



Title	Adaptive Incremental Mixture Markov chain Monte Carlo
Authors(s)	Marie, Florian, Friel, Nial, Mira, Antonietta, Raftery, Adrian E.
Publication date	2019-06-07
Publication information	Marie, Florian, Nial Friel, Antonietta Mira, and Adrian E. Raftery. "Adaptive Incremental Mixture Markov Chain Monte Carlo." Taylor & Francis, June 7, 2019. https://doi.org/10.1080/10618600.2019.1598872 .
Publisher	Taylor & Francis
Item record/more information	http://hdl.handle.net/10197/10865
Publisher's statement	This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (http://creativecommons.org/licenses/by-nc-nd/4.0/), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.
Publisher's version (DOI)	10.1080/10618600.2019.1598872

Downloaded 2026-05-01 23:44:07

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



© Some rights reserved. For more information

Adaptive Incremental Mixture Markov chain Monte Carlo

Florian Maire^{1,2}, Nial Friel^{1,2}, Antonietta Mira^{3,4}, and Adrian E. Raftery⁵

¹*School of Mathematics and Statistics, University College Dublin*

²*The Insight Centre for Data Analytics*

³*Data Science Center, ICS, Università della Svizzera italiana*

⁴*DISAT, Università dell'Insubria*

⁵*Department of Statistics, University of Washington*

June 1, 2018

Abstract

We propose Adaptive Incremental Mixture Markov chain Monte Carlo (AIMM), a novel approach to sample from challenging probability distributions defined on a general state-space. While adaptive MCMC methods usually update a parametric proposal kernel with a global rule, AIMM locally adapts a semiparametric kernel. AIMM is based on an independent Metropolis-Hastings proposal distribution which takes the form of a finite mixture of Gaussian distributions. Central to this approach is the idea that the proposal distribution adapts to the target by locally adding a mixture component when the discrepancy between the proposal mixture and the target is deemed to be too large. As a result, the number of components in the mixture proposal is not fixed in advance. Theoretically, we prove that there exists a process that can be made arbitrarily close to AIMM and that converges to the correct target distribution. We also illustrate that it performs well in practice in a variety of challenging situations, including high-dimensional and multimodal target distributions.

keywords: Adaptive MCMC, Bayesian inference, Independence Sampler, Importance weight, Local adaptation.

1 Introduction

We consider the problem of sampling from a target distribution defined on a general state space. While standard simulation methods such as the Metropolis-Hastings algorithm (Metropolis et al., 1953; Hastings, 1970) and its many variants have been extensively studied, they can be inefficient in sampling from complex distributions such as those that arise in modern applications. For example, the practitioner is often faced with the issue of sampling from

distributions which contain some or all of the following: multimodality, very non-elliptical high density regions, heavy tails, and high-dimensional support. In these cases, standard Markov chain Monte Carlo (MCMC) methods often have difficulties, leading to long mixing times and potentially to Monte Carlo estimators with large asymptotic variance. The adaptive MCMC framework originally developed by Gilks and Wild (1992), Gilks et al. (1998) and Haario et al. (1999) can help overcome these problems. The theoretical properties of these algorithms have been extensively analyzed by Andrieu and Moulines (2006) and Roberts and Rosenthal (2007). Adaptive MCMC methods improve the convergence of the chain by tuning its transition kernel on the fly using knowledge of the past trajectory of the process. This learning process causes a loss of the Markovian property and the resulting stochastic process is therefore no longer a Markov chain.

Most of the adaptive MCMC literature to date has focused on updating an initial parametric proposal distribution. For example, the Adaptive Metropolis–Hastings algorithm (Haario et al., 1999, 2001), hereafter referred to as AMH, adapts the covariance matrix of a Gaussian proposal kernel, used in a Random Walk Metropolis–Hastings algorithm. The Adaptive Gaussian Mixtures algorithm (Giordani and Kohn, 2010; Luengo and Martino, 2013), hereafter referred to as AGM, adapts a mixture of Gaussian distributions, used as the proposal in an Independent Metropolis–Hastings algorithm.

When knowledge on the target distribution is limited, the assumption that a *good* proposal kernel can be found in a specific parametric family may lead to suboptimal performance. Indeed, a practitioner using these methods must choose, sometimes arbitrarily, (i) a parametric family and (ii) an initial set of parameters to start the sampler. However, poor choices of (i) and (ii) may hamper the adaptation and result in slow convergence.

In this paper, we introduce a novel adaptive MCMC method, called Adaptive Incremental Mixture Markov chain Monte Carlo (AIMM). This algorithm belongs to the general Adaptive Independent Metropolis class of methods, developed in Holden et al. (2009), in that AIMM adapts an independence proposal density. However, the adaptation process here is quite different from others previously explored in the literature. Although our objective remains to reduce the discrepancy between the proposal and the target distribution along the chain, AIMM proceeds without any global parameter updating scheme, as is the case with adaptive MCMC methods belonging to the framework developed in Andrieu and Moulines (2006).

Our idea is instead to add a *local* probability mass to the current proposal kernel, when a large discrepancy area between the target and the proposal is encountered. The local probability mass is added through a Gaussian kernel located in the region that is not sufficiently supported by the current proposal. The decision to increment the proposal kernel is based on the importance weight function that is implicitly computed in the Metropolis

acceptance ratio. We stress that, although seemingly similar to Giordani and Kohn (2010) and Luengo and Martino (2013), our adaptation scheme is local and semiparametric, a subtle difference that has important theoretical and practical consequences. In particular, in contrast to AGM, the approach which we develop does not assume a fixed number of mixture components. The AIMM adaptation strategy is motivated by the quantitative bounds achieved when approximating a density using the recursive mixture density estimation algorithm proposed in Li and Barron (2000). The main difference is that while in AIMM the sequence of proposals is driven by the importance weight function at locations visited by the AIMM Markov chain, it is constructed in Li and Barron (2000) by successively minimizing the entropy with respect to the target distribution, an idea which was also put forward in Cappé et al. (2008). As in most adaptive MCMC, AIMM requires the specification of some tuning parameters and we provide a default setting that works well in the examples we have considered. We also present a faster version of AIMM that guarantees that the proposal is not too costly to evaluate, a situation which occurs when the number of components in the incremental mixture gets large, especially if the state space is high dimensional.

Proving the ergodicity of adaptive MCMC methods is often made easier by expressing the adaptation as a stochastic approximation of the proposal kernel parameters (Andrieu and Moulines, 2006). AIMM cannot be analysed in this framework as the adaptation does not proceed with a global parameter update step. We do not study the ergodicity of the process in the framework developed by Holden et al. (2009), and also used in Giordani and Kohn (2010), as the Doeblin condition required there is essentially equivalent to assuming that the importance function is bounded above. This condition is generally hard to meet in most practical applications, unless the state space is finite or compact. Instead, our ergodicity proof relies on the seminal work by Roberts and Rosenthal (2007), which shows that (i) if the process adapts less and less (*diminishing adaptation*) and (ii) if any Markov kernel used by the process to transition reaches stationarity in a bounded time (*containment*), then the adaptive process is ergodic. We show that AIMM can be implemented in a way that guarantees diminishing adaptation, while the containment condition remains to be proven on a case by case basis, depending on the target distribution. Moreover, we show using the recent theory developed in Craiu et al. (2015) and Rosenthal and Yang (2017) that there exists a process, which can be made arbitrarily close to the AIMM process, which is ergodic for the target distribution at hand.

The paper is organized as follows. We start in Section 2 with a pedagogical example which shows how AIMM succeeds in addressing the pitfalls of some adaptive MCMC methods. In Section 3, we formally present AIMM, and study its theoretical properties in Section 4. Section 5 illustrates the performance of AIMM on three different challenging tar-

get distributions in high dimensions, involving two heavy-tailed distributions and a bimodal distribution. AIMM is also compared with other adaptive and non-adaptive MCMC methods. In Section 6 we discuss the connections with importance sampling methods, particularly Incremental Mixture of Importance Sampling (Raftery and Bao, 2010), from which AIMM takes inspiration.

2 Introductory example

We first illustrate some of the potential shortcomings of the adaptive methods mentioned in the introduction and outline how AIMM addresses them. We consider the following pedagogical example where the objective is to sample efficiently from a one-dimensional target distribution. Appendix C gives more details of the various algorithms which we compare with AIMM and specifies the parameters of these algorithms.

Example 1. *Consider the target distribution*

$$\pi_1 = (1/4)\mathcal{N}(-10, 1) + (1/2)\mathcal{N}(0, .1) + (1/4)\mathcal{N}(10, 1).$$

For this type of target distribution it is known that Adaptive Metropolis–Hastings (AMH) (Haario et al., 2001) mixes poorly since the three modes of π_1 are far apart, a problem faced by many non-independent Metropolis algorithms. Thus an adaptive independence sampler such as the Adaptive Gaussian Mixture (AGM) (Luengo and Martino, 2013) is expected to be more efficient. AGM uses the history of the chain to adapt the parameters of a mixture of Gaussians on the fly to match the target distribution. For AGM, we consider here two families of proposals: the set of mixtures of two and three Gaussians, referred to as AGM–2 and AGM–3, respectively. By contrast, our method, Adaptive Incremental Mixture MCMC (AIMM) offers more flexibility in terms of the specification of the proposal distributions and in particular does not set a fixed number of mixture components.

We are particularly interested in studying the tradeoff between the convergence rate of the Markov chain and the asymptotic variance of the MCMC estimators, which are known to be difficult to control simultaneously (Mira, 2001; Rosenthal, 2003). Table 1 gives information about the asymptotic variance, through the Effective Sample Size (ESS), and about the convergence of the chain, through the mean squared error (MSE) of the tail-event probability $\pi_1(X > 5)$. Further details of these performance measures are given in Appendix C (supplementary material).

The ability of AMH to explore the state space appears limited as it regularly fails to visit the three modes in their correct proportions after 20,000 iterations (see the MSE of

Table 1: Example 1, results for π_1 : Effective Sample Size (ESS) (the larger the better), Mean Squared Error of $\pi_1(X > 5)$ (MSE) (the smaller the better). Estimates obtained through 100 independent runs of the four methods, each of 20,000 iterations, discarding the first 10,000 iterations for burn in. For the ESS statistic, the mean and variance are provided.

	ESS	MSE ($\times 10^4$)
AMH	(.08, .004)	6030
AGM-2	(.12, .001)	120
AGM-3	(.51, .091)	76
AIMM	(.47, .004)	7

$\pi_1(X > 5)$ in Table 1). The efficiency of AGM depends strongly on the number of mixture components of the proposal family. In fact AGM-2 is comparable to AMH in terms of the average ESS, indicating that the adaptation failed in most runs. AIMM and AGM-3 both achieve a similar exploration of the state space and AGM-3 offers a slightly better ESS. An animation corresponding to this toy example can be found at [this link](#).

From this example we conclude that AGM can be extremely efficient provided some initial knowledge on the target distribution is available, for example, the number of modes, the location of the large density regions and so on. If a mismatch between the family of proposal distributions and the target occurs, inference can be jeopardized. Since misspecifications are typical in real world models where one encounters high-dimensional, multi-modal distributions and other challenging situations, it leads one to question the efficiency of AGM samplers. On the other hand, AMH seems more robust to *a priori* lack of knowledge of the target, but the quality of the mixing of the chain remains a potential issue, especially when high density regions are disjoint.

Initiated with an naive independence proposal, AIMM adaptively builds a proposal that approximates the target iteratively by adding probability mass to locations where the proposal is not well supported relative to the target; see Section 3.6 for details. As a result, very little, if any, information regarding the target distribution is needed, hence making AIMM robust to realistic situations. Extensive experimentation in Section 5 confirms this point.

3 Adaptive Incremental Mixture MCMC

We consider target distributions π defined on a measurable space $(\mathsf{X}, \mathcal{X})$ where X is an open subset of \mathbb{R}^d ($d > 0$) and \mathcal{X} is the Borel σ -algebra of X . Unless otherwise stated, the distributions we consider are dominated by the Lebesgue measure and we therefore denote the distribution and the density function (w.r.t the Lebesgue measure) by the same symbol. In this section, we introduce the family of adaptive algorithms referred to AIMM,

the acronym for Adaptive Incremental Mixture MCMC.

3.1 Transition kernel

AIMM belongs to the general class of Adaptive Independent Metropolis algorithms originally introduced by Gåsemeyr (2003) and studied by Holden et al. (2009). AIMM generates a stochastic process $\{X_n, n \in \mathbb{N}\}$ that induces a collection of independence proposals $\{Q_n, n \in \mathbb{N}\}$. At iteration n , the process is at X_n and attempts a move to $\tilde{X}_{n+1} \sim Q_n$ which is accepted with probability α_n . In what follows, $\{\tilde{X}_n, n \in \mathbb{N}^*\}$ denotes the sequence of proposed states that are either accepted or rejected.

More formally, AIMM produces a time inhomogeneous process whose transition kernel K_n (at iteration n) is the standard Metropolis–Hastings (M–H) kernel with independence proposal Q_n and target π . For any $(x, A) \in (\mathbf{X}, \mathcal{X})$, K_n is defined by

$$K_n(x; A) := \int_A Q_n(dx') \alpha_n(x, x') + \delta_x(A) \int_{\mathbf{X}} Q_n(dx') (1 - \alpha_n(x, x')). \quad (1)$$

In (1), α_n denotes the usual M–H acceptance probability for independence samplers, namely

$$\alpha_n(x, x') := 1 \wedge \frac{W_n(x')}{W_n(x)}, \quad (2)$$

where $W_n := \pi/Q_n$ is the importance weight function defined on \mathbf{X} . Central to our approach is the idea that the discrepancy between π and Q_n , as measured by W_n , can be exploited in order to adaptively improve the independence proposal. This has the advantage that W_n is computed as a matter of course in the M–H acceptance probability (2).

3.2 Incremental process

We assume that available knowledge about π allows one to construct an initial proposal kernel Q_0 , from which it is straightforward to sample. When π is a posterior distribution, a default choice for Q_0 could be the prior distribution. The initial proposal Q_0 is assumed to be relatively flat, in the spirit of a defensive distribution (Hesterberg, 1995). Initiated at Q_0 , a sequence of proposals $\{Q_n, n \in \mathbb{N}\}$ is produced by our algorithm. In particular, the proposal kernel adapts by adding probability mass where the discrepancy between Q_n and π is deemed too large.

The adaptation mechanism is driven by the random sequence $\{W_n(\tilde{X}_{n+1}), n \in \mathbb{N}\}$, which monitors the discrepancy between π and the proposals $\{Q_n, n \in \mathbb{N}\}$ at the proposed states $\{\tilde{X}_n, n \in \mathbb{N}^*\}$. Let $\bar{W} > 0$ be a user-defined parameter such that the proposal kernel Q_n is incremented upon the event \mathcal{E}_n defined as:

$$\mathcal{E}_n := \left\{ W_n(\tilde{X}_{n+1}) > \bar{W}, \quad \tilde{X}_{n+1} \sim Q_n \right\}. \quad (3)$$

The event \mathcal{E}_n exposes subsets of \mathbf{X} where the proposal Q_n does not support π *well enough*. The parameter \bar{W} controls the tolerated discrepancy between the proposal and the target distribution.

3.3 Proposal kernel

The proposal kernel Q_n takes the form of a mixture

$$Q_n = \omega_n Q_0 + (1 - \omega_n) \sum_{\ell=1}^{M_n} \beta_\ell \phi_\ell / \sum_{\ell=1}^{M_n} \beta_\ell, \quad (4)$$

where $\{\omega_n\}_n$ is a sequence of nonincreasing weights such that $\omega_0 = 1$ and $\omega_n > 0$ (see Section 4), M_n is the number of components added to the mixture up to iteration n and $\{\phi_1, \phi_2, \dots, \phi_{M_n}\}$ are the incremental mixture components created up to iteration n . To each incremental component ϕ_ℓ is attached a weight proportional to $\beta_\ell > 0$ (see Section 3.4).

3.4 Increment design

Upon the occurrence of \mathcal{E}_n , a new component $\phi_{M_{n+1}}$ is created. Specifically, $\phi_{M_{n+1}}$ is a multivariate Gaussian distribution with mean $\mu_{M_{n+1}} := \tilde{X}_{n+1}$ and with covariance matrix $\Sigma_{M_{n+1}}$ defined by

$$\Sigma_{M_{n+1}} := \text{Cov} \left\{ \mathfrak{N}(\tilde{X}_{n+1} | X_1, \dots, X_n) \right\}, \quad (5)$$

where for a set of vectors V , $\text{Cov}(V)$ denotes the empirical covariance matrix of V and $\mathfrak{N}(\tilde{X}_{n+1} | V)$ is a subset of V defined as a neighborhood of \tilde{X}_{n+1} . In (5), $\mathfrak{N}(\tilde{X}_{n+1} | X_1, \dots, X_n)$ is the neighborhood of \tilde{X}_{n+1} defined as

$$\mathfrak{N}(\tilde{X}_{n+1} | X_1, \dots, X_n) := \left\{ X_i \in (X_1, \dots, X_n) : D_M(X_i, \tilde{X}_{n+1}) \leq \tau \rho_n \pi(\tilde{X}_{n+1}) \right\}, \quad (6)$$

where $\tau \in (0, 1)$ is a user-defined parameter controlling the neighborhood range and ρ_n the number of accepted proposals up to iteration n . In (6), $D_M(X_i, \tilde{X}_{n+1})$ denotes the Mahalanobis distance between X_i and \tilde{X}_{n+1} which relies on a covariance matrix Σ_0 . When π is a posterior distribution, Σ_0 could be, for example, the prior covariance matrix. Finally, a weight is attached to the new component $\phi_{M_{n+1}}$ proportional to

$$\beta_{M_{n+1}} := \pi(\tilde{X}_{n+1})^\gamma, \quad \gamma \in (0, 1), \quad (7)$$

where $\gamma \in (0, 1)$ is another user-defined parameter.

Note that in Eqs (6) and (7), one can possibly plug the unnormalized probability density function instead of π . Indeed, in Eq. (6) the upper bound is proportional to τ which may be set so as to account for the normalizing constant. In Eq. (7), the weights are eventually self-normalized (see Eq. (4)), hence unknown normalizing constants are not problematic here neither.

3.5 AIMM

Algorithm 1 summarizes AIMM. Note that during an initial phase consisting of N_0 iterations, no adaptation is made. This serves the purpose of gathering sample points required to produce the first increment. Also, we have denoted by N the Markov chain length which may depend on the dimension of \mathbf{X} and the computational resources available. In any case, we recommend setting $N \gg N_0$.

Algorithm 1 Adaptive Incremental Mixture MCMC

- 1: **Input:** user-defined parameters: $\overline{W}, \gamma, \tau, \{\omega_n\}_n, N_0, N$
a proposal distribution: Q_0
- 2: **Initialize:** $X_0 \sim Q_0, W_0 = \pi/Q_0, M_0 = 0$ and $\omega_0 = 1$
- 3: **for** $n = 0, 1, \dots, N$ **do**
- 4: Propose $\tilde{X}_{n+1} \sim Q_n = \omega_n Q_0 + (1 - \omega_n) \sum_{\ell=1}^{M_n} \beta_\ell \phi_\ell / \sum_{\ell=1}^{M_n} \beta_\ell$
- 5: Draw $U_n \sim \text{Unif}(0, 1)$ and set

$$\begin{cases} X_{n+1} = \tilde{X}_{n+1} & \text{if } U_n \leq \left(1 \wedge W_n(\tilde{X}_{n+1})/W_n(X_n)\right) \\ X_{n+1} = X_n & \text{otherwise} \end{cases}$$

- 6: **if** $n > N_0$ and $\mathcal{E}_n := \{W_n(\tilde{X}_{n+1}) > \overline{W}\}$ is true **then**
- 7: Update $M_{n+1} = M_n + 1$
- 8: Increment the proposal kernel with the new component

$$\phi_{M_{n+1}} = \mathcal{N}\left(\tilde{X}_{n+1}, \text{Cov}(\mathfrak{N}(\tilde{X}_{n+1} | X_1, \dots, X_n))\right)$$

- 9: weighted by $\beta_{M_{n+1}} = \pi(\tilde{X}_{n+1})^\gamma$
Update $W_{n+1} = \pi/Q_{n+1}$
 - 10: **else**
 - 11: Update $W_{n+1} = W_n$ and $M_{n+1} = M_n$
 - 12: **end if**
 - 13: **end for**
 - 14: **Output:** $\{X_n, n \leq N\}, \{Q_n, n \leq N\}$
-

3.6 Example 1 (ctd.)

For the toy example in Section 2 we used the following parameters:

$$\overline{W} = 1, \gamma = 0.5, \tau = 0.5, N_0 = 1,000, \omega_n = (1 + M_n/10)^{-1}, Q_0 = \mathcal{N}(0, 10).$$

Figure 1 shows the different states of the chain that, before the first increment takes place, are all sampled from Q_0 . The proposed state \tilde{X}_{n+1} activates the increment process when the

condition $\{W_n(\tilde{X}_{n+1}) > \overline{W}\}$ is satisfied for the first time after the initial phase ($N_0 = 1,000$) is completed. At \tilde{X}_{n+1} there is a large discrepancy between the current proposal, Q_n , and the target, π_1 . The neighborhood of \tilde{X}_{n+1} is identified and defines ϕ_1 , the first component of the incremental mixture.

The top panel of Figure 2 shows the sequence of proposal distributions up to $n = 20,000$ iterations. AIMM has incremented the proposal 15 times. The first proposals are slightly bumpy and as more samples become available, the proposal distribution gets closer to the target. Thus, even without any information about π_1 , AIMM is able to increment its proposal so that the discrepancy between Q_n and π_1 vanishes. This is confirmed by Figure 3 which reports the acceptance rate (averaged on a moving window) and the number of components in the mixture throughout the algorithm. Finally, the proposal kernels obtained after 20,000 iterations of AGM-2 and AGM-3 are reported at the bottom panel of Figure 2 and are in line with their performances as given in Table 1.

The AGM-3 proposal density declines to zero in π_1 's low density regions, which is an appealing feature in this simple example. However, AGM shrinks its proposal density in locations that have not been visited by the chain. This can be problematic if the sample space is large and it takes time to reasonably explore it (see Section 5).

4 Ergodicity of AIMM

4.1 Notation

Let $(\Omega, \mathcal{A}, \mathbb{P})$ be the probability space generated by AIMM. For any event $A \in \mathcal{A}$, \mathbb{P}_A denotes the probability measure \mathbb{P} conditionally on A . With some abuse of notations, for any $x \in \mathbf{X}$, \mathbb{P}_x refers to \mathbb{P} conditionally on the event $\{X_0 = x\}$.

Let ν and μ be two probability measures defined on $(\mathbf{X}, \mathcal{X})$. Recall that the Kullback–Leibler (KL) divergence between ν and μ is defined as

$$\text{KL}(\nu, \mu) = \int_{\mathbf{X}} \log \frac{\nu(x)}{\mu(x)} d\nu(x).$$

Moreover, the total variation distance between ν and μ can be written as

$$\|\nu - \mu\| = \sup_{A \in \mathcal{X}} |\nu(A) - \mu(A)| = \frac{1}{2} \int_{\mathbf{X}} |\mu(x) - \nu(x)| d\rho(x),$$

where the latter equality holds if μ and ν are both dominated by a common measure ρ . Finally, for a sequence of random variables $\{X_n\}_{n>0}$, the notations $X_n = o_p(1)$ stands for




Figure 1: Example 1: Illustration of one AIMM increment for the target π_1 . Top: states X_1, \dots, X_n of the AIMM chain, proposed new state \tilde{X}_{n+1} activating the increment process (*i.e.* satisfying $W_n(\tilde{X}_{n+1}) \geq \bar{W}$) and neighborhood of \tilde{X}_{n+1} , $\mathfrak{N}(\tilde{X}_{n+1} | X_1, \dots, X_n)$. Bottom: target π_1 , defensive kernel Q_0 , first increment ϕ_1 and updated kernel Q_1 plotted on a logarithmic scale.

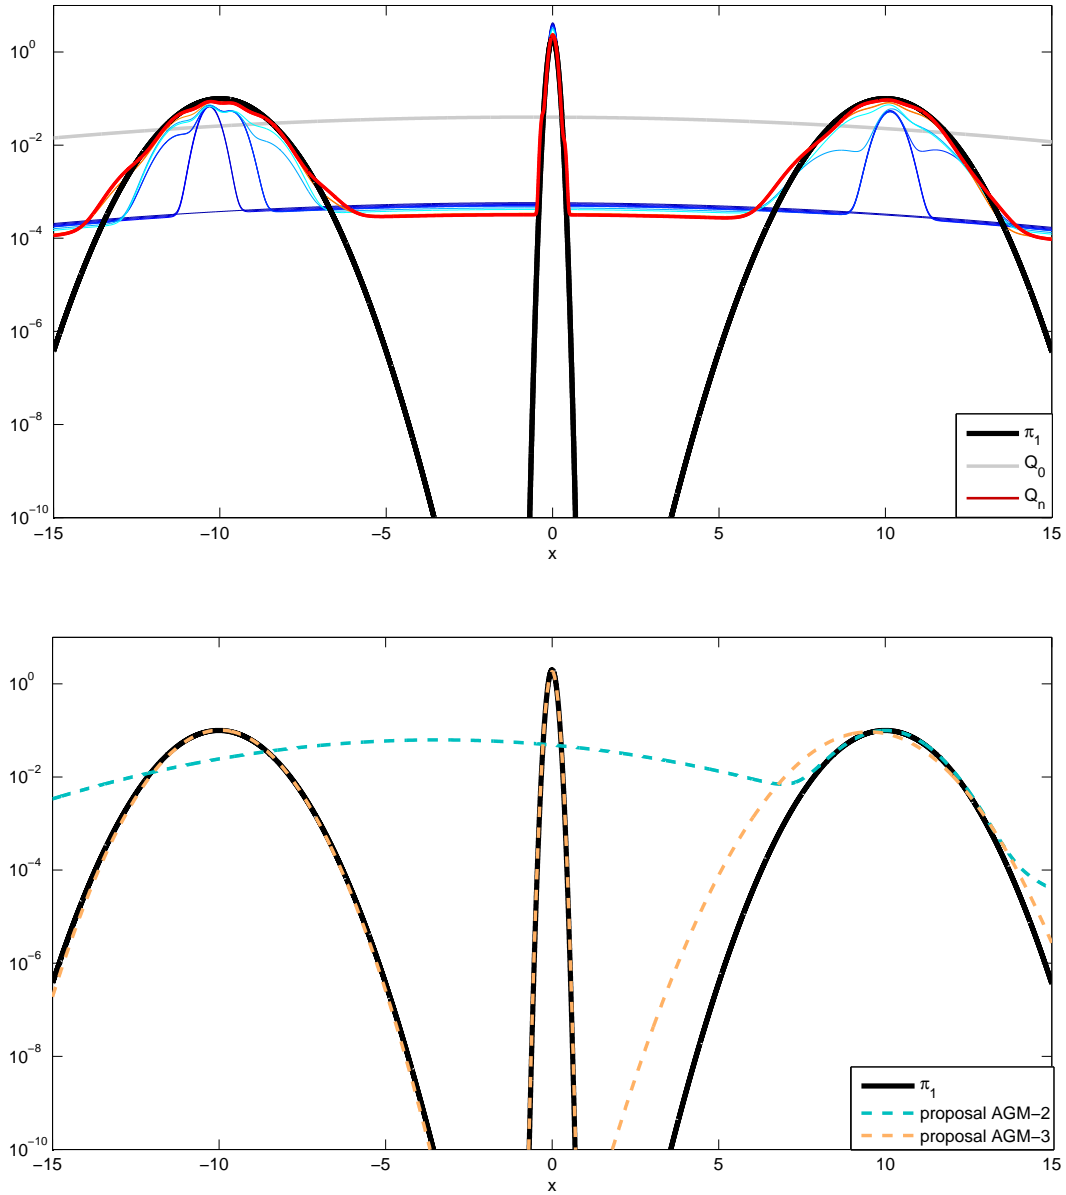


Figure 2: Example 1: Illustration of AIMM and AGM sampling from the target π_1 for $n = 20,000$ iterations. Top: sequence of proposals from Q_0 to Q_n produced by AIMM. Bottom: proposals produced by AGM-2 and AGM-3. Both plots are in log-lin scale.




Figure 3: Example 1: Effect of AIMM increments on the sampling efficiency for the target π_1 - acceptance rate and number of increments throughout $n = 20,000$ iterations of AIMM.

the convergence of $\{X_n\}_{n>0}$ to zero in probability p (i.e. with respect to the probability measure p), and $X_n = O_p(1)$ means that $\{X_n\}_{n>0}$ is bounded in probability p .

4.2 Ergodicity of adaptive Markov chains

In this section, we establish assumptions under which the AIMM process is ergodic, *i.e.*

$$\forall x \in \mathsf{X}, \quad \lim_{n \rightarrow \infty} \|\mathbb{P}_x(X_n \in \cdot) - \pi\| = 0. \quad (8)$$

We use the theoretical framework developed in Roberts and Rosenthal (2007). In particular, Theorem 2 in Roberts and Rosenthal (2007) states that if the AIMM process satisfies the *Diminishing adaptation* and *Containment* conditions (see below) then the process is ergodic and Eq. (8) holds.

Condition 1. *Diminishing adaptation.*

For all $x \in \mathsf{X}$, the stochastic process $\{\Delta_n, n \in \mathbb{N}\}$, defined as

$$\Delta_n := \sup_{x \in \mathsf{X}} \|K_{n+1}(x, \cdot) - K_n(x, \cdot)\| \quad (9)$$

converges to zero in \mathbb{P}_x -probability, *i.e.* $\Delta_n = o_{\mathbb{P}_x}(1)$.

Condition 2. *Containment.*

For all $\epsilon > 0$ and $x \in \mathsf{X}$, the stochastic process $\{M_n(\epsilon), n \in \mathbb{N}\}$, defined as

$$C_n(\epsilon) := \inf\{N \in \mathbb{N}, \|K_n^N(X_n, \cdot) - \pi\| < \epsilon\} \quad (10)$$

is bounded in \mathbb{P}_x -probability, *i.e.* $C_n(\epsilon) = O_{\mathbb{P}_x}(1)$.

Even though containment is not a necessary condition for ergodicity (see Fort et al. (2011)) and in fact seems to hold in most practical situations (see *e.g.* Rosenthal et al. (2011)), it remains a challenging assumption to establish rigorously for the setup considered in this paper, *i.e.* a broad class of target distributions defined on a not necessarily finite or compact state space. In the next section we show that diminishing adaptation holds for the AIMM process described at Section 3, up to minor implementation adjustments, while containment needs to be proven on a case by case basis. We also show that there exists a version of AIMM, conceptually different from the process described at Section 3 that can be made arbitrarily close to it, which is ergodic for any absolutely continuous target distribution π .

4.3 Main ergodicity results

We study two variants of AIMM.

4.3.1 Proposal with an unlimited number of increments

The first version of AIMM is similar to the algorithm presented in Section 3. For this version of the algorithm, we prove only that the diminishing adaptation assumption holds under some minor assumptions. Indeed, proving the containment condition is challenging without further assumptions on π , *e.g.* compact support (Craiu et al., 2009) or tail properties (see Bai et al. (2009) in the context of adaptive random walk Metropolis samplers). Moreover, most proof techniques establishing containment require the space of the adapting parameter to be compact, something which does not hold in this version of AIMM as the number of incremental components can, theoretically, grow to infinity.

Proposition 1. *Assume that there exist three positive constants (δ, η, λ) such that*

A1. *The covariance matrix of any component of the mixture satisfies $\det \Sigma_k > \delta > 0$.*

A2. *The (unnormalized) incremental mixture weights are defined as:*

$$\beta_{M_n+1} = \frac{\eta + \pi(\tilde{X}_{n+1})^\gamma}{(1 + \eta)^{M_n+1}} \quad \text{and} \quad \omega_n = \frac{1}{1 + \sum_{k=1}^{M_n} \beta_k} \vee \lambda. \quad (11)$$

A3. *The initial kernel Q_0 is subexponential or satisfies $Q_0(x) \propto \exp\{\psi(x)\}$ where $\psi = o(x^2)$ ($x \rightarrow \infty$), *i.e.* Q_0 has heavier tail than a Gaussian, *e.g.* multivariate Laplace or t -distribution.*

A4. There is a parameter $0 < \underline{W} < \overline{W}$ such that the mixture increments upon the event $\mathcal{E}_n \cup \mathcal{F}_n$,

$$\mathcal{F}_n := \{W_n(\tilde{X}_{n+1}) < \underline{W}\}$$

and when it increments upon \mathcal{F}_n , the new component is equal to Q_0 and the corresponding weight is defined as in Eq. (7).

Then, the AIMM process sampled using conditions **A1–A4** satisfies diminishing adaptation, i.e. Eq. (9) holds.

The proof is given in Appendix A.

Remark 1. Proposition 1 holds for parameters (δ, η, λ) arbitrarily close to zero. Also, Assumption **A1** can be enforced simply by expanding the neighborhood of \tilde{X}_{n+1} such that when $\det \text{Cov}\{\mathfrak{N}(\tilde{X}_{n+1})\} < \delta$,

$$\Sigma_{M_{n+1}} := \text{Cov}\{X_1^*, \dots, X_k^*\}$$

where $\{X_1^*, \dots, X_n^*\}$ is a permutation of $\{X_1, \dots, X_n\}$ such that $D_M(\tilde{X}_{n+1}, X_i^*) \leq D_M(\tilde{X}_{n+1}, X_{i+1}^*)$ and $k := \inf\{i \leq n, \text{Cov}(X_1^*, \dots, X_i^*) \geq \delta\}$.

Remark 2. The event \mathcal{F}_n is the counterpart of \mathcal{E}_n and exposes subsets of X where the proposal puts too much probability mass at locations of low π -probability. In this case, the rationale for setting $\phi_{M_{n+1}} = Q_0$ is to reduce the probability mass of the proposal locally by increasing the weight associated to the initial proposal Q_0 , assumed to be vague.

4.3.2 Proposal with adaptation on a compact set

The second version of AIMM that we study here represents somewhat a slight conceptual departure from the original algorithm presented in Section 3. We stress that the assumptions **A1–A4** from Section 4.3.1 are not required here.

Proposition 2. Assume that:

B1. The AIMM process $\{X_n, n \in \mathbb{N}\}$ has bounded jumps, i.e. there exists $D > 0$ such that for all $n \in \mathbb{N}$

$$\mathbb{P}\{\|X_n - X_{n+1}\| \leq D\} = 1. \quad (12)$$

B2. There is a compact set $\mathcal{K} \subset \mathsf{X}$ such that the weight $\{\omega_n\}_n$ in the proposal (4) is replaced by $\{\bar{\omega}_n^{\mathcal{K}}\}_n$

$$\bar{\omega}_n^{\mathcal{K}} = \omega_n \vee \mathbb{1}_{X_n \notin \mathcal{K}}, \quad (13)$$

i.e. if the process is such that $X_n \notin \mathcal{K}$ then the proposed state is generated from Q_0 with probability one. Conversely, if the process is such that $X_n \in \mathcal{K}$, then the proposed state is generated as explained in Section 3, i.e. from Q_n (4). We denote this proposal by $\bar{Q}_n^{\mathcal{K}}$.

B3. *The number of incremental components in the adaptive proposal is capped, i.e. there is a finite $M \in \mathbb{N}$ such that $\mathbb{P}(M_n \leq M) = 1$ and the mean μ_n and covariance Σ_n of each component are defined on a compact space.*

*If the AIMM process presented in Section 3 satisfies **B1–B3**, then it is ergodic.*

This result is a consequence of Theorem 2 of Roberts and Rosenthal (2007) combined with the recent developments in Craiu et al. (2015) and Rosenthal and Yang (2017). The proof is given in Appendix B.

We now explain how, in practice, a version of AIMM compatible with the assumptions of Proposition 2 can be constructed and made arbitrarily close to the version of AIMM presented in Section 3. First, fix an arbitrarily large constant $D < \infty$. Assumption **B1** holds by construction if $\tilde{X}_{n+1} \sim \bar{Q}_n^{\mathcal{K}}$ is automatically rejected when $\|\tilde{X}_{n+1} - X_n\| > D$. Assumption **B2** holds by construction if $\bar{Q}_n^{\mathcal{K}}$ is used instead of Q_n in the AIMM algorithm presented in Section 3. Finally, Assumption **B3** is satisfied by slightly modifying the adaptation mechanism proposed in Section 3. Define two arbitrarily large constants $L > 0$ and $M \in \mathbb{N}$ and increment the proposal upon the event $\bar{\mathcal{E}}_n := \mathcal{E}_n \cap \{M_n < M\}$ and in the following way:

$$\bar{\mu}_{M_{n+1}} := \tilde{X}_{n+1}|_L, \quad \bar{\Sigma}_{M_{n+1}} := \text{Cov}\{\bar{\mathfrak{N}}_L(\tilde{X}_{n+1} | X_1, \dots, X_n)\},$$

where for any vector $x \in \mathbb{R}^d$ and any $L > 0$

$$x|_L = \{\{x_1 \wedge L\} \vee -L, \dots, \{x_d \wedge L\} \vee -L\} \quad \text{and} \quad \bar{\mathfrak{N}}_L = \left\{ X_i|_L, X_i \in \mathfrak{N}(\tilde{X}_{n+1} | X_1, \dots, X_n) \right\}.$$

The definition of the unnormalized weight $\beta_{M_{n+1}}$ (Eq. (7)) is unchanged.

5 Simulations

In this Section, we consider three target distributions:

- π_2 , the banana shape distribution used in Haario et al. (2001);
- π_3 , the ridge like distribution used in Raftery and Bao (2010);
- π_4 , the bimodal distribution used in Raftery and Bao (2010).

Each of these distributions has a specific feature, resulting in a diverse set of challenging targets: π_2 is heavy tailed, π_3 has a narrow, nonlinear and ridge-like support and π_4 has two distant modes. We consider π_2 and π_4 in different dimensions: $d \in \{2, 10, 40\}$ for π_2 and $d \in \{4, 10\}$ for π_4 . We compare AIMM (Algorithm 1) with several other algorithms that are

briefly described in Appendix C using some performance indicators that are also defined, explained and justified in Appendix C. We used the following default parameters for AIMM:

$$\bar{W} = d, \quad \gamma = 0.5, \quad \tau = 0.5, \quad \kappa = 0.1, \quad N_0 = 1,000\sqrt{d}.$$

For each scenario, we implemented AIMM with different thresholds \bar{W} valued around d to monitor the tradeoff between adaptation and computational efficiency of the algorithm. In our experience, the choice $\bar{W} = d$ has worked reasonably well in a wide range of examples. As the dimension of the state space increases a higher threshold is required, too many kernels being created otherwise. However, a satisfactory choice of \bar{W} may vary depending on the target distribution and the available computational budget.

5.1 Banana-shaped target distribution

Example 2. Let $\mathbf{X} = \mathbb{R}^d$, $\pi_2(x) = \Psi_d(f_b(x); m, S)$ where $\Psi_d(\cdot; m, S)$ is the d -dimensional Gaussian density with mean m and covariance matrix S , and let $f_b : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be the mapping defined, for any $b \in \mathbb{R}$, by

$$f_b : \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_d \end{pmatrix} \rightarrow \begin{pmatrix} x_1 \\ x_2 + b x_1^2 - 100b \\ x_3 \\ \vdots \\ x_d \end{pmatrix}.$$

We consider π_2 in dimensions $d = 2$, $d = 10$ and $d = 40$ and refer to the marginal of π_2 in the i -th dimension as $\pi_2^{(i)}$. The parameters $m = \mathbf{0}_d$ and $S = \text{diag}([100, \mathbf{1}_{d-1}])$ are held constant. The target is banana-shaped along the first two dimensions, and is challenging since the high density region has a thin and wide ridge-like shape with narrow but heavy tails. Unless otherwise stated, we use $b = 0.1$ which accentuates these challenging features of π_2 . We first use the banana-shaped target π_2 in dimension $d = 2$ to study the influence of AIMM’s user-defined parameters on the sampling efficiency.

5.1.1 Influence of the defensive distribution

With limited prior knowledge of π , choosing Q_0 can be challenging. Here we consider three defensive distributions Q_0 , represented by the black dashed contours in Figure 4. The first two are Gaussian, respectively, located close to the mode and in a nearly zero probability area and the last one is the uniform distribution on the set $\mathfrak{S} = \{x \in \mathbf{X}, x_1 \in (-50, 50), x_2 \in (-100, 20)\}$. Table 2 and Figure 4 show that AIMM is robust with respect to the choice of Q_0 . Even when Q_0 yields a significant mismatch with π_2 (second case), the incremental

mixture reaches high density areas, progressively uncovering higher density regions. The price to pay for a poorly chosen Q_0 is that more components are needed in the incremental mixture to match π_2 , as shown by Table 2. The other statistics reported in Table 2 are reasonably similar for the three choices of Q_0 .

Table 2: Example 2: π_2 target, $d = 2$ - Influence of the initial proposal Q_0 on AIMM outcome after $n = 100,000$ MCMC iterations (replicated 20 times). M_n is the number of components created by AIMM, ESS is the effective sample size, ACC is the acceptance rate, KL is the KL divergence between π_2 and the chain distribution, JMP is the average distance between two consecutive states of the chain and EFF is a time-normalized ESS.

Q_0	M_n	ESS	ACC	KL	JMP	EFF ($\times 10^{-4}$)
Gaussian on a high density region	41	.19	.35	.59	197	11
Gaussian on a low density region	157	.23	.37	.71	245	4
Uniform on \mathfrak{S}	39	.24	.38	.62	320	11

5.1.2 Influence of the threshold

Even though we recommend the default setting $\bar{W} = d$, we analyse the performances of AIMM for different values of \bar{W} . We consider the same setup as in the previous subsection with the uniform defensive distribution Q_0 and report the outcome of the AIMM algorithm in Table 3a. As expected, the lower the threshold \bar{W} , the larger the number of kernels and the better is the mixing. In all cases, the adaptive transition kernel eventually stabilizes (faster if \bar{W} is large), as illustrated by Figure 5a, resulting from the fact that the event $\mathcal{E}_n = \{W_n(\tilde{X}_{n+1}) > \bar{W}, \tilde{X}_{n+1} \sim Q_n\}$ occurs less often as n increases. As \bar{W} decreases, the KL divergence between π_2 and the chain reduces while the CPU cost increases since more components are created. Therefore, the sampling efficiency is best for an intermediate threshold such as $\log \bar{W} = .75$. Finally, when the threshold is too low, the distribution of the chain converges more slowly to π_2 ; see e.g. Table 3a where the KL indicator (defined as the K-L divergence between π_2 and the sample path of the AIMM chain; see Appendix C) is larger for $\log \bar{W} = 0.25$ than for $\log \bar{W} = 0.5$. Indeed when $\bar{W} \ll 1$, too many kernels are created to support high density areas and this slows down the exploration process of the lower density regions.

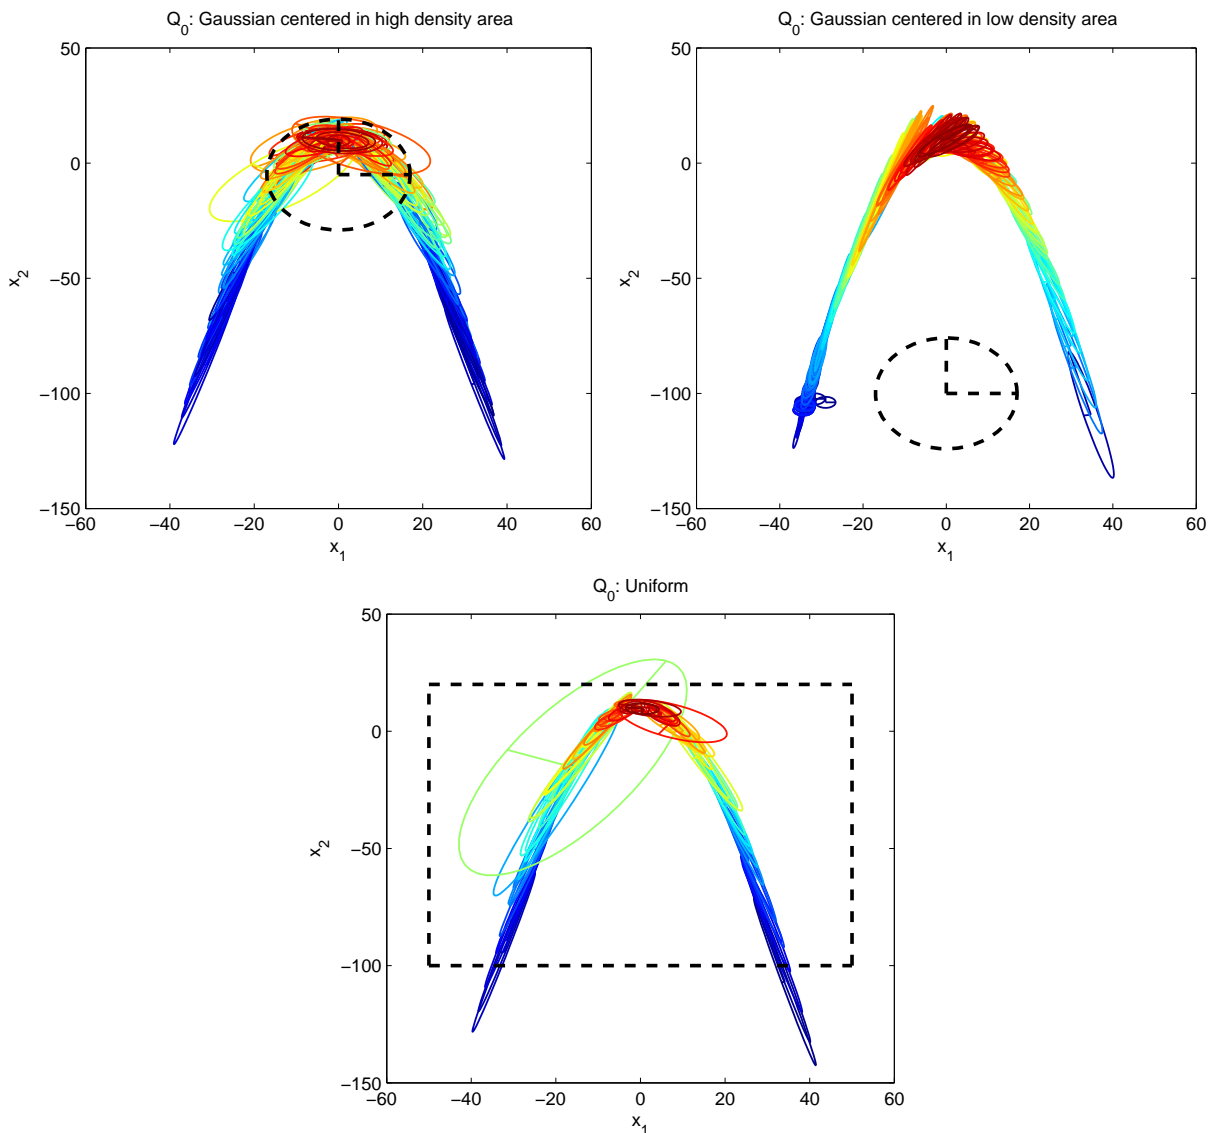




Figure 4: Example 2: π_2 target, $d = 2$ - Incremental mixture created by AIMM for three different initial proposals. Top row: Q_0 is a Gaussian density (the region inside the black dashed ellipses contains 75% of the Gaussian mass) centred on a high density region (left) and a low density region (right). Bottom row: Q_0 is Uniform (with support corresponding to the black dashed rectangle). The components $\phi_1, \dots, \phi_{M_n}$ of the incremental mixture obtained after $n = 100,000$ MCMC iterations are represented through (the region inside the ellipses contains 75% of each Gaussian mass). The color of each ellipse illustrates the corresponding component's relative weight β_ℓ (from dark blue for lower weights to red).



(a) Evolution of the number of kernels M_n created by AIMM for different thresholds.



(b) Evolution of the KL divergence between π_2 and the incremental proposal Q_n , plotted in lin-log scale.

Figure 5: Example 2: π_2 target, $d = 2$ - AIMM's incremental mixture design after $n = 100,000$ MCMC iterations.

Table 3: Example 2: π_2 target, $d = 2$ - Influence of the threshold \overline{W} on AIMM and f-AIMM outcomes after $n = 100,000$ MCMC iterations (replicated 20 times).

(a) AIMM

$\log \overline{W}$	M_n	ESS	ACC	KL	JMP	CPU	EFF ($\times 10^{-4}$)
10	0	.03	.01	14.06	10	45	2
2	21	.16	.27	.72	169	154	10
1.5	39	.24	.38	.62	235	245	10
1	79	.37	.53	.54	331	330	12
.75	149	.48	.64	.53	286	658	7
.5	317	.64	.75	.51	451	2,199	3
.25	1162	.71	.87	.54	505	6,201	1

(b) f-AIMM

$\log \overline{W}$	M_{\max}	ESS	ACC	KL	JMP	CPU	EFF ($\times 10^{-4}$)
1.5	25	.29	.51	.59	268	255	11
.75	100	.53	.69	.53	409	421	13
.5	200	.67	.80	.51	474	1,151	6

5.1.3 Speeding up AIMM

The computational efficiency of AIMM depends, to a large extent, on the number of components added in the proposal distribution, M_n . For this reason we consider a slight modification of the original AIMM algorithm to limit the number of possible components, thereby improving the computational speed of the algorithm. This variant of the algorithm is outlined below and will be referred to as *fast* AIMM and denoted by f-AIMM:

- Let M_{\max} be the maximal number of components allowed in the incremental mixture proposal. If $M_n > M_{\max}$, only the last M_{\max} added components are retained, in a moving window fashion. This *truncation* has two main advantages, (i) approximately linearizing the computational burden once M_{\max} is reached, and (ii) forgetting the first, often transient components used to jump to high density areas (see *e.g.* the loose ellipses in Figure 4, especially the very visible ones in the bottom panel).
- The threshold \overline{W} is adapted at the beginning of the algorithm in order to get the incremental process started. If the initial proposal Q_0 misses π , then $\mathbb{E}_0\{W_0(\tilde{X})\} \ll 1$ and the initial threshold \overline{W}_0 should be set as low as possible, to start the adaptation. However, as soon as the proposal is incremented, $\mathbb{E}_n\{W_n(\tilde{X})\}$ increases and leaving the threshold at its initial level will result in $\mathbb{P}_n\{W_n(\tilde{X}) > \overline{W}\} \approx 1$, i.e. in adding too many irrelevant components. The threshold adaptation produces a sequence of thresholds $\{\overline{W}_n\}_n$ such that $Q_n\{W_n(X) > \overline{W}_n\} \approx 10^{-3}$. Since sampling from Q_n (a mixture of

Gaussian distributions), can be performed routinely, \overline{W}_n is derived from a Monte Carlo estimation. As little precision is required, the Monte Carlo approximation should be rough in order to limit the computational burden generated by the threshold adaptation. Also, those samples used to set \overline{W}_n can be recycled and serve as proposal states for the AIMM chain so that this automated threshold mechanism comes for free. The threshold adaptation stops when the adapted threshold \overline{W}_n reaches a neighborhood of the prescribed one \overline{W} , here defined as $|\overline{W}_n - \overline{W}| < 1$.

Table 3 shows that f-AIMM outperforms the original AIMM, with some setups being nearly twice as efficient; compare, e.g. AIMM and f-AIMM with $\log \overline{W} = .5$ in terms of efficiency (last column). Beyond efficiency, comparing KL for a given number of mixture components (M_n), shows that π_2 is more quickly explored by f-AIMM than by AIMM.

5.1.4 Comparison with other samplers

We now compare f-AIMM with four other samplers of interest: the adaptive algorithms AMH and AGM and their non-adaptive counterparts, random walk Metropolis-Hastings (RWMH) and Independent Metropolis (IM). We consider π_2 in dimensions $d = 2, 10, 40$. Table 4 summarizes the simulation results and emphasizes the need to compare the asymptotic variance related statistics (ESS, ACC, JMP) jointly with the convergence related statistic (KL). Indeed, from the target in dimension $d = 2$, one can see that AGM yields the best ESS but is very slow to converge to π_2 as it still misses a significant portion of the state space after $n = 200,000$ iterations; see the KL column in Table 4 and Appendix C (especially Figure 1 therein) where we shed light on AGM’s performance in this situation.

AIMM seems to yield the best tradeoff between convergence and variance in all dimensions. Inevitably, as d increases, ESS shrinks and KL increases for all algorithms but AIMM still maintains a competitive advantage over the other approaches. It is worth noting that, if computation is not an issue, AIMM is able to reach a very large ESS (almost one), while having converged to the target. In such a situation, AIMM is essentially simulating *i.i.d.* draws from π_2 .

Figure 6 compares the kernel density approximation of the (Markov) chain marginal distribution of the first component for the five algorithms, with the true marginal $\pi_2^{(1)} = \mathcal{N}(0, 10)$. AIMM converges to π_2 faster than the other samplers, which need much more than $n = 200,000$ iterations to converge. The bumpy shape of the AGM and IM marginal densities is caused by the accumulation of samples in some locations and reflects the slow convergence of these algorithms.

The other four algorithms struggle to visit the tail of the distribution, see Table 5. Thus AIMM is doing better than (i) the non-independent samplers (RWMH and AMH) that

manage to explore the tail of a target distribution but need a large number of transitions to return there, and (ii) the independence samplers (IM and AGM) that can quickly explore different areas but fail to venture into the lower density regions. An animation corresponding to the exploration of π_2 in dimension 2 by AIMM, AGM and AMH can be found at this link.

5.2 Ridge like example

Example 3. Let $X = \mathbb{R}^6$ and $\pi_3(x) \propto \Psi_6(x; \mu_i, \Gamma_i) \Psi_4(g(x); \mu_o, \Gamma_o)$, where $\Psi_d(\cdot; m, S)$ is the d -dimensional Gaussian density with mean m and covariance matrix S . The target parameters $(\mu_i, \mu_o) \in \mathbb{R}^6 \times \mathbb{R}^4$ and $(\Gamma_i, \Gamma_o) \in \mathcal{M}_6(\mathbb{R}) \times \mathcal{M}_4(\mathbb{R})$ are, respectively, known means and covariance matrices and g is a nonlinear deterministic mapping $\mathbb{R}^6 \rightarrow \mathbb{R}^4$ defined as:

$$g(x_1, \dots, x_6) = \begin{cases} \prod_{i=1}^6 x_i, \\ x_2 x_4, \\ x_1 / x_5, \\ x_3 x_6. \end{cases}$$

In this context, $\Psi_6(\cdot; \mu_i, \Gamma_i)$ can be regarded as a prior distribution and $\Psi_4(\cdot; \mu_o, \Gamma_o)$ as a likelihood, the observations being some functional of the hidden parameter $x \in \mathbb{R}^6$.

Such a target distribution often arises in physics. Similar target distributions often arise in Bayesian inference for deterministic mechanistic models and in other areas of science, engineering and environmental sciences; see *e.g.* Poole and Raftery (2000); Bates et al. (2003). They are hard to sample from because the probability mass is concentrated around thin curved manifolds. We compare f-AIMM with RWMH, AMH, AGM and IM first in terms of their mixing properties; see Table 6. We also ensure that the different methods agree on the mass localisation by plotting a pairwise marginal; see Figure 7. In this example, f-AIMM clearly outperforms the four other methods in terms of both convergence and variance statistics. We observe in Figure 7, that f-AIMM is the only sampler able to discover a secondary mode in the marginal distribution of the second and fourth target components (X_2, X_4) .

5.3 Bimodal distribution

Example 4. In this last example, π_4 is a posterior distribution defined on the state space $X = \mathbb{R}^d$, where $d \in \{4, 10\}$. The likelihood is a mixture of two d -dimensional Gaussian

Table 4: Example 2: π_2 target, $d \in \{2, 10, 40\}$ - Comparison of f-AIMM with RWMH, AMH, AGM and IM. Statistics were estimated from 20 replications of the five corresponding (Markov) chains after $n = 200,000$ iterations. AGM fails to give sensible results for $d = 40$. †: the ESS column are $\times 10^{-3}$; ‡: the EFF column are $\times 10^{-6}$.

(a) π_2 in dimension $d = 2$

	$\log \bar{W}$	M_{\max}	ESS †	ACC	KL	JMP	CPU	EFF ‡
f-AIMM	1.5	25	290	.51	.59	268	355	810
f-AIMM	.5	200	670	.80	.51	474	1,751	380
RWMH	–	–	1	.30	3.75	2	48	17
AMH	–	–	4	.23	2.69	7	60	63
AGM	–	100	800	.78	29.2	155	2,670	230
IM	–	–	6	.01	10.1	8	92	63

(b) π_2 in dimension $d = 10$

	$\log \bar{W}$	M_{\max}	ESS †	ACC	KL	JMP	CPU	EFF ‡
f-AIMM	3	50	110	.27	1.46	135	1,160	94
f-AIMM	2.5	150	170	.37	.84	205	1,638	100
RWMH	–	–	1	.22	9.63	1.1	61	15
AMH	–	–	1.4	.17	5.29	2.5	143	10
AGM	–	100	600	.67	119	153	4,909	120
IM	–	–	.7	10^{-3}	97	.9	316	2

(c) π_2 in dimension $d = 40$

	$\log \bar{W}$	M_{\max}	ESS †	ACC	KL	JMP	CPU	EFF ‡
f-AIMM	5	200	4	.04	12.1	24	4,250	.7
f-AIMM	4	1,000	17	.08	5.9	61	10,012	1.8
RWMH	–	–	.8	.24	11.6	.8	71	11
AMH	–	–	.8	.13	6.7	1.5	1,534	.7
IM	–	–	.5	10^{-4}	147.1	.03	1,270	.6

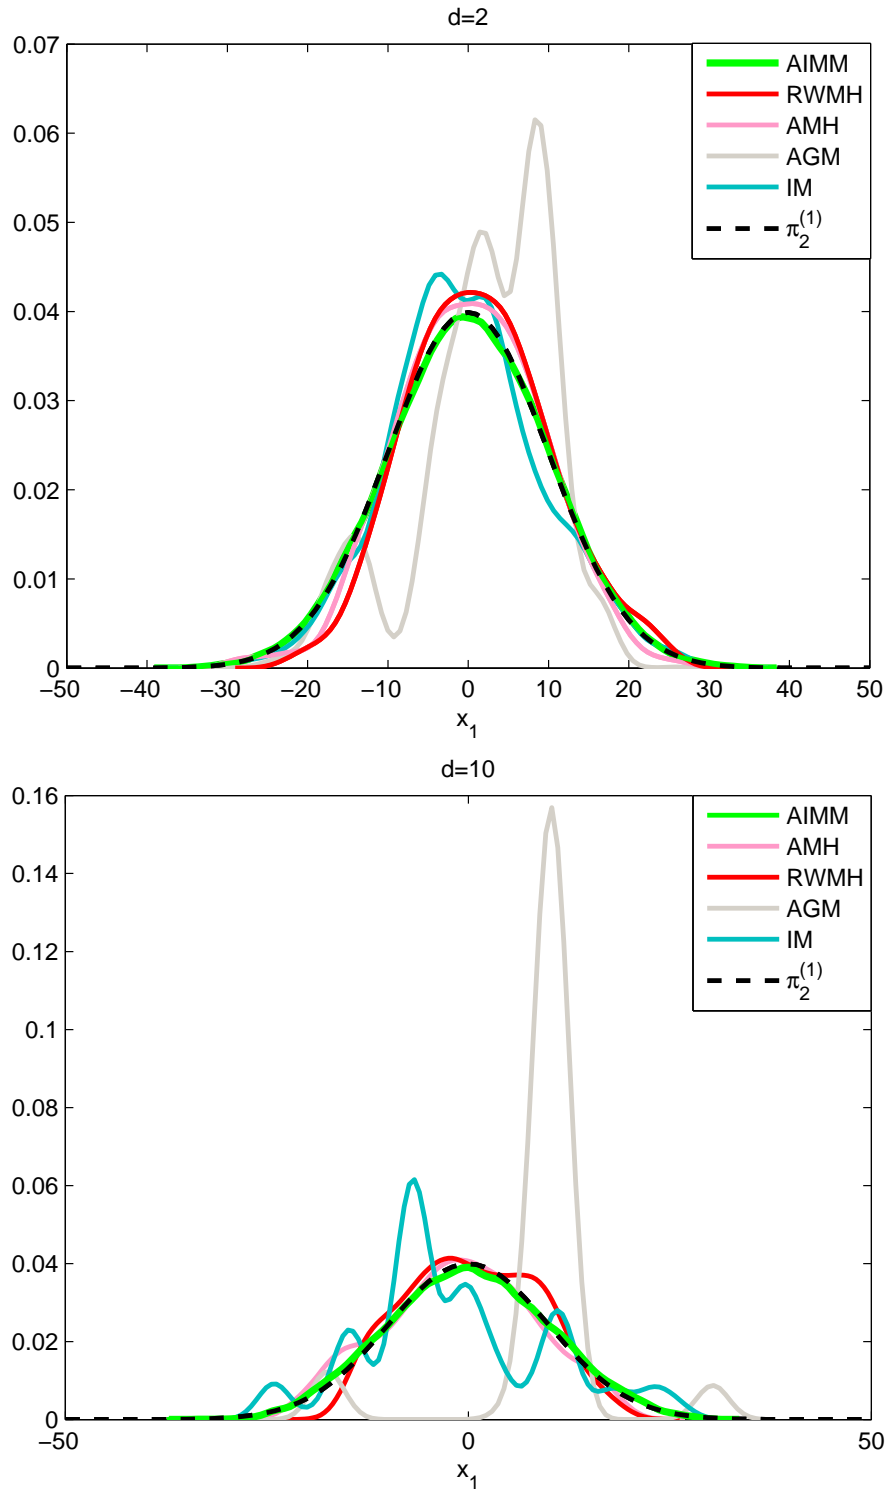


Figure 6: Example 2: π_2 target, $d \in \{2, 10\}$ - Kernel approximation of the marginal density $\pi_2^{(1)}$ provided by $n = 200,000$ samples of the five (Markov) chains AIMM, AMH, RWMH, AGM and IM for the banana shape target example in dimensions $d = 2$ (top) and $d = 10$ (bottom).

Table 5: Example 2: π_2 target, $d \in \{2, 10, 40\}$ - Tail exploration achieved by the five different samplers. Tail events are defined as $\{\text{TAI}_1\} = \{X^{(2)} < -28.6\}$ and $\{\text{TAI}_2\} = \{X^{(2)} < -68.5\}$ so that $\pi_2\{\text{TAI}_1\} = 0.05$ and $\pi_2\{\text{TAI}_2\} = 0.005$. $\{\text{RET}_1\}$ and $\{\text{RET}_2\}$ corresponds to the expected returning time to the tail events $\{\text{TAI}_1\}$ and $\{\text{TAI}_2\}$, respectively. Estimation of these statistics was achieved with $n = 200,000$ iterations of the corresponding samplers, repeated 20 times.

		M_{\max}	TAI_1	RET_1	TAI_2	RET_2
$d = 2$	f-AIMM	25	.05	43	.003	557
	f-AIMM	200	.05	23	.005	281
	RWMH	–	.04	2,945	.002	69,730
	AMH	–	.04	3,676	.004	74,292
	AGM	100	.01	1,023	$5.0 \cdot 10^{-5}$	155,230
	IM	–	.05	1,148	.004	8,885
	<hr/>					
$d = 10$	f-AIMM	50	.04	98	.001	5,463
	f-AIMM	150	.05	57	.002	1,708
	RWMH	–	.03	30,004	0	∞
	AMH	–	.04	4,081	.001	166,103
	AGM	100	.04	120,022	.001	180,119
	IM	–	.05	13,140	.003	103,454
	<hr/>					
$d = 40$	f-AIMM	200	.01	1,350	0	∞
	f-AIMM	1000	.05	175	$2.0 \cdot 10^{-4}$	16,230
	RWMH	–	.002	158,000	0	∞
	AMH	–	0	∞	0	∞
	IM	–	.006	28,552	0	∞

Table 6: Example 3: π_3 target, $d = 6$ - Comparison of f-AIMM with the four other samplers after $n = 200,000$ iterations (replicated 10 times).

	\bar{W}	M_{\max}	ACC	ESS	CPU	EFF	JMP
f-AIMM	10	20	.049	.015	274	$5.2 \cdot 10^{-5}$.08
f-AIMM	1	70	.169	.089	517	$1.7 \cdot 10^{-4}$.27
f-AIMM	.1	110	.251	.156	818	$1.9 \cdot 10^{-4}$.38
<hr/>							
RWMH	–	–	.23	.001	109	$9.2 \cdot 10^{-6}$.007
AMH	–	–	.42	.002	2,105	$9.5 \cdot 10^{-7}$.004
AGM-MH	–	100	.009	.002	2,660	$3.4 \cdot 10^{-6}$.003
IM	–	–	.003	.001	199	$5.5 \cdot 10^{-6}$.003




Figure 7: Example 3: π_3 target, $d = 6$ - Samples $\{(X_{2,n}, X_{4,n}), n \leq 200,000\}$ from the five different algorithms (from left to right, top to bottom): AIMM in two settings $M_{\max} = 20$ and $M_{\max} = 70$, AGM, IM, AMH and RWMH.

distributions with weight $\lambda = 0.5$, mean and covariance matrix as follows

$$\begin{aligned}\mu_1 &= \mathbf{0}_d, & \Sigma_1 &= AR_d(-.95), \\ \mu_2 &= \mathbf{1}_d, & \Sigma_2 &= AR_d(.95),\end{aligned}$$

where for all $\rho > 0$, $AR_d(\rho)$ is the d -dimensional first order autoregressive matrix whose coefficients are:

$$\text{for } 1 \leq (i, j) \leq d, \quad m_{i,j}(\rho) = \rho^{\max(i,j)-1}.$$

The prior is the uniform distribution $\mathcal{U}([-3, 12]^d)$. For f-AIMM and IM, Q_0 is set as this prior distribution, while for AGM, the centers of the Gaussian components in the initial proposal kernel are drawn according to the same uniform distribution. For AMH, the initial covariance matrix Σ_0 is set to be the identity matrix.

Figures 8 and 9 illustrate the experiment in dimensions $d = 4$ and $d = 10$. On the one hand, in both setups $M_{\max} = 100$ and $M_{\max} = 200$, f-AIMM is more efficient at discovering and jumping from one mode to the other. Allowing more kernels in the incremental mixture proposal will result in faster convergence; compare the settings with $M_{\max} = 100$ and $M_{\max} = 200$ in Figure 10. On the other hand, because of the distance between the two modes, RWMH visits only one while AMH visits both but in an unbalanced fashion. Figure 9 displays the samples from the joint marginal (X_1, X_2) obtained through f-AIMM (with two different values of M_{\max}), AMH and RWMH. Increasing the dimension from $d = 4$ to $d = 10$ makes AMH unable to visit the two modes after $n = 200,000$ iterations. As for AGM, the isolated samples reflect a failure to explore the state space. These facts are confirmed in Table 7 which highlights the better mixing efficiency of f-AIMM relative to the four other algorithms and in Table 8 which shows that f-AIMM is the only method which, after $n = 200,000$ iterations, visits the two modes with the correct proportion.

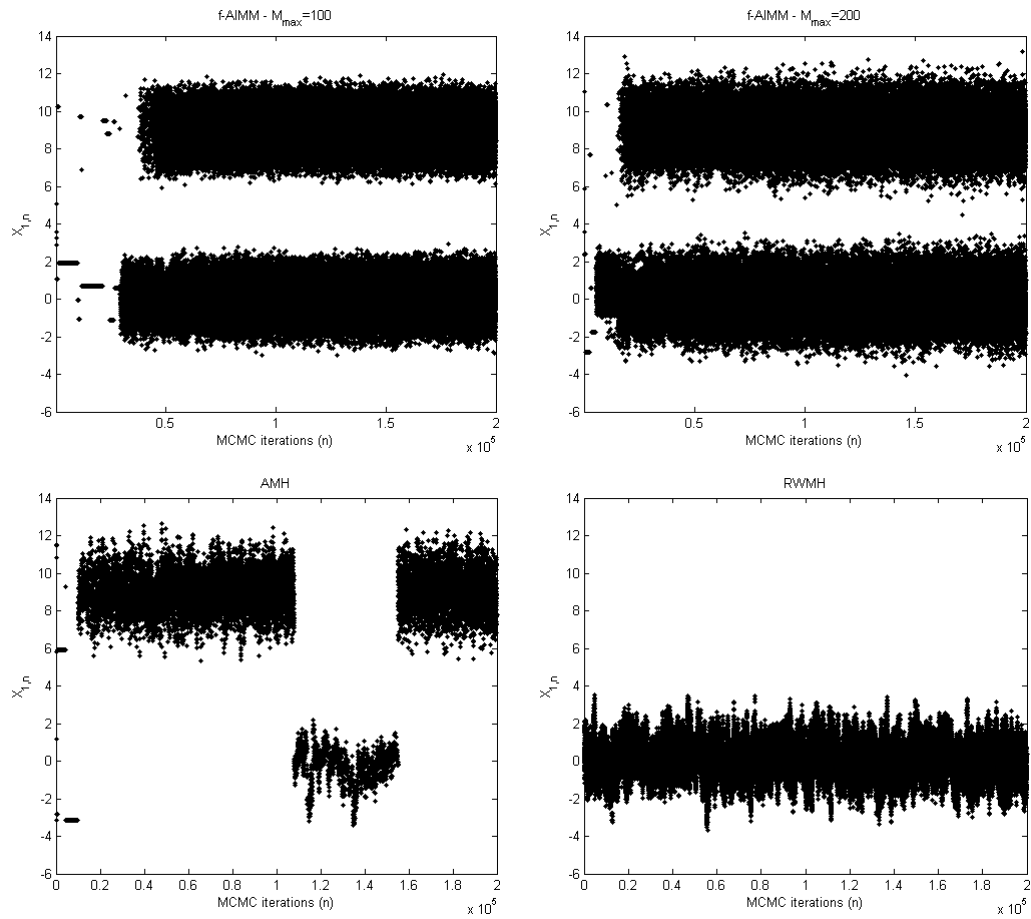


Figure 8: Example 4: π_4 target, $d = 4$ - Samples $\{X_{1,n}, n \leq 200,000\}$ from four different algorithms (from left to right, top to bottom): f-AIMM in two settings $M_{\max} = 100$ and $M_{\max} = 200$, AMH and RWMH.




Figure 9: Example 4: π_4 target, $d = 10$ - Samples $\{(X_{1,n}, X_{2,n}), n \leq 200,000\}$ from from four different algorithms (from left to right, top to bottom): f-AIMM in two settings $M_{\max} = 100$ and $M_{\max} = 200$, AGM and AMH.




Figure 10: (example 4) π_4 target, $d \in \{4, 10\}$ - Incremental mixture (colored ellipses) and uniform initial proposal (black dashed rectangle) designed by f-AIMM in the bimodal target π_4 , in dimensions $d = 4$ (top row) and $d = 10$ (bottom row). Distributions are projected onto the joint marginal space (X_1, X_2) . The colors stand for the weights $\{\beta_\ell\}_{\ell=1}^{M_n}$ attached to each component. The region inside the ellipses contains 75% of each Gaussian mass.

Table 7: Example 4: π_4 target, $d \in \{4, 10\}$ - Comparison of f-AIMM with RWMH, AMH, AGM and IM after $n = 200,000$ iterations (replicated 100 times).

(a) π_4 in dimension $d = 4$

	\bar{W}	M_{\max}	ACC	ESS	CPU	EFF	JMP
f-AIMM	5	100	.69	.30	408	$7.3 \cdot 10^{-4}$	68
RWMH	–	–	.23	$3.1 \cdot 10^{-3}$	93	$3.4 \cdot 10^{-5}$.09
AMH	–	–	.26	$2.7 \cdot 10^{-2}$	269	$1.0 \cdot 10^{-4}$.64
AGM-MH	–	100	$3.0 \cdot 10^{-3}$	$5.0 \cdot 10^{-4}$	3,517	$1.4 \cdot 10^{-7}$.29
IM	–	–	$3.3 \cdot 10^{-4}$	$5.9 \cdot 10^{-4}$	81	$7.3 \cdot 10^{-6}$.05

(b) π_4 in dimension $d = 10$

	\bar{W}	M_{\max}	ACC	ESS	CPU	EFF	JMP
f-AIMM	20	100	.45	.18	1,232	$1.4 \cdot 10^{-4}$	116.4
f-AIMM	10	200	.64	.25	1,550	$1.6 \cdot 10^{-4}$	200.1
RWMH	–	–	.38	$7.2 \cdot 10^{-4}$	476	$1.7 \cdot 10^{-6}$.07
AMH	–	–	.18	$1.3 \cdot 10^{-3}$	2,601	$5.0 \cdot 10^{-7}$.57
AGM-MH	–	100	$2.6 \cdot 10^{-4}$	$5.0 \cdot 10^{-4}$	4,459	$1.1 \cdot 10^{-7}$	$7.1 \cdot 10^{-3}$
IM	–	–	$6.4 \cdot 10^{-5}$	$6.1 \cdot 10^{-4}$	473	$1.2 \cdot 10^{-6}$	$4.1 \cdot 10^{-3}$

Table 8: Example 4: π_4 target, $d \in \{4, 10\}$ - Mean Square Error (MSE) of the mixture parameter λ , for the different algorithms (replicated 100 times).

	\bar{W}	M_{\max}	MSE(λ), $d = 4$	MSE(λ), $d = 10$
f-AIMM	5	100	.0001	.06
f-AIMM	20	100	.0006	.09
f-AIMM	10	200	.0024	.01
RWMH	–	–	.25	.25
AMH	–	–	.15	.25
AGM-MH	–	100	.22	.25
IM	–	–	.03	.20

6 Discussion

Although implicitly evaluated in an (non adaptive) independence Metropolis–Hastings transition, the information conveyed by the ratio of importance weights is lost because of the threshold set to one in the M–H acceptance probability. Indeed, while at X_n and given two realizations of the proposal Q , say \tilde{X}_{n+1} and \tilde{Y}_{n+1} , the two events $\{W(\tilde{X}_{n+1}) > W(X_n)\}$ and $\{W(\tilde{Y}_{n+1}) \gg W(X_n)\}$ result in the same transition, *i.e.* the proposed move is accepted with probability one, regardless of the magnitude difference between $W(\tilde{X}_{n+1})/W(X_n)$ and $W(\tilde{Y}_{n+1})/W(X_n)$. This paper aims to design an adaptive MCMC algorithm that makes use of this information by incrementing the independence M–H proposal distribution in the latter case and not in the former.

The general methodology, referred to as Adaptive Incremental Mixture MCMC (AIMM), presented and studied in this paper is a novel adaptive Markov chain Monte Carlo method to sample from challenging distributions. Theoretically, we establish under very mild assumptions, that if it only adapts on a compact set and has bounded jumps, this algorithm generates an ergodic process for any distribution of interest. We show that these conditions can always be enforced by using a specific implementation. In simpler implementations where those conditions may not be verified, the algorithm is nevertheless shown to work well in practice. We provide an even more efficient algorithm, referred to as f-AIMM, that guarantees that the incremental proposal evaluation does not slow down the algorithm. This algorithm can be seen as a series of AIMM algorithms where Gaussian components are progressively dropped. As a consequence, f-AIMM is compatible with the theoretical framework developed at Section 4. In particular, provided that it only adapts on a compact set and has bounded jumps, this algorithm is invariant for any target distribution. We illustrate its performance in a variety of challenging sampling scenarios.

Compared to other existing adaptive MCMC methods, AIMM needs less prior knowledge of the target. Its strategy of incrementing an initial naive proposal distribution with Gaussian kernels leads to a fully adaptive exploration of the state space. Conversely, we have shown that in some examples the adaptiveness of some other MCMC samplers may be compromised when an unwise choice of parametric family for the proposal kernel is made. The performance of AIMM depends strongly on the threshold \overline{W} which controls the adaptation rate. This parameter should be set according to the computational budget available. Based on our simulations, AIMM consistently yielded the best tradeoff between fast convergence and low variance.

The adaptive design of AIMM was inspired by Incremental Mixture Importance Sampling (IMIS) (Raftery and Bao, 2010). IMIS iteratively samples and weights particles according

to a sequence of importance distributions that adapt over time. The adaptation strategy is similar to that in AIMM: given a population of weighted particles, the next batch of particles is simulated by a Gaussian kernel centered at the particle having the largest importance weight. However IMIS and AIMM are structurally different since the former is an adaptive importance sampling method while the latter is an adaptive MCMC algorithm. Comparing them on a fair basis is difficult. In particular, the Effective Sample Size estimators for the two methods stem from different approximations. The computational efficiency of IMIS suffers from the fact that, at each iteration, the whole population of particles must be reweighted in order to maintain the consistency of the importance sampling estimator. By contrast, at each transition, AIMM evaluates only the importance weight of the new proposed state. However, since IMIS calculates the importance weight of large batches of particles, it acquires a knowledge of the state space more quickly than AIMM which accepts/rejects one particle at a time.

We therefore expect AIMM to be more efficient in situations where the exploration of the state space requires a large number of increments of the proposal and IMIS to be more efficient for short run times. To substantiate this expectation, we have compared the performance of AIMM and IMIS on the last example of Section 5 in dimension 4. Figure 11 reports the estimation of the probability $\pi_4(X_1 < -2)$ obtained through both methods for different run times. For short run times IMIS benefits from using batches of particles and gains a lot of information on π_4 in a few iterations. On the other hand, AIMM provides a more accurate estimate of $\pi_4(X_1 < -2)$ after about 150 seconds. Figure 12 illustrates the outcome of AIMM and IMIS after running them for 2,000 seconds. The mixture of incremental kernels obtained by AIMM is visually more appealing than the sequence of proposals derived by IMIS, reinforcing the results from Figure 11.

AIMM can be regarded as a transformation of IMIS, a particle-based inference method, into an adaptive Markov chain. This transformation could be applied to other adaptive importance sampling methods, thus designing Markov chains that might be more efficient than their importance sampling counterparts. In a Bayesian context, AIMM could, in addition to sampling the posterior distribution, be used to estimate intractable normalizing constants and marginal likelihoods via importance sampling. Indeed, the incremental mixture proposal produced by AIMM is an appealing importance distribution since it approximates the posterior and is straightforward to sample from.

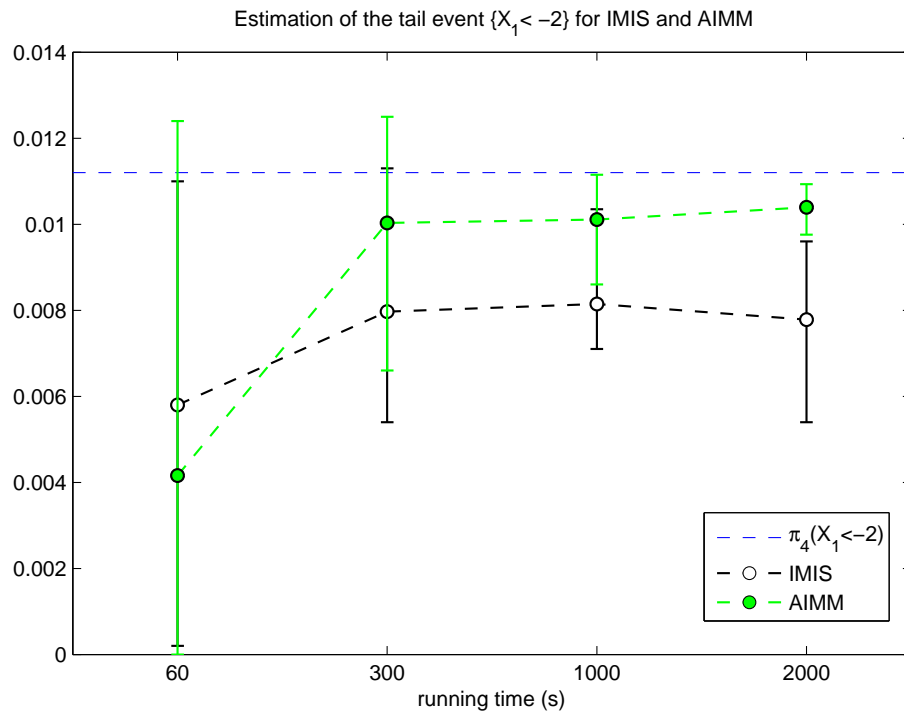


Figure 11: Example 4: π_4 target, $d = 4$ - Estimation of the probability $\pi_4(X_1 < -2)$ (true value, 0.0112) obtained through AIMM and IMIS for different runtimes (60; 300; 1,000; 2,000, in seconds).

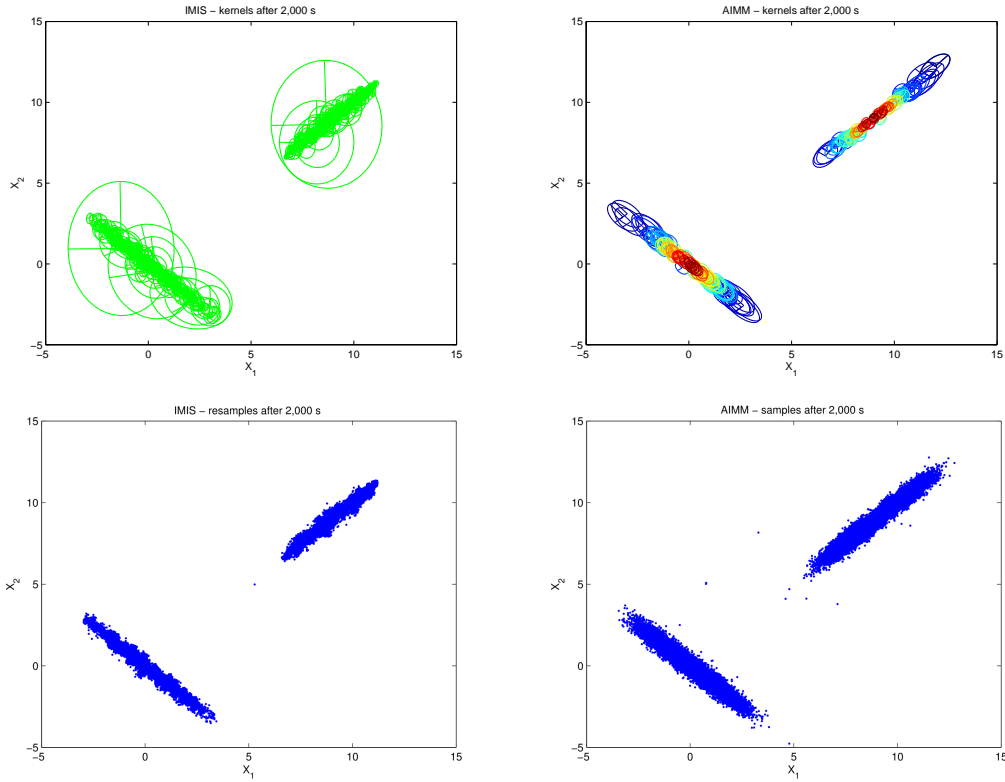


Figure 12: Example 4: π_4 target, $d = 4$, runtime of 2,000 seconds - Top: Gaussian kernels created by IMIS and AIMM (colours stand for weights assigned to the kernels and the region inside the ellipses contains 75% of each Gaussian mass). Bottom: resamples from IMIS (using the incremental proposal obtained after 2,000 seconds) and samples of the AIMM chain.

Acknowledgements

The Insight Centre for Data Analytics is supported by Science Foundation Ireland under grant SFI/12/RC/2289. Nial Friel’s research was also supported by a Science Foundation Ireland grant: 12/IP/1424. Adrian E. Raftery’s research was supported by NIH grants R01 HD054511 and R01 HD070936, Science Foundation Ireland grant 11/W.1/I2079, and the Center for Advanced Study in the Behavioral Sciences at Stanford University. Antonietta Mira’s research was supported by a Swiss National Science Foundation grant.

References

- Andrieu, C. and É. Moulines (2006). On the ergodicity properties of some adaptive MCMC algorithms. *Annals of Applied Probability* 16, 1462–1505.
- Bai, Y., G. O. Roberts, and J. S. Rosenthal (2009). On the containment condition for adaptive Markov chain Monte Carlo algorithms.
- Bates, S. C., A. Cullen, and A. E. Raftery (2003). Bayesian uncertainty assessment in multicompartment deterministic simulation models for environmental risk assessment. *Environmetrics* 14(4), 355–371.
- Cappé, O., R. Douc, A. Guillin, J.-M. Marin, and C. P. Robert (2008). Adaptive importance sampling in general mixture classes. *Statistics and Computing* 18(4), 447–459.
- Craiu, R. V., L. Gray, K. Łatuszyński, N. Madras, G. O. Roberts, J. S. Rosenthal, et al. (2015). Stability of adversarial Markov chains, with an application to adaptive MCMC algorithms. *The Annals of Applied Probability* 25(6), 3592–3623.
- Craiu, R. V., J. Rosenthal, and C. Yang (2009). Learn from thy neighbor: Parallel-chain and regional adaptive MCMC. *Journal of the American Statistical Association* 104(488), 1454–1466.
- Fort, G., E. Moulines, and P. Priouret (2011). Convergence of adaptive and interacting Markov chain Monte Carlo algorithms. *The Annals of Statistics* 39(6), 3262–3289.
- Gåsemyr, J. (2003). On an adaptive version of the Metropolis-Hastings algorithm with independent proposal distribution. *Scandinavian Journal of Statistics* 30(1), 159–173.
- Gilks, W. R., G. O. Roberts, and S. K. Sahu (1998). Adaptive Markov chain Monte Carlo through regeneration. *Journal of the American Statistical Association* 93, 1045–1054.

- Gilks, W. R. and P. Wild (1992). Adaptive rejection sampling for Gibbs sampling. *Applied Statistics* 41, 337–348.
- Giordani, P. and R. Kohn (2010). Adaptive independent Metropolis–Hastings by fast estimation of mixtures of normals. *Journal of Computational and Graphical Statistics* 19, 243–259.
- Haario, H., E. Saksman, and J. Tamminen (1999). Adaptive proposal distribution for random walk Metropolis algorithm. *Computational Statistics* 14, 375–396.
- Haario, H., E. Saksman, and J. Tamminen (2001). An adaptive Metropolis algorithm. *Bernoulli* 7, 223–242.
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* 57, 97–109.
- Hesterberg, T. (1995). Weighted average importance sampling and defensive mixture distributions. *Technometrics* 37, 185–194.
- Holden, L., R. Hauge, and M. Holden (2009). Adaptive independent Metropolis-Hastings. *Annals of Applied Probability* 19, 395–413.
- Li, J. Q. and A. R. Barron (2000). Mixture density estimation. In *Advances in neural information processing systems*, pp. 279–285.
- Liu, J. S. (1996). Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statistics and Computing* 6, 113–119.
- Luengo, D. and L. Martino (2013). Fully adaptive Gaussian mixture Metropolis-Hastings algorithm. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6148–6152.
- Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller (1953). Equation of state calculations by fast computing machines. *Journal of Chemical Physics* 21, 1087–1092.
- Mira, A. (2001). Ordering and improving the performance of Monte Carlo Markov chains. *Statistical Science*, 340–350.
- Pasarica, C. and A. Gelman (2010). Adaptively scaling the Metropolis algorithm using expected squared jumped distance. *Statistica Sinica* 20, 343–364.

- Poole, D. and A. E. Raftery (2000). Inference for deterministic simulation models: the Bayesian melding approach. *Journal of the American Statistical Association* 95(452), 1244–1255.
- Raftery, A. E. and L. Bao (2010). Estimating and projecting trends in HIV/AIDS generalized epidemics using incremental mixture importance sampling. *Biometrics* 66, 1162–1173.
- Roberts, G. O. and J. S. Rosenthal (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science* 16, 351–367.
- Roberts, G. O. and J. S. Rosenthal (2007). Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms. *Journal of Applied Probability* 44, 458–475.
- Roberts, G. O. and J. S. Rosenthal (2009). Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics* 18, 349–367.
- Rosenthal, J. S. (2003). Asymptotic variance and convergence rates of nearly-periodic Markov chain Monte Carlo algorithms. *Journal of the American Statistical Association* 98, 169–177.
- Rosenthal, J. S. et al. (2011). Optimal proposal distributions and adaptive MCMC. *Handbook of Markov Chain Monte Carlo* 4.
- Rosenthal, J. S. and J. Yang (2017). Ergodicity of combocontinuous adaptive MCMC algorithms. *Methodology and Computing in Applied Probability*, 1–17.
- Tran, M.-N., M. K. Pitt, and R. Kohn (2014). Adaptive Metropolis–Hastings sampling using reversible dependent mixture proposals. *Statistics and Computing* 26, 1–21.

Appendix A Proof of Proposition 1

In this appendix, we use the following notations

- \tilde{Q}_n is the incremental part of the proposal Q_n at iteration n , *i.e.* $\tilde{Q}_n \propto \sum_{k=1}^{n-1} \beta_k \phi_k$
- at iteration n , only $M_n \leq n-1$ components among $\phi_1, \dots, \phi_{n-1}$ have a non-zero weight
- \mathcal{E}_n is the event that the proposal increments at iteration n of AIMM, *i.e.* $\beta_n > 0$ almost surely
- \mathcal{E} is the event that the proposal increments infinitely often *i.e.* $\mathcal{E} = \bigcap_n \bigcup_{k>n} \mathcal{E}_k$

The proof of Proposition 1 is achieved by first showing that the KL divergence between \tilde{Q}_n and \tilde{Q}_{n+1} goes to zero. This allows to prove that, in the total variation norm, $\{Q_n\}_n$ adapts less and less and finally that the convergence to zero of the difference between two consecutive transition kernels $K_n - K_{n+1}$ occurs uniformly, provided that some implementation details are satisfied. We first state the following lemma and a corollary that will be needed to establish the diminishing adaptation for the AIMM process.

Lemma 1. *The sum of the weights $\bar{\beta}_n$ is not bounded in $\mathbb{P}_{\mathcal{E}}$ -probability, i.e. $\bar{\beta}_n \xrightarrow{\mathbb{P}_{\mathcal{E}}} \infty$.*

Proof. Let $M > 0$, and note that

$$\begin{aligned} \mathbb{P}_{\mathcal{E}}\{\bar{\beta}_{n+1} < M\} &= \mathbb{P}_{\mathcal{E}}\{\bar{\beta}_{n+1} < M \cap \bar{\beta}_n < M\} + \mathbb{P}_{\mathcal{E}}\{\bar{\beta}_{n+1} < M \cap \bar{\beta}_n \geq M\} \\ &= \mathbb{P}_{\mathcal{E}}\{\bar{\beta}_{n+1} < M \mid \bar{\beta}_n < M\} \mathbb{P}_{\mathcal{E}}\{\bar{\beta}_n < M\} \leq \mathbb{P}_{\mathcal{E}}\{\bar{\beta}_n < M\} \end{aligned} \quad (14)$$

since $\mathbb{P}_{\mathcal{E}}\{\bar{\beta}_{n+1} < M \cap \bar{\beta}_n \geq M\} \leq \mathbb{P}_{\mathcal{E}}\{\bar{\beta}_{n+1} < M \mid \bar{\beta}_n \geq M\} \leq \mathbb{P}_{\mathcal{E}}\{\bar{\beta}_{n+1} < \bar{\beta}_n \mid \bar{\beta}_n \geq M\} = 0$. Note that the inequality in (14) is strict on \mathcal{E}_n . Indeed,

$$\mathbb{P}_{\mathcal{E}}\{\beta_{n+1} \geq M - \bar{\beta}_n \mid \bar{\beta}_n < M, \mathcal{E}_n\} > 0 \Rightarrow \mathbb{P}_{\mathcal{E}}\{\bar{\beta}_{n+1} < M \mid \bar{\beta}_n < M, \mathcal{E}_n\} < 1, \quad (15)$$

while

$$\mathbb{P}_{\mathcal{E}}\{\bar{\beta}_{n+1} < M \mid \bar{\beta}_n < M, \bar{\mathcal{E}}_n\} = 1. \quad (16)$$

Combining (14), (15) and (16) shows that $\{\mathbb{P}_{\mathcal{E}}\{\bar{\beta}_n < M\}\}_n$ is a decreasing sequence. In particular, it is strictly decreasing an infinite number of times and therefore $\mathbb{P}_{\mathcal{E}}\{\bar{\beta}_n < M\} \rightarrow 0$ which yields $\bar{\beta}_n \xrightarrow{\mathbb{P}_{\mathcal{E}}} \infty$. \square

Corollary 1. *On \mathcal{E} , the incremental component weight i.e. $\beta_n/\bar{\beta}_n$ goes to zero in probability.*

Proof. Let $\epsilon > 0$ and note that since π is absolutely continuous we have $\sup_{x \in \mathcal{X}} \pi(x) =: \bar{\pi} < \infty$ and

$$\mathbb{P}_{\mathcal{E}} \left\{ \frac{\pi(\tilde{X}_n)^\gamma}{\bar{\beta}_n} > \epsilon \right\} \leq \mathbb{P}_{\mathcal{E}} \left\{ \frac{\bar{\pi}^\gamma}{\bar{\beta}_n} > \epsilon \right\}. \quad (17)$$

Using Lemma 1, the RHS of (17) goes to zero as $n \rightarrow \infty$ which concludes the proof. \square

A.1 Diminishing adaptation in the KL metric and total variation norm

Proposition 3. *Assume that almost surely*

A1. the covariance matrix of any component of the mixture satisfies $\det \Sigma_k > \delta > 0$

Consider the sequence $\{\tilde{D}_n\}_{n>0}$, defined as the KL divergence between the incremental components of two consecutive proposals, i.e. $\tilde{D}_n = \text{KL}\{\tilde{Q}_n, \tilde{Q}_{n+1}\}$. Then \tilde{D}_n converges to zero in probability, i.e. $\tilde{D}_n = o_{\mathbb{P}}(1)$.

Proof. For all $\epsilon > 0$, we have

$$\mathbb{P}\{\tilde{D}_n > \epsilon\} = \mathbb{P}\{\tilde{D}_n > \epsilon \cap \bar{\mathcal{E}}\} + \mathbb{P}\{\tilde{D}_n > \epsilon \cap \mathcal{E}\}. \quad (18)$$

We will show that both terms in the RHS of Eq. (18) converge to 0.

For the first term, note that $\bar{\mathcal{E}}$ is the event that the mixture will eventually stop incrementing. Therefore for all $\omega \in \bar{\mathcal{E}}$, there exists $\nu \in \mathbb{N}$ such that for all $n > n_0$, $\tilde{Q}_n = \tilde{Q}_{n+1}$. Hence, for any $n > n_0$, $\omega \notin \{\tilde{D}_n > \epsilon\}$. This implies that $\{\tilde{D}_n > \epsilon\}$ and $\bar{\mathcal{E}}$ are eventually mutually exclusive events, leading to the first term in the RHS of Eq. (18) converging to 0.

We now show that the second term in the RHS of Eq. (18) goes to zero. The case where $\mathbb{P}(\mathcal{E}) = 0$ is trivial. Indeed, $\mathbb{P}(\mathcal{E}) = 0 \Rightarrow \mathbb{P}\{\tilde{D}_n > \epsilon \cap \mathcal{E}\} = 0$ and the proof is completed. In the following, we consider the case where $0 < \mathbb{P}(\mathcal{E}) \leq 1$ and show that $\mathbb{P}\{\tilde{D}_n > \epsilon \cap \mathcal{E}\}$ converge to 0 as well. First, note that we have

$$\mathbb{P}\{\tilde{D}_n > \epsilon \cap \mathcal{E}\} \leq \mathbb{P}\{\tilde{D}_n > \epsilon \mid \mathcal{E}\}. \quad (19)$$

Defining $\bar{\beta}_n = \sum_{k \leq n} \beta_k$, we now show that the RHS of Eq. (19) goes to zero in probability.

$$\begin{aligned} \tilde{D}_n &= \int d\tilde{Q}_n \log \frac{\tilde{Q}_n}{\tilde{Q}_{n+1}} \leq \int d\tilde{Q}_n \log \frac{\tilde{Q}_n}{\tilde{Q}_{n+1}} = \int d\tilde{Q}_n \log \frac{\sum_{k=1}^{n-1} \beta_k \phi_k}{\sum_{k=1}^n \beta_k \phi_k} \times \frac{\bar{\beta}_n}{\bar{\beta}_{n-1}} \\ &\leq \int d\tilde{Q}_n \left| \log \frac{\sum_{k=1}^{n-1} \beta_k \phi_k}{\sum_{k=1}^n \beta_k \phi_k} \right| + \int d\tilde{Q}_n \left| \log \frac{\bar{\beta}_n}{\bar{\beta}_{n-1}} \right| = \int d\tilde{Q}_n \log \frac{\sum_{k=1}^n \beta_k \phi_k}{\sum_{k=1}^{n-1} \beta_k \phi_k} + \log \frac{\bar{\beta}_n}{\bar{\beta}_{n-1}}, \end{aligned}$$

where the last equality holds since for all $x \in \mathbf{X}$, $\sum_{k=1}^{n-1} \beta_k \phi_k(x) \leq \sum_{k=1}^n \beta_k \phi_k(x)$. Since we assume that π is absolutely continuous, there exists a finite constant $\bar{\pi} > 0$ such that for all $x \in \mathbf{X}$, the pdf of π satisfies $\pi(x) < \bar{\pi}$, which implies that $\beta_n < \bar{\pi}^\gamma$. Using the log inequality, this yields:

$$\tilde{D}_n \leq \bar{\pi}^\gamma \left\{ \int d\tilde{Q}_n \frac{\phi_n}{\sum_{k=1}^{n-1} \beta_k \phi_k} + \frac{1}{\bar{\beta}_{n-1}} \right\} = \frac{\bar{\pi}^\gamma}{\bar{\beta}_{n-1}} \left\{ 1 + \int d\phi_n \right\},$$

We now write that by inclusion of events,

$$\mathbb{P}_{\mathcal{E}}\{\tilde{D}_n > \epsilon\} \leq \mathbb{P}_{\mathcal{E}} \left\{ \frac{2\bar{\pi}^\gamma}{\bar{\beta}_{n-1}} > \epsilon \right\} = \mathbb{P}_{\mathcal{E}} \left\{ \bar{\beta}_{n+1} \leq \frac{2\bar{\pi}^\gamma}{\epsilon} \right\}.$$

The proof is completed by applying Lemma 1. □

Proposition 4. *Assume that almost surely*

A1. *the covariance matrix of any component of the mixture satisfies $\det \Sigma_k > \delta > 0$*

The sequence of random variables $\{D_n := \text{KL}\{Q_n, Q_{n+1}\}\}_{n>0}$ goes to zero in probability.

Proof. By straightforward algebra, we have:

$$\begin{aligned}
D_n &= \int dQ_n \log \frac{Q_n}{Q_{n+1}} = \omega_n \int dQ_0 \log \frac{Q_n}{Q_{n+1}} + (1 - \omega_n) \int d\tilde{Q}_n \log \frac{Q_n}{Q_{n+1}} \\
&\leq \omega_n \int dQ_0 \log \frac{Q_0 + (1 - \omega_n)\tilde{Q}_n}{\omega_{n+1}Q_0} + (1 - \omega_n) \int d\tilde{Q}_n \log \frac{\omega_n Q_0 + (1 - \omega_n)\tilde{Q}_n}{(1 - \omega_{n+1})\tilde{Q}_{n+1}} \\
&= \omega_n \int dQ_0 \log \frac{Q_0}{\omega_{n+1}Q_0} \left\{ 1 + \frac{(1 - \omega_n)\tilde{Q}_n}{Q_0} \right\} + (1 - \omega_n) \int d\tilde{Q}_n \log \frac{(1 - \omega_n)\tilde{Q}_n}{(1 - \omega_{n+1})\tilde{Q}_{n+1}} \left\{ 1 + \frac{\omega_n Q_0}{(1 - \omega_n)\tilde{Q}_n} \right\} \\
&\leq \omega_n \log \frac{1}{\omega_{n+1}} + \omega_n \int dQ_0 \log \left\{ 1 + \frac{(1 - \omega_n)\tilde{Q}_n}{Q_0} \right\} + (1 - \omega_n)\tilde{D}_n + (1 - \omega_n) \int d\tilde{Q}_n \log \left\{ 1 + \frac{\omega_n Q_0}{(1 - \omega_n)\tilde{Q}_n} \right\} \\
&\leq \omega_n \log \frac{1}{\omega_{n+1}} + (1 - \omega_n)\tilde{D}_n + \omega_n \int dQ_0 \frac{(1 - \omega_n)\tilde{Q}_n}{Q_0} + (1 - \omega_n) \int d\tilde{Q}_n \frac{\omega_n Q_0}{(1 - \omega_n)\tilde{Q}_n} \\
&= \omega_n \log \frac{1}{\omega_{n+1}} + (1 - \omega_n)\tilde{D}_n + \omega_n(1 - \omega_n) + (1 - \omega_n)\frac{\omega_n}{(1 - \omega_n)} \leq \tilde{D}_n + v_n,
\end{aligned}$$

where we have used the log inequality $\log(1 + x) < x$ for all $x > 0$ and defined $v_n := -\omega_n \log \omega_{n+1} + 2\omega_n$. The proof is completed by noting that

$$\mathbb{P}\{D_n > \epsilon\} \leq \mathbb{P}\{\tilde{D}_n + v_n > \epsilon\} \quad (20)$$

and applying Proposition 3 to get that $\tilde{D}_n = o_{\mathbb{P}}(1)$ which, combined with the fact that $v_n = o(1)$, shows that the RHS of Eq. (20) converges to zero. \square

Corollary 2. *Assume that almost surely*

A1. *the covariance matrix of any component of the mixture satisfies $\det \Sigma_k > \delta > 0$*

Then the difference between two consecutive proposals converges in total variation to zero in probability, i.e. $\|Q_n - Q_{n+1}\| = o_{\mathbb{P}}(1)$.

Proof. This result follows from combining Proposition 3 together with Pinsker inequality,

$$\|Q_n - Q_{n+1}\| \leq \sqrt{\frac{1}{2} \text{KL}\{Q_n, Q_{n+1}\}}. \quad (21)$$

Indeed, for all $\epsilon > 0$, we have using Eq. (21) that

$$\mathbb{P}\{\|Q_n - Q_{n+1}\|^2 > \epsilon\} \leq \mathbb{P}\{\text{KL}\{Q_n, Q_{n+1}\} > 2\epsilon\},$$

which shows that $\|Q_n - Q_{n+1}\|^2 = o_{\mathbb{P}}(1)$. The proof is completed by using the continuous mapping theorem since $x \mapsto \sqrt{x}$ is continuous on $[0, 1]$ and therefore preserves the limit. \square

A.2 Diminishing adaptation in the sup norm

We want to prove that

$$\sup_{x \in \mathsf{X}} \|K_{n+1}(x, \cdot) - K_n(x, \cdot)\|,$$

converges to zero in probability, which corresponds to the diminishing assumption of Roberts and Rosenthal (2007). We first establish the following result.

Proposition 5. *Assume that almost surely*

A1 *the covariance matrix of any component of the mixture satisfies $\det \Sigma_k > \delta > 0$*

Then the difference between two consecutive proposal probability density functions converge uniformly to zero in probability i.e. $\sup_{x \in \mathsf{X}} |Q_{n+1}(x) - Q_n(x)| = o_{\mathbb{P}}(1)$.

Proof. Let $\epsilon > 0$, we will show that $\lim_n \mathbb{P}\{\sup_{x \in \mathsf{X}} |Q_{n+1}(x) - Q_n(x)| > \epsilon\} = 0$. Using similar arguments to those in the proof of Proposition 3, we have that

$$\begin{aligned} \mathbb{P} \left\{ \sup_{x \in \mathsf{X}} |Q_{n+1}(x) - Q_n(x)| > \epsilon \right\} &= \mathbb{P} \left\{ \sup_{x \in \mathsf{X}} |Q_{n+1}(x) - Q_n(x)| > \epsilon \cap \mathcal{E} \right\} \\ &\quad + \underbrace{\mathbb{P} \left\{ \sup_{x \in \mathsf{X}} |Q_{n+1}(x) - Q_n(x)| > \epsilon \cap \bar{\mathcal{E}} \right\}}_{\searrow 0}, \end{aligned}$$

and

$$\mathbb{P} \left\{ \sup_{x \in \mathsf{X}} |Q_{n+1}(x) - Q_n(x)| > \epsilon \cap \mathcal{E} \right\} \leq \mathbb{P}_{\mathcal{E}} \left\{ \sup_{x \in \mathsf{X}} |Q_{n+1}(x) - Q_n(x)| > \epsilon \right\}.$$

Now by straightforward algebra we have that

$$\begin{aligned} |Q_{n+1}(x) - Q_n(x)| &= \left| (\omega_{n+1} - \omega_n)Q_0(x) + \frac{1 - \omega_{n+1}}{\bar{\beta}_{n+1}} \sum_{k=1}^{M_{n+1}} \beta_k \phi_k(x) - \frac{1 - \omega_n}{\bar{\beta}_n} \sum_{k=1}^{M_n} \beta_k \phi_k(x) \right| \\ &\leq \frac{\kappa}{(1 + \kappa n)(1 + \kappa(n + 1))} Q_0(x) + \sum_{k=1}^{M_n} \beta_k \phi_k(x) \left| \frac{1 - \omega_{n+1}}{\bar{\beta}_{n+1}} - \frac{1 - \omega_n}{\bar{\beta}_n} \right| + \frac{\beta_{M_{n+1}}}{\bar{\beta}_{n+1}} \phi_{M_{n+1}}(x) \\ &\leq \frac{\kappa}{(1 + \kappa n)^2} Q_0(x) + (1 - \omega_n) \left| 1 - \frac{(1 - \omega_{n+1})\bar{\beta}_n}{(1 - \omega_n)\bar{\beta}_{n+1}} \right| \tilde{Q}_n(x) + \frac{\beta_{M_{n+1}}}{\bar{\beta}_{n+1}} \phi_{M_{n+1}}(x) \\ &\leq \frac{\kappa}{(1 + \kappa n)^2} \bar{Q}_0 + \left| 1 - \frac{1 - \omega_{n+1}}{1 - \omega_n} \left(1 - \frac{\beta_{M_{n+1}}}{\bar{\beta}_{n+1}} \right) \right| \bar{Q}_n + \frac{\beta_{M_{n+1}}}{\bar{\beta}_{n+1}} \bar{\phi}_{M_{n+1}} =: R_n, \quad (22) \end{aligned}$$

where for any real-valued function $f : \mathsf{X} \rightarrow \mathbb{R}$, we have defined $\bar{f} := \sup_{x \in \mathsf{X}} f(x)$. Since $\sup_{x \in \mathsf{X}} |Q_{n+1}(x) - Q_n(x)| \leq R_n$, we have that

$$\mathbb{P}_{\mathcal{E}} \left\{ \sup_{x \in \mathsf{X}} |Q_{n+1}(x) - Q_n(x)| > \epsilon \right\} \leq \mathbb{P}_{\mathcal{E}} \{R_n > \epsilon\}.$$

The proof is concluded by noting that R_n is the sum of three random variables, each of which converges to zero in $\mathbb{P}_{\mathcal{E}}$ -probability:

- $\kappa\overline{Q_0}/(1 + \kappa n)^2 = o(1)$,
- the middle term in (22) is upper bounded by

$$\left| 1 - \frac{1 - \omega_{n+1}}{1 - \omega_n} \left(1 - \frac{\beta_{M_{n+1}}}{\bar{\beta}_{n+1}} \right) \right| \leq \left| 1 - \frac{1 - \omega_{n+1}}{1 - \omega_n} \right| + \frac{\beta_{M_{n+1}}}{\bar{\beta}_{n+1}} \frac{1 - \omega_{n+1}}{1 - \omega_n}$$

and noting that $1 - \omega_{n+1}/1 - \omega_n \rightarrow 1$, the first terms converges to zero surely and, using Corollary 1, the second one converges to zeros in $\mathbb{P}_\mathcal{E}$ -probability,

- the last term converges to zero in $\mathbb{P}_\mathcal{E}$ -probability, again using Corollary 1.

□

Proposition 6. *Assume that*

A1. the covariance matrix of any component of the mixture satisfies $\det \Sigma_k > \delta > 0$

A2. the incremental mixture (unnormalized) weights are defined as:

$$\beta_n = \frac{\eta + \pi(\tilde{X}_n)^\gamma}{(1 + \eta)^n} \mathbf{1}_{\mathcal{E}_n} \quad \text{and} \quad \omega_n = \frac{1}{1 + \sum_{k=1}^n \beta_k}, \quad (23)$$

where $\eta > 0$ is a constant, typically $\eta \ll 1$.

Then, $\{Q_n\}_n$ converges uniformly in probability to a distribution Q , i.e. $\sup_{x \in \mathcal{X}} |Q_n(x) - Q(x)| = o_{\mathbb{P}}(1)$

Proof. Uniform convergence being equivalent to uniformly Cauchy, we show that $\{Q_n\}_n$ is uniformly Cauchy, in probability, i.e.

$$\sup_{x \in \mathcal{X}} \sup_{\ell \in \mathbb{N}} |Q_{n+\ell}(x) - Q_n(x)| = o_{\mathbb{P}}(1). \quad (24)$$

Using the same reasoning that in the proof of Proposition 3, it is sufficient to prove that the convergence occurs in $\mathbb{P}_\mathcal{E}$ -probability only. Similarly to Proposition 5, we write

$$|Q_{n+\ell} - Q_n| \leq |\omega_{n+\ell} - \omega_n| \overline{Q_0} + \left| \frac{1 - \omega_{n+\ell}}{\bar{\beta}_{n+\ell}} - \frac{1 - \omega_n}{\bar{\beta}_n} \right| \sum_{k=1}^n \beta_k \overline{\phi_k} + \frac{1}{\bar{\beta}_{n+\ell}} \sum_{k=n+1}^{n+\ell} \beta_k \overline{\phi_k},$$

where we have used the same notations than in the proof of Proposition 5.

First, we note that

$$|\omega_{n+\ell} - \omega_n| = \frac{\omega_0(\bar{\beta}_{n+\ell} - \bar{\beta}_n)}{(\omega_0 + \bar{\beta}_n)(\omega_0 + \bar{\beta}_{n+\ell})} \leq \frac{\bar{\beta}_{n+\ell} - \bar{\beta}_n}{\omega_0} \quad (25)$$

and since

$$\bar{\beta}_{n+\ell} - \bar{\beta}_n = \sum_{k=n+1}^{n+\ell} \frac{\eta + \pi(\tilde{X}_k)^\gamma}{(1+\eta)^k} \leq \frac{\eta + \bar{\pi}^\gamma}{(1+\eta)^n} \sum_{k=1}^{\ell} \frac{1}{(1+\eta)^k} = \frac{\eta + \bar{\pi}^\gamma}{(1+\eta)^n} \frac{1 - \{1/(1+\eta)\}^\ell}{\eta}, \quad (26)$$

clearly $|\omega_{n+\ell} - \omega_n| \leq (\eta + \bar{\pi}^\gamma)(1+\eta)^{-n}/\eta\omega_0$. Second, we have

$$\left| \frac{1 - \omega_{n+\ell}}{\bar{\beta}_{n+\ell}} - \frac{1 - \omega_n}{\bar{\beta}_n} \right| = \left| \frac{1}{\omega_0 + \bar{\beta}_{n+\ell}} - \frac{1}{\omega_0 + \bar{\beta}_n} \right| \leq \frac{\bar{\beta}_{n+\ell} - \bar{\beta}_n}{\omega_0^2} \leq \frac{(\eta + \bar{\pi}^\gamma)}{\eta\omega_0^2(1+\eta)^n}$$

Moreover, by assumption **A2**, we have that $\bar{\phi}_k \leq 1/(2\pi)^{d/2}\sqrt{\delta}$, so

$$\sum_{k=1}^n \beta_k \bar{\phi}_k \leq \frac{1}{(2\pi)^{d/2}\sqrt{\delta}} \bar{\beta}_n.$$

Since

$$\bar{\beta}_n \leq \eta + \bar{\pi}^\gamma \sum_{k=1}^n \frac{1}{(1+\eta)^k} = \frac{\eta + \bar{\pi}^\gamma}{\eta} \{1 - (1/(1+\eta))^n\} < \frac{\eta + \bar{\pi}^\gamma}{\eta},$$

we have that

$$\left| \frac{1 - \omega_{n+\ell}}{\bar{\beta}_{n+\ell}} - \frac{1 - \omega_n}{\bar{\beta}_n} \right| \sum_{k=1}^n \beta_k \bar{\phi}_k \leq \frac{(\eta + \bar{\pi}^\gamma)^2}{\eta^2 \omega_0^2 (2\pi)^{d/2} \sqrt{\delta}} \frac{1}{(1+\eta)^n}. \quad (27)$$

Finally, using Eq. (26), we have

$$\frac{1}{\bar{\beta}_{n+\ell}} \sum_{k=n+1}^{n+\ell} \beta_k \bar{\phi}_k \leq \frac{\bar{\beta}_{n+\ell} - \bar{\beta}_n}{\omega_0} \leq \frac{\eta + \bar{\pi}^\gamma}{\eta\omega_0(1+\eta)^n}. \quad (28)$$

Combining Eqs. (25), (26), (27) and (28), we have that

$$|Q_{n+\ell}(x) - Q_n(x)| \leq Z_n := \frac{1}{(1+\eta)^n} \frac{\eta + \bar{\pi}^\gamma}{\eta\omega_0} \left\{ 1 + \bar{Q}_0 + \frac{\eta + \bar{\pi}^\gamma}{\eta\omega_0} \right\}, \quad (29)$$

so that for all $\epsilon > 0$,

$$\mathbb{P}_{\mathcal{E}} \left\{ \sup_{x,\ell} |Q_{n+\ell}(x) - Q_n(x)| > \epsilon \right\} \leq \mathbb{P}_{\mathcal{E}} \{Z_n > \epsilon\}$$

which converges to zero as $\{Z_n\}_n$ is a deterministic sequence that converges to zero. The proof is completed. \square

Proposition 7. *Assume that*

A1. *the covariance matrix of any component of the mixture satisfies $\det \Sigma_k > \delta > 0$*

A2. the incremental mixture (unnormalized) weights are defined as:

$$\beta_n = \frac{\eta + \pi(\tilde{X}_n)^\gamma}{(1 + \eta)^n} \mathbf{1}_{\mathcal{E}_n} \quad \text{and} \quad \omega_n = \frac{\omega_0}{\omega_0 + \sum_{k=1}^n \beta_k} \vee \lambda, \quad (30)$$

where $\eta > 0$ and $\lambda \in (0, 1)$ are some constants, typically $(\eta, \lambda) \ll 1$.

A3. the initial kernel Q_0 is subexponential or satisfies $Q_0(x) \propto \exp\{\psi(x)\}$ where $\psi = o(x^2)$ ($x \rightarrow \infty$), i.e. Q_0 has heavier tail than a Gaussian, e.g. multivariate Laplace or t -distribution.

Then, the function $\{H_n := 1/Q_n\}_n$ converges uniformly in probability to a function H , i.e. $\sup_{x \in \mathcal{X}} |H_n(x) - H(x)| = o_{\mathbb{P}}(1)$

Proof. As in the proof of Proposition 6, we show that $\{H_n\}_n$ is uniformly Cauchy convergent in probability. Under Assumption **A2**, we have that for all $x \in \mathcal{X}$ and $\ell \in \mathbb{N}$, for n sufficiently large, $\omega_n = \lambda$ and

$$Q_n = \lambda Q_0 + (1 - \lambda) \frac{\sum_{k=1}^n \beta_k \phi_k}{\beta_n}.$$

$$\begin{aligned} |H_{n+\ell} - H_n| &= \frac{|Q_{n+\ell} - Q_n|}{Q_{n+\ell} Q_n} \\ &= \frac{(1 - \lambda) \left| \sum_{k=1}^{n+\ell} \beta_k \phi_k / \bar{\beta}_{n+\ell} - \sum_{k=1}^n \beta_k \phi_k / \bar{\beta}_n \right|}{\left(\lambda Q_0 + (1 - \lambda) \sum_{k=1}^n \beta_k \phi_k / \bar{\beta}_n \right) \left(\lambda Q_0 + (1 - \lambda) \sum_{k=1}^{n+\ell} \beta_k \phi_k / \bar{\beta}_{n+\ell} \right)} \\ &\leq \frac{(1 - \lambda)}{\lambda^2 Q_0^2} \left| \sum_{k=1}^n \beta_k \phi_k \left(\frac{1}{\bar{\beta}_{n+\ell}} - \frac{1}{\bar{\beta}_n} \right) + \frac{1}{\bar{\beta}_{n+\ell}} \sum_{k=n+1}^{n+\ell} \beta_k \phi_k \right| \\ &\leq \frac{(1 - \lambda)}{\lambda^2} \left\{ \left| \frac{1}{\bar{\beta}_{n+\ell}} - \frac{1}{\bar{\beta}_n} \right| \sum_{k=1}^n \beta_k \frac{\phi_k}{Q_0^2} + \frac{1}{\bar{\beta}_{n+\ell}} \sum_{k=n+1}^{n+\ell} \beta_k \frac{\phi_k}{Q_0^2} \right\} \\ &\leq \frac{(1 - \lambda)}{\lambda^2 \bar{\beta}_n^2} \left\{ \sum_{k=n+1}^{n+\ell} \beta_k \sum_{k=1}^n \beta_k \frac{\phi_k}{Q_0^2} + \bar{\beta}_n \sum_{k=n+1}^{n+\ell} \beta_k \frac{\phi_k}{Q_0^2} \right\}. \end{aligned}$$

Now by assumption **A3**. the function ϕ_k/Q_0^2 is bounded for all $x \in \mathcal{X}$ and there exists $M < \infty$ such that for any Gaussian pdf ϕ , $\phi/Q_0^2 < M$. This yields:

$$\sup_{x \in \mathcal{X}} |H_{n+\ell}(x) - H_n(x)| \leq 2 \frac{(1 - \lambda) M}{\lambda^2 \bar{\beta}_n} \sum_{k=n+1}^{n+\ell} \beta_k \leq Y_n := 2 \frac{(1 - \lambda) M}{\lambda^2 \bar{\beta}_n} \frac{\eta + \bar{\pi}^\gamma}{\eta(1 + \eta)^n}.$$

Upon \mathcal{E} , for n sufficiently large $\bar{\beta}_n > 0$ almost surely and $Y_n \rightarrow 0$. The proof is completed by noting that

$$\mathbb{P}_{\mathcal{E}} \left\{ \sup_{x, \ell} |H_{n+\ell}(x) - H_n(x)| > \epsilon \right\} \leq \mathbb{P}_{\mathcal{E}} \{Y_n > \epsilon\} \rightarrow 0.$$

□

Corollary 3. *Assume that assumptions **A1**, **A2** and **A3** of Proposition 7 hold, then*

$$\sup_{x \in \mathsf{X}} \left| 1 - \frac{Q_n(x)}{Q(x)} \right| = o_{\mathbb{P}}(1).$$

Proof. We have

$$\sup_{x \in \mathsf{X}} \left| 1 - \frac{Q_n(x)}{Q(x)} \right| = \sup_{x \in \mathsf{X}} Q_n(x) \left| \frac{1}{Q_n(x)} - \frac{1}{Q(x)} \right| \leq \sup_{x \in \mathsf{X}} Q_n(x) \sup_{x \in \mathsf{X}} \left| \frac{1}{Q_n(x)} - \frac{1}{Q(x)} \right|,$$

and the proof is completed by noting that, for a sufficiently large n , we have

$$Q_n \leq \lambda \overline{Q_0} + (1 - \lambda) \frac{1}{(2\pi)^{d/2} \sqrt{\delta}}$$

and applying Proposition 7. □

We can now prove Proposition 1

Proposition 1. *Assume that **A1**–**A3** of Proposition 7 hold and that in addition, we have:*

A4. *the event \mathcal{E}_n upon which the mixture increments is defined with $\underline{W} > 0$*

Then, AIMM satisfies diminishing adaptation, i.e.

$$\sup_{x \in \mathsf{X}} \|K_{n+1}(x, \cdot) - K_n(x, \cdot)\| = o_{\mathbb{P}}(1).$$

Proof. From Proposition 6, we know that there exists a probability density function Q such that $\{Q_n\}_n$ converges uniformly to Q . Denoting by K the MH transition kernel with proposal Q , we have that

$$\|K_{n+1}(x, \cdot) - K_n(x, \cdot)\| \leq \|K_n(x, \cdot) - K(x, \cdot)\| + \|K_{n+1}(x, \cdot) - K(x, \cdot)\|.$$

To establish diminishing adaptation, it is thus sufficient to prove that $\sup_{x \in \mathsf{X}} \|K_n(x, \cdot) - K(x, \cdot)\| = o_{\mathbb{P}_{\mathcal{E}}}(1)$. In particular, we can write

$$\begin{aligned} \|K_n(x, \cdot) - K(x, \cdot)\| &= \int |(Q(y) - Q_n(y))\alpha(x, y) - Q_n(y)(\alpha_n(x, y) - \alpha(x, y))| dy + \\ &\quad \left| \int \alpha(x, y)(Q(y) - Q_n(y)) dy - \int Q_n(dy)(\alpha_n(x, y) - \alpha(x, y)) \right| \\ &\leq 2 \sup_{x \in \mathsf{X}} |Q(x) - Q_n(x)| + 2 \int Q_n(dy) |\alpha_n(x, y) - \alpha(x, y)|. \end{aligned} \quad (31)$$

Now note that for all $(x, y) \in \mathsf{X}^2$,

$$|\alpha_n(x, y) - \alpha(x, y)| \leq \left| \frac{\pi(y)Q_n(x)}{\pi(x)Q_n(y)} - \frac{\pi(y)Q(x)}{\pi(x)Q(y)} \right|$$

we have

$$\begin{aligned}
\int Q_n(\mathrm{d}y) |\alpha_n(x, y) - \alpha(x, y)| &\leq \int \frac{\pi(y)}{\pi(x)} \left| Q_n(x) - \frac{Q(x)Q_n(y)}{Q(y)} \right| \mathrm{d}y \\
&\leq \int \frac{\pi(y)}{\pi(x)} \left| Q_n(x) - Q(x) + Q(x) \left(\frac{Q_n(y)}{Q(y)} - 1 \right) \right| \mathrm{d}y \\
&\leq \int \frac{\pi(y)}{\pi(x)} |Q_n(x) - Q(x)| \mathrm{d}y + \int \frac{W(y)}{W(x)} |Q_n(y) - Q(y)| \mathrm{d}y \\
&\leq \int \frac{\pi(y)}{W(x)} \left| 1 - \frac{Q_n(x)}{Q(x)} \right| \mathrm{d}y + \int \frac{W(y)}{W(x)} |Q_n(y) - Q(y)| \mathrm{d}y \leq \frac{2}{\underline{W}} \sup_{x \in \mathsf{X}} \left| 1 - \frac{Q_n(x)}{Q(x)} \right|. \quad (32)
\end{aligned}$$

Indeed, the proposal limit Q must satisfy $\underline{W} < W := \pi/Q < \overline{W}$. Combining Eqs. (31) and (32) yields

$$\sup_{x \in \mathsf{X}} \|K_n(x, \cdot) - K(x, \cdot)\| \leq 2 \left(\sup_{x \in \mathsf{X}} |Q_n(x) - Q(x)| + \frac{1}{\underline{W}} \sup_{x \in \mathsf{X}} \left| 1 - \frac{Q_n(x)}{Q(x)} \right| \right),$$

and applying Proposition 6 and Corollary 3 completes the proof. \square

Appendix B Proof of Proposition 2

We show that the AIMM version presented in Section 4.3.2 satisfies diminishing adaptation and containment and is thus ergodic (Roberts and Rosenthal, 2007, Theorem 2).

B.1 Diminishing Adaptation

First we note that diminishing adaptation is implied by Assumption **B3**. Indeed, using the same approach that in Appendix A and introducing the quantity $\overline{\Delta}_n := \sup_{x \in \mathsf{X}} \|\overline{P}_{n+1}(x, \cdot) - \overline{P}_n(x, \cdot)\|$, we have for all $\epsilon > 0$

$$\mathbb{P}\{\overline{\Delta}_n > \epsilon\} = \mathbb{P}\{\overline{\Delta}_n > \epsilon \mid \mathcal{E}\} \mathbb{P}\{\mathcal{E}\} + \mathbb{P}\{\overline{\Delta}_n > \epsilon \mid \mathcal{E}^c\} \mathbb{P}\{\mathcal{E}^c\}, \quad (33)$$

where for readability we have denoted by \mathcal{E}^c the complementary event of \mathcal{E} ($\mathcal{E} = \bigcap_n \bigcup_{k>n} \mathcal{E}_k$). Now on \mathcal{E}^c , there exists a $n_0 < \infty$ such that for all $n > n_0$ the mixture does not increment anymore. In such a case, $\mathbb{P}\{\overline{\Delta}_n > \epsilon \mid \mathcal{E}^c\} = 0$ as soon as $n > n_0$ since for all $n > n_0$, $\mathbb{P}\{M_n = M_{n_0}\} = 1$. This implies that $\mathbb{P}\{\overline{P}_{n+1}(x, \cdot) = \overline{P}_n(x, \cdot)\} = 1$ which is straightforward if ω_n is a function of n through M_n . Now by definition of $\overline{\mathcal{E}} = \bigcap_n \bigcup_{k>n} \overline{\mathcal{E}}_k$ ($\overline{\mathcal{E}}_n = \mathcal{E}_n \cap \{M_n < M\}$), upon \mathcal{E} , there exists an $n_1 < \infty$ such that $M_{n_1} = M$ and for all $n > n_1$, $\mathbb{P}\{M_n = M\} = 1$. Similarly to the reasoning with \mathcal{E}^c , we conclude that for all $n > n_1$, $\mathbb{P}\{\overline{\Delta}_n > \epsilon \mid \mathcal{E}\} = 0$. Diminishing adaptation follows from Eq. 33 by writing that for all $\epsilon > 0$ and $n > \max(n_0, n_1)$, $\mathbb{P}\{\overline{\Delta}_n > \epsilon\} = 0$.

B.2 Containment

We now show that under the Assumptions of Proposition 2, the AIMM process satisfies the containment condition *i.e.* $C_n(\epsilon)$ defined at Eq. (10) is bounded in probability. We use Theorem 1 from Rosenthal and Yang (2017), which extends the seminal development made in Craiu et al. (2015) and in particular Theorem 21 therein. In Rosenthal and Yang (2017), adaptive MCMC are analysed in a parametric adaptation framework, *i.e.* under the assumption that there exists a set $\mathcal{Y} \subset \mathbb{R}^p$ such that any adaptive proposal belongs to a family of distributions parameterized by \mathcal{Y} . Theorem 1 (Rosenthal and Yang, 2017) identifies six assumptions, that we now enumerate for completeness, that are sufficient to guarantee the containment condition. We define for all $r > 0$, $\mathcal{K}_r := \{x \in \mathcal{X}, \min_{y \in \mathcal{K}} \|x - y\| \leq r\}$.

RY-1 \mathcal{Y} is compact.

RY-2 The process evolves through a π -Harris ergodic transition kernel that does not allow moves larger than some fixed distance $D > 0$.

RY-3 The process only adapts within a compact set \mathcal{K} .

RY-4 Q_0 is bounded above by $Q_0(dy) \leq \zeta \nu_*(dy)$ where ν_* is a probability measure concentrated on $\mathcal{K}_{2D} \setminus \mathcal{K}_D$.

RY-5 There exists two positive constants (ϵ, δ) such that the fixed kernel P , *i.e.* the M-H kernel with proposal Q_0 , is bounded below by $P(x, dy) \geq \epsilon \text{Leb}(dy)$ whenever $(x, y) \in J$ with $\|x - y\| \leq \delta$ and $J \supset \mathcal{K}_{2D} \setminus \mathcal{K}_D$.

RY-6 The function $(x, y, \gamma) \mapsto Q_\gamma(x, y)$ is combocontinuous, *i.e.* it is a finite combination of continuous functions (which generalizes the concept of piecewise continuous functions where some pieces are not necessarily intervals).

We now show that all those assumptions are satisfied by the AIMM process described at Section 4.3.2.

In the context of Proposition 2, we define \mathcal{Y} as the following union of sets

$$\mathcal{Y} := \cup_{\ell=0}^M \{A^\ell, B^\ell, C^\ell\}, \quad (34)$$

such that $A := \{\beta \in \mathbb{R}, 0 \leq \beta \leq \bar{\pi}^\gamma\}$, $B := \{\mu \in \mathbb{R}^d, -L \leq \mu_i \leq L\}$ and $C := \{\Sigma \in \mathcal{M}_d(\mathbb{R}), -L \leq \Sigma_{i,j} \leq L\}$. Indeed for all $n \in \mathbb{N}$, there exists $\gamma = \{\beta_1, \dots, \beta_\ell, \mu_1, \dots, \mu_\ell, \Sigma_1, \dots, \Sigma_\ell\}$ with $\ell \leq M$ such that $\bar{Q}_n^\mathcal{K} \equiv \bar{Q}_\gamma^\mathcal{K} := \bar{\omega}_n^\mathcal{K} Q_0 + (1 - \bar{\omega}_n^\mathcal{K}) \sum_{l=1}^\ell \beta_l \Phi(\cdot | \mu_l, \Sigma_l) / \sum_{l=1}^\ell \beta_l$ and $\gamma \in \mathcal{Y}$. Clearly, \mathcal{Y} is a finite union of compact sets and is thus compact and **RY-1** holds. By construction, the second part of **RY-2** and **RY-3** hold. Moreover, it can be noted that for

any $D > 0$, the Metropolis algorithm that consists in proposing $\tilde{X}_{n+1} \sim \bar{Q}_n^\mathcal{K}$ and accepting with probability $\mathbb{1}_{\|\tilde{X}_{n+1} - X_n\| \leq D} \bar{\alpha}^\mathcal{K}(X_n, \tilde{X}_{n+1})$, with $\bar{\alpha}^\mathcal{K}$ is defined as α (2) up to replacing Q_n by $\bar{Q}_n^\mathcal{K}$, is still π -invariant. This follows from noting that this algorithm would generate a π -reversible Markov chain. Since it is also ϕ -irreducible and aperiodic, $\bar{P}_\gamma^\mathcal{K}$ is Harris ergodic for π , for any $\gamma \in \mathcal{Y}$. This checks **RY-2**. **RY-4** is straightforward by taking Q_0 as any distribution dominated by the Lebesgue measure, *e.g.* multivariate Gaussian, multivariate T-student, etc. Assumption **RY-5** follows from Q_0 and π being absolutely continuous distributions on $(\mathbf{X}, \mathcal{X})$. Finally, we note that for a fixed $\gamma \in \mathcal{Y}$, $(x, y) \mapsto \bar{Q}_\gamma^\mathcal{K}(x, y)$ is piecewise continuous. In fact this function is continuous on the two complementary sets \mathcal{K} and $\bar{\mathcal{K}}$. Moreover, we note that for a fixed $y \in \mathbf{X}$, $(\gamma, x) \mapsto \bar{Q}_\gamma^\mathcal{K}(x, y) = \sum_{l=1}^M \mathbb{1}_{\ell(\gamma)=l} \bar{Q}_\gamma^\mathcal{K}(x, y)$, where $\ell(\gamma)$ is the number of components in the mixture $\bar{Q}_\gamma^\mathcal{K}$. Since for all $\ell \in \{0, \dots, M\}$, $x \in \mathbf{X}$, the mapping $(x, \gamma) \mapsto \bar{Q}_\gamma^\mathcal{K}(x, y)$ is piecewise continuous, $(\gamma, x) \mapsto \bar{Q}_\gamma^\mathcal{K}(x, y)$ is combocontinuous and **RY-6** holds. The proof is completed.

Appendix C Simulation details

C.1 Competing algorithms

- Adaptive Metropolis–Hastings (AMH) (Haario et al., 2001): a benchmark adaptive MCMC algorithm with non-independence proposal kernel

$$Q_n(x_n, \cdot) = \begin{cases} \Psi_d(\cdot; x_n, \Sigma_0) & \text{if } n \leq N_0 \\ p\Psi_d(\cdot; x_n, \Sigma_0) + (1-p)\Psi_d(\cdot; x_n, \Sigma_n) & \text{otherwise} \end{cases}$$

where $\Psi_d(\cdot; m, S)$ is the d -dimensional Gaussian distribution with mean m , covariance matrix S , Σ_0 is an initial covariance matrix and N_0 the number of preliminary iterations where no adaptation is made. We consider the version of AMH proposed in Roberts and Rosenthal (2009), in which there is a strictly positive probability ($p = .05$) of proposing through the initial Gaussian random walk. The covariance matrix of the adaptive part of the proposal is written $\Sigma_n = s_d \Gamma_n$, where Γ_n is defined as the empirical covariance matrix from all the previous states of the X_1, \dots, X_n , and $s_d = 2.4^2/d$ is a scaling factor. The parameters N_0 and Σ_0 were set to be equal to the corresponding AIMM parameters.

- Adaptive Gaussian Mixture Metropolis–Hastings (AGM) (Luengo and Martino, 2013): an adaptive MCMC with an independence proposal kernel defined as a mixture of Gaussian distributions

$$Q_n = \sum_{\ell=1}^M \omega_{\ell,n} \Psi_d(\cdot; \mu_{\ell,n}, \Lambda_{\ell,n}). \quad (35)$$

In AGM, the number of components M is fixed. The proposal distribution Q_n is parameterised by the vector $\Omega_n = \{\omega_{\ell,n}, \mu_{\ell,n}, \Lambda_{\ell,n}\}_{\ell=1}^M$. Given the new state X_{n+1} of the chain, the next parameter Ω_{n+1} will follow from a deterministic update ; see (Luengo and Martino, 2013, section 4) for more details. In the implementation, we set the number of components M identical to the corresponding AIMM parameter M_{\max} .

- For the sake of comparing with non-adaptive methods, we also include the Independence Sampler (IM) (see *e.g.* Liu (1996)) and the random walk Metropolis–Hastings (RWMH) (see *e.g.* Roberts and Rosenthal (2001)) using the Matlab build-in function *mhsample* with default settings.

C.2 Performance indicators

The quality of a sampler is based on its ability to explore the state space fully and quickly without getting trapped in specific states (mixing). The KL divergence between π and the chain distribution (if available) and the *Effective Sample Size* are indicators that allow us to assess those two properties. Related statistics such as the chain *Jumping distance* and the *Acceptance rate* of the sampler are also reported. We also provide the chain *Efficiency*, which penalizes the Effective Sample Size by the computational burden generated by the sampler.

- For a stochastic process $\{X_k, k \leq n\}$, we recall that when sampling from π is feasible, the KL divergence between the target and the chain distribution can be approximated by

$$\text{KL} = \frac{1}{L} \sum_{\ell=1}^L \log \left\{ \frac{\pi(Z_\ell)}{\mathcal{L}(Z_\ell | X_{1:n})} \right\}, \quad Z_\ell \sim \pi, \quad (36)$$

where $\mathcal{L}(\cdot | X_{1:n})$ is the kernel density estimation of the stochastic process (obtained using the routine `ksdensity` provided in Matlab in the default setting). We stress that using approximated values for KL are not an issue here as we are first and foremost interested in comparing different samplers, with lower KL being the better.

- The jumping distance measures how fast the sampler explores the state space, see Pesarica and Gelman (2010). For a stochastic process $\{X_k, k \leq n\}$ it is estimated by:

$$\text{JMP} = \frac{1}{n-1} \sum_{k=1}^{n-1} \|X_{k+1} - X_k\|^2, \quad (37)$$

and the larger the squared distance the better.

- The mixing rate of a of a one dimensional chain ($d = 1$) is classically evaluated with the Effective Sample Size (ESS), which is approximated by

$$\text{ESS} = 1 / \left\{ 1 + 2 \sum_{t=1}^T \hat{\rho}_t \right\}, \quad (38)$$

where $\hat{\rho}_t$ denotes the empirical lag t covariance estimated using the sample path of $X_{1:n}$, $T = \min(1000, t_0)$ and t_0 is the smaller lag such that $\hat{\rho}_{t_0+\ell} < .01$, for all $\ell > 0$. $\text{ESS} \in (0, 1)$ ($\text{ESS}=1$ corresponds to *i.i.d.* samples) and the higher ESS , the better the algorithm. When $d > 1$, ESS is set as the minimum ESS among the d marginal ESS 's.

- The tradeoff between the computational efficiency and the precision is estimated by EFF

$$\text{EFF} = \text{ESS} / \tau, \quad (39)$$

where n is the total number of iterations performed by the sampler and τ the CPU time (in second) required to achieve the N transitions (Tran et al., 2014).

C.3 AGM targeting π_2

Although similar to AIMM, AGM adapts a Gaussian mixture proposal to the target, the distribution of the resulting chain remains far from π_2 after 200,000 iterations. Figure 13 gives a hint to understand why AGM is not efficient in sampling from π_2 . By construction, AGM adapts locally a component of the proposal, provided that some samples are available in its vicinity. As a result, components initially located where the probability mass is non-negligible will adapt well to the target but the others will see their weight shrinking to zero (in dimension two, out of one hundred initial kernels having the same weight, the ten kernels with highest weight hold 0.99 of the total mixture weight after 200,000 iterations). This gives those kernels with initial poor location a high inertia, which prevents them moving away from low density regions to reach regions that are yet to be supported by the proposal. AIMM's opportunistic increment process explores the state space more efficiently.

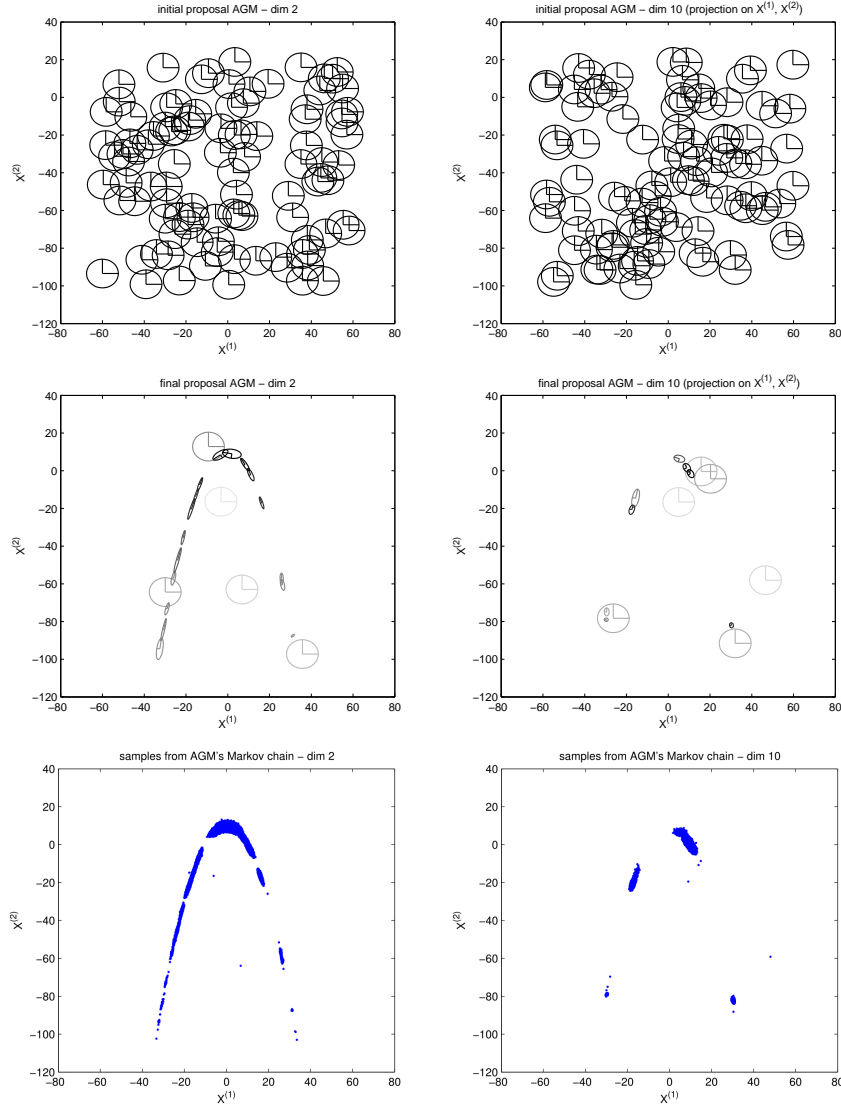


Figure 13: (example 2) π_2 target, $d \in \{2, 10\}$ - AGM sampling from π_2 in dimensions $d = 2$ (left column) and $d = 10$ (right column, projection onto the first two dimensions). First and second rows: initial mixtures with centers drawn uniformly and adapted mixtures after $n = 200,000$ iterations. The grey levels on the ellipses stand for the weights of the components (from white to black for larger weights) and the region inside the ellipses contains 75% of each Gaussian mass. Third row: samples $\{X_{1,2}, X_{2,n}, n \leq 200,000\}$ from AGM.