



A new rejection sampling method for truncated multivariate Gaussian random variables restricted to convex sets

Hassan Maatouk, Xavier Bay

► To cite this version:

Hassan Maatouk, Xavier Bay. A new rejection sampling method for truncated multivariate Gaussian random variables restricted to convex sets. Ronald Cools and Dirk Nuyens. Monte Carlo and Quasi-Monte Carlo Methods , 163, pp.521-530, 2016, .

HAL Id: hal-01063978

<https://hal.archives-ouvertes.fr/hal-01063978>

Submitted on 15 Sep 2014

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.

A New Rejection Sampling Method for Truncated Multivariate Gaussian Random Variables Restricted to Convex Sets

Hassan Maatouk and Xavier Bay

Abstract Statistical researchers have shown increasing interest in generating truncated multivariate normal distributions. In this paper, we only assume that the acceptance region is convex and we focus on rejection sampling. We propose a new algorithm that outperforms crude rejection method for the simulation of truncated multivariate Gaussian random variables. The proposed algorithm is based on a generalization of Von Neumann's rejection technique which requires the determination of the mode of the truncated multivariate density function. We provide a theoretical upper bound for the ratio of the target probability density function over the proposal probability density function. The simulation results show that the method is especially efficient when the probability of the multivariate normal distribution of being inside the acceptance region is low.

1 Introduction

The need for simulation of truncated multivariate normal distributions appears in many fields, like Bayesian inference for truncated parameter space [11] and [12], Gaussian processes for computer experiments subject to inequality constraints [6], [9] and [10] and regression models with linear constraints (see e.g. [13] and [27]). In general, we have two types of methods. The first ones are based on Markov chain Monte Carlo (MCMC) simulation [4], [18] and [24], as the Gibbs sampling [3], [13], [15], [17], [19], [23] and [25]. They provide samples from an approximate distribution which converges asymptotically to the true one. The second ones are exact simulation methods based on rejection sampling (Von Neumann [26]) and its extensions, [7], [16] and [18]. In this paper, we focus on the second type of methods.

Hassan Maatouk · Xavier Bay
Ecole Nationale Supérieure des Mines de St-Etienne, 158 Cours Fauriel, Saint-Etienne, France
e-mail: hassan.maatouk@mines-stetienne.fr, e-mail: bay@emse.fr

Recently, researchers in statistics have used an adaptive rejection technique with Gibbs sampling [13], [14], [20], [21] and [23]. Let us mention that in one dimension rejection sampling with a high acceptance rate has been developed by Robert [23], and Geweke [13]. In [23] Robert developed simulation algorithms for one-sided and two-sided truncated normal distributions. Its rejection algorithm is based on the uniform distribution. The multidimensional case where the acceptance region is a convex subset of \mathbb{R}^d is based on the same algorithm using the Gibbs sampling to reduce the simulation problem to a sequence of one-dimensional simulations. In this case, the method requires the determination of slices of the convex acceptance region. Also, Geweke [13] proposed an exponential rejection sampling to simulate a truncated normal variable. The multidimensional case is deduced by using the Gibbs algorithm. In one-dimension, Chopin [5] designed an algorithm that is computationally faster than alternative algorithms. A multidimensional rejection sampling to simulate a truncated Gaussian vector outside arbitrary ellipsoids has been developed by Ellis and Maitra [8]. For higher dimensions, Philippe and Robert [22] developed a simulation method of a Gaussian distribution restricted to positive quadrants. Also, Botts [1] improves an accept-reject algorithm to simulate positive multivariate normal distributions.

In this article, we develop a new rejection technique to simulate a truncated multivariate normal distribution restricted to any convex subset of \mathbb{R}^d . The method only requires the determination of the mode of the probability density function (pdf) restricted to the acceptance region. We provide a theoretical upper bound for the ratio of the target probability density function over the proposal probability density function.

The article is organized as follows. In Section 2, we recall the rejection method. Then, we present our new method, called *rejection sampling from the mode* (RSM) and we give the main theoretical results and the associated algorithm. In Section 3, we compare RSM with existing rejection algorithms.

2 Multivariate Normal Distribution

2.1 The General Rejection Method

Let f be a probability density function (pdf) defined on \mathbb{R}^d . Von Neumann [26] proposed the rejection method, using the notion of dominating density function. Suppose that g is another density function close to f such that for some finite constant $c \geq 1$, called rejection constant,

$$f(x) \leq cg(x), \quad x \in \mathbb{R}^d. \quad (1)$$

The acceptance/rejection method is an algorithm for generating random samples from f by drawing from the proposal pdf g and the uniform distribution.

Theorem 1 (Rejection Sampling Algorithm, Von Neumann [26]). *Suppose that f and g are two pdfs such that $f(x) \leq cg(x)$ for all x in the support of f . Then the random variable X resulting from the following algorithm is distributed according to f .*

1. Generate X with density g .
2. Generate U uniformly on $[0, 1]$. If $cg(X)U \leq f(X)$, accept X ; otherwise, go back to step 1.

Furthermore it can be shown that the acceptance rate is equal to $1/c$. In practice it is crucial to get a small c .

Notice that the rejection sampling algorithm is immediately extended to pseudo-density functions (i.e. positive function with finite integral), avoiding the computation of normalizing constant.

Corollary 1. *Let \mathcal{C} be a subset of \mathbb{R}^d and \tilde{f} and \tilde{g} be two pseudo-density functions on \mathcal{C} such that $\tilde{f}(x) \leq k\tilde{g}(x)$. Then the algorithm in Lemma 1 is still valid if the inequality condition $cg(X)U \leq f(X)$ is replaced by*

$$k\tilde{g}(X)U \leq \tilde{f}(X). \quad (2)$$

The rejection constant is $c = k \frac{\int_{\mathcal{C}} \tilde{g}(t) dt}{\int_{\mathcal{C}} \tilde{f}(t) dt}$.

Proof. We have $\tilde{f}(x) \leq k\tilde{g}(x)$, and so

$$f(x) = \frac{\tilde{f}(x)}{\int_{\mathcal{C}} \tilde{f}(t) dt} \leq c \frac{\tilde{g}(x)}{\int_{\mathcal{C}} \tilde{g}(t) dt} = cg(x), \quad (3)$$

with $c = k \frac{\int_{\mathcal{C}} \tilde{g}(t) dt}{\int_{\mathcal{C}} \tilde{f}(t) dt}$. The condition $cg(X)U \leq f(X)$ is equivalent to $k\tilde{g}(X)U \leq \tilde{f}(X)$. \square

2.2 Rejection Sampling from the Mode

Suppose that X has multivariate normal distribution with probability density function:

$$f(x | \mu, \Sigma) = \frac{1}{(2\pi)^{d/2} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(x - \mu)^\top \Sigma^{-1}(x - \mu)\right), \quad x \in \mathbb{R}^d \quad (4)$$

where $\mu = E[X]$ and Σ is the covariance matrix, assumed to be invertible.

We consider a convex subset \mathcal{C} of \mathbb{R}^d representing the acceptance region. We assume that μ does not belong to \mathcal{C} , which is a hard case for crude rejection sampling. Furthermore, as explained in Remark 1 (see below) the proposed method is

not different from crude rejection sampling if $\mu \in \mathcal{C}$. Without loss of generality, let $\mu = 0$. Our aim is to simulate the multivariate normal distribution X restricted to the convex set \mathcal{C} . The idea is twofold. Firstly, we determine the mode μ^* corresponding to the maximum of the probability density function f restricted to \mathcal{C} . It is the solution of the following convex optimization problem:

$$\mu^* = \arg \min_{x \in \mathcal{C}} \frac{1}{2} x^\top \Sigma^{-1} x. \quad (5)$$

Secondly, let g be the pdf obtained from f by shifting the center to μ^* :

$$g(x | \mu^*, \Sigma) = \frac{1}{(2\pi)^{d/2} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(x - \mu^*)^\top \Sigma^{-1}(x - \mu^*)\right). \quad (6)$$

Then we prove in the next two theorems that g can be used as a proposal pdf for rejection sampling on \mathcal{C} , and we derive the optimal constant.

Theorem 2. *Let \tilde{f} and \tilde{g} be the pseudo-density functions defined as*

$$\tilde{f}(x) = f(x | 0, \Sigma) \mathbb{1}_{x \in \mathcal{C}} \text{ and } \tilde{g}(x) = g(x | \mu^*, \Sigma) \mathbb{1}_{x \in \mathcal{C}},$$

where f and g are defined in (4) and (6). Then there exists k such that $\tilde{f}(x) \leq k\tilde{g}(x)$ for all x in \mathcal{C} and the smallest value of k is:

$$k^* = \exp\left(-\frac{1}{2}(\mu^*)^\top \Sigma^{-1} \mu^*\right). \quad (7)$$

Proof. Let us start with the one-dimensional case. Without loss of generality, we suppose that $\mathcal{C} = [\mu^*, +\infty[$, where μ^* is positive and $\Sigma = \sigma^2$. In this case, the condition $\tilde{f}(x) \leq k\tilde{g}$ is written

$$\forall x \geq \mu^*, e^{-\frac{x^2}{2\sigma^2}} \leq k e^{-\frac{(x-\mu^*)^2}{2\sigma^2}}, \quad (8)$$

and so

$$k^* = e^{\frac{(\mu^*)^2}{2\sigma^2}} \max_{x \geq \mu^*} e^{-\frac{x\mu^*}{\sigma^2}} = e^{\frac{(\mu^*)^2}{2\sigma^2}} e^{-\min_{x \geq \mu^*} \frac{x\mu^*}{\sigma^2}} = e^{-\frac{(\mu^*)^2}{2\sigma^2}}. \quad (9)$$

In the multidimensional case, we have $k^* = \max_{x \in \mathcal{C}} e^{\frac{1}{2}(\mu^*)^\top \Sigma^{-1} \mu^* - x^\top \Sigma^{-1} \mu^*}$. Since $\mu^* \in \mathcal{C}$, we only need to show that

$$\forall x \in \mathcal{C}, x^\top \Sigma^{-1} \mu^* \geq (\mu^*)^\top \Sigma^{-1} \mu^*. \quad (10)$$

The angle between the gradient vector $\Sigma^{-1} \mu^*$ of the function $\frac{1}{2} x^\top \Sigma^{-1} x$ at the mode μ^* and the vector $(x - \mu^*)$ is acute for all x in \mathcal{C} since \mathcal{C} is convex (see Figure 1). Therefore, $(x - \mu^*)^\top \Sigma^{-1} \mu^*$ is non-negative for all x in \mathcal{C} . \square

By now, we can write algorithm 1 as follows:

Theorem 3 (RSM Algorithm). *Let \tilde{f} and \tilde{g} be the pseudo-density functions defined as*




Fig. 1 Scalar product between the gradient vector $\Sigma^{-1}\mu^*$ of the function $\frac{1}{2}x^{\top}\Sigma^{-1}x$ at μ^* and $(x - \mu^*)$. The blue lines are the level curves of the function $x \mapsto \frac{1}{2}x^{\top}\Sigma^{-1}x$.

$$\tilde{f}(x) = f(x | 0, \Sigma)\mathbb{1}_{x \in \mathcal{C}} \quad \text{and} \quad \tilde{g}(x) = g(x | \mu^*, \Sigma)\mathbb{1}_{x \in \mathcal{C}},$$

where f and g are defined by (4), (5) and (6). Then the random vector X resulting from the following algorithm is distributed according to \tilde{f} .

1. Generate X with pseudo-density \tilde{g} .
2. Generate U uniformly on $[0, 1]$. If $U \leq \exp((\mu^*)^{\top}\Sigma^{-1}\mu^* - X^{\top}\Sigma^{-1}\mu^*)$, accept X ; otherwise go back to step 1.

Proof. We applied Corollary 1 with the optimal constant k^* of Theorem 2. The inequality condition (2) is equivalent to

$$U \leq e^{\frac{1}{2}(\mu^*)^{\top}\Sigma^{-1}\mu^*} e^{-\frac{1}{2}X^{\top}\Sigma^{-1}X} e^{\frac{1}{2}(X-\mu^*)^{\top}\Sigma^{-1}(X-\mu^*)}, \quad (11)$$

which is equivalent to

$$U \leq \exp\left((\mu^*)^{\top}\Sigma^{-1}\mu^* - X^{\top}\Sigma^{-1}\mu^*\right). \quad (12)$$

□

Remark 1. In practice, we use a crude rejection method to simulate X with pseudo-density \tilde{g} in the RSM algorithm. So if $\mu \in \mathcal{C}$, RSM degenerates to crude rejection sampling since $\mu^* = \mu$ and $f = g$. Therefore, the method RSM can be seen as a generalization of rejection sampling.

Remark 2. Our method requires only the maximum likelihood of the pdf restricted to the acceptance region. It is the mode of the truncated multivariate normal distribution. The numerical calculation of it is a standard problem in the minimization of positive quadratic forms subject to linear constraints, see e.g. [2].

3 Performance Comparisons

To investigate the performance of the RSM algorithm, we consider a zero-mean bivariate Gaussian random vector x with covariance matrix Σ , equal to $\begin{pmatrix} 4 & 2.5 \\ 2.5 & 2 \end{pmatrix}$. Assume that the convex set $\mathcal{C} \in \mathbb{R}^2$ is defined by the inequality constraints:

$$-10 \leq x_2 \leq 0 \quad \text{and} \quad x_1 \geq -15, \quad 5x_1 - x_2 + 15 \leq 0. \quad (13)$$

It is the acceptance region used in Figures 2 and 3. By minimizing a quadratic form subject to linear constraints, we find the mode

$$\mu^* = \arg \min_{x \in \mathcal{C}} \frac{1}{2} x^\top \Sigma^{-1} x \approx (-3.4, -2.0), \quad (14)$$

and then we compare crude rejection sampling to RSM.




Fig. 2 Crude rejection sampling using 2000 simulations. The acceptance rate is 3%.




Fig. 3 Rejection sampling from the mode using 2000 simulations. The acceptance rate is 21%.

In Figure 2, we use crude rejection sampling in 2000 simulations of a $\mathcal{N}(0, \Sigma)$. Given the number of points in \mathcal{C} (red points), it is clear that the algorithm is not efficient. The reason is that the mean of the bivariate normal distribution is outside the acceptance region. In Figure 3, we first simulate from the shifted distribution centered at the mode with same covariance matrix Σ (step one of the RSM algorithm). Now in the second step of the RSM algorithm, we have two types of points (red and

black ones) in the convex set \mathcal{C} . The black points are in \mathcal{C} but do not respect the inequality constraint (12). The red points are in \mathcal{C} , and respect (12). We observe that RSM outperforms crude rejection sampling, with acceptance rate of 21% against 3%.

Table 1 Comparison between crude rejection sampling and RSM when the probability to be inside the acceptance region becomes low. The acceptance region is $\mathcal{C} = [\mu^-, +\infty[$.

μ^-	Average of acceptance rate with crude rejection sampling (%)	Average of acceptance rate with RSM (%)	Gain
0.5	30.8	34.9	1.1
1	15.8	26.2	1.6
1.5	6.7	20.5	3.0
2	2.2	16.8	7.4
2.5	0.6	14.2	23.1
3	0.1	12.2	92.0
3.5	0.0	10.6	455.6
4	0.0	9.3	2936.7
4.5	0.0	8.4	14166.0

The performance of the method appears when the probability to be inside the acceptance region is low. In Table 1, we consider the one dimensional case $d = 1$ and we only change the position of μ^- . From the last column, we observe that our algorithm outperforms crude rejection sampling. For instance, the proposed algorithm is approximately 14000 times faster than the crude rejection sampling when the acceptance region is $[4.5, +\infty[$. Note also that the acceptance rate remains stable for large μ^- (near 10%) for the RSM method whereas it decreases rapidly to zero for crude rejection sampling.

Table 2 Comparison of average acceptance rate between Robert's method [23] and RSM under the variability of the distance between μ^- and μ^+ . The acceptance region is $\mathcal{C} = [\mu^-, \mu^+]$, where μ^- is fixed to 1.

$\mu^+ - \mu^-$	Robert's method (%)	Rejection sampling from the mode (%)	Gain
0.5	77.8	18.0	0.2
1	56.4	21.2	0.3
2	35.0	27.4	0.7
5	11.6	28.2	2.4
10	7.0	28.4	4.0

Robert [23] also proposed a rejection sampling method in the one dimensional case. To compare the acceptance rates of RSM with Robert's method, we consider a standard normal variable truncated between μ^- and μ^+ with μ^- fixed to 1. In Robert's method, the average acceptance rate is high when the acceptance interval is small (see Table 2.2 in [23]). In the proposed algorithm, simulating from shifted distributions (first step in the RSM algorithm) leads to the fact that the average acceptance rate is more important when the acceptance interval is large. As expected, the performance of the proposed algorithm appears when we have a large gap between μ^- and μ^+ , as shown in Table 2.

Table 3 Comparison between crude rejection sampling and RSM with respect to the dimension d . The acceptance region is $\mathcal{C} = [\mu^-, +\infty[^d$.

Dimension d	μ^-	Average of acceptance rate with crude rejection sampling (%)	Average of acceptance rate with RSM (%)	Gain
1	2.33	1.0	15.0	15.0
2	1.29	1.0	5.2	5.2
3	0.79	1.0	2.5	2.5
4	0.48	1.0	1.5	1.5
5	0.25	1.0	1.2	1.2

Now we investigate the influence of the problem dimension d . We simulate a standard multivariate normal distribution X restricted to $\mathcal{C} = [\mu^-, +\infty[^d$, where μ^- is chosen such that $P(X \in \mathcal{C}) = 0.01$. The mean of the multivariate normal distribution is outside the acceptance region. Simulation of truncated normal distributions in multidimensional cases is a difficult problem for rejection algorithms. From Table 3, we can remark that when the dimension increases, the parameter μ^- tends to zero. Hence, the mode $\mu^* = (\mu^-, \dots, \mu^-)$ tends to the zero-mean of the Gaussian vector X . And so, the acceptance rate of the proposed method converges to the acceptance rate of the crude rejection sampling. As an additional example when μ^- is fixed to 1.35 and $d = 5$, the RSM algorithm is 135 times faster than the crude rejection sampling. In that case, the probability of the normal distribution X being inside the acceptance region is low.

4 Conclusion

In this paper, we develop a new rejection technique, called RSM, to simulate a truncated multivariate normal distribution restricted to any convex set. Our method only requires to find the mode of the target probability density function restricted to the convex acceptance region. The proposal density function in the RSM algorithm is the shifted target distribution centered at the mode. We provide a theoretical formula of the optimal constant such that the proposal density function is as close as

possible to the target density. Note that the RSM algorithm is easy to implement. An illustrative example to compare RSM with crude rejection sampling is included. The simulation results show that using rejection sampling from the mode is more efficient than crude rejection sampling. Comparisons with Robert's method in the one dimensional case is discussed. The RSM method outperforms Robert's method when the acceptance interval is large and the probability of the normal distribution to be inside is low. The proposed rejection method has been applied in the case where the acceptance region is a convex subset of \mathbb{R}^d , and can be extended to non-convex regions by using the convex hull.

Acknowledgements The authors wish to thank Olivier Roustant, Laurence Grammont and Yann Richet for helpful discussions, as well as the participants of MCQMC conference.

References

1. Botts, C.: An accept-reject algorithm for the positive multivariate normal distribution. *Computational Statistics* **28**(4), 1749–1773 (2013)
2. Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press, New York, NY, USA (2004)
3. Breslaw, J.: Random sampling from a truncated multivariate normal distribution. *Applied Mathematics Letters* **7**(1), 1 – 6 (1994)
4. Casella, G., George, E.I.: Explaining the Gibbs sampler. *The American Statistician* **46**(3), 167–174 (1992)
5. Chopin, N.: Fast Simulation of Truncated Gaussian Distributions. *Statistics and Computing* **21**(2), 275–288 (2011)
6. Da Veiga, S., Marrel, A.: Gaussian process modeling with inequality constraints. *Annales de la faculté des sciences de Toulouse* **21**(3), 529–555 (2012)
7. Devroye, L.: *Non-Uniform Random Variate Generation*. Springer-Verlag (1986)
8. Ellis, N., Maitra, R.: Multivariate Gaussian Simulation Outside Arbitrary Ellipsoids. *Journal of Computational and Graphical Statistics* **16**(3), 692–708 (2007)
9. Emery, X., Arroyo, D., Peláez, M.: Simulating Large Gaussian Random Vectors Subject to Inequality Constraints by Gibbs Sampling. *Mathematical Geosciences* pp. 1–19 (2013)
10. Freulon, X., Fouquet, C.: Conditioning a Gaussian model with inequalities. In: A. Soares (ed.) *Geostatistics Tróia '92, Quantitative Geology and Geostatistics*, vol. 5, pp. 201–212. Springer Netherlands (1993)
11. Gelfand, A.E., Smith, A.F.M., Lee, T.M.: Bayesian Analysis of Constrained Parameter and Truncated Data Problems Using Gibbs Sampling. *Journal of the American Statistical Association* **87**(418), 523–532 (1992)
12. Geweke, J.: Exact Inference in the Inequality Constrained Normal Linear Regression Model. *Journal of Applied Econometrics* **1**(2), 127–141 (1986)
13. Geweke, J.: Efficient Simulation from the Multivariate Normal and Student-t Distributions Subject to Linear Constraints and the Evaluation of Constraint Probabilities. In: *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, pp. 571–578 (1991)
14. Gilks, W.R., Wild, P.: Adaptive Rejection Sampling for Gibbs Sampling. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **41**(2), 337–348 (1992)
15. Griffiths, W.E.: A Gibbs sampler for the parameters of a truncated multivariate normal distribution. Department of Economics - Working Papers Series 856, The University of Melbourne (2002)

16. Hörmann, W., Leydold, J., Derflinger, G.: Automatic Nonuniform Random Variate Generation. *Statistics and Computing*. Springer (2004)
17. Kotecha, J.H., Djuric, P.: Gibbs sampling approach for generation of truncated multivariate Gaussian random variables. In: *IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 3, pp. 1757–1760 (1999)
18. Laud, P.W., Damien, P., Shively, T.S.: Sampling Some Truncated Distributions Via Rejection Algorithms. *Communications in Statistics - Simulation and Computation* **39**(6), 1111–1121 (2010)
19. Li, Y., Ghosh, S.K.: Efficient sampling method for truncated multivariate normal and Student t-distribution subject to linear inequality constraints URL http://www.stat.ncsu.edu/information/library/papers/mimeo2649_Li.pdf
20. Martino, L., Miguez, J.: An adaptive accept/reject sampling algorithm for posterior probability distributions. In: *Statistical Signal Processing, 2009. SSP '09. IEEE/SP 15th Workshop on*, pp. 45–48 (2009)
21. Martino, L., Miguez, J.: A novel rejection sampling scheme for posterior probability distributions. In: *ICASSP 2009. IEEE International Conference on Acoustics, Speech and Signal Processing.*, pp. 2921–2924 (2009)
22. Philippe, A., Robert, C.P.: Perfect simulation of positive Gaussian distributions. *Statistics and Computing* **13**(2), 179–186 (2003)
23. Robert, C.P.: Simulation of truncated normal variables. *Statistics and Computing* **5**(2) (1995)
24. Robert, C.P., Casella, G.: *Monte Carlo Statistical Methods*. Springer-Verlag (2004)
25. Rodriguez-Yam, G., Davis, R.A., Scharf, L.L.: Efficient Gibbs Sampling of Truncated Multivariate Normal with Application to Constrained Linear Regression (2004). URL <http://www.stat.columbia.edu/~rdavis/papers/CLR.pdf>
26. Von Neumann, J.: Various Techniques Used in Connection with Random Digits. *J. Res. Nat. Bur. Stand.* **12**, 36–38 (1951)
27. Jun-wu YU, G.I.T.: Efficient Algorithms for Generating Truncated Multivariate Normal Distributions. *Acta Mathematicae Applicatae Sinica, English Series* **27**(4), 601 (2011)