Image-Based Rendering

CSE 274, Lecture 7: IBR as Sampled Data

Ravi Ramamoorthi

http://www.cs.ucsd.edu/~ravir


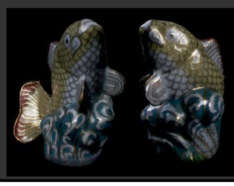
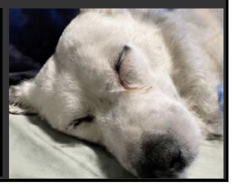

1

To Do

2

Motivation

- IBR is not just view synthesis (4D)
- Broader trend of sampled data (data-driven)
- *Data for lighting, BRDFs, motion, textures etc.*
  - *Modern Generative AI is essentially advanced data-driven texture synthesis*
- Precomputed Light Transport or Radiance Transfer extends this even to synthetic scenes
- All of these remain active areas of research

3

Reflection Maps

Blinn and Newell, 1976

4

Environment Maps

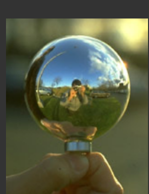


Miller and Hoffman, 1984

5

Environment Maps

Interface, Chou and Williams (ca. 1985)

6

Reflection Maps in the Movies

- From history, pauldebevec.com/ReflectionMapping
- First movie, Flight of the Navigator 1986

7

8

Rendering with Natural Light

Rendering with Natural Light, Debevec 98

9

Fiat Lux (Debevec 99)

- https://www.youtube.com/watch?v=vgJuzml0dIw
- https://www.pauldebevec.com/Probes/

10

Data-Driven BRDFs

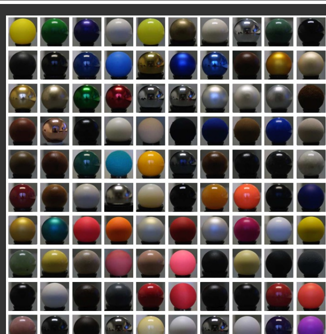

A Data-Driven Reflectance Model. Matusik et al. 03 (MERL 100 BRDF Database)

11

Motion Capture: “Signature” of Actor

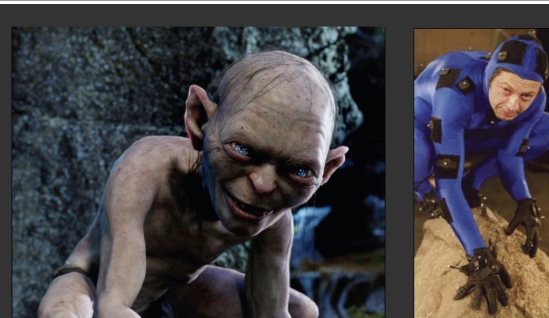

12

Facial MoCap

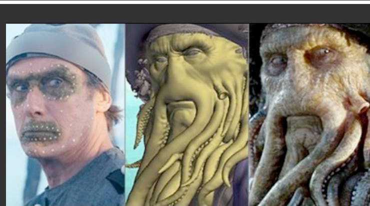

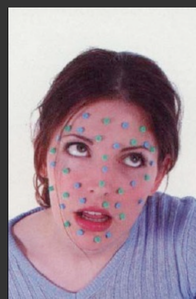

13

Weather Forecast for Dummies

- Let's predict weather:
  - Given today's weather only, we want to know tomorrow's
  - Suppose weather can only be {Sunny, Cloudy, Raining}
- The "Weather Channel" algorithm:
  - Over a long period of time, record:
    - How often S followed by R
    - How often S followed by S
    - Etc.
  - Compute percentages for each state:
    - P(R|S), P(S|S), etc.
  - Predict the state with highest probability!
  - It's a Markov Chain

14

Markov Chain

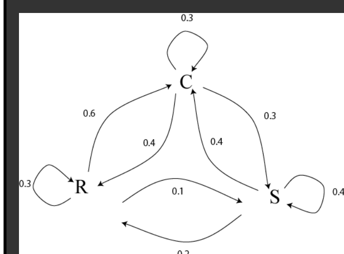

$$\begin{pmatrix} 0.3 & 0.6 & 0.1 \\ 0.4 & 0.3 & 0.3 \\ 0.2 & 0.4 & 0.4 \end{pmatrix}$$

What if we know today and yesterday's weather?

15

Text Synthesis

- [Shannon,'48] proposed a way to generate English-looking text using N-grams:
  - Assume a generalized Markov model
  - Use a large text to compute prob. distributions of each letter given N-1 previous letters
  - Starting from a seed repeatedly sample this Markov chain to generate new letters
  - Also works for whole words

WE NEED TO EAT CAKE

16

Mark V. Shaney (Bell Labs)

- Results (using `alt.singles` corpus):
  - *"As I've commented before, really relating to someone involves standing next to impossible."*
  - *"One morning I shot an elephant in my arms and kissed him."*
  - *"I spent an interesting evening recently with a grain of salt"*

17

Texture

- Texture depicts spatially repeating patterns
- Many natural phenomena are textures

radishes

rocks

yogurt

18

Texture Synthesis

- Goal of Texture Synthesis: create new samples of a given texture
- Many applications: virtual environments, hole-filling, texturing surfaces

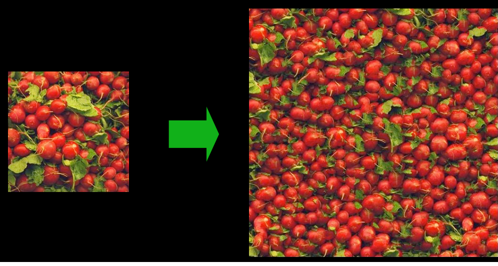

19

Today: Text to Image

- https://www.youtube.com/watch?v=GYyP7Ova8KA&list=PLWfDJ5nla8UpwShx-lzLJqcp575fKpsSO&index=25
- https://www.youtube.com/watch?v=lyodbLwb2lY&list=PLWfDJ5nla8UpwShx-lzLJqcp575fKpsSO&index=26
- (From Steve Seitz 5 minute videos, also check out language model videos at) https://www.youtube.com/playlist?list=PLWfDJ5nla8UpwShx-lzLJqcp575fKpsSO

20

Motivation

- IBR is not just view synthesis (4D)
- Broader trend of sampled data (data-driven)
- Data for lighting, BRDFs, motion, textures etc.
  - Modern Generative AI is essentially advanced data-driven texture synthesis
- *Precomputed Light Transport or Radiance Transfer extends this even to synthetic scenes*
- All of these remain active areas of research

21

Precomputed Rendering

- Image-Based Rendering: Use measured data (real photographs) and interpolate for realistic real-time
- Why not apply to real-time rendering?
  - Precompute (offline) some information (images) of interest
  - Must assume something about scene is constant to do so
  - Thereafter real-time rendering. Often accelerate hardware
- Easier and harder than conventional IBR
  - Easier because synthetic scenes give info re geometry, reflectance (but CG rendering often longer than nature)
  - Harder because of more complex effects (lighting from all directions for instance, not just changing view)
- Representations and Signal-Processing crucial

22

My General Philosophy

- This general line of work is a large data management and signal-processing problem
- Precompute high-dimensional complex data
- Store efficiently (find right mathematical represent.)
- Render in real-time
  - Worry about systems issues like caching
  - Good signal-processing: use only small amount of data but guarantee high fidelity
- Many insights into structure of lighting, BRDFs, …
  - Not just blind interpolation; signal processing

23

Precomputation-Based Relighting

- Analyze precomputed images of scene

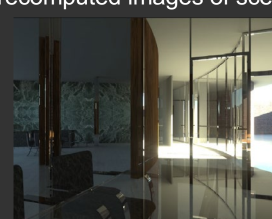

Jensen 2000

24

Precomputation-Based Relighting

- Analyze precomputed images of scene

Jensen 2000

25

Assumptions

- Static geometry
- Precomputation

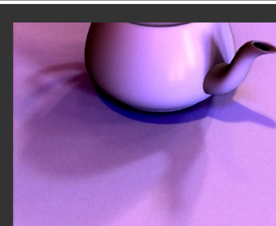

- Real-Time Rendering (relight all-frequency effects)
  - Exploit linearity of light transport for this
  - Later, change viewpoint as well

26

Why is This Hard?

- Plain graphics hardware supports only simple (point) lights, BRDFs (Phong) without any shadows
- Shadow maps can handle point lights (hard shadows)
- Environment maps complex lighting, BRDFs but no shadows
- IBR can often do changing view, fixed lighting

- How to do complex shadows in complex lighting?
- With dynamically changing illumination and view?

27

Relighting as a Matrix-Vector Multiply

$$\begin{bmatrix} P_1 \\ P_2 \\ P_3 \\ \vdots \\ P_N \end{bmatrix} = \begin{bmatrix} T_{11} & T_{12} & \cdots & T_{1M} \\ T_{21} & T_{22} & \cdots & T_{2M} \\ T_{31} & T_{32} & \cdots & T_{3M} \\ \vdots & \vdots & \ddots & \vdots \\ T_{N1} & T_{N2} & \cdots & T_{NM} \end{bmatrix} \begin{bmatrix} L_1 \\ L_2 \\ \vdots \\ L_M \end{bmatrix}$$

28

Relighting as a Matrix-Vector Multiply

$$\begin{bmatrix} P_1 \\ P_2 \\ P_3 \\ \vdots \\ P_N \end{bmatrix} = \begin{bmatrix} T_{11} & T_{12} & \cdots & T_{1M} \\ T_{21} & T_{22} & \cdots & T_{2M} \\ T_{31} & T_{32} & \cdots & T_{3M} \\ \vdots & \vdots & \ddots & \vdots \\ T_{N1} & T_{N2} & \cdots & T_{NM} \end{bmatrix} \begin{bmatrix} L_1 \\ L_2 \\ \vdots \\ L_M \end{bmatrix}$$

Output Image (Pixel Vector)

Input Lighting (Cubemap Vector)

Precomputed Transport Matrix

29

Matrix Columns (Images)

$$\begin{bmatrix} T_{11} & T_{12} & \cdots & T_{1M} \\ T_{21} & T_{22} & \cdots & T_{2M} \\ T_{31} & T_{32} & \cdots & T_{3M} \\ \vdots & \vdots & \ddots & \vdots \\ T_{N1} & T_{N2} & \cdots & T_{NM} \end{bmatrix}$$

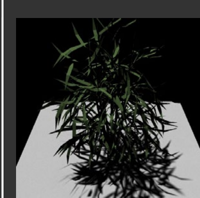

30

Precompute: Ray-Trace Image Cols

$$\begin{bmatrix} T_{11} & T_{12} & \cdots & T_{1M} \\ T_{21} & T_{22} & \cdots & T_{2M} \\ T_{31} & T_{32} & \cdots & T_{3M} \\ \vdots & \vdots & \ddots & \vdots \\ T_{N1} & T_{N2} & \cdots & T_{NM} \end{bmatrix}$$

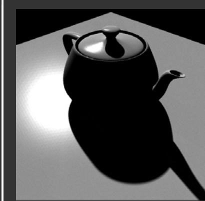

31

Precompute 2: Rasterize Matrix Rows

$$\begin{bmatrix} T_{11} & T_{12} & \cdots & T_{1M} \\ T_{21} & T_{22} & \cdots & T_{2M} \\ T_{31} & T_{32} & \cdots & T_{3M} \\ \vdots & \vdots & \ddots & \vdots \\ T_{N1} & T_{N2} & \cdots & T_{NM} \end{bmatrix}$$

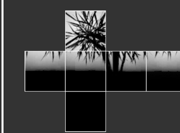

32

Problem Definition

Matrix is Enormous
- 512 x 512 pixel images
- 6 x 64 x 64 cubemap environments

Full matrix-vector multiplication is intractable
- On the order of 10^{10} operations *per frame*

How to relight quickly?

33

Outline

- Motivation and Background
- *Compression methods*
  - *Low frequency linear spherical harmonic approximation*
  - Factorization and PCA
  - Local factorization and clustered PCA
  - Non-linear wavelet approximation
- Changing view as well as lighting

34

Precomputed Radiance Transfer

- Better light integration and transport
  - dynamic, area lights
  - self-shadowing
  - interreflections
- For diffuse and glossy surfaces
- At real-time rates
- Sloan et al. 02 (most cited rendering paper in last 20 years 1000+, widely used in games, movie production: Spherical Harmonic Lighting)

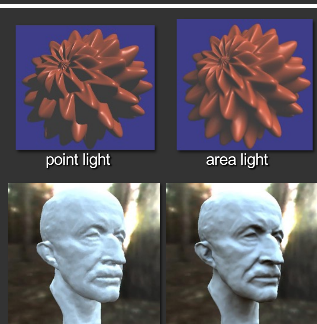

point light | area light

area lighting, no shadows | area lighting, shadows

35

Precomputation: Spherical Harmonics

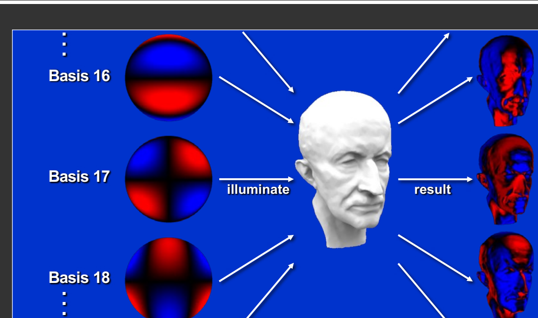

36

Diffuse Transfer Results

No Shadows/Inter | Shadows | Shadows+Inter

37

Arbitrary BRDF Results

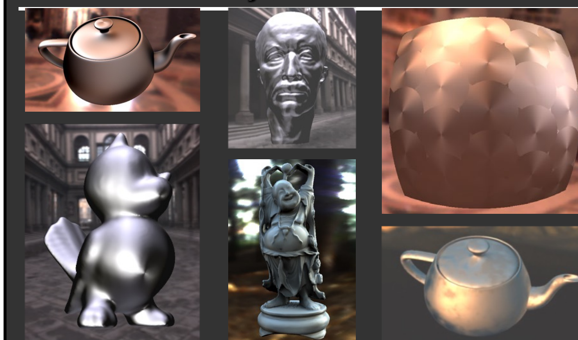

Anisotropic BRDFs | Other BRDFs | Spatially Varying

38

Precomputed Lighting (Avatar 2010)

39

Relighting as a Matrix-Vector Multiply

$$\begin{bmatrix}P_1\\P_2\\P_3\\\vdots\\P_N\end{bmatrix} = \begin{bmatrix}T_{11} & T_{12} & \cdots & T_{1M}\\T_{21} & T_{22} & \cdots & T_{2M}\\T_{31} & T_{32} & \cdots & T_{3M}\\\vdots & \vdots & \ddots & \vdots\\T_{N1} & T_{N2} & \cdots & T_{NM}\end{bmatrix}\begin{bmatrix}L_1\\L_2\\\vdots\\L_M\end{bmatrix}$$

40

Idea of Compression

- The vector is projected onto low-frequency components (say 25). Size greatly reduced.
- Hence, only 25 matrix columns
- But each pixel still treated separately (still have 300000 matrix rows for 512 x 512 image)
- Actually, for each pixel, dot product of matrix row (25 elems) and lighting vector (25 elems) in hardware
- Good technique (common in games, movies) but useful only for broad low-frequency lighting

41

Outline

- Motivation and Background
- *Compression methods*
  - Low frequency linear spherical harmonic approximation
  - *Factorization and PCA*
  - *Local factorization and clustered PCA*
  - Non-linear wavelet approximation
- Changing view as well as lighting
  - Clustered PCA
  - Factored BRDFs
  - Triple Product Integrals

42

PCA or SVD factorization

• SVD:

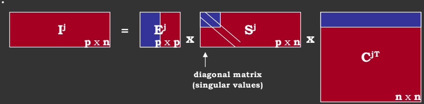

• Applying Rank **b**:

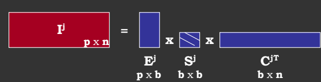

• Absorbing S^j values into C^{jT}:

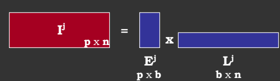

43

Idea of Compression

- Represent matrix (rather than light vector) compactly
- Can be (and is) combined with low frequency vector
- Useful in broad contexts.
  - BRDF factorization for real-time rendering (reduce 4D BRDF to 2D texture maps) McCool et al. 01 etc
  - Surface Light field factorization for real-time rendering (4D to 2D maps) Chen et al. 02, Nishino et al. 01
  - Factorization of Orientation Light field for complex lighting and BRDFs (4D to 2D) Latta et al. 02
- **Not too useful for general precomput. relighting**
  - **Transport matrix not low-dimensional!!**

44

Local or Clustered PCA

- Exploit local coherence (in say 16x16 pixel blocks)
  - Idea: light transport is locally low-dimensional. Why?
  - Even though globally complex
  - See Mahajan et al. 07 for theoretical analysis
- Original idea: Each triangle separately
  - Example: Surface Light Fields 3D subspace works well
  - Vague analysis of size of triangles
  - Instead of triangle, 16x16 image blocks [Nayar et al. 04]
- Clustered PCA [Sloan et al. 2003]
  - Combines two widely used compression techniques: Vector Quantization or VQ and Principal Component Analysis
  - For complex geometry, no need for parameterization / topology

45

Image-Based Rendering

Practical Case

Human Face

Zickler, Enrique, Ramamoorthi, Belhumeur 05, 06

46

Outline

- Motivation and Background
- *Compression methods*
  - Low frequency linear spherical harmonic approximation
  - Factorization and PCA
  - Local factorization and clustered PCA
  - *Non-linear wavelet approximation*
- Changing view as well as lighting

47

Sparse Matrix-Vector Multiplication

Choose data representations with mostly zeroes

Vector: Use *non-linear wavelet approximation* on lighting

Matrix: Wavelet-encode transport rows

$$\begin{bmatrix} T_{11} & T_{12} & \cdots & T_{1M} \\ T_{21} & T_{22} & \cdots & T_{2M} \\ T_{31} & T_{32} & \cdots & T_{3M} \\ \vdots & \vdots & \ddots & \vdots \\ T_{N1} & T_{N2} & \cdots & T_{NM} \end{bmatrix} \begin{bmatrix} L_1 \\ L_2 \\ \vdots \\ L_M \end{bmatrix}$$

48

Haar Wavelet Basis

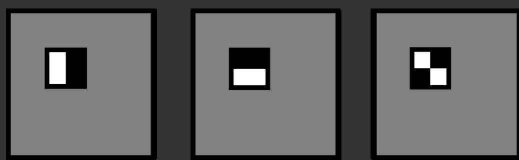

49

Non-linear Wavelet Approximation

Wavelets provide dual space / frequency locality
- Large wavelets capture low frequency area lighting
- Small wavelets capture high frequency compact features

Non-linear Approximation
- Use a dynamic set of approximating functions (*depends on each frame's lighting*)
- By contrast, linear approx. uses fixed set of basis functions (like 25 lowest frequency spherical harmonics)
- We choose 10's - 100's from a basis of 24,576 wavelets

50

Non-linear Wavelet Light Approximation

Wavelet Transform

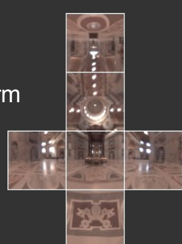

51

Non-linear Wavelet Light Approximation

$$\begin{bmatrix} 0 \\ L_2 \\ 0 \\ 0 \\ 0 \\ L_6 \\ \vdots \\ 0 \end{bmatrix}$$

Non-linear Approximation

Retain 0.1% – 1% terms

52

Error in Lighting: St Peter's Basilica

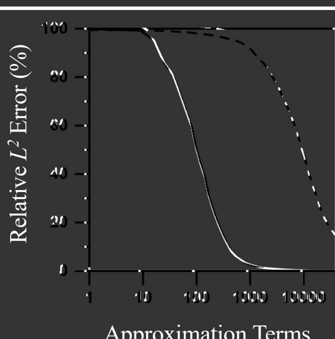

Ng, Ramamoorthi, Hanrahan 03

53

Output Image Comparison

Top: Linear Spherical Harmonic Approximation
Bottom: Non-linear Wavelet Approximation

25	200	2,000	20,000

54

Video: Real Time Relighting

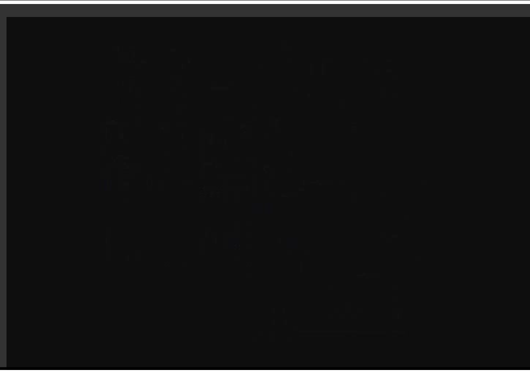

55

Outline

- Motivation and Background
- Compression methods
  - Low frequency linear spherical harmonic approximation
  - Factorization and PCA
  - Local factorization and clustered PCA
  - Non-linear wavelet approximation
- *Changing view as well as lighting*

56

Changing Only The View

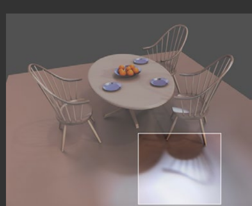

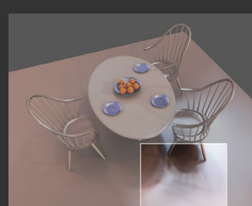

57

Problem Characterization

6D Precomputation Space

- Distant Lighting (2D)
- View (2D)
- Rigid Geometry (2D)

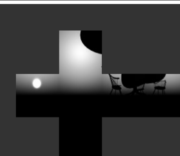

With ~ 100 samples per dimension

~ 10^{12} samples total!! : Intractable computation, rendering

58

Clustered PCA

- Use low-frequency light and view variation (Order 4 spherical harmonic = 25 for both; total = 25*25=625)
- 625 element vector for each vertex
- Apply CPCA directly (Sloan et al. 2003)
- Does not easily scale to high frequencies
  - Really cubic complexity (number of vertices, illumination directions or harmonics, and view directions or harmonics)
- Practical real-time method on GPU

59

Factored BRDFs

- Sloan et al. 04, Wang et al. 04: All-frequency effects
- Combines lots of things: BRDF factorization, CPCA, nonlinear approx. with wavelets
- Idea: Factor BRDF to depend on incident, outgoing
  - Incident part handled with view-independent relighting
  - Then linearly combine based on outgoing factor
- Effectively, break problem into a few subproblems that can be solved view-independently and added up
  - Can apply nonlinear wavelet approx. to each subproblem
  - And CPCA to the matrices for further compression

60

Factorization Approach

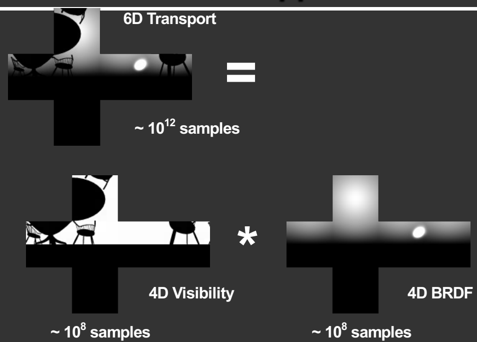

61

Triple Product Integral Relighting

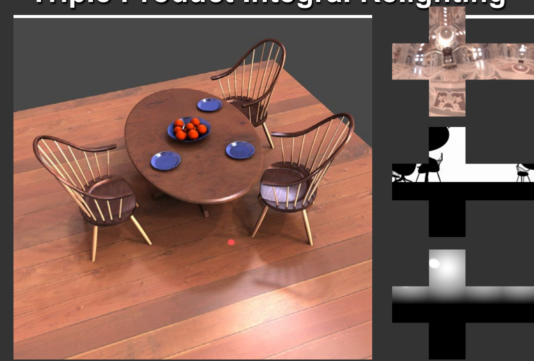

62

Relit Images (3-5 sec/frame)

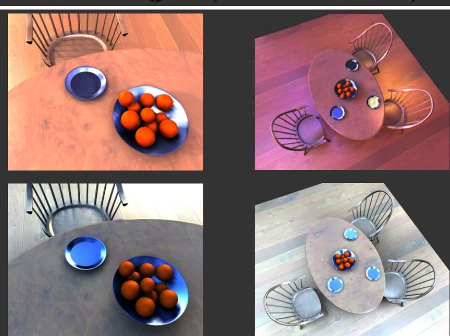

63

Summary

- Really a big data compression and signal-processing problem
- Apply many standard methods
  - PCA, wavelet, spherical harmonic, factor compression
- And invent new ones
  - VQPCA, wavelet triple products
- Guided by and gives insights into properties of illumination, reflectance, visibility
  - How many terms enough? How much sparsity?

64

Subsequent Work

- My survey 2009 (lecture only covers 2002-2004)
- Varied lighting/view. What about dynamic scenes, BRDFs
  - Much subsequent work [Zhou et al. 05, Ben-Artzi et al. 06]. But still limited for dynamic scenes
- Must work on GPU to be practical
- Sampling on object geometry remains a challenge
- Near-Field Lighting has had some work, remains a challenge
- Applications to lighting design, direct to indirect transfer
- New basis functions and theory
- Newer methods do not require precompute, various GPU tricks
- So far, low-frequency spherical harmonics used in games, all-frequency techniques have had limited applicability

65

Analytic SH Gradients

Analytic Spherical Harmonic Gradients for Real-Time Rendering with Many Polygonal Area Lights

Lifan Wu[1], Guangyan Cai[1], Shuang Zhao[2], Ravi Ramamoorthi[1]
[1] UC San Diego, [2] UC Irvine

NO AUDIO

66