Advanced Computer Graphics (Fall 2009)

CS 294, Rendering Lecture 4: Monte Carlo Integration

Ravi Ramamoorthi

http://inst.eecs.berkeley.edu/~cs294-13/fa09

Acknowledgements and many slides courtesy:
Thomas Funkhouser, Szymon Rusinkiewicz and Pat Hanrahan

Motivation

Rendering = integration
- Reflectance equation: Integrate over incident illumination
- Rendering equation: Integral equation

Many sophisticated shading effects involve integrals
- Antialiasing
- Soft shadows
- Indirect illumination
- Caustics

Example: Soft Shadows

$$E(x) = \int_{H^2} L_i(x,\omega)\cos\theta\, d\omega$$

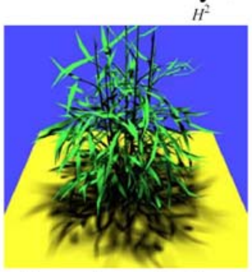

Challenges
- **Visibility and blockers**
- **Varying light distribution**
- **Complex source geometry**

Source: Agrawala. Ramamoorthi, Heirich, Moll, 2000

Monte Carlo

- Algorithms based on statistical sampling and random numbers
- Coined in the beginning of 1940s. Originally used for neutron transport, nuclear simulations
  - Von Neumann, Ulam, Metropolis, …
- Canonical example: 1D integral done numerically
  - Choose a set of random points to evaluate function, and then average (expectation or statistical average)

Monte Carlo Algorithms

Advantages
- Robust for complex integrals in computer graphics (irregular domains, shadow discontinuities and so on)
- Efficient for high dimensional integrals (common in graphics: time, light source directions, and so on)
- Quite simple to implement
- Work for general scenes, surfaces
- Easy to reason about (but care taken re statistical bias)

Disadvantages
- Noisy
- Slow (many samples needed for convergence)
- Not used if alternative analytic approaches exist (but those are rare)

Outline

- Motivation
- *Overview, 1D integration*
- Basic probability and sampling
- Monte Carlo estimation of integrals

Integration in 1D

$$\int_0^1 f(x)dx = ?$$

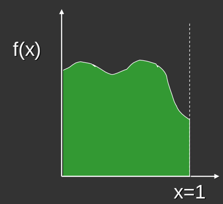

Slide courtesy of Peter Shirley

We can approximate

$$\int_0^1 f(x)dx \approx \int_0^1 g(x)dx$$

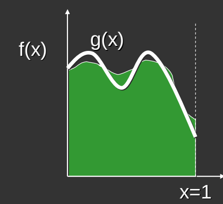

Standard integration methods like trapezoidal rule and Simpsons rule

Advantages:
- Converges fast for ***smooth*** integrands
- Deterministic

Disadvantages:
- Exponential complexity in many dimensions
- Not rapid convergence for discontinuities

Slide courtesy of Peter Shirley

Or we can average

$$\int_0^1 f(x)dx = E(f(x))$$

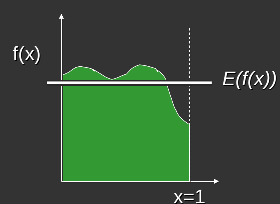

Slide courtesy of Peter Shirley

Estimating the average

$$\int_0^1 f(x)dx = \frac{1}{N}\sum_{i=1}^{N} f(x_i)$$

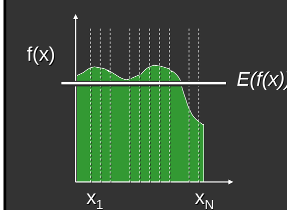

Monte Carlo methods (random choose samples)

Advantages:
- Robust for discontinuities
- Converges reasonably for large dimensions
- Can handle complex geometry, integrals
- Relatively simple to implement, reason about

Slide courtesy of Peter Shirley

Other Domains

$$\int_a^b f(x)dx = \frac{b-a}{N}\sum_{i=1}^{N} f(x_i)$$

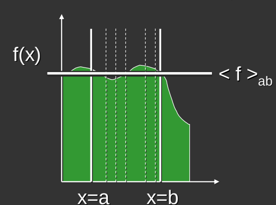

Slide courtesy of Peter Shirley

Multidimensional Domains

Same ideas apply for integration over …
- Pixel areas
- Surfaces
- Projected areas
- Directions
- Camera apertures
- Time
- Paths

$$\int_{UGLY} f(x)dx = \frac{1}{N}\sum_{i=1}^{N} f(x_i)$$

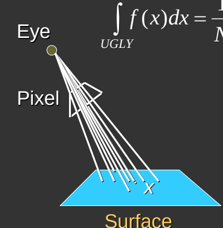

Outline

- Motivation
- Overview, 1D integration
- *Basic probability and sampling*
- Monte Carlo estimation of integrals

Random Variables

- Describes possible outcomes of an experiment
- In discrete case, e.g. value of a dice roll [x = 1-6]
- Probability p associated with each x (1/6 for dice)
- Continuous case is obvious extension

Expected Value

- Expectation Discrete: $E(x) = \sum_{i=1}^{n} p_i x_i$

  Continuous: $E(x) = \int_0^1 p(x) f(x)\, dx$

- For Dice example:

$$E(x) = \sum_{i=1}^{n} \frac{1}{6} x_i = \frac{1}{6}(1+2+3+4+5+6) = 3.5$$

Continuous Probability Distributions

PDF $p(x)$

$p(x) \geq 0$

CDF $P(x)$

$$P(x) = \int_0^x p(x) dx$$

$$P(x) = \Pr(X < x) \qquad P(1) = 1$$

$$\Pr(\alpha \leq X \leq \beta) = \int_\alpha^\beta p(x)\, dx$$

$$= P(\beta) - P(\alpha)$$

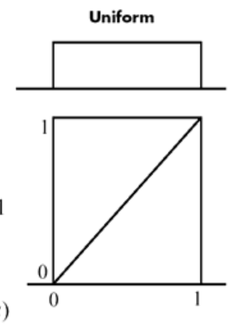

CS348B Lecture 6 Pat Hanrahan, Spring 2004

Sampling Techniques

Problem: how do we generate random points/directions during path tracing?

- Non-rectilinear domains
- Importance (BRDF)
- Stratified

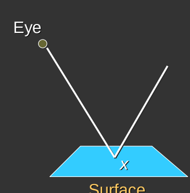

Generating Random Points

Uniform distribution:

- Use random number generator

Generating Random Points

Specific probability distribution:

- Function inversion
- Rejection
- Metropolis

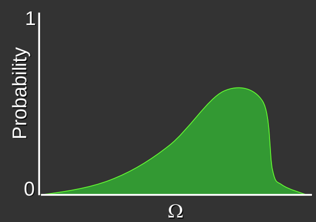

Common Operations

Want to ***sample*** probability distributions

- Draw samples distributed according to probability
- Useful for integration, picking important regions, etc.

Common distributions

- Disk or circle
- Uniform
- Upper hemisphere for visibility
- Area luminaire
- Complex lighting like an environment map
- Complex reflectance like a BRDF

Sampling Continuous Distributions

Cumulative probability distribution function

$$P(x) = \Pr(X < x)$$

Construction of samples

Solve for $X = P^{-1}(U)$

Must know:

1. **The integral of** $p(x)$
2. **The inverse function** $P^{-1}(x)$

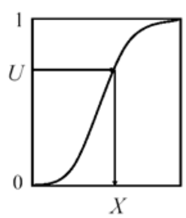

CS348B Lecture 6 — Pat Hanrahan, Spring 2004

Generating Random Points

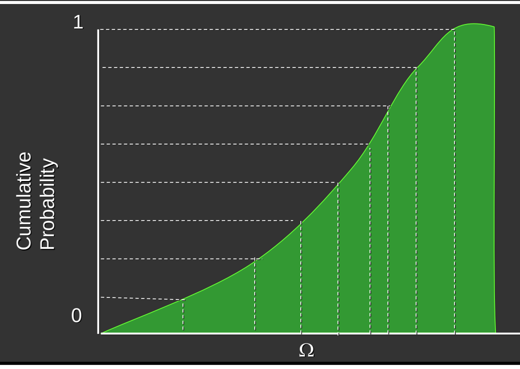

Example: Power Function

Assume

$$p(x) = (n+1)x^n \qquad \int_0^1 x^n dx = \left.\frac{x^{n+1}}{n+1}\right|_0^1 = \frac{1}{n+1}$$

$$P(x) = x^{n+1}$$

$$X \sim p(x) \Rightarrow X = P^{-1}(U) = \sqrt[n+1]{U}$$

Trick

$$Y = \max(U_1, U_2, \cdots, U_n, U_{n+1})$$

$$\Pr(Y < x) = \prod_{i=1}^{n+1} \Pr(U < x) = x^{n+1}$$

CS348B Lecture 6 — Pat Hanrahan, Spring 2004

Sampling a Circle

$$A = \int_0^{2\pi}\int_0^1 r\,dr\,d\theta = \int_0^1 r\,dr \int_0^{2\pi} d\theta = \left.\left(\frac{r^2}{2}\right)\right|_0^1 \theta\Big|_0^{2\pi} = \pi$$

$$p(r,\theta)\,dr\,d\theta = \frac{1}{\pi} r\,dr\,d\theta \Rightarrow p(r,\theta) = \frac{r}{\pi}$$

$$p(r,\theta) = p(r)p(\theta)$$

$$p(\theta) = \frac{1}{2\pi} \qquad \theta = 2\pi U_1$$

$$P(\theta) = \frac{1}{2\pi}\theta$$

$$p(r) = 2r \qquad r = \sqrt{U_2}$$

$$P(r) = r^2$$

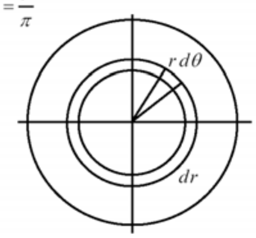

CS348B Lecture 6 — Pat Hanrahan, Spring 2004

Sampling a Circle

WRONG ≠ Equi-Areal

$\theta = 2\pi U_1$

$r = U_2$

RIGHT = Equi-Areal

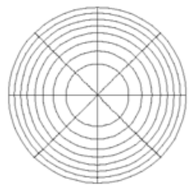

$\theta = 2\pi U_1$

$r = \sqrt{U_2}$

CS348B Lecture 6 — Pat Hanrahan, Spring 2004

Rejection Sampling

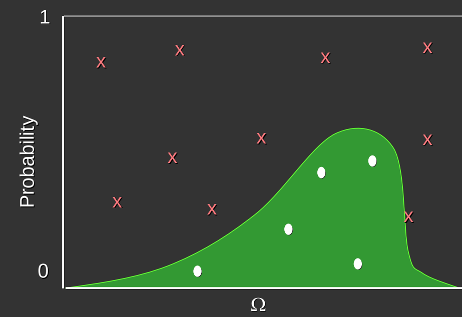

Rejection Methods

$$I = \int_0^1 f(x)\,dx$$

$$= \iint_{y<f(x)} dx\,dy$$

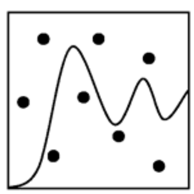

Algorithm

Pick U_1 **and** U_2

Accept U_1 **if** $U_2 < f(U_1)$

Wasteful? Efficiency = Area / Area of rectangle

CS348B Lecture 6 — Pat Hanrahan, Spring 2004

Sampling a Circle: Rejection

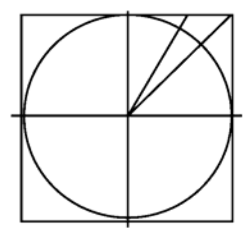

```
do {
   X=1-2*U1
   Y=1-2*U2
while( X²+ Y² >1 )
```

May be used to pick random 2D directions

Circle techniques may also be applied to the sphere

CS348B Lecture 6 — Pat Hanrahan, Spring 2004

Outline

- Motivation
- Overview, 1D integration
- Basic probability and sampling
- *Monte Carlo estimation of integrals*

Monte Carlo Path Tracing

Big diffuse light source, 20 minutes

Motivation for rendering in graphics: Covered in detail in next lecture

Jensen

Monte Carlo Path Tracing

1000 paths/pixel

Jensen

Estimating the average

$$\int_0^1 f(x)dx = \frac{1}{N}\sum_{i=1}^{N} f(x_i)$$

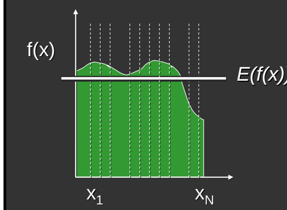

Monte Carlo methods (random choose samples)

Advantages:
- Robust for discontinuities
- Converges reasonably for large dimensions
- Can handle complex geometry, integrals
- Relatively simple to implement, reason about

Slide courtesy of Peter Shirley

Other Domains

$$\int_a^b f(x)dx = \frac{b-a}{N}\sum_{i=1}^{N} f(x_i)$$

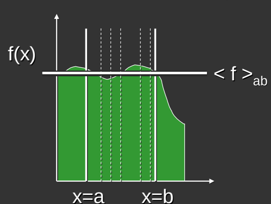

Slide courtesy of Peter Shirley

More formally

Definite integral $I(f) \equiv \int_0^1 f(x)\,dx$

Expectation of f $E[f] \equiv \int_0^1 f(x)p(x)\,dx$

Random variables $X_i \sim p(x)$

$Y_i = f(X_i)$

Estimator $F_N = \frac{1}{N}\sum_{i=1}^{N} Y_i$

Unbiased Estimator

$$E[F_N] = I(f)$$

Properties

$$E[\sum_i Y_i] = \sum_i E[Y_i]$$

$$E[aY] = aE[Y]$$

$$E[F_N] = E[\frac{1}{N}\sum_{i=1}^{N} Y_i]$$

$$= \frac{1}{N}\sum_{i=1}^{N} E[Y_i] = \frac{1}{N}\sum_{i=1}^{N} E[f(X_i)]$$

$$= \frac{1}{N}\sum_{i=1}^{N}\int_0^1 f(x)p(x)\,dx$$

$$= \frac{1}{N}\sum_{i=1}^{N}\int_0^1 f(x)\,dx$$

$$= \int_0^1 f(x)\,dx$$

Assume uniform probability distribution for now

CS348B Lecture 6 Pat Hanrahan, Spring 2004

Direct Lighting – Directional Sampling

$$E(x) = \int_{\Omega} L(x,\omega)\cos\theta\, d\omega$$

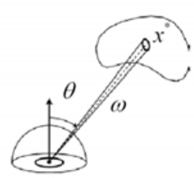

Ray intersection $x^*(x,\omega)$

Sample ω **uniformly by** Ω

$$Y_i = L(x^*(x,\omega_i), -\omega_i)\cos\theta\, 2\pi$$

CS348B Lecture 6 Pat Hanrahan, Spring 2004

Direct Lighting – Area Sampling

$$E(x) = \int_{\Omega} L_i(x,\omega)\cos\theta \, d\omega = \int_{A'} L_o(x',\omega') V(x,x') \frac{\cos\theta\cos\theta'}{|x-x'|^2} dA'$$

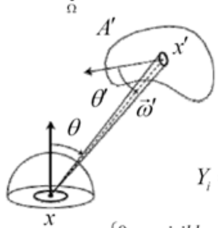

Ray direction $\omega' = x - x'$

Sample x' uniformly by A'

$$Y_i = L_o(x_i',\omega_i') V(x,x_i') \frac{\cos\theta \cos\theta_i'}{|x-x_i'|^2} A$$

$$V(x,x') = \begin{cases} 0 & \neg visible \\ 1 & visible \end{cases}$$

CS348B Lecture 6 Pat Hanrahan, Spring 2004

Variance

$$Var[f(x)] = \frac{1}{N}\sum_{i=1}^{N}[f(x_i) - E(f(x))]^2$$

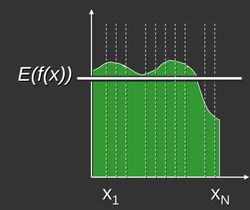

Variance

Definition

$$V[Y] \equiv E[(Y - E[Y])^2]$$
$$= E[Y^2 - 2YE[Y] + E[Y]^2]$$
$$= E[Y^2] - E[Y]^2$$

Properties

$$V[\sum_i Y_i] = \sum_i V[Y_i]$$

$$V[aY] = a^2 V[Y]$$

Variance decreases with sample size

$$V[\frac{1}{N}\sum_{i=1}^{N} Y_i] = \frac{1}{N^2}\sum_{i=1}^{N} V[Y_i] = \frac{1}{N} V[Y]$$

CS348B Lecture 6 Pat Hanrahan, Spring 2004

Variance for Dice Example?

- Work out on board (variance for single dice roll)

Variance

$$Var[E(f(x))] = \frac{1}{N} Var[f(x)]$$

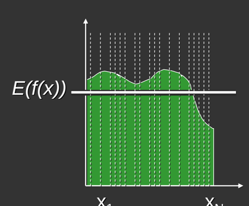

Variance decreases as 1/N
Error decreases as 1/sqrt(N)

Variance

- Problem: variance decreases with 1/N
  - Increasing # samples removes noise slowly

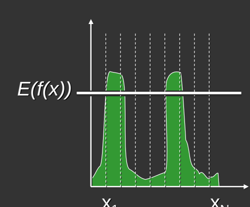

Variance Reduction

Efficiency measure

$$Efficiency \propto \frac{1}{Variance \bullet Cost}$$

Techniques

- **Importance sampling**
- **Sampling patterns: stratified, …**

Variance Reduction Techniques

- Importance sampling
- Stratified sampling

$$\int_0^1 f(x)dx = \frac{1}{N}\sum_{i=1}^{N} f(x_i)$$

Importance Sampling

Put more samples where f(x) is bigger

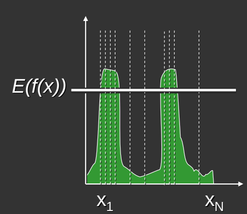

$$\int_\Omega f(x)dx = \frac{1}{N}\sum_{i=1}^{N} Y_i$$

$$Y_i = \frac{f(x_i)}{p(x_i)}$$

Importance Sampling

- This is still unbiased

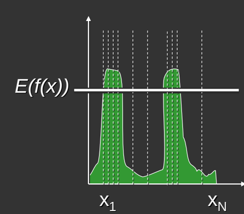

$$E[Y_i] = \int_\Omega Y(x)p(x)dx$$

$$= \int_\Omega \frac{f(x)}{p(x)}p(x)dx$$

$$= \int_\Omega f(x)dx$$

for all N

Importance Sampling

- Zero variance if p(x) ~ f(x)

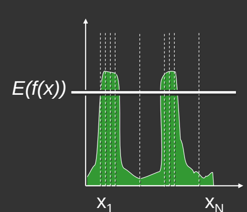

$$p(x) = cf(x)$$

$$Y_i = \frac{f(x_i)}{p(x_i)} = \frac{1}{c}$$

$$Var(Y) = 0$$

Less variance with better importance sampling

Stratified Sampling

- Estimate subdomains separately

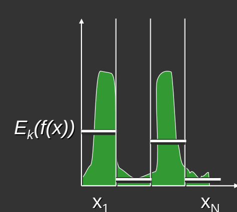

Arvo

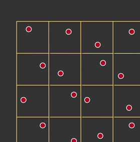

Stratified Sampling

- This is still unbiased

$$F_N = \frac{1}{N}\sum_{i=1}^{N} f(x_i)$$
$$= \frac{1}{N}\sum_{k=1}^{M} N_i F_i$$

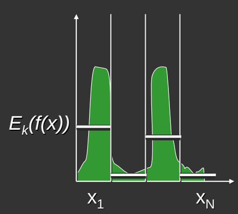

Stratified Sampling

- Less overall variance if less variance in subdomains

$$Var[F_N] = \frac{1}{N^2}\sum_{k=1}^{M} N_i Var[F_i]$$

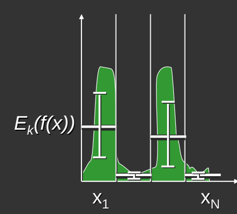

More Information

- Veach PhD thesis chapter (linked to from website)
- Course Notes (links from website)
  - *Mathematical Models for Computer Graphics*, Stanford, Fall 1997
  - *State of the Art in Monte Carlo Methods for Realistic Image Synthesis*, Course 29, SIGGRAPH 2001