JOURNAL OF COMPUTATIONAL PHYSICS 88, 484–489 (1990)

Note

An Algorithm for the Generation of Random Numbers with Density $C \exp(-\lambda|x|^{\nu})$

1. INTRODUCTION

The numerical simulation of statistical processes requires simple and efficient algorithms to generate random numbers with given distributions. In principle, one can always generate numbers x with any desired distribution by a "direct method" writing $x = F^{-1}(u)$, where F^{-1} is the inverse of the cumulative distribution function and u is uniformly distributed in the unit interval. However, in most cases an adequate numerical fit to F^{-1} is not available, and one must resort to general algorithms based on acceptance–rejection tests. The performance of such algorithms depends, in turn, on the possibility of generating efficiently the candidates for the test with a distribution which resembles closely the one in which one is interested on.

A notable example in which these limitations can be circumvented is that of the normal distribution, for which the well-known polar method [1] provides a direct algorithm which avoids the use of F^{-1} in favour of a simple transformation of two uniform random variables. The simplicity of the method makes it a very convenient and popular algorithm, at least when CPU time is not a crucial limitation.

In this paper we present an extension of the polar method to the case of exponential density functions of the form

$$f_{\nu}(x) = C_{\nu} e^{-\lambda|x|^{\nu}}; \qquad C_{\nu} = \frac{\nu \lambda^{1/\nu}}{2\Gamma(1/\nu)}, \tag{1.1}$$

where $\nu \geqslant 1$ is a real constant. Density functions of this form are of interest in a variety of statistical processes and numerical simulations. In particular, the case with $\nu = 4$ arises in the Monte Carlo simulation of the scalar sector of four-dimensional field theories [2], while larger values of ν are relevant for the simulation of lower dimensional theories. To extend the polar method to density functions of this form one must solve a non linear differential equation which has no closed form analytic solution for values of ν other than 1 or 2. Yet, the solutions are easily approximated by power expansions which allow a simple implementation of the method for any value of ν.

The method is presented in Section 2, while in Section 3 we discuss briefly its performance relative to other alternative algorithms.

0021-9991/90 $3.00

2. Derivation of the Method

As is well known, in the polar method for sampling the normal distribution [1] two independent normally distributed random numbers x and y (with zero mean and variance equal to 2) are obtained from

$$x=\sqrt{-\log u_1}\,\sin(2\pi u_2); \qquad y=\sqrt{-\log u_1}\,\cos(2\pi u_2), \tag{2.1}$$

where u_1 and u_2 are uniformly distributed in the unit interval.

For integer values of ν, $\nu>1$, one can extend this method to the density function f_ν in Eq. (1.1) by constructing adequately generalized polar coordinates in ν dimensions—generated from ν uniformly distributed random numbers—whose projection onto the cartesian axes yield ν independent random numbers each with density f_ν. Although one can show that such a construction exists, it generalizes Eqs. (2.1) into transformations of increasing complexity which are unlikely to be of any practical use. Nevertheless, one can use that formal construction as a guideline and consider at the end only one of the random variables obtained in that way while the rest is discarded. At the same time, one can analytically continue the result to non integer values of ν.

Proceeding in this way (and setting for simplicity $\lambda=1$ in Eq. (1.1)), for the sampling of f_ν one is lead to consider, for any $\nu>1$, the following transformation

$$x=(-\log u_1)^{1/\nu}S_\nu(u_2), \tag{2.2}$$

where $S_\nu(u)$ satisfies

$$\frac{dS_\nu}{du}=\frac{2\pi}{\nu\sin(\pi/\nu)}(1-|S_\nu|^\nu)^{1/\nu}; \qquad S_\nu\left(\frac{1}{2}\right)=0. \tag{2.3}$$

The functions S_ν are plotted in Fig. 1 for various values of ν. They are antisymmetric about $u=\frac{1}{2}$, and satisfy $S_\nu(1)=-S_\nu(0)=1$. To prove that this transformation produces the desired results, we compute directly the corresponding probability density function. Assume first that $x\leqslant 0$, so that $S_\nu\leqslant 0$ and $u_2\leqslant\frac{1}{2}$. From Eq. (2.2) follows that the cumulative distribution function $F(x)$ receives contributions from all $u_2\in[0,\frac{1}{2}]$ and, for each u_2, from all $u_1\in[0,\bar{u}_1]$, where $\bar{u}_1=\exp[-x/S_\nu(u_2)]^\nu$. Thus, for $F(x)$ we have

$$F(x)=\int_0^{1/2}du_2\int_0^{\bar{u}_1}du_1=\int_0^{1/2}du_2\exp[-x/S_\nu(u_2)]^\nu=\frac{\nu\sin(\pi/\nu)}{2\pi}\int_{-1}^{0}ds\,\frac{e^{-(x/s)^\nu}}{(1-|s|^\nu)^{1/\nu}},$$

where in the last expression we used Eq. (2.3) to change the integration variable from u_2 to $s=S_\nu(u_2)$. Differentiating with respect to x, and changing again the integration variable to $t=|x|^\nu(1/|s|^\nu-1)$, we get for the density function

$$f(x)=\frac{\nu\sin(\pi/\nu)}{2\pi}e^{-|x|^\nu}\int_0^\infty dt\,t^{-1/\nu}e^{-t}=\frac{\nu\sin(\pi/\nu)}{2\pi}\Gamma(1-1/\nu)\,e^{-|x|^\nu}$$

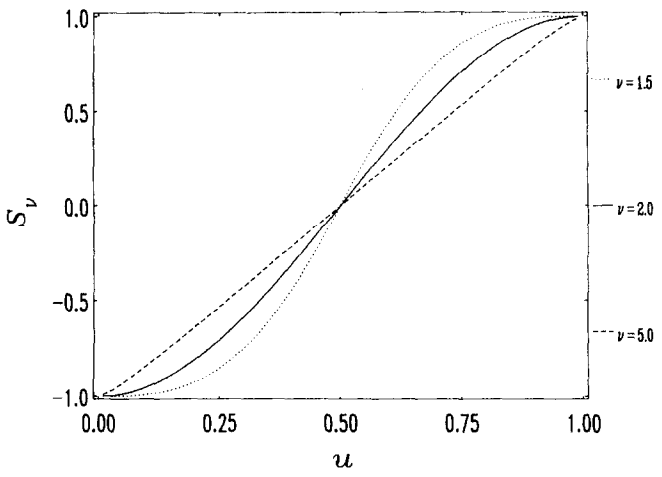

FIG. 1. The functions S_ν for various values of ν.

which is the same as $f_\nu(x)$ in Eq. (1.1). The calculation is done in a similar way for $x \geqslant 0$.

For $\nu \to 1$, $S_\nu(u)$ converges to a random sign, and (2.2) reduces to $x = F^{-1}$, F being the cumulative distribution function. For $\nu = 2$ the solution to Eq. (2.3) is $S_\nu(u) = \sin \pi(u - \frac{1}{2})$, which is equivalent to the first transformation in Eq. (2.1) of the polar method. For other integer values of ν, Eq. (2.3) can only be integrated to give the inverse function $u = u(s)$, but this cannot be inverted in closed analytic form. Similarly, for non integer values of ν, Eq. (2.3) cannot be integrated in closed form. The non-existence of a closed form expression for S_ν is of course no surprise, and for practical purposes is irrelevant provided one can find a convenient and efficient algorithm for the numerical evaluation of S_ν. In that respect, at this point we offer only a preliminary analysis based on power expansions, whose purpose is to allow a first evaluation of the method, although they possibly do not provide an optimal implementation.

Although tedious, it is straightforward to obtain power expansions for $S_\nu(u)$ directly from Eq. (2.3), either about $u = \frac{1}{2}$ ($S_\nu = 0$) or about $u - \frac{1}{2} = \pm\frac{1}{2}$ ($S_\nu = \pm 1$). One finds, about $u = \frac{1}{2}$,

$$S_\nu(u) = x\left\{1 + \sum_{k=1}^{\infty} \alpha_k y^k\right\} \tag{2.4a}$$

where

$$x = \frac{2\pi}{\nu \sin(\pi/\nu)}\left(u - \frac{1}{2}\right); \qquad y = x^\nu \tag{2.4b}$$

while, about $u-\frac{1}{2}=\pm\frac{1}{2}$,

$$S_{\nu}(u)=\operatorname{sign}\left(u-\frac{1}{2}\right)\left\{1-\sum_{k=1}^{\infty}\beta_k z^k\right\} \tag{2.5a}$$

where

$$z=\left\{\frac{2\pi(\nu-1)}{\nu^{2-1/\nu}\sin(\pi/\nu)}\left(\frac{1}{2}-\left|u-\frac{1}{2}\right|\right)\right\}^{\nu/(\nu-1)}. \tag{2.5b}$$

The coefficients α_k and β_k are listed in Table I for $k\leqslant 4$ and any value of ν. Truncating these series to a given order n, one finds that S_{ν} can be approximated by the first series for $|u-\frac{1}{2}|\leqslant c(\nu)+a$, and by the second for $|u-\frac{1}{2}|\geqslant c(\nu)-a$, where $c(\nu)$ is a slowly varying function of ν that separates the domains of each approximation, and $a>0$ is a finite interval of overlap of the two domains. Using $n=4$ in a single precision FORTRAN code, the error of this approximation does not exceed one part in 10^5 for $\nu\geqslant 1.2$, and this is certainly sufficient for most statistical simulation experiments. Since both approximations overlap in a fairly wide interval, the choice of the transition point which separates both approximation is not critical. For $n=4$ we found it convenient to fit it to

$$c(\nu)=0.5-1.16\,e^{-1.2\nu}-0.175\,e^{-0.08\nu}\qquad(\nu\geqslant 1.2). \tag{2.6}$$

TABLE I

Coefficients α_k and β_k of the Power Series Expansions for $S_{\nu}(u)$ (Eqs. (2.4), (2.5))

$$\alpha_1=\frac{1}{\nu}\frac{1}{\nu+1}$$

$$\alpha_2=\frac{1}{2!\,\nu^2}\frac{\nu^2-2\nu-1}{2\nu^2+3\nu+1}$$

$$\alpha_3=\frac{1}{3!\,\nu^3}\frac{4\nu^5-11\nu^4-2\nu^3+13\nu^2+7\nu+1}{6\nu^4+17\nu^3+17\nu^2+7\nu+1}$$

$$\alpha_4=\frac{1}{4!\,\nu^4}\frac{36\nu^8-124\nu^7+27\nu^6+206\nu^5-15\nu^4-140\nu^3-71\nu^2-14\nu-1}{24\nu^6+98\nu^5+159\nu^4+130\nu^3+56\nu^2+12\nu+1}$$

$$\beta_1=1$$

$$\beta_2=-\frac{1}{2!}\frac{\nu-1}{2\nu-1}$$

$$\beta_3=\frac{1}{3!}\frac{\nu^4-\nu^3-3\nu^2+4\nu-1}{12\nu^3-16\nu^2+7\nu-1}$$

$$\beta_4=-\frac{1}{4!}\frac{4\nu^6-20\nu^5+32\nu^4-11\nu^3-11\nu^2+7\nu-1}{96\nu^5-200\nu^4+164\nu^3-66\nu^2+13\nu-1}$$

3. Discussion

As a first evaluation of the method we briefly discuss its performance relative to two other alternatives, namely an accept reject algorithm and Forsythe's general method for the sampling of exponential densities [3]. In doing this, we implemented our method using the series (2.4) and (2.5) truncated at $k=4$, which delivers a performance roughly independent of the value of v although it is subject to further optimization.

To sample the density f_v with an efficient accept–reject algorithm, one can use candidates with a gaussian density if $v \geqslant 2$, or with exponential density if $v < 2$. However, in either case one must do some preliminary numerical work, for each chosen value of v, in order to adjust the variance of the candidates so as to optimize the acceptance ratio. The resulting optimized acceptance ratio equals one at $v = 1$ and 2, and decreases as v departs from those values (reaches 0.76 as $v \to 2^-$, and 0.48 as $v \to \infty$). CPU time measurements show that for $v > 4$ our method is faster, by a factor approaching 1.6 as $v \to \infty$, while for v close to 1 or 2 (from above) the accept–reject algorithm is more efficient, by a factor approaching 1.3. Thus, compared with this alternative, in its present form our method is competitive and provides a modest but significant improvement of the efficiency over a wide range of values of v.

The density functions f_v to which the present method applies are a particular case of the more general form $f(x) = \exp(-h(x))$, with $h'(x) > 0$, for which Forsythe's method [3] provides a generic sampling algorithm. As in the previous case, this method requires preliminary numerical work to generate tables of numerical constants which depend on the function $h(x)$. More importantly, the use of those tables makes it quite sensitive to round off errors. When applied to the density functions f_v, we found that Forsythe's method coded in single precision is roughly 10% faster than ours, but it can only reproduce the first few moments of the distribution (for example, in the range $2 \leqslant v \leqslant 6$, experiments with 10^4 samples exhibit deviations which are significant already at the fourth moment). In contrast to this, our method was tested in high statistics experiments, with over 10^6 samples, and no significant deviations were observed for a wide range of values of v.

Summarizing, we have presented a method for the sampling of density functions of the form $C\exp(-\lambda|x|^v)$, which is based on a generalization of the polar method for the gaussian distribution. As in the gaussian case, the desired samples are obtained as a functional transformation of two uniformly distributed random numbers. Although one of the functions involved in the transformation is not elementary, its numerical evaluation does not seem to present any difficulties and, in particular, low-order polynomials provide an approximation adequate for most statistical applications. A preliminary evaluation of the method shows that its overall performance compares favorably with other standard alternatives, regarding both speed and precision. A more detailed analysis leading to an optimized numerical implementation of the method may further improve its performance.

ACKNOWLEDGMENTS

This work was partially supported by DIUC under Grants 50/86 and 48/87, and by FONDECYT under Grants 8042/86 and 677/87. We thank I. Huerta for useful discussions.

REFERENCES

1. G. E. P. BOX AND M. E. MULLER, *Ann. Math. Statist.* **29**, 610 (1958).
2. B. FREEDMAN, D. SMOLENSKY, AND D. WEINGARTEN, *Phys. Lett. B* **113**, 481 (1982); D. J. E. CALLAWAY AND R. PETRONZIO, *Nucl. Phys. B* **240** [FS12], 577 (1984); M. G. DO AMARAL AND R. C. SHELLARD, *Phys. Lett. B* **171**, 285 (1986).
3. G. E. FORSYTHE, *Math. Comput.* **26**, 817 (1972).

RECEIVED: December 13, 1988

N. BRALIĆ
R. ESPINOSA
C. SAAVEDRA

Facultad de Física
Pontificia Universidad Católica de Chile
Casilla 6177, Santiago 22, Chile