Week 3 Quizzes Review

Computer Graphics and Imaging
UC Berkeley CS184
Summer 2020

Q3

1 Point

(Select all that are correct) In recursive ray tracing, secondary rays are traced from the primary ray's intersection point until

- [ ] A non-specular surface is hit
- [ ] The camera is hit
- [ ] The maximum level of recursion is reached
- [ ] The secondary ray leaves the camera field of view

Recursive Ray Tracing

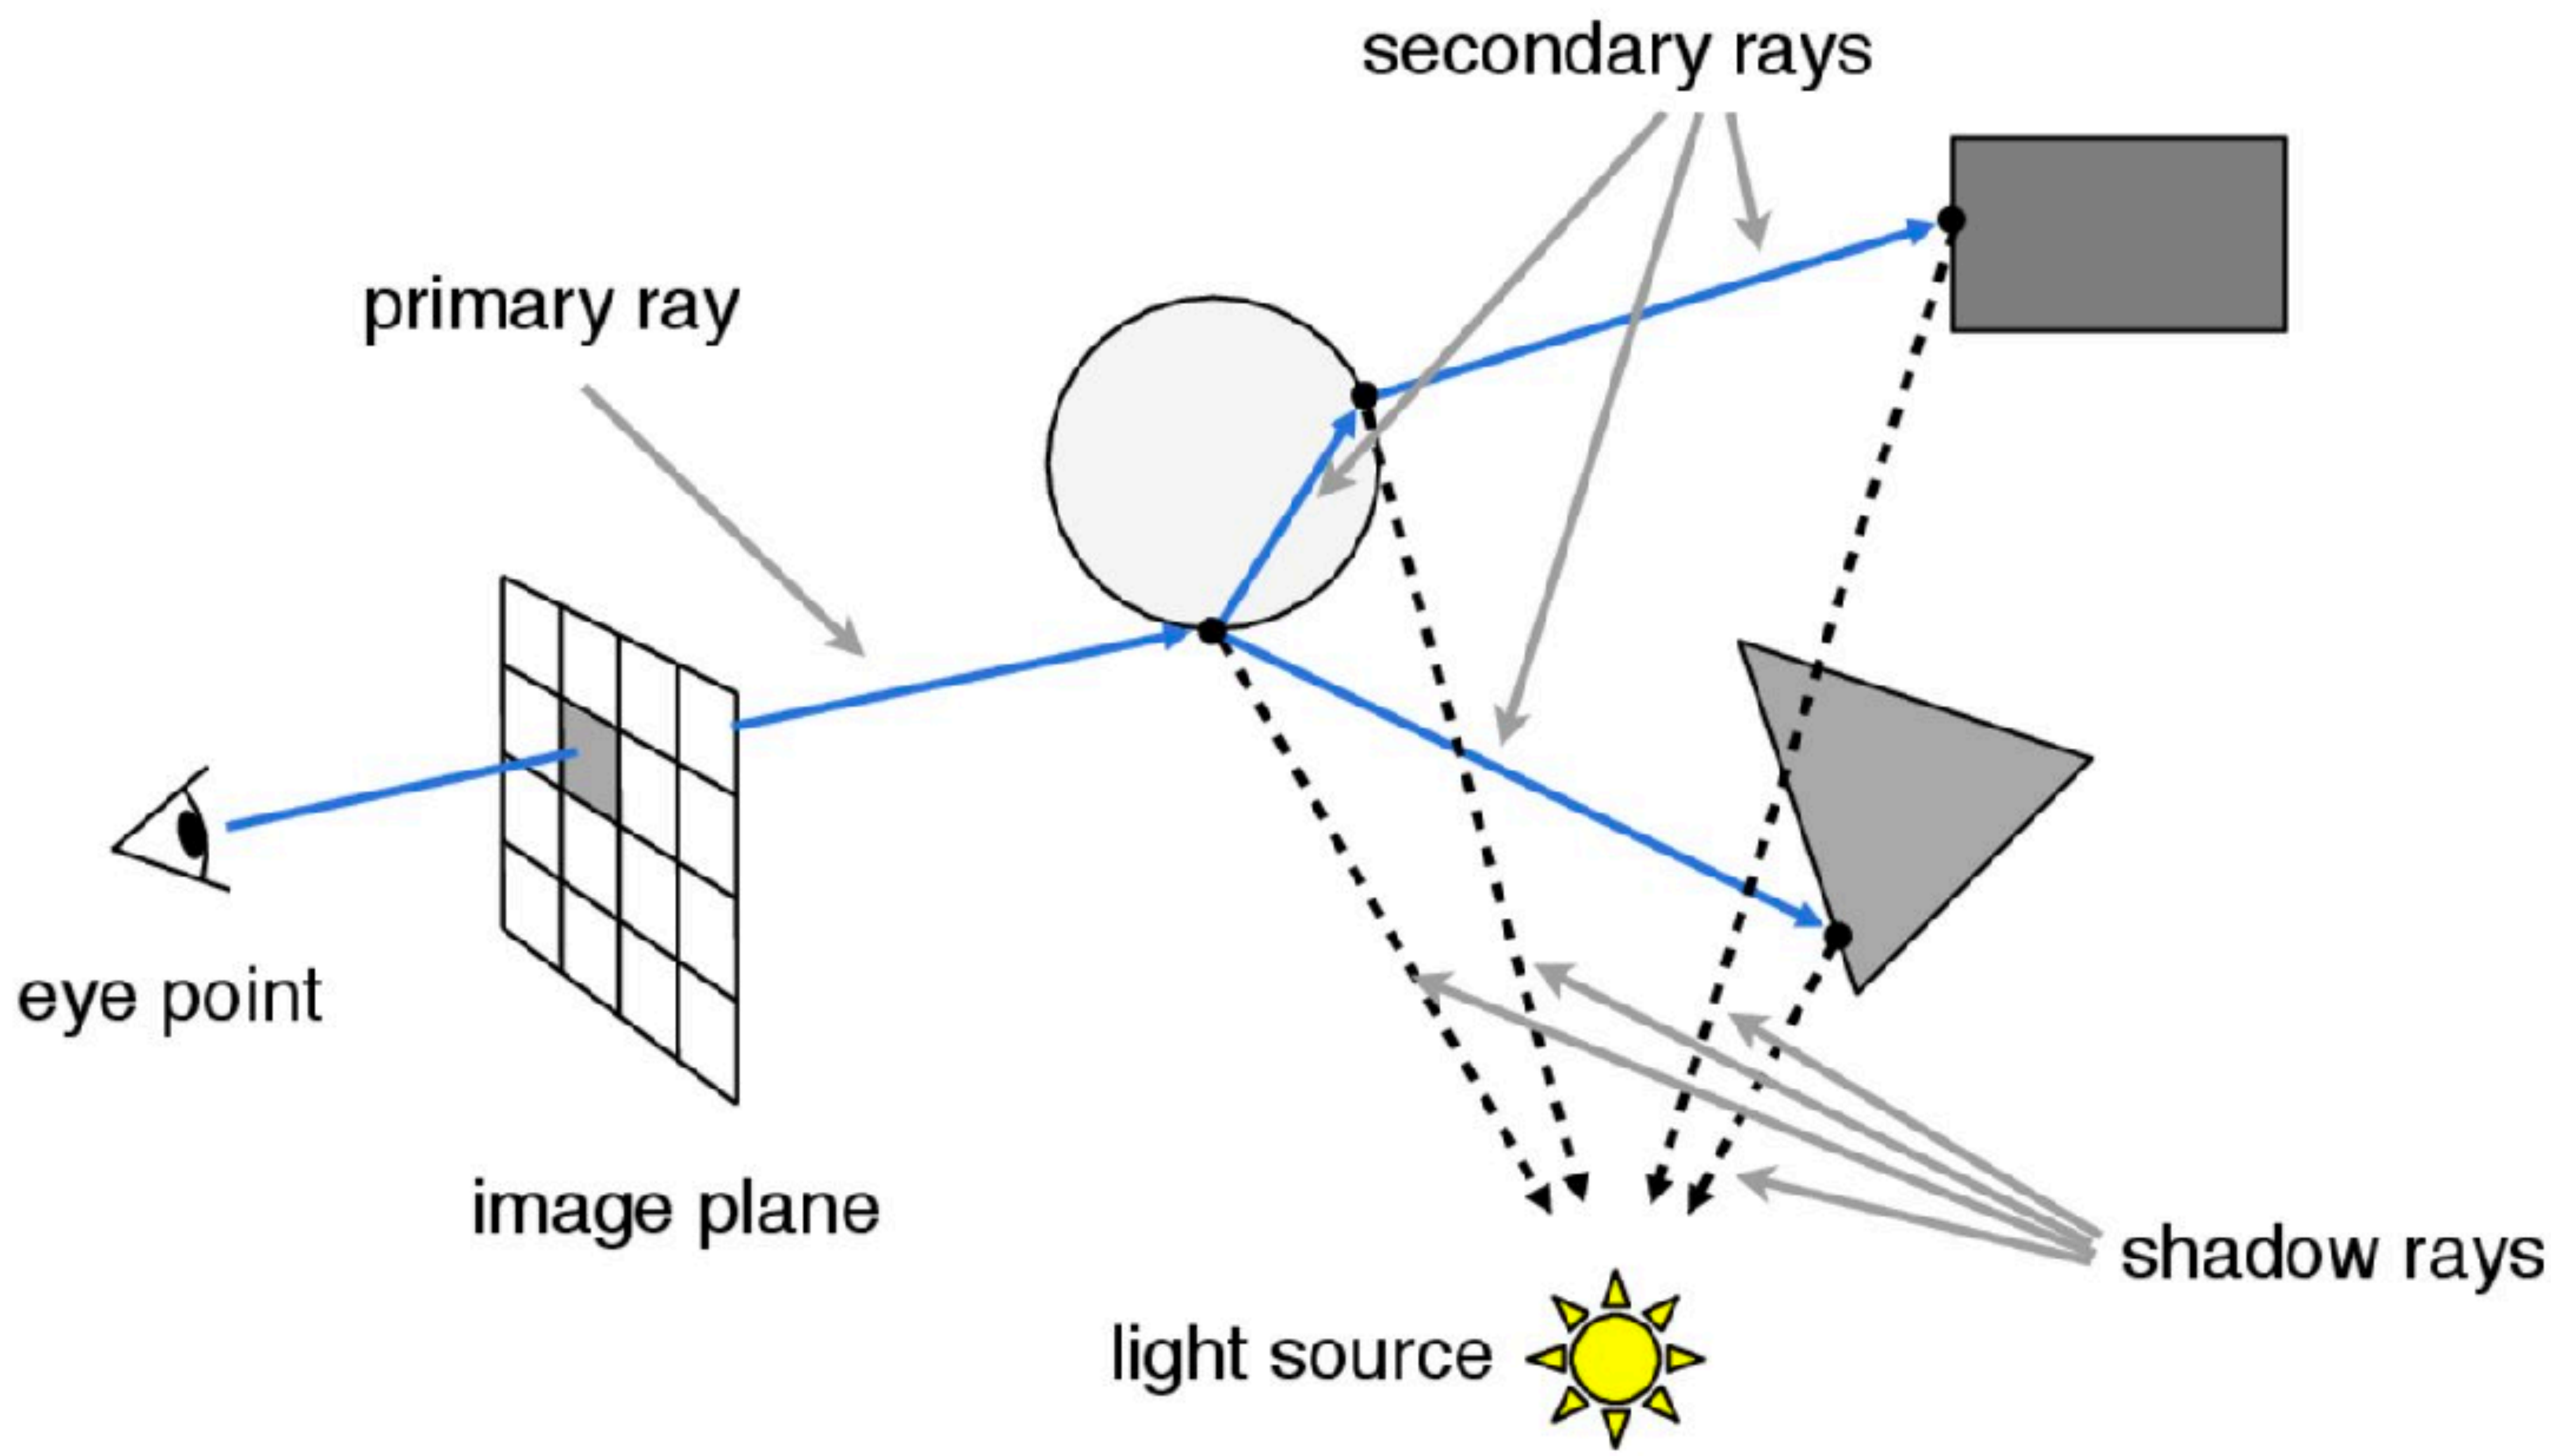

- Trace secondary rays recursively until hit a non-specular surface (or max desired levels of recursion)
- At each hit point, trace shadow rays to test light visibility (no contribution if blocked)
- Final pixel color is weighted sum of contributions along rays, as shown
- Gives more sophisticated effects (e.g. specular reflection, refraction, shadows), but we will go much further to derive a physically-based illumination model

Q3

1 Point

(Select all that are correct) In recursive ray tracing, secondary rays are traced from the primary ray's intersection point until

- [x] A non-specular surface is hit
- [ ] The camera is hit
- [x] The maximum level of recursion is reached
- [ ] The secondary ray leaves the camera field of view

Q6

1 Point

In order to intersect a ray with a sphere, we must solve a quadratic equation. Select all true statements:

- [ ] If the discriminant is negative, the ray will intersect the sphere 3 times
- [ ] A ray can intersect a sphere at most two times
- [ ] If there are two solutions to the equation for t, one negative and one positive, then the ray origin is inside the sphere

Ray Intersection With Sphere

Ray: $\mathbf{r}(t) = \mathbf{o} + t\,\mathbf{d}, \ 0 \leq t < \infty$

Sphere: $\mathbf{p} : (\mathbf{p} - \mathbf{c})^2 - R^2 = 0$

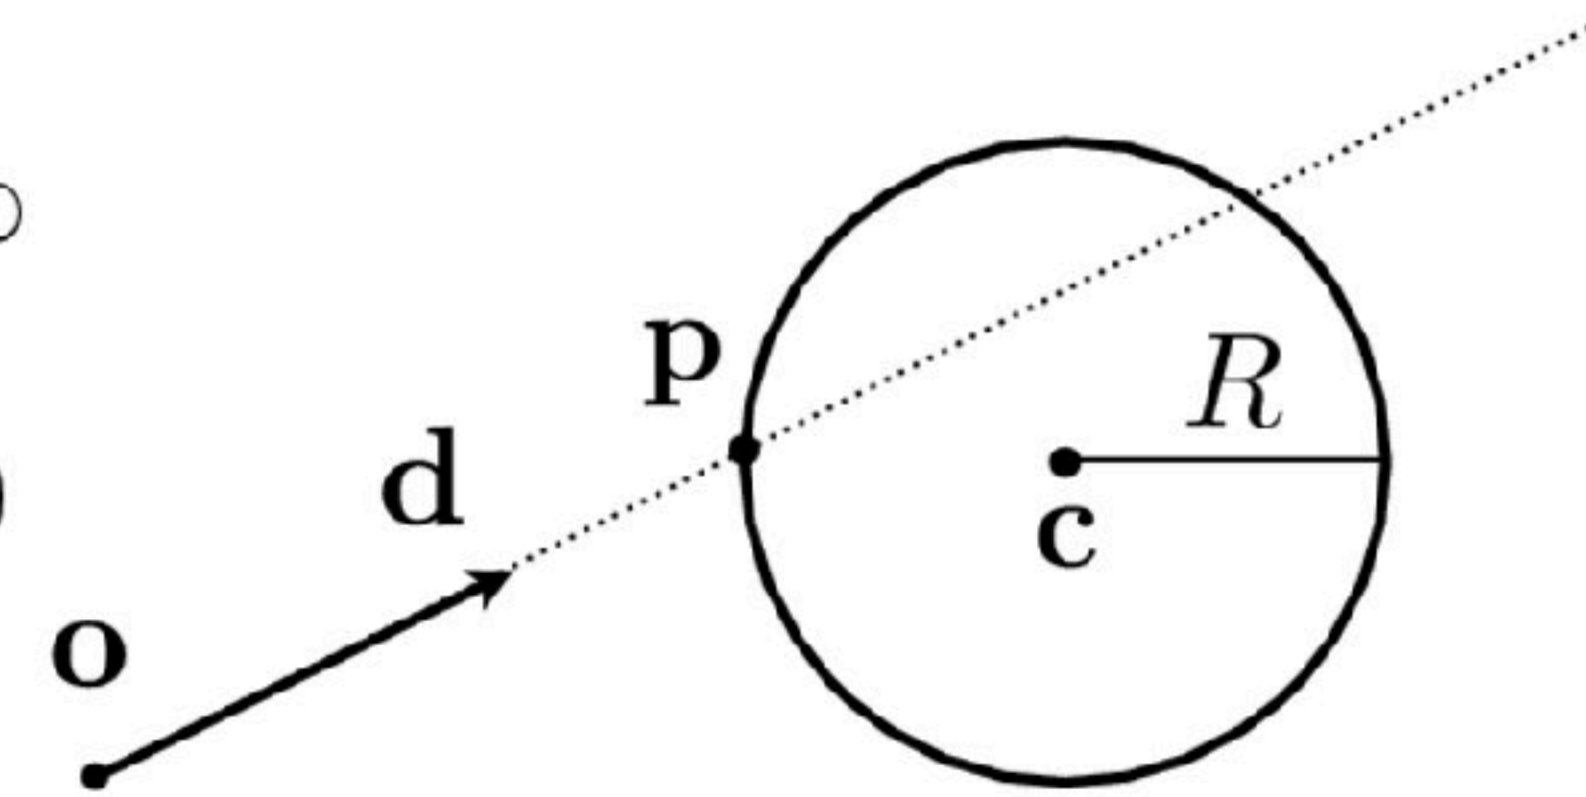

Solve for intersection:

$$(\mathbf{o} + t\,\mathbf{d} - \mathbf{c})^2 - R^2 = 0$$

$$a\,t^2 + b\,t + c = 0, \text{ where}$$

$$a = \mathbf{d} \cdot \mathbf{d}$$

$$b = 2(\mathbf{o} - \mathbf{c}) \cdot \mathbf{d}$$

$$c = (\mathbf{o} - \mathbf{c}) \cdot (\mathbf{o} - \mathbf{c}) - R^2$$

$$t = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

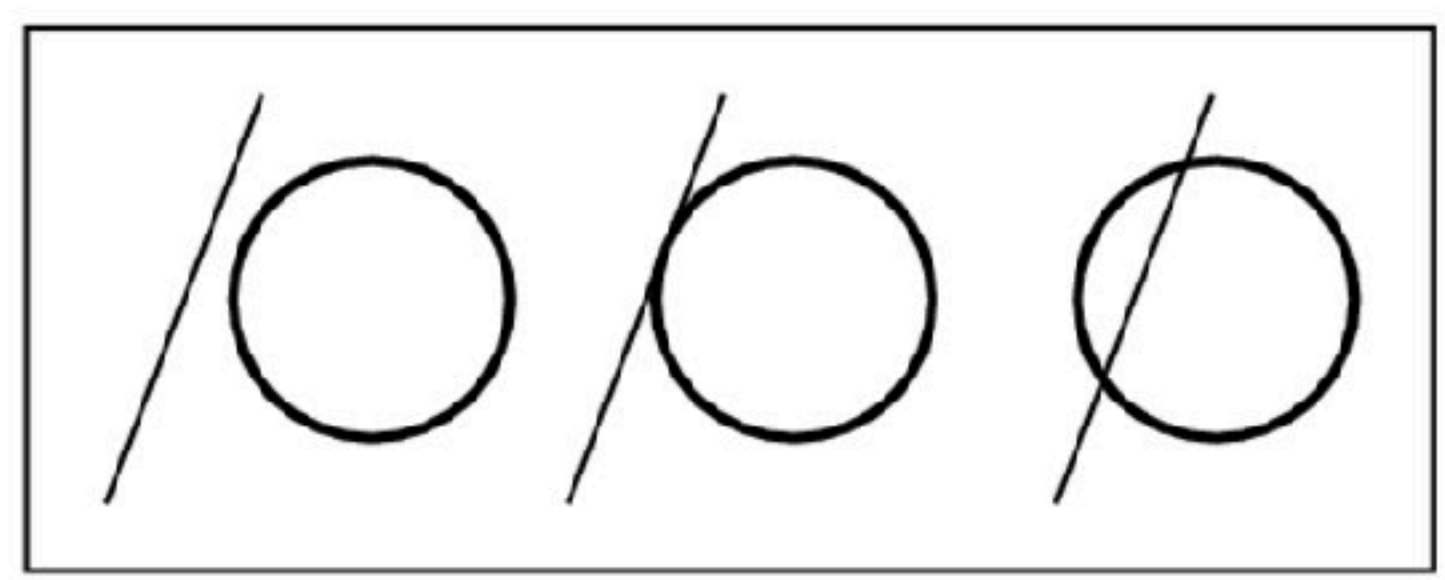

Q6

1 Point

In order to intersect a ray with a sphere, we must solve a quadratic equation. Select all true statements:

- [ ] If the discriminant is negative, the ray will intersect the sphere 3 times
- [x] A ray can intersect a sphere at most two times
- [x] If there are two solutions to the equation for t, one negative and one positive, then the ray origin is inside the sphere

Q5

1 Point

(Select all that apply) Which of the following is true about BVHs?

- [ ] Leaf nodes store a bounding box and a list of objects
- [ ] Internal nodes store a bounding box and a list of child nodes
- [ ] Each level of the BVH recursively splits the set of objects into five subsets
- [ ] In a given level, the bounding boxes are not allowed to overlap

Bounding Volume Hierarchy (BVH)

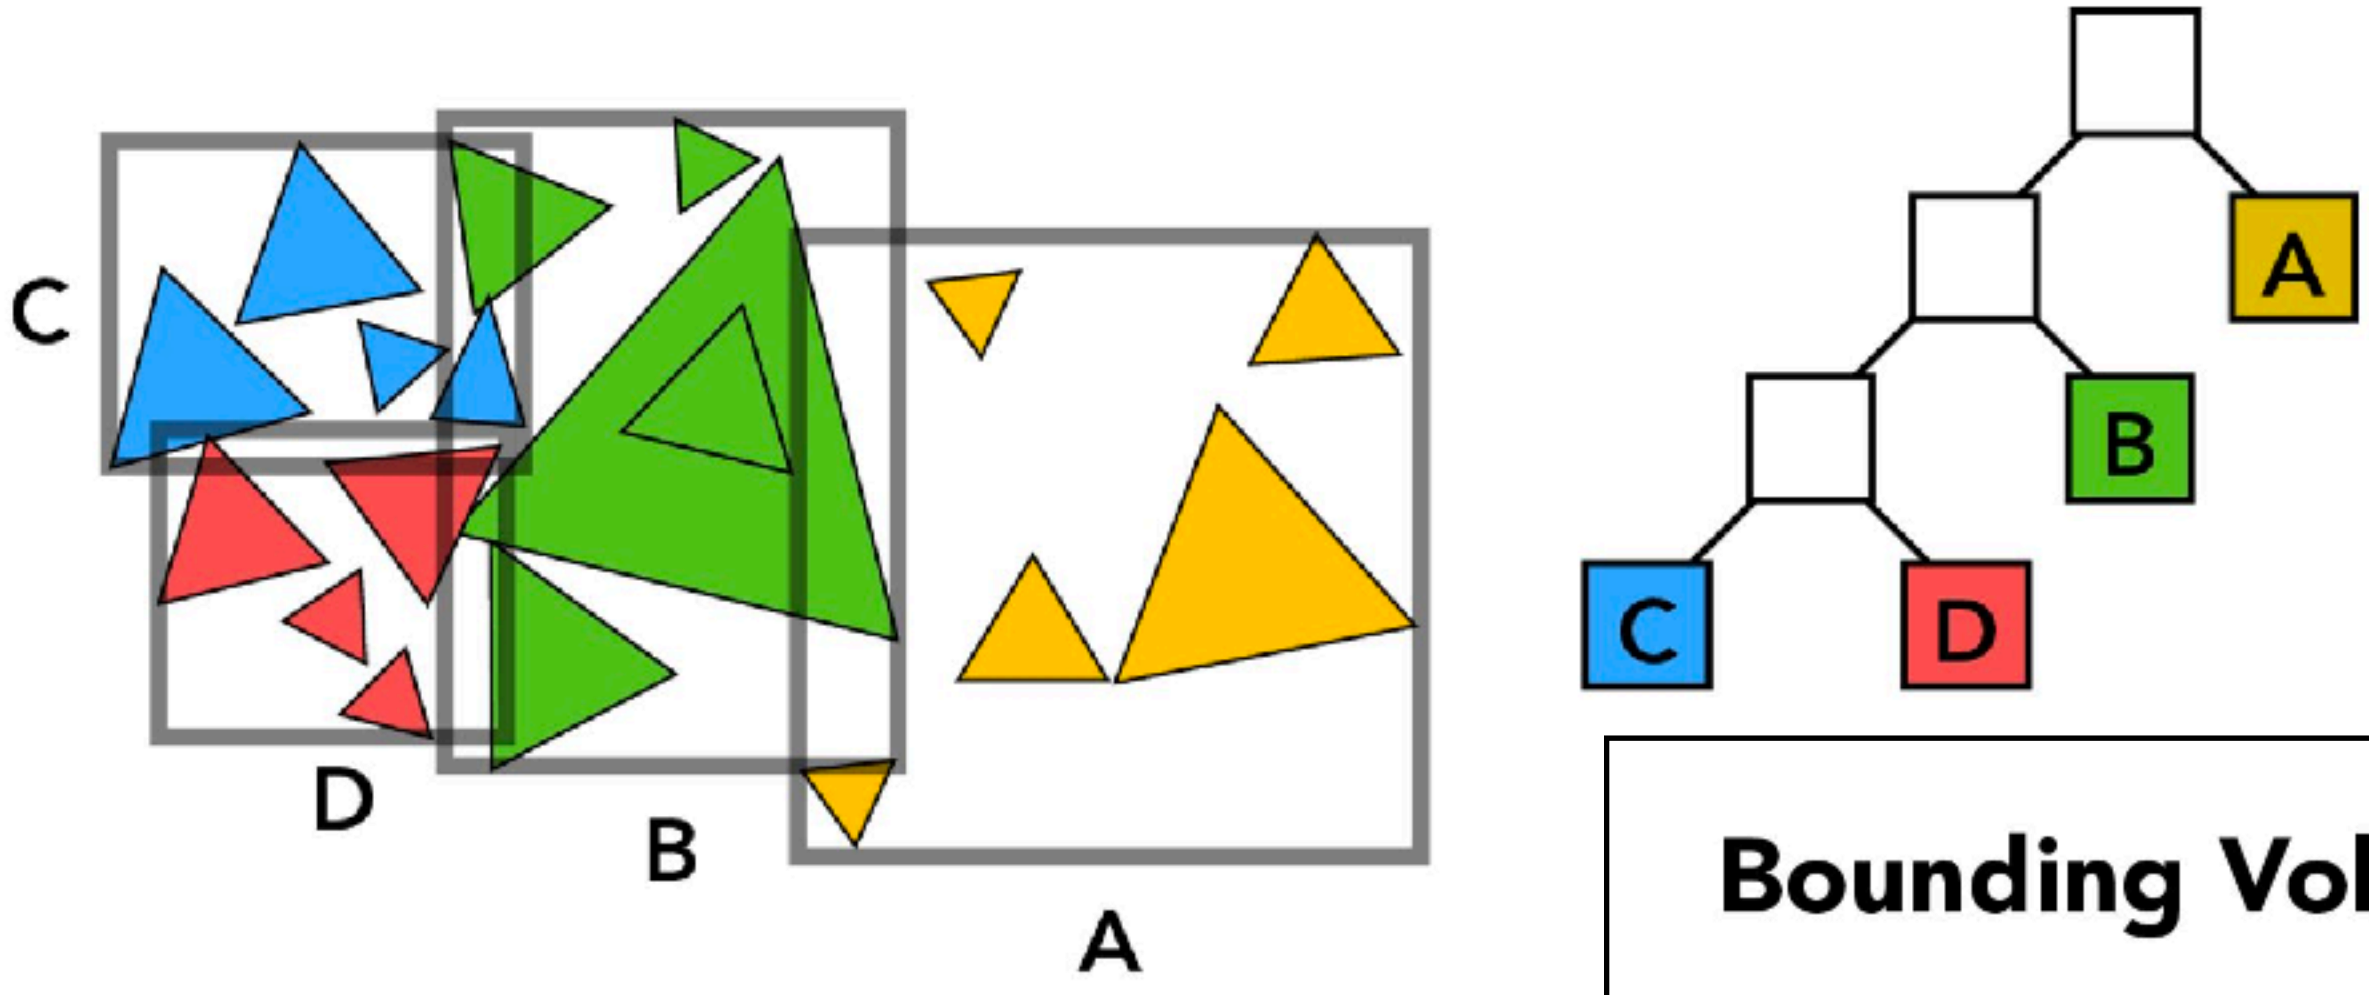

Bounding Volume Hierarchy (BVH)

Internal nodes store

- Bounding box
- Children: reference to child nodes

Leaf nodes store

- Bounding box
- List of objects

Nodes represent subset of primitives in scene

- All objects in subtree

Q5

1 Point

(Select all that apply) Which of the following is true about BVHs?

- [x] Leaf nodes store a bounding box and a list of objects
- [x] Internal nodes store a bounding box and a list of child nodes
- [ ] Each level of the BVH recursively splits the set of objects into five subsets
- [ ] In a given level, the bounding boxes are not allowed to overlap

Q6

1 Point

Which of the following are true?

- [ ] Irradiance times surface area is power
- [ ] According to Lambert's law, irradiance remains constant as a surface changes angle relative to a light source
- [ ] To calculate irradiance at a point, you can integrate cosine-weighted incident radiance over the hemisphere

Irradiance

Definition: The irradiance (illuminance) is the power per unit area incident on a surface point.

$$E(\mathbf{x}) \equiv \frac{\mathrm{d}\Phi(\mathbf{x})}{\mathrm{d}A}$$

$$\left[\frac{\mathrm{W}}{\mathrm{m}^2}\right] \quad \left[\frac{\mathrm{lm}}{\mathrm{m}^2} = \mathrm{lux}\right]$$

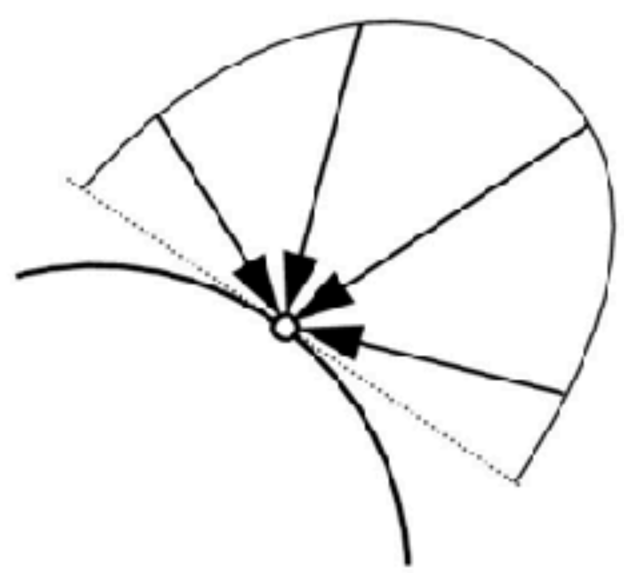

Lambert's Cosine Law

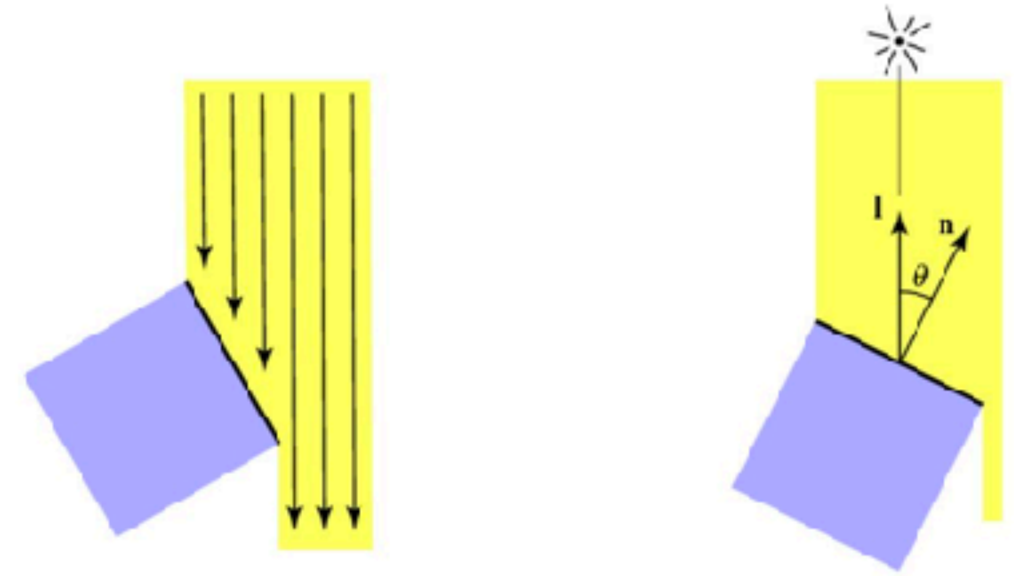

Top face of cube receives a certain amount of power	Top face of 60° rotated cube receives half power	In general, power per unit area is proportional to $\cos\theta = l \cdot n$
$E = \frac{\Phi}{A}$	$E = \frac{1}{2}\frac{\Phi}{A}$	$E = \frac{\Phi}{A}\cos\theta$

Irradiance at surface is proportional to cosine of angle between light direction and surface normal.

Irradiance from the Environment

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

$$dE(\mathrm{p},\omega) = L_i(\mathrm{p},\omega)\cos\theta\,\mathrm{d}\omega$$ ← Contribution to irradiance from light arriving from direction ω

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

Light meter

Hemisphere: H^2

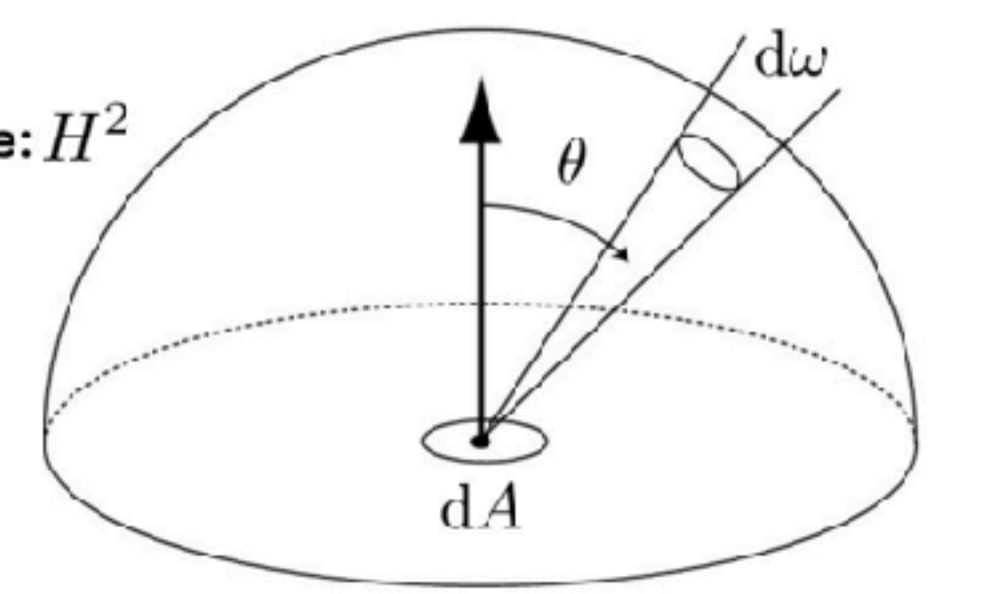

Q6

1 Point

Which of the following are true?

- [x] Irradiance times surface area is power
- [ ] According to Lambert's law, irradiance remains constant as a surface changes angle relative to a light source
- [x] To calculate irradiance at a point, you can integrate cosine-weighted incident radiance over the hemisphere

Q7

1 Point

Which of the following are true?

- [ ] Radiance is measured in lux
- [ ] Radiance is constant along a ray in a vacuum
- [ ] Incident and exitant radiance at a point on a surface are aways equal

Radiometric & Photometric Terms & Units

Physics		Radiometry	Units	Photometry	Units
Energy	Q	Radiant Energy	Joules (W·sec)	Luminous Energy	Lumen·sec
Flux (Power)	Φ	Radiant Power	W	Luminous Power	Lumen (Candela sr)
Angular Flux Density	I	Radiant Intensity	W/sr	Luminous Intensity	Candela (Lumen/sr)
Spatial Flux Density	E	Irradiance (in) Radiosity (out)	W/m²	Illuminance (in) Luminosity (out)	Lux (Lumen/m²)
Spatio-Angular Flux Density	L	Radiance	W/m²/sr	Luminance	Nit (Candela/m²)

"Thus one nit is one lux per steradian is one candela per square meter is one lumen per square meter per steradian. Got it?" — *James Kajiya*

Radiance

Light Traveling Along A Ray

1. Radiance is the fundamental field quantity that describes the distribution of light in an environment
   - Radiance is the quantity associated with a ray
   - Rendering is all about computing radiance
2. Radiance is invariant along a ray in a vacuum

Incident & Exiting Surface Radiance Differ!

Need to distinguish between incident radiance and exitant radiance functions at a point on a surface

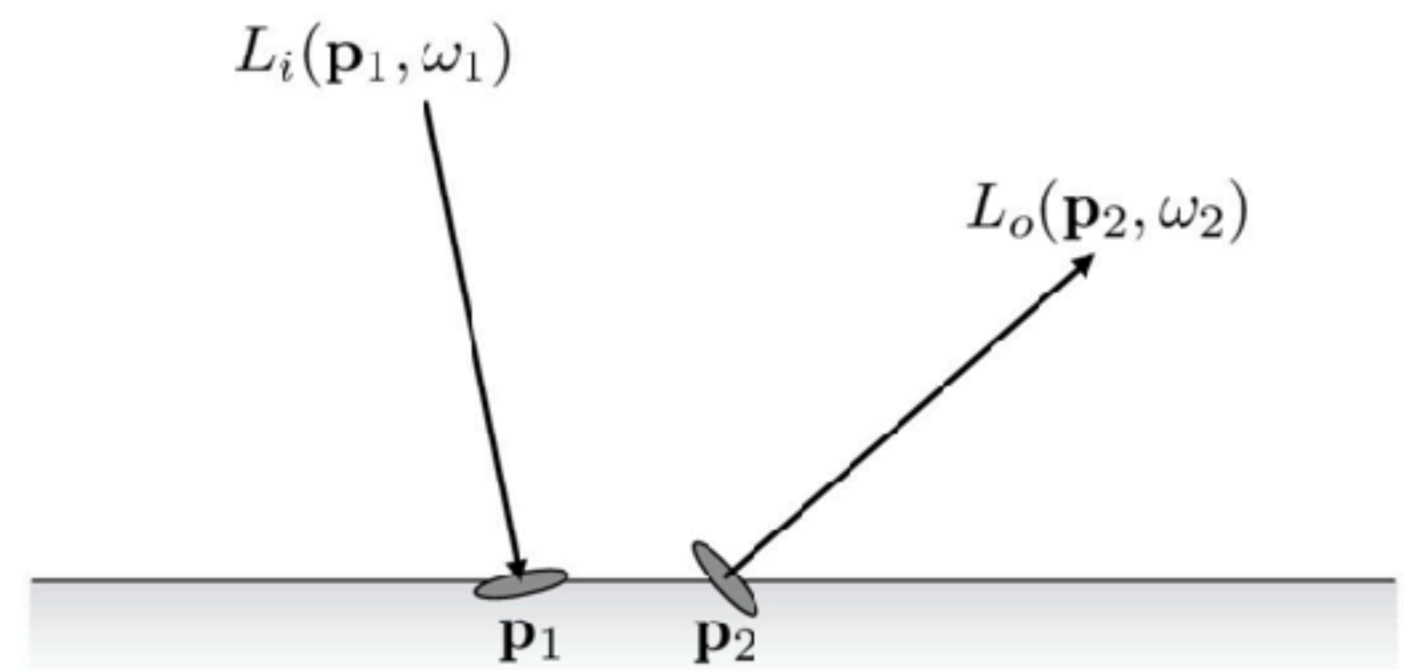

In general: $L_i(\mathbf{p}, \omega) \neq L_o(\mathbf{p}, \omega)$

Q7

1 Point

Which of the following are true?

- [ ] Radiance is measured in lux
- [x] Radiance is constant along a ray in a vacuum
- [ ] Incident and exitant radiance at a point on a surface are aways equal

Q1

1 Point

What does the "curse of dimensionality" refer to?

- ○ Monte Carlo integration cannot be used for high dimensional integrals
- ○ Storing high dimensional data is more expensive
- ○ Number of samples required for numerical integration increases exponentially with dimension
- ○ Trying to evaluate high dimensional integrals is bad luck

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

Q1

1 Point

What does the "curse of dimensionality" refer to?

- ○ Monte Carlo integration cannot be used for high dimensional integrals
- ○ Storing high dimensional data is more expensive
- ◉ Number of samples required for numerical integration increases exponentially with dimension
- ○ Trying to evaluate high dimensional integrals is bad luck

Q3

1 Point

Select all the true statements.

- [ ] Using importance sampling *cannot* increase variance
- [ ] Using importance sampling does not change the expected value of an estimator
- [ ] The best possible importance sampling distribution is the actual function being integrated (normalized)
- [ ] For estimating the integral of f a sample $X \sim p$, an unbiased Monte Carlo estimator would be $f(X) + p(X)$

Importance Sampling - Ideal Probability Distribution

Idea: concentrate selection of random samples in parts of domain where function is large ("high value samples")

Fact: sampling according to $f(x)$ itself would be optimal (though in pratice can only take an educated guess)

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

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

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

Recall definition of variance:

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

$$E[\tilde{f}^2] = \int \left[\frac{f(x)}{p(x)}\right]^2 p(x)\,\mathrm{d}x$$

$$= \int \left[\frac{f(x)}{f(x)/E[f]}\right]^2 \frac{f(x)}{E[f]}\,\mathrm{d}x$$

$$= E[f] \int f(x)\,\mathrm{d}x$$

$$= E^2[f]$$

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

Importance Sampling

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

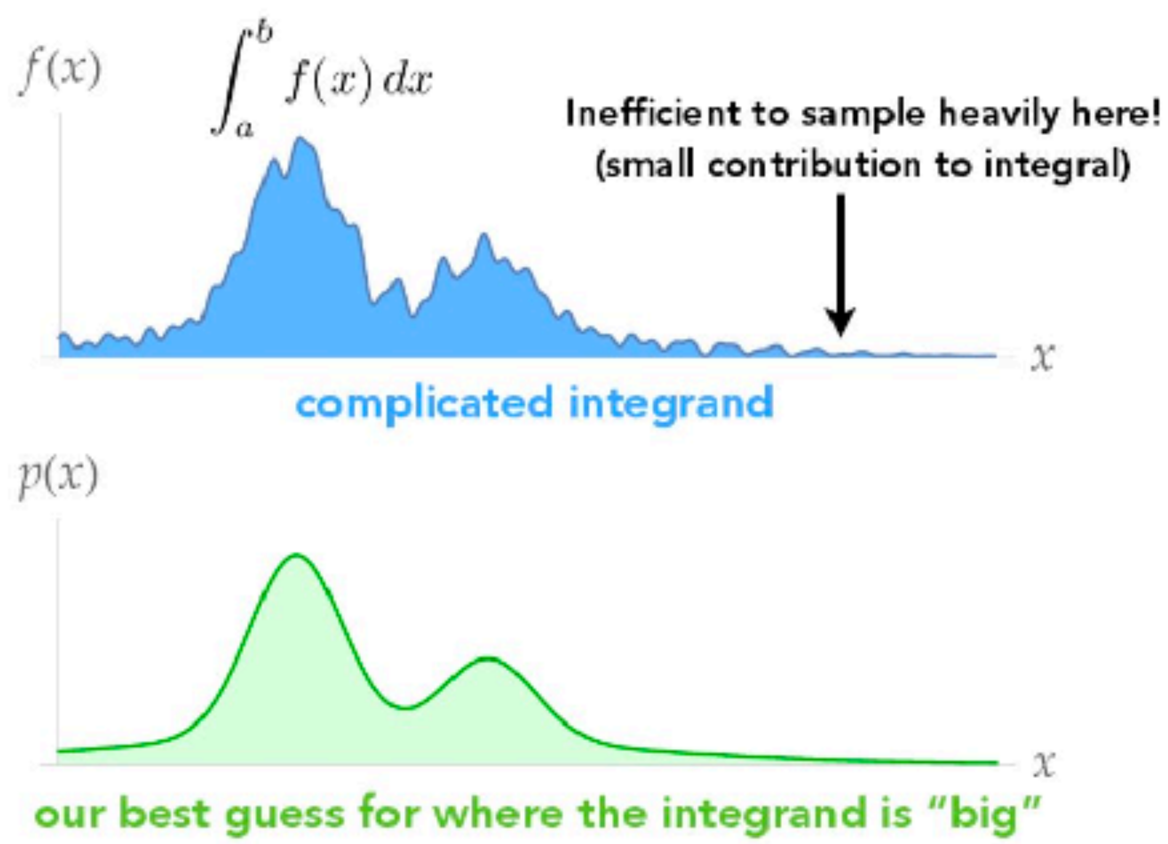

Note: p(x) must be non-zero where f(x) is non-zero

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."

Q3

1 Point

Select all the true statements.

- [ ] Using importance sampling *cannot* increase variance
- [x] Using importance sampling does not change the expected value of an estimator
- [x] The best possible importance sampling distribution is the actual function being integrated (normalized)
- [ ] For estimating the integral of f a sample $X \sim p$, an unbiased Monte Carlo estimator would be $f(X) + p(X)$

Q4

1 Point

Which of the following is *not* an example of an integral found in computer graphics?

- ○ Motion blur
- ○ Intersecting a ray with a triangle
- ○ Calculating the shadow of an area light
- ○ Antialiasing with area sampling

2D Integral: Recall Antialiasing By Area Sampling

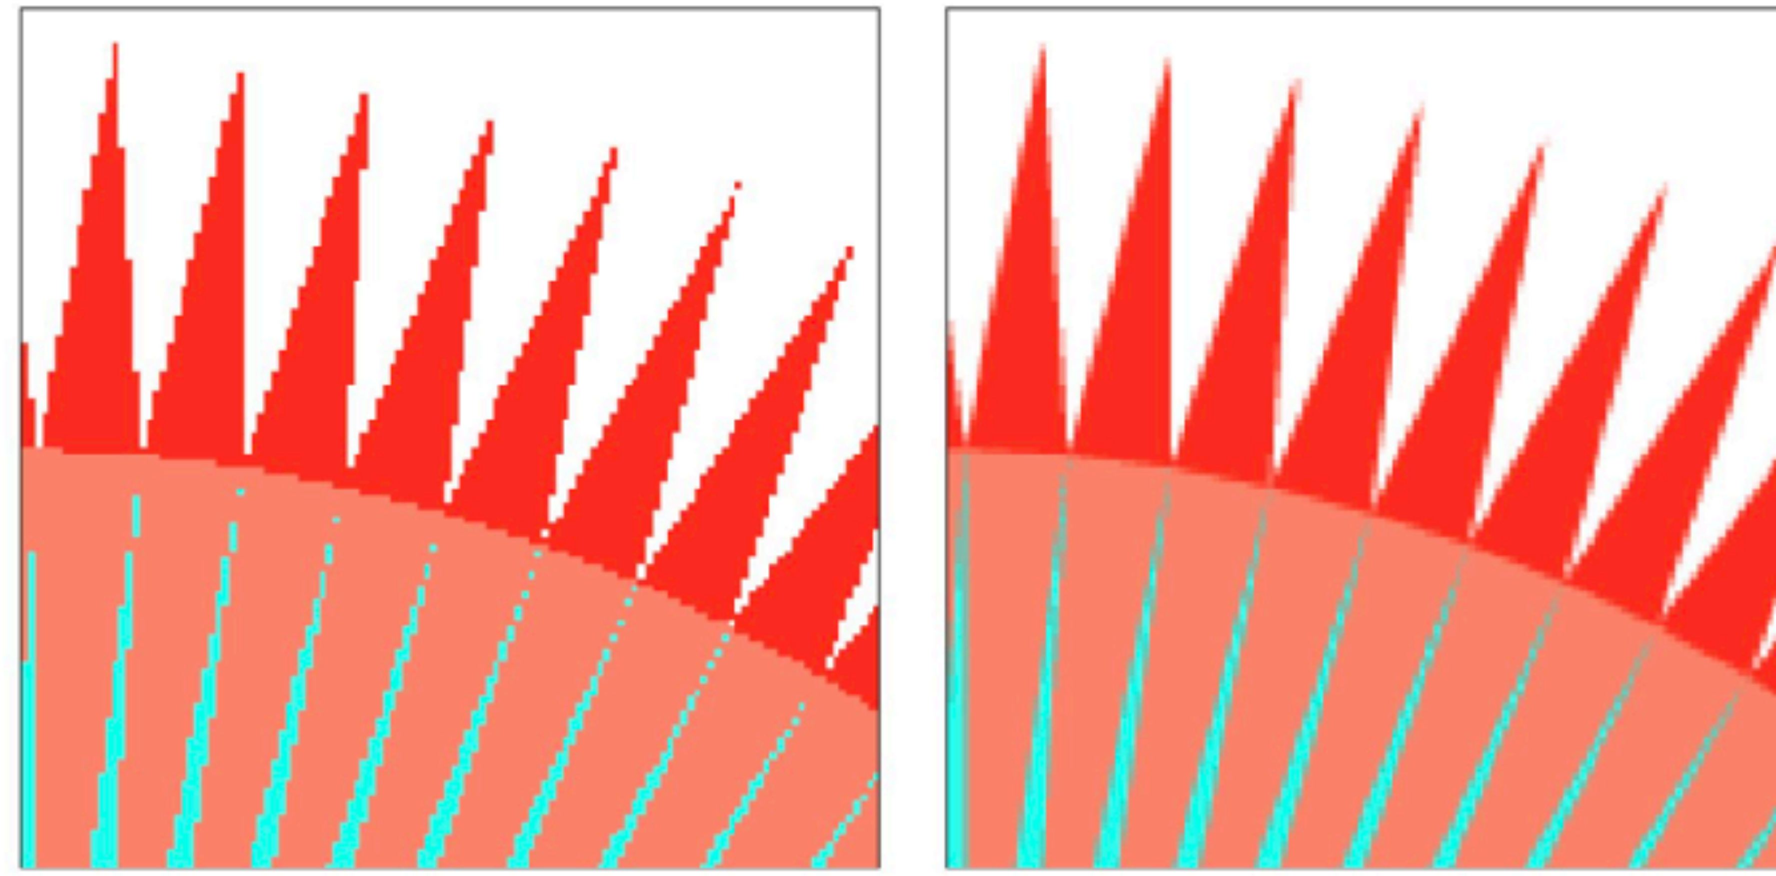

Point sampling

Area sampling

Integrate over 2D area of pixel

3D Integral: Motion Blur

Cook et al. "1984"

Integrate over area of pixel and over exposure time.

Example: Monte Carlo Lighting

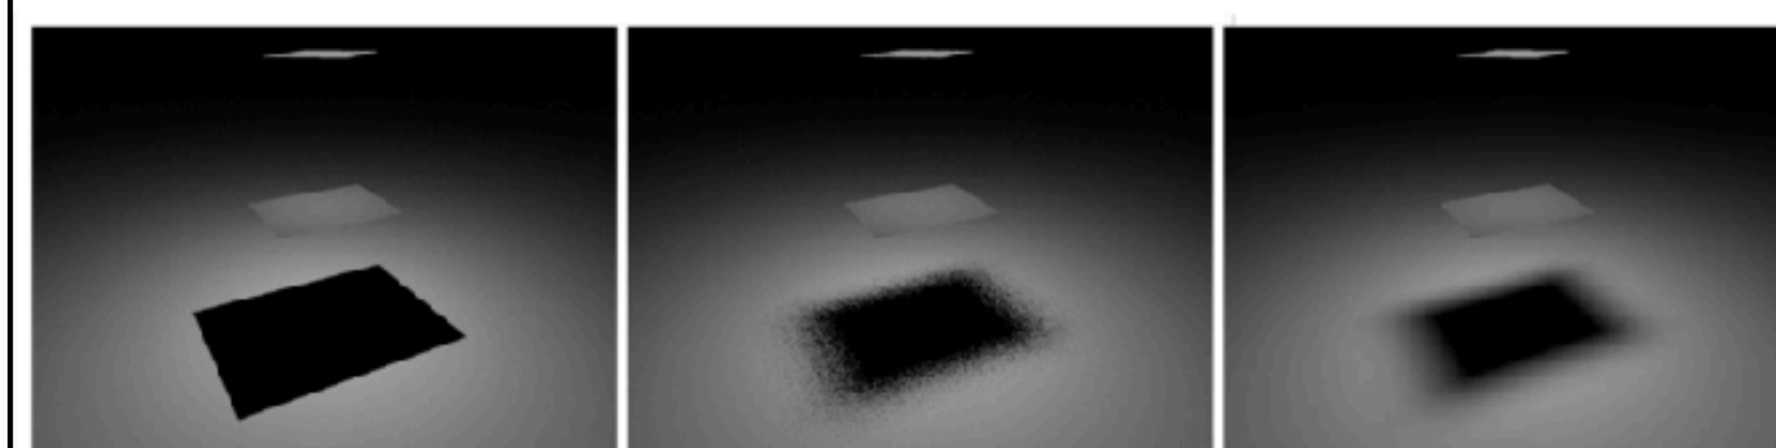

Sample center of light

Sample random point on light

True answer

Q4

1 Point

Which of the following is *not* an example of an integral found in computer graphics?

- ○ Motion blur
- ◉ Intersecting a ray with a triangle
- ○ Calculating the shadow of an area light
- ○ Antialiasing with area sampling

Q5

1 Point

The **variance** of a basic Monte Carlo estimator (uniform sampling) of a d dimensional integral using N random samples is proportional to

- ○ N
- ○ $1/N$
- ○ $1/N^{1/d}$
- ○ $1/d$

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

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

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

Q5

1 Point

The **variance** of a basic Monte Carlo estimator (uniform sampling) of a d dimensional integral using N random samples is proportional to

- ○ N
- ◉ $1/N$
- ○ $1/N^{1/d}$
- ○ $1/d$

Q6

1 Point

To make a soft shadow less noisy, I should use importance sampling to

- ○ direct more random rays toward the camera
- ○ direct more random rays toward the light source
- ○ weight the light source rays more heavily in my Monte Carlo estimate

Solid Angle Sampling vs Area Sampling

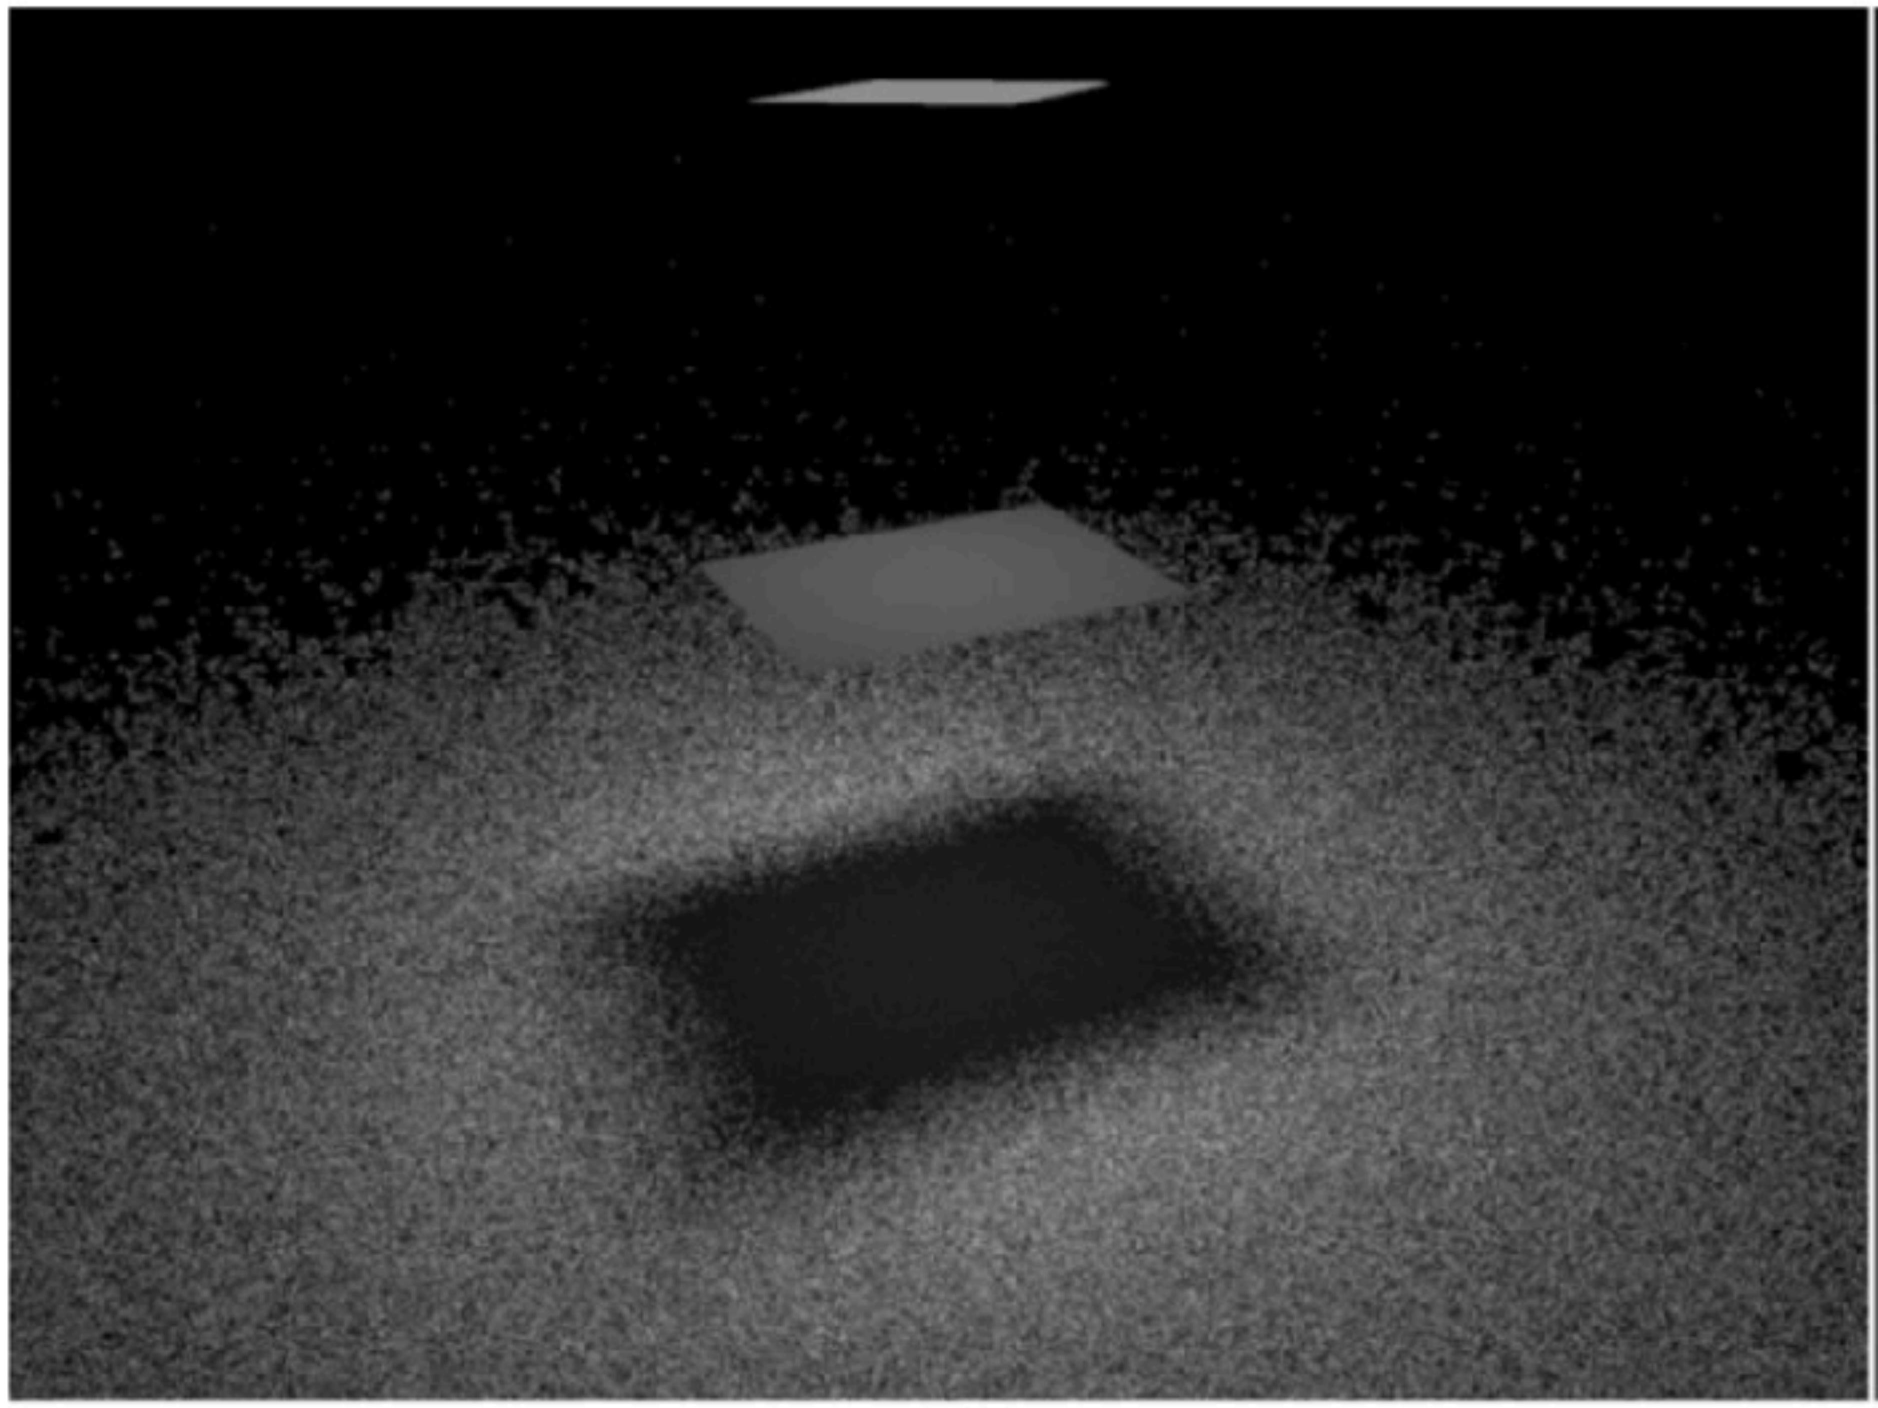

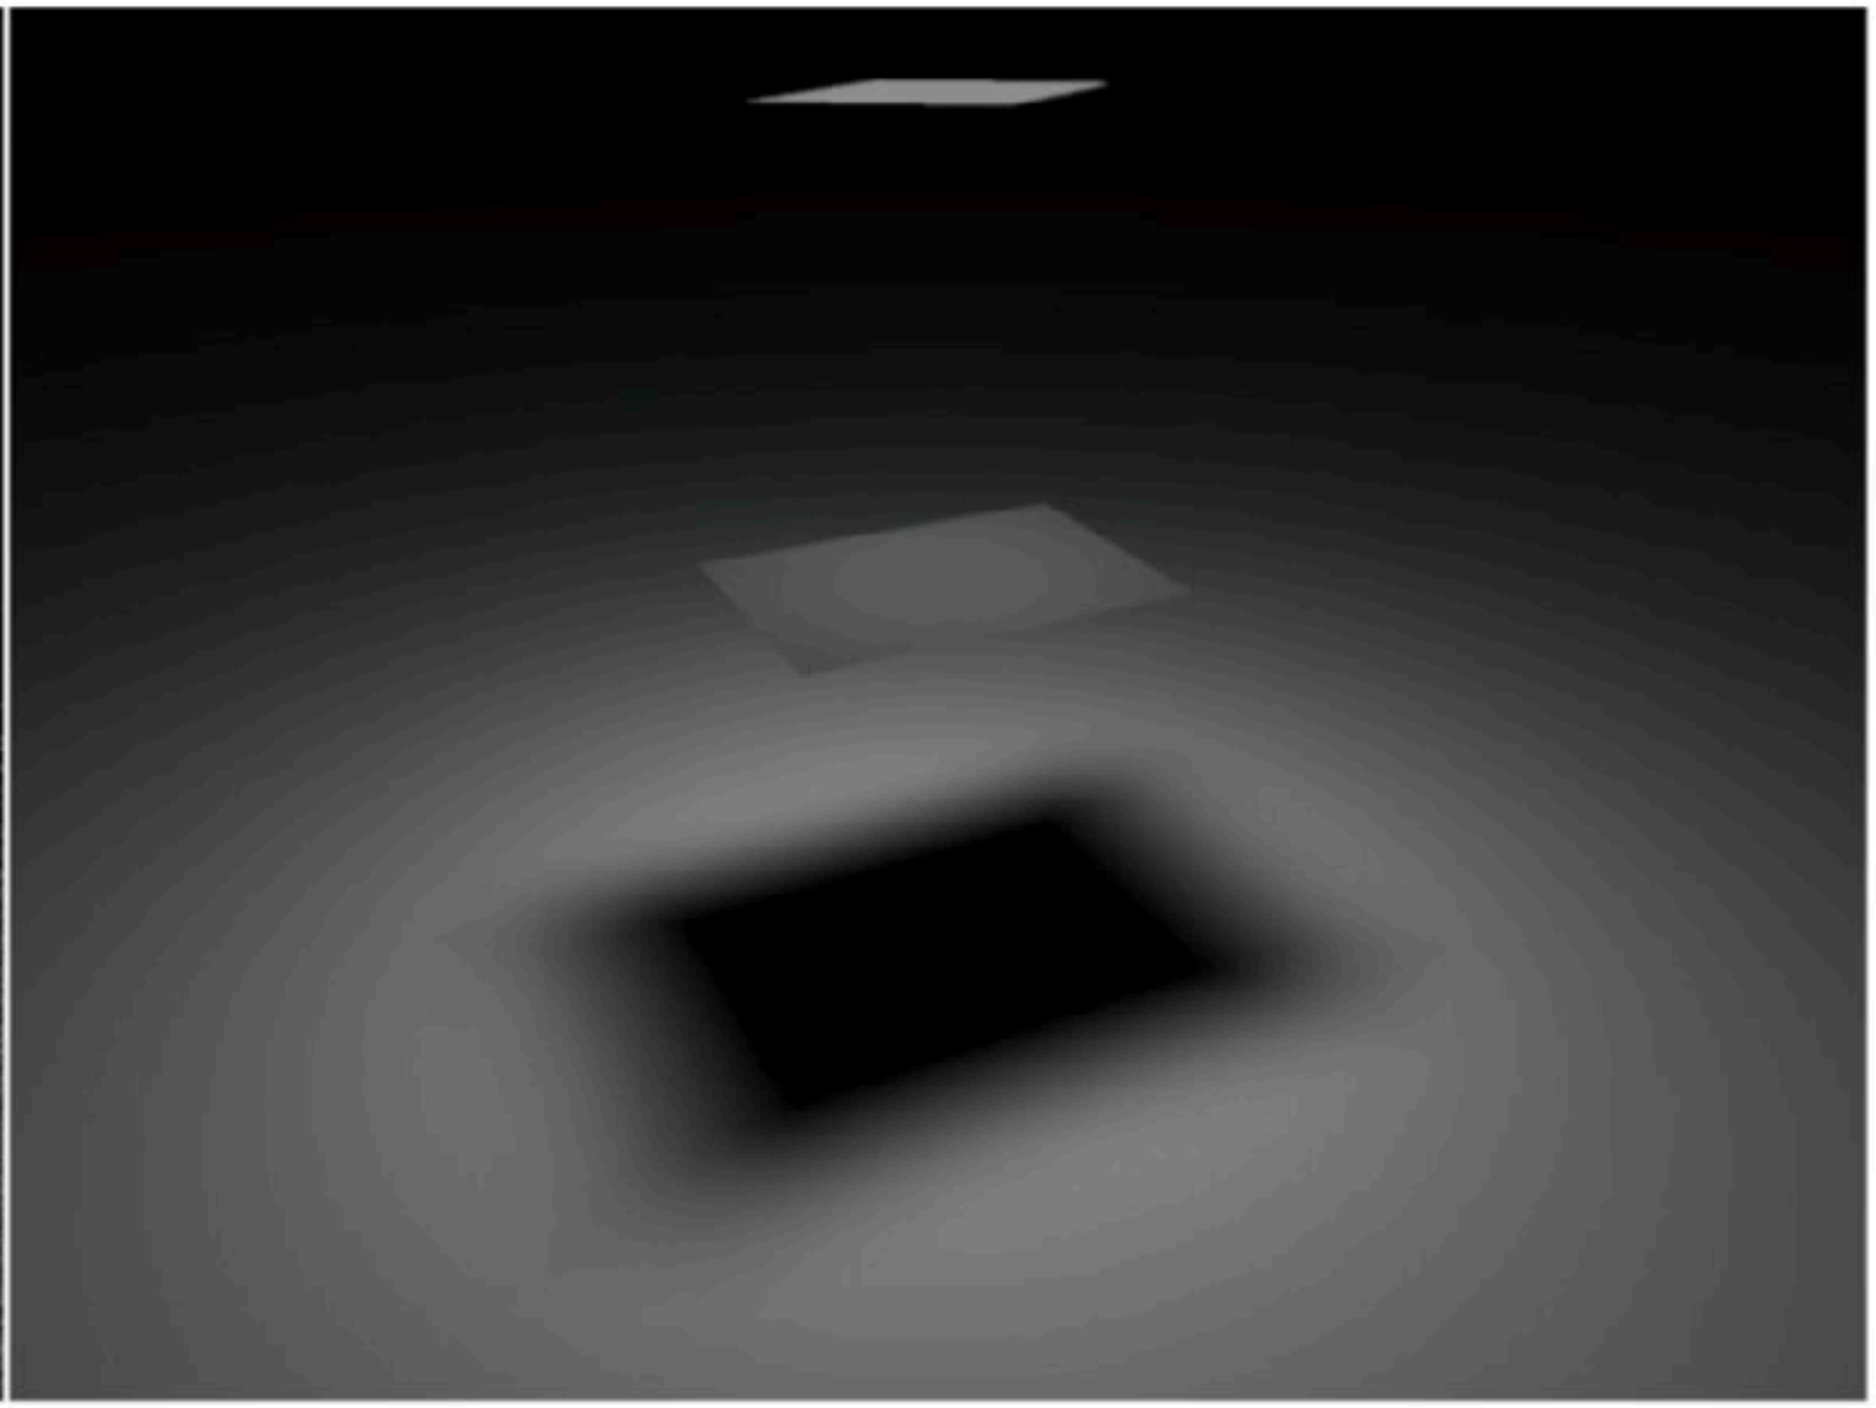

Solid angle sampling

100 random directions on hemisphere

Area sampling

100 random points on area of light source

Q6

1 Point

To make a soft shadow less noisy, I should use importance sampling to

- ○ direct more random rays toward the camera
- ◉ direct more random rays toward the light source
- ○ weight the light source rays more heavily in my Monte Carlo estimate

Q7

1 Point

I independently sample two uniform random values U_1, U_2 from the interval $[0, 1]$. Which of the following will give me a uniform random sample from the 2D unit disk?

- $(U_1/\sqrt{U_1^2 + U_2^2}, U_2/\sqrt{U_1^2 + U_2^2})$
- $(\sqrt{U_1}\cos(2\pi U_2), \sqrt{U_1}\sin(2\pi U_2))$
- $(U_1\cos(2\pi U_2), U_1\sin(2\pi U_2))$
- $(\sqrt{U_1}\cos(2\pi U_1), \sqrt{U_2}\sin(2\pi U_2))$
- $(U_1\cos(2\pi U_1), U_2\sin(2\pi U_2))$

Need to Sample Uniformly in Area

Incorrect
Not Equi-areal

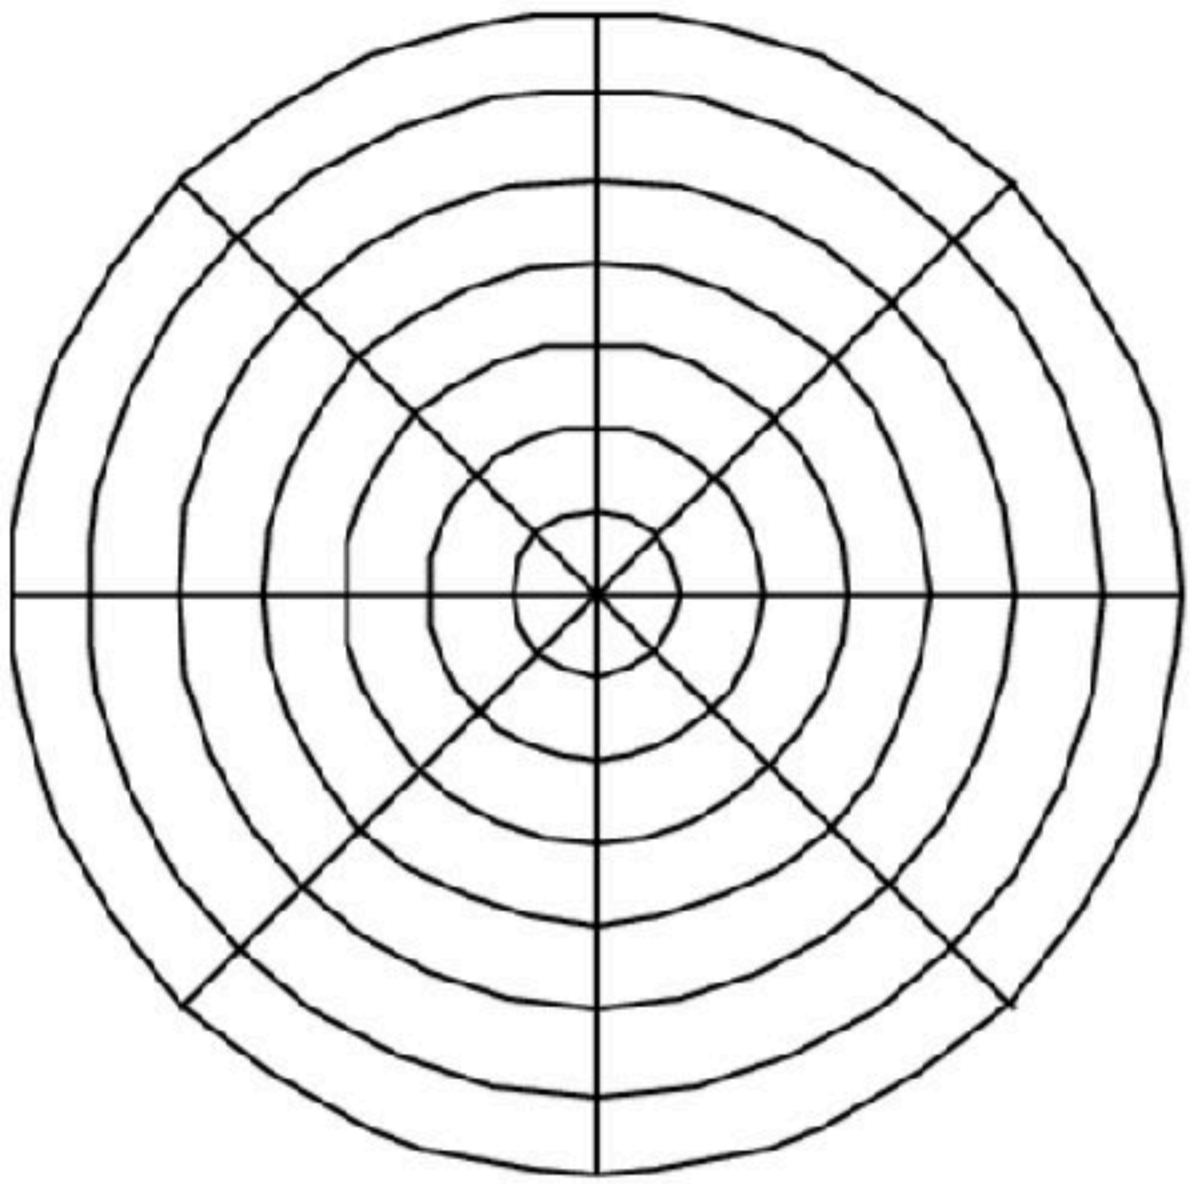

$$\theta = 2\pi\xi_1$$

$$r = \xi_2$$

Correct
Equi-areal

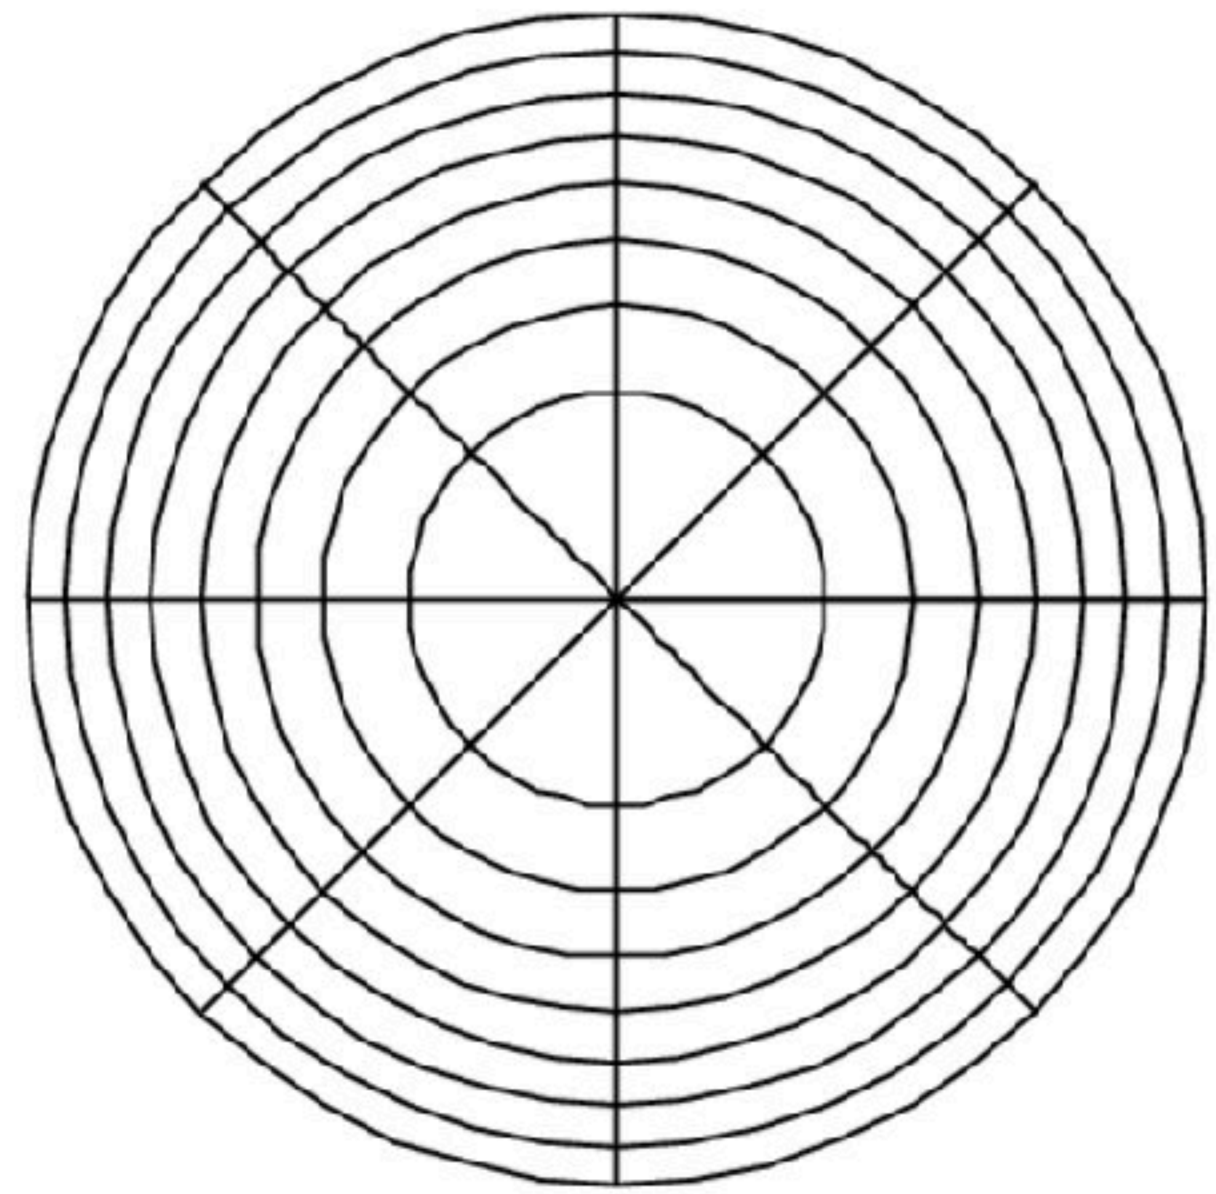

$$\theta = 2\pi\xi_1$$

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

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

Q7

1 Point

I independently sample two uniform random values U_1, U_2 from the interval $[0, 1]$. Which of the following will give me a uniform random sample from the 2D unit disk?

- ○ $(U_1/\sqrt{U_1^2+U_2^2}, U_2/\sqrt{U_1^2+U_2^2})$
- ◉ $(\sqrt{U_1}\cos(2\pi U_2), \sqrt{U_1}\sin(2\pi U_2))$
- ○ $(U_1\cos(2\pi U_2), U_1\sin(2\pi U_2))$
- ○ $(\sqrt{U_1}\cos(2\pi U_1), \sqrt{U_2}\sin(2\pi U_2))$
- ○ $(U_1\cos(2\pi U_1), U_2\sin(2\pi U_2))$

Q4

1 Point

Select all true answers.

- [ ] The light transport operator is linear
- [ ] The light transport operator is a composition of the reflection and refraction operators
- [ ] An operator maps a function to a function

Operators Are Higher-Order Functions

Functions:

$$f, g : (x, \omega) \to \mathbb{R}$$

Operators are higher-order functions:

$$P : ((x, \omega) \to \mathbb{R}) \to ((x, \omega) \to \mathbb{R})$$

$$P(f) = g$$

- Take a function and transform it into another function

Linear Operators

- Linear operators act on functions like matrices act on vectors

$$h(x) = (L(f))(x)$$

- They are linear in that:

$$L(af + bg) = aL(f) + bL(g)$$

- Examples of linear operators:

$$H(f)(x) = \int h(x, x')\, f(x')\, \mathrm{d}x'$$

$$D(f)(x) = \frac{\delta f}{\delta x}(x)$$

Rendering Equation in Operator Notation

$$L_o(\mathrm{p}, \omega_o) = L_e(\mathrm{p}, \omega_o) + \int_{H^2} f_r(\mathrm{p}, \omega_i \to \omega_o)\, L_o(tr(\mathrm{p}, \omega_i), -\omega_i)\, \cos\theta_i\, \mathrm{d}\omega_i$$

$$L_o = L_e + (R \circ T)(L_o)$$

Define full one-bounce light transport operator: $K = R \circ T$

$$\boxed{L_o = L_e + K(L_o)}$$

Q4

1 Point

Select all true answers.

- [x] The light transport operator is linear
- [ ] The light transport operator is a composition of the reflection and refraction operators
- [x] An operator maps a function to a function

Q8

1 Point

The BRDF completely describes material reflectance properties, telling us the amount of ___ in direction ω_r due to the amount of ___ from direction ω_i.

- outgoing radiance, incoming radiance
- outgoing irradiance, incoming irradiance
- outgoing radiance, incoming irradiance
- outgoing radiant intensity, incoming radiance

BRDF

Definition: The bidirectional reflectance distribution function (BRDF) represents how much light is reflected into each outgoing direction ω_r from each incoming direction

NB: ω_i points away from surface rather than into surface, by convention.

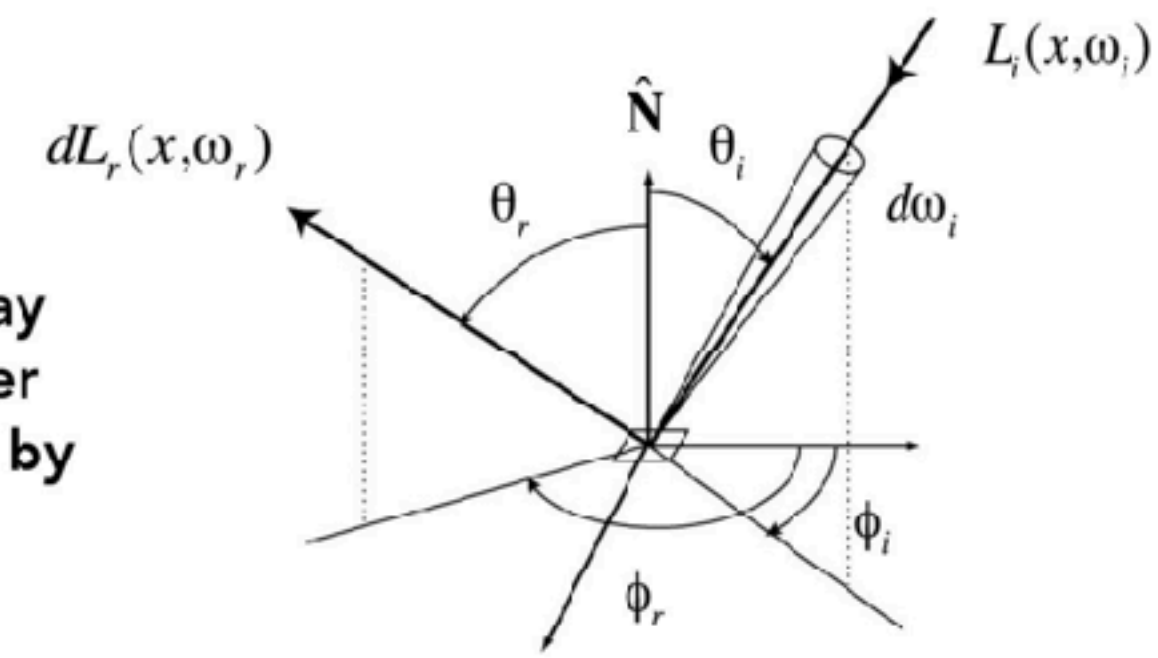

$$f_r(\omega_i \rightarrow \omega_r) = \frac{\mathrm{d}L_r(\omega_r)}{\mathrm{d}E_i(\omega_i)} = \frac{\mathrm{d}L_r(\omega_r)}{L_i(\omega_i)\cos\theta_i\,\mathrm{d}\omega_i} \quad \left[\frac{1}{\mathrm{sr}}\right]$$

Reflection at a Point

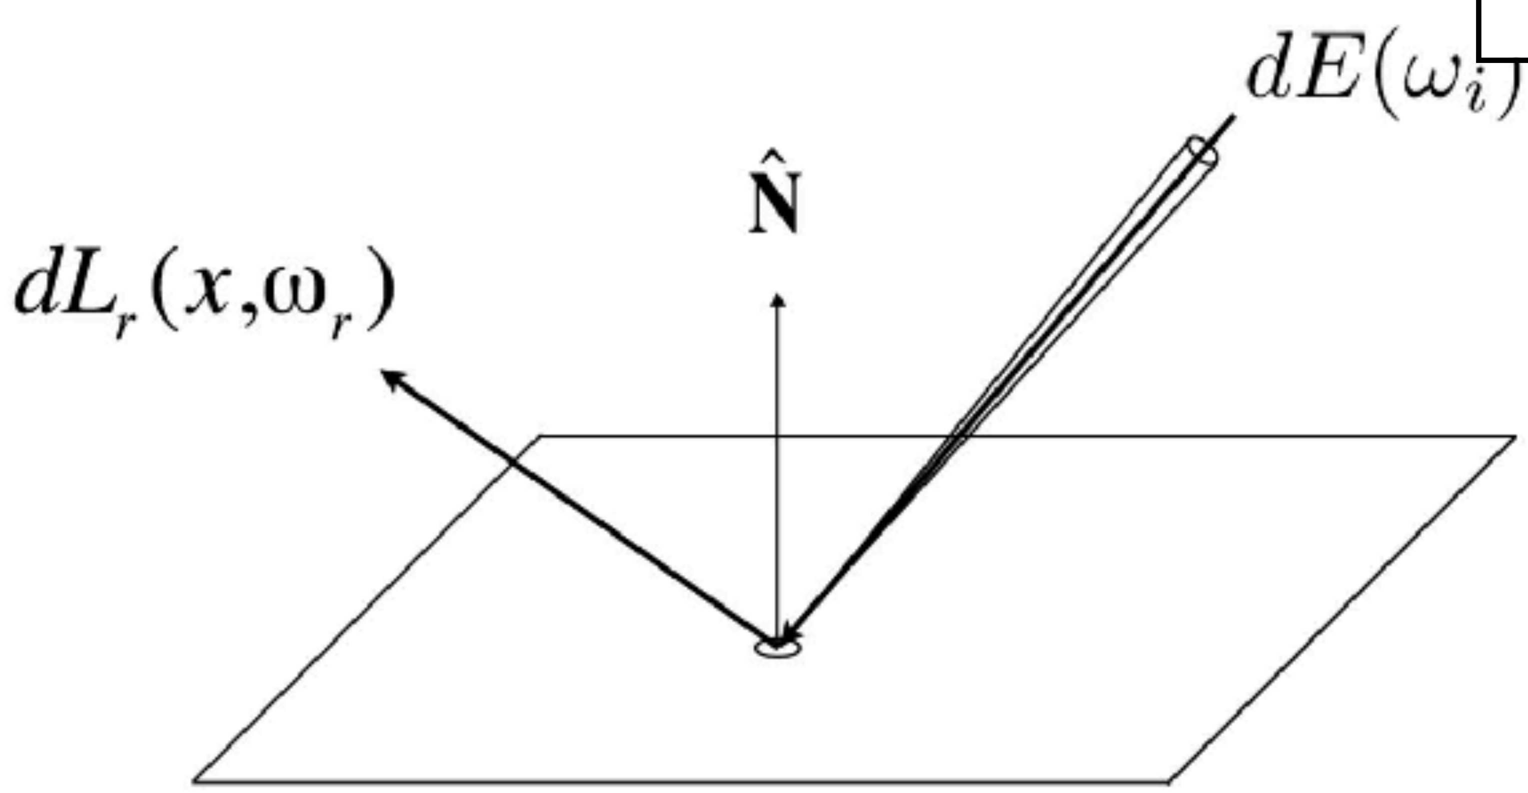

Differential irradiance incoming: $dE(\omega_i) = L(\omega_i)\, cos\,\theta_i\, d\omega_i$

Differential radiance exiting (due to $dE(\omega_i)$) $dL_r(\omega_r)$

Q8

1 Point

The BRDF completely describes material reflectance properties, telling us the amount of ____ in direction ω_r due to the amount of ____ from direction ω_i.

- ○ outgoing radiance, incoming radiance
- ○ outgoing irradiance, incoming irradiance
- ◉ outgoing radiance, incoming irradiance
- ○ outgoing radiant intensity, incoming radiance

Q2

1 Point

In the context of path tracing, Russian Roulette allows us to

- ○ delete occluded objects from the scene in a preprocessing step
- ○ randomly eliminate light sources that do not contribute brightness
- ○ estimate which parts of the scene require tracing more indirect illumination rays
- ○ evaluate an infinite dimensional integral with a finite unbiased estimator

Russian Roulette: Unbiased Random Termination

New estimator: evaluate original estimator with probability $p_{\rm rr}$, reweighted. Otherwise ignore.

$$\text{Let } X_{\rm rr} = \begin{cases} \frac{X}{p_{\rm rr}}, & \text{with probability } p_{\rm rr} \\ 0, & \text{otherwise} \end{cases}$$

Same expected value as original estimator:

$$E[X_{\rm rr}] = p_{\rm rr}\, E\left[\frac{X}{p_{\rm rr}}\right] + (1 - p_{\rm rr})\, E[0] = E[X]$$

Want to choose $p_{\rm rr}$ considering Monte Carlo efficiency

- Terminate if expensive and/or low contribution
- In path tracing, expensive to recursively trace path. Increase termination probability if brdf is low in next bounce direction

Q2

1 Point

In the context of path tracing, Russian Roulette allows us to

- ○ delete occluded objects from the scene in a preprocessing step
- ○ randomly eliminate light sources that do not contribute brightness
- ○ estimate which parts of the scene require tracing more indirect illumination rays
- ◉ evaluate an infinite dimensional integral with a finite unbiased estimator

Q3

1 Point

Select all true statements.

- [ ] Exactly solving the rendering equation would require tracing ray paths with arbitrarily many bounces
- [ ] For any valid physically-based scene model, the total energy of the radiance function decreases after applying the light transport operator K
- [ ] For $M > N$, every pixel in an image rendered using only M-bounce paths will be less bright than the equivalent pixel in an N-bounce rendering

Rendering Equation Solution

$$
\begin{aligned}
L &= (I - K)^{-1}(L_e) \\
&= (I + K + K^2 + K^3 + \cdots)(L_e) \\
&= L_e + K(L_e) + K^2(L_e) + K^3(L_e) + \cdots
\end{aligned}
$$

Emitted (L_e), **1-bounce** ($K(L_e)$), **2-bounce** ($K^2(L_e)$), **3-bounce** ($K^3(L_e)$)

Intuitive: Sum of successive bounces of light

This calculates the steady-state surface light field over the scene.

Solution Intuition

For scalar functions, recall:

$$\frac{1}{1-x} = 1 + x + x^2 + x^3 + \cdots$$

converges for $-1 < x < 1$

Similarly, for operators, it is true that

$$(I-K)^{-1} = \frac{1}{I-K} = I + K + K^2 + K^3 + \cdots$$

(Neumann series)

converges for $||K|| < 1$

where $||K|| < 1$ means that the "energy" of the radiance function decreases after applying K. This is intuitively true for valid scene models based on energy dissipation (though not trivial to prove, see Veach & Guibas).

Q3

1 Point

Select all true statements.

- [x] Exactly solving the rendering equation would require tracing ray paths with arbitrarily many bounces
- [x] For any valid physically-based scene model, the total energy of the radiance function decreases after applying the light transport operator K
- [ ] For $M > N$, every pixel in an image rendered using only M-bounce paths will be less bright than the equivalent pixel in an N-bounce rendering

Q5

1 Point

Which of the indicated points is **not** primarily lit by **indirect** illumination?

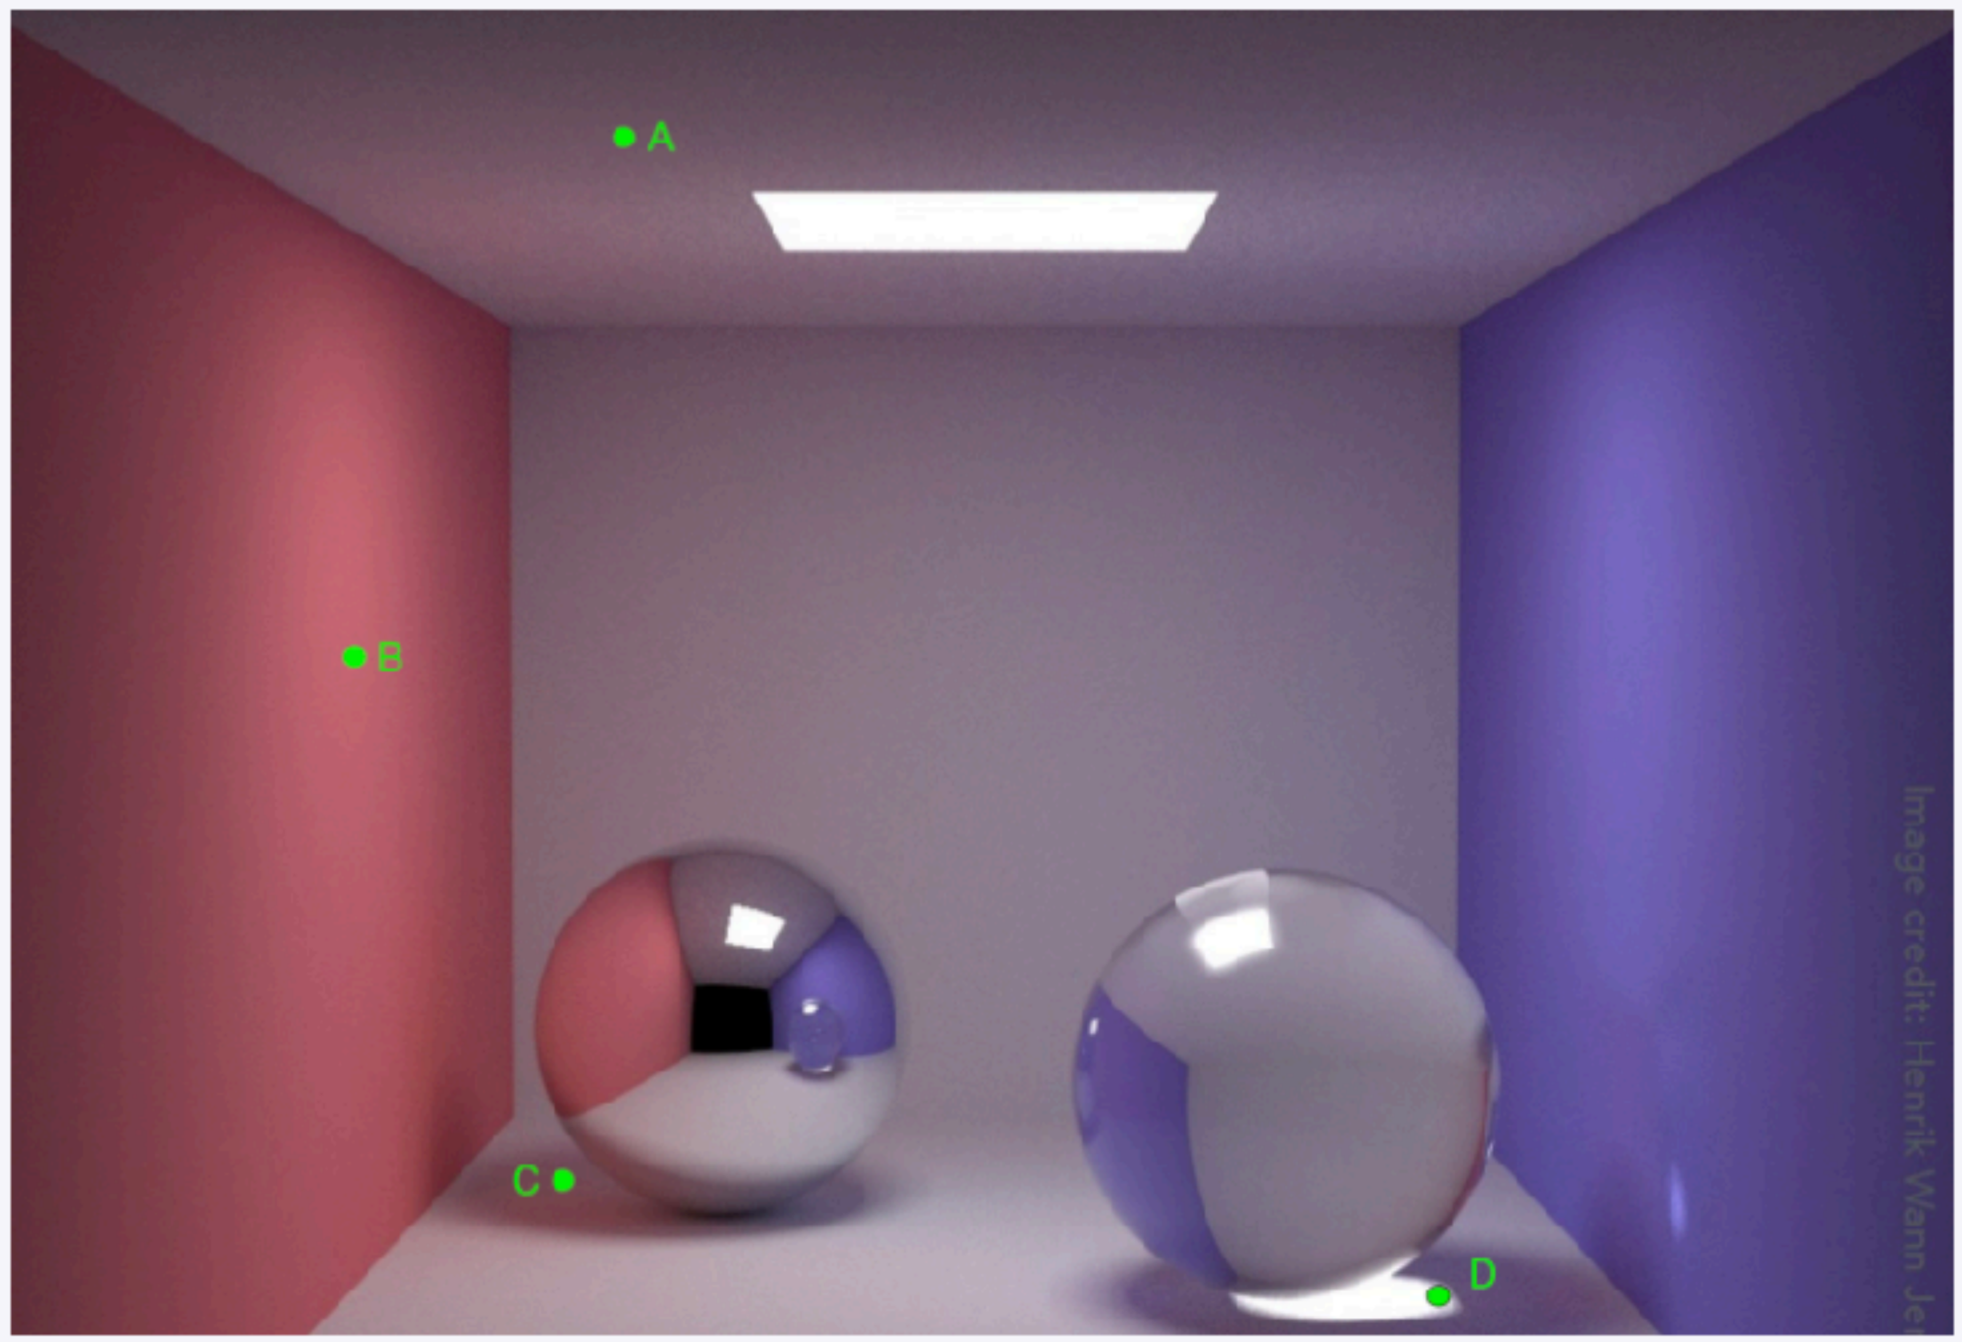

- A
- B
- C
- D

Q5

1 Point

Which of the indicated points is **not** primarily lit by **indirect** illumination?

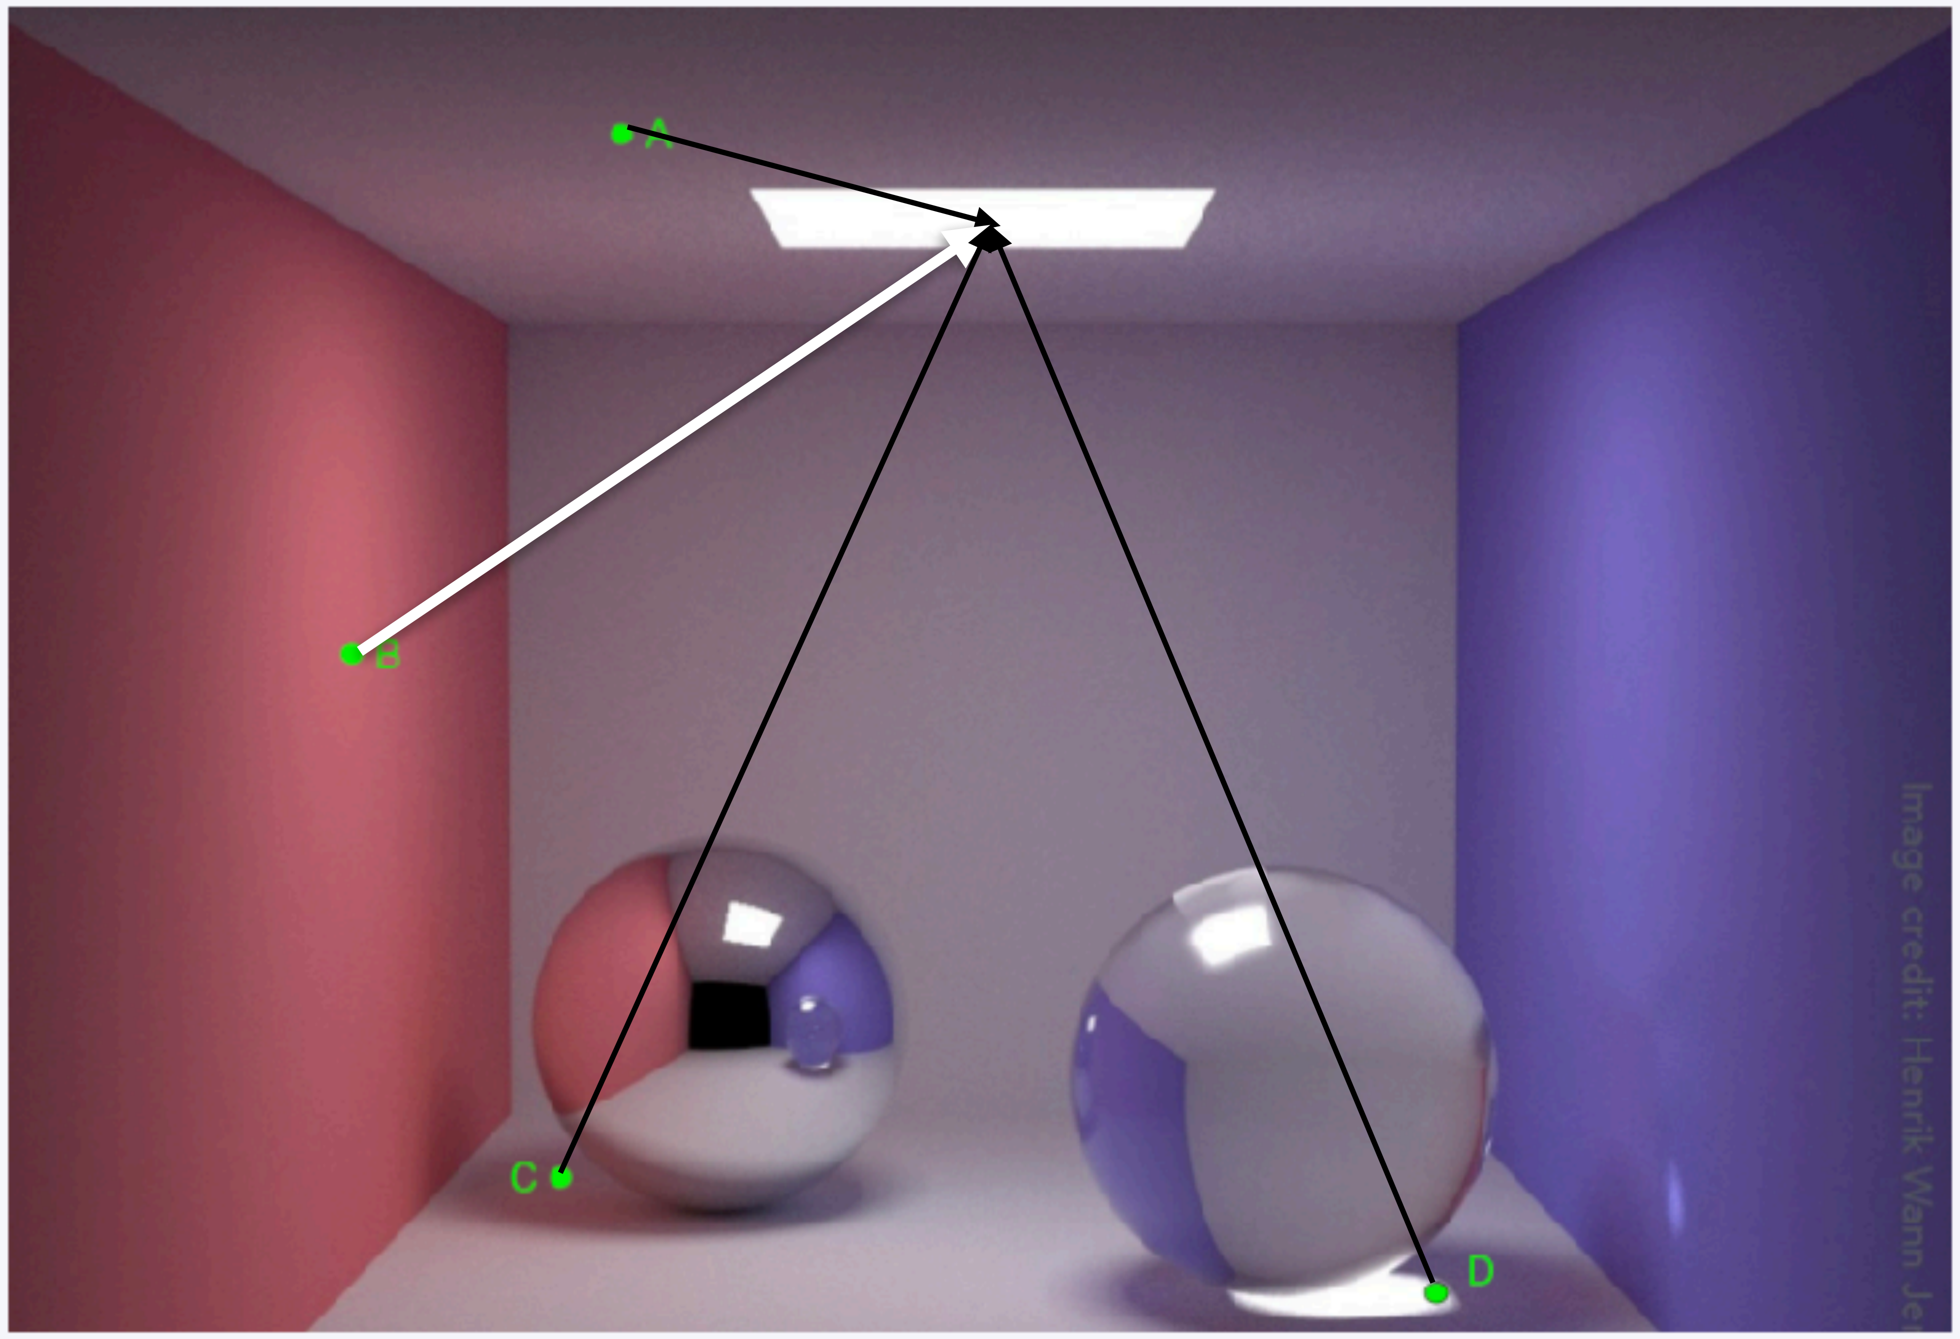

- A
- B
- C
- D

Q5

1 Point

Which of the indicated points is **not** primarily lit by **indirect** illumination?

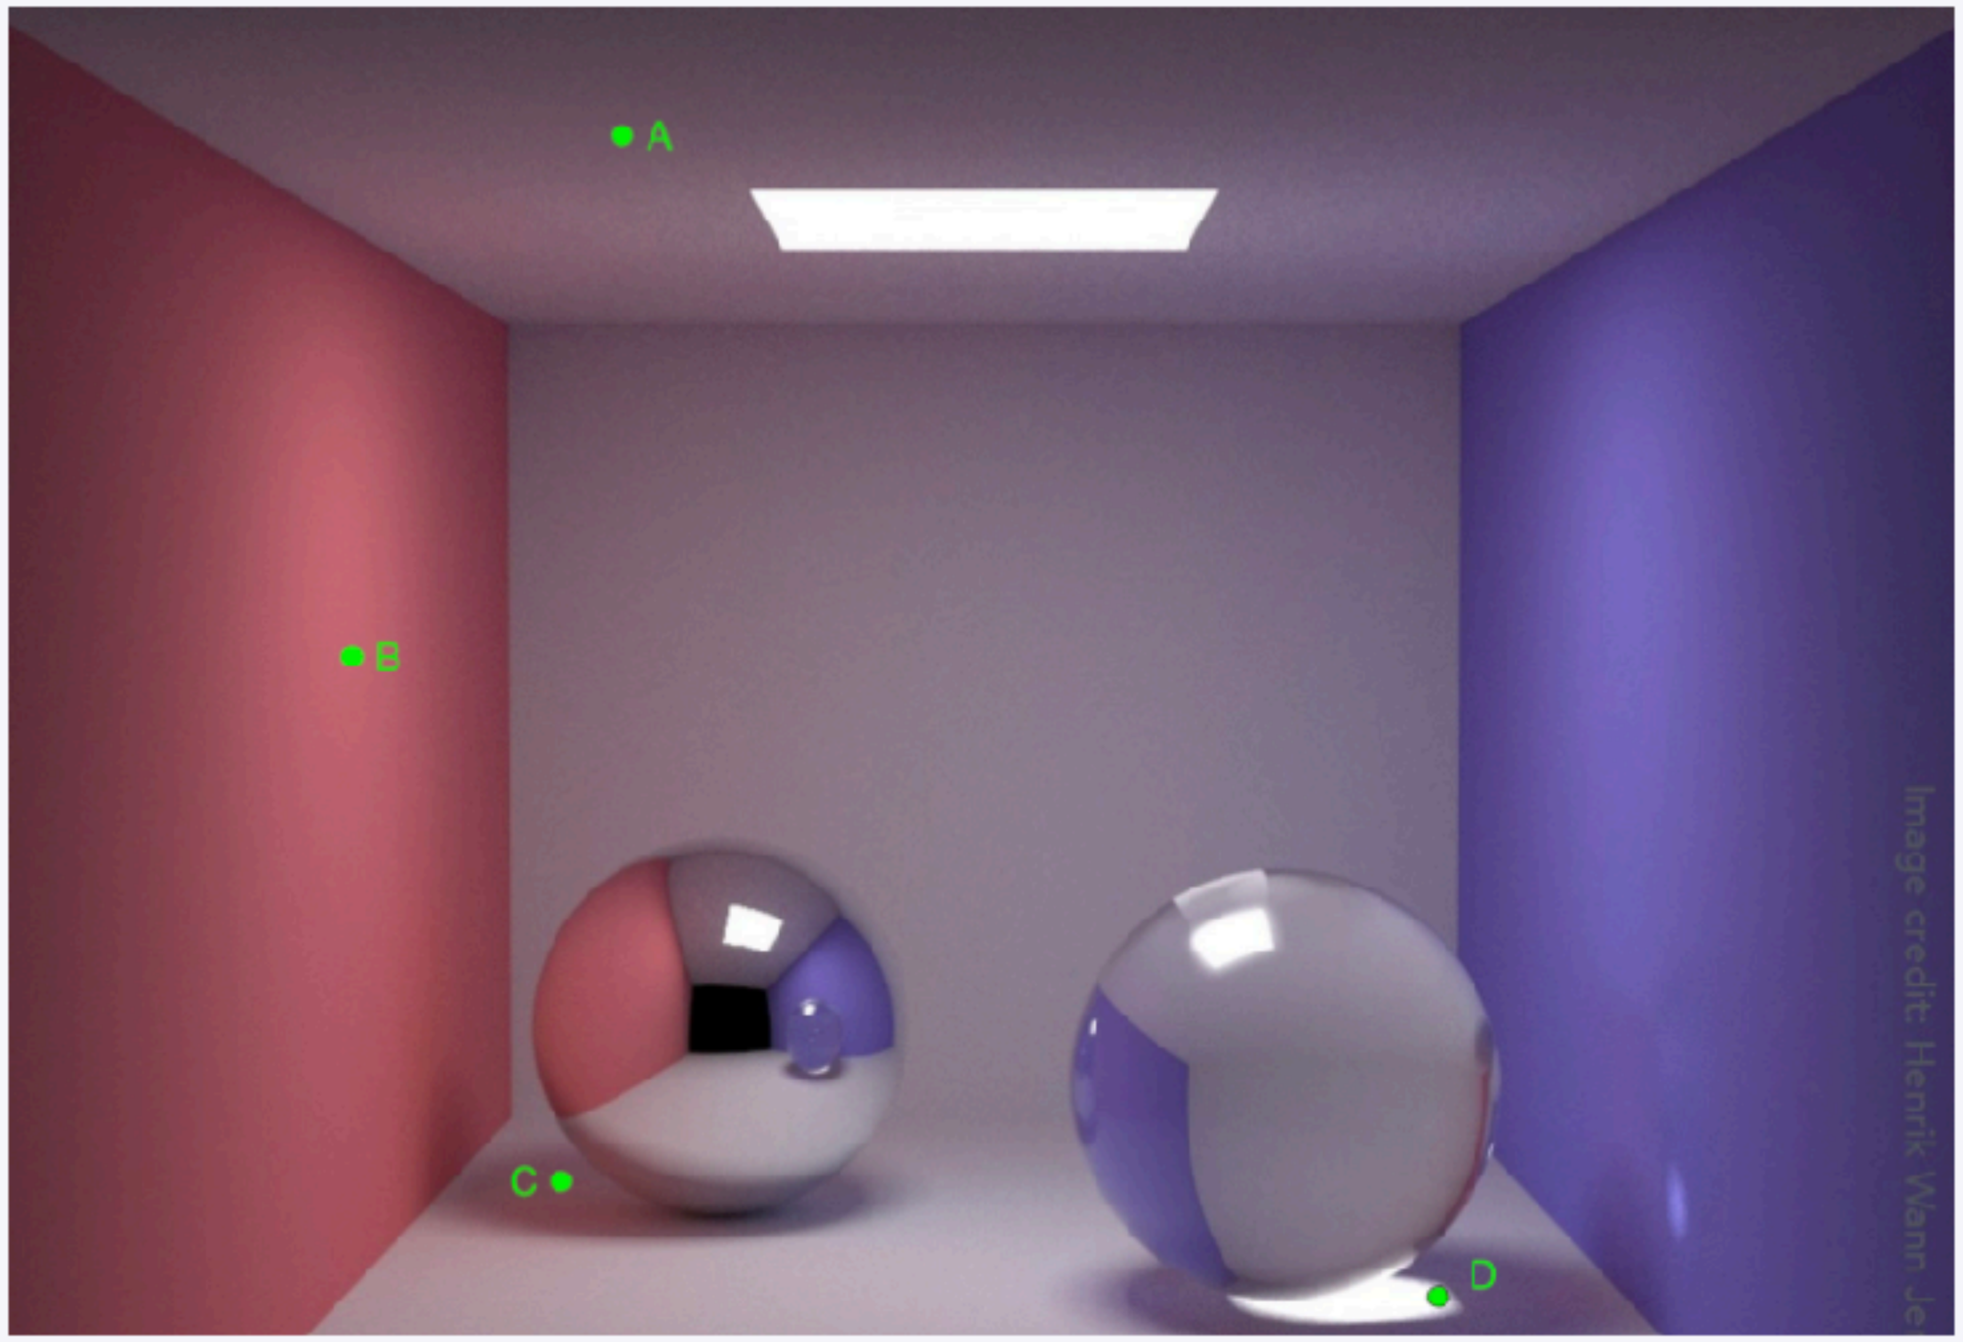

- ○ A
- ◉ B
- ○ C
- ○ D