Inverse rendering and relighting applications

Nick Michiels

Jeroen Put

Tom Haber

Relighting

- $B(x, \omega_o) = \int_\Omega \boxed{V(x, \omega_i)} \boxed{\rho(x, \omega_i, \omega_o)(\omega_i \cdot n(x))} \boxed{L(x, \omega_i)} \, d\omega_i$

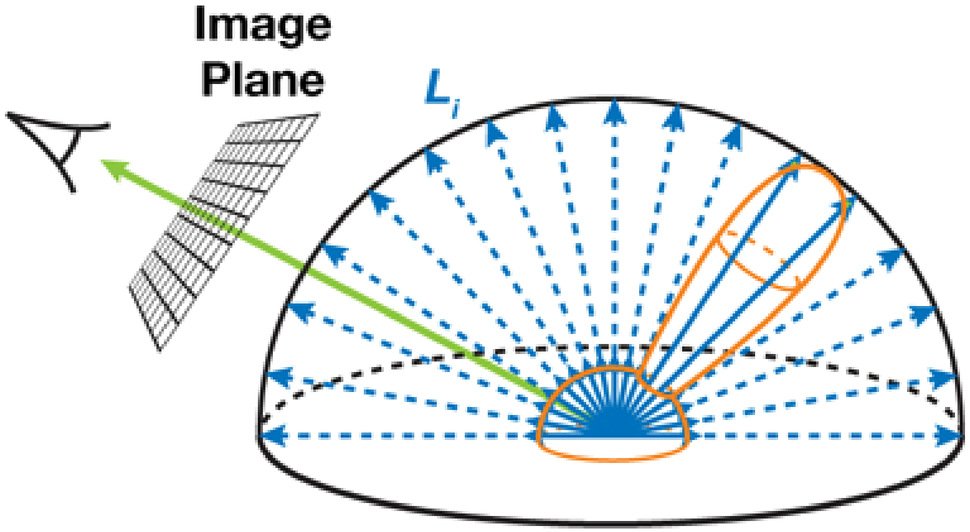

Source: http://http.developer.nvidia.com/GPUGems3/gpugems3_ch20.html

universiteit hasselt | EDM | iMinds

Relighting

- $B(x, \omega_o) = \int_{\Omega} \boxed{V(x, \omega_i)} \boxed{\rho(x, \omega_i, \omega_o)(\omega_i \cdot n(x))} \boxed{L(x, \omega_i)} \, d\omega_i$

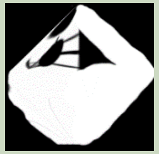

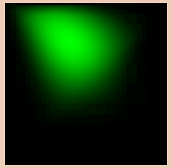

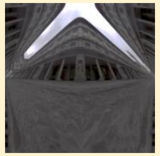

universiteit hasselt | EDM | iMinds

From 2D to the spherical domain

- Octahedron parameterisation

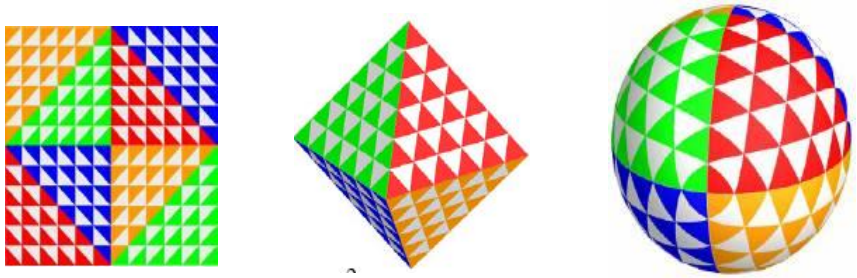

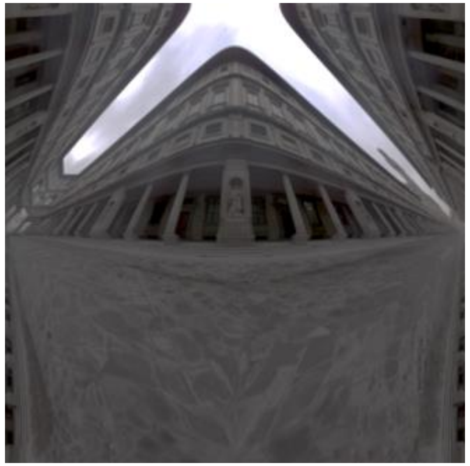

Relighting

- $B(x, \omega_o) = \int_\Omega \boxed{V(x, \omega_i)} \boxed{\rho(x, \omega_i, \omega_o)(\omega_i \cdot n(x))} \boxed{L(x, \omega_i)} \, d\omega_i$

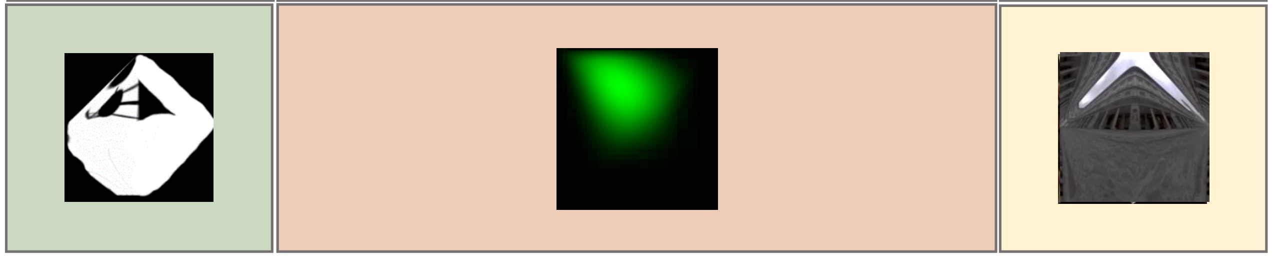

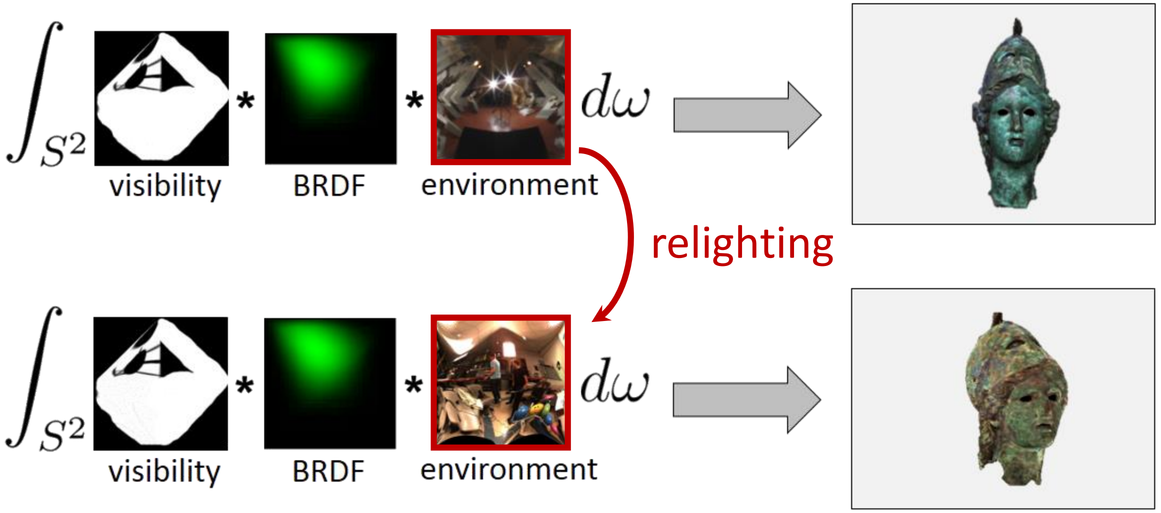

Inverse Rendering

- Computer vision <-> computer graphics
  - Reflected light = convolution of lighting and BRDF
  - Inverse rendering = deconvolution

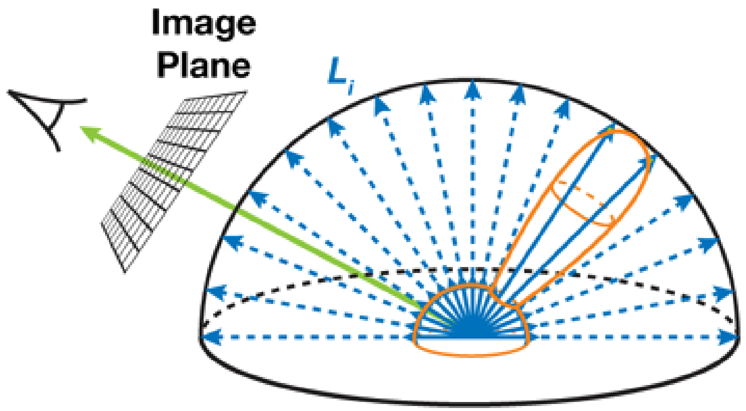

Source: http://http.developer.nvidia.com/GPUGems3/gpugems3_ch20.html

$$B(x,\omega_o) = \int_{\Omega} V(x,\omega_i)\, \rho(x,\omega_i,\omega_o)(\omega_i \cdot n(x))L(x,\omega_i)\, d\omega_i$$

(cube, size, x_0, y_0, z_0, θ_{XY}, θ_{XZ}, θ_{YZ}, ...)

(sphere, radius, x_1, y_1, z_1, ...)

Computer Graphics ↓ ↑ Computer Vision

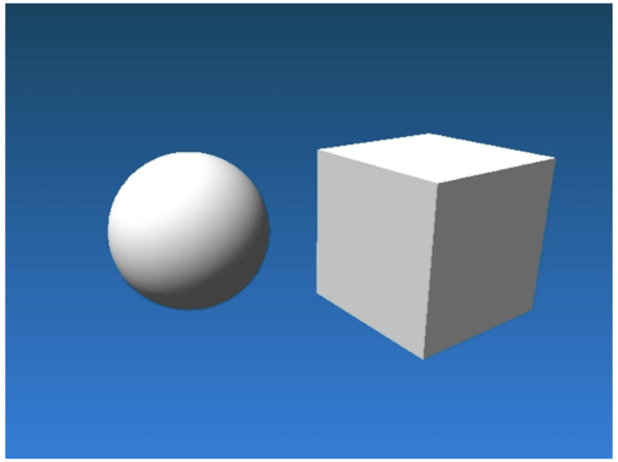

universiteit hasselt | EDM | iMinds

Goal

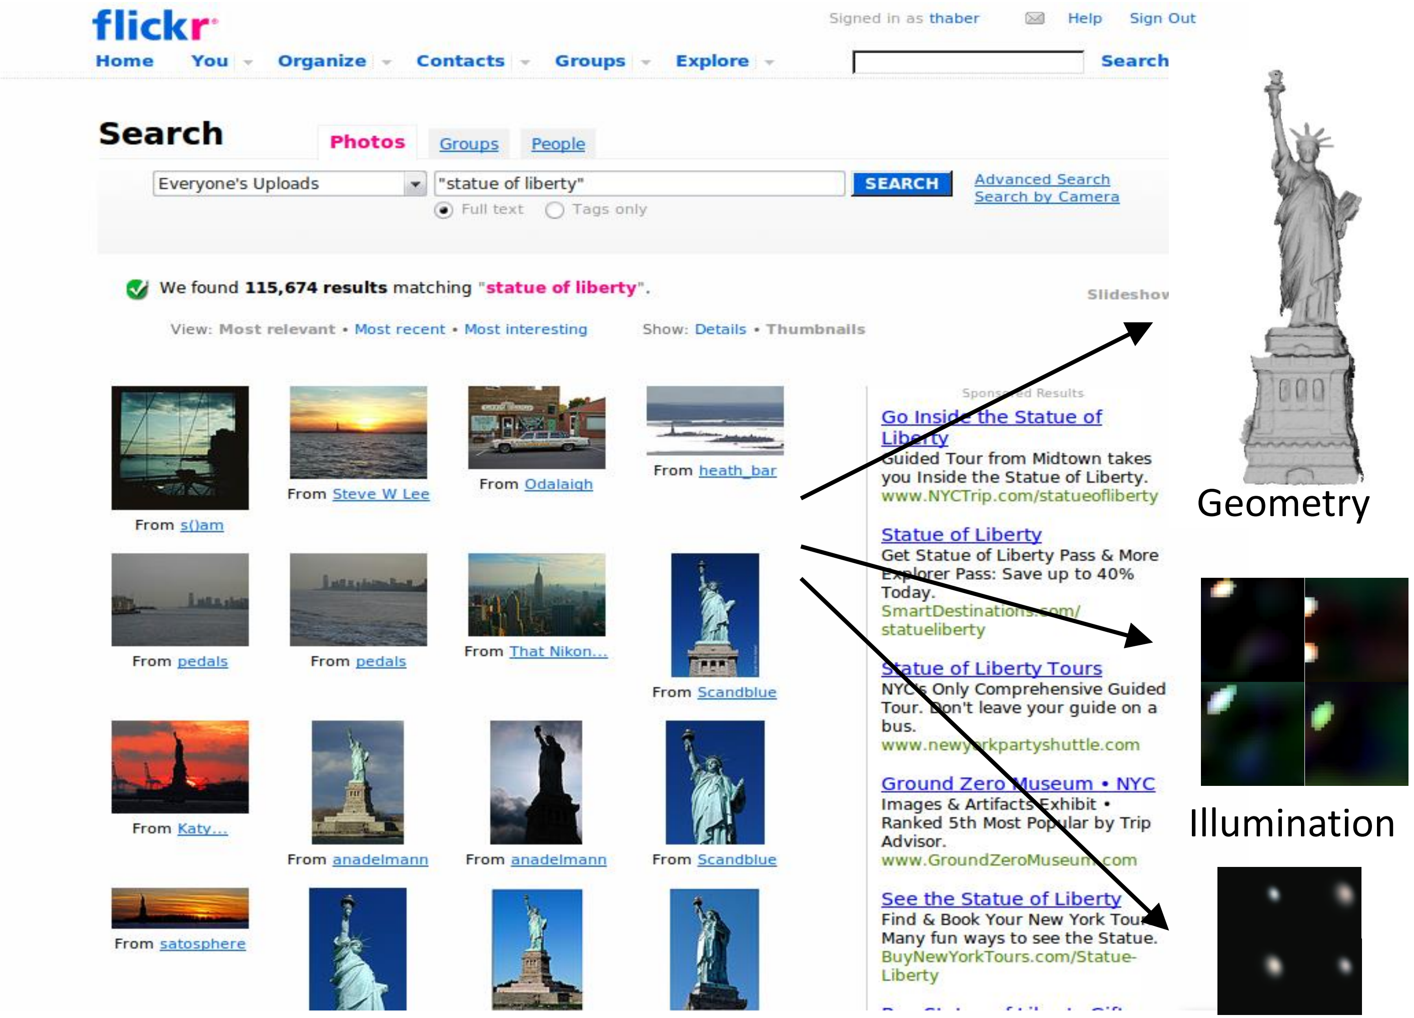
flickr
Signed in as thaber
Help
Sign Out
Home
You
Organize
Contacts
Groups
Explore
Search
Search
Photos
Groups
People
Everyone's Uploads
"statue of liberty"
SEARCH
Advanced Search
Search by Camera
Full text
Tags only
We found 115,674 results matching "statue of liberty".
View: Most relevant • Most recent • Most interesting
Show: Details • Thumbnails
From s()am
From Steve W Lee
From Odalaigh
From heath_bar
From pedals
From pedals
From That Nikon...
From Scandblue
From Katy...
From anadelmann
From anadelmann
From Scandblue
From satosphere
Go Inside the Statue of Liberty
Guided Tour from Midtown takes you Inside the Statue of Liberty.
www.NYCTrip.com/statueofliberty
Statue of Liberty
Get Statue of Liberty Pass & More
Explorer Pass: Save up to 40% Today.
SmartDestinations.com/statueliberty
Statue of Liberty Tours
NYC's Only Comprehensive Guided Tour. Don't leave your guide on a bus.
www.newyorkpartyshuttle.com
Ground Zero Museum • NYC
Images & Artifacts Exhibit • Ranked 5th Most Popular by Trip Advisor.
www.GroundZeroMuseum.com
See the Statue of Liberty
Find & Book Your New York Tour. Many fun ways to see the Statue.
BuyNewYorkTours.com/Statue-Liberty
Geometry
Illumination

BRDFs

Results





Results

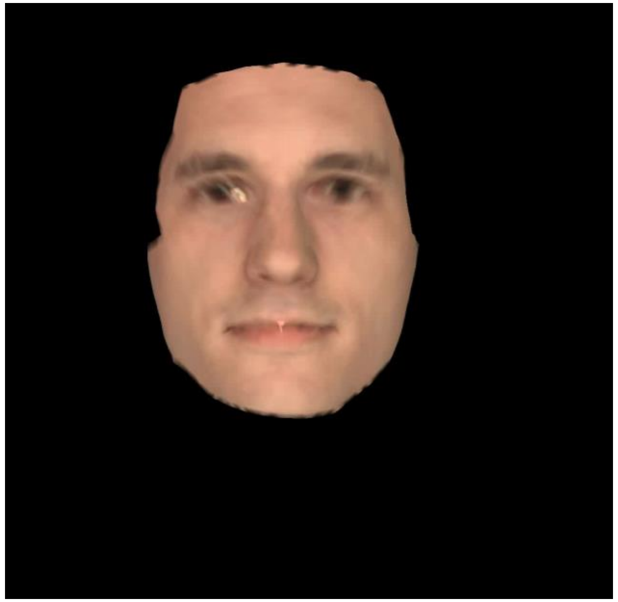

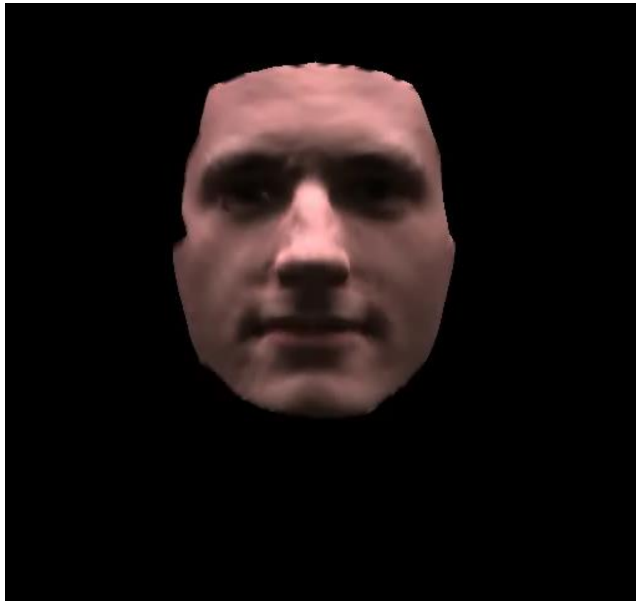

Why is it difficult?

- Image <-> array of pixels

38	36	34	33	29	25	25
88	87	85	84	80	75	72
117	116	117	117	115	113	111
119	118	119	119	120	120	119
126	125	123	122	123	125	124
127	127	126	124	123	125	127
111	110	115	111	110	113	116
88	90	88	88	88	89	91
72	74	73	73	74	74	74
59	61	62	63	64	64	64
56	58	52	54	56	57	57

Colour constancy

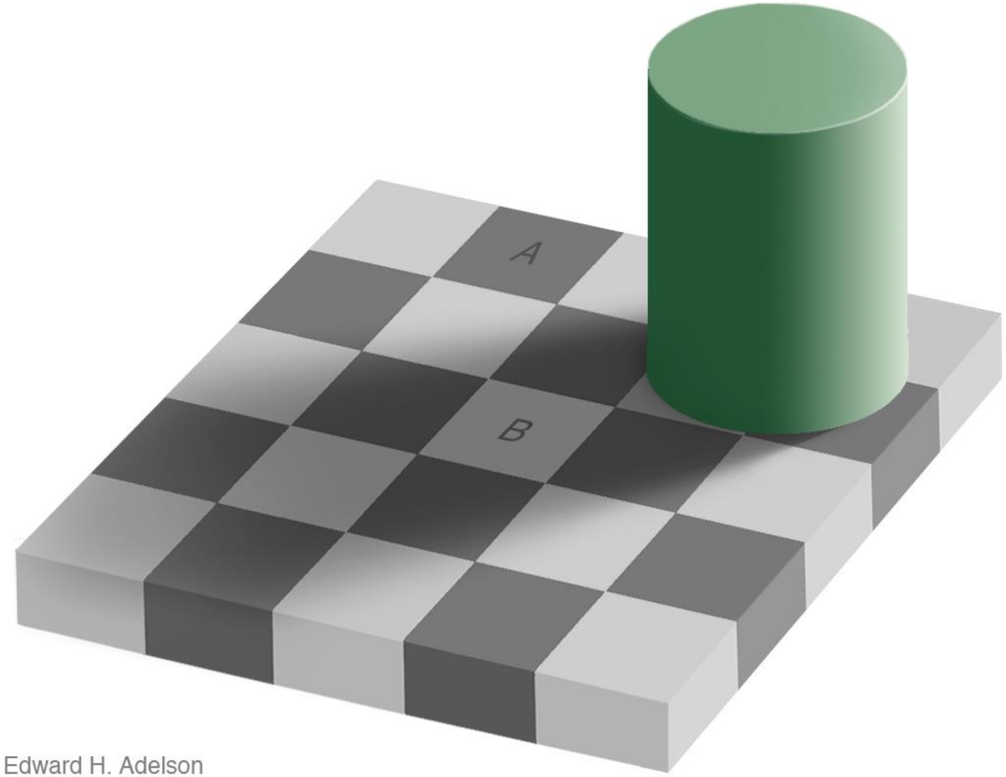

Edward H. Adelson

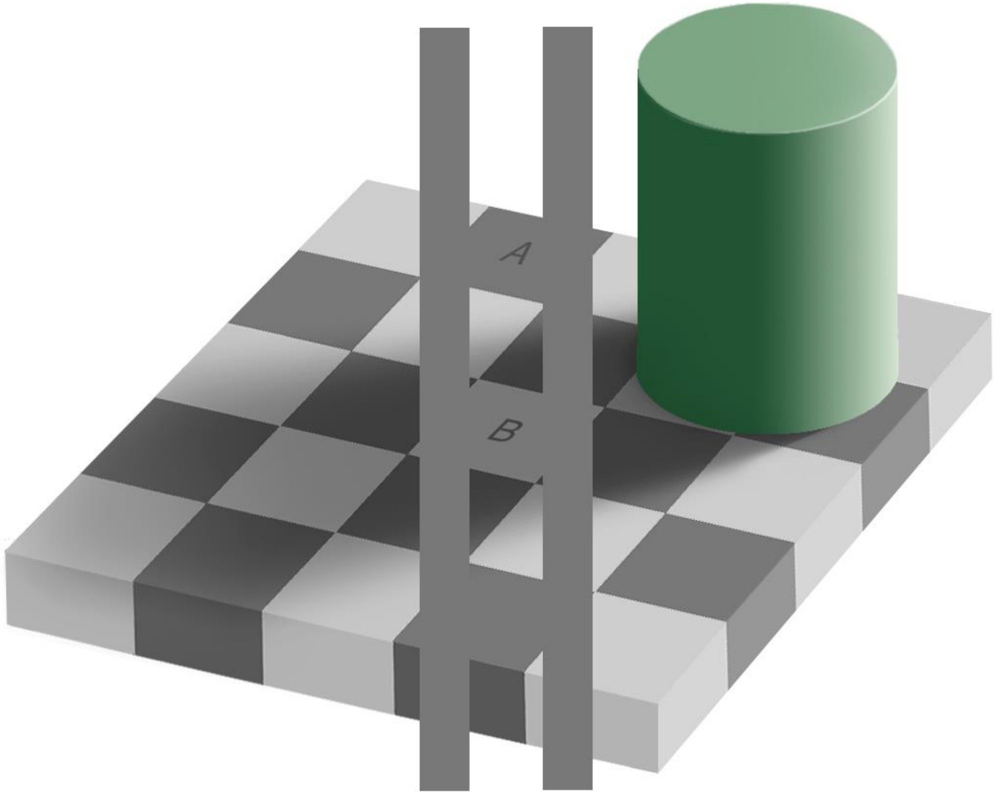

universiteit hasselt | EDM | iMinds

Brightness illusions

Why is it difficult?

- Material-lighting ambiguity

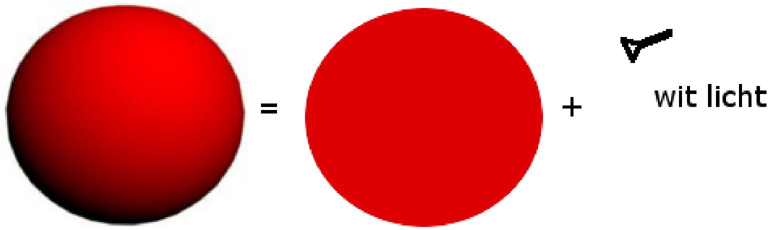

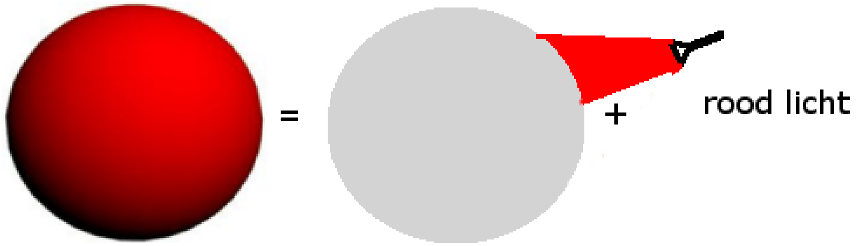

Why is it difficult?

- Problem has a high dimensionality

- Material info is 6D
  - Incident light direction (2D)
  - Viewing direction (2D)
  - Position (2D)

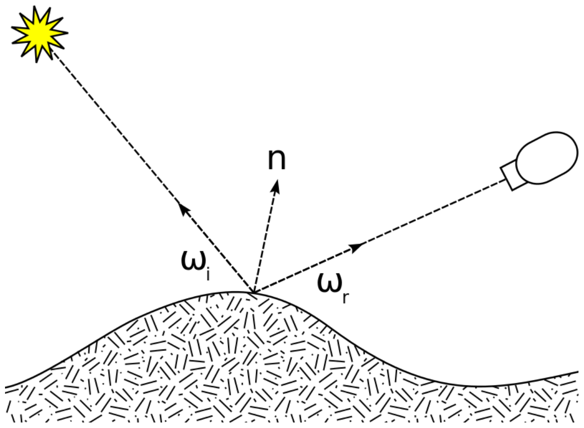

So what can we expect to recover?

- Depends on the assumptions we make (Ramamoorthi, 2001)
  - BRDF known / lighting unknown
  - Lighting known / BRDF unknown
  - Both factors unknown
- Studies various cases
  - Factorised lighting
  - Spherical harmonics used as mathematical tool
  - ~ Fourier series on the sphere
  - Recoverable frequencies proved
- Ringing / Low frequency only

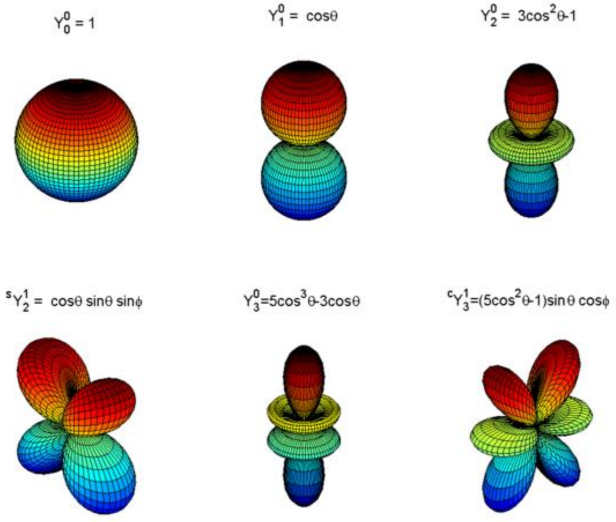

Spherical harmonics

General conclusion

- In order to have high frequencies in the result, you need to have high frequencies in the BRDF and illumination factor
- Estimating high frequency BRDF requires high frequency lights
- Estimating high frequency lights requires high frequency BRDF
- Unsolved problem: solution only unique up to a scale factor
  - smoothness factor

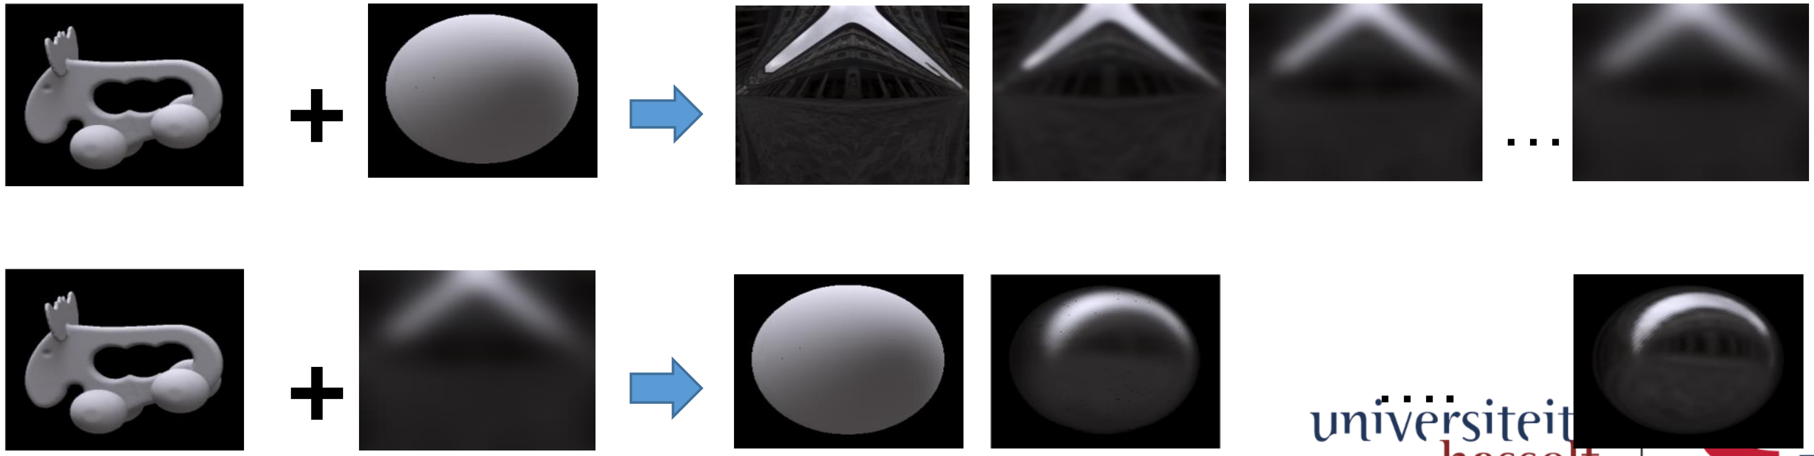

Choosing a good representation

- $B(x,\omega_o) = \int_\Omega V(x,\omega_i)\,\rho(x,\omega_i,\omega_o)(\omega_i \cdot n(x))\,L(x,\omega_i)\,d\omega_i$

$V(x,\omega_i)$	$\rho(x,\omega_i,\omega_o)(\omega_i \cdot n(x))$	$L(x,\omega_i)$
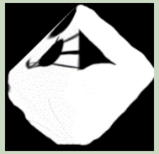	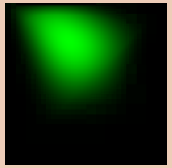	
$V(x,\omega) = \sum_i V_i \Psi_i(\omega)$	$\rho(\omega) = \sum_j \rho_j \Psi_j(\omega)$	$L(\omega) = \sum_k L_k \Psi_k(\omega)$

$$B(x,\omega_o) = \sum_i \sum_j \sum_k V_i \tilde{L}_j \rho_k C_{ijk} \quad with\ C_{ijk} = \int_\Omega \Psi_i(\omega_i)\Psi_j(\omega_i)\Psi_k(\omega_i)\,d\omega_i$$

universiteit hasselt | EDM | iMinds

Triple product integral

$$
\begin{aligned}
B &= \int_{S^2} L(\omega)\, V(\omega)\, \tilde{\rho}(\omega)\, d\omega \\
&= \int_{S^2} \left(\sum_i L_i \Psi_i(\omega)\right)\left(\sum_j V_j \Psi_j(\omega)\right)\left(\sum_k \tilde{\rho}_k \Psi_k(\omega)\right) d\omega \\
&= \sum_i \sum_j \sum_k L_i\, V_j\, \tilde{\rho}_k \int_{S^2} \Psi_i(\omega)\Psi_j(\omega)\Psi_k(\omega)\, d\omega \\
&= \sum_i \sum_j \sum_k L_i\, V_j\, \tilde{\rho}_k\, C_{ijk}
\end{aligned}
$$

Triple Product Wavelet Integrals for All-Frequency Relighting, Ren Ng, Ravi Ramamoorthi, Pat Hanrahan, SIGGRAPH 2004

Choosing a good representation

- Wavelets
  - 2D-haar
  - Spherical / high-order wavelets

- Spherical Radial basis functions (SRBFs)
  - ~ mixture of Gaussians

- Eigenbases

2D Haar wavelets

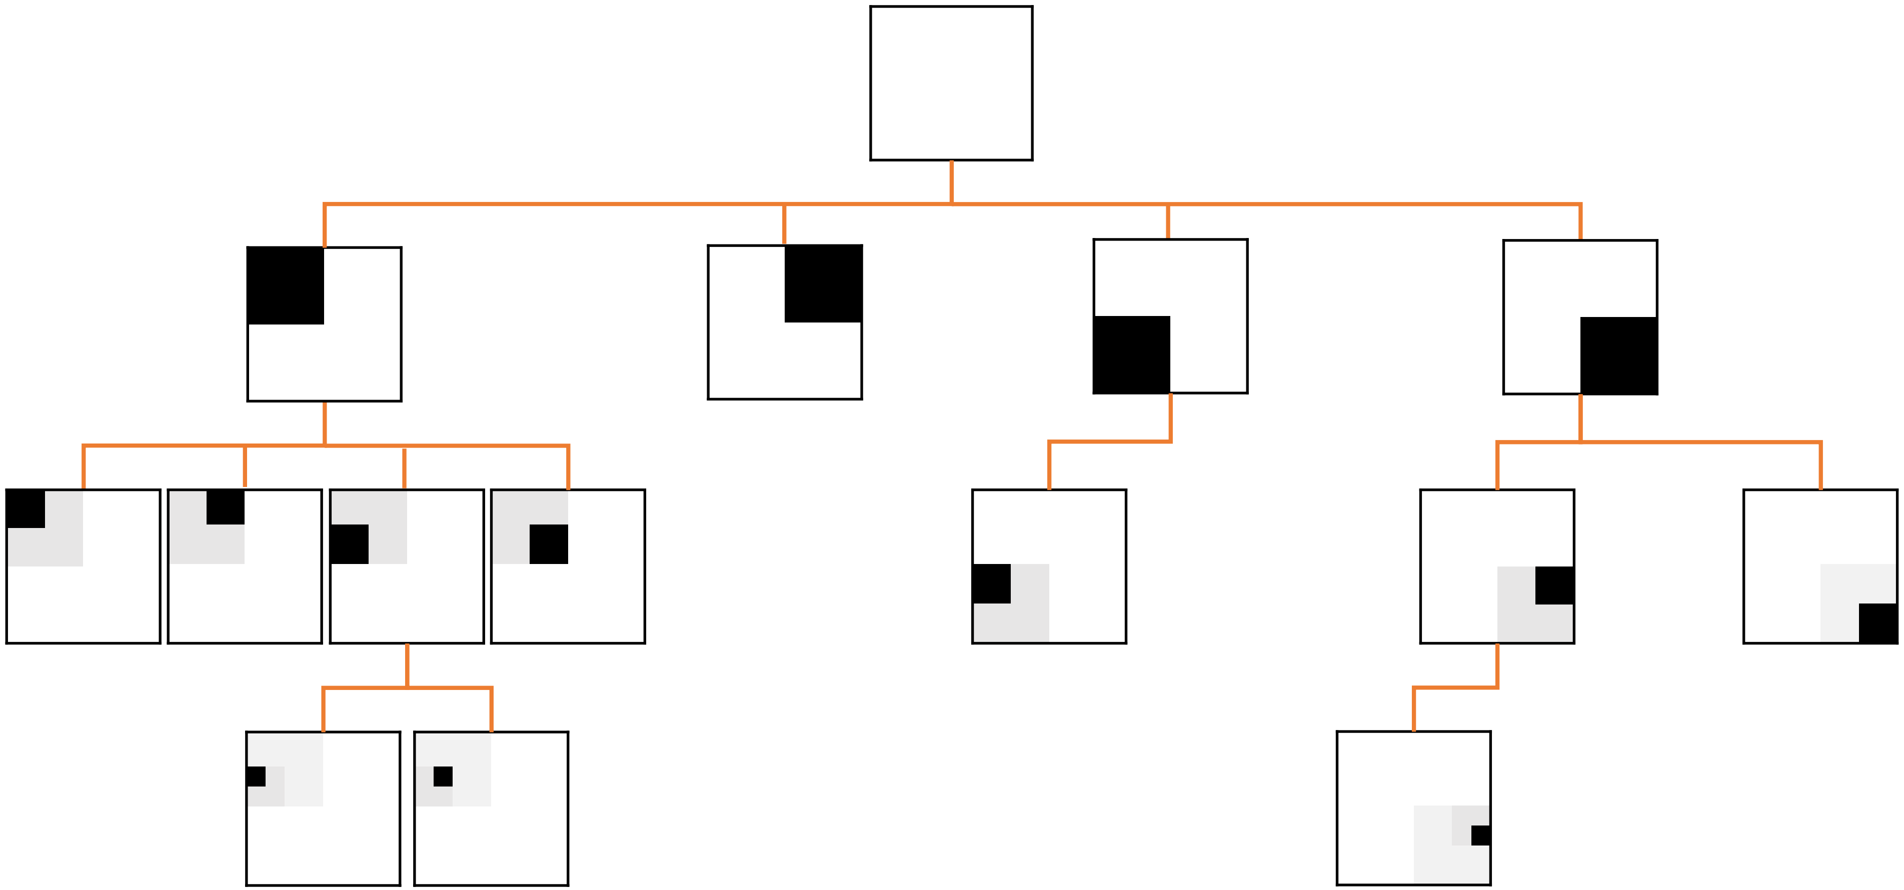

Haar wavelet representation

- Compact
  - Localised in spatial and frequency domain
- Supports high frequencies
- Fast triple product (Ng. et al, 2004)

$$B(x,\omega_o) = \sum_i \sum_j \sum_k V_i \tilde{L}_j \rho_k \boxed{\int_\Omega \Psi_i(\omega_i)\Psi_j(\omega_i)\Psi_k(\omega_i)d\omega_i}$$

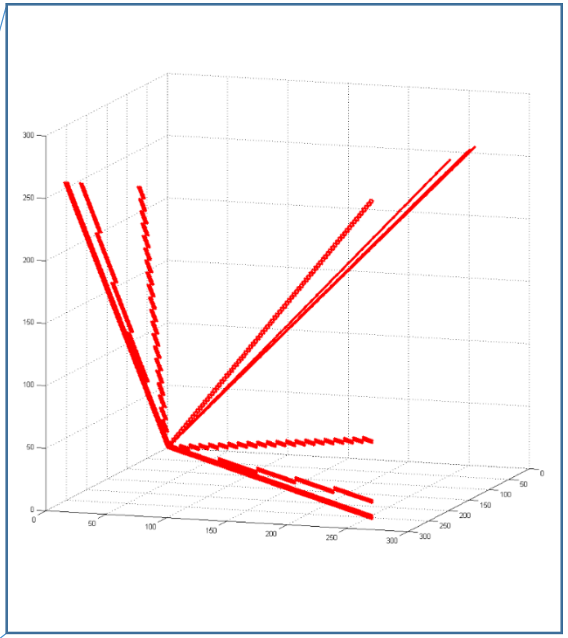

- Reasonably fast rotation method (Wang et al, 2006)

What is $\tilde{L}$?

- Rotated version of lighting
- Wavelet coefficients rotated by precalculated sparse rotation matrices
- O(N), N = number of sparse wavelet coefficients

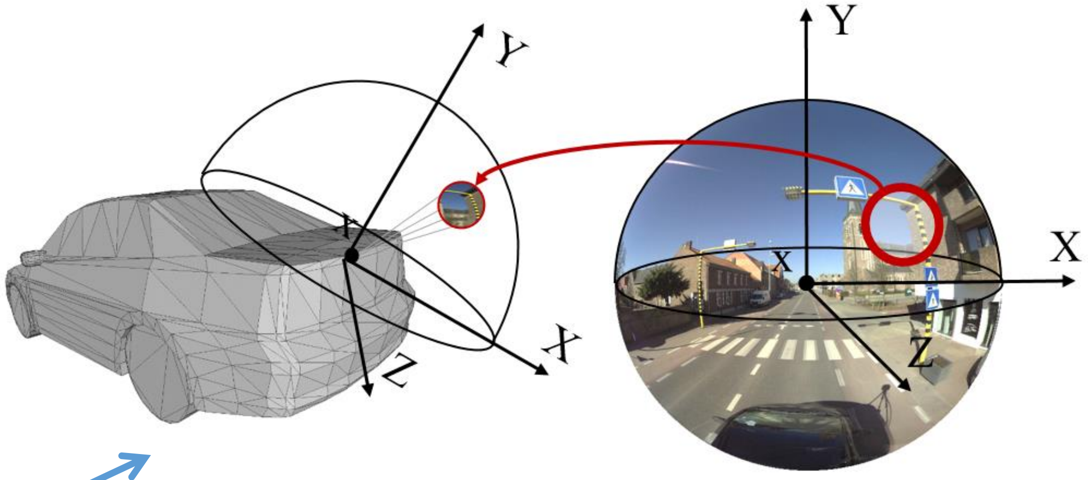

=

Precalculated rotation matrices

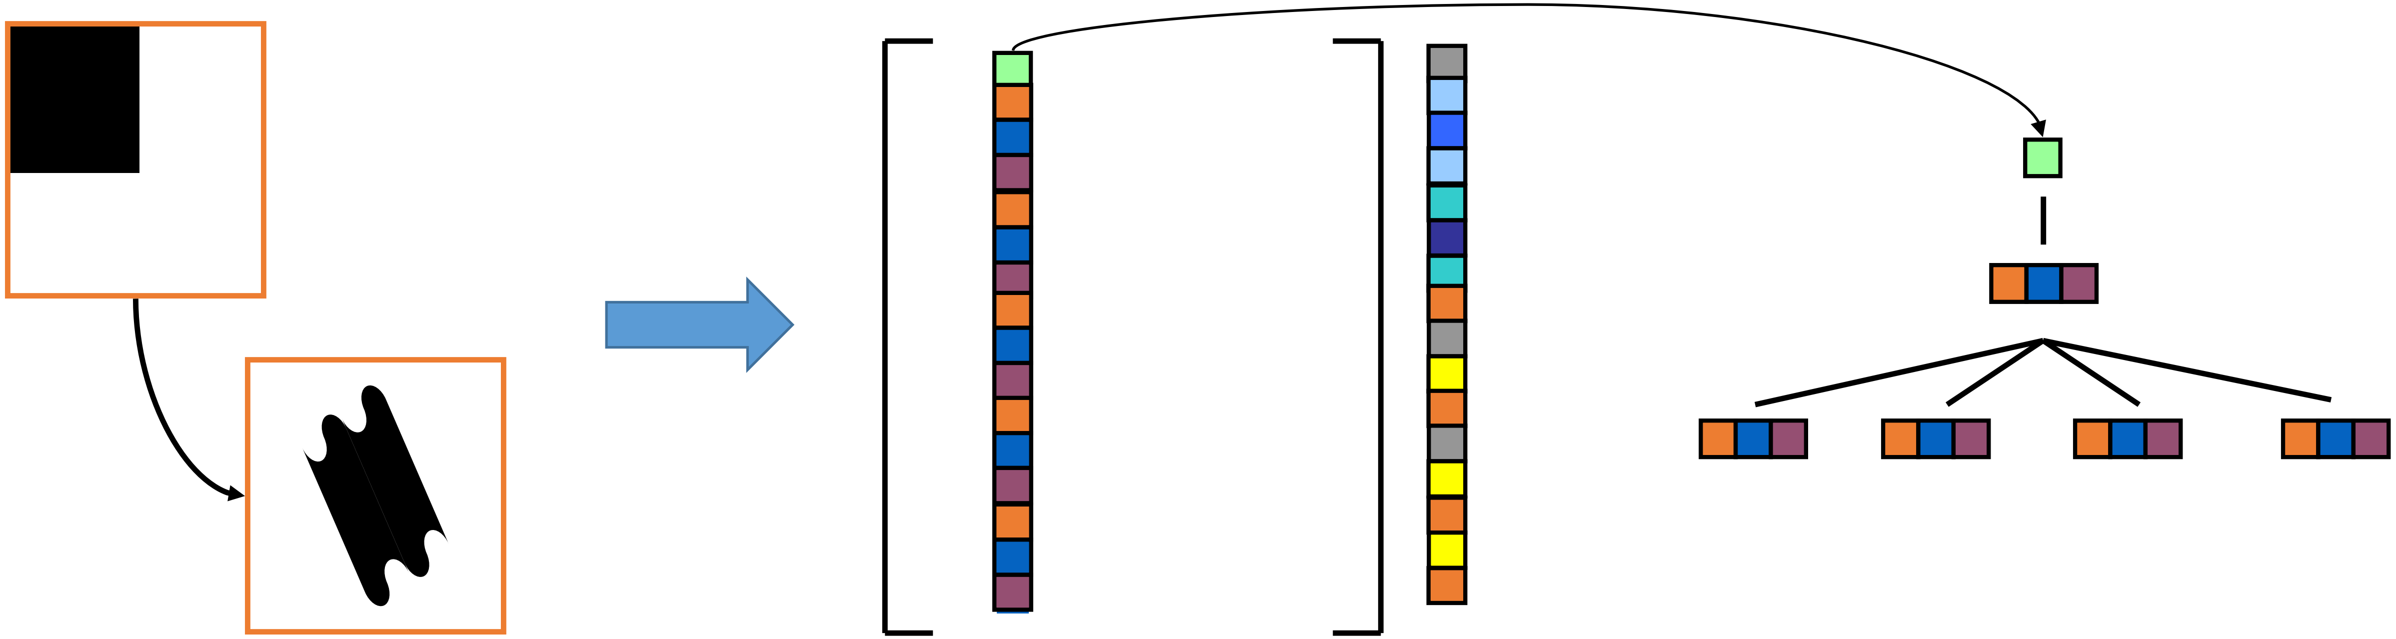

High-level algorithm overview

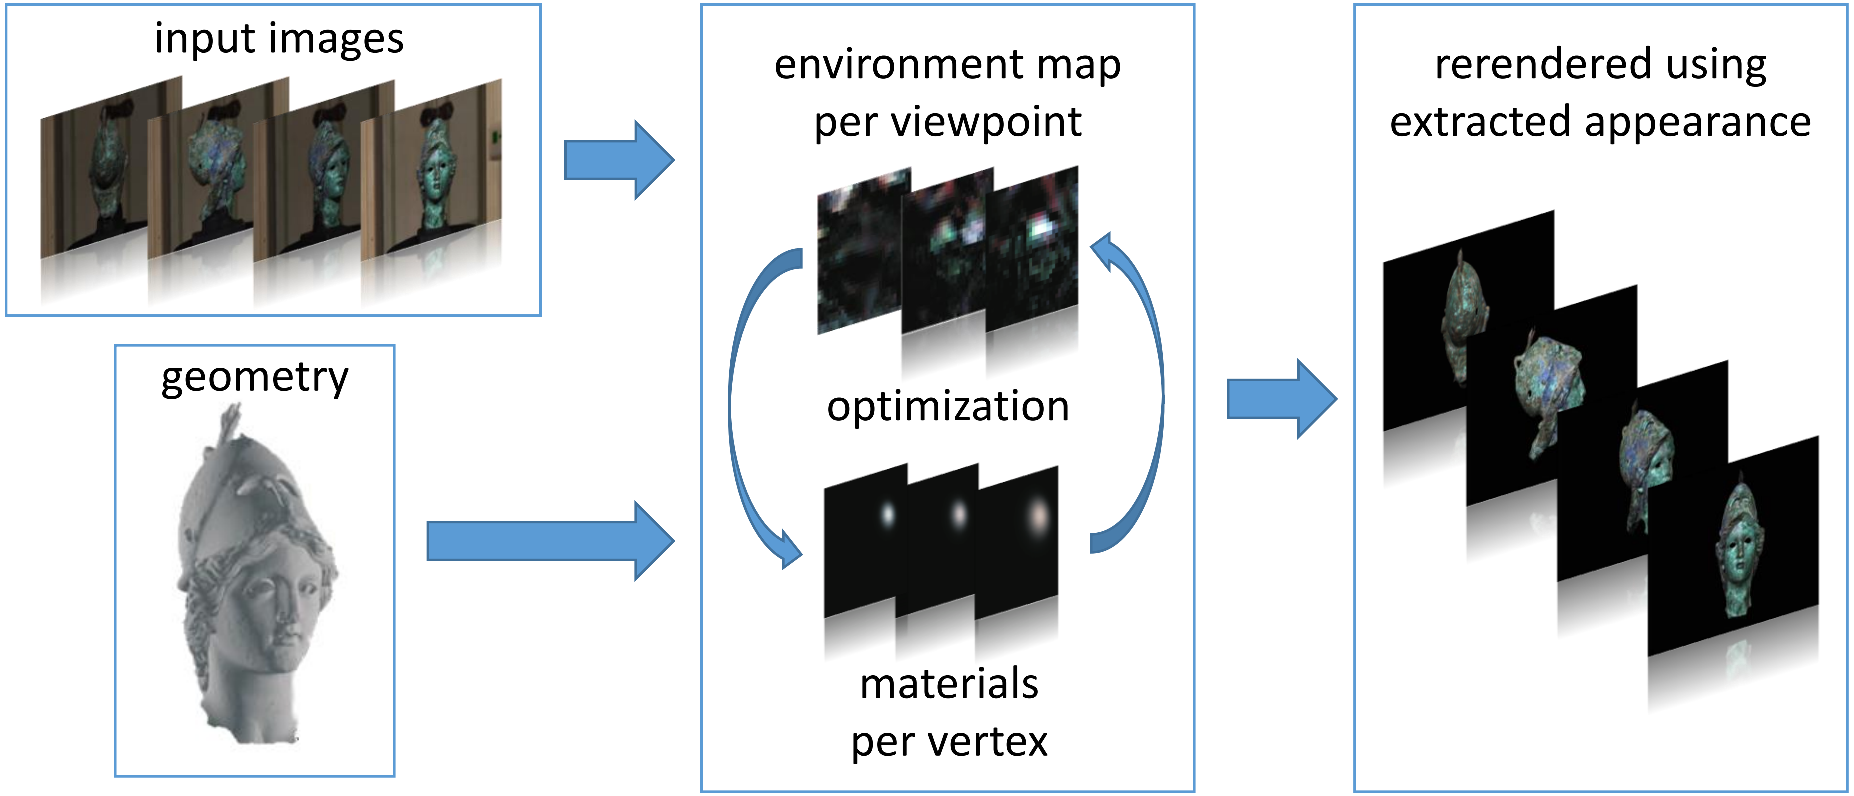

Estimation of Haar
wavelet coefficients

Problem statement

$$L(\vec{x}, w_o) = \sum_k \sum_l \sum_m C_{klm}\, \hat{\rho}_k\, V_l\, \tilde{L}_m$$ (Discrete triple product integral)

$$C_{klm} = \int_\Omega \Psi_k \Psi_l \Psi_m d\omega$$ (Triple product binding coefficients)

$$T_{\vec{x},\, km} = \sum_l C_{klm} V_l$$ (Transfer function)

$$\forall y_p : L_p = \rho_{\vec{x}} T_{\vec{x}} \tilde{L}$$ (Bilinear problem in matrix notation)

$$O = \sum_{p=1}^{N} \alpha_p (y_p - \rho_{\vec{x}} T_{\vec{x}} \tilde{L})^2$$ (Optimisation objective function)

How do we solve the problem?

- Fundamentally underconstrained!
- Approximate best solution with Quadratic Optimisation
  - Fast primal-dual interior point solver : OOQP (Gertz, 2003)

$$O = \sum_{p=1}^{N} \alpha_p (y_p - \rho_{\vec{x}} T_{\vec{x}} \tilde{L})^2 \Leftrightarrow \min_x \|Mx - Y\|_2$$

$$\Leftrightarrow \left\{ \begin{array}{l} \min_x \frac{1}{2} x^T Q x + c^T x \\ \cancel{Ax = b}, \quad \boxed{Cx \geq d} \end{array} \right.$$

Not used

Pixels positive

$M \in \mathbb{R}^{pixels\ x\ coeffs}$

$Q \in \mathbb{R}^{coeffs\ x\ coeffs}$ $\qquad Q = M^T M$

$c \in \mathbb{R}^{coeffs}$ $\qquad c = M^T Y$

$x \in \mathbb{R}^{coeffs\ x\ 1}$ $\qquad Y \in \mathbb{R}^{pixelsx1}$

Solution conditions

- The Quadratic problem is convex and has a unique global optimum
  - Can be proved
- However, local minimum still possible in bilinear problem
  - Alternating between L and ρ suboptimal

Trade-offs

- Static lighting vs varying lighting
- Single- vs multi-view
- Haber et al. [HFB+09]
  - All frequency wavelet framework
  - Incident illumination per image
  - Reflectance per surface point

Future work

- Support for local lighting
  - Most techniques require light at infinity
  - Hard to estimate lights inside the scene
  - different lighting information per pixel

    = too much data!

- Support for indirect lighting
  - Currently ignored in equations

Papers to read

- A signal processing framework for inverse rendering, Ramamoorthi, 2001
- Relighting objects from image collections, Haber et al, 2009

Expect high-level questions about the algorithms/processes on exam

Questions?