tree $\mathcal{T}$ on $\mathcal{S}$. Edges that should be removed in the true partition are marked in black. (b) Boxplot of numbers of extra iterations until hit for SCM under different numbers of trials and weight functions when SNR = 10.

Table 7: (SCM) Median of H and T_H (see text for definition; T_H in seconds) over 50 replicates. Entry with "Fail" indicates that chains never hit the target state in more than half of the replicated datasets.

		N	1	5	10	100	500	1000
H	SNR $= 10$	w_{ord}	10124	1933	1058	4638	Fail	Fail
		w_{sqrt}		1822	851	144	66	78
		w_{min}		1661	840	117	62	50
		w_{max}		2260	918	304	244	195
	SNR $= 3$	w_{ord}	23950	5083	2958	1046	Fail	Fail
		w_{sqrt}		2860	2064	235	84	126
		w_{min}		4443	2545	234	236	121
		w_{max}		6470	2810	394	462	301
T_H	SNR $= 10$	w_{ord}	906.72	164.28	90.75	1291.13	Fail	Fail
		w_{sqrt}		149.48	74.15	45.98	114.68	253.31
		w_{min}		182.53	94.21	42.74	117.88	206.83
		w_{max}		235.46	96.17	85.17	274.57	465.76
	SNR $= 3$	w_{ord}	2058.61	524.33	318.72	245.42	Fail	Fail
		w_{sqrt}		330.36	229.80	62.23	113.15	299.45
		w_{min}		486.18	264.62	58.71	233.03	288.35
		w_{max}		721.71	303.91	97.47	438.06	628.02

B.4 MTM algorithm on multimodal target distributions

In this section, we analyze the performance of the MTM algorithm with different choices of weight functions and N on the multimodal target distribution. Following [54], we generate a multimodal dataset in the context of a Bayesian variable selection problem.

Data generation. We let sample size $n = 1000$ and number of variables $p = 5000$. Each row of design matrix is independently sampled from $\boldsymbol{x}_i \overset{\text{iid}}{\sim} \mathsf{N}(0, \Sigma)$ for $i = 1, \ldots, n$ where $\Sigma = \text{diag}(\Sigma_{20}, \ldots, \Sigma_{20})$ is block-diagonal. Each block Σ_{20} has dimension 20×20, and $(\Sigma_{20})_{jk} = \exp(-|j-k|/3)$. We generate true coefficient $\boldsymbol{\beta}^{\text{true}}$ by first sampling 100 indices $j_1, \ldots, j_{100}$ uniformly at random (without replacement) from $[p]$ and let $\beta_{j_\ell} \overset{\text{iid}}{\sim} \mathsf{N}(0, \sigma_\beta^2)$ for $\ell = 1, \ldots, 100$, and $\beta_k = 0$ if $k \notin \{j_1, \ldots, j_{100}\}$. Then the response vector $\boldsymbol{y}$ is generated from $\boldsymbol{y} \sim \mathsf{N}(\boldsymbol{X}\boldsymbol{\beta}^{\text{true}}, \mathbf{I}_n)$. We consider three settings of $\sigma_\beta = 0.3, 0.4, 0.5$ to simulate the coefficients and data. For each setting, we simulate 20 datasets.

MCMC setup. We use the same BVS model described in Section 4. Hyperparameters are specified as $\mathscr{G} = p = 5000$, $\kappa = 1$, and $s_{\max} = 100$. For each dataset, we run a chain of 10,000 iteration for MTM with $N = 50, 100, 500, 1000, 2000, 5000$ using four different weight functions. The first 2000 iterations are discarded since the behavior of the chain (e.g. acceptance ratio) during the burn-in stage may be different from the behavior of the chain which entered stationarity; see also Figure 7. for trace plots. Algorithms are all initialized with null model $\gamma_0 = \mathbf{0}$. All simulation studies are performed on a Linux cluster with Intel(R) Xeon(R) Gold 6132 CPU @ 2.60GHz and 96GB memory.

Since the target distribution is no longer unimodal, hitting iteration H and wall-clock hitting time T_H are not appropriate metrics to compare mixing performance. Instead, we use three different metrics to evaluate the quality of the mixing: 1) acceptance ratio, 2) the number of unique states visited by the chain, denoted by #(unique γ), and 3) ESS/Time, where ESS is the effective sample size calculated from the hamming distances $d_{\text{H}}(\hat{\gamma}_{\max}, \gamma_t)$, $t = 2001, \ldots, 10000$ from the maximum posterior state $\hat{\gamma}_{max}$ found in a chain, Time is wall-clock time usage, measured in seconds. The results are summarized in Table 8.

Figure 7: Examples of MTM trace plot using different weight functions and the number of trials N. (Left) Simulated data with $\sigma_\beta = 0.3$, (Right) Simulated data with $\sigma_\beta = 0.5$. Each row corresponds to $N = 50, 500, 5000$. All chains are initialized at the null model $\gamma_0 = \mathbf{0}$.

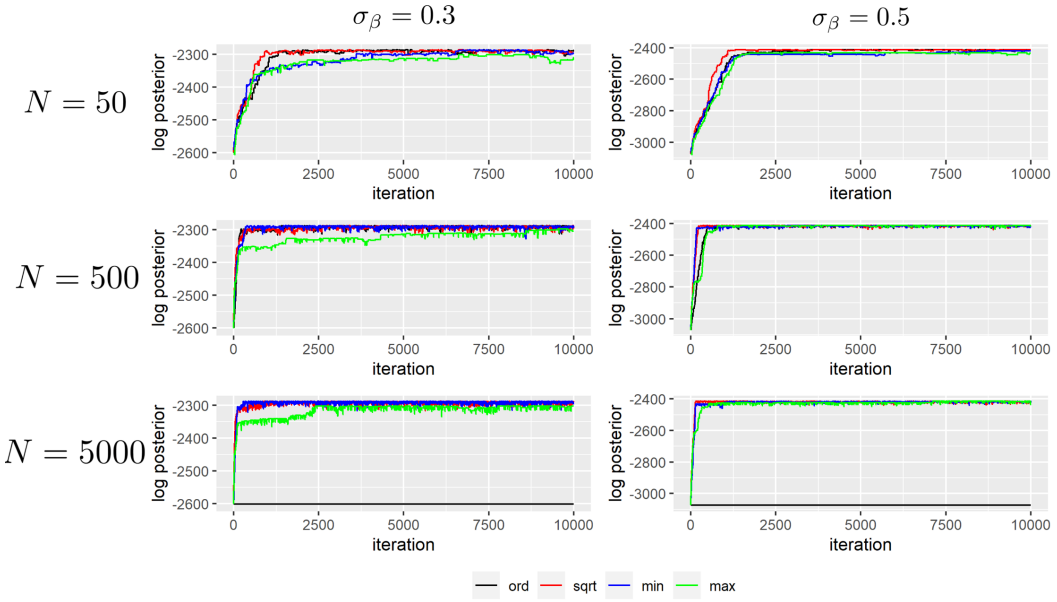

Figure 7 shows that as N increases, chains generally move faster towards the high posterior states which suggests that discarding the first 2000 samples is reasonable. One exception is when $N = 5000$,

Table 8: Multimodal posterior simulation results based on the acceptance rate, the number of unique states visited by a chain, and the effective sample size divided by the running time. All statistics are based on chains with a length of 8000 (2000 burn-in), and are averaged over 20 datasets.

		N	50	100	500	1000	2000	5000
$\sigma_\beta = 0.3$	Acc. Rate	w_{ord}	0.041	0.076	0.290	0.436	0.509	0.051
		w_{sqrt}	0.039	0.069	0.191	0.257	0.314	0.385
		w_{min}	0.042	0.075	0.299	0.430	0.576	0.717
		w_{max}	0.008	0.011	0.021	0.031	0.043	0.067
	#(unique γ)	w_{ord}	233.4	425.5	1398.3	2052.6	2130.8	161.8
		w_{sqrt}	240.2	385.1	993.2	1314.7	1549.6	1885.2
		w_{min}	255.1	408.4	1538.8	2056.4	2684.8	3168.8
		w_{max}	53.0	71.0	118.1	158.9	211.8	311.0
	ESS/Time	w_{ord}	7.04	8.71	9.32	6.12	2.27	0.02
		w_{sqrt}	5.11	8.08	7.14	4.62	2.48	1.14
		w_{min}	5.35	7.11	9.38	7.17	4.20	1.94
		w_{max}	1.20	1.12	0.49	0.41	0.13	0.06
$\sigma_\beta = 0.4$	Acc. Rate	w_{ord}	0.036	0.069	0.260	0.389	0.417	0.017
		w_{sqrt}	0.034	0.057	0.172	0.233	0.285	0.350
		w_{min}	0.036	0.064	0.261	0.404	0.541	0.674
		w_{max}	0.007	0.009	0.018	0.026	0.040	0.059
	#(unique γ)	w_{ord}	211.4	392.6	1302.9	1782.8	1852.1	72.4
		w_{sqrt}	205.8	316.1	932.5	1277.8	1470.2	1767.4
		w_{min}	219.8	361.6	1310.4	1919.6	2565.6	3067.3
		w_{max}	50.2	57.4	103.6	135.3	197.6	279.4
	ESS/Time	w_{ord}	6.13	8.06	6.31	6.58	0.52	0.03
		w_{sqrt}	4.79	6.81	6.06	4.04	2.10	1.09
		w_{min}	5.42	7.73	7.26	5.07	3.83	1.87
		w_{max}	1.04	0.80	0.35	0.23	0.11	0.05
$\sigma_\beta = 0.5$	Acc. Rate	w_{ord}	0.028	0.051	0.209	0.337	0.263	0.001
		w_{sqrt}	0.024	0.046	0.138	0.190	0.242	0.306
		w_{min}	0.026	0.055	0.209	0.337	0.481	0.633
		w_{max}	0.005	0.007	0.015	0.024	0.036	0.056
	#(unique γ)	w_{ord}	146.2	261.6	924.7	1382.0	976.5	5.6
		w_{sqrt}	132.2	246.4	648.8	877.4	1054.0	1323.6
		w_{min}	138.2	291.0	958.2	1445.6	2001.9	2444.1
		w_{max}	37.4	48.5	83.8	124.6	177.7	257.7
	ESS/Time	w_{ord}	5.93	9.05	9.85	7.77	0.32	0.01
		w_{sqrt}	6.75	6.83	7.33	5.05	3.27	1.17
		w_{min}	6.73	6.45	8.73	5.15	5.41	2.54
		w_{max}	1.97	1.00	0.53	0.39	0.15	0.07

the chain using ordinary weight function barely moves to another state, whereas the chain using locally balanced weight functions exhibits better mixing properties than $N = 50$ and $N = 500$.

Finally, Table 8 summarizes the multimodal simulation results. Similar to the previous unimodal results, the performance of w_{ord} deteriorates as N being large, especially when σ_β is large. When N is moderate, the weighting functions w_{ord} and w_{min} has a better mixing property than w_{sqrt}, while w_{max} being the worst. The inferior performance of w_{max} is also observed at Table 2. If the current state x is one of the local modes, we may have $\max\{1, \pi(x^\star)/\pi(y)\} \gg \max\{1, \pi(y)/\pi(x)\}$ where y is the proposed state, and $x^\star$ is one of the trials of y in Step 3 of Algorithm 1, which makes the acceptance probability very small so that the chain is stuck at x when using w_{max}. On the other hand, the magnitude of the difference between $h(\pi(x^\star)/\pi(y))$ and $h(\pi(y)/\pi(x))$ will be reduced if we use $w_{\text{sqrt}}, w_{\text{min}}$ so that the chain can traverse among local modes. To summarize, the simulation results suggest that w_{min} would be the best choice of the weight function if the multimodality exists, since it not only traverses the multimodal posterior efficiently, but also is robust to large N. We envision that there are a number of ways to improve the mixing under the multimodal posterior by combining it with techniques such as annealing or tempering [9].

C Real data applications

C.1 GWAS dataset for Bayesian variable selection

We consider a genome-wide association study (GWAS) dataset on glaucoma studied in [54] with sample size $n = 5418$ and number of genetic variants $p = 7255$. The response variable $\mathbf{y} \in \mathbb{R}^{5418}$ is the standardized cut-to-disk ratio measurements averaged over two eyes. We use the BVS model described in Section 4.1, with hyperparameters $\mathscr{G} = 100$ and $\kappa = 0.8$. Since the "true" state is not available, we compare the acceptance rate and the number of unique states visited, averaging over 5 chains. From Table 9, it is clear that the performance of w_{ord} deteriorates significantly as N grows whereas w_{sqrt}, w_{min}, w_{max} does not. We also report the posterior inclusion probabilities of the top 10 genetic variants in Tables 12, 13, 14, 15, 16, and 17 in Appendix E. All results generally agree with the result of [54], except that when we use w_{ord} with $N = 5000$, the chain is stuck at local modes and fails to find the significant genetic variants.

Table 9: GWAS dataset analysis results, averaged over 5 chains with random seeds.

	N	50	100	500	1000	2000	5000
	iteration	10^6	5×10^5	10^5	5×10^4	2×10^4	10^4
Acc. Rate	w_{ord}	0.4014	0.5081	0.1370	0.0471	0.0259	0.0085
	w_{sqrt}	0.3407	0.4812	0.7571	0.8325	0.8777	0.9251
	w_{min}	0.4136	0.5851	0.8252	0.8797	0.9138	0.9455
	w_{max}	0.2335	0.3404	0.6161	0.7199	0.7930	0.8698
#(unique states)	w_{ord}	199442	126238	6796	1172	259	43
	w_{sqrt}	169242	119396	37481	20588	8669	4563
	w_{min}	205459	145372	40993	21848	9080	4696
	w_{max}	115615	84124	30227	17591	7717	4211

C.2 Single-cell RNA dataset for structure learning

We consider a gene expression dataset on Alzheimer's disease used in [10] with the sample size $n = 1666$ and the number of genes $p = 73$. The goal is to learn the underlying directed acyclic graph (DAG) model among the p genes. Due to acyclicity, each DAG has at least one ordering of the nodes. For example, the ordering for the DAG $a \to b \leftarrow c$ can be either (a, c, b) or (c, a, b). A popular Bayesian structure learning strategy is to use MCMC sampling to first learn the marginal posterior distribution on the order space and then find one or multiple best DAGs for each sampled ordering.

We use an MTM implementation of the order MCMC sampler proposed in [10], which aims to learn the posterior distribution on the order space $\mathbb{S}^p$, the permutation group on $\{1, \ldots, p\}$. The size of our model space $\mathbb{S}^p$ is equal to $73! \approx 4.5 \times 10^{105}$. For each weight function and each setting, we simulate 30 chains, initialized at $(1, \ldots, 73)$. It is clear from Table 10 that the acceptance probability with ordinary weight function w_{ord} significantly deteriorates, which is consistent with our theory. We can see this tendency more clearly in the log-posterior trace plots for all weight functions in Figure 8.

Table 10: The single-cell RNA database for Alzheimer's disease analysis results, averaged over 30 chains with random seeds. The number in the parenthesis is the standard error.

	N	5	50
	iteration	5×10^2	2×10^2
Acc. Rate	w_{ord}	0.7187 (0.004)	0.0012 (0.000)
	w_{sqrt}	0.8029 (0.004)	0.9186 (0.003)
	w_{min}	0.8329 (0.002)	0.9506 (0.001)
	w_{max}	0.6643 (0.006)	0.6806 (0.008)
#(unique orderings)	w_{ord}	361.1 (3.1)	1.333333 (0.1)
	w_{sqrt}	402.3 (2.2)	184.3 (0.8)
	w_{min}	416.7 (1.8)	191.8 (0.6)
	w_{max}	332.0 (2.9)	136.2 (1.6)

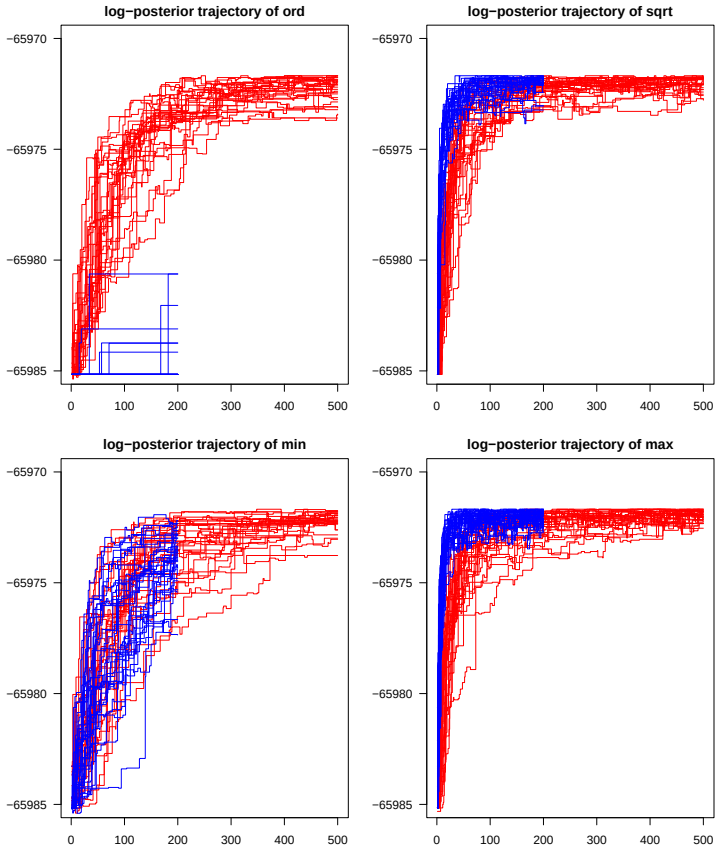

Figure 8: Log-posterior trace plots. Red trajectories indicate MTM with the number of trial $N = 5$ with 500 iterations, and blue trajectories indicate MTM with $N = 50$ with 200 iterations.

D Additional discussion

D.1 On parallelization (vectorization)

As discussed in *Remark 3* in Section 3.1, the overall theoretical computational complexity of the MTM algorithm until the convergence remains the same as a usual MH algorithm. However, MTM enables parallel computations when evaluating N weight functions and hence leads to a significant practical computational gain as evidenced by the reduced wall-clock hitting time reported in Table 1. Under the random walk proposal $\mathbf{K}_{\mathrm{RW}}$, the evaluation of weight functions is equivalent to the evaluation of target distribution at N states $y_1, \ldots, y_N$. Here we clarify that the scope of parallelism we consider is the instruction-level parallelism [41], also called vectorization. Thanks to modern When task-level parallelism, assigning a set of independent tasks in parallel across several processors, is employed to the MTM algorithm within each MCMC iteration, it suffers from communication overhead unless the evaluation of target distributions takes extremely long. Thanks to the optimized linear algebra libraries such as BLAS [7], the easiest way to achieve instruction-level parallelism is to convert the problem of evaluating target distribution at multiple states $\pi(y_1), \ldots, \pi(y_N)$ to a series of matrix multiplication problems.

Here we outline the computational strategy to simultaneously calculate $\pi(y_1), \ldots, \pi(y_N)$ for BVS and SBM. For BVS, since only one variable is added or deleted in the proposal, the Cholesky rank-1 update [17, 44] is utilized to get $\pi(y_1), \ldots, \pi(y_N)$ from $\pi(x)$. To be specific, assume $\gamma_1' \ldots, \gamma_N'$ are obtained by adding a variable from γ. By (14), the evaluation of $\pi(\gamma_j \,|\, \boldsymbol{y})$, $j = 1, \ldots, N$ corresponding to evaluating $\mathrm{SSR}(\gamma_j)$, $j = 1, \ldots, N$ from $\mathrm{SSR}(\gamma)$ saved from the previous iteration. We refer [49, Appendix B] for details of vectorization procedure with Cholesky rank-1 update. For SBM, let $A_i \in \{0,1\}^p$ be the ith column of adjacency matrix and $Z \in \{0,1\}^{p \times K}$ be one-hot encoded partition matrix such that $Z_{i,k} = 1$ if $z_i = k$ and 0 otherwise. Since a node is assigned to another block one at a time, the calculation of $\pi(y_1) \ldots, \pi(y_N)$ given current state $\pi(x)$ can be done by counting the change of the number of edges between blocks; see (15). Letting $A_J \in \{0,1\}^{p \times N}$ where column A_j corresponds to the jth proposal, the matrix-matrix multiplication $Z^\top A_J$ allows to calculate $\pi(y_1) \ldots, \pi(y_N)$ simultaneously from $\pi(x)$. In addition, if the graph is sparse, then sparse matrix multiplication algorithms can be utilized for further speedup.

D.2 On state space

Our state space of interest is finite (so discrete), but the proposed locally balanced MTM algorithm is also applicable to continuous state spaces which will be shown shortly. We choose to focus on the discrete case since the theory on continuous state spaces is usually developed under very different frameworks (and likewise, the theory on continuous spaces often cannot be readily applied to discrete ones). Indeed, developing MCMC theory or methodology on discrete spaces is often regarded as more challenging than on continuous ones [48, Section 1], due to the lack of gradient information and a widely accepted theoretical framework supported by statistical theory (for comparison, on continuous spaces, one often assumes log-concavity or asymptotic normality of the target posterior distribution).

To some extent, the proposed MTM method is conceptually similar to MALA (Metropolis adjusted Langevin algorithm) or HMC (Hamiltonian Monte Carlo) on continuous spaces in that MTM evaluates the "gradient" by a random search of neighboring states. This suggests that for continuous-state-space problems where the gradient of log-posterior cannot be easily evaluated (e.g. Bayesian inverse problems and Gaussian process regression models), the proposed MTM method can be quite useful.

We conclude this section with a simulation study that shows the weight function proposed in Proposition 2 can lead to an improved MTM algorithm on continuous spaces. Suppose our target distribution is the 10-dimensional Gaussian distribution $\mathsf{N}(0, \mathbf{I}_{10})$. We set our proposal distribution $q(\cdot|x) = \mathsf{N}_{10}(x, 10^{-2}\mathbf{I}_{10})$, initialize the chains at $x_0 = (10, 10, \ldots, 10)$, and run 10,000 iterations for each chain. The result is summarized in Table 11, where for each setting we repeat the simulation 30 times. The advantage of the weight functions considered in Proposition 2 over w_{ord} is substantial. We present the log-posterior traceplots in Figure 9 and the MCMC sample trajectories in Figure 10.

Table 11: Sampling from 10-d standard Gaussian distribution with 10,000 iterations using MTM. Averaged over 30 chains with random seeds. The number in the parenthesis is the standard error.

	N	MH ($N=1$)	10	100	1,000
Acc. Rate	w_{ord}	0.8604 (0.001)	0.7488 (0.002)	0.2971 (0.012)	0.0394 (0.008)
	w_{sqrt}		0.9667 (0.000)	0.9897 (0.000)	0.9967 (0.000)
	w_{min}		0.9656 (0.000)	0.9890 (0.000)	0.9963 (0.000)
	w_{max}		0.9593 (0.000)	0.9861 (0.000)	0.9950 (0.000)

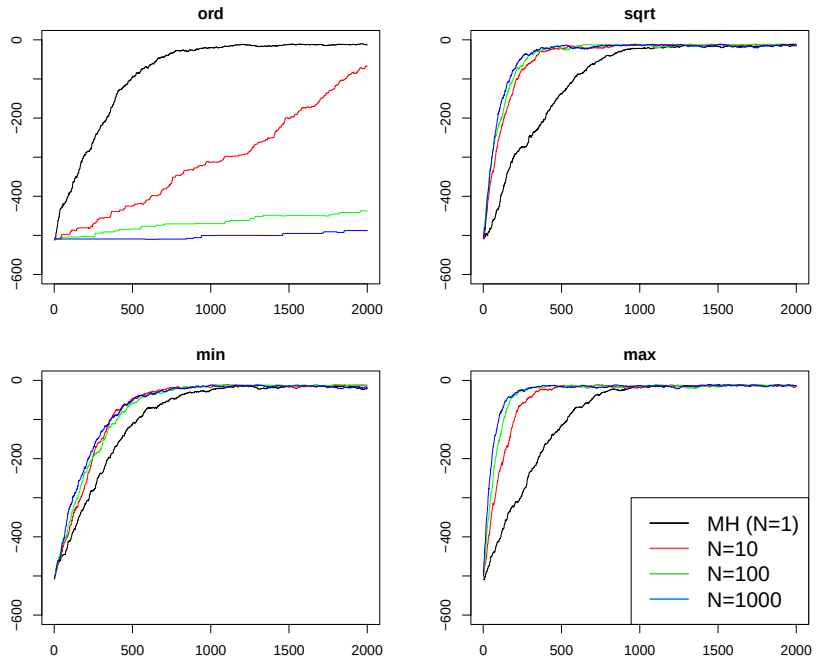

Figure 9: Log-posterior trace plots for 4 different weight functions $w_{\text{ord}}, w_{\text{sqrt}}, w_{\text{min}}$ and w_{max}. Different colors indicate a different number of trials as specified in the legend.

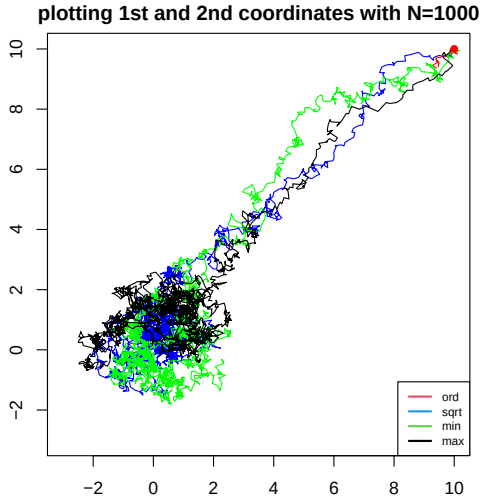

Figure 10: MCMC sample trajectories under 10-dimensional continuous target distribution $\mathsf{N}(\mathbf{0}, \mathbf{I}_{10})$ with different weight functions, initialized at $x_0 = (10, 10, \ldots, 10)$. Note that the chain with the ordinary weight function w_{ord} gets stuck at its early stage, whereas the chains with the other weight functions move to the region with a high posterior, which supports our claim.

E Additional tables

Here we report genetic variants with the top 10 highest posterior inclusion probabilities found from BVS model with MTM algorithm, under the different choices of $N = 50$, 100, 500, 1000, 2000, 5000 and weight functions $w_{\text{ord}}, w_{\text{sqrt}}, w_{\text{min}}$ and w_{max}.

Table 12: Genetic variants with top 10 posterior inclusion probability (PIP), obtained from MTM algorithm with $N = 50$ and averaged over 5 chains. Blue are genetic variants reported by [54].

	w_{ord}		w_{sqrt}		w_{min}		w_{max}	
	Name	PIP	Name	PIP	Name	PIP	Name	PIP
1	rs1063192	0.987	rs10483727	0.997	rs1063192	0.998	rs1063192	0.982
2	rs653178	0.979	rs1063192	0.979	rs653178	0.976	rs10483727	0.976
3	rs10483727	0.972	rs653178	0.978	rs10483727	0.974	rs653178	0.972
4	rs2275241	0.908	rs2275241	0.914	rs2275241	0.915	rs2275241	0.903
5	rs319773	0.806	rs319773	0.798	rs4557053	0.802	rs319773	0.827
6	rs4557053	0.77	rs4557053	0.773	rs319773	0.777	rs4557053	0.752
7	rs2369705	0.667	rs2369705	0.678	rs2369705	0.671	rs2369705	0.667
8	rs10491971	0.619	rs10491971	0.639	rs10491971	0.63	rs10491971	0.633
9	rs3177954	0.598	rs3177954	0.604	rs3177954	0.595	rs3177954	0.602
10	rs11087973	0.567	rs11087973	0.581	rs11087973	0.576	rs3843894	0.533

Table 13: Genetic variants with top 10 posterior inclusion probability (PIP), obtained from MTM algorithm with $N = 100$ and averaged over 5 chains. Blue are genetic variants reported by [54].

	w_{ord}		w_{sqrt}		w_{min}		w_{max}	
	Name	PIP	Name	PIP	Name	PIP	Name	PIP
1	rs10483727	1	rs10483727	0.993	rs10483727	0.988	rs10483727	1
2	rs1063192	0.996	rs1063192	0.99	rs653178	0.97	rs1063192	0.982
3	rs653178	0.973	rs653178	0.973	rs1063192	0.953	rs653178	0.975
4	rs2275241	0.918	rs2275241	0.922	rs2275241	0.914	rs2275241	0.925
5	rs319773	0.795	rs4557053	0.806	rs319773	0.802	rs319773	0.793
6	rs4557053	0.795	rs319773	0.78	rs4557053	0.792	rs4557053	0.765
7	rs2369705	0.673	rs2369705	0.663	rs2369705	0.662	rs2369705	0.63
8	rs10491971	0.626	rs10491971	0.624	rs10491971	0.617	rs11087973	0.604
9	rs3177954	0.618	rs3177954	0.62	rs3177954	0.598	rs10491971	0.602
10	rs11087973	0.583	rs11087973	0.597	rs11087973	0.587	rs11040978	0.563

Table 14: Genetic variants with top 10 posterior inclusion probability (PIP), obtained from MTM algorithm with $N = 500$ and averaged over 5 chains. Blue are genetic variants reported by [54].

	w_{ord}		w_{sqrt}		w_{min}		w_{max}	
	Name	PIP	Name	PIP	Name	PIP	Name	PIP
1	rs1063192	1	rs1063192	1	rs1063192	1	rs10483727	1
2	rs653178	0.992	rs10483727	0.983	rs10483727	1	rs653178	0.925
3	rs2275241	0.943	rs653178	0.979	rs653178	0.973	rs2275241	0.915
4	rs319773	0.857	rs2275241	0.922	rs2275241	0.887	rs319773	0.839
5	rs10483727	0.8	rs319773	0.791	rs319773	0.792	rs4557053	0.828
6	rs4557053	0.787	rs4557053	0.774	rs4557053	0.791	rs1063192	0.799
7	rs10491971	0.658	rs10491971	0.67	rs2369705	0.65	rs2369705	0.764
8	rs3843894	0.579	rs2369705	0.647	rs3177954	0.622	rs10491971	0.757
9	rs587409	0.553	rs3177954	0.624	rs10491971	0.615	rs12133371	0.617
10	rs3177954	0.546	rs11087973	0.613	rs11087973	0.561	rs11087973	0.605

Table 15: Genetic variants with top 10 posterior inclusion probability (PIP), obtained from MTM algorithm with $N = 1000$ and averaged over 5 chains. Blue are genetic variants reported by [54].

	w_{ord}		w_{sqrt}		w_{min}		w_{max}	
	Name	PIP	Name	PIP	Name	PIP	Name	PIP
1	rs653178	0.959	rs1063192	0.991	rs1063192	1	rs10483727	1
2	rs2275241	0.952	rs653178	0.982	rs10483727	1	rs653178	0.999
3	rs319773	0.918	rs10483727	0.939	rs653178	0.994	rs2275241	0.886
4	rs10483727	0.8	rs2275241	0.927	rs2275241	0.918	rs2369705	0.842
5	rs2369705	0.737	rs4557053	0.866	rs319773	0.791	rs319773	0.841
6	rs10491971	0.713	rs319773	0.859	rs4557053	0.779	rs1063192	0.804
7	rs4557053	0.704	rs2369705	0.651	rs10491971	0.652	rs3177954	0.787
8	rs3177954	0.688	rs10491971	0.637	rs2369705	0.646	rs10491971	0.768
9	rs12133371	0.615	rs3177954	0.6	rs3177954	0.613	rs12133371	0.656
10	rs1063192	0.6	rs11087973	0.574	rs11087973	0.588	rs3843894	0.65

Table 16: Genetic variants with top 10 posterior inclusion probability (PIP), obtained from MTM algorithm with $N = 2000$ and averaged over 5 chains. Blue are genetic variants reported by [54].

	w_{ord}		w_{sqrt}		w_{min}		w_{max}	
	Name	PIP	Name	PIP	Name	PIP	Name	PIP
1	rs2275241	1	rs653178	0.983	rs1063192	0.998	rs1063192	1
2	rs653178	1	rs10483727	0.922	rs653178	0.971	rs10483727	0.997
3	rs319773	0.819	rs2275241	0.857	rs2275241	0.916	rs653178	0.986
4	rs11087973	0.805	rs1063192	0.829	rs319773	0.794	rs2275241	0.948
5	rs1063192	0.8	rs4557053	0.824	rs4557053	0.771	rs319773	0.87
6	rs4557053	0.755	rs10491971	0.775	rs10483727	0.724	rs2369705	0.675
7	rs1460509	0.737	rs319773	0.763	rs10491971	0.714	rs4557053	0.634
8	rs3843894	0.729	rs2369705	0.697	rs3177954	0.672	rs3177954	0.527
9	rs3858886	0.648	rs11040978	0.651	rs2369705	0.628	rs2567344	0.519
10	rs10483727	0.61	rs587409	0.612	rs3843894	0.576	rs11634375	0.472

Table 17: Genetic variants with top 10 posterior inclusion probability (PIP), obtained from MTM algorithm with $N = 5000$ and averaged over 5 chains. Blue are genetic variants reported by [54].

	w_{ord}		w_{sqrt}		w_{min}		w_{max}	
	Name	PIP	Name	PIP	Name	PIP	Name	PIP
1	rs2151280	0.4	rs1063192	1	rs1063192	1	rs10483727	1
2	rs10483727	0.367	rs10483727	1	rs10483727	1	rs653178	0.997
3	rs12457539	0.29	rs653178	0.983	rs653178	0.972	rs319773	0.901
4	rs10508818	0.255	rs319773	0.804	rs2275241	0.821	rs3177954	0.844
5	rs3858886	0.231	rs2369705	0.791	rs4557053	0.732	rs2275241	0.668
6	rs7995962	0.207	rs2275241	0.772	rs319773	0.709	rs12457539	0.649
7	rs12125527	0.2	rs3177954	0.732	rs3177954	0.662	rs1063192	0.6
8	rs2738755	0.2	rs587409	0.652	rs12133371	0.593	rs587409	0.6
9	rs6661853	0.2	rs10491971	0.636	rs11087973	0.569	rs4924156	0.51
10	rs9869577	0.2	rs4557053	0.611	rs587409	0.565	rs4236601	0.498