Recyclage dans les méthodes d'acceptation-rejet

George Casella et Christian P. Robert
Cornell University, Ithaca, NY
URA CNRS 1378, Université de Rouen and CREST, Insee, Paris

BU-1281-M April 1995

Statistiques/*Statistics*

Recyclage dans les méthodes d'acceptation-rejet

George Casella et Christian P. Robert
Cornell University, Ithaca, NY
URA CNRS 1378, Université de Rouen and CREST, Insee, Paris

`Version 1.0 -- April 14, 1995`
`file ar4.tex`

Résumé – La méthode d'acceptation-rejet, à la base de nombreux algorithmes de simulation, rejette une partie des variables aléatoires qu'elle produit. Nous montrons dans cette note comment un recyclage des valeurs rejetées permet de diminuer la variance de l'approximation d'une intégrale arbitraire, sans pour autant augmenter les temps de calcul.

Recycling rejected values in accept-reject methods

Abstract – *Accept-reject algorithms generate more random variables than they use. We show in this note how the rejected values can significantly contribute to a reduction of the variance of the empirical estimate of an arbitrary integral by proposing alternate weights for each observation. Moreover, the resulting improvement does not require additional computing time.*

Version française abrégée. — Lorsqu'une intégrale

$$I = \int h(x) f(x) dx$$

est approchée par la méthode de Monte-Carlo, on utilise en général un échantillon $x_1, \ldots,$ x_n tiré suivant la loi de densité f pour estimer I par la moyenne empirique

$$\delta^{AR} = \frac{1}{n} \sum_{i=1}^{n} h(x_i).$$

La simulation étant souvent opérée au travers de l'algorithme d'acceptation-rejet, elle induit la génération supplémentaire de $y_1, \ldots, y_t$ qui ne sont pas pris en compte dans l'estimation de I. En effet, la méthode d'acceptation-rejet est fondée sur le résultat suivant (voir [2], [5] ou [6]) : s'il existe une densité g et une constante M telles que $f \leq Mg$, et si la variable x résulte de la génération de couples (z_i, u_i), distribués respectivement suivant g et la loi uniforme sur $[0,1]$, jusqu'à ce que $u_i \leq f(z_i)/Mg(z_i)$, x est distribuée suivant f si le support de f est inclus dans celui de g. Chaque génération d'une valeur suivant cette méthode produit en moyenne $(M-1)$ rejets. Nous noterons $(z_1, \ldots, z_{n+t})$ l'échantillon global et $\rho = 1/M$ la probabilité d'acceptation d'une variable générée suivant g.

De par leur rejet, les valeurs rejetées apportent une information sur f et donc sur I qui est négligée par la méthode d'acceptation-rejet. Nous allons formaliser le recyclage des valeurs rejetées en montrant qu'une pondération adéquate des $h(y_i)$ permet de réduire la variance de l'estimateur de I. Remarquons tout d'abord que δ^{AR} peut également s'écrire

$$\delta^{AR} = \frac{1}{n} \sum_{i=1}^{n+t} h(z_i) \mathbb{I}_{\{u_i \leq f(z_i)/Mg(z_i)\}}$$

où $\mathbb{I}$ désigne la fonction indicatrice. On notera ω_i le rapport $f(z_i)/Mg(z_i)$. Se fondant sur cette représentation, [1] propose une amélioration uniforme (au sens de la variance) de δ^{AR} en intégrant les u_i conditionnellement aux z_i mais l'estimateur résultant induit une augmentation conséquente du temps de calcul. L'alternative que nous proposons ici consiste à n'intégrer chaque fonction indicatrice que conditionnellement au z_i correspondant, donc à proposer l'estimateur

$$\begin{aligned}
\delta^{IAR} &= \frac{1}{n} \sum_{1}^{n+t} \mathbb{E}[\mathbb{I}_{u_i \leq \omega_i} | z_i] h(z_i) = \frac{1}{n} \sum_{1}^{n+t} P(u_i \leq \omega_i | z_i) h(z_i) \\
&= \frac{1}{n} \left(\sum_{1}^{n+t-1} \left[1 + \frac{t(g(z_i) - \rho f(z_i))}{(n+t-1)(1-\rho) f(z_i)} \right]^{-1} h(z_i) + h(x_n) \right)
\end{aligned}$$

que l'on obtient par calcul de la loi marginale du couple (z_i, u_i).

On montre alors que, conditionnellement à t, l'estimateur δ^{IAR} a une variance inférieure à celle de δ^{AR}, les deux estimateurs étant sans biais. Il est donc préférable d'utiliser δ^{IAR} puisque la pondération des $h(z_i)$ ne fait intervenir que le rapport ω_i, déjà calculé pour déterminer l'acceptation de z_i. Les simulations menées dans les cas d'une loi de Student simulée à partir d'une loi de Cauchy et d'une loi gamma simulée à partir d'une loi gamma à indice entier montrent que l'amélioration apportée par δ^{IAR} peut être substantielle.

1. The Accept–Reject Algorithm.– The accept-reject method is undoubtedly one of the most commonly used simulation methods (see [2] and [5]). It produces samples from a given distribution with density f through the simulation from another distribution with density g such that $f(x) \leq Mg(x)$, based on the following result:

Lemma 1.1. – *Given a sequence* $z_1, z_2, \ldots$ *of i.i.d. random variables with density* g *and an independent sequence* $u_1, u_2, \ldots$ *of i.i.d. uniform* $\mathcal{U}([0,1])$ *variables, the random variable* x *equal to the first* z_i *such that* $u_i \leq f(z_i)/M\, g(z_i)$ *is distributed according to* f. *Moreover, the number* N *of random variables* z_i *required to produce* x *is a geometric* $\mathcal{G}eo(\rho)$ *random variable, with* $\rho = 1/M$.

This mathematically straightforward result has been exploited in clever ways to produce most of the random number generators of standard distributions at very little cost in term of computing time. In particular, the standard random generators have acceptance probabilities uniformly bounded away from 0 and usually very close to 1 (see [2]). However, with the development of more powerful simulations devices such as Markov chain Monte-Carlo methods ([2], [6]), accept-reject algorithms are increasingly used in non-standard setups with less accurate bounds M and a corresponding increase of the rejection rate. When the sole goal of the simulation is to provide an approximation of the expectation

$$I = \mathbb{E}_f[h(x)] = \int h(x)f(x)dx,$$

rather than a truly i.i.d. sample, the rejected values can be put to use through an estimator which improves upon the original accept-reject estimator.

This improvement is to be understood as a post-processing of the accept-reject algorithm output, the simulation part *per se* being fixed. The modification proposed in this paper is therefore a statistical rather than an algorithmic refining of the original accept-reject method. We denote by $(x_1, \ldots, x_n)$ the accept-reject sample and by $(z_1, \ldots, z_{n+t})$ the overall sample. In particular, $z_{n+t} = x_n$ and t is distributed as a negative binomial $\mathcal{N}eg(n, \rho)$ random variable on $\mathbb{N}$ (i.e. $P(t = 0) = \rho^n$). If we denote by ω_i the ratio $f(z_i)/Mg(z_i)$, it is straightforward to compute the marginal density of z_i $(i \neq n+t)$ conditionally on t,

$$m(z) = \frac{n-1}{n+t-1} f(z) + \frac{t}{n+t-1} \frac{g(z) - \rho f(z)}{1-\rho} \tag{1.1}$$

and the distribution of u_i conditional on z_i and t is

$$u_i | z_i, t \sim \frac{(1-\rho)(n-1)\mathbb{I}_{[0,\omega_i]}(u_i) + \rho t \mathbb{I}_{[\omega_i,1]}(u_i)}{(1-\rho)(n-1)\omega_i + \rho t(1-\omega_i)}.$$

The distribution of z_{n+t} is f and, when $t = 0$, every z_i is also distributed according to f. A lengthy but straightforward computation provides in addition the density of (u_i, u_j, z_i, z_j) $(1 \leq i, j < n+t, i \neq j)$,

$$\frac{g(z_i)g(z_j)}{(n+t-1)(n+t-2)} \left[\mathbb{I}_{u_i \leq \omega_i} \mathbb{I}_{u_j \leq \omega_j} \frac{(n-1)(n-2)}{\rho^2} + \left\{ \mathbb{I}_{u_i \leq \omega_i} \mathbb{I}_{u_j > \omega_j} + \right. \right. \tag{1.2}$$
$$\left. \left. \mathbb{I}_{u_j \leq \omega_j} \mathbb{I}_{u_i > \omega_i} \right\} \frac{(n-1)t}{\rho(1-\rho)} + \mathbb{I}_{u_i > \omega_i} \mathbb{I}_{u_j > \omega_j} \frac{t(t-1)}{(1-\rho)^2} \right],$$

which is of use in the following improvement over the accept-reject estimator.

2. Improving on the Accept–Reject Estimator.– An improvement upon the standard accept-reject estimator,

$$\delta^{AR} = \frac{1}{n}\sum_{i=1}^{n} h(x_i) = \frac{1}{n}\sum_{i=1}^{n+t} \mathbb{I}_{u_i \le \omega_i} h(z_i)$$

has been proposed in [1] by integrating out the uniform random variables u_i. While the improvement directly follows from the partial integration, which reduces the variance while conserving the unbiasedness, the resulting estimator is quite complex with weights depending on the whole sample and of complexity of order n^2.

A general alternative method, which allows the inclusion of the rejected values, is importance sampling. In fixed sample sizes setups, this method weights each observation z_i by the ratio $f(z_i)/(n+t)g(z_i)$ (see [5]). In our case, the random nature of t and the marginal distribution (1.1) imply the modification of the importance sampling weights into

$$\frac{f(z_i)}{(n+t)m(z_i)} = \frac{n+t-1}{(n+t)(n-1)}\left[1 + \frac{t(g(z_i) - \rho f(z_i))}{(n-1)(1-\rho)f(z_i)}\right]^{-1}$$

for $i \neq n+t$ and into $1/(n+t)$ for $i = n+t$. A very similar estimator can be derived from δ^{AR} by calculating a termwise conditional expectation, conditioning each term on z_i:

$$\begin{aligned}
\delta^{IAR} &= \frac{1}{n}\sum_{1}^{n+t} \mathbb{E}[\mathbb{I}_{u_i \le \omega_i} | z_i] h(z_i) \\
&= \frac{1}{n}\left(\sum_{i=1}^{n+t-1}\left[1 + \frac{t(g(z_i) - \rho f(z_i))}{(n-1)(1-\rho)f(z_i)}\right]^{-1} h(z_i) + h(z_{n+t})\right) \\
&= \frac{1}{n}\left(h(z_{n+t} + \frac{n-1}{n+t-1}\sum_{i=1}^{n+t-1}\frac{f(z_i)}{m(z_i)}h(z_i)\right)
\end{aligned} \tag{1.3}$$

where the last equality follows from (1.1). The difference from the modified importance sampling estimator is that the random ratio $\frac{n+t-1}{(n+t)(n-1)}$ is replaced by $1/n$, which is constant. We will thus consider δ^{IAR} as a potential substitute for δ^{AR}. Since both estimators are unbiased, domination of δ^{AR} by δ^{IAR} in term of squared error loss follows from an ordering of the variances.

Proposition 2.1. – *For every function h,*

$$\mathrm{var}(\delta^{AR}) \ge \mathrm{var}(\delta^{IAR})$$

conditional on the number of rejected variables, t.

Proof. – Since both estimators are unbiased, we can assume without loss of generality that $\mathbb{E}_f[h(x)] = 0$. Note that this implies that $\mathbb{E}[\mathbb{I}_{u_i \le \omega_i} h(z_i)] = 0$, where the expectation is taken with respect to the marginal distribution $m(z_i)$.

The variance of δ^{AR} is derived from the joint distribution (1.2), as

$$\begin{aligned}
\operatorname{var}(\delta^{AR}) &= \frac{1}{n^2} \sum_{i,j=1}^{n+t} \mathbb{E}[\mathbb{I}_{u_i \le \omega_i} \mathbb{I}_{u_j \le \omega_j} h(z_i) h(z_j)] \\
&= \frac{1}{n^2} \left\{ \mathbb{E}_f[h^2(x)] + \sum_{i=1}^{n+t-1} \left[\sum_{1 \le j \ne i \le n+t-1} \mathbb{E}[\mathbb{I}_{u_i \le \omega_i} \mathbb{I}_{u_j \le \omega_j} h(z_i) h(z_j)] \right.\right. \\
&\qquad \left.\left. + \mathbb{E}[\mathbb{I}_{u_i \le \omega_i} h^2(z_i)] \right] \right\} \\
&= \frac{1}{n^2} \left\{ \mathbb{E}_f[h^2(x)] + \sum_{i=1}^{n+t-1} \left(\sum_{j \ne i} \frac{(n-1)(n-2)\mathbb{E}_f[h(x)]^2}{(n+t-1)(n+t-2)} + \mathbb{E}[\mathbb{I}_{u_i \le \omega_i} h^2(z_i)] \right) \right\} \\
&= \frac{1}{n^2} \left\{ \mathbb{E}_f[h^2(x)] + (n+t-1) \mathbb{E}[\mathbb{I}_{u_i \le \omega_i} h^2(z_i)] \right\},
\end{aligned}$$

where we have used the fact that the last variable z_{n+t} is independent of the other variables and

$$\begin{aligned}
\mathbb{E}[\mathbb{I}_{u_i \le \omega_i} \mathbb{I}_{u_j \le \omega_j} h(z_i) h(z_j)] &= \int \frac{\omega_i \omega_j}{\rho^2} \frac{(n-1)(n-2)}{(n+t-1)(n+t-2)} h(z_i) h(z_j) g(z_i) g(z_j) dz_i dz_j \\
&= \int \frac{(n-1)(n-2)}{(n+t-1)(n+t-2)} h(z_i) h(z_j) f(z_i) f(z_j) dz_i dz_j = 0.
\end{aligned}$$

If we denote the weight function in δ^{IAR} by

$$b(z) = \left(1 + \frac{t(g(z) - \rho f(z))}{(n-1)(1-\rho) f(z)} \right)^{-1} = \frac{n-1}{n+t-1} \frac{f(z)}{m(z)},$$

the variance of δ^{IAR} conditional on t can be expressed as

$$\frac{1}{n^2} \left\{ \mathbb{E}_f[h^2(x)] + \sum_{i=1}^{n+t-1} \left[\mathbb{E}[b^2(z_i) h^2(z_i)] + \sum_{1 \le j \ne i \le n+t-1} \mathbb{E}[b(z_i) h(z_i) b(z_j) h(z_j)] \right] \right\}.$$

where again $\mathbb{E}[b(z_i)h(z_i)] = 0$, where the expectation is taken with respect to the marginal distribution $m(z_i)$. For $i \neq j$,

$$\begin{aligned}
\mathbb{E}[b(z_i)h(z_i)b(z_j)h(z_j)] &= \frac{1}{(n+t-1)(n+t-2)} \int b(z_1)b(z_2)h(z_1)h(z_2)\times \\
&\left\{(n-1)f(z_1)\left[(n-2)f(z_2) + t\frac{g(z_2)-\rho f(z_2)}{1-\rho}\right] + \right. \\
&\left. t\frac{g(z_1)-\rho f(z_1)}{1-\rho}\left[(n-1)f(z_2) + (t-1)\frac{g(z_2)-\rho f(z_2)}{1-\rho}\right]\right\} dz_1 dz_2 \\
&= \frac{(n-1)^2}{(n+t-1)(n+t-2)}\left[\int h(z)f(z)dz\right]^2 \\
&- \frac{1}{(n+t-1)(n+t-2)} \int b(z_1)b(z_2)h(z_1)h(z_2)\times \\
&\left\{(n-1)f(z_1)f(z_2) + t\frac{g(z_1)-\rho f(z_1)}{1-\rho}\frac{g(z_2)-\rho f(z_2)}{1-\rho}\right\} dz_1 dz_2 \\
&= -\frac{1}{(n+t-1)(n+t-2)}\left\{(n-1)\mathbb{E}_f[b(x)h(x)]^2\right. \\
+ \quad & \left. t\mathbb{E}_f\left[b(x)h(x)\frac{g(x)-\rho f(x)}{(1-\rho)f(x)}\right]^2\right\},
\end{aligned}$$

The covariance between $b(z_i)h(z_i)$ and $b(z_j)h(z_j)$ is therefore negative uniformly in h. Moreover, for $i < n+t$, $\mathbb{E}[b^2(z_i)h^2(z_i)]$ is bounded by

$$\mathbb{E}[b^2(z_i)h^2(z_i)] = \mathbb{E}[\{\mathbb{E}\mathbb{I}_{u_i \leq \omega_i}\}^2 h^2(z_i)] \leq \mathbb{E}[\mathbb{I}_{u_i \leq \omega_i} h^2(z_i)]$$

by the Cauchy-Schwarz inequality. Therefore,

$$\mathrm{var}(\delta^{IAR}) \leq \mathrm{var}(\delta^{AR}|t).$$

■■

3. Magnitude of the improvement.– The above result is somehow surprising: a standard estimation rule can be uniformly improved by a simple alternative. Although there are some precedents of such improvements using Rao-Blackwellisation, they are usually associated with the Gibbs sampler ([3], [4]). Rao-Blackwellisation has been applied to accept-reject sampling ([1]), but the resulting estimator has a complex form, leading to the concern that the complexity of the calculation might outweigh the potential improvement (see [1] for a discussion). But here the form of the dominating estimator is so simple that one may wonder about the practical value of such an improvement. We therefore conclude this paper with an example that quantifies the amount of improvement brought by δ^{IAR} upon δ^{AR}.

Example 3.1. – We simulate a Student $\mathcal{T}_3(0,1)$ distribution from a Cauchy $\mathcal{C}(0,1)$ distribution, i.e.

$$f(x) = \frac{\Gamma(2)}{\sqrt{3\pi}\Gamma(3/2)} \left(1 + x^2/3\right)^{-2}, \qquad g(x) = \frac{1}{\pi} \left(1 + x^2\right)^{-1}.$$

In this case, M is equal to $3\sqrt{3}/4$ and the Cauchy random variables can be directly simulated by inverting the cdf of the Cauchy distribution. The following table compares the performances in terms of squared error loss of δ^{AR} and δ^{IAR} for the functions

$$h_1(x) = x, \qquad h_2(x) = x^2 \quad \text{and} \quad h_3(x) = \mathbb{I}_{x \geq 1.96}$$

and for different values of n:

	t	10	25	50	75
δ^{AR}	h_1	.148	.0617	.0309	.0195
	h_2	.863	.315	.167	.116
	h_3	$3.92 10^{-3}$	$1.52 10^{-3}$	$8.15 10^{-4}$	$5.36 10^{-5}$
δ^{IAR}	h_1	.105	.0405	.0202	.0126
	h_2	.211	.066	.033	.022
	h_3	$2.19 10^{-3}$	$7.64 10^{-4}$	$3.74 10^{-4}$	$2.58 10^{-5}$

Table 3.1. – Monte-Carlo evaluation of the mean squared error of δ^{AR} and δ^{IAR} (1,000 simulations).

The improvement brought by δ^{IAR} is therefore significant since δ^{IAR} decreases the error by at least 30%, with decreases up to 80% in the case of h_2.

Acknowledgment. This research was supported by NSF Grant No. DMS93-05547, and NSF/CNRS Grant No.INT-9216784. This is technical report BU-1281-M in the Biometrics Unit, Cornell University, Ithaca NY 14853.

Références bibliographiques

[1] G. Casella et C.P. Robert, Rao-Blackwellisation of sampling schemes, *Biometrika*, 1995, à paraître.

[2] L. Devroye, *Non-Uniform Random Variate Generation*, New York, Springer-Verlag.

[3] A. Gelfand et A.F.M. Smith, Sampling based approaches to calculating marginal densities, *Journal of the American Statistical Society*, 85, 1990, p. 398–409.

[4] J. S. Liu, W.H. Wong, et A. Kong, Covariance structure of the Gibbs sampler with applications to the comparison of estimators and augmentation schemes, *Biometrika*, 81, 1994, 27-40.

[5] B. Ripley, *Stochastic Simulation*, New York, J. Wiley, 1987.

[6] C.P. Robert, *L'Analyse Statistique Bayésienne*, Paris, Economica, 1992.