



HAL
open science

Monte Carlo Methods in Statistics

Christian Robert

► **To cite this version:**

| Christian Robert. Monte Carlo Methods in Statistics. 2009. hal-00412639

HAL Id: hal-00412639

<https://hal.science/hal-00412639v1>

Preprint submitted on 2 Sep 2009

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.

Monte Carlo Methods in Statistics

CHRISTIAN ROBERT*

Université Paris Dauphine and CREST, INSEE

September 2, 2009

Monte Carlo methods are now an essential part of the statistician's toolbox, to the point of being more familiar to graduate students than the measure theoretic notions upon which they are based! We recall in this note some of the advances made in the design of Monte Carlo techniques towards their use in Statistics, referring to Robert and Casella (2004, 2010) for an in-depth coverage.

The basic Monte Carlo principle and its extensions

The most appealing feature of Monte Carlo methods [for a statistician] is that they rely on sampling and on probability notions, which are the bread and butter of our profession. Indeed, the foundation of Monte Carlo approximations is identical to the validation of empirical moment estimators in that the average

$$\frac{1}{T} \sum_{t=1}^T h(x_t), \quad x_t \sim f(x), \quad (1)$$

is converging to the expectation $\mathbb{E}_f[h(X)]$ when T goes to infinity. Furthermore, the precision of this approximation is exactly of the same kind as the precision of a statistical estimate, in that it usually evolves as $O(\sqrt{T})$. Therefore, once a sample x_1, \dots, x_T is produced according to a distribution density f , all standard statistical tools, including bootstrap, apply to this sample (with the further appeal that more data points can be produced if deemed necessary). As illustrated by Figure 1, the variability due to a single Monte Carlo experiment must be accounted for, when drawing conclusions about its output and evaluations

*Professor of Statistics, CEREMADE, Université Paris Dauphine, 75785 Paris cedex 16, France. Supported by the Agence Nationale de la Recherche (ANR, 212, rue de Bercy 75012 Paris) through the 2009-2012 project ANR=08-BLAN-0218 Big'MC. Email: xian@ceremade.dauphine.fr. The author is grateful to Jean-Michel Marin for helpful comments.

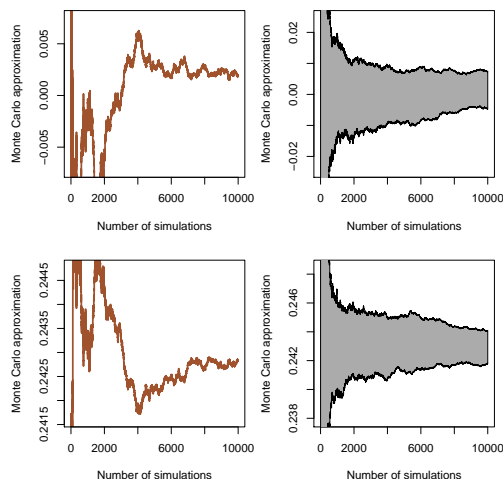


Figure 1: Monte Carlo evaluation (1) of the expectation $\mathbb{E}[X^3/(1+X^2+X^4)]$ as a function of the number of simulation when $X \sim \mathcal{N}(\mu, 1)$ using (left) one simulation run and (right) 100 independent runs for (top) $\mu = 0$ and (bottom) $\mu = 2.5$.

of the overall variability of the sequence of approximations are provided in Kendall et al. (2007). But the ease with which such methods are analysed and the systematic resort to statistical intuition explain in part why Monte Carlo methods are privileged over numerical methods.

The representation of integrals as expectations $\mathbb{E}_f[h(X)]$ is far from unique and there exist therefore many possible approaches to the above approximation. This range of choices corresponds to the importance sampling strategies (Rubinstein 1981) in Monte Carlo, based on the obvious identity

$$\mathbb{E}_f[h(X)] = \mathbb{E}_g[h(X)f(X)/g(X)]$$

provided the support of the density g includes the support of f . Some choices of g may however lead to appallingly poor performances of the resulting Monte

Carlo estimates, in that the variance of the resulting empirical average may be infinite, a danger worth highlighting since often neglected while having a major impact on the quality of the approximations. From a statistical perspective, there exist some natural choices for the importance function g , based on Fisher information and analytical approximations to the likelihood function like the Laplace approximation (Rue et al. 2008), even though it is more robust to replace the normal distribution in the Laplace approximation with a t distribution. The special case of Bayes factors (Robert and Casella 2004)

$$B_{01}(x) = \int_{\Theta} f(x|\theta)\pi_0(\theta)d\theta \bigg/ \int_{\Theta} f(x|\theta)\pi_1(\theta)d\theta,$$

which drive Bayesian testing and model choice, and of their approximation has led to a specific class of importance sampling techniques known as *bridge sampling* (Chen et al. 2000) where the optimal importance function is made of a mixture of the posterior distributions corresponding to both models (assuming both parameter spaces can be mapped into the same Θ). We want to stress here that an alternative approximation of marginal likelihoods relying on the use of *harmonic means* (Gelfand and Dey 1994, Newton and Raftery 1994) and of direct simulations from a posterior density has repeatedly been used in the literature, despite often suffering from infinite variance (and thus numerical instability). Another potentially very efficient approximation of Bayes factors is provided by Chib’s (1995) representation, based on parametric estimates to the posterior distribution.

MCMC methods

Markov chain Monte Carlo (MCMC) methods have been proposed many years (Metropolis et al. 1953) before their impact in Statistics was truly felt. However, once Gelfand and Smith (1990) stressed the ultimate feasibility of producing a Markov chain with a given stationary distribution f , either via a Gibbs sampler that simulates each conditional distribution of f in its turn, or via a Metropolis–Hastings algorithm based on a proposal $q(y|x)$ with acceptance probability [for a move from x to y]

$$\min \{1, f(y)q(x|y)/f(x)q(y|x)\},$$

then the spectrum of manageable models grew immensely and almost instantaneously.

Due to parallel developments at the time on graphical and hierarchical Bayesian models, like generalised

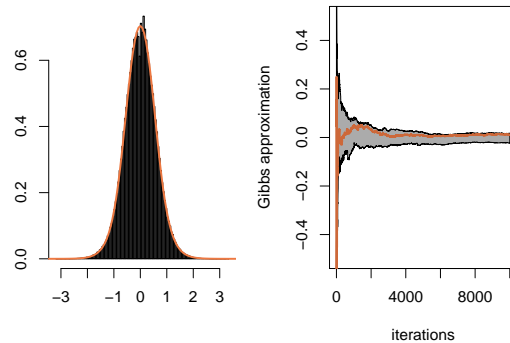


Figure 2: (left) Gibbs sampling approximation to the distribution $f(x) \propto \exp(-x^2/2)/(1+x^2+x^4)$ against the true density; (right) range of convergence of the approximation to $\mathbb{E}_f[X^3] = 0$ against the number of iterations using 100 independent Gibbs runs of the Gibbs sampler, along with a single Gibbs run.

linear mixed models (Zeger and Karim 1991), the wealth of multivariate models with available conditional distributions (and hence the potential of implementing the Gibbs sampler) was far from negligible, especially when the availability of latent variables became quasi universal due to the slice sampling representations (Damien et al. 1999, Neal 2003). (Although the adoption of Gibbs samplers has primarily taken place within Bayesian statistics, there is nothing that prevents an artificial augmentation of the data through such techniques.)

For instance, if the density $f(x) \propto \exp(-x^2/2)/(1+x^2+x^4)$ is known up to a normalising constant, f is the marginal (in x) of the joint distribution $g(x, u) \propto \exp(-x^2/2)\mathbb{I}(u(1+x^2+x^4) \leq 1)$, when u is restricted to $(0, 1)$. The corresponding slice sampler then consists in simulating

$$U|X = x \sim \mathcal{U}(0, 1/(1+x^2+x^4))$$

and

$$X|U = u \sim \mathcal{N}(0, 1)\mathbb{I}(1+x^2+x^4 \leq 1/u),$$

the later being a truncated normal distribution. As shown by Figure 2, the outcome of the resulting Gibbs sampler perfectly fits the target density, while the convergence of the expectation of X^3 under f has a behaviour quite comparable with the iid setting.

While the Gibbs sampler first appears as *the* natural solution to solve a simulation problem in complex models if only because it stems from the true target

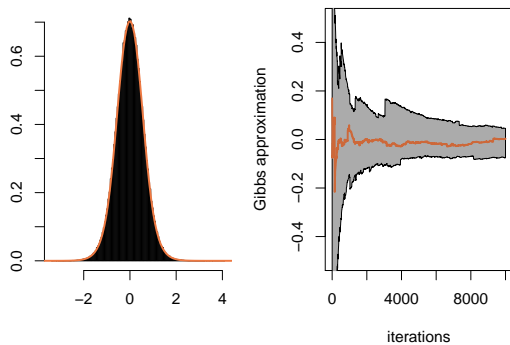


Figure 3: (left) Random walk Metropolis–Hastings sampling approximation to the distribution $f(x) \propto \exp(-x^2/2)/(1+x^2+x^4)$ against the true density for a scale of 1.2 corresponding to an acceptance rate of 0.5; (right) range of convergence of the approximation to $\mathbb{E}_f[X^3] = 0$ against the number of iterations using 100 independent runs of the Metropolis–Hastings sampler, along with a single Metropolis–Hastings run.

f , as exhibited by the widespread use of BUGS Lunn et al. (2000), which mostly focus on this approach, the infinite variations offered by the Metropolis–Hastings schemes offer much more efficient solutions when the proposal $q(y|x)$ is appropriately chosen. The basic choice of a random walk proposal $q(y|x)$ being then a normal density centred in x can be improved by exploiting some features of the target as in Langevin algorithms (see Robert and Casella 2004, section 7.8.5) and Hamiltonian or hybrid alternatives (Duane et al. 1987, Neal 1999) that build upon gradients. More recent proposals include particle learning about the target and sequential improvement of the proposal (Douc et al. 2007, Rosenthal 2007, Andrieu et al. 2010). Figure 3 reproduces Figure 2 for a random walk Metropolis–Hastings algorithm whose scale is calibrated towards an acceptance rate of 0.5. The range of the convergence paths is clearly wider than for the Gibbs sampler, but the fact that this is a generic algorithm applying to any target (instead of a specialised version as for the Gibbs sampler) must be borne in mind.

Another major improvement generated by a statistical imperative is the development of variable dimension generators that stemmed from Bayesian model choice requirements, the most important example being the reversible jump algorithm in Green (1995) which had a significant impact on the study of graphical models (Brooks et al. 2003).

Some uses of Monte Carlo in Statistics

The impact of Monte Carlo methods on Statistics has not been truly felt until the early 1980’s, with the publication of Rubinstein (1981) and Ripley (1987), but Monte Carlo methods have now become invaluable in Statistics because they allow to address optimisation, integration and exploration problems that would otherwise be unreachable. For instance, the calibration of many tests and the derivation of their acceptance regions can only be achieved by simulation techniques. While integration issues are often linked with the Bayesian approach—since Bayes estimates are posterior expectations like

$$\int h(\theta)\pi(\theta|x) d\theta$$

and Bayes tests also involve integration, as mentioned earlier with the Bayes factors—, and optimisation difficulties with the likelihood perspective, this classification is by no way tight—as for instance when likelihoods involve unmanageable integrals—and all fields of Statistics, from design to econometrics, from genomics to psychometry and environmics, have now to rely on Monte Carlo approximations. A whole new range of statistical methodologies have entirely integrated the simulation aspects. Examples include the bootstrap methodology (Efron 1982), where multi-level resampling is not conceivable without a computer, indirect inference (Gouriéroux et al. 1993), which construct a pseudo-likelihood from simulations, MCEM (Cappé and Moulines 2009), where the E-step of the EM algorithm is replaced with a Monte Carlo approximation, or the more recent approximated Bayesian computation (ABC) used in phylogenetics (Beaumont et al. 2002), where the likelihood is not manageable but the underlying model can be simulated from.

In the past fifteen years, the collection of real problems that Statistics can [afford to] handle has truly undergone a quantum leap. Monte Carlo methods and in particular MCMC techniques have forever changed the emphasis from “closed form” solutions to algorithmic ones, expanded our impact to solving “real” applied problems while convincing scientists from other fields that statistical solutions were indeed available, and led us into a world where “exact” may mean “simulated”. The size of the data sets and of the models currently handled thanks to those tools, for example in genomics or in climatol-

ogy, is something that could not have been conceived 60 years ago, when Ulam and von Neumann invented the Monte Carlo method.

References

- ANDRIEU, C., DOUCET, A. and HOLENSTEIN, R. (2010). Particle Markov chain Monte Carlo (with discussion). *J. Royal Statist. Society Series B*, **72**. (to appear).
- BEAUMONT, M., ZHANG, W. and BALDING, D. (2002). Approximate Bayesian computation in population genetics. *Genetics*, **162** 2025–2035.
- BROOKS, S., GIUDICI, P. and ROBERTS, G. (2003). Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions (with discussion). *J. Royal Statist. Society Series B*, **65** 3–55.
- CAPPÉ, O. and MOULINES, E. (2009). On-line expectation-maximization algorithm for latent data models. *J. Royal Statist. Society Series B*, **71**(3) 593–613.
- CHEN, M., SHAO, Q. and IBRAHIM, J. (2000). *Monte Carlo Methods in Bayesian Computation*. Springer-Verlag, New York.
- CHIB, S. (1995). Marginal likelihood from the Gibbs output. *J. American Statist. Assoc.*, **90** 1313–1321.
- 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.
- DOUC, R., GUILLIN, A., MARIN, J.-M. and ROBERT, C. (2007). Convergence of adaptive mixtures of importance sampling schemes. *Ann. Statist.*, **35**(1) 420–448.
- DUANE, S., KENNEDY, A. D., PENDLETON, B. J., and ROWETH, D. (1987). Hybrid Monte Carlo. *Phys. Lett. B*, **195** 216–222.
- EFRON, B. (1982). *The Jackknife, the Bootstrap and Other Resampling Plans*, vol. 38. SIAM, Philadelphia.
- GELFAND, A. and DEY, D. (1994). Bayesian model choice: asymptotics and exact calculations. *J. Royal Statist. Society Series B*, **56** 501–514.
- GELFAND, A. and SMITH, A. (1990). Sampling based approaches to calculating marginal densities. *J. American Statist. Assoc.*, **85** 398–409.
- GOURIÉROUX, C., MONFORT, A. and RENAULT, E. (1993). Indirect inference. *J. Applied Econom.*, **8** 85–118.
- GREEN, P. (1995). Reversible jump MCMC computation and Bayesian model determination. *Biometrika*, **82** 711–732.
- KENDALL, W., MARIN, J.-M. and ROBERT, C. (2007). Confidence bands for Brownian motion and applications to Monte Carlo simulations. *Statistics and Computing*, **17** 1–10.
- LUNN, D., THOMAS, A., BEST, N., and SPIEGELHALTER, D. (2000). WinBUGS – a Bayesian modelling framework: concepts, structure, and extensibility. *Statistics and Computing*, **10** 325–337.
- METROPOLIS, N., ROSENBLUTH, A., ROSENBLUTH, M., TELLER, A. and TELLER, E. (1953). Equations of state calculations by fast computing machines. *J. Chem. Phys.*, **21** 1087–1092.
- NEAL, R. (1999). *Bayesian Learning for Neural Networks*, vol. 118. Springer-Verlag, New York. Lecture Notes.
- NEAL, R. (2003). Slice sampling (with discussion). *Ann. Statist.*, **31** 705–767.
- NEWTON, M. and RAFTERY, A. (1994). Approximate Bayesian inference by the weighted likelihood bootstrap (with discussion). *J. Royal Statist. Society Series B*, **56** 1–48.
- RIPLEY, B. (1987). *Stochastic Simulation*. John Wiley, New York.
- ROBERT, C. and CASELLA, G. (2004). *Monte Carlo Statistical Methods*. 2nd ed. Springer-Verlag, New York.
- ROBERT, C. and CASELLA, G. (2010). *Introducing Monte Carlo Methods with R*. Springer-Verlag, New York.
- ROSENTHAL, J. (2007). AMCM: An R interface for adaptive MCMC. *Comput. Statist. Data Analysis*, **51** 5467–5470.
- RUBINSTEIN, R. (1981). *Simulation and the Monte Carlo Method*. John Wiley, New York.
- RUE, H., MARTINO, S. and CHOPIN, N. (2008). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations (with discussion). *J. Royal Statist. Society Series B*, **71** (2) 319–392.
- ZEGER, S. and KARIM, R. (1991). Generalized linear models with random effects; a Gibbs sampling approach. *J. American Statist. Assoc.*, **86** 79–86.