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

CSE 274 [Fall 2021], Lecture 9

Ravi Ramamoorthi

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

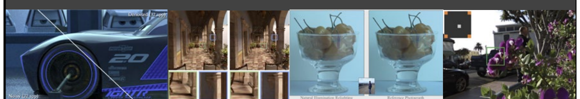

Applications

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

- Introduce concepts of sparsity, coherence, compressive sensing for reconstruction

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

Illuminate subject from many incident directions

Example Images

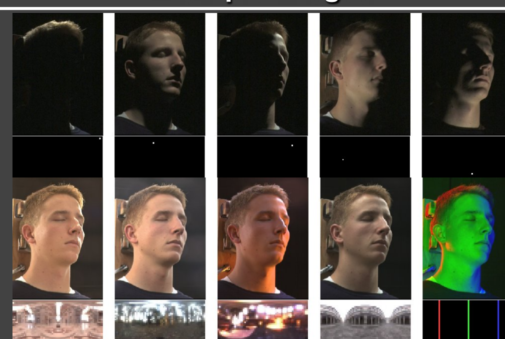

Motivation: Image-based Relighting

Sample Lighting Directions

Motivation: Image-based Relighting

Sample Lighting Directions

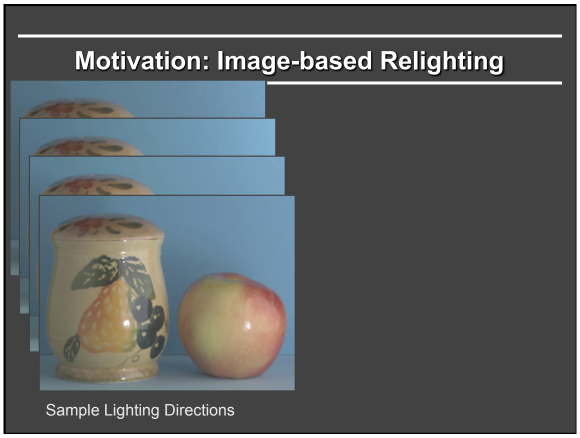
Motivation: Image-based Relighting
Sample Lighting Directions

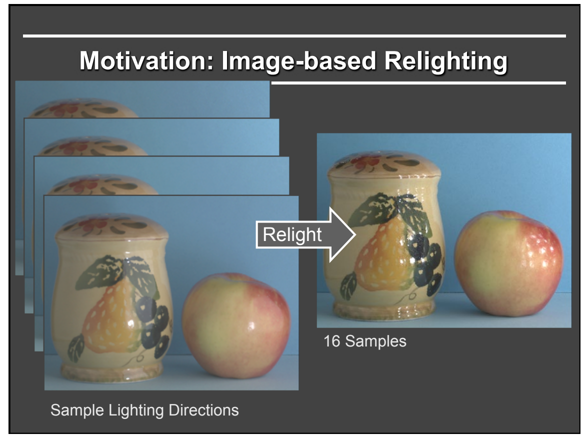
Motivation: Image-based Relighting
Relight
16 Samples
Sample Lighting Directions

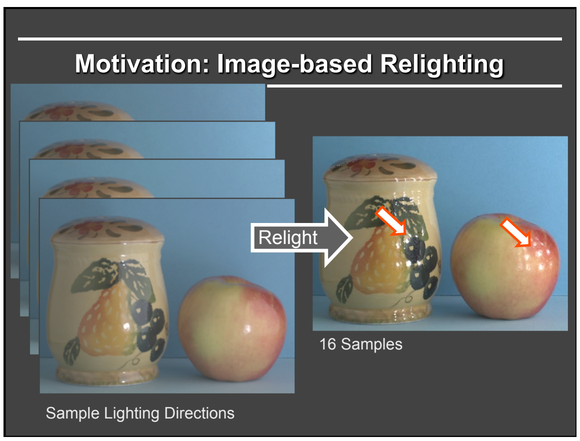
Motivation: Image-based Relighting
Relight
16 Samples
Sample Lighting Directions

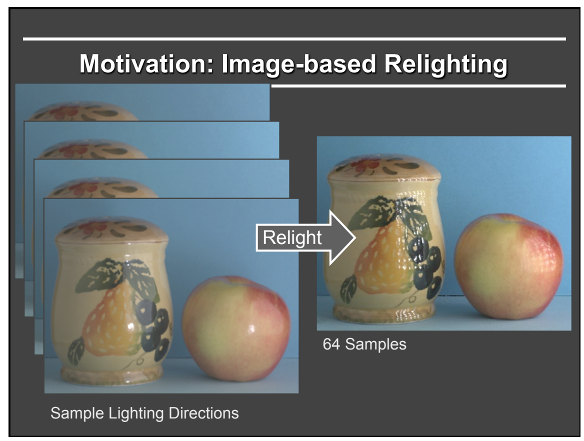
Motivation: Image-based Relighting
Relight
64 Samples
Sample Lighting Directions

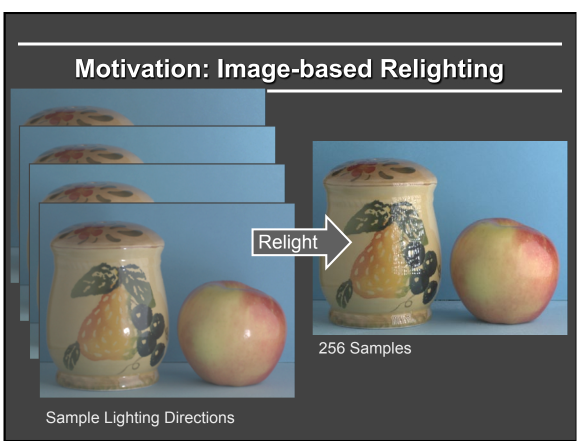
Motivation: Image-based Relighting
Relight
256 Samples
Sample Lighting Directions

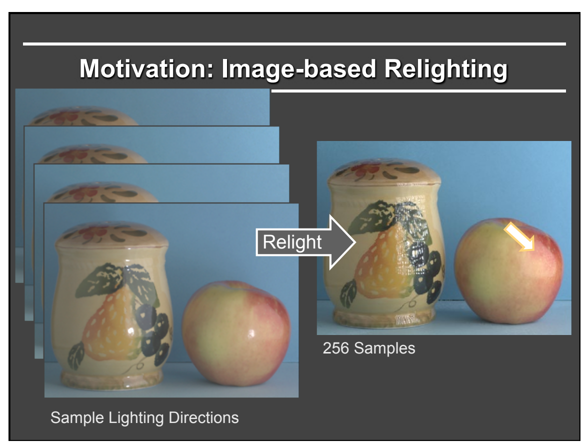
Motivation: Image-based Relighting
Relight
256 Samples
Sample Lighting Directions

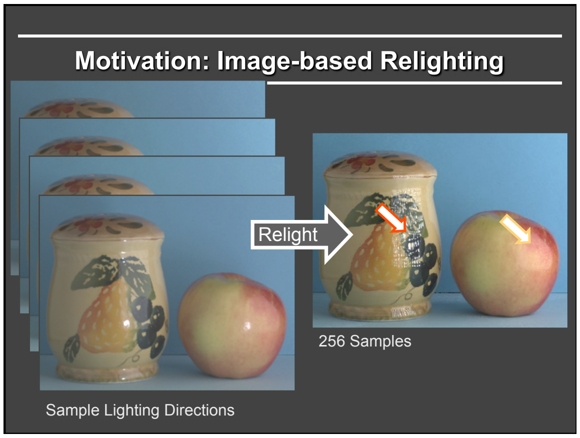
Motivation: Image-based Relighting
Relight
256 Samples
Sample Lighting Directions

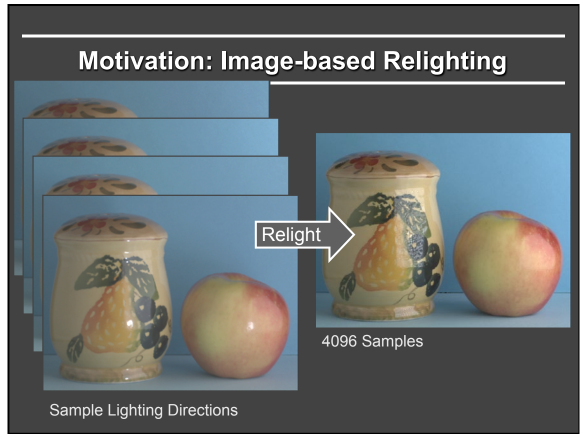
Motivation: Image-based Relighting
Relight
4096 Samples
Sample Lighting Directions

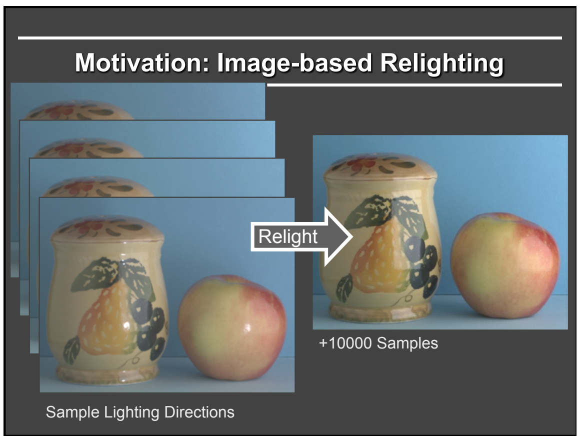
Motivation: Image-based Relighting
Relight
+10000 Samples
Sample Lighting Directions

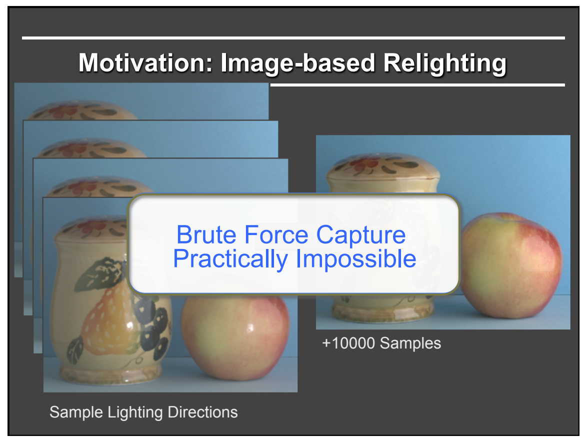
Motivation: Image-based Relighting
Brute Force Capture
Practically Impossible
+10000 Samples
Sample Lighting Directions

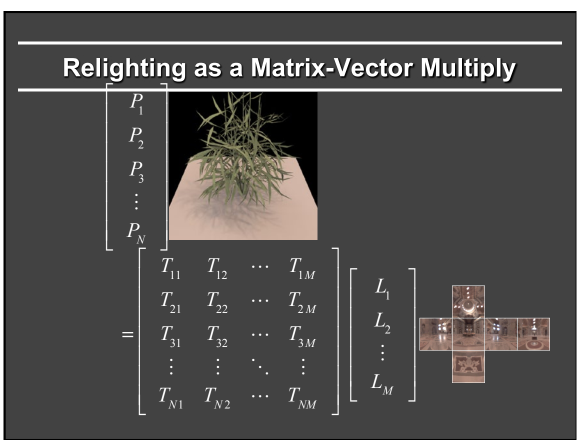
Relighting as a Matrix-Vector Multiply
P_1 P_2 P_3 $\vdots$ P_N
$=$
T_{11} T_{12} $\cdots$ T_{1M}
T_{21} T_{22} $\cdots$ T_{2M}
T_{31} T_{32} $\cdots$ T_{3M}
T_{N1} T_{N2} $\cdots$ T_{NM}
L_1 L_2 $\vdots$ L_M

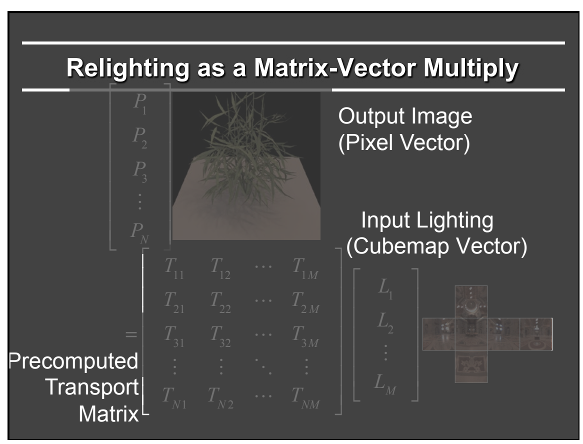
Relighting as a Matrix-Vector Multiply
Output Image
(Pixel Vector)
Input Lighting
(Cubemap Vector)
Precomputed
Transport
Matrix

Matrix Columns (Images)

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

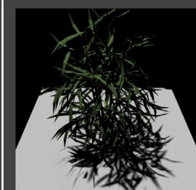

(Pre)compute: Ray-Trace Image Cols

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

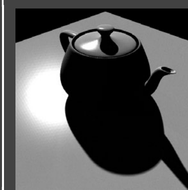

(Pre)compute 2: Rasterize Matrix Rows

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

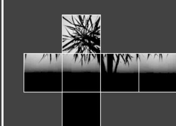

Outline

- *Matrix Row-Column Sampling (Many Lights)*

(clustering for matrix completion of light transport)
- Compressive Sensing for Light Transport
- Matrix Completion

Hasan, Pellacini, Bala SIGGRAPH 07

Complex Illumination: A Challenge

Conversion to Many Lights

- Area, indirect, sun/sky

Courtesy Walter et al., Lightcuts, SIGGRAPH 05/06

A Matrix Interpretation

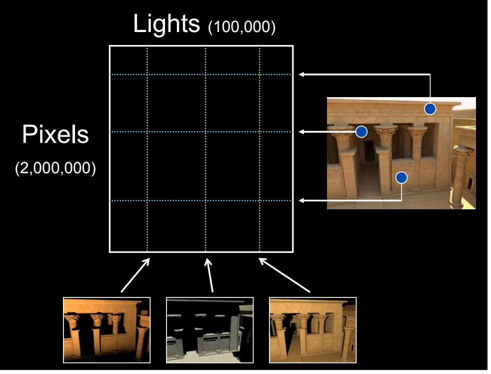

Problem Statement

- Compute sum of columns

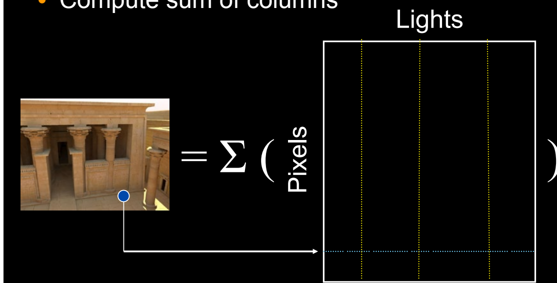

- Note: We don't have the matrix data

Image as a Weighted Column Sum

- The following is possible:

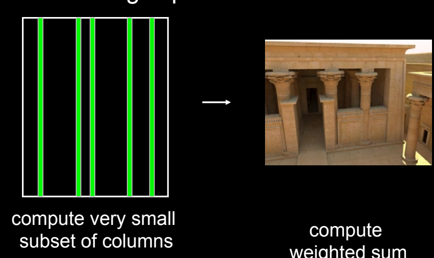

- Use rows to choose a good set of columns!

Exploration and Exploitation

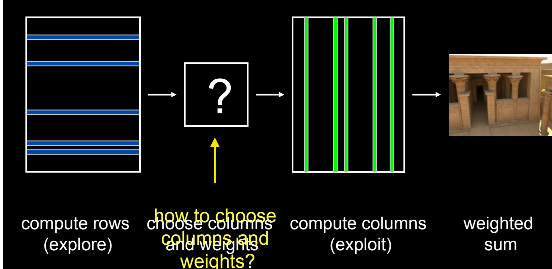

Reduced Matrix

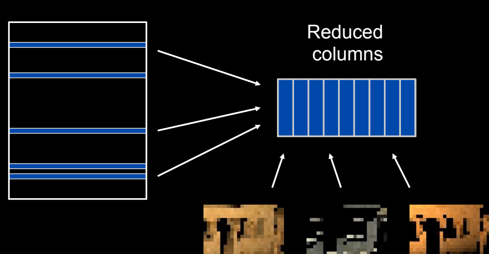

Clustering Approach

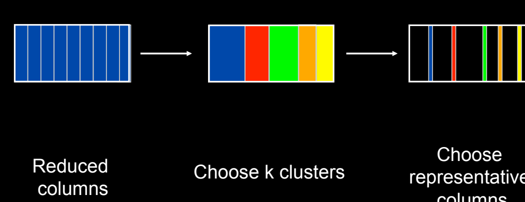

Reduced to Full

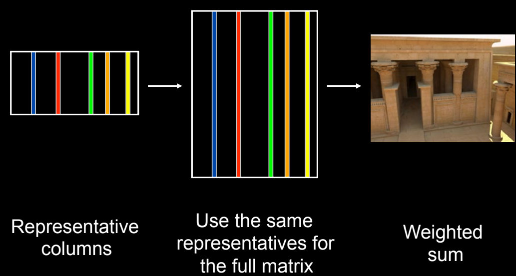

Full Algorithm

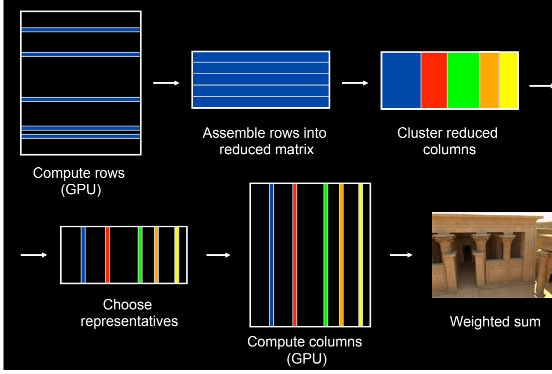

Results

- We show 5 scenes:

- Show reference and 5x difference image
- All scenes have 100,000+ lights
- Timings
  - NVidia GeForce 8800 GTX
  - Light / surface sample creation not included

Results: Kitchen

- 388k polygons
- Mostly indirect illumination
- Glossy surfaces
- Indirect shadows

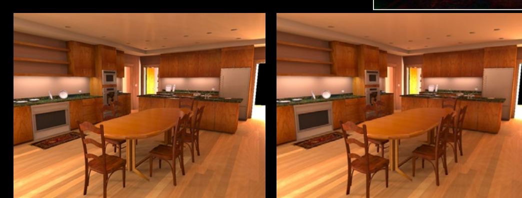

Our result: 13.5 sec (432 rows + 864 columns)

Reference: 13 min (using all 100k lights)

Results: Temple

- 2.1m polygons
- Mostly indirect & sky illumination
- Indirect shadows

Our result: 16.9 sec (300 rows + 900 columns)

Reference: 20 min (using all 100k lights)

Results: Trees

- 328k polygons
- Complex incoherent geometry

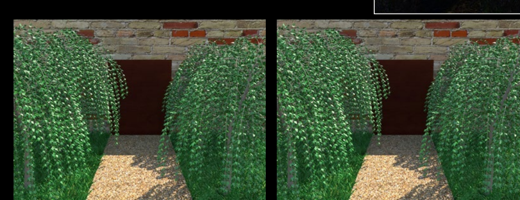

Our result: 2.9 sec (100 rows + 200 columns)

Reference: 14 min (using all 100k lights)

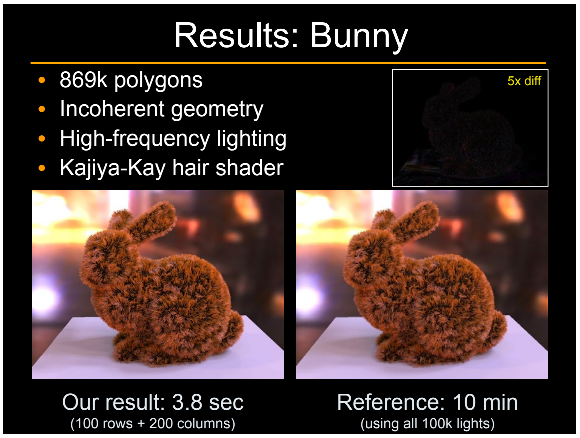
Results: Bunny
869k polygons
Incoherent geometry
High-frequency lighting
Kajiya-Kay hair shader
5x diff
Our result: 3.8 sec
(100 rows + 200 columns)
Reference: 10 min
(using all 100k lights)

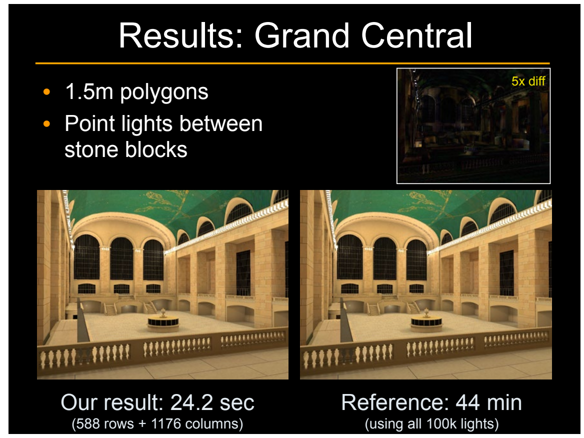
Results: Grand Central
1.5m polygons
Point lights between stone blocks
5x diff
Our result: 24.2 sec
(588 rows + 1176 columns)
Reference: 44 min
(using all 100k lights)

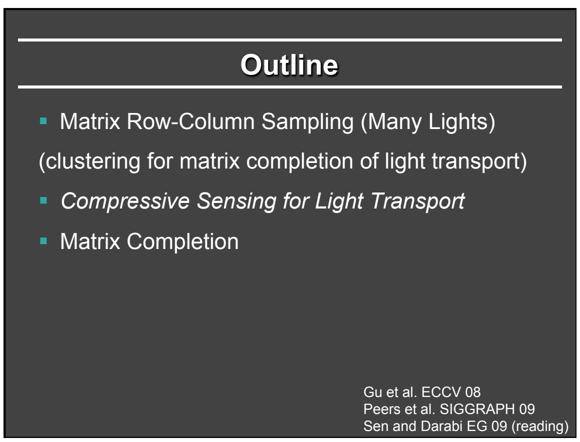
Outline
Matrix Row-Column Sampling (Many Lights)
(clustering for matrix completion of light transport)
Compressive Sensing for Light Transport
Matrix Completion
Gu et al. ECCV 08
Peers et al. SIGGRAPH 09
Sen and Darabi EG 09 (reading)

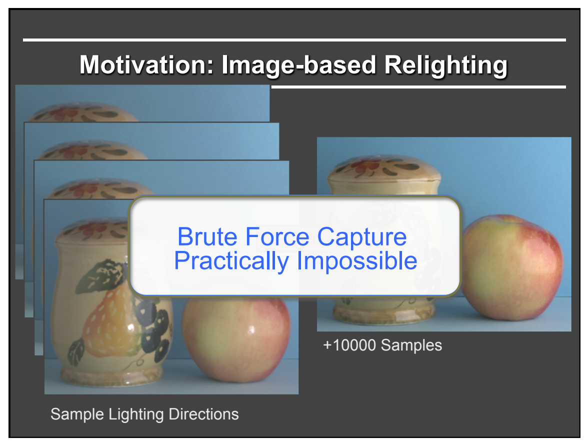
Motivation: Image-based Relighting
Brute Force Capture
Practically Impossible
+10000 Samples
Sample Lighting Directions

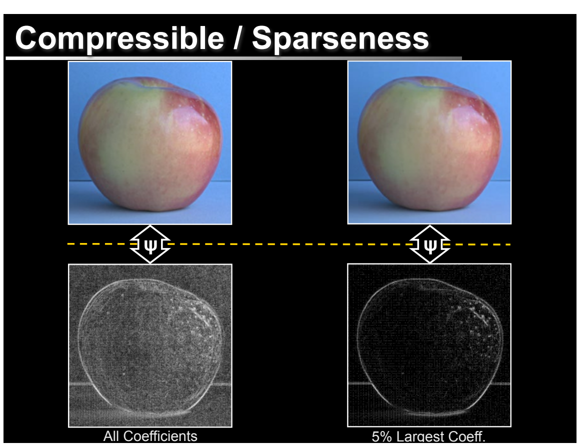
Compressible / Sparseness
Ψ
Ψ
All Coefficients
5% Largest Coeff.

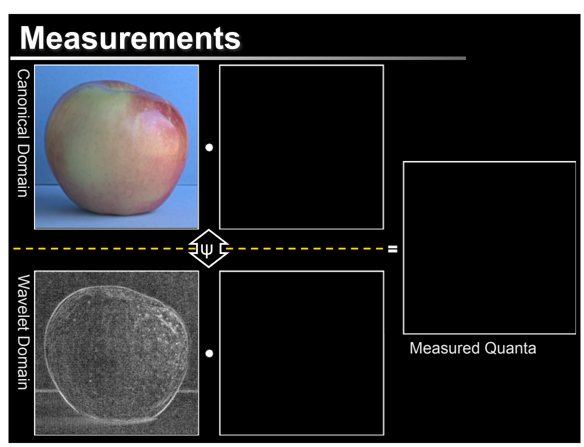
Measurements
Canonical Domain
Wavelet Domain
Ψ
Measured Quanta

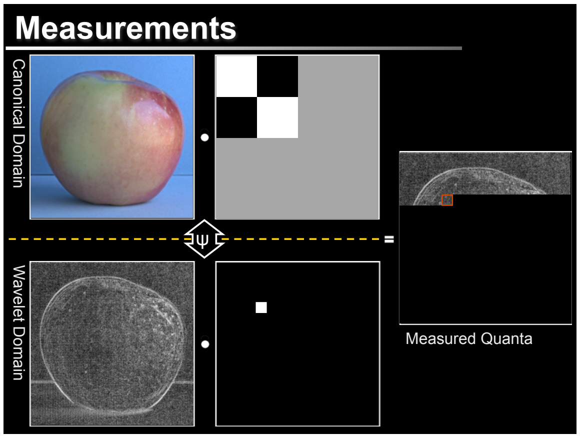
Measurements
Canonical Domain
Wavelet Domain
Ψ
=
Measured Quanta

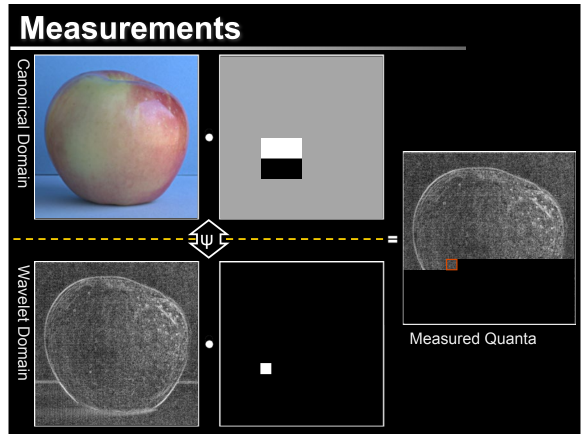
Measurements
Canonical Domain
Wavelet Domain
Ψ
=
Measured Quanta

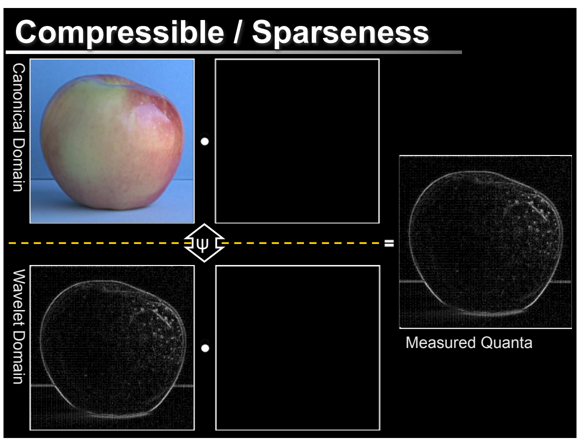
Compressible / Sparseness
Canonical Domain
Wavelet Domain
Ψ
=
Measured Quanta

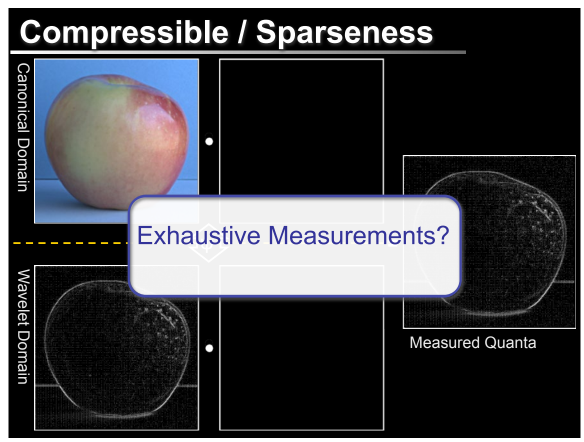
Compressible / Sparseness
Canonical Domain
Wavelet Domain
Exhaustive Measurements?
Measured Quanta

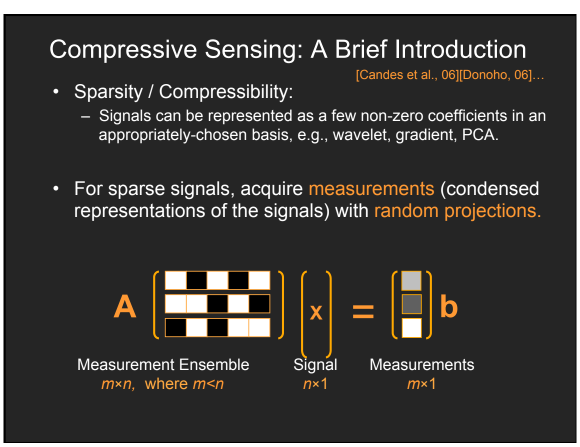
Compressive Sensing: A Brief Introduction
[Candes et al., 06][Donoho, 06]…
• Sparsity / Compressibility:
– Signals can be represented as a few non-zero coefficients in an appropriately-chosen basis, e.g., wavelet, gradient, PCA.
• For sparse signals, acquire measurements (condensed representations of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

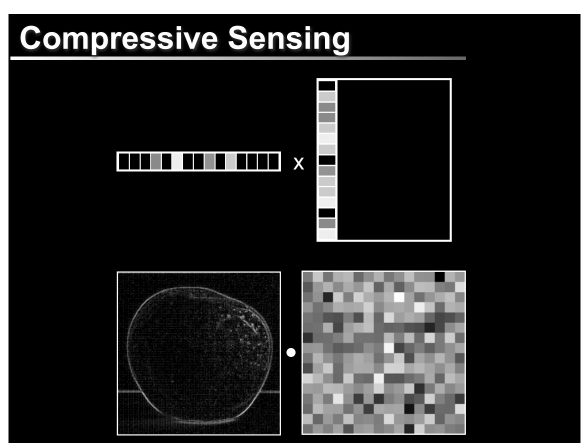
Compressive Sensing
x

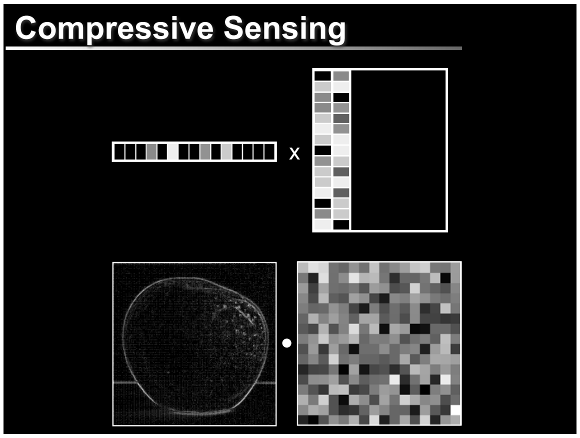
Compressive Sensing
x

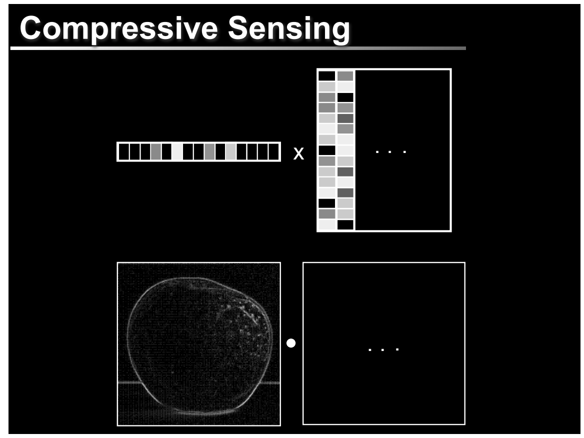
Compressive Sensing
x
. . .
. . .

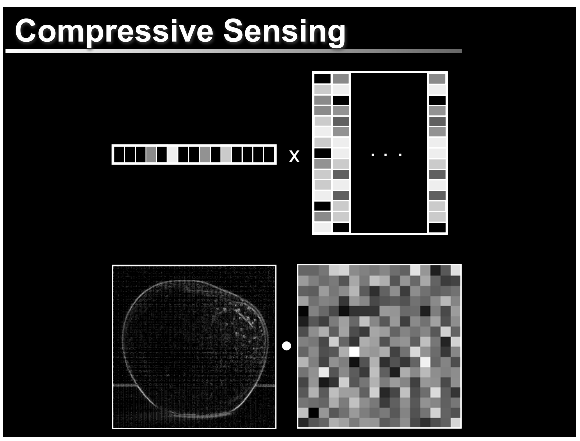
Compressive Sensing
x
. . .

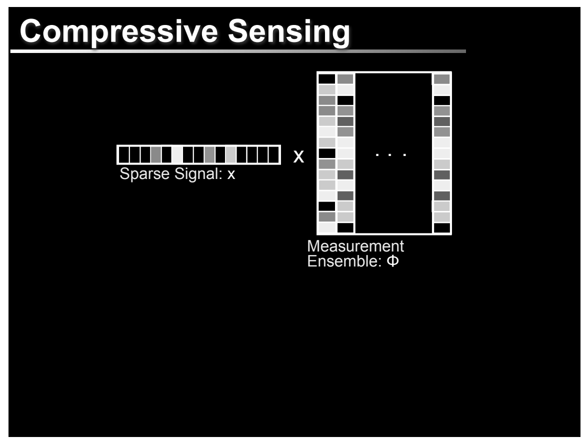
Compressive Sensing
x
. . .
Sparse Signal: x
Measurement
Ensemble: Φ

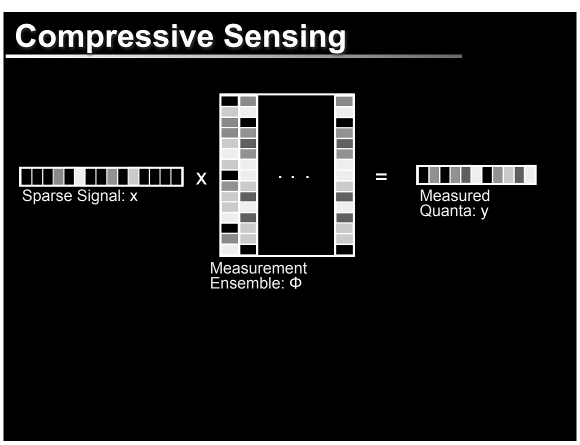
Compressive Sensing
x
. . .
=
Sparse Signal: x
Measured
Quanta: y
Measurement
Ensemble: Φ

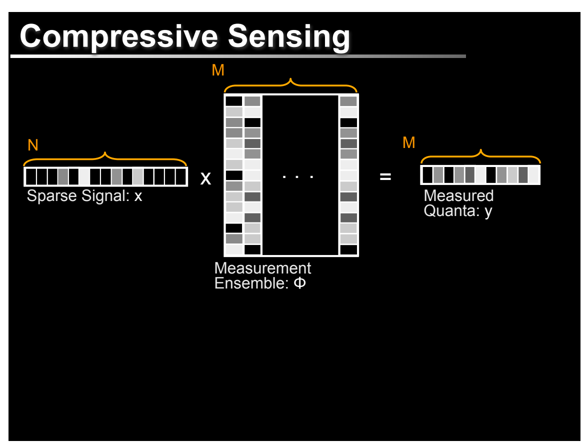
Compressive Sensing
M
N
M
x
. . .
=
Sparse Signal: x
Measured
Quanta: y
Measurement
Ensemble: Φ

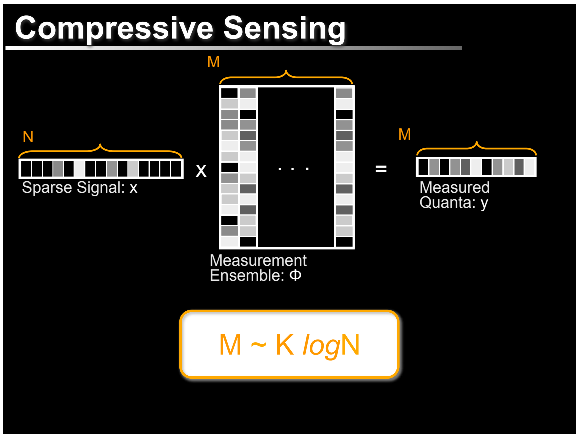
Compressive Sensing
N
M
M
Sparse Signal: x
x
. . .
=
Measured
Quanta: y
Measurement
Ensemble: Φ
M ~ K logN

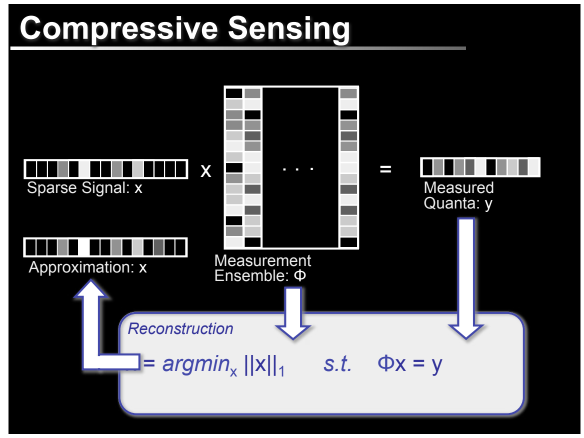
Compressive Sensing
Sparse Signal: x
x
. . .
=
Measured
Quanta: y
Approximation: x
Measurement
Ensemble: Φ
Reconstruction
= argminx ||x||1 s.t. Φx = y

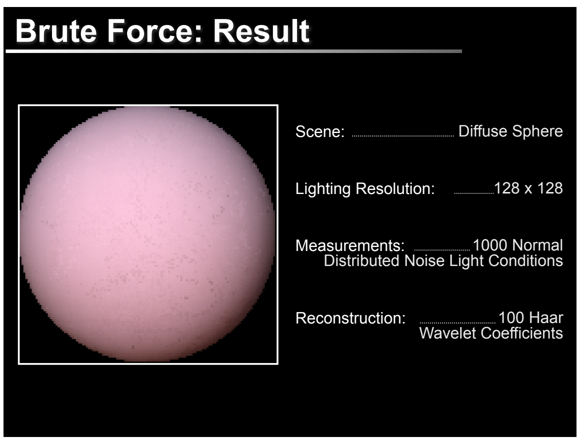
Brute Force: Result
Scene: Diffuse Sphere
Lighting Resolution: 128 x 128
Measurements: 1000 Normal
Distributed Noise Light Conditions
Reconstruction: 100 Haar
Wavelet Coefficients

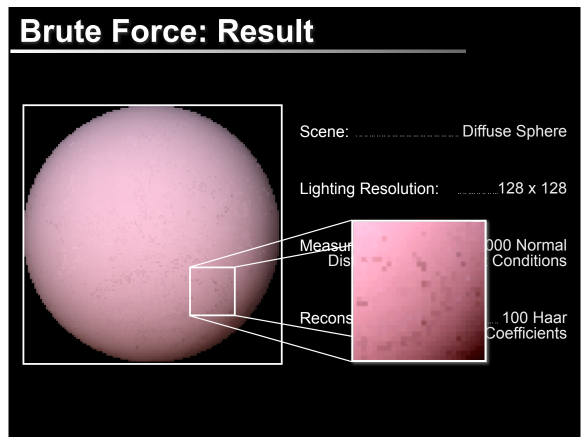
Brute Force: Result
Scene: Diffuse Sphere
Lighting Resolution: 128 x 128

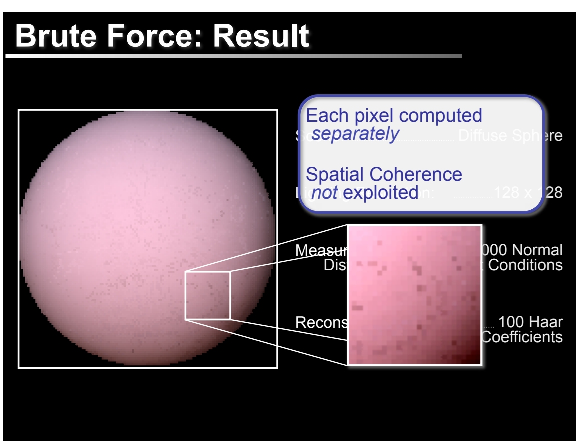
Brute Force: Result
Each pixel computed
separately
Spatial Coherence
not exploited

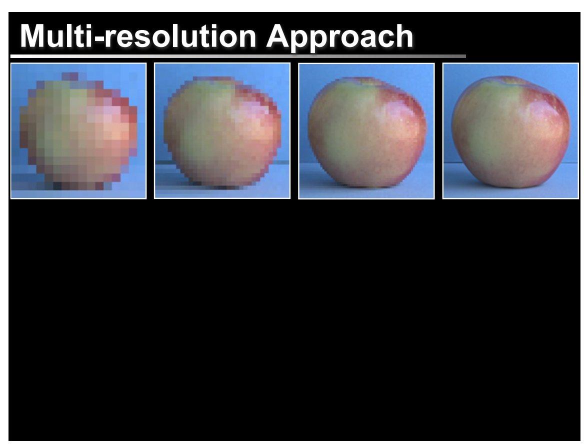
Multi-resolution Approach

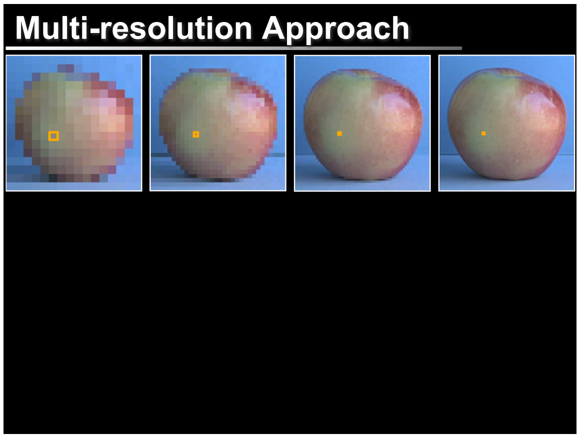
Multi-resolution Approach

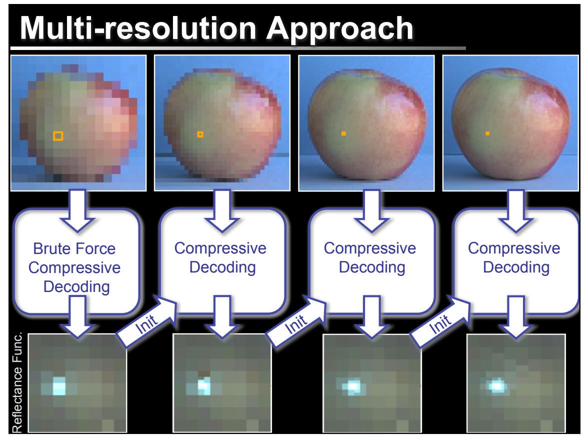
Multi-resolution Approach
Brute Force Compressive Decoding
Compressive Decoding
Compressive Decoding
Compressive Decoding
Init
Init
Init
Reflectance Func.

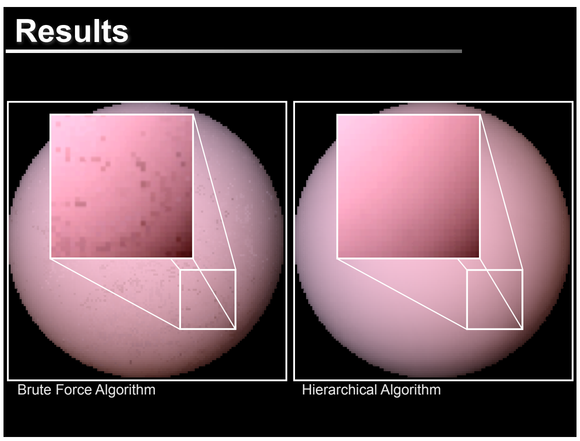
Results
Brute Force Algorithm
Hierarchical Algorithm

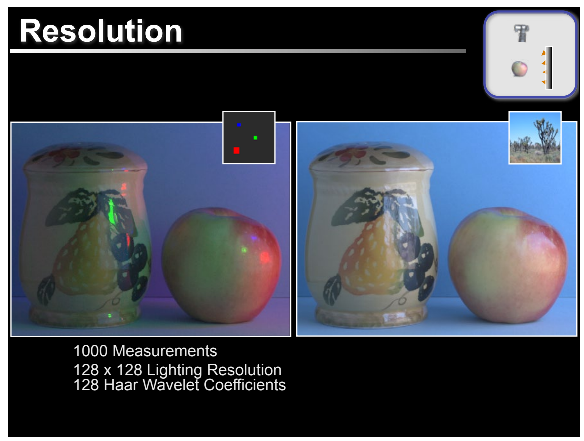
Resolution
1000 Measurements
128 x 128 Lighting Resolution
128 Haar Wavelet Coefficients

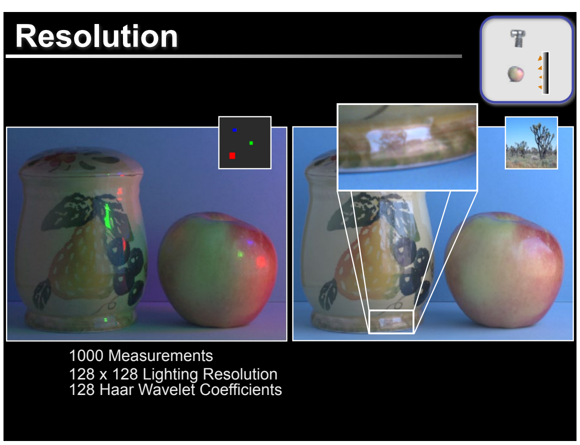
Resolution
1000 Measurements
128 x 128 Lighting Resolution
128 Haar Wavelet Coefficients

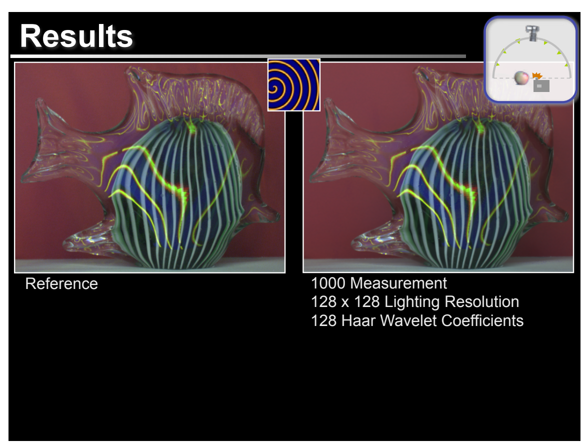
Results
Reference
1000 Measurement
128 x 128 Lighting Resolution
128 Haar Wavelet Coefficients

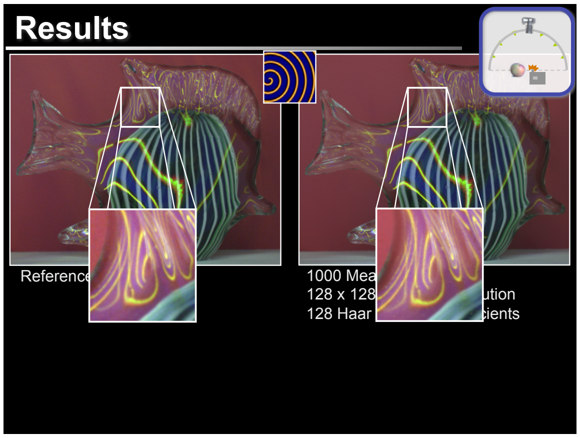
Results
Reference
1000 Mea
128 x 128
128 Haar
ution
ients

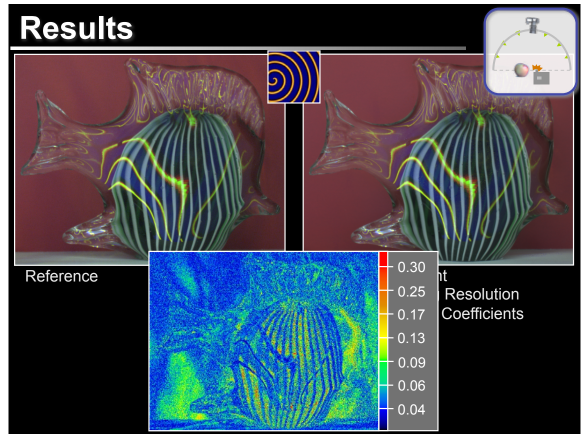
Results
Reference
0.30
0.25
0.17
0.13
0.09
0.06
0.04
nt
Resolution
Coefficients

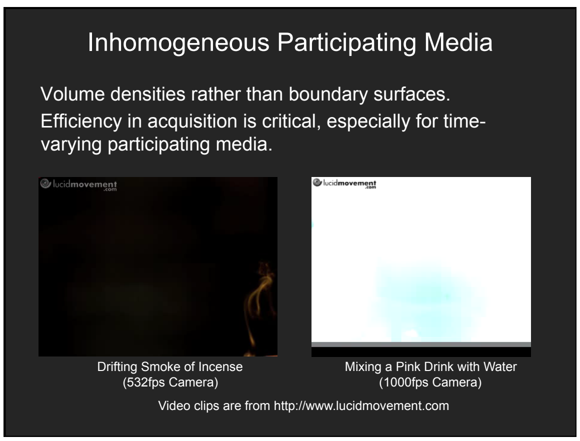
Inhomogeneous Participating Media
Volume densities rather than boundary surfaces.
Efficiency in acquisition is critical, especially for time-varying participating media.
Drifting Smoke of Incense
(532fps Camera)
Mixing a Pink Drink with Water
(1000fps Camera)
Video clips are from http://www.lucidmovement.com

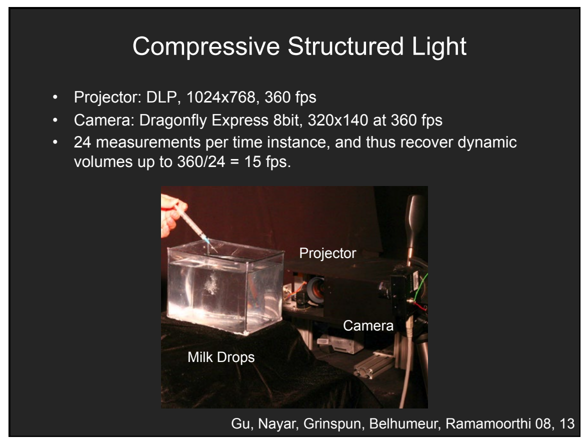
Compressive Structured Light
• Projector: DLP, 1024x768, 360 fps
• Camera: Dragonfly Express 8bit, 320x140 at 360 fps
• 24 measurements per time instance, and thus recover dynamic volumes up to 360/24 = 15 fps.
Projector
Camera
Milk Drops
Gu, Nayar, Grinspun, Belhumeur, Ramamoorthi 08, 13

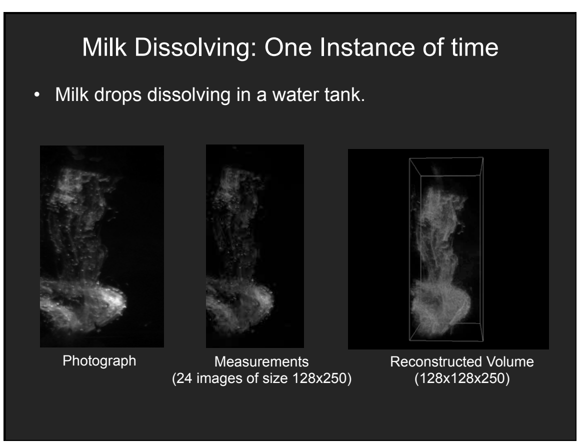
Milk Dissolving: One Instance of time
• Milk drops dissolving in a water tank.
Photograph
Measurements
(24 images of size 128x250)
Reconstructed Volume
(128x128x250)

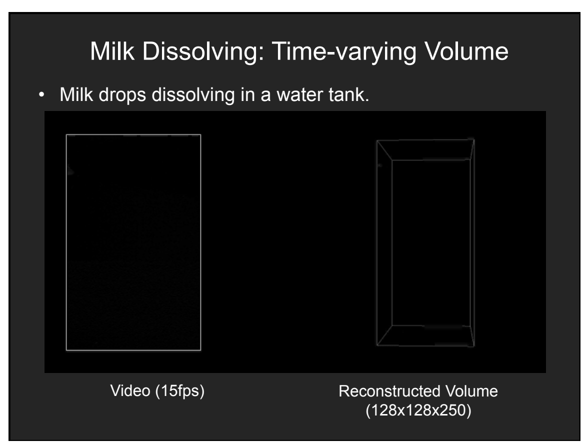
Milk Dissolving: Time-varying Volume
• Milk drops dissolving in a water tank.
Video (15fps)
Reconstructed Volume
(128x128x250)

Outline

- Matrix Row-Column Sampling (Many Lights)

(clustering for matrix completion of light transport)

- Compressive Sensing for Light Transport
- *Matrix Completion*
  - *Extension to compressive sensing: Low rank matrices*
  - *Minimize matrix norm (rank), given some entries*
  - *Combine many ideas seen previously*

Huo et al. SIGGRAPH Asia 16

Outline

- *Matrix Completion*
  - *Extension to compressive sensing: Low rank matrices*
  - *Minimize matrix norm (rank), given some entries*
  - *Combine many ideas seen previously*

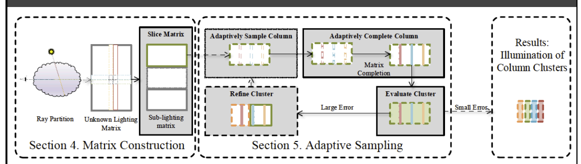

Huo et al. SIGGRAPH Asia 16

Results (Participating Media)

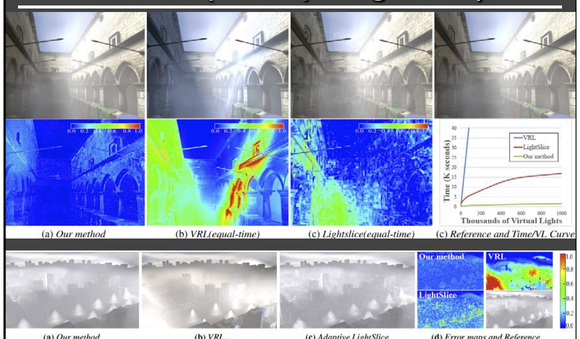

Summary

- Light Transport for Acquisition, Many Light Rendering
- Compressive Sensing for projected patterns
- Matrix Completion for many light rendering

- Leverages popular ideas in applied math
- Consider all forms of coherence
- Think about modern extensions with deep learning