

1 Inversion Method

Given a uniform random variable U in the interval $[0, 1]$, we can generate a random variable from any other one dimensional distribution using its cumulative distribution function: $X = F^{-1}(U)$. This is how we choose sample points when running a ray tracing algorithm.

1. What function of U will return a sample from the exponential distribution (with parameter λ)? This distribution has density $p_\lambda(x) = \lambda e^{-\lambda x}$, and is defined for $x \geq 0$.

Solution: First, we need to calculate the CDF:

$$F_\lambda(x) = \int_0^x p_\lambda(t) dt = \int_0^x \lambda e^{-\lambda t} dt = -e^{-\lambda t} \Big|_0^x = 1 - e^{-\lambda x}$$

Set the CDF equal to U :

$$U = 1 - e^{-\lambda x}$$

Solving for x :

$$\begin{aligned} e^{-\lambda x} &= 1 - U \\ -\lambda x &= \ln(1 - U) \\ x &= -\frac{\ln(1 - U)}{\lambda} \end{aligned}$$

Thus, the inverse function is given by:


$$F_\lambda^{-1}(x) = -\frac{\ln(1 - x)}{\lambda}$$

and we can return:

$$-\frac{\ln(1 - U)}{\lambda}$$

to sample from the exponential distribution.


2. For $\lambda = 1$, the plot below shows the CDF, Inverse CDF, and PDF. If we sample $U = 0.75$, determine the corresponding sampled value using the inverse method. Mark the corresponding point on the CDF curve in the graph.



Solution: Using the inverse CDF function:

$$x = -\ln(1 - 0.75) = -\ln(0.25) \approx 1.386.$$

Thus, the sampled value is approximately $x \approx 1.386$.



3. What does the x-axis represent for the blue CDF curve?

Solution: The x-axis represents the possible values that can be sampled from the exponential distribution. The CDF gives the probability that a randomly drawn value from this distribution is less than or equal to a given x . Therefore, the x-axis represents the range of values that an exponential random variable can take.

2 Unbiased Estimators

1. Let $f : [-2, 2] \times [-2, 2] \rightarrow \mathbb{R}$ be a function. You have a machine that allows you to sample $2n$ values independently and uniformly from the interval $[-2, 2]$. Construct an unbiased Monte Carlo estimator for

$$F = \int_{-2}^2 \int_{-2}^2 f(x, y) dx dy.$$

Solution: Draw n random samples $X_1, \dots, X_n$ and $Y_1, \dots, Y_n$. Then we take

$$\begin{aligned} \langle F_n \rangle &:= \frac{1}{n} \sum_{i=1}^n \frac{f(X_i, Y_i)}{p(X_i, Y_i)}, \text{ where } p(X_i, Y_i) = p(X_i)p(Y_i) = \left(\frac{1}{4}\right) \left(\frac{1}{4}\right) = \frac{1}{16} \\ &= \frac{1}{n} \sum_{i=1}^n \frac{f(X_i, Y_i)}{\frac{1}{16}} \\ &= \frac{16}{n} \sum_{i=1}^n f(X_i, Y_i). \end{aligned}$$

3 Boxing

Preston owns an unpolished wooden cube that scatters light diffusely. In other words, Preston's cube is a *Lambertian* surface with a bidirectional reflectance distribution function (BRDF) of $f_r(\mathbf{p}, \omega_i, \omega_o) = \frac{\rho}{\pi}$.

Recall that reflected radiance, $L_r(\mathbf{p}, \omega_o)$, is equal to $L_e(\mathbf{p}, \omega_o) + \int_{H^2} f_r(\mathbf{p}, \omega_i, \omega_o) L_i(\mathbf{p}, \omega_i) \cos \theta_i d\omega_i$.

1. What is the emitted radiance, L_e , of the cube?

Solution: L_e is zero, since the wooden cube does not emit any light.

2. L_i is the incoming radiance. Assume L_i is uniform over the hemisphere, H^2 , that surrounds point $\mathbf{p}$, located on the top face of the cube. Solve for $L_r(\mathbf{p}, \omega_o)$.

(Hint: re-parameterize your integral to be in terms of θ and ϕ , instead of ω .)

Solution:

$$\begin{aligned}
 L_r(\mathbf{p}, \omega_o) &= L_e(\mathbf{p}, \omega_o) + \int_{H^2} f_r(\mathbf{p}, \omega_i, \omega_o) L_i(\mathbf{p}, \omega_i) \cos \theta_i d\omega_i \\
 &= 0 + \int_{H^2} \frac{\rho}{\pi} L_i(\mathbf{p}, \omega_i) \cos \theta_i d\omega_i \\
 &= \frac{\rho}{\pi} L_i \int_{H^2} \cos \theta_i d\omega_i \\
 &= \frac{\rho}{\pi} L_i \int_0^{2\pi} \int_0^{\pi/2} \cos \theta_i \sin \theta_i d\theta_i d\phi_i \\
 &= \frac{\rho}{\pi} L_i \left(\int_0^{2\pi} d\phi_i \right) \left(\int_0^{\pi/2} \cos \theta_i \sin \theta_i d\theta_i \right) \\
 &= \frac{\rho}{\pi} L_i (2\pi) \left(\frac{1}{2} \right) \tag{*} \\
 &= \rho L_i.
 \end{aligned}$$

In line (*), we have used the substitution $u = \sin \theta \implies du = \cos \theta d\theta$, so that

$$\int_0^{\pi/2} \cos \theta \sin \theta d\theta = \int_{\theta=0}^{\pi/2} u du = \frac{u^2}{2} \Big|_{\theta=0}^{\pi/2} = \frac{\sin^2 \theta}{2} \Big|_{\theta=0}^{\pi/2} = \frac{1}{2}.$$

3. How does the reflected radiance depend on ω_o ? How does the reflected radiance depend on ρ ?

Solution: It doesn't depend on ω_o ! That outgoing radiance does not depend on direction is what makes the surface Lambertian.

ρ scales the incoming radiance. It's known as the *albedo* of a surface.

4. Preston adjusts his light source so that incoming radiance is no longer uniform. Now, he wants to use a Monte Carlo estimator to approximate $L_r(\mathbf{p}, \omega_i)$. He samples n directions over the hemisphere from $p(\omega)$, a distribution that is proportional to the BRDF.

Construct the Monte Carlo estimator.

Solution:

$$\begin{aligned}\hat{L}_r(\mathbf{p}, \omega_o) &= \frac{1}{n} \sum_{j=1}^n \frac{\frac{\rho}{\pi} L_i(\mathbf{p}, \omega_{i,j}) \cos \theta_{i,j}}{\frac{1}{2\pi}} \\ &= \frac{1}{n} \sum_{j=1}^n \frac{2\rho L_i(\mathbf{p}, \omega_{i,j}) \cos \theta_{i,j}}{1} \\ &= \frac{2\rho}{n} \sum_{j=1}^n L_i(\mathbf{p}, \omega_{i,j}) \cos \theta_{i,j}\end{aligned}$$

5. In practice, cosine-weighted hemisphere sampling results in better convergence. When $p(\omega) = \frac{\cos \theta}{\pi}$, what is the Monte Carlo estimator for n samples?

Solution:

$$\begin{aligned}\hat{L}_r(\mathbf{p}, \omega_o) &= \frac{1}{n} \sum_{j=1}^n \frac{\frac{\rho}{\pi} L_i(\mathbf{p}, \omega_{i,j}) \cos \theta_{i,j}}{\frac{\cos \theta_{i,j}}{\pi}} \\ &= \frac{1}{n} \sum_{j=1}^n \rho L_i(\mathbf{p}, \omega_j)\end{aligned}$$


6. Conceptually, why does cosine-weighted hemisphere sampling outperform uniform sampling over a hemisphere?

Solution: Cosine-weighted sampling improves convergence by aligning with the integrand's $\cos \theta$ term, reducing variance compared to uniform sampling, which wastes samples in less significant directions.


4 Tracing Outside of the Box

The Monte Carlo estimator derived earlier is great for direct illumination from the light source (1-bounce). For *indirect* illumination, it's not enough to sample directions — we need to sample *paths*.

1. Again, Preston's cube has a BRDF of $f_r(\mathbf{p}, \omega_i, \omega_o) = \frac{\rho}{\pi}$. Perform path-tracing by drawing multiple 2-bounce paths for point $\mathbf{p}$. Label ω_o . Light can scatter off walls. Assume the walls are also Lambertian.



2. Draw multiple 3-bounce paths. Label ω_o .



3. Suppose we trace n non-occluded 2-bounce paths from camera, to $\mathbf{p}$, to $\mathbf{p}_j$, to the light source, where $j = 1, 2, \dots, n$. Express the Monte Carlo estimator for outgoing radiance at $\mathbf{p}$ in terms of outgoing radiance at $\mathbf{p}_j$. Assume incoming directions to $\mathbf{p}$, $\omega_{i,j}$, are sampled from $p(\omega)$.

Solution: We are sampling directions $\omega_{i,j}$ from which light is incoming. Here, we are given that these directions intersect with the scene at points $\mathbf{p}_j$. That is, light travels from (outgoing) $\mathbf{p}_j$ to (incoming) $\mathbf{p}$, in the direction $-\omega_{i,j}$.

$$L_o(\mathbf{p}, \omega_o) = \frac{1}{n} \sum_{j=1}^n \frac{f_r(\mathbf{p}, \omega_{i,j}, \omega_o) L_o(\mathbf{p}_j, -\omega_{i,j}) \cos \theta_{i,j}}{p(\omega)}$$

$$= \frac{1}{n} \sum_{j=1}^n \frac{\frac{\rho}{\pi} L_o(\mathbf{p}_j, -\omega_{i,j}) \cos \theta_{i,j}}{p(\omega)}$$

5 Ray or Nay?

With Russian Roulette, we randomly terminate each ray with probability $1 - p_{\text{rr}}$ (equivalently, continue with probability p_{rr}). Additionally, if the original estimator was X , we update it to

$$X_{\text{rr}} = \begin{cases} X/p_{\text{rr}} & \text{with probability } p_{\text{rr}} \\ 0 & \text{else,} \end{cases}$$

so that $\mathbb{E}[X_{\text{rr}}] = \mathbb{E}[X]$.

For ray tracing, X might be the estimated incoming radiance for a given (p, ω_i) .

1. Suppose $p_{\text{rr}} = \frac{4}{5}$, and let $N \geq 0$ be a random variable representing the number of bounces.

- (i) What is $\mathbb{P}(N = 0)$?

Solution: The ray must terminate immediately, with probability $1 - p_{\text{rr}} = \frac{1}{5}$.

- (ii) What is $\mathbb{P}(N = 2)$?

Solution: The ray must bounce twice, then terminate. This occurs with probability $p_{\text{rr}}^2(1 - p_{\text{rr}}) = \left(\frac{4}{5}\right)^2 \cdot \frac{1}{5} = \frac{16}{125}$.

- (iii) What is the expected value of N ? (Hint: Recall that if $Z \sim \text{Geometric}(p)$, then $\mathbb{P}(Z = z) = (1 - p)^{z-1}p$ for $z \geq 1$, and $\mathbb{E}[Z] = \frac{1}{p}$.)

Solution: In general, for $n \geq 0$, we have $\mathbb{P}(N = n) = p_{\text{rr}}^n(1 - p_{\text{rr}})$. Equivalently, we can try to match the form of the probability mass function in the hint: for $n \geq 1$,

$$\mathbb{P}(N + 1 = n) = \mathbb{P}(N = n - 1) = p_{\text{rr}}^{n-1}(1 - p_{\text{rr}}).$$

So we can match the distributions, $N + 1 \sim \text{Geometric}(1 - p_{\text{rr}})$! Taking expectations,

$$\mathbb{E}[N] + 1 = \mathbb{E}[N + 1] = \frac{1}{1 - p_{\text{rr}}} = 5,$$

so $\mathbb{E}[N] = 4$.

2. If $p_{\text{rr}} = \frac{1}{5}$, then is the expected value of N ?

Solution: Following the same logic as in part [iii](#), we have

$$\mathbb{E}[N] + 1 = \frac{1}{1 - p_{\text{rr}}} = \frac{1}{1 - \frac{1}{5}} = \frac{5}{4},$$

so $\mathbb{E}[N] = \frac{1}{4}$.

3. In general, how does increasing p_{rr} affect the expected number of bounces $\mathbb{E}[N]$?

Solution: Increasing p_{rr} increases $\mathbb{E}[N]$.