Lecture 12:

Monte Carlo Integration

Computer Graphics and Imaging
UC Berkeley CS184/284A

Reminder: Quadrature-Based Numerical Integration

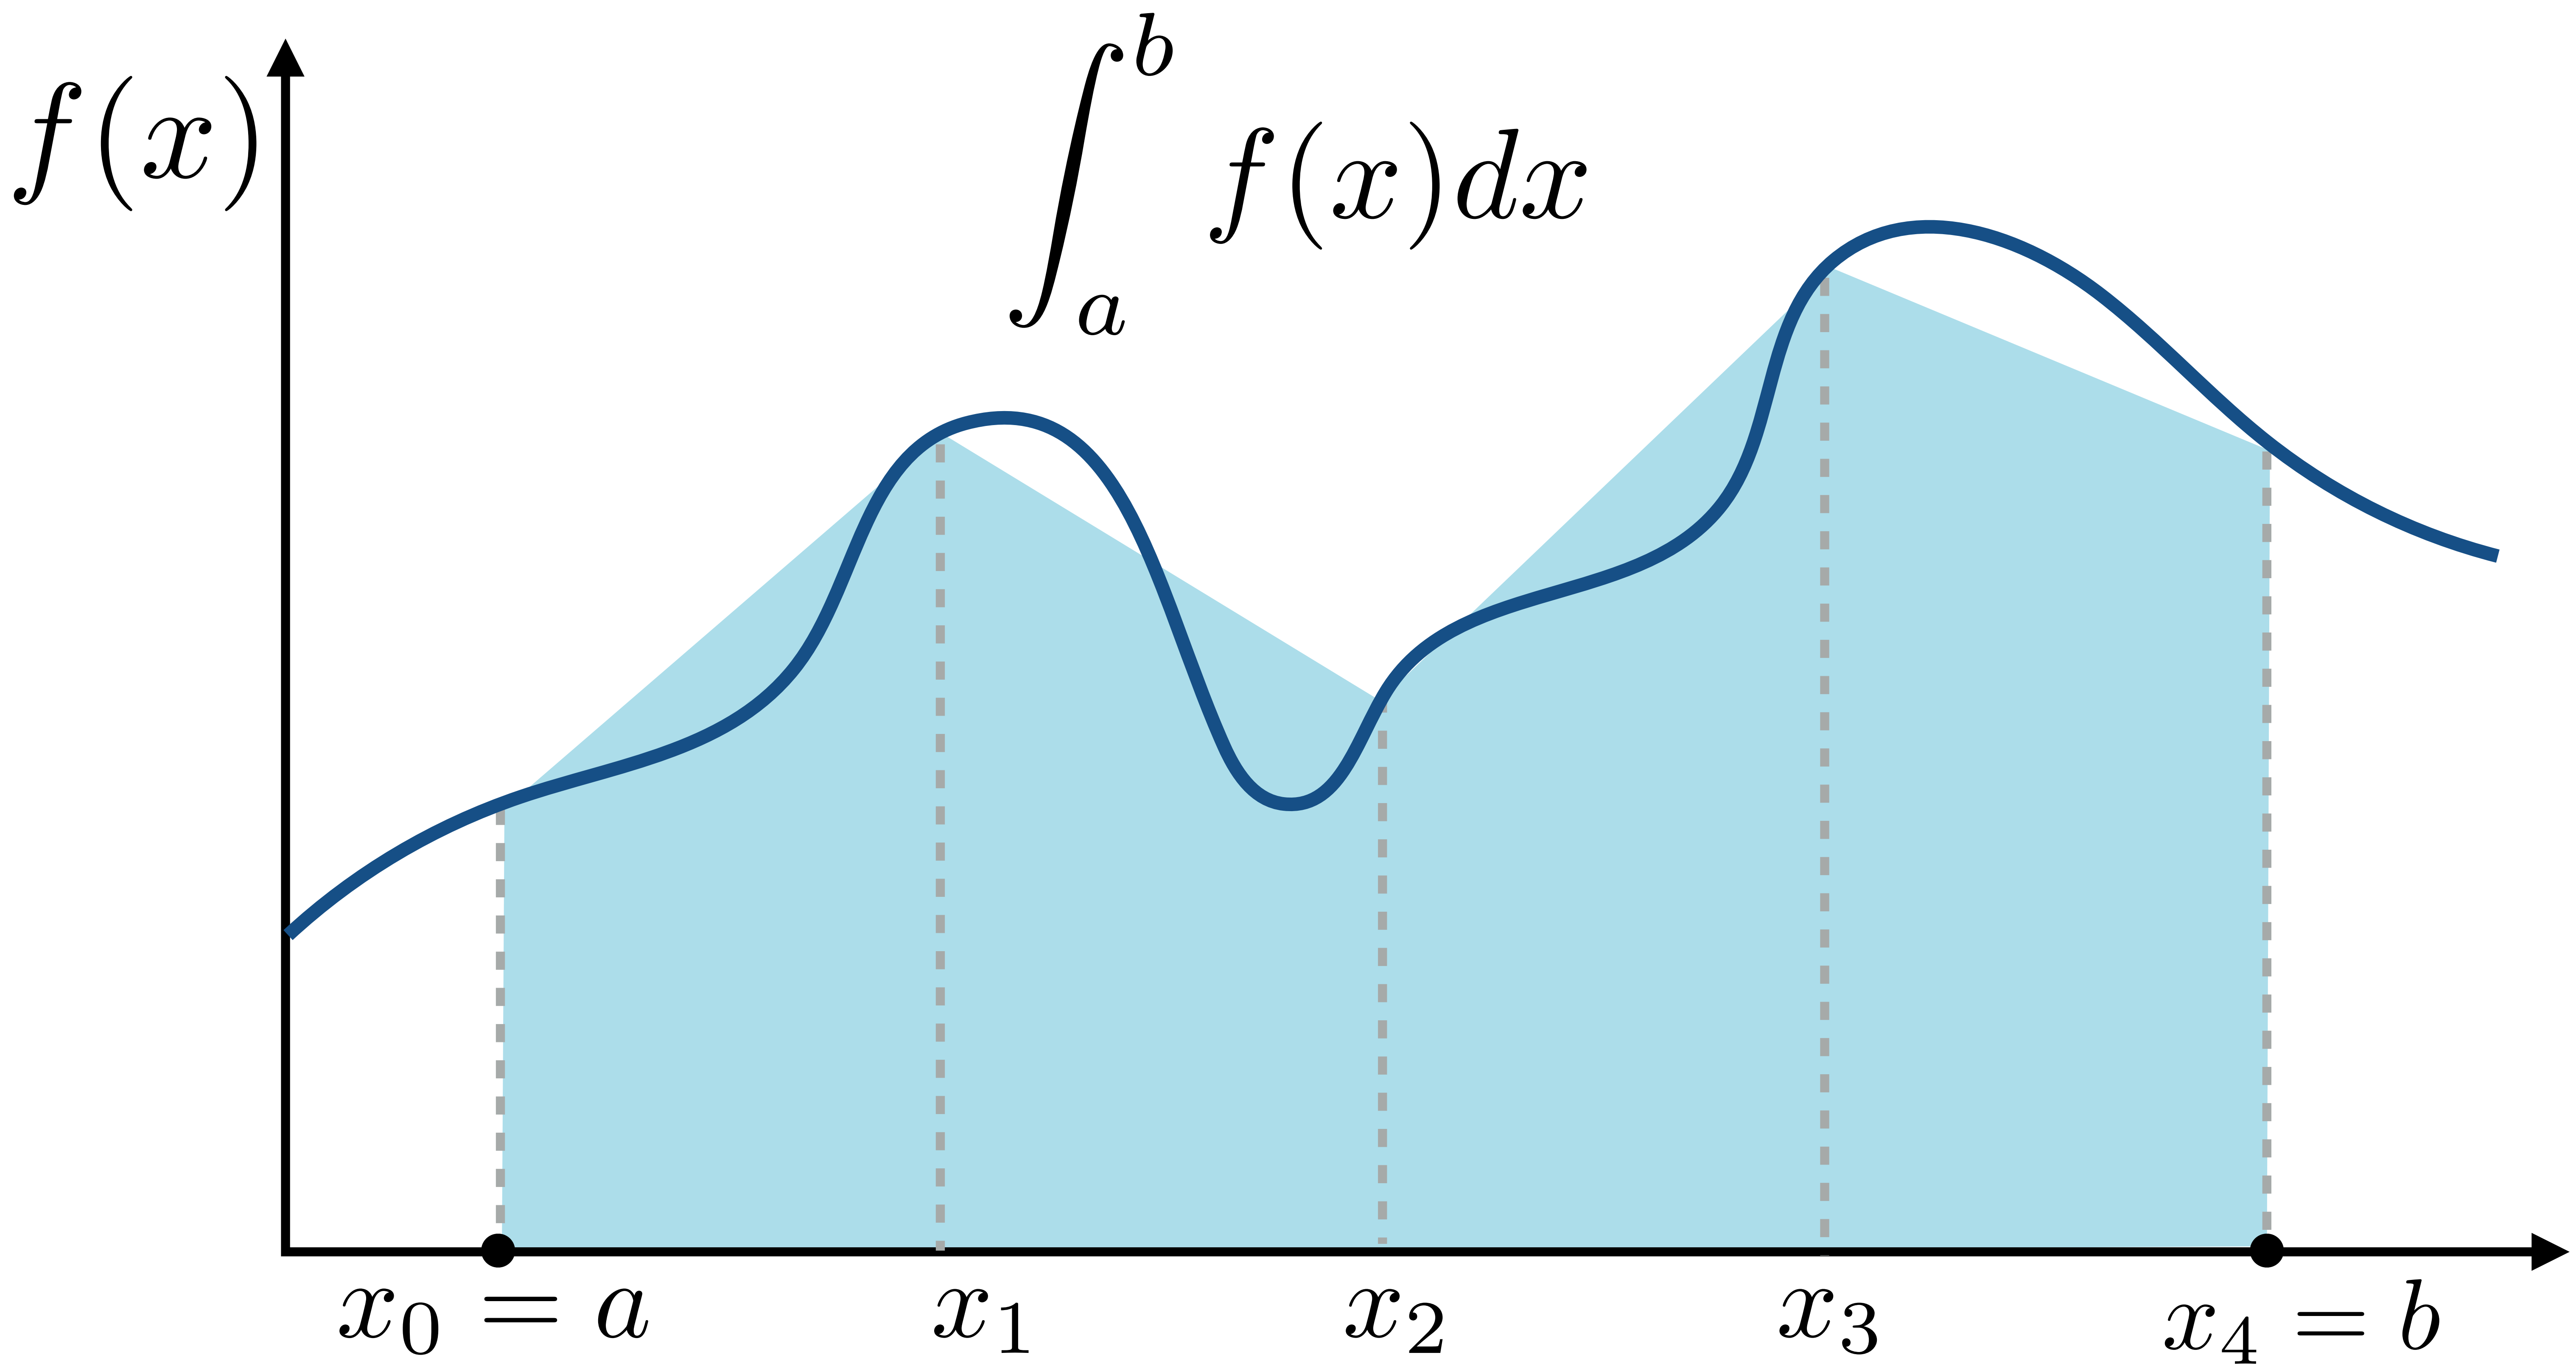

E.g. trapezoidal rule - estimate integral assuming function is piecewise linear

Multi-Dimensional Integrals
(Rendering Examples)

2D Integral: Recall Antialiasing By Area Sampling

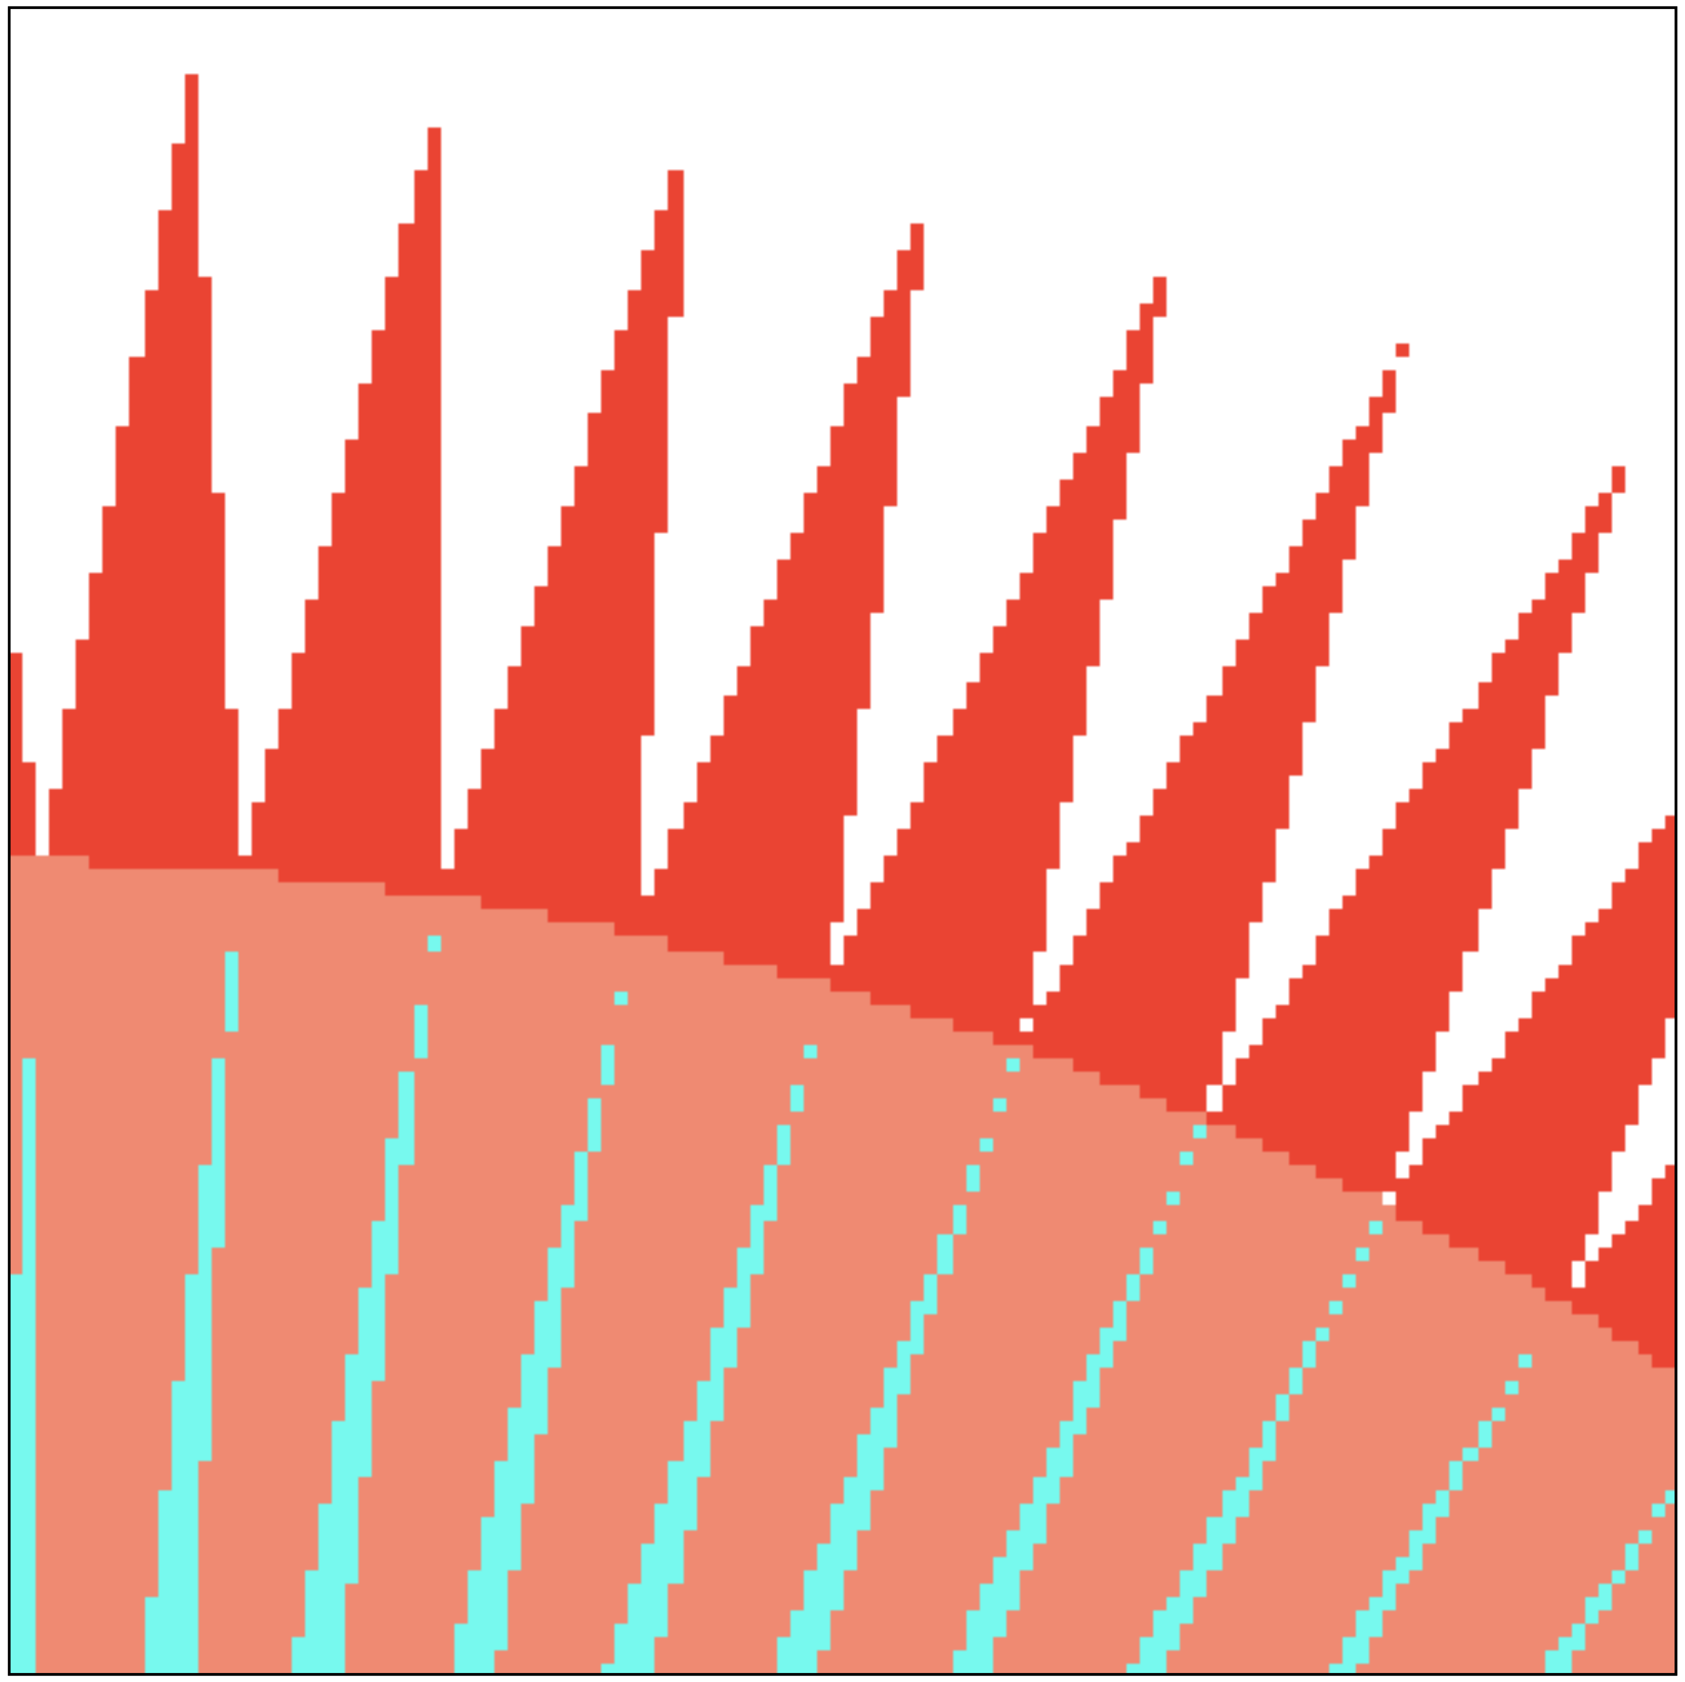

Point sampling

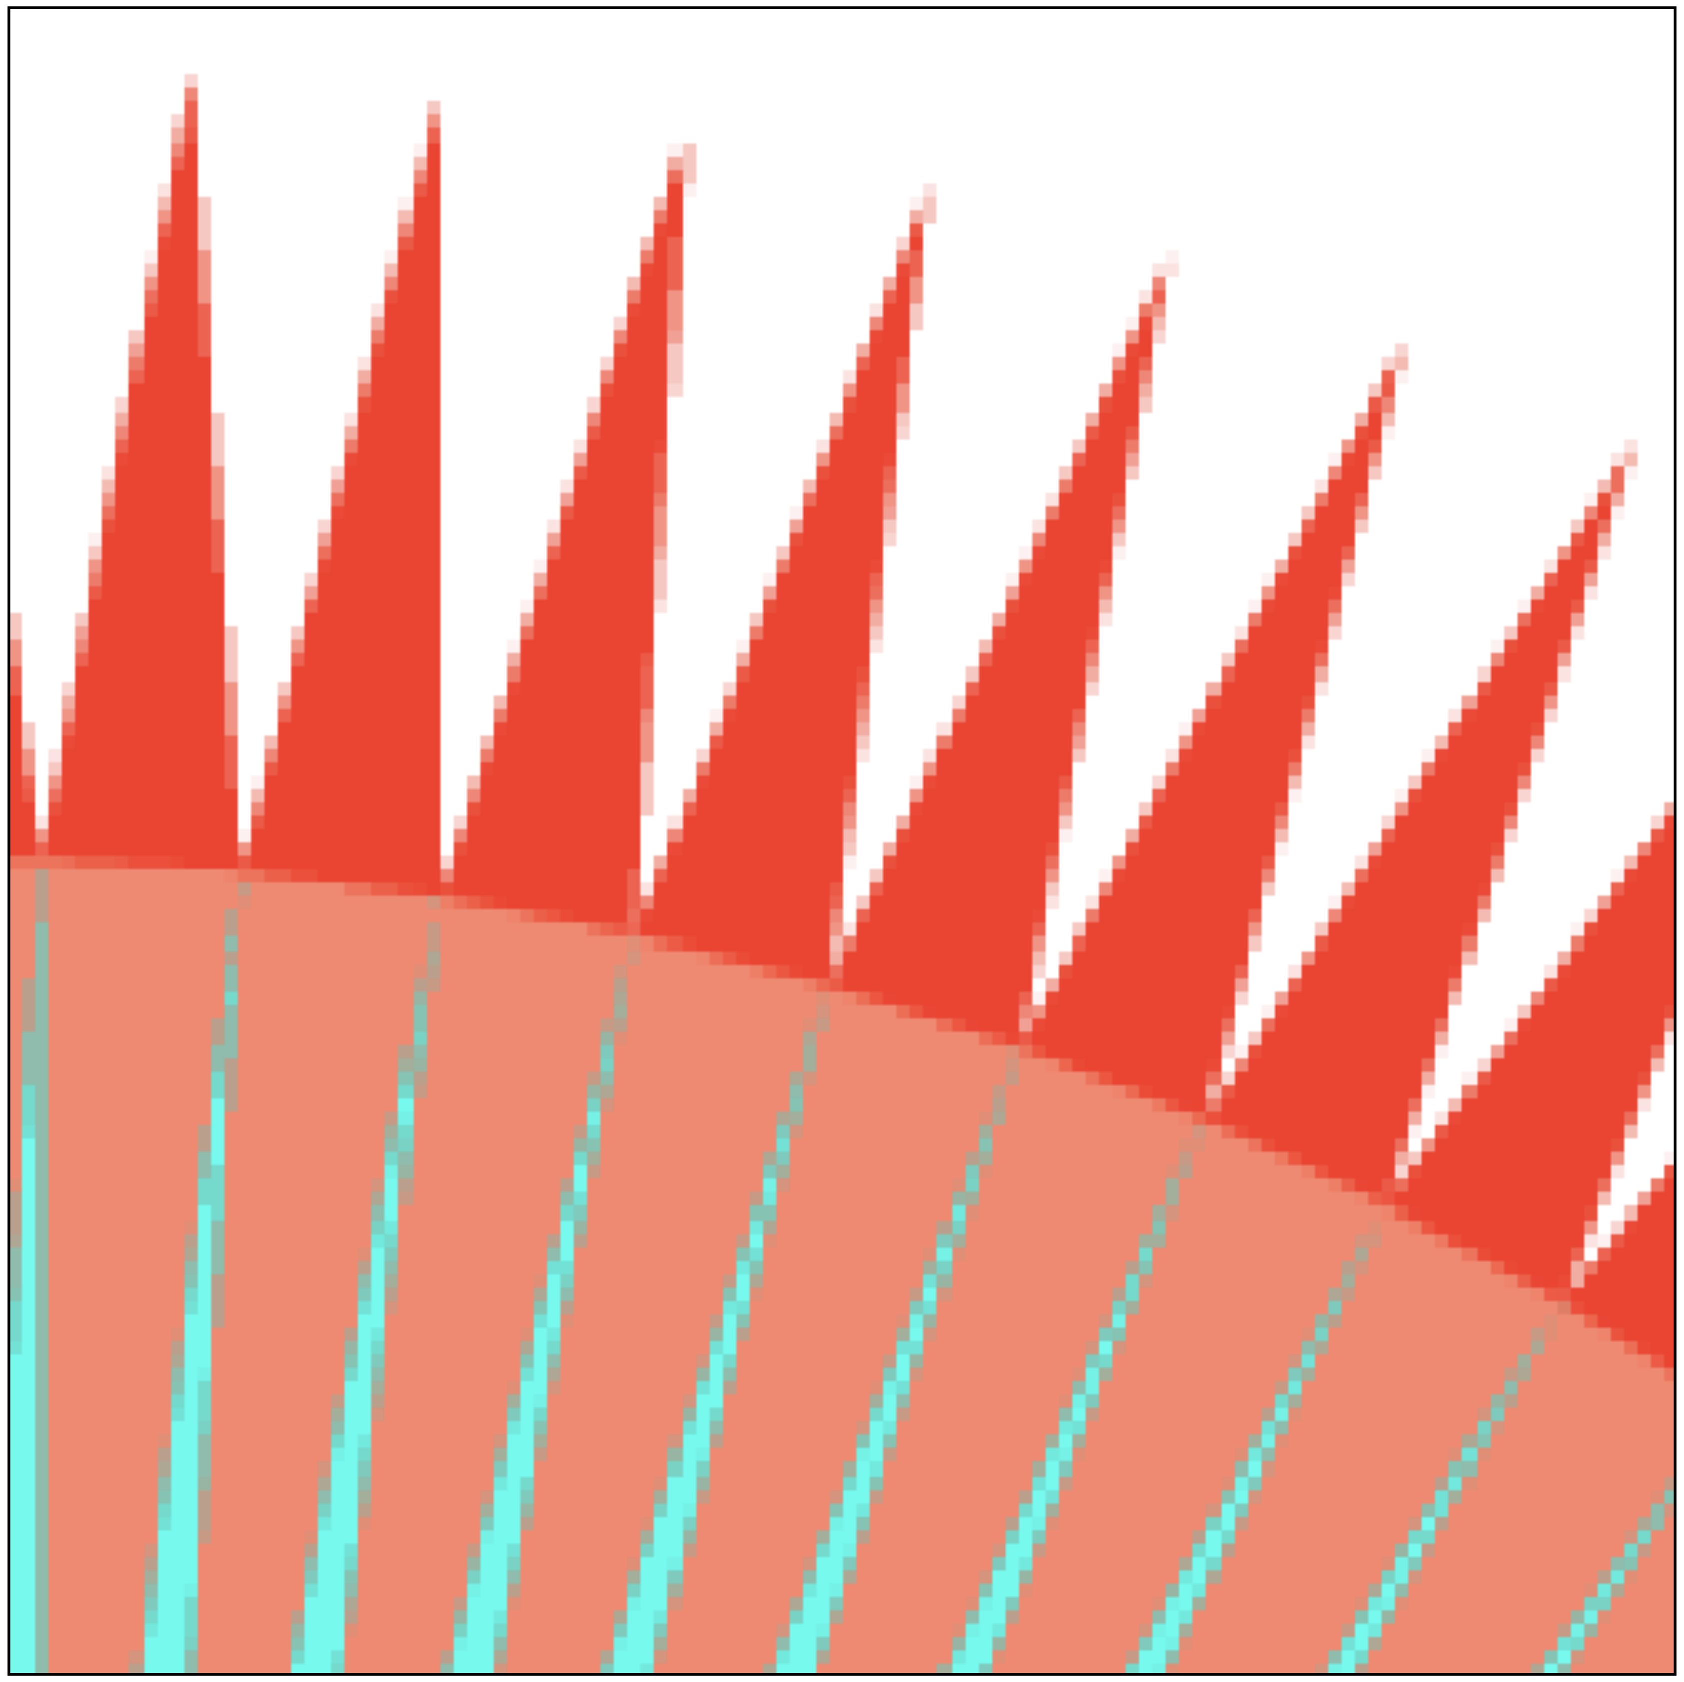

Area sampling

Integrate over 2D area of pixel

2D Integral: Irradiance from the Environment

Computing flux per unit area on surface, due to incoming light from all directions.

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

← **Contribution to irradiance from light arriving from direction ω**

Light meter

Hemisphere: H^2

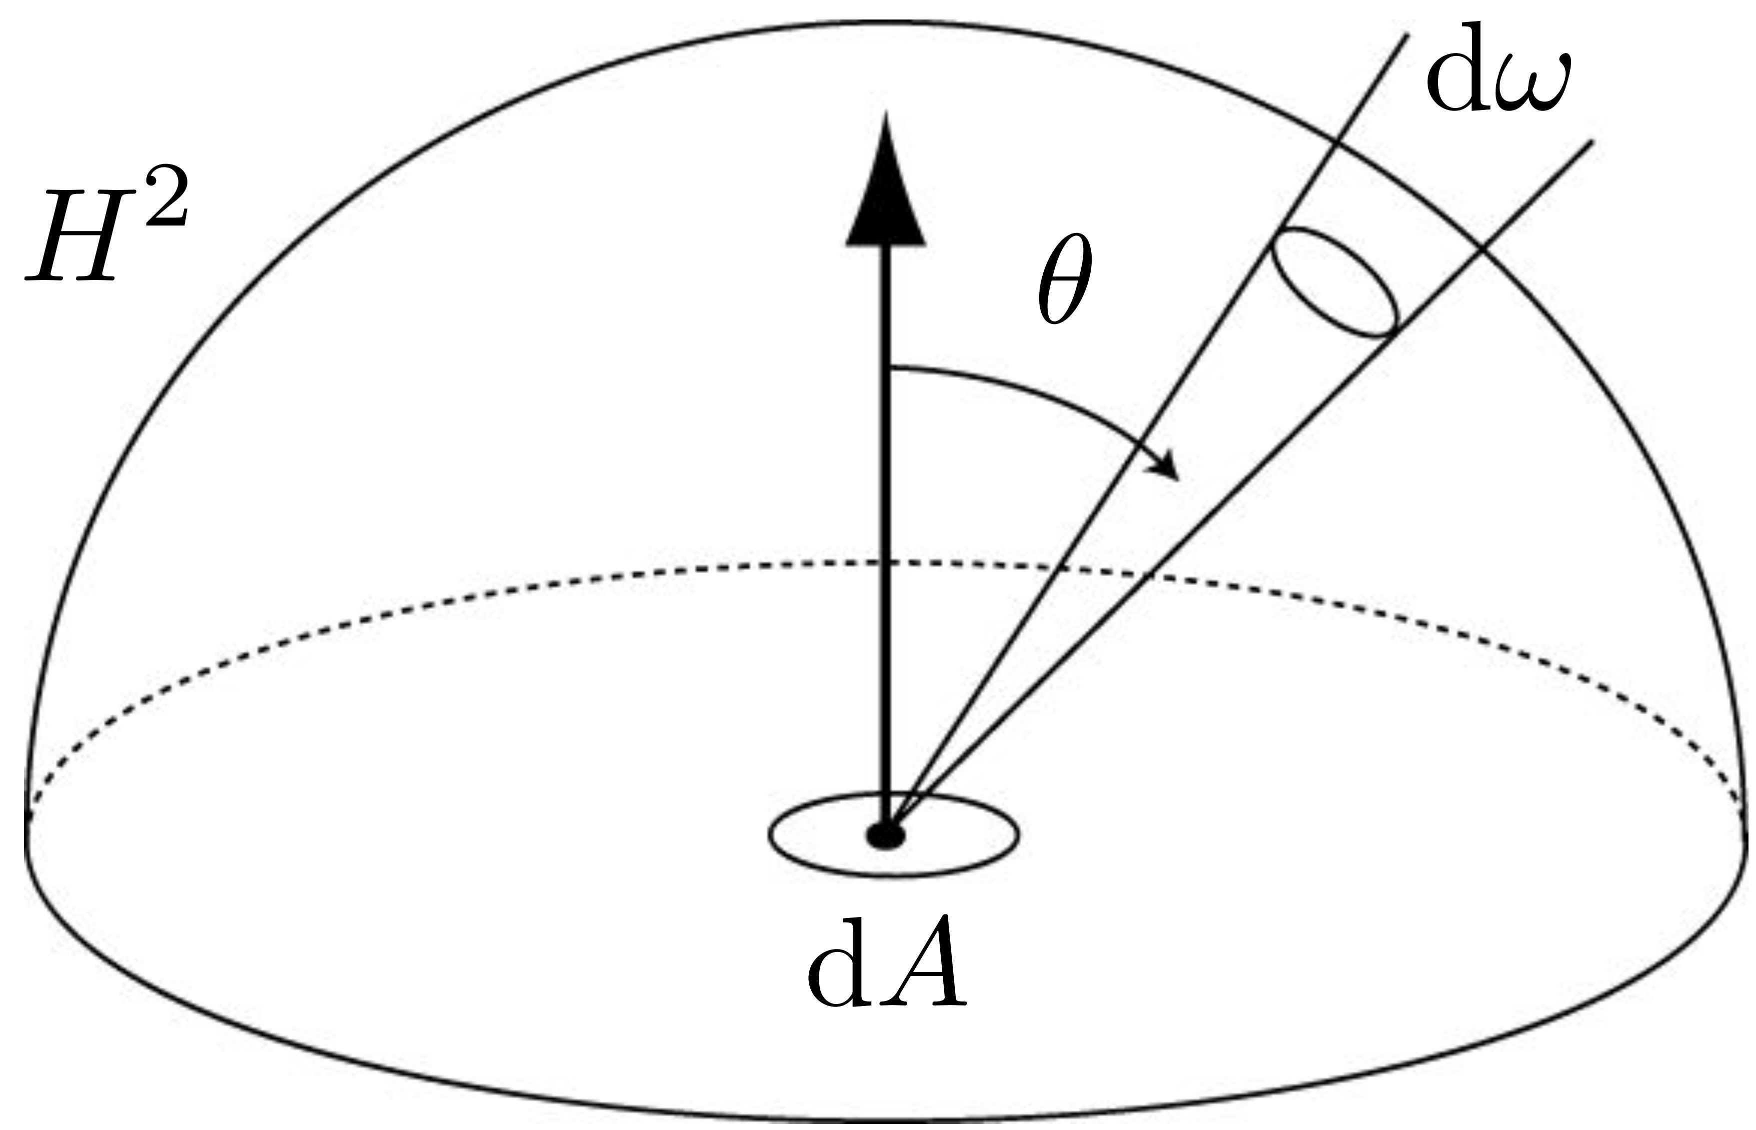

3D Integral: Motion Blur

Integrate over area of pixel and over exposure time.

Cook et al. “1984”

5D Integral: Real Camera Pixel Exposure

Credit: lycheng99, http://flic.kr/p/x4DEZh

Integrate over 2D lens pupil, 2D pixel, and over exposure time

5D Integral: Real Camera Pixel Exposure

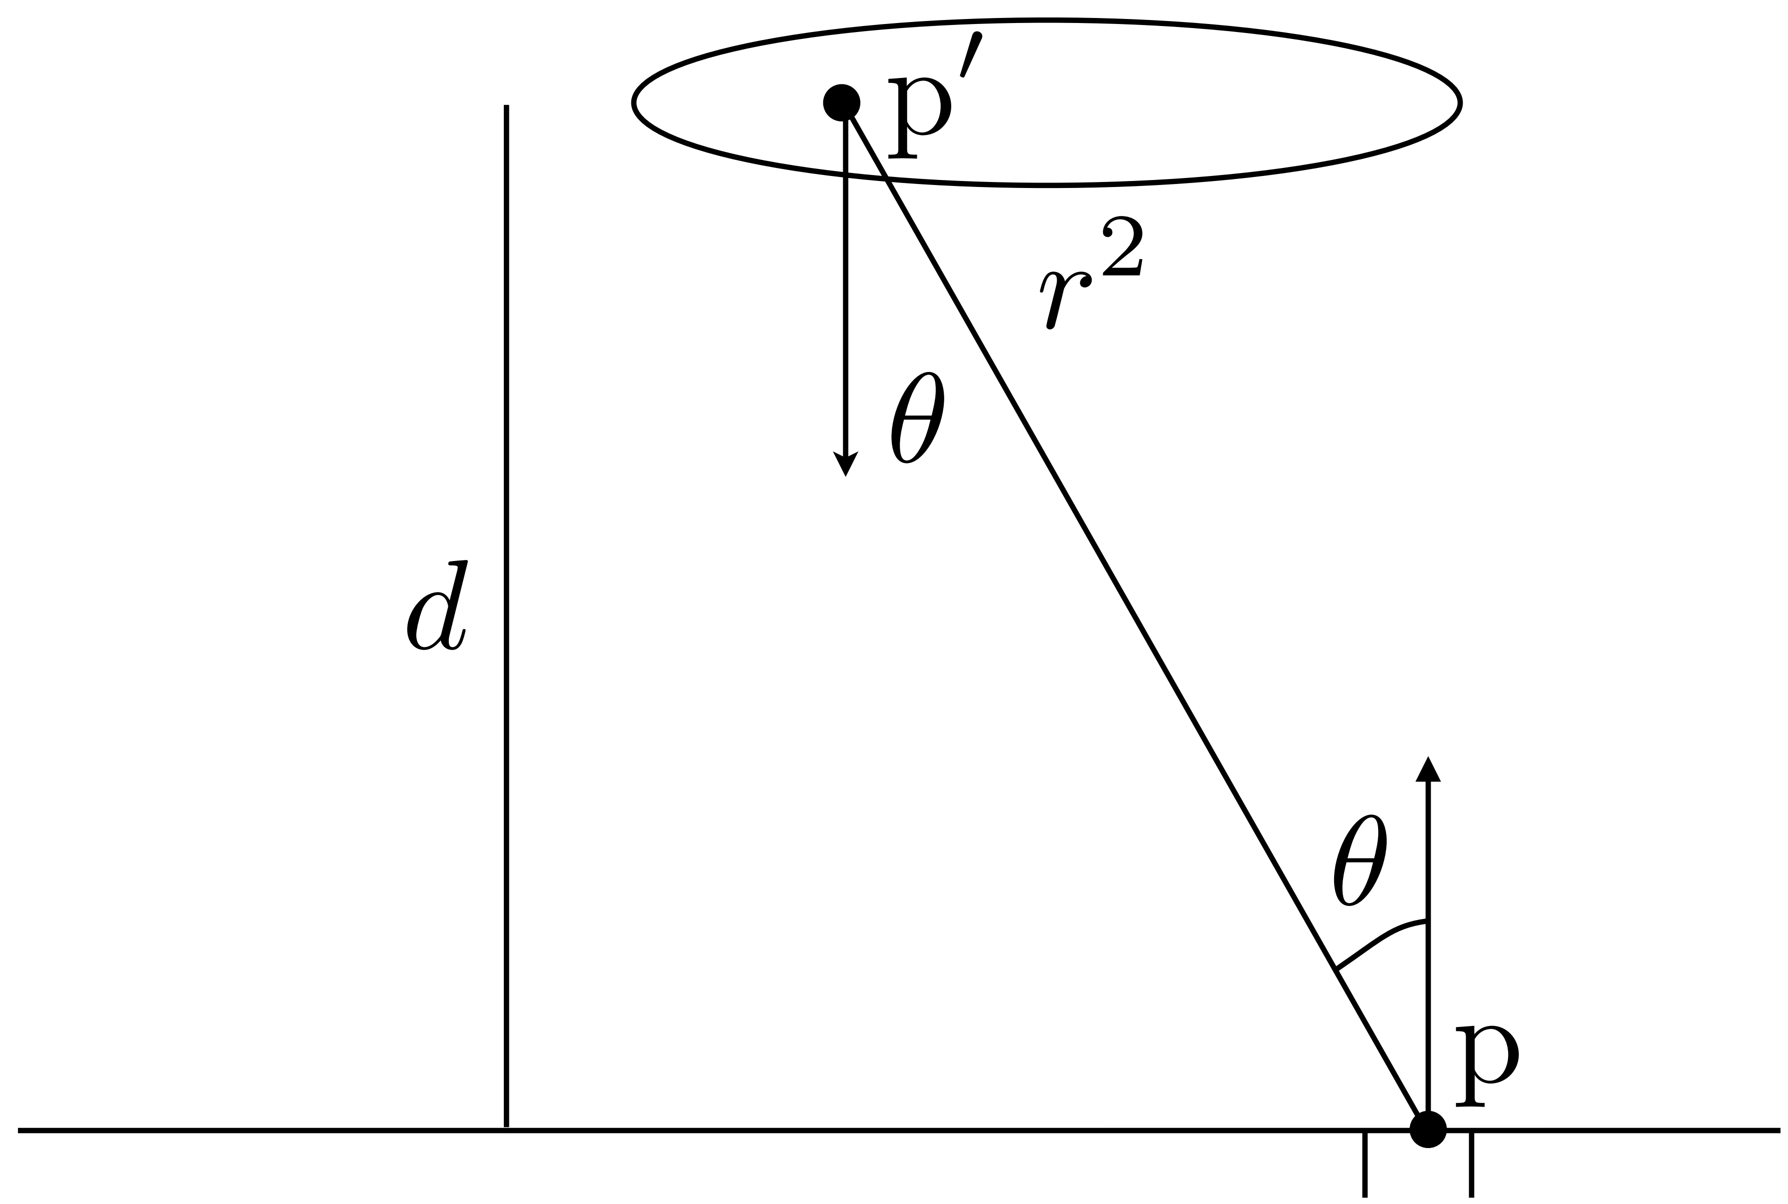

$$Q_{\text{pixel}} = \frac{1}{d^2} \int_{t_0}^{t^1} \int_{A_{\text{lens}}} \int_{A_{\text{pixel}}} L(\mathrm{p}' \rightarrow \mathrm{p}, t) \cos^4 \theta \, \mathrm{d}p \, \mathrm{d}p' \, \mathrm{d}t$$

The Curse of Dimensionality

High-Dimensional Integration

Complete set of samples:

$$N = \underbrace{n \times n \times \cdots \times n}_{d} = n^d$$

- **"Curse of dimensionality"**

Numerical integration error:

$$\text{Error} \sim \frac{1}{n} = \frac{1}{N^{1/d}}$$

Random sampling error:

$$\text{Error} = \text{Variance}^{1/2} \sim \frac{1}{\sqrt{N}}$$

In high dimensions, Monte Carlo integration requires fewer samples than quadrature-based numerical integration

Global illumination = infinite-dimensional integrals

Example: Discrete vs Monte Carlo - Shadows

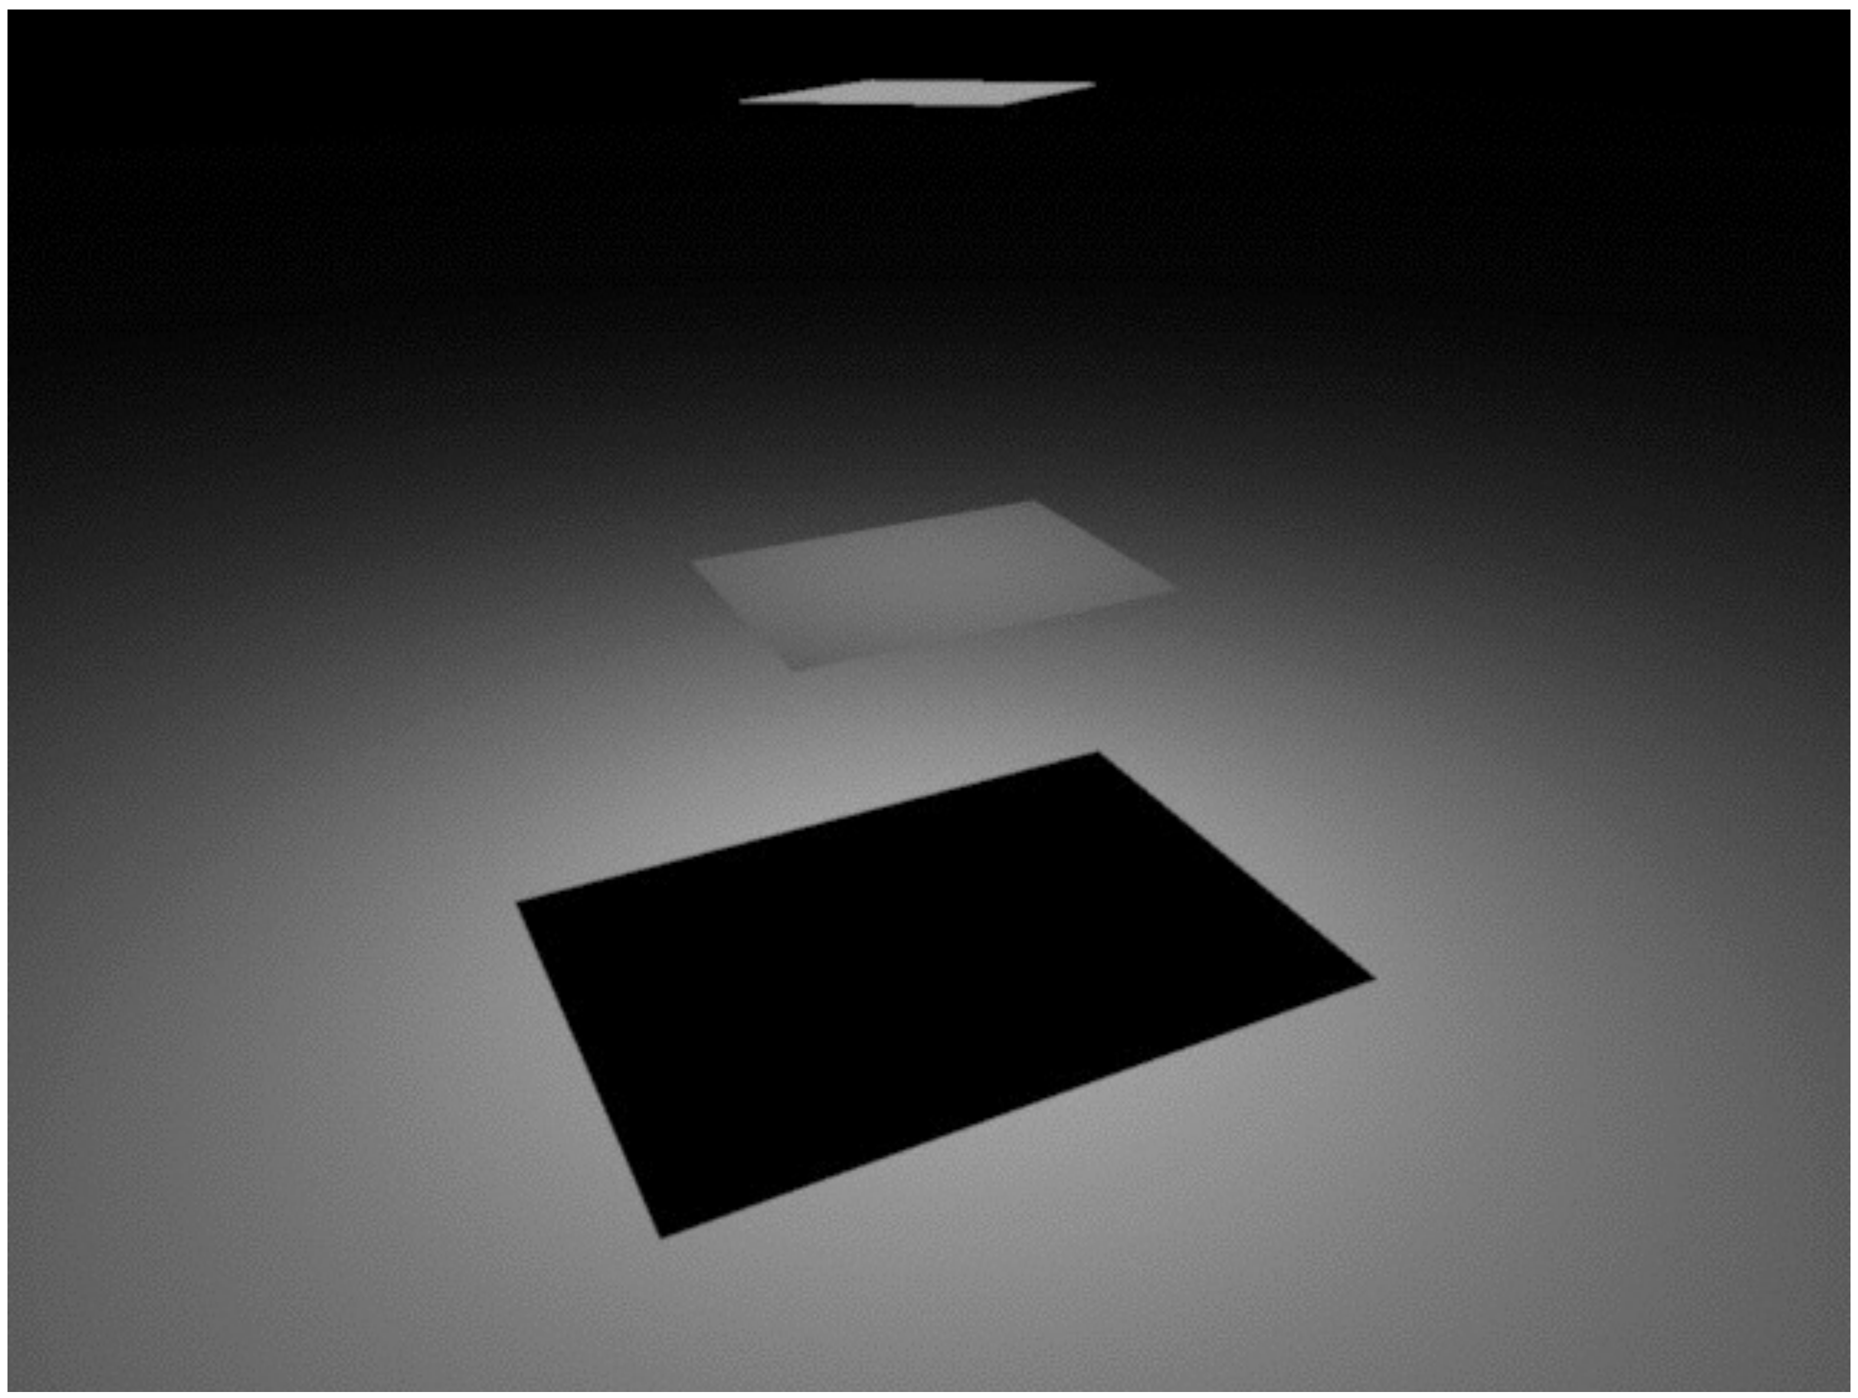

1 sample per pixel
Sample center of light

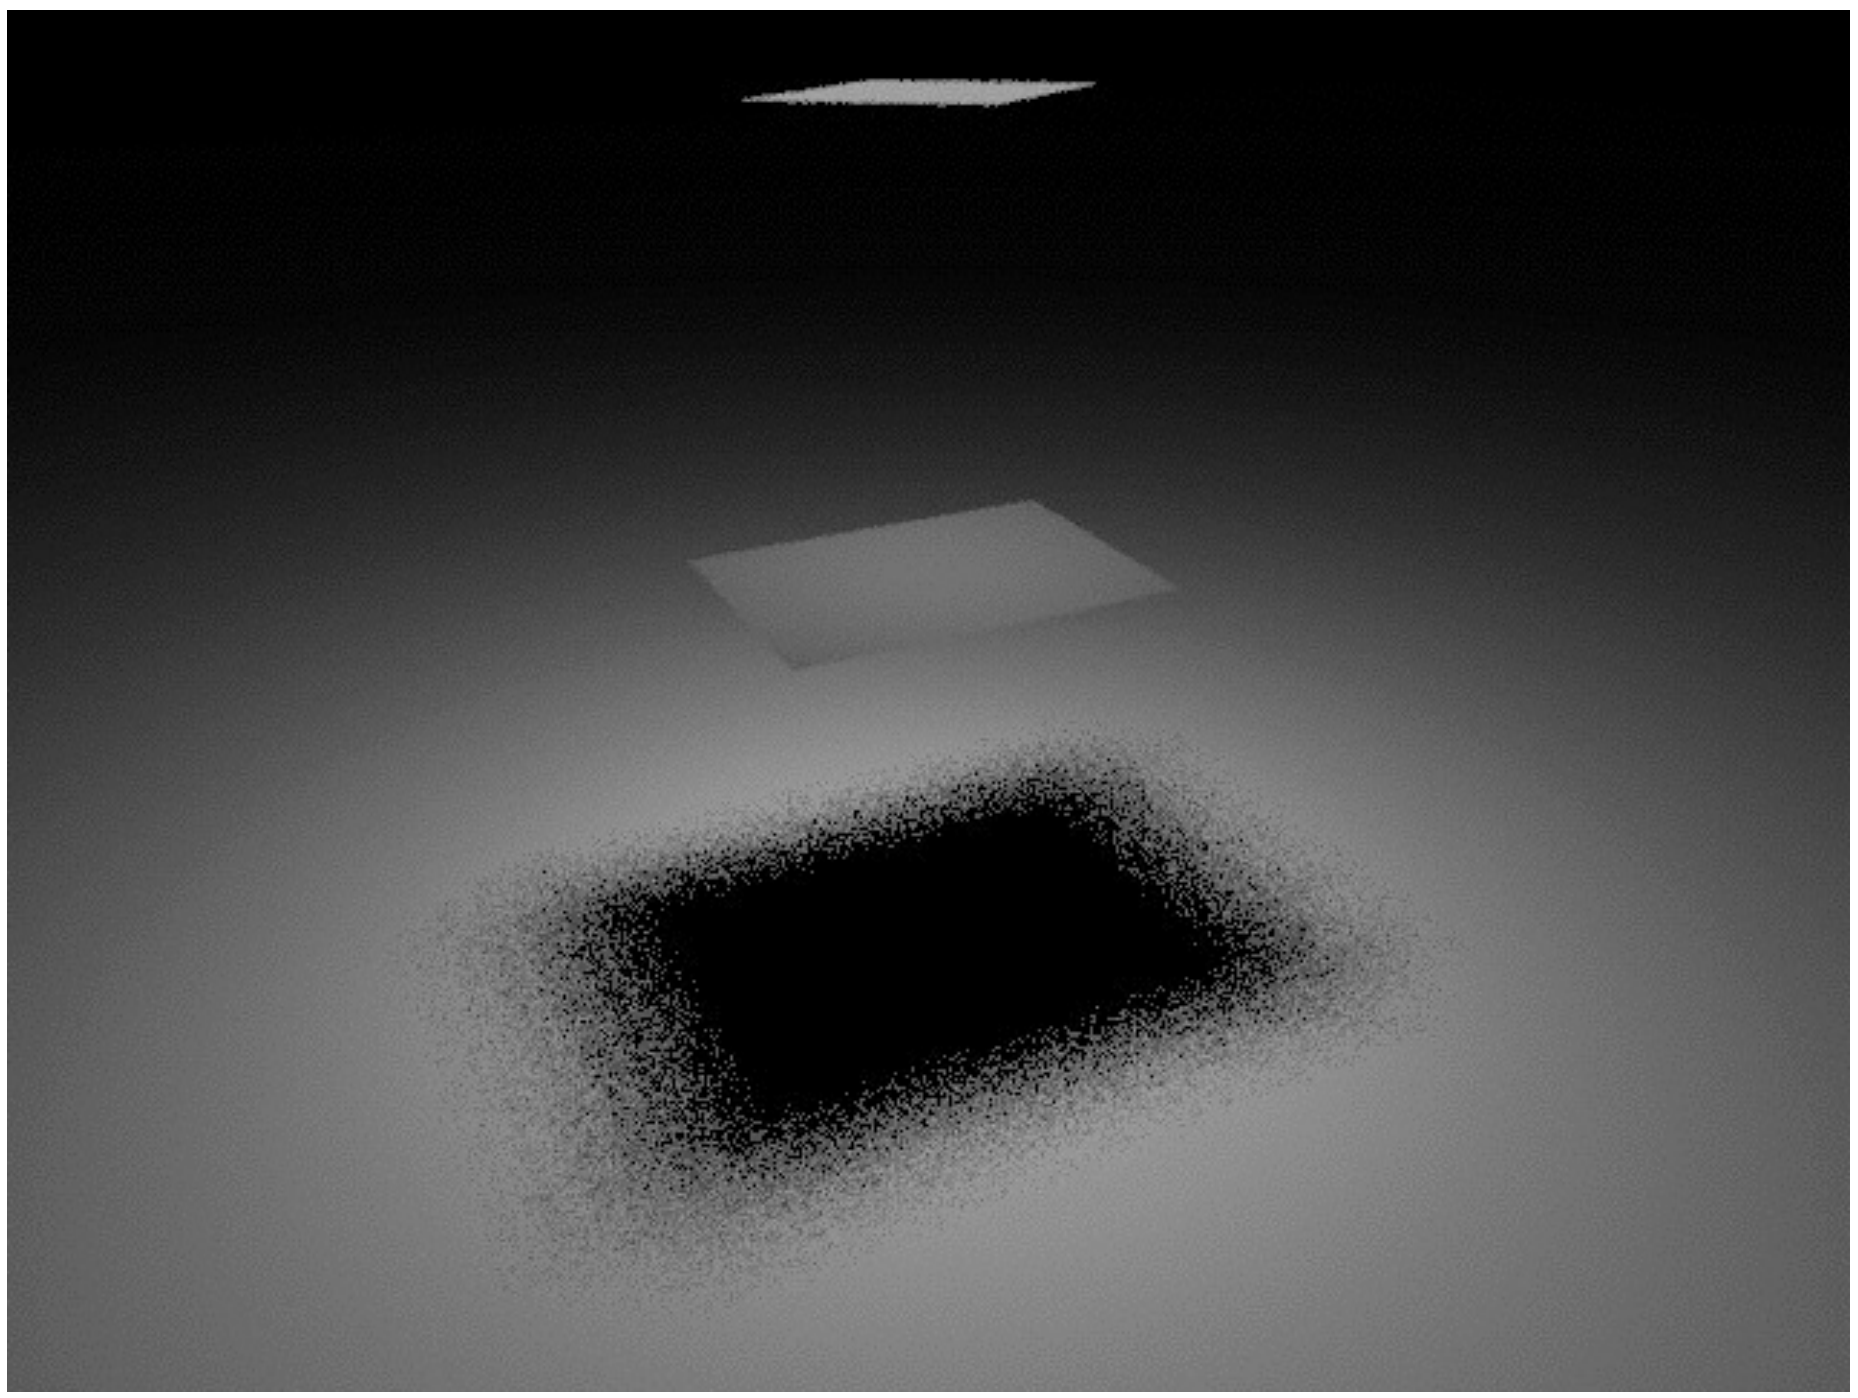

1 sample per pixel
Sample random point on light

Example: Discrete vs Monte Carlo - Shadows

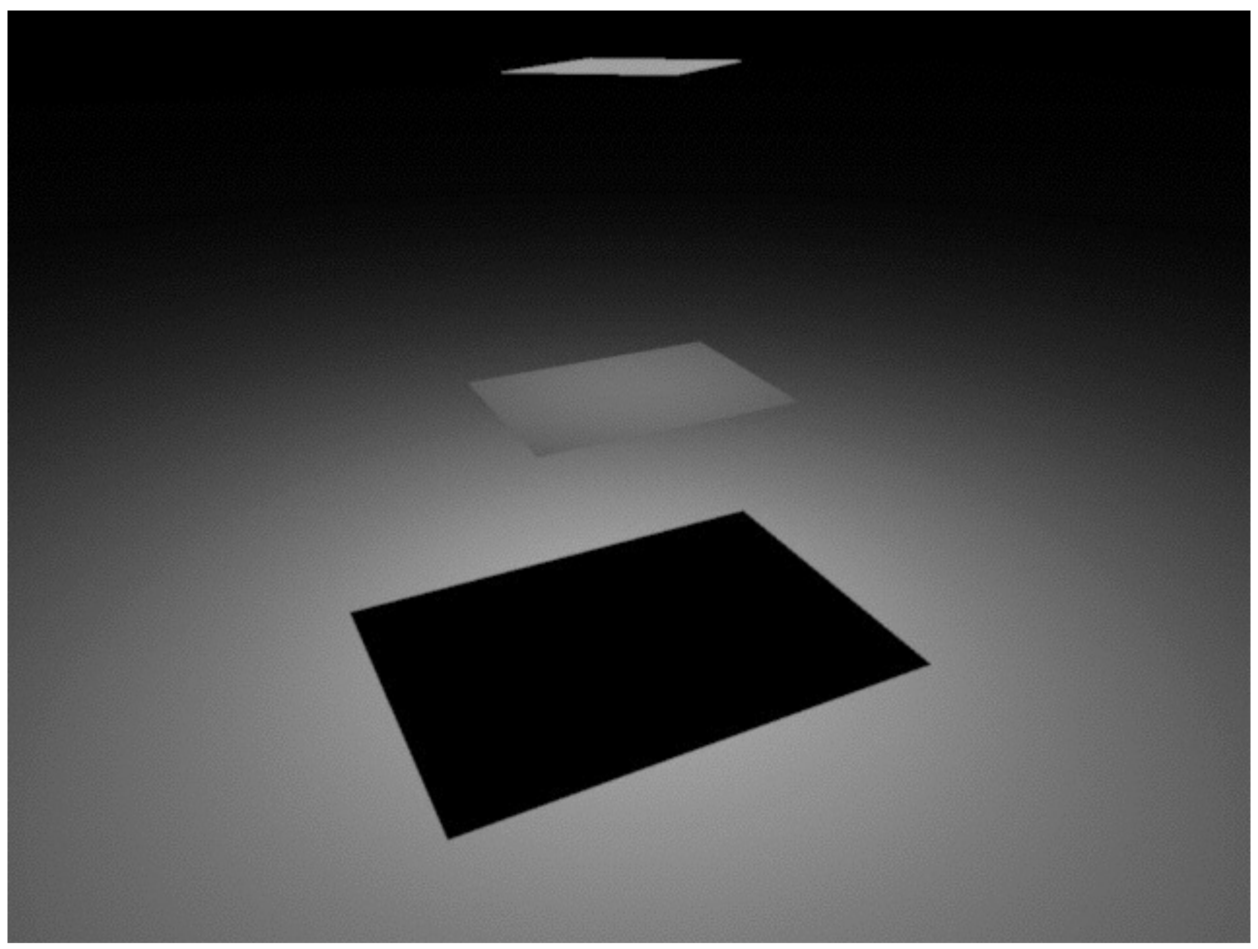

Sample center of light

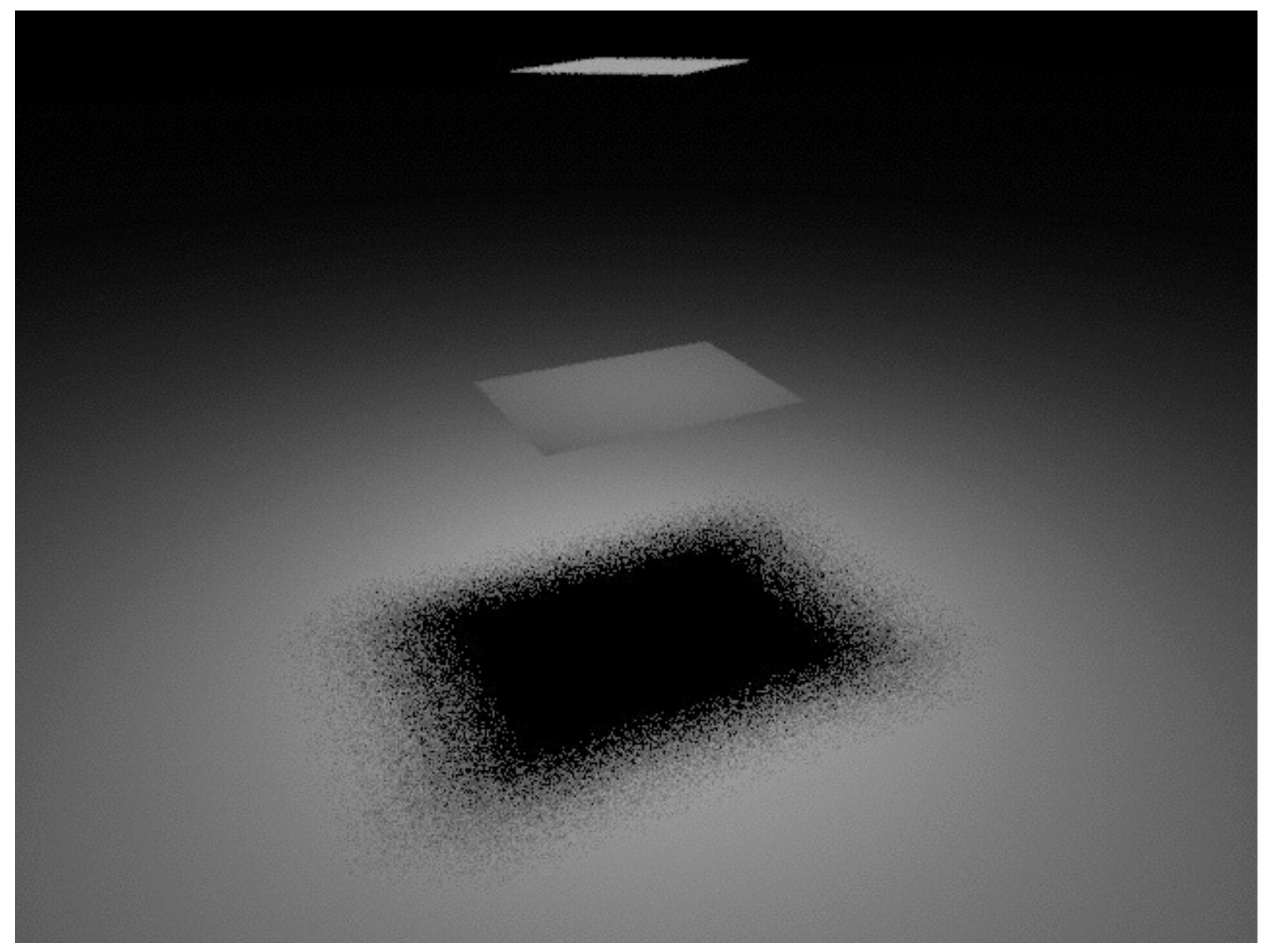

Sample random point on light

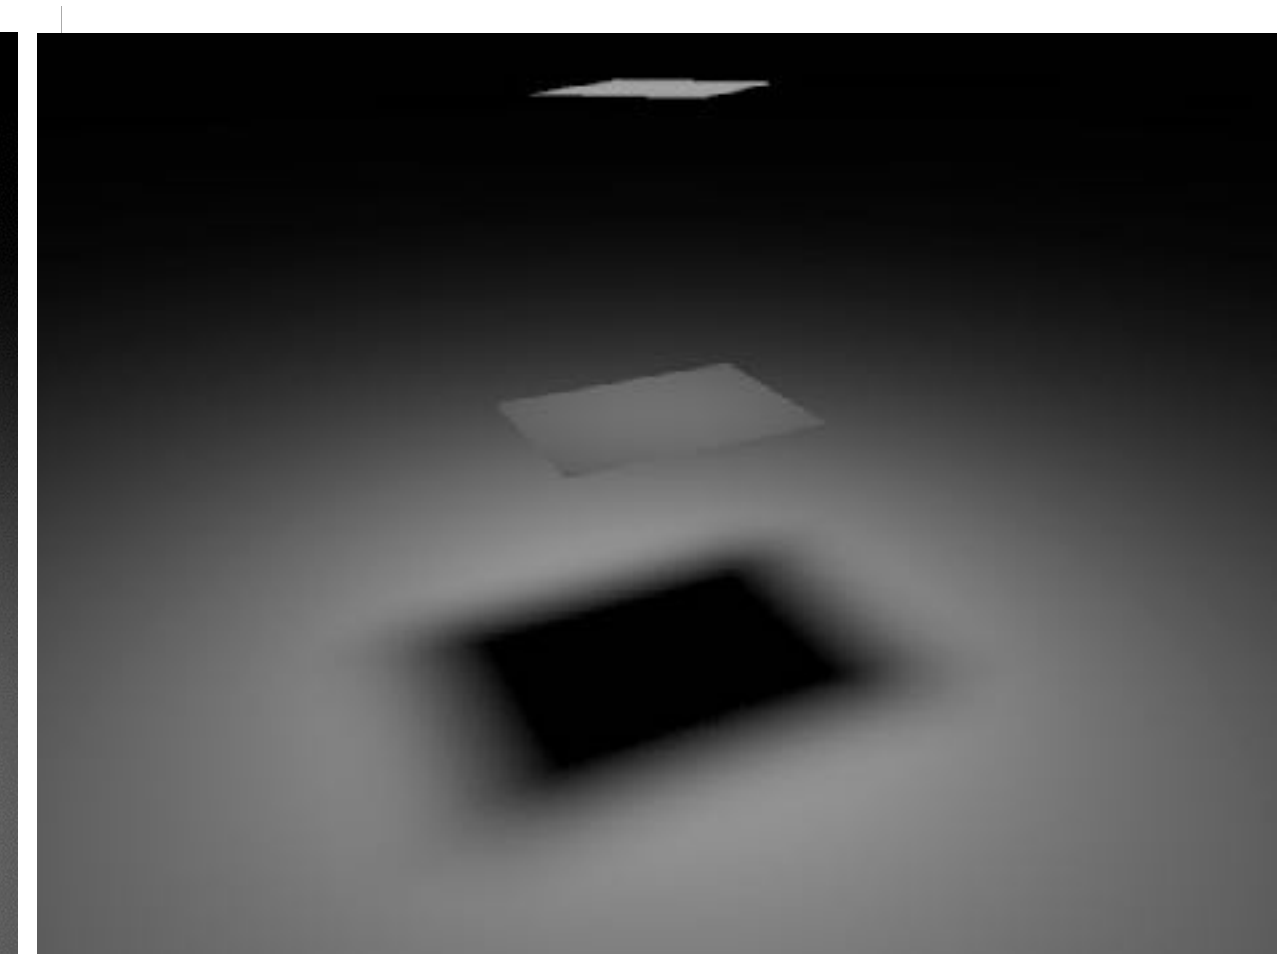

True answer

Overview: Monte Carlo Integration

Idea: estimate integral based on random sampling of function

Advantages:

- General and relatively simple method
- Requires only function evaluation at any point
- Works for very general functions, including discontinuities
- Efficient for high-dimensional integrals — avoids "curse of dimensionality"

Disadvantages:

- Noise. Integral estimate is random, only correct "on average"
- Can be slow to converge — need a lot of samples

Probability 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

Example: uniform PDF: all values over a domain are equally likely

e.g. A six-sided die

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

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

Probability Distribution Function (PDF)

n **discrete values** x_i

With probability p_i

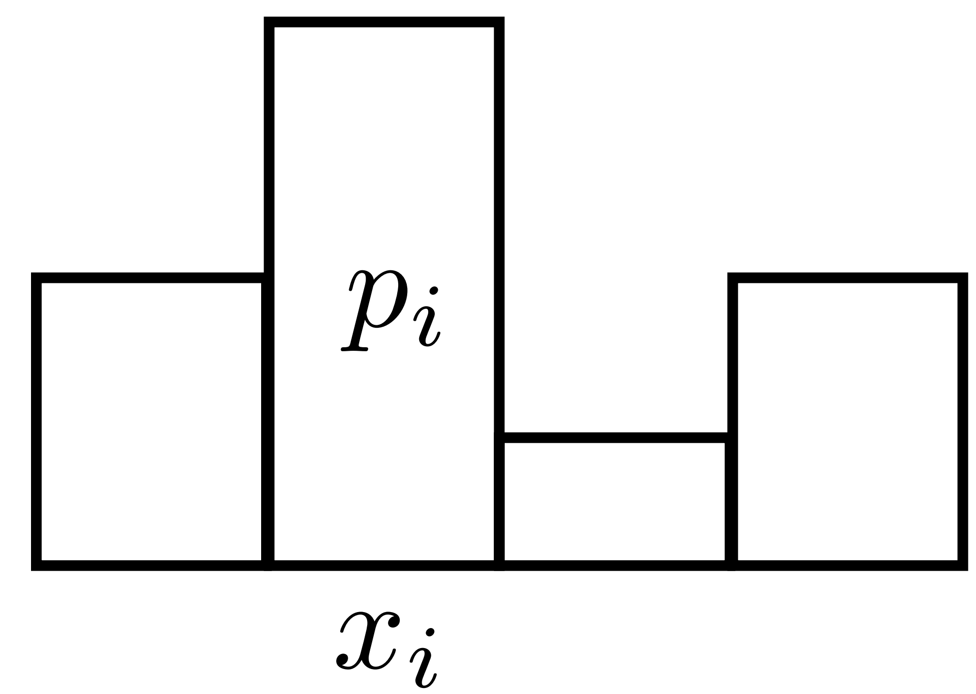

Requirements of a probability distribution:

$$p_i \geq 0$$

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

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

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

Expected Value of a Random Variable

The average value that one obtains if repeatedly drawing samples from the random distribution.

X **drawn from distribution with**

n **discrete values** x_i

with probabilities p_i

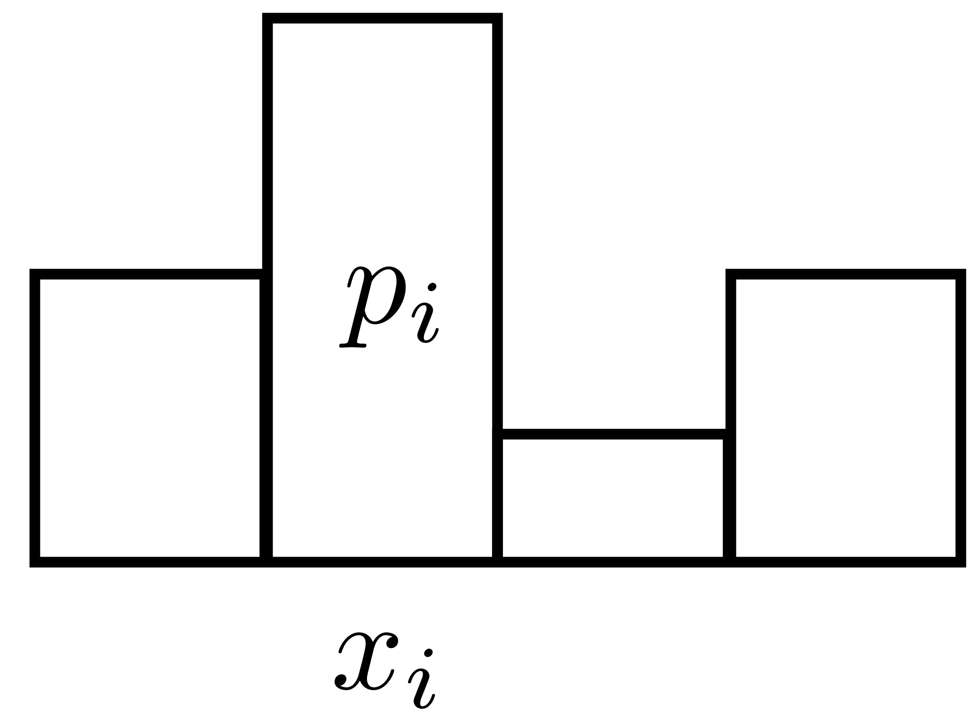

Expected value of X: $E[X] = \sum_{i=1}^{n} x_i p_i$

Die example: $E[X] = \sum_{i=1}^{n} \frac{i}{6}$

$= (1 + 2 + 3 + 4 + 5 + 6)/6 = 3.5$

Continuous Probability Distribution Function

$X \sim p(x)$

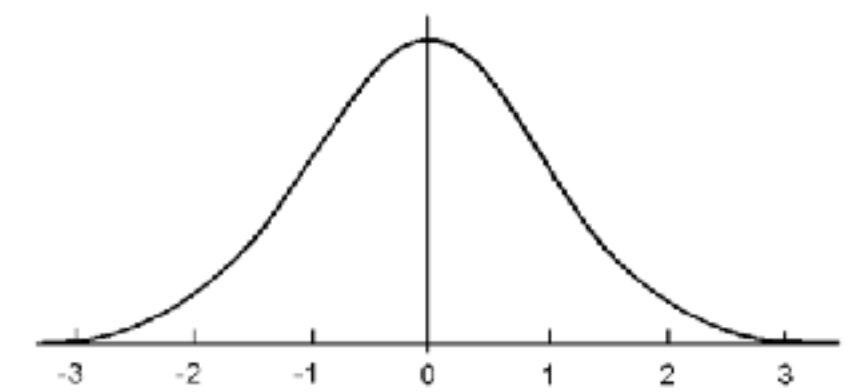

A random variable X that can take any of a continuous set of values, where the relative probability of a particular value is given by a continuous probability density function p(x).

Conditions on p(x): $p(x) \geq 0 \text{ and } \int p(x)\,dx = 1$

Expected value of X: $E[X] = \int x\,p(x)\,dx$

Function of a Random Variable

A function Y of a random variable X is also a random variable:

$$X \sim p(x)$$
$$Y = f(X)$$

Expected value of a function of a random variable:

$$E[Y] = E[f(X)] = \int f(x)\, p(x)\, dx$$

Monte Carlo Integration

Monte Carlo Integration

Simple idea: estimate the integral of a function by averaging random samples of the function's value.

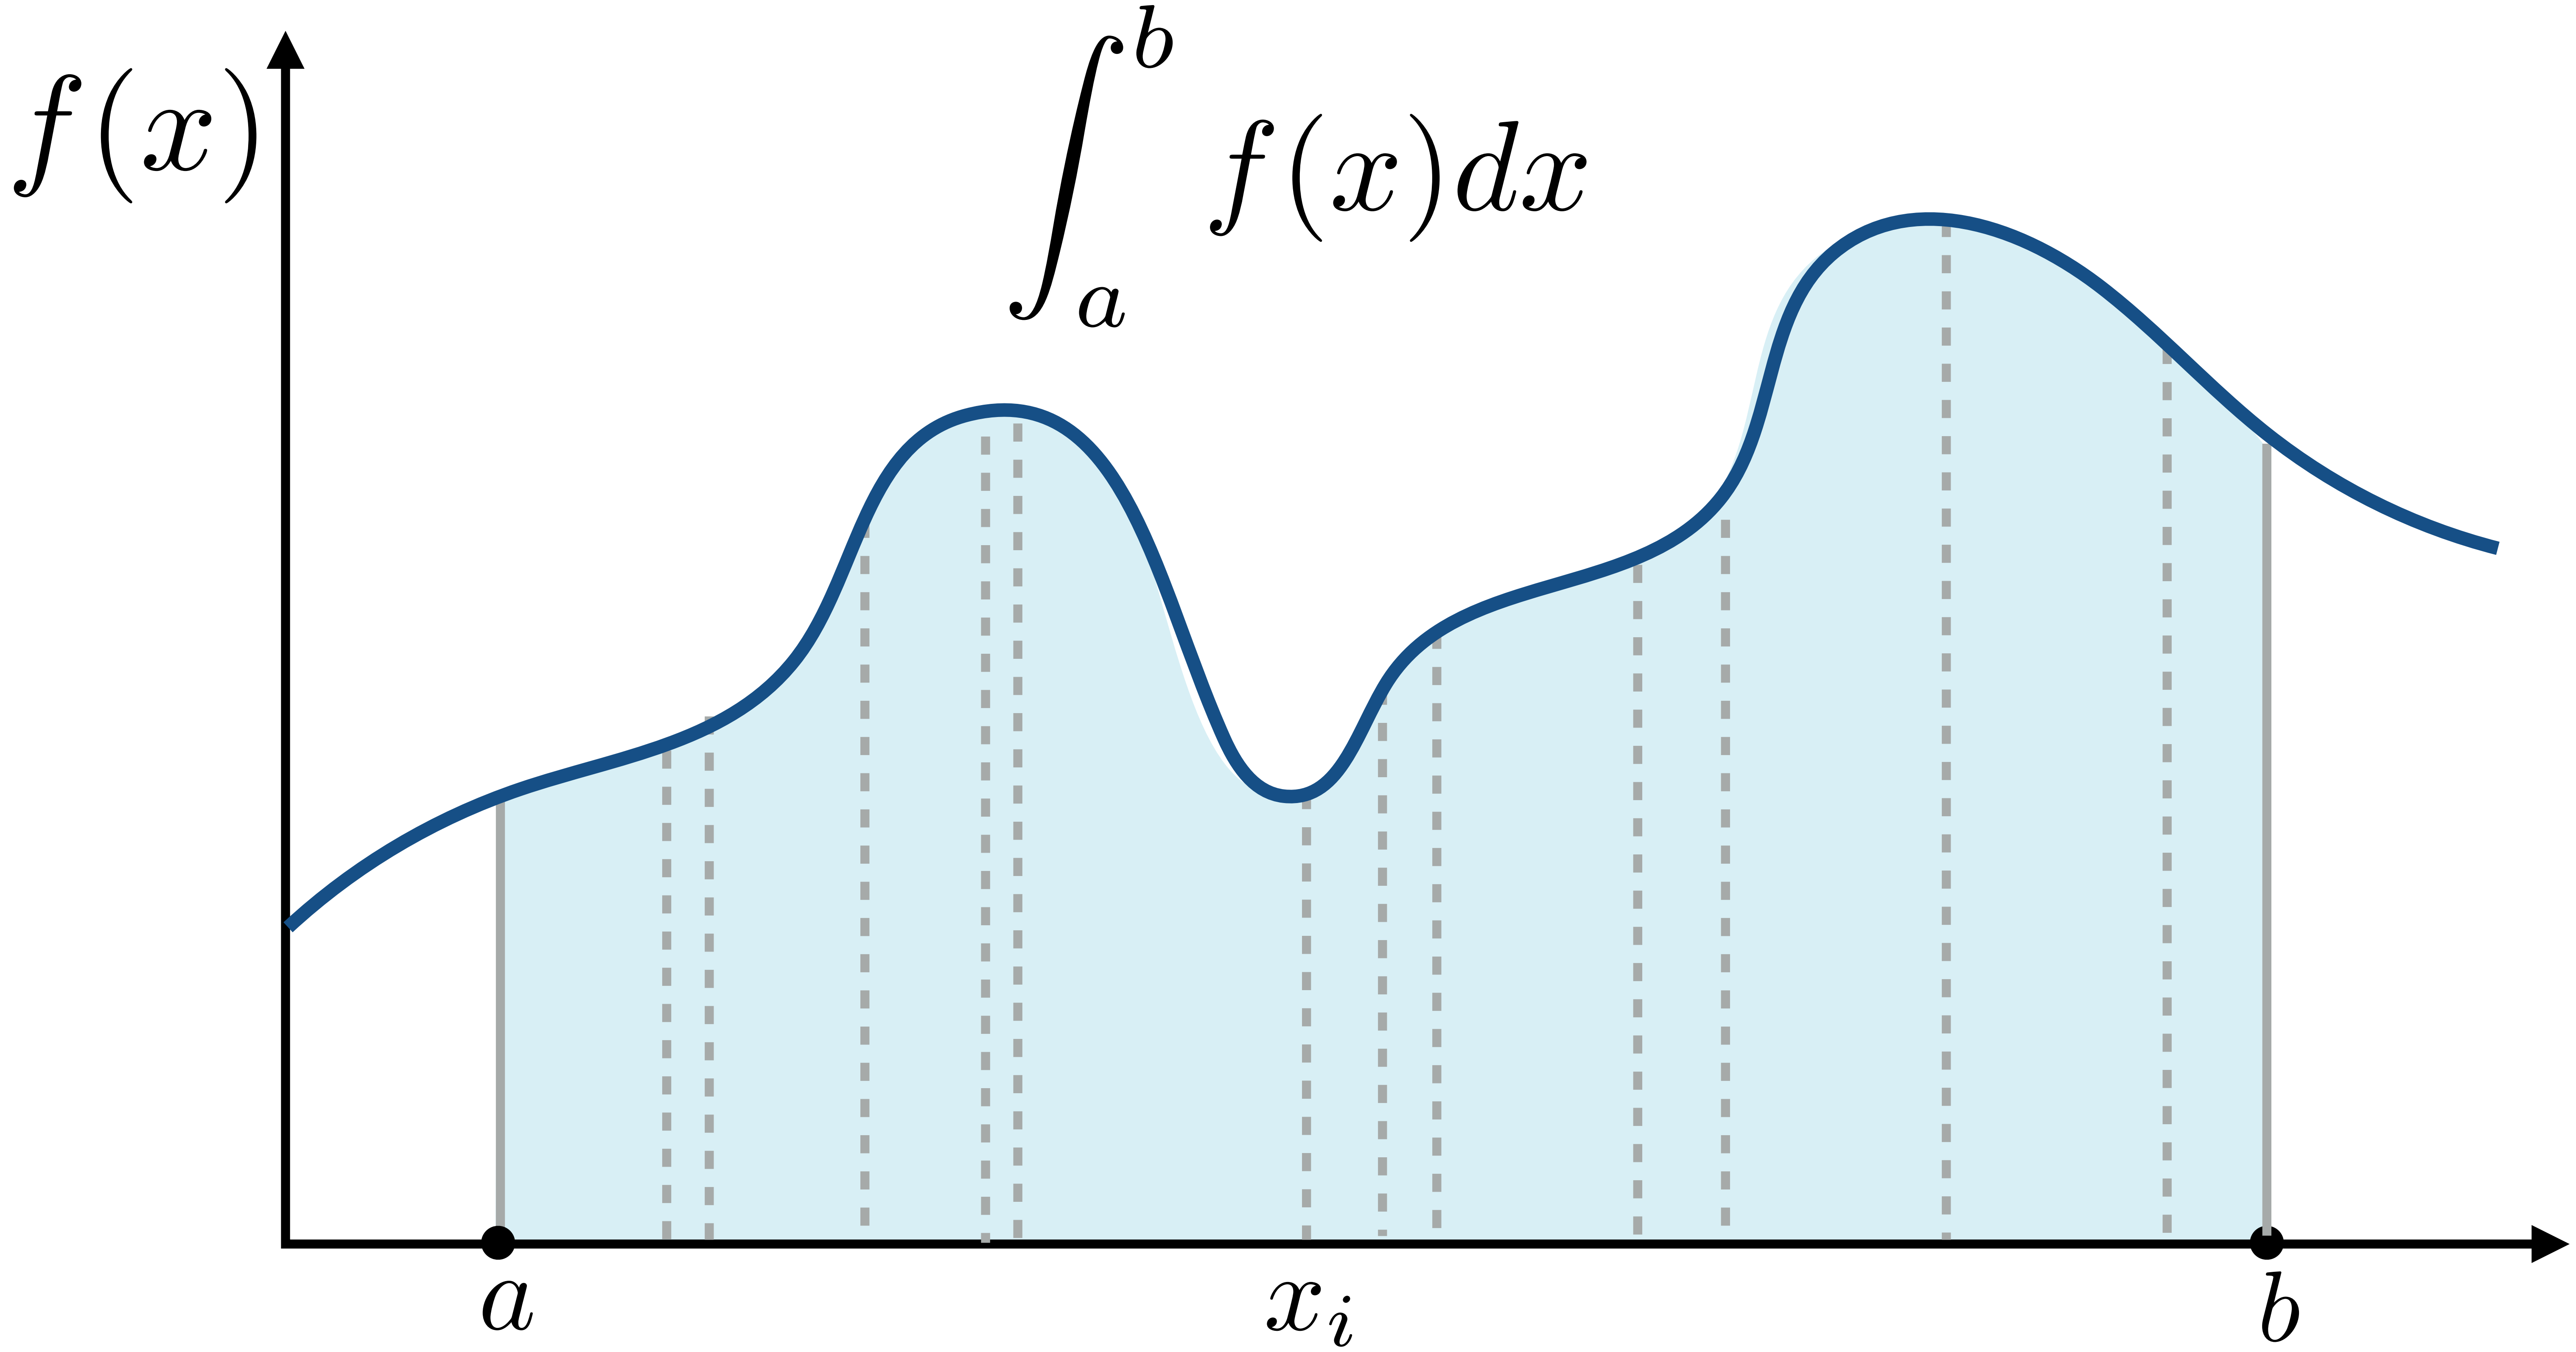

Monte Carlo Integration

Let us define the Monte Carlo estimator for the definite integral of given function $f(x)$

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

Random variable $$X_i \sim p(x)$$

Note: p(x) must be nonzero for all x where f(x) is nonzero

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

Example: Basic Monte Carlo Estimator

The basic Monte Carlo estimator is a simple special case where we sample with a uniform random variable

Uniform random variable

$$X_i \sim p(x) = C \quad \textbf{(constant)}$$

$$\int_a^b p(x)\,dx = 1$$

$$\implies \int_a^b C\,dx = 1$$

$$\implies C = \frac{1}{b-a}$$

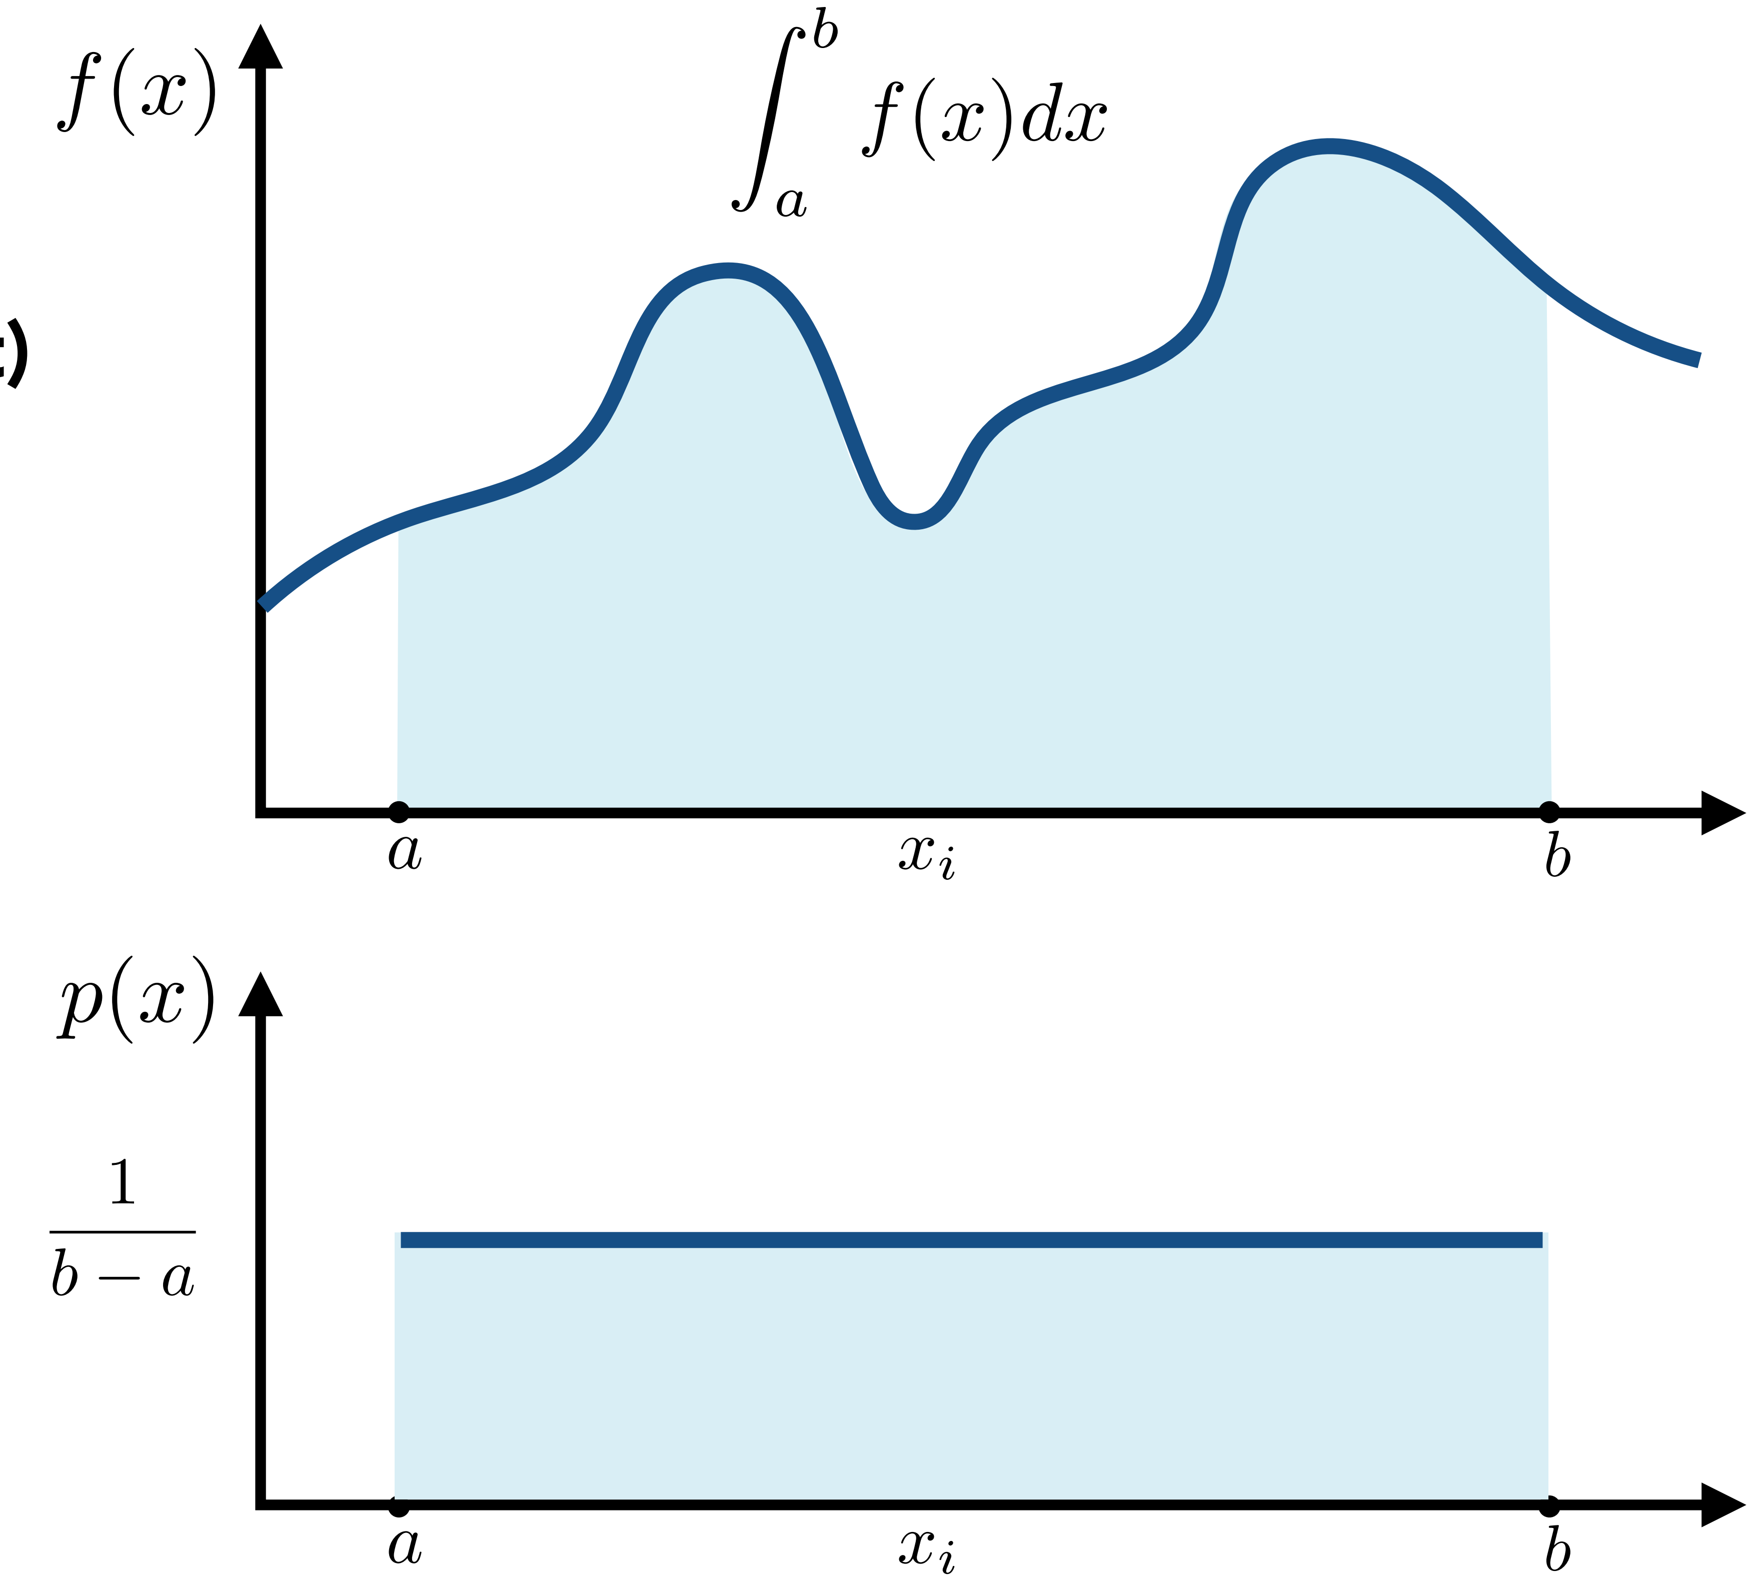

Example: Basic Monte Carlo Estimator

The basic Monte Carlo estimator is a simple special case where we sample with a uniform random variable

Basic Monte Carlo estimator (derivation)

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

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

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

Example: Basic Monte Carlo Estimator

Let us define the Monte Carlo estimator for the definite integral of given function $f(x)$

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

Uniform random variable $$X_i \sim p(x) = \frac{1}{b-a}$$

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

Unbiased Estimator

Definition: A randomized integral estimator is *unbiased* if its expected value is the desired integral.

Fact: the general and basic Monte Carlo estimators are unbiased (proof on next slide)

Why do we want unbiased estimators?

Proof That Monte Carlo Estimator Is Unbiased

$$
\begin{aligned}
E[F_N] &= E\left[\frac{1}{N}\sum_{i=1}^{N}\frac{f(X_i)}{p(X_i)}\right] \\
&= \frac{1}{N}\sum_{i=1}^{N} E\left[\frac{f(X_i)}{p(X_i)}\right] \\
&= \frac{1}{N}\sum_{i=1}^{N}\int_a^b \frac{f(x)}{p(x)}\,p(x)\mathrm{d}x \\
&= \frac{1}{N}\sum_{i=1}^{N}\int_a^b f(x)\,\mathrm{d}x \\
&= \int_a^b f(x)\,\mathrm{d}x
\end{aligned}
$$

Properties of expected values:

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

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

The expected value of the Monte Carlo estimator is the desired integral.

Variance of a Random Variable

Definition

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

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

$$V\left[\sum_{i=1}^{N} Y_i\right] = \sum_{i=1}^{N} V[Y_i] \qquad\qquad V[aY] = a^2\, V[Y]$$

More Random Samples Reduces Variance

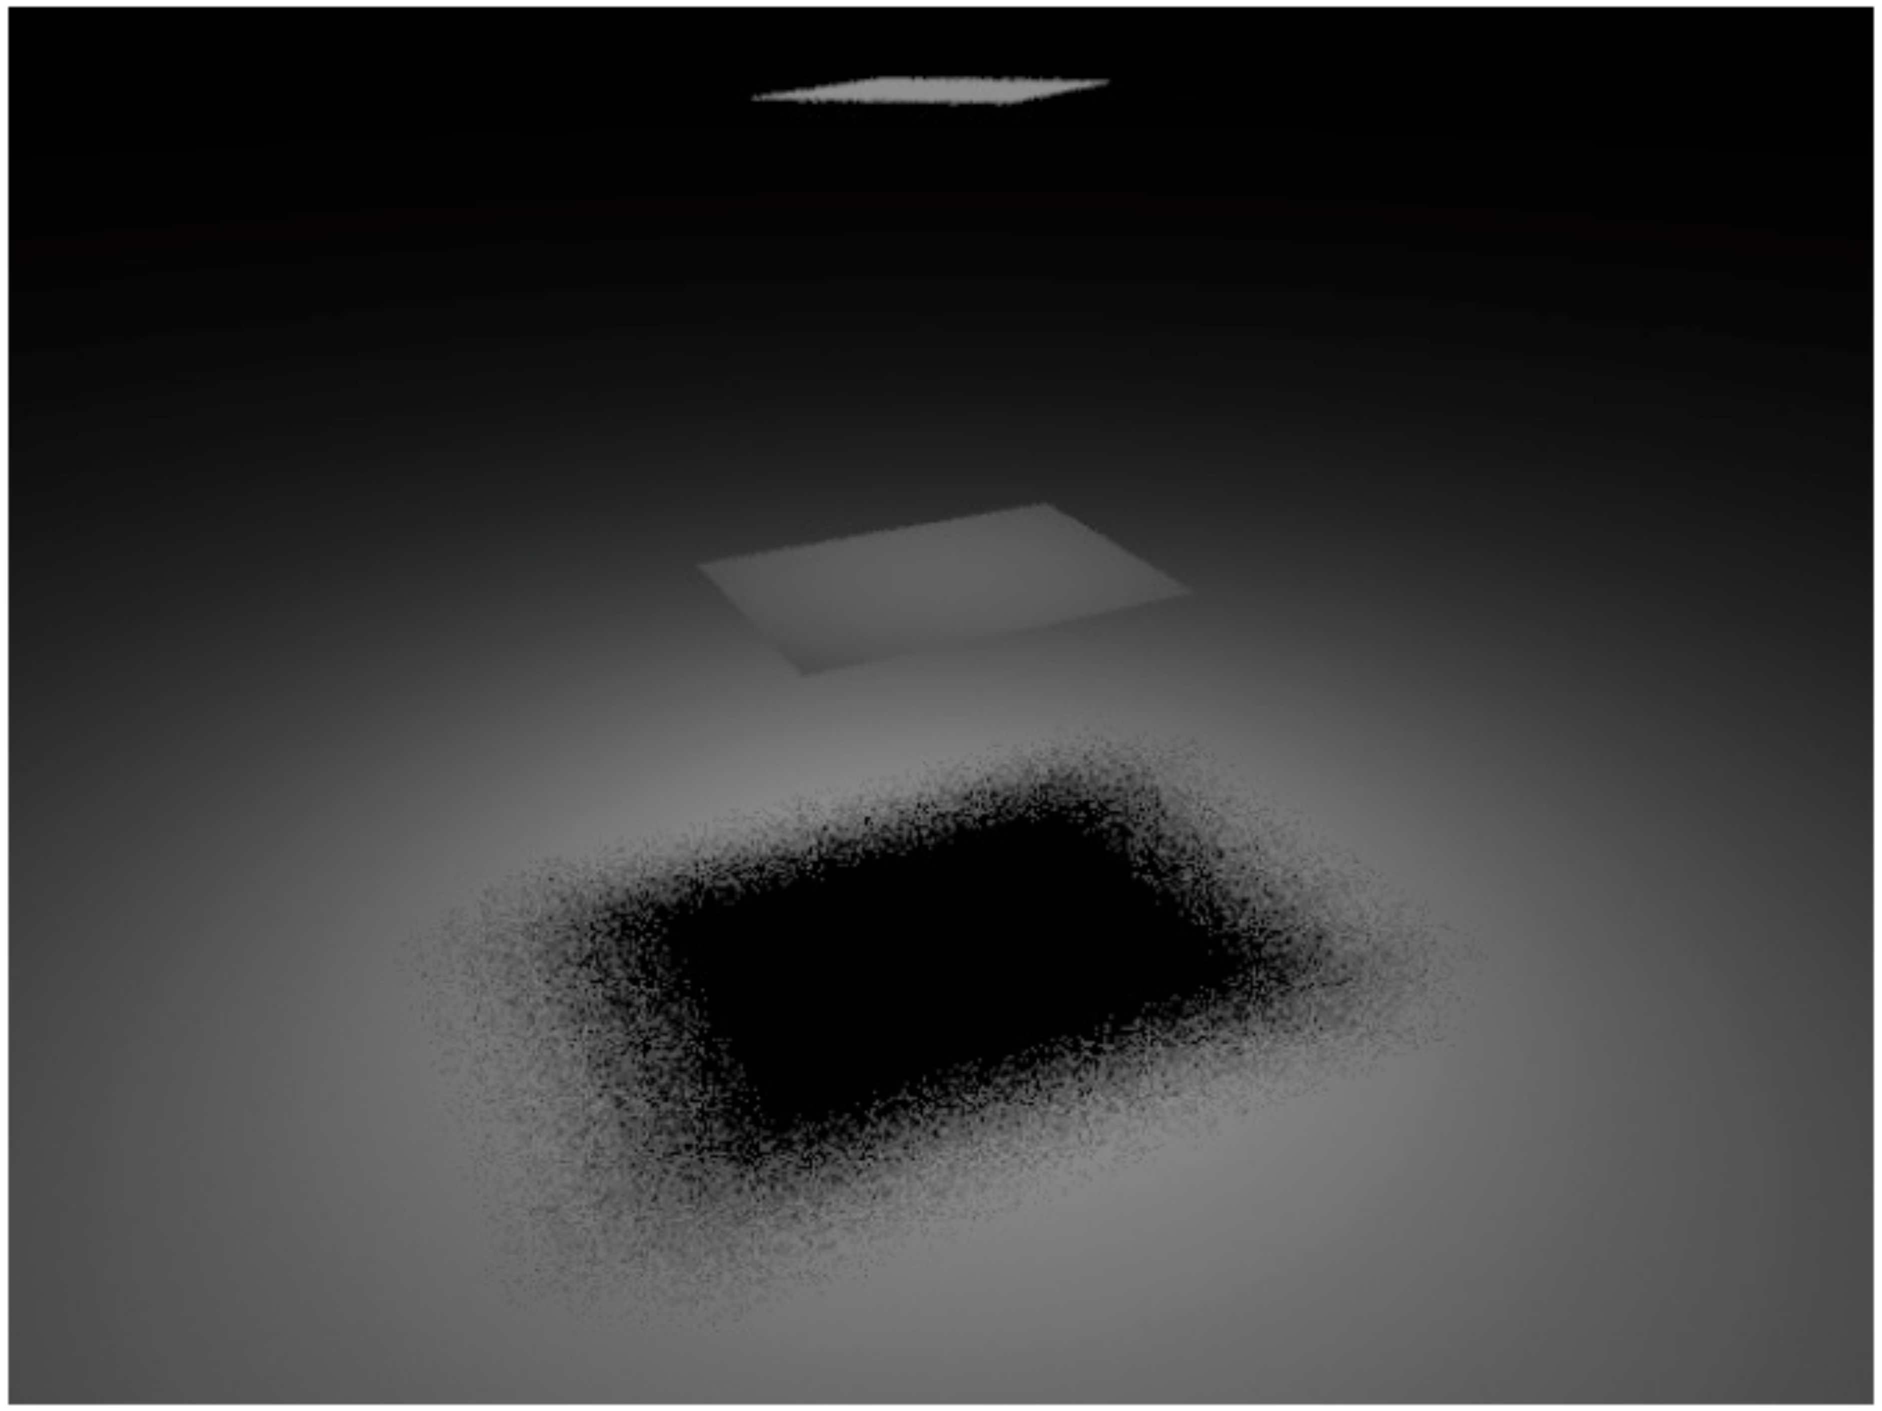

1 shadow ray

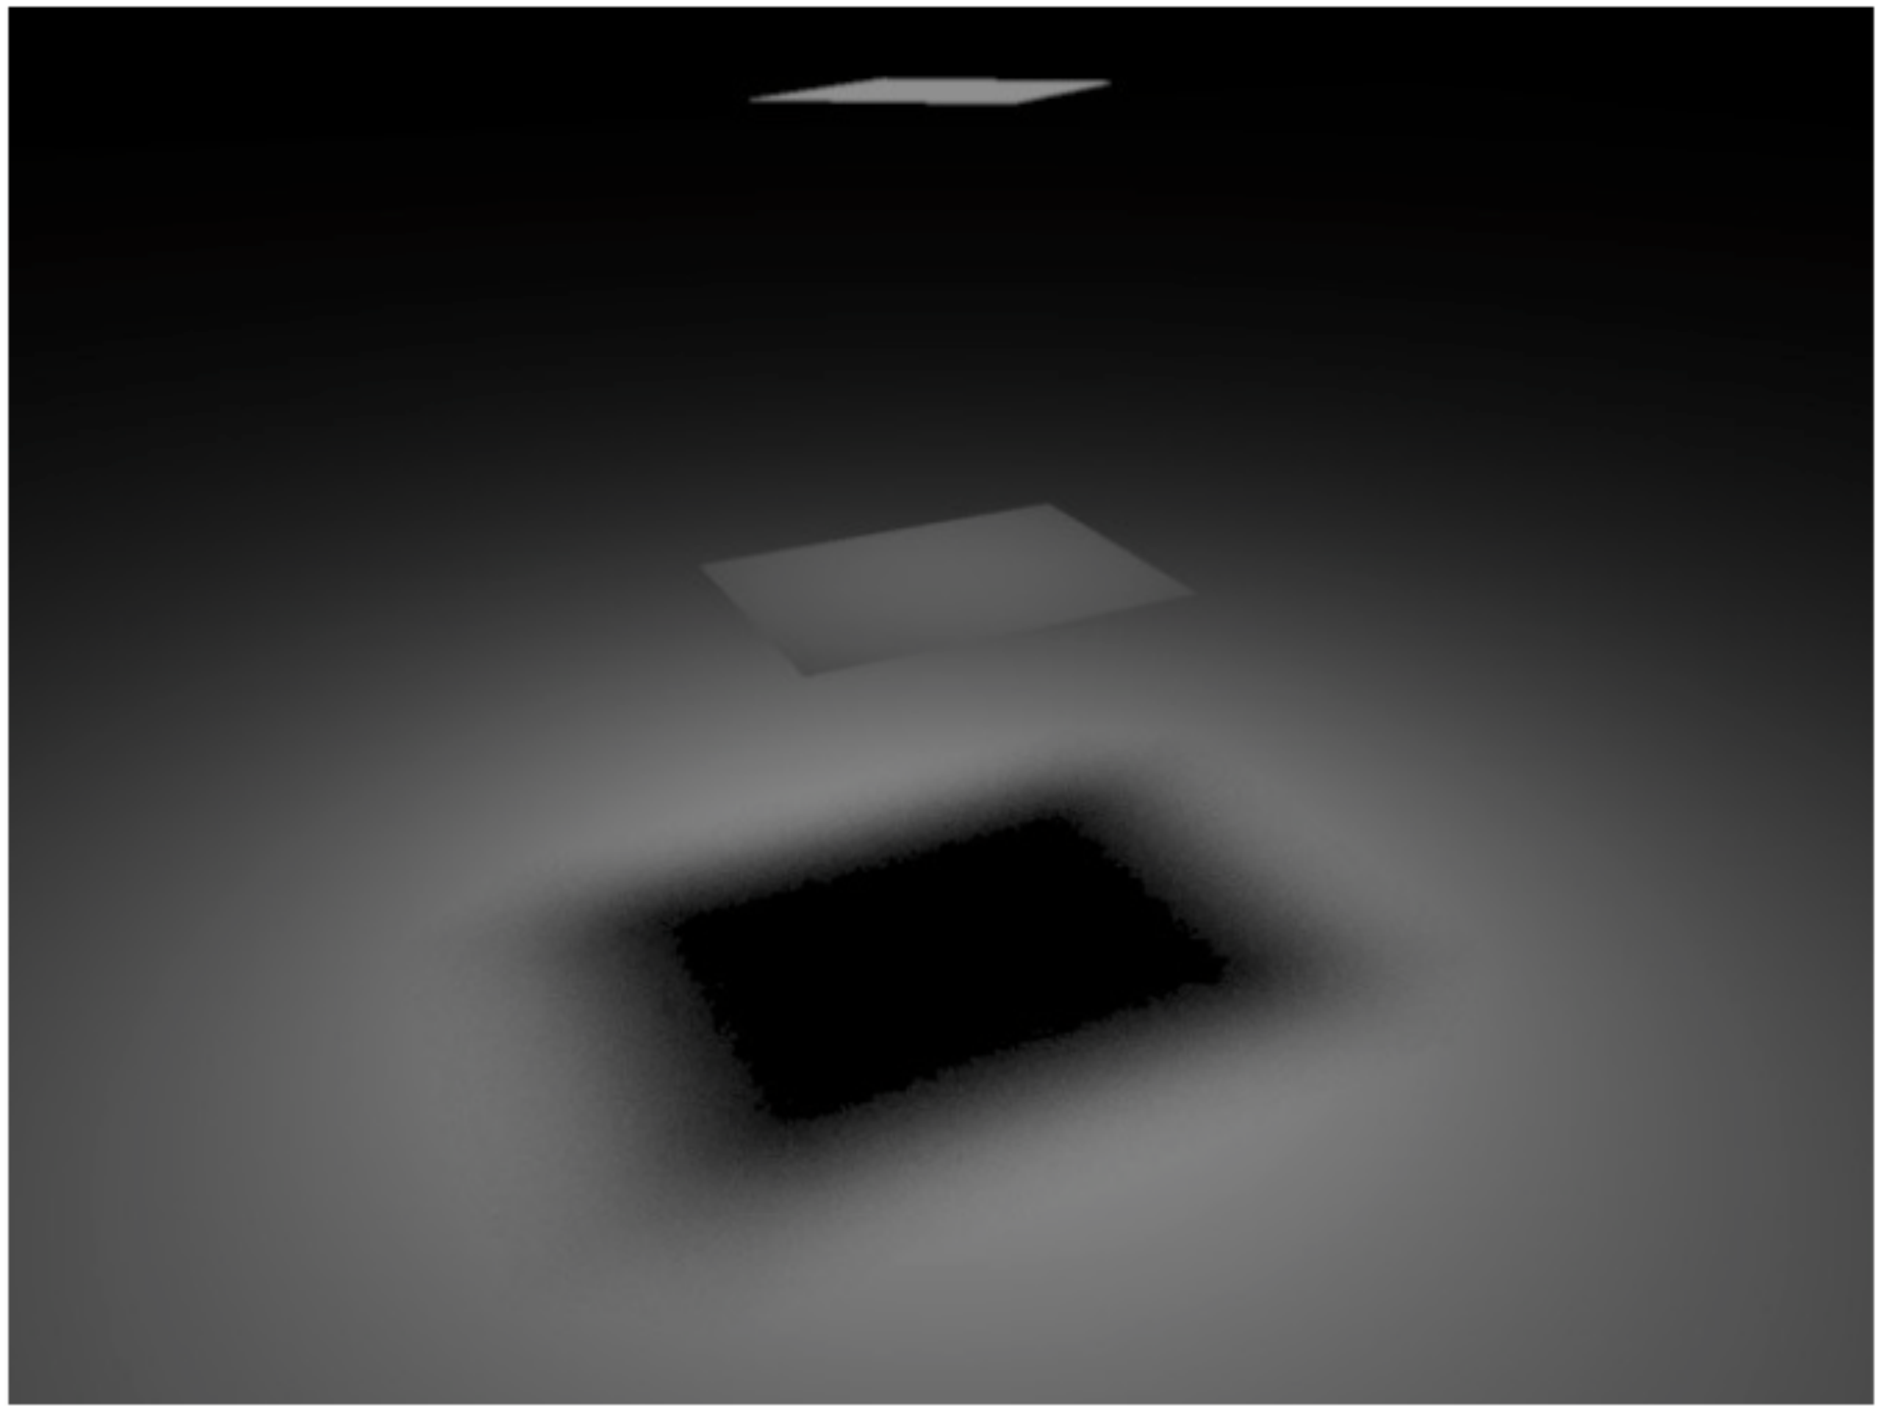

16 shadow rays

Definite Integral Can Be N-Dimensional

Example in 3D:

$$\int_{x_0}^{x_1} \int_{y_0}^{y_1} \int_{z_0}^{z_1} f(x, y, z)\, dx\, dy\, dz$$

Uniform 3D random variable*

$$X_i \sim p(x, y, z) = \frac{1}{x_1 - x_0} \frac{1}{y_1 - y_0} \frac{1}{z_1 - z_0}$$

Basic 3D MC estimator*

$$F_N = \frac{(x_1 - x_0)(y_1 - y_0)(z_1 - z_0)}{N} \sum_{i=1}^{N} f(X_i)$$

*** Generalizes to arbitrary N-dimensional PDFs**

Example: Monte Carlo Estimate Of Direct Lighting Integral

Direct Lighting (Irradiance) Estimate

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

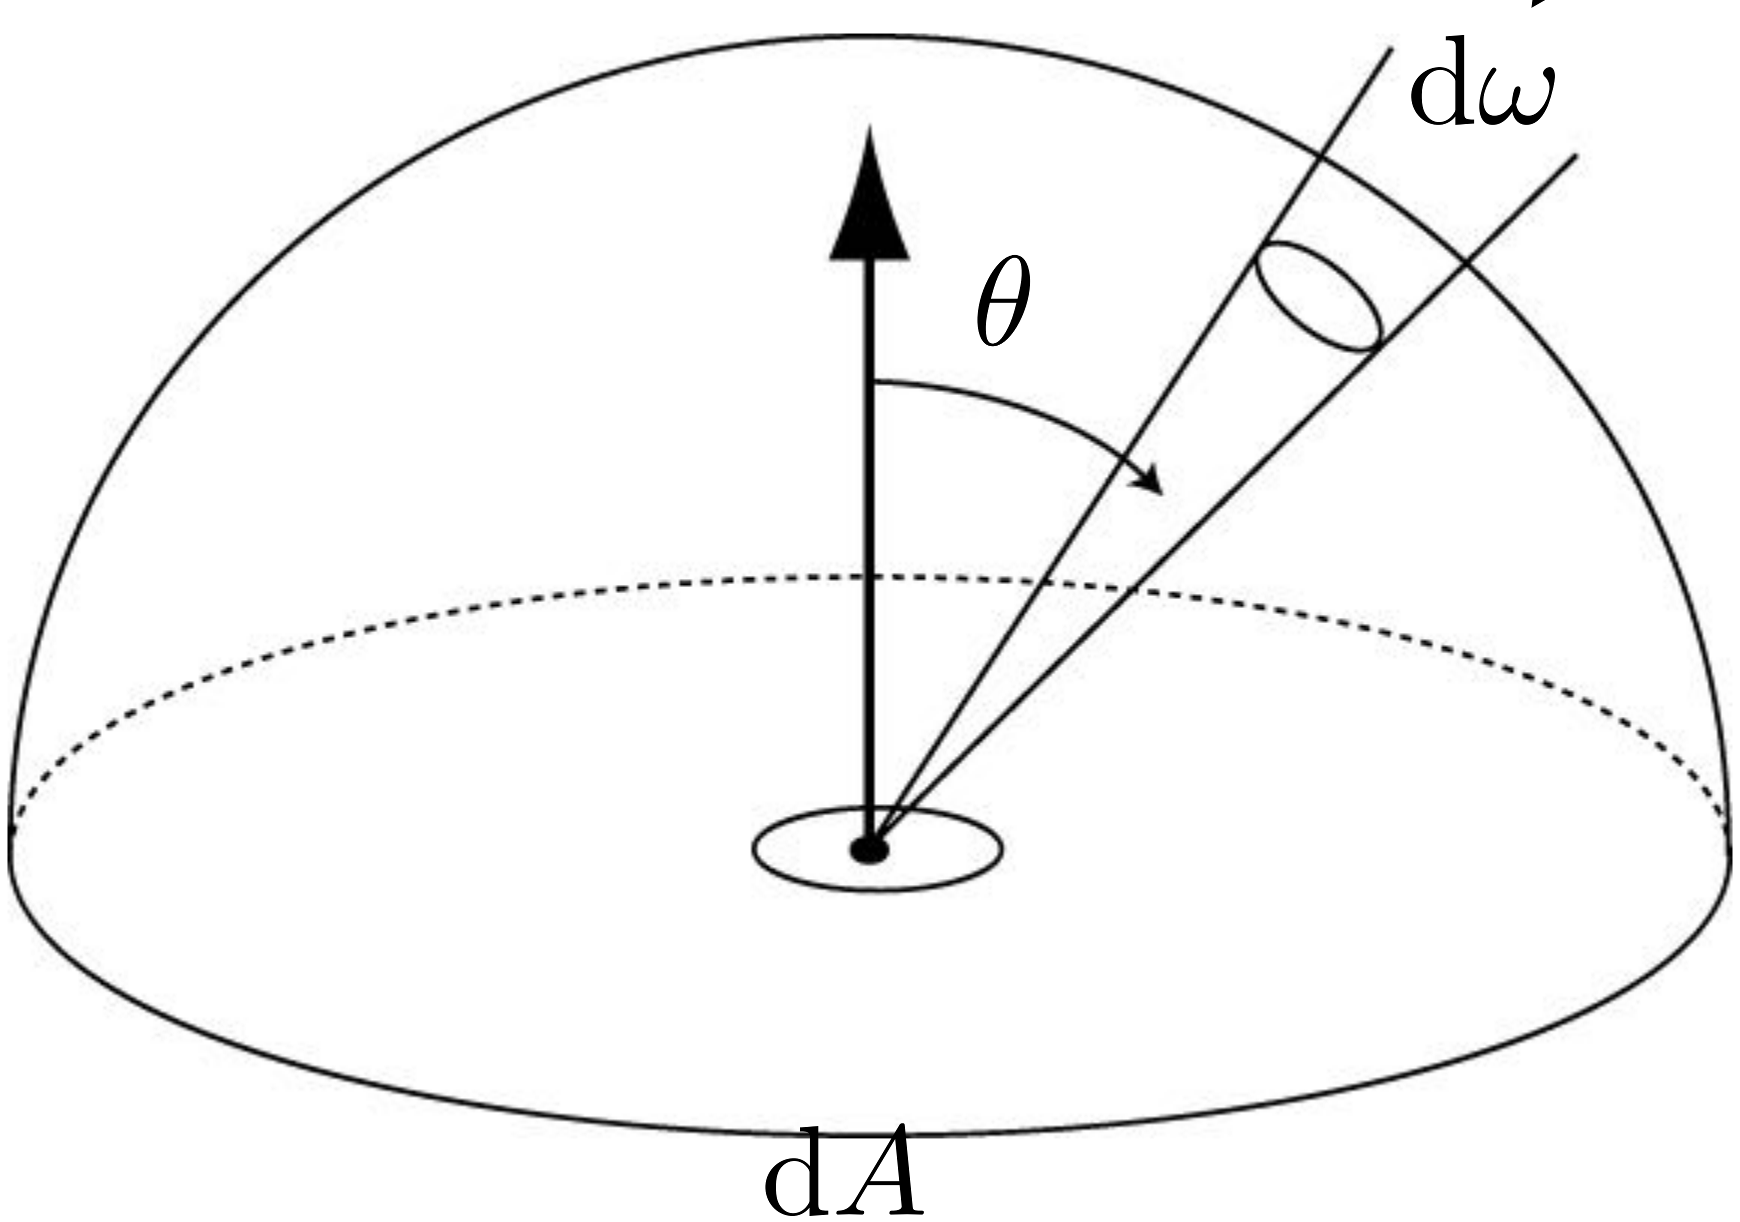

Idea: sample directions over hemisphere uniformly in solid angle

Estimator:

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

$$Y_i = f(X_i)$$

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

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

Direct Lighting (Irradiance) Estimate

Sample directions over hemisphere uniformly in solid angle

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

Given surface point *p*

A ray tracer evaluates radiance along a ray

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

Direct Lighting - Solid Angle Sampling

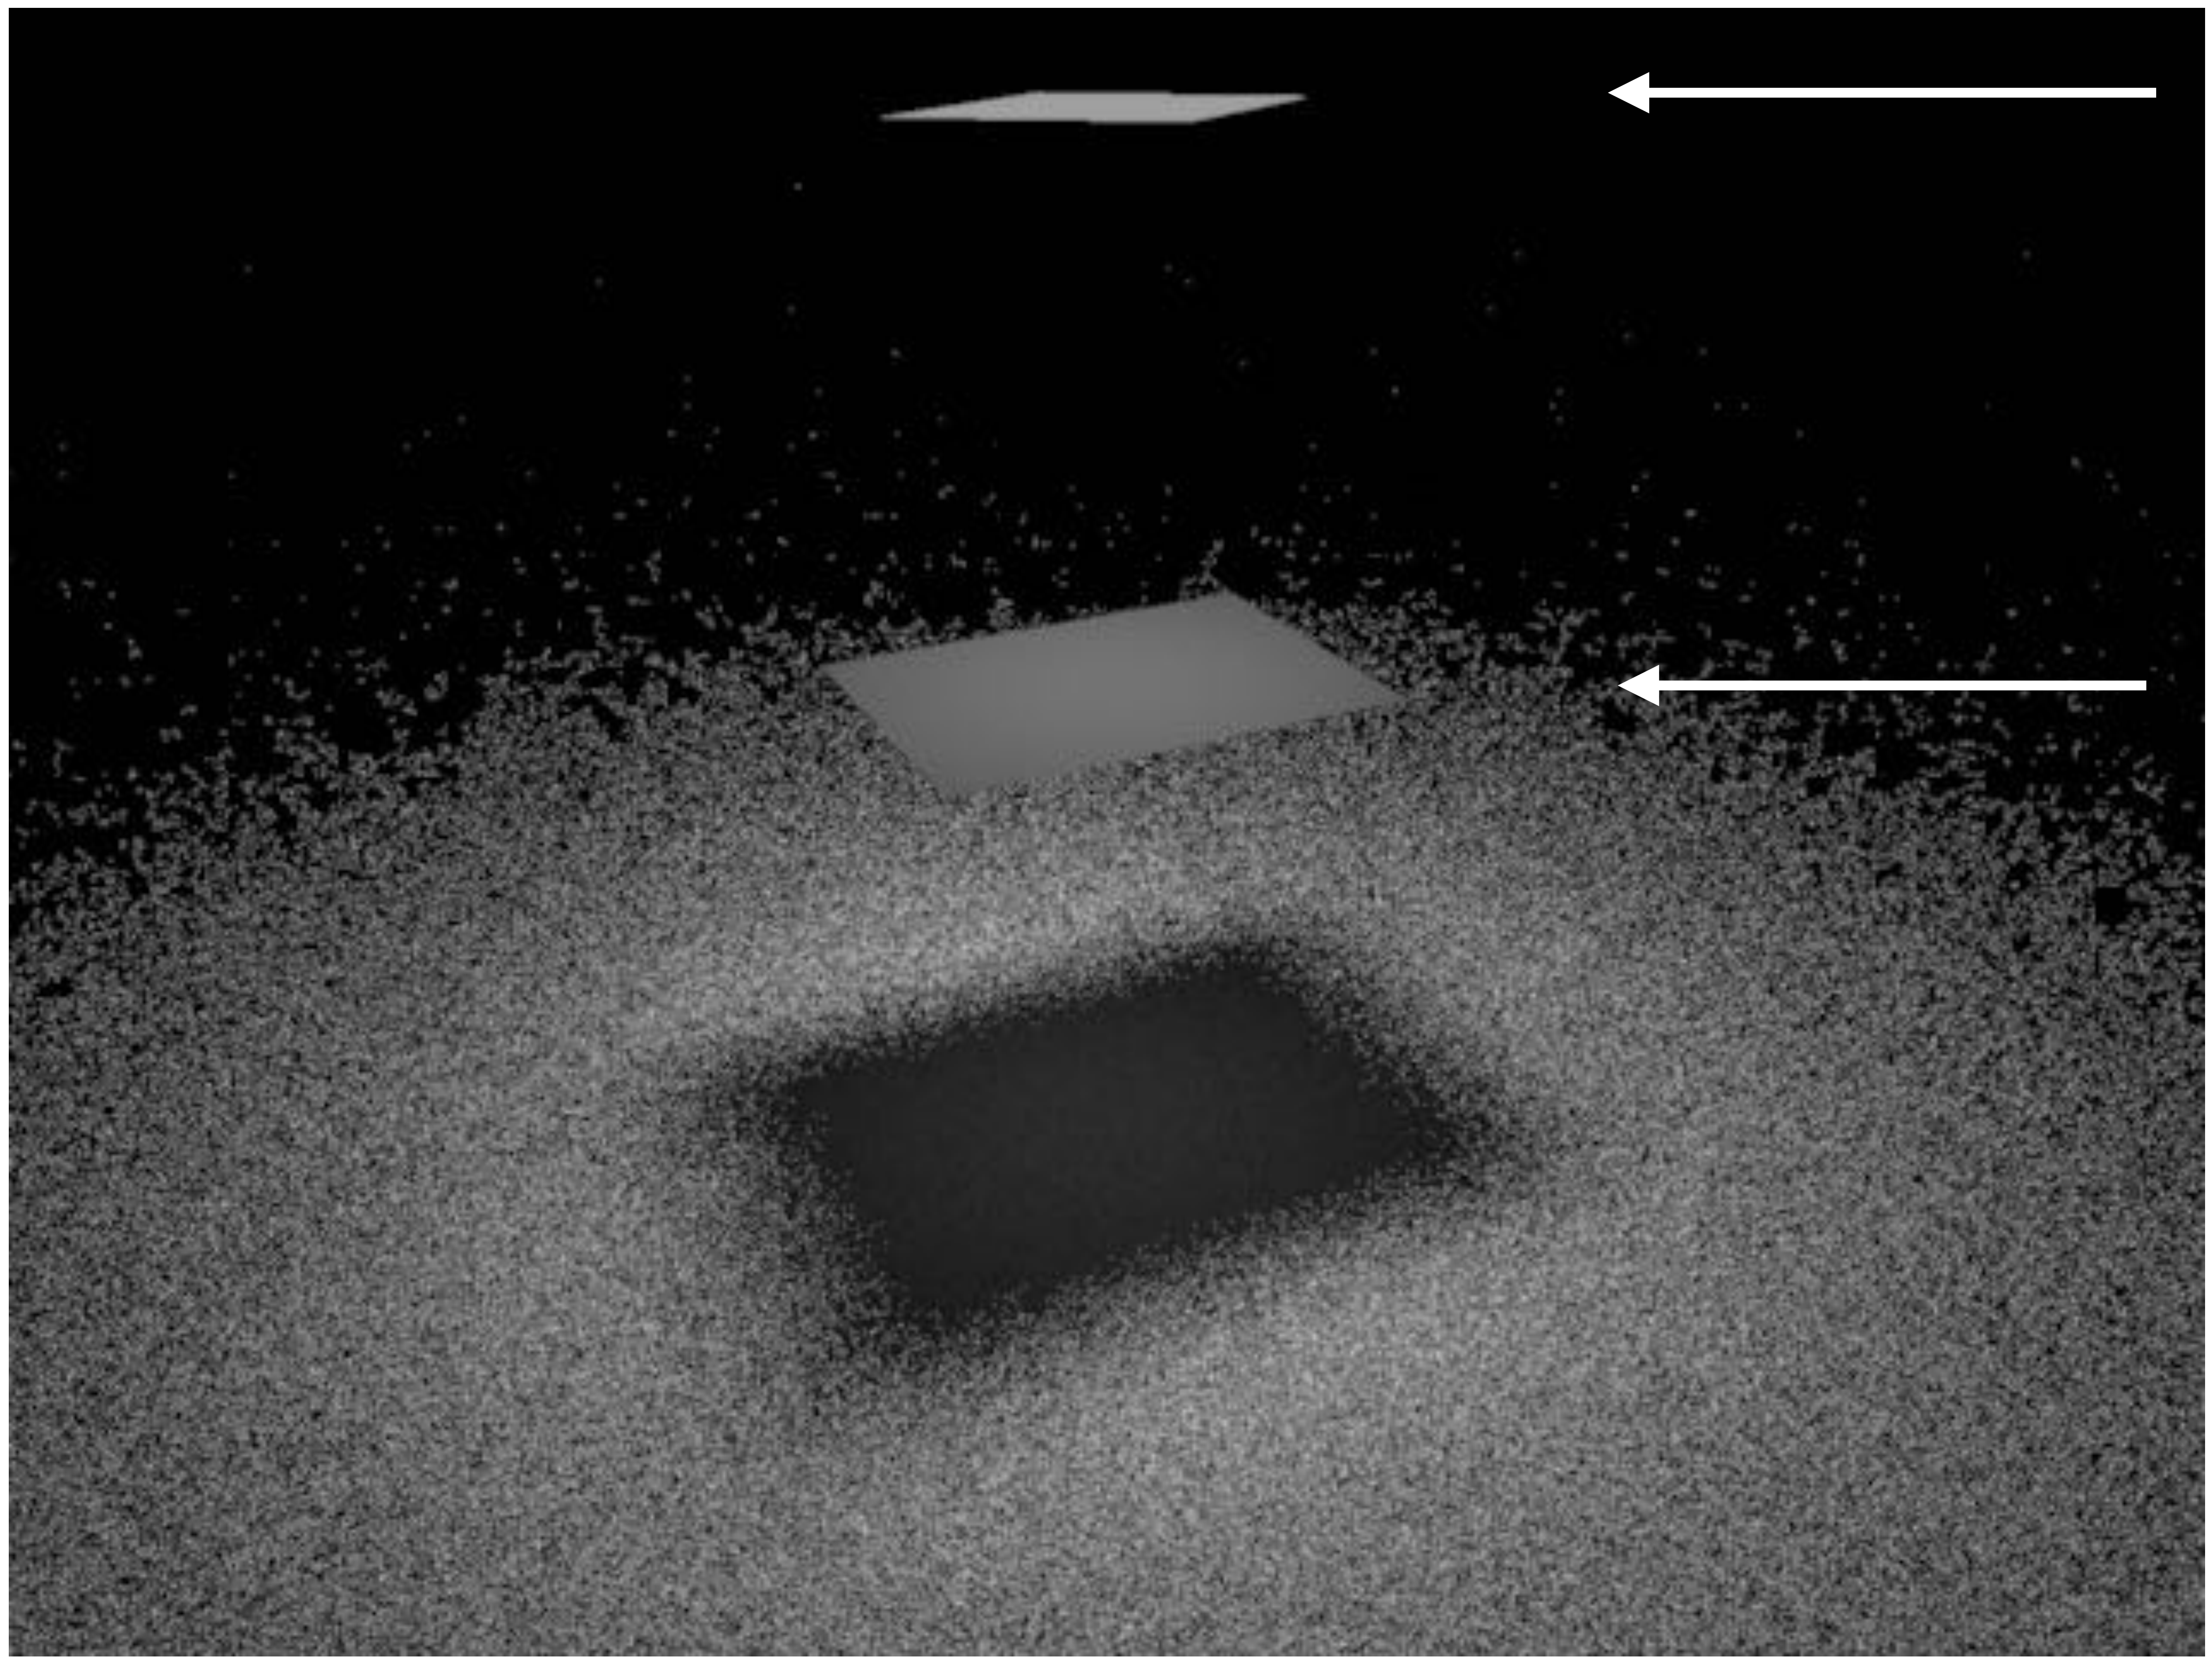

Trace 100 rays per pixel

Hemispherical Solid Angle Sampling 100 rays

(random directions drawn uniformly from hemisphere)

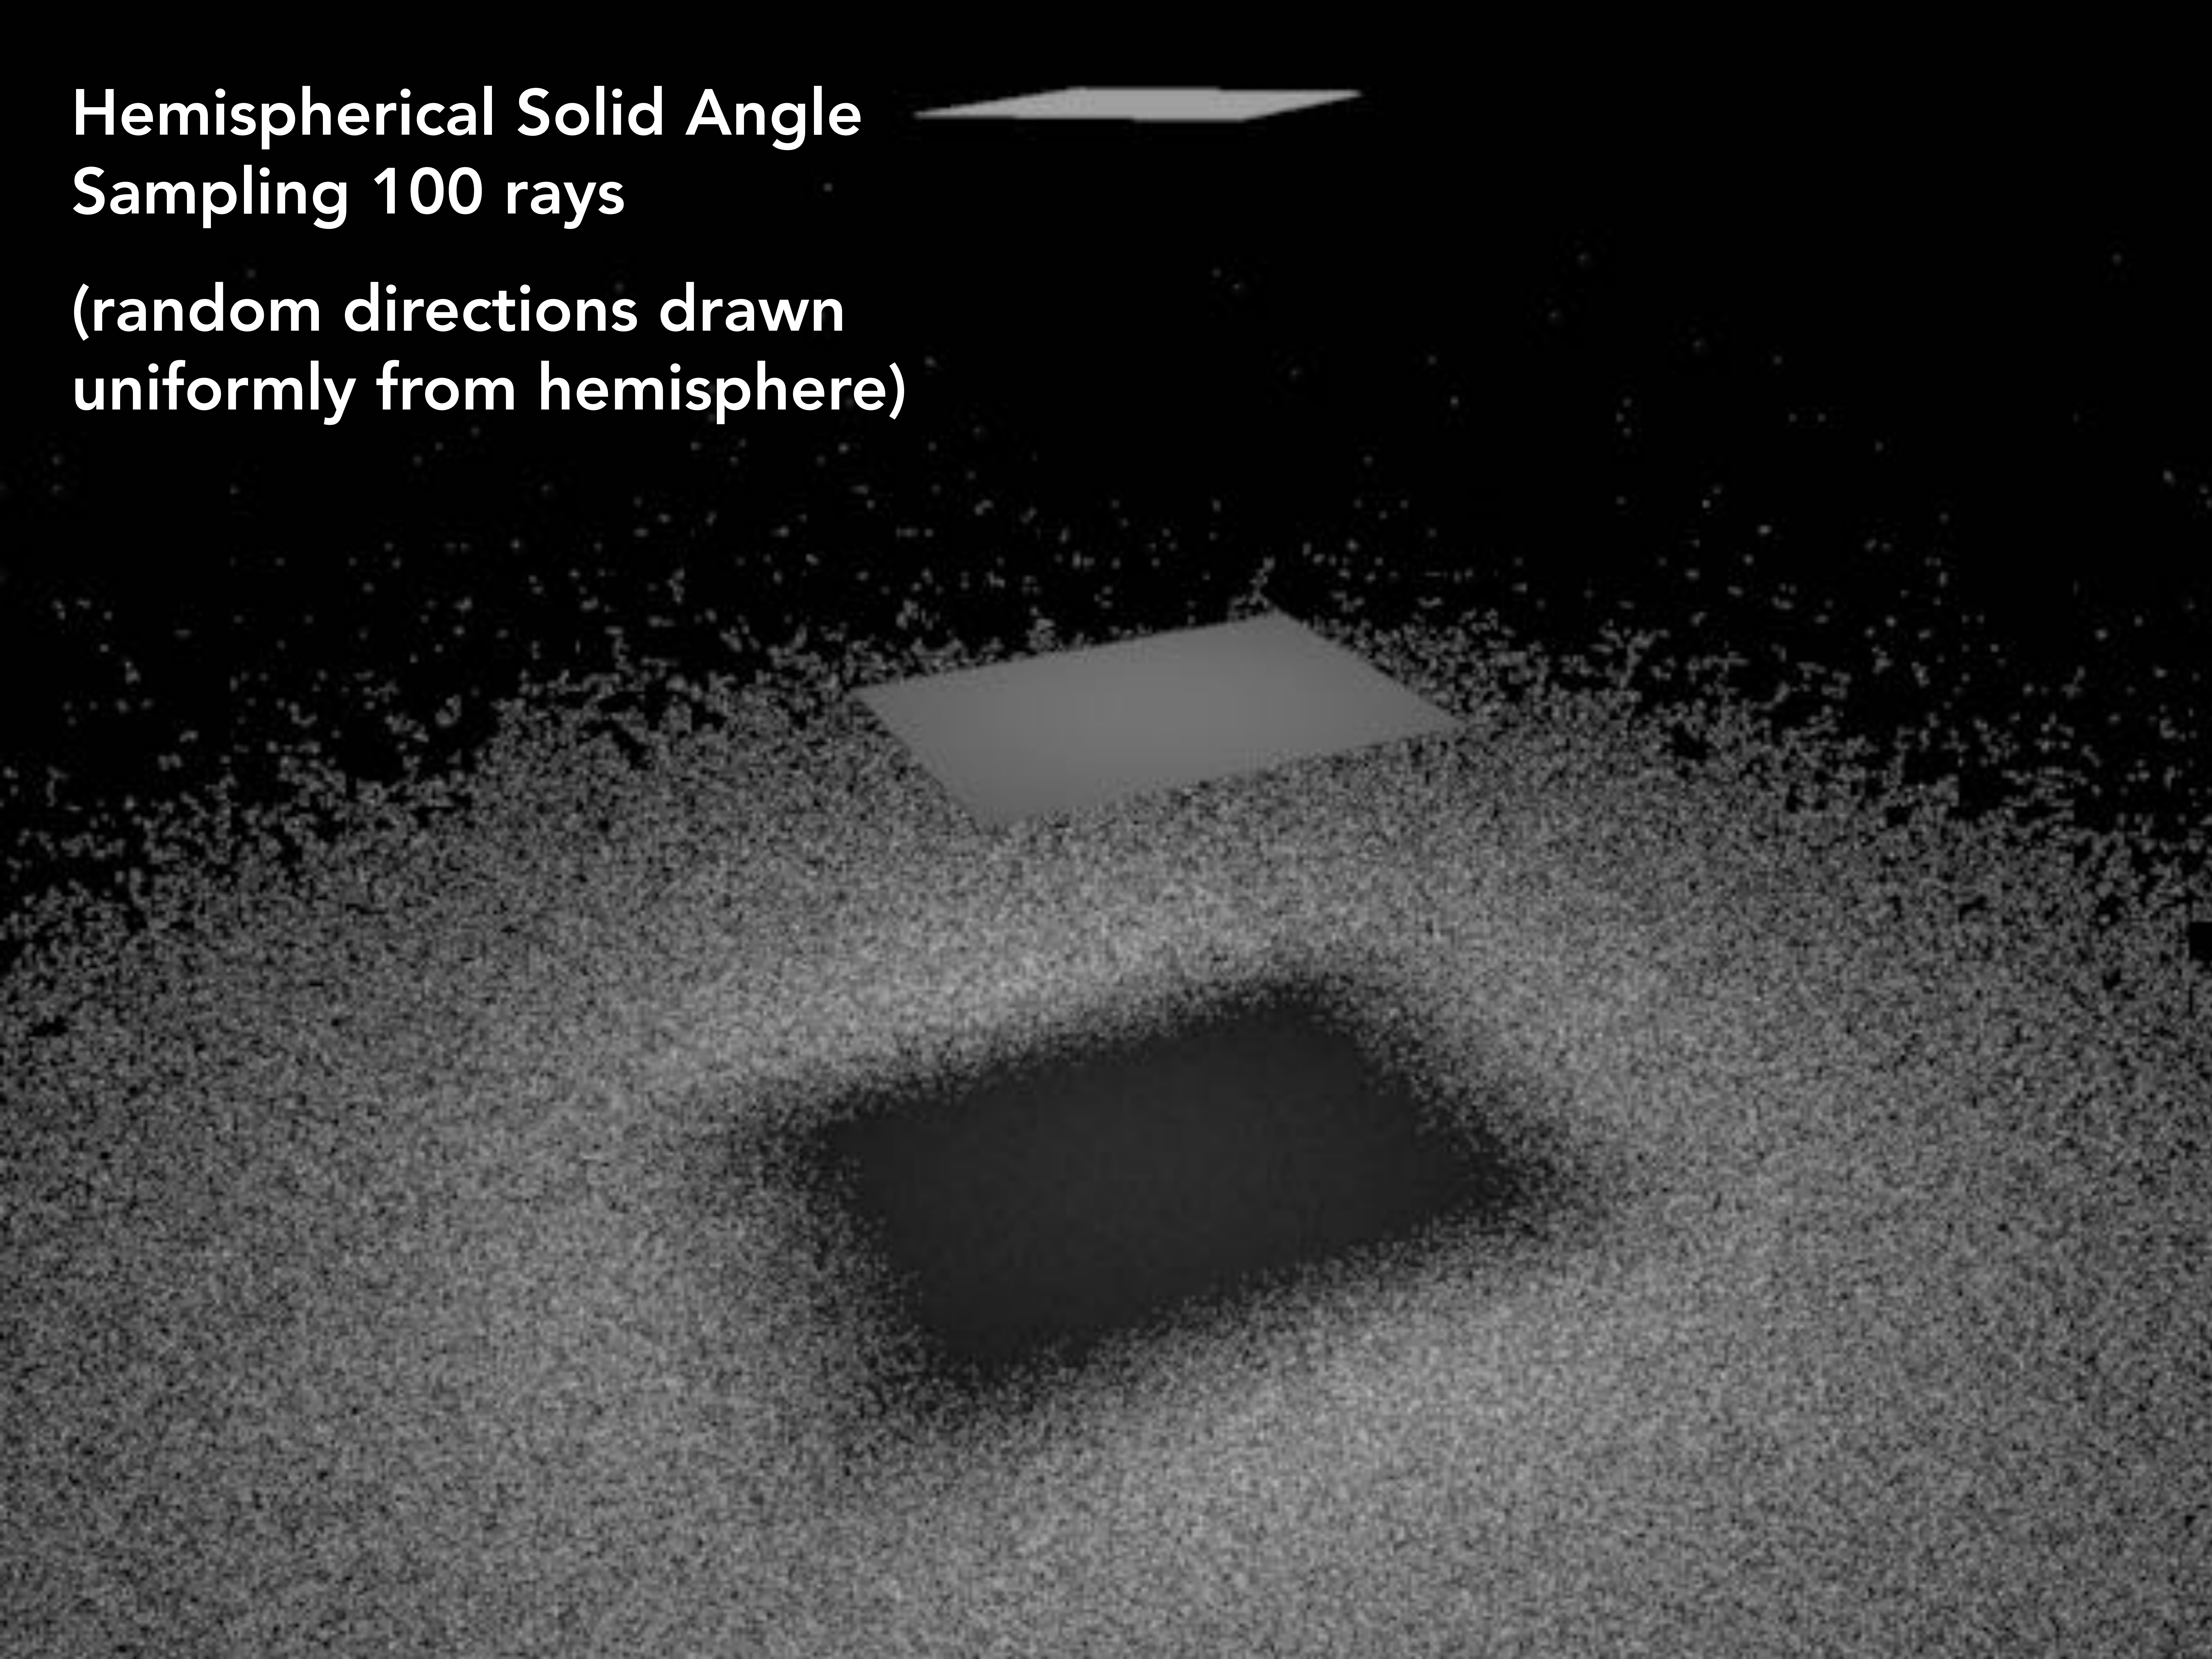

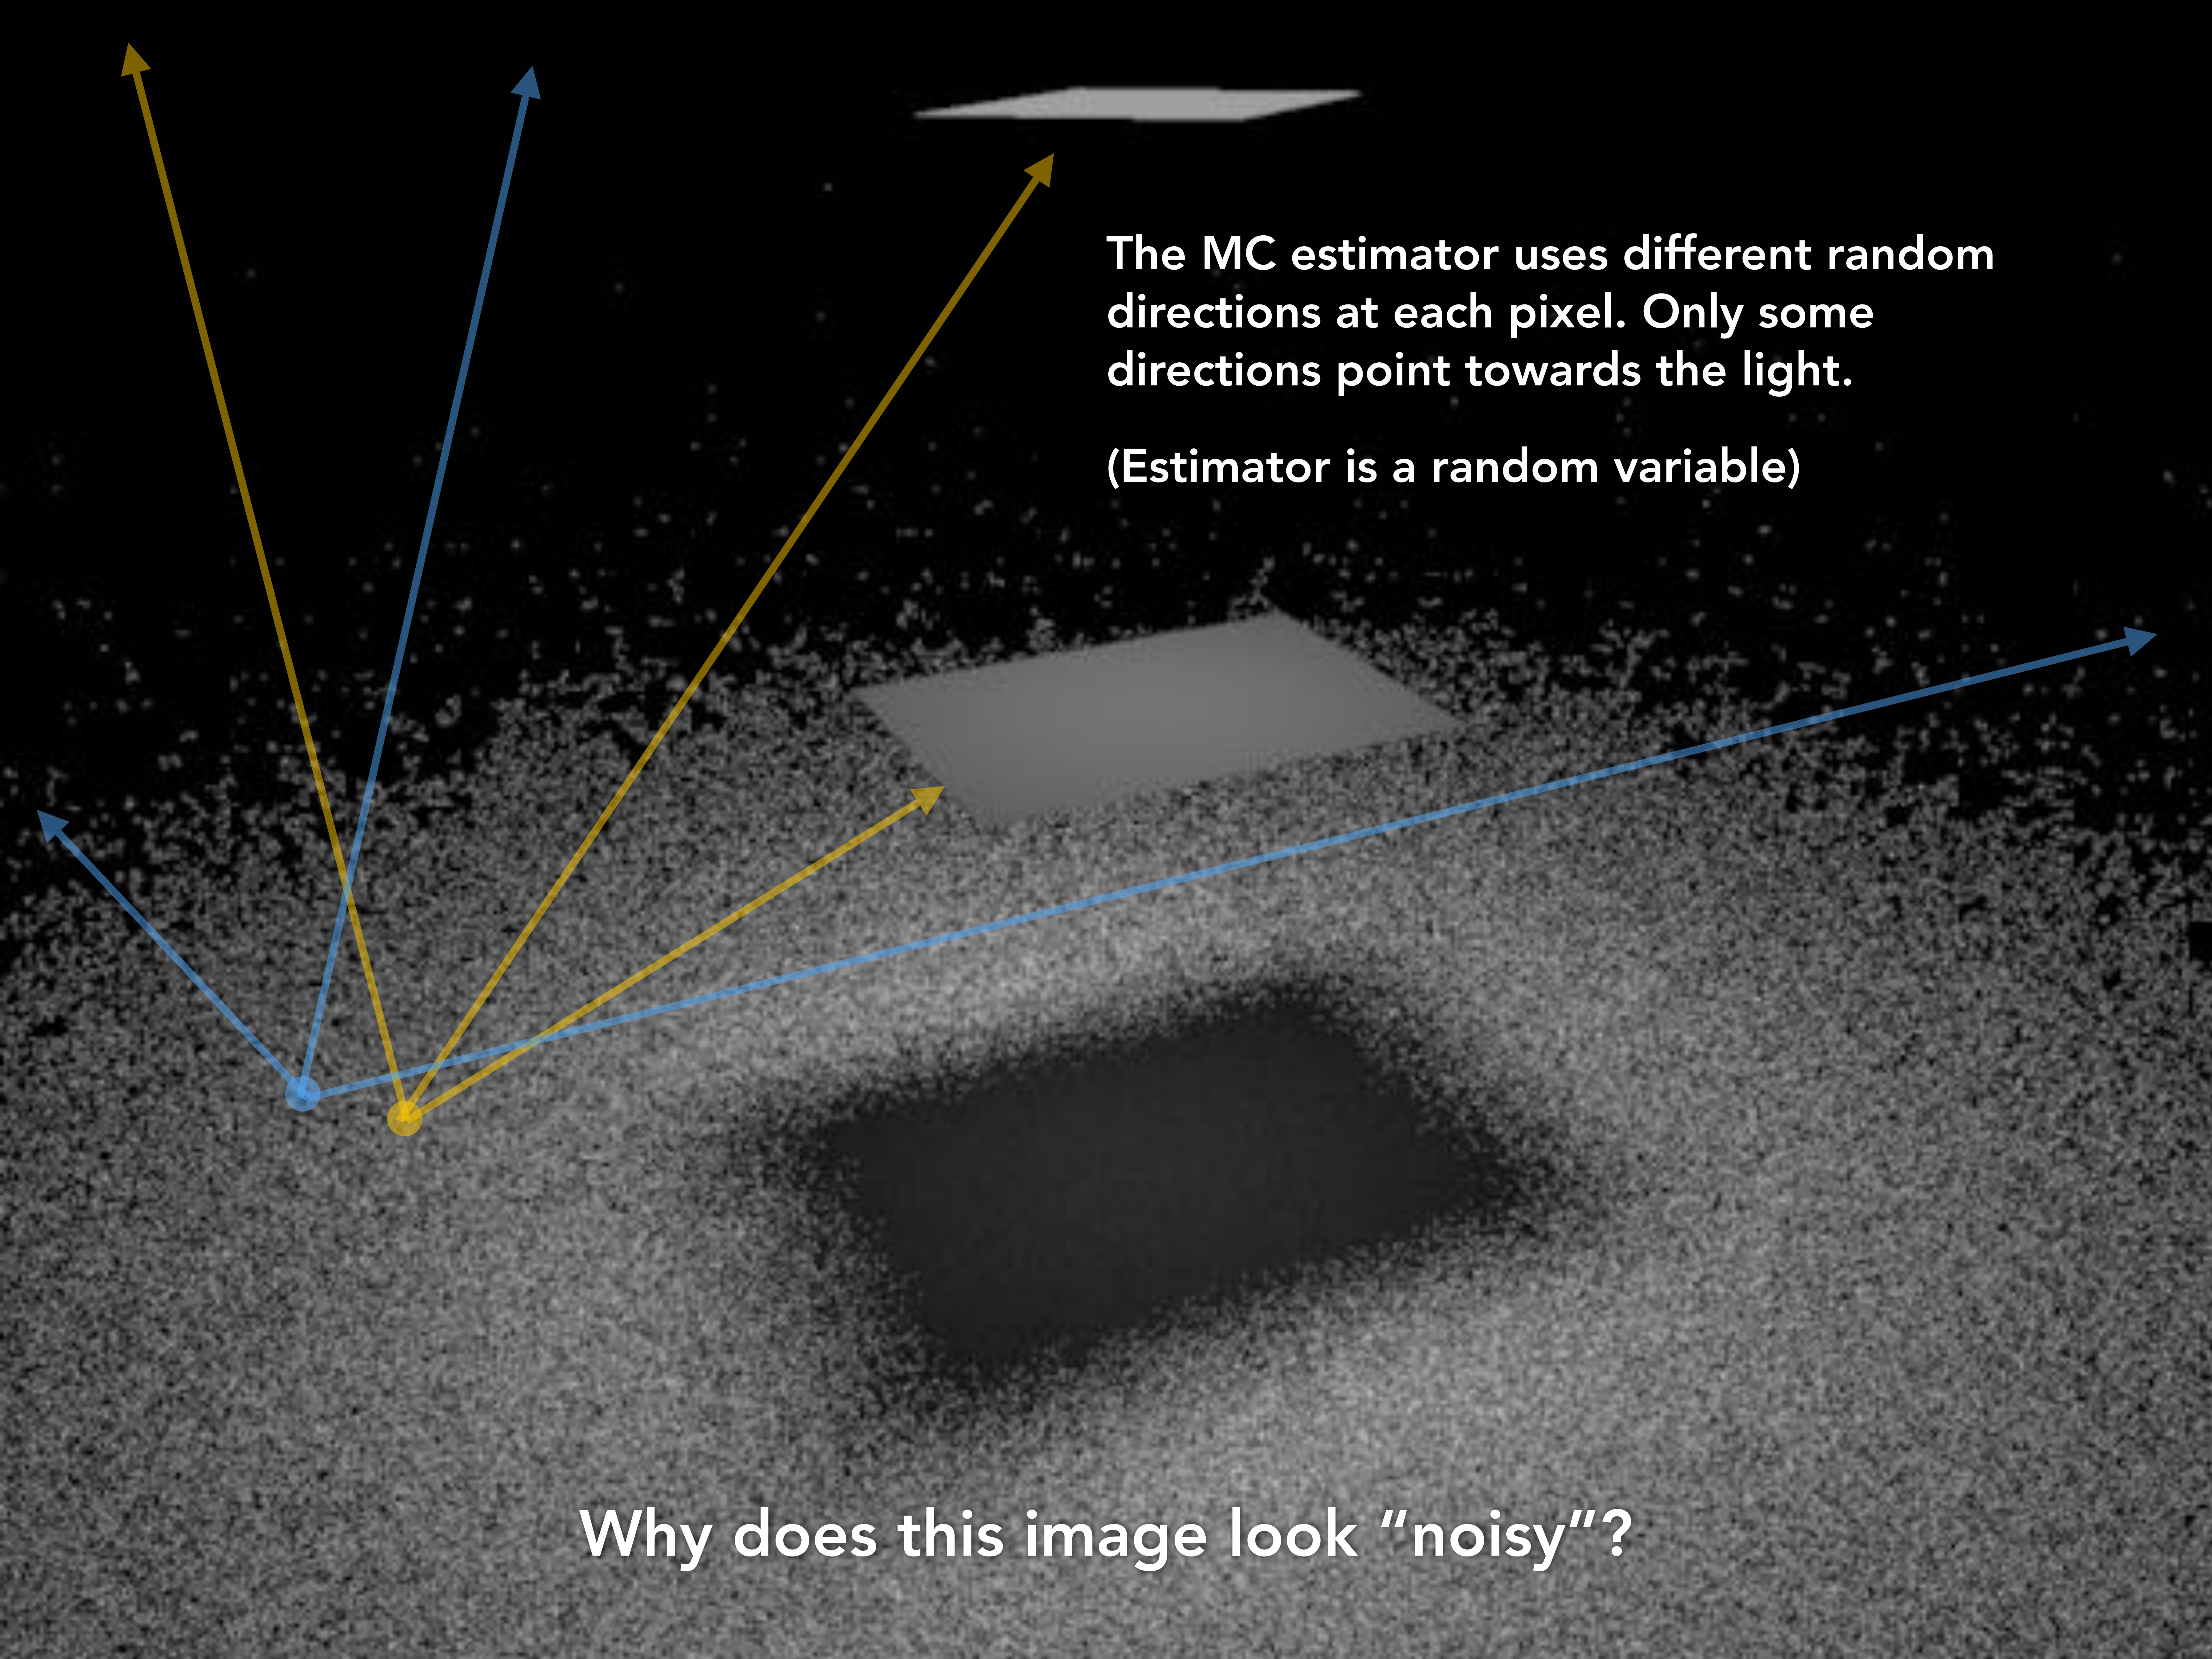
The MC estimator uses different random directions at each pixel. Only some directions point towards the light.
(Estimator is a random variable)
Why does this image look “noisy”?

Observation: incoming radiance is zero for most directions in this scene

Idea: integrate only over the area of the light (directions where incoming radiance could be non-zero)

Importance Sampling

Importance Sampling

Simple idea: sample the integrand according to how much we expect it to contribute to the integral.

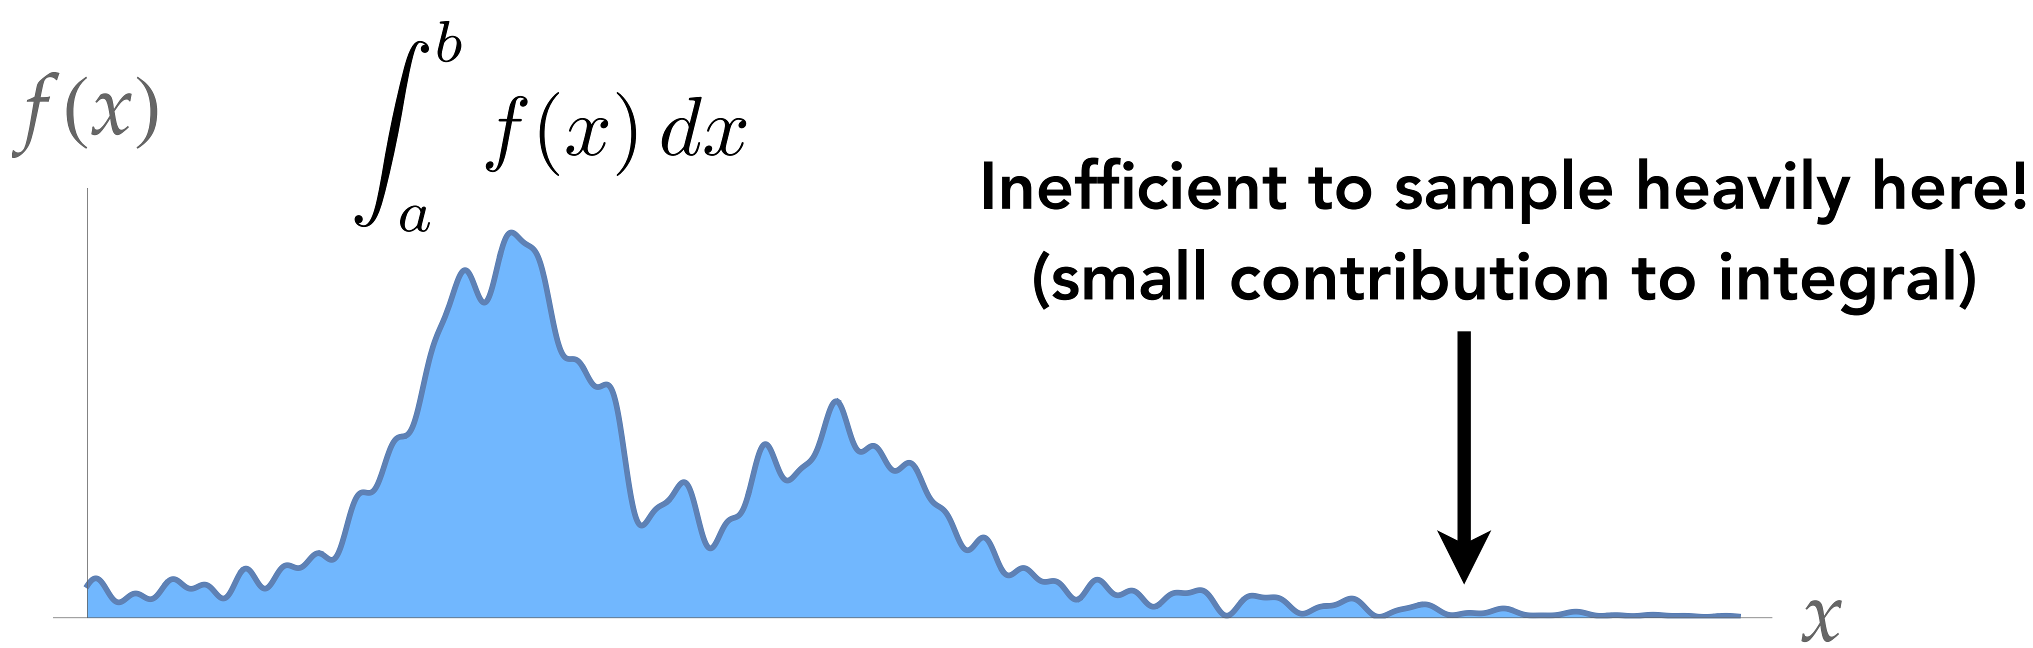

complicated integrand

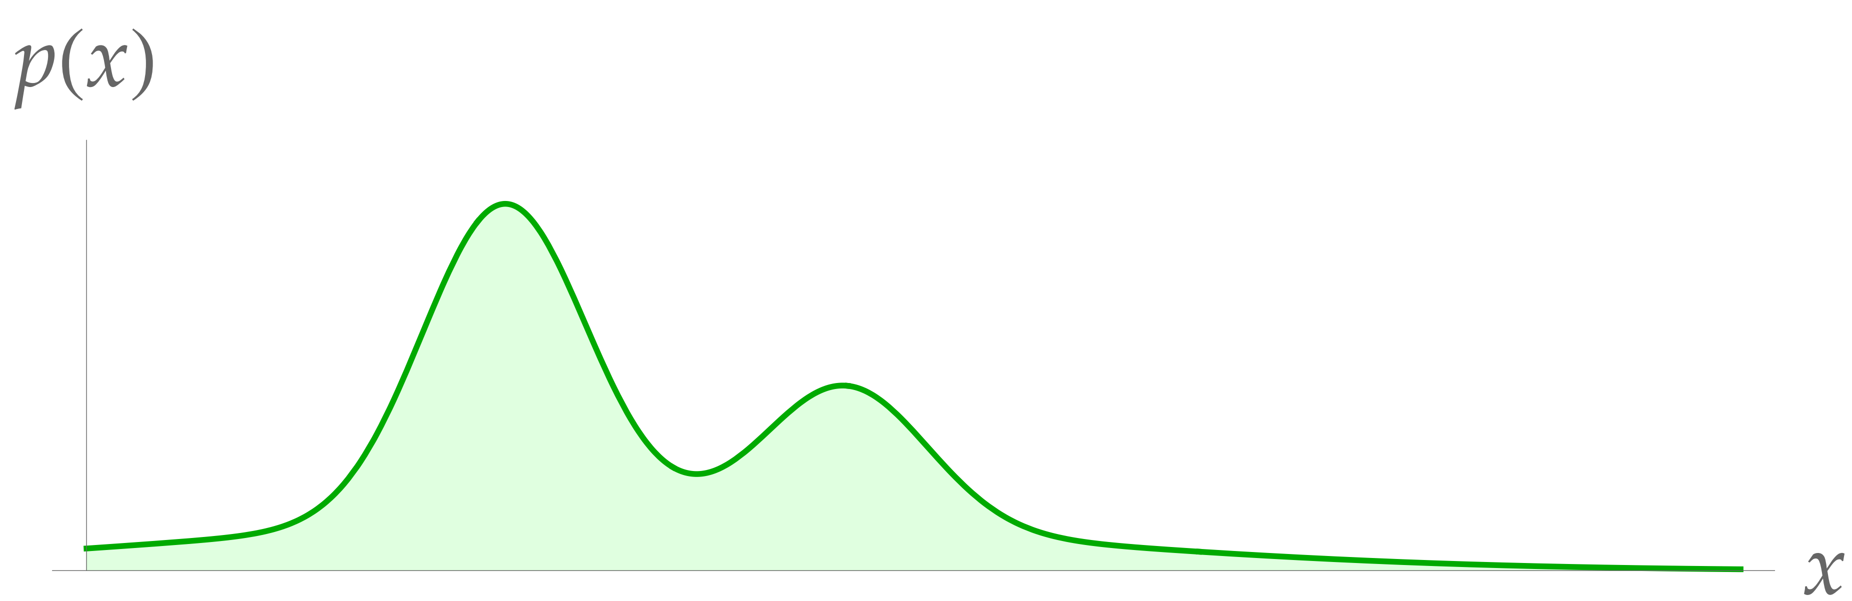

our best guess for where the integrand is "big"

Note: p(x) must be non-zero where f(x) is non-zero

Basic Monte Carlo:

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

(x_i are sampled *uniformly*)

Importance-Sampled Monte Carlo:

$$\frac{1}{n}\sum_{i=1}^{n} \frac{f(x_i)}{p(x_i)}$$

(x_i are sampled proportional to *p*)

"If I sample x less frequently, each sample should count for more."

Effect of Sampling Distribution "Fit"

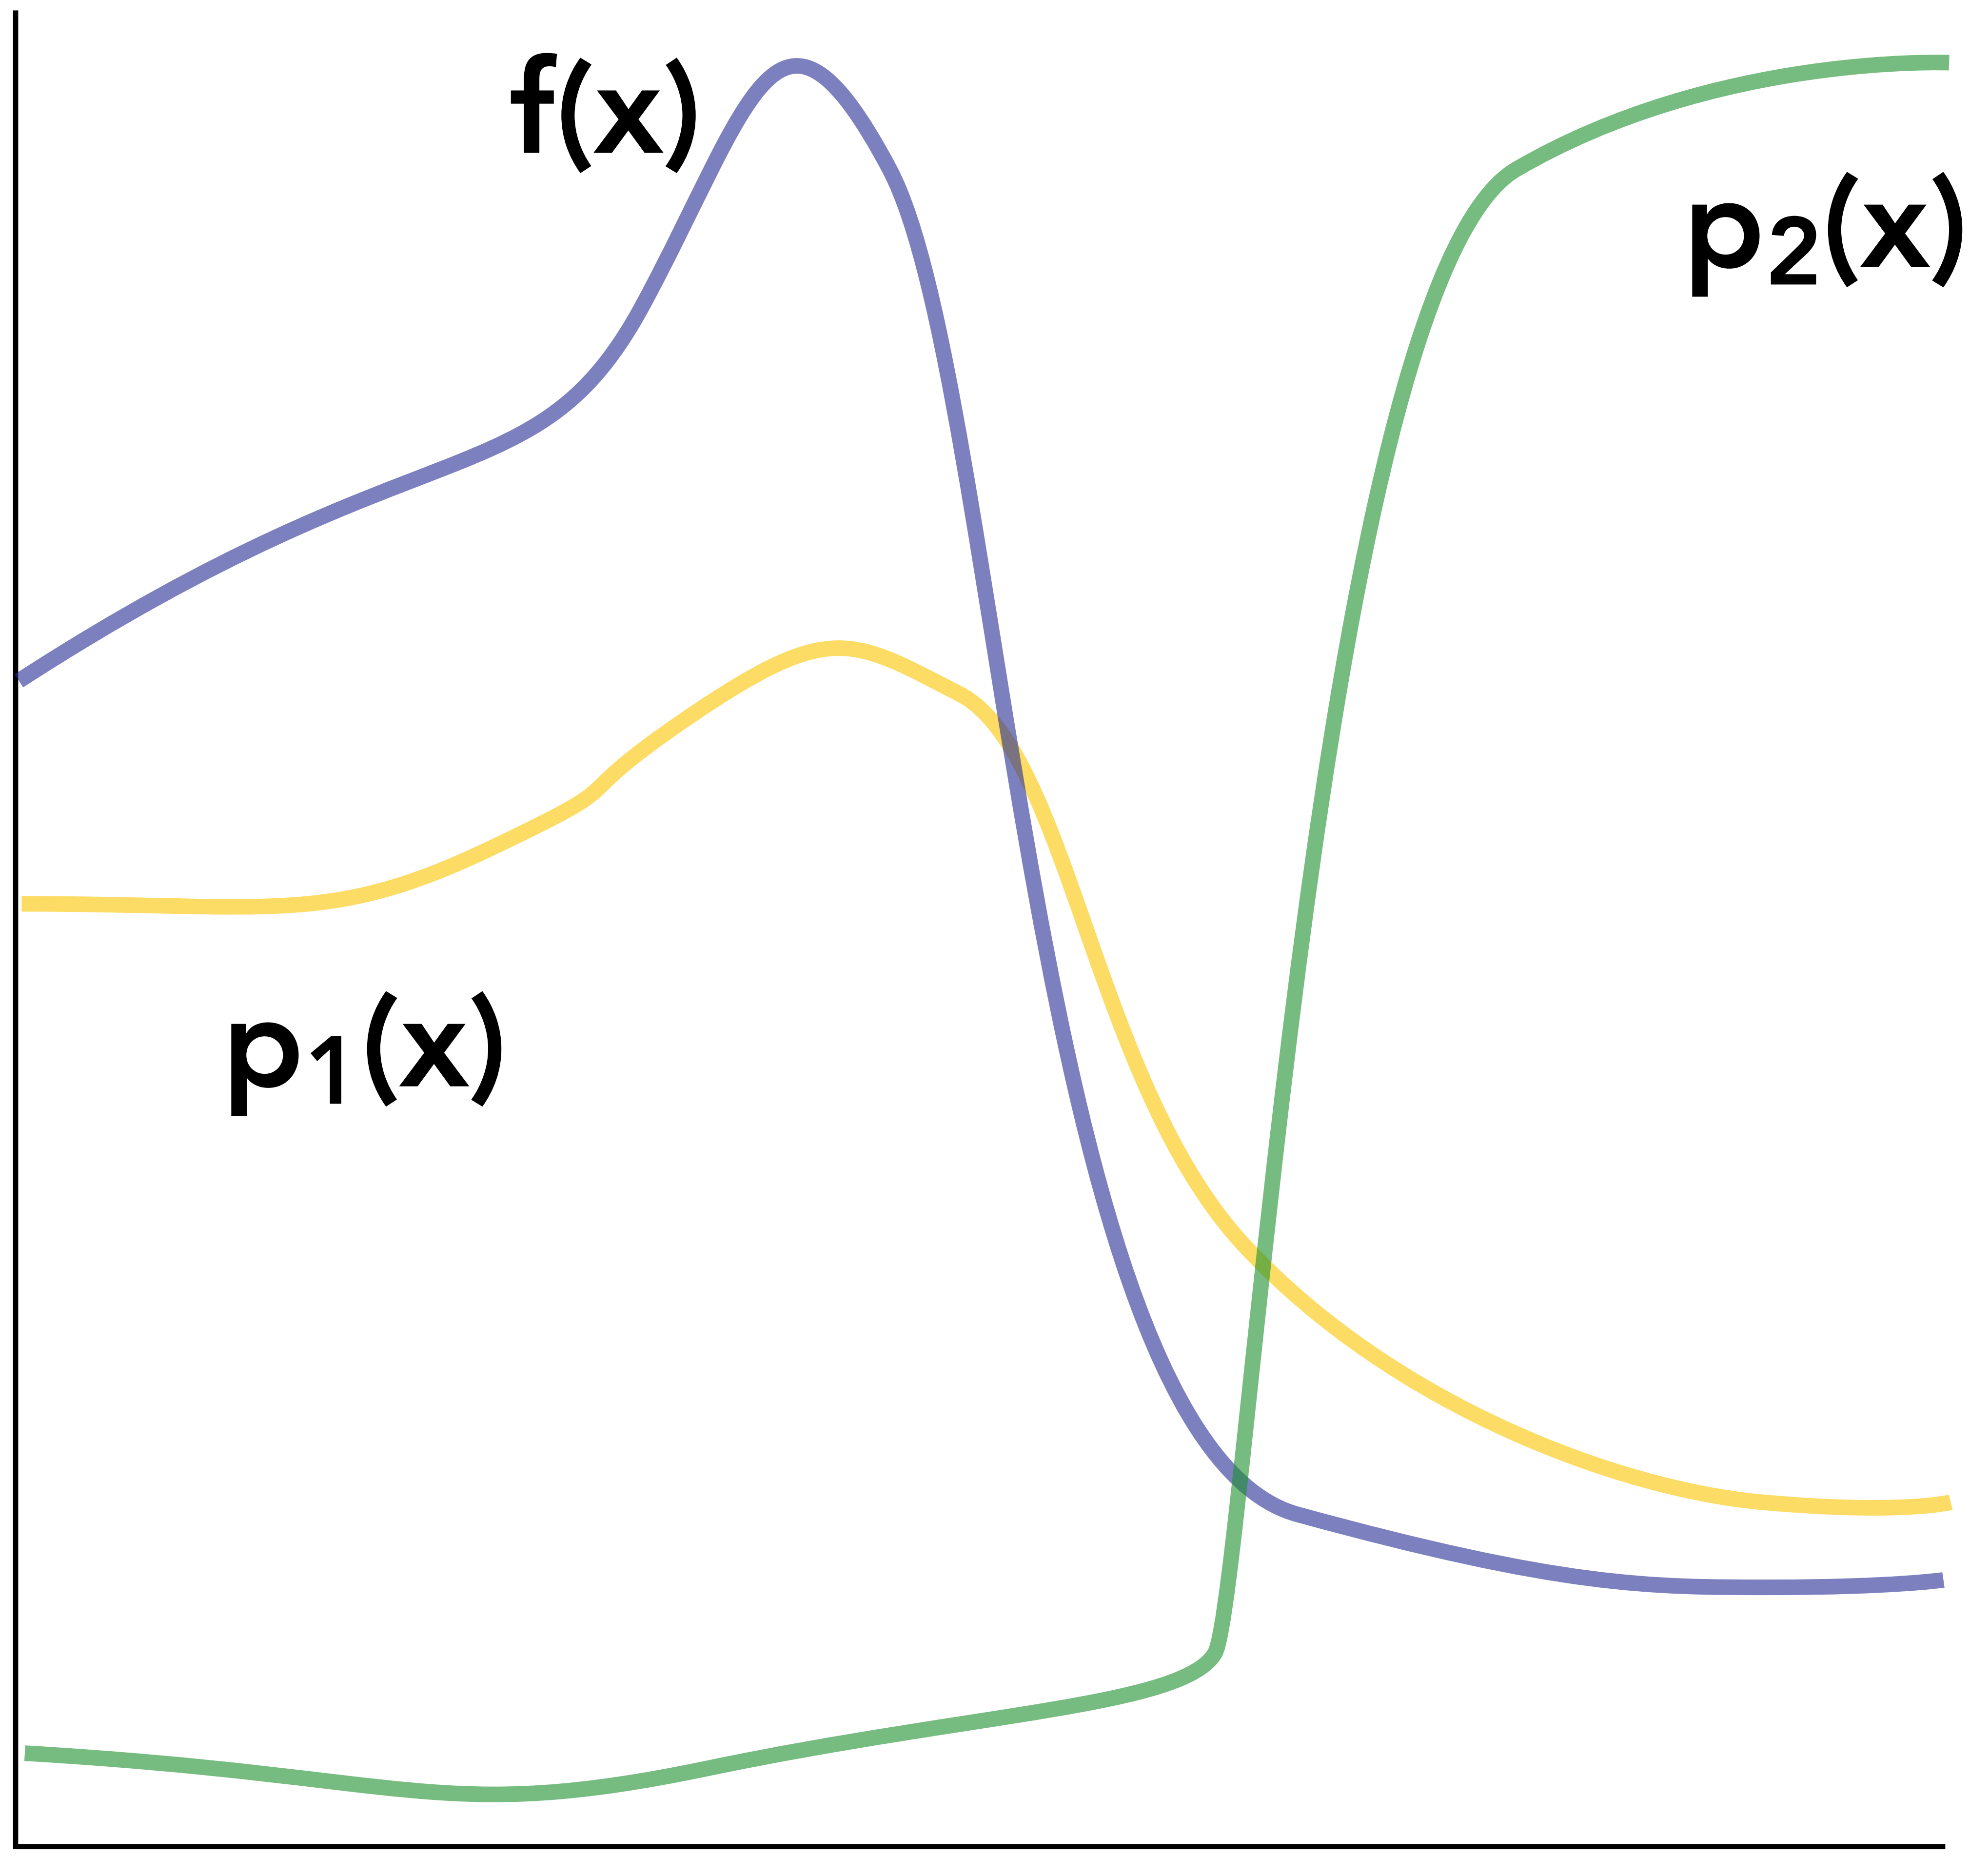

What is the behavior of $f(x)/p_1(x)$? $f(x)/p_2(x)$?

How does this impact the variance of the estimator?

Solid Angle Sampling vs Light Area Sampling

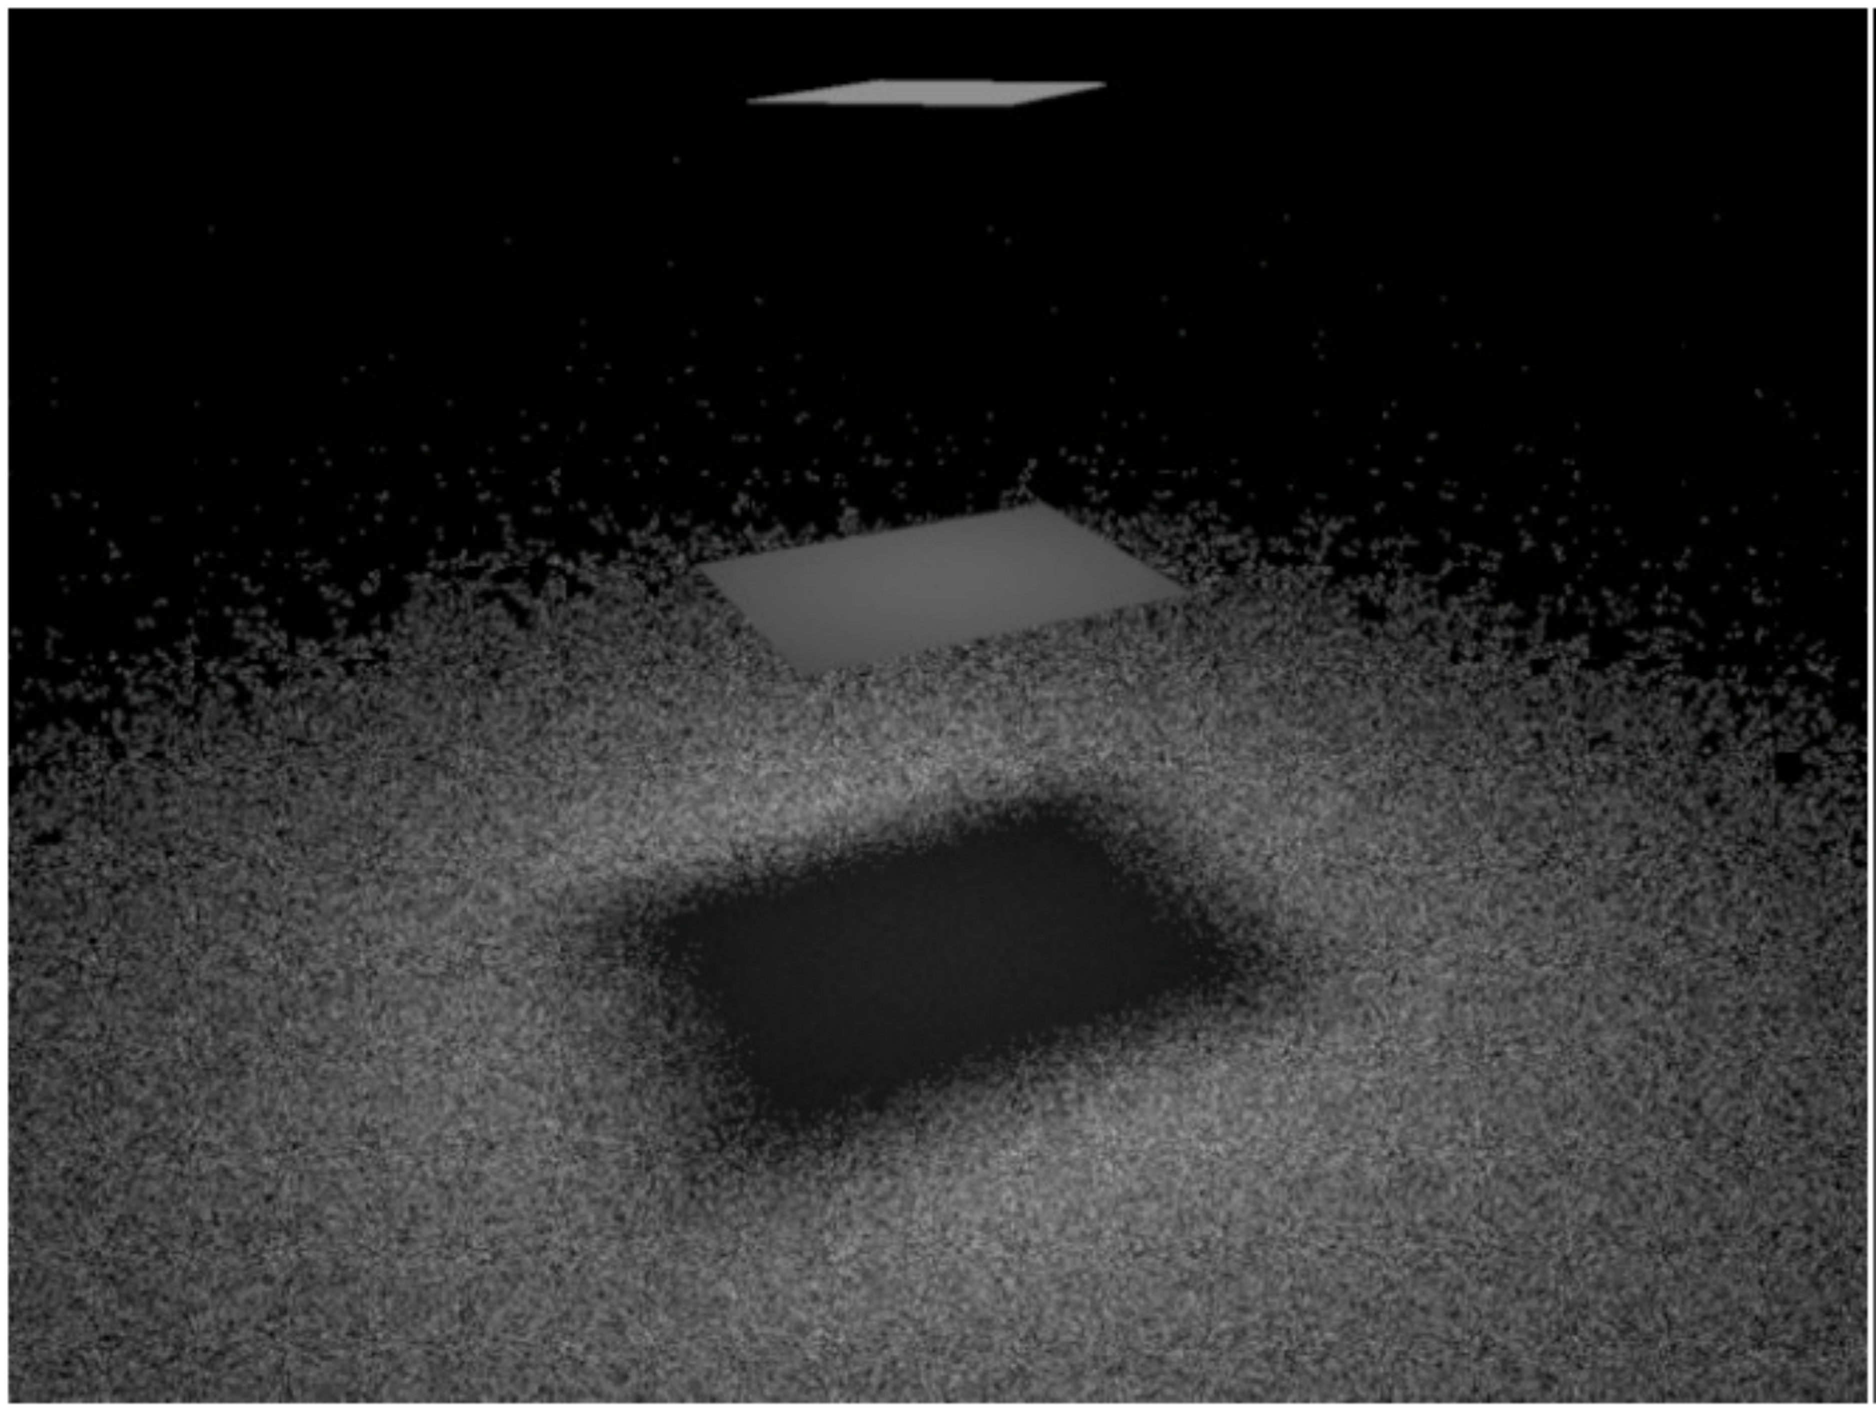

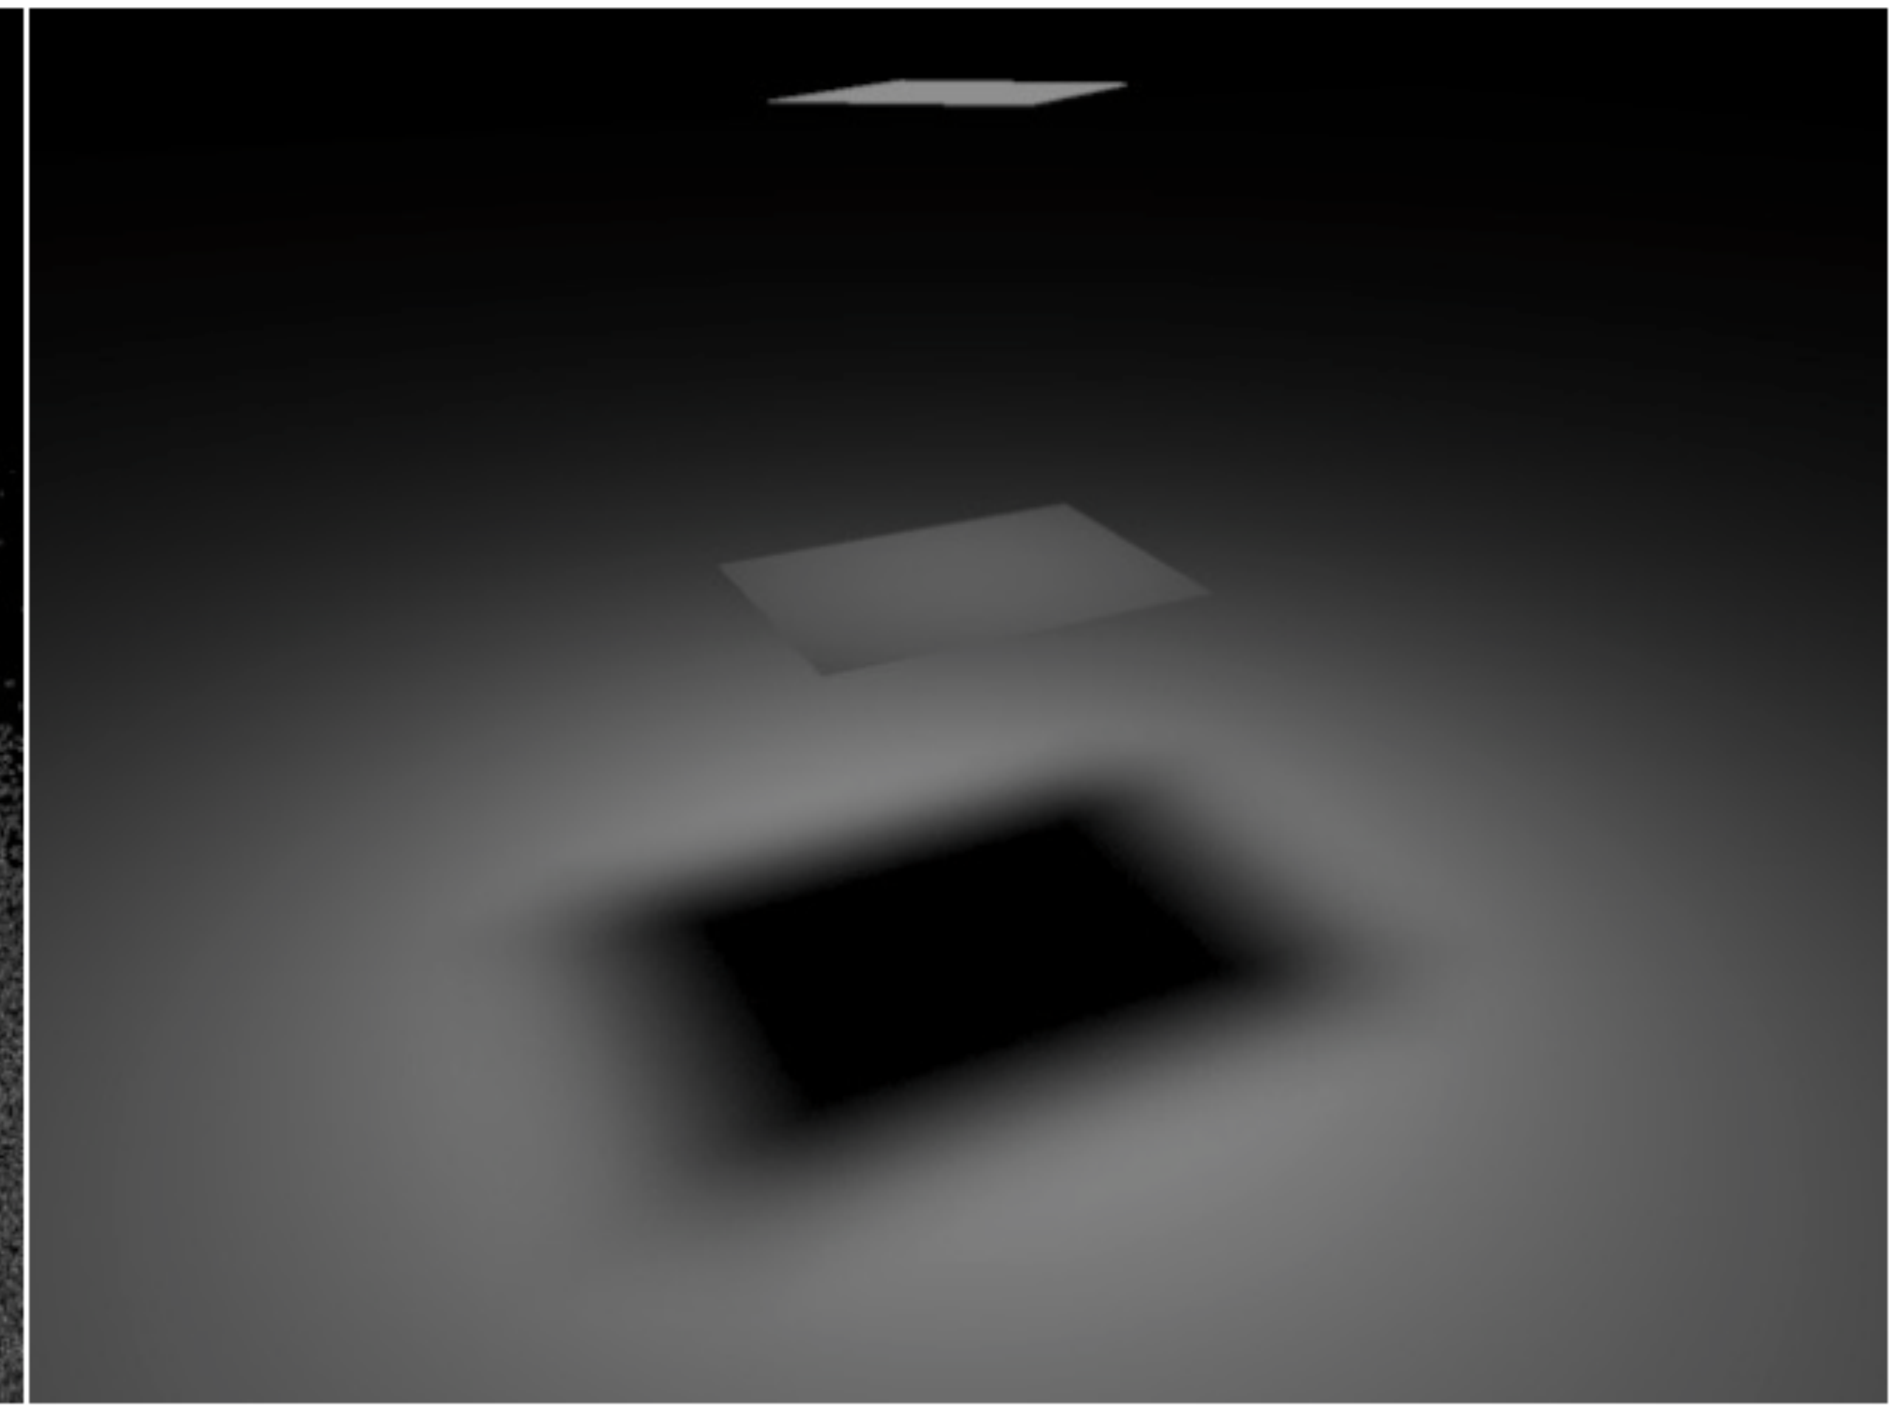

Sampling solid angle

100 random directions on hemisphere

Sampling light source area

100 random points on area of light source

Sampling Light Source Area
100 rays

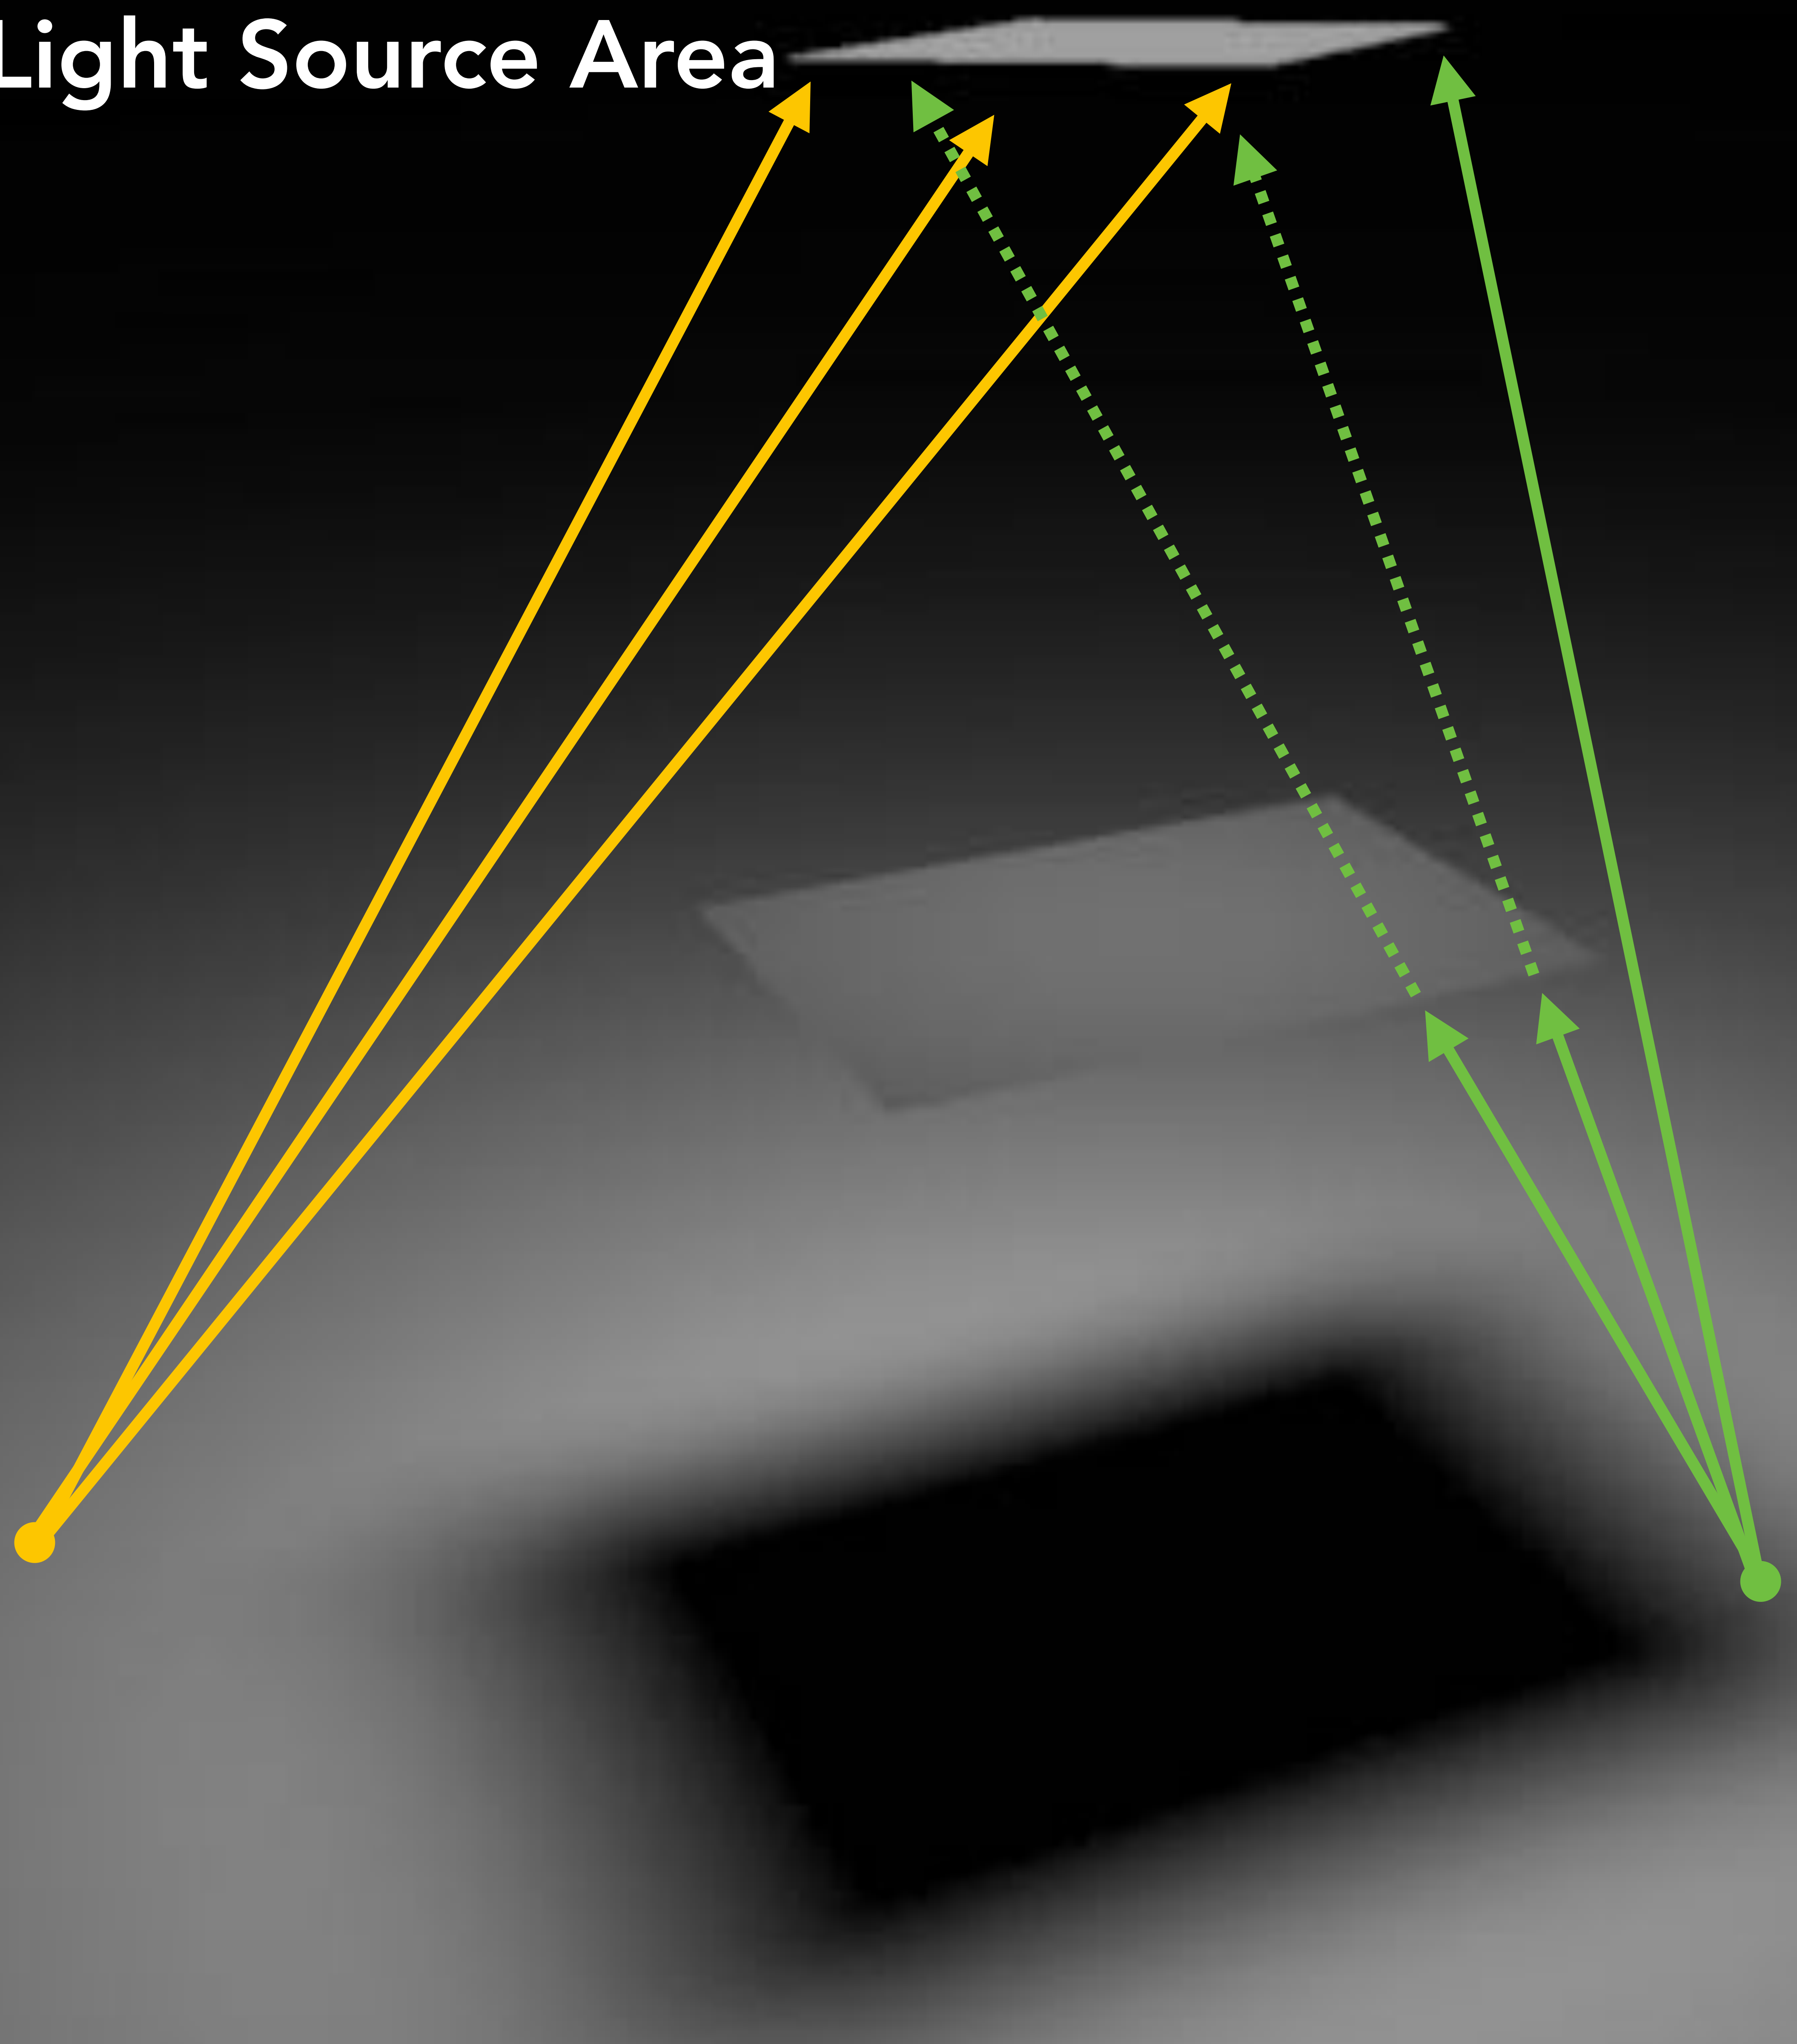

If no occlusion is present, all directions chosen in computing estimate “hit” the light source.

(Choice of direction only matters if portion of light is occluded from surface point *p*.)

Changing Basis of Integration: Sampling Hemisphere

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

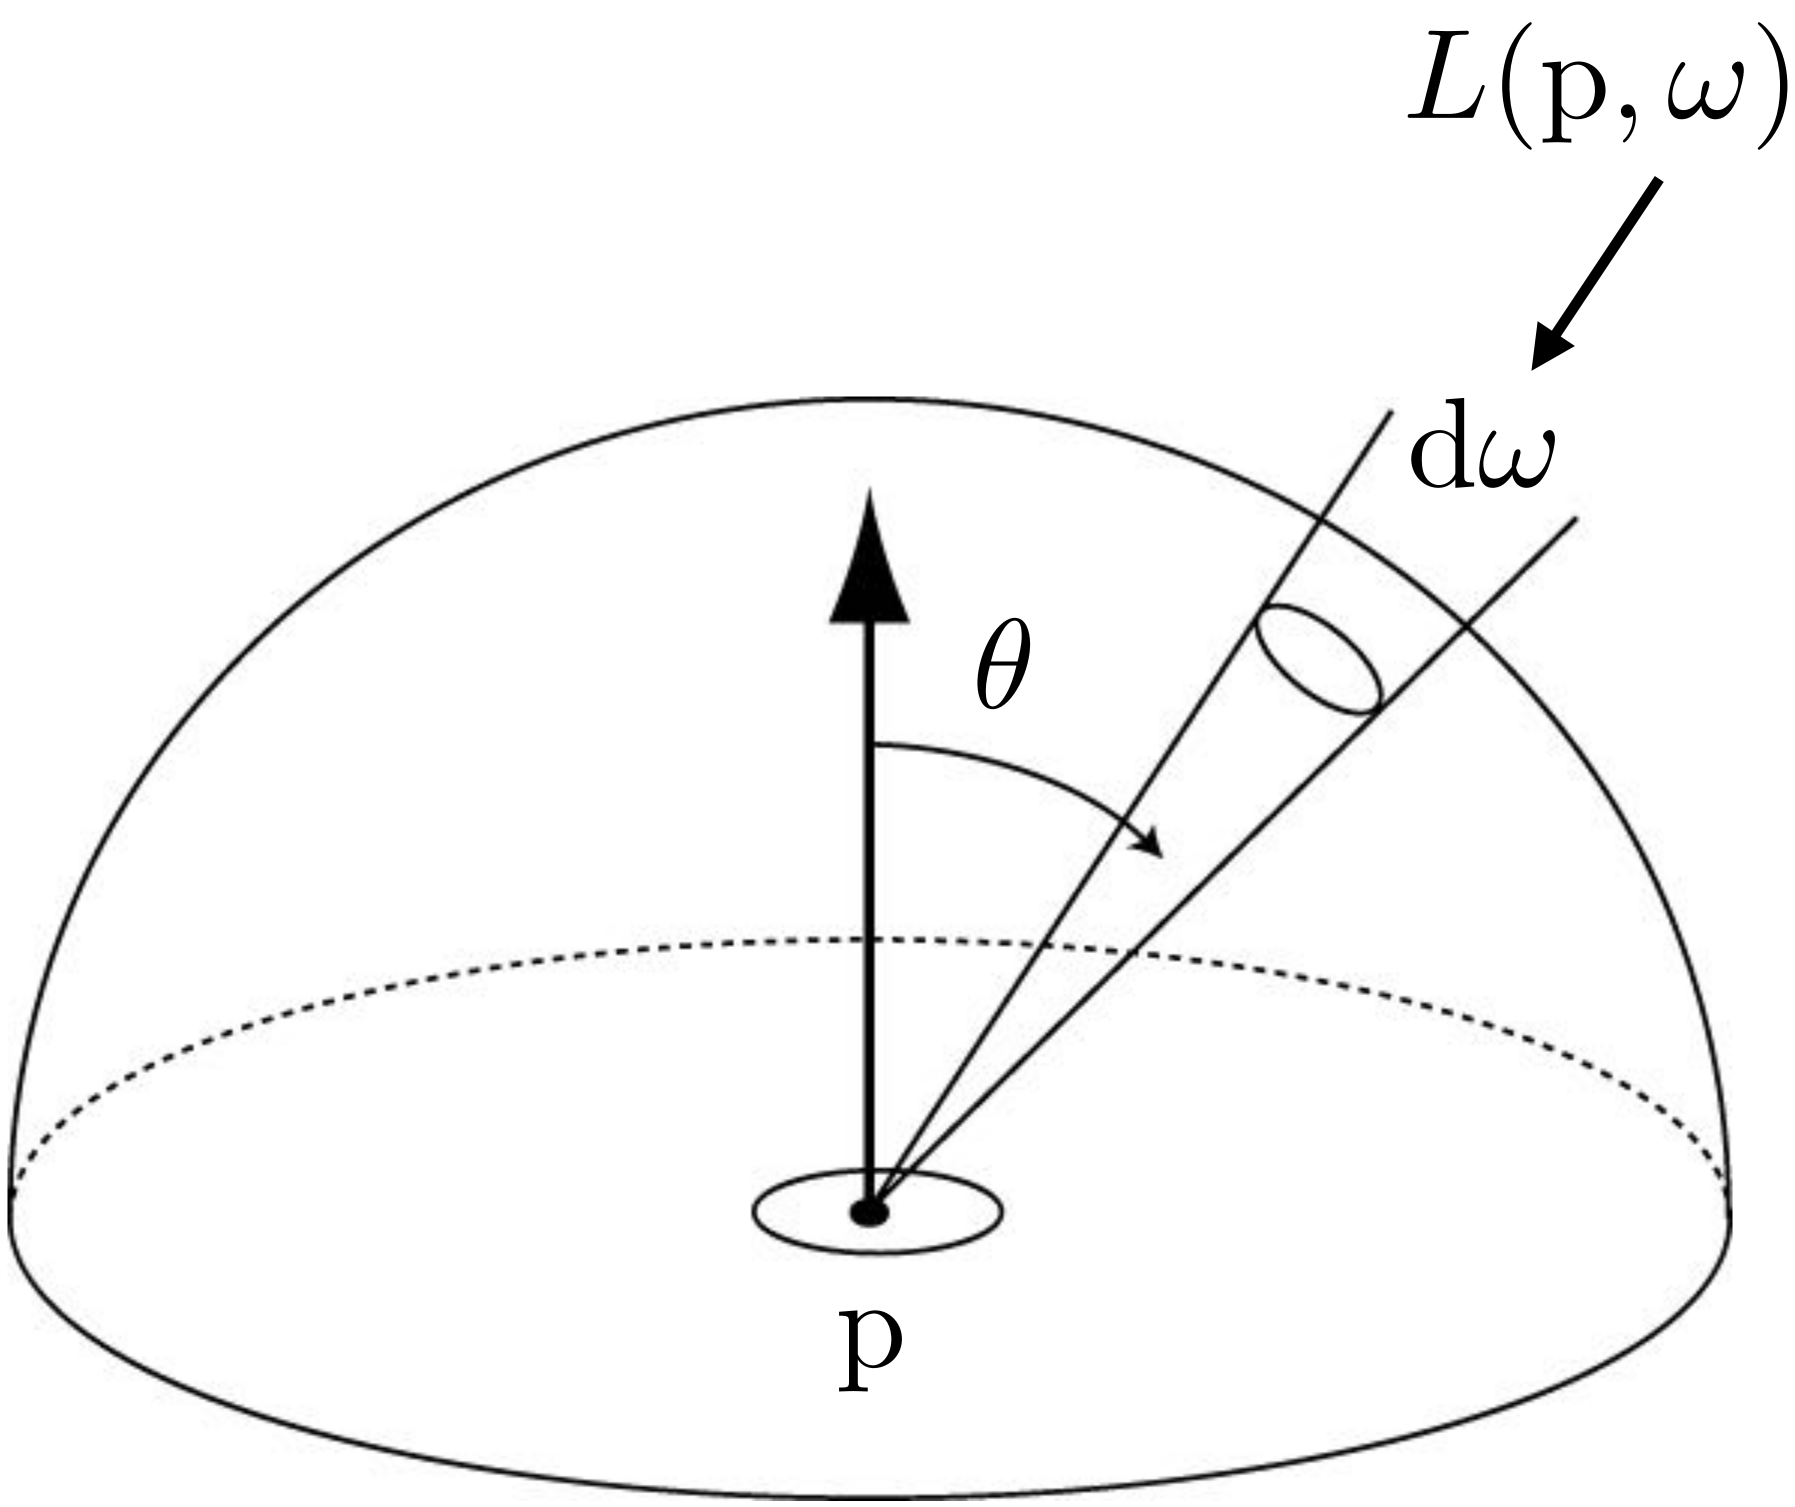

Changing Basis of Integration: Sampling Light Source Area

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

← Change of variables to integral over area of light *

$$dw = \frac{dA' cos\, \theta'}{|\mathrm{p}' - \mathrm{p}|^2}$$

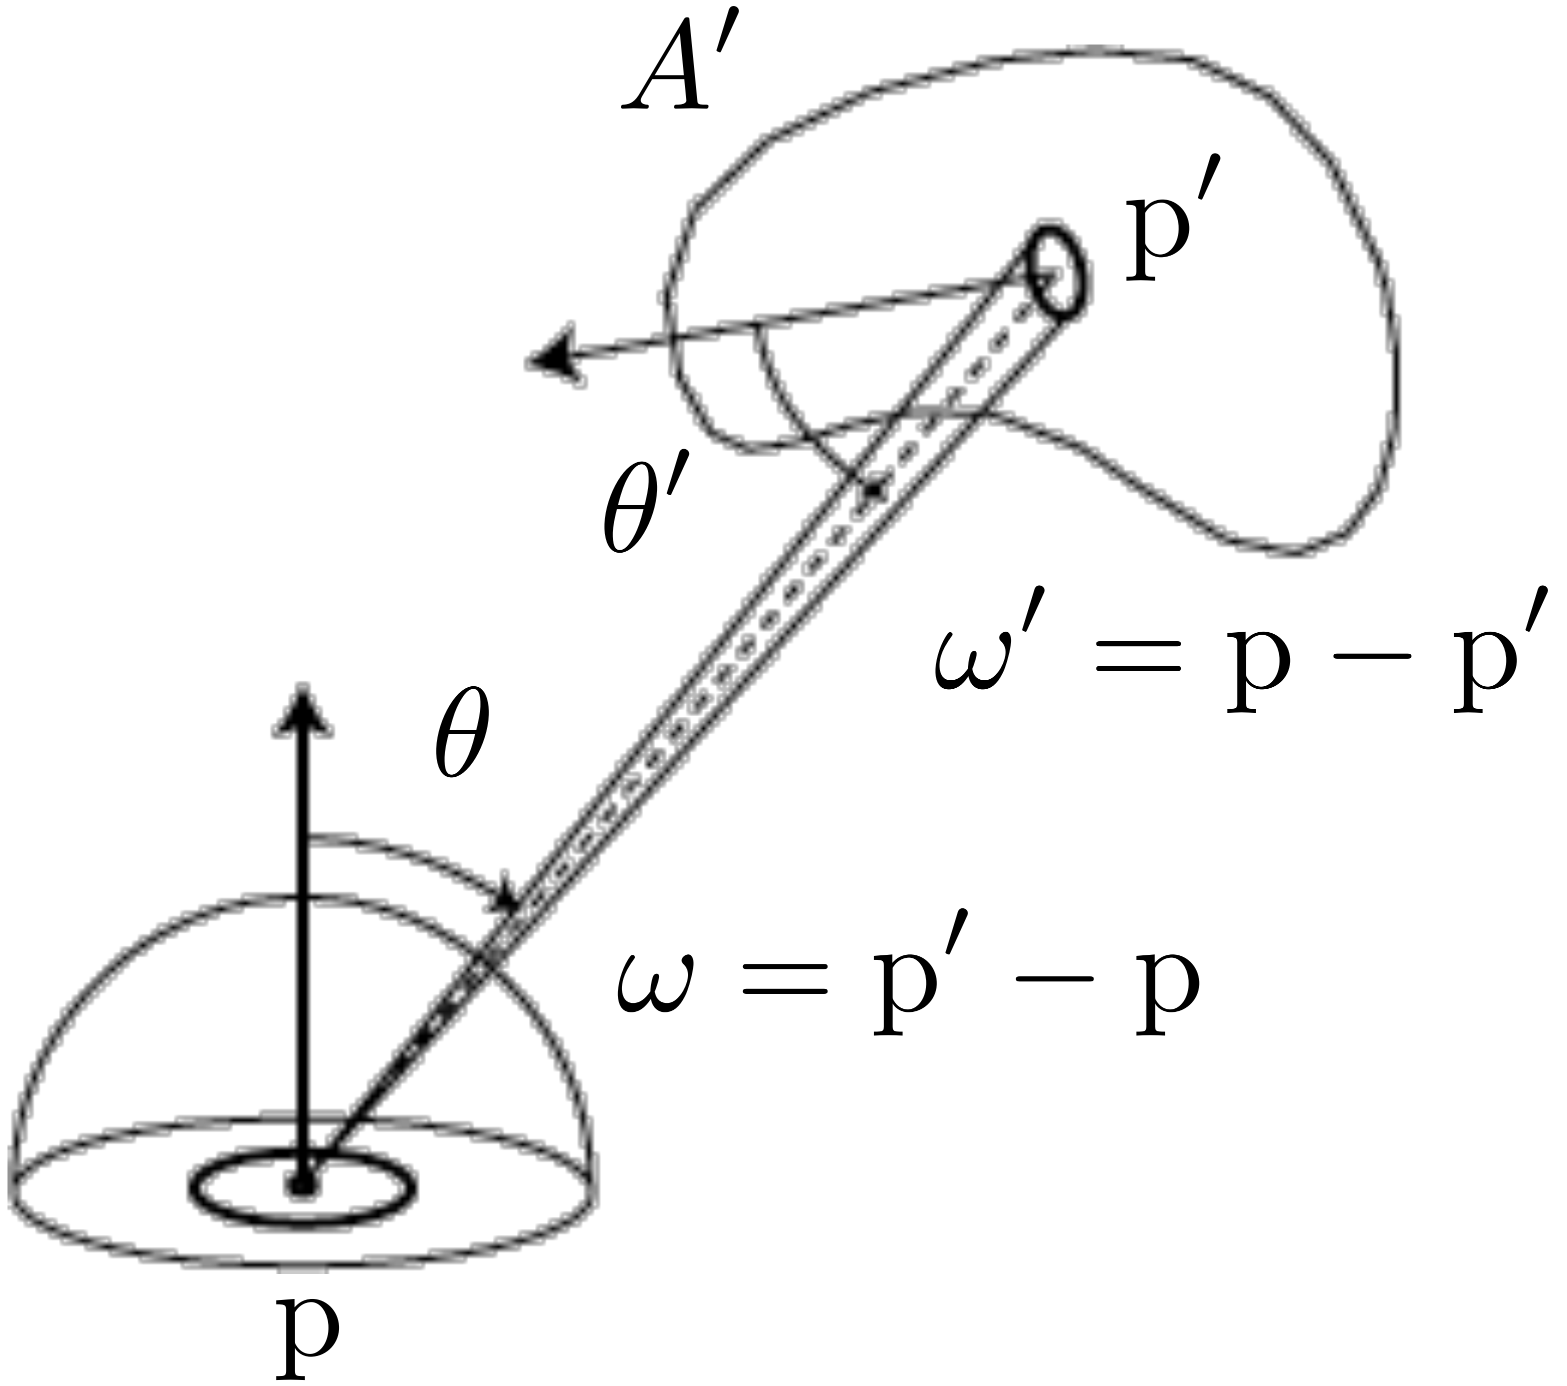

Binary visibility function:
1 if p' is visible from p, 0 otherwise
(accounts for light occlusion)

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

Monte Carlo Estimate by Sampling Light Source Area

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

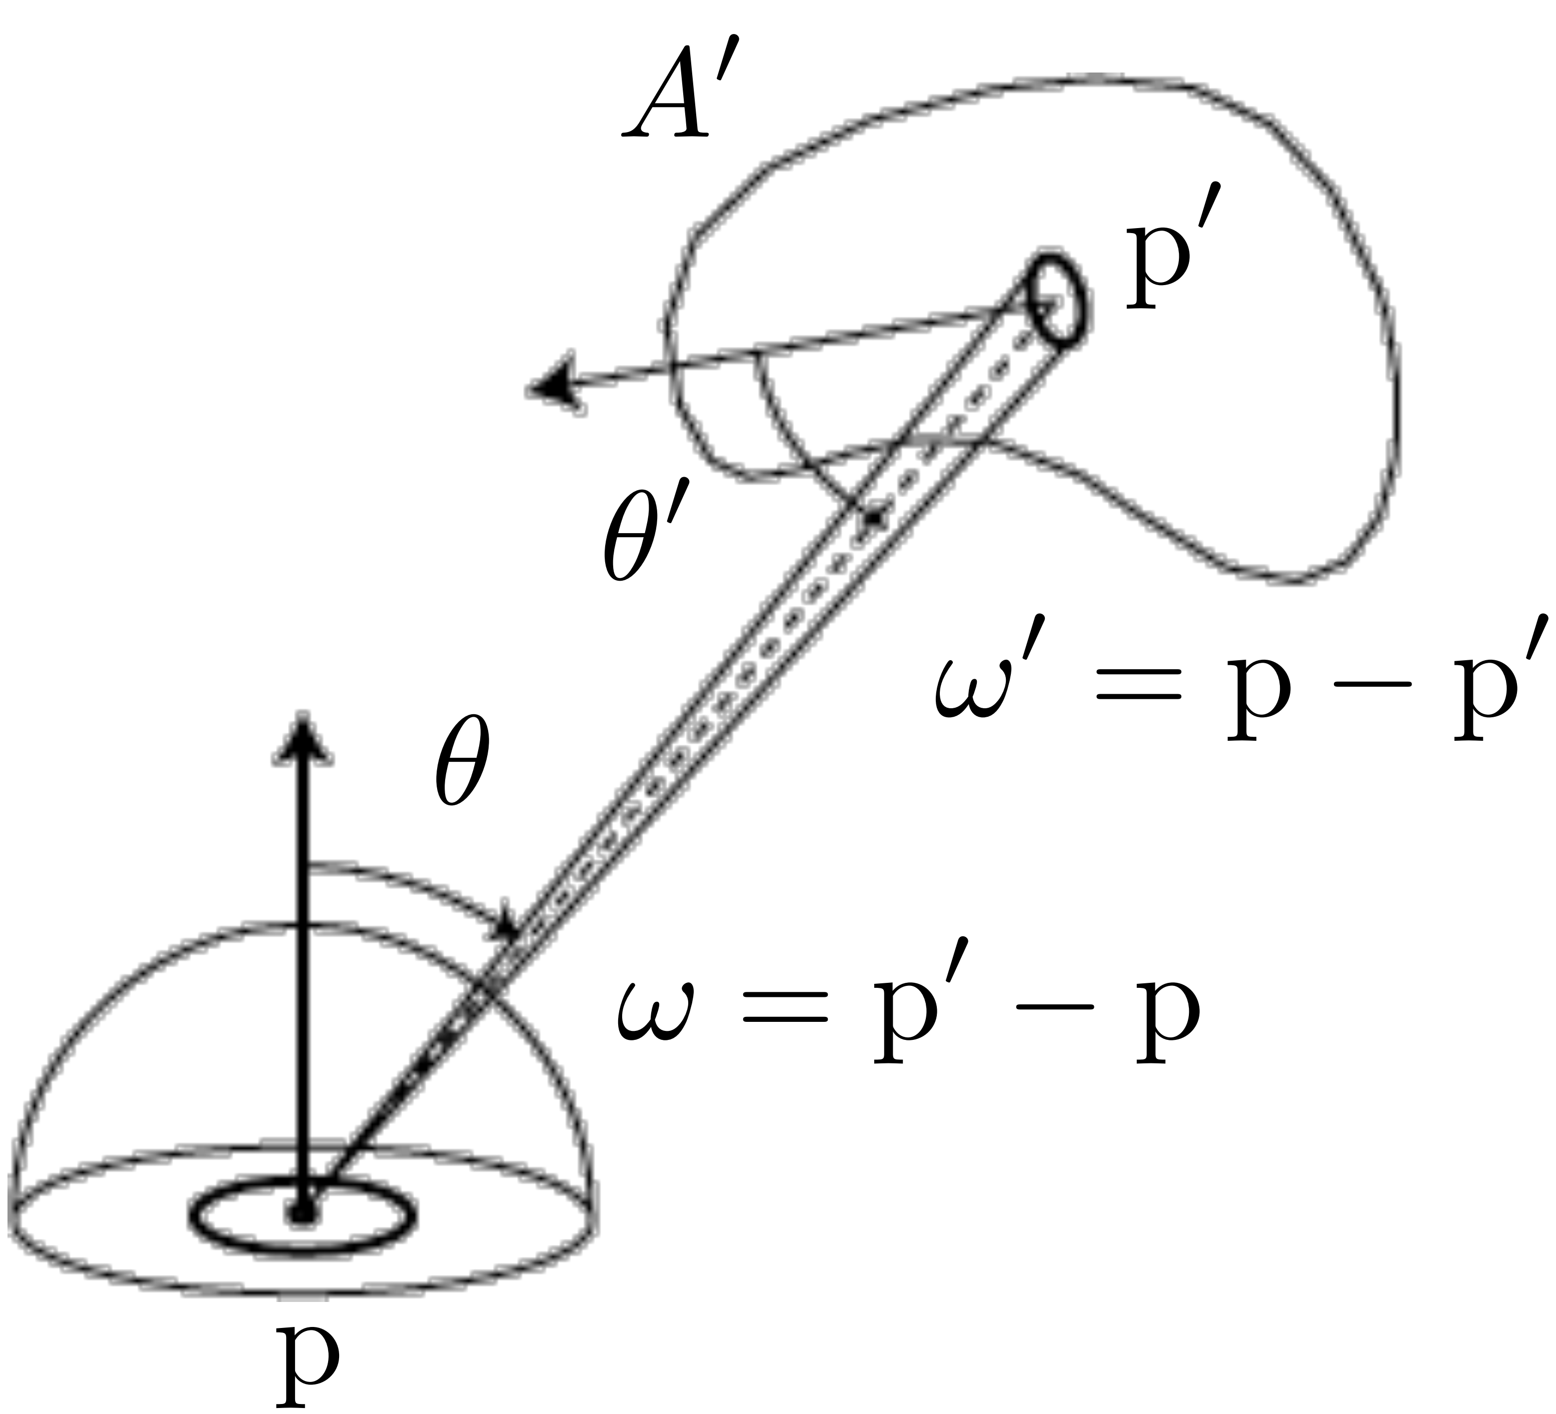

Randomly sample light source area A' (assume uniformly over area)

$$\int_{A'} p(\mathrm{p}')\, \mathrm{d}A' = 1$$

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

Monte Carlo Estimator

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

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

Solid Angle Sampling vs Light Area Sampling

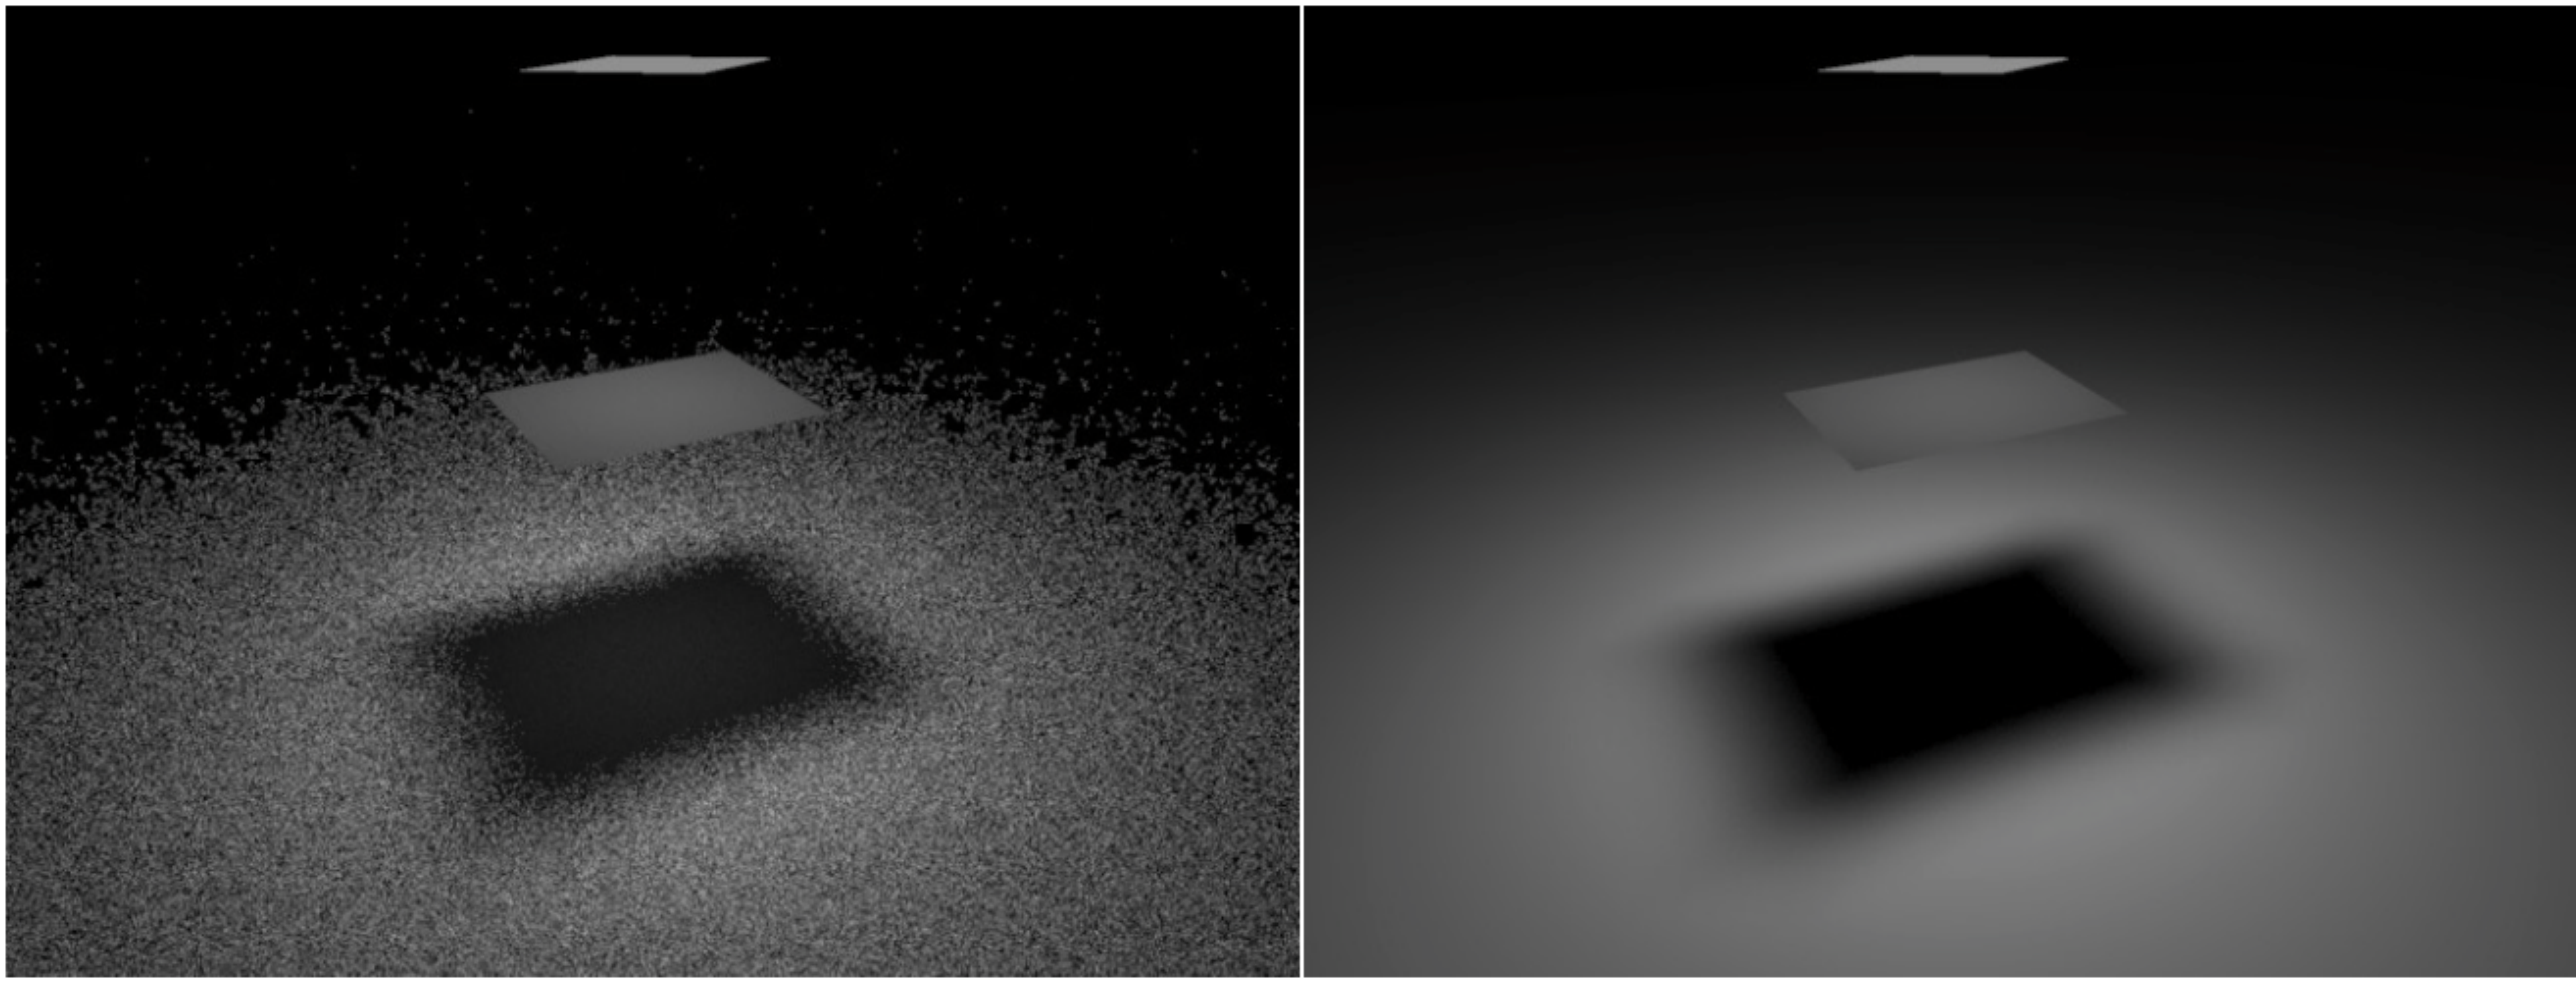

Sampling solid angle

100 random directions on hemisphere

Sampling light source area

100 random points on area of light source

How to Draw Samples From a Desired Probability Distribution? One Approach: Inversion Method

Task: Draw A Random Value From a Given PDF

Task:

Given a PDF for a discrete random variable, probability p_i for each value x_i,

Draw a random value X from this PDF.

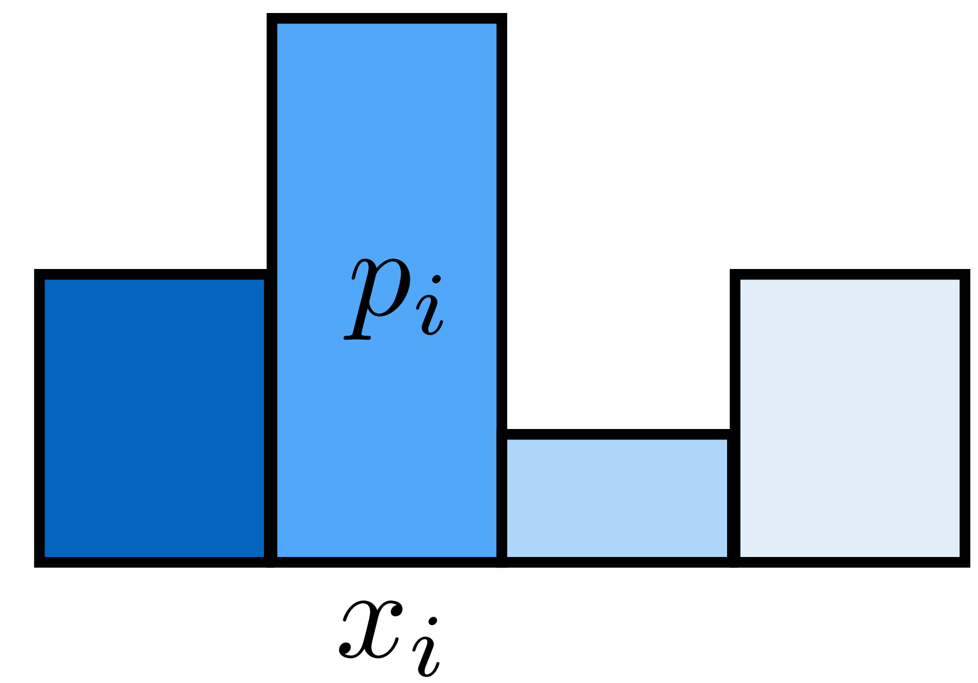

Step 1:

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

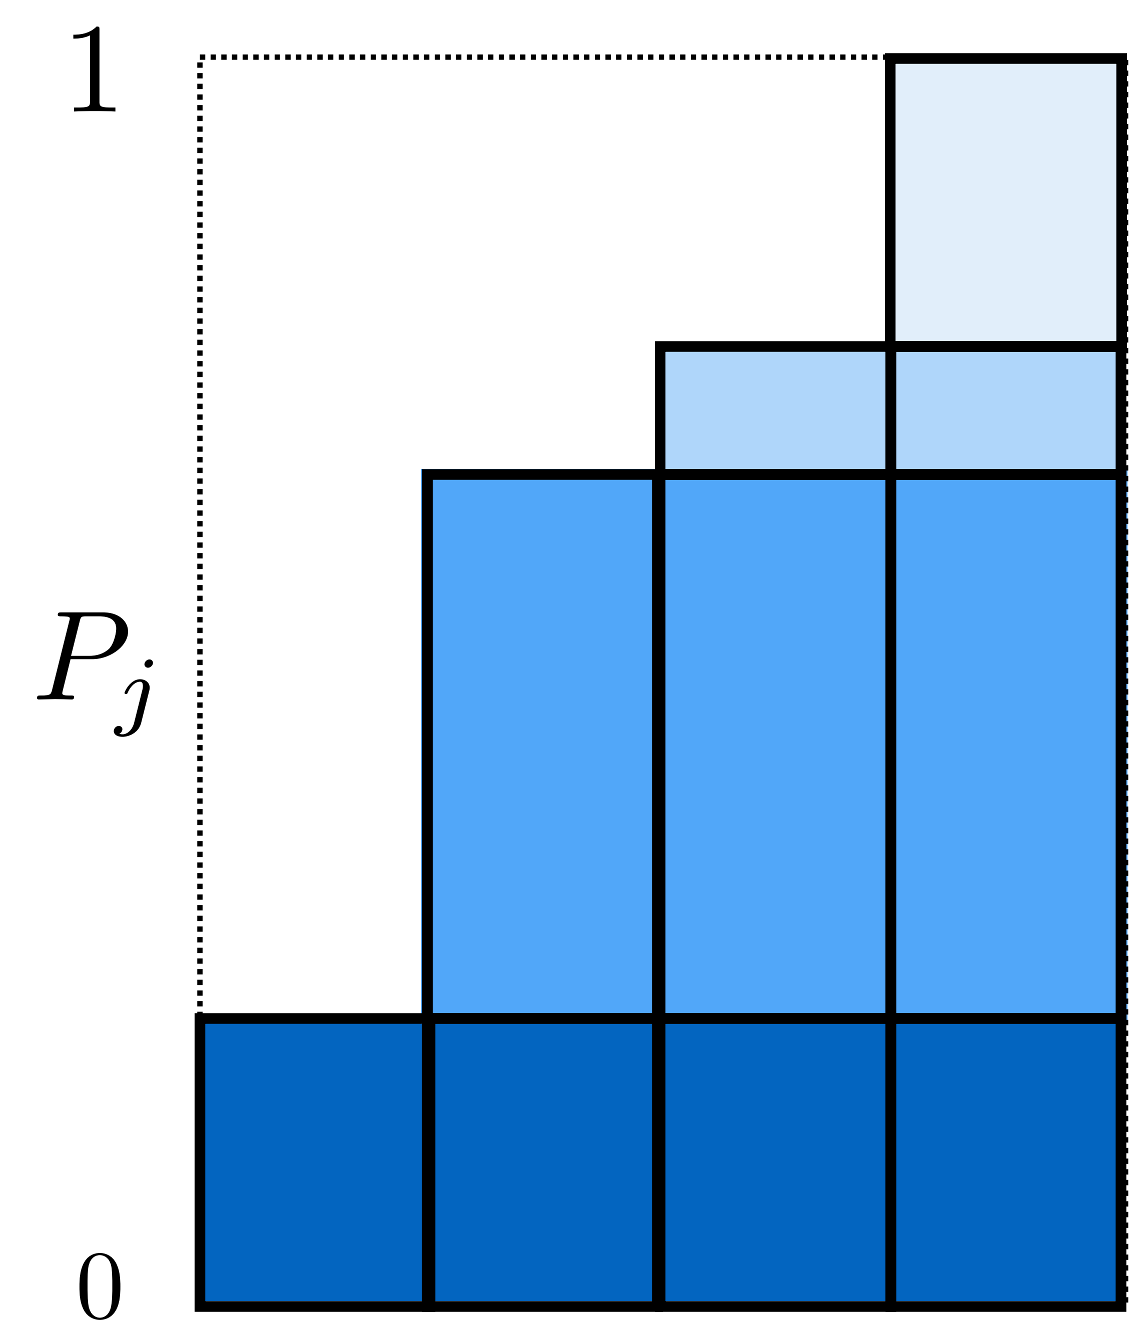

Note: must have

$$0 \leq P_i \leq 1$$

$$P_n = 1$$

Task: Draw A Random Value From a Given PDF

Task:

Given a PDF for a discrete random variable, probability p_i for each value x_i,

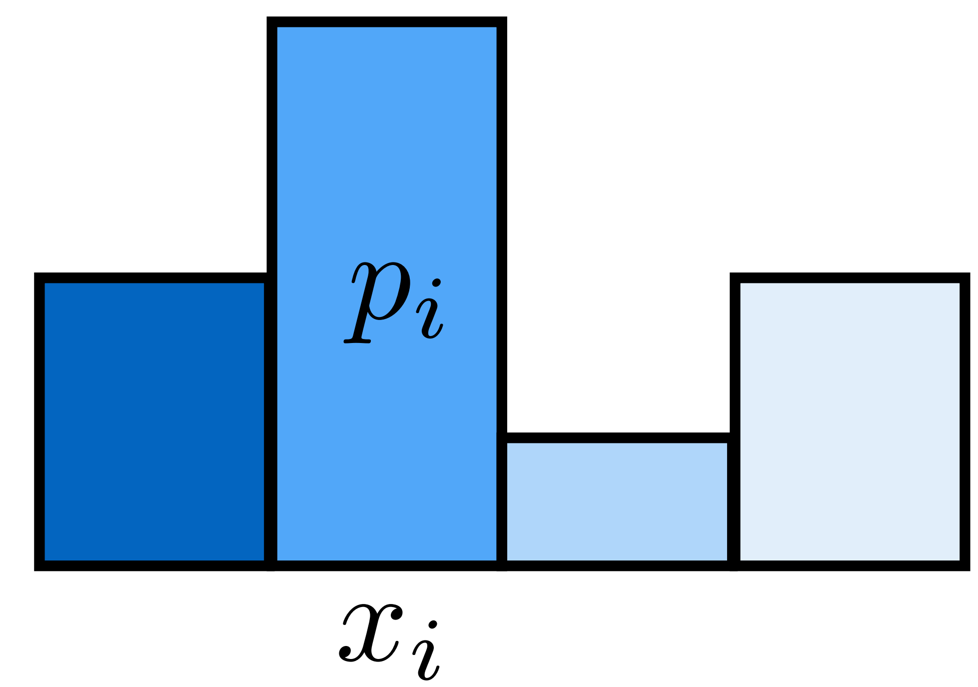

Draw a random value X from this PDF.

Step 2:

Given a uniform random variable $\xi \in [0, 1)$

choose $X = x_i$
such that $P_{i-1} < \xi \leq P_i$

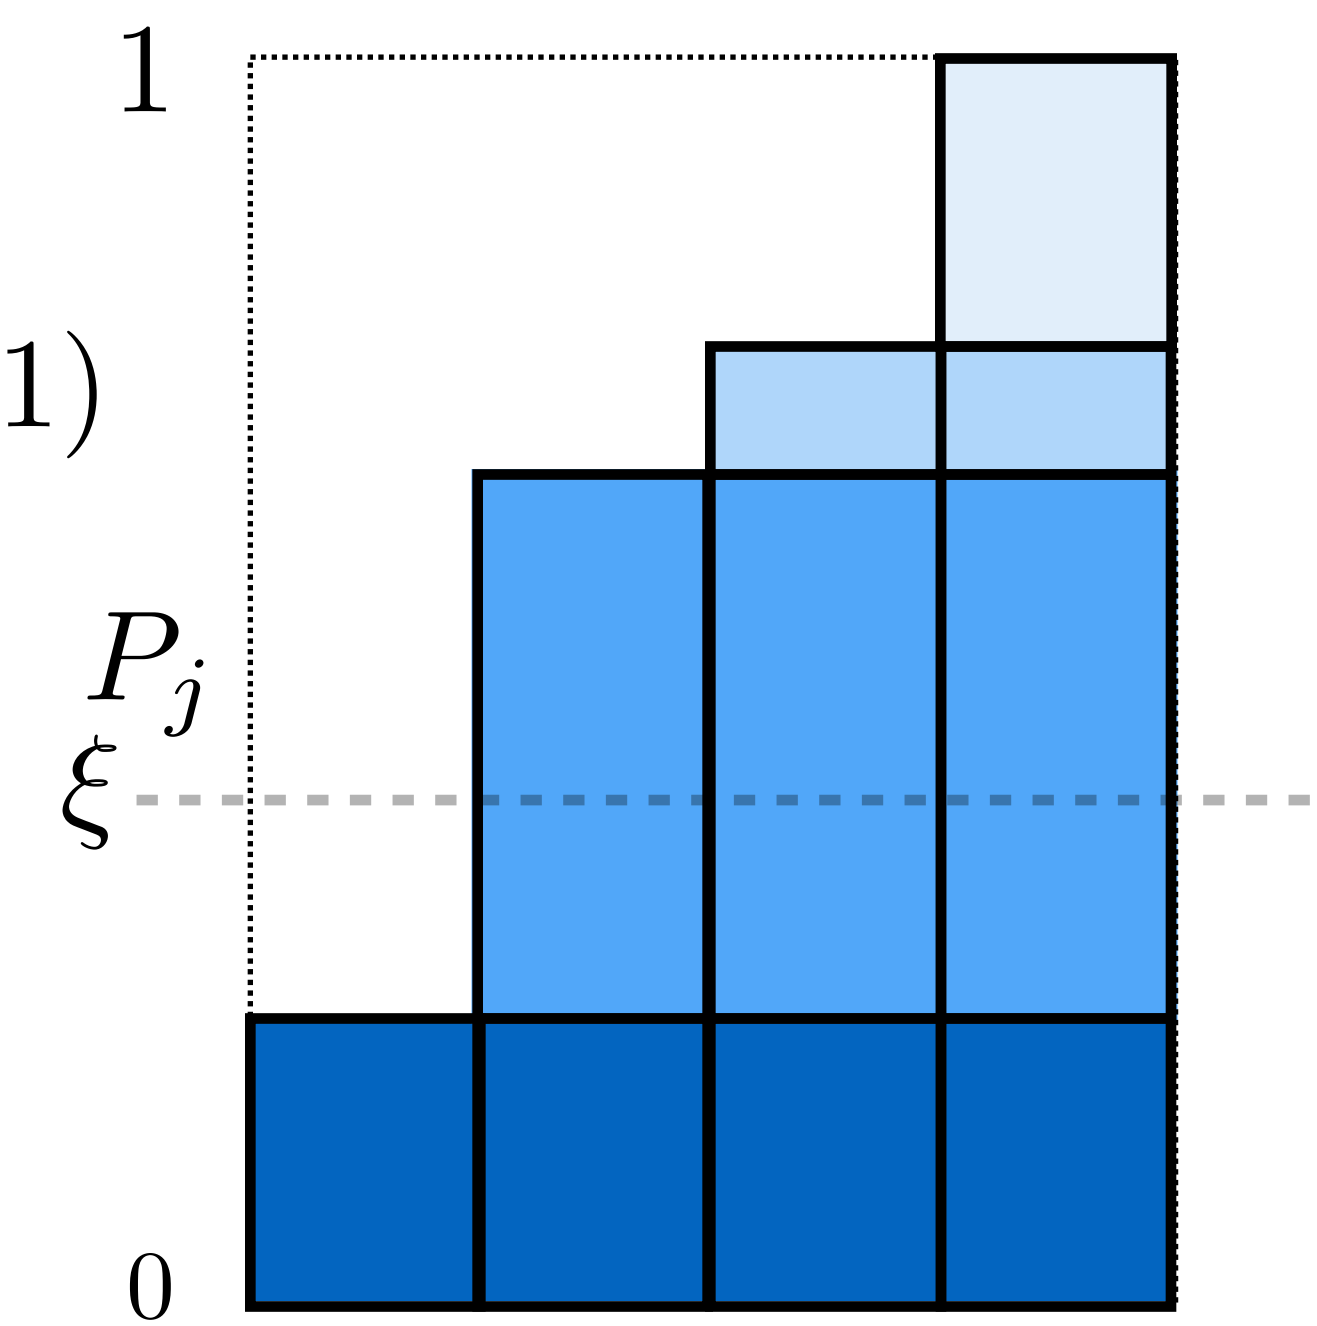

How to compute?   Binary search.

Continuous Probability Distribution

PDF $p(x)$

$$p(x) \geq 0$$

CDF $P(x)$

$$P(x) = \int_0^x p(x)\,\mathrm{d}x$$

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

$$P(1) = 1$$

$$\Pr(a \leq X \leq b) = \int_a^b p(x)\,\mathrm{d}x$$

$$= P(b) - P(a)$$

Uniform distribution on unit interval

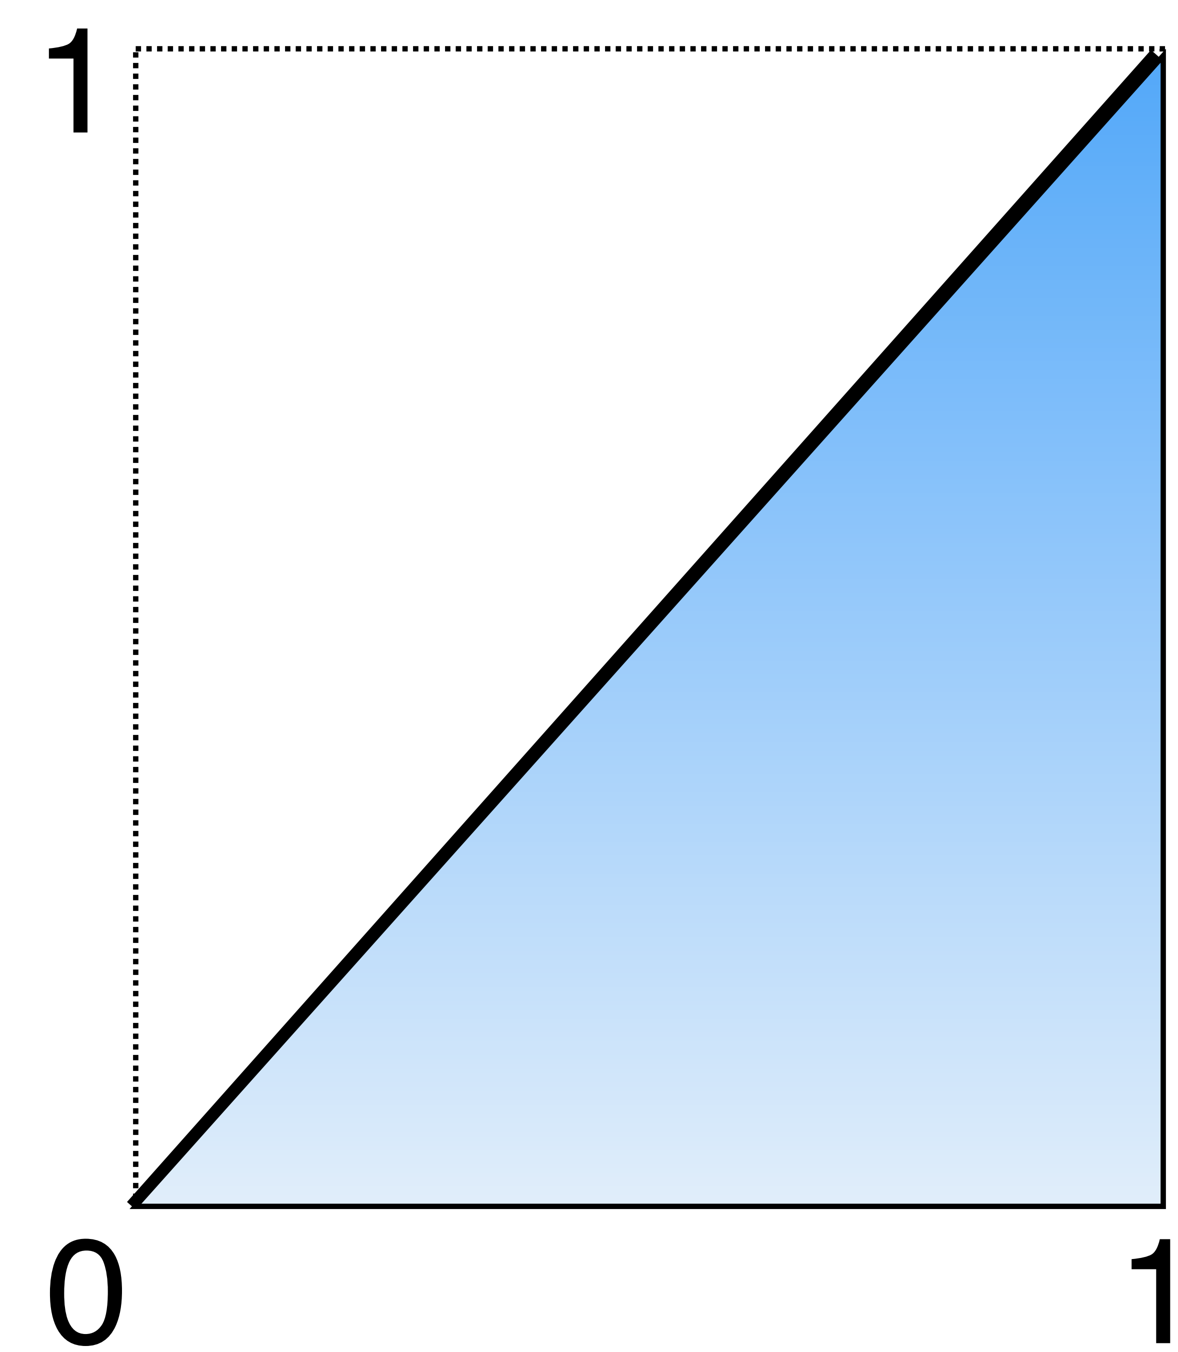

Sampling Continuous Probability Distributions

Called the "inversion method"

Cumulative probability distribution function

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

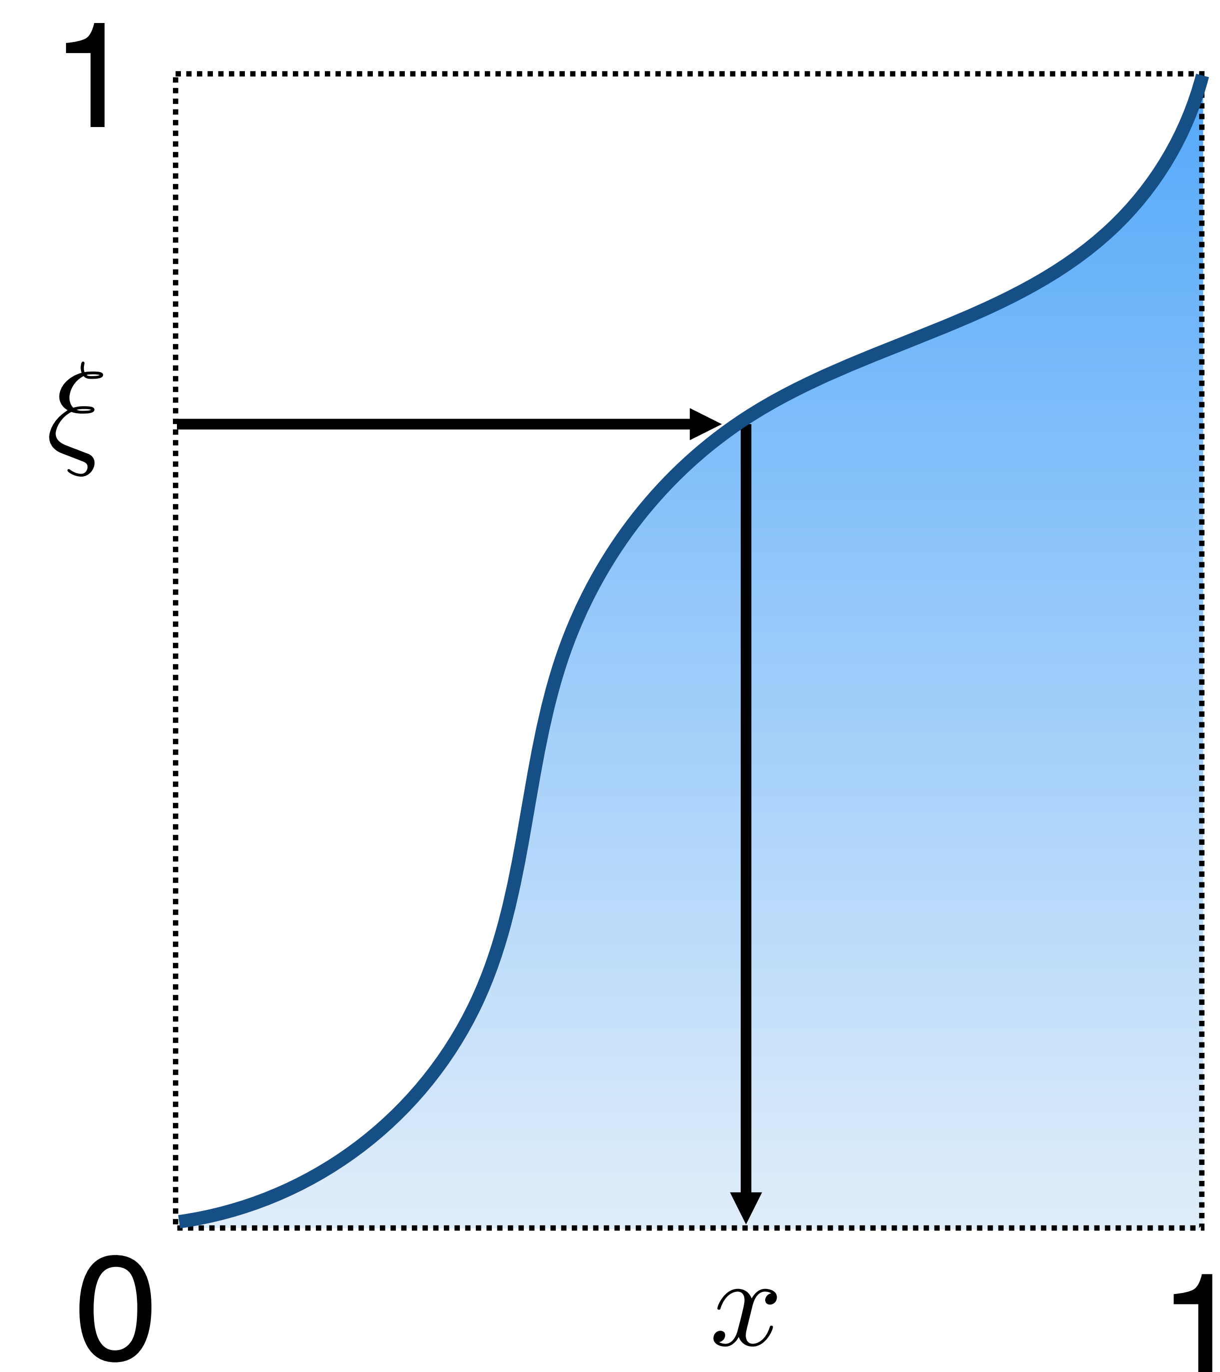

Construction of samples:

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: Sample Proportional to x^2

Given:

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

Want to sample according to this graph:

Example: Sample Proportional to x^2

Given:

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

Want to sample according to this graph:

Step 0: compute PDF by normalizing

$$p(x) = c\, f(x) = c\, x^2$$

$$\text{Also } 1 = \int_0^2 p(x)\, dx = \int_0^2 c\, x^2\, dx = \left. \frac{cx^3}{3} \right|_0^2 = \frac{8c}{3}$$

$$\implies c = \frac{3}{8}$$

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

Example: Sample Proportional to x^2

Given:

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

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

Want to sample according to this

Step 1: Compute CDF:

$$P(x) = \int_0^x p(x)\,\mathrm{d}x$$

$$= \frac{x^3}{8}$$

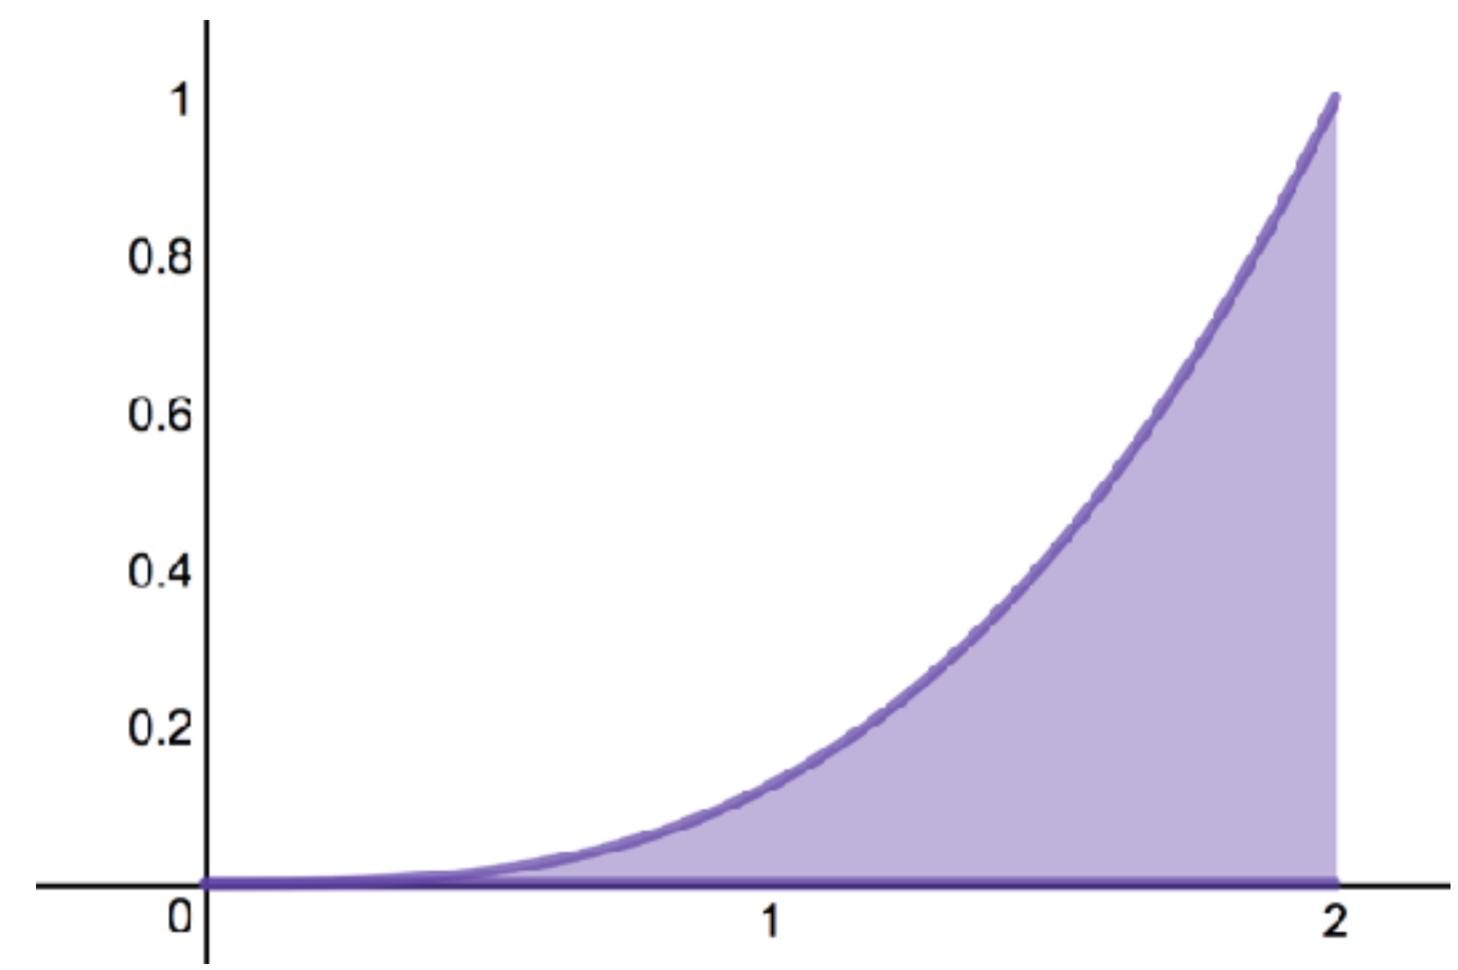

Example: Sample Proportional to x²

Given:

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

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

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

Step 2: Sample from $p(x)$

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

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

Applying the inversion method

Remember ξ is uniform random number in [0,1)

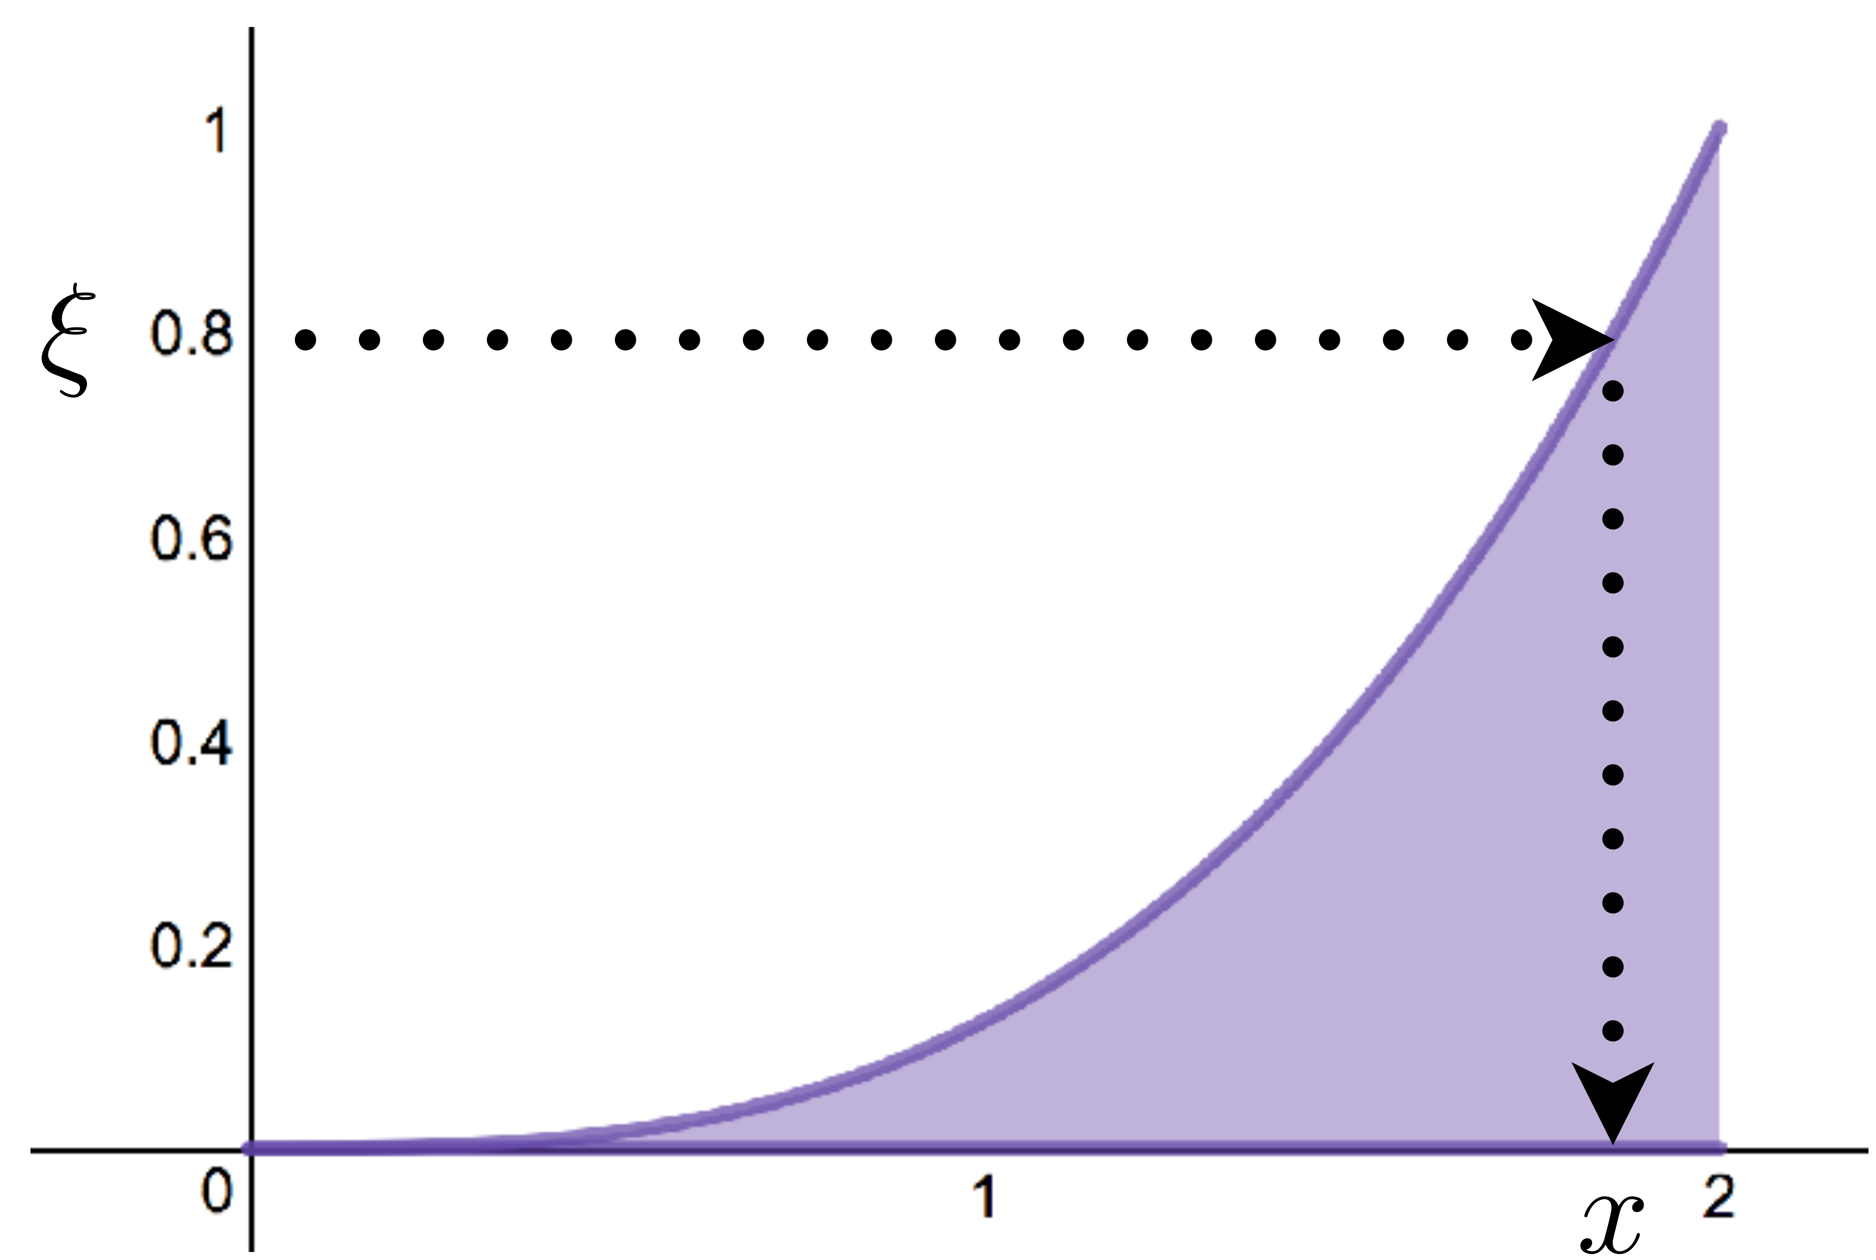

Things to Remember

Monte Carlo integration

- Unbiased estimators
- Good for high-dimensional integrals
- Estimates are visually noisy and need many samples
- Importance sampling can reduce variance (noise) if probability distribution "fits" underlying function

Sampling random variables

- Inversion method, rejection sampling
- Sampling in 1D, 2D, disks, hemispheres

Acknowledgments

Many thanks to Kayvon Fatahalian, Matt Pharr, and Pat Hanrahan, who created the majority of these slides. Thanks also to Keenan Crane.

Extra

Pseudo-Random Number Generation

Example: cellular automata #30

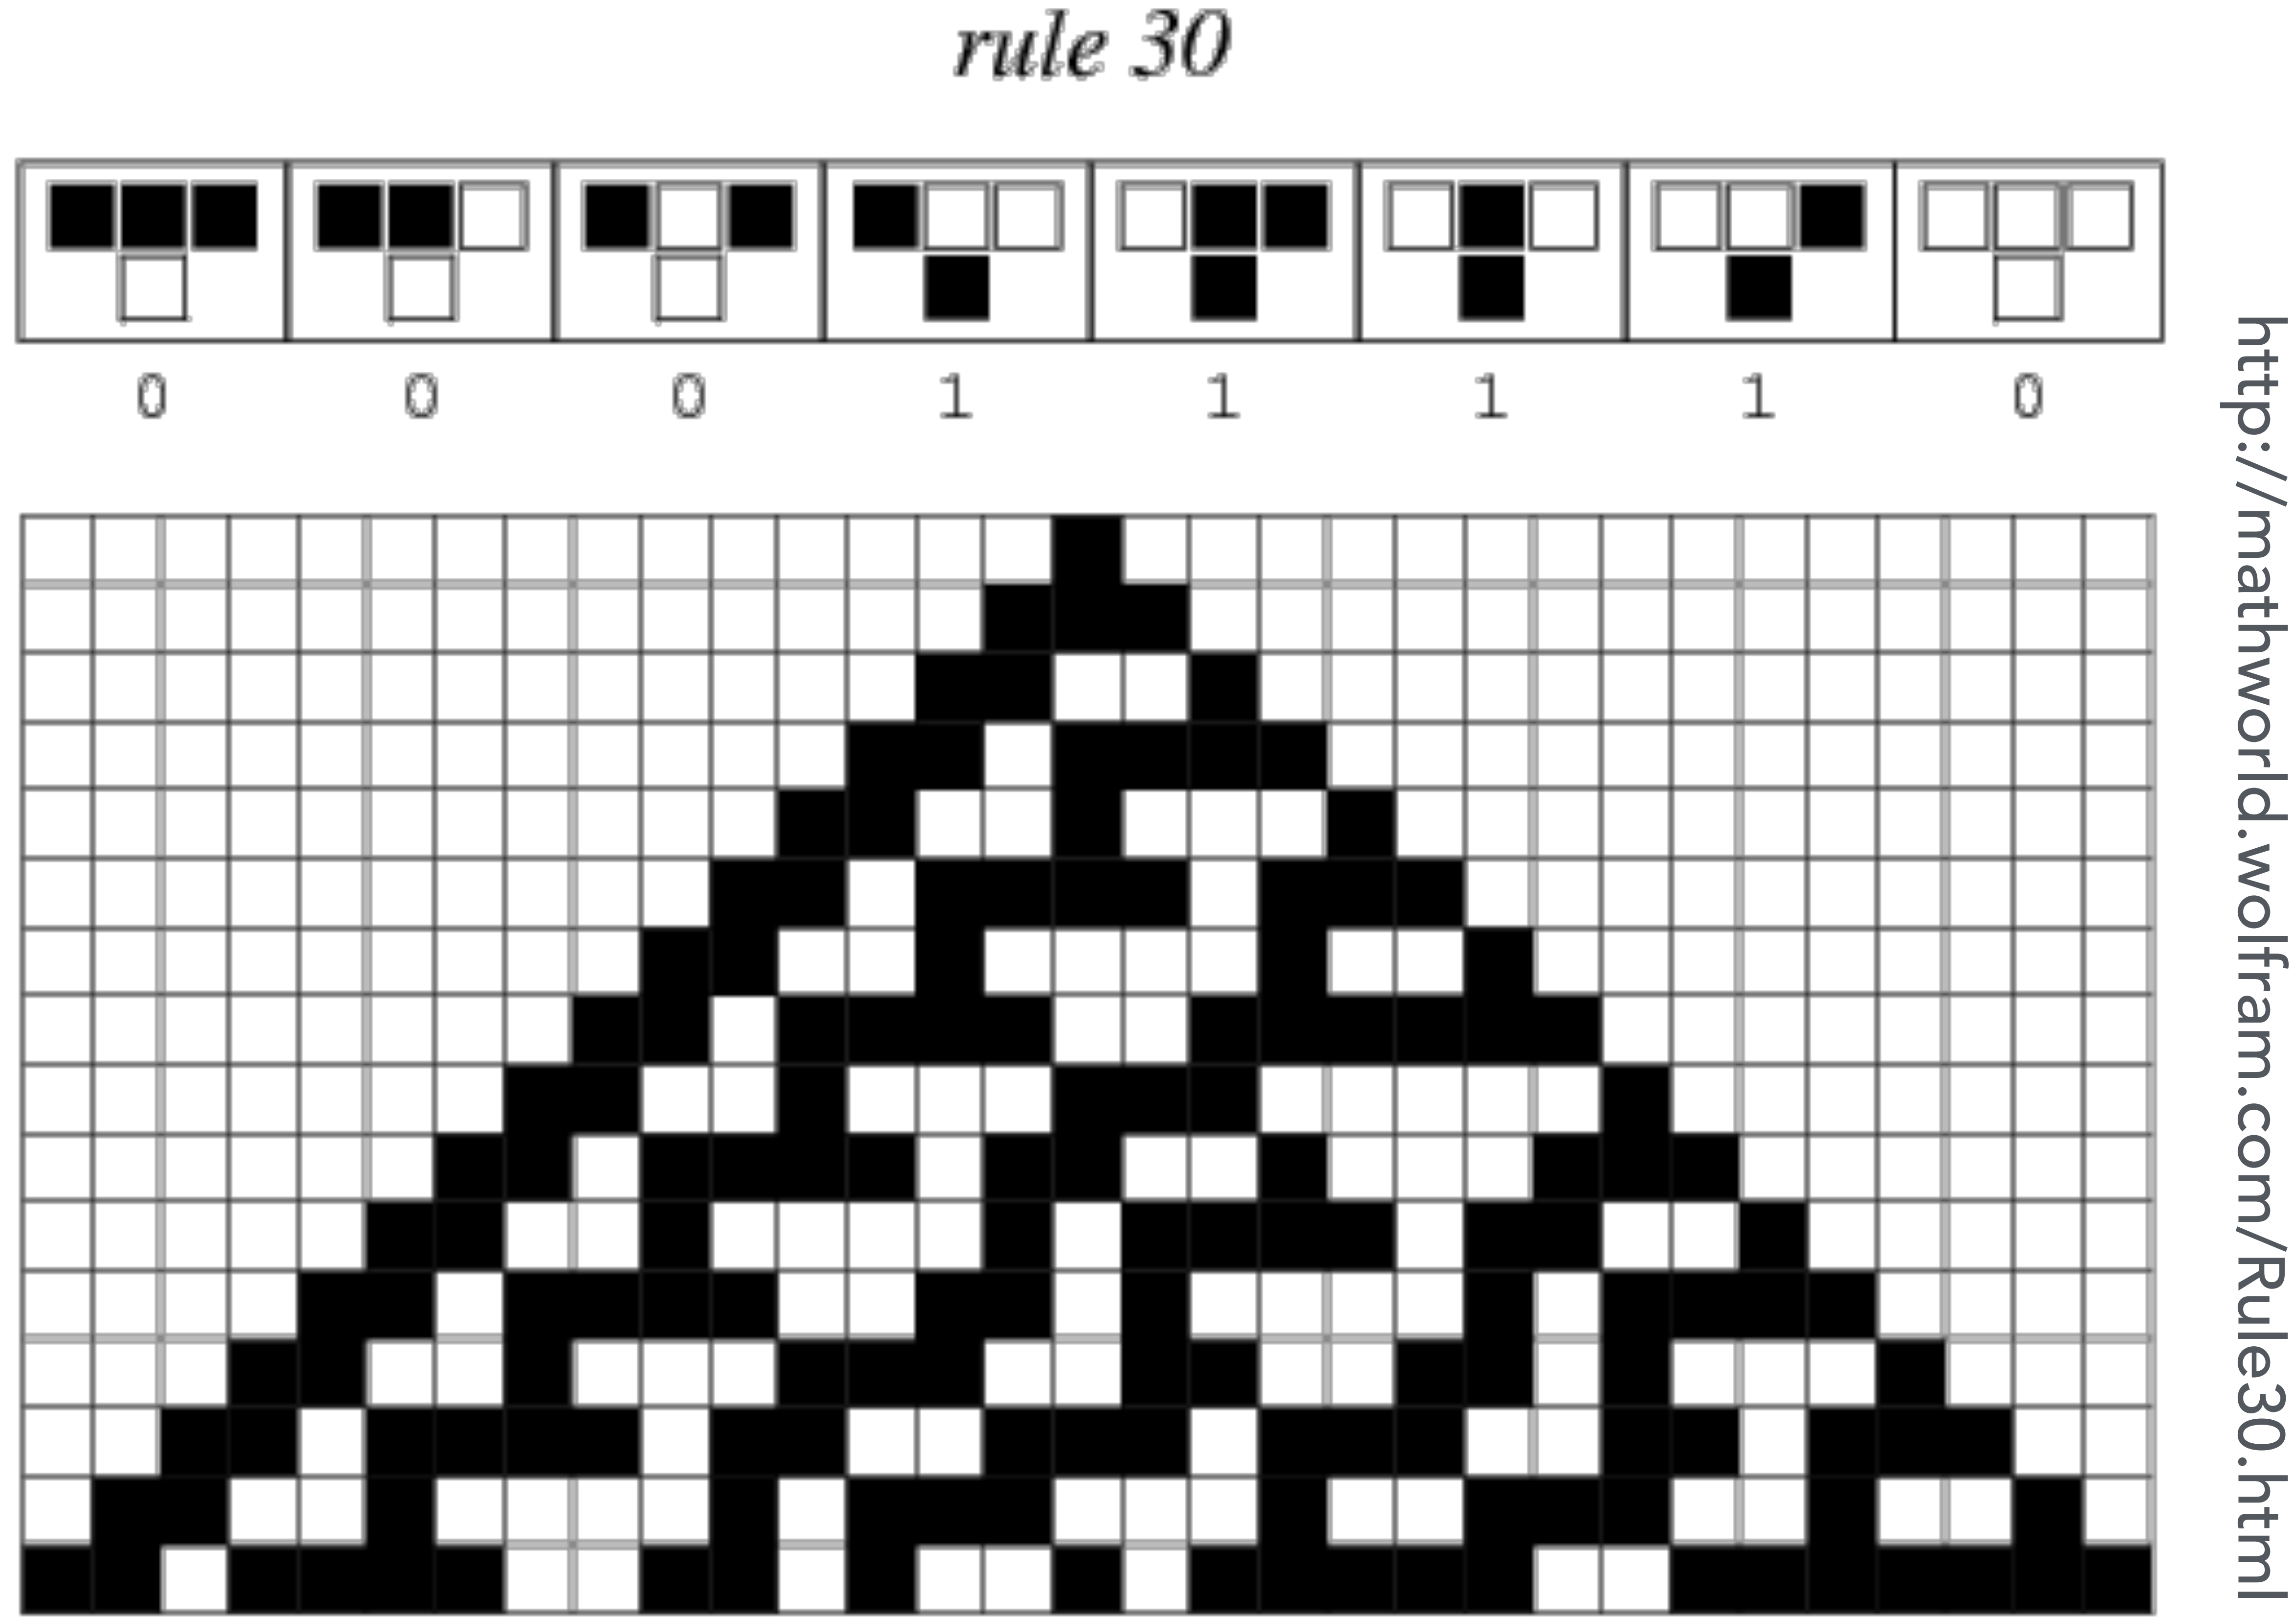

Pseudo-Random Number Generation

Example: cellular automata #30

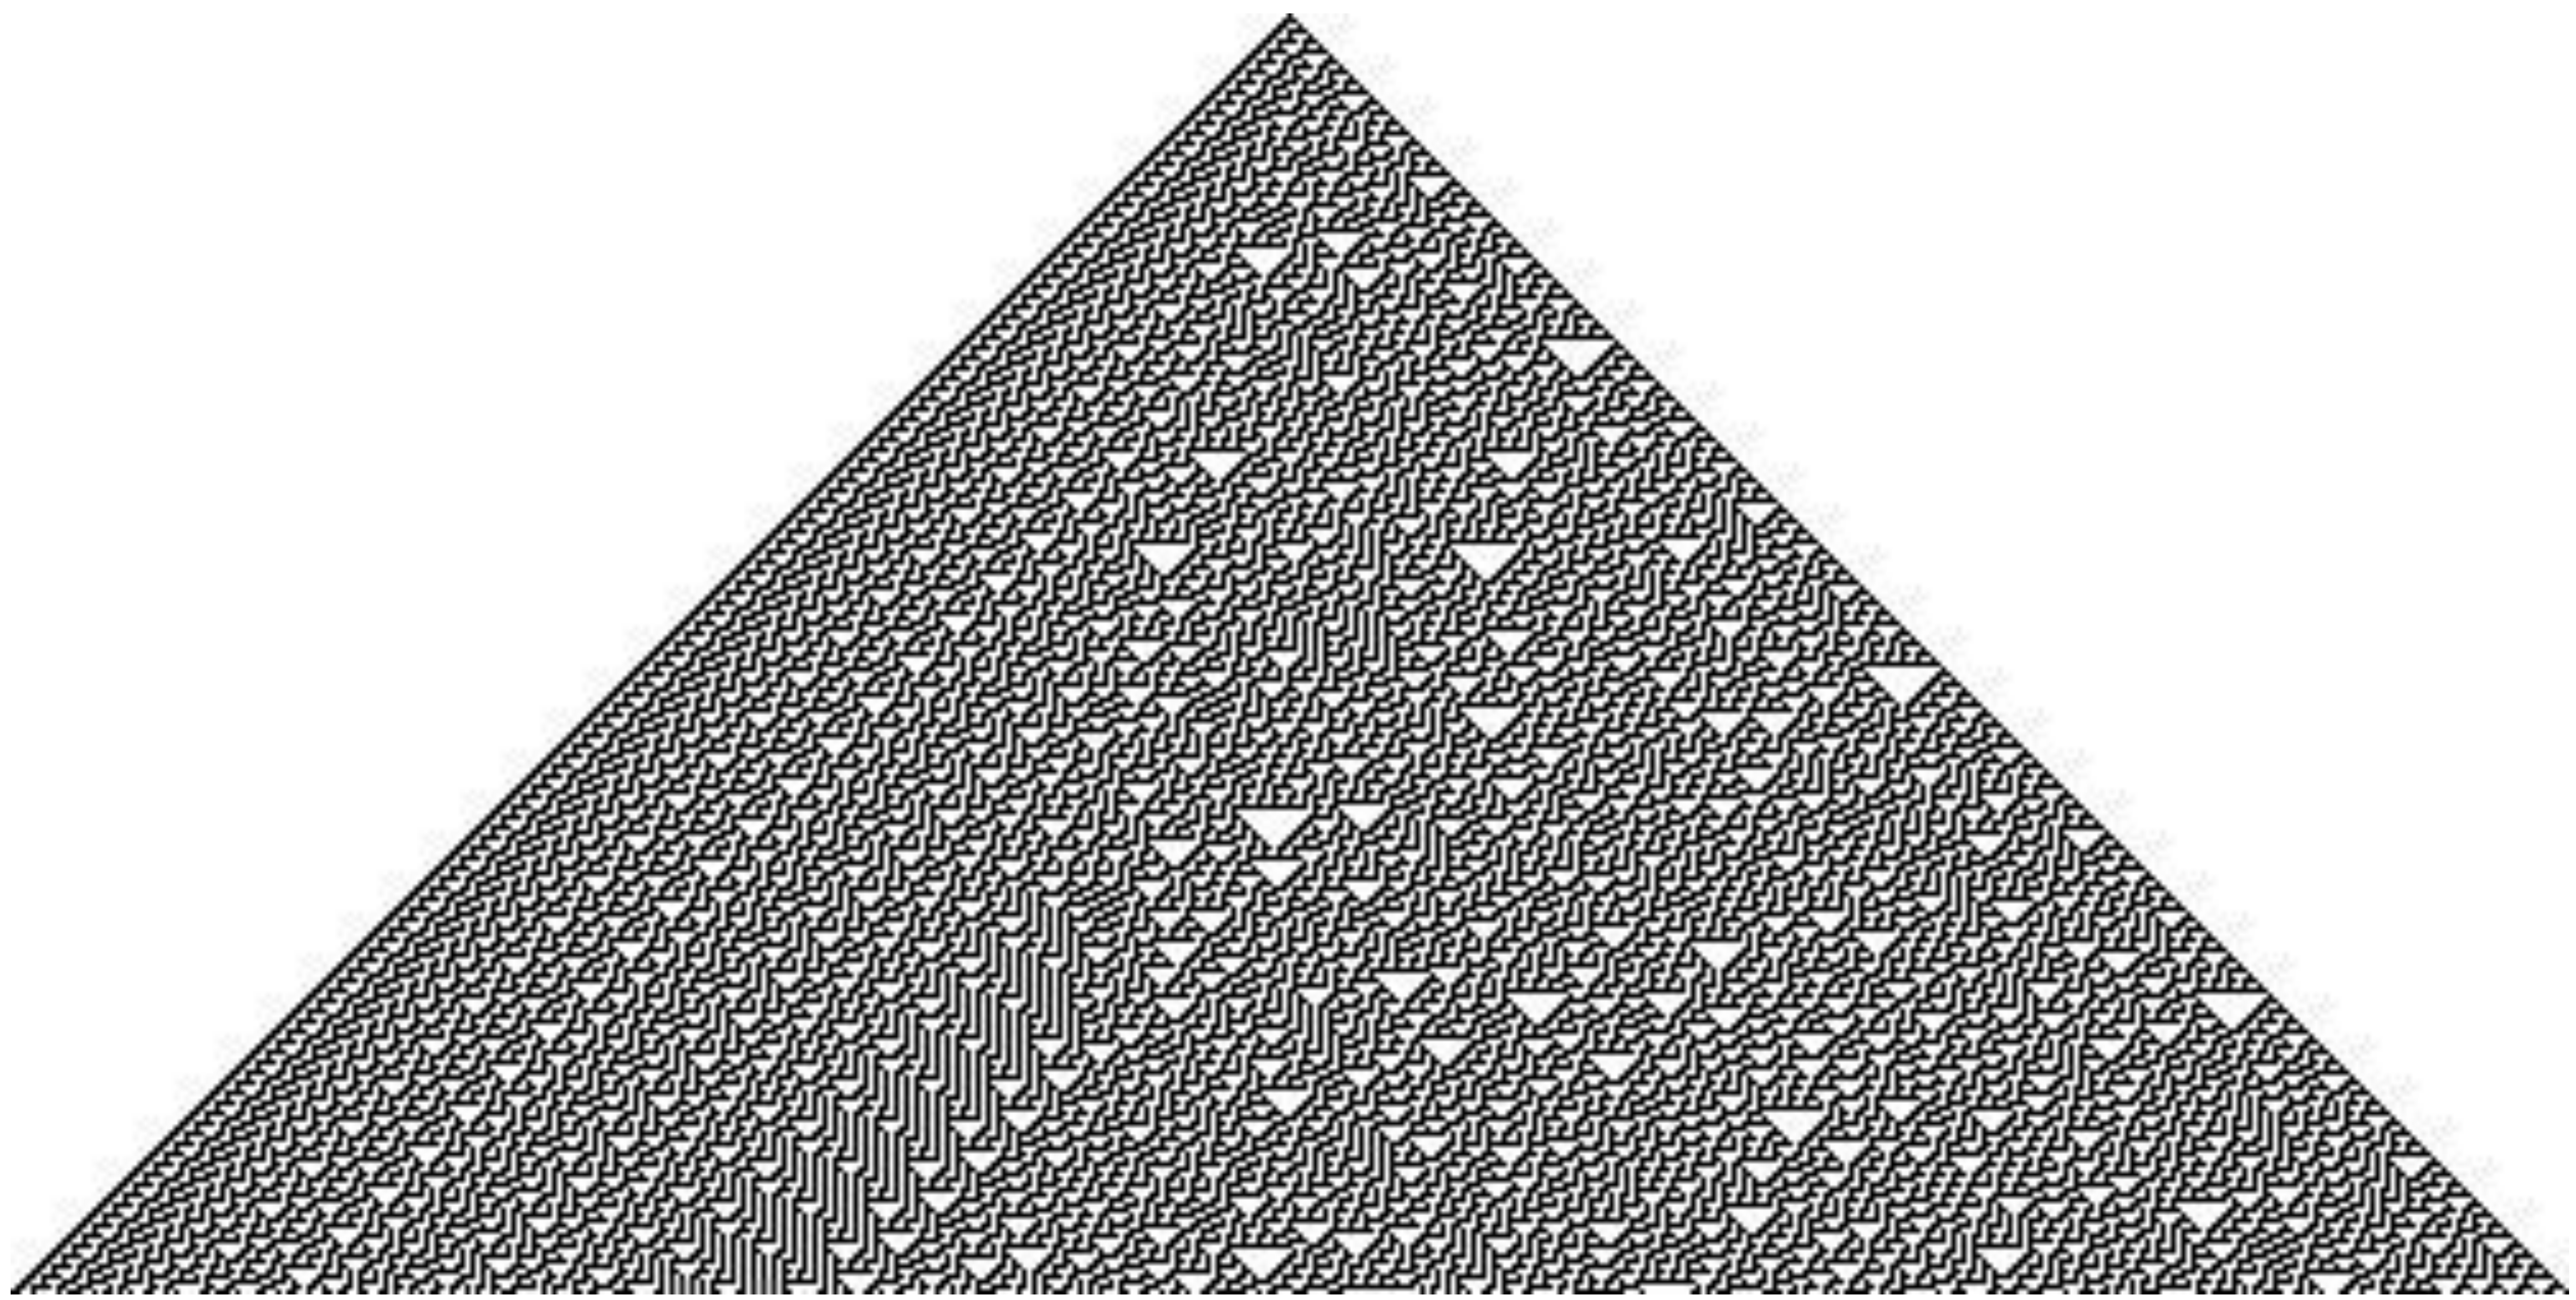

http://mathworld.wolfram.com/Rule30.html

Center line values are a high-quality random bit sequence
Once used for random number generator in Mathematica

Pseudo-Random Number Generation

Credit: Richard Ling

Random Sampling of Disks & Hemispheres

Sampling Unit Circle: Simple but Wrong Method

θ = uniform random angle between 0 and 2π

r = uniform random radius between 0 and 1

Return point: $(r\cos\,\theta, r\sin\,\theta)$

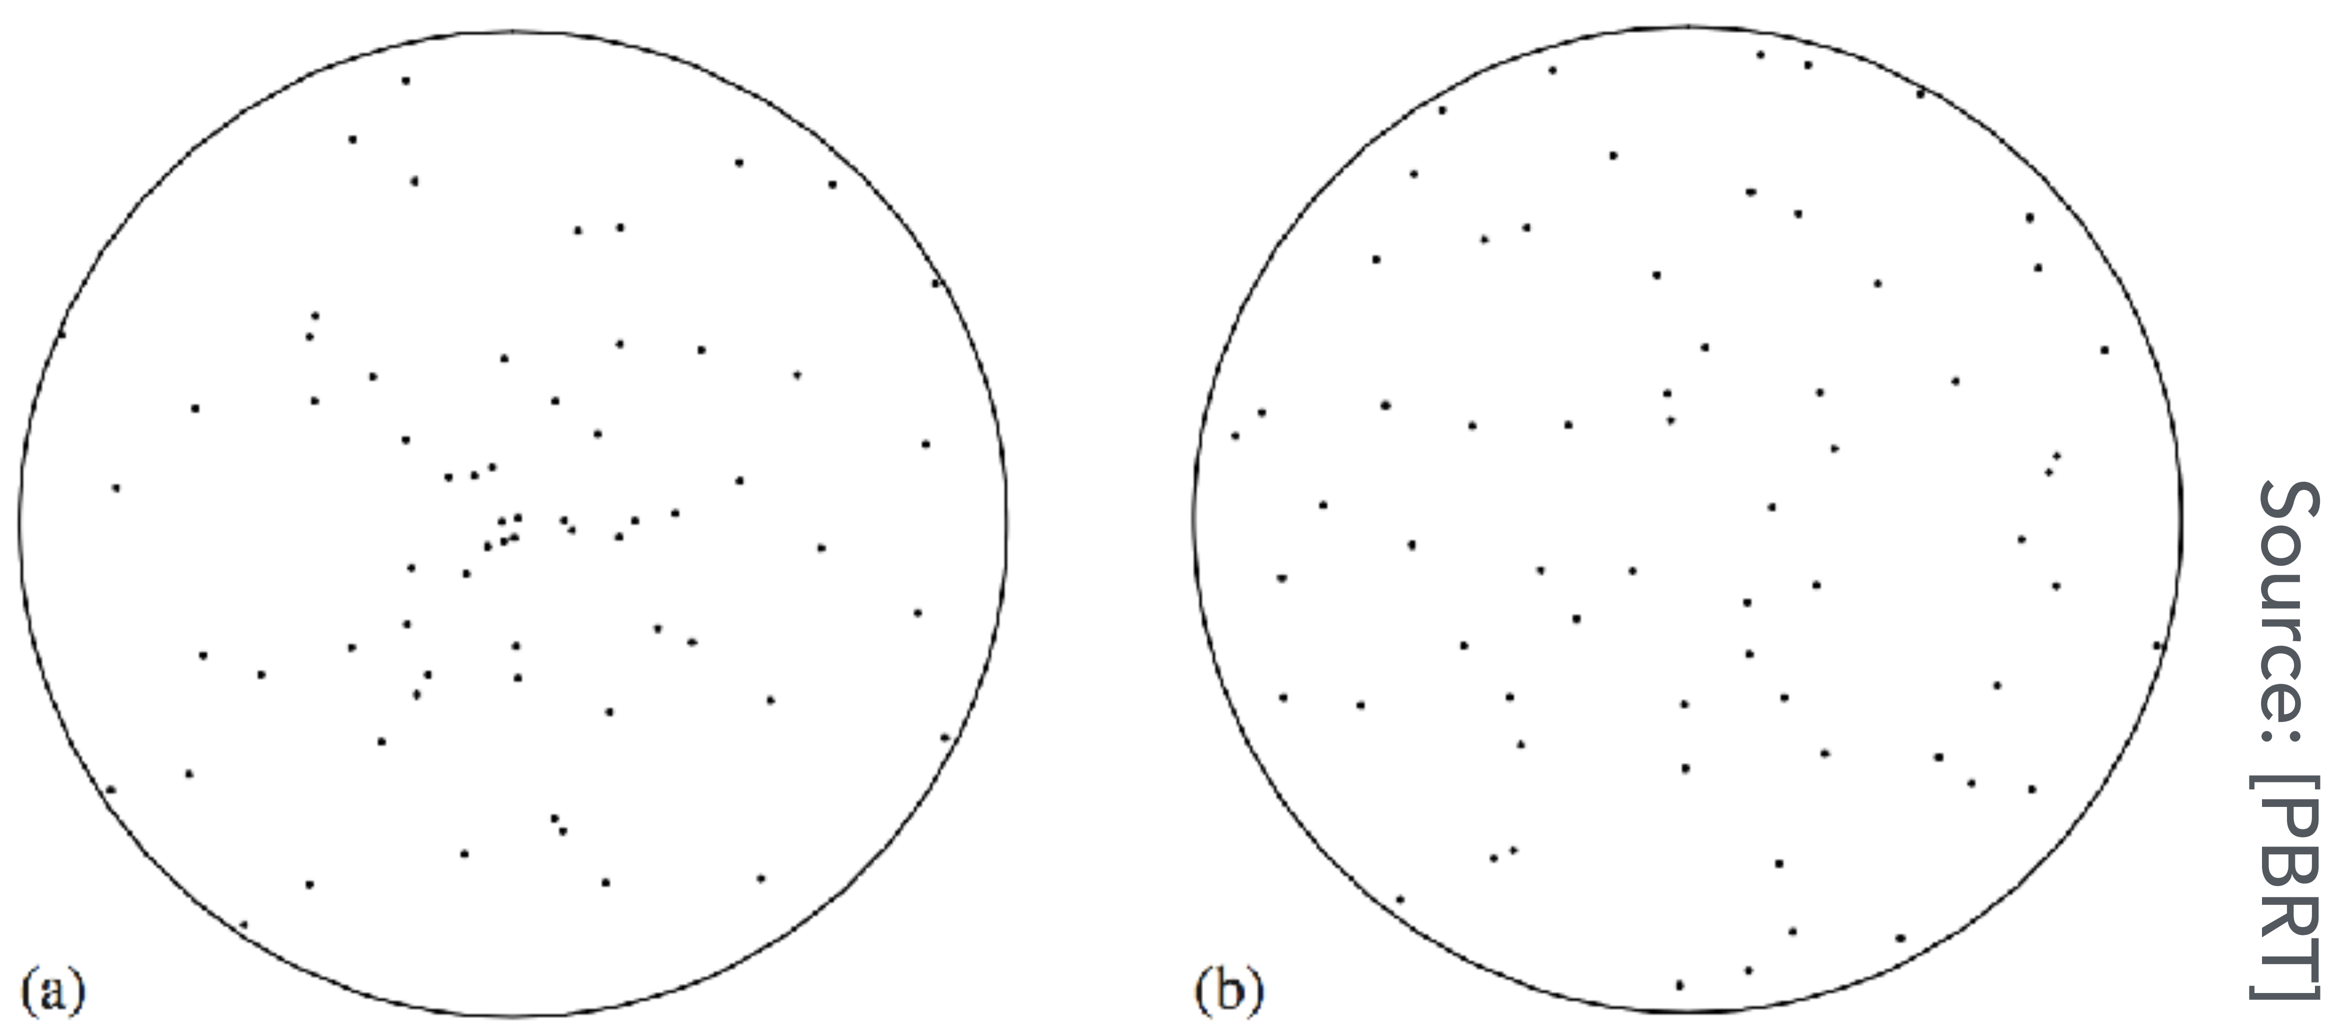

Result Reference

Need to Sample Uniformly in Area

Incorrect Not Equi-areal	Correct Equi-areal
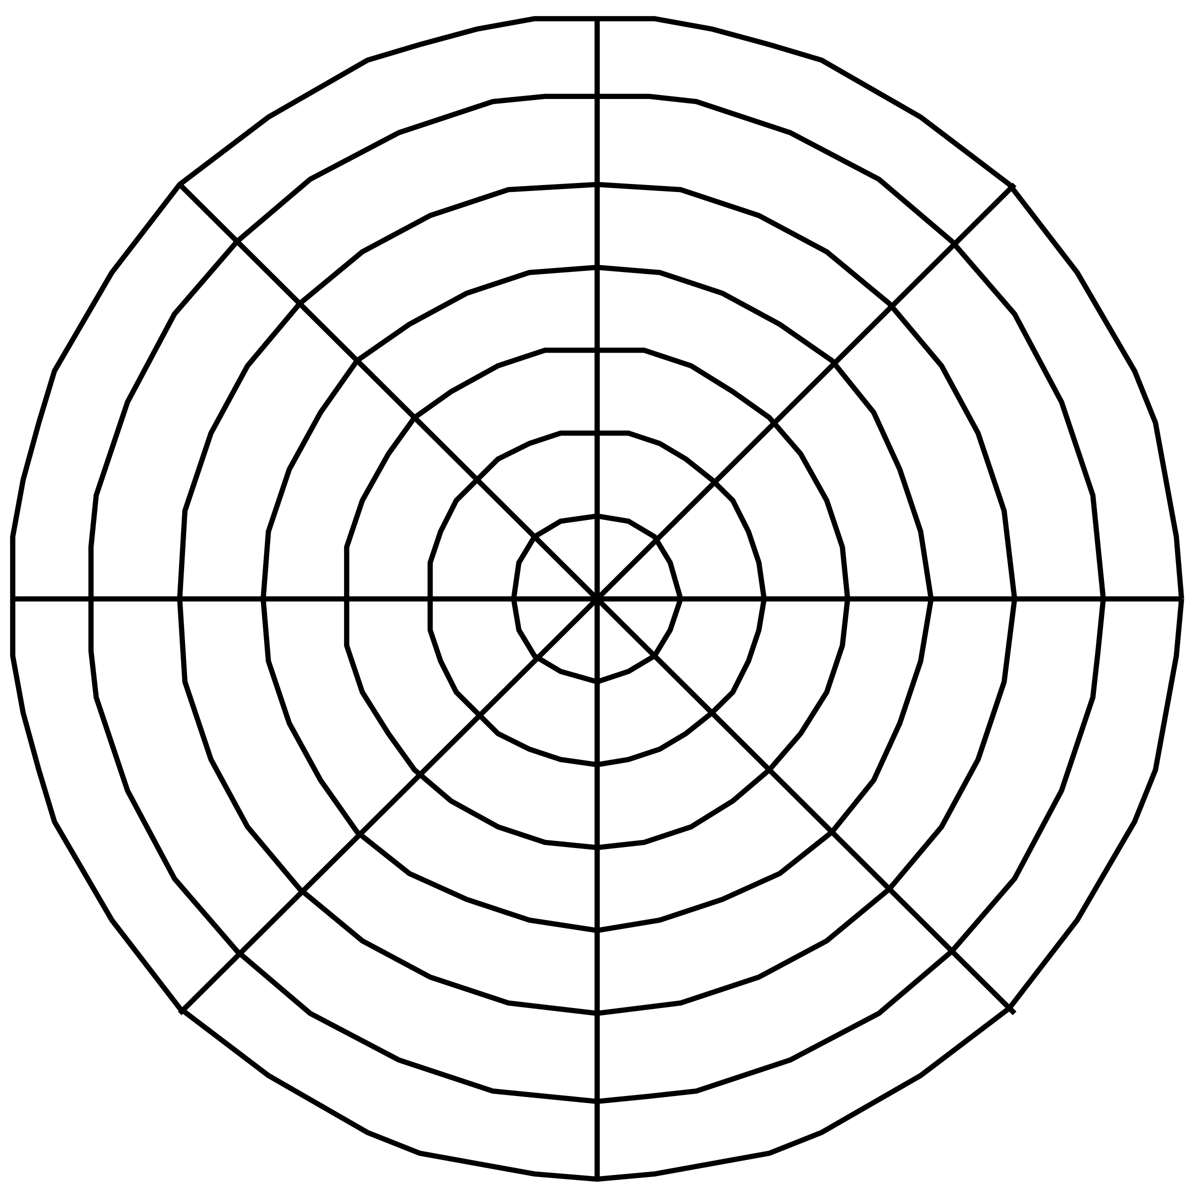	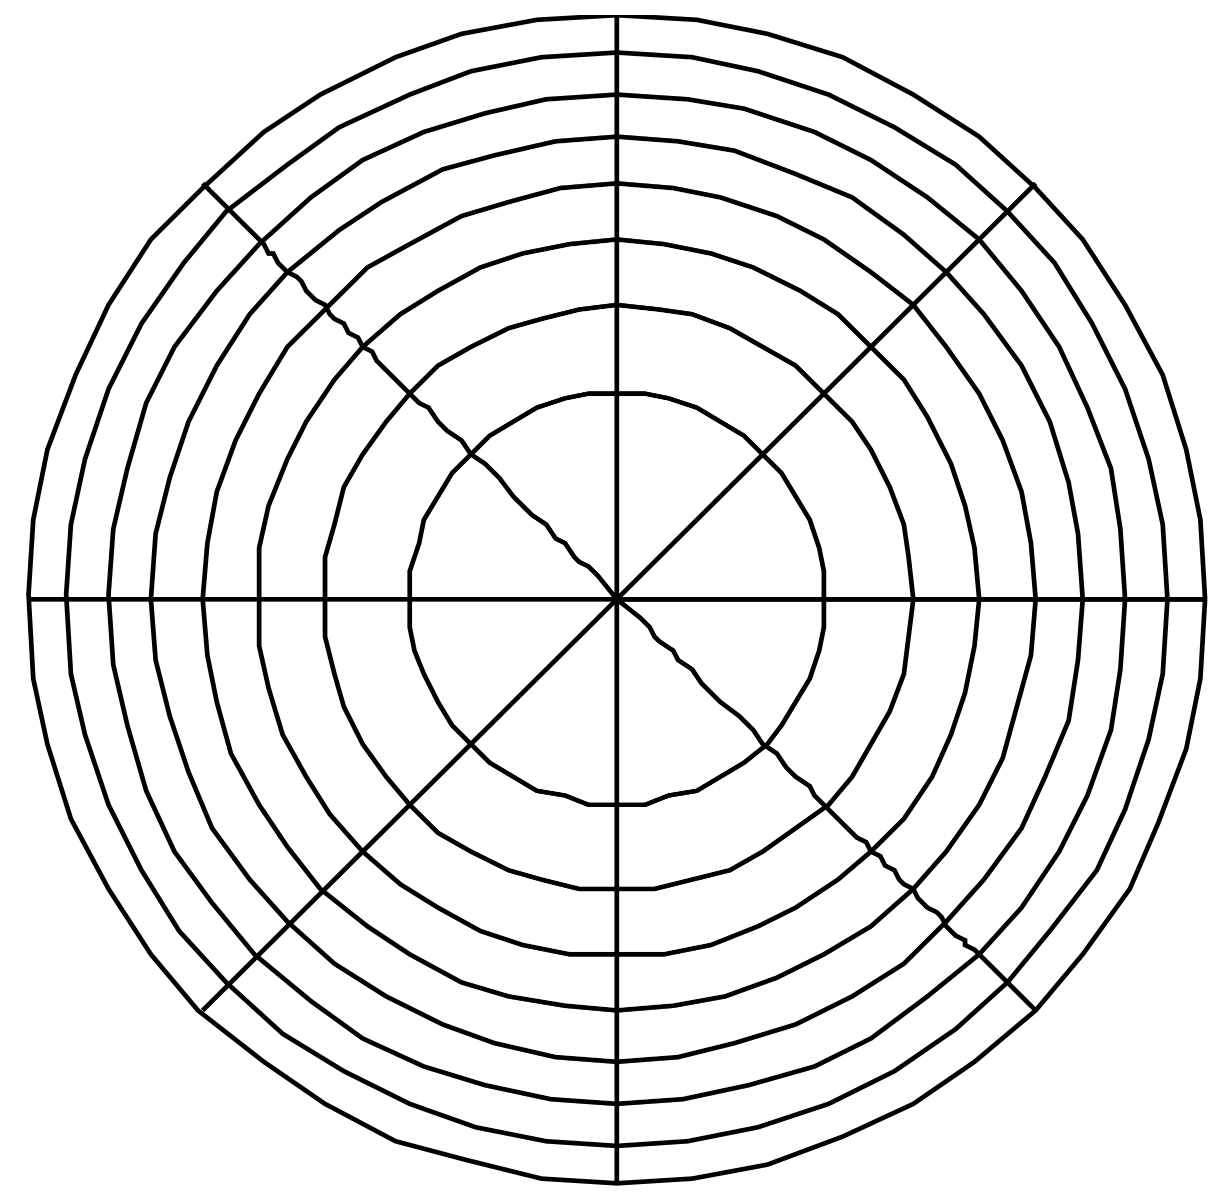
$\theta = 2\pi\xi_1$	$\theta = 2\pi\xi_1$
$r = \xi_2$	$r = \sqrt{\xi_2}$

*** See Shirley et al. p.331 for full explanation using inversion method**

Rejection Sampling Circle's Area

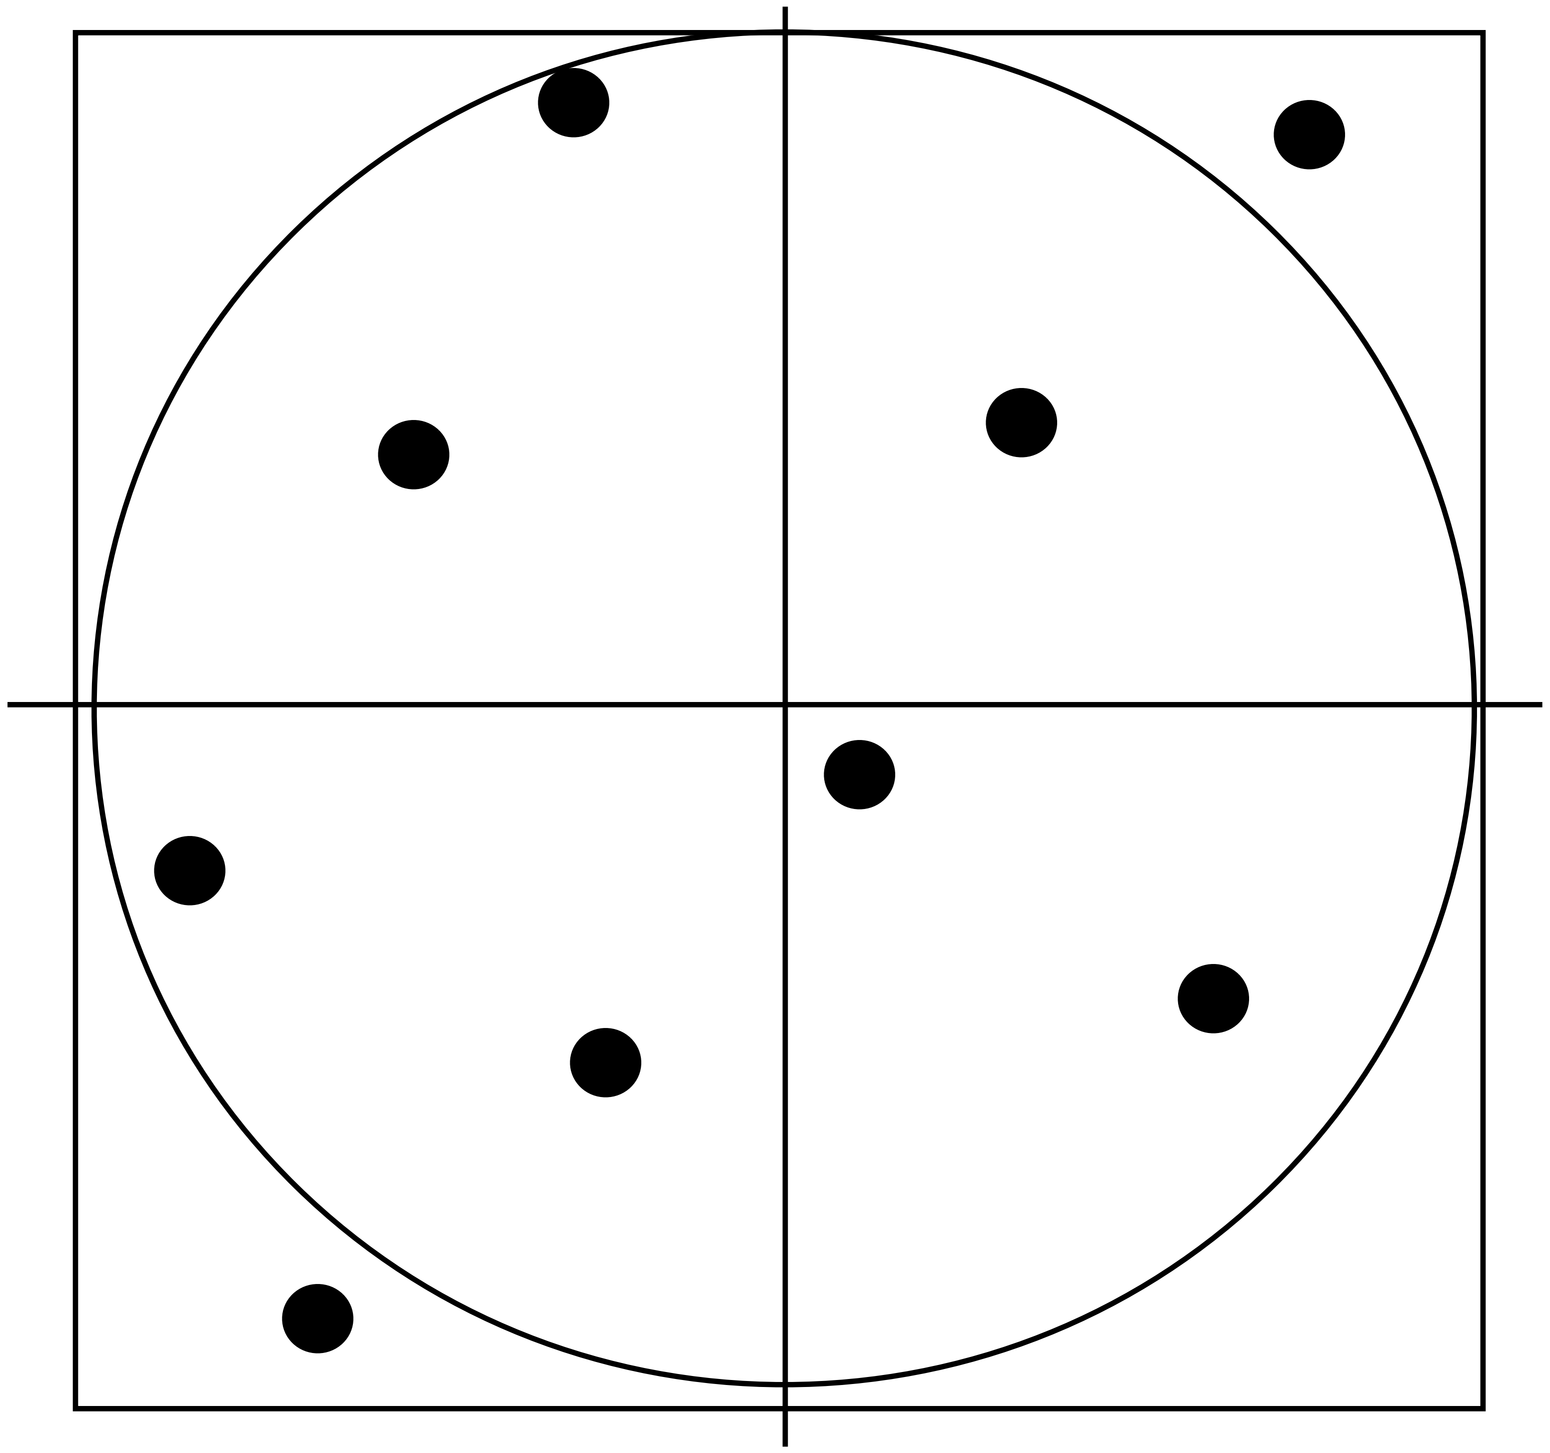

```
do {
   x = 1 - 2 * rand01();
   y = 1 - 2 * rand01();
} while (x*x + y*y > 1.);
```

Efficiency of technique: area of circle / area of square

Uniform Sampling of Hemisphere

Generate random direction on hemisphere (all dirs. equally likely)

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

Direction computed from uniformly distributed point on 2D square:

$$(\xi_1, \xi_2) \rightarrow (\sqrt{1-\xi_1^2}\cos(2\pi\xi_2), \sqrt{1-\xi_1^2}\sin(2\pi\xi_2), \xi_1)$$

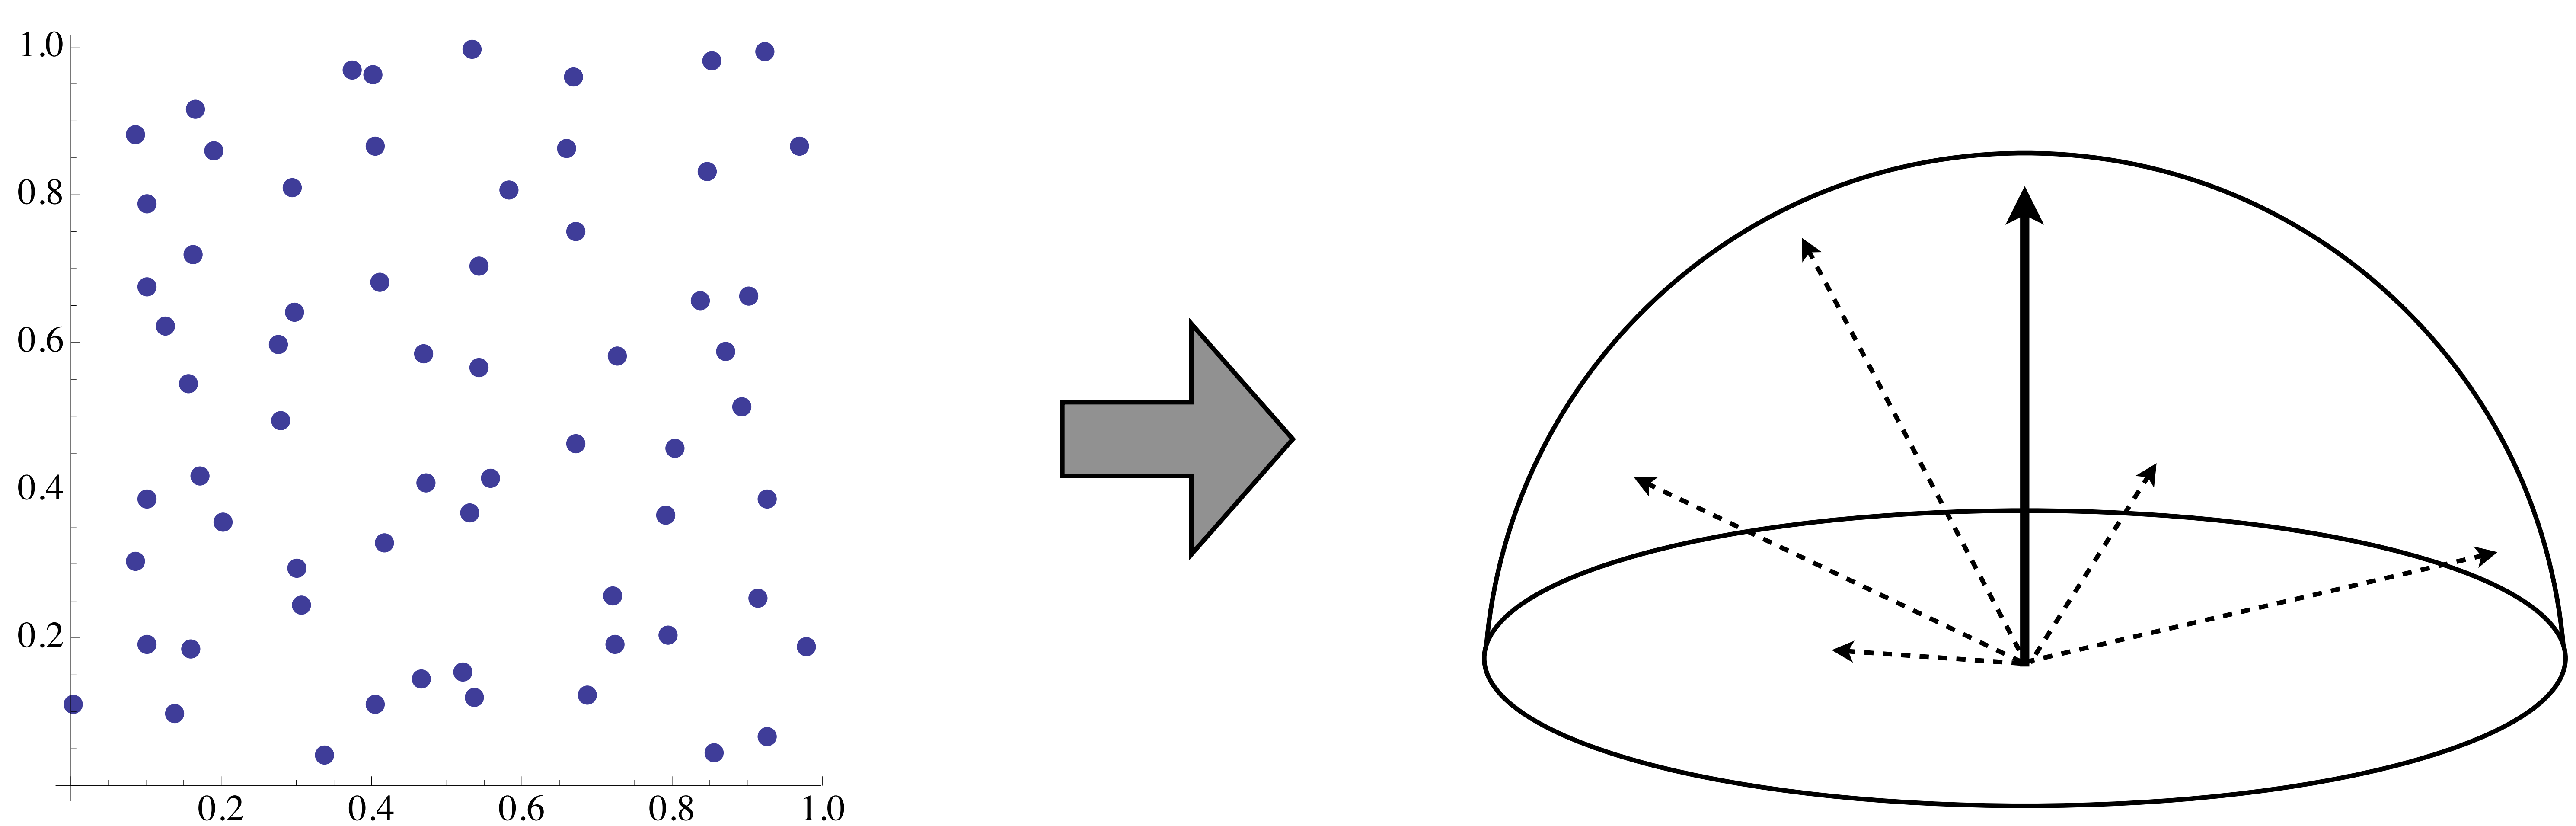

Full derivation: see PBRT 3rd Ed. 13.6.1