

10. Shading

Reading

- ♦ Watt, sections 4.1-4.6

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.)

Phun with 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, ad nauseum

Our problem

We're going to build up to an *approximation* of reality called the **Phong illumination model**.

It has the following characteristics:

- ♦ *not* physically based
- ♦ 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.

Iteration zero

Given:

- ♦ a point P on a surface
- ♦ visible through pixel p

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 has some special-purpose uses, but not really good for drawing a scene.

Iteration one

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

$$I = k_a I_a$$

- ♦ 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**.

Physically, what is "ambient" light?

Wavelength dependence

Really, k_e , k_a , and I_a are functions over all wavelengths λ .

Ideally, we would do the calculation on these functions. For the ambient shading equation, we would start with:

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

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

Traditionally, though, 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: the figures above are intuitive, but not strictly (physically) correct.

Diffuse reflectors, cont.

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



Iteration two

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

$$I = k_e + k_a I_a + k_d I_l$$

$$= k_e + k_a I_a + k_d I_l (\mathbf{N} \cdot \mathbf{L})_+$$

where:

- ◆ k_d is the **diffuse reflection coefficient**
- ◆ I_l is the 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 $\max\{0, x\}$

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

Properties:

- ◆ Specular reflection depends on the viewing direction \mathbf{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 \mathbf{N} , so


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

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} glossy surface

Iteration three

The next update to the Phong shading model is then:

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

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)

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 must drop off inversely with the square of the distance.

We can incorporate this effect by multiplying I_i by $1/d^2$.

Sometimes, this distance-squared dropoff is considered too "harsh." A common alternative is:

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

with user-supplied constants for a , b , and c .

Iteration four

Since light is additive, we can handle multiple lights by taking the sum over every light.

Our equation is now:

$$I = k_e + k_a I_a + \sum_i f(d_i) I_{li} [k_d (\mathbf{N} \cdot \mathbf{L}_i)_+ + k_s (\mathbf{V} \cdot \mathbf{R}_i)_+^{n_s}]$$

This is the Phong illumination model.

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 [0,100]
- ♦ Try $k_a + k_d + k_s < 1$
- ♦ Use a small k_d (~0.1)

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

Choosing the parameters



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? Unfortunately not...


Smooth surfaces are often approximated by polygonal facets, because:

- ◆ Graphics hardware generally wants polygons (esp. triangles).
- ◆ Sometimes it's 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:


1. Compute normals at the vertices.
2. Shade only the vertices.
3. Interpolate the resulting vertex colors.



Gouraud interpolation, cont'd

Gouraud interpolation has significant limitations.

1. If the polygonal approximation is too coarse, we can miss specular highlights.



2. We will encounter Mach banding.

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:

1. Compute normals at the vertices.
2. Interpolate normals and normalize.
3. Shade using the 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.