

Lecture 15:

Numerical Integration

(with a focus on Monte Carlo integration)

Computer Graphics


CMU 15-462/15-662, Spring 2016

Review: fundamental theorem of calculus

$$\int_a^b f(x) dx = F(b) - F(a)$$


$$f(x) = \frac{d}{dx} F(x)$$

$$\int_a^x f(t) dt = F(x) - F(a)$$




Definite integral as “area under curve”

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




Simple case: constant function

$$\int_a^b C dx = (b - a)C$$



Affine function: $f(x) = cx + d$


$$\int_a^b f(x) dx = \frac{1}{2}(f(a) + f(b))(b - a)$$



Piecewise affine function

Sum of integrals of individual affine components

$$\int_a^b f(x) dx = \frac{1}{2} \sum_{i=0}^{n-1} (x_{i+1} - x_i) (f(x_i) + f(x_{i+1}))$$



Piecewise affine function


If N-1 segments are of equal length: $h = \frac{b - a}{n - 1}$

$$\int_a^b f(x) dx = \frac{h}{2} \sum_{i=0}^{n-1} (f(x_i) + f(x_{i+1}))$$


$$= h \left(\sum_{i=1}^{n-1} f(x_i) + \frac{1}{2} (f(x_0) + f(x_n)) \right)$$

Weighted combination
of measurements.

$$= \sum_{i=0}^n A_i f(x_i)$$



Polynomials?



Aside: interpolating polynomials


Consider $n+1$ measurements of a function $f(x)$

$$f(x_0), f(x_1), f(x_2), \dots, f(x_n)$$

There is a unique degree $\leq n$ polynomial that interpolates the points:

$$\begin{aligned} p(x) &= \sum_{i=0}^n f(x_i) \prod_{j \neq i, j=0}^n \left(\frac{x - x_j}{x_i - x_j} \right) \\ &= \sum_{i=0}^n f(x_i) l_i(x) \end{aligned}$$

Note: $l_i(x)$ is 1 at x_i and 0 at all other measurement points



Gaussian quadrature theorem

If $f(x)$ is a polynomial of degree of up to $2n+1$, then its integral over $[a,b]$ is computed exactly by a weighted combination of $n+1$ measurements in this range.


$$\int_a^b f(x)dx = \sum_{i=0}^n A_i f(x_i) \qquad A_i = \int_a^b l_i(x)dx$$

Where are these points?

Roots of degree $n+1$ polynomial $q(x)$ where:

$$\int_a^b x^k q(x)dx = 0 \qquad 0 \leq k \leq n$$

Arbitrary function $f(x)$?




Trapezoidal rule

Approximate integral of $f(x)$ by assuming function is piecewise linear

For equal length segments: $h = \frac{b - a}{n - 1}$

$$\int_a^b f(x) dx = h \left(\sum_{i=1}^{n-1} f(x_i) + \frac{1}{2} (f(x_0) + f(x_n)) \right)$$




Trapezoidal rule

Consider cost and accuracy of estimate as $n \rightarrow \infty$ (or $h \rightarrow 0$)

Work: $O(n)$

Error can be shown to be: $O(h^2) = O(\frac{1}{n^2})$
(for $f(x)$ with continuous second derivative)



Integration in 2D

Consider integrating $f(x, y)$ using the trapezoidal rule
(apply rule twice: when integrating in x and in y)

$$\begin{aligned}\int_{a_y}^{b_y} \int_{a_x}^{b_x} f(x, y) dx dy &= \int_{a_y}^{b_y} \left(O(h^2) + \sum_{i=0}^n A_i f(x_i, y) \right) dy && \text{First application of rule} \\ &= O(h^2) + \sum_{i=0}^n A_i \int_{a_y}^{b_y} f(x_i, y) dy \\ &= O(h^2) + \sum_{i=0}^n A_i \left(O(h^2) + \sum_{j=0}^n A_j f(x_i, y_j) \right) && \text{Second application} \\ &= O(h^2) + \sum_{i=0}^n \sum_{j=0}^n A_i A_j f(x_i, y_j)\end{aligned}$$

Errors add, so error still: $O(h^2)$

But work is now: $O(n^2)$

(n x n set of measurements)

Must perform much more work in 2D to get same error bound on integral!


In K-D, let $N = n^k$

Error goes as: $O\left(\frac{1}{N^{2/k}}\right)$

Recall: camera measurement equation from last time

$$Q = \frac{1}{d^2} \int_{t_0}^{t^1} \int_{A_{\text{lens}}} \int_{A_{\text{film}}} L(p' \rightarrow p, t) \cos^4 \theta \, dp \, dp' \, dt'$$

5D integral!



(Rendering requires computation of infinite dimensional integrals. Coming soon in class!)

Monte Carlo Integration

Slides credit: a majority of these slides were created by Matt Pharr and Pat Hanrahan

Monte Carlo numerical integration

So far we've discussed techniques that use a fixed set of sample points (e.g., uniformly spaced, or obtained by finding roots of polynomial (Gaussian quadrature))

- Estimate value of integral using random sampling of function 
 - Value of estimate depends on random samples used
 - But algorithm gives the correct value of integral “on average”
- Only requires function to be evaluated at random points on its domain
 - Applicable to functions with discontinuities, functions that are impossible to integrate directly
- Error of estimate is independent of the dimensionality of the integrand
 - Depends on the number of random samples used: $O(n^{1/2})$

Recall previous trapezoidal rule example: $O(n^{-1/k})$
(dropping the n^2 for simplicity)

Review: random variables

X **random variable. Represents a distribution of potential values**

$X \sim p(x)$ **probability density function (PDF). Describes relative probability of a random process choosing value x**

Uniform PDF: all values over a domain are equally likely

e.g., for an unbiased die

X takes on values 1,2,3,4,5,6


$$p(1) = p(2) = p(3) = p(4) = p(5) = p(6)$$



Discrete probability distributions

n discrete values x_i

With probability p_i



Requirements of a PDF:

$$p_i \geq 0$$

$$\sum_{i=1}^n p_i = 1$$

Six-sided die example: $p_i = \frac{1}{6}$

Think: p_i is the probability that a random measurement of X will yield the value x_i
 X takes on the value x_i with probability p_i

Cumulative distribution function (CDF)



(For a discrete probability distribution)

Cumulative PDF: $P_j = \sum_{i=1}^j p_i$

where:

$$0 \leq P_i \leq 1$$

$$P_n = 1$$



How do we generate samples of a discrete random variable (with a known PDF?)



Sampling from discrete probability distributions

To randomly select an event,
select x_i if

$$P_{i-1} < \xi \leq P_i$$



Uniform random variable $\in [0, 1)$



Continuous probability distributions

PDF $p(x)$

$$p(x) \geq 0$$

CDF $P(x)$

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



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

$$P(1) = 1$$

$$\begin{aligned} \Pr(a \leq X \leq b) &= \int_a^b p(x) \, dx \\ &= P(b) - P(a) \end{aligned}$$

Uniform distribution

(for random variable X defined on $[0,1]$ domain)



Sampling continuous random variables using the inversion method

Cumulative probability distribution function


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

Construction of samples:

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

Must know the formula for:

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




Example: applying the inversion method

Given:

$$f(x) = x^2 \quad x \in [0, 2]$$

Relative density of probability
of random variable taking on
value x over $[0, 2]$ domain



Compute PDF:

$$\begin{aligned} 1 &= \int_0^2 c f(x) dx \\ &= c(F(2) - F(0)) \end{aligned}$$

$$= c \frac{1}{3} 2^3$$

$$= \frac{8c}{3} \longrightarrow c = \frac{3}{8},$$

$$p(x) = \frac{3}{8} x^2$$


Probability density function
(integrates to 1)

Example: applying the inversion method

Given:


$$f(x) = x^2 \quad x \in [0, 2]$$

$$p(x) = \frac{3}{8}x^2$$



Compute CDF:

$$\begin{aligned} P(x) &= \int_0^x p(x) \, dx \\ &= \frac{x^3}{8} \end{aligned}$$



Example: applying the inversion method

Given:

$$f(x) = x^2 \quad x \in [0, 2]$$


$$p(x) = \frac{3}{8}x^2$$

$$P(x) = \frac{x^3}{8}$$

Sample from $p(x)$

$$\xi = P(x) = \frac{x^3}{8}$$

$$x = \sqrt[3]{8\xi}$$



How do we uniformly sample the unit circle?

(Choose any point $P=(p_x, p_y)$ in circle with equal probability)


Uniformly sampling unit circle: first try

- θ = uniform random angle between 0 and 2π
- r = uniform random radius between 0 and 1
- Return point: $(r \cos \theta, r \sin \theta)$


This algorithm does not produce the desired uniform sampling of the area of a circle. Why?

Because sampling is not uniform in area!

Points farther from center of circle are less likely to be chosen



$$\theta = 2\pi\xi_1 \quad r = \xi_2$$



$$p(r, \theta) dr d\theta \sim r dr d\theta$$

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

Sampling a circle (via inversion in 2D)

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

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

$$p(r, \theta) = p(r)p(\theta) \quad \leftarrow r, \theta \text{ independent}$$

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


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

$$p(r) = 2r$$

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


$$\theta = 2\pi \xi_1$$

$$r = \sqrt{\xi_2}$$



Uniform area sampling of a circle


WRONG
Not Equi-areaal



$$\theta = 2\pi\xi_1$$

$$r = \xi_2$$



RIGHT
Equi-areaal





$$\theta = 2\pi\xi_1$$

$$r = \sqrt{\xi_2}$$

Shirley's mapping




Distinct cases for eight octants



$$r = \xi_1$$
$$\theta = \frac{\pi \xi_2}{4r}$$

Uniform sampling via rejection sampling




Generate random point within unit circle

```
do {  
    x = 1 - 2 * rand01();  
    y = 1 - 2 * rand01();  
} while (x*x + y*y > 1.);
```


Efficiency of technique: area of circle / area of square

Aside: approximating the area of a circle



```
inside = 0
for (i = 0; i < N; ++i) {
    x = 1 - 2 * rand01();
    y = 1 - 2 * rand01();
    if (x*x + y*y < 1.)
        ++inside;
}
A = inside * 4 / N;
```

Rejection sampling to generate 2D directions




**Goal: generate random directions
in 2D with uniform probability**

```
x = 1 - 2 * rand01();  
y = 1 - 2 * rand01();
```

```
r = sqrt(x*x+y*y);  
x_dir = x/r;  
y_dir = y/r;
```

This algorithm is not correct! What is wrong?

Rejection sampling to generate 2D directions



Goal: generate random directions in 2D with uniform probability

```
do {  
    x = 1 - 2 * rand01();  
    y = 1 - 2 * rand01();  
} while (x*x + y*y > 1.);  
  
r = sqrt(x*x+y*y);  
x_dir = x/r;  
y_dir = y/r;
```

Monte Carlo integration

■ Definite integral

What we seek to estimate

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

■ Random variables

X_i is the value of a random sample
drawn from the distribution $p(x)$

Y_i is also a random variable.

$$X_i \sim p(x)$$

$$Y_i = f(X_i)$$

■ Expectation of f

$$E[Y_i] = E[f(X_i)] = \int_a^b f(x) p(x) dx$$

■ Estimator

Monte Carlo estimate of $\int_a^b f(x) dx$

$$F_N = \frac{b - a}{N} \sum_{i=1}^N Y_i$$

Assuming samples X_i drawn from uniform
pdf. I will provide estimator for arbitrary
PDFs later in lecture.

Basic unbiased Monte Carlo estimator

**Unbiased estimator:
Expected value of
estimator is the integral
we wish to evaluate.**

$$E[F_N] = E \left[\frac{b-a}{N} \sum_{i=1}^N Y_i \right]$$
$$= \frac{b-a}{N} \sum_{i=1}^N E[Y_i] = \frac{b-a}{N} \sum_{i=1}^N E[f(X_i)]$$

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

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

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

**Assume uniform
probability density for now**

$$X_i \sim U(a, b)$$

$$p(x) = \frac{1}{b-a}$$

Properties of expectation:

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


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

Direct lighting estimate

Uniformly-sample hemisphere of directions with respect to solid angle

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

$$E(p) = \int L(p, \omega) \cos \theta \, d\omega$$



Estimator:

$$X_i \sim p(\omega)$$

$$Y_i = f(X_i)$$

$$Y_i = L(p, \omega_i) \cos \theta_i$$

$$F_N = \frac{2\pi}{N} \sum_{i=1}^N Y_i$$

Direct lighting estimate

Uniformly-sample hemisphere of directions with respect to solid angle

$$E(p) = \int L(p, \omega) \cos \theta \, d\omega$$

Given surface point p

A ray tracer evaluates radiance along a ray
(see `Raytracer::trace_ray()` in `raytracer.cpp`)

For each of N samples:

Generate random direction: ω_i

Compute incoming radiance arriving L_i at p from direction: ω_i

Compute incident irradiance due to ray: $dE_i = L_i \cos \theta_i$

Accumulate $\frac{2\pi}{N} dE_i$ into estimator


**Hemispherical solid angle
sampling, 100 sample rays
(random directions drawn
uniformly from hemisphere)**




Light source



**Occluder
(blocks light)**



Why is the image in the previous slide “noisy”?



Incident lighting estimator uses different random directions in each pixel. Some of those directions point towards the light, others do not.

(Estimator is a random variable)

Idea: don't need to integrate over entire hemisphere of directions (incoming radiance is 0 from most directions)


**Only integrate over the area of the light
(directions where incoming radiance is non-zero)**

Direct lighting: area integral

$$E(p) = \int L(p, \omega) \cos \theta \, d\omega \quad \leftarrow \text{Integral over directions}$$

$$E(p) = \int_{A'} L_o(p', \omega') V(p, p') \frac{\cos \theta \cos \theta'}{|p - p'|^2} \, dA' \quad \leftarrow \text{Change of variables to integral over area of light}^*$$

$$d\omega = \frac{dA}{|p' - p|^2} = \frac{dA' \cos \theta}{|p' - p|^2}$$



Binary visibility function:
1 if p' is visible from p , 0 otherwise
(accounts for light occlusion)

Outgoing radiance from light point p , in direction w' towards p


Direct lighting: area integral

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

Sample shape uniformly by area A'


$$\int_{A'} p(p') dA' = 1$$

$$p(p') = \frac{1}{A'}$$



Direct lighting: area integral

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



Probability:


$$p(p') = \frac{1}{A'}$$

Estimator

$$Y_i = L_o(p'_i, \omega'_i) V(p, p'_i) \frac{\cos \theta_i \cos \theta'_i}{|p - p'_i|^2}$$

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

**Light source area
sampling, 100 sample rays**



**If no occlusion is present, all directions chosen in computing estimate “hit” the light source.
(Choice of direction only matters if portion of light is occluded from surface point p.)**

Variance

■ Definition

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

■ Variance decreases linearly with number of samples


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

Properties of variance:


$$V\left[\sum_{i=1}^N Y_i\right] = \sum_{i=1}^N V[Y_i]$$

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

1 area light sample
(high variance in irradiance estimate)



16 area light samples
(high variance in irradiance estimate)



Comparing different techniques

- **Variance in an estimator manifests as noise in rendered images**
- **Estimator efficiency measure:**

$$\text{Efficiency} \propto \frac{1}{\text{Variance} \times \text{Cost}}$$

- **If one integration technique has twice the variance as another, then it takes twice as many samples to achieve the same variance**
- **If one technique has twice the cost of another technique with the same variance, then it takes twice as much time to achieve the same variance**

“Biasing”

- We previously used a uniform probability distribution to generate samples in our estimator
- Idea: change the distribution—bias the selection of samples

$$X_i \sim p(x)$$

- However, for estimator to remain unbiased, must change the estimator to:

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

- Note: “biasing” selection of random samples is different than creating a biased estimator
 - Biased estimator: expected value of estimator does not equal integral it is designed to estimate (not good!)

General unbiased Monte Carlo estimator

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

$$X_i \sim p(x)$$

Special case where X_i drawn from uniform distribution:

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

$$X_i \sim U(a, b)$$

$$p(x) = \frac{1}{b-a}$$

Biased sample selection, but unbiased estimator

■ **Probability:** $X_i \sim p(x)$

■ **Estimator:** $Y_i = \frac{f(X_i)}{p(X_i)}$

$$\begin{aligned} E[Y_i] &= E \left[\frac{f(X_i)}{p(X_i)} \right] \\ &= \int \frac{f(x)}{p(x)} p(x) \, dx \\ &= \int f(x) \, dx \end{aligned}$$

Importance sampling

Idea: bias selection of samples towards parts of domain where function we are integrating is large (“most useful samples”)

Sample according to $f(x)$

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

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

**If PDF is proportional to f
then variance is 0!**

Recall definition of variance:

$$V[\tilde{f}] = E[\tilde{f}^2] - E^2[\tilde{f}]$$

$$\begin{aligned} E[\tilde{f}^2] &= \int \left[\frac{f(x)}{p(x)} \right]^2 p(x) dx \\ &= \int \left[\frac{f(x)}{f(x)/E[f]} \right]^2 \frac{f(x)}{E[f]} dx \\ &= E[f] \int f(x) dx \\ &= E^2[f] \end{aligned}$$

$$\rightarrow V[\tilde{f}] = 0 \text{ ?!?$$


Importance sampling example

Cosine-weighted hemisphere sampling in irradiance estimate:

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

$$\int_{\Omega} f(\omega) \, d\omega \approx \frac{1}{N} \sum_i^N \frac{f(\omega)}{p(\omega)} = \frac{1}{N} \sum_i^N \frac{L_i(\omega) \cos \theta}{\cos \theta / \pi} = \frac{\pi}{N} \sum_i^N L_i(\omega)$$

**Idea: bias samples toward directions where $\cos \theta$ is large
(if L is constant, then these are the directions that contribute most)**



Summary: Monte Carlo integration

■ Monte Carlo estimator

- Estimate integral by evaluating function at random sample points in domain

$$F_N = \frac{1}{N} \sum_{i=1}^N \frac{f(X_i)}{p(X_i)} \approx \int_a^b f(x) dx$$

■ Useful in rendering due to estimate high dimension integrals

- Faster convergence in estimating high dimensional integrals than non-randomized quadrature methods
- Suffers from noise due to variance in estimate

■ Importance sampling

- Reduce variance by biasing choice of samples to regions of domain where value of function is large
- Intuition: pick samples that will “contribute most” to estimate

What you should know:

- How do we use the trapezoidal rule to integrate a function?
- How does work increase with dimensionality of our function?
 - This is why we typically use Monte Carlo integration in graphics!
- Give a high level overview of the process of Monte Carlo integration
- What is a probability density function (PDF)?
- What is a Cumulative Distribution Function (CDF)?
- The Inversion Method can be used to correctly draw a sample from a PDF.
 - Sketch the overall step by step process for using the Inversion Method.
 - Work through how to use it to sample area of a circle
 - Work through how to use it to sample solid angles from a hemisphere
- What is rejection sampling? Show how to use rejection sampling to sample area of a circle, volume of a sphere, directions on a sphere, and solid angles from a hemisphere.
- Use one of the sampling methods we discussed to correctly accumulate incident irradiance on a surface patch.