Physically-Based Reflectance for Games

8:40 - 9:15: Reflectance

Naty Hoffman

Reflectance

- Types of Reflectance
- Reflectance Theory
- Reflection Model Foundations

In this section, we will first discuss various types of reflectance from a more intuitive, visual or qualitative standpoint. Afterwards we will go into the quantitative aspects of reflection modeling – first covering the theory behind reflection models and then discussing the building blocks from which most reflection models are built.

Types of Reflectance

- Mirror Surfaces
  - Substances
  - Surface and Body Reflectance
- Non-Mirror Surfaces
  - Surface Roughness and Scale
  - Microscopically Rough Surfaces
  - Structured Surfaces

When discussing types of reflective surfaces, we will first discuss mirror (effectively perfectly smooth) surfaces, and then other kinds of surfaces.

When discussing mirror surfaces we will discuss how various substances (metals, homogeneous dielectrics, inhomogeneous dielectrics) reflect light differently, and the difference between surface and body (or volume) reflectance.

Mirror Surfaces

- Perfectly smooth surface
- Incident ray of light only reflected in one direction

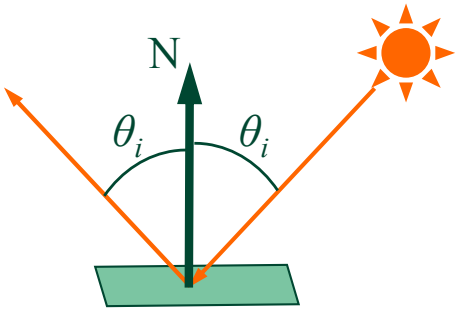

SIGGRAPH2006

The geometrically simplest surface is a perfectly smooth, or mirror surface. In this case a ray of incident (incoming) light is reflected in a single direction, which is the direction of incidence mirrored about the surface normal N.

Perfectly Smooth?

- Of course, in nature no surface can actually be perfectly smooth
- However, irregularities significantly smaller than visible light wavelengths (~400-700nm) can be ignored

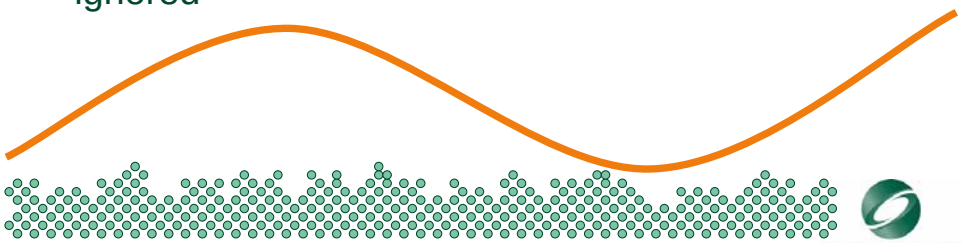

SIGGRAPH2006

Perfect smoothness is an unrealizable abstraction – if nothing else the individual atoms of the surface form ‘bumps’. However, due to the fact that visible light has wavelengths in a well-defined range (roughly 400 to 700 nanometers), any surface which only has irregularities below this scale can be considered ‘perfectly smooth’ for the purpose of reflecting visible light. Note that a surface which can be treated as ‘perfectly smooth’ for reflecting visible light may not necessarily be treated so for reflecting shorter wavelength electromagnetic radiation such as ultraviolet or X-Rays.

Mirror Surfaces

- Small, bright lights reflect as tiny bright spots
- Fine environment details can be seen in reflection

SIGGRAPH2006

What are the main visual characteristics of a mirror surface? Individual small, bright light sources are reflected in much the same way they would appear if we were to look at them directly, namely as small, extremely intense spots. In general, the environment is reflected in a recognizable manner, even fine detail can be seen in the reflection.

Substances

- Mirrors are differentiated by their substance
- We will distinguish between two main types:
  - Metals, or conductors
  - Dielectrics, or insulators
- Semiconductors are a possible third type
  - Not often seen in games
  - Visually are between metals and dielectrics

Lacking any surface structure, mirrors differ only by the substance of which they are composed. For reflectance purposes, it is useful to classify such substances into two broad groups: metals (conductors) and dielectrics (insulators), since their visual properties differ significantly. Semiconductors form a third class, which we will not discuss. They do not often appear in bulk form in game scenes, and their visual properties are often in between those of metals and dielectrics and thus can be interpolated from an understanding of both.

Reflectance

- Not all the light is reflected
  - Some is refracted under the surface
- The *Reflectance* is the ratio of reflected to incident (incoming) light
  - Must be between 0 and 1
- Varies by light wavelength
  - In games, this means it is an RGB triple
  - RGB, 0 to 1 range: is essentially a color

$$\text{Reflectance} = \frac{\text{reflected}}{\text{incident}}$$

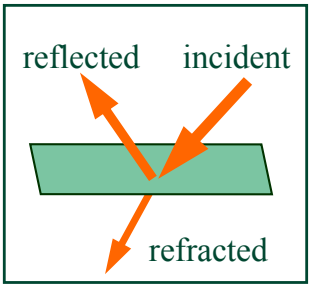

SIGGRAPH2006

Since we ignore glowing or emissive materials in this course, the reflectance must always be between 0 (none of the incident light is reflected) and 1 (all of it is reflected). Rendering systems that stress physical fidelity over performance represent spectral values (those that vary by light wavelength) as a series of dozens of spectral samples. Most production rendering systems (and all real-time rendering systems) represent such values as RGB triples (the details of the relationship between spectral values and RGB triples is outside the scope of this course).

Since a reflectance value is an RGB triple of values, each of which must be between 0 and 1, we can think of it as a color.

Reflectance

- Reflectance can also be thought of as an RGB multiplier to the incident light to produce the reflected light

reflected = Reflectance x incident

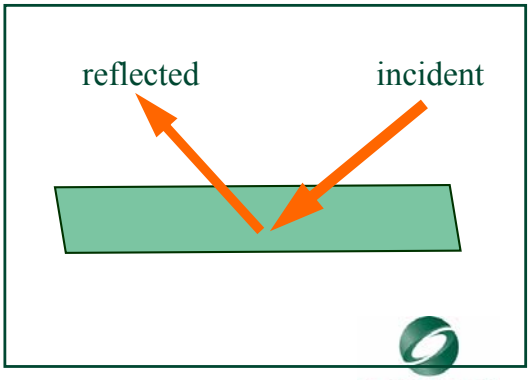

Fresnel Equations

- Surface reflectance depends on
  - Incidence angle (goes to 100% at glancing angles)
  - Refractive index (spectral)
  - ~~Polarization~~
- Surface reflectance at normal incidence is a substance characteristic

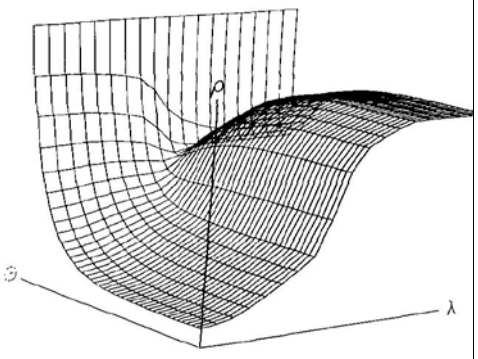

IMAGES BY R. COOK AND K. TORRANCE

For mirror surfaces, the surface reflectance (note the qualification of “reflectance” with “surface” in this case, which shall become clearer in a few slides) obeys the Fresnel equations.

The Fresnel equations depend on the angle of incidence (angle between the incoming light and the surface normal), the refractive index of the substance (which is a spectral quantity – it varies by light wavelength), and by polarization (which we ignore in this course – we assume light is non-polarized).

Here as an example, we see a 3D graph of the surface reflectance of copper as a function of incidence angle and wavelength (over the visible spectrum). In addition, we see the spectral reflectance or reflectance color as a function of incidence angle (increasing to the right) on the color strip below the graph.

As the incidence angle increases (going to more glancing angles), the reflectance increases until (at the limit) it is 100% at all wavelengths (this is true for all substances, not just copper). Note that although the general trend is increasing reflectance with increasing angle of incidence, the increase is not monotonic across all wavelengths. There is a dip on the red side before it goes up (causing the color to shift to blue just before it goes white).

Most production rendering (including film) ignores polarization, but it can be visually significant in some cases (skylight is partially polarized).

All smooth surfaces have a 100% white surface reflectance at glancing angles, but the surface reflectance at normal incidence is a function of the refractive index and is a characteristic of the material.

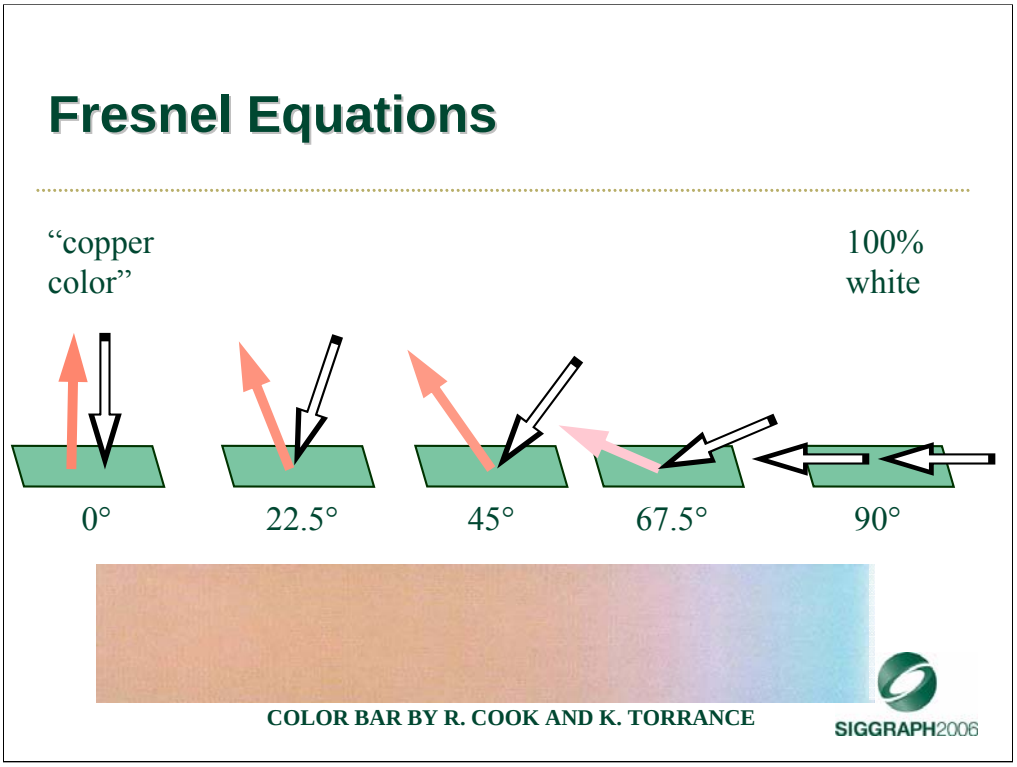

The salient fact is that the surface reflectance goes from a color which is characteristic of the substance (copper in this case) to 100% white as the light goes from hitting at normal incidence to increasingly glancing angles. The change is not linear - at first it is very gradual, then at extreme glancing angles it accelerates.

Metals

- Transmitted light absorbed; converted to heat
- High (>50%) normal-incidence surface reflectance
- Example spectral reflectances
  - Steel: 55% grey
  - Silver: 96% white
  - Gold: High-value (>0.8) yellow
  - Copper: High-value (>0.7) orange-red

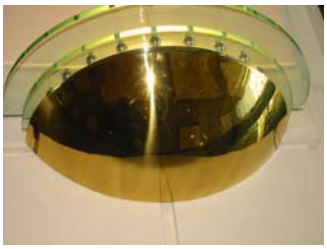

The first class of substances we will discuss are metals. Metals are both highly reflective and highly absorbent of visible light; light not reflected is absorbed quickly by the ‘sea’ of free electrons in the metal, its energy converted to heat.

Metals tend to sharply decrease in reflectance with decreasing wavelength (from red to blue) after a certain point. For most metals this point occurs above the visible range, so their reflectance is roughly constant over this range, leading to a characteristic “colorless” appearance. For other metals such as copper and gold this point occurs near the blue end of the visible range, leading to their characteristic “yellowish” or “reddish” appearance. This will color the reflections in the material.

Dielectrics

- Low (rarely above 5%) normal-incidence surface reflectance

Again, note that like other reflectance numbers, these are at normal incidence – at glancing angles it is significantly higher. Most dielectrics have a colorless surface reflection, caused by the fact that their surface reflectance does not significantly vary over the visible spectrum. Note that here again we qualify the reflectance as "surface reflectance", this will become clear in a few slides.

Homogeneous Dielectrics

- Low, colorless reflectance and low absorption
- Transparent

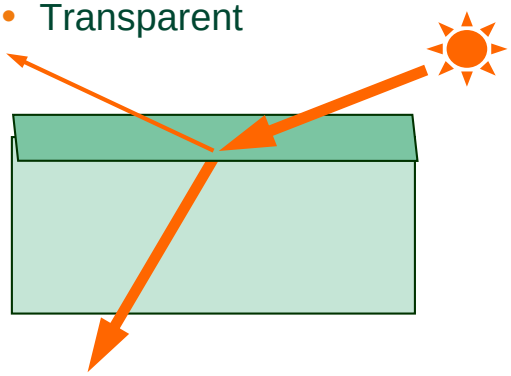

SIGGRAPH2006

We have seen that dielectrics are typically not highly reflective. In a purely homogeneous form, they are not very absorptive either – most incident light just gets transmitted through, so they are highly transparent. Examples are water, glass, crystal, transparent plastic, etc.

Inhomogeneous Dielectrics

- Transmitted light undergoes absorption and scattering until completely absorbed or re-exiting the surface

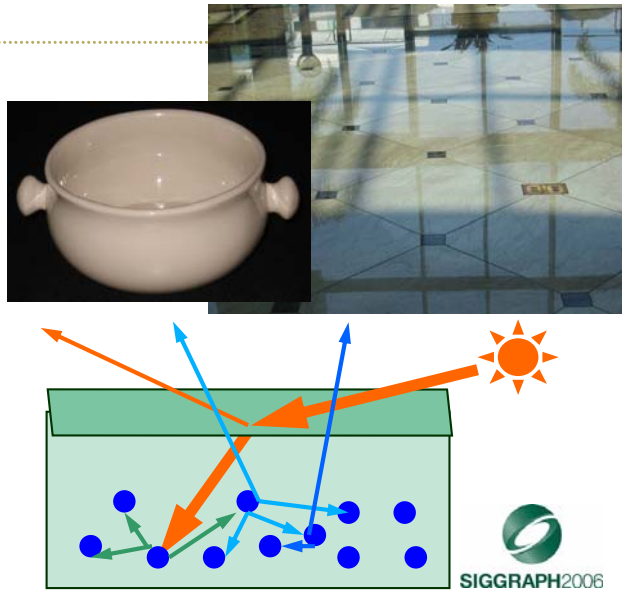

SIGGRAPH2006

Opaque dielectrics are inhomogeneous. The inhomogeneities are responsible for most of the scattering and absorption of visible light in dielectrics.

Surface vs. Volume Interactions

- Inhomogeneous dielectrics have two relevant modes of light interaction
  - Surface (reflection / refraction at the surface interface)
  - Volume (scattering / absorption by internal inhomogeneities)

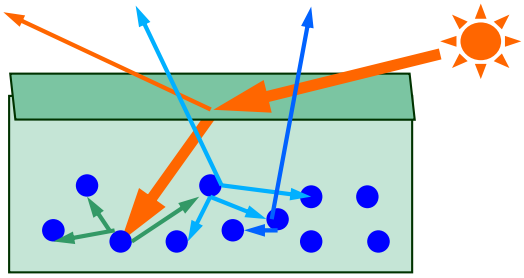

SIGGRAPH2006

Why have we specifically called out inhomogeneous dielectrics here? Volume interactions with light are not visually interesting for both metals and homogeneous dielectrics, but for opposite reasons. In metals the volume interactions with visible light are so strong that all transmitted light is quickly absorbed and none escapes. In homogeneous dielectrics the volume interactions with visible light are typically so weak that the transmitted light is not significantly attenuated.

Surface vs. Volume Interactions

- Inhomogeneous dielectrics have two relevant modes of light interaction
  - Surface (reflection / refraction at the surface interface)
  - Volume (scattering / absorption by internal inhomogeneities)

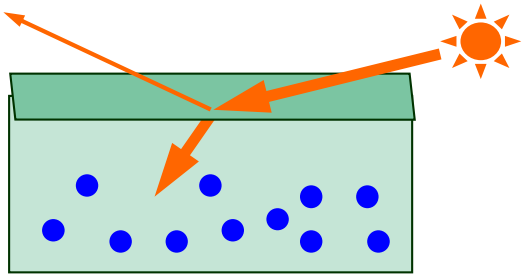

SIGGRAPH2006

Here to clarify, we look at each type of interaction separately. First the surface interaction.

Surface vs. Volume Interactions

- Inhomogeneous dielectrics have two relevant modes of light interaction
  - Surface (reflection / refraction at the surface interface)
  - Volume (scattering / absorption by internal inhomogeneities)

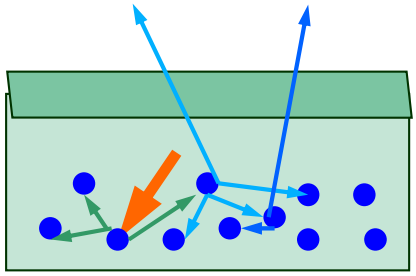

SIGGRAPH2006

Then the volume interactions (also called body or interior interactions, and often subsurface scattering although they include absorption as well as scattering phenomena). Note that both types of light-matter interaction occur where there is some kind of discontinuity in the optical properties of the medium. Note that these discontinuities can be substantial (particles of some foreign material) or structural (irregularities in molecular structure). If they are of a different substance, then their different optical properties may cause the light exiting the surface due to volume interactions to be quite different in color than that exiting due to surface interactions.

Volume Interactions as Reflectance

- 'Reflectance' is the relationship between incident and exitant light at a surface point.
- To model volume interactions as reflectance, focus on light incident / exitant on a surface point

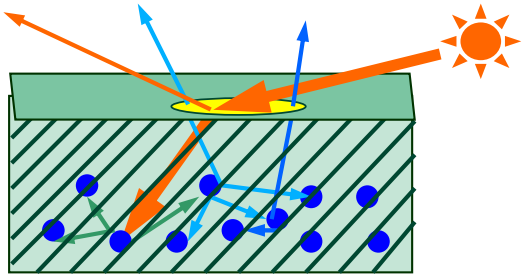

SIGGRAPH2006

It would seem that reflectance, which is inherently surface-oriented, would preclude modeling volume or interior interactions between light and matter. However, focusing on what's happening at the surface, and more specifically at a single point (or a small area that can be approximated as a single point) enables using reflectance to model volume interactions as well as surface ones. We look at the scattered light only after it has re-exited the surface, at a point close to the point of original entry (we will next try to give a feeling for what 'close' means in this context). We ignore the details of what happens in the interior.

Volume / Body Reflectance

- If light exits the surface at a large distance (compared to the scale of interest), cannot model volume interaction as reflectance

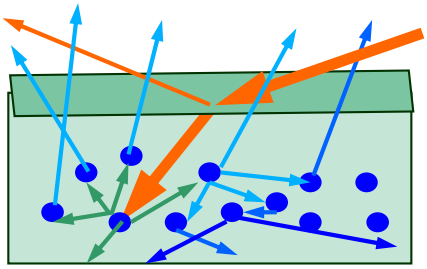

IMAGE BY H. W. JENSEN

SIGGRAPH2006

We will introduce a new term here, “volume” or “body” reflectance, as distinct from surface reflectance. Although all reflectance is modeled as occurring at a surface point, this distinguishes between reflectance originating from volume interactions and that originating from surface interactions. Here we see a case where we cannot model the volume interactions as reflectance at all, since they occur over a scale which is large compared to the scale of interest. This kind of macroscopic subsurface interaction is outside the scope of this talk.

The accompanying image illustrates this case: a small marble figurine (several inches tall) is being observed, at this scale the scattering of light within the marble occurs over visually significant distances.

Volume / Body Reflectance

- If light exits the surface at a small distance (compared to the scale of interest), can model volume interaction as *body reflectance*

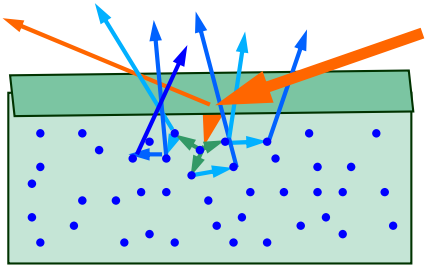

IMAGE BY H. W. JENSEN

SIGGRAPH2006

And here we see a case where the volume interactions are well modeled as reflectance, since they occur over a scale which is small compared to the scale of interest. Note also that there is a clear visual cue as to the size of an object from the scale of scattering.

The accompanying image shows an object composed of the same marble as the object in the previous image. However, the object is much larger (a large statue). At this scale, the scattering of light within the marble occurs over visually insignificant distances. The difference between the two images lies in the scale of observation, not the material properties.

Body Reflectance of Inhomogeneous Dielectrics

- Although most dielectrics have very low surface reflectance, body reflectance varies widely
  - Fresh snow: about 80%
  - Bright white paint: about 70%
  - Weathered concrete: about 35%
  - Stone: average about 20%
  - Earth (soil): average about 15%
  - Coal: close to 0%




Body reflectance can be seen as a the result of a ‘race’ between scattering and absorption - whether the light is absorbed before it has been scattered enough to re-exit the surface. The more frequently the light is scattered, the shorter the average distance traveled before re-exiting the surface. Over a shorter distance the light has less opportunity to be absorbed.

Body vs. Surface Reflectance

- Relative contribution of surface and body reflectance in inhomogeneous dielectrics varies with angle of incidence
  - Only light which was not reflected at surface is available to contribute to body reflectance
  - As angle of incidence increases, contribution of surface reflectance increases and that of body reflectance decreases

In a sense, surface reflectance has “first dibs” on the incident light since it occurs first.

Body vs. Surface Reflectance

IMAGE BY E. LAFORTUNE

Body vs. Surface Reflectance

IMAGE BY E. LAFORTUNE

SIGGRAPH2006

Body vs. Surface Reflectance

IMAGE BY E. LAFORTUNE

SIGGRAPH2006

Types of Reflectance

- Mirror Surfaces
  - Substances
  - Surface and Body Reflectance
- Non-Mirror Surfaces
  - Surface Roughness and Scale
  - Microscopically Rough Surfaces
  - Structured Surfaces

Now we will discuss surfaces which are not perfectly smooth – they have some kind of microscopic roughness or structure which affects their reflectance. We will discuss the relationship between surface roughness and scale, surfaces which have random (rough) microgeometry, and then surfaces which have structured microgeometry and how this affects their reflectance.

Non-Mirror Surfaces

- If the surface is not perfectly smooth, then the reflectance behavior is affected by the surface geometry as well as its substance
- Smoothness / Roughness is scale-dependent

Remember that ‘perfectly smooth’ for visual purposes means ‘smooth down to a scale below the wavelength of visible light’.

Surface Roughness and Scale

- Example:
  - Smooth at visible scale

SIGGRAPH2006

Smoothness is scale-dependent is several ways. To be relevant to visual phenomena, it has to take place on a scale at or above that of a wavelength of visible light. To be seen as contributing to a surface's reflectance behavior and not its shape, it has to take place at a scale below that at which the scene is being rendered, or in other words below visible scales. Also, a surface can have varying roughness at different scales and this may further affect reflectance behavior.

Here we see a surface which appears to be smooth.

Surface Roughness and Scale

- Example:
  - Smooth at visible scale
  - Rough below visible scale

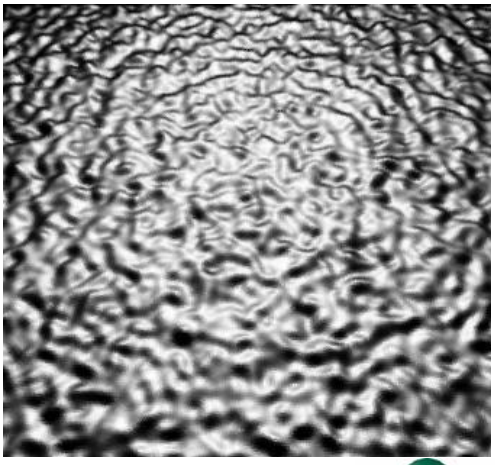

IMAGE BY S. WESTIN

SIGGRAPH2006

If we pick a small spot on the surface and magnify it, we will see that the surface exhibits considerable roughness on the microscopic scale.

Surface Roughness and Scale

- Example:
  - Smooth at visible scale
  - Rough below visible scale
  - At smaller scale still, smooth again

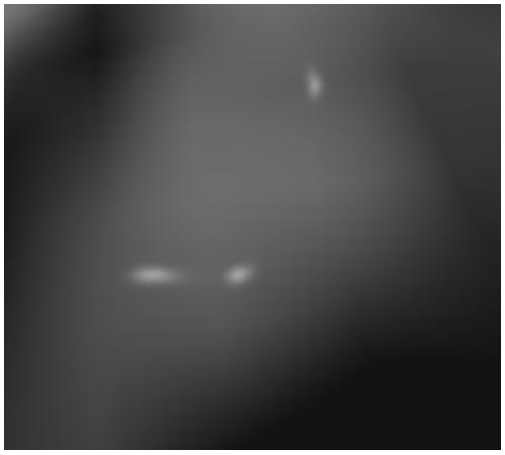

SIGGRAPH2006

Zooming in again on a single bump on the surface, we see the surface appears to be locally smooth. This scale (although quite small) is still much larger than a single wavelength of visible light.

Surface Roughness and Scale

- Example:
  - Smooth at visible scale
  - Rough below visible scale
  - At smaller scale still, smooth again
  - Finally rough at atomic scale (below visible light wavelength)

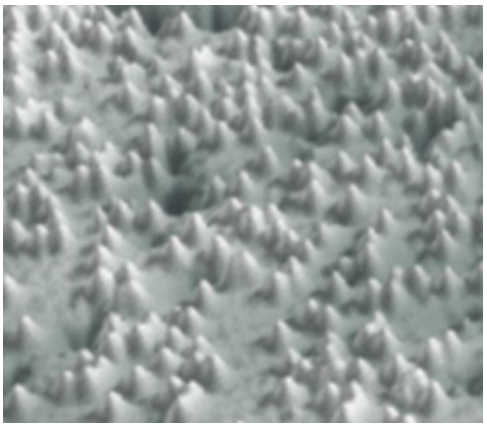

SIGGRAPH2006

Finally we see the surface at the atomic scale. It is quite rough but this will have no affect of the reflection of visible light since this scale is much smaller than a single wavelength, exhibiting considerable roughness.

Microscopically Rough Surfaces

- These surfaces have microscopic facets with a continuous distribution of surface normals
  - These facets ‘spread’ the light reflecting from them
- In most surfaces distribution of micro-normals is not uniform, but peaks at the macroscopic surface normal

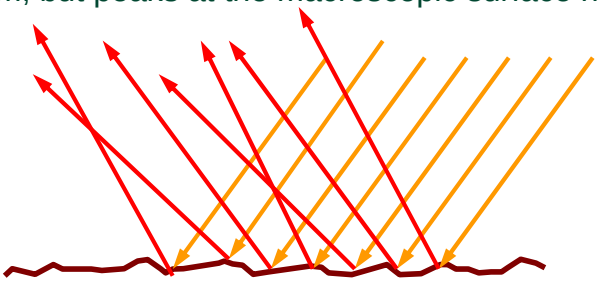

SIGGRAPH2006

This category includes all surfaces which are not mirror surfaces. Although these surfaces appear smooth at the visible scale, they contain micro-scale surface elements with varying surface normals. For most surfaces the distribution of surface normals is a smooth distribution over the hemisphere which peaks strongly at the macroscopic surface normal. Microscopically rough surfaces can exhibit a continuum of roughness, here we see a relatively smooth surface where the reflected light is only spread out a little bit.

Microscopically Rough Surfaces

- These surfaces have microscopic facets with a continuous distribution of surface normals
  - These facets ‘spread’ the light reflecting from them
- In most surfaces distribution of micro-normals is not uniform, but peaks at the macroscopic surface normal

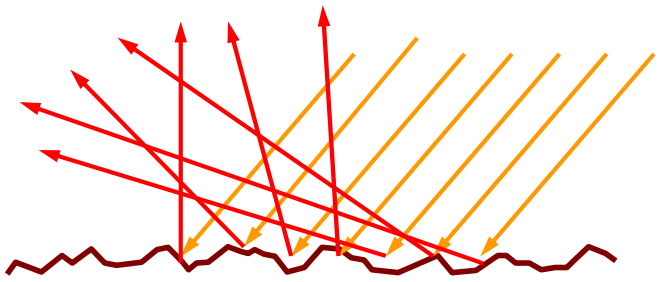

SIGGRAPH2006

And here we see a rougher surface where the reflected light is spread out more. Later in the course we will quantify the relationship between the surface roughness and the distribution of the reflected light.

Microscopically Rough Surfaces

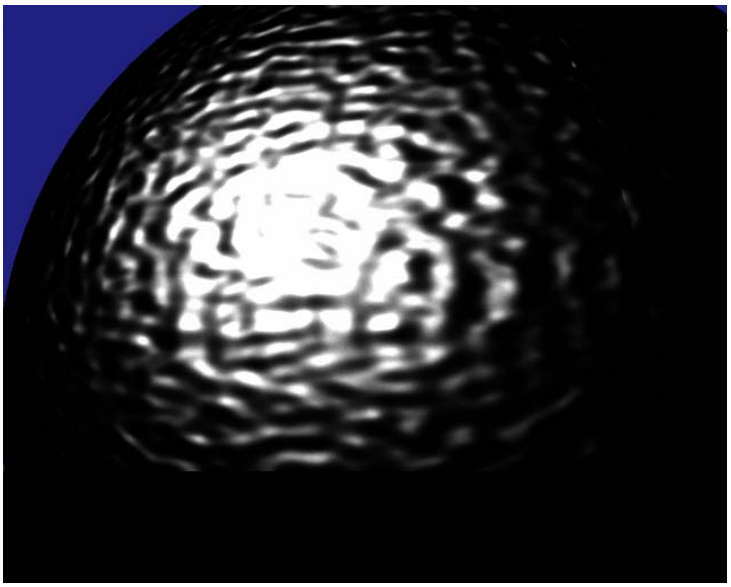

SIGGRAPH2006

Another way to think about this is to look at a bumpy curved surface. In this image, each bump has an individual highlight.

Microscopically Rough Surfaces

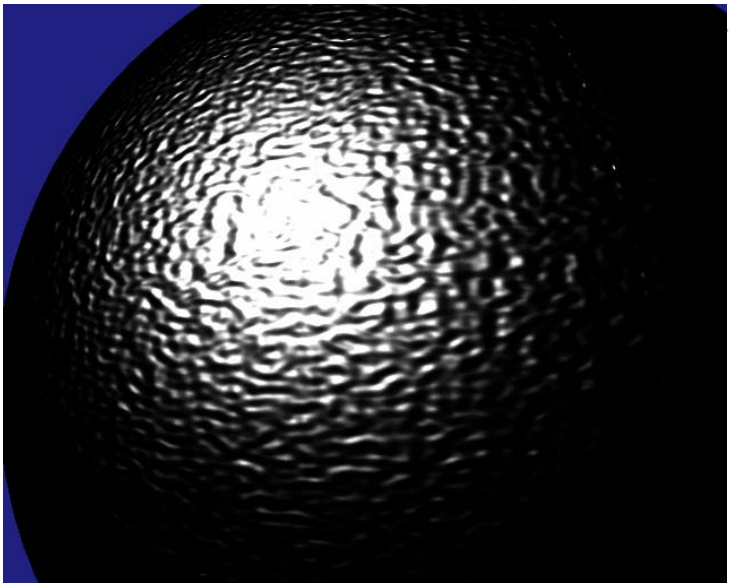

SIGGRAPH2006

When we make the bumps smaller (or look at the surface on a larger scale), the individual highlights becomes more difficult to see but still affect the overall appearance of the material. Note that the normal distribution remains the same throughout this image sequence, only the scale changes.

Microscopically Rough Surfaces

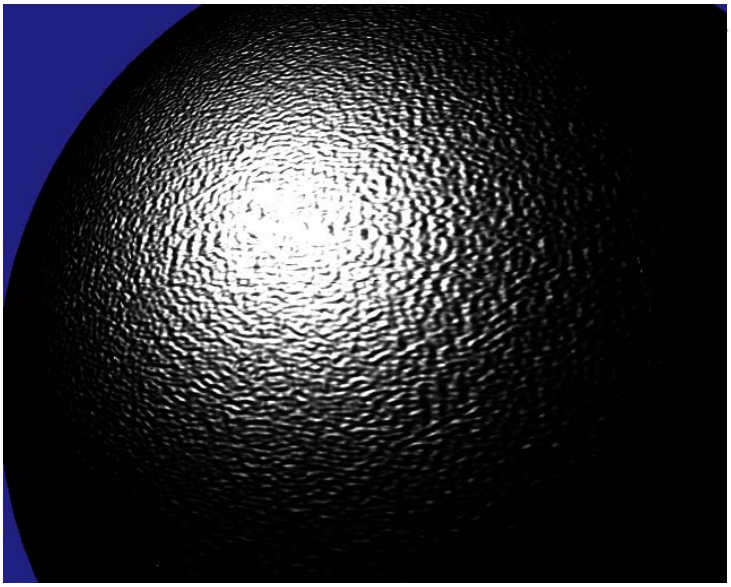

SIGGRAPH2006

Here the bumps are smaller still.

Microscopically Rough Surfaces

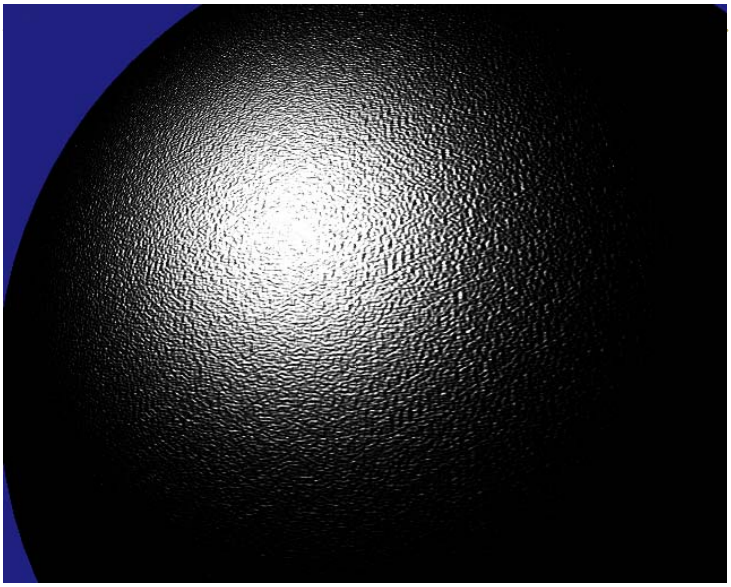

SIGGRAPH2006

At this scale, an overall pattern can be clearly seen where the individual bump highlights are clustered most densely at the center and become more spread out toward the outside.

Microscopically Rough Surfaces

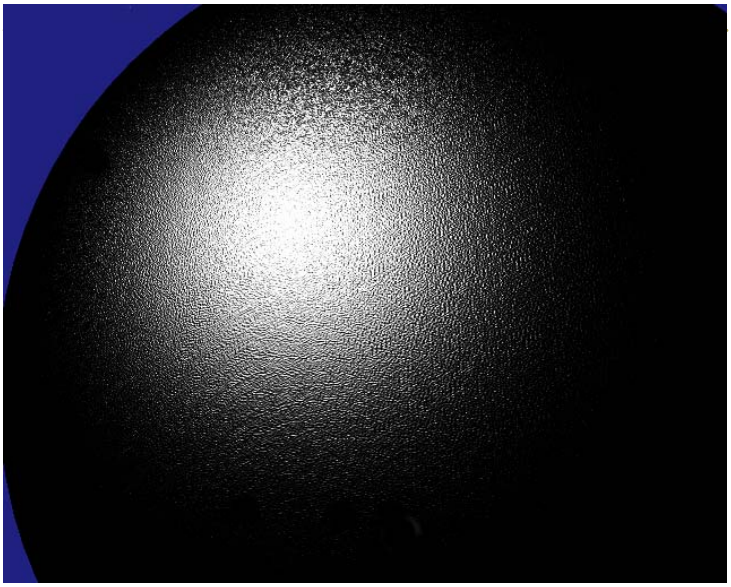

SIGGRAPH2006

At this scale the individual highlights are almost invisible, and the overall pattern (which looks very much like a single large highlight) stands out.

Microscopically Rough Surfaces

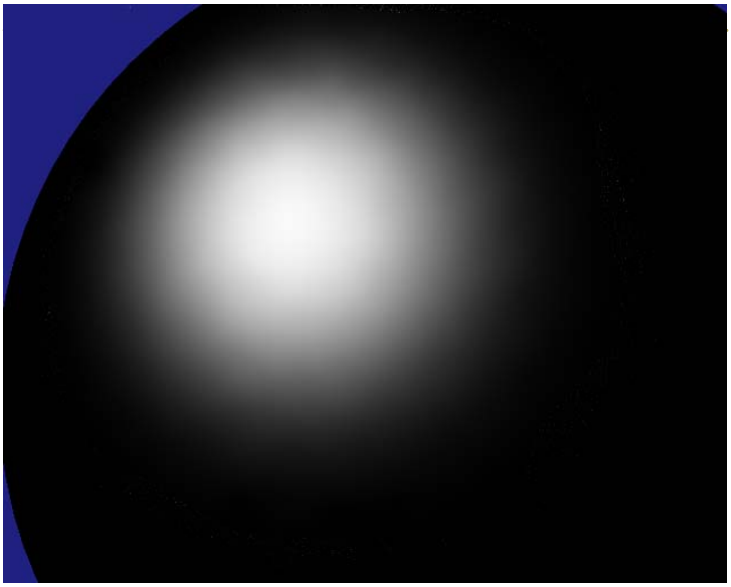

SIGGRAPH2006

Finally at this scale the bumps are not visible at all and the surface appears smooth. The effect of the bumps can be seen in the falloff of the single large highlight.

Microscopically Rough Surfaces

- Microscopic roughness also causes:
  - Shadowing

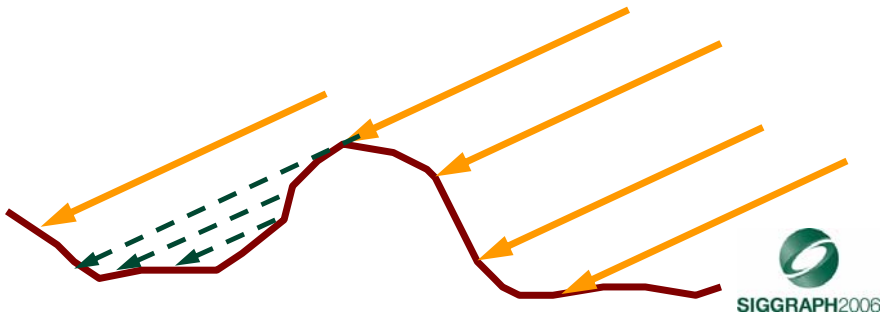

Besides causing the surface to have a continuous distribution of normals (rather than a single surface normal), the micro-geometry affects reflectance in other ways.

'Shadowing' refers to some of the micro-facets blocking light from others. The shadowing depends on the exact shape of the micro-geometry, but not directly on the normal distribution. In this image, the area with the black dashed arrows is shadowed.

Microscopically Rough Surfaces

- Microscopic roughness also causes:
  - Shadowing
  - Masking

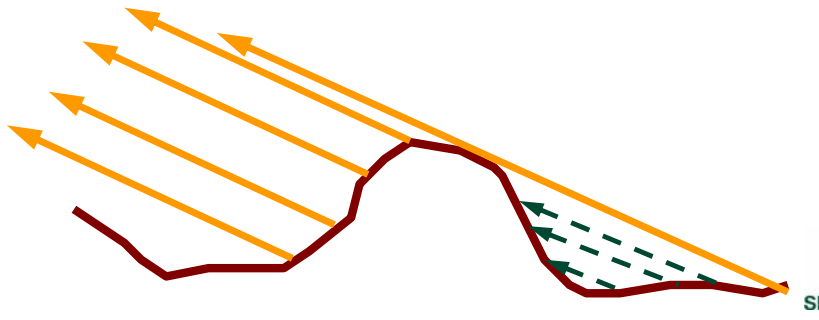

SIGGRAPH2006

'Masking' refers to some micro-facets obscuring others from the view position. This also depends on properties of the microgeometry other than the normal distribution. In this image, the area with the black dashed arrows is masked.

Microscopically Rough Surfaces

- Microscopic roughness also causes:
  - Shadowing
  - Masking
  - Interreflections

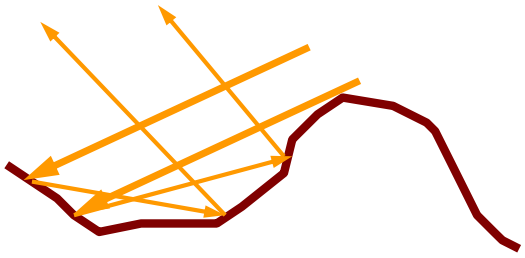

SIGGRAPH2006

Finally, some of the light that was shadowed or masked is reflected off the surface. Light may undergo several bounces before it reaches the eye.

Microscopically Rough Surfaces

- The reflectance of a surface is a function both of its surface physics and its microscale structure

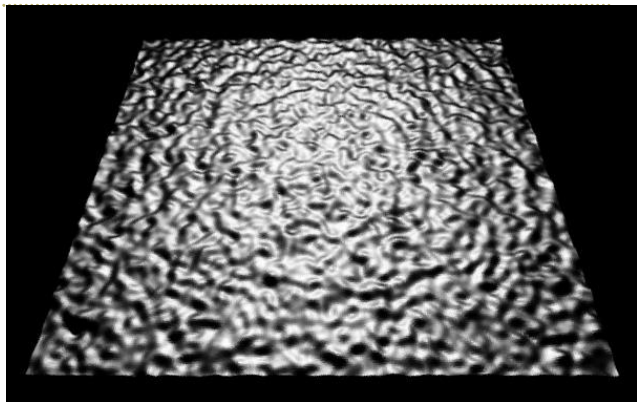

IMAGE BY S. WESTIN

SIGGRAPH2006

Just to re-emphasize. This is one of the most important concepts from this part of the course.

Microscopically Rough Surfaces

- Small, bright lights reflect as highlights
- Blurry reflection of environment

SIGGRAPH2006

These surfaces commonly exhibit glossy reflections. The highlights' size and shape mostly depends on the distribution of microfacet normals.

Microscopically Rough Metals

- Reflectance is dominated by the primary surface reflection
- In rougher surfaces, secondary reflections create an additional diffuse reflection
  - More strongly colored

Secondary reflections from rough colored metals are more saturated than the primary reflection since each bounce effectively multiplies the spectral reflectance into the result again. Burnished gold is an example of this, where the gold has a richer yellow color due to being rougher.

Microscopically Rough Dielectrics

- Highlights tend to be weaker in dielectric surfaces due to lower surface reflectance
  - For the same reason, secondary reflections are less noticeable
  - Diffuse mostly due to body reflectance

As seen for smooth dielectric surfaces, rough dielectric surfaces also exhibit a tradeoff between body and surface reflections at glancing angles (the only difference is that here the surface reflection is blurred).

Extremely Rough Dielectrics

- Distribution of normals is nearly uniform
- Diffuse reflection, with some retroreflection

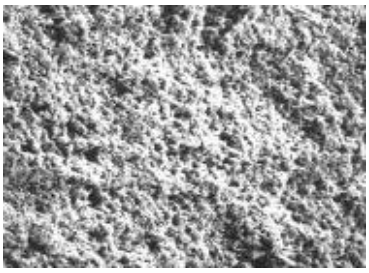

IMAGE BY T. HIMLAN

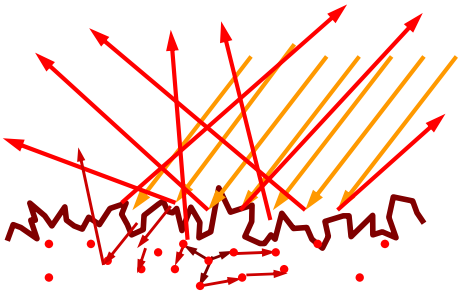

The roughest surfaces have micro-normals uniformly spread over the hemisphere so that light reflected from the surface is spread evenly in all directions.

For dielectrics, this erases the coherence which is the important visual distinction of surface reflections over body reflections, and the two can be lumped together into a single diffuse reflection.

So although for metals, extremely rough surfaces can be treated as simply an extreme case of microscopically rough surfaces, for dielectrics it is useful to treat these surfaces as a distinct case.

Rough dielectrics include surfaces such as dust, rough clay, unpolished stone and concrete. These are surfaces that we think of as “matte” or “diffuse”. Such surfaces tend to appear flat (e.g. the full moon) and contrary to expectations, do not usually obey Lambert’s law.

Extremely Rough Dielectrics

- Retroreflective tendencies caused by foreshortening and occlusion of shadowed parts of surface when eye and light directions are similar

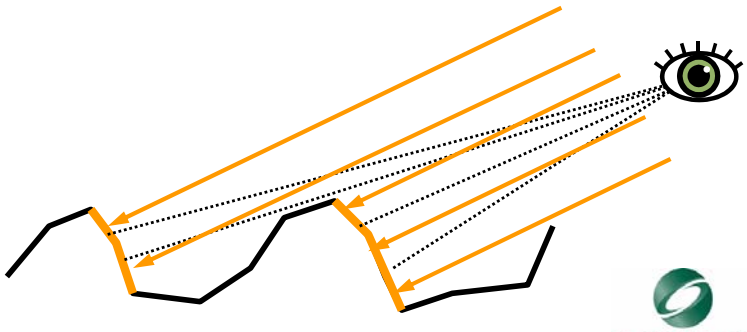

SIGGRAPH2006

If we imagine each facet being essentially Lambertian (which is a reasonable approximation to what is going on in these surfaces), then when the light and view directions are similar the surfaces we can see are the ones which are more strongly lit.

Extremely Rough Dielectrics

- And by foreshortening and occlusion of lit parts of surface when eye direction and light direction are very different

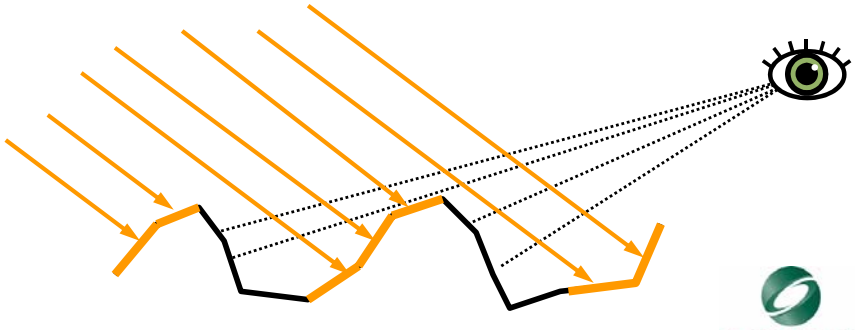

SIGGRAPH2006

When the light and view directions are very different the surfaces we can see are the ones which are more dimly lit.

Extremely Rough Dielectrics

- Lights reflect diffusely
- No environment details visible

Structured Surfaces

- So far the surfaces we have seen have been smooth or had random, unstructured microgeometry
- Surfaces with regular, structured microgeometry exhibit interesting reflective characteristics

These surfaces may be metallic or dielectric; since the microstructure is the interesting thing about them we will not distinguish the two cases.

Anisotropic Surface

- When the surface microstructure is anisotropic (not the same in all directions), then the reflectance exhibits directionality
- Examples: surfaces such as wood, brushed metal which have an oriented structure

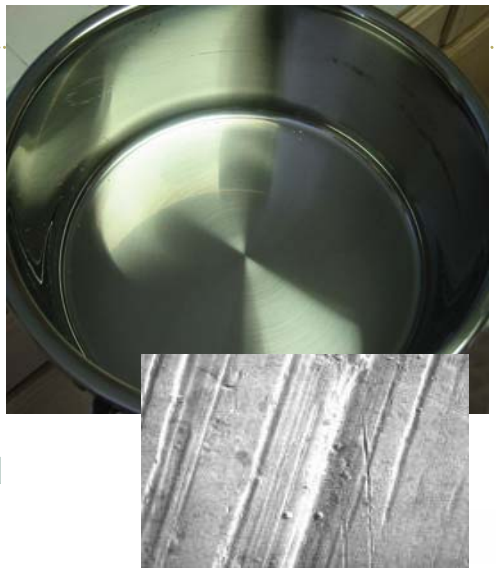

PHOTOMICROGRAPH BY T. HIMLAN

One way in which the microgeometry can be structured is if it is not the same in all directions – if it is anisotropic. This causes the reflection to exhibit directional behavior.

Anisotropic Surface

- Highlights are stretched
- Environment is directionally blurred

Retroreflective Surface

- Microgeometry is constructed to reflect most of the light back in the incident direction

SIGGRAPH2006

Most strongly retroreflective materials are artificial materials designed for use in street signs, projection screens, etc. Some rare natural surfaces also exhibit strong retroreflective behavior (such as cat's eyes).

Other Structured Surfaces

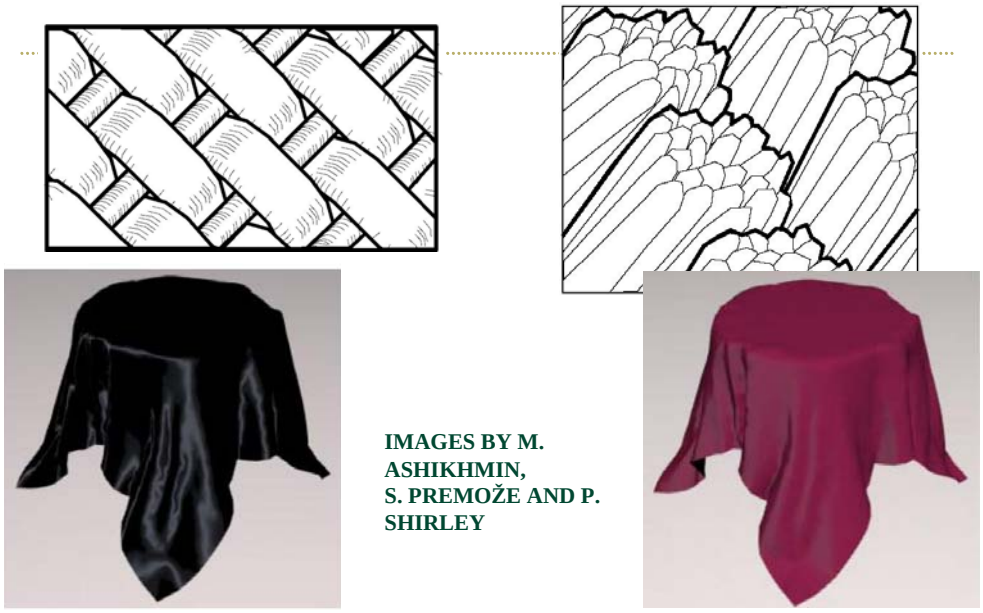

IMAGES BY M. ASHIKHMIN, S. PREMOŽE AND P. SHIRLEY

Fabrics are usually structured surfaces, and many of them exhibit interesting reflective properties.

Reflectance

- Types of Reflectance
- Reflectance Theory
- Reflection Model Foundations

We will now discuss the quantitative theory behind reflection models.

Reflectance Theory

- Radiometric Theory
- The BRDF

Here we will lay the theoretical groundwork needed to understand reflectance from a physical standpoint. First we will discuss radiometric theory, and then the concept of the BRDF.

Radiometric Theory

- Radiometry: measurement of radiant energy
- Specifically electromagnetic radiant energy
- In graphics, specifically energy in the visible portion of the EM spectrum (~380nm-780nm)

Radiometric Quantities

- Radiant Flux Φ
  - Total radiant power
  - Power
  - Watts

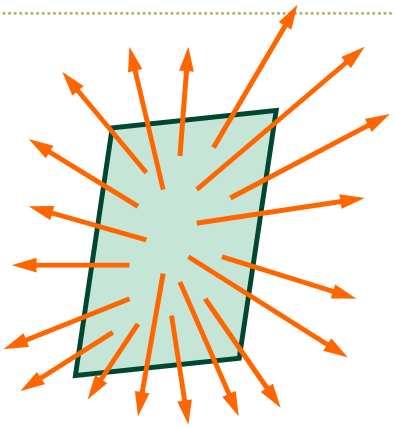

SIGGRAPH2006

In this example, we are looking at a window. The power of all the light pouring through all parts of the window, in all directions, is radiant flux.

Radiometric Quantities

- Radiant Exitance (Radiosity) B
  - Flux area density, exitant from a point
  - Power per surface area
  - Watts / meter^2

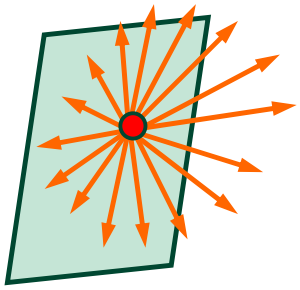

SIGGRAPH2006

The area density of the power *exitant* (coming out of) a single point on the window is radiant exitance (also called radiosity).

Radiometric Quantities

- Irradiance E
  - Flux area density, incident to a point
  - Power per surface area
  - Watts / meter^2

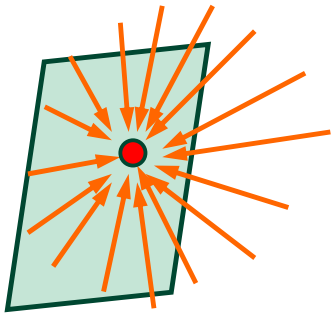

SIGGRAPH2006

Irradiance is very similar to radiosity, but it measures *incident* (incoming) light rather than exitant light.

Radiometric Quantities

- Radiant Intensity I
  - Flux directional density, in a given direction radiant power
  - Power per solid angle
  - Watts / steradian

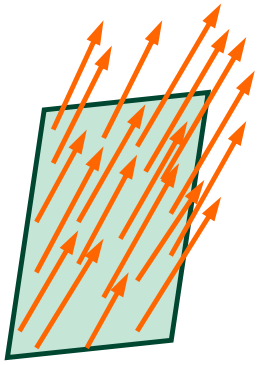

SIGGRAPH2006

In contrast to radiant exitance (the surface density of flux at a point), radiant intensity is the directional density of flux in a given direction.

A solid angle is a 3D extension of the concept of an angle – it is a sheaf of directions in 3-space. Just as an angle can be measured in radians as the length of an arc on a unit circle, a solid angle can be measured as the area of a patch on a unit sphere. Solid angles are measured in steradians, of which there are 4π in a sphere.

Radiant intensity is particularly relevant when examining a very distant (or very small) light source, where the variation over direction is more important than over surface area.

Radiometric Quantities

- Radiance L
  - Radiant power in a single ray
  - Power per (projected) surface area per solid angle
  - Watts / (meter^2 steradian)

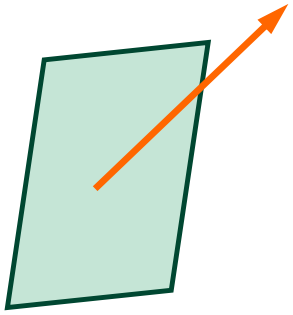

SIGGRAPH2006

Unlike radiosity and irradiance, radiance is not tied to a surface, but to a ray. Radiance is important in rendering because the final pixel RGB color is derived directly from the radiance in the rays through the pixel sample locations.

Flux is defined over all points of a surface in all directions. Radiosity is the surface flux density out of a single point in all directions, and radiant intensity is the directional flux density in a single direction, through all points. Radiance combines the specificities of radiosity and radiant intensity – it measures light through a single point in a single direction, in other words a single ray of light.

Radiometric Quantities

- Radiance L
  - Radiant power in a single ray
  - Power per (**projected**) surface area per solid angle
  - Watts / (meter^2 steradian)

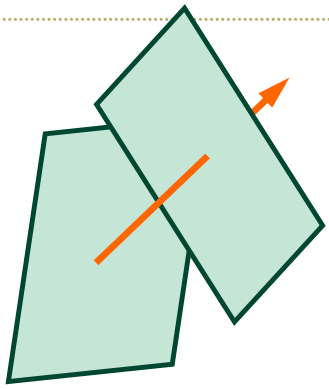

SIGGRAPH2006

The surface area here (unlike the other quantities) is projected surface area, or surface area perpendicular to the ray direction.

Light Direction

- Until this slide, the light arrows have pointed in the direction the light is going
- From now on, we will use ‘light direction’ vectors which point towards the light

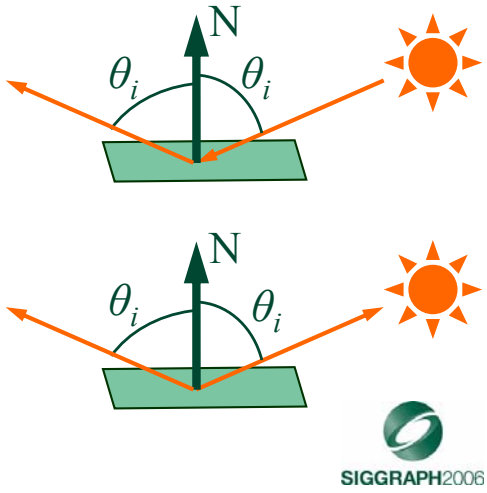

SIGGRAPH2006

Until now, we have been showing the physics of light bouncing around. From now on, we are talking about the math and implementation and there it is both customary and convenient to have the ‘light direction’ pointing **to** the light.

Radiance and Irradiance

$$dE = L_i(\omega_i)\cos\theta_i d\omega_i$$

$$E = \int_{\Omega_H} L_i(\omega_i)\cos\theta_i d\omega_i$$

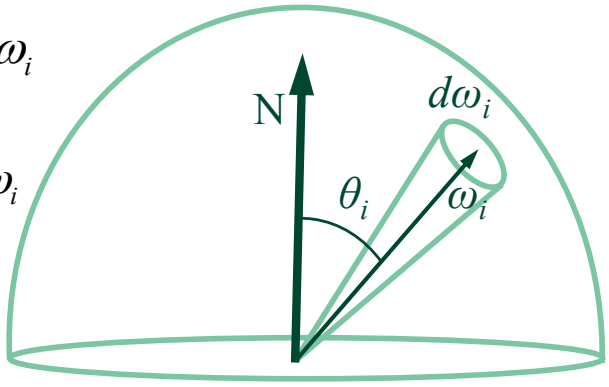

Cosine terms always clamped!

SIGGRAPH2006

Let us look at a surface point, and how it is illuminated by an infinitesimal patch of incident directions with solid angle $d\omega_i$. Since this patch of incident directions is infinitesimal, it can be accurately represented by a single incident direction ω_i and we can assume that the incident radiance from all directions in this patch is a constant $L_i(\omega_i)$. This patch contributes an amount of irradiance equal to $L_i(\omega_i)$ times $d\omega_i$, times the cosine of the angle θ_i between the incident direction and the surface normal. If we integrate this over the hemisphere Ω_H (centered on the normal), we get the total irradiance. Note that for a light source to illuminate a surface (contribute to its irradiance), it needs to both have a non-zero radiance and subtend a non-zero solid angle. Finally, note that the cosine here is only valid over the hemisphere and is assumed to be clamped to 0 outside it. This is true for all cosine terms we will see in this talk.

The (Clamped) Cosine Factor

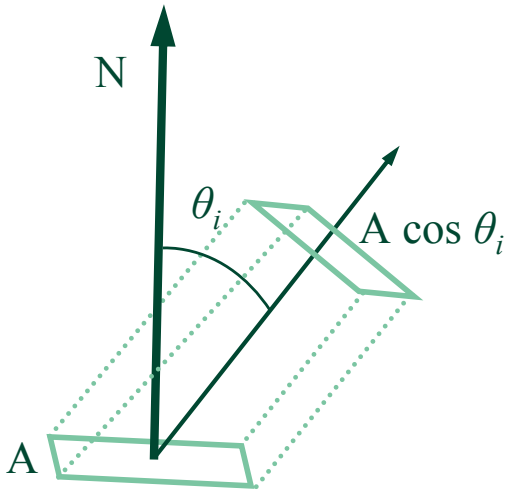

SIGGRAPH2006

Where did the (clamped) cosine factor come from? The cosine is there because radiance is defined relative to an area perpendicular to the ray, and irradiance is defined relative to an area parallel to the surface. Another way of looking at it is that the same radiance, coming in at a more oblique angle, contributes a smaller amount to the irradiance because it is 'spread out' more. We can also see here why it is clamped – if the incident direction is under the surface then there is no contribution.

Reflectance Theory

- Radiometric Theory
- The BRDF

Now we will discuss the concept of the BRDF.

The BRDF

- Bidirectional Reflectance Distribution Function
  - Ratio of irradiance to reflected radiance

$$f_r(\omega_i, \omega_e) = \frac{dL_e(\omega_e)}{dE(\omega_i)}$$

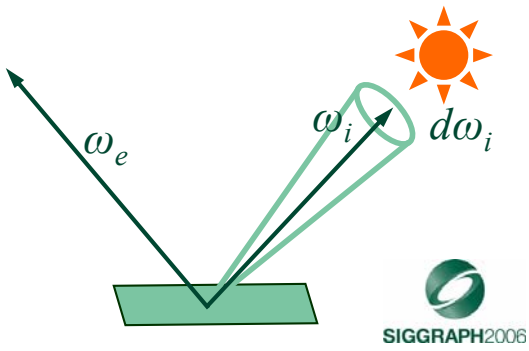

SIGGRAPH2006

ω_i is the direction to the incident irradiance, and ω_e is the direction to the exitant reflected radiance. For every such pair of directions, the BRDF gives us the ratio between irradiance and exitant radiance. Since the incident direction and the excitant direction each have two degrees of freedom (a common parameterization is to use two angles: elevation θ relative to the surface normal and rotation φ about the normal), the BRDF is a function of four scalar variables in the general case.

Since the BRDF is radiance (power/(area x solid angle)) divided by irradiance (power/area), its units are inverse solid angle, or steradians^{-1}.

The BRDF

- Incident, excitant directions defined in the surface's *local frame* (4D function)

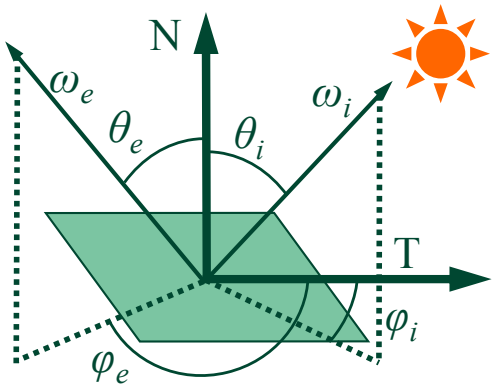

SIGGRAPH2006

Directions defined relative to the local surface normal and tangent.

The BRDF

- Isotropic BRDFs are 3D functions

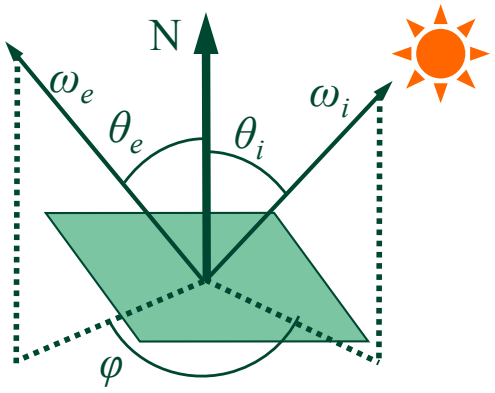

SIGGRAPH2006

For most surfaces, the relation of the incident and exitant directions to a local surface tangent doesn't matter (these surfaces are isotropic and have no local preferred direction). So instead of the rotations between each of these two directions and a tangent vector, the BRDF depends on the rotation between the two directions, which removes one degree of freedom. Note that the tangent vector is no longer needed.

The Reflection Equation

$$f_r(\omega_i, \omega_e) = \frac{dL_e(\omega_e)}{dE(\omega_i)} \qquad dE = L_i(\omega_i)\cos\theta_i d\omega_i$$

From the definition of a BRDF and the relation between irradiance and radiance, we get:

The Reflection Equation

$$f_r(\omega_i, \omega_e) = \frac{dL_e(\omega_e)}{dE(\omega_i)} \qquad dE = L_i(\omega_i)\cos\theta_i d\omega_i$$

$$L_e(\omega_e) = \int_{\Omega_H} f_r(\omega_i, \omega_e) L_i(\omega_i)\cos\theta_i d\omega_i$$

The reflection equation. This means to get the exitant radiance in a direction ω_e, we need to integrate the incident radiance, times the BRDF, times the cosine of the angle with the normal, over all incoming directions in the hemisphere around the surface normal.

The BRDF

- The laws of physics impose certain constraints on a BRDF
- To be physically plausible, it must be:
  - Reciprocal
  - Energy-Conserving

Reciprocity

- More properly, *Helmholtz Reciprocity*

$$f_r(\omega_i, \omega_e) = f_r(\omega_e, \omega_i)$$

All this means is that the surface must reflect light the same way in both directions – if incoming and outgoing directions are changed, the reflectance must remain the same.

Energy Conservation

- Directional-Hemispherical Reflectance: $R(\omega_i)$
- $R(\omega_i)$ must be less or equal to 1 for all ω_i and for all wavelengths (i.e. for R, G, B)

$$R(\omega_i) = \frac{dB}{dE(\omega_i)} = \int_{\Omega_H} f_r(\omega_i, \omega_e) \cos\theta_e \, d\omega_e$$

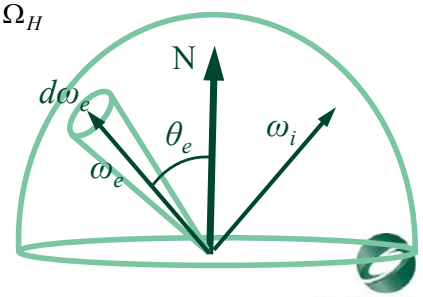

SIGGRAPH2006

The directional-hemispherical reflectance is the ratio of differential radiant exitance to differential irradiance. It tells us how much of the incoming radiant energy from a given direction is absorbed and how much is reflected.

The reflected light energy cannot be more than the incident light energy, which means that the directional-hemispherical reflectance must be less or equal to 1. 1 means a perfectly reflective surface with no absorption at the given incident angle.

Bihemispherical Reflectance / Albedo

$$\rho = \frac{B}{E} = \frac{1}{\pi} \int_{\Omega_H} R(\omega_i) \cos\theta_i d\omega_i =$$

$$\frac{1}{\pi} \int_{\Omega_H} \int_{\Omega_H} f_r(\omega_i, \omega_e) \cos\theta_i \cos\theta_e d\omega_i d\omega_e$$

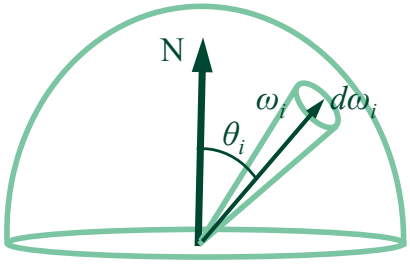

The bihemispherical reflectance (also called albedo) is the ratio between radiant exitance and irradiance. Like the directional-hemispherical reflectance, it is between 0 and 1, where 1 indicates a surface with no absorption. It is an overall measure of reflectivity – how much radiant flux striking the material from all angles is reflected vs. how much is absorbed. It can also be seen as the cosine-weighted average of the directional-hemispherical reflectance. It is computed by integrating the BRDF over all incident and exitant directions. The $1/\pi$ is a normalization factor, this is the first of many times we will see it. The reason it is so ubiquitous is that integrating the cosine factor over the hemisphere yields pi.

Reflectance

- Remember that reflectance values such as R and ρ are restricted to a range of 0 to 1
  - 0 is perfectly absorbent
  - 1 is perfectly reflective
- This does not apply to the BRDF!
  - For example, in the center of a tight highlight the BRDF value will be quite high

A reflectance value is a ratio of reflected light to incident light. By definition, it must be between 0 and 1 (at least for a non-emissive surface).

In contrast, the BRDF is a distribution function – if the distribution it describes is highly non-uniform (which it will be for a smooth surface) then it can have arbitrarily high values in certain directions.

The Reflection Equation with Point / Directional Lights

- Important case for real-time rendering
  - Lighting from infinitely small / distant light sources
  - Ambient / indirect lighting is computed separately
  - Point lights characterized by intensity I and position
  - Directional lights characterized by direction and I/d^2

$$L_{point\,e}(\omega_e) = \sum_l \frac{I_l}{d_l^2} f_r(\omega_l, \omega_e)\cos\theta_l$$

Point and directional lights are abstractions, since they imply infinite energy density. However, they are useful approximations for highly localized or distant light sources. In many scenes, these contribute most of the lighting. These lights are best characterized by radiant intensity rather than by radiant exitance or radiance. A directional light is a point light which is so distant that its direction and distance can be taken as constant over the rendered scene.

The result shown here can be derived by considering a small uniformly glowing sphere with radius r_l and flux Φ_l, illuminating a relatively distant surface point. Since the subtended solid angle is extremely small, we assume that radiance, BRDF and cosine factor are constant and can be taken out of the rendering equation integral, which leaves the solid angle. For a small distant sphere, the solid angle can be approximated as $\pi(r_l/d_l)^2$ where d_l is the distance from the sphere center to the surface point. This results in: $L_e = L_l\pi(r_l/d_l)^2\cos\theta_l f_r(\omega_e,\omega_l)$.

For this sphere, $B_l = \Phi_l/(4\pi r_l^2) = \pi L_l$, from which follows that $L_l = \Phi_l/(4\pi^2 r_l^2)$. Combined with the previous result, we get $L_e = (\Phi_l/(4\pi d_l^2))\cos\theta_l f_r(\omega_e,\omega_l)$. Since r_l is no longer present in the result, it will hold as a limit when r_l goes to 0.

For a point light source, $I_l = \Phi_l/(4\pi)$ which combined with the previous result gives us the result shown in the slide above.

Note: although I_l/d_l^2 seems to be power / (area x solid angle), it is actually power / area, the same as irradiance and radiant exitance (solid angles, being dimensionless quantities, often confuse unit analysis).

Scale and the BRDF

- Reflectance Models are intimately tied to scale
- The complexity of many BRDFs originates in statistically modeled subpixel structure

We've seen a little bit of this in the earlier discussion of microscopically rough materials.

For example, the BRDF of an object like a tree varies depending on whether we are looking at a single leaf, an entire tree, or a forest. In each case, the structure covered by a single pixel is very different.

Scale and the BRDF

- The same features may be modeled as
  - Geometry
  - Bump maps
  - BRDFs
- Depending on scale
- Smooth transitions between these representations pose interesting issues

Reflectance

- Types of Reflectance
- Reflectance Theory
- Reflection Model Foundations

We will now discuss the building blocks from which most reflection models are built.

Reflection Model Foundations

- Lambert
- Fresnel Equations
- Microfacet Theory

In this section, we will discuss the basic building blocks used to build most reflection models. First we will discuss the Lambert BRDF, then Fresnel's equations and finally microfacet theory.

Lambert

- Constant BRDF (cosine factor not part of BRDF!)
- Pure Lambert impossible
- Usually used for dielectric body reflectance term

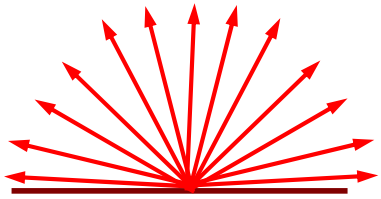

SIGGRAPH2006

A Lambertian surface reflects the same radiance in all directions. Remember that the well-known Lambertian cosine factor is actually part of the reflectance equation and not the BRDF. A perfectly Lambertian surface is an abstraction that does not exist in nature.

The Lambertian BRDF is view-independent – changing the angle of exitance has no effect on the value. This means that a Lambertian surface will not change in brightness when viewed at different angles. Due to reciprocity, only the constant Lambertian BRDF can have this property (since reciprocity implies that a BRDF which independent of the angle of exitance must also be independent of the angle of incidence).

Lambert

- The Lambertian BRDF is a constant value over all incident and exitant directions
- BRDF equal to bihemispherical reflectance (albedo) over π

$$\rho = \frac{1}{\pi} \int_{\Omega_H} \int_{\Omega_H} f_{rLambert} \cos\theta_i \cos\theta_e d\omega_i d\omega_e = \pi f_{rLambert}$$

$$f_{rLambert} = \frac{\rho}{\pi}$$

Given that the Lambertian BRDF is a constant, what is it's value? We can calculate the bihemispherical reflectance, or albedo, and see that the BRDF is equal to the bihemispherical reflectance over pi. This is useful, since the bihemispherical reflectance is an intuitive number, between 0 and 1, which tells us how reflective the surface is overall. Note since the bihemispherical reflectance is wavelength-dependent, it is an RGB triple, usually thought of as the material's diffuse color.

Reflection Model Foundations

- Lambert
- Fresnel Equations
- Microfacet Theory

Now we will discuss Fresnel's equations.

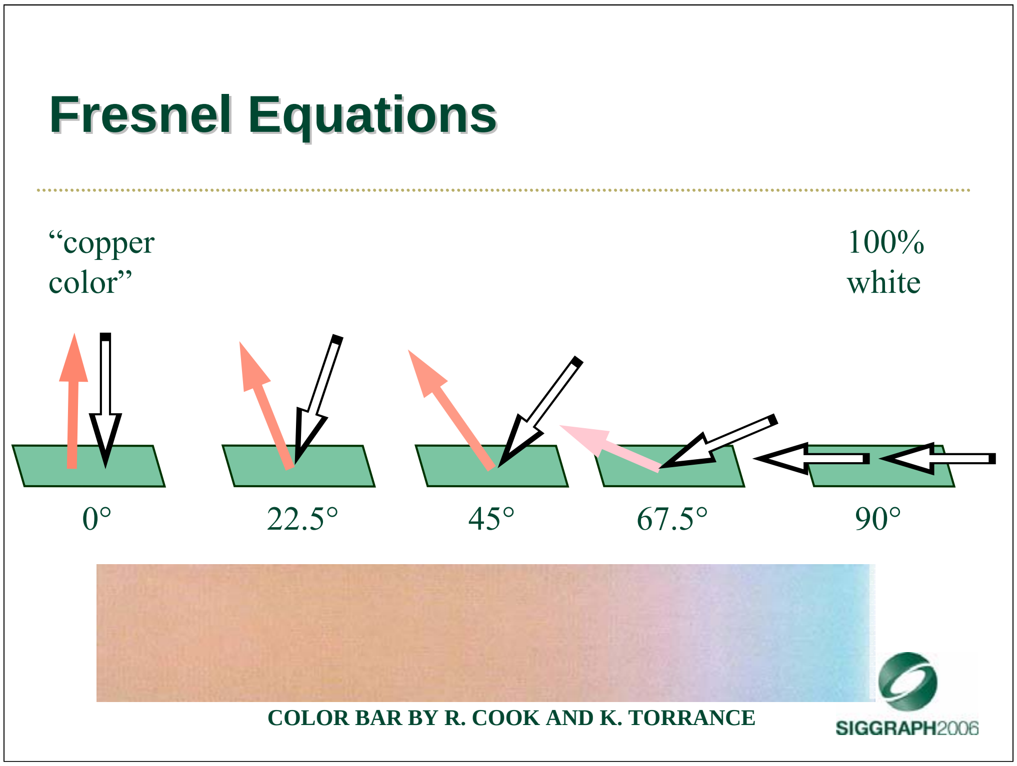

This is a recap of a previous slide. Note that the color shift is not a simple interpolation from the normal incidence reflectance color to 100% white – there is a shift to blue just before it goes to 100% white.

Fresnel Equations

- Full equations expensive to compute, need (possibly complex) refractive index data for different wavelengths
- Schlick approximation:

$$R_F(\theta_i) = R_F(0) + (1 - R_F(0))(1 - \cos\theta_i)^5$$

- $R_F(0)$ is the directional-hemispherical reflectance at normal incidence

The Schlick approximation is accurate to within a few percent, is much cheaper to compute, and has a much more intuitive parameter: $R_F(0)$, which is the reflectance at normal incidence. Note that the reflectances here are all directional-hemispherical.

$R_F(0)$ is commonly thought of as the materials' specular color. It is relatively high for metals, and low for dielectrics. The cosine factor (like all others in this course) is clamped to zero.

This approximation is quite good, though it does miss some color-shifts at glancing angles.

Fresnel Equations

- $R_F(0)$ is high for metals
  - Steel ~0.55, Silver, Aluminum ~0.95
  - Colored metals
    - Gold: from ~0.6 for blue to ~0.9 for red
    - Copper: from ~0.4 for blue to ~0.85 for red
- $R_F(0)$ lower for dielectrics
  - Water, glass, plastic, etc. ~0.04
  - Diamond ~0.15

Colorless metals have $R_F(0)$ which is almost constant over the visible spectrum. Colored metals tend to have higher $R_F(0)$ for longer wavelengths. Dielectrics are usually colorless and have low $R_F(0)$.

Reflection Model Foundations

- Lambert
- Fresnel Equations
- Microfacet Theory

Finally in this section, we discuss microfacet theory.

Microfacet Theory

- Models some effects of microgeometry
  - Single-bounce surface reflectance
  - Shadowing
  - Masking
- Doesn't model
  - Interreflections
  - Body reflectance

Microfacet theory is a useful framework for modeling microscopically rough surfaces. It does not model diffuse effects such as surface interreflections and body reflectance, and for this reason microfacet BRDFs usually add a separate term (most often Lambertian) to model the diffuse reflectance.

Microfacet Theory

- Surface modeled as flat microfacets
  - Each microfacet is a Fresnel mirror
- Normal Distribution Function (NDF) $p(\omega)$
  - $p(\omega)d\omega$ is fraction of facets which have normals in the solid angle $d\omega$ around ω
  - ω is defined in the surfaces' local frame
  - There are various different options for $p(\omega)$

We will look at various options for modeling the NDF later. As we shall see, the NDF is the most visually important feature of a microfacet model.

Microfacet Theory

- For given ω_i and ω_e, only facets oriented to reflect ω_i into ω_e are active
- These facets have normals at ω_h

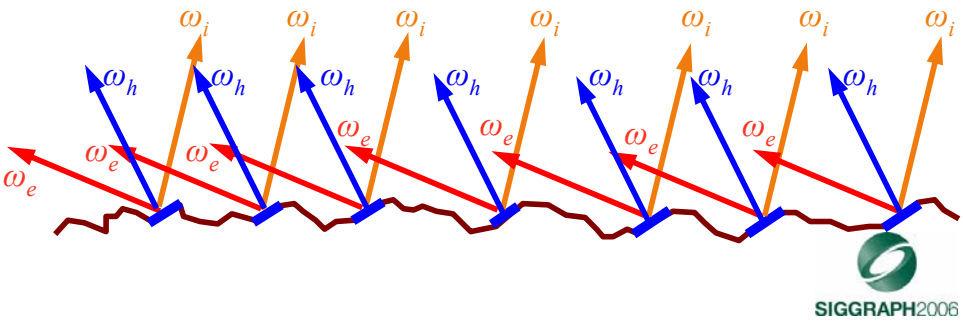

Since each facet is a perfectly smooth mirror, it has to be oriented exactly right to reflect ω_i into ω_e to participate in the reflectance at all. ω_h is the **halfway direction**, which is the direction exactly half-way between ω_i and ω_e.

Microfacet Theory

- Fraction of active microfacets is $p(\omega_h)d\omega$
- Active facets have reflectance $R_F(\alpha_h)$
- α_h is the angle between ω_i (or ω_e) and ω_h

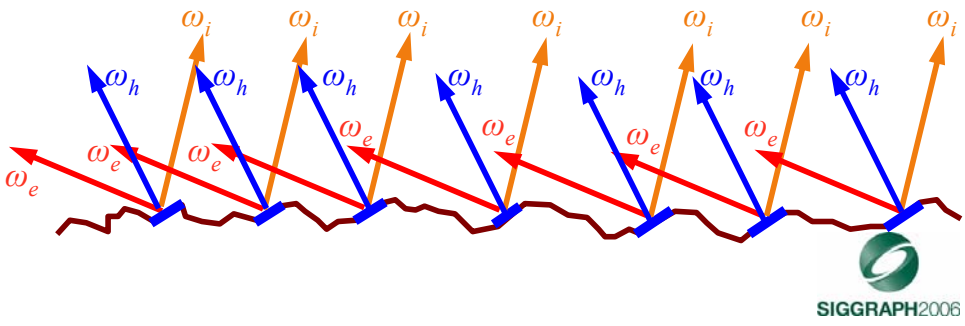

Only the microfacets with normals facing along ω_h can participate in the reflection. We can find the proportion of microfacets which have their normals in an infinitesimal patch of directions (solid angle $d\omega$) around ω_h by using the NDF. We assume the microfacets are themselves mirrors, so we can use Fresnel to get their reflectance.

Half-angle Vector

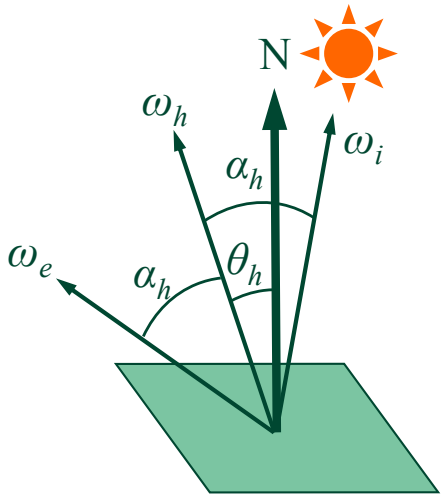

SIGGRAPH2006

Here is an illustration of some of the various angles relating to the halfway direction which appear in BRDFs.

Half-angle Vector

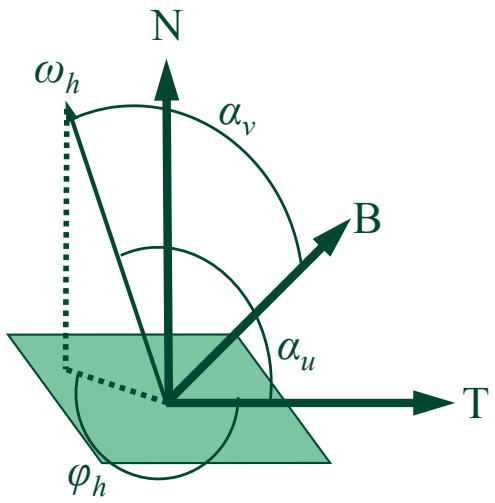

SIGGRAPH2006

Here are some more angles relating to the halfway direction, these are used in anisotropic BRDFs. In addition to the surface normal N which we have seen before, this diagram also includes the surface tangent T and bitangent B.

Microfacet Theory

$$f_r(\omega_i, \omega_e) = \frac{p(\omega_h) G(\omega_i, \omega_e) R_F(\alpha_h)}{4 K_p \cos\theta_i \cos\theta_e}$$

- $G(\omega_i,\omega_e)$ (the *geometry factor*) is the fraction of microfacets which are not masked or shadowed
- K_p is a constant with a value dependent on the microgeometry structure

From the previously seen relations and equations, we can derive this form for the microfacet BRDF (detailed derivations can be found in the literature). We will see a few different options for modeling the geometry factor later in this talk. Note that the geometry factor may contain parts that will cancel out or modify the cosine factors in the denominator.

Microfacet Theory

- $p(\omega)$ is the most important parameter
- Controls roughness, anisotropy, exact shape of highlight
- A variety of different functions can be used
- It can even be painted into a texture
  - Analogous to painting the highlight

The NDF is the most important parameter. We will learn more about hand-painted NDFs later in the course.

Isotropic Normal Distribution Functions

- "Phong" $p_{Ph}(\theta) = \frac{n+1}{2\pi}\cos^n\theta$
- Gaussian $p_G(\theta) = k_G e^{-(c_G\theta)^2}$
- Beckmann $p_B(\theta) = \frac{1}{m^2\cos^4\theta} e^{-\left\{\frac{\tan\theta}{m}\right\}^2}$
- Trowbridge-Reitz $p_{TR}(\theta) = k_{TR}\left(\frac{c_{TR}^2}{\cos^2\theta\left(c_{TR}^2-1\right)+1}\right)^2$

Here are some examples of isotropic normal distribution functions that appear in the literature. These include an NDF derived from the Phong reflectance model, the Gaussian distribution, as well as distributions from papers by Beckmann and Trowbridge and Reitz. Since they are isotropic, they are all parameterized by θ, the angle between the microscopic and microscopic surface normals. They each have parameters which control the smoothness of the surface (n for Phong, c_G for Gaussian, m for Beckmann and c_{TR} for Trowbridge-Reitz). In addition, the Gaussian and Trowbridge-Reitz NDFs have normalization factors (k_G and k_{TR} respectively) which are not shown here and need to be computed.

Microfacet vs. Reflection BRDFs

- Some BRDFs have a form superficially quite similar to microfacet BRDFs, but which is actually fundamentally different
- This difference is most apparent in the shape of highlights
- For reasons which will soon become clear, we call these “reflection BRDFs”

Microfacet BRDF

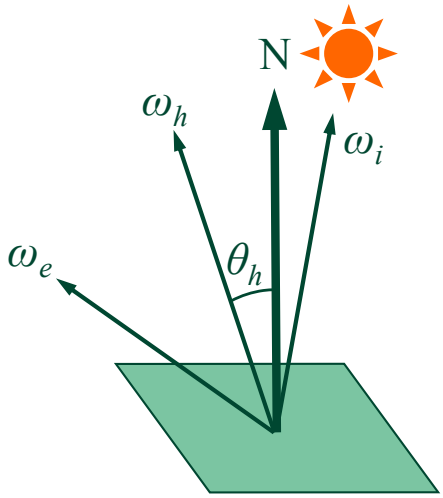

SIGGRAPH2006

Some BRDFs use a formulation which seems superficially similar to microfacet BRDFs, but which is actually significantly different.

Here we show again the halfway direction ω_h. The value of the microfacet normal distribution function in this direction gives us a first approximation of the reflectance, and determines the size and shape of the highlight. For an isotropic NDF, this is a function of the angle θ_h with the macroscopic surface normal, which is therefore an important BRDF parameter.

Reflection BRDF

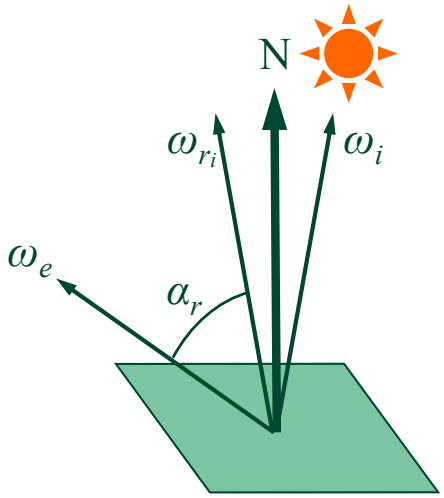

SIGGRAPH2006

Reflection BRDFs (such as the Phong BRDF) use the reflection direction ω_{r_i}, which is the incident direction ω_i reflected about the surface normal. Instead of θ_h, they are parameterized by α_r (the angle between ω_{r_i} and the exitant direction ω_e). This angle has no clear physical interpretation.

Reflection BRDF

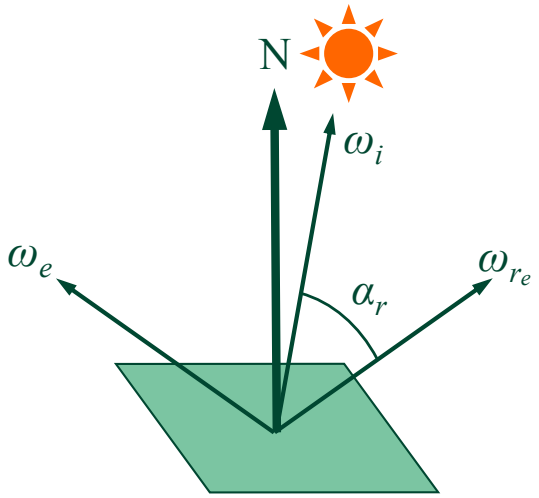

SIGGRAPH2006

Note that α_r is also sometimes described as the angle between ω_i and ω_{re}, the reflection of ω_e about the surface normal – the two descriptions are equivalent.

Center of Highlight

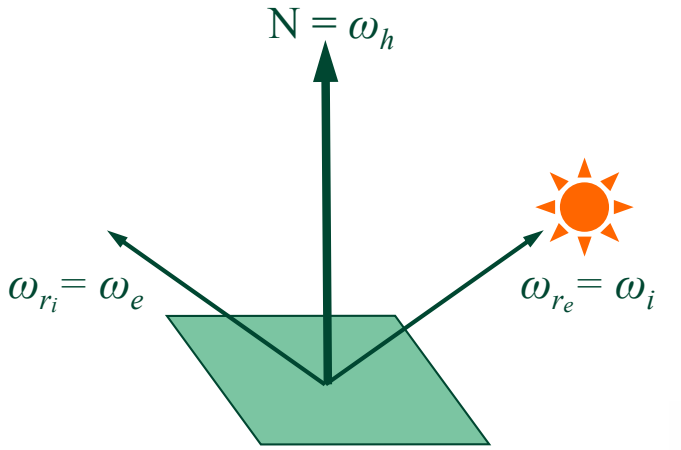

SIGGRAPH2006

Note that θ_h and α_r are equal to 0 in the same situation – when ω_e is equal to ω_{r_i} (or ω_i is equal to ω_{r_e}), then ω_h is at the surface normal. This means that both types of BRDF will place the center of the highlight at the same location. Physically, this is when the microfacets which are oriented the same as the overall surface (which is almost always the peak of the microfacet distribution) are oriented to reflect the incident light into the exitant direction. Although the center of the highlight will be in the same location, the two types of BRDF will not have the same shape of highlight.

Shape of Planar Highlight (Distant Light Source)

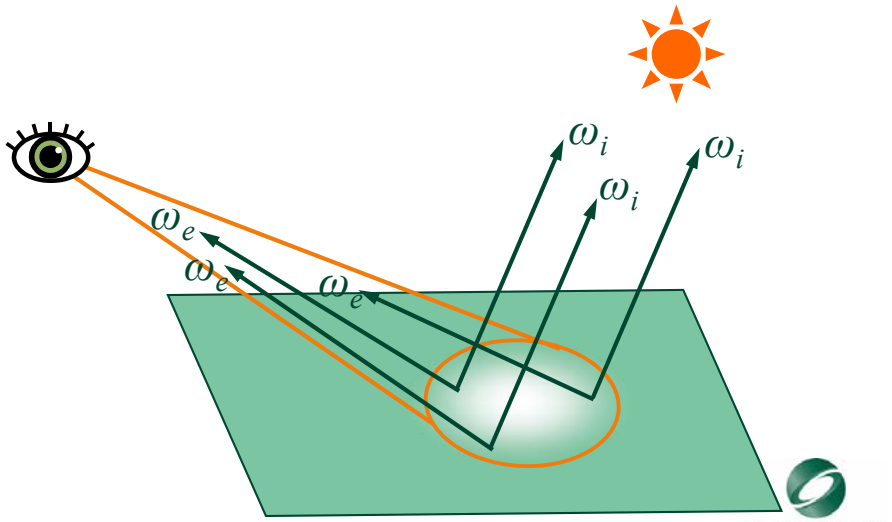

The shape of the highlight depends on various factors, including how the surface normal curves away from or towards the eye across the highlight. To simplify the comparison between the two we will look at the case where the highlight depends most simply and directly on the BRDF: a highlight of a planar surface. We will more specifically focus on the case of a distant light source to simplify the analysis, but a local light source is qualitatively similar.

Different points on the highlight have the same direction of incidence, but different directions of exitance. The highlight is defined by a cone from the eye, within which the BRDF takes high values. Since the surface is planar, the local frame is the same throughout, so this cone will be the mirror image of a cone…

Shape of Planar Highlight (Distant Light Source)

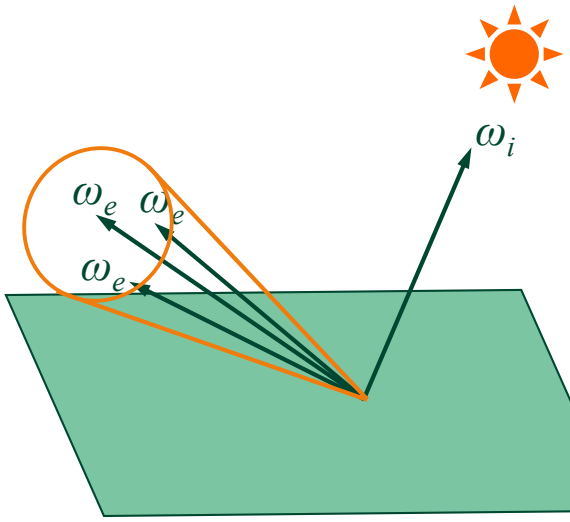

SIGGRAPH2006

Which is formed by the exitant directions. Since the two are mirror images of each other, they will have the same shape, which is the shape of the highlight.

Shape of Planar Highlight: Reflection BRDF

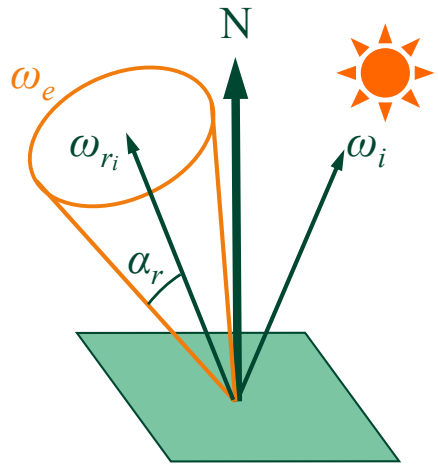

SIGGRAPH2006

With a reflection BRDF, regardless of the angle of incidence, exitant directions ω_e with equal values of α_r will form circular cones around ω_{r_i}. This means that highlights on planar surfaces with reflection BRDFs are circular.

Shape of Planar Highlight: Microfacet BRDF

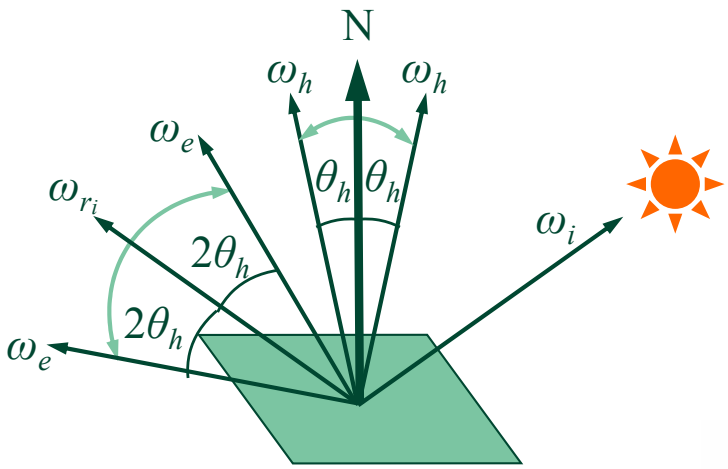

Let us now compare this circle with the set of exitant directions with a constant θ_h angle. With a microfacet BRDF, changing the halfway direction by θ_h in the plane formed by ω_i and the surface normal will cause the equivalent exitant directions to change by $2\theta_h$, regardless of the angle of incidence.

Shape of Planar Highlight: Microfacet BRDF

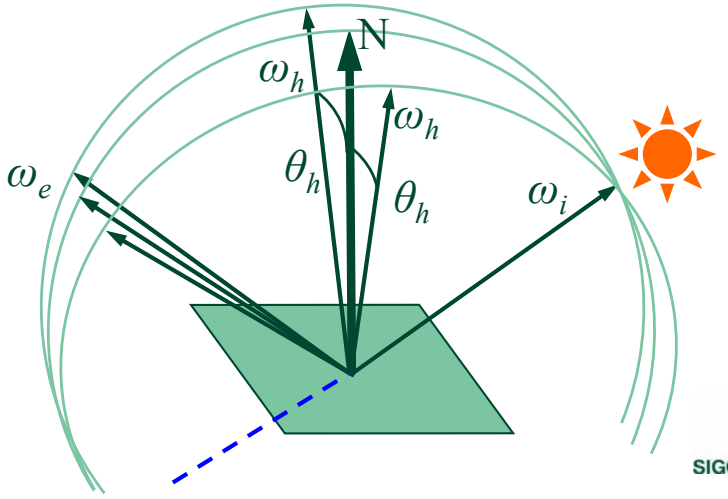

However, changing the halfway direction by θ_h in a direction perpendicular to the plane formed by ω_i and the surface normal will cause the equivalent exitant directions to change by a smaller angle. This angle will decrease as the angle of incidence increases. This can be visualized as rotating the pale green circular arc around the dashed blue axis (which is collinear with ω_i) – the points on the arc closer to the axis will move less.

Shape of Planar Highlight: Microfacet BRDF

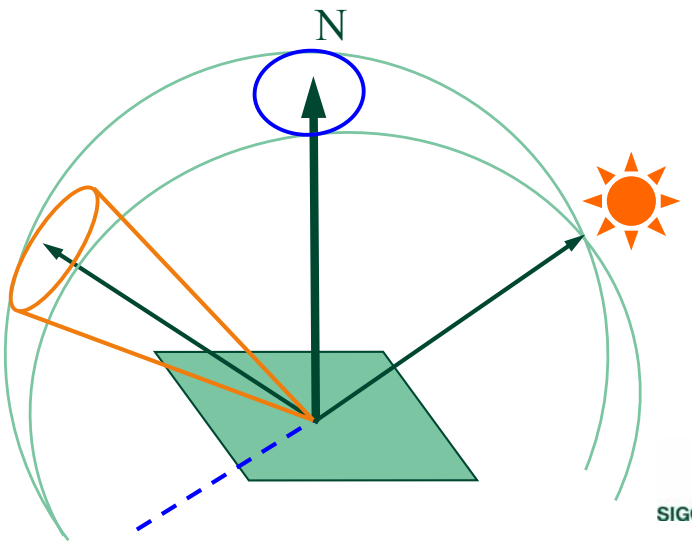

Since the variation in the angle of exitance remains the same in the plane of incidence, but decreases perpendicular to it with increasing incidence angle, this means that the highlights on planar surfaces are not circular, but narrow.

Shape of Planar Highlight: Microfacet BRDF

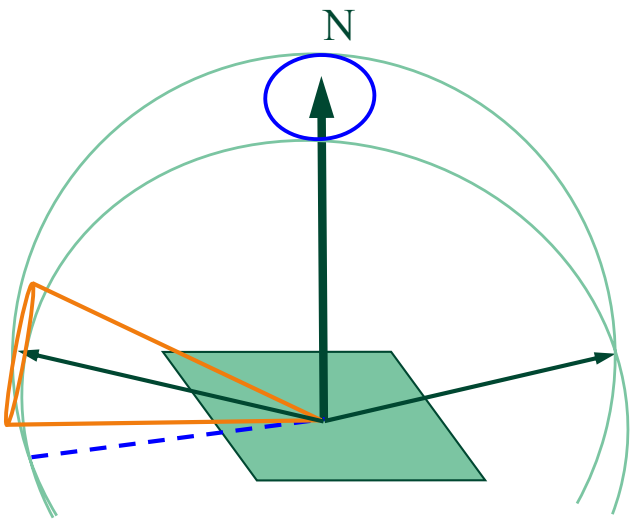

As the angle of incidence increases, the highlights become increasingly narrow. This is a very different visual prediction than that of a reflection BRDF.

Shape of Planar Highlight: Reality

IMAGES BY E. RUSKIN

As we can see from these photographs, reality matches the prediction of the microfacet BRDFs. Highlights at glancing angles off planar (or almost planar) surfaces are narrow and not circular.