



Computing Bayes: From Then 'Til Now 1

Gael M. Martin, David T. Frazier, Christian P. Robert

► To cite this version:

Gael M. Martin, David T. Frazier, Christian P. Robert. Computing Bayes: From Then 'Til Now 1. 2022. ⟨hal-03911805⟩

HAL Id: hal-03911805

<https://hal.science/hal-03911805v1>

Preprint submitted on 23 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Computing Bayes: From Then ‘Til Now¹

Gael M. Martin², David T. Frazier and Christian P. Robert

Abstract. This paper takes the reader on a journey through the history of Bayesian computation, from the 18th century to the present day. Beginning with the one-dimensional integral first confronted by Bayes in 1763, we highlight the key contributions of: Laplace, Metropolis (and, importantly, his co-authors!), Hammersley and Handscomb, and Hastings, all of which set the foundations for the computational revolution in the late 20th century — led, primarily, by Markov chain Monte Carlo (MCMC) algorithms. A very short outline of 21st century computational methods — including pseudo-marginal MCMC, Hamiltonian Monte Carlo, sequential Monte Carlo, and the various ‘approximate’ methods — completes the paper.

Key words and phrases: History of Bayesian computation, Laplace approximation, Metropolis-Hastings algorithm, importance sampling, Markov chain Monte Carlo, pseudo-marginal methods, Hamiltonian Monte Carlo, sequential Monte Carlo, approximate Bayesian methods.

1. THE BEGINNING

December 23 1763, London: Richard Price reads to the Royal Society a paper penned by a past Fellow, the late Thomas Bayes:

‘An Essay Towards Solving a Problem in the Doctrine of Chances.’

With that reading, the concept of ‘inverse probability’ — *Bayesian inference* as we know it now — has its first public airing.

To our modern eyes, the problem tackled by Bayes in his essay is a simple one: If one performs n independent Bernoulli trials, with a probability, θ , of success on each trial, what is the probability — given n outcomes — of θ lying between two values, a and b ? The answer Bayes offered is equally simple to re-cast in modern terminology. Define $Y_i|\theta \sim i.i.d. \text{Bernoulli}(\theta)$, $i = 1, 2, \dots, n$; record the observed sequence of successes ($Y_i = 1$) and failures ($Y_i = 0$) as $\mathbf{y} = (y_1, y_2, \dots, y_n)'$; denote by $p(\mathbf{y}|\theta)$ the likelihood function for θ ; and invoke a Uniform prior, $p(\theta)$, on the interval $(0, 1)$. Bayes sought:

$$(1) \quad \mathbb{P}(a < \theta < b|\mathbf{y}) = \int_a^b p(\theta|\mathbf{y})d\theta,$$

Monash University, Melbourne^{presep=}. *Monash University, Melbourne*^{presep=}. *Université Paris Dauphine and University of Warwick*^{presep=}.

¹Material that appeared in an earlier paper, entitled: ‘Computing Bayes: Bayesian Computation from 1763 to the 21st Century’ (Martin, Frazier and Robert, 2020), now appears (in amended form) in two separate papers: the current one, in which a historical overview of and timeline for all computational developments is presented; and a second one (Martin, Frazier and Robert, 2021), which provides a detailed review of 21st century approximate methods. Martin and Frazier have been supported by Australian Research Council Discovery Grants DP170100729 and DP200101414, and the Australian Centre of Excellence in Mathematics and Statistics. Frazier has also been supported by Australian Research Council Discovery Early Career Researcher Award DE200101070. Robert has been partly supported by a senior chair (2016-2021) from l’Institut Universitaire de France and by a Prairie chair from the Agence Nationale de la Recherche (ANR-19-P3IA-0001).

²Corresponding author

where $p(\theta|\mathbf{y})$ denotes the posterior probability density function (pdf) for θ ,

$$(2) \quad p(\theta|\mathbf{y}) = \frac{p(\mathbf{y}|\theta)p(\theta)}{p(\mathbf{y})},$$

$p(\mathbf{y}) = \int_0^1 p(\mathbf{y}|\theta)p(\theta)d\theta$ defines the marginal likelihood, and the scale factor $[p(\mathbf{y})]^{-1}$ in (2) ensures that $p(\theta|\mathbf{y})$ integrates to one. Given the Bernoulli assumption for $Y|\theta$, the Uniform prior on θ , and defining $x = \sum_{i=1}^n y_i$, $p(\theta|\mathbf{y})$ has a closed-form representation as the Beta pdf,

$$(3) \quad p(\theta|\mathbf{y}) = [B(x+1, n-x+1)]^{-1} \theta^x (1-\theta)^{n-x},$$

where $B(x+1, n-x+1) = \Gamma(x+1)\Gamma(n-x+1)/\Gamma(n+2) = \int_0^1 \theta^x (1-\theta)^{n-x} d\theta$ is the Beta function, and $\Gamma(x)$ is the Gamma function. Bayesian inference — namely, quantification of uncertainty about an unknown θ , conditioned on known data, \mathbf{y} — thus first emerges as the analytical solution to a particular inverse probability problem.¹

Bayes, however, did not seek the pdf in (3) *per se*. Rather, he wished to evaluate the probability in (1) which, for either $a \neq 0$ or $b \neq 1$, involved evaluation of the incomplete Beta function. Except for when either x or $(n-x)$ was small, a closed-form solution to (1) eluded Bayes. Hence, despite the analytical availability of $p(\theta|\mathbf{y})$ via (2) — ‘Bayes’ rule’ as it is now known — the quantity that was of interest to Bayes needed to be estimated, or *computed*. The quest for a computational solution to a Bayesian problem was thus born.

2. PRELIMINARIES

2.1 Why Do We Need Numerical Computation?

Bayes’ probability of interest in (1) can, of course, be expressed as a posterior expectation, $\mathbb{E}(\mathbb{I}_{[a,b]}|\mathbf{y}) = \int_{\Theta} \mathbb{I}_{[a,b]} p(\theta|\mathbf{y}) d\theta$, where $\mathbb{I}_{[a,b]}$ is the indicator function on the interval $[a, b]$. Generalizing at this point to any problem with unknown $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_p)' \in \Theta$ and joint posterior pdf $p(\boldsymbol{\theta}|\mathbf{y})$, most Bayesian quantities of interest are posterior expectations of some function $g(\boldsymbol{\theta})$ and, hence, can be expressed as,

$$(4) \quad \mathbb{E}(g(\boldsymbol{\theta})|\mathbf{y}) = \int_{\Theta} g(\boldsymbol{\theta}) p(\boldsymbol{\theta}|\mathbf{y}) d\boldsymbol{\theta}.$$

In addition to posterior probabilities like that of Bayes, familiar examples include posterior moments, marginal posterior densities and moments, predictive densities and posterior expectations of loss functions. Moreover, conditioning explicitly on the assumed model \mathcal{M} , the marginal likelihood of the model is the *prior* expectation,

$$(5) \quad p(\mathbf{y}|\mathcal{M}) = \int_{\Theta} p(\mathbf{y}|\boldsymbol{\theta}, \mathcal{M}) p(\boldsymbol{\theta}|\mathcal{M}) d\boldsymbol{\theta}.$$

The ratio of (5) to the corresponding quantity for an alternative model defines the *Bayes factor* for use in model choice. (See Berger, 1985, Koop, 2003, Geweke, 2005, and Robert, 2007, for textbook expositions).

The key point to note is that analytical solutions to (4) and (5) are usually unavailable. Indeed, Bayes’ original problem highlights that a solution to (4) can elude us *even* when the posterior pdf itself has a closed form. Typically, the posterior is known only up to a constant of proportionality, as

$$(6) \quad p(\boldsymbol{\theta}|\mathbf{y}) \propto p(\mathbf{y}|\boldsymbol{\theta})p(\boldsymbol{\theta}),$$

exceptions to this including when $p(\mathbf{y}|\boldsymbol{\theta})$ is from the exponential family, and either a natural conjugate, or convenient noninformative prior is adopted (as in Bayes’ problem). Knowledge

¹Bayes cast his problem in physical terms: as one in which balls were rolled across a square table, or plane. Over time his pictorial representation of the problem has come to be viewed as a ‘billiard table’, despite Bayes making no overt reference to such an item in his essay. For this, and other historical anecdotes, see Stigler (1986a) and Fienberg (2006).

of $p(\theta|\mathbf{y})$ only up to the integrating constant implies a lack of closed-form solution for (4), no matter what the form of $g(\theta)$. A lack of knowledge of the integrating constant automatically implies that the marginal likelihood for the assumed model in (5) is unavailable. Situations where the likelihood function itself does not have a closed form render the analytical solution of (4) and (5) an even more distant dream. In short, the implementation of all forms of Bayesian analysis relies heavily on numerical computation.

2.2 The Structure of this Review

In writing this review we have made two key decisions: *i)* to describe all methods using a common notation; and *ii)* to place the evolution of computational methods in a historical context, making use of pictorial timelines in the process. In so doing, we are able to present a coherent chronological narrative about Bayesian computation up to the present day. Specifically, all methods can be seen to be, in essence, attempting to compute integrals like (4) and (5); the use of a common notation makes that clear. However, important details of those integrals have changed over time: the dimension of θ (i.e. the number of ‘unknowns’), the dimension of \mathbf{y} (i.e. the ‘size’ of the data), the nature of the integrand itself and, most critically, the available computing technology. Computation has evolved accordingly, and the chronological ordering helps make sense of that evolution. Hence, whilst computational methods can — and often are — grouped according to the category into which they fall, e.g. *deterministic*, *simulation-based*, *approximate* etc., the over-arching structure that we adopt here is one of chronology, as understanding *when* a computational method has appeared aids in the appreciation of *why*, and *how*.

We begin, in Section 3, by returning to Bayes’ integral in (1), briefly reiterating the nature of the particular computational problem it presented. We then use this as a springboard for pinpointing four particular points in time during the two centuries (or so) subsequent to 1763: 1774, 1953, 1964 and 1970. These time points correspond, in turn, to four publications — by Laplace, Metropolis *et al.*, Hammersley and Handscomb, and Hastings, respectively — in which computational methods that produce estimates of integrals like that of Bayes, were proposed. Whilst only the method of Laplace was explicitly set within the context of inverse probability (or Bayesian inference), all four methods of computing integrals can be viewed as harbingers of what was to come in Bayesian computation *per se*.

In Section 4, we look at Bayesian computation in the late 20th century, during which time the inexorable rise in the speed, and accessibility, of computers led to the pre-eminence of *simulation-based* computation. Whilst important new developments arose that exploited the principles of importance sampling (Kloek and van Dijk, 1978; Geweke, 1989; Gordon, Salmond and Smith, 1993), the computational engine was well and truly fuelled by Markov chain Monte Carlo (MCMC) algorithms, allied with the concept of ‘data augmentation’; with Besag (1974), Geman and Geman (1984), Tanner and Wong (1987) and Gelfand and Smith (1990) being seminal contributions.

Section 5 then provides a potted summary of a ‘second computational revolution’ in the 21st century — born in response to the increased complexity, and size, of the empirical problems being analyzed via Bayesian means. We begin by briefly noting the important modifications and refinements of MCMC that have occurred since its initial appearance, including pseudo-marginal methods, Hamiltonian up-dates, adaptive sampling, and coupling; plus highlighting the renaissance of importance sampling that has occurred under the auspices of sequential Monte Carlo (SMC) methods. We then very briefly outline the main *approximate* methods that evolved to tackle intractable problems: approximate Bayesian computation (ABC), Bayesian synthetic likelihood (BSL), variational Bayes (VB) and integrated nested Laplace approximation (INLA).

Throughout the paper we explicitly reference only computational solutions to the posterior expectation in (4) which, of course, covers the use of computation in Bayesian prediction, via the appropriate choice of $g(\theta)$. For coverage of the use of computational methods to solve the prior expectation in (5) we refer the reader to Ardia *et al.* (2012) and Llorente *et al.* (2021).

Finally, in order to provide the reader with a quick ‘snapshot’ of the key personages involved in the evolution of computation (and the relevant seminal publications), Figures 1, 2

and 3 provide pictorial timelines coinciding approximately with the time periods covered in Sections 3, 4 and 5, respectively.

3. SOME EARLY CHRONOLOGICAL SIGNPOSTS: 1763-1970




FIG 1. From left to right: 1) The only known portrait (albeit unauthenticated) of the Reverend Thomas Bayes (1702-1761); 2) A portrait of Pierre-Simon Laplace (1749-1827); 3) A snapshot of the game of Solitaire. Stanislaw Ulam speculated that the proportion of wins in repeated random games played by a computer could be used as an estimate of the probability of winning the game — thereby giving birth to the idea of estimation by Monte Carlo simulation; 4) A diagram illustrating the ‘collisions of rigid spheres’ used in the description of the behaviour of interacting particles in *Metropolis et al. (1953)*; 5) Equivalent representations of the expectation of a function of a random variable x , $f(x)$. This equivalence underlies the use of repeated draws of x from $q(x)$ to produce an ‘importance sampling’ estimate of $\mathbb{E}[f(x)]$; 6) The quantity used in *Hastings (1970)* to define a Markov chain, based on the transition matrix $\mathbf{Q} = \{q_{ij}\}$, with invariant distribution $\boldsymbol{\pi} = (\pi_0, \pi_1, \dots, \pi_S)$ over states $0, 1, \dots, S$.

3.1 1763: Bayes’ Integral

Bayes’ desire was to evaluate the probability in (1). As noted above, for either $a \neq 0$ or $b \neq 1$, this required evaluation of the incomplete Beta function. For either x or $(n - x)$ small, Bayes proposed a Binomial expansion and term-by-term integration to give an exact solution (his ‘Rule 1’). However, for x and $(n - x)$ both large, this approach was infeasible: prompting Bayes (and, subsequently, Price himself; *Price, 1764*) to resort to producing upper and lower bounds for (1) using quadrature. Indeed, Stigler (1986a) speculates that the inability to produce an approximation to (1) that was sufficiently accurate may explain Bayes’ reluctance to publish his work and, perhaps, the lack of attention it received subsequent to its (posthumous) presentation by Price in 1763 and publication the following year in *Bayes (1764)*.²

Whilst the integral that Bayes wished to compute was a very particular one, it was representative of the general hurdle that needed to be overcome if the principle of inverse probability were to be a useful practical tool. In brief, inference about θ was expressed in probabilistic terms and, hence, required either the direct computation of probability intervals, or the computation of distributional moments of some sort. Ironically, the choice of the Bernoulli

²On November 10, 1763, Price sent an edited and annotated version of Bayes’ essay to the Secretary of the Royal Society, with his own Appendix added. Price read the essay to the Society on December 23, as noted earlier. The essay and appendix were subsequently published in 1764, in the *Philosophical Transactions of the Royal Society of London*. The front matter of the issue appears here: <https://royalsocietypublishing.org/cms/asset/f005dd95-c0f8-45b2-8347-0296a93c4272/front.pdf>. The publication has been reprinted since, including in *Barnard and Bayes (1958)*, with a biographical note by G.A. Barnard. Further historical detail on the important role played by Price in the dissemination of Bayes’ ideas can be found in *Hooper (2013)* and *Stigler (2018)*. As the submission of Bayes’ essay by Price, and his presentation to the Royal Society occurred in 1763, and Volume 53 of the *Philosophical Transactions* in which the essay appears is ‘For the Year 1763’, Bayes’ essay is often dated 1763. We follow *Stigler (1986a)* in using the actual publication date of 1764.

model, possibly the simplest process for generating data ‘forward’ (conditional on θ) that Bayes could have assumed, exacerbated this problem, given that the ‘inversion’ problem does not possess the simplicity of the generative problem. What was required was a solution that was, in large measure, workable no matter what the nature of the generative model, and the first solution came via the 1774 ‘*Mémoire sur la probabilité des causes par les événements*’ by Pierre Laplace.

3.2 1774: Laplace and His Method of Asymptotic Approximation

Laplace envisaged an experiment in which n tickets were drawn with replacement from an urn containing a given proportion of white and black tickets. Recasting his analysis in our notation, θ is the probability of drawing a white ticket, $\mathbf{y} = (y_1, y_2, \dots, y_n)'$ denotes the sequence of observed white tickets ($Y = 1$) and black tickets ($Y = 0$) associated with n independent draws of the random variable $Y|\theta \sim \text{Bernoulli}(\theta)$, and $x = \sum_{i=1}^n y_i$ is the number of white tickets drawn. Laplace’s aim was to show that, for arbitrary w : $\mathbb{P}(|\frac{x}{n} - \theta| < w|\mathbf{y}) = \mathbb{P}(\frac{x}{n} - w < \theta < \frac{x}{n} + w|\mathbf{y}) \rightarrow 1$ as $n \rightarrow \infty$. That is, Laplace wished to demonstrate *posterior consistency*: concentration of the posterior onto the true proportion of white tickets in the urn, $\theta_0 = \lim_{n \rightarrow \infty} \frac{x}{n}$. Along the way, however, he stumbled upon the same problem as had Bayes: computing the following probability of a Beta random variable,

$$(7) \quad \mathbb{P}(a < \theta < b|\mathbf{y}) = \frac{\int_a^b \theta^x (1 - \theta)^{n-x} d\theta}{B(x+1, n-x+1)},$$

with $a = \frac{x}{n} - w \neq 0$ and $b = \frac{x}{n} + w \neq 1$. Laplace’s genius (allied with the power of asymptotics!) was to recognize that the exponential of the integrand in (7) has the bulk of its mass in the region of its mode, as n gets large, and that the integral can be computed in closed form in this case. This enabled him to prove (in modern notation) that $\mathbb{P}(|\theta - \theta_0| > w|\mathbf{y}) = o_p(1)$, where p denotes the probability law of $Y|\theta_0$.

The route he took to this proof, however, involved approximating the Beta posterior with a Normal distribution, which (under regularity) is an approach that can be used to provide a large sample approximation of virtually *any* posterior probability. Specifically, expressing an arbitrary posterior probability as

$$(8) \quad \mathbb{P}(a < \theta < b|\mathbf{y}) = \int_a^b p(\theta|\mathbf{y}) d\theta = \int_a^b e^{nf(\theta)} d\theta,$$

where $f(\theta) = \log[p(\theta|\mathbf{y})]/n$, a second-order Taylor series approximation of $f(\theta)$ around its mode, $\hat{\theta}$, yields the *Laplace asymptotic approximation*

$$(9) \quad \mathbb{P}(a < \theta < b|\mathbf{y}) \approx e^{nf(\hat{\theta})} \sqrt{2\pi\sigma^2} \{\Phi[\frac{b-\hat{\theta}}{\sigma}] - \Phi[\frac{a-\hat{\theta}}{\sigma}]\},$$

where $\sigma^2 = -[nf''(\hat{\theta})]^{-1}$ and $\Phi(\cdot)$ denotes the standard Normal cumulative distribution function.³

With (9), Laplace had thus devised a general way of implementing inverse probability: probabilistic statements about an unknown parameter, θ , conditional on data generated from any (regular) model, could now be made, at least up to an error of approximation. Whilst his focus was solely on the computation of a specific posterior probability, and in a single parameter setting, his method was eventually used to approximate general posterior expectations of the form in (4) (Lindley, 1980; Tierney and Kadane, 1986; Tierney, Kass and Kadane, 1989) and, indeed, applied as an integral approximation method in its own right (De Bruijn, 1961).

³See Tierney and Kadane (1986) and Robert and Casella (2004) for further elaboration; and Ghosal, Ghosh and Samanta (1995) and van der Vaart (1998) for more formal demonstrations of the conditions under which a posterior distribution converges in probability to a Normal distribution, and the so-called Bernstein-von Mises theorem — the modern day version of Laplace’s approximation — holds.

The approach also underpins the modern INLA technique to be mentioned in Section 5.2 (Rue, Martino and Chopin, 2009).⁴

Meanwhile, it would take 170-odd years for the *next* major advance in the computation of probability integrals to occur; an advance that would eventually transform the way in which problems in inverse probability could be tackled. This development was based on a new form of thinking and, critically, required a platform on which such thinking could operate: namely, machines that could *simulate* repeated random draws of θ from $p(\theta|\mathbf{y})$, or from some representation thereof. Given a sufficient number of such draws, and the correct use of them, an estimate of (4) could be produced that — unlike the Laplace approximation — would be accurate for any sample size, n , and would require less analytical input. This potential to accurately estimate (4) for essentially any problem, and any given sample size, was the catalyst for a flourishing of Bayesian inference in the late 20th century and beyond. The 1953 publication in the *Journal of Chemical Physics* by Metropolis, Rosenbluth, Rosenbluth, Teller and Teller: ‘*Equation of State Calculations by Fast Computing Machines*’, was a first major step in this journey.⁵

3.3 1953: Monte Carlo Simulation and the Metropolis Algorithm

The convergence of the idea of simulating random draws from a probability distribution, and the production of such draws by computing machines, occurred in the scientific hothouse of the Los Alamos Laboratory, New Mexico, in the 1940s and 1950s; the primary impetus being the need to simulate physical processes, including neutrons in the fissile material in atomic bombs. We refer the reader to Liu (2001), Hitchcock (2003), Gubernatis (2005) and Robert and Casella (2011) for reviews of this period, including details of the various personalities who played a role therein.⁶ Our focus here is simply on the nature of the problem that was at the heart of Metropolis et al. (1953), the solution proposed, and the ultimate importance of that solution to Bayesian computation.

In short, the authors wished to compute an expectation of the form,

$$(10) \quad \mathbb{E}(g(\mathbf{x})) = \int_{\mathcal{X}} g(\mathbf{x})p(\mathbf{x})d\mathbf{x},$$

where $p(\mathbf{x})$ denotes the so-called Boltzmann distribution of a set, \mathbf{x} , of N particles on \mathbb{R}^2 . (See Robert and Casella, 2011, Section 2.1, for all details.) Two particular characteristics of (10) are relevant to us here: *i*) the integral is of very high dimension, $2N$, with N large; and *ii*) $p(\mathbf{x})$ is generally known only up to its integrating constant. The implication of *i*) is

⁴Stigler (1975, Section 2) states that he has found no documentary evidence that Laplace’s ideas on inverse probability, as presented in the 1774 publication, including his own statement of ‘Bayes’ theorem’ in (2), were informed by Bayes’ earlier ideas. See Stigler (1986a, Chapter 3) for discussion of Laplace’s later extensions of Bayes’ theorem to the case of a non-Uniform prior, and see Stigler (1975), Stigler (1986a), Stigler (1986b) and Fienberg (2006) on matters of attribution. The first recorded reference to Bayes’ prior claim to inverse probability is in the preface, written by Condorcet, to Laplace’s later 1781 publication: ‘*Mémoire sur les probabilités*’.

⁵With reference to the mechanical simulation of a random variable, we acknowledge the earlier 1870s’ invention of the *quincunx* by Francis Galton. This machine used the random dispersion of metal shot to illustrate (amongst other things) draws from a hierarchical Normal model and regression to the mean. Its use can thus be viewed as the first illustration of the conjugation of a Normal likelihood and a Normal prior. See Stigler (1986a) for more details, including Galton’s graphical illustration of his machine in a letter to his cousin (and Charles Darwin’s son), George Darwin.

⁶We make particular mention here of John and Klara von Neumann, and Stanislaw Ulam, with the latter co-authoring the 1949 publication in the *Journal of the American Statistical Association*: ‘*The Monte Carlo Method*’ with Nicholas Metropolis. We also note the controversy concerning the respective contributions of the five authors of the 1953 paper (who included two married couples). On this particular point, we refer the reader to the informative 2005 article by Gubernatis, in which Marshall Rosenbluth gives a bird’s eye account of who did what, and when. The article brings to light the important roles played by both Adriana Rosenbluth and Mici Teller. Finally, we alert the reader to the proceedings of a symposium on the Monte Carlo method held at the University of California in July, 1949, edited by A.S. Householder of the Oak Ridge National Laboratory, Tennessee. Researchers at the Oak Ridge laboratory were also involved in the early exploration of Monte Carlo methods.

that a basic rectangular integration method, based on L grid-points in each of the $2N$ directions, is infeasible, having a computational burden of L^{2N} or, equivalently, an approximation error of $O(L^{-1/2N})$ (Kloek and van Dijk, 1978). The implication of *ii*) is that a Monte Carlo (MC) estimate of (10), based on M *i.i.d.* direct draws from $p(\mathbf{x})$, $\mathbf{x}^{(i)}$, $i = 1, 2, \dots, M$: $\hat{E}_{MC}(g(\mathbf{x})) = \frac{1}{M} \sum_{i=1}^M g(\mathbf{x}^{(i)})$, with approximation error of $O(M^{-1/2})$ independent of dimension, is not available.⁷

Features *i*) and *ii*) — either individually or in tandem — broadly characterize the posterior expectations in (4) that are the focus of this review. Hence the relevance to Bayesian computation of the solution offered by Metropolis et al. (1953) to the non-Bayesian problem in (10); and we describe their solution with direct reference to (4) and the notation used therein.

Specifically, the authors advocate computing an integral such as (4) via the simulation of a *Markov chain*: $\theta^{(i)}$, $i = 1, 2, \dots, M$, with *invariant distribution* $p(\theta|\mathbf{y})$. The draw at iteration $i + 1$ in the chain is created by taking the value at the i th iteration, $\theta^{(i)}$, and perturbing it according to a random walk: $\theta^c = \theta^{(i)} + \delta\epsilon$, where each element of ϵ is drawn independently from $U(-1, 1)$, and δ ‘tunes’ the algorithm.⁸ The ‘candidate’ draw θ^c is accepted as draw $\theta^{(i+1)}$ with probability:

$$(11) \quad \alpha = \min\{p^*(\theta^c|\mathbf{y})/p^*(\theta^{(i)}|\mathbf{y}), 1\},$$

where p^* is a kernel of p . Using the theory of reversible Markov chains, it can be shown (see, for example, Tierney, 1994) that use of (11) to determine the $(i + 1)$ th value in the chain does indeed produce a dependent sequence of draws with invariant distribution $p(\theta|\mathbf{y})$. Hence, subject to convergence to $p(\theta|\mathbf{y})$ (conditions for which were verified by the authors for their particular problem) these draws can be used to estimate (4) as the sample mean,

$$(12) \quad \overline{g(\theta)} = \frac{1}{M} \sum_{i=1}^M g(\theta^{(i)}),$$

and an appropriate weak law of large numbers and central limit theorem invoked to prove the \sqrt{M} -consistency and limiting normality of the estimator. (See Geyer, 2011a, for details.)

Due to the (positive) autocorrelation in the Markov chain, the variance of the *Metropolis estimator* (as it would become known) is larger than that of the (infeasible) MC estimate in (11), computed as in (12), but using *i.i.d.* draws from $p(\theta|\mathbf{y})$, namely:

$$(13) \quad \sigma_{MC}^2 = \text{Var}(g(\theta))/M,$$

expressed here for the case of scalar $g(\theta)$. However, as is clear from (11), the Metropolis MCMC algorithm requires knowledge of $p(\theta|\mathbf{y})$ only up to the normalizing constant, and does *not* require direct simulation from $p(\theta|\mathbf{y})$ itself. It is this particular feature that would lend the technique its great power in the decades to come.⁹

3.4 1964: Hammersley and Handscomb: Importance Sampling

The obviation of the need to directly sample from $p(\theta|\mathbf{y})$ also characterizes importance sampling, and underlies its eventual importance in solving difficult Bayesian computational problems. Nevertheless, Hammersley and Handscomb (1964) did not emphasize this feature but, rather, focussed on the ability of IS to produce variance reduction in simulation-based

⁷The authors actually make mention of a *naïve* Monte Carlo method, based on *Uniform* sampling over the $2N$ dimensional space, followed by a reweighting of the Uniform draws by a kernel of $p(\mathbf{x})$. The idea is dismissed, however, as ‘not practical’. In modern parlance, whilst this method would yield an $O(M^{-1/2})$ approximation error, the constant term within the order would be large, since the Uniform distribution used to produce draws of \mathbf{x} differs substantially from the *actual* distribution of \mathbf{x} , $p(\mathbf{x})$.

⁸Metropolis et al. (1953) actually implemented their algorithm one element of θ at a time, as a harbinger of the Gibbs sampler to come. See Robert and Casella (2011) for more details.

⁹Dongarra and Sullivan (2000) rank the Metropolis algorithm as one of the 10 algorithms “with the greatest influence on the development and practice of science and engineering in the 20th century”.

estimation of integrals.¹⁰ Again, the focus was not on Bayesian integrals, but we describe the method in that setting.¹¹

In brief, given an ‘importance’ (or ‘proposal’) density, $q(\boldsymbol{\theta}|\mathbf{y})$, that preferably mimics $p(\boldsymbol{\theta}|\mathbf{y})$ well, and M *i.i.d.* draws, $\boldsymbol{\theta}^{(i)}$, from $q(\boldsymbol{\theta}|\mathbf{y})$, an IS estimate of (4) is $\overline{g(\boldsymbol{\theta})}_{IS} = \frac{1}{M} \sum_{i=1}^M g(\boldsymbol{\theta}^{(i)})w(\boldsymbol{\theta}^{(i)})$, where $w(\boldsymbol{\theta}^{(i)}) = p(\boldsymbol{\theta}^{(i)}|\mathbf{y})/q(\boldsymbol{\theta}^{(i)}|\mathbf{y})$. In the typical case where $p(\boldsymbol{\theta}^{(i)}|\mathbf{y})$ is available only up to the integrating constant, and $w(\boldsymbol{\theta}^{(i)})$ cannot be evaluated as a consequence, the estimate is modified as

$$(14) \quad \overline{g(\boldsymbol{\theta})}_{IS} = \sum_{i=1}^M g(\boldsymbol{\theta}^{(i)})w(\boldsymbol{\theta}^{(i)}) / \sum_{j=1}^M w(\boldsymbol{\theta}^{(j)}),$$

with the weights re-defined as $w(\boldsymbol{\theta}^{(j)}) = p^*(\boldsymbol{\theta}^{(i)}|\mathbf{y})/q^*(\boldsymbol{\theta}^{(i)}|\mathbf{y})$, for kernels, $p^*(\boldsymbol{\theta}^{(i)}|\mathbf{y})$ and $q^*(\boldsymbol{\theta}^{(i)}|\mathbf{y})$, of $p(\boldsymbol{\theta}|\mathbf{y})$ and $q(\boldsymbol{\theta}|\mathbf{y})$ respectively. Once again, and under regularity conditions pertaining to the *importance density* $q(\boldsymbol{\theta}|\mathbf{y})$, asymptotic theory can be invoked to prove that (14) is a \sqrt{M} -consistent estimator of $\mathbb{E}(g(\boldsymbol{\theta})|\mathbf{y})$ (Geweke, 1989). A judicious choice of $q(\boldsymbol{\theta}|\mathbf{y})$ is able to yield a sampling variance that is less than (13) in some cases, as befits the original motivation of IS as a variance reduction method. (See Geweke, 1989, and Robert and Casella, 2004, for discussion.) Critically however, like the Metropolis method, (14) serves as a feasible estimate of $\mathbb{E}(g(\boldsymbol{\theta})|\mathbf{y})$ when $p(\boldsymbol{\theta}|\mathbf{y})$ cannot be easily simulated; hence the significance of IS in Bayesian computation. Moreover, its maintenance of independent draws, allied with its re-weighting of draws from an approximating density, has led to the emergence of IS as a vehicle for implementing SMC algorithms, to be referenced in Section 5.1.

3.5 1970: Hastings and his Generalization of the Metropolis Algorithm

The final publication that we pinpoint during the 200-odd year period subsequent to 1763, is the 1970 *Biometrika* paper, ‘*Monte Carlo Sampling Methods Using Markov Chains and Their Applications*’, by Wilfred Keith Hastings. Whilst Metropolis et al. (1953) proposed the use of MCMC sampling to compute particular integrals in statistical mechanics, it was the Hastings paper that elevated the concept to a general one, and introduced it to the broader statistics community. Included in the paper is also the first mention of what would become known as the *Metropolis-within-Gibbs sampler* (Robert and Casella, 2011). Once again, the author’s focus was not a Bayesian integral *per se*; however we describe the method in that context.

In contrast to Metropolis and co-authors, Hastings (1970) acknowledges up-front that the need to know $p(\boldsymbol{\theta}|\mathbf{y})$ only up to the integrating constant is a compelling feature of an MCMC-based estimate of (4). Hastings also generalizes the acceptance probability in (11) to one that accommodates a general ‘candidate’ distribution $q(\boldsymbol{\theta}|\mathbf{y})$ from which $\boldsymbol{\theta}^c$ is drawn, as:

$$(15) \quad \alpha = \frac{p^*(\boldsymbol{\theta}^c|\mathbf{y})/q(\boldsymbol{\theta}^{(i)}|\boldsymbol{\theta}^c, \mathbf{y})}{p^*(\boldsymbol{\theta}^{(i)}|\mathbf{y})/q(\boldsymbol{\theta}^c|\boldsymbol{\theta}^{(i)}, \mathbf{y})} \wedge 1$$

which clearly collapses to (11) when $q(\boldsymbol{\theta}|\mathbf{y})$ is symmetric (in $\boldsymbol{\theta}^c$ and $\boldsymbol{\theta}^{(i)}$), as in the original random walk proposal of Metropolis et al. (1953). Importantly, the more general algorithm allows for a targeted choice of $q(\boldsymbol{\theta}|\mathbf{y})$ that reduces the need for tuning and which can, potentially, reduce the degree of dependence in the chain and, hence, the variance of the estimate of $\mathbb{E}(g(\boldsymbol{\theta})|\mathbf{y})$. Hastings formalizes the standard error of this estimate using time series theory,

¹⁰One could in fact argue that a similar aim motivated Metropolis and co-authors, given that they drew a sharp contrast (in effect) between the efficiency of their method and that of the naïve Monte Carlo technique based on Uniform sampling.

¹¹Whilst Hammersley and Handscomb (1964) provides a textbook exposition of importance sampling, the concept actually appeared in published form much earlier, initially under the name of ‘quota sampling’. See Kahn (1949) and Goertzel and Kahn (1950) for example, plus Andral (2022) for a recent exploration into the historical origins of importance sampling.

explicitly linking, for the first time, the autocorrelation in the Markov draws to the efficiency of the MCMC-based estimate of (4). Crucially, the author tackles the issue of dimension by advocating the treatment of one element (or several elements) of θ at a time, conditional on all remaining elements.

In summary, all of the important ingredients from which the huge smorgasbord of future MCMC algorithms would eventually be constructed — for the *express* purpose of solving Bayesian problems — were now on the table, via this particular paper.

4. THE LATE 20TH CENTURY: GIBBS SAMPLING & THE MCMC REVOLUTION




FIG 2. From left to right: 1) Equation (3.3) in [Kloek and van Dijk \(1978\)](#), where $M(\theta) = g(\theta)\kappa(\theta|\mathbf{Y}, \mathbf{Z})p(\theta)/I(\theta)$, with $\kappa(\theta|\mathbf{Y}, \mathbf{Z})$ the likelihood function, \mathbf{Y} and \mathbf{Z} matrices of observations on endogenous and predetermined variables respectively, $p(\theta)$ the prior, $g(\theta)$ defining the parameter function of interest, and $I(\theta)$ the importance density. The first application of importance sampling to compute a posterior expectation; 2) Figure 2 in [Geman and Geman \(1984\)](#); Gibbs sampling is used to restore image (d) from the degraded image (b); 3) Figure 1 in [Tierney and Kadane \(1986\)](#), illustrating the accuracy of a Laplace approximation of the marginal posterior for the parameter ρ ; 4) The exploitation of ‘data augmentation’ to estimate the marginal posterior of θ as an average of conditional posteriors given m draws of the latent state vector, \mathbf{z} ; Page 530 in [Tanner and Wong \(1987\)](#); 5) The sequence of conditional distributions that define the Gibbs sampler in [Gelfand and Smith \(1990\)](#); 6) An illustration from Figure 2 in [Carter and Kohn \(1994\)](#) of the rapid decline in the sample autocorrelation function for draws from their multi-state Gibbs sampling algorithm; proposed independently by [Frühwirth-Schnatter \(1994\)](#) as the forward-filtering-backward-sampling (‘FFBS’) algorithm; 7) A diagrammatic representation of the slice sampler in Figure 1 of [Roberts and Rosenthal \(1999\)](#).

Whilst the role that *could* be played by simulation in computation was thus known by the 1970s, the computing technology needed to exploit that knowledge lagged behind.¹² Over the next two decades, however, things changed. Indeed, *two* developments now went hand in hand to spawn a remarkable expansion in simulation-based Bayesian computation: *i*) the increased speed and availability of computers, including personal desktop computers ([Ceruzzi, 2003](#)), and *ii*) the collective recognition that MCMC draws from a joint posterior, $p(\theta|\mathbf{y})$, could be produced via iterative sampling from lower dimensional, and often standard, *conditional* posteriors. When allied with both the concept of *augmentation*, and an understanding of the theoretical properties of combinations of MCMC algorithms, *ii*) would lead to *Gibbs sampling* (with or without *Metropolis-Hastings* (MH) subchains) becoming the work-horse of Bayesian computation in the 1990s.

An MH algorithm ‘works’, in the sense of producing a Markov chain that converges to the required distribution $p(\theta|\mathbf{y})$, due to the form of the acceptance probability in (15) (or the

¹²Many readers may be too young to remember the punchcards! But there was a time when RAND’s 1955 *A Million Random Digits with 100,000 Normal Deviates* was more than an entry for sarcastic Amazon comments, as producing this million digits took more than two months at the time.

nested version in (11)). More formally, the algorithm, as based on candidate density $q(\boldsymbol{\theta}|\mathbf{y})$, and acceptance probability as defined in (15), defines a *transition kernel* with invariant distribution, $p(\boldsymbol{\theta}|\mathbf{y})$. The ‘Gibbs sampler’ similarly yields a Markov chain with invariant distribution, $p(\boldsymbol{\theta}|\mathbf{y})$, but via a transition kernel that is defined as the product of full conditional posteriors associated with the joint. For the simplest case of a two-dimensional vector $\boldsymbol{\theta} = (\theta_1, \theta_2)'$, the steps of the Gibbs algorithm are as follows: first, specify an initial value for $\theta_2, \theta_2^{(0)}$; second, for $i = 1, 2, \dots, M$, cycle iteratively through the two conditional distributions, drawing respectively: $\theta_1^{(i)}$ from $p_1(\theta_1^{(i)}|\theta_2^{(i-1)}, \mathbf{y})$, and $\theta_2^{(i)}$ from $p_2(\theta_2^{(i)}|\theta_1^{(i)}, \mathbf{y})$. Given the satisfaction of the required convergence conditions (which essentially place sufficient regularity on the conditionals), the draws $\boldsymbol{\theta}^{(i)} = (\theta_1^{(i)}, \theta_2^{(i)})'$, $i = 1, 2, \dots, M$, converge in distribution to the joint posterior distribution as $M \rightarrow \infty$, and can be used to produce a \sqrt{M} -consistent estimator of (4) in the form of (12). Extension to higher-dimensional problems is obvious, although decisions about how to ‘block’ the parameters, and thereby define the conditionals, now play a role (Roberts and Sahu, 1997).¹³

Gibbs thus exploits the simplicity yielded by conditioning: whilst joint and marginal posterior distributions are usually complex in form, (full) conditional posteriors are often standard and, hence, able to be simulated from directly. While one may find hints in both Hastings (1970) and Besag (1974), this point was first made clearly by Geman and Geman (1984), who also coined the phrase ‘Gibbs sampling’ because their problem used Gibbs random fields in image restoration (named, in turn, after the physicist, Josiah Willard Gibbs). However, the later paper by Gelfand and Smith (1990) is generally credited with bringing this transformational idea to the attention of the statistical community, and illustrating its broad applicability.

The idea of Gibbs sampling overlapped with a related proposal by Tanner and Wong (1987): that of ‘augmenting’ the set of unknowns ($\boldsymbol{\theta}$ in our notation) with latent data, $\mathbf{x} = (x_1, x_2, \dots, x_n)'$, to yield conditionals — $p(\boldsymbol{\theta}|\mathbf{x}, \mathbf{y})$ and $p(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})$ — that facilitate the production of a simulation-based estimate of $p(\boldsymbol{\theta}|\mathbf{y})$; with $p(\boldsymbol{\theta}|\mathbf{x}, \mathbf{y})$, in particular, often being standard. The melding of these two ideas, i.e. sampling via conditionals *per se*, and yielding more tractable conditionals through the process of augmentation, enabled the analysis of complex models that had thus far eluded Bayesian treatment, due to their dependence on high-dimensional vectors of latent variables; selected examples being: Polson, Carlin and Stoffer (1992), Carter and Kohn (1994), Frühwirth-Schnatter (1994) and Jacquier, Polson and Rossi (1994). However, it also led to the realization that *auxiliary* latent variables could be judiciously introduced into a model for the sole purpose of producing tractable conditional posteriors over the augmented space, thereby opening up a whole range of additional models to a Gibbs-based solution (e.g. Albert and Chib, 1993; Diebolt and Robert, 1994; Higdon, 1998; Kim, Shephard and Chib, 1998; Damien, Wakefield and Walker, 1999). The *slice sampler* (Roberts and Rosenthal, 1999; Neal, 2003) is one particularly notable, and generic, way of generating an MCMC algorithm via this principle of auxiliary variable augmentation.

Of course, in most high-dimensional problems — and in particular those in which latent variables feature — certain conditionals remain nonstandard, such that direct simulation from them is not possible. Critically though, the reduced dimension renders this a simpler problem than sampling from the joint itself: via either the inverse cumulative distribution function technique (Devroye, 1986) — approximated in the ‘Griddy Gibbs’ algorithm of Ritter and Tanner (1992) — or by embedding an MH algorithm within the outer Gibbs loop (a so-called ‘Metropolis-within-Gibbs’ algorithm).¹⁴

¹³The Gibbs sampler can be viewed as (and in some expositions is presented as) a special case of a ‘multiple-block’ MH sampler, in which the candidate values for each block of parameters are drawn directly from their full conditional distributions and the acceptance probability in (each blocked version of) (15) is equal to one. (See, for example, Chib, 2011). See also Tran (2018) for further discussion of this point — conducted in the context of a *generalized* MH framework, in which many of the MCMC algorithms mentioned in Sections 5.1.1 are also nested.

¹⁴We refer the reader to: Besag and Green (1993), Smith and Roberts (1993) and Chib and Greenberg (1996) for early reviews of MCMC sampling; Casella and George (1992) and Chib and Greenberg (1995) for descriptions of the Gibbs and MH algorithms (respectively) that are useful for practitioners; Robert (2015), Betancourt (2018) and Dunson and Johndrow (2019) for more recent reviews; and Andrieu, Doucet and Robert (2004) and Robert and Casella (2011) for historical accounts of MCMC sampling.

5. THE 21ST CENTURY: A SECOND COMPUTATIONAL REVOLUTION!

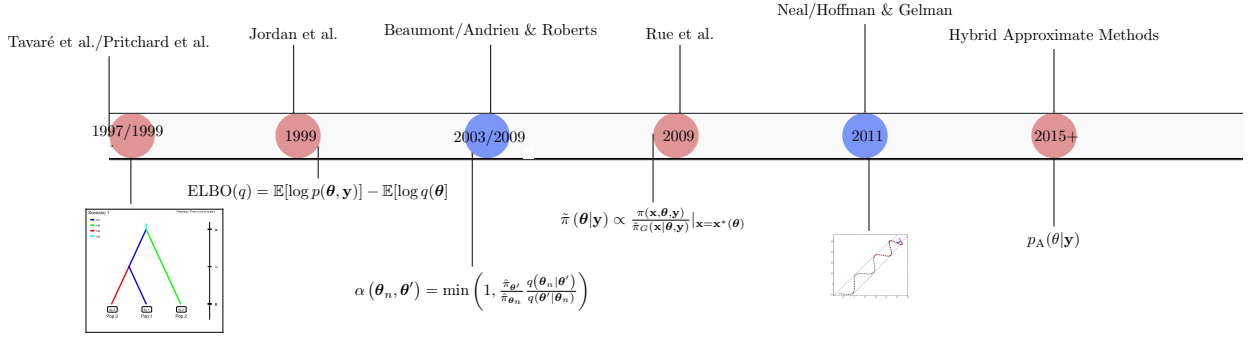


FIG 3. From left to right: 1) An example of a phylogenetic tree, which represents one of the earliest types of model to which approximate Bayesian computation (ABC) was applied; Tavaré et al. (1997) and Pritchard et al. (1999) were the first to apply the method of ABC to intractable models in population genetics; 2) The ‘evidence lower bound’ (ELBO), maximized with respect to the variational density, $q(\theta)$, in order to produce a variational Bayes approximation to the posterior $p(\theta|y)$. Jordan, Ghahramani and Jaakkola (1999) were the first to apply the method to high-dimensional graphical models; 3) The acceptance probability used in a pseudo-marginal Metropolis Hastings algorithm, as based on an estimate of the likelihood function defining the kernel of the posterior, π , and a proposal distribution q ; 4) The approximation of the posterior $\pi(\theta|y)$ based on integrated nested Laplace approximation (INLA), with $\hat{\pi}_G(\mathbf{x}|\theta, y)$ a Gaussian approximation of $\pi(\mathbf{x}|\theta, y)$ and $\mathbf{x}^*(\theta)$ the mode of $p(\mathbf{x}, \theta, y)$ (for a given θ); Equation (3) in the seminal INLA paper of Rue, Martino and Chopin (2009); 5) A trajectory generated during one iteration of the No-U-Turn Sampler; an extension of Hamiltonian Monte Carlo; Figure 2 in Hoffman and Gelman (2014); 6) The myriad of methods for producing an approximate posterior, $p_A(\theta|y)$, that are ‘hybrids’ of various individual approximate methods, are reviewed in Section 3.3 of Martin, Frazier and Robert (2021).

All of the methodological and theoretical developments referenced above, allied with advances in computer hardware and software, spawned a huge flowering of simulation-based Bayesian computation — across all disciplines — in the final decades of the 20th century, with variants of MCMC leading the charge. (See, for example, Brooks et al., 2011, and Geweke, Koop and van Dijk, 2011, for extensive coverages of the empirical problems to which MCMC algorithms (in particular) have been applied.) However, the early algorithms did stumble in the case of so-called ‘intractable’ Bayesian problems, namely: 1) a DGP that cannot readily be expressed as a probability density or mass function (the ‘unavailable likelihood’ problem); 2) a very large dimension for θ (the ‘high-dimensional’ problem); and/or 3) a very large dimension for y (the ‘big-data’ problem). See Rue, Martino and Chopin (2009), Green et al. (2015), Bardenet, Doucet and Holmes (2017), Blei, Kucukelbir and McAuliffe (2017), Betancourt (2018), Robert et al. (2018), Johndrow et al. (2019), Jahan, Ullah and Mengersen (2020), and Martin, Frazier and Robert (2021) for relevant discussions.

In response, a wealth of solutions have been proposed. For the purpose of this paper it is convenient to categorize these solutions as either ‘exact’ or ‘approximate’. Exact solutions still invoke MCMC or IS principles to compute (4) in cases where one or more form of intractability obtains. That is, the goal of such methods is to still estimate the posterior expectation in (4) ‘exactly’, at least up to an order $O(M^{-1/2})$, where M is the number of draws that defines the simulation scheme, and which can — in principle — be made arbitrarily large.¹⁵ Approximate solutions, on the other hand, use computation to target only an *approximation* — of one sort or another — of the expectation in (4).

¹⁵We note here so-called ‘quasi-Monte Carlo’ integration schemes, which aim for exactness at a faster rate than $O(M^{-1/2})$. See Lemieux (2009) for a review of such methods, Chen, Dick and Owen (2011) for the extension to quasi-MCMC algorithms, and Gerber and Chopin (2015) for an entry on sequential quasi-Monte Carlo.

5.1 Exact Solutions to Intractable Bayesian Problems

5.1.1 *Advances in MCMC.* We begin with four reminders about MCMC algorithms:

1. *First:* MCMC methods avoid the need to simulate from a joint posterior directly by simulating the unknowns indirectly, via another distribution (or set of distributions) from which simulation is feasible. However, such methods still require the evaluation of the DGP as a probability density function or a probability mass function: either in the computation of the acceptance probability in (11) or (15) in a Metropolis-type algorithm, or in the implementation of any Gibbs-based algorithm, in which the conditional posteriors are required either in full form or at least up to a scale factor.
2. *Second:* an MCMC algorithm is just that — a *Markov chain Monte Carlo* algorithm. As such, an MCMC scheme — by design — produces a *local* exploration of the target posterior, with the location in the parameter space of any simulated draw being dependent on the location of the previous draw, in a manner that reflects the specific structure of the algorithm. Most notably, an MCMC algorithm with a high degree of dependence will potentially be slow in exploring the high mass region of the target posterior (or the ‘target set’, in the language of [Betancourt, 2018](#)), with this problem usually being more severe the larger is the dimension of the parameter space. Looked at through another lens: for M MCMC draws, the greater the degree of (typically positive) dependence in those draws, the less efficient is the MCMC-based estimate of (4), relative to an estimate based on M *i.i.d.* draws from the target. This loss of efficiency is measured by the so-called *inefficiency factor* (IF), defined (in the case of scalar $g(\theta)$) as the ratio of the *MCMC standard error*, σ_{MCMC} with $\sigma_{MCMC}^2 = \text{Var}(g(\theta))[1 + 2\sum_{l=1}^{\infty} \rho_l]/M$ to the standard error associated with M *i.i.d.* draws, σ_{MC}^2 , with σ_{MC}^2 as given in (13), where ρ_l is the lag- l autocorrelation of the draws of $g(\theta)$ over the history of the chain. This ratio, in turn, defines the *effective sample size* of the MCMC algorithm, $ESS = M/[1 + 2\sum_{l=1}^{\infty} \rho_l]$. Improving the efficiency of an MCMC algorithm, for any given value of M , thus equates to increasing ESS to its maximum possible value of M by reducing the dependence in the draws.
3. *Third:* an MCMC algorithm is also a *Markov chain Monte Carlo* algorithm. That is, under appropriate regularity it produces a \sqrt{M} -consistent estimate of (4), whatever the degree of dependence in the chain, with the dependence affecting the constant term implicit in the $O(M^{-1/2})$ rate of convergence, but not the rate itself. Hence, in principle, any MCMC algorithm, no matter how inherently inefficient, can produce an estimate of (4) that is arbitrarily accurate, simply through an increase in M . However, an increase in M entails an increase in *computational cost*, measured, say, by *computing clock-time*. The extent of this increase depends, in turn, on the (per-iteration) cost of generating a (proposal/candidate) draw and, with an MH step, the cost of calculating the acceptance probability. Both component costs will (for any algorithm) clearly increase with the number of unknowns that need to be simulated, and assessed, at each iteration. Either cost, or both, will also increase with the sample size, given the need for pointwise evaluation of the likelihood function across the elements of y .
4. *Fourth:* the very concept of efficiency is relevant only if the Markov chain is (asymptotically in M) *unbiased*, which depends critically on draws being produced from the correct invariant distribution. That is, the production of an accurate MCMC-based estimate of (4) depends, not just on reducing the degree of dependence in the chain, or on increasing the number of draws, but on ensuring that the chain *actually* explores the target set, and thereby avoids bias in the estimation of (4).¹⁶

¹⁶It is acknowledged in the literature that MCMC algorithms produce potentially strong biases in their initial phase of ‘convergence’ to the typical set from an initial point in the parameter space. However, under appropriate regularity, such biases are transient, and their impact on the estimation of (4) able to be eliminated by discarding a sufficiently large number of ‘burn-in’ or ‘warm-up’ draws from the computation. (See [Robert and Casella, 2004](#), and [Gelman and Shirley, 2011](#), for textbook discussions of convergence, including diagnostic methods.) Some of

With reference to Point 1: so-called *pseudo-marginal* MCMC methods can be used to obviate the problem of the DGP either being unavailable or being computationally challenging, by inserting within an MCMC algorithm an *unbiased* estimate of $p(\mathbf{y}|\boldsymbol{\theta})$, thereby retaining $p(\boldsymbol{\theta}|\mathbf{y})$ as the invariant distribution of the chain (Beaumont, 2003; Andrieu and Roberts, 2009). The use of an estimate of the likelihood in an MH algorithm prompted use of the term ‘pseudo-marginal’ MH (PMMH) by Andrieu and Roberts (2009) although, as noted, this replacement still yields a chain with an invariant distribution equal to $p(\boldsymbol{\theta}|\mathbf{y})$ when the estimate is unbiased; hence the method remains ‘exact’. When a likelihood estimate is produced specifically via the use of *particle filtering* in a state space model (SSM), the term *particle* MCMC has also been coined (Andrieu, Doucet and Holenstein, 2011).¹⁷

Whilst unbiasedness of the likelihood estimate is required for a general PMMH algorithm to ‘work’, the variance of the estimate also affects the performance of the sampler and, hence, the simulation efficiency of any estimate of (4) that is produced. However, improving the precision of the likelihood estimator comes at a computational cost, and an ‘optimal’ number of draws that balances computational cost with an acceptable mixing of the chain needs to be sought. See Pitt et al. (2012), Doucet et al. (2015) and Deligiannidis, Doucet and Pitt (2018) for discussion of the optimal structuring and tuning of pseudo-marginal algorithms.

With reference to Points 2 to 4: many other advances in MCMC (some of which also exploit pseudo-marginal principles) aim to increase the effectiveness with which an algorithm explores the high mass region of the target posterior and, hence, the accuracy with which (4) is estimated, by doing one (or more) of three things: reducing dependence in the chain, reducing the computational cost per iteration of the chain (thus enabling more draws to be produced), or eliminating bias. In particular, with reference to the taxonomy of ‘intractable’ problems given earlier, focus is increasingly directed towards developing algorithms that *scale well*, in terms of the dimension of the data and/or the number of unknowns. With our goal of brevity in mind, we simply list the main contenders here, including certain key references or reviews, deflecting both to those papers, and to the broad overviews of modern developments in MCMC in Green et al. (2015), Robert et al. (2018) and Dunson and Johndrow (2019) for all details. Of particular note is the recent survey on Bayesian methods for ‘Big Data’ in Jahan, Ullah and Mengersen (2020) (Section 5.1 being most pertinent), which describes the precise manner in which certain of the methods cited below (and others) tackle the problem of scale. We categorize the methods according to whether improved performance is achieved (primarily): *i*) via the exploitation of more geometric information about the target posterior; *ii*) by better choice of proposals; *iii*) by the use of parallel, batched, subsample, coupled or ensemble sampling methods; or *iv*) by the explicit use of variance reduction methods.

- i*) Hamiltonian Monte Carlo (HMC) (Neal, 2011a; Carpenter et al., 2017; Betancourt, 2018); no U-turn sampling (NUTS) (Hoffman and Gelman, 2014);¹⁸ Metropolis-Adjusted Langevin algorithm (Roberts et al., 1996; Roberts and Rosenthal, 1998); stochastic gradient MCMC (Nemeth and Fearnhead, 2019); piecewise deterministic Markov processes (Bierkens et al., 2018; Fearnhead et al., 2018; Bierkens, Fearnhead and Roberts, 2019).
- ii*) Optimal scaling of random-walk MH (Roberts, Gelman and Gilks, 1997); adaptive sampling (Nott and Kohn, 2005; Roberts and Rosenthal, 2009; Rosenthal, 2011); MCMC

the more recent literature is concerned with removing this transitory bias after a finite number of iterations; e.g. Jacob, O’Leary and Atchadé (2020). Other literature is concerned with ensuring that an MCMC algorithm does not yield a bias that is *non-transitory* due to the inability of the algorithm to effectively explore the target set at all (within a meaningful time frame); see e.g. Betancourt (2018).

¹⁷Whilst not a pseudo-marginal method, particle filtering has also been used to provide an estimate of $p(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})$ in a Gibbs scheme for an SSM — so-called ‘particle Gibbs’ (Andrieu, Doucet and Holenstein, 2011).

¹⁸As described in Neal (2011a), simulation methods based on Hamiltonian dynamics can actually be viewed as having as long a history as MCMC itself. The more modern manifestations of HMC, however, including NUTS, can be viewed as Markov chain algorithms that simply explore the parameter space more effectively than (say) a default random walk scheme. The probabilistic programming platform Stan (Carpenter et al., 2017) enables implementation of NUTS.

with ordered overrelaxation (Neal, 1998); simulated tempering, parallel tempering and tempered transition methods (Geyer, 1991; Marinari and Parisi, 1992; Neal, 1996; Gramacy, Samworth and King, 2010; Geyer, 2011b; Tawn, Roberts and Rosenthal, 2020); delayed rejection sampling (Tierney and Mira, 1998); delayed acceptance sampling (Christen and Fox, 2005; Golightly, Henderson and Sherlock, 2015; Wqvist et al., 2018; Banterle et al., 2019); multiple try MCMC (Liu, Liang and Wong, 2000; Bédard, Douc and Moulines, 2012; Martino, 2018; Luo and Tjelmeland, 2019); tailored randomized block MH (Chib and Ramamurthy, 2010); tempered Gibbs sampling (Zanella and Roberts, 2019); quasi-stationary Monte Carlo and subsampling (Pollock et al., 2020).

- iii) Parallelized MCMC (Jacob, Robert and Smith, 2011; Wang and Dunson, 2013); subposterior (batched) methods (Neiswanger, Wang and Xing, 2013; Scott et al., 2016); subsampling methods based on pseudo-marginal MCMC (Bardenet, Doucet and Holmes, 2017; Quiroz et al., 2018; Quiroz et al., 2019); perfect sampling (Propp and Wilson, 1996; Casella, Lavine and Robert, 2001; Craiu and Meng, 2011; Huber, 2016); unbiased MCMC via coupling (Glynn and Rhee, 2014; Glynn, 2016; Middleton et al., 2018; Jacob, O’Leary and Atchadé, 2020); unbiased MCMC using pseudo-marginal principles (Lyne et al., 2015); ensemble MCMC (Iba, 2000; Cappé et al., 2004; Neal, 2011b).
- iv) Rao-Blackwellization (Casella and Robert, 1996; Robert and Casella, 2004; Douc and Robert, 2011); antithetic variables (Frigessi, Gasemyr and Rue, 2000; Craiu and Meng, 2005); control variates (Dellaportas and Kontoyiannis, 2012; Baker et al., 2019); thinning (Owen, 2017).

5.1.2 The role of SMC. Sequential Monte Carlo (SMC) methods — which involve the sequential application of importance sampling steps — have also evolved over the 21st century. Most notably, they have expanded from being ‘particle filters’ designed for the sequential analysis of state space models, (Gordon, Salmond and Smith, 1993), into a broader set of techniques used to perform sequential tasks in non-state space settings; we refer the reader to Naesseth et al. (2019) and Chopin and Papaspiliopoulos (2020) for extensive reviews. From the perspective of dealing with intractable problems, their role in particle MCMC (as described above) has been particularly critical, and in playing that role they have retained their status as ‘exact’ methods. However, they have also played an important role in the context of certain of the approximate methods to be discussed below, via: ABC-SMC algorithms (Sisson, Fan and Tanaka, 2007; Beaumont et al., 2009), ‘ABC filtering’ (Jasra et al., 2012; Dean et al., 2014; Calvet and Czellar, 2015; Jasra, 2015), and VB with intractable likelihood (Tran, Nott and Kohn, 2017).

5.2 Approximation Solutions to Intractable Bayesian Problems

In contrast to exact methods of computation, when applying an approximate method investigators do not seek exactness, other than perhaps claiming asymptotic (in n) validity under certain conditions. That is, for finite n at least, such methods only ever provide a numerical solution to some *approximation* of (4). The advantage of such methods, however, is that they can yield ‘reasonable’ solutions, and often quickly, to empirical problems that would test the limits of exact methods or, indeed, be otherwise infeasible.

It is convenient, for the purpose of this review, to categorize the main approximate methods according to whether their primary goal is to obviate the need to evaluate the DGP (i.e. to solve the so-called *doubly-intractable* problem), or to tackle problems of dimension (either in the data, or the unknowns or, typically, in both). This categorization corresponds to a distinction between *simulation-based* approximate methods and approximate methods that are primarily based on *optimization*, and it is those categories that we use for the sub-section headings below.

5.2.1 Simulation-based methods: ABC and BSL. In computing (4), both ABC and BSL replace the posterior in the integrand, $p(\boldsymbol{\theta}|\mathbf{y})$, with a posterior that conditions only on some *function* of the observed data, \mathbf{y} . Typically, that function is a small-dimensional vector of summary statistics, $\boldsymbol{\eta}(\mathbf{y})$, that are chosen to be ‘informative’ about $\boldsymbol{\theta}$, in which case the posterior that is targeted is represented as $p(\boldsymbol{\theta}|\boldsymbol{\eta}(\mathbf{y}))$. Using M draws from $p(\boldsymbol{\theta}|\boldsymbol{\eta}(\mathbf{y}))$, $\boldsymbol{\theta}^{(i)}$, the sample mean, $\bar{g}(\boldsymbol{\theta}) = (1/M) \sum_{i=1}^M g(\boldsymbol{\theta}^{(i)})$, is used to estimate the expectation in (4). The key here is that producing draws from $p(\boldsymbol{\theta}|\boldsymbol{\eta}(\mathbf{y}))$ — and the precise way in which this is undertaken differs between the two methods — entails only simulation of the DGP, *not* its evaluation; hence the suitability of ABC and BSL for problems in which evaluation of the DGP is infeasible or, at the very least, computationally challenging. What is lost of course, is quantification of uncertainty about $\boldsymbol{\theta}$ that is informed by the full data set; inference is only ‘partial’, with the choice of summaries that are ‘close’ (in some sense) to being sufficient being an important goal. We refer the reader to [Marin et al. \(2011\)](#), [Price et al. \(2018\)](#), [Beaumont \(2019\)](#) and [Sisson, Fan and Beaumont \(2019\)](#), for details and extensive referencing.

5.2.2 Optimization methods: VB and INLA. VB and INLA target alternative approximations to (4). Both methods are particularly beneficial when the dimension of the unknowns, or the dimension of the data, or both, are very large, usually due to the presence in the model of a high number of latent, or ‘local’, parameters, \mathbf{x} , in addition to the smaller set of ‘global’ (or ‘hyperparameters’) parameters, $\boldsymbol{\theta}$.

Using the principle of the calculus of variations, VB produces an approximation of $p(\boldsymbol{\theta}|\mathbf{y})$, $q^*(\boldsymbol{\theta})$ say, that is ‘closest’ to $p(\boldsymbol{\theta}|\mathbf{y})$ within a chosen family of densities. Depending on the form of $q^*(\boldsymbol{\theta})$, the expectation defined with respect to this *variational posterior* may be available in closed-form or, at least, able to be estimated using simple Monte Carlo sampling from $q^*(\boldsymbol{\theta})$. INLA, on the other hand, adapts the approximation method of Laplace to a high-dimensional setting - allied with low dimensional deterministic integration - to produce an approximation of (4). Both methods rely critically on modern techniques of optimization, for the purpose of minimizing the ‘distance’ between $p(\boldsymbol{\theta}|\mathbf{y})$ and $q^*(\boldsymbol{\theta})$ in the case of VB, and for the purpose of producing the mode of the high-dimensional \mathbf{x} , for use in the series of (nested) Laplace approximations that underpin INLA. See [Ormerod and Wand \(2010\)](#), [Blei, Kucukelbir and McAuliffe \(2017\)](#) and [Zhang et al. \(2018\)](#), for reviews of VB; and [Rue, Martino and Chopin \(2009\)](#), [Rue et al. \(2017\)](#), [Martino and Riebler \(2019\)](#), [Van Niekerk et al. \(2019\)](#) and [Wood \(2019\)](#) for all details on INLA.

Finally, see [Martin, Frazier and Robert \(2021\)](#) for a detailed outline and comparison of *all* approximate methods, including additional discussion on ‘hybrid’ methods that the meld features of more than one approximate technique, with the goal of tackling multiple instances of intractability.

6. POSTSCRIPT

Our journey with Bayesian computation began in 1763: with a posterior probability defined in terms of a scalar θ , whose solution challenged Bayes. We now end our journey in 2022: having referenced papers that tackle posterior distributions defined over thousands, possibly millions of unknowns, and computational problems with a degree of complexity — and scale — to match. Along the way, we have seen the huge variety of imaginative computational solutions that have been brought to bear on all such problems, over the span of 250 years. Moreover, Bayesian computation is also beginning to confront — and adapt to — the reality of misspecified DGPs ([Wang and Blei, 2019](#); [Frazier, Robert and Rousseau, 2020](#); [Frazier and Drovandi, 2021](#)), and the generalizations beyond the standard likelihood-based up-date that are evolving ([Bornn, Shephard and Solgi, 2019](#); [Schmon, Cannon and Knoblauch, 2020](#); [Frazier et al., 2021](#); [Knoblauch, Jewson and Damoulas, 2022](#)). The future of the paradigm in the 21st century thus seems assured. And with this, the 18th century Bayes (and his loyal champion, Price) would no doubt be duly impressed!

REFERENCES

- ALBERT, J. H. and CHIB, S. (1993). Bayesian analysis of binary and polychotomous response data. *J. American Statist. Assoc.* **88** 669–679.
- ANDRAL, C. (2022). An attempt to trace the birth of importance sampling. <https://hal.archives-ouvertes.fr/hal-03696061>.
- ANDRIEU, C., DOUCET, A. and ROBERT, C. (2004). Computational advances for and from Bayesian analysis. *Statist. Science* **19** 118–127.
- ANDRIEU, C., DOUCET, A. and HOLENSTEIN, R. (2011). Particle Markov chain Monte Carlo. *J. Royal Statist. Society Series B* **72(2)** 269–342. With discussion.
- ANDRIEU, C. and ROBERTS, G. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *Ann. Statist.* **37** 697–725.
- ARDIA, D., BAŞTÜRK, N., HOOGERHEIDE, L. and VAN DIJK, H. K. (2012). A comparative study of Monte Carlo methods for efficient evaluation of marginal likelihood. *Computational Statistics and Data Analysis* **56** 3398–3414.
- BAKER, J., FEARNHEAD, P., FOX, E. and NEMETH, C. (2019). Control variates for stochastic gradient MCMC. *Statist. Comp.* **29** 599–615.
- BANTERLE, M., GRAZIAN, C., LEE, A. and ROBERT, C. P. (2019). Accelerating Metropolis-Hastings algorithms by delayed acceptance. *Foundations of Data Science* **1** 103–128.
- BARDENET, R., DOUCET, A. and HOLMES, C. (2017). On Markov chain Monte Carlo methods for tall data. *J. Machine Learning Res.* **18** 1515–1557.
- BARNARD, G. A. and BAYES, T. (1958). Studies in the History of Probability and Statistics: IX. Thomas Bayes’s Essay Towards Solving a Problem in the Doctrine of Chances. *Biometrika* **45** 293–315.
- BAYES, T. (1764). An Essay Towards Solving a Problem in the Doctrine of Chances. *Philosophical Transactions of the Royal Society of London for 1773* **53** 370–418.
- BEAUMONT, M. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics* **164** 1139–1160.
- BEAUMONT, M. A. (2019). Approximate Bayesian Computation. *Annual Review of Statistics and Its Application* **6** 379–403.
- BEAUMONT, M. A., CORNUET, J. M., MARIN, J. M. and ROBERT, C. P. (2009). Adaptive approximate Bayesian Computation. *Biometrika* **96(4)** 983–990.
- BÉDARD, M., DOUC, R. and MOULINES, E. (2012). Scaling analysis of multiple-try MCMC methods. *Stochastic Processes and their Applications* **122** 758–786.
- BERGER, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*, Second ed. Springer-Verlag, New York.
- BESAG, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *J. Royal Statist. Society Series B* **36** 192–326. With discussion.
- BESAG, J. and GREEN, P. J. (1993). Spatial statistics and Bayesian computation. *J. Royal Statist. Society Series B* **55** 25–37. With discussion.
- BETANCOURT, M. (2018). A Conceptual Introduction to Hamiltonian Monte Carlo. <https://arxiv.org/abs/1701.02434v2>.
- BIERKENS, J., FEARNHEAD, P. and ROBERTS, G. (2019). The Zig-Zag process and super-efficient sampling for Bayesian analysis of big data. *Ann. Statist.* **47** 1288–1320.
- BIERKENS, J., BOUCHARD-CÔTÉ, A., DOUCET, A., DUNCAN, A. B., FEARNHEAD, P., LIENART, T., ROBERTS, G. and VOLLMER, S. J. (2018). Piecewise deterministic Markov processes for scalable Monte Carlo on restricted domains. *Statist. Prob. Letters* **136** 148–154.
- BLEI, D. M., KUCUKELBIR, A. and MCAULIFFE, J. D. (2017). Variational inference: A review for statisticians. *J. American Statist. Assoc.* **112** 859–877.
- BORNN, L., SHEPHARD, N. and SOLGI, R. (2019). Moment conditions and Bayesian non-parametrics. *J. Royal Statist. Society Series B* **81** 5–43.
- BROOKS, S., GELMAN, A., JONES, G. and MENG, X. L. (2011). *Handbook of Markov Chain Monte Carlo*. Taylor & Francis.
- CALVET, L. E. and CZELLAR, V. (2015). Accurate methods for approximate Bayesian computation filtering. *J. Finan. Econometrics* **13** 798–838.
- CAPPÉ, O., GUILLIN, A., MARIN, J. M. and ROBERT, C. P. (2004). Population Monte Carlo. *J. Comput. Graph. Statist.* **13** 907–929.
- CARPENTER, B., GELMAN, A., HOFFMAN, M., LEE, D., GOODRICH, B., BETANCOURT, M., BRUBAKER, M., GUO, J., LI, P. and RIDDELL, A. (2017). Stan: A Probabilistic Programming Language. *Journal of Statistical Software, Articles* **76**.
- CARTER, C. K. and KOHN, R. (1994). On Gibbs sampling for state space models. *Biometrika* **81** 541–553.
- CASELLA, G. and GEORGE, E. I. (1992). An introduction to Gibbs sampling. *American Statist.* **46** 167–174.
- CASELLA, G., LAVINE, M. and ROBERT, C. P. (2001). Explaining the Perfect Sampler. *American Statist.* **55** 299–305.
- CASELLA, G. and ROBERT, C. (1996). Rao-Blackwellisation of sampling schemes. *Biometrika* **83** 81–94.

- CERUZZI, P. E. (2003). *A History of Modern Computing*, second ed. MIT Press.
- CHEN, S., DICK, J. and OWEN, A. B. (2011). Consistency of Markov chain quasi-Monte Carlo on continuous state spaces. *Ann. Statist.* **39** 673–701.
- CHIB, S. and GREENBERG, E. (1995). Understanding the Metropolis–Hastings algorithm. *American Statist.* **49** 327–335.
- CHIB, S. and GREENBERG, E. (1996). Markov Chain Monte Carlo Simulation Methods in Econometrics. *Econometric Theory* **12** 409–431.
- CHIB, S. and RAMAMURTHY, S. (2010). Tailored randomized block MCMC methods with application to DSGE models. *Journal of Econometrics* **155** 19–38.
- CHOPIN, N. and PAPASPILIOPOULOS (2020). *An introduction to sequential Monte Carlo*. Springer.
- CHRISTEN, J. A. and FOX, C. (2005). Markov chain Monte Carlo using an approximation. *J. Comput. Graph. Statist.* **14** 795–810.
- CRAIU, R. V. and MENG, X.-L. (2005). Multiprocess parallel antithetic coupling for backward and forward Markov chain Monte Carlo. *Ann. Statist.* **33** 661–697.
- CRAIU, R. V. and MENG, X.-L. (2011). Perfection within reach: exact MCMC sampling. *Handbook of Markov Chain Monte Carlo* 199–226. Chapman & Hall/CRC. Eds. Brooks, S., Gelman, A., Jones, G., Meng, X.-L.
- DAMIEN, P., WAKEFIELD, J. and WALKER, S. (1999). Gibbs sampling for Bayesian non-conjugate and hierarchical models by using auxiliary variables. *J. Royal Statist. Society Series B* **61** 331–344.
- DE BRUIJN, N. (1961). *Asymptotic methods in analysis*. 2nd Edn., Sect. 4.2. Amsterdam: North-Holland.
- DEAN, T. A., SINGH, S. S., JASRA, A. and PETERS, G. W. (2014). Parameter estimation for hidden Markov models with intractable likelihoods. *Scandinavian Journal of Statistics* **41** 970–987.
- DELIGIANNIDIS, G., DOUCET, A. and PITT, M. K. (2018). The correlated pseudomarginal method. *J. Royal Statist. Society Series B* **80** 839–870.
- DELLAPORTAS, P. and KONTIOYIANNIS, I. (2012). Control variates for estimation based on reversible Markov chain Monte Carlo samplers. *J. Royal Statist. Society Series B* **74** 133–161.
- DEVROYE, L. (1986). *Non-Uniform Random Variate Generation*. Springer-Verlag, New York.
- DIEBOLT, J. and ROBERT, C. P. (1994). Estimation of Finite Mixture Distributions by Bayesian Sampling. *J. Royal Statist. Society Series B* **56** 363–375.
- DONGARRA, J. and SULLIVAN, F. (2000). Guest editors’ introduction: The top 10 algorithms. *Computing in Science & Engineering* **2** 22–23.
- DOUC, R. and ROBERT, C. P. (2011). A vanilla Rao–Blackwellization of Metropolis–Hastings algorithms. *Ann. Statist.* **39** 261–277.
- DOUCET, A., PITT, M. K., DELIGIANNIDIS, G. and KOHN, R. (2015). Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika* **102** 295–313.
- DUNSON, D. and JOHNDRON, J. (2019). The Hastings algorithm at fifty. *Biometrika* **107** 1–23.
- FEARNHEAD, P., BIERKENS, J., POLLOCK, M. and ROBERTS, G. O. (2018). Piecewise Deterministic Markov Processes for Continuous-Time Monte Carlo. *Statist. Sci.* **33** 386–412.
- FIENBERG, S. E. (2006). When did Bayesian inference become “Bayesian”? *Bayesian Analysis* **1**(1) 1–40.
- FRAZIER, D. T. and DROVANDI, C. (2021). Robust approximate Bayesian inference with synthetic likelihood. *Journal of Computational and Graphical Statistics* **30** 958–976.
- FRAZIER, D. T., ROBERT, C. P. and ROUSSEAU, J. (2020). Model misspecification in approximate Bayesian computation: consequences and diagnostics. *J. Royal Statist. Society Series B*.
- FRAZIER, D. T., LOAIZA-MAYA, R., MARTIN, G. M. and KOO, B. (2021). Loss-based variational Bayes prediction. *arXiv preprint arXiv:2104.14054*.
- FRIGESSI, A., GASEMYR, J. and RUE, H. (2000). Antithetic Coupling of Two Gibbs Sampler Chains. *The Annals of Statistics* **28** 1128–1149.
- FRÜHWIRTH-SCHNATTER, S. (1994). Data Augmentation and Dynamic Linear Models. *J. Time Ser. Anal.* **15** 183–202.
- GELFAND, A. E. and SMITH, A. F. M. (1990). Sampling based approaches to calculating marginal densities. *J. Amer. Statist. Assoc.* **85** 398–409.
- GELMAN, A. and SHIRLEY, K. (2011). Inference from simulations and monitoring convergence. *Handbook of Markov chain Monte Carlo* 163–174. Chapman & Hall/CRC. Eds. Brooks, S., Gelman, A., Jones, G., Meng, X.-L.
- GEMAN, S. and GEMAN, D. (1984). Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* **6** 721–741.
- GERBER, M. and CHOPIN, N. (2015). Sequential quasi Monte Carlo. *J. Royal Statist. Society Series B* **77** 509–579.
- GEWEKE, J. (1989). Bayesian inference in econometric models using Monte Carlo integration. *Econometrica* **57** 1317–1340.
- GEWEKE, J. (2005). *Contemporary Bayesian Econometrics and Statistics*. Wiley.
- GEWEKE, J., KOOP, G. and VAN DIJK, H. (2011). *The Oxford Handbook of Bayesian Econometrics*. OUP.
- GEYER, C. (1991). Markov Chain Monte Carlo maximum likelihood. In *Computing Science and Statistics: Proceedings of the 32nd Symposium on the Interface* (E. M. KERAMIGAS, ed.) 156–163. Interface Foundation, Fairfax.

- GEYER, C. J. (2011a). Introduction to Markov chain Monte Carlo. *Handbook of Markov chain Monte Carlo* 3–48. Chapman & Hall/CRC. Eds. Brooks, S., Gelman, A., Jones, G., Meng, X.-L.
- GEYER, C. J. (2011b). Importance Sampling, Simulated Tempering, and Umbrella Sampling. *Handbook of Markov chain Monte Carlo* 295–311. Chapman & Hall/CRC. Eds. Brooks, S., Gelman, A., Jones, G., Meng, X.-L.
- GHOSAL, S., GHOSH, J. K. and SAMANTA, T. (1995). On convergence of posterior distributions. *The Annals of Statistics* **23** 2145–2152.
- GLYNN, P. W. (2016). Exact Simulation vs. Exact Estimation. *Proceedings of the 2016 Winter Simulation Conference WSC16* 193–205.
- GLYNN, P. W. and RHEE, C.-H. (2014). Exact estimation for Markov chain equilibrium expectations. *J. Appl. Probab.* **51** 377–389.
- GOERTZEL, G. and KAHN, H. (1950). Monte Carlo Methods for Shield Computation. *U.S. Atomic Energy Commission, Technical Information Division*.
- GOLIGHTLY, A., HENDERSON, D. A. and SHERLOCK, C. (2015). Delayed acceptance particle MCMC for exact inference in stochastic kinetic models. *Statist. Comp.* **25** 1039–1055.
- GORDON, N., SALMOND, J. and SMITH, A. F. M. (1993). A novel approach to non-linear/non-Gaussian Bayesian state estimation. *IEEE Proceedings on Radar and Signal Processing* **140** 107–113.
- GRAMACY, R., SAMWORTH, R. and KING, R. (2010). Importance tempering. *Statist. Comp.* **20** 1–7.
- GREEN, P., LATUSZYNSKI, K., PEREYRA, M. and ROBERT, C. (2015). Bayesian computation: a summary of the current state, and samples backwards and forwards. *Statist. Comp.* **25** 835–862.
- GUBERNATIS, J. E. (2005). Marshall Rosenbluth and the Metropolis algorithm. *Physics of Plasmas* **12** 057303.
- HAMMERSLEY, J. M. and HANDSCOMB, D. C. (1964). *Monte Carlo Methods*. John Wiley, New York.
- HASTINGS, W. K. (1970). Monte Carlo sampling methods using Markov chains and their application. *Biometrika* **57** 97–109.
- HIGDON, D. M. (1998). Auxiliary Variable Methods for Markov Chain Monte Carlo with Applications. *Journal of the American Statistical Association* **93** 585–595.
- HITCHCOCK, D. H. (2003). A history of the Metropolis–Hastings algorithm. *American Statist.* **57** 254–257.
- HOFFMAN, M. D. and GELMAN, A. (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* **15** 1593–1623.
- HOOPER, M. (2013). Richard Price, Bayes’ theorem, and God. *Significance* **10** 36–39.
- HUBER, M. L. (2016). *Perfect simulation*. Chapman & Hall/CRC.
- IBA, Y. (2000). Population-based Monte Carlo algorithms. *Trans. Japanese Soc. Artificial Intell.* **16** 279–286.
- JACOB, P. E., O’LEARY, J. and ATCHADÉ, Y. F. (2020). Unbiased Markov chain Monte Carlo methods with couplings. *J. Royal Statist. Society Series B* **82** 1–32. With discussion.
- JACOB, P., ROBERT, C. P. and SMITH, M. A. (2011). Using parallel computation to improve independent Metropolis–Hastings based estimation. *J. Comput. Graph. Statist.* **20** 616–635.
- JACQUIER, R., POLSON, N. G. and ROSSI, P. E. (1994). Bayesian analysis of stochastic volatility models. *J. Business and Economic Statistics* **12** 371–389. With discussion.
- JAHAN, F., ULLAH, I. and MENGENSEN, K. (2020). A review of Bayesian statistical approaches for Big Data. In *Case Studies in Applied Bayesian Science* (K. MENGENSEN, P. PUDLO and C. P. ROBERT, eds.) 17–44. Springer.
- JASRA, A. (2015). Approximate Bayesian computation for a class of time series models. *International Statistical Review* **83** 405–435.
- JASRA, A., SINGH, S. S., MARTIN, J. S. and MCCOY, E. (2012). Filtering via Approximate Bayesian computation. *Statist. Comp.* **22** 1223–1237.
- JOHNDROW, J. E., SMITH, A., PILLAI, N. and DUNSON, D. B. (2019). MCMC for imbalanced categorical data. *J. American Statist. Assoc.* **114** 1394–1403.
- JORDAN, M., GHAHRAMANI, Z. and JAAKKOLA, T. (1999). An Introduction to Variational Methods for Graphical Models. *Machine Learning* **6** 183–233.
- KAHN, H. (1949). Stochastic (Monte Carlo) Attenuation Analysis. *Technical Report, RAND Corporation* R-163.
- KIM, S., SHEPHARD, N. and CHIB, S. (1998). Stochastic volatility: likelihood inference and comparison with ARCH Models. *The Review of Economic Studies* **65** 361–393.
- KLOEK, T. and VAN DIJK, H. K. (1978). Bayesian estimates of equation system parameters: an application of integration by Monte Carlo. *Econometrica* **46** 1–19.
- KNOBLAUCH, J., JEWSON, J. and DAMOULAS, T. (2022). An Optimization-centric View on Bayes’ Rule: Reviewing and Generalizing Variational Inference. *Journal of Machine Learning Research* **23** 1–109.
- KOOP, G. M. (2003). *Bayesian Econometrics*. John Wiley & Sons Inc.
- LEMIEUX, C. (2009). *Monte Carlo and quasi-Monte Carlo sampling*. Springer Science & Business Media.
- LINDLEY, D. V. (1980). Approximate Bayesian methods. In *Bayesian statistics: Proceedings of the first international meeting held in Valencia (Spain), May 28 to June 2, 1979*.
- LIU, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer Verlag, New-York.
- LIU, J. S., LIANG, F. and WONG, W. H. (2000). The Multiple-Try Method and Local Optimization in Metropolis Sampling. *J. American Statist. Assoc.* **95** 121–134.

- LLORENTE, F., MARTINO, L., DELGADO, D. and LOPEZ-SANTIAGO, J. (2021). Marginal likelihood computation for model selection and hypothesis testing: an extensive review. <https://arXiv:2005.08334>.
- LUO, X. and TJELMELAND, H. (2019). A multiple-try Metropolis-Hastings algorithm with tailored proposals. *Computational Statistics* **34** 1109–1133.
- LYNE, A.-M., GIROLAMI, M., ATCHADÉ, Y., STRATHMANN, H. and SIMPSON, D. (2015). On Russian Roulette Estimates for Bayesian Inference with Doubly-Intractable Likelihoods. *Statist. Sci.* **30** 443–467.
- MARIN, J., PUDLO, P., ROBERT, C. and RYDER, R. (2011). Approximate Bayesian computational methods. *Statist. Comp.* **21**(2) 279–291.
- MARINARI, E. and PARISI, G. (1992). Simulated tempering: A new Monte Carlo scheme. *Europhysics Letters* **19** 451–458.
- MARTIN, G. M., FRAZIER, D. T. and ROBERT, C. P. (2020). Computing Bayes: Bayesian Computation from 1763 to the 21st Century. <https://arXiv:2004.06425>.
- MARTIN, G. M., FRAZIER, D. T. and ROBERT, C. P. (2021). Approximating Bayes in the 21st Century. <https://arXiv:2112.10342>.
- MARTINO, L. (2018). A review of multiple try MCMC algorithms for signal processing. *Digital Signal Processing* **75** 134–152.
- MARTINO, S. and RIEBLER, A. (2019). Integrated nested Laplace approximations (INLA). <https://arXiv:1907.01248>.
- METROPOLIS, N., ROSENBLUTH, A. W., ROSENBLUTH, M. N., TELLER, A. H. and TELLER, E. (1953). Equations of state calculations by fast computing machines. *J. Chem. Phys.* **21** 1087–1092.
- MIDDLETON, L., DELIGIANNIDIS, G., DOUCET, A. and JACOB, P. E. (2018). Unbiased Markov chain Monte Carlo for intractable target distributions. <https://arXiv:1807.08691>.
- NAESSETH, C. A., LINDSTEN, F., SCHÖN, T. B. et al. (2019). Elements of sequential Monte Carlo. *Foundations and Trends in Machine Learning* **12** 307–392.
- NEAL, R. (1996). Sampling from multimodal distributions using tempered transitions. *Statist. Comp.* **6** 353–366.
- NEAL, R. (1998). Suppressing Random Walks in Markov Chain Monte Carlo Using Ordered Overrelaxation. *Learning in Graphical Models* 205–228. Chapman & Hall/CRC. Ed. Jordan, M. I.
- NEAL, R. (2003). Slice sampling. *Ann. Statist.* **31** 705–767. With discussion.
- NEAL, R. (2011a). MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo* 113–162. Chapman & Hall/CRC. Eds. Brooks, S., Gelman, A., Jones, G., Meng, X.-L.
- NEAL, R. (2011b). MCMC using ensembles of states for problems with fast and slow variables such as Gaussian process regression. <https://arXiv:1101.0387>.
- NEISWANGER, W., WANG, C. and XING, E. (2013). Asymptotically Exact, Embarrassingly Parallel MCMC. <https://arXiv:1311.4780>.
- NEMETH, C. and FEARNHEAD, P. (2019). Stochastic gradient Markov chain Monte Carlo. <https://arXiv:1907.06986>.
- NOTT, D. J. and KOHN, R. (2005). Adaptive sampling for Bayesian variable selection. *Biometrika* **92** 747–763.
- ORMEROD, J. T. and WAND, M. P. (2010). Explaining variational approximations. *American Statist.* **64** 140–153.
- OWEN, A. B. (2017). Statistically efficient thinning of a Markov chain sampler. *J. Comput. Graph. Statist.* **26** 738–744.
- PITT, M. K., DOS SANTOS SILVA, R., GIORDANI, P. and KOHN, R. (2012). On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *J. Econometrics* **171** 134–151.
- POLLOCK, M., FEARNHEAD, P., JOHANSEN, A. M. and ROBERTS, G. O. (2020). Quasi-stationary Monte Carlo and the ScaLE algorithm. *J. Royal Statist. Society Series B* **82** 1167–1221. (With discussion.).
- POLSON, N. G., CARLIN, B. P. and STOFFER, D. S. (1992). A Monte Carlo Approach to Nonnormal and Non-linear State-Space Modeling. *J. American Statist. Assoc.* **87** 493–500.
- PRICE, R. (1764). A demonstration of the second rule in the essay towards the solution of a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London* **54** 296–325.
- PRICE, L. F., DROVANDI, C. C., LEE, A. and NOTT, D. J. (2018). Bayesian synthetic likelihood. *J. Comput. Graph. Statist.* **27** 1–11.
- PRITCHARD, J. K., SEIELSTAD, M. T., PEREZ-LEZAUN, A. and FELDMAN, M. W. (1999). Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Mol. Biol. Evol.* **16** 1791–1798.
- PROPP, J. G. and WILSON, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. In *Proceedings of the Seventh International Conference on Random Structures and Algorithms (Atlanta, GA, 1995)* **9** 223–252. [MR1611693 \(99k:60176\)](https://arXiv:1611.06933)
- QUIROZ, M., TRAN, M.-N., VILLANI, M. and KOHN, R. (2018). Speeding up MCMC by delayed acceptance and data subsampling. *J. Comput. Graph. Statist.* **27** 12–22.
- QUIROZ, M., KOHN, R., VILLANI, M. and TRAN, M.-N. (2019). Speeding up MCMC by efficient data subsampling. *J. American Statist. Assoc.* **114** 831–843.
- RITTER, C. and TANNER, M. A. (1992). Facilitating the Gibbs sampler: The Gibbs stopper and the Griddy-Gibbs sampler. *J. American Statist. Assoc.* **87** 861–868.
- ROBERT, C. (2007). *The Bayesian Choice*. Springer-Verlag, New York.
- ROBERT, C. P. (2015). *The Metropolis–Hastings Algorithm* In *Wiley StatsRef: Statistics Reference Online* 1–15. Wiley StatsRef: Statistics Reference Online. American Cancer Society.

- ROBERT, C. and CASELLA, G. (2004). *Monte Carlo Statistical Methods*, second ed. Springer-Verlag, New York.
- ROBERT, C. and CASELLA, G. (2011). A history of Markov Chain Monte Carlo—Subjective recollections from incomplete data. *Statist. Science* **26** 102–115.
- ROBERT, C. P., ELVIRA, V., TAWN, N. and WU, C. (2018). Accelerating MCMC algorithms. *Wiley Interdisciplinary Reviews: Computational Statistics* **10** e1435.
- ROBERTS, G. O., GELMAN, A. and GILKS, W. R. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *Ann. Applied Probab.* **7** 110–120.
- ROBERTS, G. O. and ROSENTHAL, J. S. (1998). Optimal scaling of discrete approximations to Langevin diffusions. *J. Royal Statist. Society Series B* **60** 255–268.
- ROBERTS, G. O. and ROSENTHAL, J. S. (1999). Convergence of slice sampler Markov chains. *J. Royal Statist. Society Series B* **61** 643–660.
- ROBERTS, G. O. and ROSENTHAL, J. S. (2009). Examples of adaptive MCMC. *J. Comput. Graph. Statist.* **18** 349–367.
- ROBERTS, G. O. and SAHU, S. K. (1997). Updating schemes, covariance structure, blocking and parametrisation for the Gibbs sampler. *J. Royal Statist. Society Series B* **59** 291–317.
- ROBERTS, G. O., TWEEDIE, R. L. et al. (1996). Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* **2** 341–363.
- ROSENTHAL, J. S. (2011). Optimal proposal distributions and adaptive MCMC. *Handbook of Markov Chain Monte Carlo* 93–111. Chapman & Hall/CRC. Eds. Brooks, S., Gelman, A., Jones, G., Meng, X-L.
- RUE, H., MARTINO, S. and CHOPIN, N. (2009). Approximate Bayesian Inference for Latent Gaussian Models Using Integrated Nested Laplace Approximations. *J. Royal Statist. Society Series B* **71** 319–392.
- RUE, H., RIEBLER, A., SØRBYE, S. H., ILLIAN, J. B., SIMPSON, D. P. and LINDGREN, F. K. (2017). Bayesian Computing with INLA: A Review. *Annual Review of Statistics and Its Application* **4** 395–421.
- SCHMON, S. M., CANNON, P. W. and KNOBLAUCH, J. (2020). Generalized posteriors in approximate Bayesian computation. *arXiv preprint arXiv:2011.08644*.
- SCOTT, S. L., BLOCKER, A. W., BONASSI, F. V., CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (2016). Bayes and big data: The consensus Monte Carlo algorithm. *International Journal of Management Science and Engineering Management* **11** 78–88.
- SISSON, S. A., FAN, Y. and TANAKA, M. M. (2007). Sequential Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci. USA* **104** 1760–1765.
- SISSON, S. A., FAN, Y. and BEAUMONT, M. (2019). *Handbook of Approximate Bayesian Computation*. Chapman & Hall/CRC.
- SMITH, A. F. M. and ROBERTS, G. O. (1993). Bayesian computation via the Gibbs sampler and related Markov chain Monte Carlo methods. *J. Royal Statist. Society Series B* **55** 3–24. With discussion.
- STIGLER, S. M. (1975). Studies in the History of Probability and Statistics. XXXIV Napoleonic Statistics: The Work of Laplace. *Biometrika* **62** 503–517.
- STIGLER, S. M. (1986a). *The History of Statistics*. Belknap, Cambridge.
- STIGLER, S. M. (1986b). Memoir on Inverse Probability. *Statistical Science* **1** 359–363.
- STIGLER, S. M. (2018). Richard Price, the First Bayesian. *Statist. Sci.* **33** 117–125.
- TANNER, M. A. and WONG, W. (1987). The calculation of posterior distributions by data augmentation. *J. American Statist. Assoc.* **82** 528–550. With discussion.
- TAVARÉ, S. D., BALDING, D. J., GRIFFITH, R. C. and DONNELLY, P. (1997). Inferring coalescence times from DNA sequence data. *Genetics* **145** 505–518.
- TAWN, N. G., ROBERTS, G. O. and ROSENTHAL, J. S. (2020). Weight-preserving simulated tempering. *Statist. Comp.* **30** 27–41.
- TIERNEY, L. (1994). Markov chains for exploring posterior distributions. *Ann. Statist.* **22** 1701–1762. With discussion and a rejoinder by the author. [MR96m:60150](#)
- TIERNEY, L. and KADANE, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *J. American Statist. Assoc.* **81** 82–86.
- TIERNEY, L., KASS, R. E. and KADANE, J. B. (1989). Fully exponential Laplace approximations to expectations and variances of non-positive functions. *J. American Statist. Assoc.* **84** 710–716.
- TIERNEY, L. and MIRA, A. (1998). Some adaptive Monte Carlo methods for Bayesian inference. *Statistics in Medicine* **18** 2507–2515.
- TRAN, K. T. (2018). A Common Derivation for Markov Chain Monte Carlo Algorithms with Tractable and Intractable Targets. *arXiv:1607.01985*.
- TRAN, M.-N., NOTT, D. J. and KOHN, R. (2017). Variational Bayes with intractable likelihood. *J. Comput. Graph. Statist.* **26** 873–882.
- VAN DER VAART, A. (1998). *Asymptotic Statistics*. Cambridge University Press.
- VAN NIEKERK, J., BAKKA, H., RUE, H. and SCHENK, O. (2019). New frontiers in Bayesian modeling using the INLA package in R. *arXiv preprint arXiv:1907.10426*.
- WANG, Y. and BLEI, D. (2019). Variational Bayes under Model Misspecification. In *Advances in Neural Information Processing Systems* 13357–13367.
- WANG, X. and DUNSON, D. B. (2013). Parallel MCMC via Weierstrass Sampler. <https://arXiv:1312.4605>.

- WIKVIST, S., PICCHINI, U., FORMAN, J. L., LINDORFF-LARSEN, K. and BOOMSMA, W. (2018). Accelerating delayed-acceptance Markov chain Monte Carlo algorithms. <https://arXiv:1806.05982>.
- WOOD, S. (2019). Simplified integrated nested Laplace approximation. *Biometrika* **107** 223–230.
- ZANELLA, G. and ROBERTS, G. (2019). Scalable importance tempering and Bayesian variable selection. *J. Royal Statist. Society Series B* **81** 489–517.
- ZHANG, C., BÜTEPAGE, J., KJELLSTRÖM, H. and MANDT, S. (2018). Advances in variational inference. *IEEE transactions on pattern analysis and machine intelligence* **41** 2008–2026.