


EON: A Practical Energy-Preserving Rough Diffuse BRDF

Jamie Portsmouth 
Autodesk

Peter Kutz 
Adobe

Stephen Hill 
Lucasfilm




Figure 1. Lambert model (left), the classic “qualitative” Oren–Nayar QON model (center), and our new EON model (right), all at maximal roughness and the same single-scattering albedo color. The top and bottom rows show front and side-lit cases respectively. The EON model compensates for the energy loss of the QON model, producing similar brightness to the Lambert model combined with the characteristic “flattening” and enhanced backscattering of the Oren–Nayar model.

Abstract

We introduce the *energy-preserving Oren–Nayar* (EON) model for reflection from rough surfaces. Unlike the popular qualitative Oren–Nayar model (QON) and its variants, our model is energy preserving via analytical energy compensation. We include self-contained GLSL source code for efficient evaluation of the new model and importance sampling based on a novel technique we term *Clipped Linearly Transformed Cosine* (CLTC) sampling.

1. Introduction

The original rough diffuse surface reflectance model of Oren and Nayar [1994] has been highly influential in computer graphics. They noted that real materials do not necessarily exhibit Lambertian behavior, but instead appear “flatter” with less edge darkening and lower contrast. To explain this, they introduced a V-cavity–based microfacet model (in the geometrical optics limit), parameterized by a scalar roughness parameter σ and single-scattering albedo color ρ , where Lambertian behavior is recovered in the $\sigma \rightarrow 0$ limit. This exhibits a peak in backscattering at grazing configurations (i.e., when the incident and outgoing photon directions are both almost parallel to the surface and almost opposite in direction), which aligns with the behavior of real materials. Figure 2 shows examples of the flat and bright appearance of the various Oren–Nayar models compared to the Lambertian model, caused by this backscattering.

Oren and Nayar [1994] presented their full model as well as a more approximate “qualitative” version. The full version (henceforth the *full Oren–Nayar* model, or fullON) accounts for both single and double scattering (due to interreflections) on the microfacet surface. Figure 3 shows the fullON BRDF as the blue curve, where the backscattering peak is most prominent in the grazing configuration in the upper right panel.

The cheaper “qualitative” variant of the model—henceforth the *qualitative Oren–Nayar* model or QON for short—ignores the double scattering interreflections. We review the mathematical form of the QON model in Section 2.1. This QON model has been widely used in the industry as the fullON model is more complex to implement

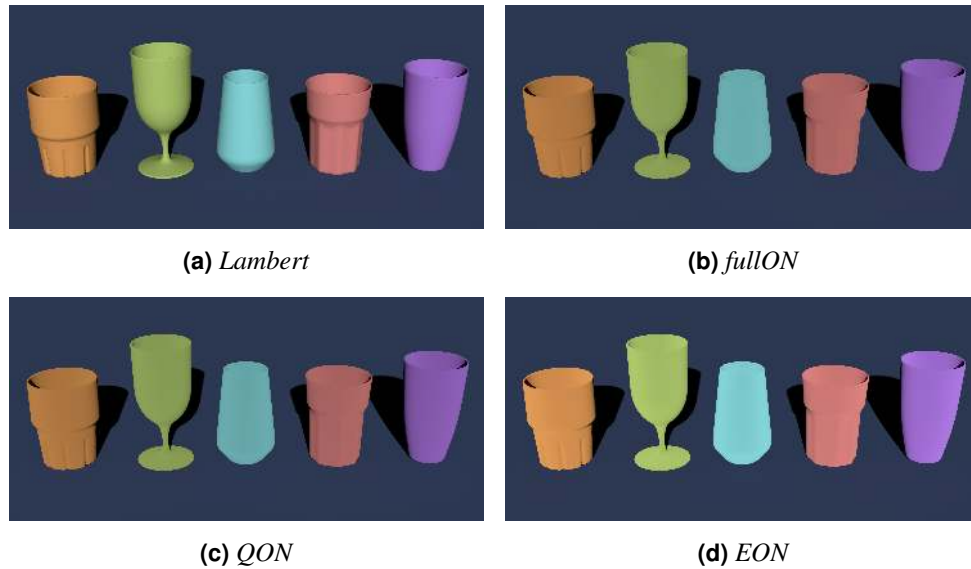


Figure 2. Backscattering from smooth (top left) and maximally rough diffuse BRDFs.

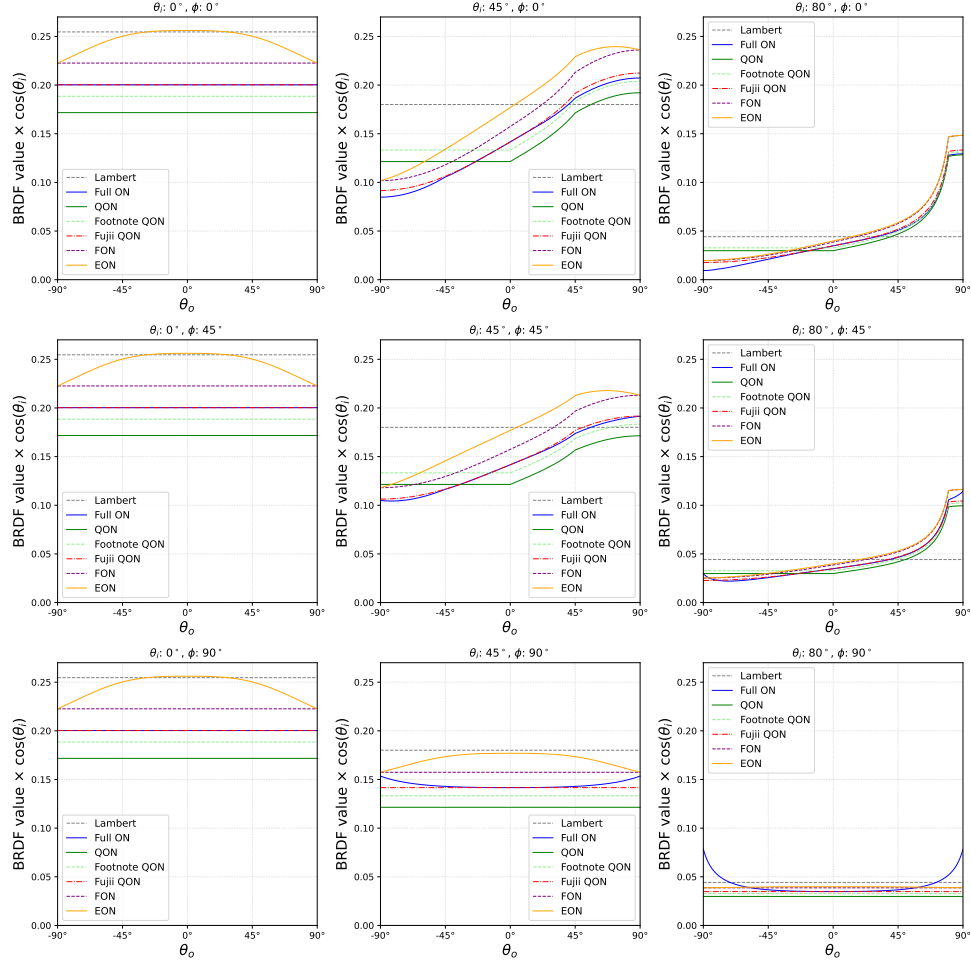


Figure 3. Comparison of the various Oren–Nayar models discussed in Section 1, plotted for various input and output angles at 80% albedo and 50% roughness (corresponding to BRDF input parameters $\sigma = \pi/4$ or $r = 1/2$ depending on the model). Note that the BRDF values are multiplied by a foreshortening factor of $\cos \theta_i$, where θ_i is the zenith angle of the light direction. These plots align with those in the article by Fujii [2013], which uses the same conventions as those in the original paper by Oren and Nayar [1994].

and costly to evaluate (e.g., in MaterialX [Smythe and Stone 2018], OSR [Gritz et al. 2010], MDL [Kettner et al. 2015], and Arnold [Georgiev et al. 2018]).

Both of the original Oren–Nayar models, fullON and QON, are reciprocal but not energy preserving. Indeed, both models violate energy preservation (i.e., lose energy) and energy conservation (i.e., gain energy) for certain configurations of directions and roughness values. In addition, the QON variant (the dark green curve in Figure 3) is overly bright when backlit and exhibits a subtle discontinuity artifact in the shading, which appears as a dark ring when rendering a sphere [Fujii 2013].

A modification to the original models of Oren–Nayar was later proposed by Fujii [2013], who noted that the fullON model can be reasonably approximated via simple modification of the coefficients of the QON model (this *Fujii QON* model is the dot-dashed bright red curve in Figure 3). This model improves upon QON—incorporating a tinted interreflection component, eliminating excessive forward scattering, and eliminating the dark ring artifact—but still suffers from the energy loss issue of the fullON and QON models. Fujii thus proposed a second, simplified qualitative model that we term the *Fujii Oren–Nayar* model or FON, which we discuss in detail in Section 2.2. The FON model is the dashed darker red curve in Figure 3. This model was designed to closely match the shape of the fullON model at lower computational cost while reducing energy loss and fixing the discontinuity artifact. While this model is brighter than fullON in most cases, it is also energy conserving. However, the FON model still does not preserve energy and thus still appears to darken for high roughness values (and thus fails a white-furnace test). Fujii [2013] also presented a numerical comparison of all these models, which matches our Figure 3.¹ This plot therefore also serves as a check that our interpretation of the FON model is correct, since the statement of it was somewhat unclear.

In Section 3 we present an analytical extension to the FON model that is reciprocal and explicitly energy preserving (thus passing a white-furnace test), which we term *energy-preserving Oren–Nayar*, or EON. EON is the orange curve in Figure 3. This augments the FON model with a term representing the full multiple-scattering contribution missing from the other models. It is analytical and straightforward to implement, and a high-quality fit can be found to its functional form for improved efficiency. While EON (like FON) is not directly derived from an underlying micro-physical theory, it provides a more practical model than any of the preceding Oren–Nayar forms while preserving the qualitative appearance, which is both energy conserving *and* energy preserving (and thus brighter and more visually appealing), while being as efficient and easy to implement as the existing standard QON model. In Section 4 we provide an analytical importance sampling scheme for EON that takes into account the backscattering peak (using a technique we introduce termed *Clipped Linearly Transformed Cosine* sampling, or CLTC), significantly improving sample variance compared to cosine-weighted hemispherical sampling. Figure 4 contains renders of all of the BRDFs that are featured in Figure 3 for comparison.

¹Note, the light green curve *Footnote QON* corresponds to this footnote in Oren and Nayar [1994]:

Discrepancies caused by the lack of the interreflection component in the qualitative model can be partially compensated by replacing the constant 0.33 in coefficient A with 0.57.




Figure 4. Visual comparison of the Lambert model to the various Oren–Nayar models.

2. Existing Oren–Nayar Models

2.1. Qualitative Oren–Nayar (QON)

Oren and Nayar [1994] provided a simplified “qualitative” version of their model (abbreviated QON), which is cheaper to evaluate and thus more suitable for real-time use cases. We briefly review here the form and properties of this model, and derive a novel analytical formula for its albedo.

The QON model takes the form (for incident and outgoing rays ω_i and ω_o , respectively)

$$\mathbf{f}_q(\omega_i, \omega_o) = \frac{\boldsymbol{\rho}}{\pi} (A_q + B_q g_q(\omega_i, \omega_o)), \quad (1)$$

where $\boldsymbol{\rho}$ is an RGB color parameter² controlling the overall albedo and the coefficients

²Note that RGB color parameters such as $\boldsymbol{\rho}$ are denoted in bold throughout.

A_q, B_q are given in terms of a roughness parameter σ :

$$A_q = 1 - 0.5 \frac{\sigma^2}{\sigma^2 + 0.33},$$

$$B_q = 0.45 \frac{\sigma^2}{\sigma^2 + 0.09}.$$

Thus in the limit $\sigma \rightarrow 0$, $\mathbf{f}_q \rightarrow \rho/\pi$, which is the Lambert BRDF. In terms of the polar angles θ_i, θ_o of the incident and outgoing rays to the normal and the azimuthal angles ϕ_i, ϕ_o , the term $g_q(\omega_i, \omega_o)$ is

$$g_q(\omega_i, \omega_o) = \cos^+(\phi_i - \phi_o) \sin \alpha \tan \beta, \quad (2)$$

where³ $x^+ := \max(x, 0)$, $\alpha := \max(\theta_i, \theta_o)$, and $\beta := \min(\theta_i, \theta_o)$. Note that the QON BRDF \mathbf{f}_q is symmetrical under the interchange $\omega_i \leftrightarrow \omega_o$, and thus reciprocal. It can equivalently be written more compactly in vector form with $g_q = s/t_q$:

$$\mathbf{f}_q(\omega_i, \omega_o) = \frac{\rho}{\pi} \left(A_q + B_q \frac{s}{t_q} \right) \quad (3)$$

with the term s given by (with normal N)

$$s = \omega_i \cdot \omega_o - (N \cdot \omega_i)(N \cdot \omega_o) = \cos(\phi_i - \phi_o) \sin \theta_i \sin \theta_o$$

and term $1/t_q$ given by

$$\frac{1}{t_q} = \begin{cases} 0 & \text{if } s \leq 0 \\ 1 / \max(N \cdot \omega_i, N \cdot \omega_o) & \text{if } s > 0. \end{cases} \quad (4)$$

According to [Oren and Nayar \[1994\]](#), the σ parameter is technically an angle giving the standard deviation of the distribution of microfacet angles to the horizontal. We adopt the range $\sigma \in [0, \pi/2]$ as suggested by [Fujii \[2013\]](#).

The left panel of Figure 5 shows a render of an object (the brown shader ball) with the QON BRDF at $\sigma = \pi/2$. The right panel shows a QON white-furnace test (i.e., uniform illumination of the $\rho = 1$, $\sigma = \pi/2$ QON material, at 50 bounces) where, with an energy-preserving BRDF, the object would blend into the background and disappear at infinite bounce count. Clearly the QON model loses a significant amount of energy. The left panel of Figure 5 shows the subtle “dark ring” artifact near the top left of the brown shader ball, which was pointed out by [Fujii \[2013\]](#). This is due to the lack of continuity of the BRDF as it passes through $s = 0$.

We can verify this (numerically and analytically) by computing the directional albedo of the QON model $\mathbf{E}_q(\omega_o)$, which is the energy reflected into direction ω_o

³A notation borrowed from [Akenine-Möller et al. \[2018\]](#).




Figure 5. QON model single-bounce render and white-furnace test at $\sigma = \pi/2$. The dotted blue line in the enhanced contrast inset indicates where the subtle “dark ring” artifact occurs.

under uniform illumination. This is given by the following integral of the BRDF w.r.t. to projected solid angle element $d\omega_i^\perp$, over the positive hemisphere \mathcal{H}_+ of ω_i :

$$\mathbf{E}_q(\omega_o) = \int_{\mathcal{H}_+} \mathbf{f}_q(\omega_i, \omega_o) d\omega_i^\perp = \boldsymbol{\rho} \left(A_q + \frac{B_q}{\pi} G_q(\omega_o) \right), \quad (5)$$

where

$$G_q(\omega_o) = \int_{\mathcal{H}_+} g_q(\omega_i, \omega_o) d\omega_i^\perp. \quad (6)$$

The integral for $G_q(\omega_o)$ factors into separate integrals over θ_i and ϕ_i :

$$\int_0^{\pi/2} \sin \theta_i \cos \theta_i \sin \alpha \tan \beta d\theta_i \int_0^{2\pi} \cos^+(\phi_i - \phi_o) d\phi_i. \quad (7)$$

The θ_i integral breaks into two terms according to whether $\theta_i < \theta_o$ or $\theta_i \geq \theta_o$ (due to the definitions of α and β), which are easily evaluated. The ϕ_i integral reduces to $\int_{-\pi/2}^{\pi/2} \cos \phi_i d\phi_i = 2$. The resulting formula for G_q is

$$G_q(\theta_o) = \sin \theta_o (\theta_o - \sin \theta_o \cos \theta_o) + \frac{2}{3} \tan \theta_o (1 - \sin^3 \theta_o). \quad (8)$$

Note that this is undefined at exactly grazing angle $\theta_o = \pi/2$, but the limit is well-defined: $G_q(\theta_o) \rightarrow \pi/2$ as $\theta_o \rightarrow \pi/2$. Thus, we have derived an analytical formula for the directional albedo of the QON model.

Figure 6 shows the directional albedo $\hat{E}_q(\omega_o)$ as a function of view angle θ_o , for various σ roughnesses (the hat on \hat{E}_q denotes that the albedo is evaluated at $\boldsymbol{\rho} = 1$).

The term $\langle \hat{E}_q \rangle = \int_{\mathcal{H}_+} \hat{E}_q(\omega_o) d\omega_o^\perp / \pi$ is the *average albedo*, which can be computed analytically as

$$\langle \hat{E}_q \rangle = A_q + \left(\frac{2}{3} - \frac{64}{45\pi} \right) B_q. \quad (9)$$

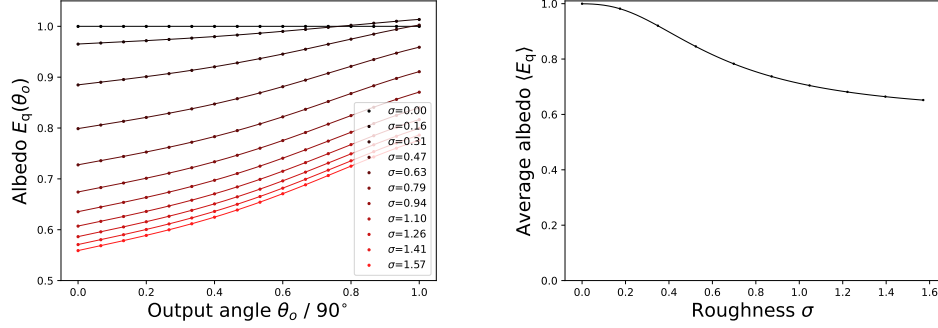


Figure 6. QON model albedos. Directional albedo \hat{E}_q (left) and average albedo $\langle \hat{E}_q \rangle$ (right), assuming $\rho = 1$ (the lines are the analytical formula given by Equations 5 and 8, and the dots are the numerically evaluated integral).

Figure 6 shows the variation of the average albedo $\langle \hat{E}_q \rangle$ as a function of σ . This corresponds to the total fraction of incident light energy that is reflected (into any direction) under uniform incident illumination. As was apparent from the white-furnace test, the QON model suffers from significant energy loss at high σ , i.e., does not preserve energy. It also violates energy conservation (i.e., the albedo exceeds 1) for grazing angles and high σ .⁴

2.2. Fujii Oren–Nayar (FON)

A modification to the qualitative Oren–Nayar model that reduces the dark ring artifact and avoids violations of energy conservation was proposed by Fujii [2013] (henceforth abbreviated FON). This has the same form as Equation 3, i.e.,

$$f_F(\omega_i, \omega_o) = \frac{\rho}{\pi} \left(A_F + B_F \frac{s}{t_F} \right), \quad (10)$$

but with modified t term (compared to Equation 4), changing the $s \leq 0$ behavior,

$$\frac{1}{t_F} = \begin{cases} 1 & \text{if } s \leq 0 \\ 1 / \max(N \cdot \omega_i, N \cdot \omega_o) & \text{if } s > 0 \end{cases} \quad (11)$$

and with modified A, B coefficients written in terms of a roughness parameter r :

$$A_F = \frac{1}{1 + \left(\frac{1}{2} - \frac{2}{3\pi}\right)r},$$

$$B_F = r A_F.$$

In this model, the roughness parameter r is no longer an angle (like the σ of fullON or QON) but instead an interpolation weight in the range $r \in [0, 1]$.

⁴It can be shown that the fullON model also violates energy conservation near grazing angles, except at high roughness. It also does not preserve energy.




Figure 7. FON model albedos: Directional albedo \hat{E}_F (left) and average albedo $\langle \hat{E}_F \rangle$ (center), assuming $\rho = 1$ (the lines are the analytical formula given by Equations 12 and 13, and the dots are the numerically evaluated integral). The right panel shows the difference between the analytical and approximate forms of the albedo given by Equations 13 and 14, respectively.

The directional albedo of the FON model, $\mathbf{E}_F(\omega_o)$, is given by

$$\mathbf{E}_F(\omega_o) = \rho \left(A_F + \frac{B_F}{\pi} G_F(\omega_o) \right) = \rho \hat{E}_F(\omega_o), \quad (12)$$

where (referring to Equation 8)

$$G_F(\theta_o) = G_q(\theta_o) - \frac{2}{3} \sin \theta_o. \quad (13)$$

This is plotted in the left panel of Figure 7, for the $\rho = 1$ case. Note that $\hat{E}_F(\omega_o) \rightarrow 1$ as $\theta_o \rightarrow \pi/2$ (for any roughness r), i.e., the model is energy preserving only at grazing view angles (and energy conserving at all other view angles). In fact, the coefficients A_F , B_F were apparently determined by requiring that $\hat{E}_F(\pi/2) = 1$ and $B_F = r A_F$.

Here we provide a good approximation to the function $G_F(\theta_o)$ that is more efficient to evaluate than the expression in Equation 13 (since the only trigonometric term is $\cos \theta_o$, which is trivial in local space), given by the sum

$$G_F(\theta_o) \approx \pi \sum_{k=1}^4 g_k (1 - \cos \theta_o)^k, \quad (14)$$

with coefficients g_k as provided in the following table:

g_1	g_2	g_3	g_4
0.0571085289	0.491881867	-0.332181442	0.0714429953

The difference between the exact albedo evaluated using Equation 13 and the approximate albedo evaluated using Equation 14 is less than 0.1% over the whole angular range, as shown in the right panel of Figure 7.

The average albedo of the FON model, $\langle \hat{E}_F \rangle = \int_{\mathcal{H}_+} \hat{E}_F(\omega_o) d\omega_o^\perp / \pi$, can be evaluated as

$$\langle \hat{E}_F \rangle = A_F + \left(\frac{2}{3} - \frac{28}{15\pi} \right) B_F. \quad (15)$$




Figure 8. FON model single-bounce renders at $r = 1$, with a directional light and colored albedo (left), and uniform skydome and white albedo (right).

This is plotted in Figure 7, which shows there is still approximately 20% energy loss at high roughness values. Figure 8 shows the shaderball render with the FON model at $r = 1$ and the corresponding white-furnace test. This model fixes the ring artifact and also performs better on the white-furnace test, but still loses energy at non-grazing angles.

3. Energy-Preserving Oren–Nayar Model (EON)

The FON model still suffers from significant energy loss due to modeling only a low number of scattering events between microfacets. We introduce the *energy-preserving Oren–Nayar* model (or EON for short) that fixes this issue by following the energy compensation approach of Kulla and Conty [2017] for microfacet conductor models, where the missing energy due to multiple scattering is added back via a compensation term designed to be reciprocal.

In the EON model, we formulate the total BRDF as the sum of the FON model BRDF of Equation 10 and a term representing scattering orders not captured by the FON model:

$$\mathbf{f}_{\text{EON}}(\omega_i, \omega_o) = \mathbf{f}_{\text{F}}(\omega_i, \omega_o) + \mathbf{f}_{\text{F}}^{\text{ms}}(\omega_i, \omega_o). \quad (16)$$

The multiple-scattering lobe is a reciprocal function constructed to have a constant average albedo \mathcal{F} , of the form

$$\mathbf{f}_{\text{F}}^{\text{ms}}(\omega_i, \omega_o) = \frac{\mathcal{F}}{\pi} \left(\frac{1 - \hat{E}_{\text{F}}(\omega_i)}{1 - \langle \hat{E}_{\text{F}} \rangle} \right) \left(\frac{1 - \hat{E}_{\text{F}}(\omega_o)}{1 - \langle \hat{E}_{\text{F}} \rangle} \right). \quad (17)$$

The factor \mathcal{F} corresponds to the fraction of total energy reflected due to the multiple-scattering lobe, thus for $\rho = 1$ we require that $\mathcal{F} = 1 - \langle \hat{E}_{\text{F}} \rangle$ in order that the average




Figure 9. EON model single-bounce renders at $r = 1$, with a directional light and colored albedo (left) and uniform skydome and white albedo (right).

albedo of the total BRDF equals 1.

To derive a suitable approximate form for \mathcal{F} , we follow the approach described by Jakob et al. [2014] and Kulla and Conty [2017]. We assume that the FON BRDF corresponds to single scattering only. Then, considering the overall energy balance at each bounce on the microfacet surface (approximating the incident illumination in each case as uniform), at each vertex a fraction ρ of the energy is absorbed, a fraction $\langle \hat{E}_F \rangle$ escapes the surface without further interaction (as for single scattering), and a fraction $1 - \langle \hat{E}_F \rangle$ hits the surface again. Thus we can approximate \mathcal{F} as the following geometric series of terms, where the 1-bounce term is omitted since the single-scattering BRDF accounts for that:

$$\begin{aligned} \mathcal{F} &= \underbrace{\rho \langle \hat{E}_F \rangle}_{\text{1-bounce}} + \underbrace{\rho \left(1 - \langle \hat{E}_F \rangle\right) \rho \langle \hat{E}_F \rangle}_{\text{2-bounces}} + \underbrace{\rho^2 \left(1 - \langle \hat{E}_F \rangle\right)^2 \rho \langle \hat{E}_F \rangle}_{\text{3-bounces}} + \dots \\ &= \frac{\rho^2 \langle \hat{E}_F \rangle \left(1 - \langle \hat{E}_F \rangle\right)}{1 - \rho \left(1 - \langle \hat{E}_F \rangle\right)}. \end{aligned} \quad (18)$$

As $\rho \rightarrow 1$, $\mathcal{F} \rightarrow 1 - \langle \hat{E}_F \rangle$ as required. A factor of ρ^2 occurs in the numerator, since the multiple-scattering lobe corresponds only to double scattering and higher order [Hill 2018; Kulla and Conty 2017]. The final form of the multiple-scattering lobe is thus:

$$\mathbf{f}_F^{\text{ms}}(\omega_i, \omega_o) = \frac{\rho^2}{\pi} \frac{\langle \hat{E}_F \rangle}{1 - \rho \left(1 - \langle \hat{E}_F \rangle\right)} \frac{\left(1 - \hat{E}_F(\omega_i)\right) \left(1 - \hat{E}_F(\omega_o)\right)}{\left(1 - \langle \hat{E}_F \rangle\right)}. \quad (19)$$

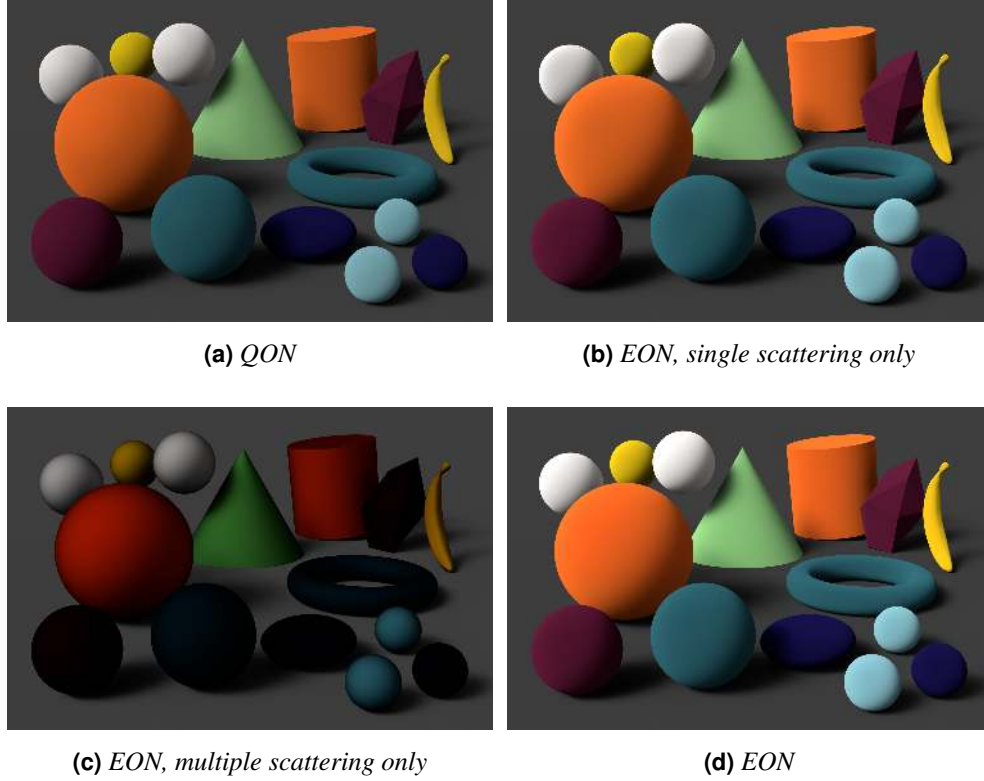


Figure 10. Comparison of QON to the single- and multiple-scattering components of EON.

The corresponding directional albedo of this multiple-scattering lobe f_F^{ms} is

$$\mathbf{E}_F^{\text{ms}}(\omega_o) = \mathcal{F} \frac{(1 - \hat{E}_F(\omega_o))}{1 - \langle \hat{E}_F \rangle}. \quad (20)$$

Figure 9 shows the shader ball render and white-furnace test with this model at $r = 1$. Since $\mathcal{F} \rightarrow 1 - \langle \hat{E}_F \rangle$ as $\rho \rightarrow 1$, the total directional albedo $E_F(\omega_o) + E_F^{\text{ms}}(\omega_o) \rightarrow 1$ also, so the EON model passes the white-furnace test. The EON model also inherits from the FON model the fix to the ring artifact of QON.

The f_F^{ms} term will produce a color shift due to the extra saturation caused by multiple scattering (as shown on objects with different albedos in Figure 10). This is physically reasonable, and in practice the extra saturation is only particularly noticeable in cases with high roughness and high, chromatic albedo. For use cases where such a color shift is not artistically desirable, the full EON model can simply be evaluated with $\rho = 1$ (i.e., white) and the result multiplied by the artist-specified color.

Listing 1 gives a self-contained GLSL implementation of the EON model, where `f_EON` returns the BRDF $\mathbf{f}_{\text{EON}}(\omega_i, \omega_o)$ value, and we also provide a function `E_EON` to compute the directional albedo of the EON BRDF.

```
const float pi = 3.14159265f;
const float rcppi = 1.0f / pi;
const float constant1_FON = 0.5f - 2.0f / (3.0f * pi);
const float constant2_FON = 2.0f / 3.0f - 28.0f / (15.0f * pi);

float E_FON_exact(float mu, float r)
{
    float AF = 1.0f / (1.0f + constant1_FON * r); // FON A coefficient
    float BF = r * AF; // FON B coefficient
    float Si = sqrt(1.0f - (mu * mu));
    float G = Si * (acos(mu) - Si * mu)
        + (2.0f / 3.0f) * ((Si / mu) * (1.0f - (Si * Si * Si)) - Si);
    return AF + (BF * rcppi) * G;
}

float E_FON_approx(float mu, float r)
{
    float mucomp = 1.0f - mu;
    const float g1 = 0.0571085289f;
    const float g2 = 0.491881867f;
    const float g3 = -0.332181442f;
    const float g4 = 0.0714429953f;
    float GoverPi = mucomp * (g1 + mucomp * (g2 + mucomp * (g3 + mucomp * g4)));
    return (1.0f + r * GoverPi) / (1.0f + constant1_FON * r);
}

// Evaluates EON BRDF value, given inputs:
// rho = single-scattering albedo parameter
// r = roughness in [0, 1]
// exact = flag to select exact or fast approx. version
// Note that this implementation assumes throughout that the directions are
// specified in a local space where the z-direction aligns with the surface normal.
vec3 f_EON(vec3 rho, float r, vec3 wi_local, vec3 wo_local, bool exact)
{
    float mu_i = wi_local.z; // Input angle cos
    float mu_o = wo_local.z; // Output angle cos
    float s = dot(wi_local, wo_local) - mu_i * mu_o; // QON s term
    float sovertF = s > 0.0f ? s / max(mu_i, mu_o) : s; // FON s/t
    float AF = 1.0f / (1.0f + constant1_FON * r); // FON A coefficient
    vec3 f_ss = (rho * rcppi) * AF * (1.0f + r * sovertF); // Single-scatter lobe
    float EFo = exact ? E_FON_exact(mu_o, r) : E_FON_approx(mu_o, r); // FON w_o albedo (exact)
    float EFi = exact ? E_FON_exact(mu_i, r) : E_FON_approx(mu_i, r); // FON w_i albedo (exact)
    float avgEF = AF * (1.0f + constant2_FON * r); // Average albedo
    vec3 rho_ms = (rho * rho) * avgEF / (vec3(1.0f) - rho * (1.0f - avgEF));
    const float eps = 1.0e-7f;
    vec3 f_ms = (rho_ms * rcppi) * max(eps, 1.0f - EFo) // Multi-scatter lobe
        * max(eps, 1.0f - EFi)
        / max(eps, 1.0f - avgEF);
    return f_ss + f_ms;
}

// Computes EON directional albedo:
vec3 E_EON(vec3 rho, float r, vec3 wi_local, bool exact)
{
    float mu_i = wi_local.z; // Input angle cos
    float AF = 1.0f / (1.0f + constant1_FON * r); // FON A coefficient
    float EF = exact ? E_FON_exact(mu_i, r) : E_FON_approx(mu_i, r); // FON w_i albedo (exact)
    float avgEF = AF * (1.0f + constant2_FON * r); // Average albedo
    vec3 rho_ms = (rho * rho) * avgEF / (vec3(1.0f) - rho * (1.0f - avgEF));
    return rho * EF + rho_ms * (1.0f - EF);
}
```

Listing 1. GLSL code for evaluation of the EON model BRDF of Section 3.

4. Importance Sampling the EON Model

The original Oren–Nayar model is typically importance-sampled using standard cosine-weighted hemispherical sampling. While this perfectly importance-samples the model in the zero-roughness (Lambertian) case, it is far from optimal in the high-roughness case near grazing angles. This is shown by the red curve in Figure 11, which is the variance of the throughput weight $w_j = (\omega_i \cdot N)^+ \mathbf{f}_{\text{EON}}/p_j(\omega_i)$ (setting $\rho = 1$) for the EON BRDF when sampled with cosine-weighted hemispherical sampling. The variance increases by a factor of over 100 at grazing angles compared to normal incidence. This is due to the cosine-weighted sampling greatly underestimating the PDF of grazing directions generated by the EON BRDF. Uniformly sampling the hemisphere produces the dotted black variance curve, which has lower variance than cosine-weighted sampling at grazing angles, but has higher variance at non-grazing angles.

Cosine-weighted hemispherical sampling amounts to sampling from the hemispherical distribution function

$$D_H(\omega_H) = \frac{(\omega_H \cdot N)^+}{\pi}. \quad (21)$$

An improved importance sampling scheme can be obtained via a technique known as *Linearly Transformed Cosines* (LTC) [Heitz et al. 2016], whereby a given direction vector ω_H sampled from this clamped cosine lobe is linearly transformed using a matrix \mathbf{M} , yielding ω_i :⁵

$$\omega_i = \frac{\mathbf{M} \omega_H}{\|\mathbf{M} \omega_H\|}. \quad (22)$$

The resulting lobe can be adjusted to best fit a cosine-weighted target BRDF by varying the matrix coefficients and can be efficiently sampled from via D_H . The PDF of the sampled direction ω_i is then given by [Heitz et al. 2016]

$$D(\omega_i) = D_H \left(\frac{\mathbf{M}^{-1} \omega_i}{\|\mathbf{M}^{-1} \omega_i\|} \right) \frac{\partial \omega_H}{\partial \omega_i}, \quad (23)$$

where the factor $\frac{\partial \omega_H}{\partial \omega_i}$ is the Jacobian of the transformation from ω_H to ω_i :

$$\frac{\partial \omega_H}{\partial \omega_i} = \frac{|\mathbf{M}^{-1}|}{\|\mathbf{M}^{-1} \omega_i\|^3}. \quad (24)$$

⁵Note that we assume the usual convention for unidirectional path tracing, where the direction of the outgoing ray ω_o is known, and the incident ray direction ω_i is the one being sampled. Both the incident ray direction ω_i and the outgoing ray direction ω_o are oriented to point away from the surface, i.e., the incident ray is in the opposite direction to incident photons, while the outgoing ray is parallel to outgoing photons. Also note that the sampling would work equally well for the reverse case, when tracing paths in the direction of light flow.




Figure 11. Statistics of the throughput weight w_j over 10^6 samples of the EON BRDF as a function of angle, using the various techniques: variance of the weight (left), the maximum weight (center), and the mean weight (right).

Rotating azimuthally into a space where the ω_o direction (of the outgoing ray) has the form $\omega_o = (\sin \theta_o, 0, \cos \theta_o)$, we take the LTC transformation matrix to have the most general form for an isotropic BRDF [Heitz et al. 2016]:

$$\mathbf{M} = \begin{bmatrix} a & 0 & b \\ 0 & c & 0 \\ d & 0 & 1 \end{bmatrix}, \quad (25)$$

where the a, c coefficients stretch the lobe radially in the x - y plane, while the b, d coefficients shear the lobe in the x - z plane. The coefficients a, b, c, d are functions of the outgoing angle θ_o and roughness r . These functions can be adjusted via a fitting process to make the LTC lobe best match the shape of the cosine-weighted EON BRDF, to better importance-sample the backscattering peak.

For our fitting we employed Mathematica’s `NonlinearModelFit`⁶ in two stages: first generating a table of best-fit LTC matrices across $\mu = \cos \theta_o$ and r , and then fitting polynomial or rational functions to this data for each of the coefficients. The goal of this second step was to offer a self-contained and portable form for practitioners by avoiding the need for a table lookup. The end result is provided by the GLSL function `ltc_coeffs` in Listing 2.

One possible simplification is to set the coefficient $d = 0$, which guarantees that the sampled directions of the LTC lobe lie in the positive hemisphere [Zeltner et al.

⁶With the default Automatic method and error function.

```
void ltc_coeffs(float mu, float r,
               out float a, out float b, out float c, out float d)
{
    a = 1.0f + r*(0.303392f + (-0.518982f + 0.111709f*mu)*mu + (-0.276266f + 0.335918f*mu)*r);
    b = r*(-1.16407f + 1.15859f*mu + (0.150815f - 0.150105f*mu)*r)/(mu*mu*mu - 1.43545f);
    c = 1.0f + r*(0.20013f + (-0.506373f + 0.261777f*mu)*mu);
    d = r*(0.540852f + (-1.01625f + 0.475392f*mu)*mu)/(-1.0743f + (0.0725628f + mu)*mu);
}
```

Listing 2. LTC lobe coefficients a, b, c, d as a function of $\mu = \cos \theta_o$ and r , for importance sampling of the EON model.

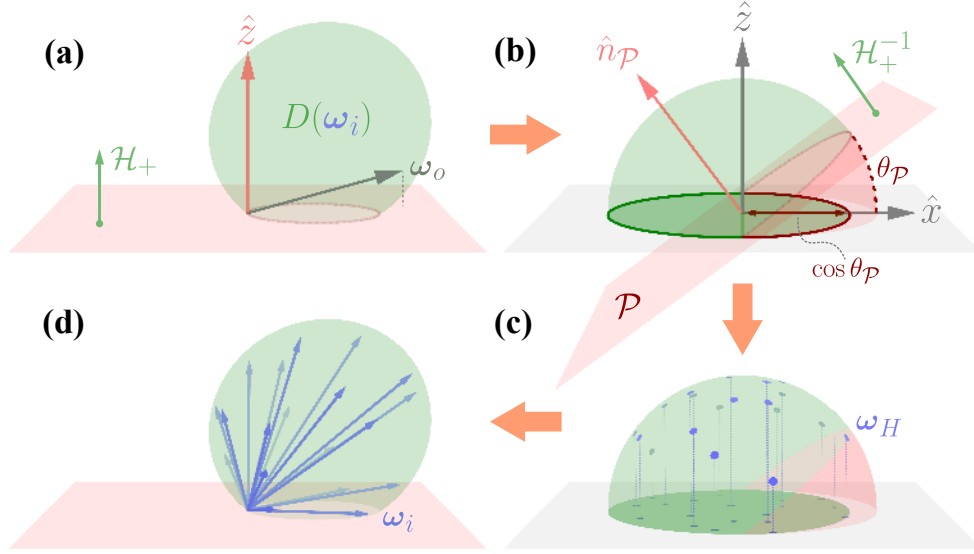


Figure 12. Geometry of CLTC importance sampling: **(a)** The LTC lobe (directions scaled by $D(\omega_i)$) for a near grazing angle ω_o and $r = 1$. We seek to restrict sample directions to the positive half-space of the x - y plane: $\omega_i \in \mathcal{H}_+$. **(b)** In the untransformed space of D_H , these directions correspond to $\omega_H \in \mathcal{H}_+^{-1}$, the positive half-space of plane \mathcal{P} with normal \hat{n}_P . It follows that we should perform cosine-weighted sampling of the upper hemisphere clipped by \mathcal{P} (shown in green). We achieve this by uniformly sampling the 2D projection of this subregion on the x - y plane and then **(c)** reprojecting back onto the upper hemisphere. **(d)** Finally, we proceed as usual by applying the linear transform to obtain the sampled output directions: $\omega_i = \mathbf{M} \omega_H$.

2022]. This produces the “LTC (non-negative)” variance curve in dotted blue in Figure 11. Unfortunately, while this is an improvement on cosine-weighted and uniform sampling, it still suffers from high variance at grazing angles, since the PDF is still very low in this region.

To significantly reduce variance, the more general LTC matrix with a nonzero d coefficient is required, which produces an LTC lobe that extends *below* the hemisphere. The variance of this is shown in the “LTC (general fit)” curve in solid blue in Figure 11. This greatly reduces the variance peak at grazing angles, but generates a rather spiky variance curve. Also, on sampling from this lobe, directions that are generated below the hemisphere are effectively wasted as they have a throughput weight of zero. This is not acceptable for renders at low sample count.

To prevent the generation of samples below the hemisphere, we modify the LTC sampling to appropriately restrict the sampling domain of the clamped cosine lobe. We refer to this as *Clipped Linearly Transformed Cosine* (CLTC) sampling. Figure 12 gives an overview of the method. This produces the CLTC green dashed variance curves in Figure 11. The idea is that directions $\omega_i \in \mathcal{H}_+$ that lie in the positive half-space of the x - y plane (Figure 12(a)) correspond to untransformed directions

$\omega_H \in \mathcal{H}_+^{-1}$ that are in the positive half-space of a plane \mathcal{P} (Figure 12(b)). The (unnormalized) normal to \mathcal{P} is⁷

$$\hat{n}_{\mathcal{P}} = \mathbf{M}^{-1} \hat{x} \times \mathbf{M}^{-1} \hat{y} = \det(\mathbf{M}^{-1}) \mathbf{M}^T \hat{z}. \quad (26)$$

In the case of our fit, $\det(\mathbf{M}^{-1}) > 0$ over all θ_o and r , so given Equation 26 and the definition of \mathbf{M} from Equation 25, $\hat{n}_{\mathcal{P}} \propto \mathbf{M}^T \hat{z} = (d, 0, 1)$. As proof that directions ω_i remain in the positive half-space under transformation by \mathbf{M}^{-1} , it is sufficient to show that $\mathbf{M}^{-1} \hat{z} \cdot \hat{n}_{\mathcal{P}} > 0$. This is the case since $\mathbf{M}^{-1} \hat{z} \cdot \hat{n}_{\mathcal{P}} = \det(\mathbf{M}^{-1})$ and, as previously stated, $\det(\mathbf{M}^{-1}) > 0$.

It follows that the ω_H region we must sample from is given by the positive hemisphere *clipped* to the positive half-space of \mathcal{P} . In other words, the full hemisphere with a *spherical lune* omitted, as shown in Figure 12(b). The opening angle of this removed lune is angle $\theta_{\mathcal{P}}$ between the plane \mathcal{P} and the x - y plane, given according to the previous description by $\tan \theta_{\mathcal{P}} = d$.

To sample from the remaining hemispherical subregion with the cosine-weighted PDF $D_H(\omega_H)$ applied, we use the Nusselt analog, in which the correct PDF is obtained by sampling uniformly from the projection of this subregion onto the x - y plane, then reprojecting the sampled 2D point back onto the upper hemisphere (Figure 12(c)).

Recall that the solid angle element $d\omega_H$ projects to an area element on the x - y plane given by $dA = (\omega_H \cdot N)^+ d\omega_H$. The area-measure PDF $p(\mathbf{x})$ of the projected points \mathbf{x} on the x - y plane must satisfy $p(\mathbf{x})dA = p(\omega_H)d\omega_H$, so $p(\mathbf{x}) = p(\omega_H)/(\omega_H \cdot N)^+$. Thus, since $p(\omega_H) \propto D_H(\omega_H)$, the area-measure PDF is constant and equal to the inverse area of the projected region.

As shown in Figure 12(b), the projected shape to sample from takes the form of a half-circle (outlined in green) and a half-ellipse (outlined in red) with semi-major axis 1 and semi-minor axis $\cos \theta_{\mathcal{P}} = 1/\sqrt{1+d^2}$. The total area of the projected shape is therefore $A = \frac{\pi}{2} (1 + \cos \theta_{\mathcal{P}})$. On reprojection, the sampling PDF for the direction ω_H is thus given by

$$p(\omega_H) = \frac{(\omega_H \cdot N)^+}{A} = \frac{2 D_H(\omega_H)}{1 + 1/\sqrt{1+d^2}}. \quad (27)$$

(Note that as $r \rightarrow 0$, $d \rightarrow 0$, so the PDF reduces to that of cosine-weighted sampling.) The final sampled direction ω_i (Figure 12(d)) is obtained by applying the linear transform of Equation 22, and $p(\omega_i) = D(\omega_i)$ is computed according to Equation 23.

A simple algorithm for uniform-sampling this union of a half-circle and half-ellipse is given by Heitz [2018], which we provide as GLSL code in Listing 3.

⁷Using the identity given in Equation 17 of the original LTC paper of Heitz et al. [2016].

```

mat3 orthonormal_basis_ltc(vec3 w)
{
    float lenSqr = dot(w.xy, w.xy);
    vec3 X = lenSqr > 0.0f ? vec3(w.x, w.y, 0.0f) * inversesqrt(lenSqr) : vec3(1,0,0);
    vec3 Y = vec3(-X.y, X.x, 0.0f); // cross(Z, X)
    return mat3(X, Y, vec3(0, 0, 1));
}

vec4 cltc_sample(vec3 wo_local, float r, float u1, float u2)
{
    float a, b, c, d; ltc_coeffs(wo_local.z, r, a, b, c, d); // Coeffs. of LTC M
    float R = sqrt(u1); float phi = 2.0f * pi * u2; // CLTC sampling
    float x = R * cos(phi); float y = R * sin(phi); // CLTC sampling
    float vz = 1.0f / sqrt(d*d + 1.0f); // CLTC sampling factors
    float s = 0.5f * (1.0f + vz); // CLTC sampling factors
    x = -mix(sqrt(1.0f - y*y), x, s); // CLTC sampling
    vec3 wh = vec3(x, y, sqrt(max(1.0f - (x*x + y*y), 0.0f))); //  $\omega_H$  sample via CLTC
    float pdf_wh = wh.z / (pi * s); // PDF of  $\omega_H$  sample
    vec3 wi = vec3(a*wh.x + b*wh.z, c*wh.y, d*wh.x + wh.z); //  $M \omega_H$  (unnormalized)
    float len = length(wi); //  $\|M \omega_H\| = 1/\|M^{-1} \omega_H\|$ 
    float detM = c*(a - b*d); // |M|
    float pdf_wi = pdf_wh * len*len*len / detM; //  $\omega_i$  sample PDF
    mat3 fromLTC = orthonormal_basis_ltc(wo_local); //  $\omega_i \rightarrow$  local space
    wi = normalize(fromLTC * wi); //  $\omega_i \rightarrow$  local space
    return vec4(wi, pdf_wi);
}

float cltc_pdf(vec3 wo_local, vec3 wi_local, float r)
{
    mat3 toLTC = transpose(orthonormal_basis_ltc(wo_local)); //  $\omega_i \rightarrow$  LTC space
    vec3 wi = toLTC * wi_local; //  $\omega_i \rightarrow$  LTC space
    float a, b, c, d; ltc_coeffs(wo_local.z, r, a, b, c, d); // Coeffs. of LTC M
    float detM = c*(a - b*d); // |M|
    vec3 wh = vec3(c*(wi.x - b*wi.z), (a - b*d)*wi.y, -c*(d*wi.x - a*wi.z)); //  $\text{adj}(M) \omega_i$ 
    float lenSqr = dot(wh, wh); //  $\|M\| \|M^{-1} \omega_i\|$ 
    float vz = 1.0f / sqrt(d*d + 1.0f); // CLTC sampling factors
    float s = 0.5f * (1.0f + vz); // CLTC sampling factors
    float pdf = detM*detM/(lenSqr*lenSqr) * max(wh.z, 0.0f) / (pi * s); //  $\omega_i$  sample PDF
    return pdf;
}

```

Listing 3. GLSL code for CLTC importance sampling.

The green dashed CLTC curves of variance and maximum throughput weight in Figure 11 show that this method generates high outliers compared to uniform sampling, which will correspond to fireflies in the render. There is also a slight bias in the estimator as exhibited in the deviation of the mean throughput weight from 1, since there are directions where the LTC lobe is zero but the EON lobe is nonzero.

To prevent the variance spikes and bias, we finally combine the CLTC lobe sampling with a uniform hemispherical lobe via *one-sample* multiple importance sampling (MIS), producing the “CLTC + uniform (MIS)” variance curve in solid green in Figure 11. With MIS, the maximum throughput weight (center) is smoothed and reduced, and the throughput weight itself (right) is debiased.

The probability of selecting the uniform lobe as a function of θ_o and roughness was found via optimization. In practice, we minimized the maximum throughput weight across $\mu = \cos \theta_o$ for $r = 1$, and then fit a polynomial function to the result.⁸ Finally, we augmented this with an additional $r^{0.1}$ term to smoothly transition to the

⁸Again using Mathematica’s NonlinearModelFit.

```
vec3 uniform_lobe_sample(float u1, float u2)
{
    float sinTheta = sqrt(1.0f - u1*u1); float phi = 2.0f * pi * u2;
    return vec3(sinTheta * cos(phi), sinTheta * sin(phi), u1);
}

// Samples (via CLTC) from EON BRDF, given inputs:
// rho = single-scattering albedo parameter
// wo_local = direction of outgoing ray (directed away from vertex)
// r = roughness in [0, 1]
// u1, u2 = IID uniform random numbers in [0,1]
// Returns vec4(vec3(wi_local), pdf)
vec4 sample_EON(vec3 wo_local, float r, float u1, float u2)
{
    float mu = wo_local.z;
    float P_u = pow(r, 0.1f) * (0.162925f + (-0.372058f + (0.538233f - 0.290822f*mu)*mu)*mu);
    float P_c = 1.0f - P_u; // Probability of CLTC sample
    vec4 wi; float pdf_c;
    if (u1 <= P_u) {
        u1 = u1 / P_u;
        wi = uniform_lobe_sample(u1, u2); // Sample wi from uniform lobe
        pdf_c = cltc_pdf(wo_local, wi.xyz, r); // Evaluate CLTC PDF at wi
    } else {
        u1 = (u1 - P_u) / P_c;
        wi = cltc_sample(wo_local, r, u1, u2); // Sample wi from CLTC lobe
        pdf_c = wi.w;
    }
    const float pdf_u = 1.0f / (2.0f * pi);
    wi.w = P_u*pdf_u + P_c*pdf_c; // MIS PDF of wi
    return wi;
}

// PDF corresponding to the above sampling routine
float pdf_EON(vec3 wo_local, vec3 wi_local, float r)
{
    float mu = wo_local.z;
    float P_u = pow(r, 0.1f) * (0.162925f + (-0.372058f + (0.538233f - 0.290822f*mu)*mu)*mu);
    float P_c = 1.0f - P_u;
    float pdf_c = cltc_pdf(wo_local, wi_local, r);
    const float pdf_u = 1.0f / (2.0f * pi);
    return P_u*pdf_u + P_c*pdf_c;
}
```

Listing 4. GLSL code for importance sampling of the EON model, as in Section 4.

Lambertian case as $r \rightarrow 0$. Listing 4 gives the complete GLSL implementation.

Figure 13 shows how the variance and maximum throughput weight of the algorithm of Listing 4 vary as a function of the roughness parameter r , compared to regular cosine-weighted sampling (note that the throughput weight is 1 at $r = 0$ for both cosine-weighted and CLTC sampling).

As $r \rightarrow 0$, both methods produce zero variance and throughput weight (since our algorithm reduces to cosine-weighted sampling in the zero roughness limit). At nonzero roughness, our method produces both variance and maximum throughput whose peaks are much lower than the grazing peaks for cosine-weighted sampling.

Figure 14 shows the variance reduction of CLTC sampling versus cosine sampling in renders at high bounce count and low sample-per-pixel count.

Table 1 shows timings for the various BRDFs discussed, as nanoseconds per BRDF evaluation or sampling invocation (where the sampling time includes the time to evaluate the BRDF in the sampled direction also). The profiles were computed by porting the GLSL implementations of the BRDFs to C++ using the GLM library [G-




Figure 13. The effect of the roughness r on the variance (left) and maximum throughput weight (right) (of 10^6 samples) of our CLTC (plus MIS with uniform) sampling method, compared to traditional cosine-weighted sampling.

Truc Creation 2024] and compiling with maximum optimizations enabled. Timings were performed on a single thread of an AMD 3970X, averaged over several billion samples (using random roughness values, albedo values, and input directions).

Unsurprisingly, the Lambert model BRDF is unbeatably fast to evaluate, while the fullON method is slowest. The exact EON BRDF is about $1.3\times$ faster to evaluate than fullON, and approximate EON is about $7\times$ faster.

The QON and FON variants are faster to evaluate than EON, but, as discussed, they suffer from energy loss and artifacts that render them impractical. The CLTC sampling of the approximate EON model takes roughly twice as long as cosine sampling of the same model (and roughly the same time as cosine sampling of the fullON model).

However, as discussed above (see Figure 13), the variance reduction of the CLTC sampling method exceeds a factor of 100 at grazing angles, thus the time to reach a given noise level will be greatly reduced.

BRDF	Eval (ns)	Cosine Sample (ns)	CLTC Sample (ns)
Lambert	1.4	67.2	—
FON	6.3	67.2	—
QON	6.3	67.4	—
fullON	141.7	211.218	—
EON (approx)	20.8	105.8	217.0
EON (exact)	108.7	193.5	297.0

Table 1. BRDF evaluation and sampling timings.




Figure 14. EON model at maximal roughness ($r = 1$), rendered with traditional cosine-weighted sampling (left) and our CLTC sampling (right), with 50 samples per pixel and 50 bounces. The variance reduction of CLTC sampling produces a significantly smaller amount of noise and fireflies than cosine-weighted sampling.

5. Conclusion

In this paper we have introduced the *energy-preserving Oren–Nayar* model (or EON), a new variant of the classic model of [Oren and Nayar \[1994\]](#), that addresses many of the practical limitations of the original. We consider the EON model to be more practical than other existing variants of the Oren–Nayar model because it is energy preserving, reciprocal, artifact-free, easy to implement, and efficient to evaluate, making it a strong candidate for use in production rendering applications.

We also provided a practical importance sampling scheme for this model, based on a Linearly Transformed Cosine (LTC) lobe. We introduced *Clipped Linearly Transformed Cosine* (CLTC) sampling to prevent negative hemisphere sampling, and combined this with uniform hemispherical sampling via MIS to achieve a much lower variance and maximum throughput weight than standard cosine-weighted sampling (as illustrated in Figure 14).

For these reasons it has been chosen as the rough diffuse model used in the proposed industry standard OpenPBR über-shader specification [[Andersson et al. 2024](#)].

The full source code presented in the listings can be downloaded from GitHub [[Portsmouth et al. 2025](#)].

Future Work. The overall shape of our new BRDF is influenced by the underlying assumptions of the Oren–Nayar model. It would be interesting to compare to alternative physically based models, especially the recent Lambertian sphere [d’Eon 2021] and VMF diffuse [d’Eon and Weidlich 2024] models, which are based on the scattering from a collection of spherical Lambertian particles. We note here though that, as far as we are aware, those models are not explicitly energy preserving, are considerably more computationally intensive to evaluate than EON, and do not come with an importance sampling scheme more sophisticated than cosine sampling. Indeed, the CLTC sampling scheme developed here could possibly be repurposed to provide improved importance sampling for these models.

Future refinements to the EON model itself could include better aligning our new model with an underlying microfacet model. This could allow more physically accurate color shifting by separating the single, double, and higher-order scattering components, with a more physically accurate shape for the multiple-scattering lobes.

While this paper focuses on the Oren–Nayar family of diffuse BRDFs, the broader question of how best to model the appearances of the wide variety of real-world diffuse materials remains open, requiring a balance between fidelity and practical considerations such as usability and computational cost. We encourage additional research to refine these tradeoffs and further advance physically based diffuse reflectance modeling.

Acknowledgements

We thank Chris Kulla, Eugene d’Eon, Luke Emrose, and Ivo Kondapaneni for helpful discussions. In several figures we used the Standard Shader Ball asset of Mazzone and Rydalch [2023]. The dragon of Figure 1 is the Glavenus model by Axolote Gaming [2022], from myminifactory.com. The model in Figure 4 is the “Mat” model from Adobe [Marchive 2022].

References

- AKENINE-MÖLLER, T., HAINES, E., HOFFMAN, N., PESCE, A., IWANICKI, M., AND HILLAIRES, S. *Real-Time Rendering*. A K Peters/CRC Press, fourth edition, 2018. URL: <https://www.routledge.com/Real-Time-Rendering-Fourth-Edition/Akenine-Moller-Haines-Hoffman/p/book/9781138627000>. 121
- ANDERSSON, Z., EDMONDSON, P., GUERTAULT, J., HERUBEL, A., KING, A., KUTZ, P., MACHIZAUD, A., PORTSMOUTH, J., SERVANT, F., AND STONE, J. OpenPBR Surface Specification v1.1. Technical report, Academy Software Foundation (ASWF), 2024. URL: <https://academysoftwarefoundation.github.io/OpenPBR/>. 136
- AXOLOTE GAMING. Glavenus RPG scale fan-art. MyMiniFactory, 2022. URL: <https://www.myminifactory.com/object/3d-print-glavenus-rpg-scale-fan-art-224500>. 137

- D'EON, E. An analytic BRDF for materials with spherical Lambertian scatterers. In *EGSR21: 32nd Eurographics Symposium on Rendering*, pages 153–161. Eurographics Association, 2021. URL: <https://doi.org/10.1111/cgf.14348>. 137
- D'EON, E. AND WEIDLICH, A. VMF diffuse: A unified rough diffuse BRDF. *Computer Graphics Forum*, 43(4), 2024. URL: <https://doi.org/10.1111/cgf.15149>. 137
- FUJII, Y. A tiny improvement of Oren–Nayar reflectance model. Mimosa Pudica, 2013. URL: <https://mimosa-pudica.net/improved-oren-nayar.html>. 118, 119, 121, 123
- G-TRUC CREATION. OpenGL Mathematics (GLM). GitHub repository, 2024. URL: <https://github.com/g-truc/glm>. 134
- GEORGIEV, I., IZE, T., FARNSWORTH, M., MONTOYA-VOZMEDIANO, R., KING, A., LOMMEL, B. V., JIMENEZ, A., ANSON, O., OGAKI, S., JOHNSTON, E., HERUBEL, A., RUSSELL, D., SERVANT, F., AND FAJARDO, M. Arnold: A brute-force production path tracer. *ACM Transactions on Graphics*, 37(3):32:1–32:12, August 2018. URL: <https://doi.org/10.1145/3182160>. 118
- GRITZ, L., STEIN, C., KULLA, C., AND CONTY, A. Open Shading Language. In *ACM SIGGRAPH 2010 Talks, SIGGRAPH '10*, page 33:1, New York, 2010. Association for Computing Machinery. URL: <https://doi.org/10.1145/1837026.1837070>. 118
- HEITZ, E. Sampling the GGX distribution of visible normals. *Journal of Computer Graphics Techniques (JCGT)*, 7(4):1–13, November 2018. URL: <http://jcgt.org/published/0007/04/01/>. 132
- HEITZ, E., DUPUY, J., HILL, S., AND NEUBELT, D. Real-time polygonal-light shading with linearly transformed cosines. *ACM Transactions on Graphics (SIGGRAPH '16)*, 35(4):41:1–41:8, 2016. URL: <https://doi.org/10.1145/2897824.2925895>. 129, 130, 132
- HILL, S. A multi-faceted exploration (Part 2). Self Shadow, June 4, 2018. URL: <https://blog.selfshadow.com/2018/06/04/multi-faceted-part-2/>. 126
- JAKOB, W., D'EON, E., JAKOB, O., AND MARSCHNER, S. A comprehensive framework for rendering layered materials. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 33(4):118:1–118:14, July 2014. URL: <https://doi.org/10.1145/2601097.2601139>. 126
- KETTNER, L., RAAB, M., SEIBERT, D., JORDAN, J., AND KELLER, A. The Material Definition Language. In *Proceedings of the Third Workshop on Material Appearance Modeling: Issues and Acquisition, MAM '15*, page 1–4, Goslar, DEU, 2015. Eurographics Association. URL: <https://doi.org/10.2312/mam.20151195>. 118
- KULLA, C. AND CONTY, A. Revisiting physically based shading at Imageworks. In *ACM SIGGRAPH 2017 Courses, Physically Based Shading in Theory and Practice*, pages 7:1–7:8. Association for Computing Machinery, 2017. URL: <https://doi.org/10.1145/3084873.3084893>. 125, 126
- MARCHIVE, M. MAT mesh. Adobe Assets, 2022. URL: <https://substance3d.adobe.com/community-assets/assets/b06197ce339fe70d12998919945e39e964141f29>. 137

- MAZZONE, A. AND RYDALCH, C. Standard shader ball: A modern and feature-rich render test scene. In *SIGGRAPH Asia 2023 Technical Communications*, pages 14:1–14:3. Association for Computing Machinery, 2023. URL: <https://doi.org/10.1145/3610543.3626181>. 137
- OREN, M. AND NAYAR, S. K. Generalization of Lambert’s reflectance model. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, pages 239–246. Association for Computing Machinery, 1994. URL: <https://doi.org/10.1145/192161.192213>. 117, 118, 119, 120, 121, 136
- PORTSMOUTH, J., KUTZ, P., AND HILL, S. EON-diffuse GitHub repository, 2025. URL: <https://github.com/portsmouth/EON-diffuse>. 136
- SMYTHE, D. AND STONE, J. MaterialX: An open standard for network-based CG object looks. Technical Report Specification v1.36, Academy Software Foundation, 2018. URL: <https://materialx.org/assets/MaterialX.v1.36.Spec.pdf>. 118
- ZELTNER, T., BURLEY, B., AND CHIANG, M. J.-Y. Practical multiple-scattering sheen using linearly transformed cosines. In *ACM SIGGRAPH 2022 Talks*, pages 7:1–7:2. Association for Computing Machinery, 2022. URL: <https://doi.org/10.1145/3532836.3536240>. 130

Author Contact Information

Jamie Portsmouth
jamports@mac.com

Peter Kutz
peter.kutz@gmail.com

Stephen Hill
steve@selfshadow.com

Jamie Portsmouth, Peter Kutz, and Stephen Hill, EON: A Practical Energy-Preserving Rough Diffuse BRDF, *Journal of Computer Graphics Techniques (JCGT)*, vol. 14, no. 1, 116–139, 2025
<http://jcgt.org/published/0014/01/06/>

Received: 2024-07-12

Recommended: 2024-11-21

Published: 2025-03-31

Corresponding Editor: Christophe Hery

Editor-in-Chief: Eric Haines

© 2025 Jamie Portsmouth, Peter Kutz, and Stephen Hill (the Authors).

The Authors provide this document (the Work) under the Creative Commons CC BY-ND 3.0 license available online at <http://creativecommons.org/licenses/by-nd/3.0/>. The Authors further grant permission for reuse of images and text from the first page of the Work, provided that the reuse is for the purpose of promoting and/or summarizing the Work in scholarly venues and that any reuse is accompanied by a scientific citation to the Work.

