Sampling and Reconstruction of Visual Appearance: From Denoising to View Synthesis

CSE 274 [Fall 2022], Lecture 5

Ravi Ramamoorthi

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

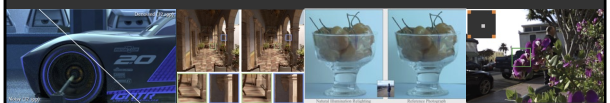

1

Applications: Sampling/Reconstruction

- *Monte Carlo Rendering*
- Light Transport Acquisition
- Light Fields and Computational Photography
- View synthesis
- Animation/Simulation (not covered in course)

- Brief overview of these applications today, and opportunities/history for sampling/reconstruction

2

Motivation

- Distribution effects (depth of field, motion blur, global illumination, soft shadows) are slow. Many dimensions sample

- Ray Tracing physically accurate but slow, not real-time
- Can we adaptively sample and filter for fast, real-time?

3

Monte Carlo Path Tracing

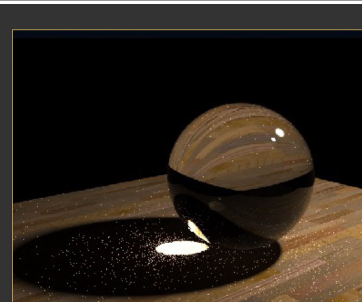

1000 paths/pixel

Jensen

4

Sampling and Reconstruction

- Monte Carlo is noisy at low sample counts
- Can we reduce time/samples by smart adaptive sampling and smart filtering/reconstruction?
- General area of Monte Carlo denoising
- Long history [Mitchell 91, Guo 98]

5

History

- Adaptive sampling old technique Mitchell et al. 87, 91,…

- But not very widely used… artifacts, can miss features
- After seminal papers in 87-91, not much follow on

6

Directional Coherence Maps

- Allocate samples to edges (Guo 98) Most of variance at those edges in the image

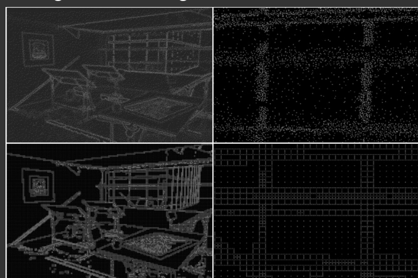

Figure 5: Comparison of the sampling patterns of adaptive stochastic sampling (top row) and the DCM (bottom row). The patterns in the left column are taken from RADIANCE work images described in Section 5. Some of the fine features are shown in zoomed views of the sampling patterns in the right column. These zoomed views correspond to the same region as the zoomed views in Fig. 4.

7

Directional Coherence Maps (Guo 98)

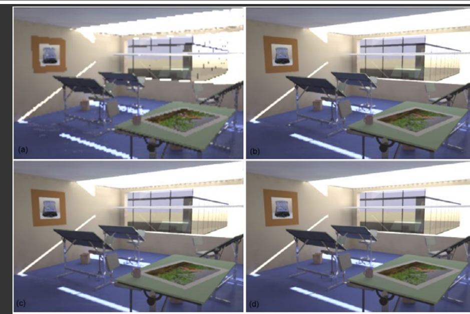

Figure 2: Progressive renderings of an office scene lit by sunlight transferred through a light shelf. (a) The approximate image at the end of the regular subdivision, with 1.6% evaluated pixels located at the corners of the 8 × 8 blocks in the work image. (b) The approximate image after boundary evaluations for all 8 × 8 edge blocks in the work image, with 5% of pixels evaluated. (c) The approximate image after evaluating about 6% of the pixels, whose locations are shown in Fig. 5 (bottom left). (d) The final image as rendered by the baseline RADIANCE system. The scene model was supplied courtesy of Greg W. Larson.

Guo 98

8

Resurgence (2008 -)

- Eurographics 2015 STAR report by Zwicker et al. [former UCSD faculty, now at Maryland]
- [Durand et al. 2005] Frequency analysis light transport. Proposed use for adaptive sampling. Not very practical

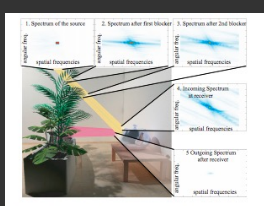

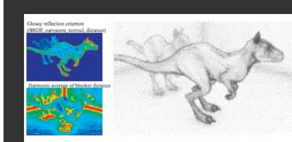

9

Multi-Dimensional Adaptive Sampling

- Hachisuka, Jarosz, … Zwicker, Jensen [MDAS 2008]
- Scenes with motion blur, depth of field, soft shadows
- Involves high-dimensional integral, converges slowly
- Exploit high-dimensional info to sample adaptively
- Sampling in 2D image plane or other dims inadequate
  - Need to consider full joint high-dimensional space

10

Multidimensional Adaptive Sampling

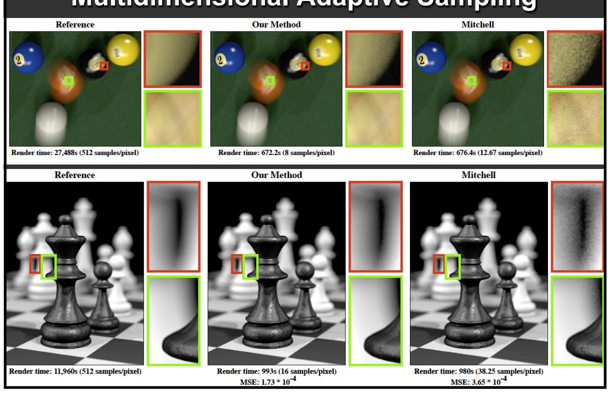

11

Multidimensional Adaptive Sampling

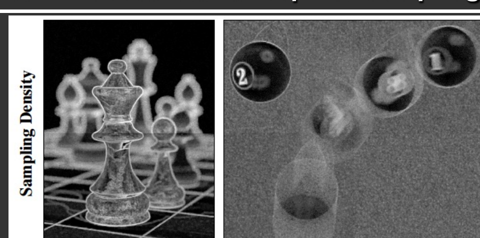

Figure 10: *Visualizations of projected sample distributions using our method for the chess scene from Figure 8 and the pool scene from Figure 7. Our adaptive sampler places samples both around high frequency image discontinuities (in focus chess piece and stationary pool ball) as well as in regions which exhibit significant motion blur or depth of field effects.*

12

Multi-Dimensional Adaptive Sampling

Motion Blur and Depth of Field 32 samples per pixel

13

A-Priori Methods

- Egan et al. 2009: Frequency Analysis and Sheared Filtering for Motion Blur; first deep use frequency anal.

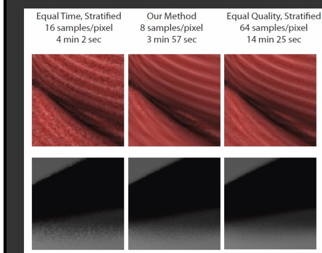

14

A-Posteriori Methods

- Adaptive Wavelet Rendering (Overbeck et al. 2009)
- Handle general effects. Sample and denoise
- Many more sophisticated methods available now; used in almost every major production rendering software

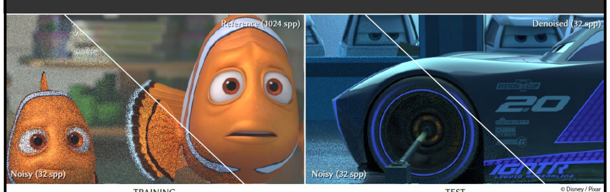

Bako et al. 17

15

Adaptive Wavelet Rendering

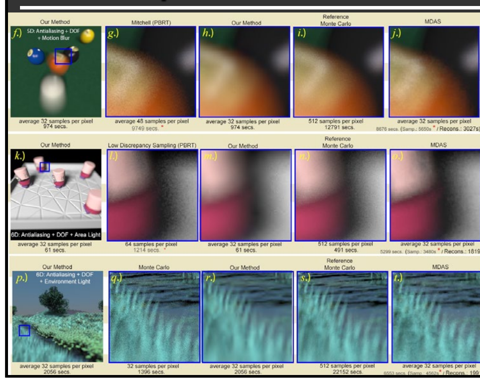

FF VIDEO

Overbeck et al 09
General high-D
effects. Simple
and fast (renders
Into wavelet dom

16

Real-Time

- Axis-Aligned Filtering (Mehta et al. 12,13,14)
- Optix plus image-space filtering
- Newer extensions to sheared filtering
- Most recent work (NVIDIA) is fully general, 1 sample per pixel, using modern machine learning methods (similar ideas relevant in offline rendering as well)
- Huge impact in real-time, video games, essential in modern real-time rendering based on deep learning

17

Real-Time MAAF Video

Multiple Axis-Aligned Filters for Rendering of Combined Distribution Effects

Online submission ID: 1000

NO AUDIO

Recurrent Autoencoder Video (Chaitanya et al. 17)

18

Applications: Sampling/Reconstruction

- Monte Carlo Rendering
- *Light Transport Acquisition*
- Light Fields and Computational Photography
- View synthesis
- Animation/Simulation (not covered in course)

- Brief overview of these applications today, and opportunities/history for sampling/reconstruction

19

Acquiring Reflectance Field of Human Face [Debevec et al. SIGGRAPH 00]

Illuminate subject from many incident directions

20

Example Images

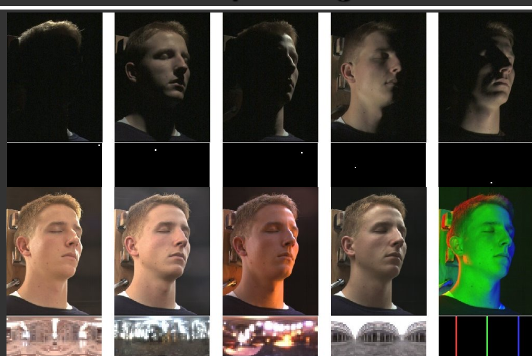

21

Motivation: Image-based Relighting

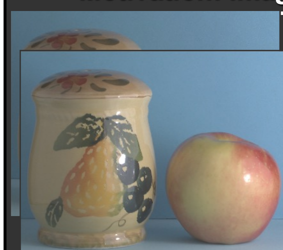

Sample Lighting Directions

22

Motivation: Image-based Relighting

Sample Lighting Directions

23

Motivation: Image-based Relighting

Sample Lighting Directions

24

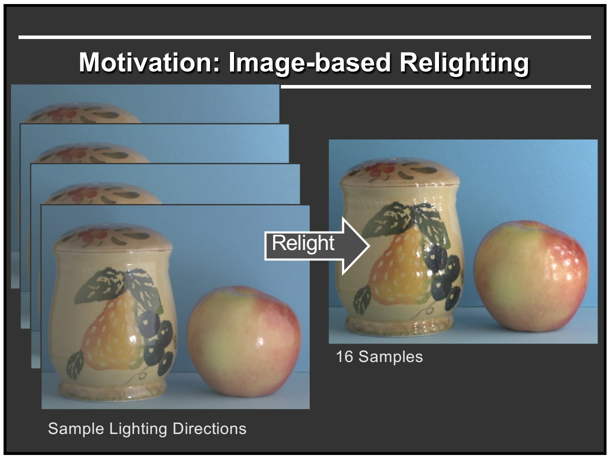

25

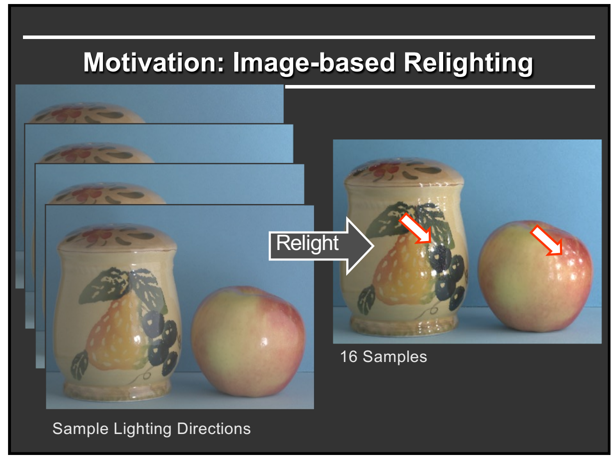

26

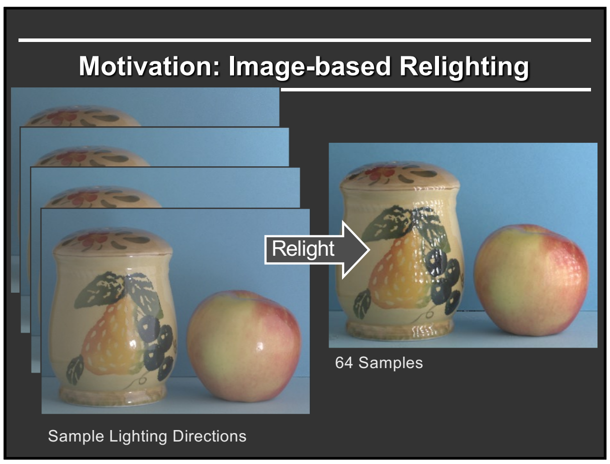

27

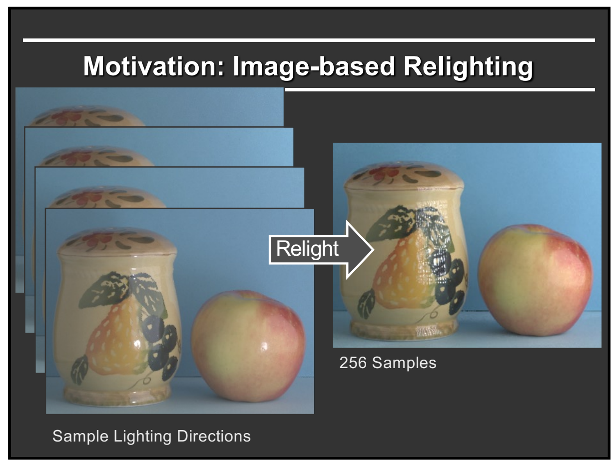

28

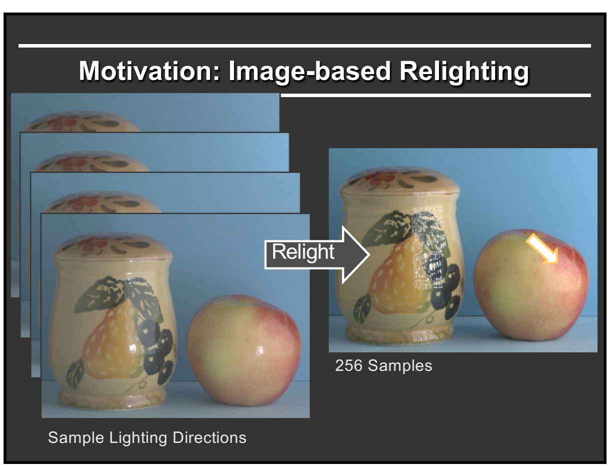

29

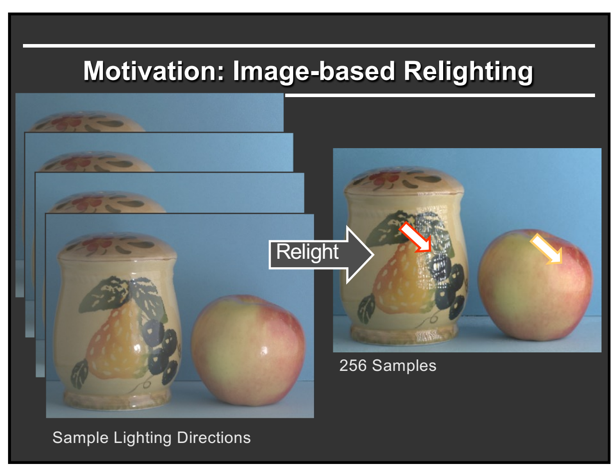

30

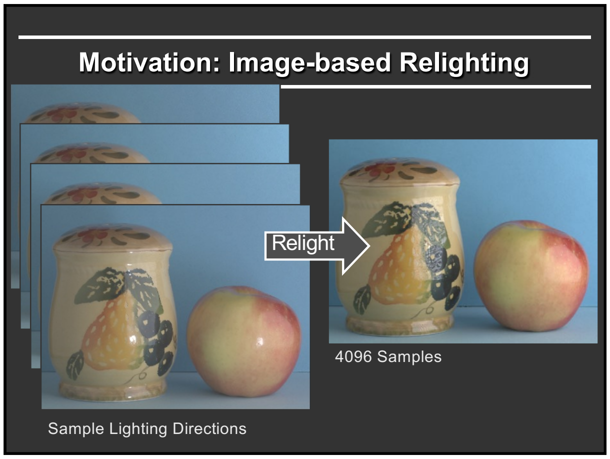

31

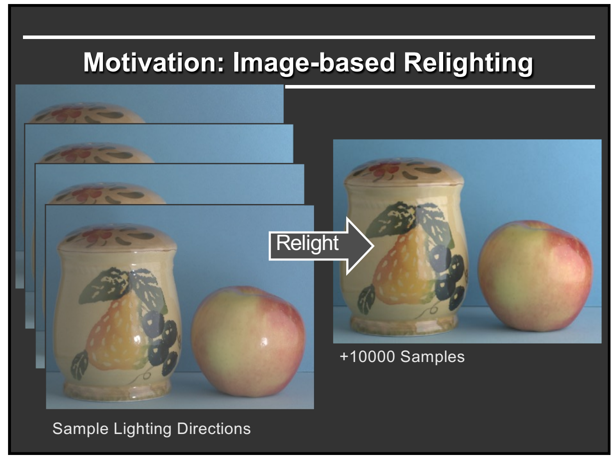

32

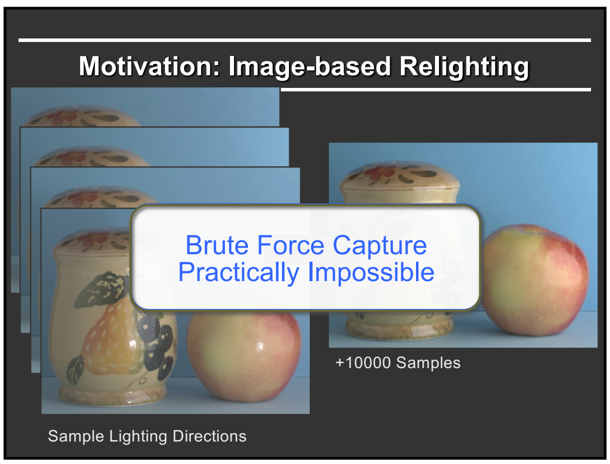

33

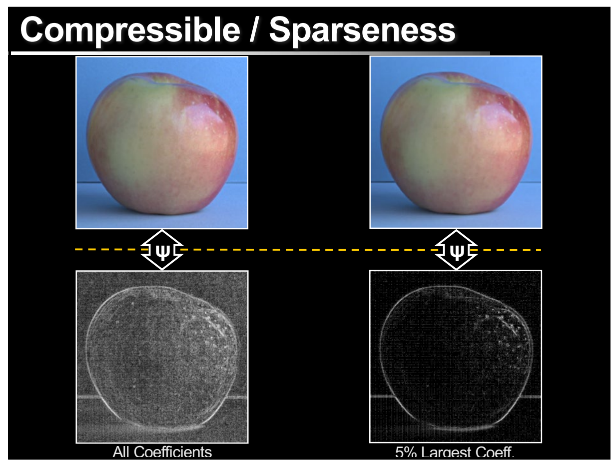

34

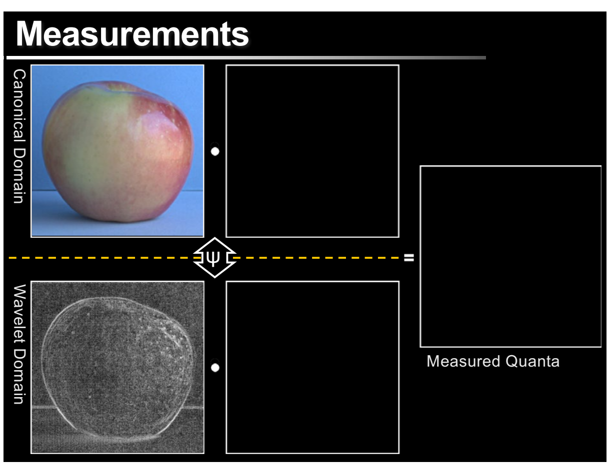

35

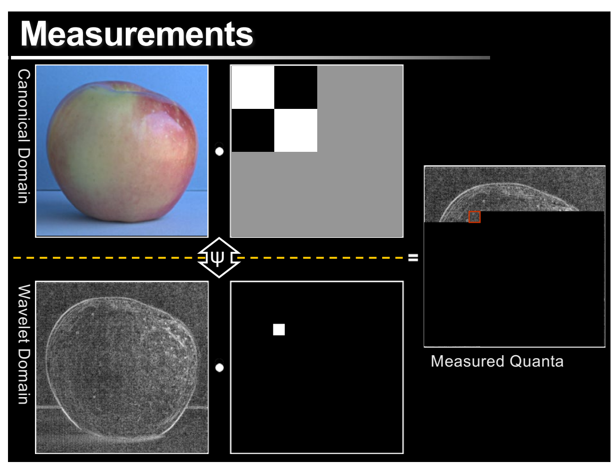

36

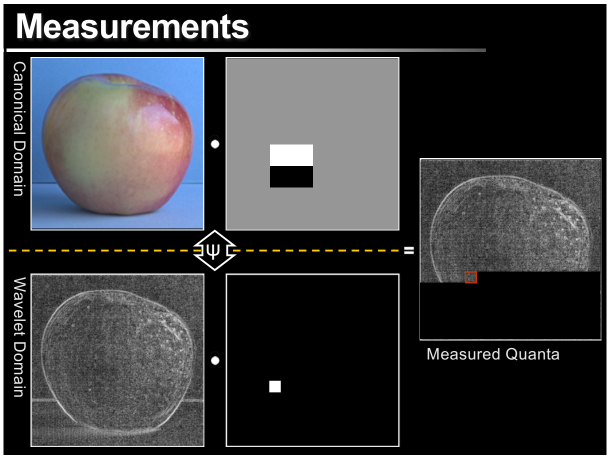

37

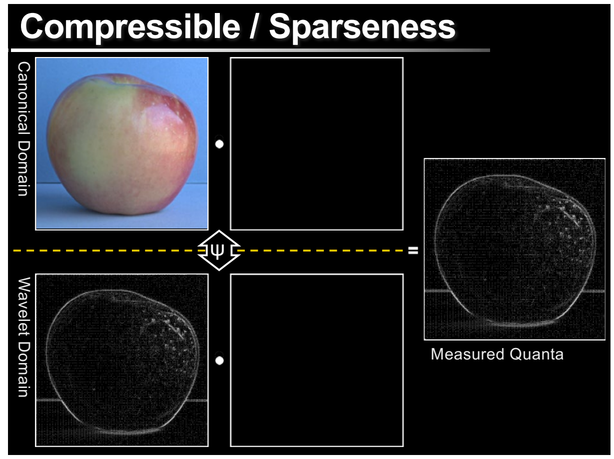

38

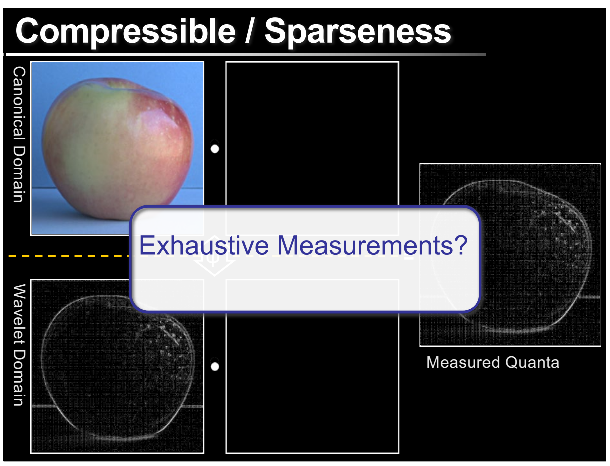

39

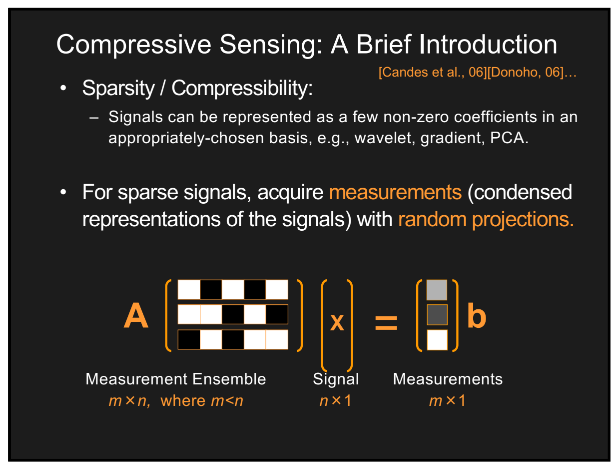

40

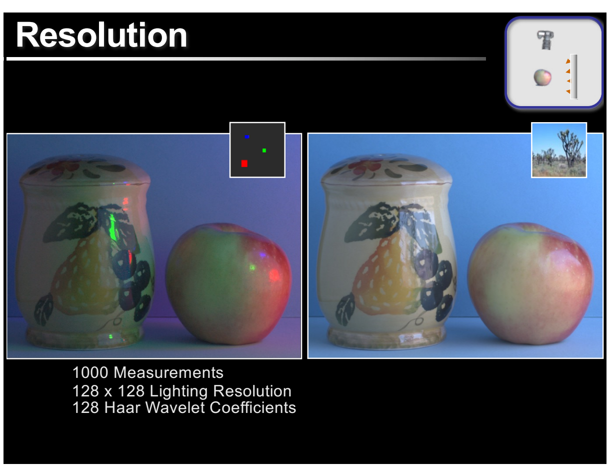

41

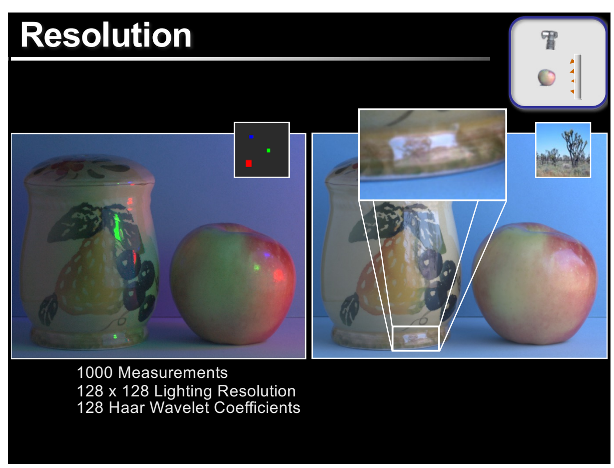

42

Results

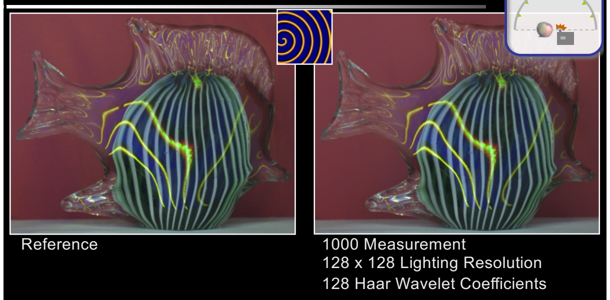

43

Results

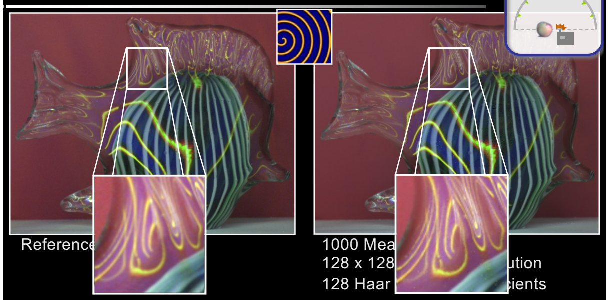

44

Results

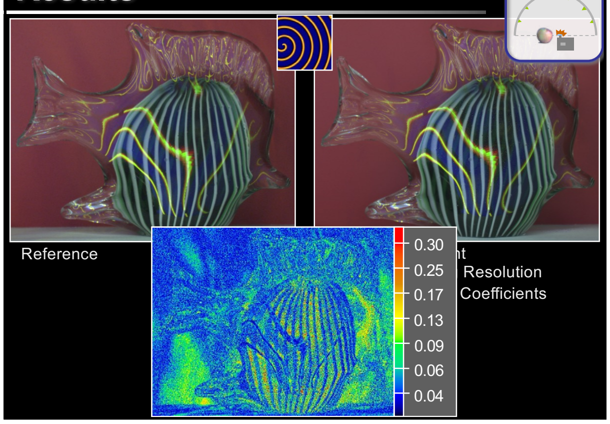

45

Applications: Sampling/Reconstruction

- Monte Carlo Rendering
- Light Transport Acquisition
- *Light Fields and Computational Photography*
- View Synthesis
- Animation/Simulation (not covered in course)

- Brief overview of these applications today, and opportunities/history for sampling/reconstruction

46

Light Field Rendering

Marc Levoy *Pat Hanrahan*

Computer Science Department
Stanford University

47

Apple's QuickTime VR

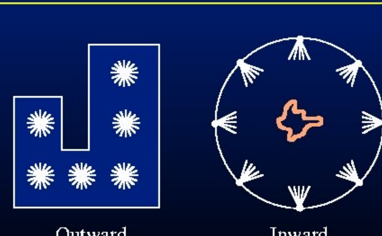

48

Generating New Views

Problem: fixed vantage point/center

One Solution: view interpolation

- Interpolating between range images (Chen and Willams, 1993)
- Correspondences and epipolar analysis (McMillan and Bishop, 1995)

-> Requires depths or correspondences:

<u>must be extracted from acquired imagery</u>

<u>relatively expensive and error-prone morph</u>

49

Light Fields

Gershun's and Moon's idea of a light field:

Radiance as a function of a ray or line: $L\,(x, y, z, \theta, \phi)$

- In "free space" (no occluders) 5D reduces to 4D
  - Exterior of the convex hull of an object
  - Interior of an environment
- Images are 2D slices
  - Insert acquired imagery
  - Extract image from a given viewpoint

50

4D Light Field

51

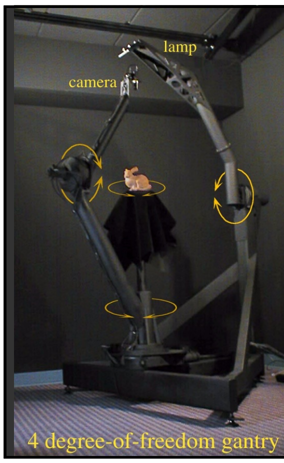

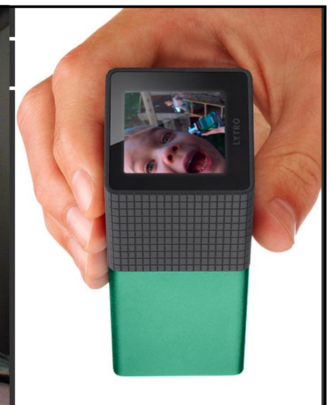

Lytro Camera

52

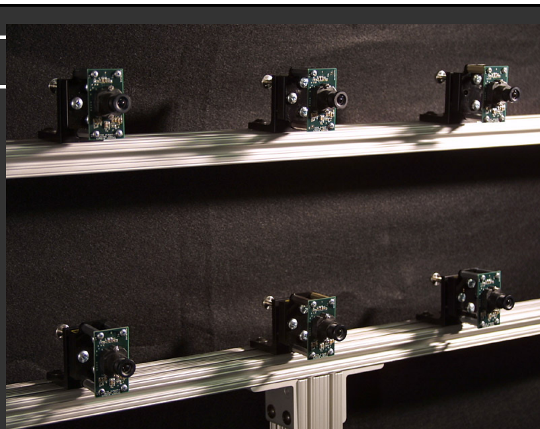

53

Light Field as a 2D Array of Image

Camera Array

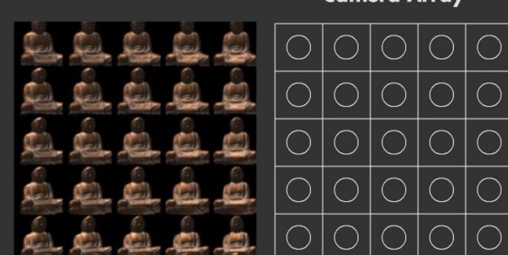

$$L(r) = L(u, v, s, t)$$

54

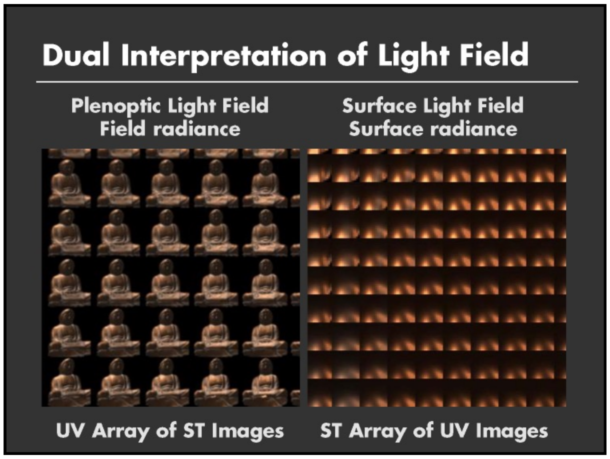

55

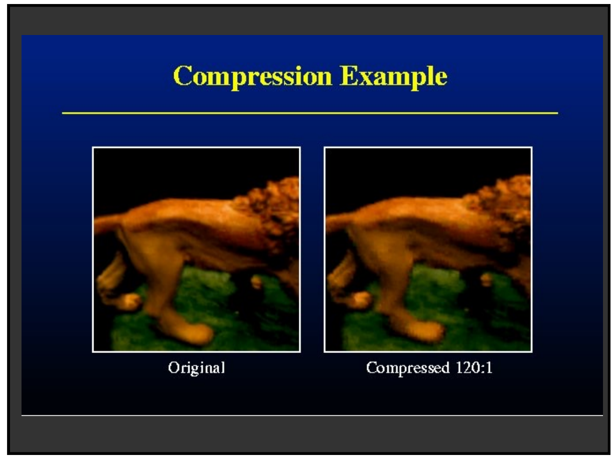

56

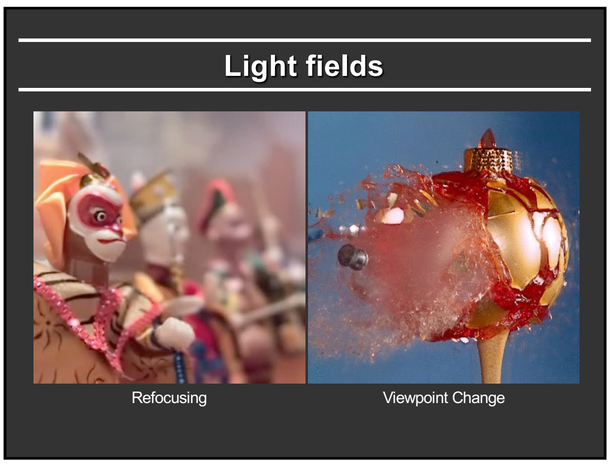

57

58

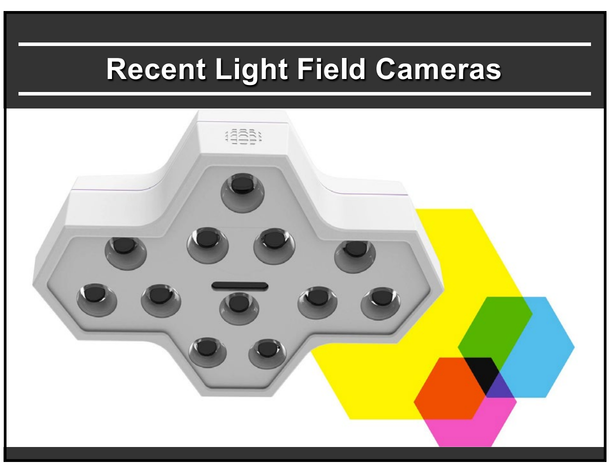

59

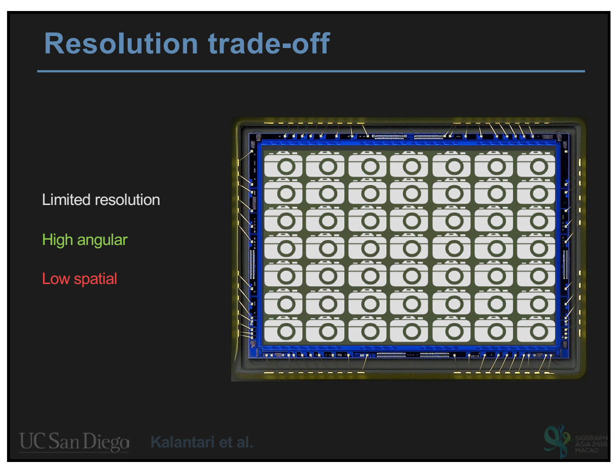

60

Solution: angular super-resolution

UC San Diego Kalantari et al.

61

Straightforward solution

- Model the process with a single CNN

UC San Diego

62

Single CNN's result

UC San Diego

63

High-level idea

- Follow the pipeline of existing techniques and break the process into two components
  Goesele et al. [2010]; Chaurasia et al. [2013]
  - Disparity estimator
  - Color predictor
- Model the components using learning
- Train both models simultaneously

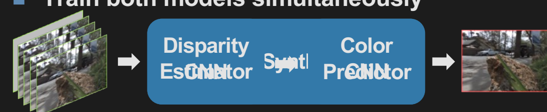

UC San Diego

64

Our result

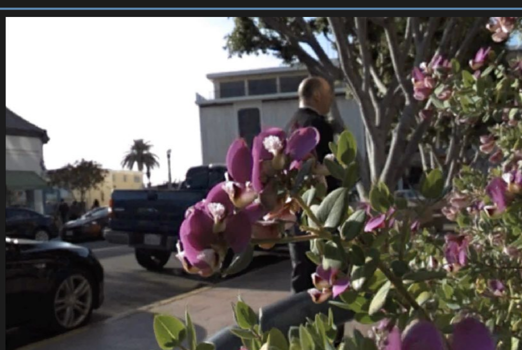

UC San Diego Kalantari et al.

65

4D RGBD Light Fields from 2D Image

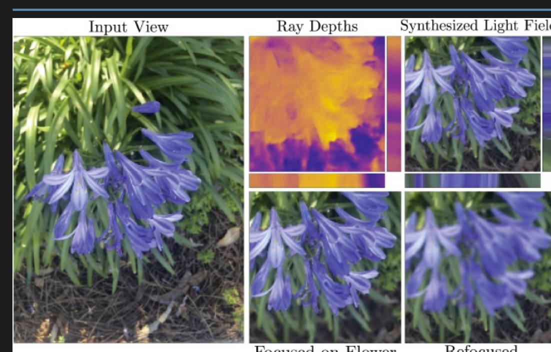

UC San Diego Srinivasan et al. ICCV 17

66

Applications: Sampling/Reconstruction

- Monte Carlo Rendering
- Light Transport Acquisition
- Light Fields and Computational Photography
- *View Synthesis*
- Animation/Simulation (not covered in course)

- Brief overview of these applications today, and opportunities/history for sampling/reconstruction

71

Light Fields with 4000x fewer views

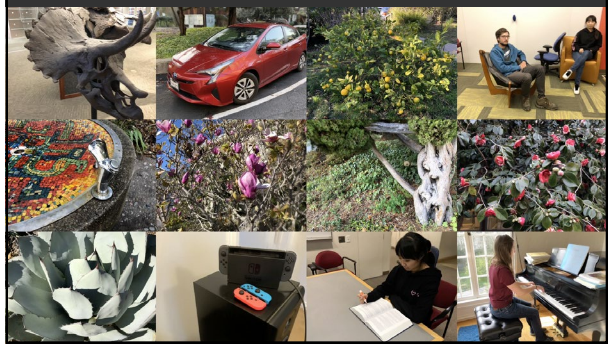

72

NEURAL RADIANCE FIELDS

73

Summary

- Brief overview of applications, some algorithms
- Will cover in greater detail in rest of course
- Biggest practical progress in Monte Carlo rendering: order of magnitude speedups
- Widely used in production, also in real-time
- Very relevant in light transport acquisition
- Sampling/Reconstruction key for light fields
- View Synthesis other major focus, huge explosion
- Many other applications: PRT, Animation, etc.

74