





Precomputed Radiance Transfer: Theory and Practice

Peter-Pike Sloan

Microsoft

Jaakko Lehtinen

Helsinki Univ. of Techn.
&
Remedy Entertainment

Jan Kautz

MIT



Light Representation

Jan Kautz

MIT



Light Representation

- Many ways of creating incident light
 - Synthetic environment map
 - HDR light probes (natural lighting)
 - Directly from analytical solutions
 - e.g. solution for a disk light source, angle τ

There are many ways to create appropriate incident lighting.



Overview

- Light Representation
 - Analytical Solutions (SH)
 - Directional, disc, and spherical light sources
 - Environment Maps (SH/Haar)
 - Light probes
 - Rotation (SH/Haar)
 - Sample on-the-fly
 - Interpolate from precomputed incident lighting (SH)
 - Use gradient for better interpolation (SH)

In this part, we will talk about several ways to create such incident lighting.

Analytical Solutions (SH)



- Directional light source
 - Coming from a single direction Ψ
 - Incident illumination is basically a Dirac peak
- Integration against basis functions boils down to evaluating basis functions in that direction: $l_i = y_i(\Psi)$
- Good for e.g. sun light


The easiest way to create a light source is to integrate the basis functions against a directional light source (basically a Dirac).

Analytical Solutions (SH)



- Disc sources (center around z):

$$d_{\tau}(\theta, \varphi) = \begin{cases} 1 & (\tau - \theta) > 0 \\ 0 & \text{otherwise} \end{cases}$$



- Integrate symbolically in Maple:

$$l_i = \int_{\varphi=0}^{2\pi} \int_{\theta=0}^{\tau} y_i(\theta, \varphi) \sin \theta d\theta d\varphi$$

A disc-shaped area source can be analytically projected into SH.

It is preferred to center the disc around the z-axis. In that case, only very few coefficients are non-zero. Here are the non-zero coefficients for the first 5 bands:

$$l(l=0, m=0) = -1.772453851 * \cos(\tau) + 1.772453851$$

$$l(l=1, m=0) = 1.534990062 * \sin(\tau)^2$$

$$l(l=2, m=0) = -1.981663648 * \cos(\tau)^3 + 1.981663648 * \cos(\tau)$$


$$l(l=3, m=0) = -2.930920064 * \cos(\tau)^4 + 3.517104076 * \cos(\tau)^2 - .5861840127$$

$$l(l=4, m=0) = -4.652691359 * \cos(\tau)^5 + 6.646701941 * \cos(\tau)^3 - 1.994010582 * \cos(\tau)$$

The light source can be rotated around using the SH rotation matrices (see later slides).

Analytical Solutions (SH)



- Disc sources: 
 - Due to SH construction any function that is rotationally symmetric around z only has non-zero coefficients for modes $m=0$ (Zonal Harmonics)
 - Efficient representation and also rotation
 - Use SH rotation matrix to move light

It's interesting to note, that the lighting vector will only have very few coefficients in that case (only for modes $m=0$), i.e. for the first 5 bands, only 5 out of the 25 coefficients will be non-zero. This also makes the rotation (see later slides) more efficient.

Analytical Solutions (SH)

- Disc sources 


$$l(1=0,m=0) = -1.772453851 \cdot \cos(\tau) + 1.772453851$$
$$l(1=1,m=0) = 1.534990062 \cdot \sin(\tau)^2$$
$$l(1=2,m=0) = -1.981663648 \cdot \cos(\tau)^3 + 1.981663648 \cdot \cos(\tau)$$
$$l(1=3,m=0) = -2.930920064 \cdot \cos(\tau)^4 + 3.517104076 \cdot \cos(\tau)^2 - .5861840127$$

These are the formulas for a disc source.

Analytical Solutions (SH)



- Spherical light sources
 - Spherical light sources appear as disc sources



- Distance of spherical light to center d and sphere radius r determines size of cap: $\sin \tau = r/d$

A special case are spherical light sources. They correspond to disc-like sources, but their size changes with the distance to the sample point, where the light is sampled. This sampling can be actually performed on-the-fly.

Analytical Solutions (SH)



- Spherical light sources
 - Lighting coeff's can be computed per-pixel/vertex
 - Store coefficients in texture map, index with τ
 - Actually compute on-the-fly
 - Rotate "disc" from pole to illumination direction
 - I.e., apply rotation matrix, see later slides

Since the formulas are very simple, the lighting coefficients can actually be computed on the fly.

Environment Maps



- Take a light probe with natural lighting



and represent it in the lighting basis function

(High-dynamic range) environment maps are a simple way of "creating" incident lighting.

Environment Maps

- Spherical Harmonics

- Integrate basis functions against fixed HDR map

$$l_i = \int L_{env}(\vec{s}) y_i(\vec{s}) d\vec{s}$$

$$l_i = \int \text{img} \cdot \text{basis}$$

- 1) Monte-Carlo integration
- 2) Precompute maps in same space (e.g. cube map) that contain the basis functions \Rightarrow big dot-product for each coeff.

This is a standard scenario. A HDR environment map is given and projected into basis functions.

If the environment map is given as a cube map, it is necessary to know the solid angle of a texel in the cube map. It is: $4 / (\sqrt{X^2 + Y^2 + Z^2})^3$, where $[X Y Z]$ is the vector to the texel (not normalized).

Environment Maps

- Haar Wavelets
 - Represent environment maps as cube maps
 - Do standard wavelet transform on each face
 - Compress/quantize data (select most significant coefficients)



Reference

4096 coeffs.

100 coeffs.

Rotating Lighting / Env.Maps



- Rotation of lighting – Spherical Harmonics
 - There is a simple linear matrix that will rotate SH coefficients directly:
$$\mathbf{I}' = \mathbf{R}_{SH} \mathbf{I}$$
 - The matrix is block diagonal
 - No spill across bands
 - For the first few SH bands, it can be computed explicitly
 - General case: recurrence formula

$$R_{SH} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & X & X & X & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & X & X & X & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & X & X & X & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & X & X & X & X & X & \dots \\ 0 & 0 & 0 & 0 & X & X & X & X & X & \dots \\ 0 & 0 & 0 & 0 & X & X & X & X & X & \dots \\ 0 & 0 & 0 & 0 & X & X & X & X & X & \dots \\ 0 & 0 & 0 & 0 & X & X & X & X & X & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

Spherical Harmonics are a steerable basis. It means that no information is lost under rotation.

There are various ways of creating this rotation matrix. There are recurrence formulas, but in case only a few bands are needed, the matrix can be created explicitly.



Rotating Lighting / Env.Maps

- Special case: disc sources
 - Remember: only $m=0$ coeff's are non-zero

$$l_i^R = R_{SH} \cdot l_i = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & X & a & X & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & X & b & X & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & X & c & X & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & X & X & d & X & X & \dots \\ 0 & 0 & 0 & 0 & X & X & e & X & X & \dots \\ 0 & 0 & 0 & 0 & X & X & f & X & X & \dots \\ 0 & 0 & 0 & 0 & X & X & g & X & X & \dots \\ 0 & 0 & 0 & 0 & X & X & h & X & X & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \cdot \begin{bmatrix} C \\ 0 \\ C \\ 0 \\ 0 \\ 0 \\ C \\ 0 \\ 0 \\ \vdots \end{bmatrix} = \begin{bmatrix} C \\ a \cdot C \\ b \cdot C \\ c \cdot C \\ d \cdot C \\ e \cdot C \\ f \cdot C \\ g \cdot C \\ h \cdot C \\ \vdots \end{bmatrix} \Rightarrow \text{becomes dot product}$$

see [Sloan05]

In case of z-rotationally symmetric sources (represented solely with Zonal Harmonics, see [Sloan05]), the rotation matrix becomes a dot-product.

Rotating Lighting / Env.Maps



SIGGRAPH2005

- Rotation of HDR map – Haar Wavelets
 - No convenient rotation matrix
 - Best solution:
 - Rotate actual environment map
 - Flatten rotated environment into cube map
 - Redo wavelet transform of each cube face
 - Potential problems: flickering
 - Avoid: blur the environment map slightly

In case of Haar Wavelets, there is no convenient rotation matrix. All one can do is to rotate in image space and then reproject.



Varying Lighting

- Up to now: lighting didn't really change
- What if character moves within scene?
 - Expect lighting on character to change
- Two solutions:
 - Sample lighting on-the-fly
 - Precompute for static scenes

Environment Maps – Sample Incident Lighting



- Sample lighting:
 - Render scene from center of object p
 - 6 times: for each cube face
 - Compute lighting coefficients (SH/Haar):

$$l_i^p = \int \text{[Crosshair Image]} \cdot \text{[Color Sphere]}$$

- No need to rotate lighting then

More local effects (SH)

- Incident light will vary within a static scene
- Borrow "Irradiance Volume" technique [Greger98]
- **Precompute** incident radiance on a grid



Courtesy of
Rafal Mantiuk

Sofar, we have assumed that lighting is infinitely far away, but that doesn't go well with interior scenes.

In that case, we can sample the incident radiance at grid points throughout the scene (assuming the scene doesn't change) and store that in SH.



More local effects (SH)

- Incident light will vary within a static scene
- Borrow "Irradiance Volume" technique [Greger98]
- **Precompute** incident radiance on a grid
 - I.e., for each sample point p_k precompute a light vector \mathbf{l}_{p_k} (assumes static lighting)
 - For new points q , linearly interpolate light from nearby sample points p_k [Mantiuk02]
- E.g., this technique is used in Max Payne 2

At run-time, one just queries the closest neighbors and interpolates the lighting coefficient vectors.

Radiance Volumes

- Note how bunny is illuminated differently




Courtesy of
Rafal Mantiuk



This gives a nice sense of an object being "in the scene". Even just the first 2 SH bands (4 coefficients) gives a great sense.

Gradients

- Can we do better than simple interpolation?
- How does incident light change when we move away from sample point?



Approximate
with gradient!

[Annen04]

Instead of just linear interpolation, one can use a gradient in addition.

Gradients

- Lighting is projected into SH as follows:

$$l_i^p = \int_{\Omega} L_s(p \leftarrow \vec{s}) y_i(\vec{s}) d\vec{s}$$

- We want to compute the gradient when we move \mathbf{d} away from p
- And reuse at run-time:

$$l_i^{p+d} = l_i^p + (\nabla l_i^p \cdot \mathbf{d})$$

Analytical Gradient

- Rewrite integral as integral over surfaces (because of implicit dependency on p)
- Take gradient: (analytical integrand)


$$\nabla l_i^p = - \int_S \nabla \left(y_i \left(\frac{\vec{\sigma}}{\|\vec{\sigma}\|} \right) \frac{\vec{n}(\underline{x}) \cdot \vec{\sigma}}{\|\vec{\sigma}\|^3} \right) L_S(p \leftarrow \underline{x}) dA$$

(Code is available from the authors)


$$\begin{aligned} \vec{n} &= \text{normal} \\ \vec{\sigma} &= \underline{x} - p \end{aligned}$$

The actual derivation is slightly involved and omitted here. It is available in the paper [Annen04]. Code is available from the authors.

Analytical Gradient – Putting it all together

- Compute incident lighting 
 - Store distance and normal as well
- Compute coefficients $l_i^p = \int \text{SH} * \text{globe}$
 - Per-pixel integration with SH basis
- Compute gradient $\nabla l_i^p = \int ((\nabla \text{SH}) * \text{globe} + \text{SH} * (\nabla \text{globe}))$
 - Per-pixel integration of simple analytical formula
 - Requires distance & normal
- For rendering:
 - Extrapolate to coefficients l_i^{p+d}

Analytical Gradient – Results



- Single Sample (extrapolate to every vertex)
 - no gradient
 - with gradient
 - ground truth
- Multiple Sample (extra/interpolate to every vertex)
 - sample locations
 - eight samples
 - eight gradients

The figure displays a 2x3 grid of renderings of a red apple. The top row shows the 'Single Sample' method: 'no gradient' shows a dark, flat apple; 'with gradient' shows a smooth gradient from dark purple to bright red; 'ground truth' shows the most realistic lighting. The bottom row shows the 'Multiple Sample' method: 'sample locations' shows a blue wireframe mesh with eight yellow dots; 'eight samples' shows a smooth gradient similar to the 'with gradient' result; 'eight gradients' shows a smooth gradient similar to the 'ground truth' result.

The gradient helps to give a much more "local feel" of the light sources.




Numerical Gradient

- Gradient can also be computed numerically
 - Use e.g. central differencing between lighting samples
 - No need to do it in 3D (for each coordinate axis)
 - Along a 1D path
 - Variation on a 2D plane

Numerical computation of the gradient is viable as well.

Numerical Gradient

- Gradient can also be computed numerically



From ATI – Ruby Demo

Demos

Given by Natalya Tatarchuk from ATI

Conclusion

- Lighting can be more than a static cube map
 - Analytical models
 - Sample on the fly
 - Precompute (interpolate)
 - Or even extrapolate using gradients

