

Distribution Ray Tracing

**Brian Curless
CSE 557
Fall 2013**

Reading

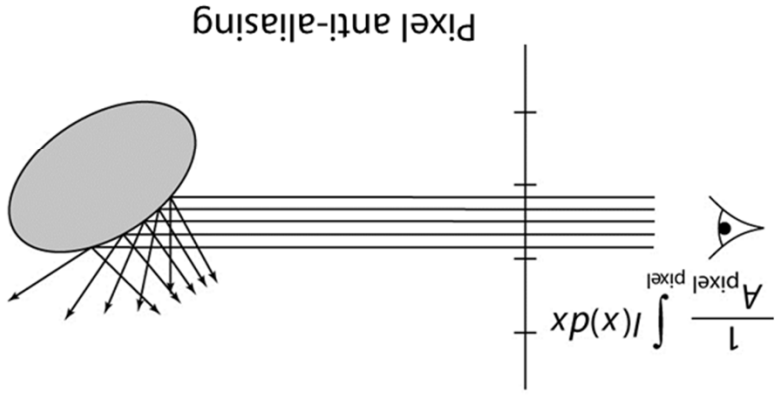

Required:

- ◆ Shirley, 13.11, 14.1-14.3

Further reading:

- ◆ A. Glassner. An Introduction to Ray Tracing. Academic Press, 1989. [In the lab.]
- ◆ Robert L. Cook, Thomas Porter, Loren Carpenter.
"Distributed Ray Tracing." Computer Graphics (Proceedings of SIGGRAPH 84). 18 (3). pp. 137-145. 1984.
- ◆ James T. Kajiya. "The Rendering Equation." Computer Graphics (Proceedings of SIGGRAPH 86). 20 (4). pp. 143-150. 1986.

Pixel anti-aliasing



All of this assumes that inter-reflection behaves in a mirror-like fashion...

BRDF

The diffuse+specular parts of the Blinn-Phong illumination model are a mapping from light to viewing directions:


$$I = I_B \left[k_p (\mathbf{N} \cdot \mathbf{L}) + k_s \left(\mathbf{N} \cdot \frac{\mathbf{L} + \mathbf{V}}{\|\mathbf{L} + \mathbf{V}\|} \right)^+ \right] = I_f f_r(\mathbf{L}, \mathbf{V})$$

The mapping function f_r is often written in terms of incoming (light) directions ω_{in} and 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...

Light reflection with BRDFs

BRDF's exhibit Helmholtz reciprocity:

$$f_r(\omega_{in} \rightarrow \omega_{out}) = f_r(\omega_{out} \rightarrow \omega_{in})$$

That means we can take two equivalent views of reflection. Suppose $\omega_{in} = \mathbf{L}$ and $\omega_{out} = \mathbf{V}$:




We can now think of the BRDF as weighting light coming in from all directions, which can be added up:

$$I(\mathbf{V}) = \int_I f_r(\mathbf{L} \rightarrow \mathbf{V}) (\mathbf{L} \cdot \mathbf{N}) d\mathbf{L}$$

Or, written more generally:

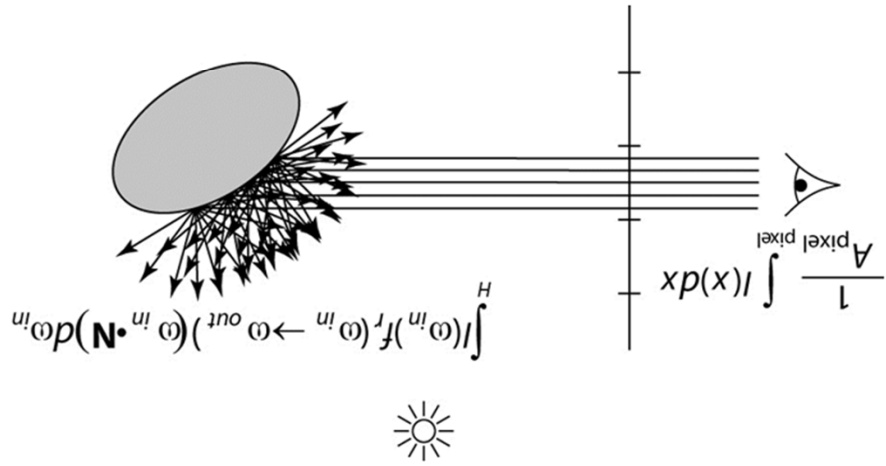
$$I(\omega_{out}) = \int_I f_r(\omega_{in} \rightarrow \omega_{out}) (\omega_{in} \cdot \mathbf{N}) d\omega_{in}$$

Reflection anti-aliasing




Reflection anti-aliasing

Pixel and reflection anti-aliasing



Pixel and reflection anti-aliasing

Full anti-aliasing



Full anti-aliasing...lots of nested integrals!

Computing these integrals is prohibitively expensive, especially after following the rays recursively.

We'll look at ways to approximate high-dimensional integrals...

Approximating integrals

Let's say we want to compute the integral of a function:

$$F = \int f(x) dx$$

If $f(x)$ is not known analytically, but can be evaluated, then we can approximate the integral by:

$$F \approx \sum_{i=1}^n f(i\Delta x) \Delta x$$


where we have sampled n times at spacing Δx . If these samples are distributed over an interval w , then

$$\Delta x = \frac{w}{n}$$

and the summation becomes:

$$F \approx \sum_{i=1}^n \frac{w}{n} f(i\Delta x)$$

Evaluating an integral in this manner is called **quadrature**.



A stochastic approach

An alternative to distributing the sample positions regularly is to distribute them **stochastically**.

Let's say the position in x is a random variable X , which is distributed according to $p(x)$, a probability density function (non-negative, integrates to unity).

Recall some of the "rules" of random variables:

$$\begin{aligned}
 E[X] &\equiv \int x p(x) dx \\
 V[X] &\equiv E[X^2] - (E[X])^2 = \int x^2 p(x) dx - \left(\int x p(x) dx \right)^2 \\
 E[kX] &= k E[X] \\
 V[kX] &= k^2 V[X]
 \end{aligned}$$

We can also show that for independent random variables X and Y :

$$\begin{aligned}
 E[X+Y] &= E[X] + E[Y] \\
 V[X+Y] &= V[X] + V[Y]
 \end{aligned}$$

A stochastic approach (cont)

We can now approximate $E[X]$ as the average of n samples:

$$E[X] \approx \frac{1}{n} \sum_{i=1}^n X_i$$

Where the X_i are independent and identically distributed (i.i.d.) random variables, each with distribution $p(x)$.

In fact, the summation is itself another random variable. What is its expected value and variance?

$$E\left[\frac{1}{n} \sum_{i=1}^n X_i\right] = \frac{1}{n} E\left[\sum_{i=1}^n X_i\right] = \frac{1}{n} \sum_{i=1}^n E[X_i] = \frac{1}{n} \sum_{i=1}^n \mu = \mu$$

$$V\left[\frac{1}{n} \sum_{i=1}^n X_i\right] = \frac{1}{n^2} V\left[\sum_{i=1}^n X_i\right] = \frac{1}{n^2} \sum_{i=1}^n V[X_i] = \frac{1}{n^2} \sum_{i=1}^n \sigma^2 = \frac{\sigma^2}{n}$$

$$\sigma = \sqrt{V[X]}$$

Integrals as expected values

Suppose now we have a function of a random variable, $g(x)$.

The expected value is:

$$E[g(X)] = \int g(x)p(x)dx$$

Getting back to our original problem of estimating an integral, can we choose $g(x)$ so that:

$$F = \int f(x)dx = E[g(X)] \quad ?$$

$$\int f(x)dx = \int g(x)p(x)dx$$

$$\int \frac{f(x)}{p(x)} p(x)dx = \int g(x)p(x)dx$$

Monte Carlo integration

Thus, given a set of samples positions, X_i , we can estimate the integral as:

$$F \approx \frac{1}{n} \sum_{i=1}^n f(X_i) \frac{1}{p(X_i)} \approx \mathbb{E} \left[\frac{f(X)}{p(X)} \right] = \mathbb{E} [g(X)] = \mathbb{E} [f]$$

This procedure is known as **Monte Carlo integration**.

What is the variance of the estimate?

$$\begin{aligned} V \left[\frac{1}{n} \sum_{i=1}^n f(X_i) \frac{1}{p(X_i)} \right] &= \frac{1}{n^2} \sum_{i=1}^n V \left[\frac{f(X_i)}{p(X_i)} \right] \\ &= \frac{1}{n^2} \sum_{i=1}^n V \left[\frac{f(X_i)}{p(X_i)} \right] \\ &= \frac{1}{n^2} \sum_{i=1}^n V \left[\frac{f(X)}{p(X)} \right] = \frac{1}{n} V \left[\frac{f(X)}{p(X)} \right] \end{aligned}$$

We want a low variance estimate. What variables and/or functions are under our control here?

$$n, p(x)$$

Uniform sampling

Suppose that the unknown function we are integrating happens to be a normalized box function

of width a :

$$f(x) = \begin{cases} 1/a & |x| \leq a/2 \\ 0 & \text{otherwise} \end{cases}$$

Suppose we now try to estimate the integral of $f(x)$

with uniform sampling over an interval of width w (i.e.,

choosing X from a uniform distribution):

$$p(x) = \begin{cases} 1/w & |x| \leq w/2 \\ 0 & \text{otherwise} \end{cases}$$


where $w \geq a$.

$$E \left[\frac{f(x)}{p(x)} \right] = E \left[\left(\frac{f(x)}{p(x)} \right)^2 \right] - E \left[\frac{f(x)}{p(x)} \right]^2$$

$$= \int_{-w/2}^{w/2} \left[\frac{f(x)}{p(x)} \right]^2 dx - \left(\int_{-w/2}^{w/2} \frac{f(x)}{p(x)} dx \right)^2$$

$$= \int_{-w/2}^{w/2} \left(\frac{1}{w} \right)^2 dx - \left(\int_{-w/2}^{w/2} \left(\frac{1}{w} \right) dx \right)^2$$

$$= \frac{1}{w} \cdot \left(\frac{1}{w} \right)^2 \cdot a - \left(\frac{1}{w} \cdot \frac{1}{w} \cdot a \right)^2$$



Importance sampling

A better approach, if $f(x)$ is non-negative, would be to choose $p(x) \sim f(x)$. In fact, this choice would be optimal.

$$p(x) = k f(x)$$

$$g(x) = \frac{f(x)}{p(x)} = \frac{f(x)}{k f(x)} = \frac{1}{k}$$

$$E[g(x)] = \frac{1}{k}$$

$$V\left[\frac{1}{k}\right] = 0$$

$$k = \frac{1}{\int f(x) dx}$$

$$\int k f(x) dx = 1 \implies \int f(x) dx = \frac{1}{k}$$

$$k \int f(x) dx = 1$$

$$k = \frac{1}{\int f(x) dx}$$

Why don't we just do that?
Requires knowledge of $\int f(x) dx$

Alternatively, we can use heuristics to guess where $f(x)$ will be large and choose $p(x)$ based on those heuristics. This approach is called **importance sampling**.


$$f(x) \approx h(x) \phi(x)$$

Stratified sampling

An improvement on importance sampling is **stratified sampling**.


The idea is that, given your probability function:

- ◆ You can break it up into bins of equal probability area (i.e., equal likelihood).
- ◆ Then choose a sample from each bin.



Summing over ray paths

We can think of this problem in terms of enumerated rays:



The intensity at a pixel is the sum over the primary rays:

$$I^{pixel} = \sum_n \frac{1}{n} I(r_i)$$


For a given primary ray, its intensity depends on secondary rays:

$$I(r_i) = \sum_j I(r_{ij}) f_r(r_{ij} \rightarrow r_i)$$

Substituting back in:

$$I^{pixel} = \sum_1^n \sum_1^j \frac{1}{n} I(r_{ij}) f_r(r_{ij} \rightarrow r_i)$$

Summing over ray paths



We can incorporate tertiary rays next:

$$I^{pixel} = \frac{1}{n} \sum_i^i \sum_j^j \sum_k^k I(r_{ijk}) f_r(r_{ijk} \leftarrow r_{ij}) f_r(r_{ij} \leftarrow r_i)$$

Each triple ijk corresponds to a ray path:

$$r_{ijk} \leftarrow r_{ij} \leftarrow r_i$$


So, we can see that ray tracing is a way to approximate a complex, nested light transport integral with a summation over ray paths (of arbitrary length!).

Problem: too expensive to sum over all paths.


Solution: choose a small number of "good" paths.

Glossy reflection revisited

Let's return to the glossy reflection model, and modify it – for purposes of illustration – as follows:




We can visualize the span of rays we want to integrate over, within a pixel:



Whitted ray tracing

Returning to the reflection example, Whitted ray tracing replaces the glossy reflection with mirror reflection:




Thus, we render with anti-aliasing as follows:

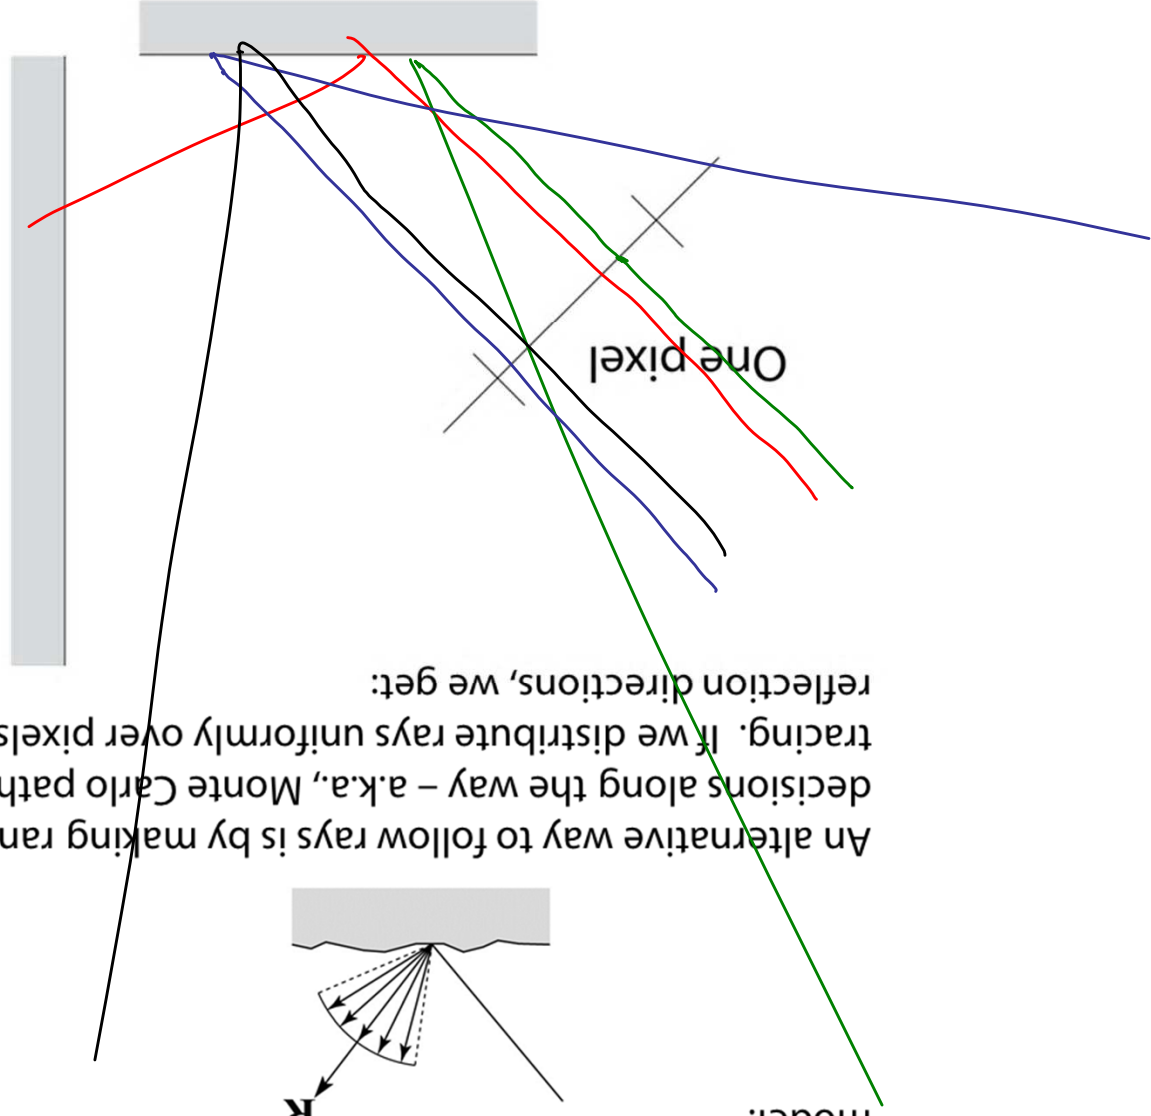


Monte Carlo path tracing

Let's return to our original (simplified) glossy reflection model:




An alternative way to follow rays is by making random decisions along the way – a.k.a., Monte Carlo path tracing. If we distribute rays uniformly over pixels and reflection directions, we get:

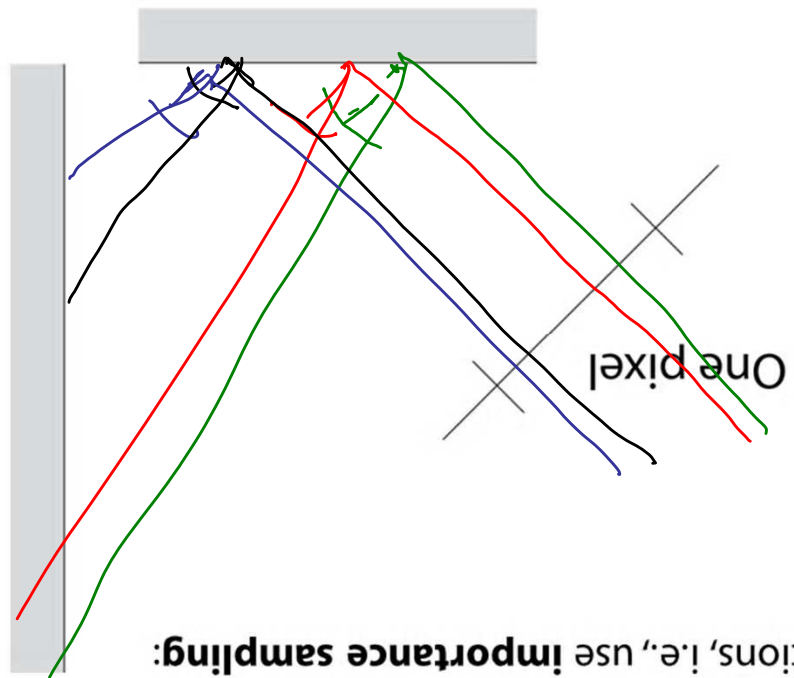


Importance sampling

The problem is that lots of samples are "wasted."
Using again our glossy reflection model:




Let's now randomly choose rays, but according to a probability that favors more important reflection directions, i.e., use **importance sampling**:




Stratified sampling

We still have a problem that rays may be clumped together. We can improve on this by splitting reflection into zones:



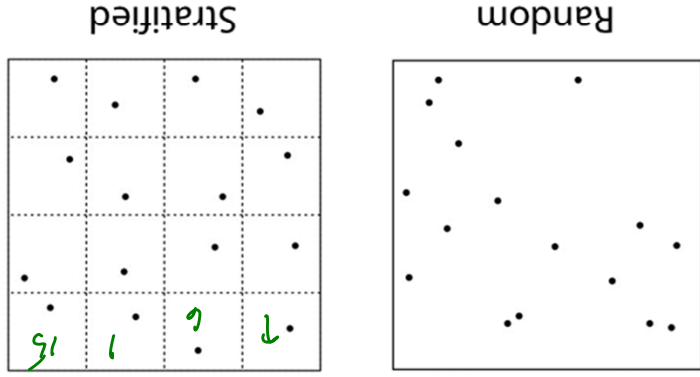
Now let's restrict our randomness to within these zones, i.e. use **stratified sampling**:



One pixel

Stratified sampling of a 2D pixel

Here we see pure uniform vs. stratified sampling over a 2D pixel (here 16 rays/pixel):



The stratified pattern on the right is also sometimes called a **jittered** sampling pattern.

One interesting side effect of these stochastic sampling patterns is that they actually injects noise into the solution (slightly grainier images). This noise tends to be less objectionable than aliasing artifacts.

Distribution ray tracing

These ideas can be combined to give a particular method called **distribution ray tracing** [Cook84]:

- ◆ uses non-uniform (jittered) samples.
- ◆ replaces aliasing artifacts with noise.
- ◆ provides additional effects by distributing rays to sample:

- Reflections and refractions
- Light source area
- Camera lens area
- Time

[This approach was originally called "distributed ray tracing," but we will call it distribution ray tracing (as in probability distributions) so as not to confuse it with a parallel computing approach.]

DRT pseudocode

TraceImage() looks basically the same, except now each pixel records the average color of jittered sub-pixel rays.


```
function traceImage (scene):  
  for each pixel (i, j) in image do  
    |i, j| → 0  
    for each sub-pixel id in (i, j) do  
      s → pixelToWorld(jitter(i, j, id))  
      p → COP  
      d → (s - p).normalize()  
      |i, j| → |i, j| + traceRay(scene, p, d, id)  
    end for  
    |i, j| → |i, j|/numSubPixels  
  end for  
end function
```

A typical choice is numSubPixels = 5*5.

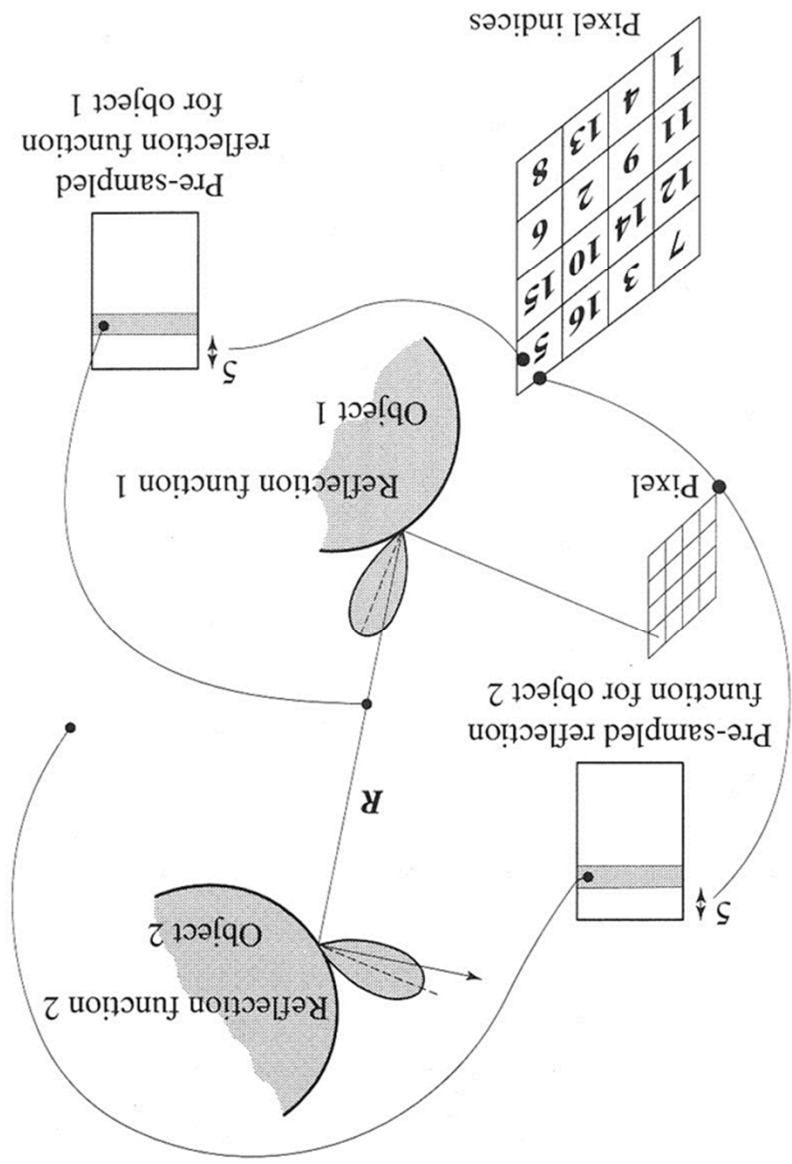
DRT pseudocode (cont'd)

Now consider `traceRay()`, modified to handle (only) opaque glossy surfaces:

```
function traceRay(scene, p, d, id):  
    (q, N, material) ← intersect(scene, p, d)  
    | ← shade(...)  
    R ← jitteredReflectDirection(N, -d, material, id)  
    | → | + material.kr * traceRay(scene, q, R, id)  
    return |  
end function
```

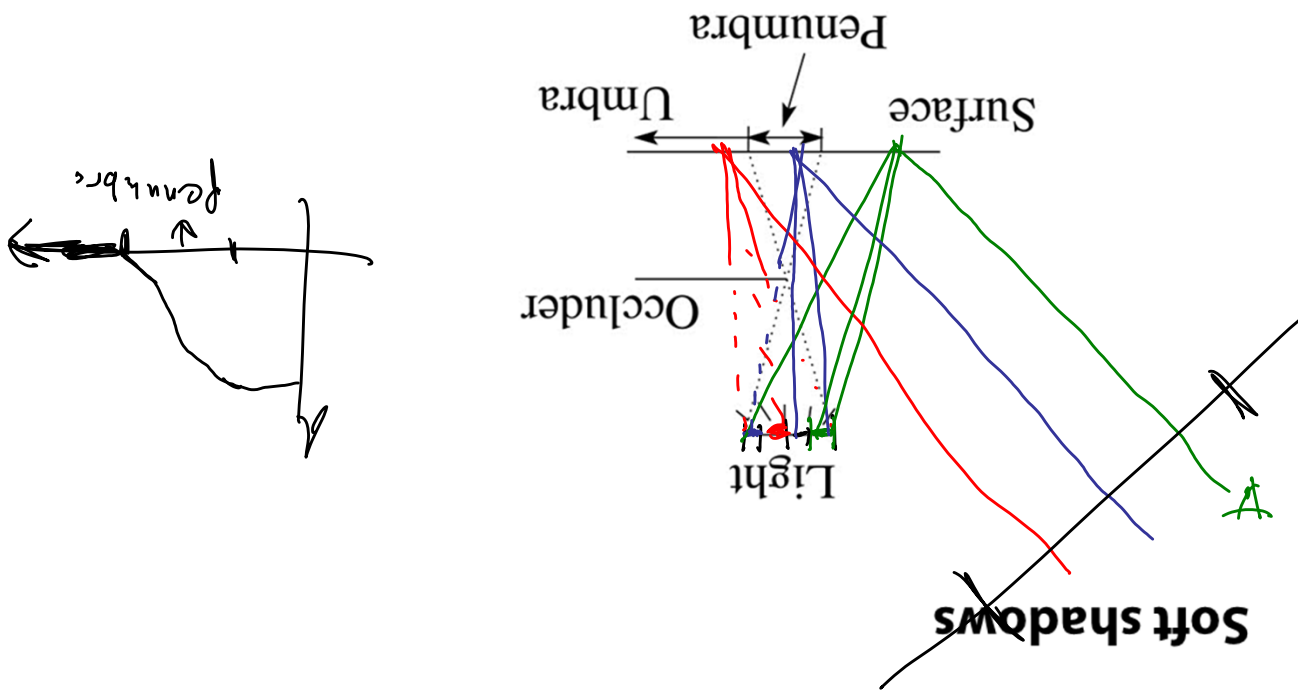


Pre-sampling glossy reflections (Quasi-Monte Carlo)






Distributing rays over light source area gives:




The pinhole camera

The first camera - "camera obscura" - known to Aristotle.




In 3D, we can visualize the blur induced by the pinhole (a.k.a., **aperture**):



Q: How would we reduce blur?



narrow aperture

Shrinking the pinhole




Q: What happens as we continue to shrink the aperture?

Shrinking the pinhole, cont'd




The pinhole camera, revisited

We can think in terms of light heading toward the image plane:




We can equivalently turn this around by following rays from the viewer:

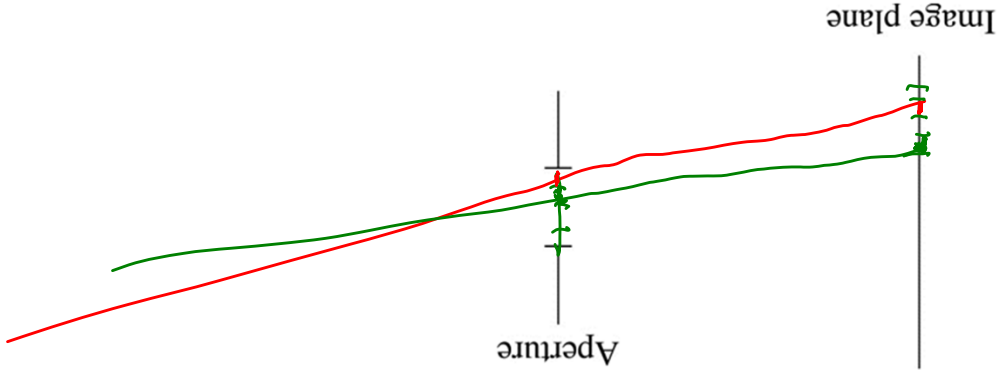


The pinhole camera, revisited

Given this flipped version:




how can we simulate a pinhole camera more accurately?



Lenses

Pinhole cameras in the real world require small apertures to keep the image in focus.

Lenses focus a bundle of rays to one point => can have larger aperture.



For a "thin" lens, we can approximately calculate where an object point will be in focus using the Gaussian lens formula:



$$\frac{1}{f} = \frac{1}{d_o} + \frac{1}{d_i}$$

where f is the **focal length** of the lens.

Depth of field

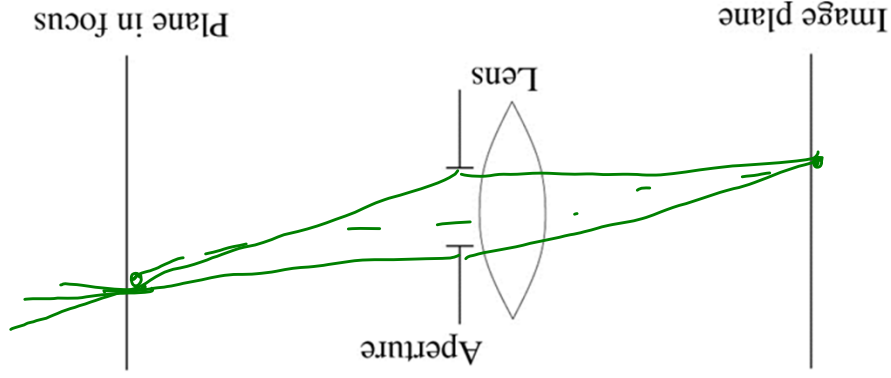
Lenses do have some limitations. The most noticeable is the fact that points that are not in the object plane will appear out of focus.

The **depth of field** is a measure of how far from the object plane points can be before appearing "too blurry:"

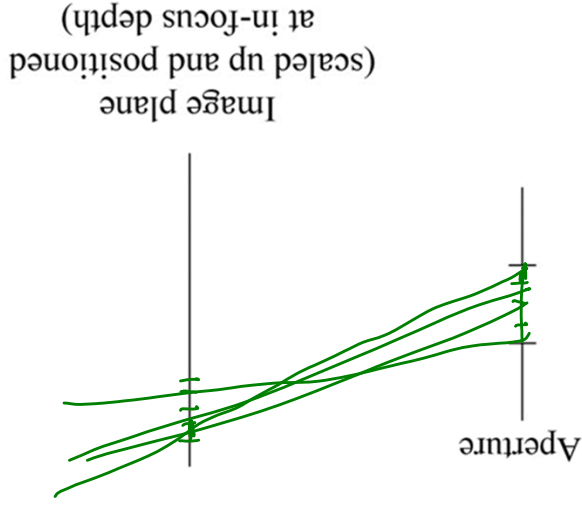


Simulating depth of field

Consider how rays flow between the image plane and the in-focus plane:



We can model this as simply placing our image plane at the in-focus location, in front of the finite aperture, and then distributing rays over the aperture (instead of the ideal center of projection):






Simulating depth of field, cont'd

Chaining the ray id's

In general, you can trace rays through a scene and keep track of their id's to handle *all* of these effects:



DRT to simulate *motion blur*

Distributing rays over time gives:

