02941 Physically Based Rendering

Monte Carlo Integration

Jeppe Revall Frisvad

June 2018

Why Monte Carlo?

- The rendering equation

$$L_o(\boldsymbol{x}, \vec{\omega}) = L_e(\boldsymbol{x}, \vec{\omega}) + \int_{2\pi} f_r(\boldsymbol{x}, \vec{\omega}', \vec{\omega}) L_i(\boldsymbol{x}, \vec{\omega}') \cos\theta \, \mathrm{d}\omega'$$

  is difficult, usually impossible, to solve analytically.
- Trapezoidal integration and Gaussian quadrature only works well for smooth low-dimensional integrals.
- The rendering equation is 5-dimensional and it usually involves discontinuities.
- There are (roughly) only three known mathematical methods for solving this type of problem:
  - Truncated series expansion
  - Finite basis (discretization)
  - Sampling (random selection)
- Monte Carlo is probably the simplest way to use sampling.

Brush up on probability

- A *random variable* $X \in A$ is a value X drawn from the sample space A of some random process.
- Applying a function $f : X \to Y$ to a random variable X results in a new random variable Y.
- A *uniform* random variable takes on all values in its sampling space with *equal probability*.
- *Probability* is the chance (represented by a real number in [0,1]) that something is the case or that an event will occur.
- The *cumulative distribution function* (cdf) is the probability that a random variable X is smaller than or equal to a value x:

$$P(x) = \Pr\{X \leq x\} \ .$$

- The *probability density function* (pdf) is the relative probability for a random variable X to take on a particular value x:

$$\mathrm{pdf}(x) = \frac{\mathrm{d}P(x)}{\mathrm{d}x} \ .$$

Properties of the probability density function

- For uniform random variables, $\mathrm{pdf}(x)$ is constant.
- Of particular interest is the continuous, uniform, random variable $\xi \in [0,1]$ which has the probability density function

$$\mathrm{pdf}(x) = \begin{cases} 1 & \text{for } x \in [0,1] \\ 0 & \text{otherwise} \end{cases}$$

- Using the pdf, we can calculate the probability that a random variable lies inside an interval:

$$\Pr\{x \in [a,b]\} = \int_a^b \mathrm{pdf}(x) \, \mathrm{d}x \ .$$

- All probability density functions have the properties:

$$\mathrm{pdf}(x) \geq 0 \quad \text{and} \quad \int_{-\infty}^{\infty} \mathrm{pdf}(x) \, \mathrm{d}x = 1 \ .$$

Expected values and variance

- The *expected value* of a random variable $X \in A$ is the average value over the distribution of values $\mathrm{pdf}(x)$:

$$E\{X\} = \int_A x \,\mathrm{pdf}(x)\,\mathrm{d}x \ .$$

- The expected value of an arbitrary function $f(X)$ is then:

$$E\{f(X)\} = \int_A f(x)\,\mathrm{pdf}(x)\,\mathrm{d}x \ .$$

- The *variance* is the expected deviation of the function from its expected value:

$$V\{f(X)\} = E\left\{(f(X) - E\{f(X)\})^2\right\} \ .$$

- The expected value operator E is linear. Thus:

$$V\{f(X)\} = E\{(f(X))^2\} - (E\{f(X)\})^2 \ .$$

Properties of variance

- The variance operator:

$$V\{f(X)\} = E\{(f(X))^2\} - (E\{f(X)\})^2 \ .$$

- $V\{\cdot\}$ is not a linear operator. For a scalar a, we have

$$V\{a\,f(X)\} = a^2 V\{f(X)\} \ .$$

- And, furthermore,

$$\begin{aligned} V\{f(X)+f(Y)\} &= E\{(f(X)+f(Y))^2\} - (E\{f(X)+f(Y)\})^2 \\ &= E\{(f(X))^2 + (f(Y))^2 + 2f(X)f(Y)\} - (E\{f(X)\})^2 - (E\{f(Y)\})^2 - 2E\{f(X)\}E\{f(Y)\} \\ &= V\{f(X)\} + V\{f(Y)\} + 2E\{f(X)f(Y)\} - 2E\{f(X)\}E\{f(Y)\} \\ &= V\{f(X)\} + V\{f(Y)\} + 2\,Cov\{f(X), f(Y)\} \end{aligned}$$

- Thus, if X and Y are uncorrelated $(Cov\{f(X), f(Y)\} = 0)$, then the variance of the sum is equal to the sum of the variances.

The Monte Carlo estimator

- The law of large numbers:

$$\Pr\left\{\frac{1}{N}\sum_{j=1}^{N} f(X_j) \to E\{f(X)\}\right\} = 1 \quad \text{for} \quad N \to \infty \ .$$

"It is certain that the estimator goes to the expected value as the number of samples goes to infinity."

- Approximating an arbitrary integral using N samples:

$$F = \int_A f(x)\,\mathrm{d}x = \int_A \frac{f(x)}{\mathrm{pdf}(x)}\,\mathrm{pdf}(x)\,\mathrm{d}x = E\left\{\frac{f(X)}{\mathrm{pdf}(X)}\right\}$$

using the law of large numbers

$$F_N = \frac{1}{N}\sum_{j=1}^{N} \frac{f(X_j)}{\mathrm{pdf}(X_j)} \ ,$$

where X_j are sampled on A and $\mathrm{pdf}(x) > 0$ for all $x \in A$.

Monte Carlo error bound

- We found the estimator:

$$F_N = \frac{1}{N}\sum_{j=1}^{N} \frac{f(X_j)}{\mathrm{pdf}(X_j)} \ .$$

- The *standard deviation* is the square root of the variance:

$$\sigma_{F_N} = (V\{F_N\})^{1/2}$$

and it is a probabilistic error bound for the estimator according to Chebyshev's inequality:

$$\Pr\{|F_N - E\{F_N\}| \geq \delta\sigma_{F_N}\} \leq \delta^{-2} \ .$$

"The error is probably not too much larger than the standard deviation."
"There is a less than 1% chance that the error is larger than 10 standard deviations."

- The rate of convergence is then the ratio between the standard deviation of the estimator σ_{F_N} and the standard deviation of a single sample σ_Y.

Monte Carlo convergence

- The standard deviation of the estimator:

$$\sigma_{F_N} = (V\{F_N\})^{1/2} = \left(V \left\{ \frac{1}{N} \sum_{j=1}^{N} Y_j \right\} \right)^{1/2} ,$$

where

$$Y_j = \frac{f(X_j)}{\mathrm{pdf}(X_j)} .$$

- Continuing (while assuming that X_j and thus Y_j are uncorrelated)

$$\begin{aligned} \sigma_{F_N} &= \left(\frac{1}{N^2} V \left\{ \sum_{j=1}^{N} Y_j \right\} \right)^{1/2} = \left(\frac{1}{N^2} \sum_{j=1}^{N} V\{Y_j\} \right)^{1/2} \\ &= \left(\frac{1}{N} V\{Y\} \right)^{1/2} = \frac{1}{\sqrt{N}} \sigma_Y . \end{aligned}$$

- Worst case: Quadruple the samples to half the error.

An estimator for the rendering equation

- The rendering equation:

$$L_o(\boldsymbol{x}, \vec{\omega}) = L_e(\boldsymbol{x}, \vec{\omega}) + \int_{2\pi} f_r(\boldsymbol{x}, \vec{\omega}', \vec{\omega}) L_i(\boldsymbol{x}, \vec{\omega}') \cos\theta \, \mathrm{d}\omega' .$$

- The Monte Carlo estimator:

$$L_N(\boldsymbol{x}, \vec{\omega}) = L_e(\boldsymbol{x}, \vec{\omega}) + \frac{1}{N} \sum_{j=1}^{N} \frac{f_r(\boldsymbol{x}, \vec{\omega}_j', \vec{\omega}) L_i(\boldsymbol{x}, \vec{\omega}_j') \cos\theta}{\mathrm{pdf}(\vec{\omega}_j')}$$

with $\cos\theta = \vec{\omega}_j' \cdot \vec{n}$, where $\vec{n}$ is the surface normal at $\boldsymbol{x}$.

- The Lambertian BRDF:

$$f_r(\boldsymbol{x}, \vec{\omega}', \vec{\omega}) = \rho_d / \pi .$$

- A good choice of pdf would be:

$$\mathrm{pdf}(\vec{\omega}_j') = \cos\theta / \pi .$$

Sampling a pdf (the inversion method)

- How to draw samples X_i from an arbitrary pdf:
  1. Compute the cdf: $P(x) = \int_{-\infty}^{x} \mathrm{pdf}(x') \, \mathrm{d}x'$.
  2. Compute the inverse cdf: $P^{-1}(x)$.
  3. Obtain a uniformly distributed random number $\xi \in [0, 1]$.
  4. Compute a sample: $X_i = P^{-1}(\xi)$.
- Example: Exponential distribution over sample space $[0, \infty)$

$$\mathrm{pdf}(x) = a e^{-ax} .$$

- Compute cdf:

$$P(x) = \int_0^x a e^{-ax'} \, \mathrm{d}x' = 1 - e^{-ax} .$$

- Invert cdf:

$$P^{-1}(x) = -\frac{\ln(1-x)}{a} .$$

- To draw samples:

$$X = -\frac{\ln(1-\xi)}{a} \quad \text{or} \quad X = -\frac{\ln \xi}{a} .$$

Sampling a pdf (the rejection method)

- Imagine a pdf which we cannot integrate to find the cdf.
- Knowing a function g with the property $\mathrm{pdf}(x) < c\, g(x)$, where $c > 1$, we can use rejection sampling with g instead of sampling the pdf directly.

- Rejection sampling is the following algorithm:
- loop forever:
  - sample X from $g(x)$ and ξ from $[0, 1]$
  - if $\xi < f(X)/(c\, g(X))$ then return X

- Rejection sampling is only a good idea if $c\, g(x)$ is a tight bound for the pdf.

Uniformly sampling a sphere

- The unit box is a (relatively) tight bound for the unit sphere.
- Rejection sampling unit directions given by points on the unit sphere:

```
Vec3f direction;
do
{
  direction[0] = 2.0f*mt_random() - 1.0f;
  direction[1] = 2.0f*mt_random() - 1.0f;
  direction[2] = 2.0f*mt_random() - 1.0f;
}
while(dot(direction, direction) > 1.0f);
direction = normalize(direction);
```

- $\text{pdf}(\vec{\omega}'_j) = \frac{1}{4\pi}$.

Sampling a 2D joint density function

- Suppose we have a joint 2D density function $\text{pdf}(x, y)$.
- To sample $\text{pdf}(x, y)$ using two independent random variables X and Y, we find the *marginal* and the *conditional* density functions.
- The *marginal density function* is

$$\text{pdf}(x) = \int \text{pdf}(x, y)\, \mathrm{d}y \ .$$

- The *conditional density function* is

$$\text{pdf}(y|x) = \frac{\text{pdf}(x, y)}{\text{pdf}(x)} \ .$$

- The inversion method is then applied to each of the marginal and conditional density functions.

Cosine-weighted hemisphere sampling

- Sampling directions according to the distribution: $\text{pdf}(\vec{\omega}'_j) = \cos\theta/\pi \quad , \quad \text{pdf}(\theta, \phi) = \cos\theta \sin\theta/\pi$.
- Compute the marginal and conditional density functions:

$$\begin{aligned} \text{pdf}(\theta) &= \int_0^{2\pi} \frac{\cos\theta}{\pi} \sin\theta \, \mathrm{d}\phi = 2\cos\theta\sin\theta \ . \\ \text{pdf}(\phi|\theta) &= \frac{\cos\theta\sin\theta/\pi}{2\cos\theta\sin\theta} = \frac{1}{2\pi} \ . \end{aligned}$$

- The cdf for the marginal density function:

$$\begin{aligned} P(\theta) &= 2\int_0^{\theta} \cos\theta' \sin\theta' \, \mathrm{d}\theta' = 2\int_1^{\cos\theta} (-\cos\theta')\, \mathrm{d}\cos\theta' = 1 - \cos^2\theta \\ P(\phi|\theta) &= \phi/(2\pi) \ . \end{aligned}$$

- Invert these to find the sampling strategy:

$$\vec{\omega}'_j = (\theta, \phi) = (\cos^{-1}\sqrt{\xi_1},\ 2\pi\xi_2) \ .$$

Ambient occlusion

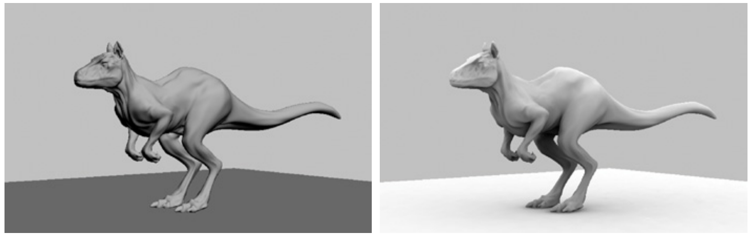

- Using the Lambertian BRDF for materials, $f_r = \rho_d/\pi$; the cosine weighted hemisphere for sampling, $\text{pdf}(\vec{\omega}'_j) = \cos\theta/\pi$; and a visibility term V for incident illumination, the Monte Carlo estimator for ambient occlusion is simply:

$$L_N(\boldsymbol{x}, \vec{\omega}) = \frac{1}{N}\sum_{j=1}^{N} \frac{f_r(\boldsymbol{x}, \vec{\omega}'_j, \vec{\omega}) L_i(\boldsymbol{x}, \vec{\omega}'_j) \cos\theta}{\text{pdf}(\vec{\omega}'_j)} = \rho_d(\boldsymbol{x}) \frac{1}{N}\sum_{j=1}^{N} V(\vec{\omega}'_j) \ .$$

Sampling a triangle mesh

- Uniformly sample a triangle (pdf $= 1/n$, where n is the number of triangle faces in the mesh).
- Uniformly sample a position on the triangle (pdf $= 1/A_\Delta$, where A_Δ is the triangle area):
  1. Sample barycentric coordinates $(u, v, w = 1 - u - v)$:

$$\begin{aligned} u &= 1 - \sqrt{\xi_1} \\ v &= (1 - \xi_2)\sqrt{\xi_1} \\ w &= \xi_2\sqrt{\xi_1} \end{aligned}$$

  2. Use the barycentric coordinates for linear interpolation of triangle vertices to obtain a point on the triangle.
  3. Use the barycentric coordinates for linear interpolation of triangle vertex normals to obtain the normal at the sampled surface point.

Soft shadows

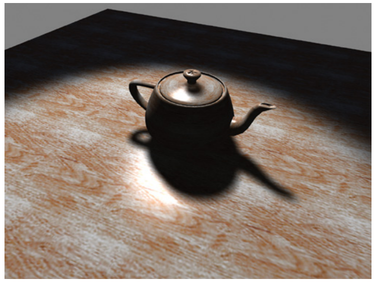

- From solid angle to area: $\mathrm{d}\omega' = \frac{\cos\theta_{\text{light}}}{r^2}\mathrm{d}A$.
- Using the Lambertian BRDF, $f_r = \rho_d/\pi$, and triangle mesh sampling of a point $\boldsymbol{x}_j$ on the light, pdf $= 1/(nA_{\Delta,j})$, the Monte Carlo estimator for area lights is:

$$L_N(\boldsymbol{x}, \vec{\omega}) = \frac{\rho_d(\boldsymbol{x})}{\pi}\frac{1}{N}\sum_{j=1}^{N} L_e(\boldsymbol{x}_j \to \boldsymbol{x})V(\boldsymbol{x}_j \leftrightarrow \boldsymbol{x})\frac{\cos\theta\cos\theta_{\text{light}}}{\|\boldsymbol{x} - \boldsymbol{x}_j\|^2}nA_{\Delta,j} \ .$$

Exercises

- Implement the triangle mesh sampling.
- Implement the Monte Carlo area light sampling.
- Implement the cosine-weighted hemispherical sampling.
- Implement the Monte Carlo ambient occlusion.