

Accelerating MCMC Algorithms

Christian P Robert

Université Paris Dauphine, PSL Research University, and Department of Statistics,
University of Warwick

Víctor Elvira

IMT Lille Douai & CRIStAL laboratory

Nick Tawn

Department of Statistics, University of Warwick

Changye Wu

Université Paris Dauphine, PSL Research University

Abstract. Markov chain Monte Carlo algorithms are used to simulate from complex statistical distributions by way of a local exploration of these distributions. This local feature avoids heavy requests on understanding the nature of the target, but it also potentially induces a lengthy exploration of this target, with a requirement on the number of simulations that grows with the dimension of the problem and with the complexity of the data behind it. Several techniques are available towards accelerating the convergence of these Monte Carlo algorithms, either at the exploration level (as in tempering, Hamiltonian Monte Carlo and partly deterministic methods) or at the exploitation level (with Rao-Blackwellisation and scalable methods).

1. INTRODUCTION

Markov chain Monte Carlo (MCMC) algorithms have been used for nearly 60 years, becoming a reference method for analysing Bayesian complex models in the early 1990's ([Gelfand and Smith, 1990](#)). The strength of this method is that it guarantees convergence to the quantity (or quantities) of interest with minimal requirements on the targeted distribution (also called *target*) behind such quantities. In that sense, MCMC algorithms are robust or universal, as opposed to the most standard Monte Carlo methods (see, e.g., [Rubinstein, 1981](#); [Robert and Casella, 2004](#)) that require direct simulations from the target distribution. This robustness may however induce a slow convergence behaviour in that the exploration of the relevant space—meaning the part of the space supporting the distribution that has a significant probability mass under that distribution—may take a long while, as the simulation usually proceeds by local jumps in the vicinity of the current position. In other words, MCMC—especially in its off-the-shelf versions like Gibbs sampling and Metropolis–Hastings algorithms—is very often myopic in that it provides a good illumination of a local area, while remaining unaware of the global support of the distribution. As with most other

simulation methods, there always exist ways of creating highly convergent MCMC algorithms by taking further advantage of the structure of the target distribution. Here, we mostly limit ourselves to the realistic situation where the target density is only known as the output of a computer code or to a setting similarly limited in its information content.

The approaches to the acceleration of MCMC algorithms can be divided in several categories, from those which improve our knowledge about the target distribution, to those that modify the proposal in the algorithm, including those that exploit better the outcome of the original MCMC algorithm. The following sections provide more details about these directions and the solutions proposed in the literature.

2. WHAT IS MCMC AND WHY DOES IT NEED ACCELERATING?

MCMC methods have a history (see, e.g. [Cappé and Robert, 2000](#)) that starts at approximately the same time as the Monte Carlo methods, in conjunction with the conception of the first computers. They have been devised to handle the simulation of complex target distributions, when complexity stems from the shape of the target density, the size of the associated data, the dimension of the object to be simulated, or from time requirements. For instance, the target density $\pi(\theta)$ may happen to be expressed in terms of multiple integrals that cannot be solved analytically,

$$\pi(\theta) = \int \omega(\theta, \xi) d\xi,$$

which requires the simulation of the entire vector (θ, ξ) . In cases when ξ is of the same dimension as the data, as for instance in latent variable models, this significant increase in the dimension of the object to be simulated creates computational difficulties for standard Monte Carlo methods, from managing the new target $\omega(\theta, \xi)$, to devising a new and efficient simulation algorithm. A Markov chain Monte Carlo (MCMC) algorithm allows for an alternative resolution of this computational challenge by simulating a Markov chain that explores the space of interest (and possibly supplementary spaces of auxiliary variables) without requiring a deep preliminary knowledge on the density π , besides the ability to compute $\pi(\theta_0)$ for a given parameter value θ_0 (if up to a normalising constant) and possibly the gradient $\nabla \log \pi(\theta_0)$. The validation of the method (e.g., [Robert and Casella, 2004](#)) is that the Markov chain is *ergodic* (e.g., [Meyn and Tweedie, 1993](#)), namely that it converges in distribution to the distribution with density π , no matter where the Markov chain is started at time $t = 0$.

The Metropolis–Hastings algorithm is a generic illustration of this principle. The basic algorithm is constructed by choosing a *proposal*, that is, a conditional density $K(\theta'|\theta)$ (also known as a *Markov kernel*), the Markov chain $\{\theta_t\}_{t=1}^\infty$ being then derived by successive simulations of the transition

$$\theta_{t+1} = \begin{cases} \theta' \sim K(\theta'|\theta_t) & \text{with probability } \left\{ \frac{\pi(\theta')}{\pi(\theta_t)} \times \frac{K(\theta_t|\theta')}{K(\theta'|\theta_t)} \right\} \wedge 1, \\ \theta_t & \text{otherwise.} \end{cases}$$

This acceptance-rejection feature of the algorithm makes it appropriate for targeting π as its stationary distribution if the resulting Markov chain $\{\theta_t\}_{t=1}^\infty$ is

irreducible, i.e., has a positive probability of visiting any region of the support of π in a finite number of iterations. (Stationarity can easily be shown, e.g., by using the so-called *detailed balance property* that makes the chain time-reversible, see, e.g., [Robert and Casella, 2004](#).)

Considering the initial goal of simulating samples from the target distribution π , the performances of MCMC methods like the Metropolis–Hastings algorithm above often vary quite a lot, depending primarily on the correspondance between the proposal K and the target π . For instance, if $K(\theta|\theta_t) = \pi(\theta)$, the Metropolis–Hastings algorithm reduces to i.i.d. sampling from the target, which is of course a formal option when i.i.d. sampling from π proves impossible to implement. Although there exist rare instances when the Markov chain $\{\theta_t\}_{t=1}^{\infty}$ leads to negative correlations between the successive terms of the chain, making it *more efficient* than regular i.i.d. sampling ([Liu et al., 1995](#)), the most common occurrence is one of positive correlation between the simulated values (sometimes uniformly, see [Liu et al., 1994](#)). This feature implies a reduced efficiency of the algorithm and hence requires a larger number of simulations to achieve the same precision as an approximation based on i.i.d. simulations (without accounting for differences in computing time). More generally, a MCMC algorithm may require a large number of iterations to escape the attraction of its starting point θ_0 and to reach stationarity, to the extent that some versions of such algorithms fail to converge in the time available (i.e., in practice if not in theory).

It thus makes sense to seek ways of accelerating (a) the convergence of a given MCMC algorithm to its stationary distribution, (b) the convergence of a given MCMC estimate to its expectation, and/or (c) the exploration of a given MCMC algorithm of the support of the target distribution. Those goals are related but still distinct. For instance, a chain initialised by simulating from the target distribution may still fail to explore the whole support in an acceptable number of iterations. While there is not an optimal and universal solution to this issue, we will discuss below approaches that are as generic as possible, as opposed to artificial ones taking advantage of the mathematical structure of a specific target distribution. Ideally, we aim at covering realistic situations when the target density is only known [up to a constant or an additional completion step] as the output of an existing computer code. Pragmatically, we also cover here solutions that require more efforts and calibration steps when they apply to a wide enough class of problems.

3. ACCELERATING MCMC BY EXPLOITING THE GEOMETRY OF THE TARGET

While there is no end in trying to construct more efficient and faster MCMC algorithms, and while this (endless) goal needs to account for the cost of devising such alternatives under limited resources budgets, there exist several generic solutions such that a given target can first be explored in terms of the geometry (or topology) of the density before constructing the algorithm. Although this type of methods somehow takes us away from our original purpose which was to improve upon an existing algorithm, they still make sense within this survey in that they allow for almost automated implementations.

3.1 Hamiltonian Monte Carlo

From the point of view of this review, Hamiltonian (or hybrid) Monte Carlo (HMC) is an auxiliary variable technique that takes advantage of a continuous time Markov process to sample from the target π . This approach comes from physics (Duane et al., 1987) and was popularised in statistics by Neal (1999, 2011) and MacKay (2002). Given a target $\pi(\theta)$, where $\theta \in \mathbb{R}^d$, an artificial auxiliary variable $\vartheta \in \mathbb{R}^d$ is introduced along with a density $\varpi(\vartheta|\theta)$ so that the joint distribution of (θ, ϑ) enjoys $\pi(\theta)$ as its marginal. While there is complete freedom in this representation, the HMC literature often calls ϑ the *momentum* of a particle located at θ by analogy with physics. Based on the representation of the joint distribution

$$\omega(\theta, \vartheta) = \pi(\theta)\varpi(\vartheta|\theta) \propto \exp\{-H(\theta, \vartheta)\},$$

where $H(\cdot)$ is called the *Hamiltonian*, Hamiltonian Monte Carlo (HMC) is associated with the continuous time process (θ_t, ϑ_t) generated by the so-called *Hamiltonian equations*

$$\frac{d\theta_t}{dt} = \frac{\partial H}{\partial \vartheta}(\theta_t, \vartheta_t) \quad \frac{d\vartheta_t}{dt} = -\frac{\partial H}{\partial \theta}(\theta_t, \vartheta_t),$$

which keep the Hamiltonian target stable over time, as

$$\frac{dH(\theta_t, \vartheta_t)}{dt} = \frac{\partial H}{\partial \vartheta}(\theta_t, \vartheta_t) \frac{d\vartheta_t}{dt} + \frac{\partial H}{\partial \theta}(\theta_t, \vartheta_t) \frac{d\theta_t}{dt} = 0.$$

Obviously, the above continuous time Markov process is deterministic and only explores a given level set,

$$\{(\theta, \vartheta) : H(\theta, \vartheta) = H(\theta_0, \vartheta_0)\},$$

instead of the whole augmented state space \mathbb{R}^{2d} , which induces an issue with irreducibility. An acceptable solution to this problem is to refresh the momentum, $\vartheta_t \sim \varpi(\vartheta|\theta_{t-})$, at random times $\{\tau_n\}_{n=1}^\infty$, where θ_{t-} denotes the location of θ immediately prior to time t , and the random durations $\{\tau_n - \tau_{n-1}\}_{n=2}^\infty$ follow an exponential distribution. By construction, continuous-time Hamiltonian Markov chain can be regarded as a specific piecewise deterministic Markov process using Hamiltonian dynamics (Davis, 1984, 1993; Bou-Rabee et al., 2017) and our target, π , is the marginal of its associated invariant distribution.

Before moving to the practical implementation of the concept, let us point out that the free cog in the machinery is the conditional density $\varpi(\vartheta|\theta)$, which is usually chosen as a Gaussian density with either a constant covariance matrix M corresponding to the target covariance or as a local curvature depending on θ in Riemannian Hamiltonian Monte Carlo (Girolami and Calderhead, 2011). Betancourt (2017) argues in favour of these two cases against non-Gaussian alternatives and Livingstone et al. (2017) analyse how different choices of kinetic energy in Hamiltonian Monte Carlo affect algorithm performances. For a fixed covariance matrix, the Hamiltonian equations become

$$\frac{d\theta_t}{dt} = M^{-1}\vartheta_t \quad \frac{d\vartheta_t}{dt} = \nabla\mathcal{L}(\theta_t),$$

which is the score function. The velocity (or momentum) of the process is thus driven by this score function, gradient of the log-target.

The above description remains quite conceptual in that there is no generic methodology for producing this continuous time process, since the Hamiltonian equations cannot be solved exactly in most cases. Furthermore, standard numerical solvers like Euler’s method create an unstable approximation that induces a bias as the process drifts away from its true trajectory. There exists however a discretisation simulation technique that produces a Markov chain and is well-suited to the Hamiltonian equations in that it preserves the stationary distribution (Betancourt, 2017). It is called the *symplectic integrator*, and one version in the independent case with constant covariance consists in the following (so-called *leapfrog*) steps

$$\begin{aligned}\vartheta_{t+\epsilon/2} &= \vartheta_t + \epsilon \nabla \mathcal{L}(\theta_t)/2, \\ \theta_{t+\epsilon} &= \theta_t + \epsilon M^{-1} \vartheta_{t+\epsilon/2}, \\ \vartheta_{t+\epsilon} &= \vartheta_{t+\epsilon/2} + \epsilon \nabla \mathcal{L}(\theta_{t+\epsilon})/2,\end{aligned}$$

where ϵ is the time-discretisation step. Using a proposal on ϑ_0 drawn from the Gaussian auxiliary target and deciding on the acceptance of the value of $(\theta_{T\epsilon}, \vartheta_{T\epsilon})$ by a Metropolis–Hastings step can limit the danger of missing the target. Note that the first two leapfrog steps induce a Langevin move on θ_t :

$$\theta_{t+\epsilon} = \theta_t + \epsilon^2 M^{-1} \nabla \mathcal{L}(\theta_t)/2 + \epsilon M^{-1} \vartheta_t,$$

thus connecting with the MALA algorithm discussed below (see Durmus and Moulines, 2017 for a theoretical discussion of the optimal choice of ϵ). Note that the leapfrog integrator is quite an appealing middleground between accuracy (as it is second-order accurate) and computational efficiency.

In practice, it is important to note that discretising the Hamiltonian dynamics introduces two free parameters, the step size ϵ and the trajectory length $T\epsilon$, both to be calibrated. As an empirically successful and popular variant of HMC, the “no-U-turn sampler” (NUTS) of Hoffman and Gelman (2014) adapts the value of ϵ based on primal-dual averaging. It also eliminates the need to choose the trajectory length T via a recursive algorithm that builds a set of candidate proposals for a number of forward and backward leapfrog steps and stops automatically when the simulated path steps back.

A further acceleration step in this area is proposed by Rasmussen (2003) (see also Fielding et al., 2011), namely the replacement of the exact target density $\pi(\cdot)$ by an approximation $\hat{\pi}(\cdot)$ that is much faster to compute in the many iterations of the HMC algorithm. A generic way of constructing this approximation is to rely on Gaussian processes, when interpreted as prior distributions on the target density $\pi(\cdot)$, which is only observed at some values of θ , $\pi(\theta_1), \dots, \pi(\theta_n)$ (Rasmussen and Williams, 2005). This solution is speeding up the algorithm, possibly by orders of magnitude, but it introduces a further approximation into the Monte Carlo approach, even when the true target is used at the end of the leapfrog discretisation, as in Fielding et al. (2011).

Stan (named after Stanislas Ullam, see Carpenter et al., 2017) is a computer language for Bayesian inference that, among other approximate techniques, implements the NUTS algorithm to remove hand-tuning. More precisely, Stan is a

probabilistic programming language in that the input is at the level of a statistical model, along with data, rather than the specifics of an MCMC algorithm. The algorithmic part is somehow automated, meaning that when models can be conveniently defined through this language, it offers an alternative to the sampler that produced the original chain. As an illustration of the acceleration brought by HMC, Figure 1, reproduced from Hoffman and Gelman (2014), shows the performance of NUTS, compared with both random-walk MH and Gibbs samplers.

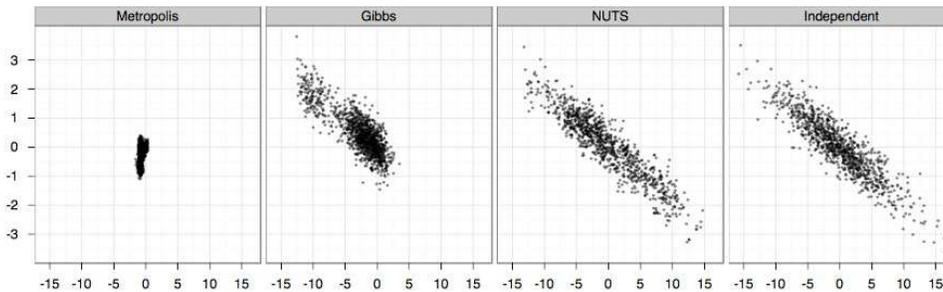


Fig 1: Comparisons between random-walk Metropolis-Hastings, Gibbs sampling, and NUTS algorithm of samples corresponding to a highly correlated 250-dimensional multivariate Gaussian target. Similar computation budgets are used for all methods to produce the 1,000 samples on display (Source: Hoffman and Gelman (2014), with permission).

4. ACCELERATING MCMC BY BREAKING THE PROBLEM INTO PIECES

The explosion in the collection and analysis of “big” datasets in recent years has brought new challenges to the MCMC algorithms that are used for Bayesian inference. When examining whether or not a new proposed sample is accepted at the accept-reject step, an MCMC algorithm such as the Metropolis-Hastings version needs to sweep over the whole data set, at each and every iteration, for the evaluation of the likelihood function. MCMC algorithms are then difficult to scale up, which strongly hinders their application in big data settings. In some cases, the datasets may be too large to fit on a single machine. It may also be that confidentiality measures impose different databases to stand on separate networks, with the possible added burden of encrypted data (Aslett et al., 2015). Communication between the separate machines may prove impossible on an MCMC scale that involves thousands or hundreds of thousands iterations.

4.1 Scalable MCMC methods

In the recent years, efforts have been made to design *scalable* algorithms, namely, solutions that manage to handle large scale targets by breaking the problem into manageable or scalable pieces. Roughly speaking, these methods can be classified into two categories (Bardenet et al., 2015): divide-and-conquer approaches and sub-sampling approaches.

Divide-and-conquer approaches partition the whole data set, denoted \mathcal{X} , into batches, $\{\mathcal{X}_1, \dots, \mathcal{X}_k\}$, and run separate MCMC algorithms on each data batch,

independently, as if they were independent Bayesian inference problems.¹ These methods then combine the simulated parameter outcomes together to approximate the original posterior distribution. Depending on the treatments of the batches selected in the MCMC stages, these approaches can be further subdivided into two finer groups: sub-posterior methods and boosted sub-posterior methods. Sub-posterior methods are motivated by the independent product equation:

$$(1) \quad \pi(\theta) \propto \prod_{i=1}^k \left(\pi_0(\theta)^{1/k} \prod_{\ell \in \mathcal{X}_i} p(x_\ell | \theta) \right) = \prod_{i=1}^k \pi_i(\theta),$$

and they target the densities $\pi_i(\theta)$ (up to a constant) in their respective MCMC steps. They thus bypass communication costs (Scott et al., 2016), by running MCMC samplers independently on each batch, and they most often increase MCMC mixing rates (in effective samples sizes produced by second), given that the sub-posterior distributions $\pi_i(\theta)$ are based on smaller datasets. For instance, Scott et al. (2016) combine the samples from the sub-posteriors, $\pi_i(\theta)$, by a Gaussian reweighting. Neiswanger et al. (2013) estimate the sub-posteriors $\pi_i(\theta)$ by non-parametric and semi-parametric methods, and they run additional MCMC samplers on the product of these estimators towards approximating the true posterior $\pi(\theta)$. Wang and Dunson (2013) refine this product estimator with an additional Weierstrass sampler, while Wang et al. (2015) estimate the posterior by partitioning the space of samples with step functions. Vehtari et al. (2014) devised an expectation propagation scheme to improve the postprocessing of the parallel samplers.

As an alternative to sampling from the sub-posteriors, boosted sub-posterior methods target instead the components

$$(2) \quad \tilde{\pi}_i(\theta) \propto \pi_0(\theta) \left(\prod_{\ell \in \mathcal{X}_i} p(x_\ell | \theta) \right)^k$$

in separate MCMC runs. Since they formally amount to repeating each batch k times towards producing pseudo data sets with the same size as the true one, the resulting boosted sub-posteriors, $\tilde{\pi}_1(\theta), \dots, \tilde{\pi}_k(\theta)$, have the same scale in variance of each component of the parameters, θ , as the true posterior, and can thus be treated as a group of estimators of the true posterior. In the subsequent combining stage, these sub-posteriors are merged together to construct a better approximation of the target distribution. For instance, Minsker et al. (2014) approximate the posterior with the geometric median of the boosted sub-posteriors, embedding them into associated reproducing kernel Hilbert spaces (rkhs), while Srivastava et al. (2015) achieve this goal using the barycentres of $\tilde{\pi}_1, \dots, \tilde{\pi}_k$, these barycentres being computed with respect to a Wasserstein distance.

In a perspective different from the above parallel scheme of divide-and-conquer approaches, sub-sampling approaches aim at reducing the number of individual datapoint likelihood evaluations operated at each iteration towards accelerating MCMC algorithms. From a general perspective, these approaches can be further

¹In order to keep the notations consistent, we still denote the target density by π , with the prior density denoted as π_0 and the sampling distribution of one observation x as $p(x|\theta)$. The dependence on the sample \mathcal{X} is not reported unless necessary.

classified into two finer classes: exact subsampling methods and approximate subsampling methods, depending on their resulting outputs. Exact subsampling approaches typically require subsets of data of random size at each iteration. One solution to this effect is taking advantage of pseudo-marginal MCMC via constructing unbiased estimators of the target density evaluated on subsets of the data (Andrieu and Roberts, 2009). Quiroz et al. (2016) follow this direction by combining the powerful debiasing technique of Rhee and Glynn (2015) and the correlated pseudo-marginal MCMC approach of Deligiannidis et al. (2015). Another direction is to use piecewise deterministic Markov processes (PDMP) (Davis, 1984, 1993), which enjoy the target distribution as the marginal of their invariant distribution. This PDMP version requires unbiased estimators of the gradients of the log-likelihood function, instead of the likelihood itself. By using a tight enough bound on the event rate function of the associated Poisson processes, PDMP can produce super-efficient scalable MCMC algorithms. The bouncy particle sampler (Bouchard-Côté et al., 2017) and the zig-zag sampler (Bierkens et al., 2016) are two competing PDMP algorithms, while Bierkens et al. (2017) unify and extend these two methods. Besides, one should note that PDMP produces a non-reversible Markov chain, which means that the algorithm should be more efficient in terms of mixing rate and asymptotic variance, when compared with reversible MCMC algorithms, such as MH, HMC and MALA, as observed in some theoretical and experimental works (Hwang et al., 1993; Sun et al., 2010; Chen and Hwang, 2013; Bierkens, 2016).

Approximate subsampling approaches aim at constructing an approximation of the target distribution. Beside the aforementioned attempts of Rasmussen (2003) and Fielding et al. (2011), one direction is to approximate the acceptance probability with high accuracy by using subsets of the data (Bardenet et al., 2014, 2015). Another solution is based on a direct modification of exact methods. The seminal work of Welling and Teh (2011), SGLD, is to exploit the Langevin diffusion

$$(3) \quad d\theta_t = \frac{1}{2} \mathbf{\Lambda} \nabla \log \pi(\theta_t) dt + \mathbf{\Lambda}^{1/2} d\mathbf{B}_t, \quad \theta_0 \in \mathbb{R}^d, t \in [0, \infty)$$

where $\mathbf{\Lambda}$ is a user-specified matrix, π is the target distribution and \mathbf{B}_t is a d -dimensional Brownian process. By virtue of the Euler-Maruyama discretisation and using unbiased estimators of the gradient of the log-target density, SGLD and its variants (Ding et al., 2014; Chen et al., 2014) often produce fast and accurate results in practice when compared with MCMC algorithms using MH steps.

Figure 2 shows the time requirements of a consensus Monte Carlo algorithm (Scott et al., 2016) compared with a Metropolis-Hastings algorithm using the whole dataset, while Figure 3 displays the saving in likelihood evaluations in confidence sampler of Bardenet et al. (2015).

4.2 Parallelisation and distributed schemes

Modern computational architectures are built with several computing units that allow for parallel processing, either fully independent or with certain communication. Although the Markovian nature of MCMC is inherently sequential and somewhat alien to the notion of parallelising, several partial solutions have been proposed in the literature for exploiting these parallel architectures. The

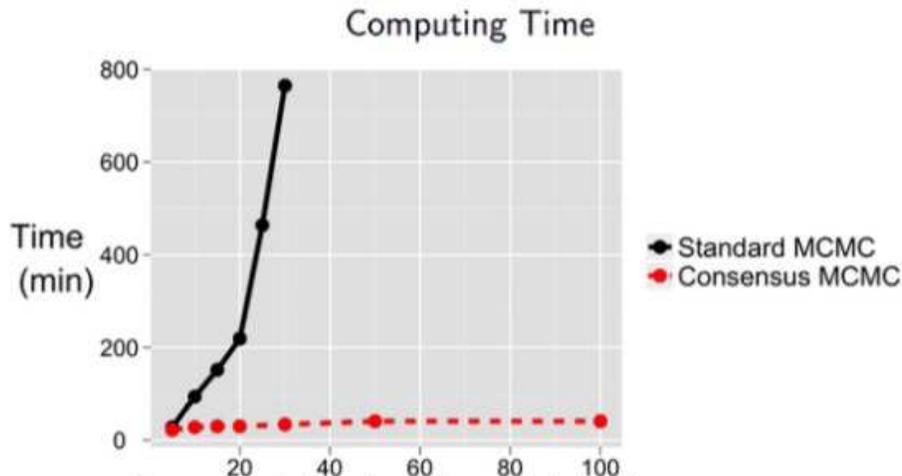


Fig 2: Elapsed time when drawing 10,000 MCMC samples with different amounts of data under the single machine and consensus Monte Carlo algorithms for a hierarchical Poisson regression. The horizontal axis represents the amounts of data. The single machine algorithm stops after 30 because of the explosion in computation budget. (Source: [Scott et al. \(2016\)](#), with permission.)

simplest approach consists in running several MCMC chains in parallel, blind to all others, until the allotted computing time is exhausted. Finally, the resulting estimators of all chains are averaged. However, this naive implementation may suffer from the fact that some of those chains have not reached their stationary regime by the end of the computation time, which then induces a bias in the resulting estimate. Ensuring that stationarity has been achieved is a difficult (if at all possible) task, although several approaches can be found in the literature ([Mykland et al., 1995](#); [Guihenneuc-Jouyaux and Robert, 1998](#); [Jacob et al., 2017](#)). At the opposite extreme, complex targets may be represented as products that involve many terms that must be evaluated, each of which can be attributed to a different thread before being multiplied all together. This strategy requires communication among processors at each MCMC step. A middle-ground version ([Jacob et al., 2011](#)) consists in running several Markov chains in parallel with periodic choices of the reference chain, all simulations being recycled through a Rao-Blackwell scheme. (See also [Calderhead, 2014](#) for a similar scheme.) The family of interacting *orthogonal* MCMC methods (O-MCMC) is proposed in [Martino et al. \(2016\)](#) with the aim of fostering better exploration of the state space, specially in high-dimensional and multimodal targets. Multiple MCMC chains are run in parallel exploring the space with random-walk proposals. The parallel chains periodically share information, also through joint MCMC steps, thus allowing an efficient combination of global (coordinated) exploration and local approximation. O-MCMC methods also allow for a parallel implementation of the Multiple Try Metropolis (MTM). In [Calderhead \(2014\)](#), a generalisation of the Metropolis-Hastings algorithm allows for a straightforward parallelisation. Each proposed point can be evaluated in a different processor at every MCMC iteration. Finally, note that the section on scalable MCMC also contains parallelisable

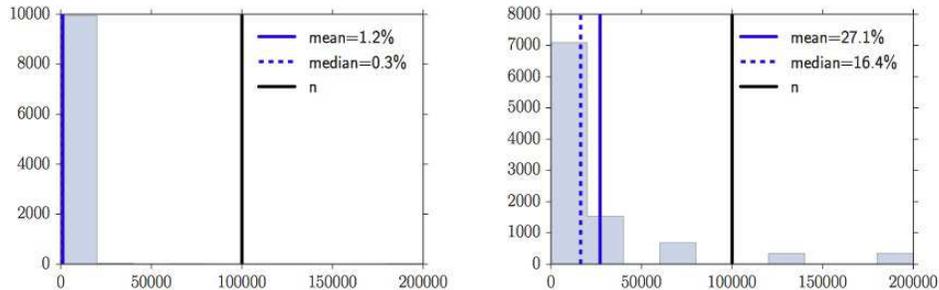


Fig 3: Percentage of numbers of data points used in each iteration of the confidence sampler with a single 2nd order Taylor approximation at θ_{MAP} . The plots describe 10,000 iterations of the confidence sampler for the posterior distribution of the mean and variance of a uni-dimensional Normal distribution with a flat prior: (*left*) 10,000 observations are generated from $\mathcal{N}(0, 1)$, (*right*) 10,000 observations are generated from $\mathcal{LN}(0, 1)$ (*Source: Bardenet et al. (2015)*, with permission).

approaches, such as the prefetching method of [Angelino et al. \(2014\)](#) (see also [Banterle et al., 2015](#) for a related approach, primarily based on an approximation of the target). A most recent endeavour called asynchronous MCMC ([Terenin et al., 2015](#)) aims at higher gains in parallelisation by reducing the amount of exchange between the parallel threads, but the notion still remains confidential at this stage.

5. ACCELERATING MCMC BY IMPROVING THE PROPOSAL

In the same spirit as the previous section, this section is stretching the purpose of this paper by considering possible modifications of the MCMC algorithm itself, rather than merely exploiting the output of a given MCMC algorithm. For instance, devising an HMC algorithm is an answer to this question even though the “improvement” is not guaranteed. Nonetheless, our argument here is that, once provided with this output, it is possible to derive new proposals in a semi-autonomous manner.

5.1 Simulated tempering

The target distribution, $\pi(\theta)$ on d -dimensional state space Θ , can exhibit multi-modality with the probability mass being located in different regions in the state space. The majority of MCMC algorithms use a localised proposal mechanism which is tuned towards local approximate optimality see, e.g., [Roberts et al. \(1997\)](#) and [Roberts and Rosenthal \(2001\)](#). By construction, these localised proposals result in the Markov chain becoming “trapped” in a subset of the state space meaning that in finite run-time the chain can entirely fail to explore other modes in the state space, leading to biased samples. Strategies to accelerate MCMC often use local gradient information and this draws the chain back towards the centre of the mode, which is the opposite of what is required in a multi-modal setting.

There is an array of methodology available to overcome issues of multi-modality in MCMC, the majority of which use state space augmentation. Auxiliary dis-

tributions that allow a Markov chain to explore the entirety of the state space are targeted and their mixing information is then passed on to aid mixing in the true target. While the sub-posteriors of the previous section can be seen as special cases of the following, the most successful and convenient implementation of these methods is to use *power-tempered target distributions*. The target distribution at inverse temperature level, β , for $\beta \in (0, 1]$ is defined as

$$\pi_\beta(\theta) = \mathfrak{K}(\beta) [\pi(\theta)]^\beta \quad \text{where} \quad \mathfrak{K}(\beta) = \left[\int [\pi(\theta)]^\beta d\theta \right]^{-1}.$$

Therefore, $\pi_1(\theta) = \pi(\theta)$. Temperatures $\beta < 1$ flatten out the target distribution allowing the chain to explore the entire state space provided the β value is sufficiently small. The simulated tempering (ST) and parallel tempering (PT) algorithms (Geyer, 1991; Marinari and Parisi, 1992) typically use the power-tempered targets to overcome the issue of multi-modality. The ST approach runs a single Markov chain on the augmented state space $\{B, \Theta\}$, where $B = \{\beta_0, \beta_1, \dots, \beta_n\}$ is a discrete collection of n inverse temperature levels with $1 = \beta_0 > \beta_1 > \dots > \beta_n > 0$. The algorithm uses a Metropolis-within-Gibbs strategy by cycling between updates in the Θ and B components of the space. For instance, a proposed temperature swap move $\beta_i \rightarrow \beta_j$ is accepted with probability

$$\min \left\{ 1, \frac{\pi_{\beta_j}(\theta)}{\pi_{\beta_i}(\theta)} \right\}$$

in order to preserve detailed balance. Note that this acceptance ratio depends on the normalisation constants $\mathfrak{K}(\beta)$ which are typically unknown, although they can sometimes be estimated, as in, e.g., Wang and Landau (2001) and Atchadé and Liu (2004). In case estimation of the marginal normalisation constants is impractical then the PT algorithm is employed. This approach simultaneously runs a Markov chain at each of the $n + 1$ temperature levels targeting the joint distribution given by $\prod_{i=0}^n [\pi(\theta_i)]^{\beta_i}$. Swap moves between chains at adjacent temperature levels are accepted according to a ratio that no longer depends on the marginal normalisation constants. Indeed, this power tempering approach has been successfully employed in a number of settings and is widely used e.g. Neal (1996), Earl and Deem (2005), Xie et al. (2010), Mohamed et al. (2012) and Carter and White (2013).

In both approaches, there is a ‘‘Goldilocks’’ principle to setting up the inverse temperature schedule. Spacings between temperature levels that are ‘‘too large’’ result in swap moves that are rarely accepted, hence delaying the transfer of hot state mixing information to the cold states. On the other hand, spacings that are too small require a large number of intermediate temperature levels, again resulting in slow mixing through the temperature space. This problem becomes even more difficult as the dimensionality of Θ increases.

Much of the historical literature suggested that a geometric spacing was optimal i.e., there exists $c \in (0, 1)$ such that $\beta_{i+1} = c\beta_i$ for $i = 0, 1, \dots, n$. However, in the case of the simulated tempering version (ST), Atchadé et al. (2011) considered the problem as an optimal scaling problem by maximising the (asymptotic in dimension) expected squared jumping distance in the B space for temperature swap moves. Under restrictive assumptions, he showed that the spacings between

consecutive inverse temperature levels should scale with dimension as $O(d^{-1/2})$ to prevent degeneracy of the swap move acceptance rate. For a practitioner the result gave guidance on optimal setup since it suggested a corresponding optimal swap move acceptance rate of 0.234 between consecutive inverse temperature levels, in accordance with [Gelman et al. \(1996\)](#). Finally, contrary to the historically recommended geometric schedule, the authors suggested that temperature schedule setup should be constructed consecutively so as to induce an approximate 0.234 swap acceptance rate between consecutive levels; which is achieved adaptively in [Miasojedow et al. \(2013\)](#). The use of expected squared jumping distance as the measure of mixing speed was justified in [Roberts and Rosenthal \(2014\)](#) where, under the same conditions as in [Atchadé et al. \(2011\)](#), it was shown that the temperature component of the ST chain has an associated diffusion process.

The target of an 0.234 acceptance rate gives good guidance to setting up the ST/PT algorithms in certain settings, but there is a major warning for practitioners following this rule for optimal setup. The assumptions made in [Atchadé et al. \(2011\)](#) and [Roberts and Rosenthal \(2014\)](#) ignore the restrictions of mixing within a temperature level, instead assuming that this can be done infinitely fast relative to the mixing within the temperature space. [Woodard et al. \(2009a\)](#), [Woodard et al. \(2009b\)](#) and [Bhatnagar and Randall \(2016\)](#) undertake a comprehensive analysis of the spectral gap of the ST/PT chains and their conclusion is rather damning of the ST/PT approaches that use power-tempered targets. Essentially, in situations where the modes have different structures, the time required to reach a given level of convergence for the ST/PT algorithms can grow exponentially in dimension. A major reason for this is that power-based tempering does not preserve the relative weights/mass between regions at the different temperature levels, see [Figure 4](#). This issue can scale exponentially in dimension. From a practical perspective, in these finite run high-dimensional non-identical modal structure settings the swap acceptance rates can be very misleading, meaning that they have limited use as a diagnostic for inter-modal mixing quality.

5.2 Adaptive MCMC

Improving and calibrating an MCMC algorithm towards a better correspondence with the intended target is a natural step in making the algorithm more efficient, provided enough information is available about this target distribution. For instance, when an MCMC sample associated with this target is available, even when it has not fully explored the range of the target, it contains some amount of information, which can then be exploited to construct new MCMC algorithms. Some of the solutions available in the literature (see, e.g. [Liang et al., 2007](#)) proceed by repeating blocks of MCMC iterations and updating the proposal K after each block, aiming at a particular optimality goal like a specific acceptance rate like 0.234 for Metropolis–Hastings steps ([Gelman et al., 1996](#)). Most versions of this method update the scale structure of a random walk proposal, based on previous realisations ([Robert and Casella, 2009](#)) or on an entire sample ([Douc et al., 2007a](#)), which turns the method into iterated importance sampling with Markovian dependence. (It can also be seen as a static version of particle filtering, [Doucet et al., 2000](#); [Andrieu and Doucet, 2002](#); [Storvik, 2002](#).)

Other adaptive resolutions bypass this preliminary and somewhat *ad hoc* con-

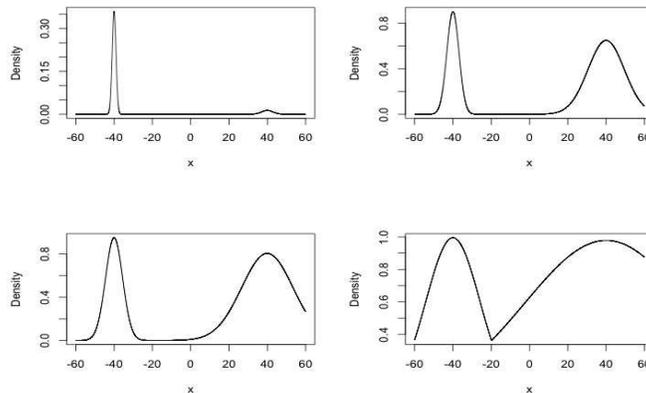


Fig 4: Un-normalised tempered target densities of a bimodal Gaussian mixture using inverse temperature levels $\beta = \{1, 0.1, 0.05, 0.005\}$ respectively. At the hot state (bottom right) it is evident that the mode centred on 40 begins to dominate the weight as β increases to ∞ even though at the cold state it was only attributable for a fraction (0.2) of the total mass.

struction and aim instead at a permanent updating within the algorithm, motivated by the idea that a continuous adaptation keeps improving the correspondence with the target. In order to preserve the validation of the method (Gelman et al., 1996; Haario et al., 1999; Roberts and Rosenthal, 2007; Saksman and Vihola, 2010), namely that the chain produced by the algorithm converges to the intended target, specific convergence results need be established, as the ergodic theorem behind standard MCMC algorithms do not apply. Without due caution (see Figure 5), an adaptive MCMC algorithm may fail to converge due to overfitting. A drawback of adaptivity is that the update of the proposal distribution relies *too much* on the earlier simulations and thus reinforces the exclusion of parts of the space that have not yet been explored.

For the validation of adaptive MCMC methods, stricter constraints must thus be imposed on the algorithm. One well-described solution (Roberts and Rosenthal, 2009) is called *diminishing adaptation*. Informally, it consists in imposing a distance between two consecutive proposal kernels to uniformly decrease to zero. In practice, this means stabilising the changes in the proposal by ridge-like factors as in the early proposal by Haario et al. (1999). A drawback of this resolution is that the decrease itself must be calibrated and may well fail to bring a significant improvement over the original proposal.

5.3 Multiple try MCMC

A completely different approach to improve the original proposal used in an MCMC algorithm is to consider a collection of proposals, built on different rationales and experiments. The *multiple try MCMC algorithm* (Liu et al., 2000; Bédard et al., 2012; Martino, 2018) follows this perspective. As the name suggests, the starting point of a multiple try MCMC algorithm is to simultaneously propose N potential moves $\theta_t^1, \dots, \theta_t^N$ of the Markov chain, instead of a single value. The proposed values θ_t^i may be independently generated according to N differ-

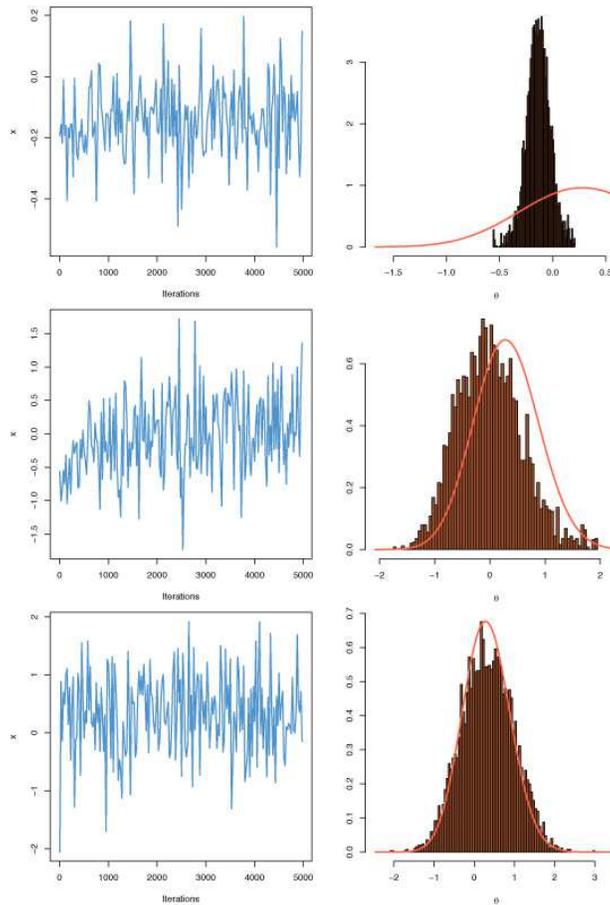


Fig 5: Markov chains produced by an adaptive algorithm where the proposal distribution is a Gaussian distribution with mean and variance computed from the past simulations of the chain. The three rows correspond to different initial distributions. The fit of the histogram of the resulting MCMC sample is poor, even for the most spread-out initial distribution (bottom) (Robert and Casella, 2004, with permission)

ent proposal densities $K_i(\cdot|\theta_t)$ that are conditional on the current value of the Markov chain, θ_t . One of the θ_t^i 's is selected based on the importance sampling weights $w_t^i \propto \pi(\theta_t^i)/K_i(\cdot|\theta_t)$. The selected value is then accepted by a further Metropolis–Hastings step which involves a ratio of normalisation constants for the importance stage, one corresponding to the selection made previously and another one created for this purpose. Indeed, besides the added cost of computing the sum of the importance weights and generating the different variates, this method faces the non-negligible drawback of requiring $N - 1$ supplementary simulations that are only used for achieving detailed balance and computing a backward summation of importance weights. This constraint may vanish when considering a collection of independent Metropolis-Hastings proposals, $q(\theta)$, but this setting is rarely realistic as it requires which make life simpler, but are less realistic since some amount of prior knowledge or experimentation to build a relevant distribution.

An alternative found in the literature is *ensemble Monte Carlo* (Iba, 2000; Cappé et al., 2008; Neal, 2011; Martino, 2018), illustrated in Figure 6 which produces a whole sample at each iteration, with target the product of the initial targets, in closer proximity with particle methods (Cappé et al., 2004; Mengersen and Robert, 2003).

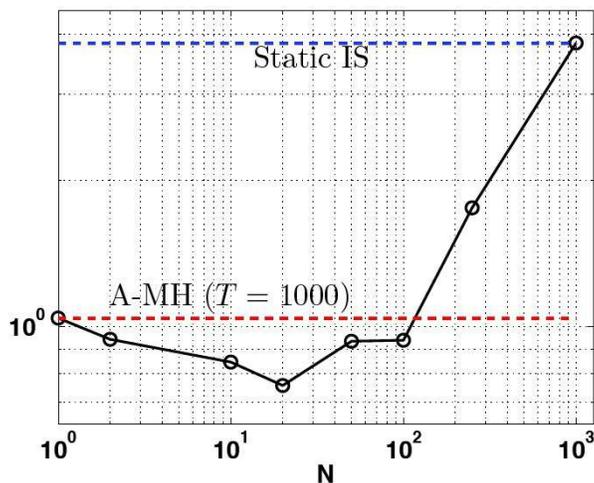


Fig 6: A comparison of an ensemble MCMC approach with a regular adaptive MCMC algorithm (lower line) and a static importance sampling approach, in terms of mean square error (MSE), for a fixed total number of likelihood evaluations, where N denotes the size of the ensemble (*Source: Martino, 2018, with permission*).

Yet another implementation of this principle is called *delayed rejection* (Tierney and Mira, 1998; Mira, 2001; Mira and Sargent, 2003), where proposals are instead considered sequentially, once the previous proposed value has been rejected. to speeding up MCMC by considering several possibilities, if sequentially. A computational difficulty with this approach is that the associated acceptance probabilities get increasingly complex as the number of delays grows, which may annihilate its appeal relative to simultaneous multiple tries. A further difficulty

is to devise the sequence of proposals in a diverse enough manner.

5.4 Multiple proposals and parameterisations

A rather basic approach to comparing proposals of MCMC algorithms is to run several in parallel and to check whether these parallel chains can be exchanged by coupling. Chains with divergent behaviour will not couple as often as chains exploring the same area. While creating multiple MCMC algorithms may seem a major challenge, automated and semi-automated schemes can be replicated as much as desired by changing the parameterisation of the target. Each change introduces a different Jacobian in the expression of the density, which means different efficiencies in the exploration of the target.

6. ACCELERATING MCMC BY REDUCING THE VARIANCE

Since the main goal of MCMC is to produce approximations for quantities of interest of the form

$$\mathfrak{J}_h = \int_{\Theta} h(\theta)\pi(\theta)d\theta,$$

an alternative (and cumulative) way of accelerating these algorithms is to improve the quality of the approximation derived from an MCMC output. That is, given an MCMC sequence $\theta_1, \dots, \theta_T$, converging to $\pi(\cdot)$, one can go beyond resorting to the basic Monte Carlo approximation

$$(4) \quad \hat{\mathfrak{J}}_h^T = 1/T \sum_{t=1}^T h(\theta_t)$$

towards reducing the variance (if not the speed of convergence) of $\hat{\mathfrak{J}}_h^T$ to \mathfrak{J}_h .

A common remark when considering Monte Carlo approximations of \mathfrak{J}_h is that the representation of the integral as an expectation is not unique (e.g. [Robert and Casella, 2004](#)). This leads to the technique of importance sampling where alternative distributions are used in replacement of $\pi(\theta)$, possibly in an adaptive manner ([Douc et al., 2007b](#)), or sequentially as in particle filters ([Del Moral et al., 2006](#); [Andrieu et al., 2011](#)). Within the framework of this essay, the outcome of a given MCMC sampler can also be exploited in several ways that lead to an improvement of the approximation of \mathfrak{J}_h .

6.1 Rao–Blackwellisation and other averaging techniques

The name ‘Rao–Blackwellisation’ was coined by [Gelfand and Smith \(1990\)](#) in their foundational Gibbs sampling paper and it has since then become a standard way of reducing the variance of integral approximations. While it essentially proceeds from the basic probability identity

$$\mathbb{E}^\pi[h(\theta)] = \mathbb{E}^{\pi_1}[\mathbb{E}^{\pi_2}\{h(\theta)|\xi\}],$$

when π can be expressed as the following marginal density

$$\pi(\theta) = \int_{\Xi} \pi_1(\xi)\pi_2(\theta|\xi)d\xi,$$

and while sufficiency does not have a clear equivalence for Monte Carlo approximation, the name stems from the Rao–Blackwell theorem ([Lehmann and Casella,](#)

1998) that improves upon a given estimator by conditioning upon a sufficient statistics. In a Monte Carlo setting, this means that (4) can be improved by a partly integrated version

$$(5) \quad \tilde{\mathcal{J}}_h^T = 1/T \sum_{t=1}^T \mathbb{E}^{\tau_2}[h(\theta)|\xi^t]$$

assuming that a second and connected sequence of simulations (ξ_t) is available and that the conditional expectation is easily constructed. For instance, Gibbs sampling (Gelfand and Smith, 1990) is often open to this Rao–Blackwell decomposition as it relies on successive simulations from several conditional distributions, possibly including auxiliary variates and nuisance parameters. In particular, a generic form of Gibbs sampling called the slice sampler (Robert and Casella, 2004) produces one or several uniform variates at each iteration.

However, a more universal type of Rao–Blackwellisation is available (Casella and Robert, 1996) for all MCMC methods involving rejection, first and foremost, Metropolis–Hastings algorithms. Indeed, first, the distribution of the rejected variables can be derived or approximated, which leads to an importance correction of the original estimator. Furthermore, the accept-reject step depends on a uniform variate, but this uniform variate can be integrated out. Namely, given a sample produced by a Metropolis–Hastings algorithm $\theta^{(1)}, \dots, \theta^{(T)}$, one can exploit both underlying samples, the proposed values $\vartheta_1, \dots, \vartheta_T$, and the uniform u_1, \dots, u_T , so that the ergodic mean can be rewritten as

$$\hat{\mathcal{J}}_h^T = 1/T \sum_{t=1}^T h(\theta^{(t)}) = 1/T \sum_{t=1}^T h(\vartheta_t) \sum_{i=t}^T \mathbb{I}_{\theta^{(i)}=\vartheta_t}.$$

The conditional expectation

$$\begin{aligned} \tilde{\mathcal{J}}_h^T &= 1/T \sum_{t=1}^T h(\vartheta_t) \mathbb{E} \left[\sum_{i=t}^T \mathbb{I}_{\theta^{(i)}=\vartheta_t} \middle| \vartheta_1, \dots, \vartheta_T \right] \\ &= 1/T \sum_{t=1}^T h(\vartheta_t) \left\{ \sum_{i=t}^T \mathbb{P}(\theta^{(i)} = \vartheta_t | \vartheta_1, \dots, \vartheta_T) \right\} \end{aligned}$$

then enjoys a smaller variance. See also Tjelmeland (2004) and Douc and Robert (2010) for connected improvements based on multiple tries. An even more rudimentary (and cheaper) version can be considered by integrating out the decision step at each Metropolis–Hastings iteration: if θ_t is the current value of the Markov chain and ϑ_t the proposed value, to be accepted (as θ_{t+1}) with probability α_t , the version

$$1/T \sum_{t=1}^T \{\alpha_t h(\vartheta_t) + (1 - \alpha_t) h(\theta_t)\}$$

should most often² bring an improvement over the basic estimate (Liu et al., 1995; Robert and Casella, 2004).

²The improvement is not universal, due to the correlation between the terms of the sum induced by the Markovian nature of the sequence $\{\theta_t\}_{t=1}^T$.

7. CONCLUSION

Accelerating MCMC algorithms may sound like a new Achilles versus tortoise paradox in that there are always methods to speed up a given algorithm. The stopping rule of this infinite regress is however that the added pain in achieving this acceleration may overcome the added gain at some point. While we have only and mostly superficially covered some of the possible directions in this survey, we thus encourage most warmly readers to keep an awareness for the potential brought by a wide array of almost cost-free accelerating solutions as well as to keep trying devising more fine-tuned improvements in every new MCMC implementation. For instance, for at least one of us, Rao-Blackwellisation is always considered at this stage. Keeping at least one such bag of tricks at one's disposal is thus strongly advised.

ACKNOWLEDGEMENTS

Christian P. Robert is grateful to Gareth Roberts, Mike Betancourt, and Julien Stoehr for helpful discussions. He is currently supported by an Institut Universitaire de France 2016–2021 senior grant. Changye Wu is currently a Ph.D. candidate at Université Paris-Dauphine and supported by a grant of the Chinese Government (CSC). Víctor Elvira acknowledges support from the *Agence Nationale de la Recherche* of France under PISCES project (ANR-17-CE40-0031-01), the Fulbright program, and the Marie Curie Fellowship (FP7/2007-2013) under REA grant agreement n. PCOFUND-GA-2013-609102, through the PRESTIGE program. The authors are quite grateful to a reviewer for his or her detailed coverage of an earlier version of the paper, which contributed to significant improvements in the presentation and coverage of the topic. All remaining errors and omissions are solely the responsibility of the authors.

REFERENCES

- ANDRIEU, C. and DOUCET, A. (2002). Particle filtering for partially observed Gaussian state space models. *J. Royal Statist. Society Series B*, **64** 827–836.
- ANDRIEU, C., DOUCET, A. and HOLENSTEIN, R. (2011). Particle Markov chain Monte Carlo (with discussion). *J. Royal Statist. Society Series B*, **72** (2) 269–342.
- ANDRIEU, C. and ROBERTS, G. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics* 697–725.
- ANGELINO, E., KOHLER, E., WATERLAND, A., SELTZER, M. and ADAMS, R. (2014). Accelerating MCMC via parallel predictive prefetching. *arXiv preprint arXiv:1403.7265*.
- ASLETT, L., ESPERANÇA, P. and HOLMES, C. (2015). A review of homomorphic encryption and software tools for encrypted statistical machine learning. [arXiv:1508.06574](https://arxiv.org/abs/1508.06574).
- ATCHADÉ, Y. F. and LIU, J. S. (2004). The Wang-Landau algorithm for Monte Carlo computation in general state spaces. *Statistica Sinica*, **20** 209–33.
- ATCHADÉ, Y. F., ROBERTS, G. and ROSENTHAL, J. (2011). Towards optimal scaling of Metropolis-coupled Markov chain Monte Carlo. *Statistics and Computing*, **21** 555–568.
- BANTERLE, M., GRAZIAN, C., LEE, A. and ROBERT, C. P. (2015). Accelerating Metropolis–Hastings algorithms by delayed acceptance. *arXiv preprint arXiv:1503.00996*.
- BARDENET, R., DOUCET, A. and HOLMES, C. (2014). Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach 405–413.
- BARDENET, R., DOUCET, A. and HOLMES, C. (2015). On Markov chain Monte Carlo methods for tall data. *J. Machine Learning Res.*, **18** 1515–1557.
- BÉDARD, M., DOUC, R. and MOULINES, E. (2012). Scaling analysis of multiple-try MCMC methods. *Stochastic Processes and their Applications*, **122** 758–786.
- BETANCOURT, M. (2017). A conceptual introduction to Hamiltonian Monte Carlo. *ArXiv e-prints*. [1701.02434](https://arxiv.org/abs/1701.02434).

- BHATNAGAR, N. and RANDALL, D. (2016). Simulated tempering and swapping on mean-field models. *Journal of Statistical Physics*, **164** 495–530.
- BIERKENS, J. (2016). Non-reversible Metropolis–Hastings. *Statistics and Computing*, **26** 1213–1228.
- BIERKENS, J., BOUCHARD-CÔTÉ, A., DOUCET, A., DUNCAN, A. B., FEARNHEAD, P., ROBERTS, G. and VOLLMER, S. J. (2017). Piecewise deterministic Markov processes for scalable Monte Carlo on restricted domains. *arXiv preprint arXiv:1701.04244*.
- BIERKENS, J., FEARNHEAD, P. and ROBERTS, G. (2016). The zig-zag process and super-efficient sampling for Bayesian analysis of big data. *arXiv preprint arXiv:1607.03188*.
- BOU-RABEE, N., SANZ-SERNA, J. M. ET AL. (2017). Randomized hamiltonian monte carlo. *The Annals of Applied Probability*, **27** 2159–2194.
- BOUCHARD-CÔTÉ, A., VOLLMER, S. J. and DOUCET, A. (2017). The bouncy particle sampler: A non-reversible rejection-free Markov chain Monte Carlo method. *Journal of the American Statistical Association*.
- CALDERHEAD, B. (2014). A general construction for parallelizing Metropolis–Hastings algorithms. *Proceedings of the National Academy of Sciences*, **111** 17408–17413.
- CAPPÉ, O., DOUC, R., GUILLIN, A., MARIN, J.-M. and ROBERT, C. (2008). Adaptive importance sampling in general mixture classes. *Statist. Comput.*, **18** 447–459.
- CAPPÉ, O., GUILLIN, A., MARIN, J.-M. and ROBERT, C. (2004). Population Monte Carlo. *J. Comput. Graph. Statist.*, **13** 907–929.
- CAPPÉ, O. and ROBERT, C. (2000). Ten years and still running! *J. American Statist. Assoc.*, **95** 1282–1286.
- CARPENTER, B., GELMAN, A., HOFFMAN, M., LEE, D., GOODRICH, B., BETANCOURT, M., BRUBAKER, M., GUO, J., LI, P. and RIDDELL, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software, Articles*, **76**.
- CARTER, J. and WHITE, D. (2013). History matching on the imperial college fault model using parallel tempering. *Computational Geosciences*, **17** 43–65.
- CASELLA, G. and ROBERT, C. (1996). Rao-Blackwellization of sampling schemes. *Biometrika*, **83** 81–94.
- CHEN, T., FOX, E. and GUESTRIN, C. (2014). Stochastic gradient hamiltonian monte carlo 1683–1691.
- CHEN, T. and HWANG, C. (2013). Accelerating reversible Markov chains. *Statistics & Probability Letters*, **83** 1956–1962.
- DAVIS, M. H. (1984). Piecewise-deterministic Markov processes: A general class of non-diffusion stochastic models. *Journal of the Royal Statistical Society. Series B (Methodological)* 353–388.
- DAVIS, M. H. (1993). *Markov Models & Optimization*, vol. 49. CRC Press.
- DEL MORAL, P., DOUCET, A. and JASRA, A. (2006). Sequential Monte Carlo samplers. *J. Royal Statist. Society Series B*, **68** 411–436.
- DELIGIANNIDIS, G., DOUCET, A. and PITT, M. K. (2015). The correlated pseudo-marginal method. *arXiv preprint arXiv:1511.04992*.
- DING, N., FANG, Y., BABBUSH, R., CHEN, C., SKEEL, R. D. and NEVEN, H. (2014). Bayesian sampling using stochastic gradient thermostats 3203–3211.
- DOUC, R., GUILLIN, A., MARIN, J.-M. and ROBERT, C. (2007a). Convergence of adaptive mixtures of importance sampling schemes. *Ann. Statist.*, **35**(1) 420–448.
- DOUC, R., GUILLIN, A., MARIN, J.-M. and ROBERT, C. (2007b). Convergence of adaptive mixtures of importance sampling schemes. *Ann. Statist.*, **35**(1) 420–448. ArXiv:0708.0711.
- DOUC, R. and ROBERT, C. (2010). A vanilla variance importance sampling via population Monte Carlo. *Ann. Statist.* To appear.
- DOUCET, A., GODSILL, S. and ANDRIEU, C. (2000). On sequential Monte-Carlo sampling methods for Bayesian filtering. *Statist. Comp.*, **10** 197–208.
- DUANE, S., KENNEDY, A. D., PENDLETON, B. J. and ROWETH, D. (1987). Hybrid Monte Carlo. *Phys. Lett. B*, **195** 216–222.
- DURMUS, A. and MOULINES, E. (2017). Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *Ann. Applied Probability*, **27** 1551–1587.
- EARL, D. J. and DEEM, M. W. (2005). Parallel tempering: Theory, Applications, and New Perspectives. *Physical Chemistry Chemical Physics*, **7** 3910–3916.
- FIELDING, M., NOTT, D. J. and LIONG, S.-Y. (2011). Efficient MCMC schemes for computationally expensive posterior distributions. *Technometrics*, **53** 16–28. <https://doi.org/10.1198/TECH.2010.09195>, URL <https://doi.org/10.1198/TECH.2010.09195>.

- GELFAND, A. and SMITH, A. (1990). Sampling based approaches to calculating marginal densities. *J. American Statist. Assoc.*, **85** 398–409.
- GELMAN, A., GILKS, W. and ROBERTS, G. (1996). Efficient Metropolis jumping rules. In *Bayesian Statistics 5* (J. Berger, J. Bernardo, A. Dawid, D. Lindley and A. Smith, eds.). Oxford University Press, Oxford, 599–608.
- GEYER, C. J. (1991). Markov chain Monte Carlo maximum likelihood. *Computing Science and Statistics*, **23** 156–163.
- GIROLAMI, M. and CALDERHEAD, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **73** 123–214.
- GUIHENNEUC-JOUYAUX, C. and ROBERT, C. P. (1998). Discretization of continuous Markov chains and Markov chain Monte Carlo convergence assessment. *Journal of the American Statistical Association*, **93** 1055–1067.
- HAARIO, H., SAKSMAN, E. and TAMMINEN, J. (1999). Adaptive proposal distribution for random walk Metropolis algorithm. *Computational Statistics*, **14(3)** 375–395.
- HOFFMAN, M. D. and GELMAN, A. (2014). The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Machine Learning Res.*, **15** 1593–1623.
- HWANG, C.-R., HWANG-MA, S.-Y. and SHEU, S.-J. (1993). Accelerating gaussian diffusions. *The Annals of Applied Probability* 897–913.
- IBA, Y. (2000). Population-based Monte Carlo algorithms. *Trans. Japanese Soc. Artificial Intell.*, **16** 279–286.
- JACOB, P. E., O’LEARY, J. and ATCHADÉ, Y. F. (2017). Unbiased Markov chain Monte Carlo with couplings. *ArXiv e-prints*. [1708.03625](https://arxiv.org/abs/1708.03625).
- JACOB, P., ROBERT, C. P. and SMITH, M. H. (2011). Using parallel computation to improve independent Metropolis–Hastings based estimation. *Journal of Computational and Graphical Statistics*, **20** 616–635.
- LEHMANN, E. and CASELLA, G. (1998). *Theory of Point Estimation (revised edition)*. Springer-Verlag, New York.
- LIANG, F., LIU, C. and CARROLL, R. (2007). Stochastic approximation in Monte Carlo computation. *JASA*, **102** 305–320.
- LIU, J., WONG, W. and KONG, A. (1994). Covariance structure of the Gibbs sampler with application to the comparison of estimators and augmentation schemes. *Biometrika*, **81** 27–40.
- LIU, J., WONG, W. and KONG, A. (1995). Covariance structure and convergence rates of the Gibbs sampler with various scans. *Journal of Royal Statistical Society*, **B 57** 157–169.
- LIU, J. S., LIANG, F. and WONG, W. H. (2000). The multiple-try method and local optimization in Metropolis sampling. *J. American Statist. Assoc.*, **95** 121–134.
- LIVINGSTONE, S., FAULKNER, M. F. and ROBERTS, G. O. (2017). Kinetic energy choice in hamiltonian/hybrid monte carlo. *arXiv preprint arXiv:1706.02649*.
- MACKAY, D. J. C. (2002). *Information Theory, Inference & Learning Algorithms*. Cambridge University Press, Cambridge, UK.
- MARINARI, E. and PARISI, G. (1992). Simulated tempering: a new Monte Carlo scheme. *EPL (Europhysics Letters)*, **19** 451.
- MARTINO, L. (2018). A Review of Multiple Try MCMC algorithms for Signal Processing. *ArXiv e-prints*. [1801.09065](https://arxiv.org/abs/1801.09065).
- MARTINO, L., ELVIRA, V., LUENGO, D., CORANDER, J. and LOUZADA, F. (2016). Orthogonal parallel MCMC methods for sampling and optimization. *Digital Signal Processing*, **58** 64–84.
- MENGERSEN, K. and ROBERT, C. (2003). Iid sampling with self-avoiding particle filters: the pinball sampler. In *Bayesian Statistics 7* (J. Bernardo, M. Bayarri, J. Berger, A. Dawid, D. Heckerman, A. Smith and M. West, eds.). Oxford University Press, Oxford.
- MEYN, S. and TWEEDIE, R. (1993). *Markov Chains and Stochastic Stability*. Springer-Verlag, New York.
- MIASOJEDOW, B., MOULINES, E. and VIHOLA, M. (2013). An adaptive parallel tempering algorithm. *Journal of Computational and Graphical Statistics*, **22** 649–664.
- MINSKER, S., SRIVASTAVA, S., LIN, L. and DUNSON, D. (2014). Scalable and robust Bayesian inference via the median posterior 1656–1664.
- MIRA, A. (2001). "on metropolis-hastings algorithms with delayed rejection". *Metron*, **59 (3-4)** 231–241.
- MIRA, A. and SARGENT, D. J. (2003). A new strategy for speeding Markov chain Monte Carlo

- algorithms. *Stat. Methods Appl.*, **12** 49–60.
- MOHAMED, L., CALDERHEAD, B., FILIPPONE, M., CHRISTIE, M. and GIROLAMI, M. (2012). Population mcmc methods for history matching and uncertainty quantification. *Computational Geosciences*, **16** 423–436.
- MYKLAND, P., TIERNEY, L. and YU, B. (1995). Regeneration in Markov chain samplers. *Journal of the American Statistical Association*, **90** 233–241.
- NEAL, R. (1999). *Bayesian Learning for Neural Networks*, vol. 118. Springer–Verlag, New York. Lecture Notes.
- NEAL, R. (2011). MCMC using Hamiltonian dynamics. In *In Handbook of Markov Chain Monte Carlo* (S. Brooks, A. Gelman, G. L. Jones and X.-L. Meng, eds.). CRC Press, New York.
- NEAL, R. M. (1996). Sampling from Multimodal Distributions using Tempered Transitions. *Statistics and Computing*, **6** 353–366.
- NEAL, R. M. (2011). MCMC Using Ensembles of States for Problems with Fast and Slow Variables such as Gaussian Process Regression. *ArXiv e-prints*. [1101.0387](https://arxiv.org/abs/1101.0387).
- NEISWANGER, W., WANG, C. and XING, E. (2013). Asymptotically exact, embarrassingly parallel MCMC. *arXiv preprint arXiv:1311.4780*.
- QUIROZ, M., VILLANI, M. and KOHN, R. (2016). Exact subsampling MCMC. *arXiv preprint arXiv:1603.08232*.
- RASMUSSEN, C. E. (2003). Gaussian processes to speed up hybrid monte carlo for expensive bayesian integrals. In *Bayesian Statistics 7* (J. Bernardo, M. Bayarri, J. Berger, A. Dawid, D. Heckerman, A. Smith and M. West, eds.). Oxford University Press, 651–659.
- RASMUSSEN, C. E. and WILLIAMS, C. K. I. (2005). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.
- RHEE, C.-H. and GLYNN, P. W. (2015). Unbiased estimation with square root convergence for sde models. *Operations Research*, **63** 1026–1043.
- ROBERT, C. and CASELLA, G. (2004). *Monte Carlo Statistical Methods*. 2nd ed. Springer-Verlag, New York.
- ROBERT, C. and CASELLA, G. (2009). *Introducing Monte Carlo Methods with R*. Springer-Verlag, New York.
- ROBERTS, G., GELMAN, A. and GILKS, W. R. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability*, **7** 110–120.
- ROBERTS, G. and ROSENTHAL, J. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science*, **16** 351–367.
- ROBERTS, G. and ROSENTHAL, J. (2007). Coupling and ergodicity of adaptive Markov Chain Monte carlo algorithms. *J. Applied Proba.*, **44**(2) 458–475.
- ROBERTS, G. and ROSENTHAL, J. (2009). Examples of adaptive MCMC. *J. Comp. Graph. Stat.*, **18** 349–367.
- ROBERTS, G. and ROSENTHAL, J. (2014). Minimising MCMC variance via diffusion limits, with an application to simulated tempering. *The Annals of Applied Probability*, **24** 131–149.
- RUBINSTEIN, R. Y. (1981). *Simulation and the Monte Carlo Method*. J. Wiley, New York.
- SAKSMAN, E. and VIHOLA, M. (2010). On the ergodicity of the adaptive Metropolis algorithm on unbounded domains. *Ann. Applied Probability*, **20**(6) 2178–2203.
- SCOTT, S. L., BLOCKER, A. W., BONASSI, F. V., CHIPMAN, H. A., GEORGE, E. I. and McCULLOCH, R. E. (2016). Bayes and big data: The consensus Monte Carlo algorithm. *International Journal of Management Science and Engineering Management*, **11** 78–88.
- SRIVASTAVA, S., CEVHER, V., DINH, Q. and DUNSON, D. (2015). Wasp: Scalable Bayes via barycenters of subset posteriors 912–920.
- STORVIK, G. (2002). Particle filters for state space models with the presence of static parameters. *IEEE Trans. Signal Process.*, **50** 281–289.
- SUN, Y., SCHMIDHUBER, J. and GOMEZ, F. J. (2010). Improving the asymptotic performance of Markov chain Monte-carlo by inserting vortices 2235–2243.
- TERENIN, A., SIMPSON, D. and DRAPER, D. (2015). Asynchronous Gibbs Sampling. *ArXiv e-prints*. [1509.08999](https://arxiv.org/abs/1509.08999).
- TIERNEY, L. and MIRA, A. (1998). Some adaptive Monte Carlo methods for Bayesian inference. *Statistics in Medicine*, **18** 2507–2515.
- TJELMELAND, H. (2004). Using all Metropolis–Hastings proposals to estimate mean values. Tech. Rep. 4, Norwegian University of Science and Technology, Trondheim, Norway.
- VEHTARI, A., GELMAN, A., SIVULA, T., JYLÄNKI, P., TRAN, D., SAHAI, S., BLOMSTEDT, P., CUNNINGHAM, J. P., SCHIMINOVICH, D. and ROBERT, C. P. (2014). Expectation propagation

- as a way of life: A framework for Bayesian inference on partitioned data. *ArXiv e-prints*. [1412.4869](#).
- WANG, F. and LANDAU, D. (2001). Determining the density of states for classical statistical models: A random walk algorithm to produce a flat histogram. *Physical Review E*, **64** 056101.
- WANG, X. and DUNSON, D. (2013). Parallelizing MCMC via weierstrass sampler. *arXiv preprint arXiv:1312.4605*.
- WANG, X., GUO, F., HELLER, K. and DUNSON, D. (2015). Parallelizing MCMC with random partition trees. *Advances in Neural Information Processing Systems* 451–459.
- WELLING, M. and TEH, Y. (2011). Bayesian learning via stochastic gradient Langevin dynamics 681–688.
- WOODARD, D. B., SCHMIDLER, S. C. and HUBER, M. (2009a). Conditions for rapid mixing of parallel and simulated tempering on multimodal distributions. *The Annals of Applied Probability* 617–640.
- WOODARD, D. B., SCHMIDLER, S. C. and HUBER, M. (2009b). Sufficient conditions for torpid mixing of parallel and simulated tempering. *Electronic Journal of Probability*, **14** 780–804.
- XIE, Y., ZHOU, J. and JIANG, S. (2010). Parallel tempering monte carlo simulations of lysozyme orientation on charged surfaces. *The Journal of chemical physics*, **132** 02B602.