

Efficient sampling from truncated bivariate Gaussians via Box-Muller transformation

L. Martino, D. Luengo and J. Míguez

Many practical simulation tasks demand procedures to draw samples efficiently from multivariate truncated Gaussian distributions. Introduced is a novel rejection approach, based on the Box-Muller transformation, to generate samples from a truncated bivariate Gaussian density with an arbitrary support. Furthermore, for an important class of support regions the new method allows exact sampling to be achieved, thus becoming the most efficient approach possible.

Introduction: The numerical simulation of many systems of practical interest demands the ability to produce Monte Carlo samples from truncated Gaussian distributions [1–3]. The simplest way to address this problem is to perform rejection sampling using the corresponding (non-truncated) Gaussian distribution as a proposal. This trivial method produces independent and identically distributed (I.I.D.) samples, but it is time consuming and computationally inefficient. For these two reasons, different methods have been introduced in the literature, e.g. using Markov chain Monte Carlo (MCMC) techniques [1, 3] or rejection sampling [4]. Unfortunately, MCMC schemes produce correlated samples, which can lead to a very slow convergence of the chain, whereas rejection methods can be computationally inefficient.

In this Letter, we introduce a novel approach, based on the Box-Muller transformation (BMT) [5], to generate I.I.D. samples from truncated bivariate Gaussian distributions. The main advantages of the proposed approach are the following: 1. it allows sampling within a generic domain $\mathcal{D} \subseteq \mathbb{R}^2$ without any restriction and 2. the inverse transformation of the BMT maps any region $\mathcal{D} \subseteq \mathbb{R}^2$ (either bounded or unbounded) into a bounded region, $\mathcal{A} \subseteq \mathcal{R} = [0, 1] \times [0, 1]$. Hence, all the procedures developed for drawing efficiently uniform random variables within bounded regions, e.g. adaptive rejection sampling or strip methods [5, 6], can always be used. Furthermore, for an important class of support regions the BMT allows us to perform exact sampling (i.e. draw I.I.D. samples from the target distribution without any rejection), which is the most efficient situation possible.

Problem formulation: The problem considered here is related to drawing samples from a truncated multivariate Gaussian distribution. In particular, in this Letter we focus on drawing samples from a bivariate truncated standard Gaussian PDF, denoted as $\mathbf{Z} \sim TN(\mathbf{0}, \mathbf{I}, \mathcal{D})$, where the support domain $\mathcal{D} \subseteq \mathbb{R}^2$ is a non-null Borel set. Note that drawing samples from a non-truncated standard Gaussian distribution, $\mathbf{Z} \sim N(\mathbf{0}, \mathbf{I})$, enables us to draw samples from an arbitrary Gaussian distribution, $\mathbf{X} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, whenever $\boldsymbol{\Sigma}$ is positive definite. More precisely, since $\boldsymbol{\Sigma}$ is positive definite, it can be expressed as $\boldsymbol{\Sigma} = \mathbf{S}\mathbf{S}^T$, using for instance the Cholesky decomposition, and the random vector $\mathbf{X} = \mathbf{S}\mathbf{Z} + \boldsymbol{\mu}$ has the desired distribution, $\mathbf{X} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Similarly, sampling from a truncated bivariate standard Gaussian distribution allows us to generate samples from an arbitrary truncated bivariate Gaussian. In this case, if $\mathbf{Z} \sim TN(\mathbf{0}, \mathbf{I}, \mathcal{D})$, then we can obtain $\mathbf{X} \sim TN(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \mathcal{D}^*)$ simply through the transformation $\mathbf{X} = \mathbf{S}\mathbf{Z} + \boldsymbol{\mu}$, with $\boldsymbol{\Sigma} = \mathbf{S}\mathbf{S}^T$ and

$$\mathcal{D} = \{\mathbf{z} \in \mathbb{R}^2 : \mathbf{z} = \mathbf{S}^{-1}(\mathbf{x} - \boldsymbol{\mu}) \quad \forall \mathbf{x} \in \mathcal{D}^*\} \quad (1)$$

Box-Muller transformation (BMT): The BMT maps a uniform random variable (RV) taking values within the rectangle $\mathcal{R} = [0, 1] \times [0, 1]$ into a standard Gaussian vector taking values on \mathbb{R}^2 . Given two independent uniform random variables inside the interval $[0, 1]$, $U_1, U_2 \sim \mathcal{U}([0, 1])$, the Box-Muller transformation is given by [5]:

$$Z_1 \triangleq M_1(U_2, U_1) = \sqrt{-2 \ln(U_1)} \sin(2\pi U_2) \quad (2)$$

$$Z_2 \triangleq M_2(U_2, U_1) = \sqrt{-2 \ln(U_1)} \cos(2\pi U_2) \quad (3)$$

The BMT transforms the RV $\mathbf{U} = (U_2, U_1)$ into another RV, $\mathbf{Z} = (Z_2, Z_1)$, following a bivariate standard Gaussian distribution, i.e. $\mathbf{Z} \sim N(\mathbf{0}, \mathbf{I})$. In the sequel, we indicate the BMT, given by (2) and (3), as $\mathbf{M} : \mathcal{R} \rightarrow \mathbb{R}^2$ with $\mathbf{M}(U_2, U_1) = (Z_2, Z_1)$. Hence, if we are able to draw uniformly a point (u'_2, u'_1) in \mathcal{R} , we can generate a transformed

point $(z'_2, z'_1) = \mathbf{M}(u'_2, u'_1)$ which is a sample from a standard bivariate Gaussian distribution.

Sampling bivariate truncated Gaussians via BMT: The BMT is a continuous transformation that establishes a one-to-one correspondence between the points in \mathcal{R} and the points in \mathbb{R}^2 , except for the point $(Z_1 = 0, Z_2 = 0)$. This point is obtained applying the BMT to any point with $u_1 = 1$ and $u_2 \in [0, 1]$. Thus, for any arbitrary support domain, $\mathcal{D} \subseteq \mathbb{R}^2$, there exists a subset, $\mathcal{A} \subseteq \mathcal{R}$, such that

$$\begin{aligned} \mathcal{D} = \mathbf{M}(\mathcal{A}) &= \{(z_2, z_1) \in \mathbb{R}^2 : (z_2, z_1) \\ &= \mathbf{M}(u_2, u_1) \text{ with } (u_2, u_1) \in \mathcal{A}\} \end{aligned}$$

Consequently, for every Borel subset $\mathcal{D} \subseteq \mathbb{R}^2$ we can obtain $\mathcal{A} \subseteq \mathcal{R}$ as

$$\mathcal{A} \triangleq \mathbf{M}^{-1}(\mathcal{D}) \quad (4)$$

Therefore, in order to draw a sample from a standard bivariate Gaussian restricted within a domain $\mathcal{D} \subseteq \mathbb{R}^2$, we can follow these two simple steps:

1. Draw a point $\mathbf{u}' = (u'_2, u'_1)$ uniformly within the set \mathcal{A} given by (4).
2. Transform \mathbf{u}' into a sample of the desired truncated Gaussian distribution using the BMT, i.e. $\mathbf{z}' = (z'_2, z'_1) = \mathbf{M}(u'_2, u'_1)$.

Note that, even if \mathcal{D} is unbounded, the set \mathcal{A} is always bounded, since $\mathcal{A} \subseteq \mathcal{R}$, which is bounded. Therefore, an efficient procedure can always be found to perform step 1 (e.g. using rejection sampling).

As an example, consider the unbounded truncation domain $\mathcal{D} = \{(z_2, z_1) : z_2 + z_1 + 0.9 \leq 0\}$, displayed in Fig. 1a. Note that the vector $\mathbf{z} = (z_2, z_1)$ can be uniquely represented by way of its polar co-ordinates, $r \triangleq \sqrt{z_1^2 + z_2^2}$ and $\theta \triangleq \tan^{-1}(z_1/z_2)$. These polar co-ordinates can be rewritten in terms of $\mathbf{u} = (u_2, u_1)$, using (2) and (3), as $r = \sqrt{-2 \ln(u_1)}$ and $\theta = 2\pi u_2$. Hence, $u_1 = \exp(-(z_1^2 + z_2^2)/2)$ and $u_2 = 1/2\pi \tan^{-1}(z_1/z_2)$. Finally, inserting the expression for the boundary of \mathcal{D} , $z_2 = -z_1 - 0.9$, into the equations of u_1 and u_2 , we can easily obtain the contour of the bounded set \mathcal{A} , shown in Fig. 1b.




Fig. 1 Example of support domain (Fig. 1a), $\mathcal{D} = \{(z_2, z_1) : z_2 + z_1 + 0.9 \leq 0\}$; we can observe that for $(z_2, z_1) \in \mathcal{D}$ we have $r \geq 0.63$ and $3\pi/4 \leq \theta \leq 7\pi/4$. The set $\mathcal{A} \subseteq \mathcal{R}$ corresponding to \mathcal{D} and the minimal rectangle \mathcal{R}_m (dashed and dotted line) embedding $\mathcal{A} \subseteq \mathcal{R}_m$ (Fig. 1b)

Uniform sampling inside \mathcal{A} : The crucial issue for this approach is sampling uniformly inside \mathcal{A} . Fortunately, when the support set \mathcal{D} is mapped into a simple region \mathcal{A} , such as a rectangle, then we can perform exact sampling (i.e. I.I.D. uniform samples can be drawn inside \mathcal{A} without any rejection). This happens, for instance, for the case in which \mathcal{D} is the sector of a circle delimited by the radii r_1 and r_2 ($r_1 < r_2$) and the angles θ_1 and θ_2 ($\theta_1 < \theta_2$), which includes a disk and a full circle as particular cases. The corresponding region \mathcal{A} is a rectangle, i.e.

$$\begin{aligned} \mathcal{A} &= \{(u_2, u_1) : e^{-r_2^2/2} \leq u_2 \leq e^{-r_1^2/2}, \theta_1/(2\pi) \\ &\leq u_1 \leq \theta_2/(2\pi)\} \end{aligned}$$

embedded inside the unit square \mathcal{R} . Hence, it is straightforward to draw samples uniformly and exactly from \mathcal{A} using two independent uniform random variables. Another relevant example of exact sampling occurs when the boundary of \mathcal{D} is a straight line passing through the origin, i.e. $az_1 + bz_2 = 0$. In this case all points in \mathcal{D} are completely expressed by the inequalities $r \geq 0$ and $\tan^{-1}\{-b/a\} \leq \theta \leq \tan^{-1}\{-b/a\} + \pi$, and \mathcal{A} is again a rectangle. Additional examples of simple \mathcal{A} regions where exact sampling can be performed include triangles and sectors of circles [5, 7].

When exact sampling cannot be performed, then we can follow the simple rejection sampling approach described in the sequel to draw uniform samples inside \mathcal{A} . Given a domain \mathcal{D} , we can find a circular sector \mathcal{C}_m that embeds \mathcal{D} (i.e. $\mathcal{D} \subseteq \mathcal{C}_m$). Now, since $\mathcal{R}_m \triangleq \mathbf{M}^{-1}(\mathcal{C}_m)$ is a rectangular region (as explained above), and since $\mathcal{R}_m = \mathbf{M}^{-1}(\mathcal{C}_m) \supseteq \mathcal{A} = \mathbf{M}^{-1}(\mathcal{D})$, we can easily generate a point \mathbf{u}' uniformly inside \mathcal{R}_m , accepting it only if it belongs to \mathcal{A} (or, if it is easier, we can check whether $\mathbf{z}' = \mathbf{M}(\mathbf{u}')$ belongs to \mathcal{D} or not). The acceptance probability of this approach, $p_a = |\mathcal{A}|/|\mathcal{R}_m|$ with $|\cdot|$ indicating the Lebesgue measure of a set, can be very close to one, as shown in the simulations.

Simulations: In a first simulation we use a rectangular domain with vertices $(0, 0)$, $(0, \Delta)$, $(\Delta, 0)$ and (Δ, Δ) . In this case, we can obtain I.I.D. samples using the following rejection techniques:

- **M1:** Trivial rejection sampler using the non-truncated Gaussian PDF as a proposal, accepting the samples that belong to \mathcal{D} . This is equivalent to drawing samples uniformly inside \mathcal{R} , accepting those that belong to \mathcal{A} .
- **M2:** Since the set \mathcal{D} is bounded ($\Delta < \infty$), we can use another trivial rejection scheme: drawing samples uniformly within \mathcal{D} and accepting them with probability $\exp(-(z_1^2 + z_2^2)/2)$.
- **M3:** A more refined rejection procedure, introduced in [6] and discussed for truncated Gaussian PDFs in [4], consists in building the plane $y = \alpha z_1 + \beta z_2 + \phi$ tangent to the paraboloid $y = (z_1^2 + z_2^2)/2$ at some arbitrary point $(z_1^*, z_2^*) \in \mathcal{D}$. This induces a proposal PDF

$$\pi(z_1, z_2) \propto \exp(-\alpha z_1 - \beta z_2 - \phi) \quad \forall (z_1, z_2) \in \mathcal{D} \quad (5)$$

which corresponds to two truncated independent exponential PDFs. We consider two variants: the plane tangent at $(z_1^*, z_2^*) = (0, 0)$ (**M3-1**) and the plane tangent at $(z_1^*, z_2^*) = (\Delta/2, \Delta/2)$ (**M3-2**).

- **M4:** This is our rejection strategy: drawing samples from the minimal rectangle, $\mathcal{R}_m = [0, \pi/2] \times [1, \exp\{-\Delta^2\}]$, that covers the set \mathcal{A} (i.e. $\mathcal{R}_m \supseteq \mathcal{A}$), and accepting them if they belong to \mathcal{A} .

We apply these four rejection techniques, computing the acceptance rate (AR), averaged over 50000 runs, against the area of the domain \mathcal{D} , Δ^2 . The results are shown in Fig. 2. Note that, in this case, **M2** and **M3-1** are equivalent. Our method, **M4**, provides the best performance for $\Delta^2 < 2.8$ (i.e. $\Delta > 1.67$). Furthermore, for $\Delta \rightarrow \infty$ our technique provides virtually exact sampling (i.e. $\text{AR} \rightarrow 1$), whereas for **M1** the AR clearly approaches 1/4 and for **M2** and **M3** the $\text{AR} \rightarrow 0$ as $\Delta \rightarrow \infty$.




Fig. 2 Acceptance rate (AR) against area of domain \mathcal{D} , Δ^2
M1 (\times marks), **M2** \equiv **M3-1** (triangles), **M3-2** (squares) and **M4** (solid line)

As a second example, we consider $\mathcal{D} = \{(z_2, z_1) : z_2 + z_1 + c \leq 0\}$, where c is a constant parameter. In this case, we can just use the trivial strategy **M1** and our method **M4**. **M2** cannot be applied, since \mathcal{D} is unbounded, and **M3** is unpractical, since it is very difficult to draw samples, using a direct method, from a bivariate exponential PDF restricted to this kind of set. For $c = 0$, our approach (**M4**) performs exact sampling (i.e. $\text{AR} = 1$), whereas **M1** provides an AR of 0.5. For $c = 0.9$, **M4** obtains an AR of 0.64, whereas **M1** provides an AR of 0.26. Finally, for $c = 2$, **M4** obtains an AR of 0.45, whereas **M1** provides an AR of 0.12. Hence, it can be seen that our approach **M4** always outperforms **M1**.

Conclusion: We have proposed a new approach to draw I.I.D. samples from truncated bivariate normal distributions that allows sampling within an arbitrary domain, without any restriction. For some truncation regions, this technique provides exact sampling without any rejection. For more general cases, it allows us to develop rejection sampling approaches that can obtain better acceptance rates than state-of-the-art methods. Furthermore, unlike most of the schemes available in the literature, our approach can always be used, regardless of whether the support is bounded or not. Future lines include extending it to arbitrary N -dimensional truncated Gaussian vectors.

Acknowledgment: This work has been partly financed by the Spanish government, through the DEIPRO project (TEC2009-14504-C02-01) and the CONSOLIDER-INGENIO 2010 programme (CSD2008-00010).

© The Institution of Engineering and Technology 2012

6 August 2012

doi: 10.1049/el.2012.2816

One or more of the Figures in this Letter are available in colour online.

L. Martino and J. Míguez (*Dep. Teoría de Señal y Comunicaciones, Univ. Carlos III de Madrid, Av. Universidad 30, Leganés 28911, Spain*)

E-mail: luca@tsc.uc3m.es

D. Luengo (*Dep. Ingeniería de Circuitos y Sistemas, Univ. Politécnica de Madrid, Carretera de Valencia km. 7, Madrid 28031, Spain*)

References

- 1 Kotecha, J.H., and Djuric, P.M.: 'Gibbs sampling approach for generation of truncated multivariate gaussian random variables'. Proc. of 25th IEEE ICASSP, Phoenix, AZ, USA, 1999
- 2 Gelfand, A.E., Smith, A.F.M., and Lee, T.M.: 'Bayesian analysis of constrained parameter and truncated data problems using Gibbs sampling', *J. Am. Stat. Assoc.*, 1992, **87**, pp. 523–532
- 3 Robert, C.P.: 'Simulation of truncated normal variables', *Stat. Comput.*, 1995, **5**, (2), pp. 121–125
- 4 Chopin, N.: 'Fast simulation of truncated Gaussian distributions', *Stat. Comput.*, 2009, **21**, (2), pp. 275–288
- 5 Devroye, L.: 'Non-uniform random variate generation' (Springer, New York, 1986)
- 6 Gilks, W.R., and Wild, P.: 'Adaptive rejection sampling for Gibbs sampling', *Appl. Stat.*, 1992, **41**, (2), pp. 337–348
- 7 Luengo, D., and Martino, L.: 'Efficient random variable generation: ratio of uniforms and polar rejection sampling', *Electron. Lett.*, 2012, **48**, (6), pp. 326–327