Lecture 6: Monte Carlo Rendering

CS 6620, Spring 2009

Kavita Bala

Computer Science

Cornell University

Monte Carlo Integration

- Numerical tool to evaluate integrals
- Use sampling
- Stochastic errors
- Unbiased
  - on average, we get the right answer!

Continuous random variable

- Expected value: $E[x] = \int_a^b x p(x) dx$

$$E[g(x)] = \int_a^b g(x) p(x) dx$$

- Variance:

$$\sigma^2[x] = \int_a^b (x - E[x])^2 p(x) dx$$

$$\sigma^2[g(x)] = \int_a^b (g(x) - E[g(x)])^2 p(x) dx$$

- Deviation: $\sigma[x], \sigma[g(x)]$

More than one sample

- Consider the weighted sum of N samples

$$g(x) = \frac{1}{N}(f(x_1) + f(x_2) + f(x_3) + \ldots f(x_N))$$

- Expected value

$$E[g(x)] = E[\frac{1}{N}\sum_i^N f(x_i)] = E[f(x)]$$

- Variance

$$\sigma^2[g(x)] = \sigma^2[\frac{1}{N}\sum_i^N f(x_i)] = \frac{1}{N}\sigma^2[f(x)]$$

- Deviation $\sigma[g(x)] = \frac{1}{\sqrt{N}}\sigma[f(x)]$

More samples

Secondary estimator

Generate N random samples x_i

Estimator: $\langle I \rangle = I_{\sec} = \frac{1}{N}\sum_{i=1}^{N} f(x^s{}_i)(b-a)$

Variance $\sigma^2_{\sec} = \sigma^2_{prim} / N$

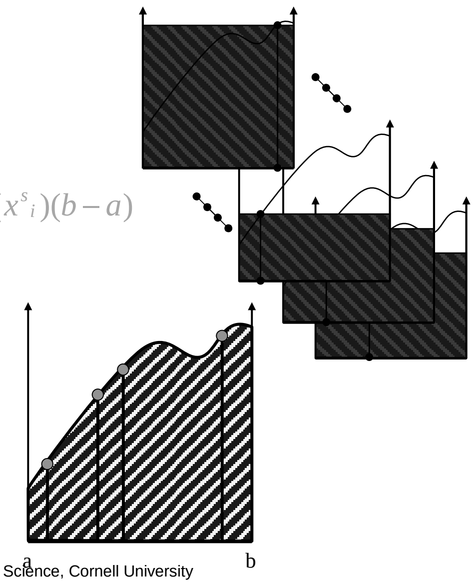

Monte Carlo Integration

- Expected value of estimator

$$E[\langle I \rangle] = E[\frac{1}{N}\sum_i^N \frac{f(x_i)}{p(x_i)}] = \frac{1}{N}\int (\sum_i^N \frac{f(x_i)}{p(x_i)}) p(x)dx$$

$$= \frac{1}{N}\sum_i^N \int (\frac{f(x)}{p(x)}) p(x)dx$$

$$= \frac{N}{N}\int f(x)dx = I$$

  – on ‘average’ get right result: **unbiased**

- Standard deviation σ is a measure of the stochastic error

$$\sigma^2 = \frac{1}{N}\int_a^b [\frac{f(x)}{p(x)} - I]^2 p(x)dx$$

Convergence Rates

- RMS error converges at a rate of O($\frac{1}{\sqrt{N}}$)
- Unbiased
- Chebychev's inequality

$$\Pr\left[\left|F - E(F)\right| \geq \sqrt{\frac{1}{\delta}}\sigma\right] \leq \delta$$

$$\Pr\left[\left|I_{estimator} - I\right| \geq \frac{1}{\sqrt{N}}\sqrt{\frac{1}{\delta}}\sigma\right] \leq \delta$$

- Strong law of large numbers

$$\Pr\left[\lim_{N\to\infty} \frac{1}{N}\sum_i^N \frac{f(x_i)}{p(x_i)} = I\right] = 1$$

Monte Carlo Integration - 2D

- MC Integration works well for higher dimensions
- Unlike quadrature

$$I = \int_a^b \int_c^d f(x, y)dxdy$$

$$\langle I \rangle = \frac{1}{N}\sum_{i=1}^{N} \frac{f(x_i, y_i)}{p(x_i, y_i)}$$

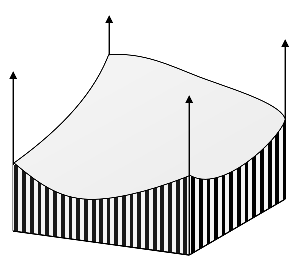

MC Advantages

- Convergence rate of O($\frac{1}{\sqrt{N}}$)
- Simple
  - – Sampling
  - – Point evaluation
  - – Can use black boxes
- General
  - – Works for high dimensions
  - – Deals with discontinuities, crazy functions,…

Importance Sampling

- Why do we sample by p(x)?
- Why not just uniformly?

- Better use of samples by taking more samples in ‘important’ regions, i.e. where the function is large

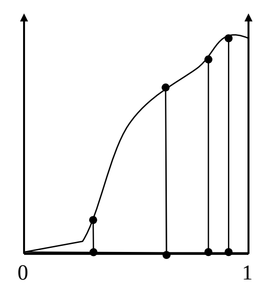

MC integration - Non-Uniform

- Some parts of the integration domain have higher importance
- Generate samples according to density function p(x)

$$\langle I \rangle = \frac{1}{N}\sum_{i=1}^{N}\frac{f(x_i)}{p(x_i)}$$

- Estimator?
- What is optimal p(x)? $p(x) \approx f(x) / \int f(x)dx$

MC integration - Non-Uniform

- Generate samples according to density function p(x)

$$p(x) \approx f(x) / \int f(x)dx$$

- Why?

$$\sigma^2 = \frac{1}{N}\int_a^b [\frac{f(x)}{p(x)} - I]^2 p(x)dx$$

$$= \frac{1}{N}\int_a^b [\frac{f(x)}{f(x)/I} - I]^2 p(x)dx = 0$$

- But…..

$$I_{estimator} = \frac{1}{N}\sum\frac{f(x)}{p(x)} = \frac{1}{N}\sum\frac{f(x)}{f(x)/I} = \frac{1}{N}\sum I = I$$

Example

- Function: $I = \int_0^4 x dx = 8$ $\quad f(x) = x$

$$\sigma^2 = \frac{1}{N}\int_a^b [\frac{f(x)}{p(x)} - I]^2 p(x)dx$$

$$p(x) = \frac{x}{8}, \sigma^2 = 0 \qquad I_{estimator} = I = 8$$

$$p(x) = \frac{1}{4}, \sigma^2 = \frac{1}{N}\int_0^4 [\frac{x}{1/4} - 8]^2 \frac{1}{4} dx = 21.3/N$$

$$p(x) = \frac{x+2}{16}, \sigma^2 = \frac{1}{N}\int_0^4 [\frac{x}{(x+2)/16} - 8]^2 \frac{x+2}{16} dx = 6.3/N$$

How to sample according to pdf?

- Consider discrete events x_i
  - with probability p_i

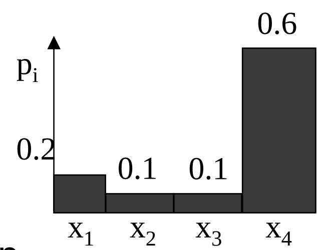

- Select x_i if:

$$p_1 + \ldots p_{i-1} < \xi < p_1 + \ldots p_{i-1} + p_i$$

$$\sum_{j=1}^{i-1} p_j < \xi < \sum_{j=1}^{i} p_j$$

$$P(x_i) = P(\xi \in [\sum_{j=1}^{i-1} p_j, \sum_{j=1}^{i} p_j])$$

$$P(a < \xi < b) = (b - a)$$

$$P(x_i) = P(\xi \in [\sum_{j=1}^{i-1} p_j, \sum_{j=1}^{i} p_j]) = \sum_{j=1}^{i} p_j - \sum_{j=1}^{i-1} p_j = p_i$$

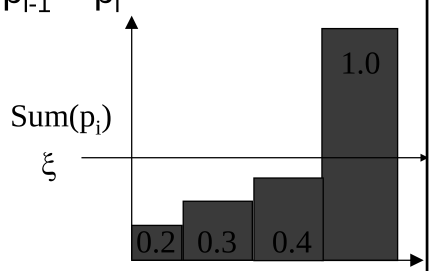

Non-Uniform Samples

- 1) Choose a normalized probability density function *p(x)*

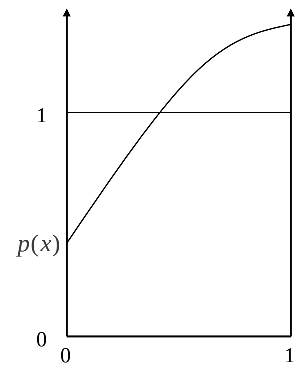

Non-Uniform Samples

- 1) Choose a normalized probability density function *p(x)*
- 2) Integrate to get a cumulative probability distribution function *P(x)*:

$$P(x) = \int_0^x p(t)dt$$

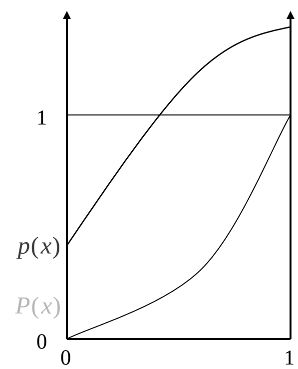

Note this is similar to computing $\sum_{j=1}^{i} p_j$

Non-Uniform Samples

- 1) Choose a normalized probability density function *p(x)*

- 2) Integrate to get a cumulative distribution function *P(x)*:

$$P(x) = \int_0^x p(t)dt$$

- 3) Invert *P*:

$$x = P^{-1}(\xi)$$

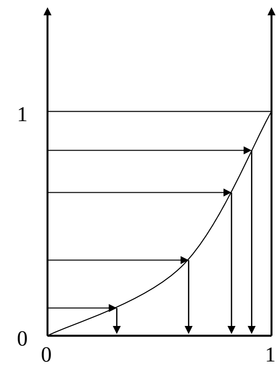

Note this is similar to going from y axis to x in discrete case!

Non-Uniform Samples

- This transforms uniform samples into non-uniform samples!

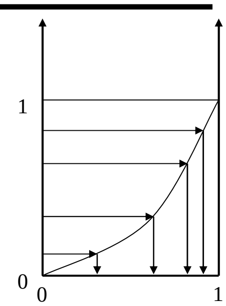

- Why? $\Pr(x \le y) = CDF(y) = \int_{-\infty}^{y} p(x)dx$

- Need:
  - CDF P(x)
  - Inverse CDF: P^{-1}

Importance Sampling

$$p(x) = \frac{f(x)}{\int_D f(x)}$$

- General principle:
  The closer the shape of p(x) is to the shape of f(x), the lower the variance
- Variance can also increase if p(x) is chosen badly

Numerical Example

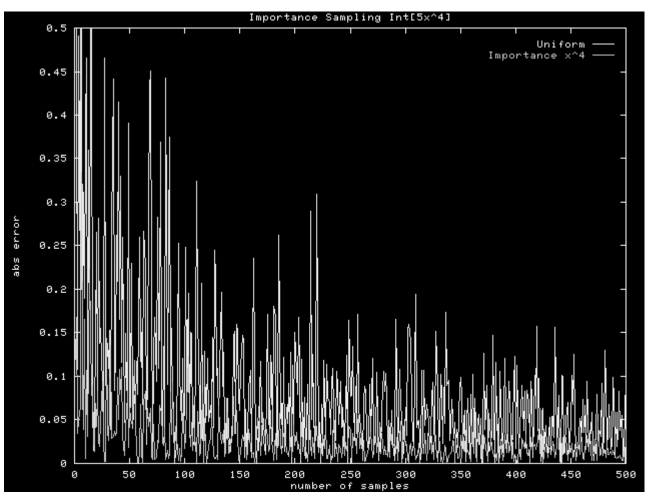

Monte Carlo Integration - 2D

- MC Integration works well for higher dimensions
- Unlike quadrature

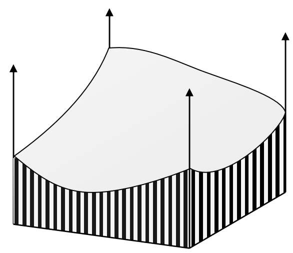

$$I = \int_a^b \int_c^d f(x, y) dx dy$$

$$\langle I \rangle = \frac{1}{N} \sum_{i=1}^{N} \frac{f(x_i, y_i)}{p(x_i, y_i)}$$

Example: How to sample p(x)

- Area of a circle: $A = \frac{1}{\pi} \int_0^{2\pi} \int_0^1 r dr d\theta = \frac{1}{\pi} \left[\frac{r^2}{2} \theta \right] = 1$

$$A = \int_0^{2\pi} \int_0^1 f(r, \theta) dr d\theta$$

$$f(r, \theta) = \frac{r}{\pi}$$

- Uniform sampling of r and θ

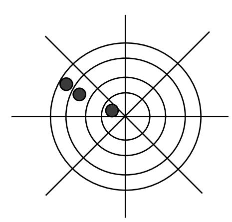

- Equal area sampling of r and θ

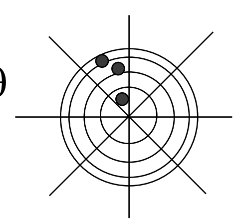

Example: How to sample p(x)

$$A = \int_0^{2\pi}\int_0^1 f(r,\theta)\,dr\,d\theta$$

$$p(r,\theta) = f(r,\theta) = \frac{r}{\pi}$$

$$CDF(r,\theta) = P(r,\theta) = \int_0^{\theta}\int_0^{r} \frac{r}{\pi}\,dr\,d\theta = r^2\,\frac{\theta}{2\pi}$$

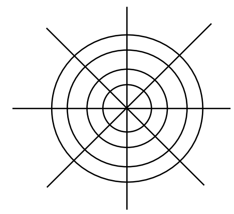

$$y = x^2 \rightarrow x = \sqrt{y} \Leftrightarrow P_r(r) = r^2, {P^{-1}}_r(r) = \sqrt{r}$$

$$x_\theta = {P^{-1}}_\theta(\xi_1) = 2\pi\xi_1$$

$$x_r = {P^{-1}}_r(\xi_2) = \sqrt{\xi_2}$$

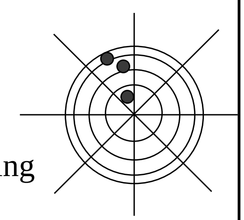

Equal area sampling

MC Integration - 2D example

- Integration over hemisphere:

$$I = \int_{\Omega} f(\Theta)\,d\omega_{\Theta}$$

$$= \int_0^{2\pi}\int_0^{\pi/2} f(\varphi,\theta)\sin\theta\,d\theta\,d\varphi$$

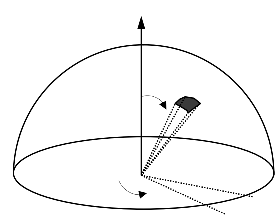

$$\langle I \rangle = \frac{1}{N}\sum_{i=1}^{N}\frac{f(\varphi_i,\theta_i)\sin\theta}{p(\varphi_i,\theta_i)}$$

Hemisphere Integration example

Irradiance due to light source:

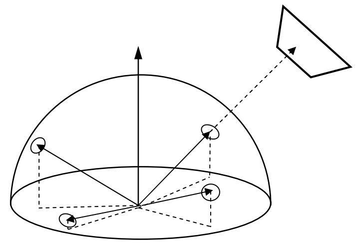

$$I = \int_{\Omega} L_{source} \cos\theta d\omega_{\Theta}$$

$$= \int_0^{2\pi}\int_0^{\pi/2} L_{source} \cos\theta \sin\theta d\theta d\varphi$$

$$p(\omega_i) = \frac{\cos\theta \sin\theta}{\pi}$$

$$\langle I \rangle = \frac{1}{N}\sum_{i=1}^{N} \frac{L_{source}(\omega_i)\cos\theta\sin\theta}{p(\omega_i)} = \frac{\pi}{N}\sum_{i=1}^{N} L_{source}(\omega_i)$$

Cosine distribution

$$f = \frac{1}{\pi}\int_0^{2\pi}\int_0^{1} \cos\theta \sin\theta d\theta d\phi$$

$$p(\theta,\phi) = \frac{\cos\theta\sin\theta}{\pi}$$

$$CDF(\theta,\phi) = \int_0^{\theta}\int_0^{\phi} \frac{\cos\theta\sin\theta}{\pi} drd\theta = (1-\cos^2\theta)\frac{\phi}{2\pi}$$

$$F(\theta) = 1-\cos^2\theta$$

$$F(\phi) = \frac{\phi}{2\pi}$$

$$\phi_i = 2\pi u_1 \qquad \theta_i = \cos^{-1}\sqrt{u_2}$$

Rejection Methods

- Pick ξ_1, ξ_2

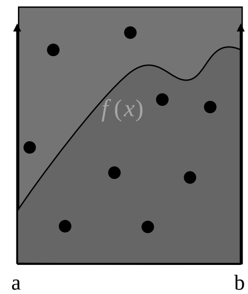

$$I = \int_a^b f(x)dx$$

- If $\xi_2 < f(\xi_1)$, select ξ_2
- Is this efficient? What determines efficiency? A(f)/A(rectangle)

MC Advantages

- Convergence rate of O($\frac{1}{\sqrt{N}}$)
- Simple
  - Sampling
  - Point evaluation
  - Can use black boxes
- General
  - Works for high dimensions
  - Deals with discontinuities, crazy functions,…