

CMSC 427

Computer Graphics¹

David M. Mount
Department of Computer Science
University of Maryland
Fall 2017

¹Copyright, David M. Mount, 2017 Dept. of Computer Science, University of Maryland, College Park, MD, 20742. These lecture notes were prepared by David Mount for the course CMSC 427, Computer Graphics, at the University of Maryland. Permission to use, copy, modify, and distribute these notes for educational purposes and without fee is hereby granted, provided that this copyright notice appear in all copies.

Lecture 1: More on Ray Tracing: Reflection and Refraction

Ray Tracing: We continue our discussion of ray tracing. Recall that this is a powerful method for rendering highly realistic images. Unlike OpenGL, it implements a global model for image generation, based on tracing the rays of light, mostly working backwards from the eye to the light sources.

Last time we discussed how rays are represented, how to generate rays based on the camera setup, and how this can be used in the context of object picking in interactive computer graphics. We also discussed how to compute the intersection of a ray with a plane. Today we consider how to intersect a ray with a sphere, and other issues such as how to handle reflection and refraction.

Ray-Sphere Intersection: Let us consider how to solve the intersection of a ray with a sphere. Suppose that the sphere is represented by giving its center point c and its radius $r > 0$ (see Fig. ??(a)). Recall that the ray $R : (p, \vec{u})$ is represented by the origin point p of the ray and the unit directional vector $\vec{u}$. Any point on the ray can be described as $p + t\vec{u}$ for some $t > 0$. The intersection problem involves determining whether the ray hits the object and if so, what is the value t of the intersection point q (see Fig. ??(c)). For the sake of lighting, we also would like to have the surface normal vector $\vec{n}$ at the contact point.




Fig. 1: Ray-sphere intersection.

We know that a point q lies on the sphere if its distance from the center of the sphere is r , that is if $\|q - c\| = r$. So the ray intersects at the value of t such that

$$\|(p + t\vec{u}) - c\| = r.$$

Notice that the quantity inside the $\|\cdot\|$ above is a vector. Let $\vec{w} = c - p$. This gives us

$$\|t\vec{u} - \vec{w}\| = r.$$

We know $\vec{u}$, $\vec{w}$, and r and we want to find t . By the definition of length using dot products we have

$$(t\vec{u} - \vec{w}) \cdot (t\vec{u} - \vec{w}) = r^2.$$

Observe that this equation is scalar valued (not a vector). We use the fact that dot-product is a linear operator, and so we can manipulate this algebraically into:

$$t^2(\vec{u} \cdot \vec{u}) - 2t(\vec{u} \cdot \vec{w}) + (\vec{w} \cdot \vec{w}) - r^2 = 0$$

This is a quadratic equation $at^2 + bt + c = 0$, where

$$\begin{aligned} a &= (\vec{u} \cdot \vec{u}) = 1 && \text{(since } \vec{u} \text{ is normalized),} \\ b &= -2(\vec{u} \cdot \vec{w}), \\ c &= (\vec{w} \cdot \vec{w}) - r^2 \end{aligned}$$

We can solve this using the quadratic formula to produce two roots. Let $\Delta = b^2 - 4ac$ denote the *determinant*. If $\Delta < 0$, then there is no root. Otherwise, using the fact that $a = 1$ we have:

$$\begin{aligned} t^- &= \frac{-b - \sqrt{b^2 - 4ac}}{2a} = \frac{-b - \sqrt{\Delta}}{2}, \\ t^+ &= \frac{-b + \sqrt{b^2 - 4ac}}{2a} = \frac{-b + \sqrt{\Delta}}{2}. \end{aligned}$$

These two values reflect the fact that the ray may hit the sphere twice (see Fig. ??(b)). The general rule in ray tracing is to *use the smaller of the two positive roots*. Thus, if $t^- > 0$ we use t^- to define the intersection point (as indicated in the figure). Otherwise, if $t^+ > 0$ we use t^+ . (What does this imply about the origin of the ray?) If both are nonpositive, then there is no intersection. (Note that there are two ways that there may be no solution, either this or because the determinant is negative. These have two different geometric interpretations. What are they?)

Note that it is not a good idea to compare floating point numbers against zero, since floating point errors are always possible. A good rule of thumb is to do all of these 3-d computations using doubles (not floats) and perform comparisons against some small value instead, e.g. “const double TINY = 1E-3”. The proper choice of this parameter is a bit of “magic”. It is usually adjusted until the final image looks okay.

(Note that this is not the most numerically accurate method for computing the intersection point. However, for most applications of ray tracing, the approximation is close enough that any errors are not noticeable to the human eye. A book on numerical analysis will discuss more accurate methods for solving the quadratic equation.)

Normal Vector: In the case of the sphere, that the normal vector is directed from the center of the sphere to point of contact. Thus, if t is the parameter value at the point of contact, the normal vector is just

$$\vec{n} = \text{normalize}(p + t\vec{u} - c).$$

Note that this vector is directed outwards. If t^- was used to define the intersection, then we are hitting the object from the outside, and so $\vec{n}$ is the desired normal. However, if t^+ was used to define the intersection, then we are hitting the object from the inside, and $-\vec{n}$ should be used instead.

Reflection: Recall the basic recursive structure to ray tracing. Shoot a ray and determine the color of the first object hit. Next, shoot a ray to the light sources to determine which of these sources illuminate this point. (If you hit another object before arriving at the light source then you are in the shadow cast by this object.) Based on the surface normal vector, apply some lighting model (e.g., the Phong model) to compute the shading at this point.

If the object hit is reflective or refractive, we next shoot additional rays (recursively) and determine their colors. (Note that lighting will be determined for these points at their points of contact, so these colors are not applied to lighting on the surface of the reflective object.) Finally, we blend the various colors together to obtain the final surface color, which is returned.

How is this blending done? Recall that in the Phong reflection model each object is associated with a color, and its coefficients of ambient, diffuse, and specular reflection, denoted ρ_a , ρ_d and ρ_s . To model the reflective component, each object will be associated with an additional parameter called the *coefficient of reflection*, denoted ρ_r . As with the other coefficients this is typically a number in the interval $[0, 1]$. Let us assume that this coefficient is nonzero. We compute the view reflection ray (which equalizes the angle between the surface normal and the view vector). Let $\vec{v}$ denote the normalized *view vector*, which points backwards along the viewing ray (see Fig. ??). Thus, if the ray is $p + t\vec{u}$, then $\vec{v} = -\text{normalize}(\vec{u})$. (This is essentially the same as the view vector used in the Phong model, but it may not point directly back to the eye because of intermediate reflections.) Let $\vec{n}$ denote the outward pointing surface *normal vector*, which we assume is also normalized. The normalized *view reflection vector* (see Fig. ??), denoted $\vec{r}_v$, is derived as follows

$$\vec{r}_v = 2(\vec{n} \cdot \vec{v})\vec{n} - \vec{v}$$




Fig. 2: Reflection.

Since the surface is reflective, we shoot the ray emanating from the surface contact point along this direction and apply the above ray-tracing algorithm recursively. Eventually, when the ray hits a nonreflective object, the resulting color is returned. This color is then factored into the Phong model, as will be described below. Note that it is possible for this process to go into an infinite loop, if say you have two mirrors facing each other. To avoid such looping, it is common to have a maximum recursion depth, after which some default color is returned, irrespective of whether the object is reflective.

Transparent objects and refraction: To model *refraction*, also called *transmission*, we maintain a coefficient of transmission, denoted ρ_t . We also need to associate each surface with two additional parameters, the *indices of refraction*² for the incident side η_i and the transmitted side, η_t . Recall from physics that the index of refraction is the ratio of the speed of

²To be completely accurate, the index of refraction depends on the wavelength of light being transmitted. This is what causes white light to be spread into a spectrum as it passes through a prism, which is called *chromatic dispersion*. Since we do not model light as an entire spectrum, but only through a triple of RGB values (which produce the same color visually, but not the same spectrum physically) it is not easy to model this phenomenon. For simplicity we assume that all wavelengths have the same index of refraction.

light through a vacuum versus the speed of light through the material. Typical indices of refraction include:

Material	Index of Refraction
Air (vacuum)	1.0
Water	1.333
Glass	1.5
Diamond	2.47

Snell's law says that if a ray is incident with angle θ_i (relative to the surface normal) then it will be transmitted with angle θ_t (relative to the opposite normal) such that

$$\frac{\sin \theta_i}{\sin \theta_t} = \frac{\eta_t}{\eta_i}$$

Let us work out the direction of the transmitted ray from this. As before let $\vec{v}$ denote the normalized view vector, directed back along the incident ray. Let $\vec{t}$ denote the unit vector along the transmitted direction, which we wish to compute (see Fig. ??(a)).




Fig. 3: Refraction. ((b) and (c) show the conditions under which total internal reflection might occur.)

The orthogonal projection of $\vec{v}$ onto the normalized normal vector $\vec{n}$ is

$$\vec{m}_i = (\vec{v} \cdot \vec{n})\vec{n} = (\cos \theta_i)\vec{n}.$$

Consider the two parallel horizontal vectors $\vec{w}_i$ and $\vec{w}_t$ in the figure. We have

$$\vec{w}_i = \vec{m}_i - \vec{v}.$$

Since $\vec{v}$ and $\vec{t}$ are each of unit length we have

$$\frac{\eta_t}{\eta_i} = \frac{\sin \theta_i}{\sin \theta_t} = \frac{\|\vec{w}_i\|/\|\vec{v}\|}{\|\vec{w}_t\|/\|\vec{t}\|} = \frac{\|\vec{w}_i\|}{\|\vec{w}_t\|}.$$

Since $\vec{w}_i$ and $\vec{w}_t$ are parallel we have

$$\vec{w}_t = \frac{\eta_i}{\eta_t} \vec{w}_i = \frac{\eta_i}{\eta_t} (\vec{m}_i - \vec{v}).$$

The projection of $\vec{t}$ onto $-\vec{n}$ is $\vec{m}_t = -(\cos \theta_t) \vec{n}$, and hence the desired refraction vector is:

$$\begin{aligned} \vec{t} &= \vec{w}_t + \vec{m}_t = \frac{\eta_i}{\eta_t} (\vec{m}_i - \vec{v}) - (\cos \theta_t) \vec{n} = \frac{\eta_i}{\eta_t} ((\cos \theta_i) \vec{n} - \vec{v}) - (\cos \theta_t) \vec{n} \\ &= \left(\frac{\eta_i}{\eta_t} \cos \theta_i - \cos \theta_t \right) \vec{n} - \frac{\eta_i}{\eta_t} \vec{v}. \end{aligned}$$

We have already computed $\cos \theta_i = (\vec{v} \cdot \vec{n})$. We can derive $\cos \theta_t$ from Snell's law and basic trigonometry:

$$\begin{aligned} \cos \theta_t &= \sqrt{1 - \sin^2 \theta_t} = \sqrt{1 - \left(\frac{\eta_i}{\eta_t} \right)^2 \sin^2 \theta_i} = \sqrt{1 - \left(\frac{\eta_i}{\eta_t} \right)^2 (1 - \cos^2 \theta_i)} \\ &= \sqrt{1 - \left(\frac{\eta_i}{\eta_t} \right)^2 (1 - (\vec{v} \cdot \vec{n})^2)}. \end{aligned}$$

What if the term in the square root is negative? This is possible if $(\eta_t/\eta_i) < \sin \theta_i$. In particular, this can only happen if $\eta_i > \eta_t$, meaning that you are already inside an object with an index of refraction greater than 1. Notice that when this is the case, Snell's law breaks down, since it is impossible to find θ_t whose sine is greater than 1. In this situation, *total internal reflection* takes place (see Fig. ??(c)). That is, the light source is not refracted at all, but is reflected back onto the incident side. (By the way, this phenomenon, combined with chromatic dispersion, is one of the reasons for the existence of rainbows.) When this happens, the refraction reduces to reflection and so we set $\vec{t} = \vec{r}_v$, the view reflection vector.

In summary, the transmission process is solved as follows.

- (1) Compute the point where the ray intersects the surface. Let $\vec{v}$ be the normalized view vector, let $\vec{n}$ be the normalized surface normal at this point, and let η_i and η_t be the indices of refraction on the incoming and outgoing sides, respectively.
- (2) Compute the angle of refraction:

$$\theta_t = \arccos \sqrt{1 - \left(\frac{\eta_i}{\eta_t} \right)^2 (1 - (\vec{v} \cdot \vec{n})^2)}.$$

If the quantity under the square root symbol is negative, process this as internal reflection, rather than transmission.

- (4) If the quantity under the square root symbol is nonnegative, compute the transmission vector

$$\vec{t} = \left(\frac{\eta_i}{\eta_t} \cos \theta_i - \cos \theta_t \right) \vec{n} - \frac{\eta_i}{\eta_t} \vec{v}.$$

The transmission ray is emitted from the contact point along this direction.