DOI: 10.1111/cgf.14058
Eurographics Symposium on Rendering 2020
C. Dachsbacher and M. Pharr
(Guest Editors)

Volume 39 (2020), Number 4

Can't Invert the CDF? The Triangle-Cut Parameterization of the Region under the Curve

E. Heitz

Unity Technologies

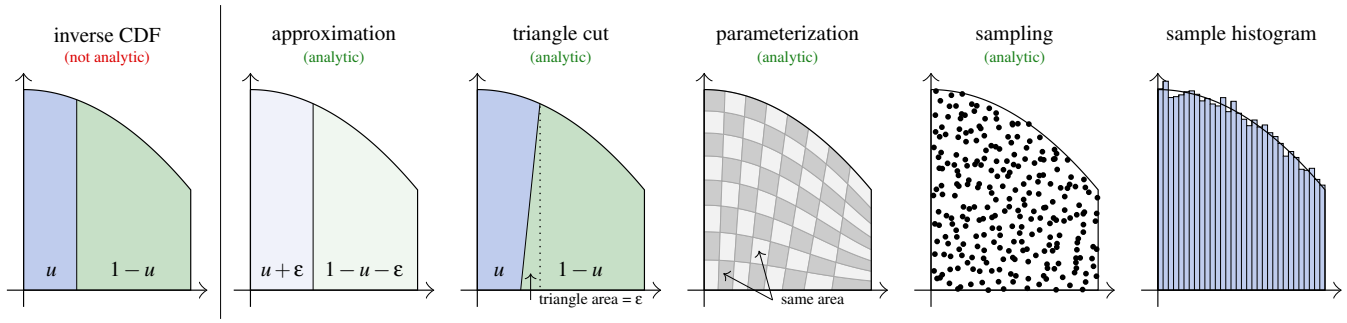

Figure 1: The triangle-cut parameterization. *Analytic sampling is typically achieved by inverting the Cumulative Distribution Function (CDF) of the target density. Intuitively, the inverse CDF partitions the region under the curve according to a uniform random number $u \in [0,1]$. In this example, the CDF is not analytically invertible. Our idea is to use the analytic inverse CDF of an approximate density and fix it by cutting a triangle such that the partitioning with respect to $u \in [0,1]$ remains correct. By sampling along the partitioning segment using a second random number, we obtain a 2D area-preserving parameterization that we use to sample points uniformly in the region under the curve of the density. The computation of these points is fully analytic and their abscissae are distributed with the target density.*

Abstract

We present an exact, analytic and deterministic method for sampling densities whose Cumulative Distribution Functions (CDFs) cannot be inverted analytically. Indeed, the inverse-CDF method is often considered the way to go for sampling non-uniform densities. If the CDF is not analytically invertible, the typical fallback solutions are either approximate, numerical, or non-deterministic such as acceptance-rejection. To overcome this problem, we show how to compute an analytic area-preserving parameterization of the region under the curve of the target density. We use it to generate random points uniformly distributed under the curve of the target density and their abscissae are thus distributed with the target density. Technically, our idea is to use an approximate analytic parameterization whose error can be represented geometrically as a triangle that is simple to cut out. This triangle-cut parameterization yields exact and analytic solutions to sampling problems that were presumably not analytically resolvable.

CCS Concepts

• ***Mathematics of computing*** → ***Stochastic processes;***

1. Introduction

Monte Carlo integration relies heavily on the generation of random variates from non-uniform distributions. The statistical literature offers a variety of techniques for this purpose [Dev86]. However, the majority of them are not recommended for Monte Carlo rendering. Indeed, it is well known that Monte Carlo estimators can greatly benefit from sample stratification [Shi91] and the rendering community actively researches *stratified sampling* techniques. The main component of stratified sampling is an *area-preserving* (for a region) or *integral-preserving* (for a density) parameterization. These parameterizations are bijective mappings between the unit square (where uniform random numbers are sampled) and the target region or density. The ability to find and evaluate these parameterizations is the key to stratified sampling and the main motivation behind this paper.

The successful research of analytic parameterizations. With stratified sampling for Monte Carlo rendering as a motivation, the computer graphics community started to revisit parameterizations of simple shapes such as triangles [Tur90], disks [SC97], cylinders and spheres [SWZ96]. Occasionally, improvements can still be found for simple shapes such as triangles [Hei19] but the community moved on to more challenging problems. To approach these new problems, Arvo promoted a general recipe for computing area-preserving parameterizations of arbitrary regions or distributions [Arv01]. It is based on the classic *inverse-CDF* sampling method, which consists of inverting the integral of the target domain or density, represented by its Cumulative Distribution Function (CDF). Today, the inverse-CDF method is largely considered the way to go for computing area-preserving parameterizations. It has been successfully used to obtain analytic solutions for the hemisphere [Arv01], Phong distributions [Arv01], spherical triangles [Arv95], convex quadrilaterals [AN07], spherical rectangles [UnFK13], and the distribution of visible normals of microfacet surfaces [Hd14, Hei18].

Stumbling on more complex problems. Unfortunately, only a fraction of functions can be analytically inverted, and the mainstream inverse-CDF method often fails to provide analytic parameterizations. It is interesting to point out that the more the community tries to address complicated problems, the less likely it seems to obtain analytic parameterizations. For instance, Gamito observed that the inverse-CDF parameterization for the solid angle of disks and cylinders is not analytic and instead falls back to a simpler proxy shape with rejection sampling, which breaks the stratification [Gam16]. Urēna et al. chose to not compromise the stratification of spherical ellipses and they use Newton iterations to make a numerical inversion of the non-analytically invertible CDFs, which is computationally expensive [GUnK*17]. Another interesting example is the area-preserving parameterization of a truncated disk as in Figure 2, which has been recently motivated by stratified sampling of projected spherical caps [UnG18, PD19]. While a truncated disk appears to be a simple shape, its CDF cannot be analytically inverted, so these previous works do not provide an analytic and exact solution. Urēna and Georgiev invert the CDF numerically and Peters and Dachsbacher designed an analytic approximation. Other communities face the same problem. For instance, elaborate phase functions for astronomy do not have analytic inverse CDFs either [Zha19]. Recent works in graphics have started to approach difficult sampling problems with machine learning [MMR*19, ZZ19]. All considered, we should not expect to obtain exact and analytic solutions when tackling more difficult sampling problems in the future.

Insights. Our contribution arises from several observations that can be made by looking at the truncated-disk example of Figure 2. First, we note that finding an area-preserving parameterization for the truncated disk is equivalent to finding one for the region under the curve of its marginal density. More generally:

→ Observation 1: *Many interesting 2D densities have a trivial mapping to the 2D region under the curve of their 1D marginal density.*

Furthermore, by looking at the region under the curve, we can see that the inverse-CDF parameterization is special.

→ Observation 2: *The inverse CDF computes a unique area-preserving parameterization of the region under the curve that is axis aligned.*

If this axis-aligned parameterization is not analytic, the inverse-CDF approach leaves no other choice than using a numerical inversion or an approximation. However, we might envision other options.

→ Observation 3: *There are an infinite number of alternative non-axis-aligned area-preserving parameterizations that we could use.*

If one of these alternative parameterizations is analytic, then we can obtain analytic stratified sampling even when the CDF is not analytically invertible. But are some of these alternative parameterizations analytic? We are not aware of relevant previous work on the topic. We thus created one of these alternative parameterizations that we call the *triangle-cut parameterization* shown in Figure 1.

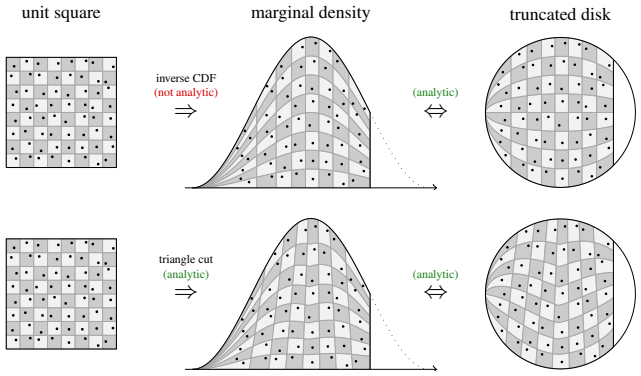

Figure 2: Stratified sampling of a truncated disk. *The parameterization of a 2D region or density can also be regarded as a parameterization of the region under the curve of its 1D marginal density. Under the region of the curve, the inverse-CDF method computes a unique parameterization that is axis aligned. In this example, this axis-aligned parameterization cannot be evaluated analytically and previous works use either a numerical inversion or an approximation [UnG18, PD19]. With the triangle cut, we relax the axis-aligned constraint to obtain an alternative area-preserving parameterization that can be evaluated analytically (Sec. 5).*

Contributions. After recalling the classic inverse-CDF parameterization in Section 2, we make the following contributions:

- In Section 3, we introduce an alternative area-preserving parameterization of the 2D region under the curve of a 1D density that we call the *triangle-cut parameterization*. We obtain it by regularizing an approximate analytic parameterization whose error can be represented as a triangle that we cut out. We show how to evaluate it analytically and provide a formal proof that, so long as the approximation meets some conditions, it is an exact area-preserving bijection.
- In Section 4, we explain how the triangle-cut parameterization can be used for sampling 1D densities and for stratified sampling of 2D densities. We apply this approach to a truncated disk (Sec. 5), the surface of a torus (Sec. 6), a polar shape (Sec. 7), a polynomial density (Sec. 8), and a subsurface-scattering model (Sec. 9).

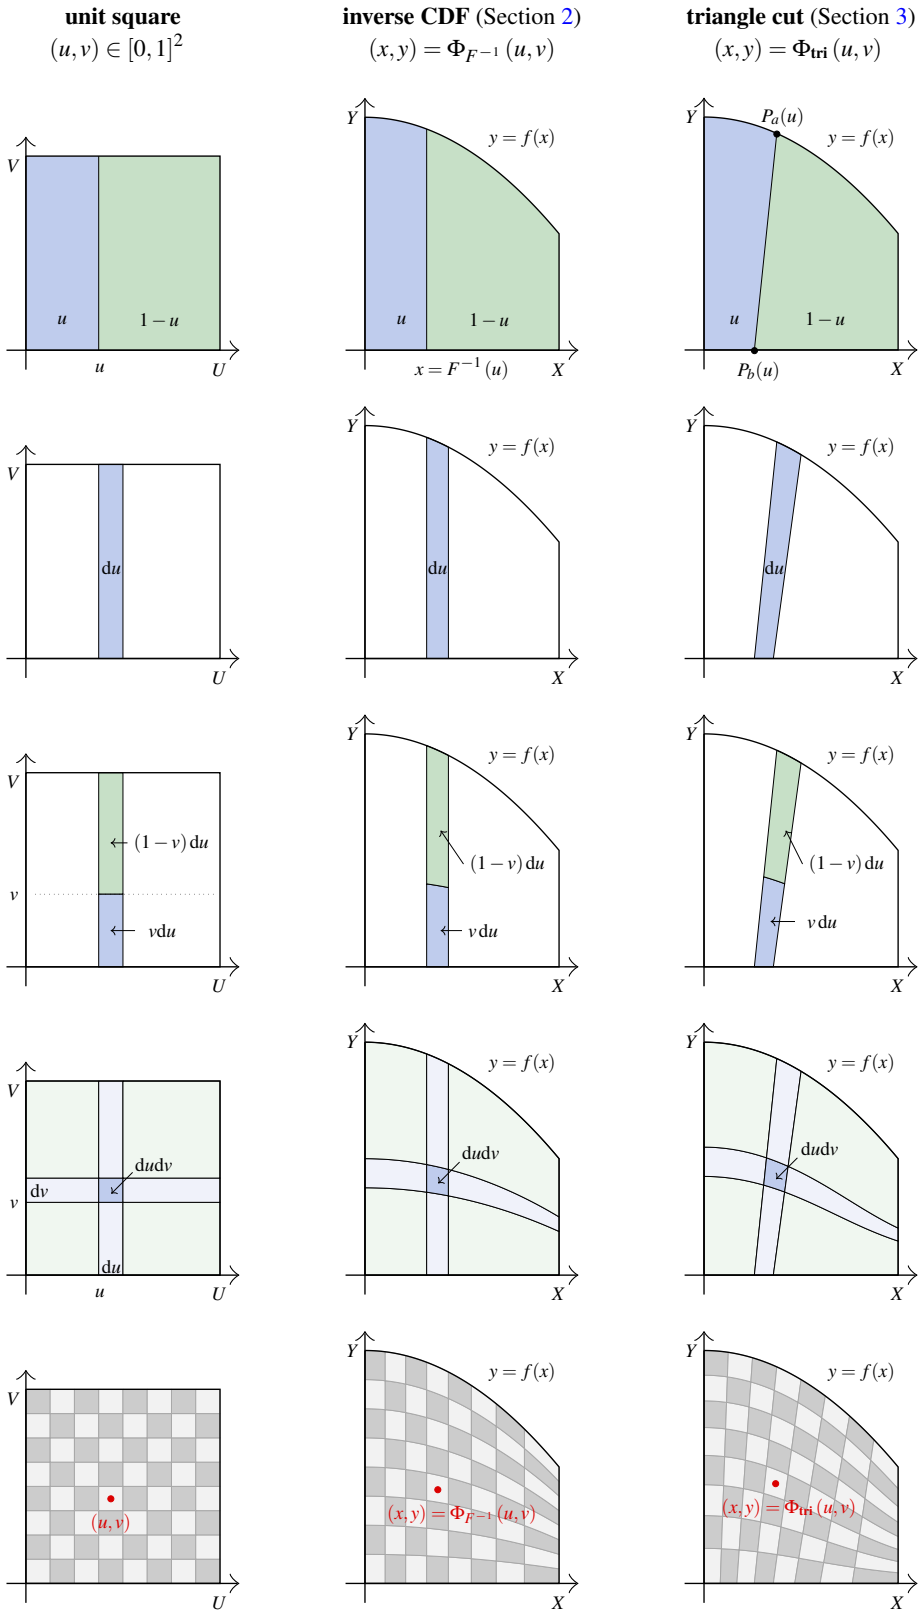

(a) Partitioning the region under the curve of density f. *The partitioning is such that the area of the left part is u and the area of the right part is* $1-u$.

(b) Infinitesimal thickness of the partitioning segment. *The parameterization maps an infinitesimal vertical slice of area* $\mathrm{d}u$ *to an infinitesimal slice of the same area.*

(c) Sampling the partitioning segment. *The partitioning segment is sampled with a distribution proportional to its infinitesimal thickness.*

(d) Jacobian of the parameterization. *The parameterization is area preserving if it maps infinitesimal regions to infinitesimal regions of the same area.*

(e) Visualization of the parameterization. *All the cells of the checkers have the same area.*

Figure 3: Area-preserving parameterizations of the 2D region under the curve of a 1D density. *The inverse CDF computes a unique area-preserving parameterization that is axis aligned. When the inverse-CDF parameterization is not analytic, we propose to compute an alternative non-axis-aligned parameterization that is also area preserving and that can be analytically computed.*

2. Background on the Inverse-CDF Parameterization

In this section, we review classic and well-known results related to the inverse CDF. The point is to review the derivation of an area-preserving parameterization in order to prepare the reader for Section 3, where we introduce our new parameterization with the same methodology, as shown in Figure 3.

Probability Distribution Function (PDF). We consider a density

$$f(x) = y, \tag{1}$$

that is a PDF, i.e. it is non-negative and integrates to exactly 1.

Cumulative Distribution Function (CDF) The CDF is the integral of the PDF and we denote it with the same capitalized letter:

$$F(x) = \int_{-\infty}^{x} f(x')\, dx'. \tag{2}$$

Inverse Cumulative Distribution Function (iCDF). The inverse CDF maps a uniform random number $u \in [0,1]$ to a point distributed with density f:

$$x = F^{-1}(u). \tag{3}$$

The inverse CDF is represented in Figure 3-(a). It partitions the region under the curve of f with a vertical segment located at x such that the area of the left part is u and the area of the right part is $1-u$. As shown in Figure 3-(b), the consequence is that it maps an infinitesimal vertical slice of area du in the uniform distribution to an infinitesimal vertical slice of the same area $\mathrm{d}u = f(x)\,\mathrm{d}x$ under the curve of f. The larger $f(x)$, the smaller dx. The density of x is thus proportional to f.

Fundamental theorem of simulation. This theorem states that uniform sampling of the region under the curve of a PDF is equivalent to sampling this PDF [MLM18]. This is what motivated us to look into area-preserving parameterizations that can be used to sample 2D points (x,y) uniformly distributed in the region under the curve of f.

Parameterizing the region under the curve. In the case of the inverse-CDF parametrization, Equation (3) yields a 1D point x distributed with density f. The second coordinate y can thus be sampled uniformly over the vertical segment, i.e. by multiplying a uniform random number $v \in [0,1]$ by $f(x)$, as shown in Figure 3-(c). The inverse-CDF parameterization of the region under the curve is defined by

$$\Phi_{F^{-1}}(u,v) = (x,y) = \left(F^{-1}(u), v\, f(x)\right), \tag{4}$$

and it is shown in Figure 3-(e).

Proof that the parameterization is area preserving. Intuitively, the parameterization is area preserving because, by construction, it maps an infinitely small region of area dudv to a region of the same area, as shown in Figure 3-(d). Formally, proving that it is area preserving means showing that the determinant of its Jacobian matrix is 1 everywhere. We compute the partial derivatives

$$\frac{\partial x}{\partial u} = \frac{1}{f(x)}, \quad \frac{\partial x}{\partial v} = 0, \quad \frac{\partial y}{\partial u} = \frac{f'(x)}{f(x)}\, v, \quad \frac{\partial y}{\partial v} = f(x), \tag{5}$$

and the determinant of the Jacobian matrix evaluates to

$$\left|J_{\Phi_{F^{-1}}}\right| = \left|\begin{pmatrix} \frac{\partial x}{\partial u} & \frac{\partial x}{\partial v} \\ \frac{\partial y}{\partial u} & \frac{\partial y}{\partial v} \end{pmatrix}\right| = 1. \tag{6}$$

Discussion. In many real-life problems, the integral of the density can be evaluated analytically (F is analytic) but cannot be analytically inverted (F^{-1} is not analytic). In the next section, we introduce an alternative parameterization that avoids non-analytic inversions.

3. The Triangle-Cut Parameterization

In this section, we introduce an alternative analytic area-preserving parameterization $\Phi_{\mathbf{tri}}$ of the region under the curve of density f that we call the *triangle-cut* parameterization.

Problem statement. Our goal is to compute an area-preserving parameterization of the region under the curve of a target density f. A classic problem to overcome is that we can easily evaluate F but not F^{-1}. Our idea is to use an approximate density g whose inverse-CDF parameterization can be computed analytically and we apply a triangle cut, a modification such that it becomes an area-preserving parameterization $\Phi_{\mathbf{tri}}$ of the region under the curve of f.

3.1. Evaluation of the Triangle-Cut Parameterization

We now show how to map two uniform random numbers $(u,v) \in [0,1]^2$ to a random point $(x,y) = \Phi_{\mathbf{tri}}(u,v)$ distributed uniformly in the region under the curve of f. We prove it in Section 3.2.

Approximate density. The first step consists of mapping the random number u to a sample x_a from the approximate density g:

$$x_\mathrm{a} = G^{-1}(u). \tag{7}$$

In Figure 4-(a), we can see that x_a is perfectly distributed with density g because it perfectly partitions the region under the curve of g. However, as shown in Figure 4-(b), the bias of x_a with respect to f can be visualized as an excess or loss of area ε in the vertical partitioning of the region under the curve, which is measured by

$$\varepsilon = u - F(x_\mathrm{a}). \tag{8}$$

Triangle cut. To fix the incorrect partitioning of Figure 4-(b), we cut out a region that has the same area as the error ε. We found that the simplest way to remove the error was to cut out a triangle, as shown in Figure 4-(c). Indeed, since we know the area ε and the height $f(x_\mathrm{a})$ of the triangle, the location x_b of the new vertex is given by

$$\varepsilon = \frac{(x_\mathrm{b} - x_\mathrm{a})\, f(x_\mathrm{a})}{2} \quad \Rightarrow \quad x_\mathrm{b} = x_\mathrm{a} + \frac{2\varepsilon}{f(x_\mathrm{a})}. \tag{9}$$

After removing the triangle, we obtain a correct partitioning of the region under the curve, which is represented in Figure 3-(a). In summary, for a given $u \in [0,1]$, the end points of the partitioning segment are

$$P_a(u) = (x_\mathrm{a}, y_\mathrm{a}) = \left(G^{-1}(u), f(x_\mathrm{a})\right), \tag{10}$$

$$P_b(u) = (x_\mathrm{b}, y_\mathrm{b}) = \left(x_\mathrm{a} + 2\,\frac{u - F(x_\mathrm{a})}{f(x_\mathrm{a})}, 0\right). \tag{11}$$

As a result, this partitioning maps an infinitesimal vertical slice of area du in the uniform distribution to an infinitesimal slice of the same area under the curve of f, as shown in Figure 3-(b).

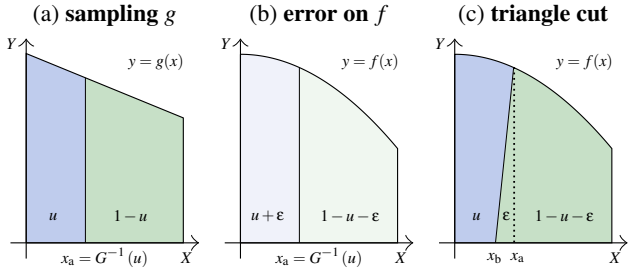

Figure 4: The triangle cut. *We generate sample x_a using an approximate density g. While x_a perfectly partitions the region under the curve of g, it produces an error under the curve of f. This error can be interpreted as an excess or loss of area ε. Our idea is to fix the partition by cutting out a triangle of the same area.*

Infinitesimal thickness of the partitioning segment. The next step consists of sampling a point on the partitioning segment with the second random number v, as illustrated in Figure 3-(c). To obtain an area-preserving parameterization, we partition the infinitesimal region of area du swept by the segment when u increases by du. In contrast to the inverse-CDF parameterization, where the infinitesimal thickness is distributed uniformly along the segment, in this case the infinitesimal region is shaped as a convex quadrilateral (see Figure 5). The thickness of a convex quadrilateral is distributed proportionally to an affine function that interpolates the infinitesimal thicknesses at the end points [AN07]. The thicknesses at the end points P_a and P_b are defined by the dot products of the derivatives of the end points:

$$\frac{\partial x_\mathrm{a}}{\partial u} = \frac{1}{g(x_\mathrm{a})}, \tag{12}$$

$$\frac{\partial y_\mathrm{a}}{\partial u} = \frac{f'(x_\mathrm{a})}{g(x_\mathrm{a})}, \tag{13}$$

$$\frac{\partial x_\mathrm{b}}{\partial u} = \frac{1}{g(x_\mathrm{a})} + 2\,\frac{(g(x_\mathrm{a}) - f(x_\mathrm{a}))\, f(x_\mathrm{a}) - f'(x_\mathrm{a})\,(u - F(x_\mathrm{a}))}{f(x_\mathrm{a})^2\, g(x_\mathrm{a})}, \tag{14}$$

$$\frac{\partial y_\mathrm{b}}{\partial u} = 0, \tag{15}$$

with the normal of the segment $N = (n_x, n_y)$:

$$(n_x, n_y) = \frac{(y_\mathrm{a} - y_\mathrm{b}, x_\mathrm{b} - x_\mathrm{a})}{\|(y_\mathrm{a} - y_\mathrm{b}, x_\mathrm{b} - x_\mathrm{a})\|} = \frac{\left(f(x_a), 2\frac{(u - F(x_a))}{f(x_a)}\right)}{\left\|\left(f(x_a), 2\frac{(u - F(x_a))}{f(x_a)}\right)\right\|}. \tag{16}$$

By expanding and simplifying the dot products

$$w_a = (n_x, n_y) \cdot \left(\frac{\partial x_\mathrm{a}}{\partial u}, \frac{\partial y_\mathrm{a}}{\partial u}\right), \tag{17}$$

$$w_b = (n_x, n_y) \cdot \left(\frac{\partial x_\mathrm{b}}{\partial u}, \frac{\partial y_\mathrm{b}}{\partial u}\right), \tag{18}$$

we obtain the infinitesimal thicknesses

$$w_a \propto f(x_\mathrm{a})^2 + 2\,(u - F(x_\mathrm{a}))\, f'(x_\mathrm{a}), \tag{19}$$

$$w_b \propto 2\, f(x_\mathrm{a})\, g(x_\mathrm{a}) - \left[f(x_\mathrm{a})^2 + 2\,(u - F(x_\mathrm{a}))\, f'(x_\mathrm{a})\right], \tag{20}$$

whose proportionality factor is the *same* and cancels out in the affine thickness density function shown in Figure 5-(right):

$$w(t) = 2\,\frac{t\, w_a + (1-t)\, w_b}{w_a + w_b} \quad \text{with} \quad t \in [0,1]. \tag{21}$$

An important condition to note is that the thicknesses should be non-negative: $w_a, w_b \geq 0$. Intuitively, this means that the segment should always move forward and never backward. This is the second condition to be verified by g that we present in Section 3.2.

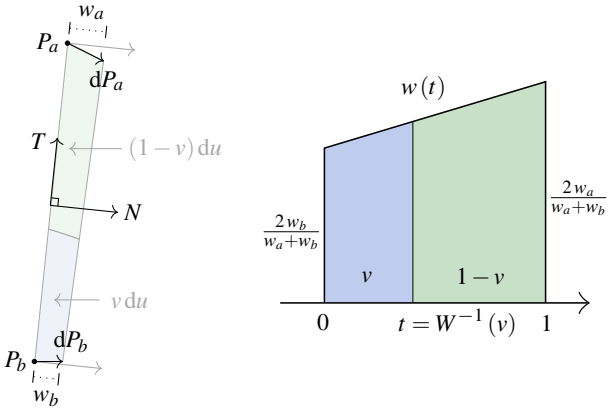

Figure 5: Infinitesimal thickness of the partitioning segment. *This is the partitioning segment of Figure 3-(c). The infinitesimal thickness of the moving partitioning segment is an affine function defined by the thicknesses at the end points. They are the dot products of the derivatives of the end points and the normal of the segment.*

Sampling the thickness density. We use the inverse CDF of w to sample t using the second random number v. This requires inverting a second-order polynomial. Muller's formulation provides a numerically stable solution to this problem [Mul56]:

$$t = W^{-1}(v) = \frac{v\,(w_a + w_b)}{w_b + \sqrt{(1-v)\, w_b^2 + v\, w_a^2}}. \tag{22}$$

Sampling the partitioning segment. Finally, we obtain the target point $(x, y) = \Phi_{\mathbf{tri}}(u, v)$ on the partitioning segment by interpolating P_a and P_b with t:

$$\Phi_{\mathbf{tri}}(u, v) = t\, P_a + (1-t)\, P_b, \tag{23}$$

which concludes the derivation of the triangle-cut parameterization illustrated in Figure 3-(e).

Evaluation of the triangle-cut parameterization. The implementation provided in Listing 1 summarizes the evaluation of $(x, y) = \Phi_{\mathbf{tri}}(u, v)$. The advantage is that it requires single calls to f, F, f', g and G^{-1}, which is what makes it competitive compared to a numerical inverse CDF that requires several calls to a potentially costly F.

Listing 1: Implementation of the triangle-cut parameterization.

```
float f(float x);       // target PDF
float F(float x);       // target CDF
float fprime(float x);  // target PDF derivative
float g(float x);       // approximate PDF
float iG(float u);      // approximate iCDF

// maps (u,v) uniformly distributed in [0,1]^2 to
// (x,y) uniformly distributed in the region under the curve of f
void trianglecut_param(float u, float v, float& x, float& y)
{
    // sample x_a with approximate PDF, Eq. (7)
    float x_a = iG(u);

    // eval functions at x_a
    float f_x_a = f(x_a);
    float F_x_a = F(x_a);
    float fprime_x_a = fprime(x_a);
    float y_a = f_x_a;
    float g_x_a = g(x_a);

    // compute x_b with triangle cut, Eq. (9)
    float x_b = x_a + 2.0*(u-F_x_a)/f_x_a;
    float y_b = 0;

    // compute infinitesimal thicknesses, Eq. (19) and (20)
    float w_a = f_x_a*f_x_a + 2.0*(u-F_x_a)*fprime_x_a;
    float w_b = 2.0*f_x_a*g_x_a - w_a;

    // sample thickness density, Eq. (22)
    float t = v*(w_a+w_b)/(w_b + sqrt((1-v)*w_b*w_b + v*w_a*w_a));

    // interpolate (x_a,y_a) and (x_b,y_b), Eq. (23)
    x = x_a * t + x_b * (1-t);
    y = y_a * t + y_b * (1-t);
}
```

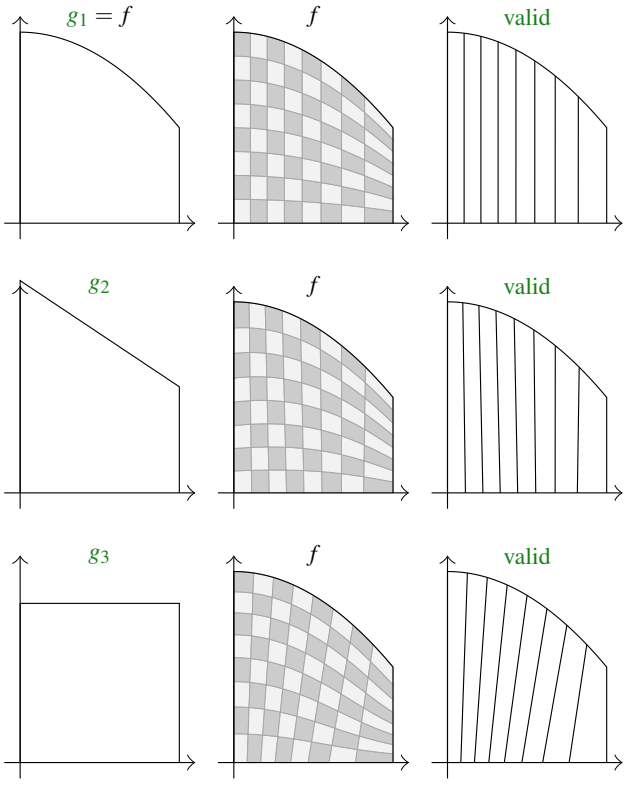

Figure 6: Valid triangle-cut parameterizations. *The cruder the approximation, the more the parameterization is tilted to compensate for the error. If $g = f$, the triangle-cut parameterization is exactly the inverse-CDF parameterization, which is axis aligned.*

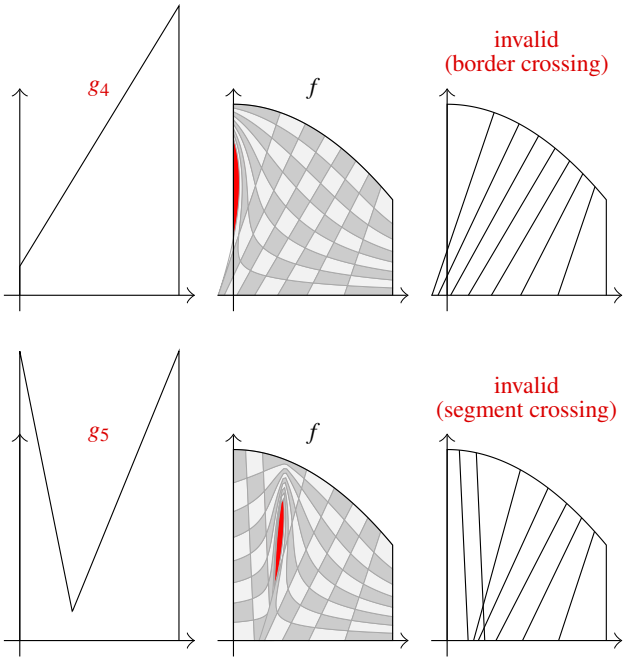

Figure 7: Invalid triangle-cut parameterizations. *If the approximate density g is not a reasonable approximation of the target density f, the triangle-cut parameterization is invalid. The invalidity can be visualized as the partitioning segments crossing each other or crossing the border of the region under the curve. Note that in the two invalid cases, the parameterization misses a part of the region under the curve (in red). Indeed, since the parameterization is area preserving, the area lost in the overlap or outside the border is also missing somewhere in the region under the curve.*

3.2. Validity of the Triangle-Cut Parameterization

Figures 6 and 7 show that not every approximate density yields a valid triangle-cut parameterization. We provide two conditions to verify when choosing g, and we prove that the resulting parameterization is valid if they are met.

Condition 1: no border crossing. We expect the parameterization to remain inside the region under the curve (the regions where f returns 0 should not be sampled at all). In Figure 7, the parameterization computed with g_4 crosses the border of the region under the curve. To prevent this, we verify that every point reached by the parameterization $(x(u,v), y(u,v)) = \Phi_{\mathbf{tri}}(u,v)$ is located in the region under the curve, i.e. that

$$y(u,v) \leq f(x(u,v)) \text{ for all } (u,v) \in [0,1]^2. \tag{24}$$

Condition 2: no segment crossing. In Figure 7, the parameterization computed with g_5 overlaps itself when the partitioning segment moves backward. This can be visualized as a crossing of its vertical partitioning segments. As already discussed after Equation (21), to prevent this we verify that the infinitesimal thicknesses of the segment are non-negative:

$$w_a(u) \geq 0 \text{ and } w_b(u) \geq 0 \text{ for all } u \in [0,1]. \tag{25}$$

Proof that the parameterization is area preserving. The evaluation presented in Section 3.1 requires Condition 2, as explained after Equation (21). If this condition is verified, the triangle-cut parameterization $\Phi_{\mathbf{tri}}$ is area preserving because, by construction, it maps any infinitesimal domain of area $\mathrm{d}u\,\mathrm{d}v$ to an infinitesimal domain of the same area, as shown in Figure 3-(d). To provide a formal proof of this result, we show that the determinant of the Jacobian matrix of $\Phi_{\mathbf{tri}}$ is 1 everywhere. First, we compute the partial derivatives of point $\Phi_{\mathbf{tri}}(u,v) = t(u,v)\,P_a(u) + (1-t(u,v))\,P_b(u)$ with respect to u and v:

$$\frac{\partial \Phi_{\mathbf{tri}}}{\partial u} = t(u,v)\,\frac{\partial P_a}{\partial u}(u) + (1-t(u,v))\,\frac{\partial P_b}{\partial u}(u) + \frac{\partial t}{\partial u}(u,v)\,(P_a(u) - P_b(u)), \tag{26}$$

$$\frac{\partial \Phi_{\mathbf{tri}}}{\partial v} = \frac{\partial t}{\partial v}(u,v)\,(P_a(u) - P_b(u)). \tag{27}$$

Note that the calculation of the Jacobian is simpler in the basis (N,T), where N is the normal of the partitioning segment and T its tangent. By projecting the partial derivatives in this basis (see Appendix A for a detailed derivation of these equations),

$$\frac{\partial \Phi_{\mathbf{tri}}}{\partial u} \cdot N = t(u,v)\,w_a(u) + (1-t(u,v))\,w_b(u), \tag{28}$$

$$\frac{\partial \Phi_{\mathbf{tri}}}{\partial u} \cdot T = \text{not required}, \tag{29}$$

$$\frac{\partial \Phi_{\mathbf{tri}}}{\partial v} \cdot N = 0, \tag{30}$$

$$\frac{\partial \Phi_{\mathbf{tri}}}{\partial v} \cdot T = \frac{1}{t(u,v)\,w_a(u) + (1-t(u,v))\,w_b(u)}, \tag{31}$$

the determinant of the Jacobian matrix simplifies to

$$\left|J_{\Phi_{\mathbf{tri}}}\right| = \left| \begin{pmatrix} \frac{\partial \Phi_{\mathbf{tri}}}{\partial u} \cdot N & \frac{\partial \Phi_{\mathbf{tri}}}{\partial v} \cdot T \\ \frac{\partial \Phi_{\mathbf{tri}}}{\partial u} \cdot N & \frac{\partial \Phi_{\mathbf{tri}}}{\partial v} \cdot T \end{pmatrix} \right| = 1. \tag{32}$$

Proof that the parameterization remains in the right domain. This is directly verified by Condition 1.

Proof that the parameterization is injective. This is directly verified by Condition 2.

Proof that the parameterization is surjective. The parameterization is surjective if every point (x,y) in the region under the curve is mapped by a point $(u,v) \in [0,1]^2$. This condition is harder to verify directly than the two previous ones. Fortunately, if Conditions 1 and 2 are met, the parameterization is automatically surjective. First, if the parameterization is area preserving and injective (no overlap), it covers a region of area exactly 1. Second, if it remains inside the region under the curve, it covers a region of area exactly 1 under the curve. Finally, since f is a PDF, the area of the region under the curve is exactly 1 and is thus entirely covered by the parameterization. Hence, any point in the region under the curve is reached by the parameterization.

Validity of the triangle-cut parameterization. In summary, if the approximate density g is such that Equation (24) and Equation (25) are verified, the resulting triangle-cut parameterization is valid, i.e. it is an *area-preserving bijection between the unit square and the region under the curve of f*. In many cases, the two conditions are hard to verify analytically but simple to verify numerically with the implementation provided in Listing 1.

4. Sampling with the Triangle-Cut Parameterization

In this section, we show how to use the triangle-cut parameterization to obtain analytic solutions to sampling problems.

4.1. Sampling 1D Densities

The evaluation of the triangle-cut parameterization yields a straightforward sampling method for the 1D target density f, as shown in Figure 8. Indeed, since the parameterization is area preserving, it maps uniform random numbers $(u,v) \in [0,1]$ to points (x,y) uniformly distributed in the region under the curve of f and the abscissae x of these points are distributed with density f. Note that this method does not preserve the stratification of the random numbers because it maps two random numbers to one. It can be compared to Marsaglia's method that also uses an approximate PDF to exactly sample a target PDF whose inverse CDF is not analytic [Mar84]. However, Marsaglia's method requires a tunable parameter and is a rejection-based method that uses an unbounded number of random numbers. In contrast, our method uses only two random numbers and no tunable parameter.

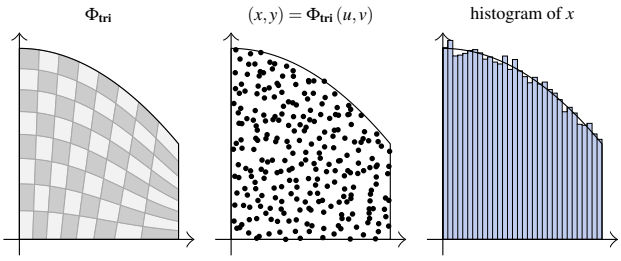

Figure 8: Sampling a 1D density with the triangle-cut parameterization. *The area-preserving parameterization uniformly distributes points in the region under the curve of f and the density of their abscissae is thus f. In this example, the histogram has 32 bins computed with 65536 random points.*

4.2. Stratified Sampling of 2D Densities.

In Section 1, we motivated the research of area-preserving parameterizations of the region under the curve of 1D densities by claiming that in some cases they can be directly used for stratified sampling of 2D densities. We now explain when this applies and how to use the triangle-cut parameterization.

Problem statement. Let us consider a 2D density $f(x_1,x_2)$ that we would like to sample using two uniform random numbers $(u,v) \in [0,1]^2$. A typical sampling approach consists of sampling x_1 from the *marginal* density and x_2 from the *conditional* density:

$$f_{\mathbf{marg}}(x_1) = \int_{-\infty}^{+\infty} f(x_1,x_2)\,\mathrm{d}x_2, \tag{33}$$

$$f_{\mathbf{cond}}(x_2|x_1) = \frac{f(x_1,x_2)}{f_{\mathbf{marg}}(x_1)}, \tag{34}$$

using their inverse CDF. The typical scenario that we are interested in is a 2D density with a non-analytic marginal inverse CDF and an analytic conditional inverse CDF

$$x_1 = F^{-1}_{\textbf{marg}}(u), \qquad \text{not analytic} \tag{35}$$

$$x_2 = F^{-1}_{\textbf{cond}}(v \mid x_1). \qquad \text{analytic} \tag{36}$$

This scenario is quite common. Indeed, many 2D distributions encountered in practical problems have a difficult marginal CDF but a simple conditional CDF, or at least can be formulated such that this is the case, using the appropriate basis, change of variable, etc.

Using the triangle-cut parameterization. Our idea is to define a triangle-cut parameterization $\Phi_{\textbf{tri}}$ of the region under the curve of the marginal distribution $f_{\textbf{marg}}$. By computing $(x,y) = \Phi_{\textbf{tri}}(u,v)$, we obtain a point whose x component is distributed with target density $f_{\textbf{marg}}$ and an interesting property is that the y component can be converted back to a new uniform random number $w \in [0,1]$:

$$w = \frac{y}{f(x)} \tag{37}$$

that is independent of x and thus can be used for something else, as if it were a freshly generated uniform random number. We use w to sample the conditional distribution analytically. In summary, we compute

$$(x_1, y_1) = \Phi_{\textbf{tri}}(u,v), \qquad \text{analytic} \tag{38}$$

$$w = \frac{y_1}{f_{\textbf{marg}}(x_1)}, \qquad \text{analytic} \tag{39}$$

$$x_2 = F^{-1}_{\textbf{cond}}(w \mid x_1). \qquad \text{analytic} \tag{40}$$

The resulting parameterization maps two random numbers (u,v) to a unique random 2D point (x_1, x_2) with density f. We can thus use this approach for stratified sampling of the 2D density f. We use it in Sections 5, 6, and 9.

Stratified Sampling of nD densities. This approach can be trivially extended to nD densities so long as the conditional distribution of the last dimension has an analytic inverse CDF.

5. Application: Stratified Sampling of a Truncated Disk

In this section, we derive a triangle-cut parameterization for a truncated disk, a geometry for which the absence of exact and analytic area-preserving parameterization was recently brought to light [UnG18, PD19].

Target density. In the configuration of Figure 2, the area of the truncated disk and its derivative are:

$$\mathcal{F}(\theta) = \theta - \cos\theta \sin\theta, \tag{41}$$

$$\mathcal{F}'(\theta) = 2\sin^2\theta. \tag{42}$$

The density is proportional to the derivative and the CDF to the area

$$f(\theta) = \frac{\mathcal{F}'(\theta)}{\mathcal{F}(\theta_0)}, \tag{43}$$

$$F(\theta) = \frac{\mathcal{F}(\theta)}{\mathcal{F}(\theta_0)}. \tag{44}$$

There is no analytic solution to $F(\theta) = u$, which prevents an analytic inverse-CDF parameterization.

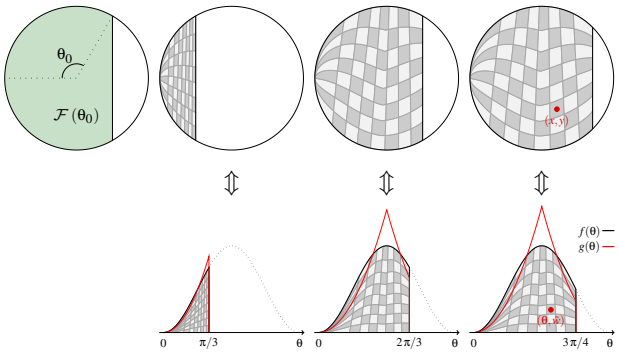

Figure 9: Truncated disk: triangle-cut parameterization. *We plot the target density f in black and the approximate density g in red, for different configurations (with plots scaled for consistency). In each configuration, all the cells of the checker cover the same area.*

Approximate density. We need an approximate density g that is close enough to f. One detail to pay attention to is that $f(0) = f(\pi) = 0$. The triangle-cut method has a singularity when f approaches zero (because the triangle width is computed by dividing by f) and thus it is important to cancel this singularity by ensuring that $\lim_{\theta\to 0} \frac{f(\theta)}{g(\theta)} = \text{constant}$. Therefore, a good choice is a function g whose Taylor expansion at 0 (and π) is of the same order as the one of f. This is why we have chosen an approximation of area $\mathcal{F}$

$$\mathcal{G}(\theta) = \begin{cases} \frac{\theta^3}{3} & \text{if } \theta \in [0, \frac{\pi}{2}], \\ \frac{\pi^3}{12} - \frac{(\pi-\theta)^3}{3} & \text{if } \theta \in (\frac{\pi}{2}, \pi]. \end{cases} \tag{45}$$

that yields the approximate PDF, CDF and iCDF:

$$g(\theta) = \frac{\mathcal{G}'(\theta)}{\mathcal{G}(\theta_0)}, \tag{46}$$

$$G(\theta) = \frac{\mathcal{G}(\theta)}{\mathcal{G}(\theta_0)}, \tag{47}$$

$$G^{-1}(u) = \begin{cases} \sqrt[3]{3u\mathcal{G}(\theta_0)} & \text{if } U \le \frac{\mathcal{G}(\frac{\pi}{2})}{\mathcal{G}(\theta_0)}, \\ \pi - \sqrt[3]{\frac{\pi^3}{4} - 3u\mathcal{G}(\theta_0)} & \text{if } \theta \in (\frac{\pi}{2}, \pi]. \end{cases} \tag{48}$$

We numerically verified that this g satisfies the two conditions of Section 3.2.

Triangle-cut parameterization. The triangle-cut parameterization maps two uniform random numbers (u,v) to an angle θ (the abscissa) distributed with density f and an ordinate $\tilde{w}$. As explained in Section 4.2, we remap the ordinate to obtain a second uniform random number $w = \tilde{w}/f(\theta)$, which we use to sample a height y uniformly in $[-\sin\theta, +\sin\theta]$. In summary:

$$(\theta, \tilde{w}) = \Phi_{\textbf{tri}}(u,v), \qquad \text{analytic} \tag{49}$$

$$w = \frac{\tilde{w}}{f(\theta)}, \qquad \text{analytic} \tag{50}$$

$$(x,y) = (\cos\theta, (-1+2w)\sin\theta), \qquad \text{analytic} \tag{51}$$

where (x,y) is a point uniformly distributed in the truncated disk. The parameterization is shown in Figure 9.

6. Application: Stratified Sampling of a Torus

In this section, we derive an area-preserving parameterization for the surface of a torus.

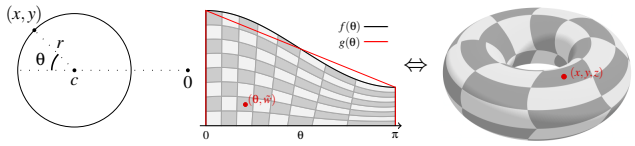

Figure 10: Torus: derivation of the triangle-cut parameterization. *(Left) The torus is a solid of revolution obtained by rotating a circle. The density of the points is proportional to their distance to the axis of rotation. In the remaining figures, we use the torus given by $c = 1$ and $r = \frac{1}{2}$. (Middle) The target density f, the approximate density g and the triangle-cut parameterization. (Right) The triangle-cut parameterization of the torus. All the cells of the checker cover the same area.*

Target density. Figure 10-(left) parameterizes a torus as a solid of revolution by rotating a circle of center $(c,0)$ and radius r:

$$P(\theta) = (c + r\cos\theta, r\sin\theta) \tag{52}$$

Since it is a revolution surface, the density of the points on the surface is proportional to their distance from the axis of rotation, which in this case is the x–axis. To simplify, we consider only the upper part of the torus, i.e. for $\theta \in [0,\pi]$:

$$f(\theta) = \frac{(c + r\cos\theta)}{c\pi}, \tag{53}$$

$$F(\theta) = \frac{(c\theta + r\sin\theta)}{c\pi}. \tag{54}$$

There is no analytic solution to $F(\theta) = u$, which prevents an analytic inverse-CDF parameterization.

Approximate density. We use an affine approximate density that interpolates the values of $f(0) = c + r$ and $f(\pi) = c - r$:

$$g(\theta) = \frac{\frac{\theta}{\pi}(c-r) + \left(1 - \frac{\theta}{\pi}\right)(c+r)}{c\pi}, \tag{55}$$

$$G(\theta) = \frac{\frac{\theta^2}{2\pi}(c-r) + \left(1 - \frac{\theta^2}{2\pi}\right)(c+r)}{c\pi}, \tag{56}$$

$$G^{-1}(u) = \pi\frac{\sqrt{(c+r)^2 + u\left((c-r)^2 - (c+r)^2\right)} - (c+r)}{-2r}. \tag{57}$$

We numerically verified that this g satisfies the two conditions of Section 3.2 for $c = 1$ and $r = \frac{1}{2}$.

Triangle-cut parameterization. The triangle-cut parameterization maps two uniform random numbers (u,v) to an angle θ (the abscissa) distributed with density f and an ordinate $\tilde{w}$. As explained in Section 4.2, we remap the ordinate to obtain a second uniform random number $w = \tilde{w}/f(\theta)$, which we use to sample a uniform rotation $\phi \in [0, 2\pi]$ around the y–axis. In summary:

$$(\theta, \tilde{w}) = \Phi_{\text{tri}}(u,v), \qquad \text{analytic} \tag{58}$$

$$w = \frac{\tilde{w}}{f(\theta)}, \quad \phi = 2\pi w, \qquad \text{analytic} \tag{59}$$

$$(x,y,z) = (t\cos\phi, r\sin\theta, t\sin\phi), t = (c + r\cos\theta), \qquad \text{analytic} \tag{60}$$

where (x,y,z) is a point uniformly distributed on the upper part of the torus ($\theta \in [0,\pi]$). We can trivially wrap this parameterization to both parts of the torus ($\theta \in [0,2\pi]$) since they are symmetric. The parameterization of the full torus is shown in Figure 10-(right).

7. Application: Stratified Sampling of a Polar Shape

We derive a triangle-cut parameterization for the polar shape defined by its radius

$$r(\theta) = 1 + \frac{1}{8}\cos(8\theta) + \frac{1}{16}\cos(16\theta), \tag{61}$$

as shown in Figure 11. The density associated with the angle of a polar shape is proportional to the squared radius, i.e. the target density is

$$f(\theta) \propto r^2(\theta). \tag{62}$$

The CDF is a sum of products of cosines and sines and a linear term that cannot be analytically inverted. We compute a triangle-cut parameterization with an approximate density that is uniform over $[0, 2\pi]$.

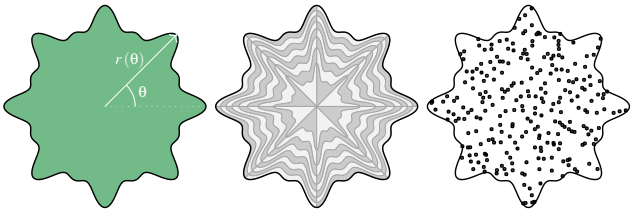

Figure 11: Polar shape: triangle-cut parameterization. *The parameterization is area preserving, i.e. each cell of the checker has the same area.*

8. Application: Stratified Sampling of a Polynomial Density

We derive a triangle-cut parameterization for the 2D polynomial density

$$f(x,y) = \begin{cases} \frac{120}{83}\left(1 + x - x^2 + x^3 - x^4 + x^5\right) y & \text{if } (x,y) \in [0,1], \\ 0 & \text{otherwise.} \end{cases} \tag{63}$$

that is shown in Figure 12. The CDF of the marginal over x is a 6th–order polynomial and cannot be analytically inverted. We compute a triangle-cut parameterization with an approximate density that is uniform over $[0,1]$.

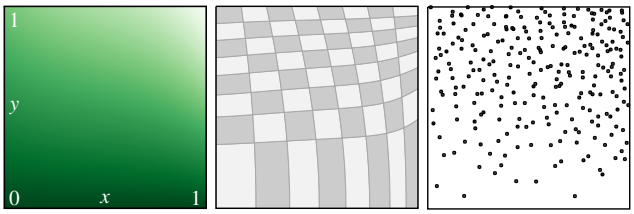

Figure 12: Polynomial density: triangle-cut parameterization. *The parameterization is integral preserving, i.e. the integral of the density is the same inside each cell of the checker.*

9. Application: Stratified Sampling of Disney's BSSRDF

In this section, we derive a triangle-cut parameterization for Disney's BSSRDF [Bur15]. The original paper claimed that the CDF is not analytically invertible and the solution remained unknown for several years until a derivation was recently found [Gol19]. We wrote this section of the paper before Golubev's blog post was brought to our attention and our recommendation to today's practitioners is to use this analytic inverse CDF. Nonetheless, we believe that it is still a worthwhile exercise to explore how the triangle-cut parameterization could have been be used as a substitute at the time when the analytic inverse CDF was still unknown. The point is to showcase the potential of our approach in general, not to make a competitive contribution to this specific application.

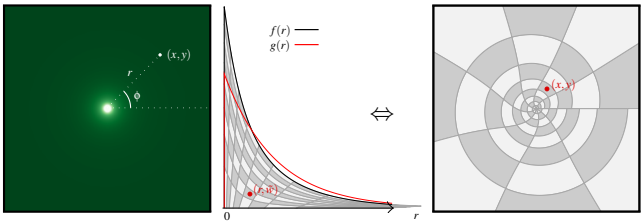

Figure 13: BSSRDF: triangle-cut parameterization. *(Left) The radially symmetric diffusion profile. (Middle) The target density f, the approximate density g and the triangle-cut parameterization. (Right) The triangle-cut parameterization of the BSSRDF. All cells of the checker cover the same area.*

Diffusion profile. Burley uses a radial diffusion profile that is the sum of two exponential distributions:

$$R_d(r) = \frac{\exp\left(-\frac{r}{d}\right) + \exp\left(-\frac{r}{3d}\right)}{8\pi d r} \tag{64}$$

It is represented in Figure 13-(left).

Target density. The target density is the radial diffusion profile multiplied by the radius:

$$f(r) = \frac{\exp\left(-\frac{r}{d}\right) + \exp\left(-\frac{r}{3d}\right)}{4d}, \tag{65}$$

$$F(r) = \frac{4d - d\exp\left(-\frac{r}{d}\right) - 3d\exp\left(-\frac{r}{3d}\right)}{4d}. \tag{66}$$

Burley claimed that there was no analytic solution to $F(r) = u$, which prevented an analytic inverse-CDF parameterization. In the following, we show how to obtain one in the absence of this analytic solution.

Approximate density. We need an approximate density g that is close enough to f. One detail to pay attention to is that $\lim_{r\to\infty} f(r) = 0$ and thus it is important to cancel this singularity by ensuring that $\lim_{r\to\infty} \frac{f(\theta)}{g(\theta)} = \text{constant}$. Therefore, a good choice is a function g whose convergence is the same order as the one of f, which is why we choose the widest exponential distribution of the sum:

$$g(r) = \frac{\exp\left(-\frac{r}{3d}\right)}{3d}, \tag{67}$$

$$G(r) = 1 - \exp\left(-\frac{r}{3d}\right), \tag{68}$$

$$G^{-1}(r) = 3d\log(1-u). \tag{69}$$

We numerically verified that this g satisfies the two conditions of Section 3.2.

Triangle-cut parameterization. The triangle-cut parameterization maps two uniform random numbers (u,v) to a radius r (the abscissa) distributed with density f and an ordinate $\tilde{w}$. As explained in Section 4.2, we remap the ordinate to obtain a second uniform random number $w = \tilde{w}/f(r)$, which we use to sample a uniform azimuthal angle $\phi \in [0, 2\pi]$. In summary:

$$(r, \tilde{w}) = \Phi_{\mathbf{tri}}(u,v), \qquad \text{analytic} \tag{70}$$

$$w = \frac{\tilde{w}}{f(r)}, \quad \phi = 2\pi w, \qquad \text{analytic} \tag{71}$$

$$(x,y) = (r\cos\phi, r\sin\phi), \qquad \text{analytic} \tag{72}$$

where (x,y) is a point distributed with the diffusion profile. The parameterization is shown in Figure 13-(right).

Convergence analysis. An important criterion for Monte Carlo rendering is whether the stratification allows for taking advantage of quasi-random low-discrepancy sequences such as Sobol's [Sob67]. In Figure 14, we render the diffusion of a laser beam in a slab with different integration techniques:

INDEPENDENT: we sample the diffusion profile with independent uniform random numbers. The convergence is the same regardless of the parameterization.

SOBOL + RANDOM LOBE: we use quasi-random 3D points $(u,v,w) \in [0,1]^3$ of the Sobol sequence. We use u to randomly choose the exponential lobe ($d' = d$ or $d' = 3d$) and we sample this lobe with (v,w) using the analytic inverse-CDF parameterization. This is the sampling technique proposed by Burley [Bur15].

SOBOL + ICDF: we use quasi-random 2D points $(u,v) \in [0,1]^2$ of the Sobol sequence and we map them to the diffusion profile with the inverse-CDF parameterization $\Phi_{F^{-1}}$ using a numerical inversion for F^{-1}.

SOBOL + TRIANGLE CUT: we use quasi-random 2D points $(u,v) \in [0,1]^2$ of the Sobol sequence and we map them to the diffusion profile with the triangle-cut parameterization $\Phi_{\mathbf{tri}}$.

We compare the convergence of the techniques in Figure 15. As expected, the variance decreases linearly with independent random numbers. The convergence speed-up expected from the Sobol sequence is suboptimal with the random-lobe technique. This is because the two lobes overlap each other and thus the stratification is not preserved. The usage of the Sobol sequence is optimal with 2D area-preserving parameterizations that make stratified sampling possible. The convergence with the inverse CDF and the triangle-cut parameterizations are equivalent.

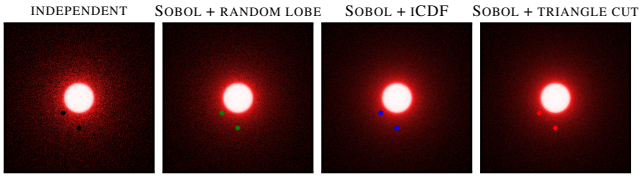

Figure 14: BSSRDF: visualization of the diffusion on a slab lit by a laser beam. *The reference image is the laser beam footprint (a uniform disk) convolved with the BSSRDF diffusion profile shown in Figure 13. For each pixel, we sample the diffusion profile. If the samples are located inside the laser beam, they add energy to the pixel estimate. These images are rendered at 512 spp, and the convergence curves of the two marked pixels are shown in Figure 15.*

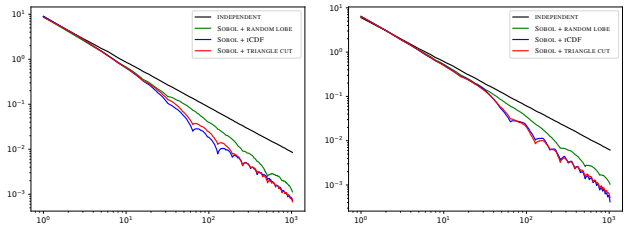

Figure 15: BSSRDF: convergence analysis of the different sampling techniques. *We compare the variance of the different techniques as a function of the number of samples, for the two marked pixels in Figure 14.*

Performance. In Table 1, we compare the accuracy and the computational time of the inverse CDF and the triangle-cut parameterizations, which have the same convergence. We compute the numerical inverse CDF by starting from a first guess (obtained with G^{-1} for a fair comparison) and refined using Newton iterations. For each generated sample $r \approx F^{-1}(u)$, we average the absolute error $|u - F(r)|$ over $u \in [0,1]$. The more iterations, the more the error decreases but the costlier the inversion becomes. Meanwhile, the triangle cut is error free (at least within the bounds of floating point precision).

Summary. The triangle-cut parameterization provides the same convergence as the inverse-CDF parameterization with quasi-random numbers but without the cost of a high-quality numerical inversion and with the guarantee that the result is error free. Hence, before the recent introduction of the analytic form of the inverse CDF [Gol19], the triangle-cut parameterization would have been the best option.

Table 1: BSSRDF: performance comparison. *We compare the errors and the timings of the triangle-cut parameterization to a numerical inverse CDF with Newton iterations. The timings are provided for the generation of* 10^7 *samples on an Intel i7-5960X CPU.*

	iCDF	iCDF	iCDF	iCDF	triangle cut
iterations	0	1	2	3	n/a
performance	0.11s	0.42s	0.78s	1.26s	0.61s
av. abs. error	6.2e-2	8.5e-3	1.4e-4	7e-8	n/a

10. Conclusion

The goal of this paper was to overcome the problem posed by non-analytically invertible CDFs. We have introduced an alternative parameterization, which we call *triangle cut*, that can be analytically evaluated, and we have proven that, under certain conditions, it is a valid area-preserving bijection of the region under the curve of the target density. We have shown, through different examples, how to apply it to sampling 1D densities, to stratified sampling of 2D shapes, and to densities with non-analytic marginal inverse CDFs and analytic conditional inverse CDFs. In the example of Sec. 9, we have seen that 2D stratified sampling with the triangle-cut parameterization has similar convergence properties to the classic inverse-CDF approach but with a fully analytic solution and guaranteed bias-free result.

The immediate potential impact of the triangle-cut parameterization is that it might allow us to revisit countless classic sampling problems from different fields that currently lack analytic solutions. We have seen that, in some cases, successfully using the triangle-cut parameterization requires some crafting. The main difficulty is in finding an approximate density that satisfies the two mandatory conditions. In Sec. 5, we crafted an approximation that has the same Taylor expansion as the target density where it evaluates to 0. In Sec. 9, we crafted an approximation that has the same convergence towards 0 in the limit to $+\infty$ as the target density.

Finally, we believe that the most important takeaway is that the world of analytic sampling techniques is not tied to the inverse CDF. The triangle cut is just one parameterization among an infinity of alternatives, and we hope that it will inspire the research of original approaches to analytic sampling.

Acknowledgements

This paper was written during an intensive and weird work-from-home period due to COVID-19. I spare a thought for people who did a job more important and more dangerous than nerding a sampling paper in a safe environment. Special thanks to Sabrina for supporting me during this period and an affectionate thought for Babs who passed away. Thanks to Laurent Belcour, Jonathan Dupuy, Kenneth Vanhoey and Stephen Hill for their help and feedback. Thanks also to the reviewers for their valuable suggestions.

References

[AN07] ARVO J., NOVINS K.: Stratified sampling of convex quadrilaterals. *J. Graphics Tools 12* (01 2007), 1–12. 2, 5

[Arv95] ARVO J.: Stratified sampling of spherical triangles. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques* (1995), SIGGRAPH 1995, pp. 437–438. 2

[Arv01] ARVO J.: Stratified sampling of 2-manifolds. In *State of the Art in Monte Carlo Ray Tracing for Realistic Image Synthesis. ACM SIGGRAPH Courses* (2001). 2

[Bur15] BURLEY B.: Extending the disney BRDF to a BSDF with integrated subsurface scattering. In *Physically Based Shading in Theory and Practice. ACM SIGGRAPH Courses* (2015). 10

[Dev86] DEVROYE L.: *Non-Uniform Random Variate Generation*. Springer, 1986. 1

[Gam16] GAMITO M. N.: Solid angle sampling of disk and cylinder lights. *Computer Graphics Forum 35*, 4 (2016), 25–36. 2

[Gol19] GOLUBEV E.: Sampling Burley's Normalized Diffusion Profiles. blog post, 2019. URL: https://zero-radiance.github.io/post/sampling-diffusion/. 10, 11

[GUnK*17] GUILLÉN I., UREÑA C., KING A., FAJARDO M., GEORGIEV I., LÓPEZ-MORENO J., JARABO A.: Area-preserving parameterizations for spherical ellipses. *Computer Graphics Forum 36*, 4 (2017), 179–187. 2

[Hd14] HEITZ E., D'EON E.: Importance sampling microfacet-based BSDFs using the distribution of visible normals. *Computer Graphics Forum 33*, 4 (2014). 2

[Hei18] HEITZ E.: Sampling the GGX distribution of visible normals. *Journal of Computer Graphics Techniques (JCGT) 7*, 4 (November 2018), 1–13. 2

[Hei19] HEITZ E.: A Low-Distortion Map Between Triangle and Square. technical report, 2019. URL: https://hal.archives-ouvertes.fr/hal-02073696v2/. 2

[Mar84] MARSAGLIA G.: The exact-approximation method for generating random variables in a computer. *Journal of the American Statistical Association 79*, 385 (1984), 218–221. 7

[MLM18] MARTINO L., LUENGO D., MIGUEZ J.: *Independent Random Sampling Methods*. 2018, pp. 65–113. 4

[MMR*19] MÜLLER T., MCWILLIAMS B., ROUSSELLE F., GROSS M., NOVÁK J.: Neural importance sampling. *ACM Trans. Graph. 38*, 5 (2019). 2

[Mul56] MULLER D. E.: A method for solving algebraic equations using an automatic computer. *Mathematical Tables and Other Aids to Computation 10*, 56 (1956), 208–215. 5

[PD19] PETERS C., DACHSBACHER C.: Sampling projected spherical caps in real time. *Proc. ACM Comput. Graph. Interact. Tech. 2*, 1 (2019). 2, 8

[SC97] SHIRLEY P., CHIU K.: A low distortion map between disk and square. *J. Graph. Tools 2*, 3 (Dec. 1997), 45–52. 2

[Shi91] SHIRLEY P.: Discrepancy as a Quality Measure for Sample Distributions. In *EG 1991-Technical Papers* (1991), Eurographics Association. 1

[Sob67] SOBOL I. M.: The distribution of points in a cube and the approximate evaluation of integrals. In *USSR Computational Mathematics and Mathematical Physics 7* (1967), pp. 86–112. 10

[SWZ96] SHIRLEY P., WANG C., ZIMMERMAN K.: Monte carlo techniques for direct lighting calculations. *ACM Trans. Graph. 15*, 1 (Jan. 1996), 1–36. 2

[Tur90] TURK G.: Generating random points in triangles. In *Graphics Gems* (1990), Academic Press, pp. 24–28. 2

[UnFK13] UREÑA C., FAJARDO M., KING A.: An Area-Preserving Parametrization for Spherical Rectangles. *Computer Graphics Forum* (2013). 2

[UnG18] UREÑA C., GEORGIEV I.: Stratified sampling of projected spherical caps. *Computer Graphics Forum 37*, 4 (2018), 13–20. 2, 8

[Zha19] ZHANG J.: On sampling of scattering phase functions. *Astronomy and Computing 29* (2019), 100329. 2

[ZZ19] ZHENG Q., ZWICKER M.: Learning to importance sample in primary sample space. *Comput. Graph. Forum 38* (2019), 169–179. 2

Appendix A: Derivation of the Projected Partial Derivatives

We explain how we obtain Equations (28), (29), (30), and (31). One of the key simplifications comes from the fact that P_a and P_b are the end points of the partitioning segment and N and T are its normal and tangent, respectively (see Figure 5). We thus have

$$(P_a - P_b)\cdot N = 0, \tag{73}$$
$$(P_a - P_b)\cdot T = \|P_a - P_b\|. \tag{74}$$

- For Equation (28), we use Equation (73) to remove $\left(\frac{\partial t}{\partial u}(u,v)\,(P_a(u) - P_b(u))\right)\cdot N = 0$ and from Equations (17) and (18), we obtain $\frac{\partial P_a}{\partial u}\cdot N = w_a$ and $\frac{\partial P_b}{\partial u}\cdot N = w_b$.
- The result of Equation (29) is not required because in the Jacobian matrix determinant, $\frac{\partial \Phi_{\mathbf{tri}}}{\partial u}\cdot T$ is multiplied by $\frac{\partial \Phi_{\mathbf{tri}}}{\partial v}\cdot N = 0$.
- For Equation (30), we use Equation (73).
- For Equation (31), we use Equation (74) to obtain $\frac{\partial \Phi_{\mathbf{tri}}}{\partial v}\cdot T = \frac{\partial t}{\partial v}(u,v)\,\|P_a - P_b\|$. Second, since $t = W^{-1}(v)$, its derivative with respect to v is inversely proportional to the density function w:

$$\frac{\partial t}{\partial v}(u,v) = \frac{1}{w(t)} = \frac{\frac{w_a+w_b}{2}}{t\,w_a + (1-t)\,w_b}. \tag{75}$$

Finally, as shown in Figure 5, the area du of the infinitesimal convex quadrilateral is

$$\begin{aligned} du &= \text{segment length}\,\frac{\text{vector1}\cdot\text{normal} + \text{vector2}\cdot\text{normal}}{2} \\ &= \|P_a - P_b\|\,\frac{(dP_a\cdot N + dP_b\cdot N)}{2}, \end{aligned} \tag{76}$$

and the length of the segment

$$\|P_a - P_b\| = \frac{2}{\frac{dP_a}{du} + \frac{dP_b}{du}} = \frac{2}{w_a + w_b} \tag{77}$$

cancels out in Equation (31):

$$\frac{\partial \Phi_{\mathbf{tri}}}{\partial v}\cdot T = \frac{\partial t}{\partial v}(u,v)\,\|P_a - P_b\| \tag{78}$$
$$= \frac{\frac{w_a+w_b}{2}}{t\,w_a + (1-t)\,w_b}\,\frac{1}{\frac{w_a+w_b}{2}} = \frac{1}{t\,w_a + (1-t)\,w_b}. \tag{79}$$