Lecture 11:

Monte Carlo Evaluation of the Reflection Equation

Interactive Computer Graphics
Stanford CS248A, Winter 2023

Last time: the reflection equation

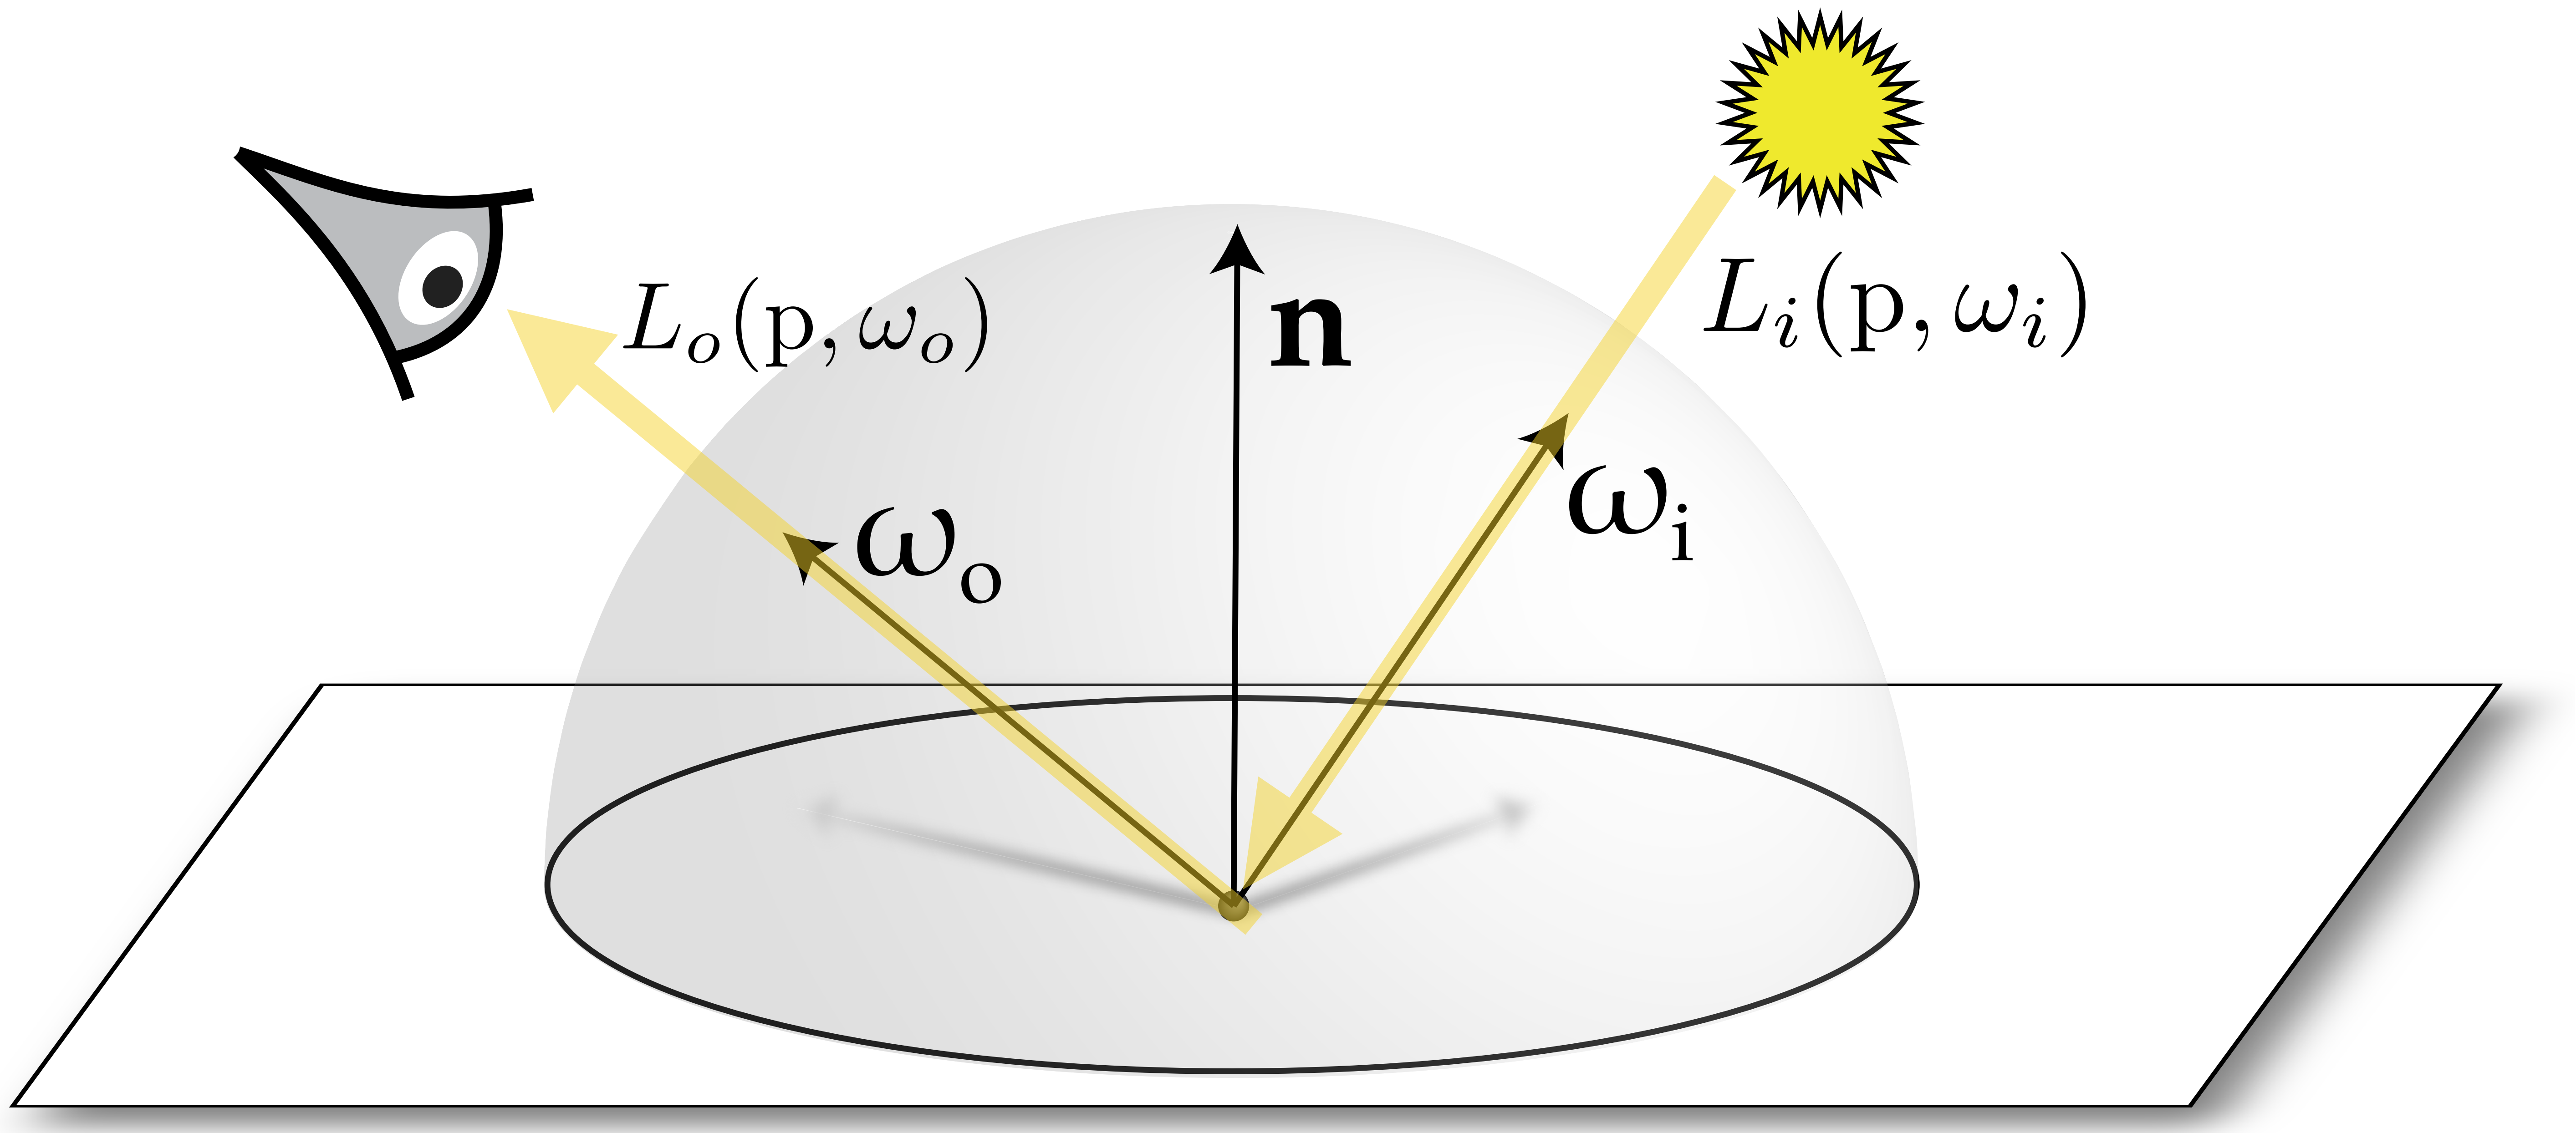

$$L_{\mathrm{o}}(\mathrm{p}, \omega_{\mathrm{o}}) = \int_{\Omega^2} \underbrace{f_{\mathrm{r}}(\mathrm{p}, \omega_{\mathrm{i}} \rightarrow \omega_{\mathrm{o}})}_{\textbf{BRDF}} \, \underbrace{L_{\mathrm{i}}(\mathrm{p}, \omega_{\mathrm{i}}) \, \cos\theta_{\mathrm{i}}}_{\textbf{Illumination}} \, \mathrm{d}\omega_{\mathrm{i}}$$

Review: radiometry and illumination

Review: differential solid angles

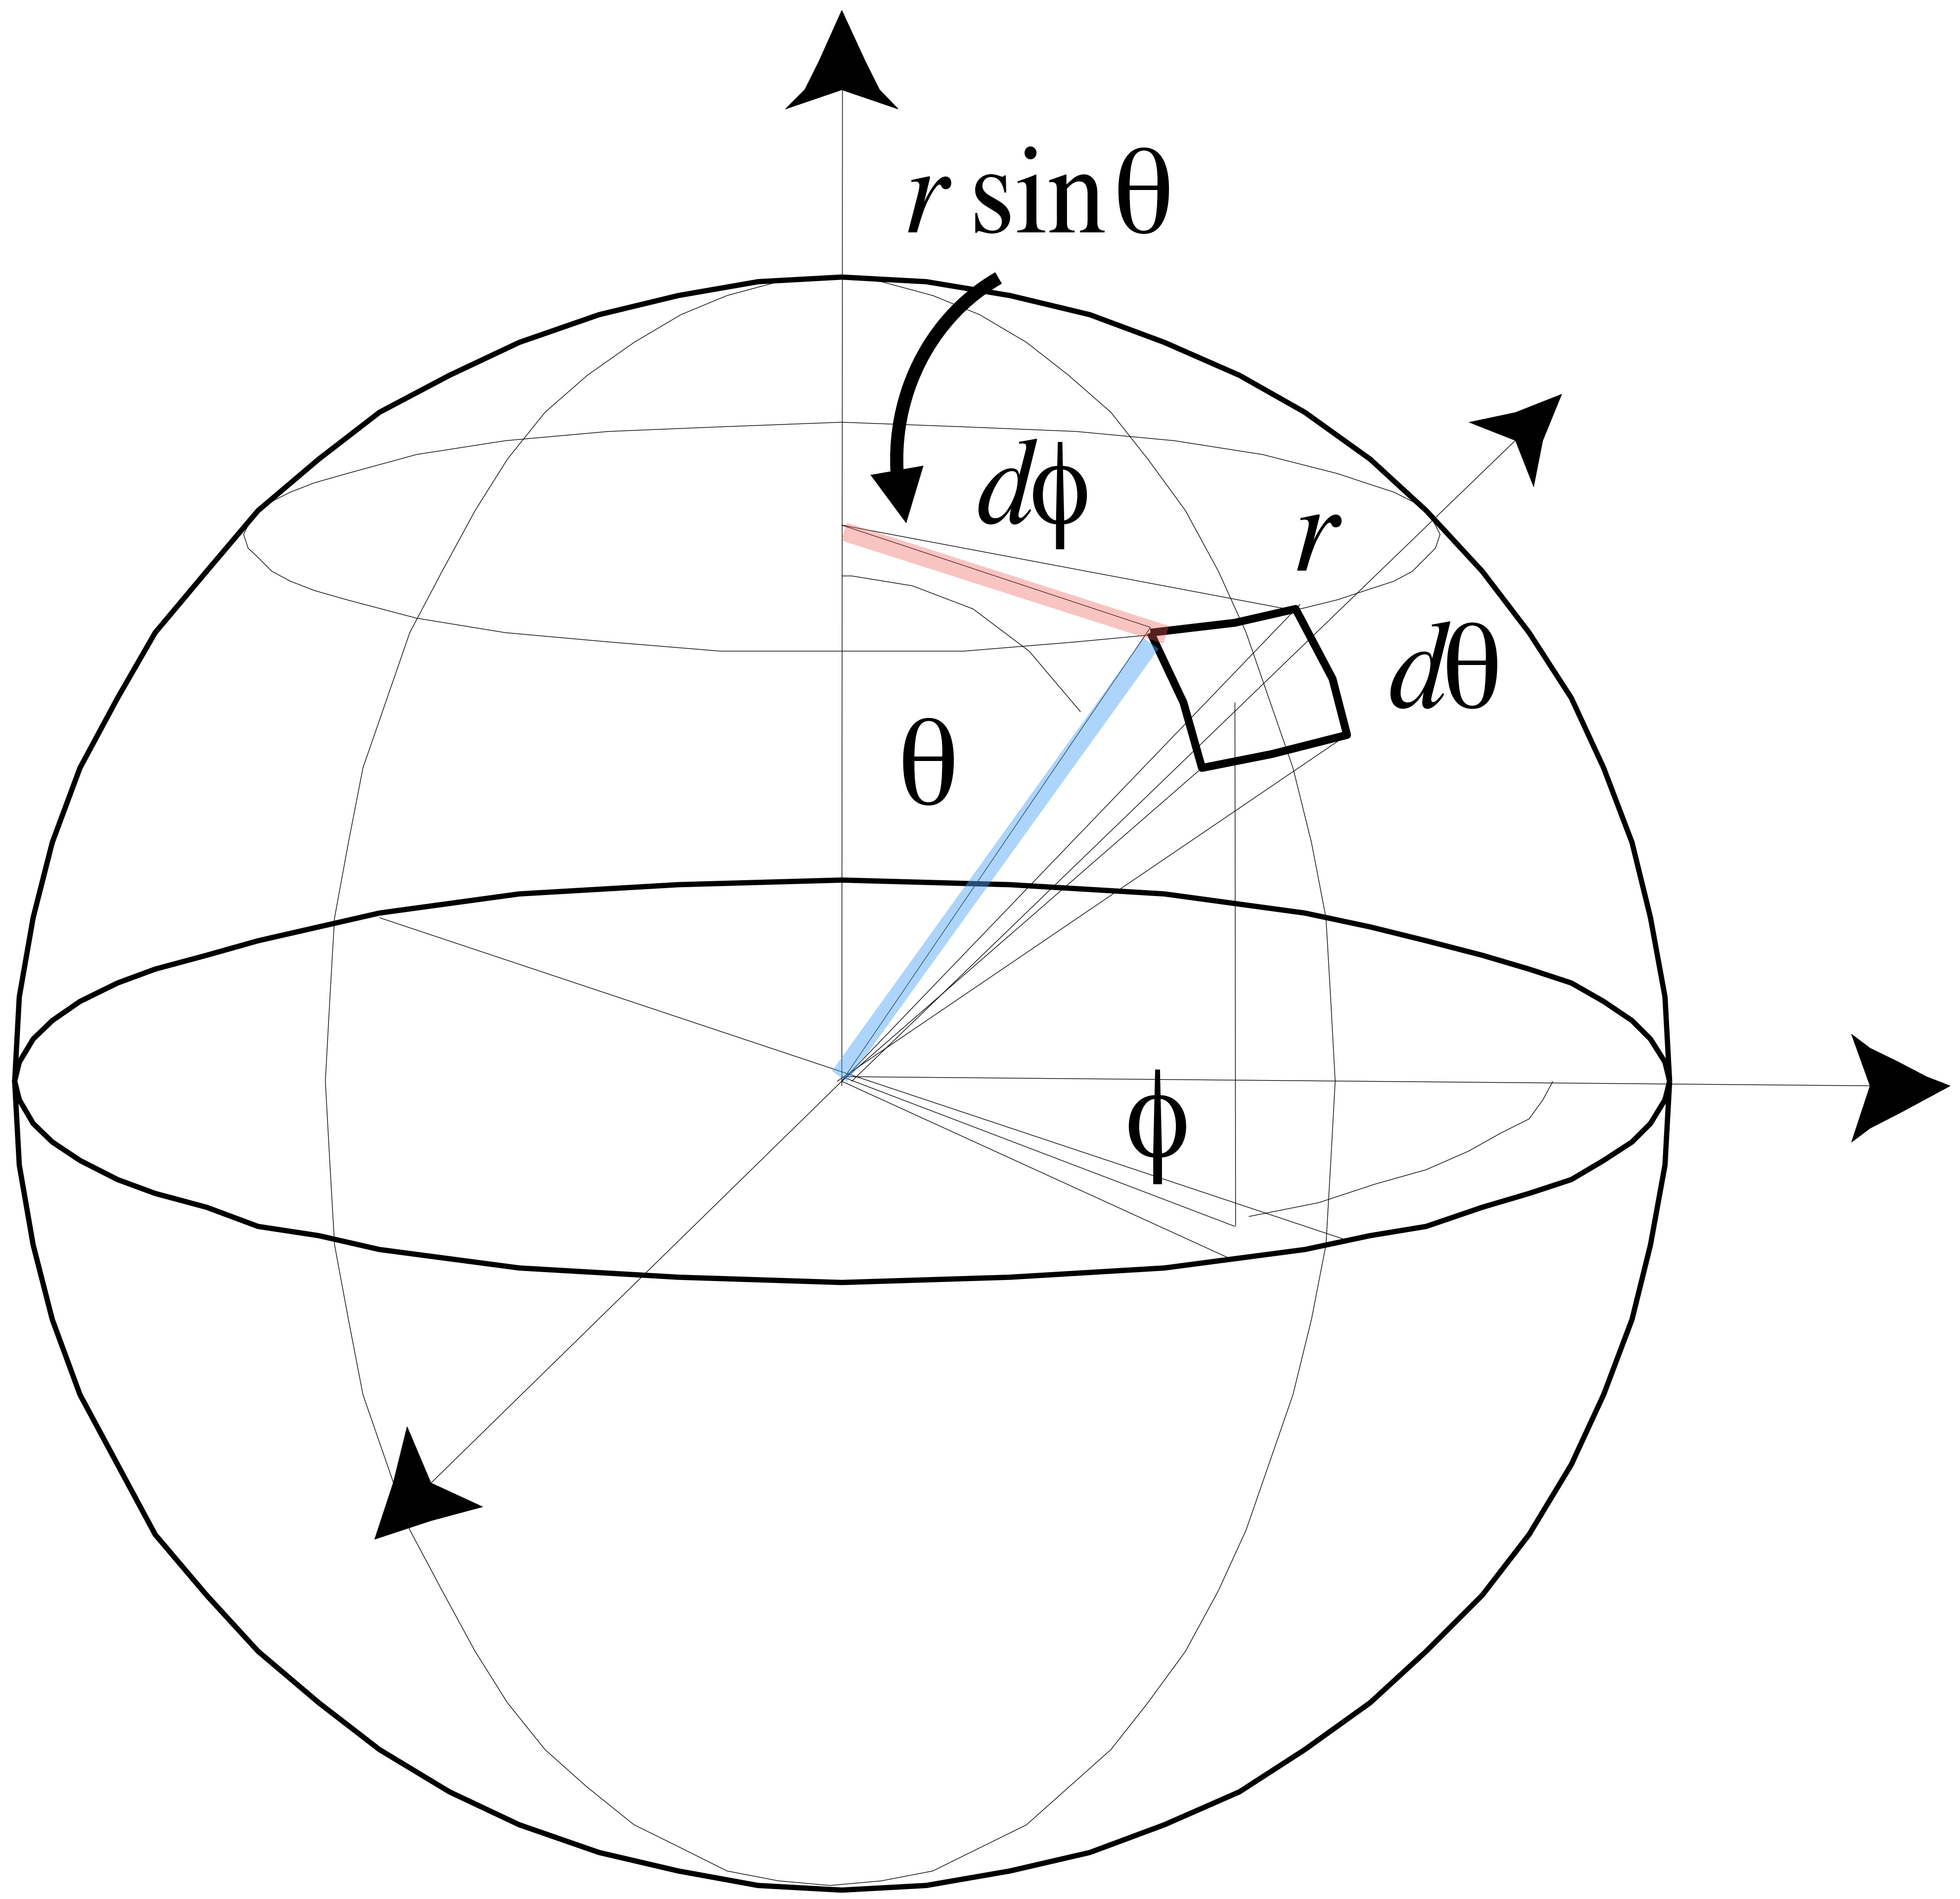

$$dA = (r\,d\theta)(r\sin\theta\,d\phi)$$
$$= r^2\sin\theta\,d\theta\,d\phi$$

$$d\omega = \frac{dA}{r^2} = \sin\theta\,d\theta\,d\phi$$

Review: radiance

Radiance (L) is energy along a ray defined by origin point *p* and direction ω

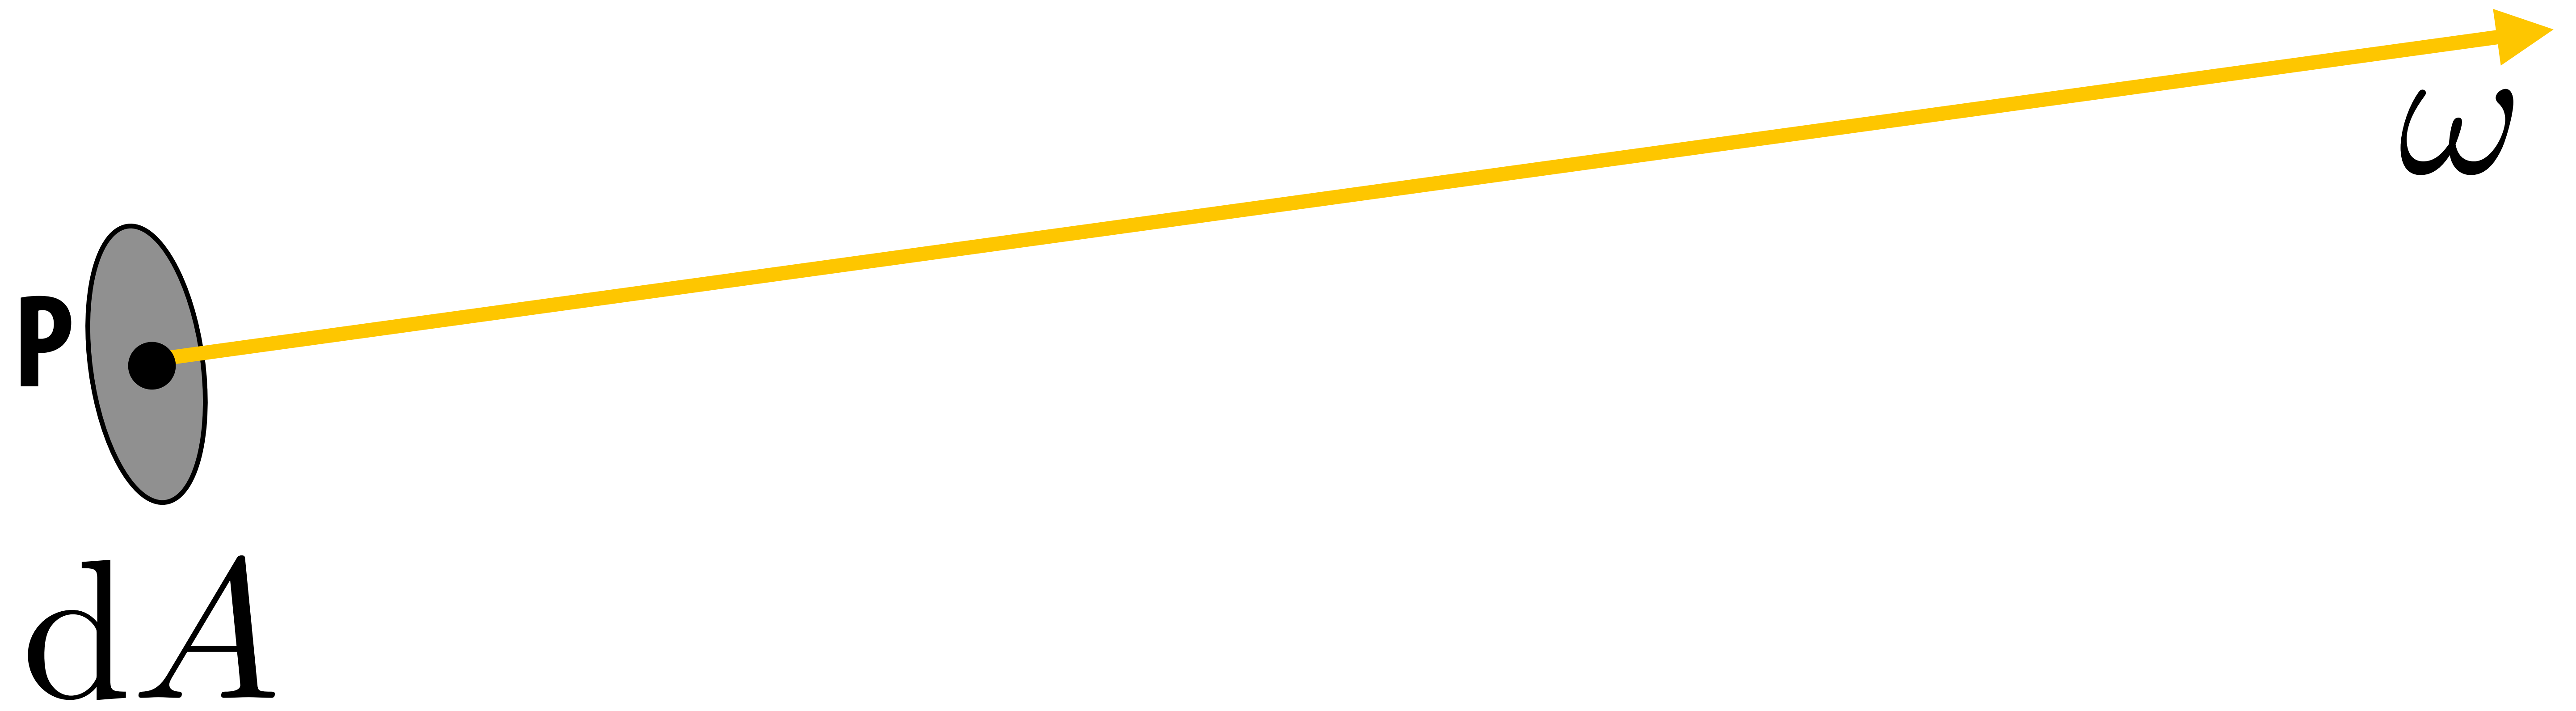

- **Radiance is the solid angle density of irradiance (irradiance per unit direction)**

  where ω denotes that the differential surface area is oriented to face in the direction

Review: irradiance = power per unit area

Irradiance at surface is proportional to cosine of angle between light direction and surface normal. (Lambert's Law)

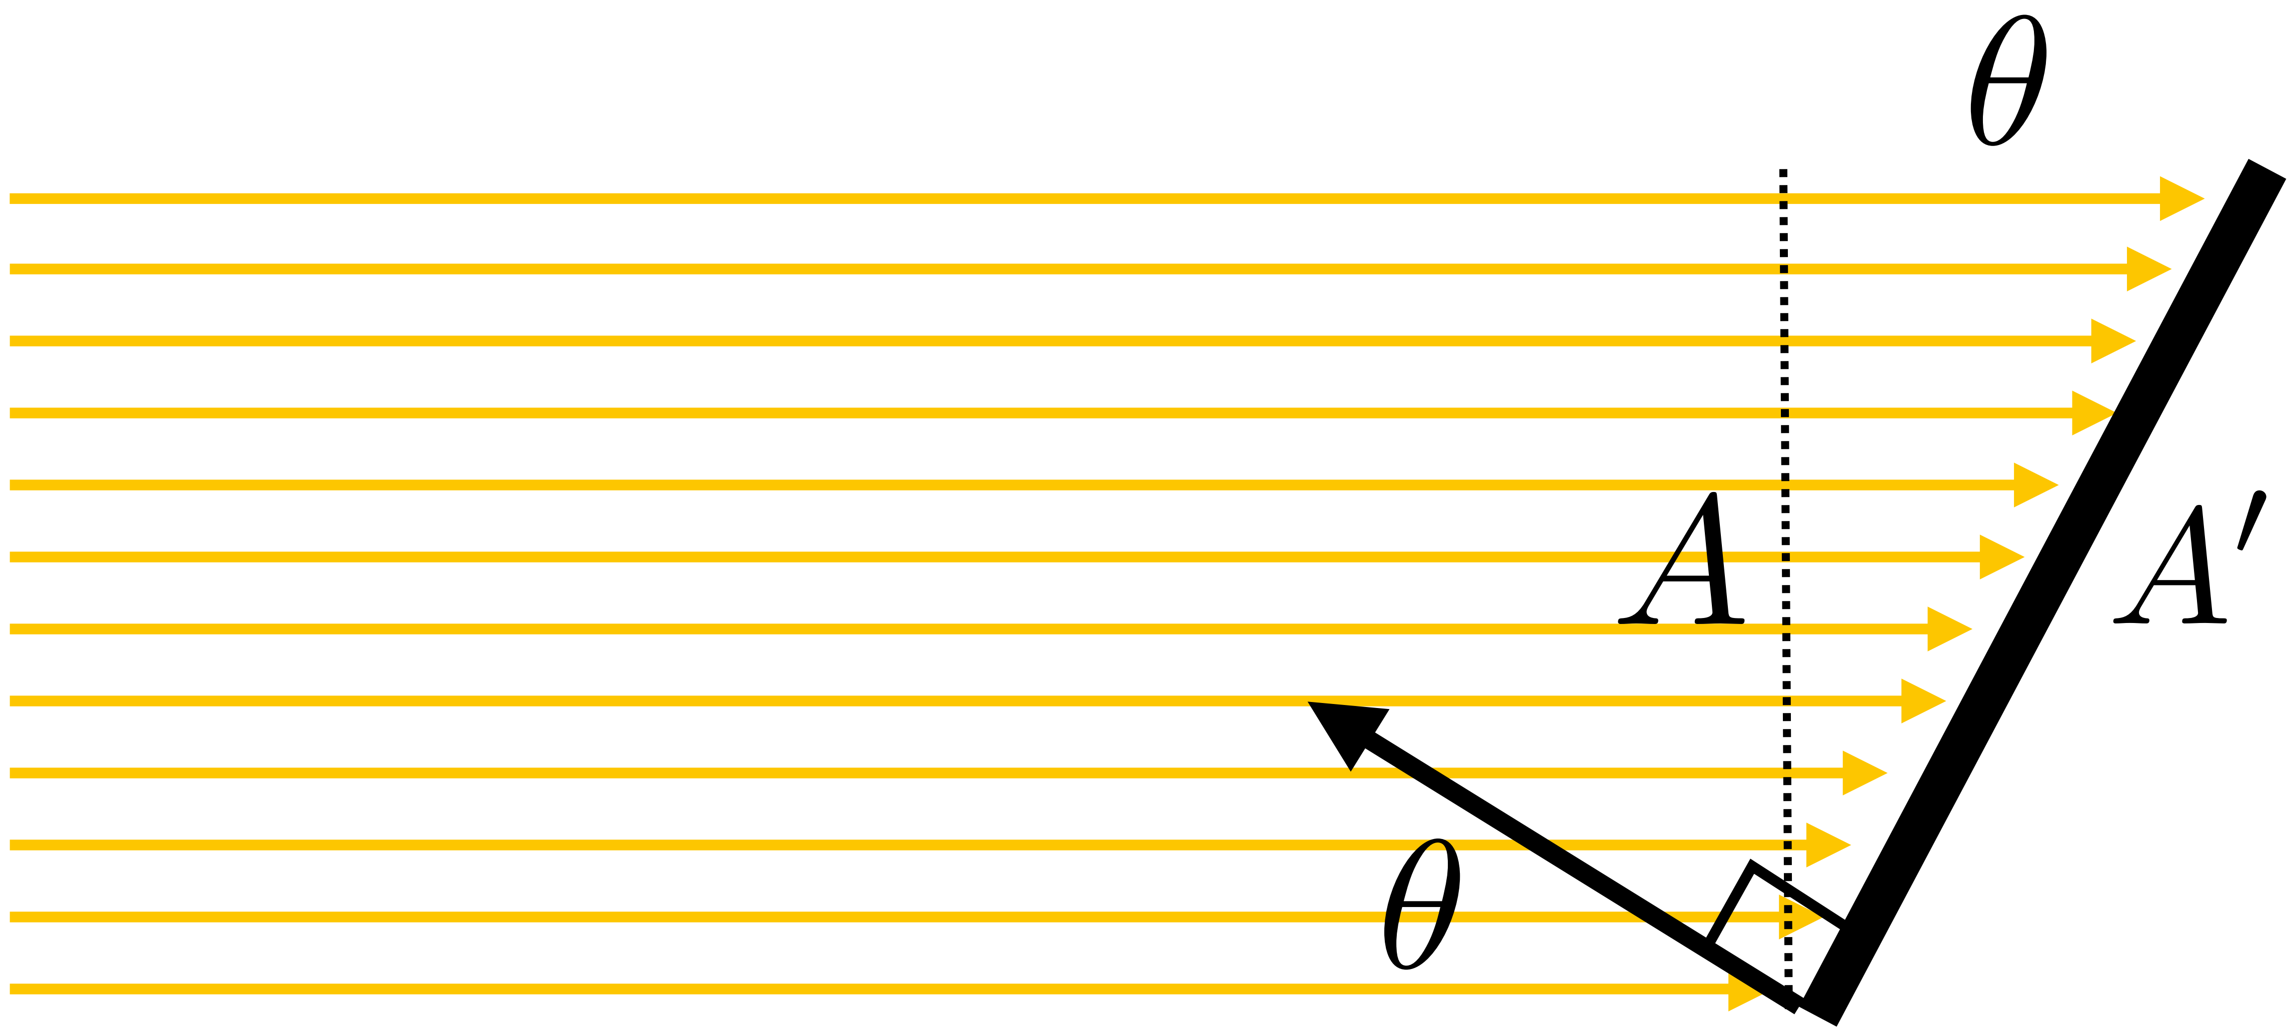

$$A = A' \cos\theta$$

$$E = \frac{\Phi}{A'} = \frac{\Phi\cos\theta}{A}$$

Review: how much light hits the surface at point p? (from multiple point light sources)

(irradiance at point P)

$$\sum_i L_i \cos\theta_i$$

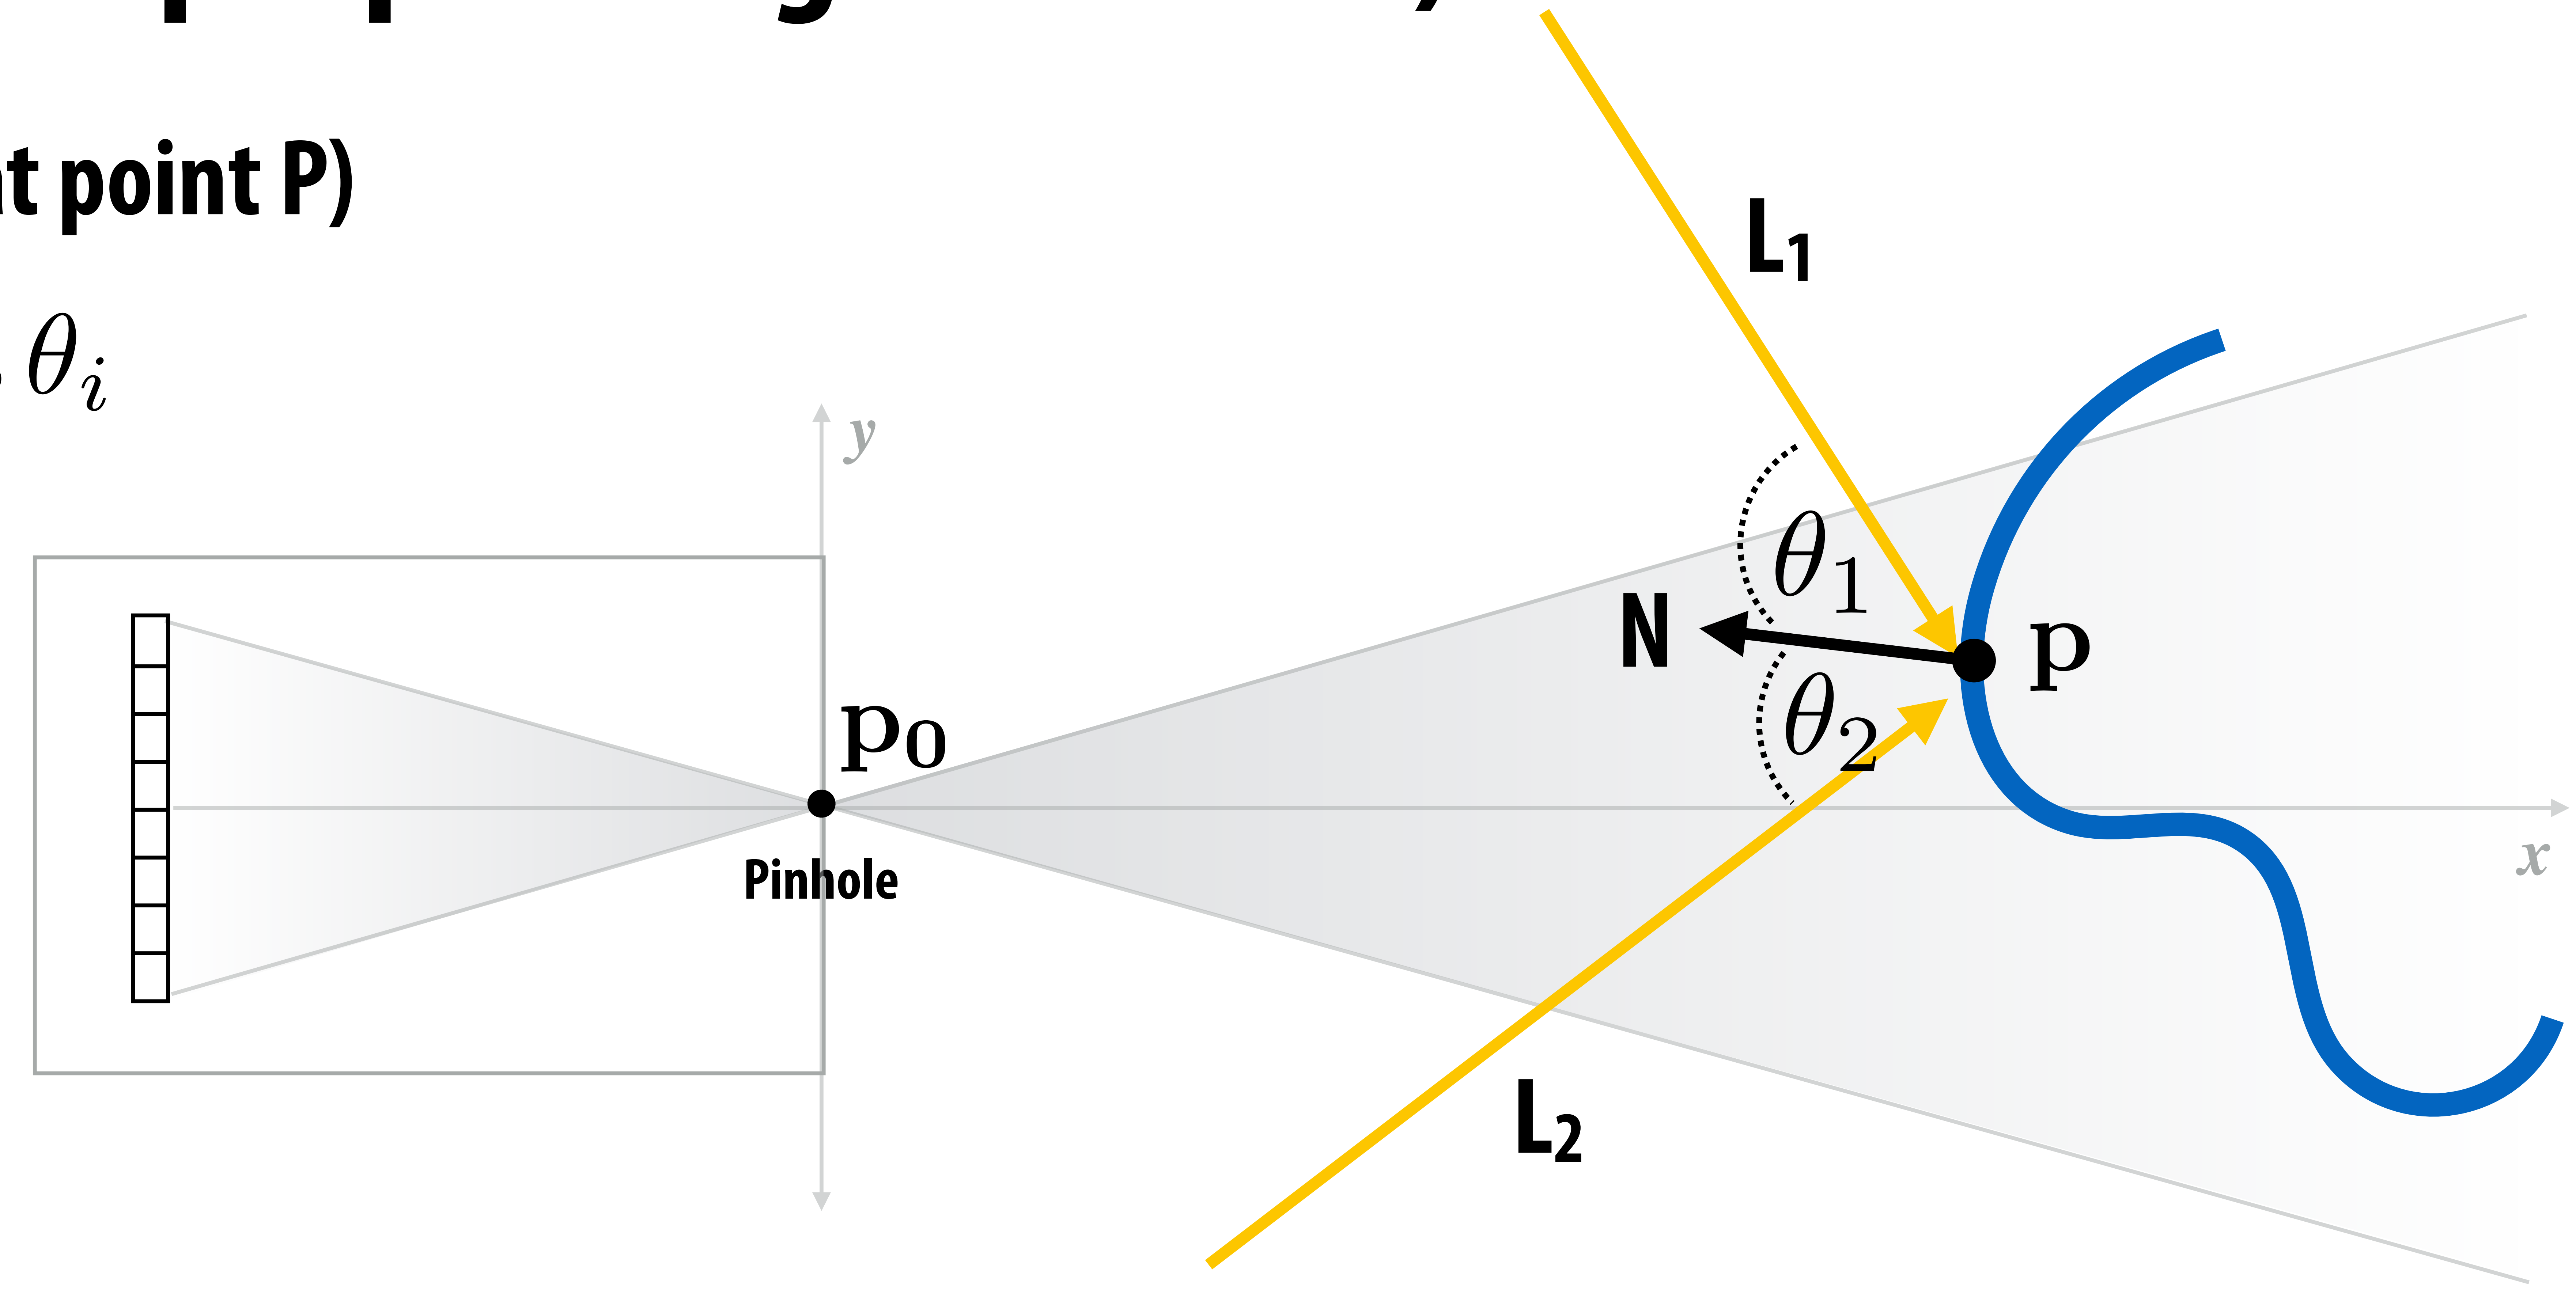

How much light hits the surface at point p?
(from light from all directions!)

(irradiance at point P)

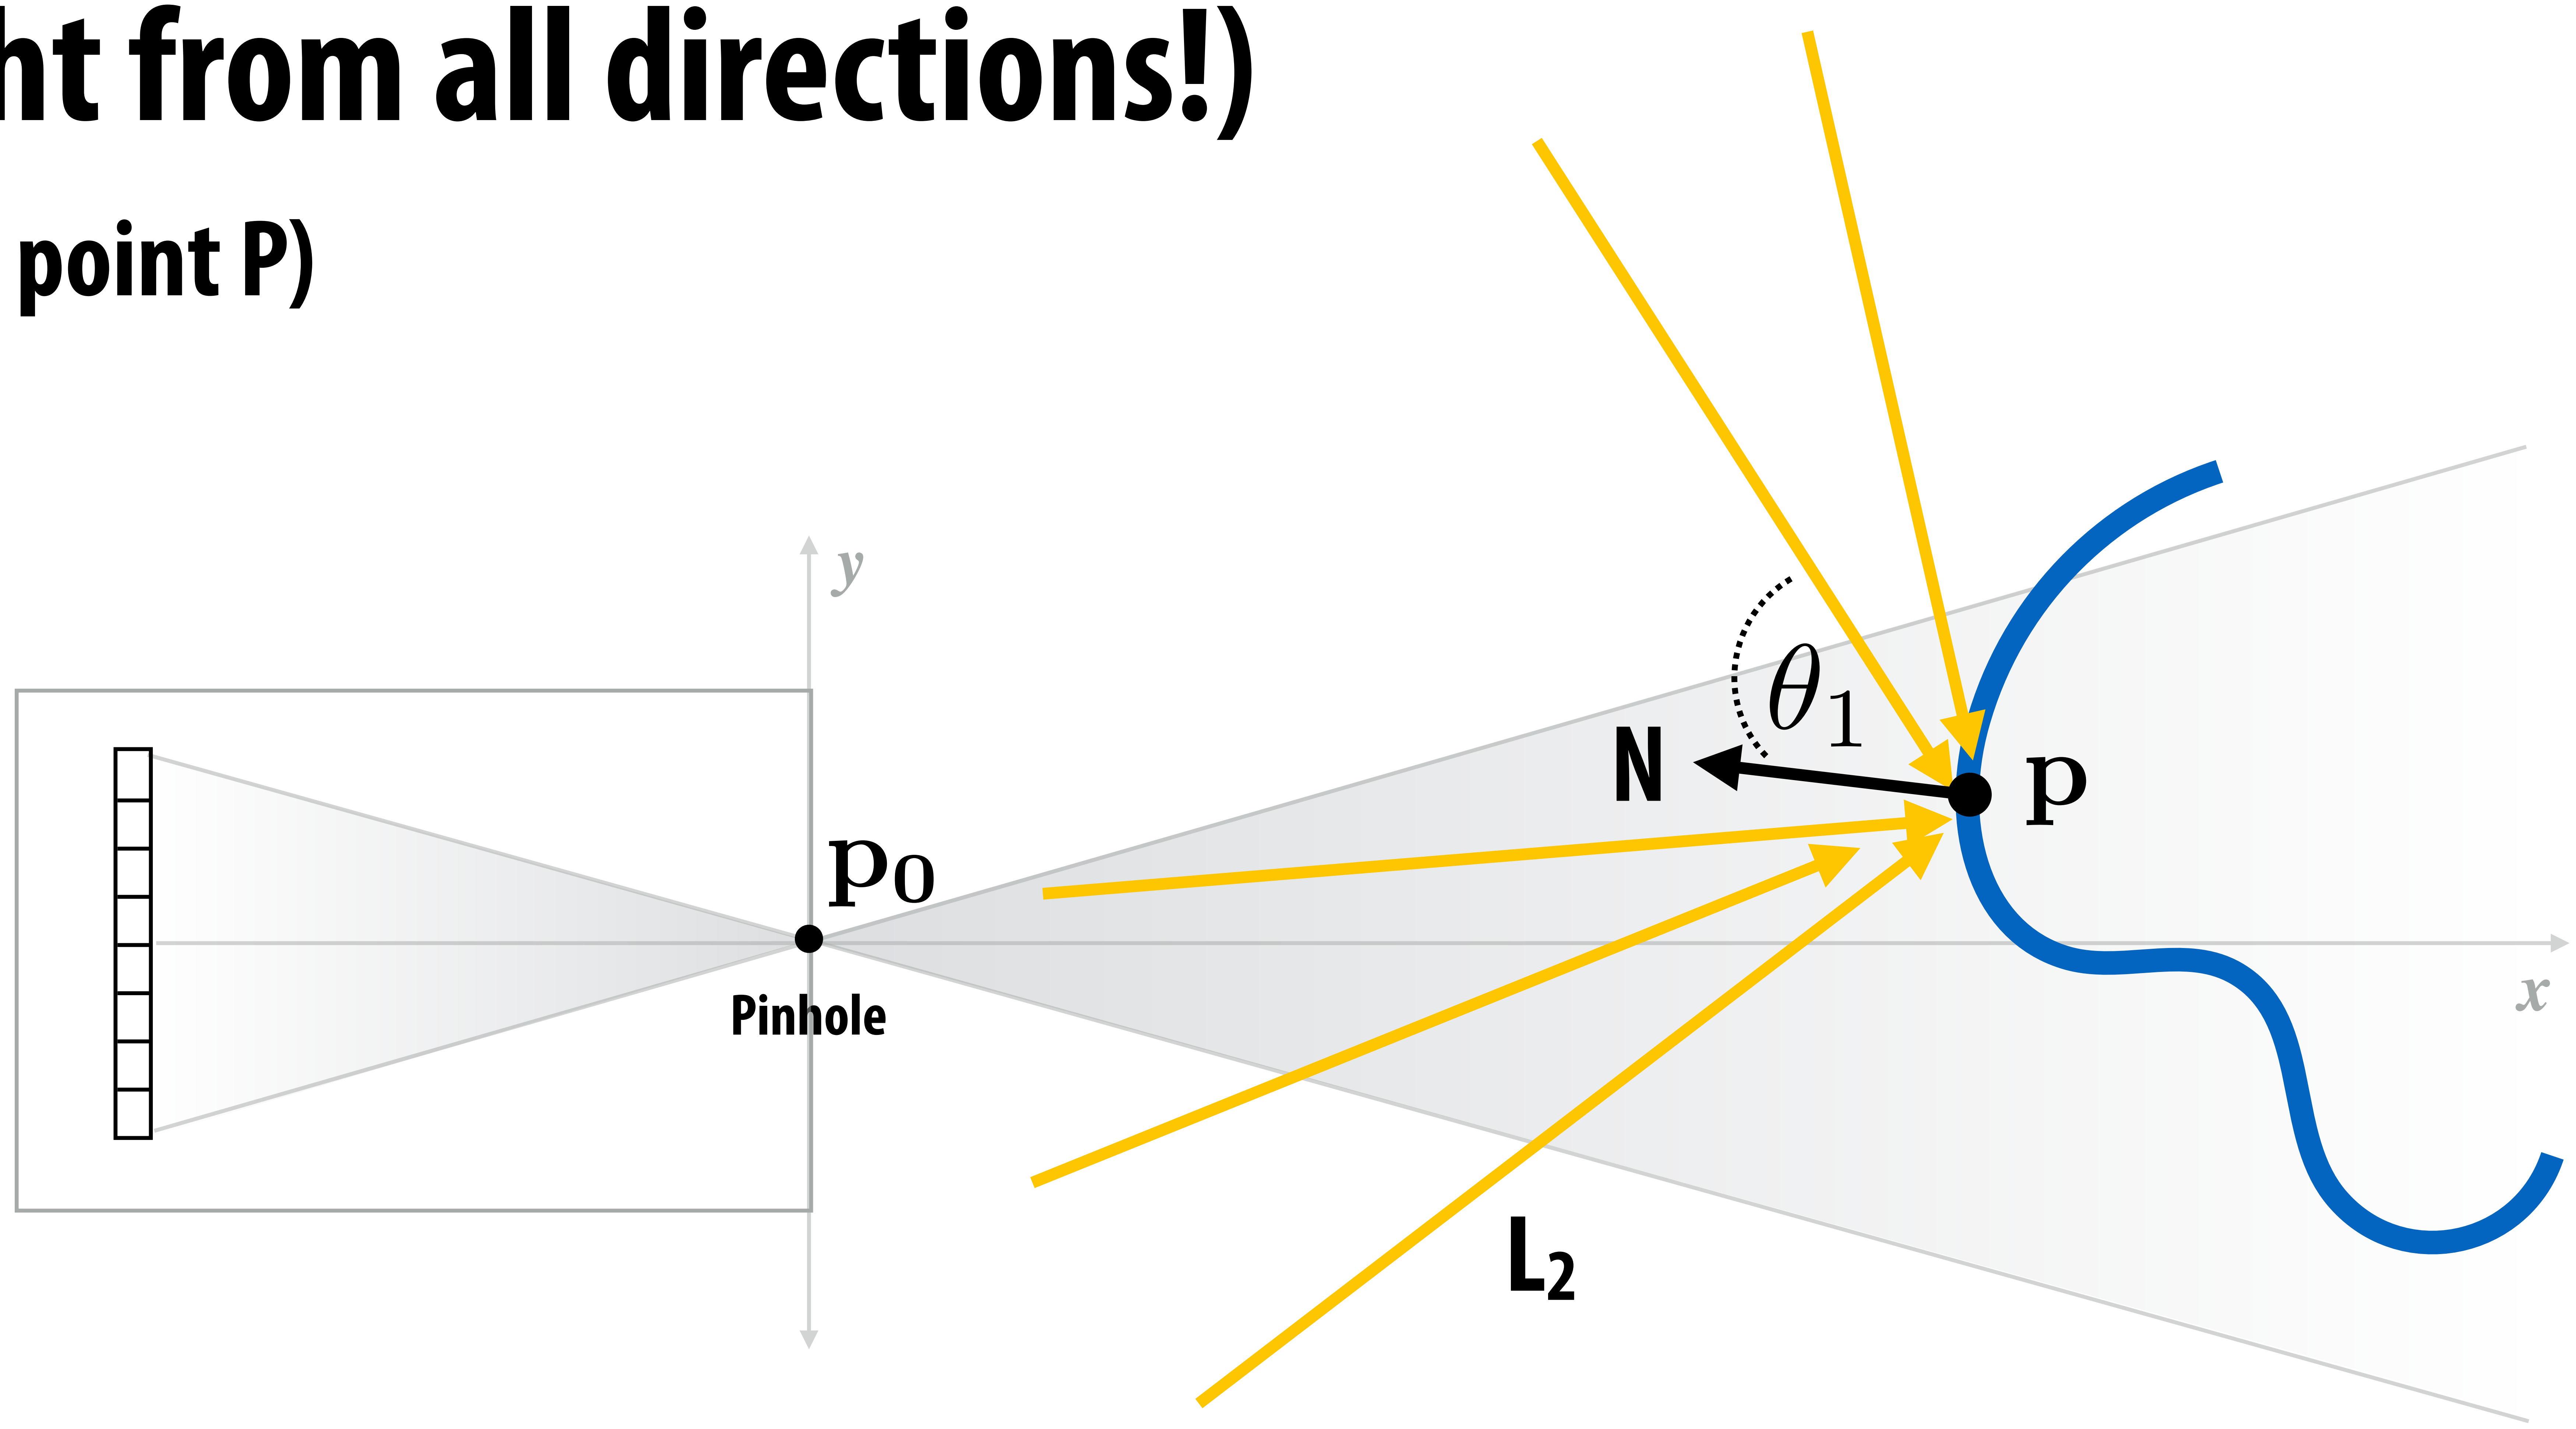

$$\int_{S^2} L_i(\omega_i) \cos\theta_i d\omega = \int_0^{2\pi} \int_0^{\pi} L_i(\omega_i) \cos\theta_i \sin\theta_i d\theta d\phi$$

Irradiance at point X from a uniform area source

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

$$= L \int_{\Omega} \cos\theta \, \mathrm{d}\omega$$

$$= L\tilde{\Omega}$$

Constant
(it's a uniform source)

Total projected solid angle

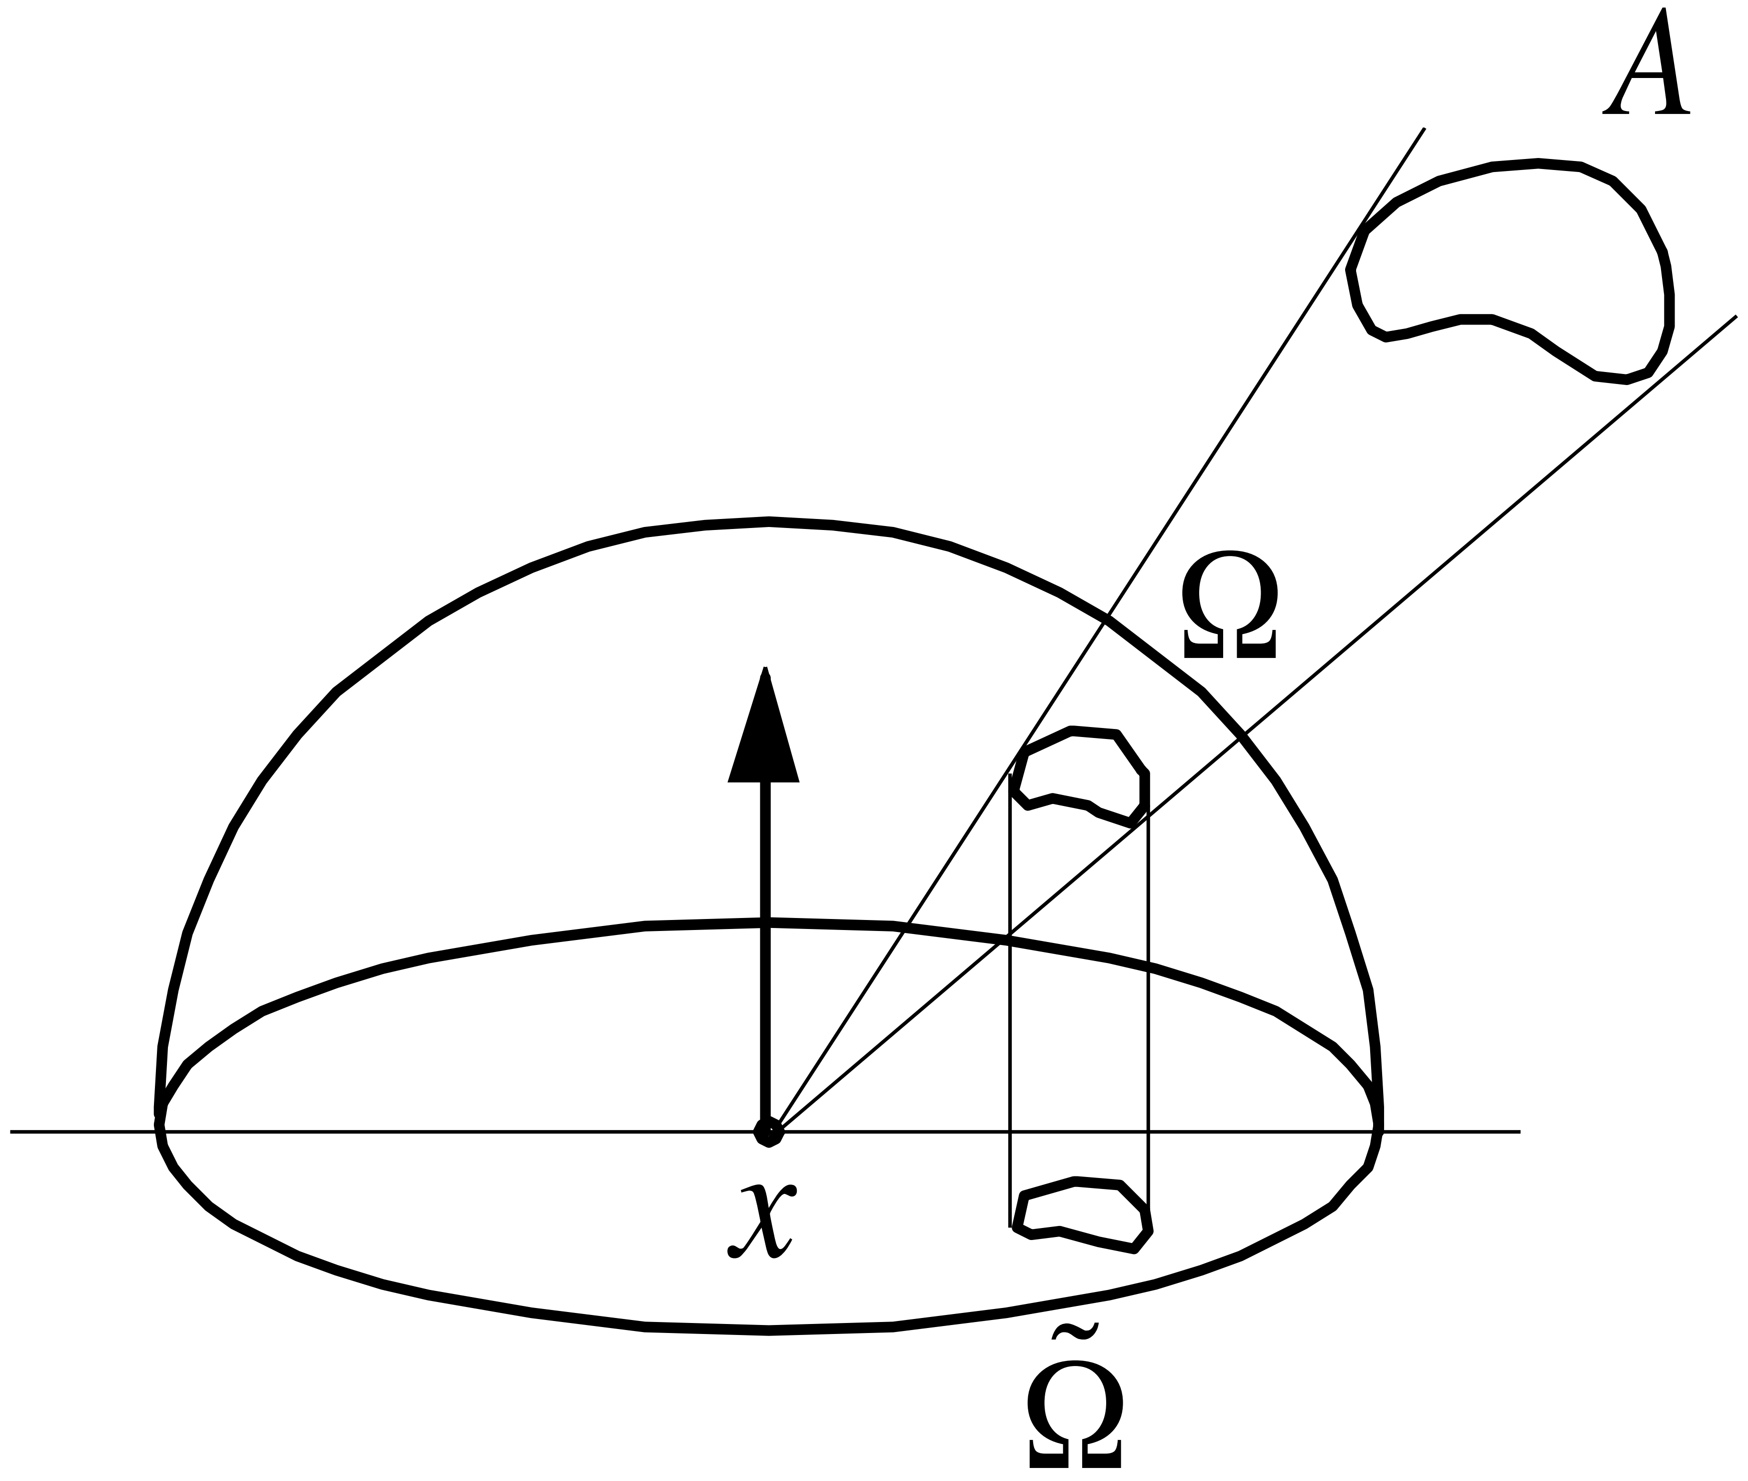

Irradiance at point X from uniform area source

$$E(x) = \int_{H^2} L_i(x,\omega)\ \cos\theta\, d\omega = \underbrace{\int_{A'} L \frac{\cos\theta\cos\theta'}{|x-x'|^2} dA'}_{\text{Reparameterization: now integrate over light source area, instead of solid angle}}$$

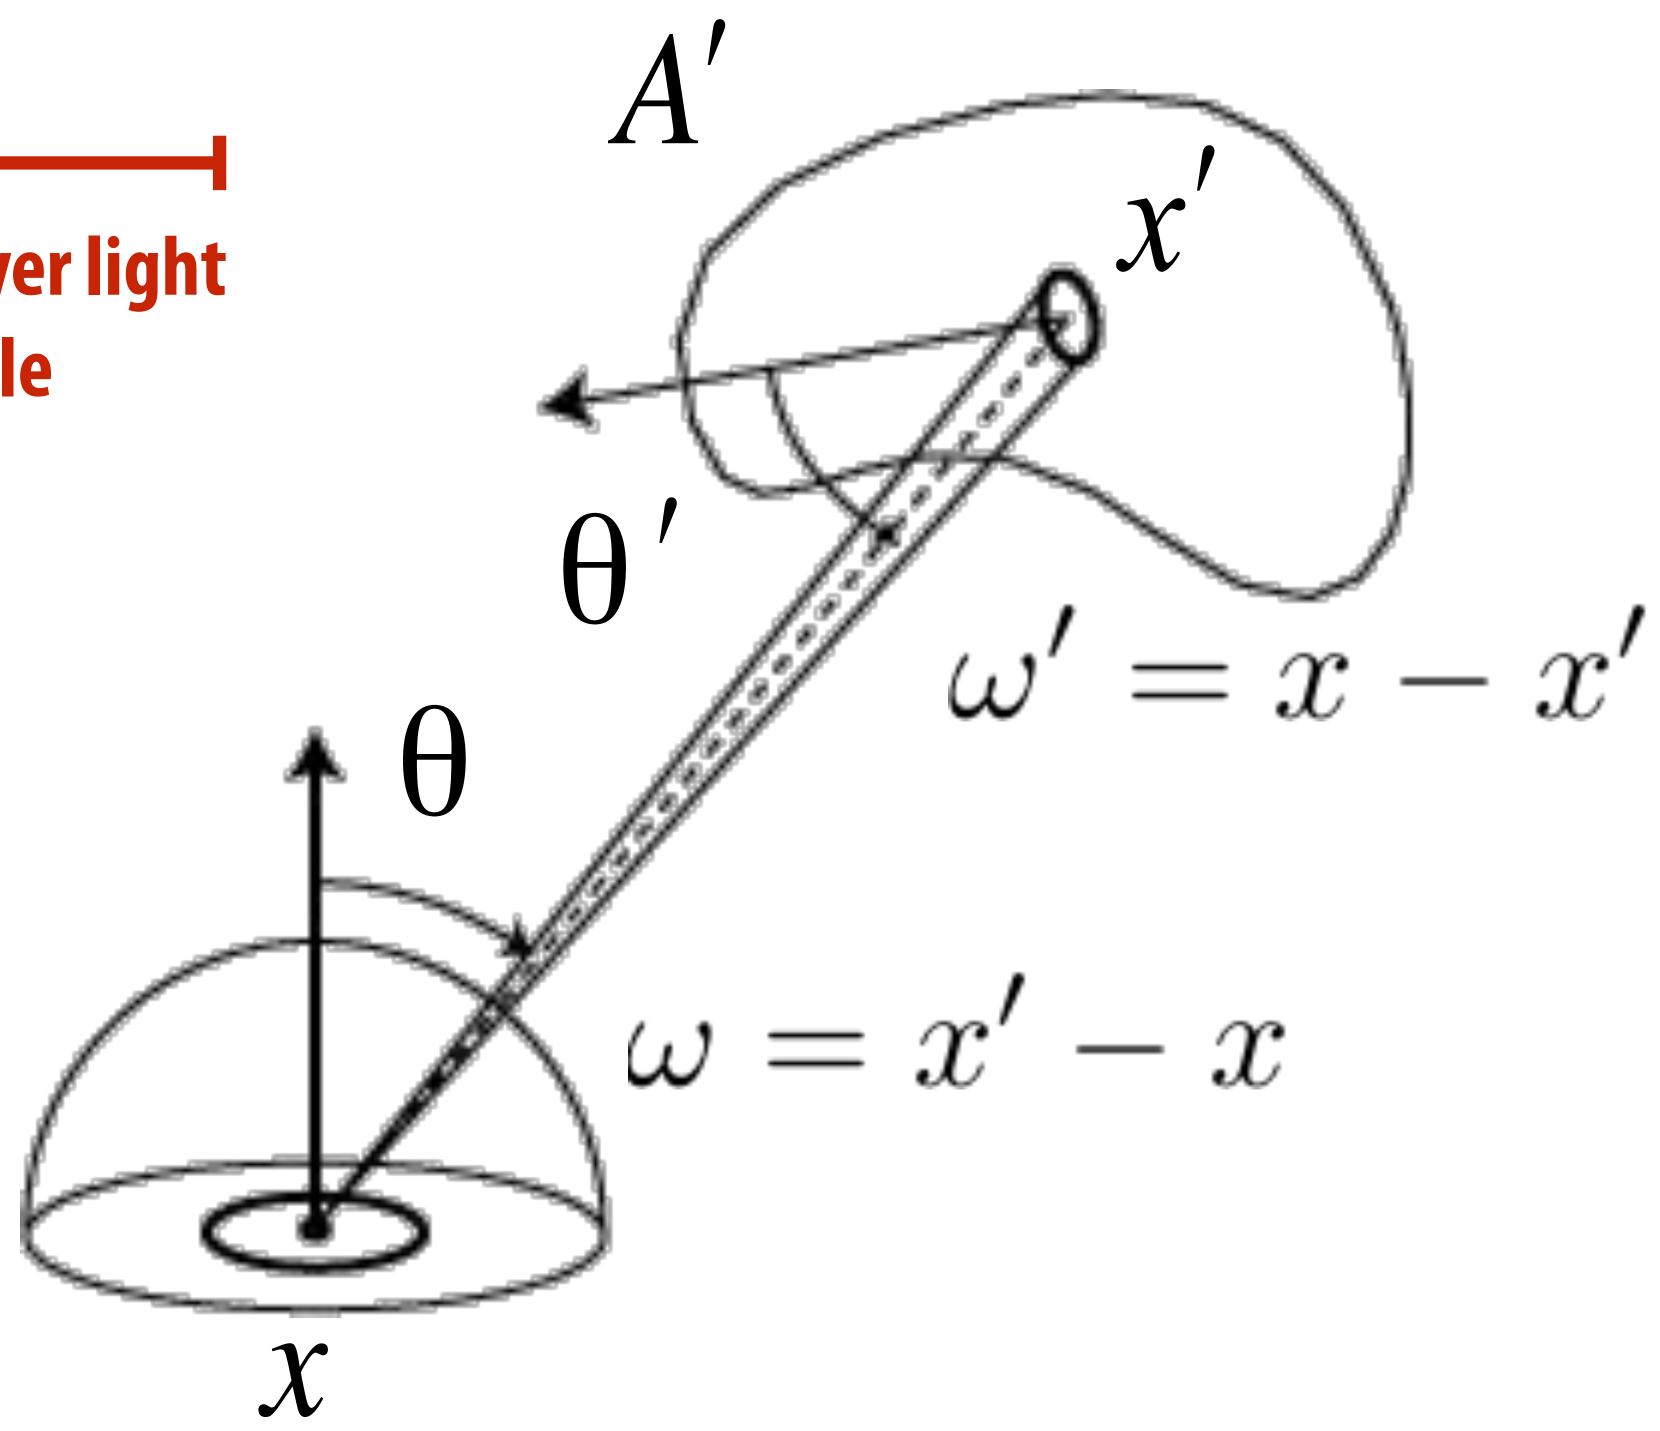

Integral reparameterization:

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

Radiance leaving light from x′ in direction w′ = radiance arriving at surface at x from w.
(assuming that w is pointing at the light)

$$L_i(x,\omega) = L_o(x',\omega') = L$$

Review: materials

Review: the BRDF

Bidirectional Reflectance-Distribution Function

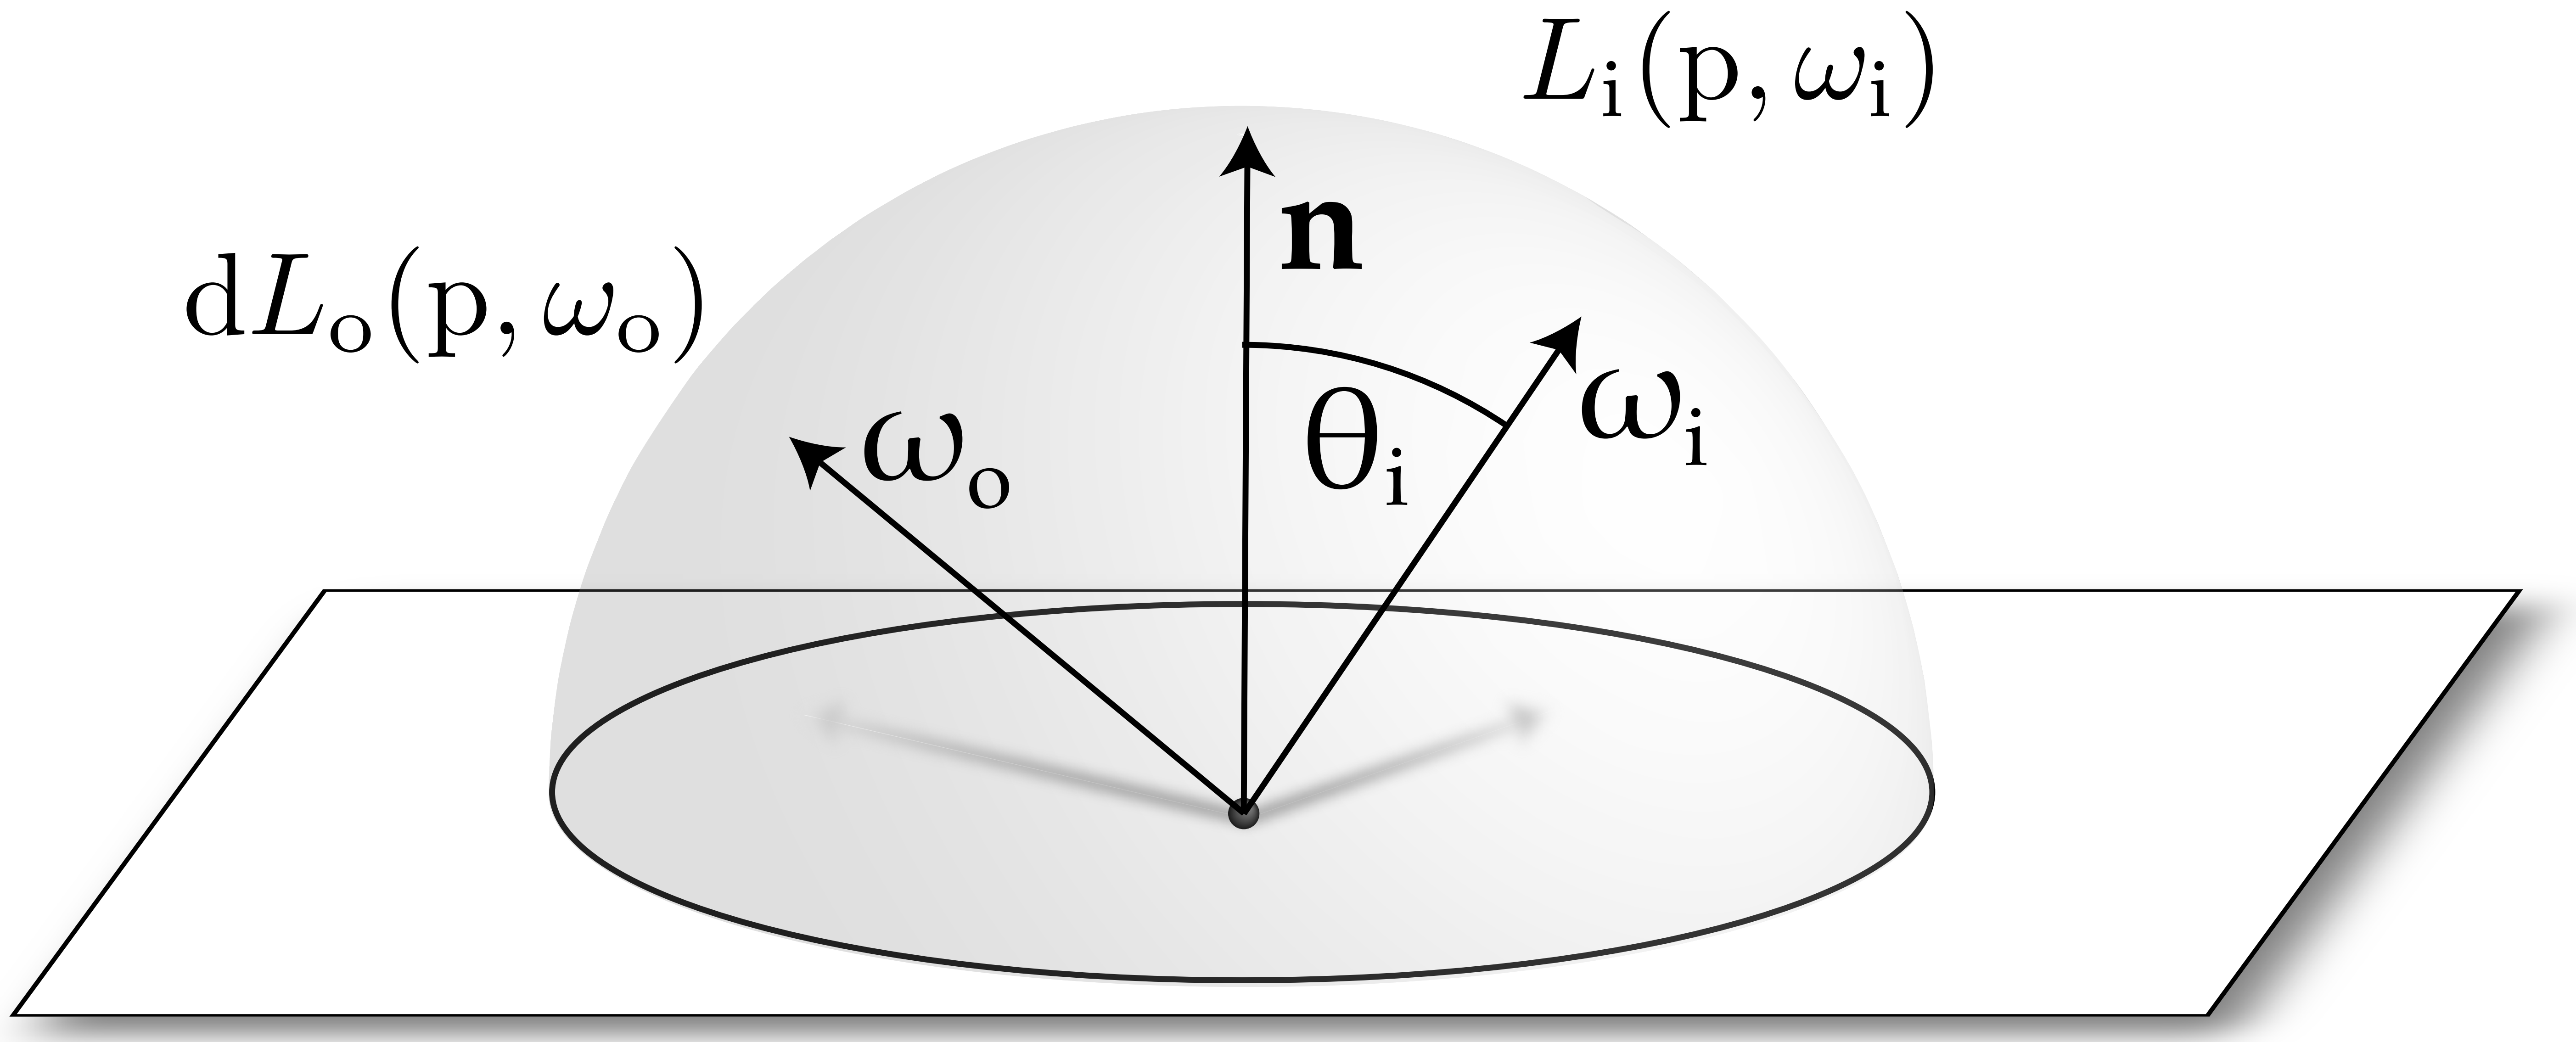

$$f_{\mathrm{r}}(\omega_{\mathrm{i}} \to \omega_{\mathrm{o}}) \equiv \frac{\mathrm{d}L_{\mathrm{o}}(\omega_{\mathrm{o}})}{\mathrm{d}E_{\mathrm{i}}(\omega_{\mathrm{i}})} \quad \left[\frac{1}{sr}\right]$$

"For a given change in incident irradiance, how much does exit radiance change"

Materials: diffuse

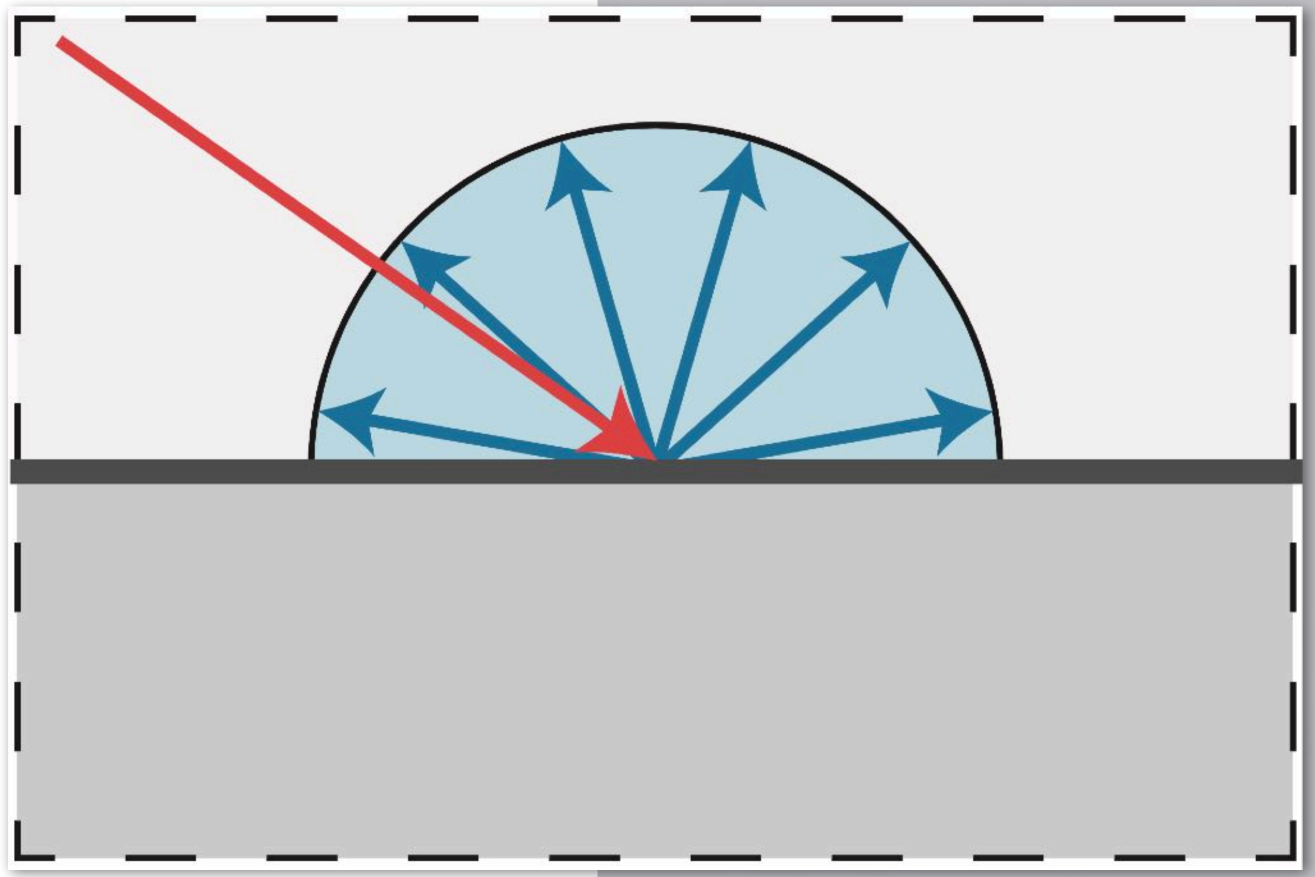

Materials: mirror

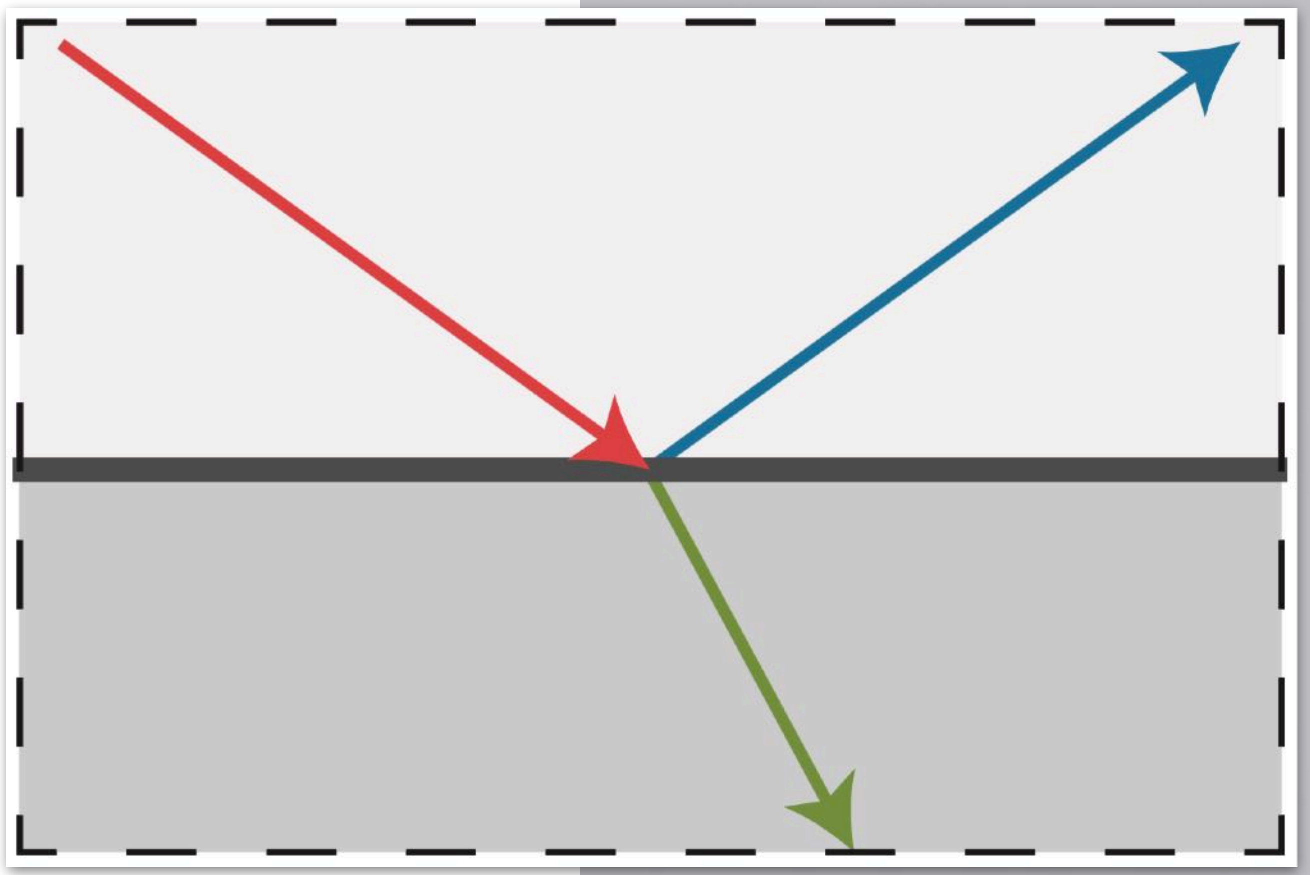

Materials: gold

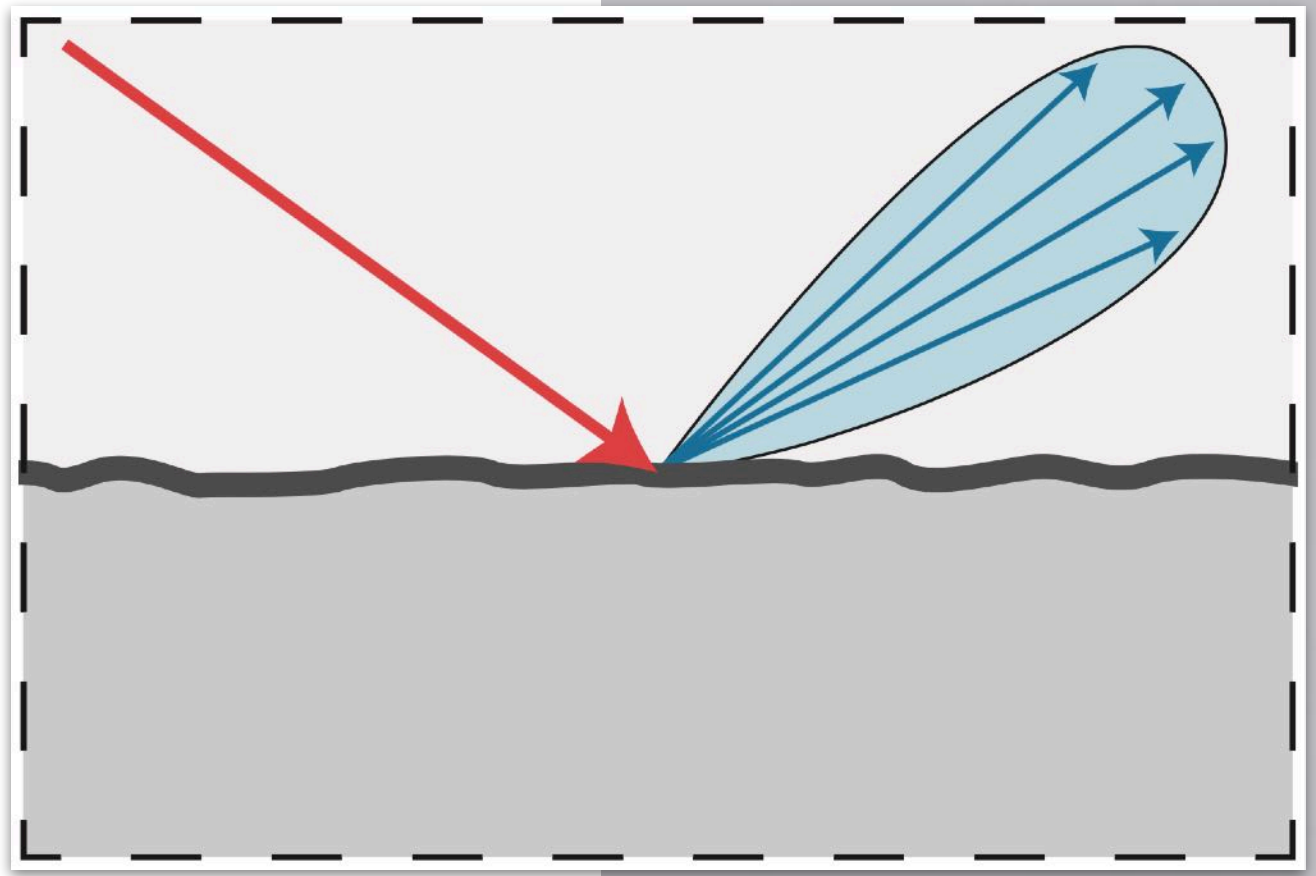

BRDF energy conservation

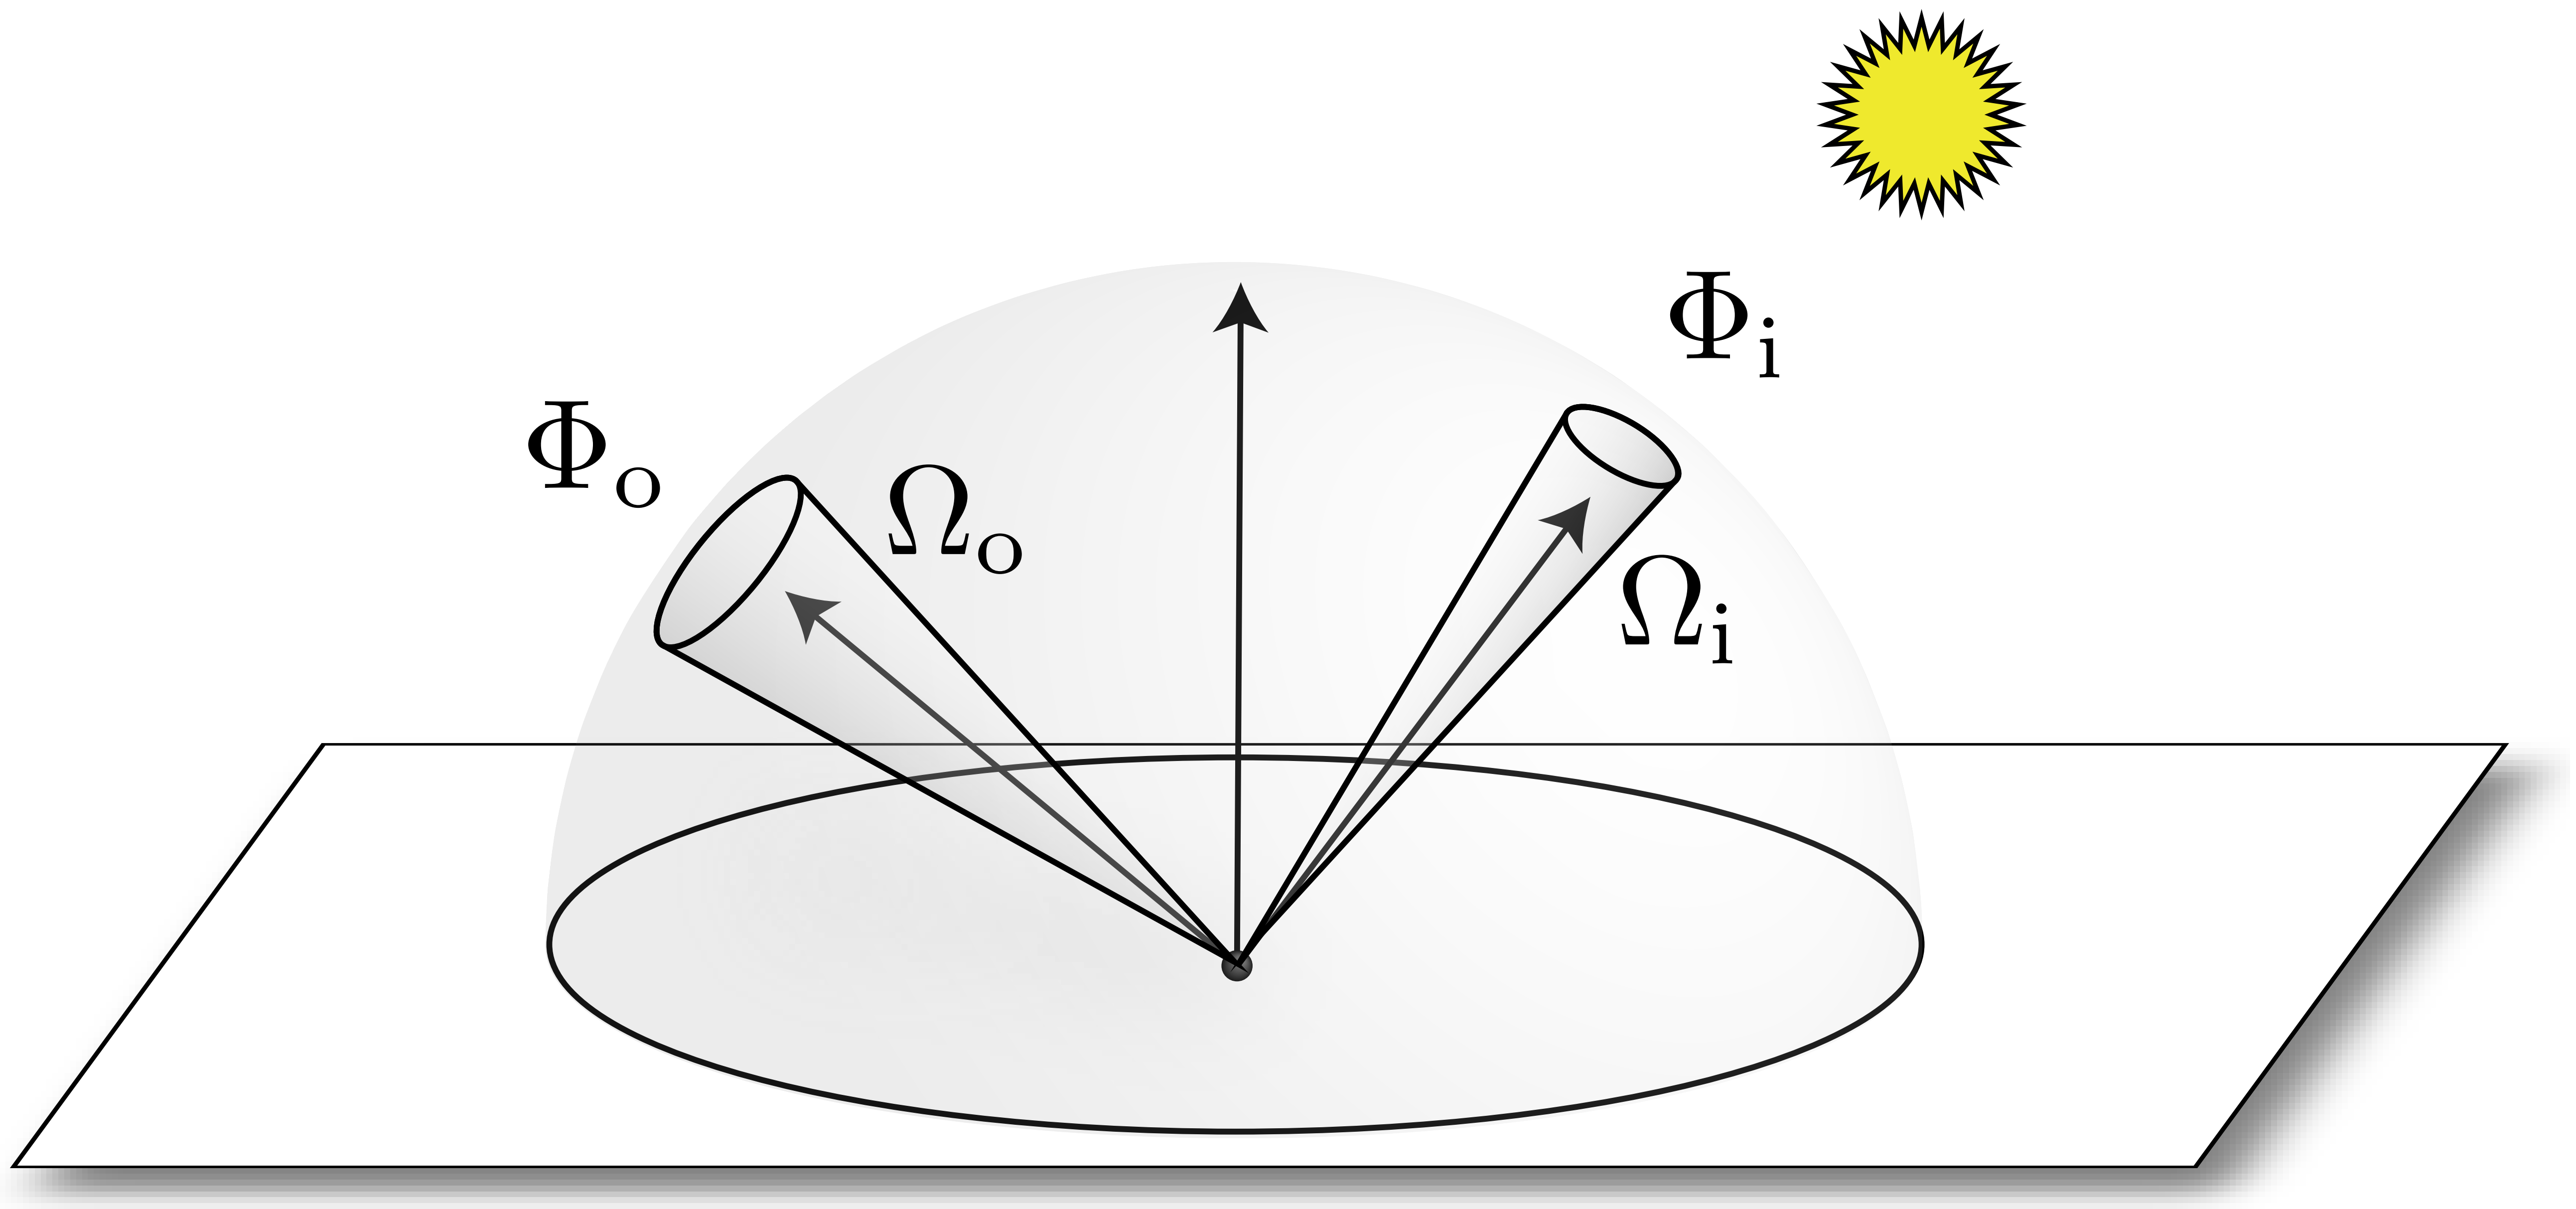

Reflectance $$\rho = \frac{\Phi_\mathrm{o}}{\Phi_\mathrm{i}} = \frac{\int_{\Omega_\mathrm{o}} L_\mathrm{o}(\omega_\mathrm{o}) \cos\theta_\mathrm{o}\, \mathrm{d}\omega_\mathrm{o}}{\int_{\Omega_\mathrm{i}} L_\mathrm{i}(\omega_\mathrm{i}) \cos\theta_\mathrm{i}\, \mathrm{d}\omega_\mathrm{i}}$$

$$0 \leq \rho \leq 1$$

Hemispherical incident radiance

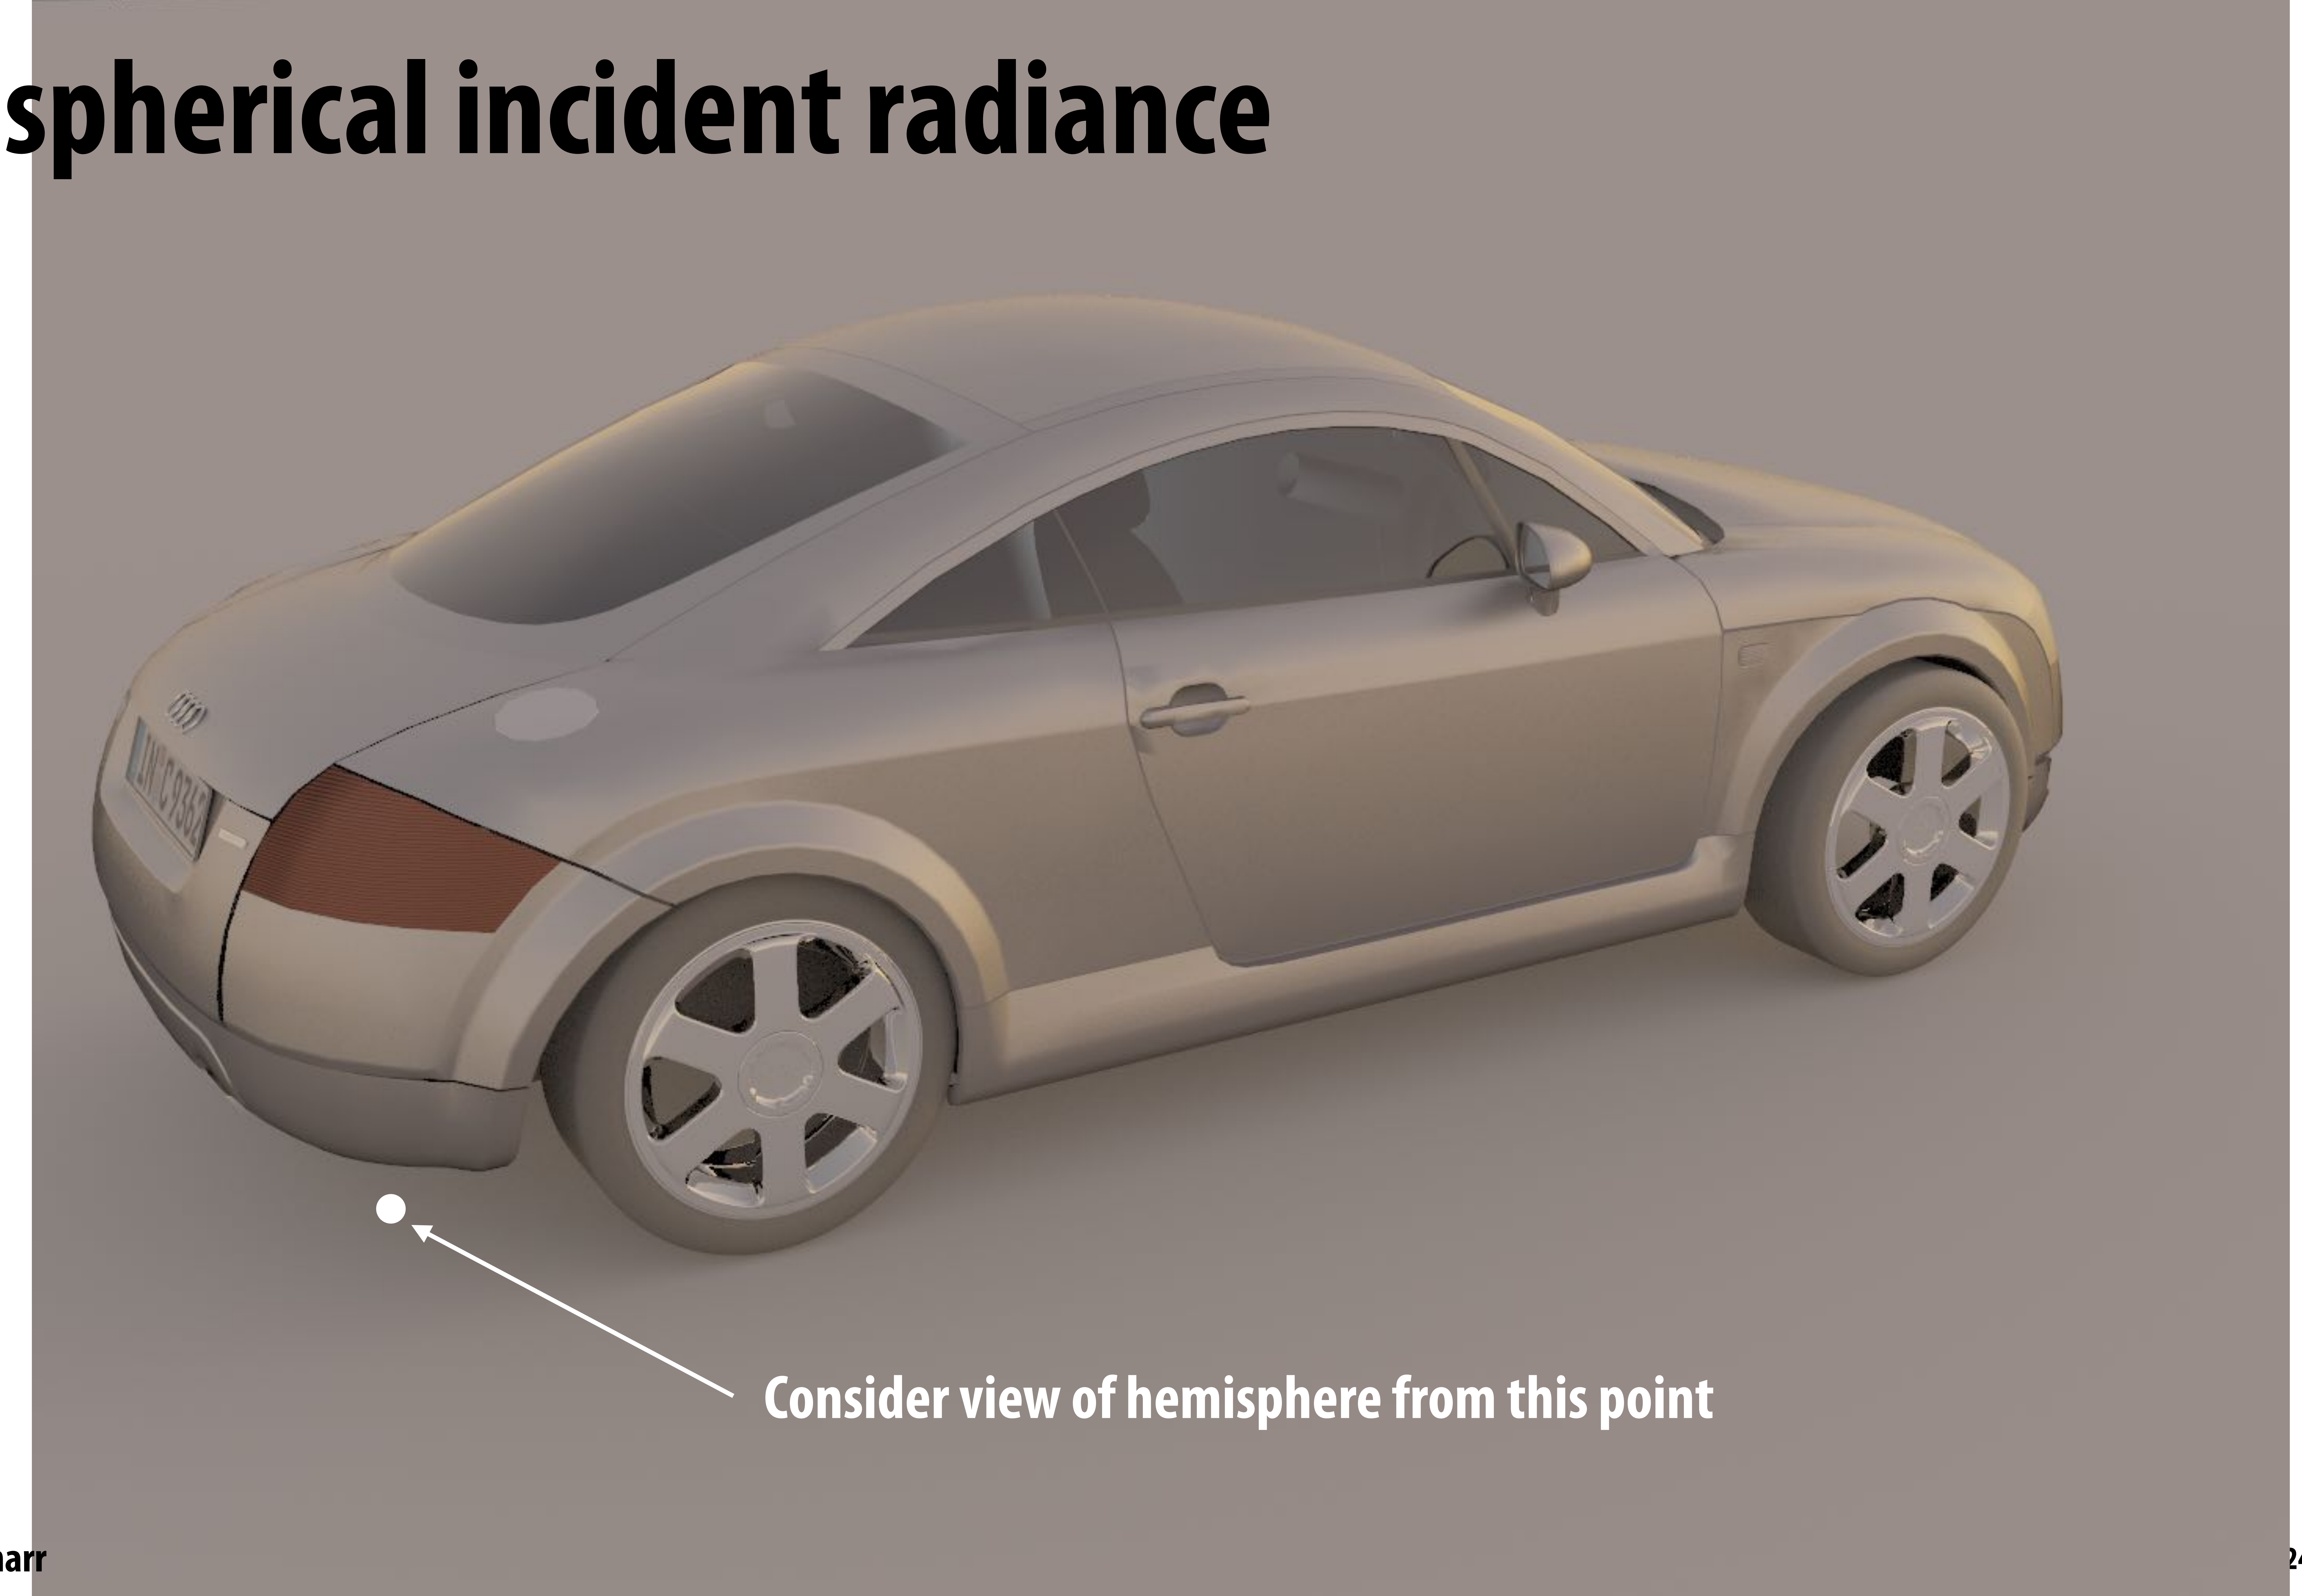

Image credit Matt Pharr

Hemispherical incident radiance

At any point on any surface in the scene, there's an incident radiance field that gives the directional distribution of illumination at the point

Image credit Matt Pharr

Ideal specular reflection

Incident radiance

Exitant radiance

Diffuse reflection

Exitant radiance is the same in all directions

Incident radiance

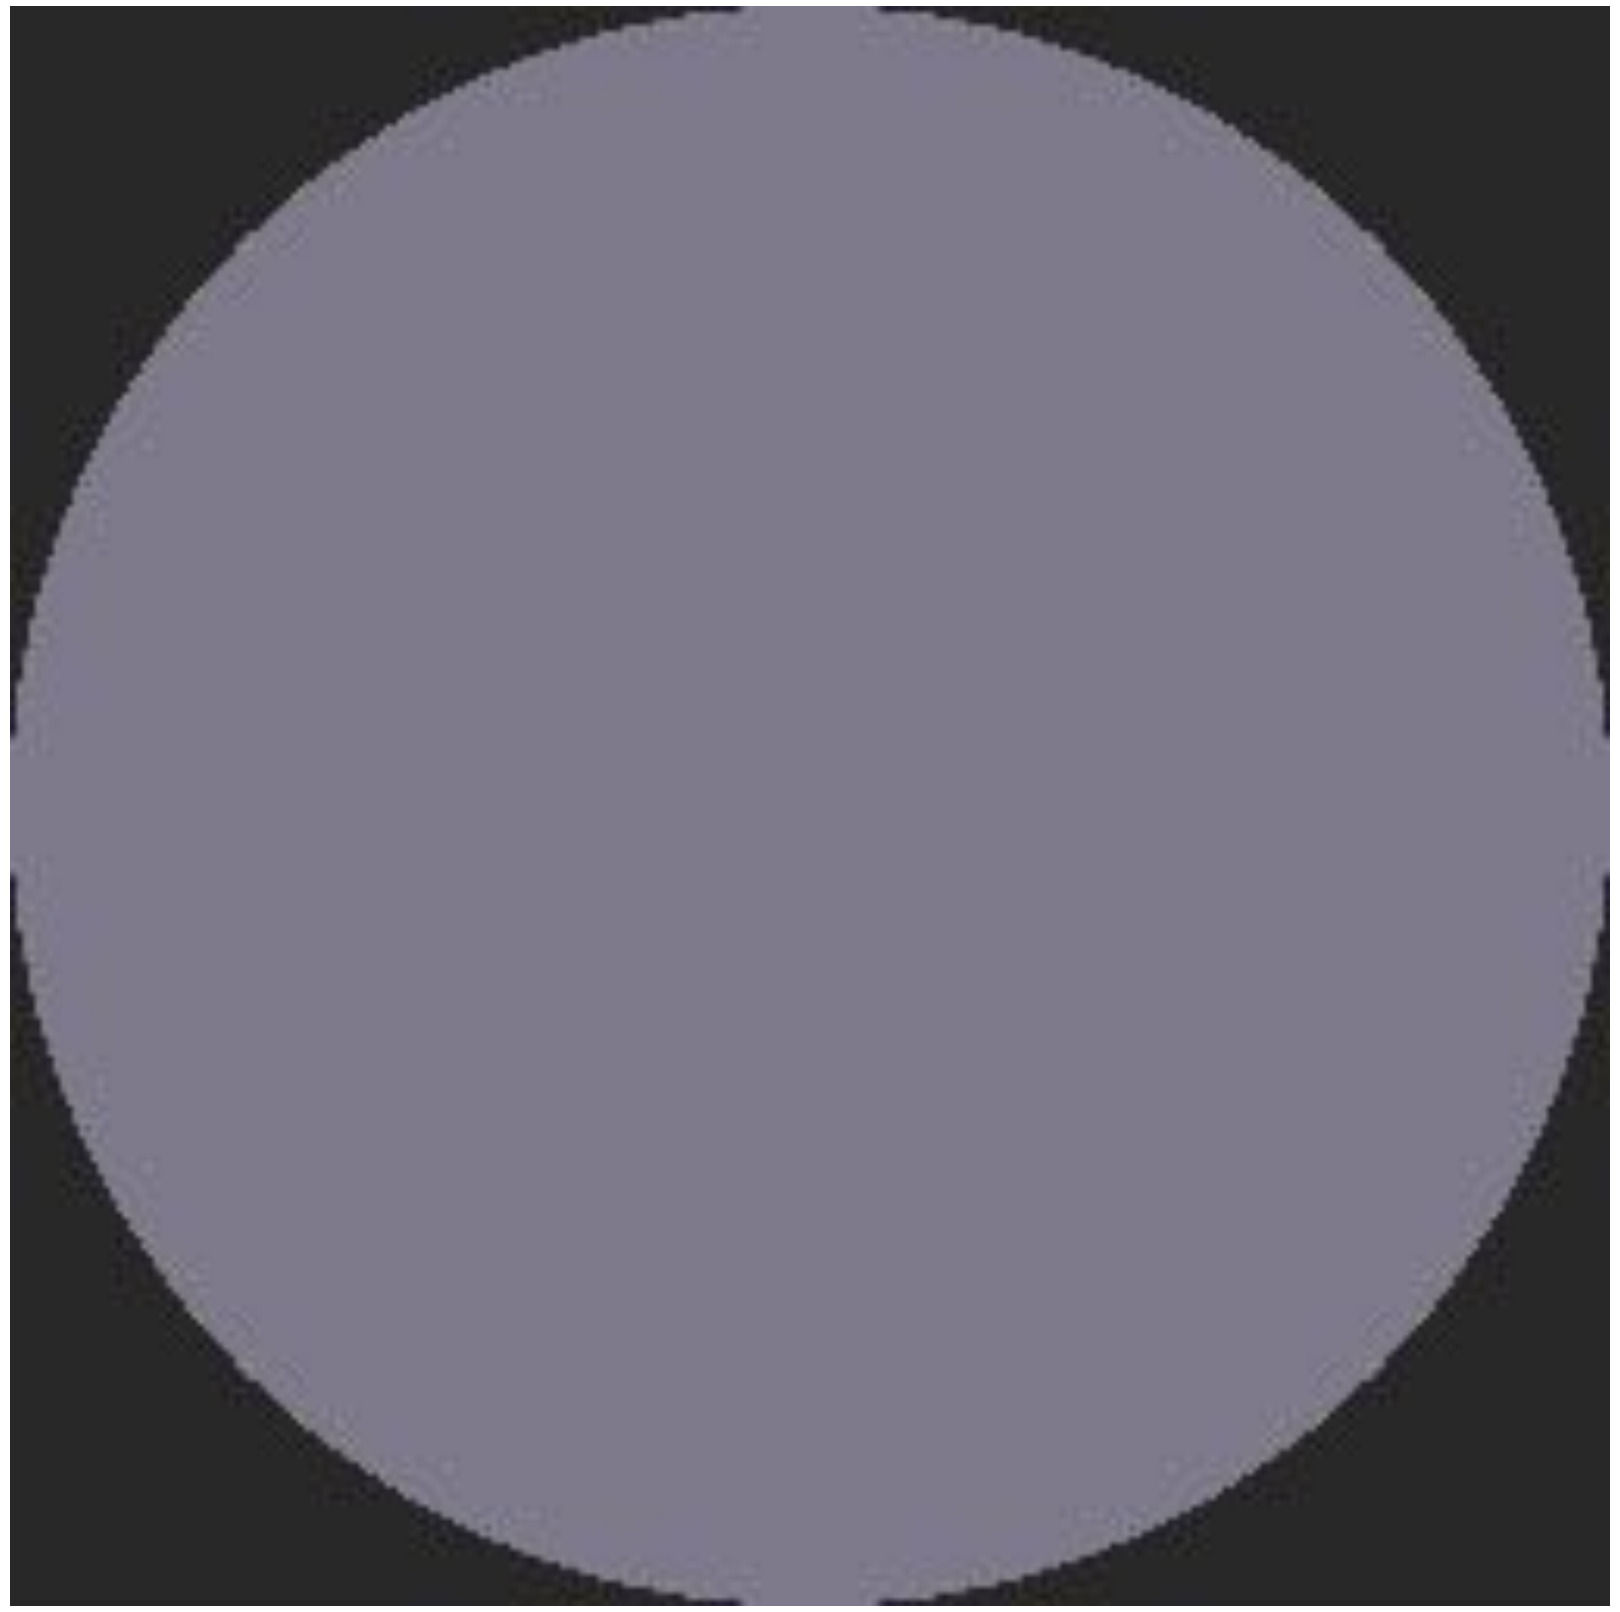

Exitant radiance

Image credit Matt Pharr

Plastic

Incident radiance

Exitant radiance

Copper

Incident radiance

Exitant radiance

Image credit Matt Pharr

BRDF for diffuse surface with albedo ρ

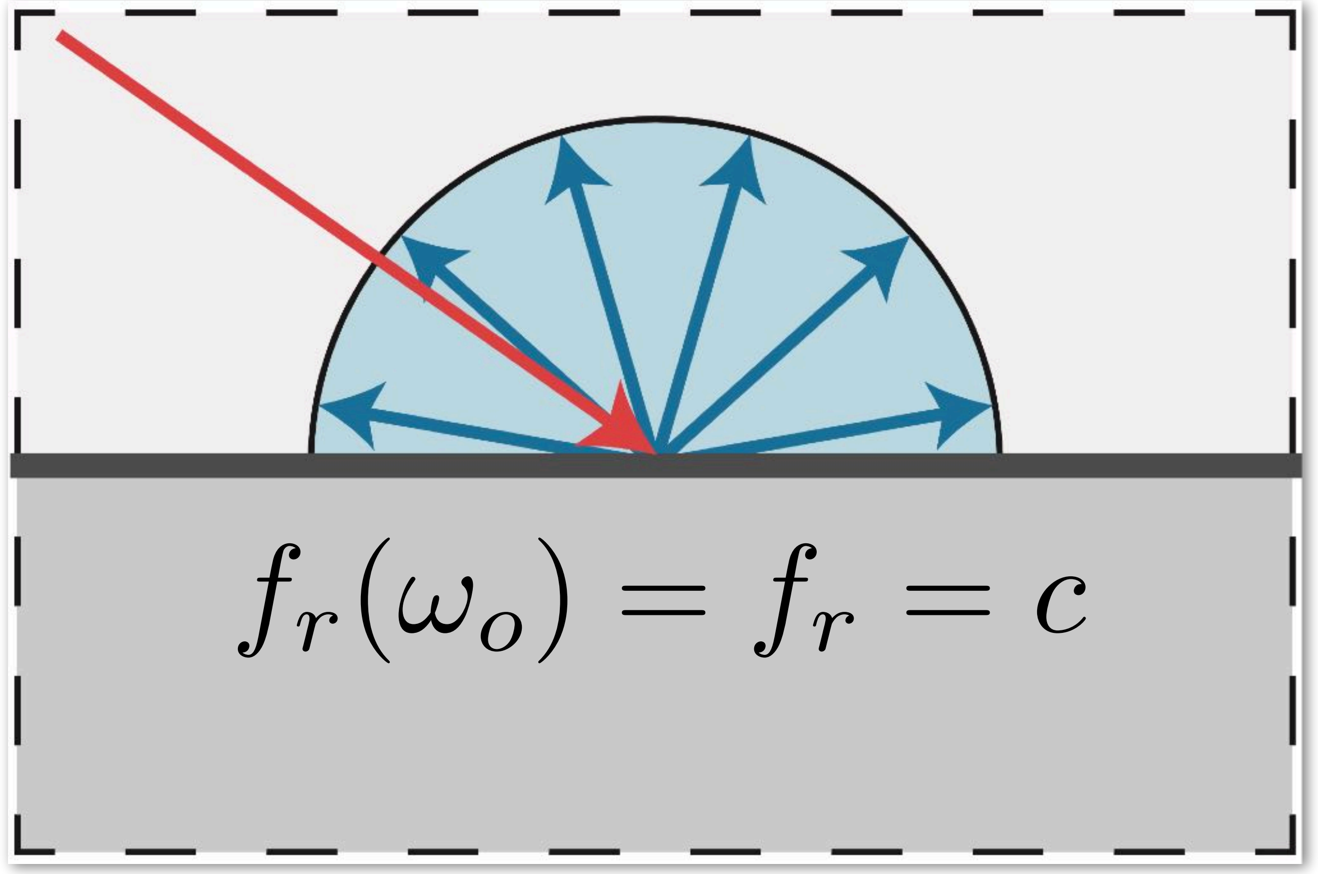

$$L_o(w_o) = \int_{H^2} f_r L_i(w_i)\ \cos\theta_i \, d\omega_i$$

$$= f_r \int_{H^2} L_i(w_i)\ \cos\theta_i \, d\omega_i$$

$$= \boxed{f_r E}$$

Radiance in outgoing direction

Let's call the overall reflectance (albedo) of the surface ρ

Total outgoing surface irradiance

$$\rho E = \int_{H^2} \boxed{f_r E} \cos\theta_o \, d\omega_o$$

$$\rho = f_r \int_{H^2} \cos\theta_o \, d\omega_o$$

$$\rho = f_r \pi$$

$$f_r = \boxed{\frac{\rho}{\pi}}$$

Given a desired ρ BRDF should be the constant $\dfrac{\rho}{\pi}$

A bit more on materials from last time
(Returning to last lecture's slides…)

Transmission
Refraction
Subsurface scattering

Numerical Integration

Many examples of needing to compute integrals already in this lecture

$$E(x) = \int_{H^2} L \, \cos\theta \, d\omega = \int_{A'} L \frac{\cos\theta \cos\theta'}{|x - x'|^2} dA'$$

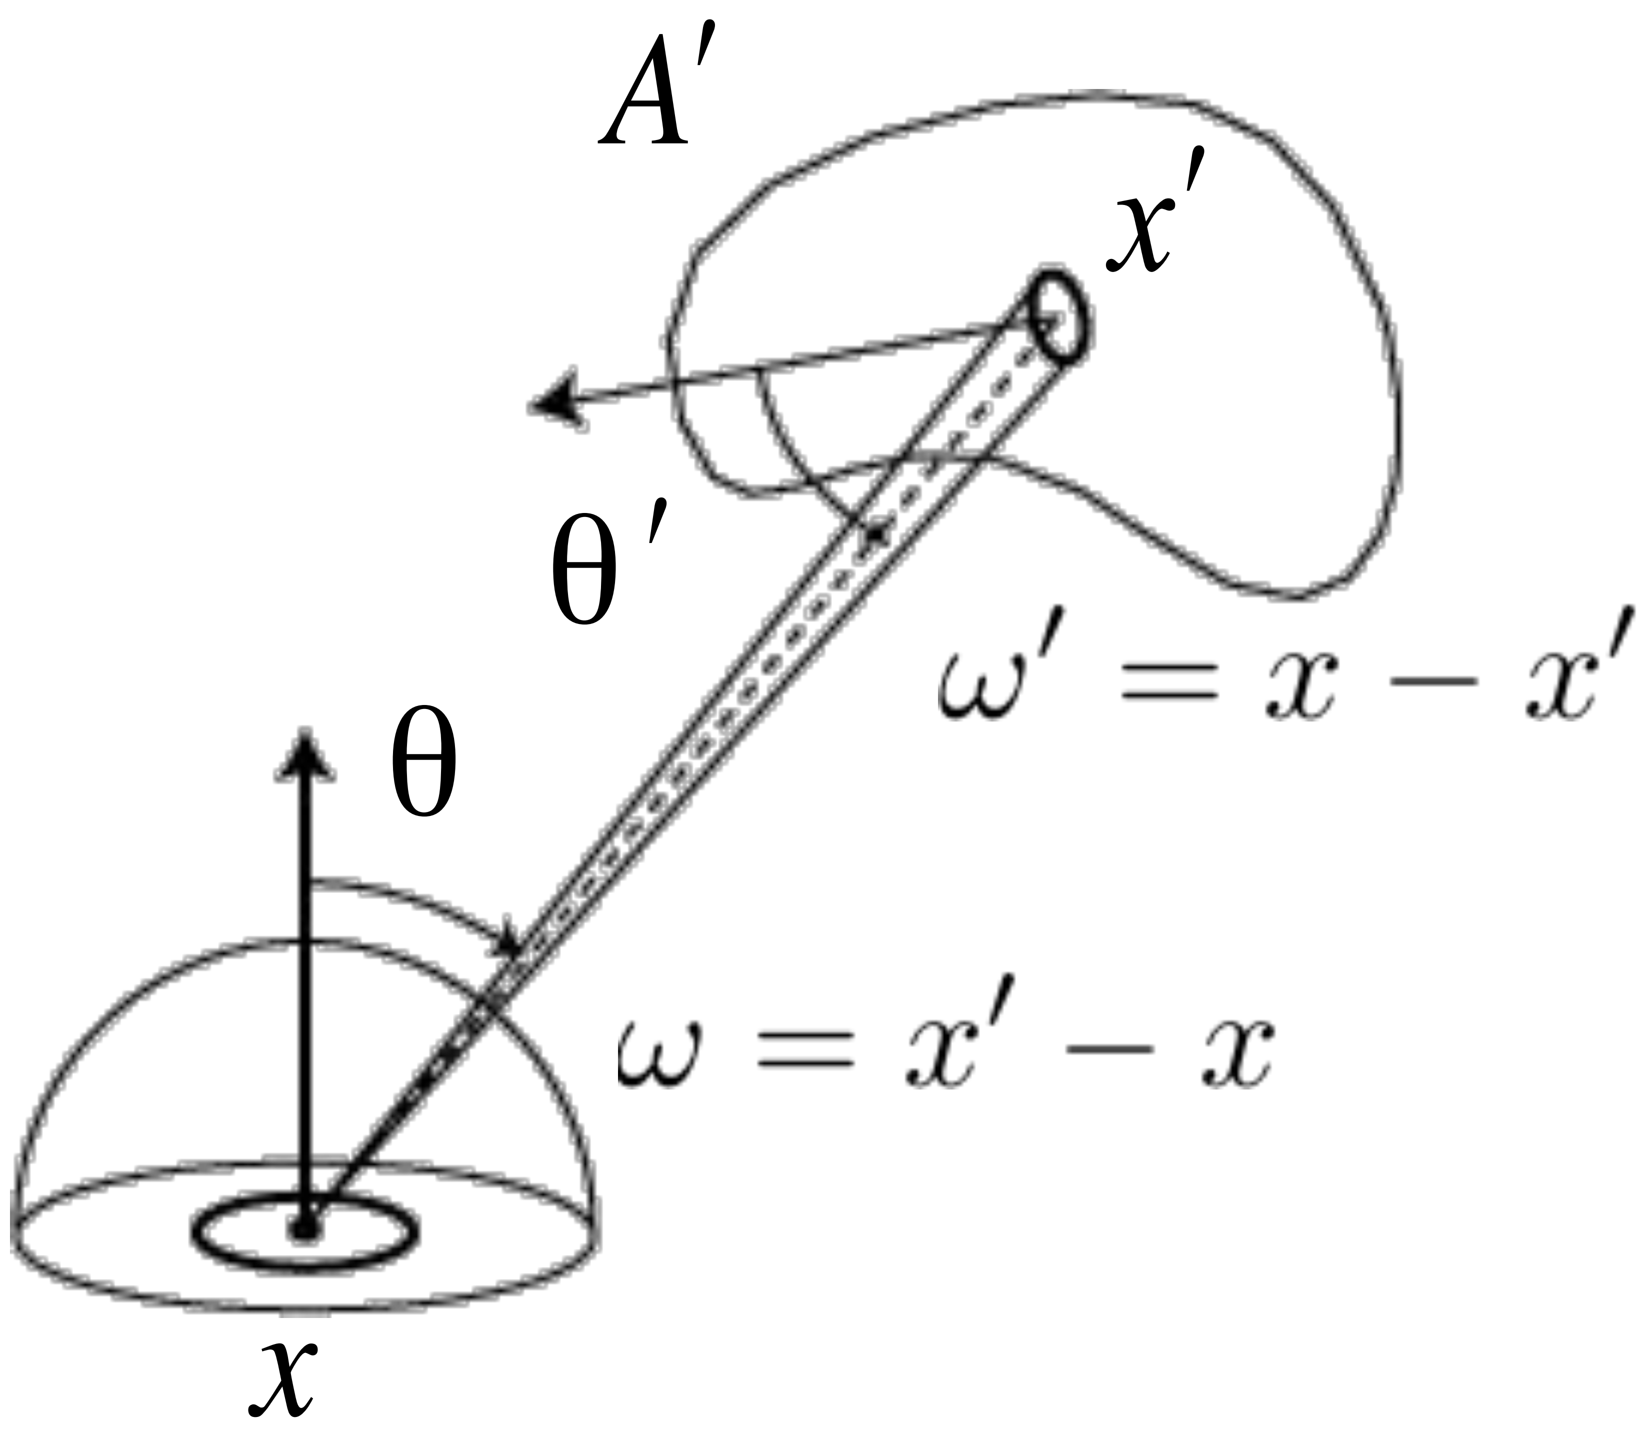

Review: fundamental theorem of calculus

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

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

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

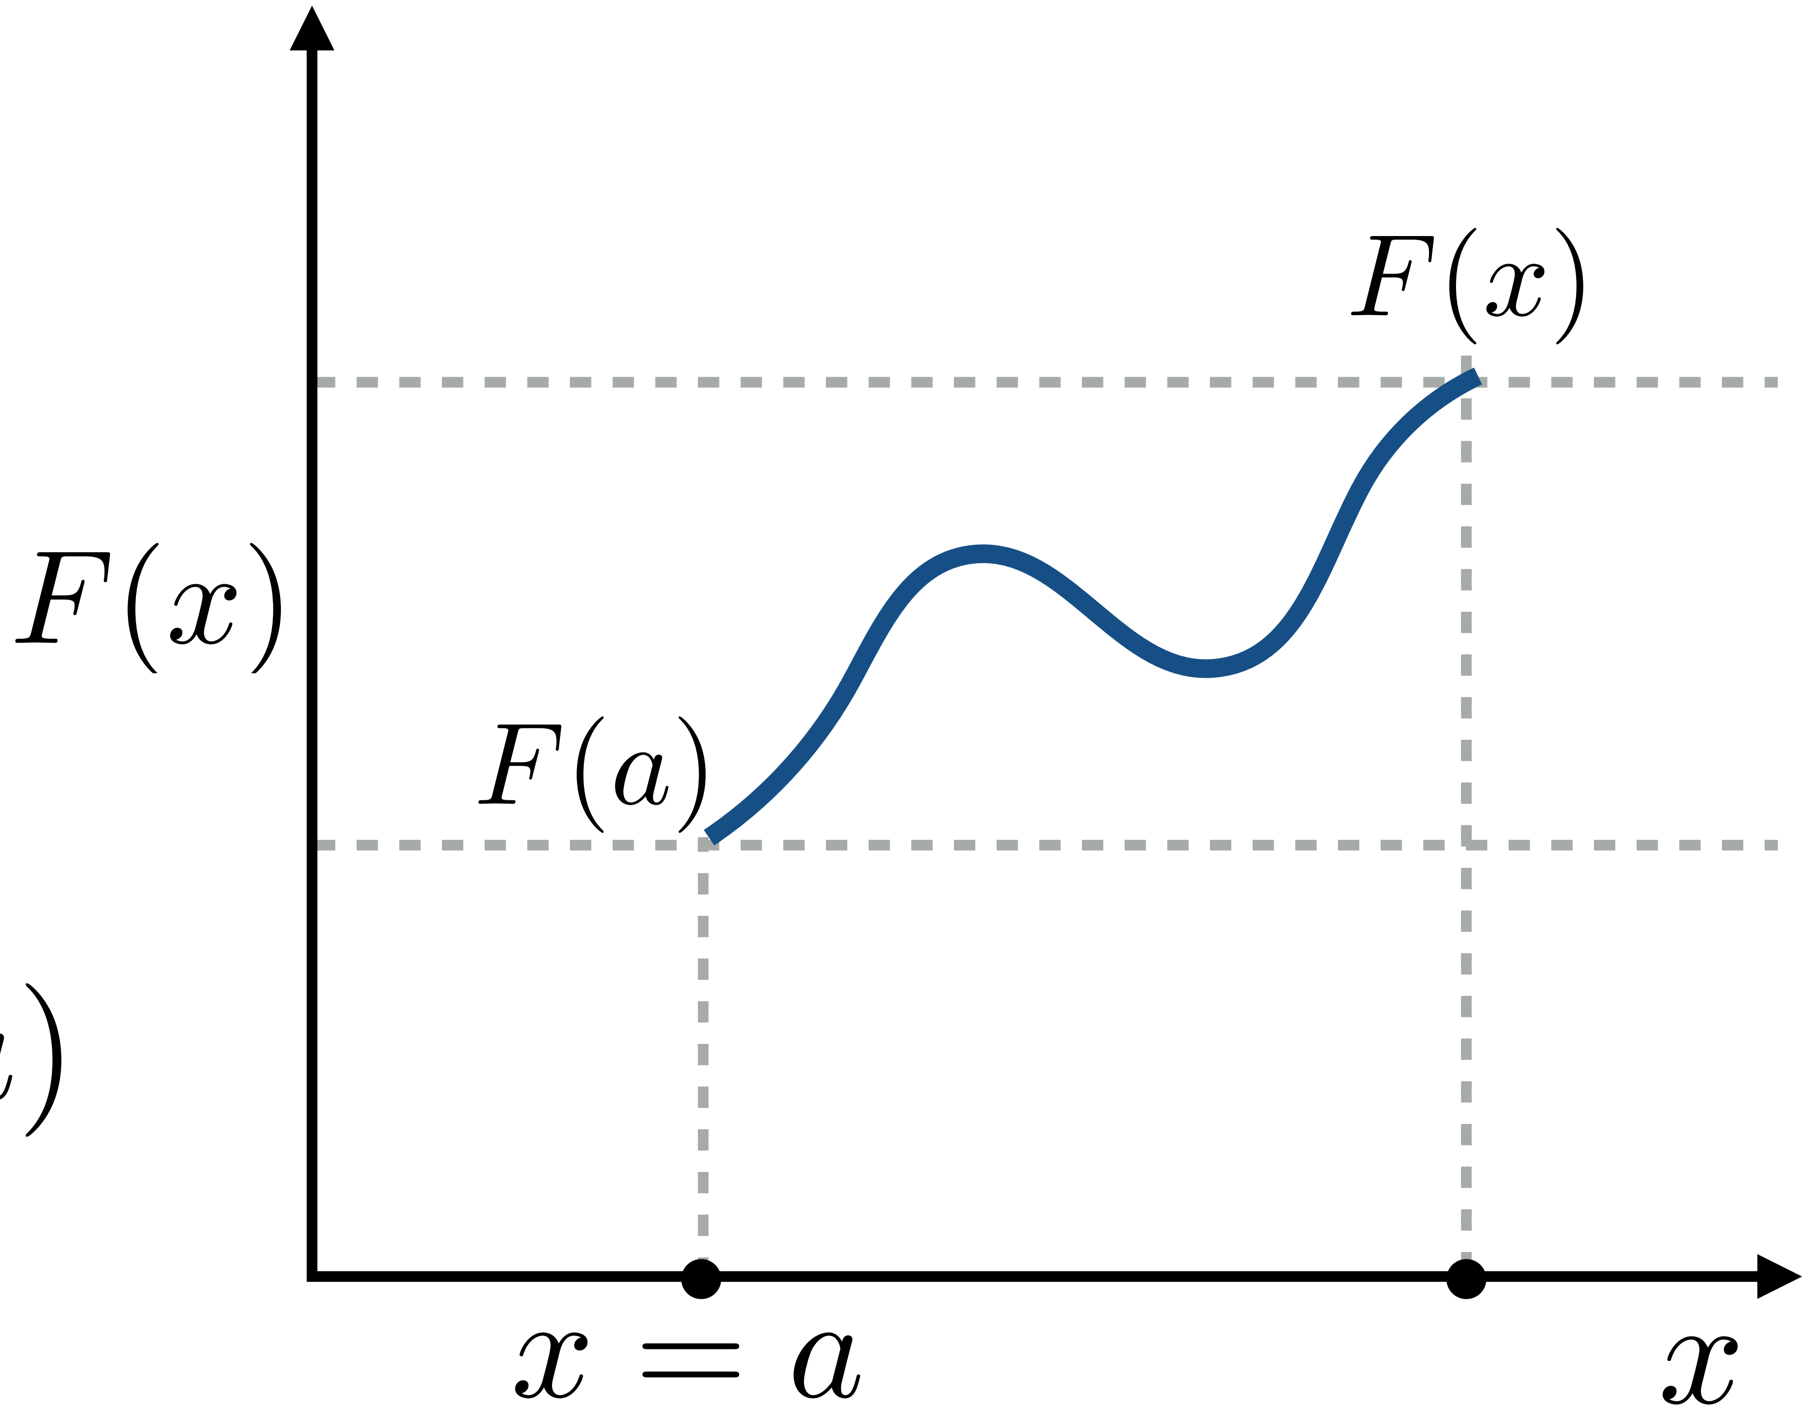

Definite integral as "area under curve"

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

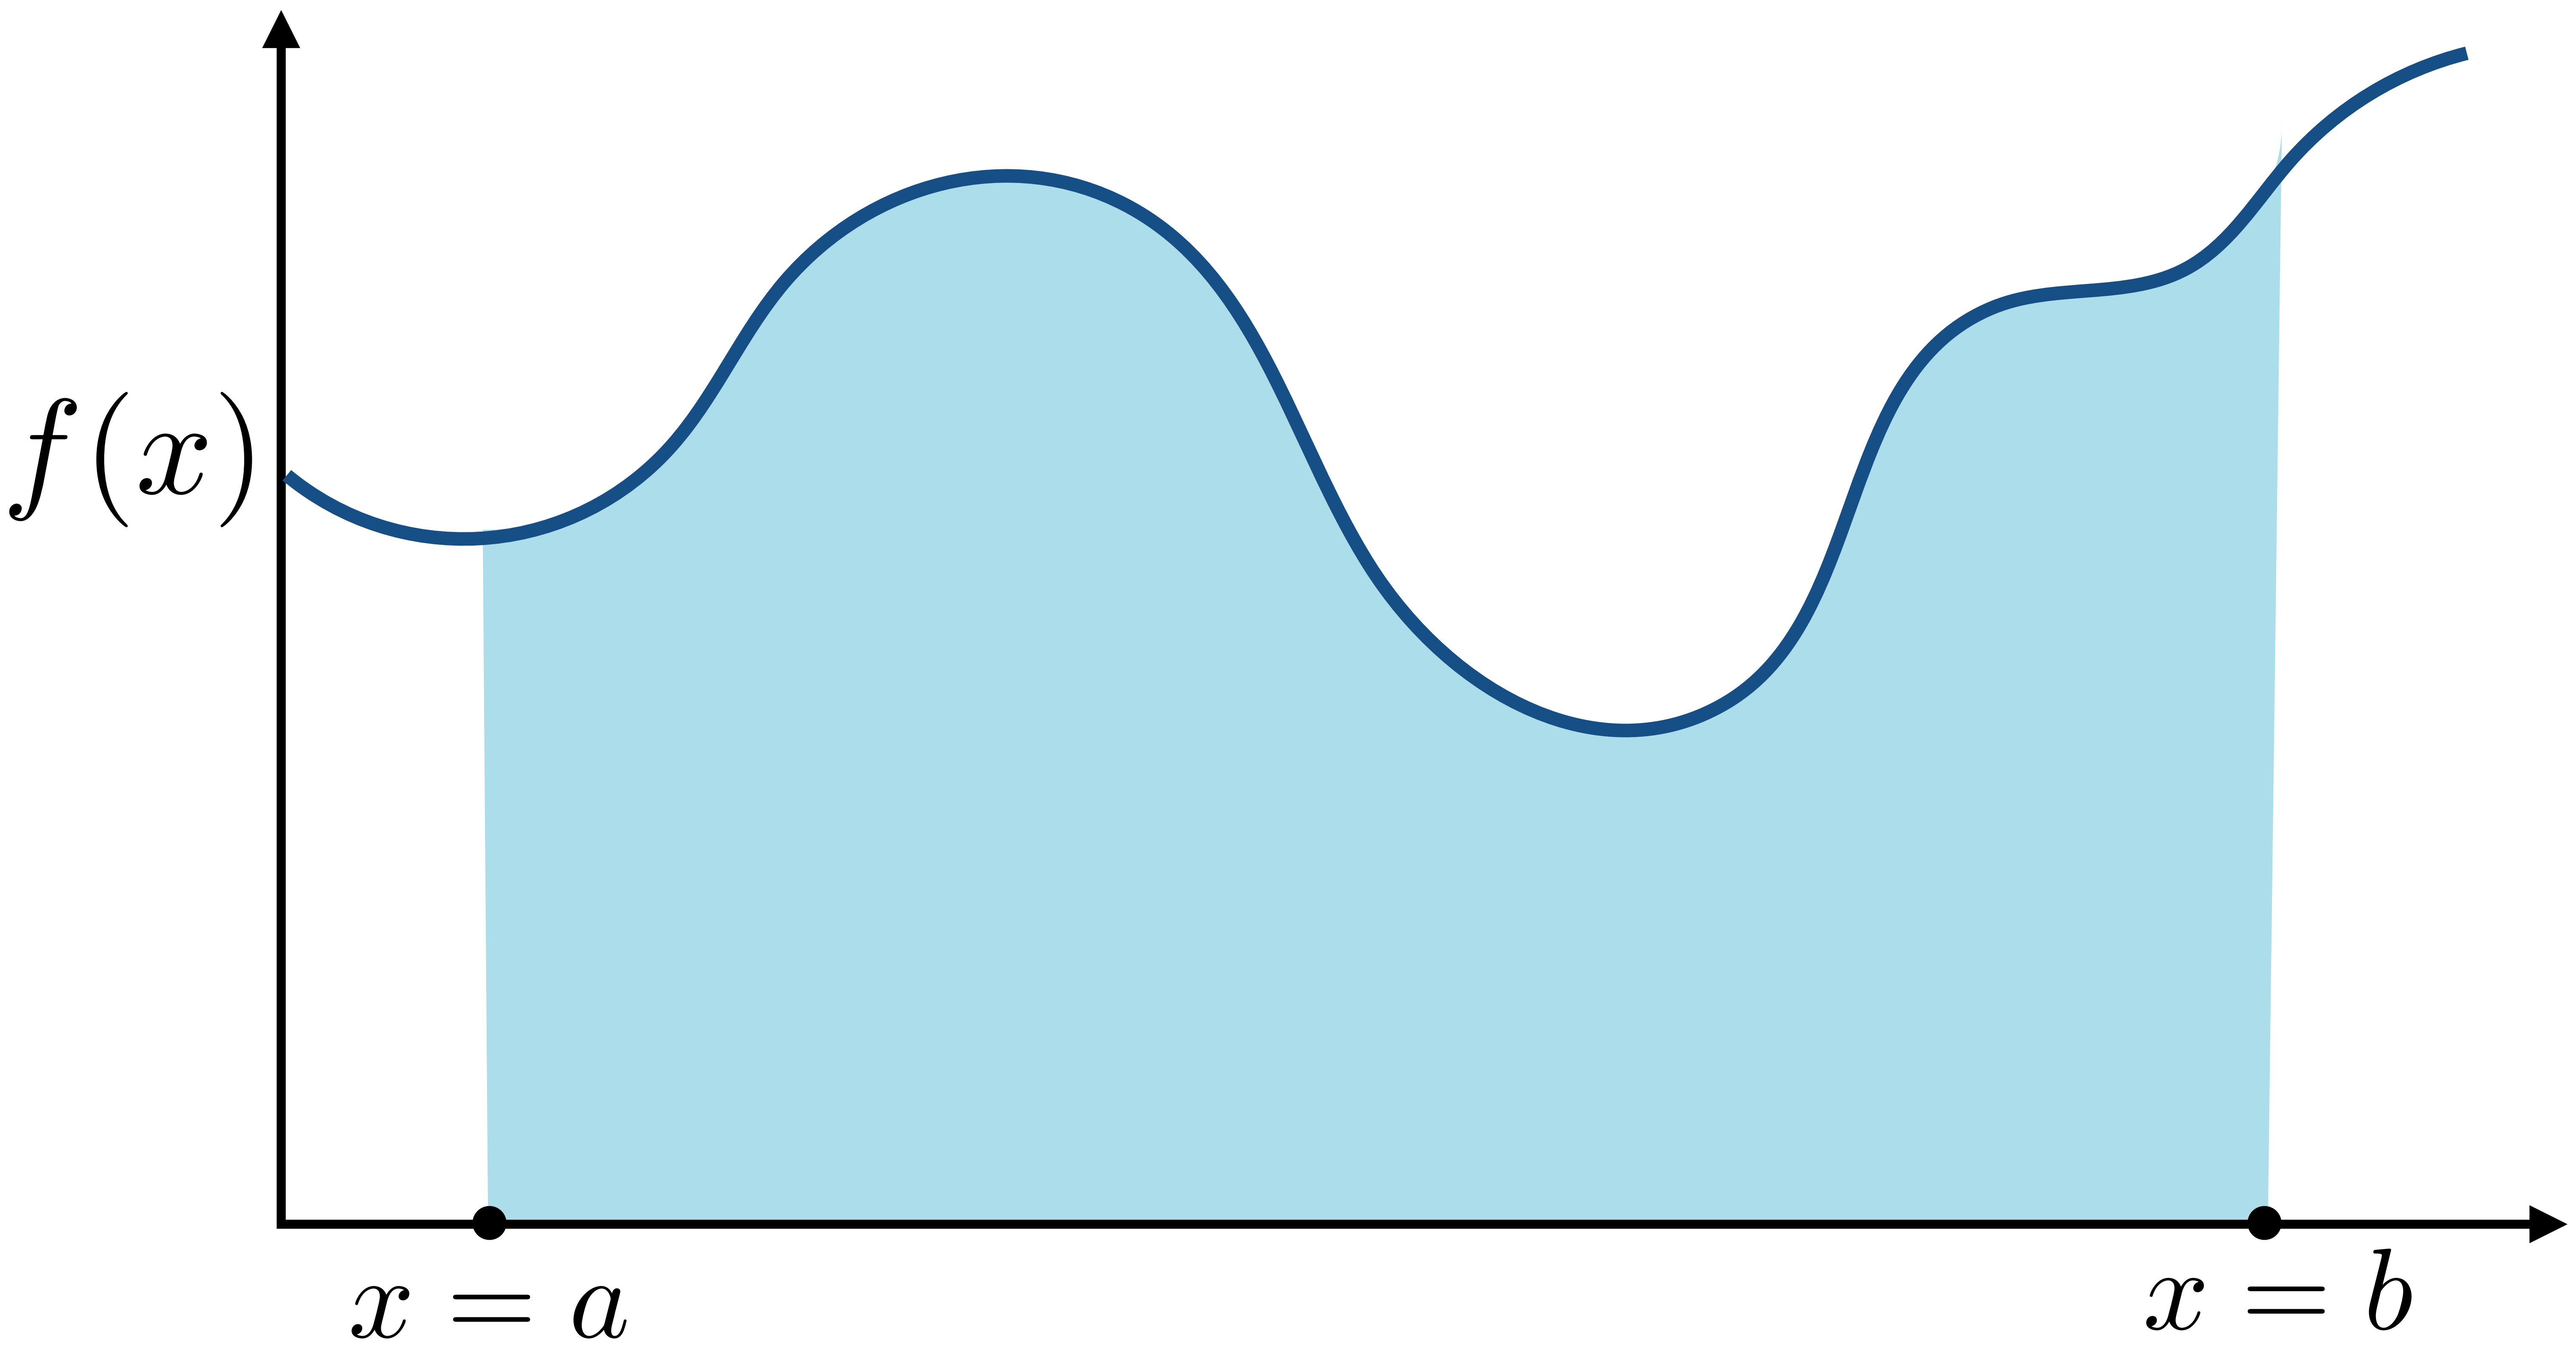

Simple case: constant function

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

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

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

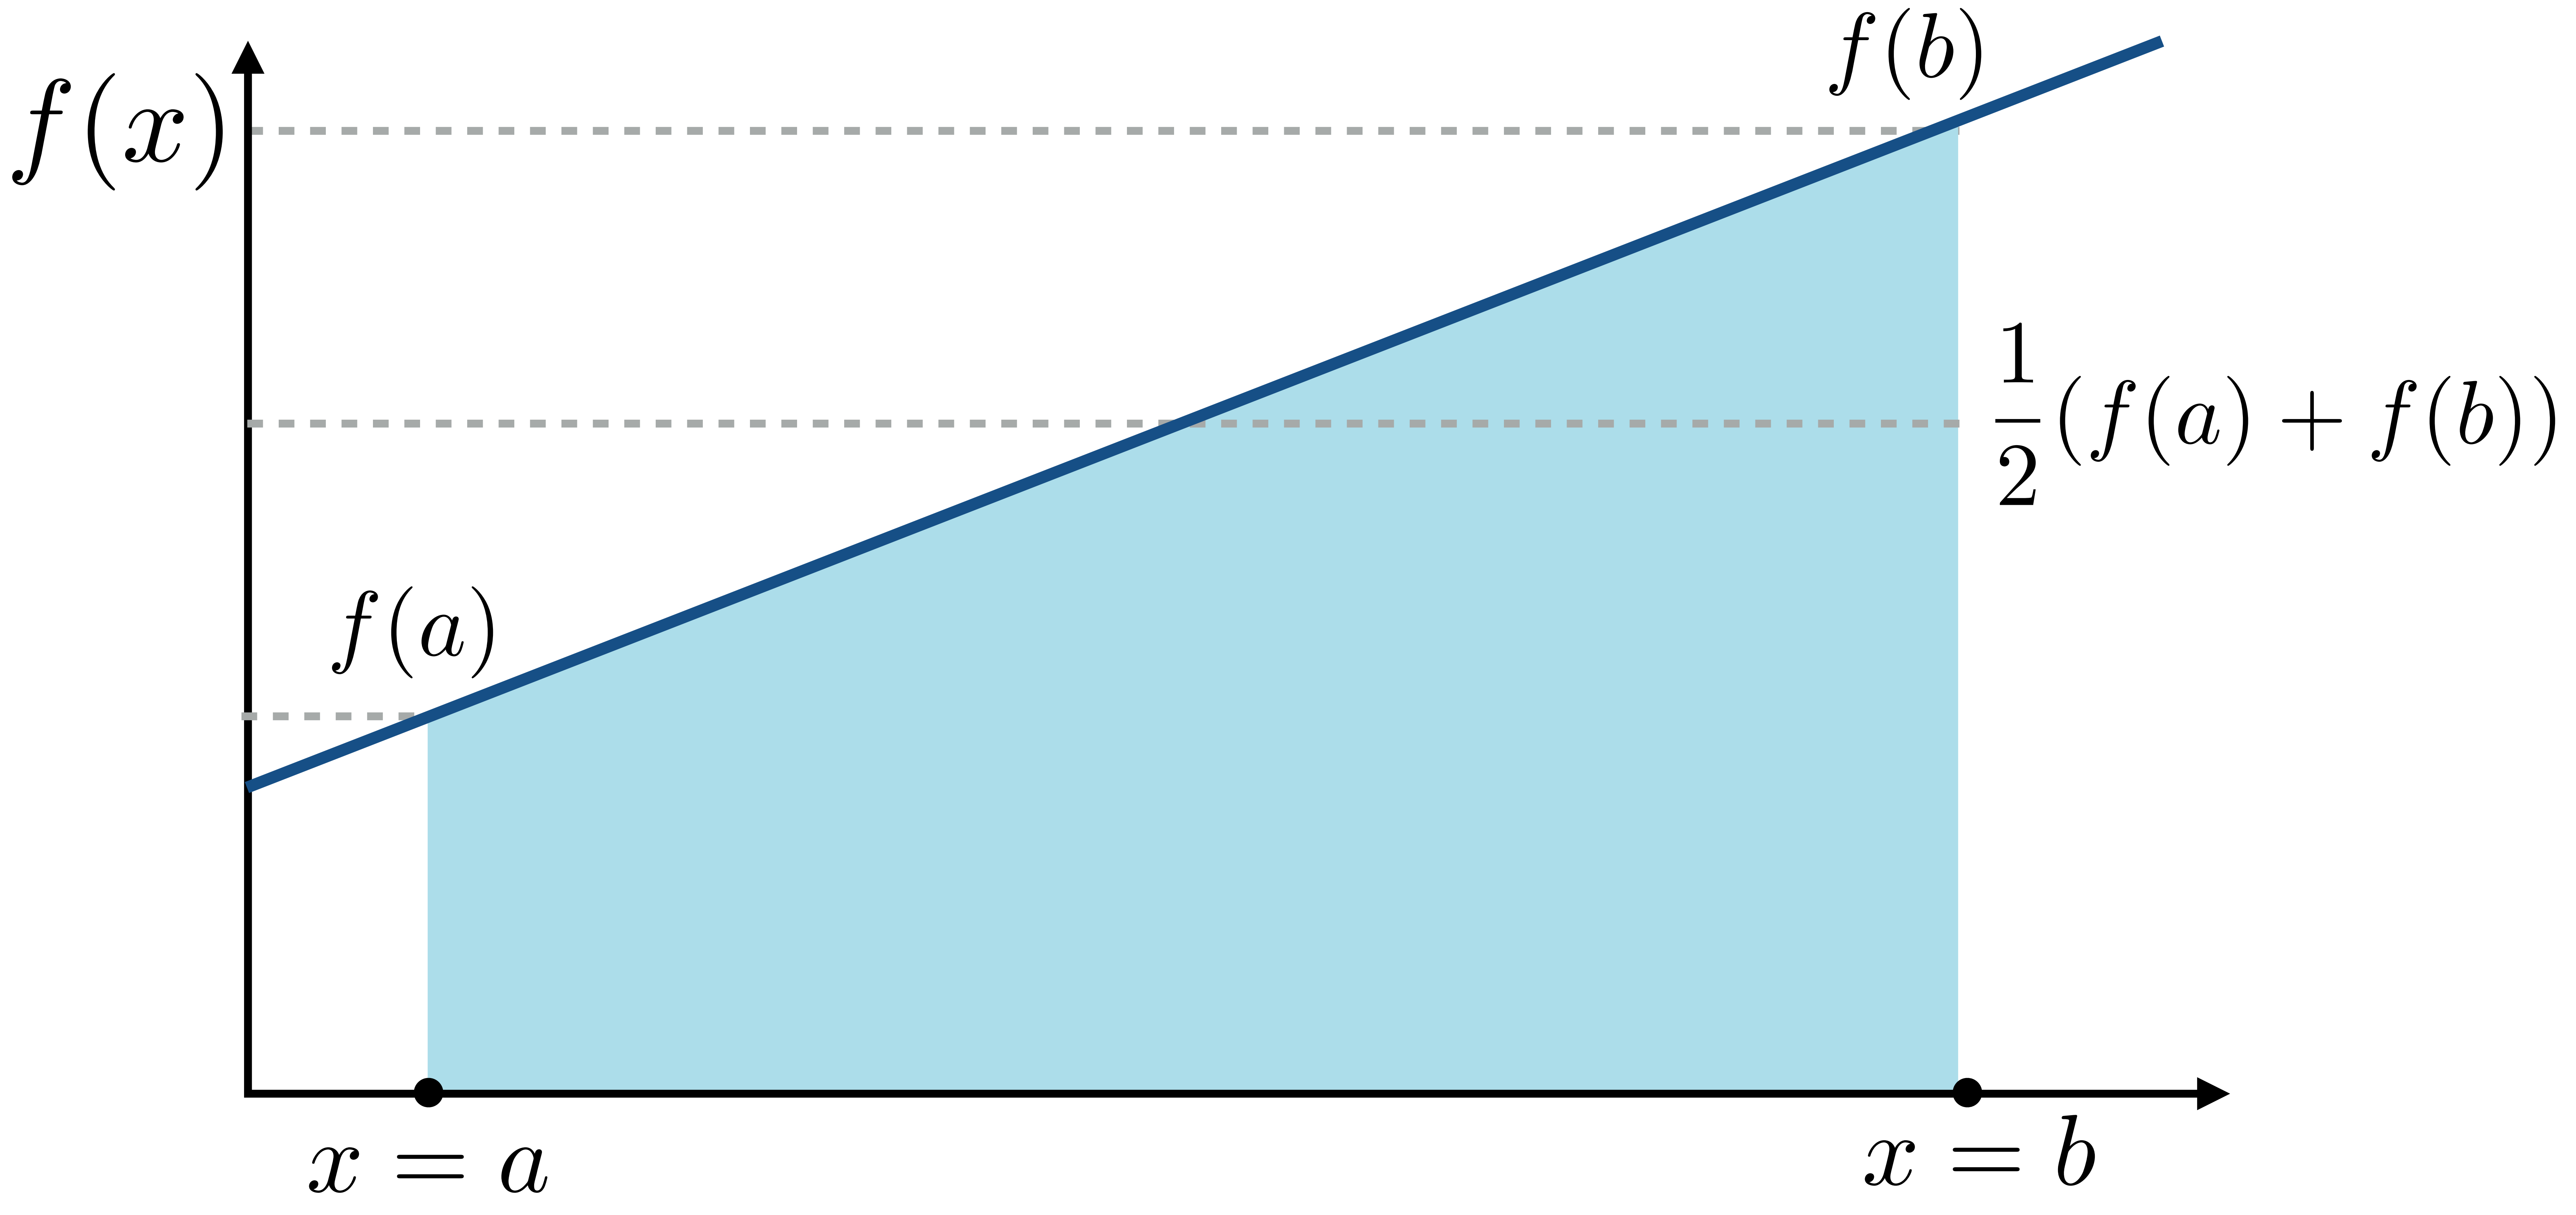

Piecewise affine function

Sum of integrals of individual affine components

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

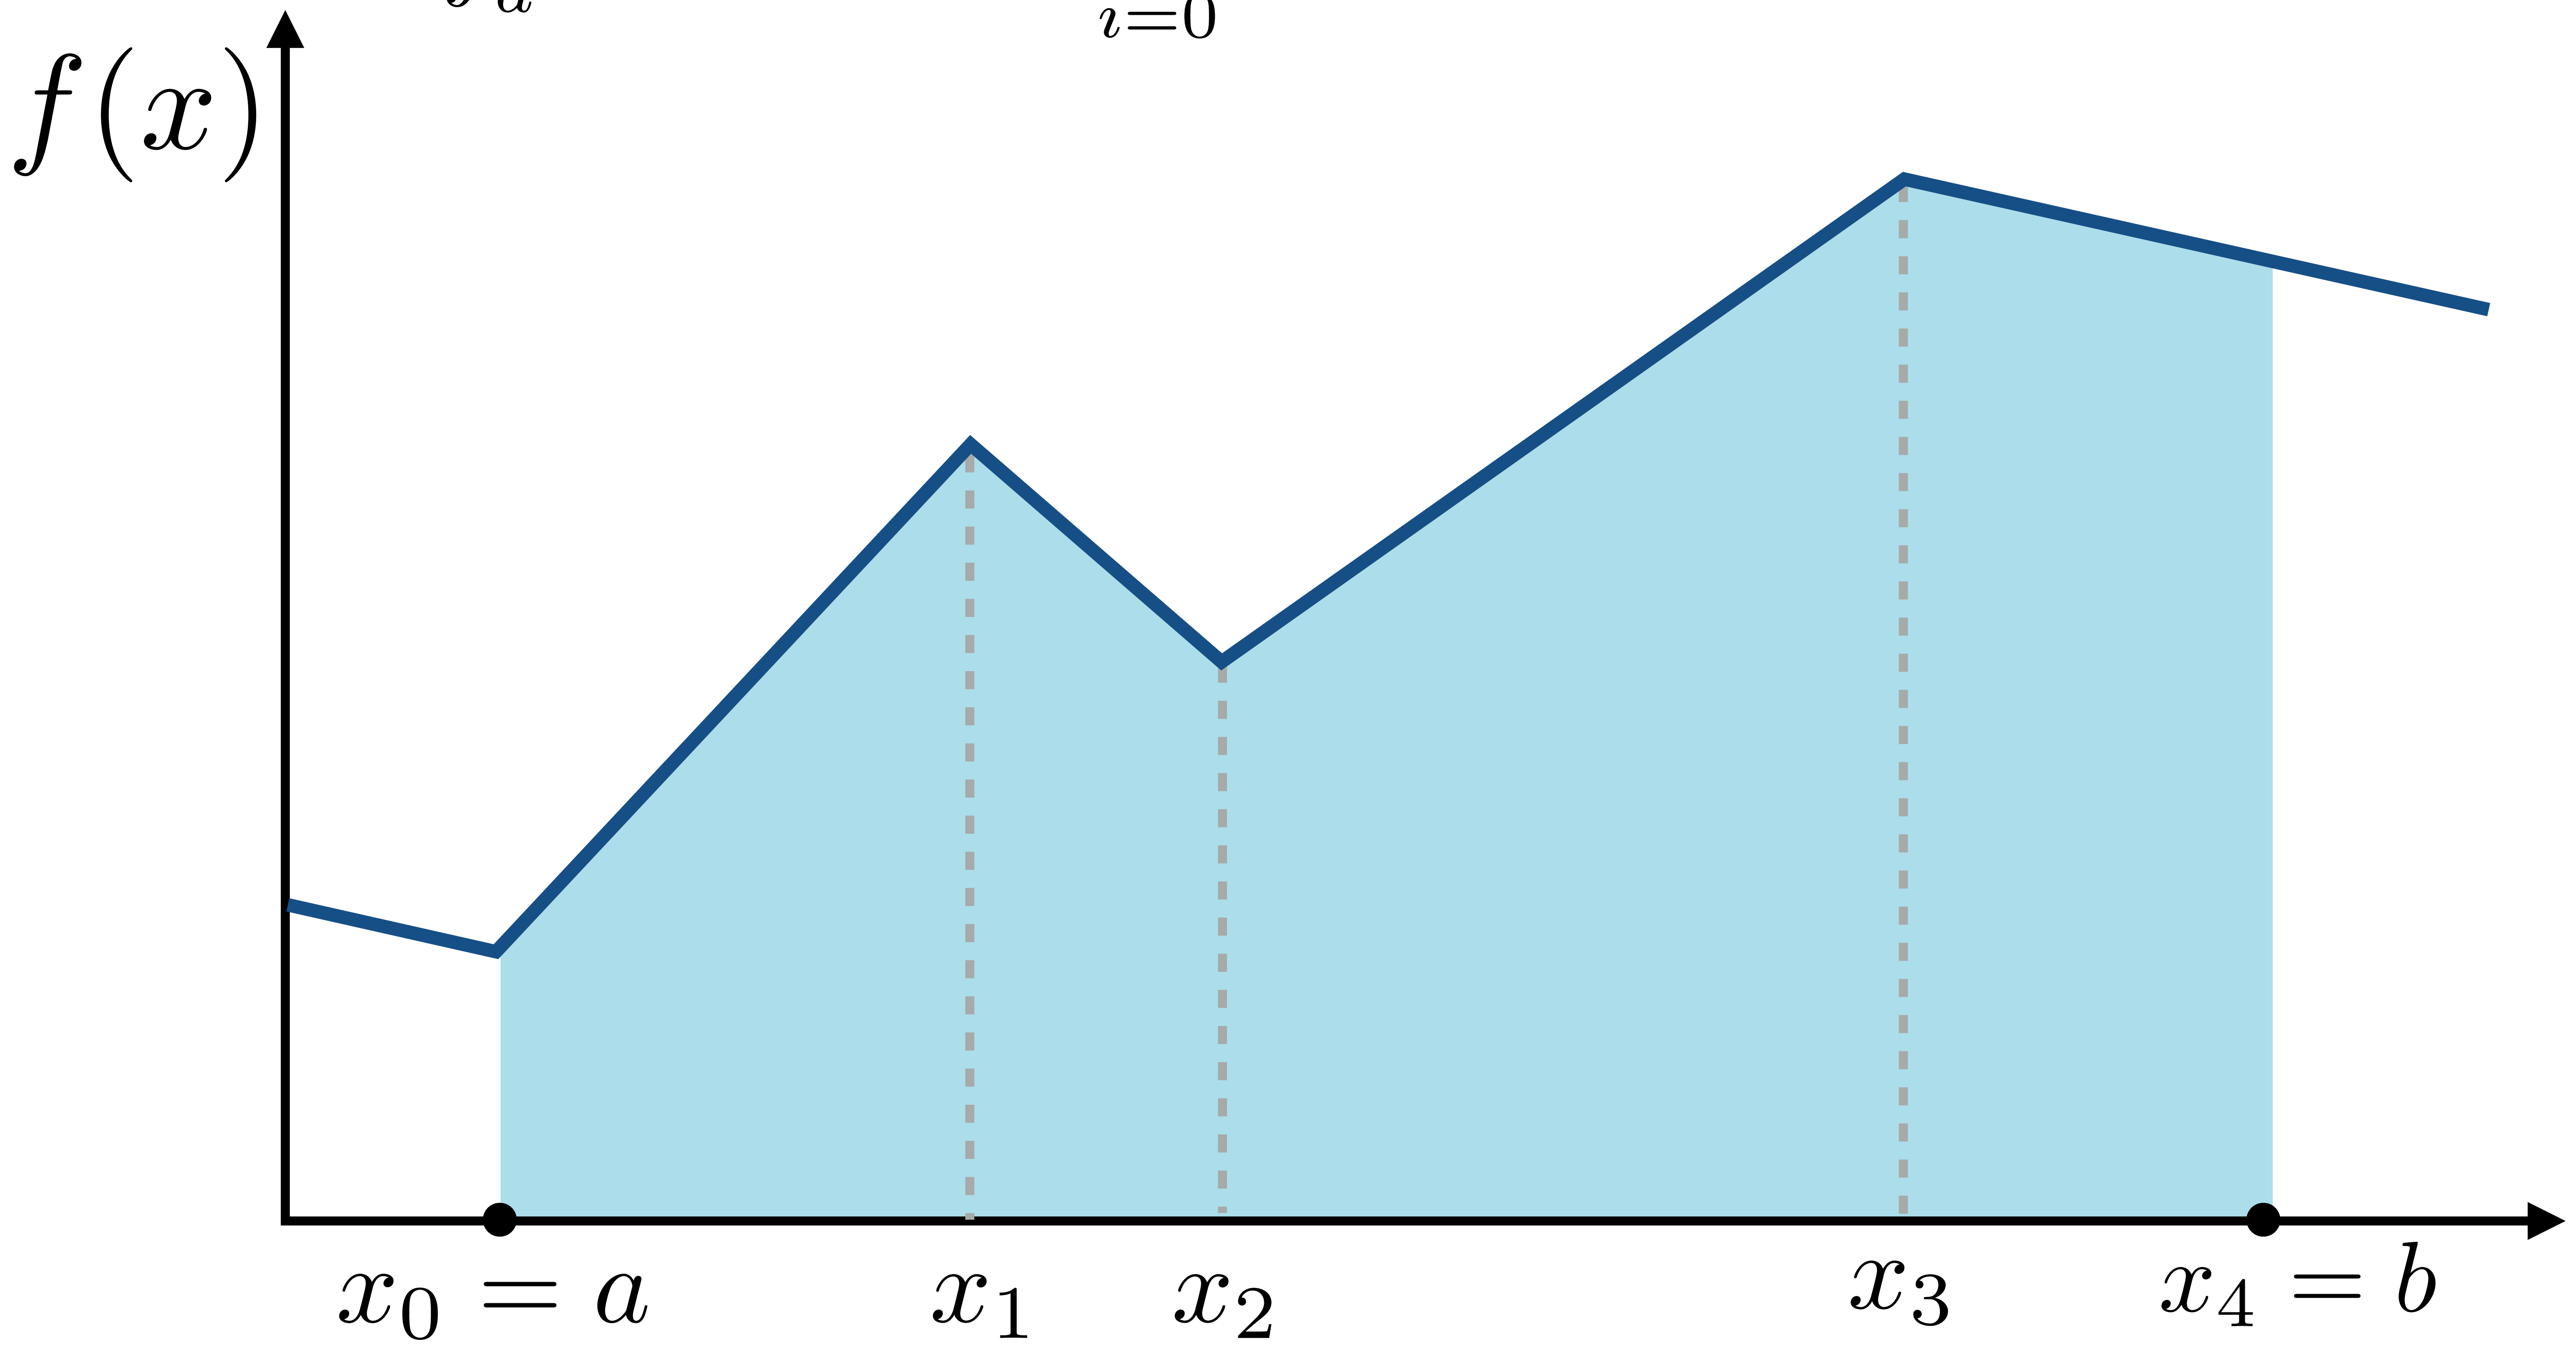

Piecewise affine function

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

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

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

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

Weighted combination of measurements.

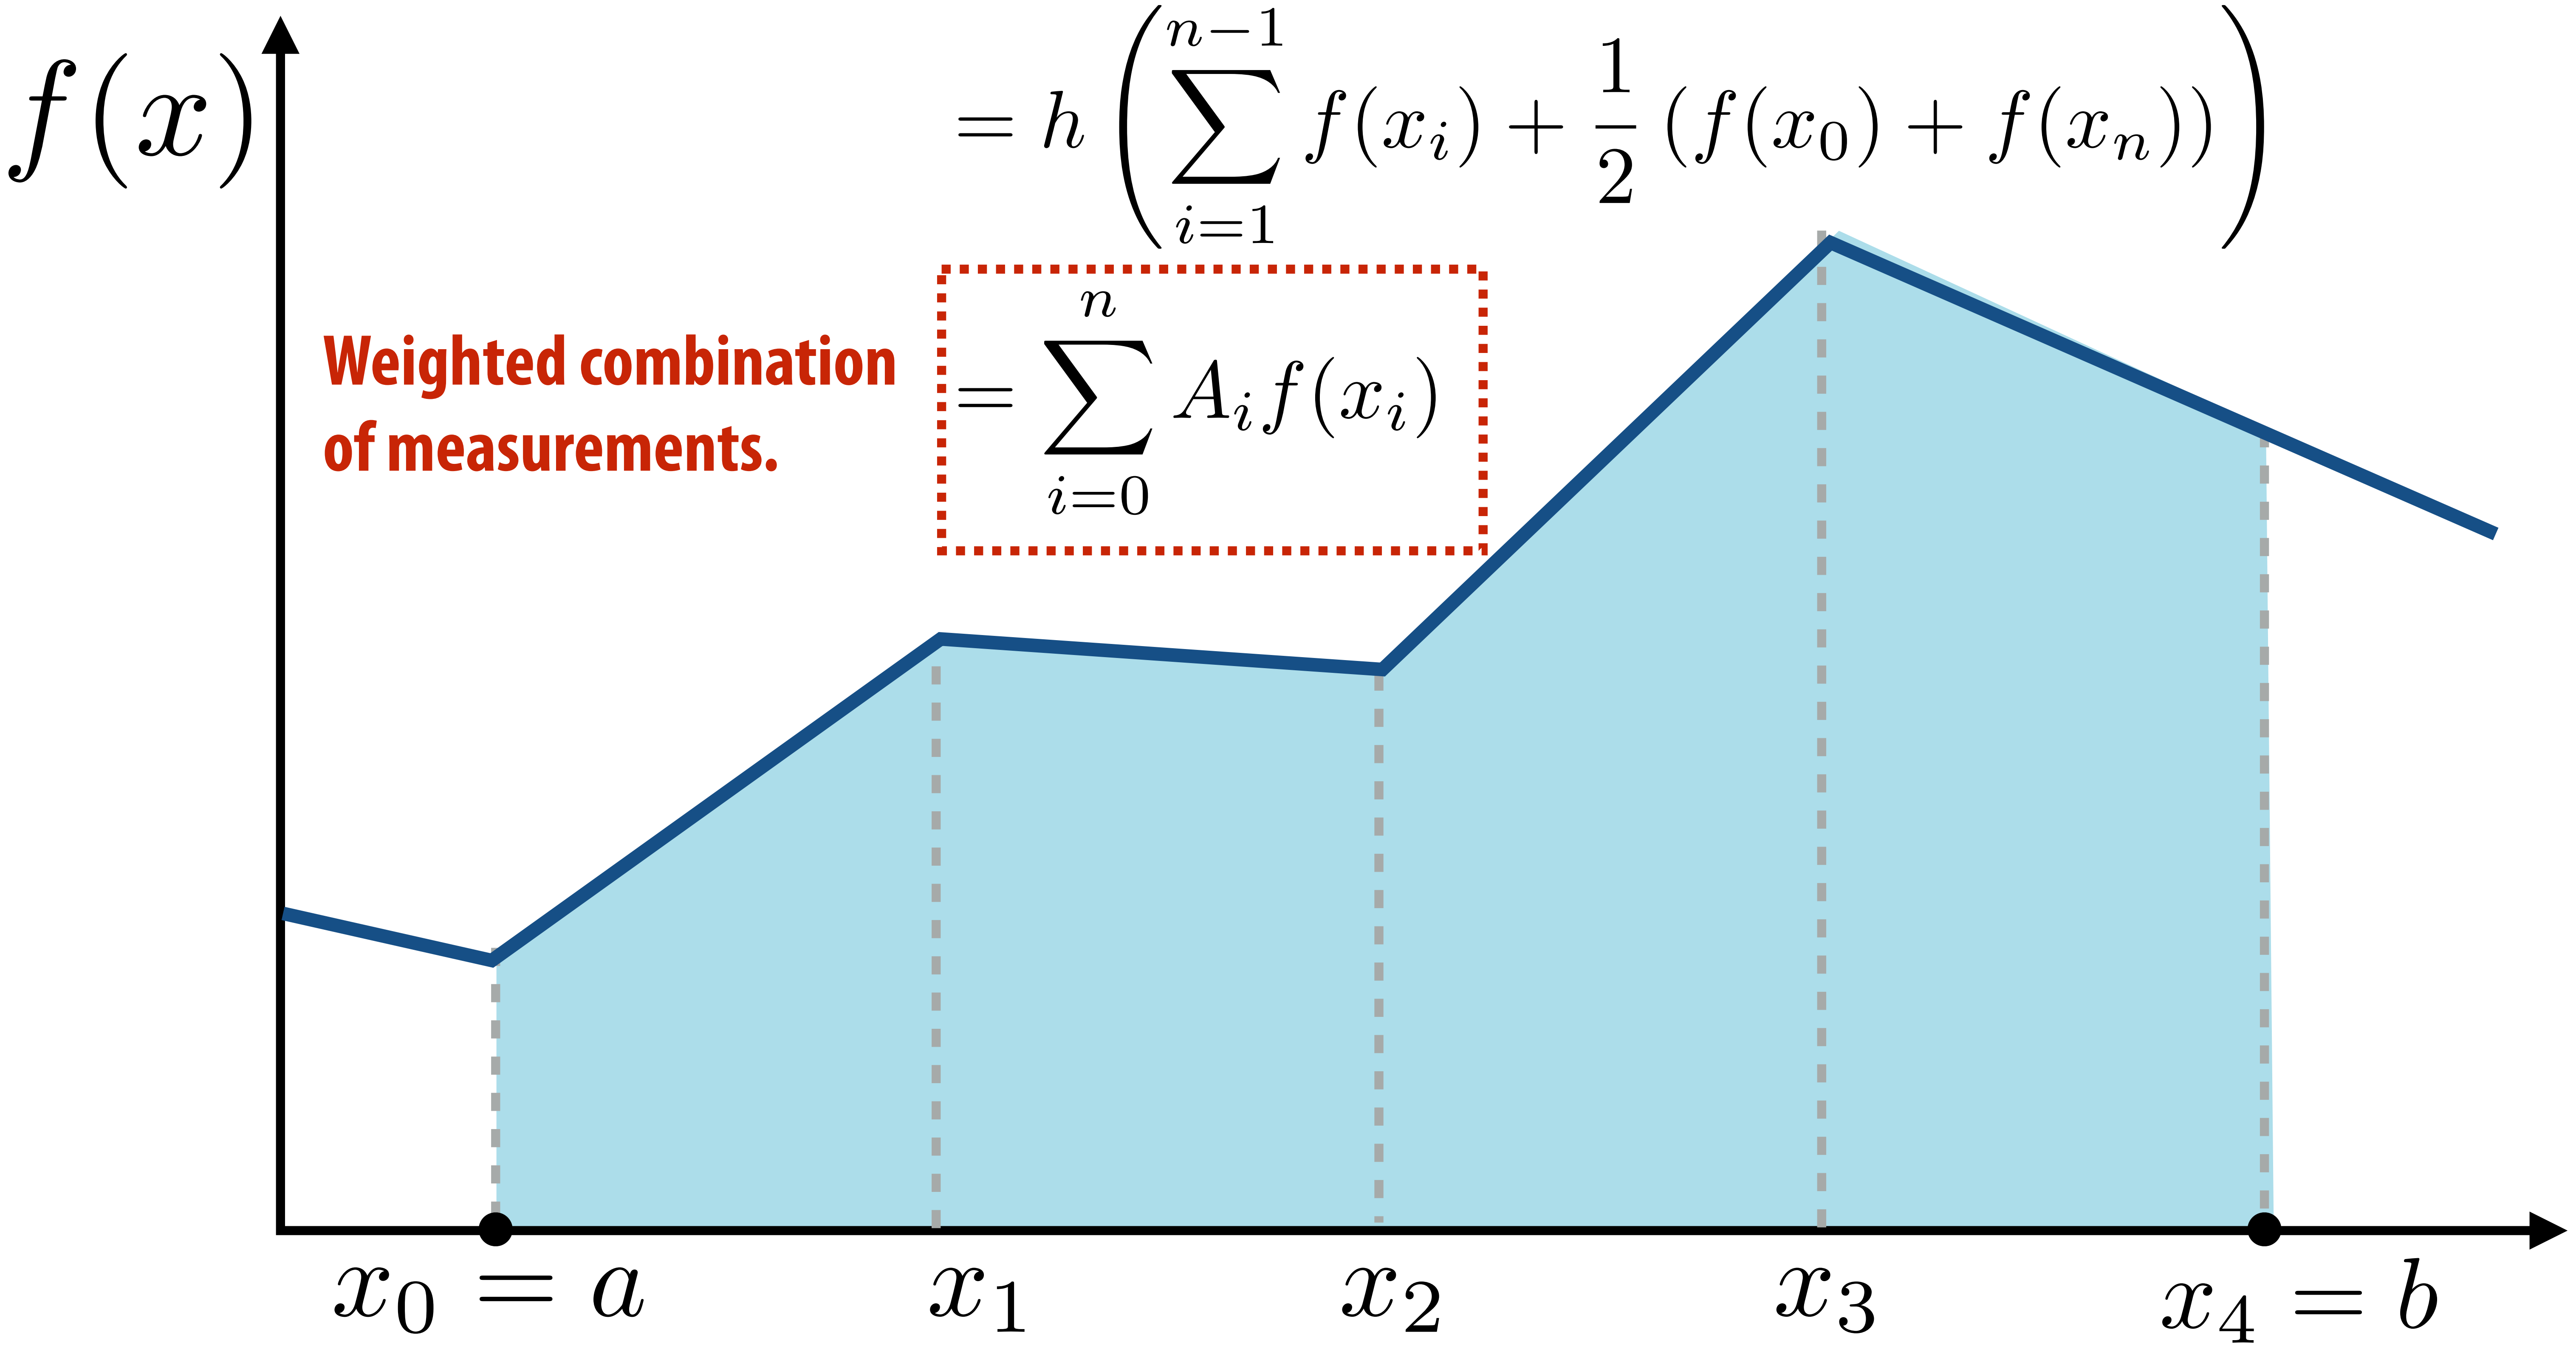

Arbitrary function f(x)?

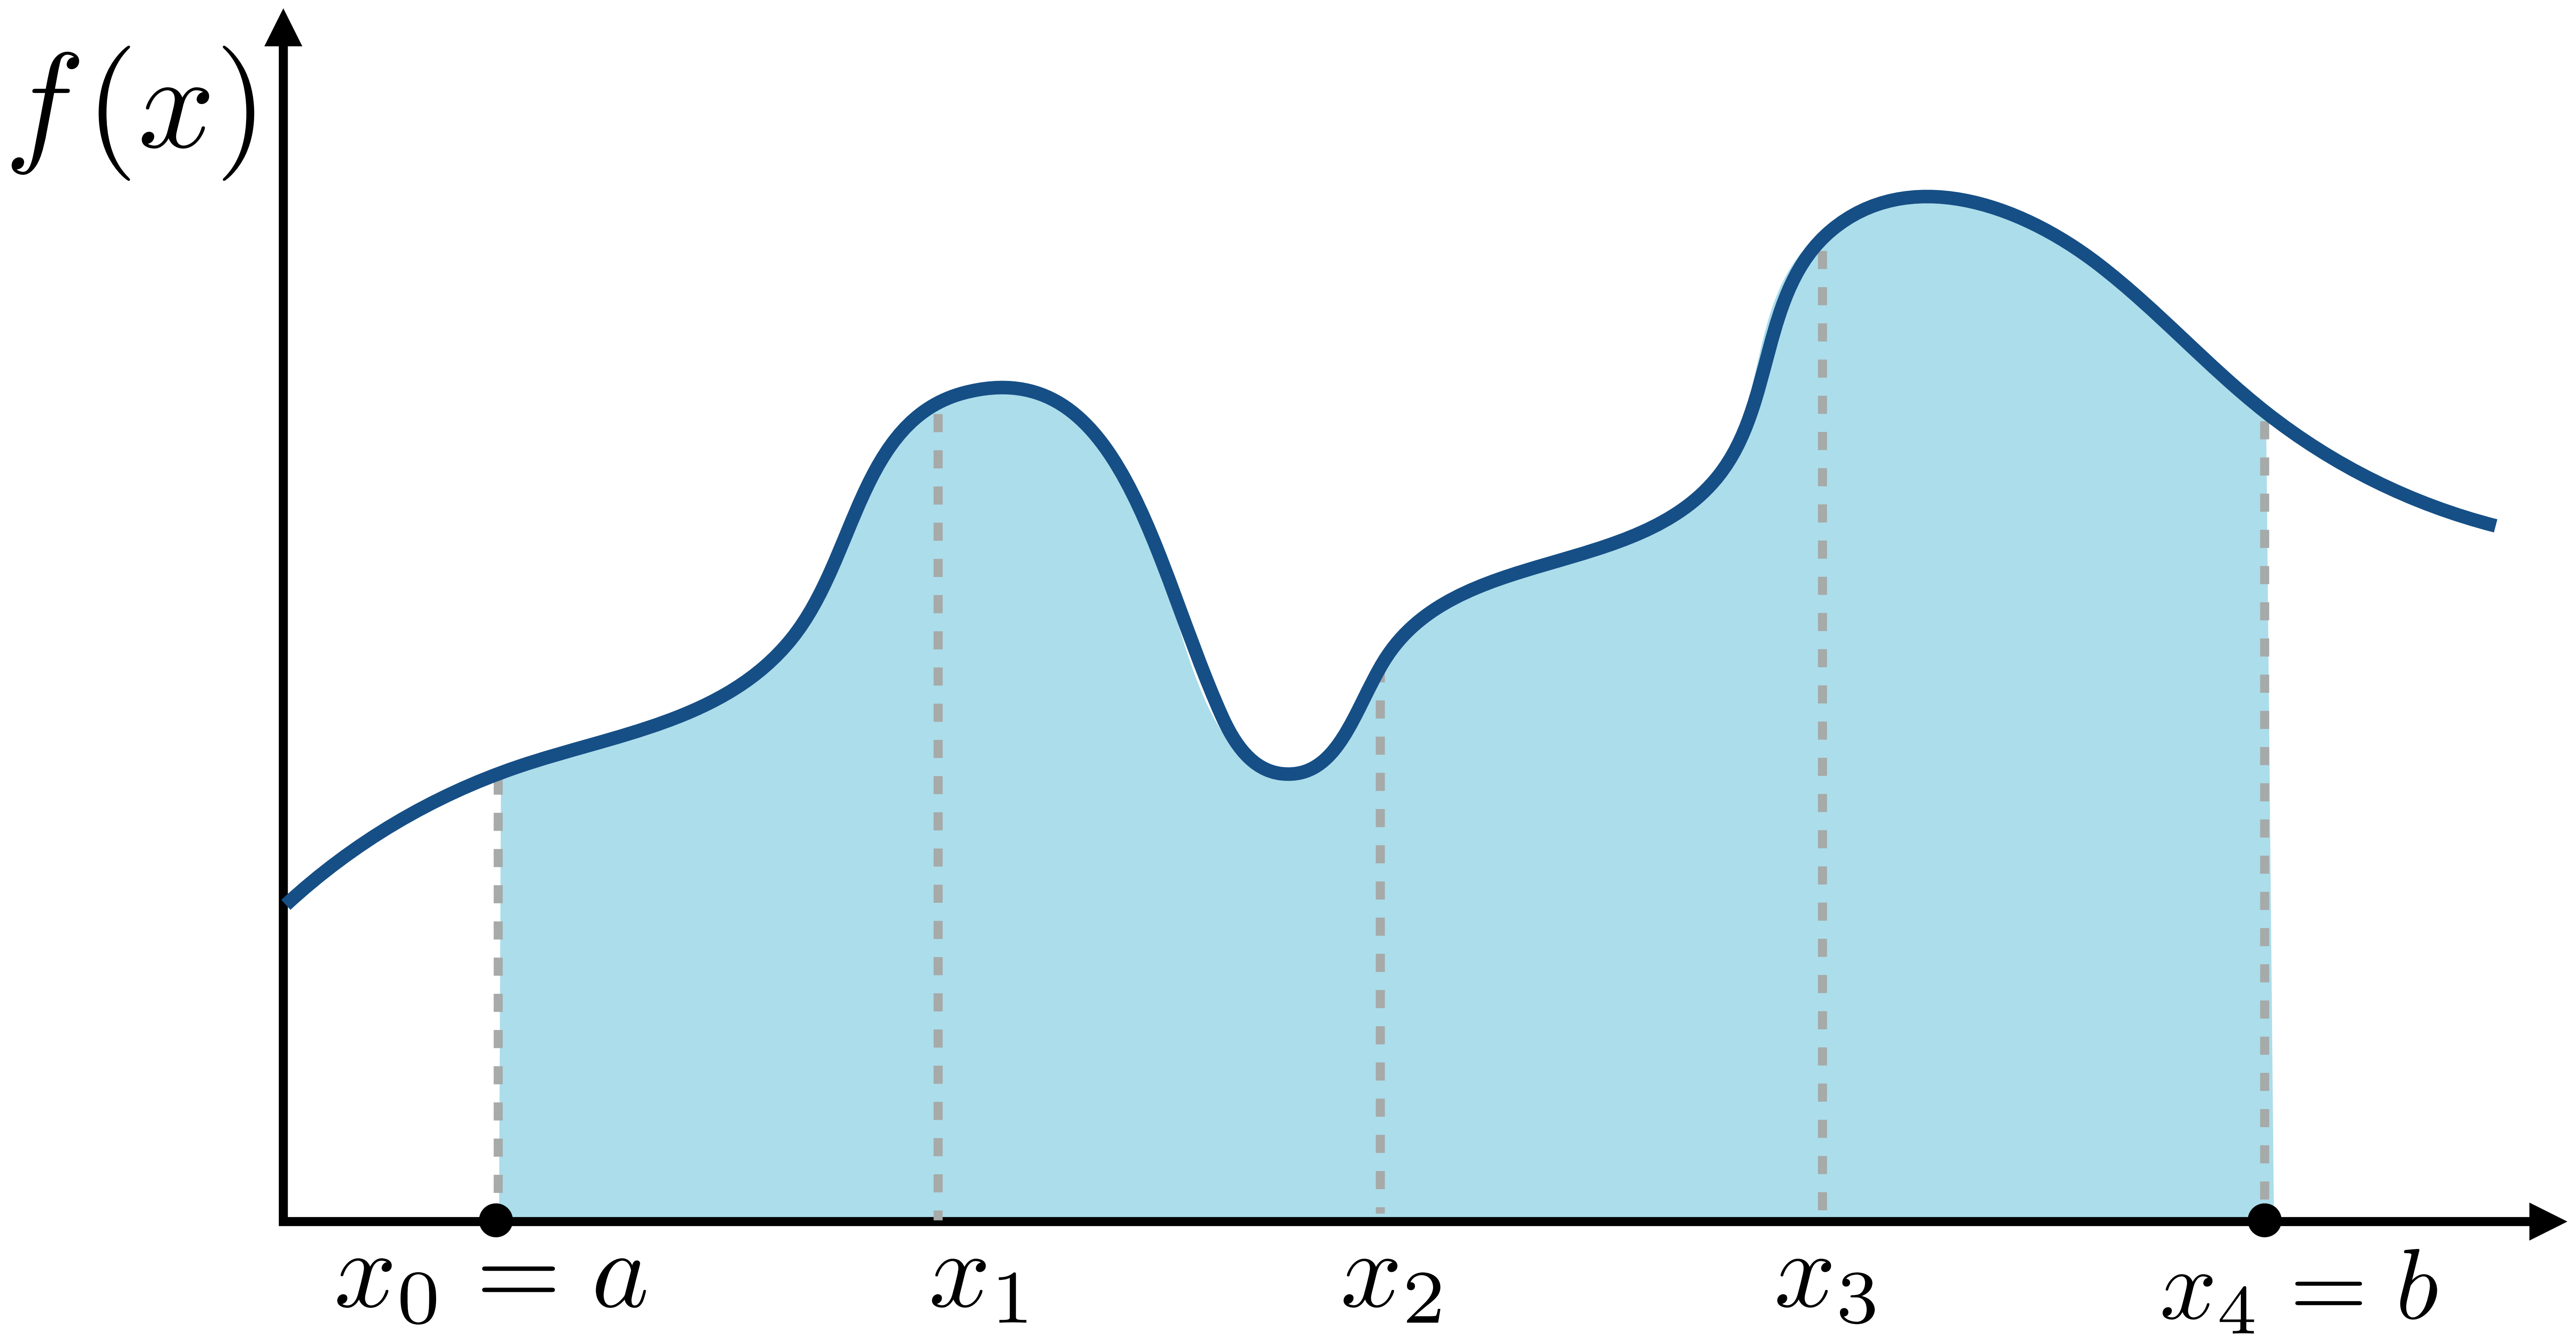

Trapezoidal rule

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

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

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

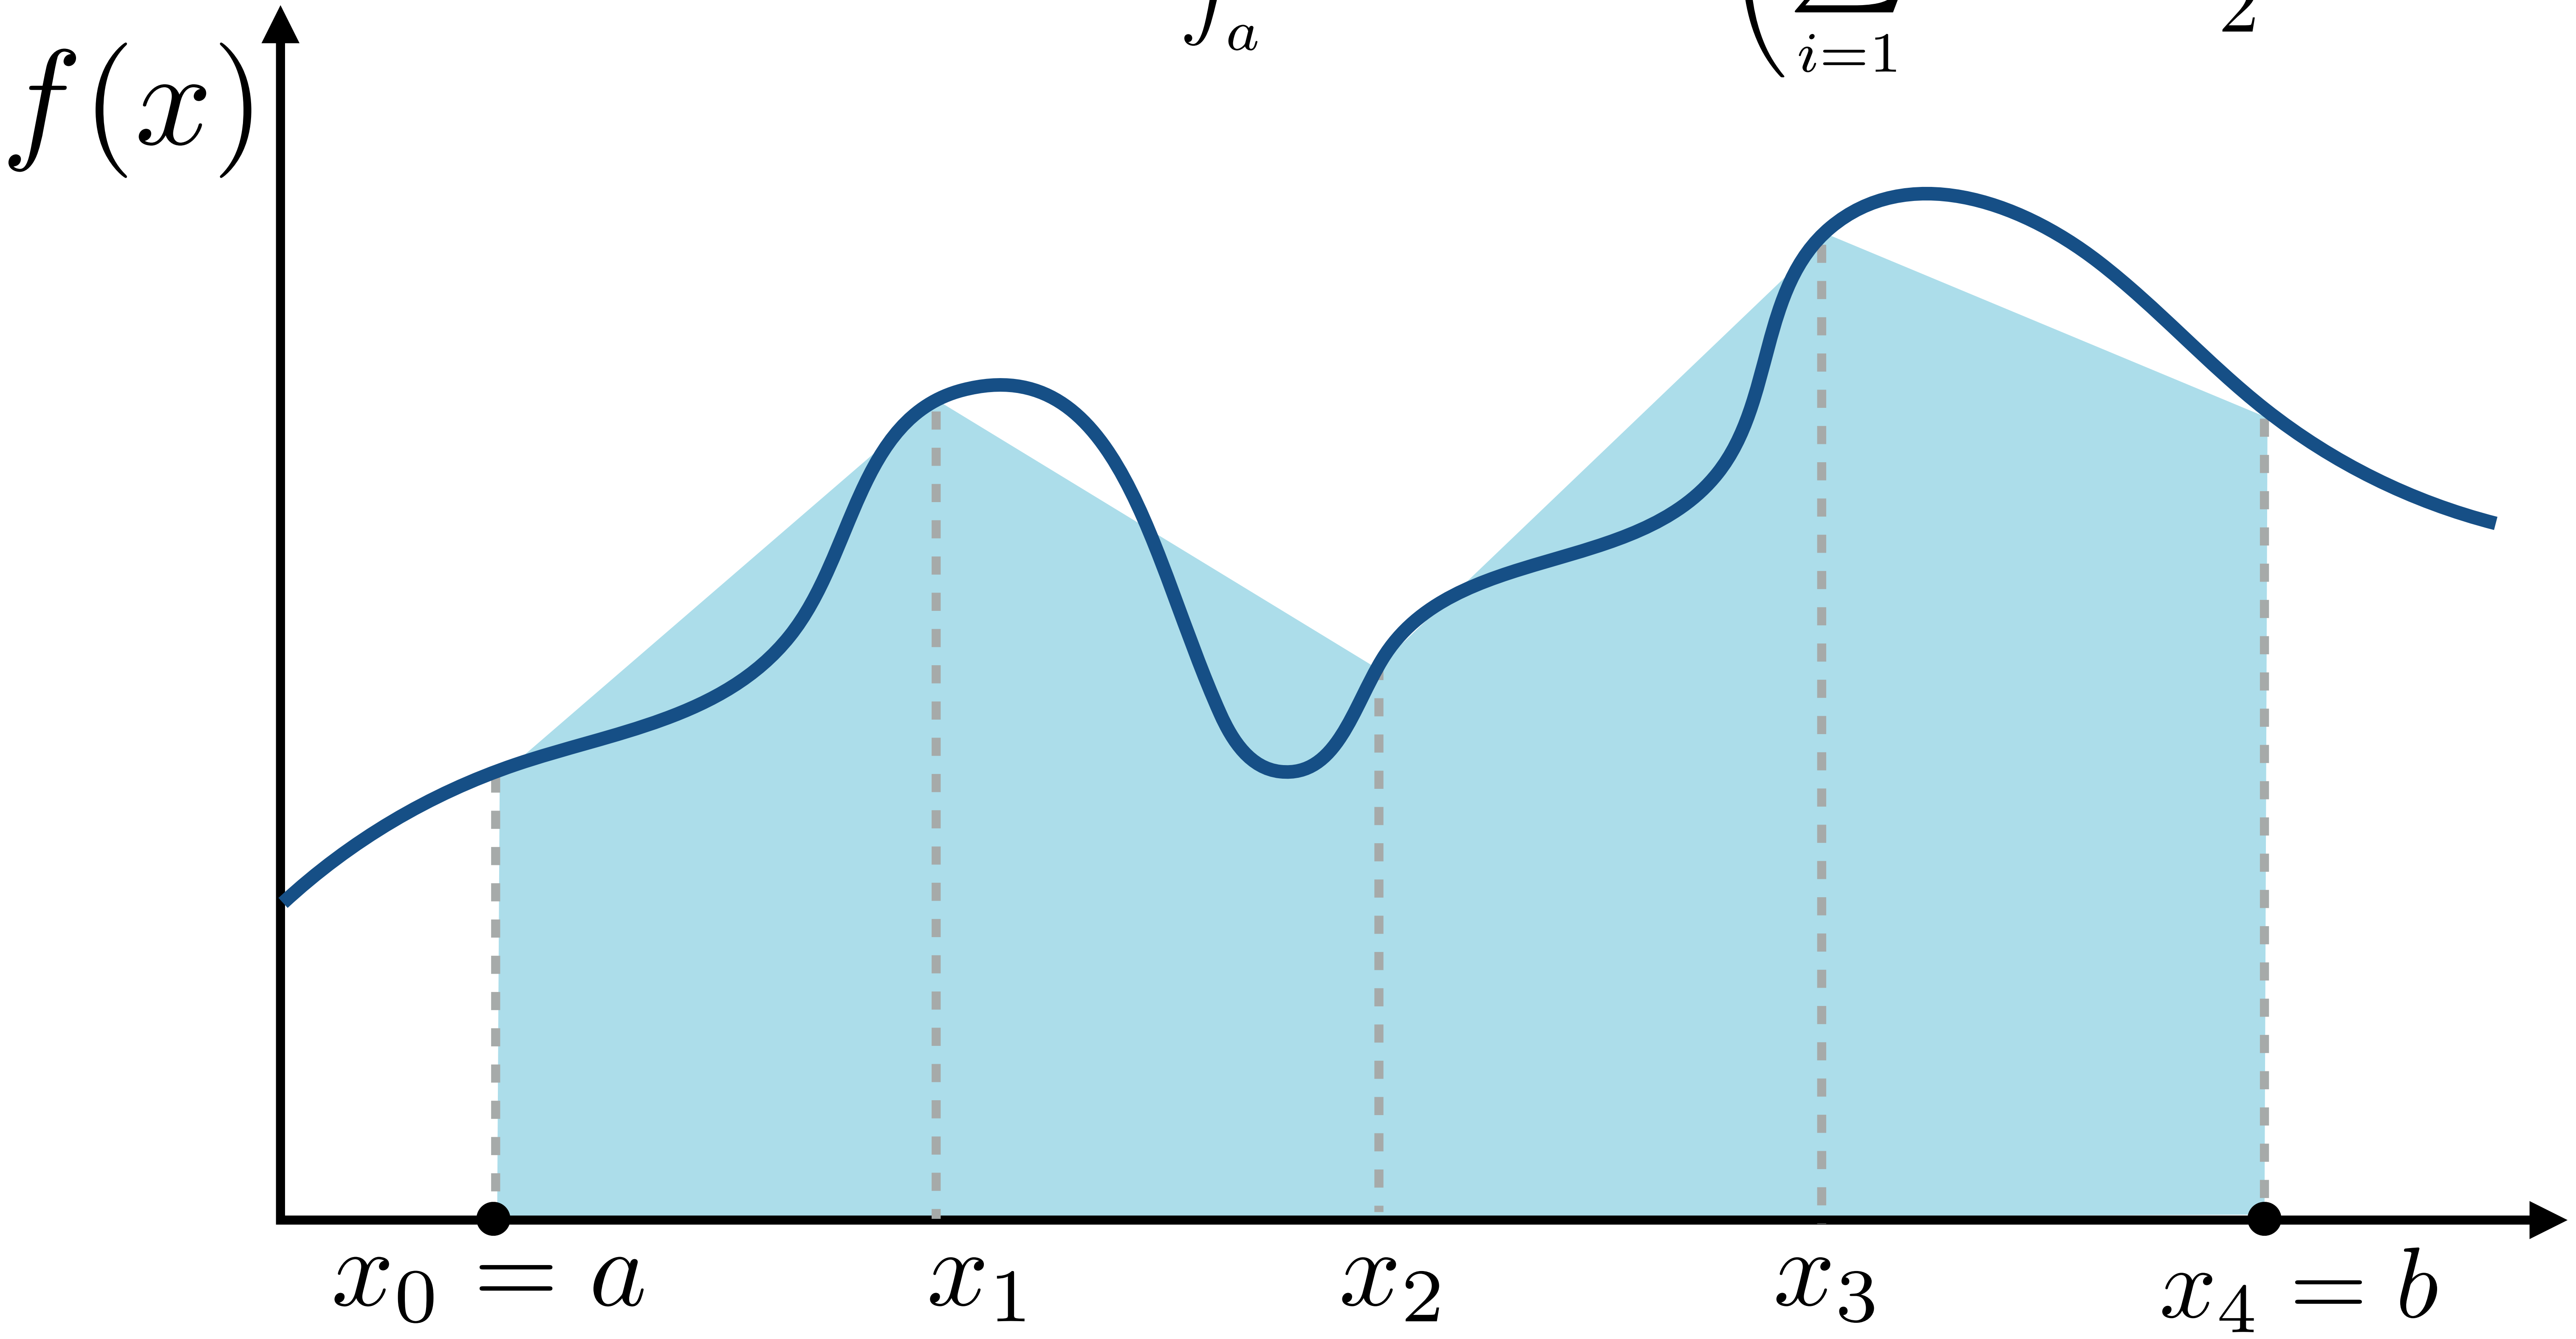

Trapezoidal rule

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

Work: $O(n)$

Error can be shown to be: $O(h^2) = O(\frac{1}{n^2})$

(for f(x) with continuous second derivative)

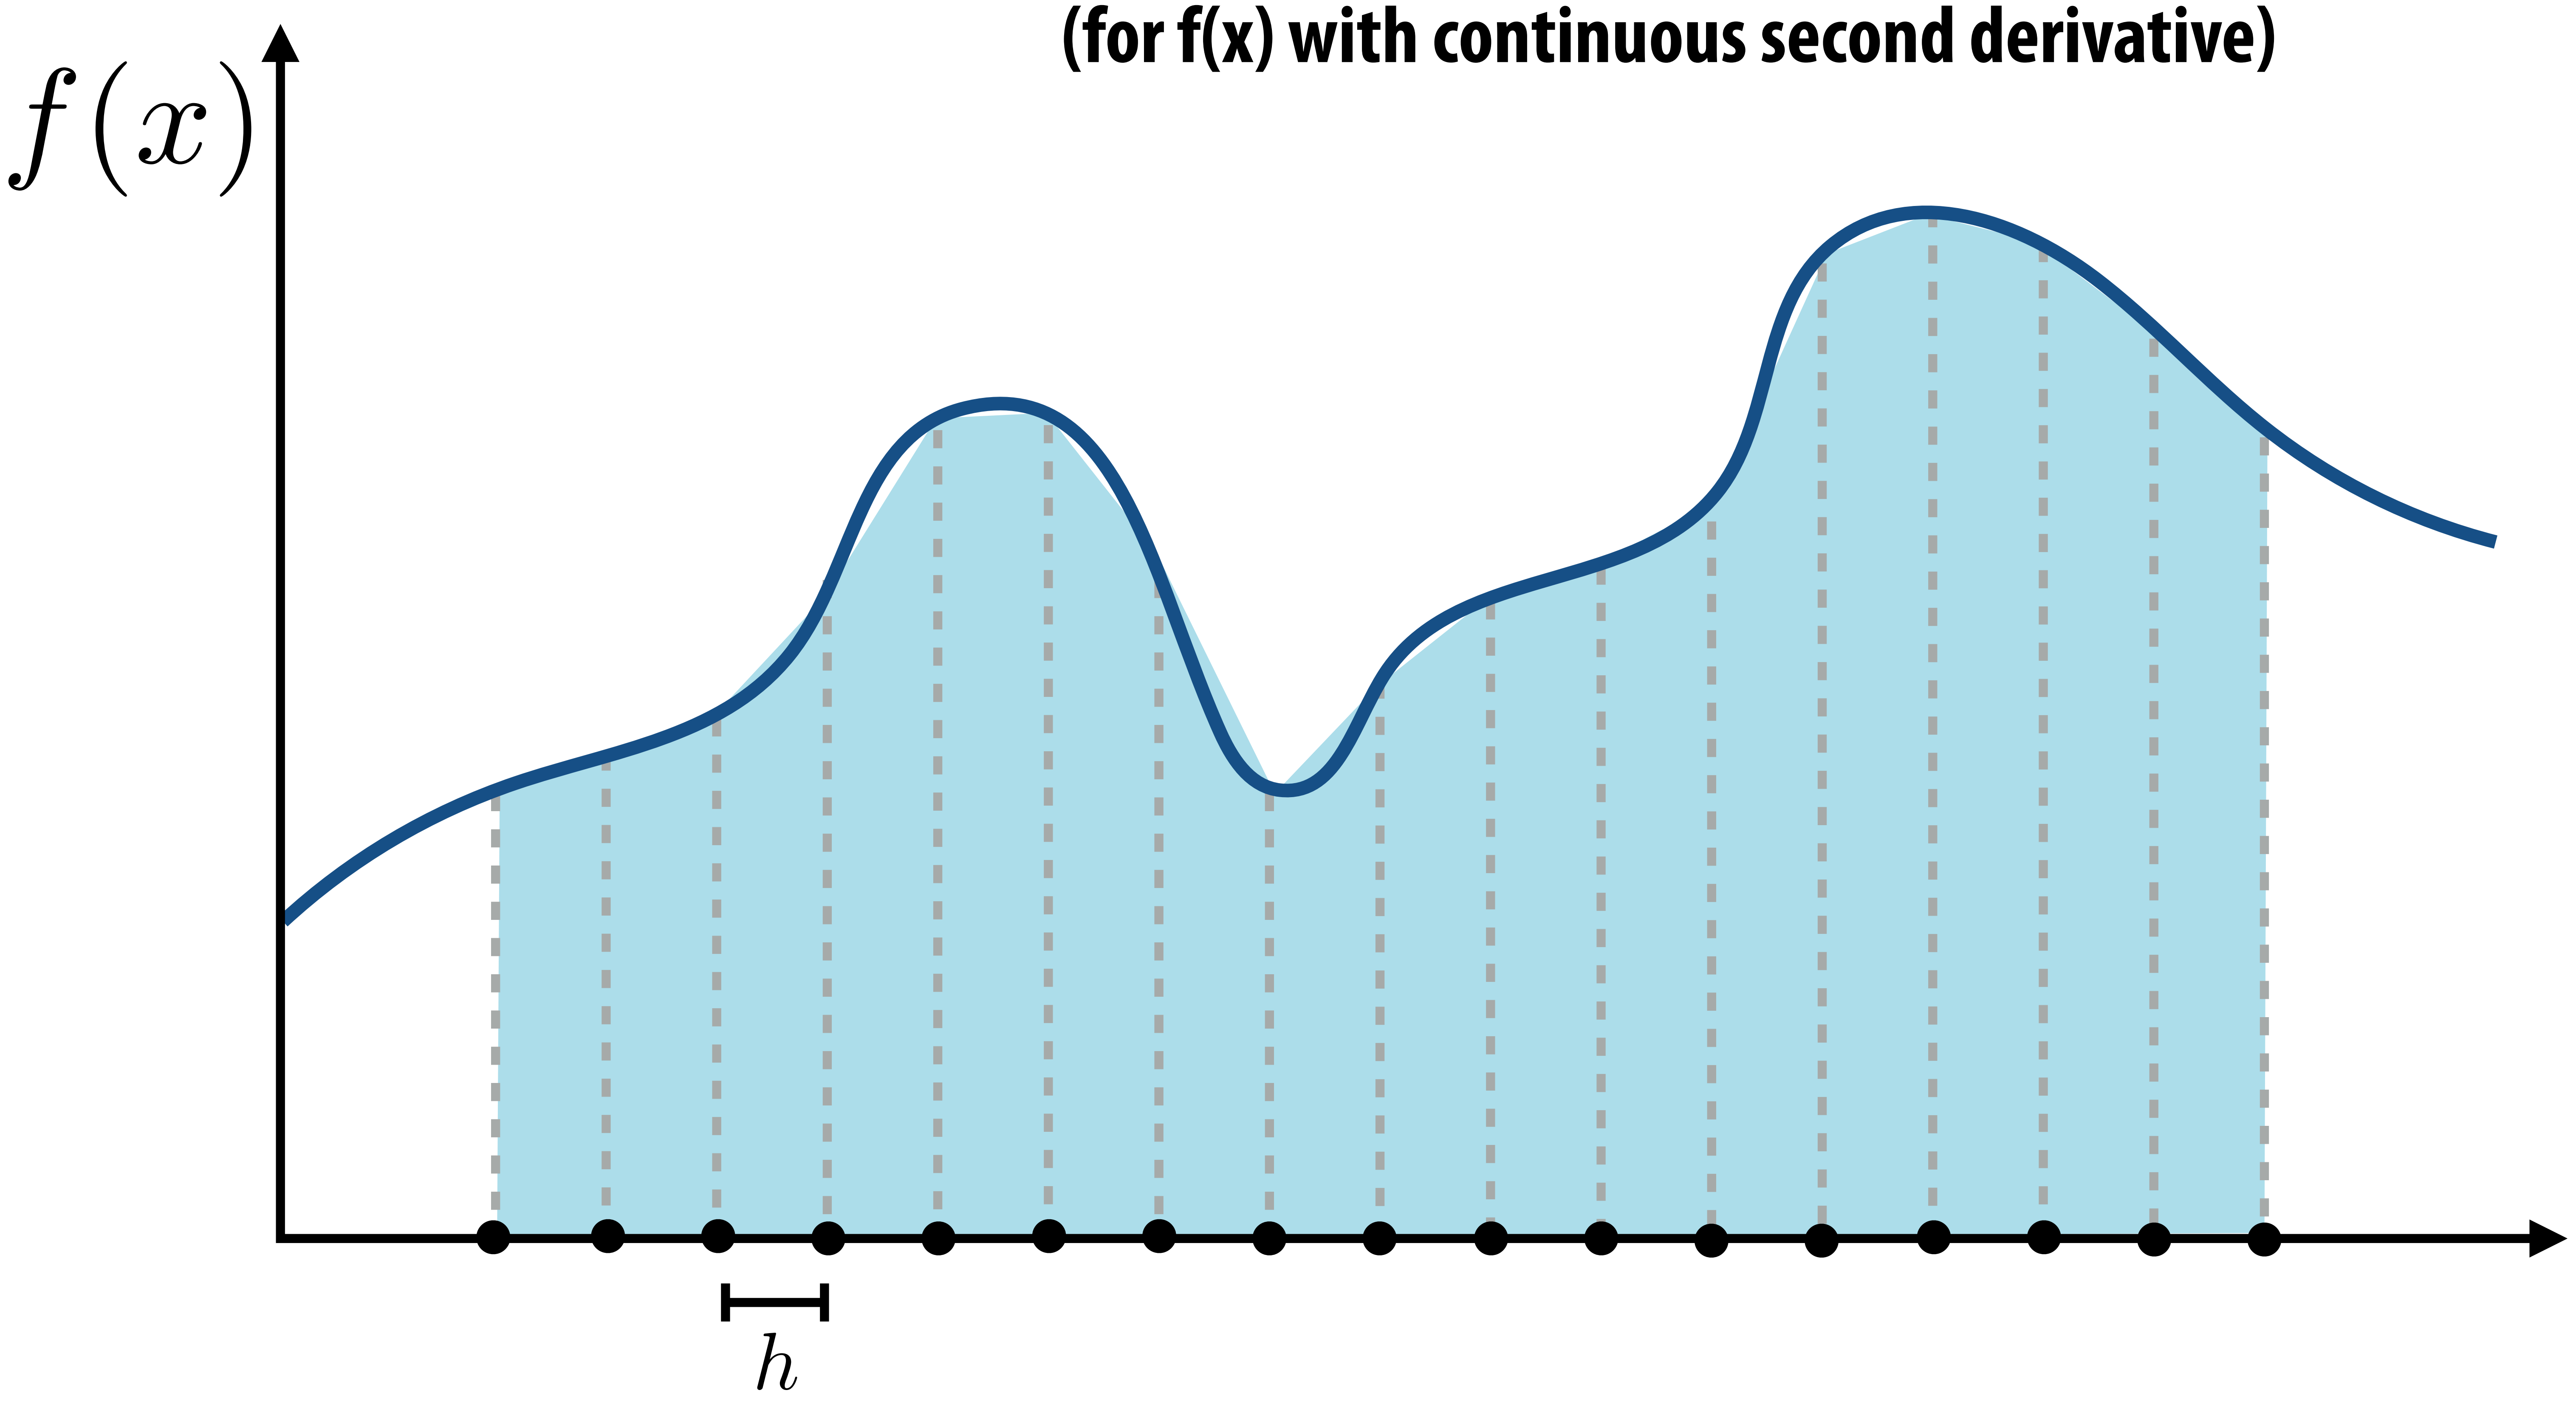

Integration in 2D

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

$$\int_{a_y}^{b_y} \int_{a_x}^{b_x} f(x,y)dxdy = \int_{a_y}^{b_y} \left(\boxed{O(h^2) + \sum_{i=0}^{n} A_i f(x_i, y)} \right) dy$$

First application of rule

$$= O(h^2) + \sum_{i=0}^{n} A_i \int_{a_y}^{b_y} f(x_i, y)dy$$

$$= O(h^2) + \sum_{i=0}^{n} A_i \left(\boxed{O(h^2) + \sum_{j=0}^{n} A_j f(x_i, y_j)} \right)$$

Second application

$$= O(h^2) + \sum_{i=0}^{n} \sum_{j=0}^{n} A_i A_j f(x_i, y_j)$$

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

Monte Carlo integration

Monte Carlo numerical integration

- 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: $O(n^{1/2})$

Monte Carlo Algorithms

- Advantages
  - Easy to implement
  - Easy to think about (but be careful of subtleties)
  - Robust when used with complex integrands (lights, BRDFs) and domains (shapes)
  - Efficient for high-dimensional integrals
  - Efficient when only need solution at a few points

- Disadvantages
  - Noisy
  - Slow (many samples needed for convergence)

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

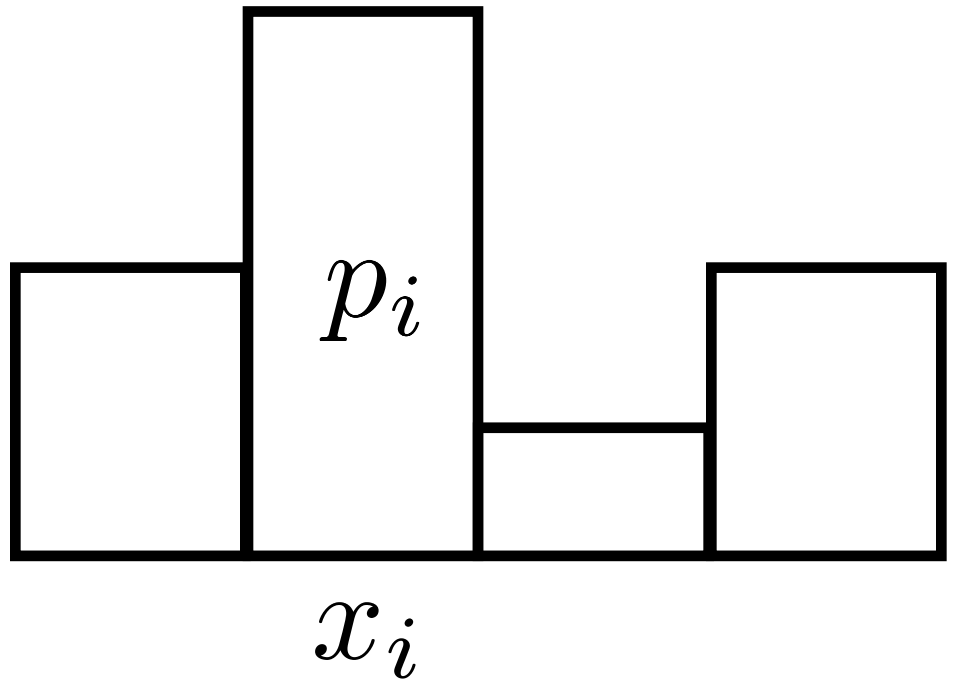

Requirements of a PDF:

$$p_i \geq 0$$

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

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:

$0 \leq P_i \leq 1$

$P_n = 1$

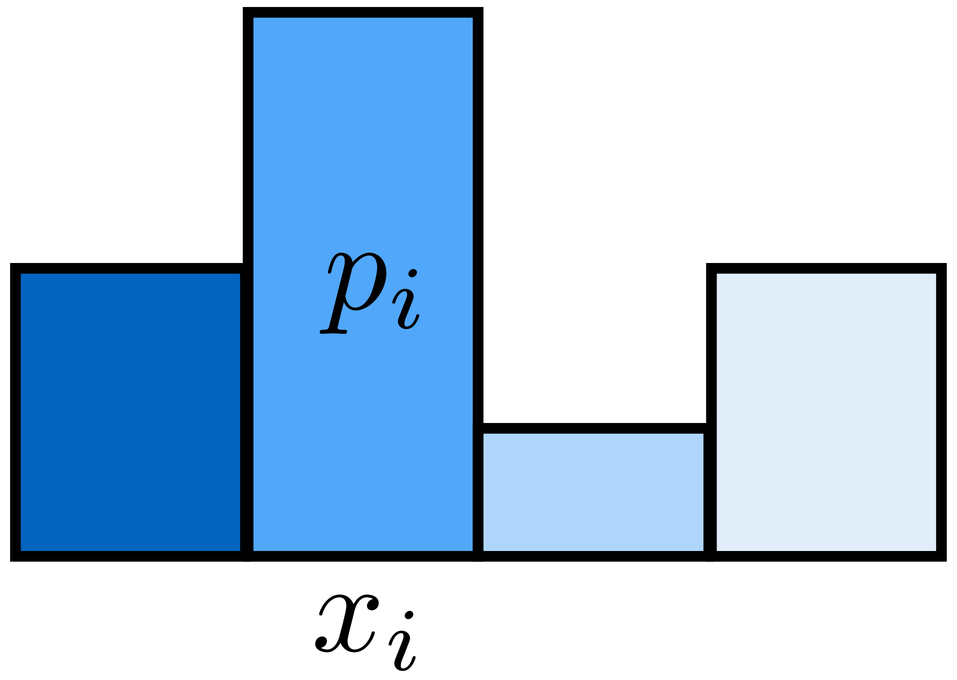

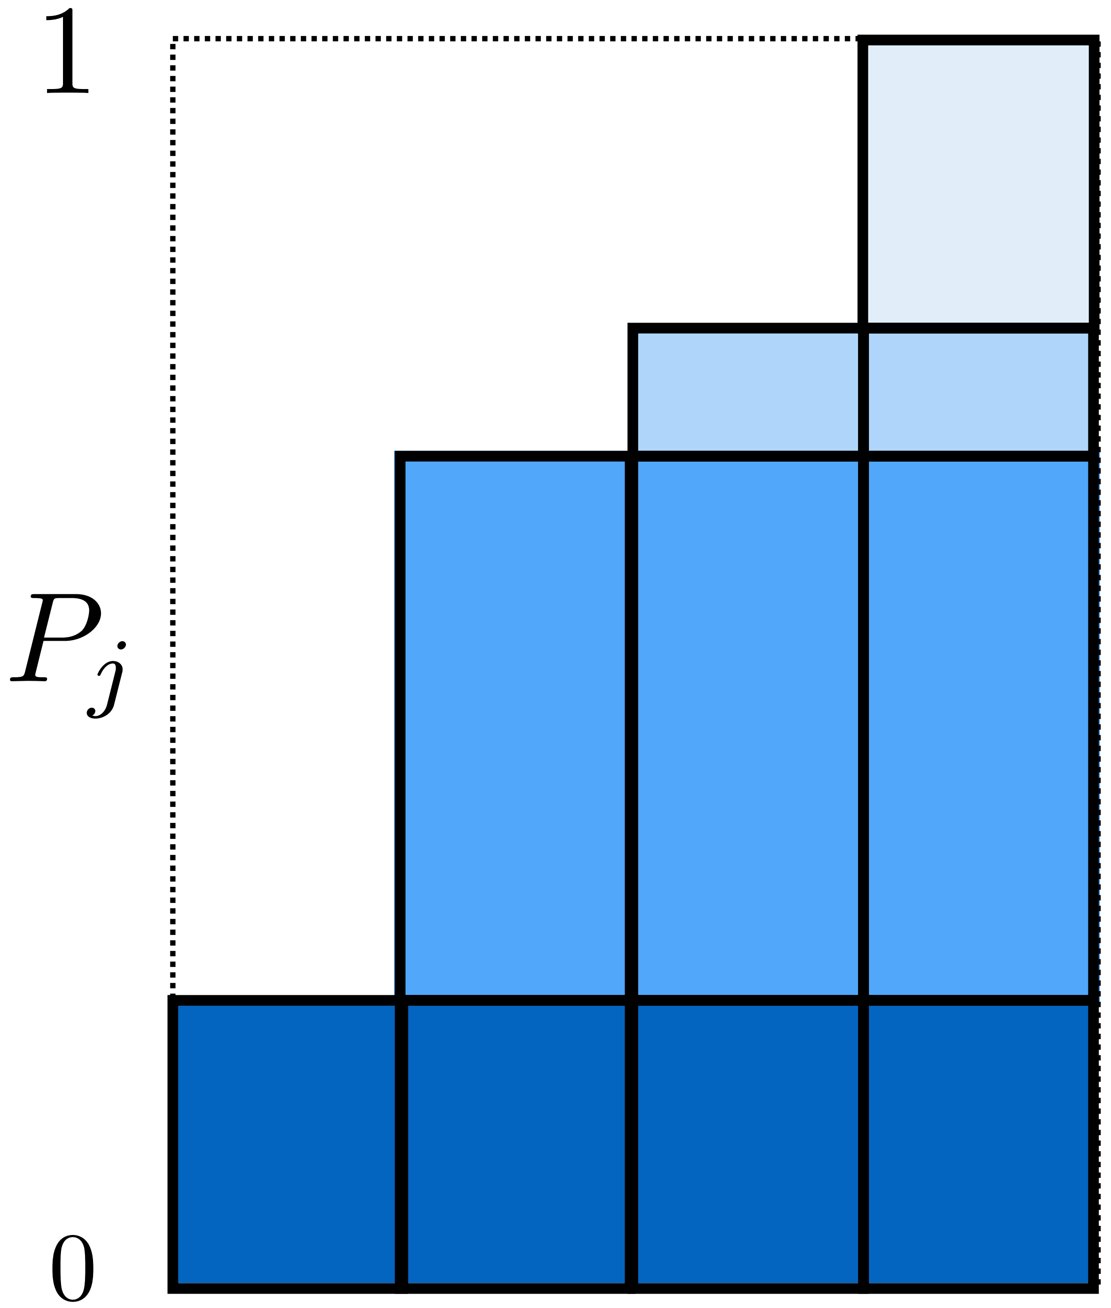

Sampling from discrete probability distributions

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

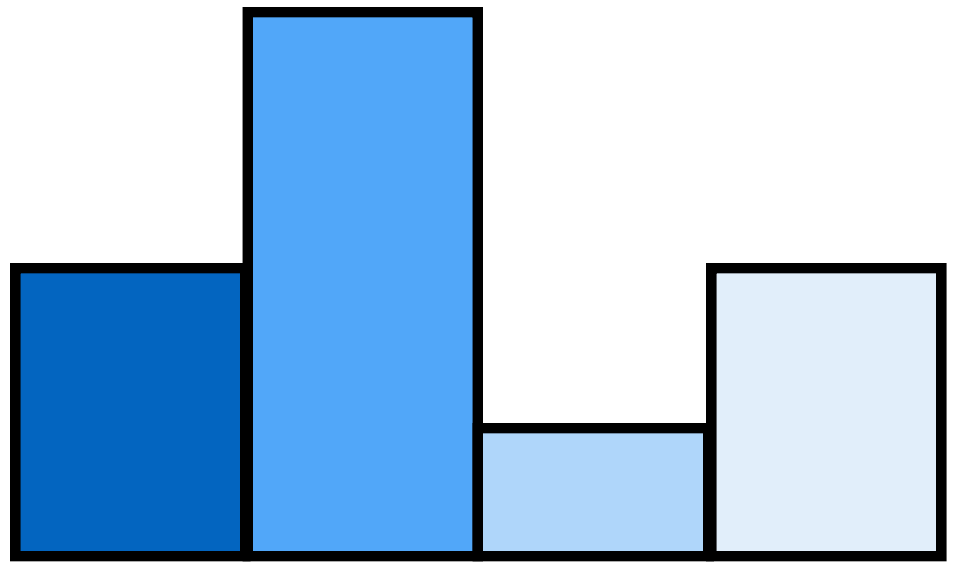

To randomly select an event, select x_i if

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

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

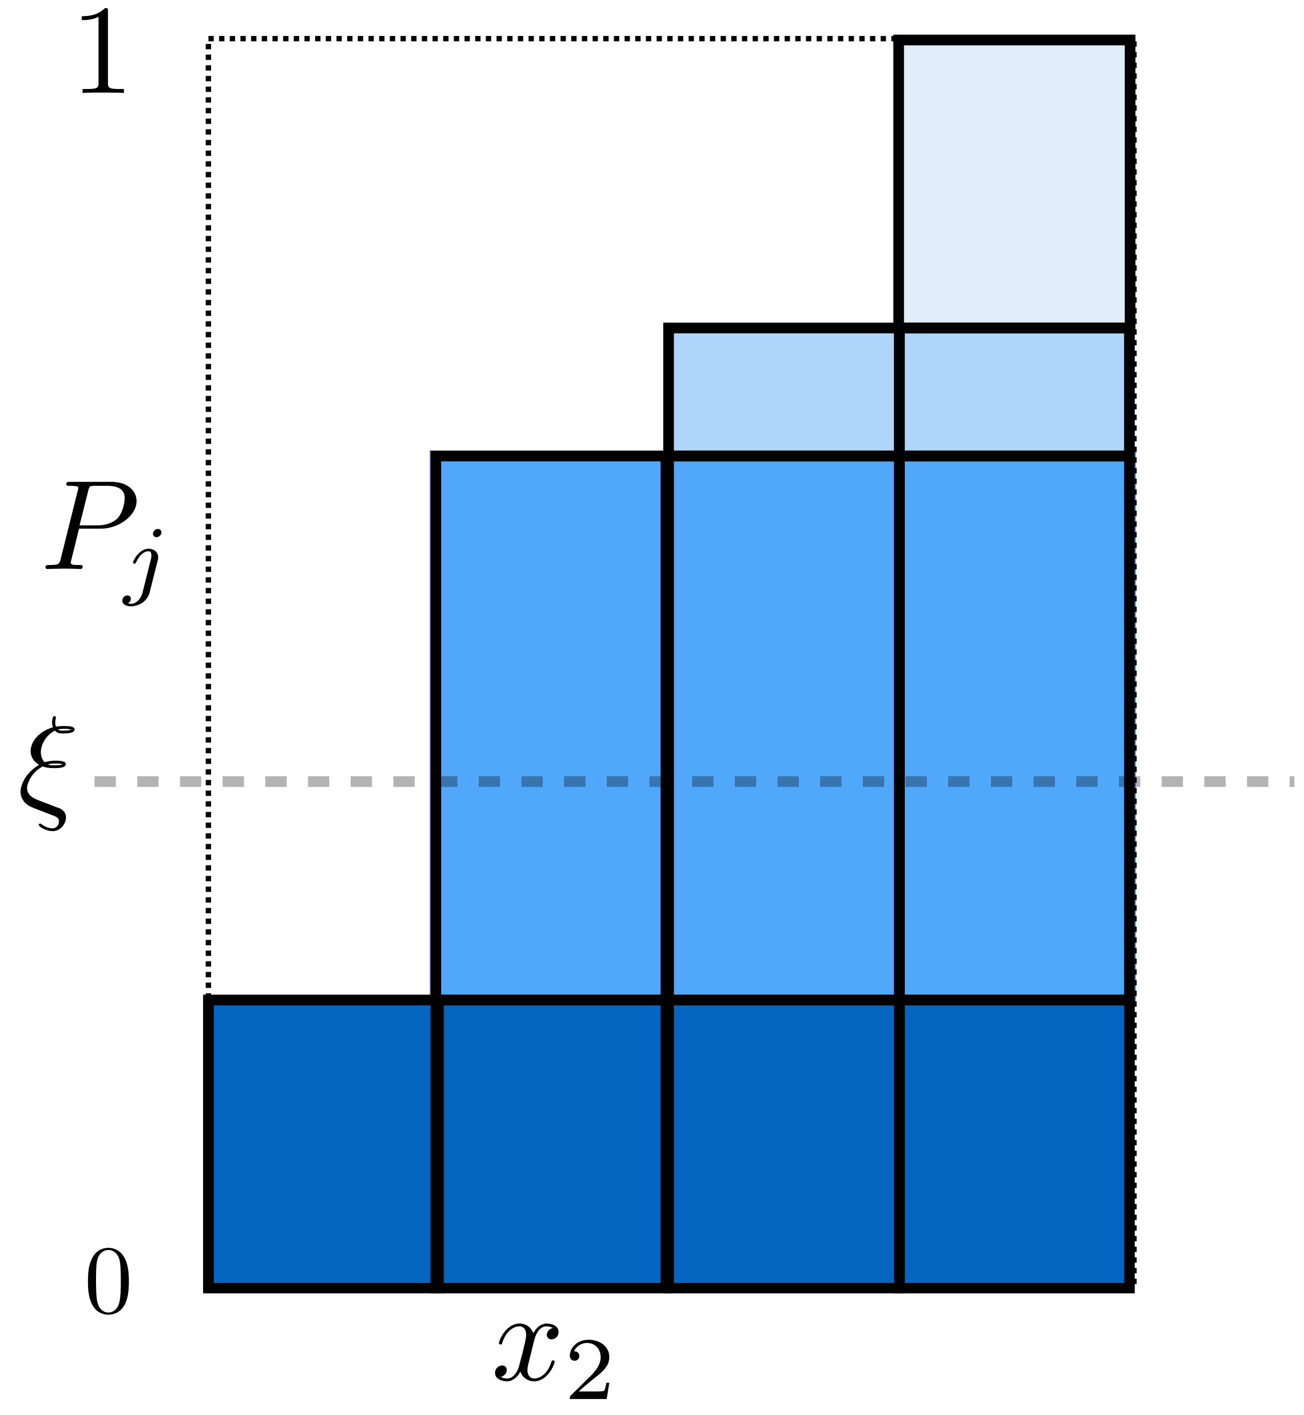

Continuous probability distributions

PDF $p(x)$

$$p(x) \geq 0$$

CDF $P(x)$

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

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

$$P(1) = 1$$

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

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

Uniform distribution: p(x) = c

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

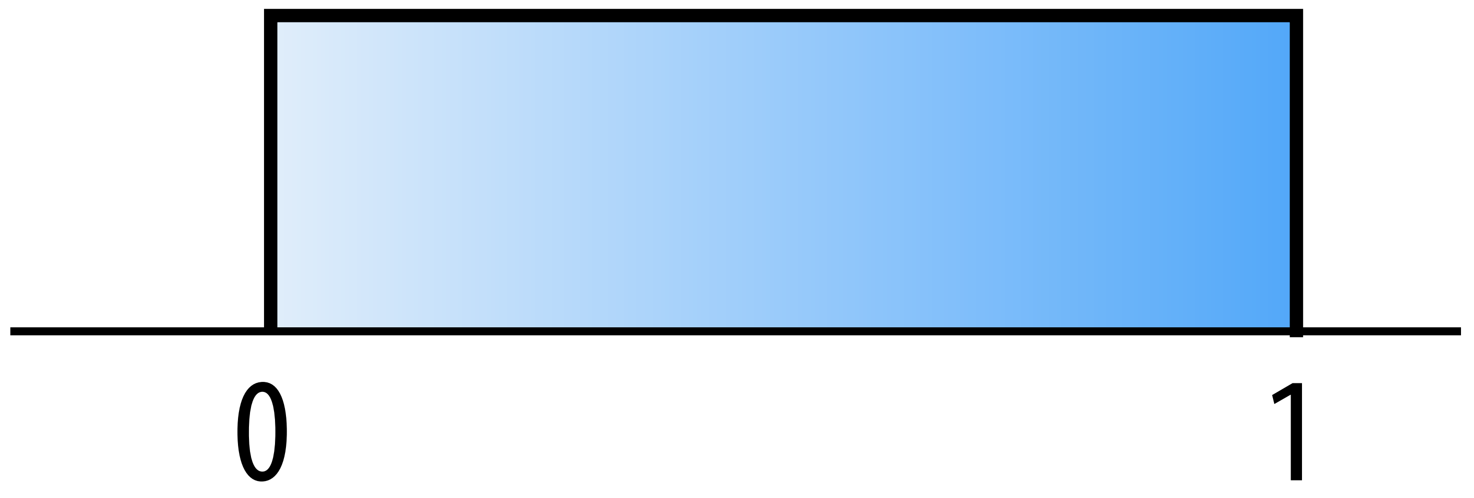

CDF P(x)

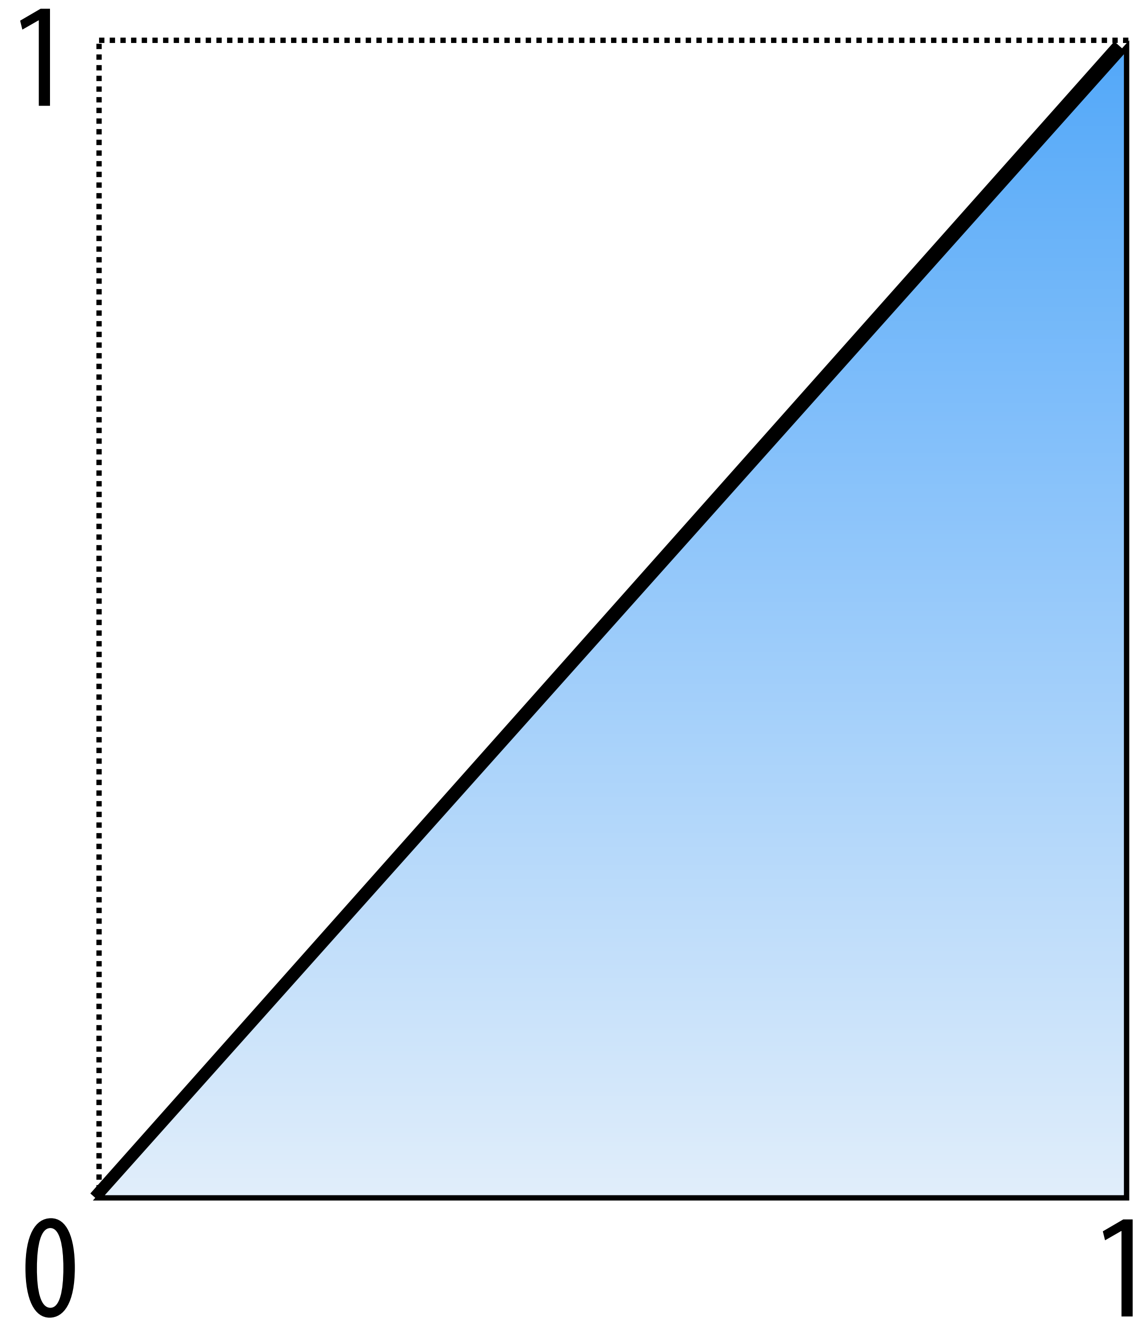

Sampling continuous random variables using the inversion method

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** $P^{-1}(x)$

CDF P(x)

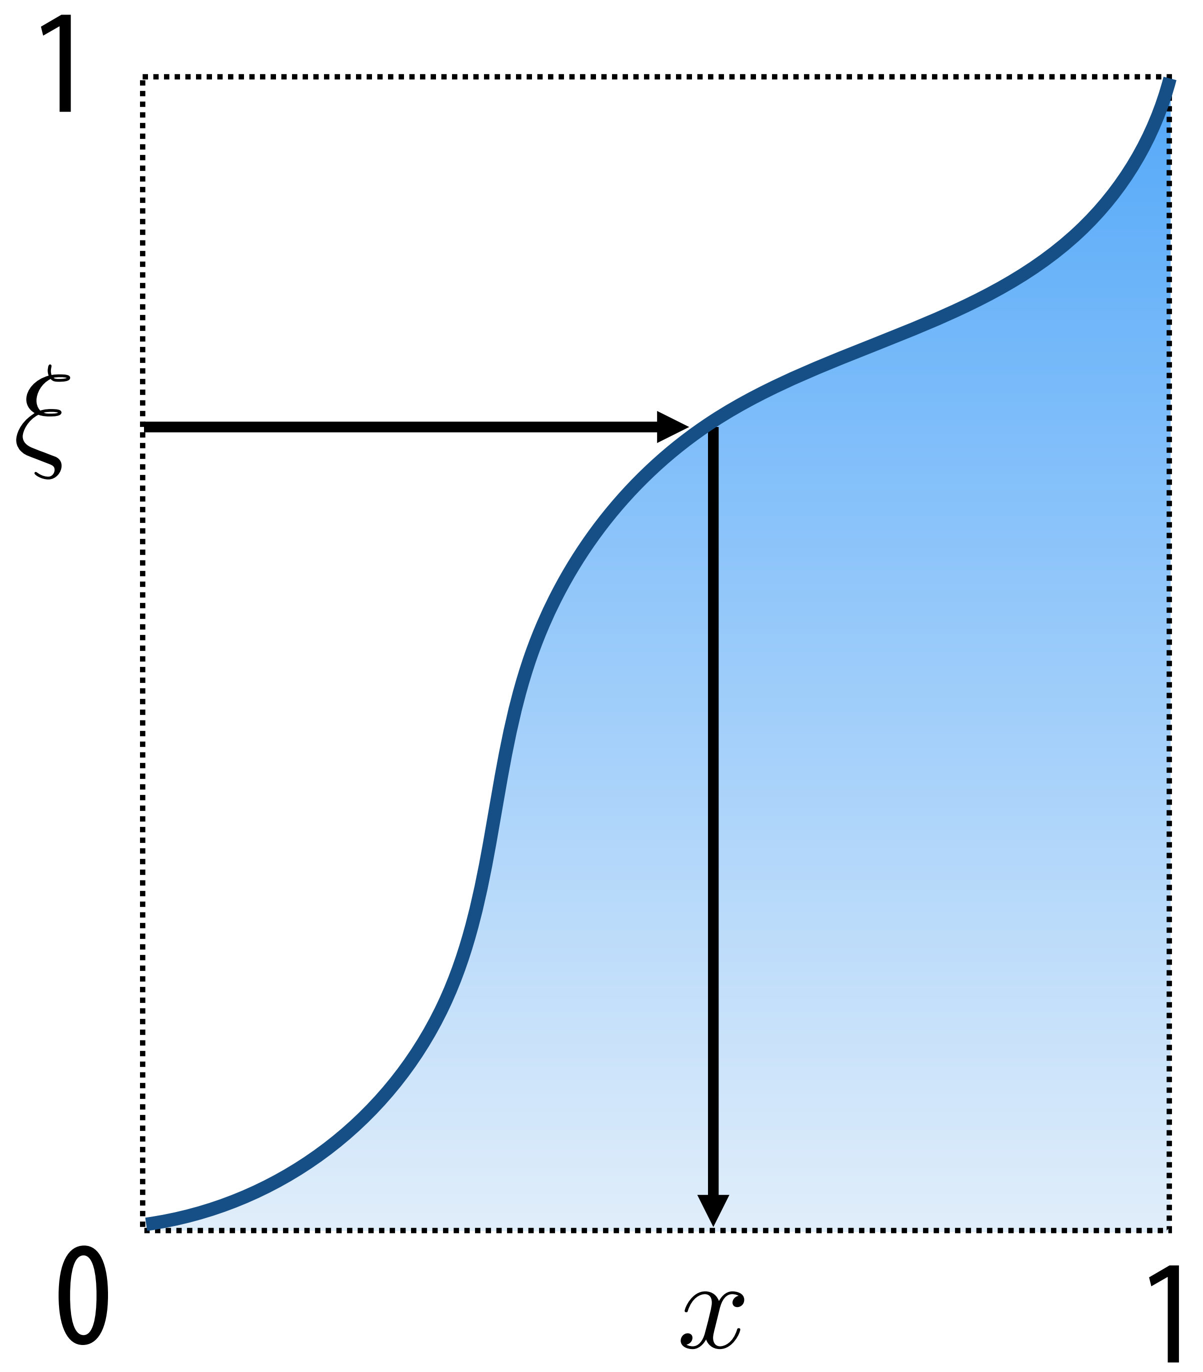

Example: applying the inversion method

Given:

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

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

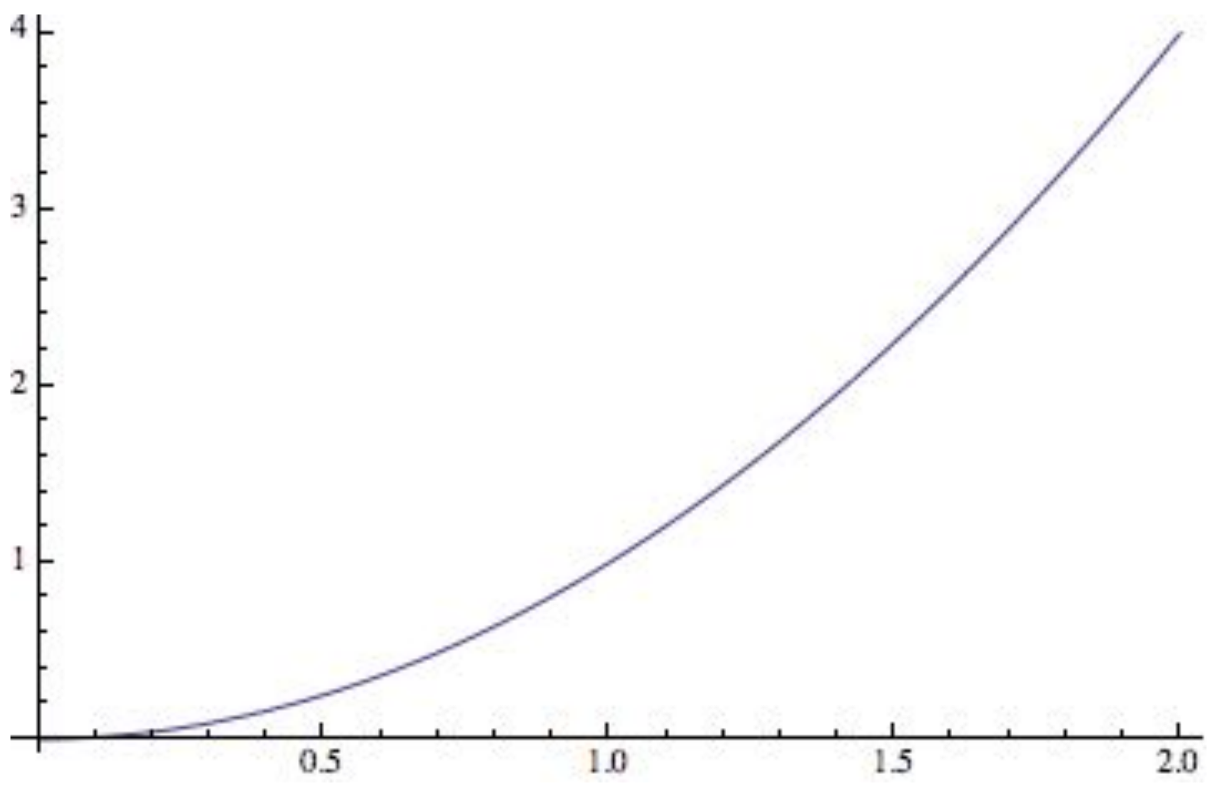

Compute PDF from f(x):

$$\begin{aligned} 1 &= \int_0^2 c\, f(x)\,\mathrm{d}x \\ &= c(F(2) - F(0)) \qquad\qquad F(x) = \frac{1}{3}x^3 \\ &= c\frac{1}{3}2^3 \\ &= \frac{8c}{3} \longrightarrow \quad c = \frac{3}{8}, \qquad p(x) = \frac{3}{8}x^2 \end{aligned}$$

Probability density function (integrates to 1)

Example: applying the inversion method

Given:

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

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

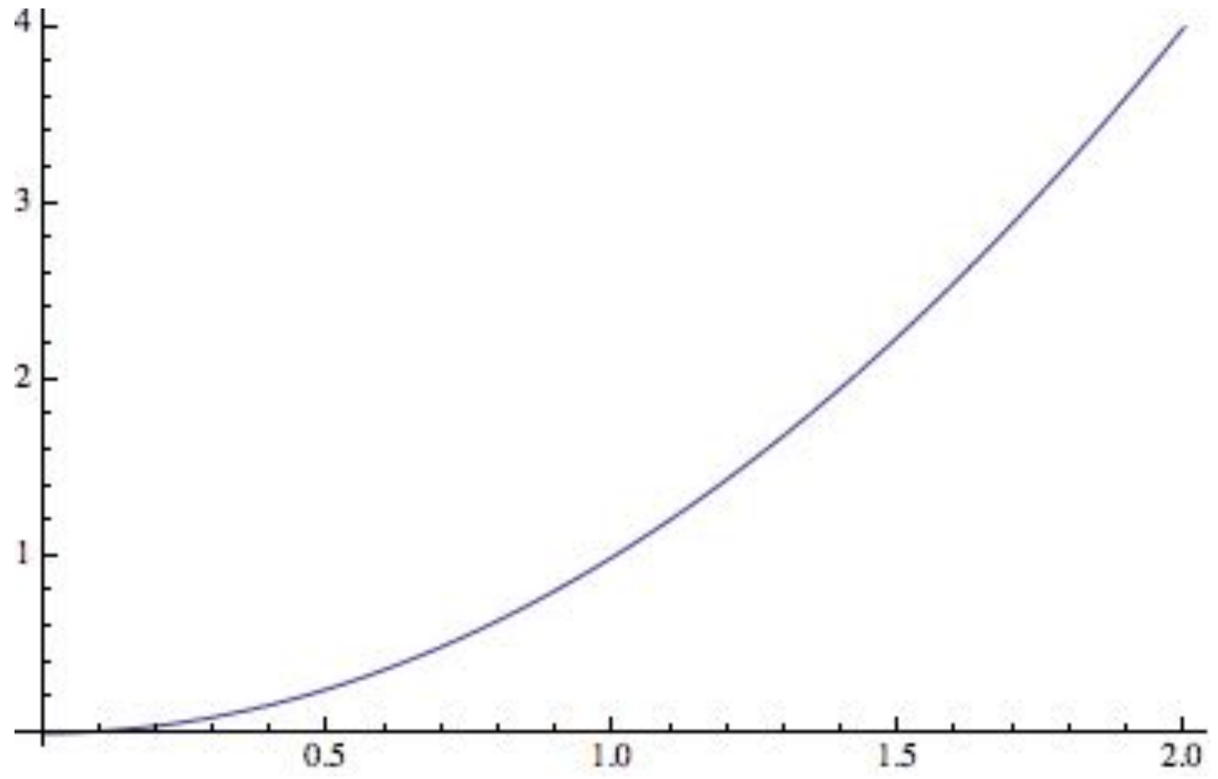

Compute CDF:

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

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

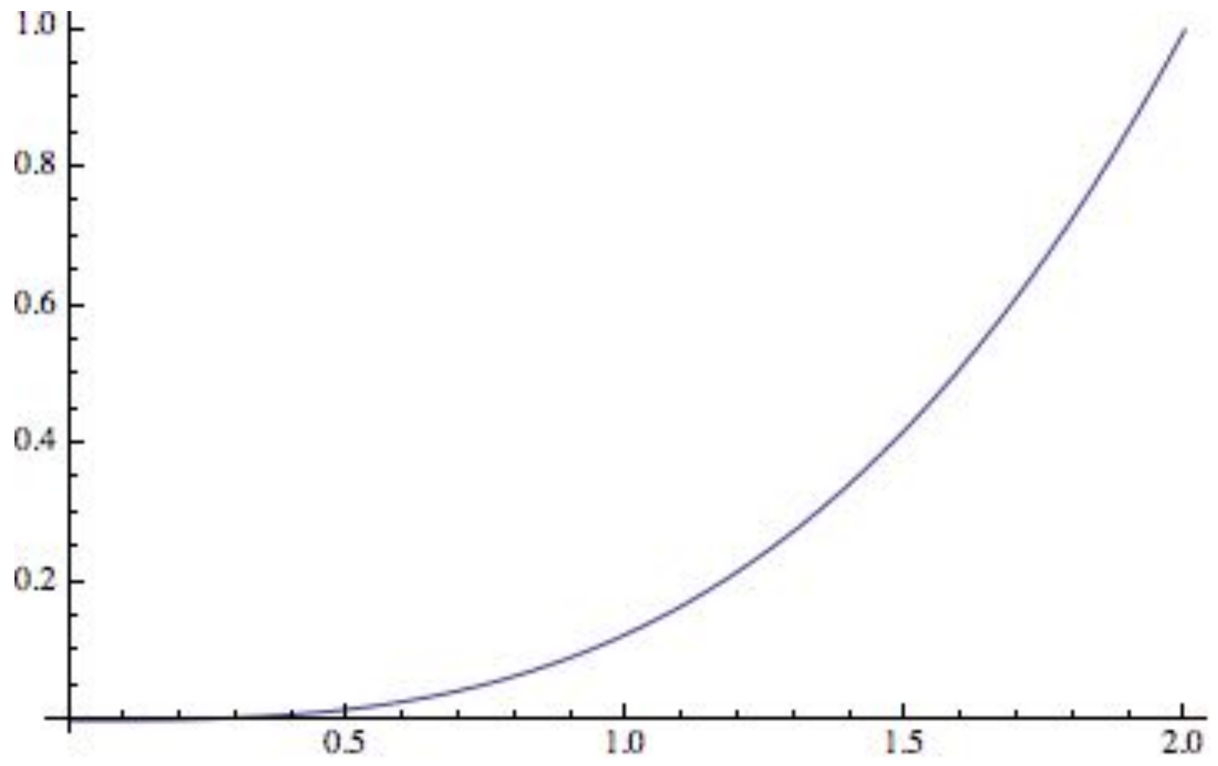

Example: applying the inversion method

Given:

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

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

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

Sample from $p(x)$

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

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

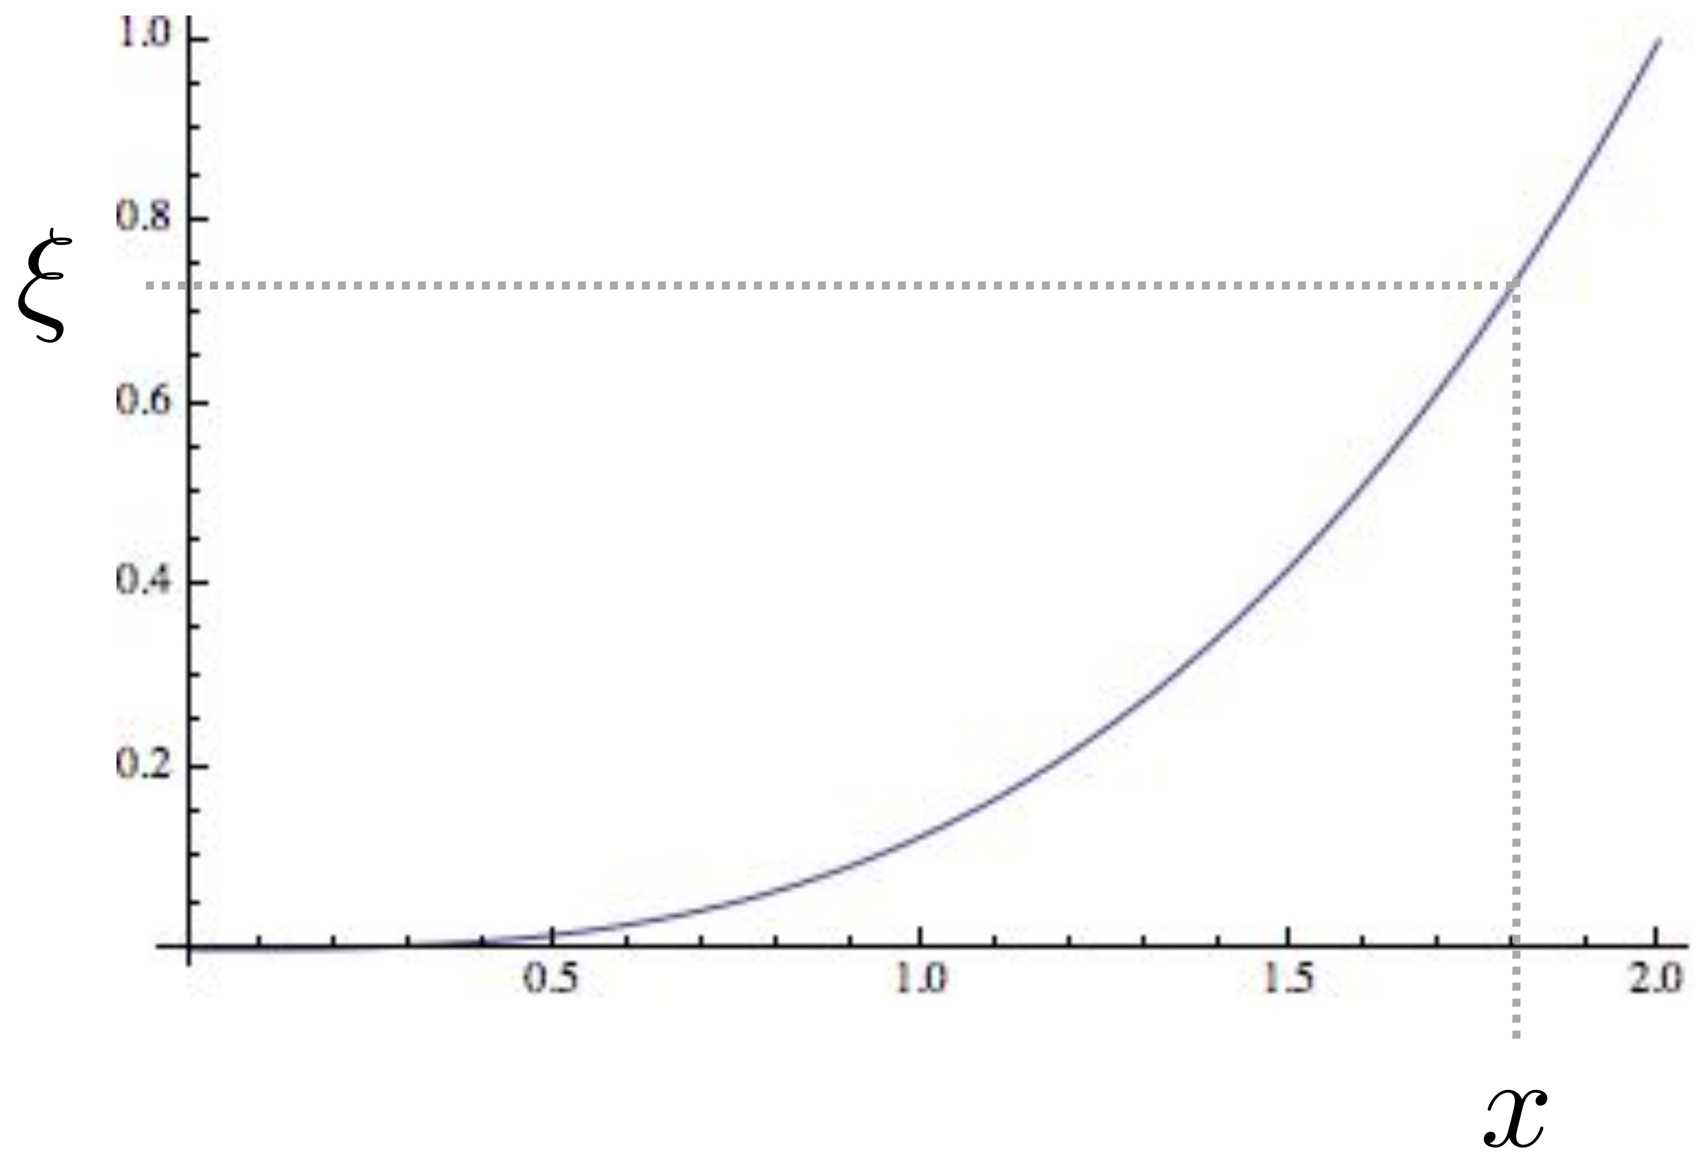

How do we uniformly sample the unit circle?

(Choose any point P=(px, py) 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 <u>does not</u> 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

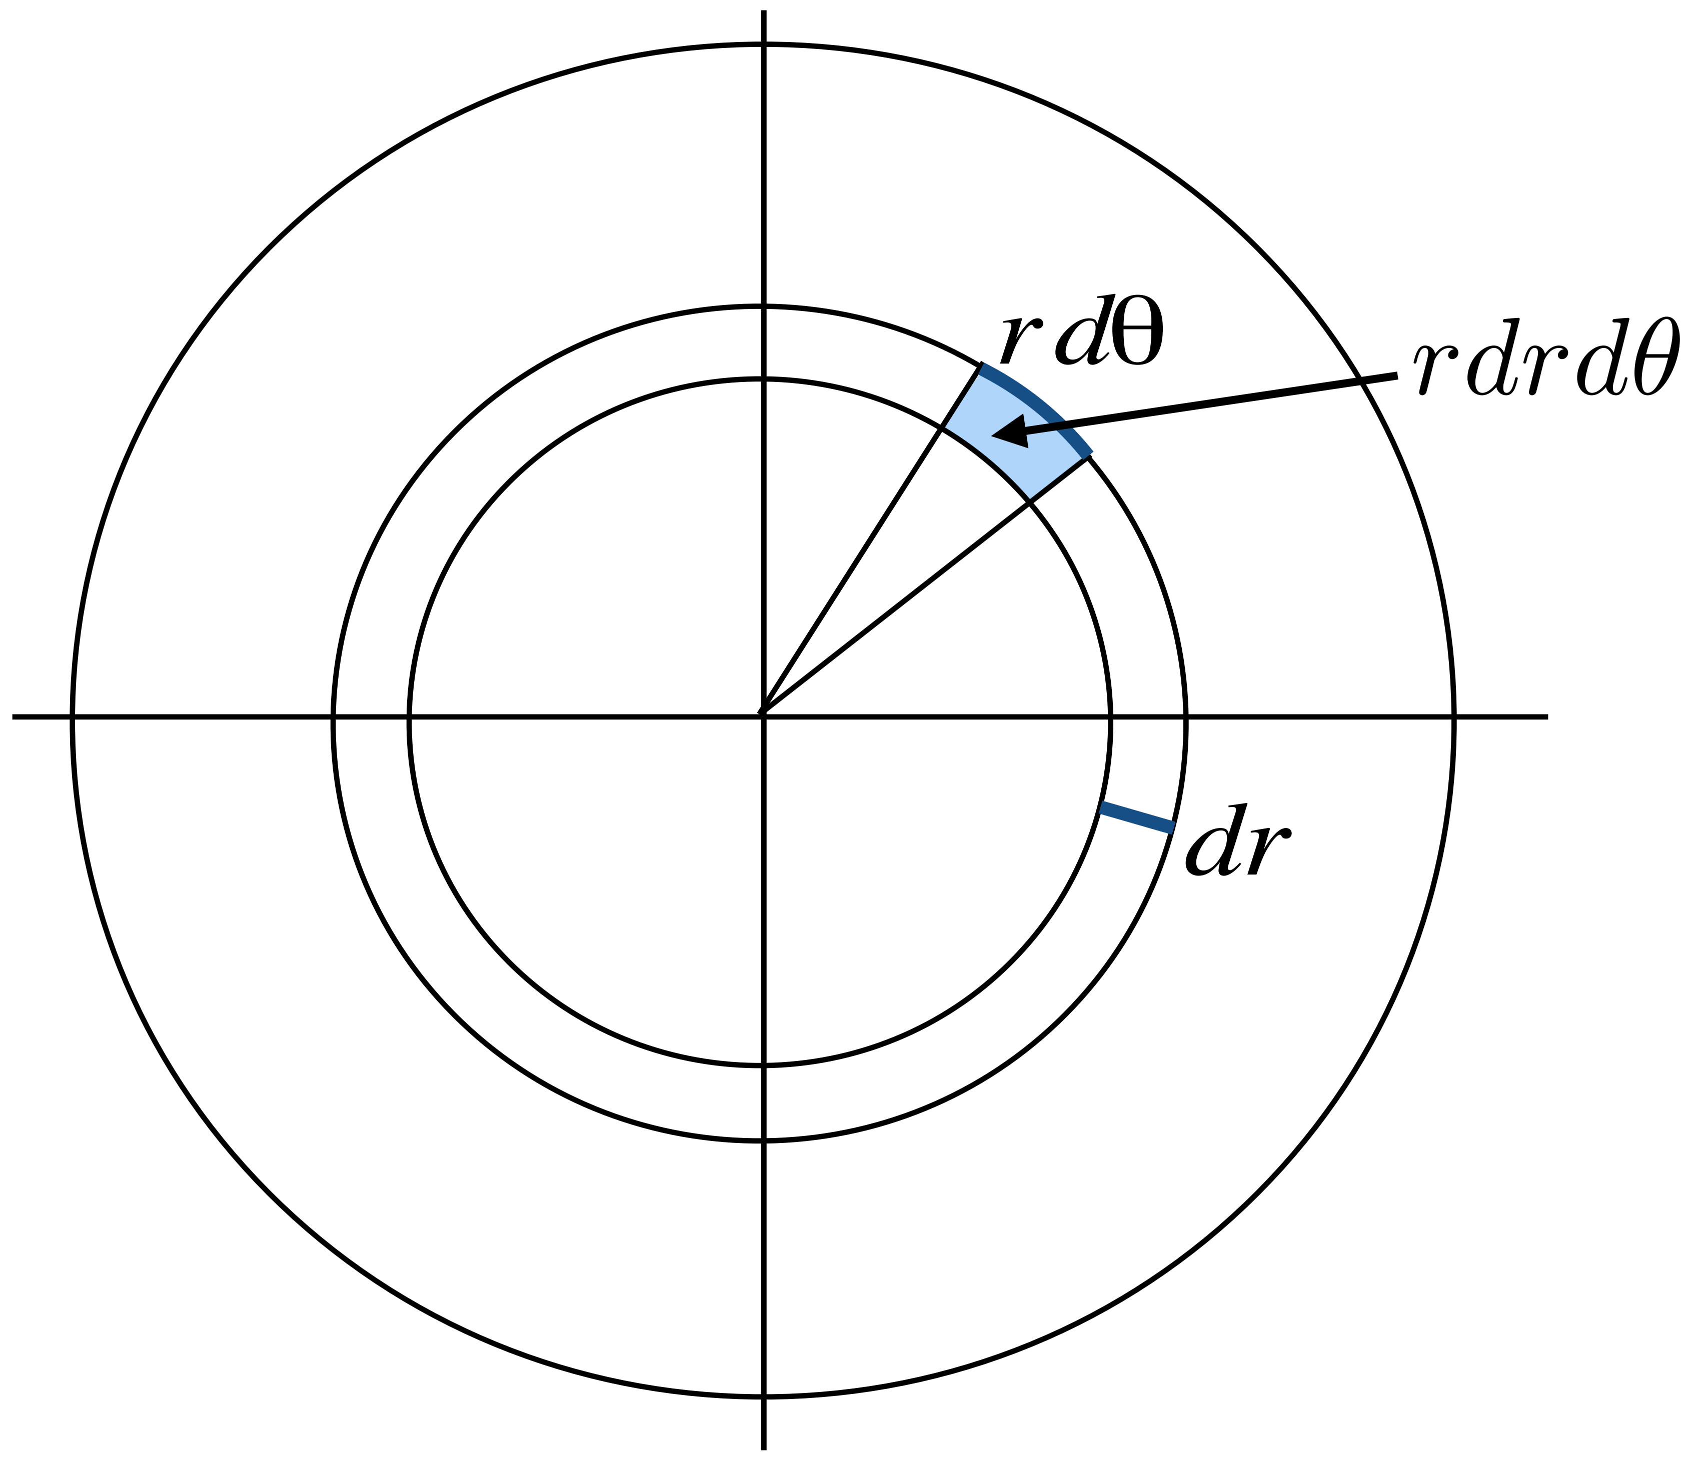

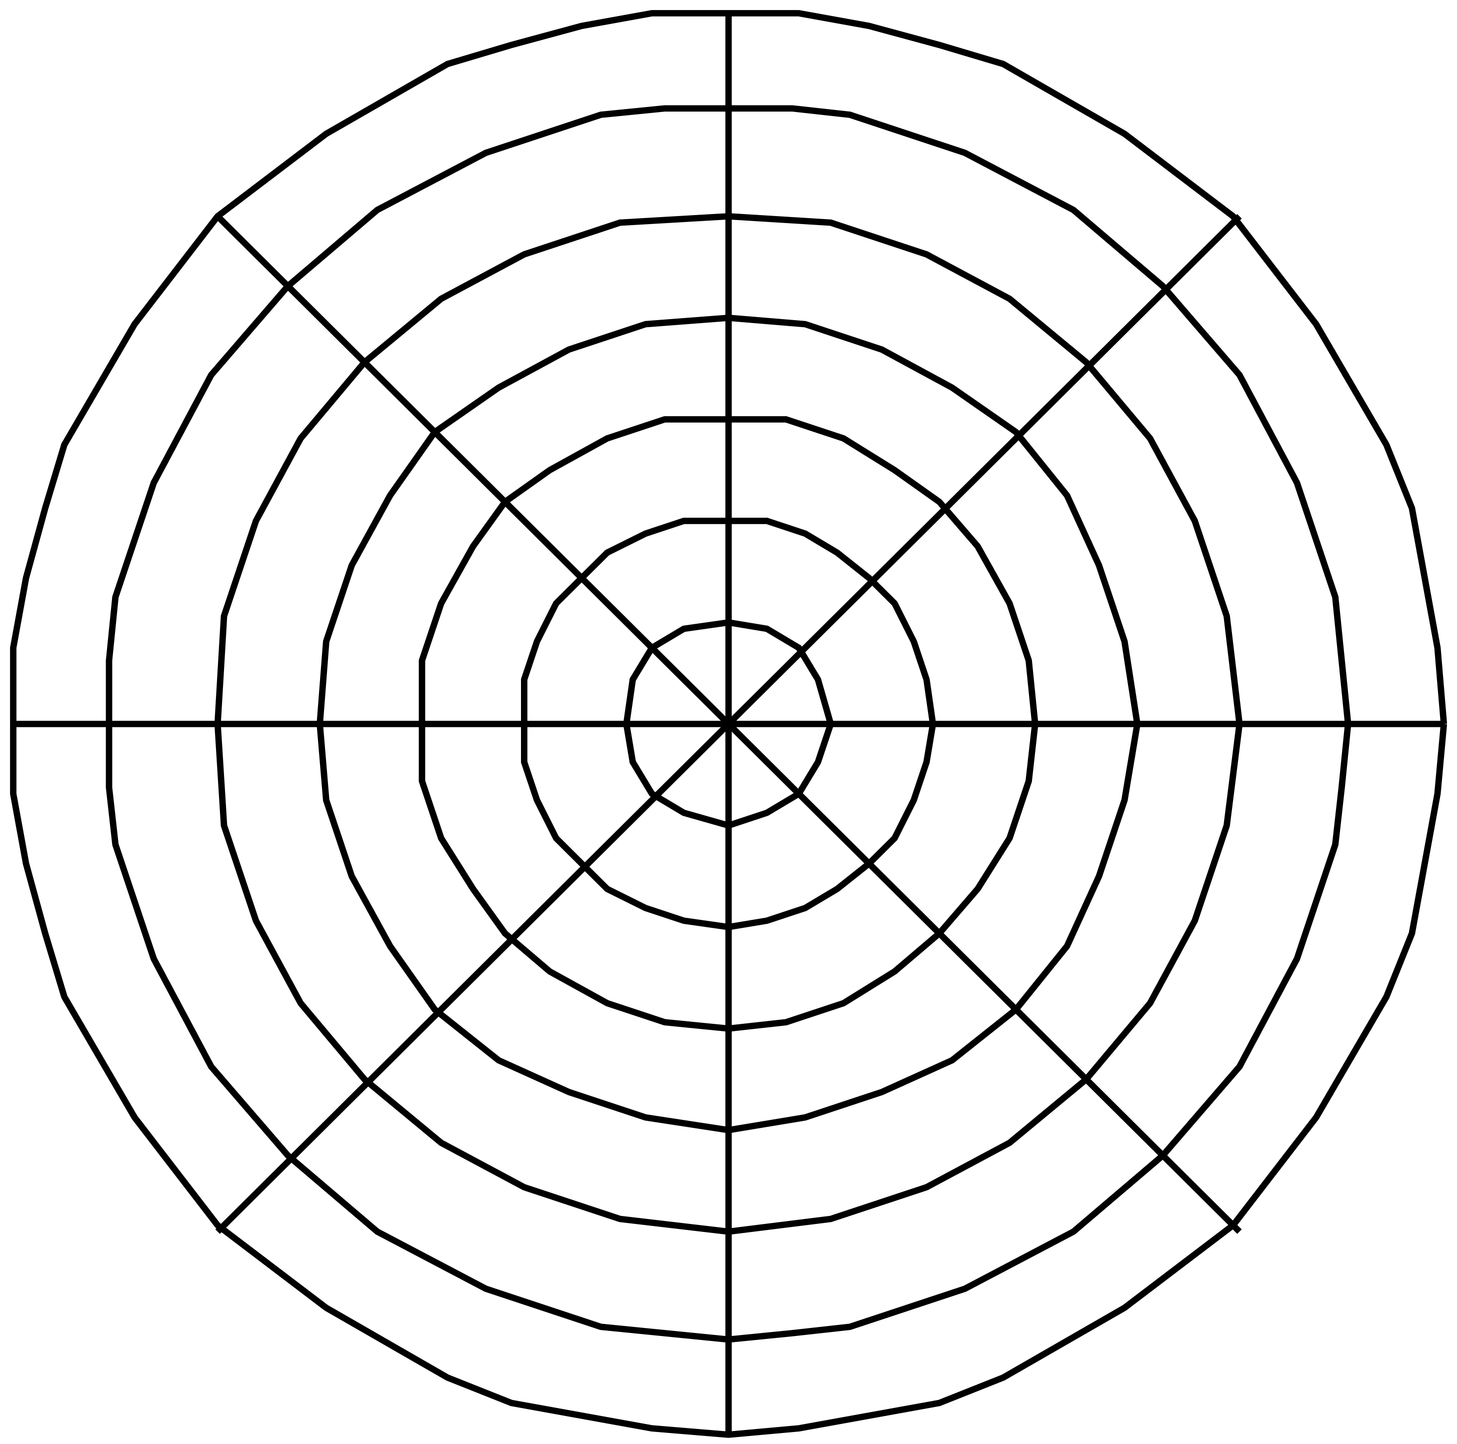

$$\theta = 2\pi\xi_1 \qquad r = \xi_2$$

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

Uniform area sampling of a circle

WRONG
Not Equi-areal

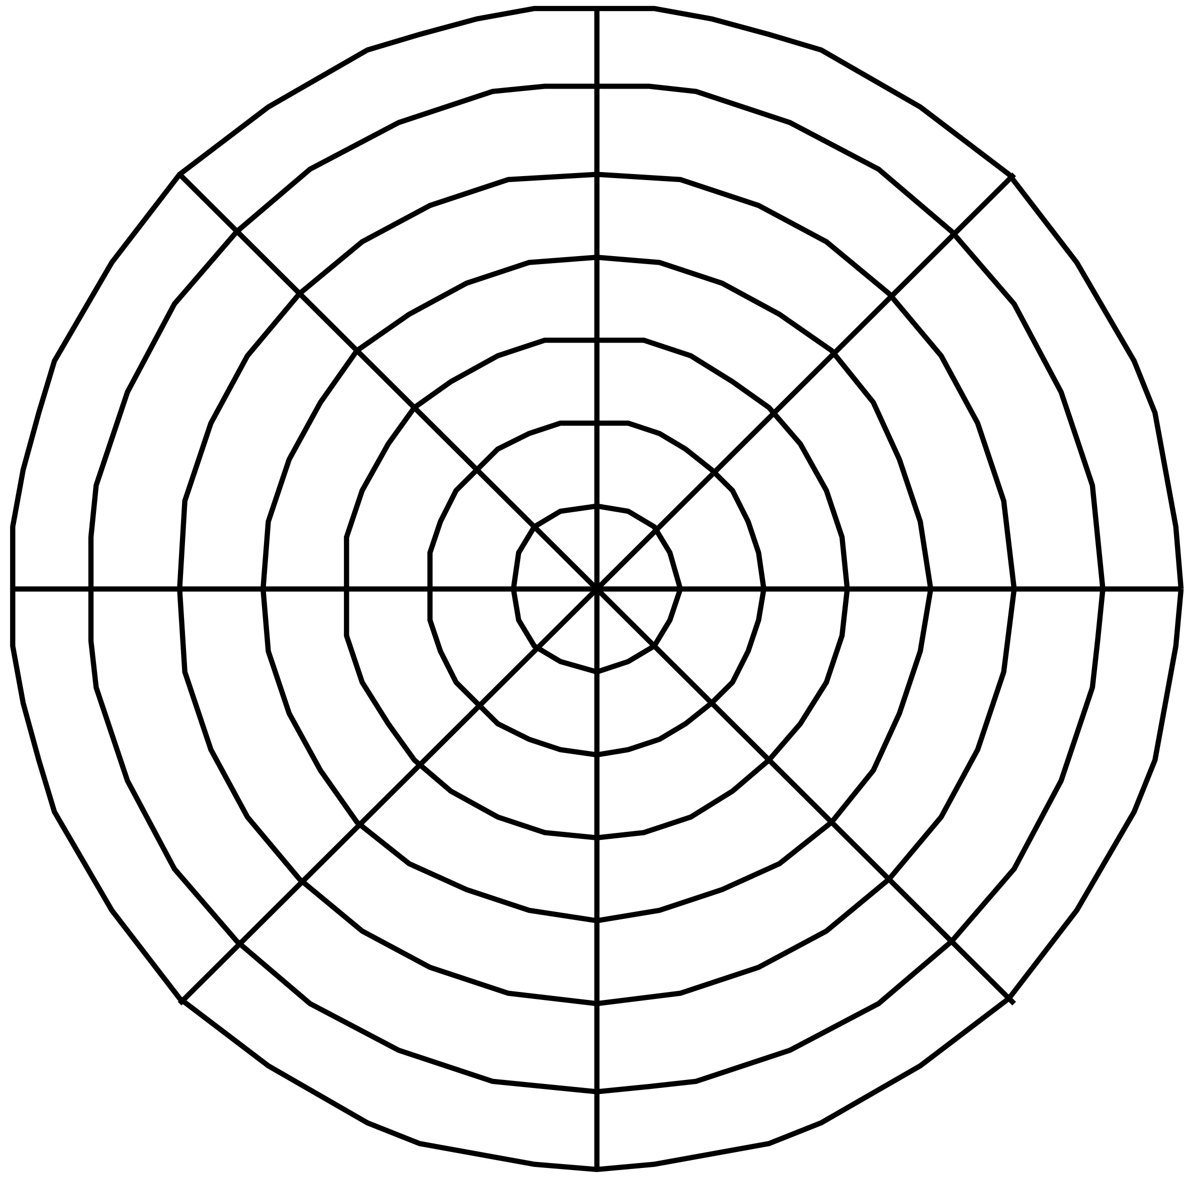

$$\theta = 2\pi\xi_1$$

$$r = \xi_2$$

RIGHT
Equi-areal

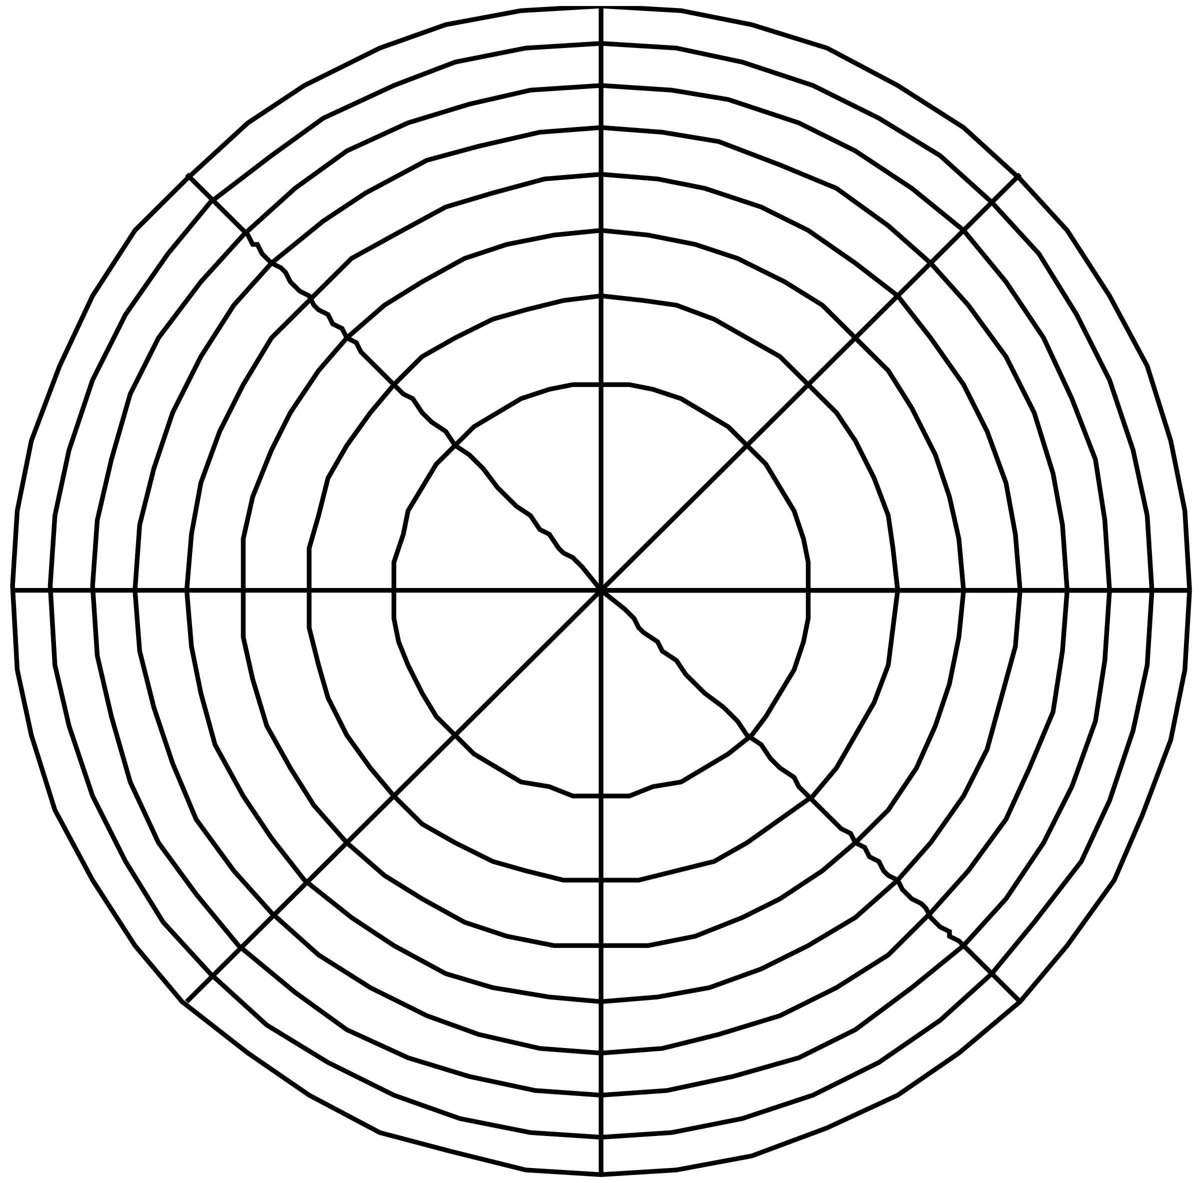

$$\theta = 2\pi\xi_1$$

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

Sampling a circle (via inversion in 2D)

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

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

$$p(r,\theta) = p(r)p(\theta)$$ ⟵ r, θ **independent**

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

$$P(\theta) = \frac{1}{2\pi}\theta \qquad \theta = 2\pi\xi_1$$

$$p(r) = 2r$$

$$P(r) = r^2 \qquad r = \sqrt{\xi_2}$$

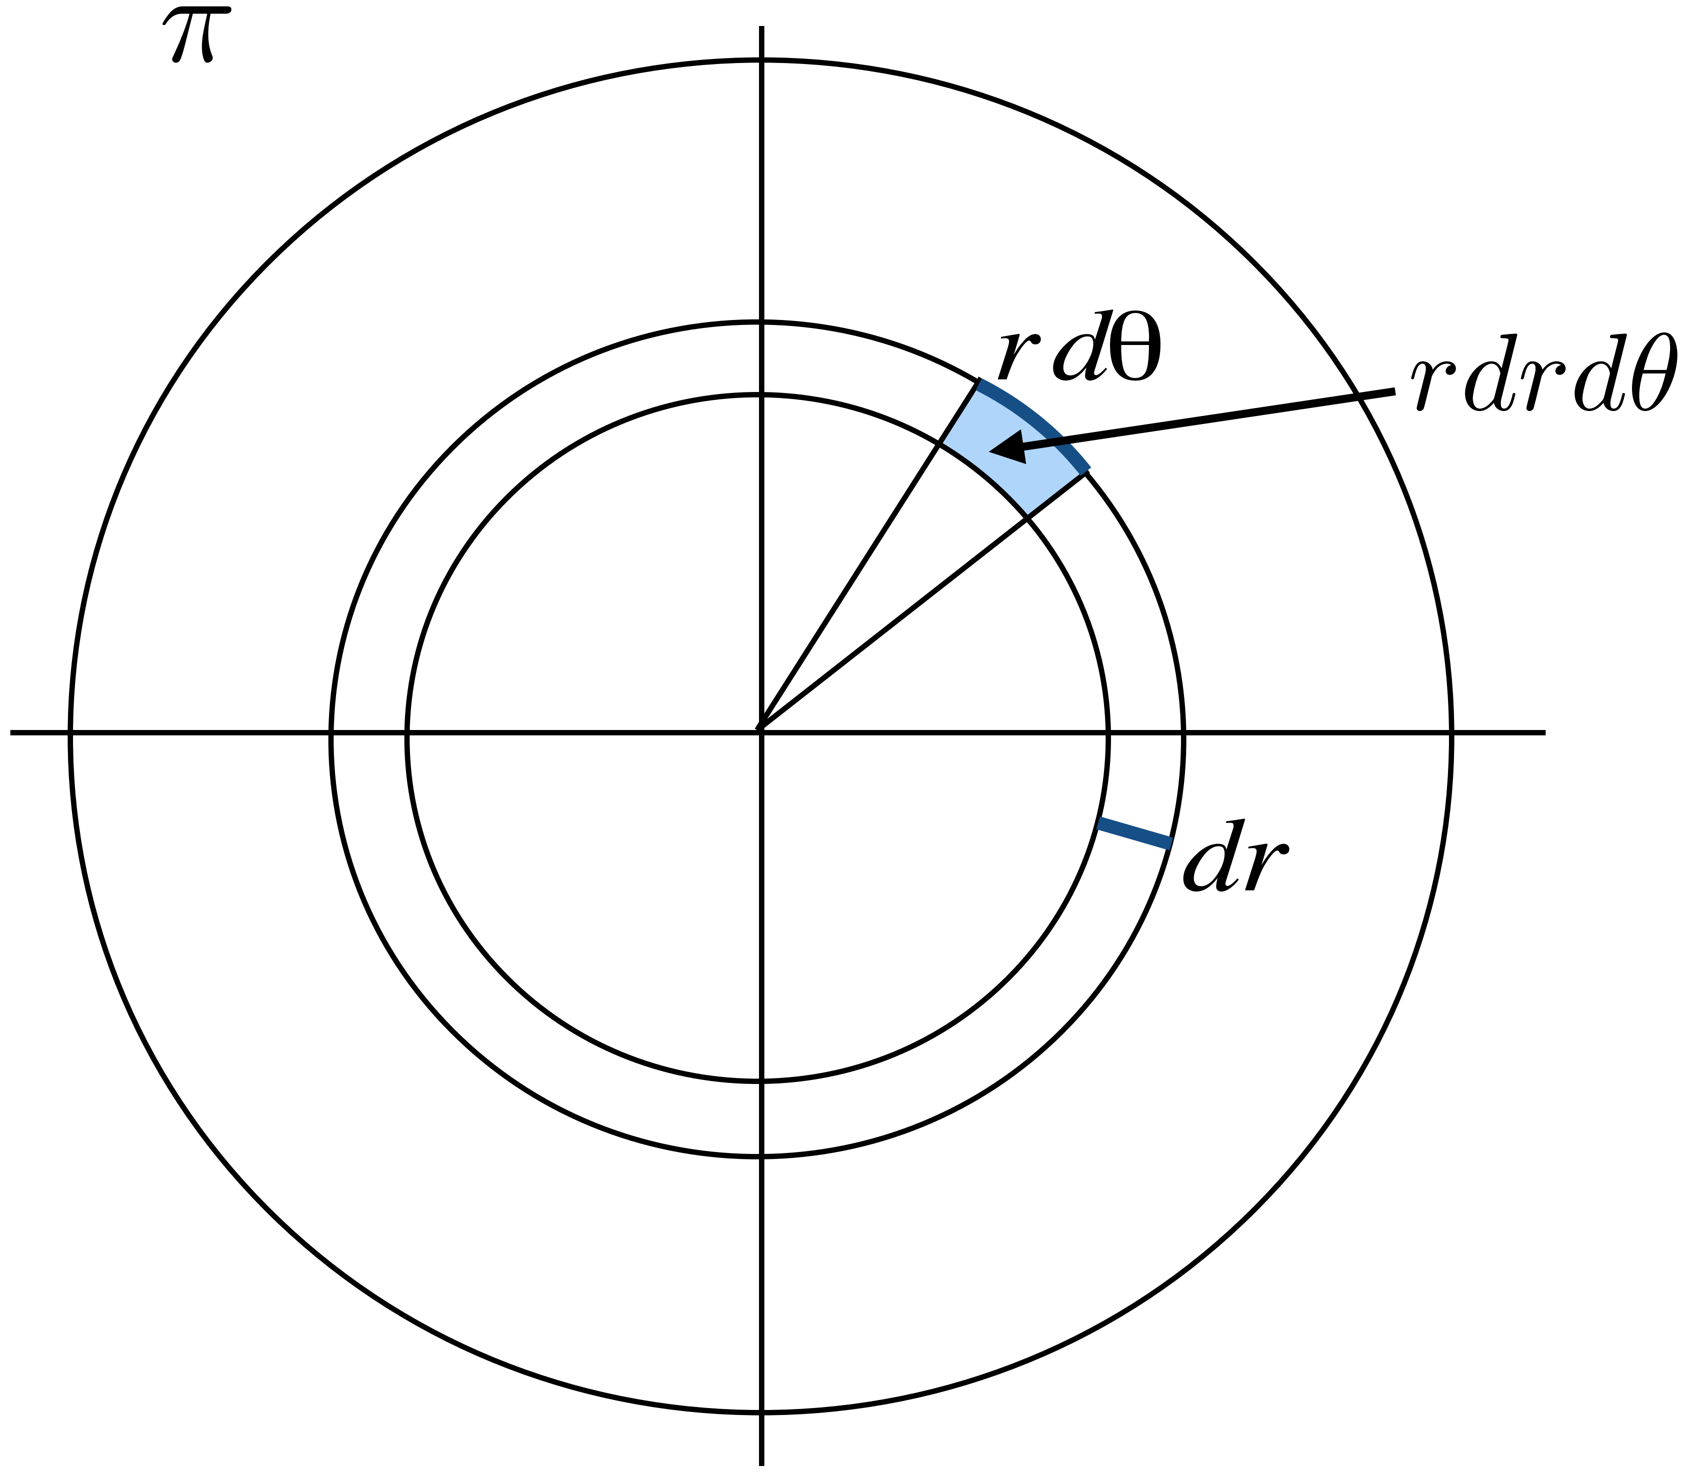

Shirley's mapping

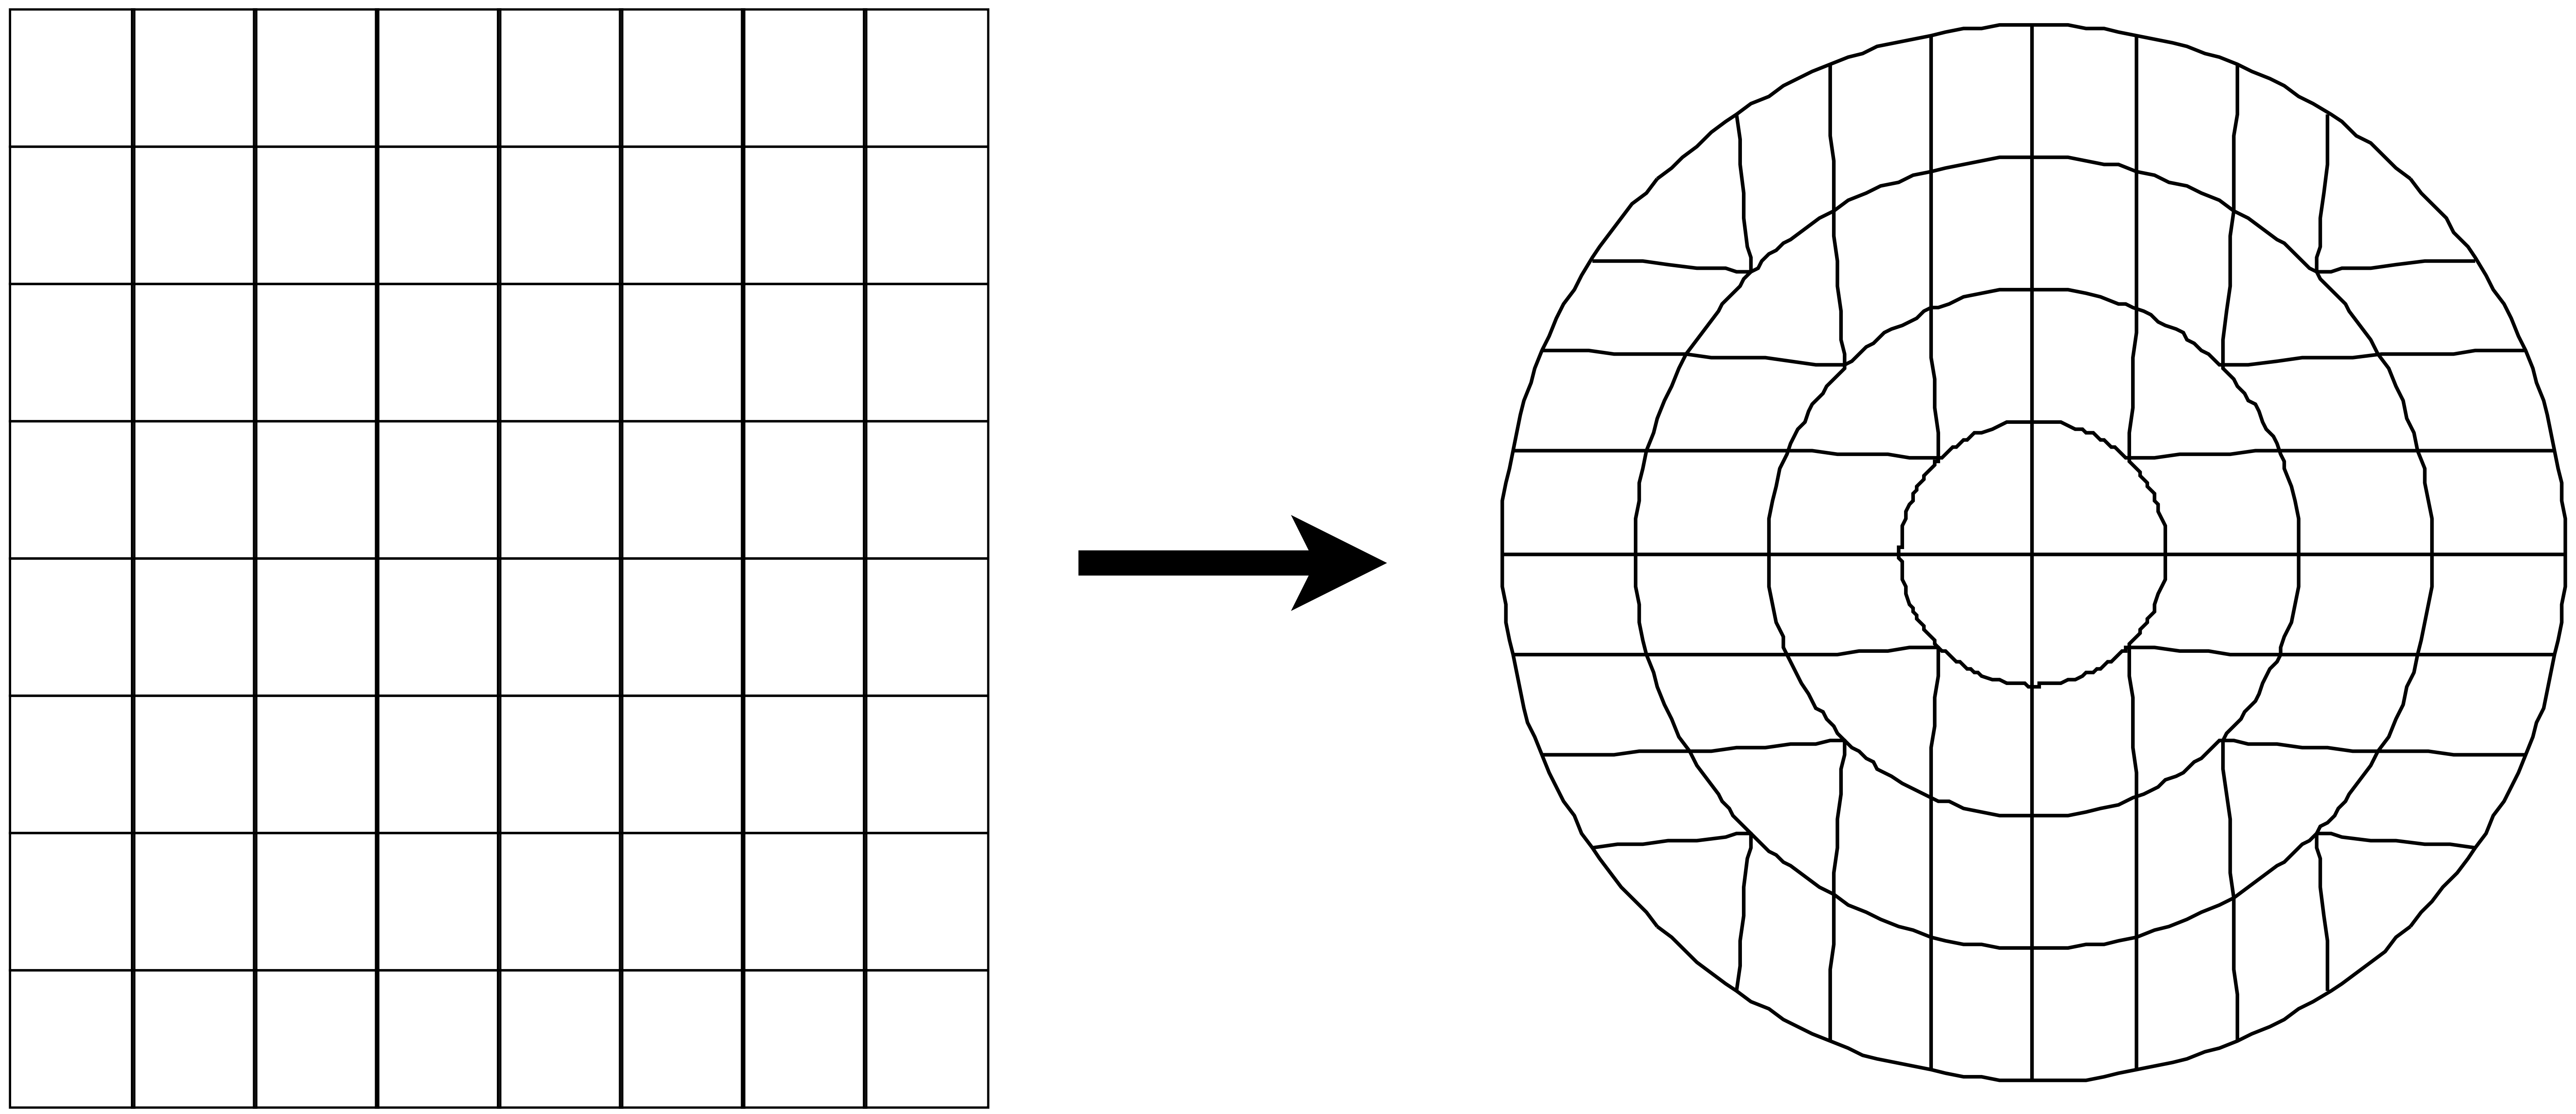

Distinct cases for eight octants

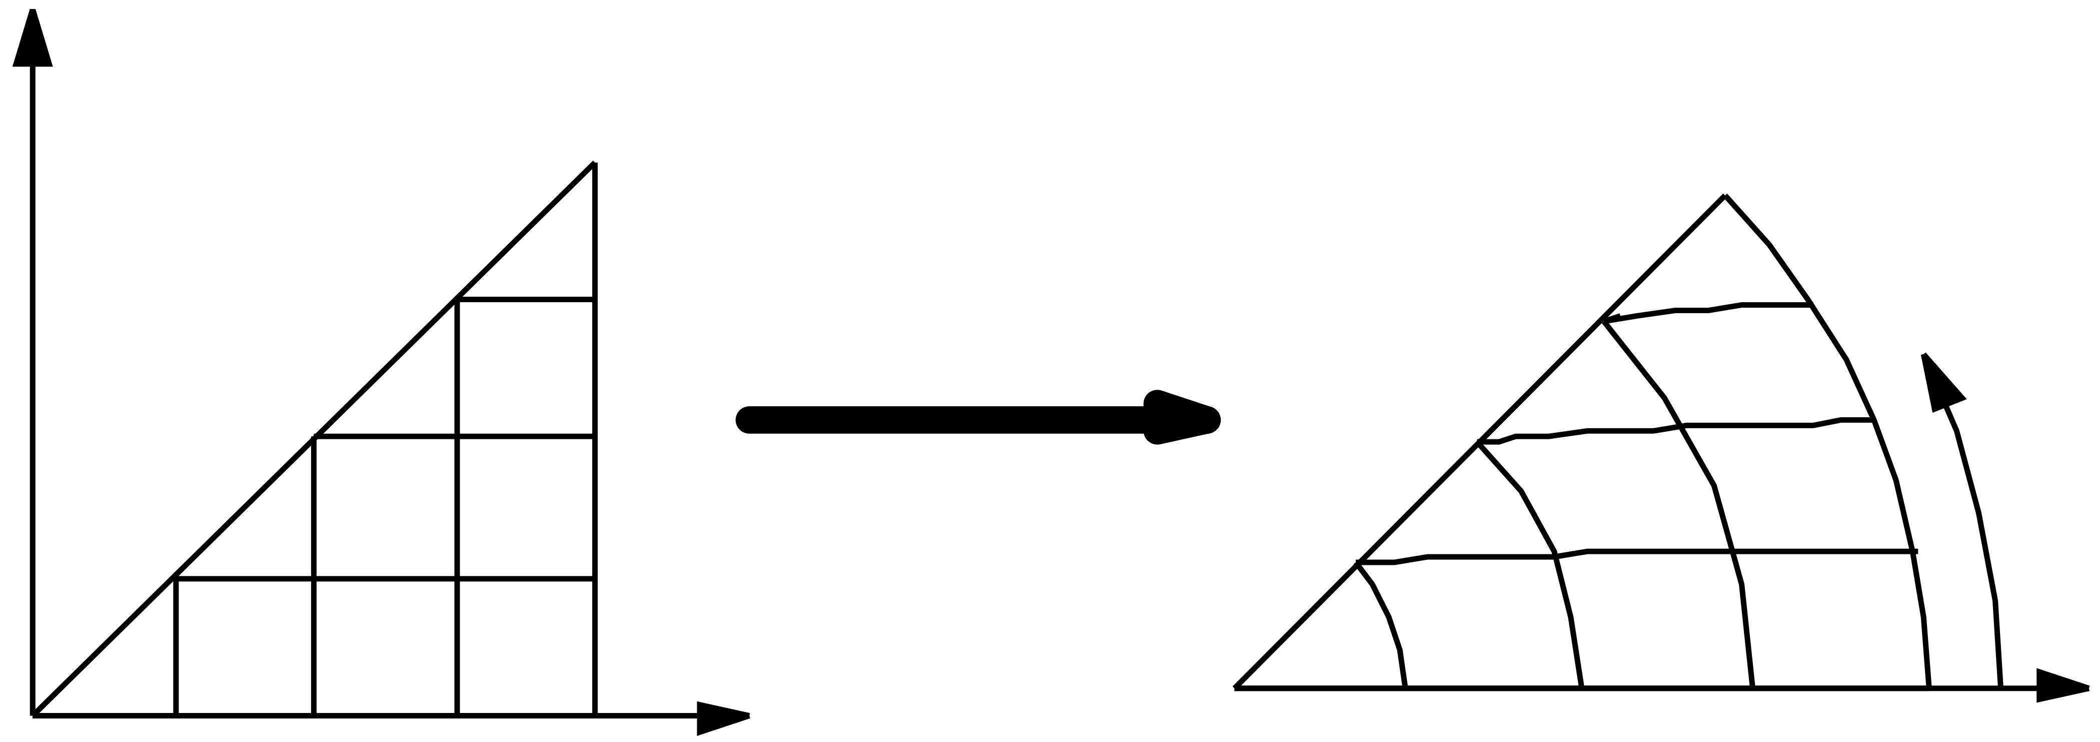

$$r = \xi_1$$

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

Uniform sampling via rejection sampling

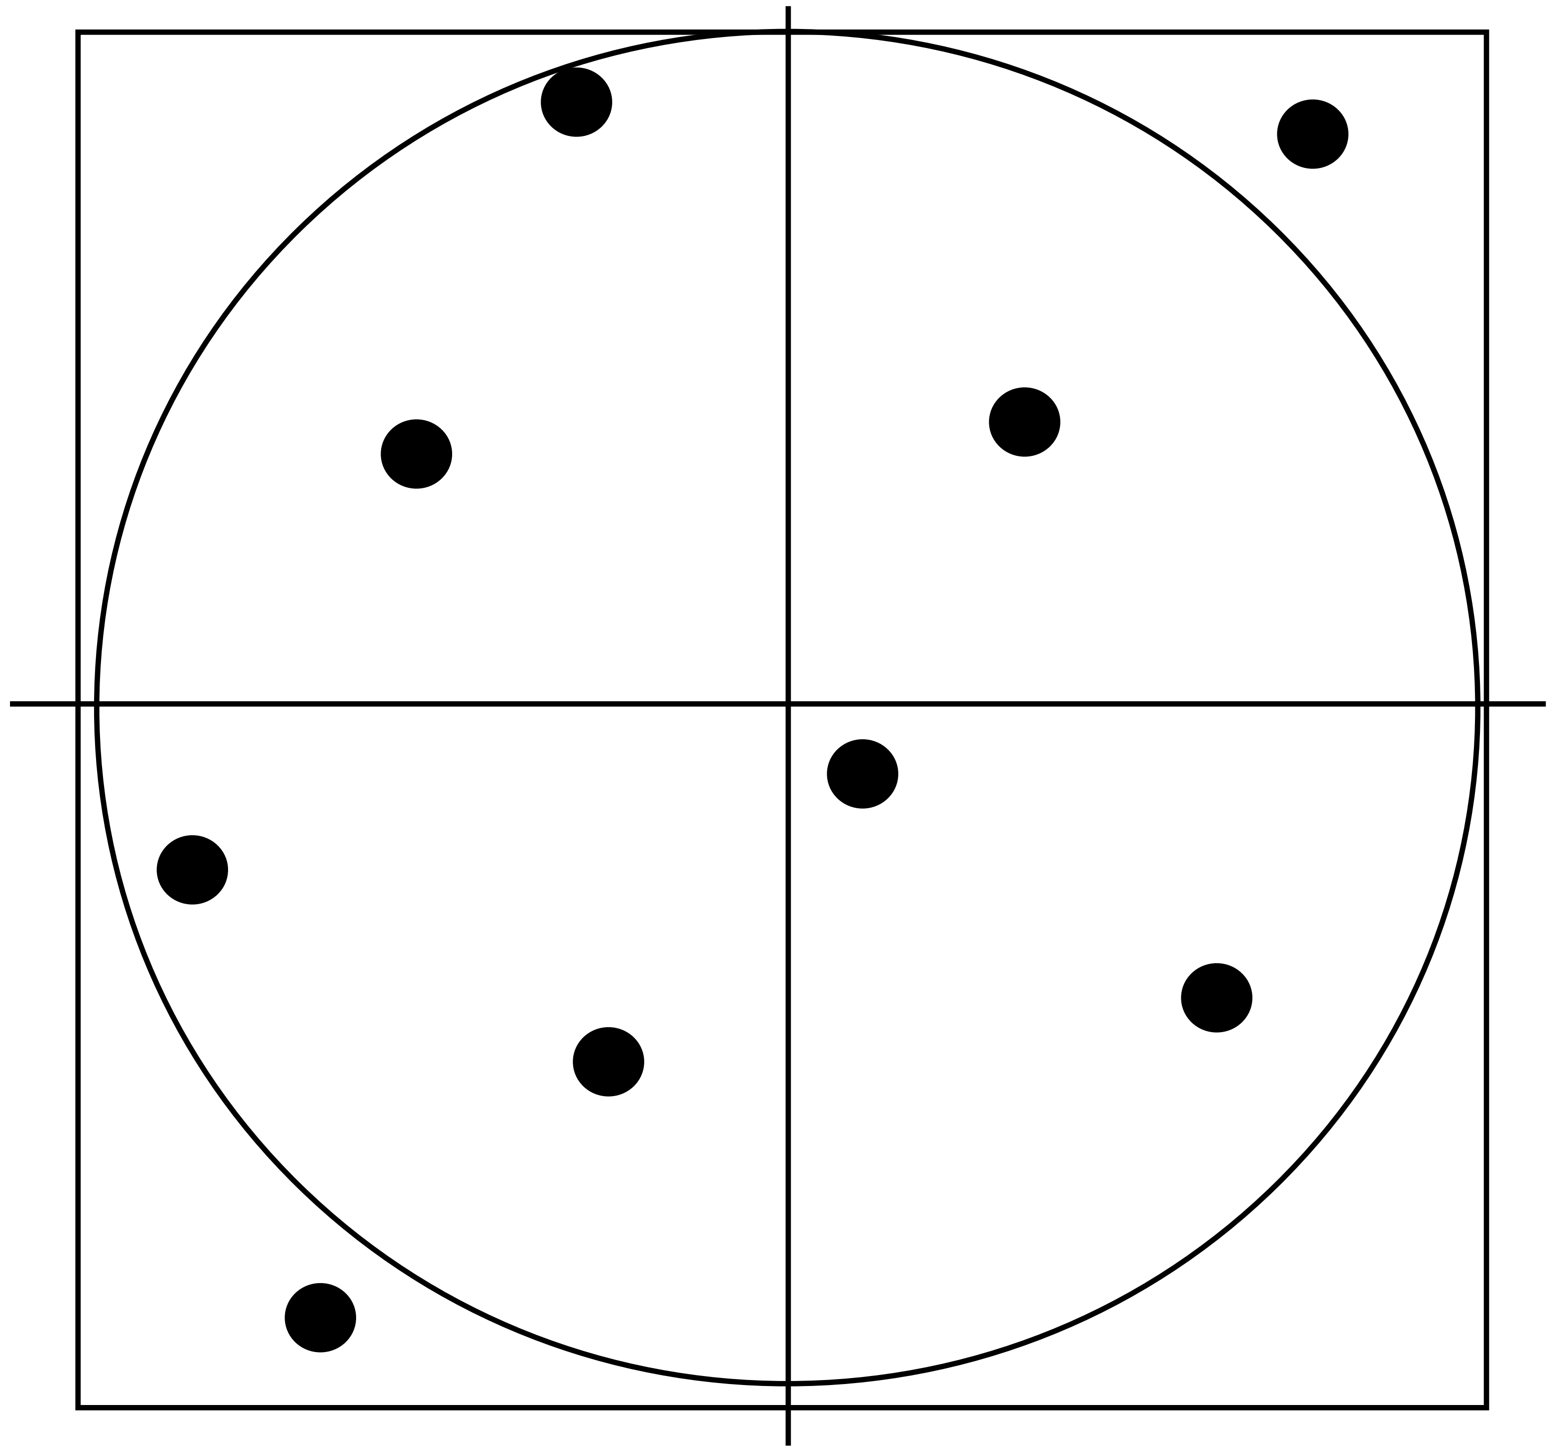

Generate random point within unit circle

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

Efficiency of technique: area of circle / area of square

Rejection sampling to generate 2D directions

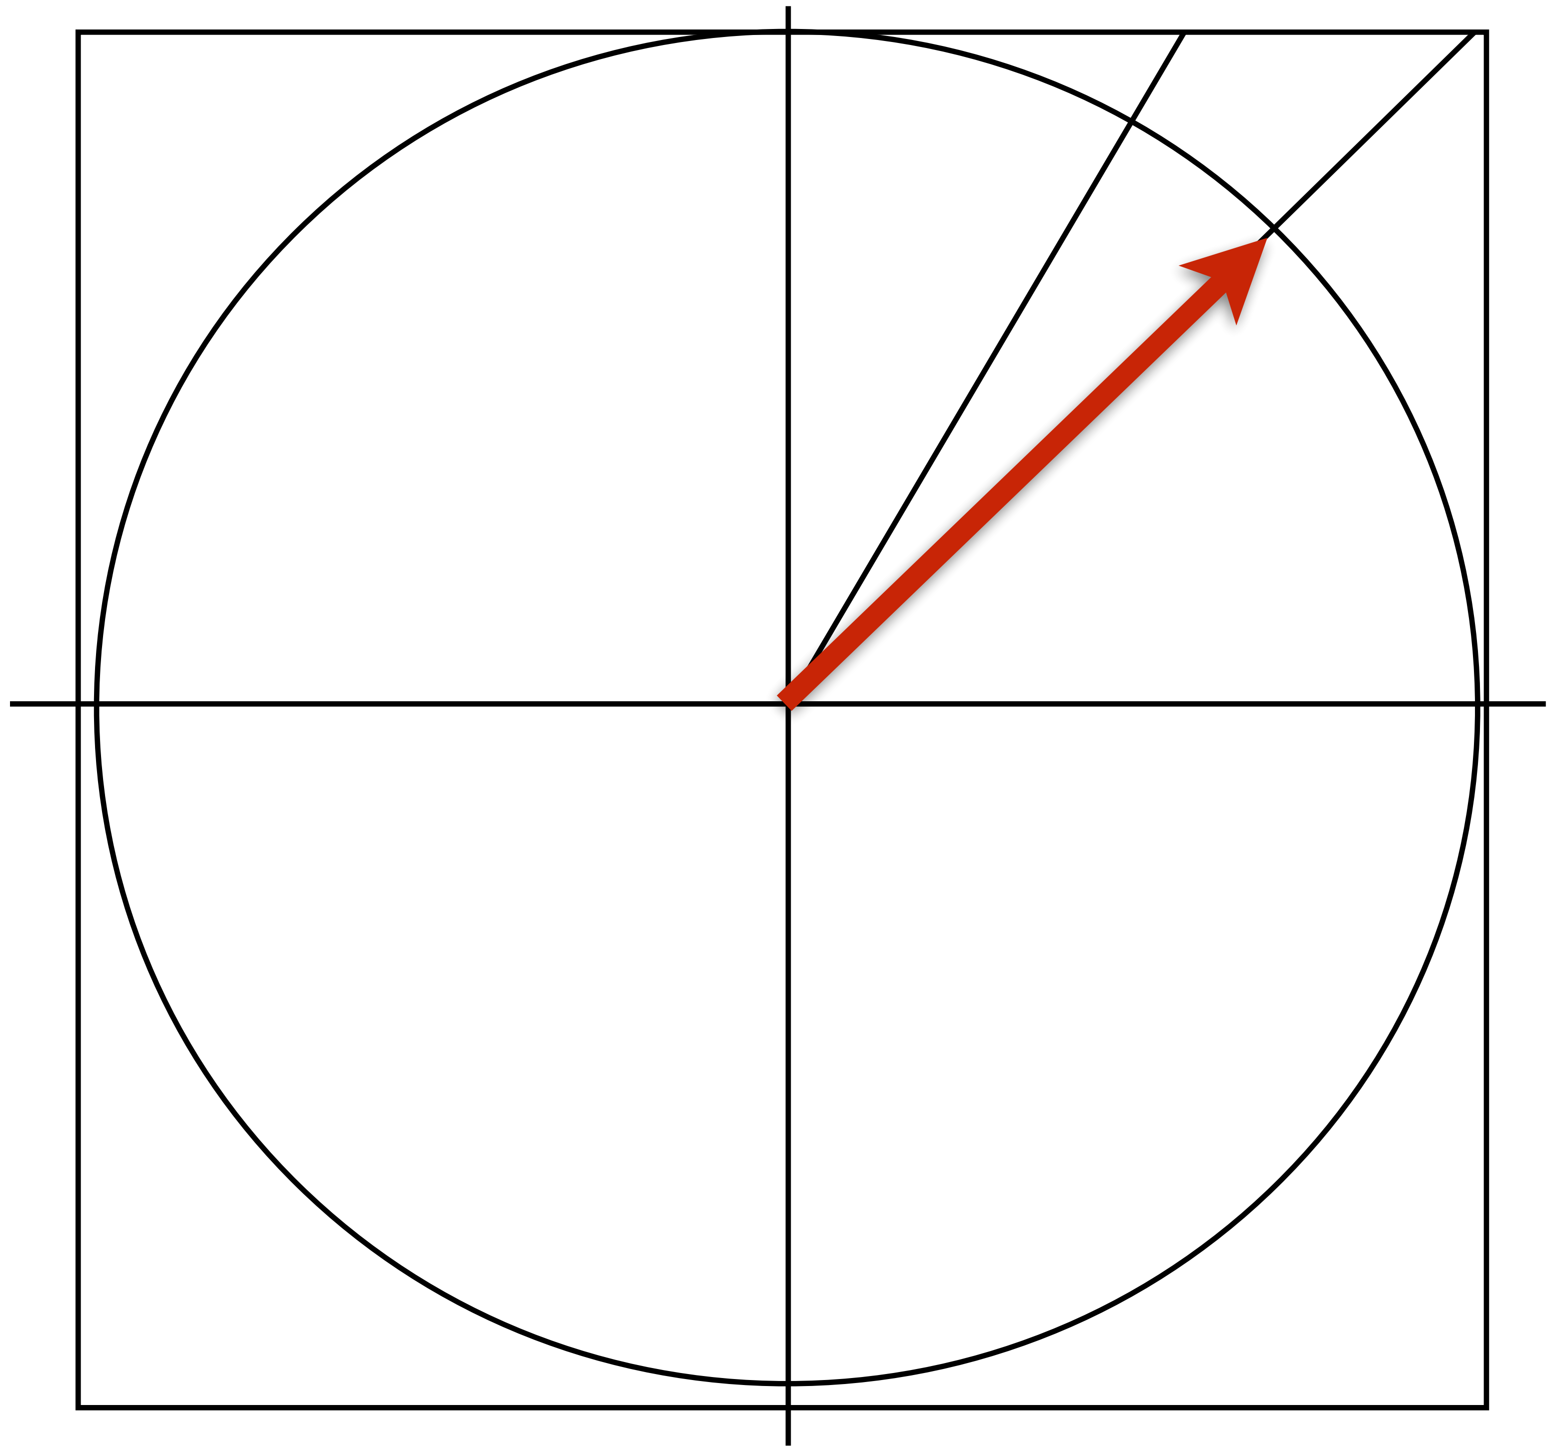

Goal: generate random directions in 2D with uniform probability

```
x = uniform(-1,1);
y = uniform(-1,1);

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

This algorithm is not correct! What is wrong?
What's a better algorithm?

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***

$$E[Y_i] = E[f(X_i)] = \int_a^b f(x)\, p(x)\, \mathrm{d}x$$

- **Estimator**

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

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

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

Basic unbiased Monte Carlo estimator

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

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

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

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

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

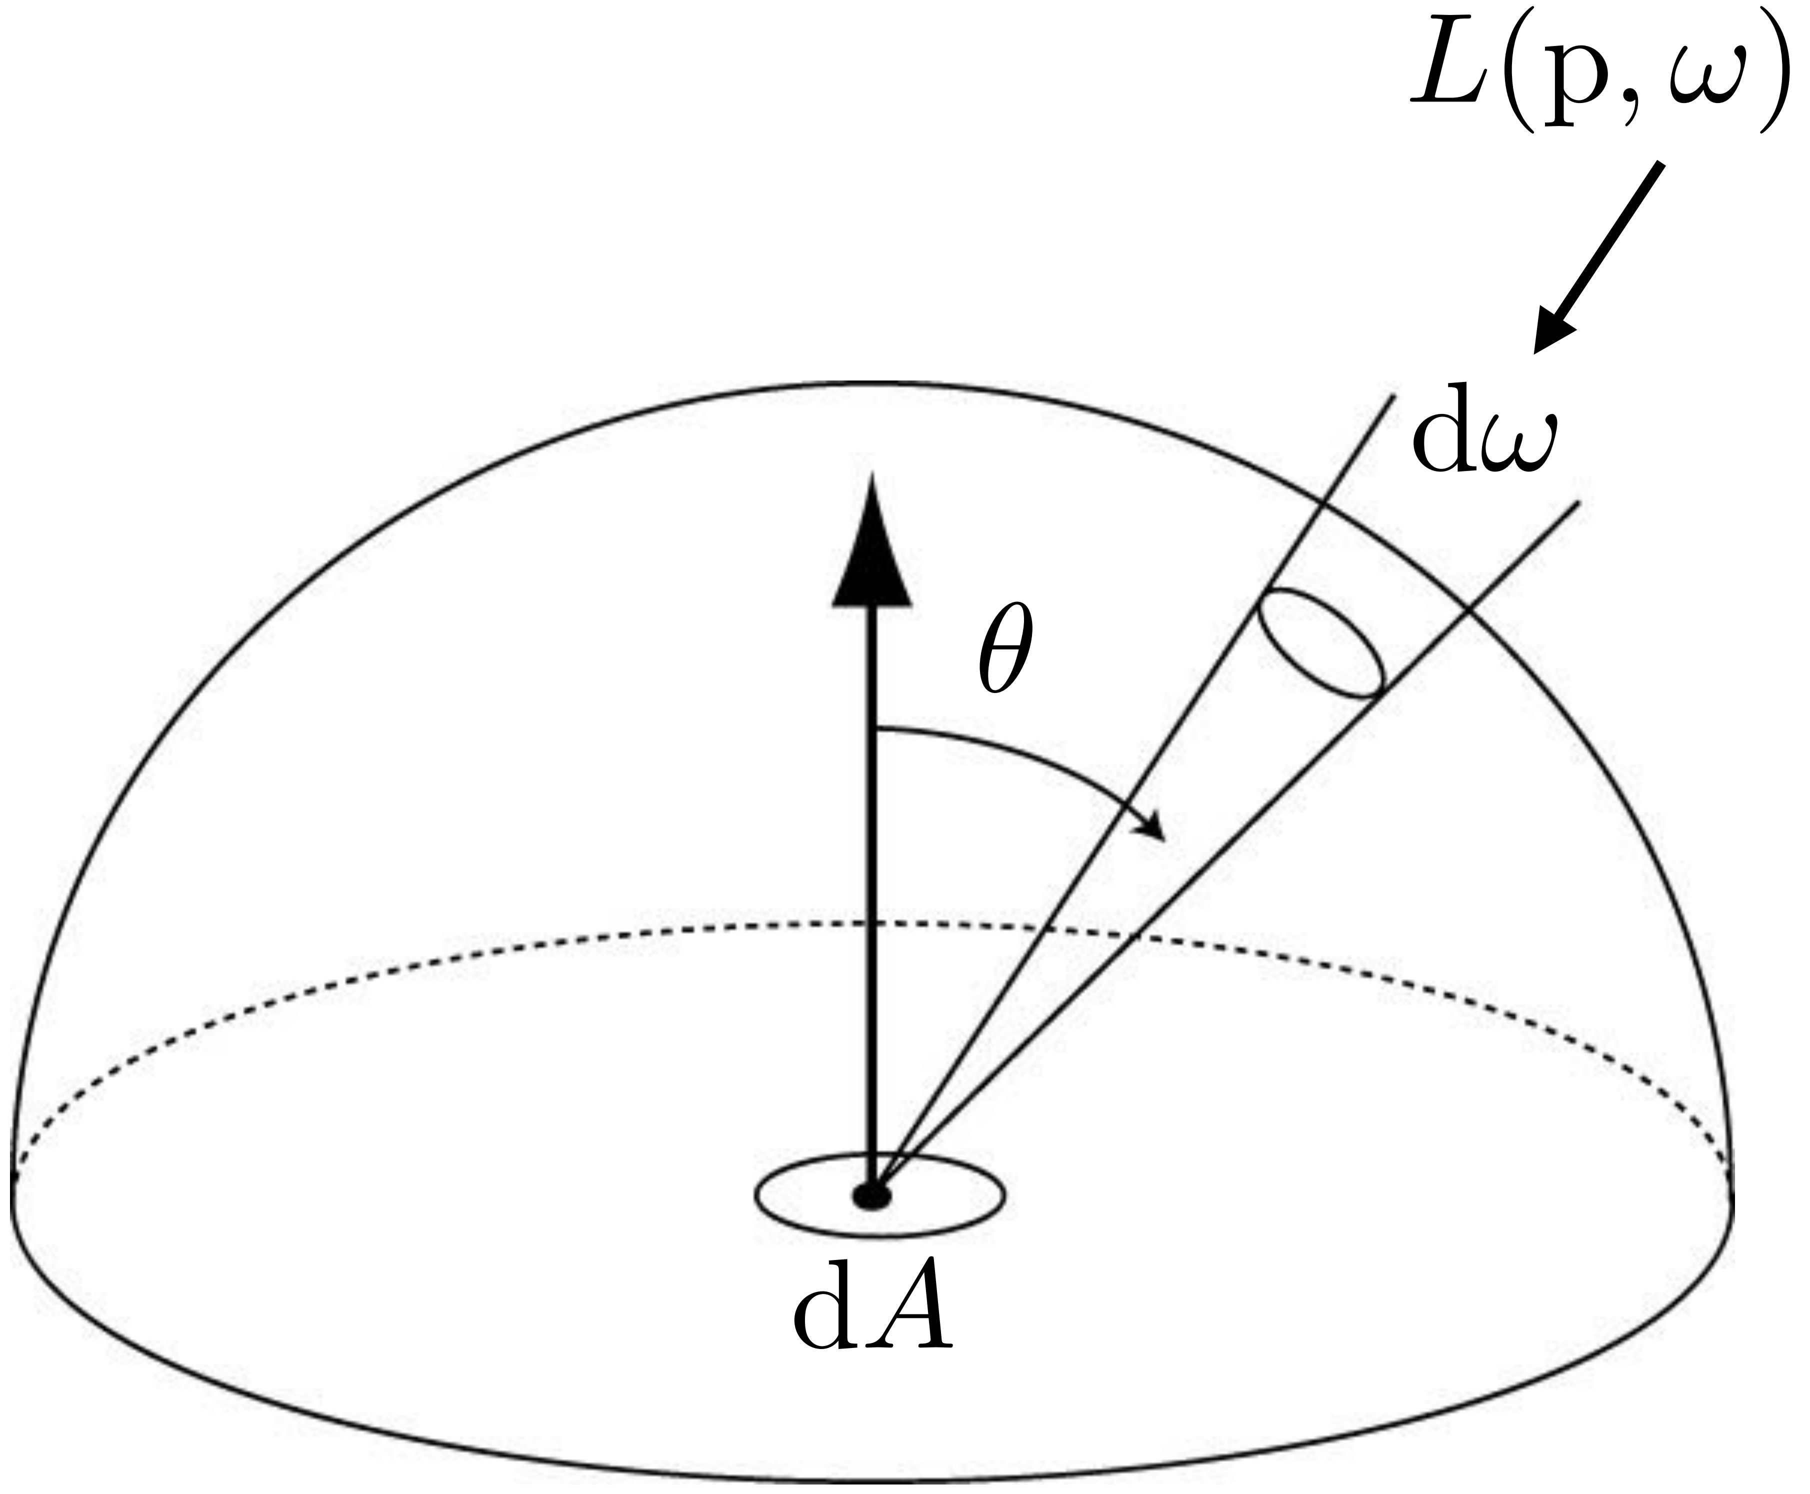

Estimator:

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

$$Y_i = f(X_i)$$

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

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

Direct lighting estimate

Uniformly-sample hemisphere of directions with respect to solid angle

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

Given surface point *p*

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**

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

Direct lighting: hemisphere sampling

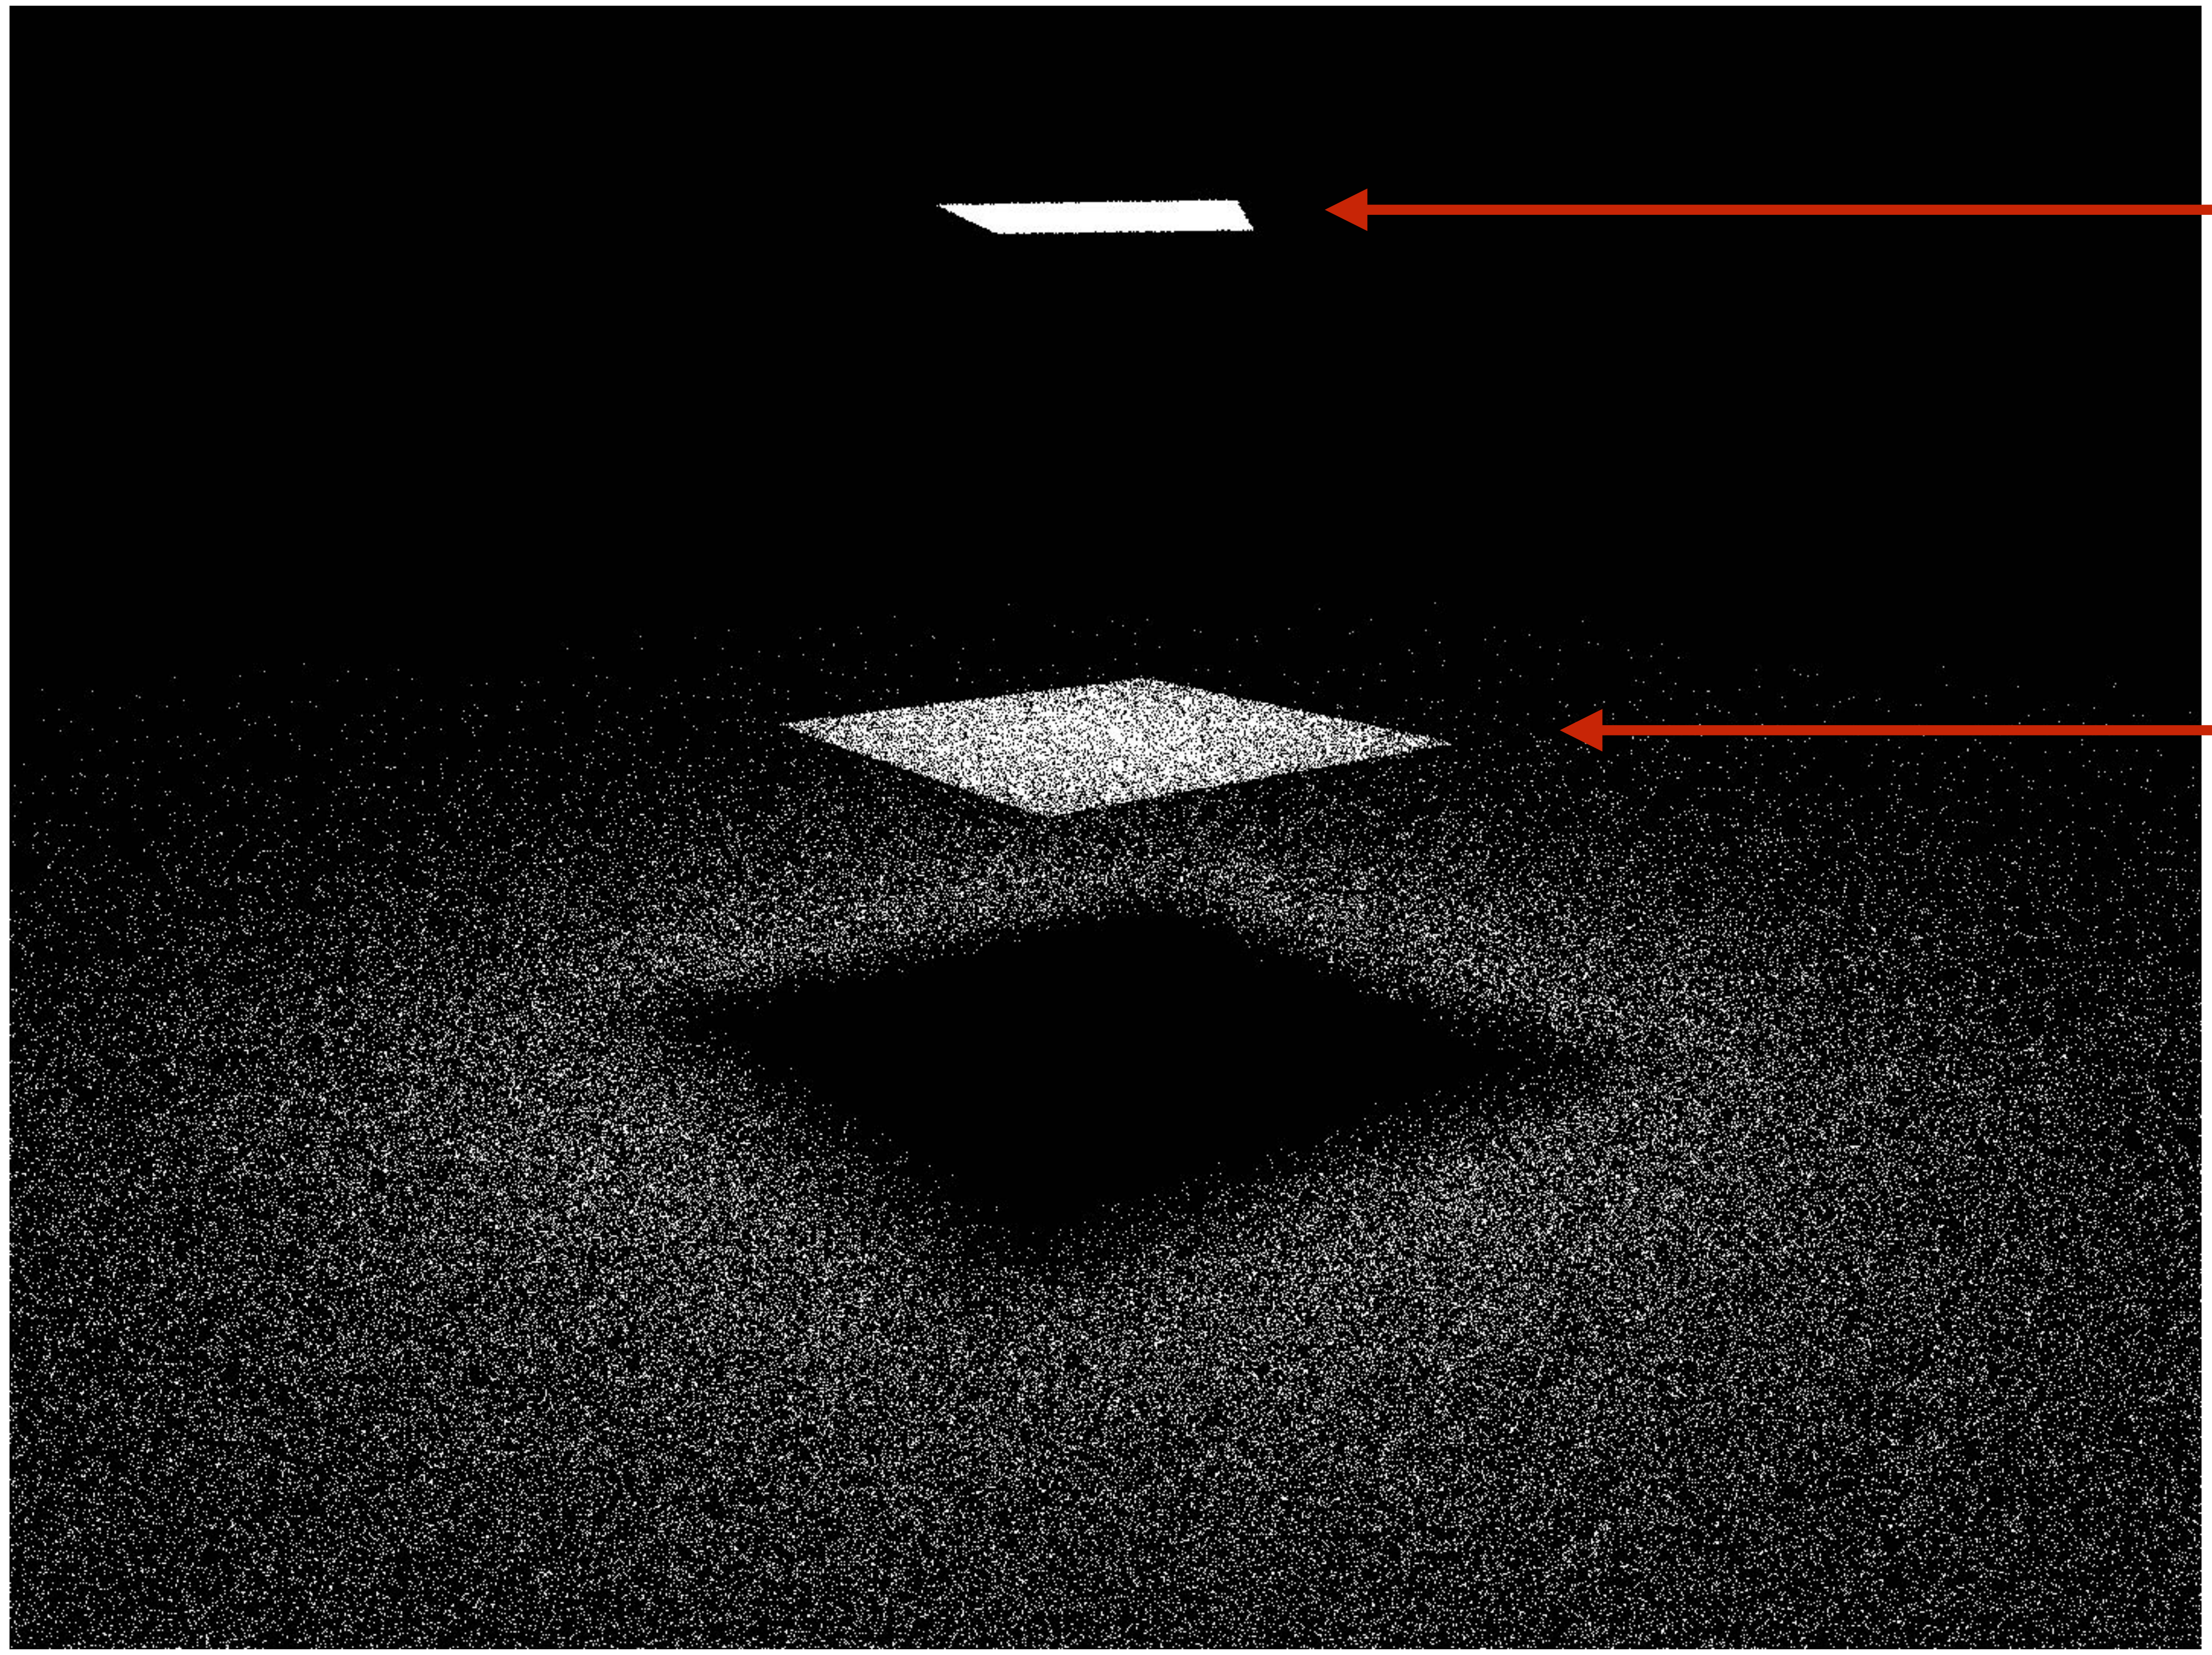

Hemisphere

16 light samples (=16 shadow rays)

Uniform hemisphere sampling

Generate random direction on hemisphere (all directions equally likely)

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

Direction computed from uniformly distributed point on 2D plane:

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

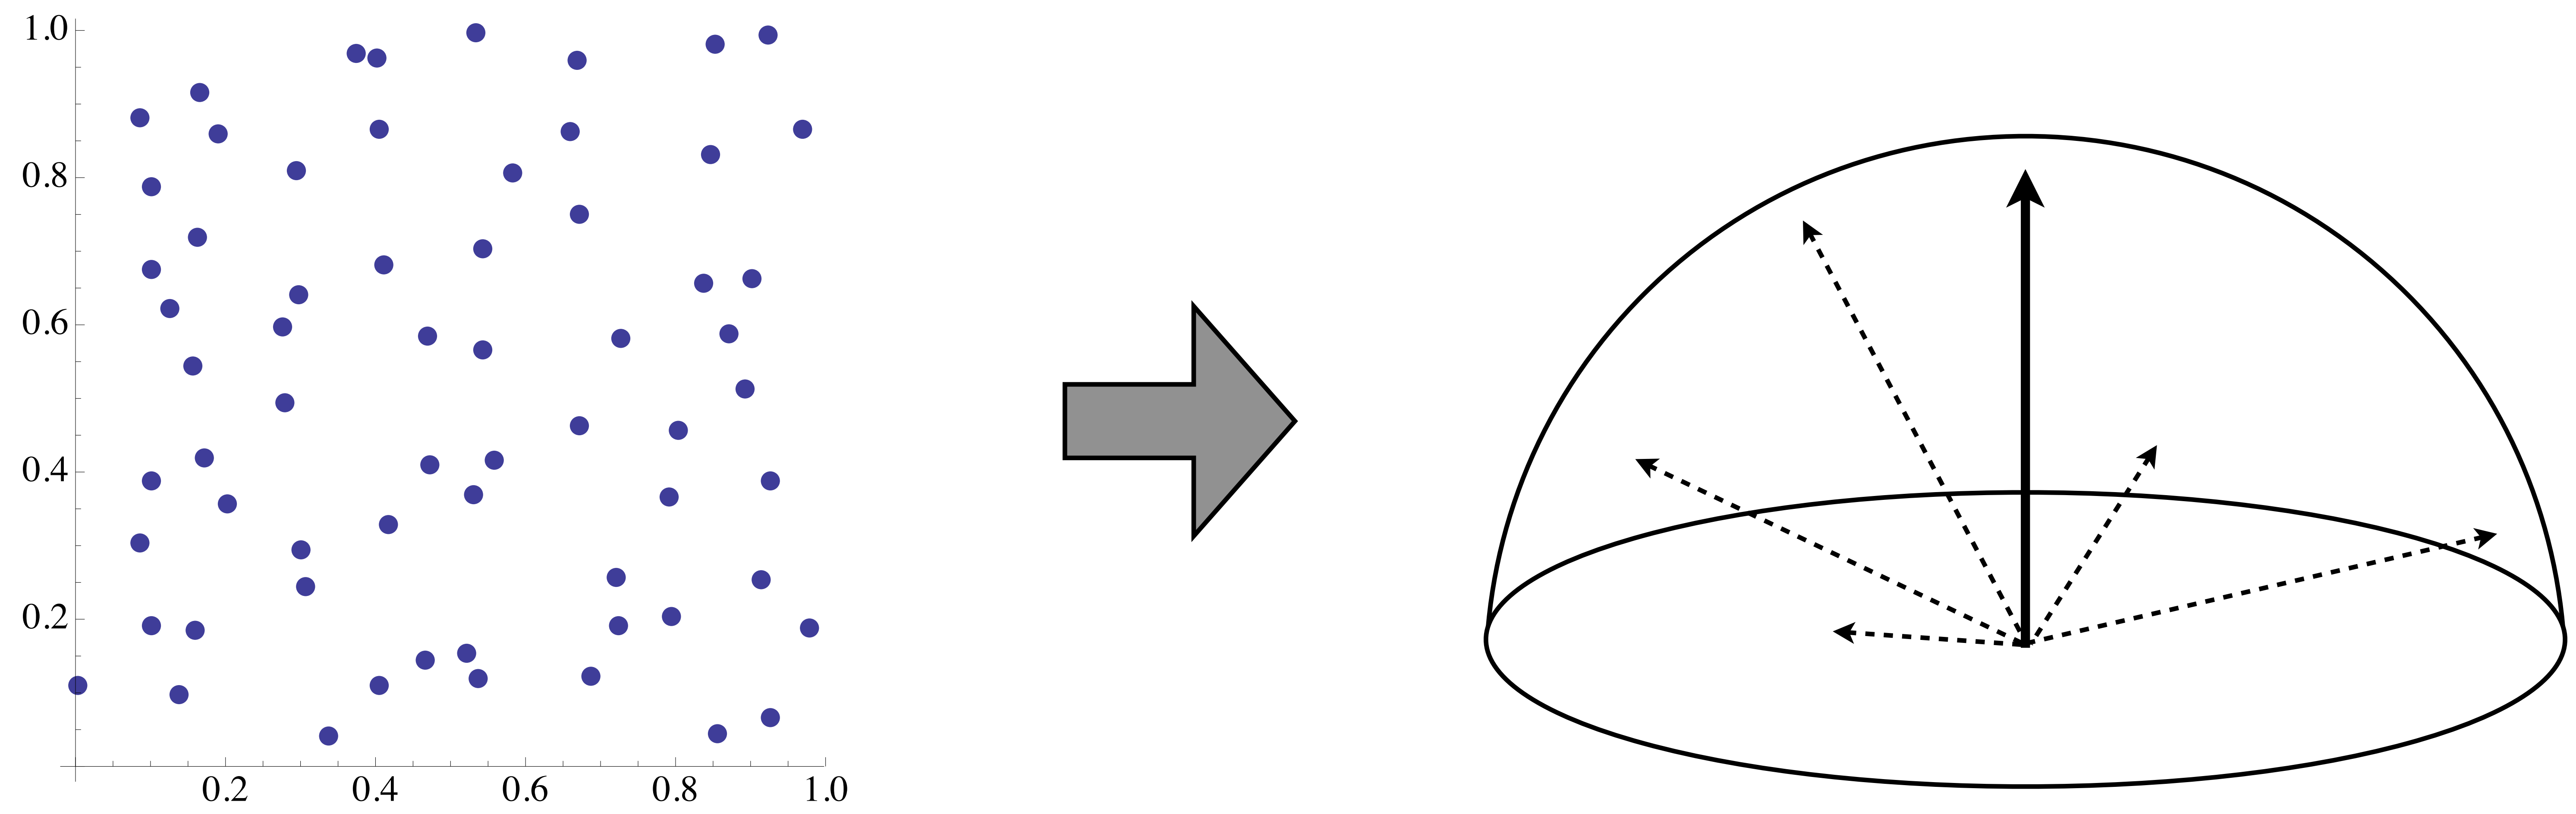

Exercise to students: derive from the inversion method

Direct lighting: area integral formulation

Consider uniformly sampling surface of light, instead of hemisphere of directions…

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

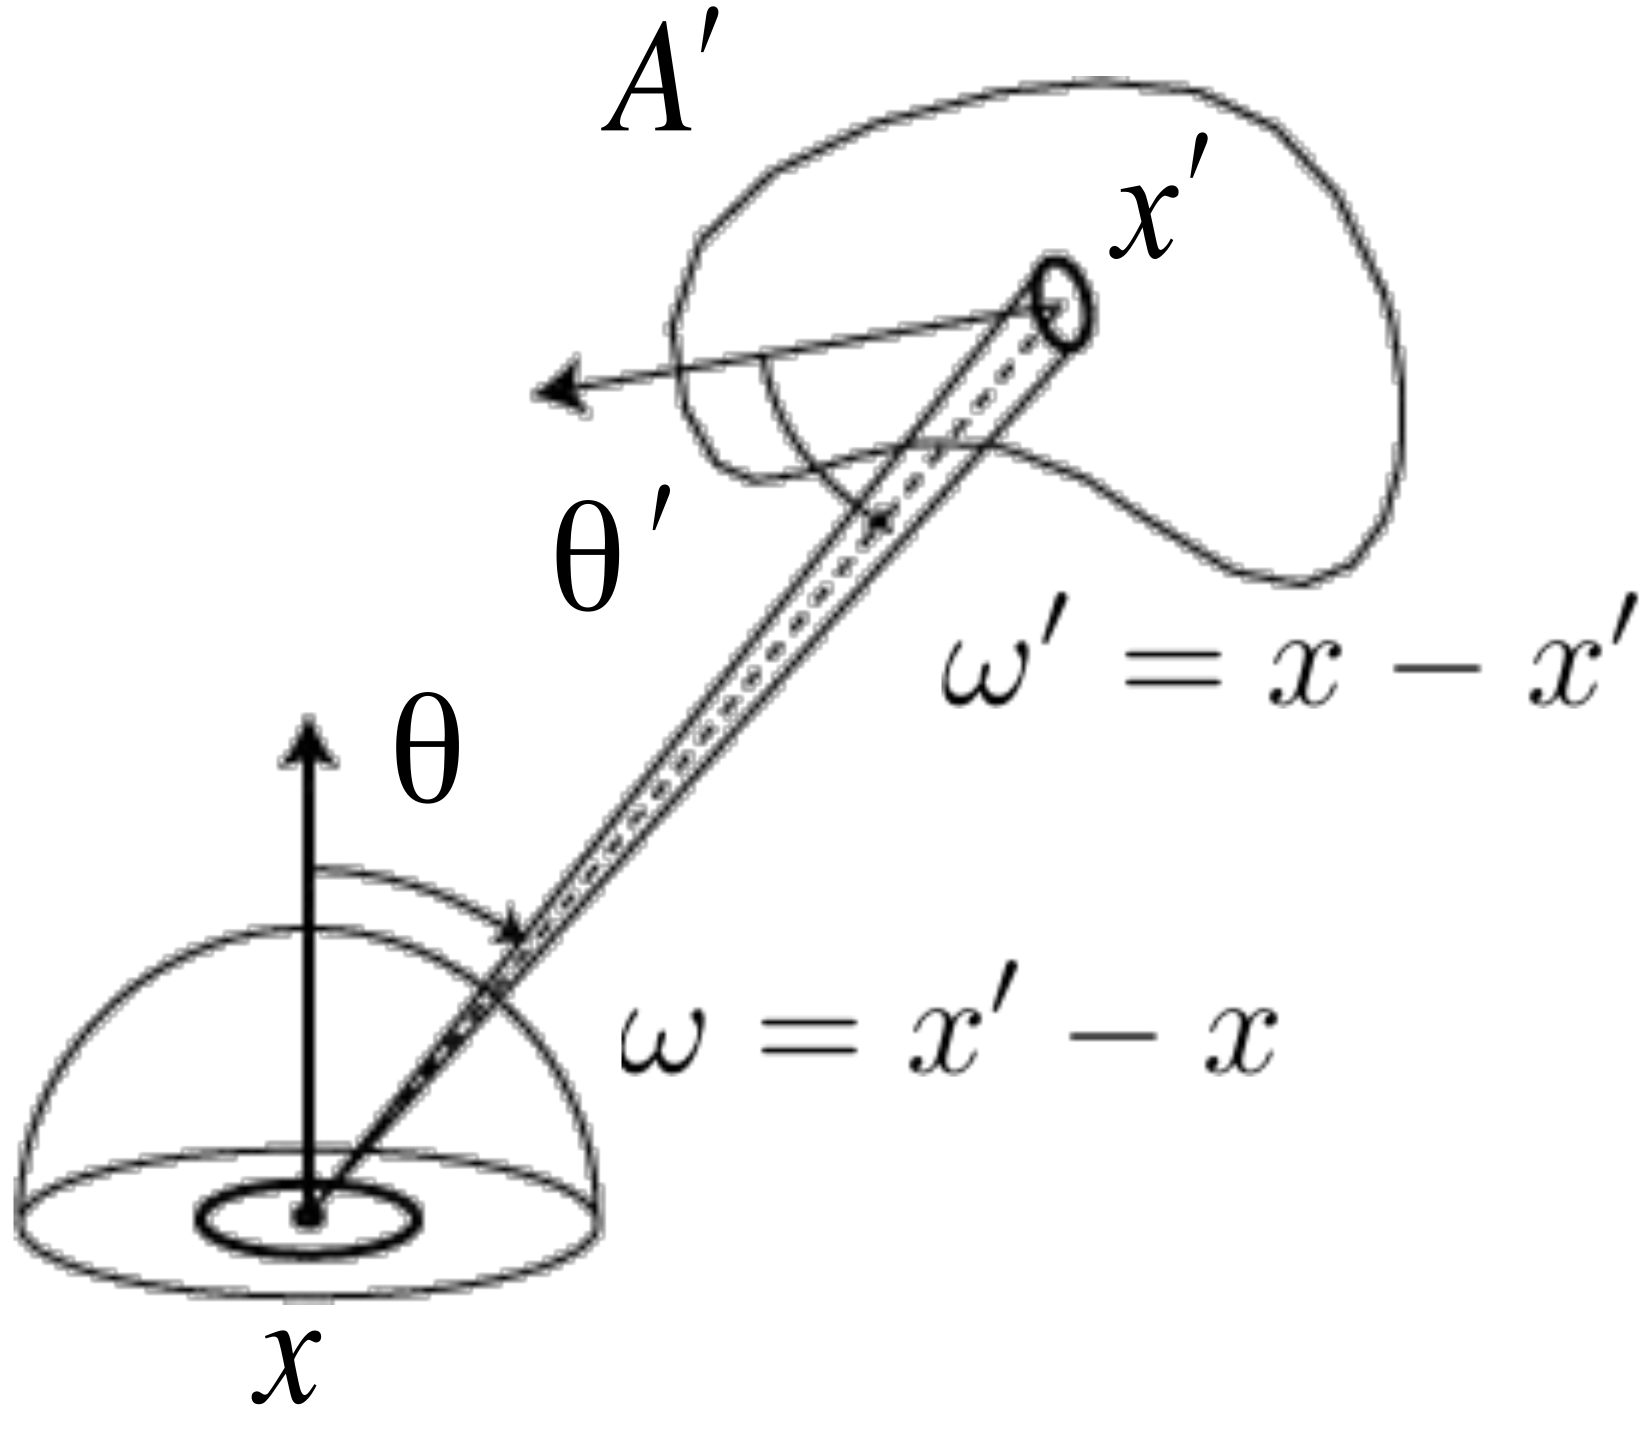

Integral

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

Visibility

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

Radiance

$$L_i(x,\omega) = L_o(x',\omega')$$

Direct lighting: area sampling

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

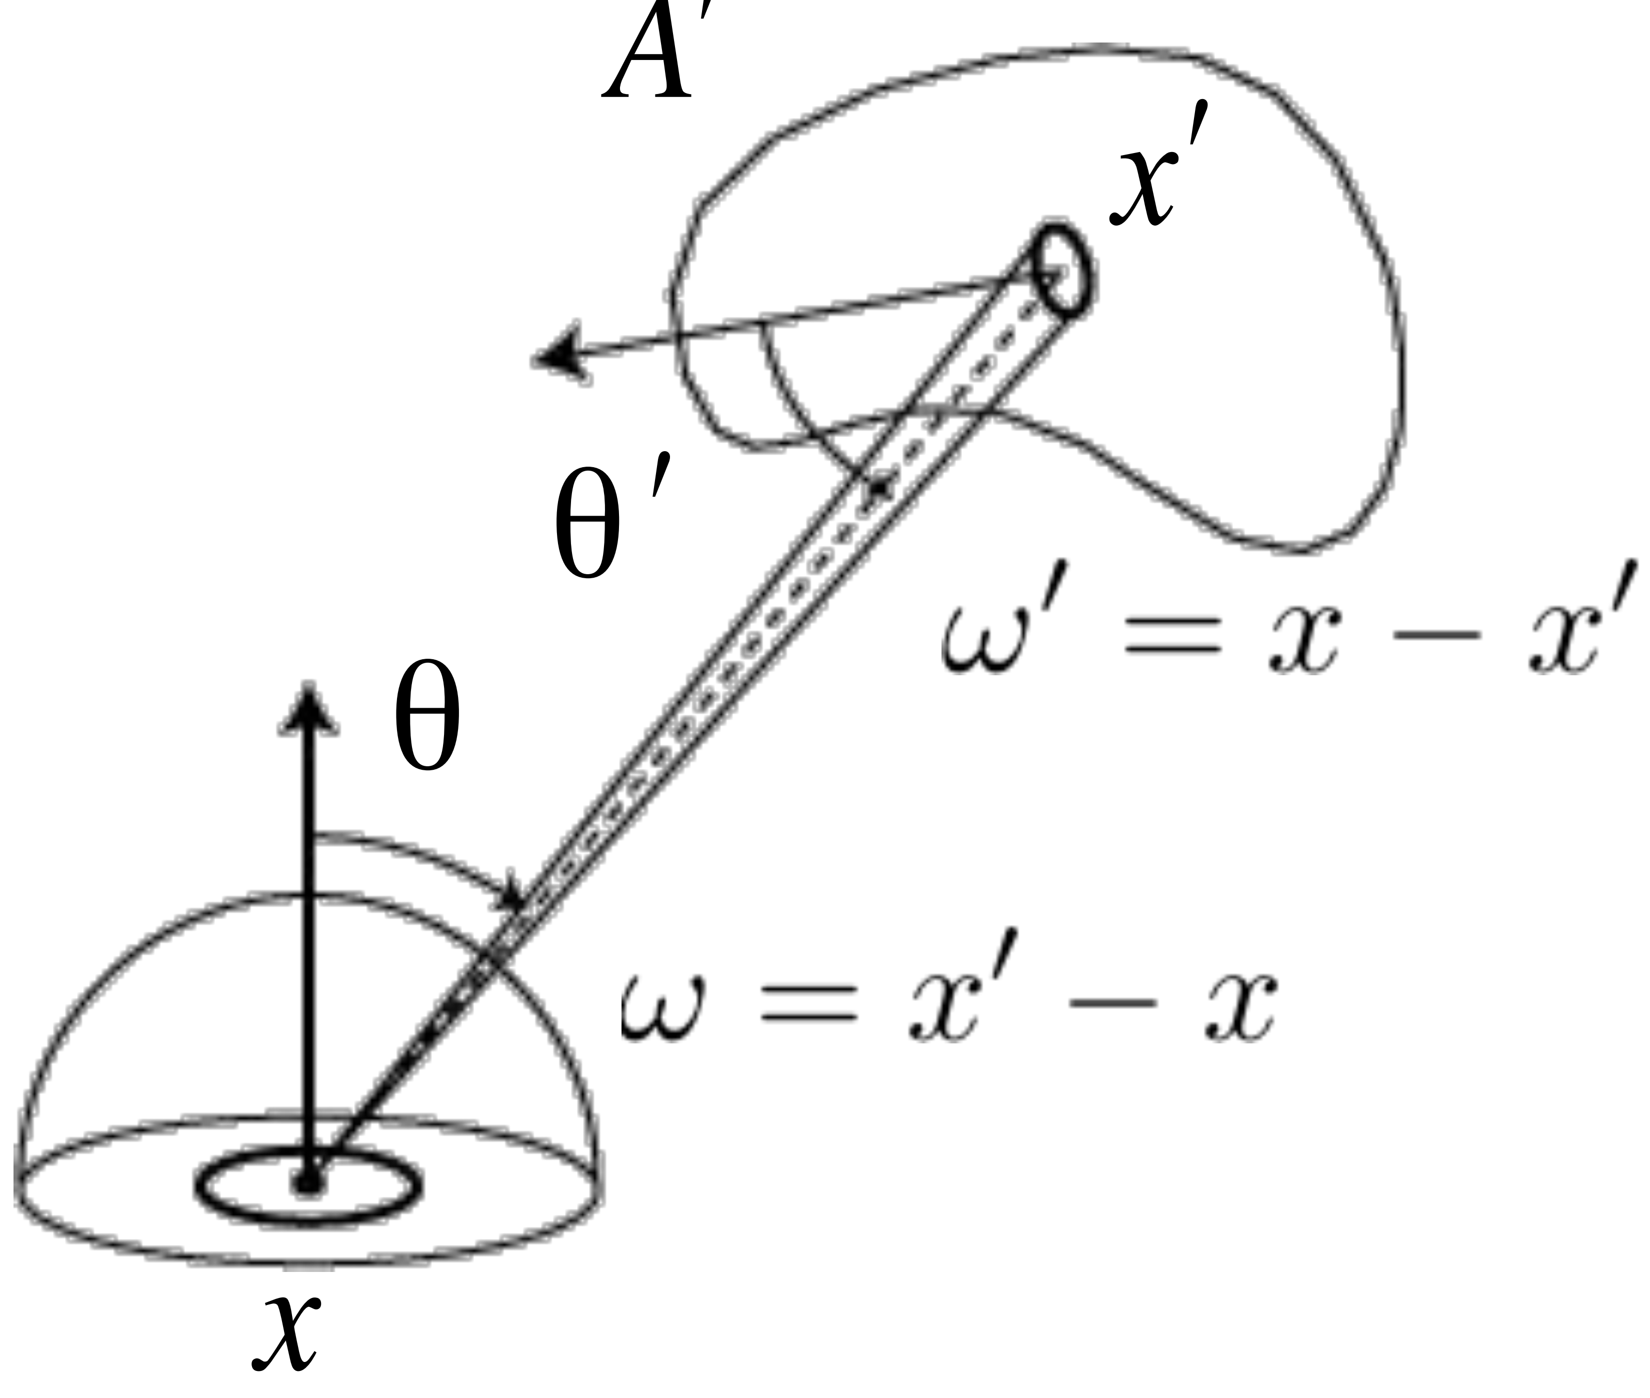

$$\int_{A'} p(x')\, dA' = 1$$

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

Sample shape uniformly by area

(Picking random points on the light)

Direct lighting: area sampling

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

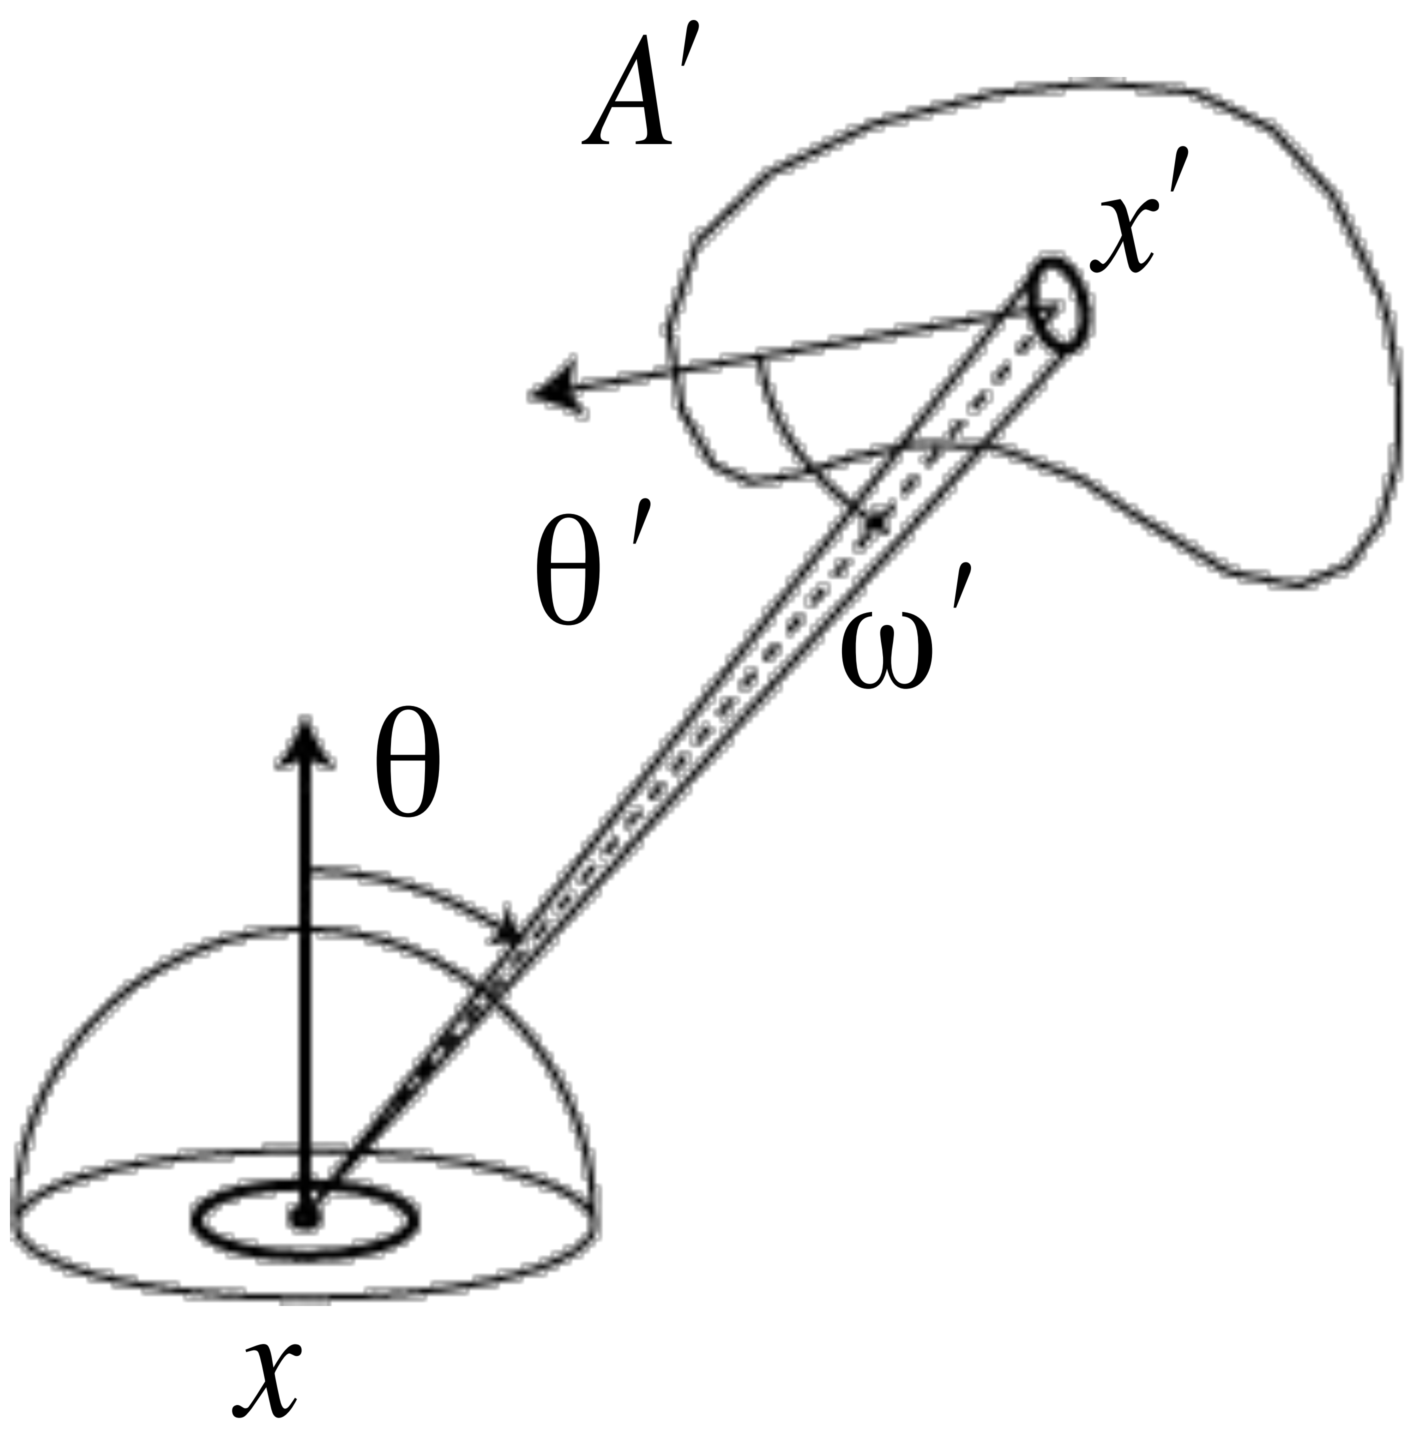

MC Estimator

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

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

Direct lighting estimate (area sampling light with area A')

Given surface point *x*

For each of N samples:

- **Generate random point x' on area light, compute direction from x to x':** ω_i
- **Compute incident irradiance due to ray from x' to x: as** $dE_i = L_o(x', -w_i)V(x, x')\frac{\cos\theta_i \cos\theta_o}{|x - x'|^2}$
- **Accumulate** $\frac{A'}{N} dE_i$ **into estimator**

How do you evaluate V()?

Direct lighting: area sampling

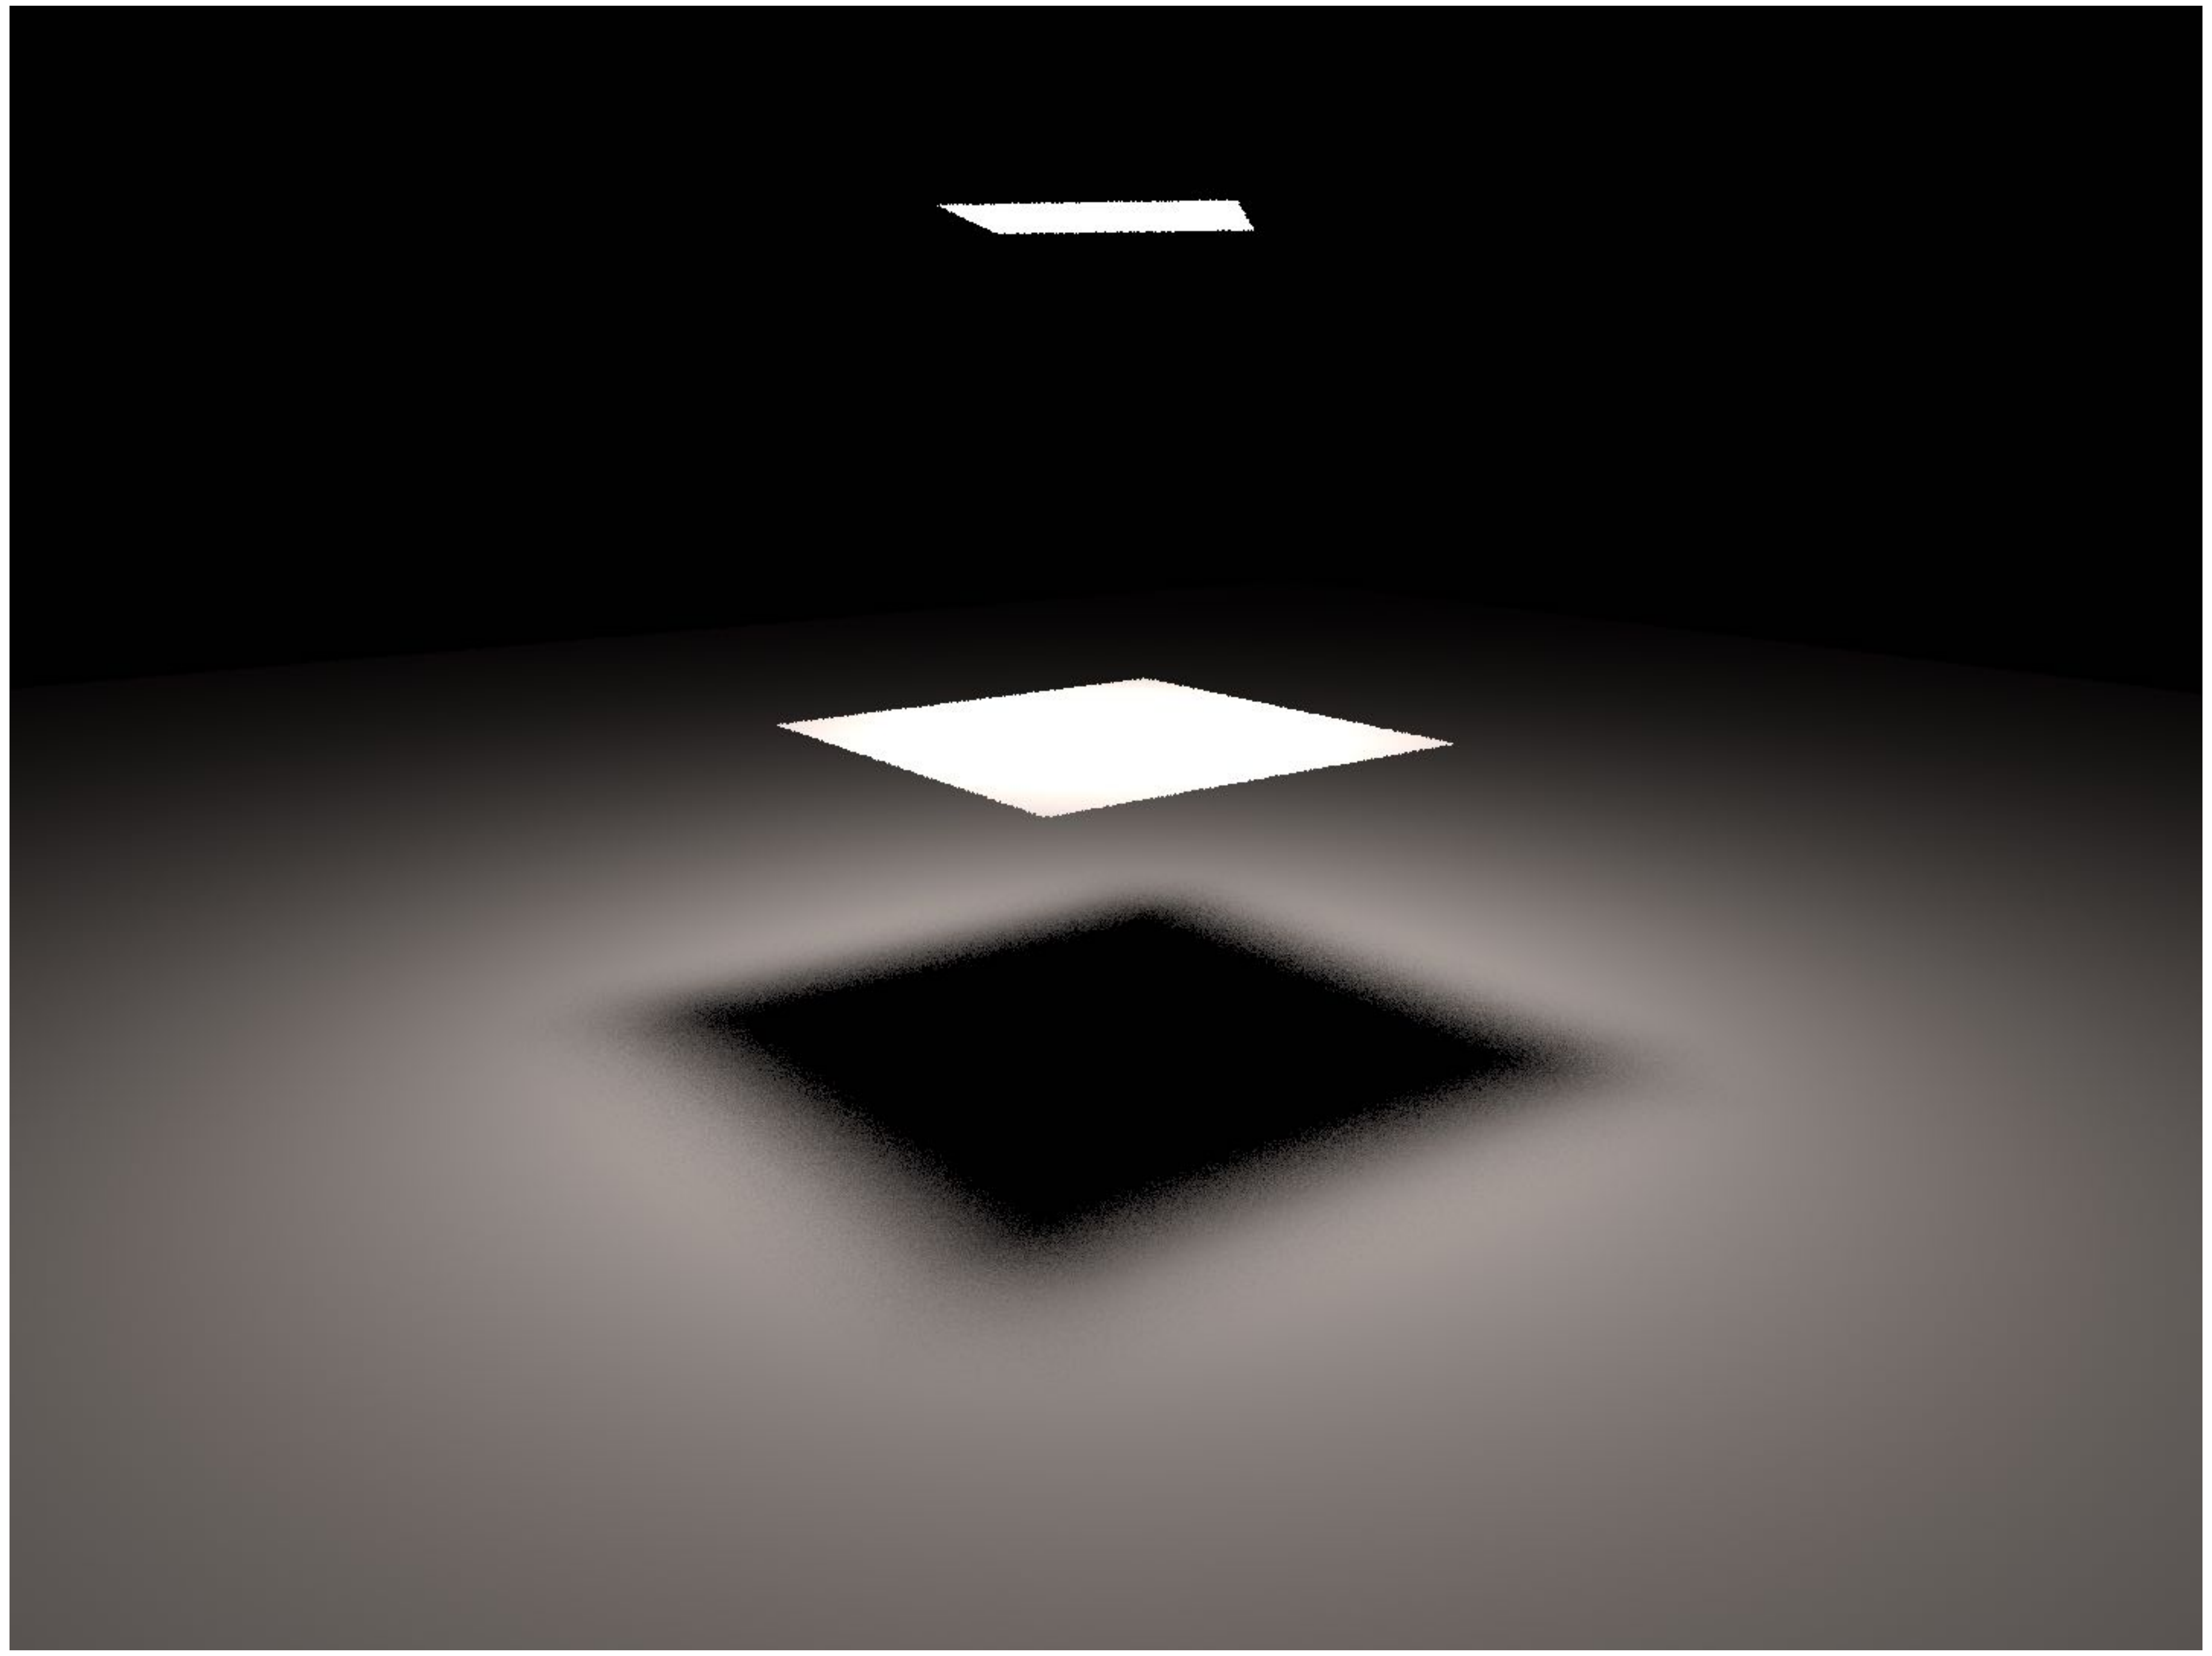

Area

16 shadow rays

Random sampling introduces noise

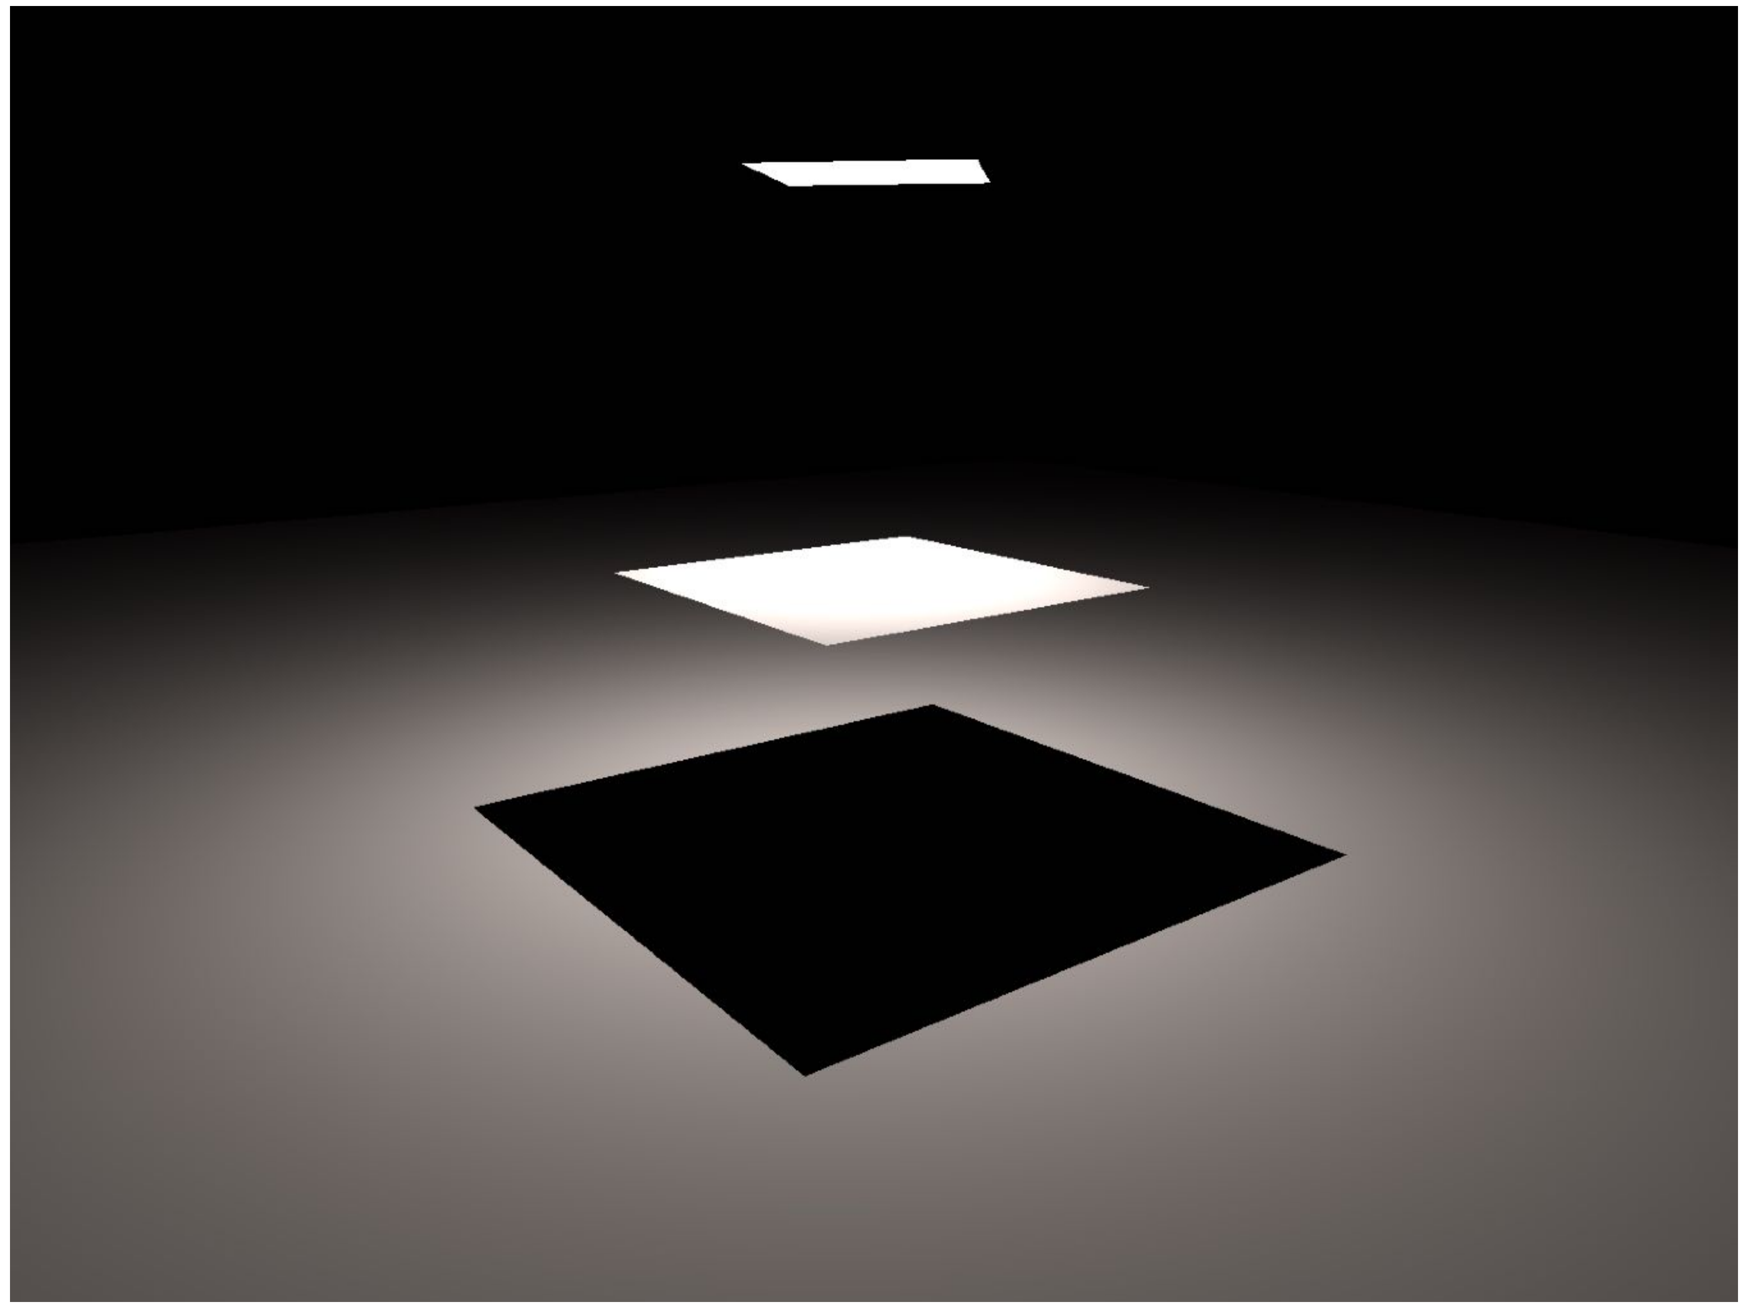

Center

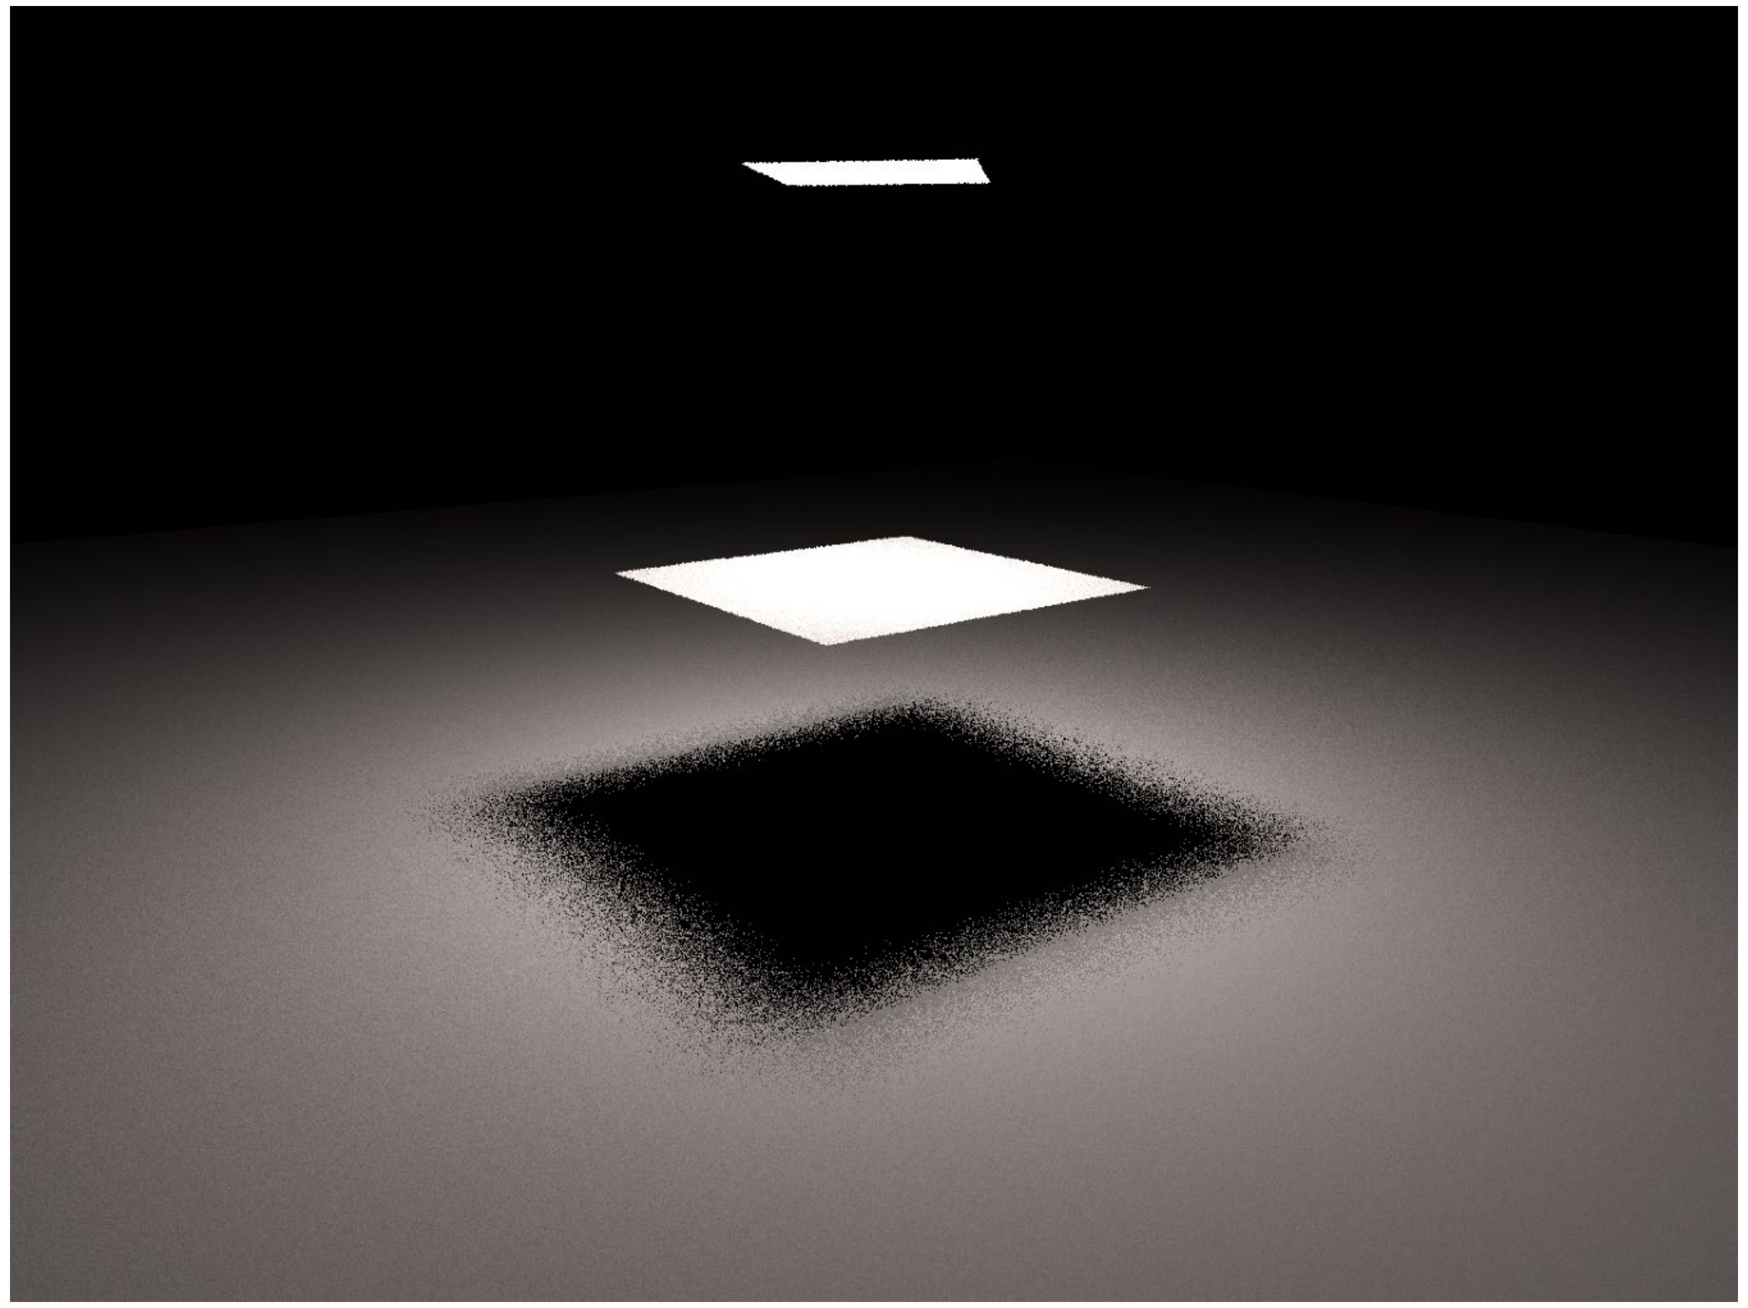

Random

1 shadow ray per eye ray

Random sampling introduces noise

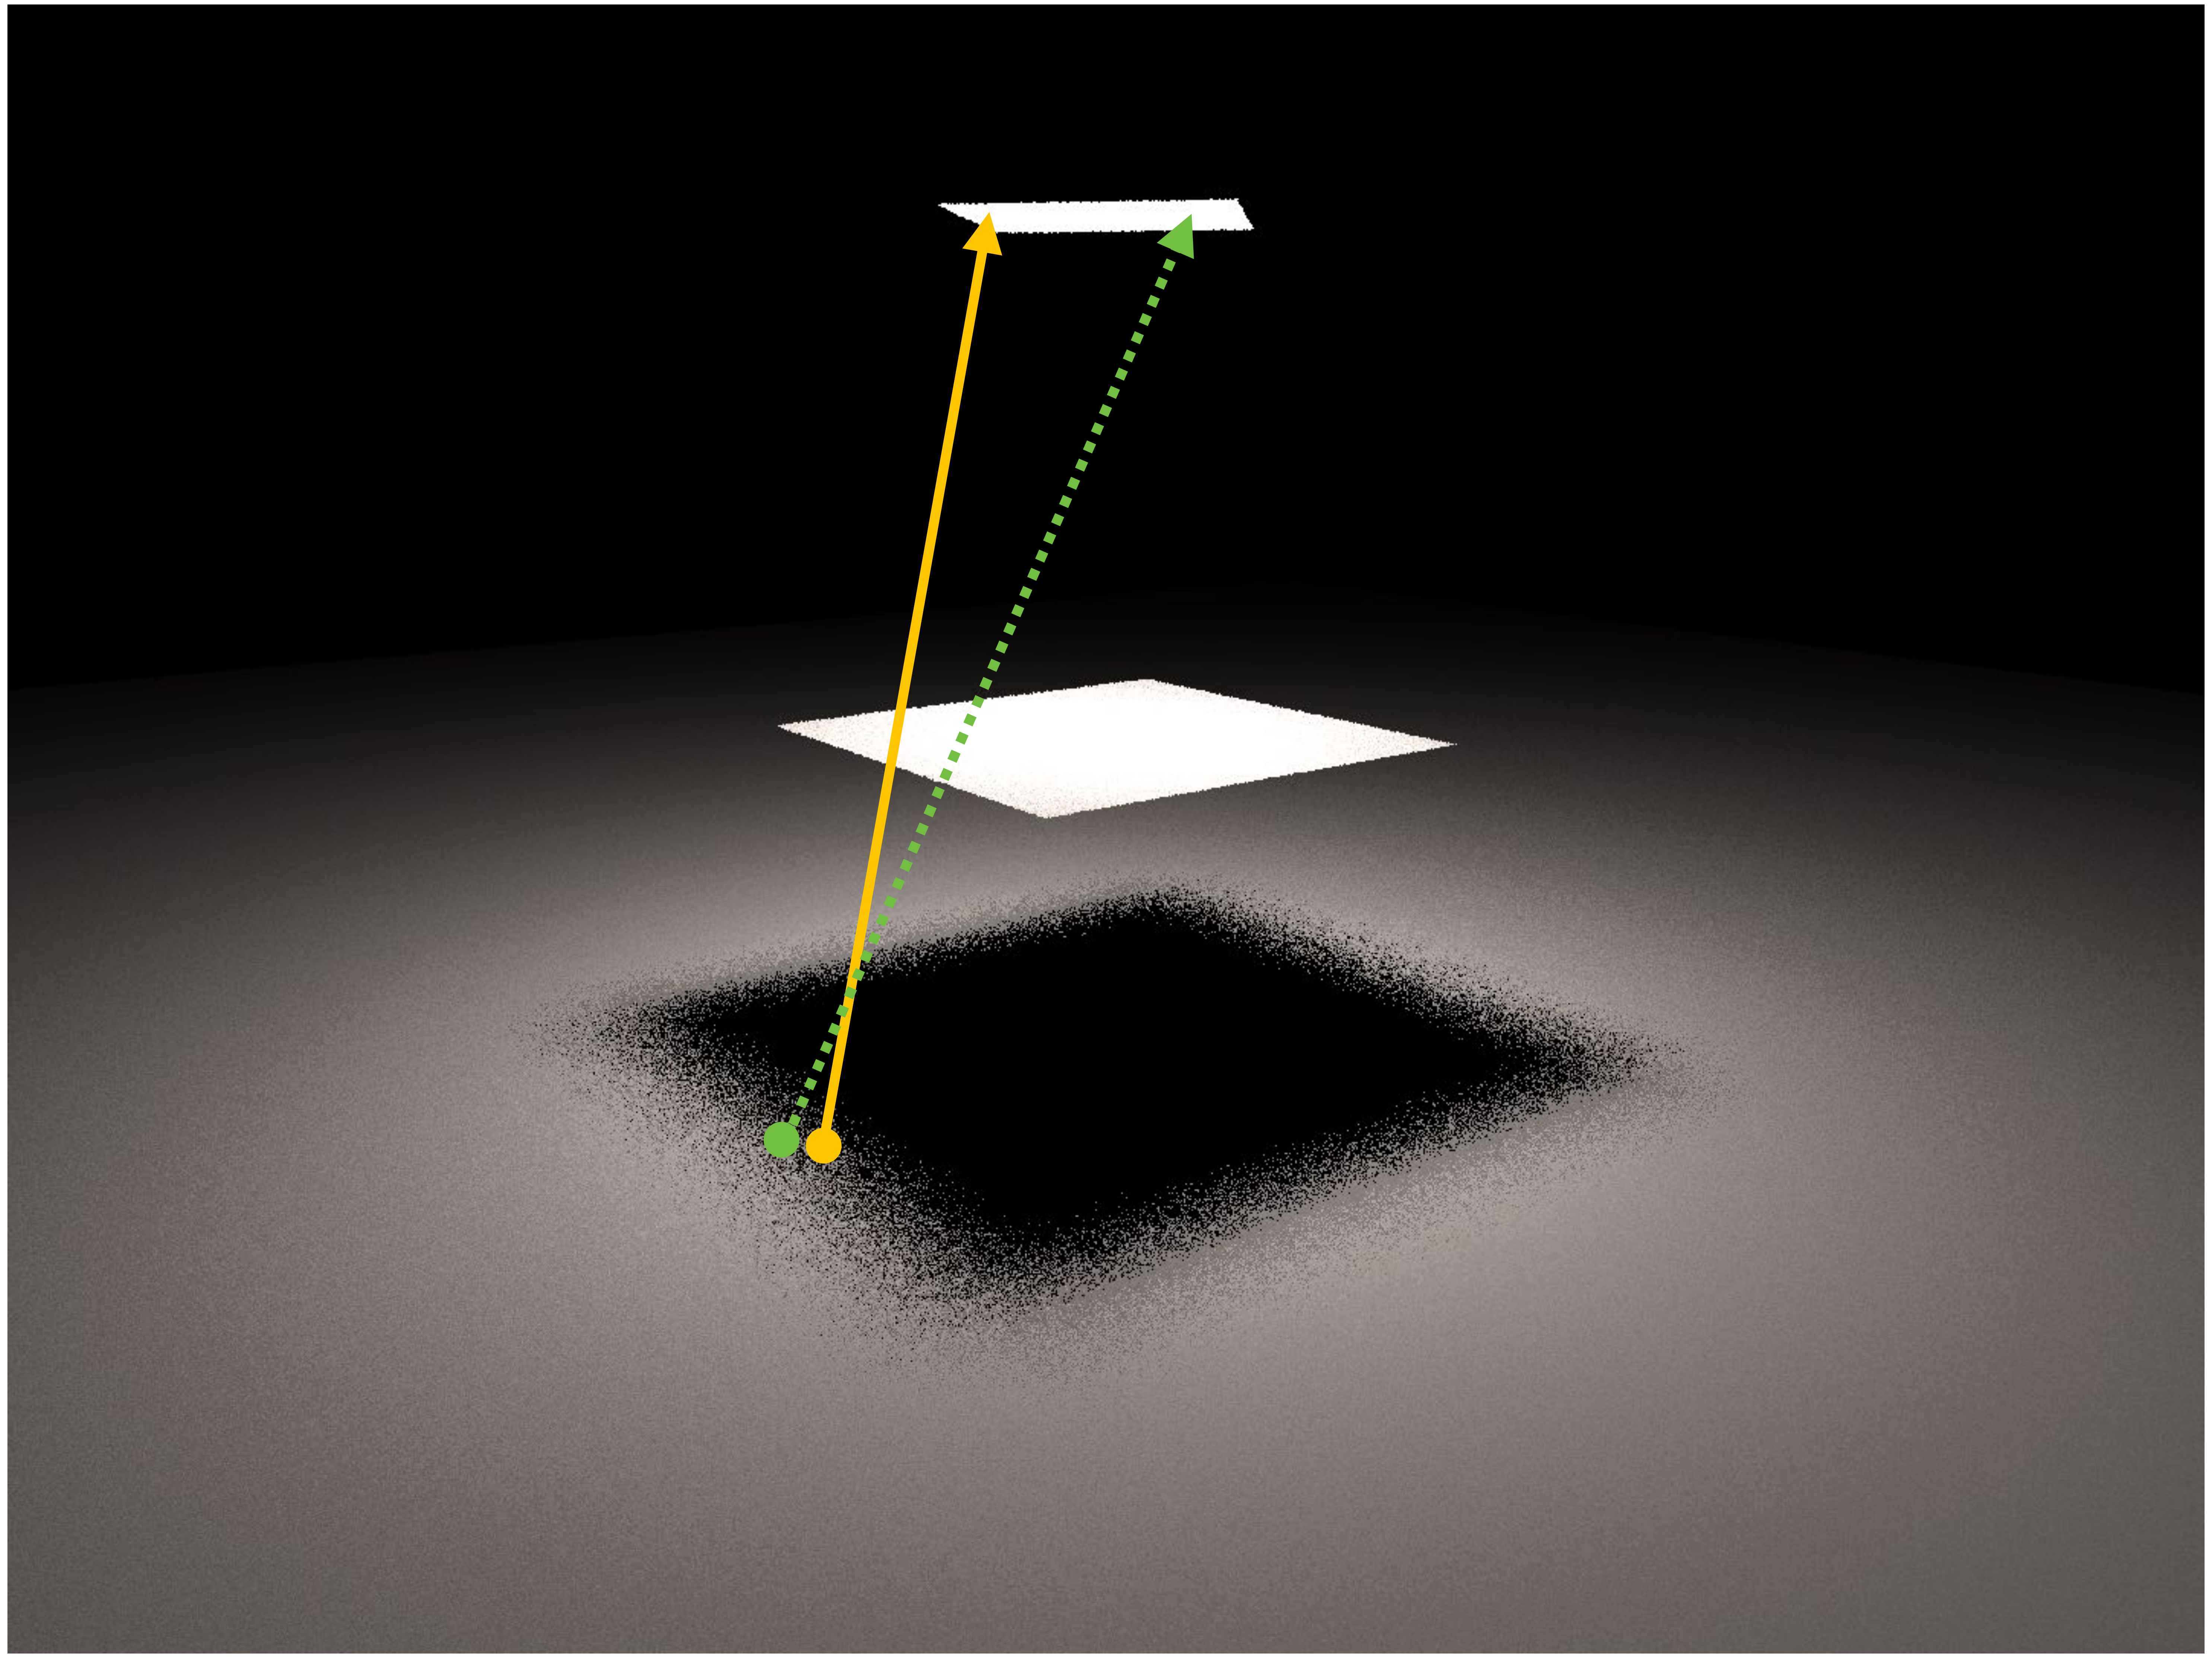

Incident lighting estimator uses different random directions when computing incident lighting for different points. Some of those directions are occluded, some are not!

(The estimator is a random variable!)

Quality improves with more rays

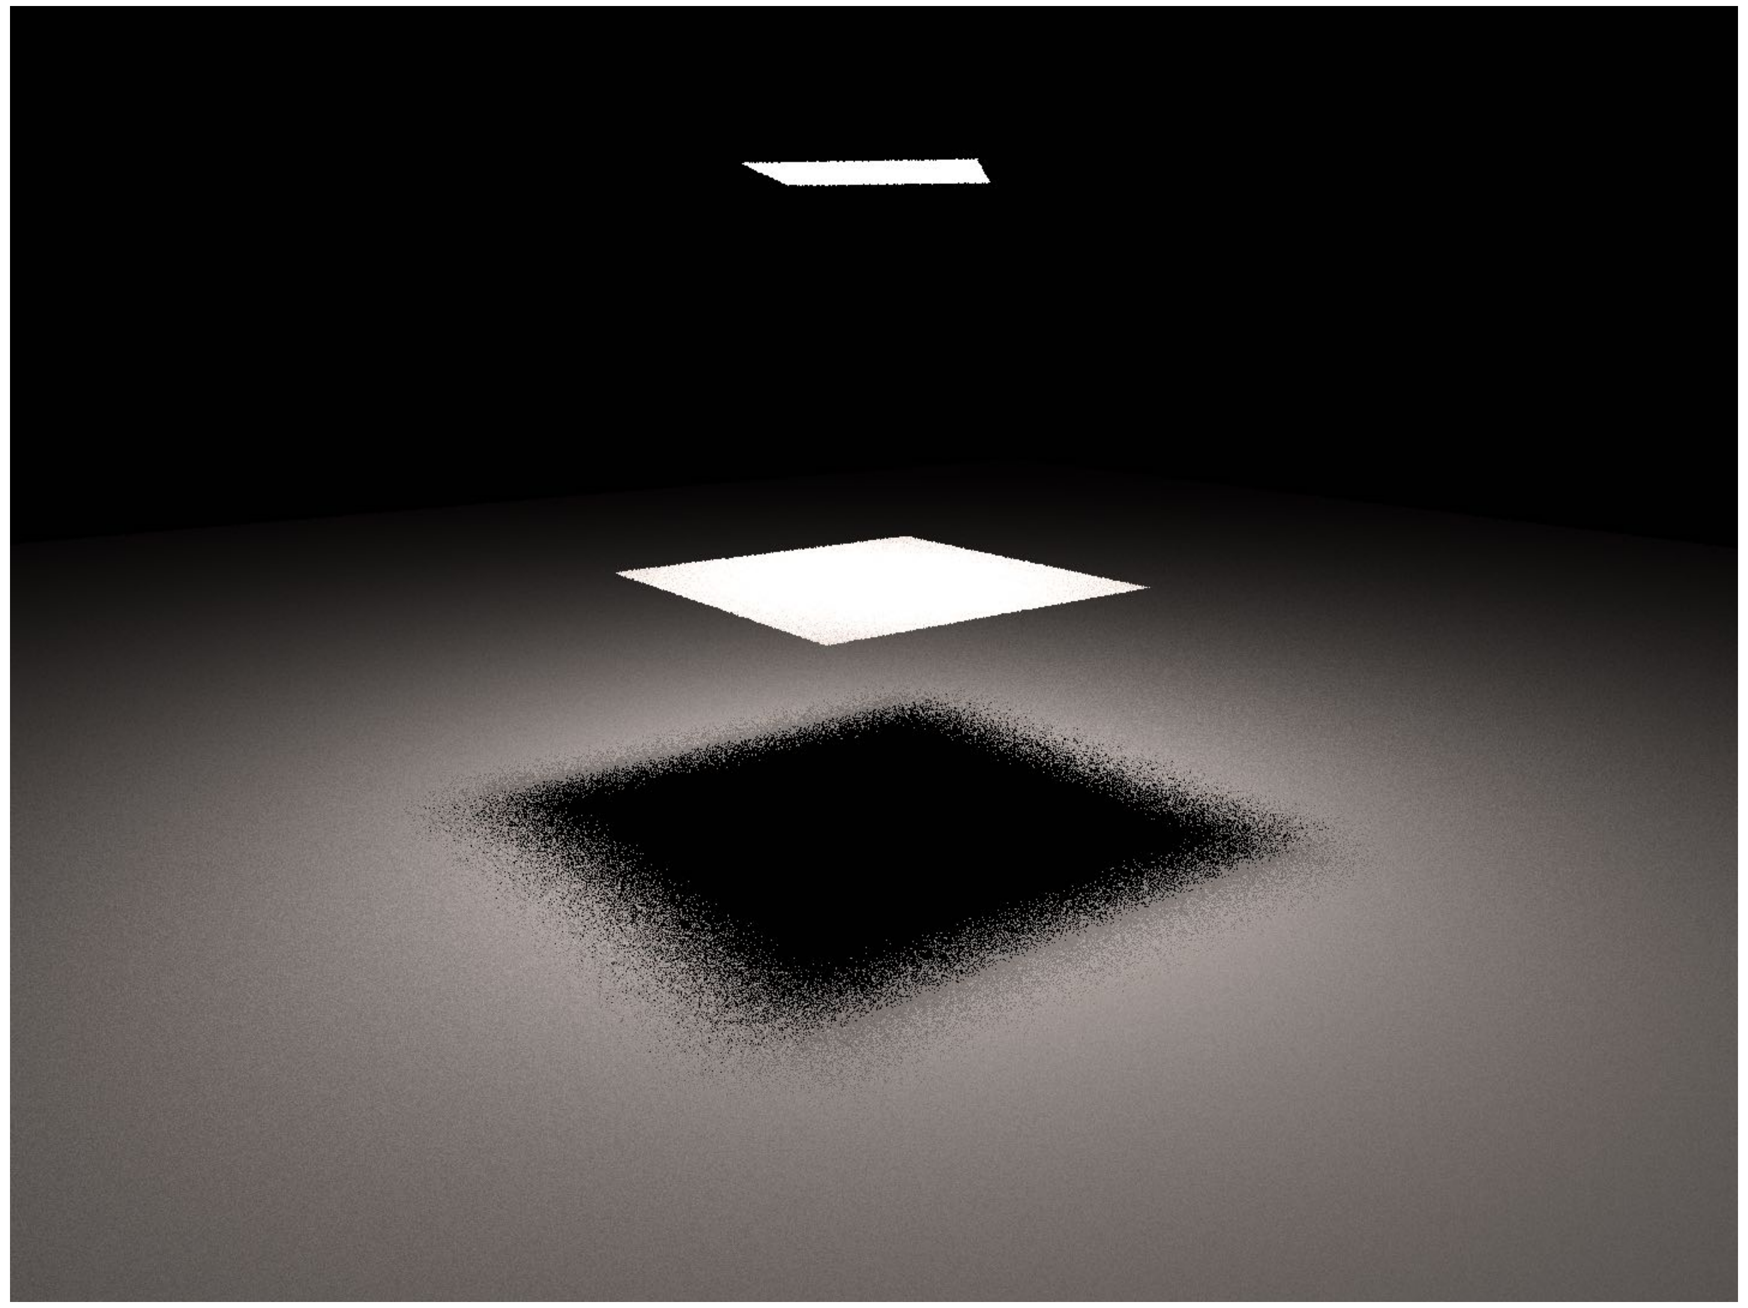

Area

1 shadow ray

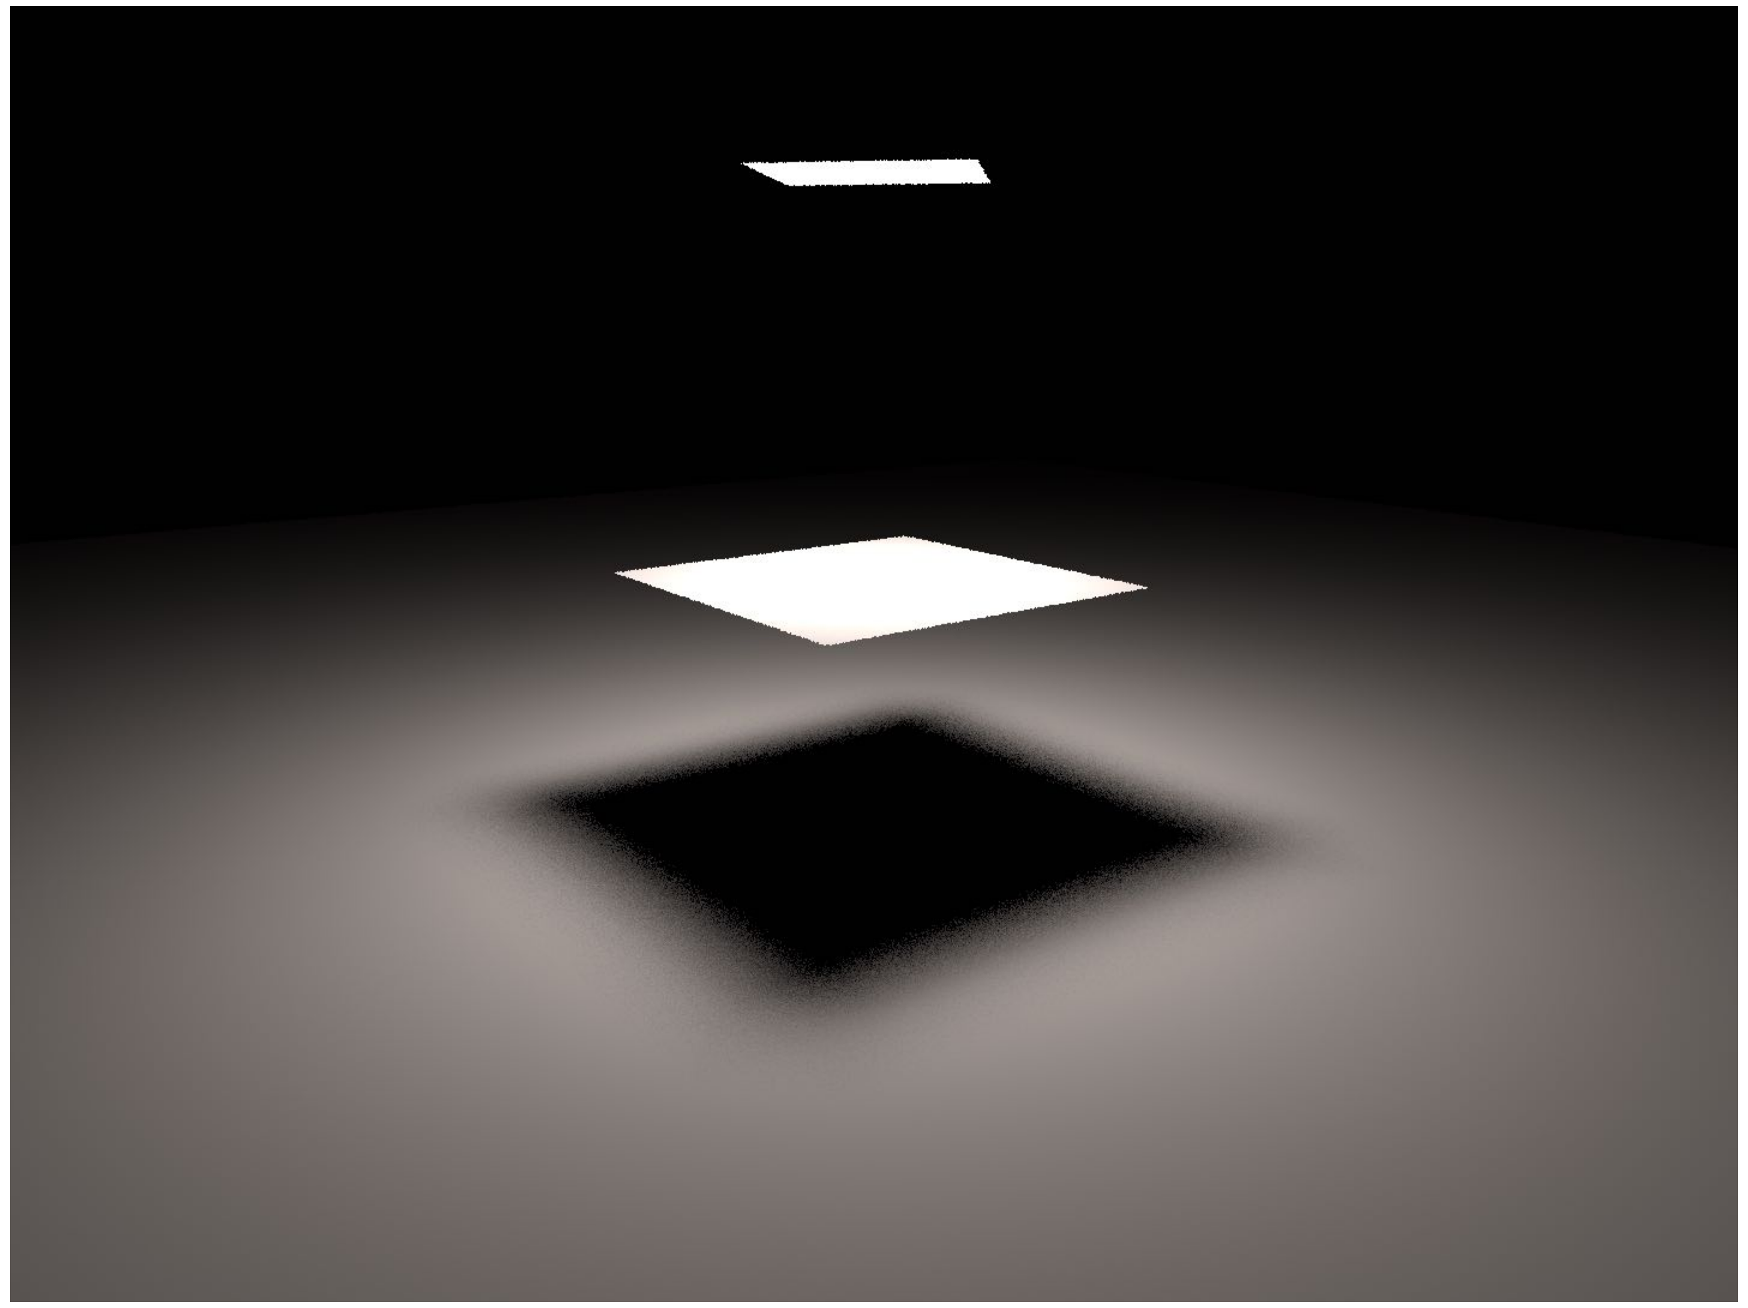

Area

16 shadow rays

Why is area better than hemisphere?

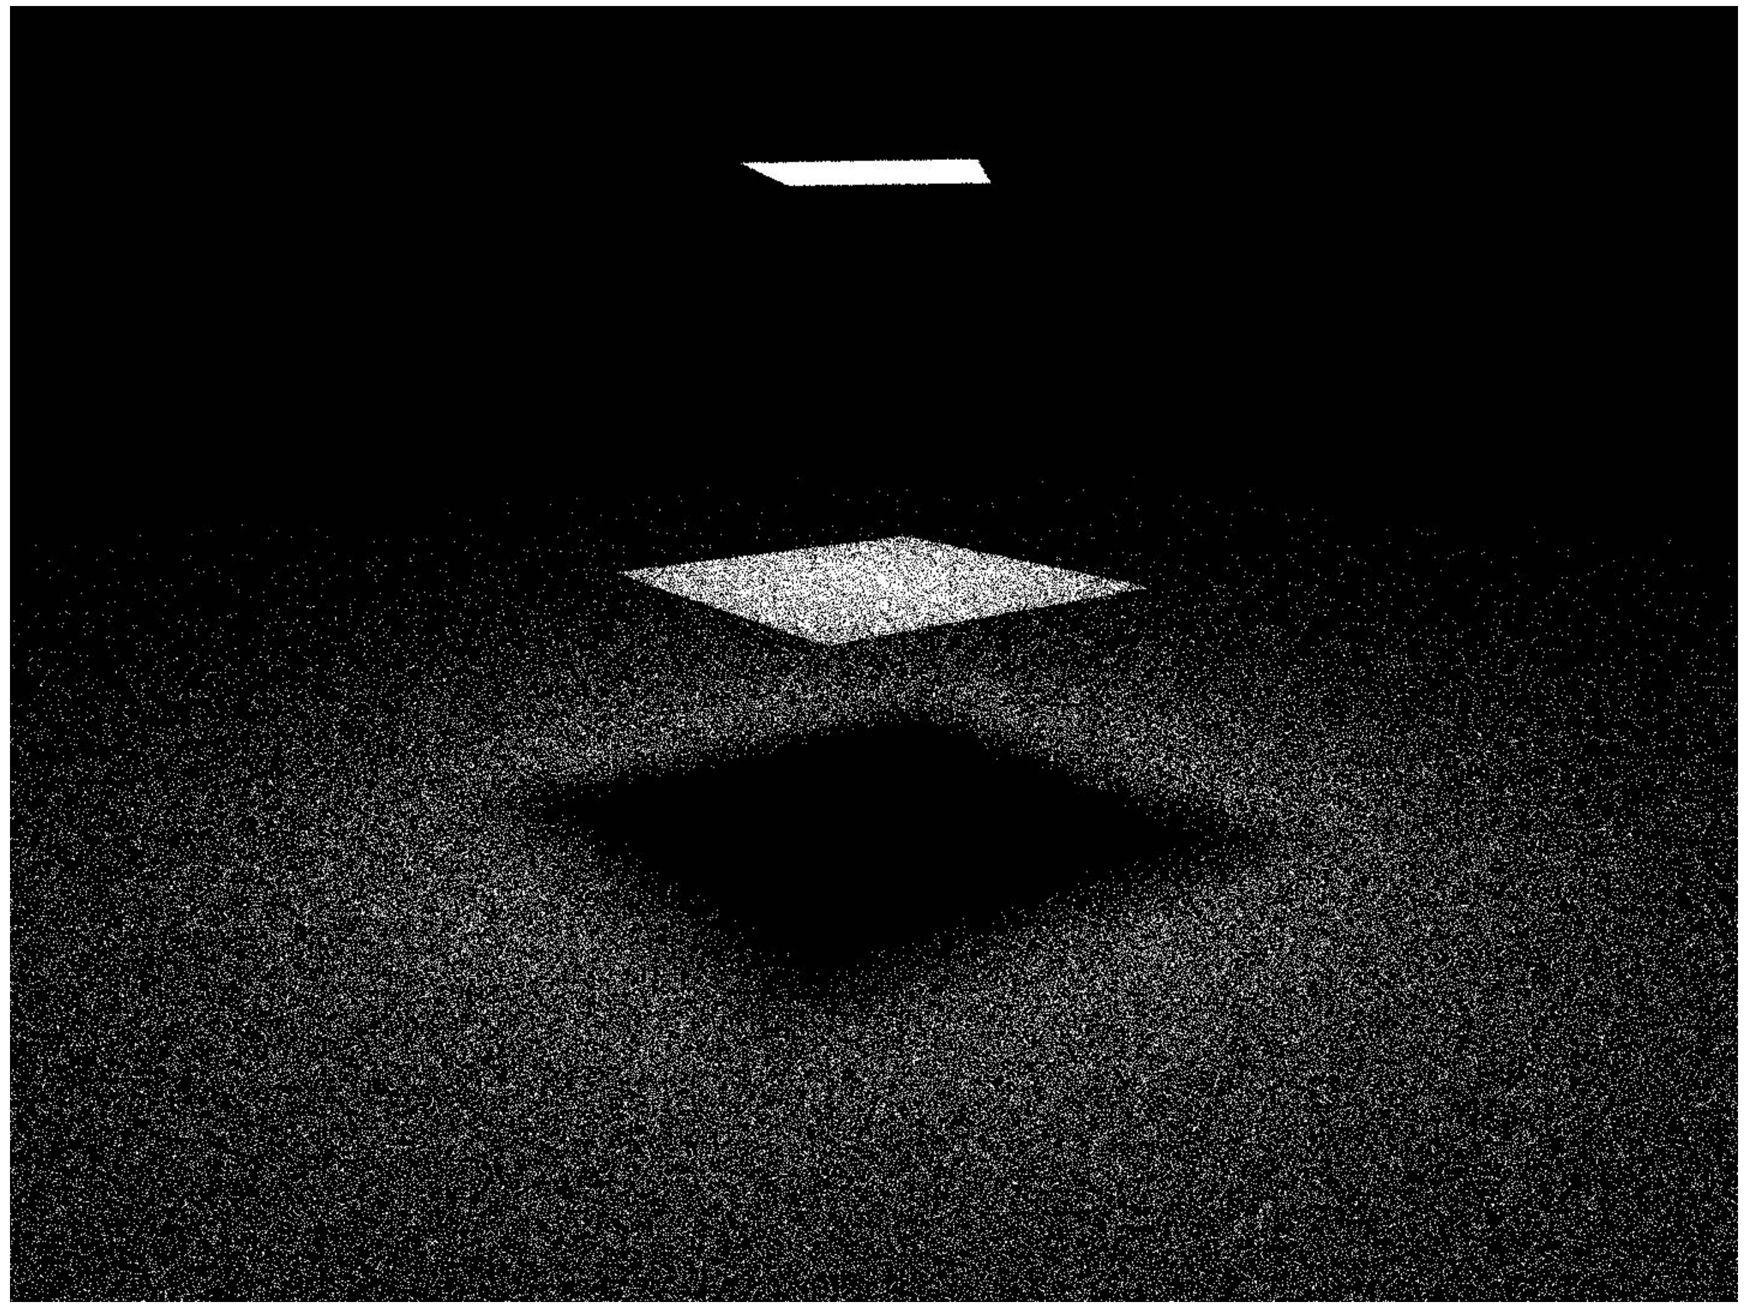

Hemisphere

16 shadow rays

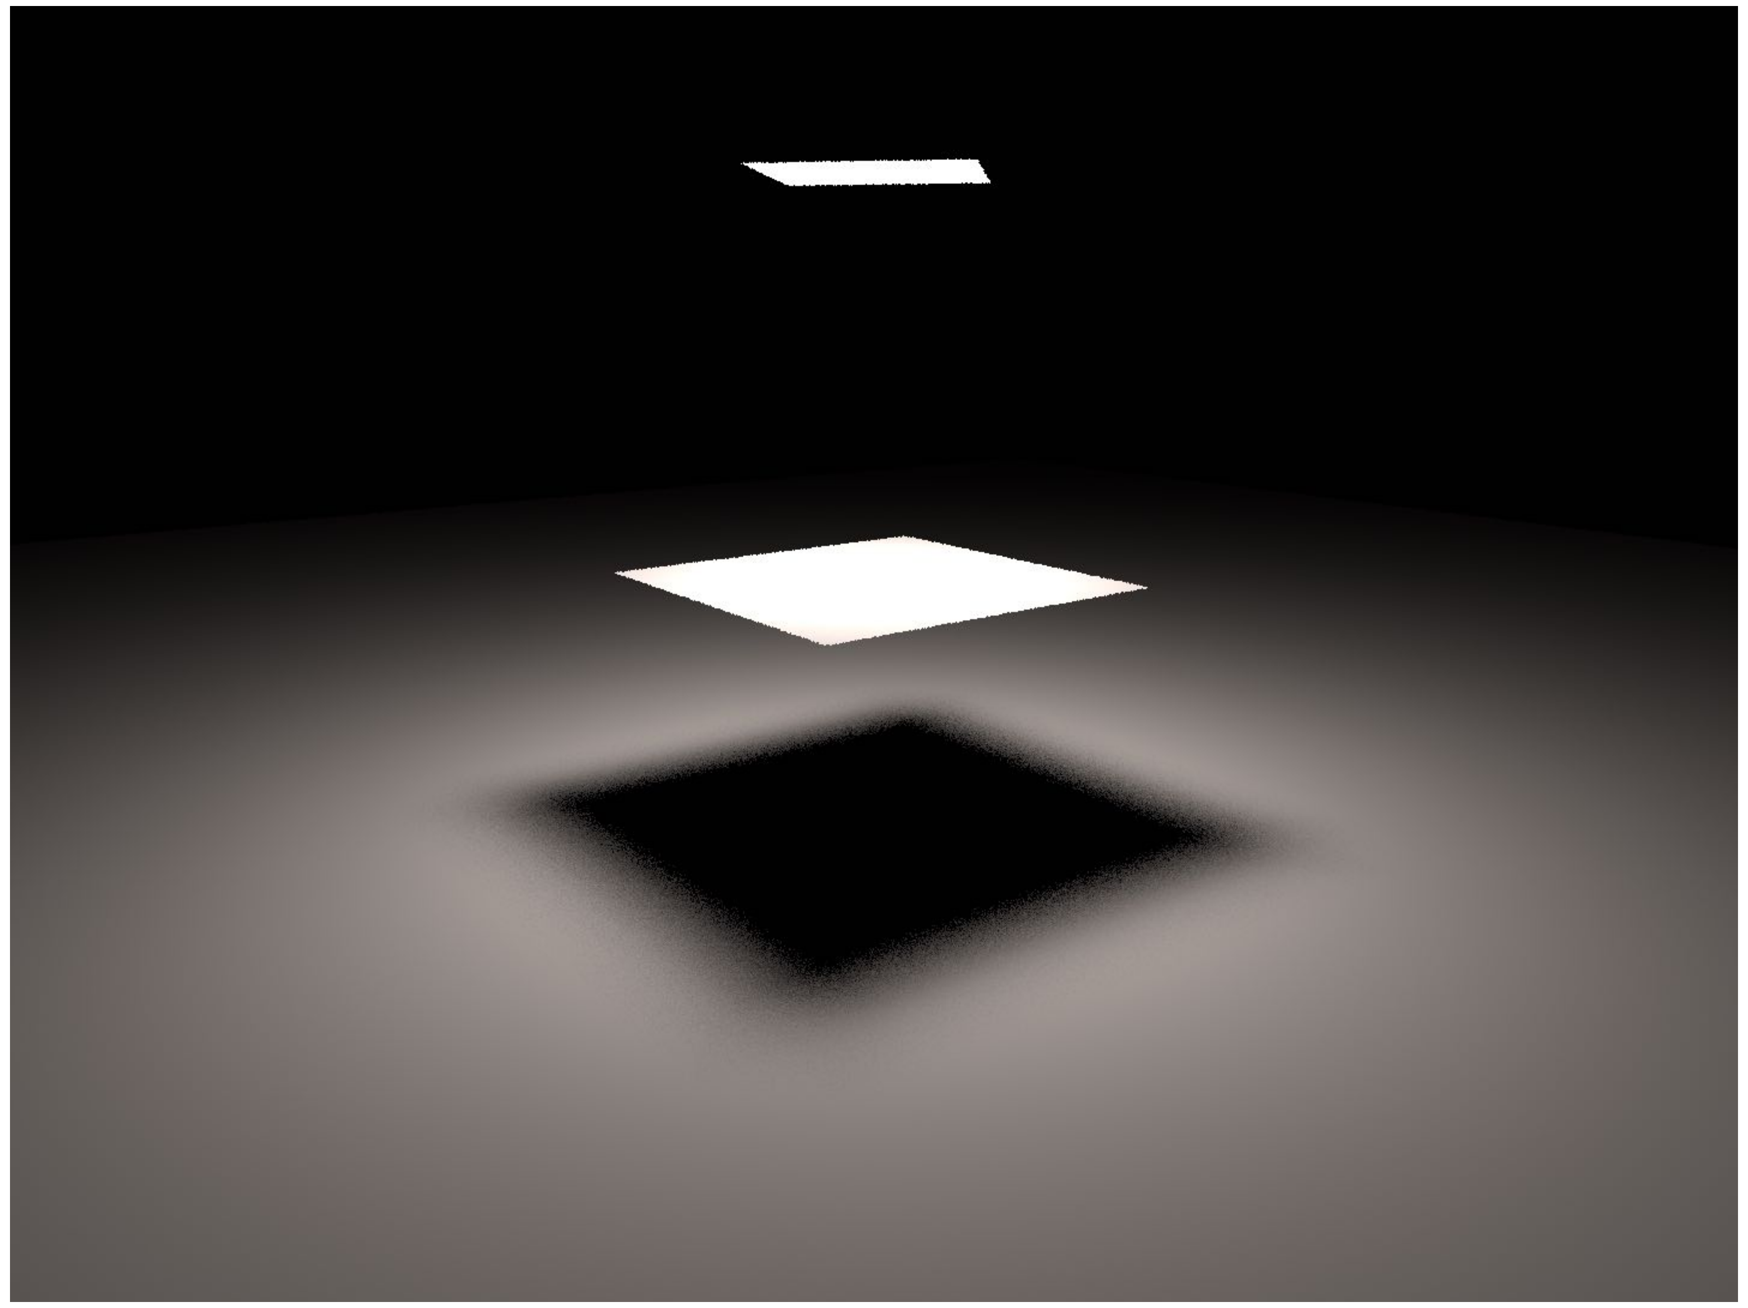

Area

16 shadow rays

Variance

- **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:

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

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

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:**

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

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

Summary: Monte Carlo integration

- **Monte Carlo estimator**
  - Estimate integral by evaluating function at random sample points in domain

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

- **The function (the estimator) is computed by a ray tracer!**

- **Useful in rendering due to estimate high dimension integrals**
  - Faster convergence in estimating high dimensional integrals than non-randomized methods
  - But it's still slow…
  - Suffers from noise due to variance in estimate (need many samples to produce good quality images)

- **Not discussed today: importance sampling = picking good samples to reduce variance**

Acknowledgements

- Thanks to Keenan Crane, Ren Ng, Pat Hanrahan and Matt Pharr for presentation resources