

Fresnel Equations Considered Harmful

N. Hoffman

Lucasfilm, United States

Abstract

Microfacet shading models in film and game production have long used a simple approximation to the Fresnel equations, published by Schlick in 1994. Recently a growing number of film studios have transitioned to using the full Fresnel reflectance equations in lieu of Schlick's approximation. This transition has been facilitated by Gulbrandsen's 2014 parameterization which uses reflectance and edge tint instead of eta and kappa.

Our recent investigations have found some unexpected drawbacks to this approach. In this presentation, we will show that in the context of RGB rendering (still by far the most common modality in film production), the Fresnel equations are actually less physically principled than Schlick's approximation. In addition, they are less robust in practice and less amenable to authoring. Most surprisingly, as commonly used the Fresnel equations result in less accurate matches to measured materials, compared to Schlick's approximation.

The presentation primarily discusses metal reflectance, since our investigations so far have focused on metals.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Rendering—Reflectance modeling

1. Introduction

Schlick's Fresnel approximation [Sch94] has long been the most commonly used Fresnel shading term in film and game production [BHS04, Che08, Mar10, Bur12, Smi12, Kar13, LdR14]. More recently, formulations based directly on the Fresnel reflectance equations [Fre23] have been increasing in popularity, used either as a direct replacement for Schlick's approximation [Lan14, XL17] or as an optional model [And13, HL17, KC17].

The first such formulations were parameterized with RGB values for η and κ (eta and kappa, the real and imaginary parts of the complex index of refraction). Currently an alternate parameterization [Gul14] is typically used, with reflectivity and edge tint RGB parameters which are effectively a "user interface" over η and κ (they are used to compute RGB η and κ values which are then applied to the Fresnel equations).

Various motivations may drive the use of Fresnel equations in rendering: physical correctness in principle, robustness in production use, accurate modeling of real-world substances, or artistic control of edge behavior. Similar motivations led us to explore this approach for our own potential future shading models. However, during those investigations we discovered that in the common case of RGB rendering, these expected improvements either do not materialize, or can be better achieved via other models.

2. Schlick's Fresnel Approximation

In 1994, Schlick [Sch94] introduced the following approximation of the Fresnel equations:

$$F(\theta) \approx r + (1 - r)(1 - \cos \theta)^5 \quad (1)$$

Where θ is the angle of incidence and r is the Fresnel reflectance at normal incidence ($\theta = 0^\circ$). In practice, Equation 1 is applied in RGB rendering, with RGB values for $F(\theta)$ and r .

To evaluate the accuracy of this approximation, we must first carefully define the ground truth it should be compared to. The ground truth for RGB Fresnel computation for a given angle θ is the spectral evaluation of the Fresnel equations at that angle, using spectral values for η and κ (for example, measured values which are available for various metals). This results in a spectral reflectance curve, which is then multiplied by the spectral power distribution of a neutral white illuminant (e.g., D60) as appropriate to the color space used. This yields the spectral power distribution of the reflected radiance, which directly corresponds to the stimulus perceived by an observer viewing the reflected light. This spectral stimulus is converted to an RGB triple using standard methods: convolution with color-matching curves and multiplication by a 3x3 matrix, with an additional chromatic adaptation step in the case where the color space white point has a different chromaticity than the reference illuminant, e.g., as in the ACEScg color

space [Aca14]. This process is shown below the dotted red line in Figure 1.

Overall Schlick's approximation is quite accurate; for the majority of metals the average error is under 2 dE (CIE2000 delta E). The full presentation includes detailed tables, plots, renderings and difference images showing the error of Schlick's approximation for various metals.

However, the accuracy of Schlick's approximation is not perfect—if possible, would it not be preferable to use the original equations instead? The rest of the presentation will go over the hoped-for advantages from using the Fresnel equations over Schlick's approximation: physical correctness in principle, robustness in production use, accurate modeling of real-world substances, and artistic control of edge behavior—and show that in practice, they do not materialize as expected.

3. Physical Correctness in Principle

It would seem that using the original Fresnel equations would inarguably be more correct than any approximation, from a theoretical standpoint at least. But the problem is that the Fresnel equations are intended to be used on individual spectral samples of η and κ , as in the ground-truth example earlier. And to derive RGB Fresnel values, the rather involved procedure shown in the bottom part (below the dotted red line) of Figure 1 must be applied. However, that is not how the Fresnel equations are used in practice. RGB rendering is still ubiquitous in production rendering—so far, only Weta Digital has adopted spectral rendering in a significant way for film production [Han17]. When RGB production renderers use the Fresnel equations, it is in the manner shown in the upper part (above the dotted red line) of Figure 1. RGB values for η and κ are derived from the spectral curves (a process with no physical meaning, for which there is no a priori correct method), and the Fresnel equations are applied directly to these values to produce an RGB result. At this point any claim to physical correctness has been forfeited—using the Fresnel equations in this way is just another approximation, like Schlick's.

But there is an important sense in which this approximation is *less* physically principled than Schlick's. RGB values are only meaningful for expressing perceptual quantities. And the Fresnel equations work on physical quantities, which need to be expressed spectrally.

The only physical rendering quantity which is also a perceptual one is radiance, which is directly processed as a stimulus by the human visual system and thus can be meaningfully expressed as an RGB color. RGB can also be meaningfully used for reflectance colors, which (as in the ground truth example) are defined indirectly as the reflected radiance from the white reference illuminant. This definition enables us to perform RGB reflectance calculations by multiplying the RGB color of the incoming light with RGB reflectance values. In the case of white (reference illuminant) light reflected directly from a surface, this method produces exact results. In most other cases, the results are at least visually plausible, if not exactly correct. Some issues do occur: the results vary between different working color spaces [Agl14] and in certain scenarios (e.g., multiple-bounce diffuse illumination with

highly saturated surfaces) RGB rendering can produce significant errors [MSHD15]. However, in practice, these errors are minor or rare, and (as mentioned above) RGB rendering is by far the most common type of production rendering.

RGB rendering tends to work fairly well as long as the RGB operations are linear. Typically, most nonlinear operations in shaders are applied to non-color data, for example evaluating specular distribution lobes with roughness values. Schlick's approximation fits into this framework; while its angle-dependent part is highly nonlinear, all RGB computations are simple adds and multiplies. The dependence of the resulting reflectance value on the parameter r is strictly linear, which is a highly desirable property. Since r can be defined as both a physical and a perceptual quantity, its expression as an RGB color in rendering equations is meaningful.

In contrast, when the full Fresnel equations are used in an RGB renderer, highly nonlinear operations are performed on RGB quantities. The Gulbrandsen (r, g) parameterization does not change this because it does not change the underlying equations.

Unlike the reflectance r parameter, the η and κ quantities in the full Fresnel equations (as well as the Gulbrandsen g parameter) have no perceptual analog. Their relationship to final rendered colors (perceptual stimuli) is very indirect and highly nonlinear. There is no principled way to compute RGB values for η and κ , and they are not amenable to painting. Gulbrandsen's (r, g) parameterization attempts to fix this issue, but fails. While r is a meaningful color, the edge tint parameter g is not (though it superficially appears to be one). Like η and κ , the value of g does not correspond to any perceptual quantity.

To summarize, in the context of RGB rendering the Fresnel equations are as much an approximation as Schlick's approximation. Schlick's is actually a *more principled* RGB approximation since it works on quantities that are both perceptually and physically meaningful.

4. Robustness in Production Use

The presentation will show that for RGB rendering, the Fresnel equations are less numerically robust than Schlick's approximation. Two cases will be shown: converting parameter values to a new color space, and blending of parameter values. Both are important operations for production rendering, and in both cases the Fresnel equations introduce significant errors when the operation is performed—errors which are not present (or significantly reduced) if Schlick's approximation is used.

5. Accurate Modeling of Real-world Substances

In the presentation, we will cover two cases: computing parameter values from measured data, and painting parameter values from photographic reference. For painting from photographic reference (which is the most common case in production), Schlick's approximation will be shown to result in more accurate material models than the Fresnel equations. Gulbrandsen's g parameter is not well-suited for matching by eye since it does not correspond to the reflected color at any angle. It also does not have a reasonable default value, and there is no straightforward way to "opt out" of using it.

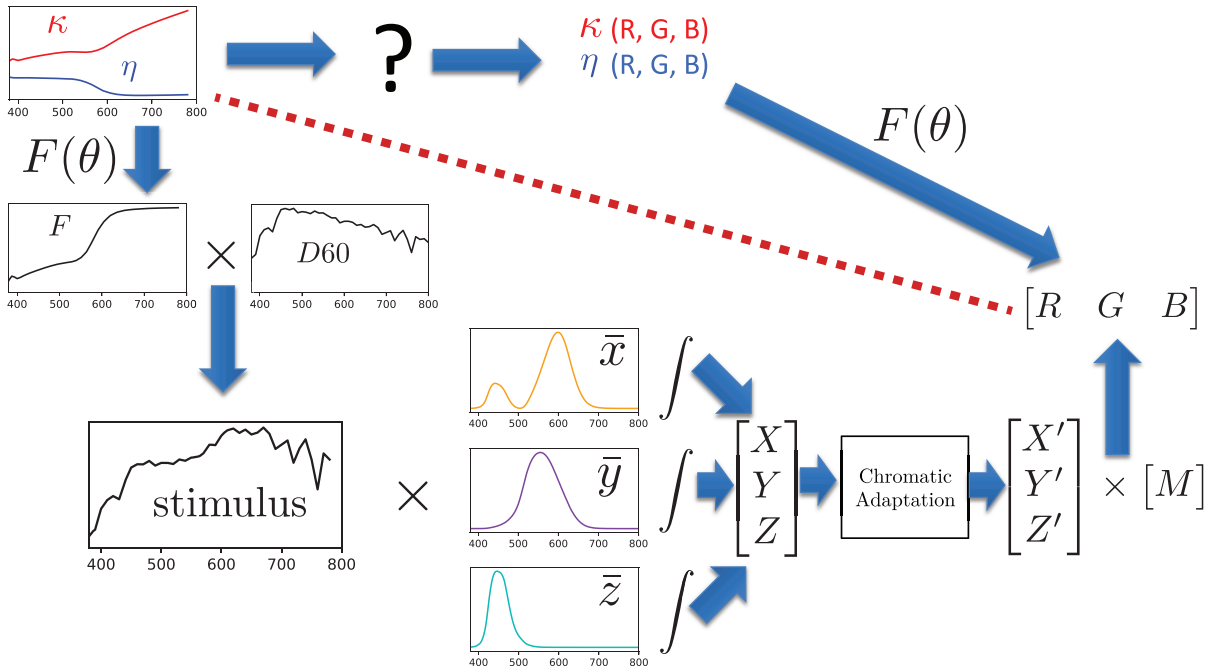


Figure 1: The Fresnel equations were derived for use on spectral values of η and κ , producing a spectral reflectance curve. To convert this to an RGB Fresnel value, this curve is first multiplied by the SPD of a reference illuminant to produce the “reflected white light” SPD. This SPD corresponds to a perceptual stimulus that can be converted to RGB tristimulus values in the standard way, including convolution with color-matching functions, possibly chromatic adaptation, and then multiplication by a 3×3 matrix. This is shown under the dotted red line. Above the line we see the way in which the Fresnel equations are actually used in production RGB rendering. RGB values for η and κ are—somehow—derived from the spectral curves, and then the Fresnel equations are applied directly to these RGB values.

The case of computing values from measured data is more complex. It is possible to achieve highly accurate results with the Fresnel equations, but only when taking great care when computing parameter values. The methods which are currently in common production use (e.g., spectral point sampling) introduce errors which are as large or larger than those resulting from Schlick’s approximation. The most accurate results are achieved via black-box numerical fitting, which is to my knowledge not currently used by any production studio. This is not surprising since the process of computing RGB values for η , κ , or Gulbrandsen’s g parameter is physically meaningless, and there is no principled method for doing so.

6. Artist-driven Control of Edge Appearance

The fact that Gulbrandsen’s g parameter does not correspond to the reflected color at any angle not only makes it difficult to match photographic reference, it also poses difficulties for creative look development. When painting a novel material, there is no simple way for an artist to predict the resulting appearance based on the value of g . The presentation will also show that the g parameter lacks two properties that are important for artist-driven control of material appearance: decoupling from other parameters, and perceptually uniform behavior over the valid range of values.

7. Complexity

Besides the flaws detailed so far, the Fresnel equations also have greatly increased computational complexity over Schlick’s approximation. This increased cost may be a concern for offline rendering times, and it is prohibitive for real-time applications. Having a shared shading model for offline and real-time rendering can be highly beneficial for multiple-use assets [CL16], and using the Fresnel equations makes this difficult.

8. An Alternative Model

In many cases Schlick’s approximation is likely to be the best choice for Fresnel RGB evaluation. However, some applications may require higher accuracy or artistic control over edge falloff. To address these cases, we have developed a reparameterization of Lazanyi’s model [LSK05]. This model extends Schlick’s approximation with an additional term, designed to compensate for most of its error:

$$F(\theta) \approx r + (1 - r)(1 - \cos\theta)^5 - a \cos\theta(1 - \cos\theta)^\alpha \quad (2)$$

We fix the value of the α parameter at 6, which puts the maximum of the absolute error term at $\theta_{\max} = \arccos\left(\frac{1}{7}\right)$. This angle is

close to the peak of the actual Schlick error term for the metals we investigated.

We also introduce a new parameter h , which maps to the a parameter thus:

$$a = \frac{r + (1-r)(1 - \cos \theta_{\max})^5 - h}{\cos \theta_{\max}(1 - \cos \theta_{\max})^6} \quad (3)$$

After substituting ($\cos \theta_{\max} = \frac{1}{7}$) we get:

$$a = \frac{823543}{46656}(r-h) + \frac{49}{6}(1-r) \quad (4)$$

The value of the h parameter is equal to the Fresnel reflectance at θ_{\max} , which is a well-defined color quantity.

In the presentation, we compare this model to Schlick's approximation as well as Fresnel's equations. We show that the reparameterized Lazanyi model has lower error than Schlick's approximation, allows for well-behaved artist control over edge behavior, and does not exhibit the nonlinearity issues of Fresnel's equations.

9. Conclusions

The use of the Fresnel equations (facilitated by Gulbrandsen's parameterization) seems at first to be an improvement over Schlick's approximation, but further analysis shows it to be worse on every important axis. In most cases Schlick's approximation is likely the best choice for RGB renderers. For cases where additional accuracy or control are needed, we present an alternative model based on Lazanyi's error term.

References

- [Aca14] ACADEMY COLOR ENCODING SYSTEM (ACES) PROJECT COMMITTEE: *ACEScg—A Working Space for CGI Render and Compositing*. Standard S-2014-004, The Academy of Motion Picture Arts and Sciences—Science and Technology Council, 2014. URL: <http://j.mp/S-2014-004>. 8
- [Agl14] AGLAND S.: CG rendering and ACES, 2014. URL: <https://nbviewer.jupyter.org/gist/sagland/3c791e79353673fd24fa>. 8
- [And13] ANDERSSON Z.: Everything you always wanted to know about mia_material. In *Physically Based Shading in Theory and Practice, ACM SIGGRAPH 2013 Courses* (2013), SIGGRAPH '13, ACM. URL: <http://selfshadow.com/publications/s2013-shading-course/>. 7
- [BHS04] BAKER D., HOFFMAN N., SLOAN P.-P.: Advanced real-time reflectance. In *GDC 2004* (2004), Game Developers Conference, CMP Media. URL: <http://renderwonk.com/publications/gdc-2004/>. 7
- [Bur12] BURLEY B.: Physically based shading at Disney. In *Practical Physically Based Shading in Film and Game Production, ACM SIGGRAPH 2012 Courses* (2012), SIGGRAPH '12, ACM. URL: <http://selfshadow.com/publications/s2012-shading-course/>. 7
- [Che08] CHEN H.: Lighting and material of Halo 3. In *GDC 2008* (2008), Game Developers Conference, CMP Media. URL: <https://www.gdcvault.com/play/253/Lighting-and-Material-of-HALO>. 7
- [CL16] CORDES R., LOBL D.: Unified shading and asset development at Lucasfilm and ILM. In *Physically Based Shading in Theory and Practice, ACM SIGGRAPH 2016 Courses* (2016), SIGGRAPH '16, ACM. URL: <http://selfshadow.com/publications/s2016-shading-course/>. 9
- [Fre23] FRESNEL A.-J.: Mémoire sur la loi des modifications que la réflexion imprime à la lumière polarisée (memoir on the law of the modifications that reflection impresses on polarized light). In *Mémoires* (1823), French Academy of Sciences. 7
- [Gul14] GULBRANDSEN O.: Artist friendly metallic Fresnel. *Journal of Computer Graphics Techniques (JCGT)* 3, 4 (December 2014), 64–72. URL: <http://jcgt.org/published/0003/04/03/>. 7
- [Han17] HANIKA J.: Manuka: Weta Digital's spectral renderer. In *Path Tracing in Production, ACM SIGGRAPH 2017 Courses* (2017), SIGGRAPH '17, ACM. URL: <https://jo.dreggn.org/path-tracing-in-production/2017/index.html>. 8
- [HL17] HERY C., LING J.: Pixar's foundation for materials: PxrSurface and PxrMarschnerHair. In *Physically Based Shading in Theory and Practice, ACM SIGGRAPH 2017 Courses* (2017), SIGGRAPH '17, ACM. URL: <http://selfshadow.com/publications/s2017-shading-course/>. 7
- [Karl13] KARIS B.: Real shading in Unreal Engine 4. In *Physically Based Shading in Theory and Practice, ACM SIGGRAPH 2013 Courses* (2013), SIGGRAPH '13, ACM. URL: <http://selfshadow.com/publications/s2013-shading-course/>. 7
- [KC17] KULLA C., CONTY A.: Revisiting physically based shading at Imageworks. In *Physically Based Shading in Theory and Practice, ACM SIGGRAPH 2017 Courses* (2017), SIGGRAPH '17, ACM. URL: <http://selfshadow.com/publications/s2017-shading-course/>. 7
- [Lan14] LANGLANDS A.: Physically based shader design in Arnold. In *Physically Based Shading in Theory and Practice, ACM SIGGRAPH 2014 Courses* (2014), SIGGRAPH '14, ACM. URL: <http://selfshadow.com/publications/s2014-shading-course/>. 7
- [LdR14] LAGARDE S., DE ROUSIERS C.: Moving Frostbite to PBR. In *Physically Based Shading in Theory and Practice, ACM SIGGRAPH 2014 Courses* (2014), SIGGRAPH '14, ACM. URL: <http://selfshadow.com/publications/s2014-shading-course/>. 7
- [LSK05] LAZÁNYI I., SZIRMAY-KALOS L.: Fresnel term approximations for metals. In *WSCG* (2005). URL: http://wscg.zcu.cz/WSCG2005/Papers_2005/Short/H29-full.pdf. 9
- [Mar10] MARTINEZ A.: Faster photorealism in Wonderland: Physically-based shading and lighting at Sony Pictures Imageworks. In *Physically-Based Shading Models in Film and Game Production, ACM SIGGRAPH 2010 Courses* (2010), SIGGRAPH '10, ACM. URL: <http://renderwonk.com/publications/s2010-shading-course/>. 7
- [MSHD15] MENG J., SIMON F., HANIKA J., DACHSBACHER C.: Physically meaningful rendering using tristimulus colours. *Computer Graphics Forum* 34, 4 (2015), 31–40. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12676>, doi:10.1111/cgf.12676. 8
- [Sch94] SCHLICK C.: An inexpensive brdf model for physically-based rendering. *Computer Graphics Forum* 13, 3 (1994), 233–246. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8659.1330233>, doi:10.1111/1467-8659.1330233. 7
- [Smi12] SMITS B.: Reflection model design for WALL-E and Up. In *Practical Physically Based Shading in Film and Game Production, ACM SIGGRAPH 2012 Courses* (2012), SIGGRAPH '12, ACM. URL: <http://selfshadow.com/publications/s2012-shading-course/>. 7

- [XL17] XIE F., LANZ J.: Physically based shading at Dream-Works Animation. In *Physically Based Shading in Theory and Practice, ACM SIGGRAPH 2017 Courses* (2017), SIGGRAPH '17, ACM. URL: <http://selfshadow.com/publications/s2017-shading-course/.7>