

Physically-Based Shading Models in Film and Game Production

SIGGRAPH 2010 COURSE NOTES

Course Organizer

NATY HOFFMAN
Activision Studio Central

Presenters

YOSHIHARU GOTANDA
tri-Ace

ADAM MARTINEZ
Sony Pictures Imageworks

BEN SNOW
ILM

Course Description

Physically grounded shading models have been known for many years, but they have only recently started to replace the "ad-hoc" models in common use for both film and game production. Compared to "ad-hoc" models, which require laborious tweaking to produce high-quality images, physically-based, energy-conserving shading models easily create materials that hold up under a variety of lighting environments. These advantages apply to both photorealistic and stylized scenes, and to game development as well as production of CG animation and computer VFX. Surprisingly, physically-based models are not more difficult to implement or evaluate than the traditional "ad-hoc" ones.

This course begins with a short explanation of the physics of light-matter interaction and how it is expressed in simple shading models. Then several speakers discuss specific examples of how shading models have been used in film and game production. In each case, the advantages of the new models are demonstrated, and drawbacks or issues arising from their usage are discussed. The course also includes descriptions of specific production techniques related to physically-based shading.

LEVEL OF DIFFICULTY: Intermediate

Intended Audience

Practitioners from the videogame, CG animation, and VFX fields, as well as researchers interested in shading models.

Prerequisites

Basic familiarity with computer graphics, illumination and shading models in particular.

Course Website

All course materials can be found at <http://renderwonk.com/publications/s2010-shading-course>

Contact

Address questions or comments to s2010course@renderwonk.com

About the Presenters

YOSHIHARU GOTANDA is the CEO and CTO of tri-Ace, Inc, which is a game development studio in Japan.

NATY HOFFMAN is a Technical Director at Activision Studio Central, where he assists Activision's worldwide studios with graphics research and development. Prior to joining Activision in 2008, Naty worked for two years on *God of War III* at SCEA Santa Monica Studio. Naty has also worked at Naughty Dog (where he had an instrumental role in the development of the ICE libraries for first-party PS3 developers), at Westwood Studios (where he was graphics lead on *Earth and Beyond*) and at Intel as a microprocessor architect, assisting in the definition of the SSE and SSE2 instruction set extensions.

ADAM MARTINEZ is a Computer Graphics supervisor for Sony Pictures Imageworks and a member of the Shading Department, which oversees all aspects of shader writing and production rendering at Imageworks. He is a pipeline developer, look development artist, and technical support liaison for productions at the studio and he is one of the primary architects of Imageworks' rendering strategy behind *2012* and *Alice In Wonderland*. Adam started his career in commercial post houses and animation boutiques in New York City as a freelance computer graphics artist. He began his work in film visual effects on the project *Cremaster 3* by artist-filmmaker Matthew Barney. Since then he has served as both effects and lighting technical director, CG supervisor and pipeline developer for various studios in the San Francisco Bay Area. At ESC Entertainment, Adam led the effects team in the creation of complex insect crowd simulation tools for *Constantine* and destruction effects for *Matrix:Revolutions*. As computer graphics supervisor for The Orphanage on *Superman Returns*, Adam oversaw the creation of a ballistics simulation and rendering system. At Lucas Animation Adam was both a rendering pipeline developer and CG concept artist for television and feature animation. Adam's primary interest is in simulation and the construction of complex, but highly usable, systems for dynamic effects and rendering. Adam has a BA from Rutgers University.

BEN SNOW studied computing and film at the University of Canberra. He started in Computer Graphics at while traveling in the U.K., then returned to Australia to set up the computer animation department for a company in Sydney. In 1994, Snow left Australia to join Industrial Light & Magic. At ILM he played a leading role in the R&D development for *Twister*, *Deep Impact*, *The Mummy* and *Pearl Harbor*. In 2002 he became visual effects supervisor on *Star Wars: Episode II—Attack of the Clones* for which he was honored with an Academy Award nomination for best achievement in visual effects. Snow also received Academy Award nominations for his work on *Pearl Harbor* and *Iron Man*. Snow went to Weta digital in October 2004 to work as a visual effects supervisor on *Peter Jackson's King Kong*. Returning to ILM in 2006, Snow visual effects supervised *Iron Man*, *Terminator Salvation*, and *Iron Man 2*. He's currently supervising ILM's work on *Pirates of the Caribbean: On Stranger Tides*.

Presentation Schedule

- 2:00–2:30 **Background: Physically-Based Shading** (*Hoffman*)
- 2:–3:00 **Practical Implementation of Physically-Based Shading Models at tri-Ace** (*Gotanda*)
- 3:00–3:30 **Crafting Physically Motivated Shading Models for Game Development** (*Hoffman*)
- 3:30–3:45 **Break**
- 3:45–4:30 **Terminators and Iron Men: Image-Based Lighting and Physical Shading at ILM** (*Snow*)
- 4:30–5:00 **Faster Photorealism in Wonderland: Physically-Based Shading and Lighting at Sony Pictures Imageworks** (*Martinez*)
- 5:00–5:15 **Conclusion, Q&A** (*Gotanda, Hoffman, Martinez*)

Background: Physically-Based Shading

by Naty Hoffman

In this section of the course notes, we will go over the fundamentals behind physically-based shading models, starting with a qualitative description of the underlying physics, followed by a quantitative description of the relevant mathematical models, and finally discussing how these mathematical models can be implemented for shading.

The Physics of Shading

The physical phenomena underlying shading are those related to the interaction of light with matter. To understand these phenomena, it helps to have a basic understanding of the nature of light.




Figure 1: Light is an electromagnetic transverse wave.

Light is an electromagnetic *transverse wave*, which means that it oscillates in directions perpendicular to its propagation (see Figure 1).

Since light is a wave, it is characterized by its *wavelength*—the distance from peak to peak. Electromagnetic wavelengths cover a very wide range but only a tiny part of this range (about 400 to 700 nanometers) is visible to humans and thus of interest for shading (see Figure 2).

The effect matter has on light is defined by a property called the *refractive index*. The refractive index is a complex number; its real part measures how the matter affects the speed of light (slowing




Figure 2: The visible spectrum.

it down relative to its speed in a vacuum) and the imaginary part determines whether the light is *absorbed* (converted to other forms of energy) as it propagates. The refractive index may vary as a function of light wavelength.

Homogeneous Media

The simplest case of light-matter interaction is light propagating through a *homogeneous medium*. This is a region of matter with uniform index of refraction (at the scale of the light wavelength; in the case of visible light this means that any variations much smaller than 100 nanometers or so don't count).



Figure 3: Light in transparent media like water and glass (left) just keeps on propagating in a straight line at the same intensity and color (right).

A *transparent* medium is one in which the complex part of the index of refraction is very low for visible light wavelengths; this means that there is no significant absorption and any light propagating through the medium just keeps on going in a straight line, unchanged. Examples of transparent media include water and glass (see Figure 3).

If a homogeneous medium does have significant absorptivity in the visible spectrum, it will absorb some amount of light passing through it. The farther the distance traveled by the light, the higher the absorption. However, the direction of the light will not change, just its intensity (and, if the absorptivity is selective to certain visible wavelengths, the color)—see Figure 4.

Note that the scale as well as the absorptivity of the medium matters. for example, water actually absorbs a little bit of visible light, especially on the red end of the spectrum. On a scale of inches this



Figure 4: Light propagating through clear, absorbent media (left) continues in a straight line, but loses intensity (and may change color) with distance (right).



Figure 5: The slight absorptivity of water becomes significant over larger distances.

is negligible (as seen in Figure 3) but it is quite significant over many feet of distance; see Figure 5.

Scattering

In homogeneous media, light always continues propagating in a straight line and does not change its direction (although its amount can be reduced by absorption). A *heterogeneous medium* has variations in the index of refraction. If the index of refraction changes slowly and continuously, then the light bends in a curve. However, if the index of refraction changes abruptly, over a short distance (compared to the light wavelength), then the light *scatters*; it splits into multiple directions. Note that scattering does not change the overall amount of light.

Microscopic particles induce an isolated “island” where the refraction index differs from surrounding regions. This causes light to scatter continuously over all possible outgoing directions (see Figure 6). Note that the distribution of scattered light over different directions is typically not uniform and depends on the type of particle. Some cause forward scattering (more light goes in the forward direction), some cause backscattering (more light goes in the reverse of the original direction), and some have complex distributions with “spikes” in certain directions.

In *cloudy media*, the density of scattering elements is sufficient to somewhat randomize light propagation direction (Figure 7). In *translucent* or *opaque media* the density of scattering elements is so high that the light direction is completely randomized (Figure 8).

Like absorption, scattering depends on scale; a medium such as clean air which has negligible scattering over distances of a few feet causes substantial light scattering over many miles (Figure 9).




Figure 6: Particles cause light to scatter in all directions.




Figure 7: Light in cloudy media (left) has its direction somewhat randomized as it propagates (right).




Figure 8: Light in translucent or opaque media (left) has its direction completely randomized as it propagates (right).



Figure 9: Even clean air causes considerable light scattering over a distance of miles.

Media Appearance

Previous sections discussed two different modes of interaction between matter and light. Regions of matter with complex-valued refraction indices cause absorption—the amount of light is lessened over distance (potentially also changing the light color if absorption occurs preferentially at certain wavelengths), but the light’s direction does not change. On the other hand, rapid changes in the index of refraction cause scattering—the direction of the light changes (splitting up into multiple directions), but the overall amount or spectral distribution of the light does not change. There is a third mode of interaction—*emission*, where new light is created from other forms of energy (the opposite of absorption). This occurs in light sources, but it doesn’t come up often in shading. Figure 10 illustrates the three modes of interaction.




Figure 10: The three modes of interaction between light and matter: absorption (left), scattering (middle), and emission (right).



Figure 11: Media with varying amounts of light absorption and scattering.

Most media both scatter and absorb light to some degree. Each medium’s appearance depends

on the relative amount of scattering and absorption present. Figure 11 shows media with various combinations of scattering and absorptivity.

Scattering at a Planar Boundary

Maxwell’s equations can be used to compute the behavior of light when the index of refraction changes, but in most cases analytical solutions do not exist. There is one special case which does have a solution, and it is of great relevance for shading. This is the case of an infinite, perfectly flat planar boundary between two volumes with different refractive indices. This is a good description of an object surface, with the refractive index of air on one side of the boundary, and the refractive index of the object on the other. The solutions to Maxwell’s equations in this special case are called the *Fresnel equations*.




Figure 12: Refractive index changes at planar boundaries cause light to scatter in two directions (*image from “Real-Time Rendering, 3rd edition” used with permission from A K Peters*).

Although real object surfaces are not infinite, in comparison with the wavelength of visible light they can be treated as such. As for being “perfectly flat”, an objection might be raised that no object’s surface can truly be flat—if nothing else, individual atoms will form pico-scale “bumps”. However, as with everything else, the scale relative to the light wavelength matters. It is indeed possible to make surfaces that are perfectly flat at the scale of hundreds of nanometers—such surfaces are called *optically flat* and are typically used for high-quality optical instruments such as telescopes.

In the special case of a planar refractive index boundary, instead of scattering in a continuous fashion over all possible directions, light splits into exactly two directions: reflection and refraction (see Figure 12).

As you can see in Figure 12, the angle of reflection is equal to the incoming angle, but the angle of refraction is different. The angle of refraction depends on the refractive index of the medium (if you are interested in the exact math, look up *Snell’s Law*). The proportions of reflected and refracted light are described by the Fresnel equations, and will be discussed in a later section.

Non-Optically-Flat Surfaces

Of course, most real-world surfaces are not polished to the same tolerances as telescope mirrors. What happens with surfaces that are not optically flat? In most cases, there are indeed irregularities present which are much larger than the light wavelength, but too small to be seen or resolved (i.e., they are smaller than the coverage area of a single pixel or shading sample). In this case, the surface behaves like a large collection of tiny optically flat surfaces. The surface appearance is the aggregate result of many points with different surface orientations—each point reflects incoming light in a slightly different direction (see Figure 13).




Figure 13: Visible reflections from non-optically flat surfaces are the aggregate result of reflections from many surface points with different orientations (*image from "Real-Time Rendering, 3rd edition" used with permission from A K Peters*).




Figure 14: On the top row, the surface is relatively smooth; the surface orientation only varies slightly, resulting in a small variance in reflected light directions and thus sharper reflections. The surface on the bottom row is rougher; different points on the surface have widely varying orientations, resulting in a high variance in reflected light directions and thus blurry reflections. Note that both surfaces appear smooth at the visible scale—the roughness difference is at the microscopic scale (*image from "Real-Time Rendering, 3rd edition" used with permission from A K Peters*).




Figure 15: When viewed macroscopically, non-optically flat surfaces can be treated as reflecting (and refracting) light in multiple directions (*image from "Real-Time Rendering, 3rd edition" used with permission from A K Peters*).

The rougher the surface is at this microscopic scale, the blurrier the reflections as the surface orientations diverge more strongly from the overall, macroscopic surface orientation (see Figure 14).

For shading purposes, it is common to treat this *microgeometry* statistically and view the surface as reflecting (and refracting) light in multiple directions (see Figure 15).

Subsurface Scattering

What happens to the refracted light? This depends on the composition of the object. Metals have very high absorption coefficients (imaginary part of refractive index) in the visible spectrum. All refracted light is immediately absorbed (soaked up by free electrons). On the other hand, non-metals (also referred to as dielectrics or insulators) behave as regular participating media once the light is refracted inside them, exhibiting the range of absorption and scattering behaviors we covered in previous sections. In most cases, some of the refracted light is scattered enough to be re-emitted out of the same surface. Both of these cases are illustrated in Figure 16.




Figure 16: In metals (on the left), all refracted light energy is immediately absorbed by free electrons; in non-metals (on the right) refracted light energy scatters until it re-emerges from the surface, typically after undergoing partial absorption (*images from "Real-Time Rendering, 3rd edition" used with permission from A K Peters*).

On the right side of Figure 16, you can see that the subsurface-scattered light (denoted with blue arrows) is emitted from various points on the surface, at varying distances from the original entrance point of the light. Figure 17 shows the relationship between these distances and the pixel size in two cases. On the upper left, the pixel is larger than the entry-to-exit subsurface scattering distances. In this case, the entry-to-exit distances can be ignored and the subsurface scattered light can be assumed to enter and exit the surface at the same point, as seen on the upper right. This allows shading to be handled as a completely local process; the outgoing light at a point only depends on incoming light at the same point. On the bottom of Figure 17, the pixel is smaller than the entry-to-exit distances. In this case, the shading of each point is affected by light impinging on other points. To capture this effect, local shading will not suffice and specialized rendering techniques need to be used. These are typically referred to as “subsurface scattering” techniques, but it is important to note that “ordinary” diffuse shading is the result of the same physical phenomena (subsurface scattering of refracted light). The only difference is the scattering distance relative to the scale of observation. This insight tells us that materials which are commonly thought of as exhibiting “subsurface scattering” behavior can be handled with regular diffuse shading at larger distances (e.g. the skin of a distant character) . On the other hand, materials which are thought of as exhibiting “regular diffuse shading” behavior will have a “subsurface scattering” appearance when viewed very close up (e.g. an extreme close-up of a small plastic toy).

The Mathematics of Shading

The measurement of electromagnetic radiation in general (including visible light) is called *radiometry*. There are various radiometric quantities used to measure light over surfaces, over directions, etc.; we




Figure 17: On the upper left, the pixel (green circle with red border) is larger than the distances traveled by the light before it exits the surface. In this case, the outgoing light can be assumed to be emitted from the entry point (upper right). On the bottom, the pixel is smaller than the scattering distances; these distances cannot be ignored if realistic shading is desired. (images from “*Real-Time Rendering, 3rd edition*” used with permission from A K Peters).

will only concern ourselves with *radiance*, which is used to quantify the magnitude of light along a single ray¹. We will use the common radiometric notation L to denote radiance; when shading a surface point, L_i denotes radiance incoming to the surface and L_o denotes outgoing radiance.

Radiance (like other radiometric quantities) is a *spectral quantity* - the amount varies as a function of wavelength. In theory, to express visible-light radiance a continuous spectral distribution needs to be stored. Dense spectral samples are indeed used in some specialized rendering applications, but for all production (film and game) rendering, RGB triples are used instead. An explanation of how these triples relate to spectral distributions can also be found in many websites and books, including *Real-Time Rendering* [20].

The BRDF

It is most commonly assumed that shading can be handled locally (as illustrated on the upper right of Figure 17). In this case, how a given surface point responds to light only depends on the incoming (light) and outgoing (view) directions. In this document, we will use \mathbf{v} to denote a unit-length vector pointing along the outgoing direction and \mathbf{l} to denote a unit-length vector pointing opposite to the incoming direction (it is convenient to have all vectors point away from the surface). The surface’s response to light is quantified by a function called the BRDF (Bidirectional Reflectance Distribution Function), which we will denote as $f(\mathbf{l}, \mathbf{v})$. Each direction (incoming and outgoing) can be parameterized with two numbers (e.g. polar coordinates), so the overall dimensionality of the BRDF is four. In many cases, rotating the light and view directions around the surface normal does not affect the BRDF. Such *isotropic BRDFs* can be parameterized with three angles (see Figure 18). In practice, the number of

¹An explanation of other radiometric quantities can be found in various texts, including Chapter 7 of the 3rd edition of *Real-Time Rendering* [20] and Dutré’s *Global Illumination Compendium* [10].

angles used to compute a given BRDF commonly varies from one to five—some commonly used angles are shown in Figure 19.




Figure 18: The BRDF depends on incoming and outgoing directions; these can be parameterized with four angles, or three in the case of isotropic BRDFs. Here \mathbf{n} is the surface normal vector, \mathbf{l} is the incoming light direction vector, \mathbf{v} is the outgoing (view) direction vector, and \mathbf{t} is a tangent vector defining a preferred direction over the surface (this is only used for *anisotropic BRDFs* where the reflection behavior changes when light and view vector are rotated around \mathbf{n}). (image from “Real-Time Rendering, 3rd edition” used with permission from A K Peters).




Figure 19: Examples of some angles which are commonly used in BRDF evaluation, in addition to those in Figure 18 (images from “Real-Time Rendering, 3rd edition” used with permission from A K Peters).

In principle, the BRDF is only defined for light and view directions above the surface; in other words, the dot products $(\mathbf{n} \cdot \mathbf{l})$ and $(\mathbf{n} \cdot \mathbf{v})$ must both be non-negative (recall that the dot product between two unit-length vectors is equal to the cosine of the angle between them; if this is negative, then the angle exceeds 90°). In production shading, situations arise when shading needs to be performed for angles outside this range (for example, normal mapping can result in normal vectors backfacing to the view vector). This is typically handled in practice by clamping the dot product to 0, but other approaches are possible [25].

The BRDF can be intuitively interpreted in two ways; both are valid. The first interpretation is that given a ray of light incoming from a certain direction, the BRDF gives the relative distribution of reflected and scattered light over all outgoing directions above the surface. The second interpretation is that for a given view direction, the BRDF gives the relative contribution of light from each incoming direction to the outgoing light. both interpretations are illustrated in Figure 20.




Figure 20: On the left side, we see one interpretation of the BRDF - that for a given outgoing (view) direction, it specifies the relative contributions of incoming light. On the right side we see an alternative interpretation - that for a given incoming light direction, it specifies the distribution of outgoing light. (image from “Real-Time Rendering, 3rd edition” used with permission from A K Peters).

The BRDF is a spectral quantity. In theory the input and output wavelengths would need to be additional BRDF inputs, increasing its dimensionality. However, in practice there is no cross-talk between the individual wavelengths²; each wavelength of outgoing light is only affected by that same wavelength in the incoming light. This means that instead of treating input and output wavelengths as BRDF inputs, we (more simply) treat the BRDF as a spectral-valued function that is multiplied with spectral-valued light colors. In production shading, this means an RGB-valued BRDF multiplied by RGB-valued light colors.

The BRDF is used in the *reflectance equation*³:

$$L_o(\mathbf{v}) = \int_{\Omega} f(\mathbf{l}, \mathbf{v}) \otimes L_i(\mathbf{l})(\mathbf{n} \cdot \mathbf{l}) d\omega_i. \quad (1)$$

Although this equation may seem a bit daunting, its meaning is straightforward: outgoing radiance equals the integral (over all directions above the surface) of incoming radiance times the BRDF and a cosine factor. If you are not familiar with integrals, you can think of them as a kind of continuous weighted average. The \otimes symbol is used here to denote component-wise vector multiplication; it is used because both BRDF and light color are spectral (RGB) vectors.

Not any arbitrary function over incoming and outgoing directions can make sense as a BRDF. It is commonly recognized that are two properties a BRDF must have to be *physically plausible*: *reciprocity* and *energy conservation*. Reciprocity simply means that the BRDF has the same value if \mathbf{l} and \mathbf{v} are swapped:

$$f(\mathbf{l}, \mathbf{v}) = f(\mathbf{v}, \mathbf{l}). \quad (2)$$

Energy conservation refers to the fact that a surface cannot reflect more than 100% of incoming light energy. Mathematically, it is expressed via the following equation:

$$\forall \mathbf{l}, \int_{\Omega} f(\mathbf{l}, \mathbf{v})(\mathbf{n} \cdot \mathbf{v}) d\omega_o \leq 1. \quad (3)$$

This means that for any possible light direction \mathbf{l} , the integral of the BRDF times a cosine factor over outgoing directions \mathbf{v} must not exceed 1.

The phenomena described by the BRDF includes (at least for non-metals) two distinct physical phenomena—surface reflection and subsurface scattering. Since each of these phenomena has different behavior, BRDFs typically include a separate term for each one. The BRDF term describing surface reflection is usually called the *specular term* and the term describing subsurface scattering is called the *diffuse term*; see Figure 21.

²There are two physical phenomena involving such crosstalk—fluorescence and phosphorescence; but they rarely occur in production shading.

³The reflectance equation is a special case of the *rendering equation* [16].




Figure 21: BRDF specular terms are typically used for surface reflection, and BRDF diffuse terms for subsurface scattering. (image from “Real-Time Rendering, 3rd edition” used with permission from A K Peters).

Surface Reflectance (Specular Term)

The basis for most physically-based specular BRDF terms is *microfacet theory*. This theory was developed to describe surface reflection from general (non-optically flat) surfaces. The basic assumption underlying microfacet theory is that the surface is composed of many *microfacets*, too small to be seen individually. Each microfacet is assumed to be optically flat. As mentioned in the previous section, an optically flat surface splits light into exactly two directions—reflection and refraction.




Figure 22: Microfacets with $\mathbf{m} = \mathbf{h}$ are oriented to reflect \mathbf{l} into \mathbf{v} —other microfacets do not contribute to the BRDF. (image from “Real-Time Rendering, 3rd edition” used with permission from A K Peters).




Figure 23: On the left, we see that some microfacets are occluded from the direction of \mathbf{l} , so they are shadowed and do not receive light (so they cannot reflect any). In the center, we see that some microfacets are not visible from the view direction \mathbf{v} , so of course any light reflected from them will not be seen. In both cases these microfacets do not contribute to the BRDF. In reality, shadowed light does not simply vanish; it continues to bounce from the microfacets and some of it does make its way into the view direction (as see on the right side). These *interreflections* are ignored by microfacet theory. (image from “Real-Time Rendering, 3rd edition” used with permission from A K Peters).

Each of these microfacets reflects light from a given incoming direction into a single outgoing direction which depends on the orientation of the microfacet normal \mathbf{m} . When evaluating a BRDF term, both the light direction \mathbf{l} and the view direction \mathbf{v} are specified. This means that of all the millions of microfacets on the surface, only those that happen to be angled just right to reflect \mathbf{l} into \mathbf{v} have any contribution to the BRDF value. In Figure 22 we can see that these *active microfacets* have

their surface normal \mathbf{m} oriented exactly halfway between \mathbf{l} and \mathbf{v} . The vector halfway between \mathbf{l} and \mathbf{v} is called the *half-vector* or *half-angle vector*; we will denote it as \mathbf{h} .

Not all microfacets for which $\mathbf{m} = \mathbf{h}$ will contribute to the reflection; some are blocked by other microfacets from the direction of \mathbf{l} (*shadowing*), from the direction of \mathbf{v} (*masking*), or from both. Microfacet theory assumes that all shadowed light is lost from the specular term; in reality, due to multiple surface reflections some of it will eventually be visible, but this is not accounted for in microfacet theory. This is not typically a major source of error in most cases (rough metal surfaces are a possible exception). The various types of light-microfacet interaction are shown in Figure 23.

With these assumptions (optically flat microfacets, no interreflections), a specular BRDF term can be derived from first principles ([1, 25]). The microfacet specular BRDF term has the following form⁴:

$$f_{\mu\text{facet}}(\mathbf{l}, \mathbf{v}) = \frac{F(\mathbf{l}, \mathbf{h})G(\mathbf{l}, \mathbf{v}, \mathbf{h})D(\mathbf{h})}{4(\mathbf{n} \cdot \mathbf{l})(\mathbf{n} \cdot \mathbf{v})} \quad (4)$$

We will go into each of the terms in more detail, but first a quick summary. $F(\mathbf{l}, \mathbf{h})$ is the Fresnel reflectance of the active microfacets as a function of the light direction \mathbf{l} and the active microfacet normal $\mathbf{m} = \mathbf{h}$. $G(\mathbf{l}, \mathbf{v}, \mathbf{h})$ is the proportion of microfacets (of the ones with $\mathbf{m} = \mathbf{h}$) which are *not* shadowed or masked, as a function of the light direction \mathbf{l} , the view direction \mathbf{v} , and the active microfacet normal $\mathbf{m} = \mathbf{h}$. $D(\mathbf{h})$ is the microfacet normal distribution function evaluated at the active microfacet normal $\mathbf{m} = \mathbf{h}$; in other words, the concentration of microfacets with normals equal to \mathbf{h} . Finally, the denominator $4(\mathbf{n} \cdot \mathbf{l})(\mathbf{n} \cdot \mathbf{v})$ is a correction factor which accounts for quantities being transformed between the local space of the microfacets and that of the overall macrosurface.

Fresnel Reflectance Term

The Fresnel reflectance term computes the fraction of light reflected from an optically flat surface. Its value depends on two things: the incoming angle (angle between light vector and surface normal) and the refractive index of the material. Since the refractive index may vary over the visible spectrum, the Fresnel reflectance is a spectral quantity—for production purposes, an RGB triple. We also know that each of the RGB values have to lie within the 0 to 1 range, since a surface cannot reflect less than 0% or more than 100% of the incoming light. Since we are only concerned with active microfacets for which $\mathbf{m} = \mathbf{h}$, the incidence angle for Fresnel reflectance is actually the one between \mathbf{l} and \mathbf{h} .

The full Fresnel equations are somewhat complex, and the required material parameter (complex refractive index sampled densely over the visible spectrum) is not particularly convenient for artists (to say the least). However, a simpler expression with more convenient parametrization can be derived by inspecting the behavior of these equations for real-world materials. With this in mind, let us inspect the graph in Figure 24.

The materials selected for this graph represent a wide variety. Despite this, some common elements can be seen. Reflectance is almost constant for incoming angles between 0° and about 45° . The reflectance changes more significantly (typically, but not always, increasing somewhat) between 45° and about 75° . Finally, between 75° and 90° reflectance always goes rapidly to 1 (white, if viewed as an RGB triple). Since the Fresnel reflectance stays close to the value for 0° over most of the range, we can think of this value $F(0^\circ)$ as the *characteristic specular reflectance* of the material. This value has all the properties of what is typically thought of as a “color”—it is composed of RGB values between 0 and 1, and it is a measure of selective reflectance of light. For this reason, we will also refer to this value as the *specular color* of the surface, denoted as \mathbf{c}_{spec} .

⁴Note that cases where one or both of the dot products in the denominator are negative or zero need to be handled, although in theory this is outside the domain over which the BRDF is defined. In practice, this is handled by clamping the dot products to a very small positive value, though some authors [25] recommend using absolute value instead of clamping.




Figure 24: Fresnel reflectance for external reflection from a variety of substances. Since copper and aluminum have significant variation in their reflectance over the visible spectrum, their reflectance is shown as three separate curves for R, G, and B. Copper’s R curve is highest, followed by G, and finally B (thus its reddish color). Aluminum’s B curve is highest, followed by G, and finally R. (*image from “Real-Time Rendering, 3rd edition” used with permission from A K Peters*).

\mathbf{c}_{spec} looks like an ideal parameter for a Fresnel reflectance approximation, and indeed Schlick [22] gives a cheap and reasonably accurate approximation that uses it:

$$F_{\text{Schlick}}(\mathbf{c}_{\text{spec}}, \mathbf{l}, \mathbf{n}) = \mathbf{c}_{\text{spec}} + (1 - \mathbf{c}_{\text{spec}})(1 - (\mathbf{l} \cdot \mathbf{n}))^5 \quad (5)$$

This approximation is widely used in computer graphics. In the special case of active microfacets, \mathbf{h} must be substituted for the surface normal \mathbf{n} :

$$F_{\text{Schlick}}(\mathbf{c}_{\text{spec}}, \mathbf{l}, \mathbf{h}) = \mathbf{c}_{\text{spec}} + (1 - \mathbf{c}_{\text{spec}})(1 - (\mathbf{l} \cdot \mathbf{h}))^5 \quad (6)$$

To know which values are reasonable to assign to \mathbf{c}_{spec} , it is instructive to look at the values of $F(0^\circ)$ for various real-world materials. These can be found in Table 1. Values are given in both linear and gamma (sRGB) space; we recommend anyone unfamiliar with the importance of computing shading in linear space and the issues involved in converting input from gamma space consult some of the articles on the topic ([14, 15]).











Material	$F(0^\circ)$ (Linear)	$F(0^\circ)$ (sRGB)	Color
Water	0.02,0.02,0.02	0.15,0.15,0.15	
Plastic / Glass (Low)	0.03,0.03,0.03	0.21,0.21,0.21	
Plastic High	0.05,0.05,0.05	0.24,0.24,0.24	
Glass (High) / Ruby	0.08,0.08,0.08	0.31,0.31,0.31	
Diamond	0.17,0.17,0.17	0.45,0.45,0.45	
Iron	0.56,0.57,0.58	0.77,0.78,0.78	
Copper	0.95,0.64,0.54	0.98,0.82,0.76	
Gold	1.00,0.71,0.29	1.00,0.86,0.57	
Aluminum	0.91,0.92,0.92	0.96,0.96,0.97	
Silver	0.95,0.93,0.88	0.98,0.97,0.95	

Table 1: Values of $F(0^\circ)$ for various materials. (*table from “Real-Time Rendering, 3rd edition” used with permission from A K Peters*).

When inspecting Table 1, several things stand out. One is that metals have significantly higher values of $F(0^\circ)$ than non-metals. Iron is a very dark metal, and it reflects more than 50% of incoming light at 0° . Recall that metals have no sub-surface reflectance; a bright specular color and no diffuse color is the distinguishing visual characteristic of metals. On the other hand diamond, one of the brightest non-metals, reflects only 17% of incoming light at 0° ; most non-metals reflect significantly less than that. Very few materials have values in the “no mans land” between 20% and 40%; these are typically semiconductors and other exotic materials which are unlikely to appear in production shading situations. The same is true for values lower than 2% (the $F(0^\circ)$ value of water). In fact, ruling out metals, gemstones, and crystals, pretty much any material you are likely to see outside a laboratory will have a narrow range of $F(0^\circ)$ values—between 2% and 5%.

Normal Distribution Term

In most surfaces, the microfacet’s orientations are not uniformly distributed. Microfacet normals closer to the macroscopic surface normal tend to appear with higher frequency. The exact distribution is defined via the *microfacet normal distribution function* $D(\mathbf{m})$. Unlike $F()$, the value of $D()$ is not restricted to lie between 0 and 1—although values must be non-negative, they can be arbitrarily large. Also unlike $F()$, the function $D()$ is not spectral or RGB-valued, but scalar-valued. In microfacet BRDF terms, $D()$ is evaluated for the direction \mathbf{h} , to help determine the concentration of active microfacets (those for which $\mathbf{m} = \mathbf{h}$). This is why the normal distribution term appears in Equation 4 as $D(\mathbf{h})$.

The function $D()$ determines the size, brightness, and shape of the specular highlight. Several different normal distribution functions appear in the graphics literature, all are somewhat Gaussian-like, with some kind of “roughness” or variance parameter (anisotropic functions typically have two variance parameters). As the surface roughness decreases, the concentration of the microfacet normals \mathbf{m} around the overall surface normal \mathbf{n} increases, and the values of $D(\mathbf{m})$ can become very high (in the limit, for a perfect mirror, the value is infinity at $\mathbf{m} = \mathbf{n}$). Walter et. al. [25] discuss the correct normalization of the distribution function, and give several examples; more examples can be found in other papers [2, 3, 19, 26].

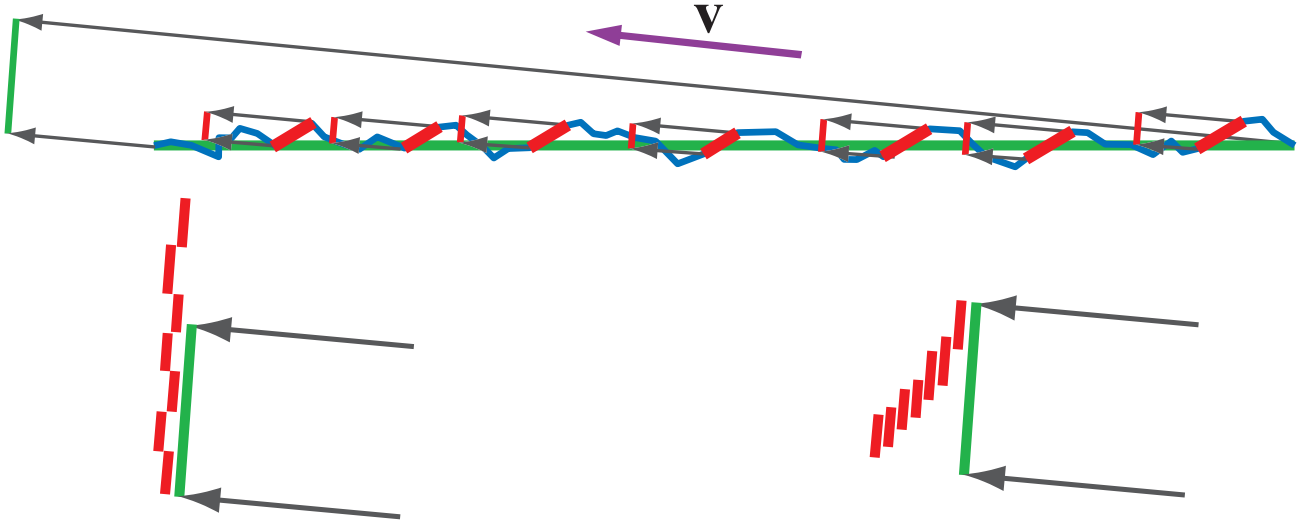


Figure 25: On the top the flat macroscopic surface is shown in green, and the rugged microscopic surface is shown in blue. The facets for which $\mathbf{m} = \mathbf{h}$ are marked in red. The projection of the macroscopic surface area (length in this 2D side illustration) onto the view direction (in other words, its foreshortened surface area) is shown as a green line on the upper left. The projected areas of the individual red microfacets are shown as separate red lines. On the bottom left the areas of the red microfacets are added up without accounting for masking, resulting in an active area greater than the total area. This is illogical, and more importantly can result in the BRDF reflecting more energy than it receives. On the right we see that the red areas are combined in a way that accounts for masking. The overlapping areas are no longer counted multiple times, and we see that the correct active area is smaller than the total area. When the viewing angle is lower, then the effect will be even more pronounced—ignoring the effects of masking could lead to the BRDF reflecting thousands of times the amount of energy received or more (the amount of reflected energy would go to infinity in the limit as the angle goes to 90°).

Shadowing-Masking Term

The shadowing and masking term $G(\mathbf{l}, \mathbf{v}, \mathbf{h})$ is also often called the *geometry term* in the BRDF literature. The function $G(\mathbf{l}, \mathbf{v}, \mathbf{m})$ represents the probability that microfacets with a given normal \mathbf{m} will be visible from both the light direction \mathbf{l} and the view direction \mathbf{v} . In the microfacet BRDF, \mathbf{m} is replaced with \mathbf{h} (for similar reasons as in the previous two terms). Since the function $G()$ represents a probability, its values are scalars and are constrained to lie between 0 and 1. As in the case of $D()$, there are various analytical expressions for $G()$ in the literature [2, 3, 7, 8, 17, 19, 25]; these are typically approximations based on some simplified model of the surface. The $G()$ function typically does not introduce any new parameters to the BRDF; it either has no parameters, or uses the roughness parameters of the $D()$ function. In many cases, the shadowing-masking term partially cancels out the $(\mathbf{n} \cdot \mathbf{l})(\mathbf{n} \cdot \mathbf{v})$ denominator in Equation 4, replacing it with some other expression such as $\max(\mathbf{n} \cdot \mathbf{l}, \mathbf{n} \cdot \mathbf{v})$.

The shadowing-masking term is essential for BRDF energy conservation—without such a term the BRDF can reflect arbitrarily more light energy than it receives. A key part of the microfacet BRDF is the ratio between the active area (combined area of the microfacets which reflect light energy from \mathbf{l} to \mathbf{v}) and the total area (of the macroscopic surface). If shadowing and masking are not accounted for, then the active area may exceed the total area, an obvious impossibility which can lead to the BRDF not conserving energy, in some cases by a huge amount (see Figure 25).

Microfacet Models

The choice of $D()$ and $G()$ functions is independent; they can be mixed and matched from different microfacet models. Most papers proposing a new microfacet BRDF model are best understood as introducing a new $D()$ and / or $G()$ function.

Once the $D()$ and $G()$ functions have been chosen, the full BRDF is determined by selecting parameter values. Microfacet BRDFs have compact parameterizations, typically only consisting of one RGB value for (\mathbf{c}_{spec}) and one scalar for roughness (two in the case of anisotropic BRDFs).

Subsurface Reflectance (Diffuse Term)

Although there are several models for subsurface local reflection in the literature, the most widely-used one by far is the *Lambertian* BRDF term. The Lambertian BRDF is actually a constant value; the well-known cosine or ($\mathbf{n} \cdot \mathbf{l}$) factor is part of the reflection equation, not the BRDF (as we saw in Equation 1). The exact value of the Lambertian BRDF is:

$$f_{\text{Lambert}}(\mathbf{l}, \mathbf{v}) = \frac{\mathbf{c}_{\text{diff}}}{\pi}. \quad (7)$$

Here \mathbf{c}_{diff} is the fraction of light which is diffusely reflected. As in the case of \mathbf{c}_{spec} , it is an RGB value with R, G, and B restricted to the 0 – 1 range, and corresponds closely to what most people think of as a “surface color”. This parameter is typically referred to as the *diffuse color*.

Non-Lambertian diffuse terms attempt to model either the trade-off between specular and diffuse terms at glancing angles [2, 3, 17, 24], or the effects of surface roughness at a scale larger than the scattering distance [21].

Implementing Shading

In the previous section, we saw the mathematical models that are typically employed to describe surface shading. In this section, we will discuss how such models are implemented in film and game production renderers.

General Lighting

In the most general case, the BRDF must be integrated against incoming light from all different directions. This includes not only primary light sources (with area) but also skylight and accurate reflections of other objects in the scene. To fully solve this, global illumination algorithms are required. Detailed descriptions of these algorithms are outside the domain of this talk; more details can be found in various references ([18, 11]), as well as Adam Martinez’s talk in this course, *Faster Photorealism in Wonderland: Physically-Based Shading and Lighting at Sony Pictures Imageworks*.

Punctual Light Sources

A far more restricted, but common production lighting environment is comprised of one or more *punctual light sources*. These are the classic computer graphics point, directional, and spot lights (more complex variants are also used [4]). Since they are infinitely small and infinitely bright, they aren’t physically realizable or realistic, but they do produce reasonable results in many cases and are computationally convenient. Punctual light sources are parameterized by the light color $\mathbf{c}_{\text{light}}$ and the light direction vector $\mathbf{l}_{\mathbf{c}}$. For artist convenience, $\mathbf{c}_{\text{light}}$ does not correspond to a direct radiometric measure of the light’s intensity; it is specified as the color a white Lambertian surface would have

when illuminated by the light from a direction parallel to the surface normal ($\mathbf{l}_c = \mathbf{n}$). Like other color quantities we have seen, $\mathbf{c}_{\text{light}}$ is spectral (RGB)-valued, but unlike them its range is unbounded.

The primary advantage of punctual light sources is that they greatly simplify the reflection equation (Equation 1), as we will show here. We will start by defining a tiny area light source centered on \mathbf{l}_c , with a small angular extent ε . This tiny area light illuminates a shaded surface point with the incoming radiance function $L_{\text{tiny}}(\mathbf{l})$. The incoming radiance function has the following two properties:

$$\forall \mathbf{l} \angle(\mathbf{l}, \mathbf{l}_c) > \varepsilon, L_{\text{tiny}}(\mathbf{l}) = 0. \quad (8)$$

$$\text{if } \mathbf{l}_c = \mathbf{n}, \text{ then } \mathbf{c}_{\text{light}} = \frac{1}{\pi} \int_{\Omega} L_{\text{tiny}}(\mathbf{l})(\mathbf{n} \cdot \mathbf{l}) d\omega_i. \quad (9)$$

The first property says that no light is incoming for any light directions which form an angle greater than ε with \mathbf{l}_c . In other words, the light does not produce any light outside its angular extent of ε . The second property follows from the definition of $\mathbf{c}_{\text{light}}$, applying Equations 1 and 7 with $\mathbf{c}_{\text{diff}} = 1$. Equation 9 still holds in the limit as ε goes to 0:

$$\text{if } \mathbf{l}_c = \mathbf{n}, \text{ then } \mathbf{c}_{\text{light}} = \lim_{\varepsilon \rightarrow 0} \left(\frac{1}{\pi} \int_{\Omega} L_{\text{tiny}}(\mathbf{l})(\mathbf{n} \cdot \mathbf{l}) d\omega_i \right). \quad (10)$$

Since $\mathbf{l}_c = \mathbf{n}$ and $\varepsilon \rightarrow 0$, we can assume $(\mathbf{n} \cdot \mathbf{l}) = 1$ which gives us:

$$\mathbf{c}_{\text{light}} = \lim_{\varepsilon \rightarrow 0} \left(\frac{1}{\pi} \int_{\Omega} L_{\text{tiny}}(\mathbf{l}) d\omega_i \right). \quad (11)$$

Note that Equation 11 is independent of the value of \mathbf{l}_c , so it is true for any valid light orientation, not just $\mathbf{l}_c = \mathbf{n}$. Simple rearrangement isolates the value of the integral in the limit:

$$\lim_{\varepsilon \rightarrow 0} \left(\int_{\Omega} L_{\text{tiny}}(\mathbf{l}) d\omega_i \right) = \pi \mathbf{c}_{\text{light}}. \quad (12)$$

Now we shall apply our tiny area light to a general BRDF, and look at its behavior in the limit as ε goes to 0:

$$L_o(\mathbf{v}) = \lim_{\varepsilon \rightarrow 0} \left(\int_{\Omega} f(\mathbf{l}, \mathbf{v}) \otimes L_{\text{tiny}}(\mathbf{l})(\mathbf{n} \cdot \mathbf{l}) d\omega_i \right) = f(\mathbf{l}_c, \mathbf{v}) \otimes \lim_{\varepsilon \rightarrow 0} \left(\int_{\Omega} L_{\text{tiny}}(\mathbf{l}) d\omega_i \right) (\mathbf{n} \cdot \mathbf{l}_c). \quad (13)$$

Substituting Equation 12 into the right part of Equation 13 gives us the final punctual light equation:

$$L_o(\mathbf{v}) = \pi f(\mathbf{l}_c, \mathbf{v}) \otimes \mathbf{c}_{\text{light}}(\underline{\mathbf{n} \cdot \mathbf{l}_c}). \quad (14)$$

Compared to the original reflectance equation, we have replaced the integral with a straightforward evaluation of the BRDF, which is much simpler to compute. Note the line under the dot product ($\underline{\mathbf{n} \cdot \mathbf{l}_c}$); this is our notation for clamping to 0. In other words, $\underline{x} \equiv \max(x, 0)$. The dot product is clamped to handle punctual lights which are behind the surface; these should have no contribution (rather than a negative contribution).

In the case of directional light sources (such as the Sun), both \mathbf{l}_c and $\mathbf{c}_{\text{light}}$ are constant over the scene. In the case of other punctual light types such as point lights and spotlights, both will vary. In reality, $\mathbf{c}_{\text{light}}$ would fall off proportionally to the inverse square distance, but in practice other falloff functions are often used.

If multiple punctual light sources are illuminating the surface, Equation 14 is computed multiple times and the results summed. Punctual light sources are rarely used by themselves, since the lack of any illumination coming from other directions is noticeable, especially with highly specular surfaces. For this reason punctual light sources are typically combined with some kind of ambient or environmental lighting; these types of lighting will be discussed below.

Ambient Lighting

Here we define ambient lighting as some numerical representation of low-frequency lighting, ranging from a single constant light color and intensity over all incoming directions to more complex representations such as spherical harmonics (SH). Often this type of lighting environments is only applied to the diffuse BRDF term; more high-frequency image-based lighting are applied to the specular term. However, it is possible to apply ambient lighting environments to the specular BRDF term. Yoshiharu Gotanda's talk in this course, *Practical Implementation of Physically-Based Shading Models at tri-Ace* gives a specular implementation for constant and SH ambient, a recent presentation by Bungie [5] discusses applying the Cook-Torrance [7, 8] specular term to SH lighting, and a ShaderX⁷ article by Schüler [23] describes an implementation of a physically-based specular term with hemispherical lighting.

Image-Based Lighting

Image-based lighting is typically done with environment maps. These maps represent distant lighting. If they are sampled at an appropriate reference position (say, at the center of an object) they can be a very good representation of reflections from distant objects. To correctly handle local shading with a general BRDF and an environment map⁵, many samples are required. Importance sampling helps to keep the number of samples to a somewhat more manageable number (at least for film rendering). Another approach that can be used, either by itself (an approximate solution, but suitable for games) or in combination with importance sampling, is environment map prefiltering. More information on importance sampling can be found at another course this year [6], as well as as Adam Martinez's talk in this course. Other aspects of image-based lighting in film production are discussed in Ben Snow's talk in this course, *Terminators and Iron Men: Image-Based Lighting and Physical Shading at ILM*, and some aspects of shading with environment maps for video games are discussed in Naty Hoffman's other talk, *Crafting Physically Motivated Shading Models for Game Development*.

Further Reading

Chapter 7 of the 3rd edition of "Real-Time Rendering" [20] provides a broad overview of physically-based shading models, going into somewhat more depth than these course notes. For even greater depth, consider reading Glassner's *Principles of Digital Image Synthesis* [12, 13], or *Digital Modeling of Material Appearance* [9] by Dorsey, Rushmeier, and Sillion.

Dutr e's free online *Global Illumination Compendium* [10] is a useful reference for BRDFs, radiometric math, and much else.

Acknowledgments

The author would like to thank A K Peters for permission to use images from the book *Real-Time Rendering, 3rd edition*, and also Paul Edelstein, Yoshiharu Gotanda and Dimitar Lazarov for many thought-provoking discussions on physically-based shading models.

⁵global illumination effects such as interreflections on an object can also be handled, but in that case you are effectively using the environment map as a light source for a global illumination renderer.

Bibliography

- [1] Ashikhmin, Michael, Simon Premože, and Peter Shirley, “A Microfacet-Based BRDF Generator,” *Computer Graphics (SIGGRAPH 2000 Proceedings)*, pp. 67–74, July 2000. <http://www.cs.utah.edu/~shirley/papers/facets.pdf> Cited on p. 13
- [2] Ashikhmin, Michael, and Peter Shirley, “An Anisotropic Phong Light Reflection Model,” Technical Report UUCS-00-014, Computer Science Department, University of Utah, June 2000. www.cs.utah.edu/research/techreports/2000/pdf/UUCS-00-014.pdf Cited on p. 15, 16, 17
- [3] Ashikhmin, Michael, Simon Premože, and Peter Shirley, “An Anisotropic Phong BRDF Model,” *journal of graphics tools*, vol. 5, no. 2, pp. 25–32, 2000. www.cs.utah.edu/~shirley/papers/jgtbrdf.pdf Cited on p. 15, 16, 17
- [4] Barzel, Ronen, “Lighting Controls for Computer Cinematography” *journal of graphics tools*, vol. 2, no. 1, pp. 1–20, 1997. Cited on p. 17
- [5] Chen, Hao, “Lighting and Material of Halo 3,” *Game Developers Conference*, March 2008. http://www.bungie.net/images/Inside/publications/presentations/lighting_material.zip Cited on p. 19
- [6] Colbert, Mark, Simon Premože, and Guillaume François, “Importance Sampling for Production Rendering,” *SIGGRAPH 2010 Course Notes*, 2010. <http://sites.google.com/site/isrendering/> Cited on p. 19
- [7] Cook, Robert L., and Kenneth E. Torrance, “A Reflectance Model for Computer Graphics,” *Computer Graphics (SIGGRAPH '81 Proceedings)*, pp. 307–316, July 1981. Cited on p. 16, 19
- [8] Cook, Robert L., and Kenneth E. Torrance, “A Reflectance Model for Computer Graphics,” *ACM Transactions on Graphics*, vol. 1, no. 1, pp. 7–24, January 1982. <http://graphics.pixar.com/library/ReflectanceModel/> Cited on p. 16, 19
- [9] Dorsey, Julie, Holly Rushmeier, and François Sillion, *Digital Modeling of Material Appearance*, Morgan Kaufmann, 2007. Cited on p. 19
- [10] Dutré, Philip, *Global Illumination Compendium*, 1999. <http://www.graphics.cornell.edu/~phil/GI> Cited on p. 9, 19
- [11] Dutré, Philip, Kavita Bala, and Philippe Bekaert, *Advanced Global Illumination*, second edition, A K Peters Ltd., 2006. Cited on p. 17
- [12] Glassner, Andrew S., *Principles of Digital Image Synthesis*, vol. 1, Morgan Kaufmann, 1995. Cited on p. 19
- [13] Glassner, Andrew S., *Principles of Digital Image Synthesis*, vol. 2, Morgan Kaufmann, 1995. Cited on p. 19
- [14] Gritz, Larry, and Eugene d’Eon, “The Importance of Being Linear,” in Hubert Nguyen, ed., *GPU Gems 3*, Addison-Wesley, pp. 529–542, 2007. http://http.developer.nvidia.com/GPUGems3/gpugems3_ch24.html Cited on p. 14
- [15] Hoffman, Naty, “Adventures with Gamma-Correct Rendering,” Renderwonk blog. <http://renderwonk.com/blog/index.php/archive/adventures-with-gamma-correct-rendering/> Cited on p. 14
- [16] Kajiyama, James T., “The Rendering Equation,” *Computer Graphics (SIGGRAPH '86 Proceedings)*, pp. 143–150, August 1986. <http://www.cs.brown.edu/courses/cs224/papers/kajiyama.pdf> Cited on p. 11

- [17] Kelemen, Csaba, and László Szirmay-Kalos, “A Microfacet Based Coupled Specular-Matte BRDF Model with Importance Sampling,” *Eurographics 2001*, short presentation, pp. 25–34, September 2001. http://www.fsz.bme.hu/~szirmay/scook_link.htm Cited on p. 16, 17
- [18] Křivánek, Jaroslav, Marcos Fajardo, Per H. Christensen, Eric Tabellion, Michael Bunnell, David Larsson, and Anton Kaplanyan, “Global Illumination Across Industries,” *SIGGRAPH 2010 Course Notes*, 2010. <http://www.graphics.cornell.edu/~jaroslav/gicourse2010/> Cited on p. 17
- [19] Kurt, Murat, László Szirmay-Kalos, and Jaroslav Křivánek, “An Anisotropic BRDF Model for Fitting and Monte Carlo Rendering,” *Computer Graphics*, vol. 44, no. 1, pp. 1–15, 2010. <http://www.graphics.cornell.edu/~jaroslav/> Cited on p. 15, 16
- [20] Akenine-Möller, Tomas, Eric Haines, and Naty Hoffman, *Real-Time Rendering*, third edition, A K Peters Ltd., 2008. <http://realtimerendering.com/> Cited on p. 9, 19
- [21] Oren, Michael, and Shree K. Nayar, “Generalization of Lambert’s Reflectance Model,” *Computer Graphics (SIGGRAPH 94 Proceedings)*, pp. 239–246, July 1994. <http://www.cs.columbia.edu/CAVE/projects/oren/> Cited on p. 17
- [22] Schlick, Christophe, “An Inexpensive BDRF Model for Physically based Rendering,” *Computer Graphics Forum*, vol. 13, no. 3, Sept. 1994, pp. 149–162. <http://dept-info.labri.u-bordeaux.fr/~schlick/DOC/eur2.html> Cited on p. 14
- [23] Schüler, Christian, “An Efficient and Physically Plausible Real Time Shading Model,” in Wolfgang Engel, ed., *ShaderX⁷*, Charles River Media, pp. 175–187, 2009. Cited on p. 19
- [24] Shirley, Peter, Helen Hu, Brian Smits, Eric Lafortune, “A Practitioners’ Assessment of Light Reflection Models,” *Pacific Graphics '97*, pp. 40–49, October 1997. <http://www.graphics.cornell.edu/pubs/1997/SHSL97.html> Cited on p. 17
- [25] Walter, Bruce, Stephen R. Marschner, Hongsong Li, Kenneth E. Torrance, “Microfacet Models for Refraction through Rough Surfaces,” *Eurographics Symposium on Rendering (2007)*, 195–206, June 2007. <http://www.cs.cornell.edu/~srm/publications/EGSR07-btdf.html> Cited on p. 10, 13, 15, 16
- [26] Ward, Gregory, “Measuring and Modeling Anisotropic Reflection,” *Computer Graphics (SIGGRAPH '92 Proceedings)*, pp. 265–272, July 1992. <http://radsite.lbl.gov/radiance/papers/sg92/paper.html> Cited on p. 15