Irradiance Gradients
in the Presence of
Media & Occlusions

Wojciech Jarosz
in collaboration with
Matthias Zwicker and Henrik Wann Jensen

University of California, San Diego
June 23, 2008

Thursday, 6 September 12

Thursday, 6 September 12

* In this talk we are interested in rendering scenes such as this one, where there is a strong connection between lighting that arrives at surfaces and lighting within participating media such as dust in the air or smoke
* In particular we are interested in efficiently computing the intricate indirect lighting arriving on the surfaces.

Previous Work

(Ir)radiance Caching Methods

- Irradiance Caching - Ward et al. ‘88
- Irradiance Gradients - Ward and Heckbert ‘92
- Radiance Caching - Křivánek et al. ’05
- Volumetric Radiance Caching - Jarosz et al. ‘08

* There has been a vast amount of work on how to compute indirect illlumination, enough to fill a whole course.
* A popular technique, which is most related to our work, is irradiance caching, which was originally developed in 1988, and has been subsequently improved.

Irradiance Caching

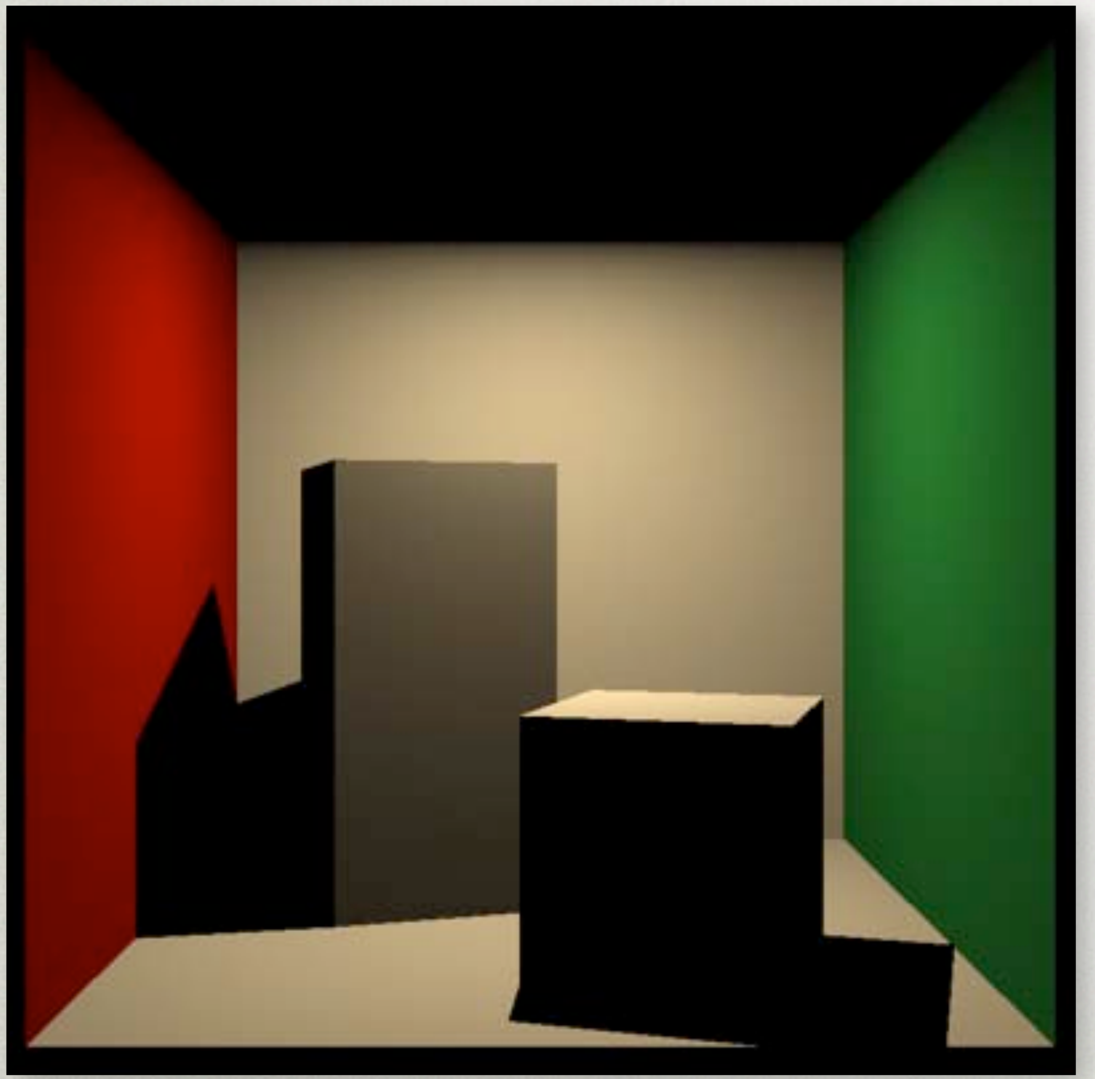

Direct
(sharp discontinuities)

Indirect
(low frequency)

Ward et al. '88

* The observation that ward made was that even though direct lighting may have sharp discontinuities, such as...
* If we just look at indirect irradiance, by handling direct lighting separately, it tends to have a very smooth appearance.

Irradiance Caching

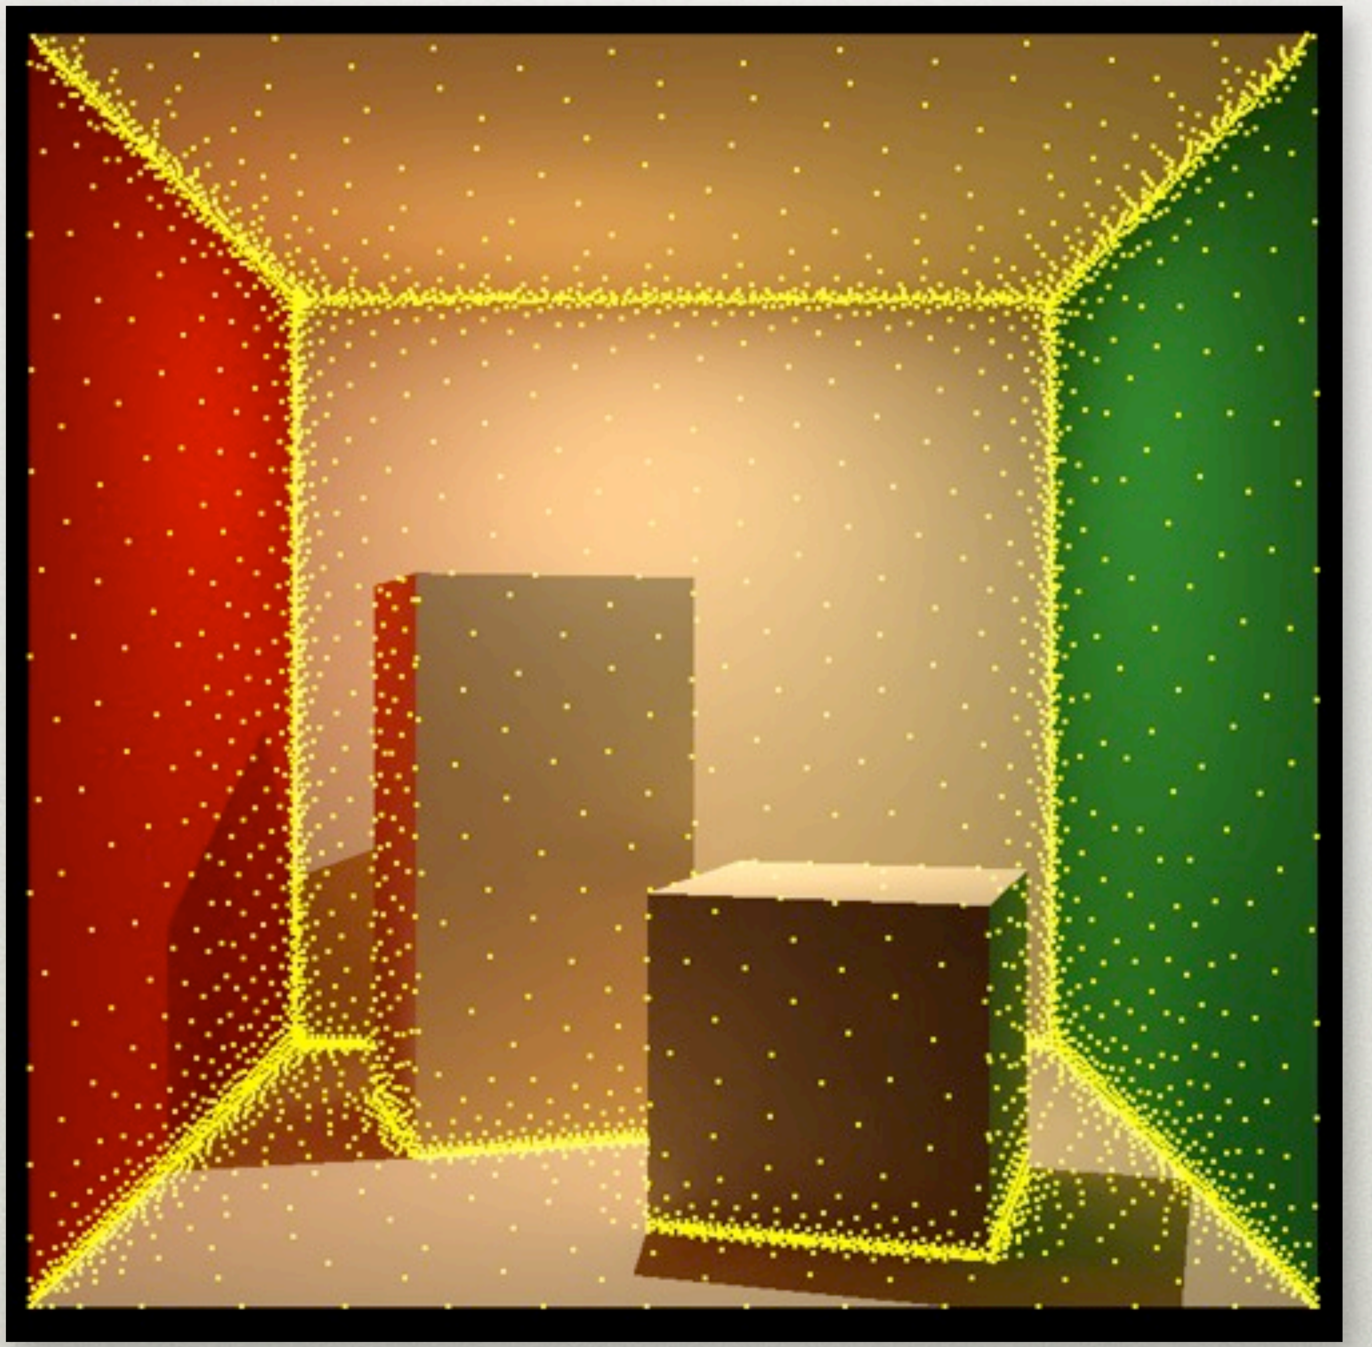

4621 samples

Ward et al. '88

Thursday, 6 September 12

* This makes it a perfect candidate for sparse sampling and interpolation.
* Irradiance caching computes indirect irradiance only at a sparse set of locations in the scene, and tries to interpolate these values as often as possible in order to gain efficiency.
* On average only about 1 out of every 50 pixels need to compute indirect lighting in this image

Previous Work

(Ir)radiance Caching Methods

- Irradiance Caching - Ward et al. '88
- Irradiance Gradients - Ward and Heckbert '92
- Radiance Caching - Křivánek et al. '05
- Volumetric Radiance Caching - Jarosz et al. '08

Interpolation Quality

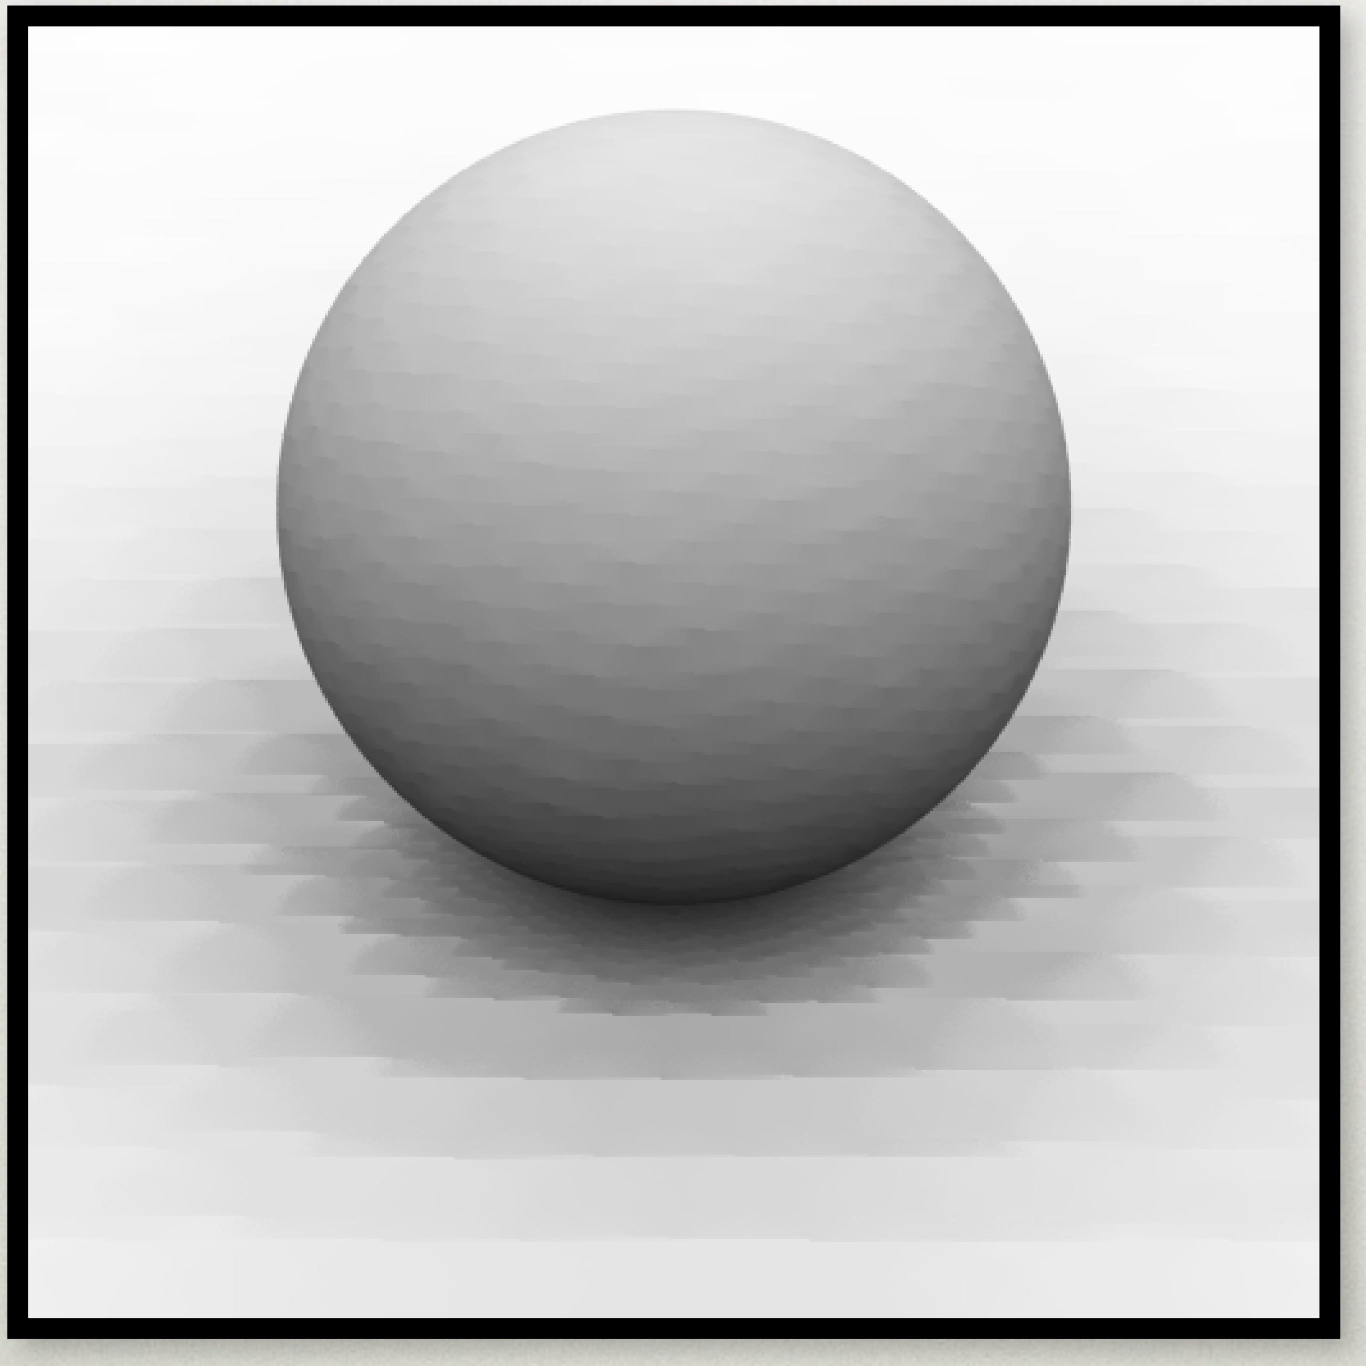

Ward and Heckbert '92

Interpolation Quality

Without Gradients

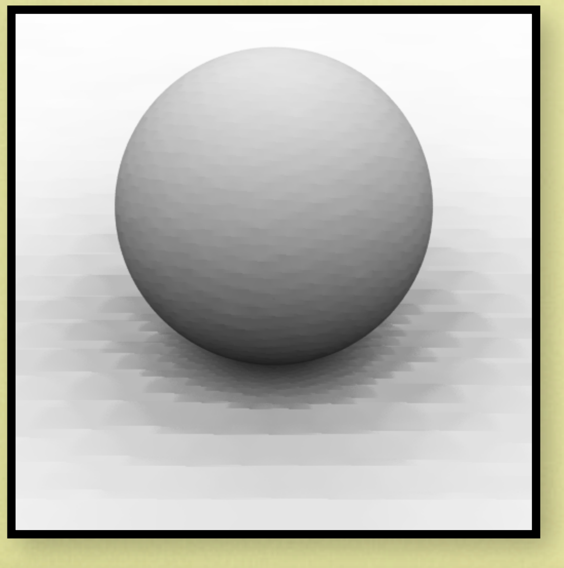

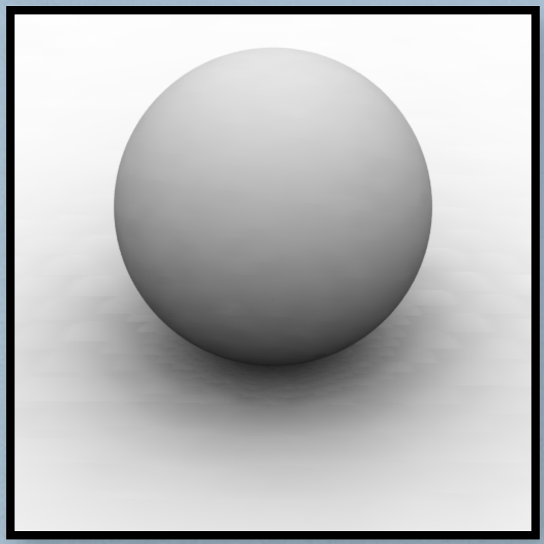

With Gradients

Ward and Heckbert '92

Interpolation Quality

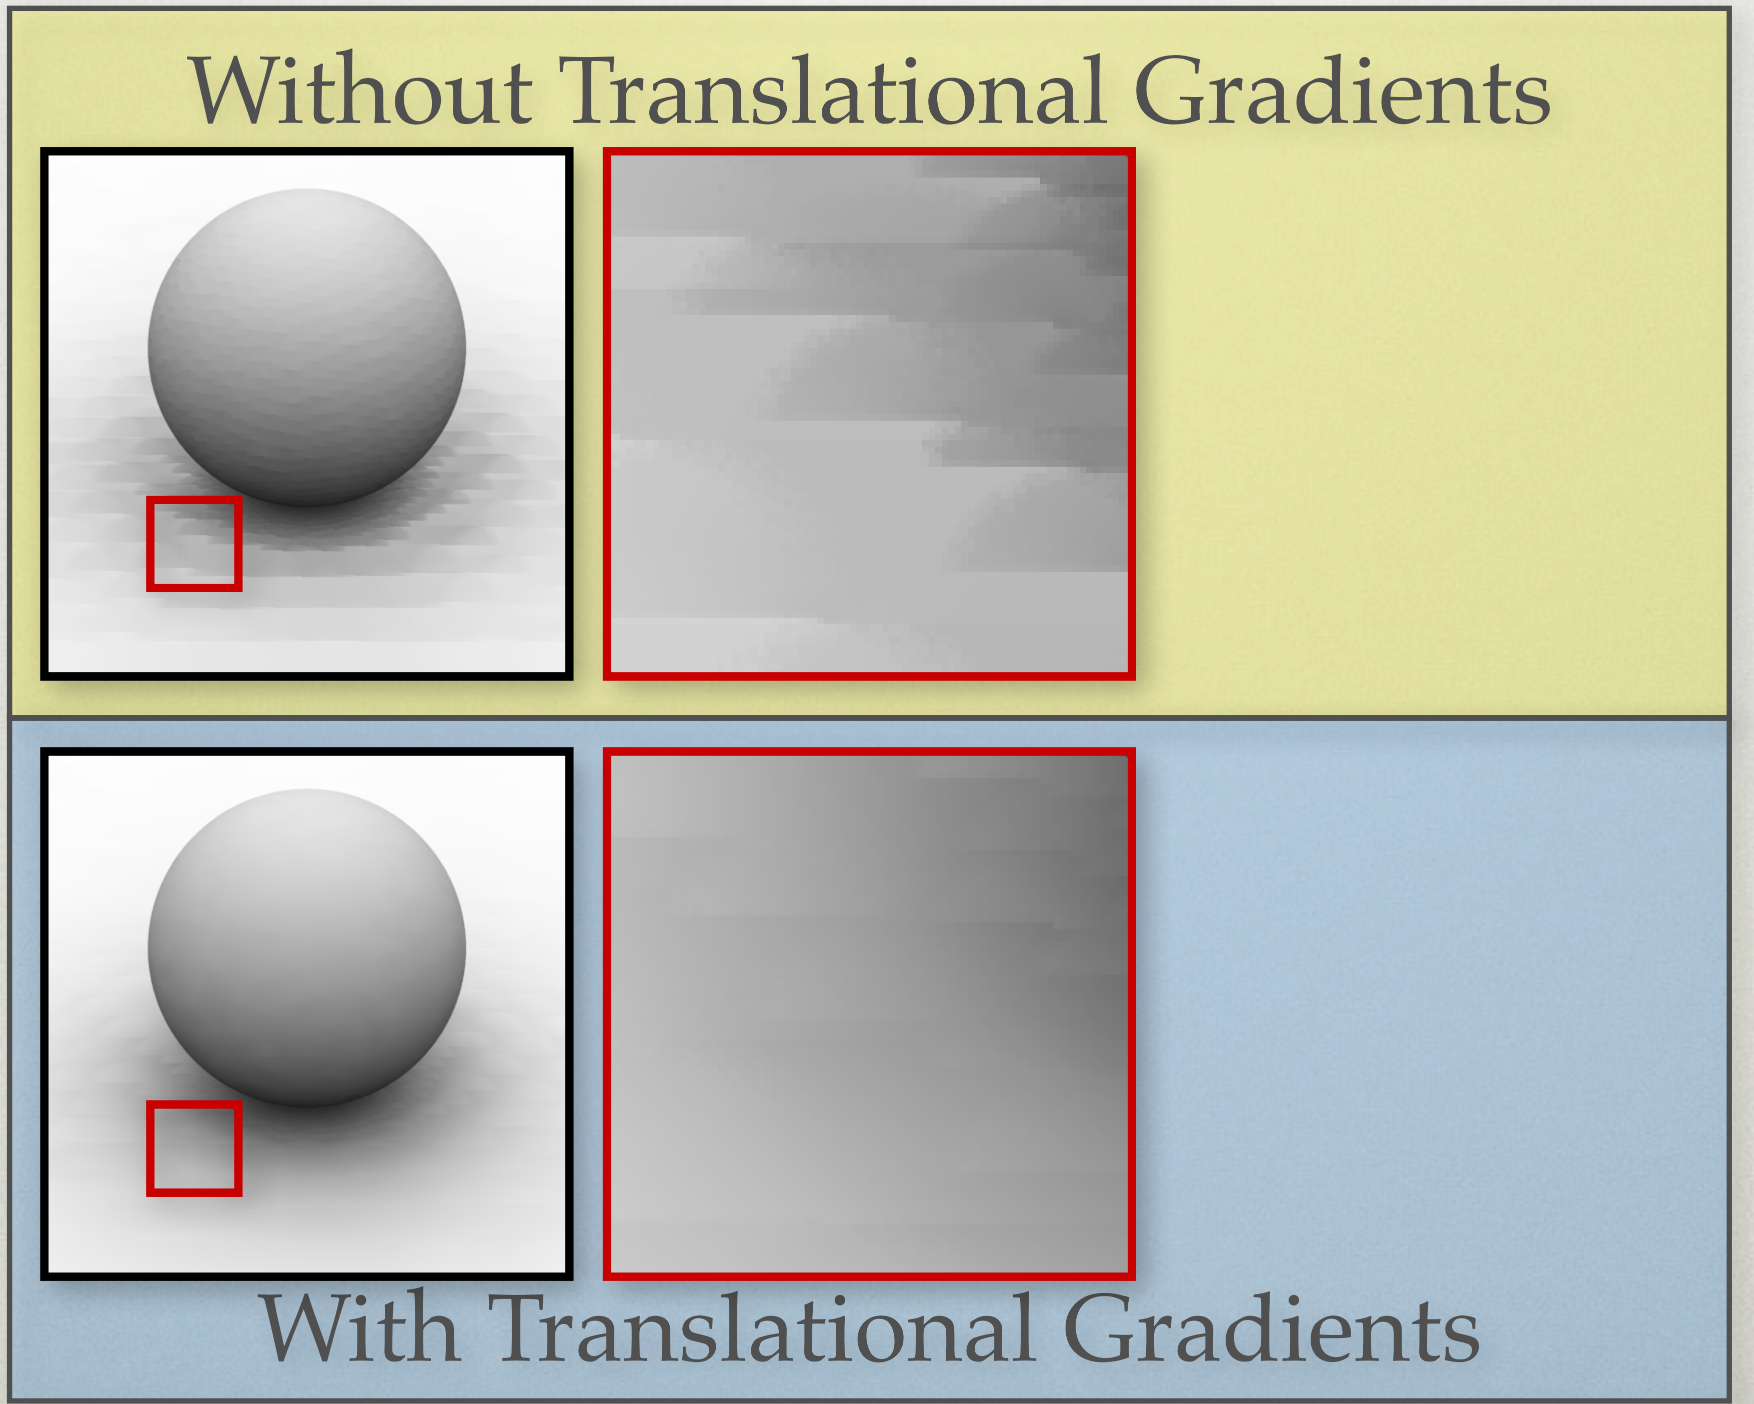

Ward and Heckbert ‘92

Interpolation Quality

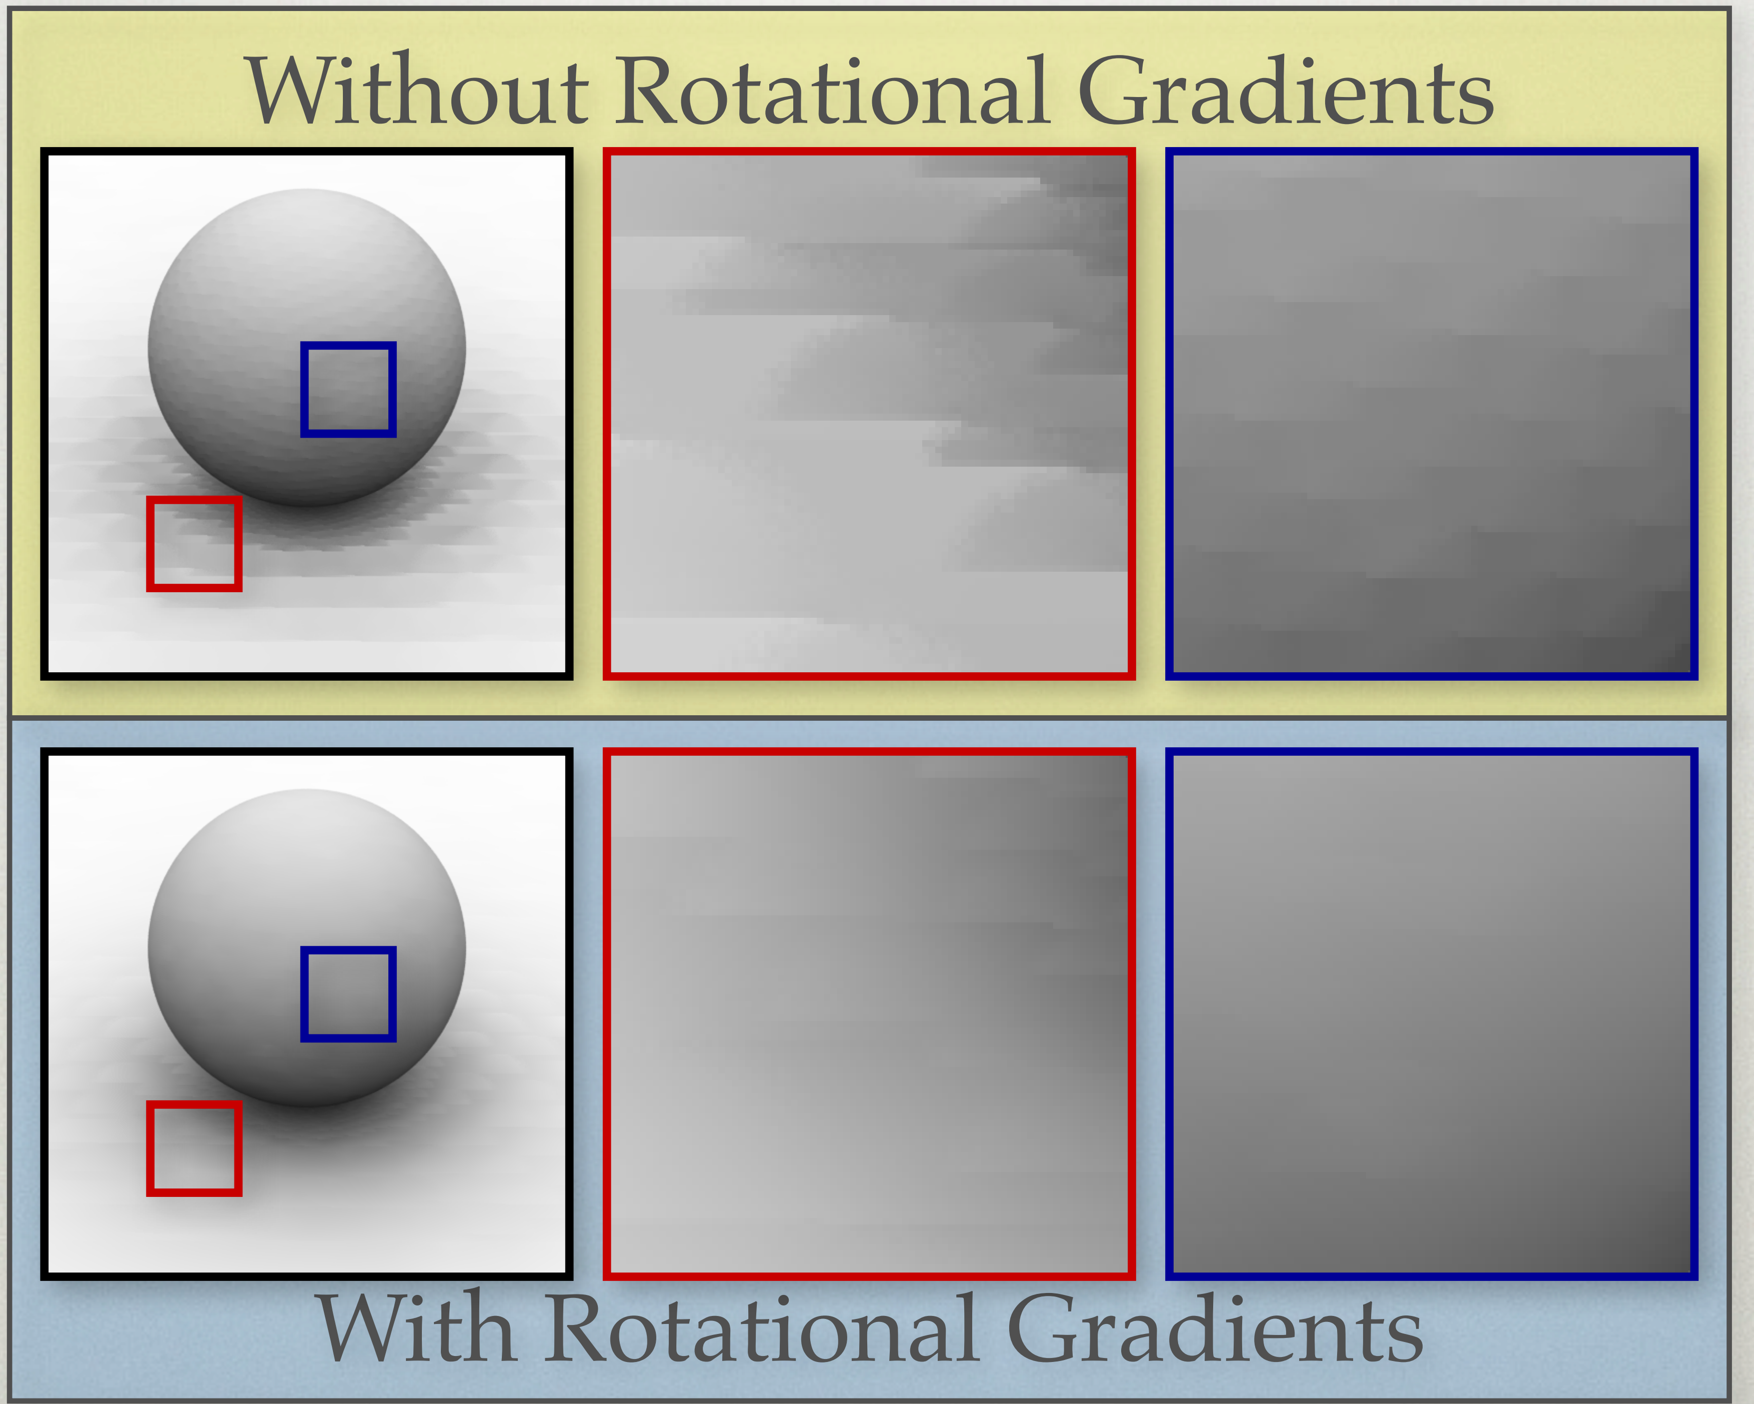

Ward and Heckbert '92

Thursday, 6 September 12

Previous Work

(Ir)radiance Caching Methods

- Irradiance Caching - Ward et al. ‘88
- Irradiance Gradients - Ward and Heckbert ‘92
- Radiance Caching - Křivánek et al. ’05
    - Support caching on glossy surfaces
- Volumetric Radiance Caching - Jarosz et al. ‘08

Previous Work

(Ir)radiance Caching Methods

- Irradiance Caching - Ward et al. '88
- Irradiance Gradients - Ward and Heckbert '92
- Radiance Caching - Křivánek et al. '05
- Volumetric Radiance Caching - Jarosz et al. '08
    - Cache radiance within volume, compute radiance gradients.

Previous Work

(Ir)radiance Caching Methods

- Irradiance Caching - Ward et al. '88
- Irradiance Gradients - Ward and Heckbert '92
- Radiance Caching - Křivánek et al. '05
- Volumetric Radiance Caching - Jarosz et al. '08

–Do not take into account participating media

* The limitations of the first three methods is that they do not account for participating media. They assume all surfaces are in a vacuum.
* This means we cannot effectively apply these gradient methods if the scene contains media

Previous Work

(Ir)radiance Caching Methods

- Irradiance Caching - Ward et al. ‘88
- Irradiance Gradients - Ward and Heckbert ‘92
- Radiance Caching - Křivánek et al. ’05
- Volumetric Radiance Caching - Jarosz et al. ’08
  - Does not take into account occlusions

* On the other hand, our previous work computes gradients within participating media, and it is possible to trivially apply this to irradiance gradients by only integrating over the hemisphere, instead of the whole sphere.
* However, the drawback of our previous approach is that it does not take into account occlusions which can lead to significant interpolation artifacts in regions with occlusion changes.

Goal

- Compute accurate gradients of irradiance on surfaces in the presence of participating media AND occlusions.

Thursday, 6 September 12

* Our goal is to fill this gap in previous work and compute accurate gradients of irradiance on surfaces in the presence of participating media AND occlusions.
* We are only interested in computing a translational gradient, since rotational gradients are not effected by participating media.

Participating Media

No Media

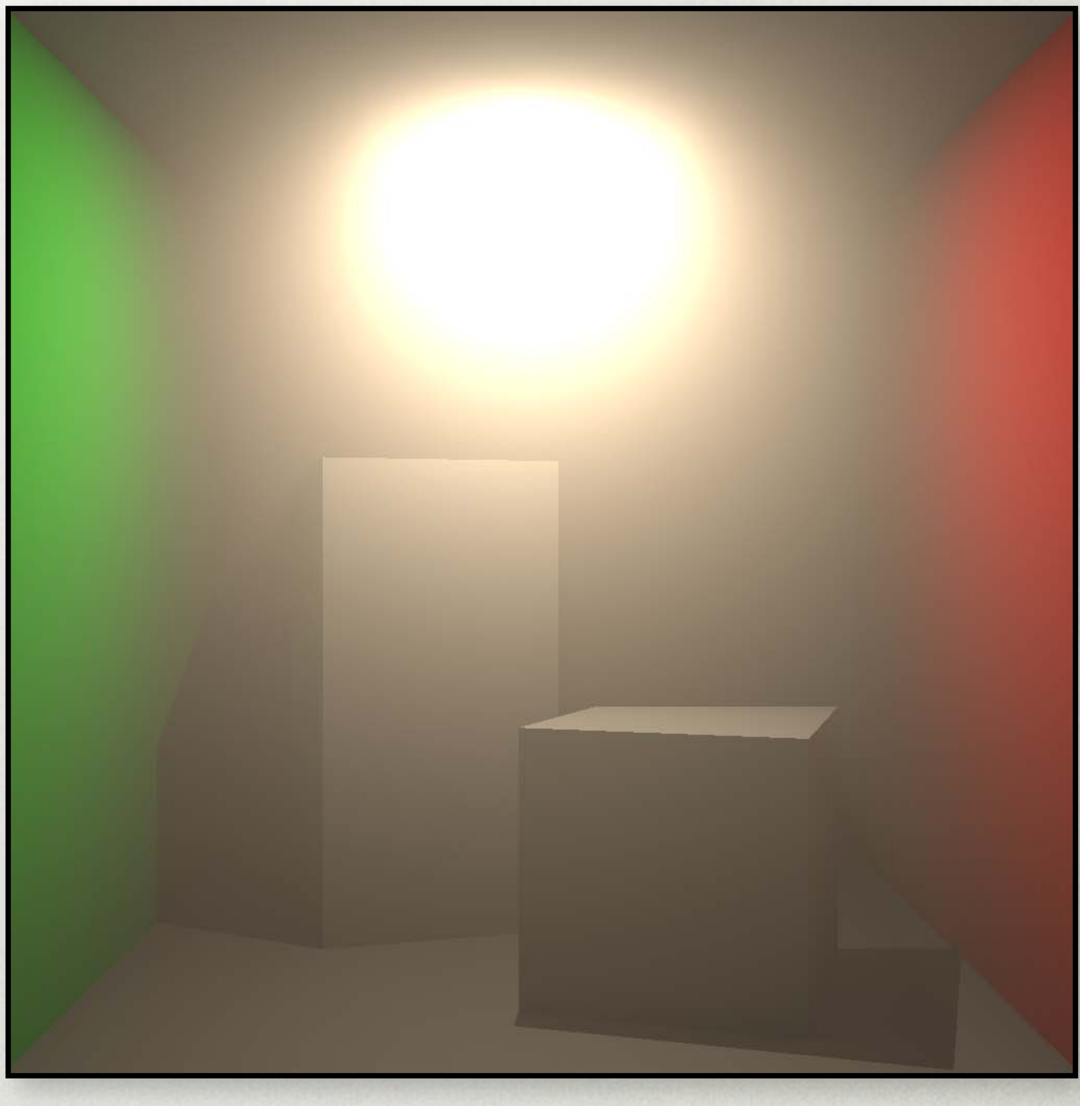

With media

* To see why this is in fact an important problem which is not adequately handled by previous methods, lets consider the classic Cornell box both with and without media.

Participating Media

No Media
(indirect irradiance)

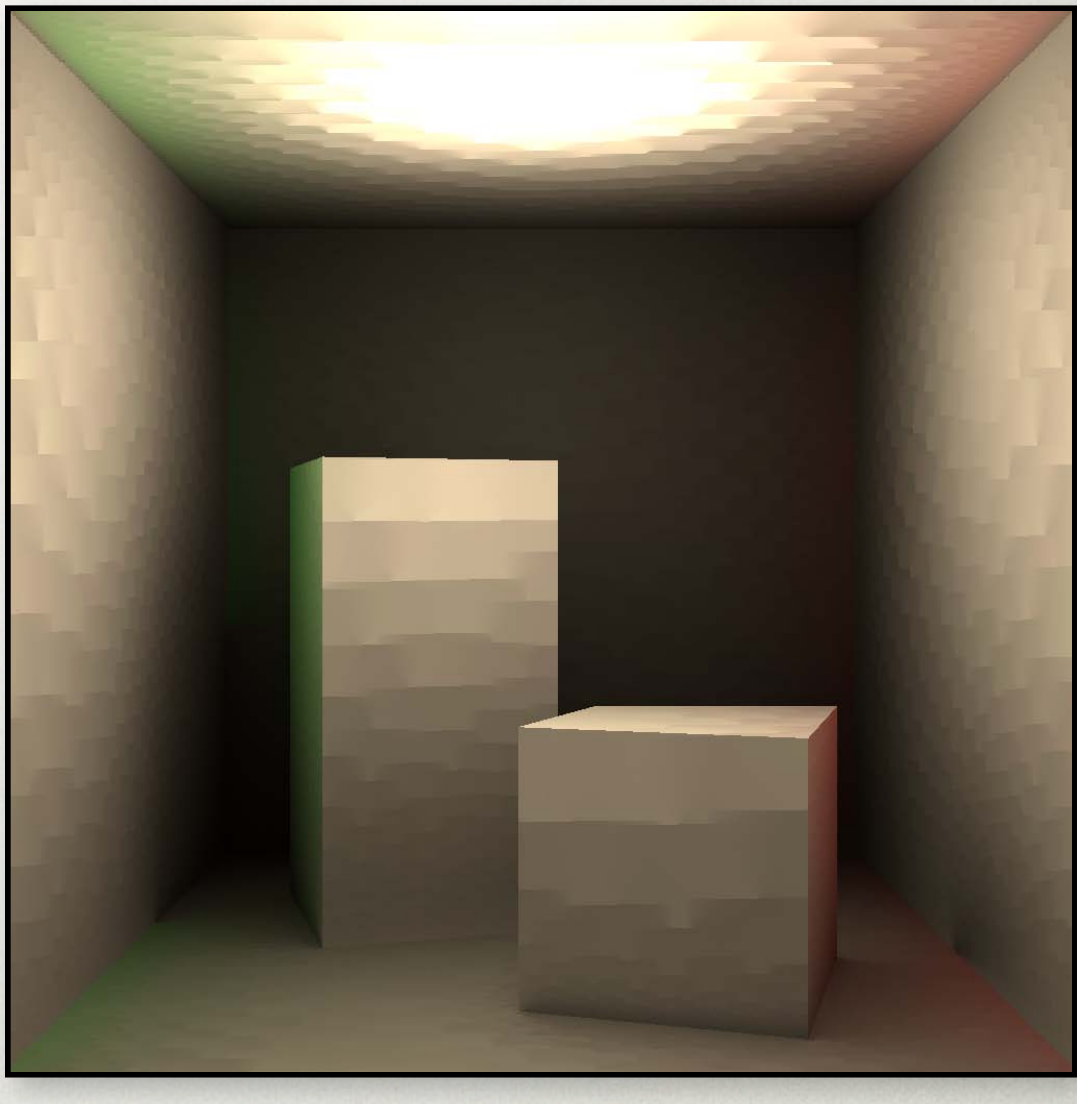

With media
(indirect irradiance)

* Irradiance caching with gradients can very effectively compute the indirect illumination if no media is present.
* However, since the gradient formulation does not account for media, significant artifacts appear if we apply these gradient computations when media is present
* The reason for this is that these scenes invalidate a major underlying assumption of irradiance gradients, which is, that surfaces are embedded within a vacuum.

Participating Media

No Media
(indirect irradiance)

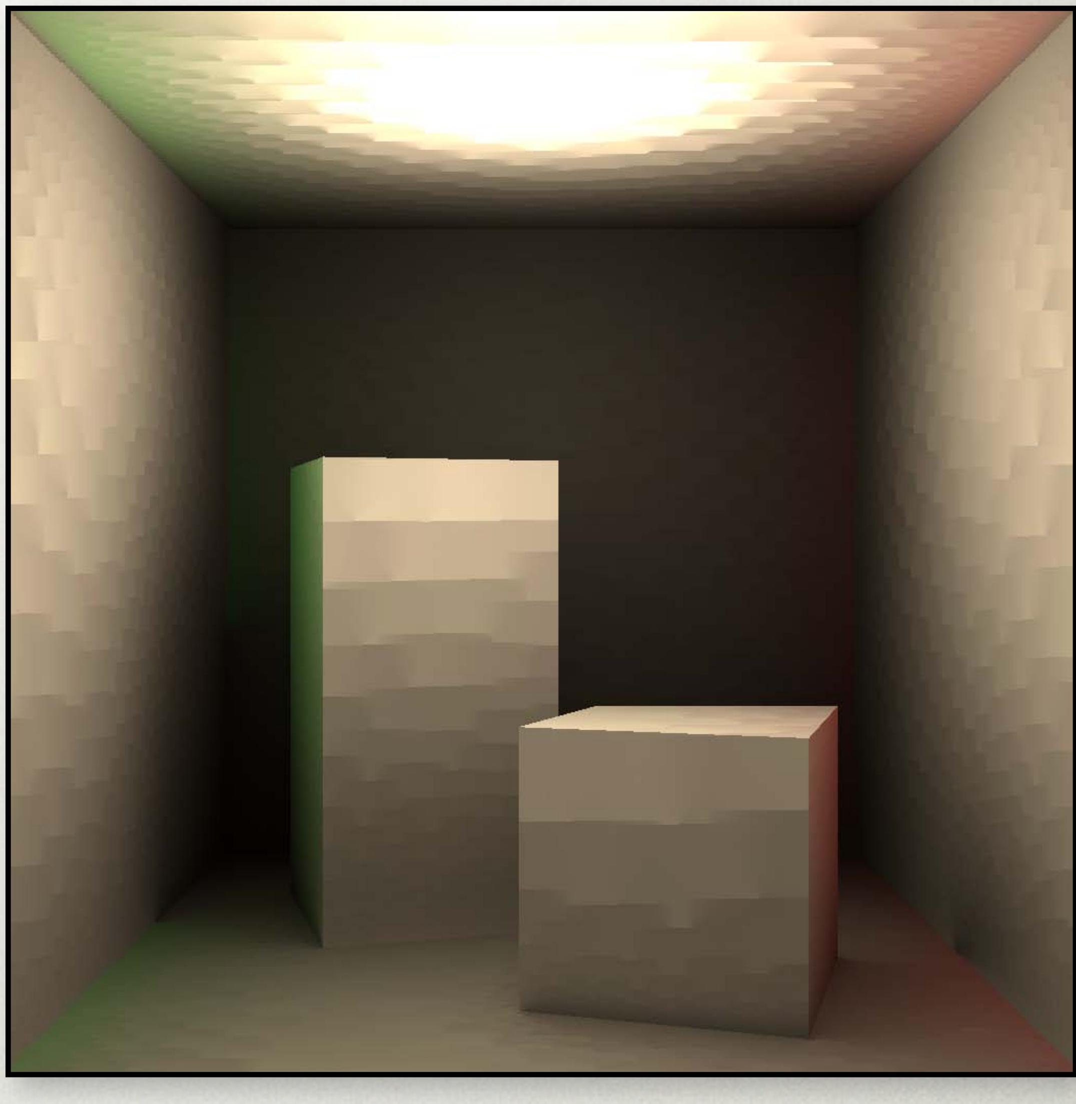

With media
(indirect irradiance)

Thursday, 6 September 12

* Irradiance caching with gradients can very effectively compute the indirect illumination if no media is present.
* However, since the gradient formulation does not account for media, significant artifacts appear if we apply these gradient computations when media is present
* The reason for this is that these scenes invalidate a major underlying assumption of irradiance gradients, which is, that surfaces are embedded within a vacuum.

Participating Media

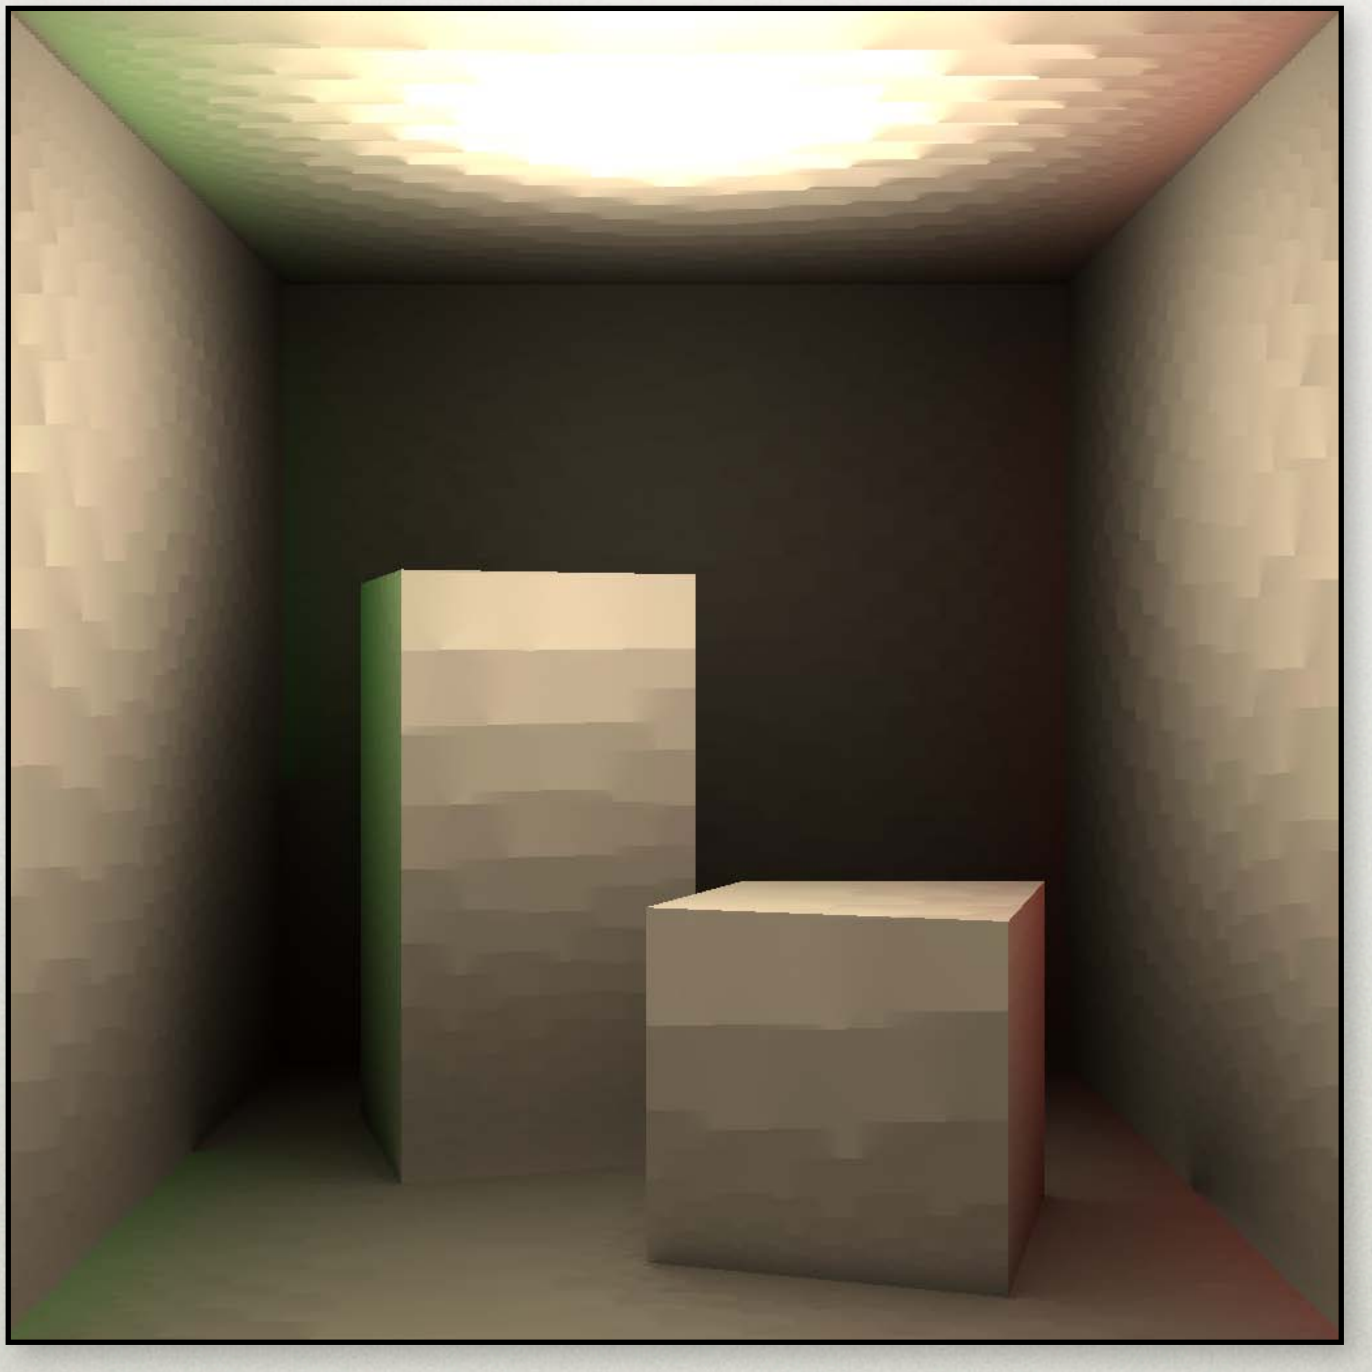

Ward and Heckbert

* Irradiance caching with gradients can very effectively compute the indirect illumination if no media is present.
* However, since the gradient formulation does not account for media, significant artifacts appear if we apply these gradient computations when media is present
* The reason for this is that these scenes invalidate a major underlying assumption of irradiance gradients, which is, that surfaces are embedded within a vacuum.

Participating Media

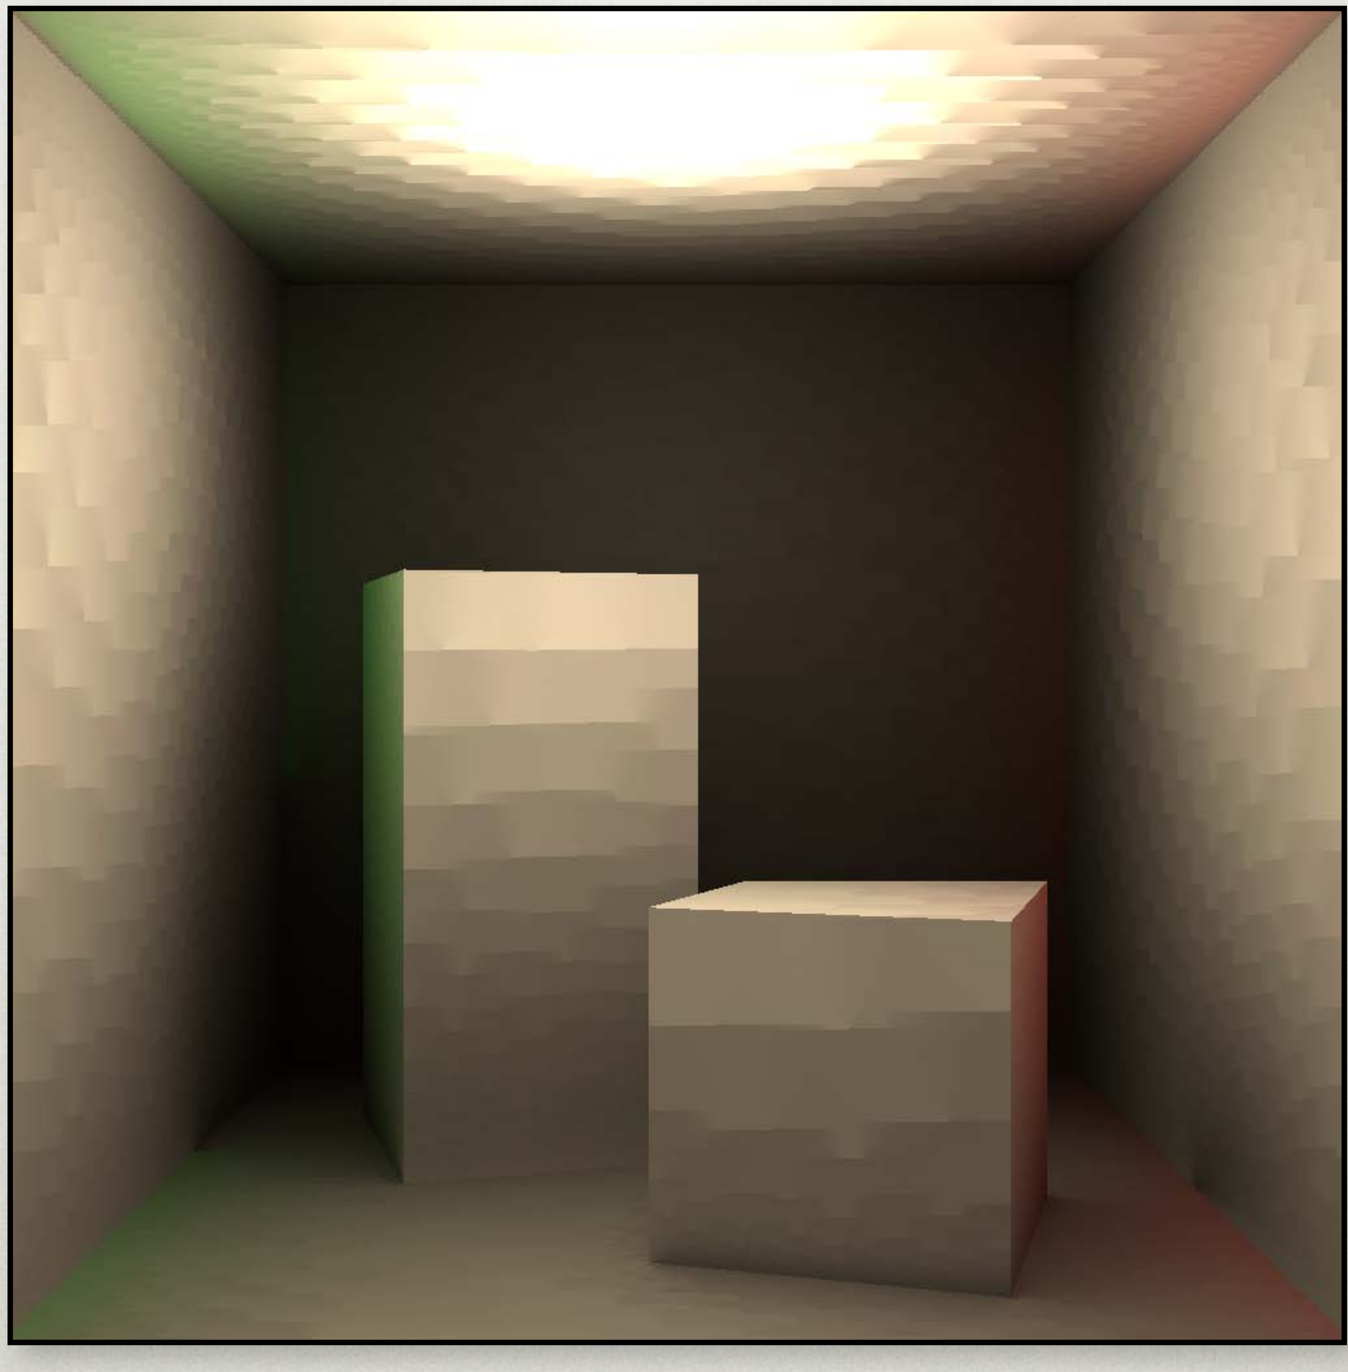

Our Gradients

Thursday, 6 September 12

* Using the techniques described in our paper, in the same amount of time, we are able to compute a much more accurate gradient which allows for higher quality interpolation.
* A key thing to note here is that the actual cache point locations are identical between these two images, just the gradient computation is changed.

Participating Media

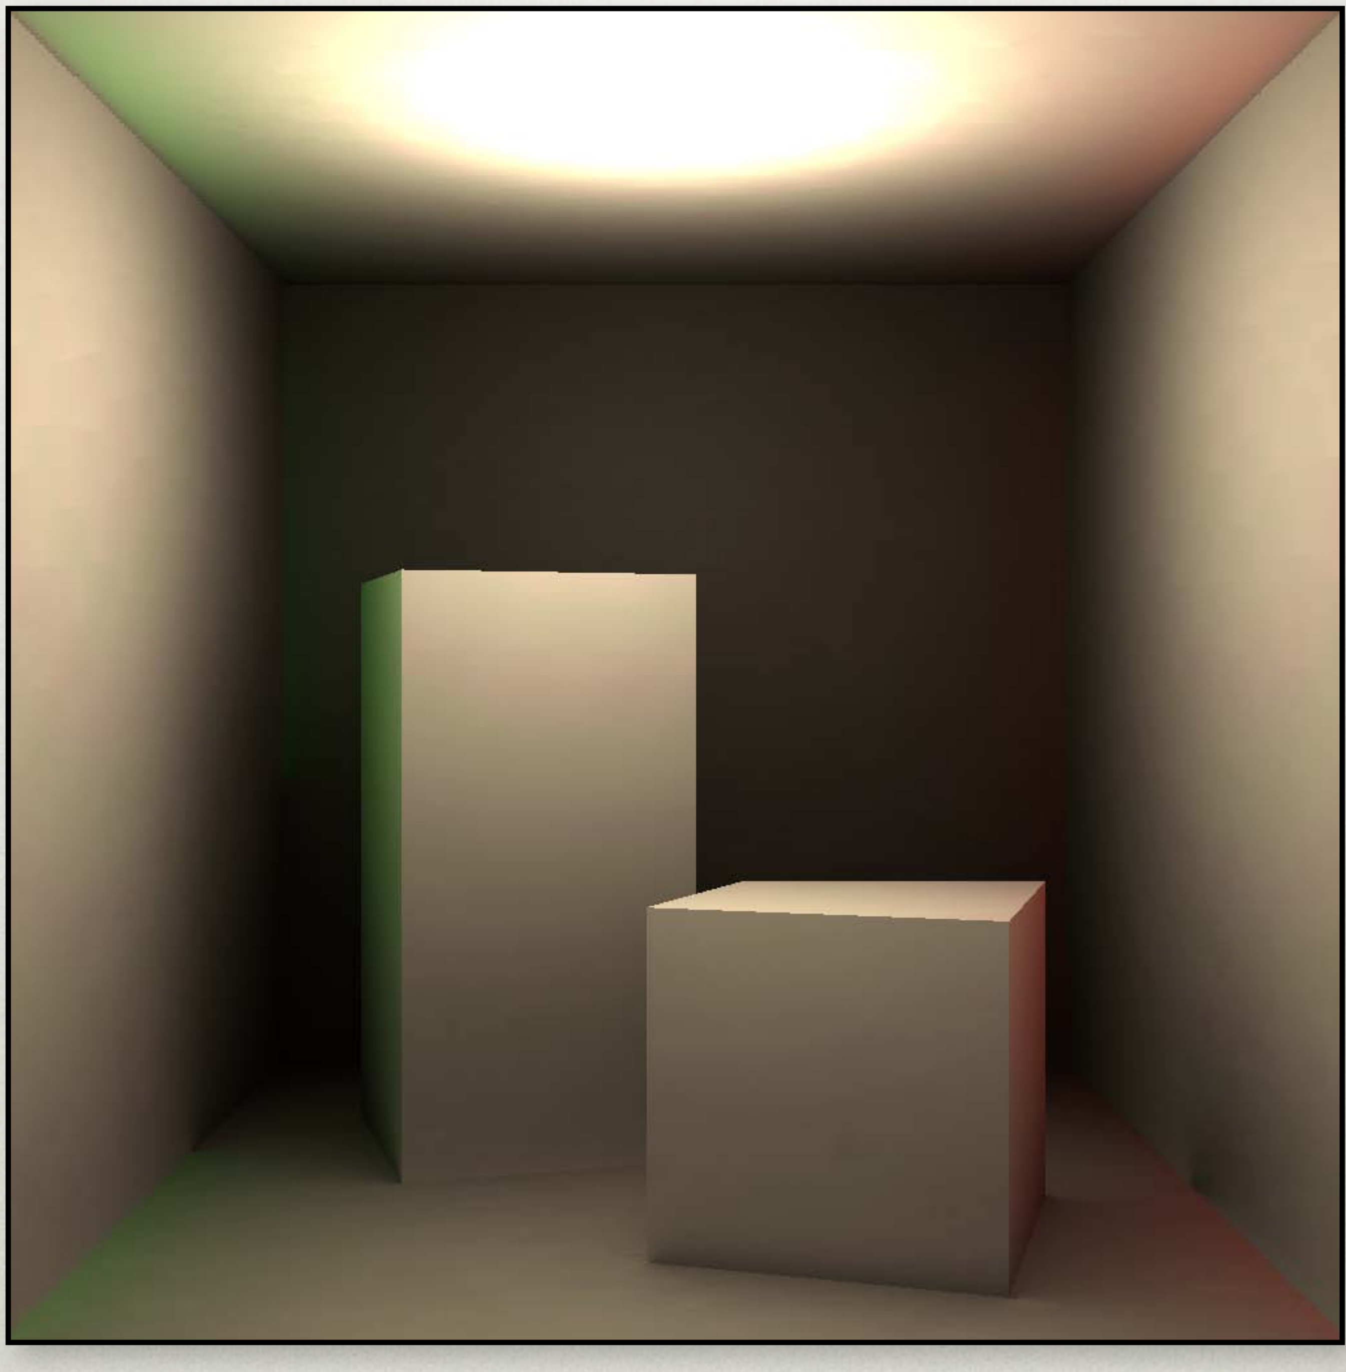

Our Gradients

Thursday, 6 September 12

* Using the techniques described in our paper, in the same amount of time, we are able to compute a much more accurate gradient which allows for higher quality interpolation.
* A key thing to note here is that the actual cache point locations are identical between these two images, just the gradient computation is changed.

Participating Media

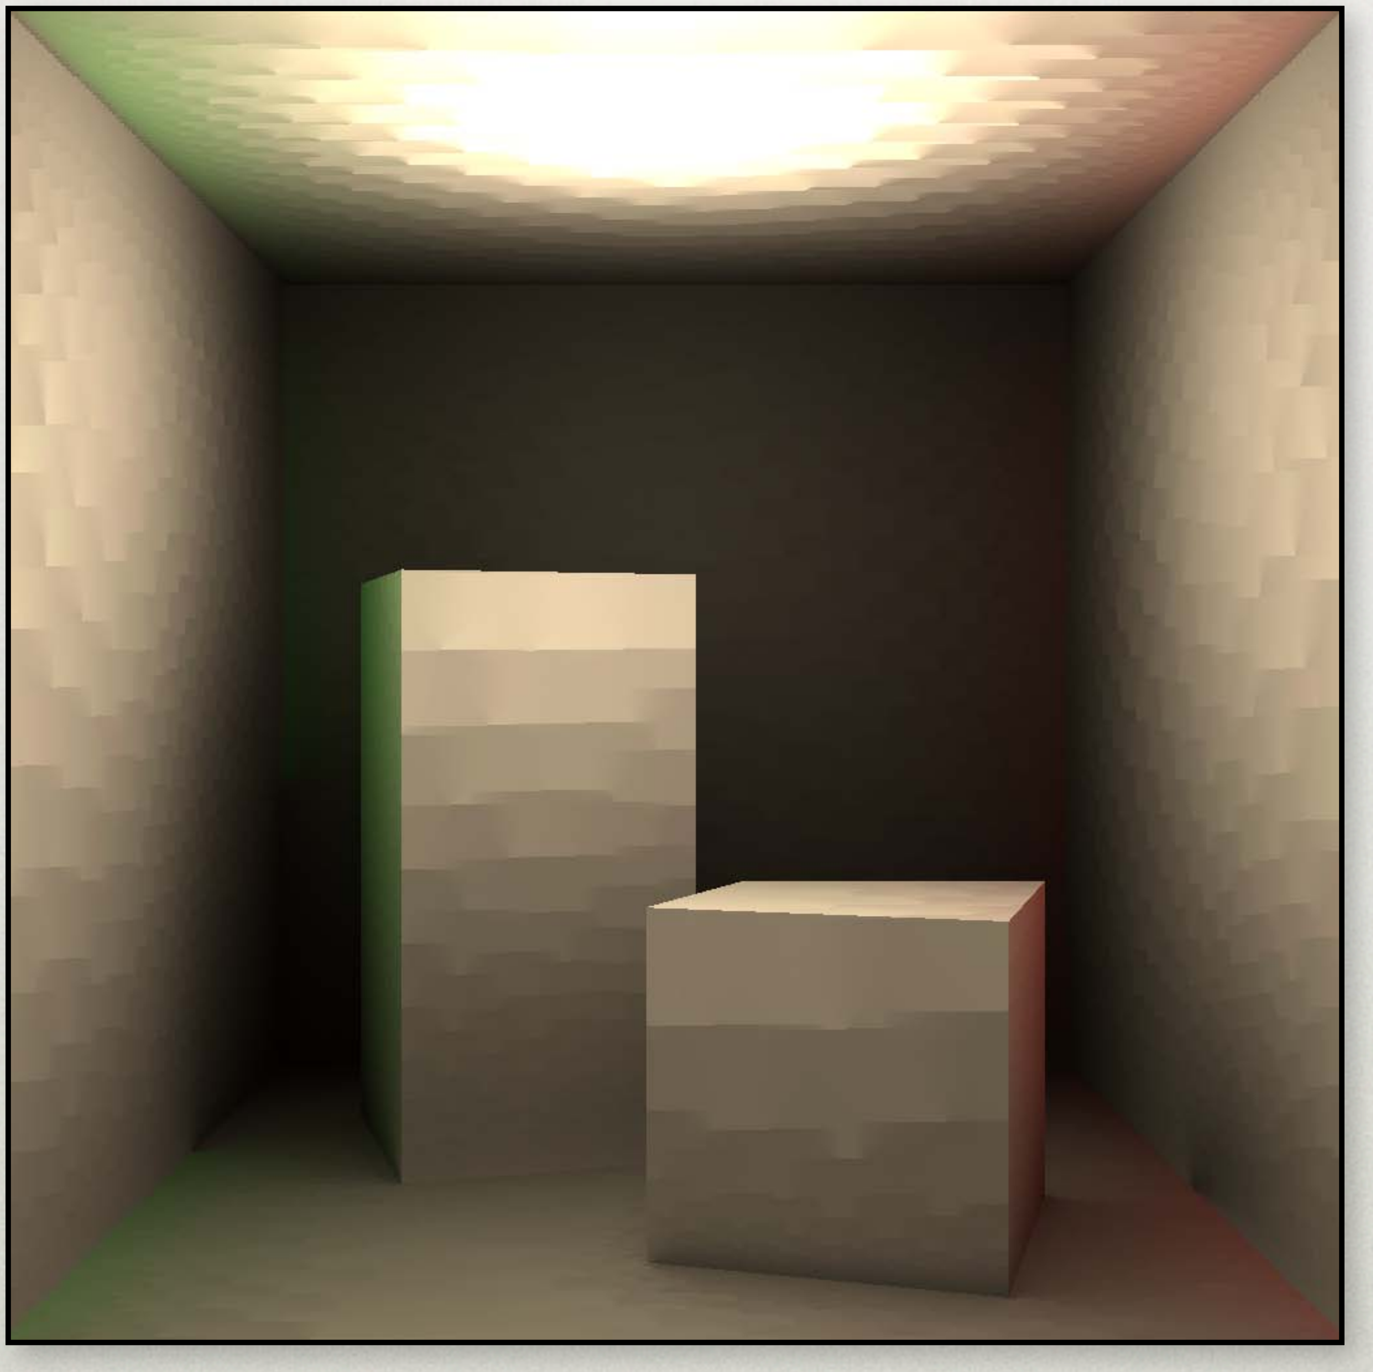

Ward and Heckbert Gradients

Thursday, 6 September 12

* Using the techniques described in our paper, in the same amount of time, we are able to compute a much more accurate gradient which allows for higher quality interpolation.
* A key thing to note here is that the actual cache point locations are identical between these two images, just the gradient computation is changed.

Volume Rendering Equation

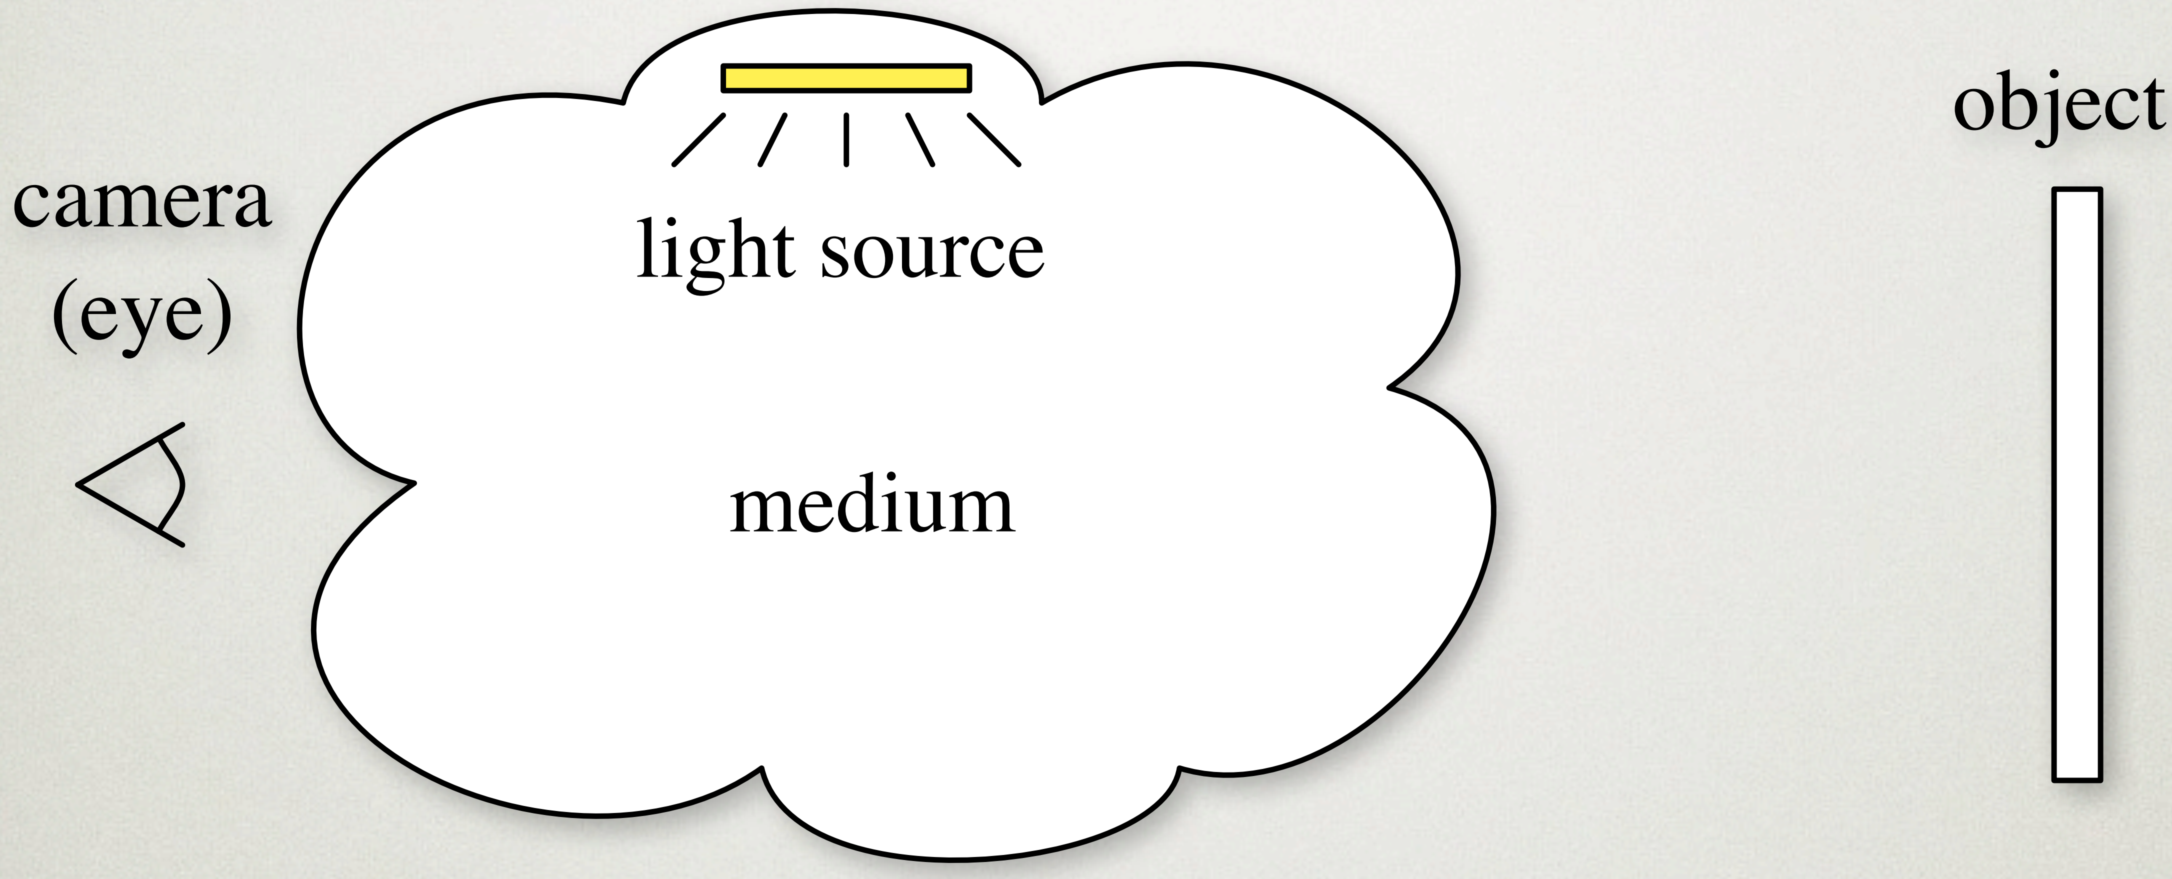

Thursday, 6 September 12

* In order to compute these gradients, we must first understand the behavior of light in the presence of media.

Volume Rendering Equation

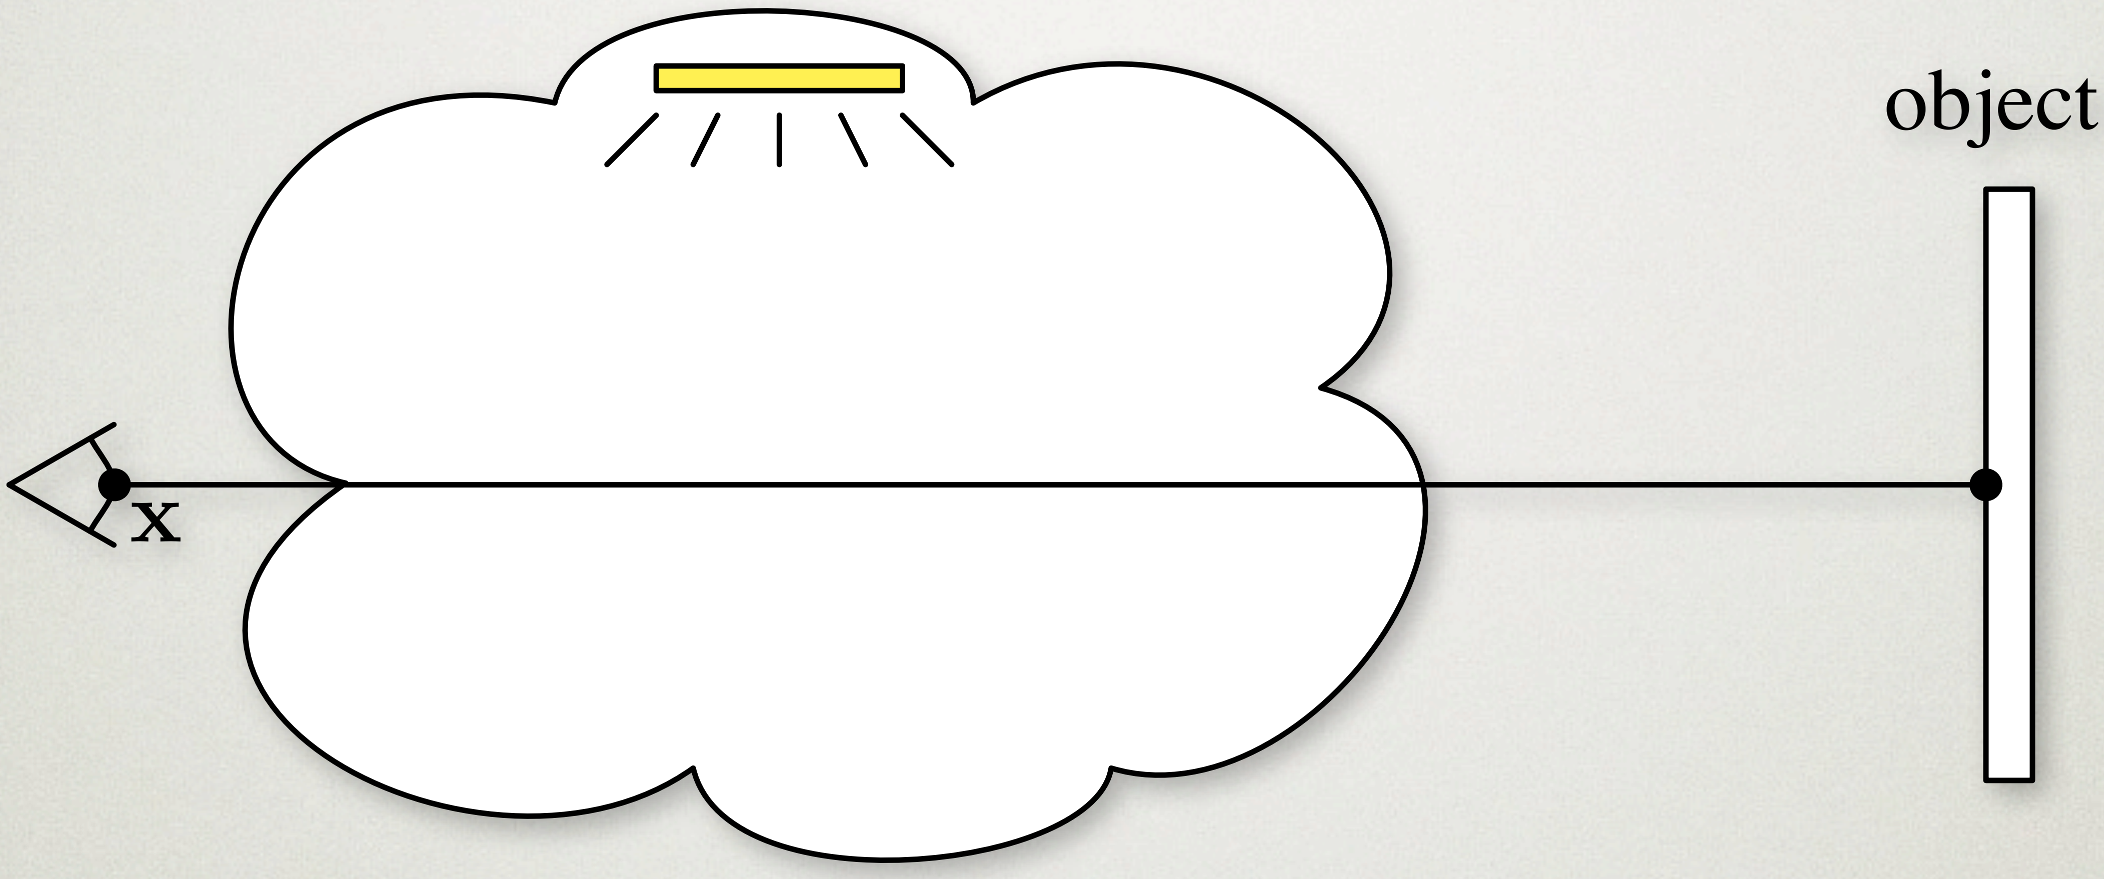

$$L(\mathbf{x}, \vec{\omega}) = L_m(\mathbf{x}, \vec{\omega}) + L_s(\mathbf{x}, \vec{\omega})$$

* The radiance, L, arriving at any location x along a ray can be expressed using the volume rendering equation.
* but at a high-level the meaning is pretty simple.
* In the presence of participating media, the radiance is the sum of two terms:

Volume Rendering Equation

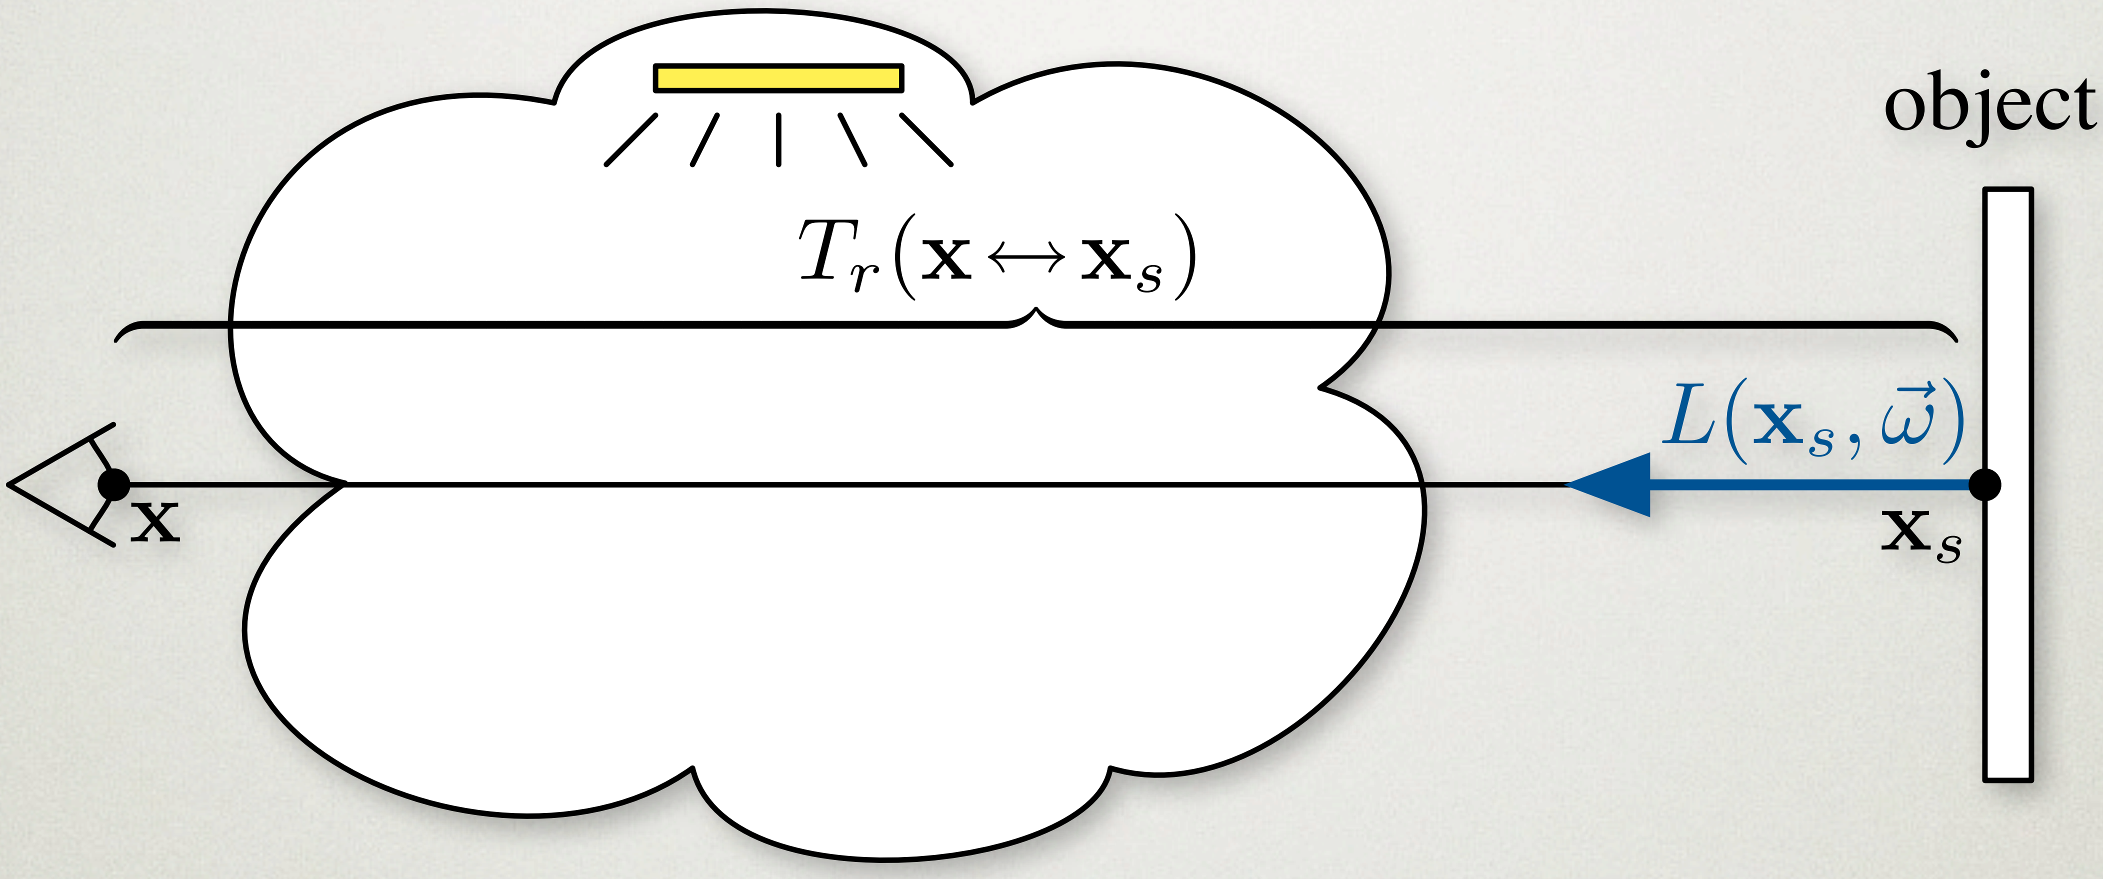

$$L(\mathbf{x}, \vec{\omega}) = L_m(\mathbf{x}, \vec{\omega}) + \boxed{L_s(\mathbf{x}, \vec{\omega})}$$

surface radiance

* the right-hand term incorporates lighting arriving from a surface
* before reaching the eye, this radiance must travel through the medium and so is attenuated by a transmission term

Volume Rendering Equation

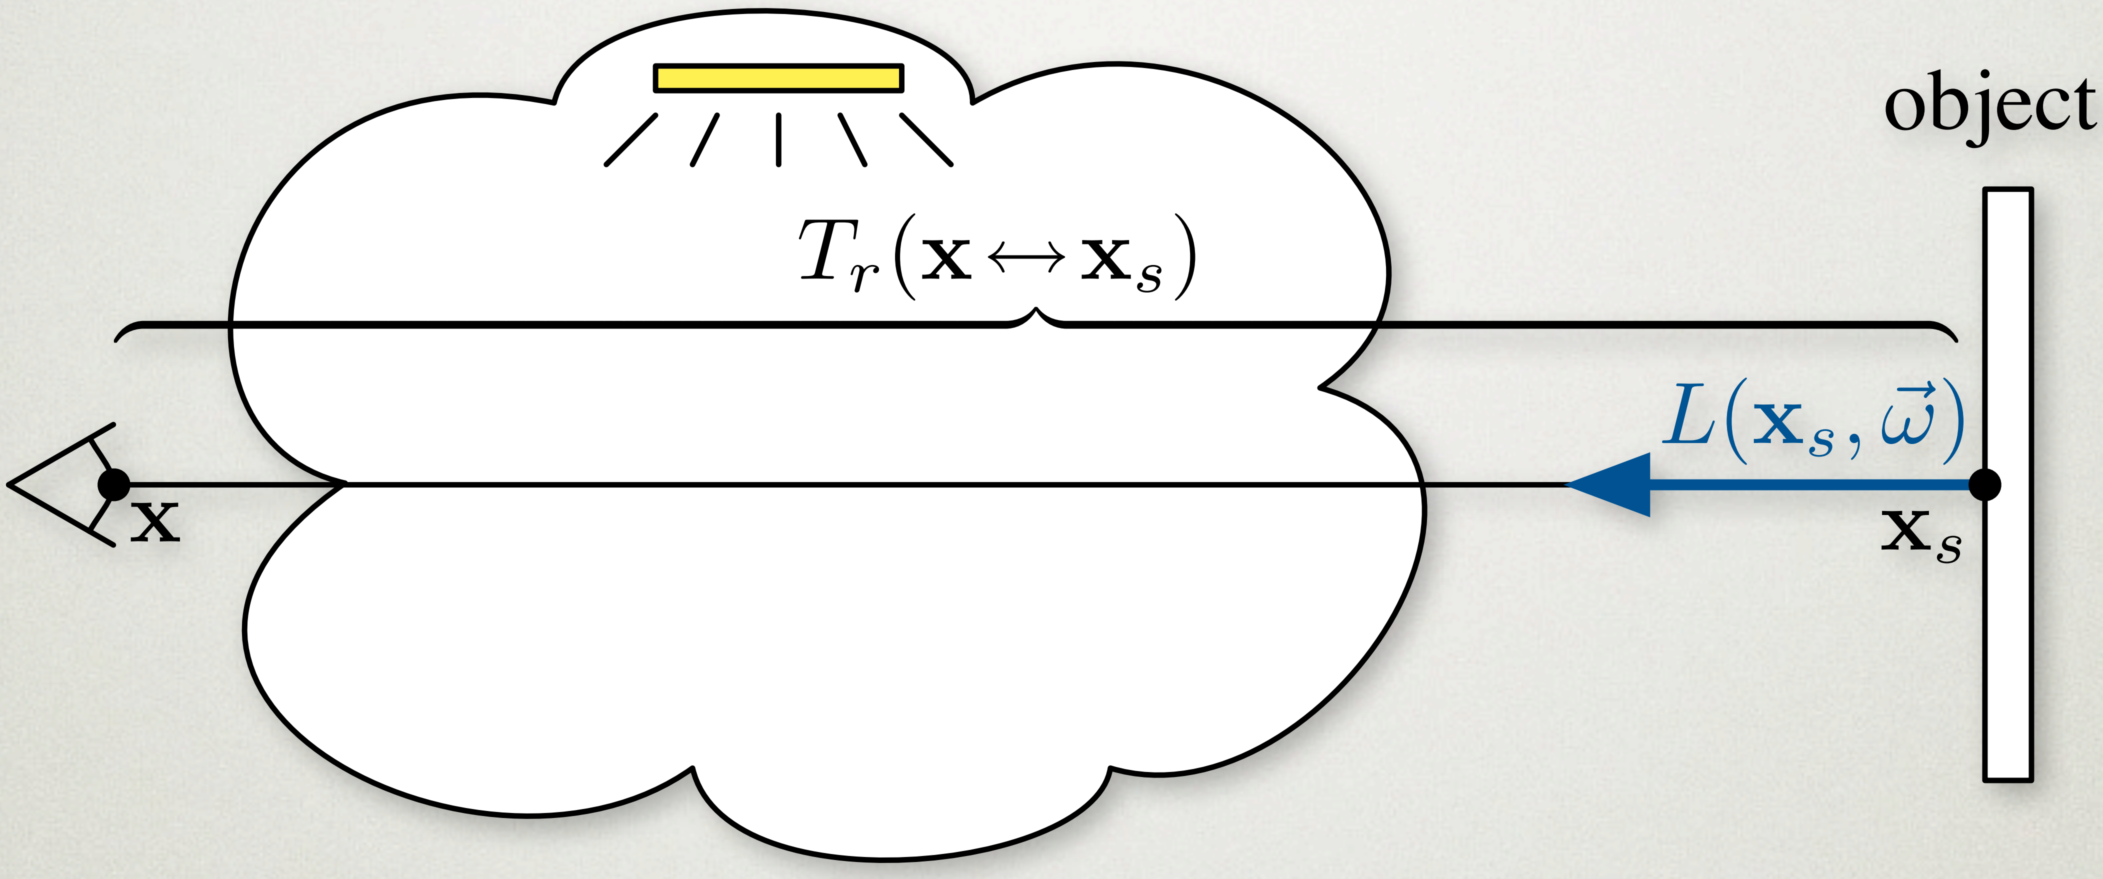

$$L(\mathbf{x}, \vec{\omega}) = L_m(\mathbf{x}, \vec{\omega}) + \boxed{L_s(\mathbf{x}, \vec{\omega})}$$

$$\boxed{L_s(\mathbf{x}, \vec{\omega})} = T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega})$$

surface radiance

Thursday, 6 September 12

* the right-hand term incorporates lighting arriving from a surface
* before reaching the eye, this radiance must travel through the medium and so is attenuated by a transmission term

Volume Rendering Equation

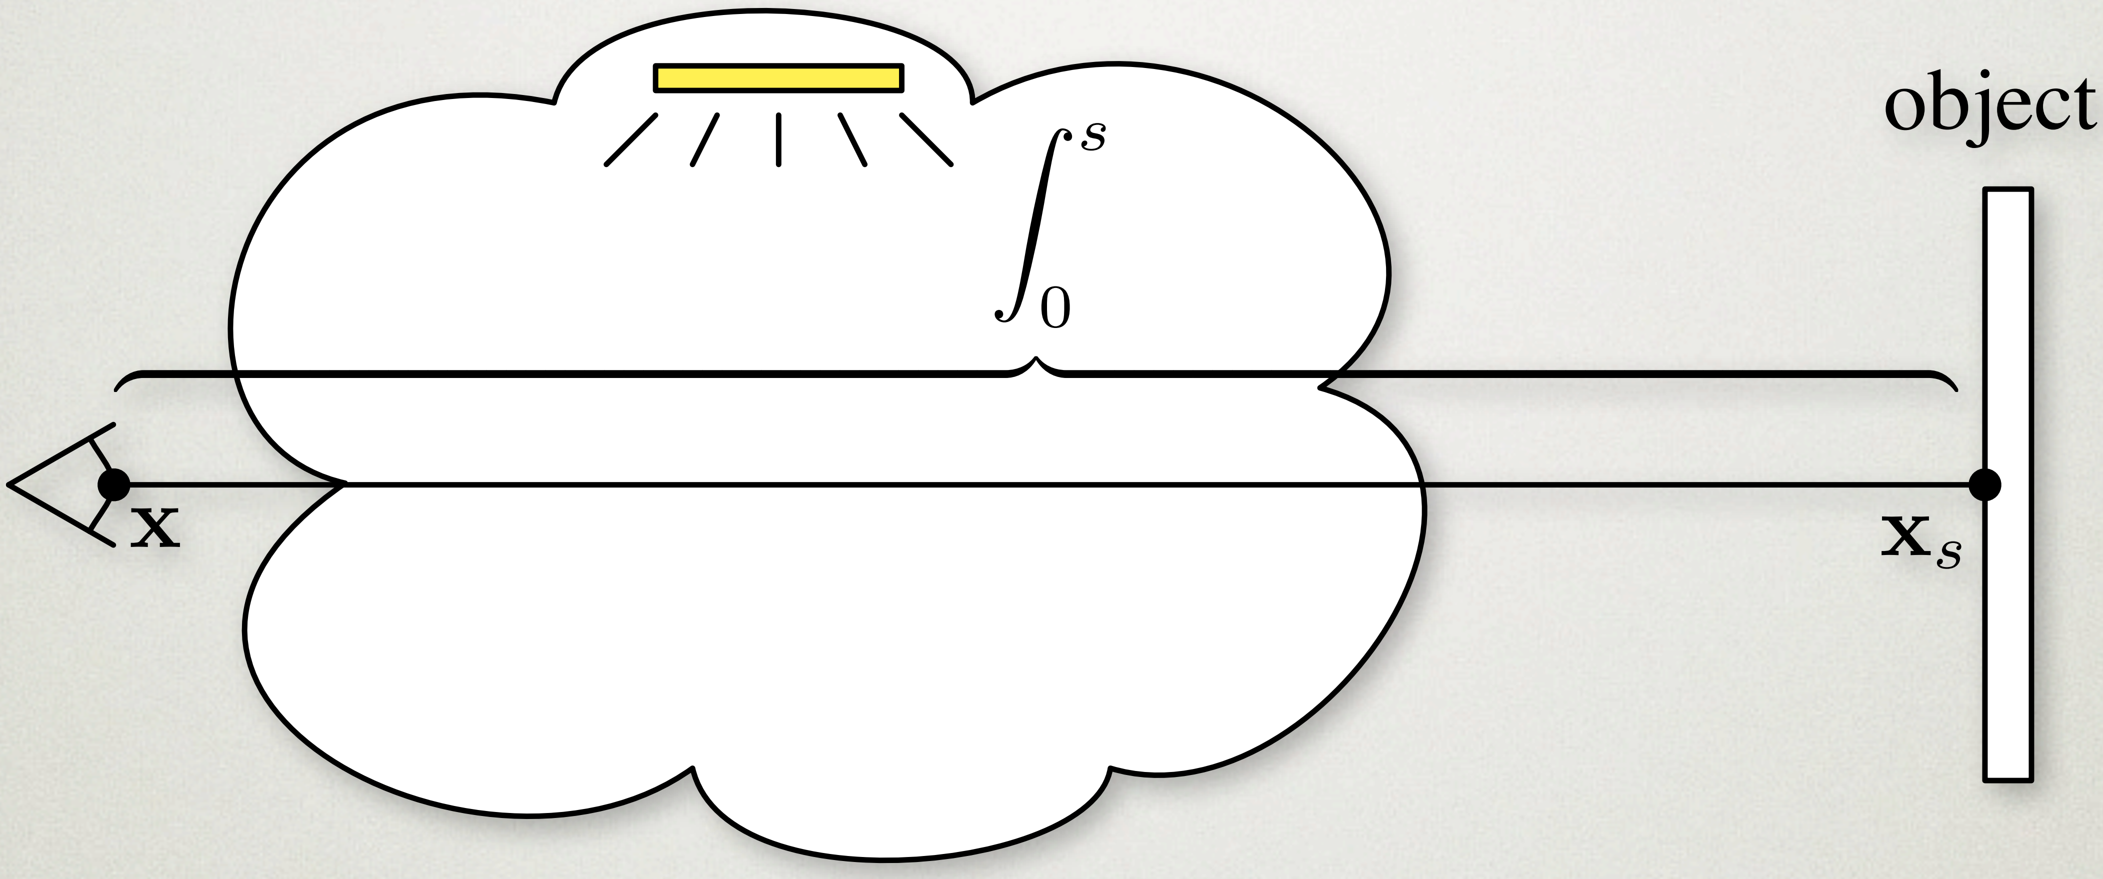

$$L(\mathbf{x}, \vec{\omega}) = L_m(\mathbf{x}, \vec{\omega}) + L_s(\mathbf{x}, \vec{\omega})$$

media radiance

* the left-hand term integrates the scattering of light from the medium along the whole length of the ray

Volume Rendering Equation

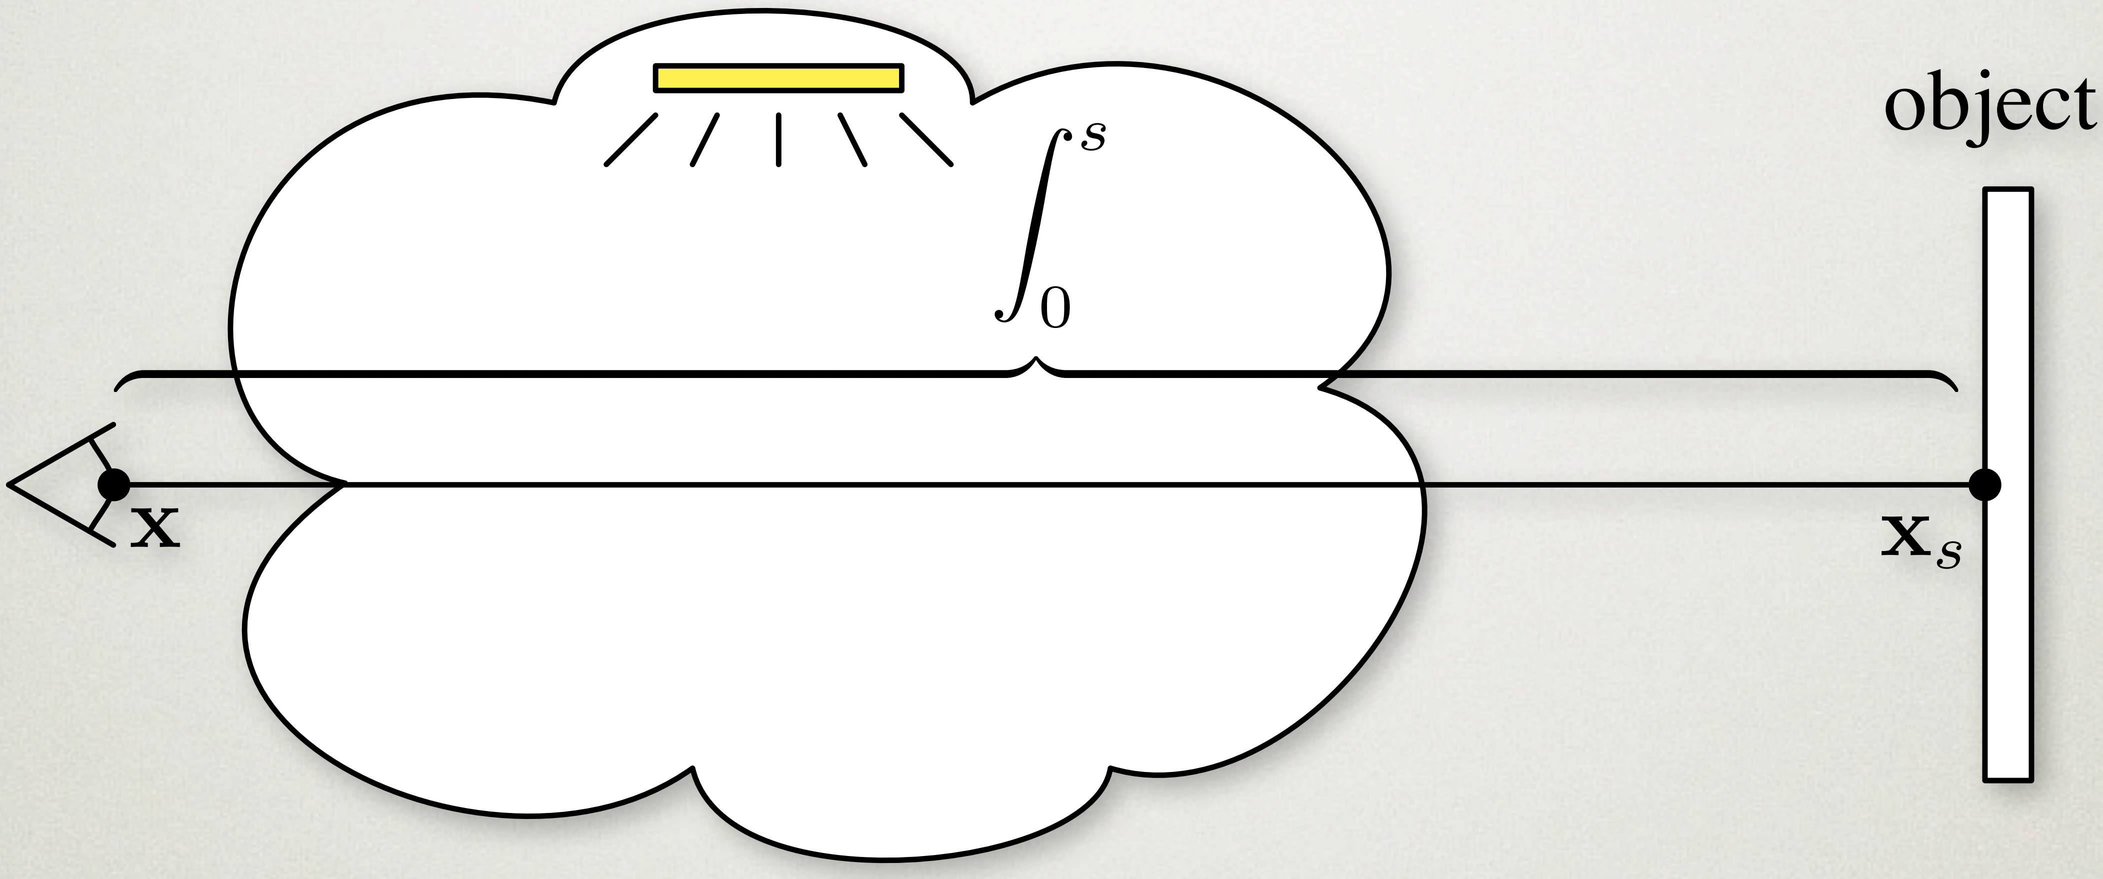

$$L(\mathbf{x},\vec{\omega}) = L_m(\mathbf{x},\vec{\omega}) + L_s(\mathbf{x},\vec{\omega})$$

$$L_m(\mathbf{x},\vec{\omega}) = \int_0^s T_r(\mathbf{x}\leftrightarrow\mathbf{x}_t)\sigma_s(\mathbf{x}_t)L_i(\mathbf{x}_t,\vec{\omega})\,dt$$

media radiance

* the left-hand term integrates the scattering of light from the medium along the whole length of the ray

Volume Rendering Equation

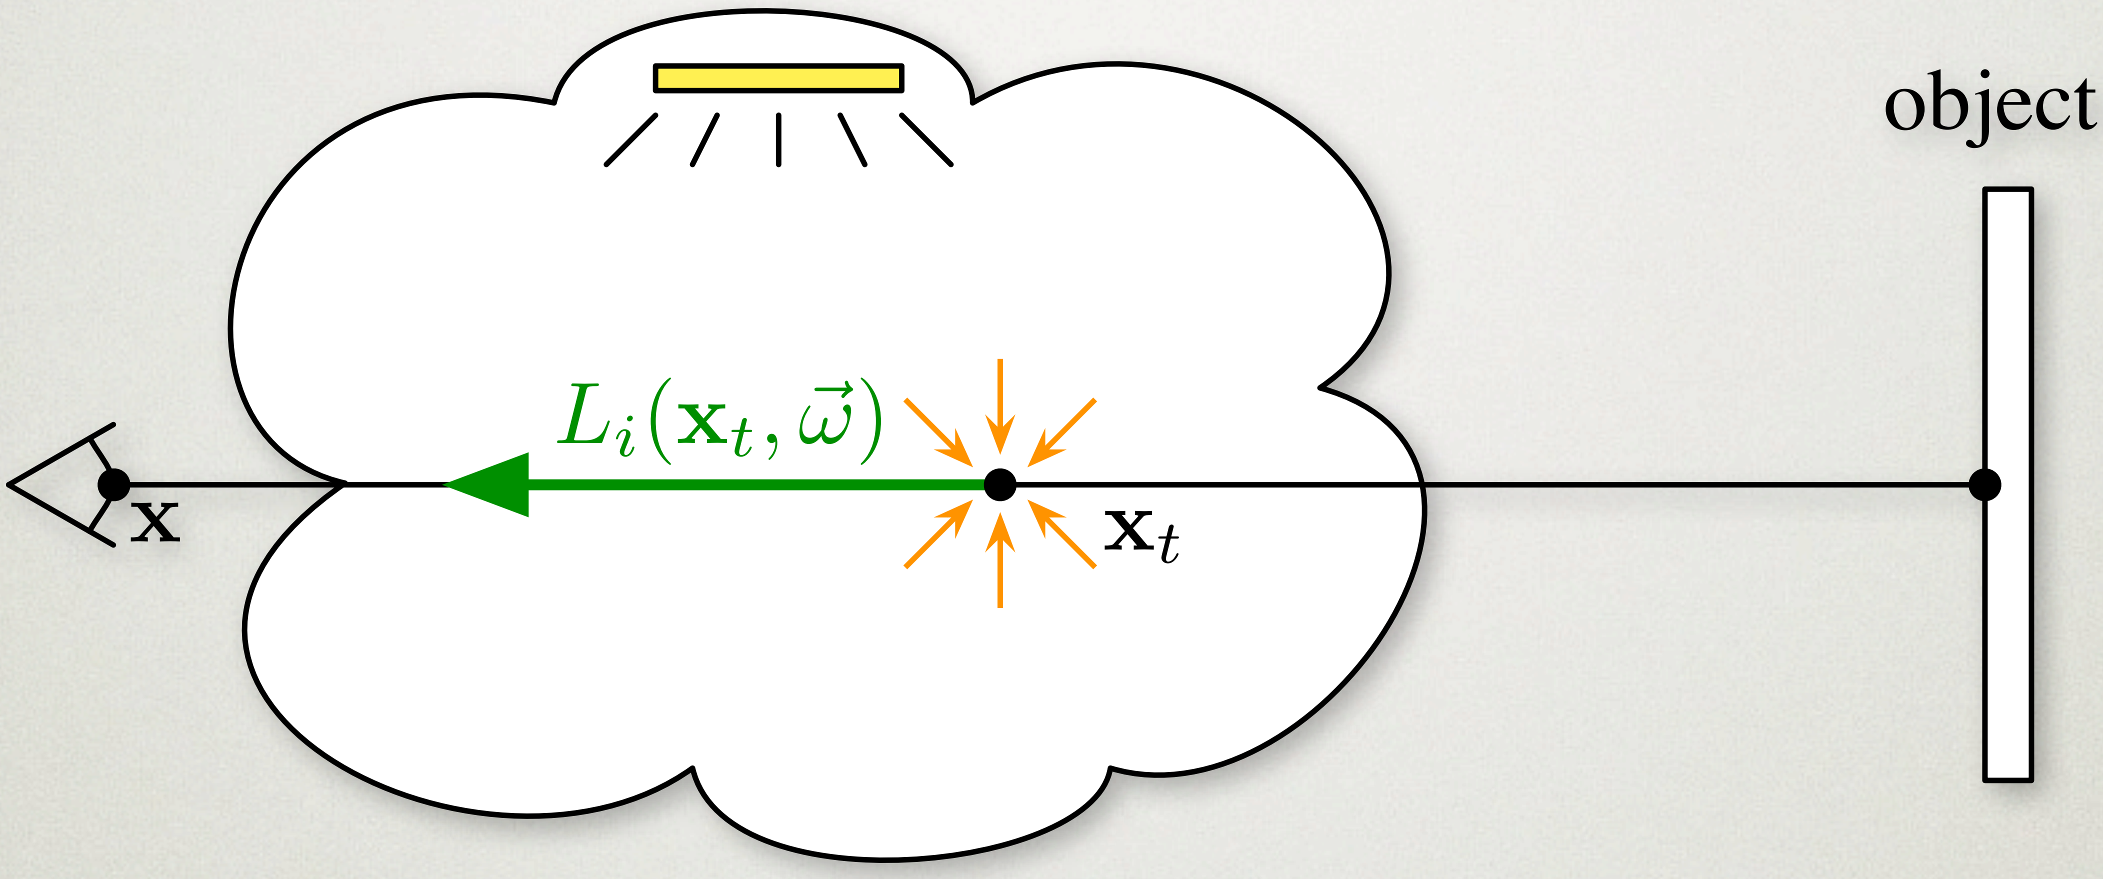

$$L(\mathbf{x}, \vec{\omega}) = L_m(\mathbf{x}, \vec{\omega}) + L_s(\mathbf{x}, \vec{\omega})$$

$$L_m(\mathbf{x}, \vec{\omega}) = \int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t)\sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega})\, dt$$

media radiance

Thursday, 6 September 12

* the main quantity that is integrated, Li, is inscattered radiance
* This represents the amount of light that reaches some point in the volume (from any other location in the scene), and then subsequently scatterers towards the eye

VOLUME RENDERING EQUATION

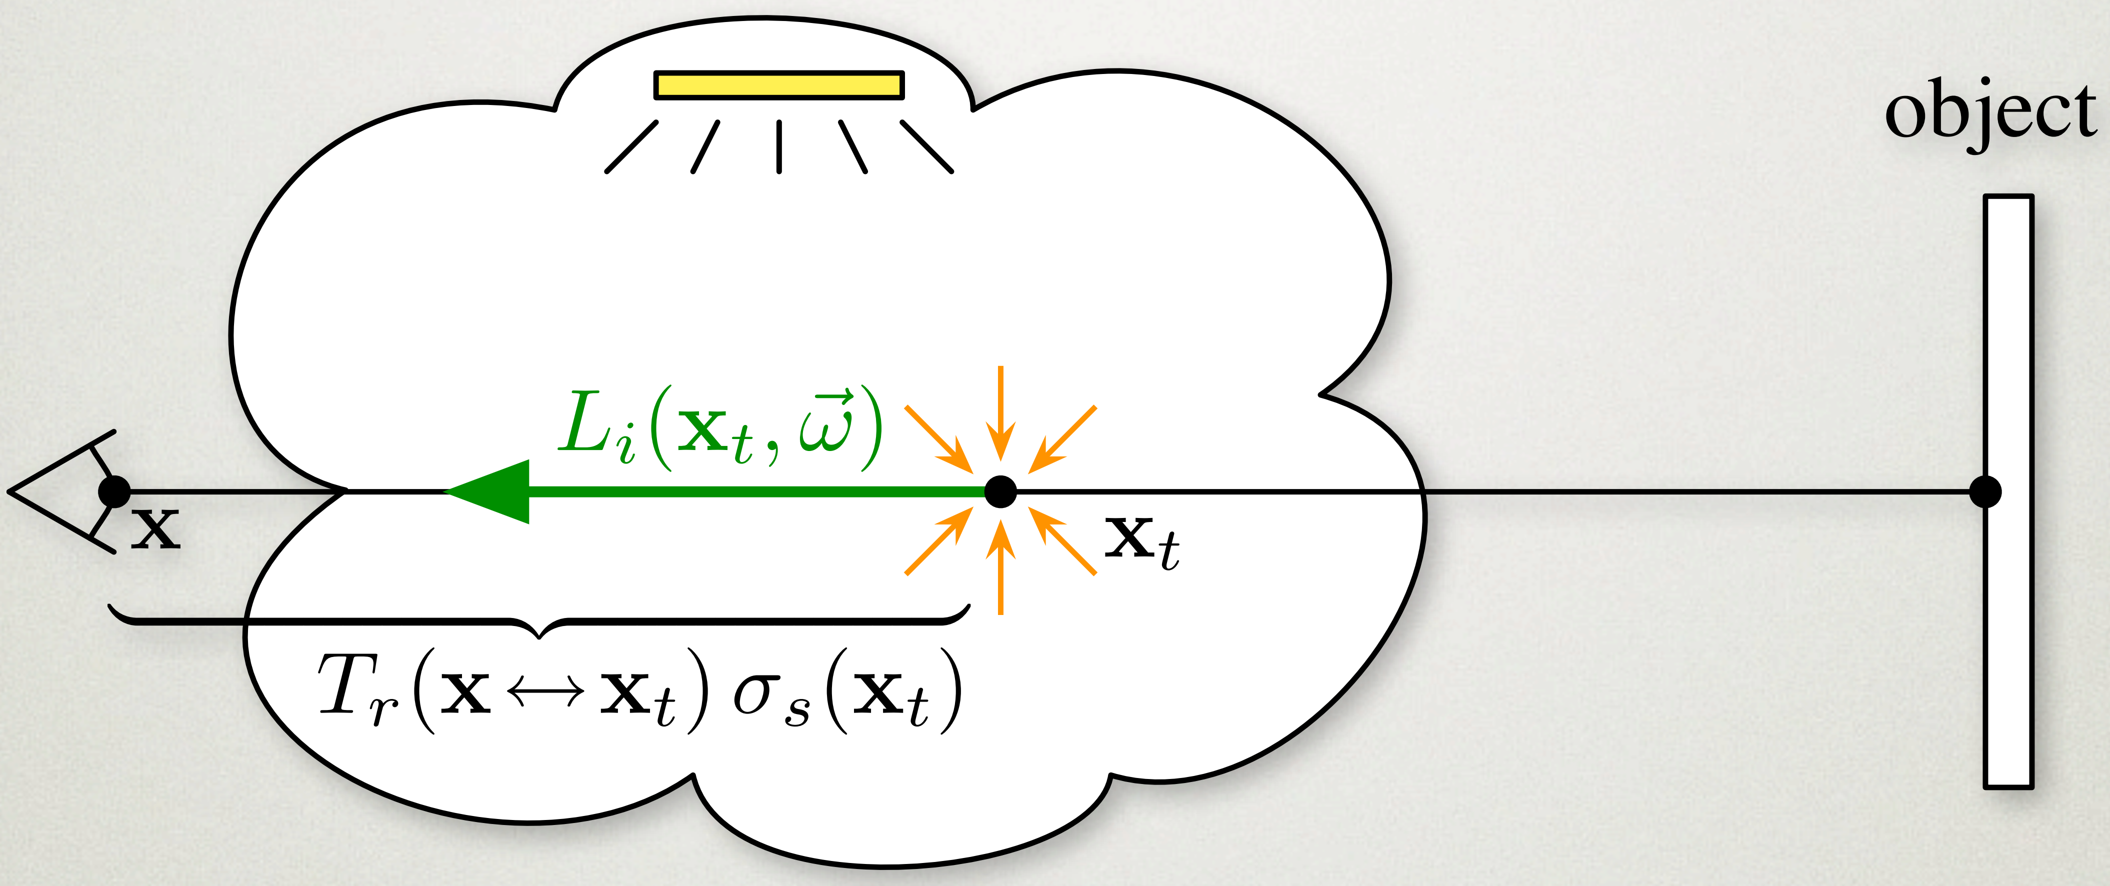

$$L(\mathbf{x}, \vec{\omega}) = L_m(\mathbf{x}, \vec{\omega}) + L_s(\mathbf{x}, \vec{\omega})$$

$$L_m(\mathbf{x}, \vec{\omega}) = \int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t)\sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega})\, dt$$

media radiance

Thursday, 6 September 12

* as this scattered light travels towards the eye it is also dissipated by extinction through the medium
* this computation is very expensive and there has been a lot of work on how to solve this problem efficiently

Contribution

- Compute translational gradients of irradiance in the presence of media

Contribution

- Compute translational gradients of irradiance in the presence of media
  - **Absorbing** media
  - **Emissive / scattering** media

Contribution

- Compute translational gradients of irradiance in the presence of media
  - **Absorbing** media
  - **Emissive**/**scattering** media
- Higher quality irradiance interpolation

Irradiance in Part. Med.

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

* Irradiance is simply the integral of the cosine weighted radiance over the hemisphere
* Since we decomposed the definition of radiance as radiance coming from surfaces and radiance coming from the media, we can perform the same decomposition on the hemispherical integral.

Irradiance in Part. Med.

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

$$L(\mathbf{x}, \vec{\omega}) = L_m(\mathbf{x}, \vec{\omega}) + L_s(\mathbf{x}, \vec{\omega})$$

* Irradiance is simply the integral of the cosine weighted radiance over the hemisphere
* Since we decomposed the definition of radiance as radiance coming from surfaces and radiance coming from the media, we can perform the same decomposition on the hemispherical integral.

Irradiance in Part. Med.

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

$$L(\mathbf{x}, \vec{\omega}) = L_m(\mathbf{x}, \vec{\omega}) + L_s(\mathbf{x}, \vec{\omega})$$

$$E(\mathbf{x}) = E_m(\mathbf{x}) + E_s(\mathbf{x})$$

$$E_s(\mathbf{x}) = \int_{\Omega} L_s(\mathbf{x}, \vec{\omega})\,(\vec{\mathbf{n}} \cdot \vec{\omega})\; d\vec{\omega}$$

$$E_m(\mathbf{x}) = \int_{\Omega} L_m(\mathbf{x}, \vec{\omega})\,(\vec{\mathbf{n}} \cdot \vec{\omega})\; d\vec{\omega}$$

* Irradiance is simply the integral of the cosine weighted radiance over the hemisphere
* Since we decomposed the definition of radiance as radiance coming from surfaces and radiance coming from the media, we can perform the same decomposition on the hemispherical integral.

Irradiance in Part. Med.

$$E(\mathbf{x}) = E_m(\mathbf{x}) + E_s(\mathbf{x})$$

* Irradiance is simply the integral of the cosine weighted radiance over the hemisphere
* Since we decomposed the definition of radiance as radiance coming from surfaces and radiance coming from the media, we can perform the same decomposition on the hemispherical integral.

Irradiance Gradient

$$E(\mathbf{x}) = E_m(\mathbf{x}) + E_s(\mathbf{x})$$

$$\nabla E(\mathbf{x}) = \nabla E_m(\mathbf{x}) + \nabla E_s(\mathbf{x})$$

Thursday, 6 September 12

* Since the total irradiance is the sum of two terms, the total irradiance gradient is just the sum of two gradient terms.
* The right hand term is the gradient due to surface irradiance and the left is the gradient due to media irradiance.
* In the remainder of the talk I will describe how we compute the two irradiance values and their corresponding gradients.

Irradiance Gradient

$$E(\mathbf{x}) = E_m(\mathbf{x}) + E_s(\mathbf{x})$$

$$\nabla E(\mathbf{x}) = \nabla E_m(\mathbf{x}) + \nabla E_s(\mathbf{x})$$

Gradient from Surfaces

* Since the total irradiance is the sum of two terms, the total irradiance gradient is just the sum of two gradient terms.
* The right hand term is the gradient due to surface irradiance and the left is the gradient due to media irradiance.
* In the remainder of the talk I will describe how we compute the two irradiance values and their corresponding gradients.

Irradiance Gradient

$$E(\mathbf{x}) = E_m(\mathbf{x}) + E_s(\mathbf{x})$$

$$\nabla E(\mathbf{x}) = \nabla E_m(\mathbf{x}) + \nabla E_s(\mathbf{x})$$

Gradient from Media

Gradient from Surfaces

* Since the total irradiance is the sum of two terms, the total irradiance gradient is just the sum of two gradient terms.
* The right hand term is the gradient due to surface irradiance and the left is the gradient due to media irradiance.
* In the remainder of the talk I will describe how we compute the two irradiance values and their corresponding gradients.

Irradiance From Surfaces

$$E_s(\mathbf{x}) = \int_\Omega L_s(\mathbf{x}, \vec{\omega})\,(\vec{\mathbf{n}} \cdot \vec{\omega})\; d\vec{\omega}$$

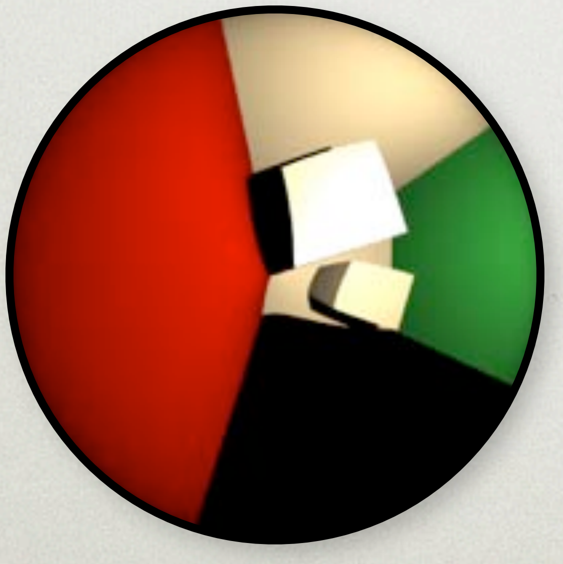

* Given the definition of the surface irradiance, we can estimate it by performing a stratified Monte Carlo integration.
* This involves subdividing the hemisphere of directions into a number of strata, or cells, and sampling the radiance using a jittered sample within each cell.
* The irradiance is just the sum of all the radiance samples weighted by their cell area and the cosine term.
* This is exactly the approach used by standard irradiance caching techniques.

Irradiance From Surfaces

$$E_s(\mathbf{x}) = \int_\Omega L_s(\mathbf{x}, \vec{\omega})\,(\vec{\mathbf{n}} \cdot \vec{\omega})\,d\vec{\omega}$$

$$E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} A_{j,k}\, L_s(\mathbf{x}, \vec{\omega}_{j,k})\,(\vec{n} \cdot \vec{\omega}_{j,k})$$

* Given the definition of the surface irradiance, we can estimate it by performing a stratified Monte Carlo integration.
* This involves subdividing the hemisphere of directions into a number of strata, or cells, and sampling the radiance using a jittered sample within each cell.
* The irradiance is just the sum of all the radiance samples weighted by their cell area and the cosine term.
* This is exactly the approach used by standard irradiance caching techniques.

Irradiance From Surfaces

$$E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} A_{j,k} \, L_s(\mathbf{x}, \vec{\omega}_{j,k}) \, (\vec{n} \cdot \vec{\omega}_{j,k})$$

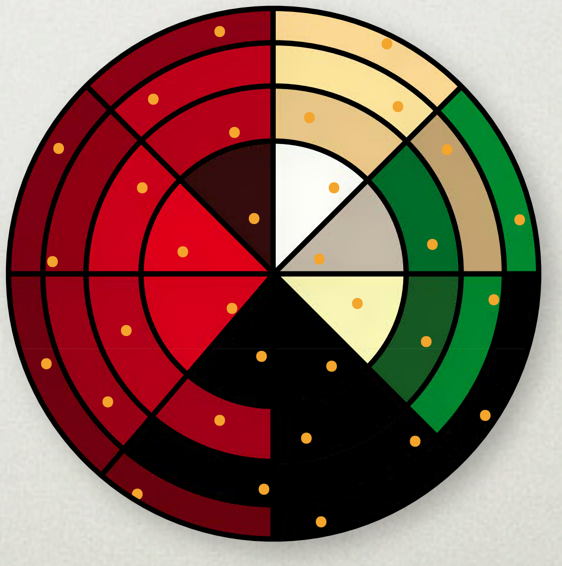

Irradiance From Surfaces

$$E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} A_{j,k}\, L_s(\mathbf{x}, \vec{\omega}_{j,k})\, (\vec{n} \cdot \vec{\omega}_{j,k})$$

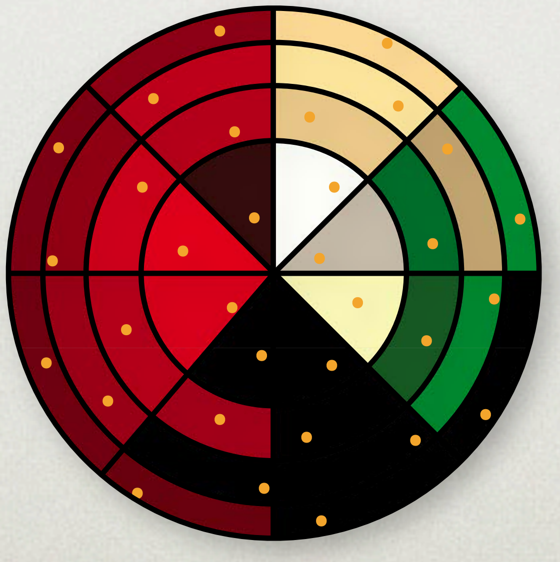

Irradiance Gradient From Surfaces

$$E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} A_{j,k}\, L_s(\mathbf{x}, \vec{\omega}_{j,k})\, (\vec{n} \cdot \vec{\omega}_{j,k})$$

$$\nabla E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} (\nabla A_{j,k}\, L^s_{j,k} + A_{j,k}\, \nabla L^s_{j,k})\, (\vec{n} \cdot \vec{\omega}_{j,k})$$

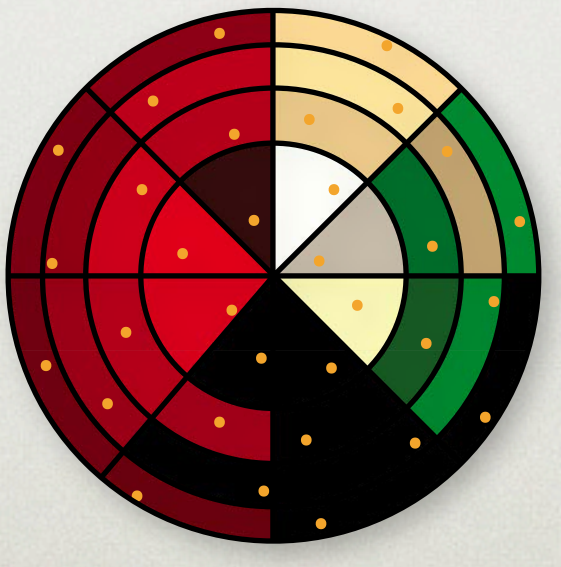

Thursday, 6 September 12

* We can compute the translational gradient of this estimate by using the product rule within the summation.

Irradiance Gradient From Surfaces

$$\nabla E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} (\nabla A_{j,k}\, L^s_{j,k} + A_{j,k}\, \nabla L^s_{j,k})\, (\vec{n} \cdot \vec{\omega}_{j,k})$$

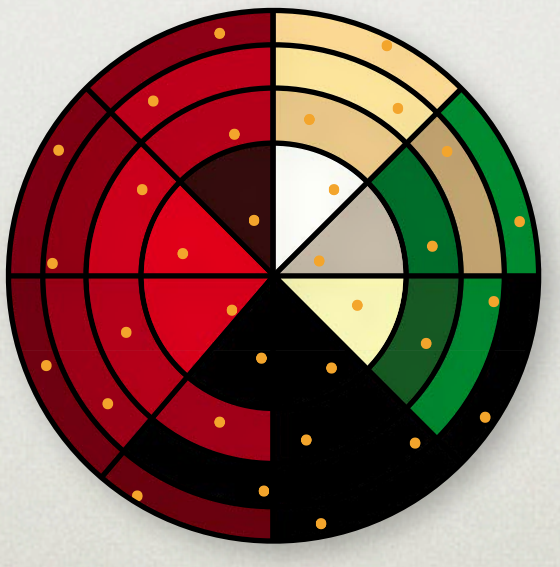

Irradiance Gradient From Surfaces

$$\nabla E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} (\nabla A_{j,k} \, L^s_{j,k} + A_{j,k} \, \nabla L^s_{j,k}) \, (\vec{n} \cdot \vec{\omega}_{j,k})$$

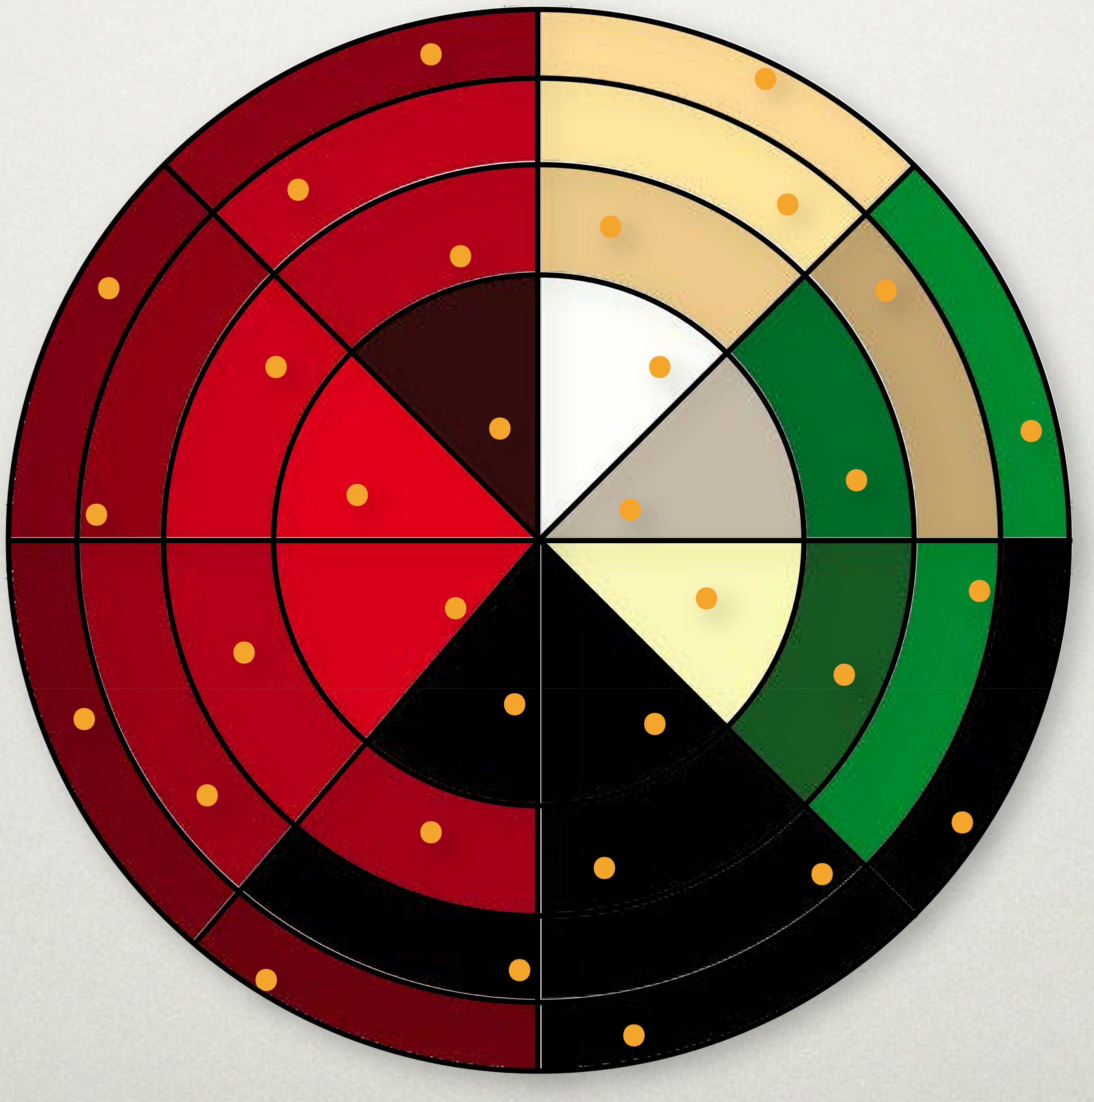

Irradiance Gradient From Surfaces

$$\nabla E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} (\nabla A_{j,k}\, L^s_{j,k} + A_{j,k}\, \nabla L^s_{j,k})\, (\vec{n} \cdot \vec{\omega}_{j,k})$$

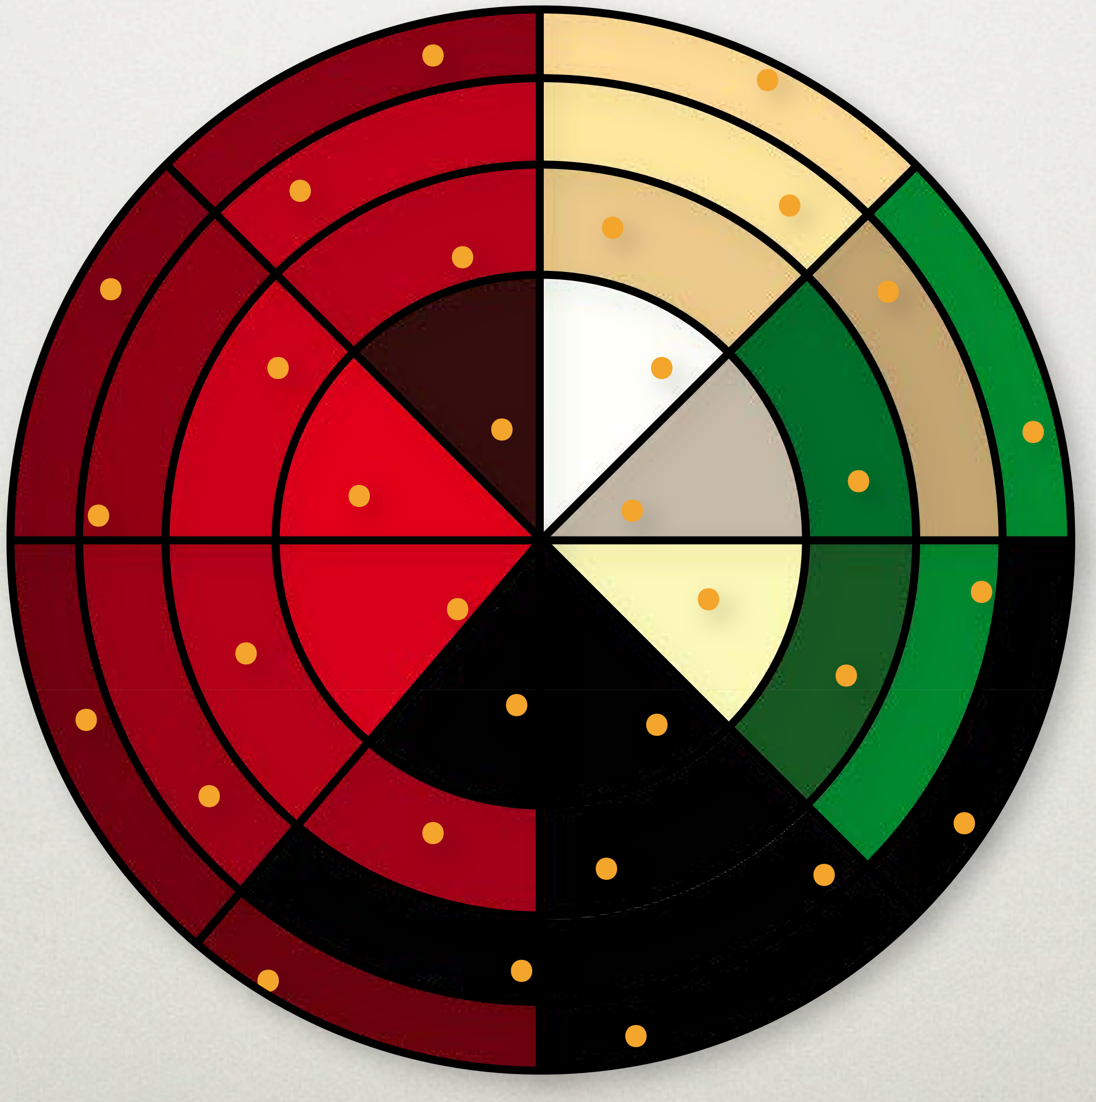

* Computing the gradient therefore involves estimating how the cell areas change due to a translation
* This term is what Ward and Heckbert derived
* Our contribution is additionally taking into account a gradient of the cell radiance.

Irradiance Gradient From Surfaces

$$\nabla E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} \left(\nabla A_{j,k} \, L^s_{j,k} + A_{j,k} \, \nabla L^s_{j,k} \right) (\vec{n} \cdot \vec{\omega}_{j,k})$$

gradient of cell area (Ward and Heckbert)

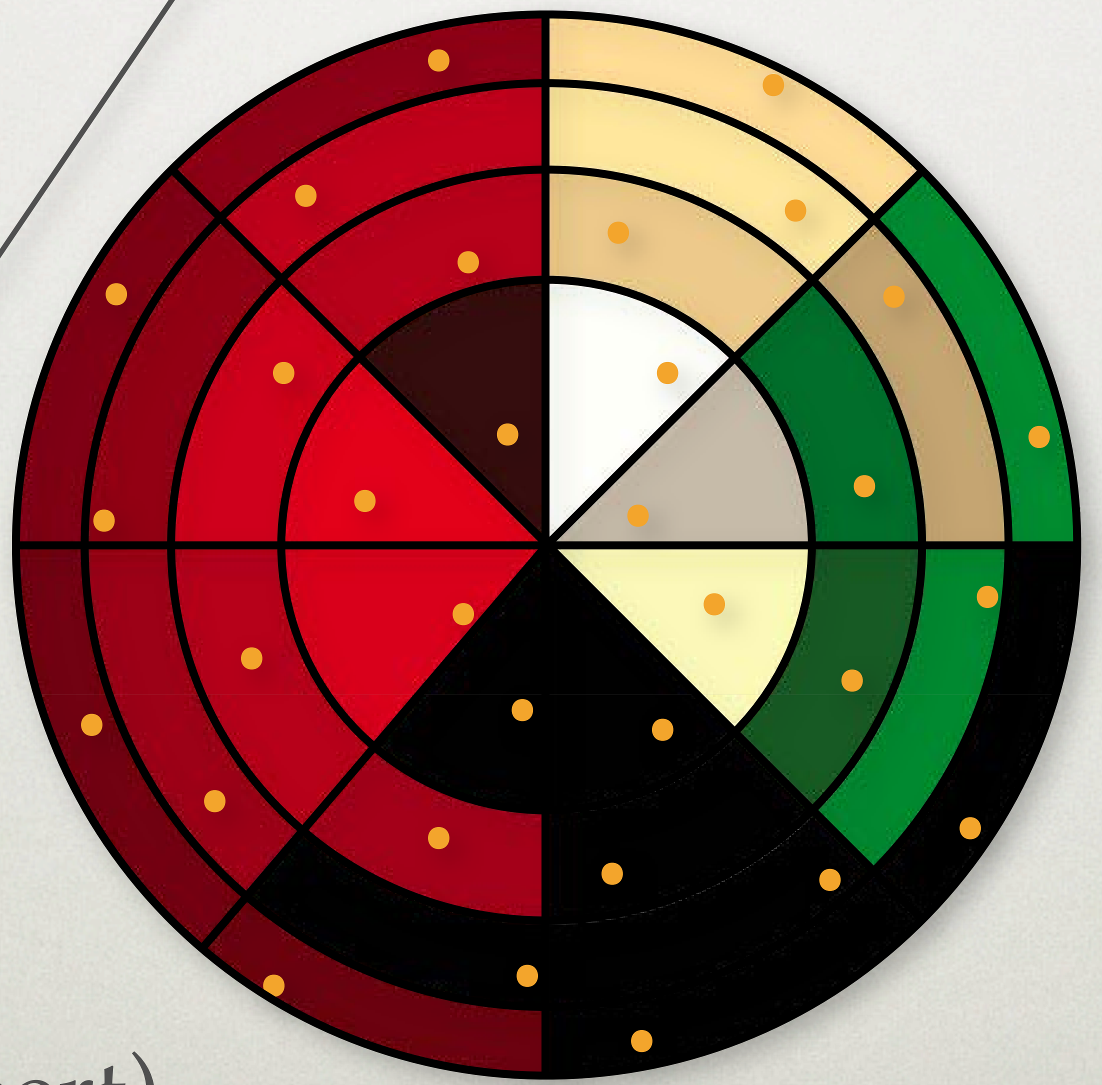

Thursday, 6 September 12

* Computing the gradient therefore involves estimating how the cell areas change due to a translation
* This term is what Ward and Heckbert derived
* Our contribution is additionally taking into account a gradient of the cell radiance.

Irradiance Gradient From Surfaces

$$\nabla E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} \left(\boxed{\nabla A_{j,k}} L^s_{j,k} + A_{j,k} \boxed{\nabla L^s_{j,k}} \right) (\vec{n} \cdot \vec{\omega}_{j,k})$$

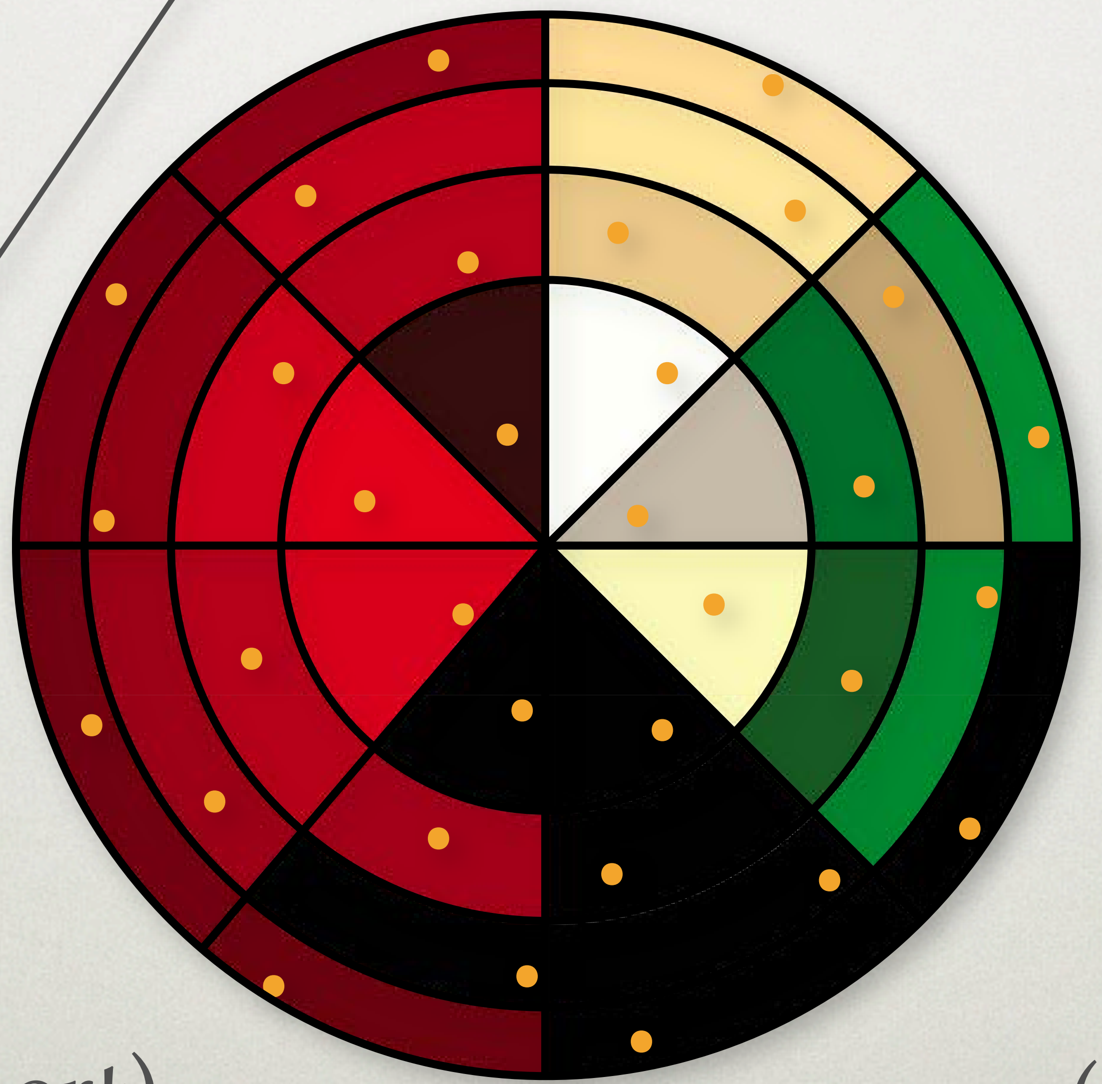

gradient of cell area (Ward and Heckbert)

gradient of cell radiance (our contribution)

* Computing the gradient therefore involves estimating how the cell areas change due to a translation
* This term is what Ward and Heckbert derived
* Our contribution is additionally taking into account a gradient of the cell radiance.

Gradient of Cell Radiance

$$\nabla E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} (\nabla A_{j,k}\, L^s_{j,k} + A_{j,k}\, \nabla L^s_{j,k})\, (\vec{n} \cdot \vec{\omega}_{j,k})$$

Thursday, 6 September 12

* In participating media, the surface radiance is the product of two terms, so its gradient can be computed using the product rule.
* We recently published a method at TOG which derives the necessary expressions for computing the gradient of the transmittance.
* The gradient of cell radiance was ignored by all previous methods. This implies that all these methods (including radiance caching for glossy surfaces) assumed that all surfaces visible during final gather are Lambertian surfaces in a vacuum.
* By incorporating this term we can not only account for participating media, but also get the added benefit of being able to handle glossy indirect reflectors.

Gradient of Cell Radiance

$$\nabla E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} (\nabla A_{j,k}\, L^s_{j,k} + A_{j,k} \boxed{\nabla L^s_{j,k}})\, (\vec{n} \cdot \vec{\omega}_{j,k})$$

$$L^s_{j,k} = T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, -\vec{\omega}_{j,k})$$

* In participating media, the surface radiance is the product of two terms, so its gradient can be computed using the product rule.
* We recently published a method at TOG which derives the necessary expressions for computing the gradient of the transmittance.
* The gradient of cell radiance was ignored by all previous methods. This implies that all these methods (including radiance caching for glossy surfaces) assumed that all surfaces visible during final gather are Lambertian surfaces in a vacuum.
* By incorporating this term we can not only account for participating media, but also get the added benefit of being able to handle glossy indirect reflectors.

Gradient of Cell Radiance

$$\nabla E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} (\nabla A_{j,k}\, L^s_{j,k} + A_{j,k} \boxed{\nabla L^s_{j,k}})\, (\vec{n} \cdot \vec{\omega}_{j,k})$$

$$L^s_{j,k} = T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, -\vec{\omega}_{j,k})$$

$$\boxed{\nabla L^s_{j,k}} = \nabla T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, -\vec{\omega}_{j,k}) + T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) \nabla L(\mathbf{x}_s, -\vec{\omega}_{j,k})$$

Thursday, 6 September 12

* In participating media, the surface radiance is the product of two terms, so its gradient can be computed using the product rule.
* We recently published a method at TOG which derives the necessary expressions for computing the gradient of the transmittance.
* The gradient of cell radiance was ignored by all previous methods. This implies that all these methods (including radiance caching for glossy surfaces) assumed that all surfaces visible during final gather are Lambertian surfaces in a vacuum.
* By incorporating this term we can not only account for participating media, but also get the added benefit of being able to handle glossy indirect reflectors.

Gradient of Cell Radiance

$$\nabla E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} (\nabla A_{j,k}\, L^s_{j,k} + A_{j,k} \boxed{\nabla L^s_{j,k}})\, (\vec{n} \cdot \vec{\omega}_{j,k})$$

$$L^s_{j,k} = T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, -\vec{\omega}_{j,k})$$

$$\boxed{\nabla L^s_{j,k}} = \boxed{\nabla T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s)} L(\mathbf{x}_s, -\vec{\omega}_{j,k}) + T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) \nabla L(\mathbf{x}_s, -\vec{\omega}_{j,k})$$

Attenuation due to media

Jarosz et al. ACM TOG '08.

* In participating media, the surface radiance is the product of two terms, so its gradient can be computed using the product rule.
* We recently published a method at TOG which derives the necessary expressions for computing the gradient of the transmittance.
* The gradient of cell radiance was ignored by all previous methods. This implies that all these methods (including radiance caching for glossy surfaces) assumed that all surfaces visible during final gather are Lambertian surfaces in a vacuum.
* By incorporating this term we can not only account for participating media, but also get the added benefit of being able to handle glossy indirect reflectors.

Gradient of Cell Radiance

$$\nabla E_s(\mathbf{x}) \approx \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} (\nabla A_{j,k}\, L^s_{j,k} + A_{j,k}\, \nabla L^s_{j,k})\, (\vec{n} \cdot \vec{\omega}_{j,k})$$

$$L^s_{j,k} = T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, -\vec{\omega}_{j,k})$$

$$\nabla L^s_{j,k} = \nabla T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, -\vec{\omega}_{j,k}) + T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) \nabla L(\mathbf{x}_s, -\vec{\omega}_{j,k})$$

Glossy indirect reflectors

Thursday, 6 September 12

* In participating media, the surface radiance is the product of two terms, so its gradient can be computed using the product rule.
* We recently published a method at TOG which derives the necessary expressions for computing the gradient of the transmittance.
* The gradient of cell radiance was ignored by all previous methods. This implies that all these methods (including radiance caching for glossy surfaces) assumed that all surfaces visible during final gather are Lambertian surfaces in a vacuum.
* By incorporating this term we can not only account for participating media, but also get the added benefit of being able to handle glossy indirect reflectors.

Hemispherical Sampling

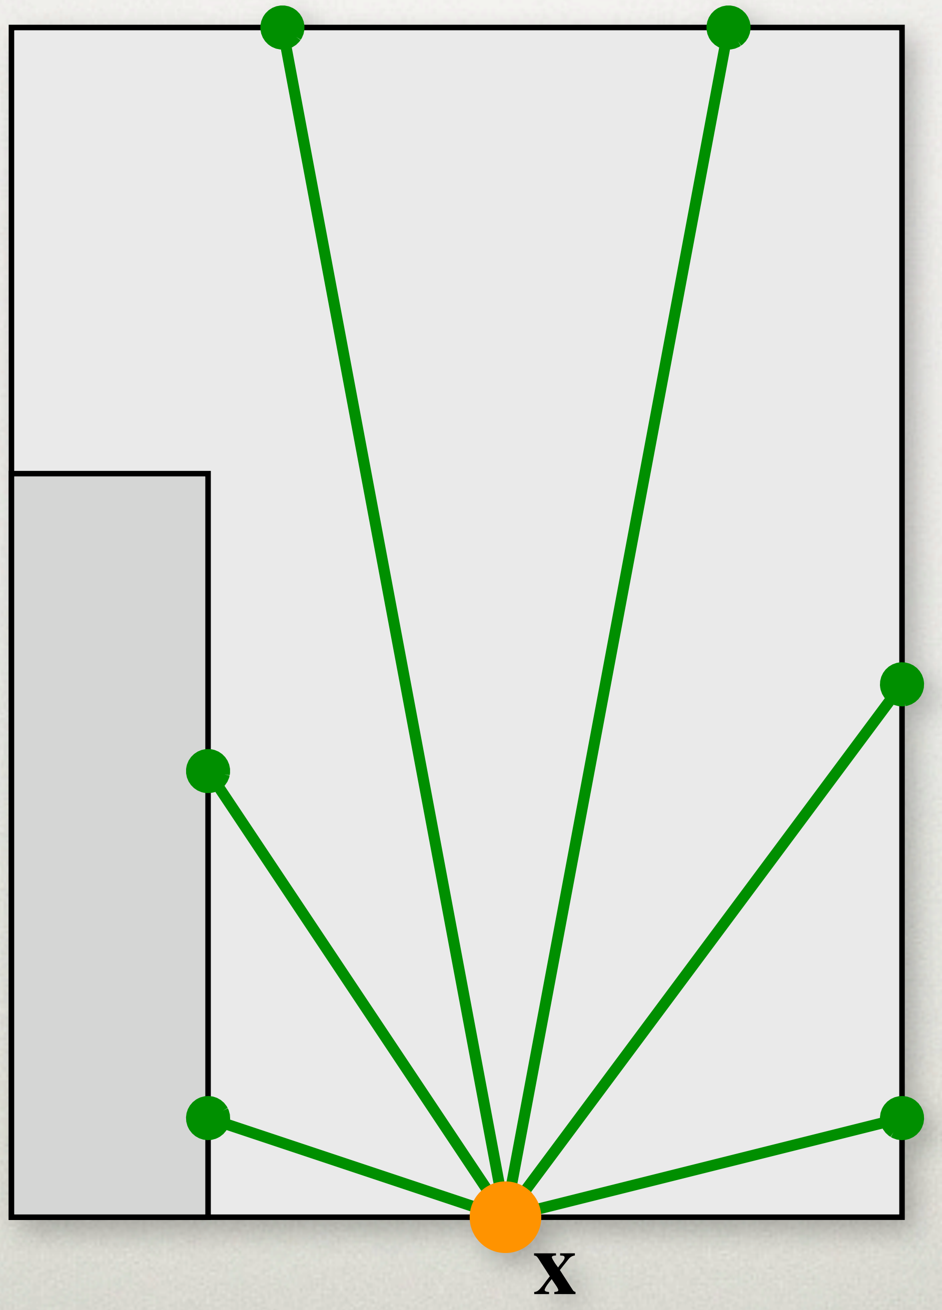

* The way this derivation can be interpreted visually, is that we start with a hemispherical sampling around some point x

Hemispherical Sampling

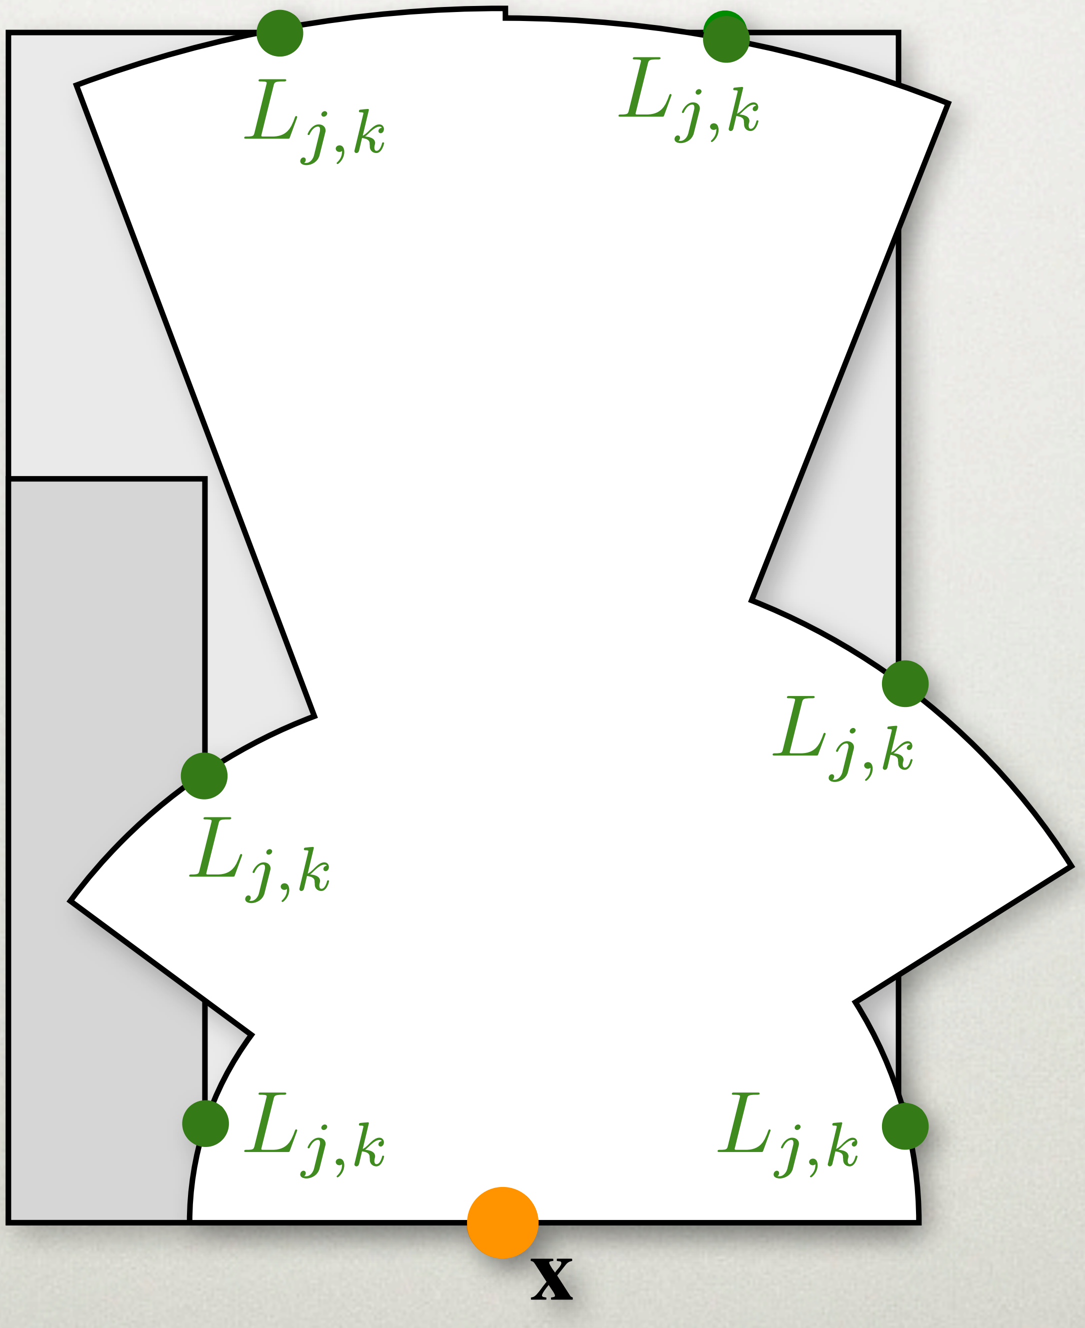

Thursday, 6 September 12

* We now know the radiance coming from each cell, and the distance to the surface within each cell, which results in a discretization of the visible environment.
* In order to compute the gradient of irradiance, we consider how the contribution from each cell will change as we move the point x along the tangent plane.

Hemispherical Sampling

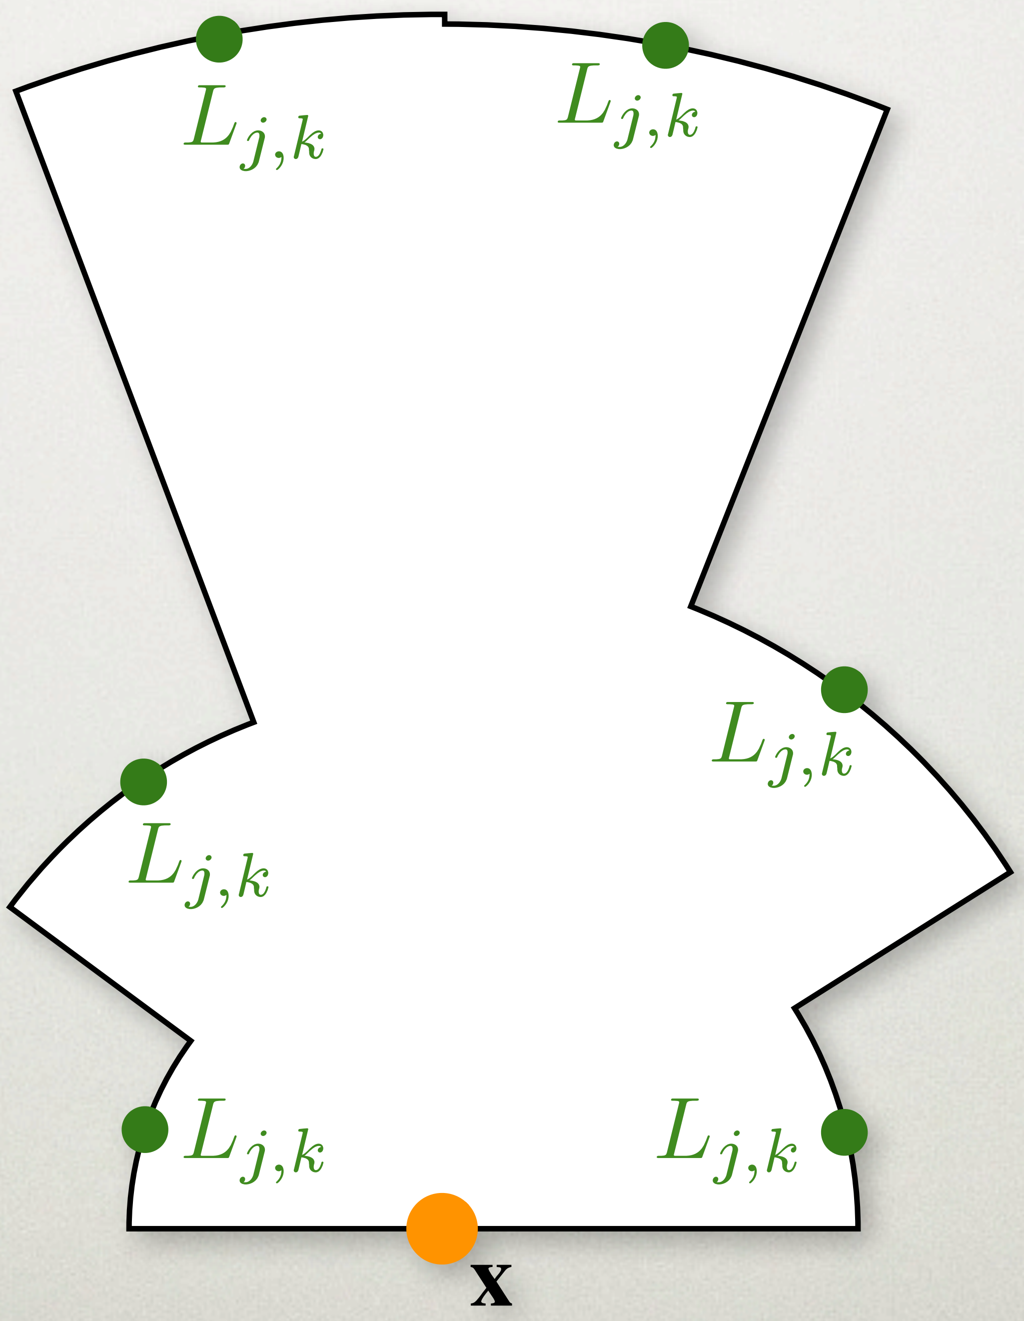

* We now know the radiance coming from each cell, and the distance to the surface within each cell, which results in a discretization of the visible environment.
* In order to compute the gradient of irradiance, we consider how the contribution from each cell will change as we move the point x along the tangent plane.

Surface Irradiance Gradient

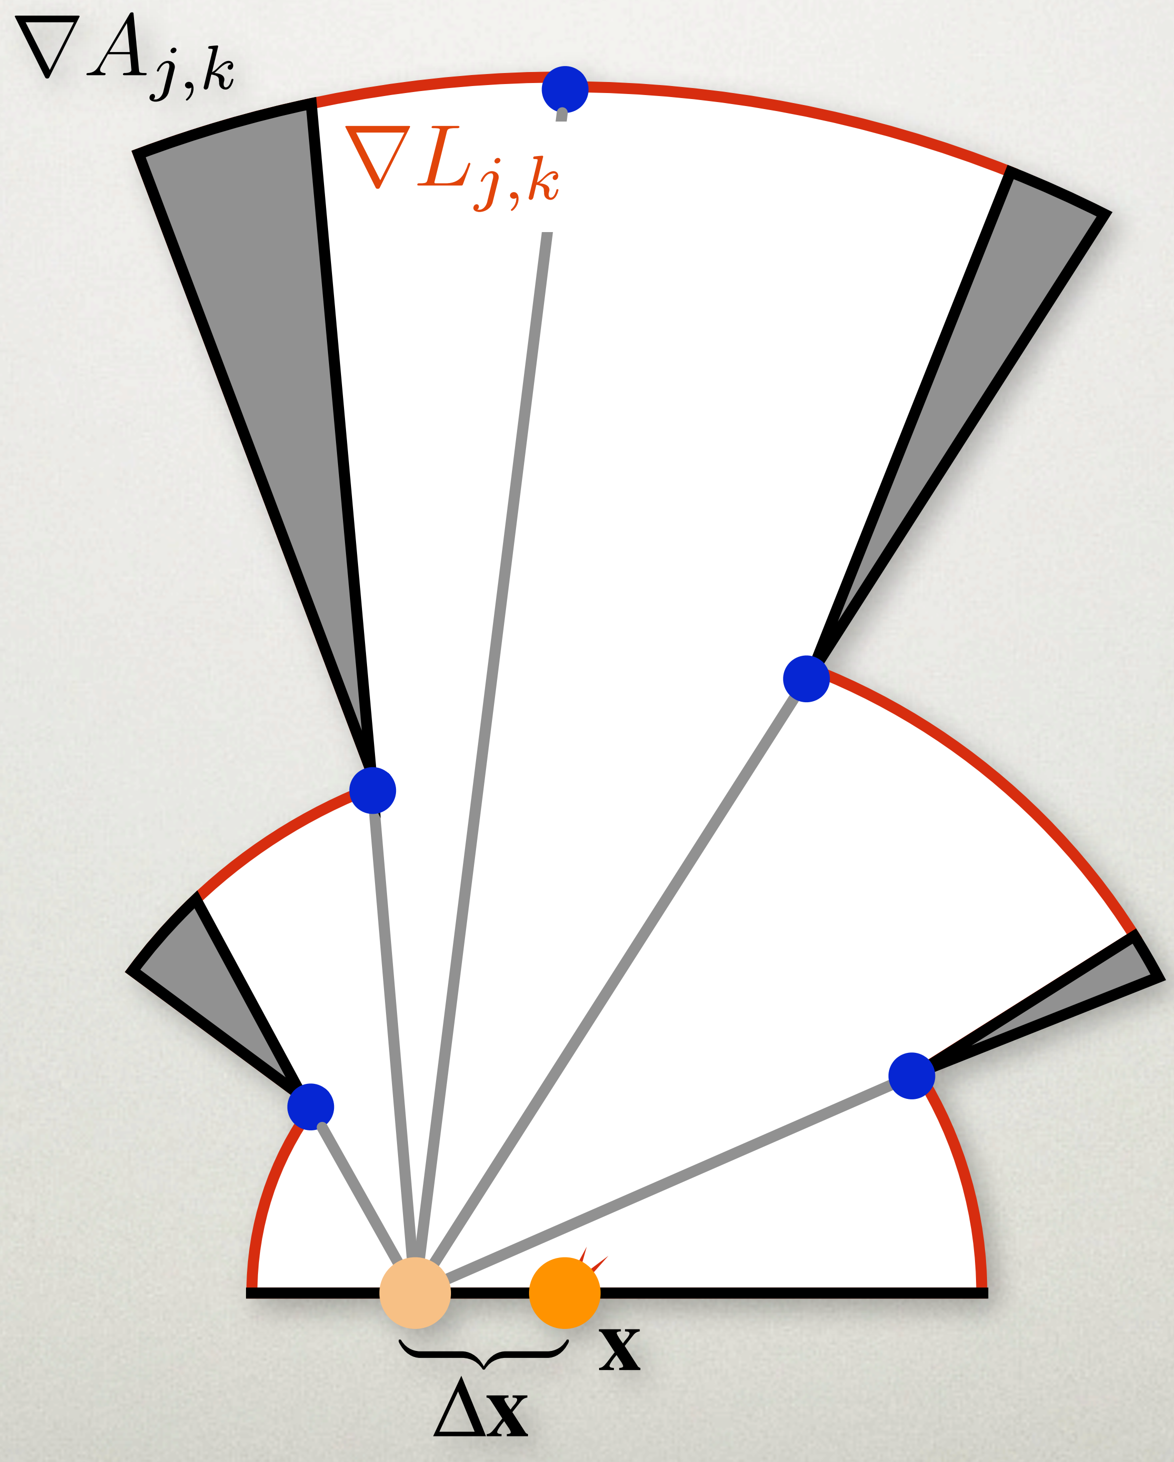

Thursday, 6 September 12

* Moving the point will result in the cell areas changing due to occlusions from neighboring surfaces (shown in grey).
* Additionally, the radiance coming from each cell may change due to changes in extinction (shown in red).

Absorbing Medium

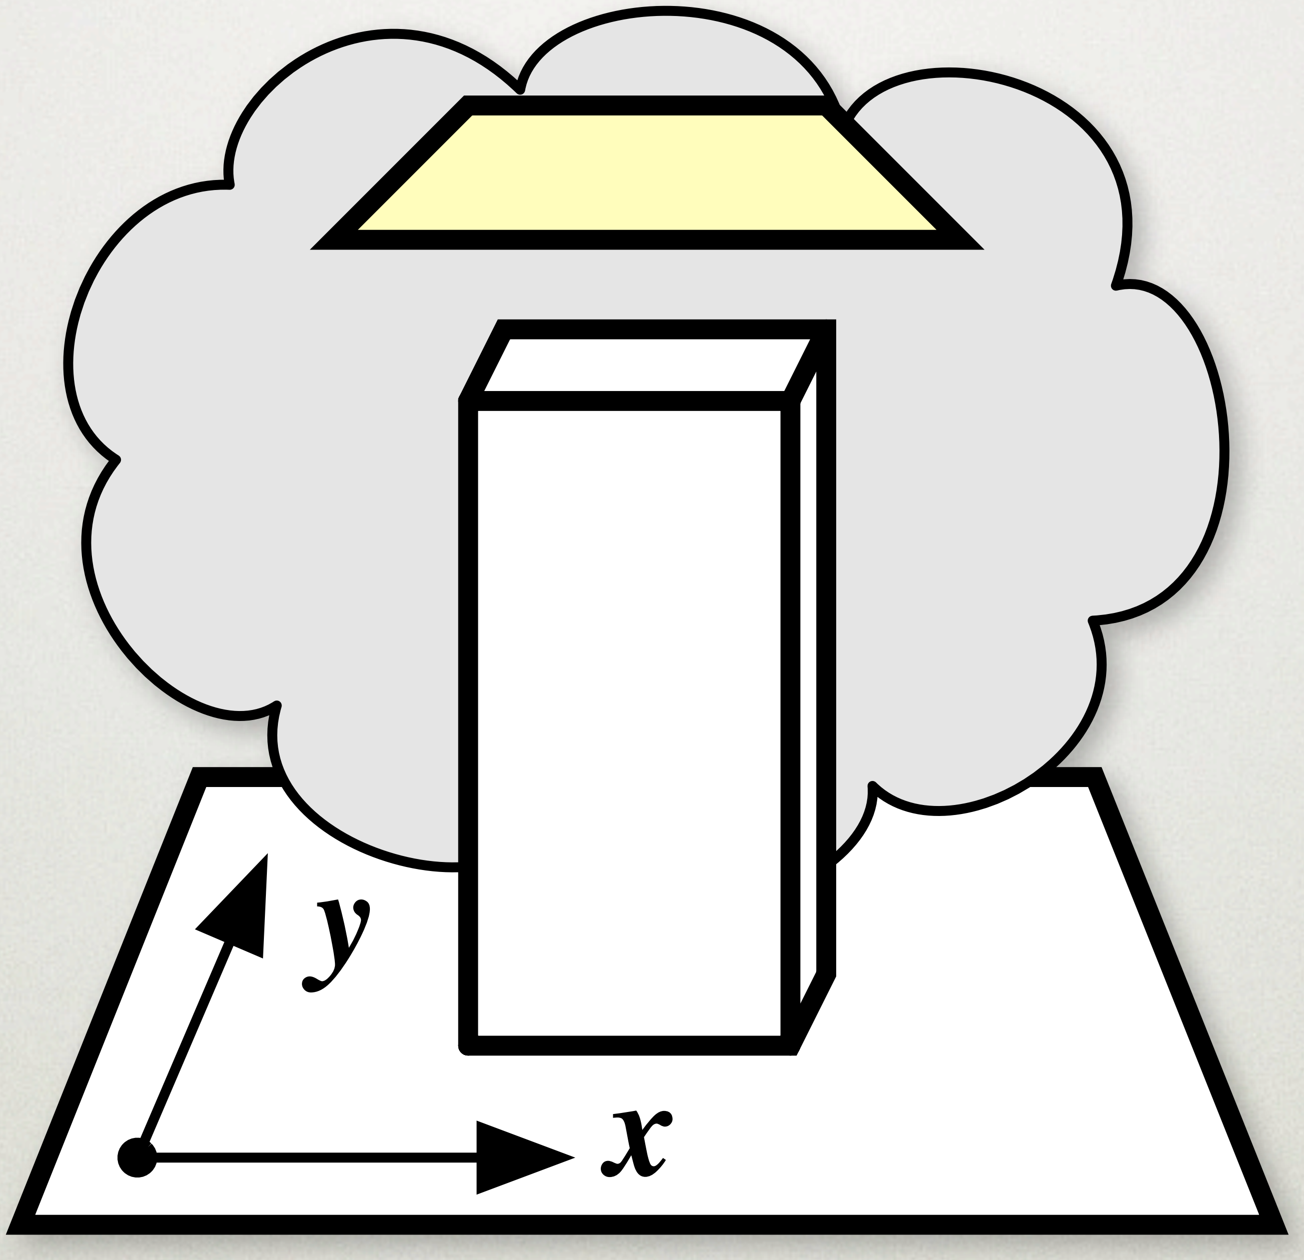

Thursday, 6 September 12

* To validate this gradient formulation we visualized the gradients within this simple synthetic scene, which contains a ground plane, an occluding block, and a polygon reflecting indirect light. The whole scene is embedded within an absorbing medium.

Absorbing Medium

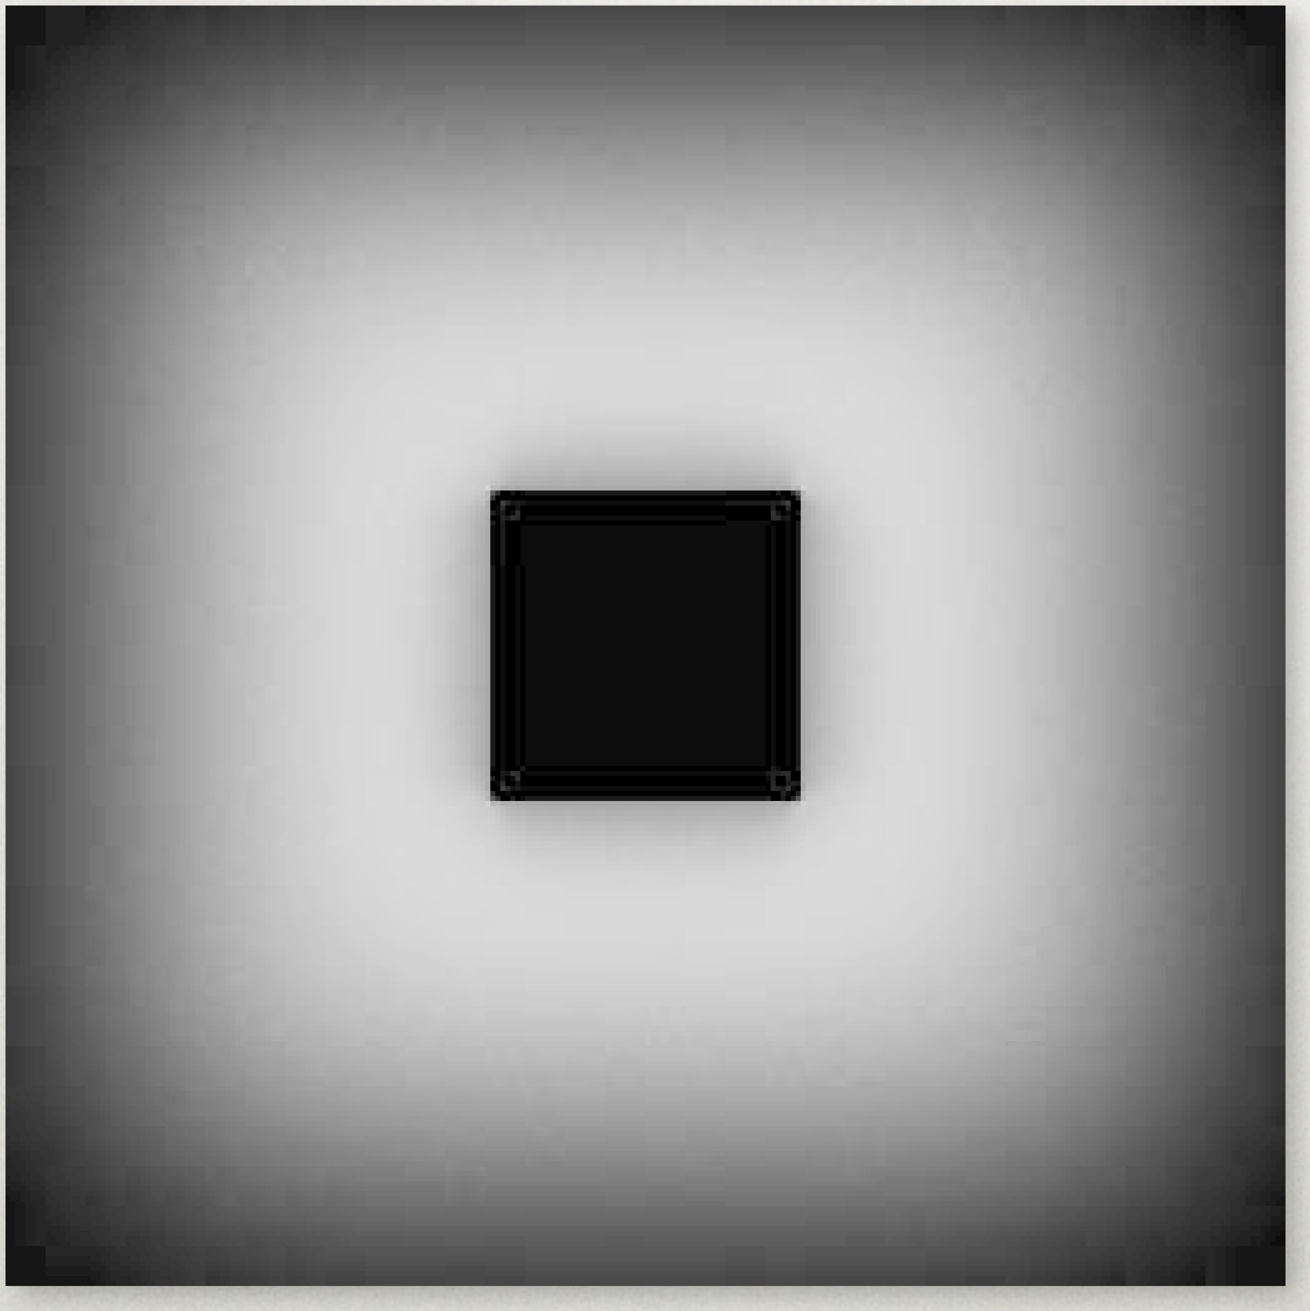

Irradiance on floor

Thursday, 6 September 12

* We can visualize the irradiance on the ground plane.

Per-Pixel Irradiance Gradient

$$\text{red} = \left|\frac{\partial E}{\partial x}\right|$$

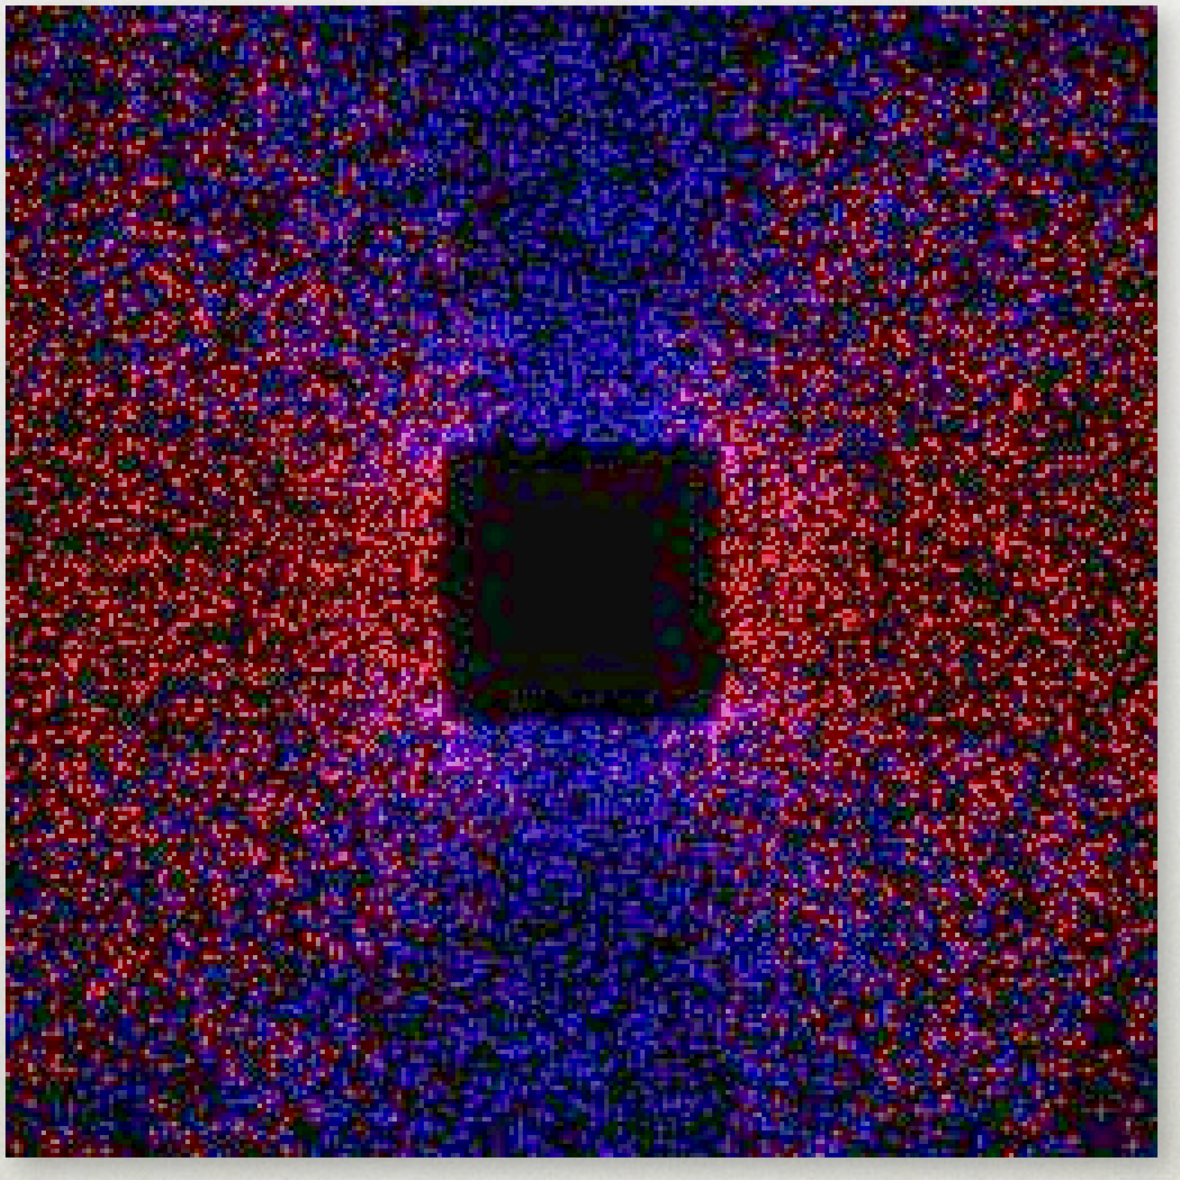

$$\text{blue} = \left|\frac{\partial E}{\partial y}\right|$$

(Finite Differences)

* We can also compute a ground truth solution to the gradient by performing finite differences along the ground plane.
* In these visualizations the absolute value of the x component of the gradient is shown in red and the y component is shown in blue. And we compute the gradient per-pixel
* This unfortunately suffers from significant noise.

Per-Pixel Irradiance Gradient

$$\text{red} = \left|\frac{\partial E}{\partial x}\right|$$

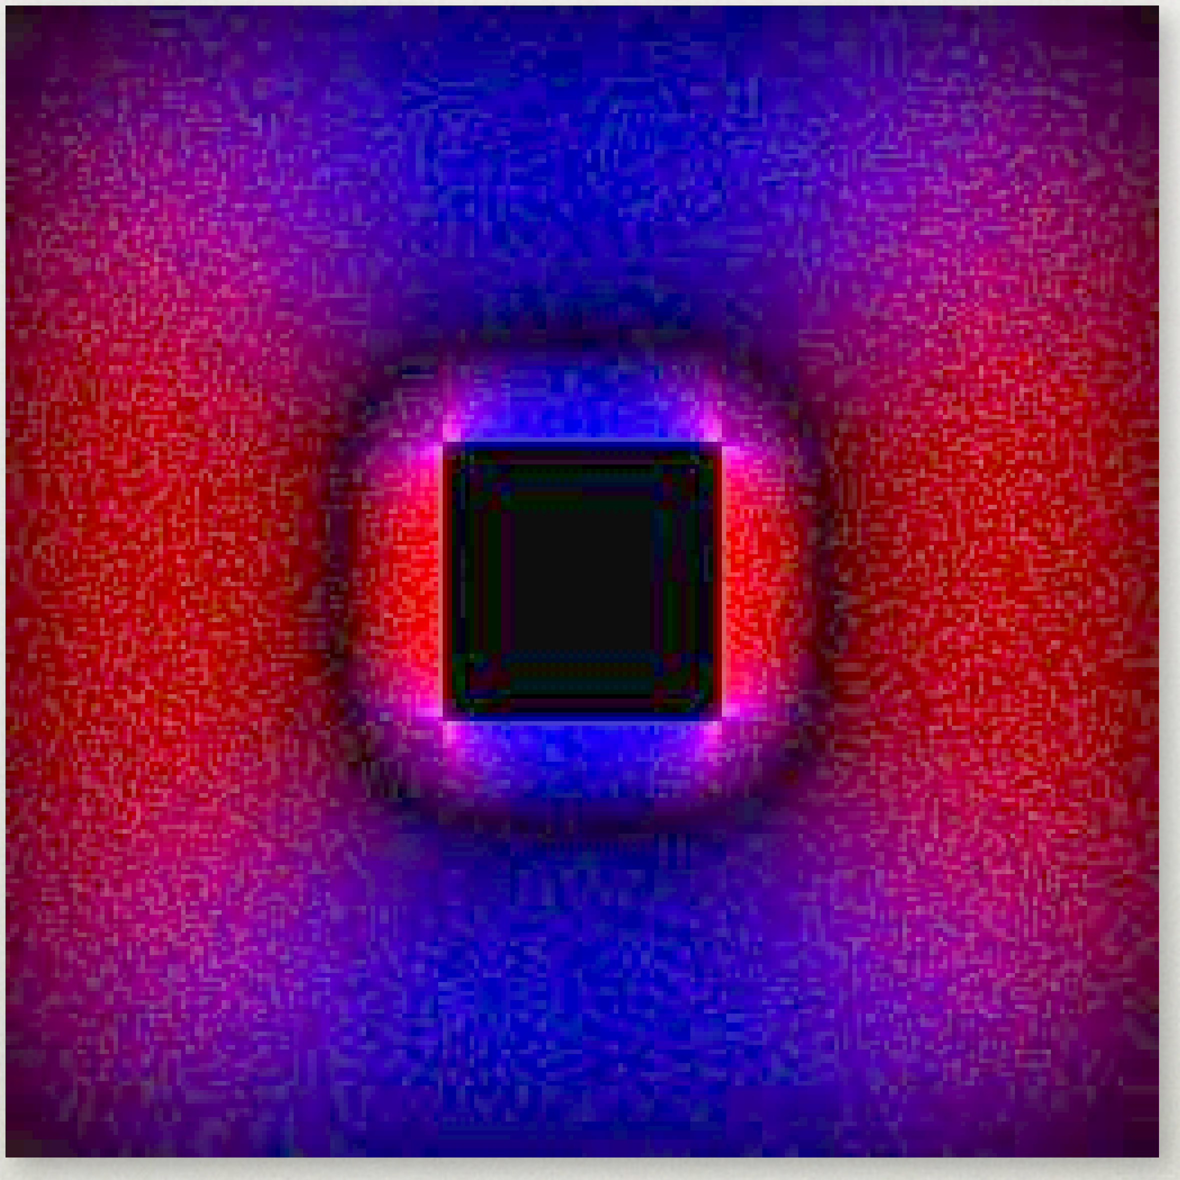

$$\text{blue} = \left|\frac{\partial E}{\partial y}\right|$$

(Finite Differences 10X)

* We can improve the quality by taking 10 times as many samples, and this starts to reveal the structure of the true gradient, however it is not a practical approach since it is very expensive

Per-Pixel Irradiance Gradient

$$\text{red} = \left|\frac{\partial E}{\partial x}\right|$$

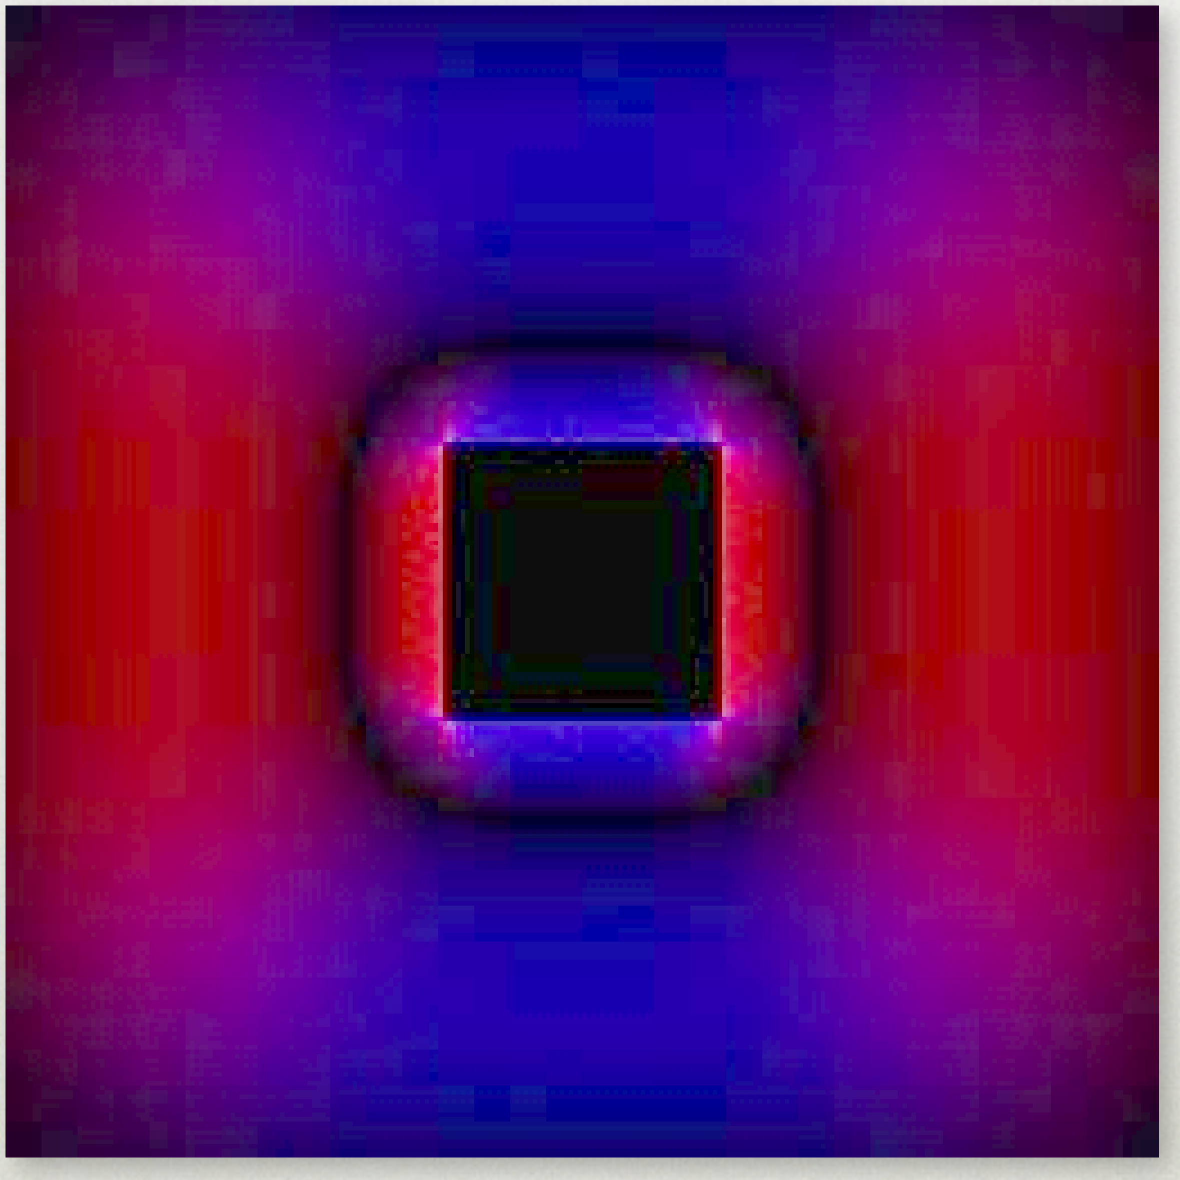

$$\text{blue} = \left|\frac{\partial E}{\partial y}\right|$$

(Our Method)

* Using our approach, we can match the behavior of this gradient, with less noise, and using only 1/10th of the number of samples.

Per-Pixel Irradiance Gradient

$$\text{red} = \left|\frac{\partial E}{\partial x}\right|$$

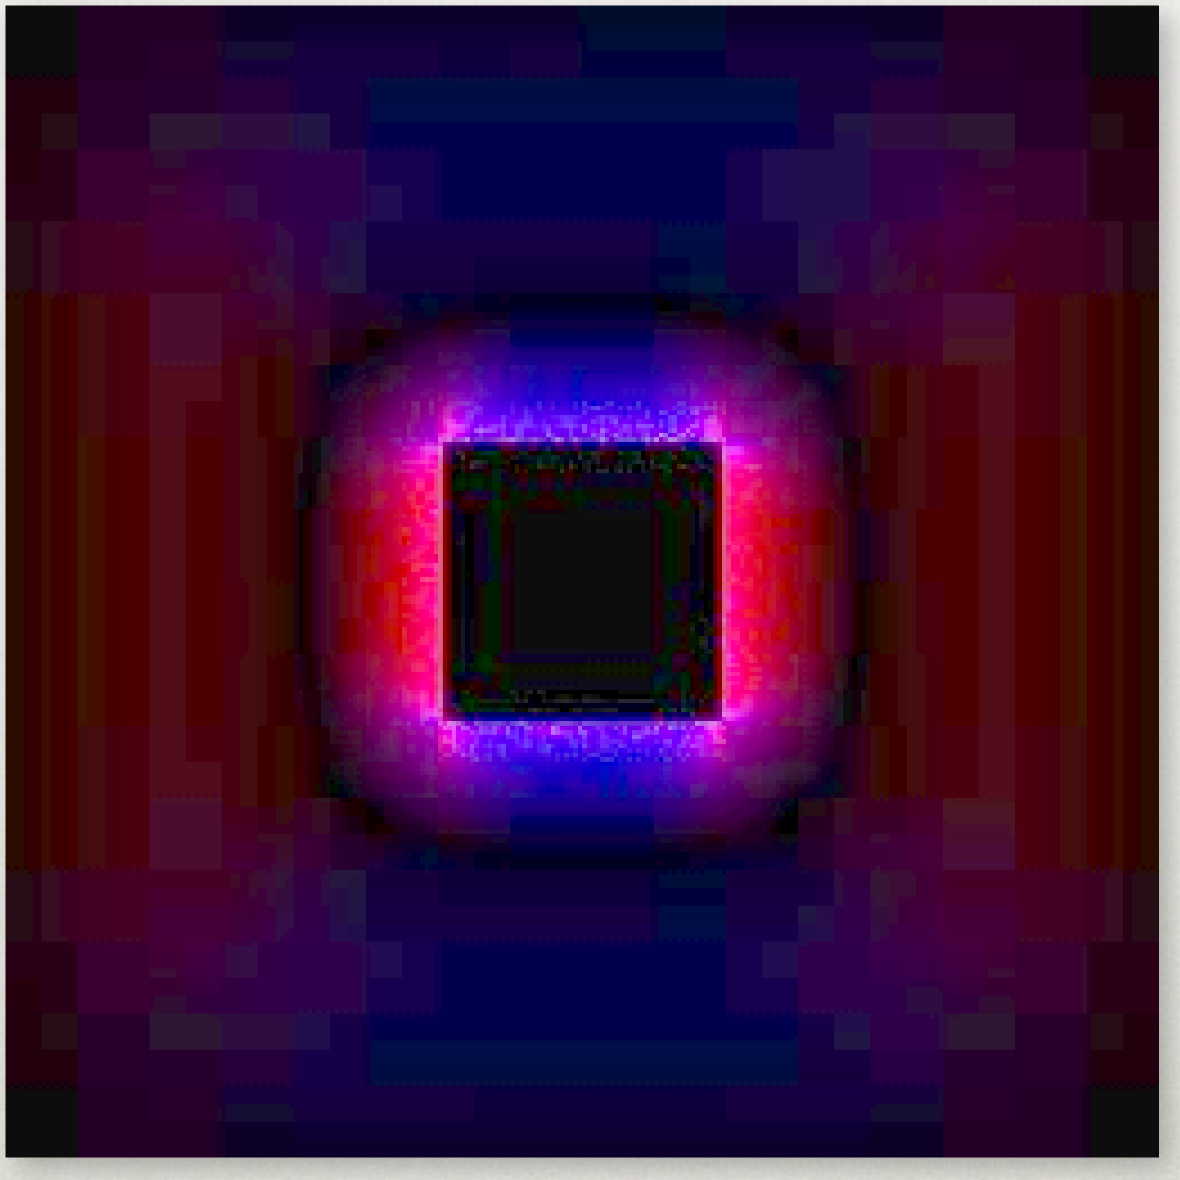

$$\text{blue} = \left|\frac{\partial E}{\partial y}\right|$$

(Ward and Heckbert)

* If we were to compute the gradient using the original Ward and Heckbert formulation, the results are significantly different than the finite difference gradients.

Gradient Comparison

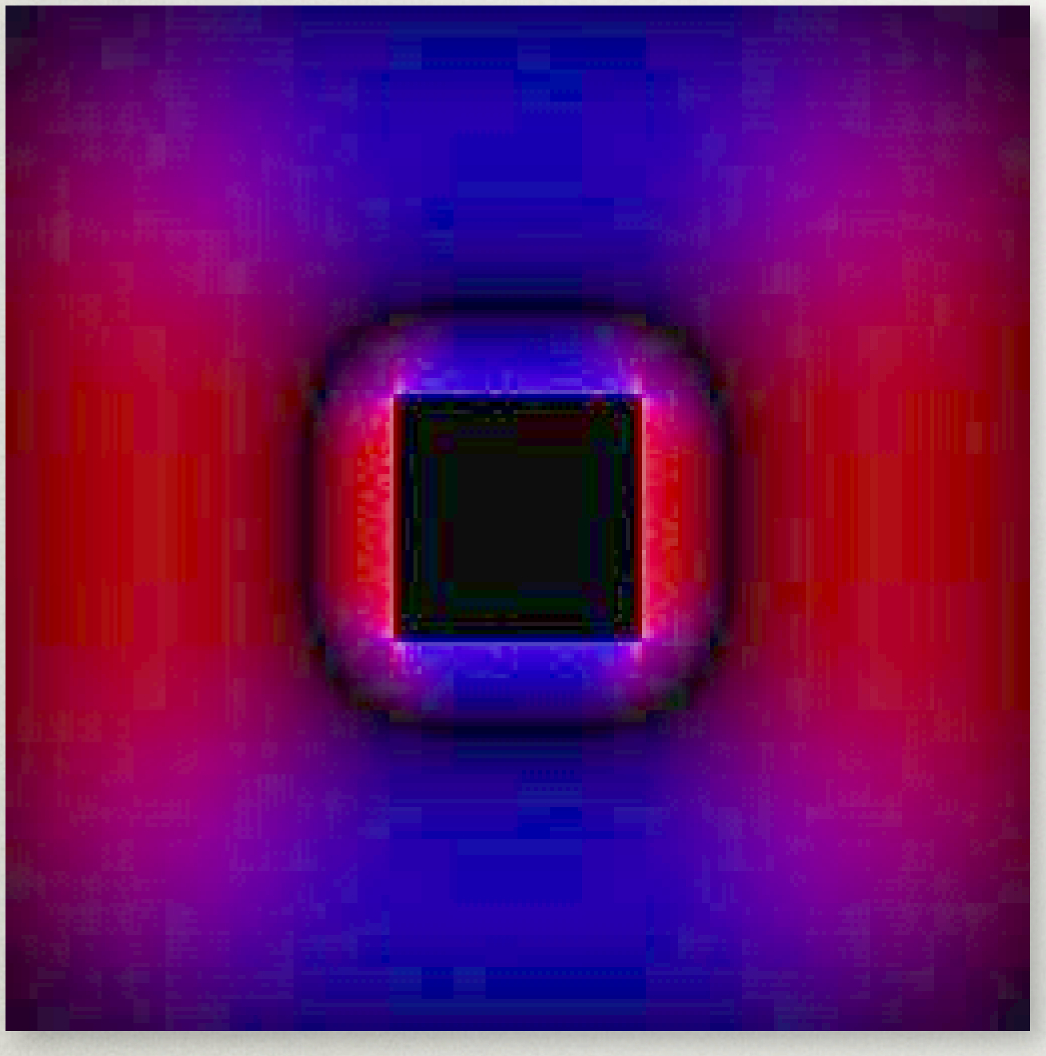

Our Method

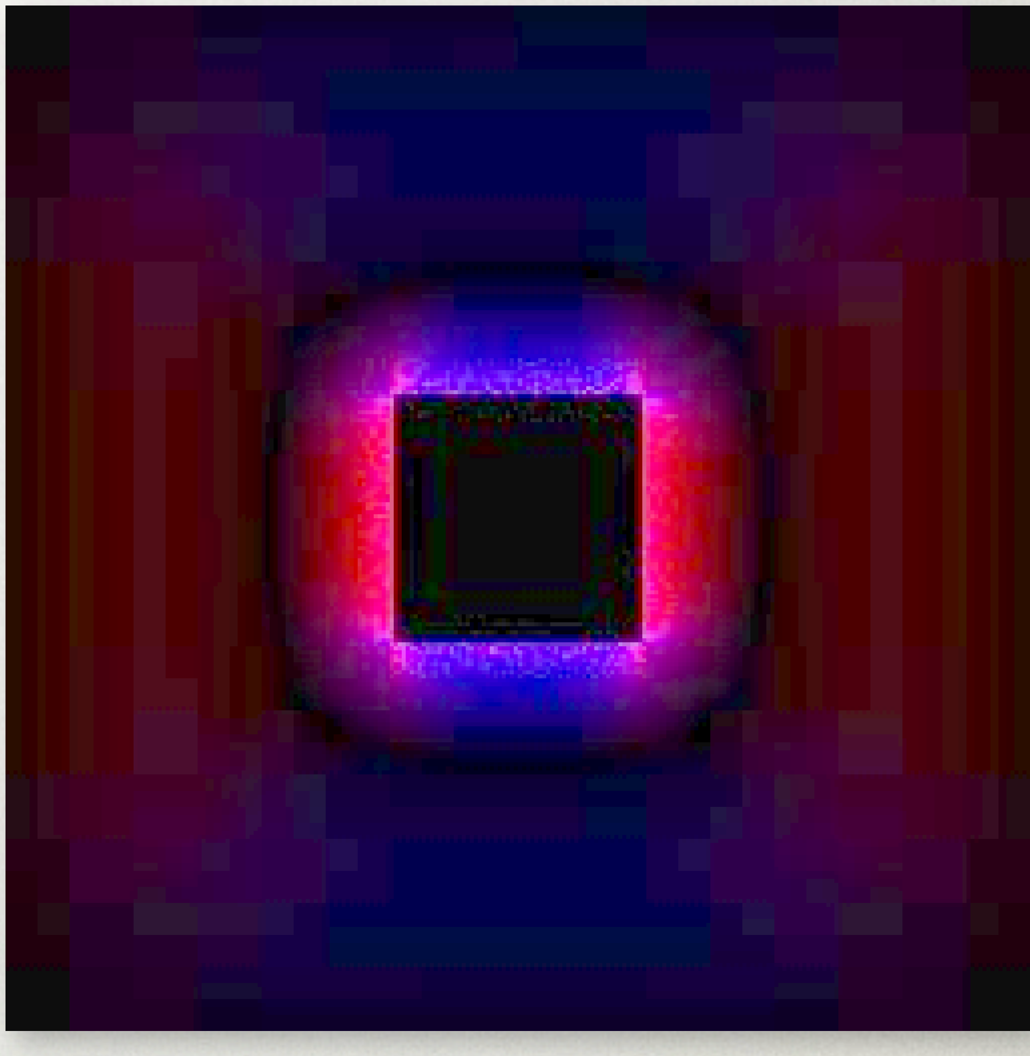

Ward and Heckbert

Why is the Ward & Heckbert gradient darker?

Thursday, 6 September 12

* If we look at these side by side we can immediately see that the Ward and Heckbert version is darker.
* For Ward & Heckbert, the radiance has an inverse squared falloff
* In participating media, which our gradients take into account, the radiance has a sharper falloff since it is also attenuated by transmittance.
* This leads to a higher gradient value.

Gradient Comparison

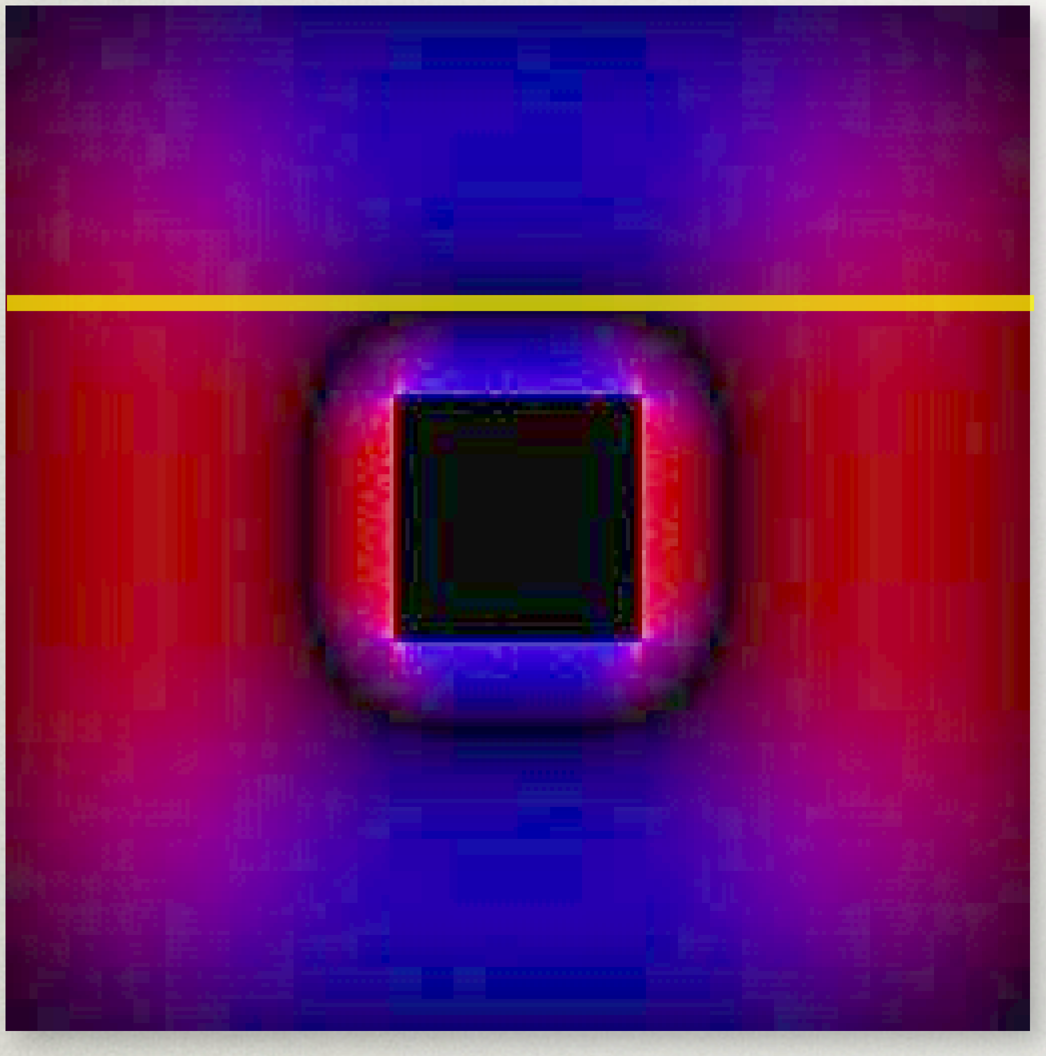

Our Method

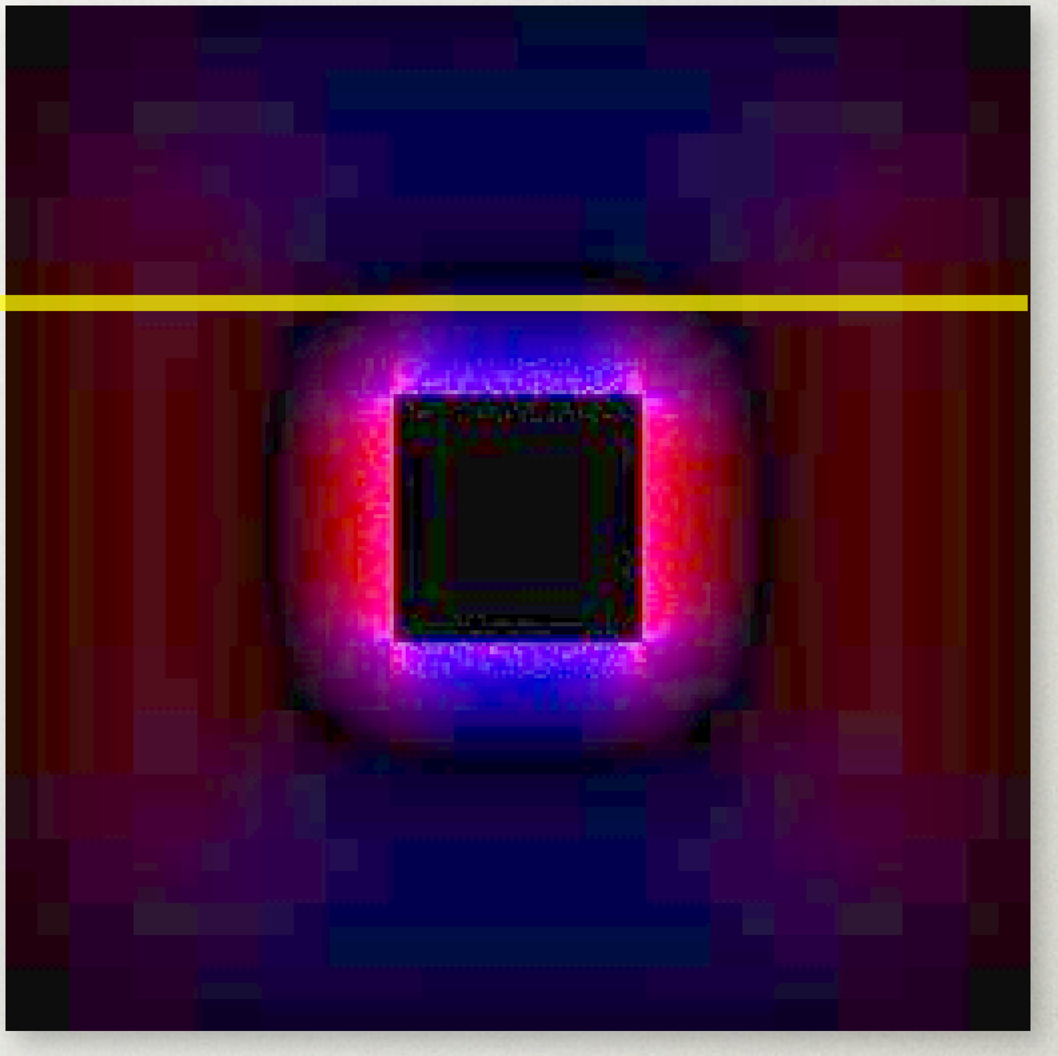

Ward and Heckbert

Why is the Ward & Heckbert gradient darker?

Thursday, 6 September 12

* If we look at these side by side we can immediately see that the Ward and Heckbert version is darker.
* For Ward & Heckbert, the radiance has an inverse squared falloff
* In participating media, which our gradients take into account, the radiance has a sharper falloff since it is also attenuated by transmittance.
* This leads to a higher gradient value.

Gradient Comparison

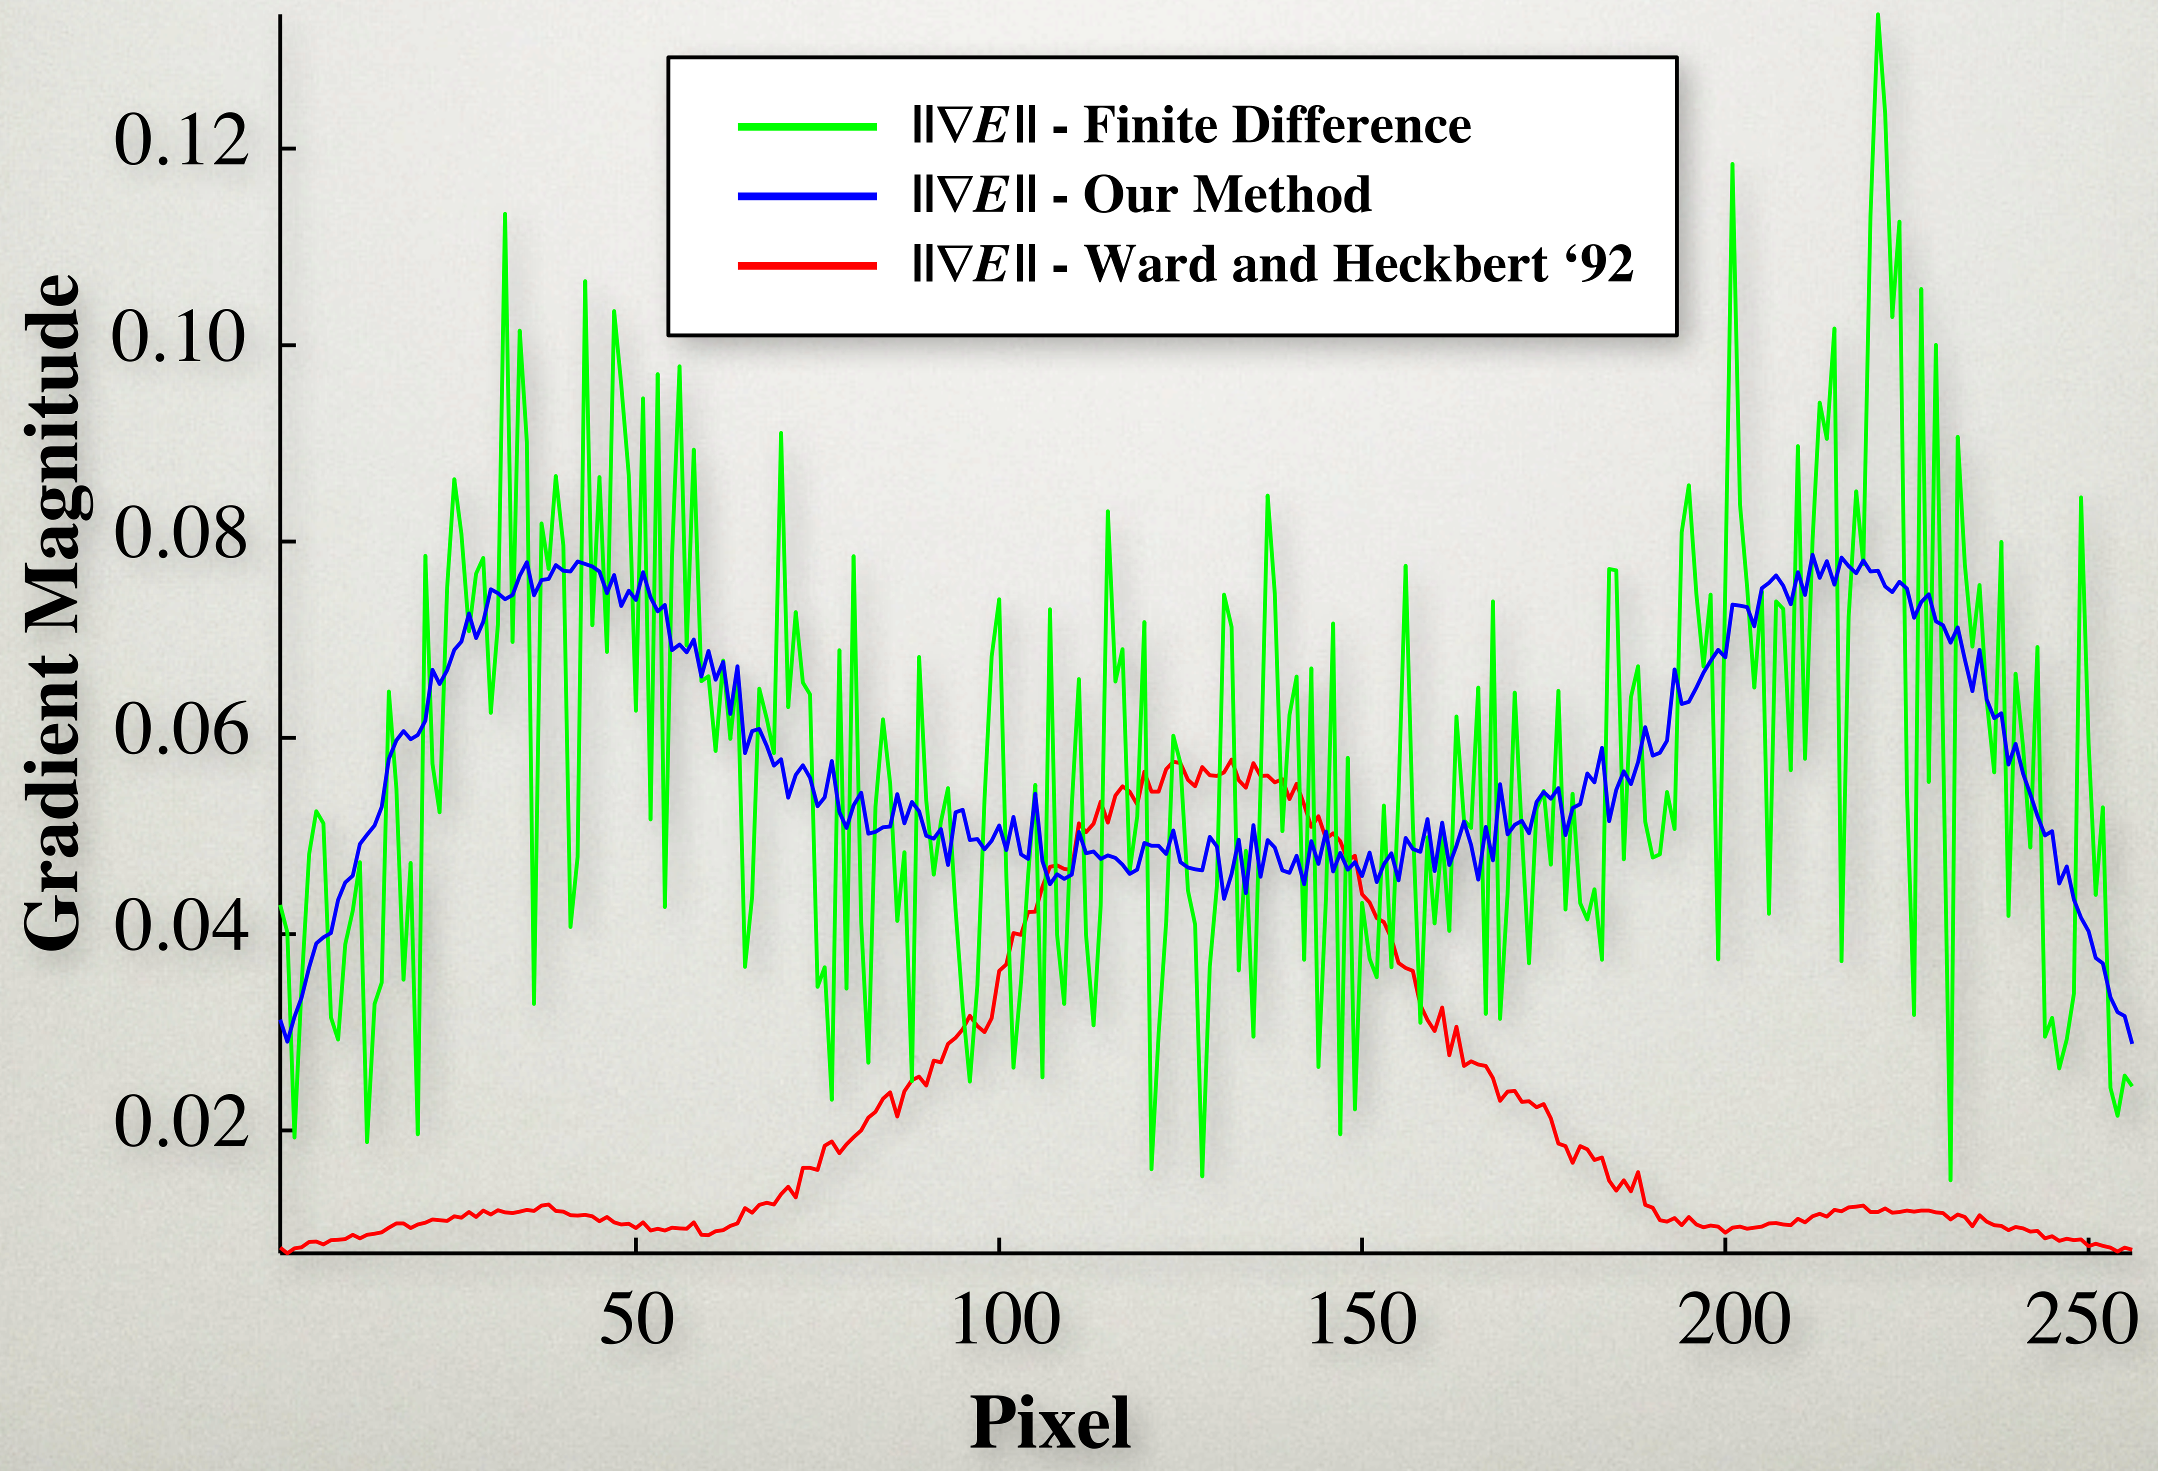

Extrapolated Irradiance

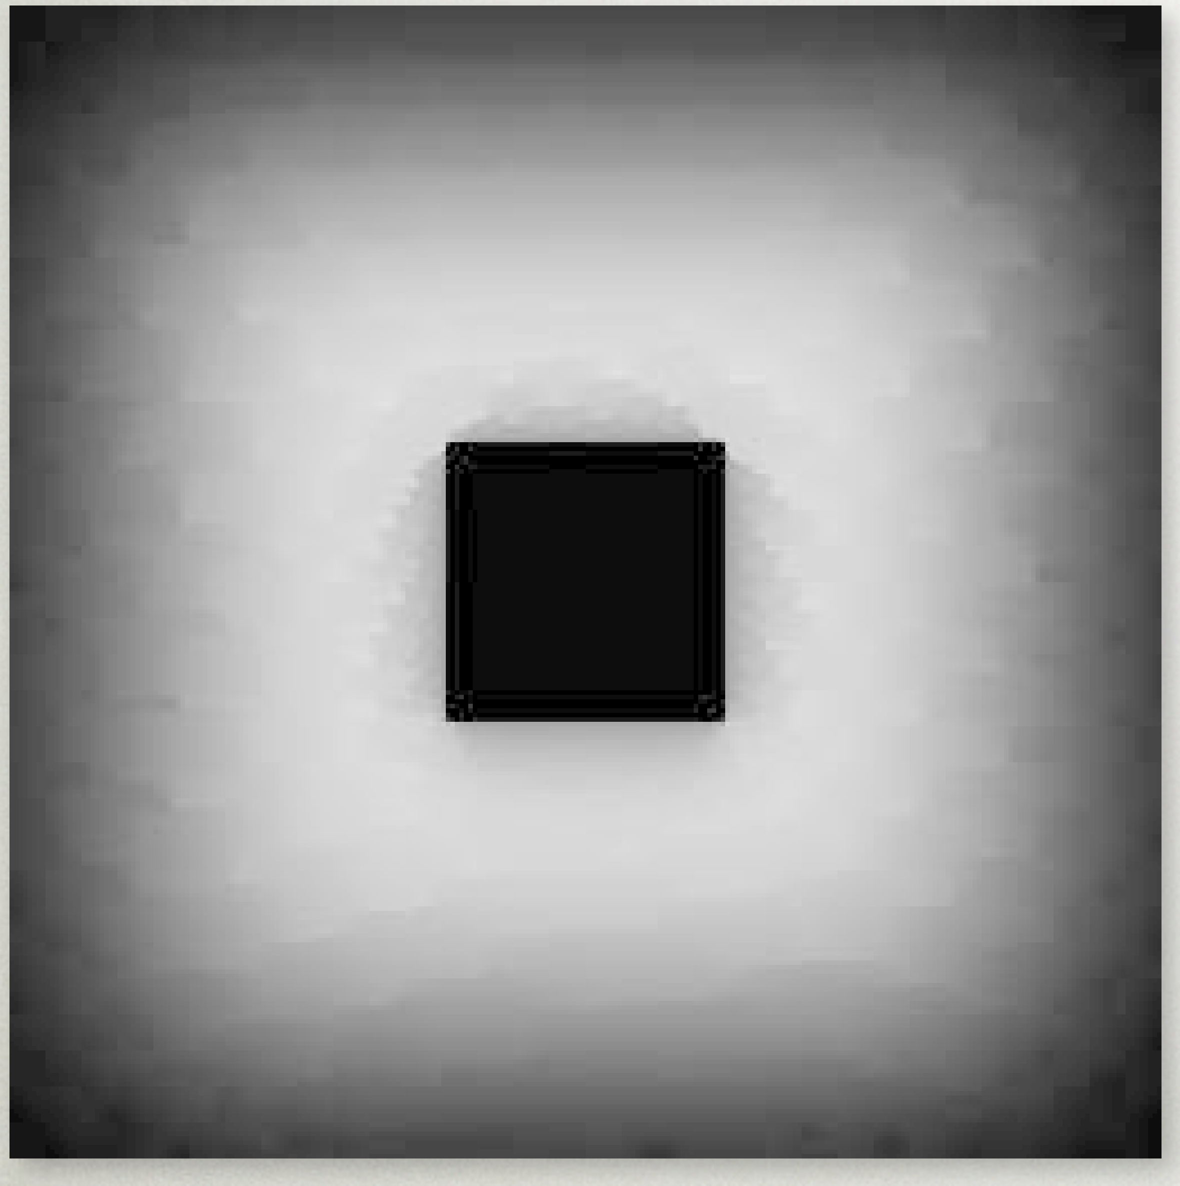

Our Method

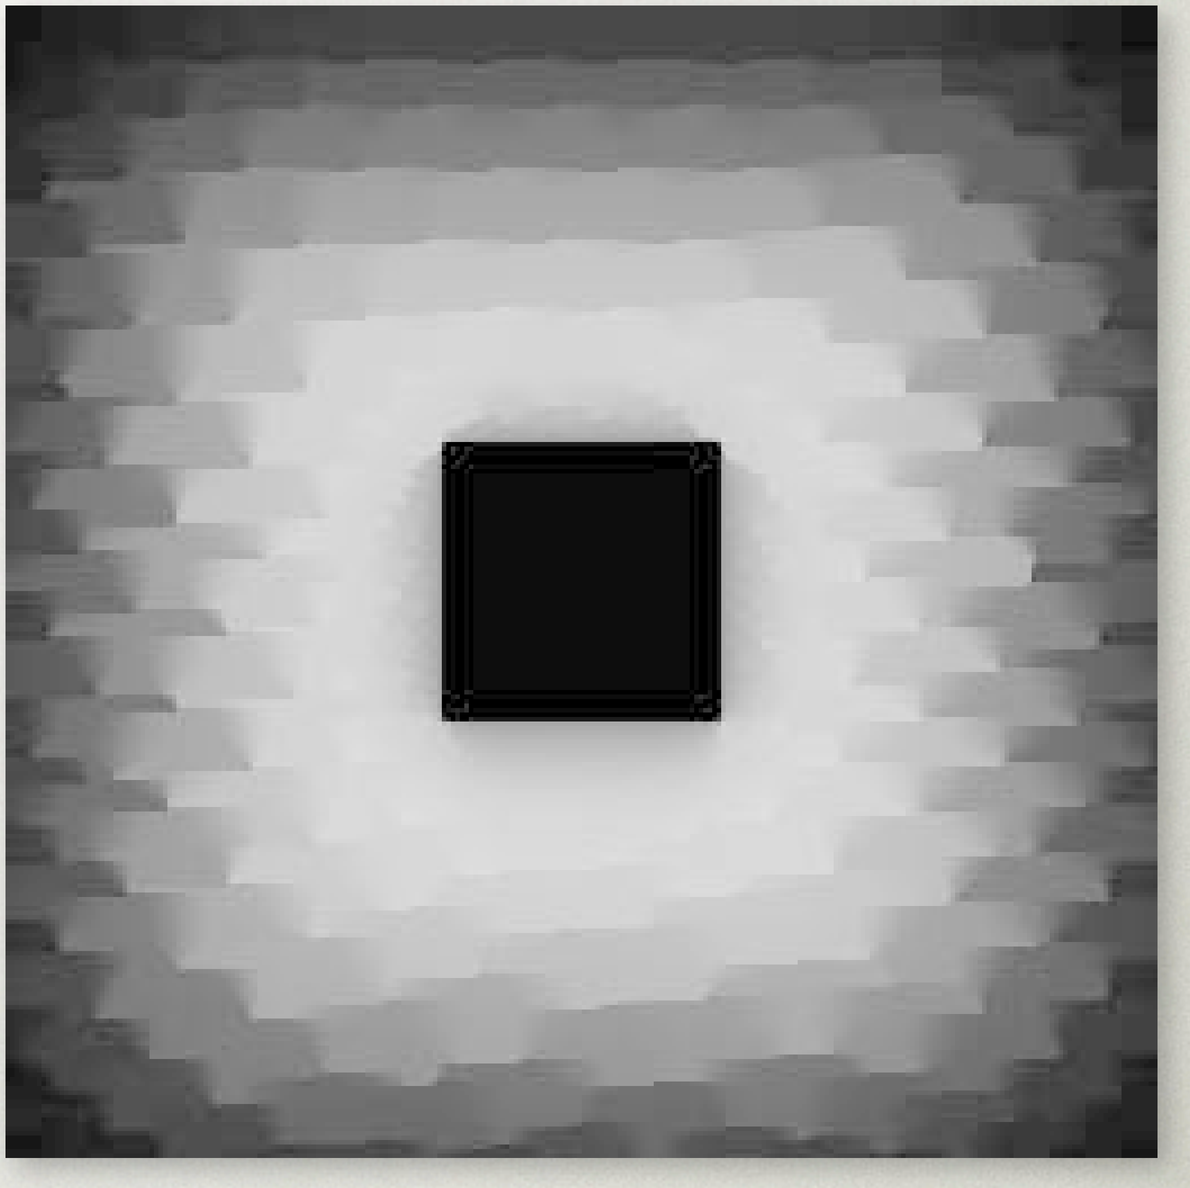

Ward and Heckbert

Same cache points

Visual Break

Thursday, 6 September 12

This is frame 352 from the Patterson film taken on October 20, 1967. It is the most famous picture of bigfoot ever taken.

Irradiance From Media

$$E_m(\mathbf{x}) = \int_{\Omega} L_m(\mathbf{x}, \vec{\omega})\,(\vec{\mathbf{n}} \cdot \vec{\omega})\, d\vec{\omega}$$

Irradiance From Media

$$E_m(\mathbf{x}) = \int_\Omega L_m(\mathbf{x}, \vec{\omega})\,(\vec{\mathbf{n}} \cdot \vec{\omega})\; d\vec{\omega}$$

where:

$$L_m(\mathbf{x}, \vec{\omega}) = \int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t)\sigma_s(\mathbf{x}_t)L_i(\mathbf{x}_t, \vec{\omega})\; dt$$

Ray Marching

$$L_m(\mathbf{x}, \vec{\omega}) = \int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t)\sigma_s(\mathbf{x}_t)L_i(\mathbf{x}_t, \vec{\omega})\ dt$$

Ray Marching

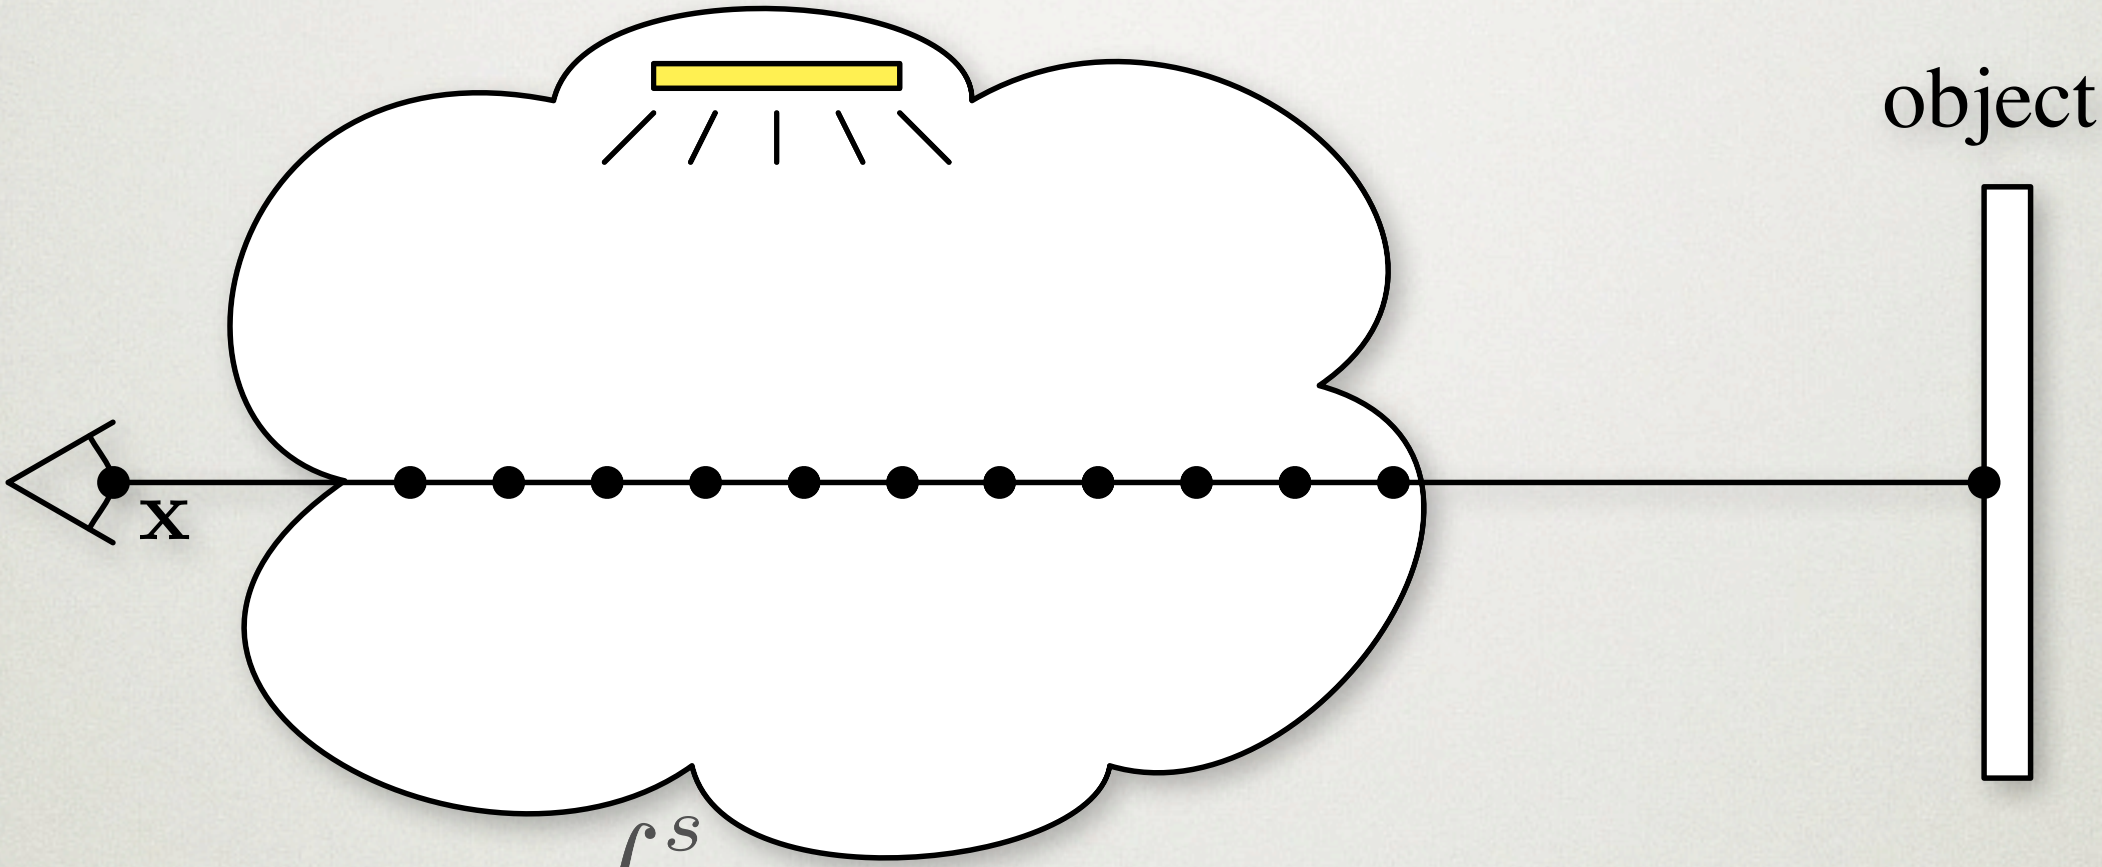

$$L_m(\mathbf{x},\vec{\omega}) = \int_0^s T_r(\mathbf{x}\leftrightarrow\mathbf{x}_t)\sigma_s(\mathbf{x}_t)L_i(\mathbf{x}_t,\vec{\omega})\ dt$$

$$L_m(\mathbf{x},\vec{\omega}) \approx \sum_{t=0}^{S-1} T_r(\mathbf{x}\leftrightarrow\mathbf{x}_t)\sigma_s(\mathbf{x}_t)L_i(\mathbf{x}_t,\vec{\omega})\Delta_t$$

Irradiance From Media

$$E_m(\mathbf{x}) = \int_\Omega L_m(\mathbf{x}, \vec{\omega})\,(\vec{\mathbf{n}} \cdot \vec{\omega})\; d\vec{\omega}$$

where:

$$L_m(\mathbf{x}, \vec{\omega}) \approx \sum_{t=0}^{S-1} T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t)\sigma_s(\mathbf{x}_t)L_i(\mathbf{x}_t, \vec{\omega})\Delta_t$$

Irradiance From Media

$$E_m(\mathbf{x}) = \int_\Omega L_m(\mathbf{x}, \vec{\omega})\,(\vec{\mathbf{n}} \cdot \vec{\omega})\; d\vec{\omega}$$

where:

$$L_m(\mathbf{x}, \vec{\omega}) \approx \sum_{t=0}^{S-1} T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t)\sigma_s(\mathbf{x}_t)L_i(\mathbf{x}_t, \vec{\omega})\Delta_t$$

- Does not have an associated “distance”

Irradiance From Media

$$E_m(\mathbf{x}) = \int_\Omega L_m(\mathbf{x}, \vec{\omega})\,(\vec{\mathbf{n}} \cdot \vec{\omega})\; d\vec{\omega}$$

where:

$$L_m(\mathbf{x}, \vec{\omega}) \approx \sum_{t=0}^{S-1} T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t)\sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega})\Delta_t$$

- Does not have an associated "distance"
- Cannot use the same gradient formulation

Hemispherical Sampling

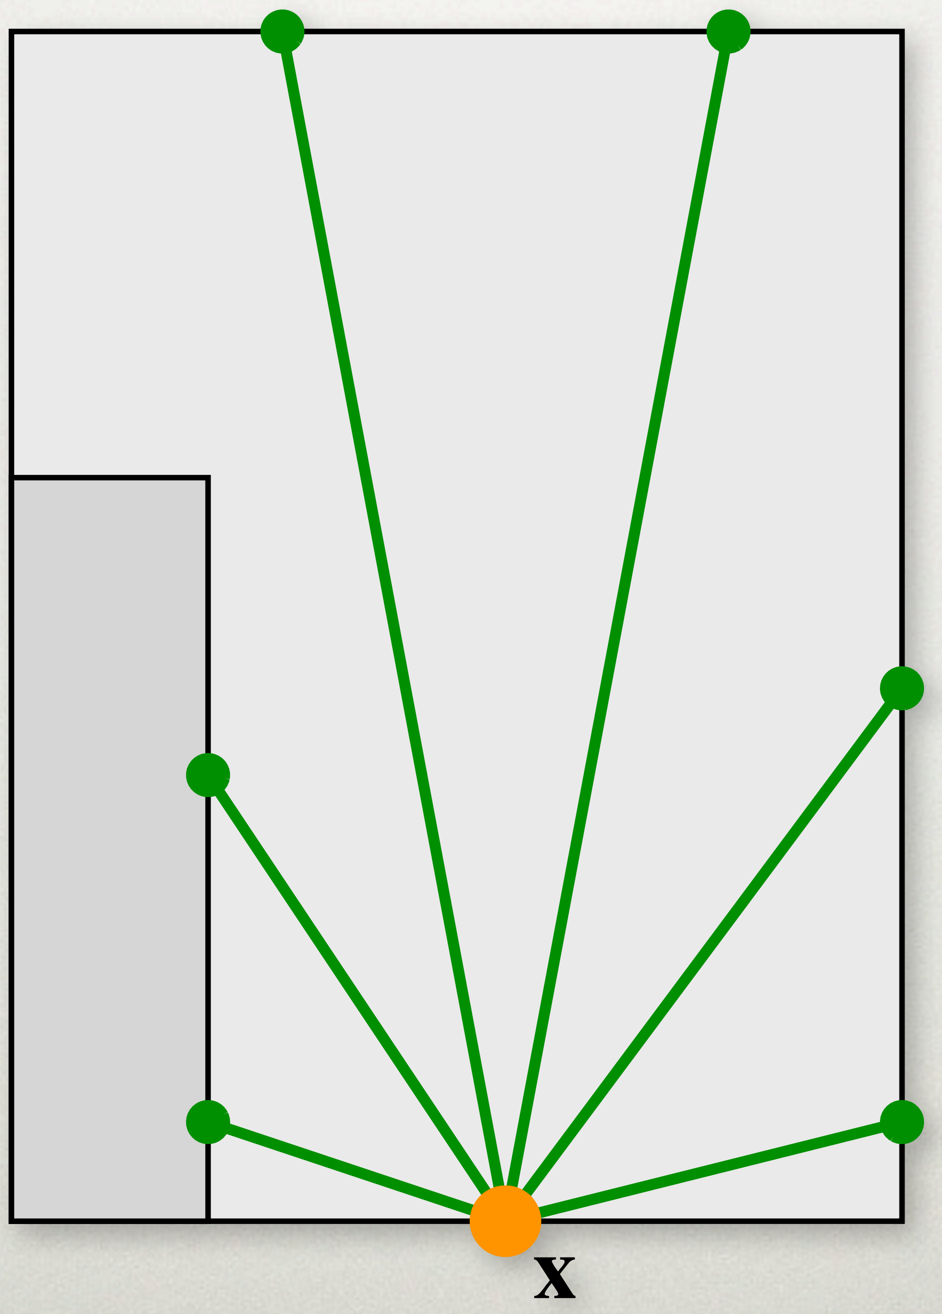

Thursday, 6 September 12

* I’ll describe the process of computing the media irradiance gradient at a high level using this 2D example
* the details of this process are in the paper.

Ray Marching

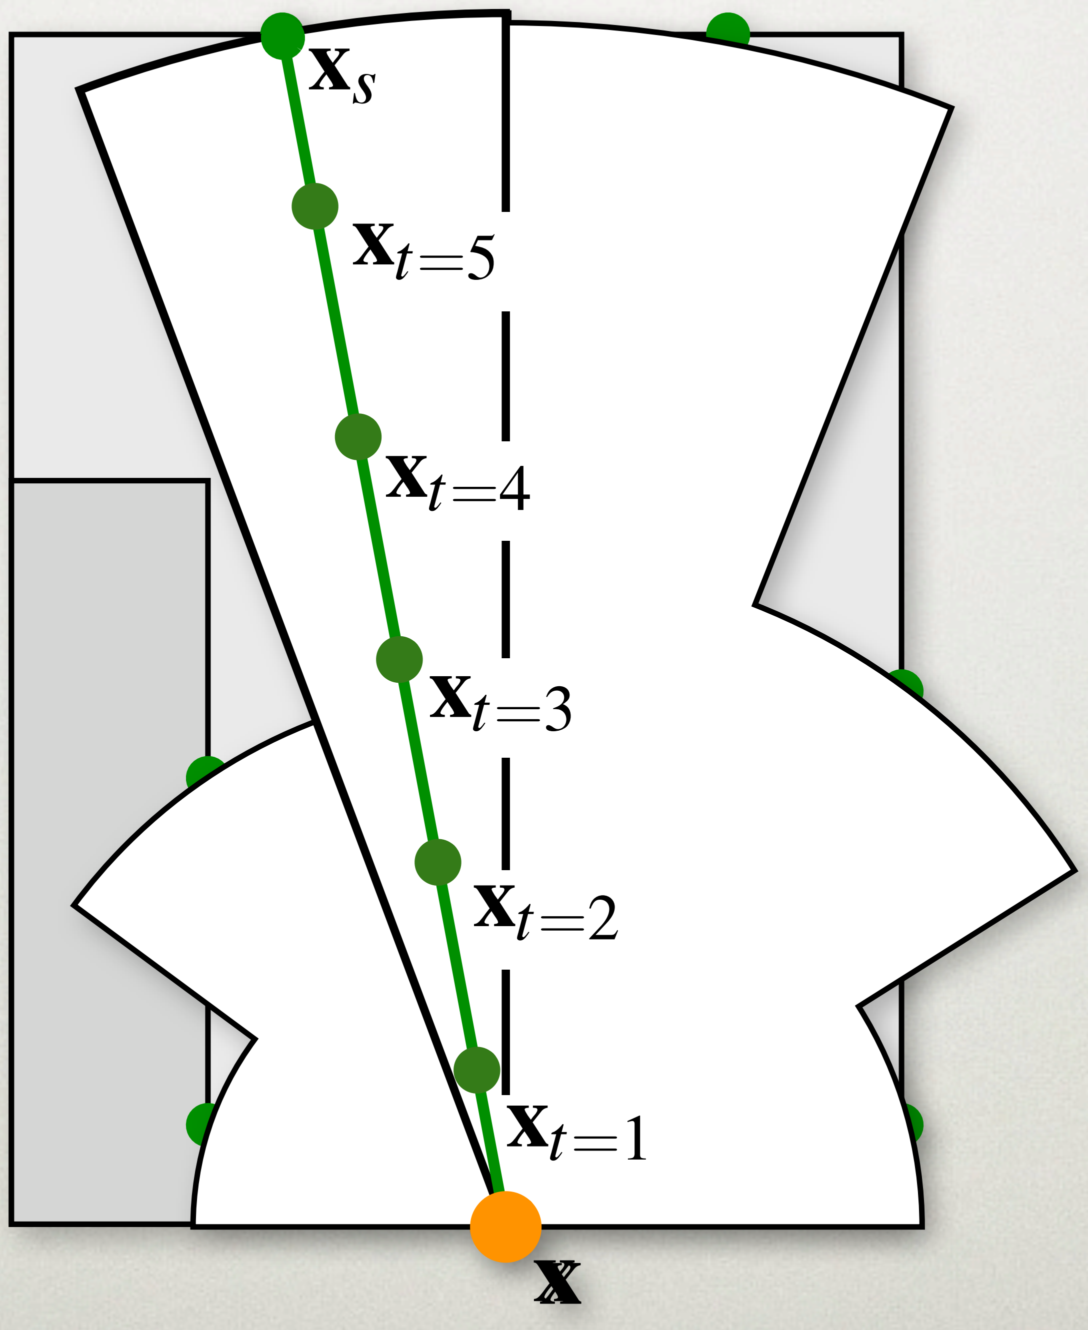

* An individual cell in this case samples the medium at multiple steps using ray marching.

Ray Marching

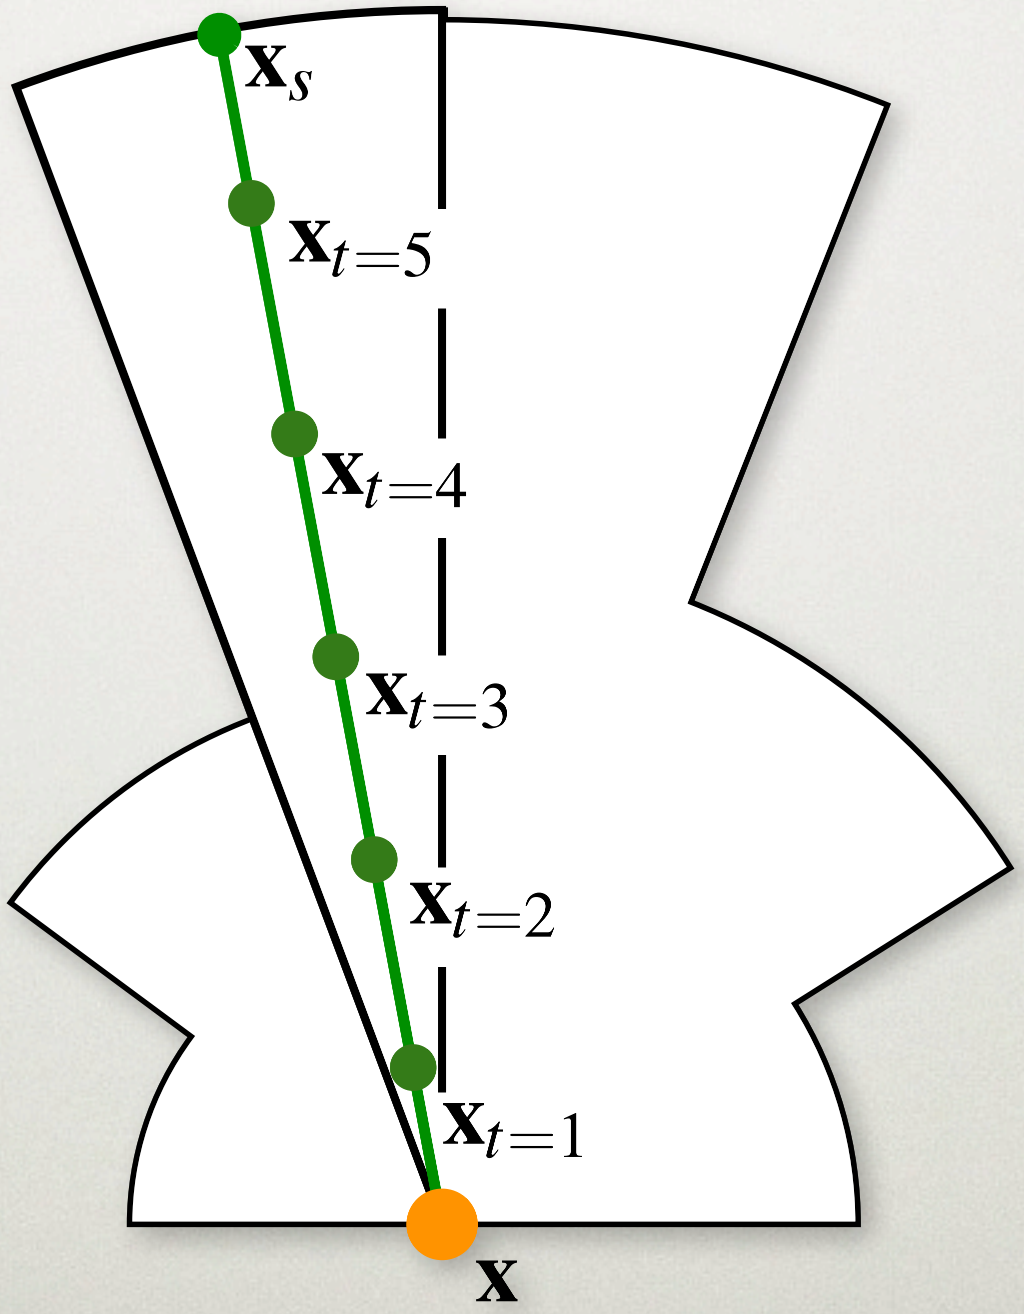

Thursday, 6 September 12

* An individual cell in this case samples the medium at multiple steps using ray marching.

Ray Marching

Medium radiance interpreted as coming from “shells”

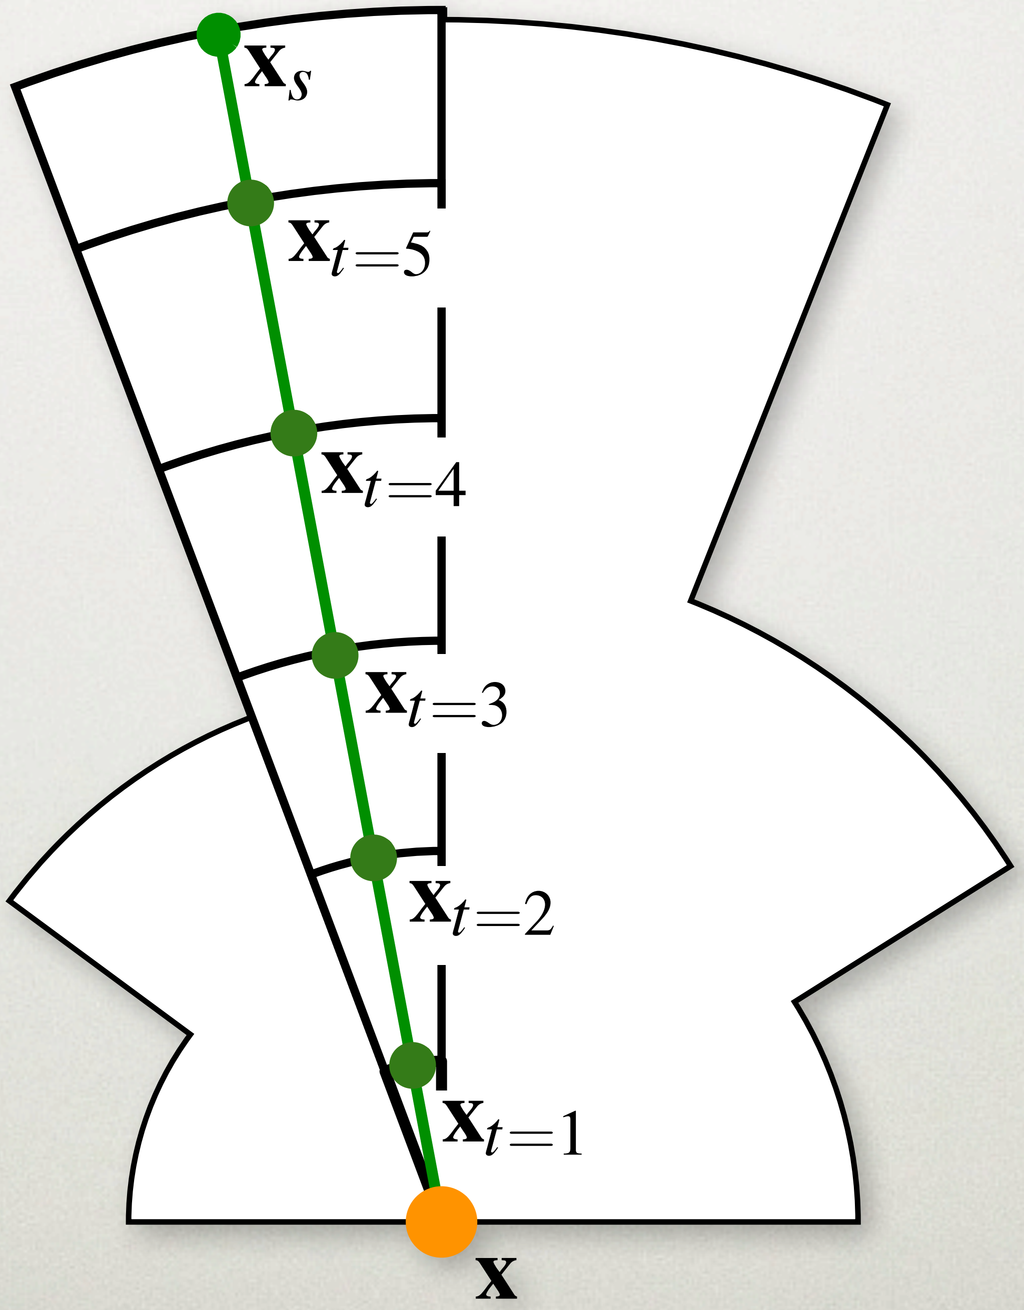

Thursday, 6 September 12

* In order to compute the contribution to the gradient for each cell, we interpret these samples as radiance come from multiple shells of different radius

Ray Marching

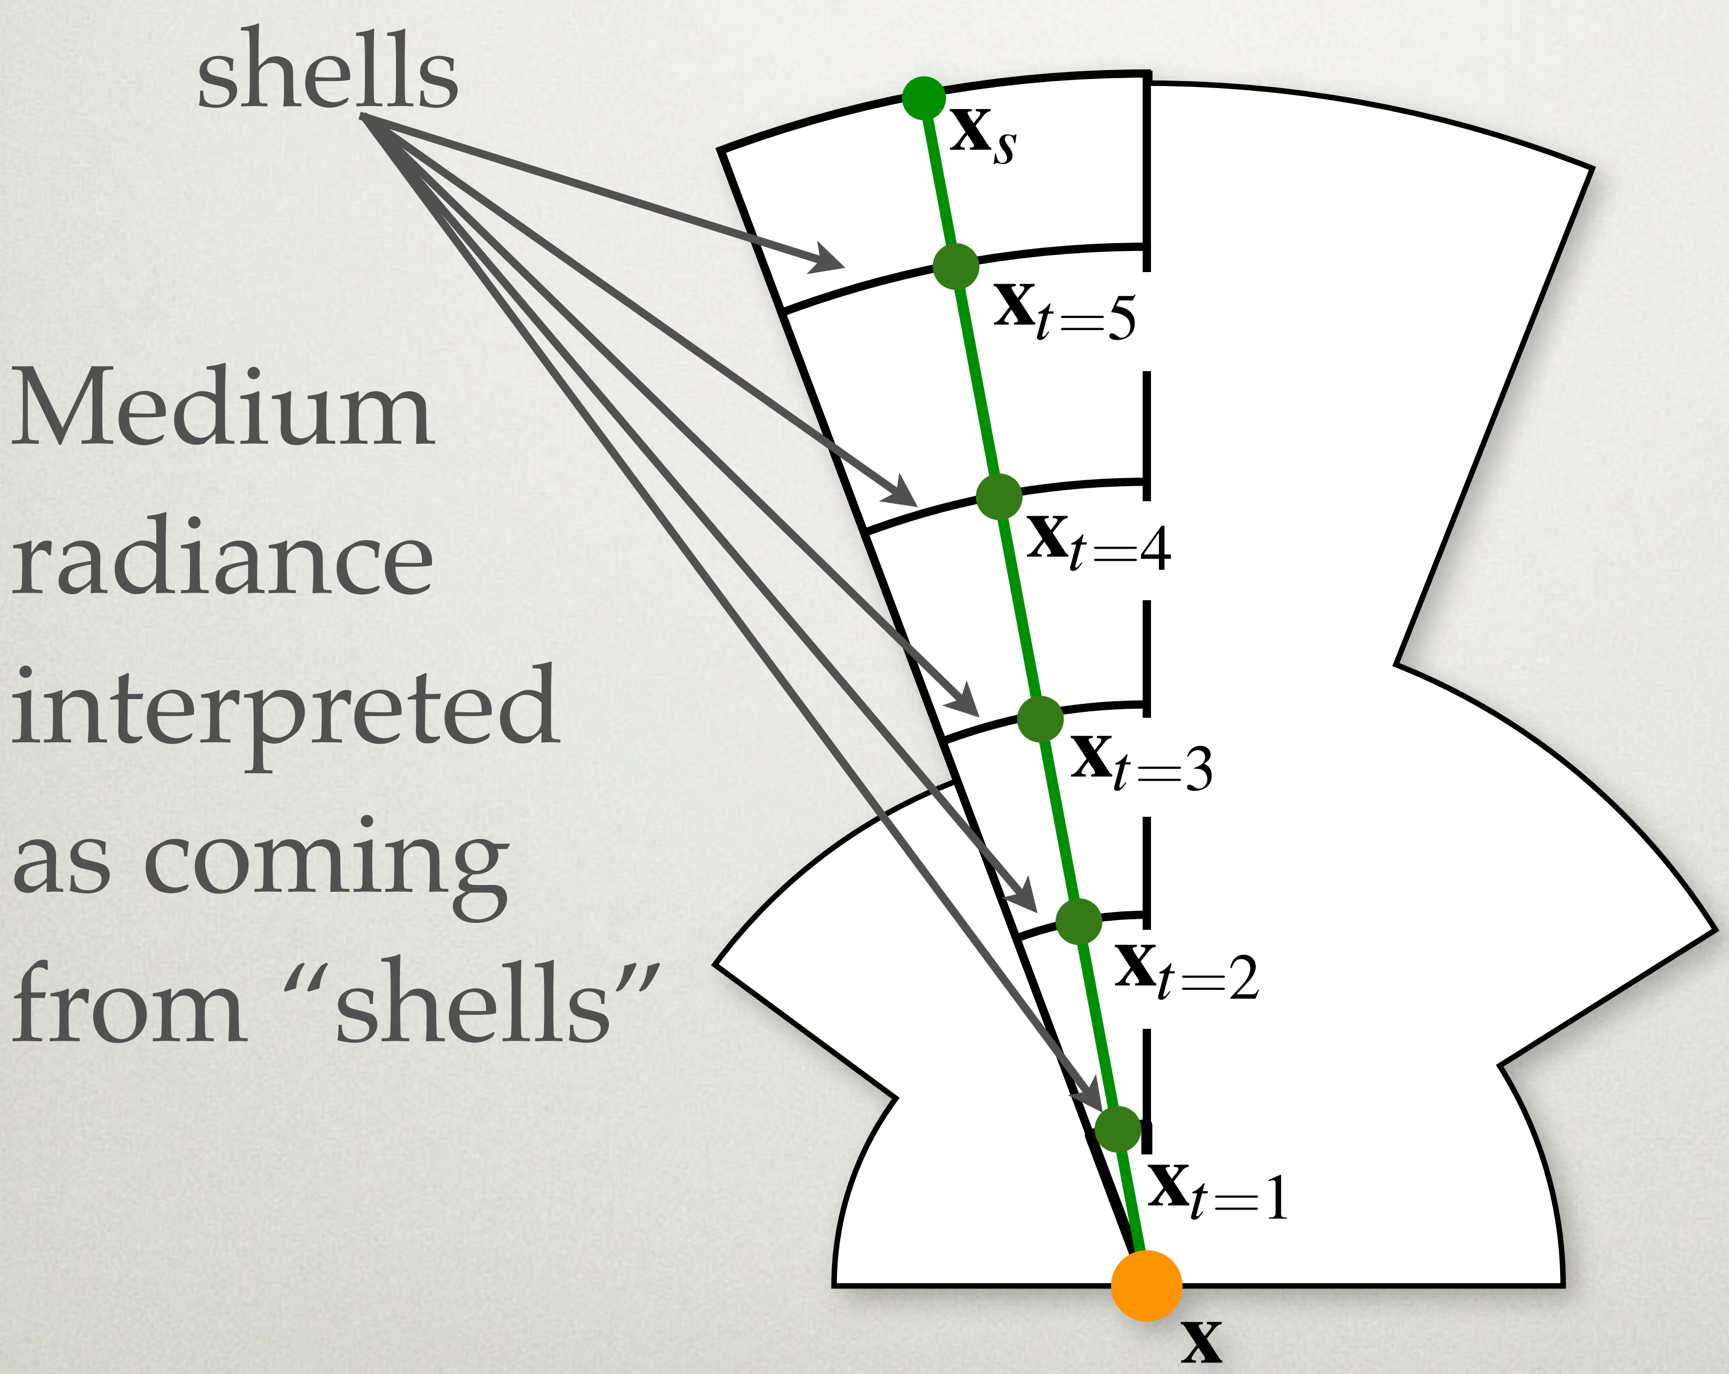

Thursday, 6 September 12

* In order to compute the contribution to the gradient for each cell, we interpret these samples as radiance come from multiple shells of different radius

Ray Marching

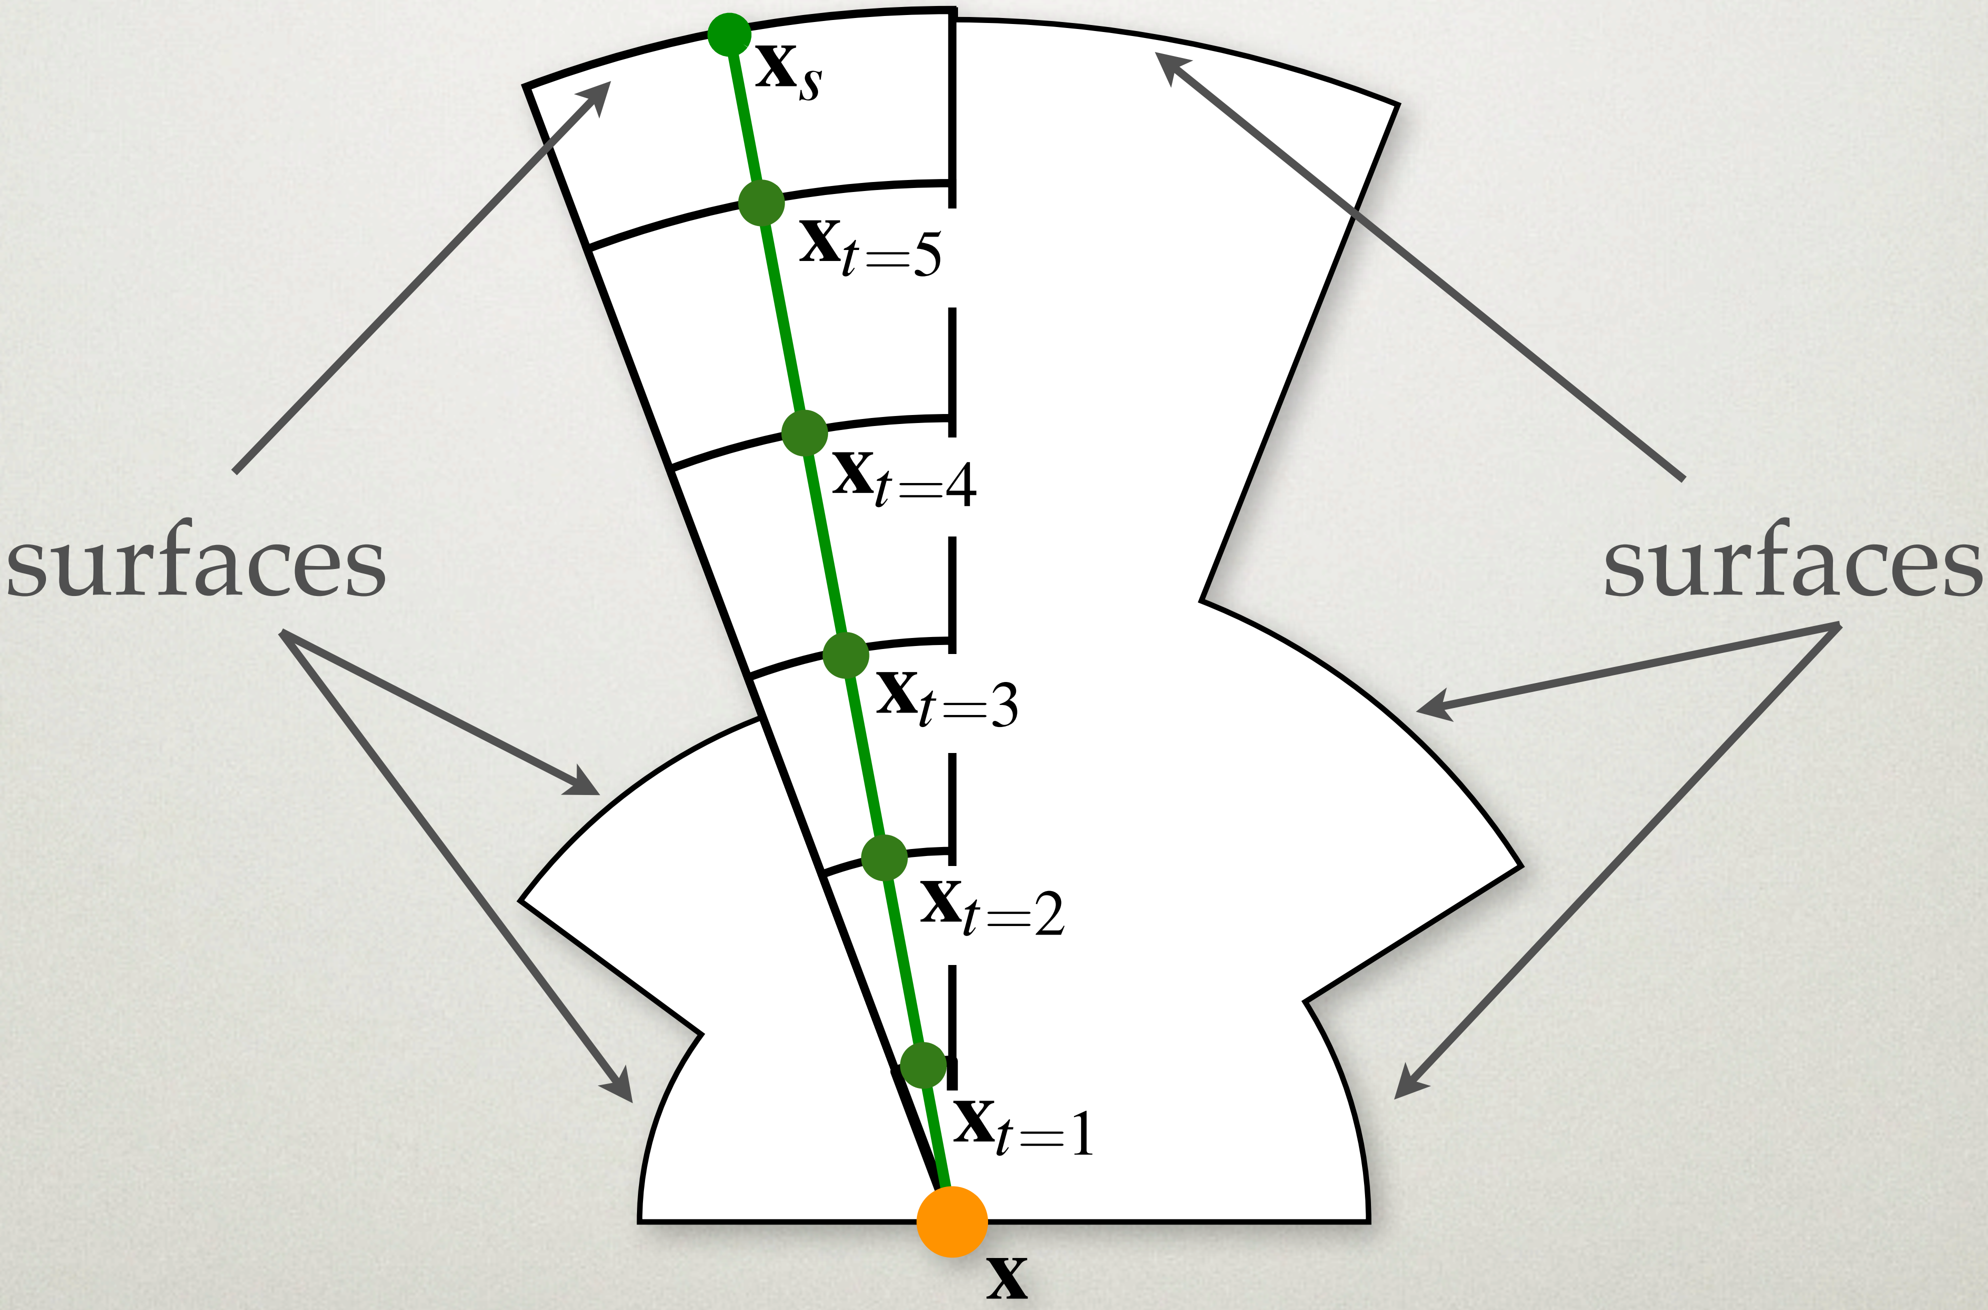

Thursday, 6 September 12

* We also want to handle occlusions from neighboring surfaces, like in the Ward and Heckbert formulation

Media Irradiance Gradient

Each shell is occluded by surfaces at a different rate as we move **x**.

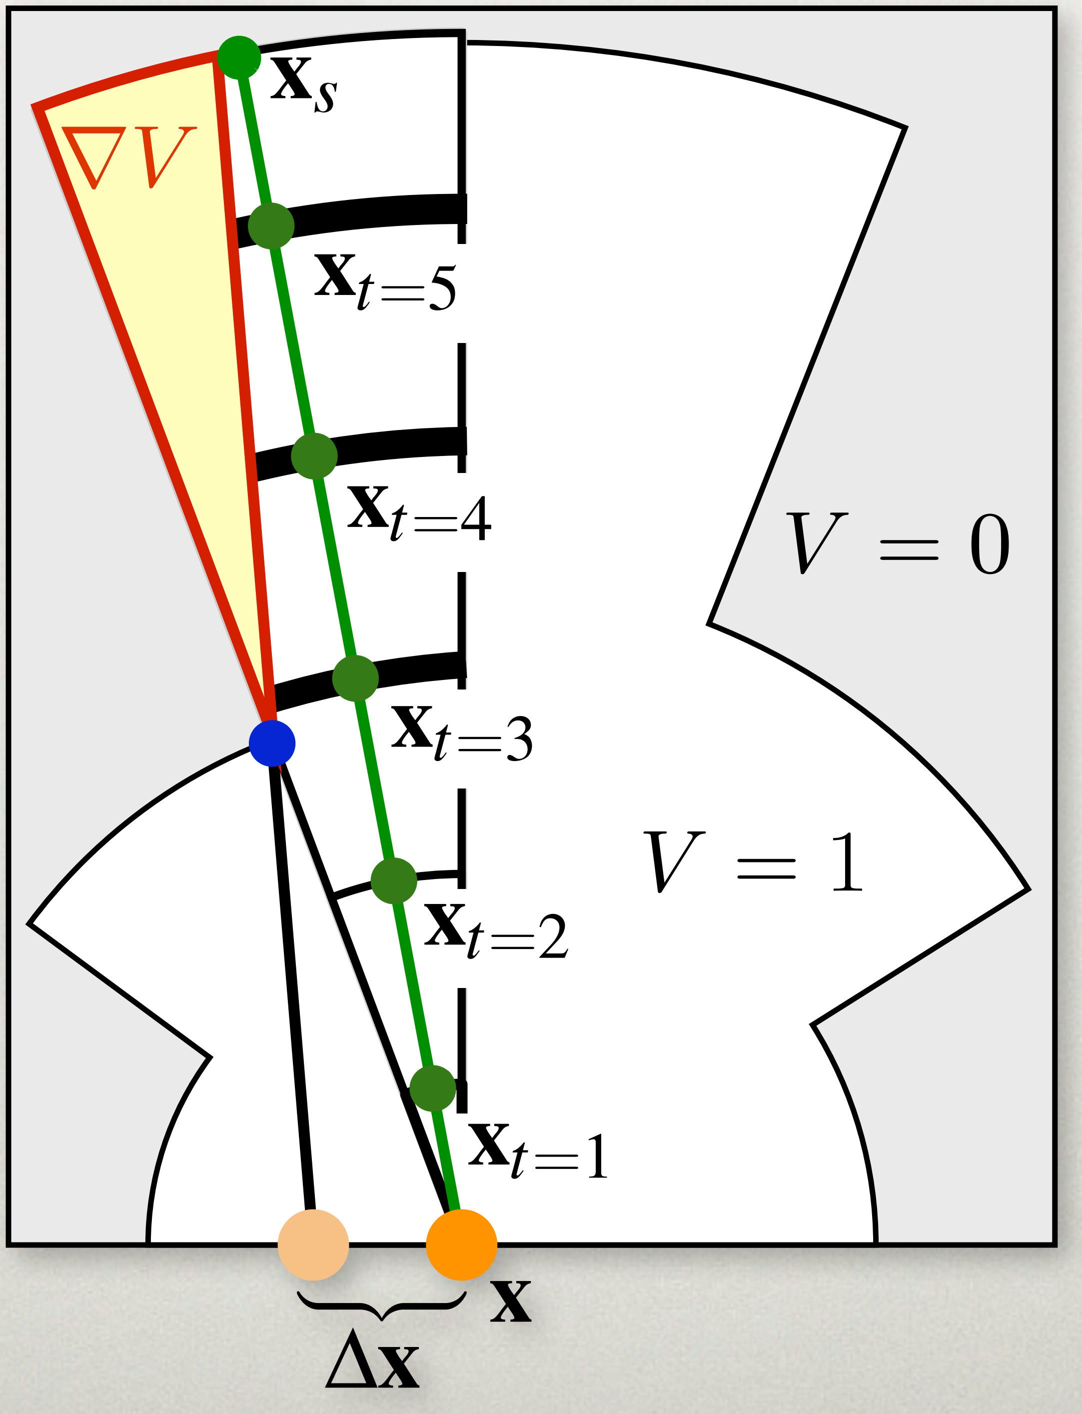

Thursday, 6 September 12

* To compute the gradient contribution of each cell, we determine how each shell may be occluded by surfaces.
* The rate of occlusion depends on the distance to the “shell,” and the distance to the neighboring surface causing the occlusion (shown in blue)
* This means that shells in front of neighboring surfaces do not get occluded with translation, and shells past surfaces get occluded faster with increased distance.
* The media irradiance gradient can be thought of as applying the Ward and Heckbert gradient formulation to estimate the change in occlusion individually for each of these shells of increasing radius.

Media Irradiance Gradient

Each shell is occluded by surfaces at a different rate as we move **x**.

Determined by distance to shell and to occluder.

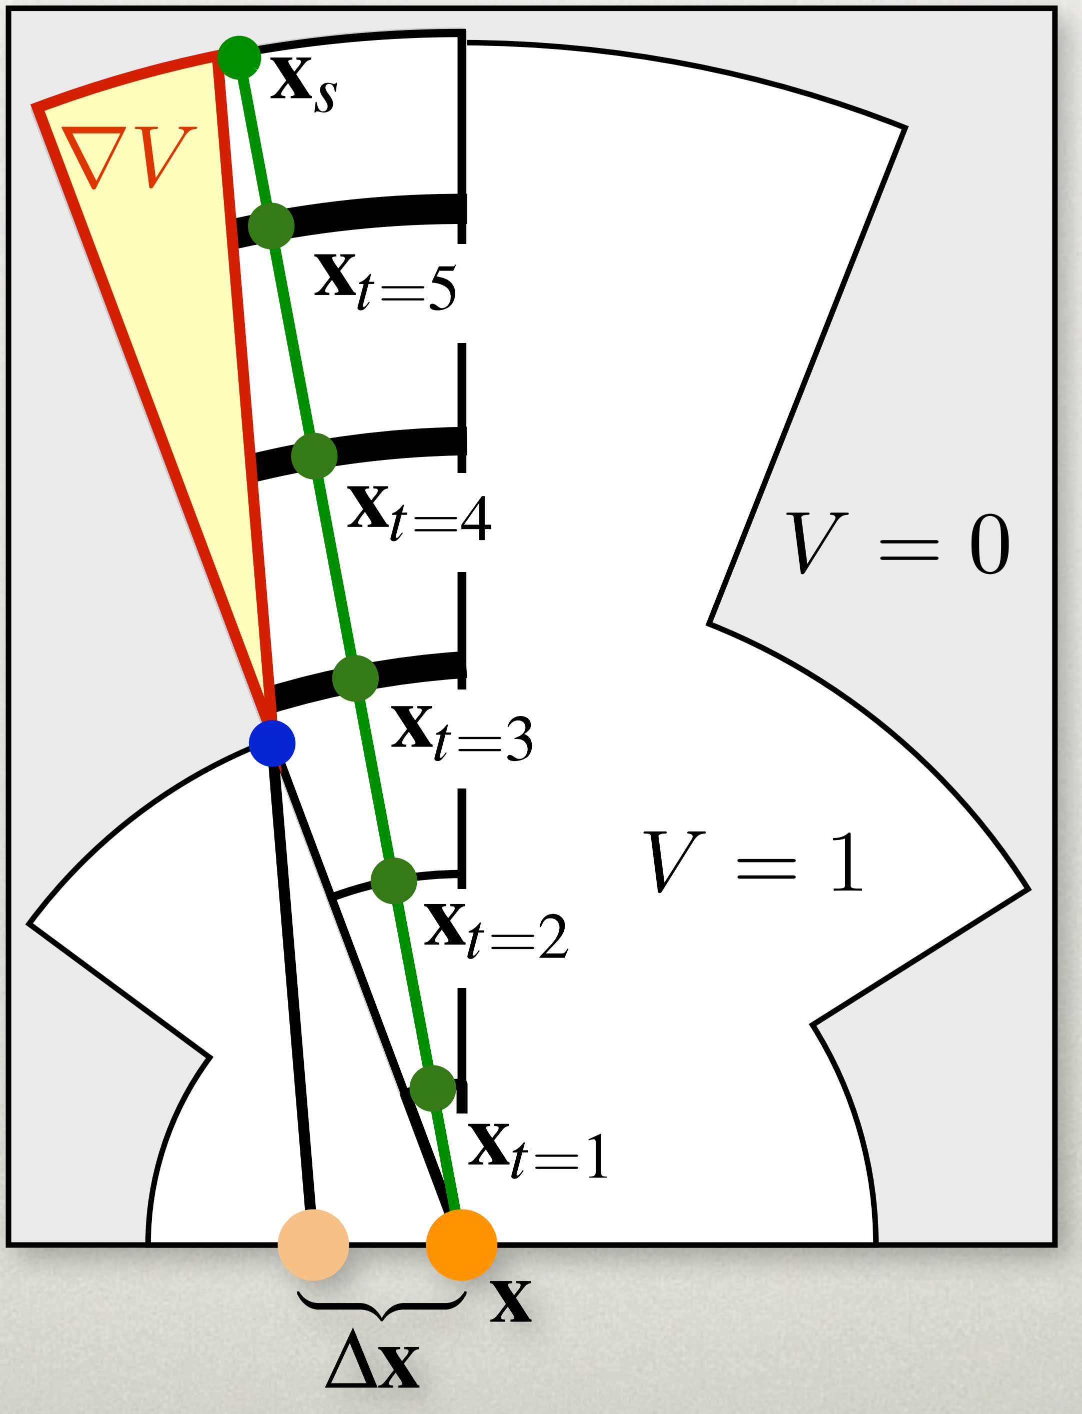

Shells do not occlude each other

* To compute the gradient contribution of each cell, we determine how each shell may be occluded by surfaces.
* The rate of occlusion depends on the distance to the “shell,” and the distance to the neighboring surface causing the occlusion (shown in blue)
* This means that shells in front of neighboring surfaces do not get occluded with translation, and shells past surfaces get occluded faster with increased distance.
* The media irradiance gradient can be thought of as applying the Ward and Heckbert gradient formulation to estimate the change in occlusion individually for each of these shells of increasing radius.

Scattering Medium

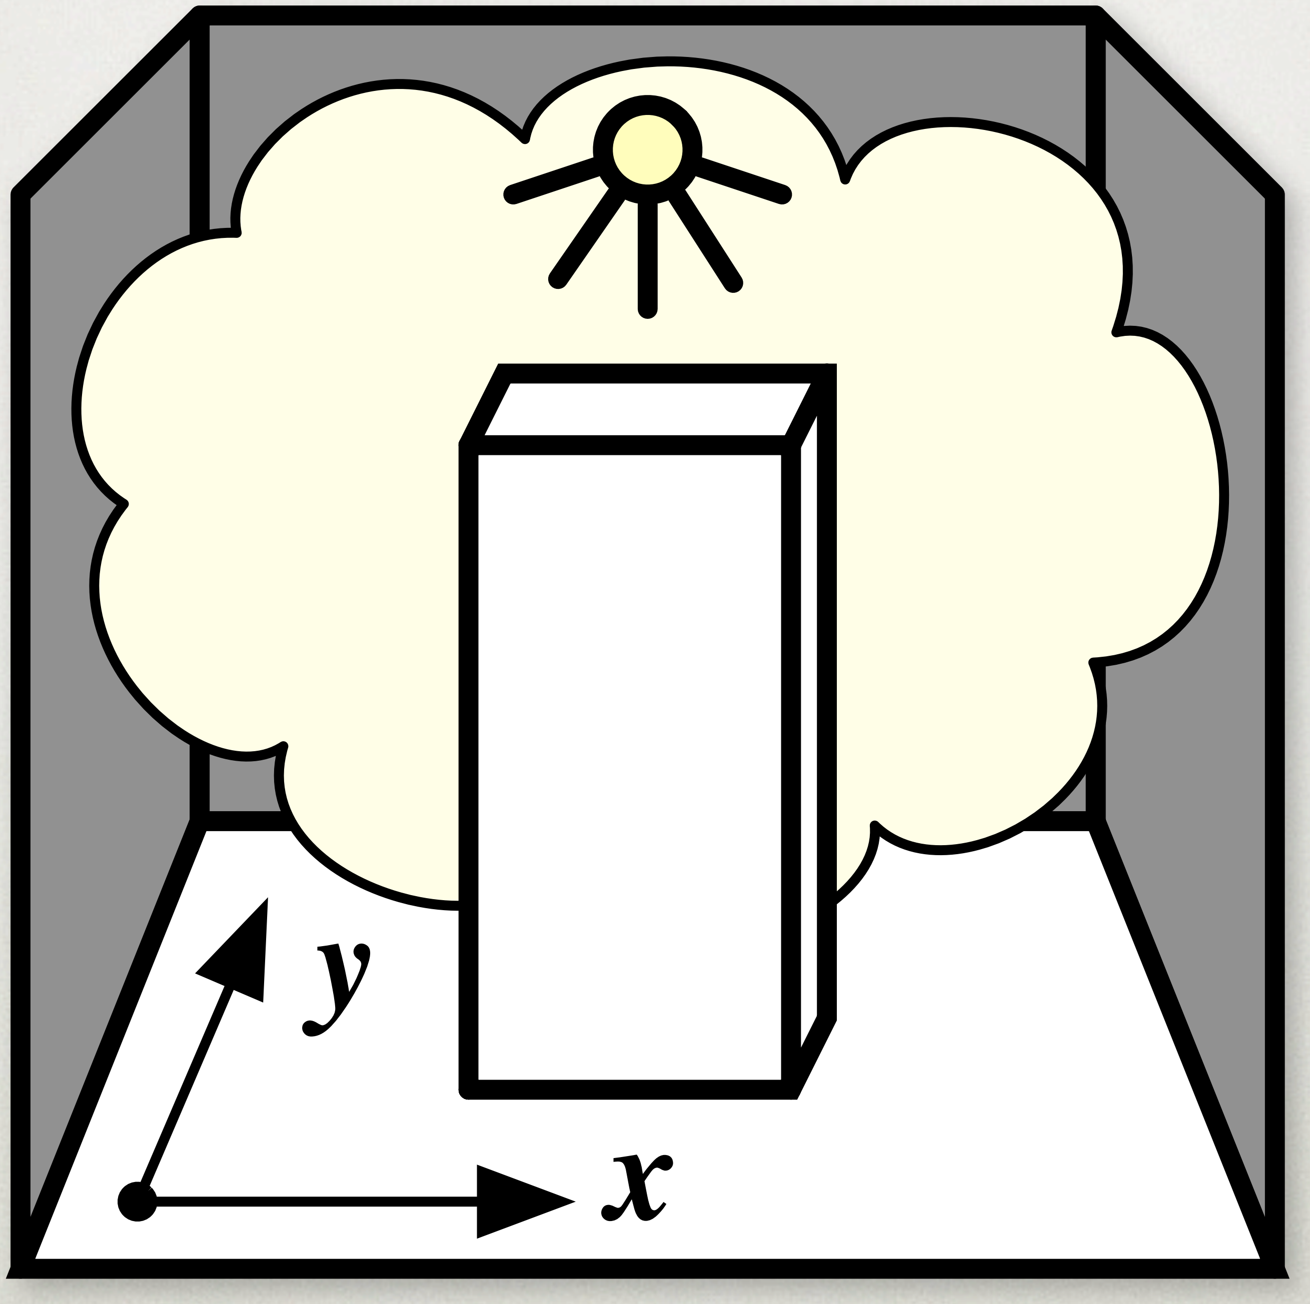

Thursday, 6 September 12

* Using a modification of the previous scene, we can validate the correctness of our media irradiance gradients.
* In this case, we use a scattering media, and a point light source.
* The scene is constructed in a way where all lighting on the ground plane has first scattered within the medium.

Per-Pixel Irradiance Gradient

$$\text{red} = \left|\frac{\partial E}{\partial x}\right|$$

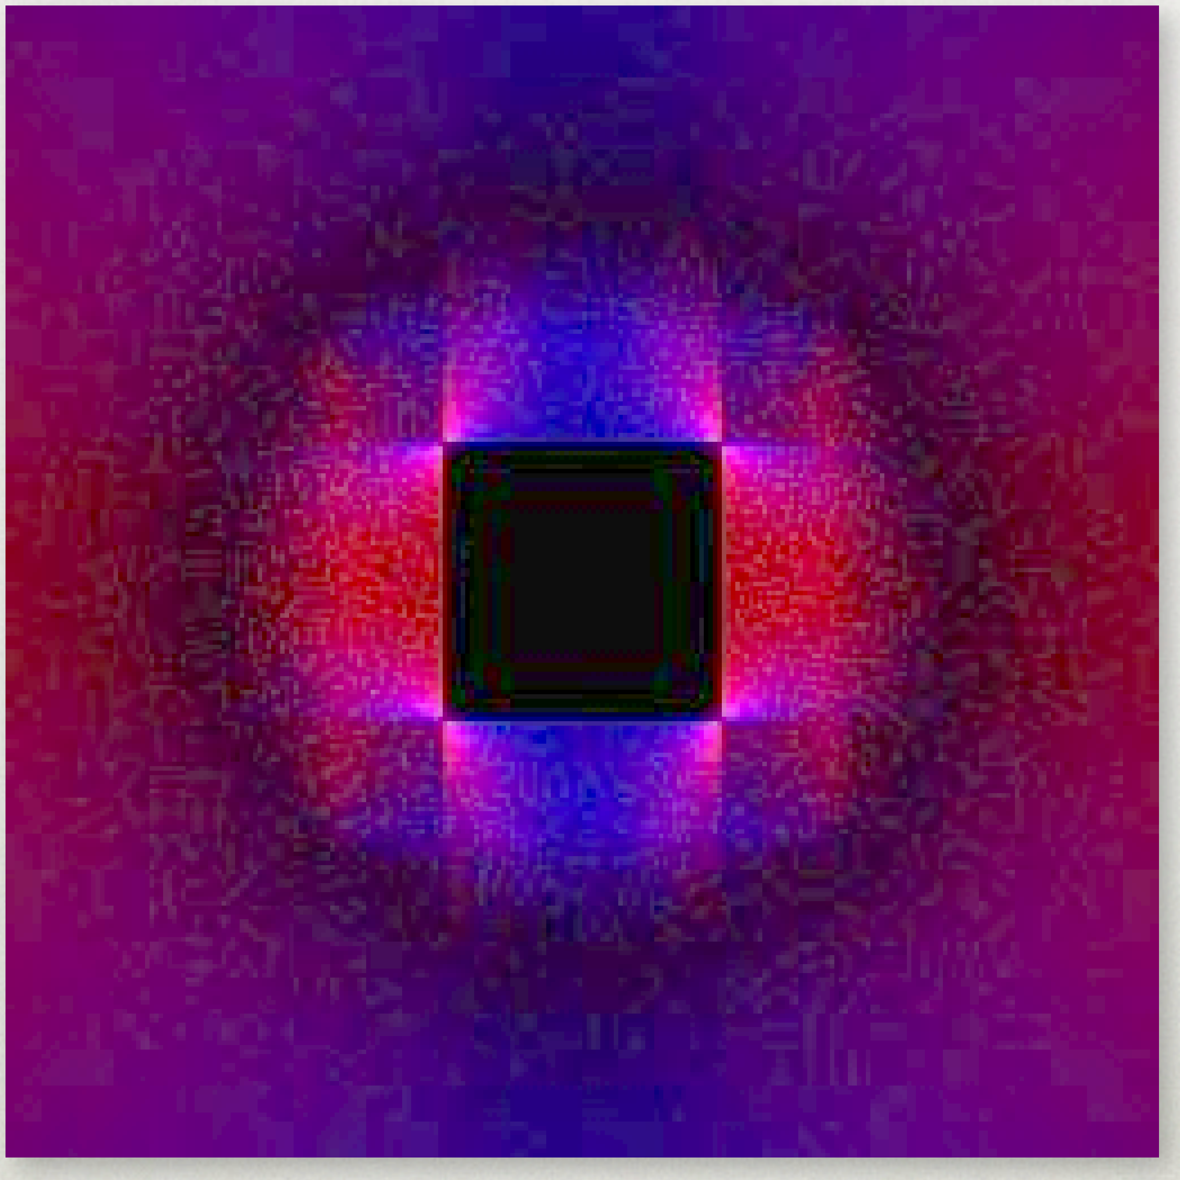

$$\text{blue} = \left|\frac{\partial E}{\partial y}\right|$$

(Finite Differences 10X)

* We compute a ground truth gradient using finite differences
* Even with a very large number of samples the finite difference gradient suffers from significant noise

Per-Pixel Irradiance Gradient

$$\text{red} = \left|\frac{\partial E}{\partial x}\right|$$

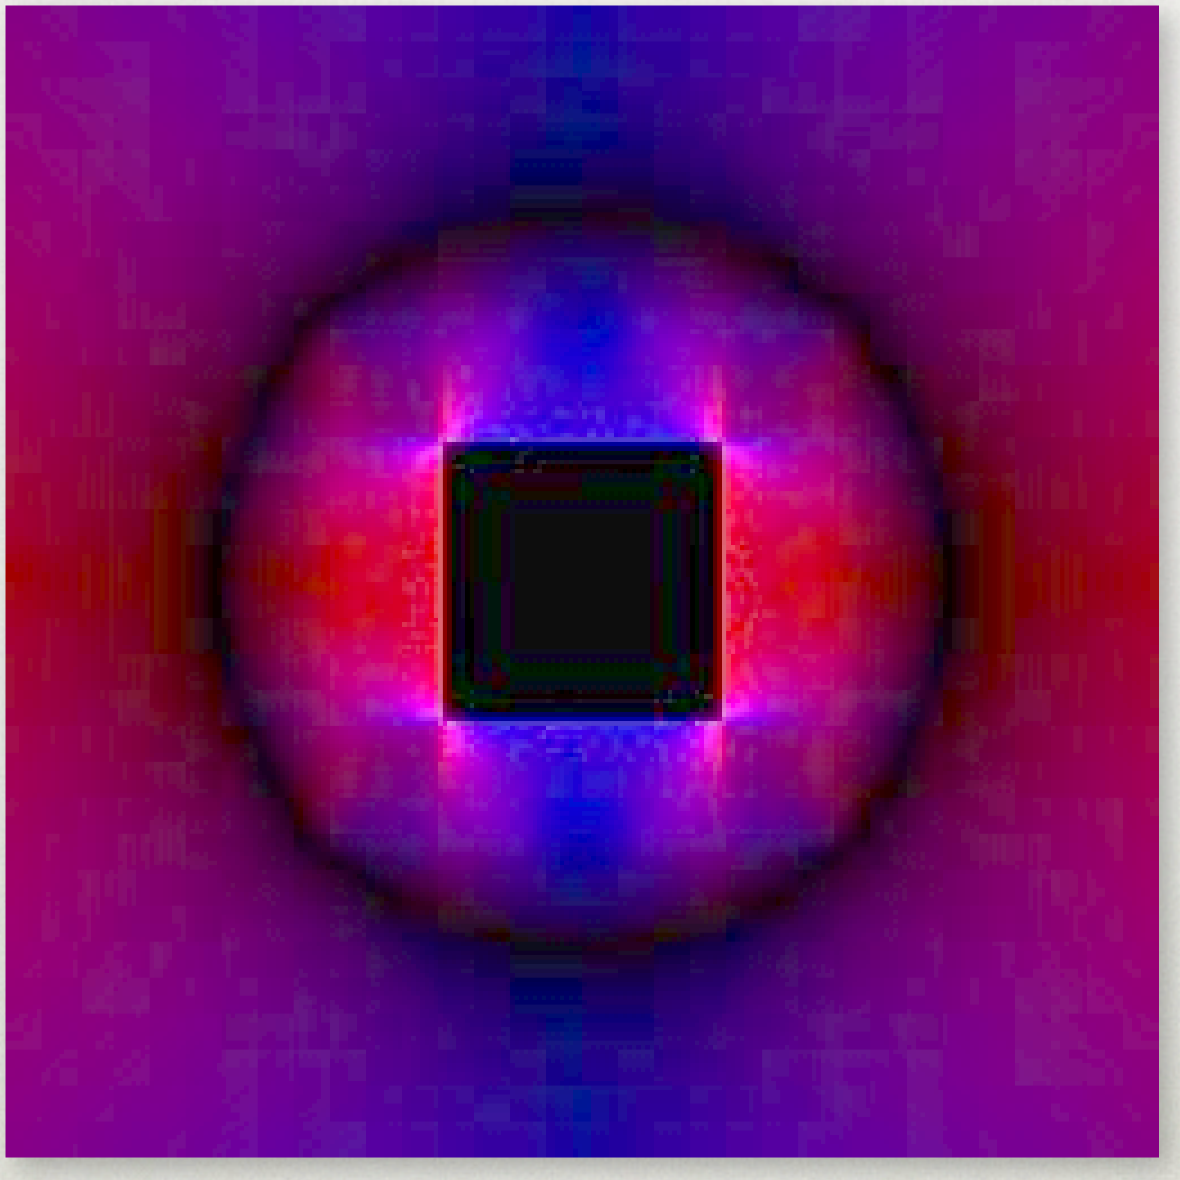

$$\text{blue} = \left|\frac{\partial E}{\partial y}\right|$$

(Our Method)

* The gradients estimated using our method match the behavior of the true gradient but have significantly less noise using only 1/10th of the number of samples

Per-Pixel Irradiance Gradient

$$\text{red} = \left|\frac{\partial E}{\partial x}\right|$$

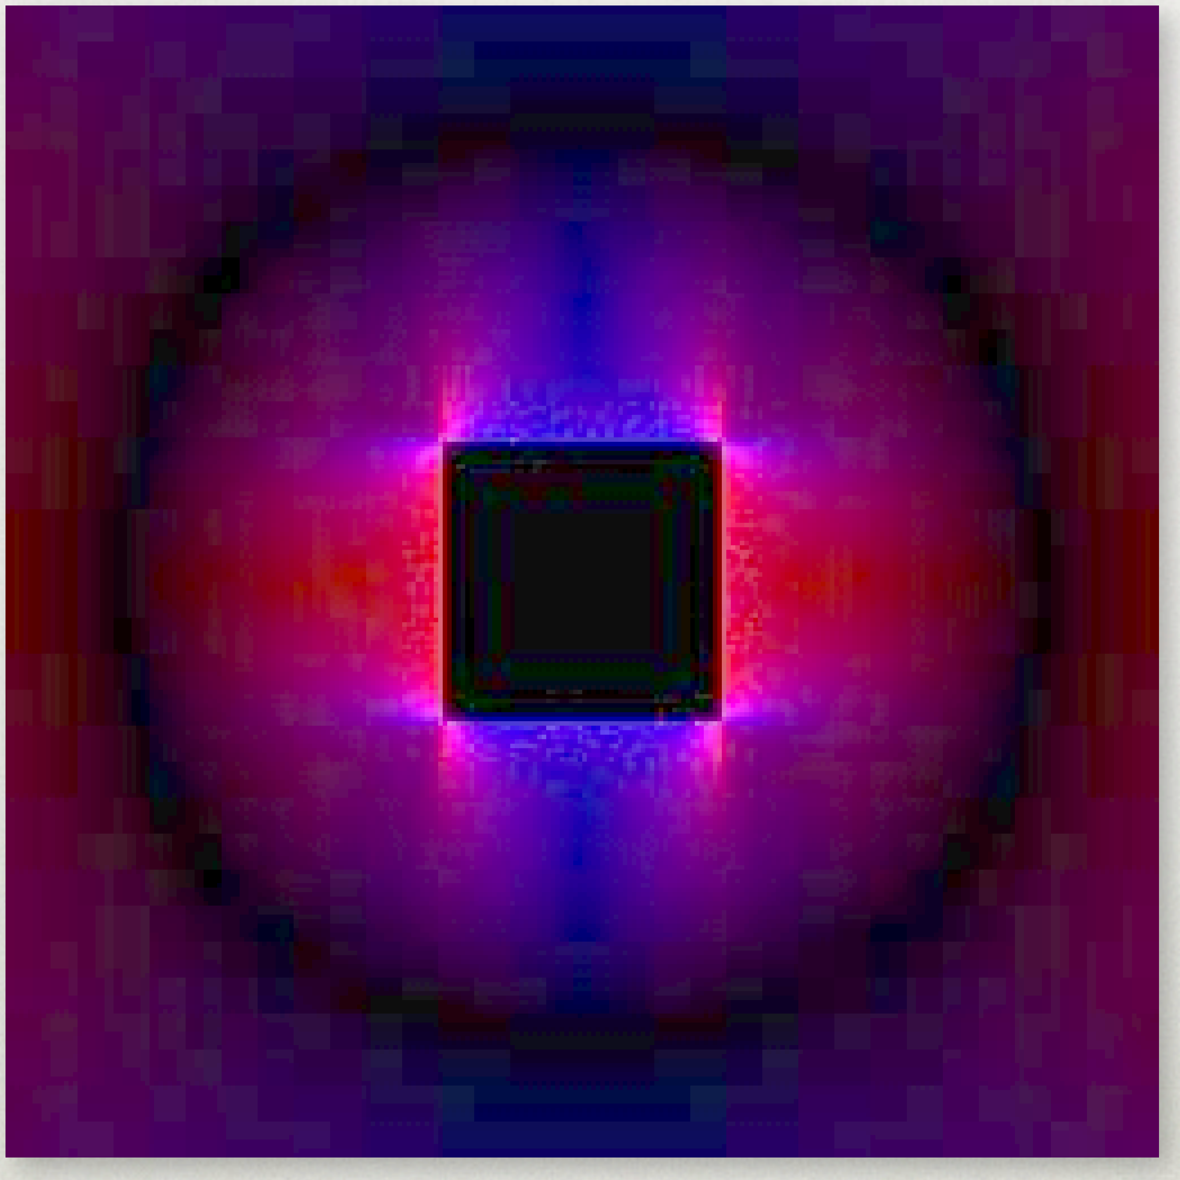

$$\text{blue} = \left|\frac{\partial E}{\partial y}\right|$$

(Ward and Heckbert)

* The Ward and Heckbert gradient formulation again significantly differs from the true gradient since it does not take into account media scattering.

Per-Pixel Irradiance Gradient

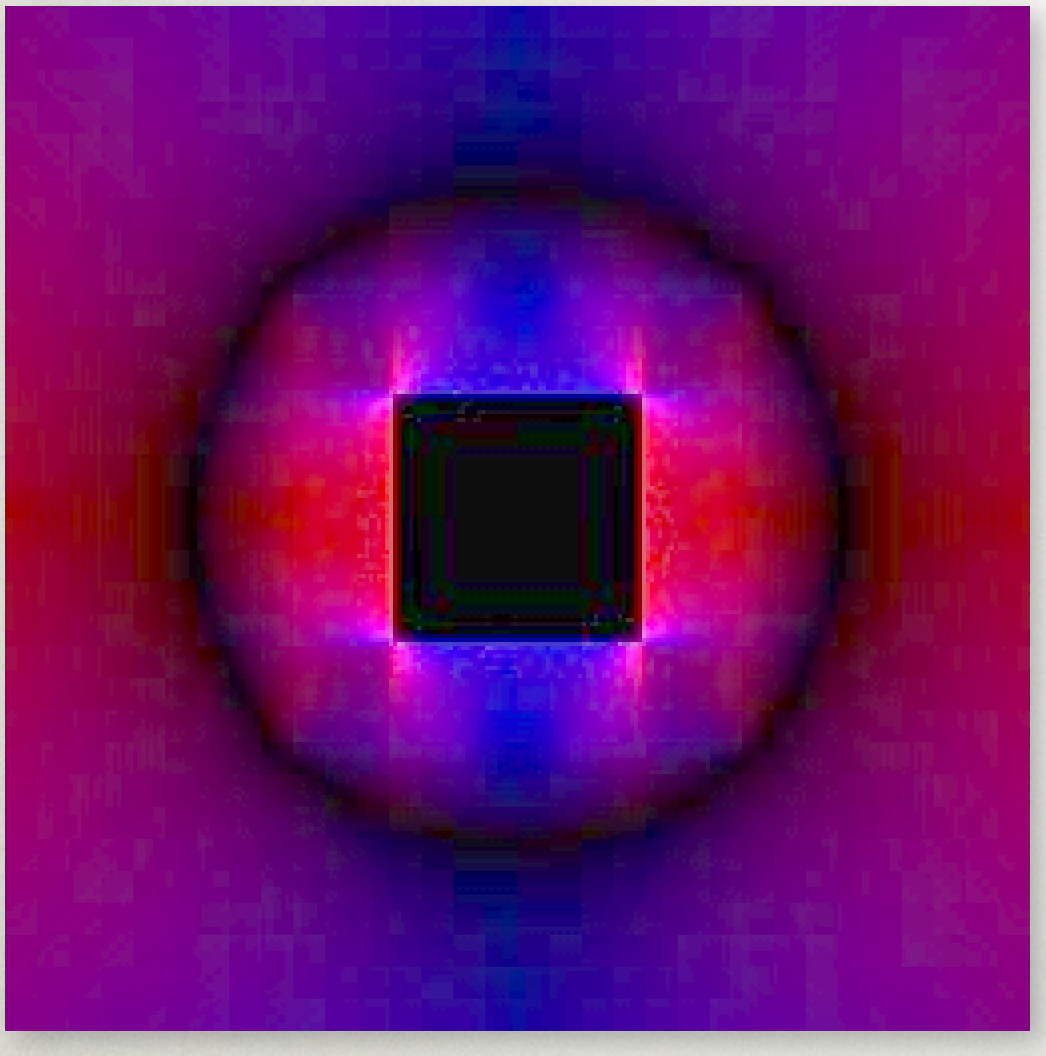

Our Method

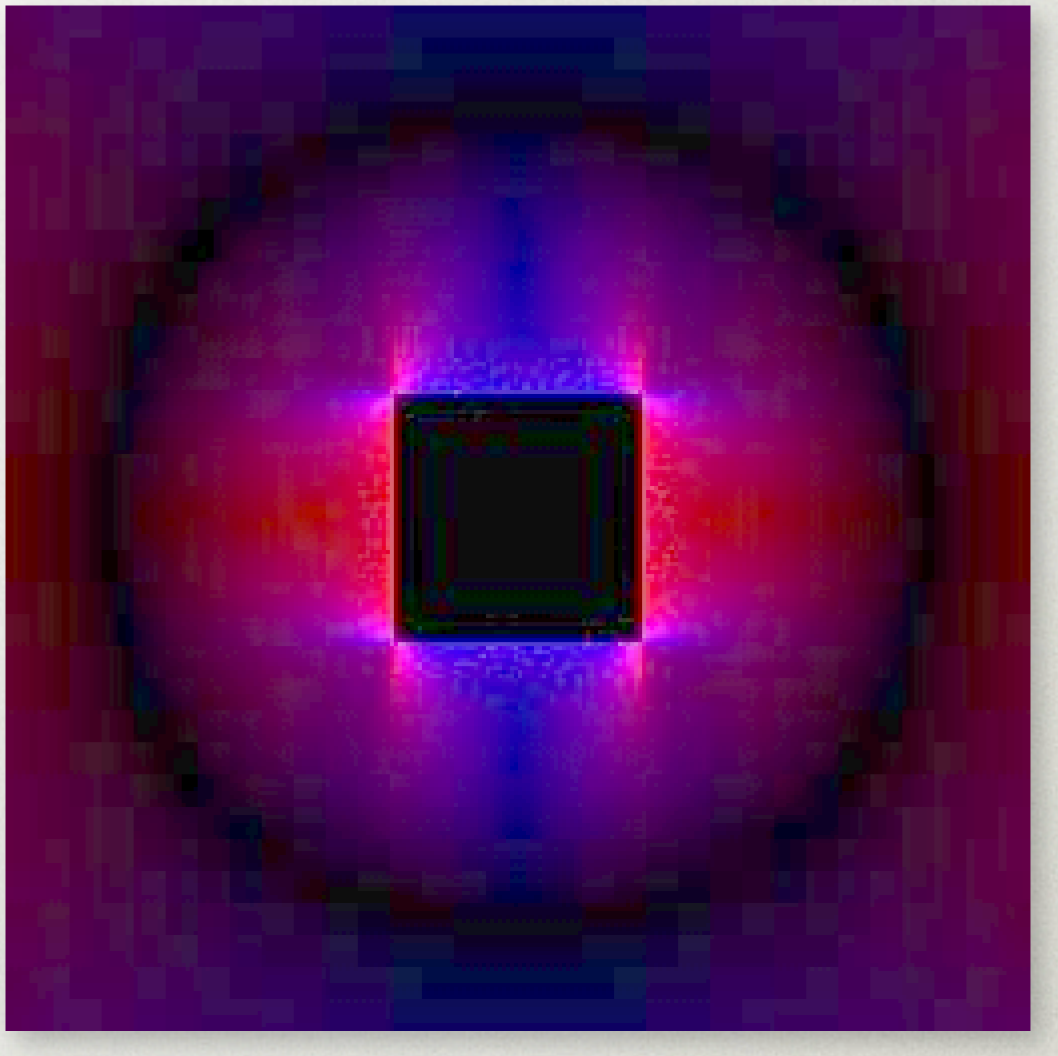

Ward and Heckbert

Why is the Ward & Heckbert gradient darker?

* In addition to having a different structure, it is also overall darker.
* This is because for Ward & Heckbert, all radiance is assumed to come from the surface past the medium.
* and since the gradient is inversely proportional to the distance, this underestimates the gradient.
* By comparing the gradients along a single scanline

Per-Pixel Irradiance Gradient

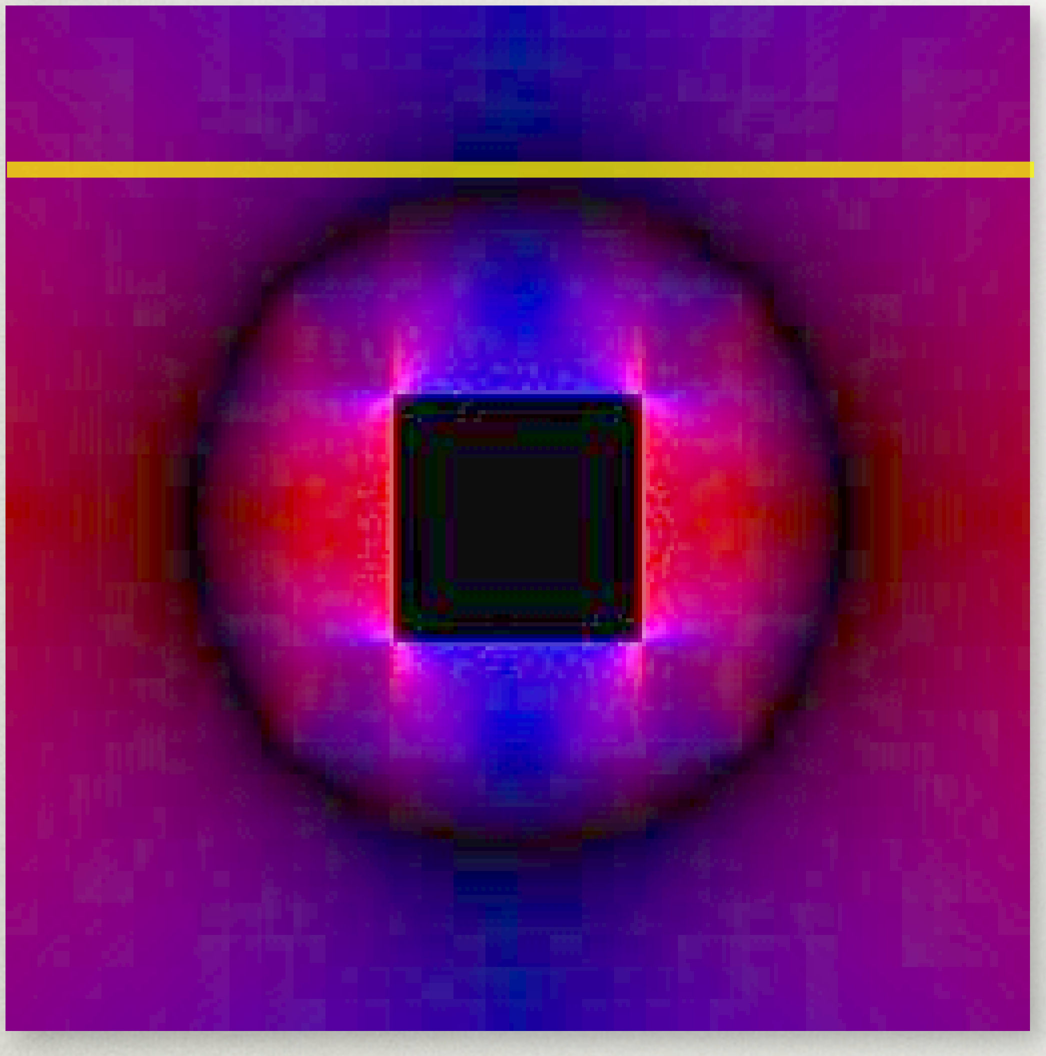

Our Method

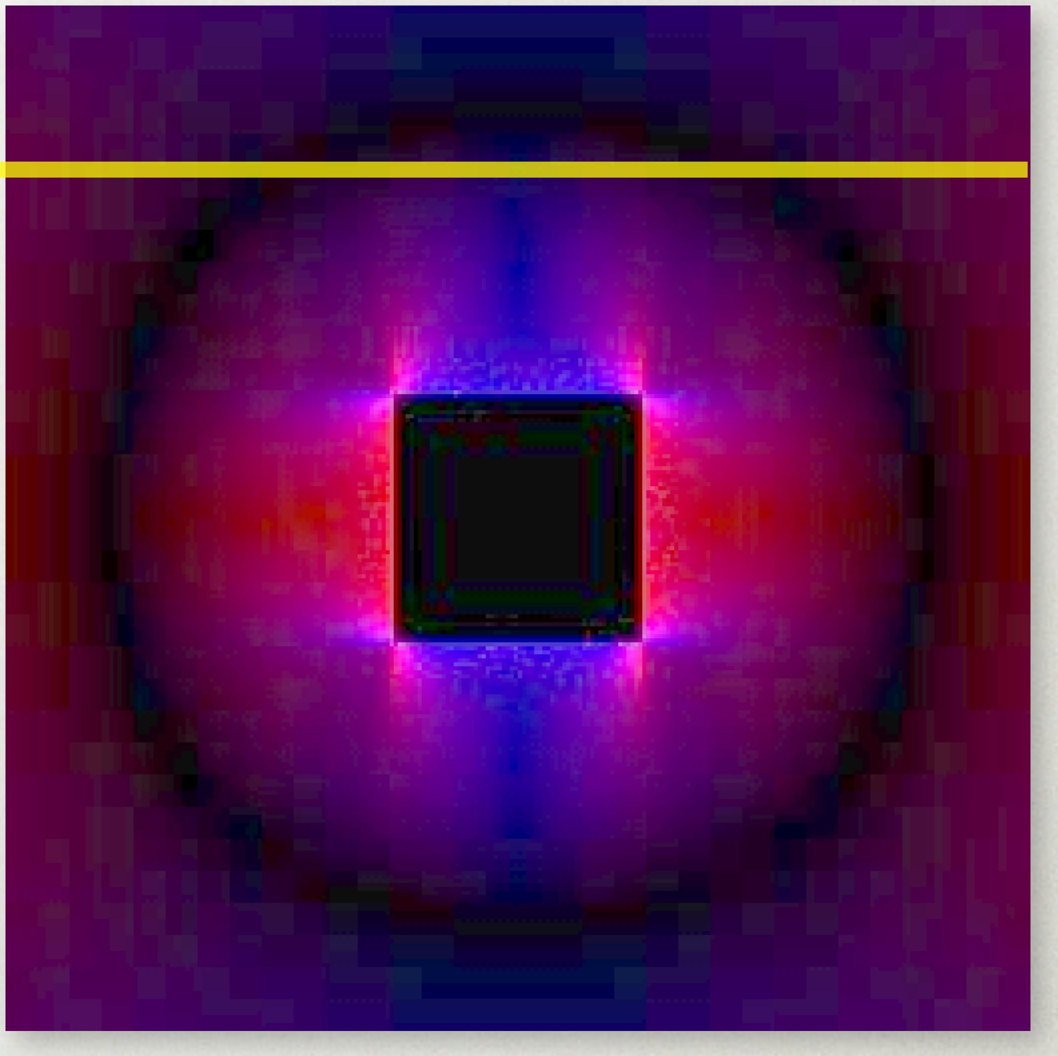

Ward and Heckbert

Why is the Ward & Heckbert gradient darker?

Thursday, 6 September 12

* In addition to having a different structure, it is also overall darker.
* This is because for Ward & Heckbert, all radiance is assumed to come from the surface past the medium.
* and since the gradient is inversely proportional to the distance, this underestimates the gradient.
* By comparing the gradients along a single scanline

Gradient Comparison

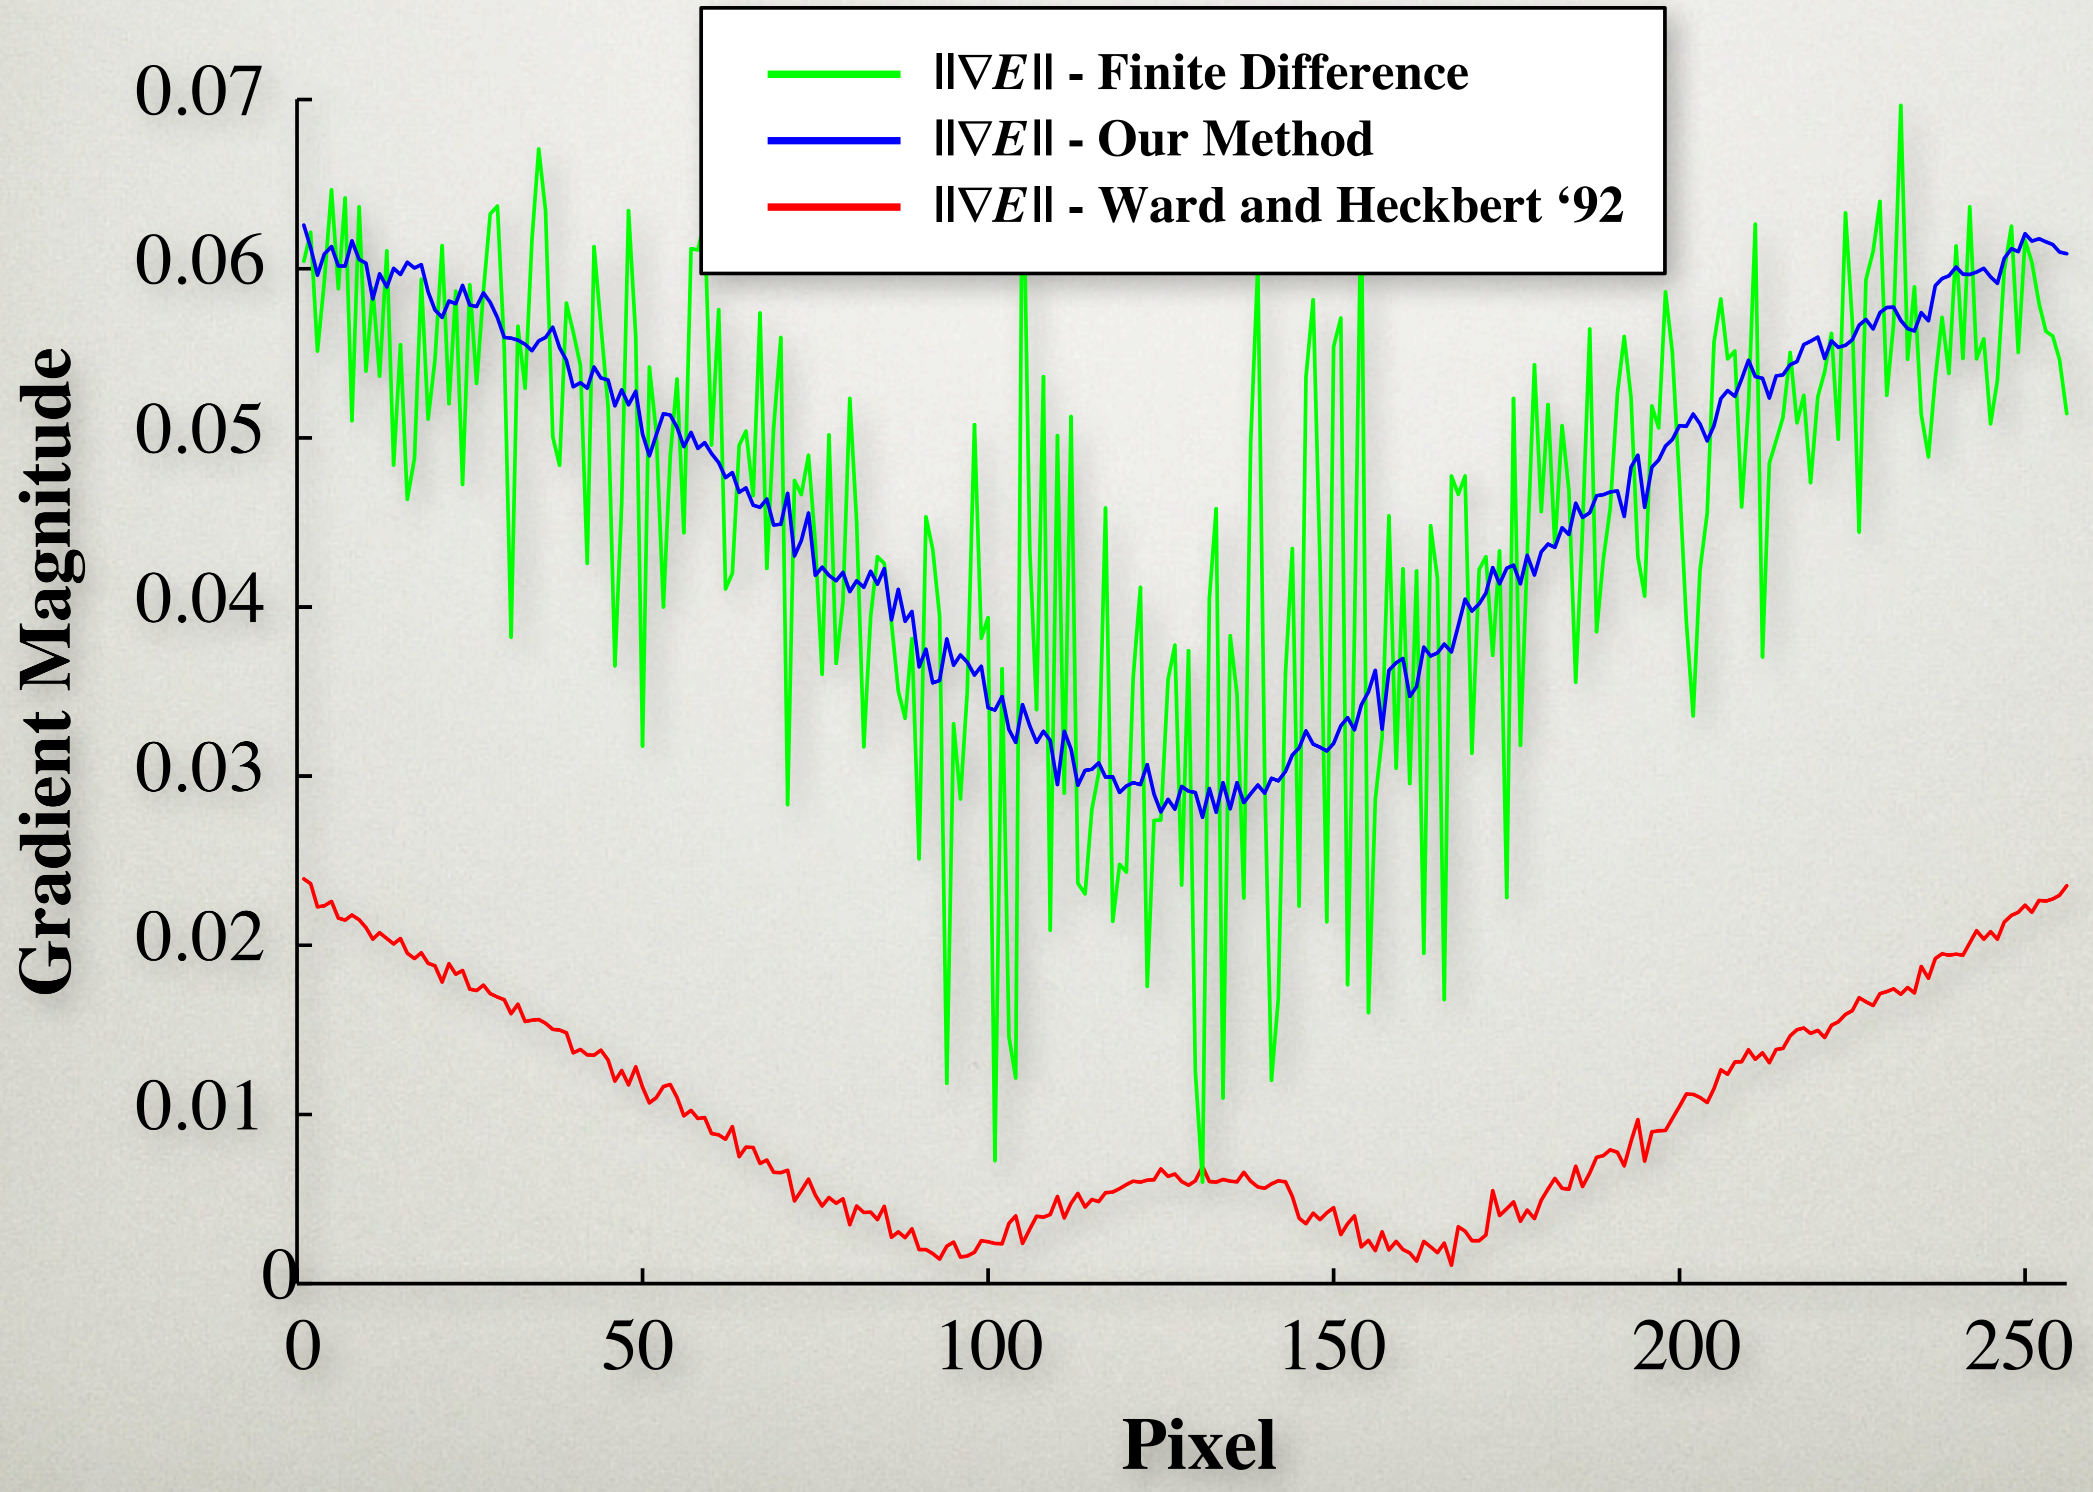

Thursday, 6 September 12

* we see that our method (shown in blue) matches the ground truth, whereas Ward and Heckbert gradients (shown in red) significantly differ from this

Extrapolated Irradiance

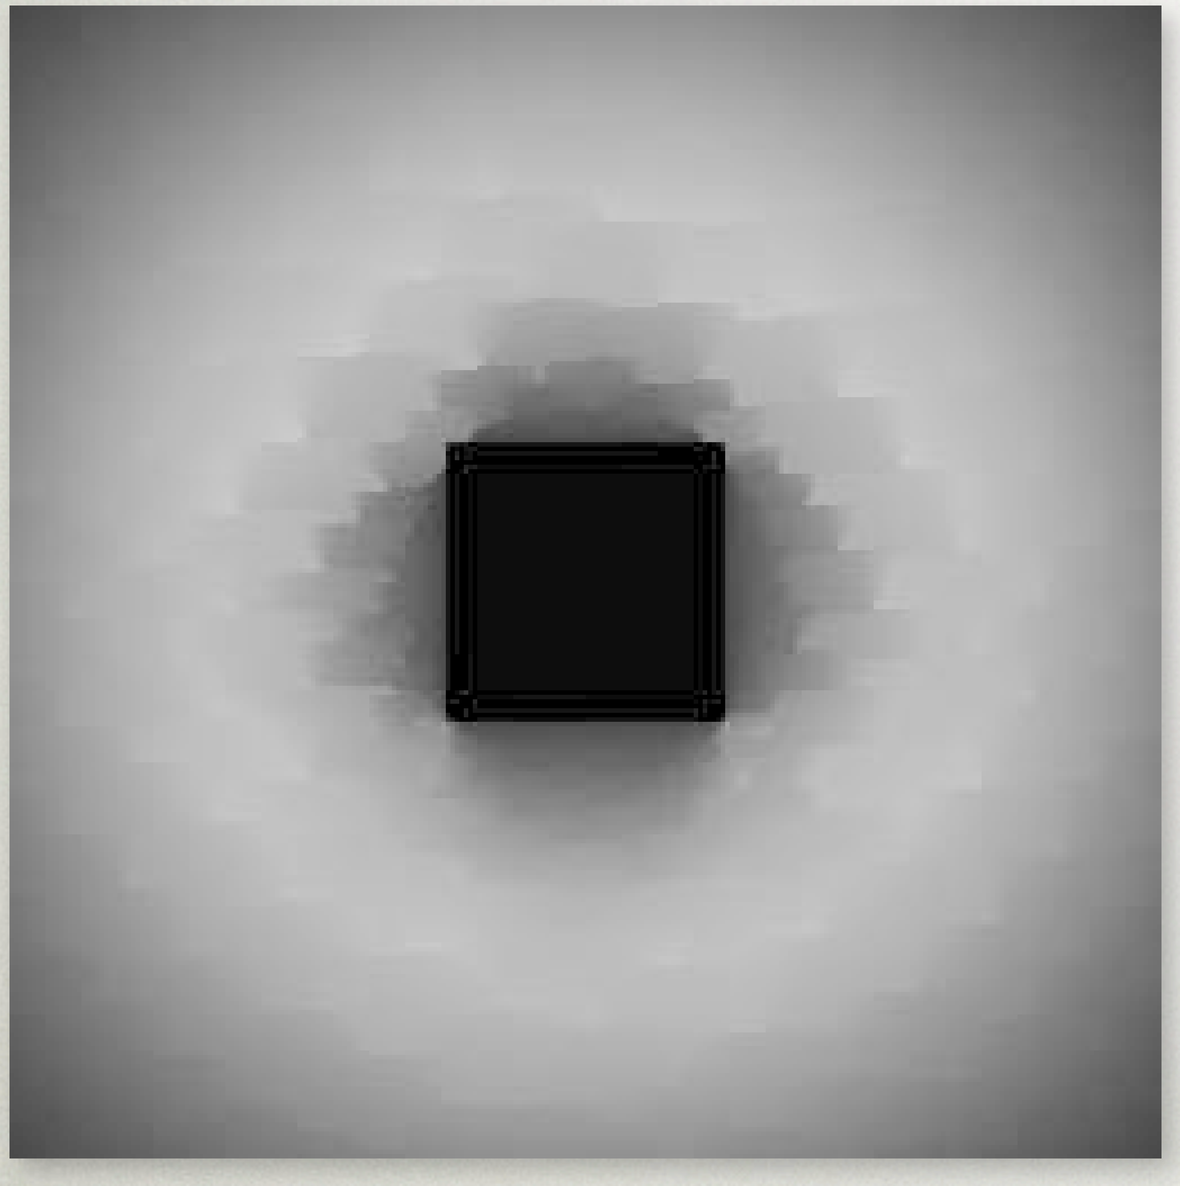

Our Method

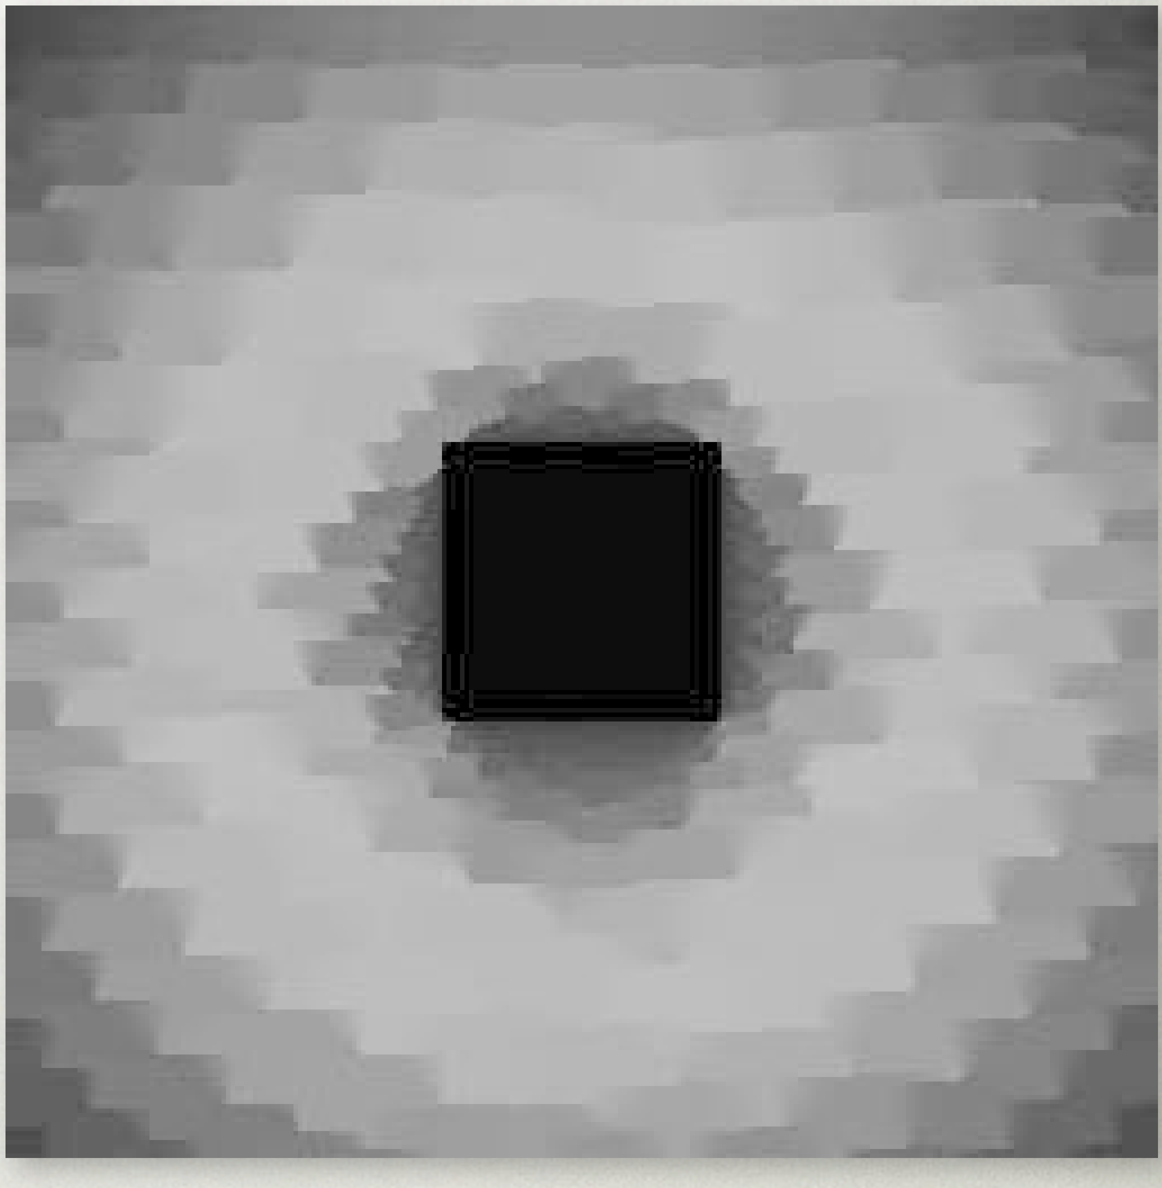

Ward and Heckbert

Same cache points

Gradient Comparison

- In a scene with no walls, Ward & Heckbert would estimate 0 gradients!

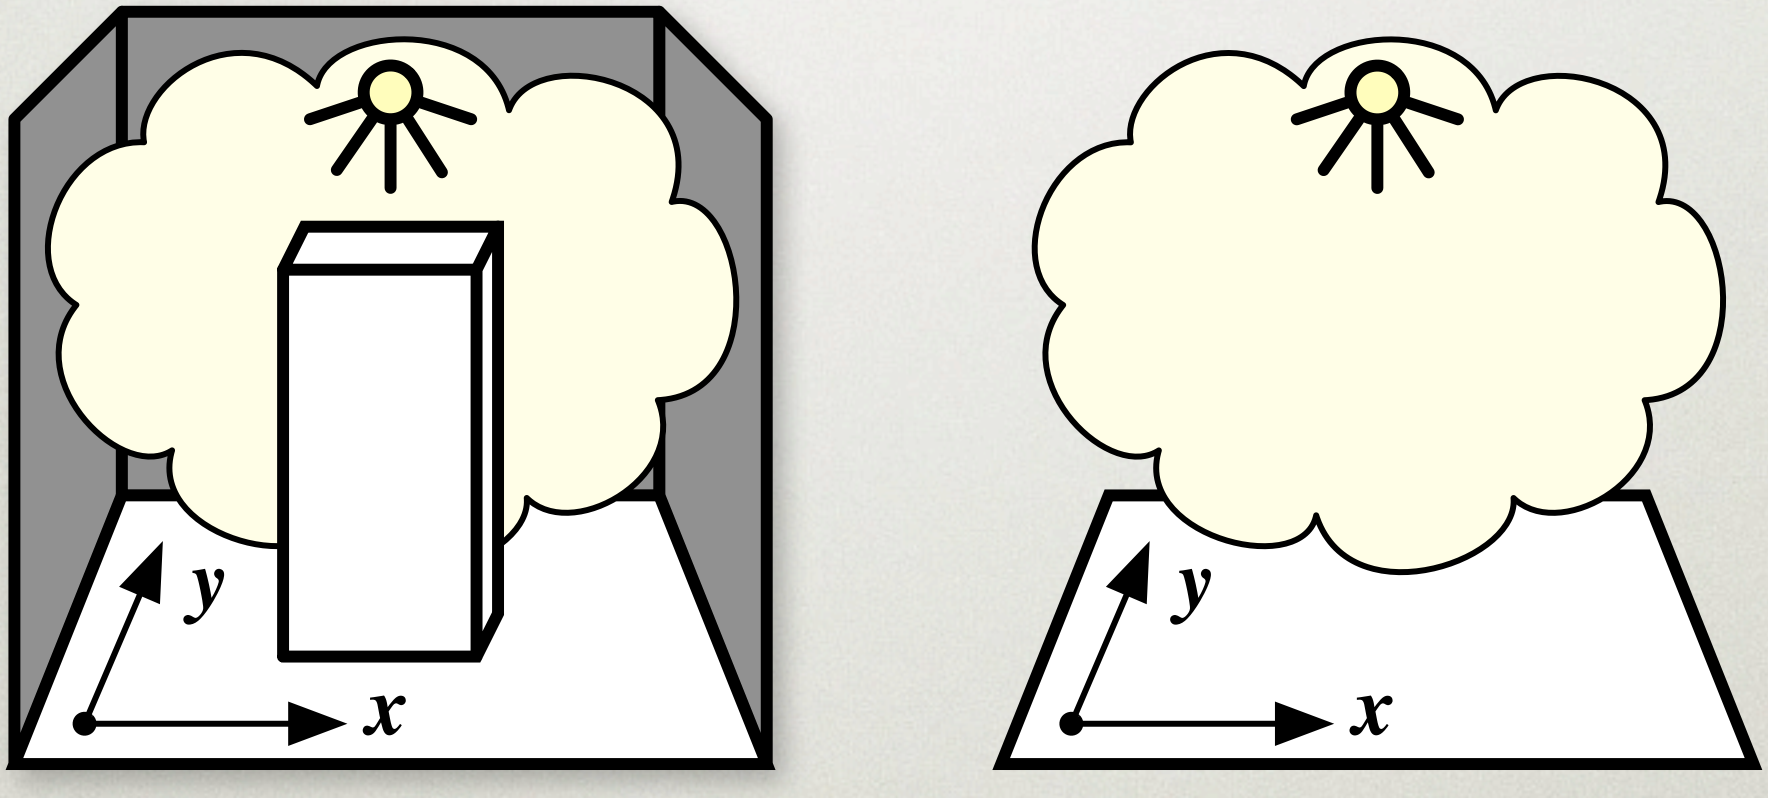

Thursday, 6 September 12

* This test scene is actually constructed to give the original Ward & Heckbert gradient a helping hand.
* If we removed the box and the walls then Ward & Heckbert's formulation would incorrectly estimate a 0 gradient everywhere, which would be of no benefit for interpolation.

Convergence

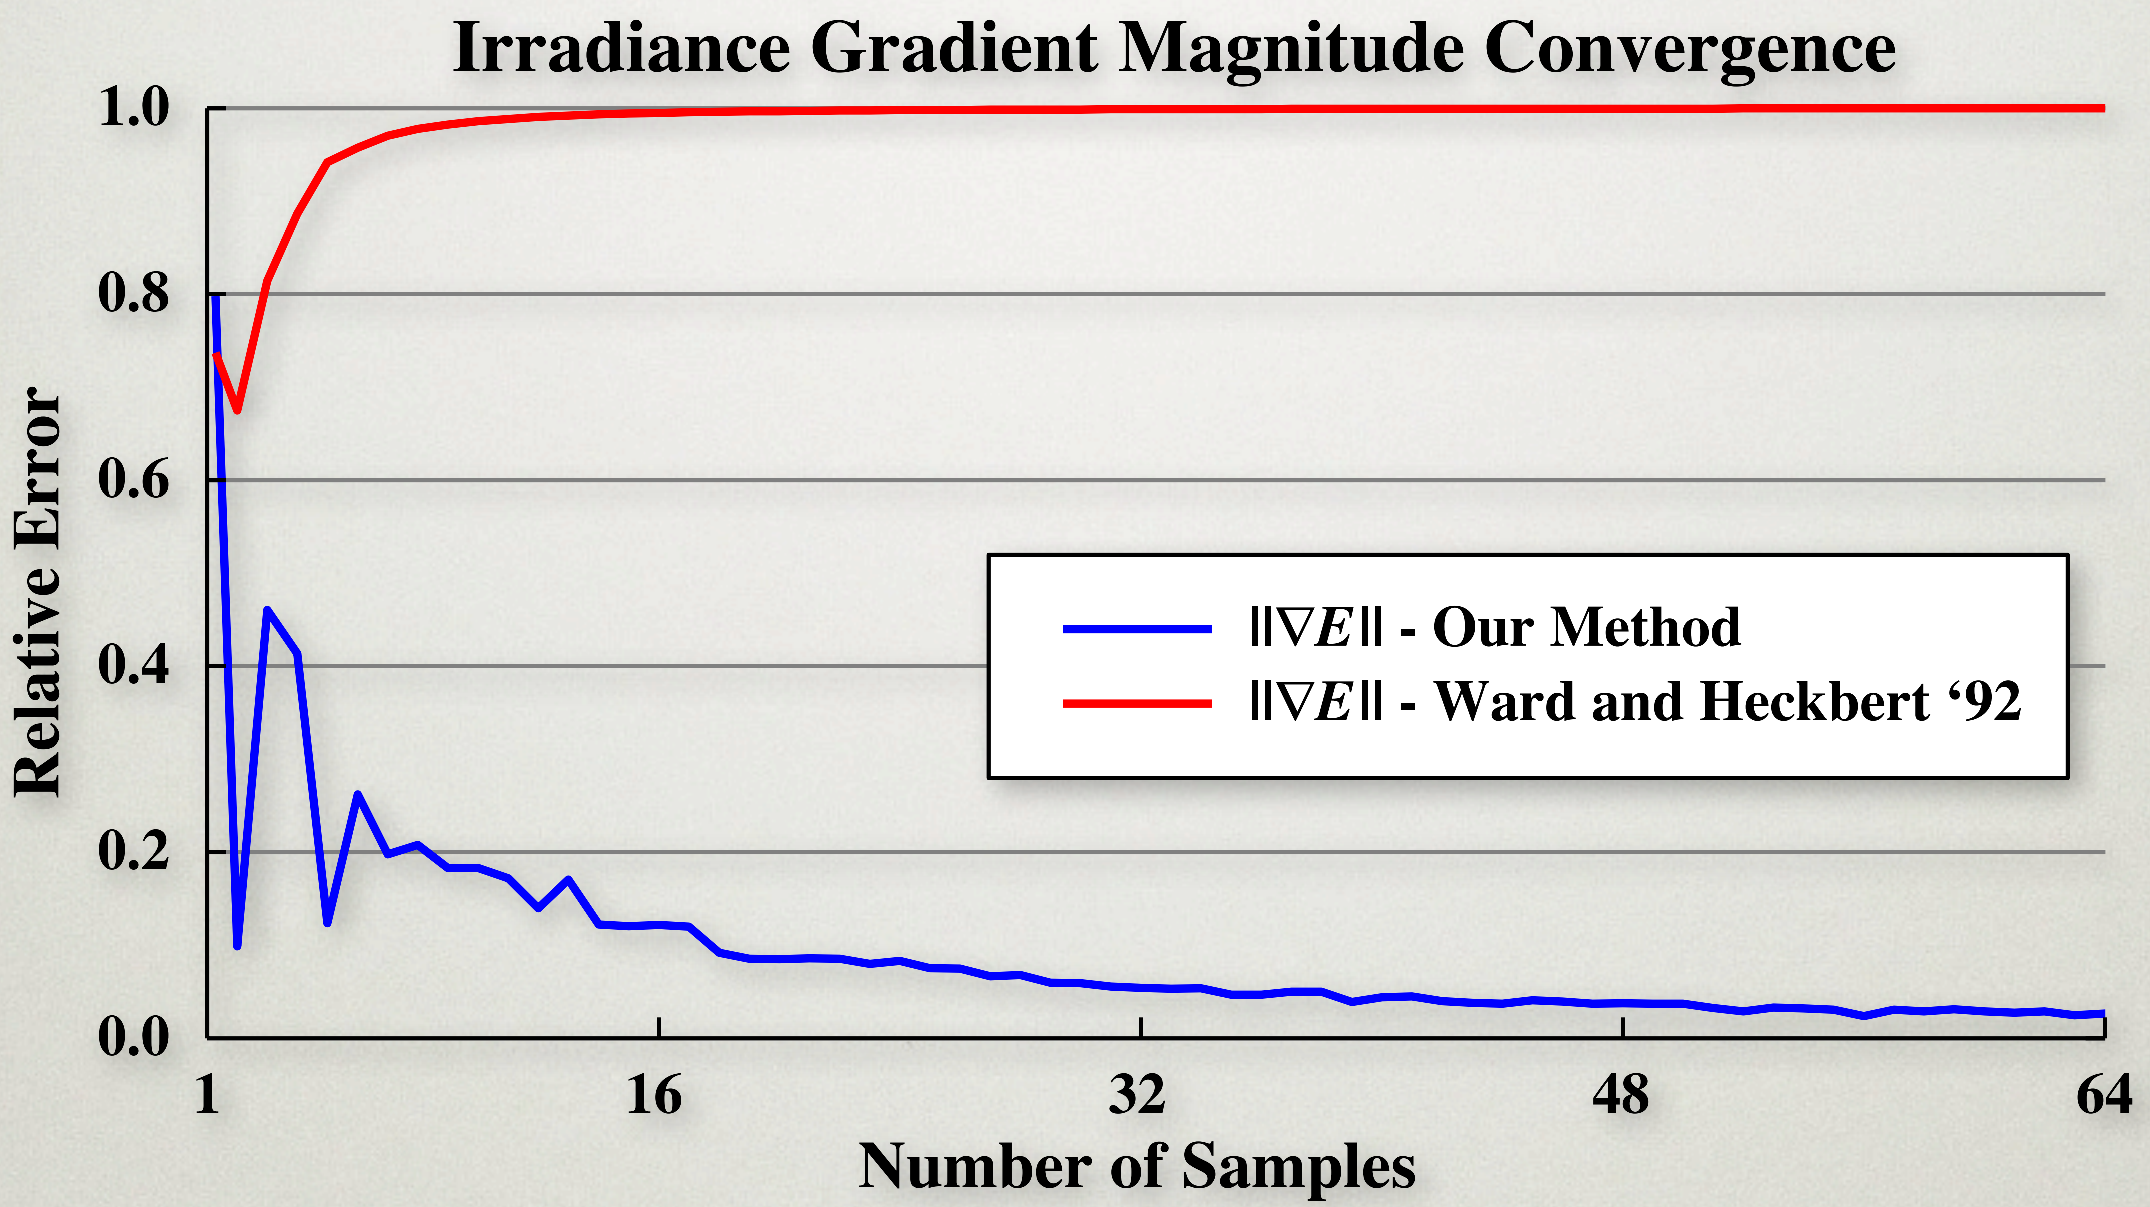

Results

- Rendered at 1K horizontal resolution
- On an Intel Core 2 Duo 2.4 GHz PC

Smoky Cornell Box

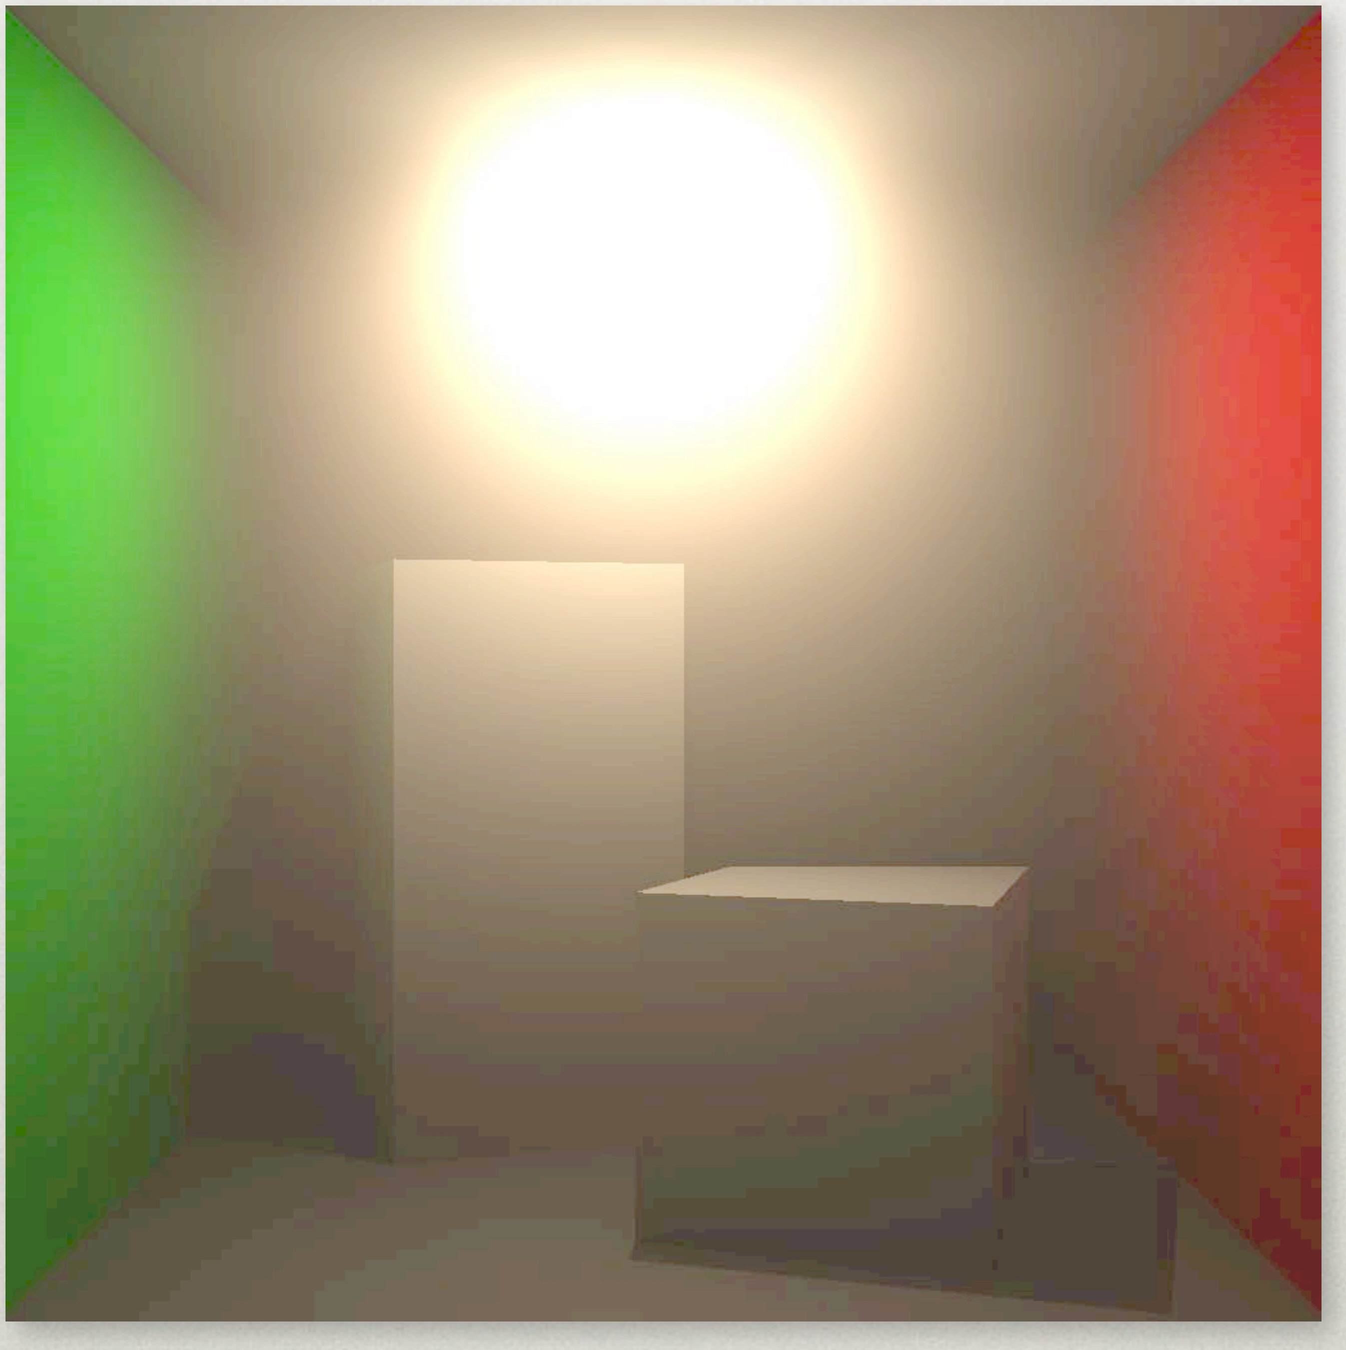

(8:14)

Smoky Cornell Box

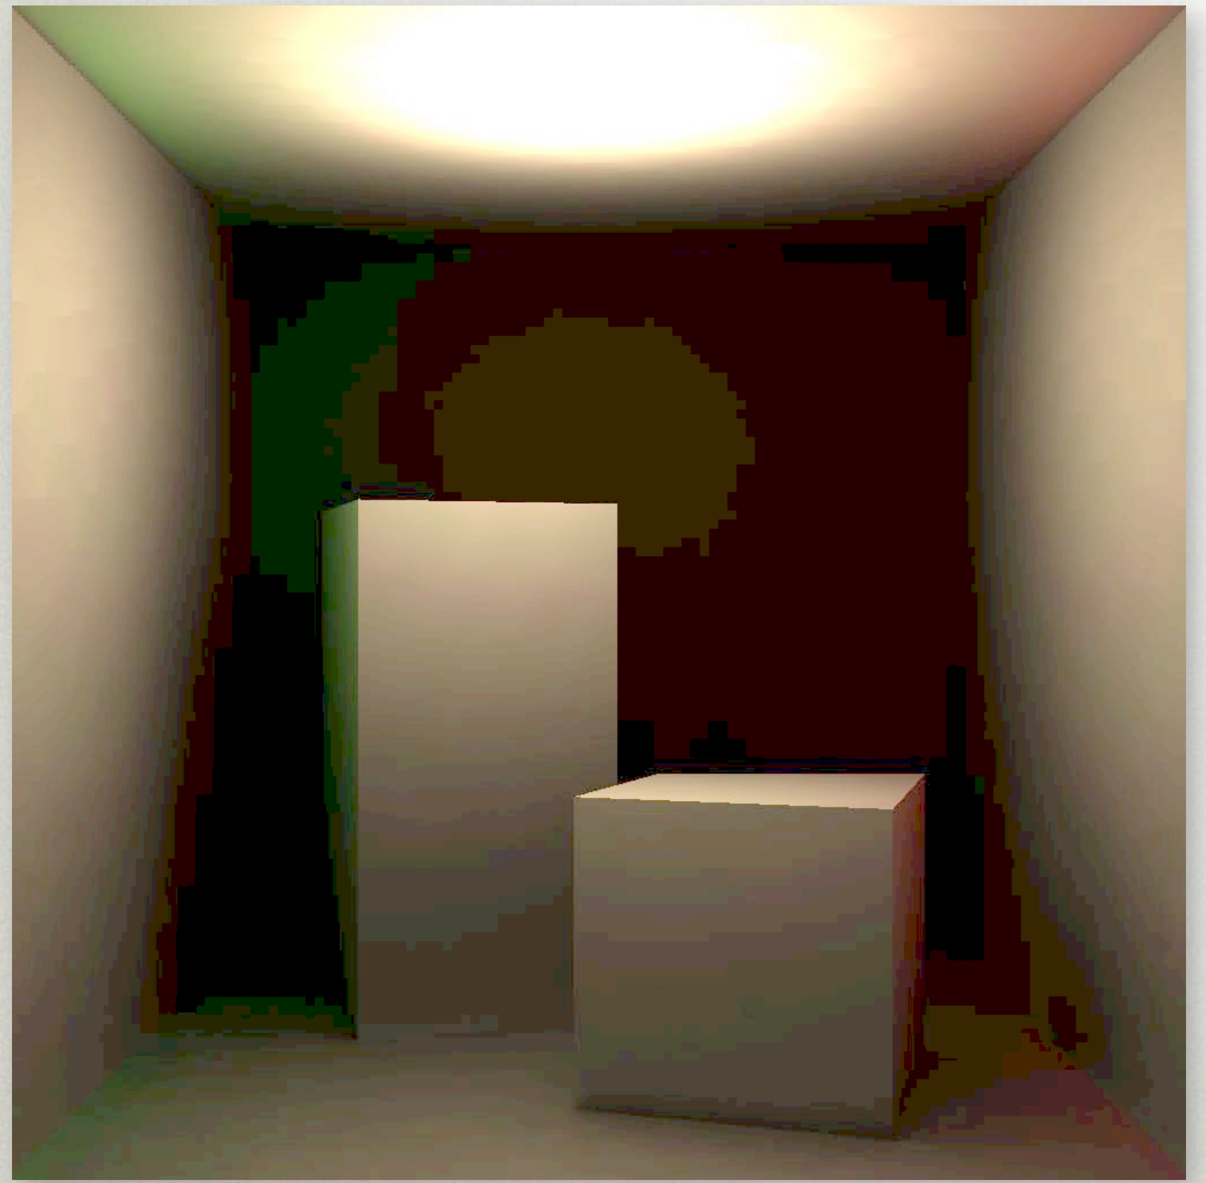

Our Method
(8:14)

Ward and Heckbert
(8:10)

Beam through Window

Our Method
(3:25)

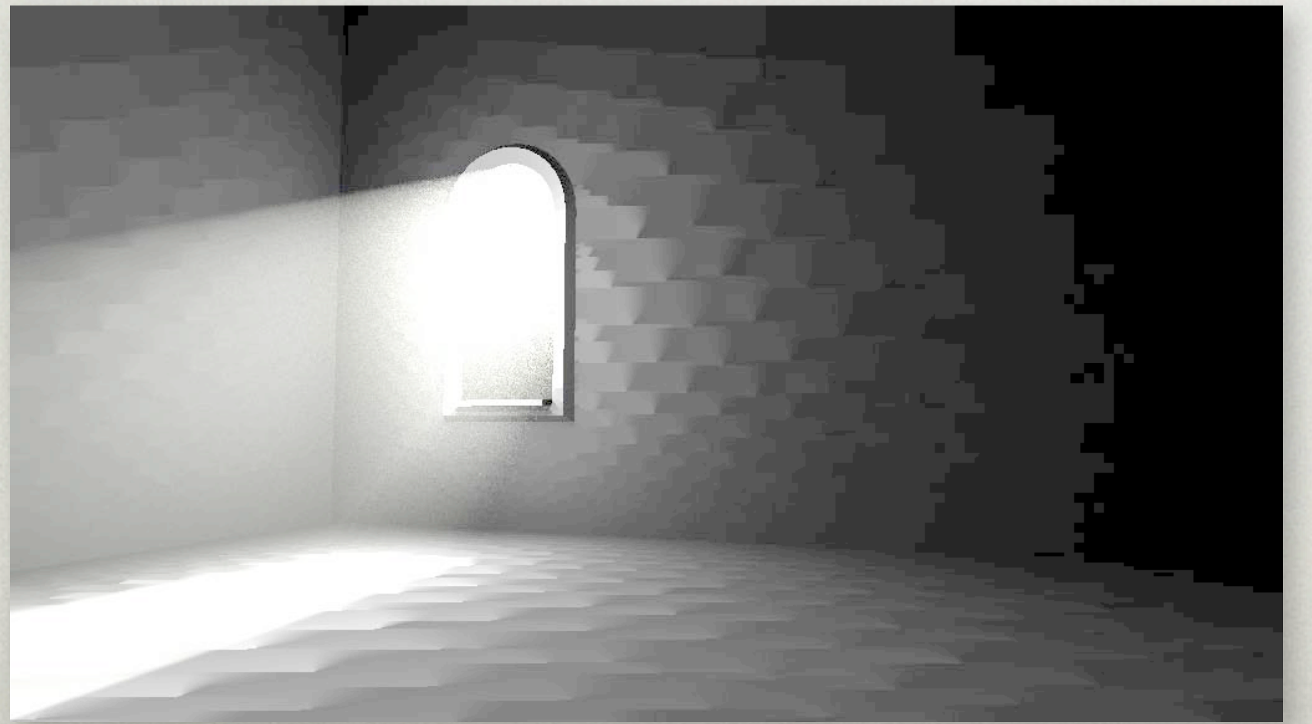

Ward and
Heckbert
(3:17)

* still based on stochastic sampling, so there may still be errors (e.g., on the floor), not as bad as previous methods

Disco Ball

(10:33)

Thursday, 6 September 12

* All illumination on the ground plane has first scattered in the medium.

Disco Ball

Our Method
(10:33)

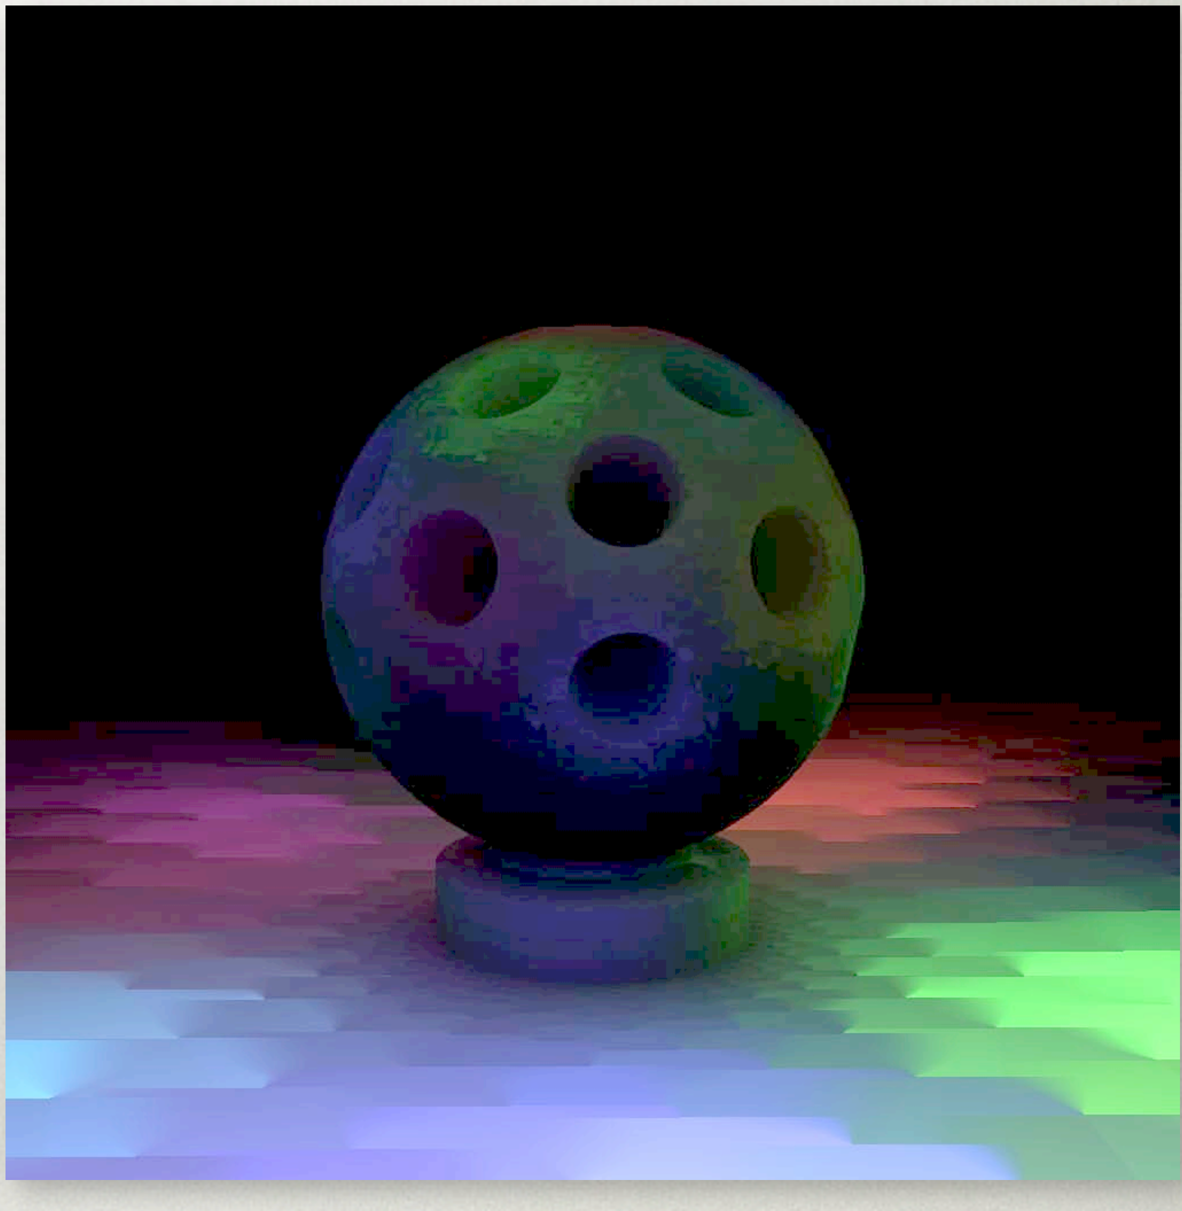

Ward and Heckbert
(10:30)

Future Work

- Error metric

Future Work

- Error metric
- *Radiance* gradients

Future Work

- Error metric
- *Radiance* gradients
- Radiance gradients *in* participating media

Conclusion

- Accurate irradiance gradients for scenes with media and occlusions
- Can be applied to the irradiance caching algorithm for higher quality interpolation

Thank You