

Ray-tracing

1. Introduction

Ray-tracing is a method for depicting three-dimensional objects realistically. It is computationally expensive, because it requires a separate computation for each pixel on the screen.

The ray-tracing technique can be used to handle many objects at once and many lighting sources. The objects can reflect light like a mirror (specular reflection), or diffusely like paper (diffuse reflection), or they can be transparent and refract (bend) light like glass, or a combination. The light can come from nearby sources or sources at infinity, it can be of any combination of colors. A perspective view is usually used.

2. The setup

To keep things simple, we shall concentrate on a single convex object with a diffusely reflective surface, illuminated by a single light source of monochrome light at infinity. The corresponding setup is shown in Figure 1. We take the viewpoint to be $\begin{bmatrix} 0 \\ 0 \\ H \end{bmatrix}$ (with $H > 0$). The object is shown on the side of the viewplane with $z < 0$, but it could actually be anywhere with $z < H$. The light source is at infinity with direction given by the unit vector \mathbf{s} .

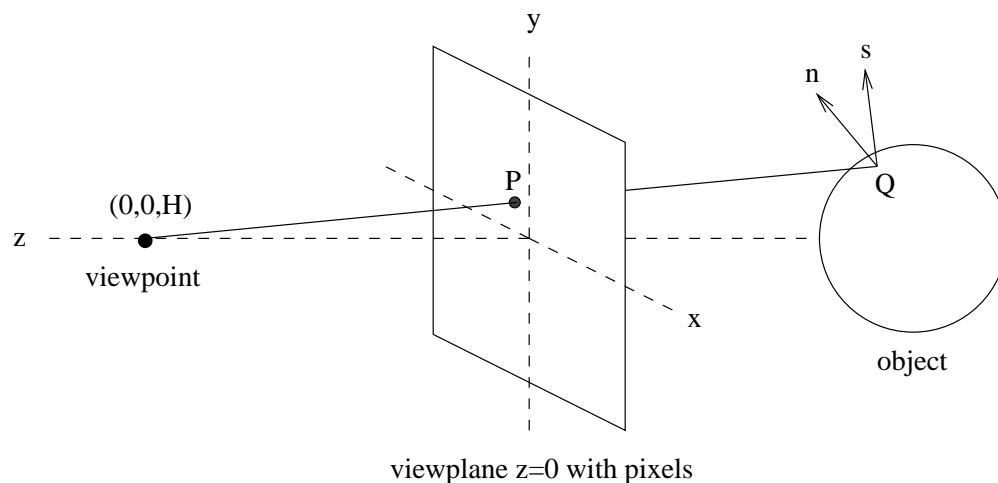


Figure 1: Setup for ray-tracing

Physically, light travels from the light source to the object and then to the viewpoint (the eye). This will not work computationally, however, since we

would have to trace a great number of rays, only a few of which would reach the eye. A much easier procedure is to trace the rays in the opposite direction, one for each pixel on the screen. The lines involved are the same either way!

2.1 *Method*: For each pixel position $P = \begin{bmatrix} a \\ b \\ 0 \end{bmatrix}$ in the viewplane (screen):

Step 1. Find the equation of the line (ray) from the viewpoint $\begin{bmatrix} 0 \\ 0 \\ H \end{bmatrix}$ to

$$P = \begin{bmatrix} a \\ b \\ 0 \end{bmatrix}. \text{ This will be } \mathbf{x} = (1-t) \begin{bmatrix} 0 \\ 0 \\ H \end{bmatrix} + t(a, b, 0), \text{ which simplifies to}$$

$$\mathbf{x} = \begin{bmatrix} at \\ bt \\ (1-t)H \end{bmatrix}.$$

Step 2. Find the point Q where this line first intersects the object, if it does at all. (If it does not intersect the object, then the brightness of the pixel is 0; stop here.)

Step 3. Find the outward unit normal \mathbf{n} at Q .

Step 4. Use “Phong shading”: The brightness of the pixel at P is $\mathbf{n} \cdot \mathbf{s}$, representing a fraction of the maximum possible brightness, provided $\mathbf{n} \cdot \mathbf{s} \geq 0$; otherwise the brightness is 0. This formula does not involve the direction of the ray itself, because diffuse reflection scatters light equally in all directions.

2.2 *Example*

Suppose the viewpoint is $\begin{bmatrix} 0 \\ 0 \\ 4 \end{bmatrix}$, the lighting direction is $\mathbf{s} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$ (straight up, which is the positive y direction), the object is the sphere of radius 3 centered at $\begin{bmatrix} 0 \\ 0 \\ -3 \end{bmatrix}$, and we are interested in the pixel at $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$, so $P =$

$\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$. The equation of the sphere is then

$$(x - 0)^2 + (y - 0)^2 + (z - (-3))^2 = 3^2, \text{ or } x^2 + y^2 + (z + 3)^2 = 9.$$

In Step 1, the line is $\mathbf{x} = \begin{bmatrix} t \\ t \\ 4 - 4t \end{bmatrix}$.

In Step 2, we simply substitute $x = t$, $y = t$, $z = 4 - 4t$ into the equation for the sphere. We get $t^2 + t^2 + (7 - 4t)^2 = 9$. Expanding and collecting terms we get $18t^2 - 56t + 40 = 0$, or after division by 2 to simplify, $9t^2 - 28t + 20 = 0$.

By the quadratic formula, the roots are $t = \frac{-(-28) \pm \sqrt{(-28)^2 - 4 \cdot 9 \cdot 20}}{2 \cdot 9}$
 $= \frac{14 \pm \sqrt{14^2 - 180}}{9} = \frac{14 \pm 4}{9}$, so $t = \frac{10}{9}$ or $t = 2$.

Because the ray starts at the viewpoint with $t = 0$, the first time it reaches the sphere is at $t = \frac{10}{9}$, which corresponds to the point $Q = \begin{bmatrix} \frac{10}{9} \\ \frac{10}{9} \\ -\frac{4}{9} \end{bmatrix}$.

In Step 3, notice that the equation for the sphere has the form $F(x, y, z) = c$, a “level surface” or “isosurface” of F . One normal is therefore the gradient, $\nabla F = \begin{bmatrix} 2x \\ 2y \\ 2(z + 3) \end{bmatrix}$, which at Q is $\begin{bmatrix} \frac{20}{9} \\ \frac{20}{9} \\ \frac{46}{9} \end{bmatrix} = \frac{2}{9} \begin{bmatrix} 10 \\ 10 \\ 23 \end{bmatrix}$. Because $F(x, y, z)$ for the sphere grows from the center out, the gradient gives an outward normal.

Now we need to normalize the gradient to get the *unit* outward normal vector \mathbf{n} . It is harmless to scale first by a positive scalar, so let’s leave off the factor of $\frac{2}{9}$ and use the outward normal $\begin{bmatrix} 10 \\ 10 \\ 23 \end{bmatrix}$. The length of this vector is

$\sqrt{10^2 + 10^2 + 23^2} = \sqrt{729} = 27$, so the *unit* outward normal is $\mathbf{n} = \begin{bmatrix} \frac{10}{27} \\ \frac{10}{27} \\ \frac{23}{27} \end{bmatrix}$.

In Step 4, $\mathbf{n} \cdot \mathbf{s} = \begin{bmatrix} \frac{10}{27} \\ \frac{10}{27} \\ \frac{23}{27} \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} = \frac{10}{27}$. Again, this is the fraction of the maximum possible brightness.

3. Variations

- If there are several objects, it is necessary to see where the ray hits them all, mathematically, and then take Q to be the first hit (lowest t). If the surface at Q is diffusely reflective, you have to trace the ray from Q in the direction of the light source to see if hits another object (so Q would be in a shadow).
- If one or more objects are not convex, then it is possible that $\mathbf{n} \cdot \mathbf{s}$ could be positive at Q and yet the ray from Q to the light source could hit another part of the same object, so you need to check for that too.

- If the surface of the object is specular-reflective (mirror-like) at Q , then it is necessary to bounce the ray off it by using the law of reflection (which involves the ray direction and \mathbf{n}). Then follow that ray for possible reflections off other objects, and so on, until you get to a diffusely reflective surface.
- A transparent medium such as glass partly reflects and partly refracts. The proportions of each depend on the angle of incidence. It is necessary to trace both the reflected ray by the method just described and the refracted ray by using Snell's Law, which involves the ray direction, optical density of the medium, and \mathbf{n} . Then keep going with each of the two rays. The refracted ray inside the medium might hit the surface again at an angle that produces total internal reflection, or it might be partly reflected and partly refracted again and then go on to hit other objects, and so on—all this to find the brightness of one pixel!
- If the surfaces are described parametrically instead of being easy objects like spheres defined by relational equations $F(x, y, z) = c$, then finding where the ray hits the surface can be much harder.
- If the lighting source is at a finite point, the method is the same except that \mathbf{s} will depend on Q . You may wish also to include a factor for the fact that brightness of a light source is inversely proportional to the square of the distance it is away.
- The methods described so far produce strong shadows. Often an ambient light source is added as well, meaning a light source that illuminates all objects uniformly, with no particular direction, the way light bouncing off white walls might. (“Ambient” means “surrounding”. In calculus you may have seen Newton's law of cooling, which can be described as telling how the temperature of a hot object approaches the “ambient temperature”, meaning the air temperature of the surrounding room.)
- If the light source is in color, then you can consider its red, green, and blue components separately and treat each as if they were monochrome.
- The Phong formula can also include any additional constant factors you want, say representing the fact that a diffusely reflecting surface may not be fully reflective. (Some surfaces could be darker than others.)
- If an object has color, in the Phong formula you can introduce an extra factor that depends on color (red, green, or blue) and expresses the property that some colors are reflected more than others.

- Phong's original formula allows a power of $\mathbf{n} \cdot \mathbf{s}$, say $(\mathbf{n} \cdot \mathbf{s})^r$, where r is a constant ≥ 1 . Because $\mathbf{n} \cdot \mathbf{s}$ is between 0 and 1, a higher power has the effect of a sharper falloff of brightness for normal directions not pointing near the light source.
- Some books use a setup with the z coordinate axis running the other way, so that the viewpoint is on the negative z axis. Such a coordinate system is left-handed, but the math is the same idea as described in §2. Another possibility is to put the viewpoint at the origin and the viewplane at $z = H$; in this case the ray from the viewpoint to P has the easy equation $\mathbf{x} = \frac{t}{P}$.

4. Problems

Problem S-1. Suppose everything is exactly as in Example 2.2, except that the sphere is centered at $\begin{bmatrix} 0 \\ 0 \\ -5 \end{bmatrix}$ (while still being of radius 3). For the same P , what is the brightness of the pixel there?

Problem S-2. Using Phong shading, find the brightness of the pixel at $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ on the viewplane $z = 0$ as seen from the viewpoint $\begin{bmatrix} 0 \\ 0 \\ 2 \end{bmatrix}$, if the surface is the sphere of radius 5 centered at $\begin{bmatrix} -1 \\ 0 \\ -6 \end{bmatrix}$ and the light source is in the direction $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$.

Problem S-3. Suppose everything is exactly as in Example 2.2, except that the object is a plane piece of paper with equation $x + y + z = -3$. For $P = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$ as before, what is the brightness of the pixel there? Does it depend on the choice of P ? (For the normal, there isn't any "outward" direction; instead, choose the direction that is on the same side of the paper as the light source.)

Problem S-4. Let the viewpoint be at the origin, let the viewplane be the $z = 10$ plane, and let $\mathbf{s} = \frac{1}{3} \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}$. Suppose the object is the sphere of radius 10 centered at $\begin{bmatrix} 0 \\ 0 \\ 20 \end{bmatrix}$. What is the brightness of the pixel at $\begin{bmatrix} 3 \\ 4 \\ 10 \end{bmatrix}$? (Or is it in shadow?)

Problem S-5. In Example 2.2, consider the same ray but suppose the sphere is shiny (specular reflection). Find parametric equations for the reflected ray.

Method: For the reflected ray, think separately about its direction and about a point through which it goes; then use the point-direction parametric form of a line. For the direction, first find a direction vector for the incoming ray. A fact, perhaps to be derived later in the course, is that the matrix of a reflection in a mirror through the origin with unit normal \mathbf{n} (as a column vector) is $I - 2\mathbf{nn}^t$; note that the column times a row is 3×3 ; and so a column vector \mathbf{v} has reflected image $(I - 2\mathbf{nn}^t)\mathbf{v} = \mathbf{v} - 2\mathbf{nn}^t\mathbf{v} = \mathbf{v} - 2\mathbf{n}(\mathbf{n} \cdot \mathbf{v}) = \mathbf{v} - 2(\mathbf{n} \cdot \mathbf{v})\mathbf{n}$. Use this information to reflect the path direction vector in the mirror through the origin with appropriate normal, obtaining the direction of the reflected ray. (Since this part of the calculation is with vectors rather than points, you don't have to worry about where the mirror is, but only the direction of its normal.)