A note on generating random variables with log-concave densities

Luc Devroye
School of Computer Science
McGill University

January 22, 2012

Abstract. We present a black-box style rejection method that is valid for generating random variables with any log-concave density, provided that one knows a mode of the density, which is known only up to a constant factor.

Keywords and phrases. Random variate generation. Simulation. Monte Carlo method. Expected time analysis. Log-concavity. Rejection method.

1991 Mathematics Subject Classifications: Primary 65C10. Secondary 65C05, 11K45, 68U20.

Address: School of Computer Science, McGill University, 3480 University Street, Montreal, Canada H3A 2K6. Research sponsored by NSERC Grant A3456. Email: `lucdevroye@gmail.com`

Introduction

In 1984, the author of this note published a black-box style rejection algorithm that can be used for any log-concave density f on the real line for which the location of a mode m is known. It is based on the inequality

$$\frac{1}{f(m)} f\left(m + x/f(m)\right) \le \min\left(1, e^{1-|x|}\right).$$

An extension to discrete distributions on the integers followed in 1987. The integral over the bounding curve is precisely 4, so that we have a uniform performance guarantee for the rejection method. That method falls apart when f is proportional to a density, and the proportionality or normalization constant in unknown. In many applications, one does not know that constant.

There have been scattered attempts at addressing this or realted problems. They come in several categories:

(i) One could apply so-called table methods, adaptive or not, to break the line up into intervals on which we can to a good job bounding f. See, e.g., the book of Devroye (1986, chapter 8) or the articles of Gilks (1992) and Gilks and Wild (1992, 1993), who developed the ARS algorithm, which stands for adaptive rejection sampling. These methods work well, but there is no universally applicable explicit bound on the performance. The original paper of Gilks and Wild (1992) required also the derivative of f at various places (in black-box style), but Gilks (1992) did not need that. See also an exercise in Devroye (1986, chapter 8.2) about adaptive rejction sampling. In any case, we are not aware of any uniform performance bounds for any of these methods. Gilks, Best and Tan (1995) published ARMS, or adaptive Metropolis rejection sampling. It applies to non-log-concave densities but is only approximate, given a finite set of resources.

(ii) One can modify the bounds of Devroye (1984) if the normalization constant is known to be between two bounds. Still, in a black-box set-up, one would have to compute these bounds, and do so with explicit guarantees. This is equivalent to estimating $\int f$ uniformly over all log-concave densities to within a constant multiplicative factor. In a distant sense, the present paper indirectly does this.

(iii) Hörmann, Leydold and Derflinger (2004) [see also Hörmann and Leydold, 2000, 2001] started an ambitious project of creating black box automated random variate generators. In some of that work, where it touches on log-concave densities, they assume the availability of the derivative of f in black-box format. For example, they automate the method of finding a threshold beyond which an exponentially decreasing cap is used, and perform an optimization along the way. Their methods are fast and efficient, but a uniform performance guarantee is still missing.

In this note, we give a simple algorithm for all log-concave densities for which the normalization constant is not known. It has a one-time set-up cost of $O(1 + |\log_2 R|)$ where R is the ratio between the used normalization constant and the unknown true normalization constant.

Simple inequalities for log-concave functions

Consider a nonincreasing nonnegative log-concave function f on $[0,\infty)$, and let $0 \le a < b < \infty$. Define the line through $(a, \log f(a))$ and $(b, \log f(b))$,

$$g(x) = \frac{b-x}{b-a} \log f(a) + \frac{x-a}{b-a} \log f(b).$$

By log-concavity,

$$\log f(x) \begin{cases} \ge g(x) & \text{if } a \le x \le b \ , \\ \le g(x) & \text{if } x \le a \text{ or } x \ge b. \end{cases}$$

These trivial inequalities suffice for our development. The inequality $f \le e^g$ will be used on $[b,\infty)$ in what follows. We will require the following integrals:

$$\int_a^b \exp(g(x))\,dx = \frac{b-a}{\log\left(\frac{f(a)}{f(b)}\right)}\ (f(a) - f(b))$$

$$\int_b^\infty \exp(g(x))\,dx = \frac{b-a}{\log\left(\frac{f(a)}{f(b)}\right)}\ f(b).$$

A monotone log-concave density on the positive halfline

Log-concave densities have either one mode m or an interval of modes. In the latter case, let m be any canonical choice in that interval. We assume throughout the paper that a mode m is known. By a mere translation, we can and do assume that $m = 0$. Let f be proportional to a density and be supported on $[0,\infty)$. Since $m = 0$, f is necessarily nonincreasing on $\mathbb{R}^+$. The following recipe leads to an acceptable algorithm—it is easy to improve on the constants, but at this moment, simplicity is our main concern.

Using linear search among the values $2^i/f(0)$, where i runs through all integers, positive and negative, starting with $i = 0$, and moving up or down by one depending upon the situation, find the largest value a of the form 2^i such that simultaneously

$$\frac{f(a)}{f(0)} \ge \frac{1}{4} \ge \frac{f(2a)}{f(0)}.$$

Since $f(x) \downarrow 0$ as $x \to \infty$ and f is monotone and log-concave, such an a exists and is unique. This search is a one-time cost that has complexity $\Theta(1 + |\log_2 a|)$. Note that $\max(a, 1/a)$ roughly measures how far off we are in terms of the normalization constant. We will provide more details about the binary search in the next section.

We use the following bounds to set up a rejection method:

$$f(x) \le h(x) \stackrel{\text{def}}{=} \begin{cases} f(0) & 0 \le x \le a, \\ f(a) & a \le x \le 2a, \\ \exp\left(\frac{2a-x}{a} \log f(a) + \frac{x-a}{a} \log f(2a)\right) & x \ge 2a. \end{cases} \qquad (1)$$

Note that if $f(2a) = 0$, then $\log f(2a) = -\infty$, and therefore, the last entry in (1) is formally zero. A random variable with a density proportional to the upper bound can easily be obtained. On $[0,a]$ and

$[a, 2a]$, it suffices to use aU and $a(1+U)$, respectively, where U is uniform $[0, 1]$. On $[2a, \infty)$, one can use $2a + aE/\log(f(a)/f(2a))$ (and thus 0 if $f(2a) = 0$), where E is a standard exponential random variate.

The integrals of the upper bound over the three pieces are required to obtain the proper mixture:

$$\int_0^a h = af(0), \int_a^{2a} h = af(a), \int_{2a}^{\infty} h = \frac{af(2a)}{\log\left(\frac{f(a)}{f(2a)}\right)}.$$

For the sake of completeness, we state the rejection algorithm in its entirety.

[f is proportional to a log-concave density with mode at 0]
set-up: find $a > 0$ such that $f(a)/f(0) \geq 1/4 \geq f(2a)/f(0)$
set $s = af(0) + af(a) + af(2a)/\log\left(\frac{f(a)}{f(2a)}\right)$
repeat generate U, V, W uniform $[0, 1]$
if $sV \leq af(0)$, then set $X \leftarrow aU$
else if $sV \leq af(0) + af(a)$, then set $X \leftarrow a + aU$
else set $X \leftarrow 2a + a\log(1/U)/\log(f(a)/f(2a))$
until $Wh(X) \leq f(X)$ where h is defined in (1)
return X

The returned X has density proportional to f. We now show that the expected number of iterations in the rejection algorithm is bounded by 5, uniformly over all log-concave densities on $\mathbb{R}^+$ with mode $m = 0$.

Proof. The expected number of iterations is

$$\frac{\int_0^{\infty} h(x)\,dx}{\int_0^{\infty} f(x)\,dx.}$$

We have $\int h = s$, where s is defined in the first line of the algorithm. For $\int f$, we use this lower bound:

$$\int_0^{\infty} f(x)\,dx \geq af(a) + \int_a^{2a} f(x)\,dx \geq af(a) + \int_a^{2a} \exp(g(x))\,dx = af(a) + \frac{a}{\log\left(\frac{f(a)}{f(2a)}\right)}\left(f(a) - f(2a)\right).$$

Log-concavity implies

$$\sqrt{f(2a)f(0)} \leq f(a),$$

and thus

$$\frac{f(2a)}{f(a)} \leq \sqrt{\frac{f(2a)}{f(0)}} \leq \sqrt{\frac{1}{4}} = \frac{1}{2}.$$

Using $f(a) - f(2a) \geq f(2a)$, we obtain the lower bound

$$\int_0^{\infty} f(x)\,dx \geq af(a) + \frac{af(2a)}{\log\left(\frac{f(a)}{f(2a)}\right)}.$$

Putting things together,

$$
\begin{aligned}
\frac{\int_0^\infty h(x)\,dx}{\int_0^\infty f(x)\,dx} &\le \frac{af(0)+af(a)+\frac{af(2a)}{\log\left(\frac{f(a)}{f(2a)}\right)}}{af(a)+\frac{af(2a)}{\log\left(\frac{f(a)}{f(2a)}\right)}} \\
&\le 1+\frac{af(0)}{af(a)+\frac{af(2a)}{\log\left(\frac{f(a)}{f(2a)}\right)}} \\
&\le 1+\frac{f(0)}{f(a)} \\
&\le 5.
\end{aligned}
$$

In other words, the algorithm is uniformly fast over all log-concave densities in our class. □

The binary search

The binary search for a can be organized so that the one-time set-up cost it entails has a small cost. To be precise, let the true density be $f^* = Rf$ with normalization constant R unknown. It is assumed that the user has adjusted f as much as possible to make R as near to one as posible. So, $\max(R, 1/R)$ measures the ignorance of the normalization constant. Let us start the binary search at $1/f(0)$, exploring the sequence $2^i/f(0)$ starting at $i = 0$ and going left or right as necessary.

With this understanding, the number of steps taken in the binary search is bounded by

$$
O(1) + |\log_2 R|,
$$

where $O(1)$ is a universal constant.

PROOF Let a be the output of the search, and the value used in the rejection algorihm. We show that

$$
\frac{2\log 4}{1+4\log 4} \le af^*(0) \le \log(4e).
$$

The fact that any a is inversely proportional to $1/f^*(0)$ to within a universally constant factor will be important. By an inequality of Devroye (1984),

$$
f^*(x/f^*(0)) \le f^*(0)\exp(1-|x|).
$$

Set $x = af^*(0)$ and note that

$$
\frac{1}{4} \le \frac{f(a)}{f(0)} \le \frac{f^*(a)}{f^*(0)} \le e^{1-af^*(0)}.
$$

This implies the second inequality. Next,

$$
\int_{2a}^\infty f^* \le \int_{2a}^\infty e^g = 2af^*(0)\frac{f^*(2a)/f^*(0)}{\log(f^*(0)/f^*(2a))} \le 2af^*(0)\frac{1}{4\log 4},
$$

where g is defined in terms of f^* rather than f. Thus,

$$
2af^*(0) \ge 1 - \int_{2a}^\infty f^* \ge 1 - \frac{2af^*(0)}{4\log 4},
$$

i.e.,

$$2af^*(0) \geq \frac{4\log 4}{1+4\log 4}.$$

This proves the leftmost inequality. So, if the binary search is started at $1/f(0) = R/f^*(0)$, then at most a constant plus $|\log_2 R|$ steps suffice to find an a as required. $\square$

In conclusion, while the expected time of the algorithm is uniformly bounded over all log-concave densities, the one time set-up, which consists of the binary search, has a non-uniform component.

Log-concave densities in general

For general unimodal f with a mode at the origin, one can split the problem in two carefully. Instead of three dominating pieces, one can take three to the right of m, and three to the left. This requires separate thresholds $a > 0$ for R^+ and $b < 0$ for R^-. The weights of the six pieces must be computed, as we did in the algorithm above. Using the notation $f \le h$ as in the previous section, we see that the rejection constant is

$$\frac{\int h}{\int f} = \frac{\int_{-\infty}^0 h + \int_0^\infty h}{\int_{-\infty}^0 f + \int_0^\infty f} \le \max\left(\frac{\int_{-\infty}^0 h}{\int_{-\infty}^0 f}, \frac{\int_0^\infty h}{\int_0^\infty f}\right) \le 5.$$

We summarize, at the risk of being too repetitive. We begin with the definition of the function h:

$$f(x) \le h(x) \stackrel{\text{def}}{=} \begin{cases} \exp\left(\frac{2b-x}{b}\log f(b) + \frac{x-b}{b}\log f(2b)\right) & x \le 2b \\ f(b) & 2b \le x \le b, \\ f(0) & b \le x \le 0, \\ f(0) & 0 \le x \le a, \\ f(a) & a \le x \le 2a, \\ \exp\left(\frac{2a-x}{a}\log f(a) + \frac{x-a}{a}\log f(2a)\right) & x \ge 2a. \end{cases} \tag{2}$$

Then the algorithm is as follows:

> [f is proportional to a log-concave density on the real line]
> set-up: find $a > 0$ such that $f(a)/f(0) \ge 1/4 \ge f(2a)/f(0)$
>   find $b < 0$ such that $f(b)/f(0) \ge 1/4 \ge f(2b)/f(0)$
> $s_+ \leftarrow af(0) + af(a) + af(2a)/\log\left(\frac{f(a)}{f(2a)}\right)$
> $s_- \leftarrow |b|f(0) + |b|f(b) + |b|f(2b)/\log\left(\frac{f(b)}{f(2b)}\right)$
> $s \leftarrow s_+ + s_-$
> repeat generate U, V, W uniform $[0,1]$
>   if $sV \le af(0)$, then set $X \leftarrow aU$
>   if $sV \le (a+|b|)f(0)$, then set $X \leftarrow b + (a-b)U$
>   else if $sV \le (a+|b|)f(0) + af(a)$, then set $X \leftarrow a + aU$
>   else if $sV \le (a+|b|)f(0) + af(a) + |b|f(b)$, then set $X \leftarrow b + bU$
>   else if $sV \le (a+|b|)f(0) + af(a) + |b|f(b) + af(2a)/\log\left(\frac{f(a)}{f(2a)}\right)$,
>   set $X \leftarrow 2a + a\log(1/U)/\log(f(a)/f(2a))$
>   else set $X \leftarrow 2b + b\log(1/U)/\log(f(b)/f(2b))$
> until $Wh(X) \le f(X)$ where h is defined in (2)
> return X

A few final remarks are in order. There is an obvious generalization to the discrete case, when the random variables are atomic and are supported on the integers. The details are left to the reader. The referee noted that in parcatice, on finite word size computers, the binary search will terminate after a finite number of steps. This would mean that certain densities cannot be dealt with. These correspond roughly to those for which the normalization constant is off by more than 2^{31} on 32-bit floating point computers. The paper was written in the assumption that all real-value arithmetic can be performed with infnite accuracy.

Acknowledgment

We thank Daniel Sheldon (Oregon State University) for prodding us to come up with an explicit solution for this problem. We also also grateful to the referee for a very enlightening report.

References

L. Devroye, "A simple algorithm for generating random variates with a log-concave density," *Computing*, vol. 33, pp. 247–257, 1984.

L. Devroye, *Non-Uniform Random Variate Generation*, Springer-Verlag, New York, 1986.

L. Devroye, "A simple generator for discrete log-concave distributions," *Computing*, vol. 39, pp. 87–91, 1987.

W. R. Gilks, "Derivative-free adaptive rejection sampling for Gibbs sampling," in: *Bayesian Statistics 4*, edited by J. Bernardo, J. Berger, A. P. Dawid, and A. F. M. Smith, Oxford University Press, 1992.

W. R. Gilks and P. Wild, *Adaptive rejection sampling for Gibbs sampling*, vol. 41, pp. 337–148, Applied Statistics, 1992.

W. R. Gilks and P. Wild, *Algorithm AS 287: Adaptive Rejection Sampling from Log-Concave density function*, vol. 41, pp. 701–709, Applied Statistics, 1993.

W. R. Gilks, N. G. Best, and K. K. C. Tan, *Adaptive rejection Metropolis sampling*, vol. 44, pp. 455–472, Applied Statistics, 1995.

W. Hörmann, J. Leydold, and G. Derflinger, *Automatic Nonuniform Random Variate Generation*, Springer-Verlag, Berlin, 2004.

J. Leydold and W. Hörmann, "Black box algorithms for generating non-uniform continuous random variates," in: *COMPSTAT 2000*, edited by W. Jansen and J. G. Bethlehem, pp. 53–54, 2000.

J. Leydold and W. Hörmann, "Universal algorithms as an alternative for generating non-uniform continuous random variates," in: *G.I. Schuler, P.D. Spanos*, in: *Monte Carlo Simulation*, pp. 177–183, 2001.