

Computer Graphics III – Monte Carlo integration Direct illumination

Jaroslav Křivánek, MFF UK

Jaroslav.Krivanek@mff.cuni.cz


Entire the lecture in 5 slides

Reflection equation




$$L_r(\mathbf{x}, \omega_o) = \int_{H(\mathbf{x})} L_i(\mathbf{x}, \omega_i) \cdot f_r(\mathbf{x}, \omega_i \rightarrow \omega_o) \cdot \cos \theta_i \, d\omega_i$$

upper hemisphere
over \mathbf{x}



Rendering = Integration of functions

$$L_r(\mathbf{x}, \omega_o) = \int_{H(\mathbf{x})} L_i(\mathbf{x}, \omega_i) f_r(\mathbf{x}, \omega_i \rightarrow \omega_o) \cos \theta_i d\omega_i$$



■ Problems


- ❑ Discontinuous integrand (visibility)
- ❑ Arbitrarily large integrand values (e.g. light distribution in caustics, BRDFs of glossy surfaces)
- ❑ Complex geometry



Incoming radiance $L_i(\mathbf{x}, \omega_i)$ for a point on the floor.

Monte Carlo integration

- General tool for estimating definite integrals



Integral:

$$I = \int f(\mathbf{x}) d\mathbf{x}$$

Monte Carlo estimate I :

$$\langle I \rangle = \frac{1}{N} \sum_{i=1}^N \frac{f(\xi_i)}{p(\xi_i)}; \quad \xi_i \propto p(\mathbf{x})$$

Works “on average”:

$$E[\langle I \rangle] = I$$

Application of MC to reflection eq: Estimator of reflected radiance

- Integral to be estimated:

$$\int_{H(\mathbf{x})} L_i(\mathbf{x}, \omega_i) f_r(\mathbf{x}, \omega_i \rightarrow \omega_o) \cos \theta_i d\omega_i$$


- PDF for cosine-proportional sampling:

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

- **MC estimator:**

$$\begin{aligned} F_N &= \frac{1}{N} \sum_{k=1}^N \frac{\text{integrand}(\omega_{i,k})}{\text{pdf}(\omega_{i,k})} \\ &= \frac{\pi}{N} \sum_{k=1}^N L_i(\mathbf{x}, \omega_{i,k}) f_r(\mathbf{x}, \omega_{i,k} \rightarrow \omega_o) \end{aligned}$$

Variance => image noise



... and now the slow way

Digression: Numerical quadrature

Quadrature formulas for numerical integration

- General formula in 1D:

$$\hat{I} = \sum_{i=1}^n w_i f(x_i)$$

- f integrand (i.e. the integrated function)
 n quadrature order (i.e. number of integrand samples)
 x_i node points (i.e. positions of the samples)
 $f(x_i)$ integrand values at node points
 w_i quadrature weights

Quadrature formulas for numerical integration

- Quadrature rules differ by the choice of node point positions x_i and the weights w_i
 - E.g. rectangle rule, trapezoidal rule, Simpson's method, Gauss quadrature, ...
- The samples (i.e. the node points) are placed deterministically

Quadrature formulas in multiple dimensions

- General formula for quadrature of a function of multiple variables:

$$\hat{I} = \sum_{i_1=1}^n \sum_{i_2=1}^n \dots \sum_{i_s=1}^n w_{i_1} w_{i_2} \dots w_{i_s} f(x_{i_1}, x_{i_2}, \dots, x_{i_s})$$

- Convergence speed of approximation error E for an s -dimensional integral is $E = O(N^{-1/s})$
 - E.g. in order to cut the error in half for a 3-dimensional integral, we need $2^3 = 8$ – times more samples
- Unusable in higher dimensions
 - **Dimensional explosion**

Quadrature formulas in multiple dimensions

- **Deterministic quadrature vs. Monte Carlo**
 - ❑ In 1D deterministic better than Monte Carlo
 - ❑ In 2D roughly equivalent
 - ❑ From 3D, MC will always perform better
- Remember, quadrature rules are NOT the Monte Carlo method

Monte Carlo

History of the Monte Carlo method

- Atomic bomb development, Los Alamos 1940
John von Neumann, Stanislav Ulam, Nicholas Metropolis
- Further development and practical applications from the early 50's


Monte Carlo method

- We simulate many random occurrences of the same type of events, e.g.:
 - Neutrons – emission, absorption, collisions with hydrogen nuclei
 - Behavior of computer networks, traffic simulation.
 - Sociological and economical models – demography, inflation, insurance, etc.

Monte Carlo – applications

- Financial market simulations
- Traffic flow simulations
- Environmental sciences
- Particle physics
- Quantum field theory
- Astrophysics
- Molecular modeling
- Semiconductor devices
- Optimization problems
- **Light transport calculations**
- ...

Example: calculation of π



Area of square: $A_s = 1$

Area of circle: $A_c = \pi$

Fraction p of random points inside circle:

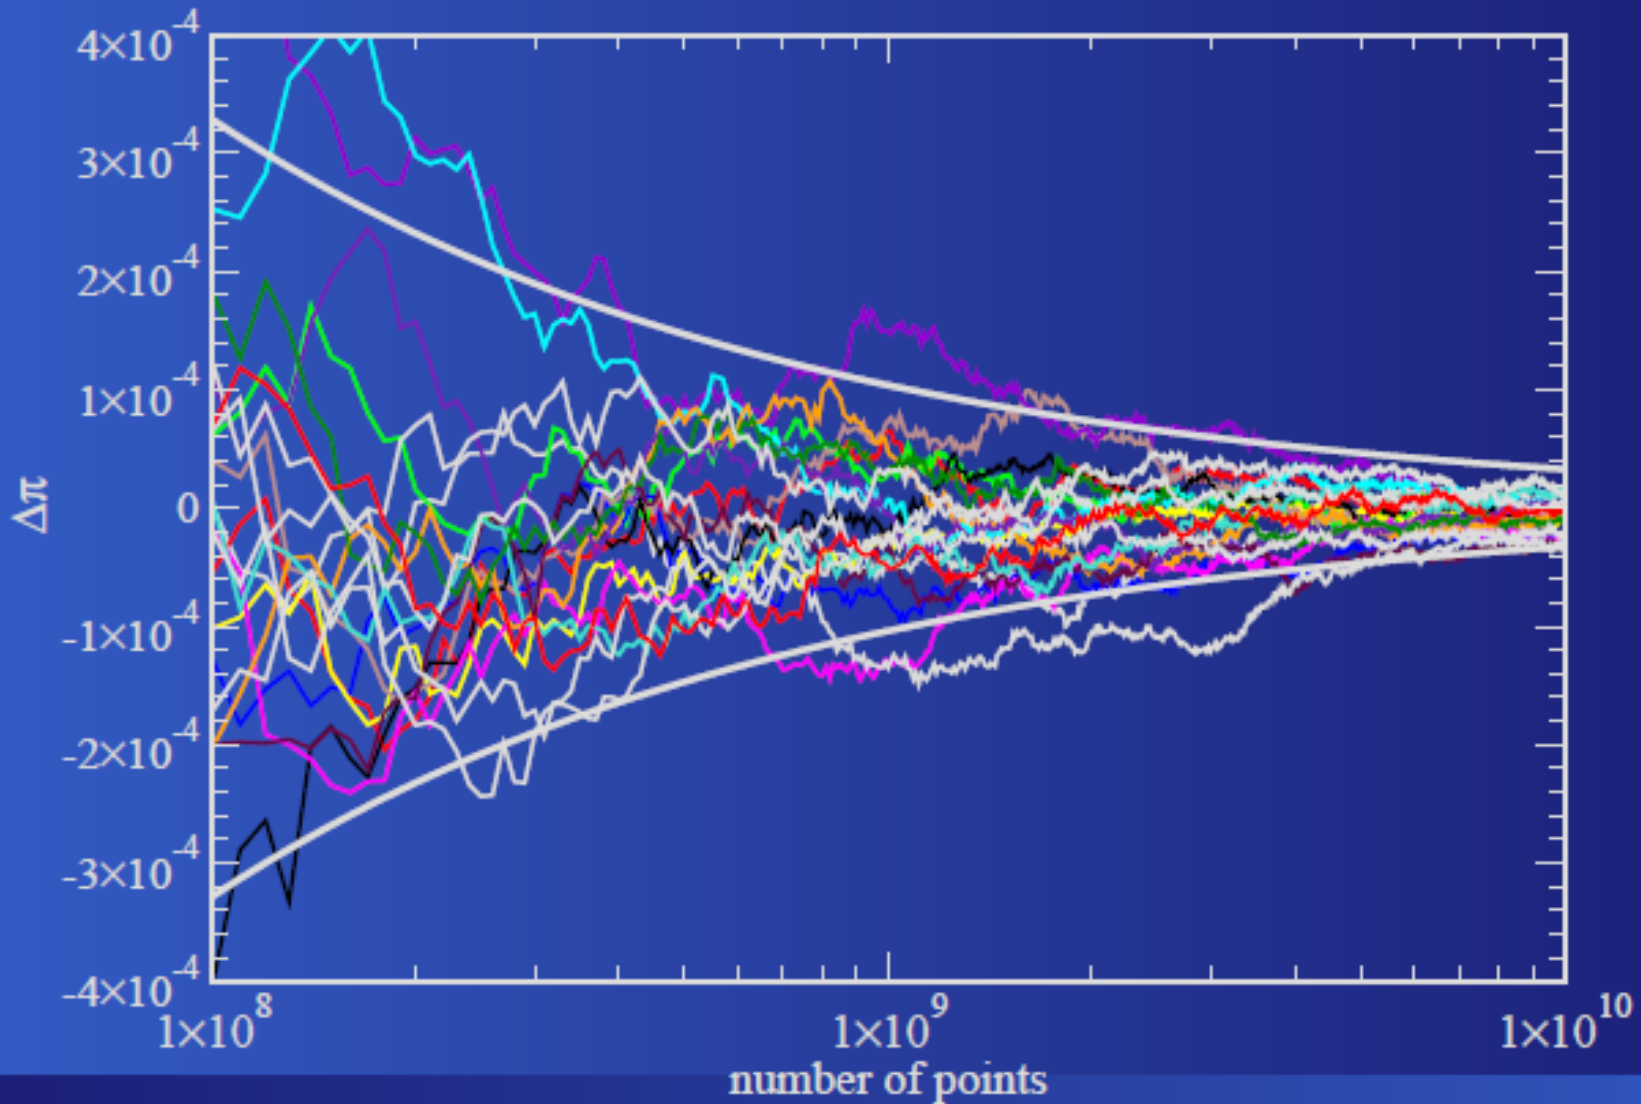
$$p = \frac{A_c}{A_s} = \frac{\pi}{4}$$

Random points: N



Random points inside circle: N_c

$$\Rightarrow \pi = \frac{4N_c}{N}$$

Calculation of π (cont'd)




Variance => image noise



Monte Carlo integration

- General tool for estimating definite integrals



Integral:

$$I = \int f(\mathbf{x}) d\mathbf{x}$$

Monte Carlo estimate I :

$$\langle I \rangle = \frac{1}{N} \sum_{i=1}^N \frac{f(\xi_i)}{p(\xi_i)}; \quad \xi_i \propto p(\mathbf{x})$$

Works “on average”:

$$E[\langle I \rangle] = I$$

Monte Carlo integration

- Samples are placed randomly (or pseudo-randomly)
- Convergence of standard error: std. dev. = $O(N^{-1/2})$
 - **Convergence speed independent of dimension**
 - **Faster than classic quadrature rules** for 3 and more dimensions
- Special methods for placing samples exist
 - Quasi-Monte Carlo
 - Faster asymptotic convergence than MC for “smooth” functions

Monte Carlo integration

■ Pros

- ❑ Simple implementation
- ❑ Robust solution for complex integrands and integration domains
- ❑ Effective for high-dimensional integrals

■ Cons

- ❑ Relatively slow convergence – halving the standard error requires four times as many samples
- ❑ In rendering: images contain noise that disappears slowly

Review – Random variables

Random variable

- X ... random variable
- X assumes different values with different probability
 - Given by the probability distribution D
 - $X \sim D$

Discrete random variable


- Finite set of values of x_i
- Each assumed with prob. p_i

$$p_i \equiv \Pr(X = x_i) \geq 0 \quad \sum_{i=1}^n p_i = 1$$


- **Cumulative distribution function**

$$P_i \equiv \Pr(X \leq x_i) = \sum_{j=1}^i p_j \quad P_n = 1$$

Probability mass function



Cumulative distribution func.



Continuous random variable

- Probability density function, **pdf**, $p(x)$

$$\Pr(X \in D) = \int_D p(x) dx$$

- In 1D:

$$\Pr(a < X \leq b) = \int_a^b p(t) dt$$

Continuous random variable

- Cumulative distribution function, **cdf**, $P(x)$

V 1D:

$$P(x) \equiv \Pr(X \leq x) = \int_{-\infty}^x p(t) dt$$

$$\Pr(X = a) = \int_a^a p(t) dt = 0!$$


Continuous random variable

Example: Uniform distribution

Probability density function (**pdf**)




Cumulative distribution function (**cdf**)



Continuous random variable

Gaussian (normal) distribution

Probability density function (**pdf**)



Cumulative distribution function (**cdf**)




Image: wikipedia

Expected value and variance

- **Expected value**

$$E[X] = \int_D \mathbf{x} p(\mathbf{x}) d\mathbf{x}$$

- **Variance**

$$\begin{aligned} V[X] &= E[(X - E[X])^2] \\ &= E[X^2 - 2XE[X] + E[X]^2] \\ &= E[X^2] - E[X]^2 \end{aligned}$$

- **Properties of variance**

$$V\left[\sum_i X_i\right] = \sum_i V[X_i] \quad (\text{if } X_i \text{ are independent})$$

$$V[aX] = a^2 V[X]$$

Transformation of a random variable

$$Y = f(X)$$


- Y is a random variable
- Expected value of Y

$$E[Y] = \int_D f(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}$$

Monte Carlo integration

Monte Carlo integration

- General tool for estimating definite integrals



Integral:

$$I = \int f(\mathbf{x}) d\mathbf{x}$$

Monte Carlo estimate I :

$$\langle I \rangle = \frac{1}{N} \sum_{i=1}^N \frac{f(\xi_i)}{p(\xi_i)}; \quad \xi_i \propto p(\mathbf{x})$$

Works “on average”:

$$E[\langle I \rangle] = I$$

Primary estimator of an integral


Integral to be estimated:

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

Let X be a random variable from the distribution with the pdf $p(x)$, then the random variable F_{prim} given by the transformation $f(X)/p(X)$ is called the **primary estimator** of the above integral.

$$F_{\text{prim}} = \frac{f(X)}{p(X)}$$

Primary estimator of an integral



Estimator vs. estimate

- **Estimator is a random variable**
 - It is defined through a transformation of another random variable
- **Estimate** is a concrete realization (outcome) of the estimator
- No need to worry: the above distinction is important for proving theorems but less important in practice

Unbiased estimator

- A general statistical estimator is called **unbiased** if – “on average” – it yields the correct value of an estimated quantity Q (without systematic error).
- More precisely:

$$E[F] = Q$$

Estimator of the quantity Q
(random variable)

Estimated quantity
(In our case, it is an integral, but in general it could be anything. It is a number, not a random variable.)

Unbiased estimator

The primary estimator F_{prim} is an unbiased estimator of the integral I .

Proof:

$$\begin{aligned} E[F_{\text{prim}}] &= \int_{\Omega} \frac{f(x)}{p(x)} p(x) dx \\ &= I \end{aligned}$$

Variance of the primary estimator

For an unbiased estimator, the error is due to **variance**:

$$\underline{V[F_{\text{prim}}]} = \underline{\sigma_{\text{prim}}^2} = E[F_{\text{prim}}^2] - E[F_{\text{prim}}]^2 = \int_{\Omega} \frac{f(x)^2}{p(x)} dx - I^2$$

(for an unbiased estimator)

If we use only a single sample, the variance is usually too high.
We need more samples in practice => secondary estimator.

Secondary estimator of an integral

- Consider N independent random variables X_i
- The estimator F_N given by the formula below is called the **secondary estimator** of I .

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

- The secondary estimator is unbiased.

Variance of the secondary estimator

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

... standard deviation is \sqrt{N} -times smaller!
(i.e. error converges with $1/\sqrt{N}$)

Properties of estimators

Unbiased estimator

- A general statistical estimator is called **unbiased** if – “on average” – it yields the correct value of an estimated quantity Q (without systematic error).
- More precisely:

$$E[F] = Q$$

Estimator of the quantity Q
(random variable)

Estimated quantity
(In our case, it is an integral, but in general it could be anything. It is a number, not a random variable.)

Bias of a biased estimator

- If

$$E[F] \neq Q$$

then the estimator is “**biased**” (cz: vychýlený).

- **Bias** is the systematic error of the estimator:

$$\beta = Q - E[F]$$

Consistency

- Consider a secondary estimator with N samples:

$$F_N = F_N(X_1, X_2, \dots, X_N)$$

- Estimator F_N is **consistent** if

$$Pr \left\{ \lim_{N \rightarrow \infty} F_N = Q \right\} = 1$$

i.e. if the error $F_N - Q$ converges to zero with probability 1.

Consistency

- Sufficient condition for consistency of an estimator:

$$\lim_{N \rightarrow \infty} \beta[F_N] = \lim_{N \rightarrow \infty} V[F_N] = 0$$

↑
bias

- Unbiasedness is not sufficient for consistency by itself (if the variance is infinite).
- But if the variance of a primary estimator finite, then the corresponding secondary estimator is necessarily consistent.

Rendering algorithms

■ **Unbiased**

- ❑ Path tracing
- ❑ Bidirectional path tracing
- ❑ Metropolis light transport

■ **Biased & Consistent**

- ❑ Progressive photon mapping

■ **Biased & not consistent**

- ❑ Photon mapping
- ❑ Irradiance / radiance caching

Mean Squared Error – MSE

(cz: Střední kvadratická chyba)

■ Definition

$$MSE[F] = E[(F - Q)^2]$$

■ Proposition

$$MSE[F] = V[F] + \beta[F]^2$$

□ Proof

$$\begin{aligned} MSE[F] &= E[(F - Q)^2] \\ &= E[(F - E[F])^2] + 2E[F - E[F]](E[F] - Q) + (E[F] - Q)^2 \\ &= V[F] + \beta[F]^2, \end{aligned}$$

Mean Squared Error – MSE

(cz: Střední kvadratická chyba)

- If the estimator F is unbiased, then

$$MSE[F] = V[F]$$

i.e. for an unbiased estimator, it is much easier to estimate the error, because it can be estimated directly from the samples $Y_i = f(X_i) / p(X_i)$.

- Unbiased **estimator of variance**

$$\hat{V}[F_N] = \frac{1}{N-1} \left\{ \left(\frac{1}{N} \sum_{i=1}^N Y_i^2 \right) - \left(\frac{1}{N} \sum_{i=1}^N Y_i \right)^2 \right\}$$

Efficiency of an estimator

- **Efficiency** of an unbiased estimator is given by:

$$\epsilon[F] = \frac{1}{V[F] T[F]}$$

variance

Calculation time
(i.e. operations count, such as number of cast rays)

MC estimators for illumination calculation

Irradiance estimator – uniform sampling

- Integral to be estimated:


$$E(\mathbf{x}) = \int_{H(\mathbf{x})} L_i(\mathbf{x}, \omega_i) \cdot \cos \theta_i \, d\omega_i$$

- PDF for uniform sampling:

$$p(\omega) = \frac{1}{2\pi}$$

- Estimator:

$$\begin{aligned} F_N &= \frac{1}{N} \sum_{k=1}^N \frac{f(\omega_{i,k})}{p(\omega_{i,k})} \\ &= \frac{2\pi}{N} \sum_{k=1}^N L_i(\mathbf{x}, \omega_{i,k}) \cdot \cos \theta_{i,k} \end{aligned}$$



Irradiance estimator – cosine-proportional sampling

- Integral to be estimated:


$$E(\mathbf{x}) = \int_{H(\mathbf{x})} L_i(\mathbf{x}, \omega_i) \cdot \cos \theta_i \, d\omega_i$$

- PDF for cosine-proportional sampling:


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

- Estimator

$$\begin{aligned} F_N &= \frac{1}{N} \sum_{k=1}^N \frac{f(\omega_{i,k})}{p(\omega_{i,k})} \\ &= \frac{\pi}{N} \sum_{k=1}^N L_i(\mathbf{x}, \omega_{i,k}) \end{aligned}$$



Irradiance estimate – light source sampling




Irradiance estimate – light source sampling

- Reformulate the reflection integral (change of variables)

$$E(\mathbf{x}) = \int_{H(\mathbf{x})} L_i(\mathbf{x}, \omega_i) \cdot \cos \theta_i \, d\omega_i$$
$$= \int_A L_e(\mathbf{y} \rightarrow \mathbf{x}) \cdot V(\mathbf{y} \leftrightarrow \mathbf{x}) \cdot \frac{\cos \theta_y \cdot \cos \theta_x}{\|\mathbf{y} - \mathbf{x}\|^2} \, dA$$

$G(\mathbf{y} \leftrightarrow \mathbf{x})$




- PDF for uniform sampling of the surface area:

$$p(\mathbf{y}) = \frac{1}{|A|}$$


- **Estimator**

$$F_N = \frac{|A|}{N} \sum_{k=1}^N L_e(\mathbf{y}_k \rightarrow \mathbf{x}) \cdot V(\mathbf{y}_k \leftrightarrow \mathbf{x}) \cdot G(\mathbf{y}_k \leftrightarrow \mathbf{x})$$

Light source vs. cosine sampling





Light source **area sampling**



Cosine-proportional sampling

Images: Pat Hanrahan

Example – Area Sampling





1 shadow ray per eye ray

Center

Random

Example – Area Sampling




16 shadow rays per eye ray

Uniform grid

Stratified random

Area light sources



1 sample per pixel

9 samples per pixel

36 samples per pixel

Direct illumination on a surface with an arbitrary BRDF

- Integral to be estimated

$$L_o(\mathbf{x}, \omega_o) = \int_A L_e(\mathbf{y} \rightarrow \mathbf{x}) \cdot f_r(\mathbf{y} \rightarrow \mathbf{x} \rightarrow \omega_o) \cdot V(\mathbf{y} \leftrightarrow \mathbf{x}) \cdot G(\mathbf{y} \leftrightarrow \mathbf{x}) dA$$

- **Estimator** based on uniform light source sampling

$$F_N = \frac{|A|}{N} \sum_{k=1}^N L_e(\mathbf{y}_k \rightarrow \mathbf{x}) \cdot f_r(\mathbf{y}_k \rightarrow \mathbf{x} \rightarrow \omega_o) \cdot V(\mathbf{y}_k \leftrightarrow \mathbf{x}) \cdot G(\mathbf{y}_k \leftrightarrow \mathbf{x})$$