Stat Comput (2015) 25:835–862
DOI 10.1007/s11222-015-9574-5

Bayesian computation: a summary of the current state, and samples backwards and forwards

Peter J. Green[1] · **Krzysztof Łatuszyński**[2] · **Marcelo Pereyra**[1] · **Christian P. Robert**[2]

Accepted: 8 May 2015 / Published online: 11 June 2015
© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract Recent decades have seen enormous improvements in computational inference for statistical models; there have been competitive continual enhancements in a wide range of computational tools. In Bayesian inference, first and foremost, MCMC techniques have continued to evolve, moving from random walk proposals to Langevin drift, to Hamiltonian Monte Carlo, and so on, with both theoretical and algorithmic innovations opening new opportunities to practitioners. However, this impressive evolution in capacity is confronted by an even steeper increase in the complexity of the datasets to be addressed. The difficulties of modelling and then handling ever more complex datasets most likely call for a new type of tool for computational inference that dramatically reduces the dimension and size of the raw data while capturing its essential aspects. Approximate models and algorithms may thus be at the core of the next computational revolution.

Keywords Bayesian analysis · MCMC algorithms · ABC techniques · Optimisation

✉ Peter J. Green
P.J.Green@bristol.ac.uk

Krzysztof Łatuszyński
K.G.Latuszynski@warwick.ac.uk

Marcelo Pereyra
Marcelo.Pereyra@bristol.ac.uk

Christian P. Robert
bayesianstatistics@gmail.com

[1] School of Mathematics, University of Bristol, Bristol, UK

[2] Department of Statistics, University of Warwick, Coventry, UK

1 Introduction

One may reasonably balk at the terms "computational statistics" and "Bayesian computation" since, from its very start, statistics has always involved some computational step to extract information, something manageable like an estimator or a prediction, from raw data. This necessarily incomplete and unavoidably biased review of the recent past, current state, and immediate future of algorithms for Bayesian inference thus first requires us to explain what we mean by computation in a statistical context, before turning to what we perceive as medium term solutions and possible deadends.

Computations are an issue in statistics whenever processing a dataset becomes a difficulty, a liability, or even an impossibility. Obviously, the computational challenge varies according to the time when it is faced: what was an issue in the nineteenth century is most likely not so any longer (take for instance the derivation of the moment estimates of a mixture of two normal distributions so painstakenly set by Pearson 1894 for estimating the ratio of "forehead" breadth–body length on a dataset of 1000 crabs or the intense algebraic derivations found in the analysis of variance of the 1950s and 1960s Searle et al. 1992).

The introduction of simulation tools in the 1940s followed hard on the heels of the invention of the computer and certainly contributed an impetus towards faster and better computers, at least in the first decade of this revolution. This shows that these tools were both needed, and unavailable without electronic calculators. The introduction of Markov chain Monte Carlo is harder to pin down as some partial versions can be traced all the way back to 1944–1945 and the Manhattan project at Los Alamos (Metropolis 1987). It is surprisingly much later, i.e., only by the early 1990s, that such methods became part of the Bayesian toolbox, that is, some time after the devising of other computer-dependent

tools like the bootstrap or the EM algorithm, and despite the availability of personal computers that considerably eased programming and experimenting (Robert and Casella 2011). It is presumably pointless to try to attribute this delay to a definite cause but a certain lack of probabilistic culture within the statistics community is probably partly to blame.

What makes this time-lag in MCMC methods becoming assimilated into the statistics community even more surprising is that fact that Bayesian inference having a significant role in statistical practice was really on hold pending the discovery of flexible computational tools that (implicitly or explicitly) delivered values for the medium–high-dimensional integrals that underpin the calculation of posterior distributions, in all but toy problems where conjugacy provided explicit answers. In fact, until Bayesians discovered MCMC, the only computational methodology that seemed to offer much chance of making practical Bayesian statistics practical was the portfolio of quadrature methods developed under Adrian Smith's leadership at Nottingham (Naylor and Smith 1982; Smith et al. 1985, 1987).

The very first article in the first issue of *Statistics and Computing*, whose quarter-century we celebrate in this special issue, was (to the journal's credit!) on Bayesian analysis, and was precisely in this direction of using clever quadrature methods to approach moderately high-dimensional posterior analysis (Dellaportas and Wright 1991). By the next (second) issue, sampling-based methods had started to appear, with three papers out of five in the issue on or related to Gibbs sampling (Verdinelli and Wasserman 1991; Carlin and Gelfand 1991; Wakefield et al. 1991).

Now, reflecting upon the evolution of MCMC methods over the 25 or so years they have been at the forefront of Bayesian inference, the focus has evolved a long way, from hierarchical models that extended the linear, mixed and generalised linear models (Albert 1988; Carlin et al. 1992; Bennett et al. 1996) which were initially the focus, and graphical models that stemmed from image analysis (Geman and Geman 1984) and artificial intelligence, to dynamical models driven by ODE's (Wilkinson 2011b) and diffusions (Roberts and Stramer 2001; Dellaportas et al. 2004; Beskos et al. 2006), hidden trees (Larget and Simon 1999; Huelsenbeck and Ronquist 2001; Chipman et al. 2008; Aldous et al. 2008) and graphs, aside with decision making in highly complex graphical models. While research on MCMC theory and methodology is still active and continually branching (Papaspiliopoulos et al. 2007; Andrieu and Roberts 2009; Łatuszyński et al. 2011; Douc and Robert 2011), progressively incorporating the capacities of parallel processors and GPUs (Lee et al. 2009; Jacob et al. 2011; Strid 2010; Suchard et al. 2010; Scott et al. 2013; Calderhead 2014), we wonder if we are not currently facing a new era where those methods are no longer appropriate to undertake the analysis of new models, and of new formulations where models are no longer completely defined. We indeed believe that imprecise models, incomplete information and summarised data will become, if not already, a central aspect of statistical analysis, due to the massive influx of data and the need to provide non-statisticians with efficient tools. This is why we also cover in this survey the notions of approximate Bayesian computation (ABC) and comment on the use of optimisation tools.

The plan of the paper is that in Sects. 2 and 3 we discuss recent progress and current issues in Markov chain Monte Carlo and ABC, respectively. In Sect. 4, we highlight some areas of modern optimisation that, through lack of familiarity, are making less impact in the mainstream of Bayesian computation than we think justified. Our Discussion in Sect. 5 raises issues about data science and relevance to applications, and looks to the future.

2 MCMC, targeting the posterior

When MCMC techniques were introduced to the mainstream statistical (Bayesian) community in 1990, they were received with skepticism that they could 1 day become the central tool of Bayesian inference. For instance, despite the assurance provided by the ergodic theorem, many researchers thought at first that the convergence of those algorithms was a mere theoretical anticipation rather than a practical reality, in contrast to traditional Monte Carlo methods, and hence that they could not be trusted to provide "exact" answers. This perspective is obviously obsolete by now, when MCMC output is considered as "exact" as regular Monte Carlo, if possibly less efficient in some settings. Nowadays, MCMC is again attracting more attention (than in the past decade, say, where developments were more about alternatives, some of which described in the following sections), both because of methodological developments linked to better theoretical tools, for instance in the handling of stochastic processes, and because of new advances in accelerated computing via parallel and cloud computing.

2.1 Basics of MCMC

The introduction of Markov chain based methods within Monte Carlo thus took a certain amount of argument to reach the mainstream statistical community, when compared with other groups who were using MCMC methods 10–30 years earlier. It may sound unlikely at the current stage of our knowledge, but using methods that (a) generated correlated output, (b) required some burnin time to remove the impact of the initial distribution and (c) did not lead to a closed form expression for asymptotic variances were indeed met with resistance at first. As often, the immense computing advantages offered by this versatile tool soon overcame the reluctance to accept those methods as similarly "exact"

as other Monte Carlo techniques, applications driving the move from the early 1990s. We reproduce below the generic version of the "all variables at once" Metropolis–Hastings algorithm (Metropolis et al. 1953; Hastings 1970; Besag et al. 1995; Robert and Casella 2011) as it (still) constitutes in our opinion a fundamental advance in computational statistics, namely that, given a computable density π (up to a normalising constant) on Θ, and a proposal Markov kernel $q(\cdot|\cdot)$, there exists a universal machine that returns a Markov chain with the proper stationary distribution, hence an associated operational MCMC algorithm.

Algorithm 1 Metropolis–Hastings algorithm (generic version)

Choose a starting value $\theta^{(0)}$
for $n = 1$ to N **do**
  Generate θ^* from a proposal $q(\cdot|\theta^{(n-1)})$
  Compute the acceptance probability

$$\rho^{(n)} = 1 \wedge \pi(\theta^*)\, q(\theta^{(n-1)}|\theta^*) \big/ \pi(\theta^{(n-1)})\, q(\theta^*|\theta^{(n-1)})$$

  Generate $u_n \sim \mathcal{U}(0,\ 1)$ and take $\theta^{(n)} = \theta^*$ if $u_n \leq \rho^{(n)}$, $\theta^{(n)} = \theta^{(n-1)}$ otherwise.
end for

The first observation about the Metropolis–Hastings is that the flexibility in choosing q is a blessing, but also a curse since the choice determines the performance of the algorithm. Hence a large part of the research on MCMC along the past 30 years (if we arbitrarily set the starting date at Geman and Geman 1984) has been on choice of the proposal q to improve the efficiency of the algorithm, and in characterising its convergence properties. This typically requires gathering or computing additional information about π and we discuss some of the fundamental strategies in subsequent sections. Algorithm 1, and its variants in which variables are updated singly or in blocks according to some schedule, remains a keystone in standard use of MCMC methodology, even though the newer Hamiltonian Monte Carlo (HMC) approach (see Sect. 2.3) may sooner or later come to replace it. While there is nothing intrinsically unique to the nature of this algorithm, or optimal in its convergence properties (other than the result of Peskun 1973 on the optimality of the acceptance ratio), attempts to bypass Metropolis–Hastings are few and limited. For instance, the birth-and-death process developed by Stephens (2000) used a continuous time jump process to explore a set of models, only to be later shown (Cappé et al. 2002) to be equivalent to the (Metropolis–Hastings) reversible jump approach of Green (1995).

Another aspect of the generic Metropolis–Hastings that became central more recently is that while the accept–reject step does overcome need to know the normalising constant, it still requires π, if unnormalised, and this may be too expensive to compute or even intractable for complicated models and large datasets. Much recent research effort has been devoted to the design and understanding of appropriate modifications that use estimators or approximations of π instead and we will take the opportunity to summarise some of the progress in this direction.

2.2 MALA and manifold MALA

Stochastic differential equations (SDEs) have been and still are informing Monte Carlo development in a number of seminal ways. A key insight is that the Langevin diffusion on Θ solving

$$d\theta_t = \frac{1}{2}\nabla \log \pi\,(\theta_t)\, dt + dB_t, \tag{1}$$

has π as its stationary and limiting distribution. Here B_t is the standard Brownian motion and ∇ denotes gradient. The crude approach of sampling an Euler discretisation (Kloeden and Platen 1992) of (1) and using it as an approximate sample from π was introduced in the applied literature (Ermak 1975; Doll and Dion 1976). The method results in a Markov chain evolving according to the dynamics

$$\begin{aligned}\theta^{(n)}|\theta^{(n-1)} &\sim Q\left(\theta^{(n-1)},\ \cdot\right)\\ &:= \theta^{(n-1)} + \frac{h}{2}\nabla \log \pi\left(\theta^{(n-1)}\right) + h^{1/2} N\,(0,\ I_{d\times d})\,,\end{aligned} \tag{2}$$

for a chosen discretisation step h. There is a delicate trade-off between accuracy of the approximation improving as $h \to 0$ and sampling efficiency [as measured, e.g., by the effective sample size (ESS)] improving when h increases. This solution was soon followed by its Metropolised version (Rossky et al. 1978) that uses the Euler approximation of (2) to produce a proposal in the Metropolis–Hastings algorithm 1, by letting $q(\cdot|\theta^{(n-1)}) := \theta^{(n-1)} + \frac{h}{2}\nabla \log \pi(\theta^{(n-1)}) + h^{1/2} N(0,\ I_{d\times d})$. While in the probability community Langevin diffusions and their equilibrium distributions had also been around for some time (Kent 1978), it was the Roberts and Tweedie (1996a) paper (motivated by Besag 1994 comment on Grenander and Miller 1994) that brought the approach to the centre of interest of the computational statistics community and sparked systematic study, development and applications of Metropolis adjusted Langevin algorithms (hence MALA) and their cousins.

There is a large body of empirical evidence that at the extra price of computing the gradient, MALA algorithms typically provide a substantial speed-up in convergence on certain types of problems. However for very light-tailed distributions the drift term may grow to infinity and cause additional instability. More precisely, for distributions

with sufficiently smooth contours, MALA is geometrically ergodic (c.f. Roberts and Rosenthal 2004) if the tails of π decay as $\exp\{-|\theta|^\beta\}$ with $\beta \in [1, 2]$, while the random walk Metropolis algorithm is geometrically ergodic for all $\beta \geq 1$ (Roberts and Tweedie 1996a; Mengersen and Tweedie 1996). The lack of geometrical ergodicity has been precisely quantified by Bou-Rabee and Hairer (2012).

Various refinements and extensions have been proposed. These include optimal scaling and choice of the discretisation step h, adaptive versions (both discussed in Sect. 2.4), combinations with proximal operators (Pereyra 2015; Schreck et al. 2013), and applications and algorithm development for the infinite-dimensional context (Pillai et al. 2012; Cotter et al. 2013). One particular direction of active research is considering a more general version of Eq. (1) with state-dependent drift and diffusion coefficient

$$d\theta_t = \left(\frac{\sigma(\theta_t)}{2}\nabla \log \pi(\theta_t) + \frac{\gamma(\theta_t)}{2}\right) dt + \sqrt{\sigma}(\theta_t)\, dB_t$$
$$\gamma_i(\theta_t) = \sum_j \frac{\partial \sigma_{ij}(\theta_t)}{\partial \theta_j}, \tag{3}$$

which also has π as invariant distribution (Xifara et al. 2014, c.f. Kent 1978). The resulting proposals are

$$q\left(\cdot|\theta^{(n-1)}\right) := \frac{h}{2}\left(\sigma\left(\theta^{(n-1)}\right)\nabla \log \pi\left(\theta^{(n-1)}\right) + \gamma\left(\theta^{(n-1)}\right)\right) + h^{1/2} N\left(0,\ \sigma\left(\theta^{(n-1)}\right)\right) + \theta^{(n-1)}.$$

Choosing appropriate σ for improved ergodicity is however nontrivial. The idea has been explored in Stramer and Tweedie (1999a, 1999b), Roberts and Stramer (2002) and more recently Girolami and Calderhead (2011) introduced a mathematically-coherent approach of relating σ to a metric tensor on a Riemannian manifold of probability distributions. The resulting algorithms are termed Manifold MALA (MMALA), Simplified MMALA (Girolami and Calderhead 2011), and position-dependent MALA (PMALA) (Xifara et al. 2014), and differ in implementation cost, depending on how precise is the use they make of versions of Eq. (3). The approach still leaves the specification of the metric to be used in the space of probability distributions to the user, however there are some natural choices. One can, for example, take the Hessian of π and replace its eigenvalues by their absolute values $\lambda_i \to |\lambda_i|$. Building the metric tensor from this spectrally-positive version of the Hessian of π and randomising the discretisation step size h results in an algorithm that is as robust as random walk Metropolis, in the sense that it is geometrically ergodic for targets with tail decay of $\exp\{-|\theta|^\beta\}$ for $\beta > 1$ (see Taylor 2014). A robustified version of such a metric has been introduced in Betancourt (2013) and termed SoftAbs. Here one approximates the absolute value of the eigenspectrum of the Hessian of π with a smooth strictly positive function $\lambda_i \to \lambda_i \frac{\exp\{\alpha\lambda_i\}+\exp\{-\alpha\lambda_i\}}{\exp\{\alpha\lambda_i\}-\exp\{-\alpha\lambda_i\}}$, where α is a smoothing parameter. The metric stabilises the behaviour of both MMALA, and HMC algorithms (discussed in the sequel), in the neighbourhoods where the signature of the Hessian changes.

2.3 Hamiltonian Monte Carlo

As with many improvements in the literature, starting with the very notion of MCMC, Hamiltonian (or hybrid) Monte Carlo (HMC) stems from Physics (Duane et al. 1987). After a slow emergence into the statistical community (Neal 1999), it is now central in statistical software like STAN (Stan Development Team 2014). For a complete account of this important flavour of MCMC, the reader is referred to Neal (2013), which inspired the description below; see also Betancourt et al. (2014) for a highly mathematical differential-geometric approach to HMC.

This method can be seen as a particular and efficient instance of auxiliary variables (see, e.g., Besag and Green 1993; Rubinstein 1981), in which we apply a deterministic-proposal Metropolis method to the augmented target. In physical terms, the idea behind HMC is to add a "kinetic energy" term to the "potential energy" (negative log-target), leading to the Hamiltonian

$$H(\theta,\ p) = -\log \pi(\theta) + p^{\mathrm{T}} M^{-1} p/2,$$

where θ denotes the object to be simulated (i.e., the parameter), p its speed or momentum and M the Hamiltonian matrix of π. In more statistical language, HMC creates an auxiliary variable p such that moving according to Hamilton's equations

$$\frac{\theta}{dt} = \frac{\partial H}{\partial p} = \frac{\partial H}{\partial p} = M^{-1} p,$$
$$\frac{dp}{dt} = -\frac{\partial H}{\partial \theta} = \frac{\partial \log \pi}{\partial \theta},$$

preserves the joint distribution with density $\exp\{-H(\theta,\ p)\}$, hence the marginal distribution of θ, that is, $\pi(\theta)$. Hence, if we could simulate exactly this joint distribution of $(\theta,\ p)$, a sample from $\pi(\theta)$ would be a by-product. However, in practice, the equation is solved approximately and hence requires a Metropolis correction. As discussed in, e.g., Neal (2013), the dynamics induced by Hamilton's equations is reversible and volume-preserving in the $(\theta,\ p)$ space, which means in practice that there is no need for a Jacobian in Metropolis updates. The practical implementation relies on a k-th order symplectic integrator (Hairer et al. 2006), most commonly on the second order *leapfrog approximation* that relies on a

small step level ϵ, updating p and θ via a modified Euler's method called the leapfrog that is reversible and being symplectic, preserves volume as well. This discretised update can be repeated for an arbitrary number of steps.

When considering the implementation via a Metropolis algorithm, a new value of the momentum p is drawn from the pseudo-prior $\propto \exp\{-p^{\mathrm{T}}M^{-1}p/2\}$ and it is followed by a Metropolis step, which proposal is driven by the leapfrog approximation to the Hamiltonian dynamics on $(\theta,\ p)$ and which acceptance is governed by the Metropolis acceptance probability. What makes the potential strength of this augmentation (or disintegration) scheme is that the value of $H(\theta,\ p)$ hardly changes during the Metropolis move, which means that it is most likely to be accepted and that it may produce a very different value of $\pi(\theta)$ without modifying the overall acceptance probability. In other words, moving along level sets is almost energy-free, but if the move proceeds for "long enough", the chain can reach far-away regions of the parameter space, thus avoid the myopia of standard MCMC algorithms. As explained in Neal (2013), this means that HMC avoids the inefficient random walk behaviour of most Metropolis–Hastings algorithms. What drives the exploration of the different values of $H(\theta,\ p)$ is therefore the simulation of the momentum, which makes its calibration both quite influential and delicate (Betancourt et al. 2014) as it depends on the unknown normalising constant of the target. (By calibration, we mean primarily the choice of the time discretisation step ε in the leapfrog approximation and of the number L of leapfrog leaps, but also the choice of the precision matrix M).

2.4 Optimal scaling and adaptive MCMC

The convergence of the Metropolis–Hastings algorithm 1 depends crucially on the choice of the proposal distribution q, as does the performance of both more complex MCMC and SMC algorithms, that often are hybrids using Metropolis–Hastings as simulation substeps.

Optimising over all implementable q appears to be a "disaster problem" due to its infinite-dimensional character, lack of clarity about what is implementable, what is not, and the fact that this optimal q must depend in a complex way on the target π to which we have only a limited access. In particular MALA provides a specific approach to constructing π-tailored proposals and HMC can be viewed as a combination of Gibbs and special Metropolis moves for an extended target.

In this optimisation context, it is thus reasonable to restrict ourselves to some parametric family of proposals q_ξ, or more generally of Markov transition kernels P_ξ, where $\xi \in \Xi$ is a tuning parameter, possibly high-dimensional.

The aim of adaptive Markov chain Monte Carlo is conceptually very simple. One expects that there is a set $\Xi_\pi \subset \Xi$ of good parameters ξ for which the kernel P_ξ converges quickly to π, and one allows the algorithm to search for Ξ_π "on the fly" and redesign the transition kernel during the simulation as more and more information about π becomes available. Thus an adaptive MCMC algorithm would apply the kernel $P_{\xi^{(n)}}$ to sample $\theta^{(n)}$ given $\theta^{(n-1)}$, where the tuning parameter $\xi^{(n)}$ is itself a random variable which may depend on the whole history $\theta^{(0)},\ldots,\theta^{(n-1)}$ and on $\xi^{(n-1)}$. Adaptive MCMC rests on the hope that the adaptive parameter $\xi^{(n)}$ will find Ξ_π, stay there essentially forever and inherit good convergence properties.

There are at least two fundamental difficulties in executing this strategy in practice. First, standard measures of efficiency of Markovian kernels, like the total variation convergence rate (c.f. Meyn and Tweedie 2009; Roberts and Rosenthal 2004), $L^2(\pi)$ spectral gap (Diaconis and Stroock 1991; Roberts 1996; Saloff-Coste 1997; Levin et al. 2009) or asymptotic variance (Peskun 1973; Geyer 1992; Tierney 1998) in the Markov chain central limit theorem will not be available explicitly, and their estimation from a Markov chain trajectory is often an even more challenging task than the underlying MCMC estimation problem itself.

Secondly, when executing an adaptive strategy and trying to improve the transition kernel on the fly, the Markov property of the process is violated, therefore standard theoretical tools do not apply, and establishing validity of the approach becomes significantly more difficult. While the approach has been successfully applied in some very challenging practical problems (Solonen et al. 2012; Richardson et al. 2010; Griffin et al. 2014), there are examples of seemingly reasonable adaptive algorithms that fail to converge to the intended target distribution (Bai et al. 2011; Łatuszyński et al. 2013), indicating that compared to standard MCMC even more care must be taken to ensure validity of inferential conclusions.

While heuristics-based adaptive algorithms have been considered already in Gilks et al. (1994), a remarkable result providing a tool to address the difficulty of optimising Markovian kernels coherently is the Roberts et al. (1997) paper on scaling the proposal variance. It considers settings of increasing dimensionality and investigates efficiency of the random walk Metropolis algorithm as a function of its average acceptance rate. More specifically, given a sequence of targets π_d on the product state space Θ^d with iid components constructed from conveniently smooth marginal f,

$$\pi_d(\theta) := \prod_{i=1}^{d} f(\theta_i), \quad \text{for} \quad d = 1, 2, \ldots, \tag{4}$$

consider a sequence of Markov chains θ_d, $d = 1, 2, \ldots$, where the chain $\theta_d = (\theta_d^{(n)})_{n=0,1,\ldots}$ is a random walk Metropolis targeting π_d with proposal increments distributed as $N(0,\ \sigma_d^2 I_{d\times d})$.

It then turns out that the only sensible scaling of the proposal as dimensionality increases is to take $\sigma_d^2 = l^2 d^{-1}$. In this regime the sequence of time-rescaled first coordinate processes

$$Z_d^{(t)} := \theta_{d,1}^{(\lfloor td \rfloor)}, \quad \text{for} \quad d = 1, 2, \ldots,$$

converges in a suitable sense to the solution Z of a SDE

$$dZ_t = h(l)^{1/2} dB_t + \frac{1}{2} h(l) \nabla \log f(Z_t)\, dt.$$

Hence maximising the speed of the above diffusion $h(l)$ is equivalent to maximising the efficiency of the algorithm as the dimension goes to infinity. Surprisingly, there is a one-to-one correspondence between the value $l_{opt} = \text{argmax}\, h(l)$ and the mean acceptance probability of 0.234.

The magic number 0.234 does not depend on f and gives a universal tuning recipe to be used for example in adaptive algorithms: choose the scale of the increment so that approximately 23 % of the proposals are accepted.

The result, established under restrictive assumptions, has been empirically verified to hold much more generally, for non iid targets and also in medium- and even low-dimensional examples with d as small as 5. It has been also combined with relative efficiency loss due to mismatch between the proposal and target covariance matrices (see Roberts and Rosenthal 2001).

The simplicity of the result and easy access to the average acceptance rate makes optimal scaling the main theoretical driver in development of adaptive MCMC algorithms, and adaptive MCMC is the main application and motivation for researching optimal scaling.

A large body of theoretical work extends optimal scaling formally to different and more general scenarios. For example Metropolis for smooth non iid targets has been addressed, e.g., by Bédard (2007), and in infinite dimensional settings by Beskos et al. (2009). Discrete and other discontinuous targets have been considered in Roberts (1998) and Neal et al. (2012). For MALA algorithms an optimal acceptance rate of 0.574 has been established in Roberts and Rosenthal (1998) and confirmed in infinite-dimensional settings in Pillai et al. (2012) along with the stepsize $\sigma_d^2 = l^2 d^{-1/3}$. Hybrid Monte Carlo (see Sect. 2.3) has been analysed in a similar spirit by Beskos et al. (2013) and Betancourt et al. (2014) concluding that any value $\in [0.6, 0.9]$ will be close to optimal and the leapfrog step size should be taken as $h = l \times d^{-1/4}$. These results not only inform about optimal tuning, but also provide an efficiency ordering on the algorithms in d-dimensions. Metropolis algorithms need $\mathcal{O}(d)$ steps to explore the state space, while MALA and HMC need respectively $\mathcal{O}(d^{1/3})$ and $\mathcal{O}(d^{1/4})$.

Further extensions include studying the transient phase before reaching stationarity (Christensen et al. 2005; Jourdain et al. 2012, 2014), the scaling of multiple-try MCMC (Bédard et al. 2012) and delayed rejection MCMC (Bédard et al. 2014), and the temperature scale of parallel tempering type algorithms (Atchadé et al. 2011b; Roberts and Rosenthal 2014). Interestingly, the optimal scaling of the discussed in Sect. 2.5 pseudo-marginal algorithms as obtained in Sherlock et al. (2014), and extended to more general settings in Doucet et al. (2012) and Sherlock (2014), suggests an acceptance rate of just 0.07.

While each of these numerous optimal scaling results gives rise, at least in principle, to an adaptive MCMC design, the pioneering and most successful algorithm is the Adaptive Metropolis of Haario et al. (2001). With its increasing popularity in applications, this has fuelled the development of the field.

Here one considers a normal increment proposal that estimates the target covariance matrix from past samples and applies appropriate dimension-dependent scaling and covariance shrinkage. Precisely, the proposal takes the form

$$q\left(\cdot|\theta^{(n-1)}\right) = N\left(\theta^{(n-1)}, C^{(n)}\right), \tag{5}$$

with the covariance matrix

$$C^{(n)} = \frac{(2.38)^2}{d}\left(\hat{\text{cov}}\left(\theta^{(0)}, \ldots, \theta^{(n-1)}\right) + \varepsilon I_{d\times d}\right), \tag{6}$$

which is efficiently computed using a recursive formula.

Versions and refinements of the adaptive Metropolis algorithm (Roberts and Rosenthal 2009; Andrieu and Thoms 2008) have served well in applications and motivated much of the theoretical development. These include, among many other contributions, adaptive Metropolis, delayed rejection adaptive Metropolis (Haario et al. 2006), regional adaptation and parallel chains (Craiu et al. 2009), and the robust version of Vihola (2012) estimating the shape of the distribution rather than its covariance matrix and hence suitable for heavy tailed targets.

Analogous development of adaptive MALA algorithms in Atchadé (2006) and Marshall and Roberts (2012) and of adaptive HMC and Riemannian Manifold Monte Carlo in Wang et al. (2013) building on the adaptive scaling theory, resulted in a similar drastic mixing improvement as the original Adaptive Metropolis.

Another substantial and still unexplored area where adaptive algorithms are applied for very high dimensional and multimodal problems is model and variable selection (Nott and Kohn 2005; Richardson et al. 2010; Lamnisos et al. 2013; Ji and Schmidler 2013; Griffin et al. 2014). These algorithms can incorporate reversible jump moves (Green 1995) and are guided by scaling limits for discrete distributions as

well as temperature spacing of parallel tempering to address multimodality. Successful implementations allow for fully Bayesian variable selection in models with over 20,000 variables for which otherwise only ad hoc heuristic approaches have been used in the literature.

To address the second difficulty with adaptive algorithms, several approaches have been developed to establish their theoretical underpinning. While for standard MCMC, convergence in total variation and law of large numbers are obtained almost trivially, and the effort concentrates on stronger results, like CLTs, geometric convergence, nonasymptotic analysis, and, maybe most importantly, comparison and ordering of algorithms, adaptive samplers are intrinsically difficult. The most elegant and theoretically-valid strategy is to change the underlying Markovian kernel at regeneration times only (Gilks et al. 1998). Unfortunately, this is not very appealing for practitioners since regenerations are difficult to identify in more complex settings and are essentially impractically rare in high dimensions. The original Adaptive Metropolis of Haario et al. (2001) has been validated (under some restrictive additional conditions) by controlling the dependencies introduced by the adaptation and using convergence results for mixingales. The approach has been further developed in Atchadé and Rosenthal (2005) and Atchadé (2006) to verify its ergodicity under weaker assumptions and apply the mixingale approach to adaptive MALA. Another successful approach (Andrieu and Moulines 2006 refined in Saksman and Vihola 2010) rests on martingale difference approximations and martingale limit theorems to obtain, under suitable technical assumptions, versions of LLN and CLTs. There are close links between analysing adaptive MCMC and stochastic approximation algorithms and in particular the adaptation step can be often written as a mean field of the stochastic approximation procedure; Andrieu and Robert (2001), Atchadé et al. (2011a) and Andrieu et al. (2015) contribute to this direction of analysis. Fort et al. (2011) develop an approach where both adaptive and interacting MCMC algorithms can be treated in the same framework. This allows addressing "external adaptation" algorithms such as the interacting tempering algorithm (a simplified version of the celebrated equi-energy sampler of Kou et al. 2006) or adaptive parallel tempering in Miasojedow et al. (2013).

We present here the rather general but fairly simple coupling approach (Roberts and Rosenthal 2007) to establishing convergence. Successfully applied to a variety of adaptive Metropolis samplers under weak regularity conditions (Bai et al. 2011), adaptive Gibbs and adaptive Metropolis within adaptive Gibbs samplers (Łatuszyński et al. 2013), it shows that two properties *Diminishing Adaptation* and *Containment* are sufficient to guarantee that an adaptive MCMC algorithm will converge asymptotically to the correct target distribution. To this end recall the total variation distance between two measures defined as $\|\nu(\cdot)-\mu(\cdot)\| := \sup_{A\in\mathcal{F}} |\nu(A)-\mu(A)|$, and for every Markov transition kernel P_ξ, $\xi \in \Xi$ and every starting point $\theta \in \Theta$ define the ε convergence function $M_\varepsilon : \Theta \times \Xi \to \mathbb{N}$ as

$$M_\varepsilon(\theta, \xi) := \inf\left\{n \geq 1 : \left\| P_\xi^{(n)}(\theta, \cdot) - \pi(\cdot)\right\| \leq \varepsilon\right\}.$$

Let $\{(\theta^{(n)}, \xi^{(n)})\}_{n=0}^{\infty}$ be the corresponding adaptive MCMC algorithm and by $A^{(n)}((\theta, \xi), \cdot)$ denote its marginal distribution at time t, i.e.,

$$A^{(n)}((\theta, \xi), B) := \mathbb{P}\left(\theta^{(n)} \in B | \theta^{(0)} = \theta, \xi^{(0)} = \xi\right).$$

The adaptive algorithm is ergodic for every starting values of θ and ξ if $\lim_{n\to\infty} \|A^{(n)}((\theta, \xi), \cdot) - \pi(\cdot)\| = 0$. The two conditions guaranteeing ergodicity are

Definition 1 (*Diminishing Adaptation*) The adaptive algorithm with starting values $\theta^{(0)} = \theta$ and $\xi^{(0)} = \xi$ satisfies Diminishing Adaptation, if

$$\lim_{n\to\infty} D^{(n)} = 0 \quad \text{in probability, where}$$
$$D^{(n)} := \sup_{\theta\in\Theta} \left\| P_{\xi^{(n+1)}}(\theta, \cdot) - P_{\xi^{(n)}}(\theta, \cdot)\right\|.$$

Definition 2 (*Containment*) The adaptive algorithm with starting values $\theta^{(0)} = \theta$ and $\xi^{(0)} = \xi$ satisfies Containment, if for all $\varepsilon > 0$ the sequence $\{M_\varepsilon(\theta^{(n)}, \xi^{(n)})\}_{n=0}^{\infty}$ is bounded in probability.

While diminishing adaptation is a standard requirement, Containment is subject to some discussion. On one hand, it may seem difficult to verify in practice; on the other, it may appear restrictive in the context of ergodicity results under some weaker conditions (c.f. Fort et al. 2011). However, it turns out (Łatuszyński and Rosenthal 2014) that if Containment is not satisfied, then the algorithm may still converge, but with positive probability it will be asymptotically less efficient than any nonadaptive ergodic MCMC scheme. Hence algorithms that do not satisfy Containment are termed AdapFail and are best avoided. Containment has been further studied in Bai et al. (2011) and is in particular implied by simultaneous geometric or polynomial drift conditions of the adaptive kernels.

Given that adaptive algorithms may be incorporated in essentially any sampling scheme, their introduction seems to be one of the most important innovations of the last two decades. However, despite substantial effort and many ingenious contributions, the theory of adaptive MCMC lags behind practice even more than may be the case in other computational areas. While theory always matters, the numerous unexpected and counterintuitive examples of transient adaptive algorithms suggest that in this area theory matters even more for healthy development.

For adaptive MCMC to become a routine tool, a clear-cut result is needed saying that under some easily verifiable conditions these algorithms are valid and perform not much worse than their nonadaptive counterpart with fixed parameters. Such a result is yet to be established and may require deeper understanding of how to construct stable adaptive MCMC, rather than aiming heavy technical artillery at algorithms currently in use without modifying them.

2.5 Estimated likelihoods and pseudo-marginals

There are numerous settings of interest where the target density $\pi(\cdot|y)$ is not available in closed form. For instance, in latent variable models, the likelihood function $\ell(\theta|y)$ is often only available as an intractable integral

$$\ell(\theta|y) = \int_{\mathcal{Z}} g(z,\ y|\theta)\mathrm{d}z,$$

which leads to

$$\pi(\theta|y) \propto \pi(\theta) \int_{\mathcal{Z}} g(z,\ y|\theta)\mathrm{d}z,$$

being equally intractable. A solution proposed from the early days of MCMC (Tanner and Wong 1987) is to consider z as an auxiliary variable and to simulate the joint distribution $\pi(\theta,\ z|y)$ on $\Theta \times \mathcal{Z}$ by a standard method, leading to simulating the marginal density $\pi(\cdot|y)$ as a by-product. However, when the dimension of the auxiliary variable z grows with the sample size, this technique may run into difficulties as induced MCMC algorithms are more and more likely to have convergence issues. An illustration of this case is provided by hidden Markov models, which have eventually to resort to particle filters as Markov chain algorithms become ineffective (Chopin 2007; Fearnhead and Clifford 2003). Another situation where the target density $\pi(\cdot|y)$ cannot be directly computed is the case of the "doubly intractable" likelihood (Murray et al. 2006a), when the likelihood function $\ell(\theta|y) \propto g(y|\theta)$ itself contains a term that is intractable, in that it makes the normalising constant

$$\mathfrak{Z}(\theta) = \int_{\mathcal{Y}} g(y|\theta)\mathrm{d}y,$$

impossible to compute. The resulting posterior writes

$$\pi(\theta|y) = \frac{\pi(\theta)g(y|\theta)}{\mathfrak{Z}(\theta)p(y)}, \quad \text{where}$$
$$p(y) = \int_{\Theta} \frac{\pi(\theta)g(y|\theta)}{\mathfrak{Z}(\theta)}\mathrm{d}\theta,$$

and consequently the Metropolis–Hastings acceptance rate becomes

$$\alpha(\theta,\ \theta') = \min\left\{1,\ \frac{\pi(\theta')g(y,\ \theta')q(\theta'|\theta)}{\pi(\theta)g(y,\ \theta)q(\theta|\theta')} \times \frac{\mathfrak{Z}(\theta)}{\mathfrak{Z}(\theta')}\right\},$$

and cannot be evaluated exactly for algorithmic purposes.

Examples of this kind abound in Markov random fields models, as for instance for the Ising model (Murray et al. 2006b; Møller et al. 2006).

Both the approaches of Murray et al. (2006a) and Møller et al. (2006) require sampling data from the likelihood $\ell(\theta|y)$, which limits their applicability. The latter uses in addition an importance sampling function and may suffer from poor acceptance rates.

Andrieu and Roberts (2009) propose a more general resolution of such problems by designing a Metropolis–Hastings algorithm that replaces the intractable target density $\pi(\cdot|y)$ with an unbiased estimator, following an idea of Beaumont (2003). The approach is termed pseudo-marginal. Rather than evaluating the posterior exactly, a positive unbiased estimate S_θ of $\pi(\theta|y)$ is utilised. More formally, a new Markov chain is constructed that evolves on the extended state space $\Theta \times \mathbb{R}_+$, where at iteration n, given the pair $(\theta^{(n-1)},\ S^{(n-1)}_{\theta^{(n-1)}})$ of parameter value and density estimate at this value, the proposal $(\theta',\ S_{\theta'})$ is obtained by sampling $\theta' \sim q(\cdot|\theta^{(n-1)})$ and obtaining $S_{\theta'}$, the estimate of $\pi(\theta'|y)$. Analogously to the standard Metropolis–Hastings step, the pair $(\theta',\ S_{\theta'})$ is accepted as $(\theta^{(n)},\ S^{(n)}_{\theta^{(n)}})$ with probability

$$\min\left\{1,\ \frac{S_{\theta'}q(\theta^{(n-1)}|\theta')}{S^{(n-1)}_{\theta^{(n-1)}}q(\theta'|\theta^{(n-1)})}\right\},$$

otherwise the proposal is rejected and the new value set as $(\theta^{(n)},\ S^{(n)}_{\theta^{(n)}}) := (\theta^{(n-1)},\ S^{(n-1)}_{\theta^{(n-1)}})$.

It is not difficult to verify that the bivariate chain on extended state space $\Theta \times \mathbb{R}_+$ enjoys the correct $\pi(\theta|y)$ marginal on Θ and the approach is valid, see Andrieu and Roberts (2009) for details (and also Andrieu and Vihola 2015 for an abstracted account).

One specific instance of constructing unbiased estimators of the posterior is presented in Girolami et al. (2013) and based on random truncations of infinite series expansions. The paper also offers an excellent overview of inference methods for intractable likelihoods.

The performance of the pseudo-marginal approach will depend on the quality of the estimators S_θ and hence stabilising them as well as understanding this relationship is an active area of current development. Often S_θ is constructed as an importance sampler based on an importance sample z. Thus in particular, the improvements from using multiple samples of z to estimate π are of interest and can be assessed from Andrieu and Vihola (2015) where the efficiency of the algorithm is studied in terms of its spectral gap and CLT asymptotic variance. Sherlock et al. (2014), Doucet et al. (2012) and Sherlock (2014), on the other hand, investigate

the efficiency as a function of the acceptance rate and variance of the noise, deriving the optimal scaling, as discussed in Sect. 2.4.

As an alternative to the above procedure of using estimates of the intractable likelihood to design a new Markov chain on an extended state space with correct marginal, one could naively use these estimates to approximate the Metropolis–Hastings accept–reject ratio and let the Markov chain evolve in the original state space. This would amount to dropping the current realisation of S_θ and obtaining a new one in each accept–reject attempt. Such a procedure is termed Monte Carlo within Metropolis (Andrieu and Roberts 2009). Unfortunately this approach does not preserve the stationary distribution, and the resulting Markov chain may even not be ergodic (Medina-Aguayo et al. 2015). If ergodic, the difference between stationary distribution, resulting from the noisy acceptance must be quantified, which is a highly nontrivial task and the bounds will rarely be tight (see also Alquier et al. 2014; Pillai and Smith 2014; Rudolf and Schweizer 2015 for related methodology and theory). The approach is however an interesting avenue since at the price of being biased, it overcomes mixing difficulties of the exact pseudo-marginal version.

Design and understanding of pseudo-marginal algorithms is a direction of dynamic methodological development that in the coming years will be further fuelled not only by complex models with intractable likelihoods, but also by the need of MCMC algorithms for Big Data. In this context the likelihood function cannot be evaluated for the whole dataset even in the *iid* case just because computing the long product of individual likelihoods is infeasible. Several Big Data MCMC approaches have been already considered in Welling and Teh (2011), Korattikara et al. (2013), Teh et al. (2014), Bardenet et al. (2014), Maclaurin and Adams (2014), Minsker et al. (2014), Quiroz et al. (2014) and Strathmann et al. (2015).

2.6 Particle MCMC

While we refrain from covering particle filters here, since others (Beskos et al. 2015) in this volume are focussing on this technique, a recent advance at the interface between MCMC, pseudo-marginals, and particle filtering is the notion of particle MCMC (or *pMCMC*), developed by Andrieu et al. (2011). This innovation is indeed rather similar to the pseudo-marginal algorithm approach, taking advantage of the state-space models and auxiliary variables used by particle filters. It differs from standard particle filters in that it targets (mostly) the marginal posterior distribution of the parameters.

The simplest setting in which pMCMC applies is one of a state-space model where a latent sequence $x_{0:T}$ is a Markov chain with joint density

$$p_0(x_0|\theta)\, p_1(x_1|x_0,\theta)\cdots p_T(x_T|x_{T-1},\theta),$$

and is associated with an observed sequence $y_{1:T}$ such that

$$y_{1:T}|x_{1:T},\theta \sim \prod_{i=1}^{T} q_i(y_i|x_i,\theta).$$

The iterations of pMCMC are MCMC-like in that, at iteration t, a new value θ' of θ is proposed from an arbitrary transition kernel $\mathfrak{h}(\cdot|\theta^{(t)})$ and then a new value of the latent series $x'_{0:T}$ is generated from a particle filter approximation of $p(x_{0:T}|\theta', y_{1:T})$. Since the particle filter returns as a by-product (Del Moral et al. 2006) an unbiased estimator of the marginal posterior of $y_{1:T}$, $\hat{q}(y_{1:T}|\theta')$, this estimator can be used as such in the Metropolis–Hastings ratio

$$\frac{\hat{q}(y_{1:T}|\theta')\pi(\theta')\mathfrak{h}(\theta^{(t)}|\theta')}{\hat{q}(y_{1:T}|\theta)\pi(\theta^{(t)})\mathfrak{h}(\theta'|\theta^{(t)})} \wedge 1.$$

Its validity follows from the general argument of Andrieu and Roberts (2009), although some additional (notational) effort is needed to demonstrate all random variables used therein are correctly assessed (see Andrieu et al. 2011; Wilkinson 2011a, the latter providing a very progressive introduction to the notions of pMCMC and particle Gibbs, which helped greatly in composing this section). Note however that the general validation of pMCMC as targetting the joint posterior of the states and parameters and of the parallel particle Gibbs sampler does not follow from pseudo-marginal arguments.

This approach is being used increasingly in complex dynamic models like those found in signal processing (Whiteley et al. 2010), dynamical systems like the PDEs in biochemical kinetics (Wilkinson 2011b) and probabilistic graphical models (Lindsten et al. 2014). An extension to approximating the sequential filtering distribution is found in Chopin et al. (2013).

2.7 Parallel MCMC

Since MCMC relies on local updating based on the current value of a Markov chain, opportunities for exploiting parallel resources, either CPU or GPU, would seem quite limited, In fact, the possibilities reach far beyond the basic notion of running independent or coupled MCMC chains on several processors. For instance, Craiu and Meng (2005) construct parallel antithetic coupling to create negatively correlated MCMC chains (see also Frigessi et al. 2000), while Craiu et al. (2009) use parallel exploration of the sample space to tune an adaptive MCMC algorithm. Jacob et al. (2011) exploit GPU facilities to improve by Rao-Blackwellisation the Monte Carlo approximations produced by a Markov chain, even though the parallelisation does not improve the convergence of the chain. See also Lee et al. (2009) and

Suchard et al. (2010) for more detailed contributions on the appeal of using GPUs towards massive parallelisation, and Wilkinson (2005) for a general survey on the topic.

Another recently-explored direction is "prefetching". Based on Brockwell (2006) this approach computes the $2^2, 2^3, \ldots, 2^k$ values of the posterior that will be needed $2, 3, \ldots, k$ sweeps ahead by simulating the possible "futures" of the Markov chain, according to whether the next k proposals are accepted or not, in parallel. Running a regular Metropolis–Hastings algorithm then means building a decision tree back to the current iteration and drawing $2, 3, \ldots, k$ uniform variates to go down the tree to the appropriate branch. As noted by Brockwell (2006), "in the case where one can guess whether or not acceptance probabilities will be 'high' or 'low', the tree could be made deeper down 'high' probability paths and shallower in the 'low' probability paths." This idea is exploited in Angelino et al. (2014), by creating "speculative moves" that consider the reject branch of the prefetching tree more often than not, based on some preliminary or dynamic evaluation of the acceptance rate. Using a fast but close-enough approximation to the true target (and a fixed sequence of uniforms) may also produce a "single most likely path" on which prefetched simulations can be run. The basic idea is thus to run simulations and costly likelihood computations on many parallel processors along a prefetched path, a path that has been prefetched for its high approximate likelihood. There are obviously instances where this speculative simulation is not helpful because the actual chain with the genuine target ends up following another path. Angelino et al. (2014) actually go further by constructing sequences of approximations for the precomputations. The proposition for the sequence found therein is to subsample the original data and use a normal approximation to the difference of the log (sub-)likelihoods. See Strid (2010) for related ideas.

A different use of parallel capabilities is found in Calderhead (2014). At each iteration of Calderhead's algorithm, N replicas are generated, rather than 1 in traditional Metropolis–Hastings. The Markov chain actually consists of N components, from which one component is selected at random as a seed for the next proposal. This approach can be seen as a special type of data augmentation (Tanner and Wong 1987), where the index of the selected component is an auxiliary variable. The neat trick in the proposal (and the reason for its efficiency gain) is that the stationary distribution of the auxiliary variable can be determined and hence used N times in updating the vector of N components. An interesting feature of this approach is when the original Metropolis–Hastings algorithm is expressed as a finite state space Markov chain on the set of indices $\{1, \ldots, N\}$. Conditional on the values of the N dimensional vector, the stationary distribution of that subchain is no longer uniform. Hence, picking N indices from the stationary helps in selecting the most appropriate images, which explains why the rejection rate decreases. The paper indeed evaluates the impact of increasing the number of proposals in terms of ESS, acceptance rate, and mean squared jump distance. Since this proposal is an almost free bonus resulting from using N processors, it sounds worth investigating and comparing with more complex parallel schemes.

Neiswanger et al. (2013) introduced the notion of embarrassingly parallel MCMC, where "embarrassing" refers to the "embarrassingly simple" solution proposed therein, namely to solve the difficulty in handling very large datasets by running completely independent parallel MCMC samplers on parallel threads or computers and using the outcomes of those samplers as density estimates, pulled together as a product towards an approximation of the true posterior density. In other words, the idea is to break the posterior as

$$p(\theta|y) \propto \prod_{i=1}^{m} p_i(\theta|y), \tag{7}$$

and to use the estimate

$$\hat{p}(\theta|y) \propto \prod_{i=1}^{m} \hat{p}_i(\theta|y),$$

where the individual estimates are obtained, say, nonparametrically. The method is then "asymptotically exact" in the weak (and unsurprising) sense of converging in the number of MCMC iterations. Still, there is a theoretical justification that is not found in previous parallel methods that mixed all resulting samples without accounting for the subsampling. And the point is made that, in many cases, running MCMC samplers with subsamples produces faster convergence. The decomposition of $p(\cdot)$ into its components is done by partitioning the iid data into M subsets and taking a power $1/m$ of the prior in each case. (This may induce issues about impropriety.) However, the subdivision is arbitrary and can thus be implemented in cases other than the fairly restrictive iid setting. Because each (subsample) nonparametric estimate involves T terms, the resulting overall estimate contains Tm terms and the authors suggest using an independent Metropolis sampler to handle this complexity. This is in fact necessary for producing a final sample from the (approximate) true posterior distribution.

In a closely related way, Wang and Dunson (2013) start from the same product representation of the target (posterior), namely, (7). However, they criticise the choice made by Neiswanger et al. (2013) to use MCMC approximations for each component of the product for the following reasons:

(1) Curse of dimensionality in the number of parameters d;
(2) Curse of dimensionality in the number of subsets m;
(3) Tail degeneration;
(4) Support inconsistency and mode misspecification.

Point 1 is relevant, but there may be ways other than kernel estimation to mix samples from the terms in the product. Point 2 is less of a clearcut drawback: while the Tm terms corresponding to a product of m sums of T terms sounds self-defeating, Neiswanger et al. (2013) use a clever device to avoid the combinatorial explosion, namely operating on one component at a time. Having non-manageable targets is not such an issue in the post-MCMC era. Point 3 is formally correct, in that the kernel tail behaviour induces the kernel estimate tail behaviour, most likely disconnected from the true target tail behaviour, but this feature is true for any non-parametric estimate, even for the Weierstrass transform defined below, and hence maybe not so relevant in practice. In fact, by lifting the tails up, the simulation from the subposteriors should help in visiting the tails of the true target. Finally, point 4 does not seem to be life-threatening. Assuming that the true target can be computed up to a normalising constant, the value of the target for every simulated parameter could be computed, eliminating those outside the support of the product and highlighting modal regions.

The Weierstrass transform of a density f is a convolution of f and of an arbitrary kernel K. Wang and Dunson (2013) propose to simulate from the product of the Weierstrass transform, using a multi-tiered Gibbs sampler. Hence, the parameter is only simulated once and from a controlled kernel, while the random effects from the convolution are related with each subposterior. While the method requires coordination between the parallel threads, the components of the target are separately computed on a single thread. The clearest perspective on the Weierstrass transform may possibly be the rejection sampling version where simulations from the subpriors are merged together into a normal proposal on θ, to be accepted with a probability depending on the subprior simulations.

VanDerwerken and Schmidler (2013) keep with the spirit of parallel MCMC papers like consensus Bayes (Scott et al. 2013), embarrassingly parallel MCMC (Neiswanger et al. 2013) and Weierstrass MCMC (Wang and Dunson 2013), namely that the computation of the likelihood can be broken into batches and MCMC run over those batches independently. The idea of the authors is to replace an exploration of the whole space operated via a single Markov chain (or by parallel chains acting independently which all have to "converge") with parallel and independent explorations of parts of the space by separate Markov chains. The motivation is that "Small is beautiful": it takes a shorter while to explore each set of the partition, hence to converge, and, more importantly, each chain can work in parallel with the others. More specifically, given a partition of the space, into sets A_i with posterior weights w_i, parallel chains are associated with targets equal to the original target restricted to those A_is. This is therefore an MCMC version of partitioned sampling. With regard to the shortcomings listed in the quote above, the authors consider that there does not need to be a bijection between the partition sets and the chains, in that a chain can move across partitions and thus contribute to several integral evaluations simultaneously. It is somewhat unclear (a) whether or not this impacts ergodicity (it all depends on the way the chain is constructed, i.e., against which target) as it could lead to an over-representation of some boundary regions and (b) whether or not it improves the overall convergence properties of the chain(s). A more delicate issue with the partitioned MCMC approach stands with the partitioning. Indeed, in a complex and high-dimension model, the construction of the appropriate partition is a challenge in itself as we often have no prior idea where the modal areas are. Waiting for a correct exploration of the modes is indeed faster than waiting for crossing between modes, provided all modes are represented and the chain for each partition set A_i has enough energy to explore this set. It actually sounds unlikely that a target with huge gaps between modes will see a considerable improvement from the partitioned version when the partition sets A_i are selected on the go, because some of the boundaries between the partition sets may be hard to reach with an off-the-shelf proposal. A last comment about this innovative paper is that the adaptive construction of the partition has much in common with Wang–Landau schemes (Wang and Landau 2001; Lee et al. 2005; Atchadé and Liu 2010; Jacob and Ryder 2014).

3 ABC and others, exactly delivering an approximation

Motivated by highly complex models where MCMC algorithms and other Monte Carlo methods were too inefficient by far, approximate methods have emerged where the output cannot be considered as simulations from the genuine posterior, even under idealised situations of infinite computing power. These methods include ABC techniques, described in more details below, but also variational Bayes (Jaakkola and Jordan 2000), empirical likelihood (Owen 2001), integrated nested Laplace approximation (INLA) (Rue et al. 2009) and other solutions that rely on pseudo-models, or on summarised versions of the data, or both. It is quite important to signal this evolution as we think that it may be a central feature of computational Bayesian statistics in the coming years. From a statistical perspective, it also induces a somewhat paradoxical situation where loss of information is balanced by improvement in precision, for a given computational budget. This perspective is not only interesting at the computational level but forces us (as statisticians) to re-evaluate in depth the nature of a statistical model and could produce a paradigm shift in the near future by giving a brand new meaning to George Box's motto that "all models are wrong".

3.1 ABC per se

It seems important to discuss ABC in this partial tour of Bayesian computational techniques as (a) they provide the only approach to their model for some Bayesians, (b) they deliver samples in the parameter space that are exact simulations from a posterior of some kind (Wilkinson 2013), $\pi^{\mathrm{ABC}}(\theta|\mathbf{y}_0)$ if not the original posterior $\pi(\theta|\mathbf{y}_0)$, where $\mathbf{y}_0$ denotes the data in this section (c) they may be more intuitive to some researchers outside statistics, as they entail simulating from the inferred model, i.e., going forward from parameter to data, rather than backward, from data to parameter, as in traditional Bayesian inference, (d) they can be merged with MCMC algorithms, and (e) they allow drawing inference directly from summaries of the data rather than the data itself.

ABC techniques play a role in the 2000s that MCMC methods did in the 1990s, in that they handle new models for which earlier (e.g., MCMC) algorithms were at a loss, in the same way the latter (MCMC) were able to handle models that regular Monte Carlo approaches could not reach, such as latent variable models (Tanner and Wong 1987; Diebolt and Robert 1994; Richardson and Green 1997). New models for which ABC unlocked the gate include Markov random fields, Kingman's coalescent for phylogeographical data, likelihood models with an intractable normalising constant, and models defined by their quantile function or their characteristic function. While the ABC approach first appeared a "quick-and-dirty" solution, to be considered only until more elaborate representations could be found, those algorithms have been progressively incorporated into the statistician's toolbox as a novel form of generic nonparametric inference handling partly-defined statistical models. They are therefore attractive as much for this reason as for being handy computational solutions when everything else fails.

A statistically intriguing feature of those methods is that they customarily require—for greater efficiency—replacing the data with (much) smaller-dimension summaries[1] or summary statistics, because of the complexity of the former. In almost every case calling for ABC, those summaries are not sufficient statistics and the method thus implies from the start a loss of statistical information, at least at a formal level, since relying on the raw data is out of the question and therefore the additional information it provides is moot. This imposed reduction of the statistical information raises many relevant questions, from the choice of summary statistics (Blum et al. 2013) to the consistency of the ensuing inference (Robert et al. 2011).

[1] Maybe due to their initial introduction in population genetics, the oxymoron 'summary statistics' is now prevalent in descriptions of ABC algorithms, included in the statistical literature, where the (linguistically sufficient) term 'statistic' would suffice.

Although it has now diffused into a wide range of applications, the technique of ABC was first introduced by and for population genetics (Tavaré et al. 1997; Pritchard et al. 1999) to handle ancestry models driven by Kingman's coalescent and with strictly intractable likelihoods (Beaumont 2010). The likelihood function of such genetic models is indeed "intractable" in the sense that, while derived from a fully defined and parameterised probability model, this function cannot be computed (at all or within a manageable time) for a single value of the parameter and for the given data. Bypassing the original example to avoid getting mired into the details of population genetics, examples of intractable likelihoods include densities with intractable normalising constants, i.e., $f(\mathbf{y}|\theta) = g(\mathbf{y}|\theta)/Z(\theta)$ such as in Potts (1952) and auto-exponential (Besag 1972) models, and pseudo-likelihood models (Cucala et al. 2009).

Example 1 A very simple illustration of an intractable likelihood is provided by Bayesian inference based on the median and median absolute deviation statistics of a sample from an arbitrary location-scale family, $y_1, \ldots, y_n \overset{\text{iid}}{\sim} \sigma^{-1} g(\sigma^{-1}\{y-\mu\})$, as the joint distribution of this statistic is not available in closed form.

The concept at the core of ABC methods can be seen as both very naïve and intrinsically related to the foundations of Bayesian statistics as *inverse probability* (Rubin 1984). This concept is that data $\mathbf{y}$ simulated conditional on values of the parameter close to the "true" value of the parameter should look more similar to the actual data $\mathbf{y}_0$ than data $\mathbf{y}$ simulated conditional on values of the parameter far from the "true" value. ABC actually involves an acceptance/rejection step in that parameters simulated from the prior are accepted only when

$$\rho(\mathbf{y}, \mathbf{y}_0) < \epsilon,$$

where $\rho(\cdot, \cdot)$ is a distance and $\epsilon > 0$ is called the tolerance. It can be shown that the algorithm exactly samples the posterior when $\epsilon = 0$, but this is very rarely achievable in practice (Grelaud et al. 2009). An algorithmic representation is as follows:

Algorithm 2 ABC (basic version)

for $t = 1$ to N **do**
  repeat
    Generate θ^* from the prior $\pi(\cdot)$
    Generate $\mathbf{y}^*$ from the model $f(\cdot|\theta^*)$
    Compute the distance $\rho(\mathbf{y}^0, \mathbf{y}^*)$
    Accept θ^* if $\rho(\mathbf{y}^0, \mathbf{y}^*) < \epsilon$
  until acceptance
end for
return N accepted values of θ^*

Calibration of the ABC method in Algorithm 2 involves selecting the distance $\rho(\cdot, \cdot)$ and deducing the tolerance from computational cost constraints. However, in realistic settings, ABC is never implemented as such because comparing raw data to simulated raw data is rarely efficient, noise dominating signal (see, e.g., Marin et al. 2012 for toy examples). It is therefore natural that one first considers dimension-reduction techniques to bypass this curse of dimensionality. For instance, if rudimentary estimates $S(y)$ of the parameter θ are available, they are good candidates. In the ABC literature, they are called *summary statistics*, a term that does not impose any constraint on their form and hence leaves open the question of performance, as discussed in Marin et al. (2012) and Blum et al. (2013). A more practical version of the ABC algorithm is shown in Algorithm 3 below, with a different output for each choice of the summary statistic. We stress in this version of the algorithm the construction of the tolerance ϵ as a quantile of the simulated distances $\rho(S(\mathbf{y}^0),\ S(\mathbf{y}^{(t)}))$, rather than an additional parameter of the method.

Algorithm 3 ABC (version with summary)

for $t = 1$ to N_{ref} **do**
  Generate $\theta^{(t)}$ from the prior $\pi(\cdot)$
  Generate $\mathbf{y}^{(t)}$ from the model $f(\cdot|\theta^{(t)})$
  Compute $d_t = \rho(S(\mathbf{y}^0),\ S(\mathbf{y}^{(t)}))$
end for
Order distances $d_{(1)} \leq d_{(2)} \leq \ldots \leq d_{(N_{ref})}$
return the values $\theta^{(t)}$ associated with the k smallest distances

An immediate question about this approximate algorithm is how much it remains connected with the original posterior distribution and in case it does not, where does it draw its legitimacy. A first remark in this connection is that it constitutes at best a convergent approximation to the posterior distribution $\pi(\theta|S(y_0))$. It can easily be seen that ABC generates outcomes from a genuine posterior distribution when the data is randomised with scale ϵ (Wilkinson 2013; Fearnhead and Prangle 2012). This interpretation indicates a decrease in the precision of the inference but it does not provide a universal validation of the method. A second perspective on the ABC output is that it is based on a non-parametric approximation of the sampling distribution (Blum 2010; Blum and François 2010), connected with both indirect inference (Drovandi et al. 2011) and k-nearest neighbour estimation (Biau et al. 2014). While a purely Bayesian non-parametric analysis of this aspect has not yet emerged, this brings an additional if cautious support for the method.

Example 2 Continuing from the previous example of a location-scale sample only monitored through the pair median plus mad statistic, we consider the special case of a normal sample $y_1, \ldots, y_n \sim \mathcal{N}(\mu, \sigma^2)$, with $n = 100$. Using a conjugate prior $\mu \sim \mathcal{N}(0, 10)$, $\sigma^{-2} \sim \mathcal{G}a(2, 5)$, we generated 10^6 parameter values, along with the corresponding pairs of summary statistics. When creating the distance $\rho(\cdot, \cdot)$, we used both following versions:

$$\rho_1\left(S\left(\mathbf{y}^0\right), S(\mathbf{y})\right) = |\mathrm{med}(\mathbf{y}^0) - \mathrm{med}(\mathbf{y}|/\mathrm{mad}(\mathrm{med}(\mathbf{Y}))) + |\mathrm{mad}(\mathbf{y}^0) - \mathrm{mad}(\mathbf{y}|/\mathrm{mad}(\mathrm{mad}(\mathbf{Y}))),$$

$$\rho_2\left(S\left(\mathbf{y}^0\right), S(\mathbf{y})\right) = |\mathrm{med}(\mathbf{y}^0) - \mathrm{med}(\mathbf{y}|/\mathrm{mad}(\mathrm{med}(\mathbf{Y}))) + |\log\mathrm{mad}(\mathbf{y}^0) - \log\mathrm{mad}(\mathbf{y}|/\mathrm{mad}(\log\mathrm{mad}(\mathbf{Y}))),$$

where the denominators are computed from the reference table in order to scale the components properly. Figure 1 shows the impact of the choice of this distance, but even more clearly the discrepancy between inference based on the ABC and the true inference on (μ, σ^2).

The discrepancy can however be completely eliminated by post-processing: Fig. 2 reproduces Fig. 1 by comparing the histograms of an ABC sample with the version corrected by Beaumont et al.'s (2002) local regression, as the latter is essentially equivalent to a regular Gibbs output.

Barber et al. (2015) studies the rate of convergence for ABC algorithms through the mean square error when approximating a posterior moment. They show the convergence rate is of order $\mathrm{O}(n^{2/q+4})$, when q is the dimension of the ABC summary statistic, associated with an optimal tolerance in $\mathrm{O}(n^{-1/4})$. Those rates are connected with the nonparametric nature of ABC, as already suggested in the earlier literature: for instance, Blum (2010), who links ABC with standard kernel density non-parametric estimation and find a tolerance (re-expressed as a bandwidth) of order $n^{-1/q+4}$ and an rmse of order $^2/_{q+4}$ as well, while Fearnhead and Prangle (2012) obtain similar rates, with a tolerance of order $n^{-1/q+2}$ for noisy ABC. See also Calvet and Czellar (2014). Similarly, Biau et al. (2014) obtain precise convergence rates for ABC interpreted as a k-nearest-neighbour estimator.

Lee and Łatuszyński (2014) have also produced precise characterisations of the geometric ergodicity or lack thereof of four ABC–MCMC algorithms:

(1) the standard ABC–MCMC (with N replicates of the simulated pseudo-data to each simulated parameter value),
(2) versions involving simulations of the replicates repeated at the subsequent step,
(3) use of a stopping rule in the generation of the pseudo data, and
(4) a "gold-standard algorithm" based on the (unavailable) measure of an ϵ ball around the data.

Based a result by Roberts and Tweedie (1996b), also used in Mengersen and Tweedie (1996), namely that an MCMC chain cannot be geometrically ergodic when there exist almost-absorbing states, they derive that (under some tech-

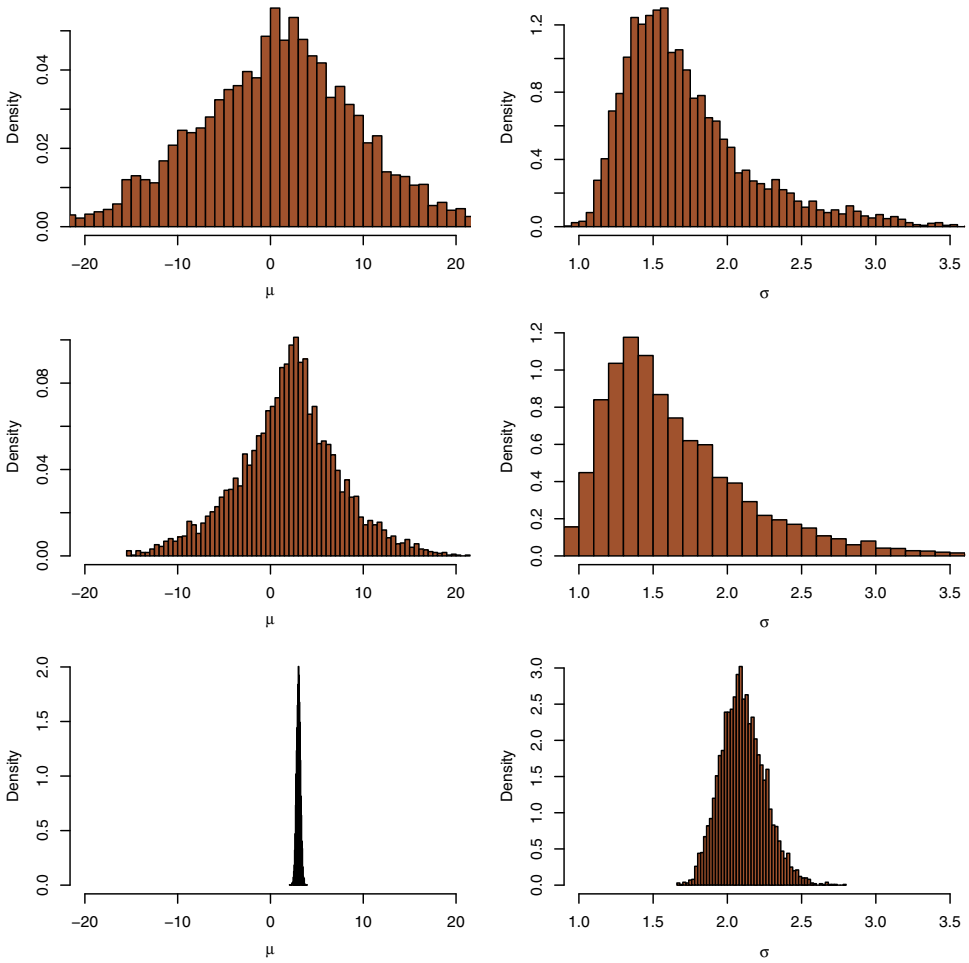

Fig. 1 Comparison of the posterior distributions on μ (*left*) and σ (*right*) when using an ABC Algorithm 3 with distance ρ_1 (*top*) and ρ_2 (*central*), and when using a standard Gibbs sampler (*bottom*). All three samples are based on the same number of subsampled parameters. The dataset is a $\mathcal{N}(3,\ 2^2)$ sample and the tolerance value ϵ corresponds to $\alpha = 0.5\,\%$ of the reference table

nical assumptions) the first two versions above cannot be variance-bounding (i.e., that the spectral gap is zero), while the last two versions can be both variance-bounding and geometrically ergodic under some appropriate conditions on the prior and the above ball measure. This result is thus rather striking in simulating a *random* number of auxiliary variables is sufficient to produce geometric ergodicity. We note that this result does not contradict the parallel result of Bornn et al. (2014), who establish that there is no efficiency gain in simulating $N > 1$ replicates of the pseudo-data, since there is no randomness involved in that approach. However, the latter result only applies to functions with finite variances.

When testing hypotheses and selecting models, the Bayesian approach relies on modelling hypotheses and model indices as part of the parameter and hence ABC naturally operates as this level as well, as demonstrated in Algorithm 4 following Cornuet et al. (2008), Grelaud et al. (2009) and Toni et al. (2009). In fields like population genetics, model choice and hypotheses validation is presumably the primary motivation for using ABC methods as exemplified in Belle et al. (2008), Cornuet et al. (2010), Excoffier et al. (2009), Ghirotto et al. (2010), Guillemaud et al. (2009), Leuenberger and Wegmann (2010), Patin et al. (2009), Ramakrishnan and Hadly (2009), Verdu et al. (2009) and Wegmann and Excoffier (2010). It is also the area that attracts most of the criticisms addressed against ABC: while some are easily dismissed (see, e.g., Templeton 2008, 2010; Beaumont et al. 2010; Berger et al. 2010), the impact of the choice of the summary statistics on the value of the posterior probability remains a delicate issue that prompted Pudlo et al. (2014) to advocate the alternative use of a posterior predictive error.

Algorithm 4 ABC (model choice)

for $i = 1$ to N **do**
  Generate $\mathfrak{M}$ from the prior $\pi(\mathcal{M} = m)$
  Generate $\theta_{\mathfrak{M}}$ from the prior $\pi_{\mathfrak{M}}(\theta_{\mathfrak{M}})$
  Generate $\mathbf{y}$ from the model $f_{\mathfrak{M}}(\mathbf{y}|\theta_{\mathfrak{M}})$
  Compute the distance $\rho\{S(\mathbf{y}),\ S(\mathbf{y}_0)\}$
  Set $\mathfrak{M}^{(i)} = \mathfrak{M}$ and $\theta^{(i)} = \theta_{\mathfrak{M}}$
end for
return the values $\mathfrak{M}^{(i)}$ associated with the k smallest distances

Indeed, Robert et al. (2011) pointed out the potential irrelevance of ABC-based posterior probabilities, due to the

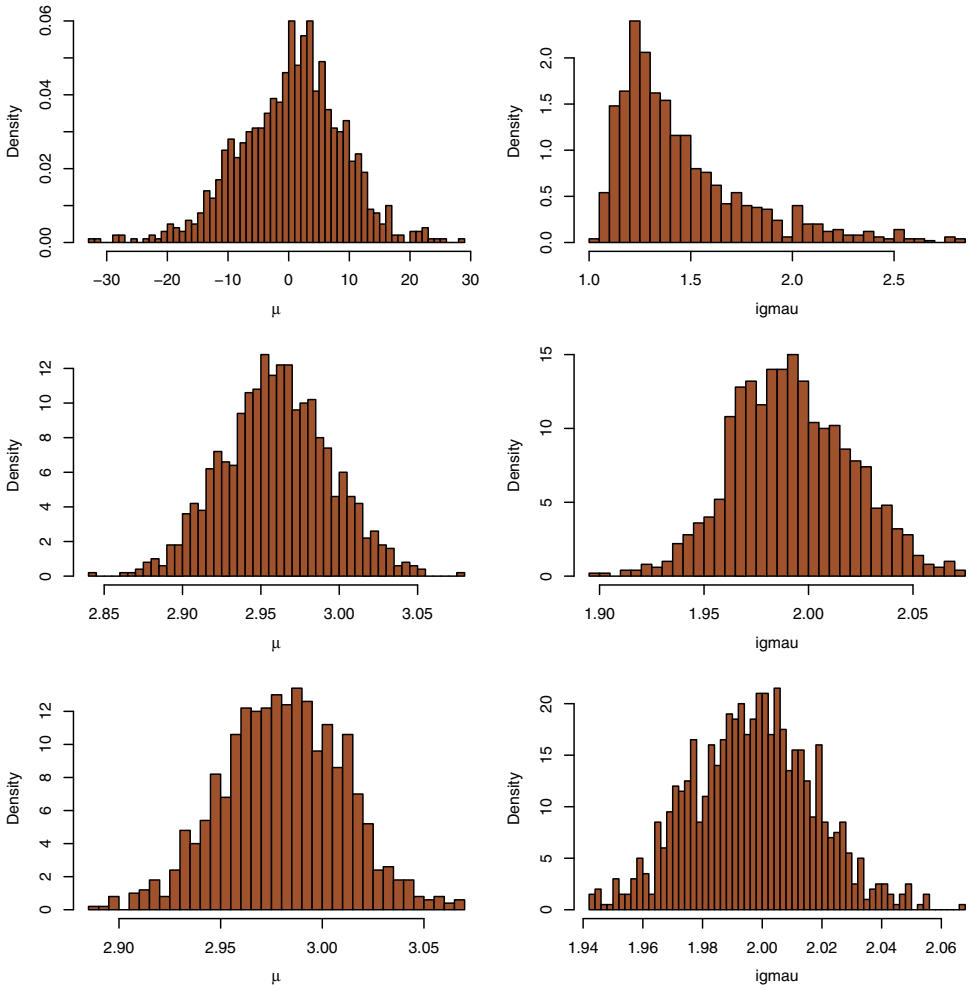

Fig. 2 Comparison of the posterior distributions on μ (*left*) and σ (*right*) when using an ABC algorithm 3 with distance ρ_1 (*top*), a post-processed version by Beaumont et al.'s (2002) local regression (*central*), and when using a standard Gibbs sampler (*bottom*). The simulation setting is the same as in Fig. 1

possible ancilarity (for model choice) of summary statistics, as also explained in Didelot et al. (2011). Marin et al. (2014) consider for instance the comparison of normal and Laplace fits on both normal and Laplace samples and show that using sample mean and sample variance as summary statistics produces Bayes factors converging to values near 1, instead of the consistent 0 and $+\infty$.

Marin et al. (2014) analyses this phenomenon with the aim of producing a necessary and sufficient consistency condition on summary statistics. Quite naturally, the summaries that are acceptable must display different behaviour under both models, in the guise of ranges of means $\mathbb{E}_\theta[S(\mathbf{y}_0)]$ that do not intersect for the two models. (In the counter-example of the normal–Laplace test, the expectations of the sample mean and variance can be recovered under both models.) This characterisation then leads to a practical asymptotic test validating summary statistics and to the realisation that a larger number of summaries helps in achieving this goal (while degrading the estimated tolerance). More importantly, it shows that the reduction of information represented by an ABC approach may prevent discriminating between models, at least when trying to recover the Bayes factor. In the end, this is a natural consequence of simplifying the description of both the data and the model, and can be found in most limited information settings.

3.2 More fish in the alphabet soup

Besides ABC, approximation techniques have spread wide and far towards analysing more complex or less completely defined models. Rather than a confusion, this multiplicity of available approximations is beneficial both to the understanding of the underlying model and to the calibration of those different methods.

Variational Bayes methods have been proposed for at least two decades to substitute exponential families $q(\theta|\lambda)$ for complex posterior distributions $\pi(\theta)$ (Jordan et al. 1999; MacKay 2002). The central notion in those methods is that the exponential family structure and a so-called mean-field representation of the approximation

$$q(\theta|\lambda) = \prod_{i=1}^{k} q_i\,(\theta_i|\lambda_i)\,,$$

allows for a sometimes closed-form minimisation of the Kullback–Leibler distance $\mathrm{KL}(q(\theta|\lambda), \pi(\theta))$ between the true target and its approximation. If not, the setting is quite congenial to the use of EM algorithms (Paisley et al. 2012). See Salimans and Knowles (2013) for a contemporary view on this approach, which offers considerable gains in terms of computing time, while being difficult to assess in terms of discrepancy with the "truth", i.e., the outcome that would result from using the genuine posterior.

Another approach that has met with considerable interest in the past 5 years is INLA (Rue et al. 2009). The method operates on latent Gaussian random fields, with likelihoods of the form

$$\prod_{i=1}^{n} f\left(x_i|\eta_i, \theta\right),$$

where the x_i's are the observables and the η_i's are latent variables. Using Laplace approximations to the marginal distributions $\pi(\theta|\mathbf{x}_0)$ and to $f(\eta|\mathbf{x}_0)$, INLA produces fast and accurate approximations of the true posterior distribution as well as of the marginal likelihood value. Thanks to the availability of a well-constructed package called R-INLA, this approach has gathered a large group of followers.

A somewhat exotic example of variational approximation is expectation-propagation (EP) (Minka 2001), which starts from an arbitrary decomposition of the target distribution

$$\pi(\theta) = \prod_{j=1}^{k} \pi_j(\theta),$$

(often inspired by a likelihood decomposition into groups of observations) and iteratively approximate each term π_j in the product by a density member of an exponential family, $\nu(\cdot|\lambda)$m using the other approximations as a marginal. Given the current approximation of $\pi(\theta)$ at iteration t,

$$\nu\left(\theta|\lambda_t\right) = \prod_{j=1}^{k} \nu_j\left(\theta|\lambda_t\right),$$

where λ_t is the current value of the hyperparameter, the tth step in the EP algorithm goes as follows:

(1) Select $1 \leq j \leq k$ at random
(2) Define the marginal

$$\nu_{-j}\left(\theta|\lambda_t\right) \propto \frac{\nu(\theta|\lambda_t)}{\nu_j(\theta|\lambda_t)};$$

(3) Update the hyperparameter λ_t by solving

$$\lambda_{t+1} = \underset{\lambda}{\operatorname{argmin}}\, \mathrm{KL}\left\{\pi_j(\theta)\nu_{-j}\left(\theta|\lambda_t\right),\ \nu(\theta|\lambda)\right\}$$

(4) Update $\nu_j(\theta|\lambda_t)$ as

$$\nu_j\left(\theta|\lambda_{t+1}\right) \propto \frac{\nu(\theta|\lambda_{t+1})}{\nu_{-j}(\theta|\lambda_t)}.$$

(In the above, KL denotes the Kullback–Leibler divergence.) The algorithm stops at stationarity. The convergence of this approach is not yet fully understood, but Barthelmé and Chopin (2014) consider EP as a practical substitute for ABC, avoiding the selection of summary statistics by using a local constraint

$$\left\|x_i - x^{\mathrm{obs}}\right\| \leq \epsilon,$$

on each element of the simulated pseudo-data vector, x^{obs} being the actual data. In addition, EP provides an approximation of the evidence. In the ABC setting, when using a Normal distribution as the exponential family default, implementing EP means computing empirical mean and empirical variance, one observation at a time, under the above tolerance constraint. Obviously, using a Normal candidate means that the final approximation will also look much like a Normal distribution, which both links with other Normal approximations like INLA and variational methods, and signals a difficulty with EP in less smooth cases, such as ridge-like or multimodal posteriors.

While different approximations keep being developed and tested, with arguments ranging from efficient programming, to avoiding simulations, to having an ability to deal with more complex structures, their drawback is the overall incapacity to assess the amount of approximation involved. Bootstrap evaluations can be attempted in the simplest cases but cannot be extended to more realistic situations.

4 Optimisation in modern Bayesian computation

Optimisation methodology for high-dimensional maximum-a-posteriori (MAP) estimation is another area of Bayesian computation that has received a lot of attention over the last years, particularly for problems related to machine learning, signal processing and computer vision. One reason for this is that for many Bayesian models optimisation is significantly more computationally tractable than integration. This has generated a lot of interest in MAP estimators, especially for applications involving very high-dimensional parameter spaces or tight computing time constraints, for which calculating other summaries of the posterior distribution is not feasible. Here we review some of the major breakthroughs in this topic, which originated mainly outside the statistics community. We focus on developments related to high-dimensional convex optimisation, though many of

the techniques discussed below are also useful for non-convex optimisation. In particular, in Sect. 4.1 we concentrate on *proximal optimisation algorithms*, a powerful class of iterative methods that exploit tools from convex analysis, monotone operator theory and theory of non-expansive mappings to construct carefully designed fixed-point schemes. We refer the reader to the excellent book by Bauschke and Combettes (2011) for the mathematics underpinning proximal optimisation algorithms, and to the recent tutorial papers by Combettes and Pesquet (2011), Cevher et al. (2014) and Parikh and Boyd (2014) for an overview of the field and applications to signal processing and machine learning.

However, we do think it is vital to insist that, at the same time as asserting that modern optimisation methodology represents a much-underused opportunity in Bayesian inference, in its raw form it inevitably fails to deliver essential elements of the Bayesian paradigm. The vision is not to deliver a point estimate of an unknown structure, but the full richness of Bayesian inference in its coherence, its proper treatment of uncertainty, its intrinsic treatment of model uncertainty, and so on. Bayesian statistics does not boil down to optimisation with penalisation (Lange et al. 2014). We need to express the uncertainty associated with decisions and estimation, stemming from the stochastic nature of the data, and our lack of knowledge about relevant mechanisms.

The challenge is to use the awesome capacity of fast optimisation in a high-dimensional parameter space to focus on local regions of that space where a combination of analytic and numerical investigation can deliver at least approximations to full posterior distributions and derived quantities. The community has barely risen to this challenge, with only isolated examples such as the discussion in Green (2015) of a problem in unlabelled shape analysis. However, the growing community of INLA (Rue et al. 2009) users may bring an heightened awareness of such possibilities, along with its efficient code (Schrödle and Held 2011; Muff et al. 2013). Another promising research area is to use mathematical and algorithmic tools from convex optimisation to design more efficient high-dimensional MCMC algorithms (Pereyra 2015).

4.1 Proximal algorithms

Similarly to many other computational methodologies that are widely used nowadays, proximal algorithms were first proposed several decades ago by Moreau (1962), Martinet (1970) and Rockafellar (1976), and regained attention recently in the context of large-scale inverse problems and "big data".

We consider the computation of maximisers of posterior densities $\pi(\boldsymbol{\theta}) = \exp\{-g(\boldsymbol{\theta})\}/\kappa$ that are high-dimensional and log-concave, which we formulate as

$$\hat{\boldsymbol{\theta}}_{MAP} = \underset{\boldsymbol{\theta}\in\mathbb{R}^n}{\operatorname{argmin}}\, g(\boldsymbol{\theta}), \tag{8}$$

where g belongs to the class $\Gamma_0(\mathbb{R}^n)$ of lower semicontinuous convex functions from $\mathbb{R}^n \to (-\infty, +\infty]$. Notice that g may be non-differentiable and take value $g(\boldsymbol{\theta}) = +\infty$, reflecting constraints in the parameter space. In order to introduce proximal algorithms we first recall the following standard definitions and results from convex analysis: we say that $\boldsymbol{\varphi} \in \mathbb{R}^n$ is a subgradient of g at $\boldsymbol{\theta} \in \mathbb{R}^n$ if it satisfies $(\mathbf{u}-\boldsymbol{\theta})^T\boldsymbol{\varphi} + g(\boldsymbol{\theta}) \leq g(\mathbf{u})$, $\forall \mathbf{u} \in \mathbb{R}^n$. The set of all such subgradients defines the subdifferential set $\partial g(\boldsymbol{\theta})$, and $\hat{\boldsymbol{\theta}}_{MAP}$ is a minimiser of g if and only if $\mathbf{0} \in \partial g(\hat{\boldsymbol{\theta}}_{MAP})$. The (convex) conjugate of $g \in \Gamma_0(\mathbb{R}^n)$ is the function $g^* \in \Gamma_0(\mathbb{R}^n)$ defined as $g^*(\boldsymbol{\varphi}) = \sup_{\mathbf{u}\in\mathbb{R}^n} \mathbf{u}^T\boldsymbol{\varphi} - g(\mathbf{u})$. The subgradients of g and g^* satisfy the property $\boldsymbol{\varphi} \in \partial g(\boldsymbol{\theta}) \Leftrightarrow \boldsymbol{\theta} \in \partial g^*(\boldsymbol{\varphi})$.

Proximal algorithms take their name from the proximity mapping, defined for $g \in \Gamma_0(\mathbb{R}^n)$ and $\lambda > 0$ as (Moreau 1962)

$$\operatorname{prox}_g^\lambda(\boldsymbol{\theta}) = \underset{\mathbf{u}\in\mathbb{R}^n}{\operatorname{argmin}}\, g(\mathbf{u}) + \|\mathbf{u}-\boldsymbol{\theta}\|^2/2\lambda. \tag{9}$$

In order to gain intuition about this mapping it is useful to analyse its behaviour when $\lambda \in \mathbb{R}^+$ is either very small or very large. In the limit $\lambda \to \infty$, the quadratic penalty term vanishes and (9) maps all points to $\hat{\boldsymbol{\theta}}_{MAP}$. In the opposite limit $\lambda \to 0$, (9) becomes the identity operator and maps $\boldsymbol{\theta}$ to itself. For finite values of λ, $\operatorname{prox}_g^\lambda(\boldsymbol{\theta})$ behaves similarly to a gradient mapping and moves points in the direction of $\hat{\boldsymbol{\theta}}_{MAP}$. Like gradients, proximity mappings have several properties that are useful for devising fixed-point methods (Bauschke and Combettes 2011).

Property 1 *The proximity mapping of g is related to its subdifferential by the inclusion $\{\boldsymbol{\theta} - \operatorname{prox}_g^\lambda(\boldsymbol{\theta})\}/\lambda \in \partial g\{\operatorname{prox}_g^\lambda(\boldsymbol{\theta})\}$, which collapses to $\nabla g\{\operatorname{prox}_g^\lambda(\boldsymbol{\theta})\}$ when $g \in \mathcal{C}^1$. As a result, for any $\lambda > 0$, the minimiser of g verifies the fixed-point equation $\boldsymbol{\theta} = \operatorname{prox}_g^\lambda(\boldsymbol{\theta})$.*

Property 2 *Proximity mappings are firmly non-expansive; that is, $\|\operatorname{prox}_g^\lambda(\boldsymbol{\theta}) - \operatorname{prox}_g^\lambda(\mathbf{u})\|^2 \leq (\boldsymbol{\theta}-\mathbf{u})^T\{\operatorname{prox}_g^\lambda(\boldsymbol{\theta}) - \operatorname{prox}_g^\lambda(\mathbf{u})\}$, $\forall \boldsymbol{\theta}, \mathbf{u} \in \mathbb{R}^n$.*

Property 3 *The proximity mappings of g and its conjugate g^* are related by Moreau's decomposition formula: $\boldsymbol{\theta} = \operatorname{prox}_g^\lambda(\boldsymbol{\theta}) + \lambda \operatorname{prox}_{g^*}^{1/\lambda}(\boldsymbol{\theta}/\lambda)$.*

The simplest proximal method to solve (8) is the *proximal point algorithm* given by the iteration

$$\boldsymbol{\theta}^{k+1} = \operatorname{prox}_g^\lambda\left(\boldsymbol{\theta}^k\right). \tag{10}$$

Every sequence $\{\boldsymbol{\theta}^k\}_{k\in\mathbb{N}}$ produced by this algorithm converges to $\hat{\boldsymbol{\theta}}_{MAP}$, even if proximity mappings are evaluated

inexactly, as long as the errors are of certain types (e.g., summable). A more general proximal point algorithm includes relaxation, i.e.,

$$\boldsymbol{\theta}^{k+1} = (1-\alpha_k)\,\boldsymbol{\theta}^k + \alpha_k \operatorname{prox}_g^{\lambda}\left(\boldsymbol{\theta}^k\right), \quad \alpha_k \in (0,\,2),$$

and with over-relaxation (i.e., $\alpha_k \in (1,\,2)$) often converges faster than (10). Notice from Property 1 that (10) can be interpreted as an implicit (backward) subgradient steepest descent to minimise g, i.e., $\boldsymbol{\theta}^{k+1} = \boldsymbol{\theta}^k - \lambda\boldsymbol{\varphi}$, with $\boldsymbol{\varphi} \in \partial g(\boldsymbol{\theta}^{k+1})$. Alternatively, proximal point algorithms can also be interpreted as explicit (forward) gradient steepest descent to minimise the *Moreau envelope* of g, $e_\lambda(\boldsymbol{\theta}) = \inf_{\mathbf{u}\in\mathbb{R}^n} g(\mathbf{u}) + \|\mathbf{u}-\boldsymbol{\theta}\|^2/2\lambda$, a convex lower bound on g that by construction is continuously differentiable and has the same minimiser as g.

Proximal point algorithms may appear of little relevance because evaluating $\operatorname{prox}_g^{\lambda}$ can be as difficult as solving (8) in the first place (notice that (9) is a convex minimisation problem similar to (8)). Surprisingly, many advanced proximal optimisation methods can in fact be shown to be either applications of this simple algorithm, or closely related to it.

Most proximal methods operate by splitting g, e.g.,

$$\hat{\boldsymbol{\theta}}_{MAP} = \operatorname*{argmin}_{\boldsymbol{\theta}\in\mathbb{R}^n} \{g_1(\boldsymbol{\theta}) + g_2(\boldsymbol{\theta})\}, \tag{11}$$

such that $g_1 \in \Gamma_0(\mathbb{R}^n)$ and $g_2 \in \Gamma_0(\mathbb{R}^n)$ have gradients or proximity mappings that are easy to compute or approximate. For example, for many Bayesian models it is possible to find a decomposition $g(\boldsymbol{\theta}) = g_1(\boldsymbol{\theta}) + g_2(\boldsymbol{\theta})$ such that g_1 is β-Lipschitz[2] differentiable and $g_2 \in \Gamma_0(\mathbb{R}^n)$, possibly non-differentiable, has a proximity mapping that can be computed efficiently with a specialised algorithm. This decomposition is useful for instance in linear inverse problems, where g_1 is often related to a Gaussian observation model involving linear operators and g_2 to a log-prior promoting a parsimonious representation (e.g., sparsity on some appropriate dictionary, low-rankness) or enforcing convex constraints (e.g., positivity, positive definiteness). For models that admit this decomposition, it is possible to compute $\hat{\boldsymbol{\theta}}_{MAP}$ efficiently with a *forward–backward* algorithm, also known as the proximal gradient algorithm

$$\boldsymbol{\theta}^{k+1} = \operatorname{prox}_{g_2}^{\lambda_n}\left(\boldsymbol{\theta}^k - \lambda_n \nabla g_1\left(\boldsymbol{\theta}^k\right)\right). \tag{12}$$

For $\lambda_n = \lambda \in (0,\,1/\beta)$ the objective function $g(\boldsymbol{\theta}^k)$ converges to $g(\hat{\boldsymbol{\theta}}_{MAP})$ with rate $O(1/k)$. If the value of the Lipschitz constant β is unknown λ_n can be found by line-search.

[2] $g_1 \in \mathcal{C}^1$ has β-Lipschitz continuous gradient if $\|\nabla g_1(\boldsymbol{\theta}) - \nabla g_1(\mathbf{u})\| \leq \beta\|\boldsymbol{\theta}-\mathbf{u}\|$, $\forall(\boldsymbol{\theta},\,\mathbf{u}) \in \mathbb{R}^N \times \mathbb{R}^N$

A remarkable property of (12) is that it can be accelerated to converge with rate $O(1/k^2)$, which is optimal for this class of problems (Nesterov 2004). This can be achieved for instance by introducing an extrapolation step

$$\begin{aligned}\boldsymbol{\theta}^{+} &= \boldsymbol{\theta}^k + \omega_k\left(\boldsymbol{\theta}^k - \boldsymbol{\theta}^{k-1}\right),\\ \boldsymbol{\theta}^{k+1} &= \operatorname{prox}_{g_2}^{\beta^{-1}}\left(\boldsymbol{\theta}^{+} - \beta^{-1}\nabla g_1\left(\boldsymbol{\theta}^{+}\right)\right),\end{aligned} \tag{13}$$

where $\{\omega_k\}_{k\in\mathbb{N}}$ is an appropriate sequence of extrapolation parameters. It was noticed by Combettes and Pesquet (2011) that several important convex optimisation algorithms can be derived as applications of the forward–backward algorithm, for example the projected gradient algorithm for minimising a Lipschitz differentiable function subject to a convex constraint (in this case the proximity mapping reduces to a projection onto the convex set). Notice that (12) can be interpreted as an implementation of the proximal point iteration (10) where $\operatorname{prox}_g^{\lambda}(\boldsymbol{\theta}^k)$ is approximated by replacing g_1 with its first order Taylor series approximation around the point $\boldsymbol{\theta}^k$.

Moreover, in some cases it may be more efficient to compute $\hat{\boldsymbol{\theta}}_{MAP}$ by solving the dual of (11), for instance if g admits a decomposition $g(\boldsymbol{\theta}) = g_1(\boldsymbol{\theta}) + g_2(L\boldsymbol{\theta})$ for some linear operator $L \in \mathbb{R}^{n\times p}$, $g_1 \in \Gamma_0(\mathbb{R}^n)$ strongly convex and $g_2 \in \Gamma_0(\mathbb{R}^p)$ with efficient proximity mapping. In this case, the Fenchel–Rockafellar theorem states that $\hat{\boldsymbol{\theta}}_{MAP}$ can be computed by solving the dual problem (Bauschke and Combettes 2011, Chap. 19)

$$\boldsymbol{\psi}^* = \operatorname*{argmin}_{\boldsymbol{\psi}\in\mathbb{R}^p} g_1^*\left(-L^T\boldsymbol{\psi}\right) + g_2^*(\boldsymbol{\psi}), \tag{14}$$

and setting $\hat{\boldsymbol{\theta}}_{MAP} = \nabla g_1^*(-L^T\boldsymbol{\psi}^*)$. This p-dimensional problem can be solved iteratively with a forward–backward algorithm $\boldsymbol{\psi}^{k+1} = \operatorname{prox}_{g_2^*}^{\lambda_n}(\boldsymbol{\psi}^k - \lambda_n \nabla g_1^*(-L^T\boldsymbol{\psi}^k))$ that can also be accelerated to converge with rate $O(1/k^2)$, and where we note that the proximity mapping of g_2^* is typically evaluated by using Property 3, and that the strong convexity of g_1 implies Lipschitz differentiability of g_1^*. Computing $\hat{\boldsymbol{\theta}}_{MAP}$ via (14) can lead to important computational savings, in particular if $p \ll n$ or if g_2 is separable and has a proximity mapping that can be computed in parallel for each element of $\boldsymbol{\theta}$ (this is generally not possible for $g_2 \circ L$). We refer the reader to Komodakis and Pesquet (2014) for an overview of recent dual and primal–dual algorithms and guidelines for parallel implementations.

Another important proximal optimisation method is the Douglas–Rachford splitting algorithm given by

$$\begin{aligned}\boldsymbol{\theta}^{k+\frac{1}{2}} &= \operatorname{prox}_{g_1}^{\lambda}\left(\boldsymbol{\theta}^k\right),\\ \boldsymbol{\theta}^{k+1} &= \boldsymbol{\theta}^k - \boldsymbol{\theta}^{k+\frac{1}{2}} + \operatorname{prox}_{g_2}^{\lambda}\left(2\boldsymbol{\theta}^{k+\frac{1}{2}} - \boldsymbol{\theta}^k\right).\end{aligned} \tag{15}$$

From a theoretical viewpoint this algorithm is more general than the forward–backward algorithm because it does not require g_1 or g_2 to be continuously differentiable. However, its practical application is limited to problems for which both g_1 and g_2 have efficient proximity mappings. Similarly to the forward–backward algorithm, (15) includes many proximal algorithms that been proposed in the literature for specific models, and can also be interpreted as an application of the proximal point algorithm.

The proximal method that is arguably most widely used in Bayesian inference is the *alternating direction method of multipliers* (ADMM), which operates by formulating (11) as a constrained optimisation problem

$$\begin{aligned} &\underset{\boldsymbol{\theta}\in\mathbb{R}^n,\, \mathbf{z}\in\mathbb{R}^n}{\operatorname{argmin}} \quad g_1(\boldsymbol{\theta}) + g_2(\mathbf{z}) \\ &\text{subject to} \quad \boldsymbol{\theta} = \mathbf{z}, \end{aligned} \tag{16}$$

and then using augmented Lagrangian techniques to express (16) as an unconstrained saddle point problem with saddle function $g_1(\boldsymbol{\theta}) + g_2(\mathbf{z}) + \lambda\boldsymbol{\varphi}^T(\boldsymbol{\theta} - \mathbf{z}) + \|\boldsymbol{\theta} - \mathbf{z}\|^2/2\lambda$ (Boyd et al. 2011). ADMM solves this problem with the iteration

$$\begin{aligned} \boldsymbol{\theta}^{k+1} &= \operatorname{prox}^{\lambda}_{g_1}\left(\mathbf{z}^k - \boldsymbol{\varphi}^k\right), \\ \mathbf{z}^{k+1} &= \operatorname{prox}^{\lambda}_{g_2}\left(\boldsymbol{\theta}^{k+1} + \boldsymbol{\varphi}^k\right), \\ \boldsymbol{\varphi}^{k+1} &= \boldsymbol{\varphi}^k + \boldsymbol{\theta}^{k+1} - \mathbf{z}^{k+1}, \end{aligned} \tag{17}$$

that also involves the proximity mappings of g_1 and g_2. This basic ADMM iteration can be tailored to specific models in many ways (e.g., to exploit decompositions of the form $g_1 = \tilde{g}_1 \circ L_1$ and $g_2 = \tilde{g}_2 \circ L_2$ so that proximal updates can be performed in parallel for all components of $\boldsymbol{\theta}$, $\mathbf{z}$ and $\boldsymbol{\varphi}$). Interestingly, ADMM can be interpreted as an application of the Douglas–Rachford algorithm to the dual of (16), and is therefore also a special case of the proximal point algorithm. For more details about the ADMM algorithm, see the recent tutorial by Boyd et al. (2011).

Furthermore, an important characteristic of proximal optimisation algorithms is that they can be massively parallelised to take advantage of parallel computer architectures. Suppose for instance that g admits the decomposition $g(\boldsymbol{\theta}) = \sum_{m=1}^{M} g_m(L_m\boldsymbol{\theta})$ with $g_m \in \Gamma(\mathbb{R}^{p_m})$ and $L_m \in \mathbb{R}^{n\times p_m}$ such that the mappings of g_m are easy to compute and $Q = \sum_{m=1}^{M} L_m^T L_m$ is invertible. Then, in a manner akin to (16), we express (8) as

$$\begin{aligned} &\underset{\mathbf{z}_1\in\mathbb{R}^n,\ldots,\mathbf{z}_M\in\mathbb{R}^n}{\operatorname{argmin}} \sum_{m=1}^{M} g_m\left(\mathbf{z}_m\right) \\ &\text{subject to} \quad \mathbf{z}_m = L_m\boldsymbol{\theta}, \quad \forall m = 1,\ldots,M, \end{aligned} \tag{18}$$

and compute $\hat{\boldsymbol{\theta}}_{MAP}$ with the following iteration

$$\begin{aligned} \boldsymbol{\theta}^{k+1} &= Q^{-1}\sum_{m=1}^{M} L_m^T\left(\mathbf{z}_m^k - \boldsymbol{\varphi}_m^k\right), \\ \mathbf{z}_m^{k+1} &= \operatorname{prox}^{\lambda}_{g_m}\left(L_m\boldsymbol{\theta}^{k+1} - \boldsymbol{\varphi}_m^k\right), \quad \forall m = 1,\ldots,M, \\ \boldsymbol{\varphi}_m^{k+1} &= \boldsymbol{\varphi}_m^k + L_m\boldsymbol{\theta}^{k+1} - \mathbf{z}_m^{k+1}, \quad \forall m = 1,\ldots,M, \end{aligned} \tag{19}$$

that can be parallelised with factor M at a coarse level (e.g., on a multi-processor system). Further parallelisation may be possible at a finer scale (e.g., on a vectorial processor such as GPU or FPGA) by taking advantage of the structure of $\operatorname{prox}^{\lambda}_{g_m}$ or by using specialised algorithms. This algorithm, known as the *simultaneous direction method of multipliers*, is also closely related to the ADMM, Douglas–Rachford and proximal point algorithms. Notice that splitting g not only allows the exploitation of parallel computer architectures, but may also significantly simplify the computation of proximity mappings; often $\operatorname{prox}^{\lambda}_{g_m}$ has a closed-form expression. Lastly, it is worth mentioning that there are other modern proximal optimisation algorithms that can be massively parallelised, for example the *generalised forward backward* algorithm (Raguet et al. 2013), the *parallel proximal* algorithms (Combettes and Pesquet 2008; Pesquet and Pustelnik 2012), and the parallel primal–dual algorithm (Combettes and Pesquet 2012).

Finally, main current topics of research in proximal optimisation include theory and methodology for: (1) randomised and stochastic algorithms that operate with estimators of gradients and proximity mappings to reduce computational complexity and allow for errors in the update rules, (2) adaptive and variable metric algorithms (e.g., Riemannian and Newton-type) that exploit the model's geometry to improve convergence speed, and (3) proximal methods for non-convex problems. We anticipate that in the future new and stronger connections will emerge between proximal optimisation and stochastic simulation, in particular through developments in stochastic optimisation and high-dimensional MCMC sampling. For example, one connection is through the integration of modern stochastic convex optimisation and Markovian stochastic approximation (Combettes and Pesquet 2014; Andrieu et al. 2015), and of proximal optimisation and high-dimensional MCMC sampling (Pereyra 2015).

4.2 Convex relaxations

Modern proximal optimisation was greatly motivated by important theoretical results on the recovery of partially-observed sparse vectors and low-rank matrices through convex minimisation (Candès et al. 2006; Candès and Tao

2009) and on *compressive sensing* (Candès and Wakin 2008). A key idea underlying these works is that of approximating a combinatorial optimisation problem, whose solution is NP-hard, with a "relaxed" convex problem that is computationally tractable, and whose solution is in some sense close to the solution of the original problem. Reciprocally, the development of modern convex optimisation has in turn generated much interest in log-concave models, convex regularisers, and "convexifications" (i.e., convex relaxations for intractable or poorly tractable models) for statistical inference problems involving high-dimensionality, large datasets and computing time constraints (Chandrasekaran et al. 2012; Chandrasekaran and Jordan 2013).

4.3 Illustrative example

For illustration, we show an application of proximal optimisation to Bayesian image resolution enhancement. The goal is to recover a high-resolution image $\boldsymbol{\theta} \in \mathbb{R}^n$ from a blurred and noisy observed image $\mathbf{y} \sim \mathcal{N}(H\boldsymbol{\theta},\ \sigma^2 \boldsymbol{I}_n)$, where $H \in \mathbb{R}^{n\times n}$ is a linear operator representing the blur point spread function of the low resolution acquisition system and σ^2 is the system's noise power. This inverse problem is ill-posed, a difficulty that Bayesian image processing methods address by exploiting prior knowledge about $\boldsymbol{\theta}$. Here we use the following hierarchical Bayesian model (Oliveira et al. 2009)

$$
\begin{aligned}
&f(\mathbf{y}|\boldsymbol{\theta}) = \left(2\pi\sigma^2\right)^{-n/2} \exp\left\{-\|\mathbf{y} - H\boldsymbol{\theta}\|_2^2/2\sigma^2\right\},\\
&\pi(\boldsymbol{\theta}|\alpha) \propto \alpha^{-n} \exp\left(-\alpha\, \|\nabla_d \boldsymbol{\theta}\|_{1-2}\right),\\
&\pi(\alpha) = e^{-\alpha} \mathbf{1}_{\mathbb{R}^+}(\alpha),
\end{aligned}
\tag{20}
$$

where $\pi(\boldsymbol{\theta}|\alpha)$ is the (improper) *total-variation* Markov random field, $\|\cdot\|_{1-2}$ denotes the composite $\ell_1 - \ell_2$ norm and ∇_d is the discrete gradient operator that computes the vertical and horizontal differences between neighbour image pixels. This prior is log-concave and models the fact that differences between neighbouring image pixels are usually very small but occasionally take large values; it is arguably the most widely used prior in modern statistical image processing. The values of H and σ^2 are typically determined during the system's calibration process and are here assumed known.

We compute the MAP estimator of $\boldsymbol{\theta}$ associated with the marginal posterior $\pi(\boldsymbol{\theta}|\mathbf{y}) = \int_0^\infty \pi(\boldsymbol{\theta},\ \alpha|\mathbf{y})\mathrm{d}\alpha$, which is unimodal but not log-concave,

$$
\begin{aligned}
\hat{\boldsymbol{\theta}}_{MAP} = &\operatorname*{argmin}_{\boldsymbol{\theta}\in\mathbb{R}^n} \|\mathbf{y} - H\boldsymbol{\theta}\|_2^2/2\sigma^2\\
&+ (n+1)\log\left(\|\nabla_d\boldsymbol{\theta}\|_{1-2} + 1\right).
\end{aligned}
\tag{21}
$$

Problem (21) is not convex, but can nevertheless be solved efficiently with proximal algorithms by using a *majorisation–minimisation* strategy. To be precise, starting from some initial condition $\boldsymbol{\theta}^{(0)}$, e.g., $\boldsymbol{\theta}^{(0)} = \mathbf{y}$, we iteratively minimise the following sequence of strictly convex majorants (Oliveira et al. 2009)

$$
\begin{aligned}
&\boldsymbol{\theta}^{(t+1)} = \operatorname*{argmin}_{\boldsymbol{\theta}\in\mathbb{R}^n} \|\mathbf{y} - H\boldsymbol{\theta}\|_2^2/2\sigma^2 + \alpha_{\text{eff}}^{(t)}\, \|\nabla_d\boldsymbol{\theta}\|_{1-2}\,,\\
&\text{with}\quad \alpha_{\text{eff}}^{(t)} = (n+1)\left(\left\|\nabla_d\boldsymbol{\theta}^{(t)}\right\|_{1-2} + 1\right).
\end{aligned}
\tag{22}
$$

Iteration (22) involves a convex subproblem that can easily be solved using most modern proximal optimisation techniques. For example, here we use the state-of-the-art ADMM algorithm *SALSA* (Afonso et al. 2011) implemented with $g_1(\boldsymbol{\theta}) = \|\mathbf{y} - H\boldsymbol{\theta}\|_2^2/2\sigma^2$, $g_2(\mathbf{u}) = \alpha_{eff}^{(t)}\|\nabla_d\mathbf{u}\|_{1-2}$, and the constraint $\boldsymbol{\theta} = \mathbf{u}$ (though we could have also used other modern algorithms Pesquet and Pustelnik 2012; Combettes and Pesquet 2012; Raguet et al. 2013). To compute the proximity mapping of g_1 we use the fact that H is block-circulant to compute matrix products and pseudo-inverses with the FFT algorithm. We compute the proximity mapping of g_2 with a highly parallelised implementation of the specialised algorithm of Chambolle (2004).

Fig. 3 Observed blurred noisy image $\mathbf{y}$

Figure 3 presents a blurred and noisy observation $\mathbf{y}$ of the popular "boats" image[3] of size 512 × 512 pixels, generated with a uniform 9 × 9 blur and a noise power of $\sigma^2 = 0.5^2$ (blurred-signal-to-noise ratio $BRSN = 10\log_{10}\{\|H\boldsymbol{\theta}_0\|_2^2/\sigma^2\} = 40\,\text{dB}$). Figure 4 below shows the MAP estimate $\hat{\boldsymbol{\theta}}_{MAP}$ obtained by solving (21) using four iterations of (22) and a total of 51 ADMM iterations. We observe that this resolution enhancement process has produced a remarkably sharp image with very noticeable fine detail. Moreover, Fig. 5 shows the magnitude of the

[3] The boat image is available for download from the SIPI image database at http://sipi.usc.edu/database/database.php?volume=misc\&image=38#top

Fig. 4 Resolution enhanced image $\hat{\boldsymbol{\theta}}_{MAP}$ obtained by solving (21) with the majorisation–minimisation strategy (22)

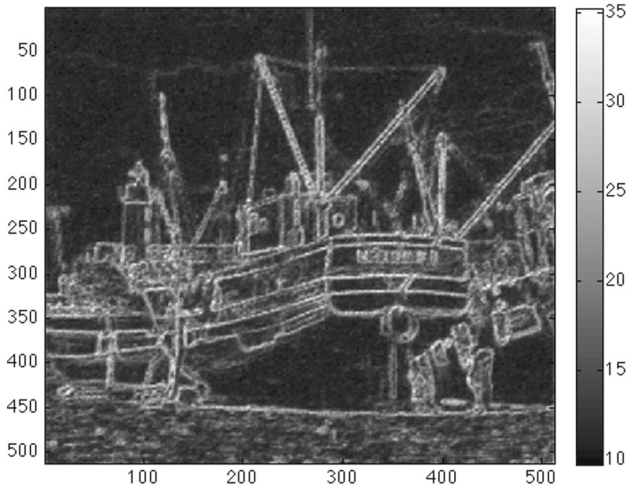

Fig. 5 Widths of pixel-wise 90 % marginal credibility intervals estimated with the proximal MCMC algorithm of Pereyra (2015)

marginal 90 % credibility regions for each pixel, as measured by the distance between the 5 and 95 % quantile estimates. These estimates were computed using the proximal Metropolis-adjusted Langevin algorithm (Pereyra 2015), which is appropriate for high-dimensional densities that are not continuously differentiable. We observe in Fig. 5 that the uncertainty is mainly concentrated at the contours and object boundaries, revealing that model is able to accurately detect the presence of sharp edges in the image but with some uncertainty about their exact location. Finally, Fig. 6 shows the convergence of the estimates $\boldsymbol{\theta}^{(t,k)}$ produced by each ADMM iteration to $\hat{\boldsymbol{\theta}}_{MAP}$ (as measured by the mean squared error $\|\boldsymbol{\theta}^{(t,k)} - \hat{\boldsymbol{\theta}}_{MAP}\|_2^2$). Notice that computing $\hat{\boldsymbol{\theta}}_{MAP}$ only required 10 s (experiment conducted on an Apple Macbook Pro computer running MATLAB 2013, a C++ implementation would certainly produce even faster results). This is remarkably fast given the high dimensionality of the problem

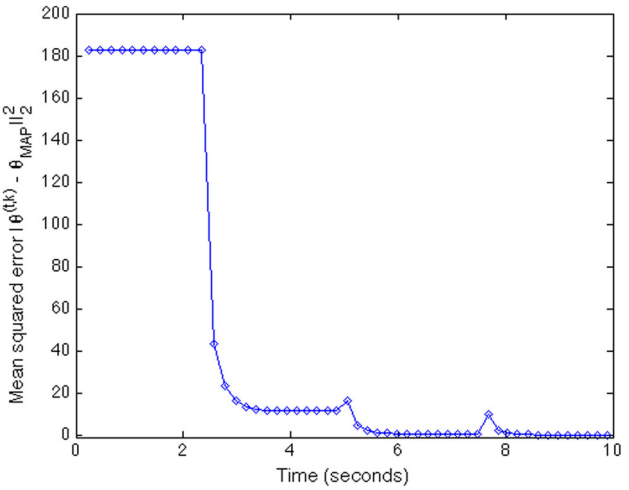

Fig. 6 Convergence of the estimate $\boldsymbol{\theta}$–$\hat{\boldsymbol{\theta}}_{MAP}$ versus computing time (s)

($n = 262, 144$). The computation of the credibility regions by MCMC sampling (20,000 samples with a thinning factor of 1000 to reduce the algorithm's memory foot-print) required 75 h.

5 Discussion

5.1 Bayesian computation in the era of data science

Is there a revolution taking place right now and have we missed the train, standing on the platform, only concerned with small-print on the train schedules—apart, that is, from the obvious but not-so-new requirement to handle massive datasets (and the mistakes that come with them)?!

As with other areas of statistical science, the Bayesian computation community has to decide whether data science is an opportunity or a threat. Inevitably if we do not treat it as an opportunity, it will become a threat. Thanks to the ubiquity of "big data" (as an over-hyped phrase mostly useful for attracting research funding, but also to at least some extent in reality), a new potentially multi-disciplinary field of data science is rapidly opening up. This field is attracting huge material resources, and will absorb much human talent. Statistical science has to be a part of this, for its own survival, but also for the sake of society. As Tim Harford has cogently argued (Harford 2014):

> Recall big data's four articles of faith. Uncanny accuracy is easy to overrate if we simply ignore false positives […]. The claim that causation has been "knocked off its pedestal" is fine if we are making predictions in a stable environment but not if the world is changing […] or if we ourselves hope to change it.

> The promise that "N = All", and therefore that sampling bias does not matter, is simply not true in most cases that count. As for the idea that "with enough data, the numbers speak for themselves" – that seems hopelessly naïve in data sets where spurious patterns vastly outnumber genuine discoveries.
> "Big data" has arrived, but big insights have not. The challenge now is to solve new problems and gain new answers – without making the same old statistical mistakes on a grander scale than ever.

It is a mistake to think that Bayes has no part to play in these developments, but more of us need to get more involved, and learn new tools, as in the way the Consensus Monte Carlo algorithm (Scott et al. 2013) exploits the Hadoop environment (White 2012) and the MapReduce programming model (Dean and Ghemawat 2008). Another direction that can prevent a potential schism between Bayesian modelling and highly complex models is to aim for modularity and local learning, that it, to abandon the goal of modelling big universes for analysing a series of small worlds, in spite of the loss of coherence, amd hence compromise to the Bayesian paradigm, that this entails. The curious case of the cut models presented in Plummer (2015) is an illustration of the potential for developing partial-information Bayesian inference tools where "small is beautiful" because this is the only viable solution.

5.2 Do we care enough about applications?

Bayesian computation began in order to answer rather practical problems—how can we perform a Bayesian analysis of these data using this model?—or the corresponding meta-problems—how can Bayesian analysis be performed generally and reliably for this class of models? The focus was applied methodology (although since the methods were new, they tended to be published in premier theory/methodology journals). Because the research community wanted to understand (the advantages, performance and limitations of) the methods they were advocating, more theoretical work started to be conducted, and, for example, many probabilists were attracted to study the Markov chains that MCMC methodologists created. The centre of mass of research activity drifted away from the original motivations, just as has happened in other areas of mathematically-rigorous computation.

At the same time, those working with data became more ambitious with regard to the scale of data, the complexity of modelling and the sophistication of analysis, all factors that have in principle (and often in fact) stimulated new developments in Bayesian computation. But to a large extent this is a rich, self-stimulating and self-supporting area of research; new applications may or may not need new computational techniques, but new techniques don't seem to need applications to justify themselves. It is apposite to ask to what extent is cutting-edge computational methodology research really delivering answers to questions that application domains are posing. And to what extent is cutting-edge computational methodology research successfully answering real questions?

We may not be unanimous about answers to these questions, except we can probably all agree they are "not entirely". We will also disagree about how much this matters, but again there may be something to agree about, that we have failed if methodological innovations disconnect completely from applications. Legitimate differences in research goals partially explain the trend in this direction, but it is fair to say that there is a big communication problem between the computational statistics community and many of the communities where Bayesian computational methods are applied. Unfortunately people in these communities do not always keep up with the state of the art in computational statistics. At the same time, statisticians are often not aware of important developments arising in other fields. (ABC is a good illustration: it took more than 5 years of development within the population genetics community before statisticians became aware the technique existed and a few more years before they realised this was proper Bayesian inference applied on approximate models.) We can perhaps blame the fact that there are not enough people working at the interface of the different communities, but life at the interface is not easy because multidisciplinary and interdisciplinary research is often seen as "marginal" by both communities and is thus difficult to publish, communicate, etc. Then there are of course problems in dissemination, related to the different writing styles, journals, computing languages, software, etc. of each community.

We strongly encourage those developing new techniques always to find a way to disseminate them in such a way that at least *somebody* else could use them, preferably someone without the ability to have invented the technique for themselves!—and advocate, of course, that successful dissemination be properly rewarded in our career structures.

In a somewhat parallel path, we have seen over the past decades the emergence of new languages and meta-languages intended to handle complexity both of problems and of solutions towards a wider audience of users. BUGS (Lunn et al. 2010) is the archetypal example of such languages and it has been successful to the extent that a large proportion of the users has a fairly limited statistical background and often even less of a computational background. However, the population of BUGS users and sympathisers is tiny compared to that of SAS or other corporate statistical systems. In this respect, we have failed to disseminate concepts like Bayesian analysis and wonderful tools like MCMC algorithms, because most people are unable to turn them into codes by themselves. (Perusing one of the numerous statistics

and machine-learning on-line forums like Cross Validated quickly exposes the methodological gap between academics and the masses!) It is unclear how novel programming developments like STAN (Stan Development Team 2014) are going to modify this picture, in that they still assume a decent understanding of both modelling and simulation issues. In that respect, network-based approaches as those covered by BUGS sound more promising towards "modelling locally to learn globally". Similarly, ABC software is either too specific, like DIYABC (Cornuet et al. 2008) which addresses only population genetic questions, or too dependent on the ability of the modeller to program simulated outcomes from the model under study.

5.3 Anticipating the future

In which of the areas we discuss do we expect a particular emphasis of effort, or significant progress, or do we see particular needs for new efforts or new directions?

One expectation is that in the future computational methodologies will be more flexible and malleable. Over the past 25 years Bayesian modelling and inference techniques have been applied successfully to thousands of problems across a wide range application domains. Each application brings its own constraints in terms of model dimensionality and complexity, data, inferences, accuracy and computing times. These constraints also vary significantly within specific applications. For example, in hyperspectral remote sensing, when a new Bayesian model is introduced it is often first explored and validated by MCMC sampling, then approximated with a variational Bayes method, and then approximated again so that it can be applied to gigabyte-large datasets by using optimisation techniques. Similarly, an interesting result revealed by a fast inference technique can be analysed more deeply with more reliable and accurate methods might. Therefore we expect that in the future the different main computational methodologies will become more adaptable and that the boundaries between them will be less well defined, with many algorithms developed that combine simulation, variational approximations and optimisation. These will be able to handle a wide spectrum of models, degrees of accuracy and computing times, as well as models that have some parts that are simple but high-dimensional and others that are more complex but that only involve low-dimensional components. This can be achieved by using approximations and optimisation to improve stochastic sampling, by using simulation within deterministic algorithms to handle specific parts of the model that are difficult to compute analytically, or in completely new and original ways.

We also anticipate that computational methodologies will continue to be challenged by larger and larger datasets. There is of course a threat that the whole field turns into a library of machine-learning techniques, with limited validation on reference learning sets and a quick turnover of methods, which would both impoverish the field and fail to reach a general audience of practitioners. We must retain a sense of the stochastic elements in data collection, data analysis, and inference, recognising uncertainty in data and models, to preserve the inductive strength of data science—seeing beyond the data we have to what it might have been, what it be next time, and where it came from.

Acknowledgments Supported in part by "SuSTaIn", EPSRC Grant EP/D063485/1, at the University of Bristol, and "i-like", EPSRC Grant EP/K014463/1, at the University of Warwick. Krzysztof Łatuszyński holds a Royal Society University Research Fellowship, and Marcelo Pereyra a Marie Curie Intra-European Fellowship for Career Development. Peter Green also holds a Distinguished Professorship at UTS, Sydney, and Christian Robert an Institut Universitaire de France chair at CEREMADE, Université Paris-Dauphine.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

Afonso, M., Bioucas-Dias, J., Figueiredo, M.: An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems. IEEE. Trans. Image Process. **20**, 681–695 (2011)

Albert, J.: Computational methods using a Bayesian hierarchical generalized linear model. J. Am. Stat. Assoc. **83**, 1037–1044 (1988)

Aldous, D., Krikun, M., Popovic, L.: Stochastic models for phylogenetic trees on higher-order taxa. J. Math. Biol. **56**, 525–557 (2008)

Alquier, P., Friel, N., Everitt, R., Boland, A.: Noisy Monte Carlo: convergence of Markov chains with approximate transition kernels. Stat. Comput. 1–19 (2014)

Andrieu, C., Doucet, A., Holenstein, R.: Particle Markov chain Monte Carlo (with discussion). J. R. Stat. Soc. B **72**(2), 269–342 (2011)

Andrieu, C., Moulines, É: On the ergodicity properties of some adaptive MCMC algorithms. Ann. Appl. Probab. **16**, 1462–1505 (2006)

Andrieu, C., Robert, C.: Controlled MCMC for optimal sampling. Technical report, Cahiers du Ceremade (2001)

Andrieu, C., Roberts, G.: The pseudo-marginal approach for efficient Monte Carlo computations. Ann. Stat. **37**, 697–725 (2009)

Andrieu, C., Tadić, V.B., Vihola, M.: On the stability of some controlled Markov chains and its applications to stochastic approximation with Markovian dynamic. Ann. Appl. Probab. **25**, 1–45 (2015)

Andrieu, C., Thoms, J.: A tutorial on adaptive MCMC. Stat. Comput. **18**, 343–373 (2008)

Andrieu, C., Vihola, M.: Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. Ann. Appl. Probab. **25**, 1030–1077 (2015)

Angelino, E., Kohler, E., Waterland, A., Seltzer, M., Adams, R.: Accelerating MCMC via parallel predictive prefetching. arXiv preprint arXiv:1403.7265 (2014)

Atchadé, Y.: An adaptive version for the Metropolis adjusted Langevin algorithm with a truncated drift. Methodol. Comput. Appl. Probab. **8**, 235–254 (2006)

Atchadé, Y., Rosenthal, J.: On adaptive Markov chain Monte Carlo algorithms. Bernoulli **11**, 815–828 (2005)

Atchadé, Y.F., Liu, J.S.: The Wang–Landau algorithm in general state spaces: applications and convergence analysis. Stat. Sin. **20**, 209–233 (2010)

Atchadé, Y., Fort, G., Moulines, E., Priouret, P.: Adaptive Markov chain Monte Carlo: theory and methods. In: Barber, D., Cemgil, A.T., Chiappa, S. (eds.) Bayesian Time Series Models, chap. 2, pp. 32–51. Cambridge University Press, Cambridge (2011a)

Atchadé, Y.F., Roberts, G.O., Rosenthal, J.S.: Towards optimal scaling of Metropolis-coupled Markov chain Monte Carlo. Stat. Comput. **21**, 555–568 (2011b)

Bai, Y., Roberts, G., Rosenthal, J.: On the containment condition for adaptive Markov chain Monte Carlo algorithms. Adv. Appl. Stat. **21**, 1–54 (2011)

Barber, S., Voss, J., Webster, M.: The rate of convergence for Approximate Bayesian computation. Electron. J. Stat. **9**, 80–105 (2015)

Bardenet, R., Doucet, A., Holmes, C.: Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach. In: Jebara, T., Xing, E.P. (eds.) Proceedings of the 31st International Conference on Machine Learning (ICML-14). JMLR Workshop and Conference Proceedings, pp 405–413 (2014)

Barthelmé, S., Chopin, N.: Expectation propagation for likelihood-free inference. J. Am. Stat. Assoc. **109**, 315–333 (2014)

Bauschke, H.H., Combettes, P.L.: Convex Analysis and Monotone Operator Theory in Hilbert Spaces. Springer, New York (2011)

Beaumont, M.: Estimation of population growth or decline in genetically monitored populations. Genetics **164**, 1139–1160 (2003)

Beaumont, M.: Approximate Bayesian computation in evolution and ecology. Annu. Rev. Ecol. Evol. Syst. **41**, 379–406 (2010)

Beaumont, M., Nielsen, R., Robert, C., Hey, J., Gaggiotti, O., Knowles, L., Estoup, A., Mahesh, P., Coranders, J., Hickerson, M., Sisson, S., Fagundes, N., Chikhi, L., Beerli, P., Vitalis, R., Cornuet, J.-M., Huelsenbeck, J., Foll, M., Yang, Z., Rousset, F., Balding, D., Excoffier, L.: In defense of model-based inference in phylogeography. Mol. Ecol. **19**(3), 436–446 (2010)

Beaumont, M., Zhang, W., Balding, D.: Approximate Bayesian computation in population genetics. Genetics **162**, 2025–2035 (2002)

Bédard, M.: Weak convergence of Metropolis algorithms for non-i.i.d. target distributions. Ann. Appl. Probab. **17**, 1222–1244 (2007)

Bédard, M., Douc, R., Moulines, E.: Scaling analysis of multiple-try MCMC methods. Stoch. Process. Appl. **122**, 758–786 (2012)

Bédard, M., Douc, R., Moulines, E.: Scaling analysis of delayed rejection MCMC methods. Methodol. Comput. Appl. Probab. **16**, 811–838 (2014)

Belle, E., Benazzo, A., Ghirotto, S., Colonna, V., Barbujani, G.: Comparing models on the genealogical relationships among Neandertal, Cro-Magnoid and modern Europeans by serial coalescent simulations. Heredity **102**, 218–225 (2008)

Bennett, J., Racine-Poon, A., Wakefield, J.: MCMC for nonlinear hierarchical models. In: Gilks, W., Richardson, S., Spiegelhalter, D. (eds.) Markov Chain Monte Carlo in Practice, pp. 339–358. Chapman and Hall, New York (1996)

Berger, J., Fienberg, S., Raftery, A., Robert, C.: Incoherent phylogeographic inference. Proc. Natl. Acad. Sci. USA **107**, E57 (2010)

Besag, J.: Comments on "Representations of knowledge in complex systems" by U. Grenander and M.I. Miller. J. R. Stat. Soc. B **56**, 591–592 (1994)

Besag, J., Green, P.: Spatial statistics and Bayesian computation (with discussion). J. R. Stat. Soc. B **55**, 25–38 (1993)

Besag, J., Green, P.J., Higdon, D., Mengersen, K.: Bayesian computation and stochastic systems (with discussion). Stat. Sci. **10**, 3–66 (1995)

Besag, J.E.: Nearest-neighbour systems and the auto-logistic model for binary data. J. R. Stat. Soc. B **34**, 75–83 (1972)

Beskos, A., Jasra, A., Muzaffer, E.A., Stuart, A.M.: Sequential Monte Carlo methods for Bayesian elliptic inverse problems. Stat. Comput. **25** (2015, in press)

Beskos, A., Papaspiliopoulos, O., Roberts, G., Fearnhead, P.: Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes (with discussion). J. R. Stat. Soc. B **68**, 333–382 (2006)

Beskos, A., Pillai, N., Roberts, G., Sanz-Serna, J.-M., Stuart, A.: Optimal tuning of the hybrid Monte Carlo algorithm. Bernoulli **19**, 1501–1534 (2013)

Beskos, A., Roberts, G., Stuart, A.: Optimal scalings for local Metropolis–Hastings chains on nonproduct targets in high dimensions. Ann. Appl. Probab. **19**, 863–898 (2009)

Betancourt, M.: A general metric for Riemannian manifold Hamiltonian Monte Carlo. In: Nielsen, F., Barbaresco, F. (eds) National Conference on the Geometric Science of Information. Lecture Notes in Computer Science, vol 8085, pp 327–334. Springer (2013)

Betancourt, M., Byrne, S., Girolami, M.: Optimizing the integrator step size for Hamiltonian Monte Carlo. arXiv preprint arXiv:1411.6669 (2014)

Betancourt, M.J., Byrne, S., Livingstone, S., Girolami, M.: The geometric foundations of Hamiltonian Monte Carlo. ArXiv e-prints. 1410.5110 (2014)

Biau, G., Cérou, F., Guyader, A.: New insights into Approximate Bayesian computation. Ann. l'IHP (Probab. Stat.) **51**, 376–403 (2014)

Blum, M.: Approximate Bayesian computation: a non-parametric perspective. J. Am. Stat. Assoc. **105**, 1178–1187 (2010)

Blum, M., François, O.: Non-linear regression models for approximate Bayesian computation. Stat. Comput. **20**, 63–73 (2010)

Blum, M., Nunes, M., Prangle, D., Sisson, S.: A comparative review of dimension reduction methods in Approximate Bayesian computation. Stat. Sci. **28**, 189–208 (2013)

Bornn, L., Pillai, N., Smith, A., Woodard, D.: A pseudo-marginal perspective on the ABC algorithm. ArXiv e-prints. 1404.6298 (2014)

Bou-Rabee, N., Hairer, M.: Nonasymptotic mixing of the MALA algorithm. IMA J. Numer. Anal. **33**(1), 80–110 (2012)

Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. Found. Trends Mach. Learn. **3**, 1–122 (2011)

Brockwell, A.: Parallel Markov chain Monte Carlo simulation by prefetching. J. Comput. Graph. Stat. **15**, 246–261 (2006)

Calderhead, B.: A general construction for parallelizing Metropolis–Hastings algorithms. Proc. Natl. Acad. Sci. USA **111**, 17408–17413 (2014)

Calvet, C., Czellar, V.: Accurate methods for Approximate Bayesian computation filtering. J. Financ. Econom. (2014, to appear)

Candès, E.J., Romberg, J.K., Tao, T.: Stable signal recovery from incomplete and inaccurate measurements. Commun. Pure Appl. Math. **59**, 1207–1223 (2006)

Candès, E.J., Tao, T.: The power of convex relaxation: near-optimal matrix completion. IEEE Trans. Inf. Theory **56**, 2053–2080 (2009)

Candès, E.J., Wakin, M.B.: An introduction to compressive sampling. IEEE Signal Process. Mag. **25**, 21–30 (2008)

Cappé, O., Robert, C., Rydén, T.: Reversible jump MCMC converging to birth-and-death MCMC, and more general continuous time samplers. J. R. Stat. Soc. B **65**, 679–700 (2002)

Carlin, B., Gelfand, A.: An iterative Monte Carlo method for nonconjugate Bayesian analysis. Stat. Comput. **1**, 119–128 (1991)

Carlin, B., Gelfand, A., Smith, A.: Hierarchical Bayesian analysis of change point problems. Appl. Stat. (C) **41**, 389–405 (1992)

Cevher, V., Becker, S., Schmidt, M.: Convex optimization for big data: scalable, randomized, and parallel algorithms for big data analytics. IEEE Signal Process. Mag. **31**, 32–43 (2014)

Chambolle, A.: An algorithm for total variation minimization and applications. J. Math. Imaging Vis. **20**, 89–97 (2004)

Chandrasekaran, V., Jordan, M.I.: Computational and statistical tradeoffs via convex relaxation. Proc. Natl. Acad. Sci. USA **110**, 1181–1190 (2013)

Chandrasekaran, V., Recht, B., Parrilo, P., Willsky, A.: The convex geometry of linear inverse problems. Found. Comput. Math. **12**, 805–849 (2012)

Chipman, H., George, E., McCulloch, R.: BART: Bayesian additive regression trees. Technical report, Acadia University. arXiv:0806.3286v1 (2008)

Chopin, N.: Inference and model choice for time-ordered hidden Markov models. J. R. Stat. Soc. B **69**(2), 269–284 (2007)

Chopin, N., Jacob, P.E., Papaspiliopoulos, O.: SMC2: an efficient algorithm for sequential analysis of state space models. J. R. Stat. Soc. B **75**, 397–426 (2013)

Christensen, O., Roberts, G., Rosenthal, J.: Scaling limits for the transient phase of local Metropolis–Hastings algorithms. J. R. Stat. Soc. B **67**, 253–268 (2005)

Combettes, P.L., Pesquet, J.-C.: A proximal decomposition method for solving convex variational inverse problems. Inverse Probl. **24**, 065014 (2008)

Combettes, P.L., Pesquet, J.-C.: Proximal splitting methods in signal processing. In: Bauschke, H.H., Burachik, R.S., Combettes, P.L., Elser, V., Luke, D.R., Wolkowicz, H. (eds.) Fixed-Point Algorithms for Inverse Problems in Science and Engineering, pp. 185–212. Springer, New York (2011)

Combettes, P.L., Pesquet, J.-C.: Primal–dual splitting algorithm for solving inclusions with mixtures of composite, Lipschitzian, and parallel-sum type monotone operators. Set-Valued Anal. **20**, 307–330 (2012)

Combettes, P.L., Pesquet, J.-C.: Stochastic Quasi-Fejér block-coordinate fixed point iterations with random sweeping. ArXiv e-prints. 1404.7536 (2014)

Cornuet, J.-M., Ravigné, V., Estoup, A.: Inference on population history and model checking using DNA sequence and microsatellite data with the software DIYABC (v1.0). BMC Bioinform. **11**, 401 (2010)

Cornuet, J.-M., Santos, F., Beaumont, M., Robert, C., Marin, J.-M., Balding, D., Guillemaud, T., Estoup, A.: Inferring population history with DIYABC: a user-friendly approach to Approximate Bayesian computation. Bioinformatics **24**, 2713–2719 (2008)

Cotter, S., Roberts, G., Stuart, A., White, D., et al.: MCMC methods for functions: modifying old algorithms to make them faster. Stat. Sci. **28**, 424–446 (2013)

Craiu, R., Rosenthal, J., Yang, C.: Learn from thy neighbour: parallel-chain and regional adaptive MCMC. J. Am. Stat. Assoc. **104**, 1454–1466 (2009)

Craiu, R.V., Meng, X.-L.: Multiprocess parallel antithetic coupling for backward and forward Markov chain Monte Carlo. Ann. Stat. **33**, 661–697 (2005)

Cucala, L., Marin, J.-M., Robert, C., Titterington, D.: Bayesian inference in k-nearest-neighbour classification models. J. Am. Stat. Assoc. **104**(485), 263–273 (2009)

Dean, J., Ghemawat, S.: MapReduce: simplified data processing on large clusters. Commun. ACM **51**, 107–113 (2008)

Del Moral, P., Doucet, A., Jasra, A.: Sequential Monte Carlo samplers. J. R. Stat. Soc. B **68**, 411–436 (2006)

Dellaportas, P., Papaspiliopoulos, O., Roberts, G.: Bayesian inference for non-Gaussian Ornstein–Uhlenbeck stochastic volatility processes. J. R. Stat. Soc. B **66**, 369–393 (2004)

Dellaportas, P., Wright, D.: Positive embedded integration in Bayesian analysis. Stat. Comput. **1**, 1–12 (1991)

Diaconis, P., Stroock, D.: Geometric bounds for eigenvalues of Markov chains. Ann. Appl. Probab. **1**, 36–61 (1991)

Didelot, X., Everitt, R., Johansen, A., Lawson, D.: Likelihood-free estimation of model evidence. Bayesian Anal. **6**, 48–76 (2011)

Diebolt, J., Robert, C.: Estimation of finite mixture distributions by Bayesian sampling. J. R. Stat. Soc. B **56**, 363–375 (1994)

Doll, J., Dion, D.: Generalized Langevin equation approach for atom/solid-surface scattering: numerical techniques for Gaussian generalized Langevin dynamics. J. Chem. Phys. **65**, 3762–3766 (1976)

Douc, R., Robert, C.: A vanilla Rao-Blackwellization of Metropolis–Hastings algorithms. Ann. Stat. **39**, 261–277 (2011)

Doucet, A., Pitt, M., Deligiannidis, G., Kohn, R.: Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. Biometrika (to appear 2015). arXiv preprint arXiv:1210.1871 (2012)

Drovandi, C., Pettitt, A., Fddy, M.: Approximate Bayesian computation using indirect inference. J. R. Stat. Soc. A **60**, 503–524 (2011)

Duane, S., Kennedy, A.D., Pendleton, B.J., Roweth, D.: Hybrid Monte Carlo. Phys. Lett. B **195**, 216–222 (1987)

Ermak, D.: A computer simulation of charged particles in solution. I. Technique and equilibrium properties. J. Chem. Phys. **62**, 4189–4196 (1975)

Excoffier, C., Leuenberger, D., Wegmann, L.: Bayesian computation and model selection in population genetics. arXiv:0901.2231 (2009)

Fearnhead, P., Clifford, P.: On-line inference for hidden Markov models via particle filters. J. R. Stat. Soc. B **65**, 887–899 (2003)

Fearnhead, P., Prangle, D.: Constructing summary statistics for Approximate Bayesian computation: semi-automatic Approximate Bayesian computation (with discussion). J. R. Stat. Soc. B **74**, 419–474 (2012)

Fort, G., Moulines, E., Priouret, P.: Convergence of adaptive and interacting Markov chain Monte Carlo algorithms. Ann. Stat. **39**, 3262–3289 (2011)

Frigessi, A., Gasemyr, J., Rue, H.: Antithetic coupling of two Gibbs sampler chains. Ann. Stat. **28**, 1128–1149 (2000)

Geman, S., Geman, D.: Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. IEEE Trans. Pattern Anal. Mach. Intell. **6**, 721–741 (1984)

Geyer, C.J.: Practical Markov chain Monte Carlo. Stat. Sci. **7**, 473–483 (1992)

Ghirotto, S., Mona, S., Benazzo, A., Paparazzo, F., Caramelli, D., Barbujani, G.: Inferring genealogical processes from patterns of Bronze-Age and modern DNA variation in Sardinia. Mol. Biol. Evol. **27**, 875–886 (2010)

Gilks, W., Roberts, G., George, E.: Adaptive direction sampling. J. R. Stat. Soc. D **43**, 179–189 (1994)

Gilks, W., Roberts, G., Sahu, S.: Adaptive Markov chain Monte Carlo through regeneration. J. Am. Stat. Assoc. **93**, 1045–1054 (1998)

Girolami, M., Calderhead, B.: Riemann manifold Langevin and Hamiltonian Monte Carlo methods. J. R. Stat. Soc. B **73**, 123–214 (2011)

Girolami, M., Lyne, A.-M., Strathmann, H., Simpson, D., Atchadé, Y.: Playing Russian roulette with intractable likelihoods. arXiv preprint. arXiv:1306.4032 (2013)

Green, P.: Reversible jump MCMC computation and Bayesian model determination. Biometrika **82**, 711–732 (1995)

Green, P.J.: Chap. 19: MAD-Bayes matching and alignment for labelled and unlabelled configurations. In: Dryden, I.L., Kent, J.T. (eds.) Geometry Driven Statistics, pp. 365–375. Wiley, Chichester (2015)

Grelaud, A., Marin, J.-M., Robert, C., Rodolphe, F., Tally, F.: Likelihood-free methods for model choice in Gibbs random fields. Bayesian Anal. **3**(2), 427–442 (2009)

Grenander, U., Miller, M.I.: Representations of knowledge in complex systems (with discussion and a reply by the authors). J. R. Stat. Soc. B **56**, 549–603 (1994)

Griffin, J., Łatuszyński, K., Steel, M.: Individual adaptation: an adaptive MCMC scheme for variable selection problems arXiv preprint. arXiv:1412.6760 (2014)

Guillemaud, T., Beaumont, M., Ciosi, M., Cornuet, J.-M., Estoup, A.: Inferring introduction routes of invasive species using approximate Bayesian computation on microsatellite data. Heredity **104**, 88–99 (2009)

Haario, H., Laine, M., Mira, A., Saksman, E.: DRAM: efficient adaptive MCMC. Stat. Comput. **16**, 339–354 (2006)

Haario, H., Saksman, E., Tamminen, J.: An adaptive Metropolis algorithm. Bernoulli **7**, 223–242 (2001)

Hairer, E., Lubich, C., Wanner, G.: Geometric numerical integration: structure-preserving algorithms for ordinary differential equations. In: Springer Series in Computational Mathematics, 2nd edn, vol 31. Springer, Berlin (2006)

Harford, T.: Big data: are we making a big mistake? Significance **11**, 14–19 (2014)

Hastings, W.: Monte Carlo sampling methods using Markov chains and their application. Biometrika **57**, 97–109 (1970)

Huelsenbeck, J.P., Ronquist, F.: MRBAYES: Bayesian inference of phylogenetic trees. Bioinformatics **17**, 754–755 (2001)

Jaakkola, T., Jordan, M.: Bayesian parameter estimation via variational methods. Stat. Comput. **10**, 25–37 (2000)

Jacob, P., Robert, C., Smith, M.: Using parallel computation to improve independent Metropolis–Hastings based estimation. J. Comput. Graph. Stat. **20**, 616–635 (2011)

Jacob, P.E., Ryder, R.J.: The Wang–Landau algorithm reaches the flat histogram criterion in finite time. Ann. Appl. Probab. **24**, 34–53 (2014)

Ji, C., Schmidler, S.C.: Adaptive Markov chain Monte Carlo for Bayesian variable selection. J. Comput. Graph. Stat. **22**, 708–728 (2013)

Jordan, M.I., Ghahramani, Z., Jaakkola, T.S., Saul, L.K.: An introduction to variational methods for graphical models. Mach. Learn. **37**, 183–233 (1999)

Jourdain, B., Lelièvre, T., Miasojedow, B.: Optimal scaling for the transient phase of the random walk Metropolis algorithm: the mean-field limit. Ann. Appl. Probab. arXiv preprint. arXiv:1210.7639 (to appear, 2012)

Jourdain, B., Lelièvre, T., Miasojedow, B.: Optimal scaling for the transient phase of Metropolis Hastings algorithms: the longtime behavior. Bernoulli **20**, 1930–1978 (2014)

Kent, J.: Time-reversible diffusions. Adv. Appl. Probab. **10**(4), 819–835 (1978)

Kloeden, P.E., Platen, E.: Numerical solution of stochastic differential equations. In: Applications of Mathematics (New York), vol 23. Springer, Berlin (1992)

Komodakis, N., Pesquet, J.-C.: Playing with duality: an overview of recent primal–dual approaches for solving large-scale optimization problems. ArXiv e-prints. 1406.5429 (2014)

Korattikara, A., Chen, Y., Welling, M.: Austerity in MCMC land: cutting the Metropolis–Hastings budget. arXiv preprint. arXiv:1304.5299 (2013)

Kou, S.C., Zhou, Q., Wong, W.H.: Equi-energy sampler with applications in statistical inference and statistical mechanics (with discussions and a rejoinder by the authors). Ann. Stat. **34**, 1581–1652 (2006)

Lamnisos, D., Griffin, J.E., Steel, M.F.: Adaptive Monte Carlo for Bayesian variable selection in regression models. J. Comput. Graph. Stat. **22**, 729–748 (2013)

Lange, K., Chi, E.C., Zhou, H.: A brief survey of modern optimization for statisticians. Int. Stat. Rev. **82**, 46–70 (2014)

Larget, B., Simon, D.L.: Markov chain Monte Carlo algorithms for the Bayesian analysis of phylogenetic trees. Mol. Biol. Evol. **16**, 750–759 (1999)

Łatuszyński, K., Kosmidis, I., Papaspiliopoulos, O., Roberts, G.O.: Simulating events of unknown probabilities via reverse time martingales. Random Struct. Algorithms **38**, 441–452 (2011)

Łatuszyński, K., Roberts, G., Rosenthal, J.: Adaptive Gibbs samplers and related MCMC methods. Ann. Appl. Probab. **23**(1), 66–98 (2013)

Łatuszyński, K., Rosenthal, J.S.: The containment condition and AdapFail algorithms. J. Appl. Probab. **51**, 1189–1195 (2014)

Lee, A., Łatuszyński, K.: Variance bounding and geometric ergodicity of Markov chain Monte Carlo kernels for Approximate Bayesian computation. Biometrika **101**, 655–671 (2014)

Lee, A., Yau, C., Giles, M., Doucet, A., Holmes, C.: On the utility of graphics cards to perform massively parallel simulation with advanced Monte Carlo methods. Arxiv preprint. arXiv:0905.2441 (2009)

Lee, H.K., Okabe, Y., Landau, D.P.: Convergence and refinement of the Wang–Landau algorithm. Technical report (2005)

Leuenberger, C., Wegmann, D.: Bayesian computation and model selection without likelihoods. Genetics **184**, 243–252 (2010)

Levin, D.A., Peres, Y., Wilmer, E.L.: Markov Chains and Mixing Times. American Mathematical Society, Providence (2009)

Lindsten, F., Jordan, M.I., Schön, T.B.: Particle Gibbs with ancestor sampling. ArXiv e-prints. 1401.0604 (2014)

Lunn, D., Thomas, A., Best, N., Spiegelhalter, D.: The BUGS Book: A Practical Introduction to Bayesian Analysis. Chapman and Hall/CRC Press Boca Raton (2010)

MacKay, D.J.C.: Information Theory, Inference and Learning Algorithms. Cambridge University Press, Cambridge (2002)

Maclaurin, D., Adams, R.P.: Firefly Monte Carlo: exact MCMC with subsets of data. arXiv preprint. arXiv:1403.5693 (2014)

Marin, J., Pillai, N., Robert, C., Rousseau, J.: Relevant statistics for Bayesian model choice. J. R. Stat. Soc. B **76**(5), 833–859 (2014)

Marin, J., Pudlo, P., Robert, C., Ryder, R.: Approximate Bayesian computational methods. Stat. Comput. **21**(2), 1–14 (2012)

Marshall, T., Roberts, G.: An adaptive approach to Langevin MCMC. Stat. Comput. **22**, 1041–1057 (2012)

Martinet, B.: Regularisation d'inéquations variationelles par approximations successives. Rev. Fran. d'Autom. Infomat. Rech. Opérationelle **4**, 154–159 (1970)

Medina-Aguayo, F.J., Lee, A., Roberts, G.O.: Stability of noisy Metropolis–Hastings. arXiv preprint. arXiv:1503.07066 (2015)

Mengersen, K., Tweedie, R.: Rates of convergence of the Hastings and Metropolis algorithms. Ann. Stat. **24**, 101–121 (1996)

Metropolis, N.: The beginning of the Monte Carlo method. Los Alamos Sci. **15**, 125–130 (1987)

Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E.: Equations of state calculations by fast computing machines. J. Chem. Phys. **21**, 1087–1092 (1953)

Meyn, S., Tweedie, R.: Markov Chains and Stochastic Stability. Cambridge University Press, Cambridge (2009)

Miasojedow, B., Moulines, E., Vihola, M.: An adaptive parallel tempering algorithm. J. Comput. Graph. Stat. **22**, 649–664 (2013)

Minka, T.: Expectation propagation for approximate Bayesian inference. In: Jack, D.K., Breese, S. (eds) UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence, pp 362–369. University of Washington, Seattle (2001)

Minsker, S., Srivastava, S., Lin, L., Dunson, D.B.: Robust and scalable Bayes via a median of subset posterior measures. arXiv preprint. arXiv:1403.2660 (2014)

Møller, J., Pettitt, A.N., Reeves, R., Berthelsen, K.K.: An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. Biometrika **93**, 451–458 (2006)

Moreau, J.-J.: Fonctions convexes duales et points proximaux dans un espace Hilbertien. C. R. Acad. Sci. Paris Sér. A Math. **255**, 2897–2899 (1962)

Muff, S., Riebler, A., Rue, H., Saner, P., Held, L.: Bayesian analysis of measurement error models using INLA. ArXiv e-prints. 1302.3065 (2013)

Murray, I., Ghahramani, Z., MacKay, D.: MCMC for doubly-intractable distributions. In: Uncertainty in Artificial Intelligence. UAI-2006 (2006a)
Murray, I., MacKay, D.J., Ghahramani, Z., Skilling, J.: Nested sampling for Potts models. In: Weiss, Y., Schölkopf, B., Platt, J. (eds.) Advances in Neural Information Processing Systems, vol. 18, pp. 947–954. MIT Press, Cambridge (2006b)
Naylor, J., Smith, A.: Application of a method for the efficient computation of posterior distributions. Appl. Stat. **31**, 214–225 (1982)
Neal, P., Roberts, G., Yuen, W.K.: Optimal scaling of random walk Metropolis algorithms with discontinuous target densities. Ann. Appl. Probab. **22**, 1880–1927 (2012)
Neal, R.: Bayesian Learning for Neural Networks. Lecture Notes, vol. 118. Springer, New York (1999)
Neal, R.: MCMC using Hamiltonian dynamics. In: Brooks, S., Gelman, A., Jones, G., Meng, X.-L. (eds.) Handbook of Markov Chain Monte Carlo, pp. 113–162. Chapman and Hall/CRC Press, Boca Raton (2013)
Neiswanger, W., Wang, C., Xing, E.: Asymptotically exact, embarrassingly parallel MCMC. arXiv preprint. arXiv:1311.4780 (2013)
Nesterov, Y.: Introductory lectures on convex optimization: a basic course. In: Applied Optimization, vol 87. Kluwer Academic Publishers, Boston (2004)
Nott, D., Kohn, R.: Adaptive sampling for Bayesian variable selection. Biometrika **92**, 747–763 (2005)
Oliveira, J., Bioucas-Dias, J., Figueiredo, M.: Adaptive total variation image deblurring: a majorization–minimization approach. Signal Process. **89**, 1683–1693 (2009)
Owen, A.B.: Empirical Likelihood. Chapman and Hall, Boca Raton (2001)
Paisley, J., Blei, D.M., Jordan, M.I.: Variational Bayesian inference with stochastic search. In: Proceedings of the 29th International Conference on Machine Learning (ICML-12), pp 1367–1374 (2012)
Papaspiliopoulos, O., Roberts, G.O., Sköld, M.: A general framework for the parametrization of hierarchical models. Stat. Sci. **22**, 59–73 (2007)
Parikh, N., Boyd, S.: Proximal algorithms. Found. Trends Optim. **1**, 123–231 (2014)
Patin, E., Laval, G., Barreiro, L., Salas, A., Semino, O., Santachiara-Benerecetti, S., Kidd, K., Kidd, J., Van Der Veen, L., Hombert, J., et al.: Inferring the demographic history of African farmers and Pygmy hunter-gatherers using a multilocus resequencing data set. PLoS Genet. **5**, e1000448 (2009)
Pearson, K.: Contribution to the mathematical theory of evolution. Proc. Trans. R. Soc. A **185**, 71–110 (1894)
Pereyra, M.: Proximal Markov chain Monte Carlo algorithms. Stat. Comput. (2015, to appear)
Peskun, P.: Optimum Monte Carlo sampling using Markov chains. Biometrika **60**, 607–612 (1973)
Pesquet, J.-C., Pustelnik, N.: A parallel inertial proximal optimization method. Pac. J. Optim. **8**, 273–305 (2012)
Pillai, N.S., Smith, A.: Ergodicity of approximate MCMC chains with applications to large data sets. arXiv preprint. arXiv:1405.0182 (2014)
Pillai, N.S., Stuart, A.M., Thiéry, A.H.: Optimal scaling and diffusion limits for the Langevin algorithm in high dimensions. Ann. Appl. Probab. **22**, 2320–2356 (2012)
Plummer, M.: Cuts in Bayesian graphical models. Stat. Comput. **25**(1), 37–43 (2015)
Potts, R.B.: Some generalized order–disorder transitions. Proc. Camb. Philos. Soc. **48**, 106–109 (1952)
Pritchard, J., Seielstad, M., Perez-Lezaun, A., Feldman, M.: Population growth of human Y chromosomes: a study of Y chromosome microsatellites. Mol. Biol. Evol. **16**, 1791–1798 (1999)
Pudlo, P., Marin, J.-M., Estoup, A., Cornuet, J.-M., Gautier, M., Robert, C.P.: ABC model choice via random forests. ArXiv e-prints. 1406.6288 (2014)
Quiroz, M., Villani, M., Kohn, R.: Speeding up MCMC by efficient data subsampling. arXiv preprint. arXiv:1404.4178 (2014)
Raguet, H., Fadili, J., Peyré, G.: A generalized forward–backward splitting. SIAM J. Imaging Sci. **6**, 1199–1226 (2013)
Ramakrishnan, U., Hadly, E.: Using phylochronology to reveal cryptic population histories: review and synthesis of 29 ancient DNA studies. Mol. Ecol. **18**, 1310–1330 (2009)
Richardson, S., Bottolo, L., Rosenthal, J.: Bayesian models for sparse regression analysis of high dimensional data. Bayesian Stat. **9**, 539–569 (2010)
Richardson, S., Green, P.: On Bayesian analysis of mixtures with an unknown number of components (with discussion). J. R. Stat. Soc. B **59**, 731–792 (1997)
Robert, C., Casella, G.: A short history of Markov chain Monte Carlo: subjective recollections from incomplete data. Stat. Sci. **26**, 102–115 (2011)
Robert, C., Cornuet, J.-M., Marin, J.-M., Pillai, N.: Lack of confidence in ABC model choice. Proc. Natl. Acad. Sci. USA **108**(37), 15112–15117 (2011)
Roberts, G., Gelman, A., Gilks, W.: Weak convergence and optimal scaling of random walk Metropolis algorithms. Ann. Appl. Probab. **7**, 110–120 (1997)
Roberts, G., Rosenthal, J.: Optimal scaling of discrete approximations to Langevin diffusions. J. R. Stat. Soc. B **60**, 255–268 (1998)
Roberts, G., Rosenthal, J.: Optimal scaling for various Metropolis–Hastings algorithms. Stat. Sci. **16**, 351–367 (2001)
Roberts, G., Rosenthal, J.: General state space Markov chains and MCMC algorithms. Probab. Surv. **1**, 20–71 (2004)
Roberts, G., Rosenthal, J.: Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms. J. Appl. Probab. **44**, 458 (2007)
Roberts, G., Rosenthal, J.: Examples of adaptive MCMC. J. Comput. Graph. Stat. **18**, 349–367 (2009)
Roberts, G., Stramer, O.: Langevin diffusions and Metropolis–Hastings algorithms. Methodol. Comput. Appl. Probab. **4**, 337–357 (2002)
Roberts, G., Tweedie, R.: Exponential convergence of Langevin distributions and their discrete approximations. Bernoulli **2**, 341–363 (1996a)
Roberts, G., Tweedie, R.: Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. Biometrika **83**, 95–110 (1996b)
Roberts, G.O.: Markov chain concepts related to sampling algorithms. In: Markov Chain Monte Carlo in Practice, p 57. Chapman and Hall, London (1996)
Roberts, G.O.: Optimal Metropolis algorithms for product measures on the vertices of a hypercube. Stoch. Stoch. Rep. **62**, 275–283 (1998)
Roberts, G.O., Rosenthal, J.S.: Minimising MCMC variance via diffusion limits, with an application to simulated tempering. Ann. Appl. Probab. **24**, 131–149 (2014)
Roberts, G.O., Stramer, O.: On inference for partially observed nonlinear diffusion models using the Metropolis–Hastings algorithm. Biometrika **88**, 603–621 (2001)
Rockafellar, R.T.: Monotone operators and the proximal point algorithm. SIAM J. Control Optim. **14**, 877–898 (1976)
Rossky, P., Doll, J., Friedman, H.: Brownian dynamics as smart Monte Carlo simulation. J. Chem. Phys. **69**, 4628–4633 (1978)
Rubin, D.: Bayesianly justifiable and relevant frequency calculations for the applied statistician. Ann. Stat. **12**, 1151–1172 (1984)
Rubinstein, R.Y.: Simulation and the Monte Carlo Method. Wiley, New York (1981)
Rudolf, D., Schweizer, N.: Perturbation theory for Markov chains via Wasserstein distance. arXiv preprint. arXiv:1503.04123 (2015)

Rue, H., Martino, S., Chopin, N.: Approximate Bayesian inference for latent Gaussian models using integrated nested Laplace approximations (with discussion). J. R. Stat. Soc. B **71**, 319–392 (2009)
Saksman, E., Vihola, M.: On the ergodicity of the adaptive Metropolis algorithm on unbounded domains. Ann. Appl. Probab. **20**, 2178–2203 (2010)
Salimans, T., Knowles, D.A.: Fixed-form variational posterior approximation through stochastic linear regression. Bayesian Anal. **8**, 837–882 (2013)
Saloff-Coste, L.: Lectures on finite Markov chains. In: Lectures on Probability Theory and Statistics (Saint-Flour, 1996). Lecture Notes in Mathematics, vol 1665, pp 301–413. Springer, Berlin (1997)
Schreck, A., Fort, G., Corff, S.L., Moulines, E.: A shrinkage-thresholding Metropolis adjusted Langevin algorithm for Bayesian variable selection. arXiv preprint. arXiv:1312.5658 (2013)
Schrödle, B., Held, L.: A primer on disease mapping and ecological regression using INLA. Comput. Stat. **26**, 241–258 (2011)
Scott, S., Blocker, A., Bonassi, F., Chipman, H., George, E., McCulloch, R.: Bayes and big data: the consensus Monte Carlo algorithm. In: EFaBBayes 250 Conference, vol 16 (2013)
Searle, S., Casella, G., McCulloch, C.: Variance Components. Wiley, New York (1992)
Sherlock, C.: Optimal scaling for the pseudo-marginal random walk Metropolis: insensitivity to the noise generating mechanism. arXiv preprint. arXiv:1408.4344 (2014)
Sherlock, C., Thiery, A.H., Roberts, G.O., Rosenthal, J.S.: On the efficiency of pseudo-marginal random walk Metropolis algorithms. Ann. Stat. **43**, 238–275 (2014)
Smith, A., Skene, A., Shaw, J., Naylor, J., Dransfield, M.: The implementation of the Bayesian paradigm. Commun. Stat. Theory Methods **14**, 1079–1102 (1985)
Smith, A., Skene, A.M., Shaw, J.E.H., Naylor, J.C.: Progress with numerical and graphical methods for practical Bayesian statistics. J. R. Stat. Soc. D **36**, 75–82 (1987)
Solonen, A., Ollinaho, P., Laine, M., Haario, H., Tamminen, J., Järvinen, H.: Efficient MCMC for climate model parameter estimation: parallel adaptive chains and early rejection. Bayesian Anal. **7**, 715–736 (2012)
Stan Development Team: STAN: A C++ Library for Probability and Sampling, Version 2.5.0. http://mc-stan.org/ (2014)
Stephens, M.: Bayesian analysis of mixture models with an unknown number of components—an alternative to reversible jump methods. Ann. Stat. **28**, 40–74 (2000)
Stramer, O., Tweedie, R.: Langevin-type models I: diffusions with given stationary distributions and their discretizations. Methodol. Comput. Appl. Probab. **1**, 283–306 (1999a)
Stramer, O., Tweedie, R.: Langevin-type models II: self-targeting candidates for MCMC algorithms. Methodol. Comput. Appl. Probab. **1**, 307–328 (1999b)
Strathmann, H., Sejdinovic, D., Girolami, M.: Unbiased Bayes for big data: paths of partial posteriors. arXiv preprint. arXiv:1501.03326 (2015)
Strid, I.: Efficient parallelisation of Metropolis–Hastings algorithms using a prefetching approach. Comput. Stat. Data Anal. **54**, 2814–2835 (2010)
Suchard, M., Wang, Q., Chan, C., Frelinger, J., Cron, A., West, M.: Understanding GPU programming for statistical computation: studies in massively parallel massive mixtures. J. Comput. Graph. Stat. **19**, 418–438 (2010)
Tanner, M., Wong, W.: The calculation of posterior distributions by data augmentation. J. Am. Stat. Assoc. **82**, 528–550 (1987)
Tavaré, S., Balding, D., Griffith, R., Donnelly, P.: Inferring coalescence times from DNA sequence data. Genetics **145**, 505–518 (1997)
Taylor, K.: PhD Thesis, University of Warwick, 2014
Teh, Y.W., Thiéry, A., Vollmer, S.: Consistency and fluctuations for stochastic gradient Langevin dynamics. arXiv preprint. arXiv:1409.0578 (2014)
Templeton, A.: Statistical hypothesis testing in intraspecific phylogeography: nested clade phylogeographical analysis vs. Approximate Bayesian computation. Mol. Ecol. **18**(2), 319–331 (2008)
Templeton, A.: Coherent and incoherent inference in phylogeography and human evolution. Proc. Natl. Acad. Sci. USA **107**(14), 6376–6381 (2010)
Tierney, L.: A note on Metropolis–Hastings kernels for general state spaces. Ann. Appl. Probab. **8**, 1–9 (1998)
Toni, T., Welch, D., Strelkowa, N., Ipsen, A., Stumpf, M.: Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. J. R. Soc. Interface **6**, 187–202 (2009)
VanDerwerken, D.N., Schmidler, S.C.: Parallel Markov chain Monte Carlo. ArXiv e-prints. 1312.7479 (2013)
Verdinelli, I., Wasserman, L.: Bayesian analysis of outlier problems using the Gibbs sampler. Stat. Comput. **1**, 105–117 (1991)
Verdu, P., Austerlitz, F., Estoup, A., Vitalis, R., Georges, M., Théry, S., Froment, A., Le Bomin, S., Gessain, A., Hombert, J.-M., Van der Veen, L., Quintana-Murci, L., Bahuchet, S., Heyer, E.: Origins and genetic diversity of pygmy hunter-gatherers from Western Central Africa. Curr. Biol. **19**, 312–318 (2009)
Vihola, M.: Robust adaptive Metropolis algorithm with coerced acceptance rate. Stat. Comput. **22**, 997–1008 (2012)
Wakefield, J., Gelfand, A., Smith, A.: Efficient generation of random variates via the ratio-of-uniforms method. Stat. Comput. **1**, 129–133 (1991)
Wang, F., Landau, D.P.: Efficient, multiple-range random walk algorithm to calculate the density of states. Phys. Rev. Lett. **86**, 2050–2053 (2001)
Wang, X., Dunson, D.: Parallelizing MCMC via Weierstrass sampler. arXiv preprint. arXiv:1312.4605 (2013)
Wang, Z., Mohamed, S., de Freitas, N.: Adaptive Hamiltonian and Riemann manifold Monte Carlo. In: Proceedings of the 30th International Conference on Machine Learning, pp 1462–1470 (2013)
Wegmann, D., Excoffier, L.: Bayesian inference of the demographic history of chimpanzees. Mol. Biol. Evol. **27**, 1425–1435 (2010)
Welling, M., Teh, Y.W.: Bayesian learning via stochastic gradient Langevin dynamics. In: Proceedings of the 28th International Conference on Machine Learning (ICML-11), pp 681–688 (2011)
White, T.: Hadoop: The Definitive Guide. O'Reilly Media, Sebastopol (2012)
Whiteley, N., Andrieu, C., Doucet, A.: Efficient Bayesian inference for switching state-space models using discrete particle Markov chain Monte Carlo methods. ArXiv e-prints. 1011.2437 (2010)
Wilkinson, D.: Parallel Bayesian computation. In: Kontoghiorghes, E.J. (ed.) Handbook of Parallel Computing and Statistics, pp. 481–512. Marcel Dekker/CRC Press, New York (2005)
Wilkinson, D.: The particle marginal Metropolis–Hastings (PMMH) particle MCMC algorithm. Darren Wilkinson's research blog. https://darrenjw.wordpress.com/2011/05/17/the-particle-marginal-metropolis-hastings-pmmh-particle-mcmc-algorithm/ (2011a)
Wilkinson, D.J.: Stochastic Modelling for Systems Biology, 2nd edn. CRC Press, New York (2011)
Wilkinson, R.: Approximate Bayesian computation (ABC) gives exact results under the assumption of model error. Stat. Appl. Genet. Mol. Biol. **12**, 129–141 (2013)
Xifara, T., Sherlock, C., Livingstone, S., Byrne, S., Girolami, M.: Langevin diffusions and the Metropolis-adjusted Langevin algorithm. Stat. Probab. Lett. **91**, 14–19 (2014)