Irradiance Caching Methods

Wojciech Jarosz

CSE 168

May 27, 2008

Outline

Outline

- Introduction
  - What is indirect illumination?

Outline

- Introduction
  - What is indirect illumination?
- What is Irradiance Caching?

Outline

- Introduction
  - What is indirect illumination?
- What is Irradiance Caching?
- Extensions to Irradiance Caching

Outline

- Introduction
  - What is indirect illumination?
- What is Irradiance Caching?
- Extensions to Irradiance Caching
- Future Work

Direct Illumination

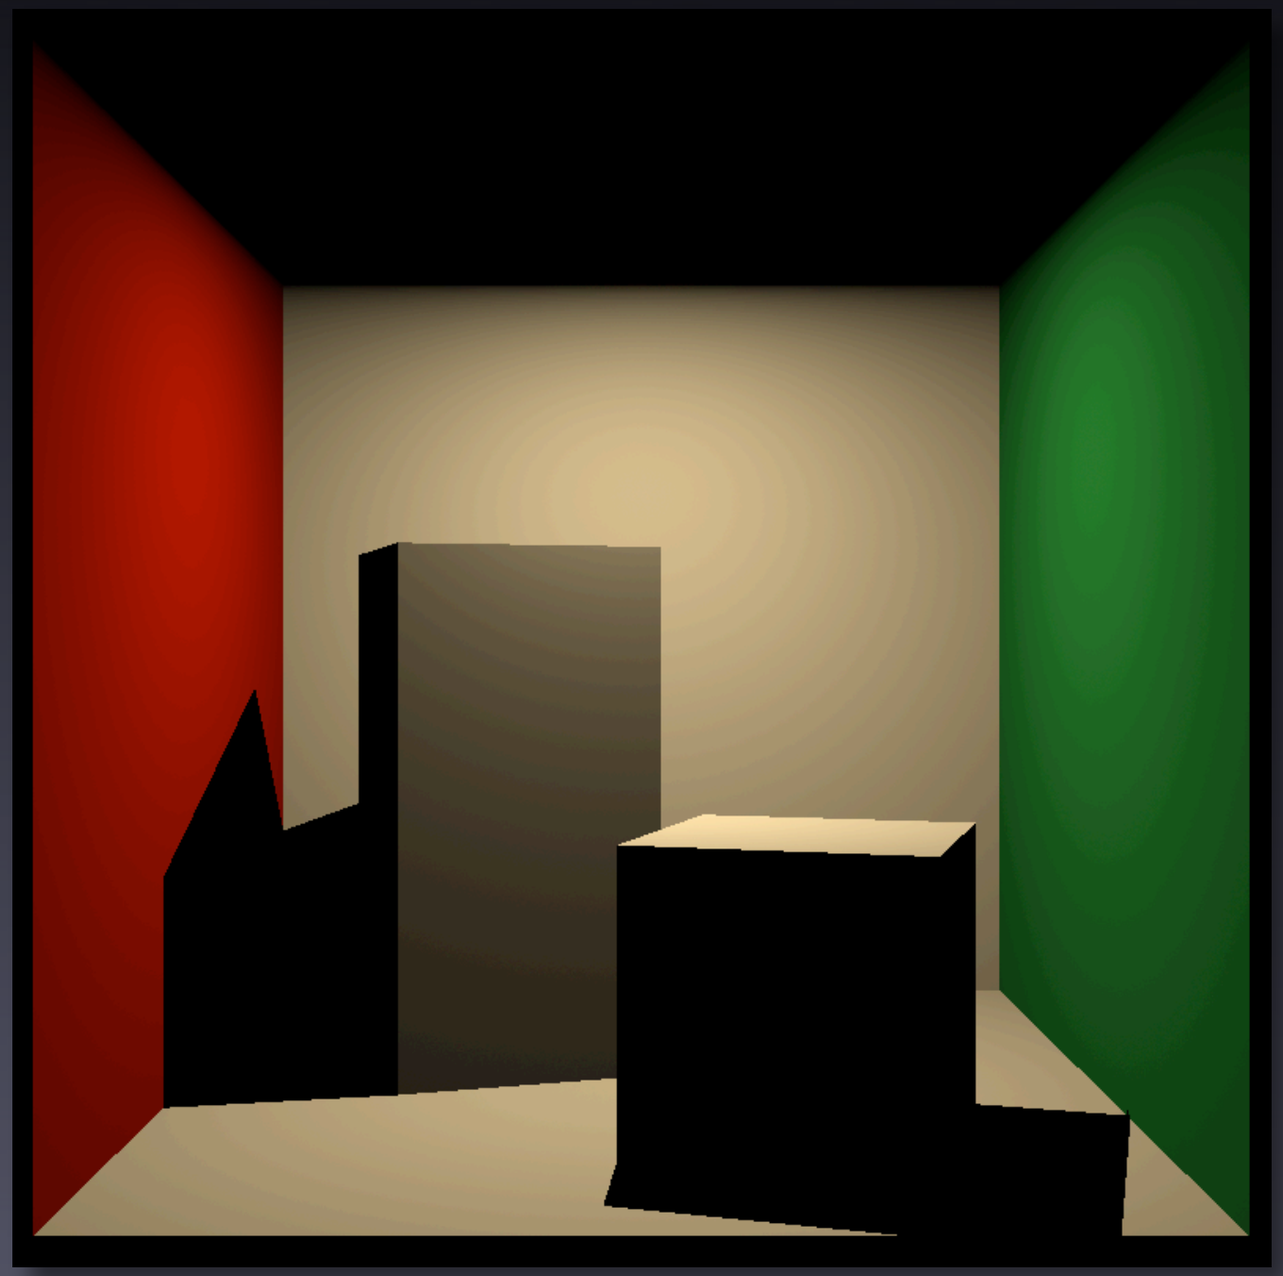

Comparison

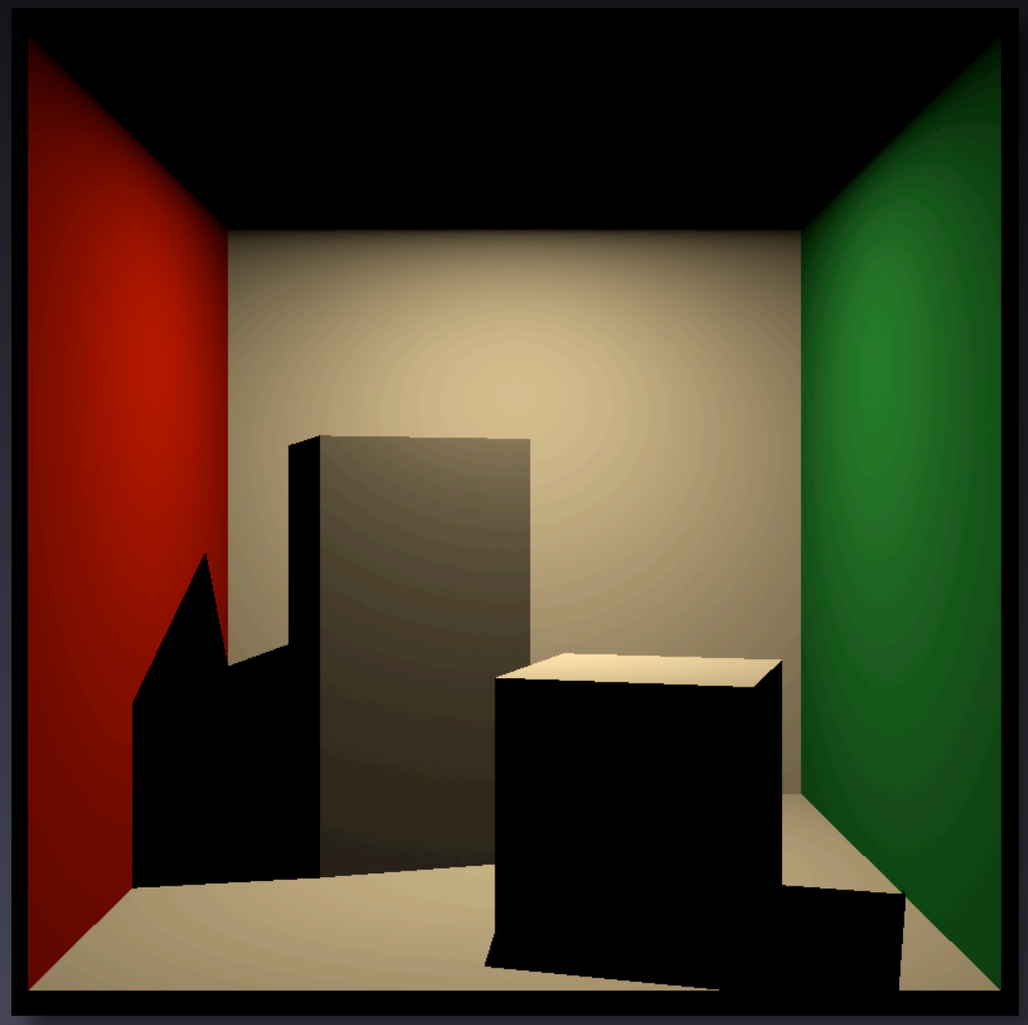

Simulated

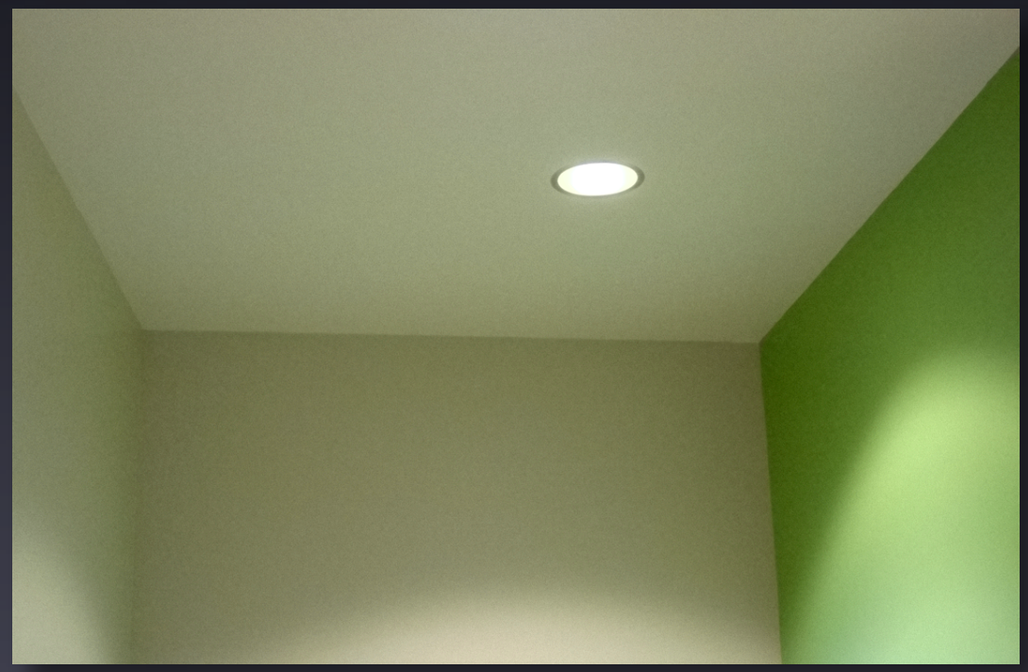

Real

Computing Indirect Illum.

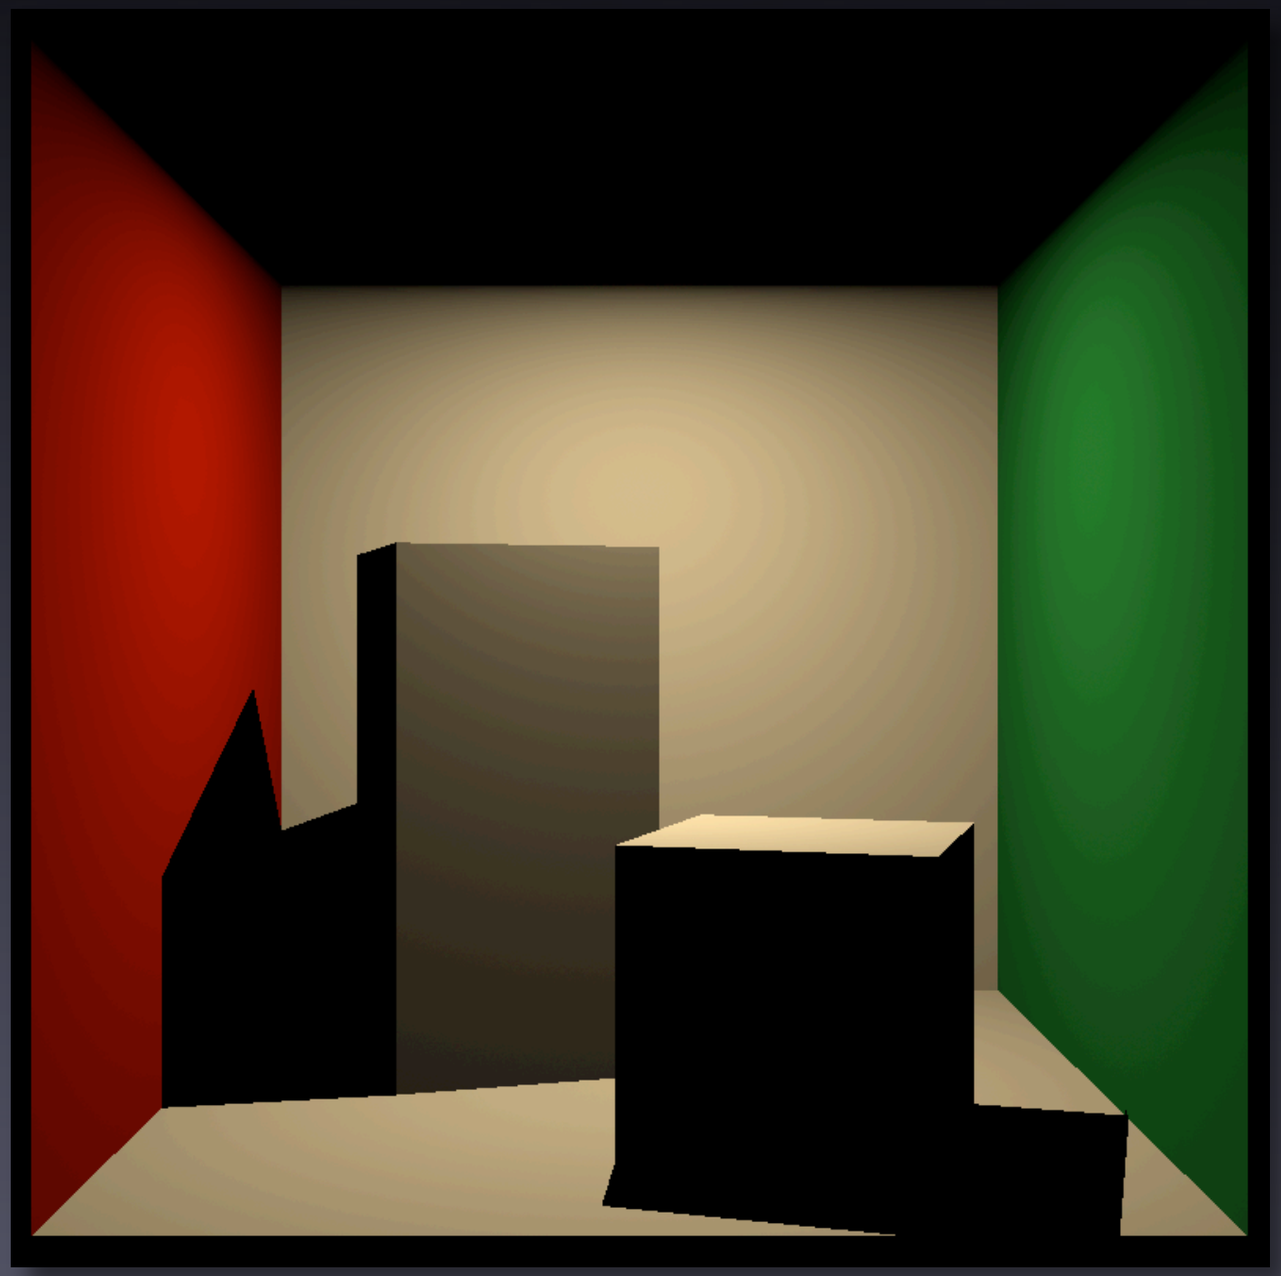

Computing Indirect Illum.

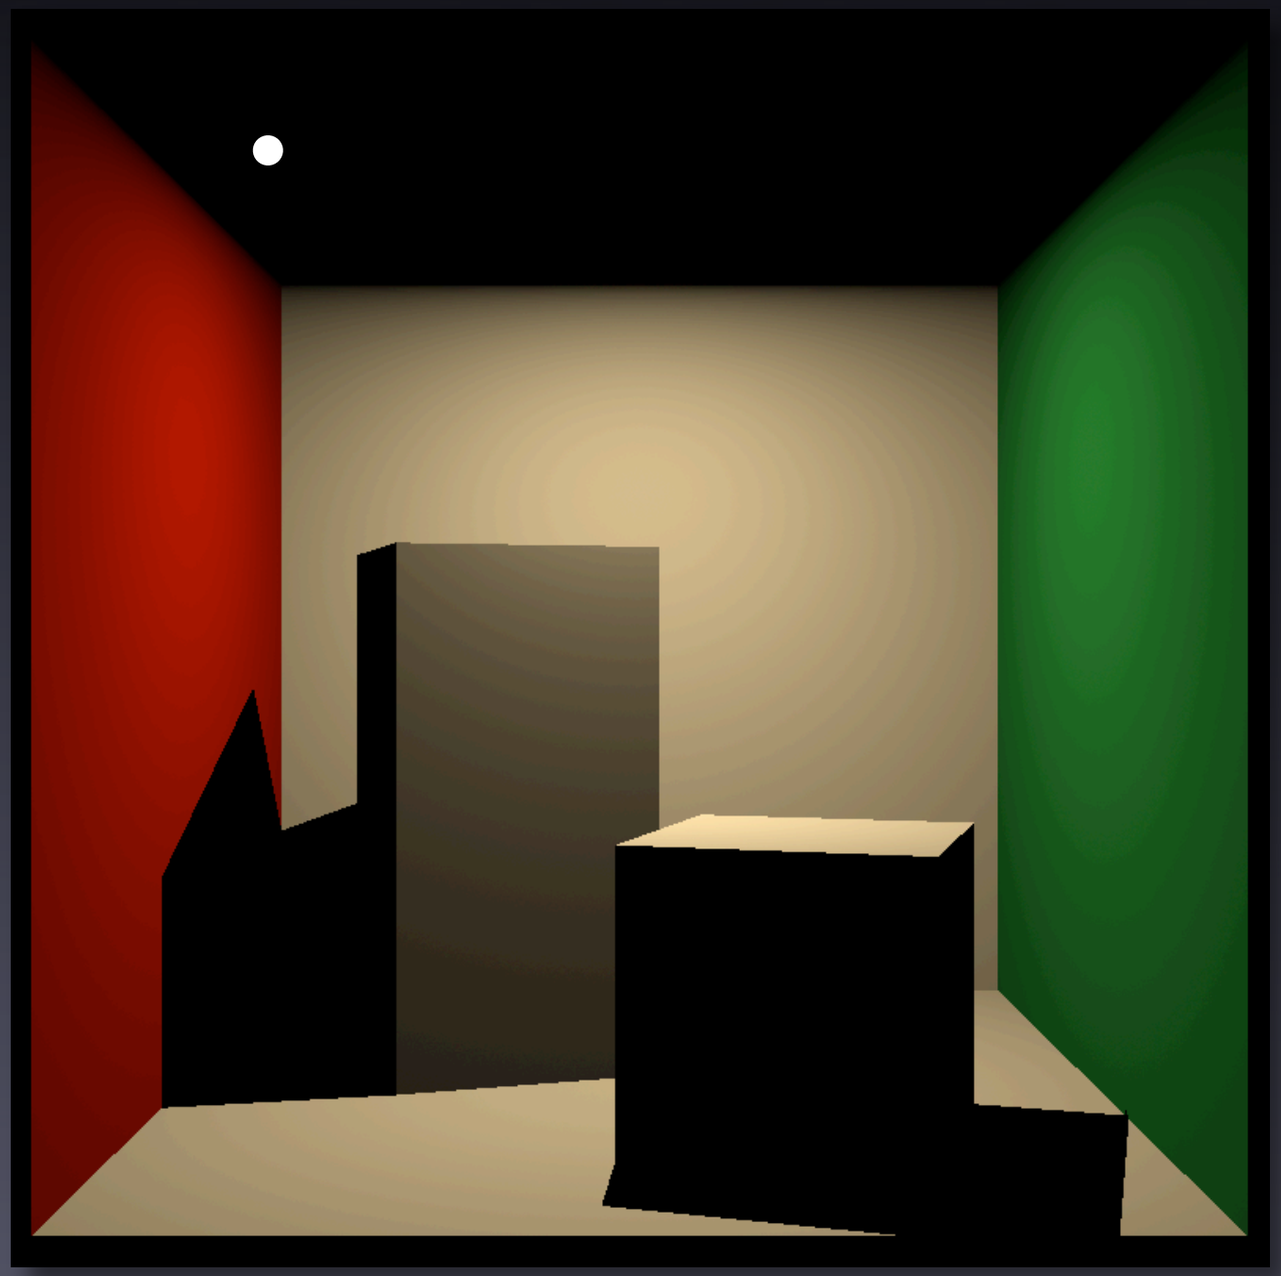

Computing Indirect Illum.

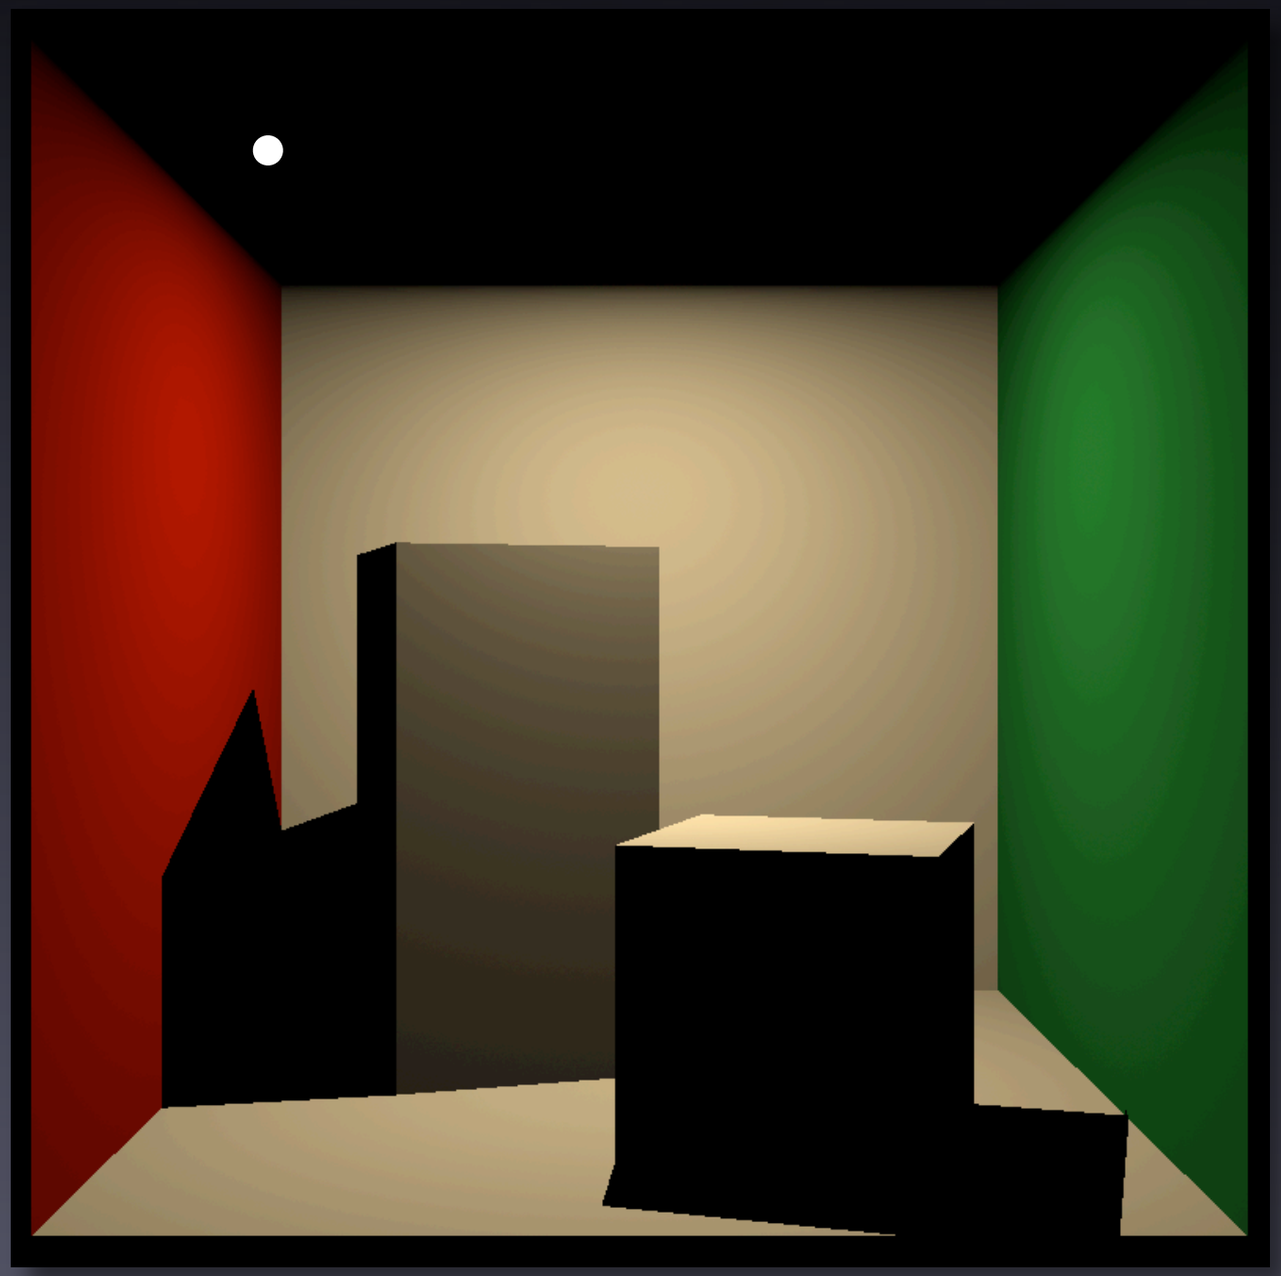

Computing Indirect Illum.

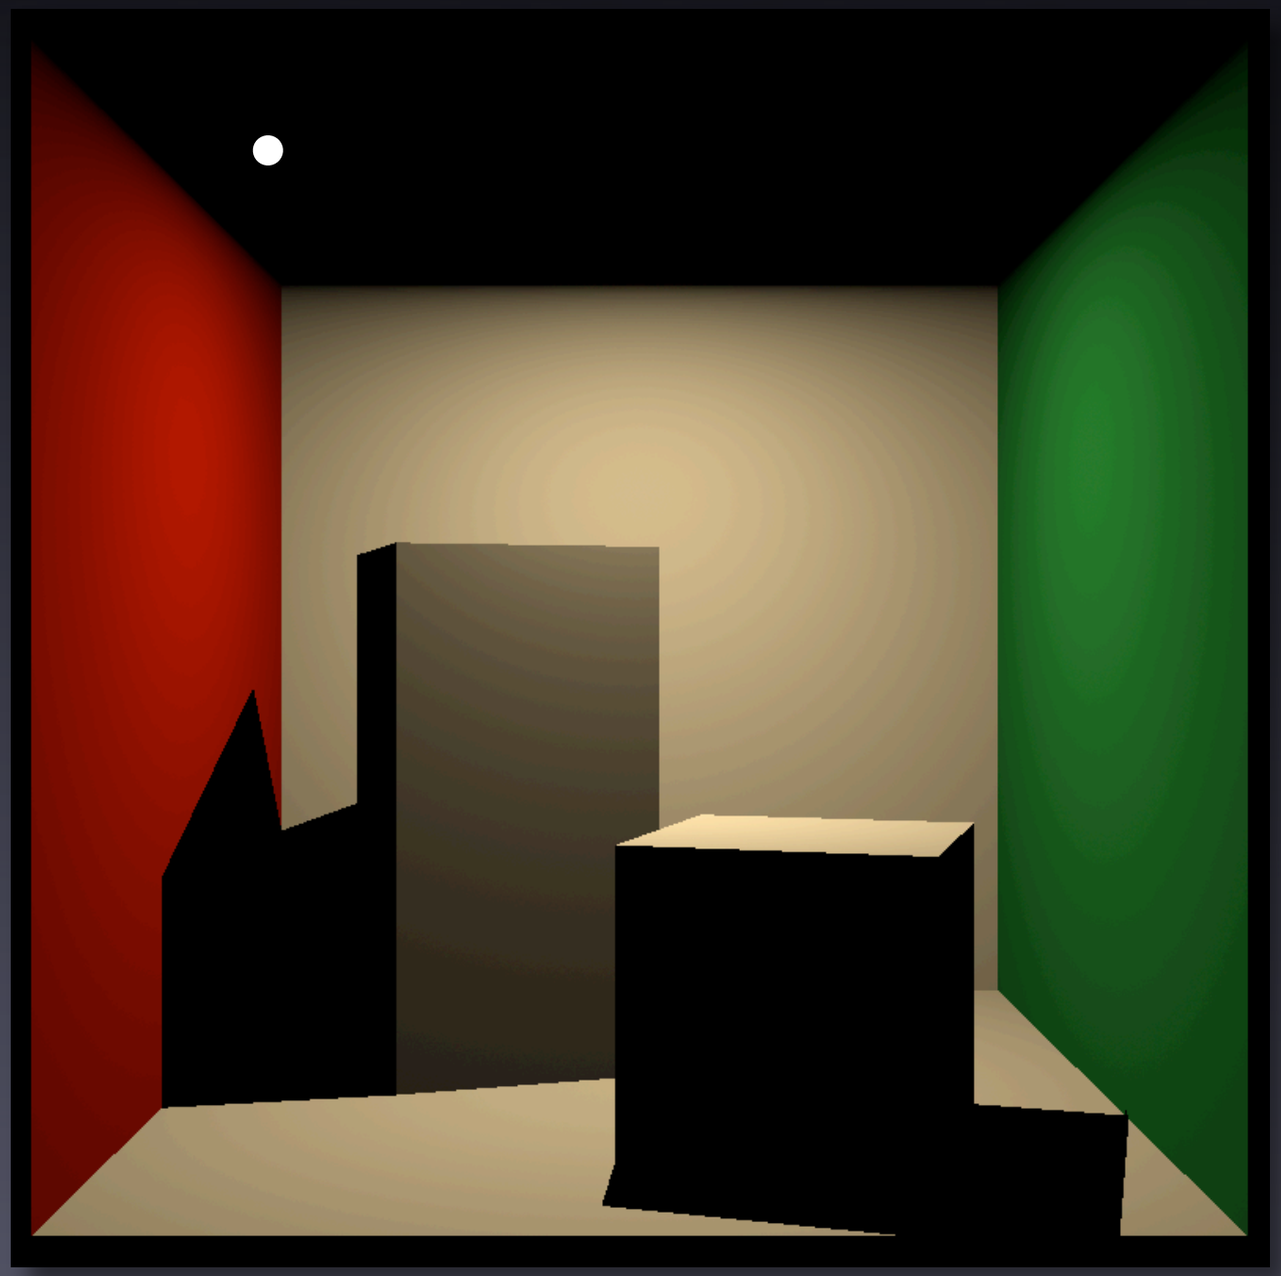

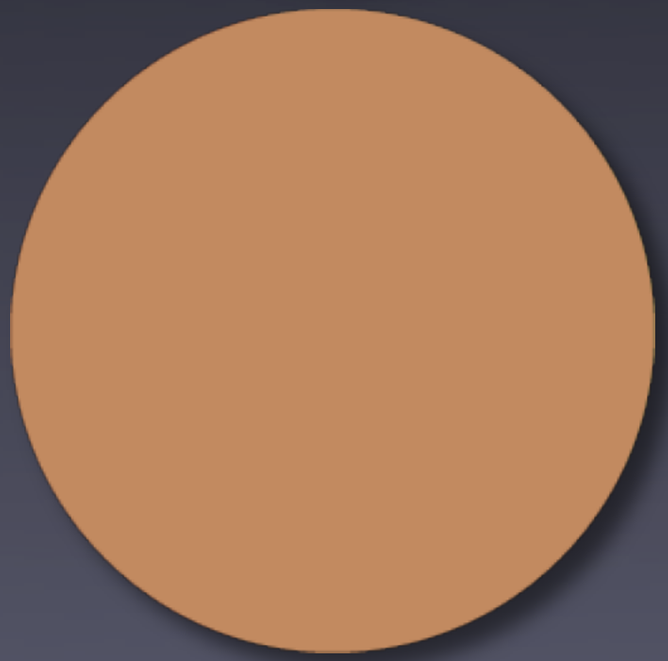

Computing Indirect Illum.

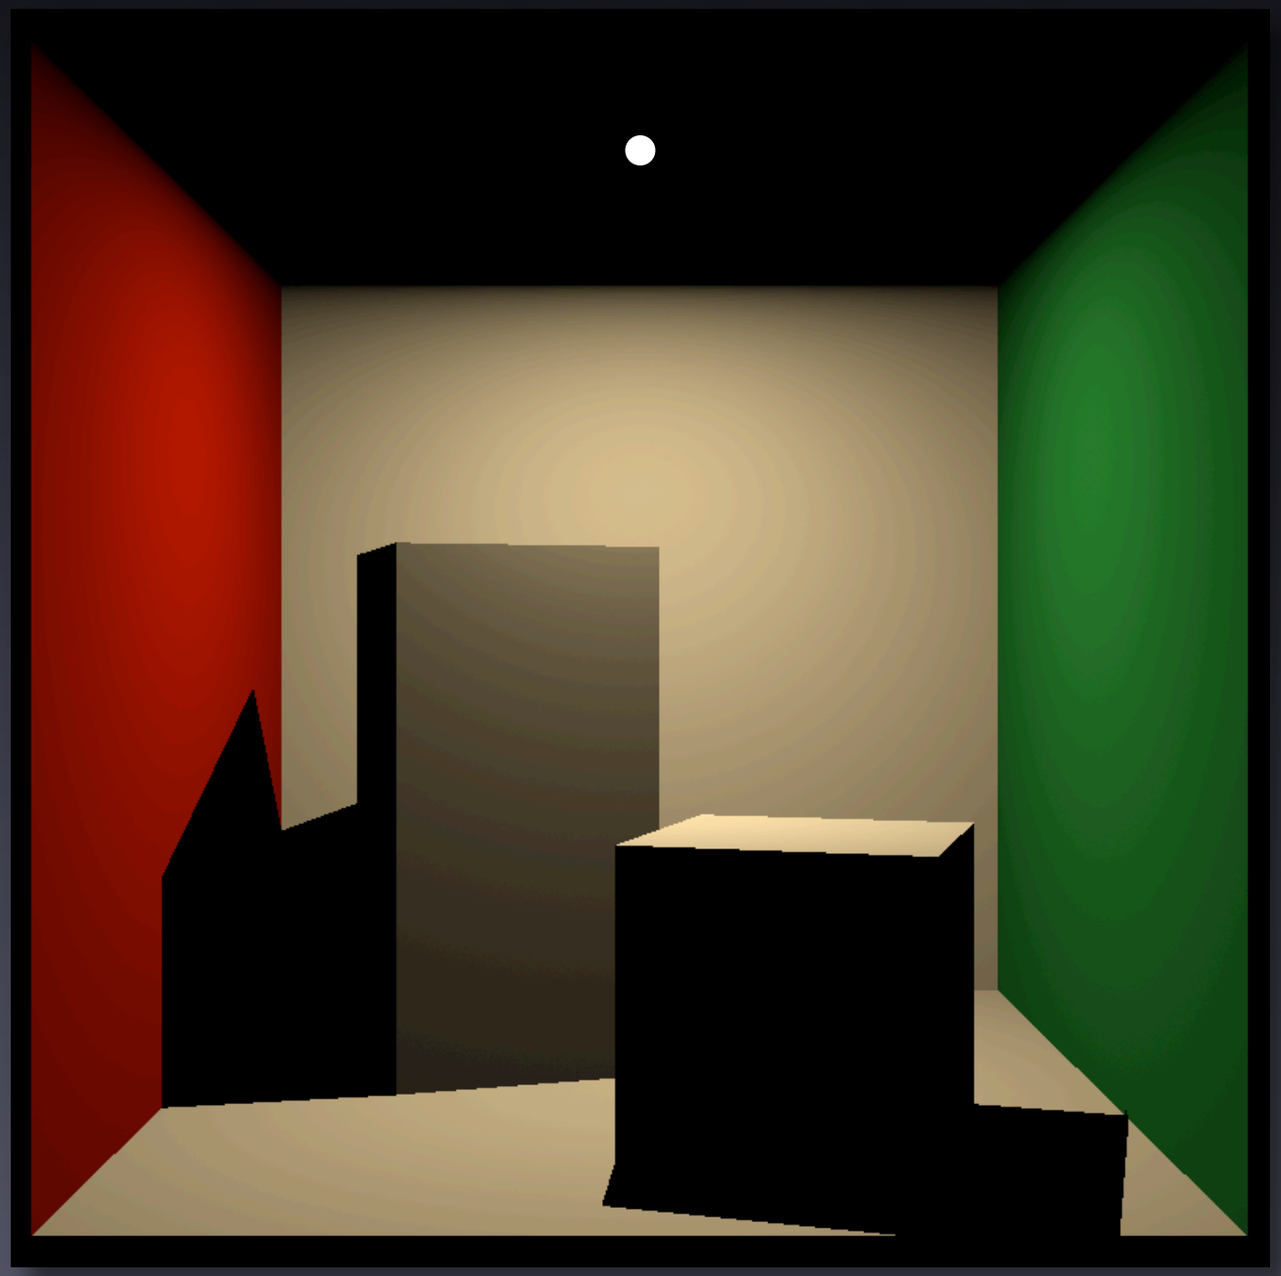

Computing Indirect Illum.

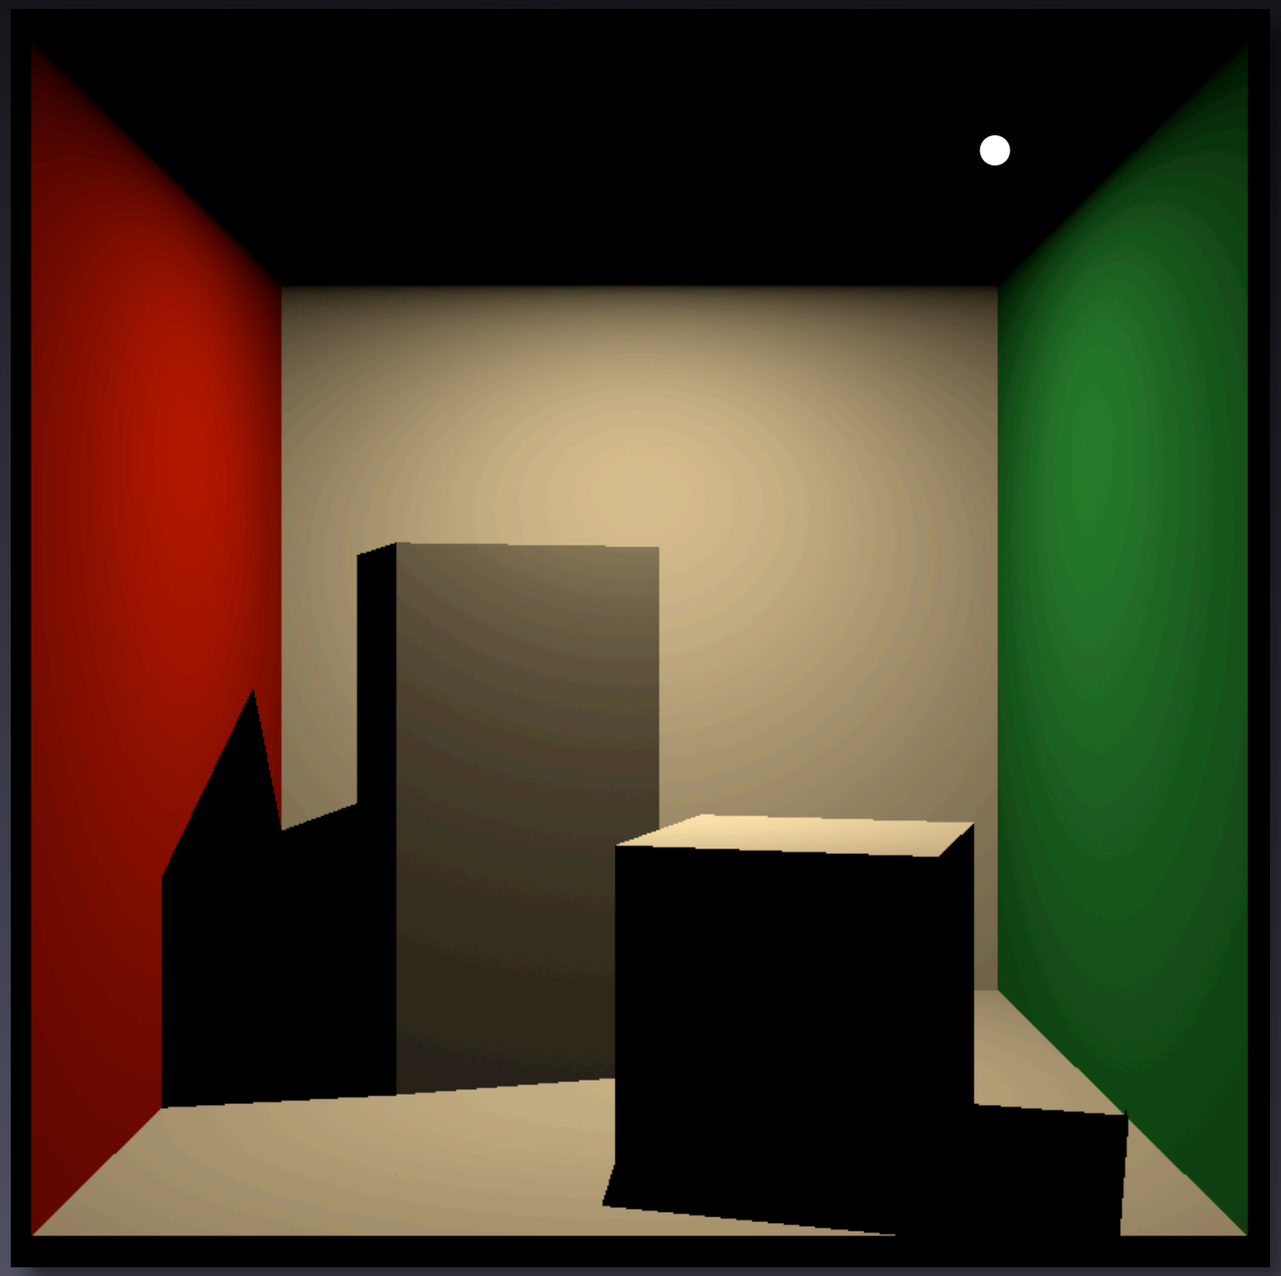

Computing Indirect Illum.

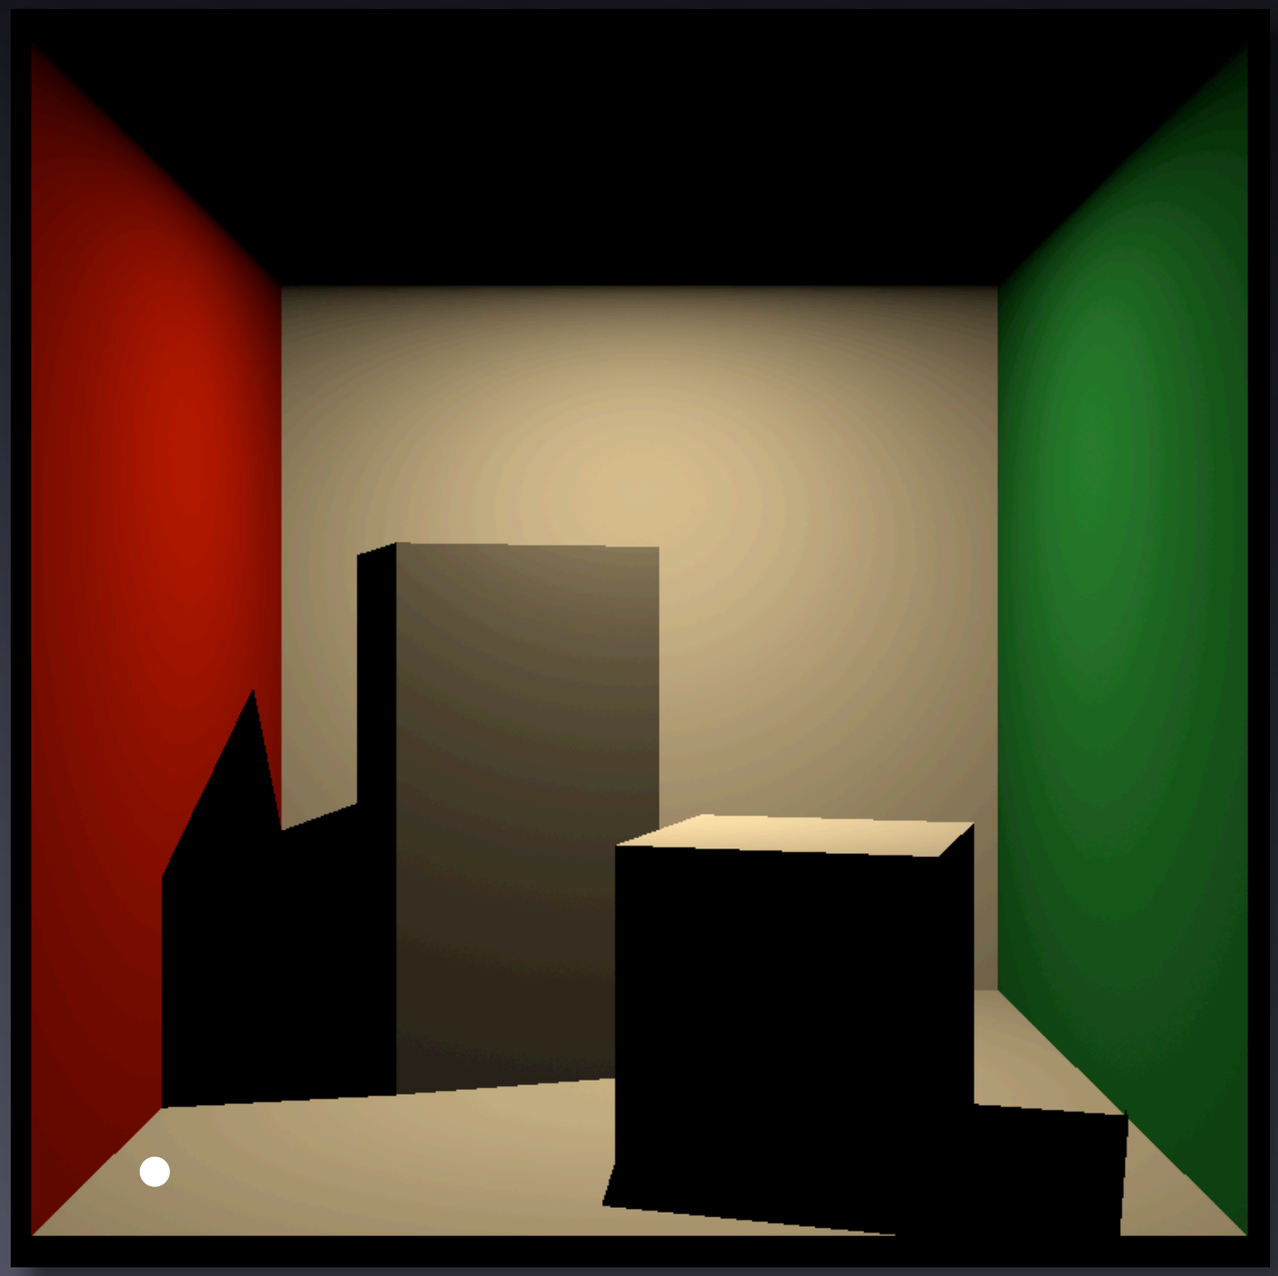

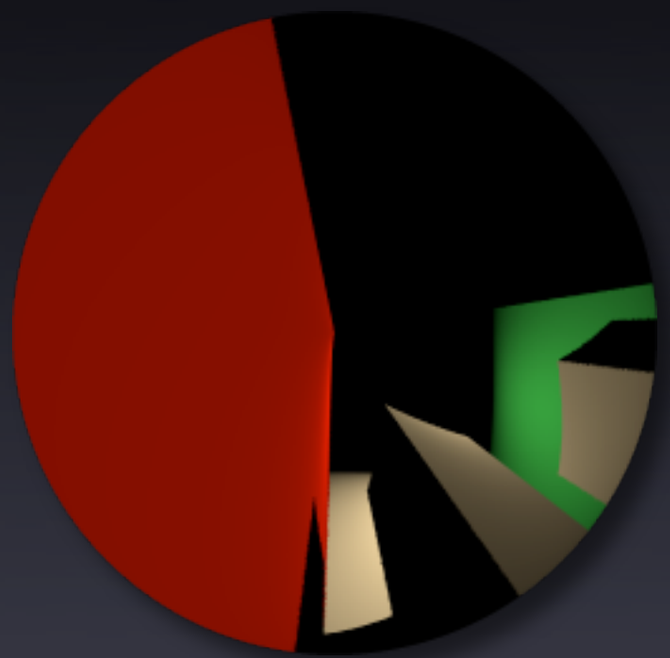

Computing Indirect Illum.

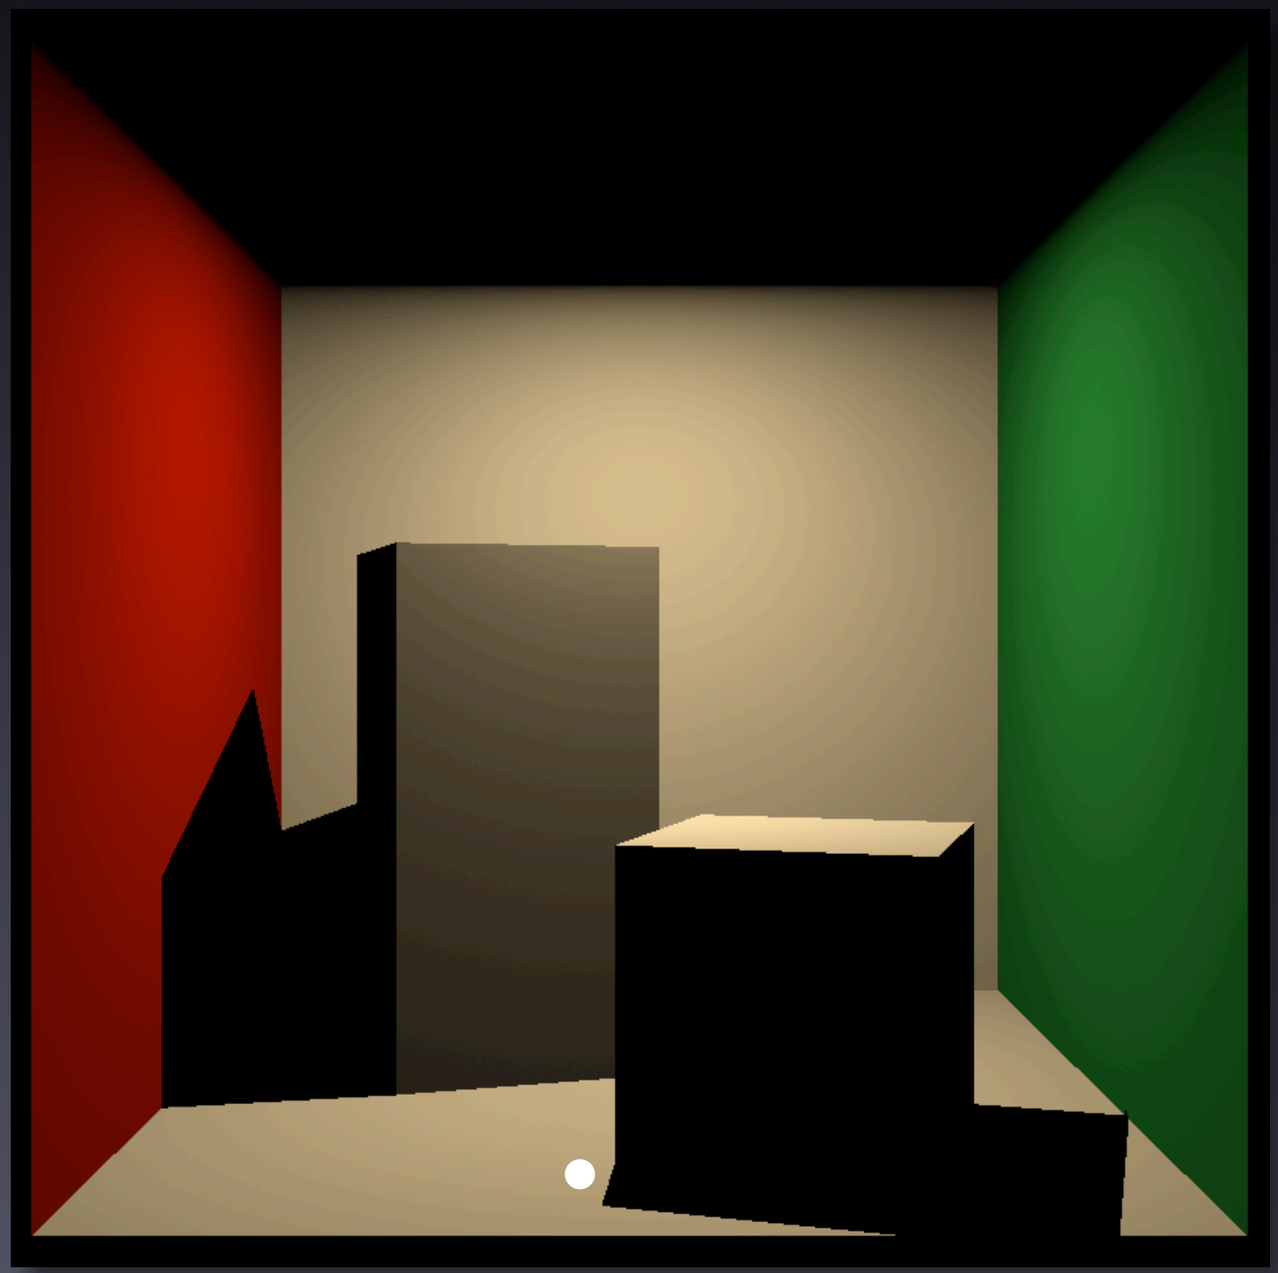

Computing Indirect Illum.

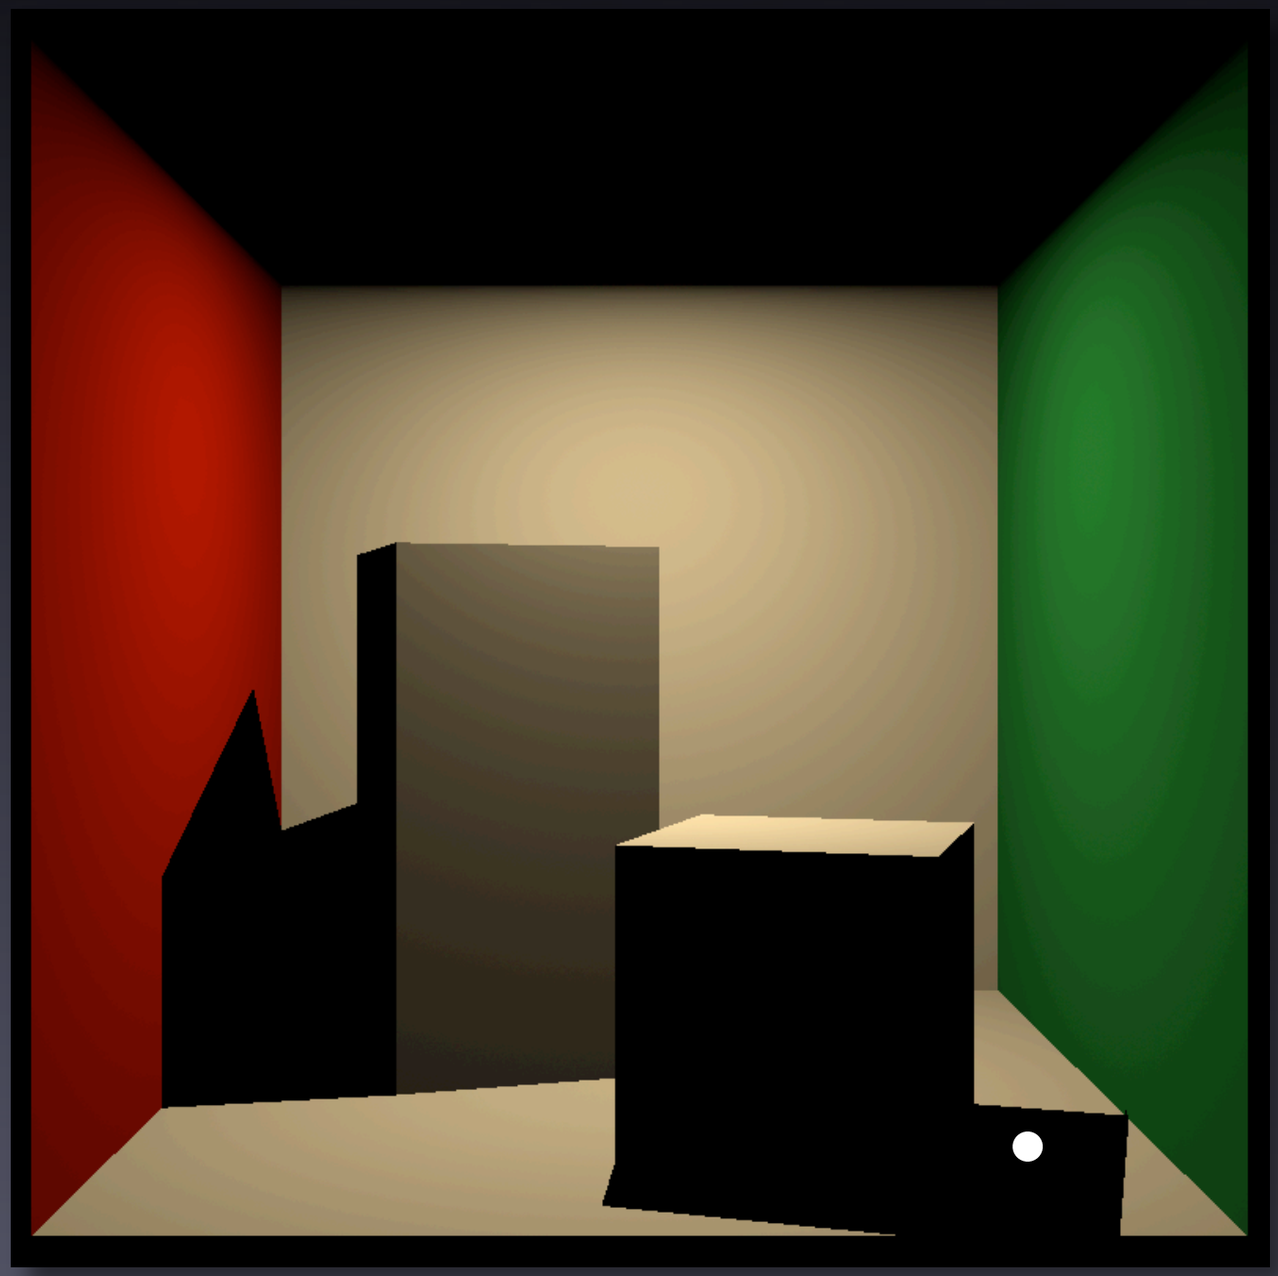

Direct Illumination

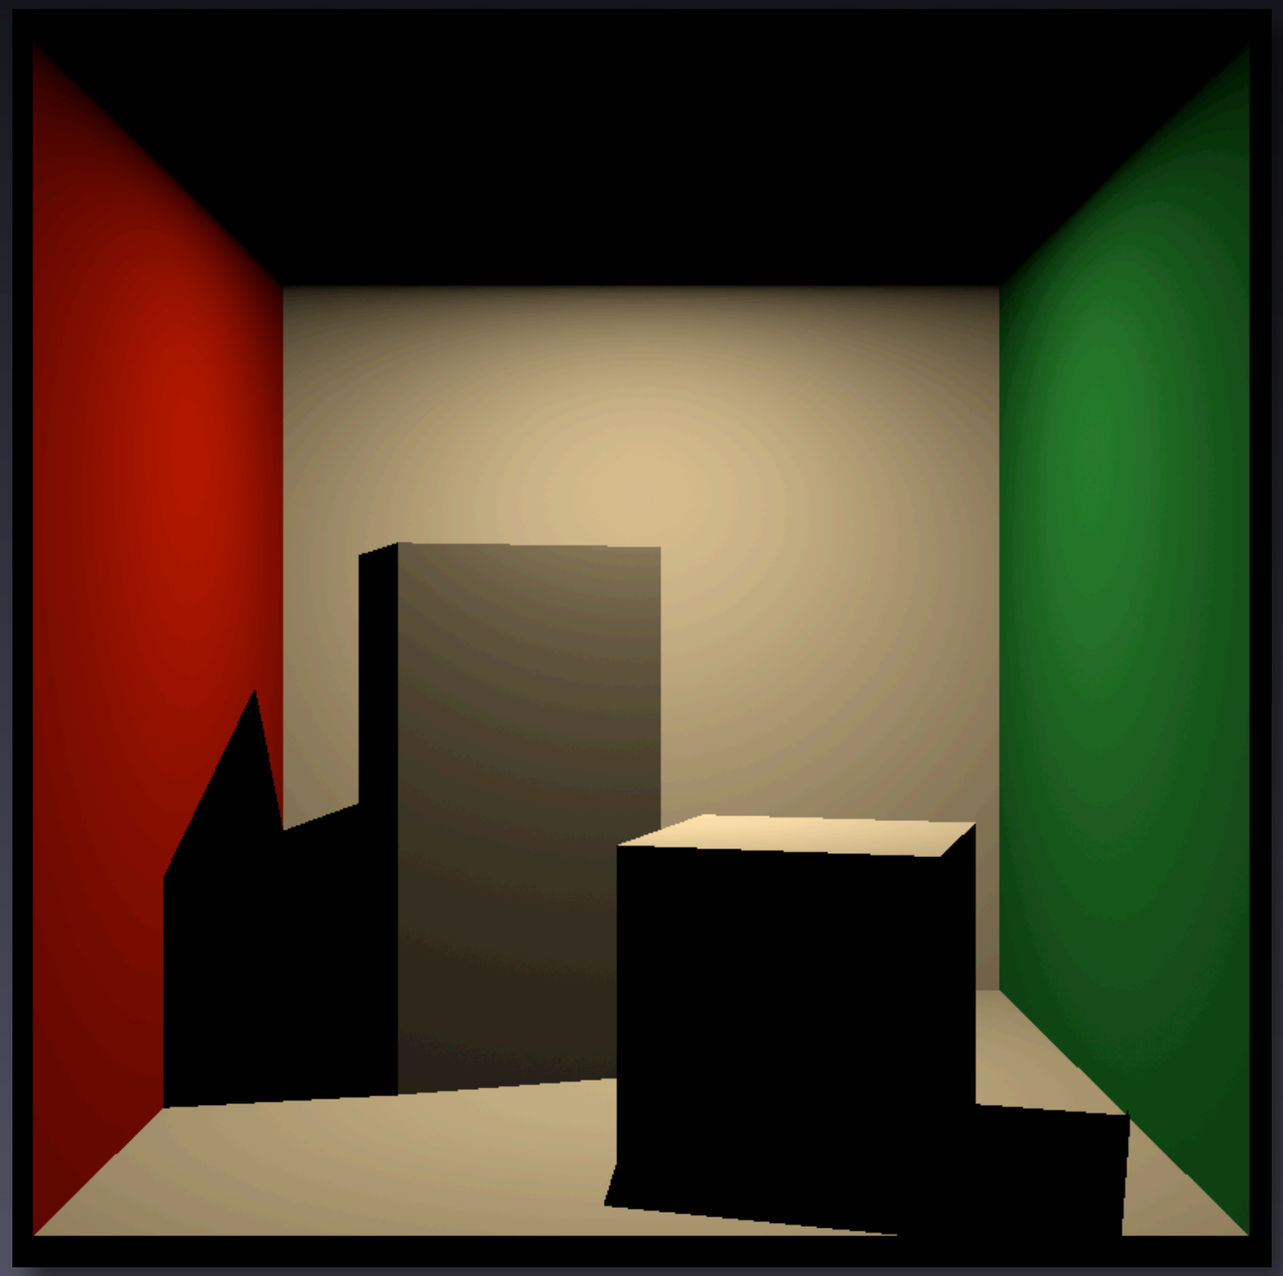

Direct + Indirect Illum.

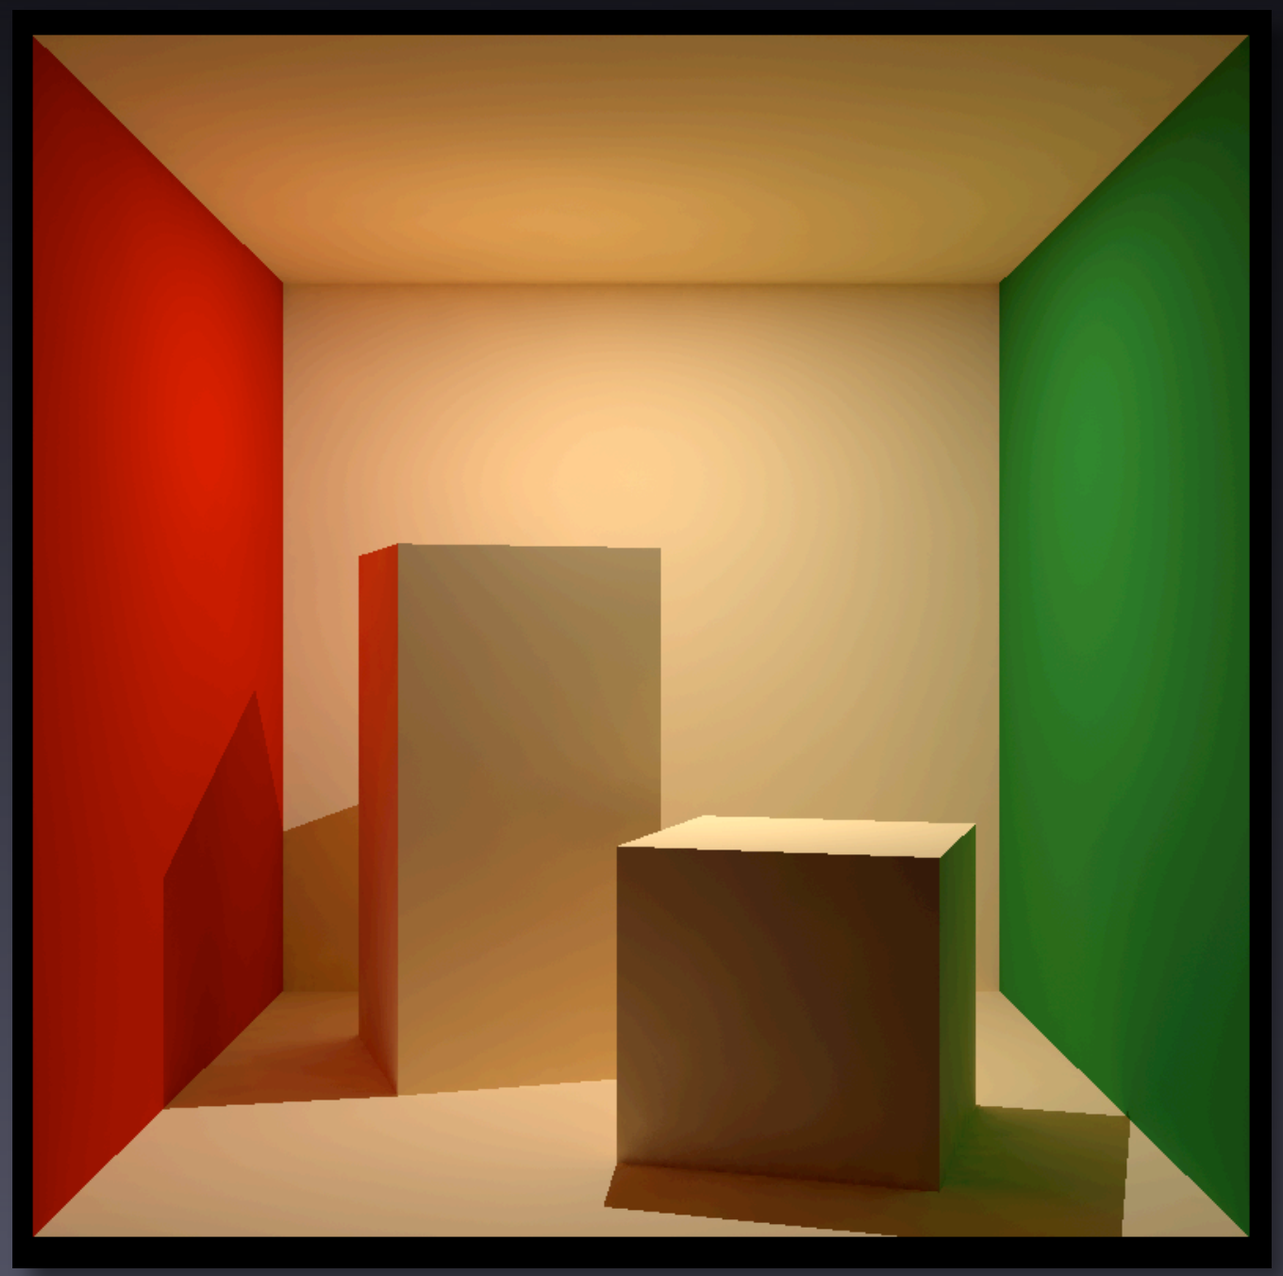

Indirect Illumination

Direct Illumination

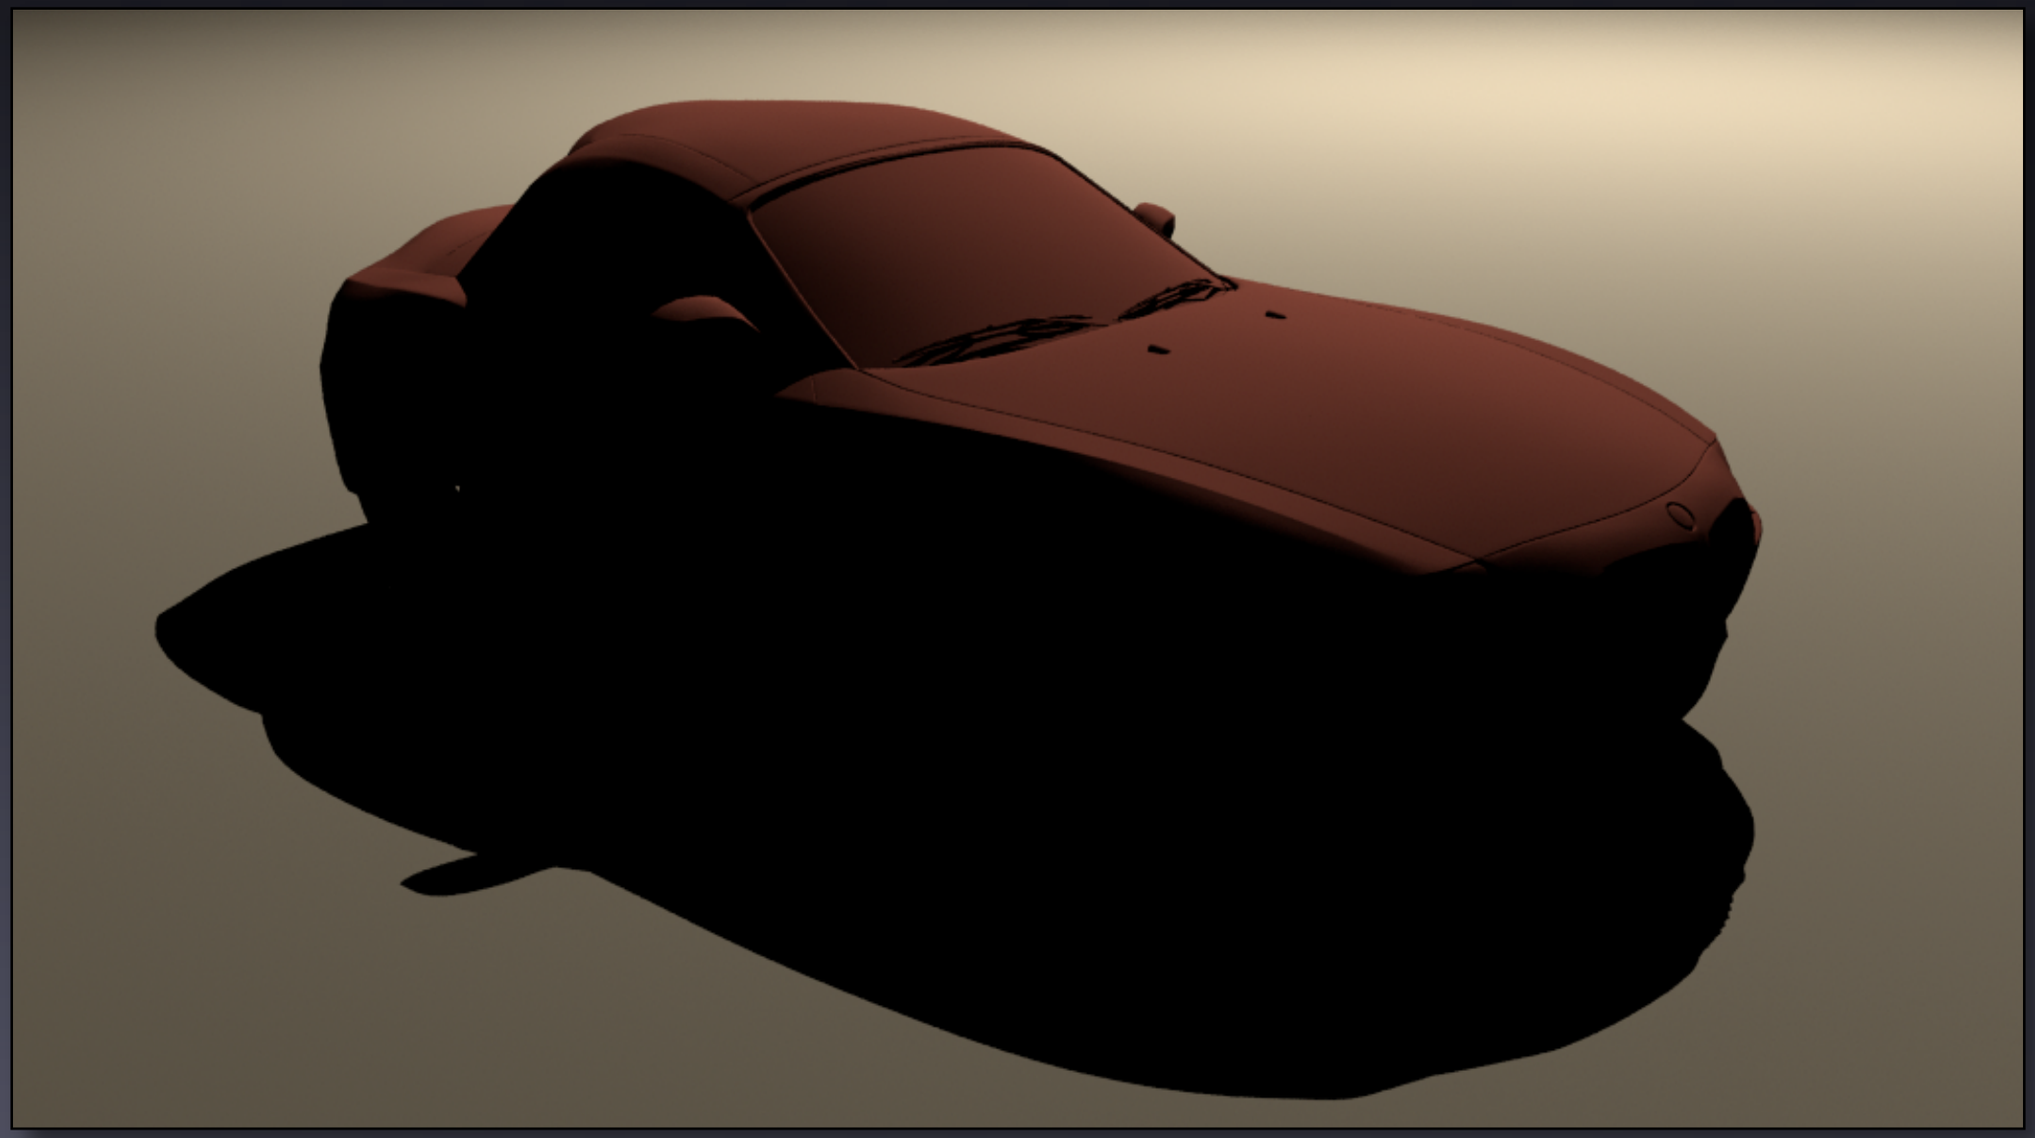

Direct Illumination

Direct + Indirect Illum.

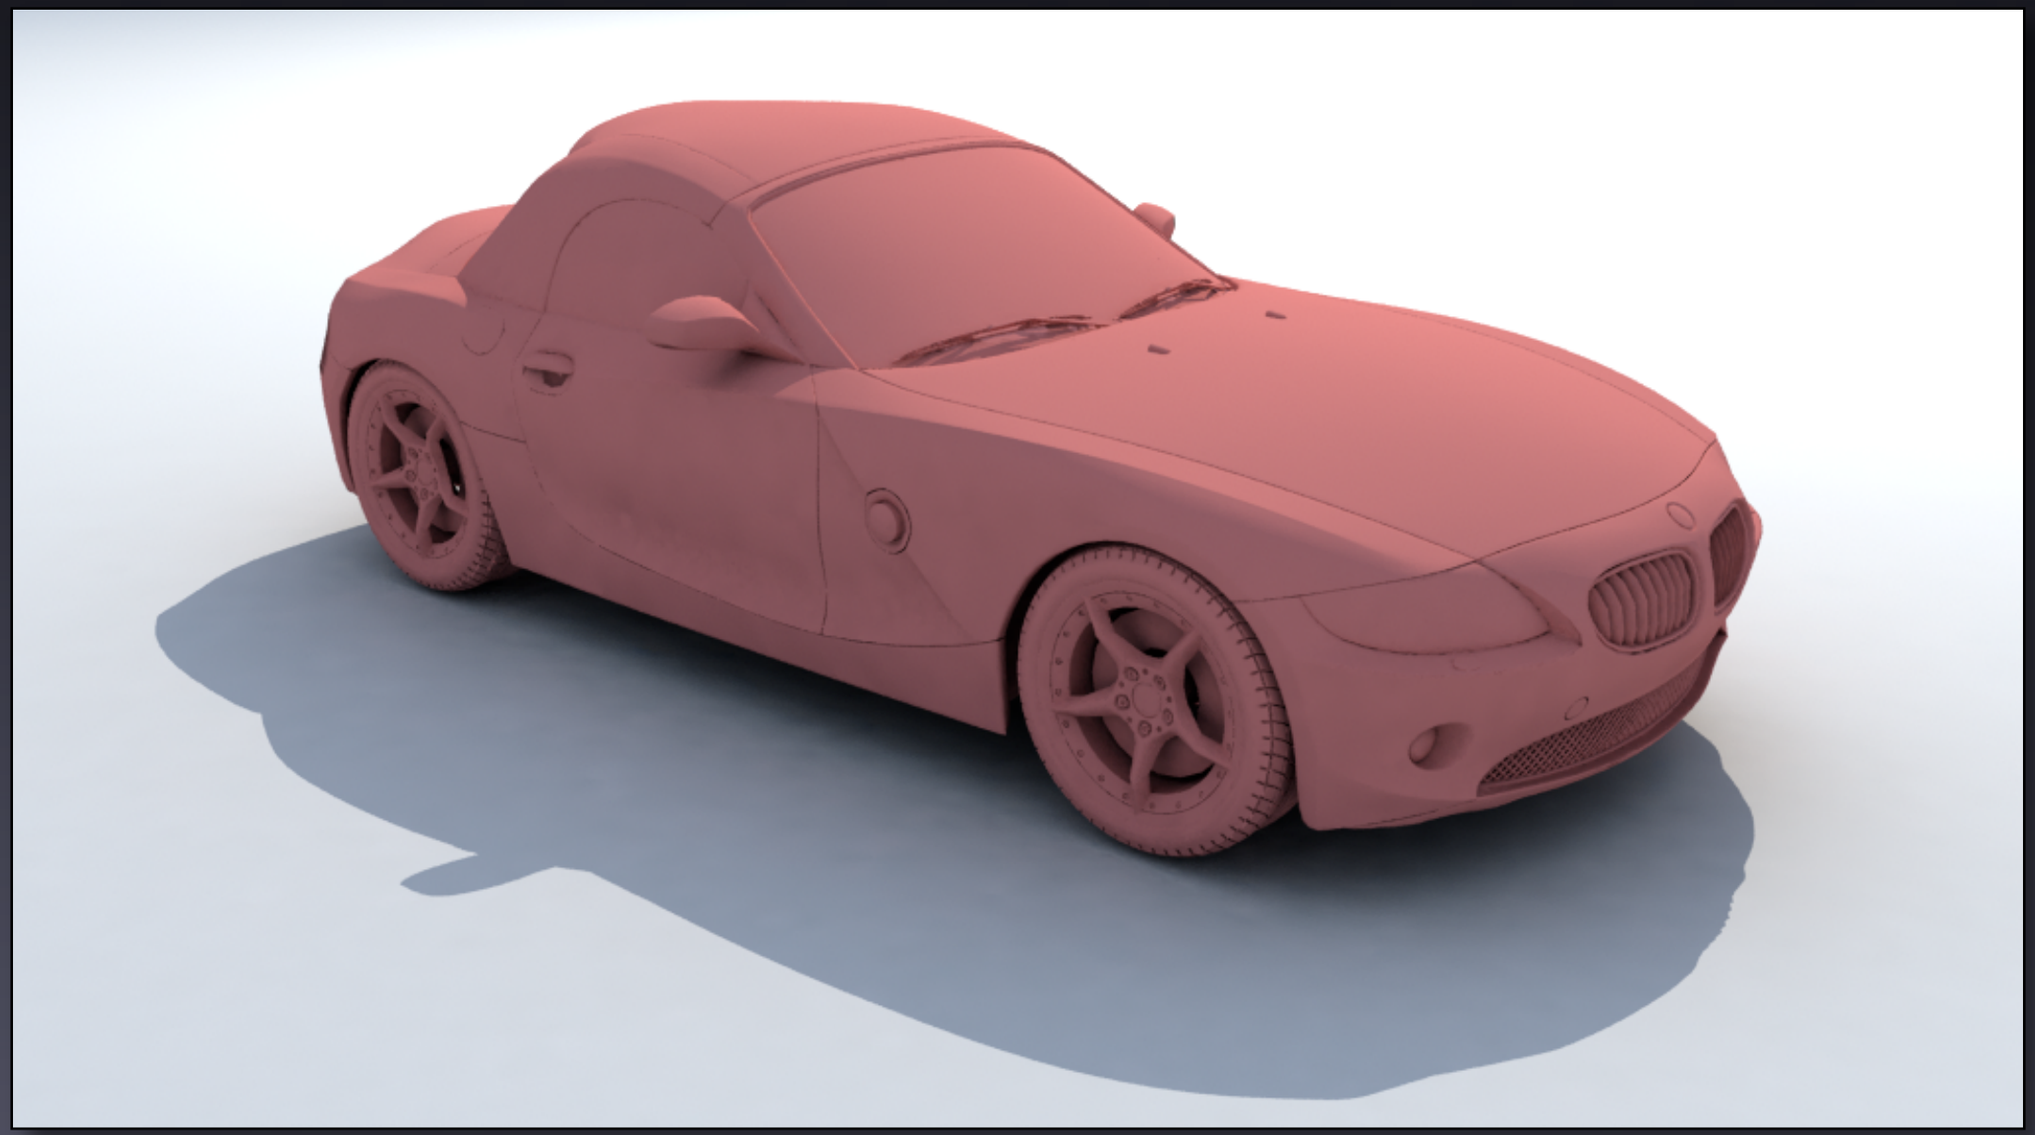

- Direct illumination is easy
- We will focus only on indirect illumination

Path Tracing

Kajiya ‘86

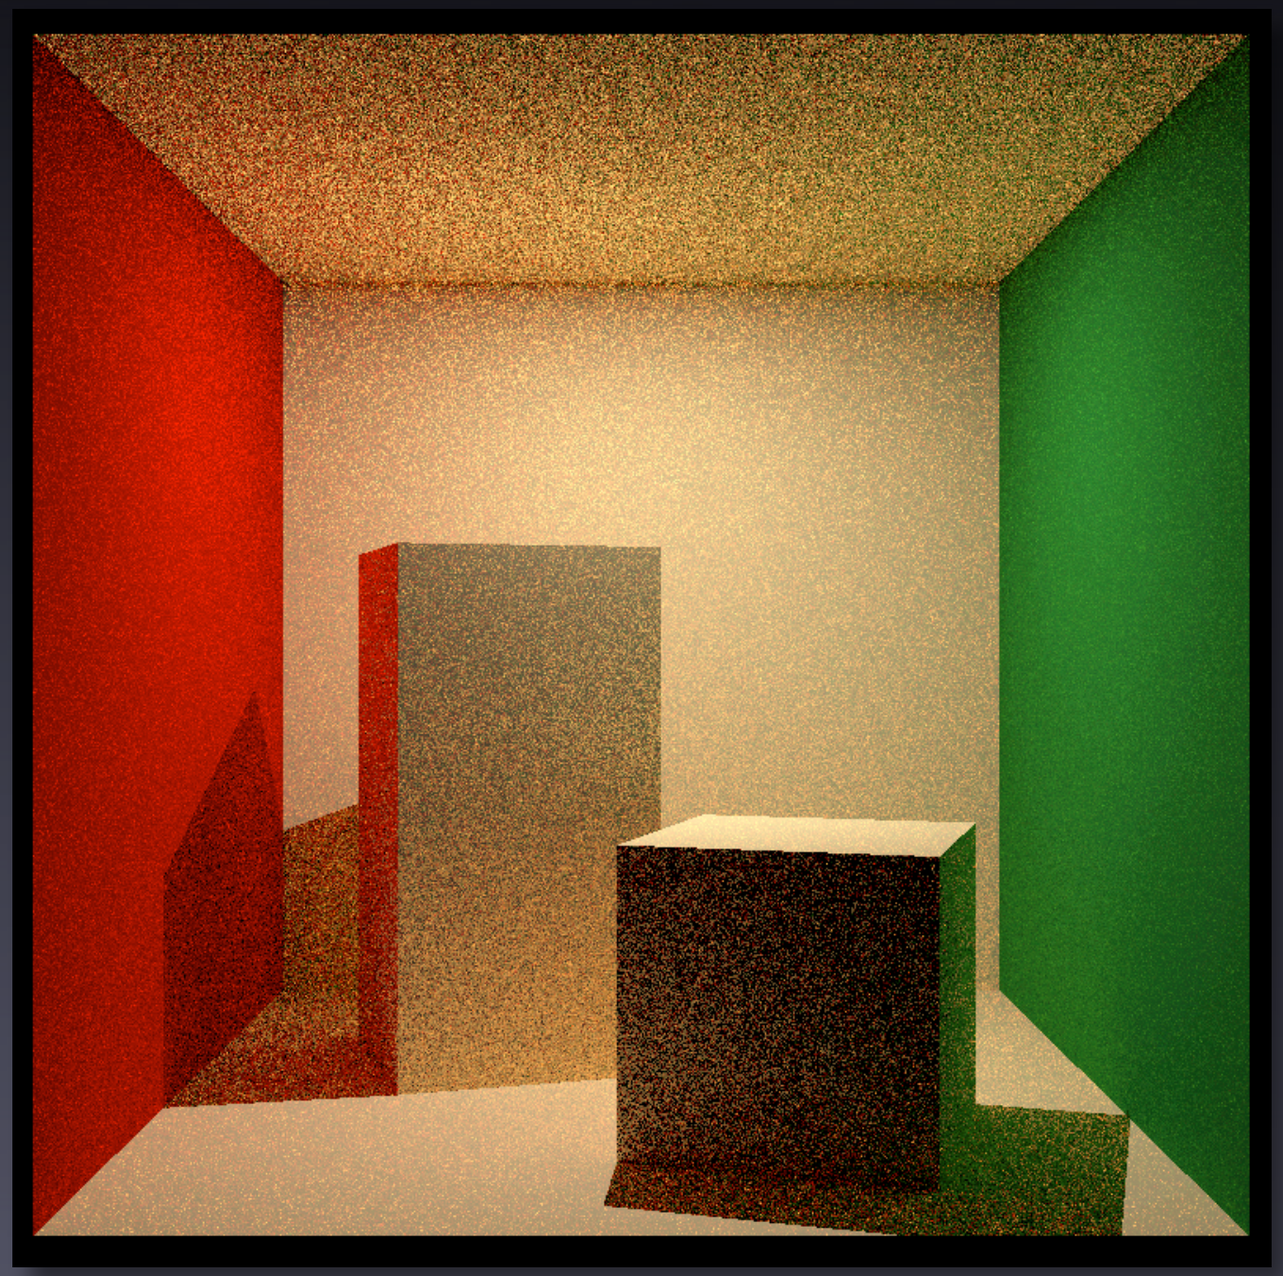

4 rays/pixel

Path Tracing

Kajiya ‘86

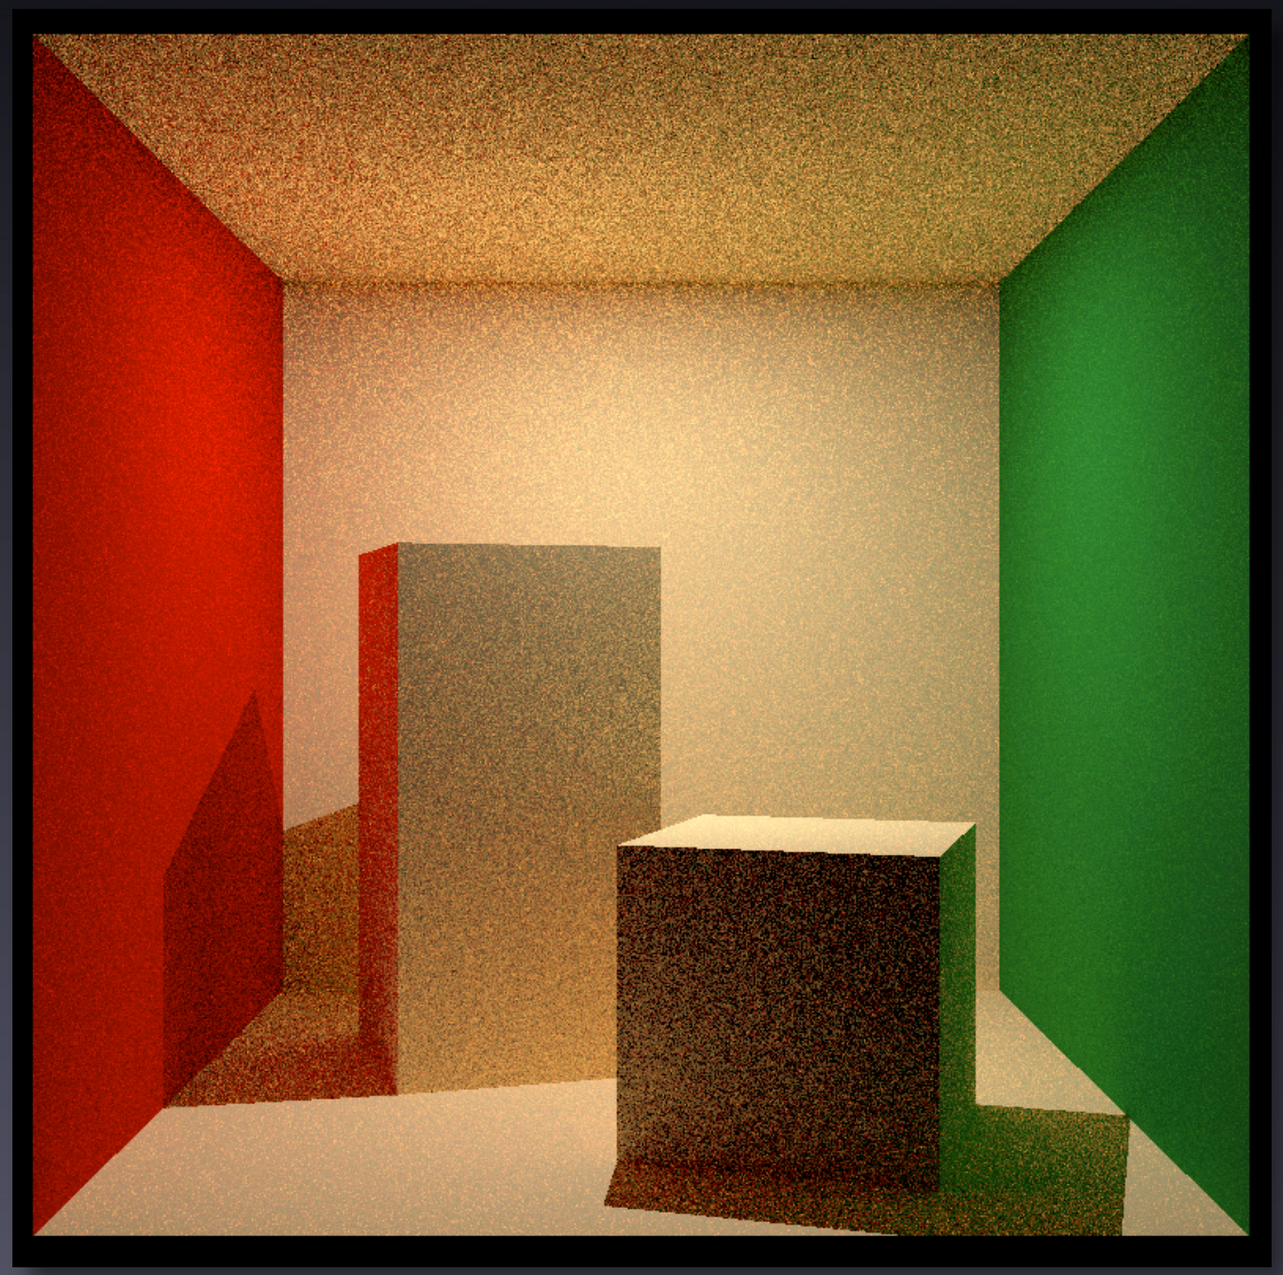

8 rays/pixel

Path Tracing

Kajiya ‘86

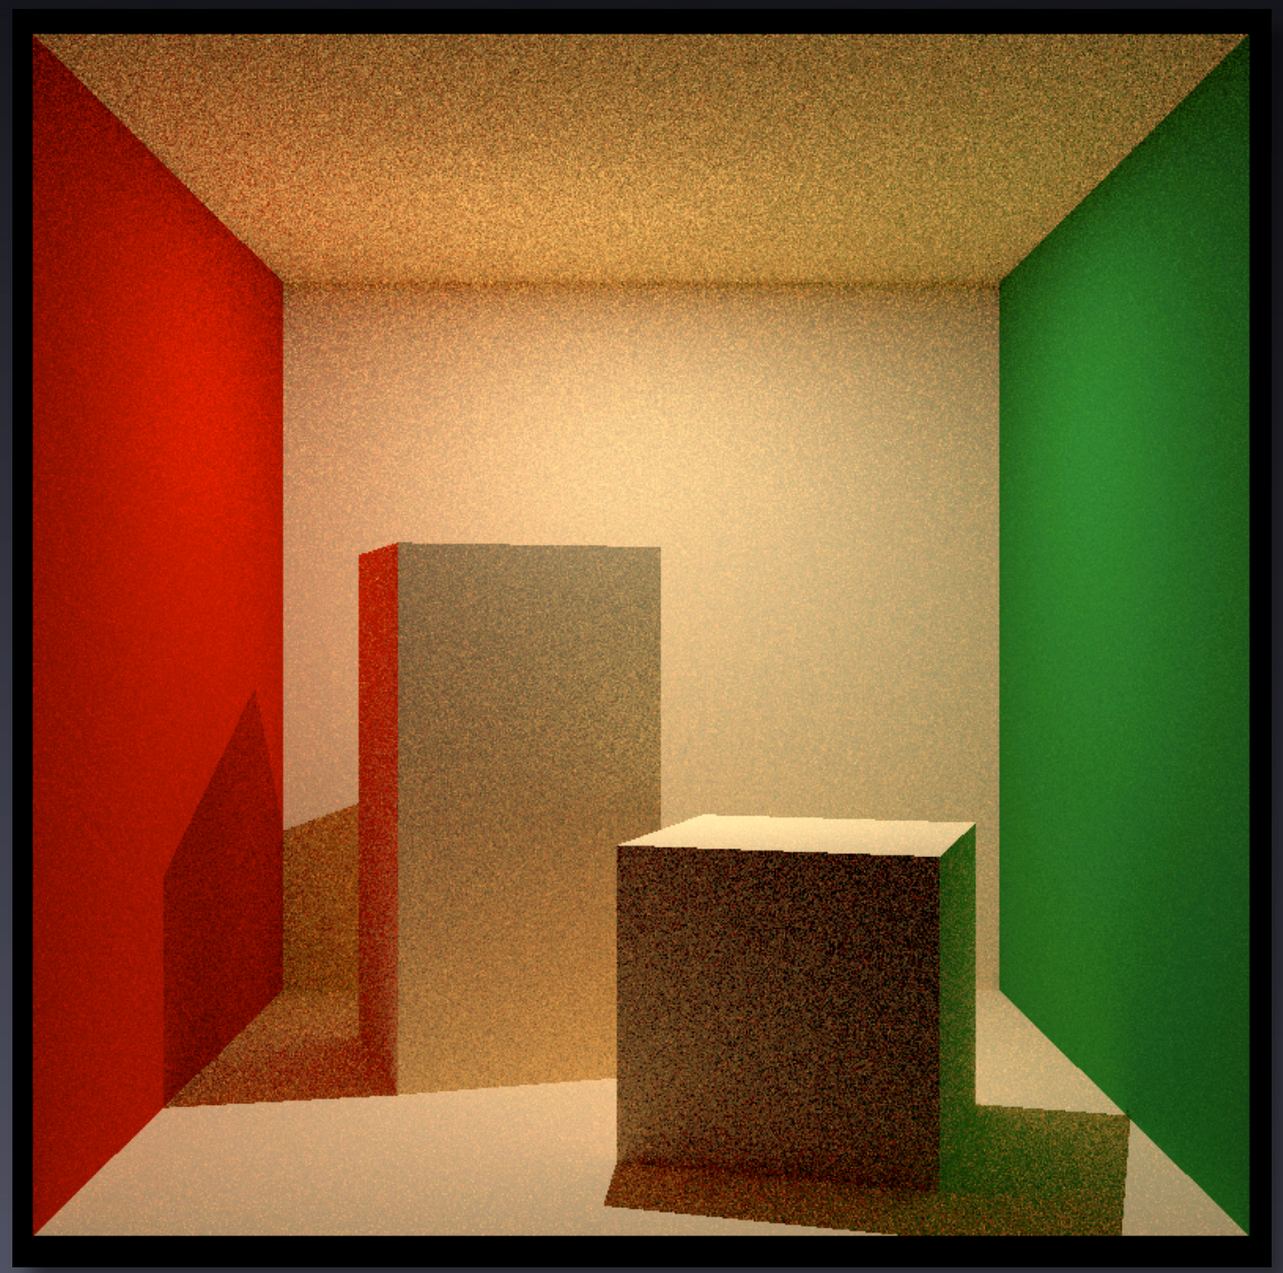

16 rays/pixel

Path Tracing

Kajiya ‘86

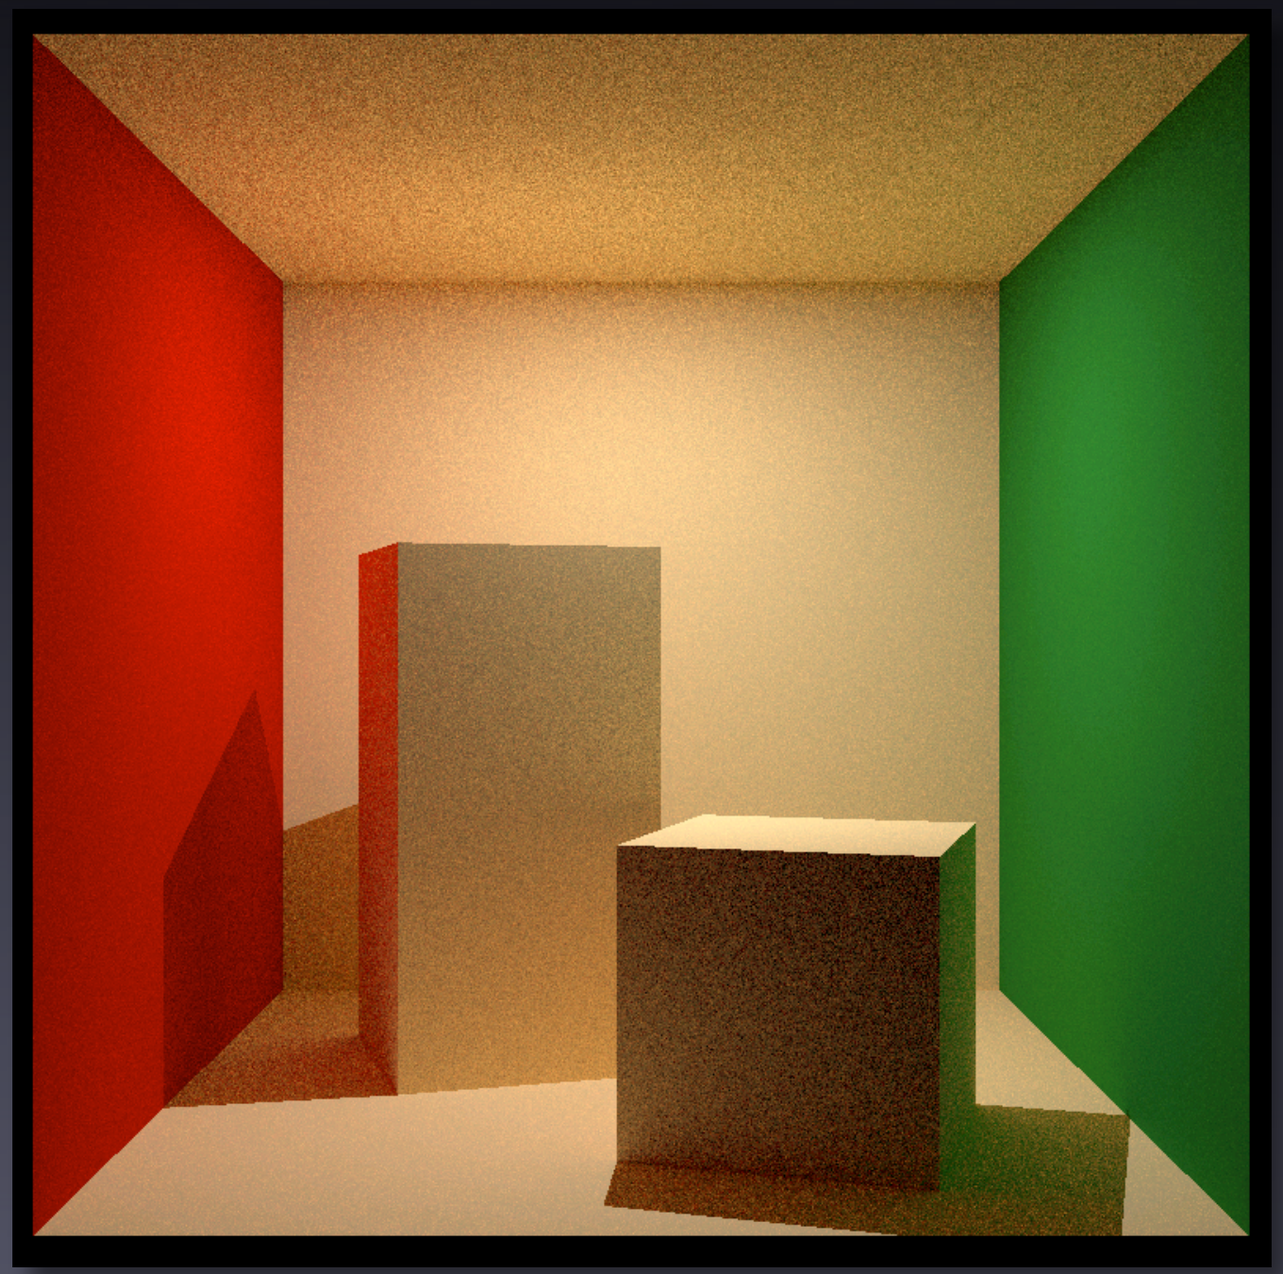

32 rays/pixel

Path Tracing

Kajiya ‘86

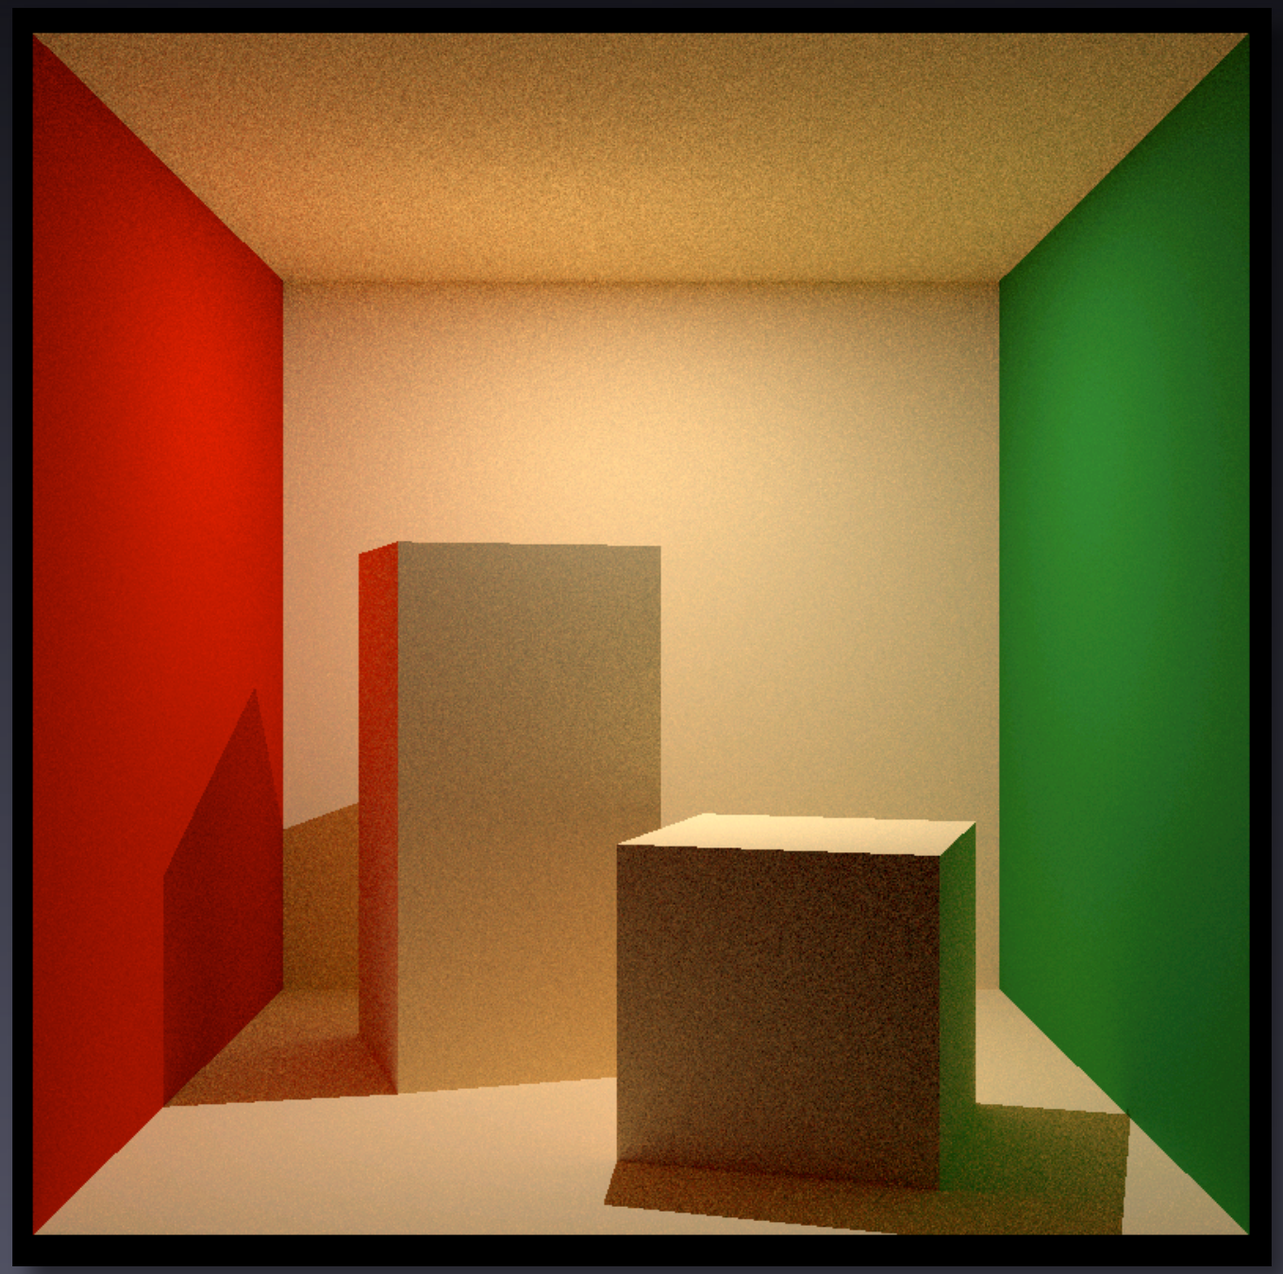

64 rays/pixel

Path Tracing

Kajiya ‘86

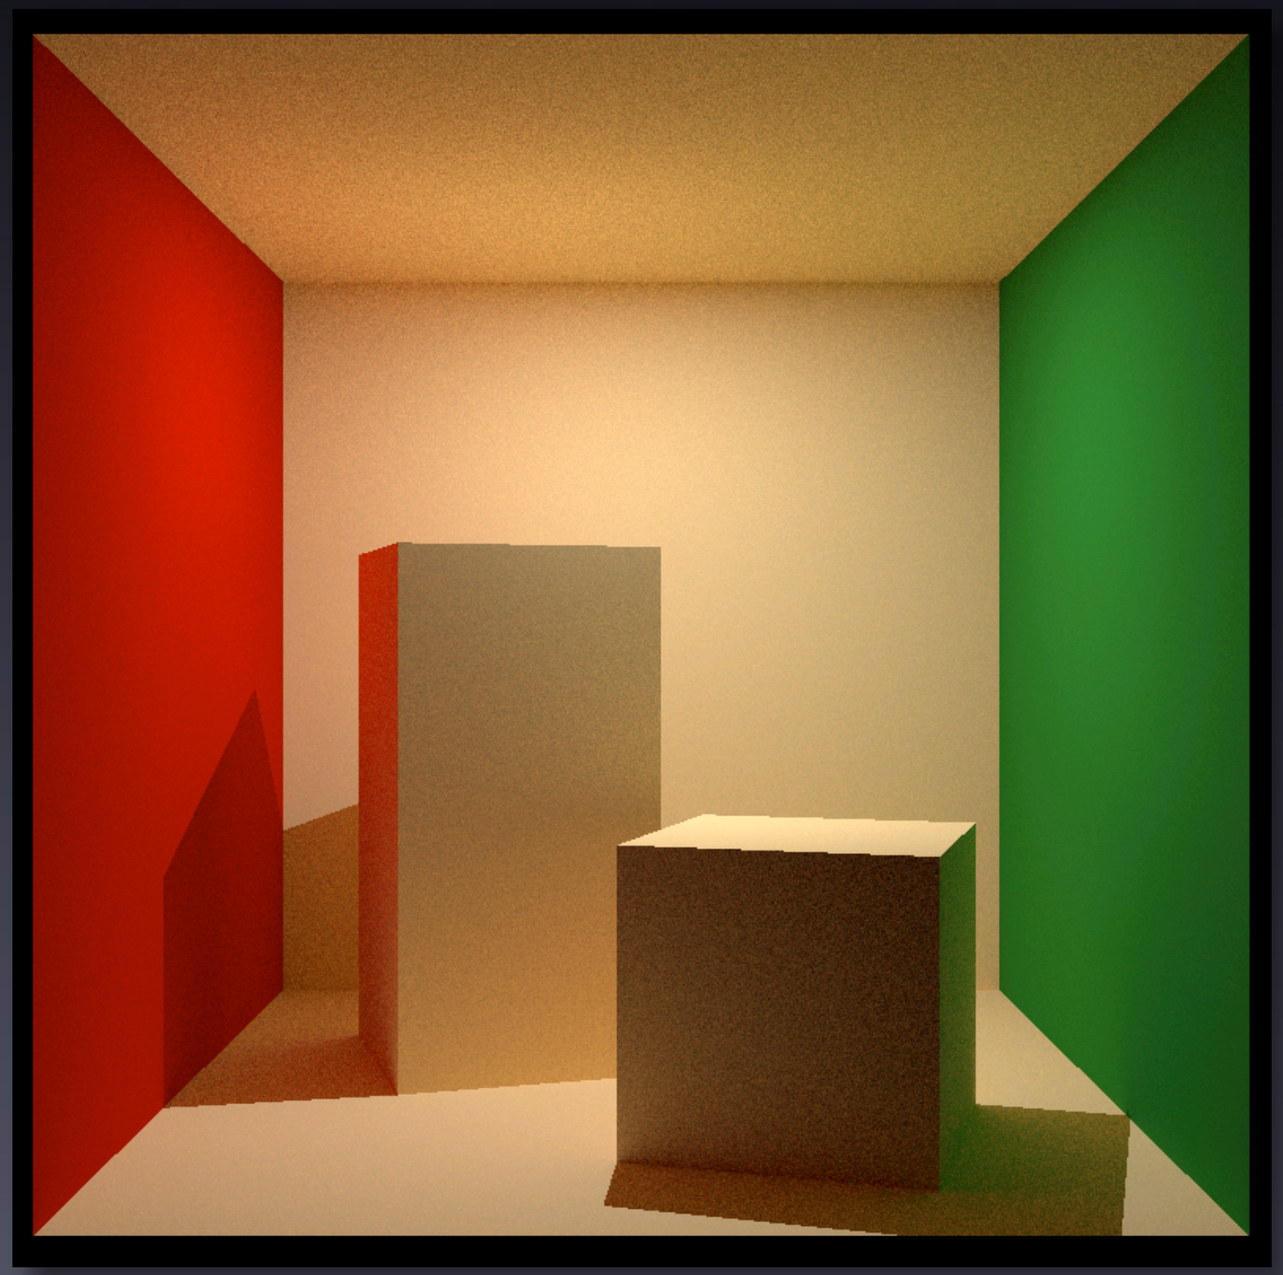

128 rays/pixel

Photon Mapping

Direct Visualization

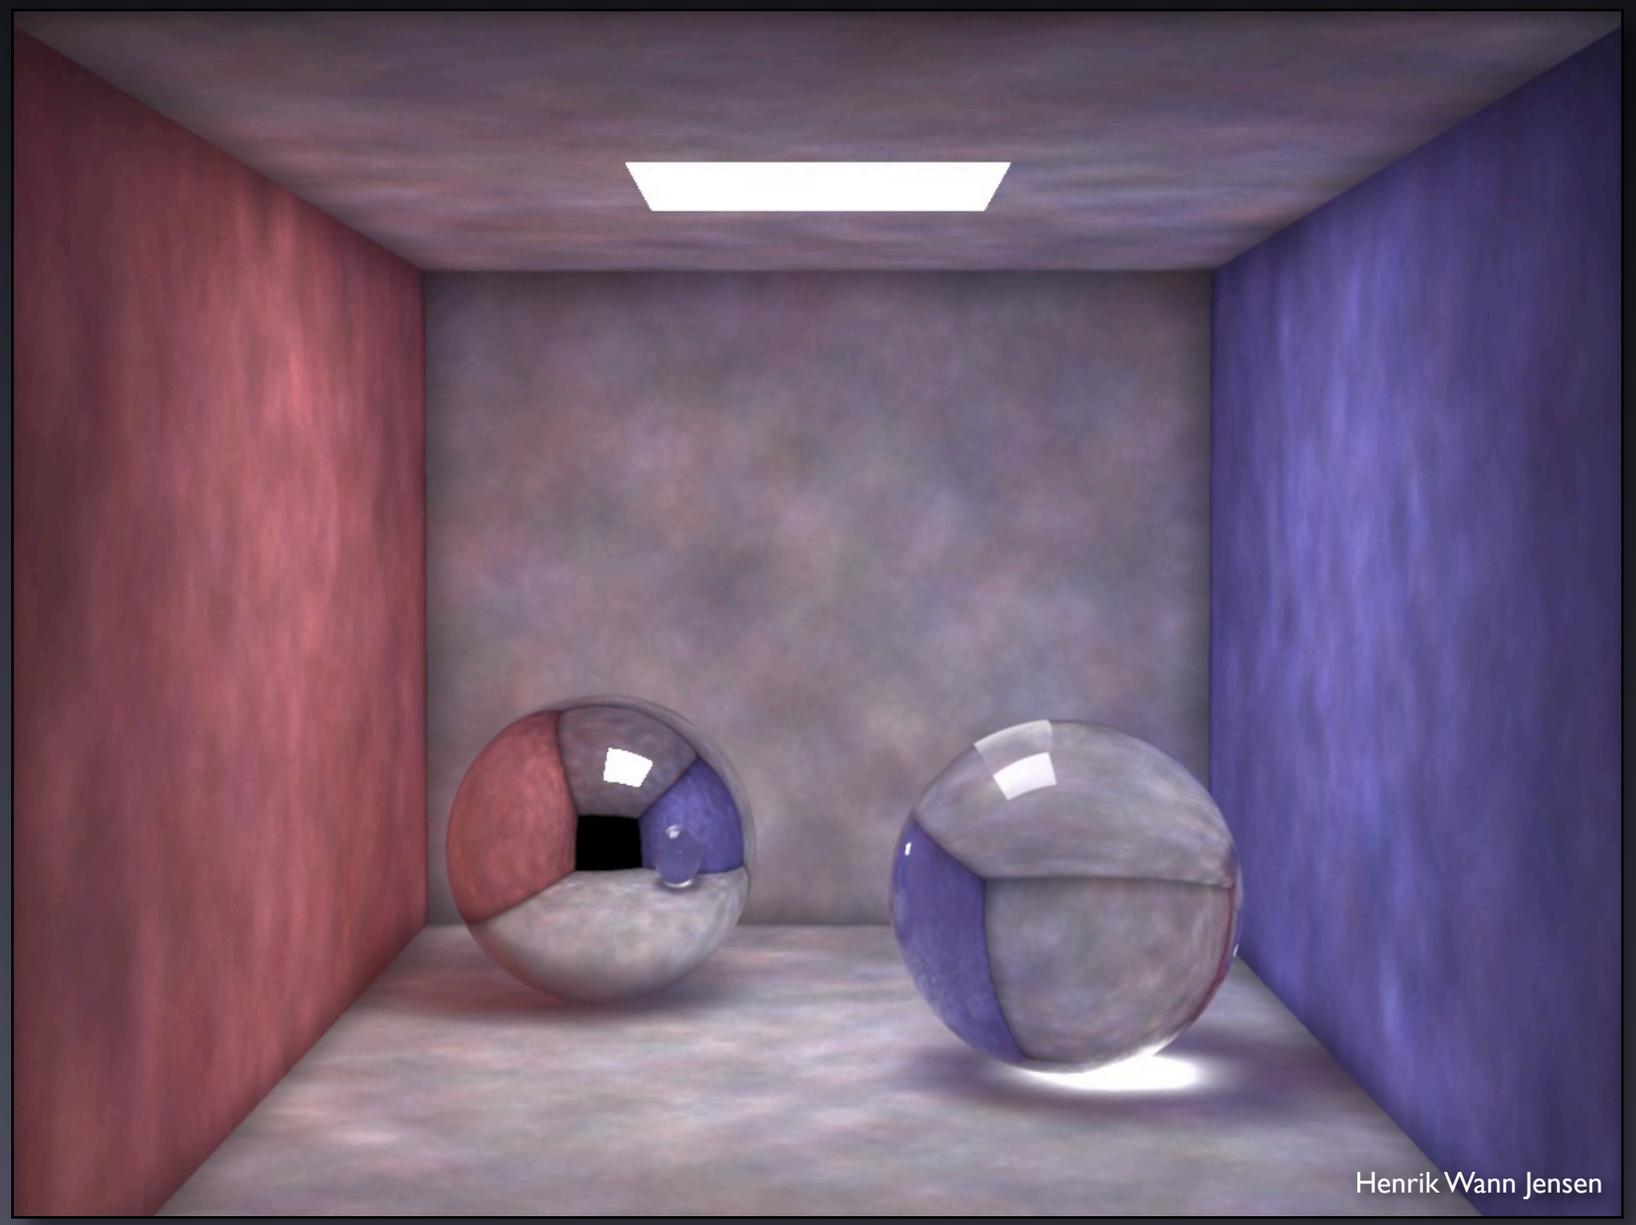

100000 photons / 50 photons in radiance estimate

Photon Mapping

Direct Visualization

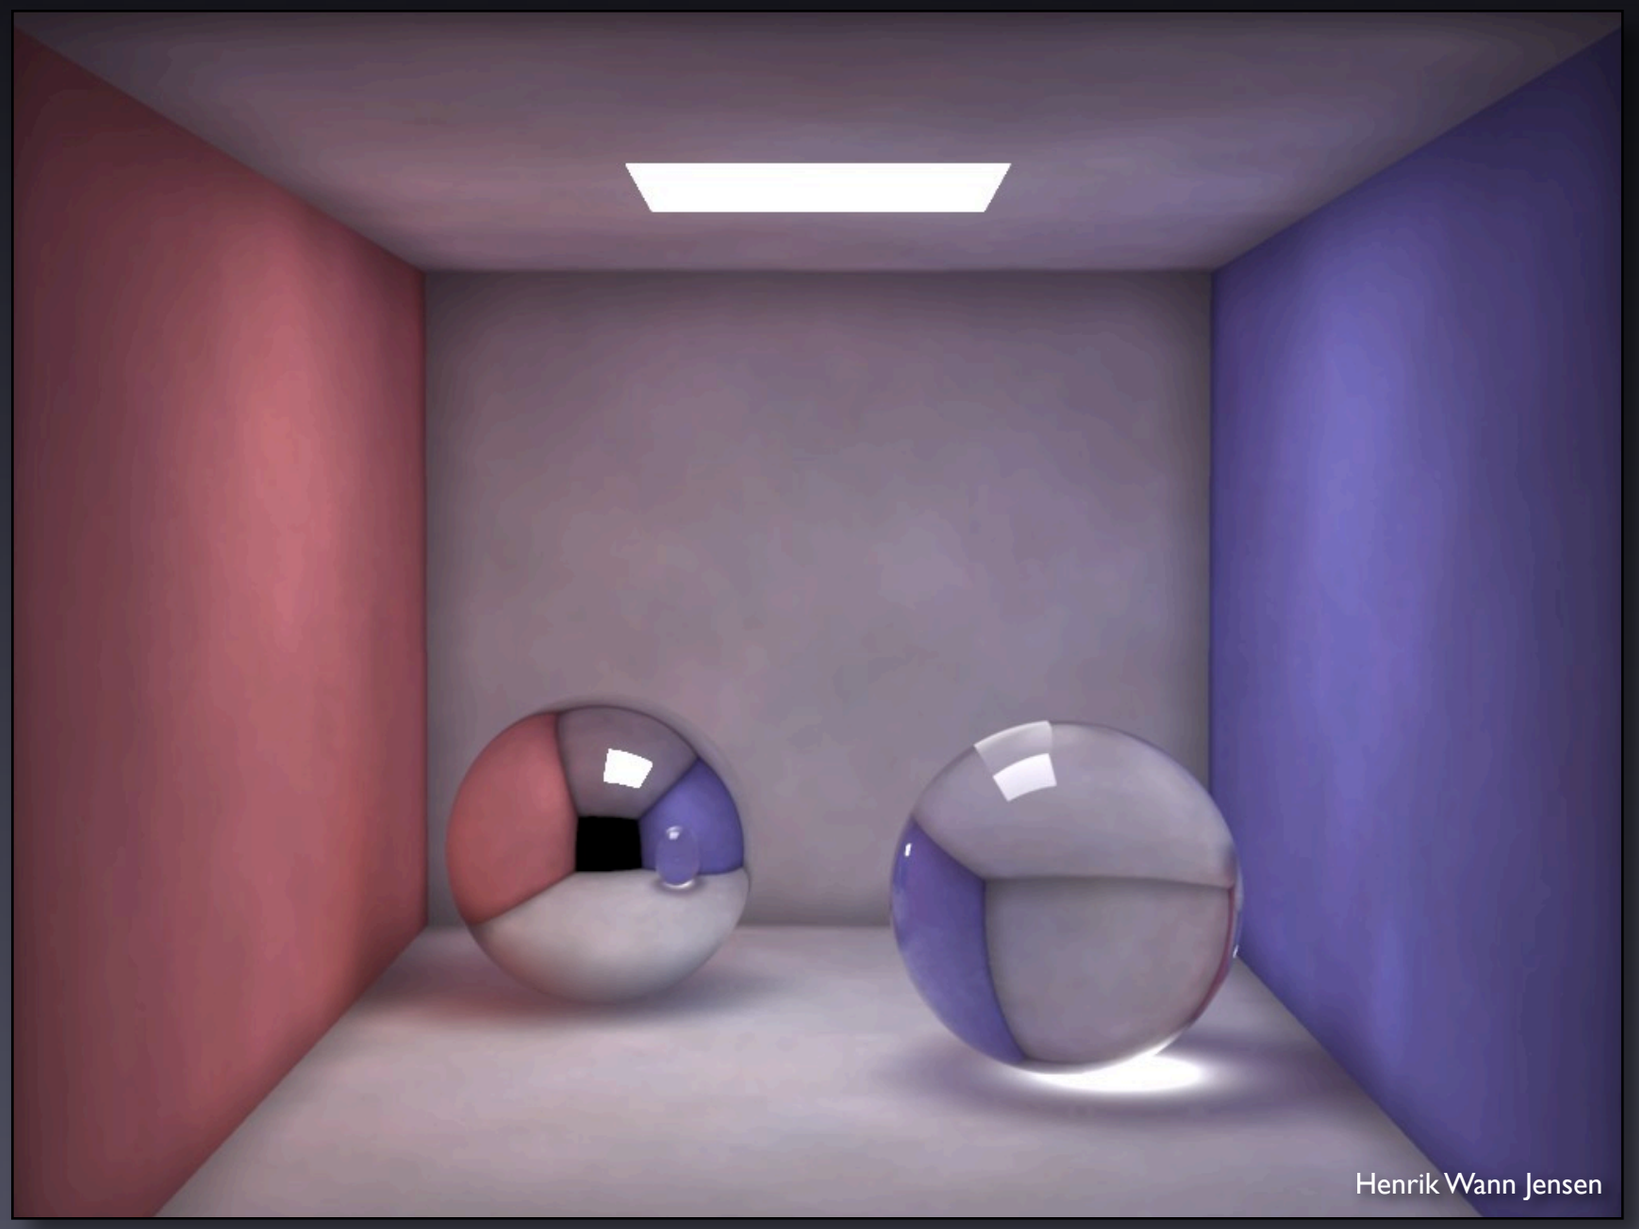

500000 photons / 500 photons in radiance estimate

Final Gather

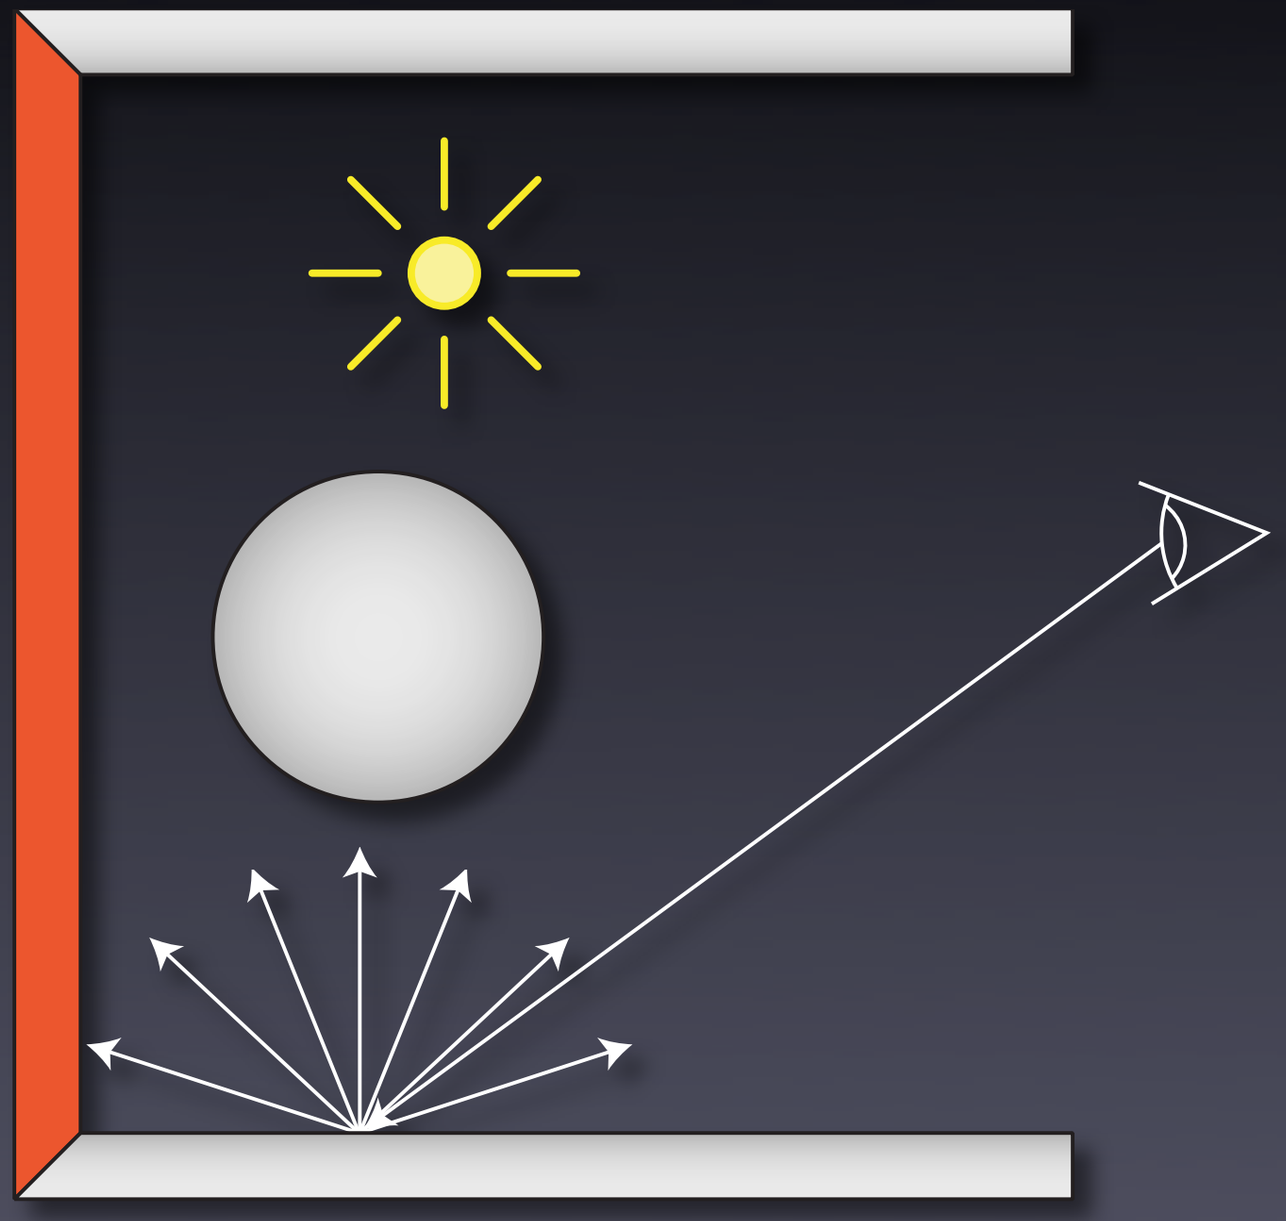

Motivation

- Tracing rays is costly.
- High-quality indirect illumination may require hundreds of gather rays per pixel for noise-free results.
- Hence, computing indirect illumination in this naive way is *very* costly.

Outline

- Introduction
- What is Irradiance Caching?
- Extensions to Irradiance Caching
- Future Work

Irradiance Caching

Ward et al. '88

- Irradiance caching was introduced by Ward et al. in 1988.

Direct Illumination

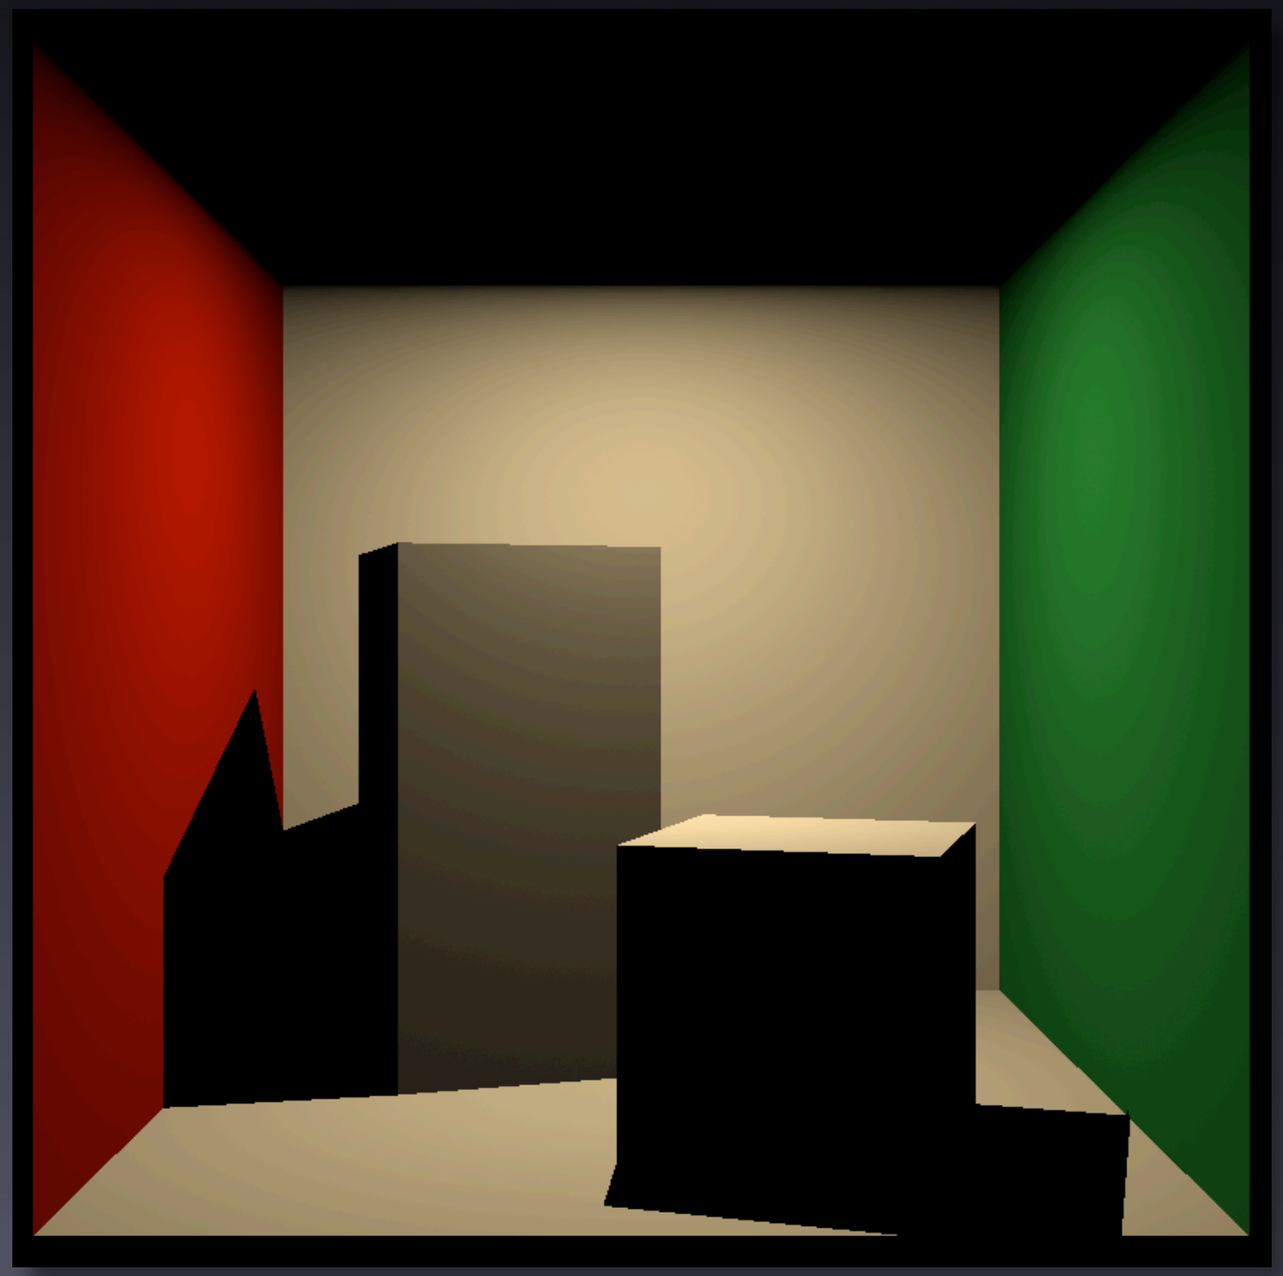

Indirect Illumination

Interpolate Indirect Illum.

Ward et al. '88

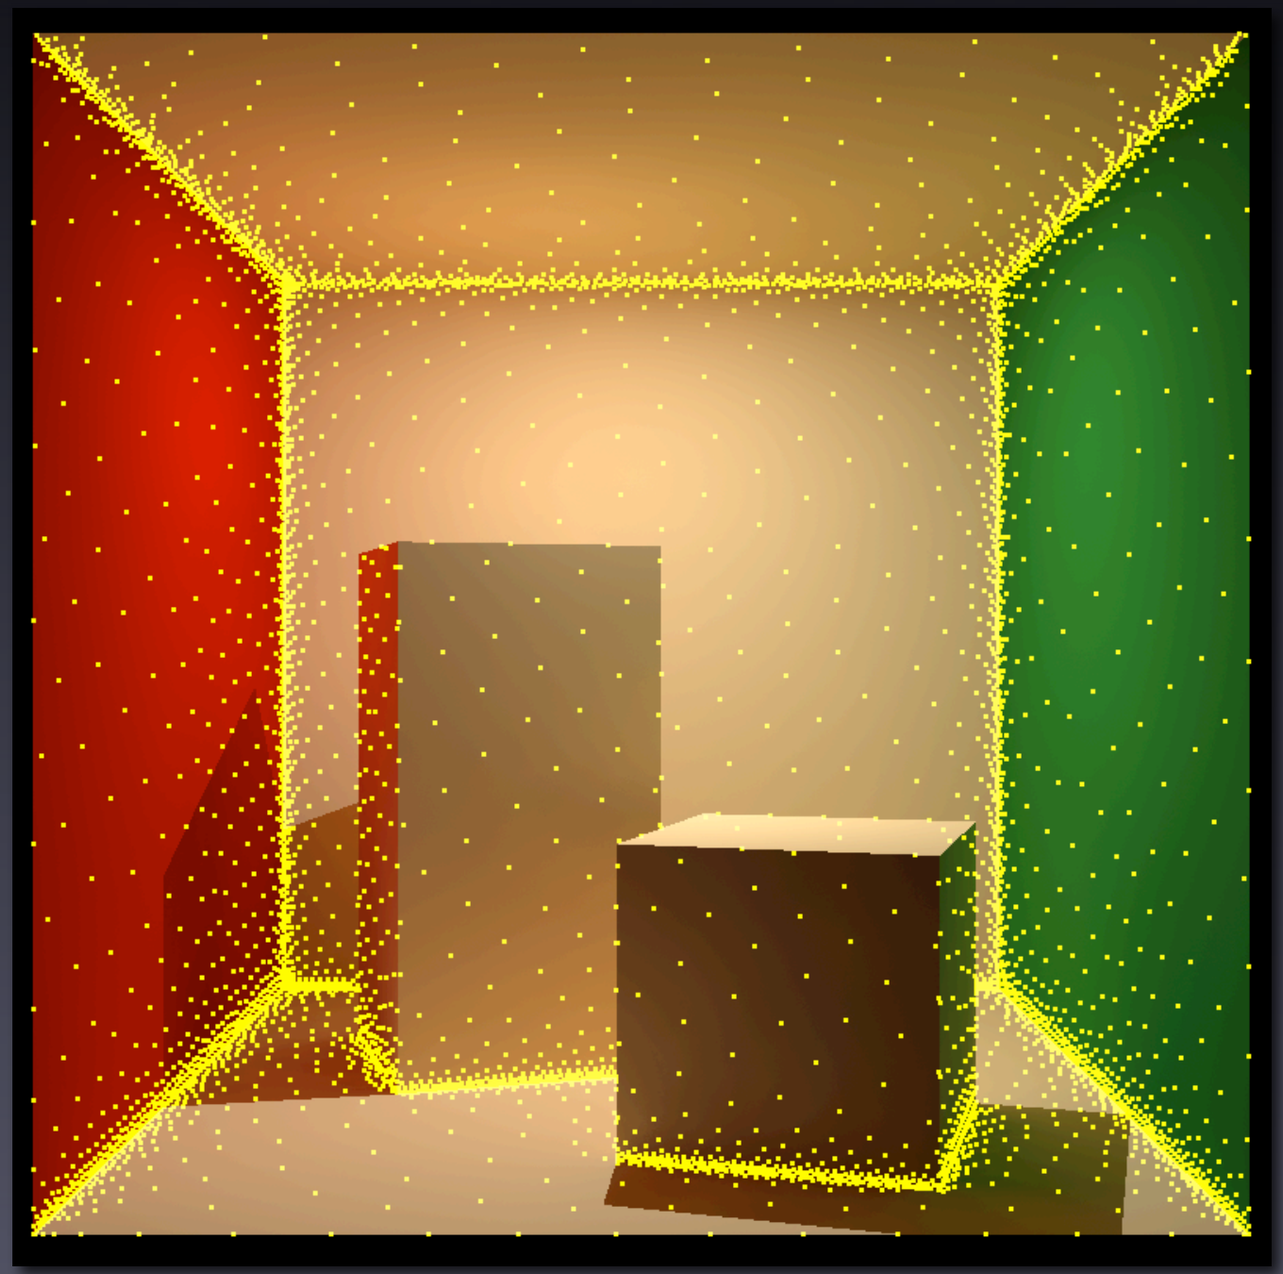

4621 samples

Irradiance Caching

Ward et al. '88

- Some questions that remain:
  - How do we compute the indirect illumination values?
  - What criterion is used to determine whether a cache point is ``near?''
  - How do we interpolate the nearby cached values?

What is Irradiance?

Ward et al. '88

$$E(\mathbf{x}) = \int_{\Omega} L_i(\mathbf{x}, \vec{\omega}_i)(\vec{\omega}_i \cdot \vec{n})\, d\vec{\omega}_i$$

What is Irradiance?

Ward et al. '88

$$E(\mathbf{x}) = \int_{\Omega} L_i(\mathbf{x}, \vec{\omega}_i)(\vec{\omega}_i \cdot \vec{n})\, d\vec{\omega}_i$$

What is Irradiance?

Ward et al. '88

$$E(\mathbf{x}) = \int_{\Omega} L_i(\mathbf{x}, \vec{\omega}_i)(\vec{\omega}_i \cdot \vec{n})\, d\vec{\omega}_i$$

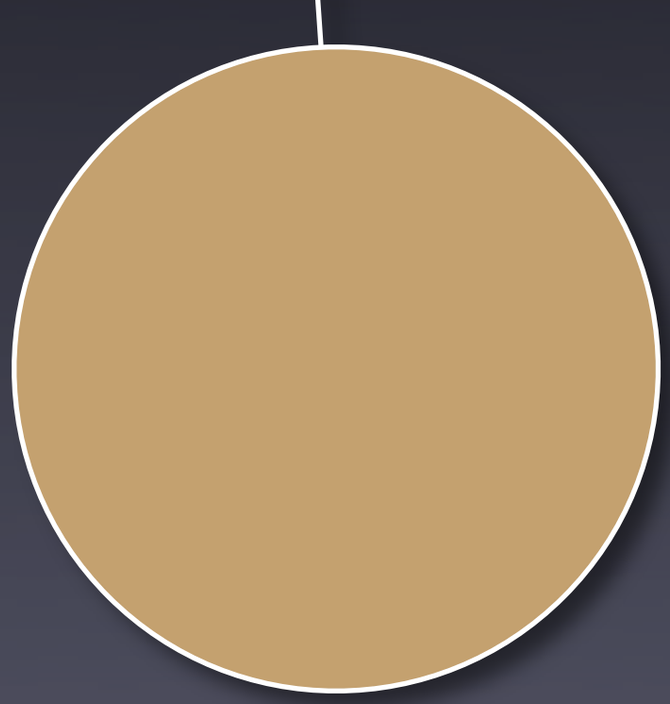

What is Irradiance?

Ward et al. '88

$$E(\mathbf{x}) = \int_{\Omega} L_i(\mathbf{x}, \vec{\omega}_i)(\vec{\omega}_i \cdot \vec{n})\, d\vec{\omega}_i$$

Nearly impossible to solve analytically!

Computing Irradiance

Ward et al. '88

Computing Irradiance

Ward et al. '88

Computing Irradiance

Ward et al. '88

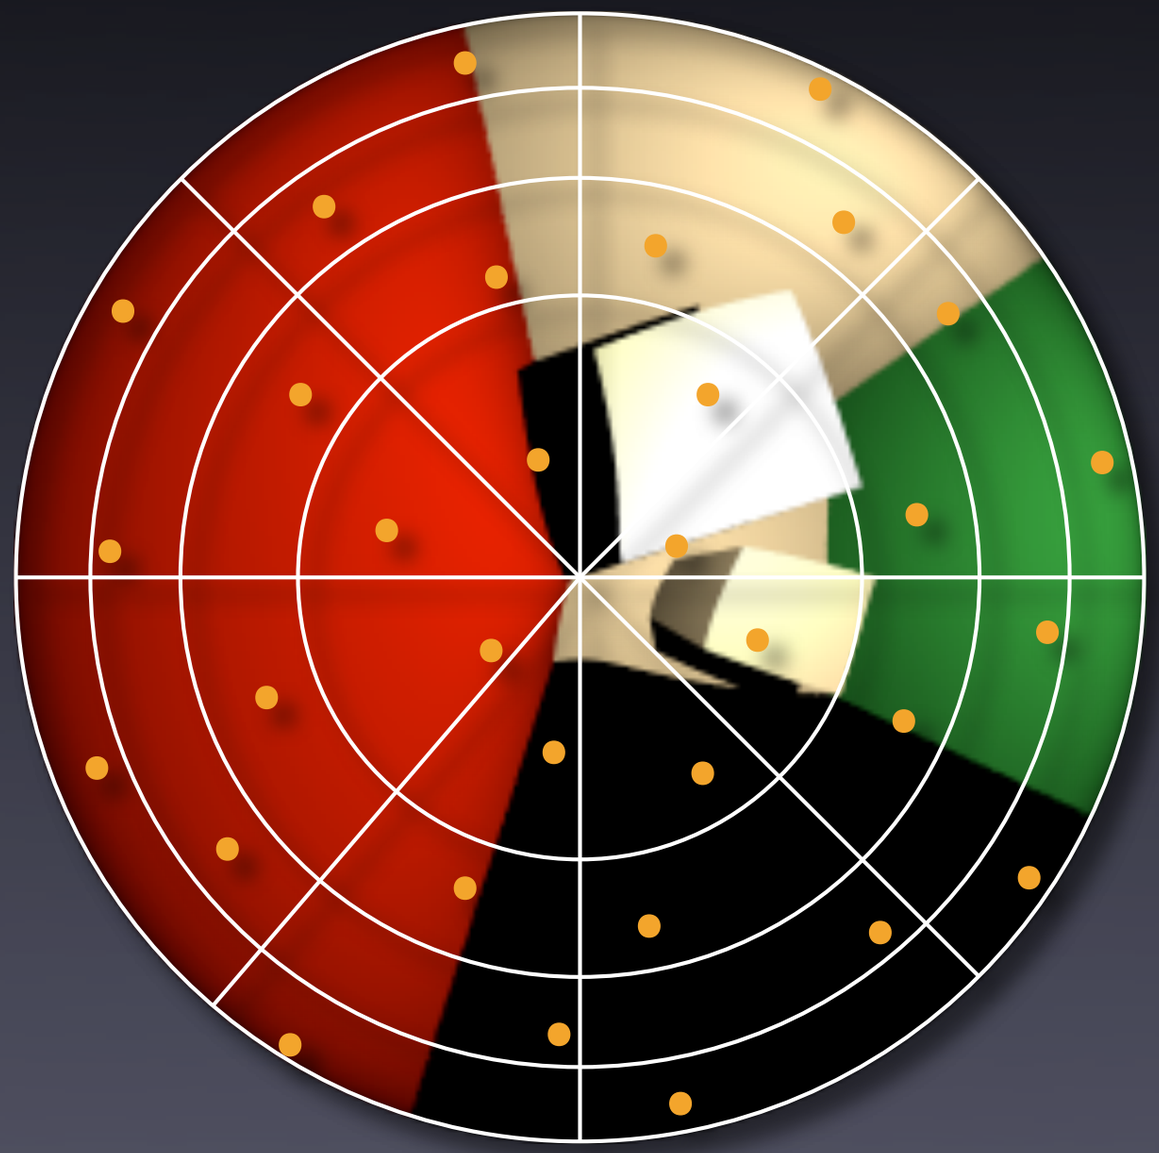

Computing Irradiance

Ward et al. '88

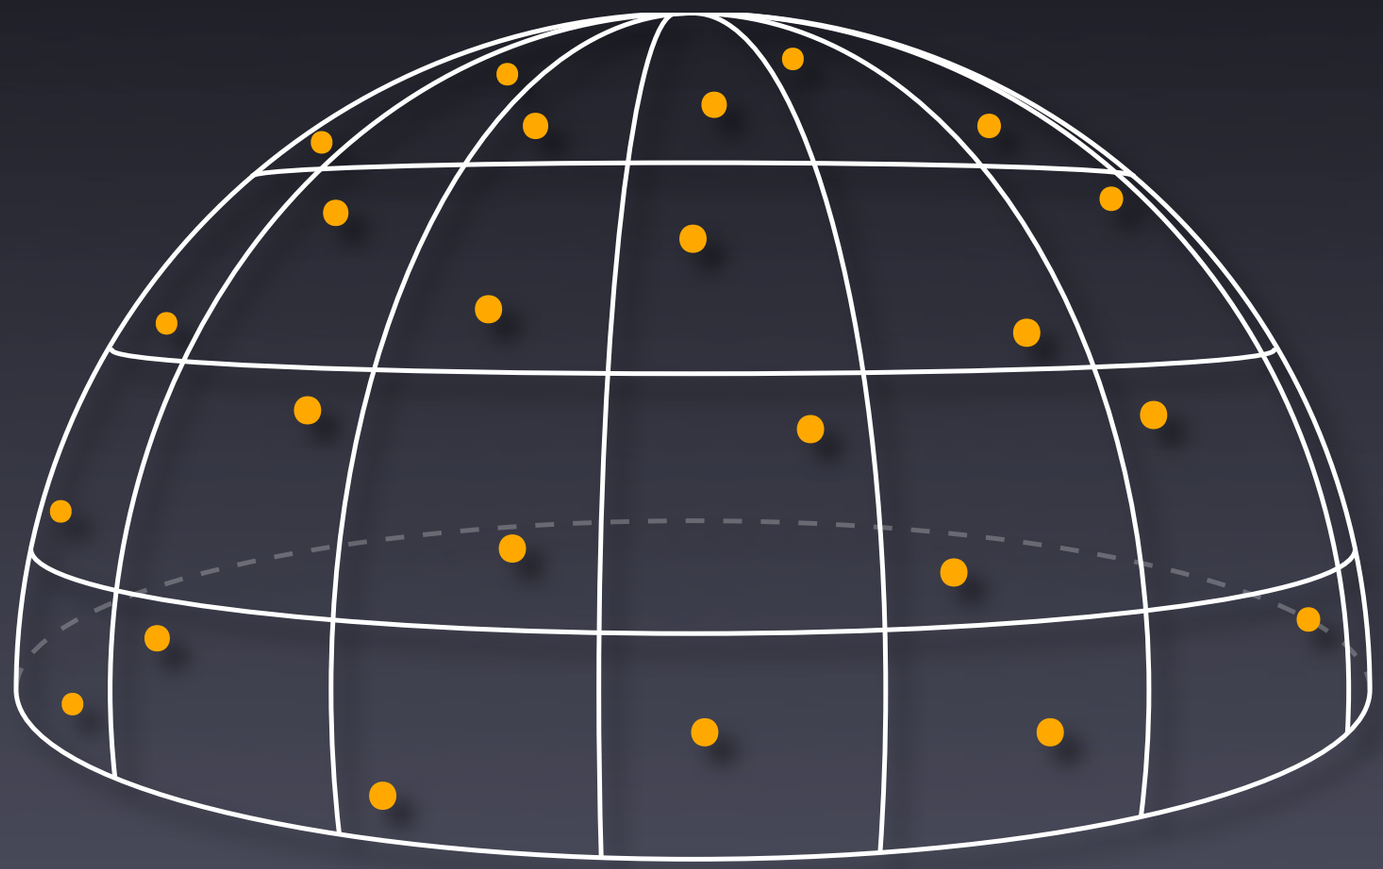

Computing Irradiance

Ward et al. '88

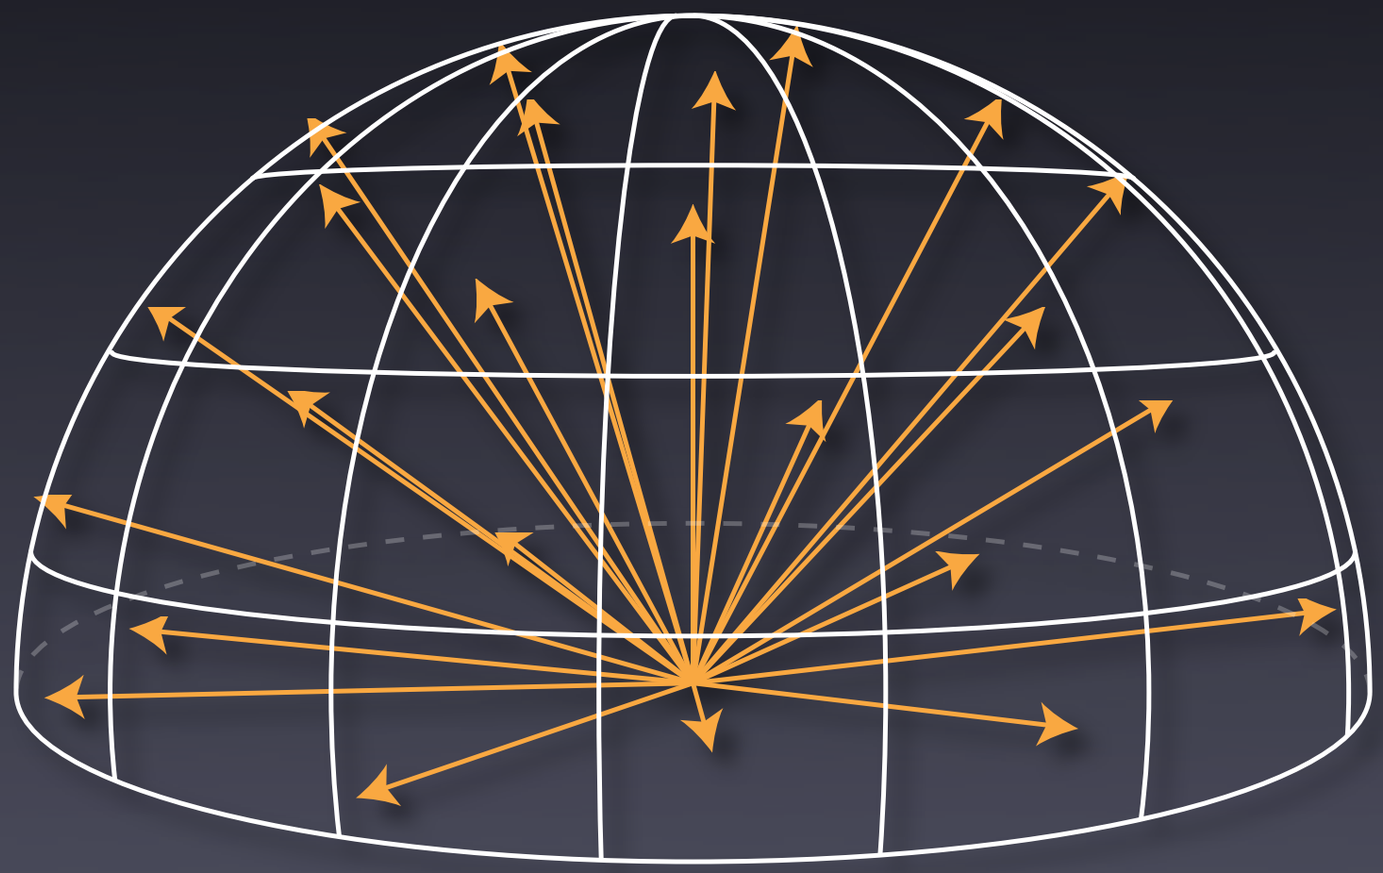

Computing Irradiance

Ward et al. '88

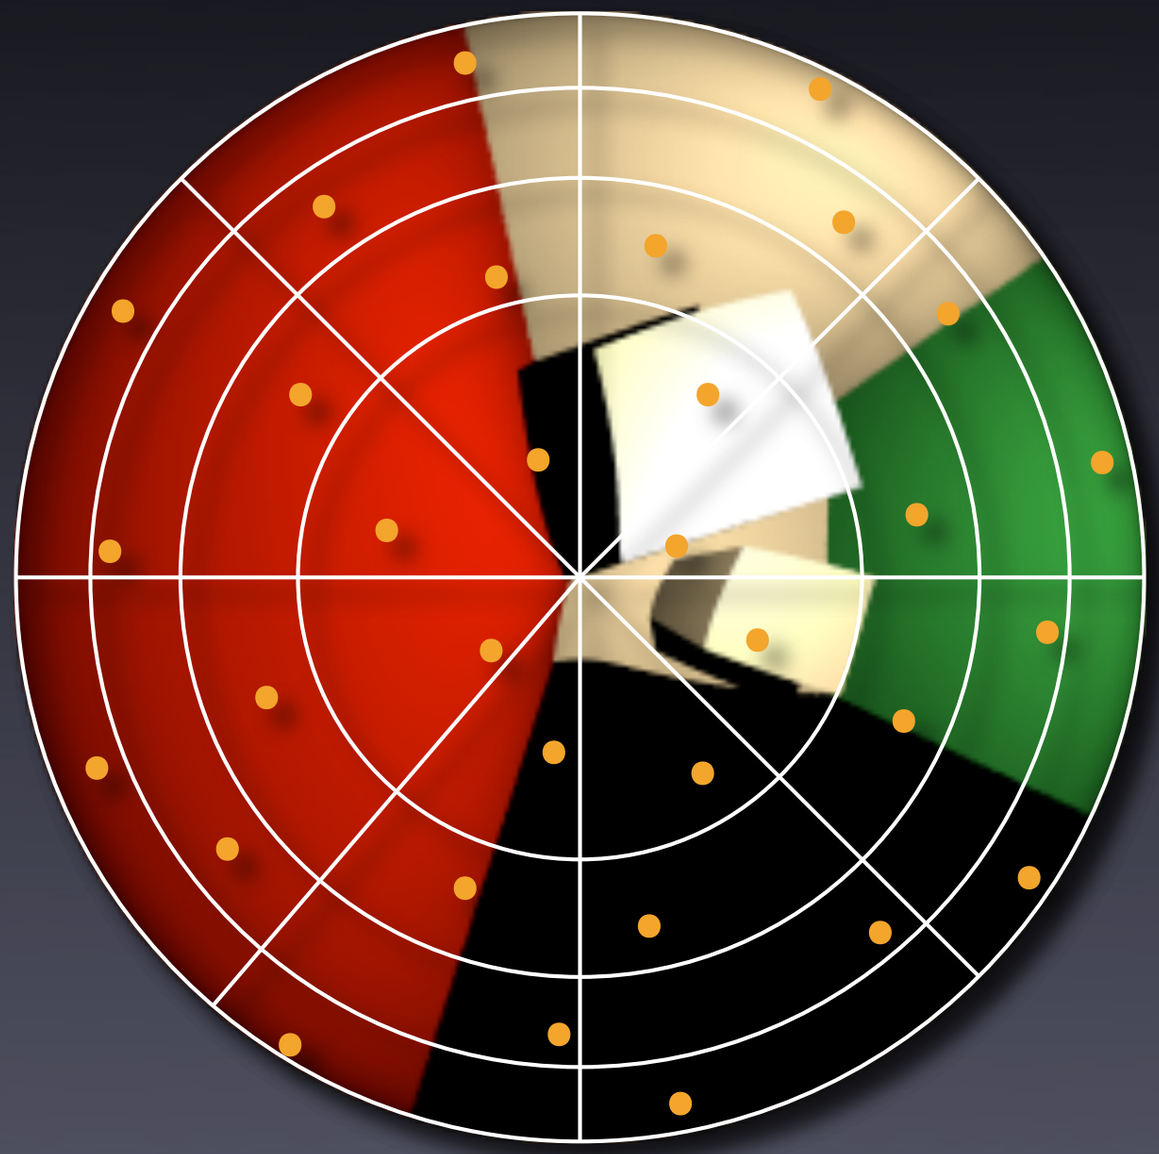

Computing Irradiance

Ward et al. '88

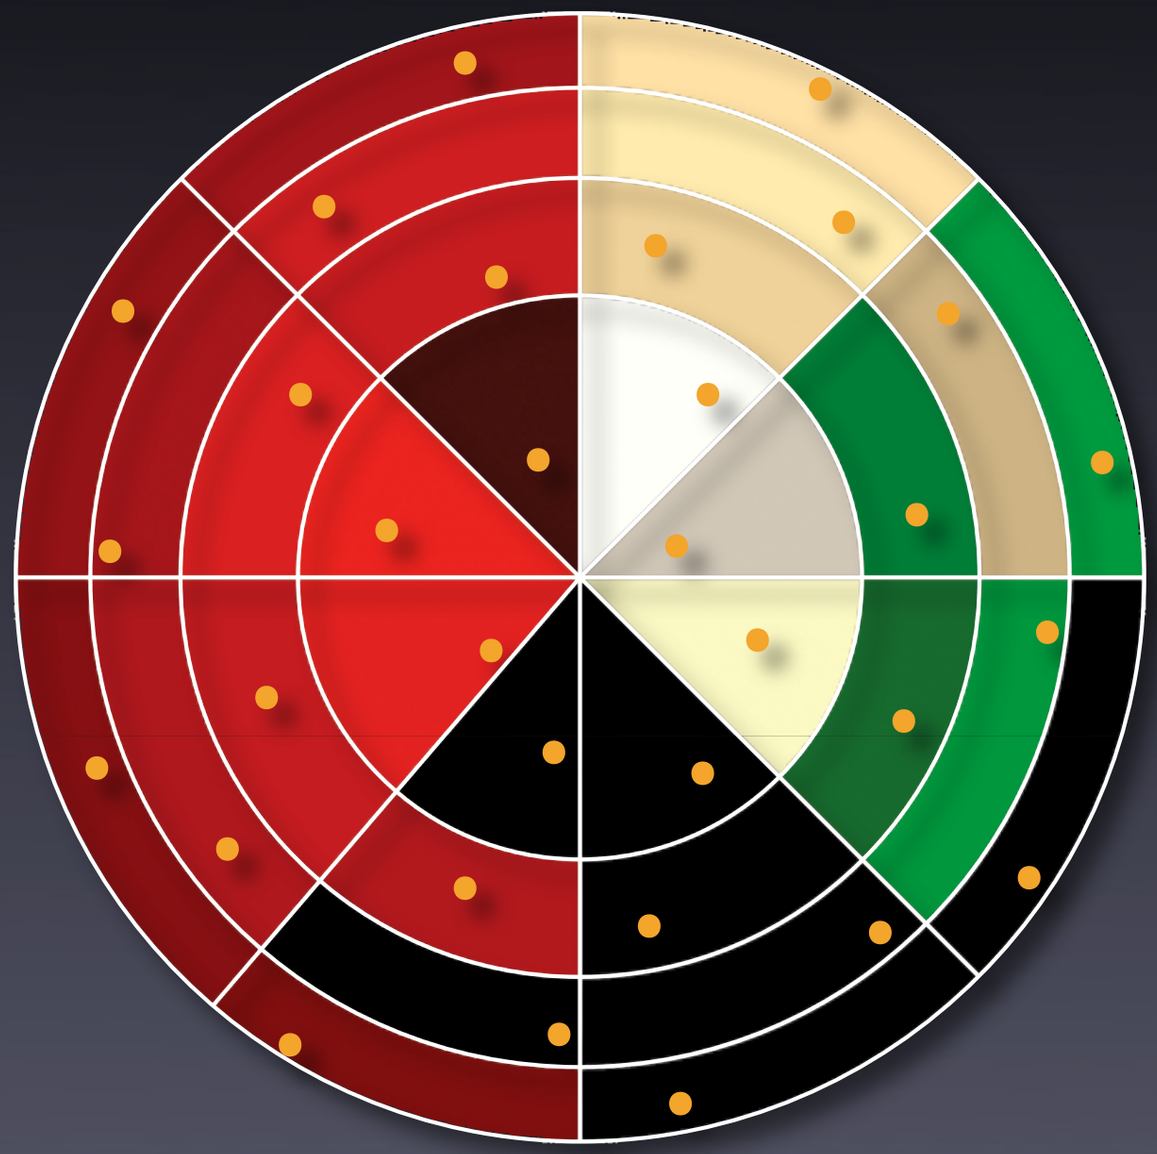

Computing Irradiance

Ward et al. '88

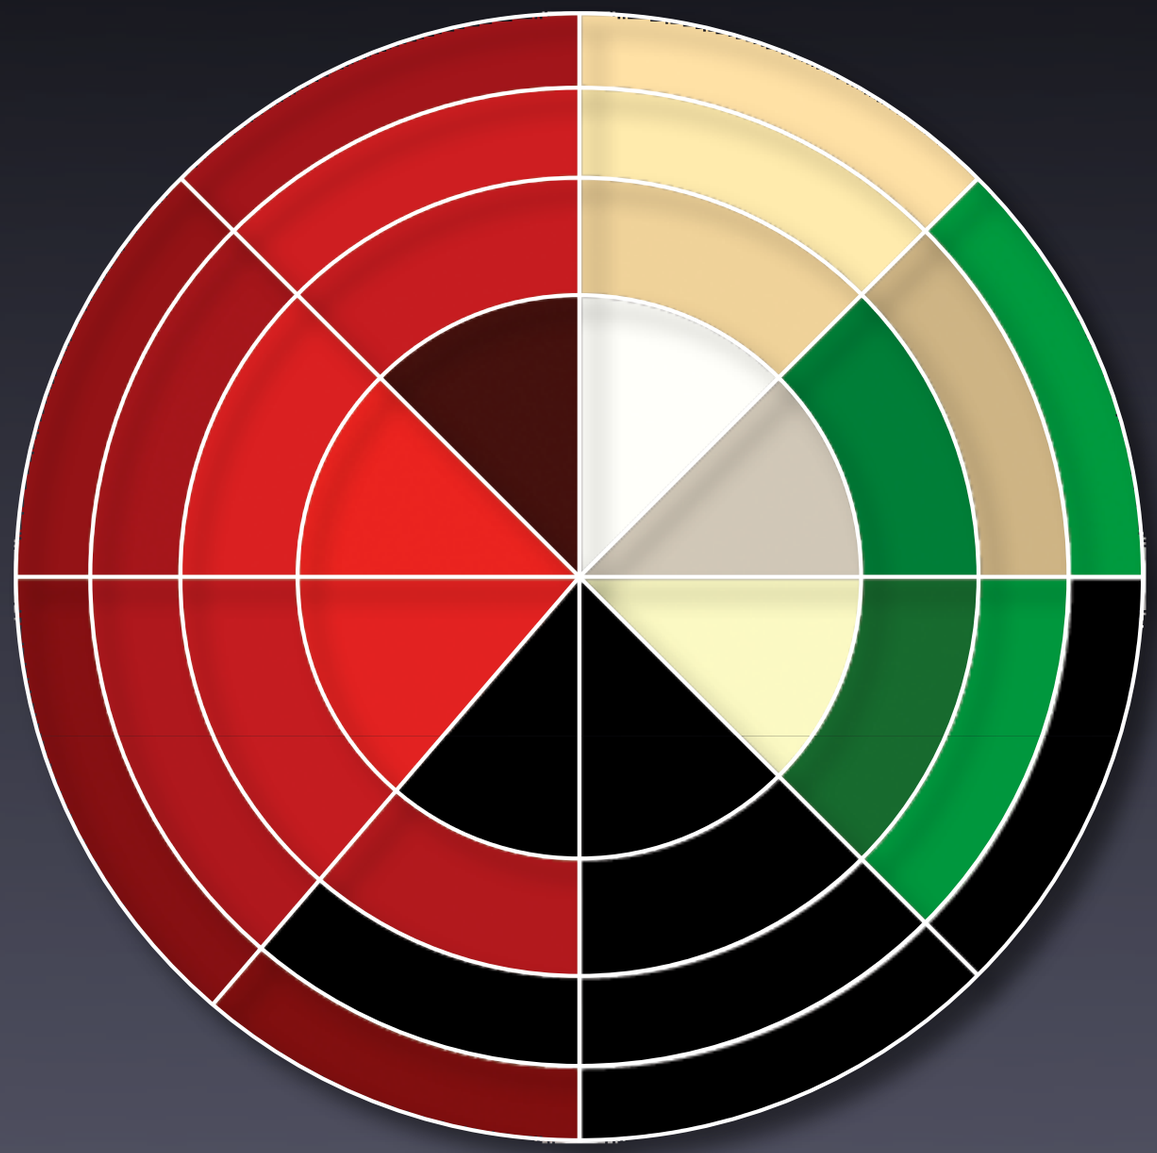

Computing Irradiance

Ward et al. '88

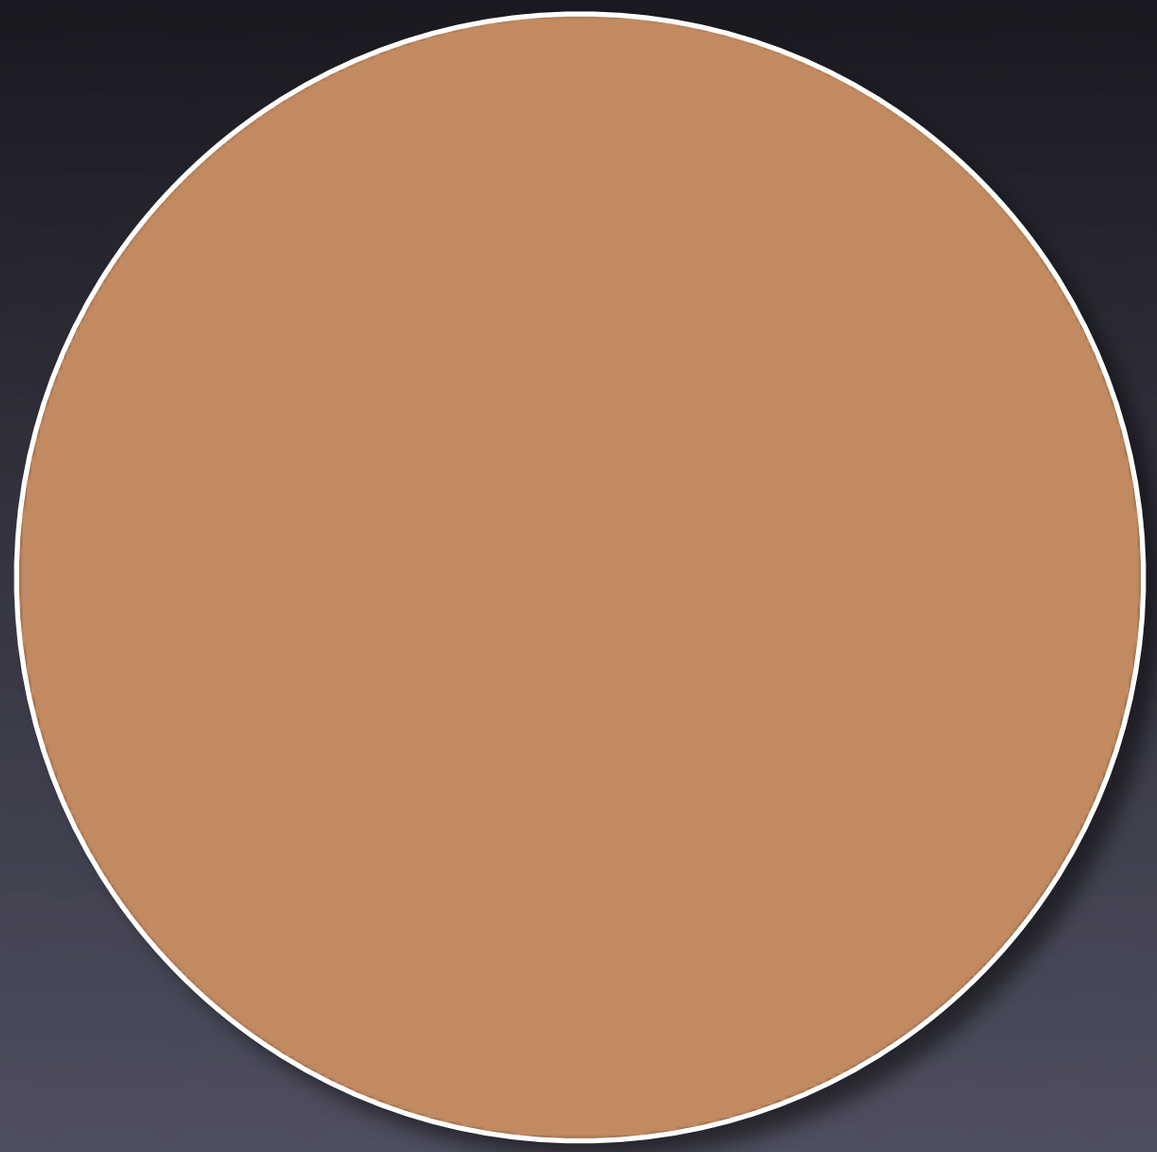

Computing Irradiance

Ward et al. '88

$$E(\mathbf{x}) = \int_{\Omega} L_i(\mathbf{x}, \vec{\omega}_i)(\vec{\omega}_i \cdot \vec{n})\, d\vec{\omega}_i$$

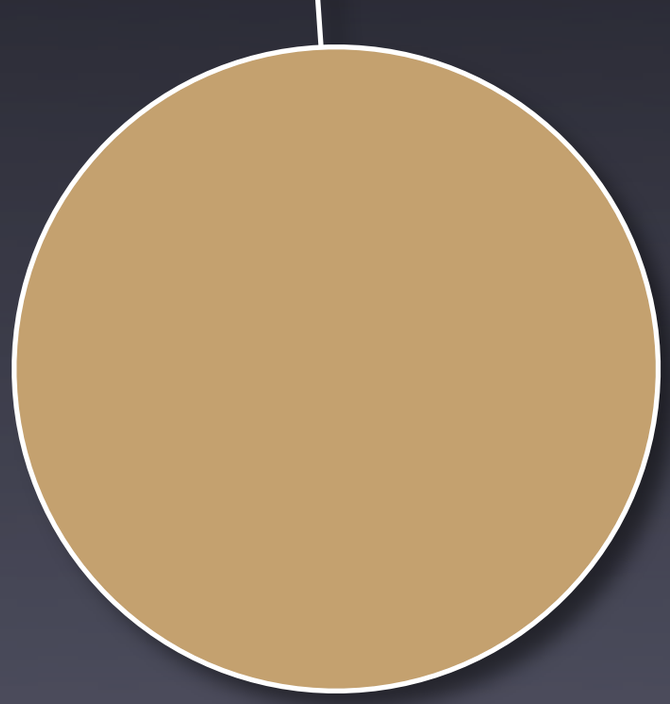

Computing Irradiance

Ward et al. '88

$$E(\mathbf{x}) = \int_{\Omega} L_i(\mathbf{x}, \vec{\omega}_i)(\vec{\omega}_i \cdot \vec{n})\, d\vec{\omega}_i$$

Computing Irradiance

Ward et al. '88

$$E(\mathbf{x}) = \int_{\Omega} L_i(\mathbf{x}, \vec{\omega}_i)(\vec{\omega}_i \cdot \vec{n})\, d\vec{\omega}_i$$

$$E(\mathbf{x}) \approx \frac{\pi}{N_1 N_2} \sum_{j=1}^{N_1} \sum_{i=1}^{N_2} L_i(\mathbf{x}, \theta_j, \phi_i),$$

where:

$$\theta_j = \sin^{-1}\left(\sqrt{\frac{j-\xi_1}{N_1}}\right), \qquad \phi_i = 2\pi\frac{i-\xi_2}{N_2}$$

Interpolating Irradiance

Ward et al. '88

- We wish to minimize the number of times we compute this costly irradiance integral.
- How far away can we reuse a cache point?

Interpolating Irradiance

Ward et al. '88

- Ward et al. derived an approximate formula for this error function.
- Error function is decomposed w.r.t. translation and orientation.
  - Translation: function of distance to cache point & average distance to surface visible at cache point.
  - Orientation: function of rotational offset

Interpolating Irradiance

Ward et al. '88

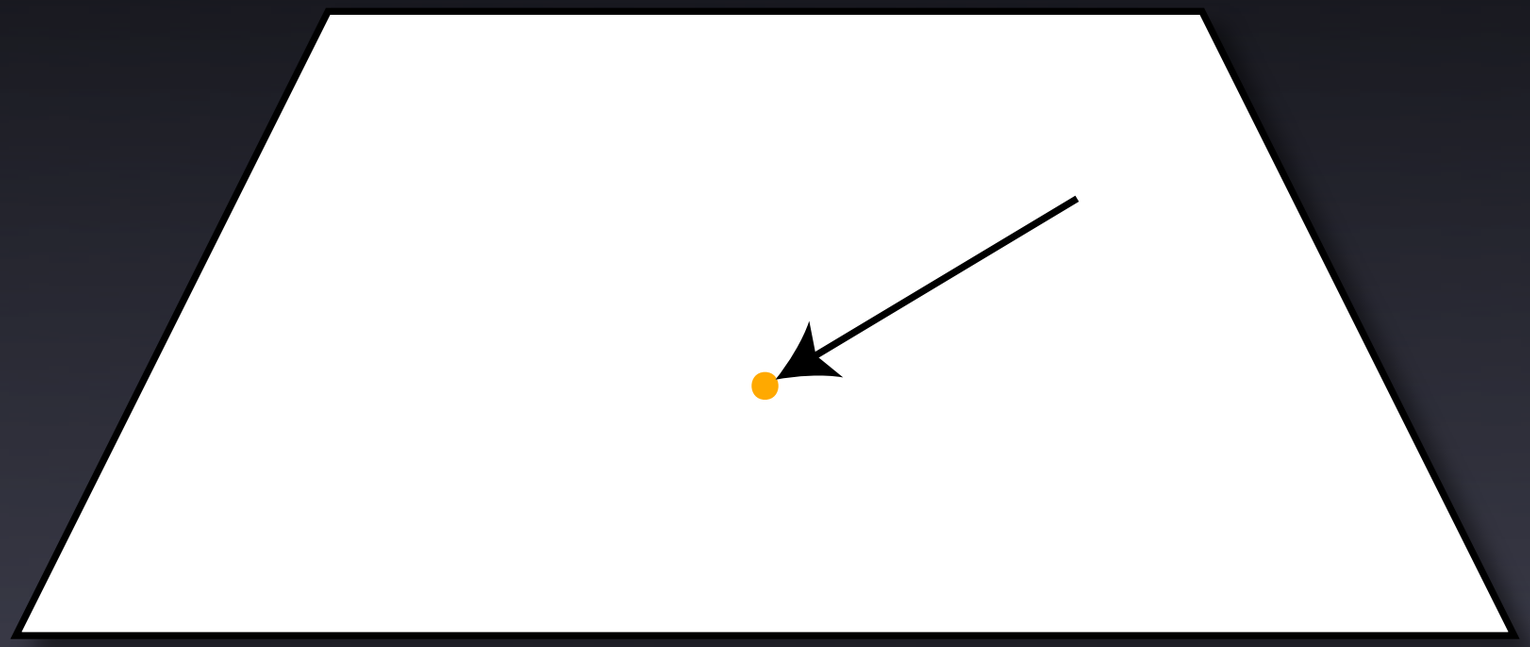

Interpolating Irradiance

Ward et al. '88

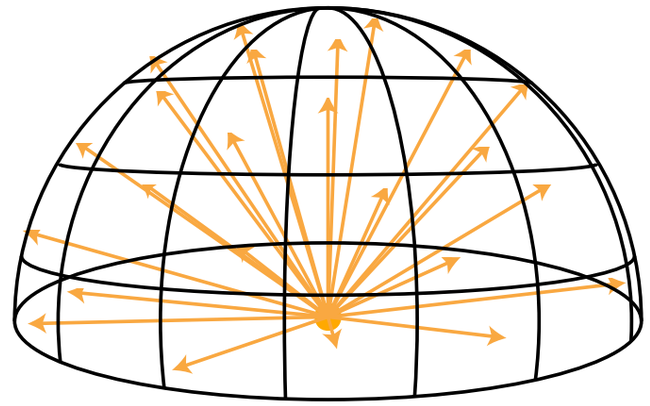

Interpolating Irradiance

Ward et al. '88

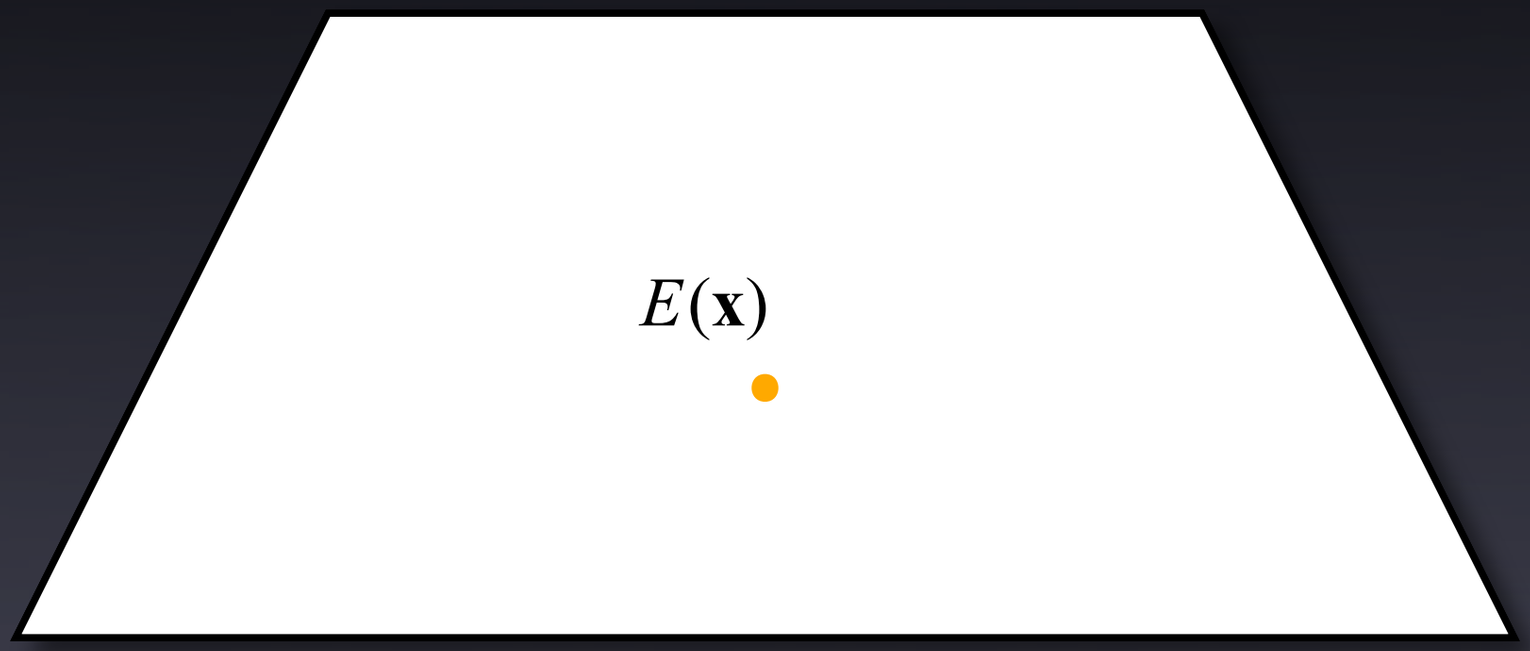

Interpolating Irradiance

Ward et al. '88

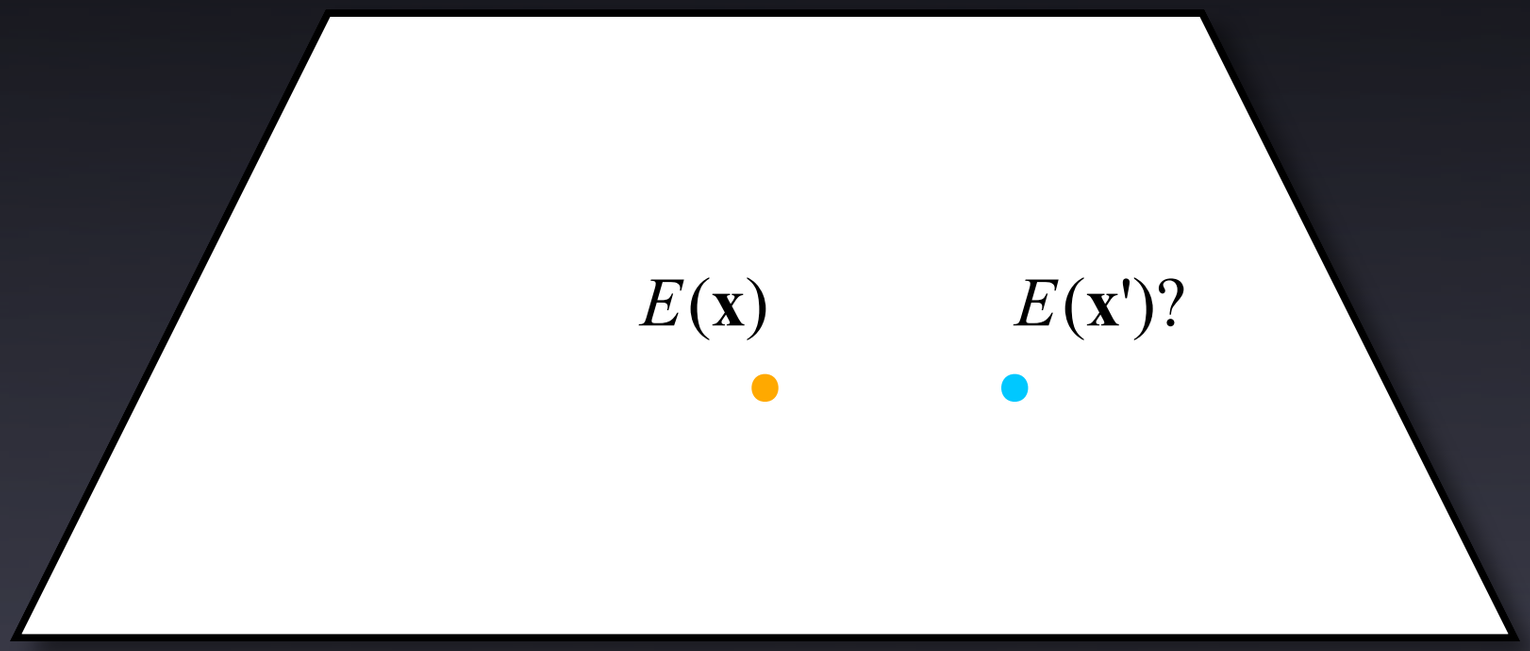

What happens if we re-use E(x) at E(x')?

Interpolating Irradiance

Ward et al. '88

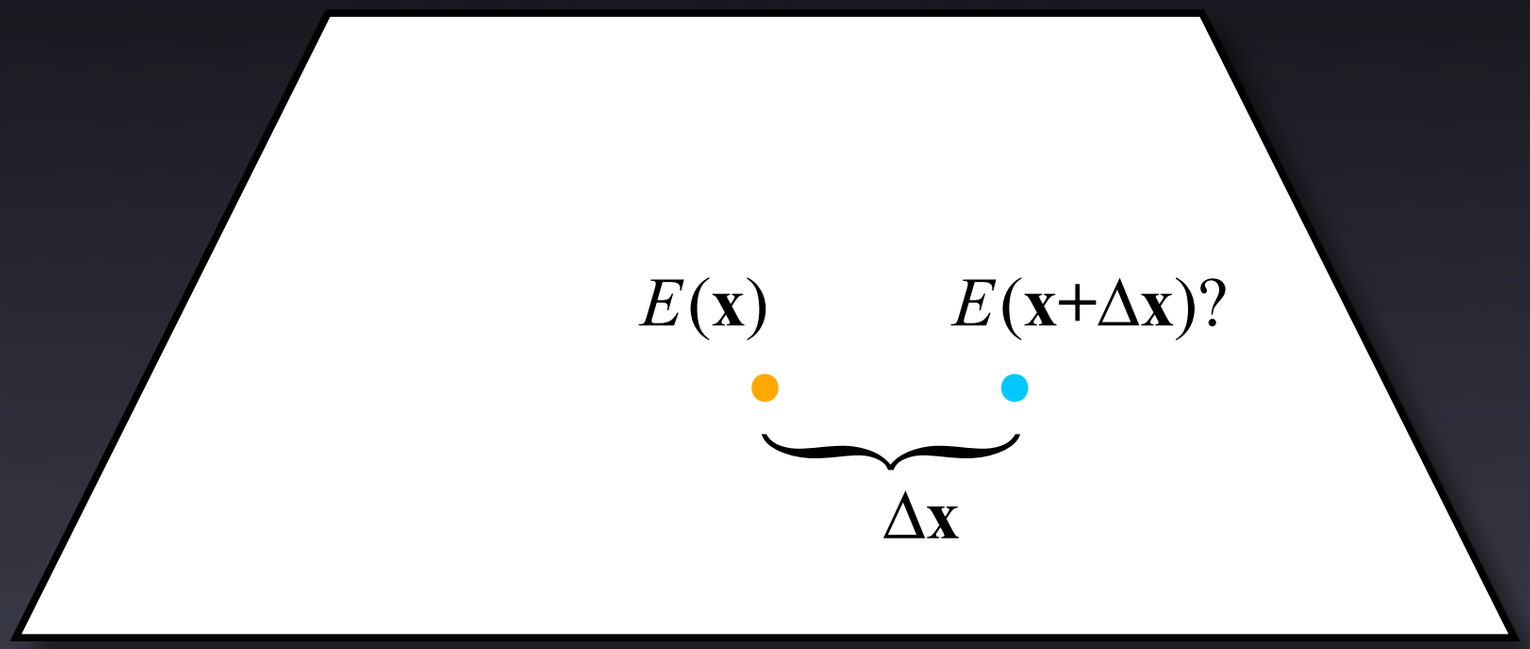

What happens if we re-use E(x) at E(x')?

Interpolating Irradiance

Ward et al. '88

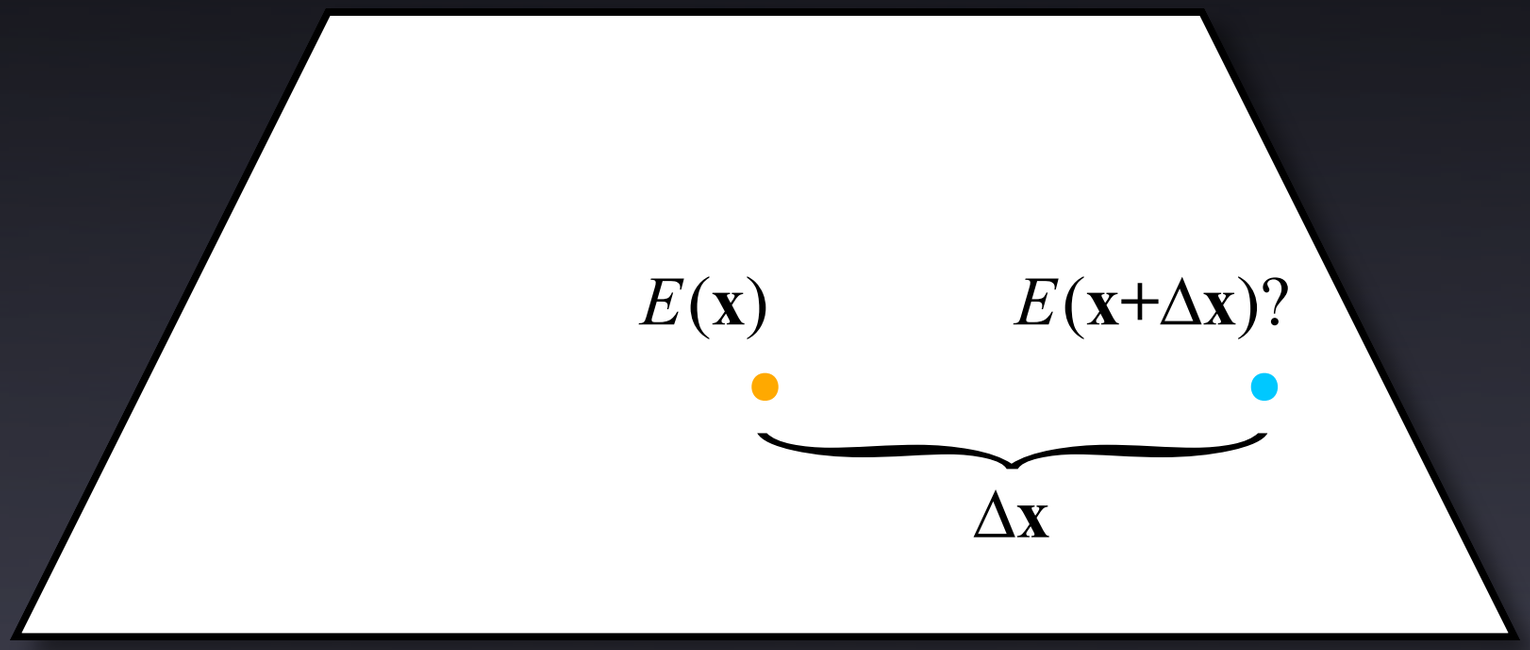

What happens if we re-use E(x) at E(x')?

Interpolating Irradiance

Ward et al. '88

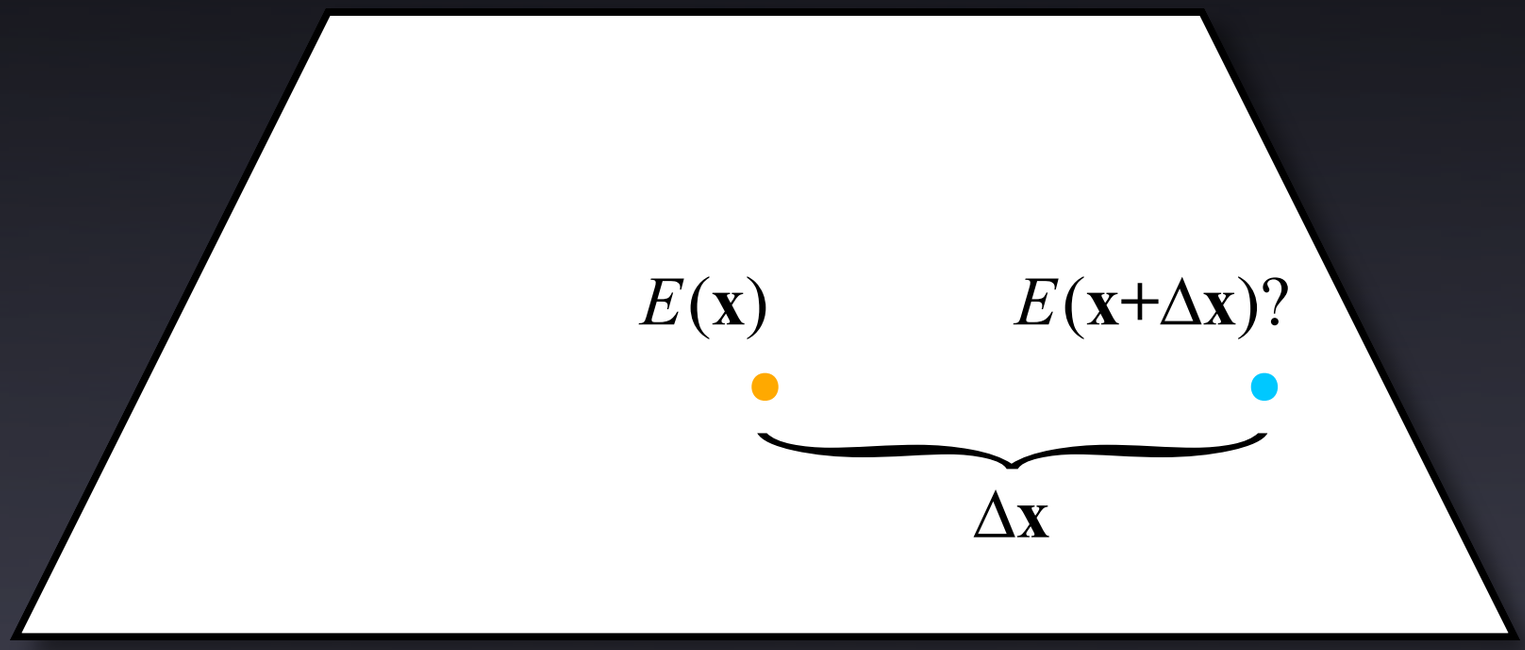

What happens if we re-use E(x) at E(x')?

$$E(x') \approx E(x) + \left(\frac{\partial E}{\partial x}(\Delta x) + \frac{\partial E}{\partial \varphi}(\Delta \varphi) \right)$$

Interpolating Irradiance

Ward et al. '88

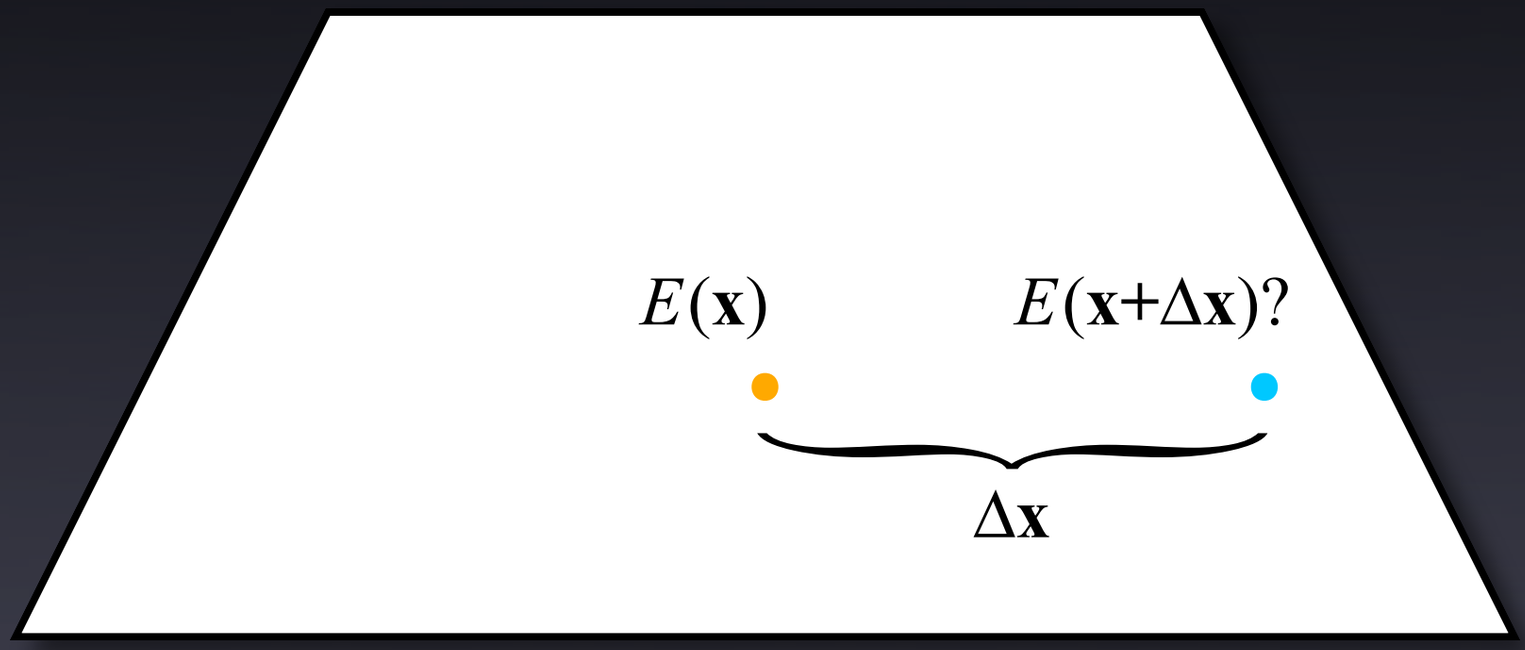

What happens if we re-use E(x) at E(x')?

error

$$E(x') \approx E(x) + \left(\frac{\partial E}{\partial x}(\Delta x) + \frac{\partial E}{\partial \varphi}(\Delta \varphi) \right)$$

Interpolating Irradiance

Ward et al. '88

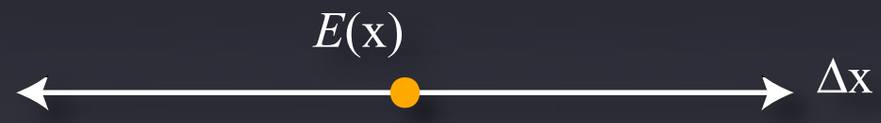

Interpolating Irradiance

Ward et al. '88

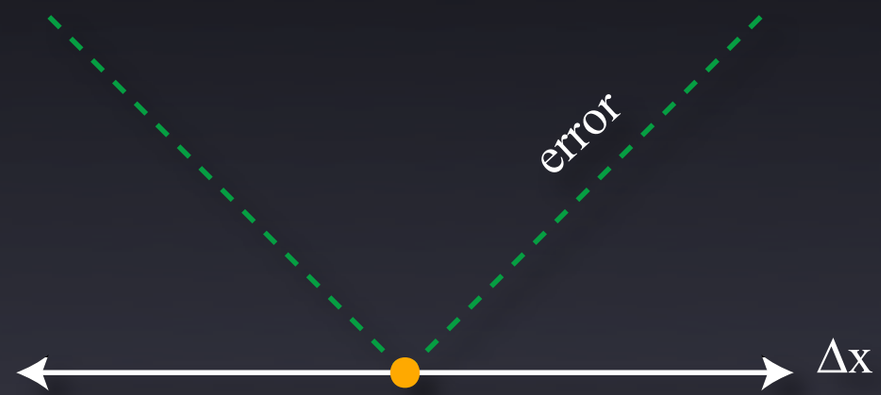

Interpolating Irradiance

Ward et al. '88

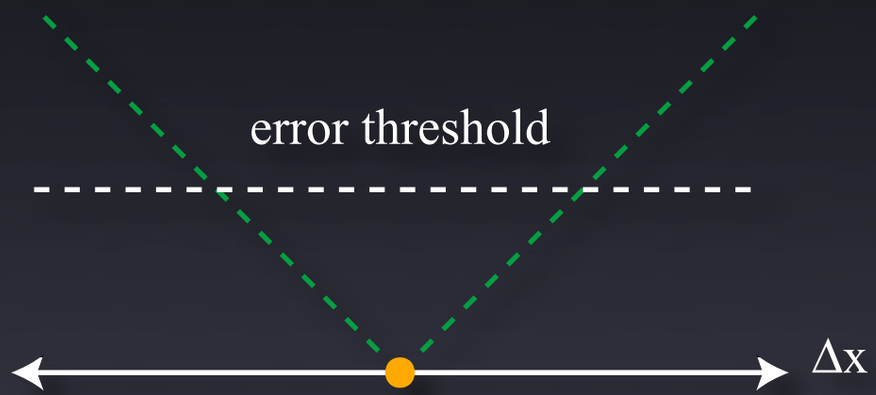

Interpolating Irradiance

Ward et al. '88

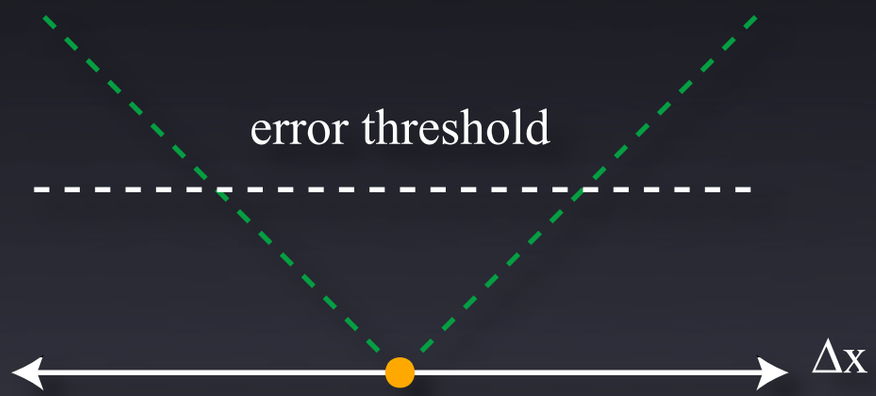

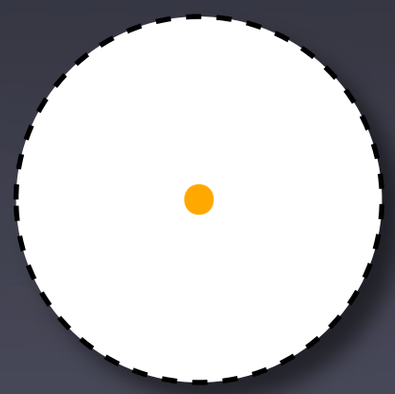

Interpolating Irradiance

Ward et al. '88

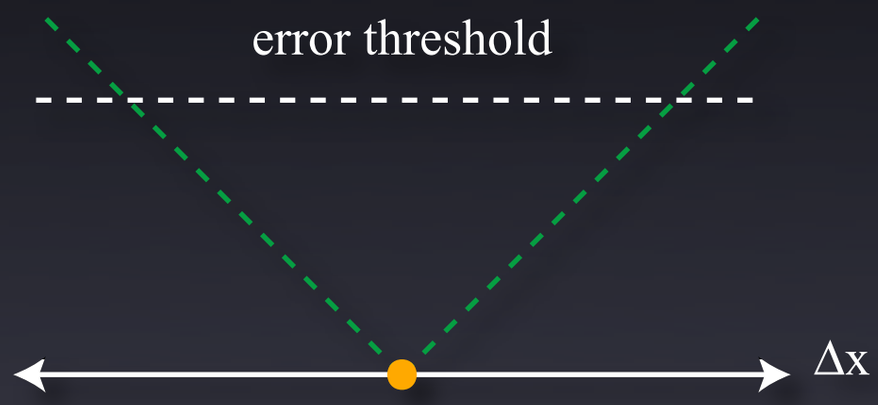

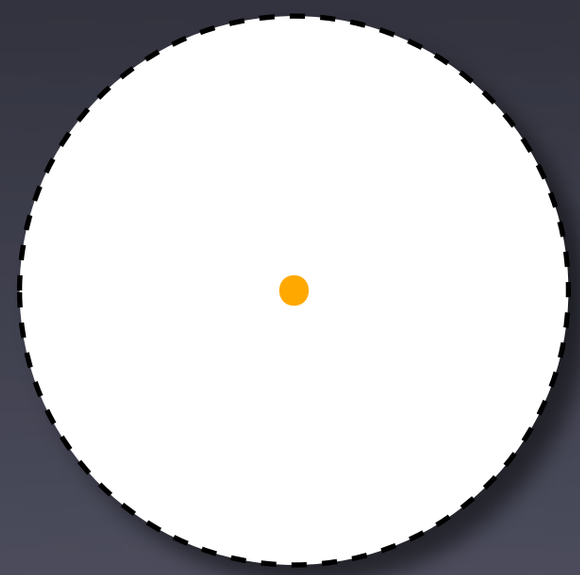

Interpolating Irradiance

Ward et al. '88

- Approximate using Taylor expansion:

$$\varepsilon \gtrsim \left| \frac{\partial E}{\partial x}(x - x_0) + \frac{\partial E}{\partial \varphi}(\varphi - \varphi_0) \right|$$

The “Split-Sphere”

Ward et al. ‘88

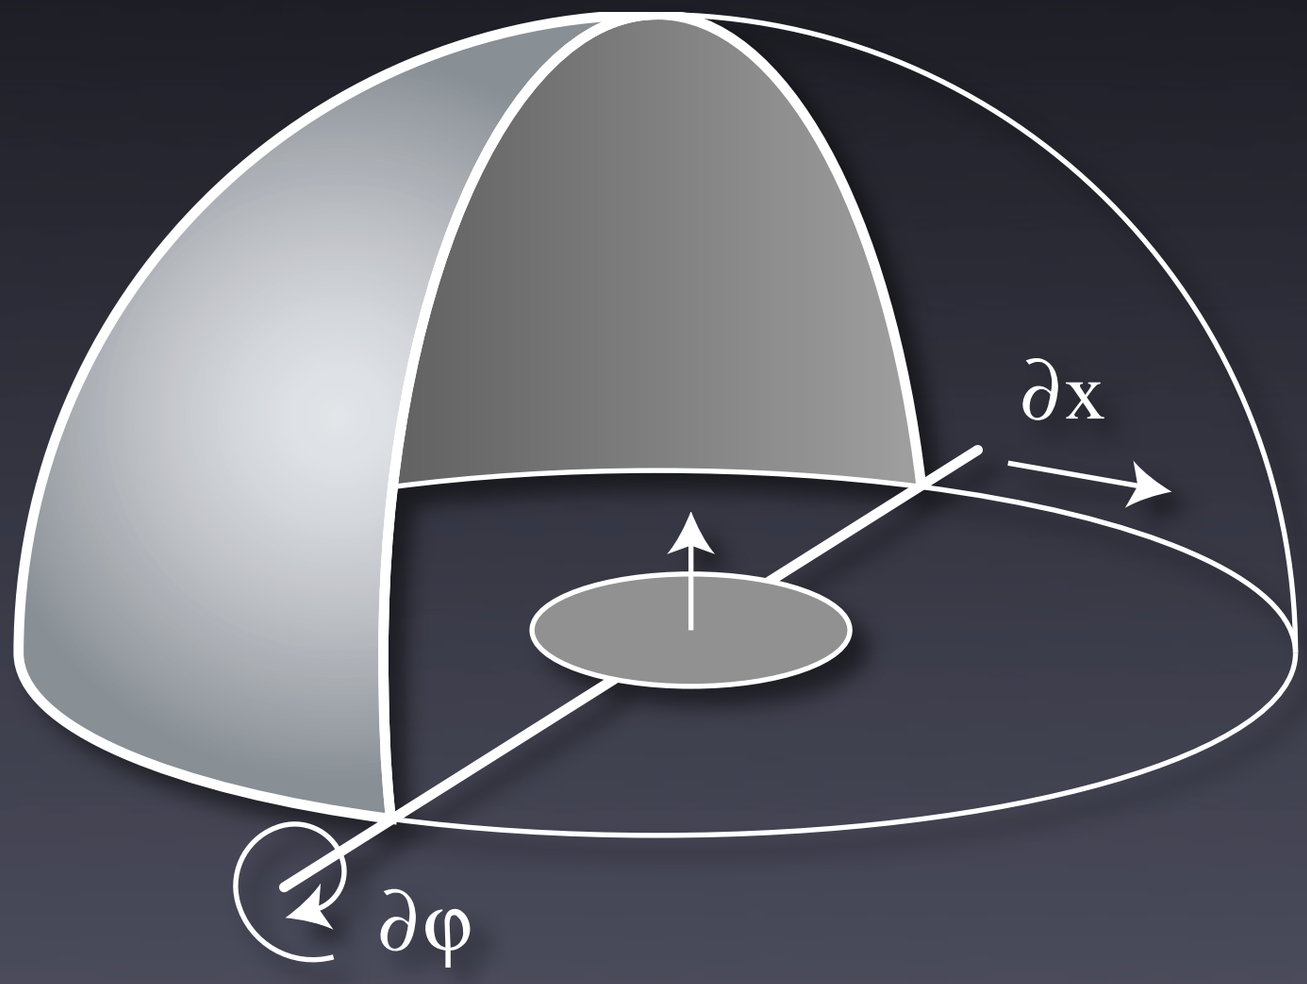

Interpolating Irradiance

Ward et al. '88

- In the "Split-Sphere" environment the error becomes:

$$\varepsilon \lesssim \left| \frac{\partial E}{\partial x}(x - x_0) + \frac{\partial E}{\partial \varphi}(\varphi - \varphi_0) \right|$$

Interpolating Irradiance

Ward et al. '88

- In the "Split-Sphere" environment the error becomes:

$$\varepsilon \lesssim E_0 \left(\frac{4}{\pi} \frac{|x - x_0|}{R} + |\varphi - \varphi_0| \right)$$

Interpolating Irradiance

Ward et al. '88

- In the "Split-Sphere" environment the error becomes:

$$\varepsilon \lesssim E_0 \left(\frac{4}{\pi} \frac{|x - x_0|}{R} + |\varphi - \varphi_0| \right)$$

"average" distance to visible surfaces in hemisphere

Interpolating Irradiance

Ward et al. '88

- At each shading location, perform a weighted average of all cached values which have an error below some threshold.
- Reciprocal of the error is used as the weight.

Interpolating Irradiance

Ward et al. '88

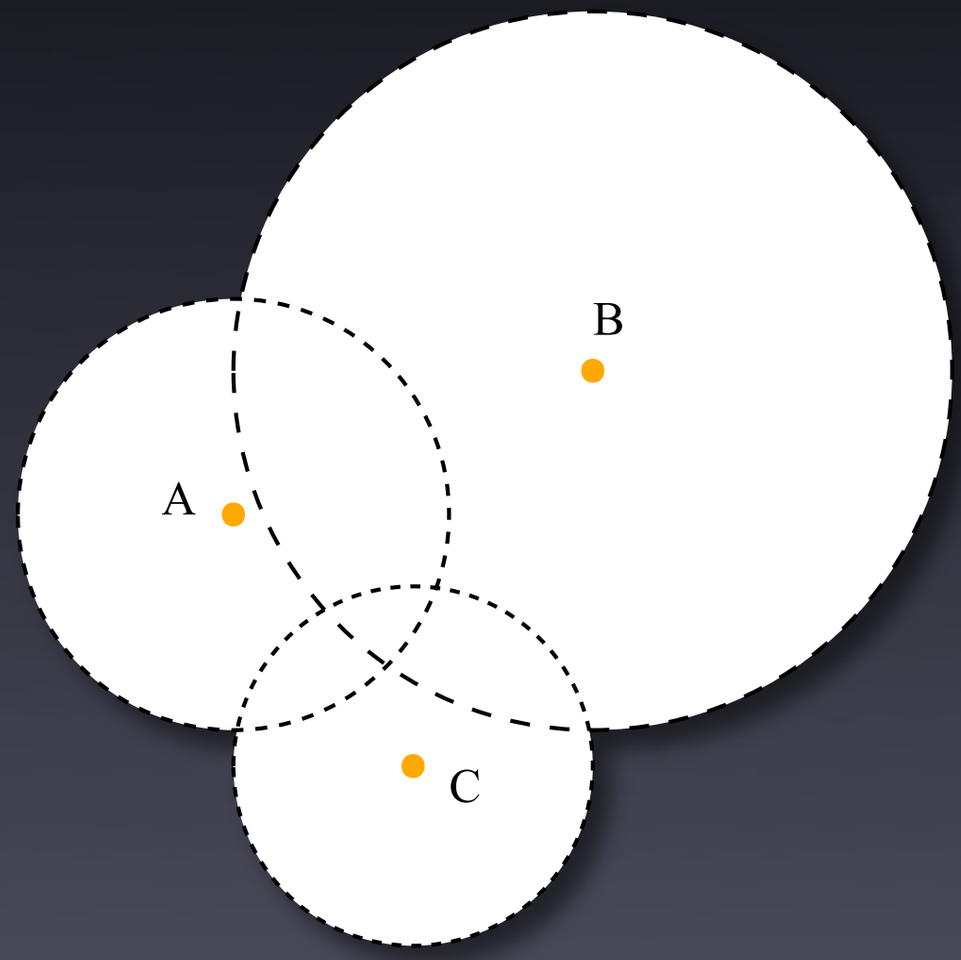

Interpolating Irradiance

Ward et al. '88

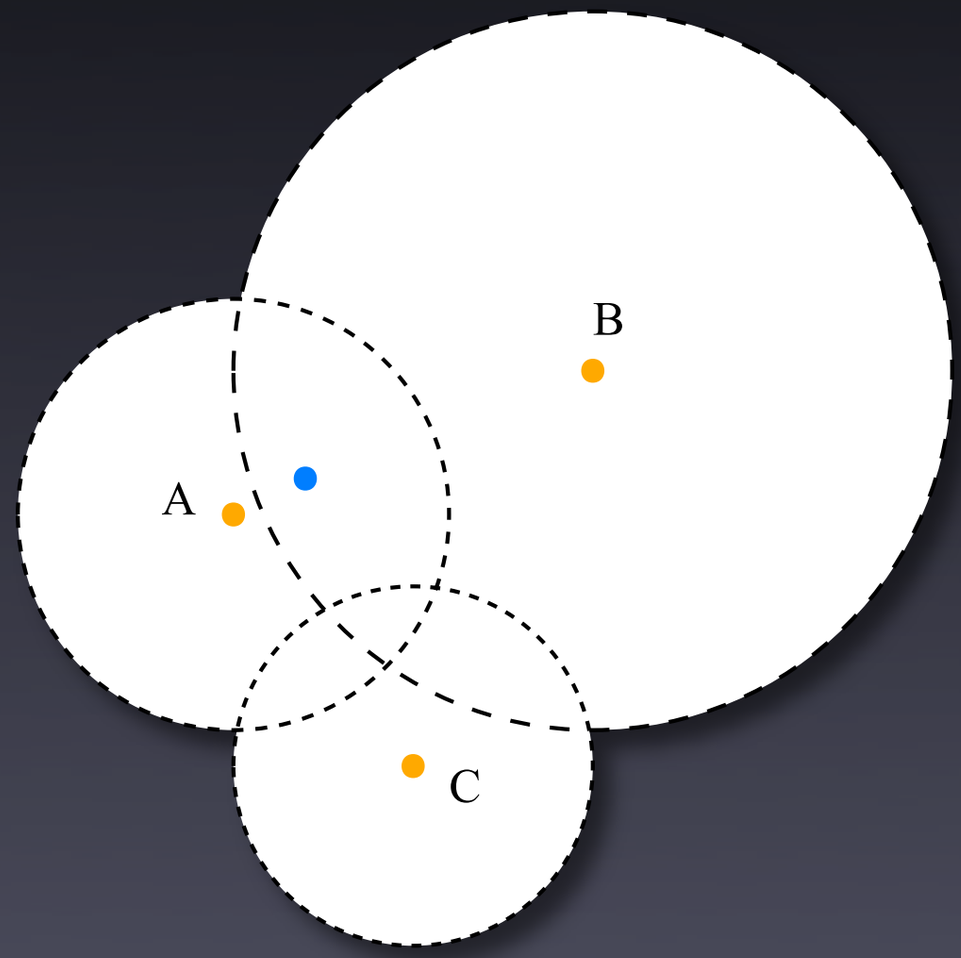

Interpolating Irradiance

Ward et al. '88

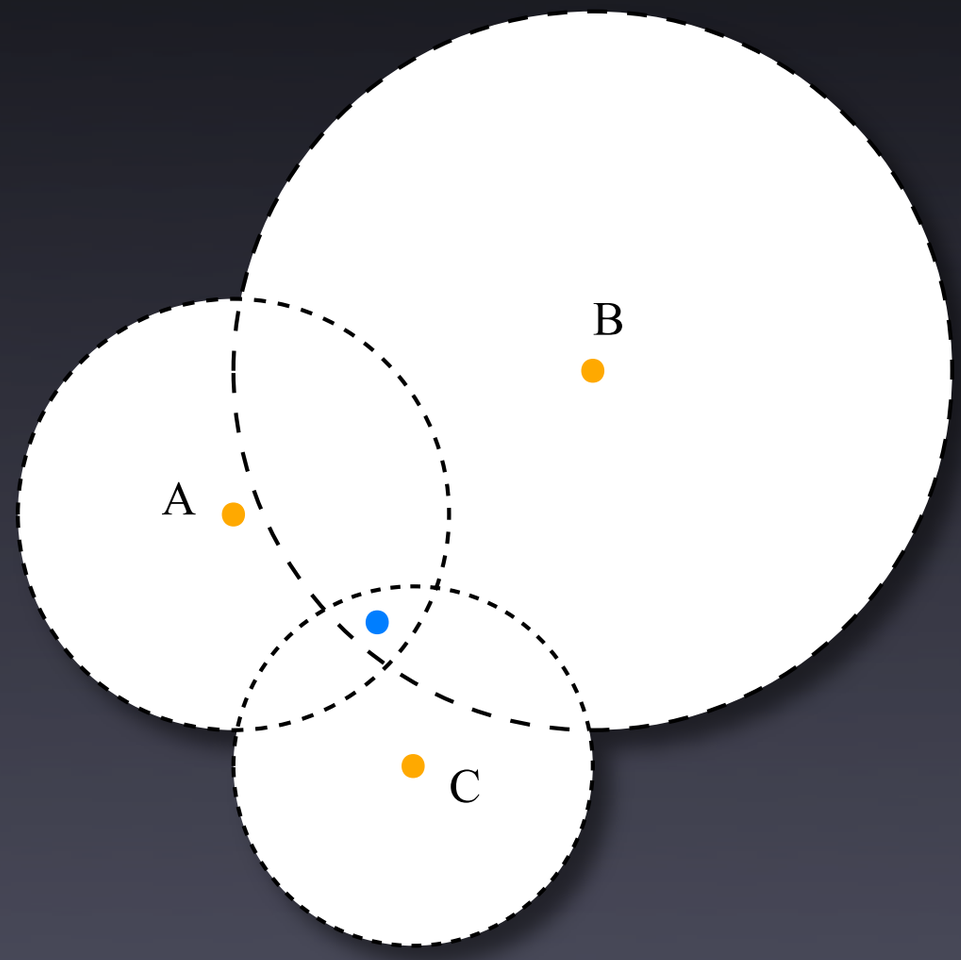

Interpolating Irradiance

Ward et al. '88

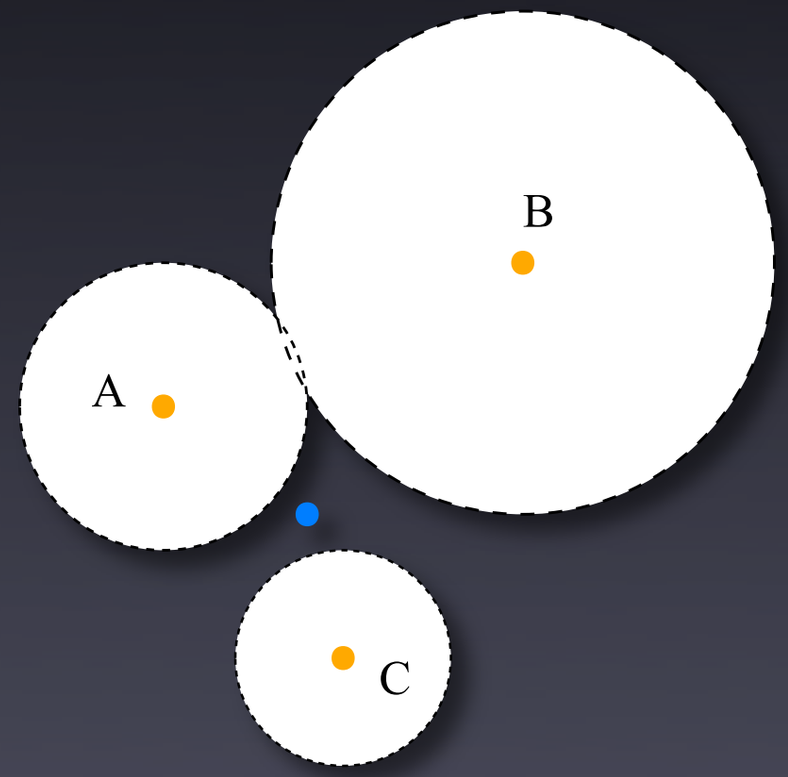

Interpolating Irradiance

Ward et al. '88

$$E(\mathbf{x}, \vec{n}) \approx \frac{\sum_{i \in S} \mathrm{w}_i(\mathbf{x}, \vec{n})\, E_i}{\sum_{i \in S} \mathrm{w}_i(\mathbf{x}, \vec{n})}$$

Interpolating Irradiance

Ward et al. '88

$$E(\mathbf{x},\vec{n}) \approx \frac{\sum_{i\in S} \mathrm{w}_i(\mathbf{x},\vec{n})\, E_i}{\sum_{i\in S} \mathrm{w}_i(\mathbf{x},\vec{n})}$$

where:

$$\mathrm{w}_i(\mathbf{x},\vec{n}) = \frac{1}{\frac{\|\mathbf{x}-\mathbf{x}_i\|}{R_i} + \sqrt{1-\vec{n}\cdot\vec{n}_i}}$$

$$S = \left\{ i : \mathrm{w}_i(\mathbf{x},\vec{n}) > \frac{1}{a} \right\}$$

Irradiance Caching

Ward et al. '88

- Pros:
  - Independent of resolution.
  - Concentrates computation in visible regions.
  - Hundreds of times faster than naive path tracing.

Irradiance Caching

Ward et al. '88

- Cons:
  - Interpolation/extrapolation can introduce visible artifacts.
  - Limited to Lambertian (matte) surfaces.
  - Still slow.

Improvements/Extensions

- Ward and Heckbert ‘92 - better interpolation
- Křivánek et al. ‘05a, ‘05b - glossy surfaces
- Tabellion and Lamorlette ‘04 - speed
- Tawara et al. ‘04 - animation
- Yee ‘00 - speed/perception
- Kato ‘02 - parallel/distributed computation
- Arikan et al. ‘05 - speed

Improvements/Extensions

- Ward and Heckbert ‘92
- Křivánek et al. ‘05a, ‘05b
- Tabellion and Lamorlette ‘04
- Tawara et al. ‘04

Improvements/Extensions

- Ward and Heckbert ‘92
- Křivánek et al. ‘05a, ‘05b
- Tabellion and Lamorlette ‘04
- Tawara et al. ‘04

Irradiance Gradients

Ward and Heckbert '92

- Improve quality by performing higher-order interpolation/extrapolation.
- Estimate a derivative to the irradiance.
- Apply this derivative to the weighted average.

Extrapolating Irradiance

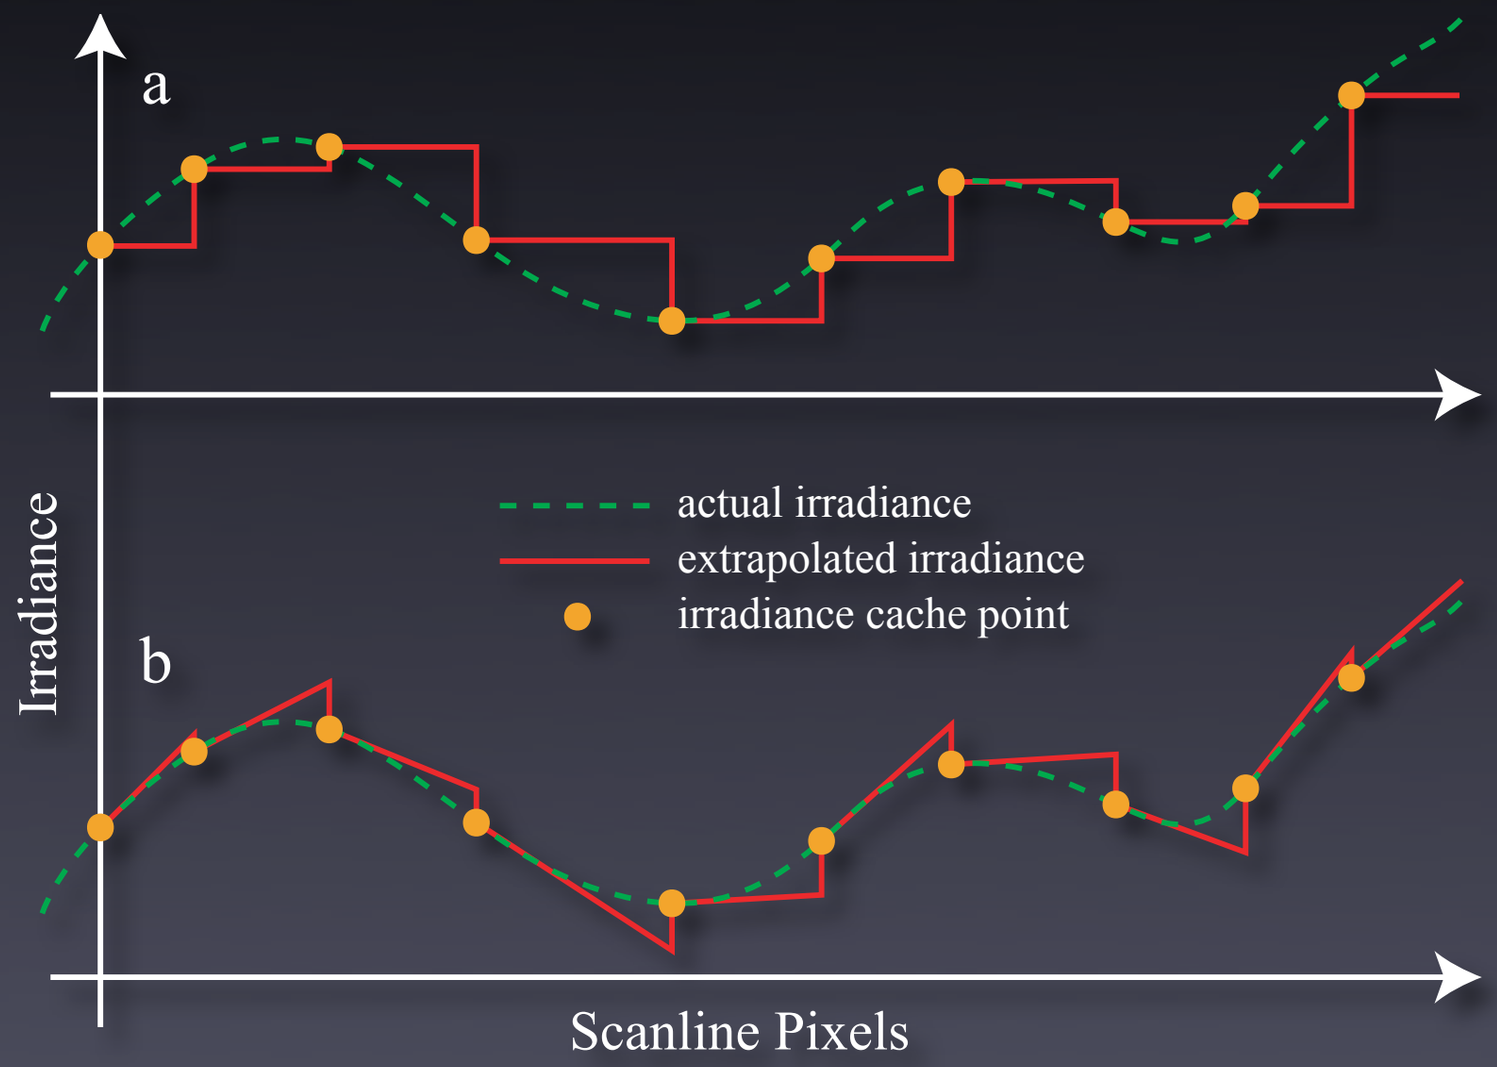

Irradiance Gradients

Ward and Heckbert '92

- Decomposed into:
  - Rotational Gradient - captures change in irradiance w.r.t. change in surface orientation
  - Translational Gradient - captures change in irradiance w.r.t. change in position

Irradiance Gradients

Ward and Heckbert '92

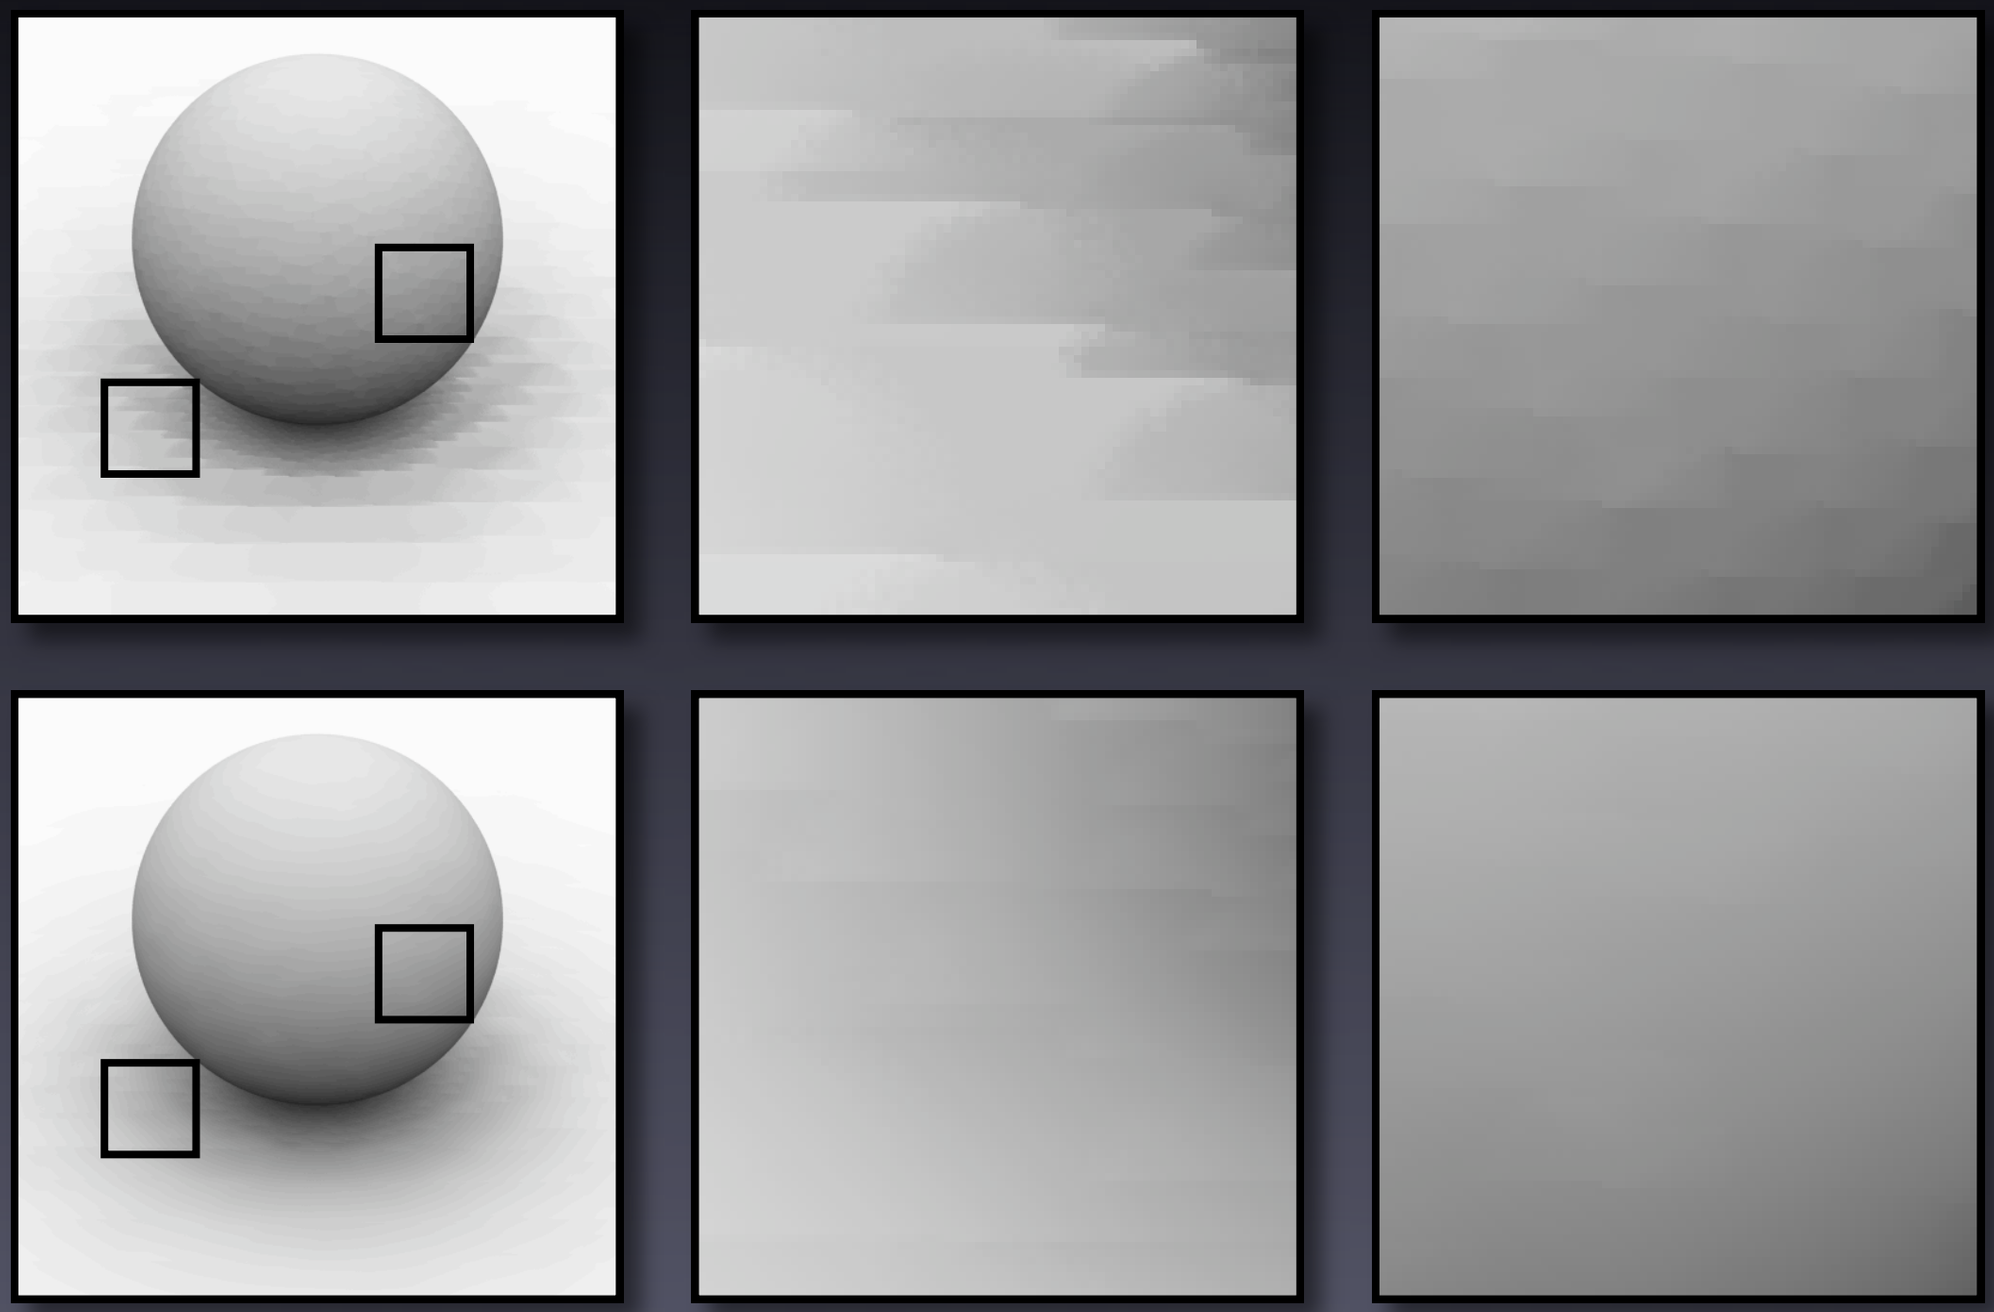

Applying Gradients

Ward and Heckbert '92

$$E(\mathbf{x},\vec{n}) \approx \frac{\sum_{i \in S} \mathrm{w}_i(\mathbf{x},\vec{n})\, E_i}{\sum_{i \in S} \mathrm{w}_i(\mathbf{x},\vec{n})}$$

Applying Gradients

Ward and Heckbert '92

$$E(\mathbf{x},\vec{n}) \approx \frac{\sum_{i\in S} \mathrm{w}_i(\mathbf{x},\vec{n})\left(E_i+(\vec{n}_i\times\vec{n})\cdot\vec{\nabla}_r E_i+(\mathbf{x}-\mathbf{x}_i)\cdot\vec{\nabla}_t E_i\right)}{\sum_{i\in S} \mathrm{w}_i(\mathbf{x},\vec{n})}$$

Applying Gradients

Ward and Heckbert '92

estimated change due to rotation and translation

$$E(\mathbf{x},\vec{n}) \approx \frac{\sum_{i\in S} \mathrm{w}_i(\mathbf{x},\vec{n})\left(E_i + (\vec{n}_i \times \vec{n})\cdot\vec{\nabla}_r E_i + (\mathbf{x}-\mathbf{x}_i)\cdot\vec{\nabla}_t E_i\right)}{\sum_{i\in S} \mathrm{w}_i(\mathbf{x},\vec{n})}$$

Improvements/Extensions

- Ward and Heckbert ‘92
- Křivánek et al. ‘05a, ‘05b
- Tabellion and Lamorlette ‘04
- Tawara et al. ‘04

Radiance Caching

Křivánek et al. '05a, '05b

- Extends irradiance caching to glossy surfaces.

Glossy Surfaces

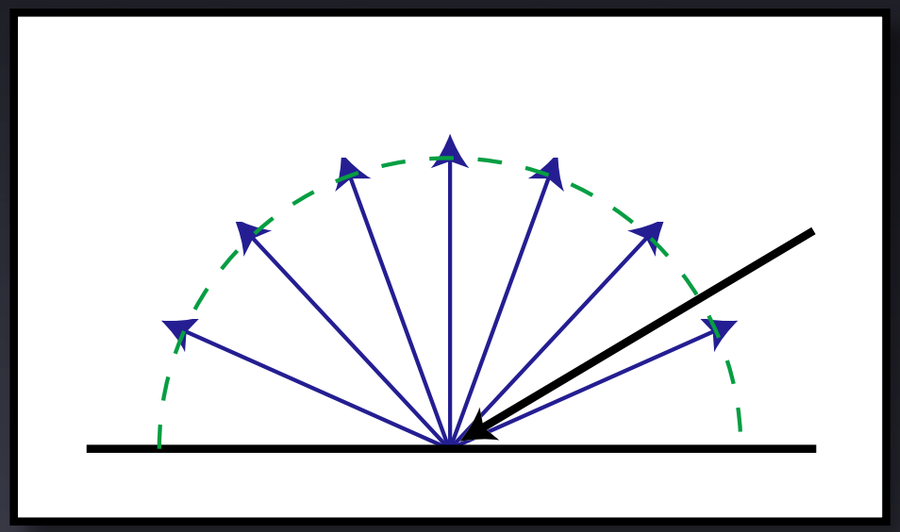

Matte surface

Glossy Surfaces

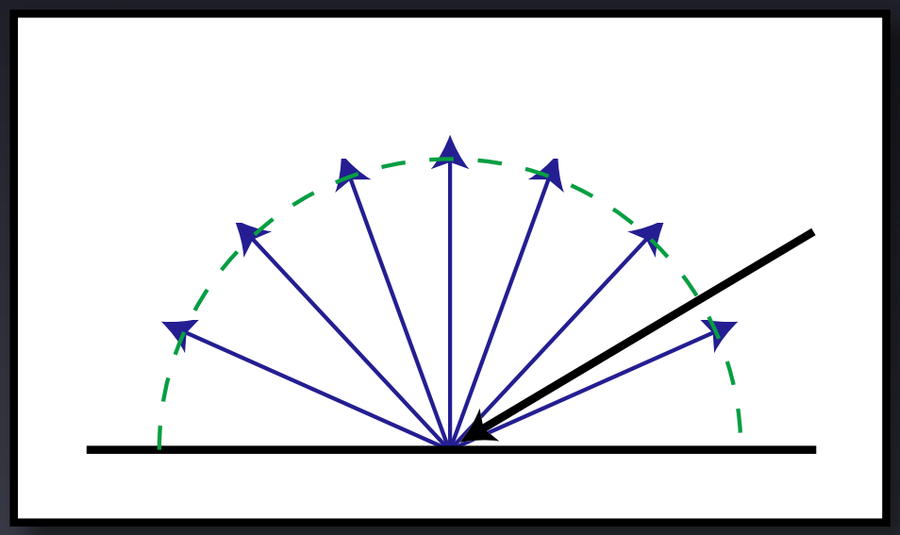

Matte surface

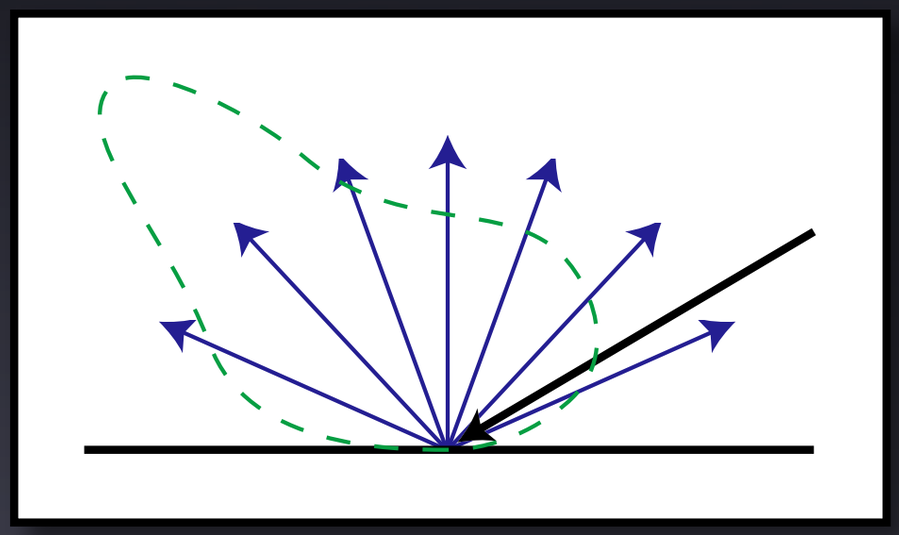

Glossy surface

Glossy Surfaces

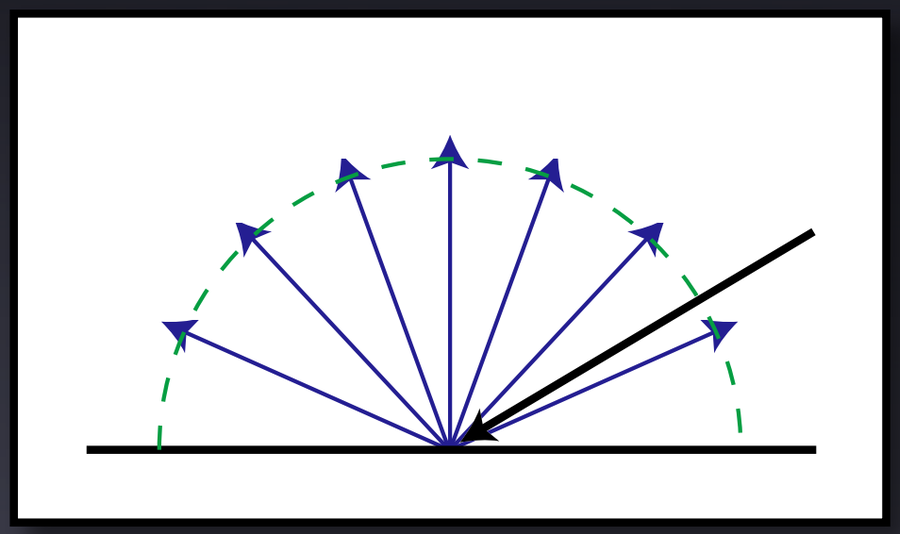

Matte surface

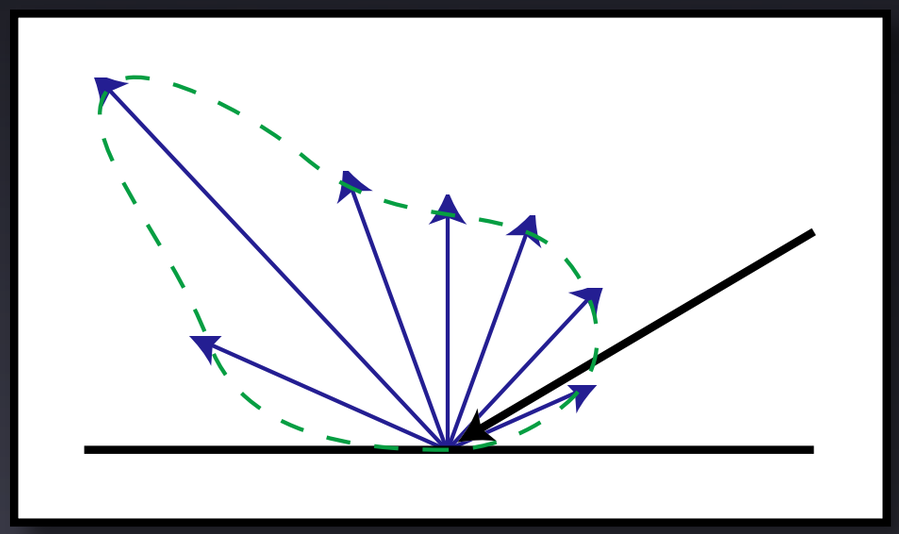

Glossy surface

Glossy Surfaces

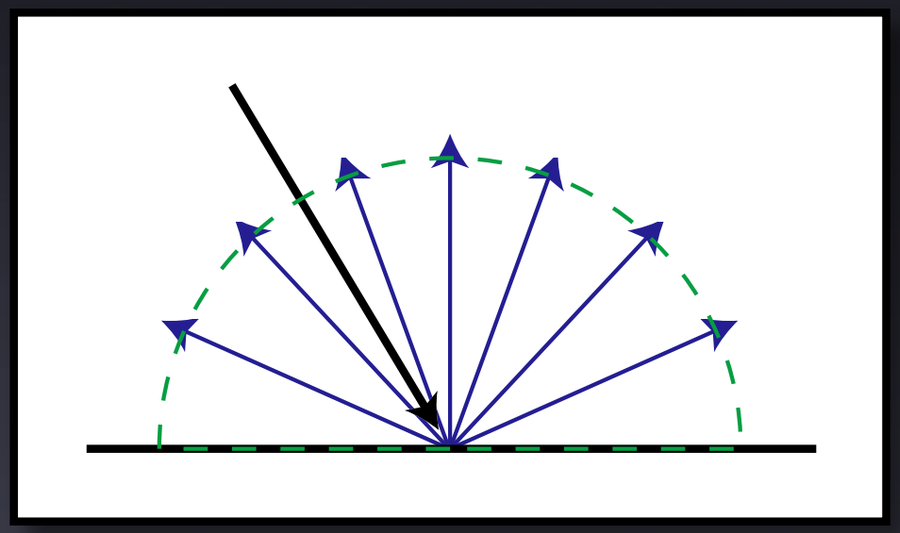

Matte surface

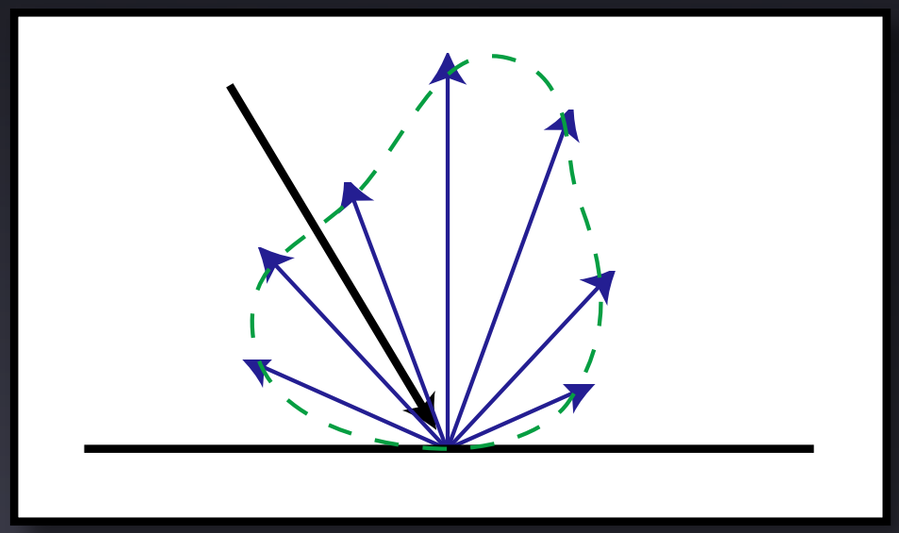

Glossy surface

Glossy Surfaces

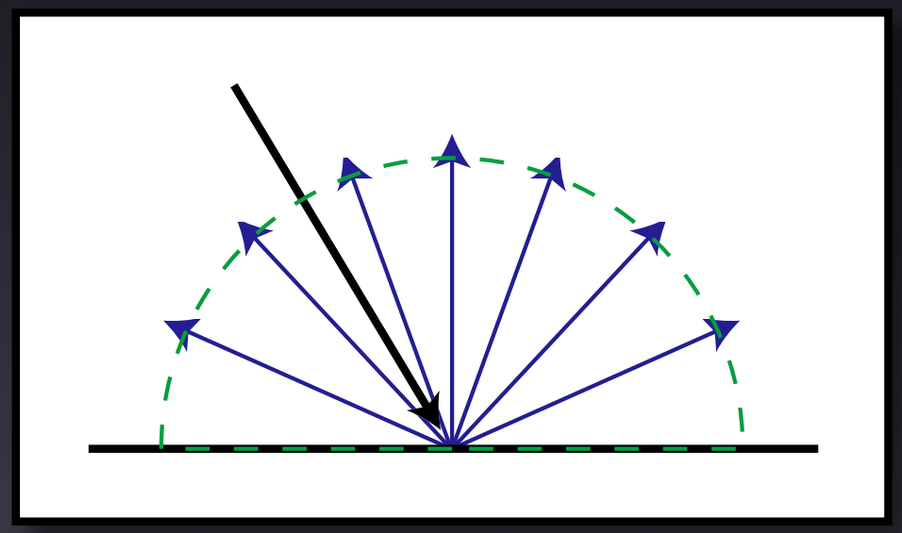

Matte surface

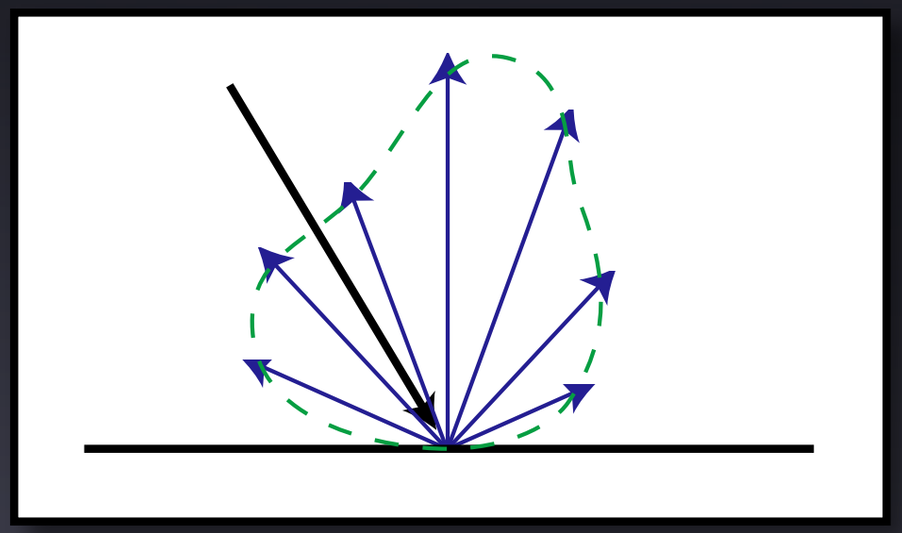

Glossy surface

Bidirectional Reflectance Distribution Function
BRDF

Radiance Caching

Křivánek et al. '05a, '05b

- Can no longer cache just the average irradiance value.
- Cache full hemispherical radiance field at sparse locations.

Radiance Interpolation

Křivánek et al. '05a, '05b

- Interpolate the radiance field from nearby locations.

Radiance Storage

Křivánek et al. '05a, '05b

- Use spherical harmonics (SH) or hemispherical harmonics (HSH).
- Generalization of Fourier series onto spherical and hemispherical domains.
- Can efficiently approximate smooth functions with just a few numbers.

(H)SH

Křivánek et al. '05a, '05b

- Projecting a function, s, onto the (H)SH basis functions results in a set of coefficients:

$$\psi_l^m = \int_0^{2\pi} \int_0^{\frac{\pi}{2}} s(\theta, \phi) H_l^m(\theta, \phi) \sin\theta \, d\theta \, d\phi$$

(H)SH

Křivánek et al. '05a, '05b

- Projecting a function, s, onto the (H)SH basis functions results in a set of coefficients:

$$\psi_l^m = \int_0^{2\pi} \int_0^{\frac{\pi}{2}} s(\theta, \phi) H_l^m(\theta, \phi) \sin\theta \, d\theta \, d\phi$$

- From these coefficients, we can approximate the original function:

$$s(\theta, \phi) \approx \sum_{l=0}^{n-1} \sum_{m=-l}^{l} \psi_l^m H_l^m(\theta, \phi)$$

Outgoing Radiance

Křivánek et al. '05a, '05b

- Incoming radiance must be convolved with the BRDF to attain outgoing radiance.
- If BRDFs are also stored in (H)SH basis this convolution is just a dot product!

Glossy Surfaces

Křivánek et al. ‘05a, ‘05b

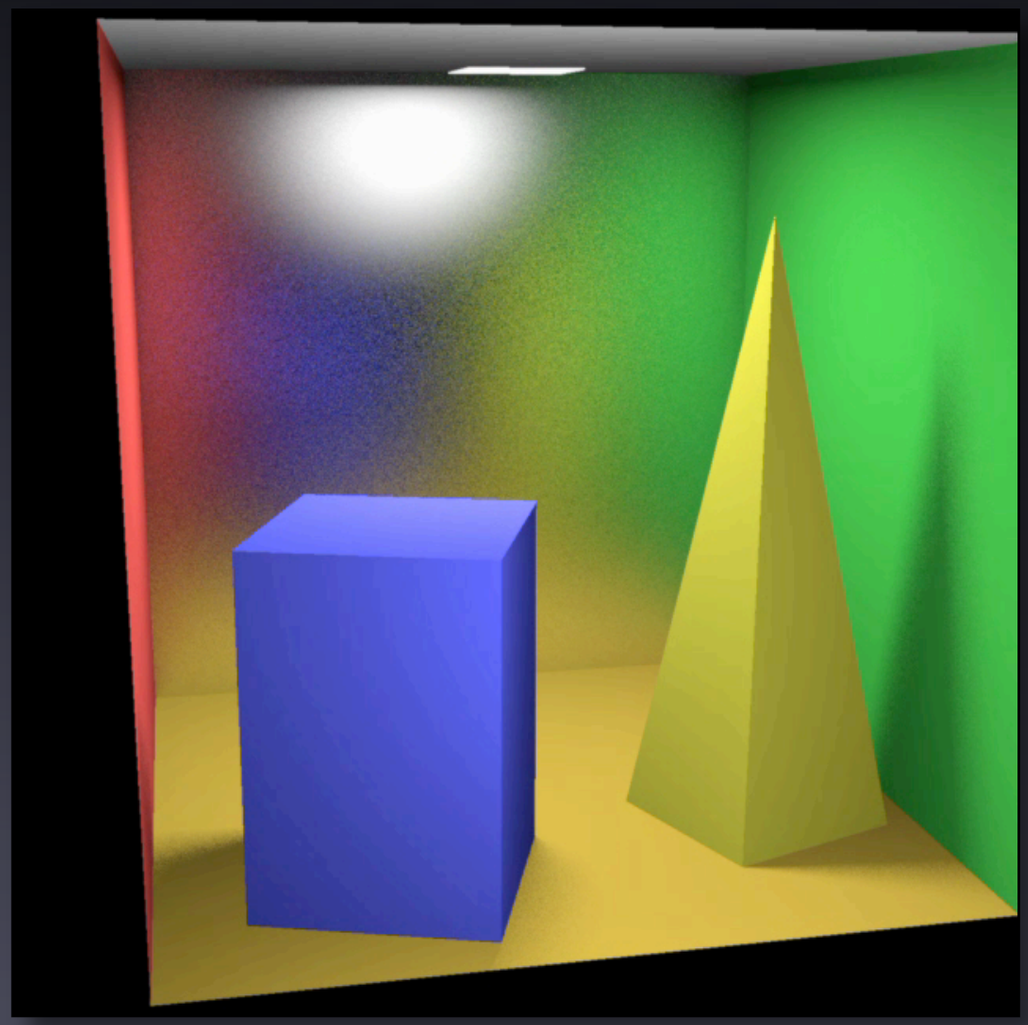

Path Tracing

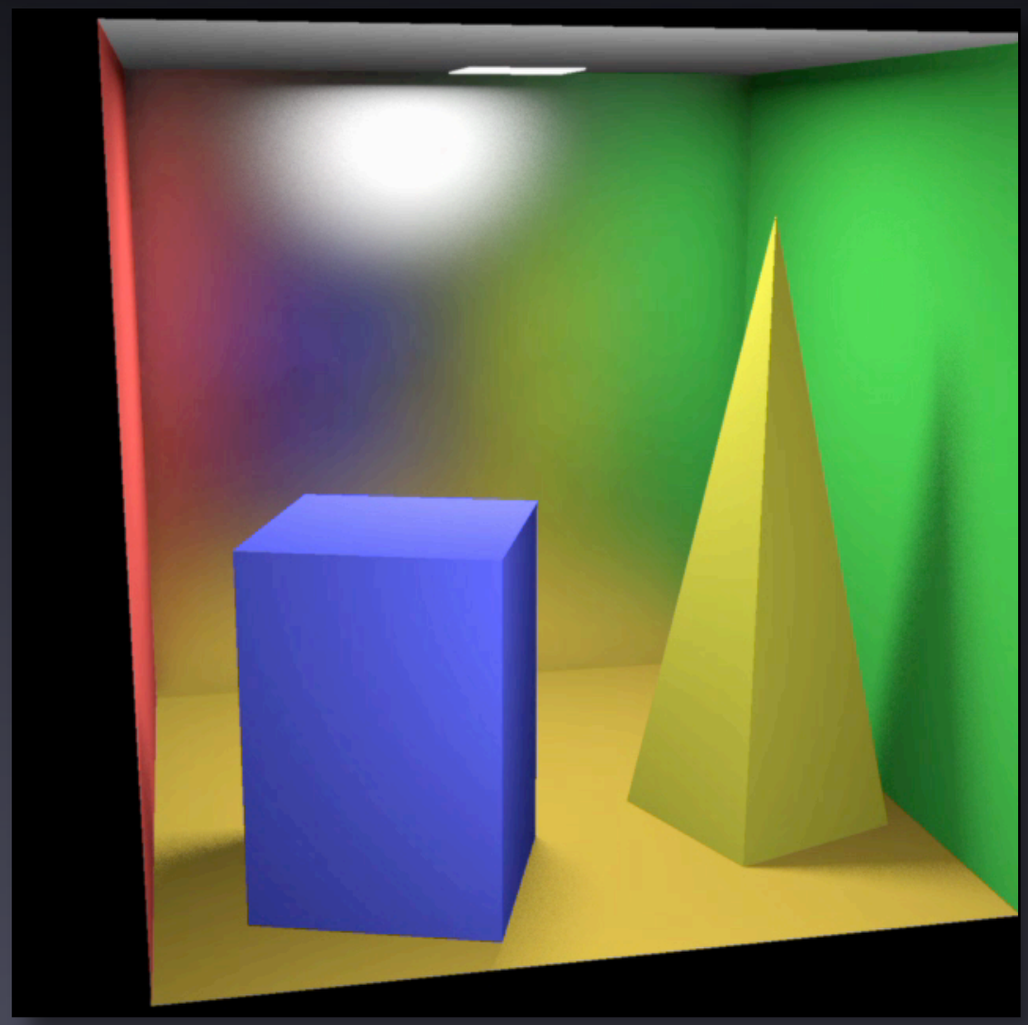

Radiance Caching

from Křivánek et al. ‘05a, ‘05b

Radiance Gradients

Křivánek et al. '05a, '05b

- Improve interpolation quality by storing gradient of incoming radiance field.

Improvements/Extensions

- Ward and Heckbert ‘92
- Křivánek et al. ‘05a, ‘05b
- Tabellion and Lamorlette ‘04
- Tawara et al. ‘04

Simplified Geometry

Tabellion and Lamorlette '04

- Apply approximations to speed up rendering.
  - Use simplified scene geometry for indirect illumination.

Simplified Geometry

Tabellion and Lamorlette ‘04

from Tabellion and Lamorlette ‘04

Simplified Geometry

Tabellion and Lamorlette ‘04

- Self intersection:

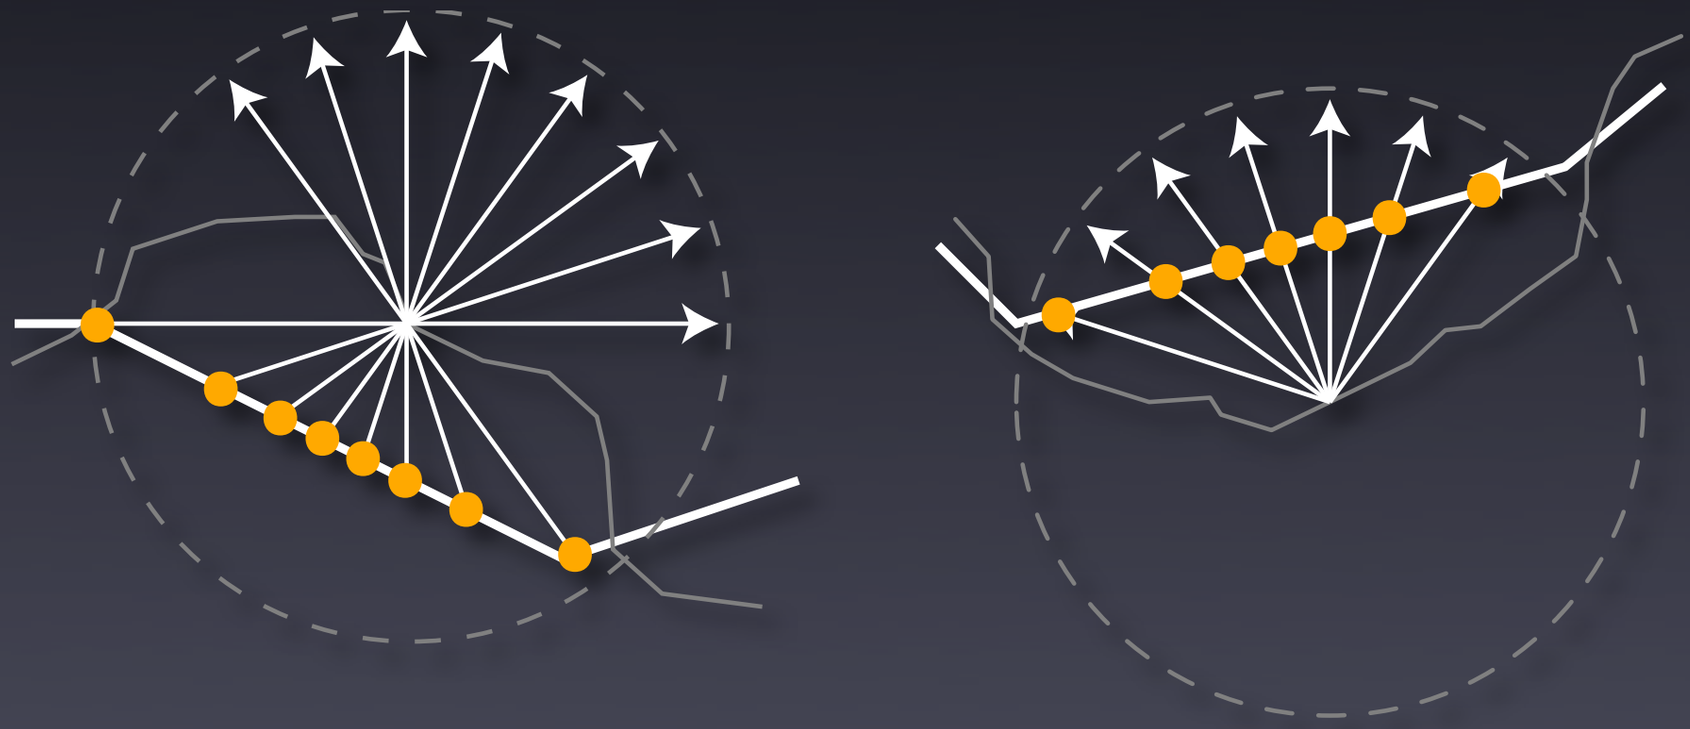

re-created from Tabellion and Lamorlette ‘04

Modified Error Metric

Tabellion and Lamorlette '04

- Introduced new error metric to reduce clumping in corners.

Over-Sampling/Clumping

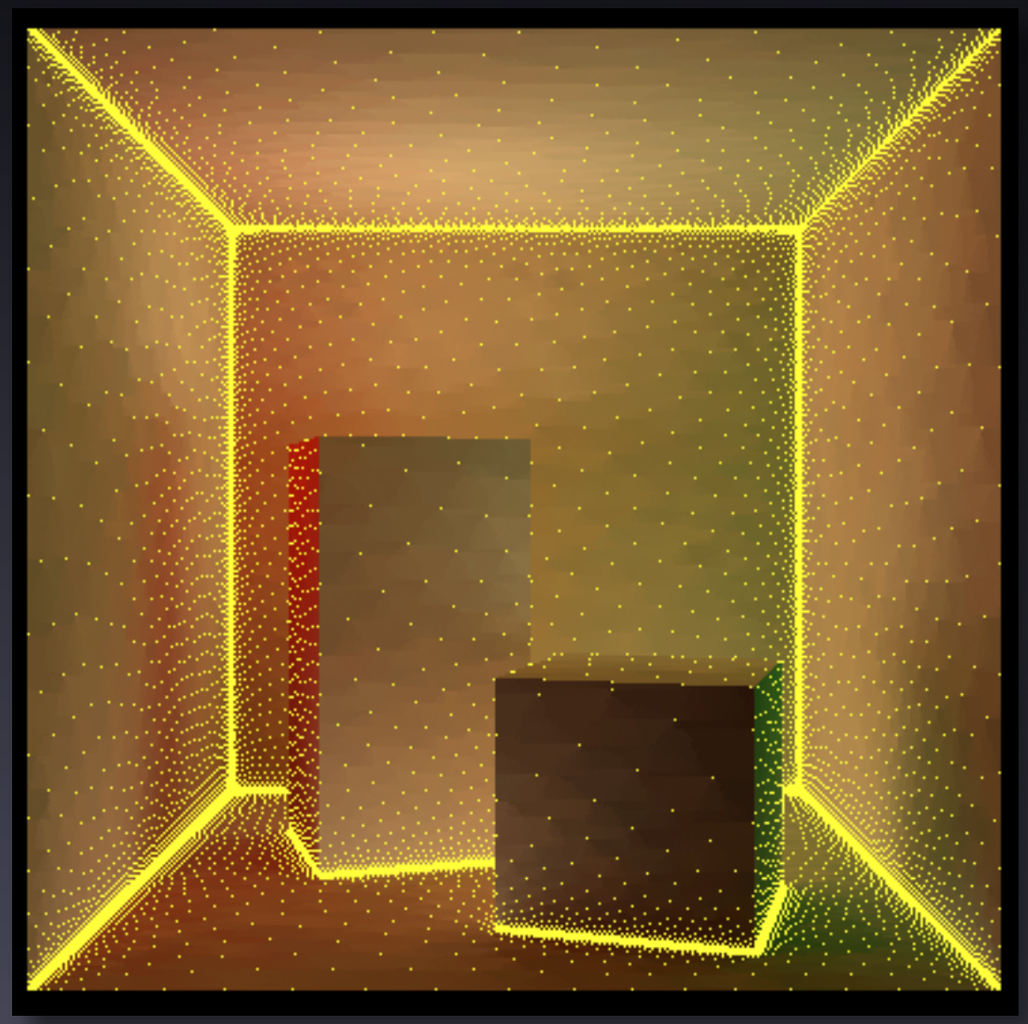

Ward Metric

Over-Sampling/Clumping

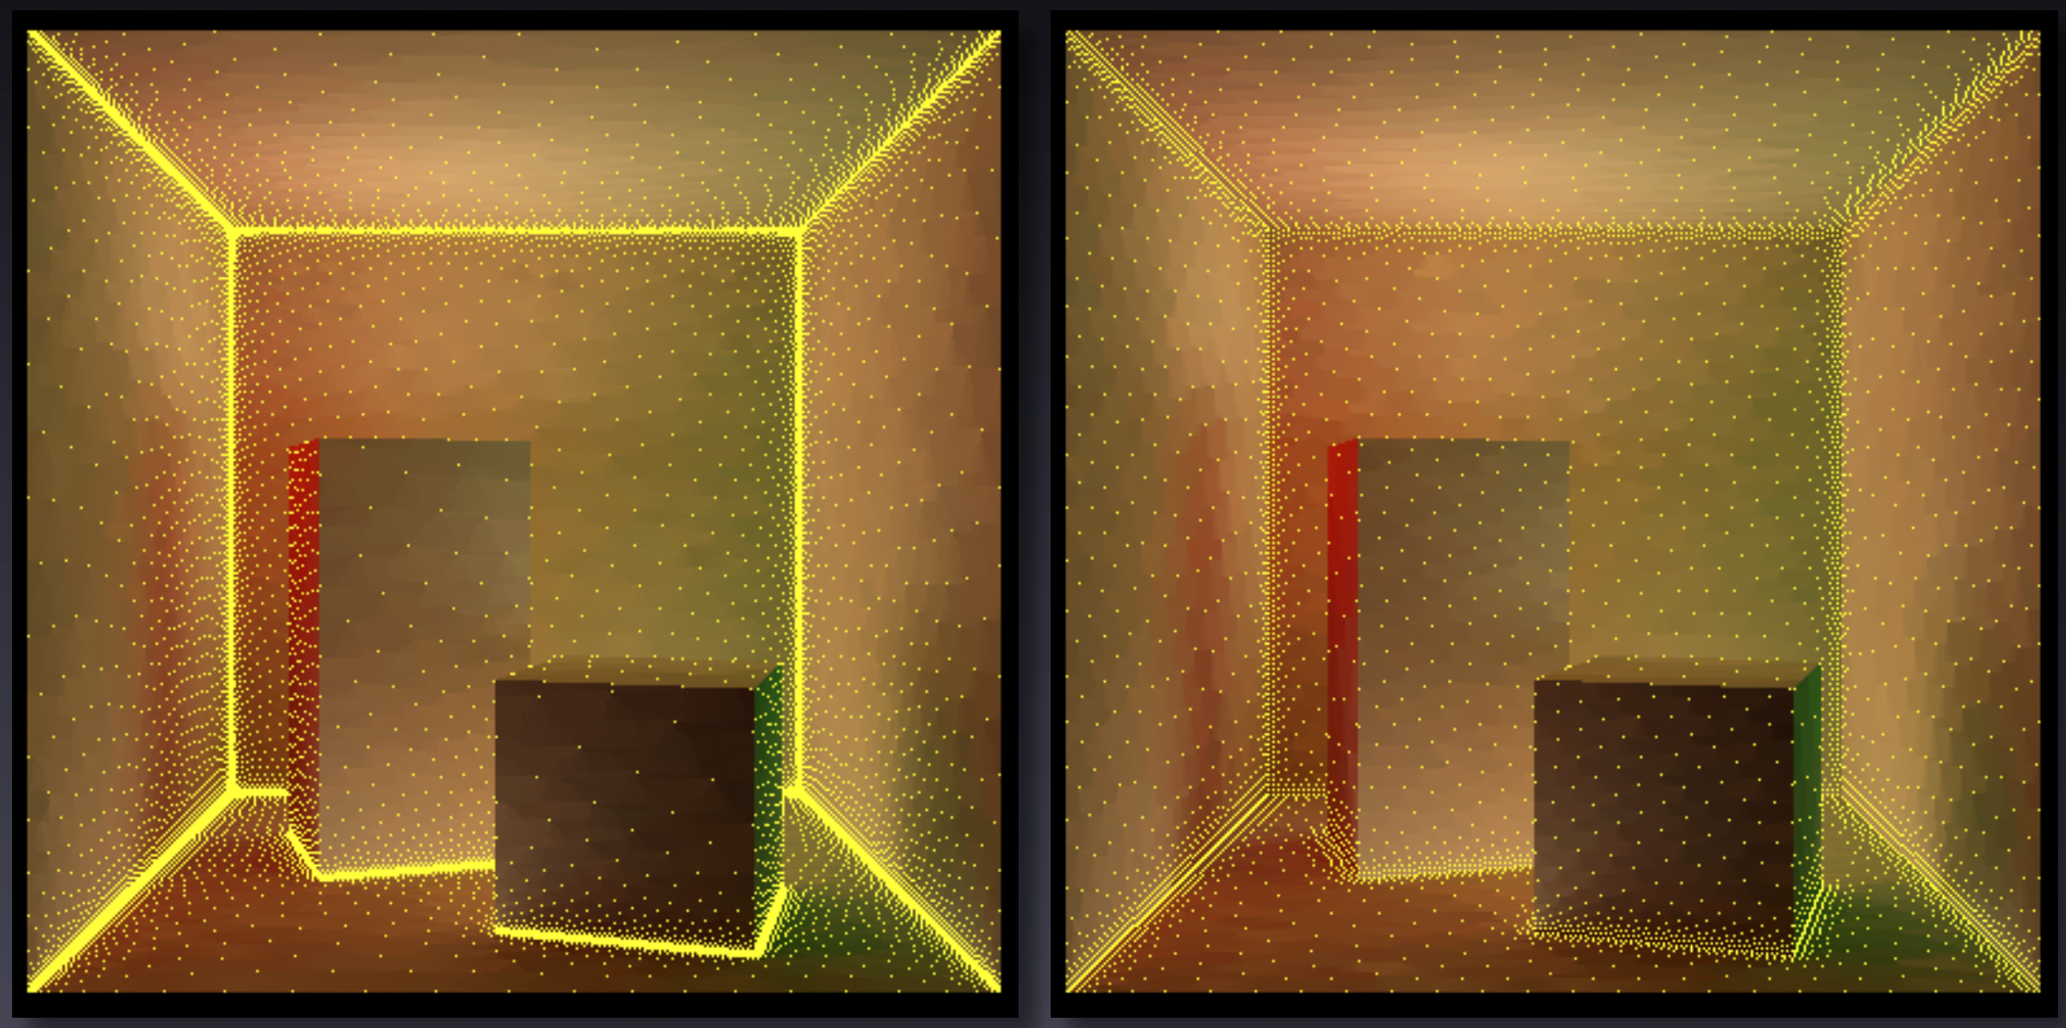

Ward Metric

Tabellion Metric

Under-Sampling

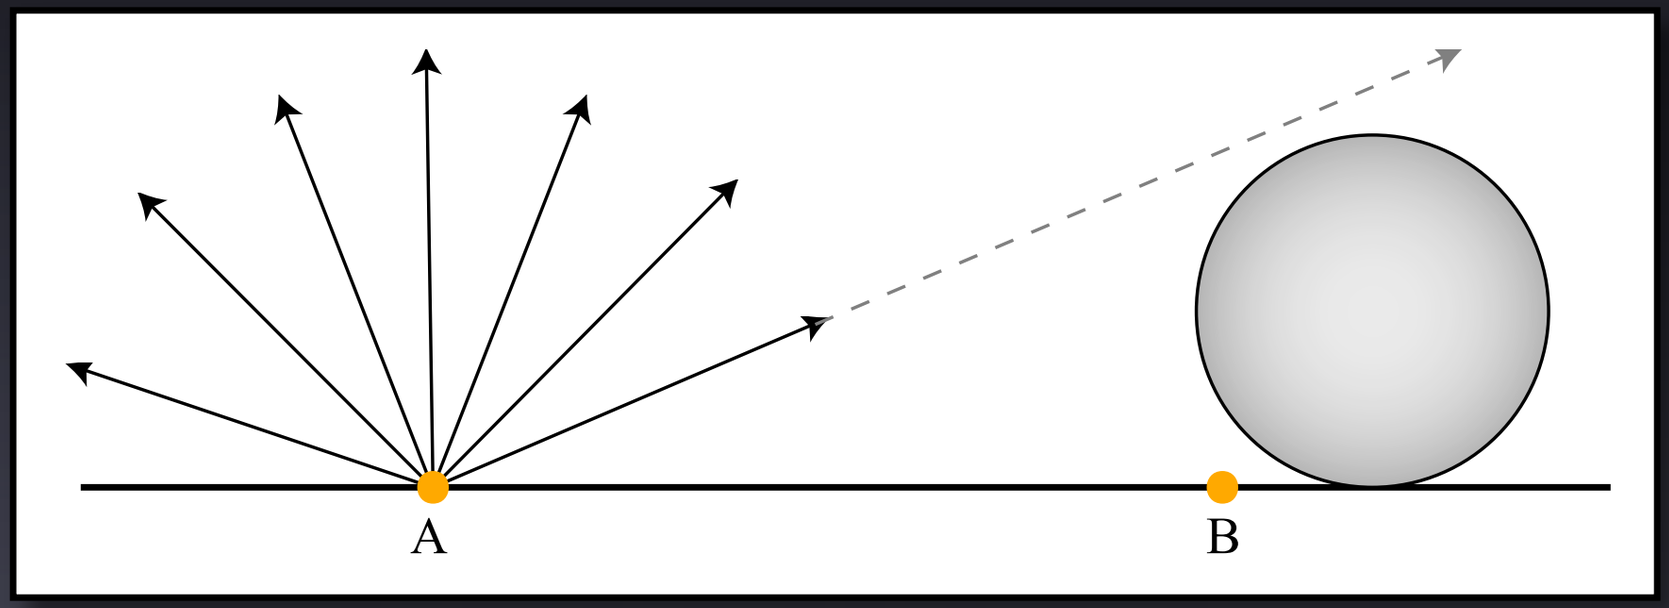

Under-Sampling

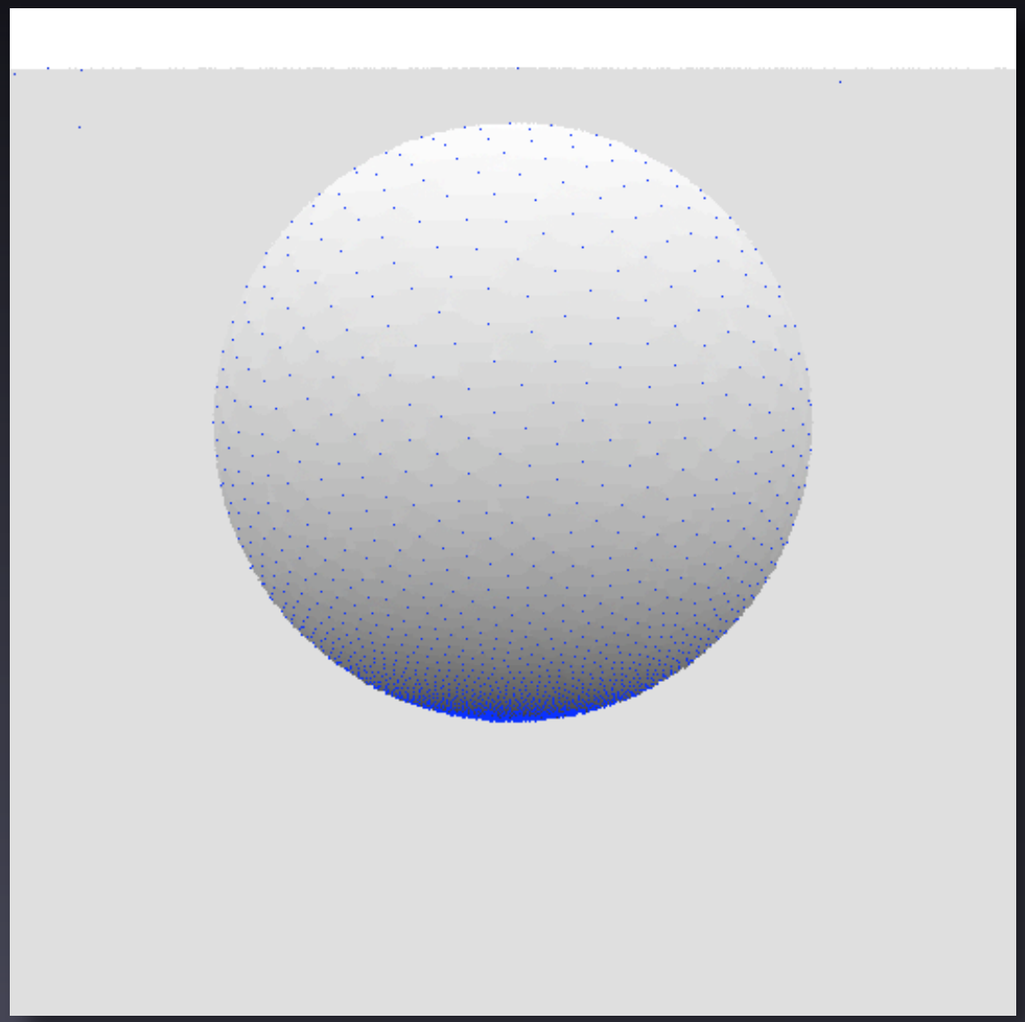

Ward Metric

Under-Sampling

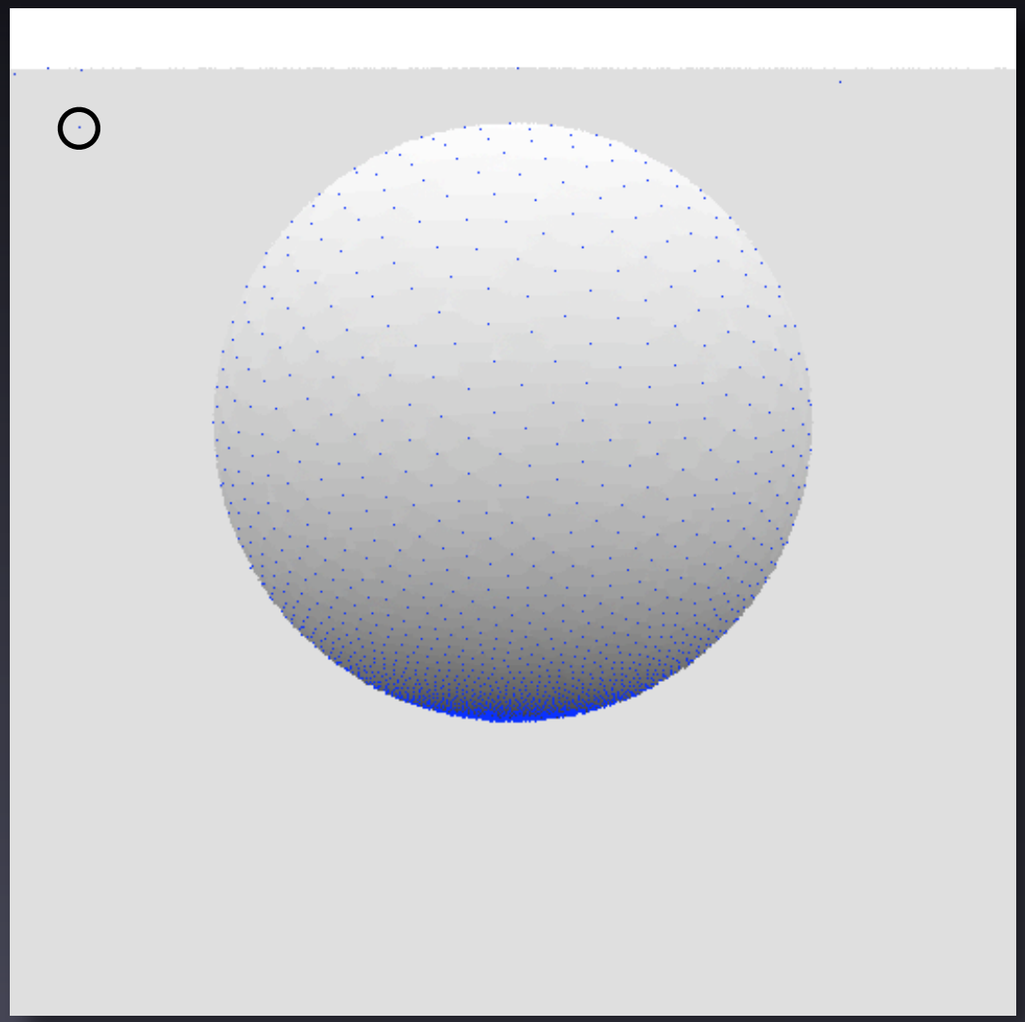

Ward Metric

Under-Sampling

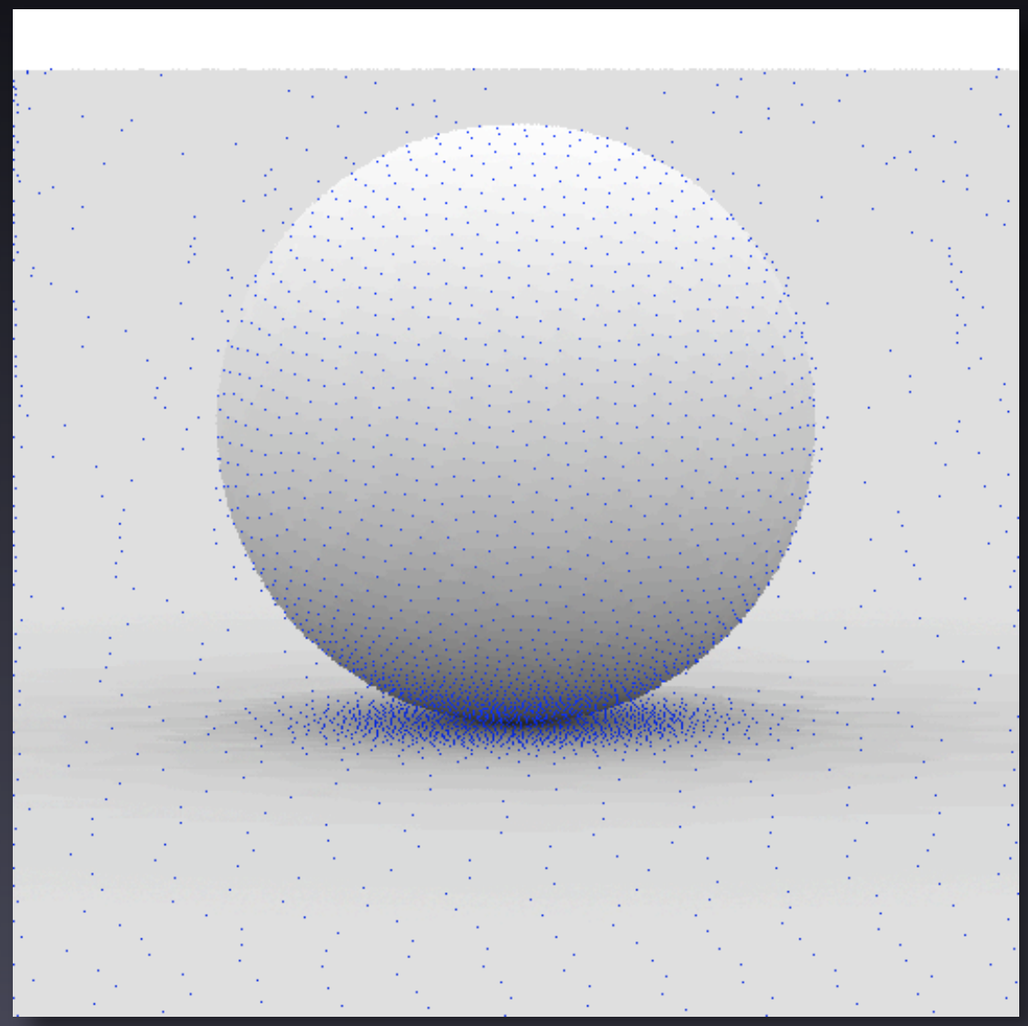

Tabellion Metric

Modified Ward Metric

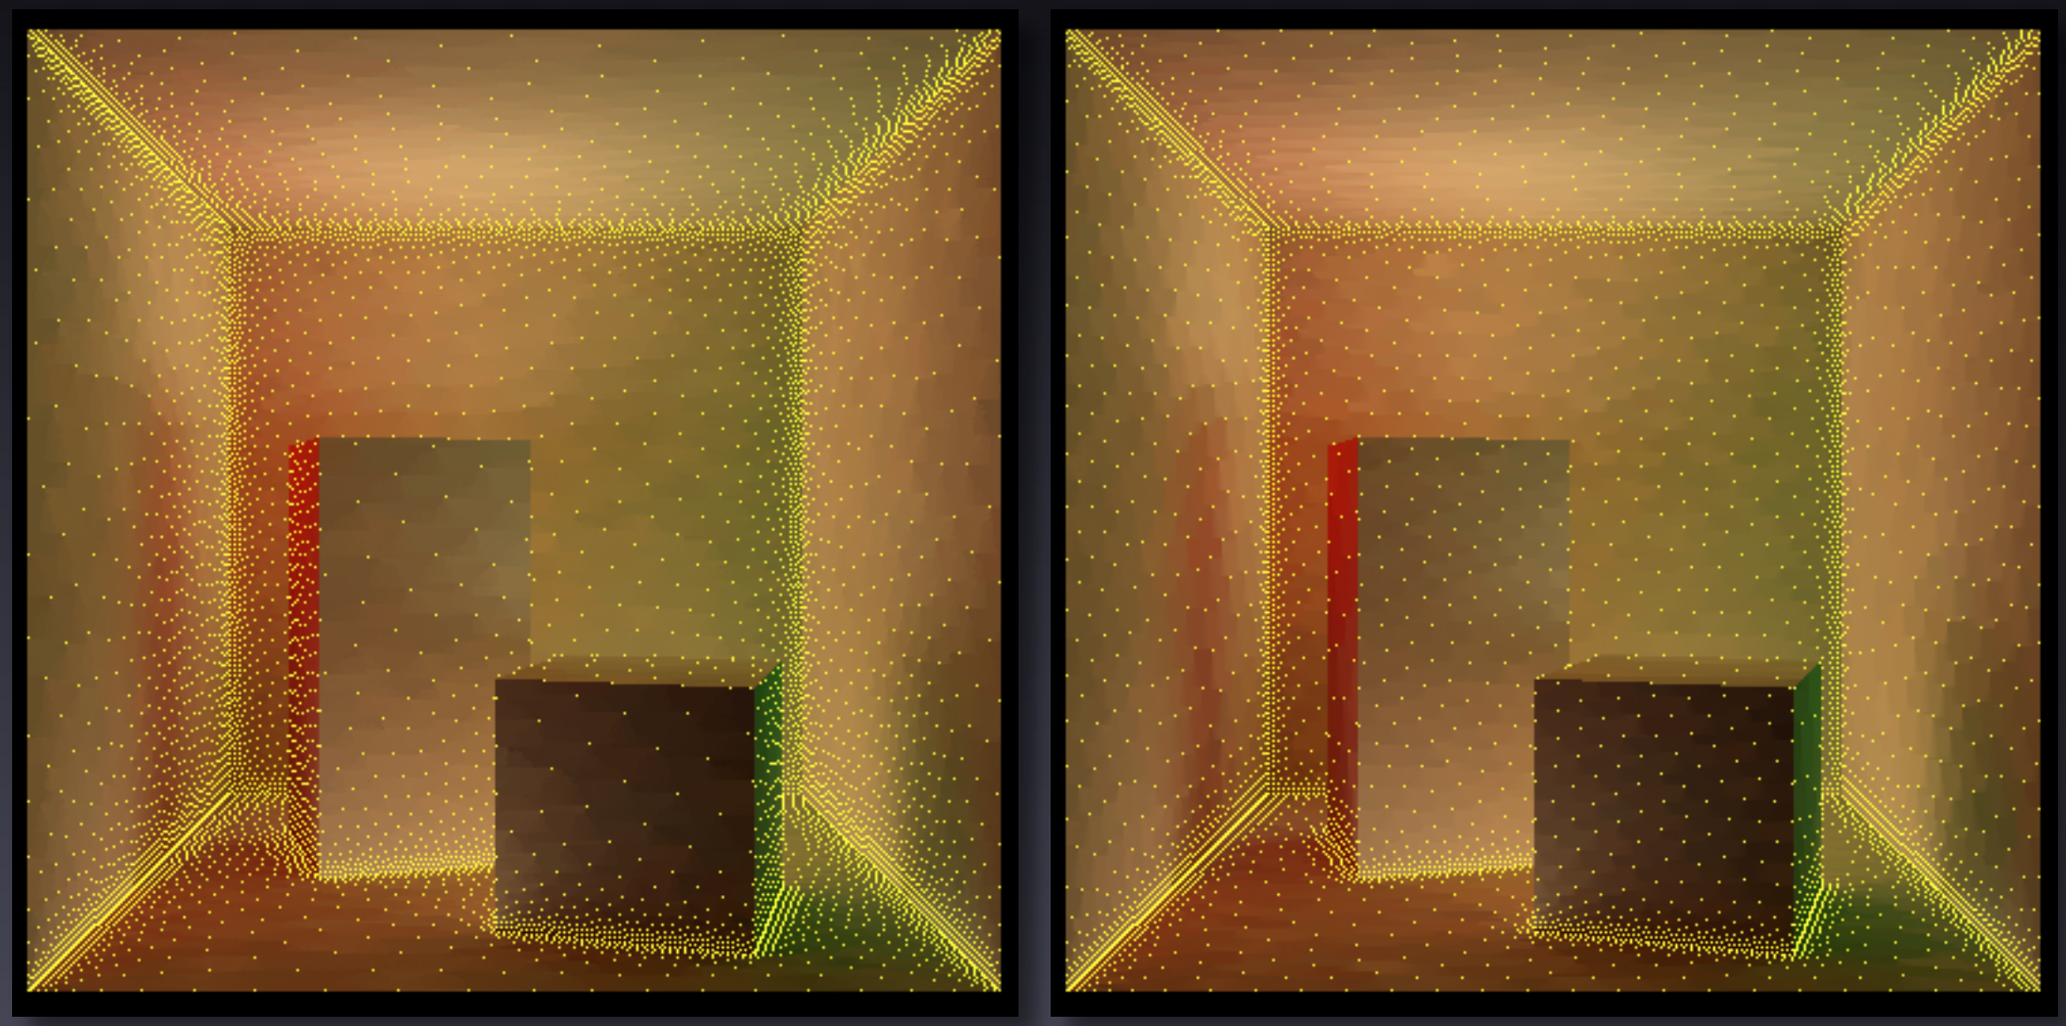

Modified Ward Metric

Tabellion Metric

Improvements/Extensions

- Ward and Heckbert ‘92
- Křivánek et al. ‘05a, ‘05b
- Tabellion and Lamorlette ‘04
- Tawara et al. ‘04

Dynamic Camera

Dynamic Camera

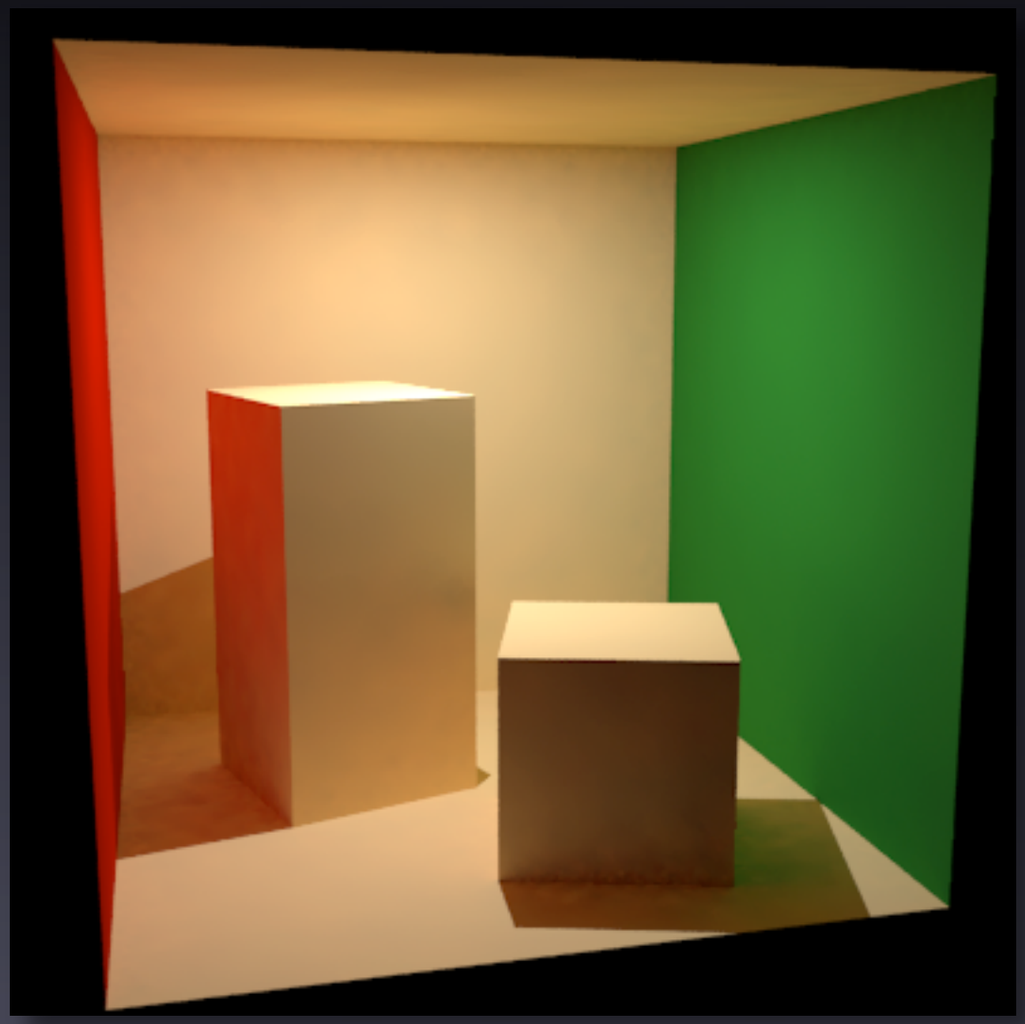

Dynamic Objects

Direct + Indirect

Indirect Illumination

Dynamic Objects

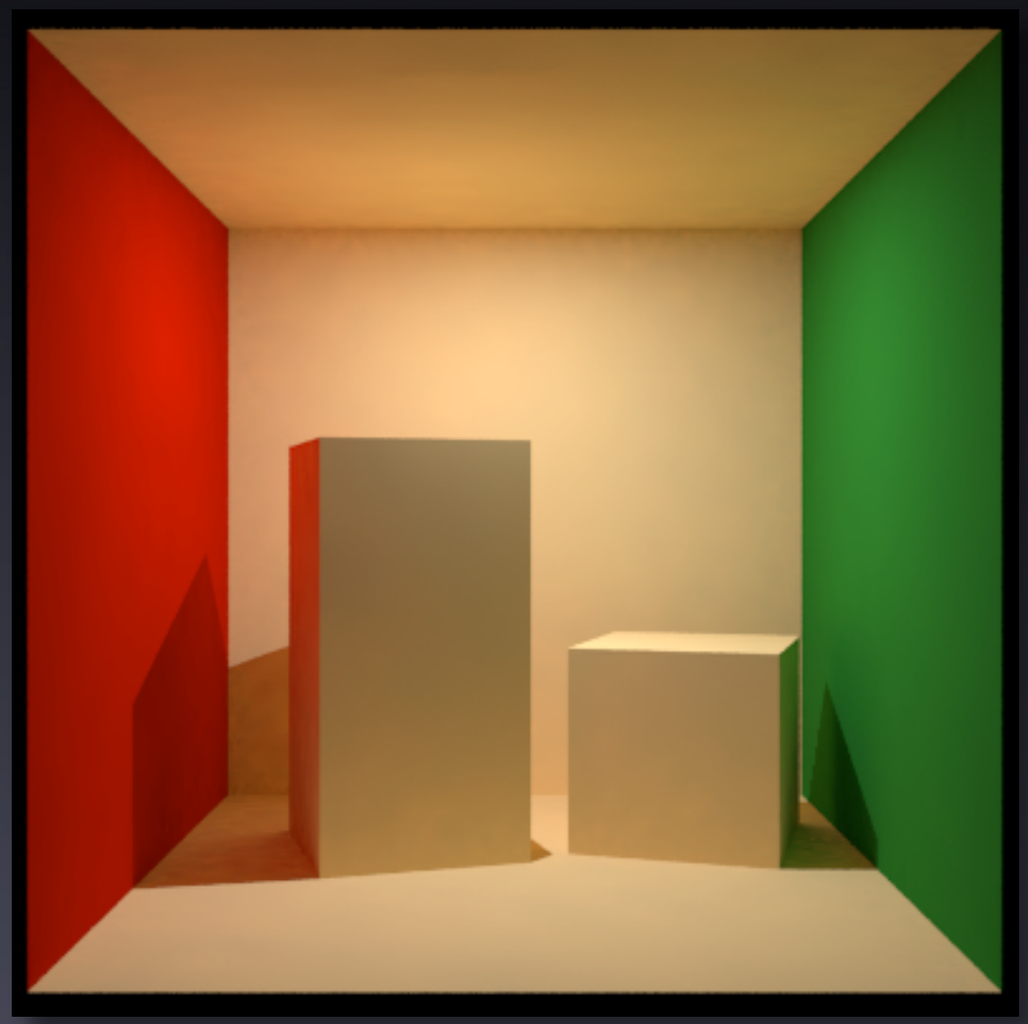

Direct + Indirect

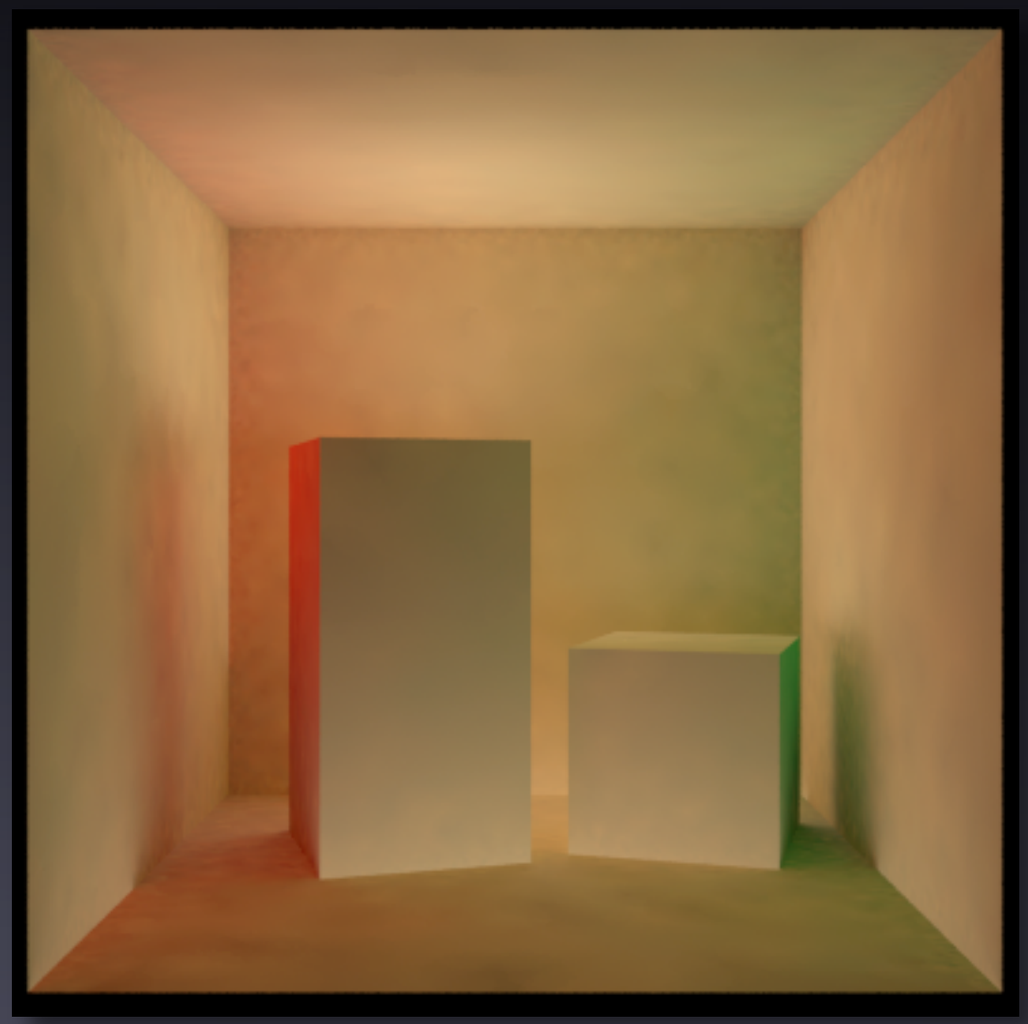

Indirect Illumination

Age-based Invalidation

Tawara et al. '04

- Only recompute a small subset of global gather rays.
- Store full stratified radiance field.
- Assign age to each stratum.
- Update a small percentage of oldest strata (10%) for frame.
- Can also give higher probability to strata which "see" moving objects.

Improvements/Extensions

- Ward and Heckbert ‘92 - better interpolation
- Křivánek et al. ‘05a, ‘05b - glossy surfaces
- Tabellion and Lamorlette ‘04 - speed
- Tawara et al. ‘04 - animation
- Yee ‘00 - speed/perception
- Kato ‘02 - parallel/distributed computation
- Arikan et al. ‘05 - speed

Next time:

Next time:

- Participating Media

Participating Media

Smoky Cornell Box

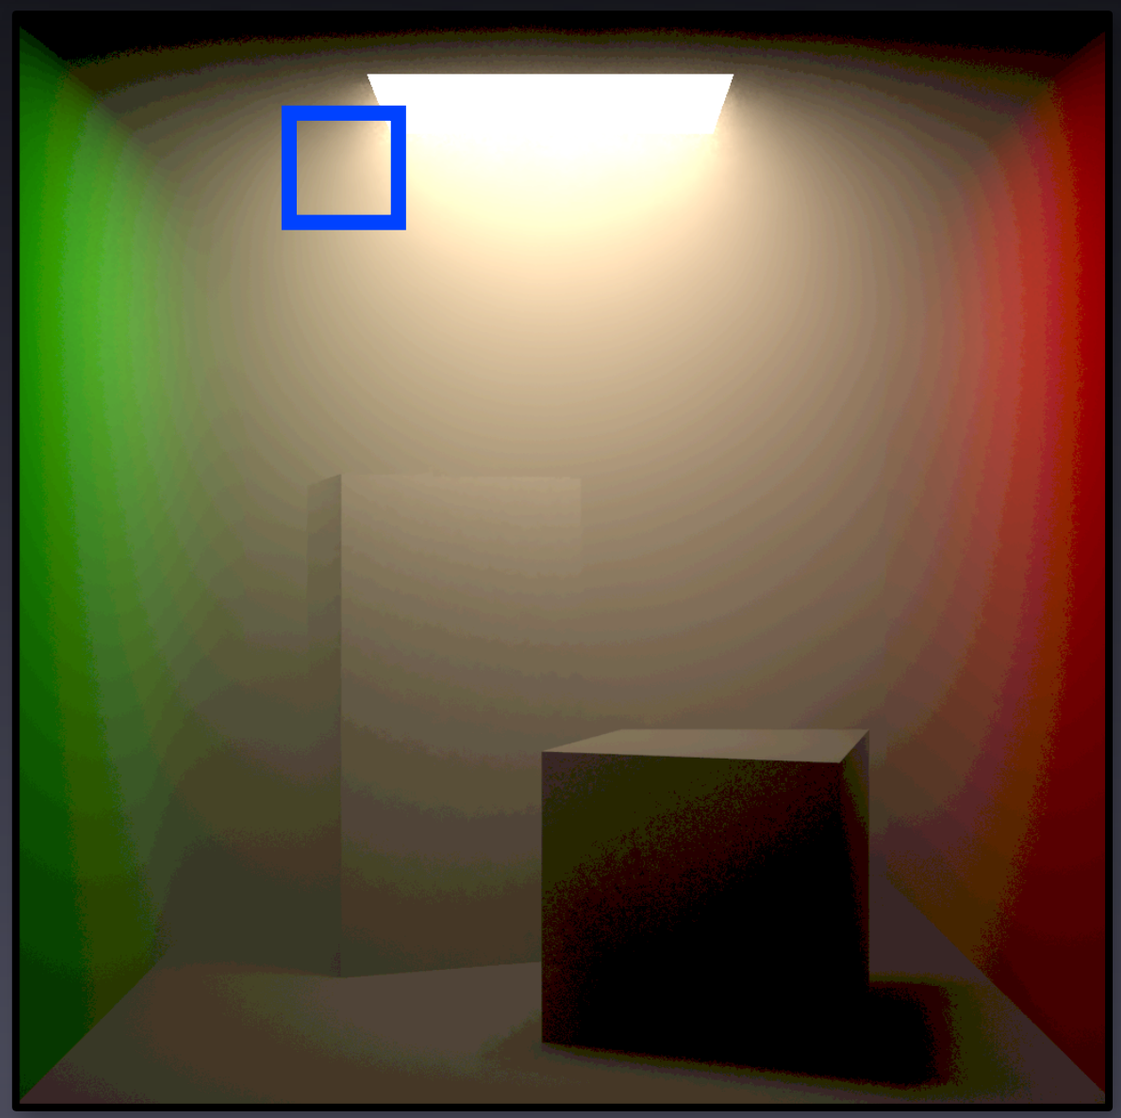

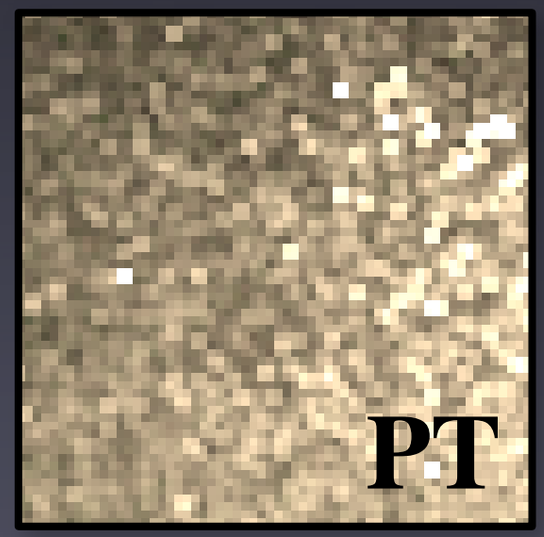

Foggy Road

Foggy Road

Questions?