Eurographics 2009

Advanced Illumination Techniques for GPU Volume Raycasting

Markus Hadwiger
VR VIS Research Center
Vienna, Austria

Patric Ljung
Siemens Corporate Research
Princeton, NJ, USA

Christof Rezk Salama
Computer Graphics Group
Institute for Vision and Graphics
University of Siegen, Germany

Timo Ropinski
Visualization and Computer Graphics Research Group,
University of Münster, Germany

Ambient Occlusion in DVR

Markus Hadwiger
VR VIS Research Center
Vienna, Austria

Patric Ljung
Siemens Corporate Research
Princeton, NJ, USA

Christof Rezk Salama
Computer Graphics Group
Institute for Vision and Graphics
University of Siegen, Germany

Timo Ropinski
Visualization and Computer
Graphics Research Group,
University of Münster, Germany

Outline

- Introduction and Concepts
- Vicinity Shading
- Dynamic Ambient Occlusion
- Local Ambient Occlusion
- Global Light Propagation
- Summary

Motivation

- Increased perception of
  - Shapes
  - Densities
  - Depth
- Issues with surface-based shading
  - Noisy data have poorly defined normals/gradients
  - Surface shading for volumetric objects???

Ambient Occlusion

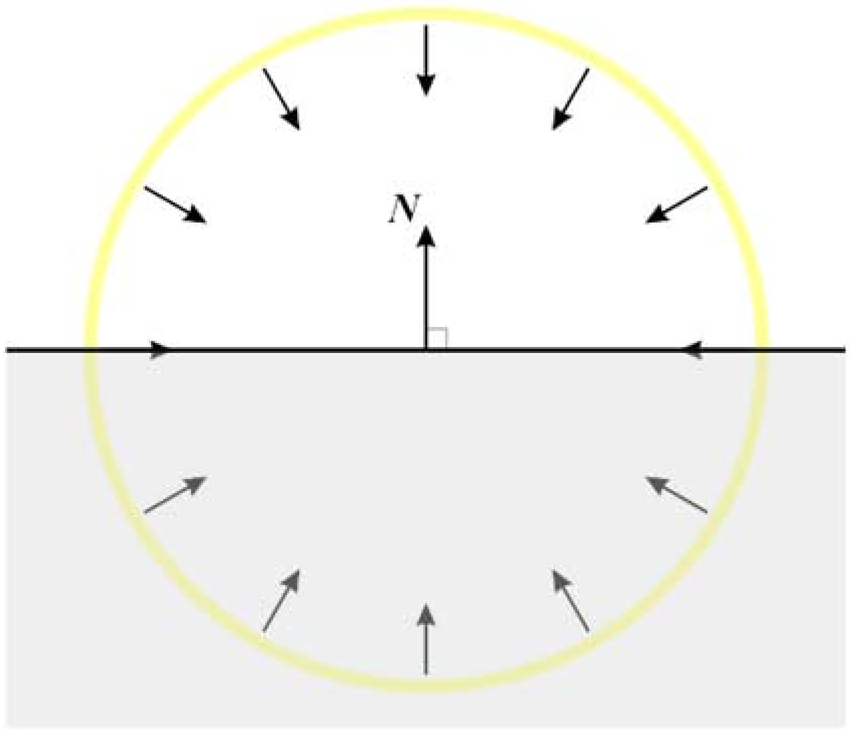

- Global Light Source
- Isotropic Incident Light
- Visible Hemisphere

Ambient Occlusion

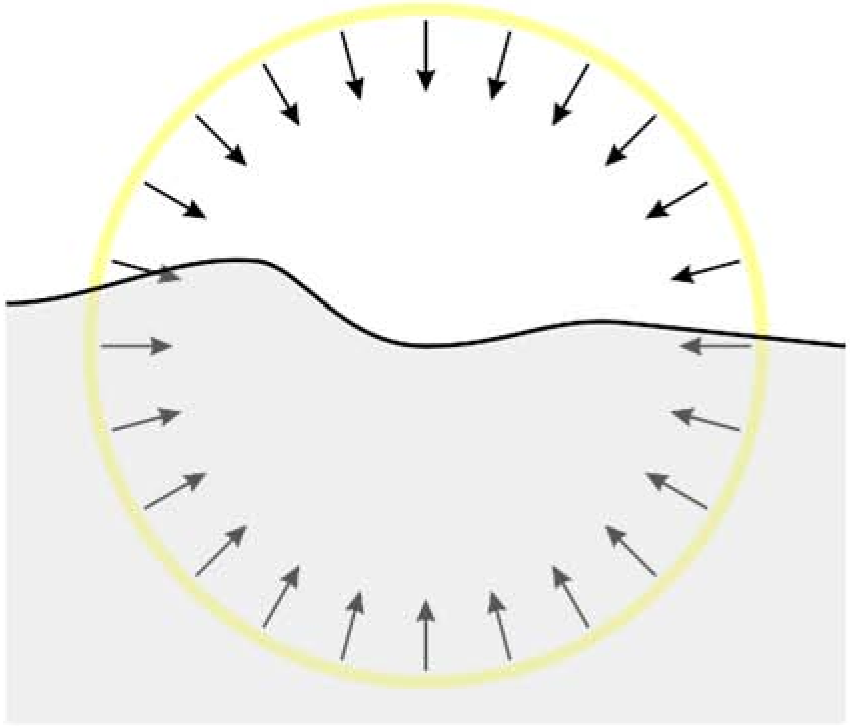

- Global Light Source
- Isotropic Incident Light
- Visible Hemisphere
- Integrate Illumination for a point on the surface
- Sampled directions or projection of occluders
- Accuracy: number of directions and/or resolution of light "cache"

Vicinity Shading

- A. James Stewart, IEEE Vis 2003

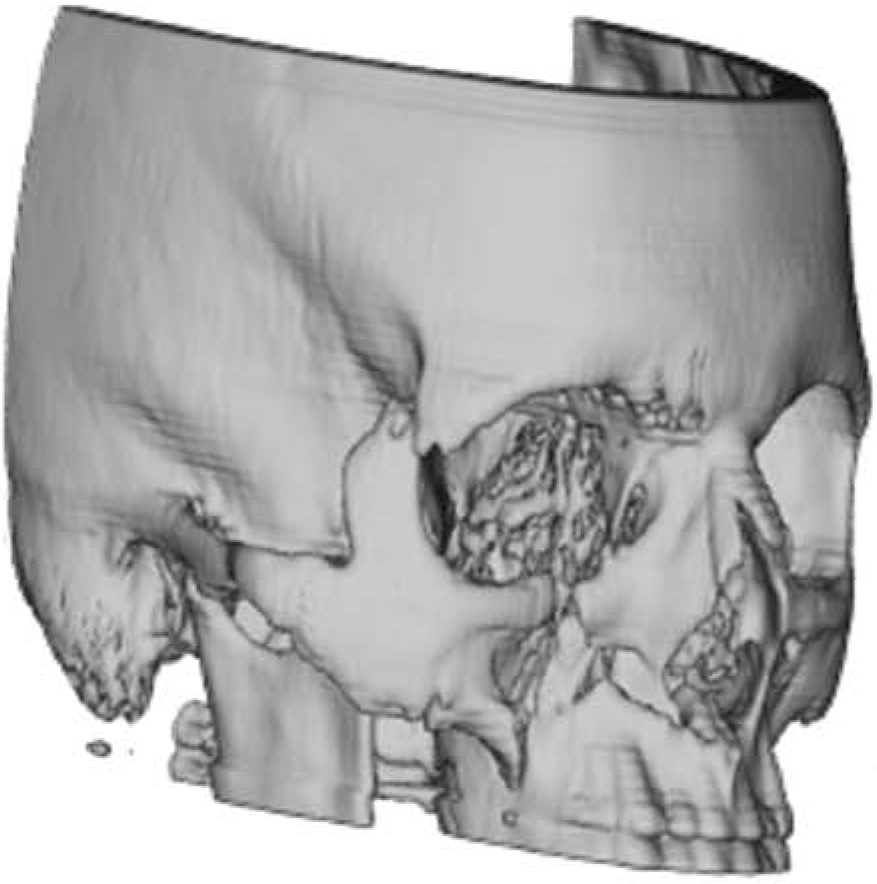

Diffuse Surface Shading

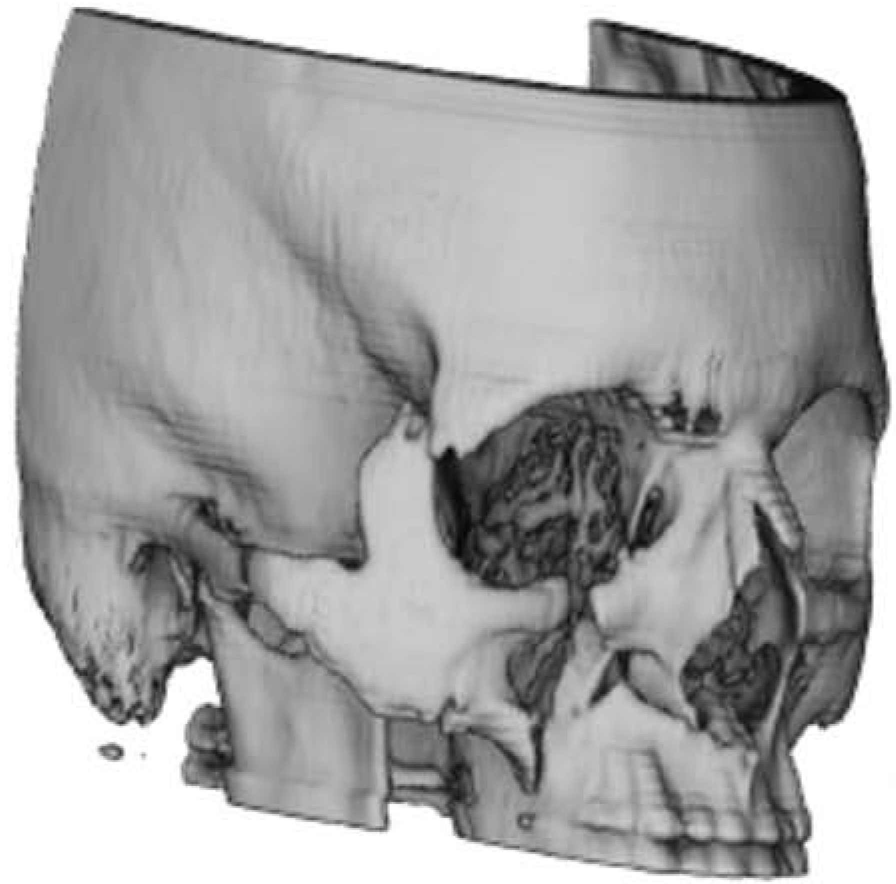

Vicinity Shading

Vicinity Shading Concepts

- Considers only iso-surface rendering
  - Does not consider transparency
- Precomputes visibility for all iso-values
  - Each processed direction traverses each voxel exactly once, using 3D Bresenham line algorithm
- Independent of vicinity radius
- Added irradiance based on distance between voxel location and occluding voxel

Vicinity Shading Performance

Data	Dims	# Directions	Time [mins]	Time/Dir. [s]
Skull	**256 x 256 x 203**	**312**	**14.3**	**2.8**
Skull	**256 x 256 x 203**	**1272**	**59.0**	**2.8**
Cortex	**128 x 512 x 256**	**1272**	**57.5**	**2.7**

Dynamic Ambient Occlusion

- Ropinski et al., Eurographics 2008

Diffuse Surface Shading

Dynamic Ambient Occlusion

Process Flow

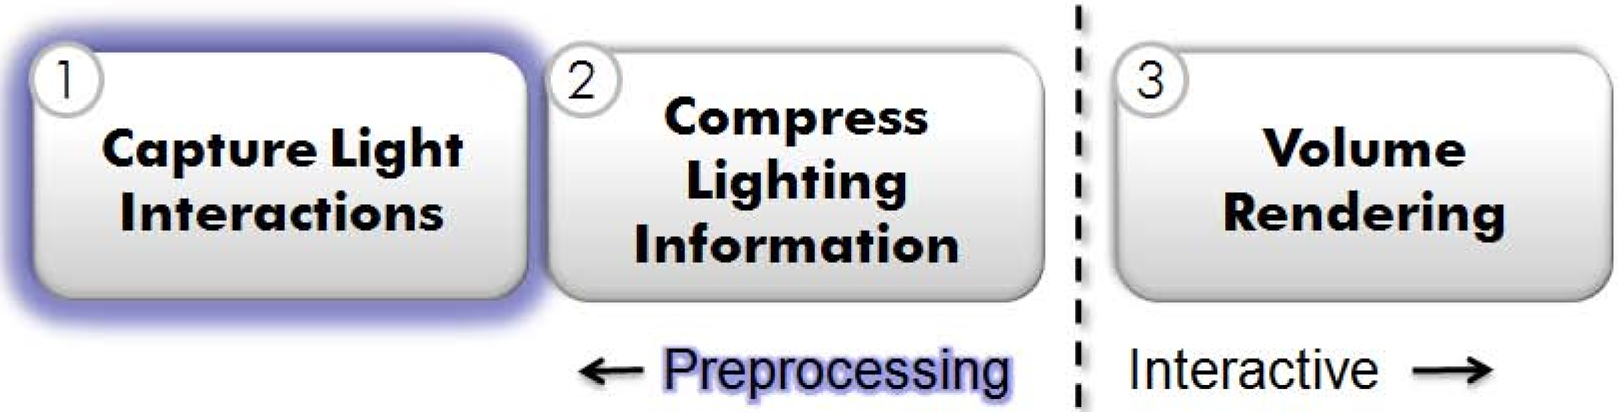

- **Light interactions**
  - … are captured for the **vicinity** only
  - … are expressed by using a **local histogram**

Process Flow

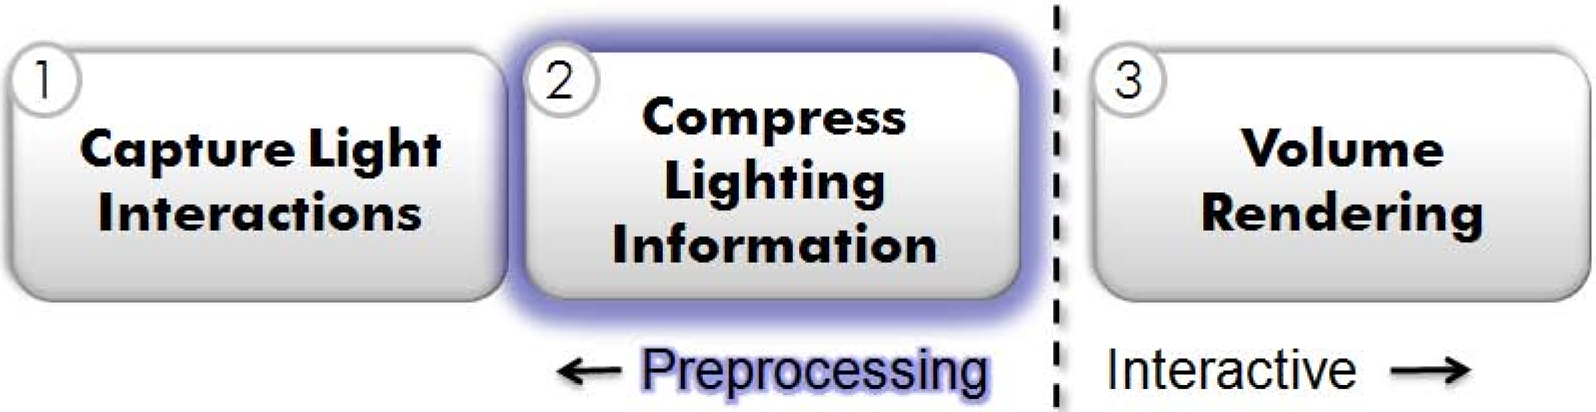

- One local histogram for each voxel results in **unmanageable data sizes**
- Light information is **clustered** to handle it interactively during rendering

Process Flow

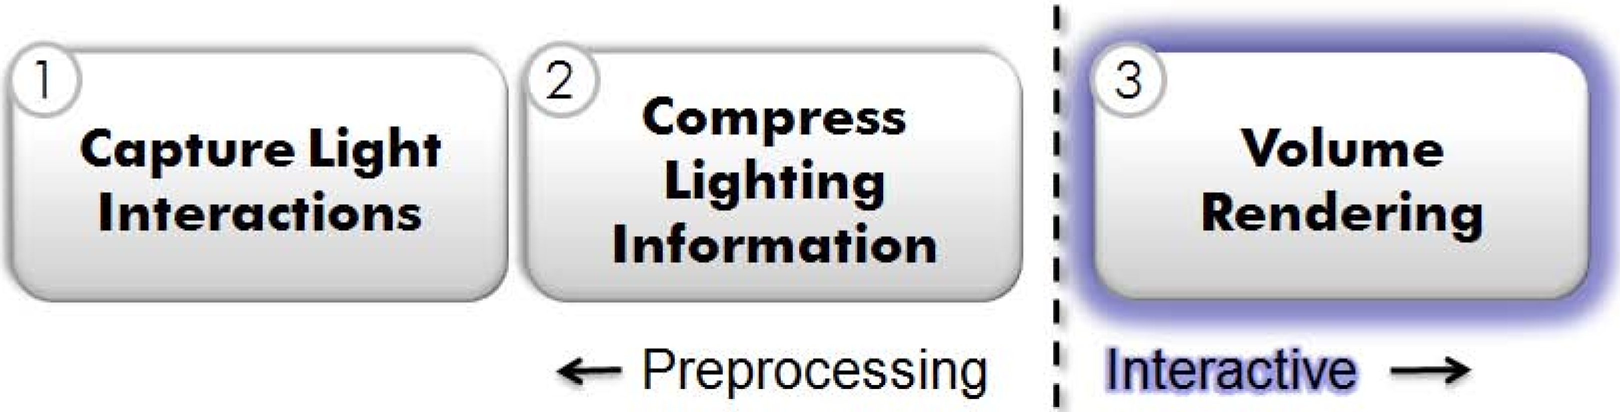

- Rendering parameters are **changed frequently**
- Representative local histograms can be **modulated interactively**
- Volume rendering requires only **two additional texture fetches**

Histogram Generation

- **Analyze vicinity of each voxel**
- **Compute a normalized local histogram with $n = 2^b$ bins**

$$LH(x) = (LH_0(x), \ldots, LH_{n-1}(x))$$

$$LH_k(x) = \sum_{\substack{\tilde{x} \in S_r(x) \\ \tilde{x} \neq x}} f_{dist}\left(\frac{|x - \tilde{x}|}{d_{min}}\right) \cdot g(f(\tilde{x}), k)$$

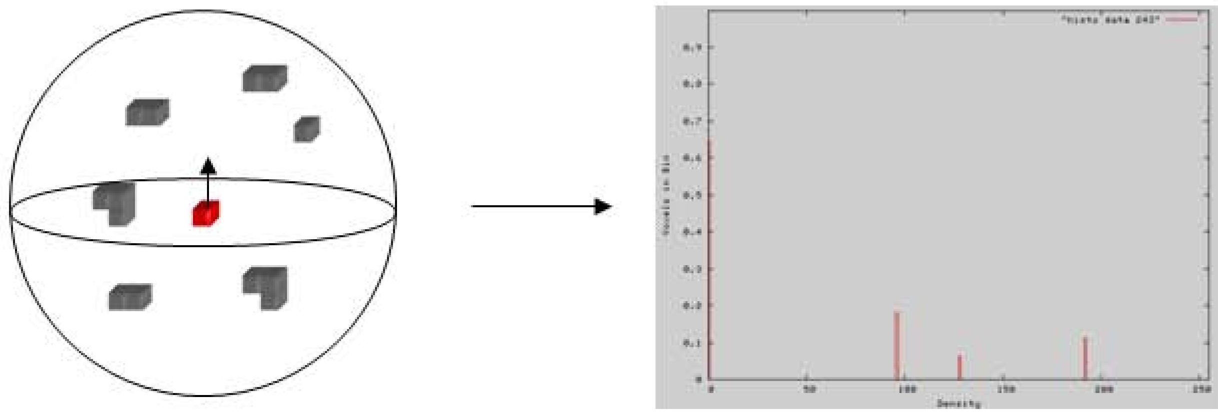

Histogram Generation

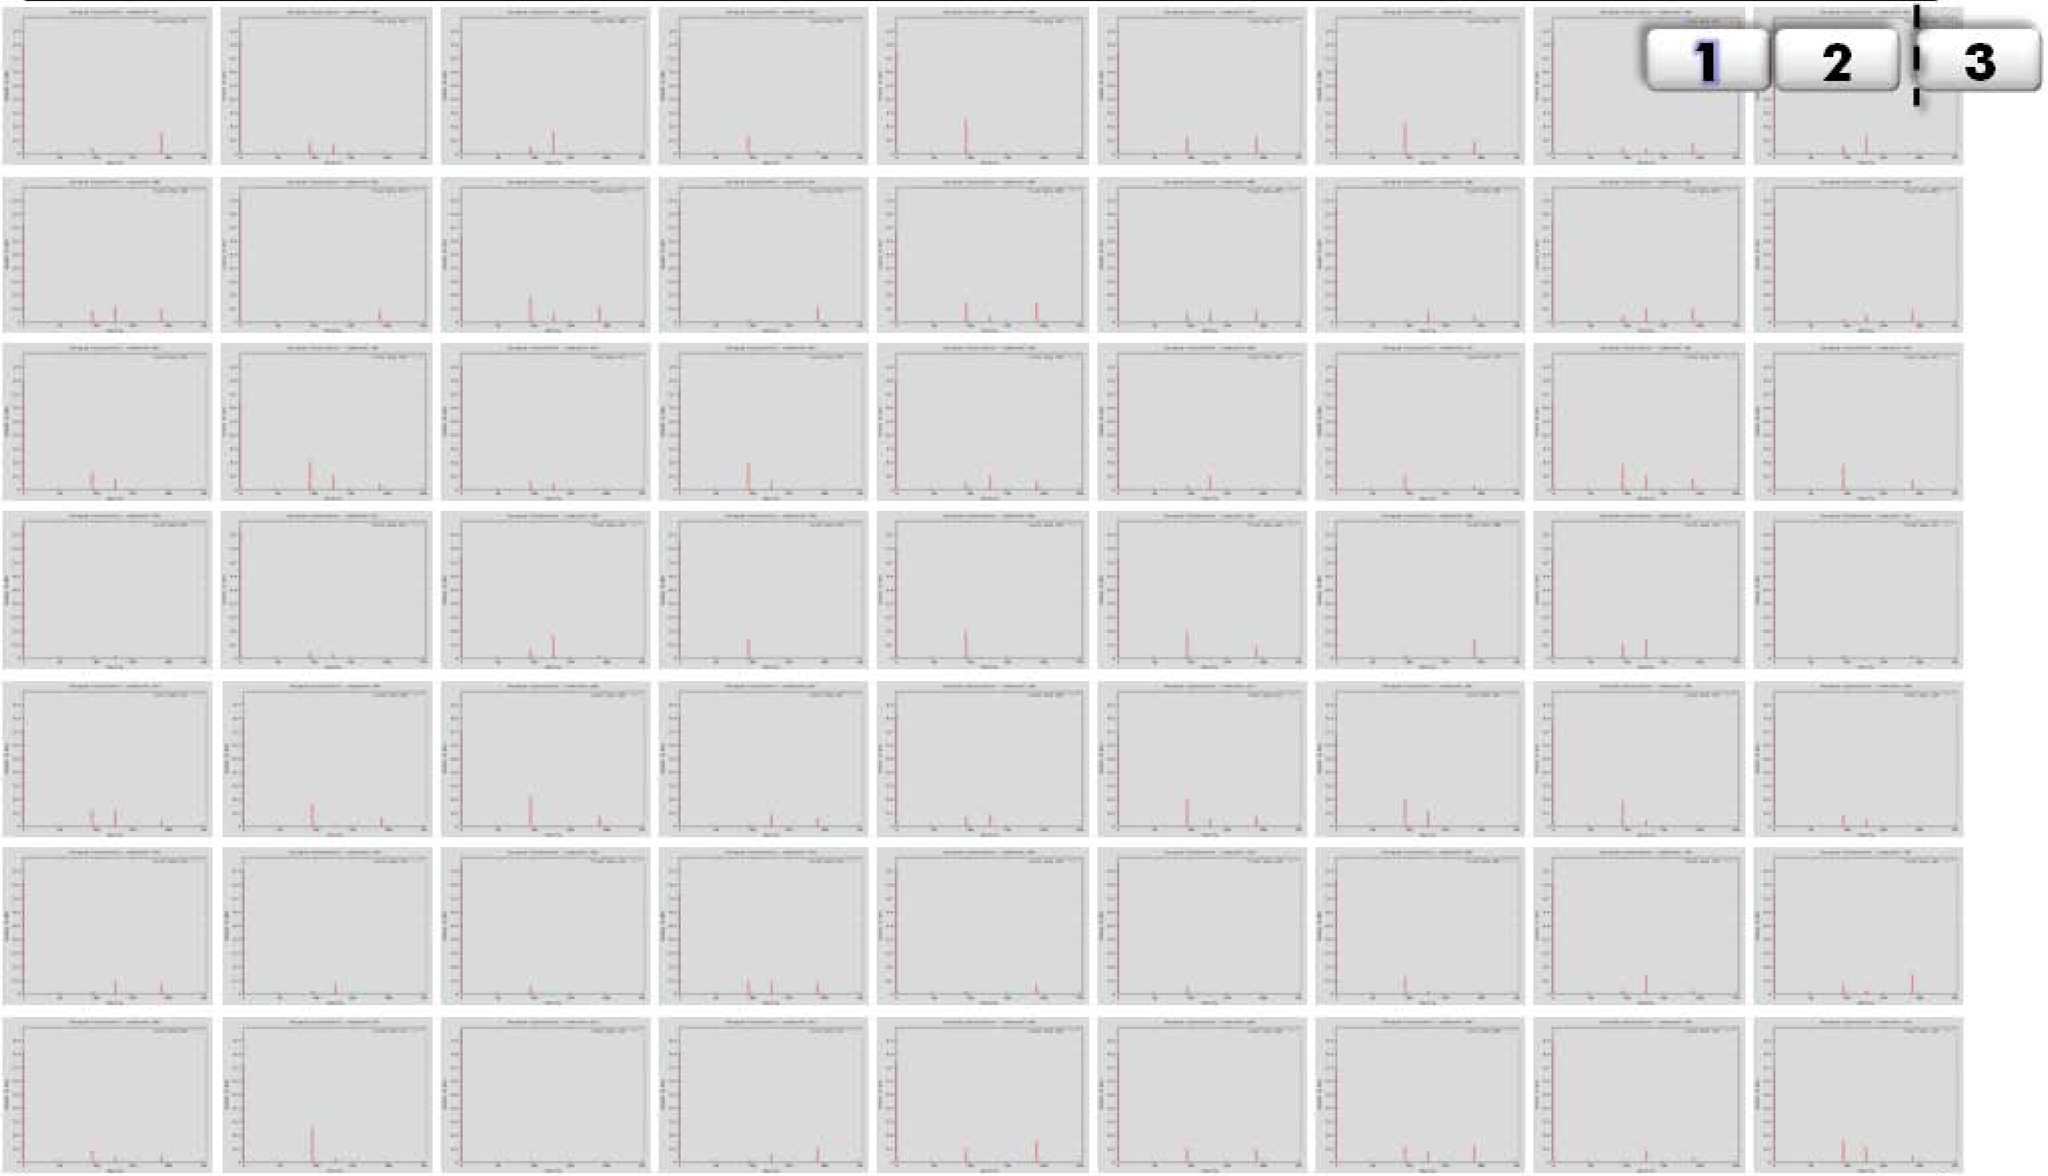

Histogram Clustering

1 2 3

- **Given**: *n* local histograms (vectors)
- **Desired**: *m* << *n* representatives (*code book*)

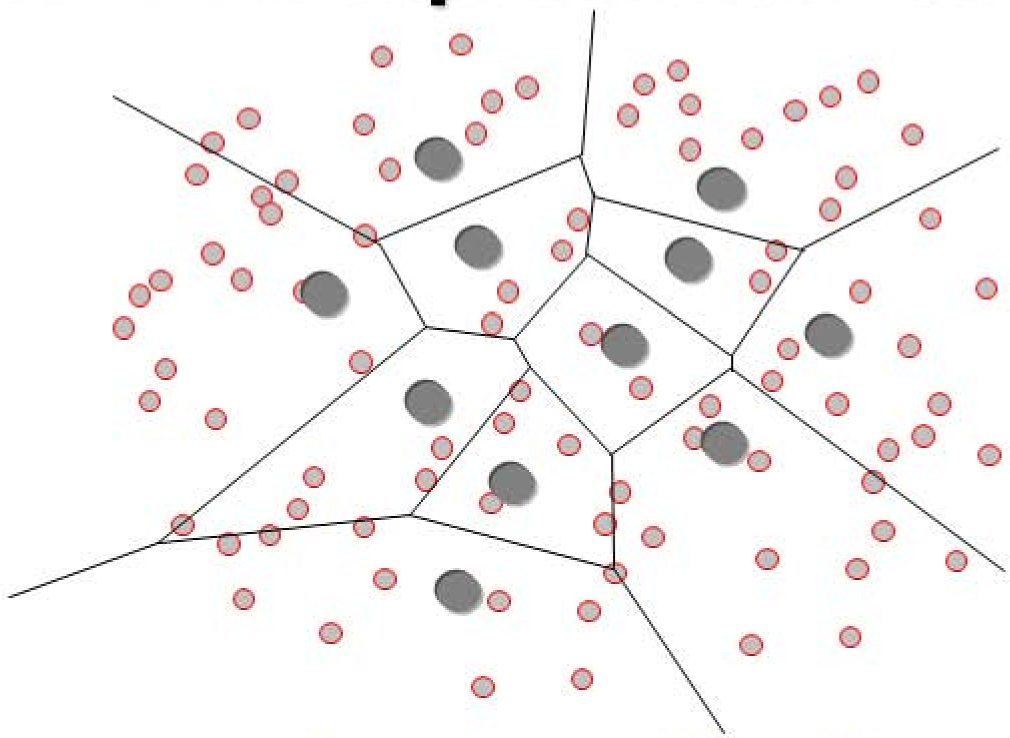

- We exploit a **vector quantization** (vq) for the clustering

Histogram Compression

1 2 3

- **Goal: generate a packed histogram with *j* << *i* bins (*i* = initial number of bins)**
- ***Iterative splitting*** **is used to reduce histogram dimensionality**
- **Example: *i = 12, j = 9***

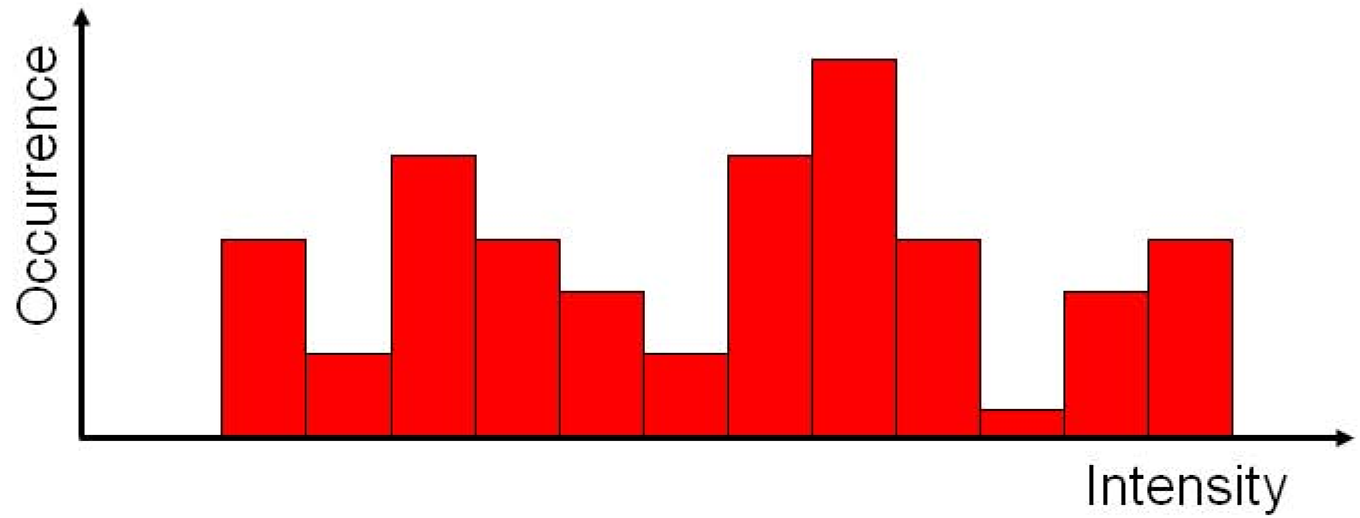

Histogram Compression

1 2 3

- Goal: generate a packed histogram with *j* << *i* bins (*i* = initial number of bins)
- ***Iterative splitting*** is used to reduce histogram dimensionality
- Example: *i = 12, j = 9*

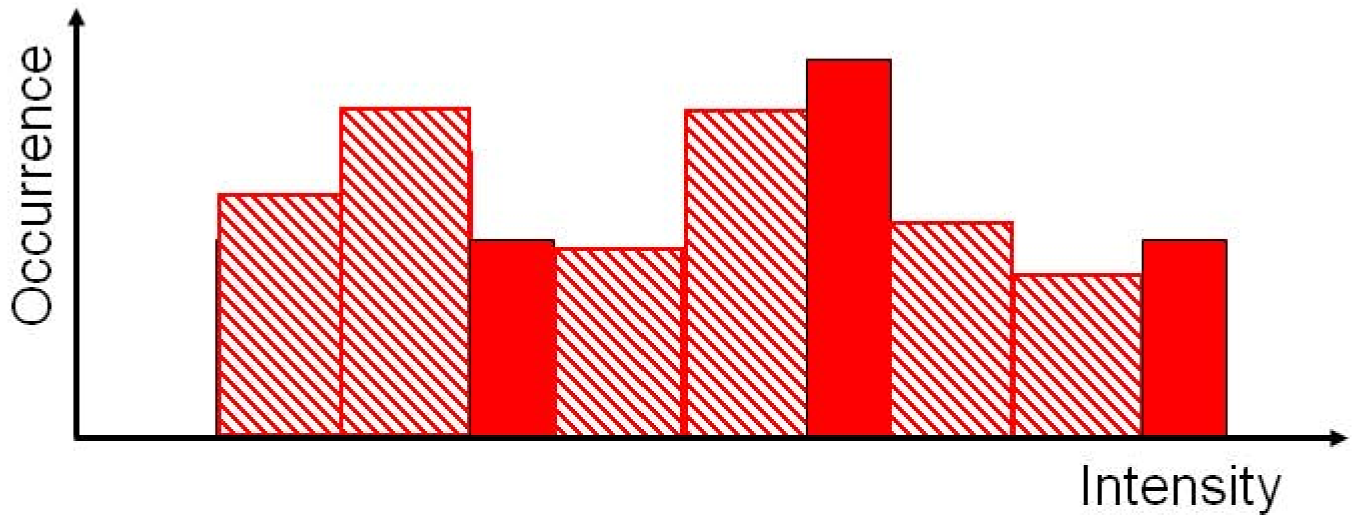

Interactive Rendering

1 2 3

- **Two additional texture fetches** required
  1. Obtain the cluster ID of the current sample x
  2. Fetch the current environment color $E_{env}(x)$

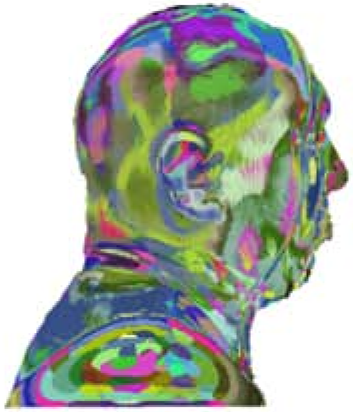

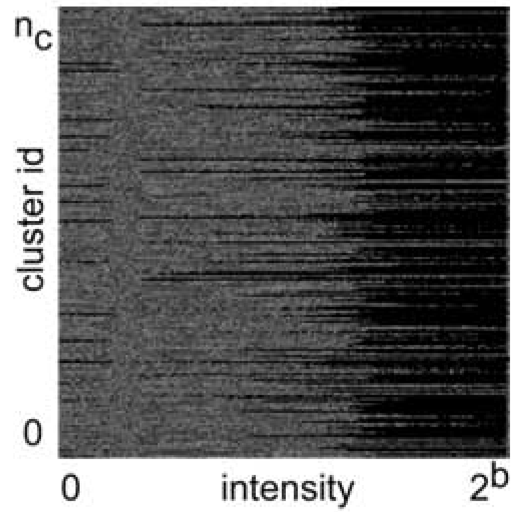

- $E_{env}(x)$ is computed by considering the current transfer function

Isosurface Shading

1 2 3

- **Combination with the transfer function**

$$O_{env}(x, \nabla\tau(f(x))) = \frac{1}{\frac{2}{3}\pi r^3} \sum_{0 \le j < 2^b} \tau_\alpha(j) \cdot LH_j(x)$$

- **Apply Phong shading by using**

$$I_a(x) = 1.0 - O_{env}(x, \nabla\tau(f(x))) \cdot Col_{iso}$$

Isosurface Shading

1 2 3

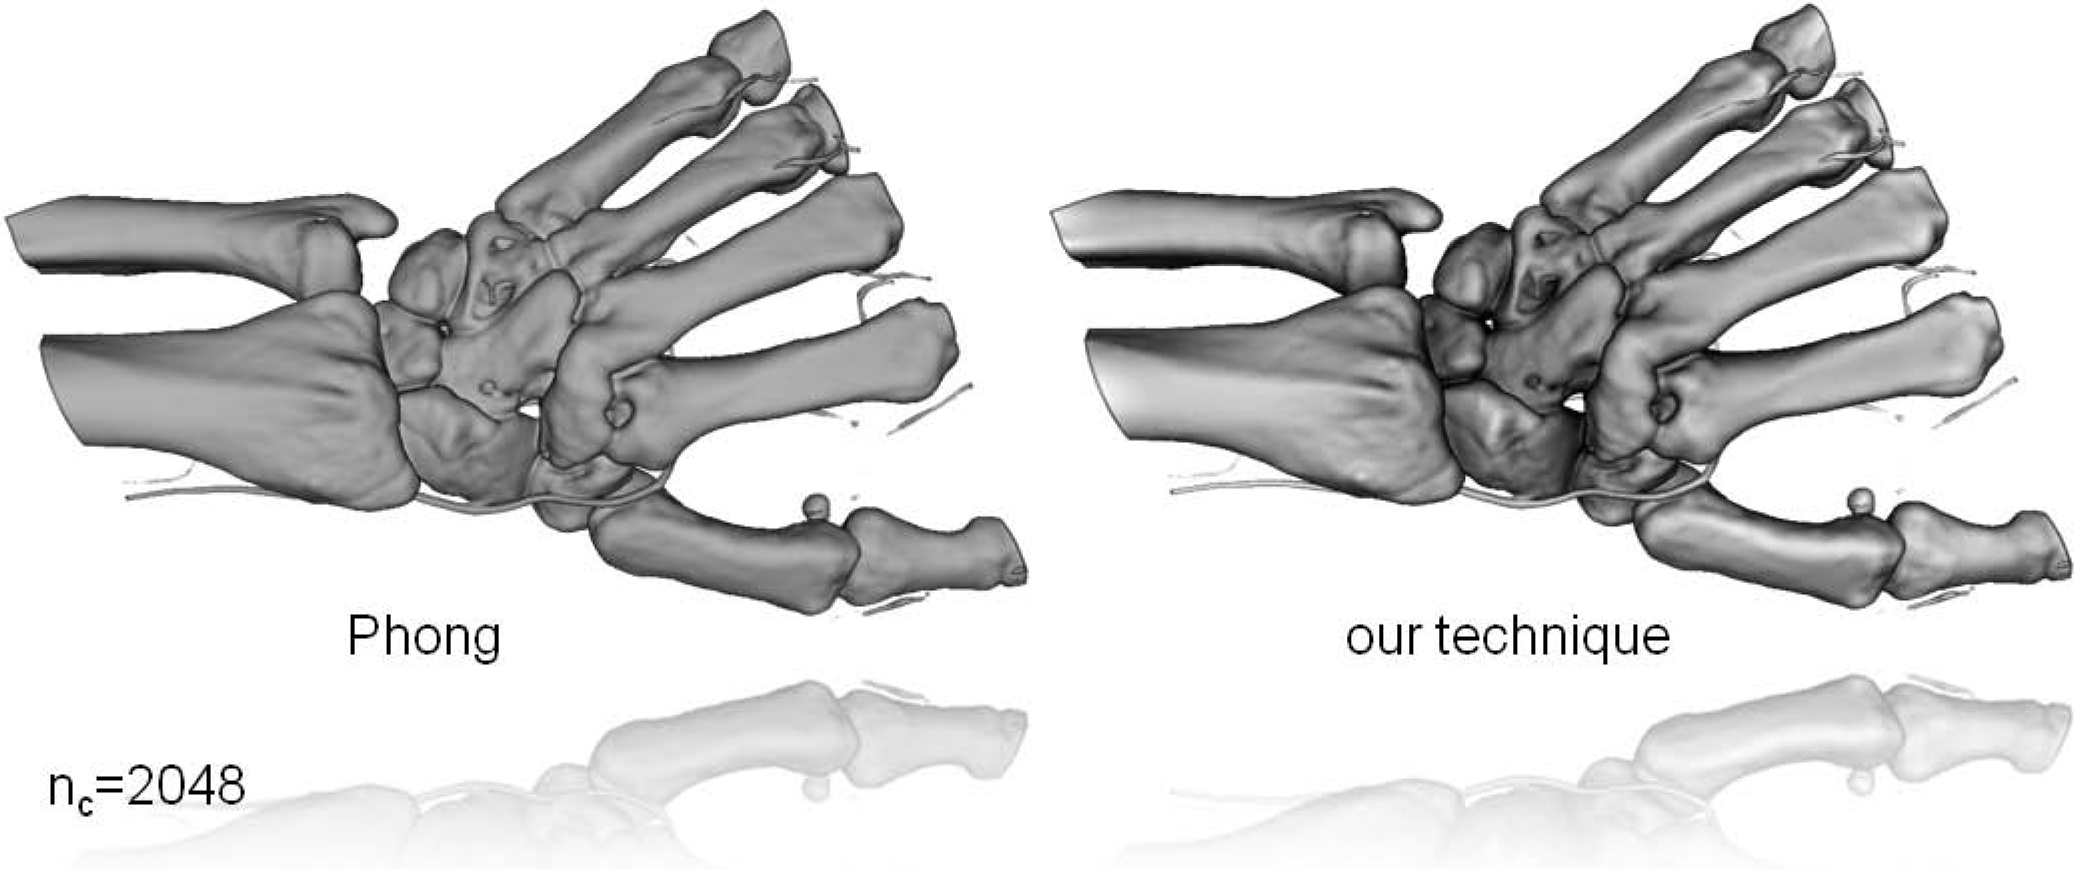

n_c=2048

Isosurface Shading

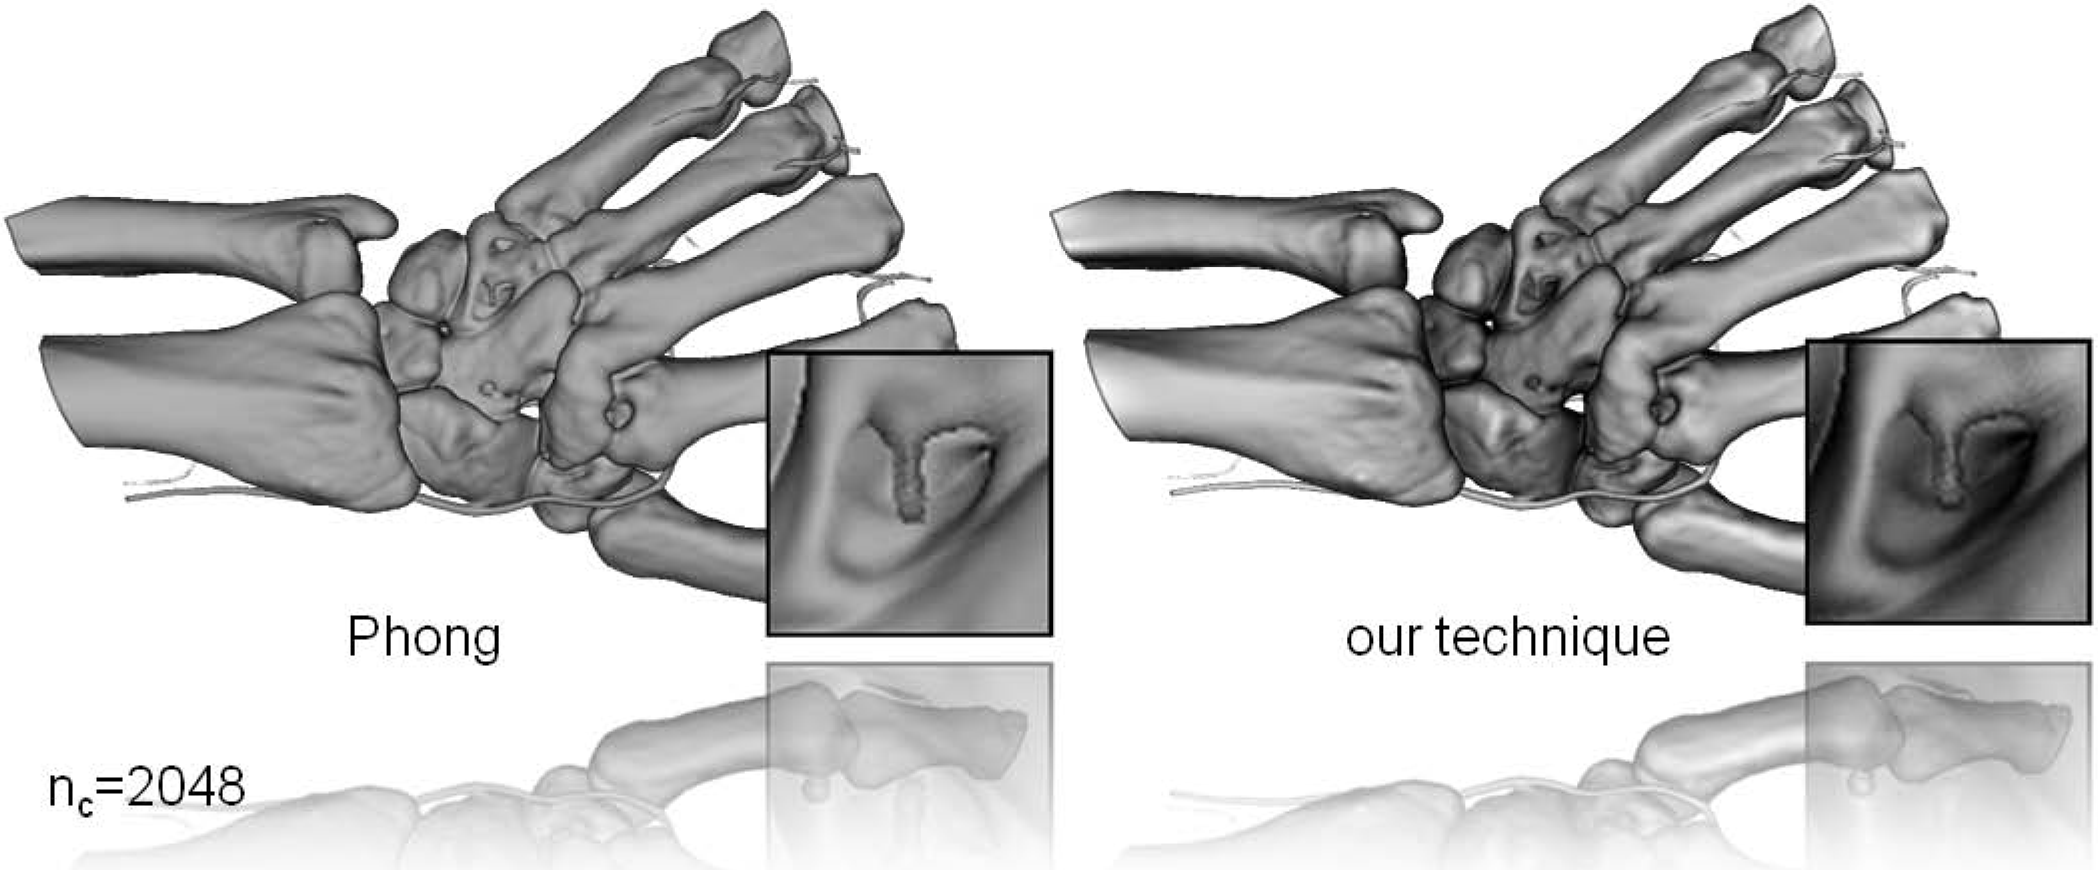

Direct Volume Rendering

1 2 3

- **Challenges**
  - **More than one hue is present**
  - **Areas with participating media may occur**

- **Combination with the transfer function**

$$E_{env}(x, \nabla\tau(f(x))) = \frac{1}{\frac{2}{3}\pi r^3} \sum_{0 \leq j < 2^b} \tau_\alpha(j) \cdot \tau_{rgb}(j) \cdot LH_j(x)$$

Application Example - DVR

1 2 3

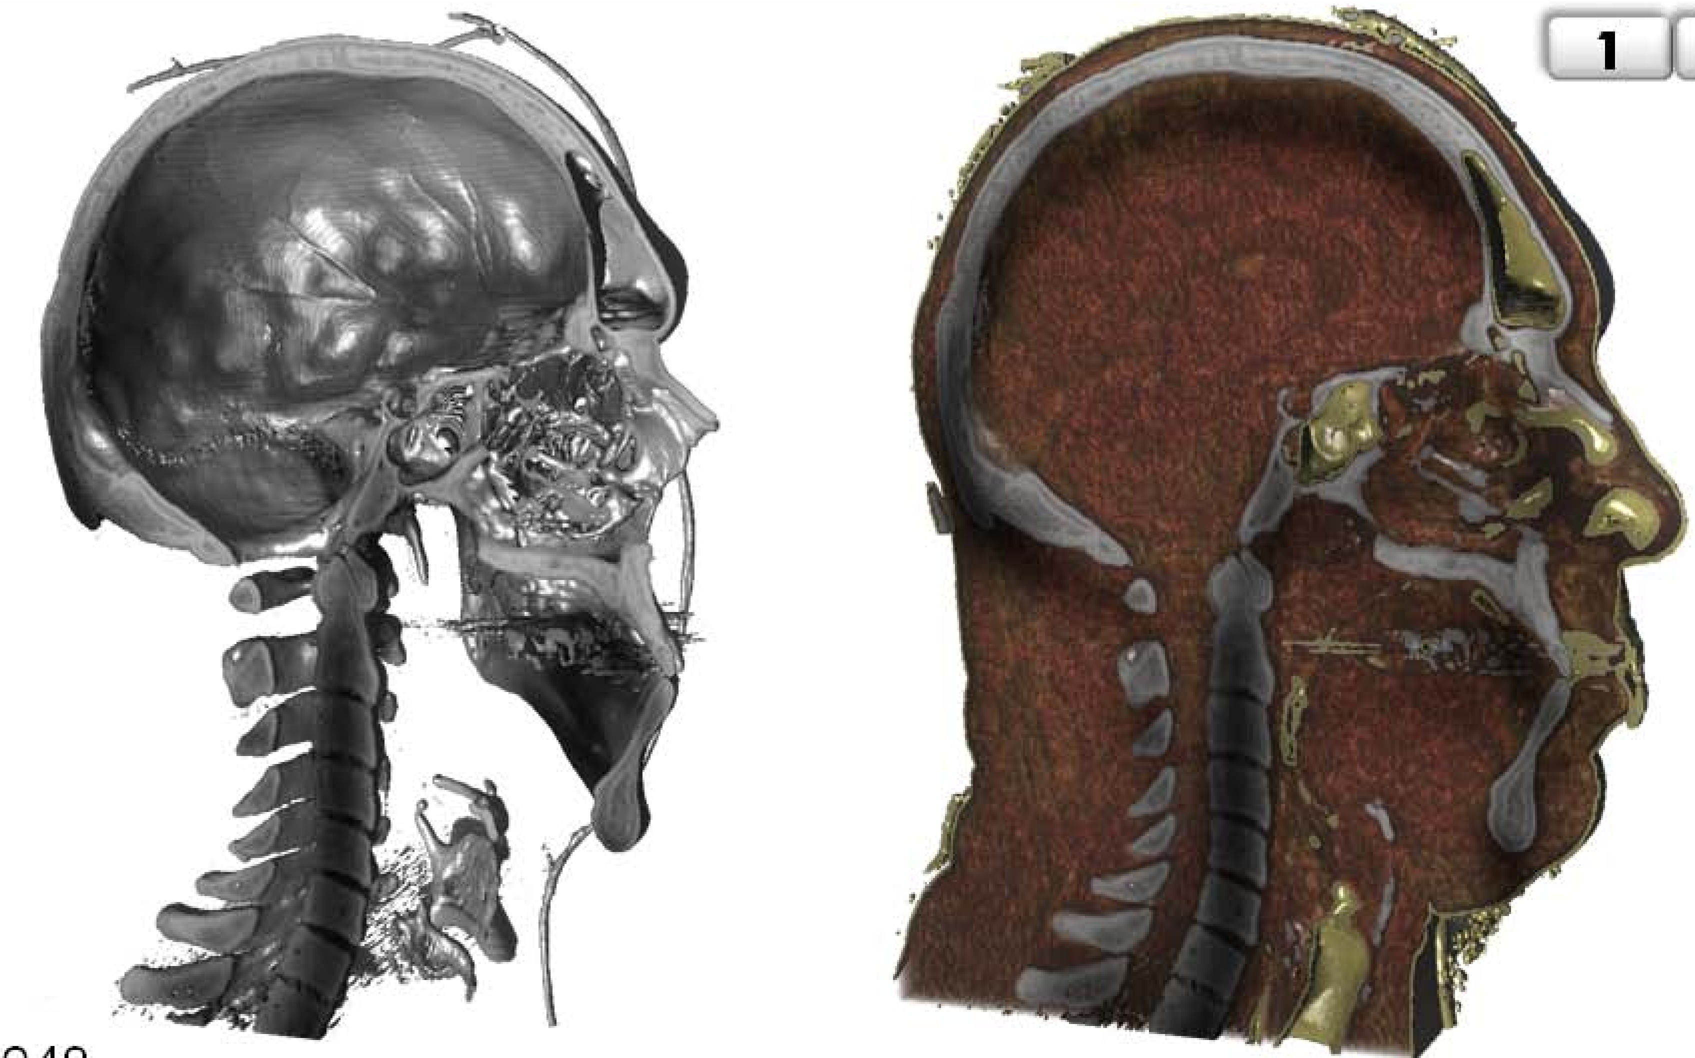

n_c=2048

Application Example - DVR

1 2 3

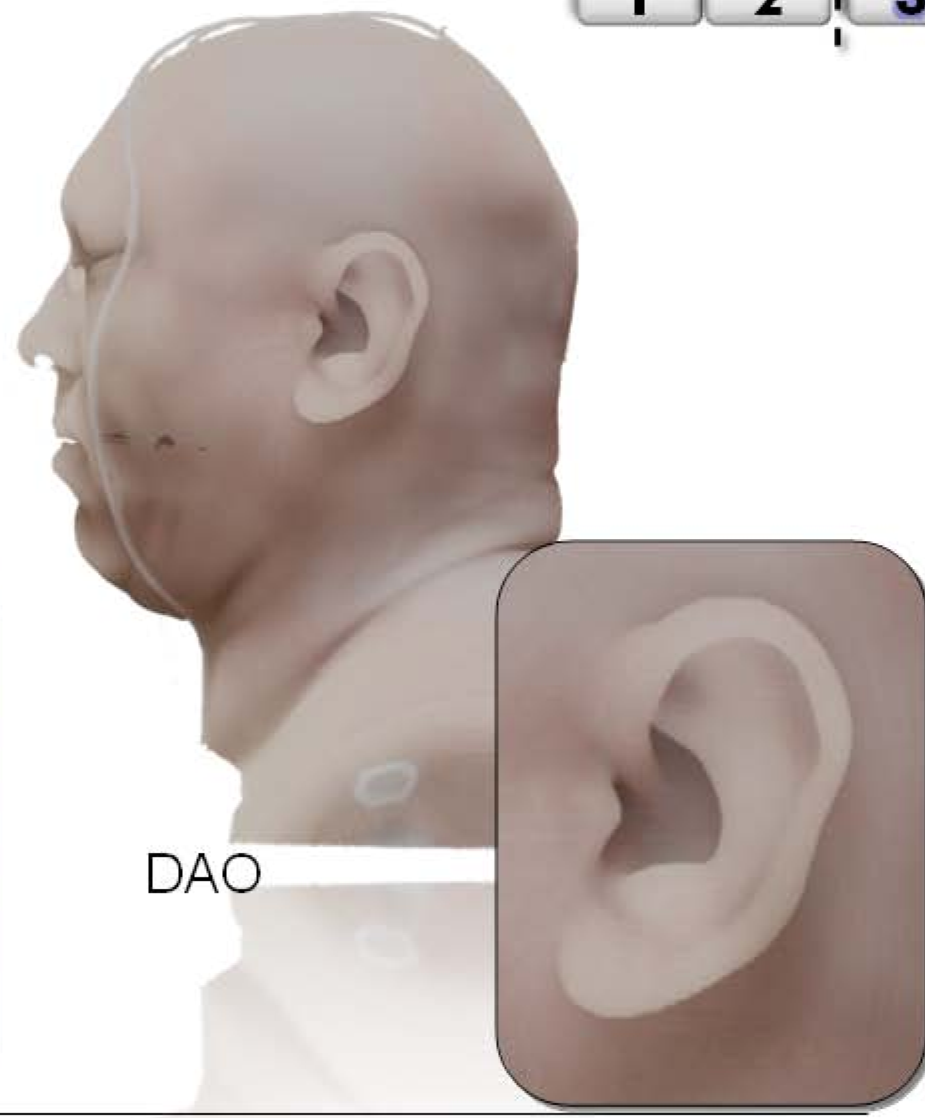

n_c=2048

DAO Performance

1 2 3

data set	size (voxel^3)	sphere size (radius)	hist. gen. (min.)	training parameters (codewords/packed dim)	training (min.)	local histo–gram size	codebook size
Cornell box	$128 \times 128 \times 128$	32	236.03	256/16	2.10	2048 MB	0.01 MB
				512/16	7.08		
				1024/16	21.90		
Visible Human head	$192 \times 192 \times 110$	12	16.63	2048/64	528.58	3960 MB	0.5 MB
		16	38.71		484.31		
		24	132.40		510.01		
	$256 \times 256 \times 147$	16	101.60	2048/64	704.21	9408 MB	0.5 MB
	$512 \times 512 \times 294$	32	3025.38		–[a]	75264 MB	0.5 MB
hand	$244 \times 124 \times 257$	20	514.31[b]	2048/64	633.80	7593 MB	0.5 MB
feet[b]	$128 \times 64 \times 128$	12	15.98	2048/64	320.81	1024 MB	0.5 MB
data sets computed using pre-processing with performance improved implementation							
cloud	$256 \times 128 \times 128$	24	6.03	2048/16	5.03	4096 MB	0.5 MB
Visible Human head	$256 \times 256 \times 147$	16	10.26	2048/64	61.60	9408 MB	0.5 MB

[a] not calculated. [b] calculated on a machine with $4 \times$ Xeon 2.8.

Local Ambient Occlusion

- **Hernell et al., Eurographics/IEEE VG 2007**


Diffuse Surface Shading


Local Ambient Occlusion

Volumetric Ambient Occlusion

- Volumetric Object
  - Semi-transparent tissues
- Global Light Source & Isotropic Incident Light

Volumetric Ambient Occlusion

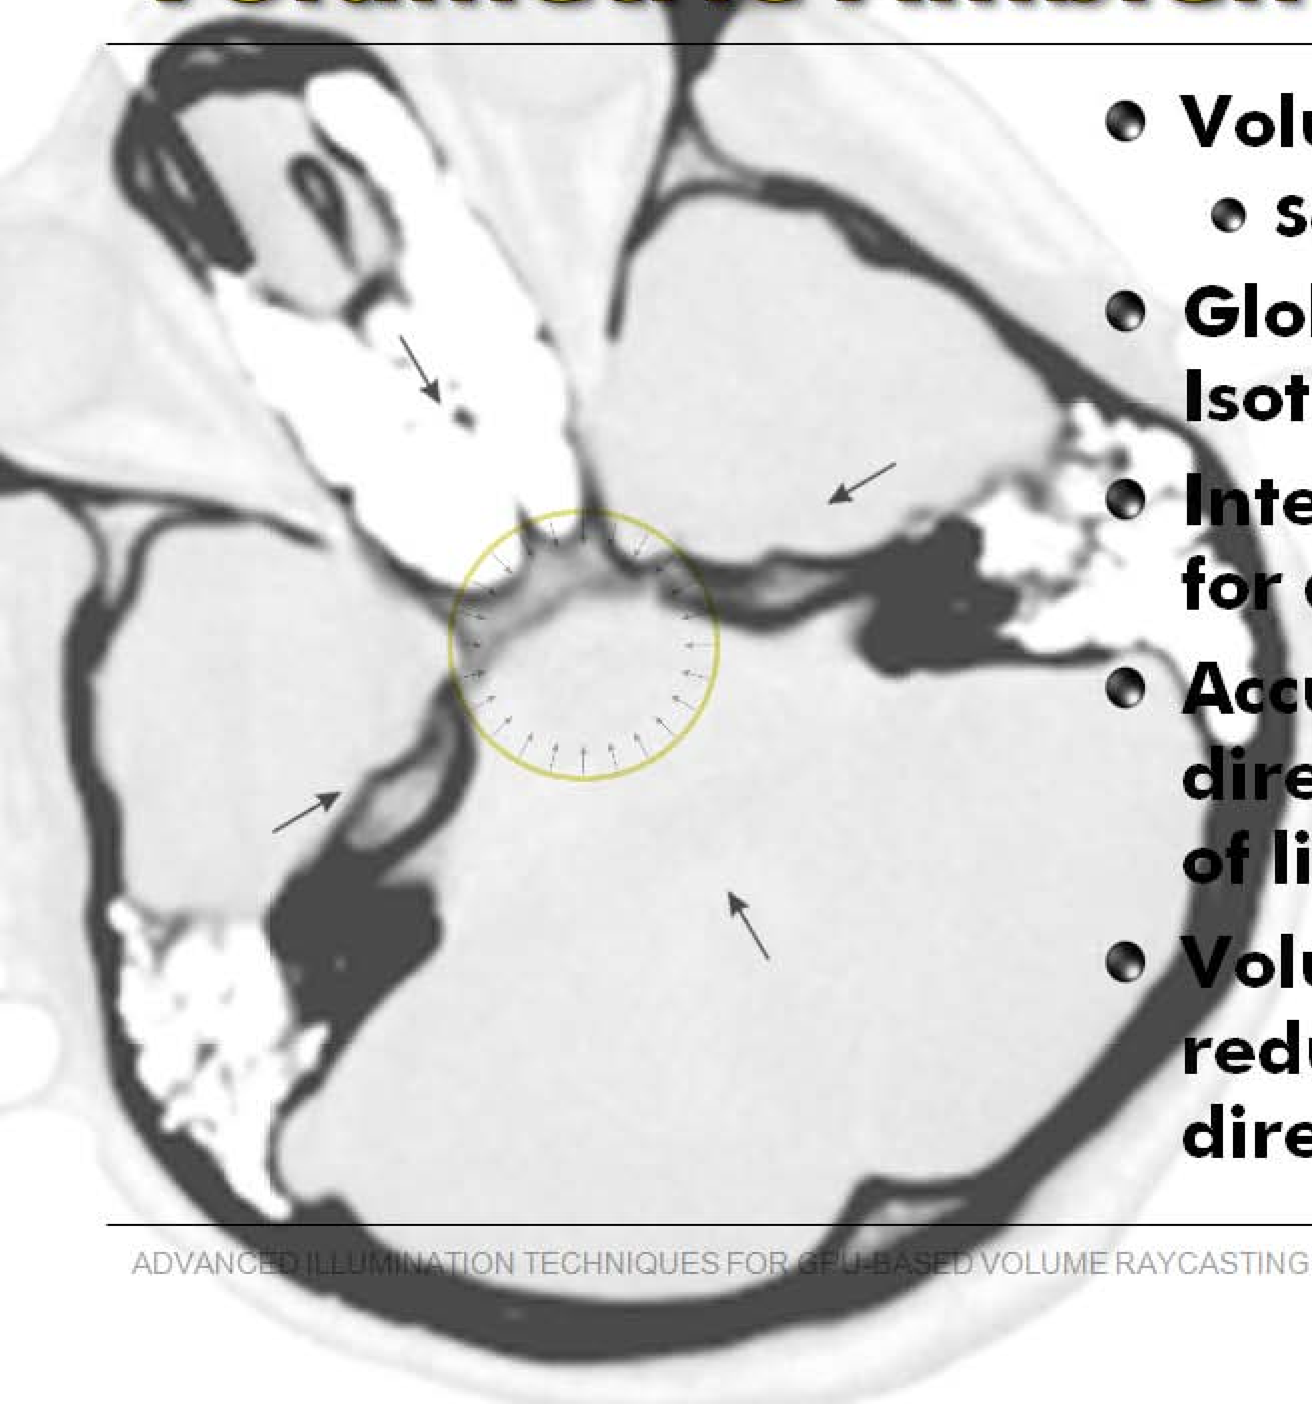

- Volumetric Object
  - Semi-transparent tissues
- Global Light Source & Isotropic Incident Light
- Integrate Illumination for a point in the volume
- Accuracy: number of directions and resolution of light “cache”
- Volumetric Light Sphere reduce required directions

Additional Features

- Exploits Multiresolution LOD data
  - Improves performance
  - Reduces memory requirement
  - No precomputation
- Interactive TF-based light emission
  - Highlighting user-specific data ranges

LAO – Single Direction

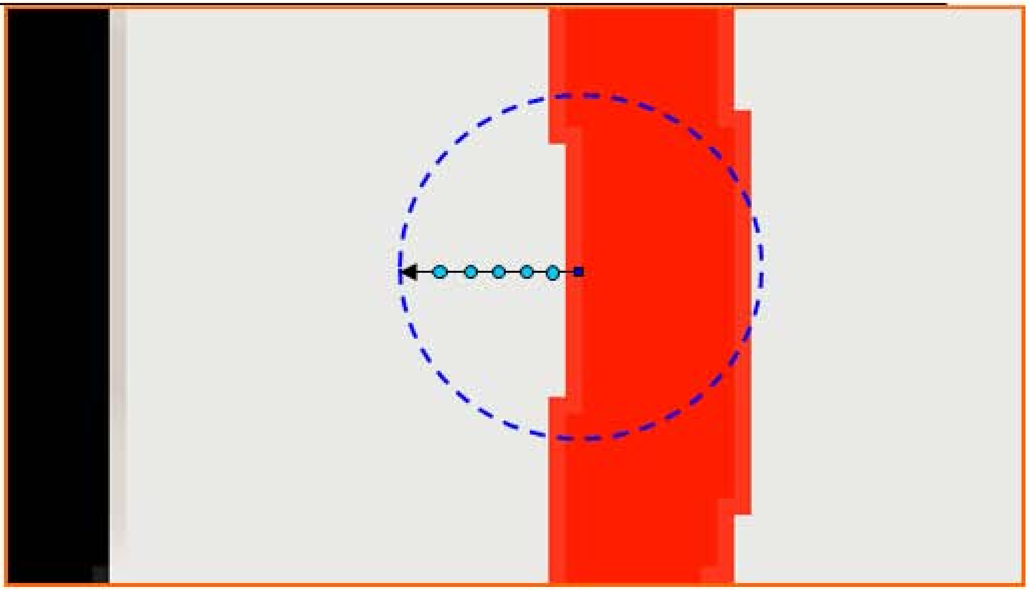

- Spherical neighborhood, Ω
- Estimate attenuation of light
- Semi-transparency

LAO – Single Direction

LAO – Multiple Directions

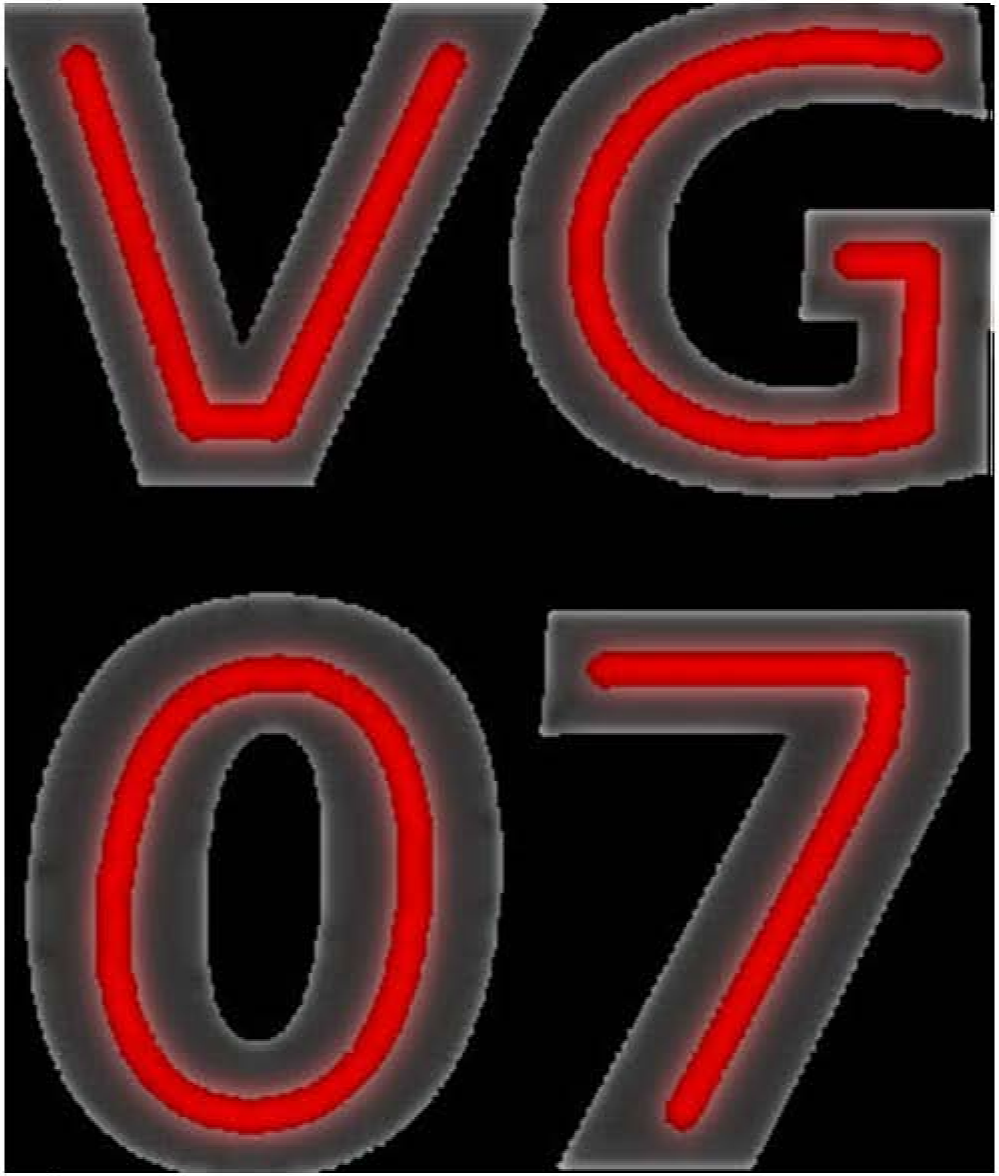

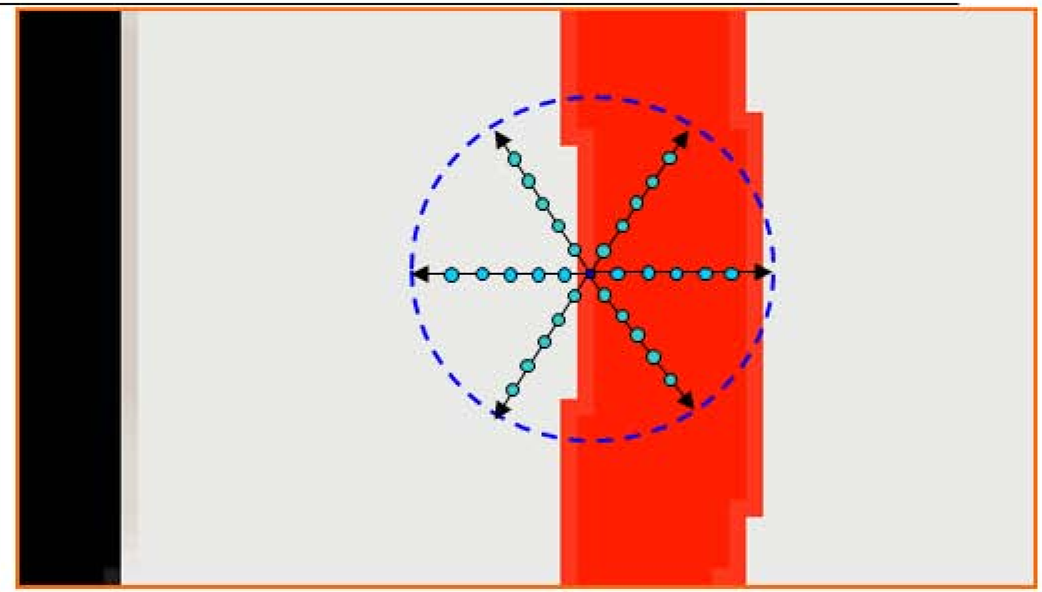

- Compute LAO in K number of directions
- Updated incrementally
  - one ray per frame

LAO Integration

- **Integration of irradiance from one direction**

$$I_k(x) = \int_a^{R_\Omega} \frac{1}{R_\Omega - a} \exp(-\int_a^s \tau(u)du)ds$$

- **Numerical evaluation**

$$I_k(x) = \sum_{m=0}^{M} \frac{1}{M} \prod_{i=0}^{m-1} (1 - \alpha_i)$$

- **Estimating light from K number of directions**

$$I(x) = I_{\text{bias}} + \frac{1}{K} \sum_k^K w_k I_k(x)$$

Volumetric Light Contribution

$$I_k(x) = \sum_{m=0}^{M} \boxed{\frac{1}{M}} \prod_{i=0}^{m-1} (1 - \alpha_i)$$

Boundary Only Contribution

Volumetric Light Contribution

LAO with Emission

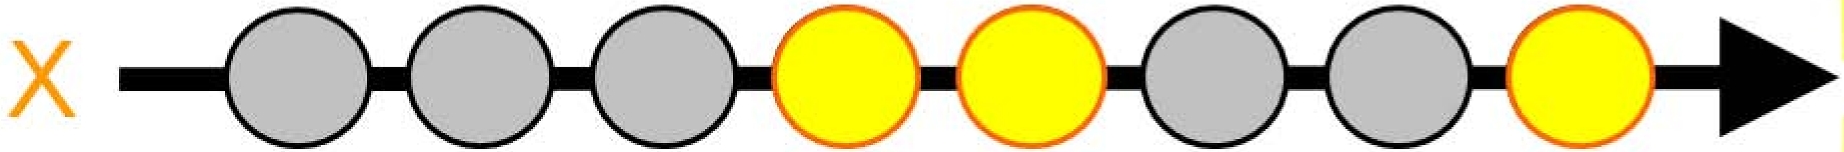

- Affect intensity and color of x
- Emissive component in the TF
- Add C_E (color light emission) to the integral

$$I_k(x) = \int_a^{R_\Omega} \frac{1}{R_\Omega - a} \exp\left(-\int_a^s \tau(u)\,du\right) ds$$

LAO with Emission

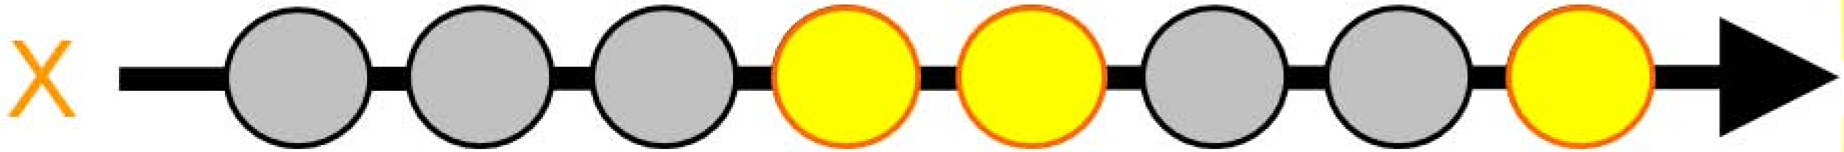

- Affect intensity and color of x
- Emissive component in the TF
- Add C_E (color light emission) to the integral

$$I_k(x) = \int_a^{R_\Omega} \frac{1 + c_E(s)}{R_\Omega - a} \exp\left(-\int_a^s \tau(u)\,du\right) ds$$

Multiresolution Volume

- Superfluous sampling in areas of low occlusion
- Multiresolution approach from Ljung et al., VolVis 2004

Multiresolution Volume

- Superfluous sampling in areas of low occlusion
- Multiresolution approach from Ljung et al., VolVis 2004
- TF-based LOD selection on flat blocking data structure

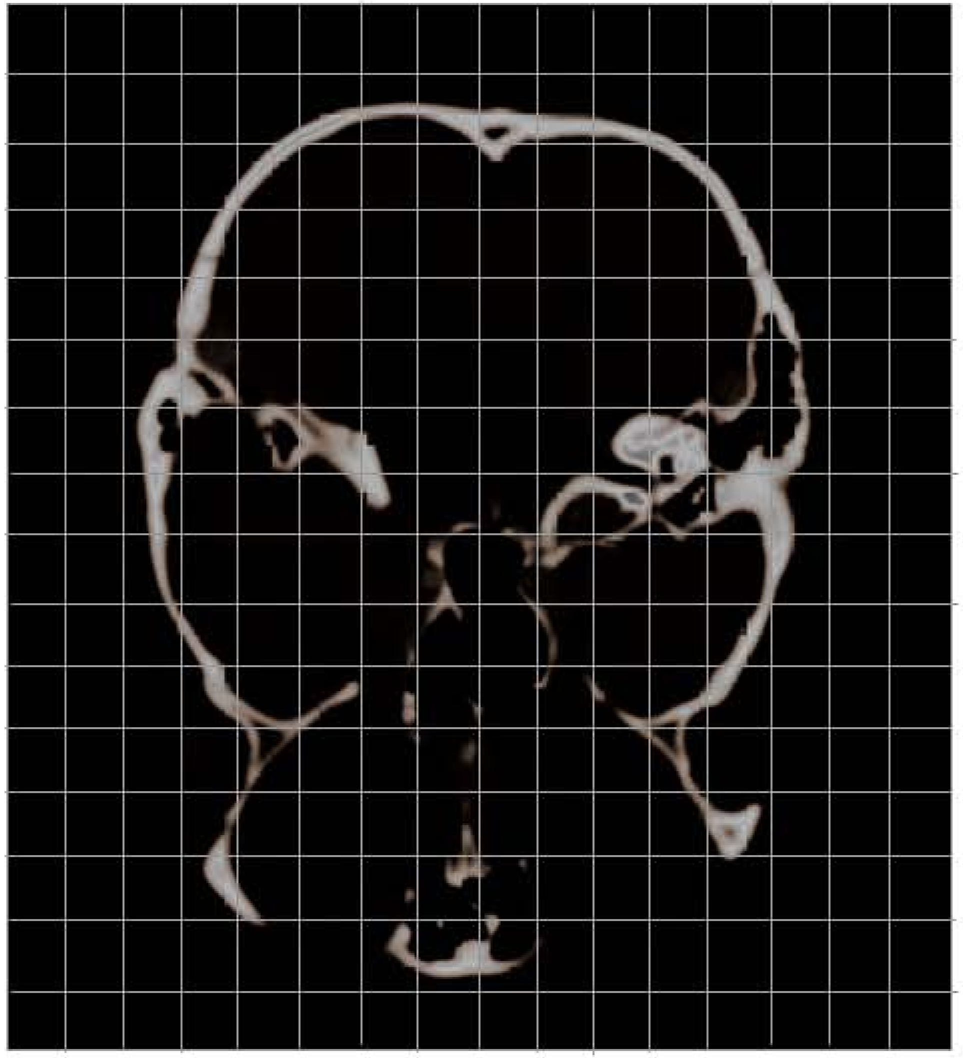

Multiresolution Volume

- Superfluous sampling in areas of low occlusion
- Multiresolution approach from Ljung et al., VolVis 2004
- TF-based LOD selection on flat blocking data structure

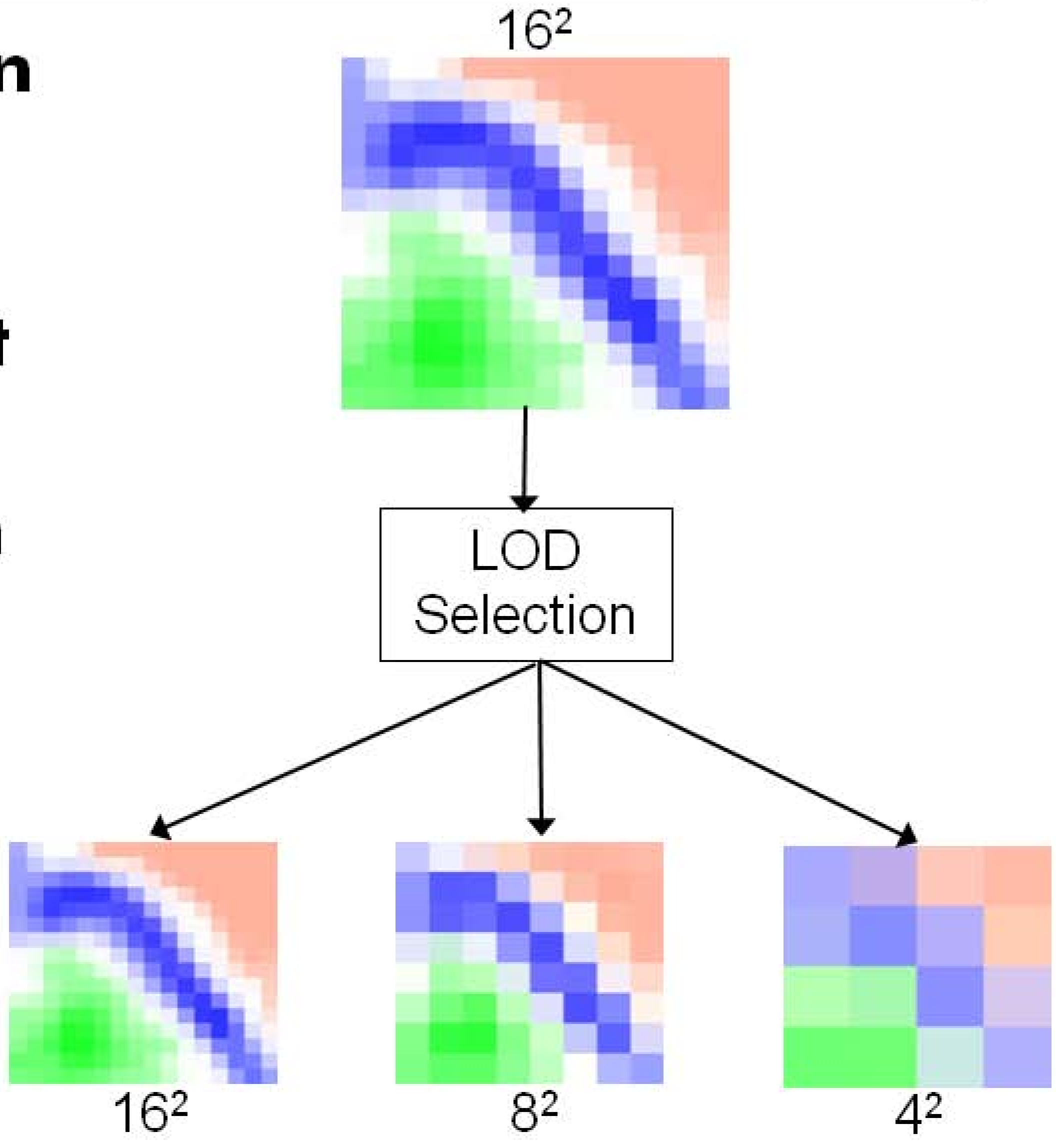

Packing LOD Volume

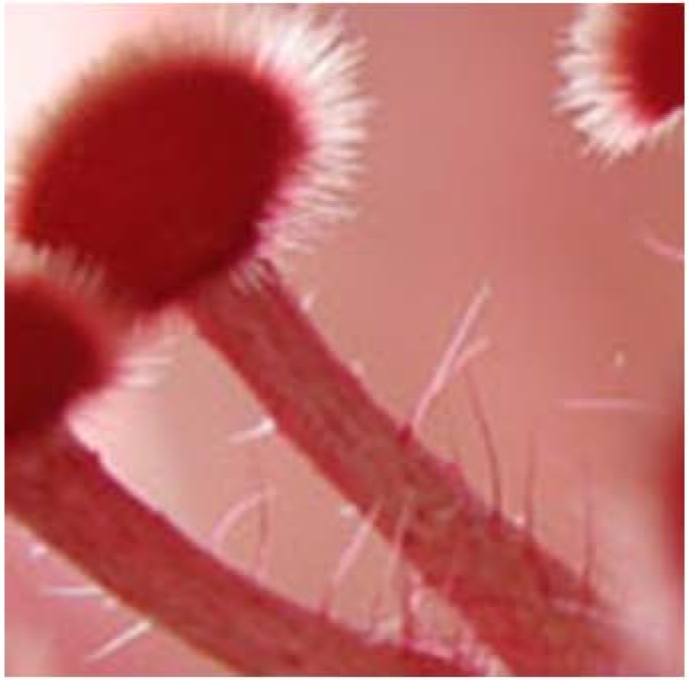

High resolution (16^2)

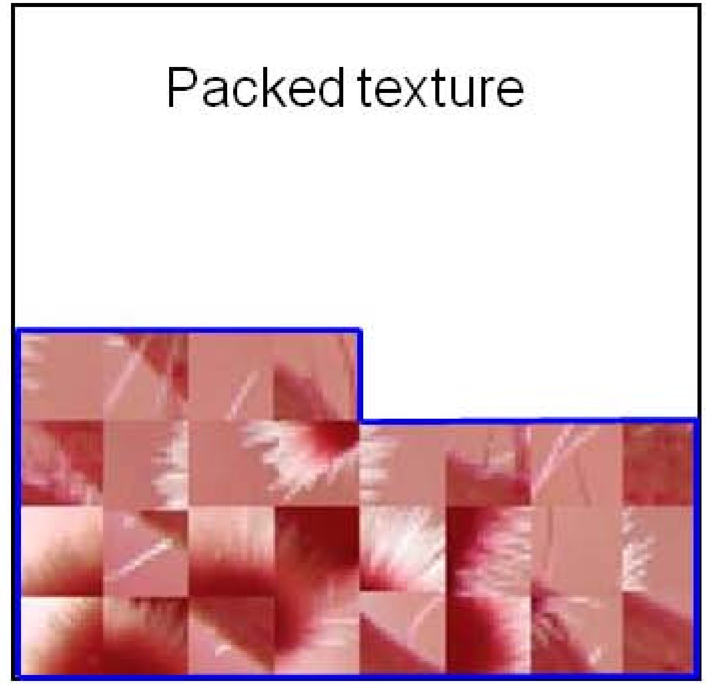

Packing LOD Volume

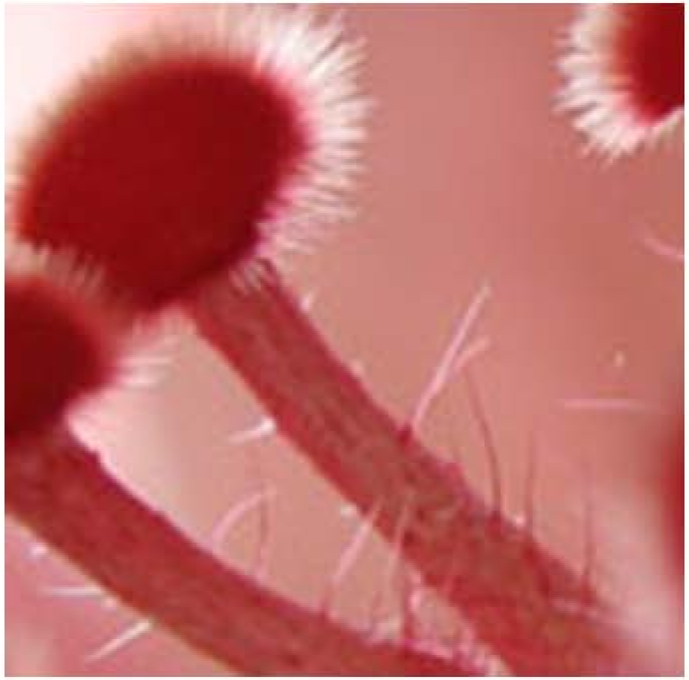

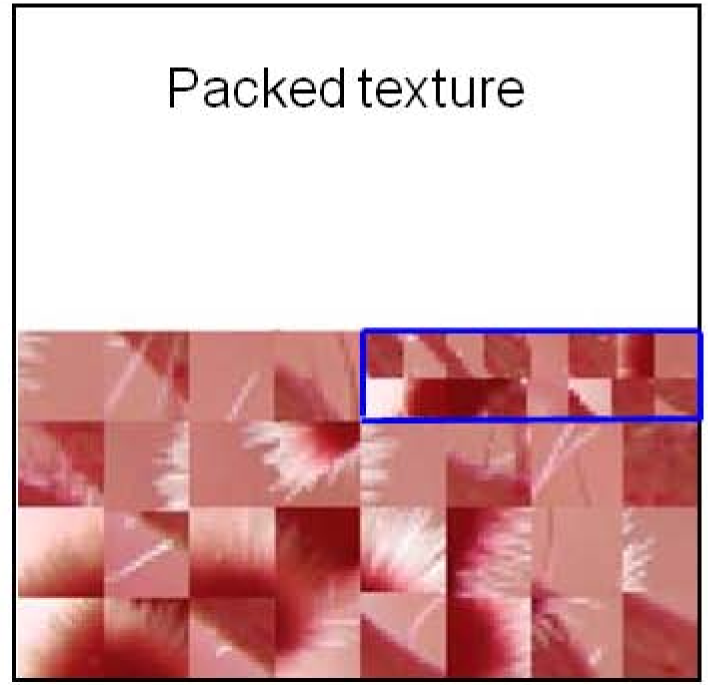

High resolution (16^2)

Middle resolution (8^2)

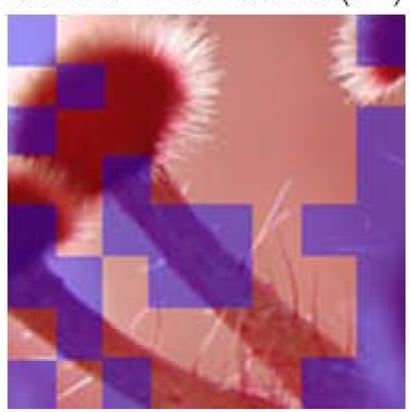

Packing LOD Volume

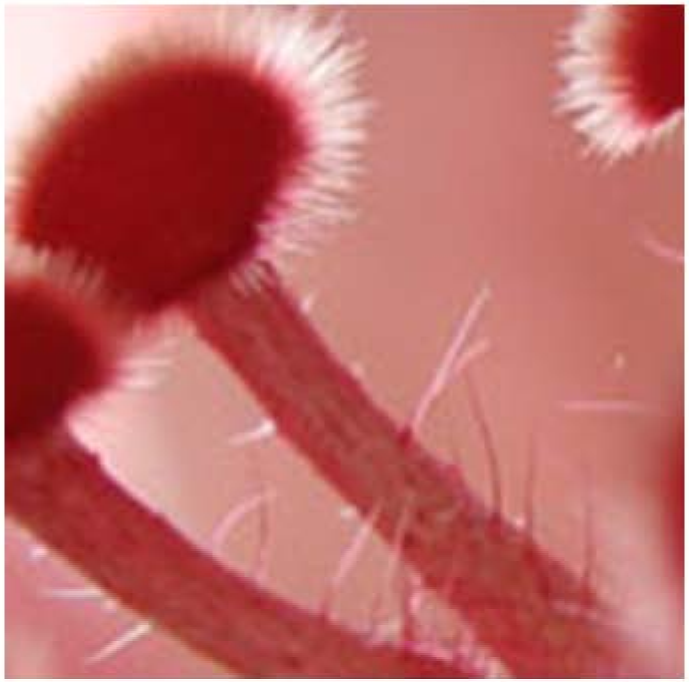

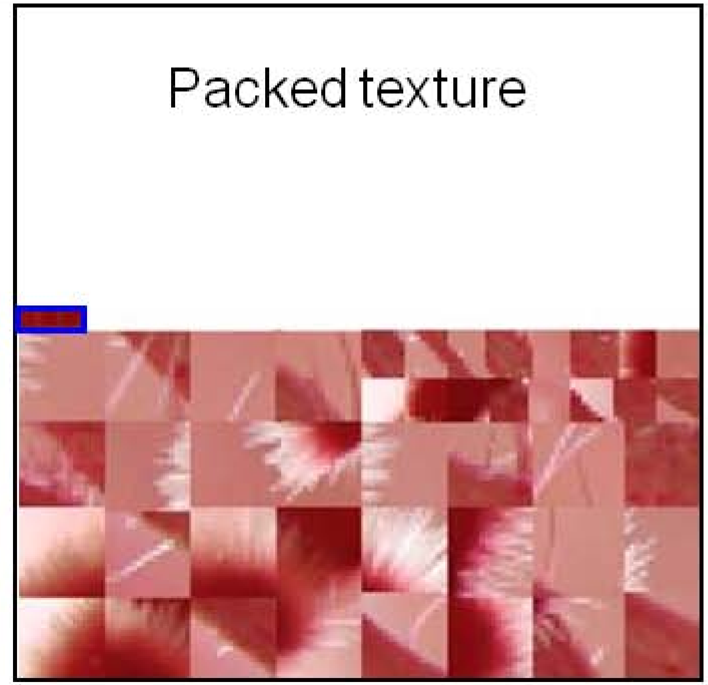

High resolution (16^2)

Middle resolution (8^2)

Low resolution (4^2)

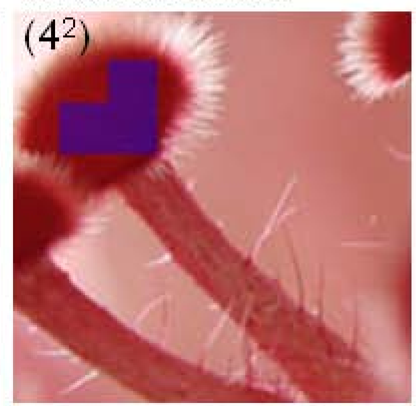

Packing LOD Volume

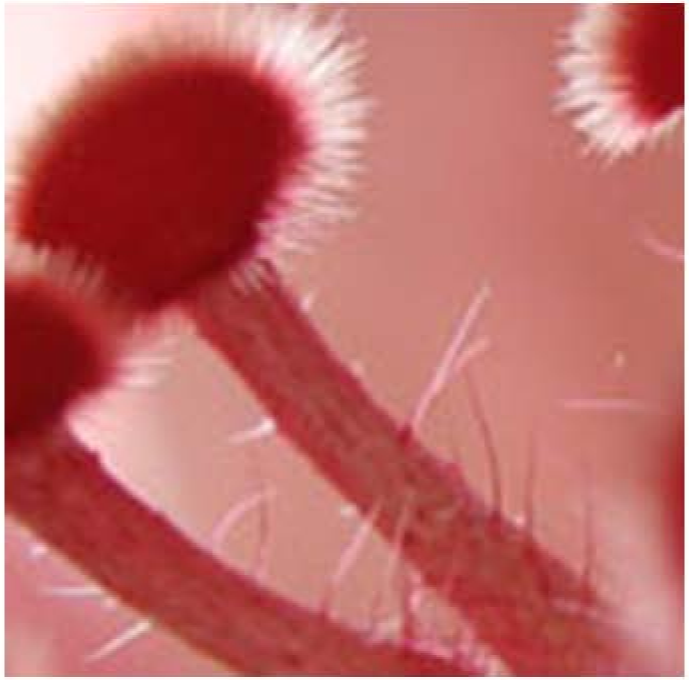

High resolution (16^2)

Middle resolution (8^2)

Low resolution (4^2)

Empty

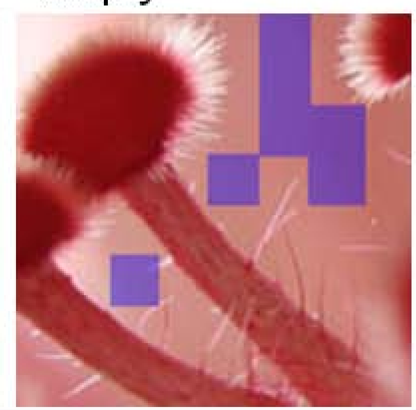

Results

Diffuse Illumination

LAO with 1 direction

LAO with 8 directions

- NVIDIA GF8800 Ultra
- 768 MB of graphics texture memory
- 51 ms/frame to update one ray in the LAO map
- 27 FPS

Results

Diffuse Illumination

LAO with 1 direction

LAO with 8 directions

- NVIDIA GF8800 Ultra
- 768 MB of graphics texture memory
- 51 ms/frame to update one ray in the LAO map
- 27 FPS

Results

Diffuse Illumination

LAO with 1 direction

LAO with 8 directions

- NVIDIA GF8800 Ultra
- 768 MB of graphics texture memory
- 51 ms/frame to update one ray in the LAO map
- 27 FPS

Results

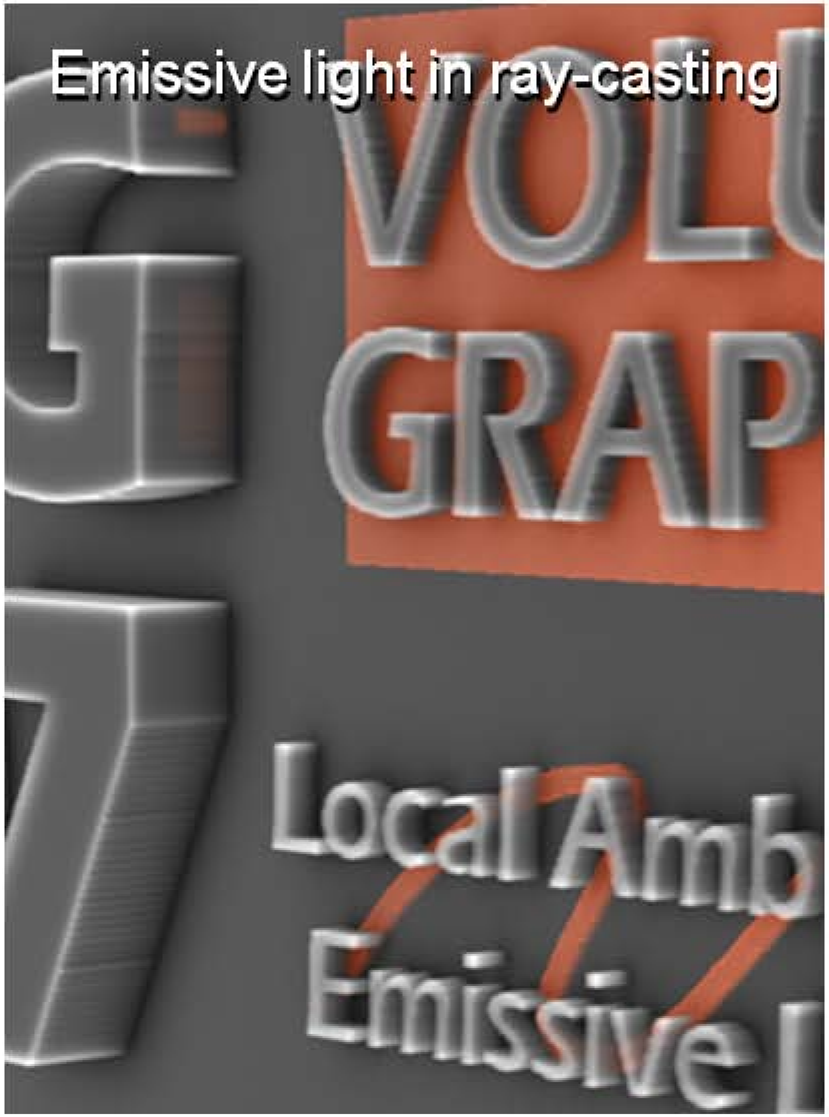
Emissive light in ray-casting

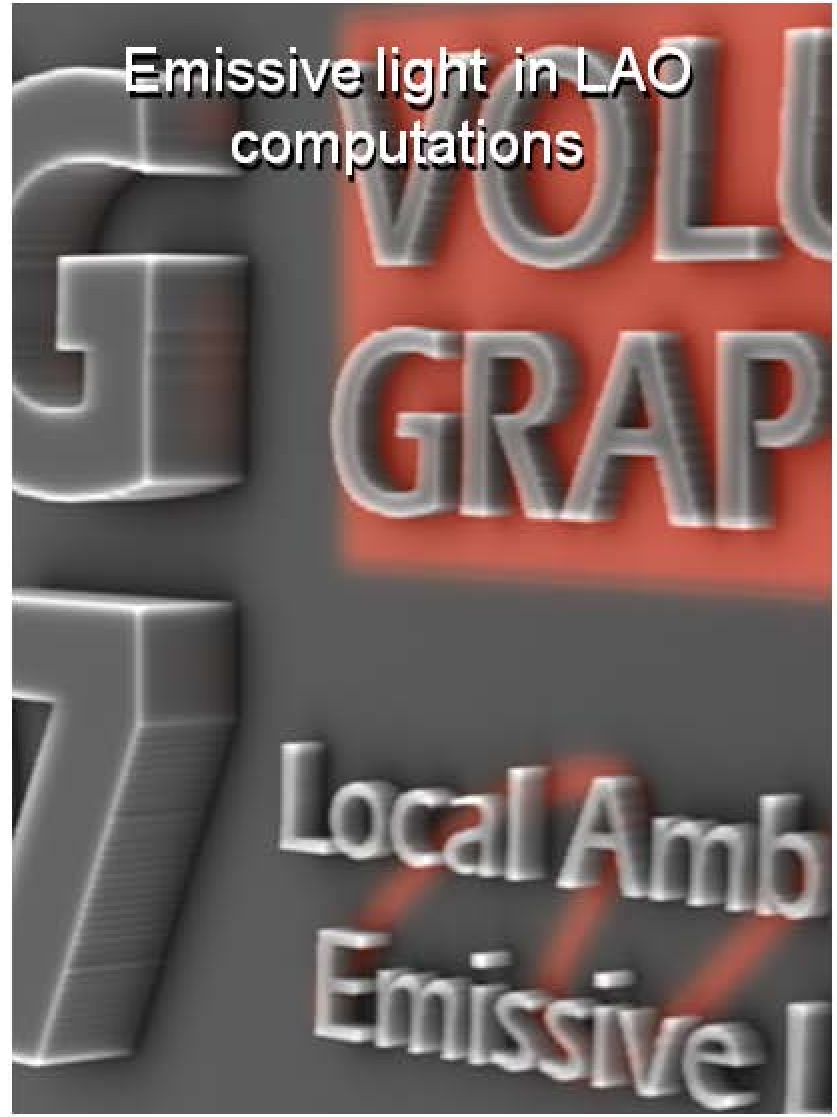
Emissive light in LAO
computations

Results

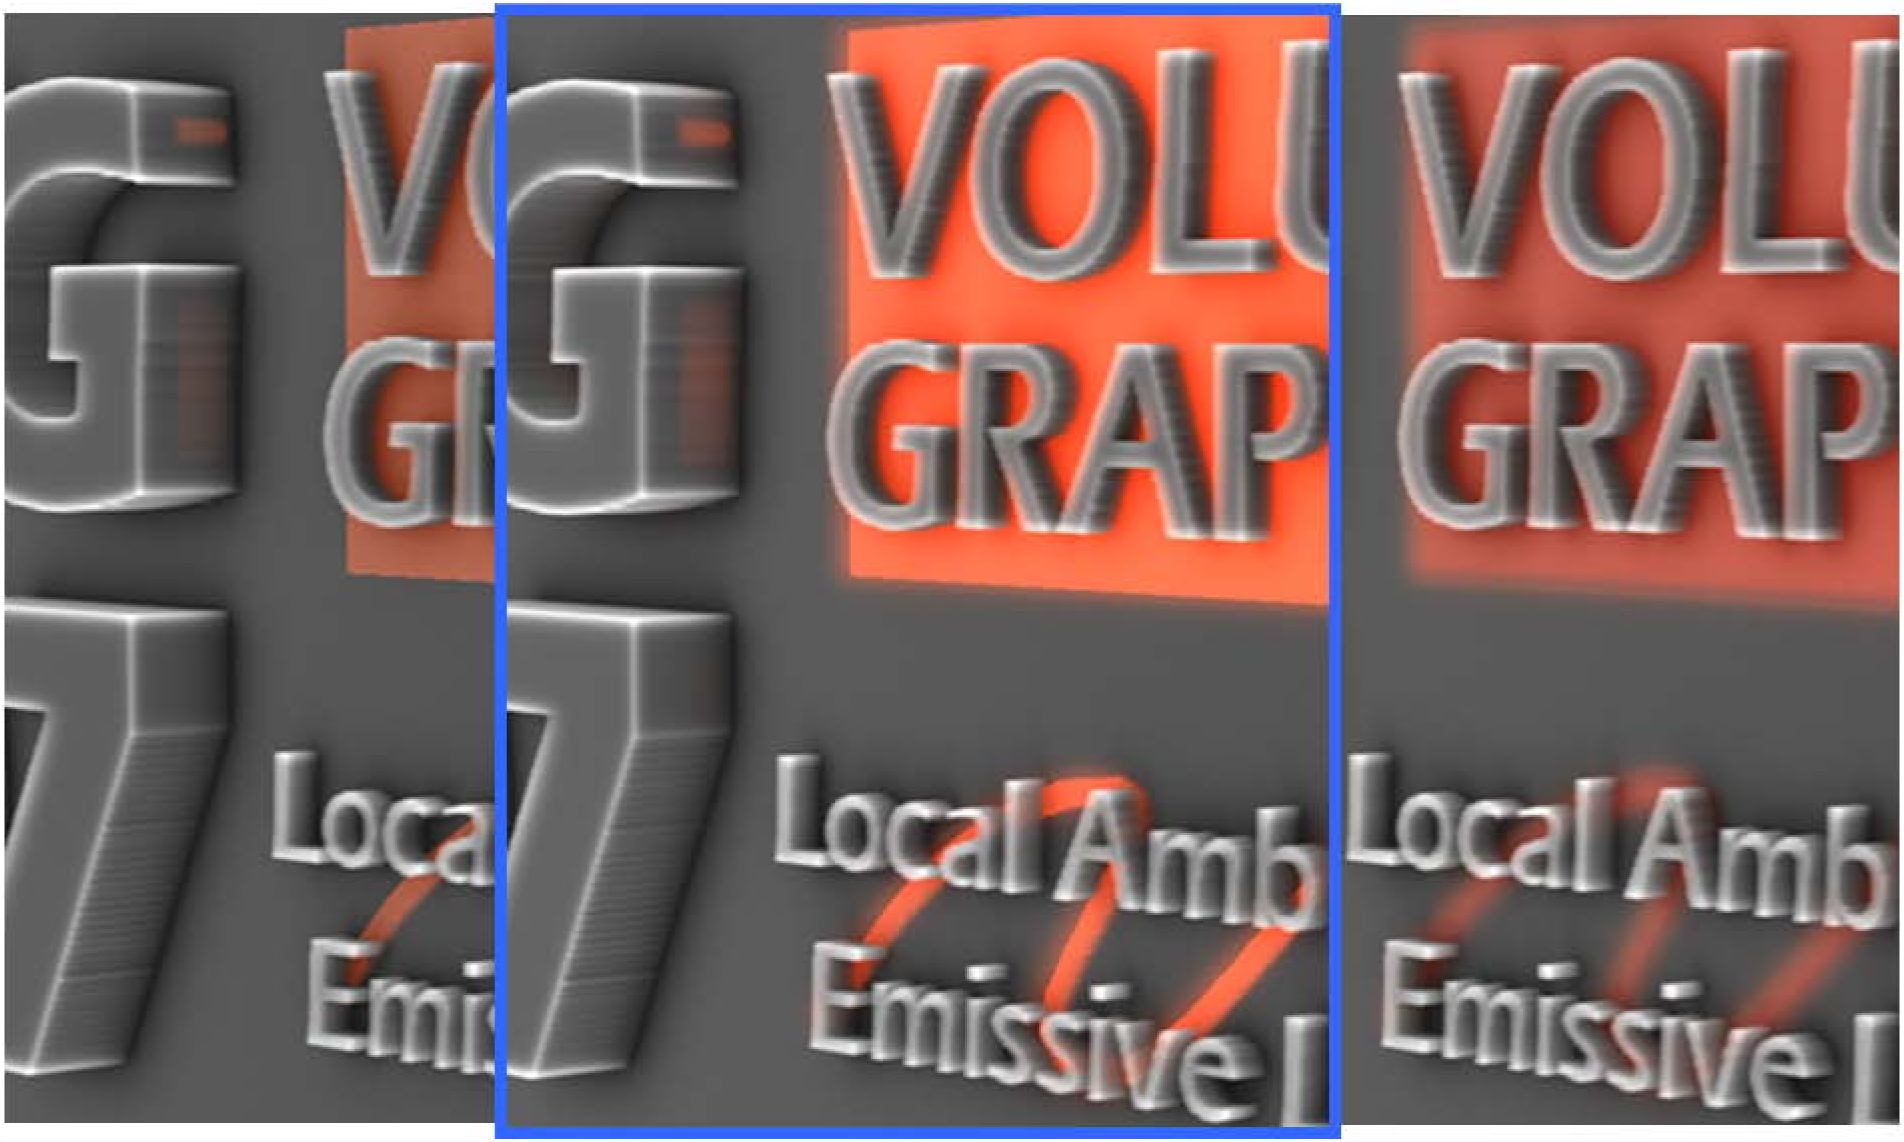

Comparison w/ Diffuse Shading

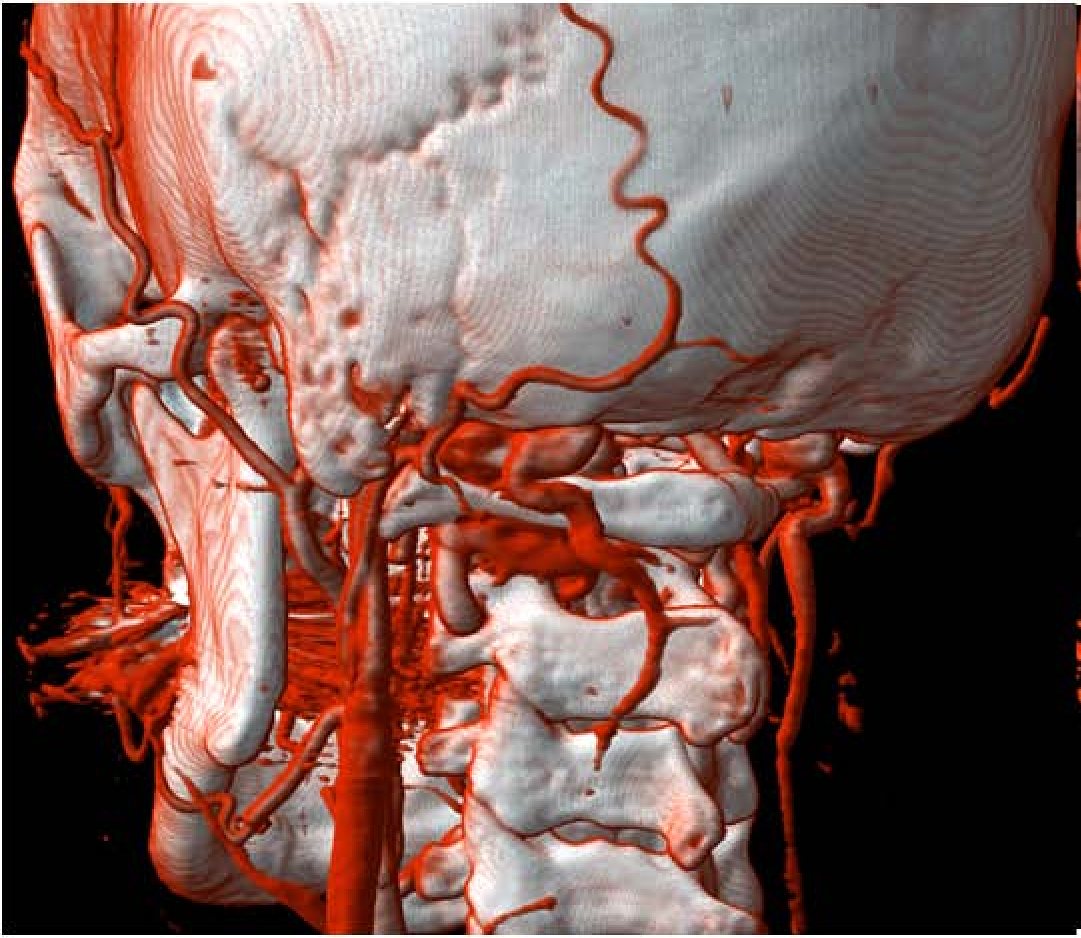

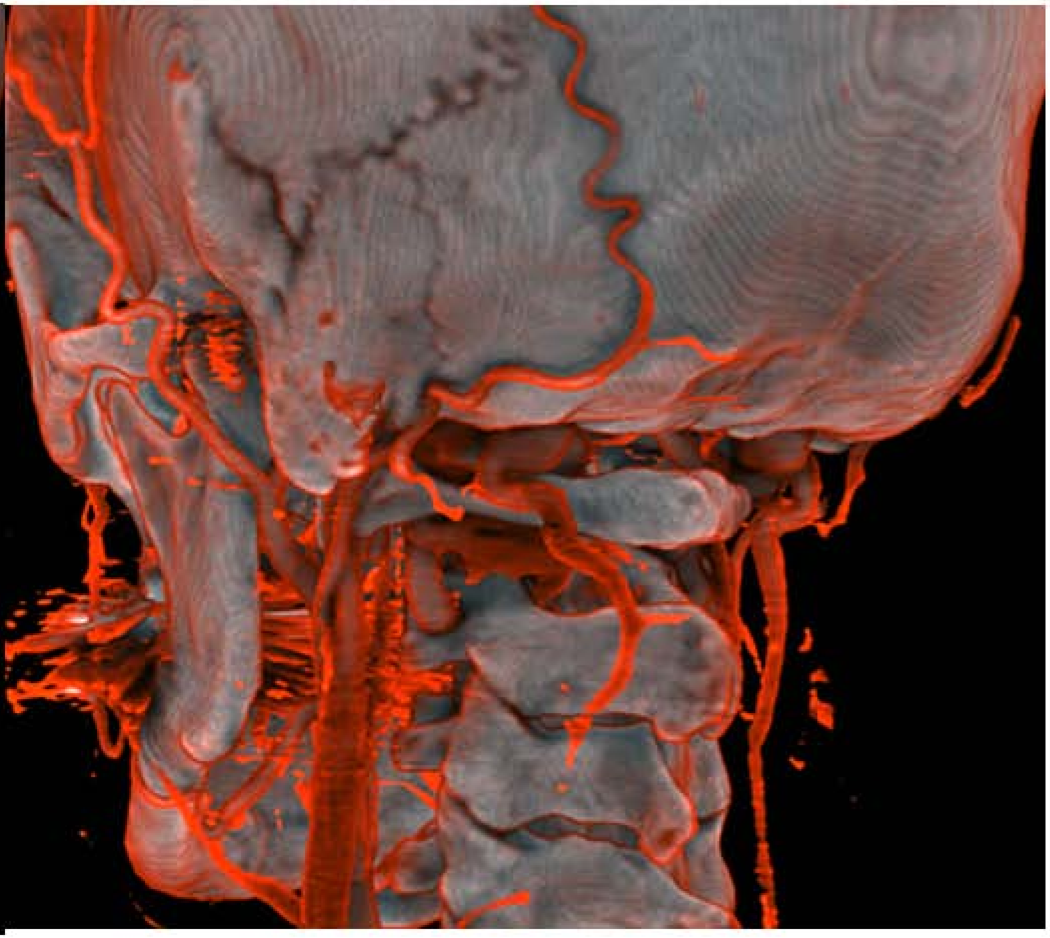

Diffuse illumination LAO, 32 rays

Comparison w/ Diffuse Shading

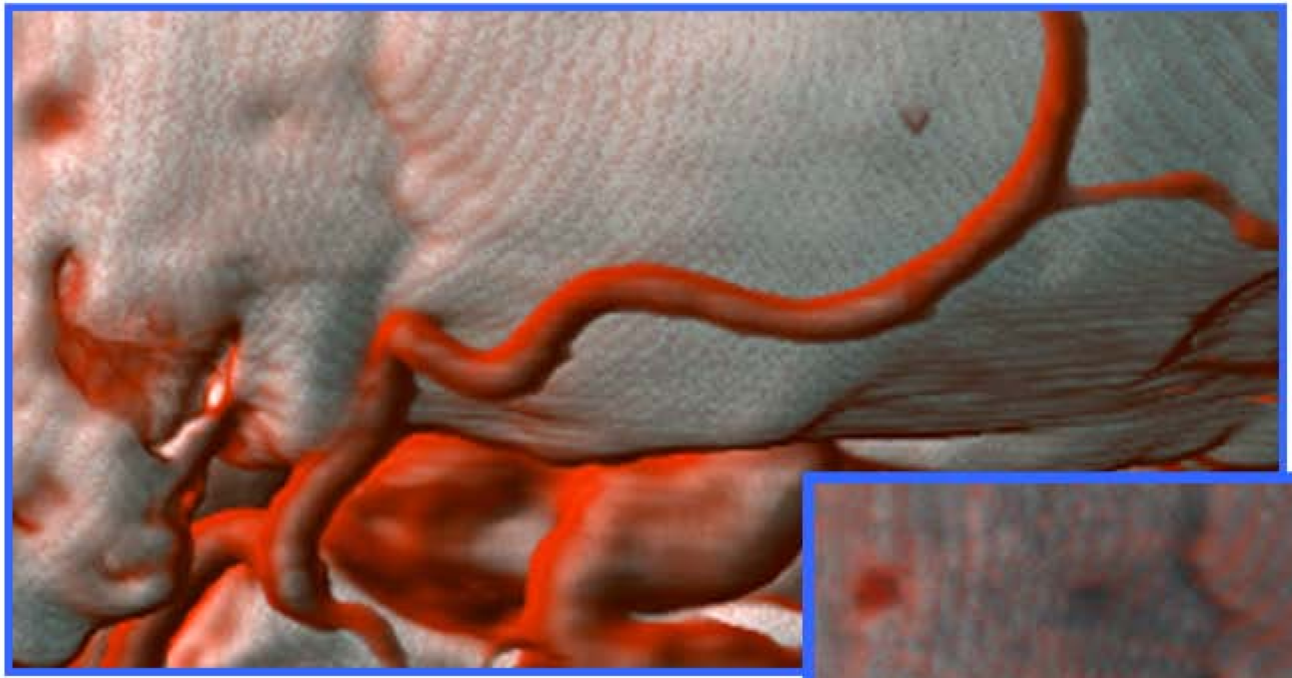

Diffuse illumination

LAO, 32 rays

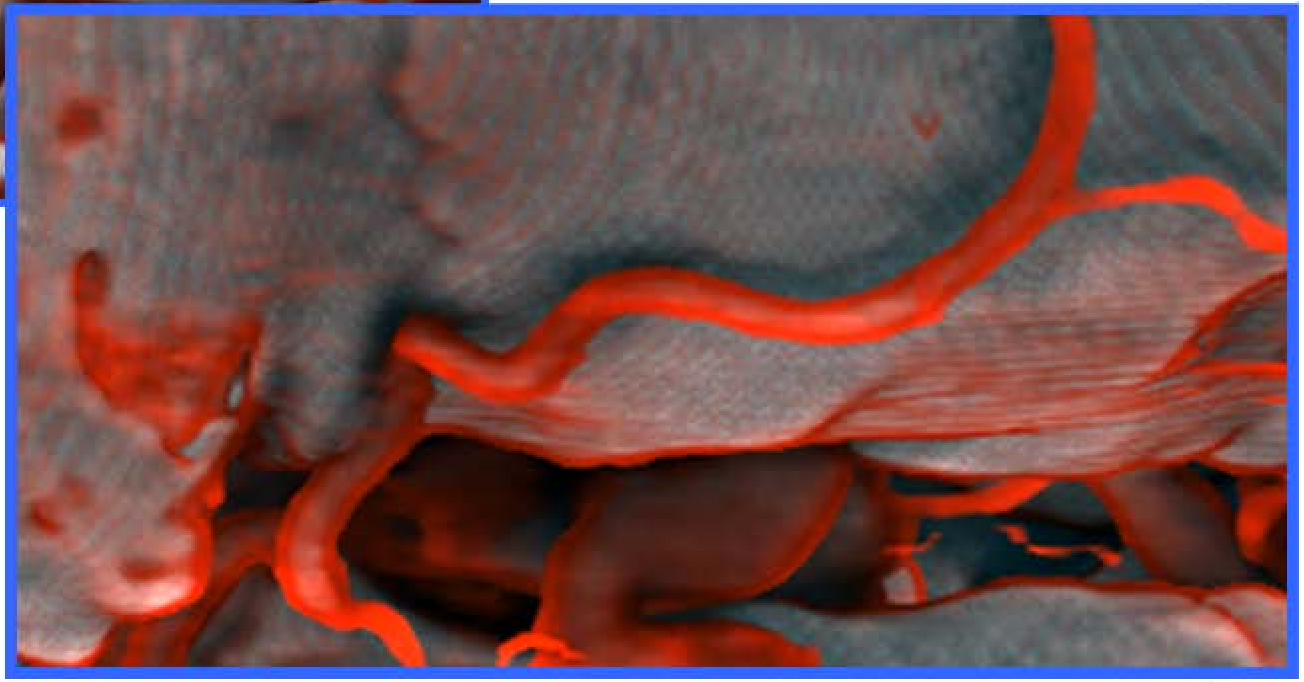

Application to Virtual Autopsy Case

(a) *Diffuse Illumination* (b) *LAO with emission* (c) *Diffuse Illumination × LAO (with emission)*

Application to Virtual Autopsy Case

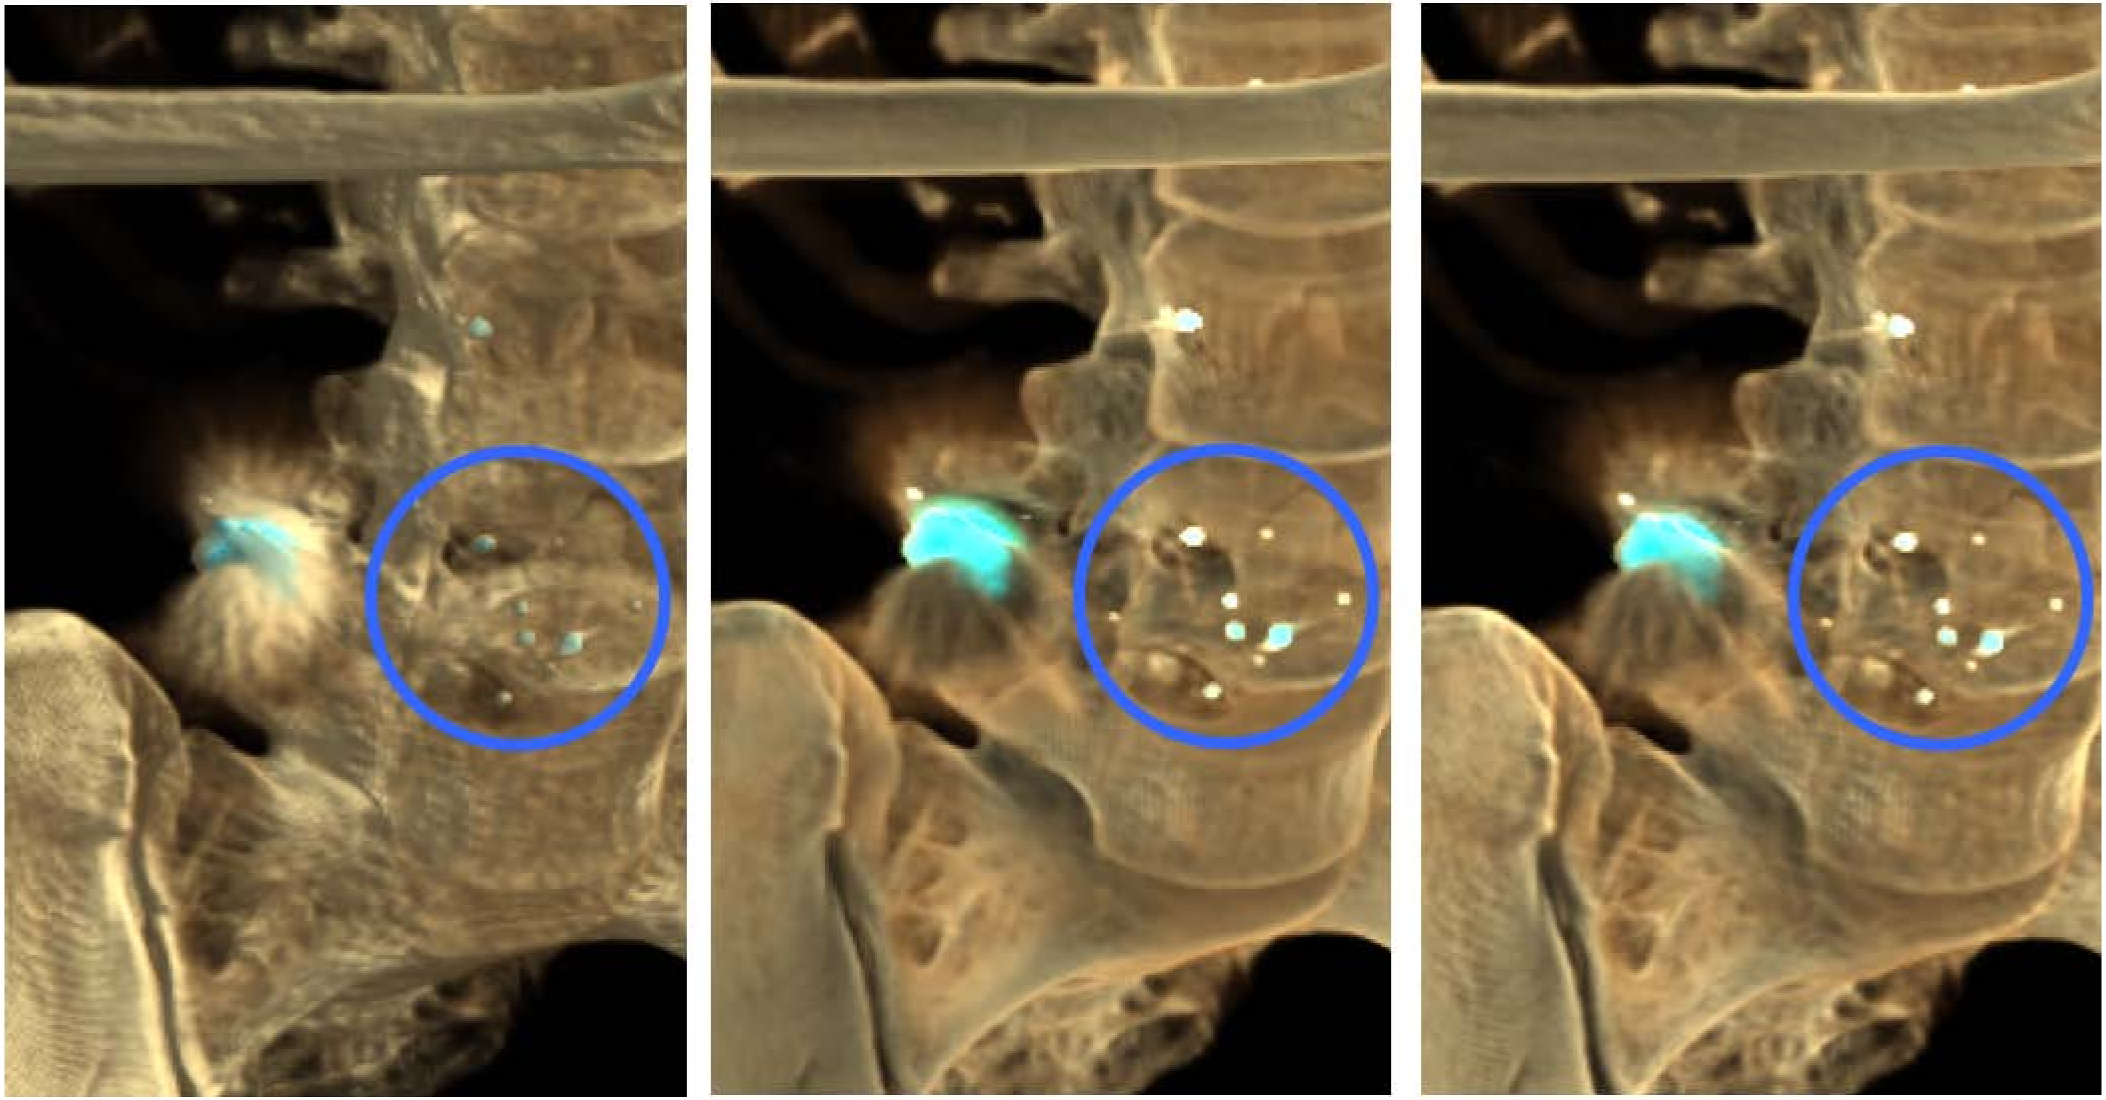

Application to Virtual Autopsy Case

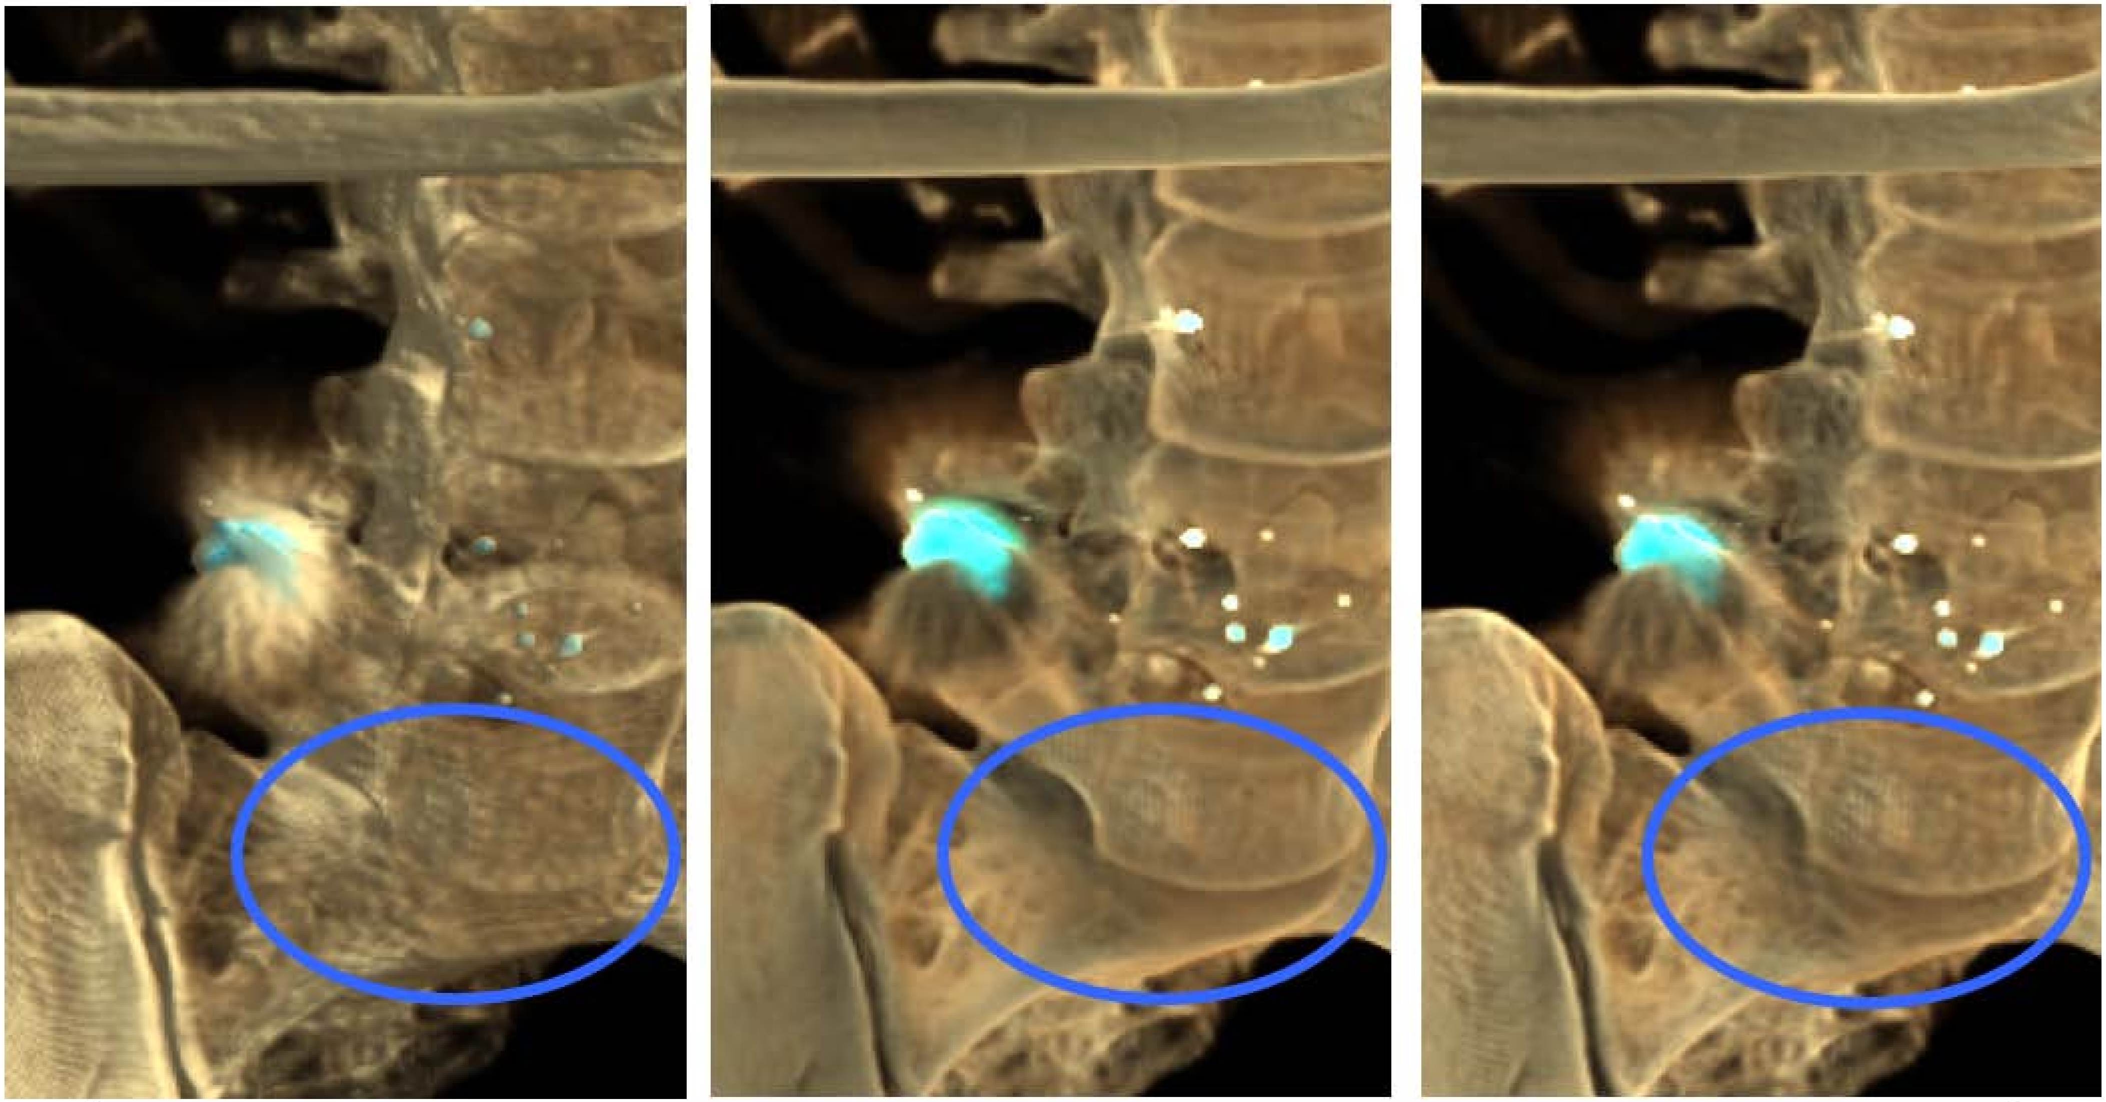

Global Light Propagation

- **Hernell et al., Volume Graphics 2008**

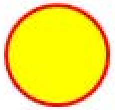

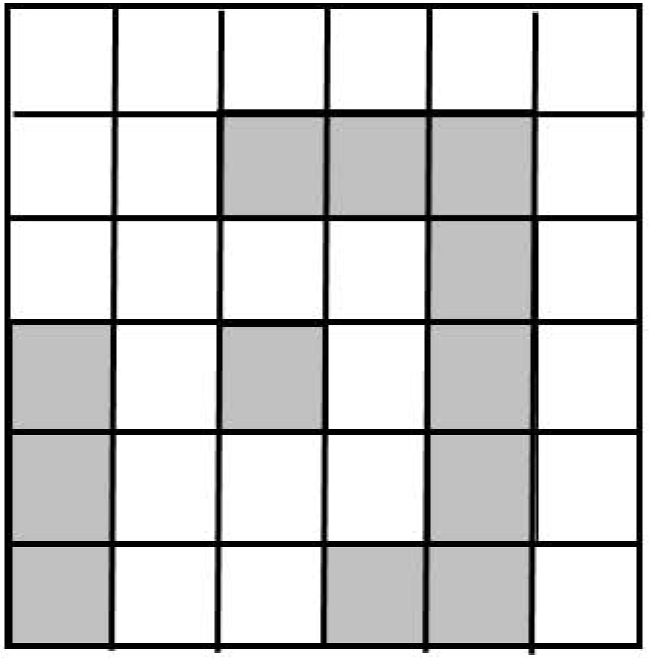

Global Light Propagation

- Hernell et al., Volume Graphics 2008
- Approximate global illumination with first order scattering
- Extends LAO with initial light propagation step
- Reuse segments of computed light attenuation

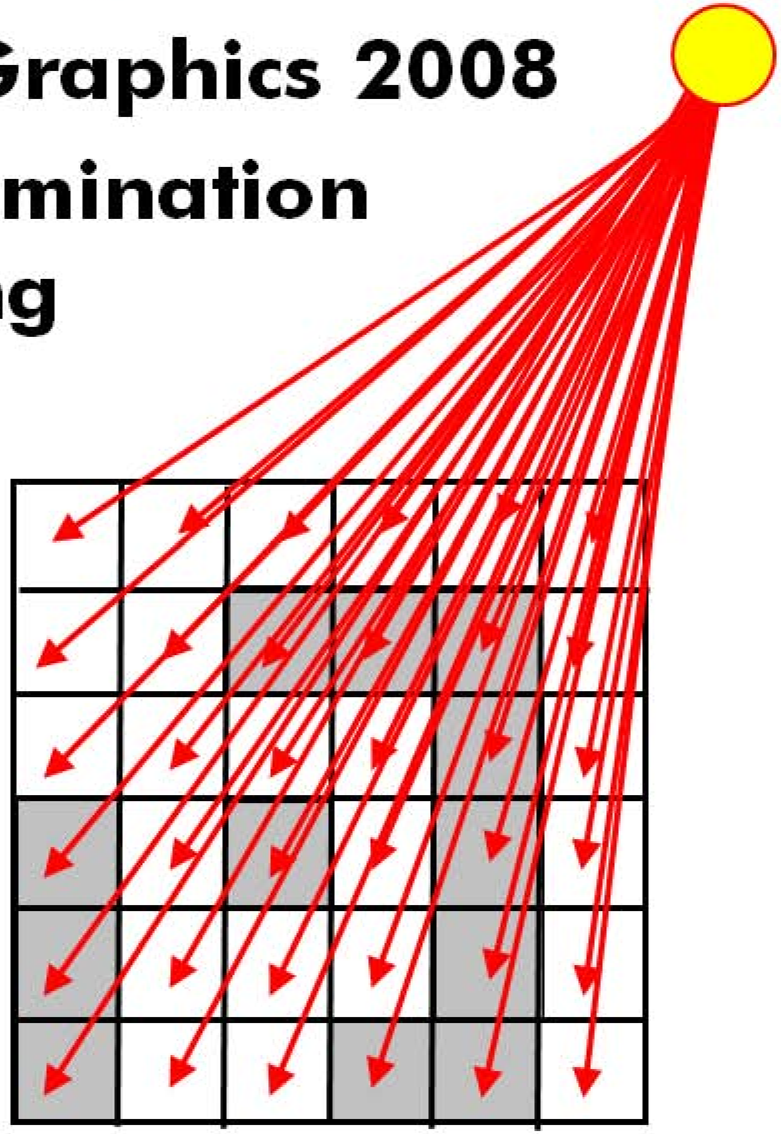

Pipeline: Algorithm Outline

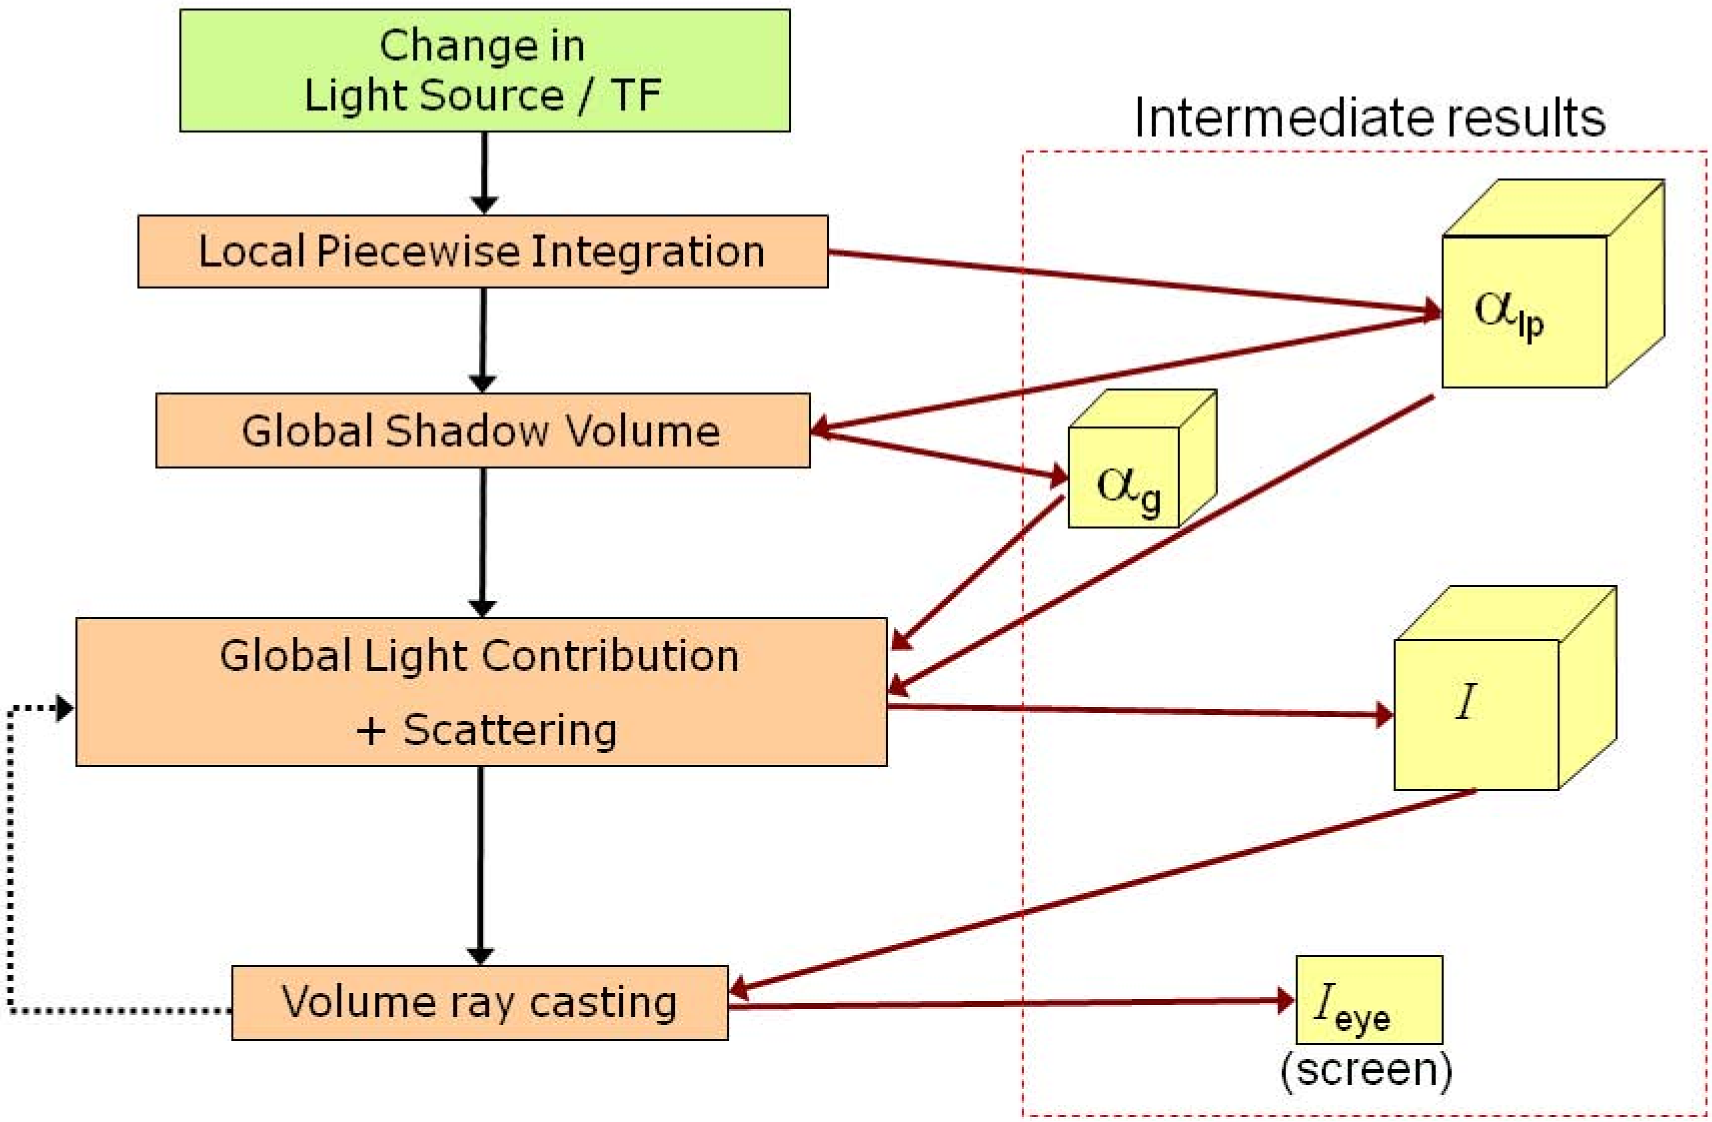

Piecewise Integration

$$T_{lp}(x) = e^{-\int_x^{x'} \tau(u)du}$$

$$\alpha_{\mathrm{lp}}(x) = 1 - \prod_{n=0}^{M}(1 - \alpha(n))$$

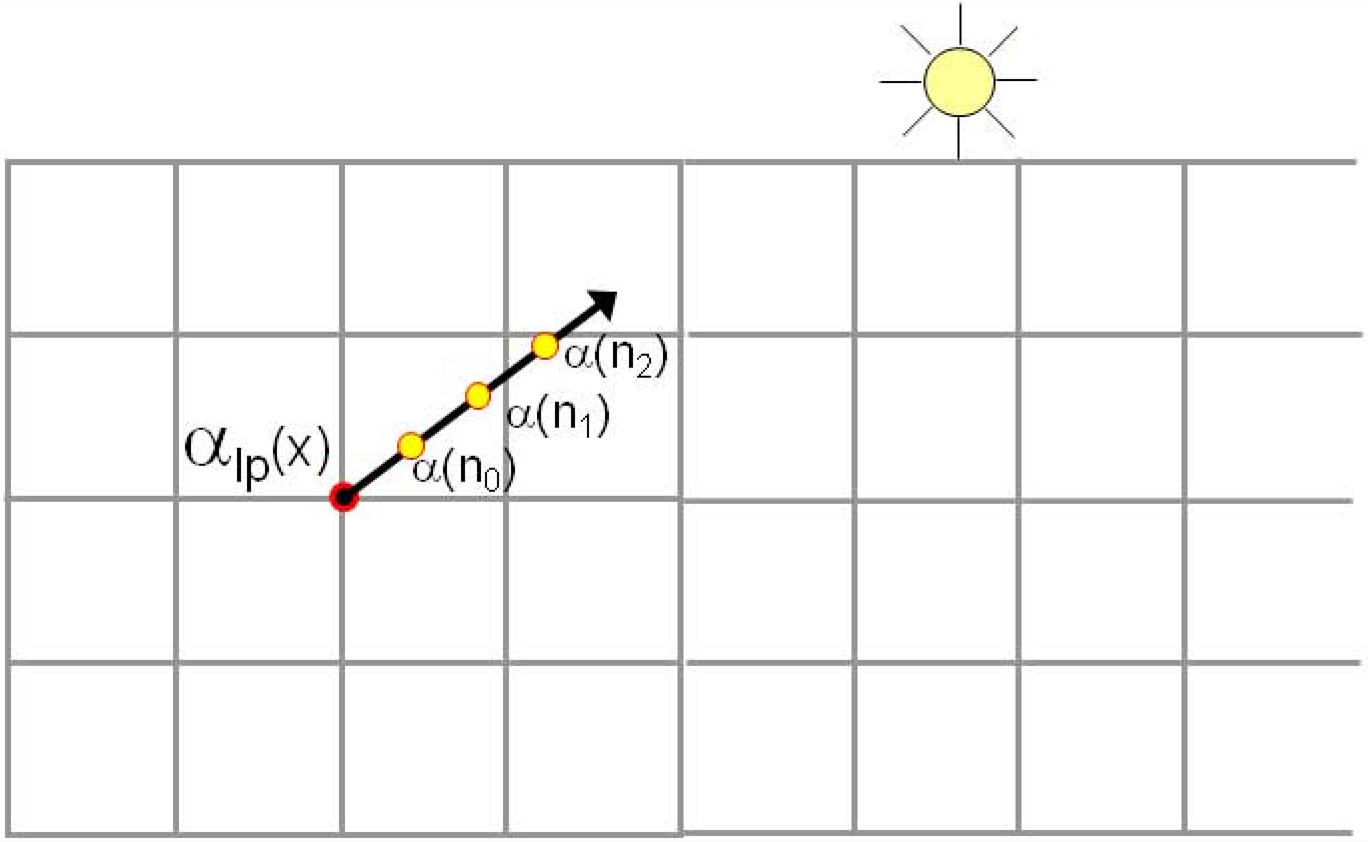

Piecewise Integration

$$T_{lp}(x) = e^{-\int_x^{x'} \tau(u)du} \qquad \alpha_{lp}(x) = 1 - \prod_{n=0}^{M}(1 - \alpha(n))$$

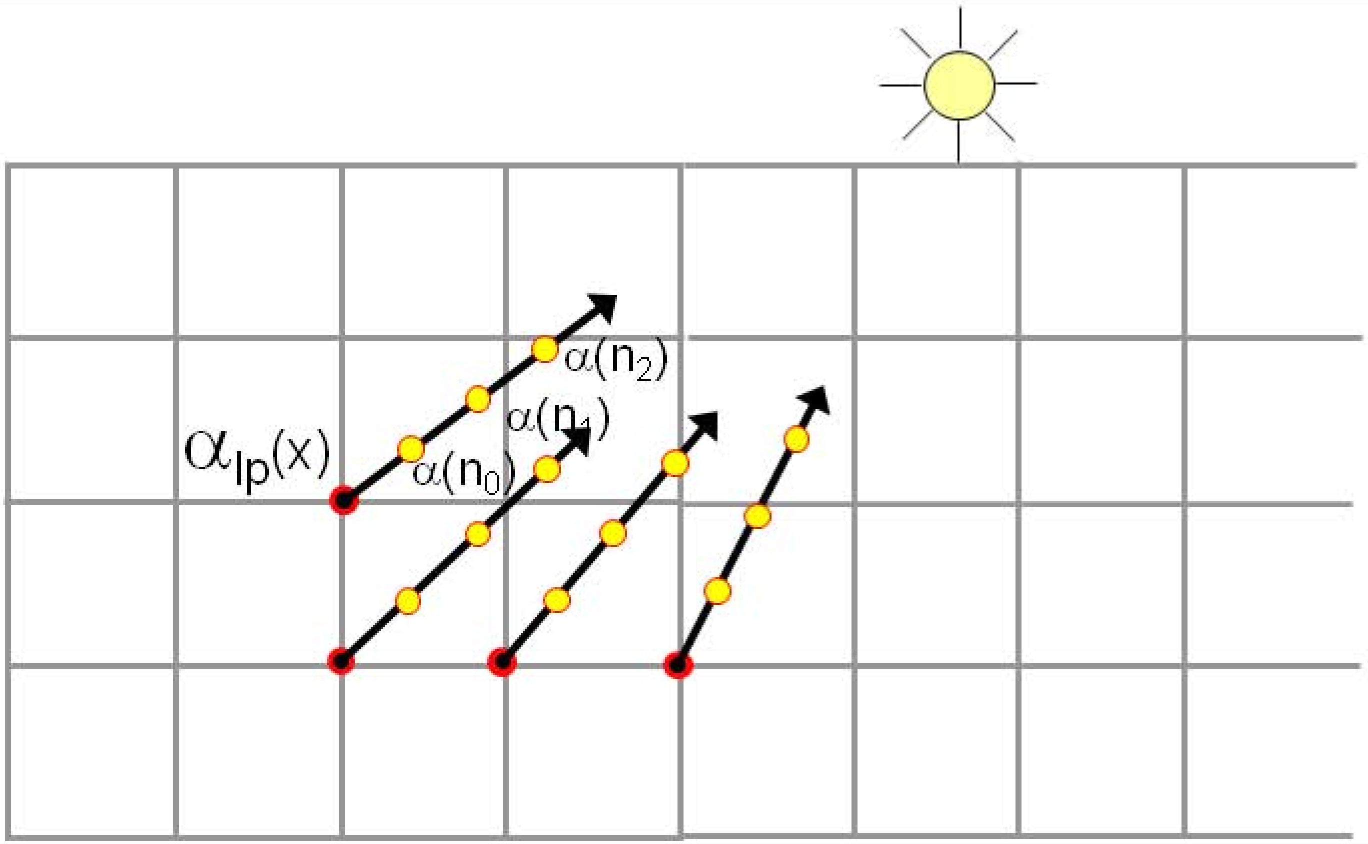

Piecewise Integration

$$\alpha_g(s_0) = 1 - \prod_{i=0}^{k} (1 - \alpha_{lp}(s_i))$$

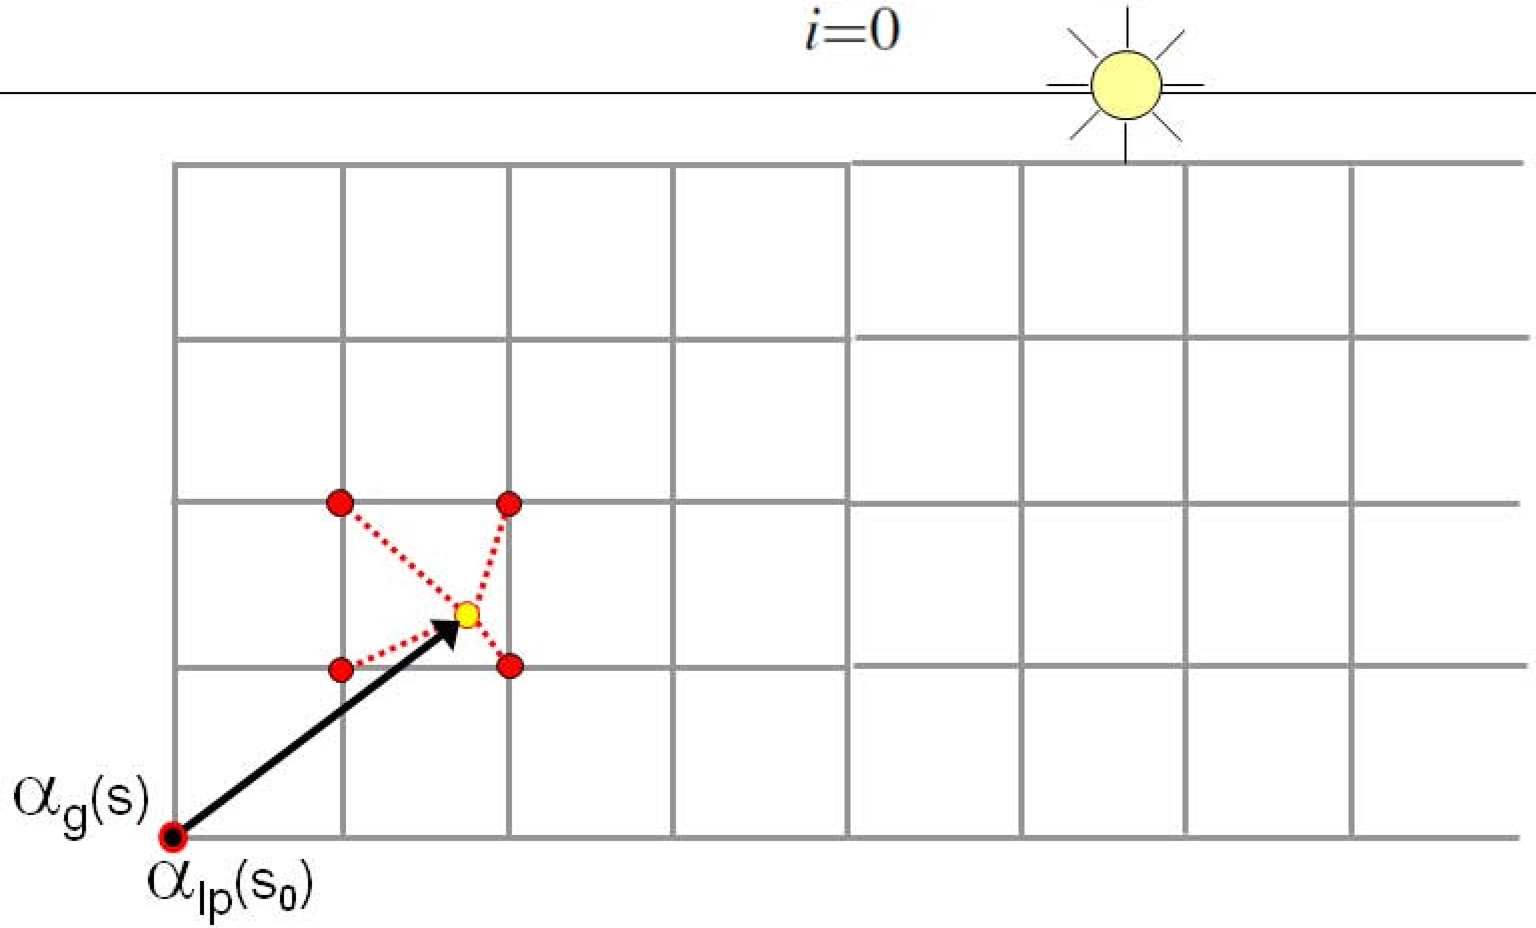

Piecewise Integration

$$\alpha_g(s_0) = 1 - \prod_{i=0}^{k}(1 - \alpha_{lp}(s_i))$$

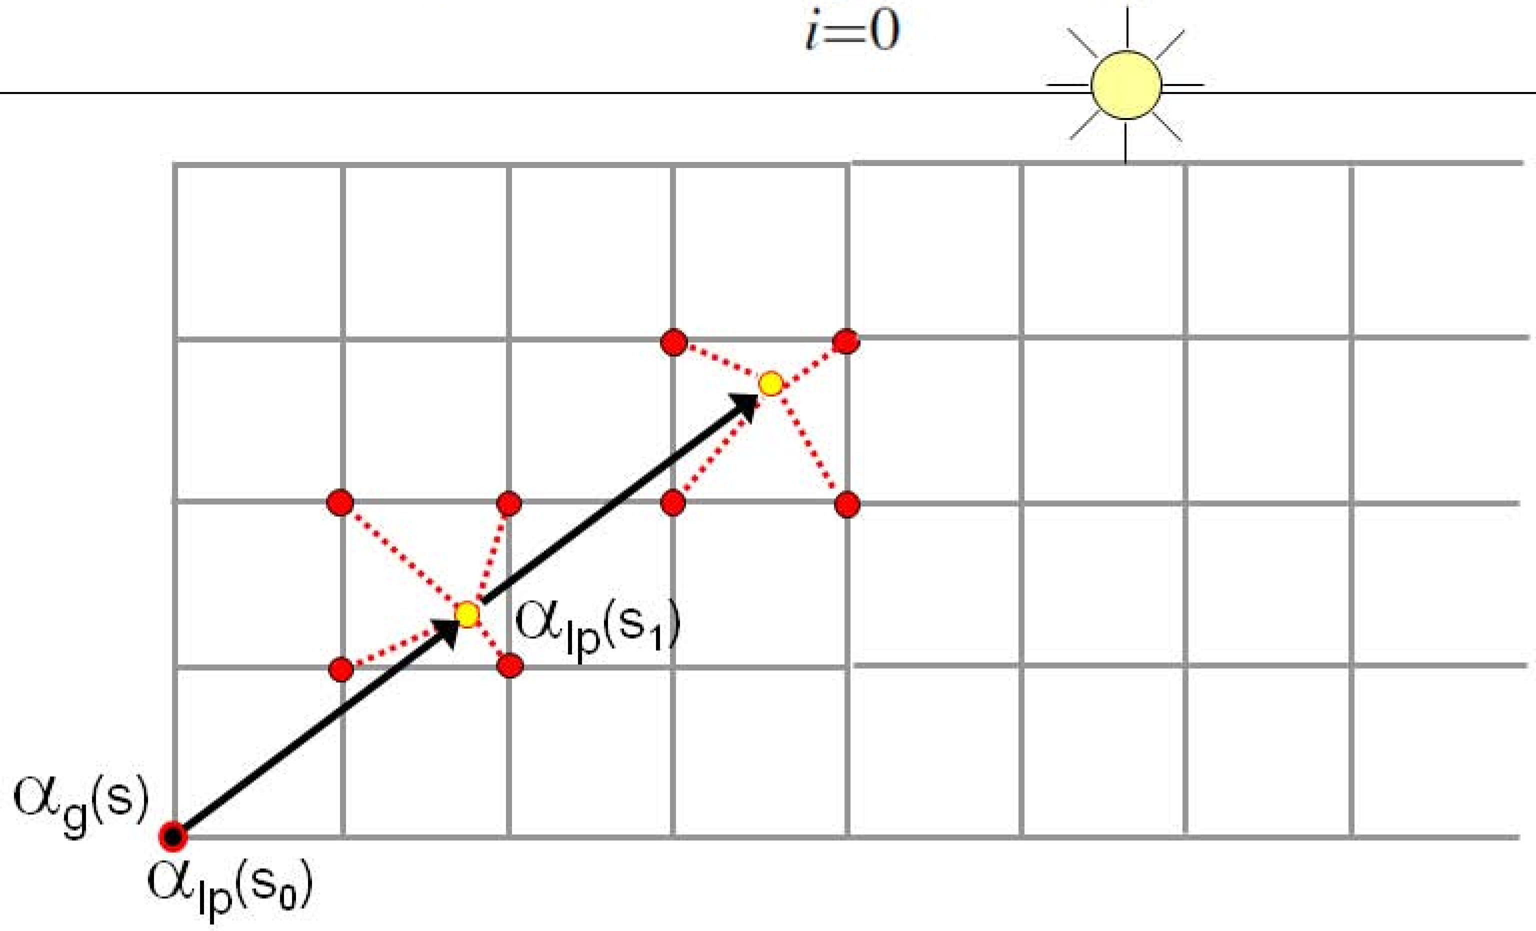

Improving Accuracy for Direct Light

$$I_d(s_0) = I_0 \cdot (1 - \alpha_g(s_1)) \cdot (1 - \alpha_{lp}(s_0))$$

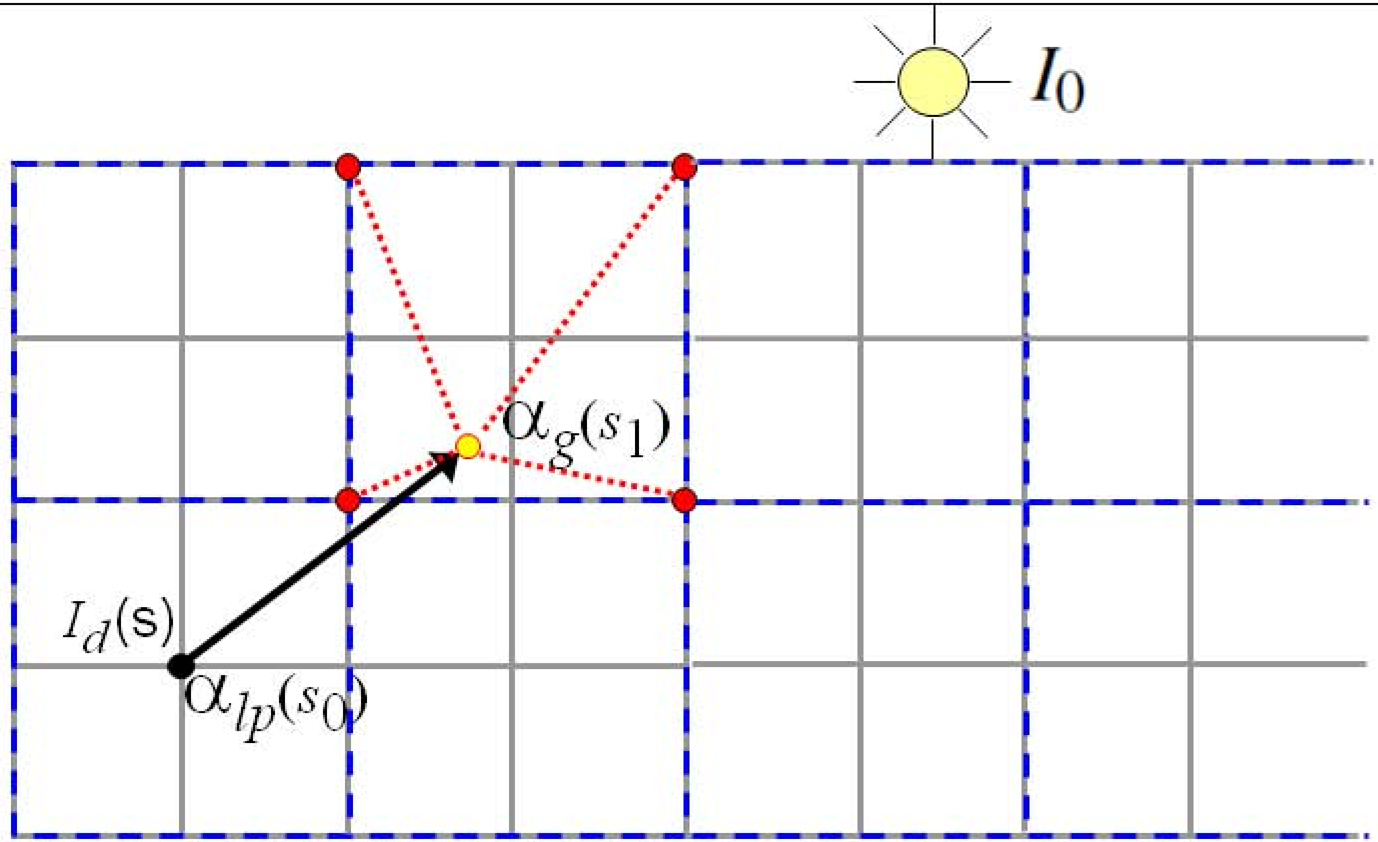

Shadow Volume Sensitivity

Original volume: 512^3 voxels, Datareduction: 8.9:1, Segment Length: 16 voxels

32^3

64^3

128^3

256^3

Global and LAO Example

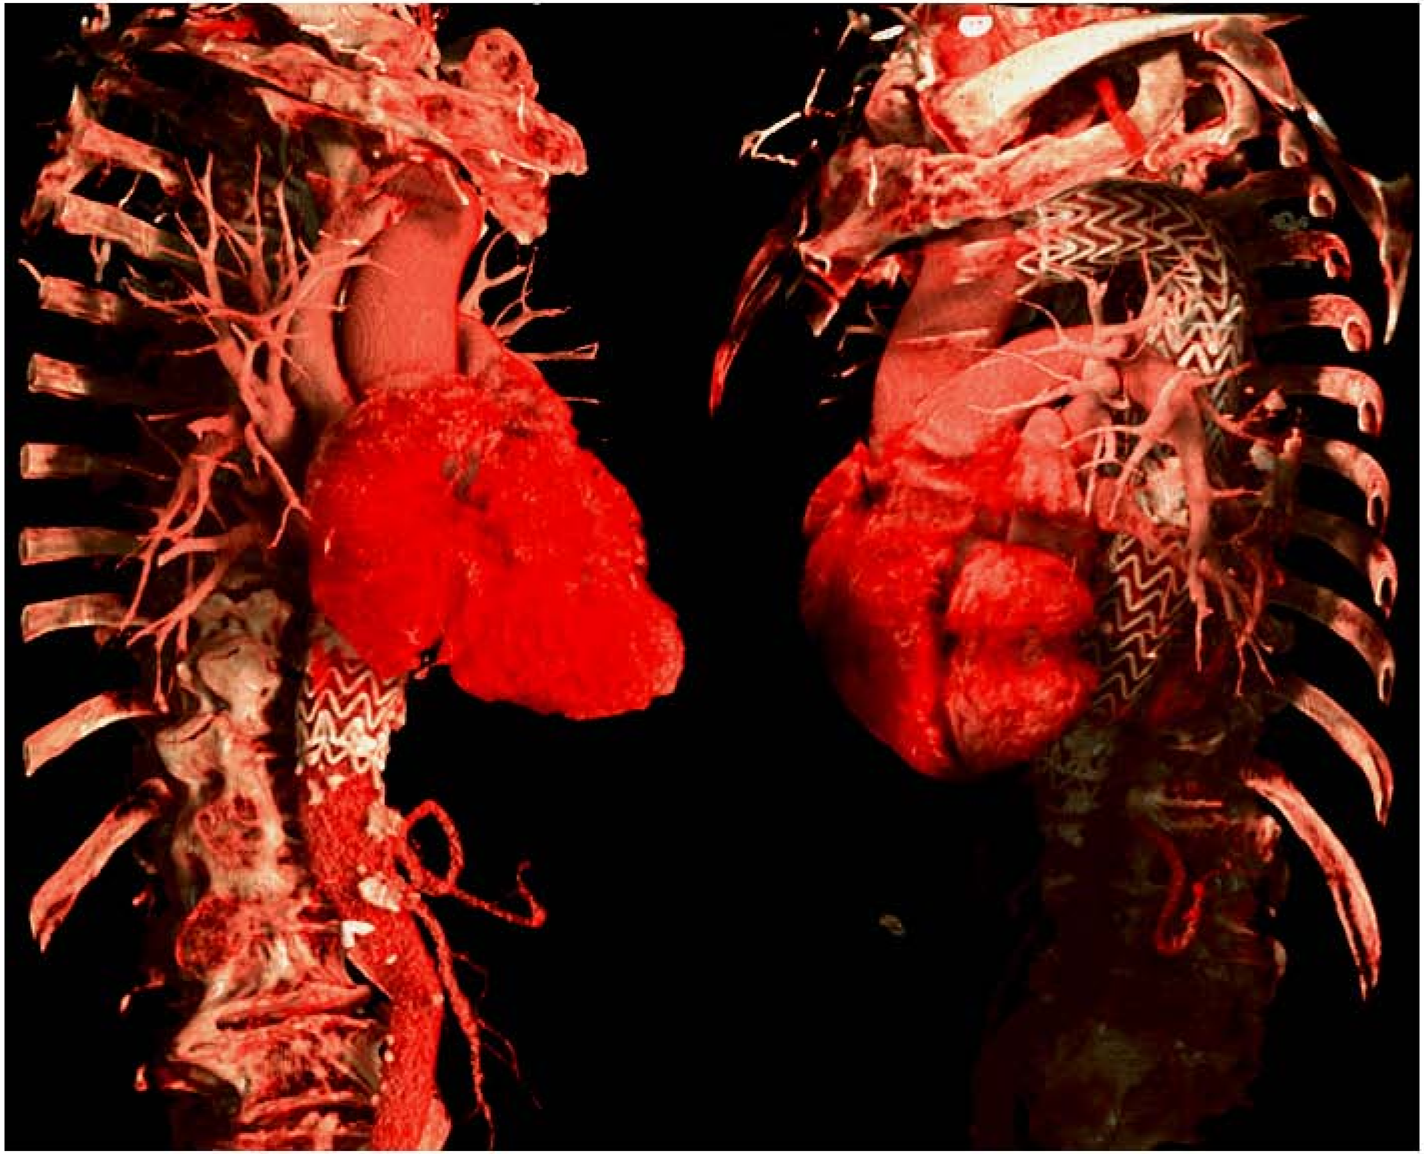

Compare w/ Diffuse Shading

Gradient based local lighting

Novel global + ambient lighting

Results: Translucency Effects

$R_\Omega = 16,\ I_{bias} = 0$ $R_\Omega = 48,\ I_{bias} = 0$

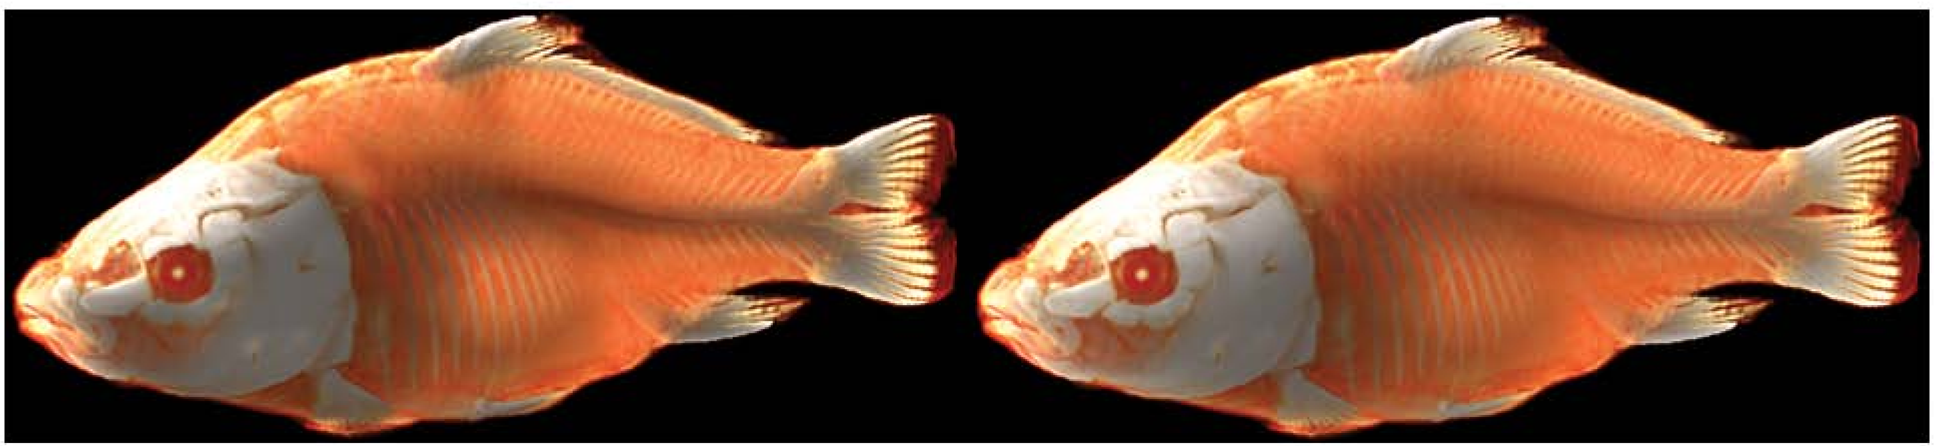

$R_\Omega = 16,\ I_{bias} = 0.2$ $R_\Omega = 48,\ I_{bias} = 0.2$

Performance Timing

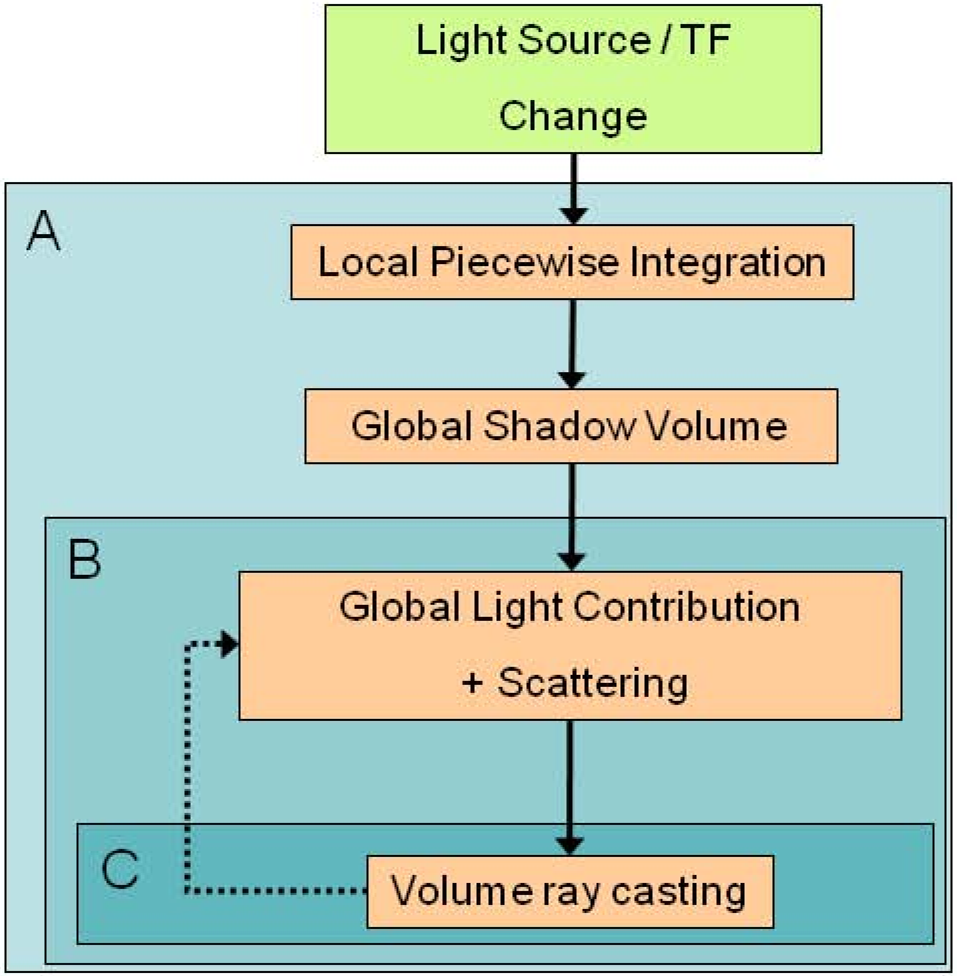

Piecewise segment length (voxels)	A			
	32^3	64^3	128^3	256^3
4	261	267	439	1428
8	284	297	373	552
16	331	339	380	641
32	436	439	463	862

milliseconds

Data reduction	A		B	C
	32^3	256^3		
8.9:1	284	552	233	68
14.8:1	178	515	145	68
22.1:1	121	403	96	48
35.2:1	81	365	62	46

milliseconds

Original volume: 512^3 voxels
Datareduction: 8.9:1
Segment Length: 16
Viewport: 1024x1024

Summary

- Advanced interactive shading techniques
- Significantly improved realism
- Volumetric and iso-surface techniques
- Complex structures and depth perception
- Color-bleeding and emission
- Global light propagation
- First order scattering

Final Views

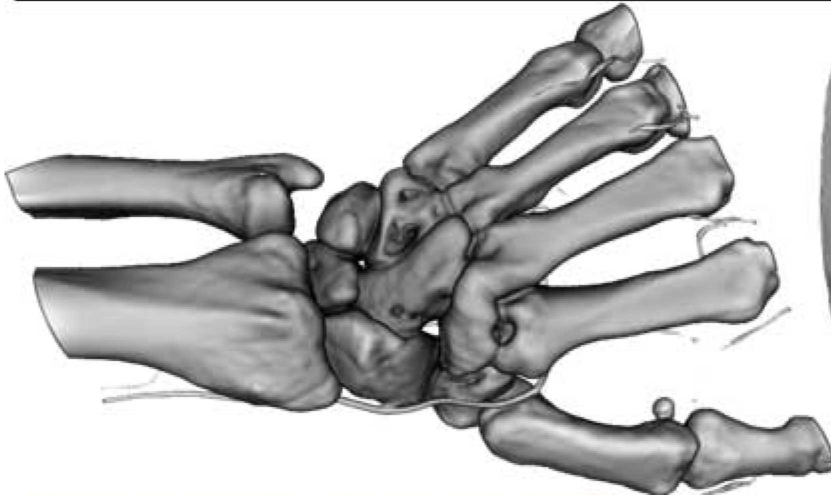

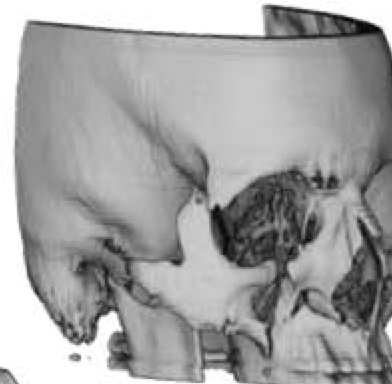

Future Work

- Improve performance
- End-user control
- Automated feature enhancement

References

- A. James Stewart. Vicinity Shading for Enhanced Perception of Volumetric Data. In *Proc. IEEE Visualization 2003*
- Kevin M. Beason, Josh Grant, David C. Banks, Brad Futch, and M. Yousuff Hussaini. Pre-Computed Illumination for Isosurfaces. In *Proc. of SPIE 2006*
- Chris Wyman, Steven Parker, Peter Shirley, and Charles Hansen. Interactive Display of Isosurfaces with Global Illumination. *IEEE TVCG Vol. 12, No. 2, 2006*
- Timo Ropinski, Jennis Meyer-Spradow, Stefan Diepenbrock, Jörg Mensmann, and Klaus Hinrichs. Interactive Volume Rendering with Dynamic Ambient Occlusion and Color Bleeding. *Computer Graphics Forum (Eurographics 2008), Volume 27, Number 2, 2008*
- Eric Penner and Ross Mitchell. Isosurface Ambient Occlusion and Soft Shadows with Filterable Occlusion Maps. In *Proc. of Eurographics/IEEE Volume Graphics 2008*

References

- Frida Hernell, Patric Ljung, and Anders Ynnerman. Efficient Ambient and Emissive Tissue Illumination using Local Occlusion in Multiresolution Volume Rendering. In *Proc. of Eurographics/IEEE Volume Graphics 2007*
- Frida Hernell, Patric Ljung, and Anders Ynnerman. Interactive Global Light Propagation in Direct Volume Rendering using Local Piecewise Integration. In *Proc. of Eurographics/IEEE Volume Graphics 2008*
- Patric Ljung, Claes Lundström, Ken Museth, and Anders Ynnerman. Transfer Function Based Adaptive Decompression for Volume Rendering of Large Medical Data Sets. In *Proc. of IEEE Volume Visualization and Graphics 2004*
- Patric Ljung, Claes Lundström, and Anders Ynnerman. Multiresolution Interblock Interpolation in Direct Volume Rendering. In *Proc. of Eurographics/IEEE Symposium on Visualization 2006*
- Patric Ljung. Adaptive Sampling in Single Pass, GPU-based Raycasting of Multiresolution Volumes. In *Proc. of Eurographics/IEEE Volume Graphics 2006*