Sampling and Reconstruction of Visual Appearance

CSE 274 [Winter 2018], Lecture 5

Ravi Ramamoorthi

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

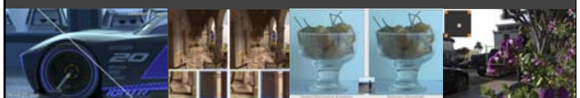

Applications: Sampling/Reconstruction

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

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

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?

Monte Carlo Path Tracing

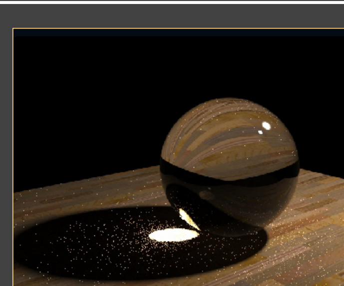

1000 paths/pixel

Jensen

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]

History

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

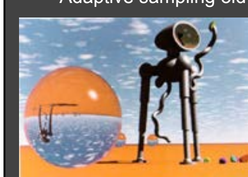

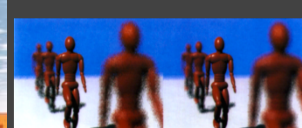

- But not very widely used… artifacts, can miss features
- After seminal papers in 87-91, not much follow on

Directional Coherence Maps

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

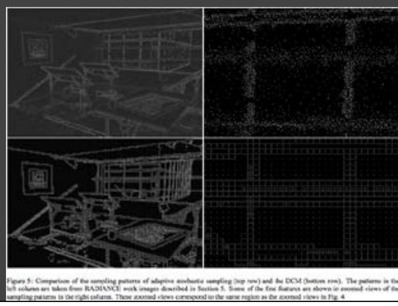

Directional Coherence Maps (Guo 98)

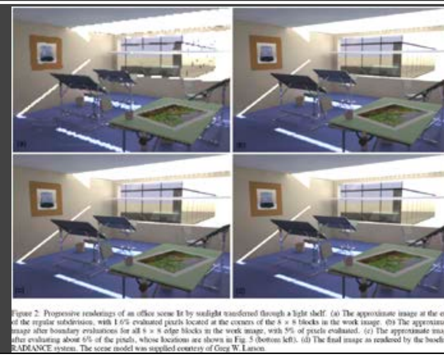

Guo 98

Resurgence (2008 -)

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

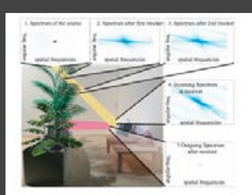

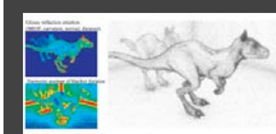

Multi-Dimensional Adaptive Sampling

- 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

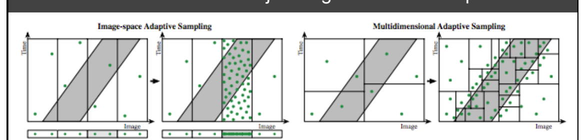

Multidimensional Adaptive Sampling

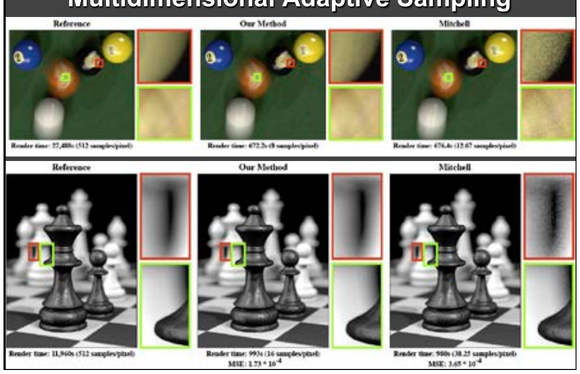

Multidimensional Adaptive Sampling

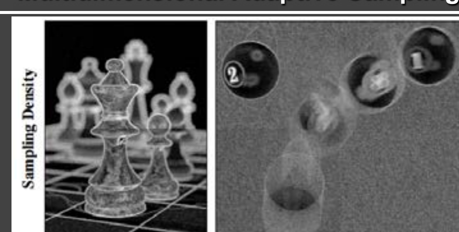

Figure 10: *Visualizations of projected sample distributions using our method for the chess scene from Figure 8 and the pool scene from Figure 7. Our adaptive sampler places samples both around high frequency image discontinuities (in focus chess piece and stationary pool ball) as well as in regions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

Multi-Dimensional Adaptive Sampling

Motion Blur and Depth of Field 32 samples per pixel

A-Priori Methods

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

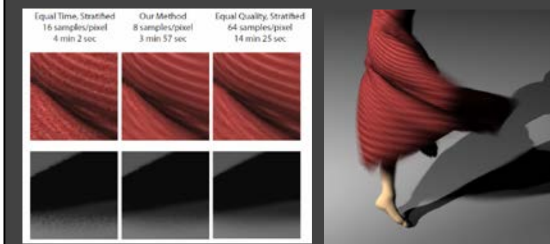

A-Posteriori Methods

- 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
- And at least one startup company

Bako et al. 17

Adaptive Wavelet Rendering

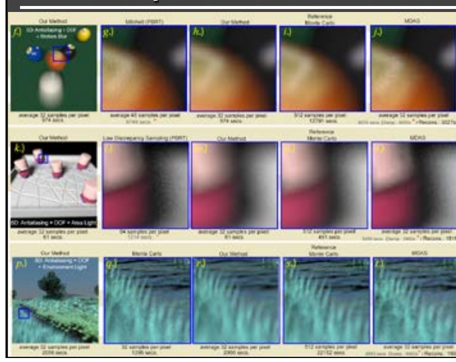

FF VIDEO

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

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)

Real-Time MAAF Video

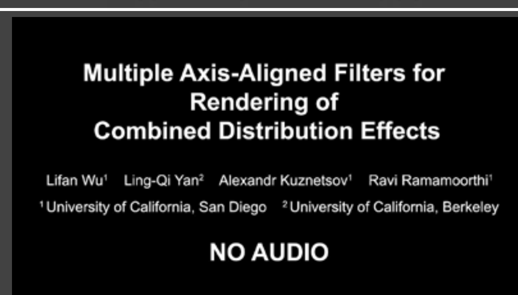

Recurrent Autoencoder Video (Chaitanya et al. 17)

Applications: Sampling/Reconstruction

- Monte Carlo Rendering (biggest application)
- *Light Transport Acquisition*
- Light Fields and Computational Photography
- Animation/Simulation (not covered in course)
- Brief overview of these applications today, and opportunities/history for sampling/reconstruction

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

Illuminate subject from many incident directions

Example Images

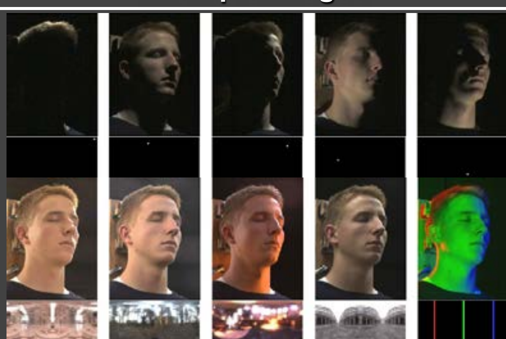

Motivation: Image-based Relighting

Sample Lighting Directions

Motivation: Image-based Relighting

Sample Lighting Directions

Motivation: Image-based Relighting

Sample Lighting Directions

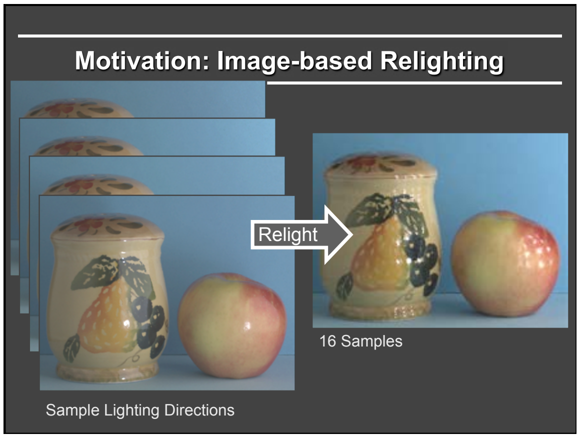
Motivation: Image-based Relighting
Relight
16 Samples
Sample Lighting Directions

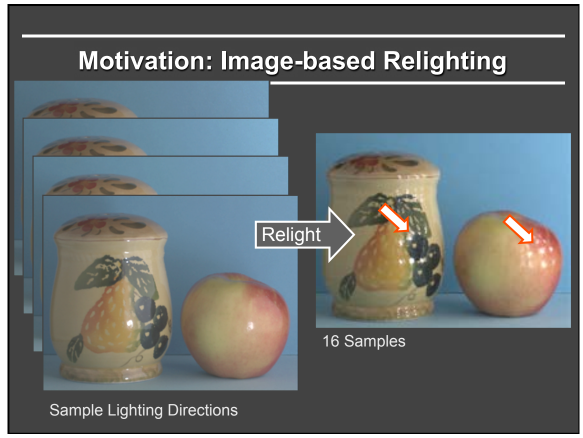
Motivation: Image-based Relighting
Relight
16 Samples
Sample Lighting Directions

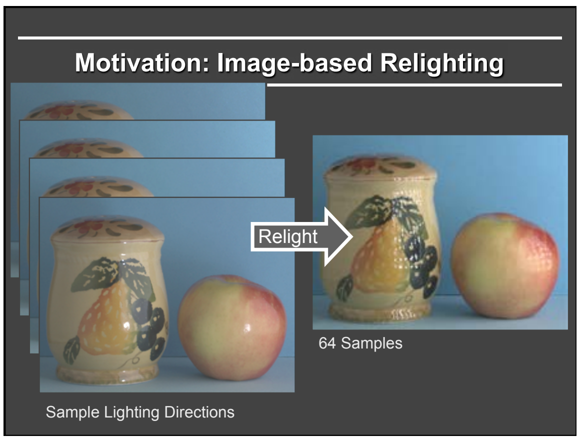
Motivation: Image-based Relighting
Relight
64 Samples
Sample Lighting Directions

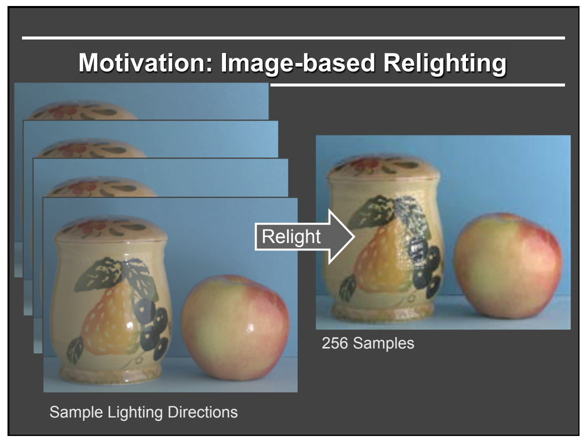
Motivation: Image-based Relighting
Relight
256 Samples
Sample Lighting Directions

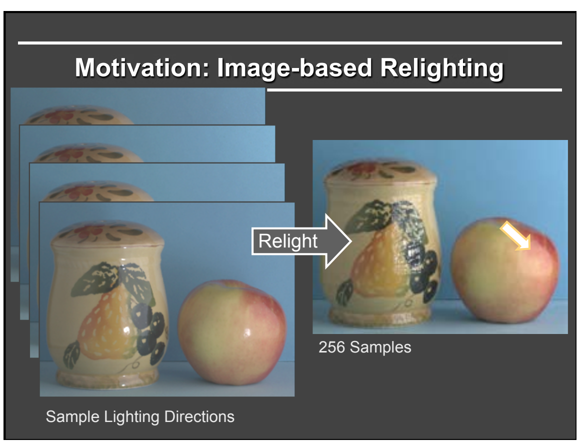
Motivation: Image-based Relighting
Relight
256 Samples
Sample Lighting Directions

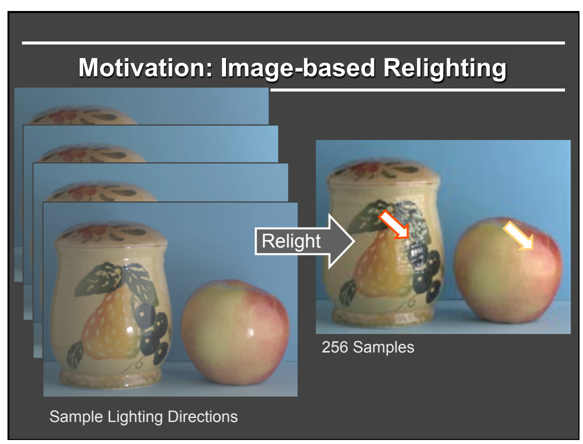
Motivation: Image-based Relighting
Relight
256 Samples
Sample Lighting Directions

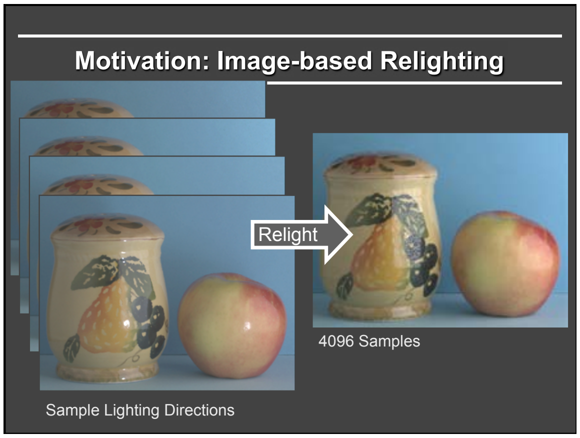
Motivation: Image-based Relighting
Relight
4096 Samples
Sample Lighting Directions

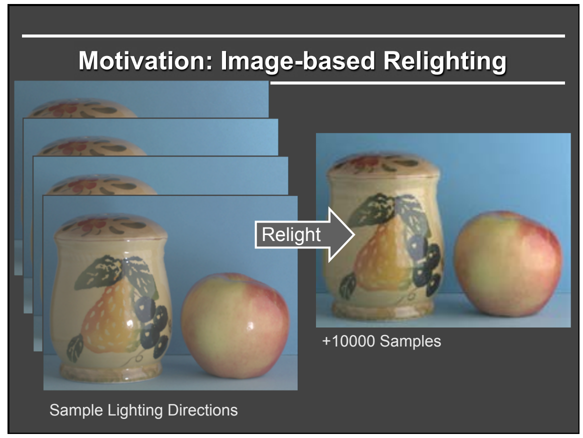
Motivation: Image-based Relighting
Relight
+10000 Samples
Sample Lighting Directions

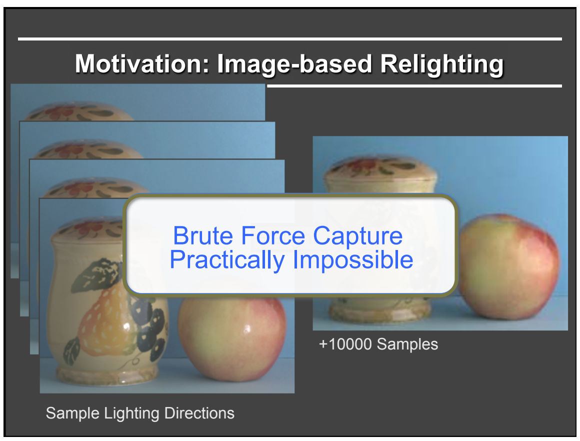
Motivation: Image-based Relighting
Brute Force Capture
Practically Impossible
+10000 Samples
Sample Lighting Directions

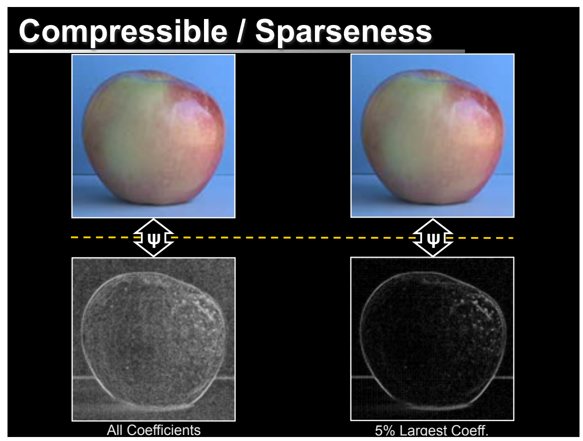
Compressible / Sparseness
All Coefficients
5% Largest Coeff.

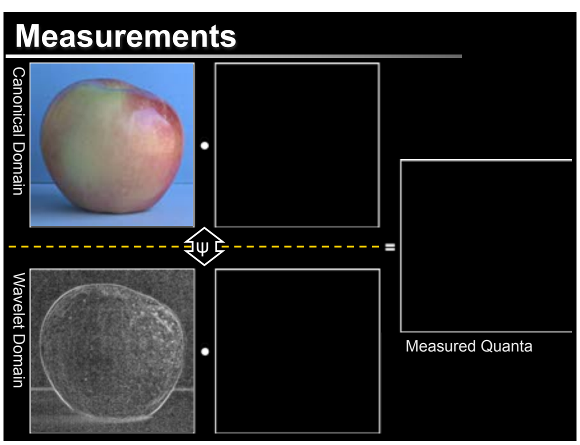
Measurements
Canonical Domain
Wavelet Domain
Measured Quanta

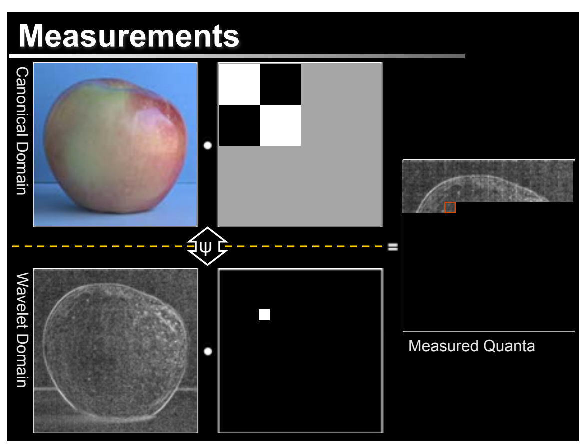
Measurements
Canonical Domain
Wavelet Domain
Measured Quanta

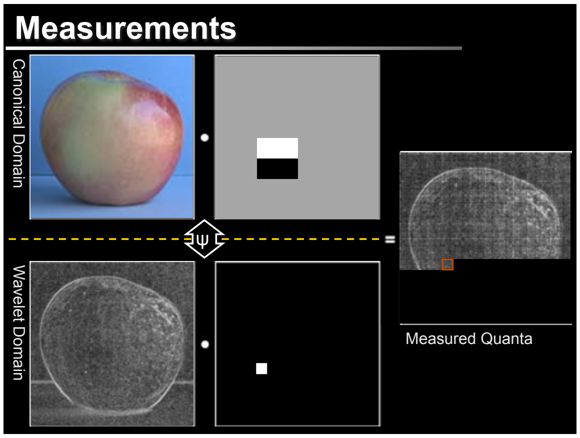
Measurements
Canonical Domain
Wavelet Domain
Ψ
=
Measured Quanta

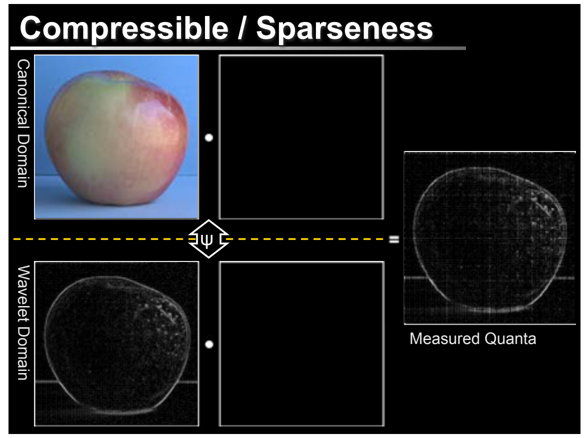
Compressible / Sparseness
Canonical Domain
Wavelet Domain
Ψ
=
Measured Quanta

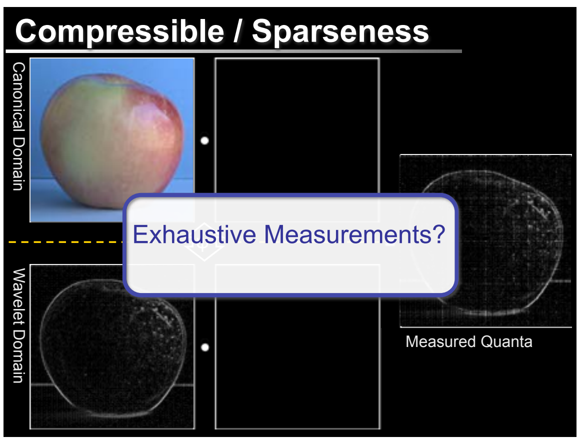
Compressible / Sparseness
Canonical Domain
Wavelet Domain
Exhaustive Measurements?
Measured Quanta

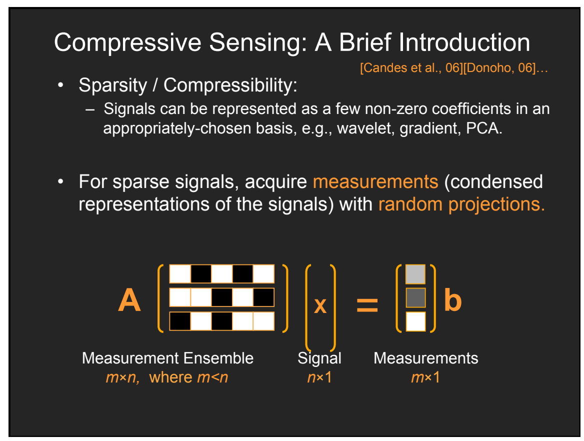
Compressive Sensing: A Brief Introduction
[Candes et al., 06][Donoho, 06]…
• Sparsity / Compressibility:
– Signals can be represented as a few non-zero coefficients in an appropriately-chosen basis, e.g., wavelet, gradient, PCA.
• For sparse signals, acquire measurements (condensed representations of the signals) with random projections.
A
x
=
b
Measurement Ensemble
m×n, where m<n
Signal
n×1
Measurements
m×1

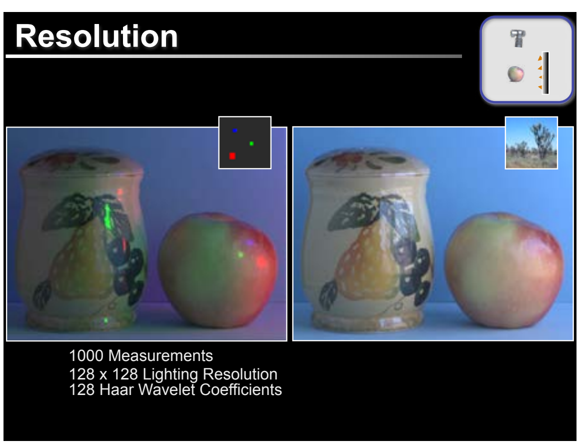
Resolution
1000 Measurements
128 x 128 Lighting Resolution
128 Haar Wavelet Coefficients

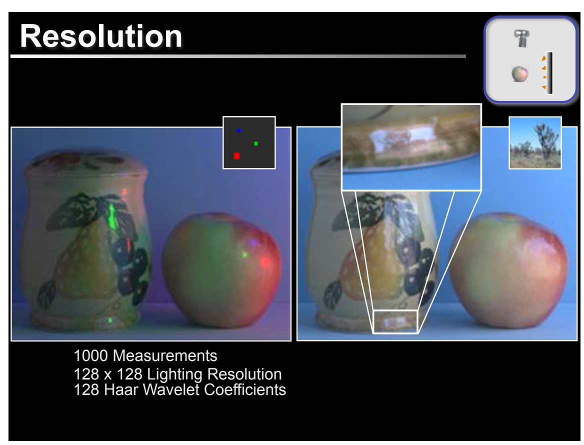
Resolution
1000 Measurements
128 x 128 Lighting Resolution
128 Haar Wavelet Coefficients

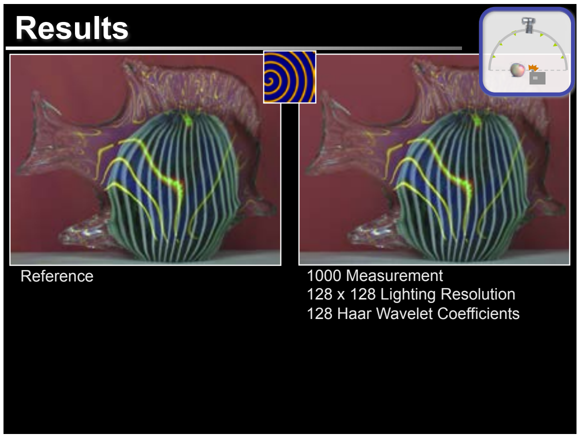

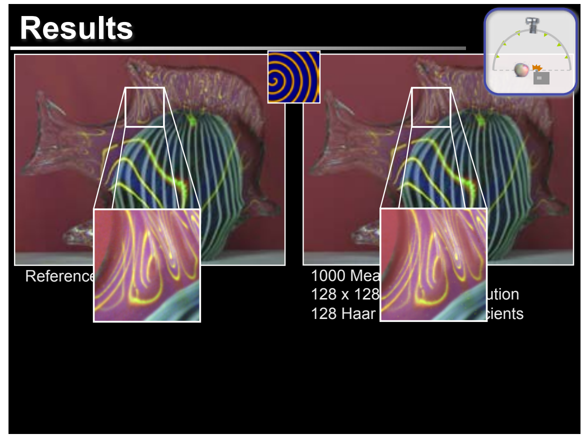

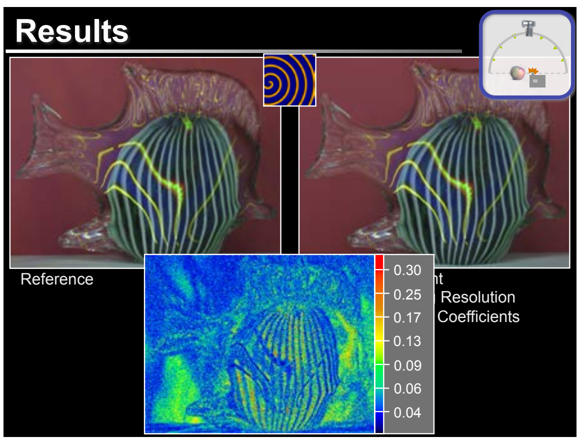

Applications: Sampling/Reconstruction

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

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

Light Field Rendering

Marc Levoy *Pat Hanrahan*

Computer Science Department
Stanford University

Apple's QuickTime VR

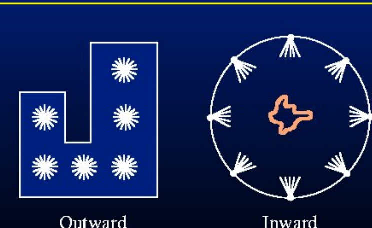

Generating New Views

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>

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

4D Light Field

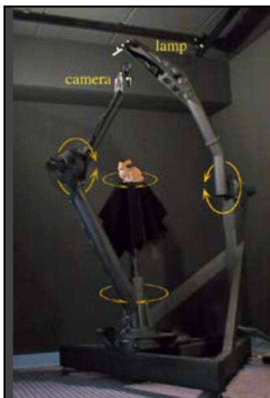

4 degree-of-freedom gantry

Lytro Camera

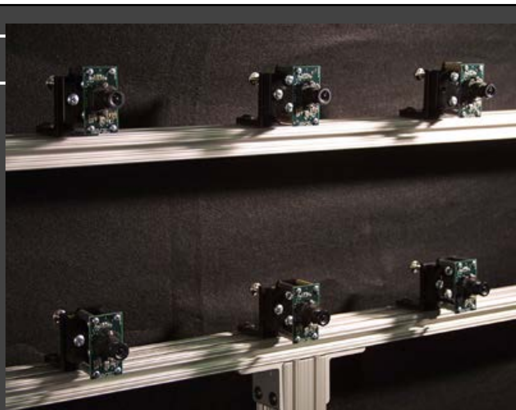

Light Field as a 2D Array of Image

Camera Array

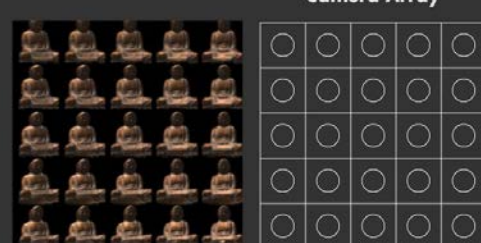

$$L(r) = L(u, v, s, t)$$

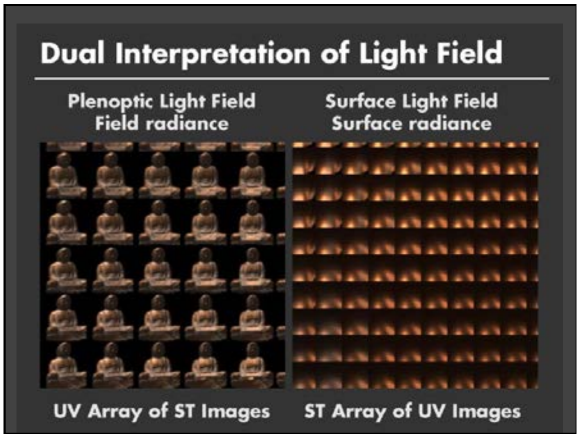
Dual Interpretation of Light Field
Plenoptic Light Field
Field radiance
Surface Light Field
Surface radiance
UV Array of ST Images
ST Array of UV Images

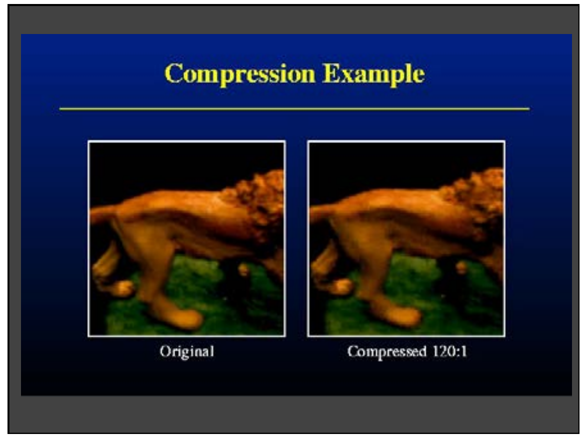
Compression Example
Original
Compressed 120:1

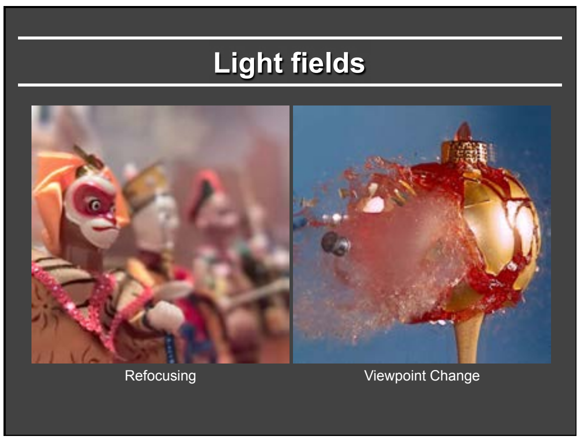
Light fields
Refocusing
Viewpoint Change


Consumer light field cameras
Lytro (first generation)
Pelican
Light
RayTrix
Lytro Illum

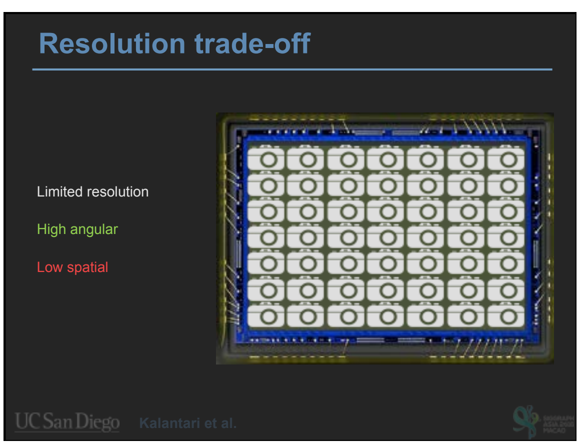
Resolution trade-off
Limited resolution
High angular
Low spatial
UC San Diego
Kalantari et al.

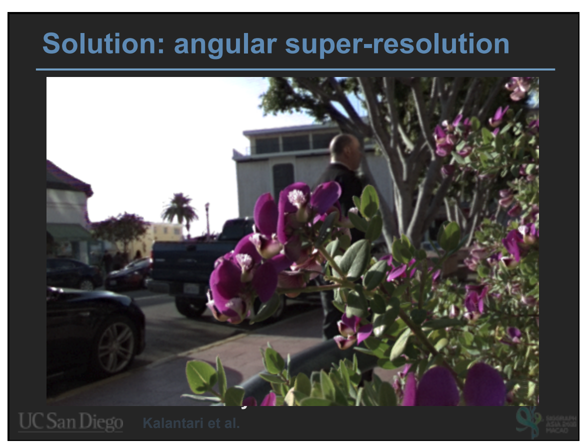
Solution: angular super-resolution
UC San Diego
Kalantari et al.

Straightforward solution

- Model the process with a single CNN

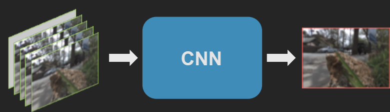

Single CNN's result

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

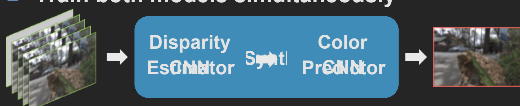

Our result

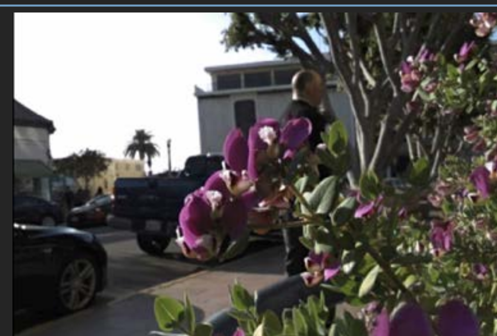

Kalantari et al.

4D RGBD Light Fields from 2D Image

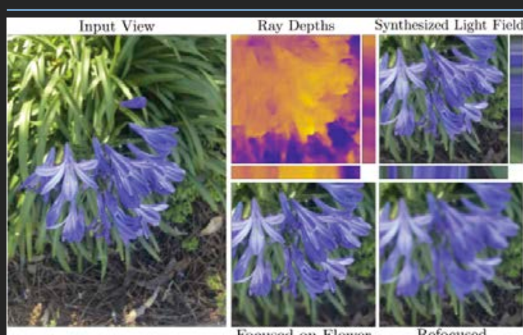

Srinivasan et al. ICCV 17

Light field video

- Consumer light field cameras limited bandwidth
- Capture low frame rate videos

Lytro Illum (3 fps video)

Wang et al. SIGGRAPH 17

Lytro video

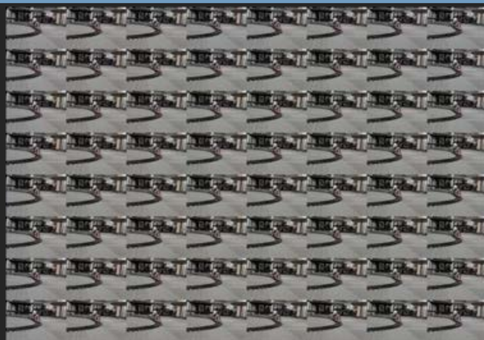

Hybrid Light Field Video System

DSLR

Lytro

30 fps

3 fps

Our result

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
- Many other applications: PRT, Animation, etc.