
Shading

CSE 457, Autumn 2003

Graphics

<http://www.cs.washington.edu/education/courses/457/03au/>

Readings and References

Readings

- Sections 6.2-6.3, *3D Computer Graphics*, Watt

Other References

- OpenGL red book, chapter 5.
- Chapter 7, *3D Computer Graphics*, Watt

Introduction

- So far, we've talked exclusively about geometry.
 - » What is the shape of an object?
 - » How do I place it in a virtual 3D space?
 - » How do I know which pixels it covers?
 - » How do I know which of the pixels I should actually draw?
- Once we've answered all those, we have to ask one more important question:
 - » To what value do I set each pixel?
- Answering this question is the job of the **shading model**.
- Of course, people also call it a lighting model, a light reflection model, a local illumination model, a reflectance model, etc., etc.

An abundance of photons

- Properly determining the right color is *really hard*.
- Look around the room. Each light source has different characteristics. Trillions of photons are pouring out every second.
- These photons can:
 - » interact with the atmosphere, or with things in the atmosphere
 - » strike a surface and
 - be absorbed
 - be reflected
 - cause fluorescence or phosphorescence.
 - » interact in a wavelength-dependent manner
 - » generally bounce around and around


Our problem

- We're going to build up to an *approximation* of reality called the **Phong illumination model**.
 - » understand reality → approximate it → better understand reality in order to understand the differences → refine the approximation → ...
- The Phong illumination model has the following characteristics:
 - » *not* a model of the actual physics of light transport and reflection
 - » gives a first-order *approximation* to physical light reflection
 - » very fast
 - » widely used
- In addition, we will assume **local illumination**, i.e., light goes: light source -> surface -> viewer.
- No interreflections, no shadows.

Setup...

- Given:
 - » a point \mathbf{P} on a surface visible through pixel p
 - » The normal \mathbf{N} at \mathbf{P}
 - » The lighting direction, \mathbf{L} , and intensity, I_ℓ , at \mathbf{P}
 - » The viewing direction, \mathbf{V} , at \mathbf{P}
 - » The shading coefficients at \mathbf{P}
- Compute the color, I , of pixel p .
- Assume that the direction vectors are normalized:

$$\|\mathbf{N}\| = \|\mathbf{L}\| = \|\mathbf{V}\| = 1$$



Iteration zero

- The simplest thing you can do is...
- Assign each polygon a single color:

$$I = k_e$$

- where
 - » I is the resulting intensity
 - » k_e is the **emissivity** or intrinsic shade associated with the object
 - » this is the color regardless of available light, but it is *not* a light itself
- This has some special-purpose uses, but not really good for drawing a scene.
- [Note: k_e is omitted in Watt.]

Iteration one

- Let's make the color at least dependent on the overall quantity of light available in the scene:

$$I = k_e + k_a I_a$$

- » k_a is the **ambient reflection coefficient**.
 - really the reflectance of ambient light
 - “ambient” light is assumed to be equal in all directions
 - » I_a is the **ambient intensity**.
-
- In the real world, what is “ambient” light?

Wavelength dependence?

- Really, k_e , k_a , and I_a are functions over all wavelengths, λ , of light ... but we are crazed approximators ...
- Ideally, we would do the calculation on these functions. For example, our equation so far could be written as:

$$I(\lambda) = k_e(\lambda) + k_a(\lambda)I_a(\lambda)$$

then we would find good RGB values to represent the spectrum $I(\lambda)$.

- Traditionally, though, k_e , k_a and I_a are represented as RGB triples, and the computation is performed on each color channel separately:

$$I_R = k_{a,R} I_{a,R}$$

$$I_G = k_{a,G} I_{a,G}$$


$$I_B = k_{a,B} I_{a,B}$$

Diffuse reflection

- Let's examine the ambient shading model:
 - » objects have different colors
 - » we can control the overall light intensity
 - what happens when we turn off the lights?
 - what happens as the light intensity increases?
 - what happens if we change the color of the lights?
- So far, objects are uniformly lit.
 - » not the way things really appear
 - » in reality, light sources are directional
- **Diffuse**, or **Lambertian** reflection will allow reflected intensity to vary with the direction of the light.

Diffuse reflectors

- Diffuse reflection occurs from dull, matte surfaces, like latex paint, or chalk.
- These **diffuse** or **Lambertian** reflectors reradiate light equally in all directions.
- Picture a rough surface with lots of tiny **microfacets**.




Diffuse reflectors

...or picture a surface with little pigment particles embedded beneath the surface (neglect reflection at the surface for the moment):

The microfacets and pigments distribute light rays in all directions.


Embedded pigments are responsible for the coloration of diffusely reflected light in plastics and paints.



Note: these figures are intuitive, but not strictly (physically) correct and we are not trying to build a physics model of what is happening on surfaces like these.

Diffuse reflectors, cont.

- The reflected intensity from a diffuse surface does not depend on the direction of the viewer. The amount of available light, though, does depend on the direction of the light source:



$$A_p = A \cos \theta$$

$$\frac{A_p}{A} = \cos \theta$$

$$N \cdot L = \cos \theta$$

so $N \cdot L$ gives a scaling value for diffuse reflected light intensity

Iteration two

- The incoming energy is proportional to _____, giving the diffuse reflection equations:

$$\begin{aligned} I &= k_e + k_a I_a + k_d I_l (\quad) \\ &= k_e + k_a I_a + k_d I_l (\quad) \end{aligned}$$


- where:
 - » k_d is the **diffuse reflection coefficient**
 - » I_ℓ is the diffuse intensity of the light source
 - » \mathbf{N} is the normal to the surface (unit vector)
 - » \mathbf{L} is the direction to the light source (unit vector)
 - » $(x)_+$ means $\max \{0, x\}$

[Note: Watt uses I_i instead of I_ℓ .]

Specular reflection

- **Specular reflection** accounts for the highlight that you see on some objects.
- It is particularly important for *smooth, shiny* surfaces, such as:
 - » metal
 - » polished stone
 - » plastics
 - » apples
 - » skin
- **Properties:**
 - » Specular reflection depends on the viewing direction V .
 - » For non-metals, the color is determined solely by the color of the light.
 - » For metals, the color may be altered (e.g., brass)

Specular reflection “derivation”





- For a perfect mirror reflector, light is reflected about N , so

$$I = \begin{cases} I_\ell & \text{if } \mathbf{V} = \mathbf{R} \\ 0 & \text{otherwise} \end{cases}$$

- For a near-perfect reflector, you might expect the highlight to fall off quickly with increasing angle ϕ .
- Also known as:
 - » **“rough specular” reflection, “directional diffuse” reflection, “glossy” reflection**

Derivation, cont.



- One way to get this effect is to take $(\mathbf{R} \cdot \mathbf{V})$, raised to a power n_s .
- As n_s gets larger,
 - » the dropoff becomes {more,less} gradual
 - » gives a {larger,smaller} highlight
 - » simulates a {more,less} mirror-like surface

Iteration three

- The next update to the Phong shading model is then:

$$I = k_e + k_a I_a + k_d I_l (N \cdot L)_+ + k_s I_l (V \cdot R)_+^{n_s}$$

- where:
 - » k_s is the **specular reflection coefficient**
 - » n_s is the **specular exponent** or **shininess**
 - » \mathbf{R} is the reflection of the light about the normal (unit vector)
 - » \mathbf{V} is viewing direction (unit vector)
- [Note: Watt uses n instead of n_s .]

Intensity drop-off with distance

- OpenGL supports different kinds of lights: point, directional, and spot.
- For point light sources, the laws of physics state that the intensity of a point light source drops off inversely with the square of the distance.
- We can incorporate this effect by multiplying I_1 by $1/d^2$.
- Sometimes, this distance-squared dropoff is considered too “harsh.” A common alternative is:

$$f_{atten}(d) = \frac{1}{a + bd + cd^2}$$

- with user-supplied constants for a , b , and c .
- [Note: not discussed in Watt.]

Iteration four

- Since light is additive, we can handle multiple lights by taking the sum over every light.
- Our equation is now a version of the Phong illumination model:

$$I = k_e + k_a I_a + \sum_j f(d_j) I_{lj} \left[k_d (N \cdot L_j)_+ + k_s (V \cdot R_j)_+^{n_s} \right]$$

Which quantities are spatial vectors?

Which are RGB triples?

Which are scalars?

Choosing the parameters

- Experiment with different parameter settings. To get you started, here are a few suggestions:
 - » Try n_s in the range [0,128]
 - » Try $k_a + k_d + k_s < 1$
 - » Use a small k_a (~0.1)

	n_s	k_d	k_s
Polished metal	large	Small, color of metal	Large, color of metal
Plastic	medium	Medium, color of plastic	Medium, white
Planet	0	varying	0

Materials in OpenGL

- The OpenGL code to specify the surface shading properties is fairly straightforward. For example:

```
GLfloat ke[] = { 0.1, 0.15, 0.05, 1.0 };
GLfloat ka[] = { 0.1, 0.15, 0.1, 1.0 };
GLfloat kd[] = { 0.3, 0.3, 0.2, 1.0 };
GLfloat ks[] = { 0.2, 0.2, 0.2, 1.0 };
GLfloat ns[] = { 50.0 };
glMaterialfv(GL_FRONT, GL_EMISSION, ke);
glMaterialfv(GL_FRONT, GL_AMBIENT, ka);
glMaterialfv(GL_FRONT, GL_DIFFUSE, kd);
glMaterialfv(GL_FRONT, GL_SPECULAR, ks);
glMaterialfv(GL_FRONT, GL_SHININESS, ns);
```

The `GL_FRONT` parameter tells OpenGL that we are specifying the materials for the front of the surface.

Only the alpha value of the diffuse color is used for blending. It's usually set to 1.

Shading in OpenGL

- The OpenGL lighting model allows you to associate different lighting intensities according to material properties they will influence.
- Thus, our original shading equation:

$$I = k_e + k_a I_a + \sum_j f(d_j) I_{lj} \left[k_d (N \cdot L_j)_+ + k_s (V \cdot R_j)_+^{n_s} \right]$$

- becomes:

$$I = k_e + k_a I_a + \sum_j f(d_j) \left[k_a I_{laj} + k_d I_{ldj} (N \cdot L_j)_+ + k_s I_{lsj} (V \cdot R_j)_+^{n_s} \right]$$

$$I = k_e + k_a I_a + \sum_j f(d_j) \left[k_a I_{la_j} + k_d I_{ld_j} (N \cdot L_j)_+ + k_s I_{ls_j} (V \cdot R_j)_+^{n_s} \right]$$

- In OpenGL this equation is specified something like:

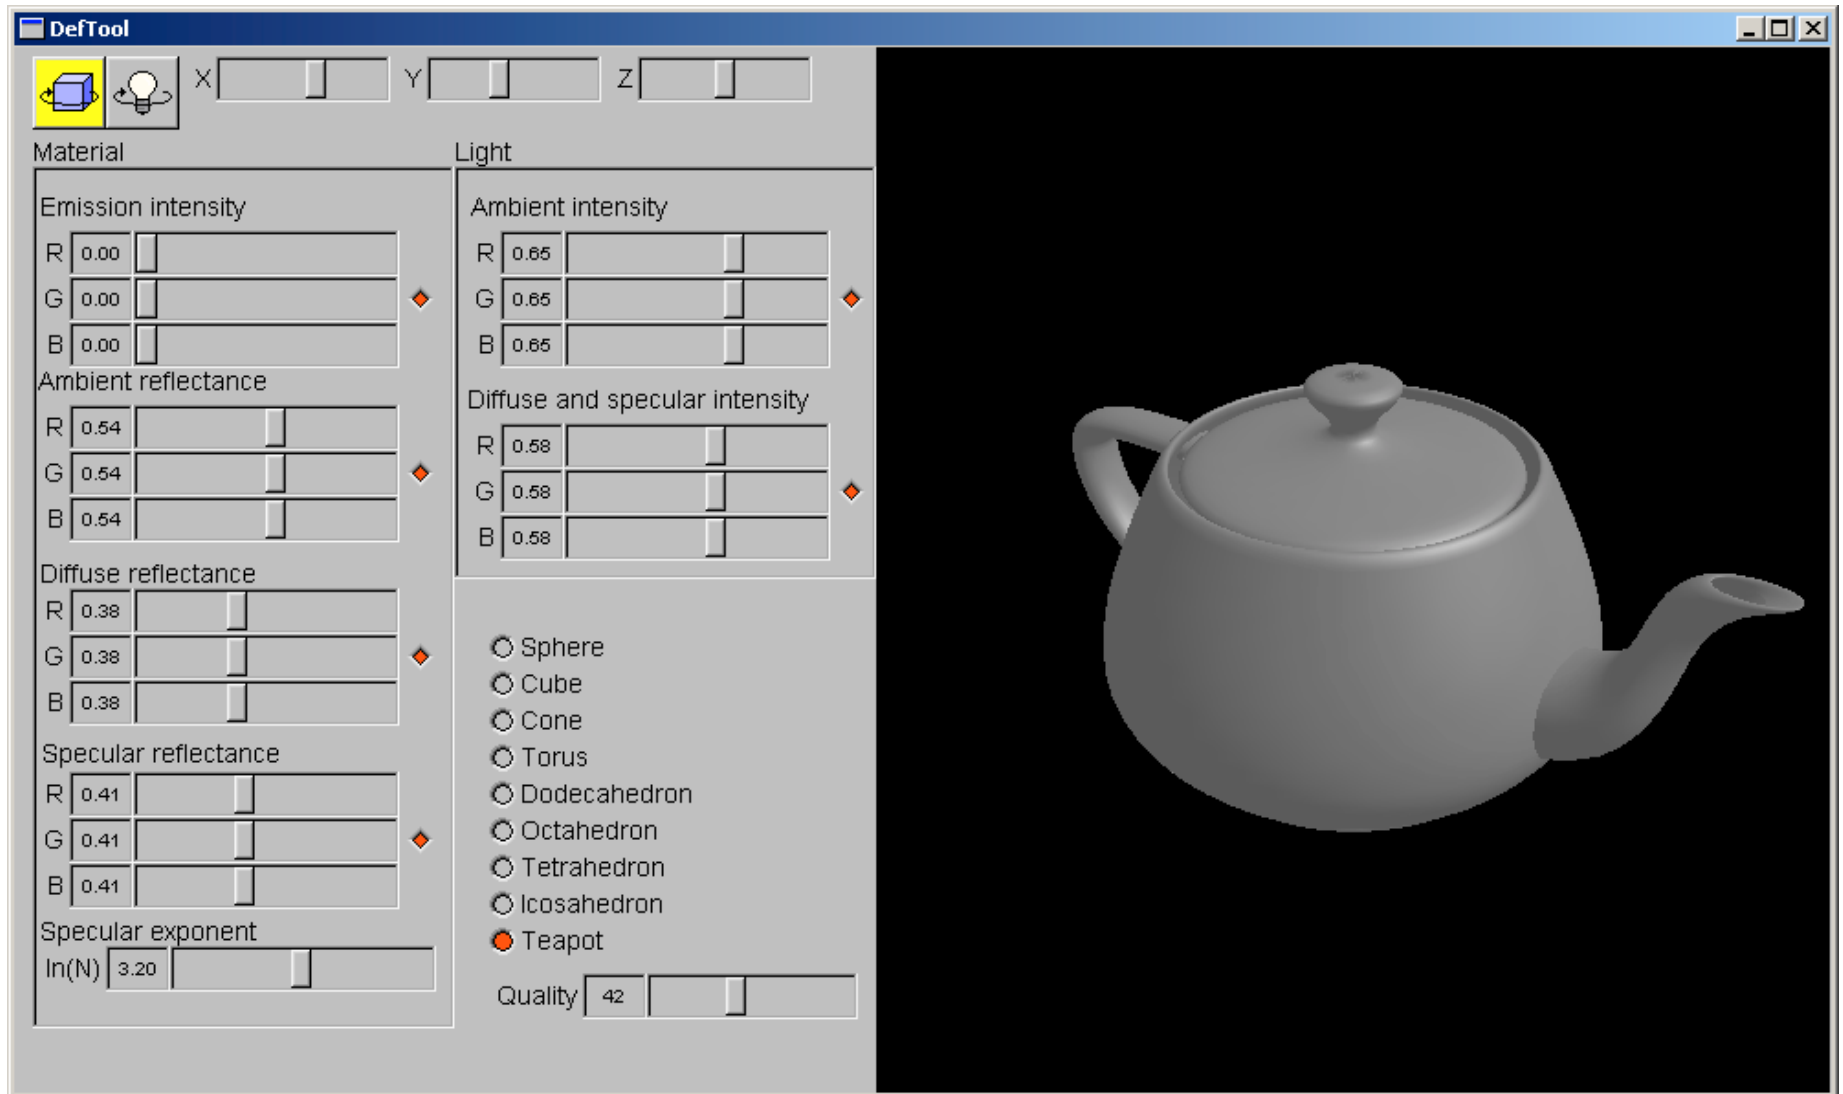
```

GLfloat Ia[] = { 0.2, 0.2, 0.2, 1.0 };
GLfloat Ia0[] = { 0.1, 0.1, 0.1, 1.0 };
GLfloat Id0[] = { 1.0, 1.0, 1.0, 1.0 };
GLfloat Is0[] = { 1.0, 1.0, 1.0, 1.0 };
GLfloat pos0[] = { 1.0, 1.0, 1.0, 0.0 };
GLfloat a[] = { 1.0 };
GLfloat b[] = { 0.5 };
GLfloat c[] = { 0.25 };

glLightModelfv(GL_LIGHT_MODEL_AMBIENT, Ia);
glLightfv(GL_LIGHT0, GL_AMBIENT, Ia0);
glLightfv(GL_LIGHT0, GL_DIFFUSE, Id0);
glLightfv(GL_LIGHT0, GL_SPECULAR, Is0);
glLightfv(GL_LIGHT0, GL_POSITION, pos0);
glLightfv(GL_LIGHT0, GL_CONSTANT_ATTENUATION, a);
glLightfv(GL_LIGHT0, GL_LINEAR_ATTENUATION, b);
glLightfv(GL_LIGHT0, GL_QUADRATIC_ATTENUATION, c);

```

- You can have as many as `GL_MAX_LIGHTS` lights in a scene. This number is system-dependent.



lighting.exe

BRDF

- The Phong illumination model is really a function that maps light from incoming (light) directions ω_{in} to outgoing (viewing) directions ω_{out} :

$$f_r(\omega_{in}, \omega_{out})$$


- This function is called the **Bi-directional Reflectance Distribution Function (BRDF)**.
- Here's a plot with ω_{in} held constant:



- BRDF's can be quite sophisticated...

Cook-Torrance-Sparrow model


One of the more common “sophisticated” BRDF’s is the Cook-Torrance-Sparrow model. It treats the surface as a set of mirrored facets with random orientations and decides whether or not light reflects from a given direction directly to the viewer.




The amount of light reflected from a facet is determined by the Fresnel coefficient, which depends in general on wavelength.



Cook and Torrance, 1982




(a) Phong model




(b) Torrance-Sparrow model

Fig. 16.43 Comparison of Phong and Torrance–Sparrow illumination models for light at a 30° angle of incidence. (By J. Blinn [BLIN77a], courtesy of the University of Utah.)




(a) Phong model



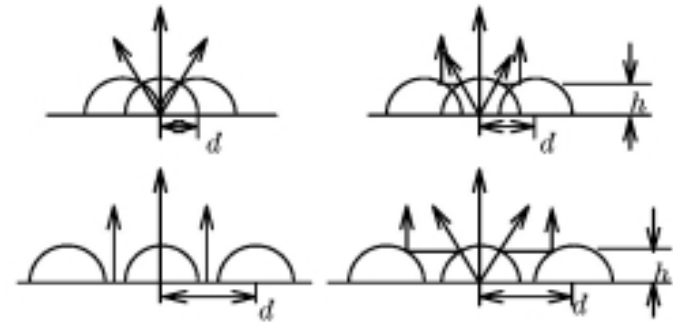
(b) Torrance-Sparrow model

Fig. 16.44 Comparison of Phong and Torrance–Sparrow illumination models for light at a 70° angle of incidence. (By J. Blinn [BLIN77a], courtesy of the University of Utah.)

Anisotropic reflection




Westin, Arvo, Torrance 1992



Poulin and Fournier 1990



Weird BRDF: the moon




Gouraud vs. Phong interpolation

- Now we know how to compute the color at a point on a surface using the Phong lighting model.
- Does graphics hardware do this calculation at every point? Typically not...
- Smooth surfaces are often approximated by polygonal facets, because:
 - » Graphics hardware generally wants polygons (esp. triangles).
 - » Sometimes it easier to write ray-surface intersection algorithms for polygonal models.
- How do we compute the shading for such a surface?

Faceted shading

- Assume each face has a constant normal:




- For a distant viewer and a distant light source, how will the color of each triangle vary?
- Result: faceted, not smooth, appearance.

Gouraud interpolation

To get a smoother result that is easily performed in hardware, we can do **Gouraud interpolation**.


Here's how it works:

- Compute normals at the vertices.
- *Shade* only at the vertices.
- *Interpolate* the resulting vertex colors to shade the polygon.



Gouraud interpolation, cont'd

- Gouraud interpolation has significant limitations.
 - » If the polygonal approximation is too coarse, we can miss specular highlights.




We will encounter Mach banding (derivative discontinuity enhanced by human eye).

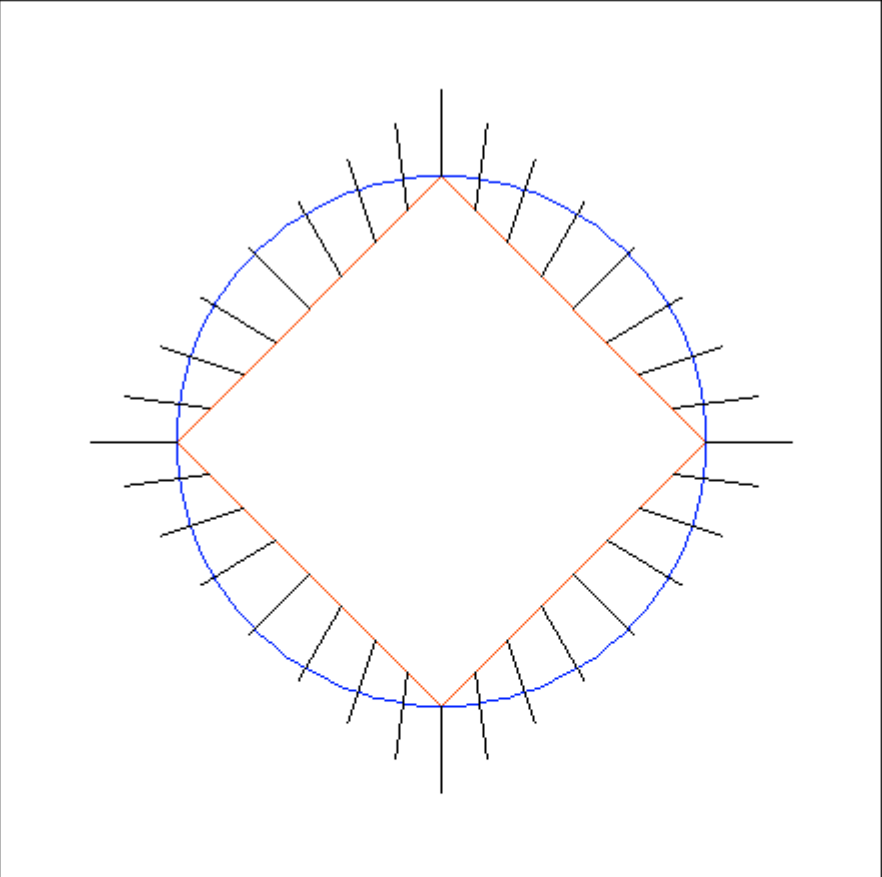
- ◆ Alas, this is usually what graphics hardware supports.
- ◆ Maybe someday soon we'll get...

Phong interpolation

- To get an even smoother result with fewer artifacts, we can perform **Phong interpolation**.
- Here's how it works:
 - Compute normals at the vertices.
 - *Interpolate* normals and normalize.
 - *Shade* using interpolated normals at each point.



Normal vectors



- Show curve
- Show curve normals

- Show faces
 - Show face normals
- 4

- Show vertex normals
- Show interpolated normals

Summary

- The most important thing to take away from this lecture is the final equation for the Phong model.
 - » What is the physical meaning of each variable?
 - » How are the terms computed?
 - » What effect does each term contribute to the image?
 - » What does varying the parameters do?
- You should also understand the differences between faceted, Gouraud, and Phong interpolated shading.