Lecture 11:

Materials (Part 2) + Monte Carlo Integration Basics

Interactive Computer Graphics
Stanford CS248A, Winter 2025

Review: radiometry and illumination

Review: differential solid angles

Sphere with radius r

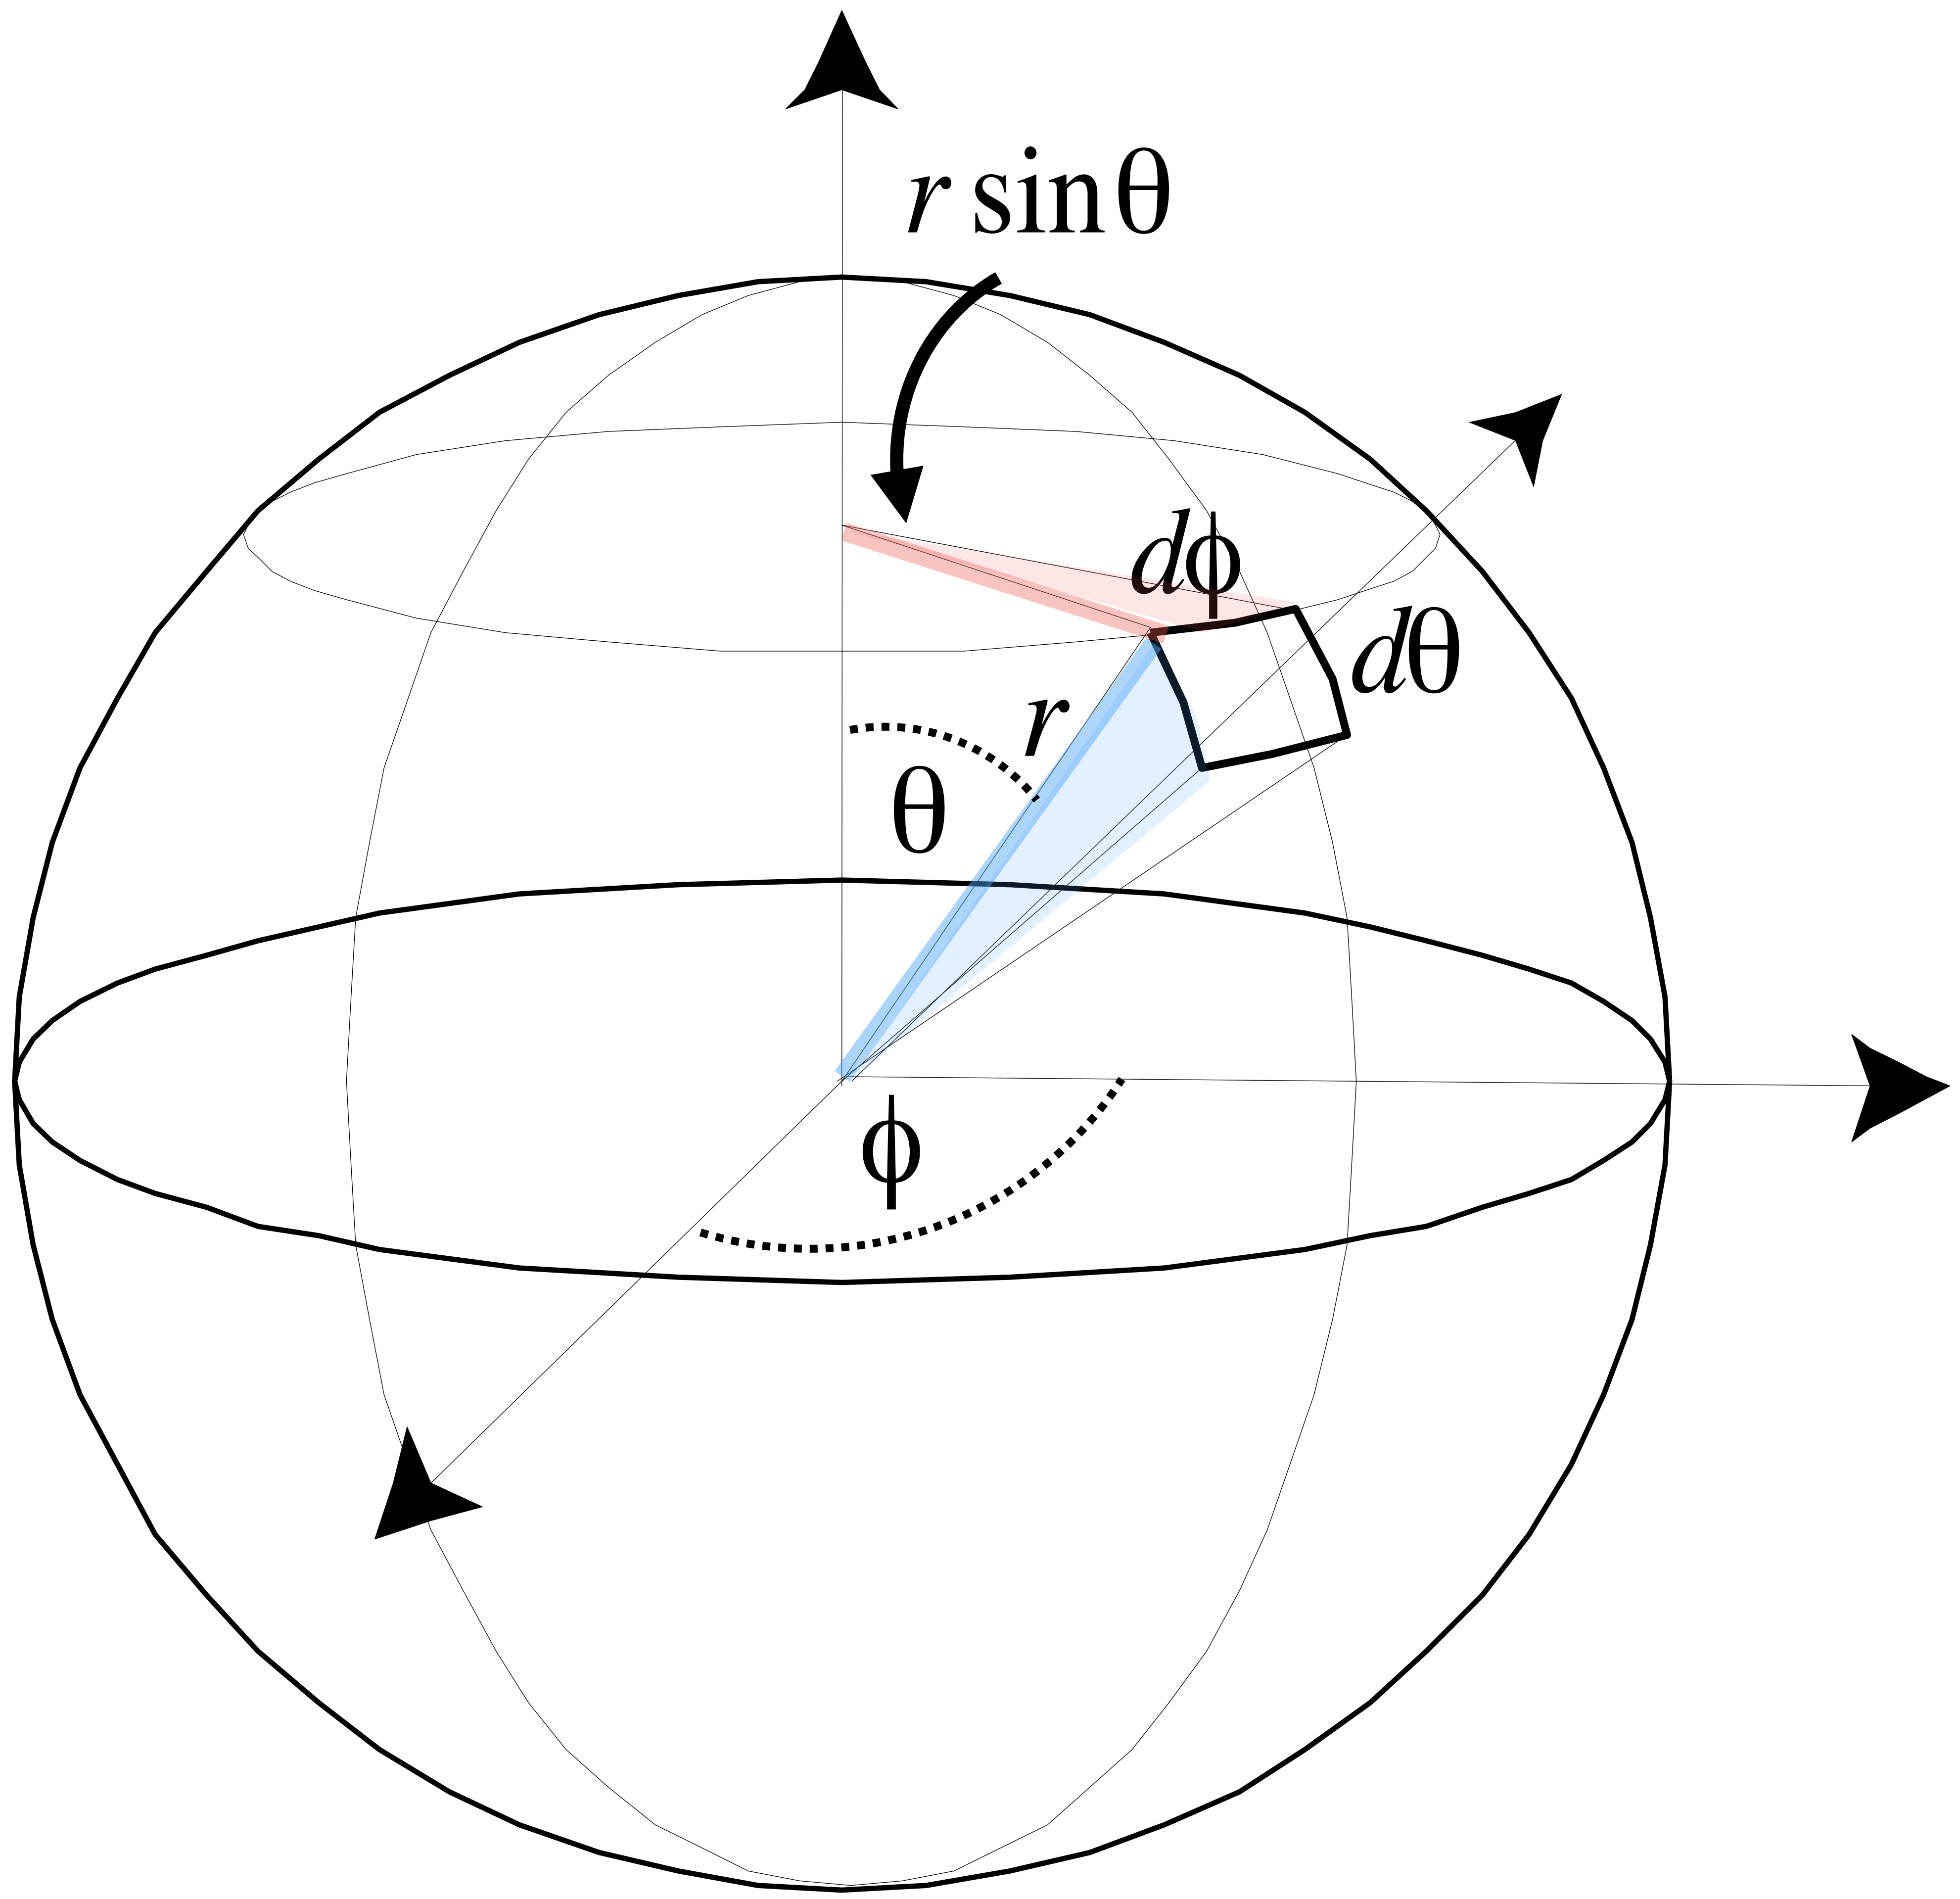

$$dA = (r\, d\theta)(r \sin\theta\, d\phi)$$
$$= r^2 \sin\theta\, d\theta\, d\phi$$

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

Review: radiance

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

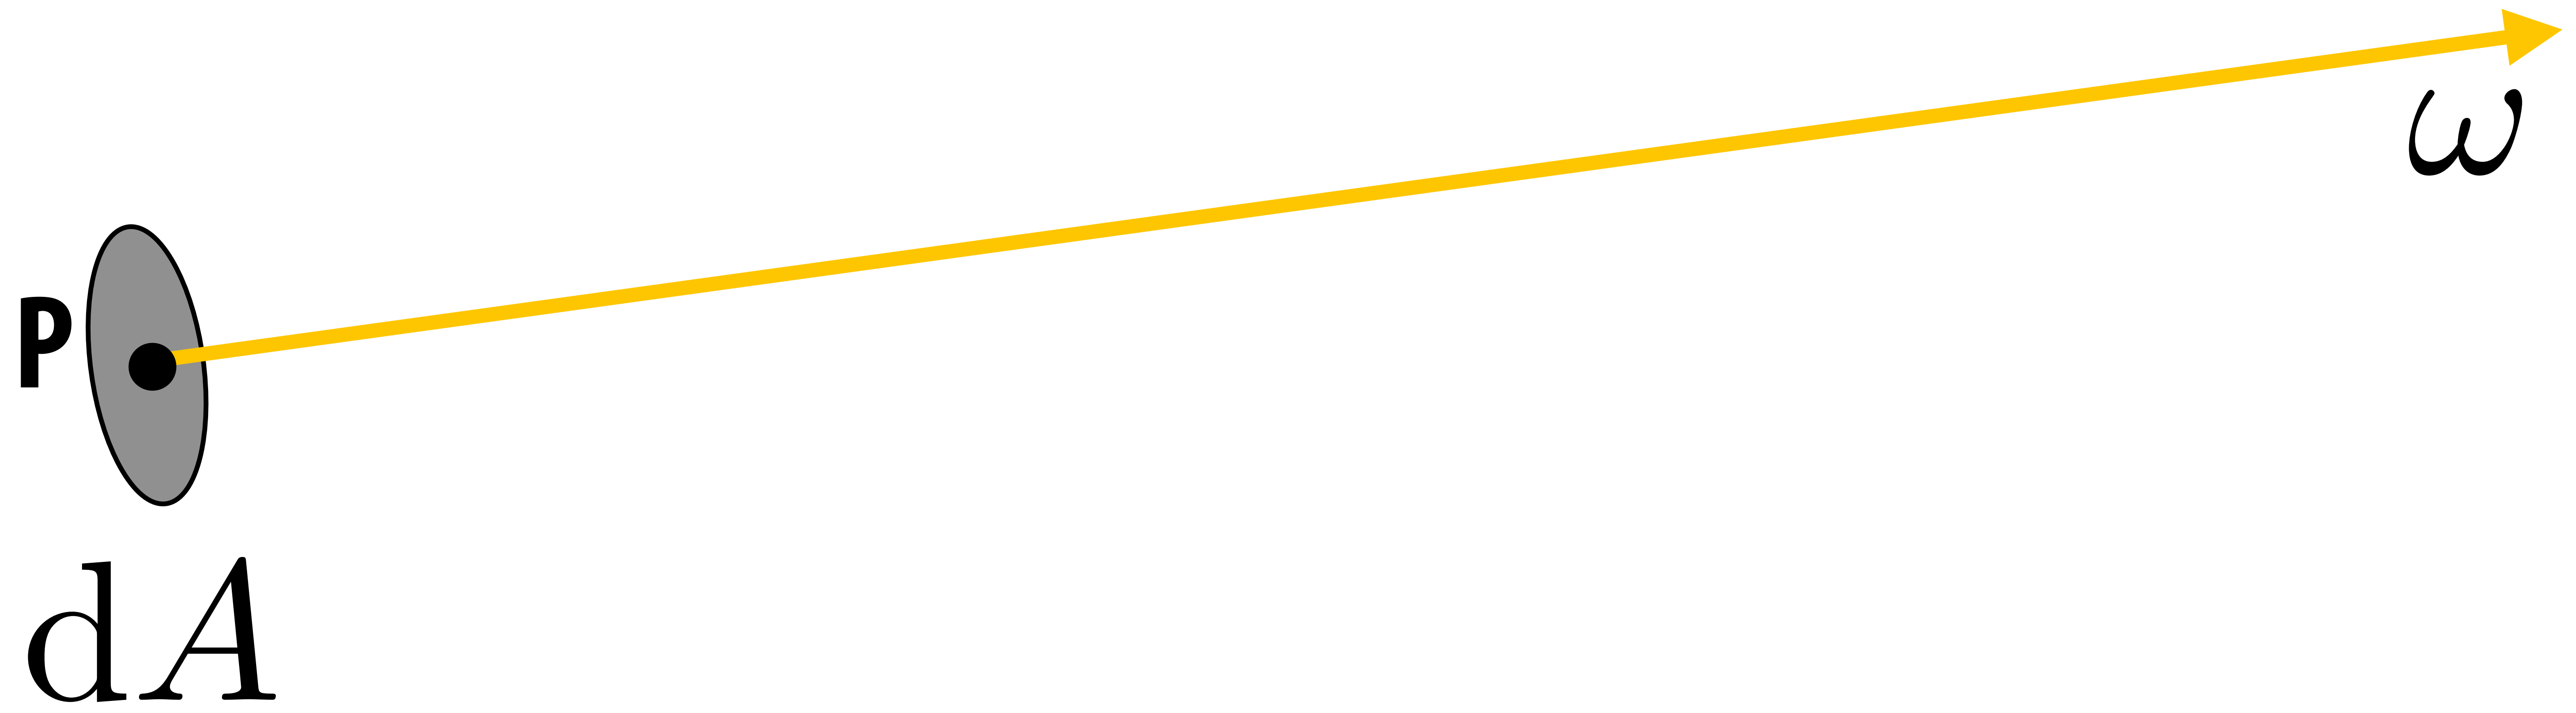

- **Radiance is the solid angle density of irradiance (irradiance per unit direction)**

  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)

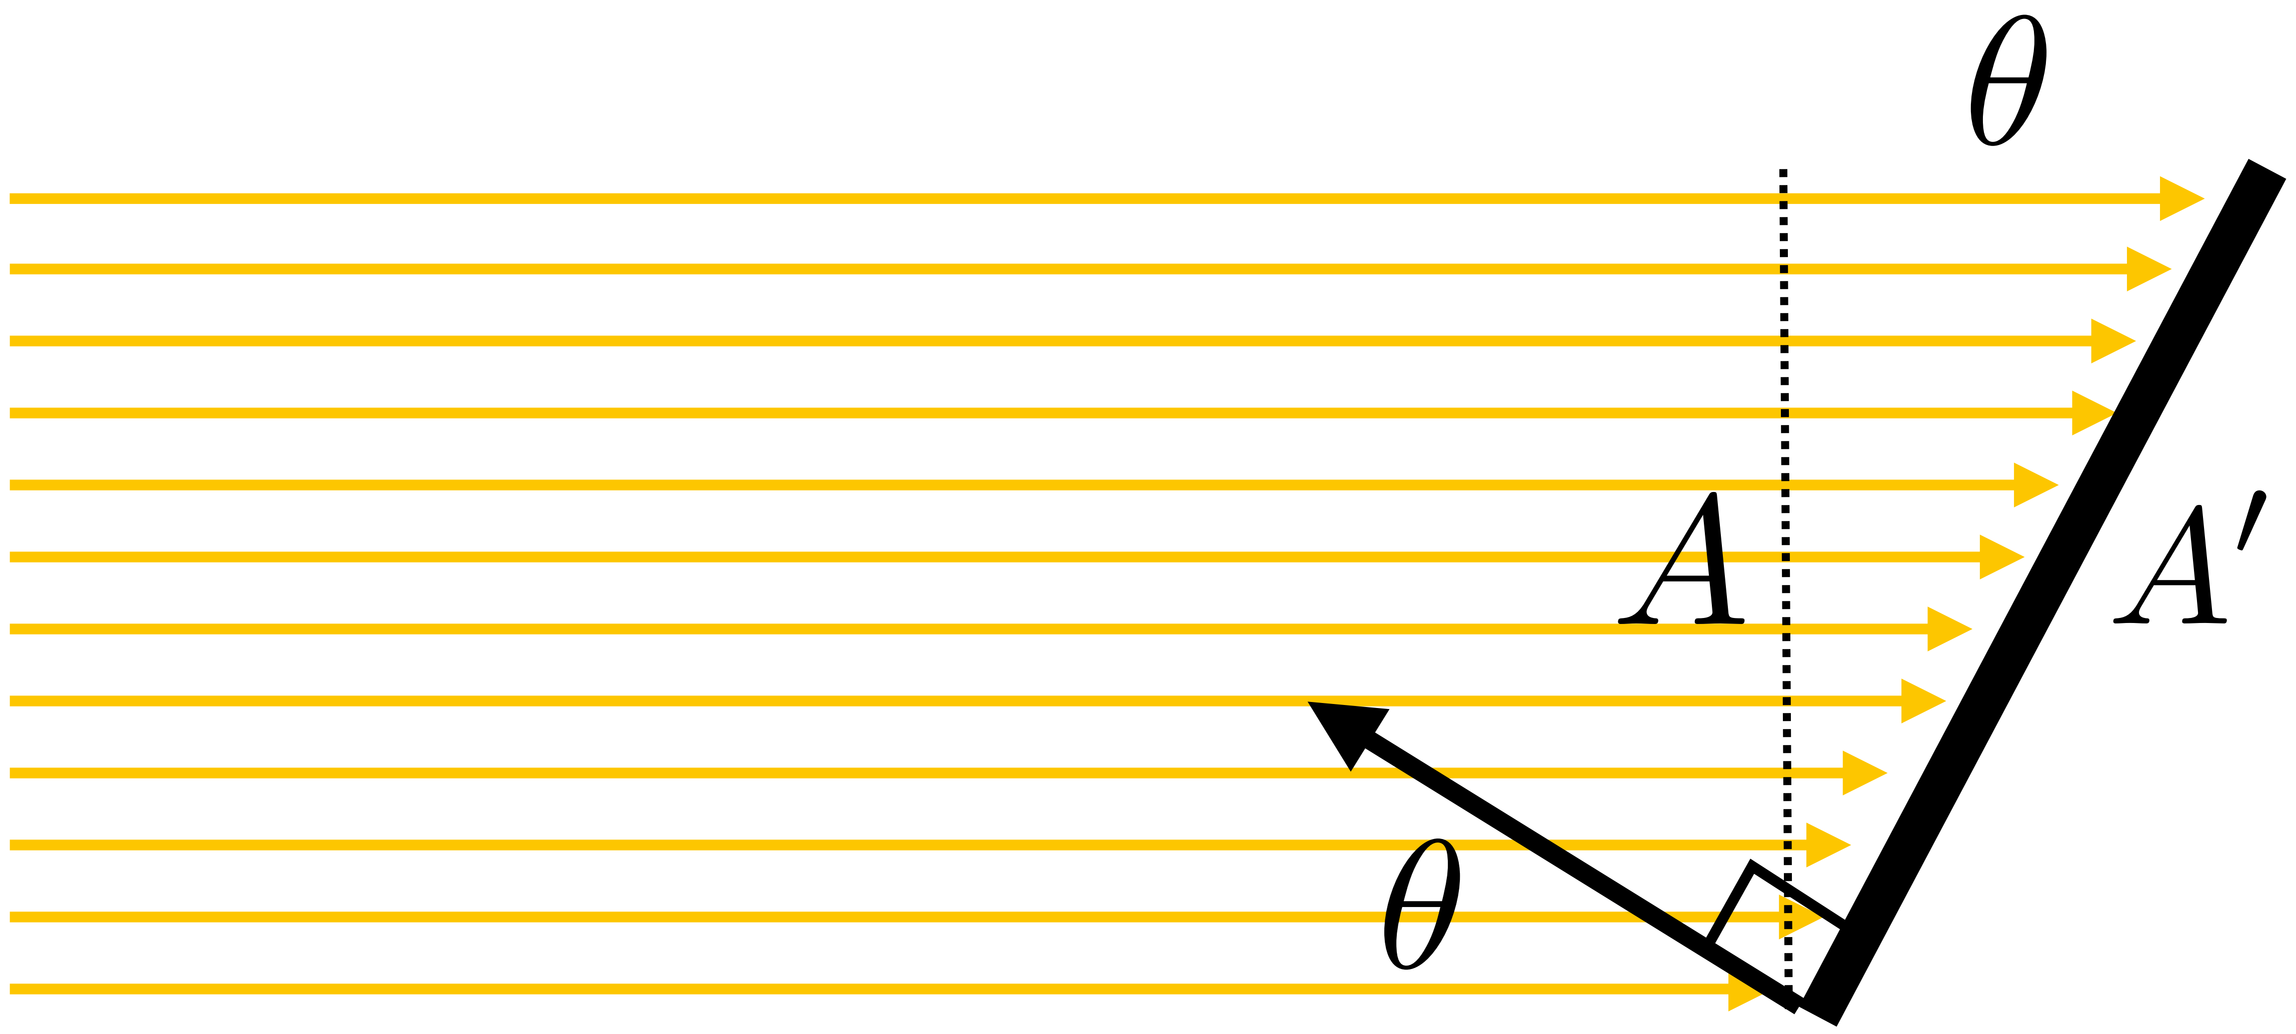

$$A = A' \cos\theta$$

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

How much light hits the surface at point p from multiple light sources? (What is irradiance at point p?)

$$\sum_i L_i \cos\theta_i$$

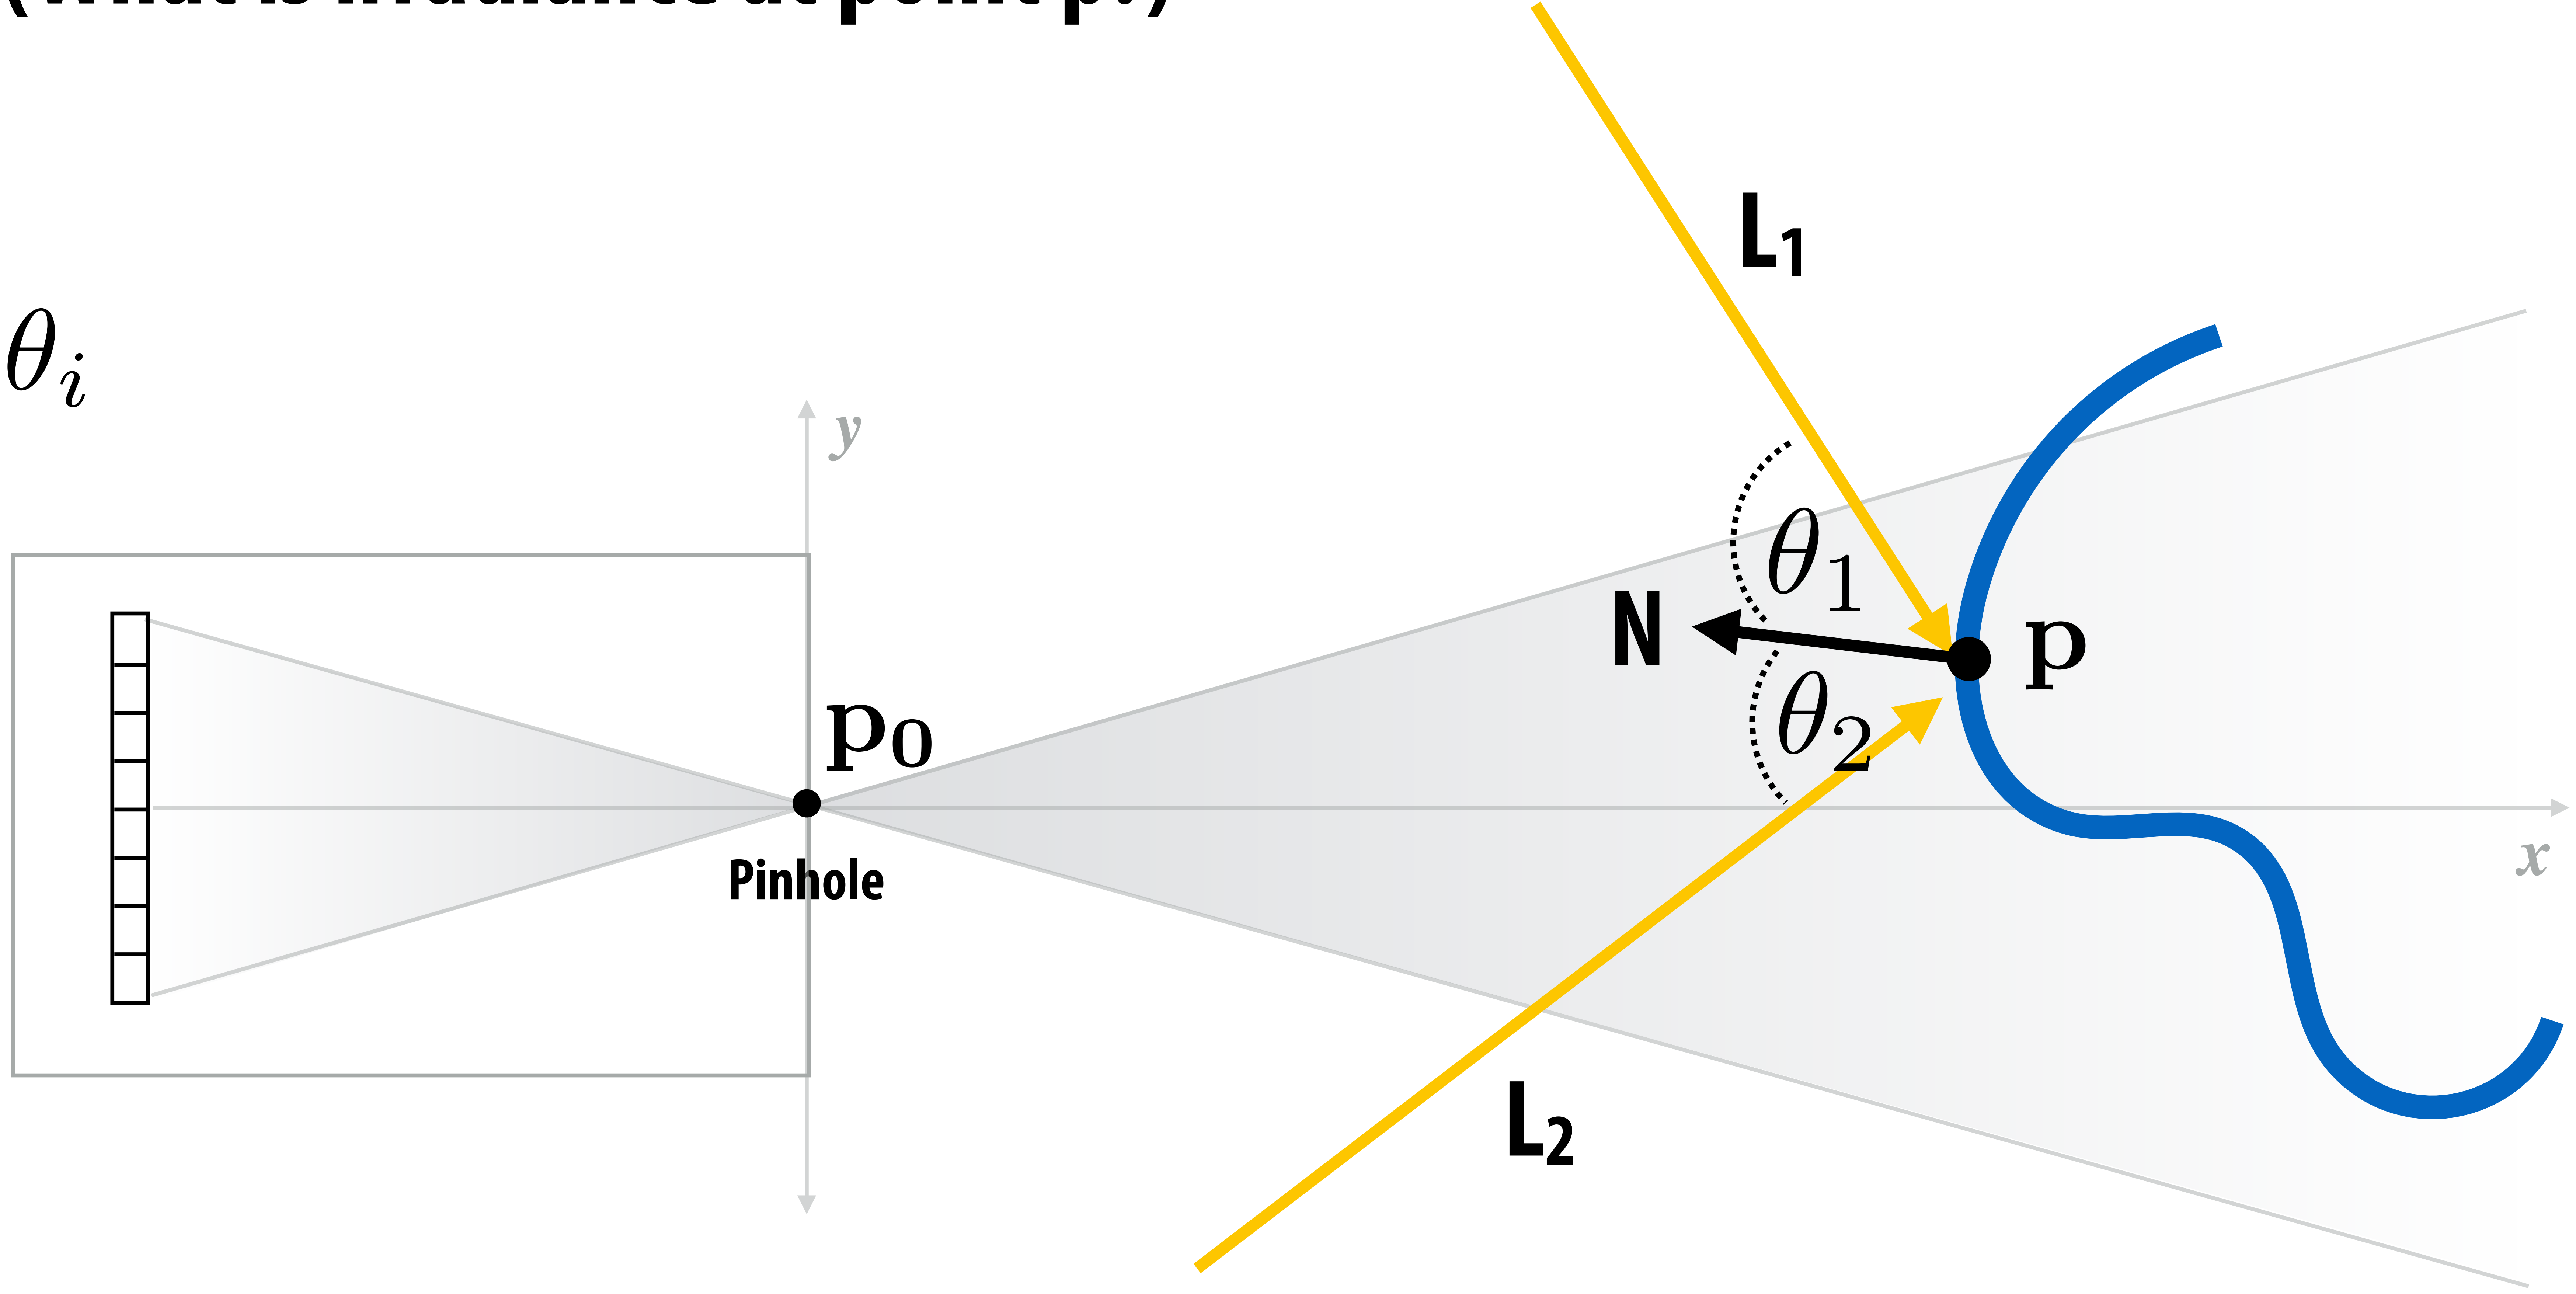

How much light hits the surface at point p from light from all directions?

(What is irradiance at point p?)

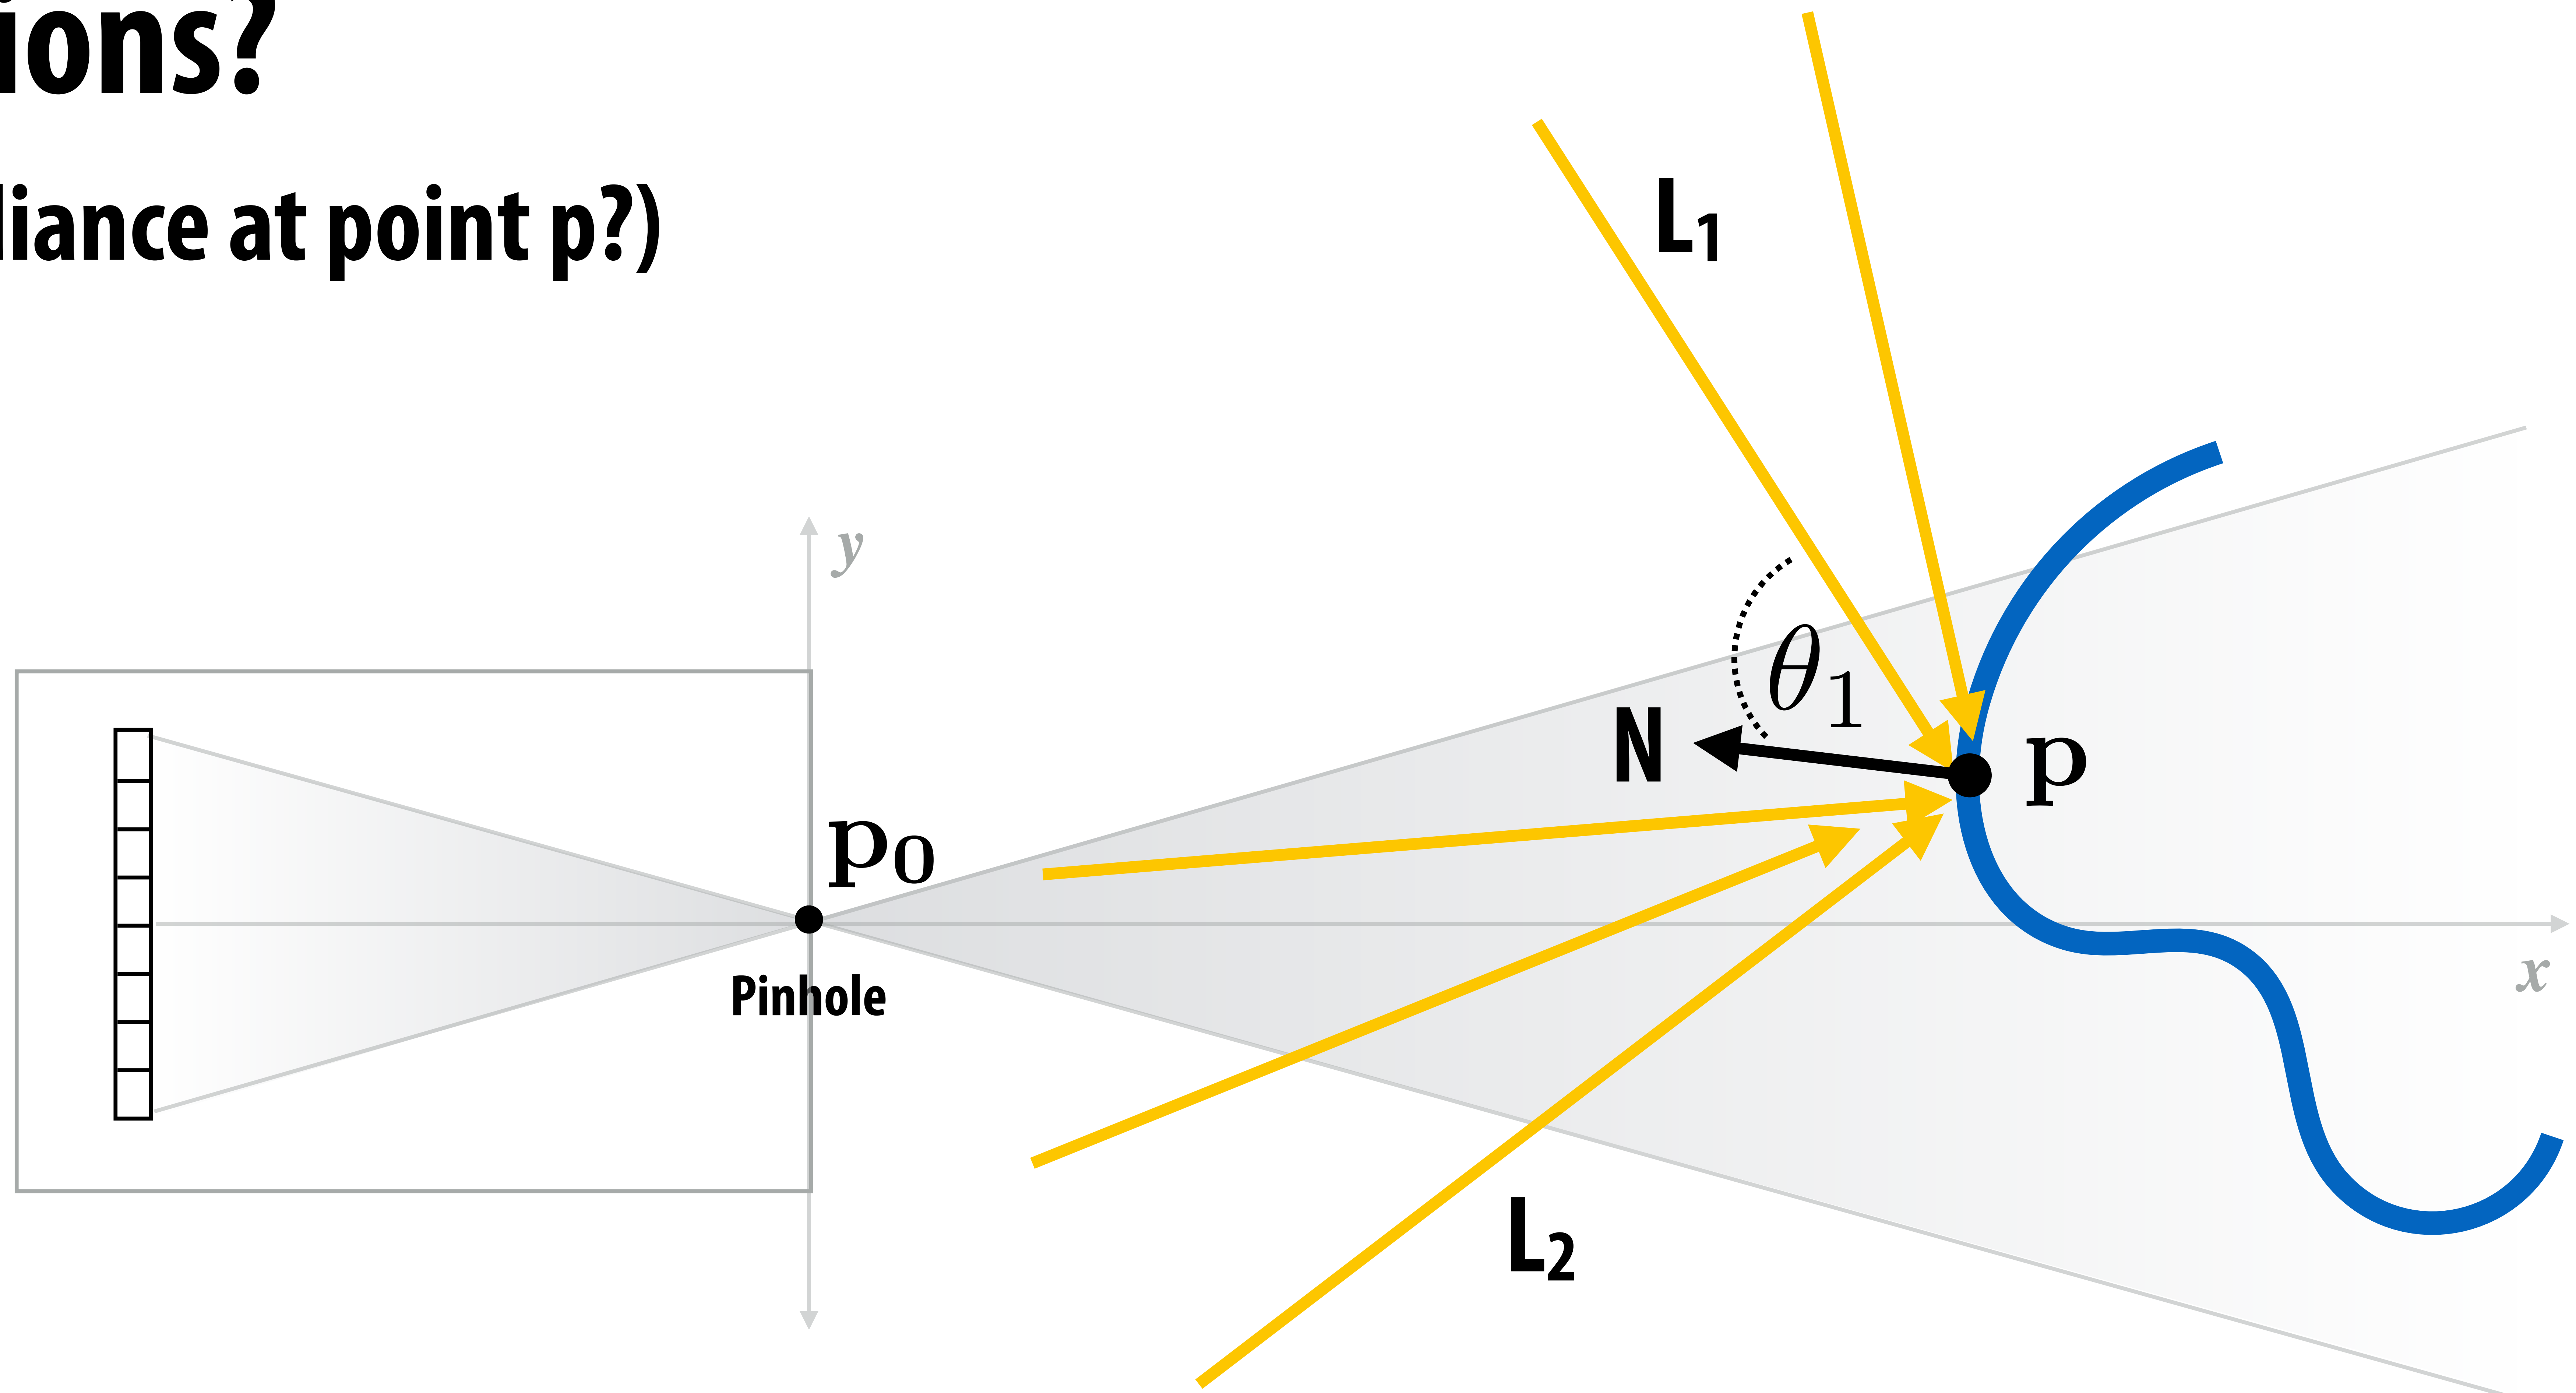

$$\int_{S^2} L(\omega_i)\cos\theta_i \, \mathrm{d}\omega_i = \int_0^{2\pi}\int_0^{\pi} L(\theta_i,\phi_i)\cos\theta_i \sin\theta_i \, \mathrm{d}\theta_i \mathrm{d}\phi_i$$

Irradiance at point X from a uniform area source

Assume single light source in scene, so incoming light is 0 except from directions toward the light

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

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

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

Constant
(it's a uniform source)

Total projected solid angle

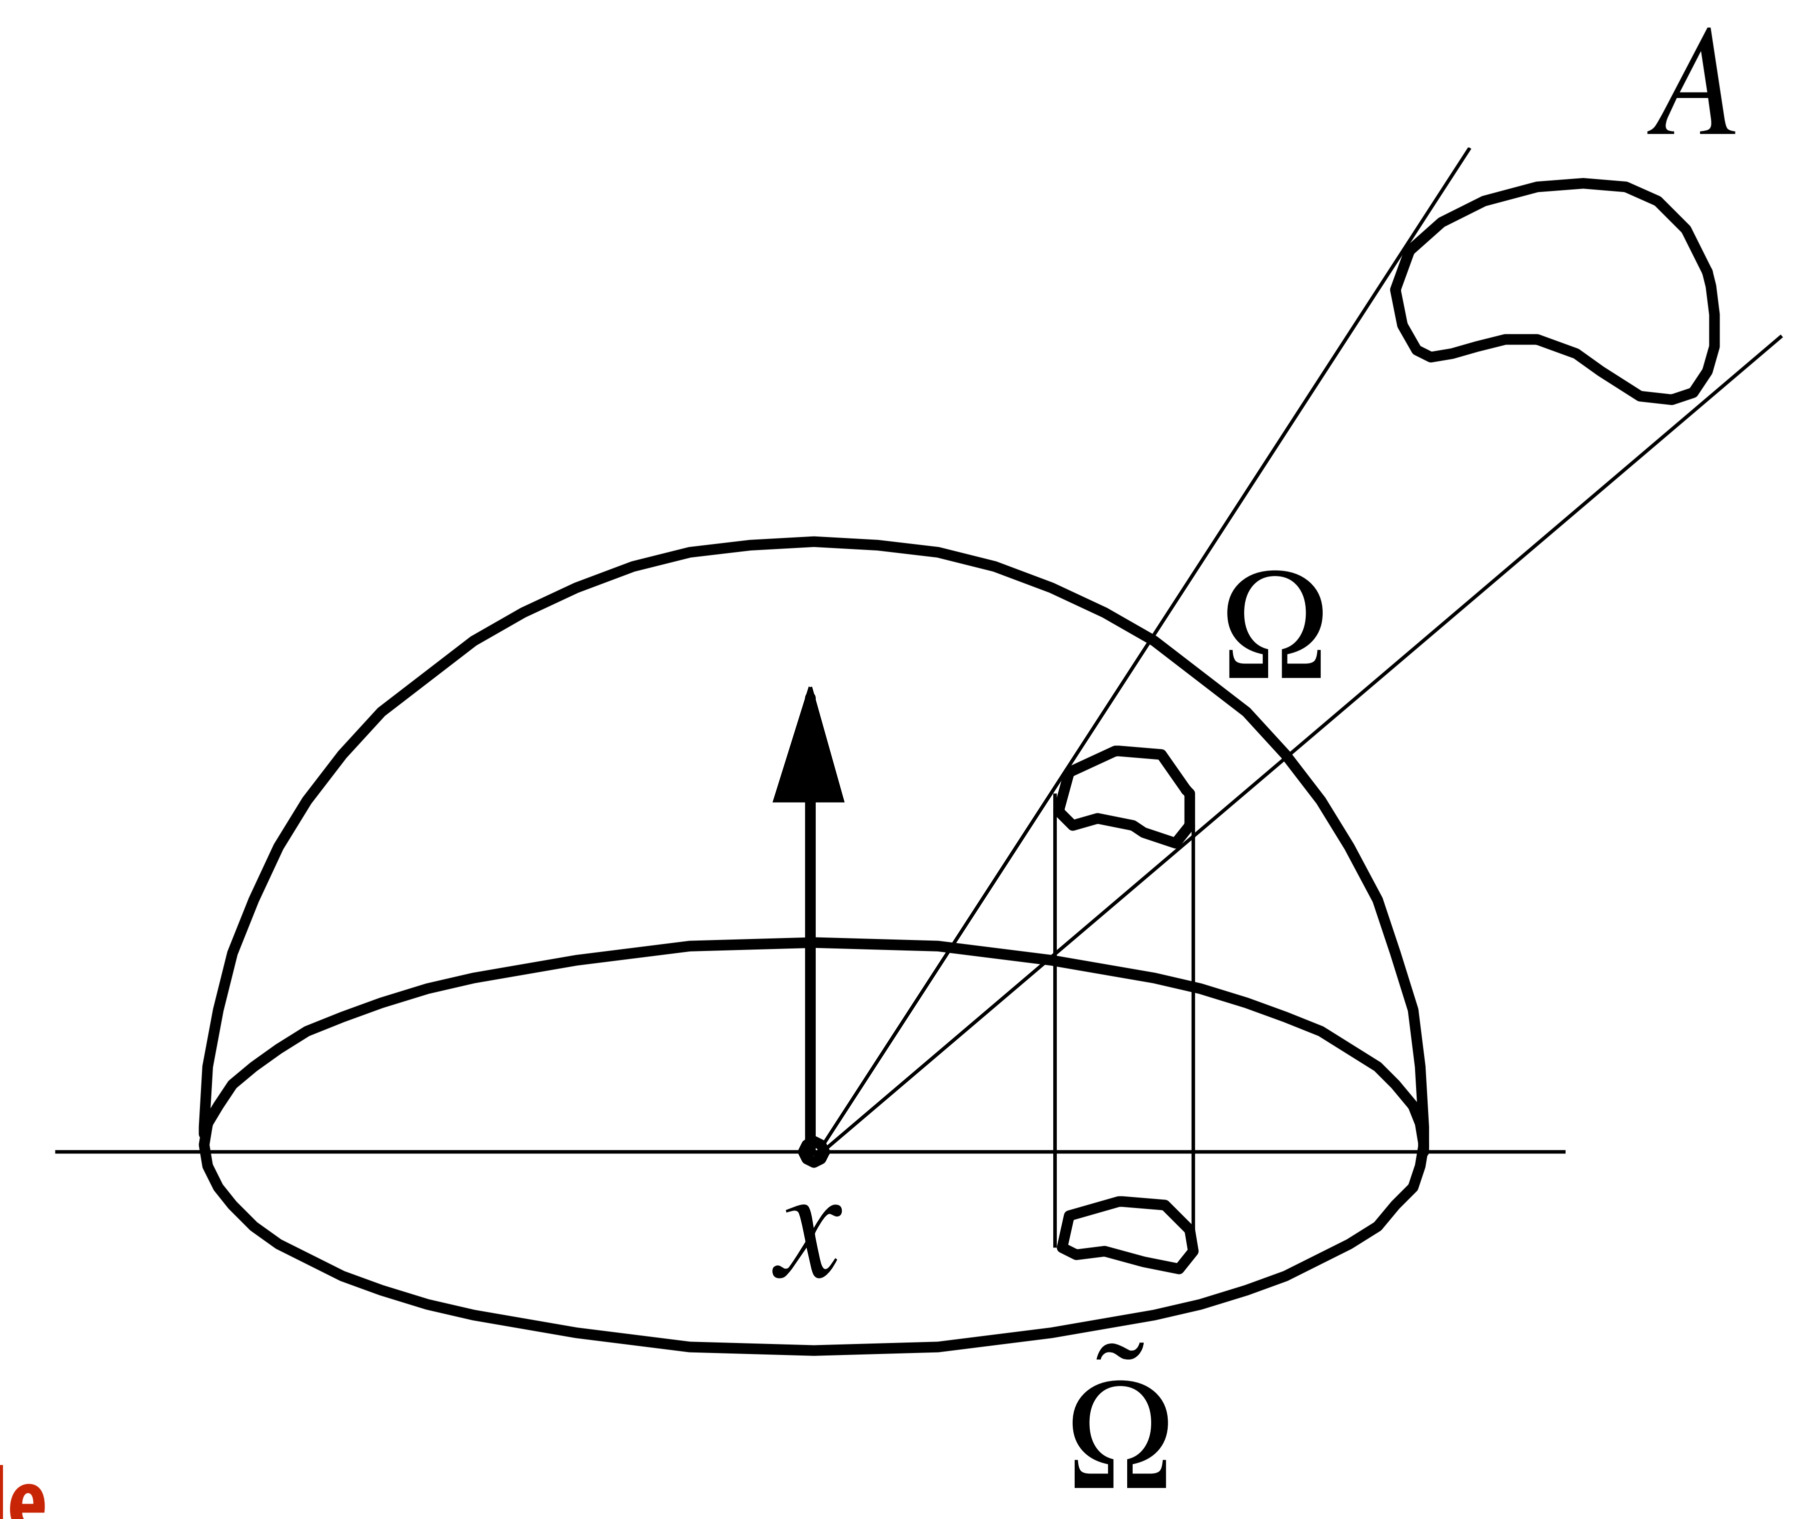

Irradiance at point X from a uniform area source

Reparameterize integral over solid angle to integral over area of light source.

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

Reparameterization: now integrate over light source area, instead of solid angle

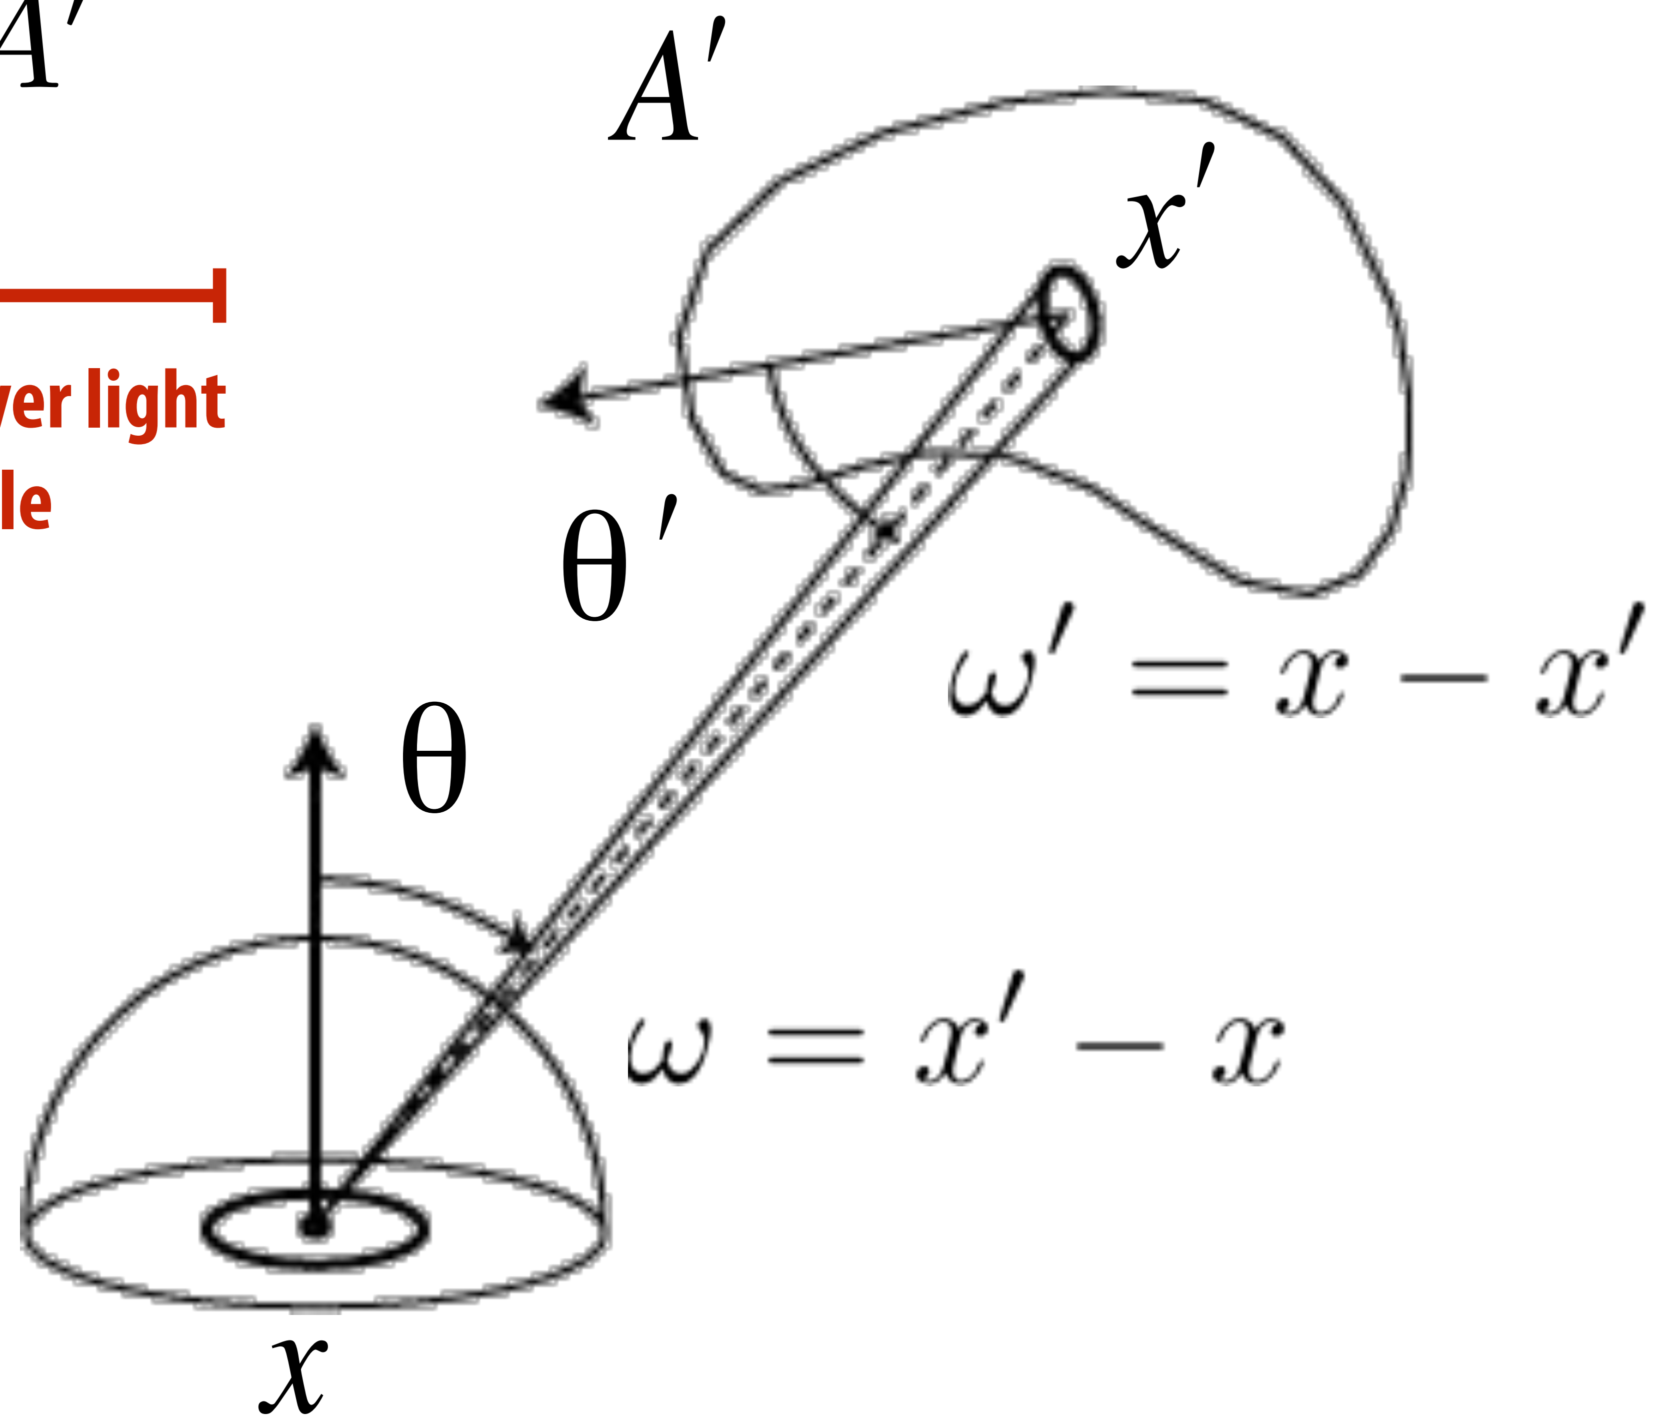

Integral reparameterization:

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

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

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

Review: the reflection equation

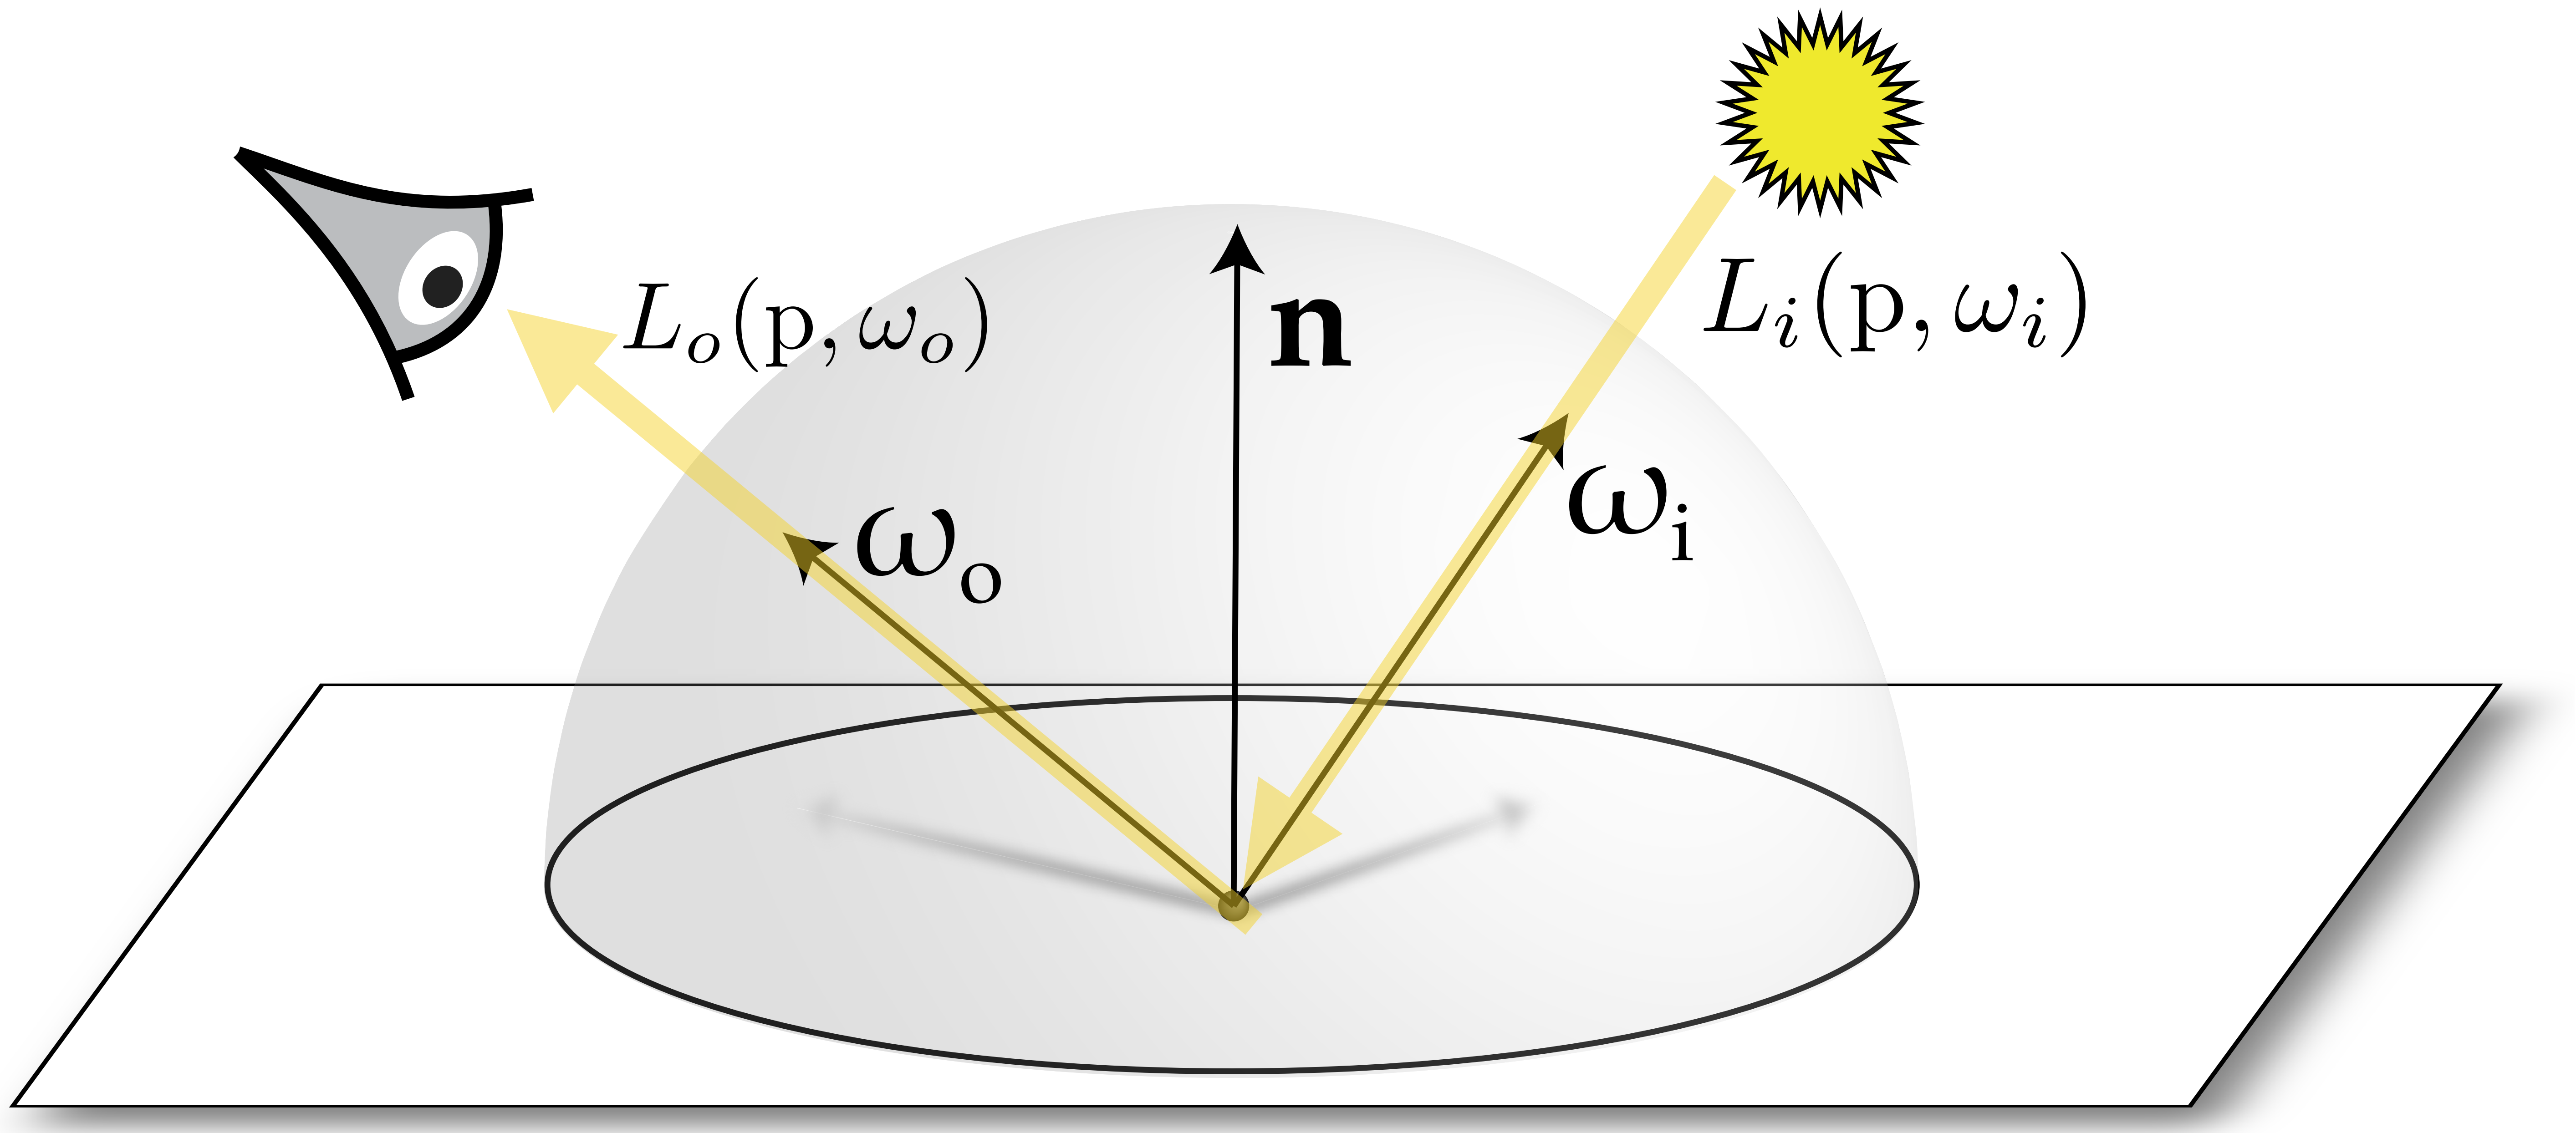

$$L_{\mathrm{o}}(\mathrm{p}, \omega_{\mathrm{o}}) = \int_{\Omega^2} \underbrace{f_{\mathrm{r}}(\mathrm{p}, \omega_{\mathrm{i}} \rightarrow \omega_{\mathrm{o}})}_{\textbf{BRDF}} \underbrace{L_{\mathrm{i}}(\mathrm{p}, \omega_{\mathrm{i}}) \cos\theta_{\mathrm{i}}}_{\textbf{Illumination}} \, \mathrm{d}\omega_{\mathrm{i}}$$

More About Materials

Last time: diffuse materials

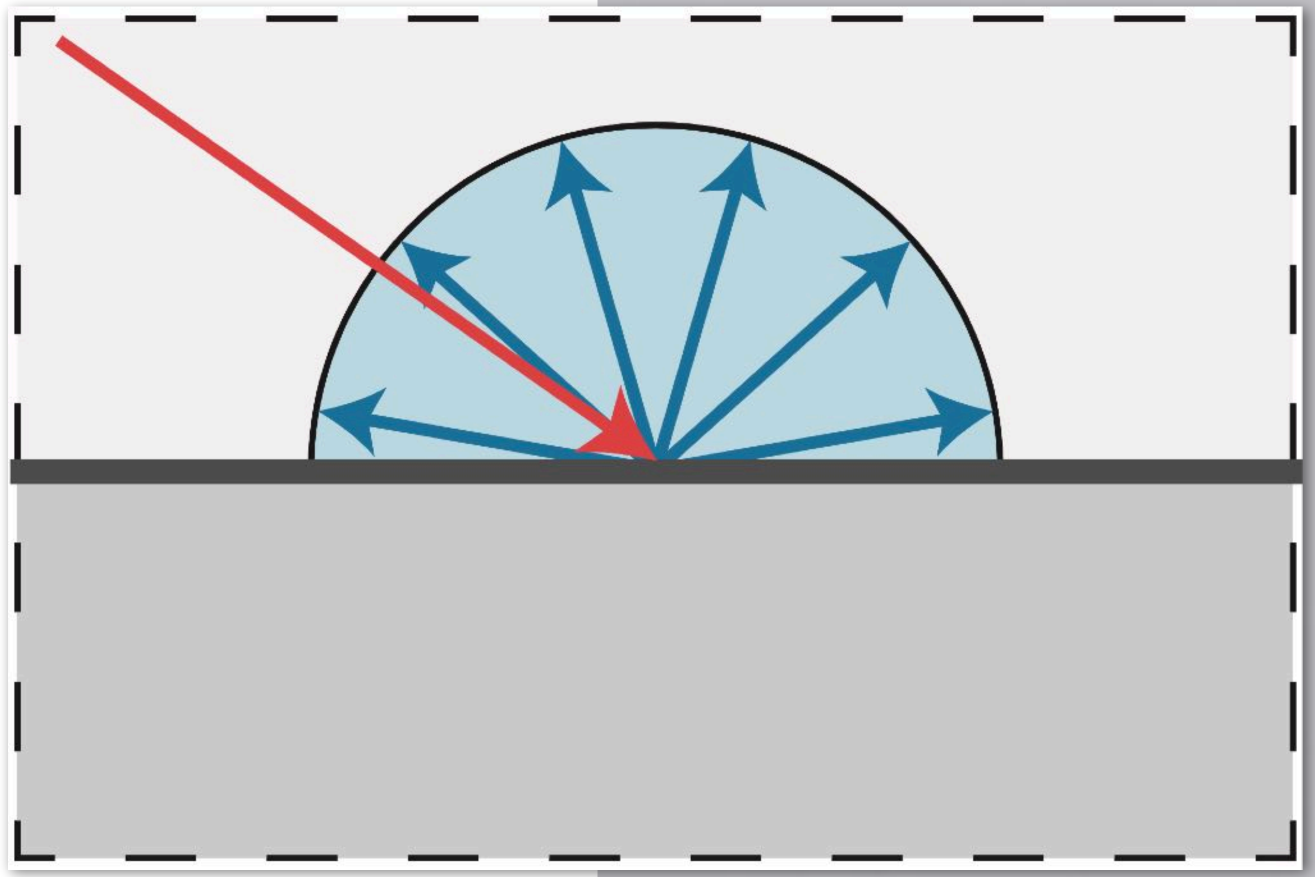

What is this material?

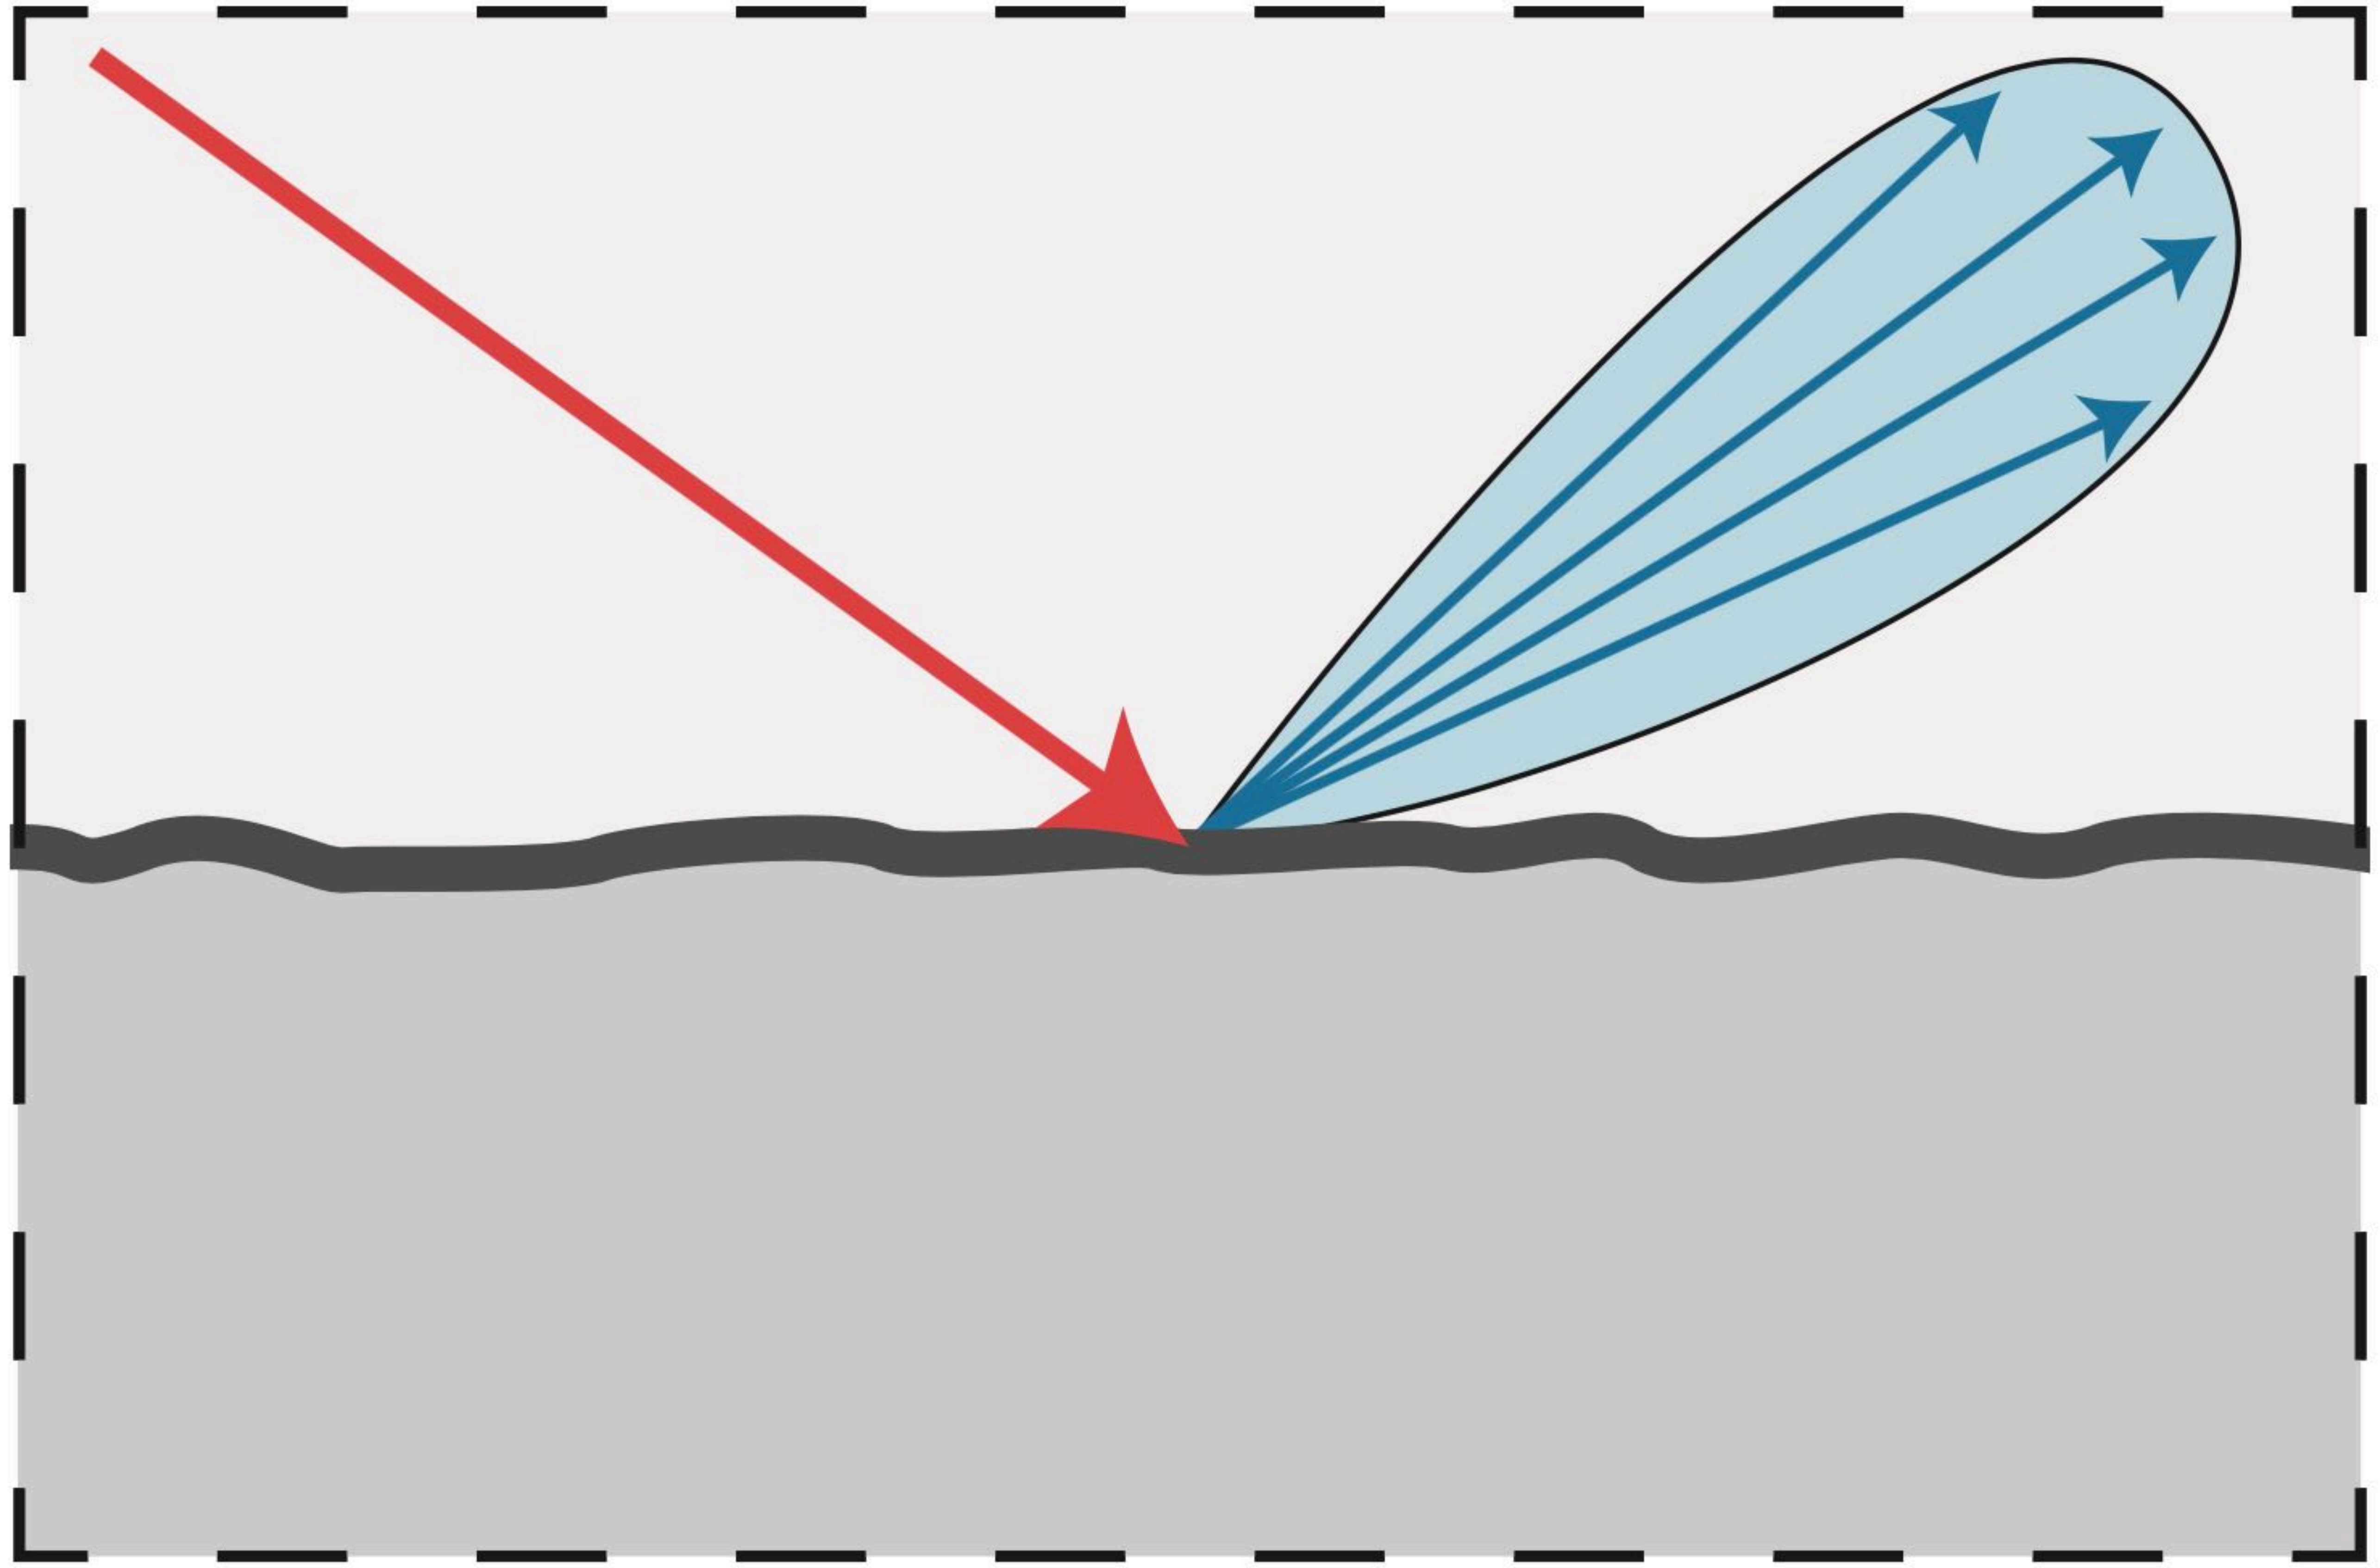

Glossy material (BRDF)

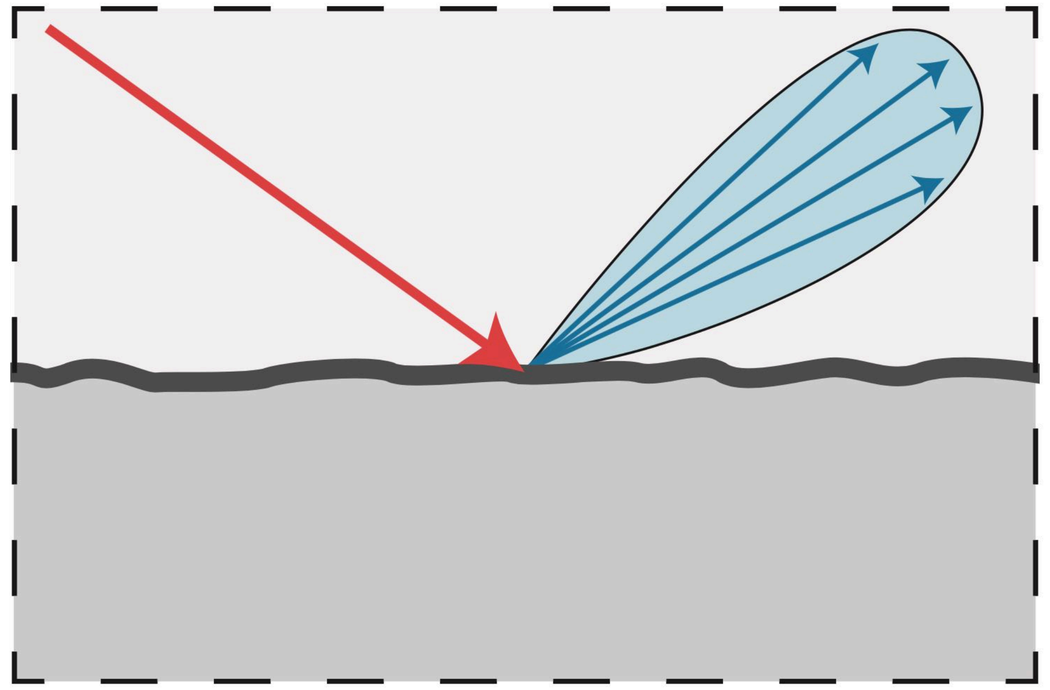

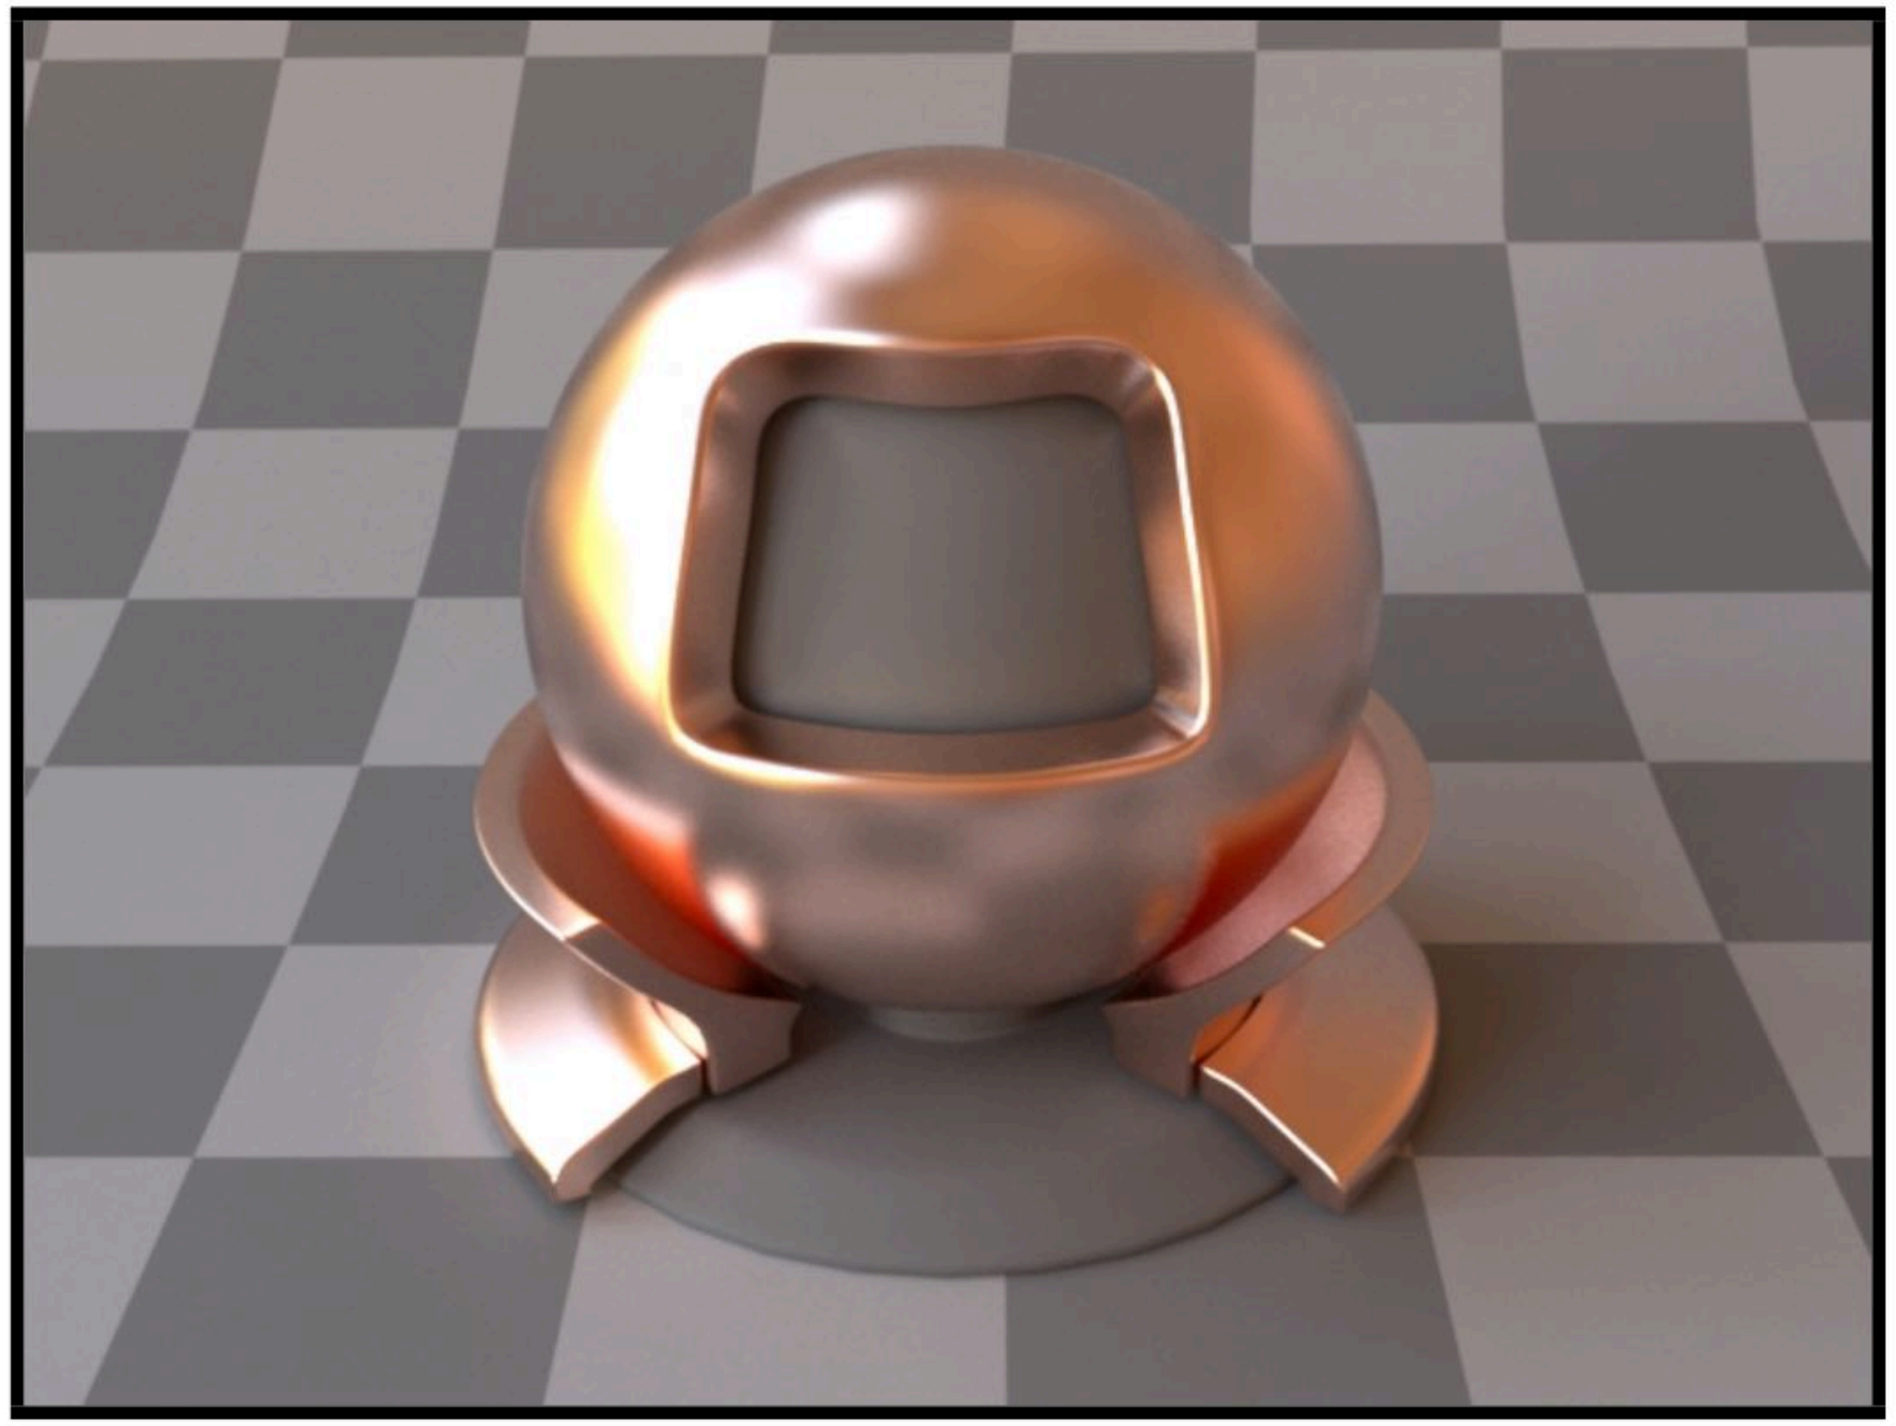

Copper

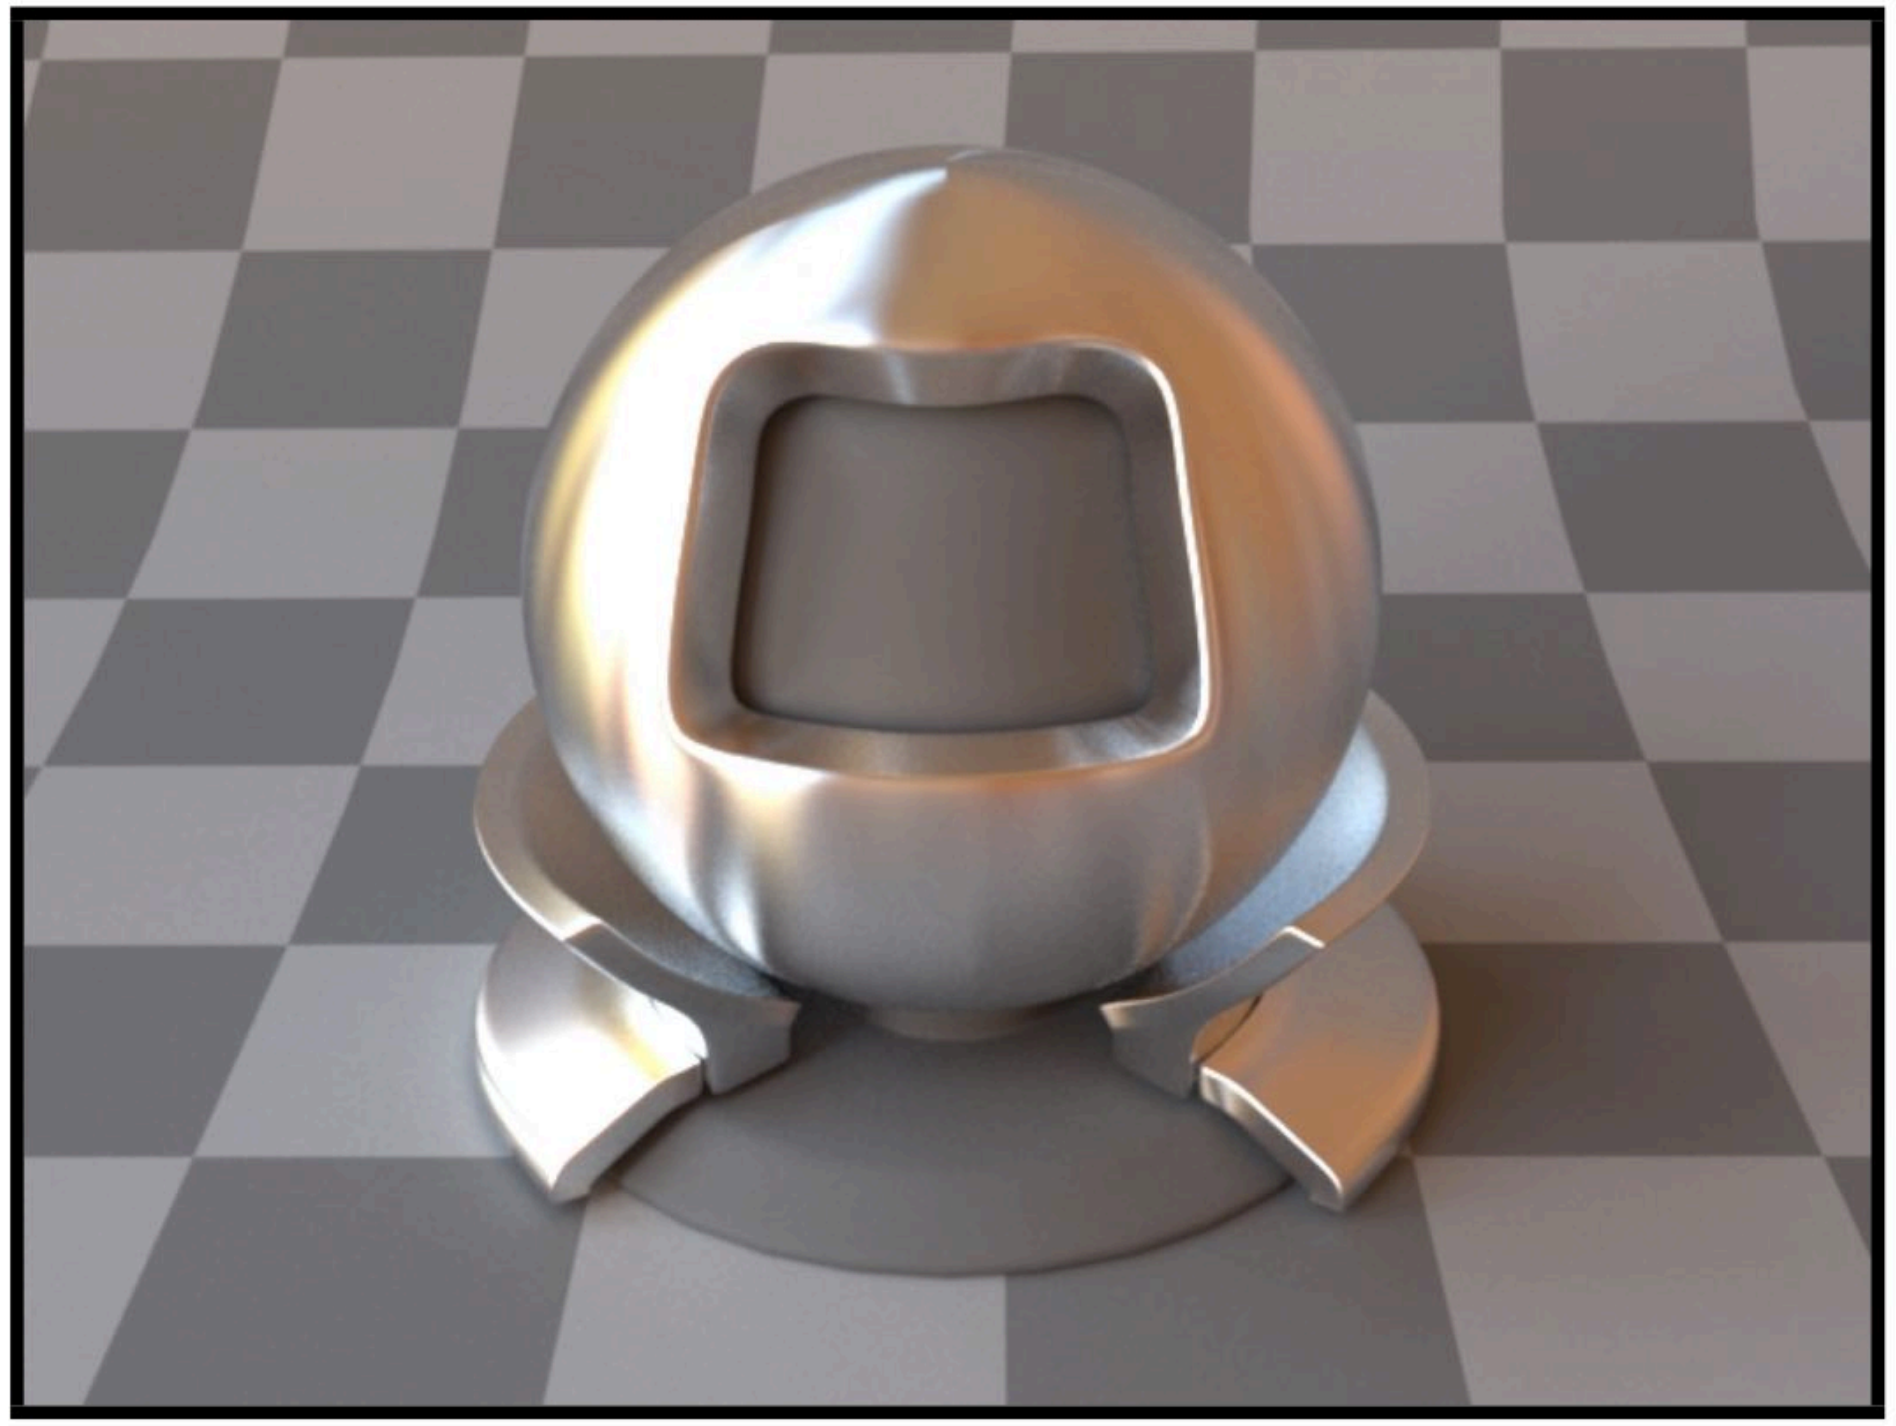

Aluminum

[Mitsuba renderer, Wenzel Jakob, 2010]

What is this material?

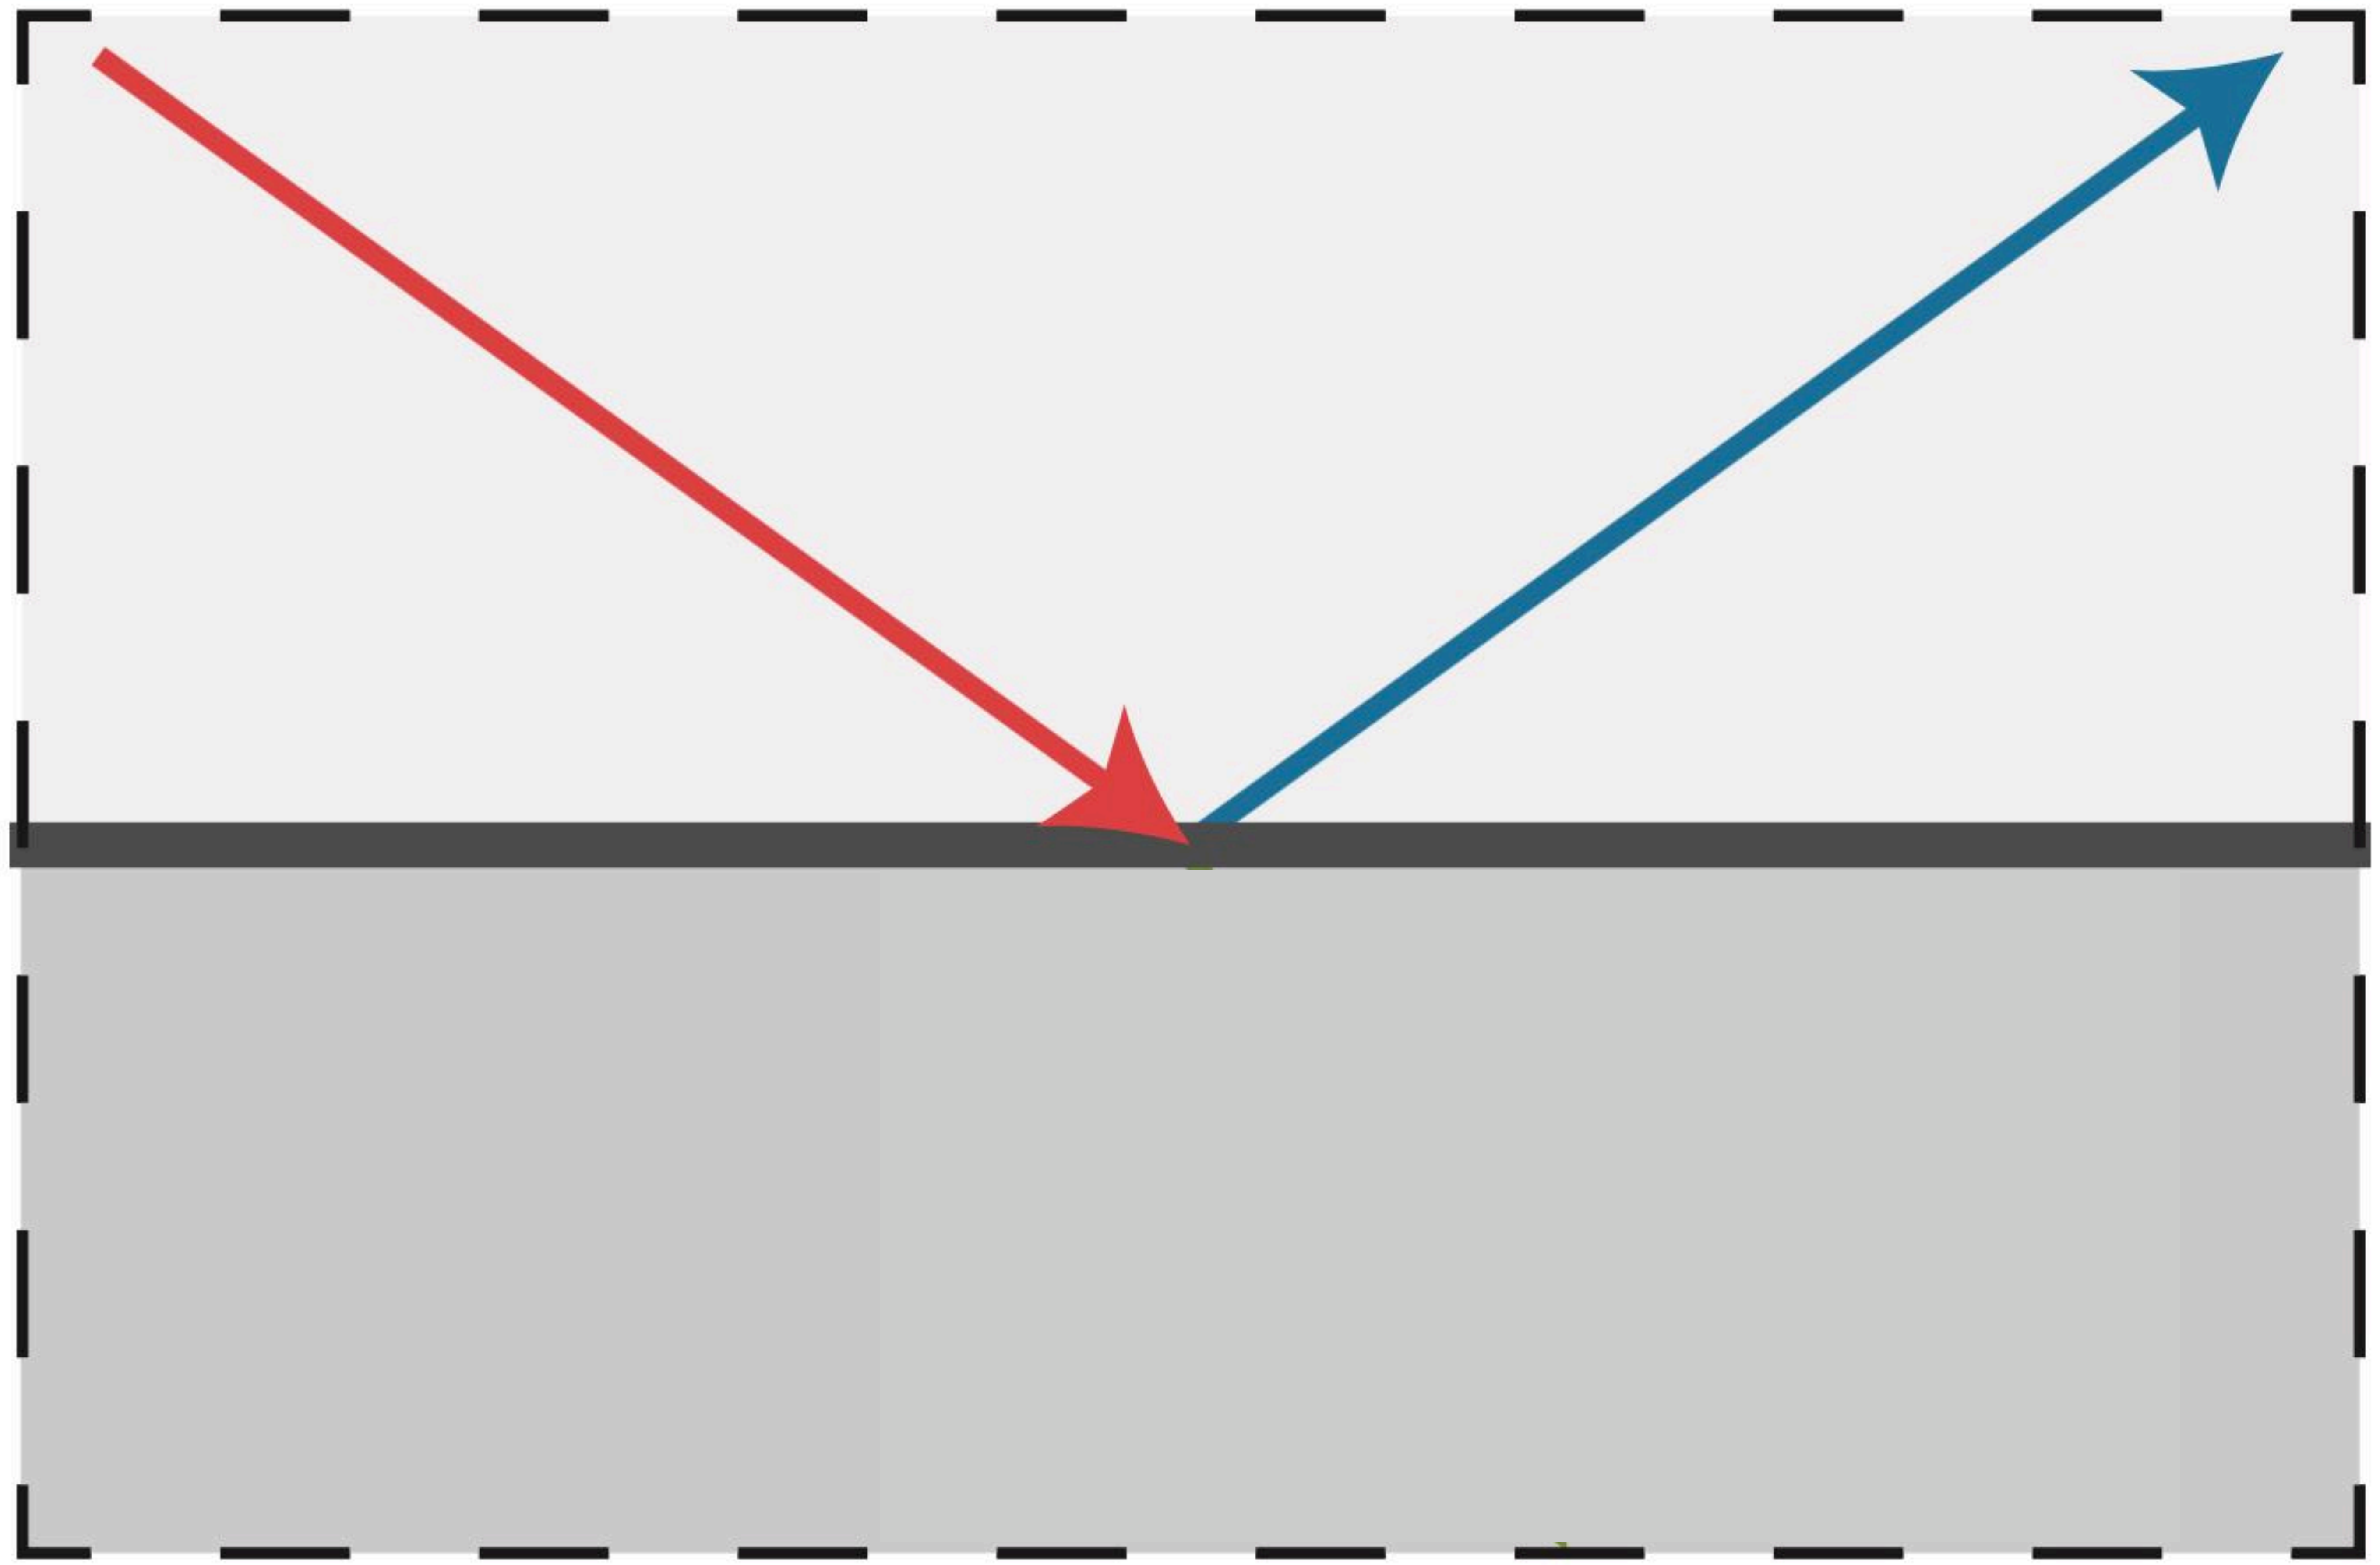

Perfect specular reflection

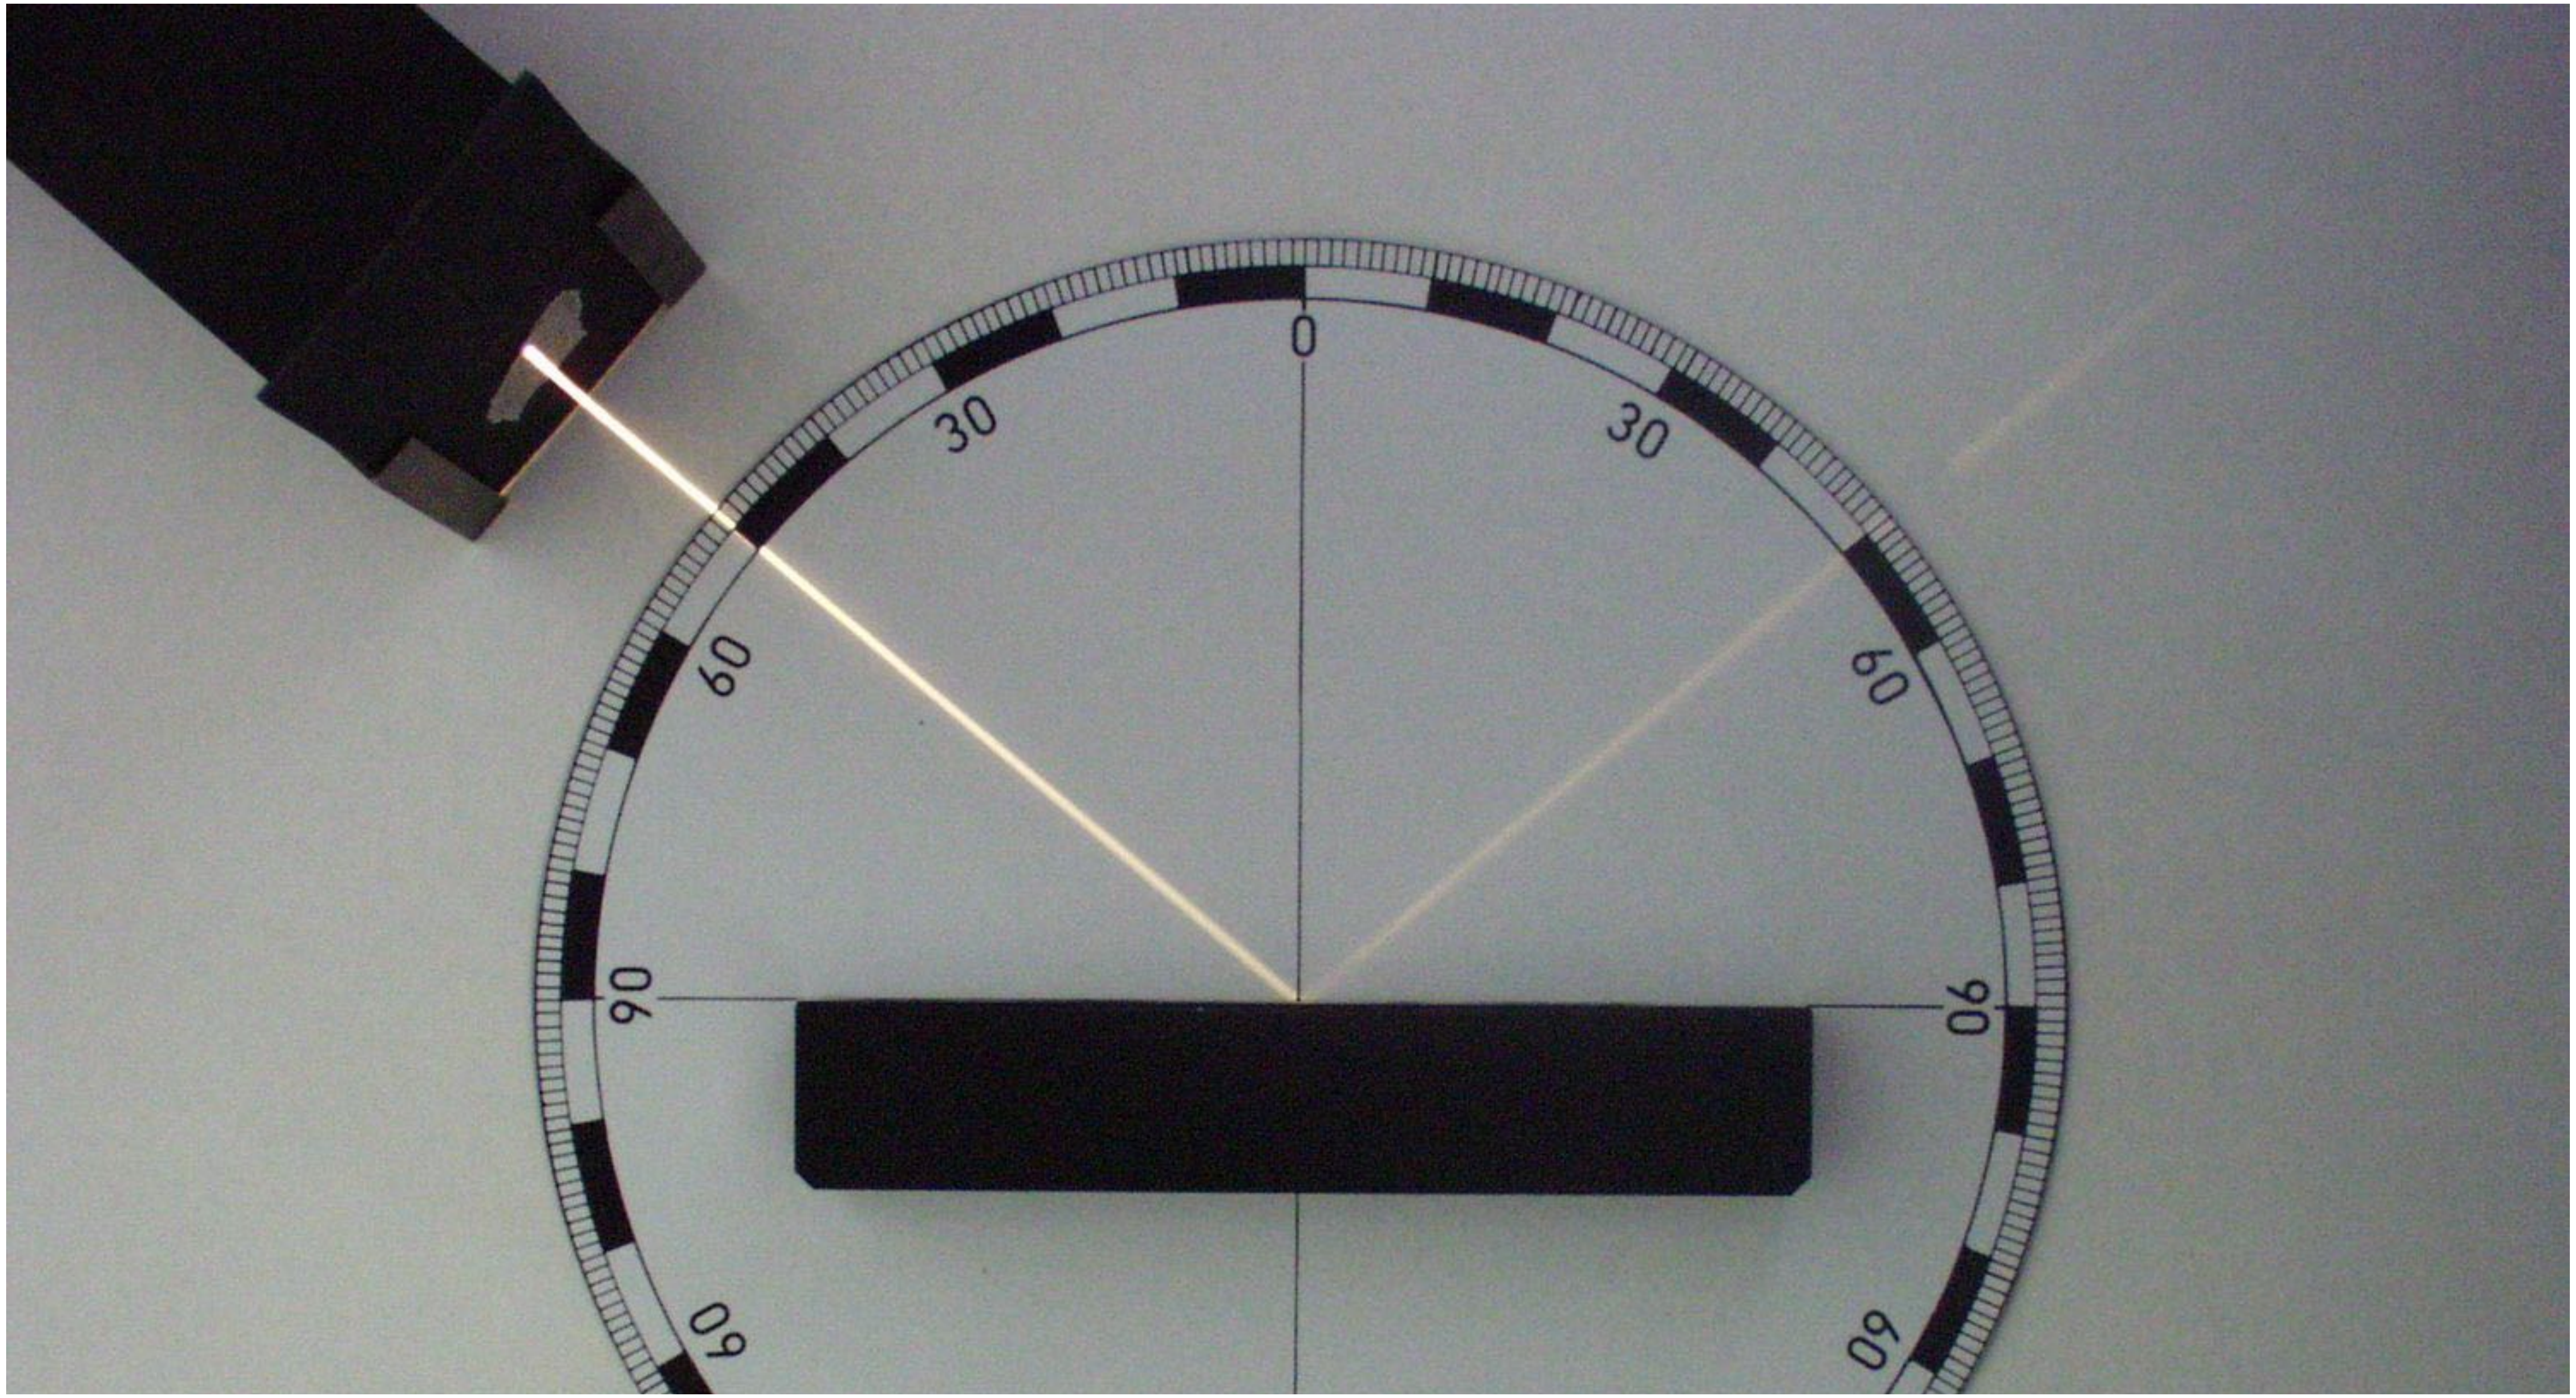

[Zátonyi Sándor]

Perfect specular reflection

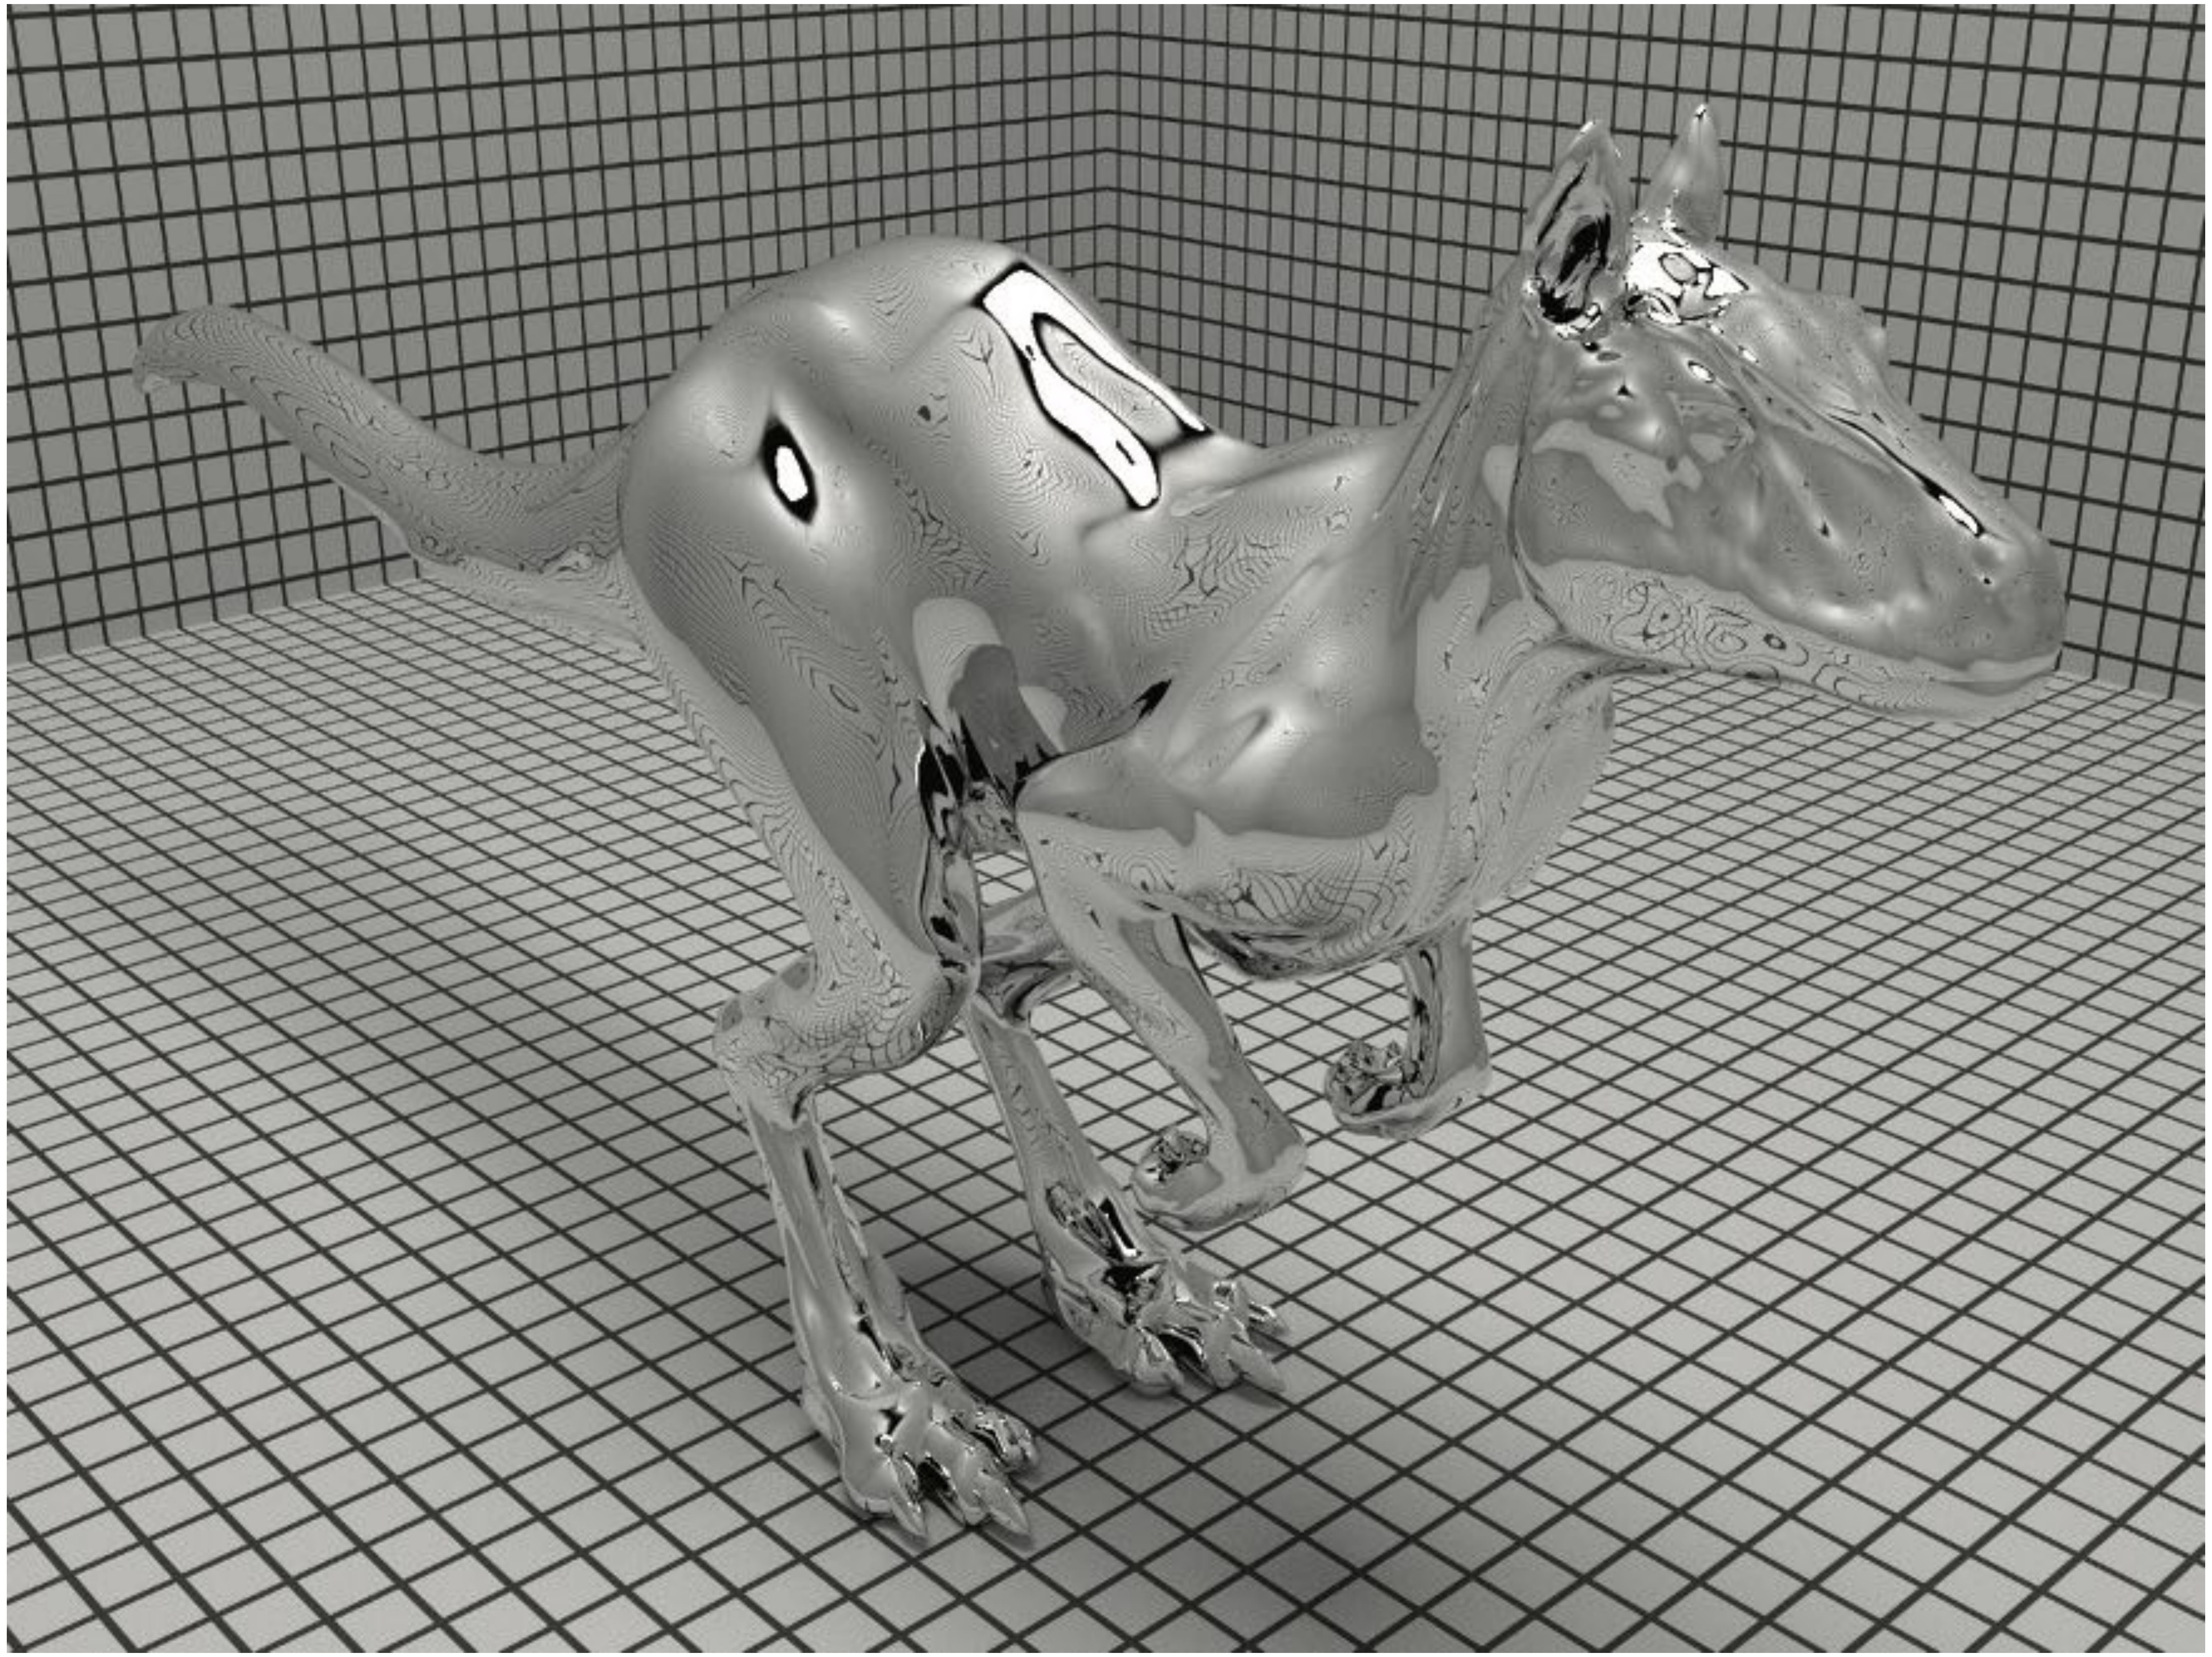

Image credit: PBRT

Calculating direction of specular reflection

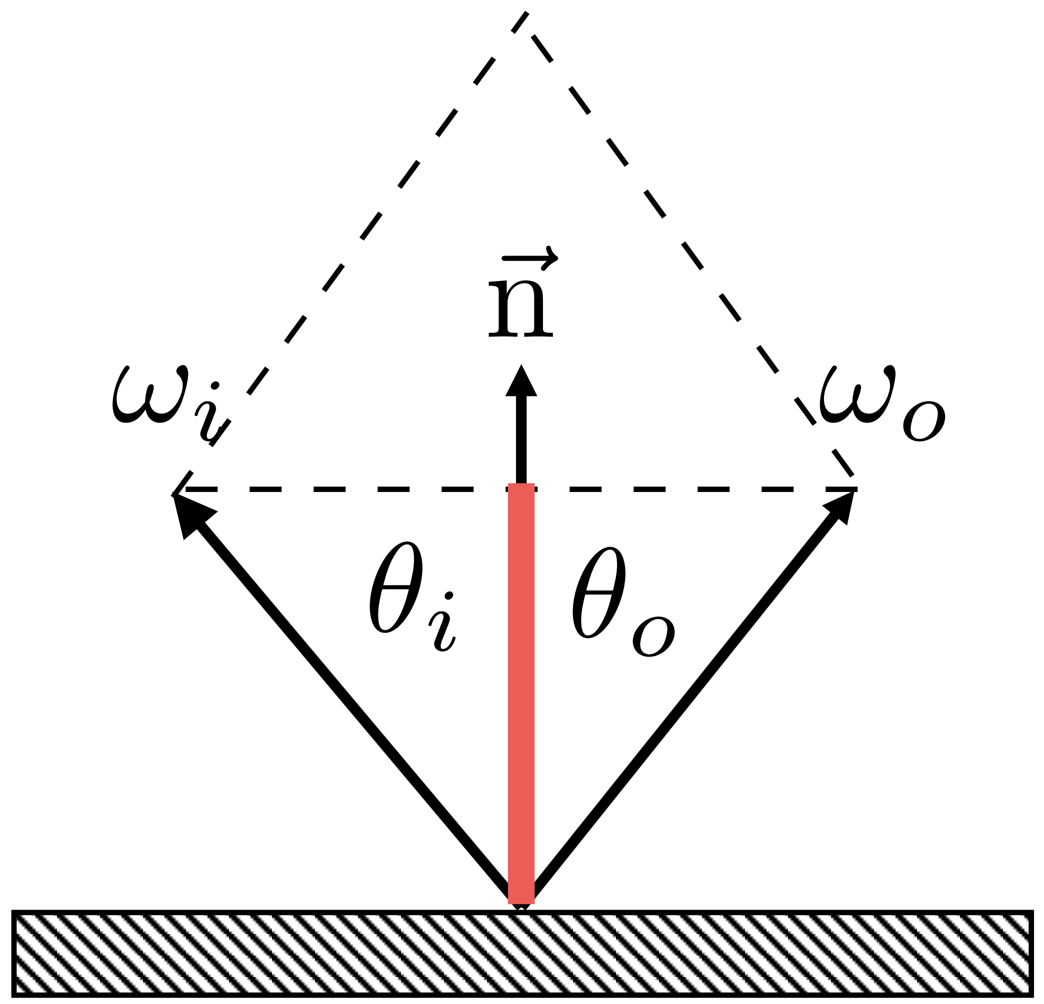

$$\theta = \theta_o = \theta_i$$

Top-down view
(looking straight down on surface)

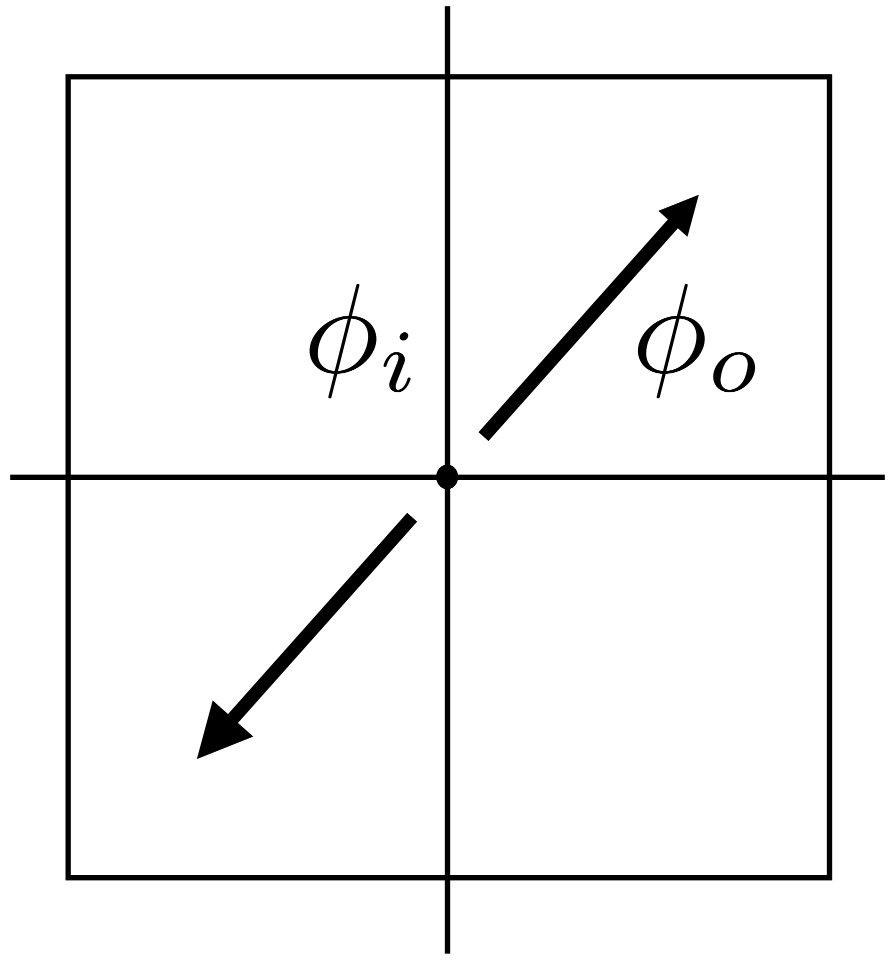

$$\phi_o = (\phi_i + \pi) \bmod 2\pi$$

$$\omega_o + \omega_i = 2\cos\theta\,\vec{n} = 2(\omega_i \cdot \vec{n})\vec{n}$$

$$\omega_o = -\omega_i + 2(\omega_i \cdot \vec{n})\vec{n}$$

Hemispherical incident radiance

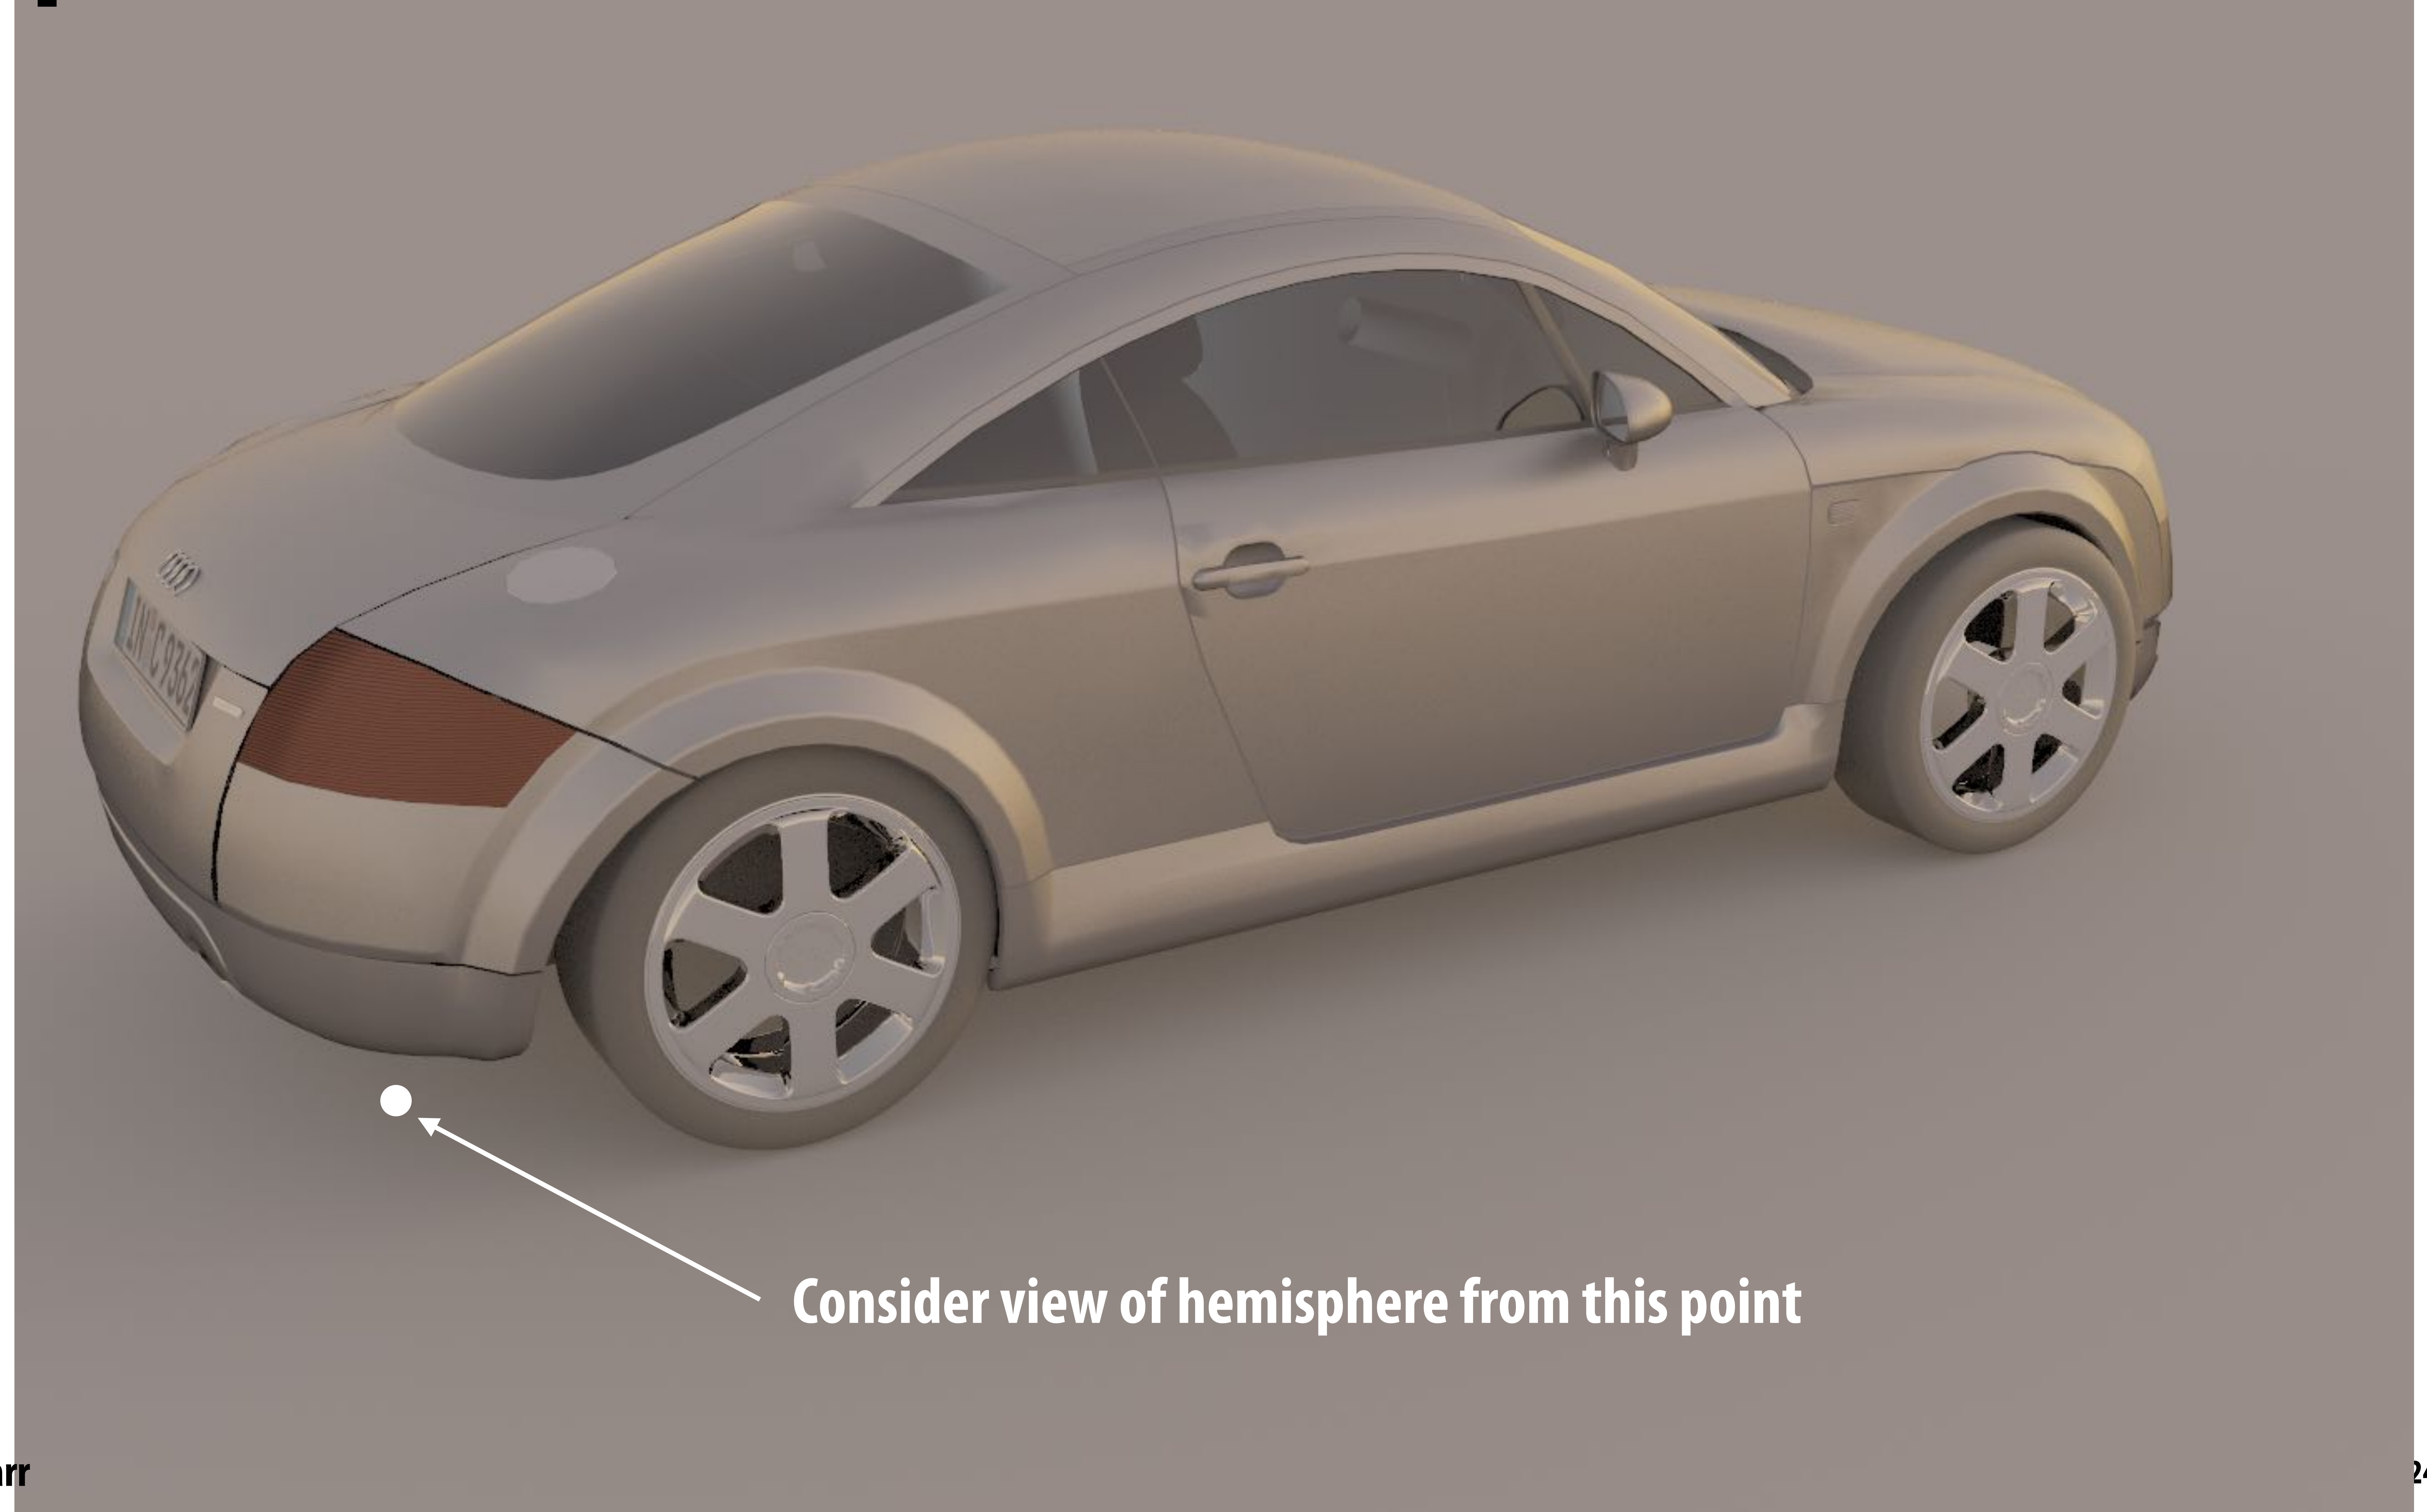

Image credit Matt Pharr

Hemispherical incident radiance

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

Diffuse reflection

Exitant radiance is the same in all directions

Incident radiance

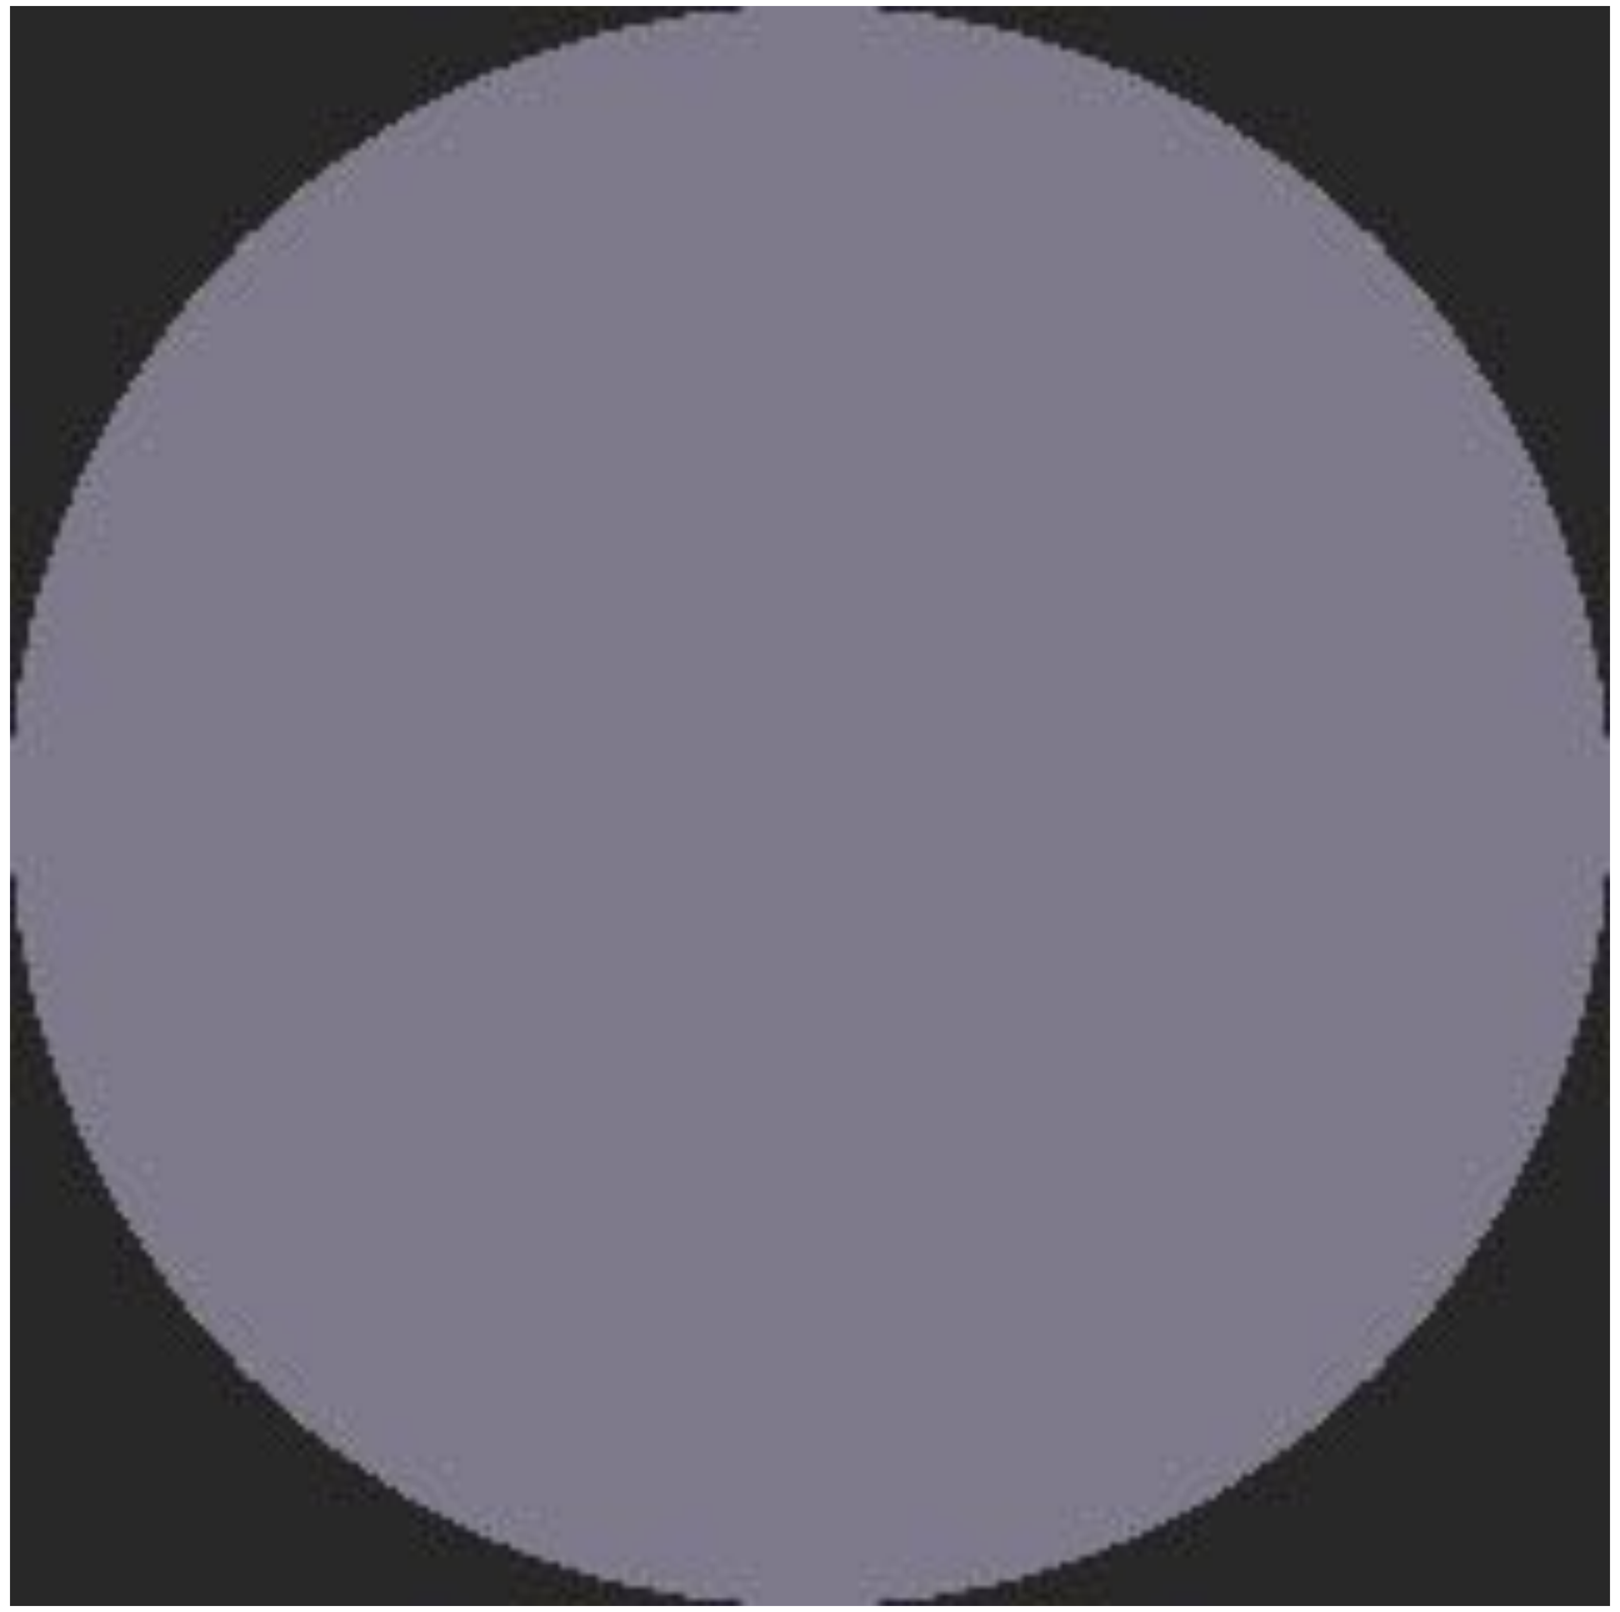

Exitant radiance

Image credit Matt Pharr

Ideal specular reflection

Incident radiance

Exitant radiance

Image credit Matt Pharr

How might you render a specular surface

- Compute direction from surface point *p* to camera = ω_o
- Given normal at *p*, compute reflection direction ω_i
- Light reflected in direction ω_o is light arriving from direction ω_i
- How do you measure light arriving from ω_i?

One idea…
look up amount in environment map!
(more on this later)

Pixel (x,y) stores radiance L from direction (ϕ, θ)

Plastic

Incident radiance

Exitant radiance

Image credit Matt Pharr

Copper

Incident radiance

Exitant radiance

Some basic reflection functions

- **Ideal specular**
  Perfect mirror

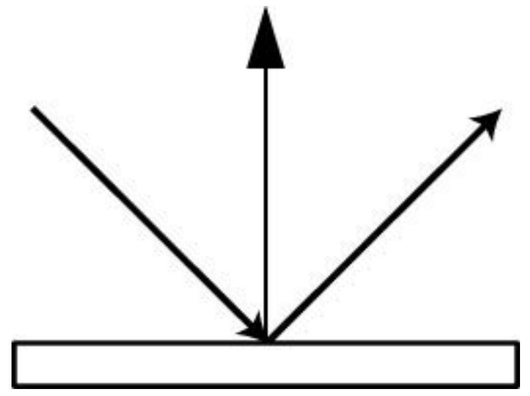
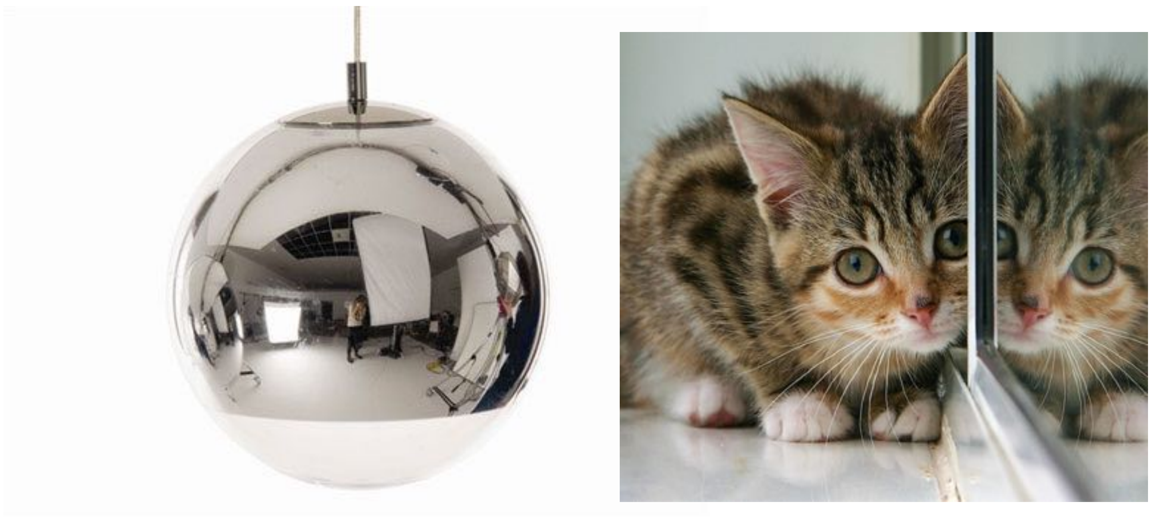

- **Ideal diffuse**
  Uniform reflection in all directions

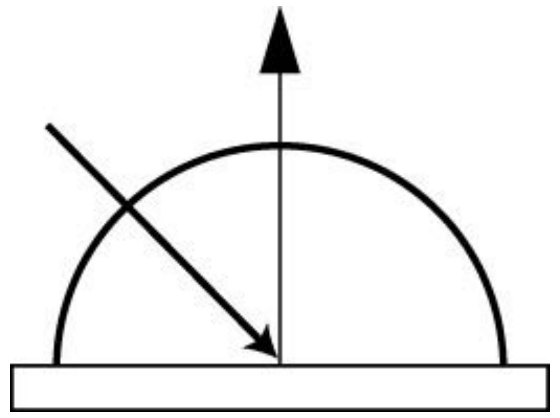


- **Glossy specular**
  Majority of light distributed in reflection direction

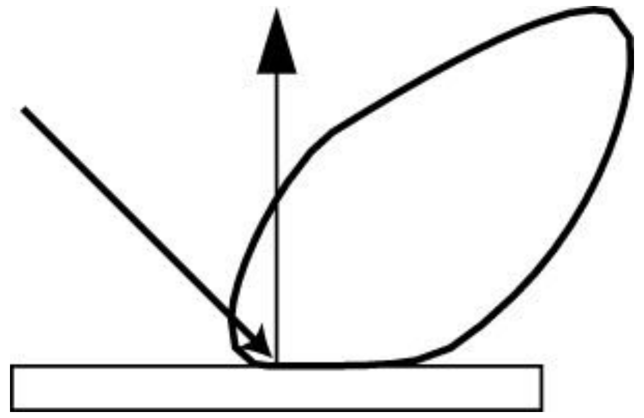


- **Retro-reflective**
  Reflects light back toward source

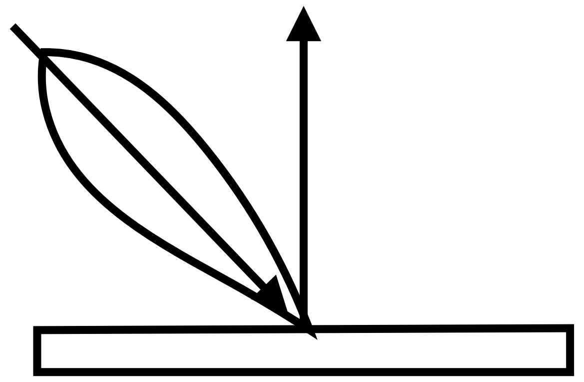
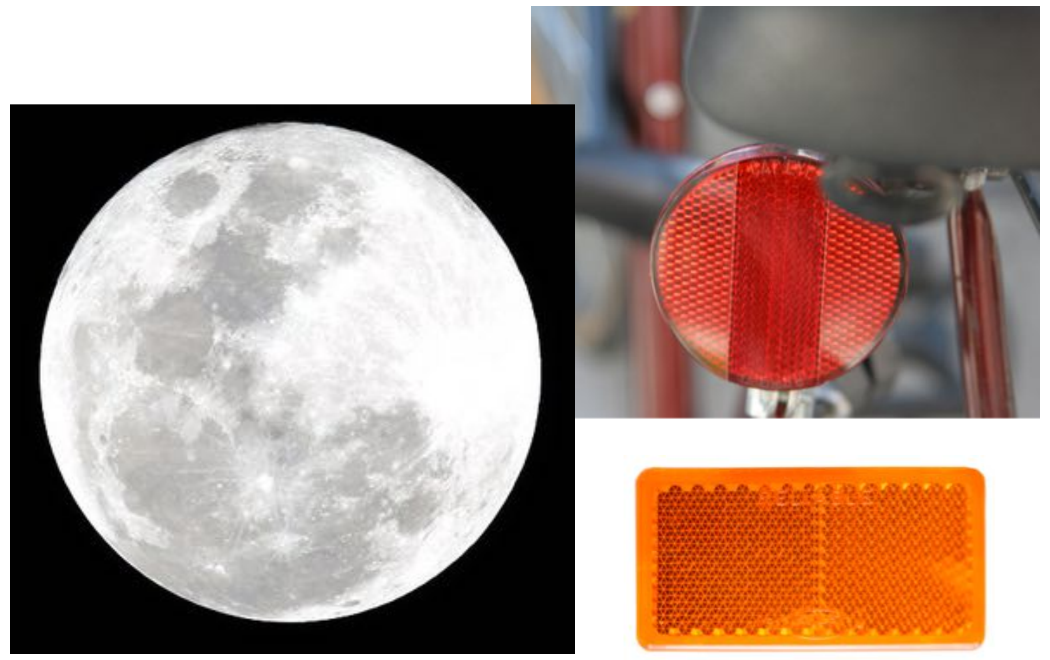

Diagrams illustrate how incoming light energy from a given direction is reflected in various directions.

More complex materials

Isotropic / anisotropic materials (BRDFs)

Key: directionality of underlying surface

Isotropic

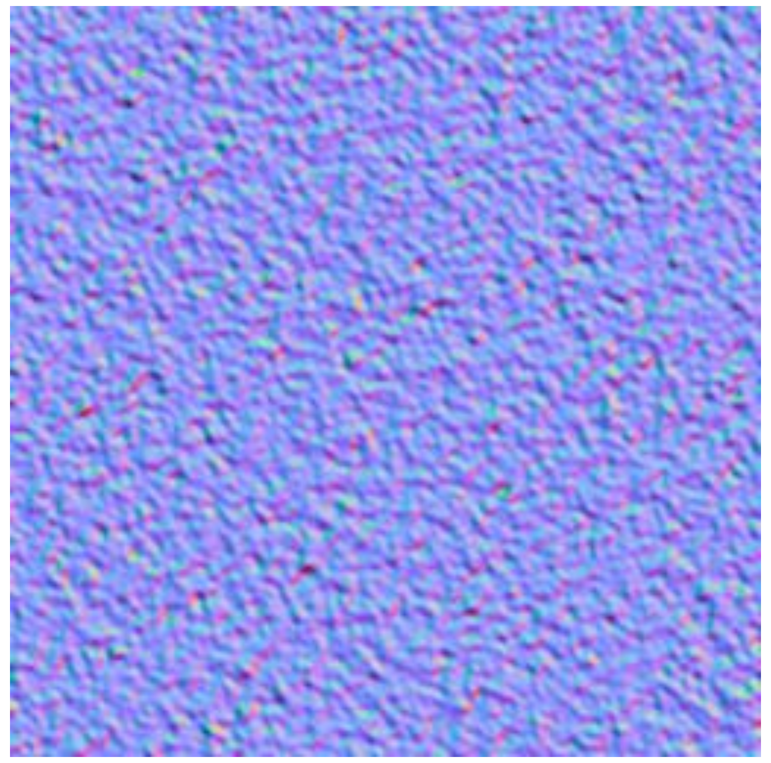

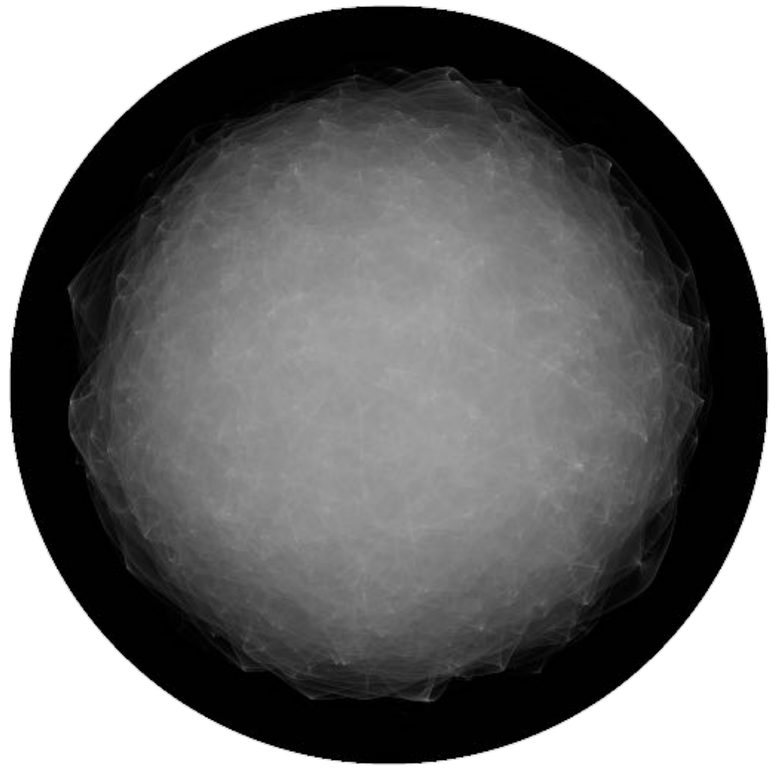

Anisotropic

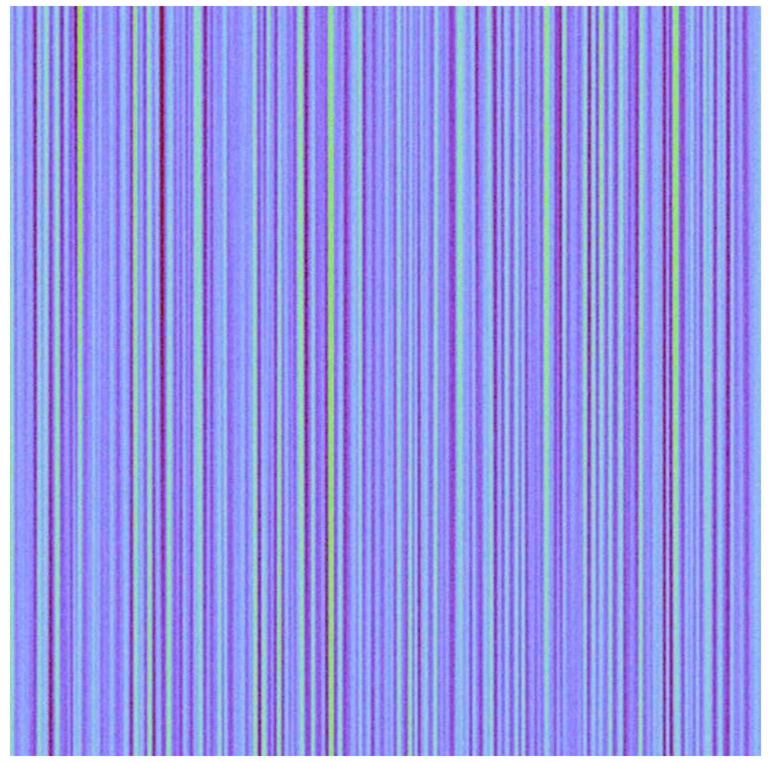

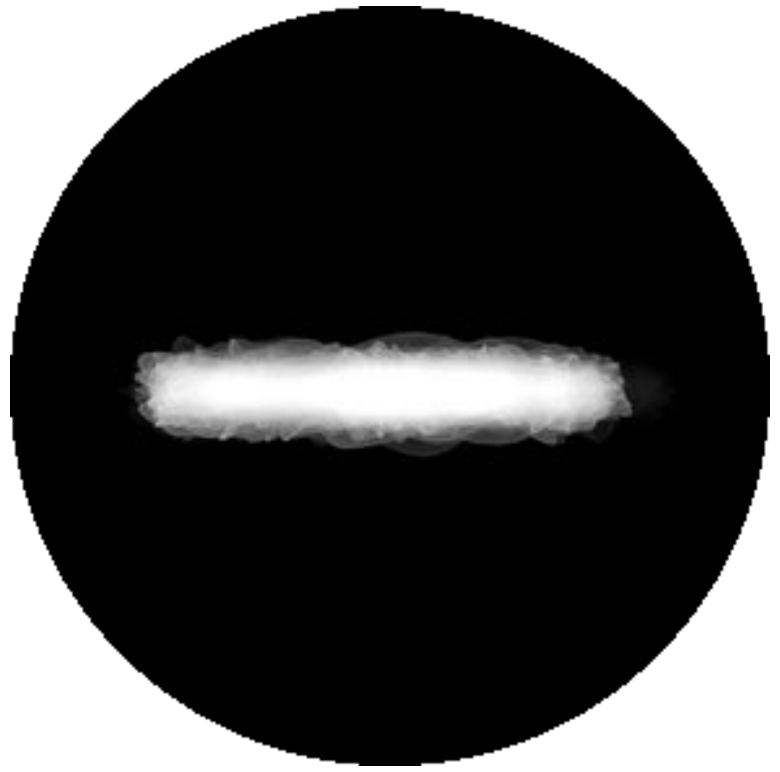

Surface (normals)

BRDF (fix wi, vary wo)

Anisotropic BRDFs

Reflection depends on azimuthal angle ϕ

$$f_r(\theta_i, \phi_i; \theta_r, \phi_r) \neq f_r(\theta_i, \theta_r, \phi_r - \phi_i)$$

Results from oriented microstructure of surface, e.g., brushed metal

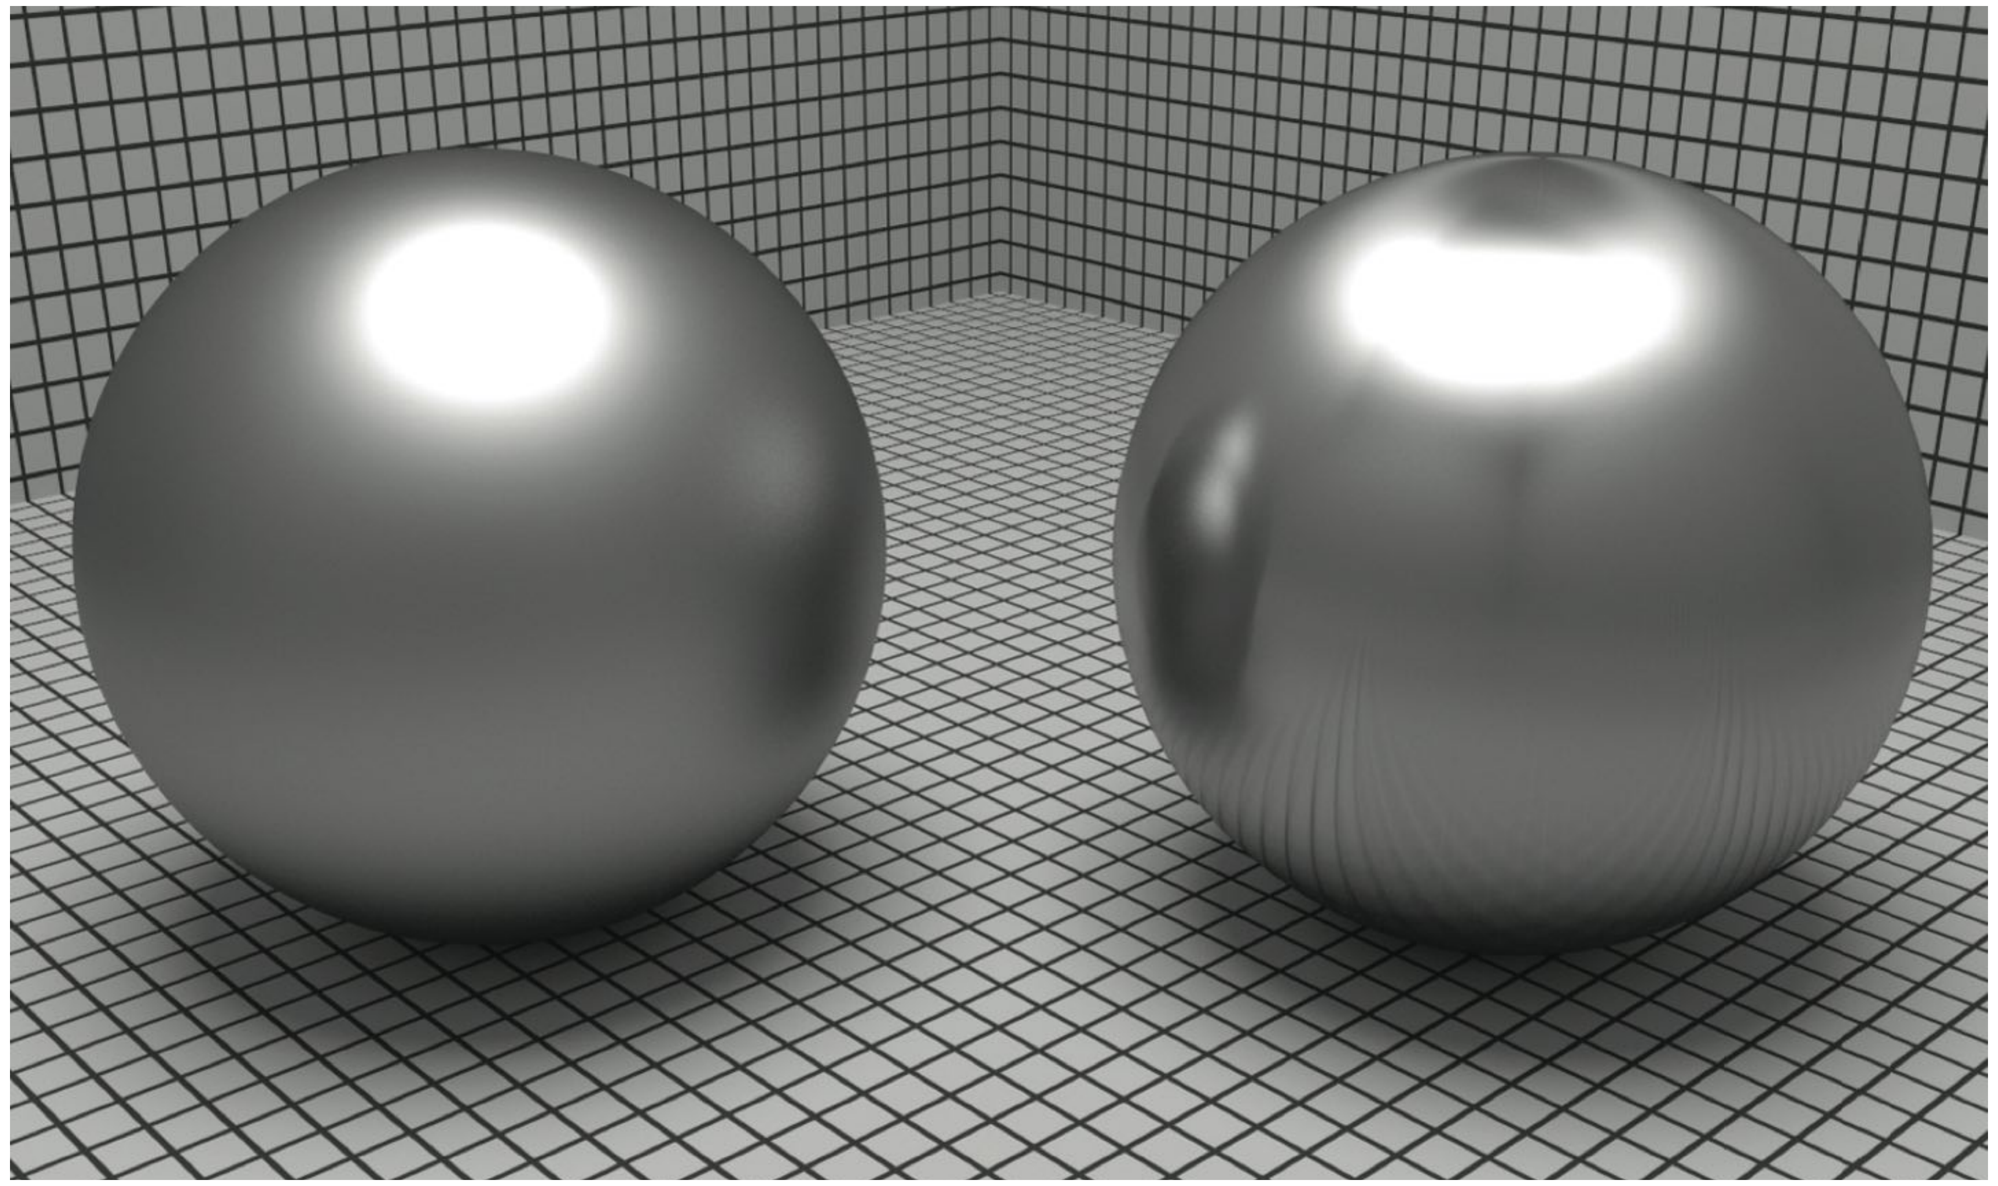

Anisotropic BRDF: Nylon

[Westin et al. 1992]

Anisotropic BRDF: Velvet

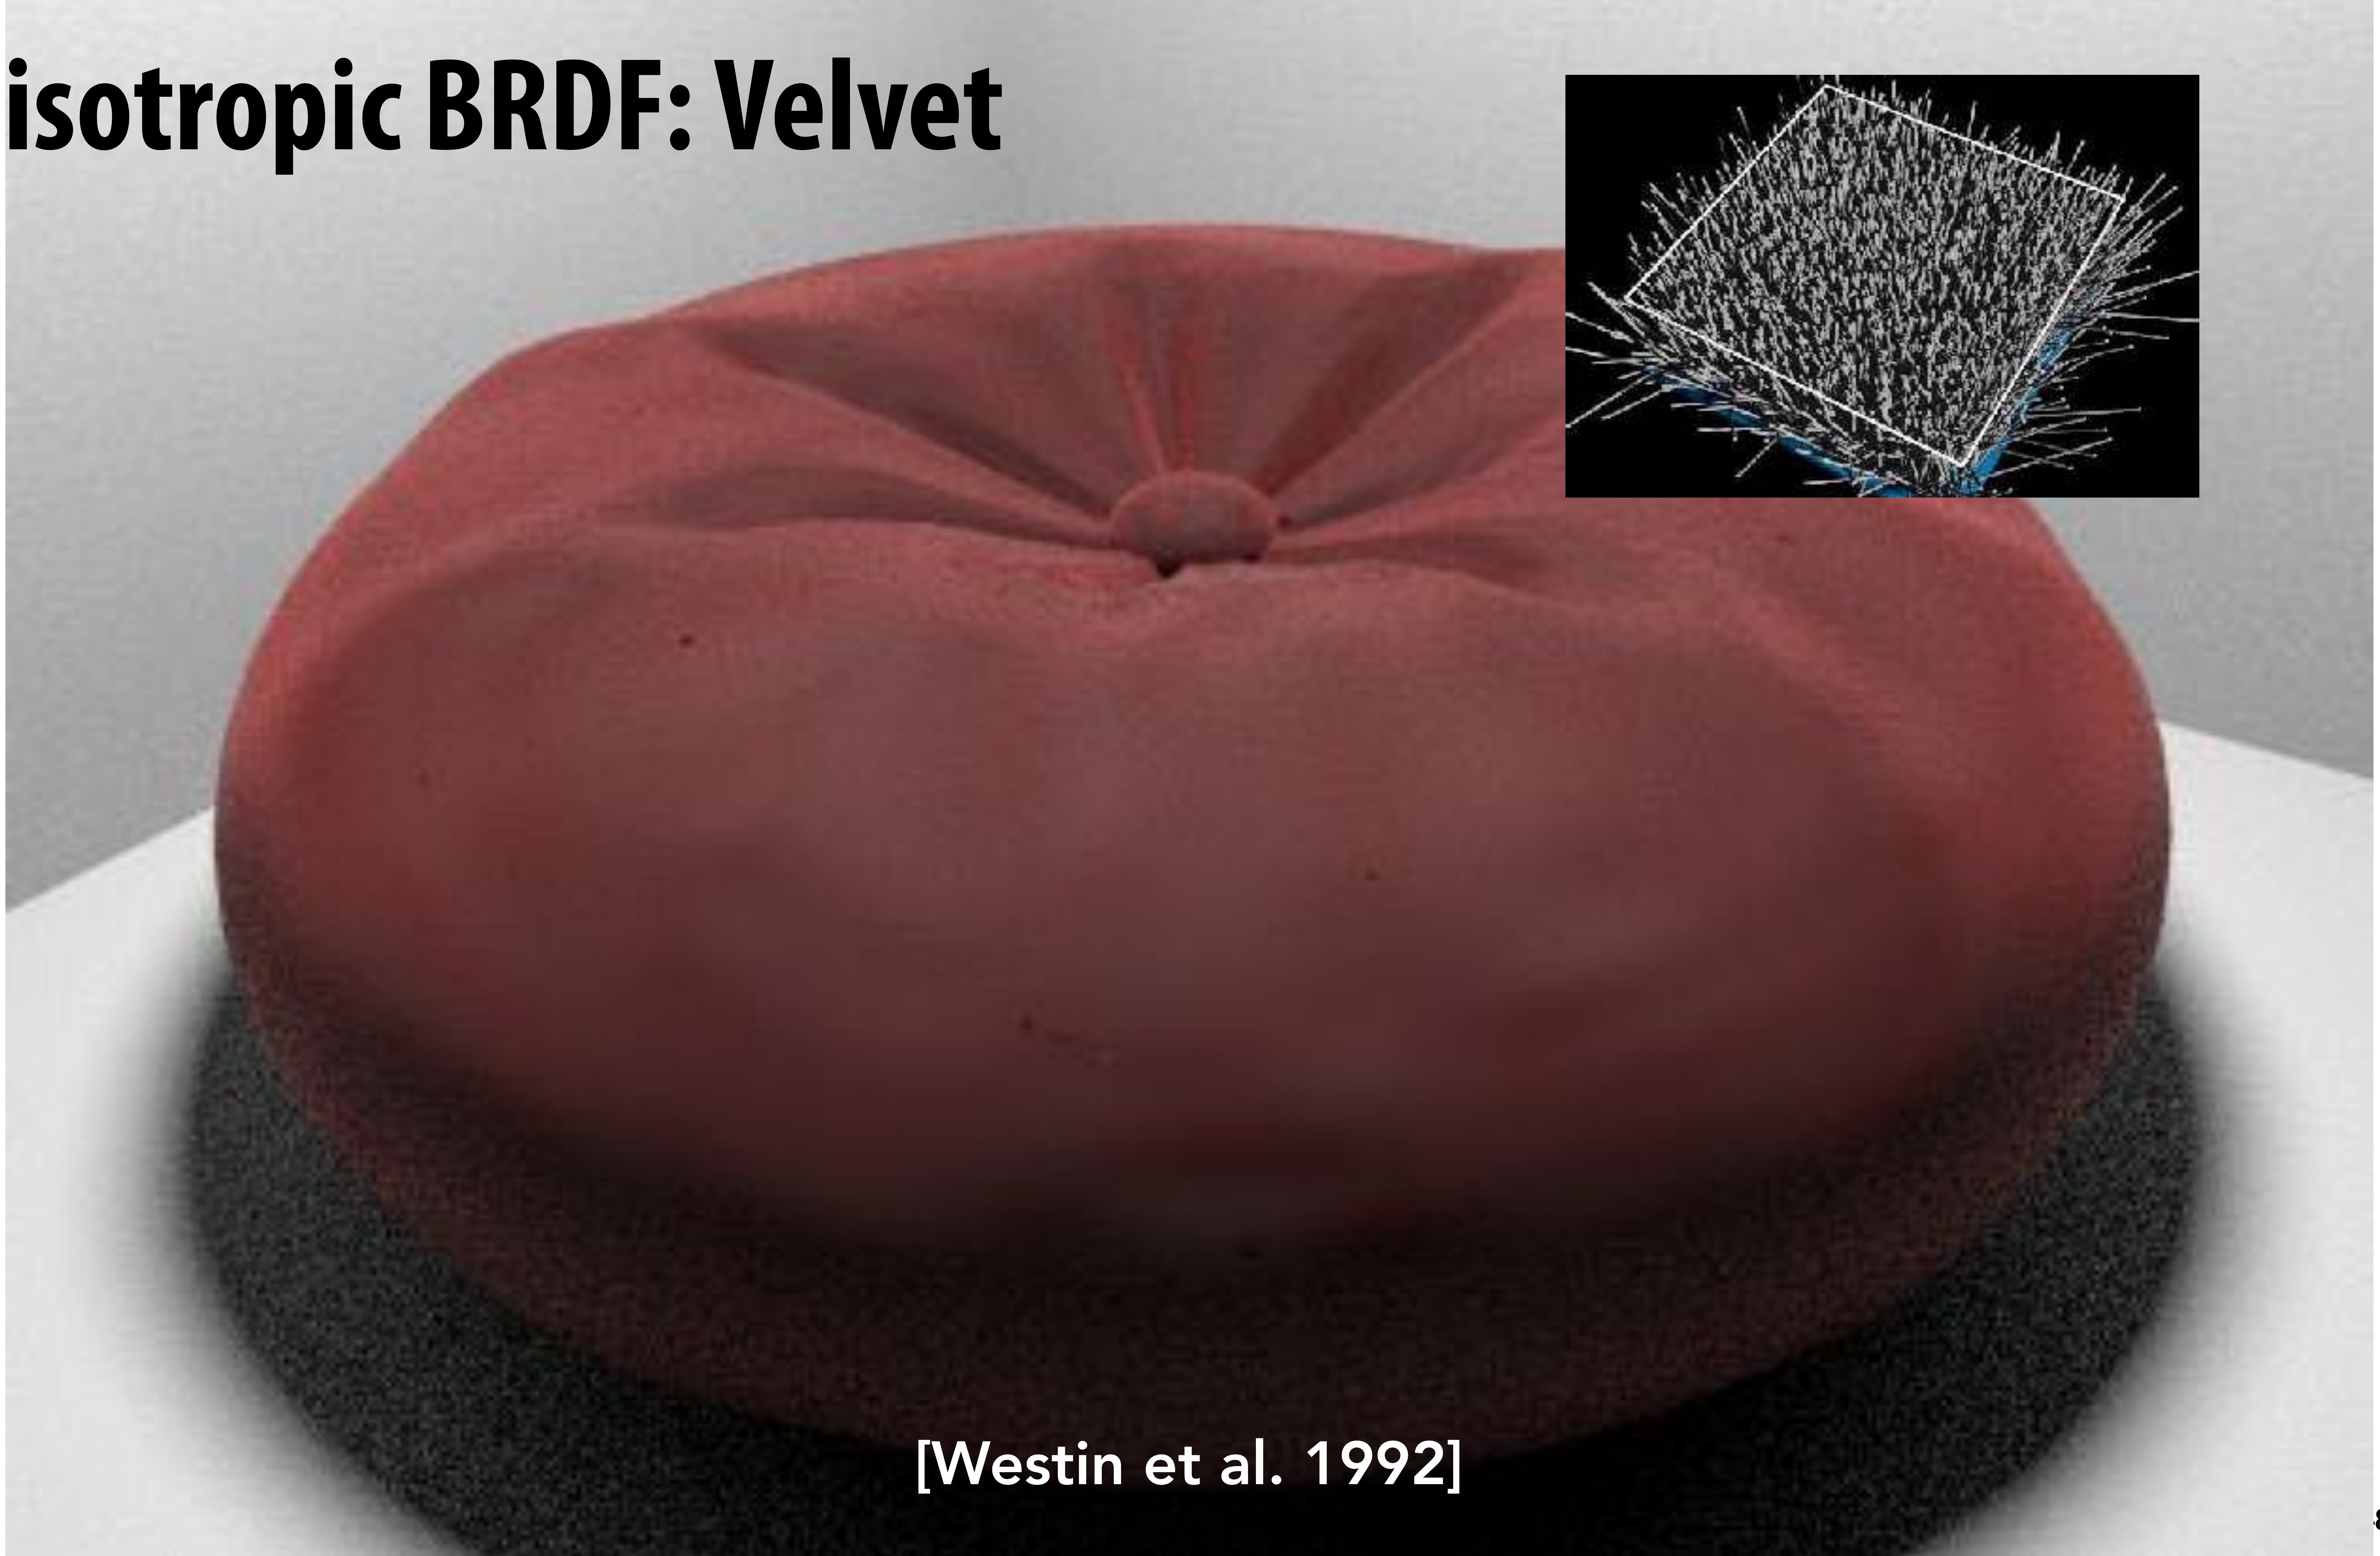

Anisotropic BRDF: Velvet

[https://www.youtube.com/watch?v=2hjoW8TYTd4]

What is this material?

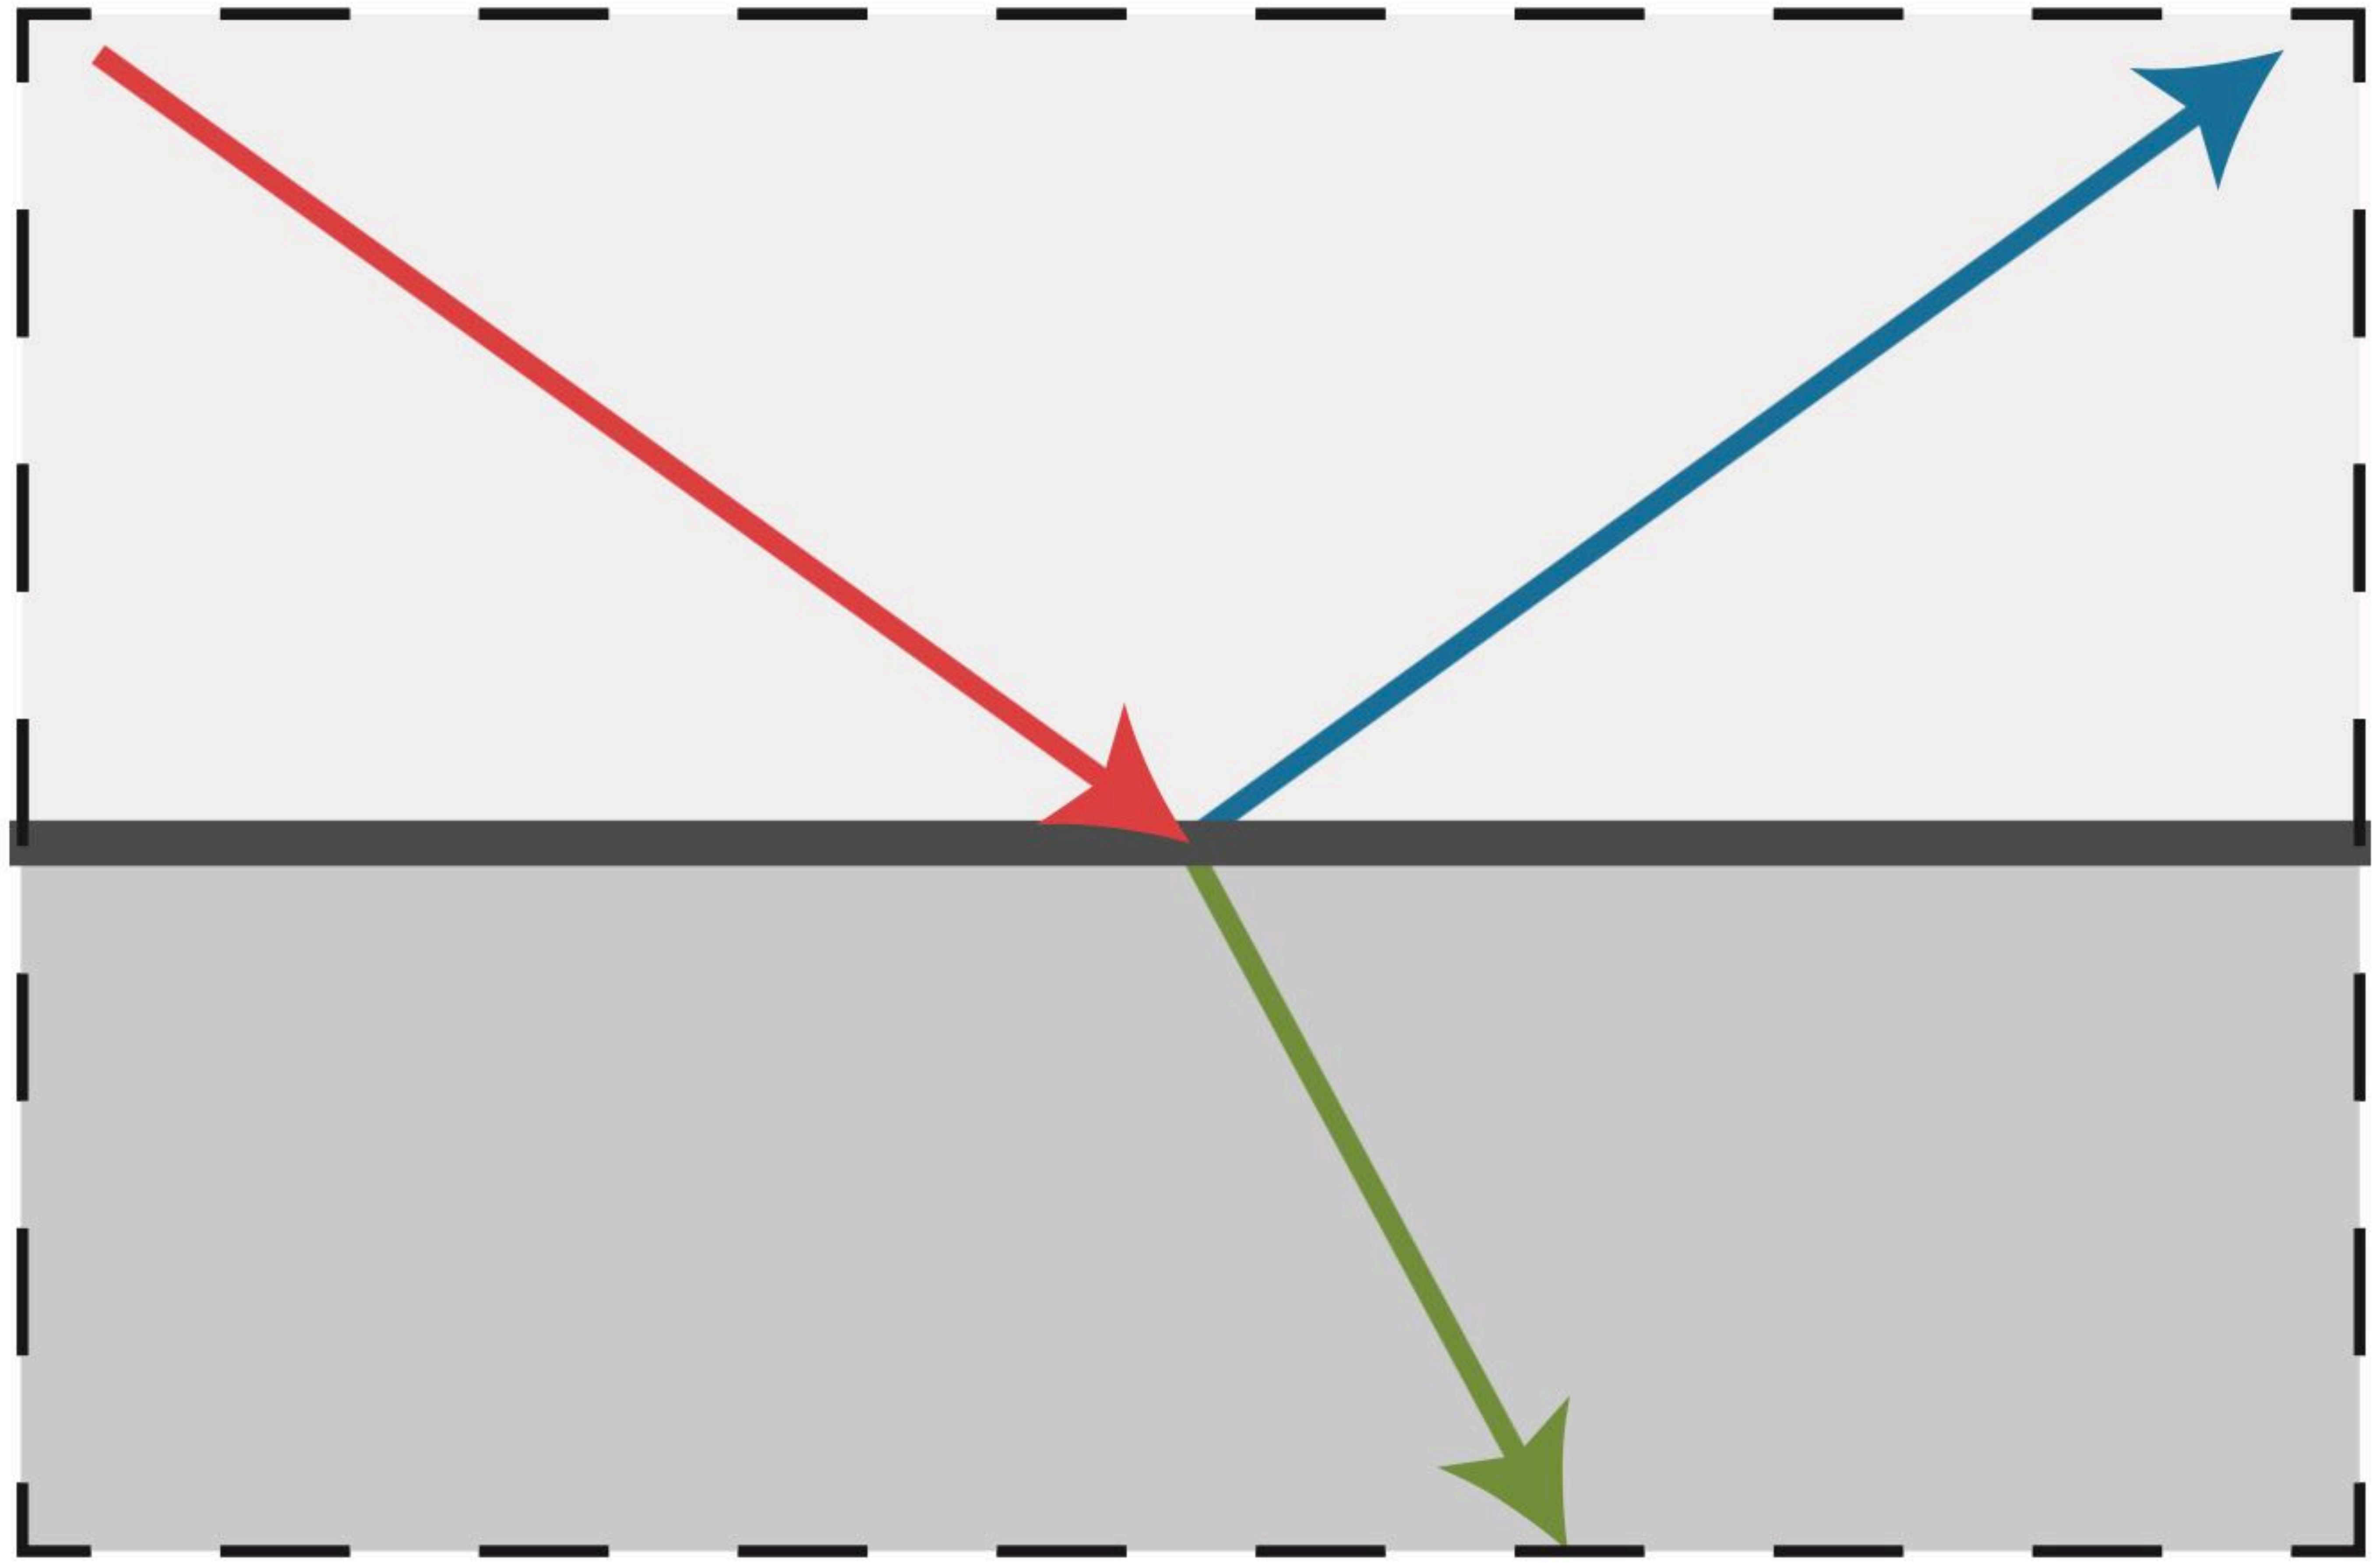

Ideal reflective / refractive material (BxDF *)

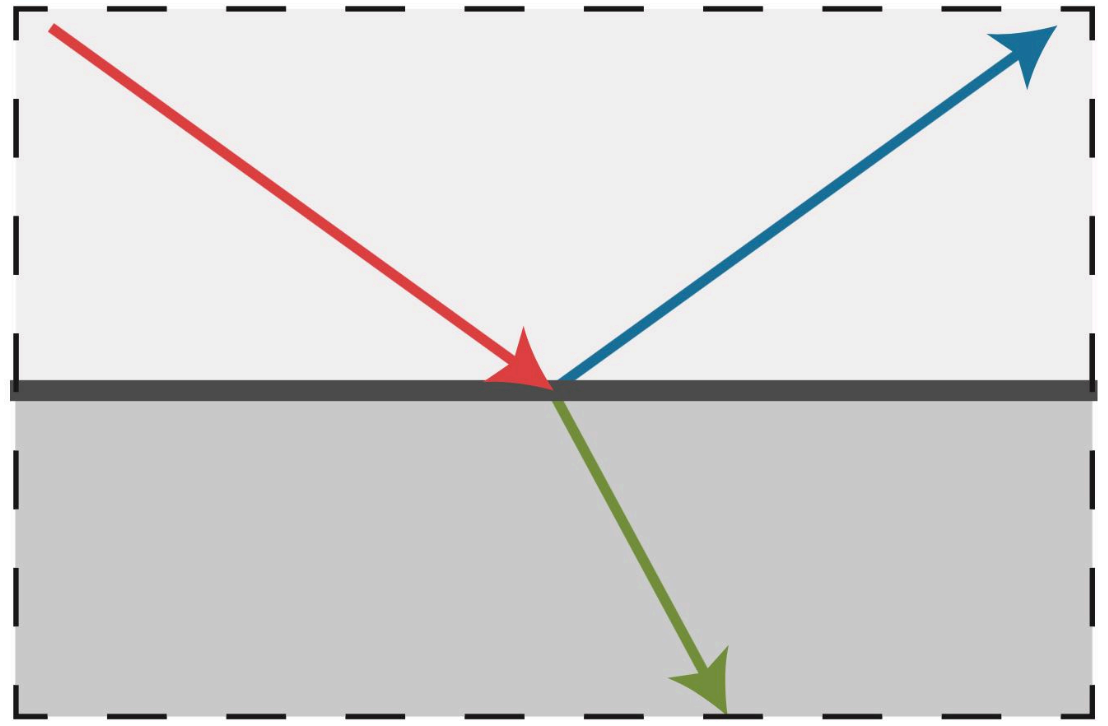

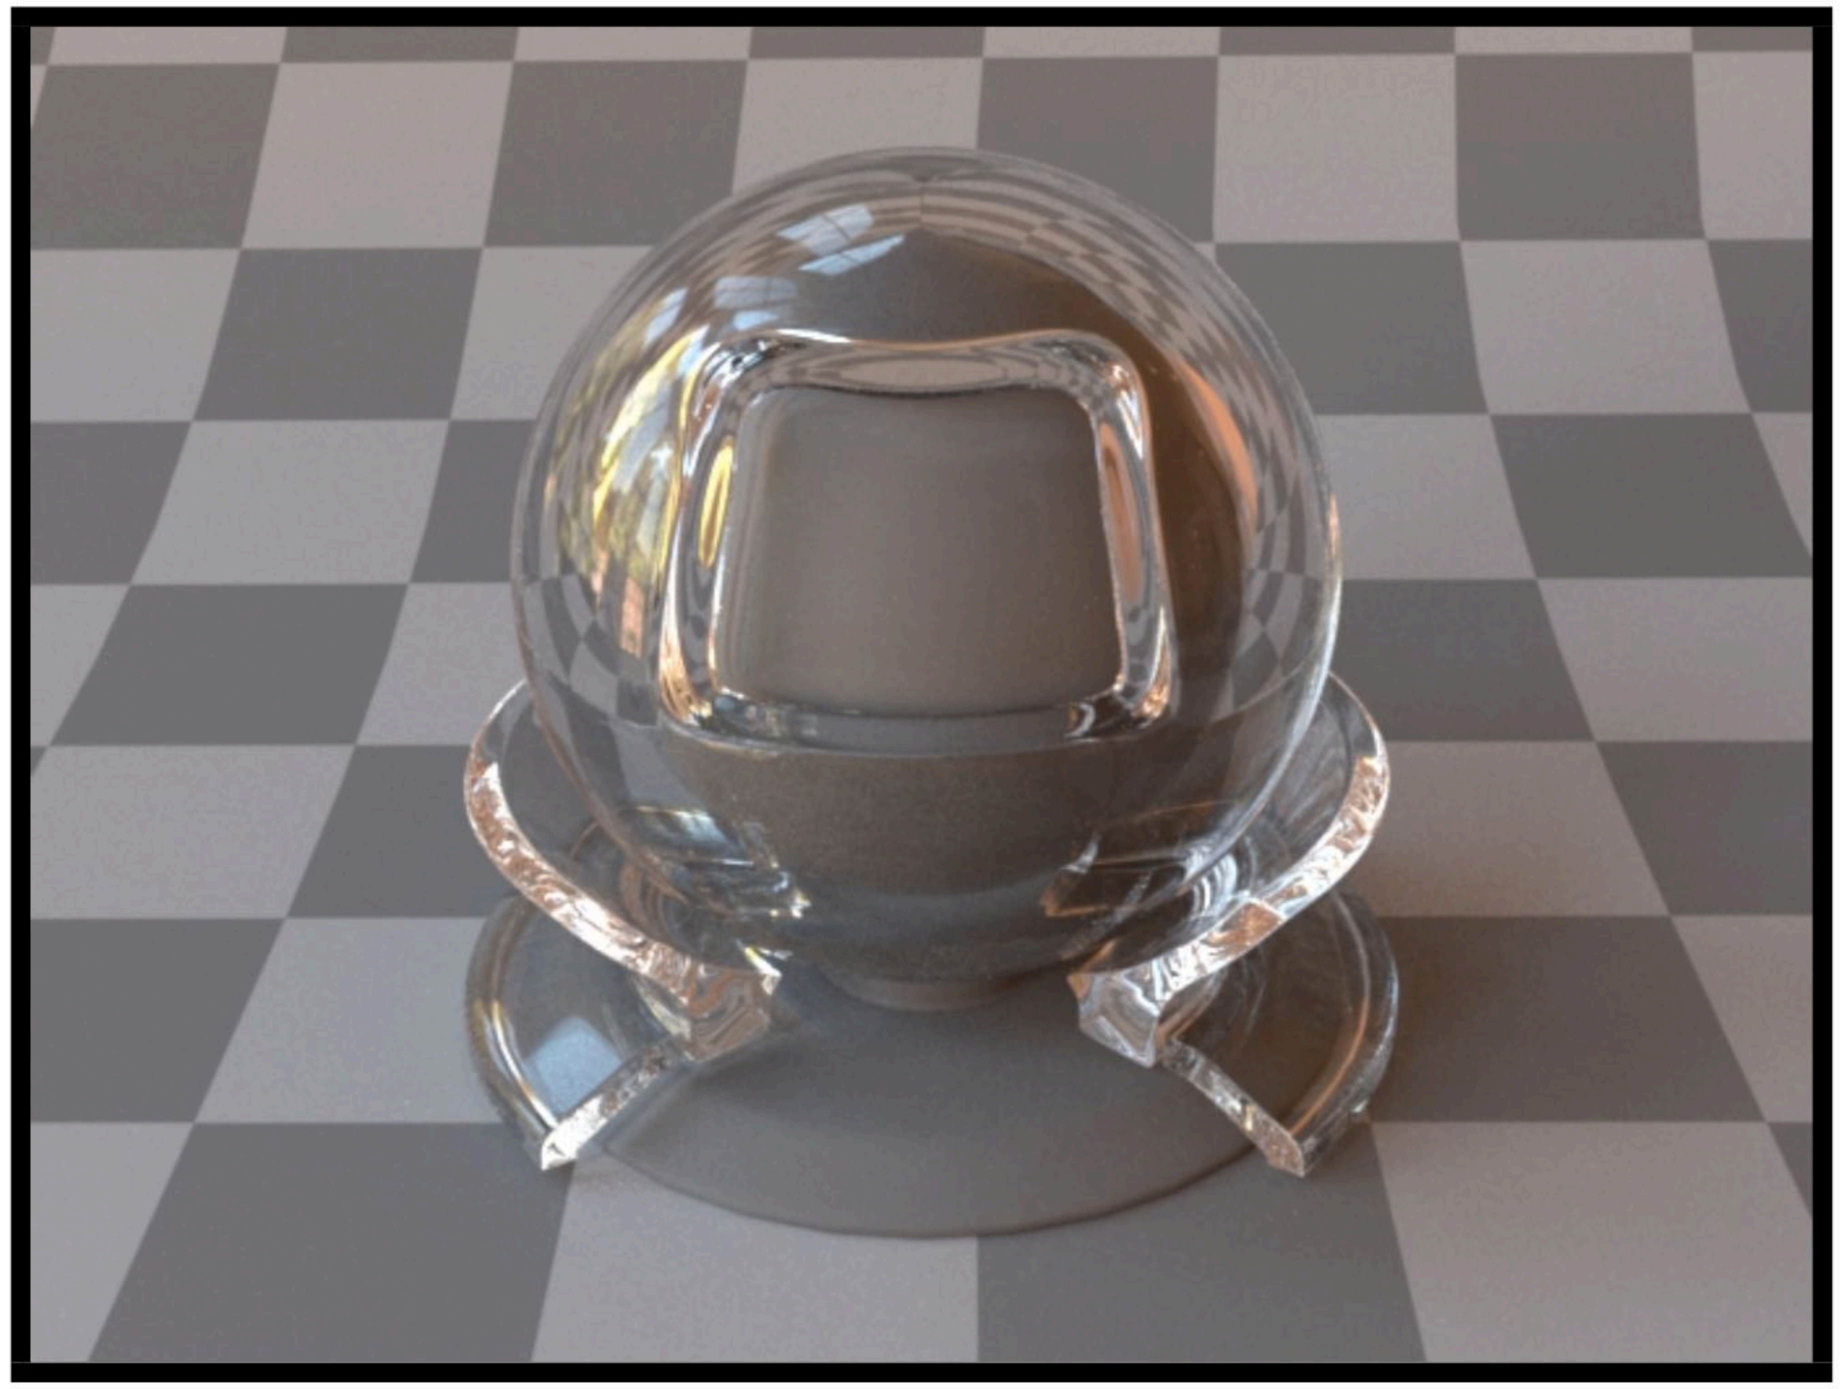

Air <-> water interface

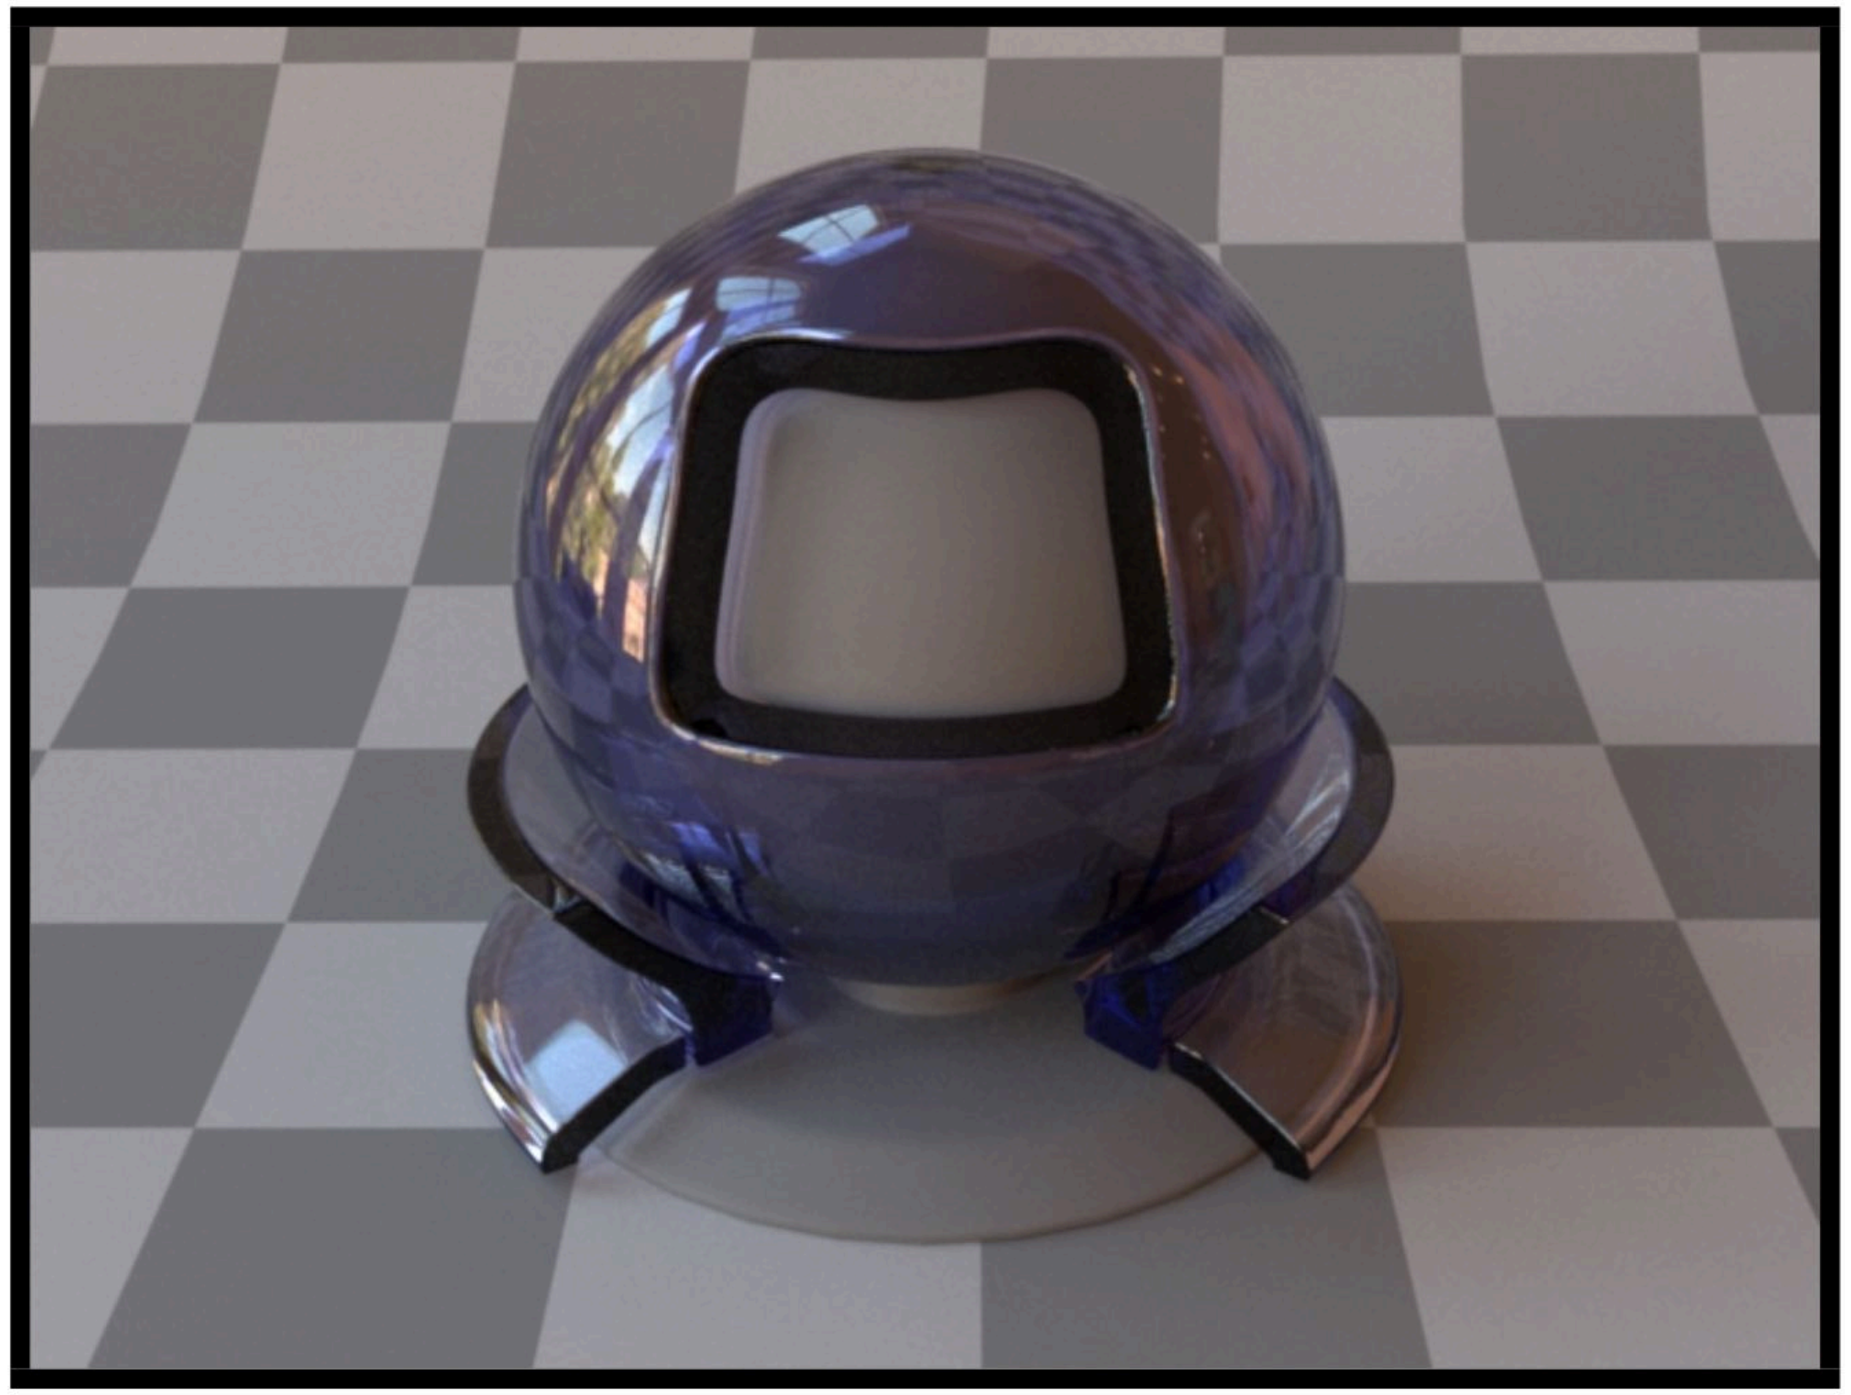

Air <-> glass interface
(with absorption)

* X stands in for reflectance "r", scattering "s", or transmission "t", etc.

[Mitsuba renderer, Wenzel Jakob, 2010]

Transmission

In addition to reflecting off surface, light may be transmitted through the surface.

Light refracts when it enters a new medium.

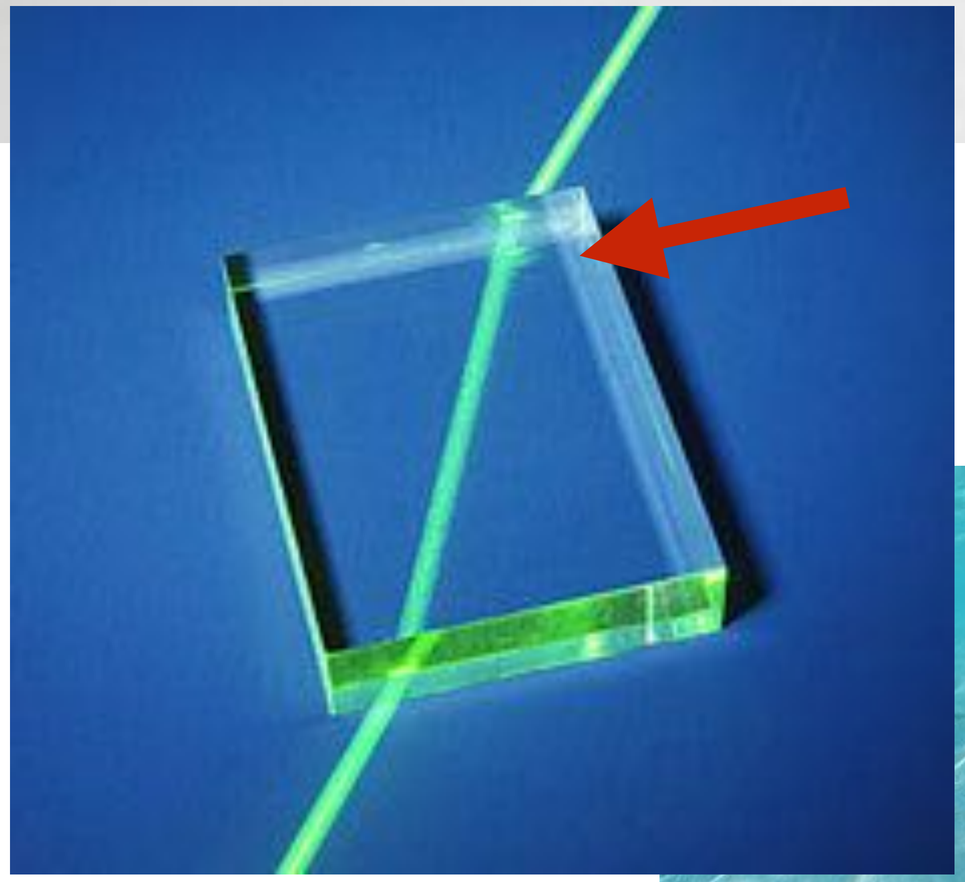

Snell's Law

Transmitted angle depends on index of refraction of medium incident ray is in and index of refraction of medium light is entering.

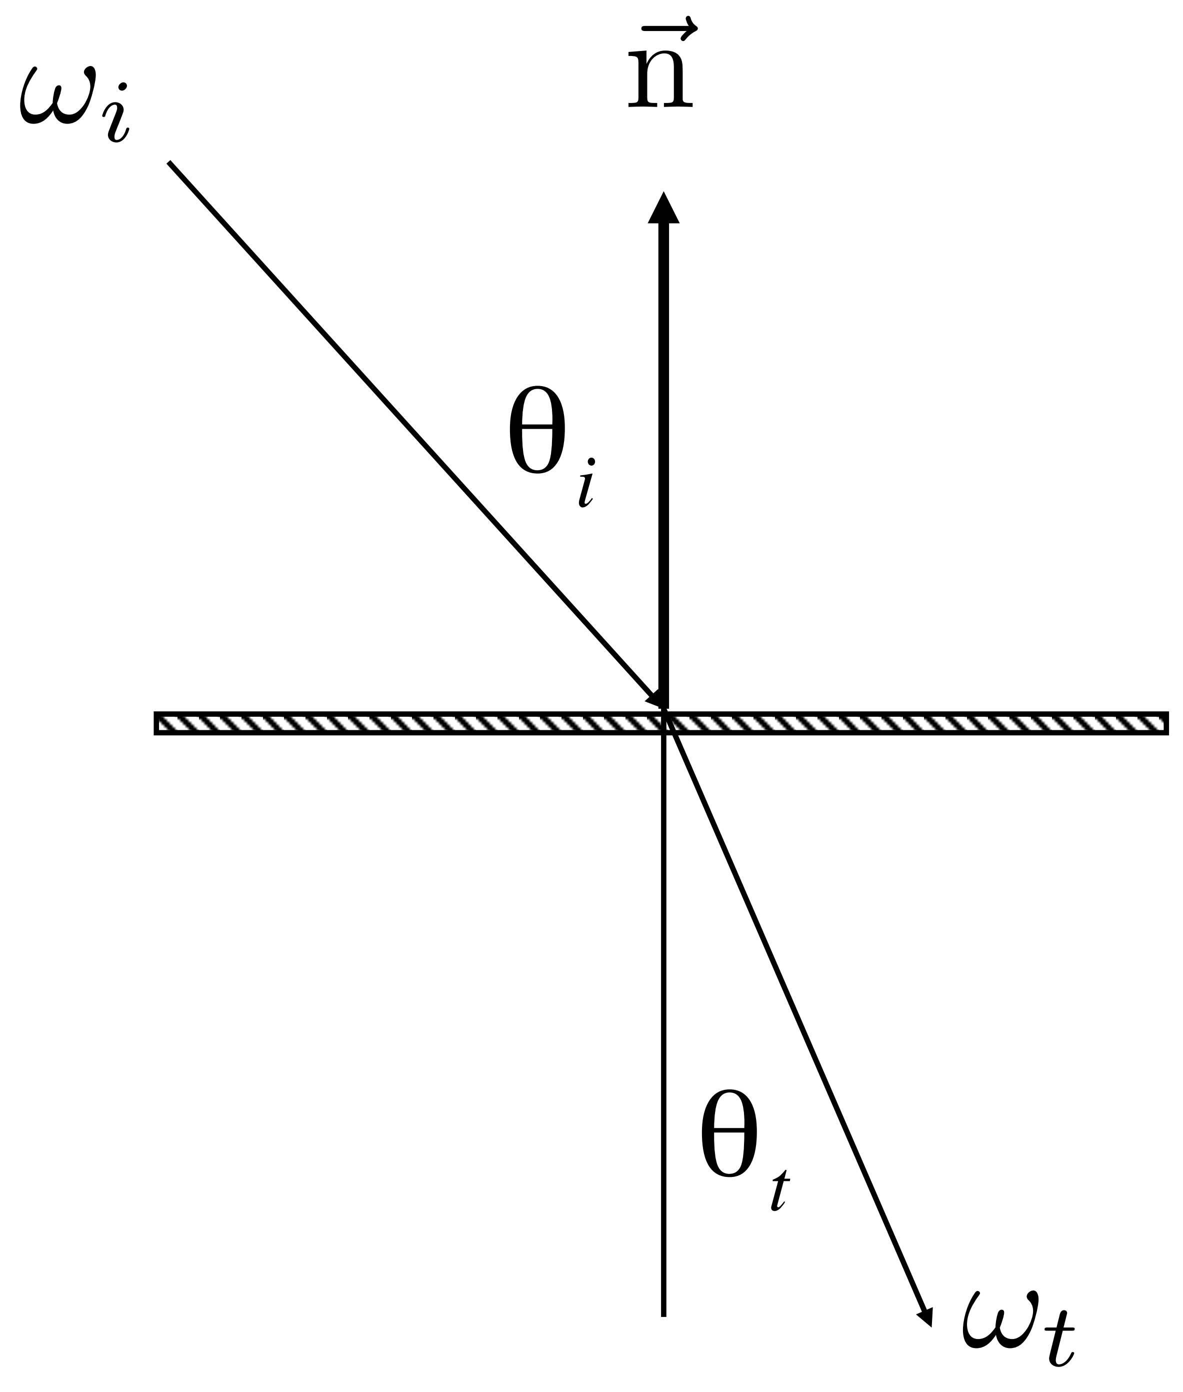

$$\eta_i \sin\theta_i = \eta_t \sin\theta_t$$

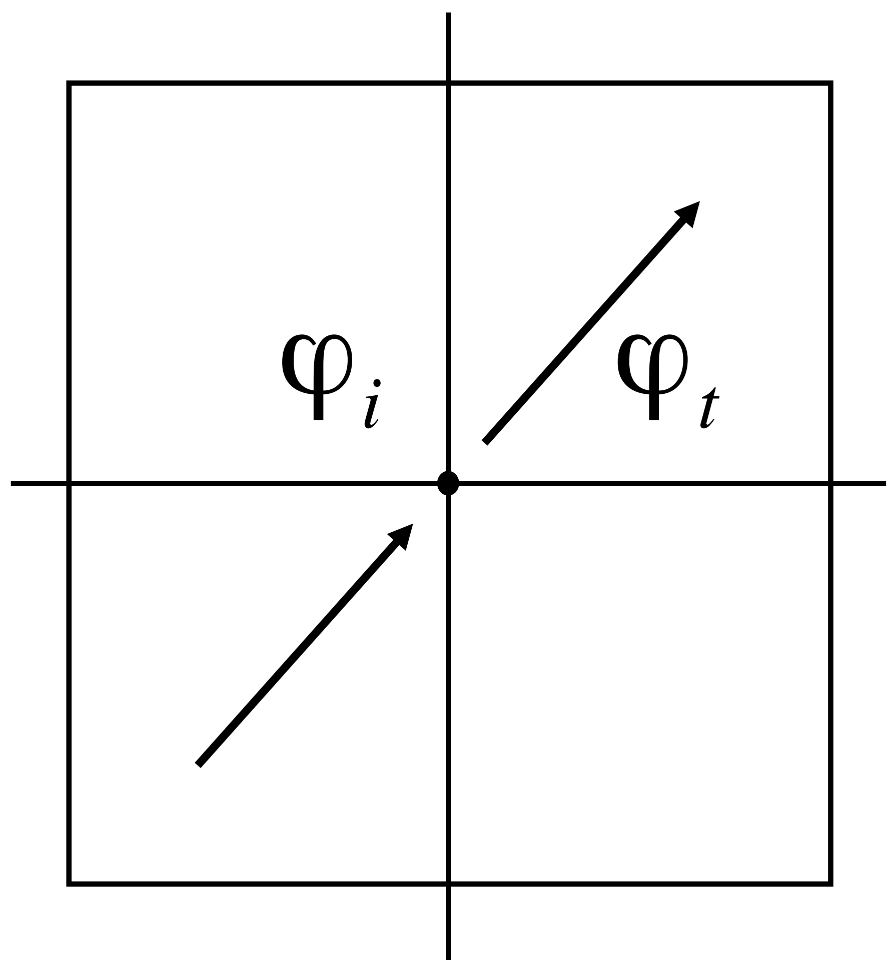

Medium	η*
Vacuum	1.0
Air (sea level)	1.00029
Water (20°C)	1.333
Glass	1.5-1.6
Diamond	2.42

*** index of refraction is wavelength dependent (these are averages)**

Fresnel reflection

For many real materials, reflectance increases w/ viewing angle

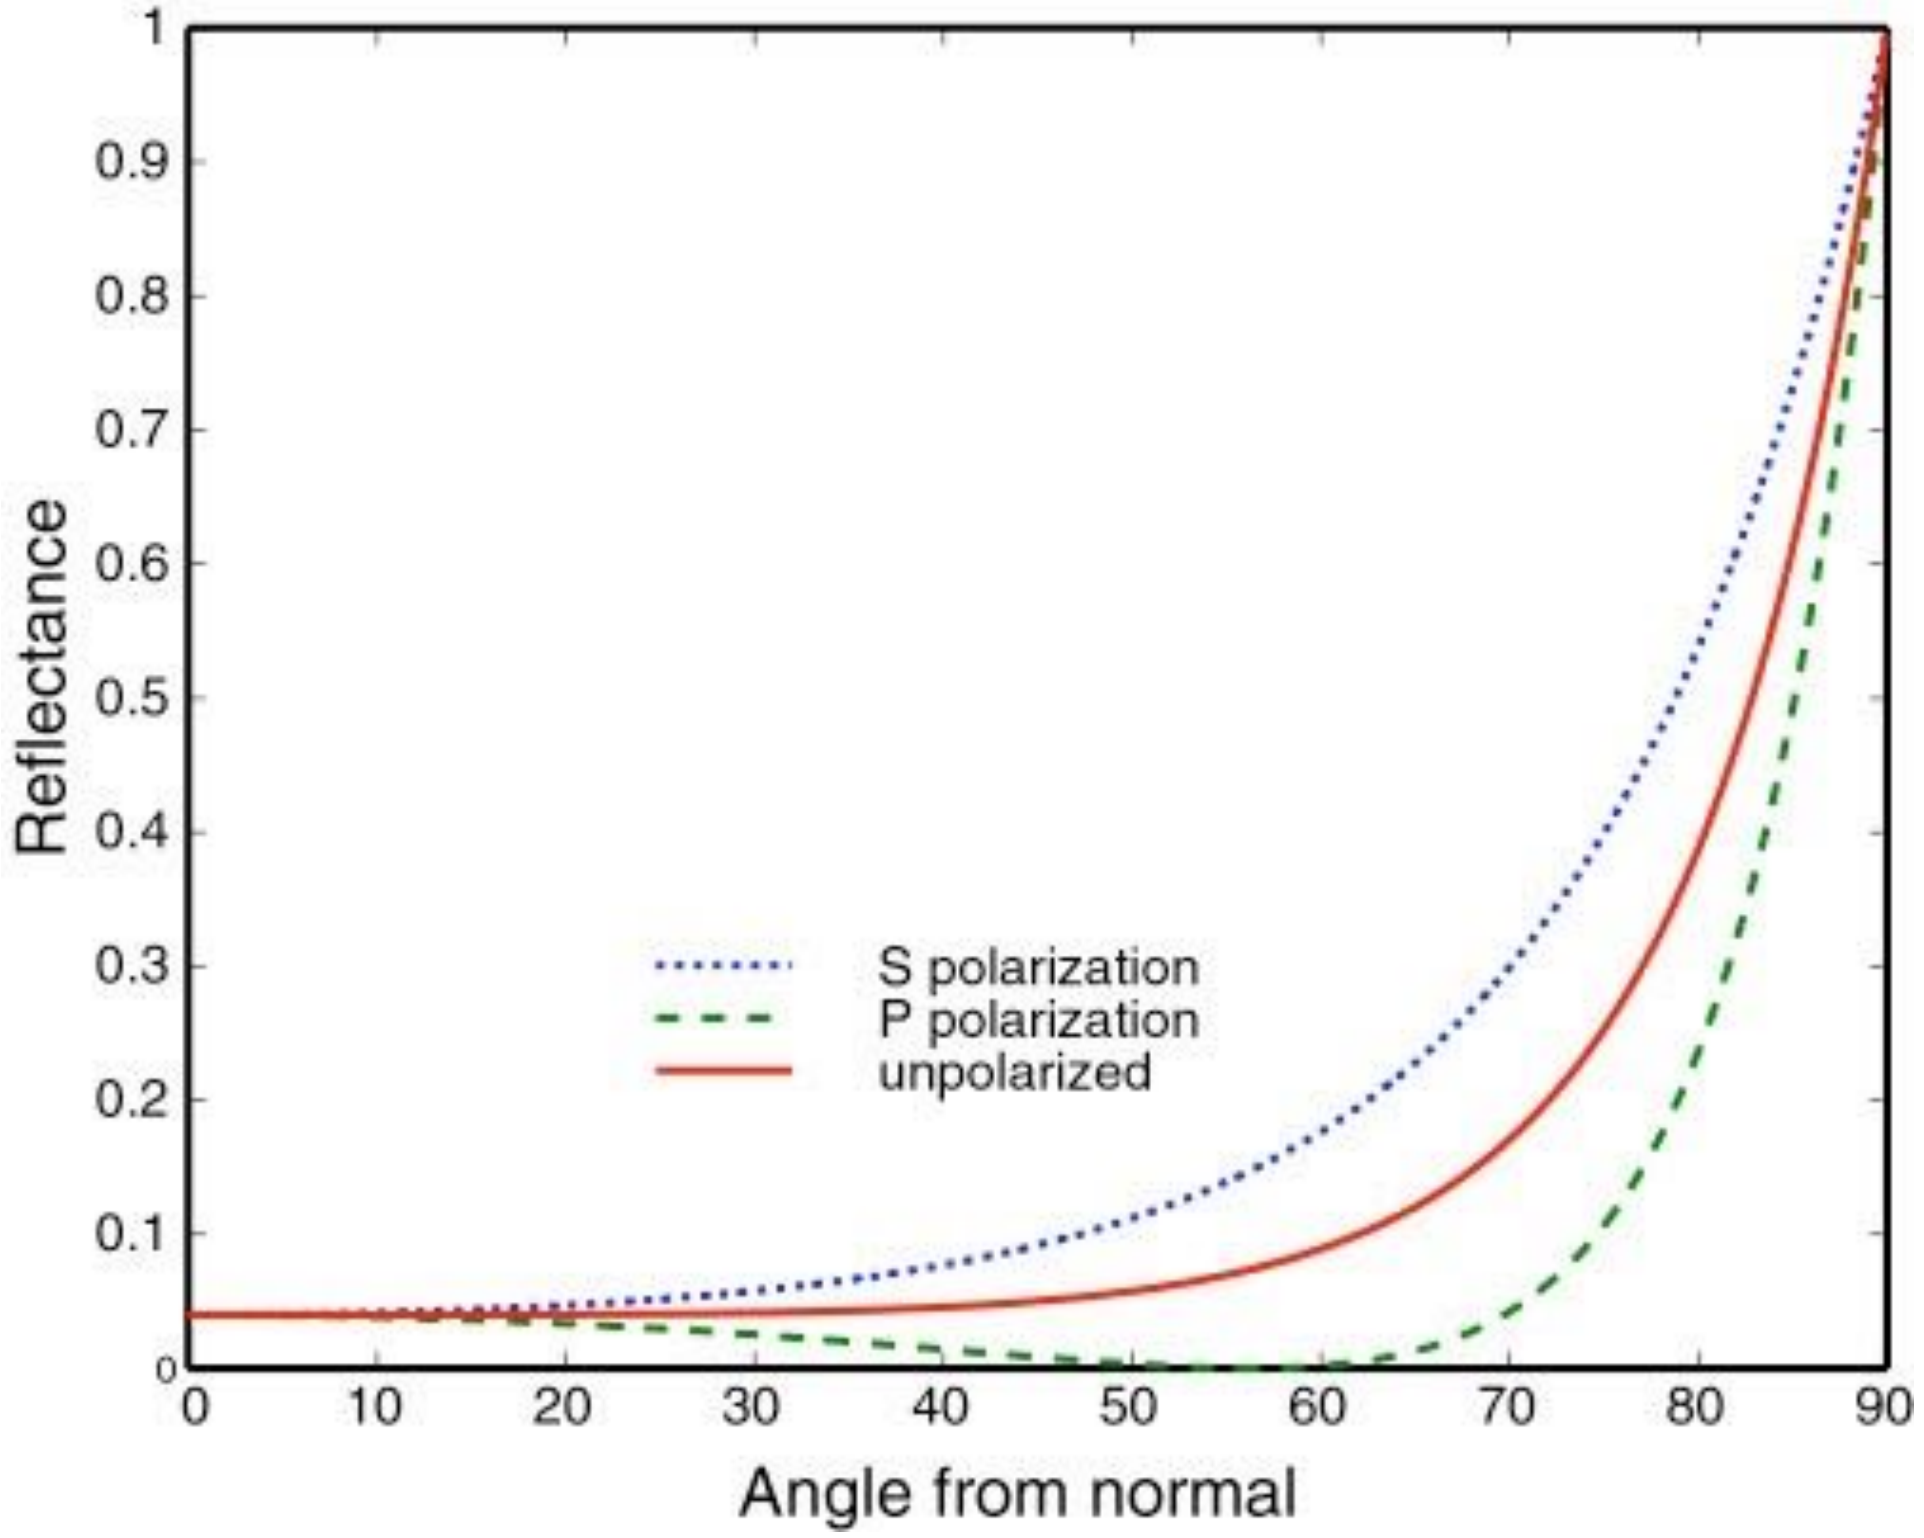

[Lafortune et al. 1997]

Snell + Fresnel: example

Subsurface scattering

- Visual characteristics of many surfaces caused by light entering at different points than it exits
  - Violates a fundamental assumption of the BRDF
  - Need to generalize scattering model (BSSRDF)

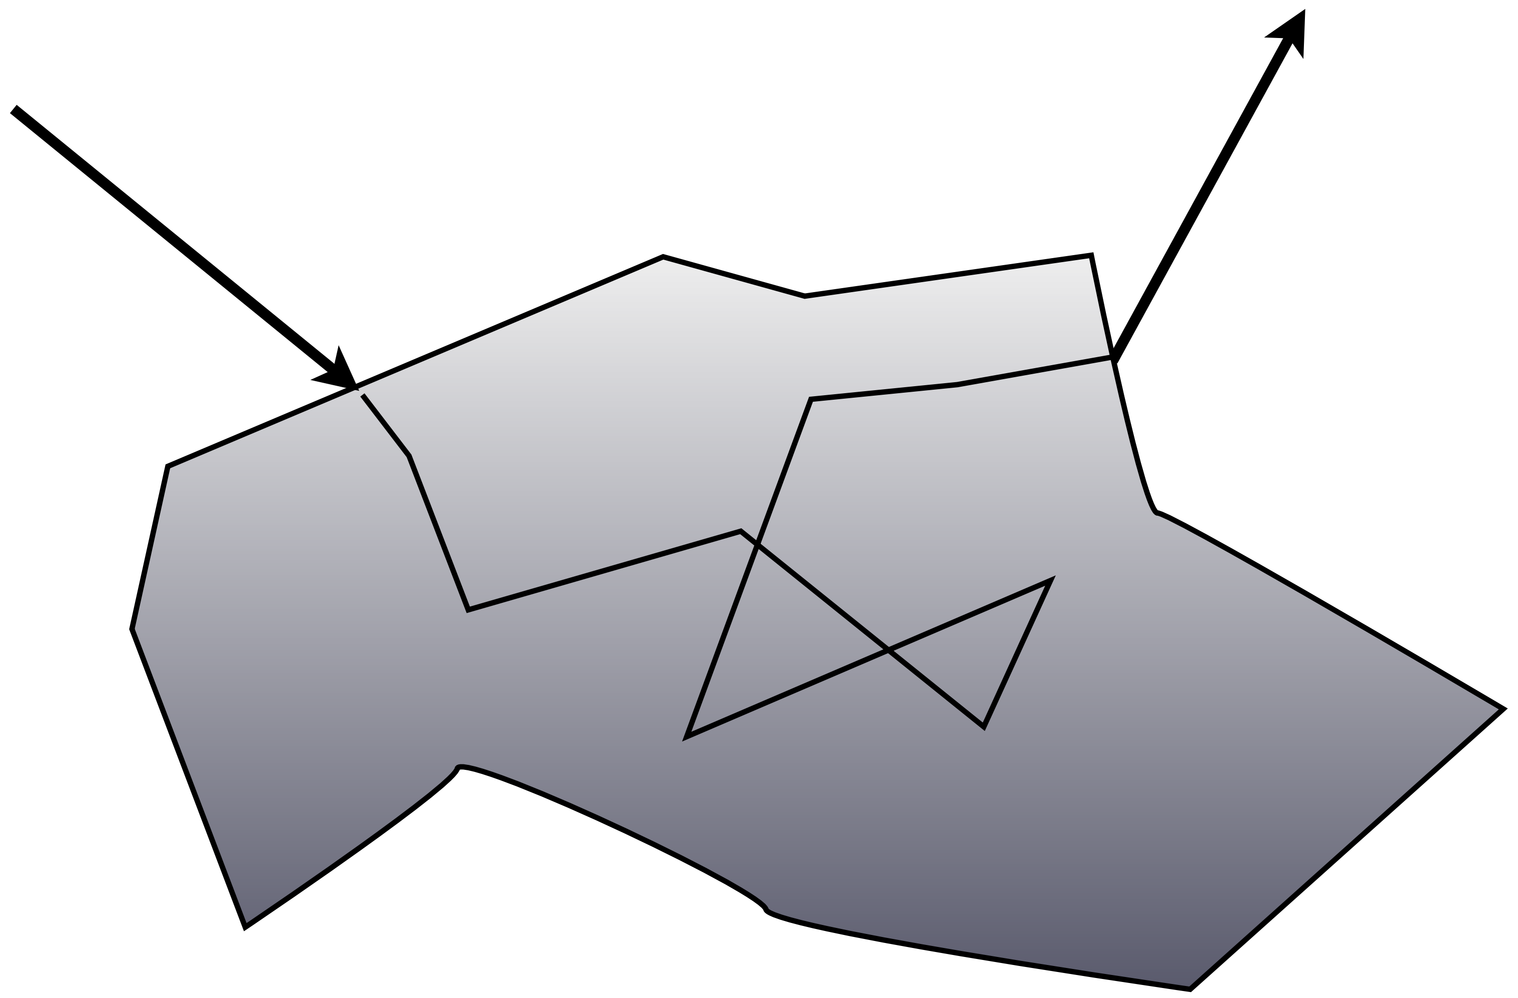

[Jensen et al 2001]

[Donner et al 2008]

* BSSRDF = bidirectional subsurface scatting reflectance distribution function

Translucent materials: skin

BRDF

BSSRDF

Parameters to Disney BRDF

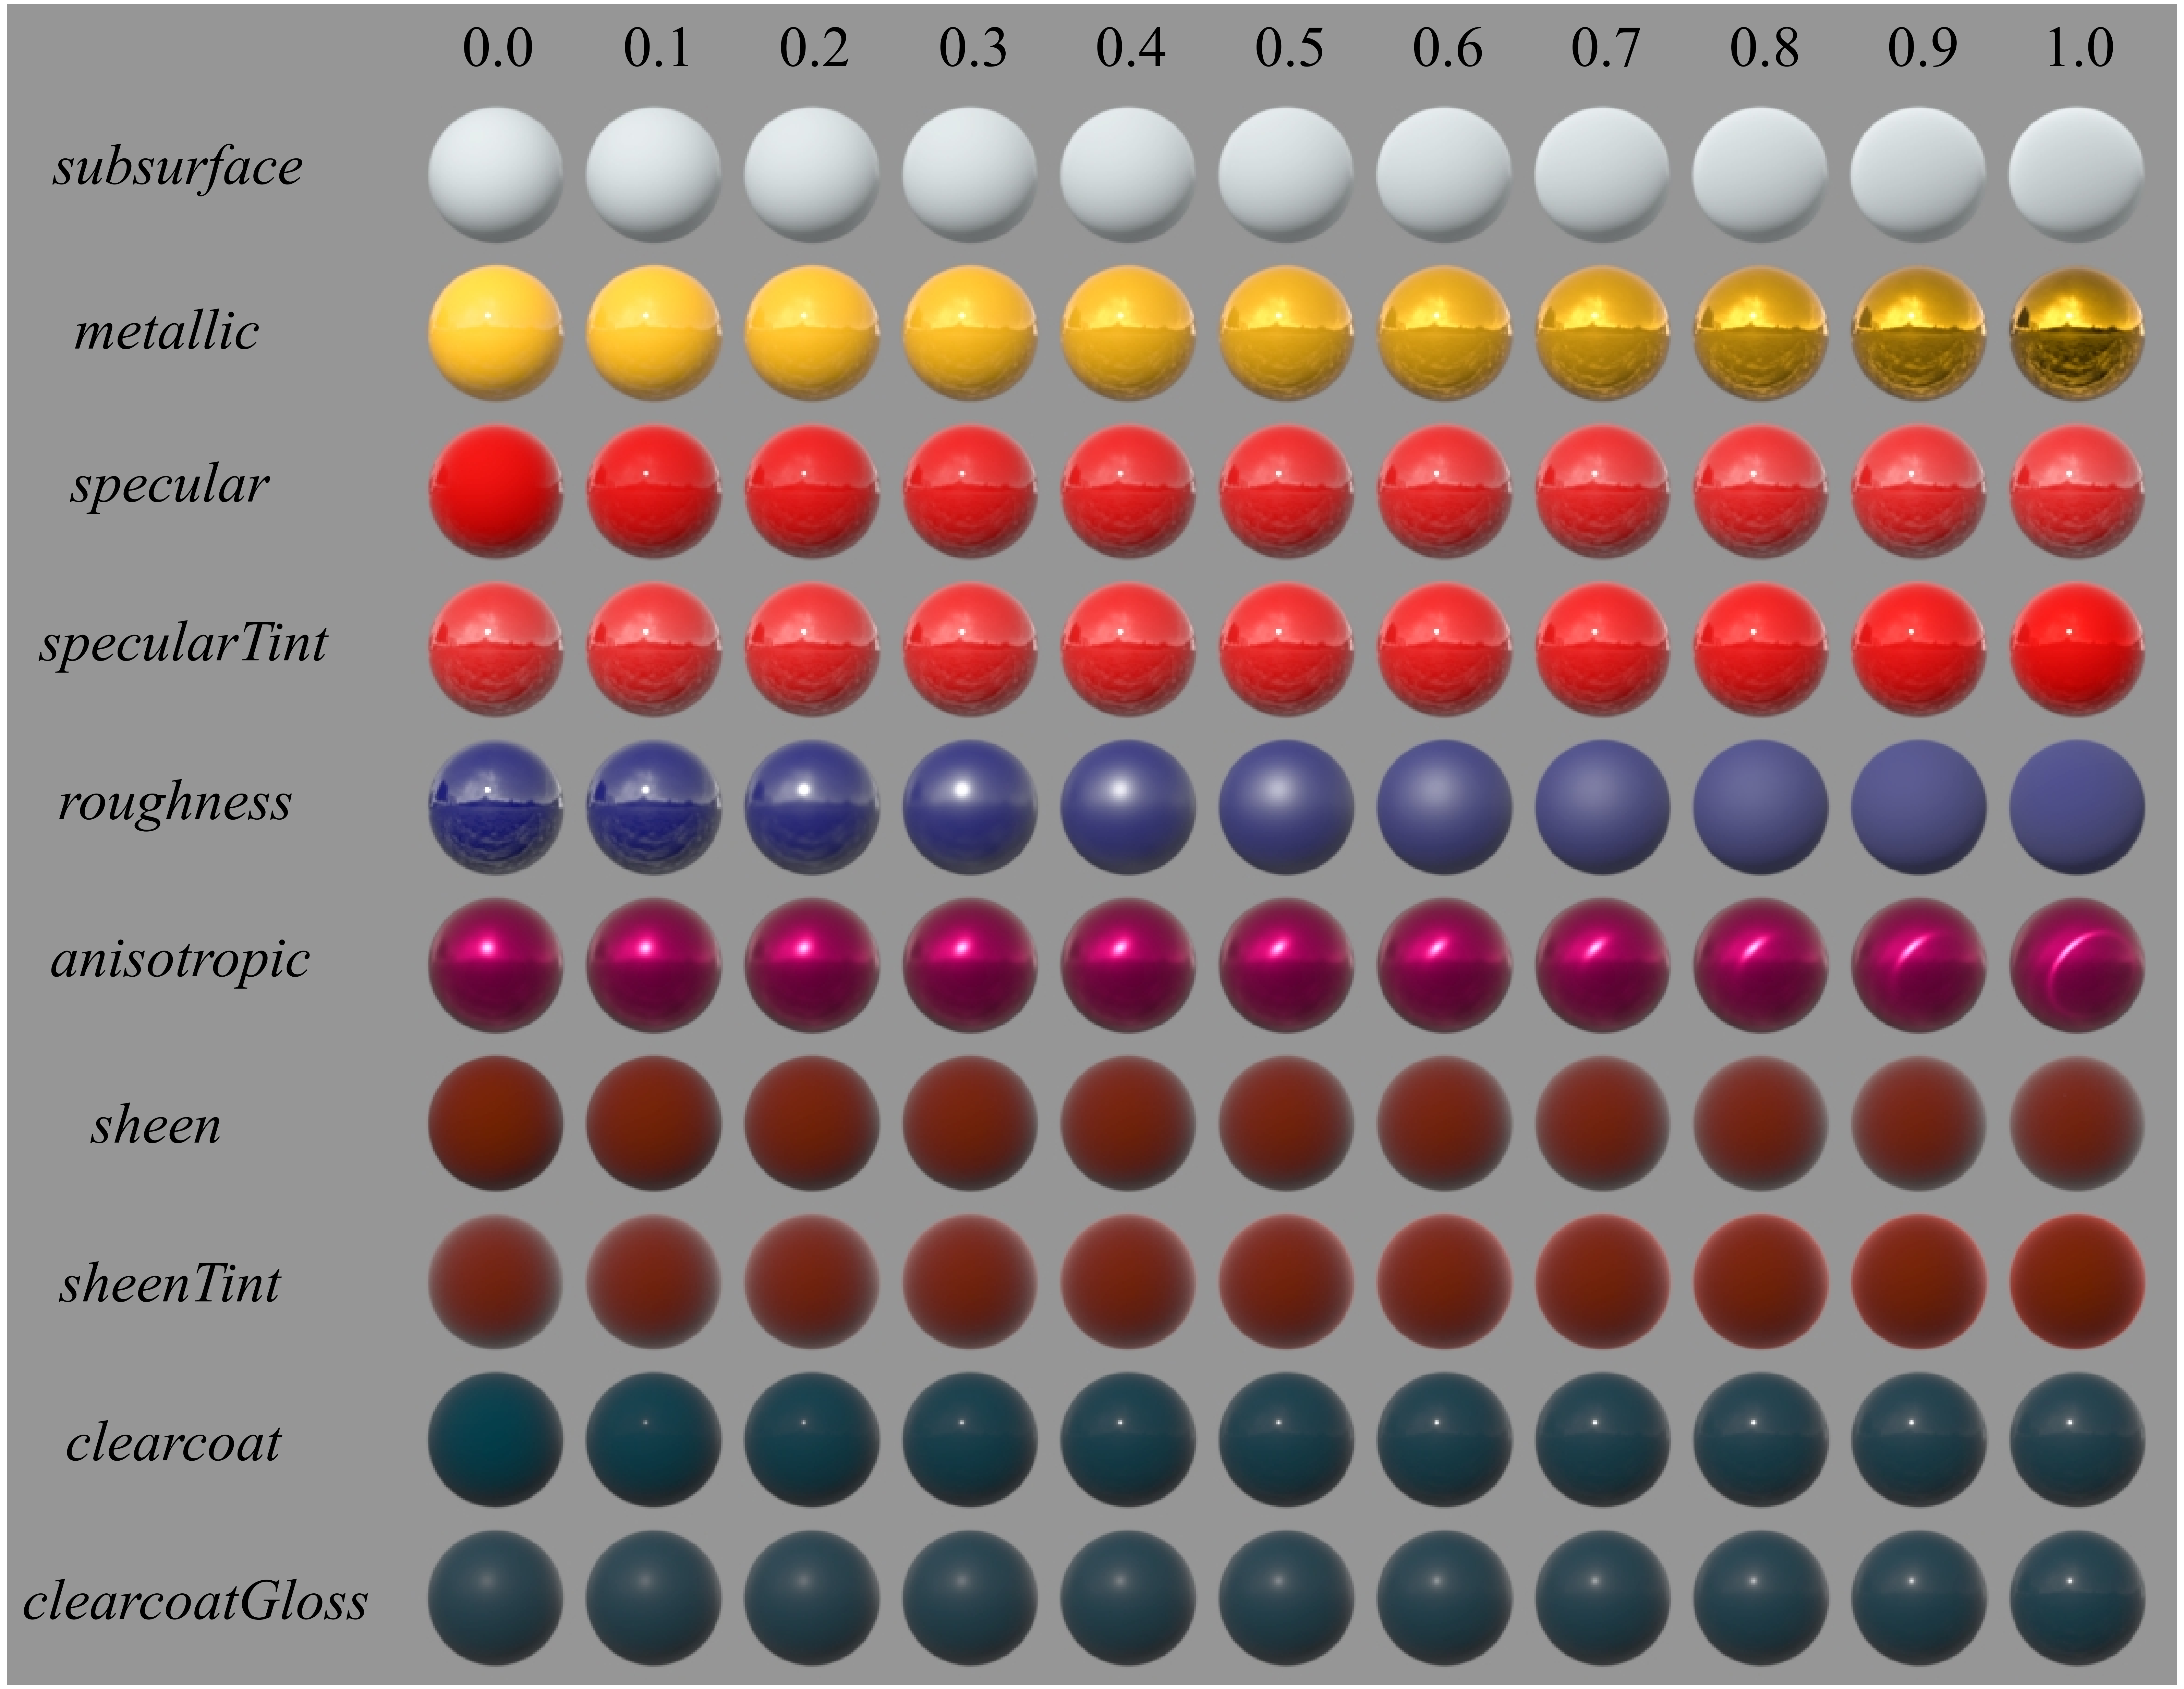

Pattern generation vs. BRDF

In practice, it is convenient to separate computation of spatially varying BRDF parameters (like albedo, shininess, etc.) from the reflectance function itself

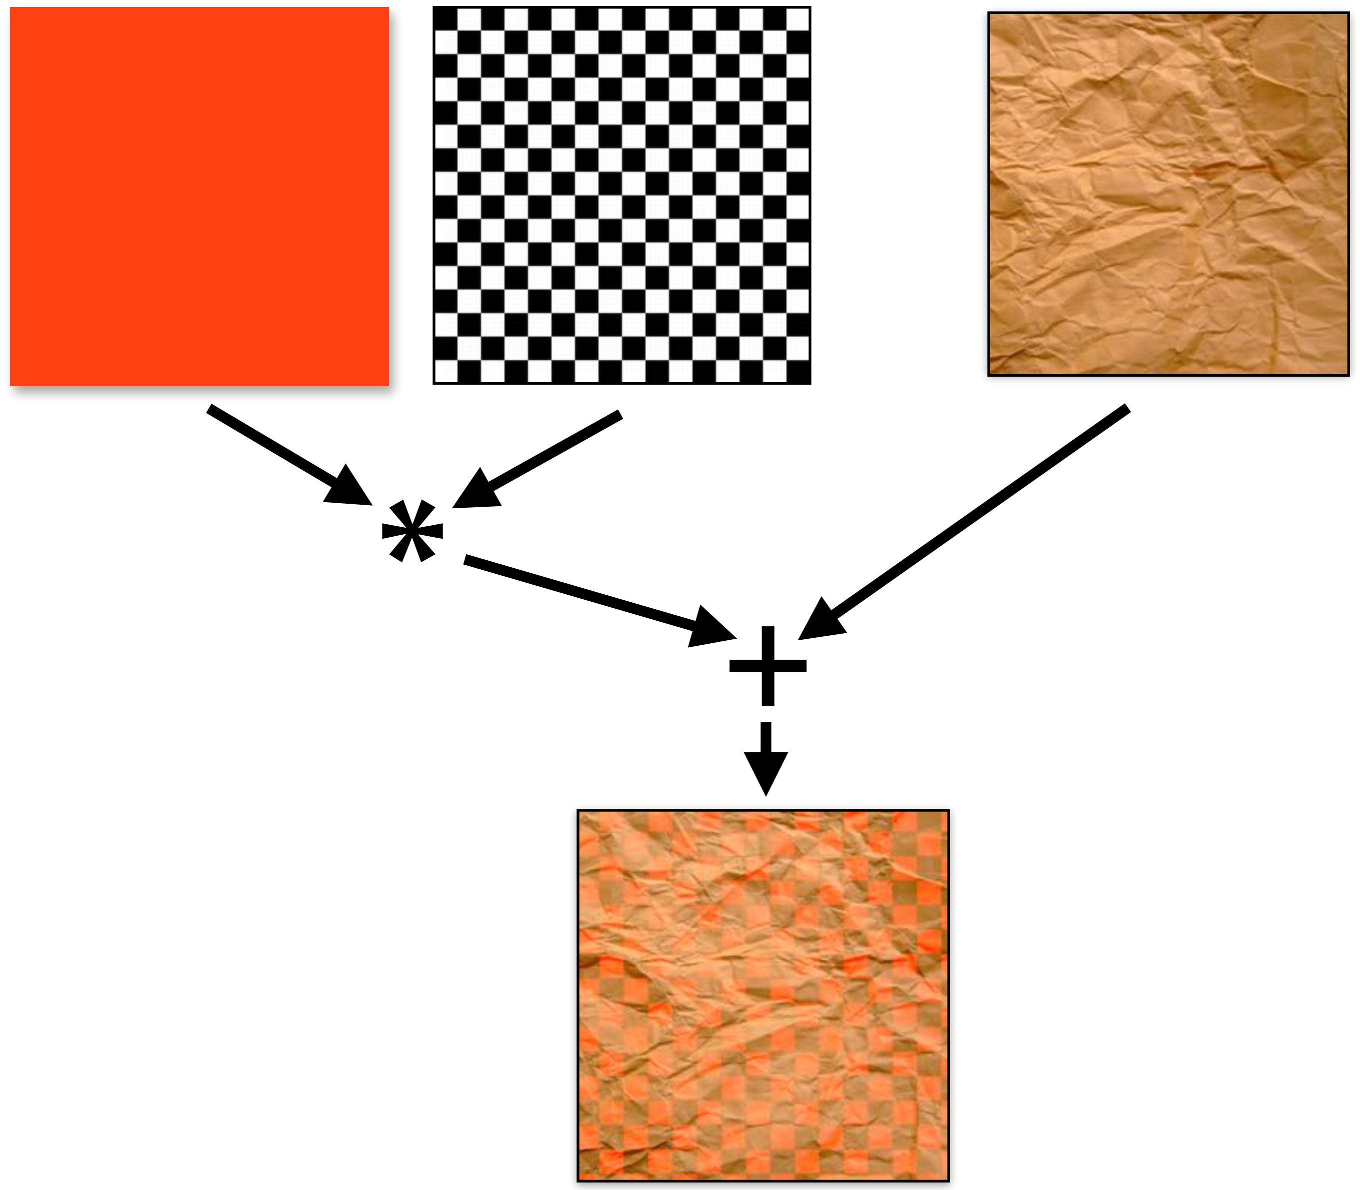

Example 1: albedo value at surface point is given by expression combining multiple textures

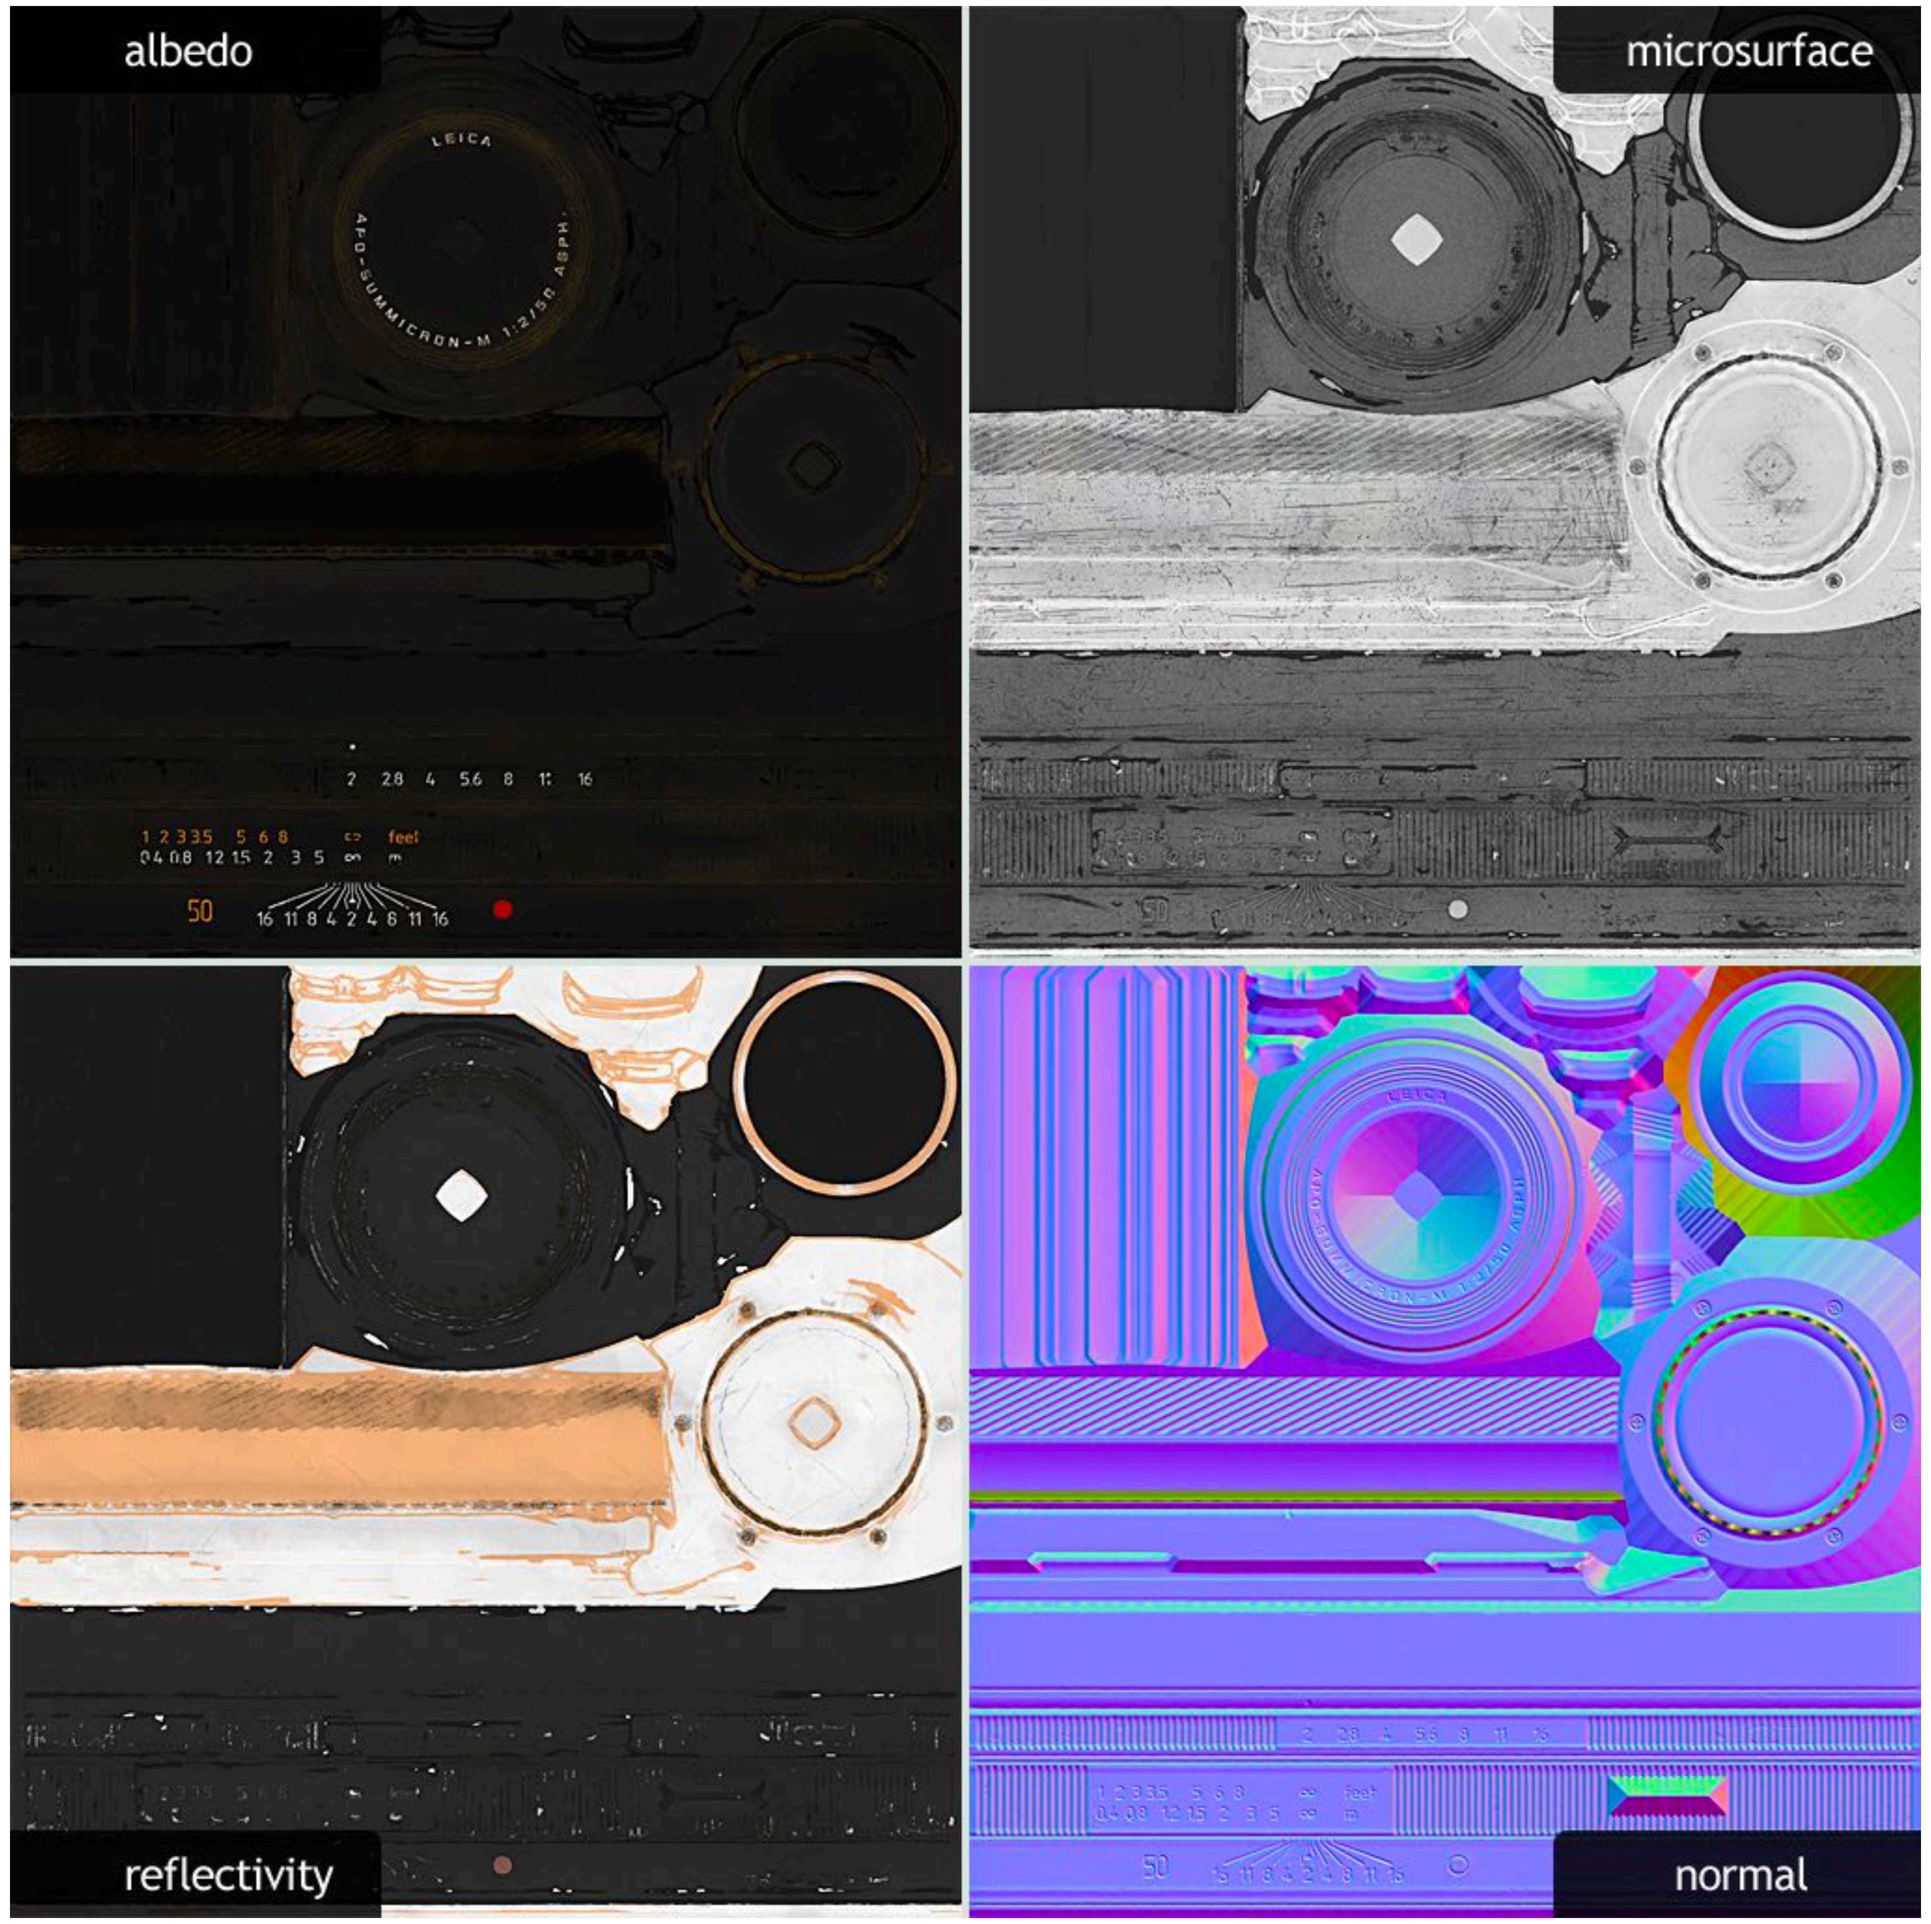

Example 2:
Different textures defining different spatially varying BRDF input parameters

Unity's shader graph

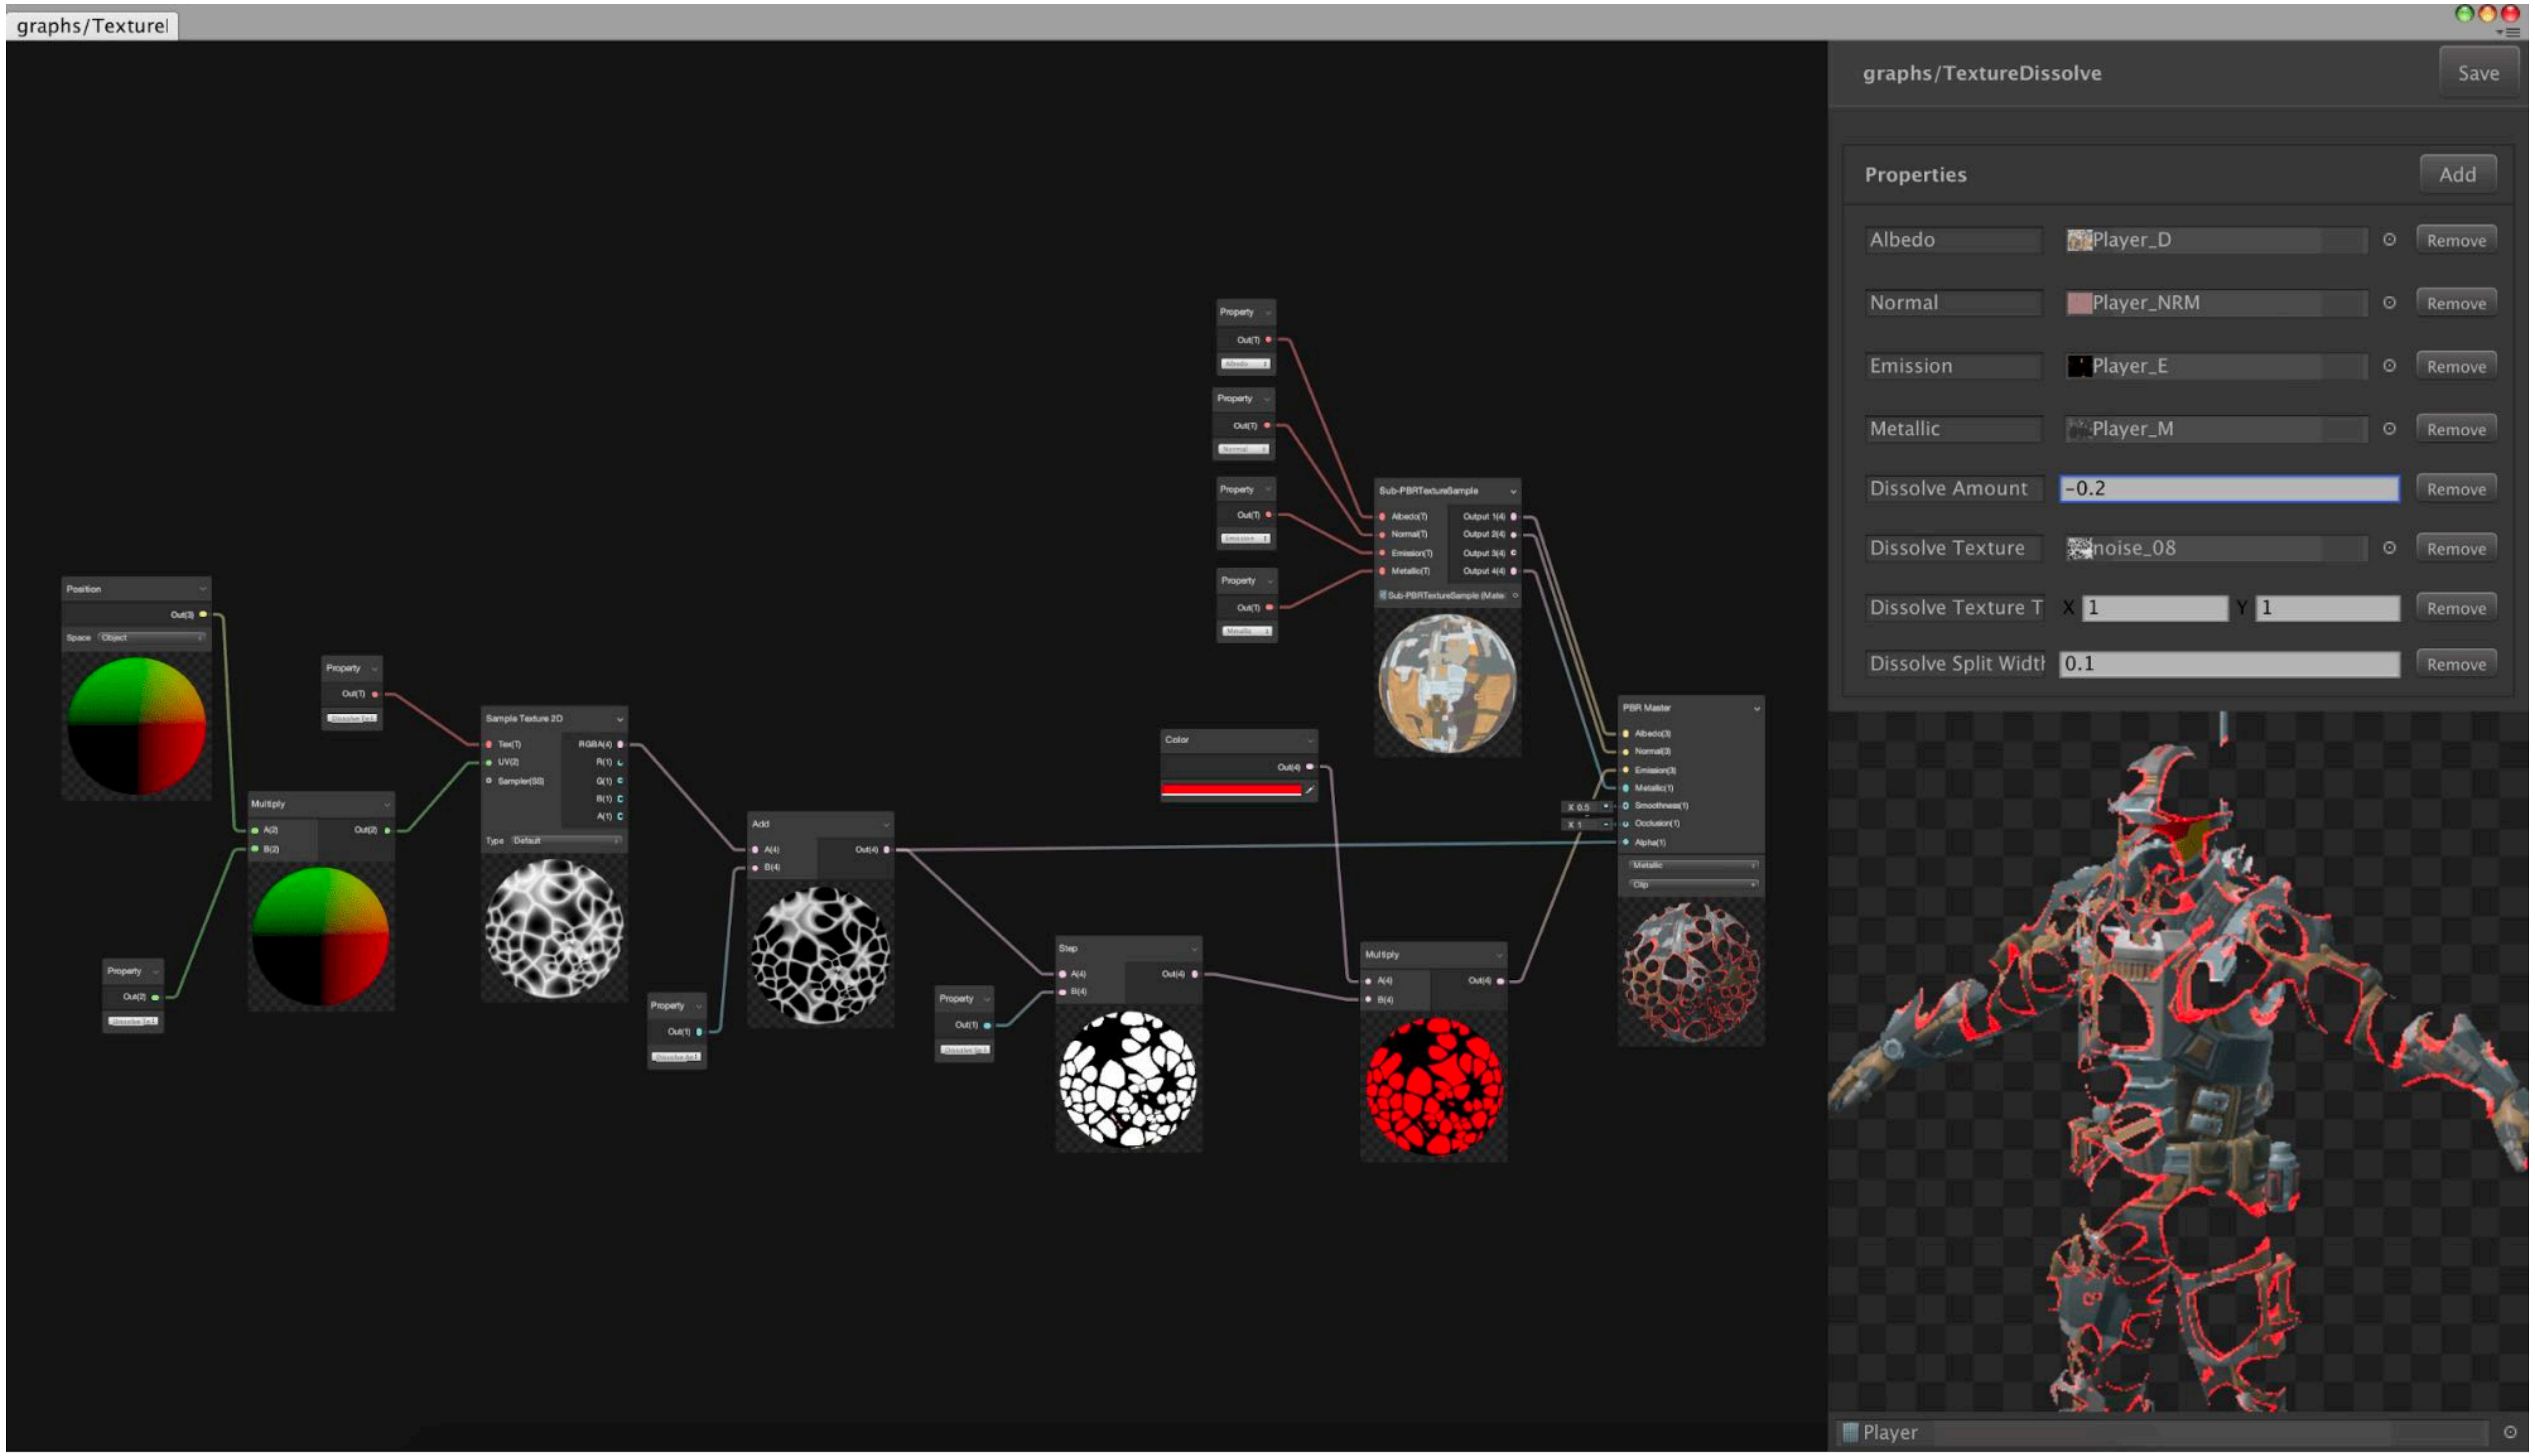

Numerical Integration

So far in this lecture, we've seen examples of needing to compute integrals

Example: computing incident irradiance at x due to a single area light source.

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

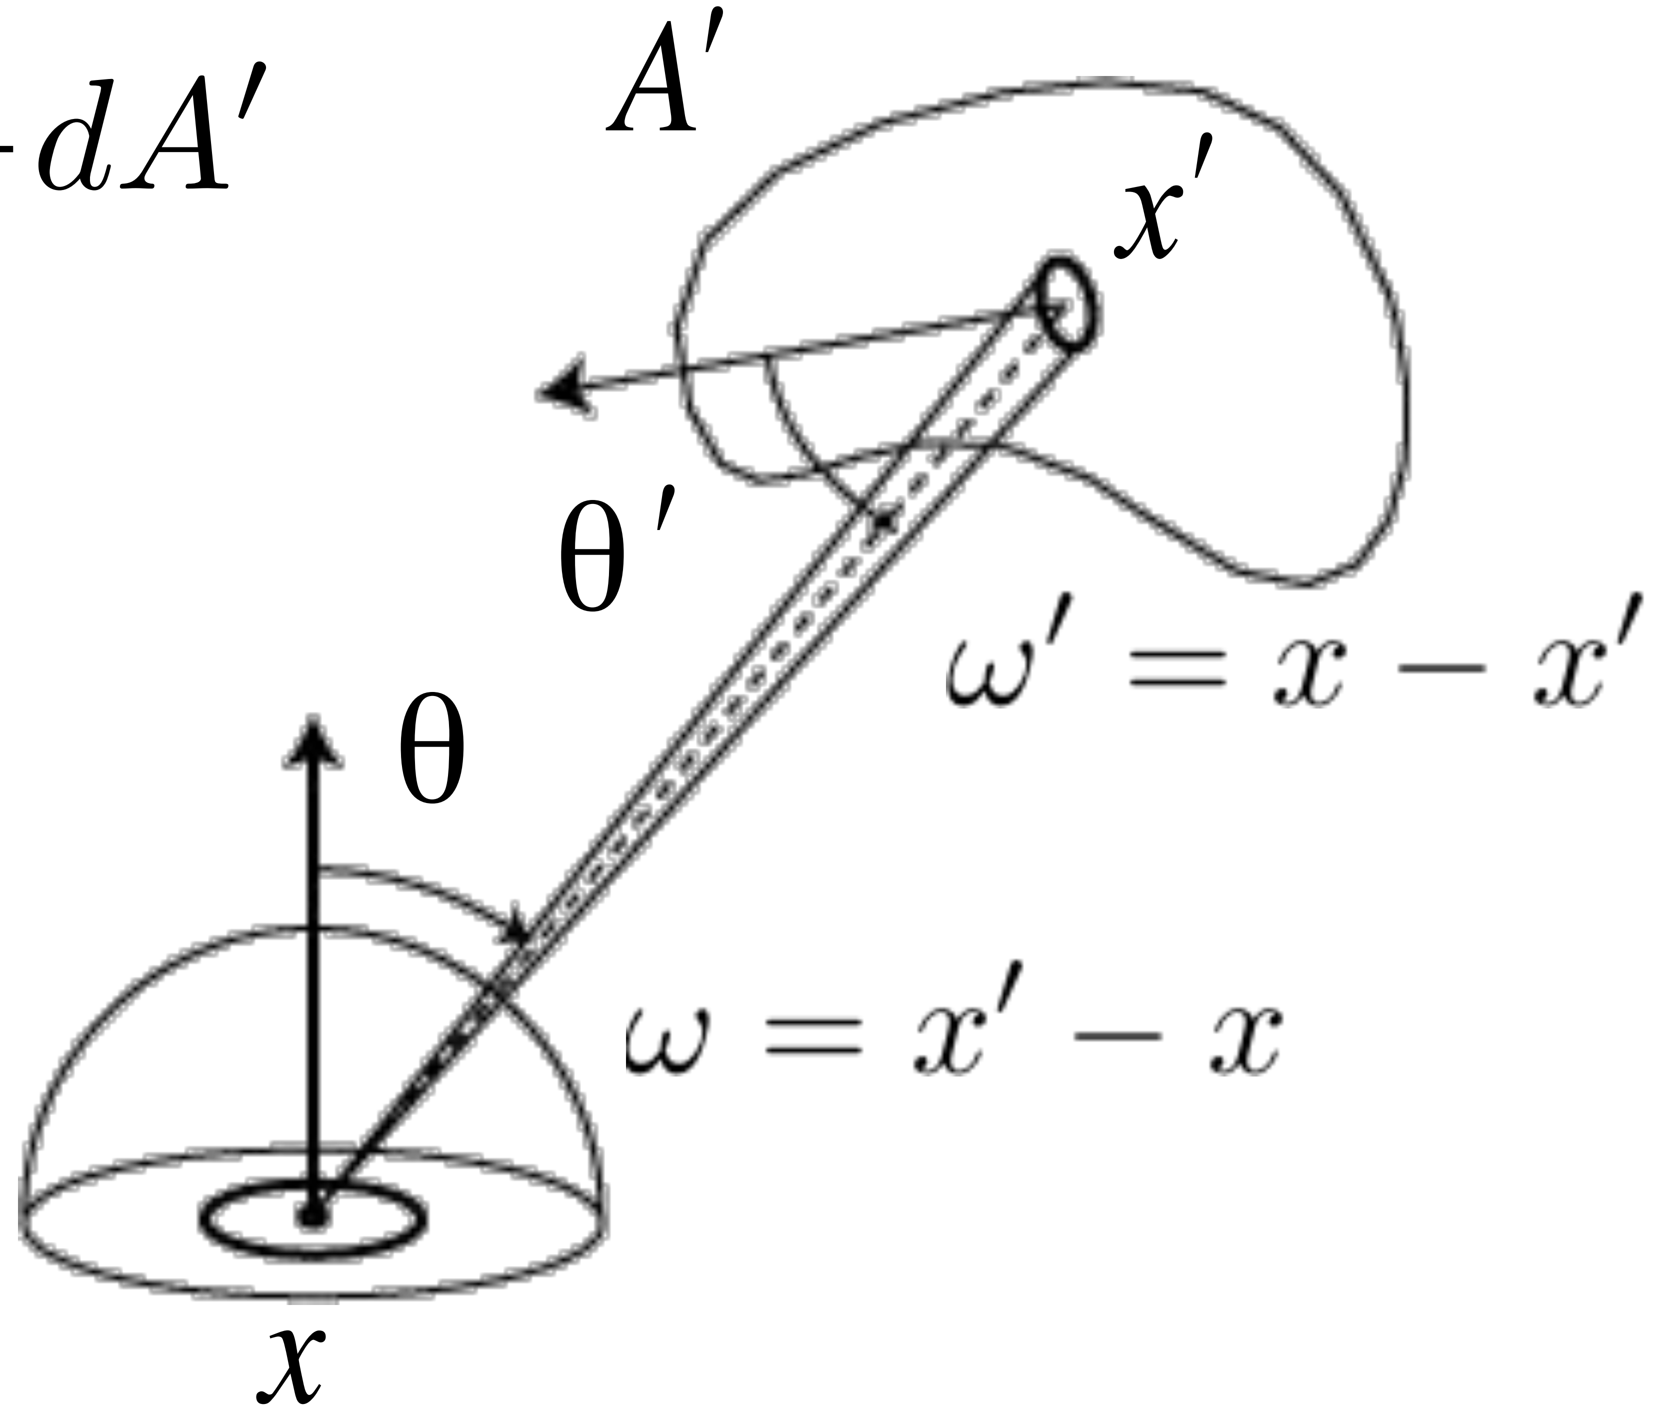

Review: fundamental theorem of calculus

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

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

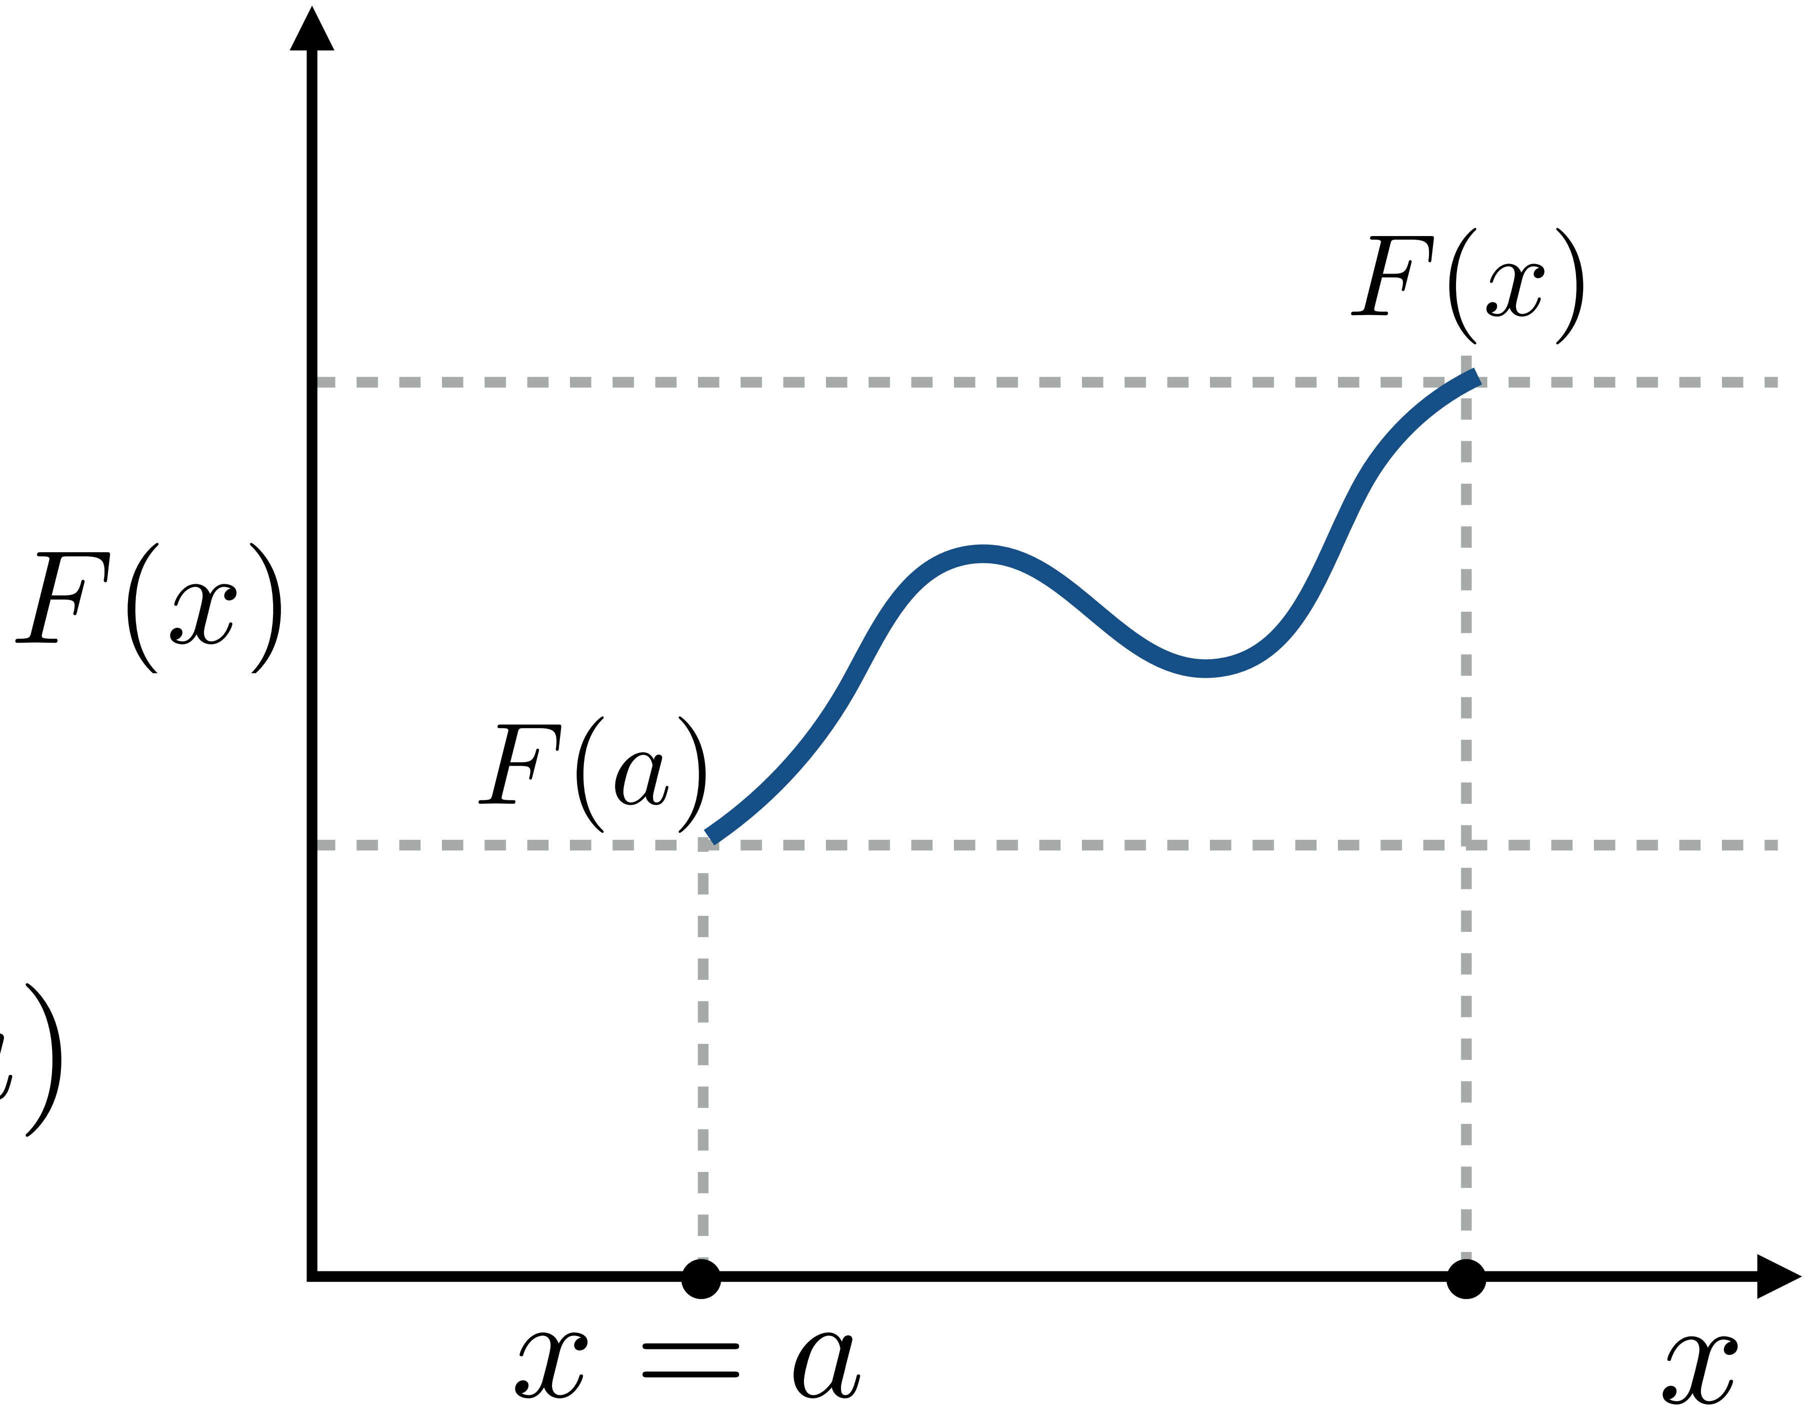

$$\int_a^x f(t)dt = F(x) - F(a)$$

Definite integral as "area under curve"

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

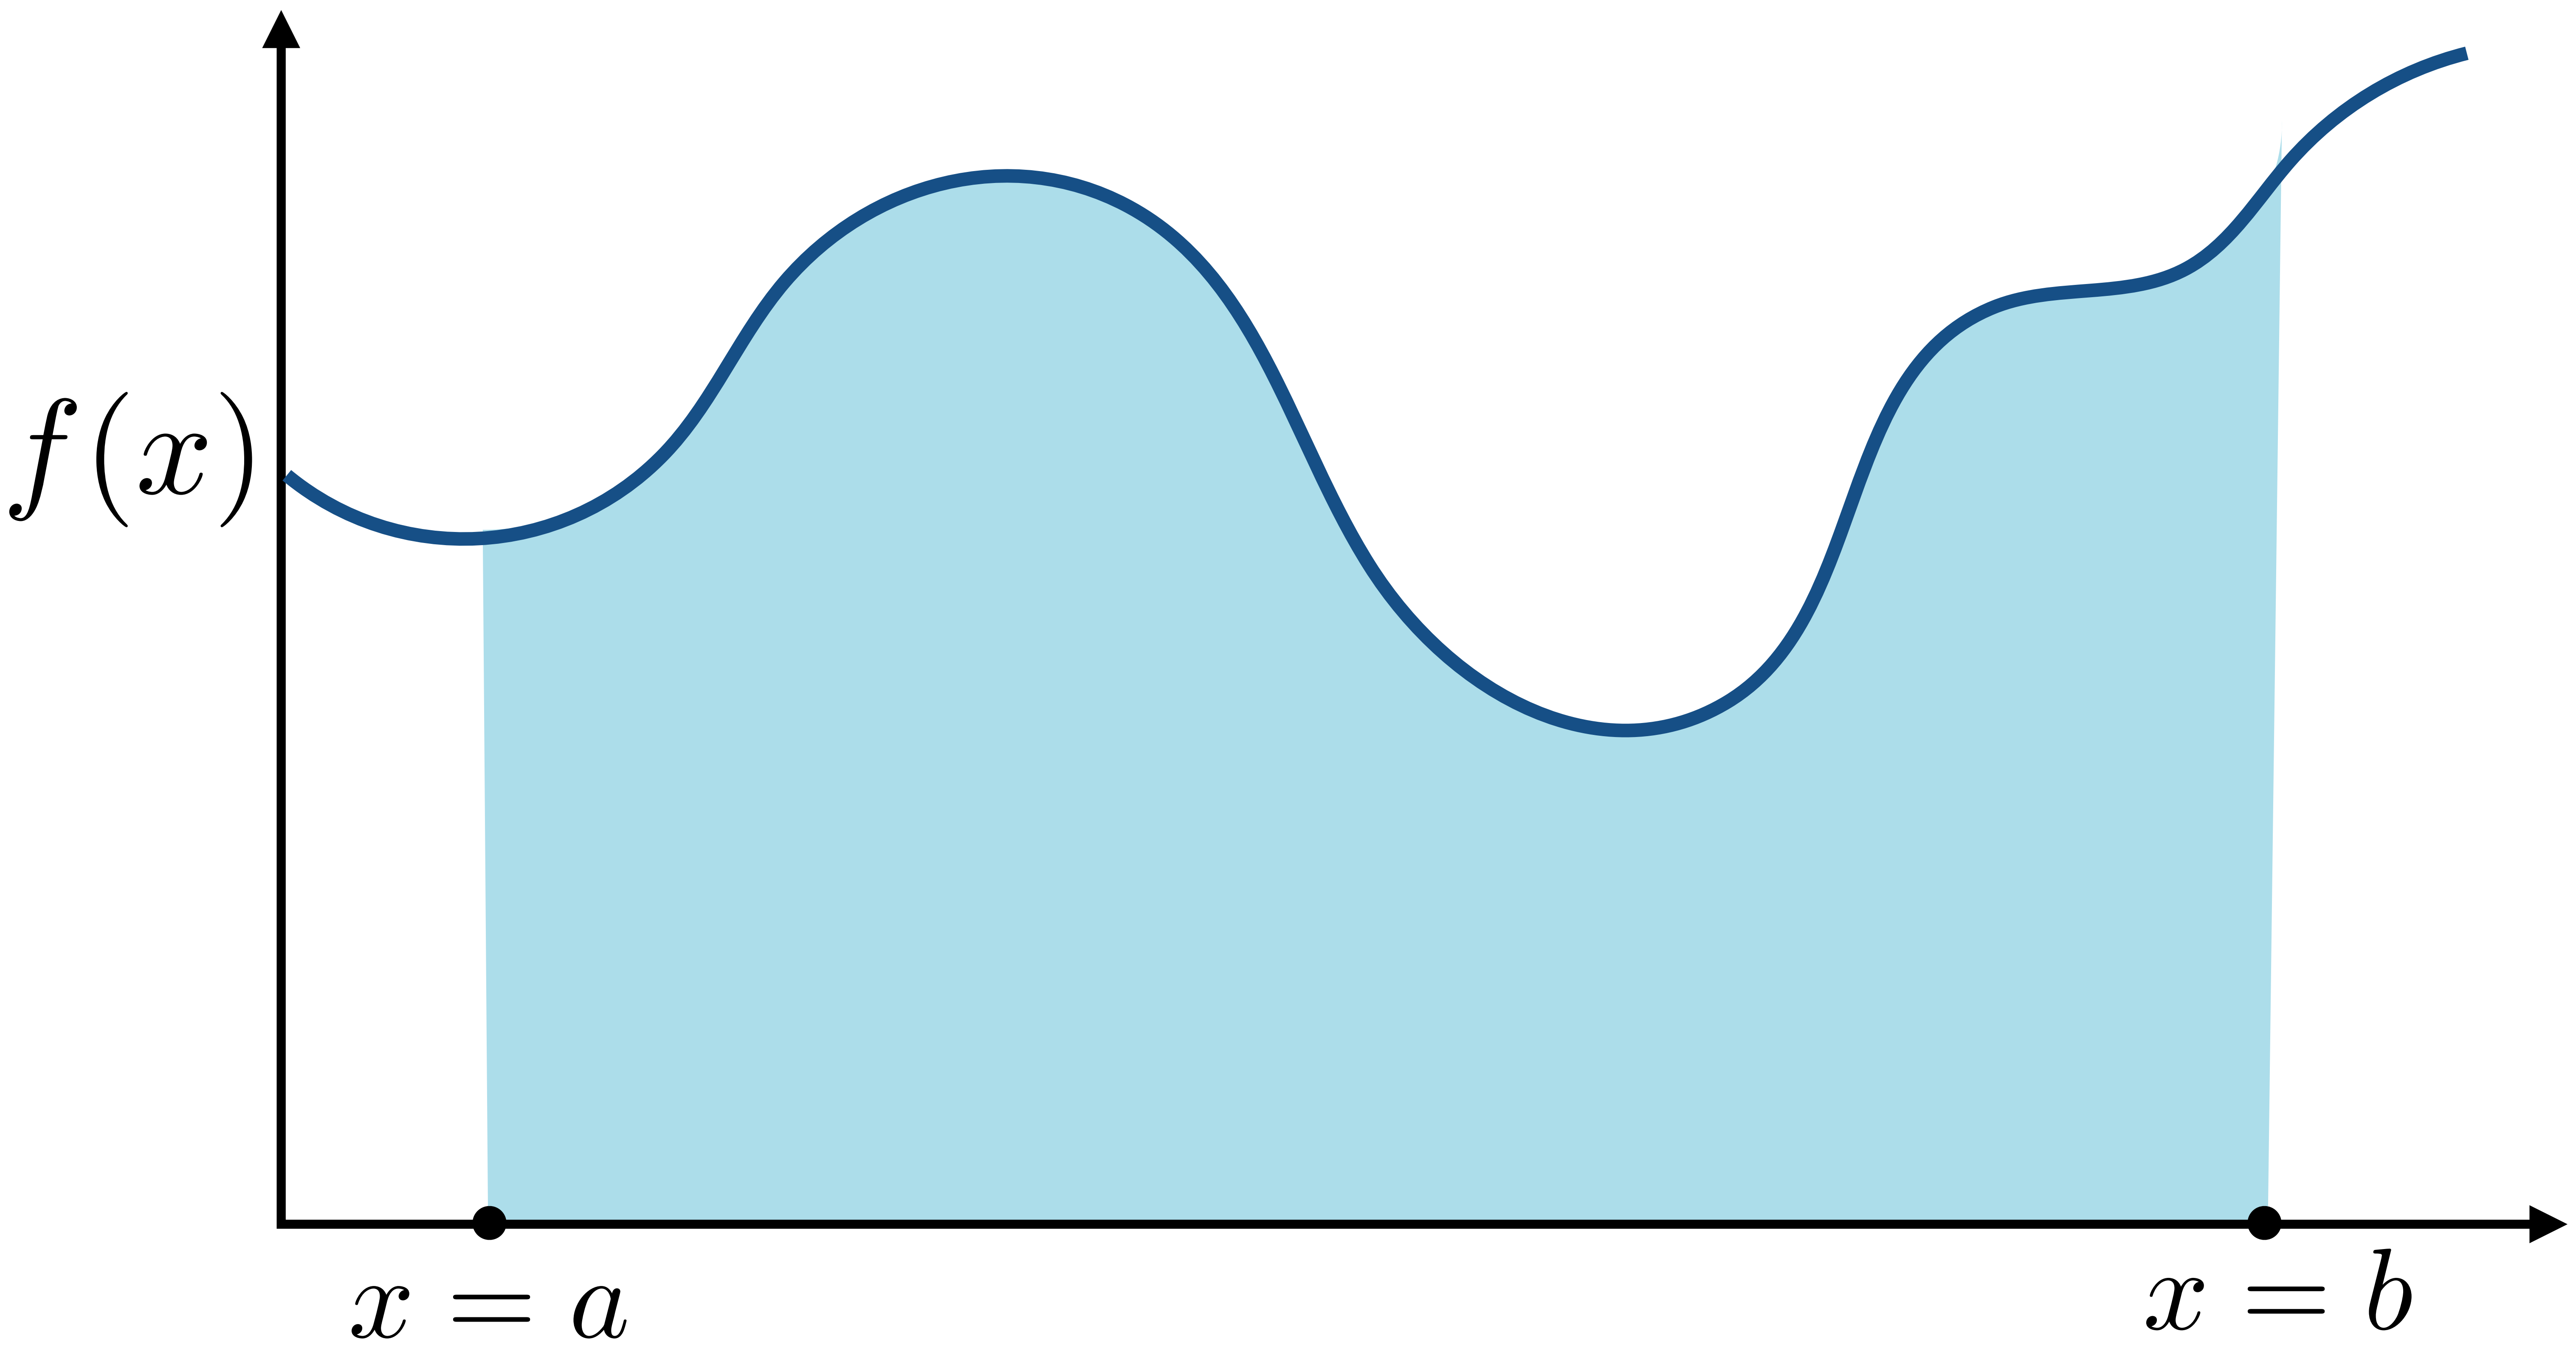

Simple case: constant function

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

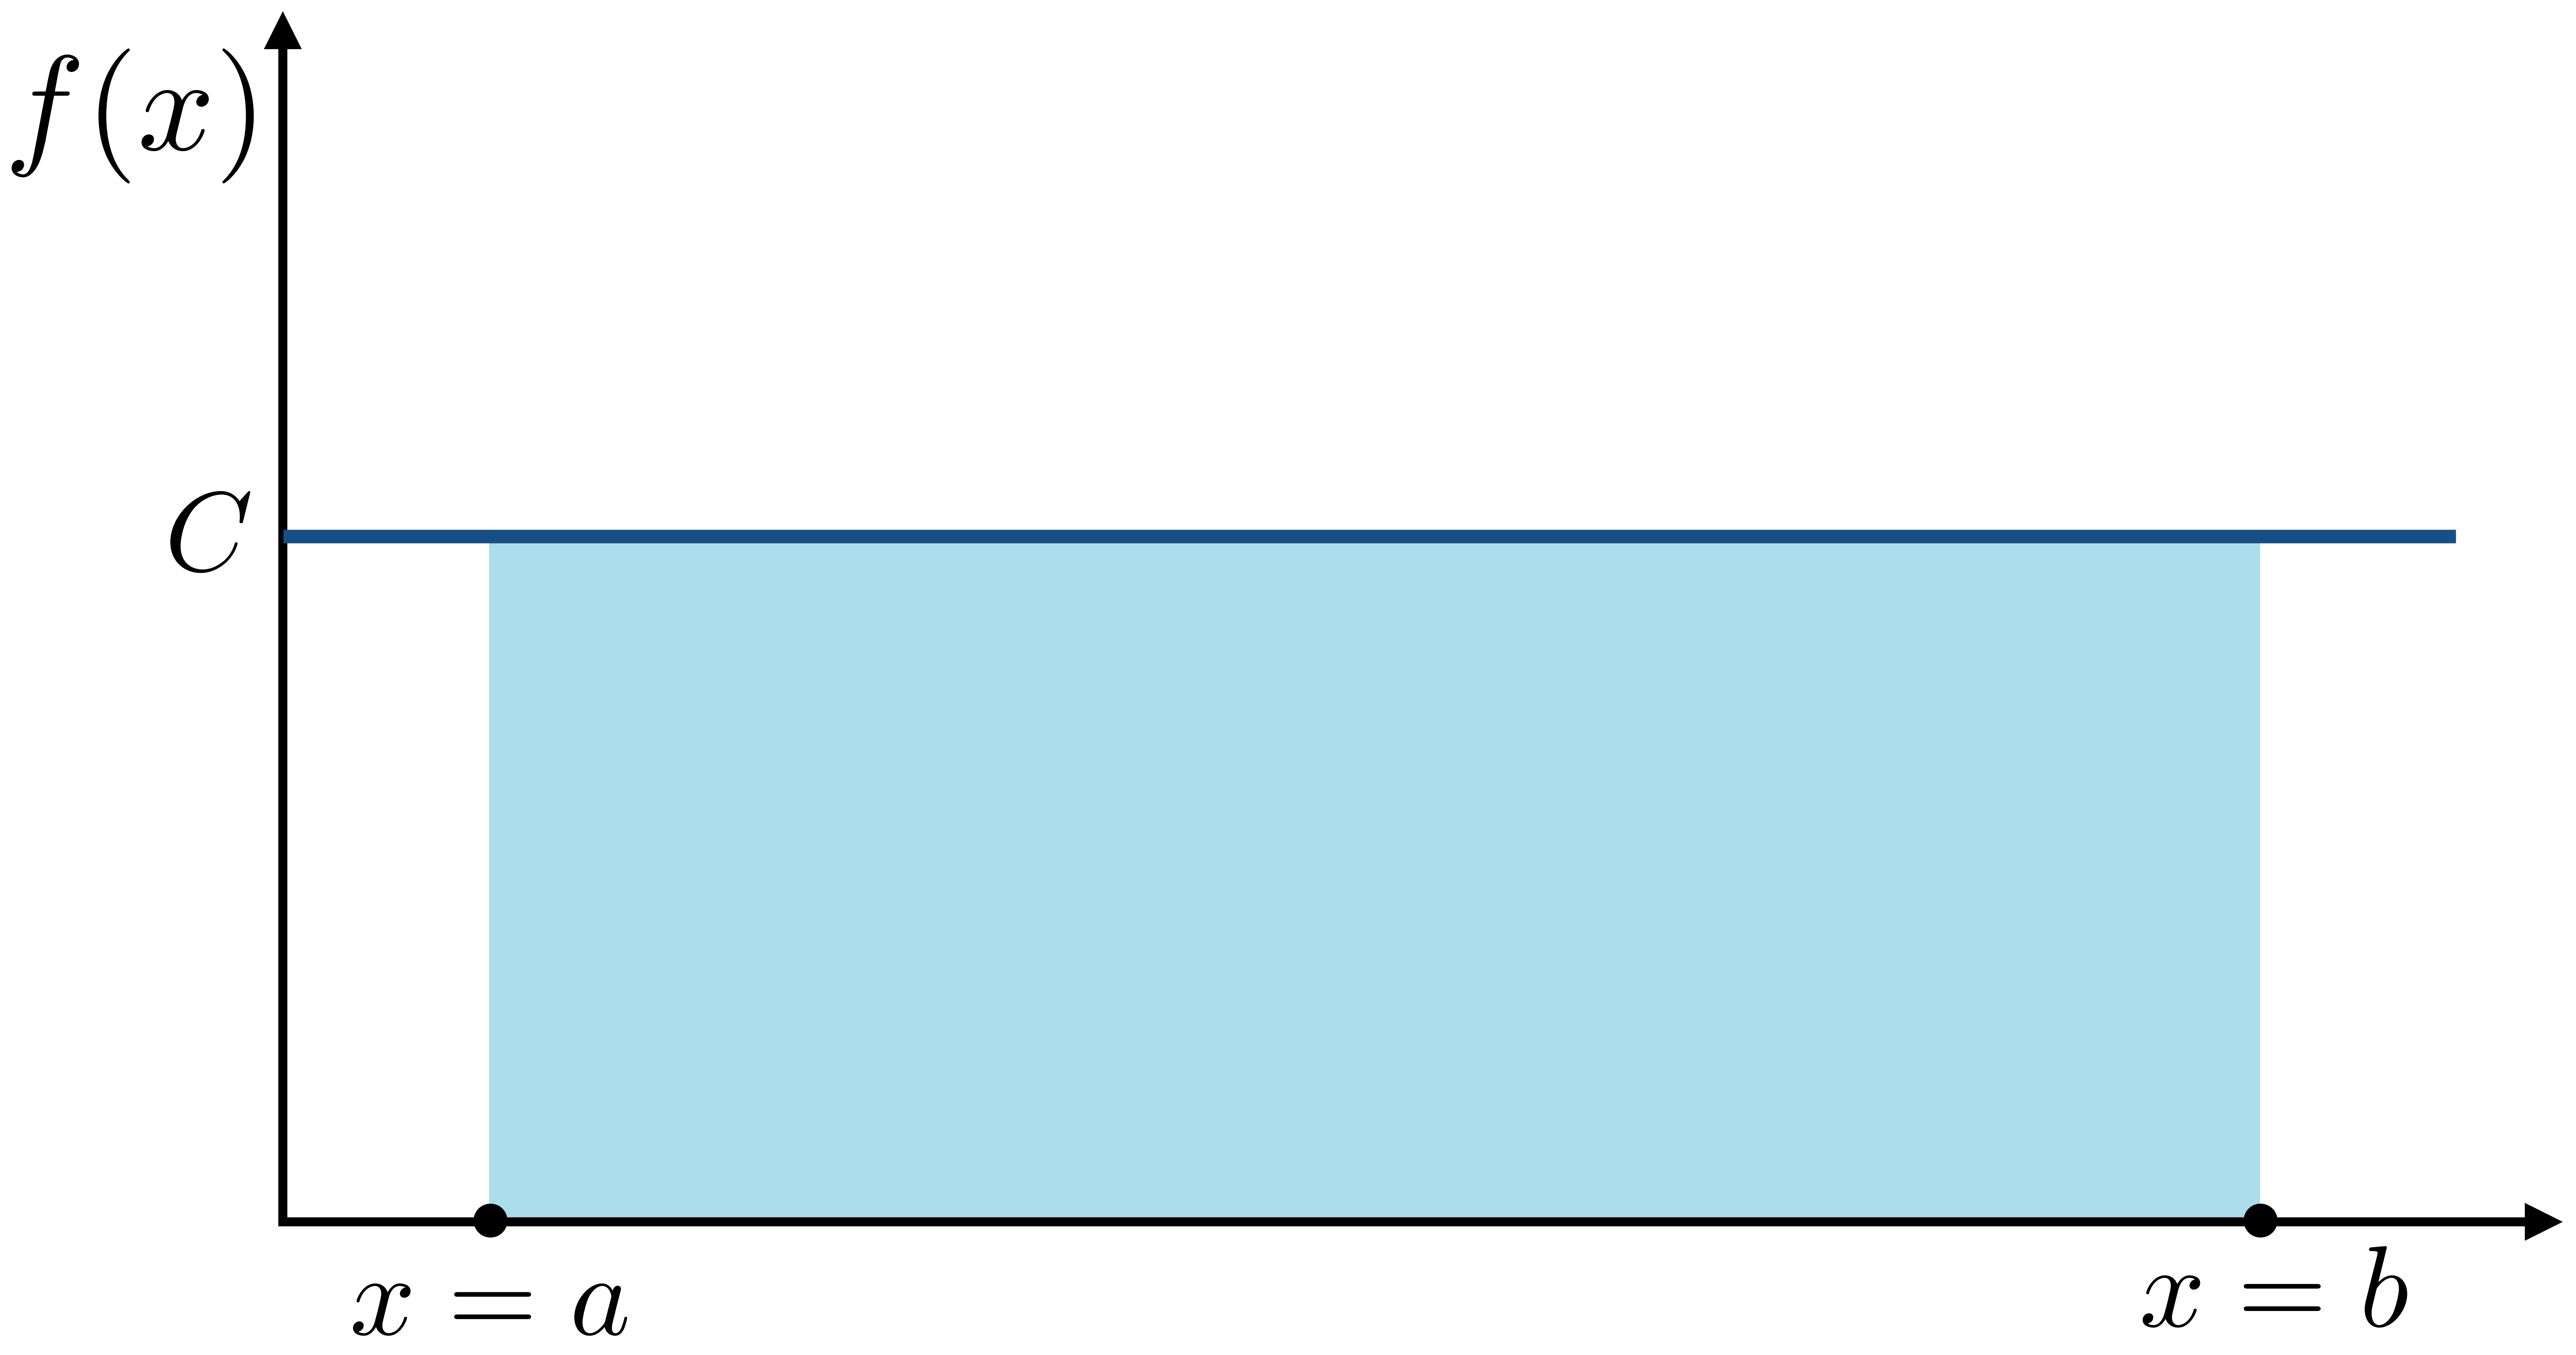

Affine function: $f(x) = cx + d$

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

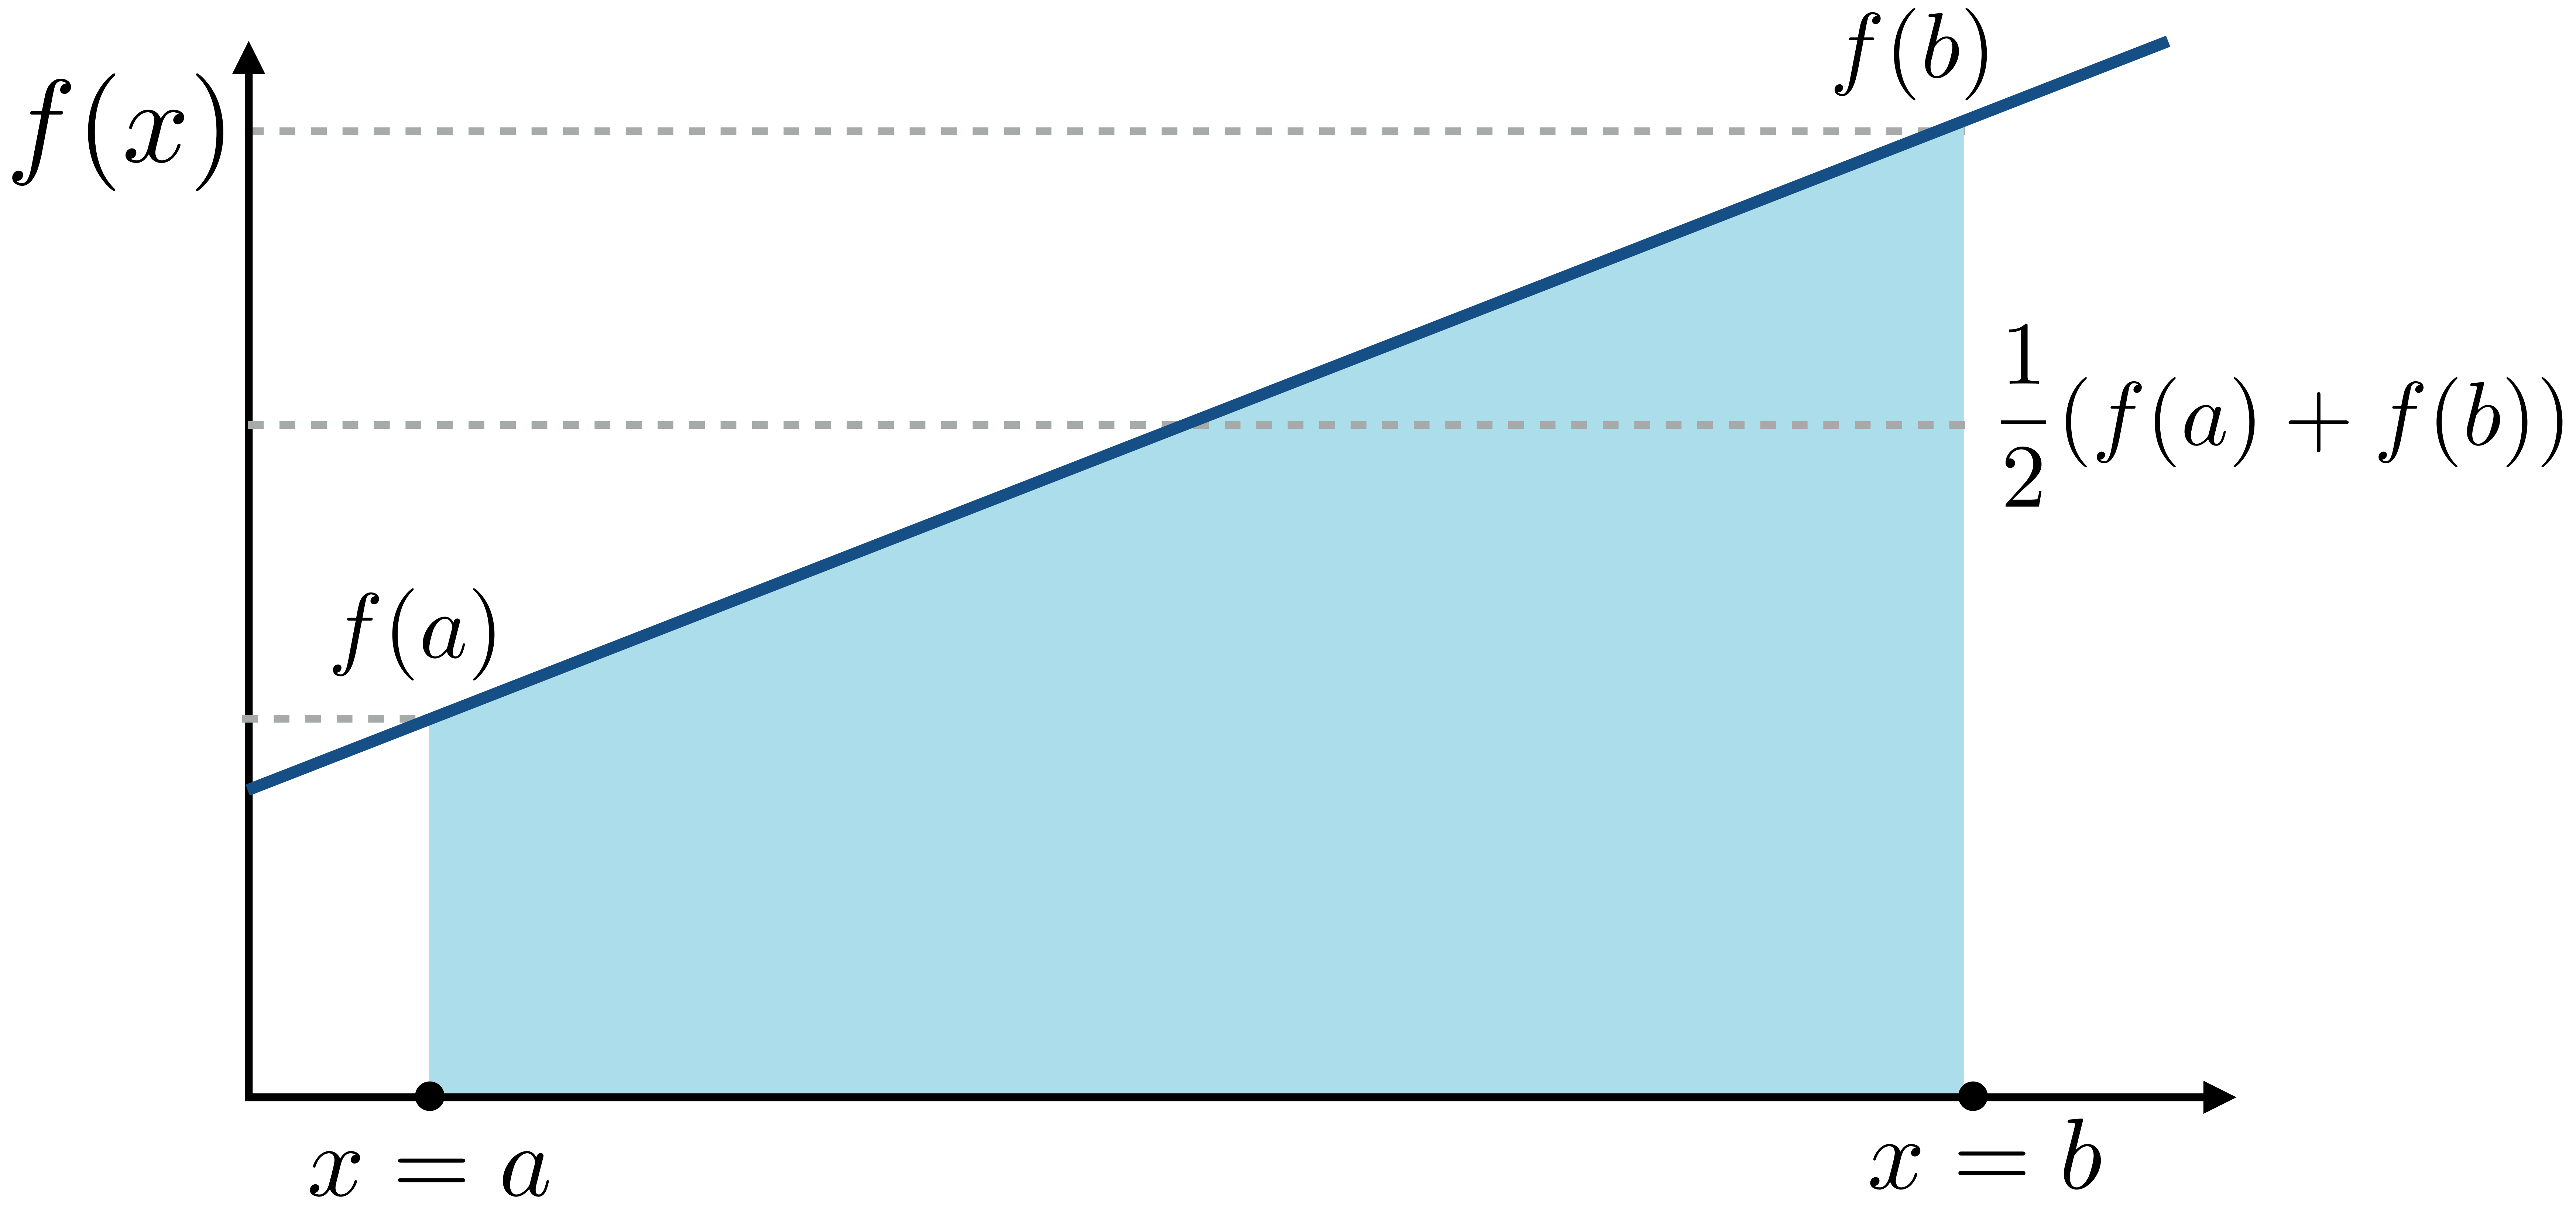

Piecewise affine function

Sum of integrals of individual affine components

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

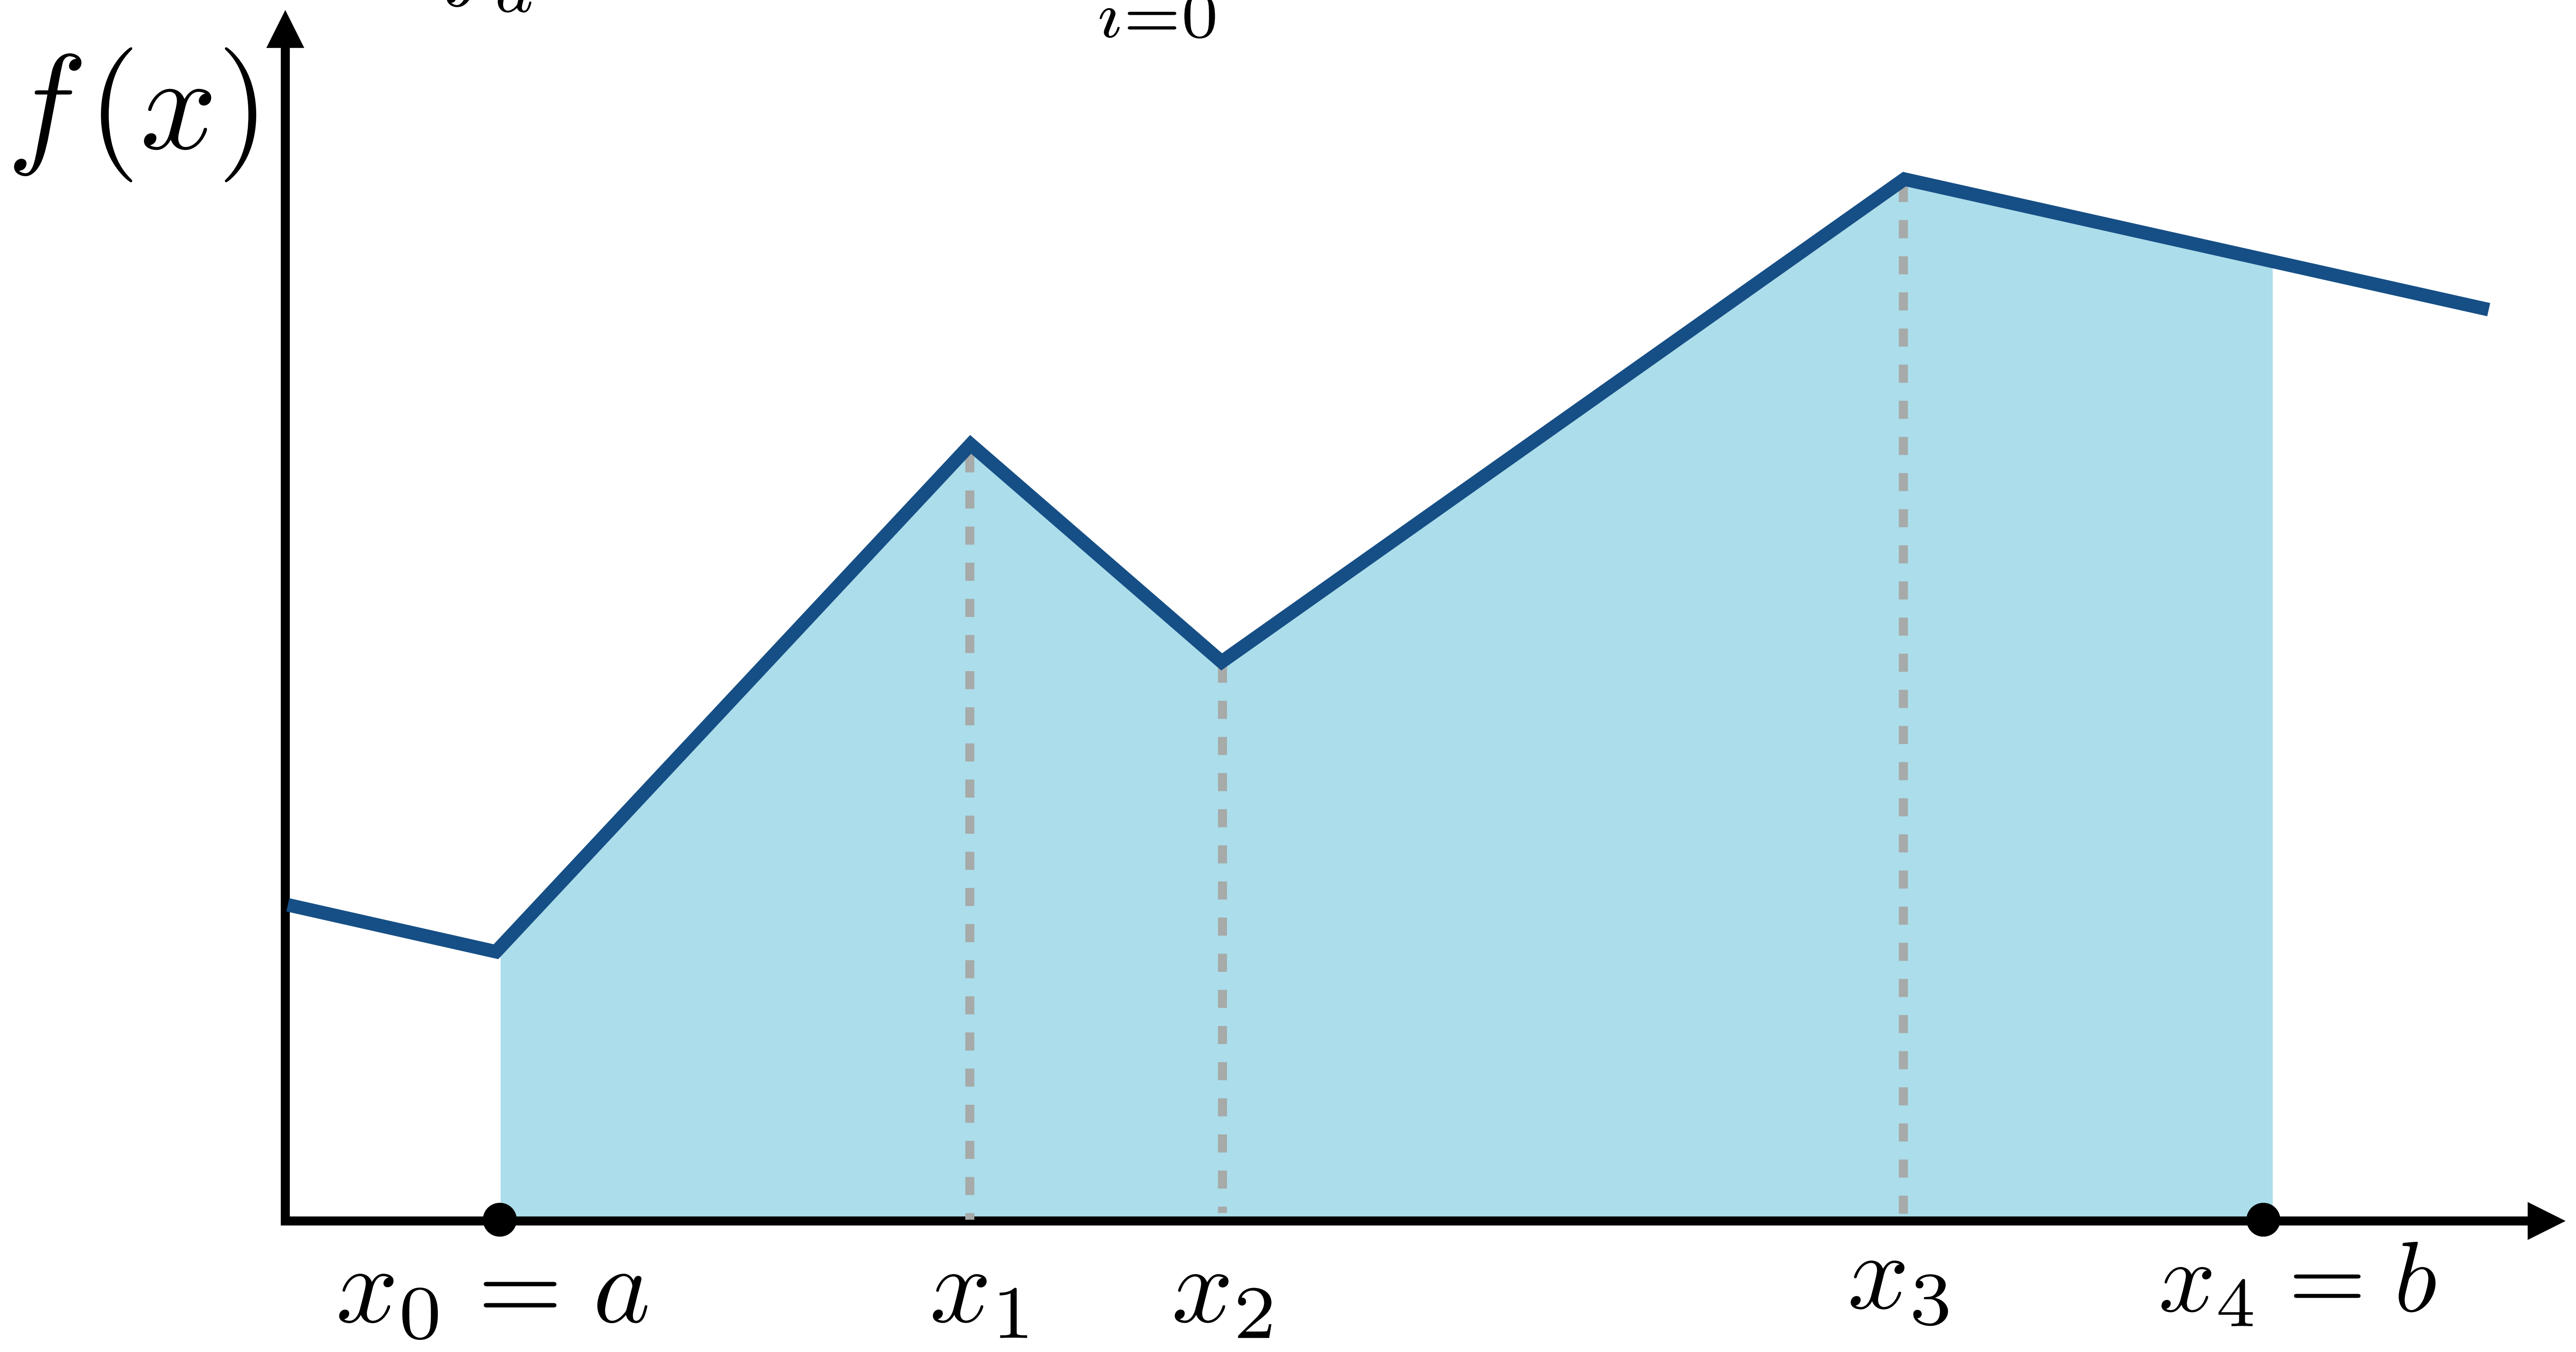

Piecewise affine function

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

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

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

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

Weighted combination of measurements.

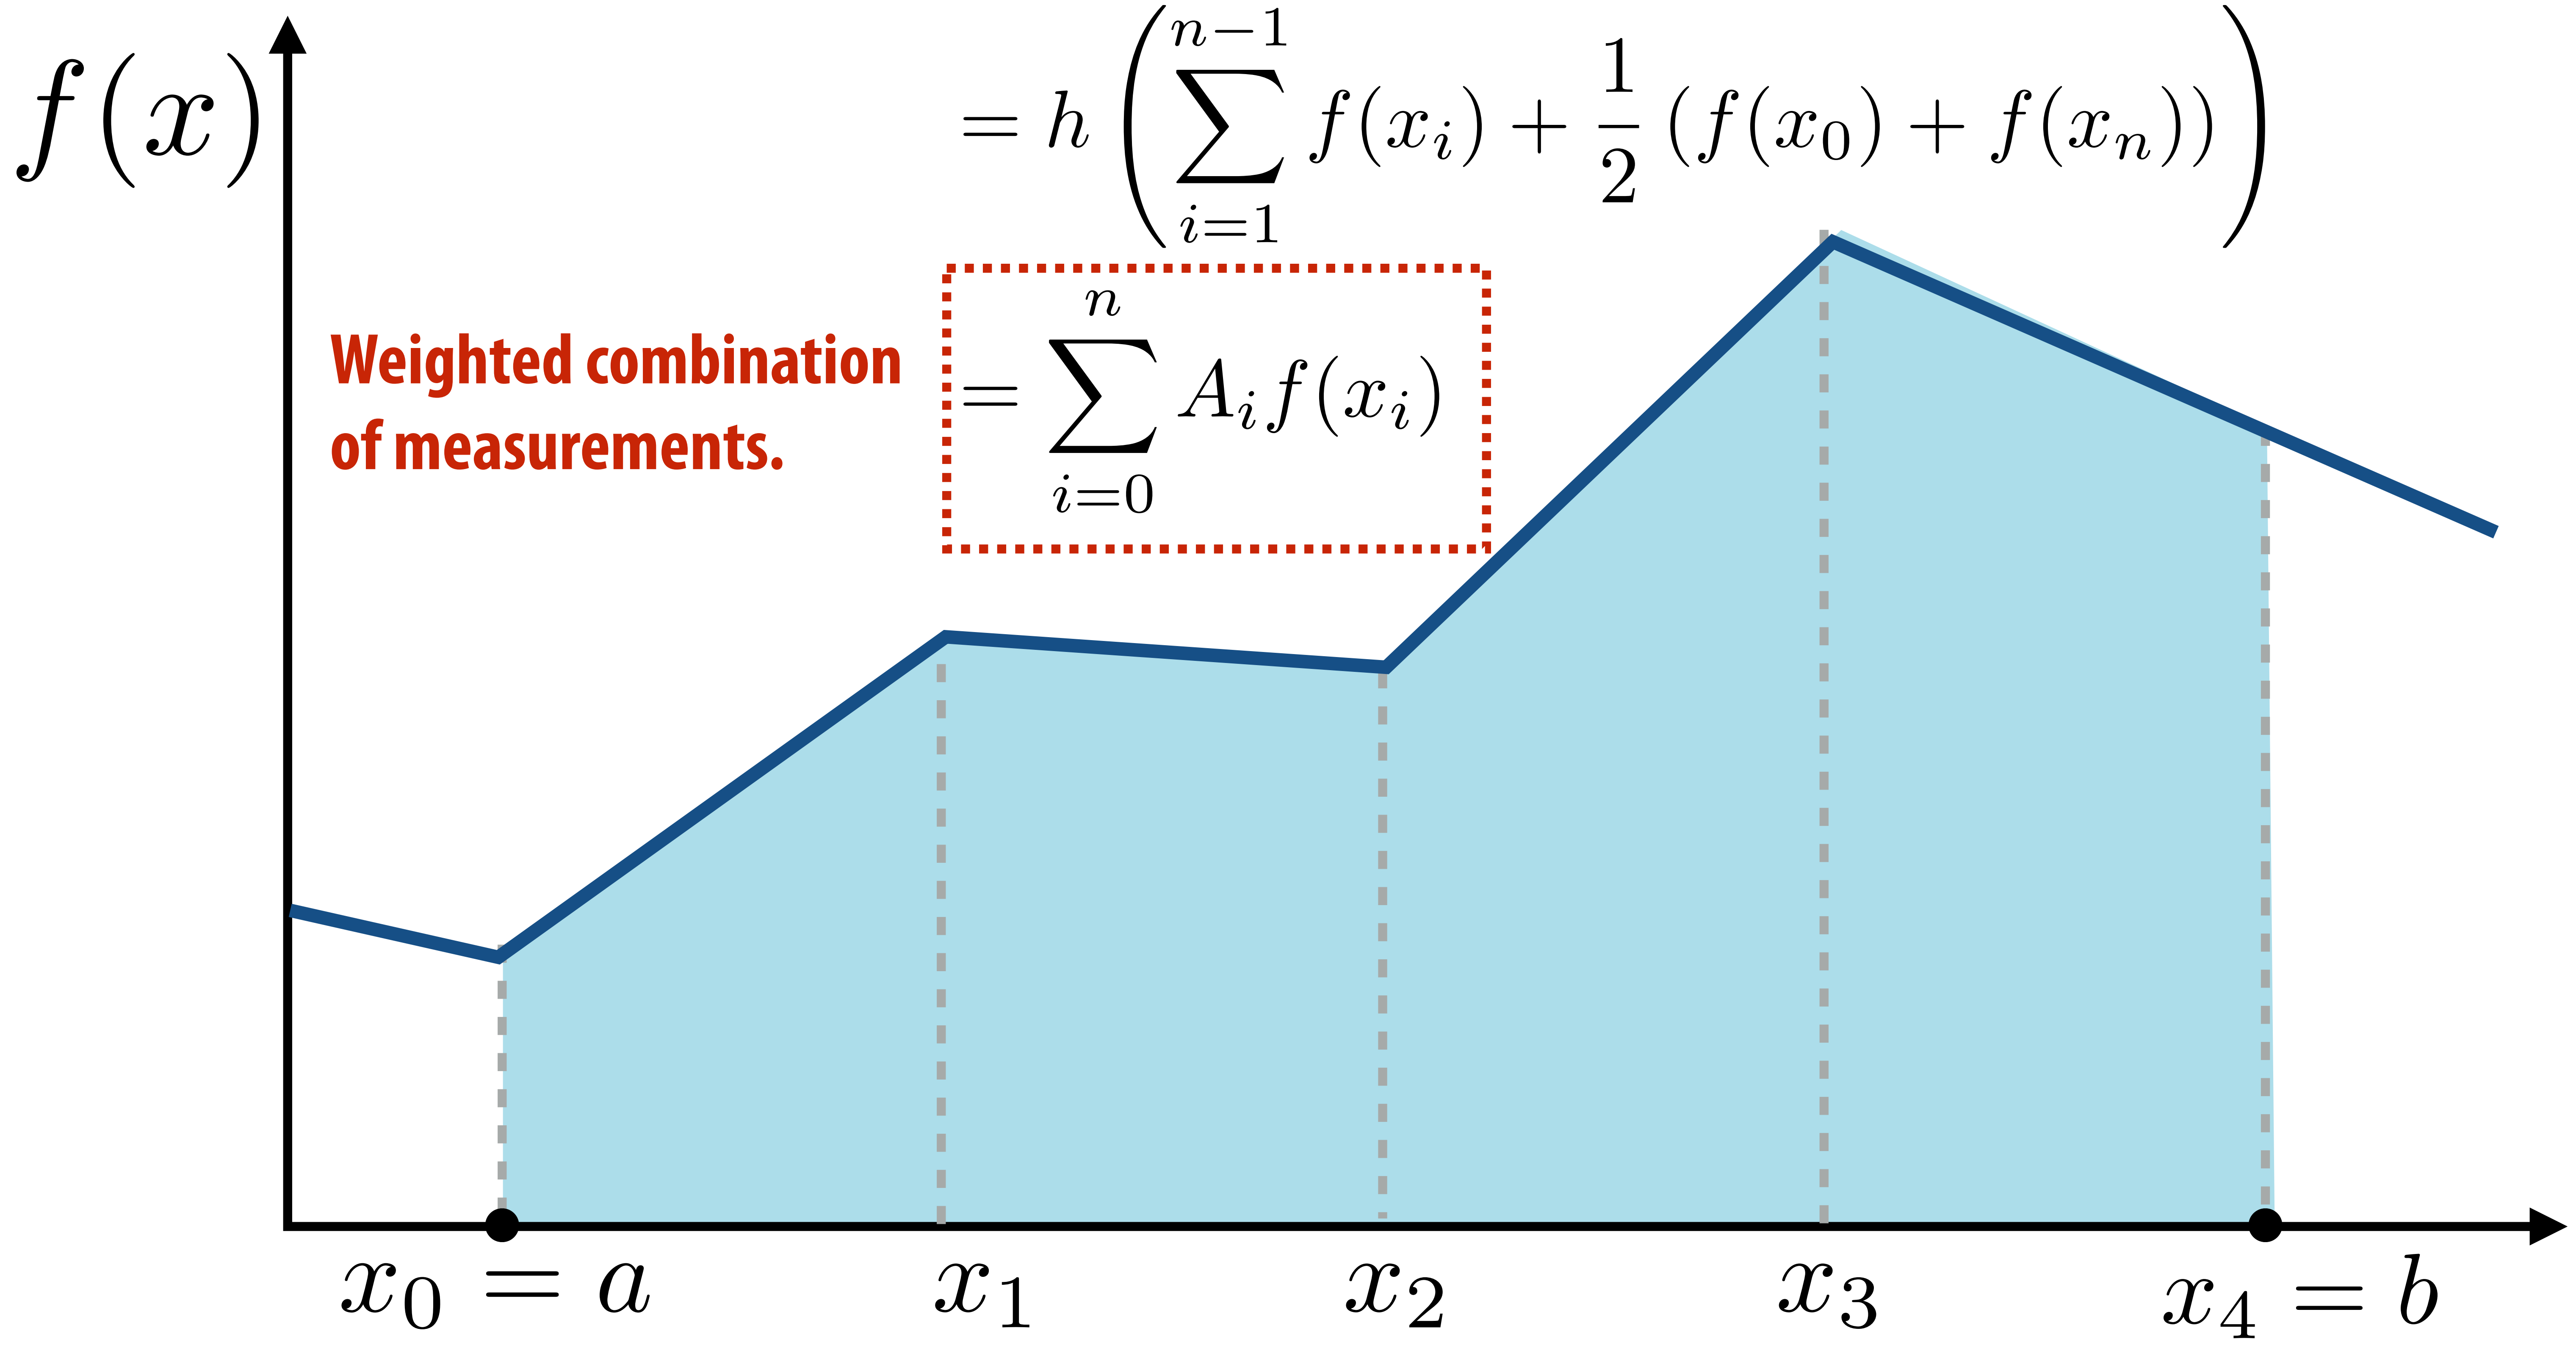

Arbitrary function f(x)?

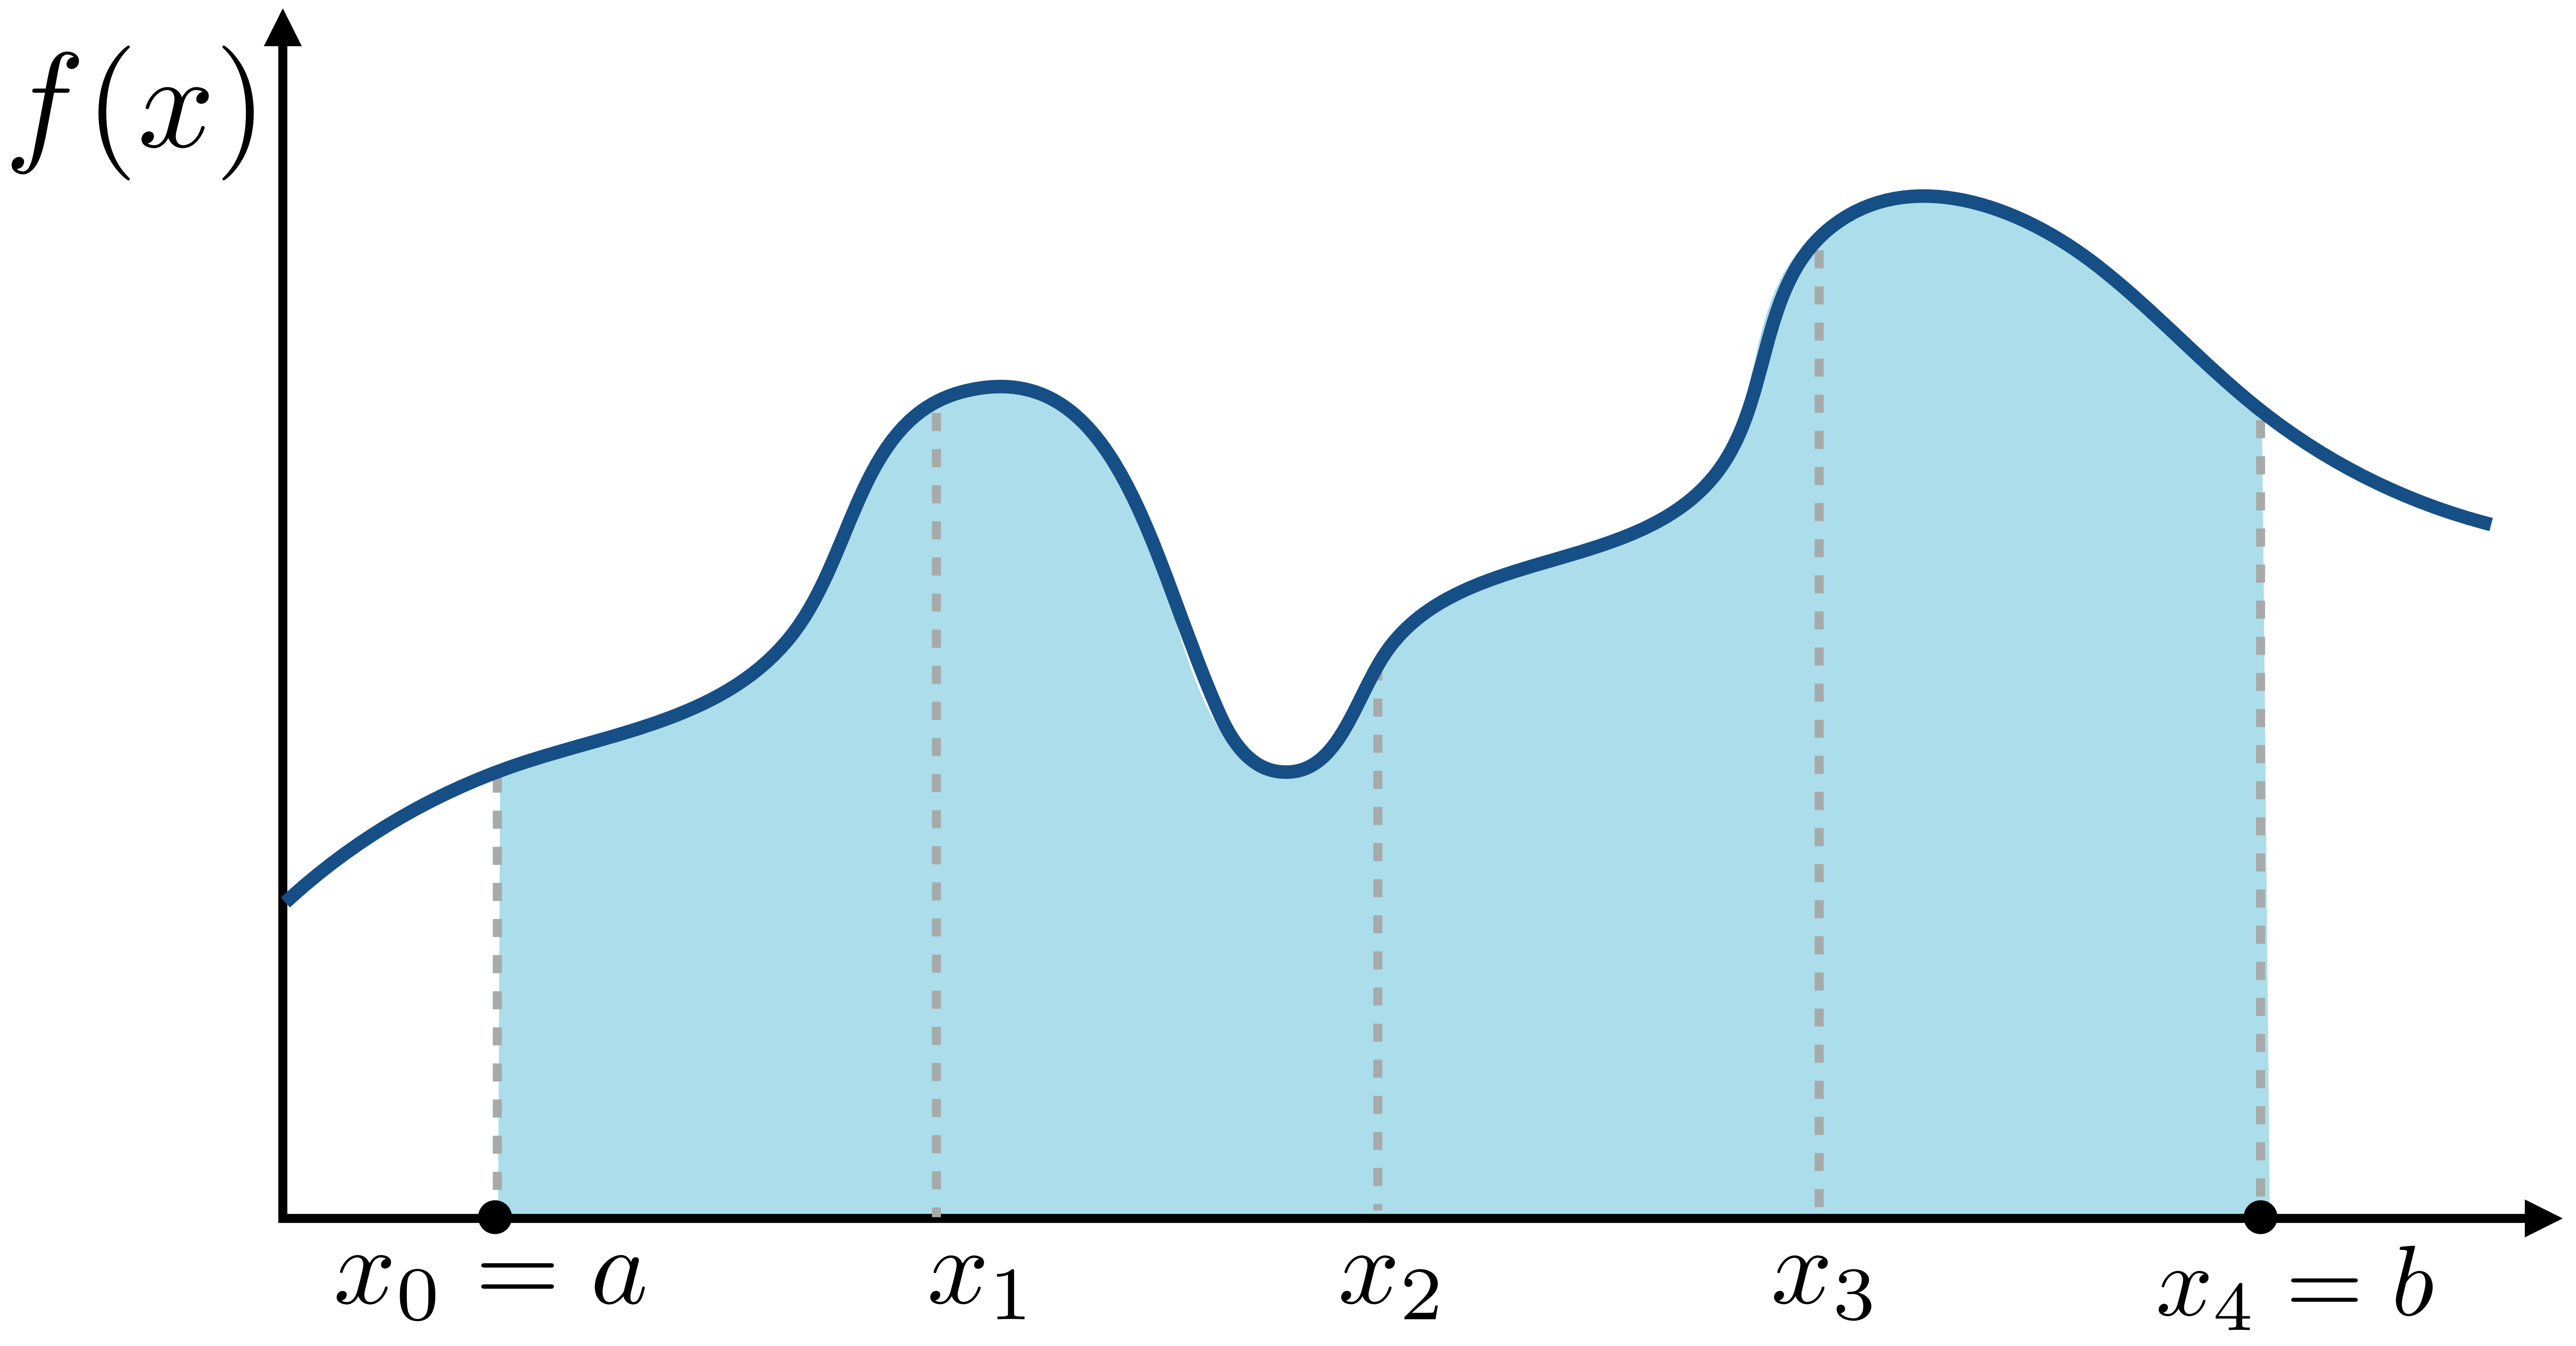

Trapezoidal rule

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

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

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

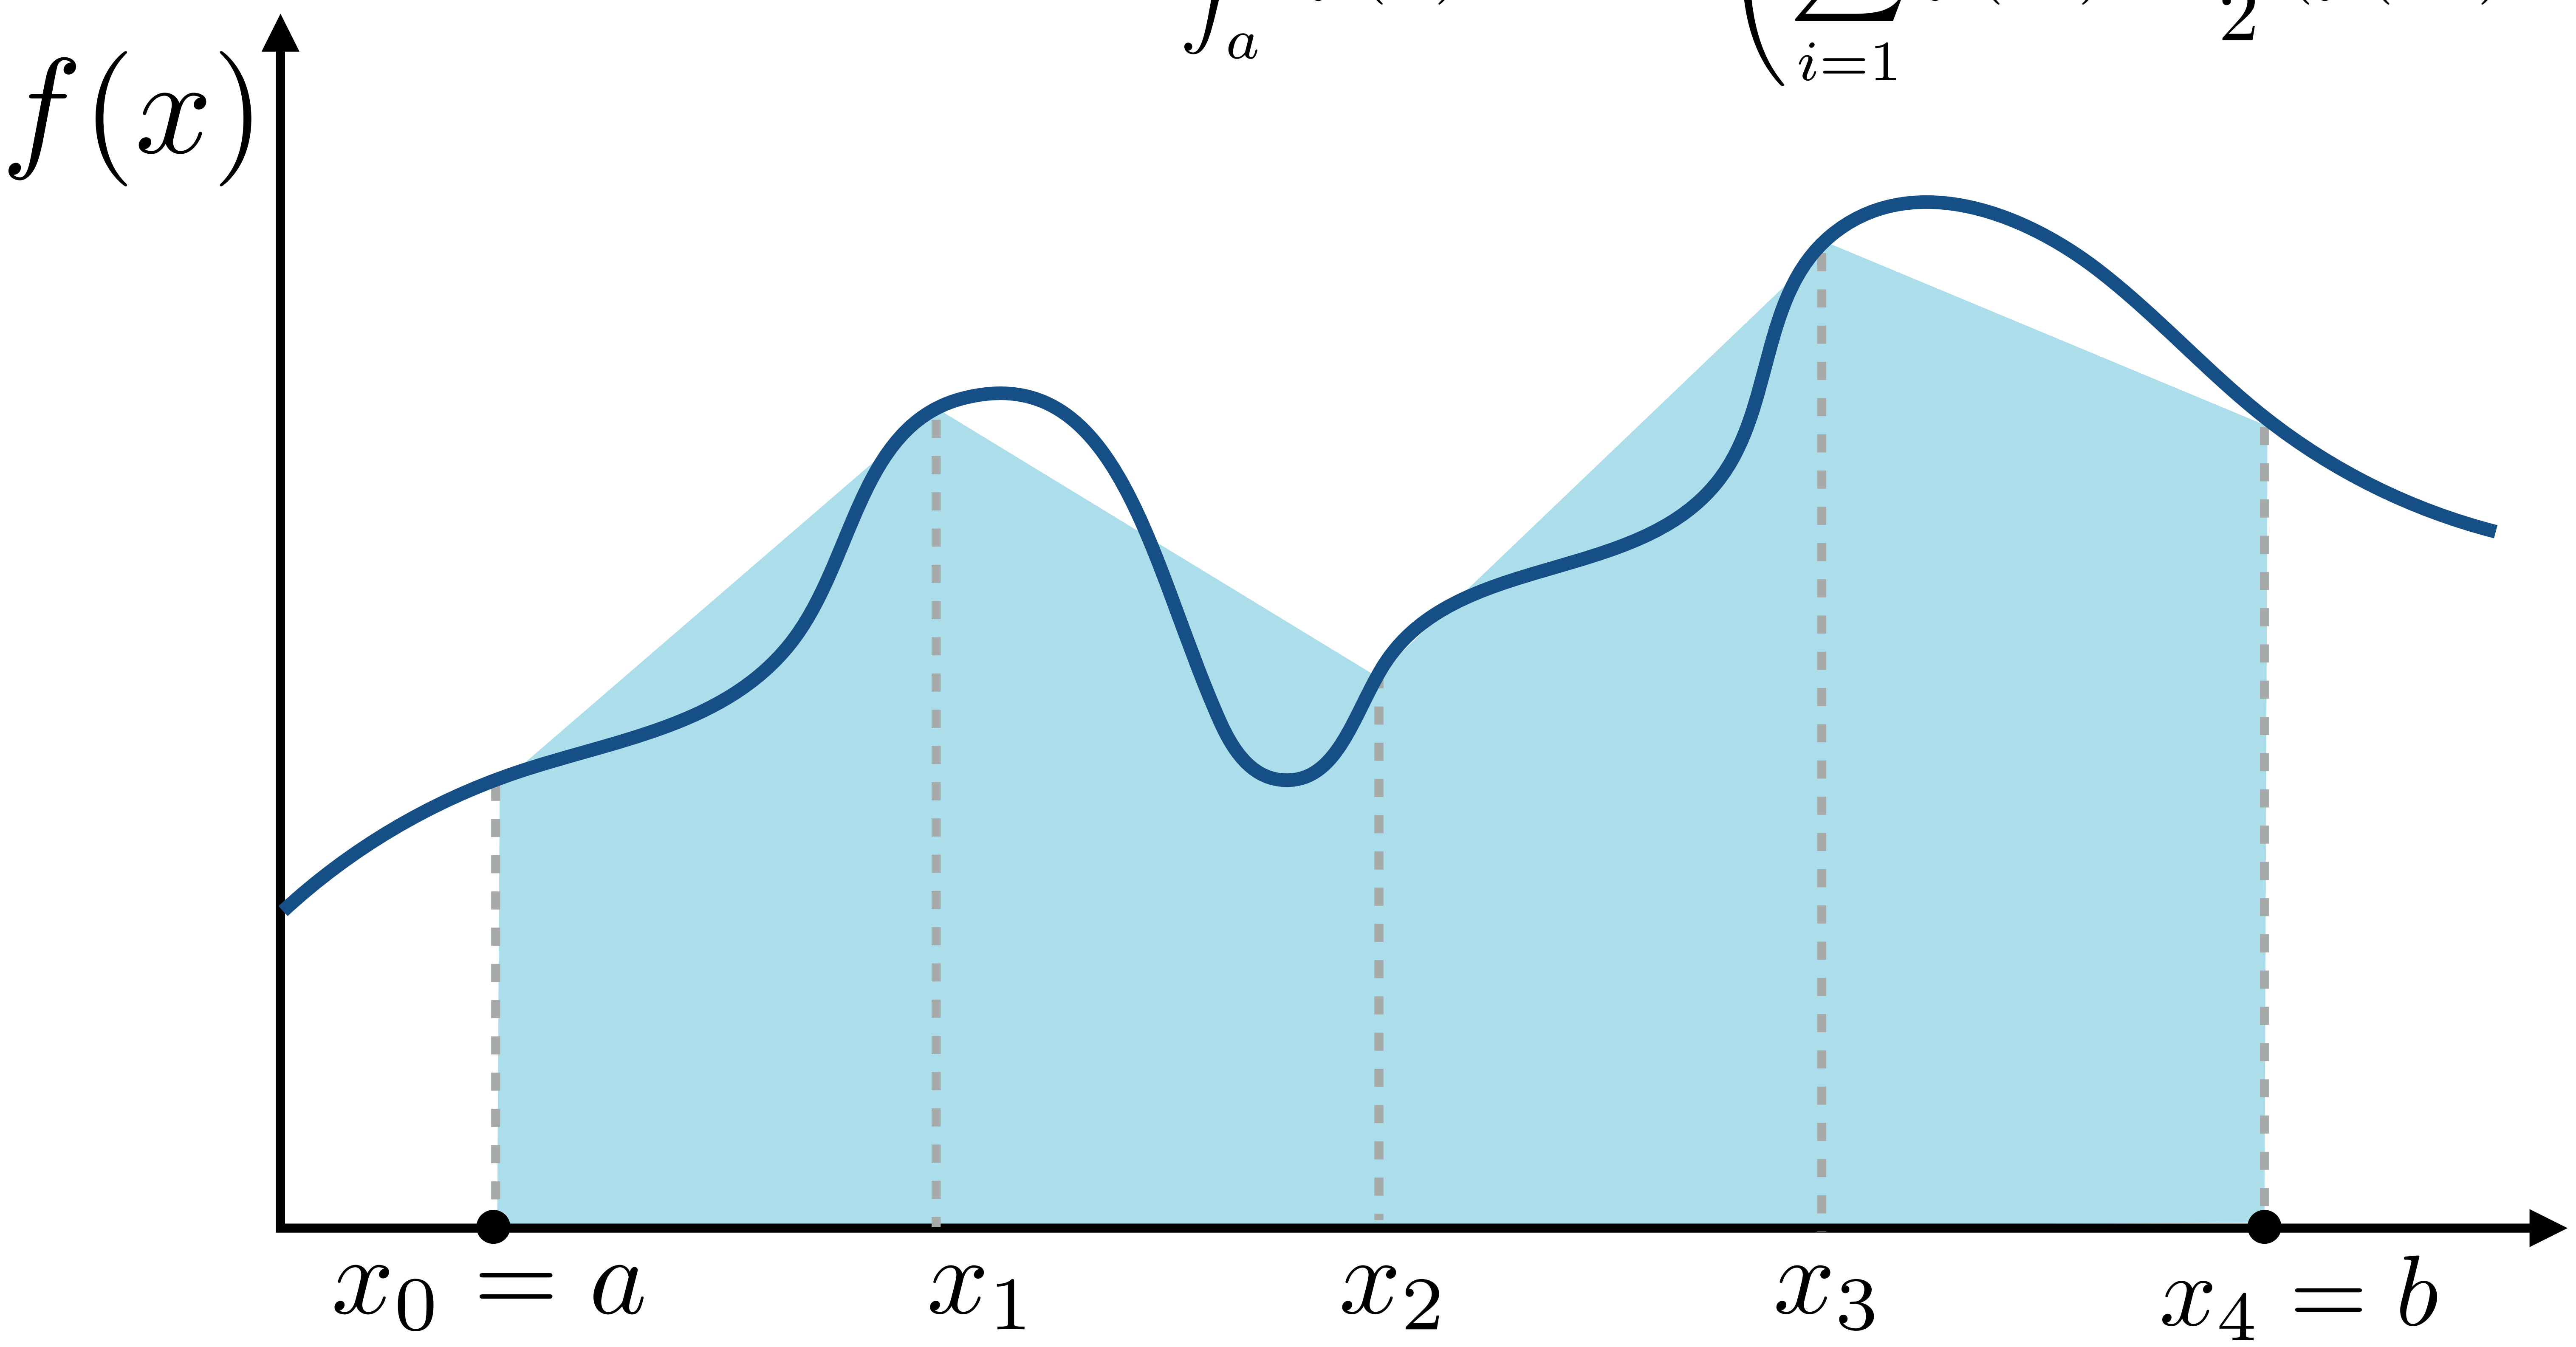

Trapezoidal rule

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)**

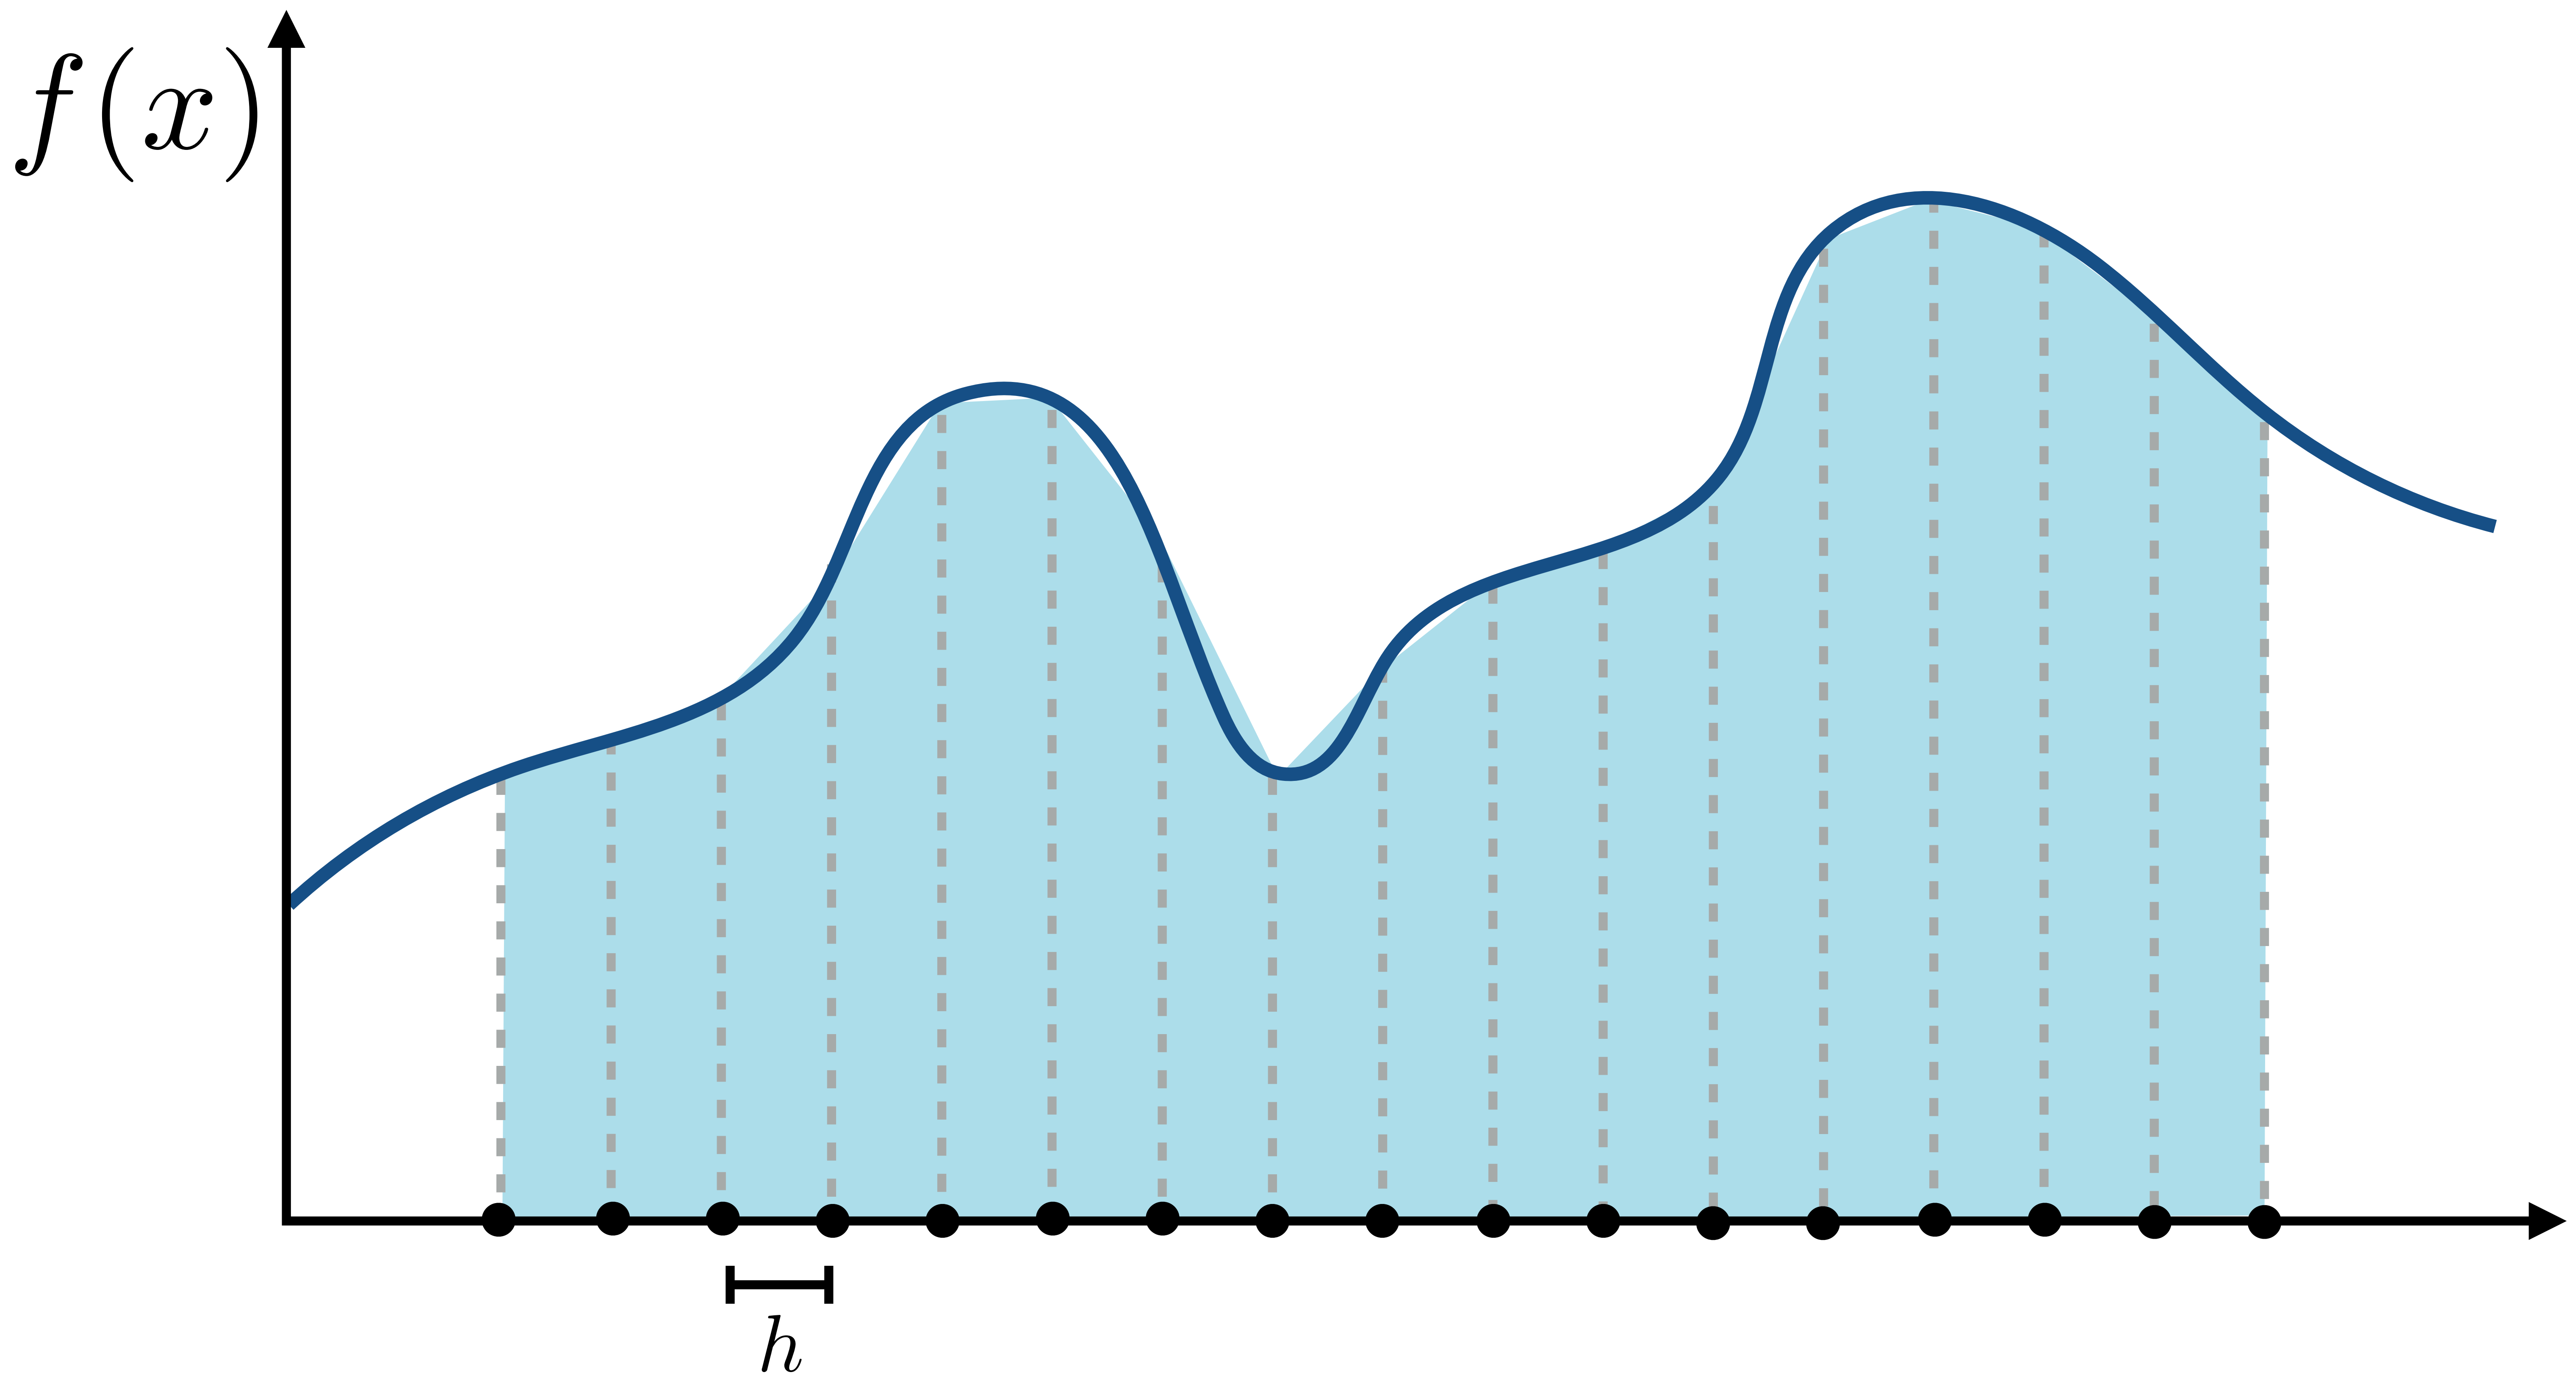

Integration in 2D

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

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

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

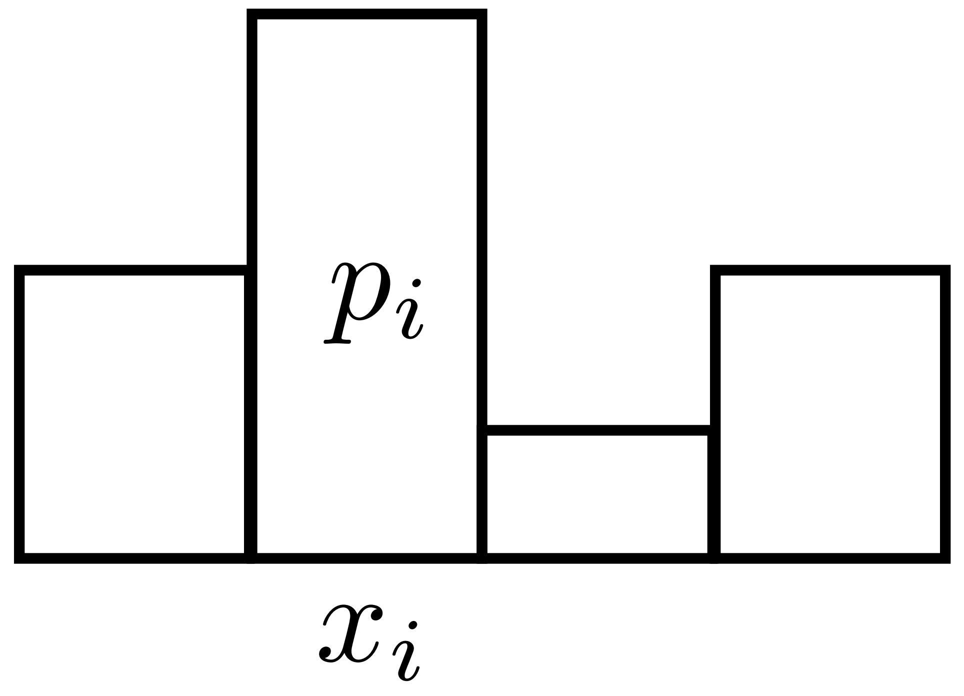

Requirements of a PDF:

$$p_i \geq 0$$

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

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:

$$0 \le P_i \le 1$$

$$P_n = 1$$

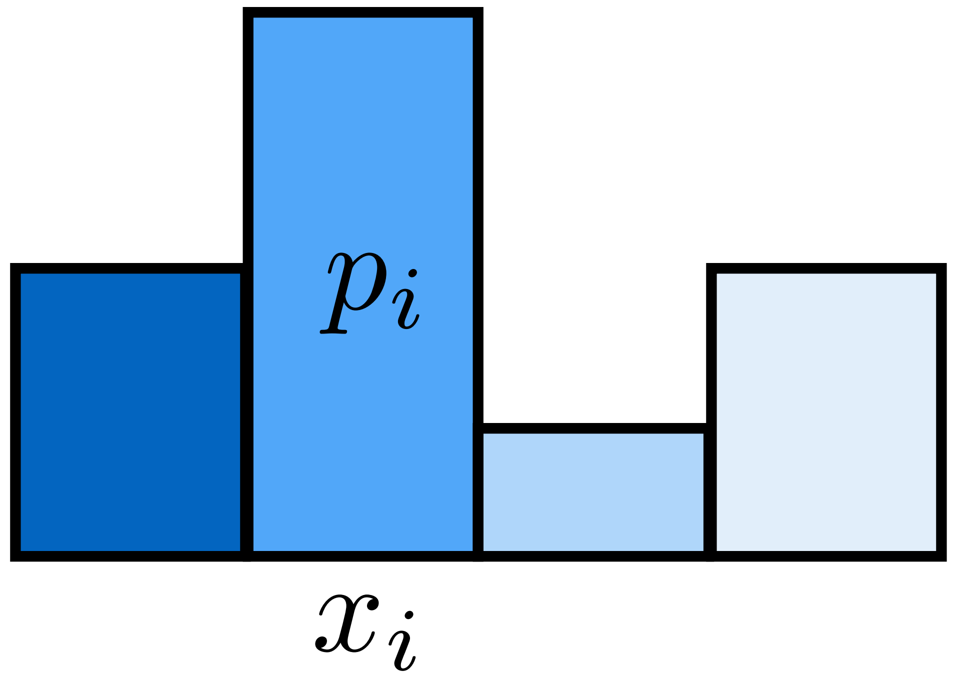

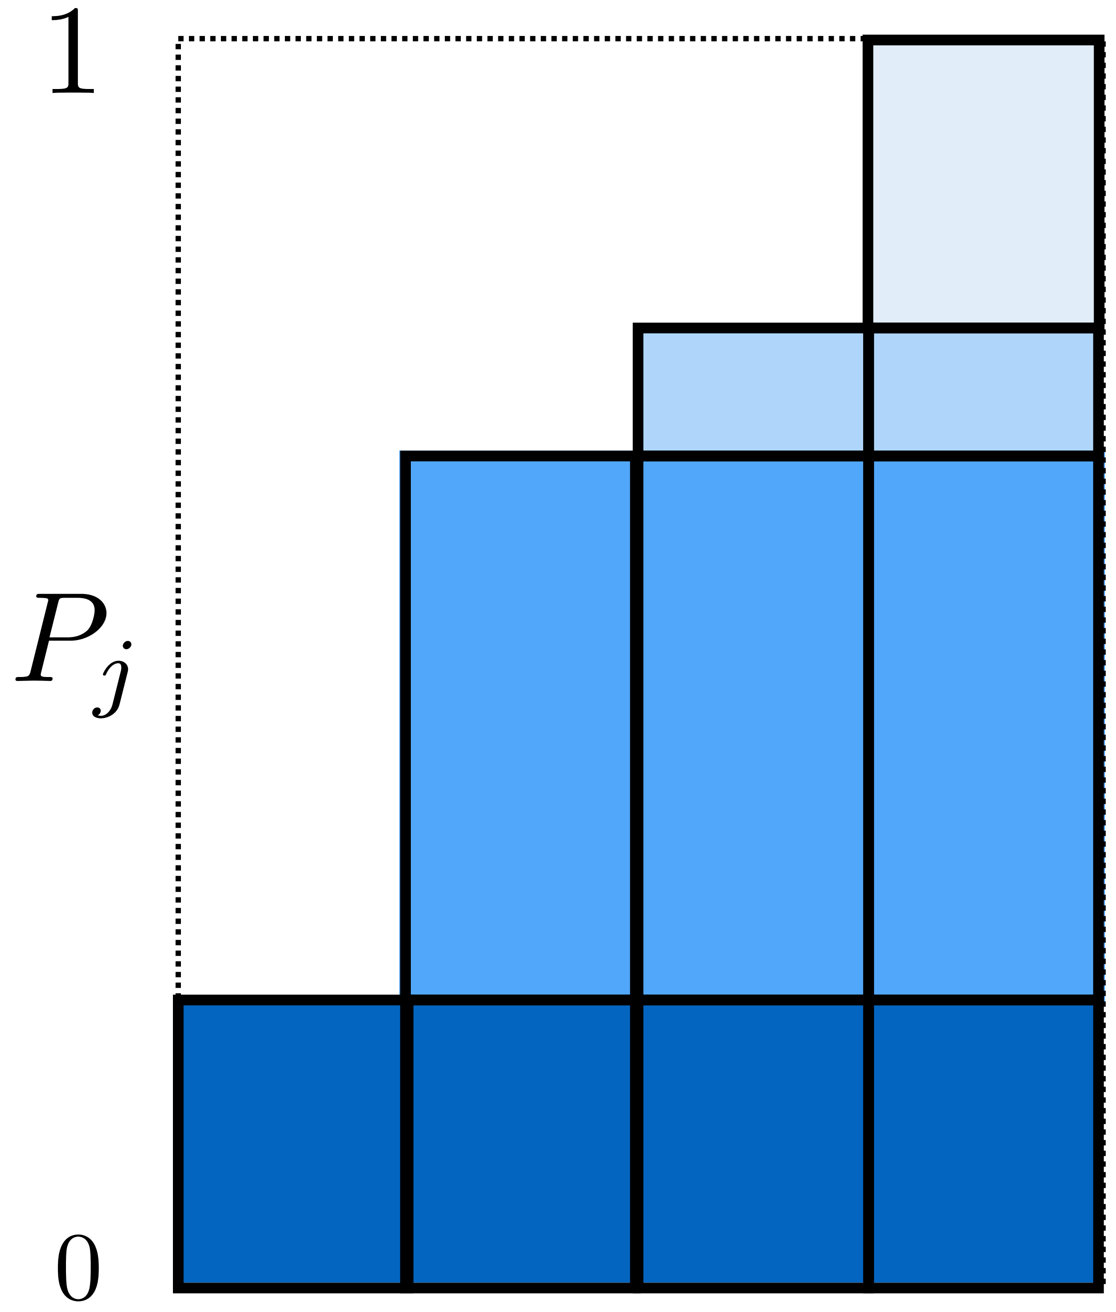

Sampling from discrete probability distributions

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

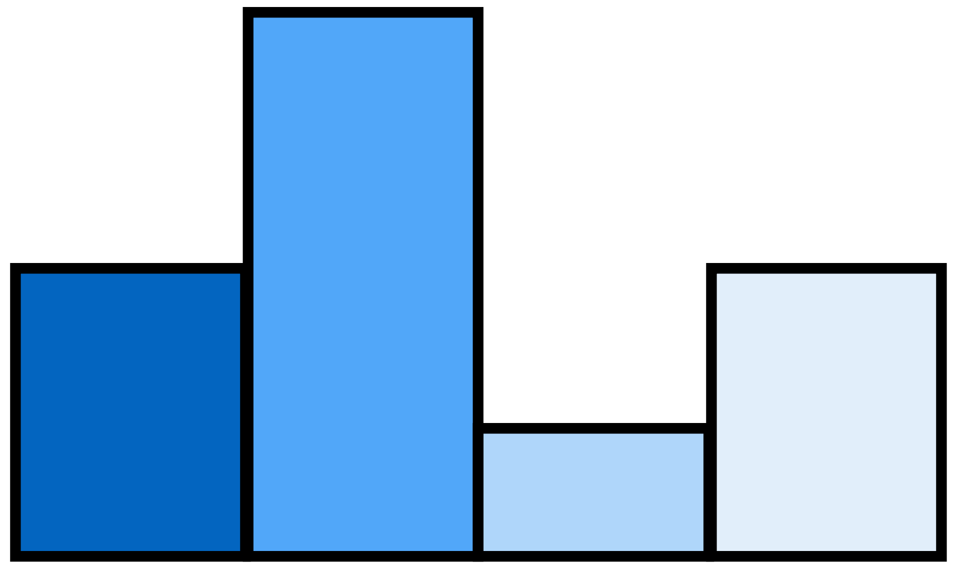

To randomly select an event, select x_i if

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

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

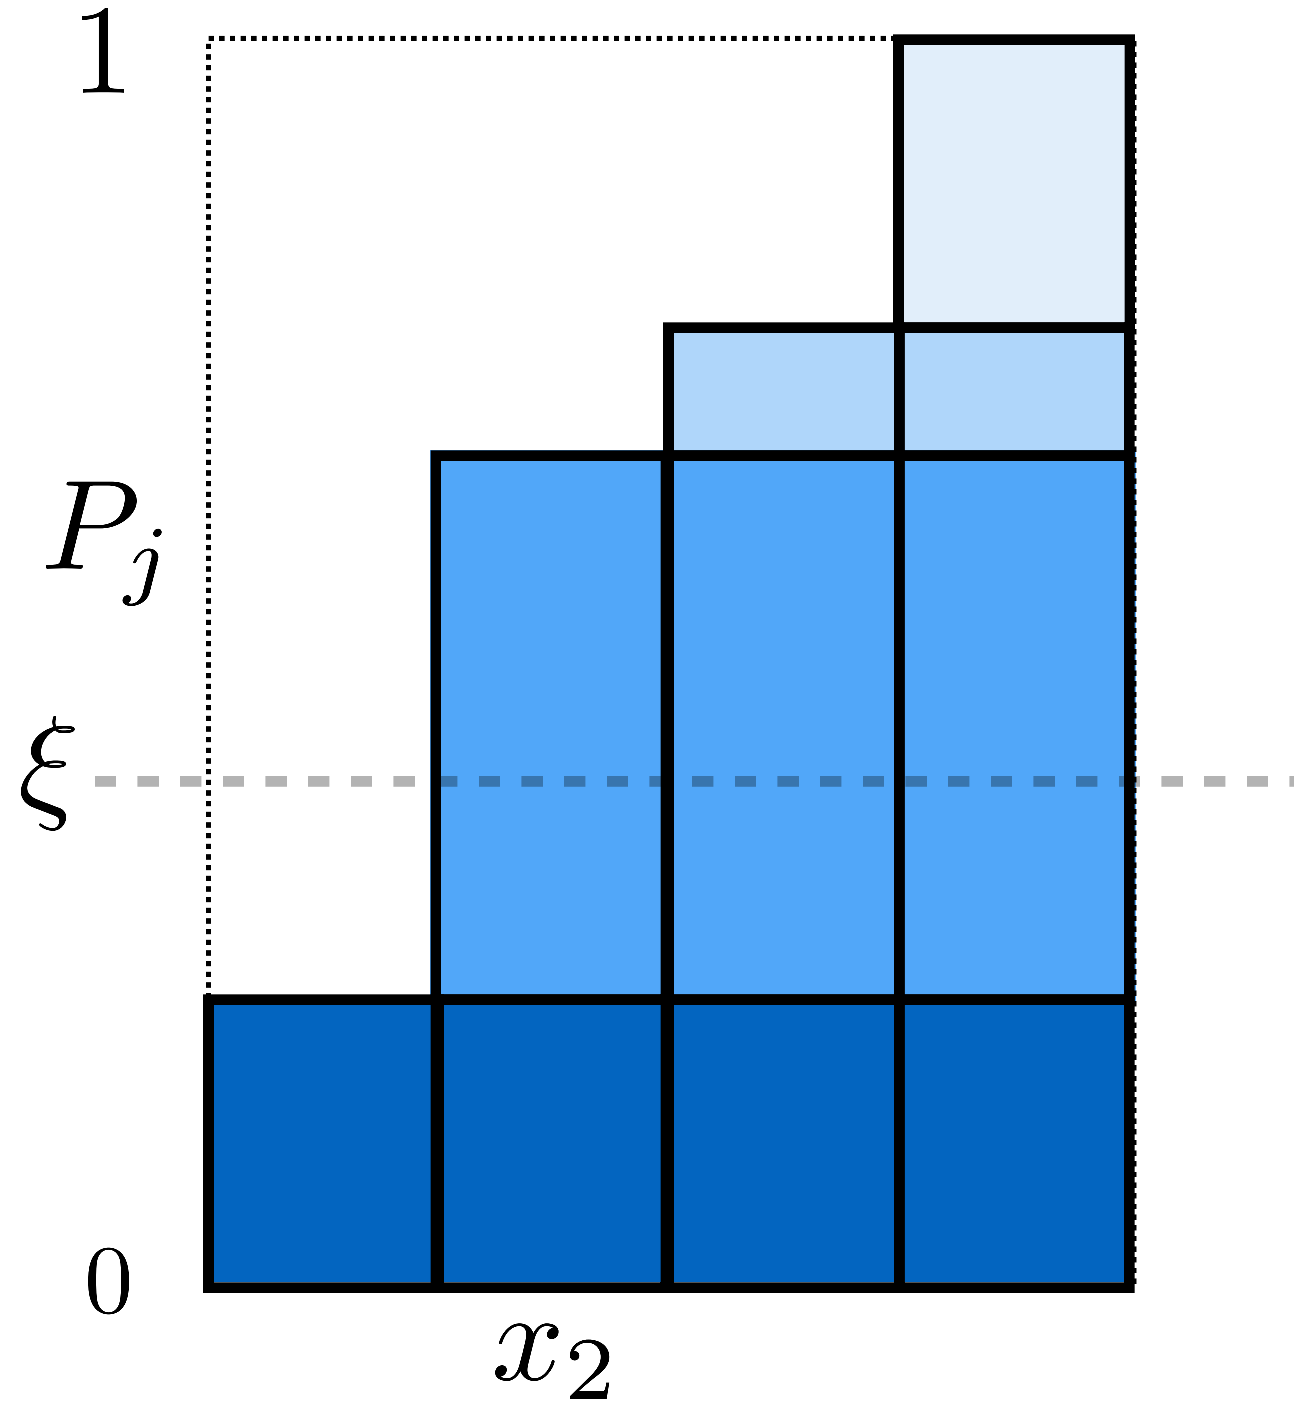

Continuous probability distributions

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)

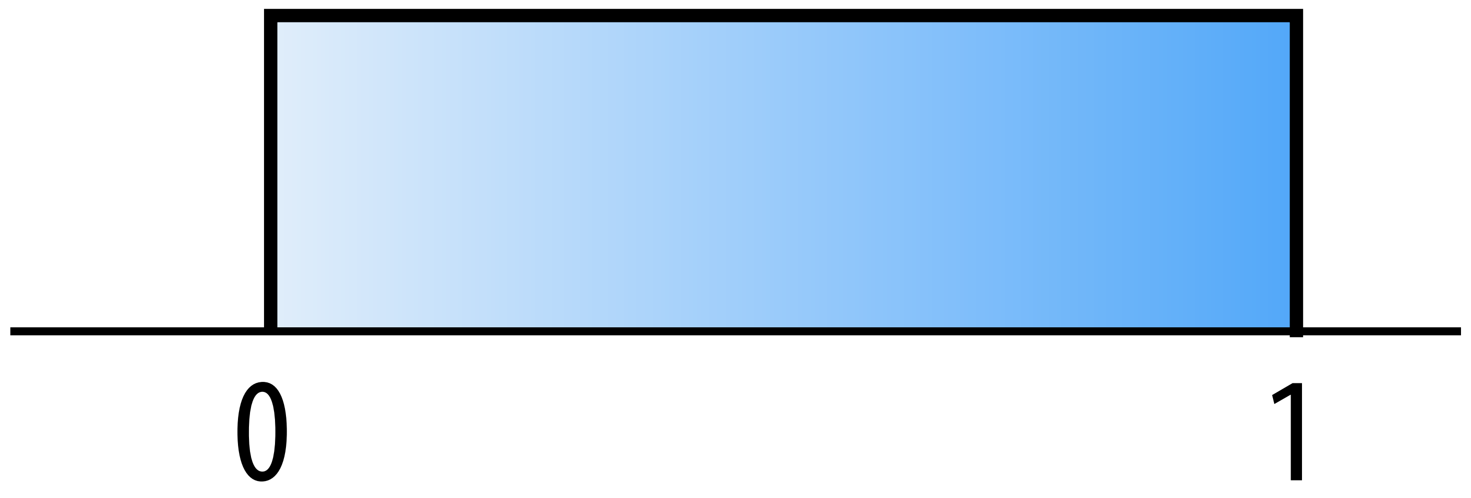

CDF P(x)

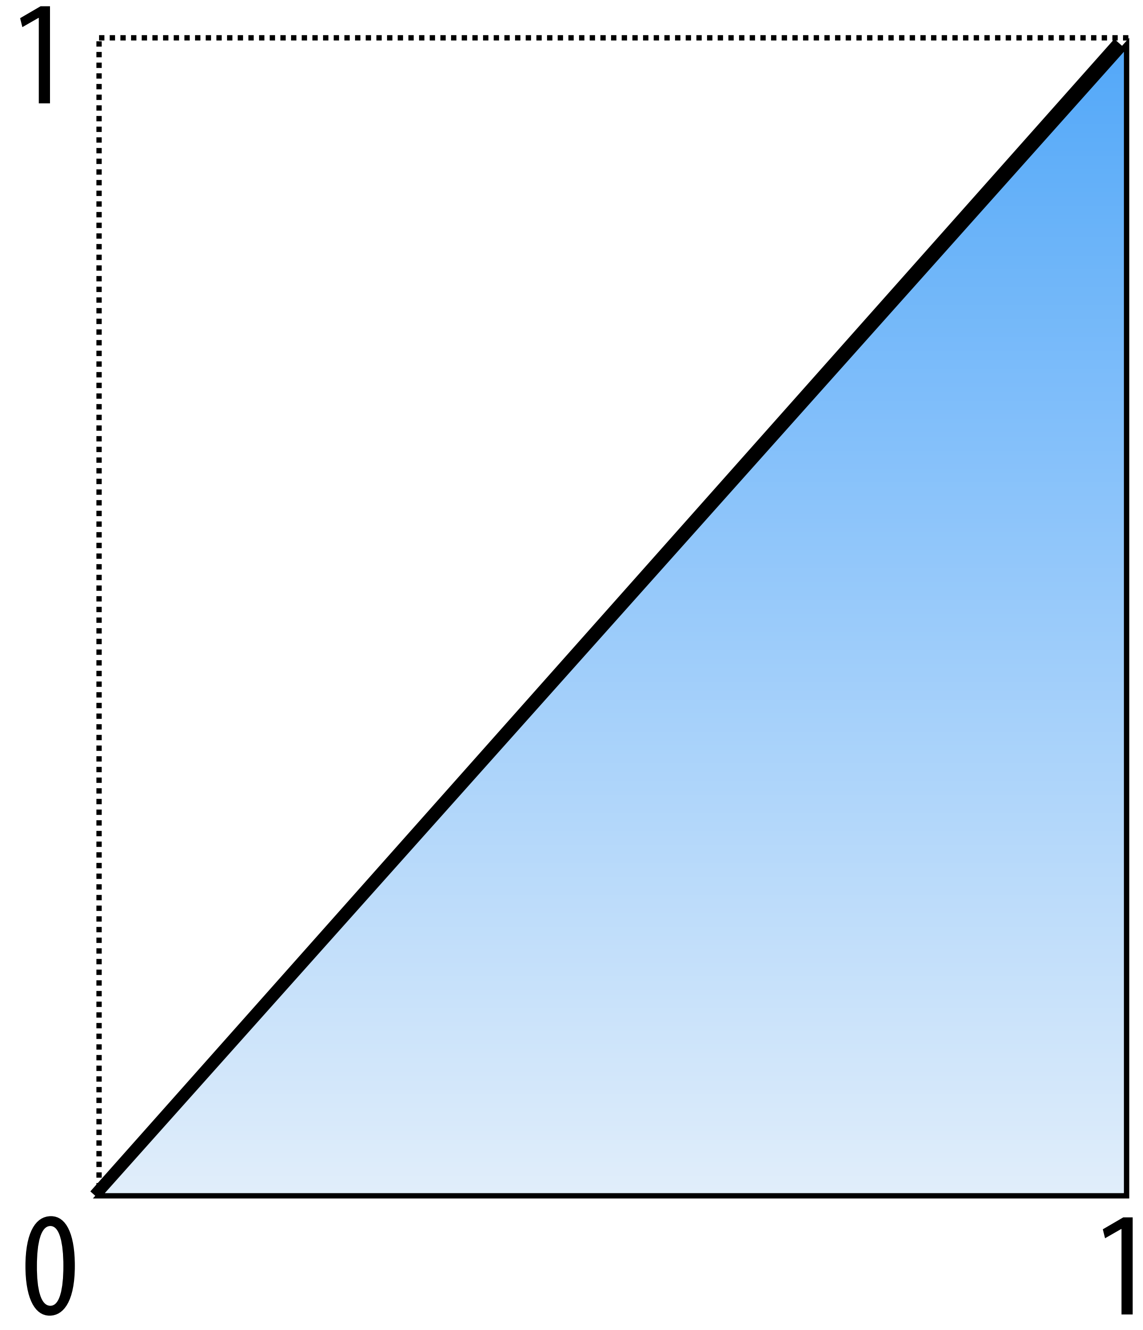

Sampling continuous random variables using the inversion method

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)

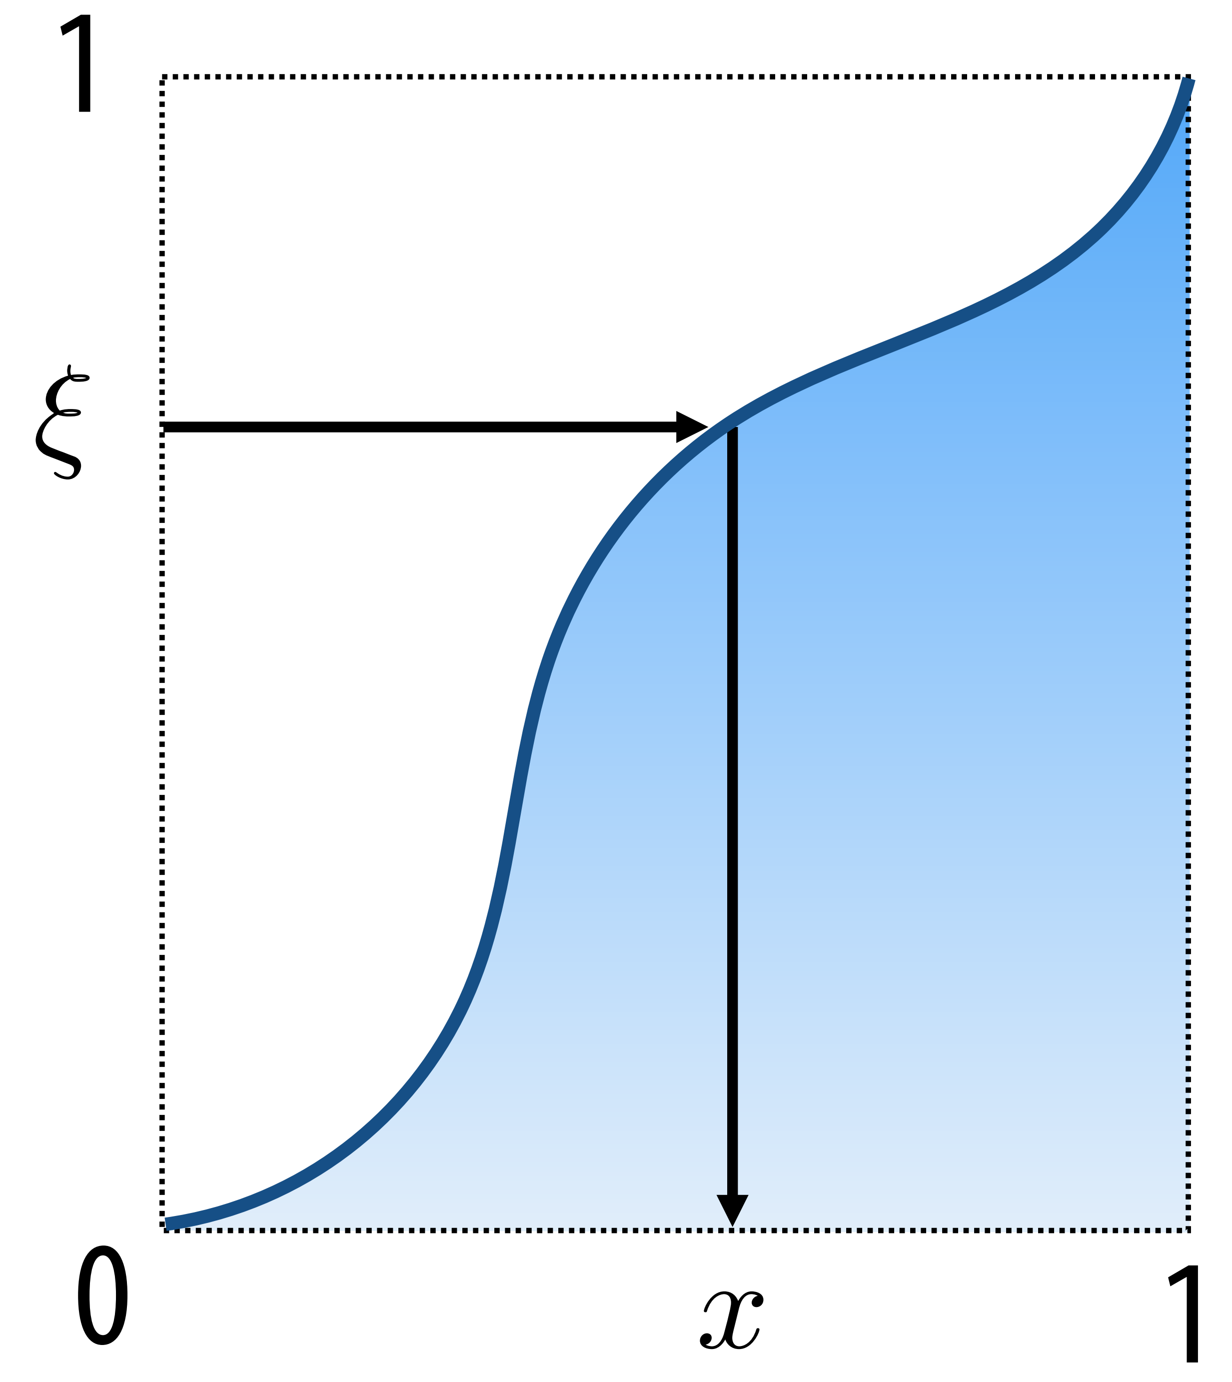

Example: applying the inversion method

Given:

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

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

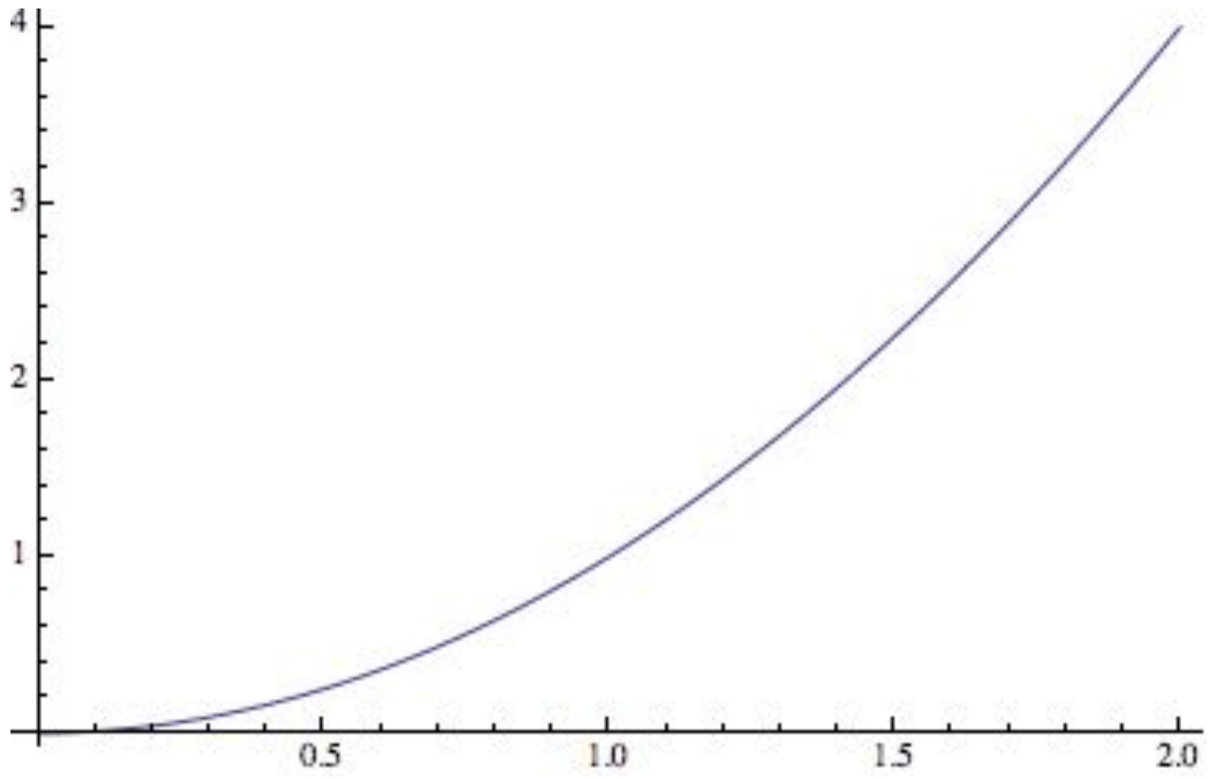

Compute PDF from f(x):

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

Probability density function (integrates to 1)

Example: applying the inversion method

Given:

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

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

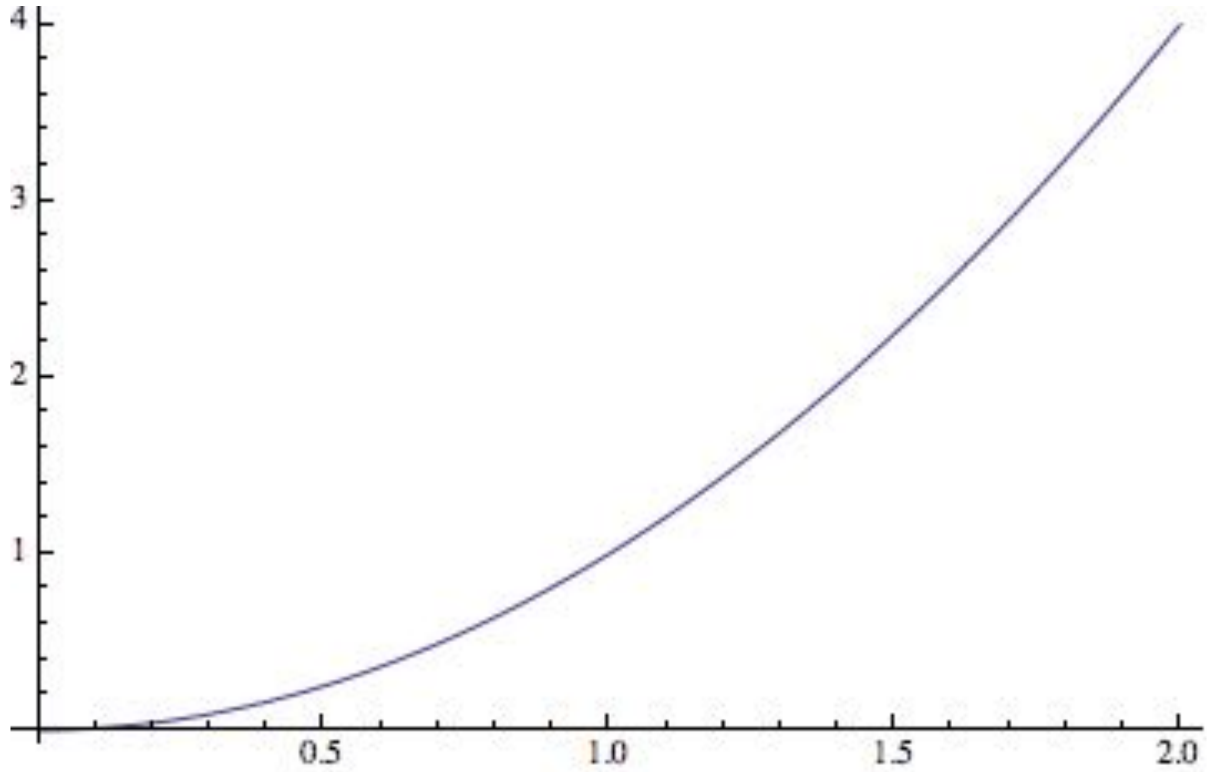

Compute CDF:

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

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

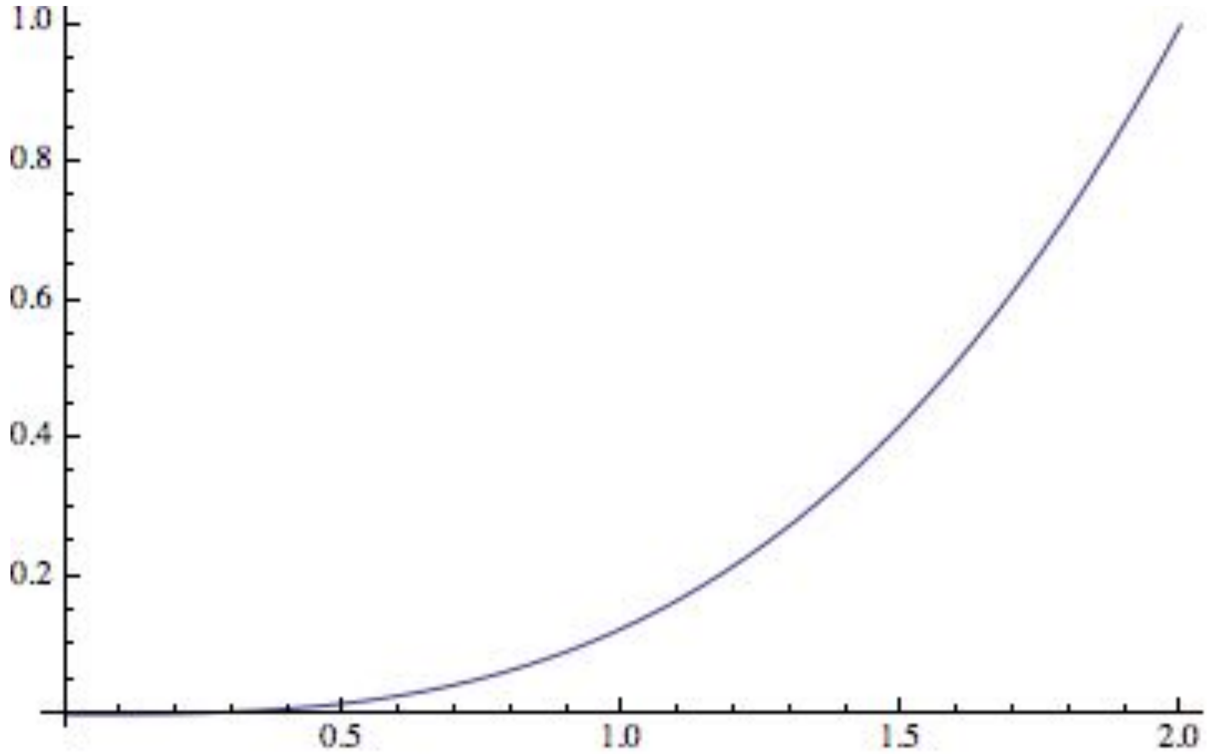

Example: applying the inversion method

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 $p(x)$

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

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

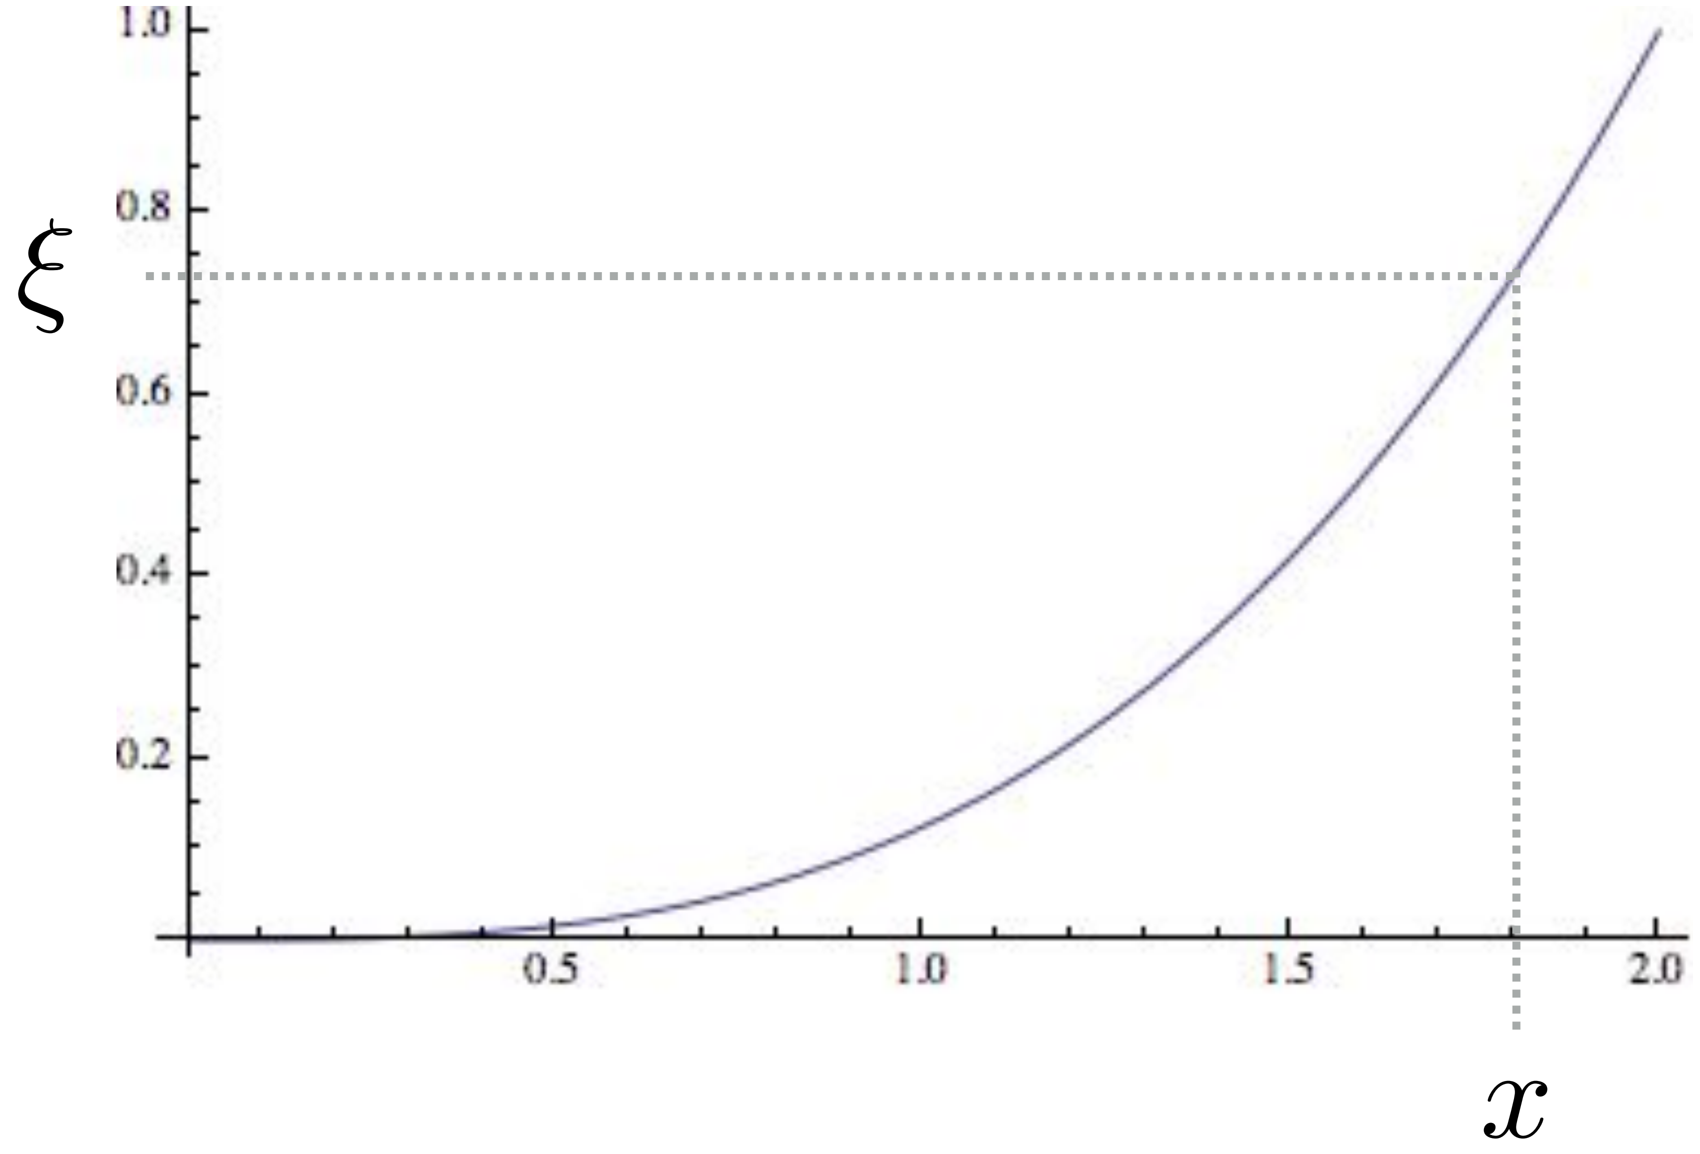

How do we uniformly sample the unit circle?

(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

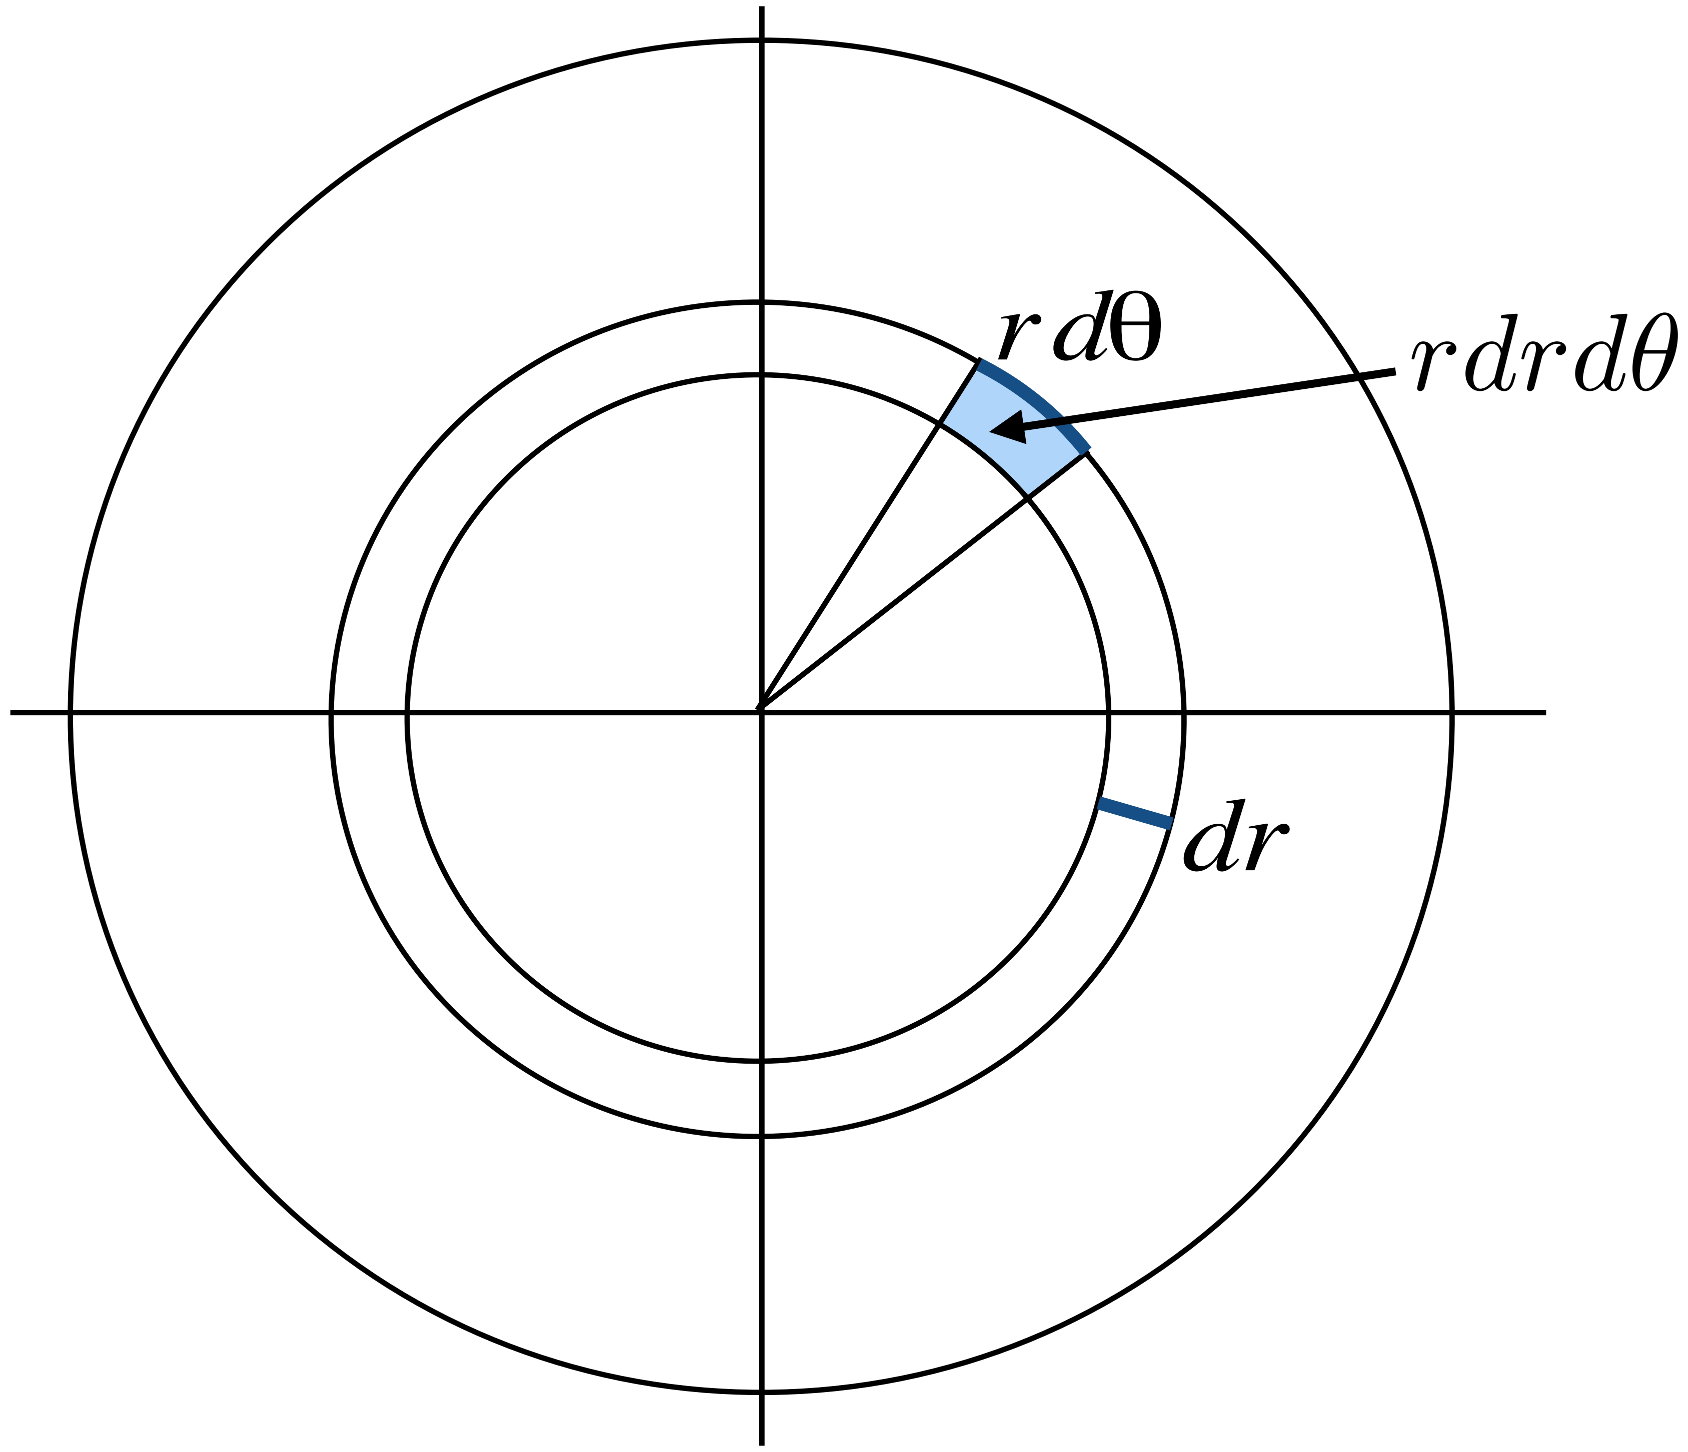

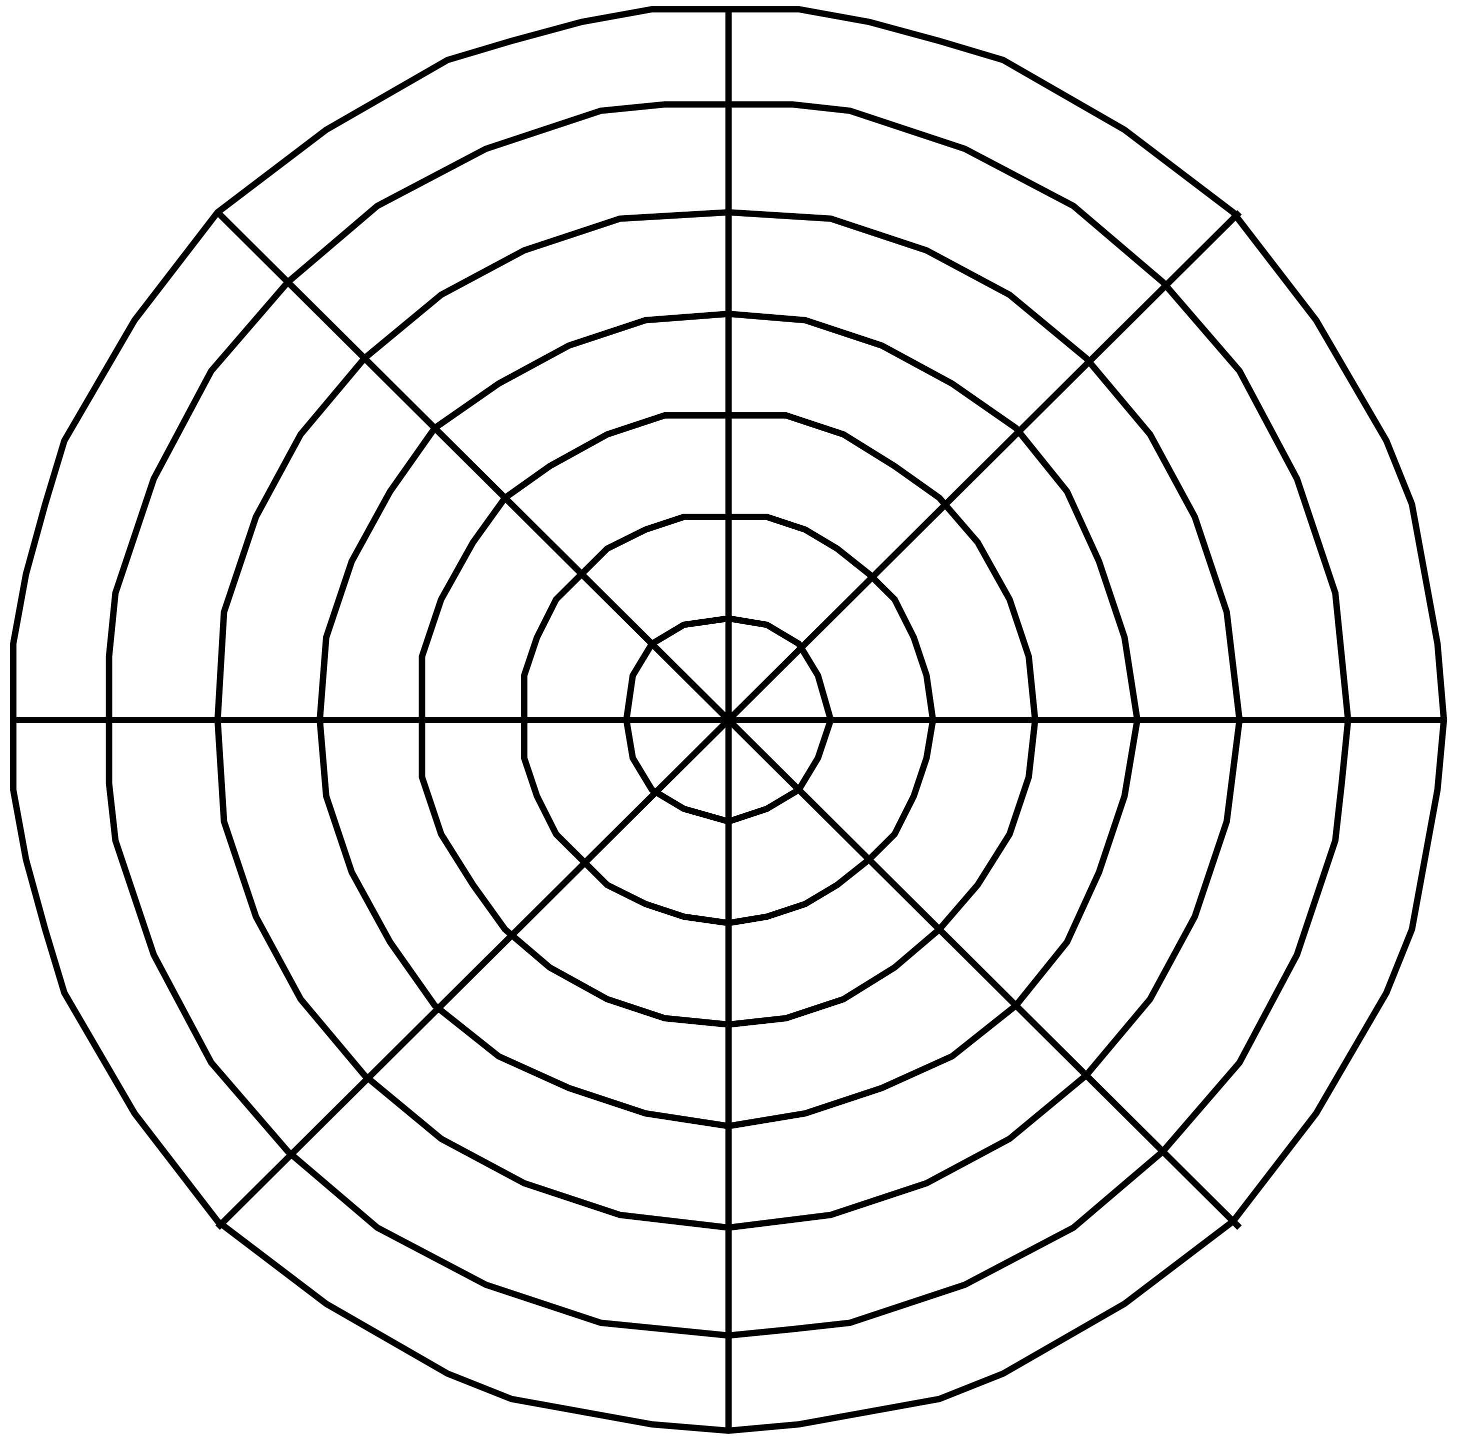

$$\theta = 2\pi\xi_1 \qquad r = \xi_2$$

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

Uniform area sampling of a circle

WRONG
Not Equi-areal

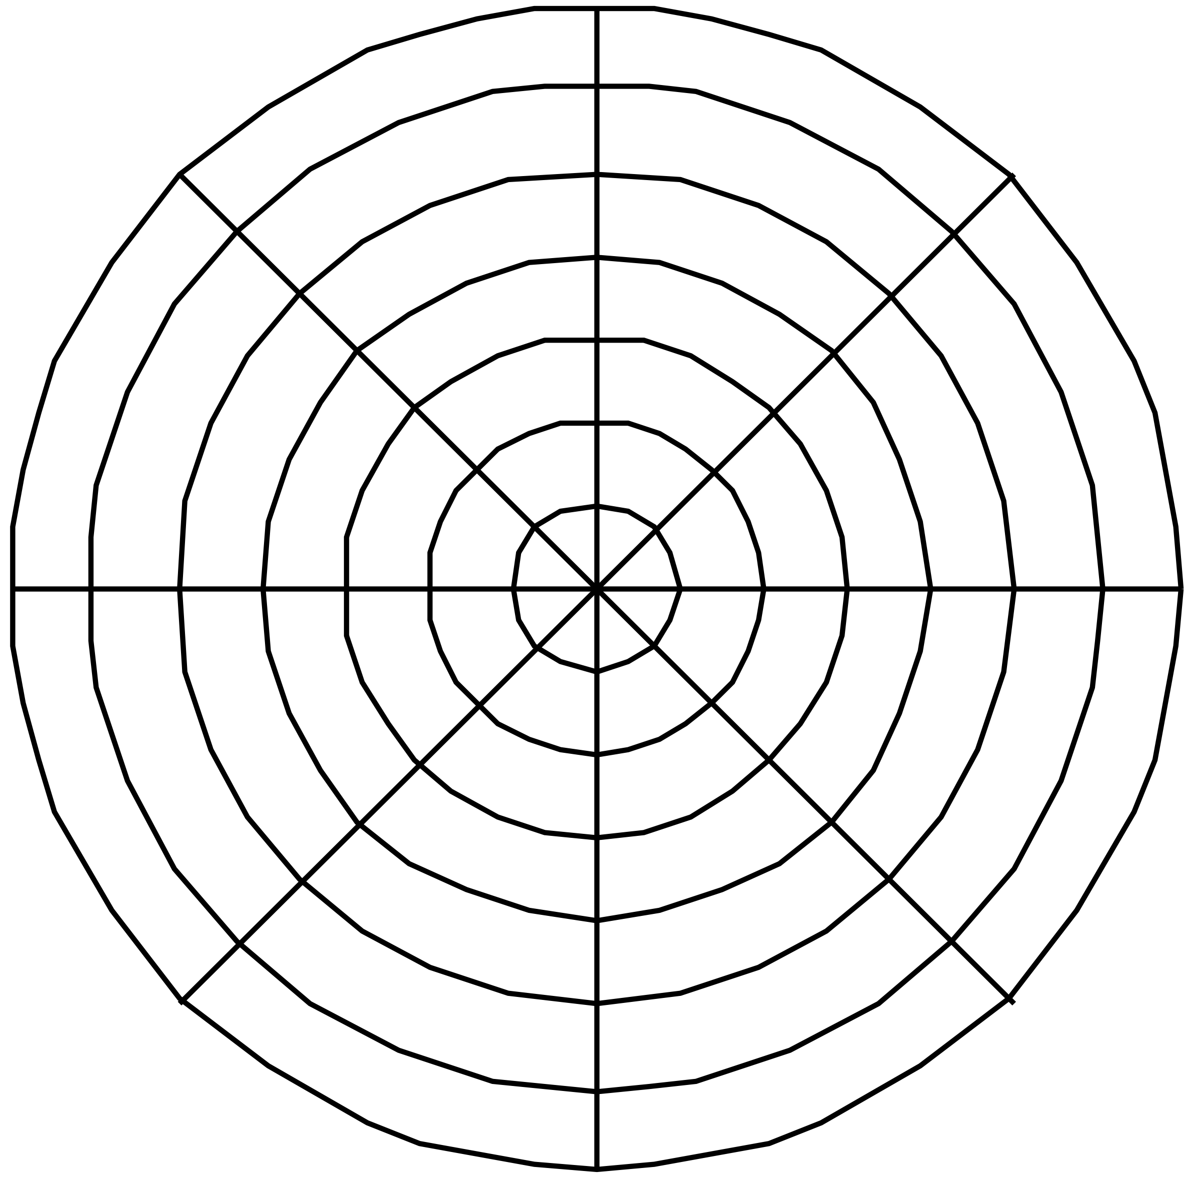

$$\theta = 2\pi\xi_1$$

$$r = \xi_2$$

RIGHT
Equi-areal

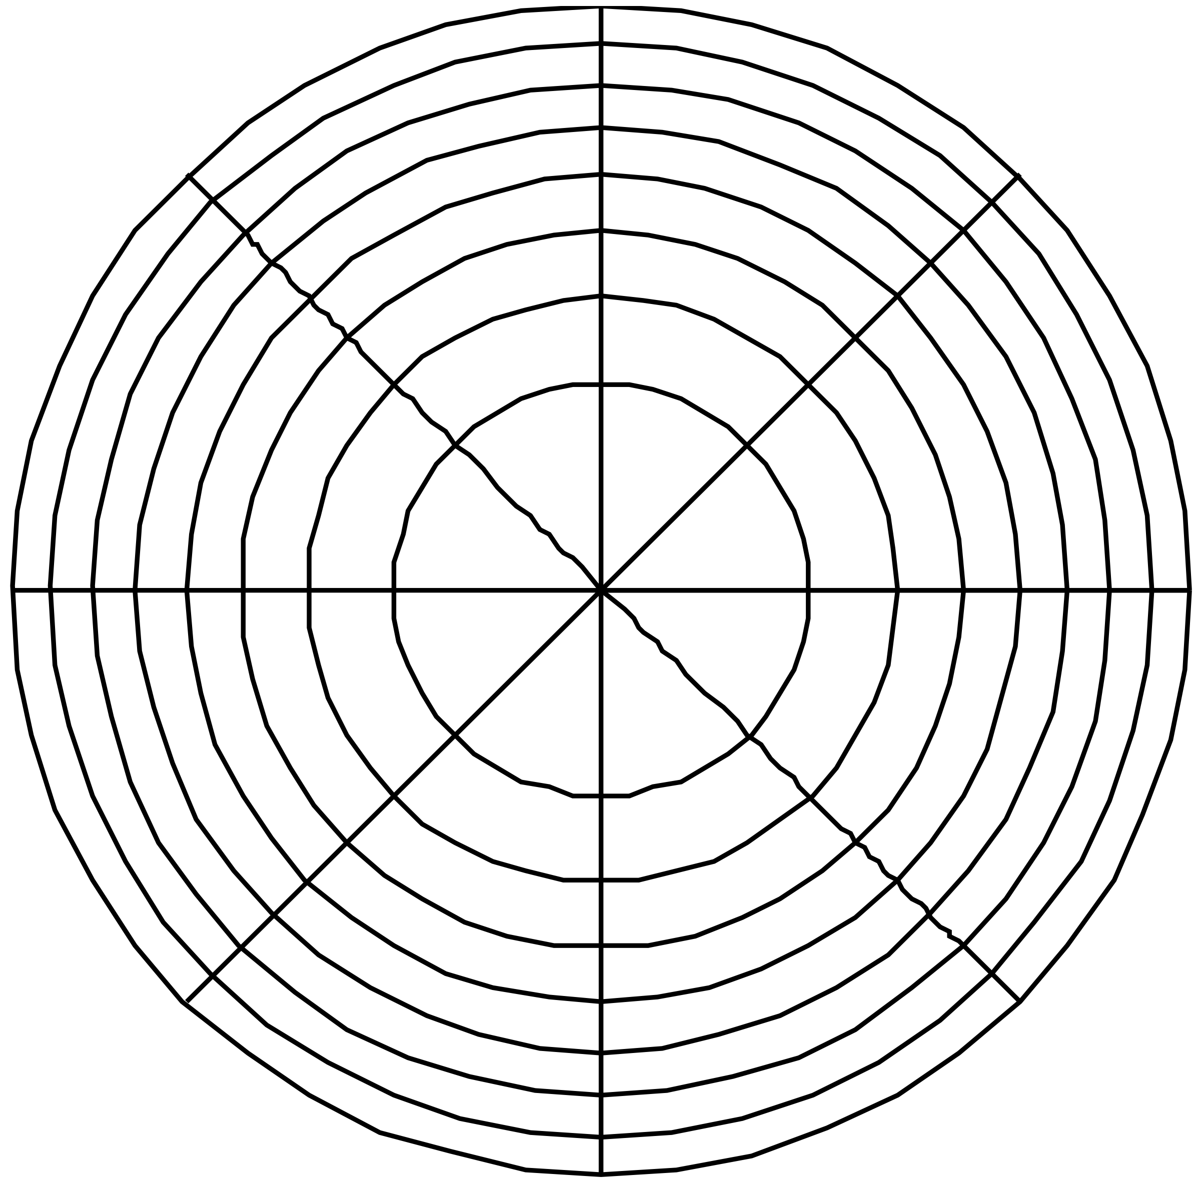

$$\theta = 2\pi\xi_1$$

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

Sampling a circle (via inversion in 2D)

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

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

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

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

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

$$p(r) = 2r$$

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

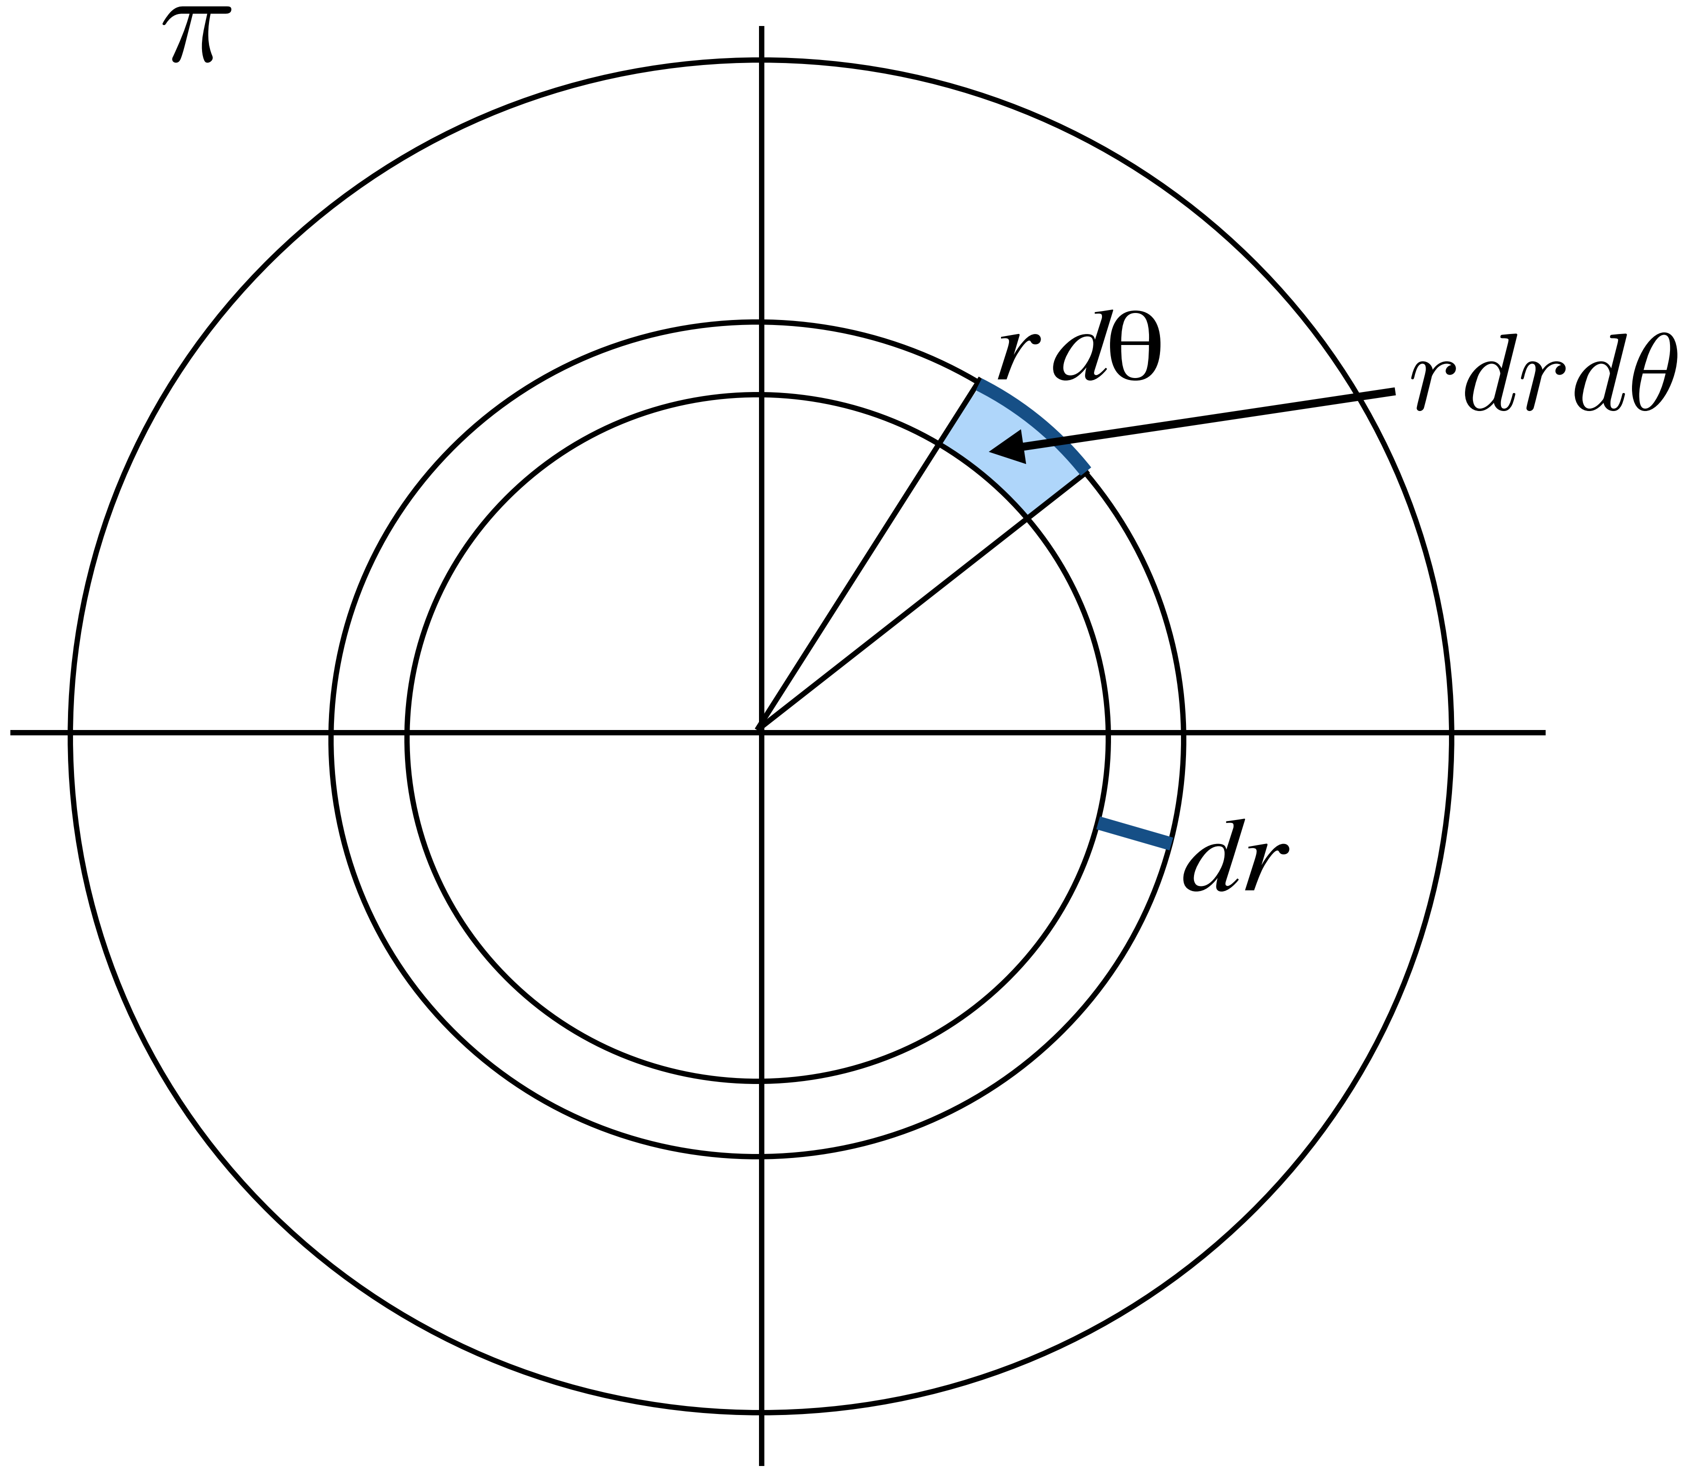

Shirley's mapping

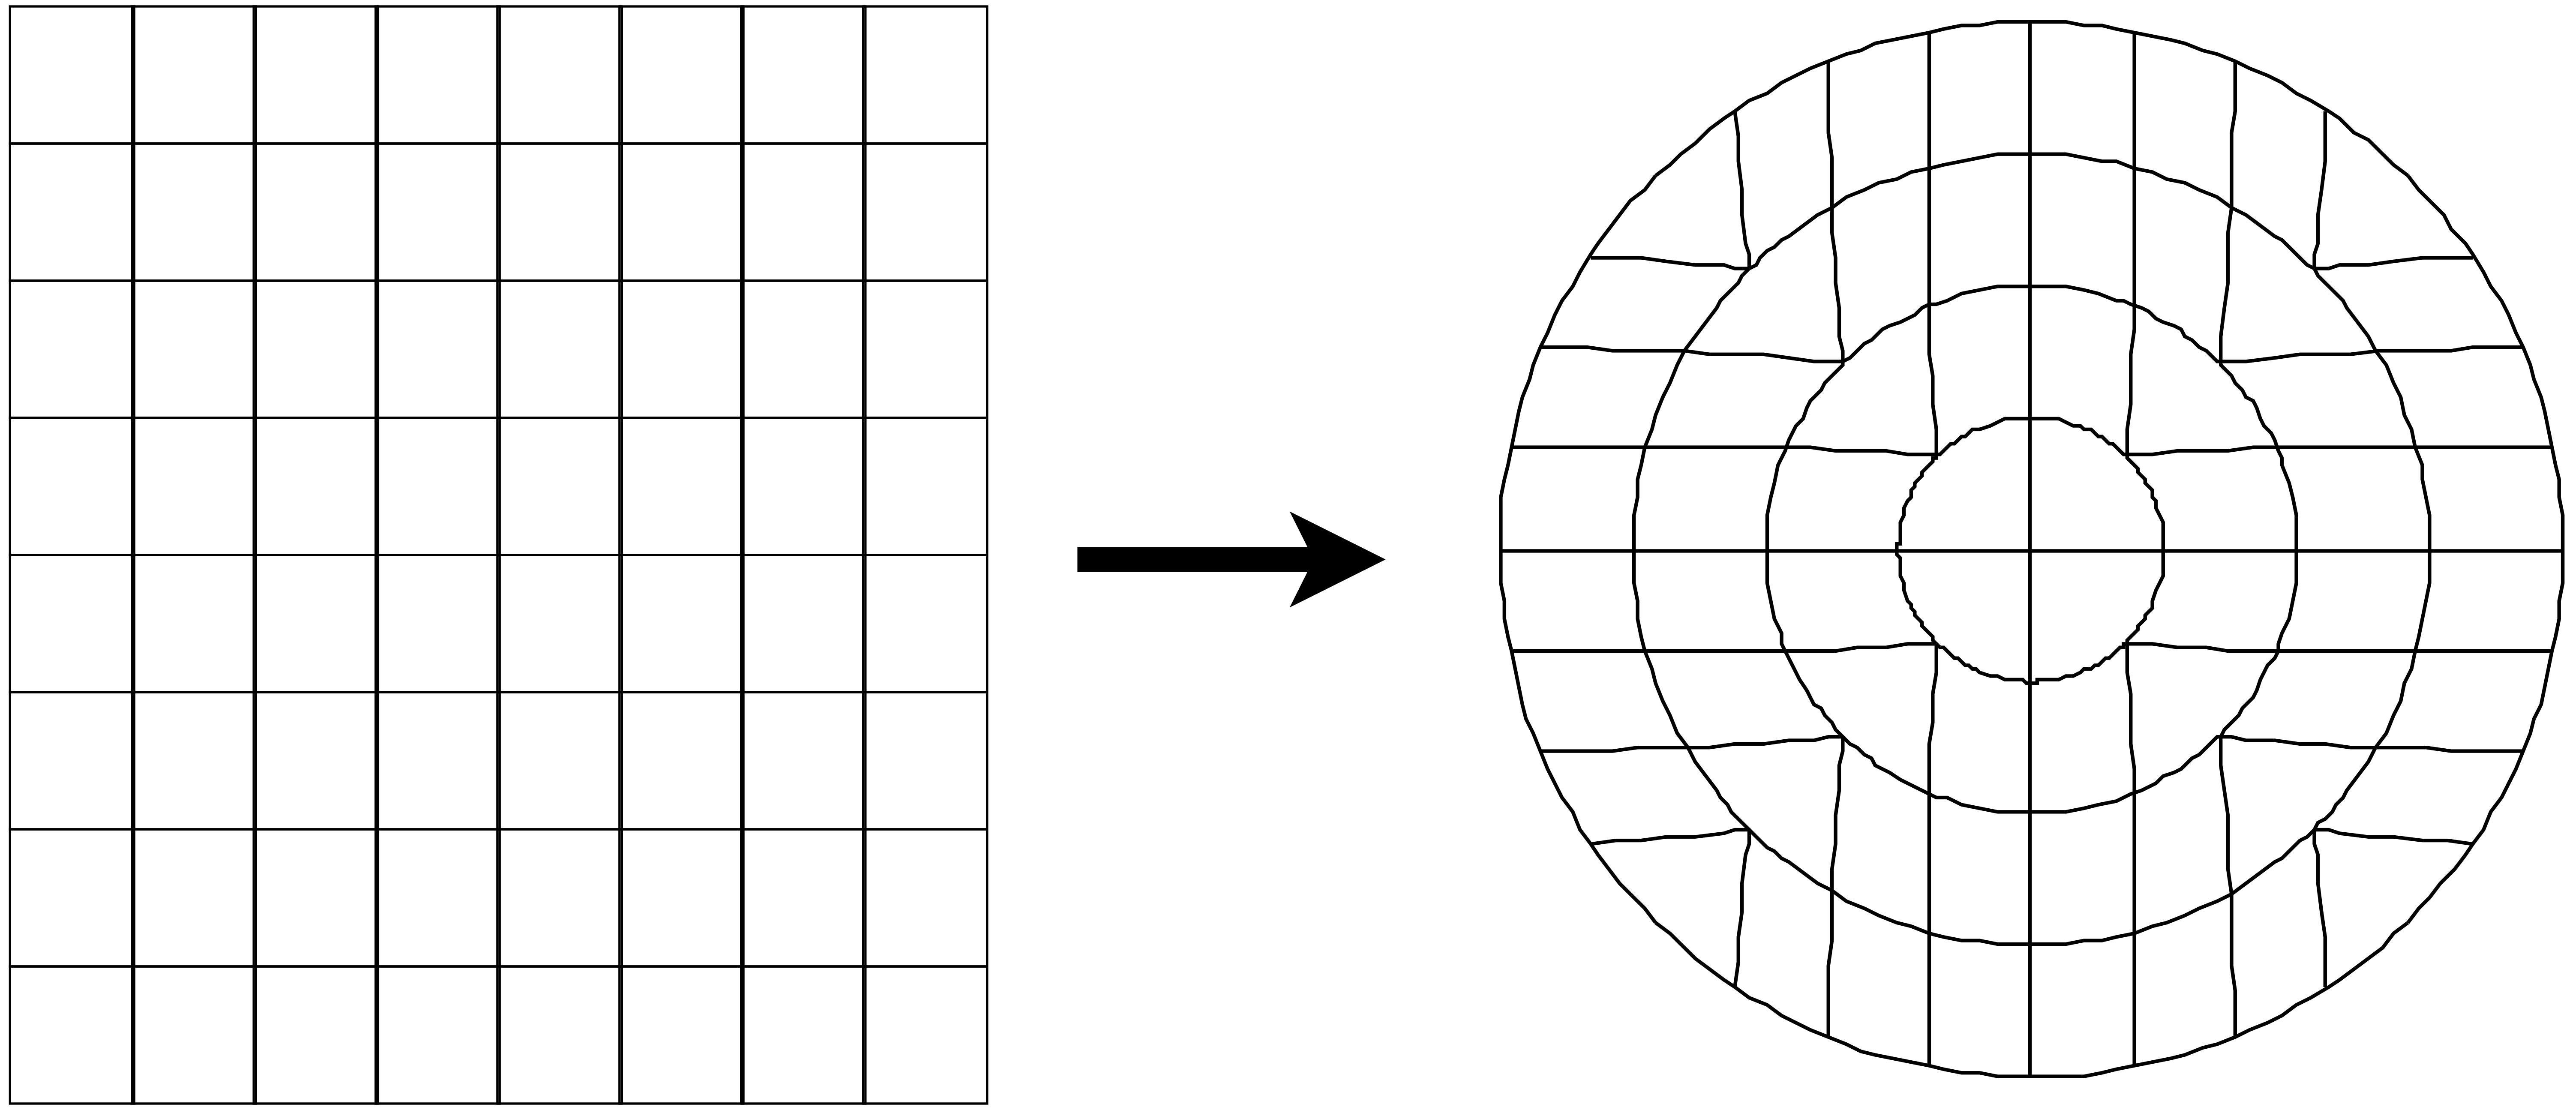

Distinct cases for eight octants

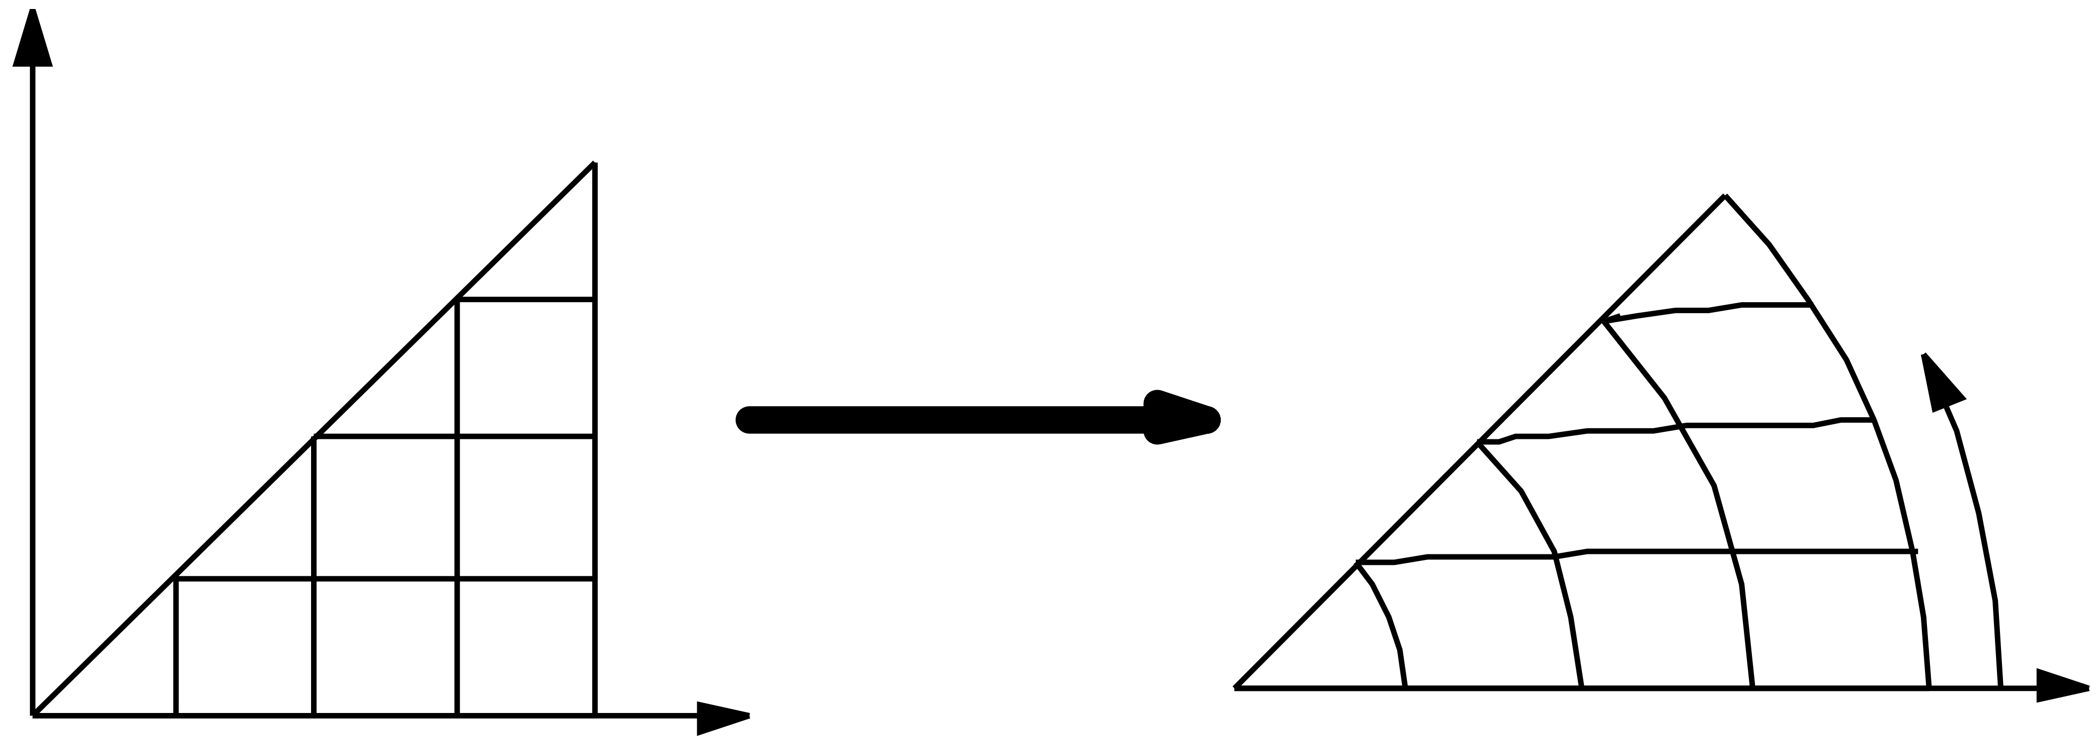

$$r = \xi_1$$

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

Uniform sampling via rejection sampling

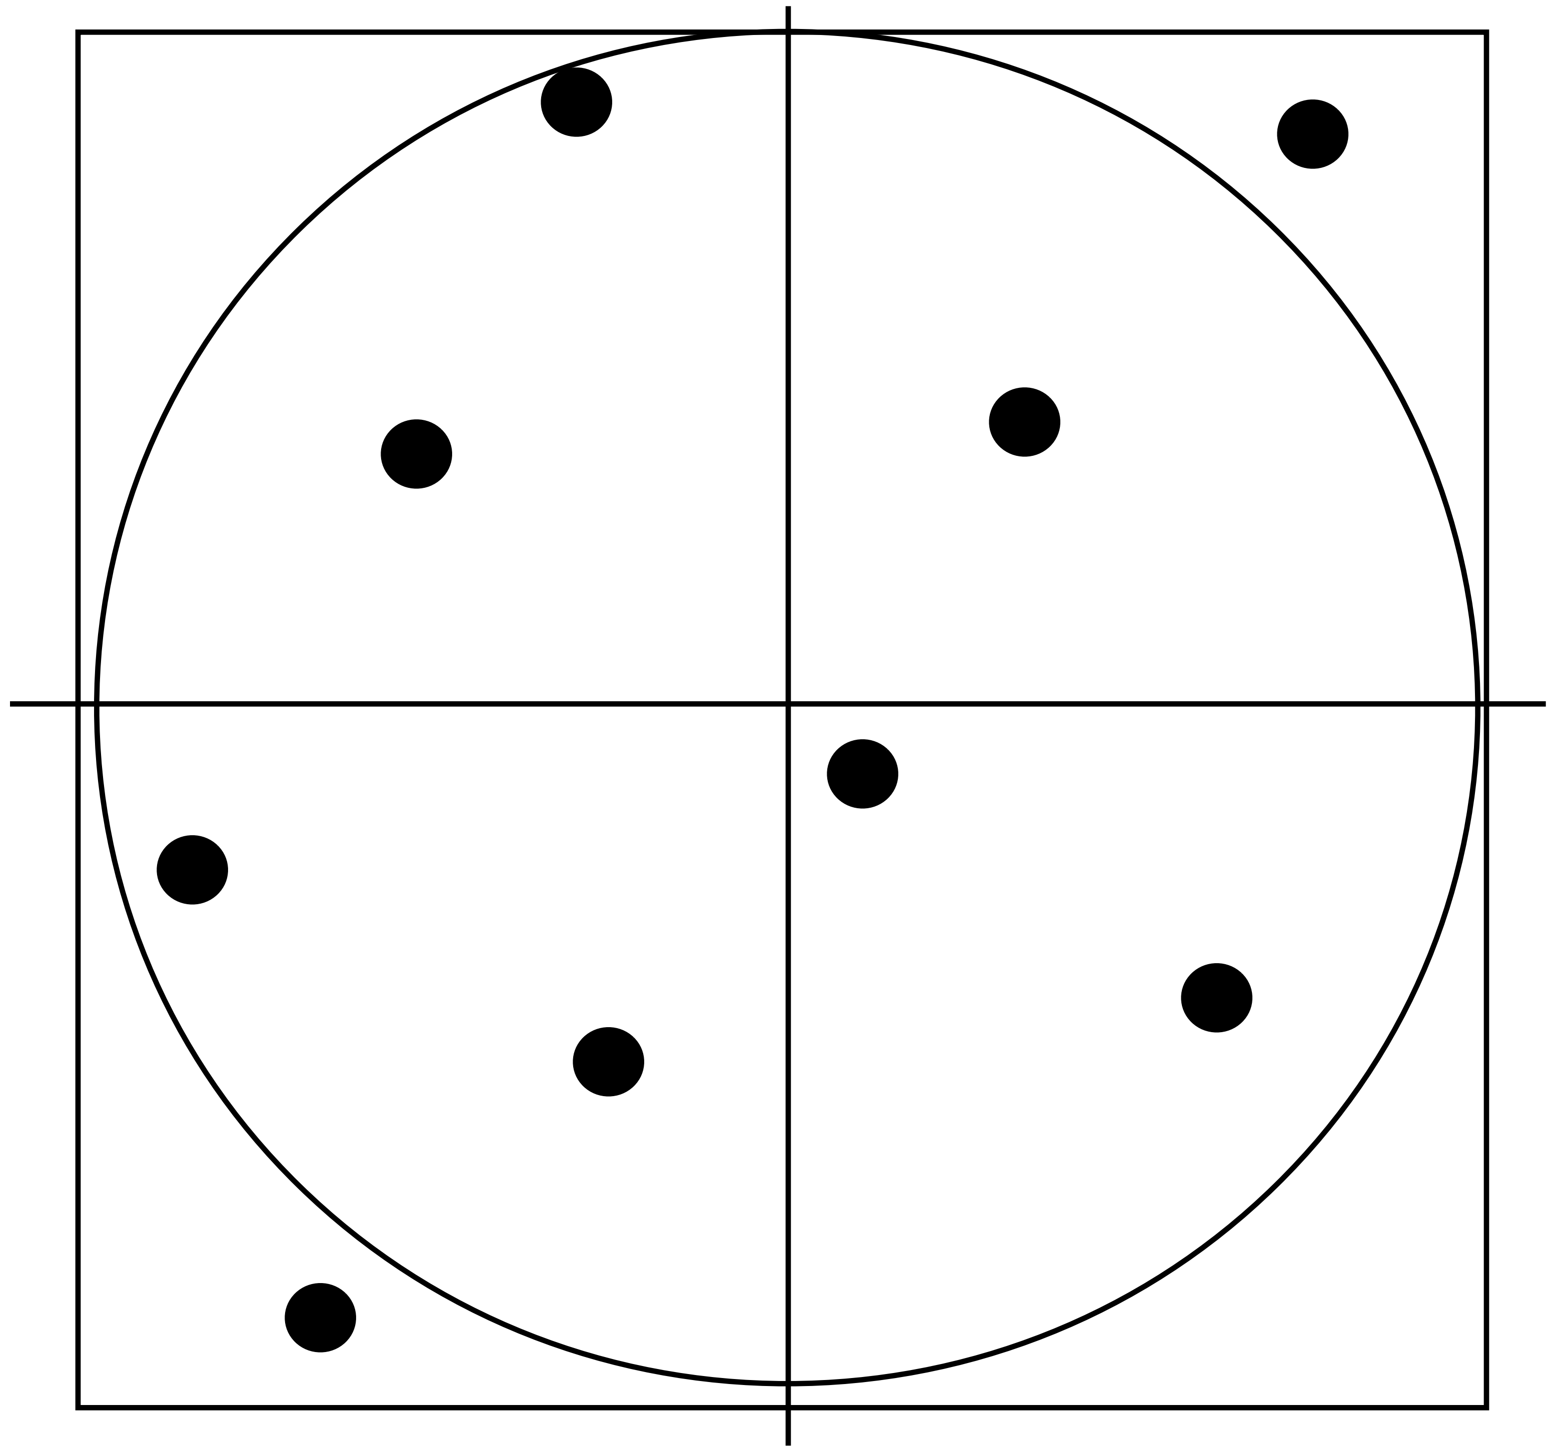

Generate random point within unit circle

```
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

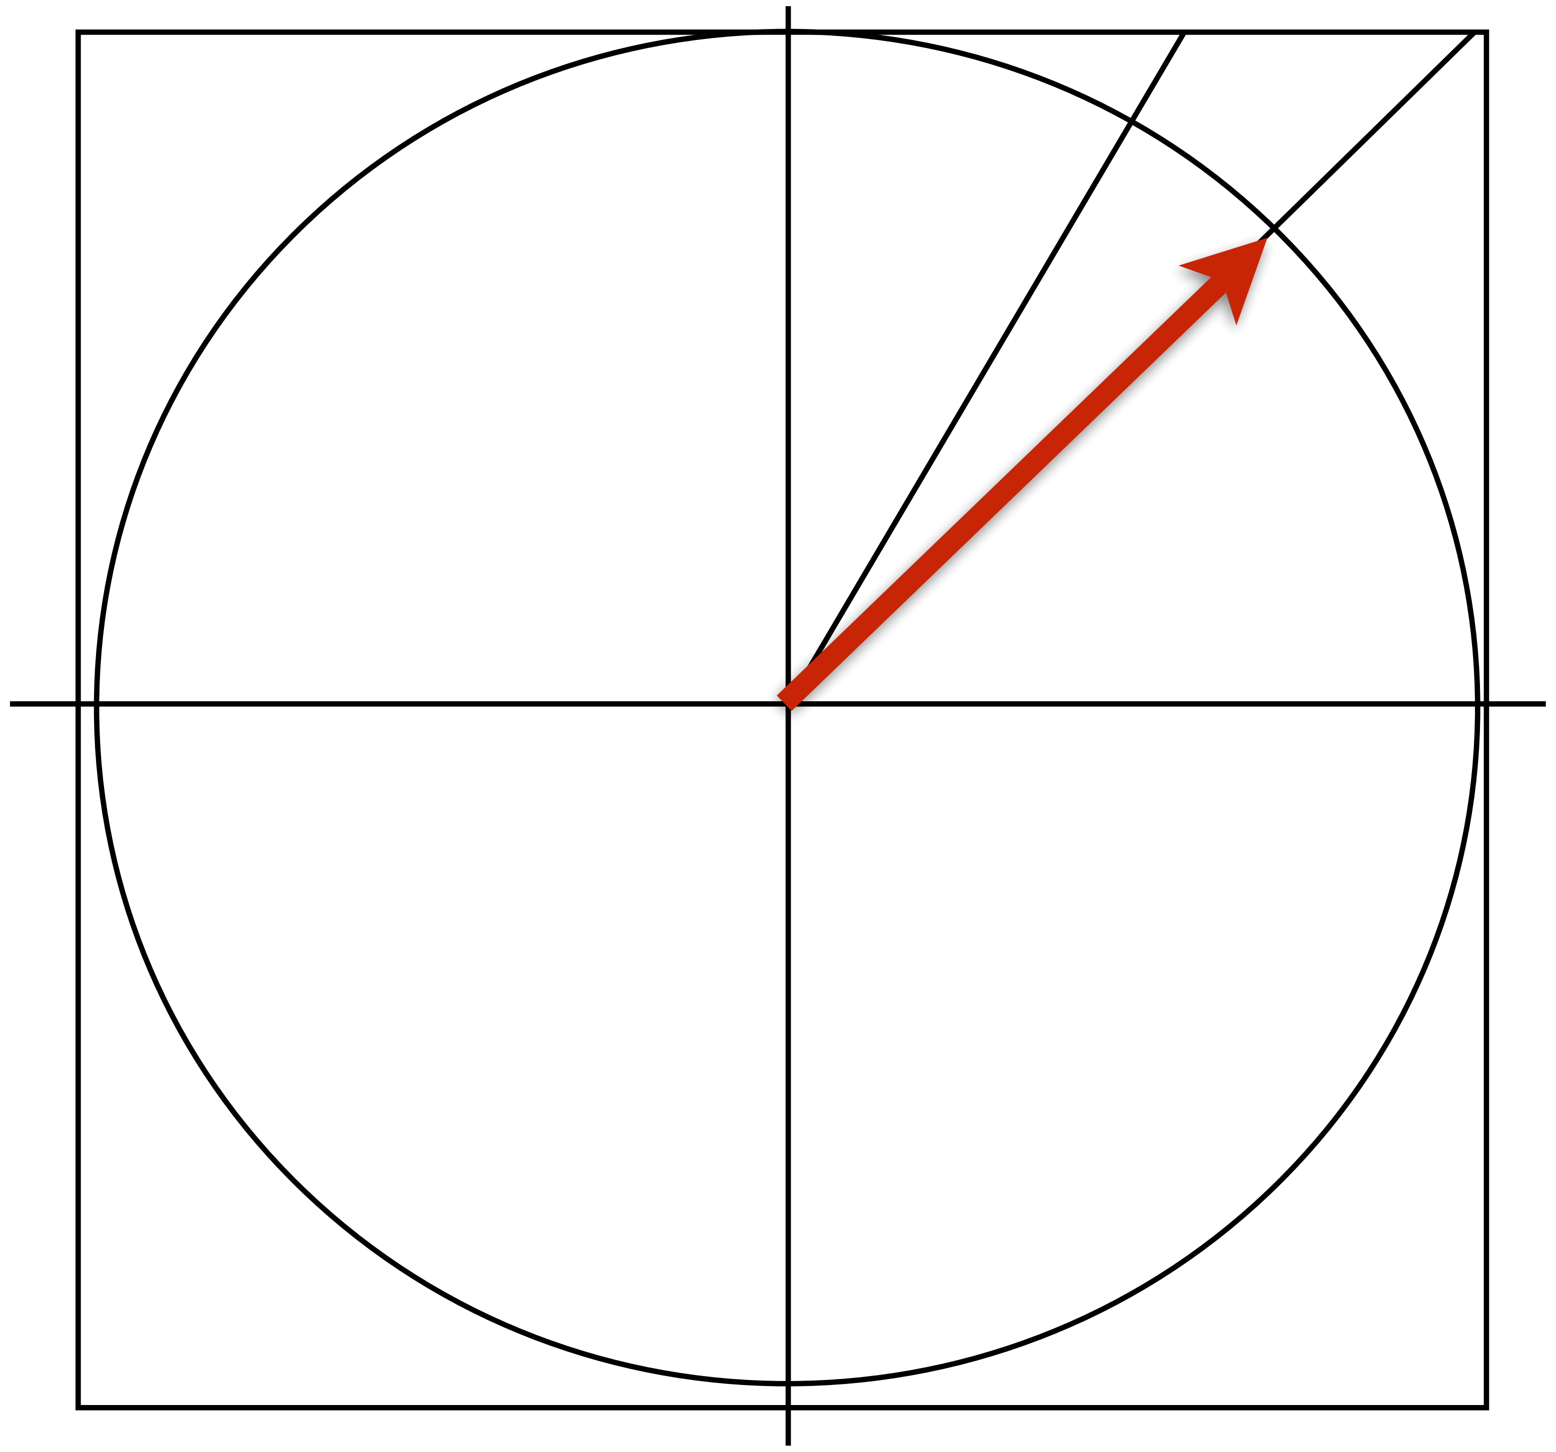

Goal: generate random directions in 2D with uniform probability

```
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?

Now back to Monte Carlo integration…

(Remember the whole point was to approximate the value of integrals numerically on a computer)

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

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

Monte Carlo integration

- **Definite integral**

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

The integral we seek to estimate

- **Random variables**

$$X_i \sim p(x)$$

$$Y_i = f(X_i)$$

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

Y_i **is also a random variable because its a function of** X_i

- **Expectation of a random variable**

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

- **Monte Carlo estimator of the integral**

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

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.

Basic unbiased Monte Carlo estimator

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

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

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

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

$$= \int_a^b f(x)\,\mathrm{d}x$$

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

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

Estimate incident irradiance by uniformly-sampling hemisphere of directions with respect to solid angle

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

We want to estimate this integral (total incident irradiance at surface point x)

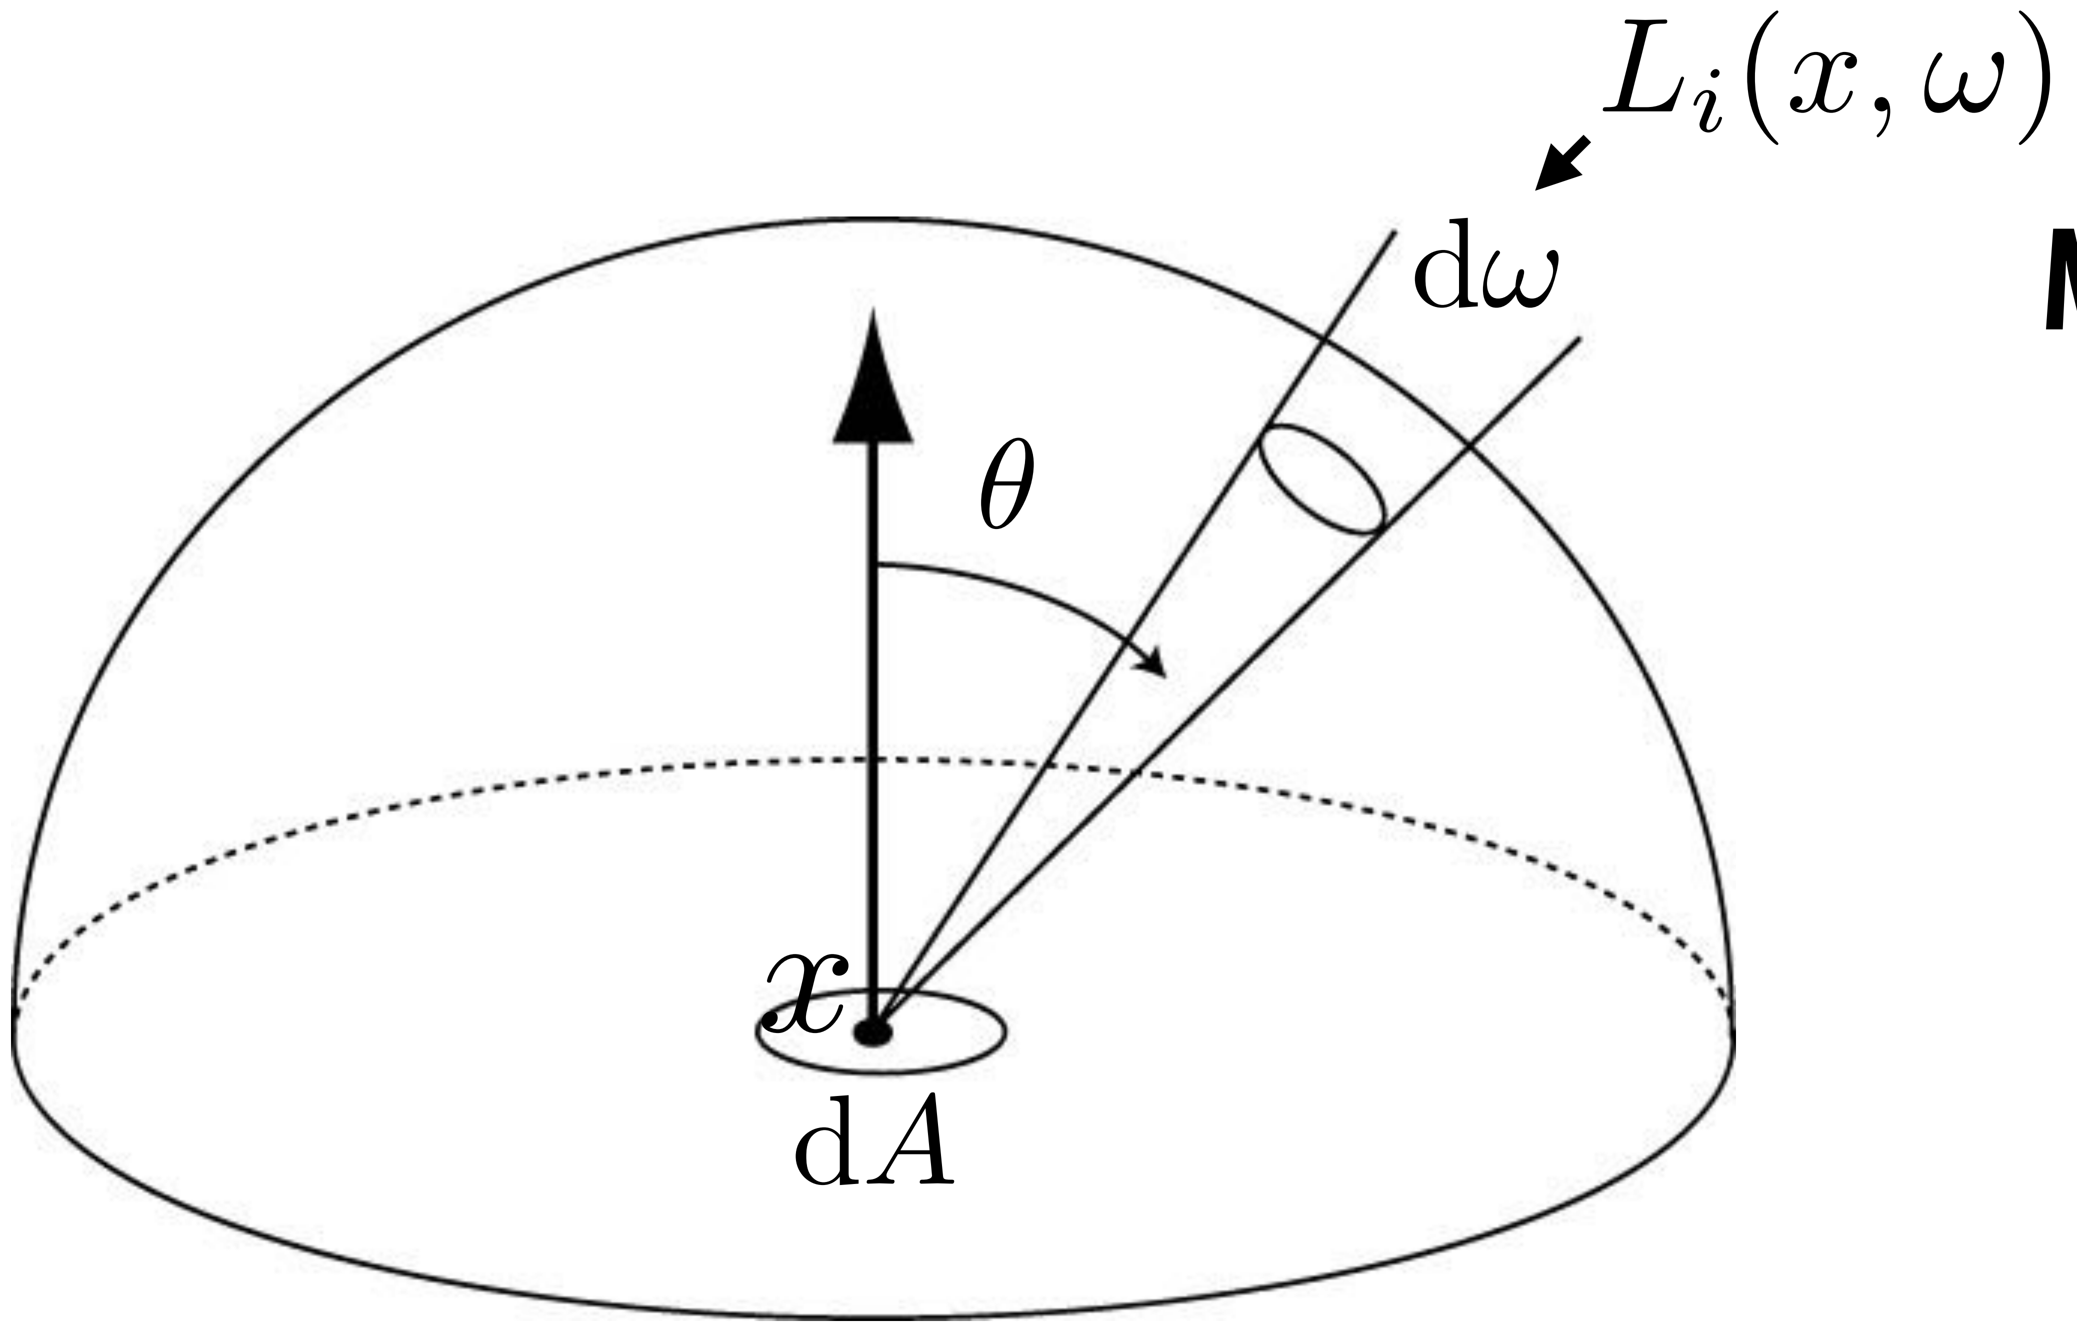

Monte Carlo estimator:

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

We sample directions (aka rays) uniformly from the hemisphere of directions
(a ray direction is a random variable)

$$Y_i = f(X_i)$$

$$Y_i = \boxed{L_i(x,\omega_i)\cos\theta_i}$$

For each ray we compute the incident irradiance on surface at x.

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

We average all these samples, and scale by the size of the domain we are sampling from.
(The hemisphere has 2π steradians)

Then the expected value of the result is the value of the integral.

Direct lighting estimate

Uniformly-sample hemisphere of directions with respect to solid angle

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

Given surface point *x*

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 *x* 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**

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:

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

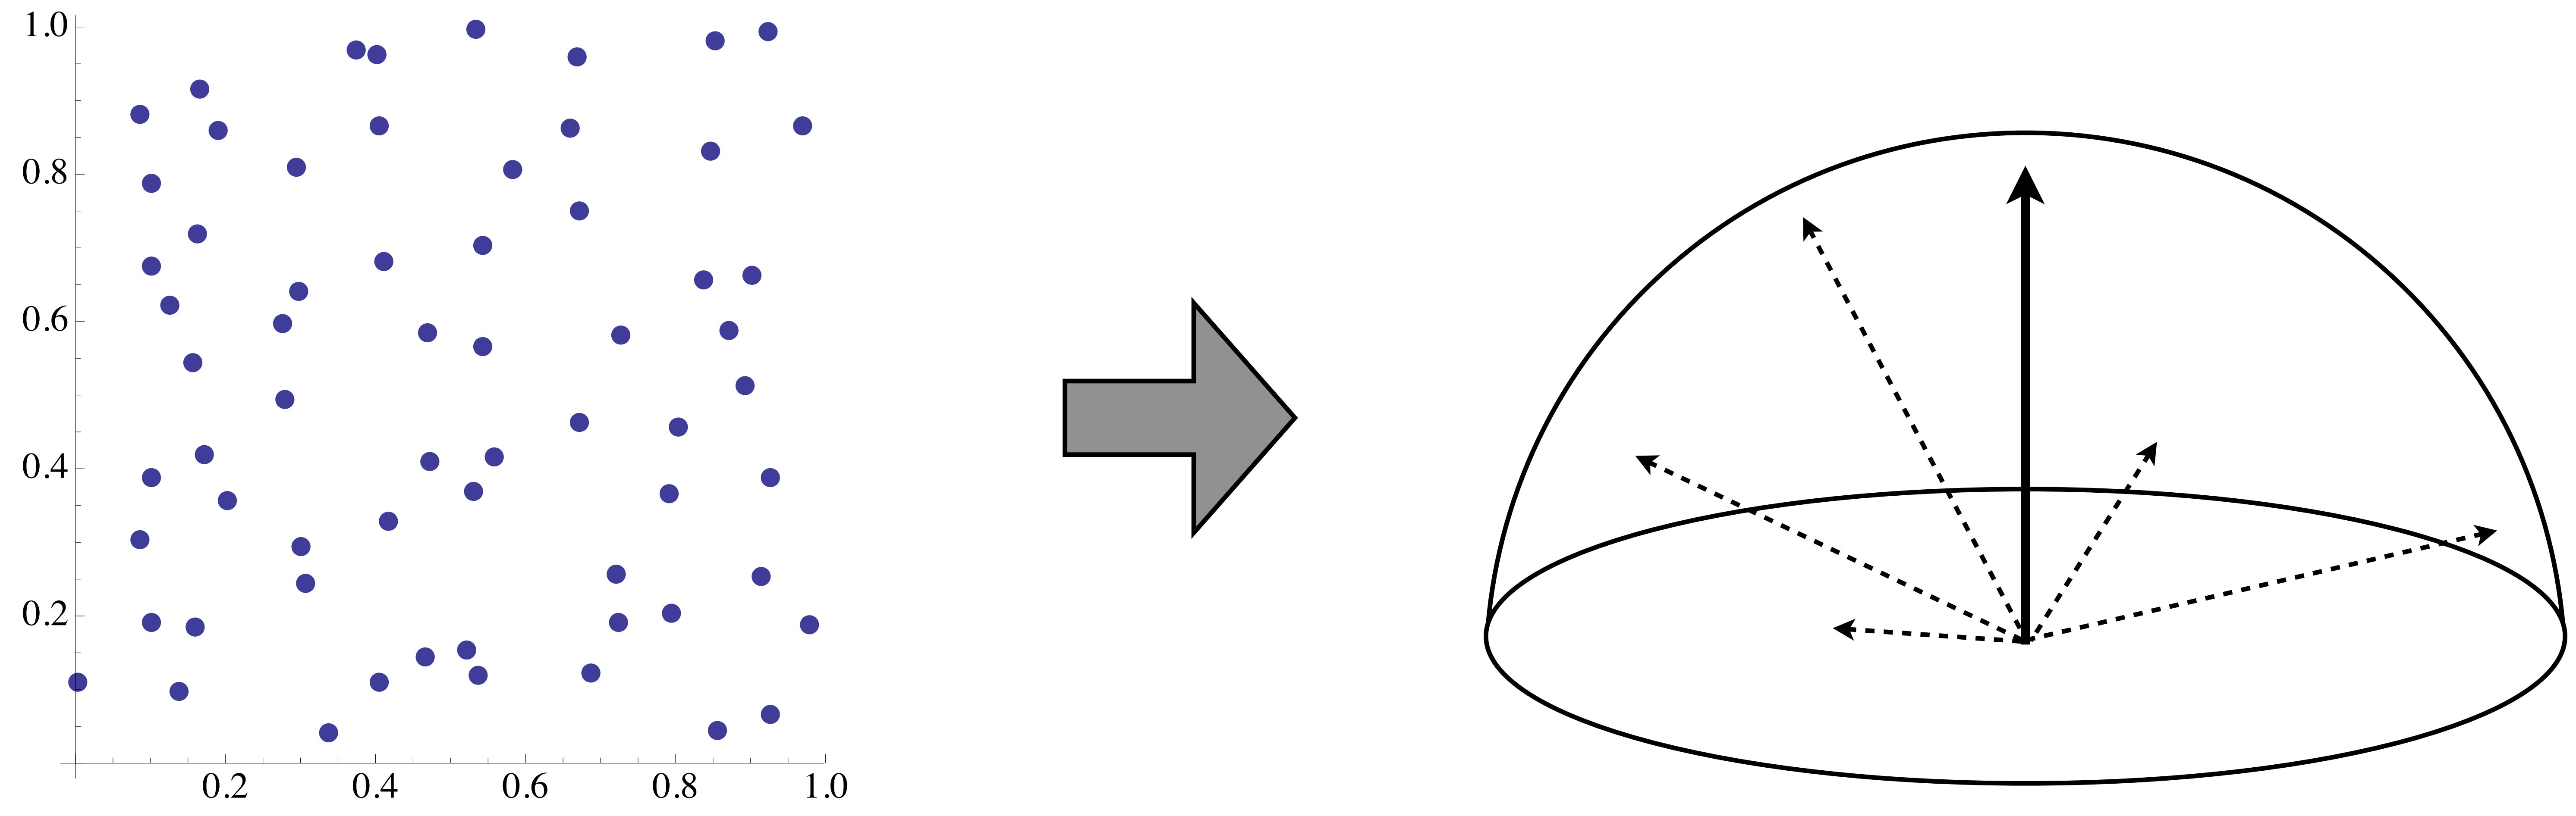

Exercise to students: derive from the inversion method

Example scene with an "area light"

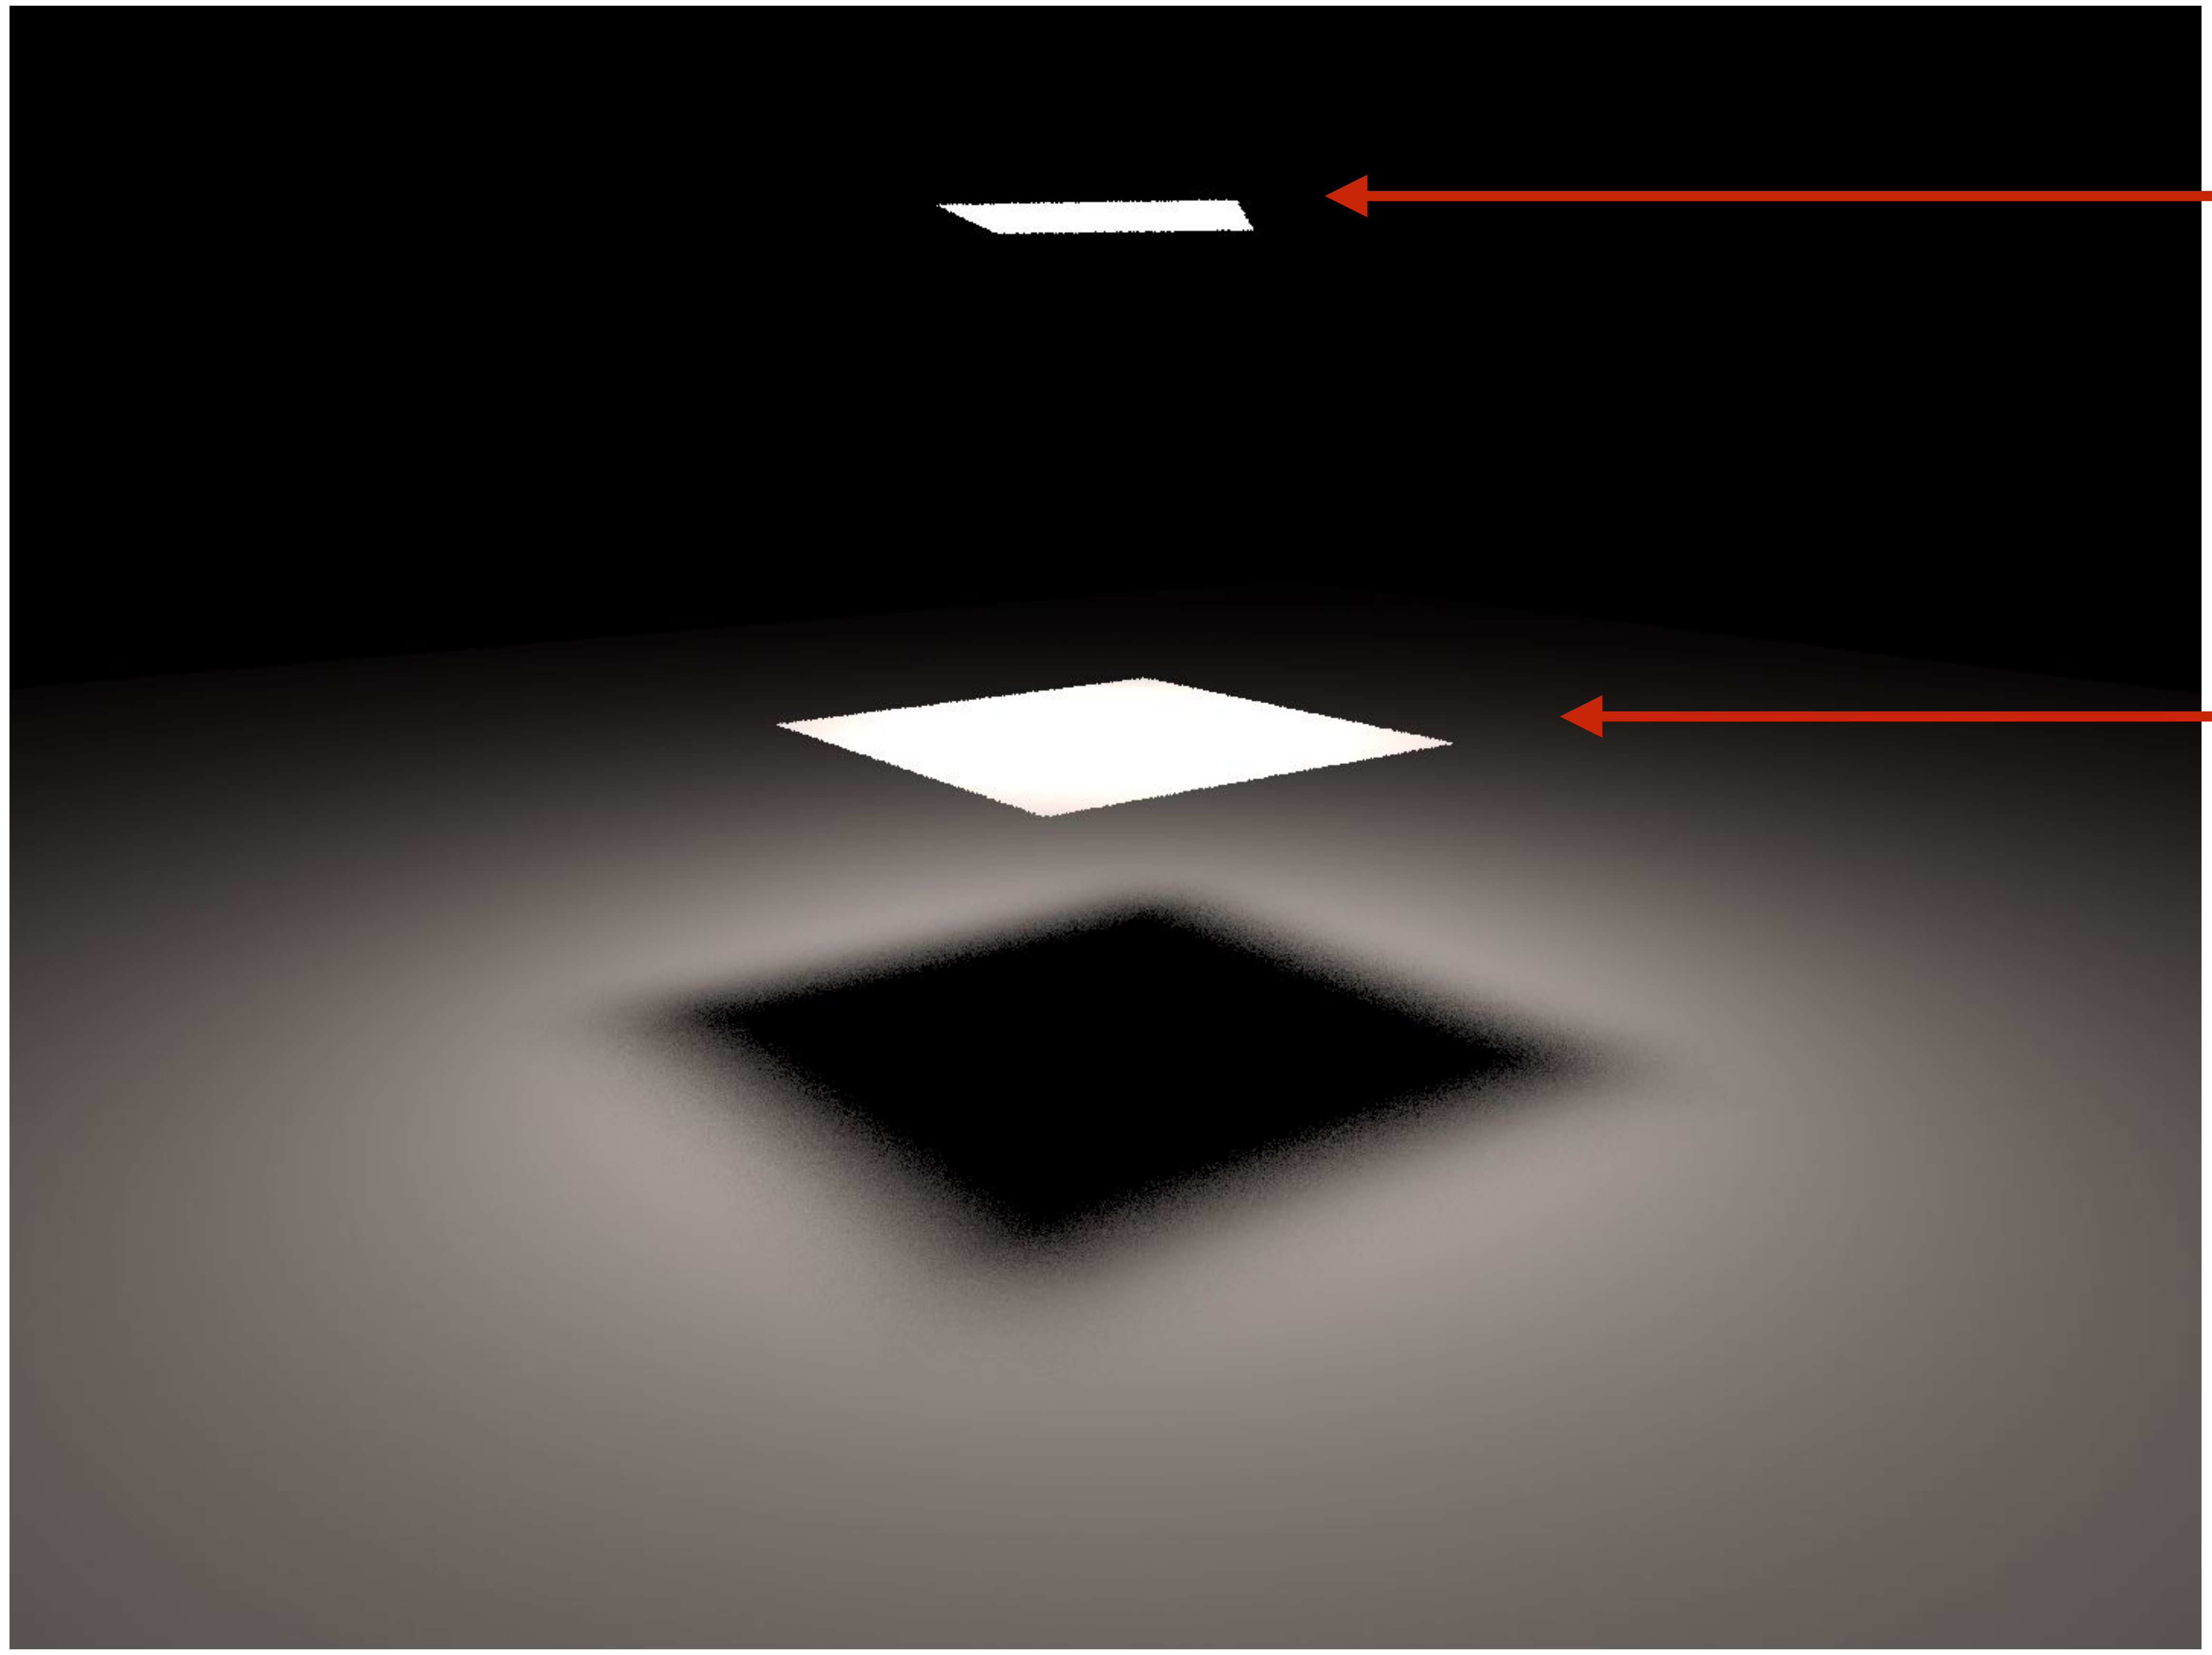

Light source

Occluder
(blocks light)

Direct lighting estimate: uniform hemisphere sampling

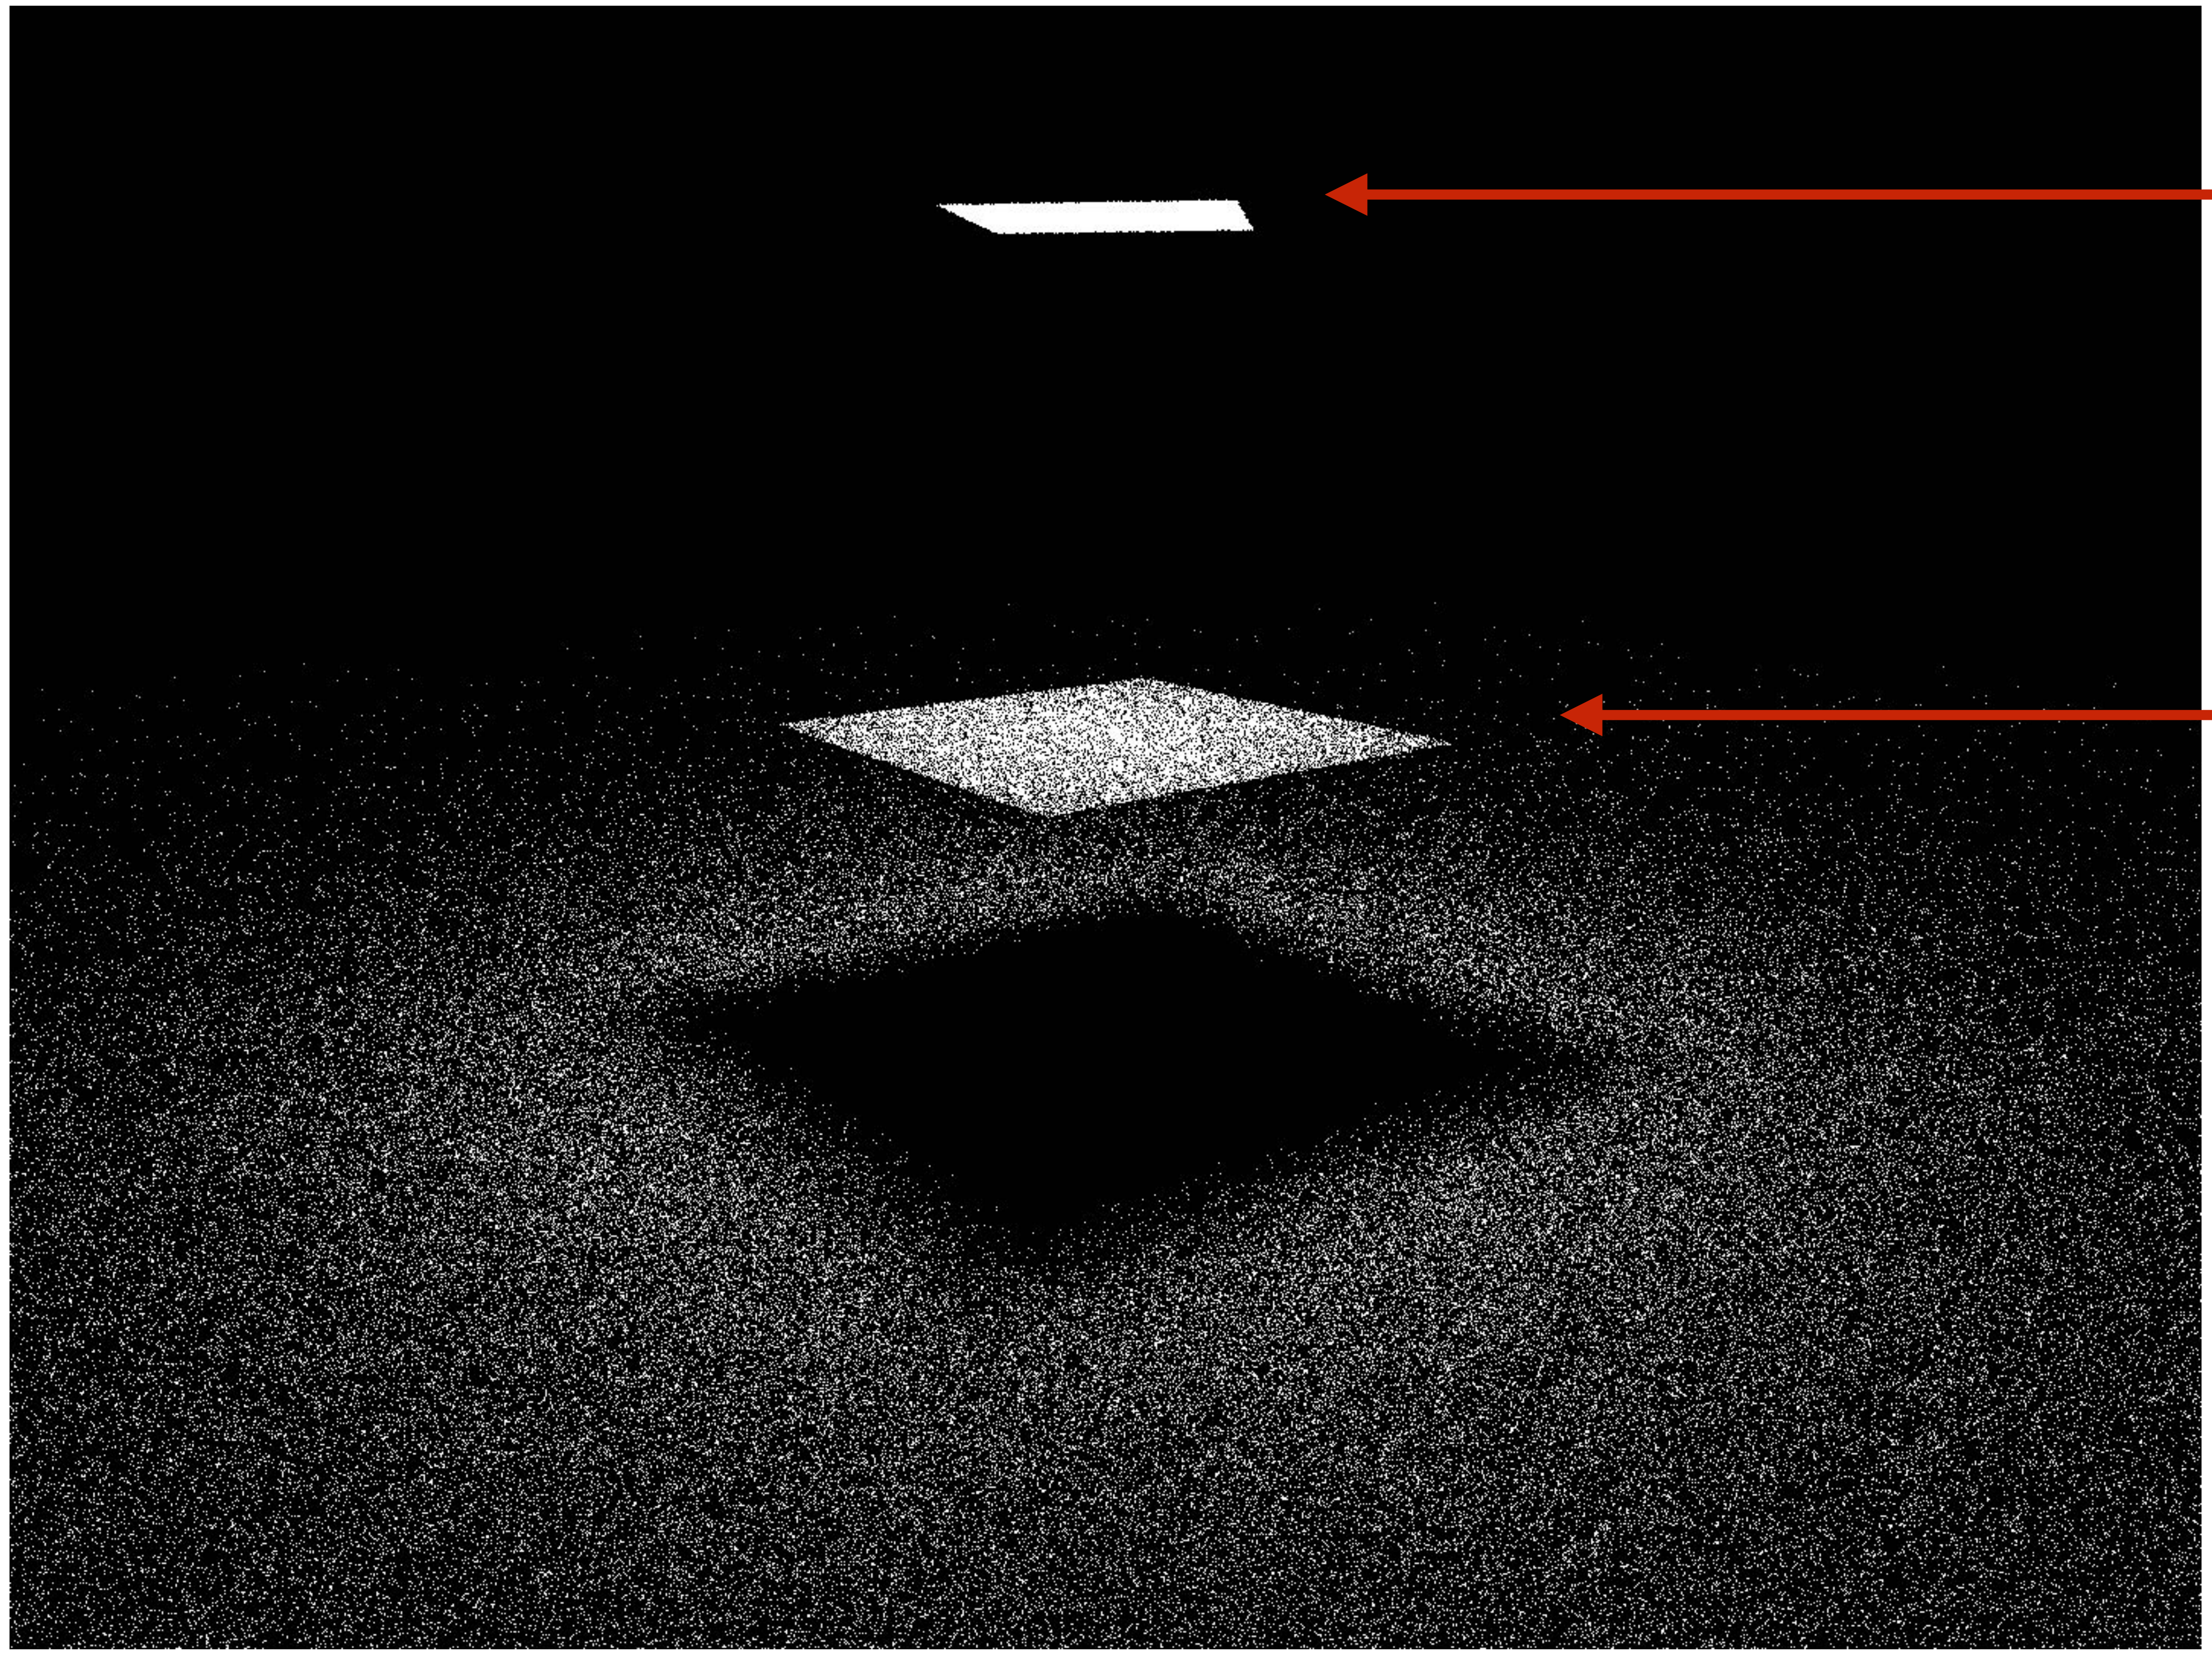

Light source

Occluder (blocks light)

16 samples to estimate incoming irradiance

Direct lighting: uniform hemisphere sampling

Incident lighting estimator uses random directions when computing incident lighting for different points. Some of those directions hit the light (and contribute illumination, some do not)

(The estimator is a random variable!)

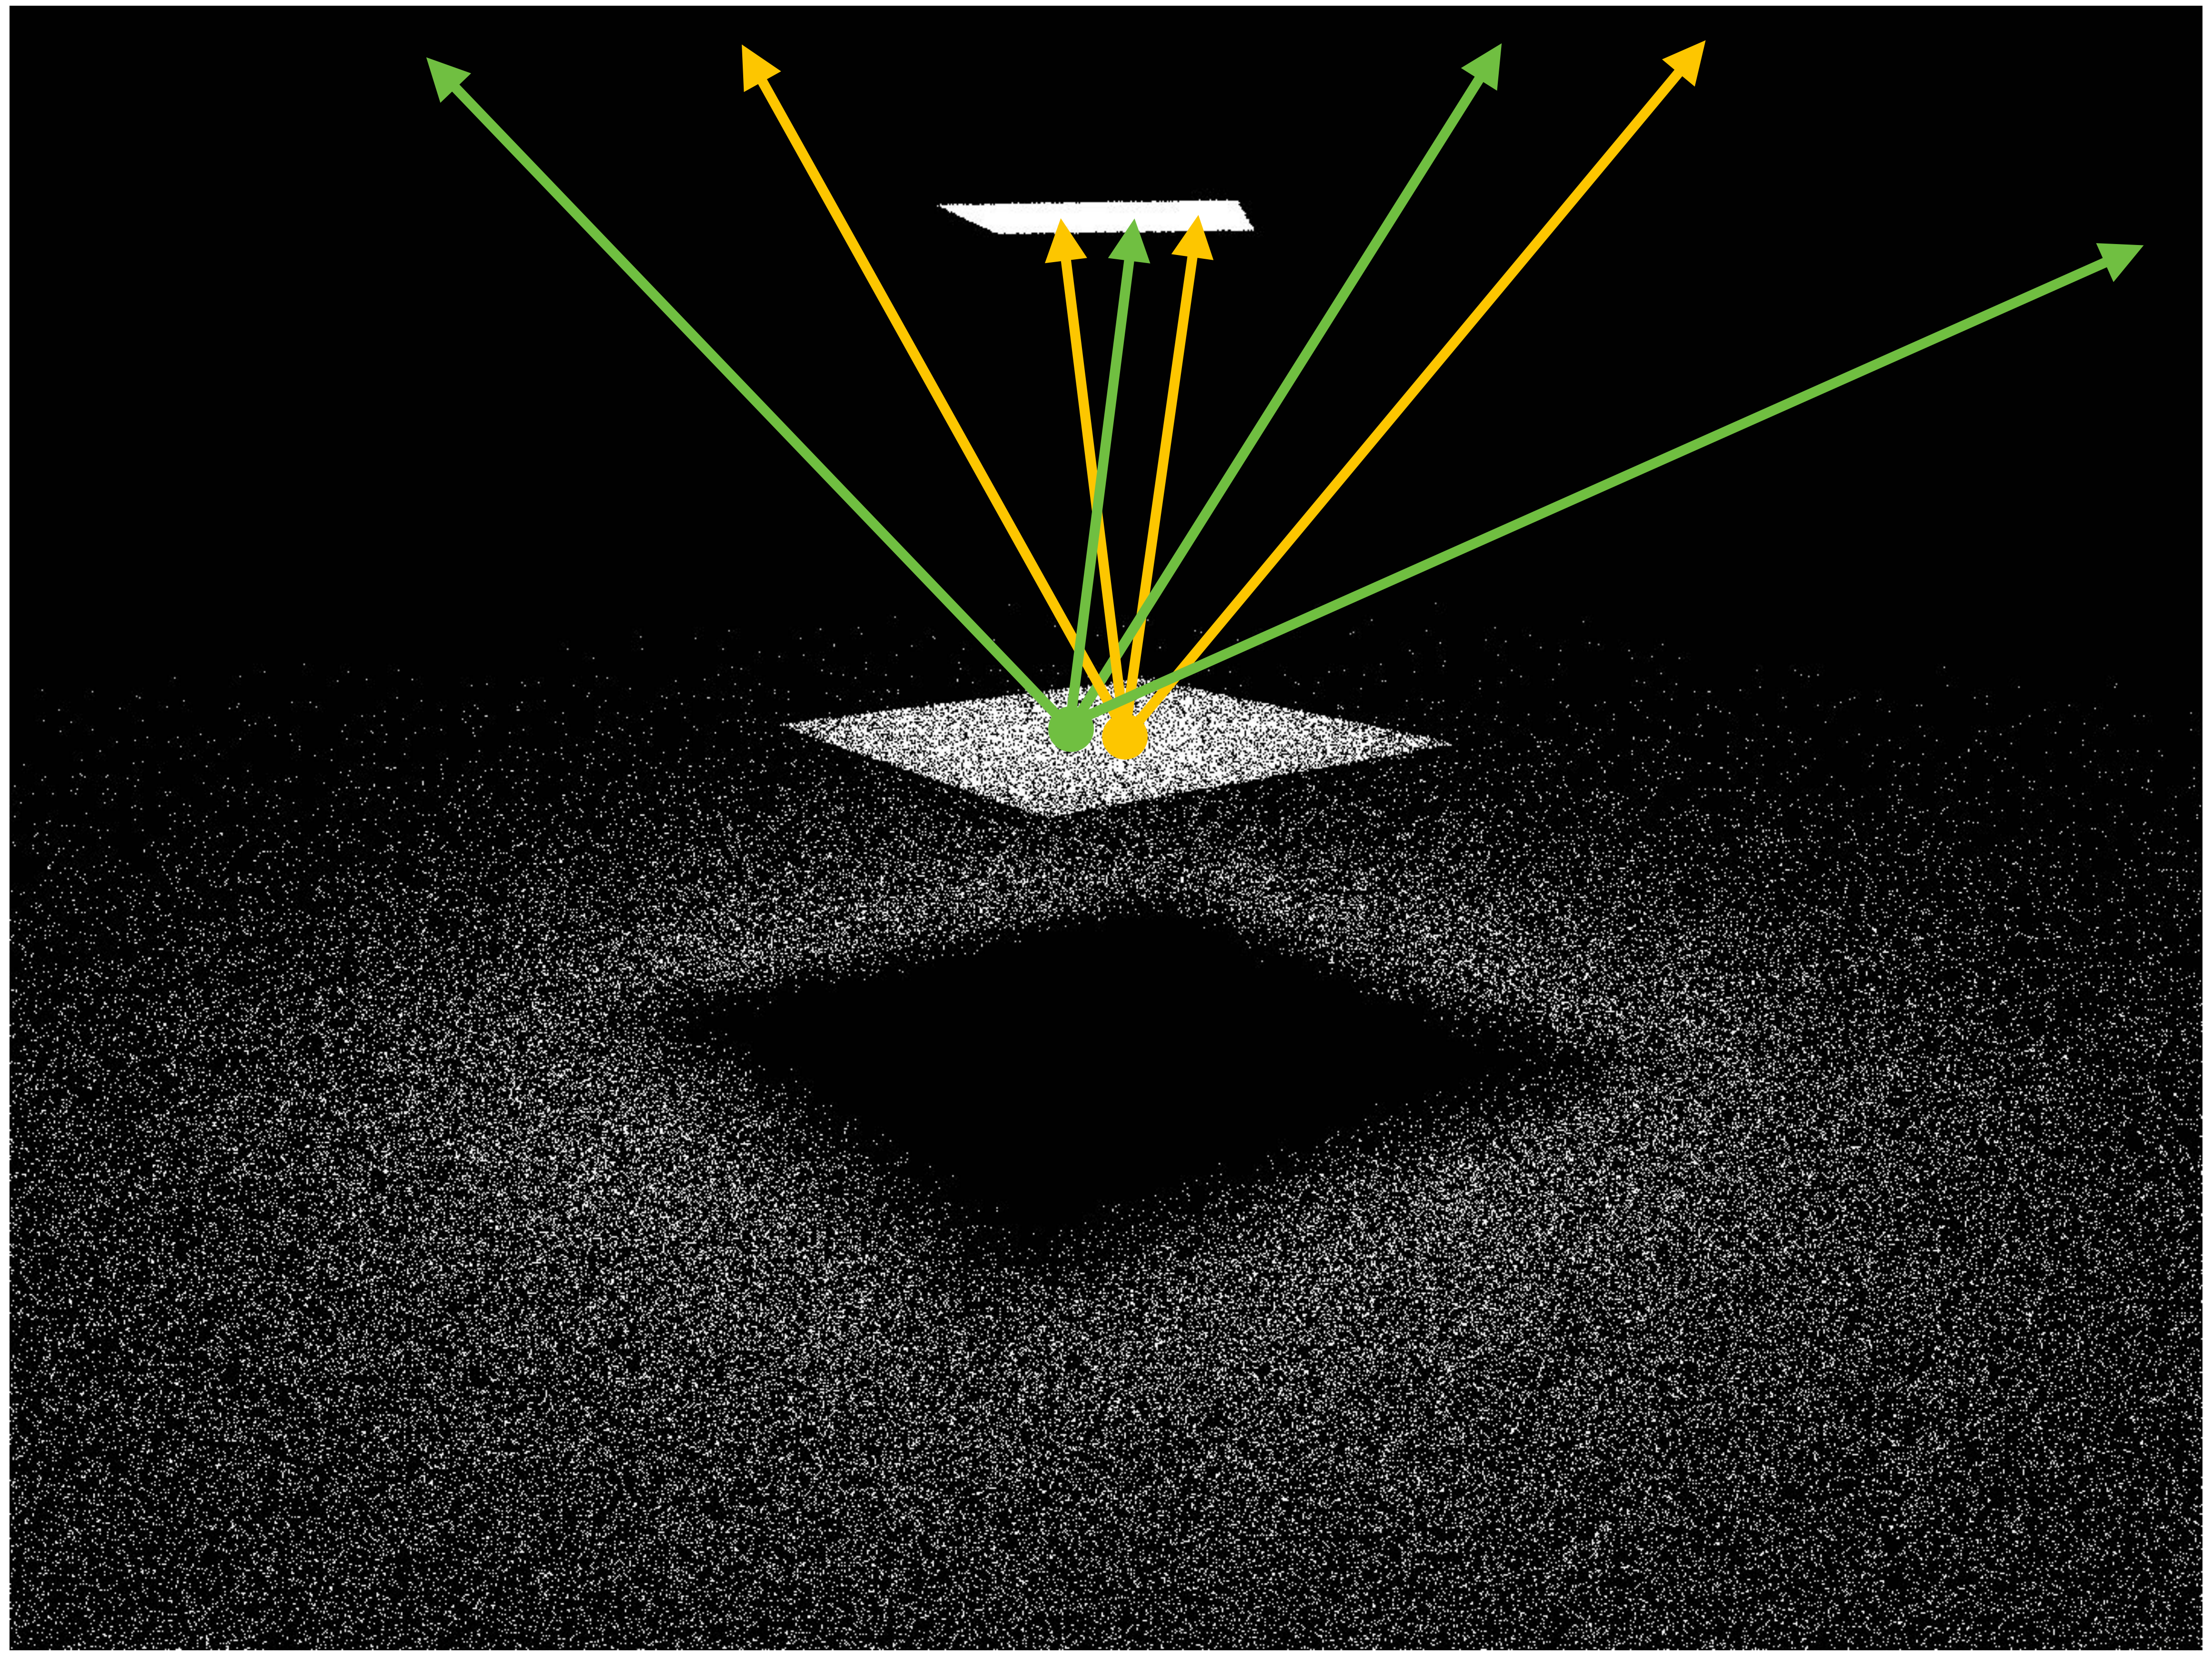

16 samples to estimate incoming irradiance
(Uniformly sampled from hemisphere)

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**

Direct lighting estimate: uniform hemisphere sampling

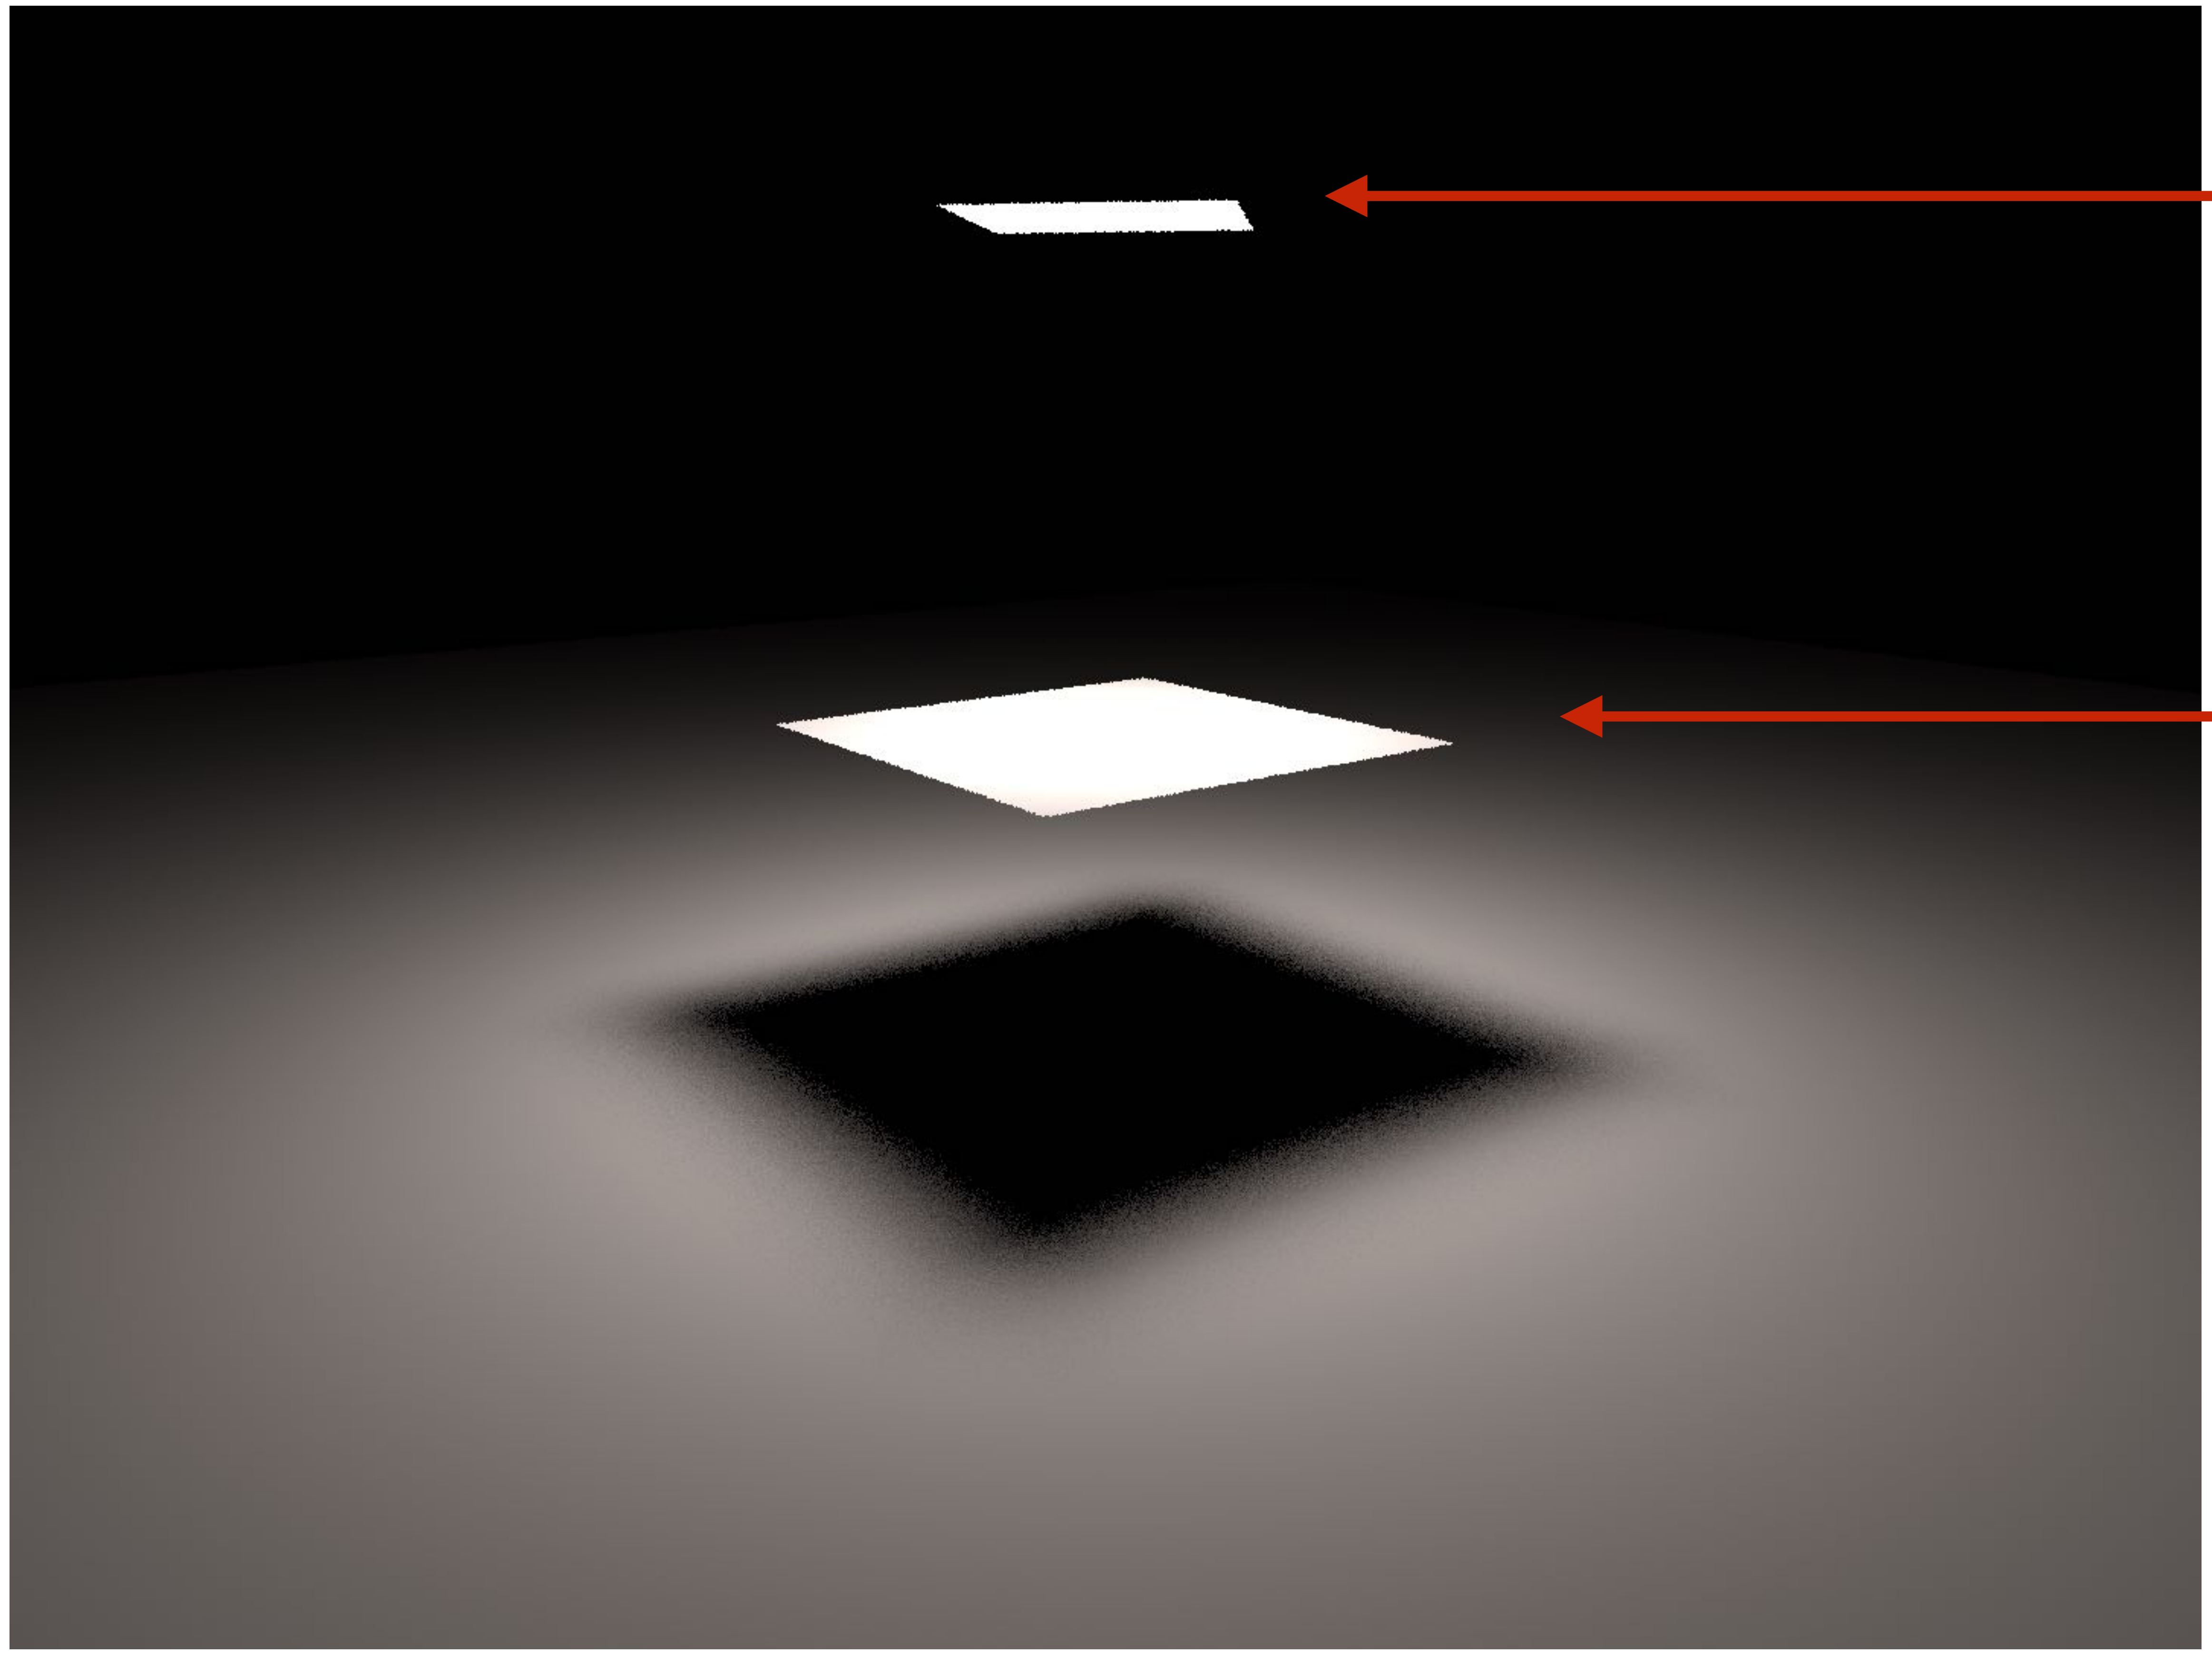

Light source

Occluder
(blocks light)

1000's of samples
(Uniformly sampled from hemisphere)

Direct lighting: only sample center of light

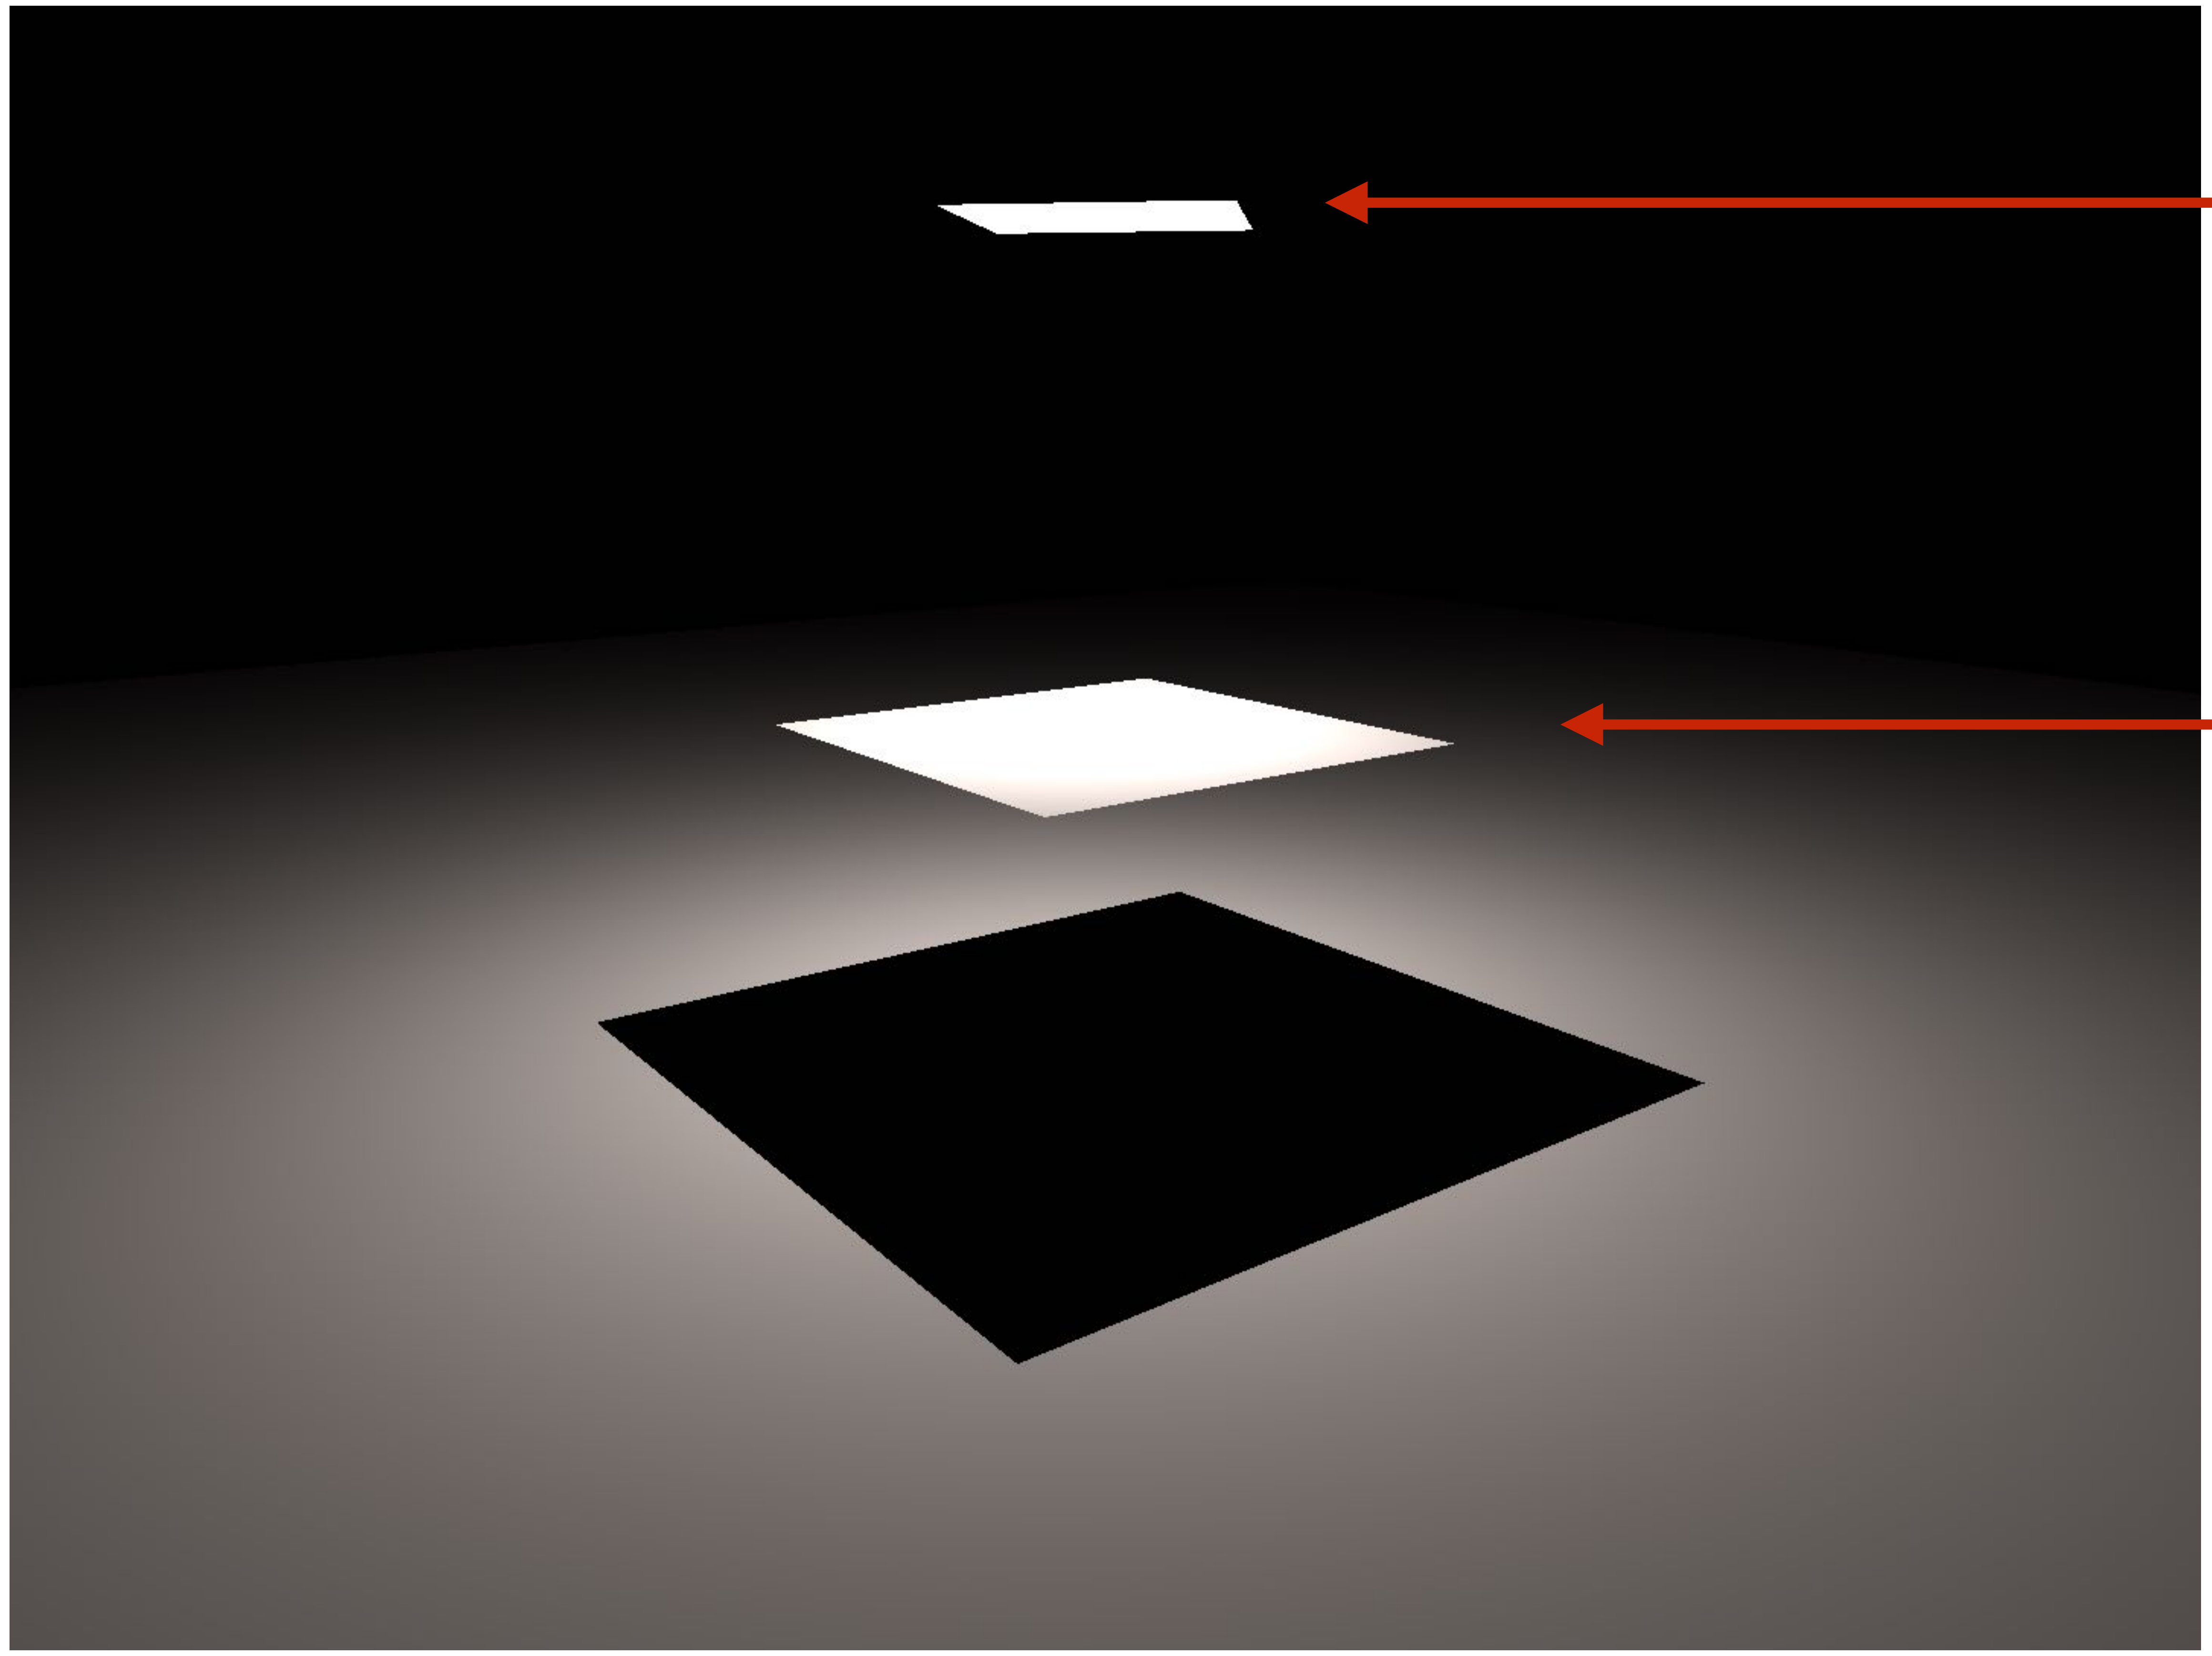

1 light sample, always sample center of light
(Notice “hard shadow”… what you’d expect from a point light source, not an area light source)

Q. Why is there no “noise”?

Summary: Monte Carlo integration

- **Monte Carlo estimator**
  - Estimate integral of function by evaluating function at N random sample points in its domain
  - For the special case of uniform sampling a N-dimensional domain Ω

Let D be the size of the integration domain

$$E[F_N] = E\left[\frac{D}{N}\sum_{i=1}^{N} f(X_i)\right] = \int_{\Omega} f(x)\,\mathrm{d}x$$

- **The estimator is computed by a ray tracer!**

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

- **Coming soon: importance sampling = picking good samples to reduce variance**

Acknowledgements

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