What is a Monte Carlo Method?

A method involving deliberate use of random numbers in a calculation that has the structure of a stochastic process (Kalos & Whitlock)

A method designed to find the parameters of a distribution from observations (values) of a random variable (Glassner)

"If you have a program that calls drand48(), there's a good chance that you're doing Monte Carlo" (Shirley)

Probability Theory Refresher

A *continuous random variable* X takes random values from a continuous space S.

A *probability density function* (p.d.f.) $p(x) : S \rightarrow \mathbb{R}^1$ satisfies:

$$p(x) \geq 0 \text{ and } \int_S p(x)\, dx = 1$$

We say that $X \sim p(x)$ if for all $S_i \subseteq S$

$$\mathit{Pr}(X \in S_i) = \int_{S_i} p(x)\, dx$$

Q: What is the pdf if X is distributed uniformly over the interval $[a, b]$?

A: $p(x) = 1/(b - a)$

Refresher cont'd

The *mean (expected) value* of X is:

$$E[X] \;=\; \int_S x\, p(x)\, dx$$

The *variance* of X is:

$$V[X] = E[(X - E[X])^2] = E[X^2] - E[X]^2$$

The *standard deviation (error)* is:

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

Example: uniform distribution over $[a, b]$:

$$\begin{aligned} E[X] &= \int_a^b \frac{x}{b-a}\, dx = \frac{b^2 - a^2}{2(b-a)} = \frac{b+a}{2} \\ V[X] &= \frac{(b-a)^2}{12} \\ \sigma[X] &= \frac{b-a}{2\sqrt{3}} \approx 0.29(b-a) \end{aligned}$$

Law of Large Numbers

If the random variables X_i are independent and identically distributed (have the same pdf), then

$$Pr\left(E(X) = \lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} X_i\right) = 1$$

I.e., the mean of N measured samples approaches the expected value of the random variable X, as N goes to infinity.

Monte Carlo Integration

Let $X_1, X_2, \ldots, X_N$ be independent random variables with pdf $p(x)$. Define

$$G_N = \frac{1}{N}\sum_{i=1}^{N} f(X_i)$$

Then

$$\begin{aligned}
E[G_N] &= E\left[\frac{1}{N}\sum_{i=1}^{N} f(X_i)\right] \\
&= \frac{1}{N}\sum_{i=1}^{N} E[f(X_i)] \\
&= E[f(X_i)] \\
&= \int f(x)\, p(x)\, dx
\end{aligned}$$

Example

Let's compute the total flux leaving a point:

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

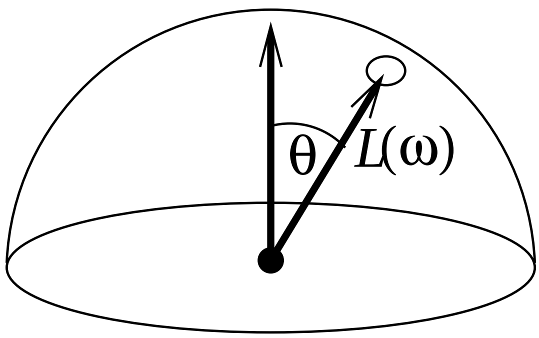

Let $f = L$ and $p = \frac{\cos\theta}{\pi}$. Then

$$\begin{aligned}\int_{\Omega} L(\omega)\ \cos\theta\, d\omega &= \pi \int_{\Omega} L(\omega)\, \frac{\cos\theta}{\pi}\, d\omega \\ &\approx \frac{\pi}{N} \sum_i L(X_i)\end{aligned}$$

where $X_i \sim \frac{\cos\theta}{\pi}$

Monte Carlo Integration

How do we compute the integral of f?

$$\begin{aligned} \int f(x)\,dx &= \int \frac{f(x)}{p(x)} p(x)\,dx \\ &= E\left[f(X)/p(X)\right] \\ &= E\left[\frac{1}{N}\sum_{i=1}^{N}\frac{f(X_i)}{p(X_i)}\right] \end{aligned}$$

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

Monte Carlo Integration

How good is the approximation?

Law of Large Numbers: approximation is accurate in the limit as $N \to \infty$.

Need to look at the variance:

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

Conclusion: standard error goes down as $\frac{1}{\sqrt{N}}$
Need to quadruple the number of samples to reduce the error by a factor of 2!

Variance Reduction: Importance

Want p to resemble $|f|$

Ideally, $p(x) = \frac{|f(x)|}{c}$. This gives 0 variance for single-signed functions f!

$$\int f(x)\,dx \approx \frac{1}{N} \sum_{i=1}^{N} \frac{cf(X_i)}{|f(X_i)|} = \pm c$$

Example: computing direct illumination

Variance Reduction: Stratification

Better coverage of the integrand: more difficult to miss important features.

Limits clumping of samples.

Effective if domain can be partitioned into strata within which the variance is smaller than the difference between their means.

Example

We'd like to integrate $f(x) = x^2 - x^3$ over the interval $[-1, 1]$:

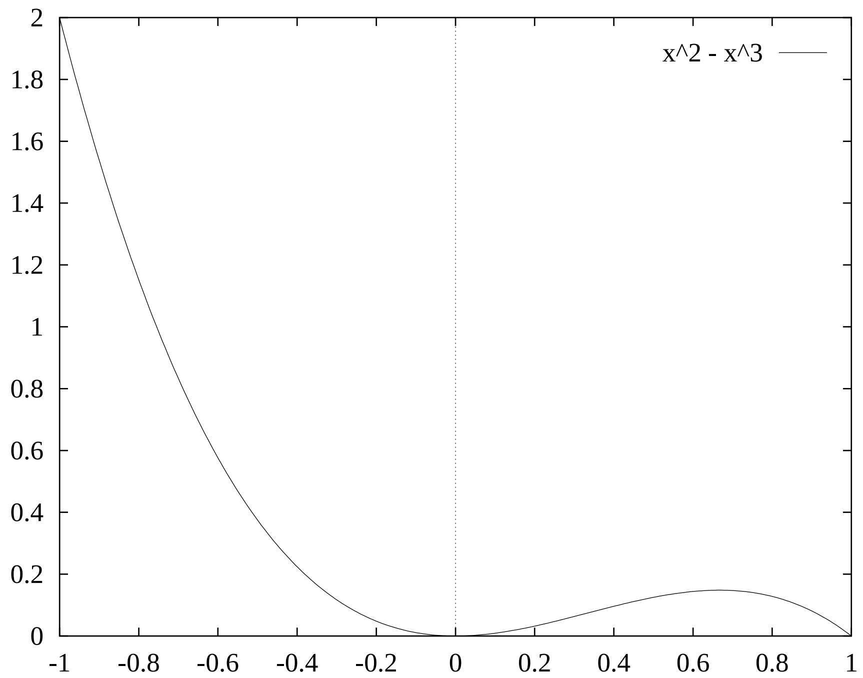

Analytical solution is 2/3. Gaussian quadrature (two-point rule) gives this solution.

Q: How well will Monte Carlo integration do?

Uniform Sampling

The pdf is $p(x) = 1/(b-a) = 1/2$.

$$\int_{-1}^{1} (x^2 - x^3)\, dx \approx \frac{1}{N} \sum_{i=1}^{N} \frac{\left(X_i^2 - X_i^3\right)}{1/2}$$

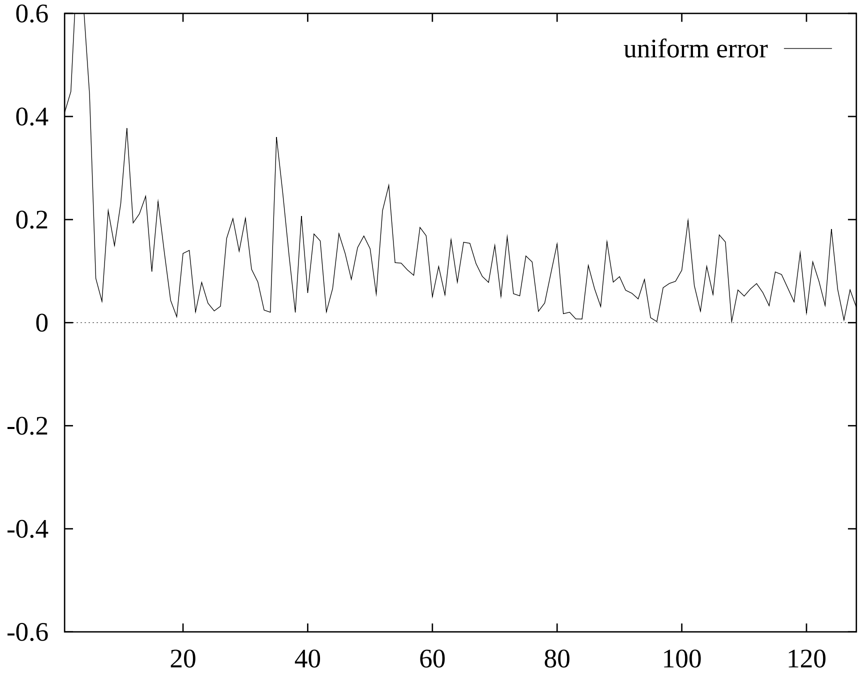

Importance Sampling

Using the pdf $p(x) = \frac{3}{2}x^2$:

$$\int_{-1}^{1} (x^2 - x^3)\, dx \approx \frac{1}{N} \sum_{i=1}^{N} \frac{\left(X_i^2 - X_i^3\right)}{3X_i^2/2}$$

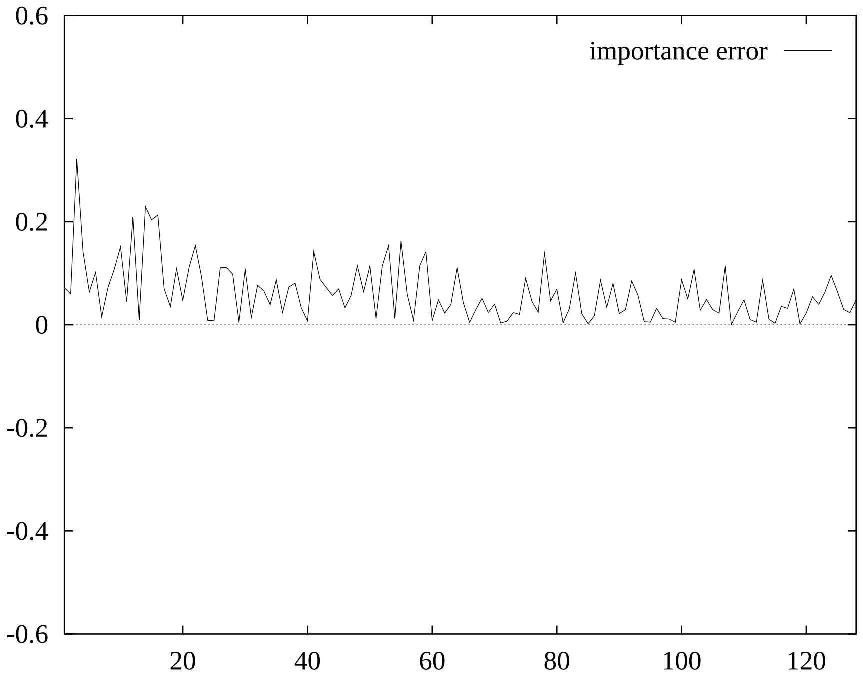

Stratified Sampling

Break $[-1, 1]$ into N equal intervals and pick a random point in each:

$$\int_{-1}^{1} (x^2 - x^3)\, dx \approx \frac{1}{N} \sum_{i=1}^{N} \frac{\left(X_i^2 - X_i^3\right)}{1/2}$$

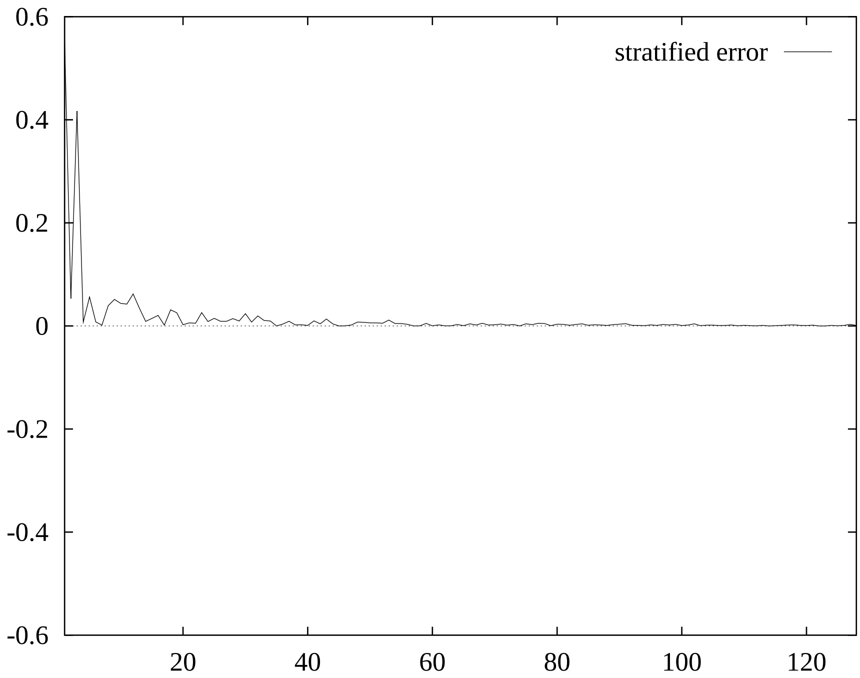

Stratified Importance Sampling

Intervals are distributed $\sim \frac{3}{2}x^2$:

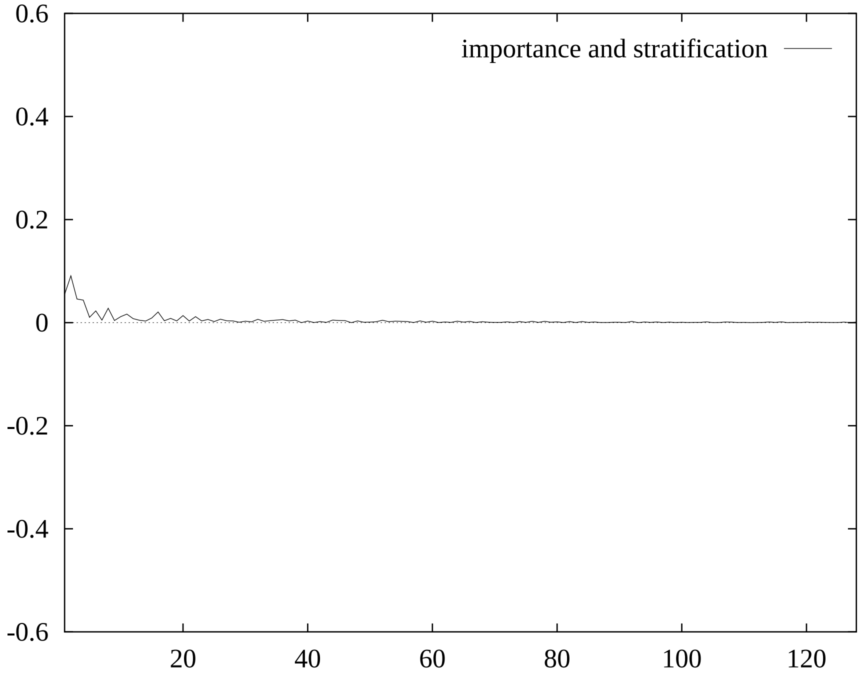

Random Number Generation

Suppose that we can generate random numbers uniformly distributed over $[0, 1]$.

Given $[a, b]$ and a pdf p, how can we generate random numbers distributed according to p over $[a, b]$?

Cumulative distribution function:

$$F(x) = \int_a^x p(y)\,dy$$

F is monotonically increasing;
$F(a) = 0$, $F(b) = 1$.

If X is uniformly distributed over $[0, 1]$, then $Y = F^{-1}(X)$ is distributed over $[a, b]$ according to p!

Random Number Generation cont'd

If X is uniformly distributed over $[0, 1]$, then $Y = F^{-1}(X)$ is such that

$$\begin{aligned} Pr(\alpha < Y < \beta) &= Pr\left(F(\alpha) < X < F(\beta)\right) \\ &= F(\beta) - F(\alpha) \\ &= \int_\alpha^\beta p(y)\, dy \end{aligned}$$

Another general technique *rejection sampling*: applicable to any distribution function, but expensive.

Particle Tracing

Repeatedly generate and track particles (representing photons) through the scene. For each particle:

Choose:
- a) wavelength
- b) location of origin
- c) direction

Repeat (until absorbed):
- a) find first surface hit
- b) determine outcome of interaction
- c) choose outgoing direction

Q: How do we render an image using particle tracing?

Solving the Rendering Equation

$$
\begin{aligned}
L(x_0) &= L^e(x_0) + \int_{x_1 \in \Omega} k(x_1, x_0) L(x_1)\, d\omega(x_1) \\
&= L^e(x_0) + \\
&\quad \int L^e(x_1) k(x_1, x_0)\, d\omega(x_1) + \\
&\quad \int L^e(x_2) k(x_2, x_1) k(x_1, x_0)\, d\omega(x_2)\, d\omega(x_1) + \cdots
\end{aligned}
$$

Can approximate $L(x_0)$ by Monte Carlo integration of each integral in the series: *Path Tracing*.

Bidirectional Path Tracing

To construct a patch of length k:

1. Generate m vertices starting from the light
2. Generate $k + 1 - m$ vertices from the eye
3. Connect ends deterministically

Estimate pixel value as sum of contributions of paths up to length k. (Do this several times and average to reduce the variance.)