Practical multiple scattering compensation for microfacet models

Emmanuel Turquin*
Industrial Light & Magic

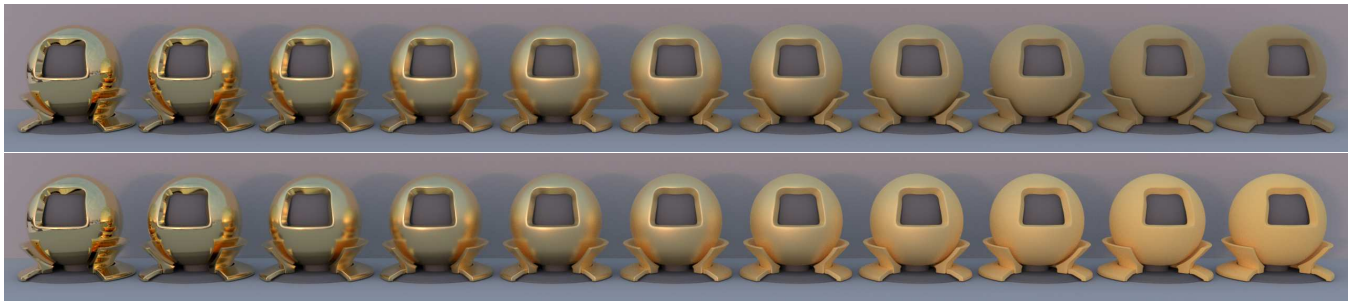

Figure 1: *Increasingly rough GGX conductor widgets.* ***Top:*** *single-scattering only.* ***Bottom:*** *with our multiple scattering compensation. Notice the boosted intensity and saturation as the surface becomes rougher, which offers a visually more satisfying and uniform gradation.*

Abstract

With the adoption by the game, animation and VFX industries of more physically grounded light transport methods and material representations, an effort has been put on *energy conservation*, to ensure both numerical stability of the computations and realistic looking images. However, this effort has been much more focused on avoiding undue light emission than ensuring no energy is lost in the process. Microfacet models have become a standard building block of surface materials for representing specular components of varying roughness; and yet, while they possess many desirable properties in addition to producing convincing results, their very design neglects an important source of scattering, that can cause a significant loss of energy. Specifically, they only model single-scattering across the microfacets, and ignore the subsequent interactions, which get more important as roughness increases. From a user standpoint, this results in an unexpected darkening of rough specular lobes, which typically has to be accounted for in *ad hoc* ways. In this document, we present and compare different approaches aiming to address this flaw and ensure *energy preservation*, including one that has been developed at ILM.

Keywords: microfacets, multiple scattering, energy compensation

1 Introduction

Since its introduction by Cook and Torrance [1982], the microfacet model has become an ubiquitous way of repesenting rough specular reflections and refractions [Walter et al. 2007], both in real-time and offline rendering. If we consider mirror-like reflecting facets, a microfacet BRDF ρ can be expressed using the well known formula:

$$\rho\left(\omega_o, \omega_i\right) \doteq \frac{F\left(\omega_o, h\right) G\left(\omega_o, \omega_i, h\right) D\left(h\right)}{4|\omega_i \cdot n||\omega_o \cdot n|} \tag{1}$$

with n being the geometric normal, h the microfacet normal, ω_i and ω_o the incoming and outgoing directions, F the Fresnel term, G the shadowing-masking function and D the microfacet distribution function. This formulation, amongst other advantages, offers a great modularity as all three terms F, G and D can be chosen between candidates with different characteristics.

*e-mail: eturquin@ilm.com

A particularly popular choice is the use of the GGX distribution as D (and its associated term G; see [Heitz 2014] on how they should relate), along with a conductor Fresnel term, to reproduce rough metals (as in Figure 1, top row). Alternatively, adding a BTDF combined with a dieletric F can replicate refractive materials such as frosted glass.

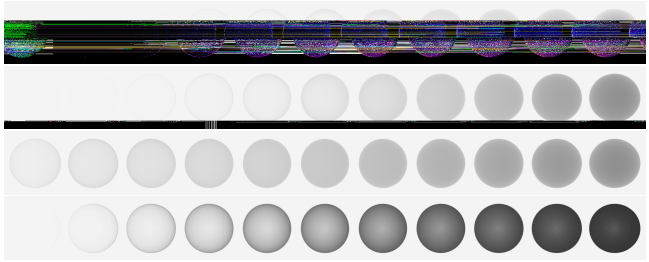

Figure 2: *White furnace test on a microfacet BRDF with $F = 1$, and roughness going from 0 (left) to 1 (right). From top to bottom: Beckmann, GGX, $GTR_{\gamma=1}$ and $STD_{\gamma=1.55}$ distributions. If multiple scattering was properly accounted for, the albedo would be 1.*

Unfortunately, this model also suffers from an important limitation: it only simulates a single interaction of the incoming ray of light on the microfacets, after which the scattered ray is either visible or not from ω_i/ω_o, as determined by G. The multiple scattering that would occur when the ray is not visible (*i.e.*, hits another microfacet, once or more) is not accounted for (see Section 3 of [Heitz 2014] for more details). As one would expect, there is no missing energy when the roughness is zero, as this corresponds to a perfectly flat surface with no possible occlusion. But the rougher the surface (or the wider the spread of microfacet orientations is), the more this lack of multiple scattering energy becomes noticeable. Looking back at Figure 1, top row, the right-most objects appear overly dark and dull, in comparison to the shinier ones. As a matter of fact, in the case of GGX and when ignoring Fresnel absorption, the loss of energy reaches about 60% for a roughness value $\alpha = 1$. It gets even worst when using distributions with longer tails, such as GTR [Burley 2012] or STD [Ribardière et al. 2017], where it can reach in excess of 90%, as illustrated by Figure 2. In a production environment, this calls for an eye-balled, manual compensation of the albedo, applied typically at the look development stage, and potentially breaking *energy conservation* in other ways. Thankfully, easier and sounder methods have been devised.

2 Previous work

The various ways of reinjecting the missing multiple scattering contribution range from heuristic and artistically controlled, to fully physically based and automated. They also vastly differ in terms of complexity, practicality and speed.

Burley [2015] introduces a simple, physically justified, user-driven *Sheen* component using the Schlick Fresnel profile $(1 - \cos\theta_o)^5$, that adds grazing forward reflection to better match observed materials, and complements the diffuse retro-reflection already proposed in [Burley 2012]. The motivation for this term is to approximately compensate for the assumed missing "multiple-scattering effects between and through microsurface features".

Figure 3: *Effect of Sheen, ranging from 0 to 1, as in [Burley 2015].*

Unfortunately, this qualitative compensation (illustrated in Figure 3) fails to take into account important factors, such as the effect of roughness variation, and thus needs to be readjusted whenever other attributes of the BSDF are edited. In Section 5.1 of those same course notes, an emphasis is put on how much of a practical issue it still is in production, while a more exact and automatic method is mentioned as future work.

In the *Multiple Scattering* sub-section of Section 7 of [Heitz 2014], a suggestion is made to "combine the knowledge of *energy conservation* and empirical observations", and to consider a new BRDF model that would be expressed as:

$$\rho\left(\omega_o, \omega_i\right) \doteq \rho_{ss}\left(\omega_o, \omega_i\right) + \rho_{ms}\left(\omega_o, \omega_i\right) \tag{2}$$

where ρ_{ss} is the usual, single-scattering only term of Equation 1, and ρ_{ms} a new multiple scattering term to be defined, that accounts for all secondary lobes. A first important constraint on ρ_{ms} is naturally imposed by *energy preservation*. Let us introduce the *albedo* as:

$$E\left(\omega_o\right) \doteq \int_{\Omega_i} \rho\left(\omega_o, \omega_i\right) |\omega_i \cdot n| d\omega_i \tag{3}$$

(similarly, E_{ss} and E_{ms}). Then, if we momentarily ignore Fresnel (*i.e.* we consider that $F = 1$), the following equality must hold:

$$E\left(\omega_o\right) = E_{ss}\left(\omega_o\right) + E_{ms}\left(\omega_o\right) = 1 \tag{4}$$

expressing the fact that without any absorption, all incoming energy is reflected back, after one or more bounce(s) on the microfacets. Put differently:

$$E_{ms}\left(\omega_o\right) = 1 - E_{ss}\left(\omega_o\right) \tag{5}$$

Given that E_{ss} is completely defined, we have different options to (pre-)compute it, or E_{ms}. Respecting Equation 5 prevents energy loss, and would naturally lead to a perfectly constant Figure 2, as desired. Now, what remains to be explored is:

- The shape of ρ_{ms}
- How to incorporate Fresnel absorption.

While Heitz [2014] doesn't offer a direct solution to either problem, he notes that:

> "The shape of ρ_{ms} could be investigated, for instance, by computing Monte-Carlo simulations on rough surface samples." (Q1)

adding right after:

> "If it turns out to be simple, then as a first approximation we could model it with an analytical function (*e.g.* as a single lobe)." (Q2)

About Fresnel, he suggests to precompute an average term F_{ss} for one bounce and rescale E_{ms} by it. If deemed not accurate enough, "perhaps the average value F_{ms} after multiple bounces could be precomputed as well". Concluding with this promising observation:

> "Since multiple scattering tends to smooth out functions, one can reasonably expect it to be efficiently represented and stored with simple analytical functions or small precomputed look-up textures." (Q3)

In a subsequent effort, Heitz et al. [2016] follow up on the idea expressed in Q1, and propose a full-on stochastic model to simulate multiple scattering, using the microflake theory for volumes and random walks on the microsurface. The model is shown to be extremely accurate, not only for conductor reflections but also dielectric scattering, and would be ideal if not for its involved nature that makes it not so straightforward to plug into an existing production rendering system (notably with the use of additional random numbers), and more importantly computationally very expensive.

An interesting and somewhat surprising take from this work (visible in Figure 15 of their paper, and partially reproduced here in Figure 4 for convenience), is that secondary lobes do not look diffuse but rather like scaled-down versions of the primary one.

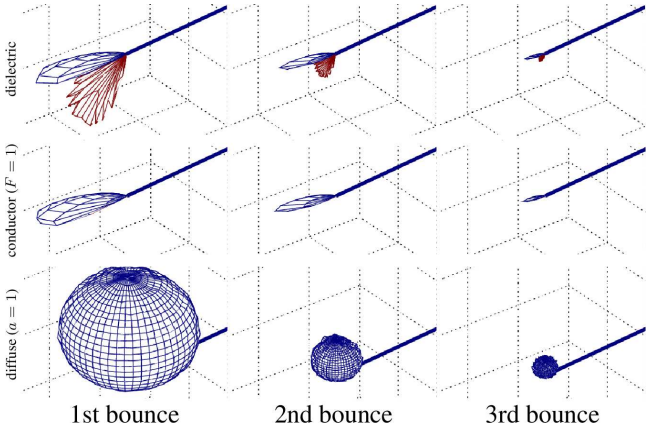

Figure 4: *Resulting lobes from a simulated microsurface using the Beckmann distribution, courtesy of Heitz et al. [2016].*

In an effort concurrent to ours, Kulla and Conty [2017] exploit observations made in Q2 and Q3 and propose a significantly less accurate, but much simpler and faster solution. They adapt the work of Kelemen and Szirmay-Kalos [2001] to the present problem, and use their diffuse-looking matte component as multiple scattering lobe, which by construction respects Equation 5:

$$\rho_{ms}\left(\omega_o, \omega_i\right) \doteq F_{ms} \frac{(1 - E_{ss}(\omega_o))(1 - E_{ss}(\omega_i))}{\pi(1 - E_{avg})} \tag{6}$$

with the Fresnel term defined as:

$$F_{ms} \doteq \frac{F_{ss} E_{avg}}{1 - F_{ss}\left(1 - E_{avg}\right)} \tag{7}$$

and with $F_{ss} = 2\int_0^1 F\left(\mu\right)\mu d\mu$ computed assuming diffuse reflection and $E_{avg} = 2\int_0^1 E_{ss}\left(\mu\right)\mu d\mu$. Note that they choose to remap F and E_{ss} with $\mu = |\omega \cdot n|$ instead of ω, which is perfectly fine as both terms are azimuthally invariant. The detailed derivation of the F_{ms} term, obtained using a geometric series, is given

in [Jakob et al. 2014], Section 5.6. The overall method relies on small 1D (α) and 2D (α, μ) look-up tables for the smoothly varying E_{avg} and E_{ss} (with dimensions of 32 and 32×32, respectively, in their implementation), which have to be precomputed for each given microfacet distribution. By design the lobe is also reciprocal (*i.e.* $\rho_{ms}(\omega_o, \omega_i) = \rho_{ms}(\omega_i, \omega_o)$), which is crucial if used in a bi-directional setting.

As an additional contribution, given that F also depends on its IOR (two wavelength-dependent values for conductors), they offer simple analytical fits of F_{ss} for various Fresnel formulations (Schlick approximation, conductor, artist-friendly conductor). They also support rough dielectric reflection and refraction, this time using a couple of tables, one for each type of interface (exterior towards interior, and the other way around). In this case, Fresnel has to be incorporated into the energy LUTs, as it now represents a ratio between reflected and refracted rays, thus adding a dimension to tables of otherwise matching resolution.

All in all, this method is very fast and robust, and can be added to a renderer with relatively minor modifications to its BSDFs. However, it elects to use a diffuse-looking ρ_{ms} lobe, hence contradicting previous results from Heitz et al. [2016], which showed that the secondary lobes are rather similar in shape to the primary one.

3 Our method

Whereas Kulla and Conty's main concern and design constraint was to preserve reciprocity, our method stems from this observed similarity of shapes between lobes, while also aiming at extreme minimalism (unless more complexity would prove beneficial to the user). This brought us to simply try and use a scaled ρ_{ss} as ρ_{ms}:

$$\rho(\omega_o, \omega_i) \doteq \rho_{ss}(\omega_o, \omega_i) + F_{ms} k_{ms}(\omega_o) \rho_{ss}(\omega_o, \omega_i) \quad (8)$$

with k_{ms} a factor to be defined, accounting for the missing energy.

In this section, we use a GGX BRDF, unless mentioned otherwise.

3.1 Conductors

3.1.1 Energy term

If we push aside the effect of Fresnel on the multiple scattering term, we can actually see ρ as a normalized version of ρ_{ss}:

$$\rho(\omega_o, \omega_i) = (1 + k_{ms}(\omega_o)) \rho_{ss}(\omega_o, \omega_i) = \frac{\rho_{ss}(\omega_o, \omega_i)}{E_{ss}(\omega_o)} \quad (9)$$

Naturally giving us:

$$k_{ms}(\omega_o) = \frac{1 - E_{ss}(\omega_o)}{E_{ss}(\omega_o)} \quad (10)$$

And by construction, such a BRDF verifies Equation 4:

$$E(\omega_o) = \int_{\Omega_i} \frac{\rho_{ss}(\omega_o, \omega_i)}{E_{ss}(\omega_o)} |\omega_i \cdot n| d\omega_i = \frac{E_{ss}(\omega_o)}{E_{ss}(\omega_o)} = 1 \quad (11)$$

Due to its smooth variation, we found like Kulla and Conty that E_{ss} can be precomputed (again, with $F = 1$) and stored in a small LUT, parameterized by outgoing angle (or cosine) and roughness. In our implementation, we too use 32×32 tables with $\cos\theta_o$ and $\sqrt{\alpha}$ (see Figure 5). For distributions with an additional γ tail-length parameter, such as GTR or STD, we extend the tables to $32 \times 32 \times 32$ with a γ-mapping tailored to each distribution.

But while the energy term is probably the most critical component of the overall compensation, the Fresnel absorption happening at each bounce on a microfacet is far from negligible, as shown in Figure 6. We still need to find a satisfying F_{ms} term.

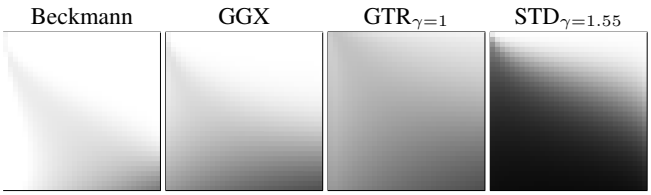

Figure 5: *E_{ss} for different microfacet models, quantized at 32×32. Left to right: $\cos\theta_o$, and top to bottom: $\sqrt{\alpha} \in [0, 1]$.*

3.1.2 Fresnel term

Before going any further, we have to acknowledge that the decision to model multiple scattering with a single lobe, and to consider F_{ms} as separable and directionally invariant, is a very coarse approximation. Within this design, we can only expect a partial match at best. However, a simple qualitative match will likely be good enough for our production needs, and for a user it makes the behaviour of the compensation more easily predictable, with a saturation directly proportional to the amount of energy being added.

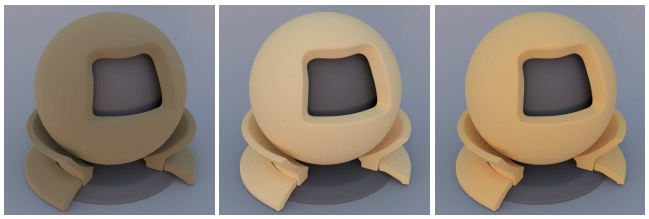

Figure 6: *Comparison between no mult. scattering compensation, compensation without and with F_{ms}, on a rough GGX conductor. Note the difference in color saturation when Fresnel is applied.*

A great option for F_{ms} is to simply reuse the term defined in Eq. 7. This function goes from being exactly equal to F_{ss} when $\alpha = 0$ (or $E_{avg} = 1$), corresponding indeed to a case where there is only single-scattering, to being close to F_{ss}^2 for $\alpha = 1$ (or $E_{avg} \simeq 0.4$).

Looking at Figure 7 left[1], we see that F_{ss}^2 would give slightly more saturation than even this formulation, which happens to be visually a bit closer to ground truth results (more on that in Section 4).

We use it with a minor alteration though, as ρ_{ss} already incorporates a directional single-scattering Fresnel coefficient F (see Equation 1), effectively replacing F_{ss} at the numerator. We can justify this as a way to give more directionality to our overall effect, counter-balancing the diffuse assumption made in the original derivation.

$$F_{ms} \doteq \frac{E_{avg}}{1 - F_{ss}(1 - E_{avg})} \quad (12)$$

The function's behaviour corroborates measurements made with the Heitz [2016] model, showing that the average random walk depth (which can be seen as an exponent for F_{ss}) varies from 1 to almost 2, when α goes from 0 to 1 (see Figure 7 right).

Exploiting these results, another (even simpler) suitable formulation for F_{ms} would be:

$$F_{ms} \doteq F_{ss}^{\alpha\sqrt{\alpha}} \quad (13)$$

We can simplify the Fresnel term further still, by observing that, as a multiplier of k_{ms} in Eq. 8, it makes more of an impact when the k_{ms} value is high, which corresponds to large roughness values.

[1] A dynamic version of these curves is available online at [Turquin 2017].

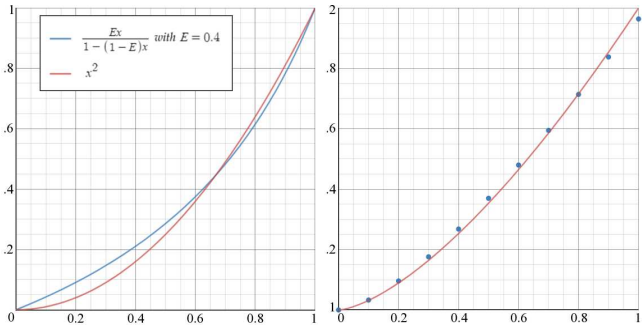

Figure 7: ***Left:*** *F_{ms} as defined in Eq. 7, for $E_{avg} = 0.4$ (i.e. $\alpha \simeq 1$), as a function of F_{ss}. We can see that it gets close to F_{ss}^2.* ***Right:*** *In blue: Measured average random walk depth for GGX, as a function of α. In red: fitting the data with $1 + \alpha\sqrt{\alpha}$.*

Since the influence of F_{ms} is most visible when $\alpha = 1$, we can strip it down to the value it takes for such a roughness:

$$F_{ms} \doteq F_{ss} \tag{14}$$

We are not completely done yet: a last simplification can be applied. F_{ss}, as defined in [Kulla and Conty 2017], is expressed as a polynomial fit with wavelength-dependent parameters η and κ, or alternatively F_0 and F_{edge} (F_0 being the color at normal angle; see [Gulbrandsen 2014] for details on this parameterization). Its evaluation has a cost, especially if targetting real-time applications.

In our tests, for strongly tinted, physically plausible conductors such as gold or copper, F_{ss} and F_0 values are so close that they can be interchangeably used with no visual difference in the results. Note that this observation only holds when using a physical Fresnel; if we pick a more artist-driven one where grazing values can be arbitrarily colored, it is recommended to stick with Eq. 14.

This leads us to the ultimate, bare-bones formula:

$$F_{ms} \doteq F_0 \tag{15}$$

and if we rewrite our BRDF defined in Eq. 8 more explicitly:

$$\rho\left(\omega_o, \omega_i\right) = \left(1 + F_0 \frac{1 - E_{ss}(\omega_o)}{E_{ss}(\omega_o)}\right) \rho_{ss}\left(\omega_o, \omega_i\right) \tag{16}$$

This is the formulation that has been used to generate the results of Figure 1, bottom row. Interestingly, no matter what Fresnel term is chosen between 12, 13, 14 and 15, the results are visually very close, and they all achieve the saturated effect we were originally looking for. It is likely that anything close to an F^2 shape towards high roughnesses would do the trick for our production needs.

3.2 Dielectrics

For dielectrics, the approach is similar with few key differences. We cannot neglect Fresnel when determining albedo normalization the way we did with conductors in Eq. 9, as it is indicative of the ratio between reflected E_{ss}^R and transmitted E_{ss}^T energies, and not just reflected *vs* absorbed anymore. What can be normalized, however, is the sum $E_{ss}^S = E_{ss}^R + E_{ss}^T$, each time an interface is hit.

If we define our single-scattering dielectric BSDF as:

$$\rho_{ss}^S\left(\omega_o, \omega_i\right) \doteq \rho_{ss}^R\left(\omega_o, \omega_i\right) + \rho_{ss}^T\left(\omega_o, \omega_i\right) \tag{17}$$

and consider that the same reflection/transmission ratio applies to multiple scattering, we can obtain our *energy preserving* BSDF:

$$\rho^S\left(\omega_o, \omega_i\right) \doteq \frac{\rho_{ss}^R\left(\omega_o, \omega_i\right) + \rho_{ss}^T\left(\omega_o, \omega_i\right)}{E_{ss}^S(\omega_o)} \tag{18}$$

Given that Fresnel is included in the E_{ss}^S term, our LUT becomes 3D, with the IOR as an additional dimension. On the other hand, this means that we do not need to worry about a separate multiple scattering Fresnel term anymore. As Kulla and Conty [2017], we compute tables separately for the two types of interface, leading to two $32 \times 32 \times 32$ tables (note that the tables would even be 4D if computed for a tail-length parametrized model), and use the bounded $F_0 \in [0, 1]$, instead of η for the extra dimension.

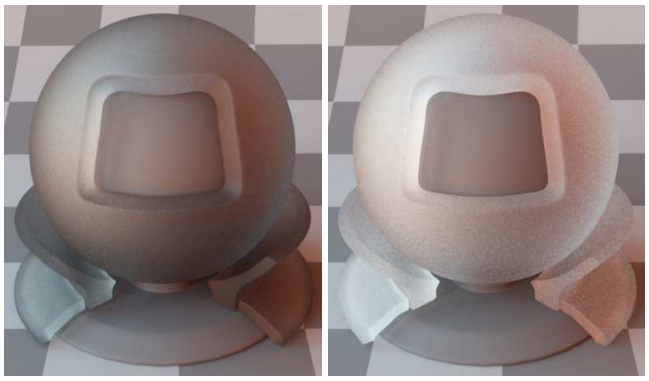

Figure 8: *Comparison between rough ($\alpha = 0.5$) GGX glass object without (left) and with (right) multiple scattering compensation.*

As can be seen in Figure 8, the effect of this compensation is quite noticeable – maybe even more so than with conductors.

4 Comparison with other approaches

The method presented in Section 3 has been implemented in *Isotropix Clarisse*, *Pixar RenderMan's RIS* and the *Mitsuba* path tracer, but could easily be added to any other renderer (including real-time rasterization engines).

It is close to Kulla and Conty's, essentially with a different choice of shape for ρ_{ms}: the same lobe as ρ_{ss} in our case, as opposed to a diffuse one in theirs. The former is closer to the ground truth (as shown in [Heitz et al. 2016]), but not reciprocal, which could be a deal breaker when the resulting BSDF is used with bidirectional light transport. However, we believe that this lack of formal reciprocity is unlikely to produce any significant artefacts.

Because our method depends only on ω_o, rather than (ω_o, ω_i), it is even simpler to integrate into an existing renderer, as the compensation can be applied as a gain to the closure, instead of having to modify the closure itself. Being simpler, it is marginally faster too.

Another beneficial side effect is that, as our result can be seen as a scaling of ρ_{ss}, it perfectly maintains its shape, meaning we can reuse the exact same sampling and evaluation routines, with the same PDF. In contrast, Kulla and Conty's lobe is a mixture of ρ_{ss} and an azimuthally invariant lobe, for which optimal analytical importance sampling cannot be achieved.

Both methods, while theoretically less elegant or accurate, and far from representing the same step forward scientifically, are markedly easier to implement than Heitz et al. [2016], and much faster at runtime. In tests performed in *Mitsuba* using their original implementation, a $7\times$ slowdown was observed on a simple conductor test, and as high as $15\times$ on a dielectric transmission test, while the visual difference appeared minimal (see results in Figure 9).

The most notable difference is the amount of saturation on conductors, with the reference looking slightly more saturated at high roughnesses, even with us using F_{ms} of Eq. 15, rather than the more justified one of Eq. 12, that saturates the result a bit less (Fig. 7 left).

Method	Shape of secondary lobes	Reciprocity	Optimal sampling	Plausibility	Simplicity	Speed	Flexibility
Heitz et al.	Similar to primary	Yes	No	A+	C	C	B
Kulla and Conty	Azimuthally invariant	Yes	No	B	A	A	A+
Ours	Same as primary	No	Yes	B+	A+	A+	A+

Table 1: *Qualitative summary of the pros and cons of each method.*

In any case, we cannot expect a perfect match, given that F_{ms} should not be separated from ρ_{ms} as we did, and would definitely vary directionally. We believe, however, that the results are close enough to warrant keeping the approach as simple as it is.

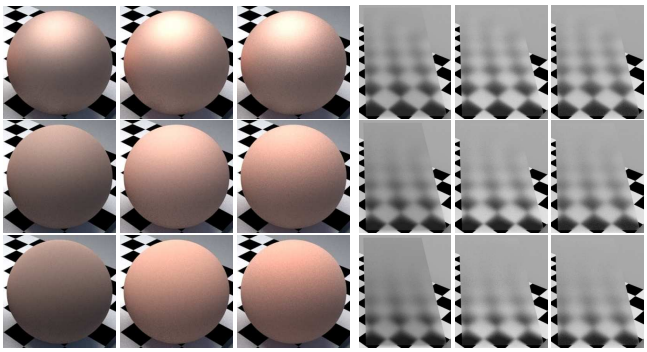

Figure 9: *Comparisons for rough GGX reflection and transmission. Left to right: no compensation, ours and [Heitz et al. 2016]. Top to bottom: medium, high and max roughness.*

Finally, both [Kulla and Conty 2017] and our method can accommodate any sort of BSDF, as they only rely on the computation of the single-scattering albedo E_{ss}, whereas [Heitz et al. 2016] imposes considerably stronger constraints, making it even more difficult to use for models other than Beckmann and GGX.

For a quick summary of all pros and cons, please refer to Table 1.

5 Conclusion and future work

We have presented different approaches to compensate for the missing multiple scattering energy in rough reflections and transmissions. For a production use, the most suitable candidates appear to be either [Kulla and Conty 2017] or the one introduced here, depending on how critical reciprocity is for the renderer it is integrated into. We are convinced that either one will be of great help to artists, and remove the need for more error-prone, *ad hoc* fixes.

Comparisons in this report have been largely of a qualitative nature. A more thorough and quantitative analysis of the differences between the three main methods is left for future work.

Furthermore, it would be interesting to find good analytical fits for all of the smooth look-up tables, which might be beneficial to a real-time implementation, by virtue of avoiding costly texture accesses.

We could also consider not separating Fresnel for conductors anymore, and instead extend the LUT to 3D (or 4D if we really want to take the imaginary part of the IOR into account), computing it with a reference algorithm such as [Heitz et al. 2016] (at least for the Beckmann and GGX distributions).

Speaking of which, as a much more ambitious future endeavour, it would be fantastic to figure out a way to turn their method into something more production friendly – the critical element being speed here, with ease of use coming second.

Acknowledgements

Big thanks are in order to the ILM Rendering R&D team, Stephen Hill and Naty Hoffman from Lucasfilm, Brent Burley and Matt Chiang from Disney Animation, and Chris Kulla from Sony Imageworks, for insightful discussions and feedback that greatly helped shape this document and figure out its trickiest aspects.

References

BURLEY, B., 2012. Physically based shading at Disney. ACM SIGGRAPH course notes, `http://blog.selfshadow.com/publications/s2012-shading-course/`.

BURLEY, B., 2015. Extending the Disney BRDF to a BSDF with integrated subsurface scattering. ACM SIGGRAPH course notes, `http://blog.selfshadow.com/publications/s2015-shading-course/`.

COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph. 1*, 1 (Jan.), 7–24.

GULBRANDSEN, O. 2014. Artist Friendly Metallic Fresnel. *Journal of Computer Graphics Techniques (JCGT) 3*, 4 (December), 64–72.

HEITZ, E., HANIKA, J., D'EON, E., AND DACHSBACHER, C. 2016. Multiple-scattering microfacet BSDFs with the Smith model. *ACM Trans. Graph. 35*, 4 (July), 58:1–58:14.

HEITZ, E. 2014. Understanding the masking-shadowing function in microfacet-based BRDFs. *Journal of Computer Graphics Techniques (JCGT) 3*, 2 (June), 48–107.

JAKOB, W., D'EON, E., JAKOB, O., AND MARSCHNER, S., 2014. A comprehensive framework for rendering layered materials. Expanded technical report, `http://www.cs.cornell.edu/projects/layered-sg14/`.

KELEMEN, C., AND SZIRMAY-KALOS, L., 2001. A Microfacet based coupled specular-matte BRDF model with importance sampling. Short presentation, Eurographics.

KULLA, C., AND CONTY, A., 2017. Revisiting Physically Based Shading at Imageworks. ACM SIGGRAPH course notes, `http://blog.selfshadow.com/publications/s2017-shading-course/`.

RIBARDIÈRE, M., BRINGIER, B., MENEVEAUX, D., AND SIMONOT, L. 2017. STD: Student's t-Distribution of slopes for microfacet based BSDFs. *Computer Graphics Forum*.

TURQUIN, E., 2017. Multiple scattering Fresnel curves. `https://www.desmos.com/calculator/0mig1qjjvf`.

WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques*, EGSR'07, Eurographics, 195–206.