

Shading

**Brian Curless
CSEP 557
Spring 2019**

Reading

Optional:


- ◆ Angel and Shreiner: chapter 5.
- ◆ Marschner and Shirley: chapter 10, chapter 17.

Further reading:


- ◆ OpenGL red book, chapter 5.

Basic 3D graphics

With affine matrices, we can now transform virtual 3D objects in their local coordinate systems into a global (world) coordinate system:




To synthesize an image of the scene, we also need to add light sources and a viewer/camera:



Pinhole camera

To create an image of a virtual scene, we need to define a camera, and we need to model lighting and shading. For the camera, we use a **pinhole camera**.



The image is rendered onto an **image plane** (usually in front of the camera).

Viewing rays emanate from the **center of projection** (COP) at the center of the pinhole.

The image of an object point **P** is at the intersection of the viewing ray through **P** and the image plane.

But is **P** visible? This is the problem of **hidden surface removal** (a.k.a., **visible surface determination**). We'll consider this problem later.


Shading

Next, we'll need a model to describe how light interacts with surfaces.

Such a model is called a **shading model**.

Other names:

- ◆ Lighting model
- ◆ Light reflection model
- ◆ Local illumination model
- ◆ Reflectance model
- ◆ BRDF



An abundance of photons

Given the camera and shading model, properly determining the right color at each pixel is *extremely hard*.

Look around the room. Each light source has different characteristics. Trillions of photons are pouring out every second.

These photons can:

- ◆ interact with molecules and particles in the air (“participating media”)
- ◆ strike a surface and
 - be absorbed
 - be reflected (scattered)
 - cause fluorescence or phosphorescence.
- ◆ interact in a wavelength-dependent manner
- ◆ generally bounce around and around

Our problem

We're going to build up to a *approximations* of reality called the **Phong and Blinn-Phong illumination models**.

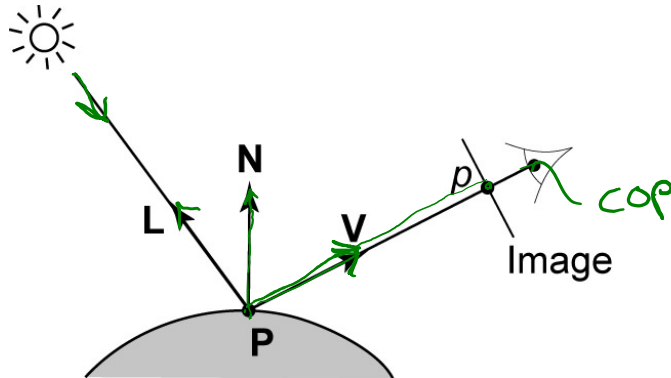
They have the following characteristics:

- ◆ *not* physically correct
- ◆ 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 (color) 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 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:

$$I = k_e + k_a I_{La}$$

- ◆ k_a is the **ambient reflection coefficient**.
 - really the reflectance of ambient light
 - “ambient” light is assumed to be equal in all directions
- ◆ I_{La} is the **ambient light intensity**.

Physically, what is “ambient” light?

poor person's interreflection

Wavelength dependence

Really, k_e , k_a , and I_{La} are functions over all wavelengths λ .

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

$$I(\lambda) = k_a(\lambda) I_{La}(\lambda)$$




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

Traditionally, though, k_a and I_{La} are represented as RGB triples, and the computation is performed on each color channel separately:

$$I^R = k_a^R I_{La}^R$$

$$I^G = k_a^G I_{La}^G$$

$$I^B = k_a^B I_{La}^B$$




Diffuse reflectors

Emissive and ambient reflection don't model realistic lighting and reflection. To improve this, we will look at **diffuse** (a.k.a., **Lambertian**) reflection.

Diffuse reflection can occur from dull, matte surfaces, like latex paint, or chalk.


These diffuse 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 in this and the previous slide 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:



$$\cos\theta = \frac{dA_p}{dA} \quad \theta = 0^\circ$$

$$dA_p = dA \cos\theta$$

$$I \sim B \cdot \cos\theta = B \cdot \hat{N} \cdot \hat{L}$$

$$B = \begin{cases} 1 & \text{if } \cos\theta \geq 0 \\ 0 & \text{else} \end{cases}$$

“Iteration two”

The incoming energy is proportional to $\cos\theta$, giving the diffuse reflection equations:

$$I = k_e + k_a I_{La} + k_d I_L B \cos\theta$$
$$= k_e + k_a I_{La} + k_d I_L B (N \cdot L)$$

where:

- ◆ k_d is the **diffuse reflection coefficient**
- ◆ I_L is the (color) 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)
- ◆ B prevents contribution of light from below the surface:

$$B = \begin{cases} 1 & \text{if } N \cdot L > 0 \\ 0 & \text{if } N \cdot L \leq 0 \end{cases}$$

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


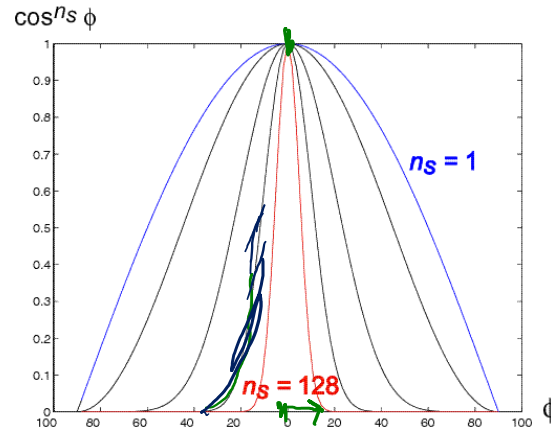
$$I = \begin{cases} I_L & \text{if } \mathbf{V} = \mathbf{R} \\ 0 & \text{otherwise} \end{cases}$$

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

Phong specular reflection



One way to get this effect is to take $(\mathbf{R} \cdot \mathbf{V})$, raised to a power n_s .

Phong specular reflection is proportional to:

$$I_{\text{specular}} \sim B(\mathbf{R} \cdot \mathbf{V})_+^{n_s}$$

where $(x)_+ \equiv \max(0, x)$.


Q: As n_s gets larger, does the highlight on a curved surface get smaller or larger?

Blinn-Phong specular reflection

A common alternative for specular reflection is the **Blinn-Phong model** (sometimes called the **modified Phong model**.)

We compute the vector halfway between \mathbf{L} and \mathbf{V} as:

$$\mathbf{H} = \frac{\frac{1}{2}(\mathbf{L} + \mathbf{V})}{\|\frac{1}{2}(\mathbf{L} + \mathbf{V})\|} = \frac{\mathbf{L} + \mathbf{V}}{\|\mathbf{L} + \mathbf{V}\|}$$



Analogous to Phong specular reflection, we can compute the specular contribution in terms of $(\mathbf{N} \cdot \mathbf{H})$, raised to a power n_s :

$$I_{\text{specular}} \sim B(\mathbf{N} \cdot \mathbf{H})_+^{n_s}$$

where, again, $(x)_+ \equiv \max(0, x)$.

“Iteration three”

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

$$\begin{aligned} I &= k_e + k_a I_{La} + k_d I_L B(\mathbf{N} \cdot \mathbf{L}) + k_s I_L B(\mathbf{N} \cdot \mathbf{H})^{n_s} \\ &= k_e + k_a I_{La} + I_L B \left[k_d (\mathbf{N} \cdot \mathbf{L}) + k_s (\mathbf{N} \cdot \mathbf{H})^{n_s} \right] \end{aligned}$$

where:


- ◆ k_s is the **specular reflection coefficient**
- ◆ n_s is the **specular exponent** or **shininess**
- ◆ \mathbf{H} is the unit halfway vector between \mathbf{L} and \mathbf{V} , where \mathbf{V} is the viewing direction.


Directional lights

The simplest form of lights supported by renderers are ambient, directional, and point. Spotlights are also supported often as a special form of point light.

We've seen ambient light sources, which are not really geometric.

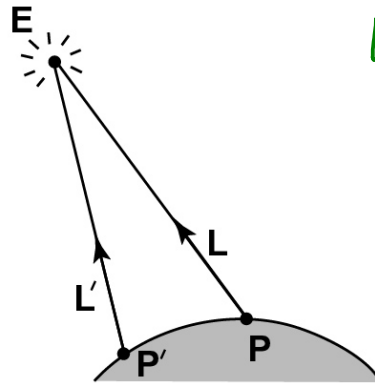
Directional light sources have a single direction and intensity associated with them.



Using affine notation, what is the homogeneous coordinate for a directional light? 

Point lights

The direction of a **point light** sources is determined by the vector from the light position to the surface point.



$$L = \frac{E - P}{\|E - P\|}$$

$$r = \|E - P\|$$

Physics tells us the intensity must drop off inversely with the square of the distance:

$$f_{\text{atten}} = \frac{1}{r^2}$$

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

$$f_{\text{atten}} = \frac{1}{ar^2 + br + c}$$


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

Using affine notation, what is the homogeneous coordinate for a point light?



Spotlights

We can also apply a *directional attenuation* of a point light source, giving a **spotlight** effect.



A common choice for the spotlight intensity is:

$$f_{\text{spot}} = \begin{cases} \frac{(\mathbf{L} \cdot \mathbf{S})^e}{ar^2 + br + c} & \alpha \leq \beta \\ 0 & \text{otherwise} \end{cases}$$

where

- ♦ \mathbf{L} is the direction to the point light.
- ♦ \mathbf{S} is the center direction of the spotlight.
- ♦ α is the angle between \mathbf{L} and \mathbf{S}
- ♦ β is the cutoff angle for the spotlight
- ♦ e is the angular falloff coefficient

Note: $\alpha \leq \beta \Leftrightarrow \cos^{-1}(\mathbf{L} \cdot \mathbf{S}) \leq \beta \Leftrightarrow \mathbf{L} \cdot \mathbf{S} \geq \cos \beta$.

“Iteration four”

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

Our equation is now (for spotlight lighting):

Use k_d instead of k_a

$$I = k_e + \sum_j k_a I_{L,j} + \frac{(\mathbf{L}_j \cdot \mathbf{S}_j)^{e_j}}{a_j r_j^2 + b_j r_j + c_j} I_{L,j} B_j \left[k_d (\mathbf{N} \cdot \mathbf{L}_j) + k_s (\mathbf{N} \cdot \mathbf{H}_j)^{n_s} \right]$$

The equation includes several handwritten annotations: a green arrow points from the text "Use k_d instead of k_a " to the k_a term in the summation; a green arrow points from the text k_d to the k_d term in the spotlight equation; red arrows point to k_e , $I_{L,j}$, B_j , k_d , and k_s ; red wavy lines underline the terms $k_d (\mathbf{N} \cdot \mathbf{L}_j)$ and $k_s (\mathbf{N} \cdot \mathbf{H}_j)^{n_s}$ in the spotlight equation.

This is the Blinn-Phong illumination model (for spotlights). Note that, in practice, we usually set $k_a = k_d$.


Which quantities are spatial vectors?

Which are RGB triples?

Which are scalars?

3D Geometry in the Graphics Hardware Pipeline


Graphics hardware applies transformations to bring the objects and lighting into the camera's coordinate system:



The geometry is assumed to be made of triangles, and the **vertices** are projected onto the image plane.

Rasterization


After projecting the vertices, graphics hardware “smears” vertex properties across the interior of the triangle in a process called **rasterization**.



Smearing the z-values and using a Z-buffer will enable the graphics hardware to determine if a point inside a triangle is visible. (More on this in another lecture.)

If we have stored colors at the vertices, then we can smear these as well.

Hardware Pipeline




A vertex shader is run for each vertex, and outputs values to be interpolated across the triangle.


The vertices are grouped into triangles (or other primitives, e.g. lines) to be rasterized. A geometry shader is possibly run to generate more primitives.

We iterate through scanlines, interpolating outputs from the vertex shader at each pixel.

A fragment shader (or pixel shader) is called at each pixel in the primitive, which gets the interpolated values and outputs a final color to the framebuffer.




GLSL: A Simple Vertex Shader



```
in vec3 position;  
in vec3 vertex_color;  
  
out vec3 color;  
  
uniform mat4 modelview;  
uniform mat4 projection;  
  
void main() {  
    color = vertex_color;  
    gl_Position = projection * modelview * vec4(position, 1.0);  
}
```


GLSL: A Simple Fragment Shader



```
in vec3 color;  
  
out vec4 frag_color;  
  
void main() {  
    frag_color = vec4(color, 1.0);  
}
```

GLSL: Storage Qualifiers

uniform: Global value that is the same across all vertices and fragments (for this draw call).

Vertex shader

- **in**: Per-vertex attributes (that were sent to the GPU)
- **out**: Values to be interpolated by the rasterizer and then passed to a fragment shader

Fragment shader

- **in**: Interpolated values of vertex shader **out**'s
- **out**: Value to be written to frame buffer

uniform's may include model/view/projection matrices, light parameters, material parameters, textures...

in's and **out**'s may include normals, positions, colors, material parameters, texture coordinates...

Shading the interiors of triangles

We will be computing colors using the Blinn-Phong lighting model.

Let's assume (as graphics hardware does) that we are working with triangles.


How should we shade the interiors of triangles?

We will consider this over the next few slides...

Per-face normals for triangle meshes

We will be shading and calculating reflections and refractions based on surface normals.

For a triangle mesh, we can make the natural assumption that each triangle has a constant normal (the normal of its supporting plane):



Recall the Blinn-Phong shading equation for a single light source (no ambient or emissive):



$$I = I_L B \left[k_d (\mathbf{N} \cdot \mathbf{L}) + k_s (\mathbf{N} \cdot \mathbf{H})_+^{n_s} \right]$$

Typically, \mathbf{L} and \mathbf{V} vary only a small amount over each triangle, if at all.

Q: If material properties (k_d, k_s, n_s) are constant over the mesh, how will shading vary within a triangle?

It won't

Faceted shading (cont'd)




[Williams and Siegel 1990]

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.



Faced shading vs. Gouraud interpolation



[Williams and Siegel 1990]


G
1.1

Gouraud interpolation artifacts


~~A~~

Gouraud interpolation has significant limitations.


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



$$f(x) - f''(x)$$




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



This is what graphics hardware does by default.

A substantial improvement is to do...




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.




Gouraud vs. Phong interpolation




[Williams and Siegel 1990]

Old pipeline: Gouraud interpolation



Default vertex processing:


$\mathbf{L} \leftarrow$ determine lighting direction
 $\mathbf{V} \leftarrow$ determine viewing direction
 $\mathbf{N} \leftarrow \text{normalize}(\mathbf{n}_e)$
 $c_{\text{blinn-phong}} \leftarrow$ shade with $\mathbf{L}, \mathbf{V}, \mathbf{N}, k_d, k_s, n_s$
 $v_i \leftarrow$ project v to image
out $c_{\text{blinn-phong}}$
out v_i



Default fragment processing:

$\text{color} \leftarrow c_{\text{blinn-phong}}^p$

Programmable pipeline: Phong-interpolated normals!



Vertex shader:

determine eye, normal, vertex in world coordinates

$v_i \leftarrow \text{project } v \text{ to image}$


out eye_w

out \mathbf{n}_w

out v_w

out v_i

$v_i^1, v_i^2, v_i^3 \rightarrow \text{triangle}$



Fragment shader:

$\mathbf{L} \leftarrow \text{determine lighting direction (using } v_w^p)$

$\mathbf{V} \leftarrow \text{normalize}(v_w - v_e^p)$

$\mathbf{N} \leftarrow \text{normalize}(\mathbf{n}_w^p)$

color $\leftarrow \text{shade with } \mathbf{L}, \mathbf{V}, \mathbf{N}, k_d, k_s, n_s$

Choosing Blinn-Phong shading 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_a (~ 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

BRDF

For more physical correctness, we would also weight the specular part by $\mathbf{N} \cdot \mathbf{L}$:


$$\begin{aligned}
 I &= I_L B \left[k_d (\mathbf{N} \cdot \mathbf{L}) + k_s (\mathbf{N} \cdot \mathbf{L}) \left(\mathbf{N} \cdot \frac{\mathbf{L} + \mathbf{V}}{\|\mathbf{L} + \mathbf{V}\|} \right)_+^{n_s} \right] \\
 &= I_L B (\mathbf{N} \cdot \mathbf{L}) \left[k_d + k_s \left(\mathbf{N} \cdot \frac{\mathbf{L} + \mathbf{V}}{\|\mathbf{L} + \mathbf{V}\|} \right)_+^{n_s} \right] \\
 &= I_L B (\mathbf{N} \cdot \mathbf{L}) f_r(\mathbf{L}, \mathbf{V})
 \end{aligned}$$

The function f_r maps incoming (light) directions ω_{in} to outgoing (viewing) directions ω_{out} :

$$f_r(\omega_{in}, \omega_{out}) \quad \text{or} \quad f_r(\omega_{in} \rightarrow \omega_{out})$$


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...


More sophisticated BRDF's



[Cook and Torrance, 1982]




Anisotropic BRDFs [Westin, Arvo, Torrance 1992]




Artistics BRDFs [Gooch]

More sophisticated BRDF's (cont'd)

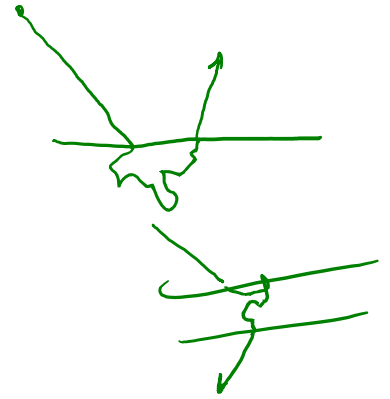


Hair illuminated from different angles [Marschner et al., 2003]



Wool cloth and silk cloth [Irawan and Marschner, 2012]

BSSRDFs for subsurface scattering



[Jensen et al., 2001]

Summary

You should understand the equation for the Blinn-Phong lighting model described in the “Iteration Four” slide:

- ◆ 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.