Capturing Light

The Penitent Magdalen, Georges de La Tour, c. 1640

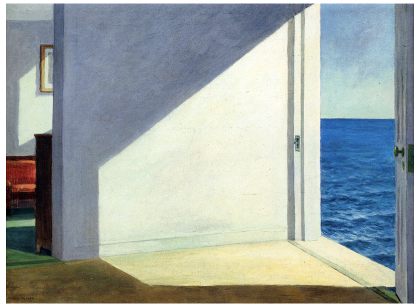

Rooms by the Sea, Edward Hopper, 1951

Some slides from M. Agrawala, F. Durand, P. Debevec, A. Efros, R. Fergus, D. Forsyth, M. Levoy, and S. Seitz

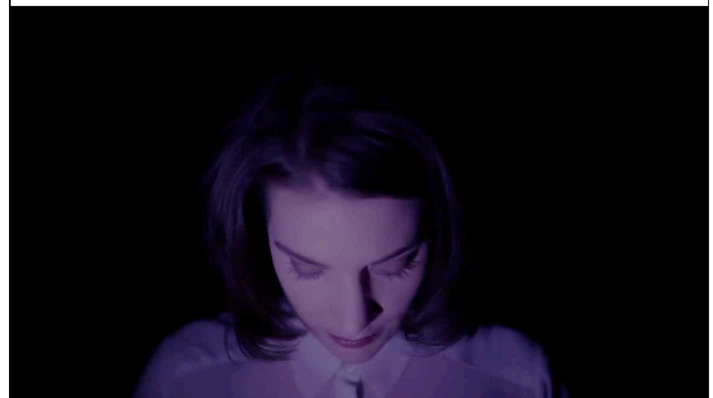

OPALE "Sparkles and Wine" 2013

The Light Field

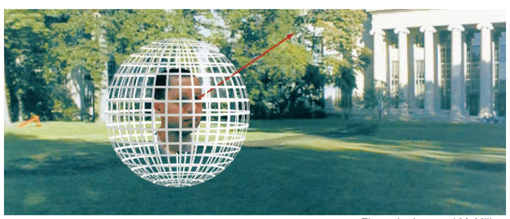

Figure by Leonard McMillan

- What is the set of all things that we can ever see?
- Answer: The **Light Field** (aka **Plenoptic Function**)
- Let's start with a stationary person and try to parameterize <u>everything</u> that she can see …

Grayscale Snapshot

$$P(q, f)$$

- is intensity of light
  - Seen from a single viewpoint
  - At a single time
  - Averaged over the wavelengths of the visible spectrum

Color Snapshot

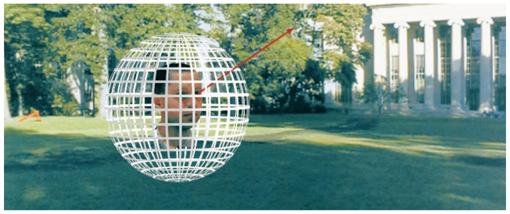

$P(q, f, l)$

- is intensity of light
  - Seen from a single viewpoint
  - At a single time
  - As a function of wavelength

A Movie

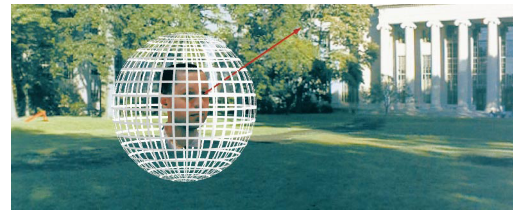

$P(q, f, l, t)$

- is intensity of light
  - Seen from a single viewpoint
  - Over time
  - As a function of wavelength

Holographic Movie

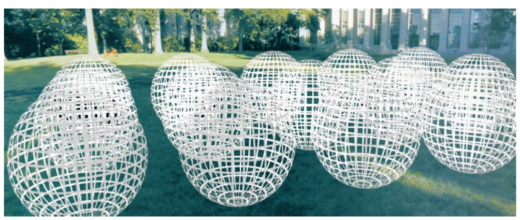

$P(q, f, l, t, V_X, V_Y, V_Z)$

- is intensity of light
  - Seen from ANY viewpoint
  - Over time
  - As a function of wavelength

The Light Field

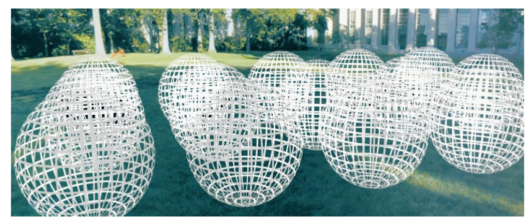

$P(q, f, l, t, V_X, V_Y, V_Z)$

- Can reconstruct every possible view, at every moment, from every position, at every wavelength
- Contains every photograph, every movie, everything that anyone has ever seen!

Sampling the Light Field

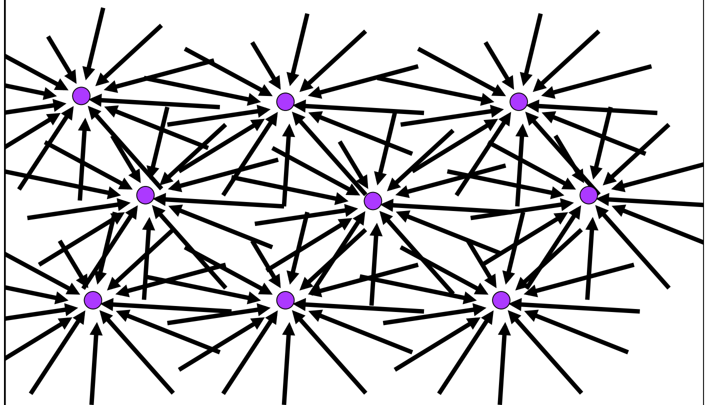

Camera

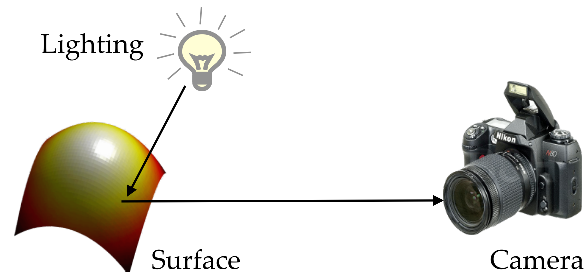

- A camera is a device for capturing and storing **samples of the Light Field**

Building Better Cameras

- Capture more rays
  - Higher density sensor arrays
  - Color cameras, multi-spectral cameras
  - Video cameras

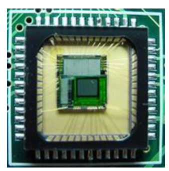

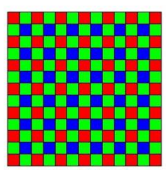

Bayer filter

Modify Optics: Wide-Angle Imaging

Multiple Cameras

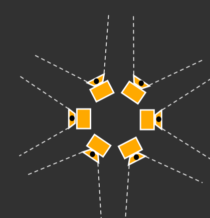

Examples: Disney 55, McCutchen 91, Nalwa 96
Swaminathan & Nayar 99, Cutler et al. 02

Catadioptric Imaging

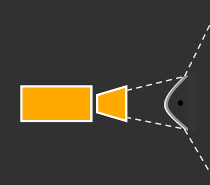

Examples: Rees 70, Charles 87, Nayar 88,
Yagi 90, Hong 91, Yamazawa 95, Bogner 95,
Nalwa 96, Nayar 97, Chahl & Srinivasan 97

Catadioptric Cameras for 360˚ Imaging

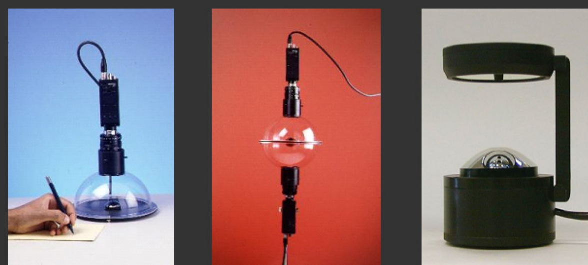

Omnidirectional Image

Catadioptric Imaging

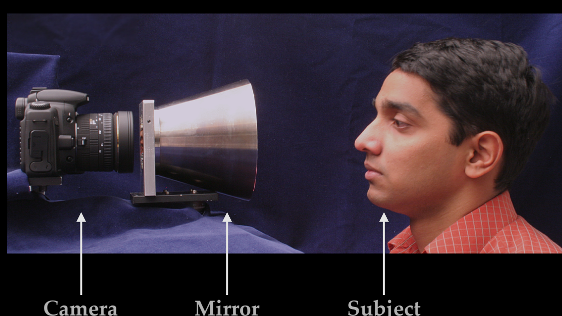

Catadioptric Imaging

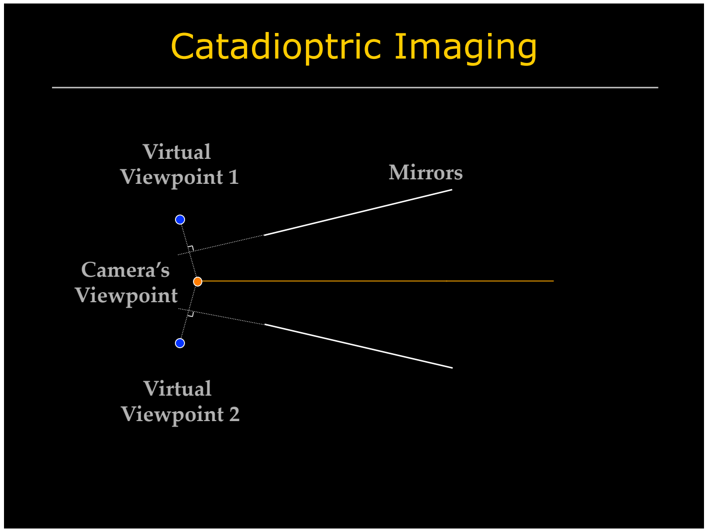
Catadioptric Imaging
Virtual Viewpoint 1
Mirrors
Camera's Viewpoint
Virtual Viewpoint 2

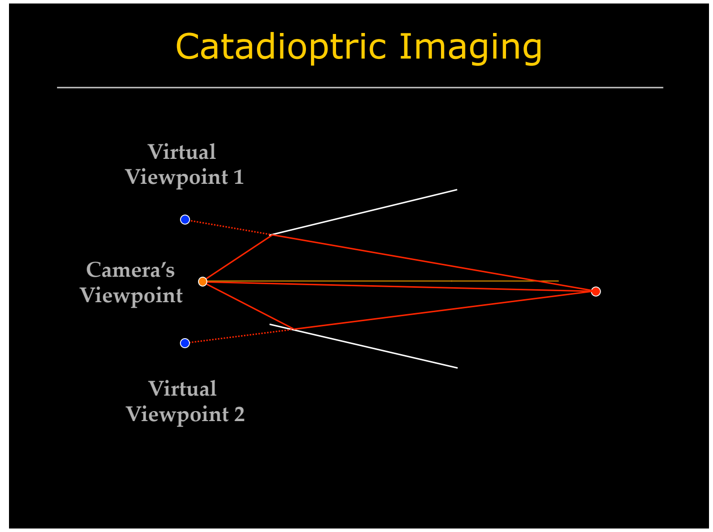
Catadioptric Imaging
Virtual Viewpoint 1
Camera's Viewpoint
Virtual Viewpoint 2

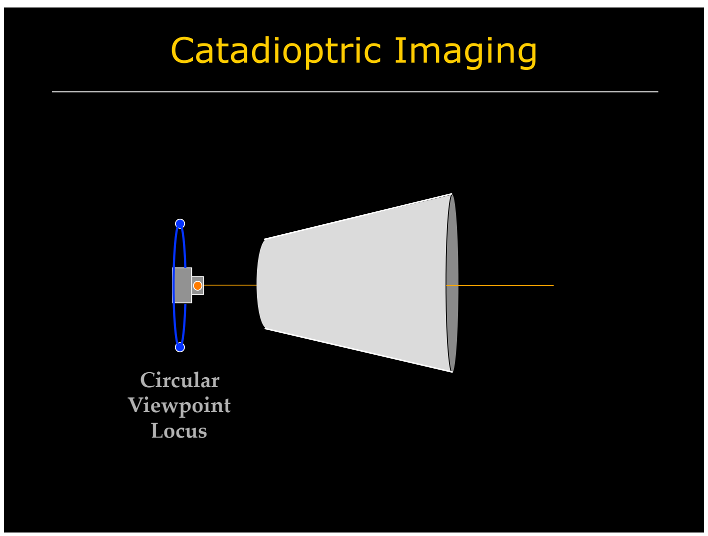
Catadioptric Imaging
Circular Viewpoint Locus

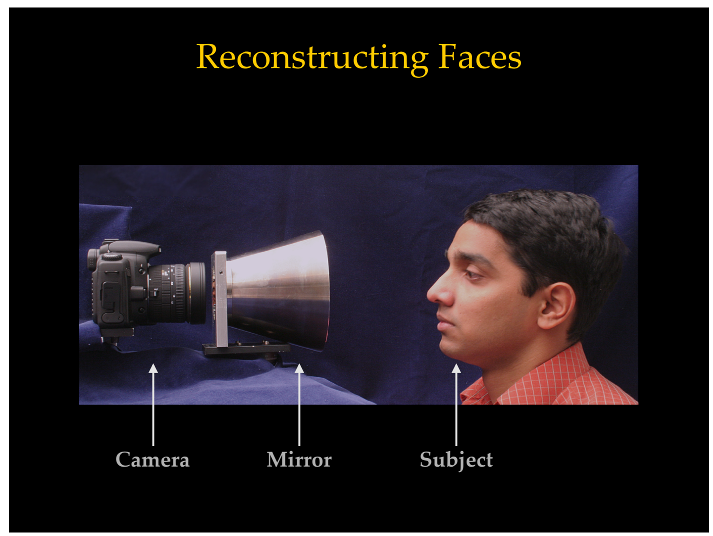
Reconstructing Faces
Camera
Mirror
Subject

Reconstructing Faces

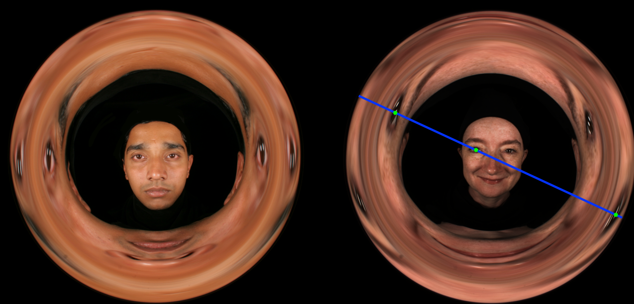

Stereo Views

3D Reconstructions

Femto Photography

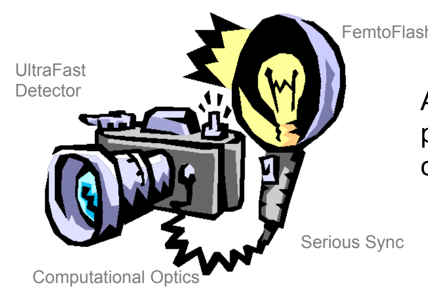

A trillion frame per second camera

Preliminary hardware prototype of the proposed Transient imaging camera

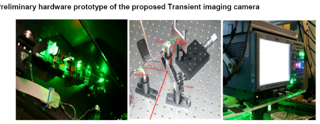

See UW research on this by Prof. Andreas Velten

http://www.youtube.com/watch?v=9xjlck6W020

We successfully built the first hardware prototype to experimentally validate the Transient light transport theory. This is only

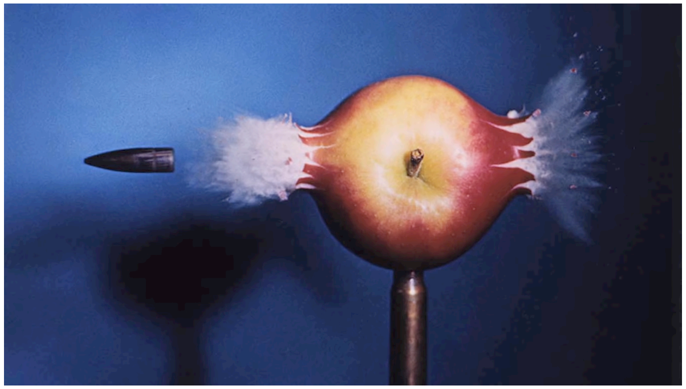

The Light Field

- How to Capture it?
- What's it good for?

Ray

- Ignoring time and color, one sample:

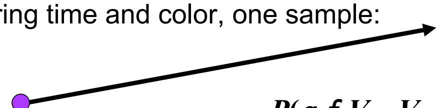

$P(q, f, V_X, V_Y, V_Z)$

- 5D
  - 3D position
  - 2D direction

Slide by Rick Szeliski and Michael Cohen

The Light Field *Surface*

4D:
- 2D direction
- 2D position
- non-dispersive medium

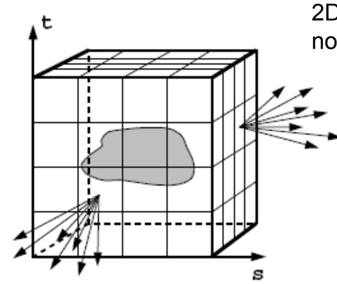

Figure 1: The surface of a cube holds all the radiance information due to the enclosed object.

Light Field - Organization

- 2D position
- 2D direction

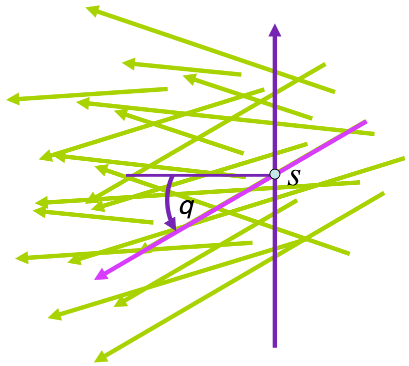

Slide by Rick Szeliski and Michael Cohen

Light Field - Organization

- 2D position
- 2D position

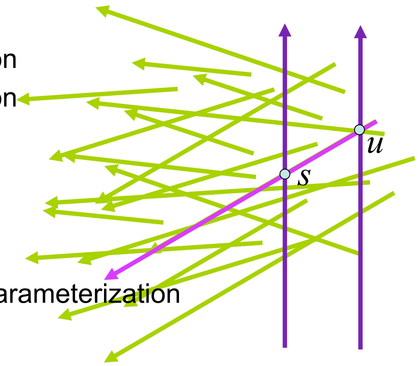

- 2 plane parameterization

Slide by Rick Szeliski and Michael Cohen

Light Field - Organization

- 2D position
- 2D position

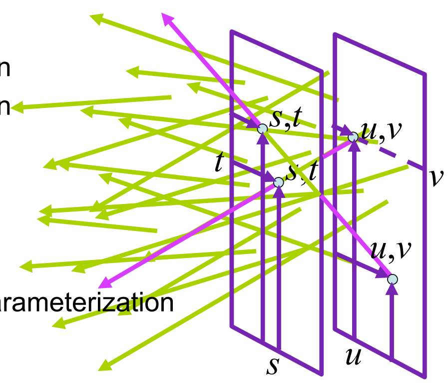

- 2 plane parameterization

Slide by Rick Szeliski and Michael Cohen

Light Field - Organization

- Hold s, t constant
- Let u, v vary
- An image

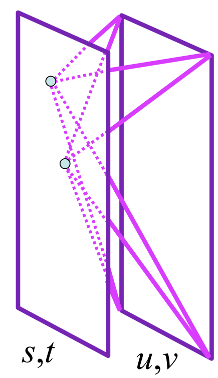

Slide by Rick Szeliski and Michael Cohen

Light Field

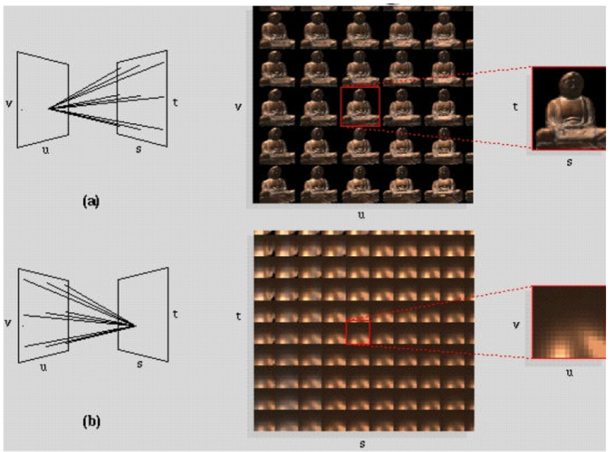

How to Capture Light Fields?

- One camera + move object (and light sources)
- Multiple cameras
- One camera + multiple microlenses

Light Field - Capture

- Idea 1
  - Move camera carefully over *s*, *t* plane
  - Gantry

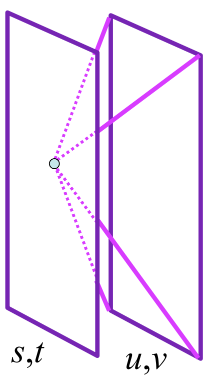

Slide by Rick Szeliski and Michael Cohen

Gantry

- Lazy Susan
  - Manually rotated
- XY Positioner
- Lights turn with lazy susan
- Correctness by construction

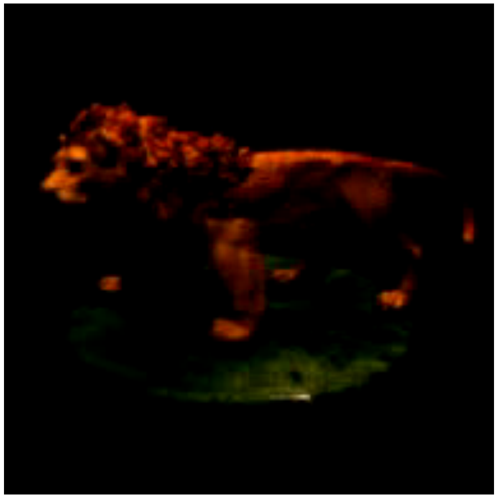

Multi-Camera Arrays

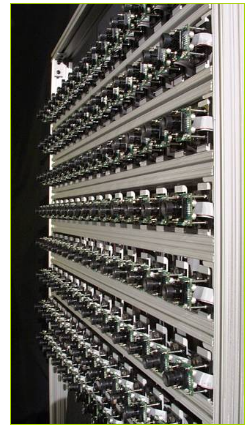

- Stanford's 640 × 480 pixels × 30 fps × 128 cameras
- synchronized timing
- continuous streaming
- flexible arrangement

Stanford Tiled Camera Array

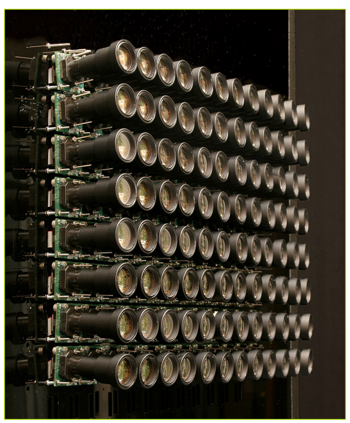

What's a Light Field Good For?

- Synthetic aperture photography
  - Seeing through occluding objects
- Refocusing
- Changing Depth of Field
- Synthesizing images from novel viewpoints

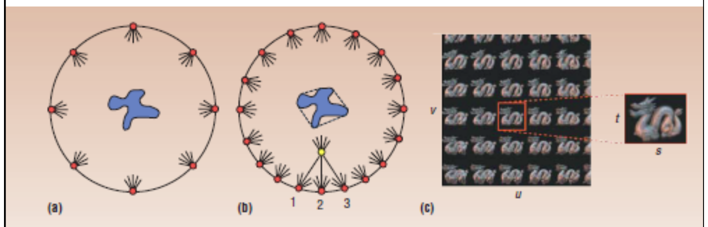

Synthetic Aperture Photography
[Vaish CVPR 2004]

45 cameras aimed at bushes

Synthetic Aperture Photography

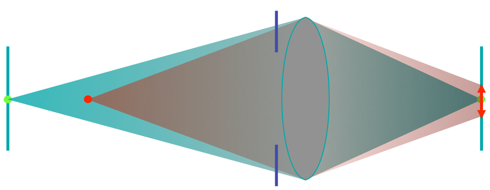

Synthetic Aperture Photography

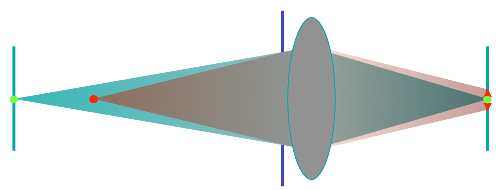

Synthetic Aperture Photography

Red point effectively disappears because it is so blurry

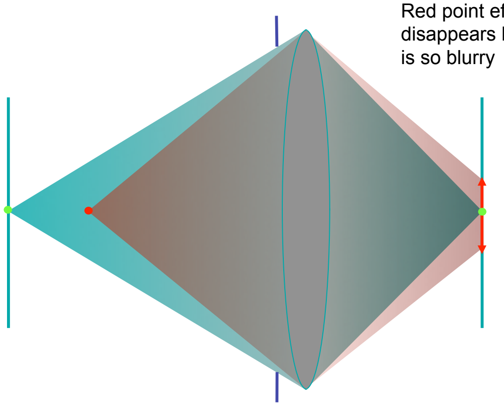

Synthetic Aperture Photography

- If aperture is larger than a foreground occluding object, then some rays from ***behind*** the object are captured

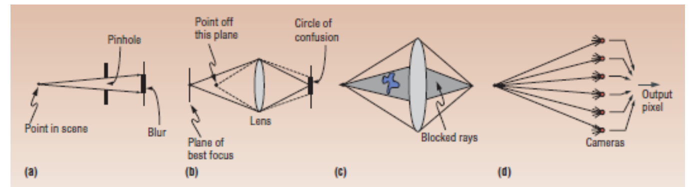

- Leonardo da Vinci observed that if you hold a needle in front of your eye, it adds haze but does not completely obscure any part of it (because your eye' s pupil is bigger than the needle)

Synthetic Aperture Photography

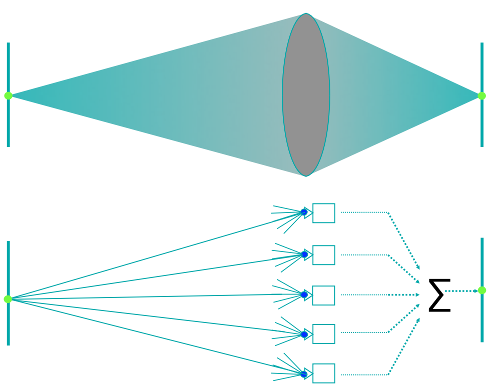

Synthetic Aperture Photography

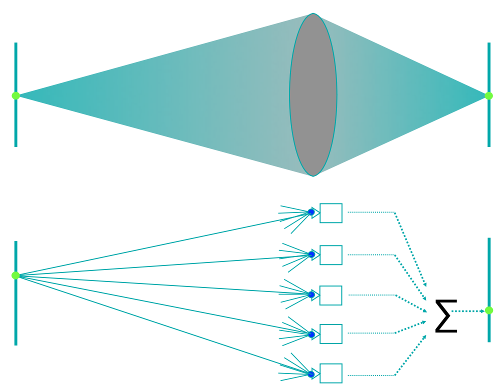

Synthetic Aperture Photography

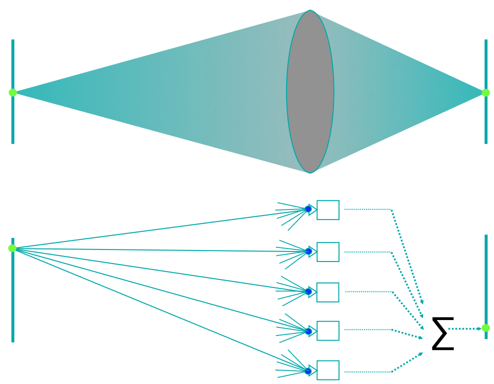

- Another way to think about synthetic aperture photography
  - take the images from **all the cameras**
  - **rectify** them to a common plane in scene (focal plane)
  - **shift** them by a certain amount
  - and **add** them together
- Objects that become aligned by the shifting process
  - will be **sharply focused**
  - objects **in front** of that plane are blurred away
  - objects **in back** of that plane are blurred away

Synthetic Aperture Photography

One image of people behind bushes

Reconstructed synthetic aperture image

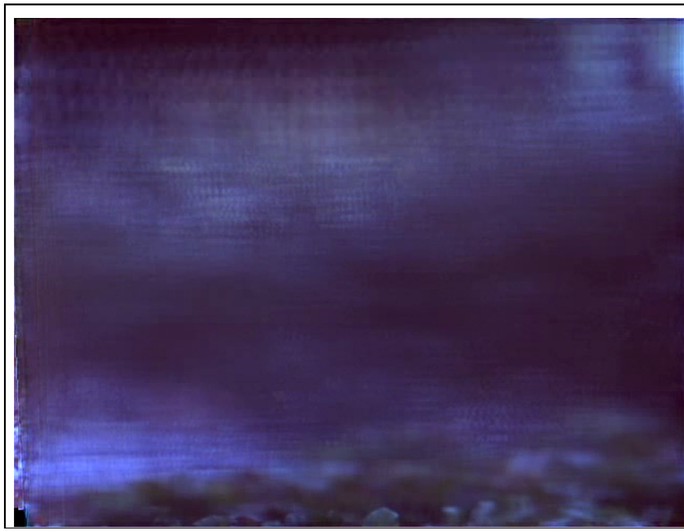

How to Capture Light Fields?

- One camera + move object (and light sources)
- Multiple cameras
- **One camera + multiple microlenses**

Light Field Photography using a Handheld Light Field Camera

Ren Ng, Marc Levoy, Mathieu Brédif, Gene Duval, Mark Horowitz and Pat Hanrahan

Proc. SIGGRAPH 2005

Source: M. Levoy

Lytro Illum Light Field Camera

- www.lytro.com

- 30-250mm lens
- 8.3x optical zoom
- f/2.0 aperture

- $280 ($1,600 MSRP)
- 40 megaray ½″ CMOS sensor
- Maximum image resolution: 2450 × 1634 (4.0 megapixels)

Conventional vs. Light Field Camera

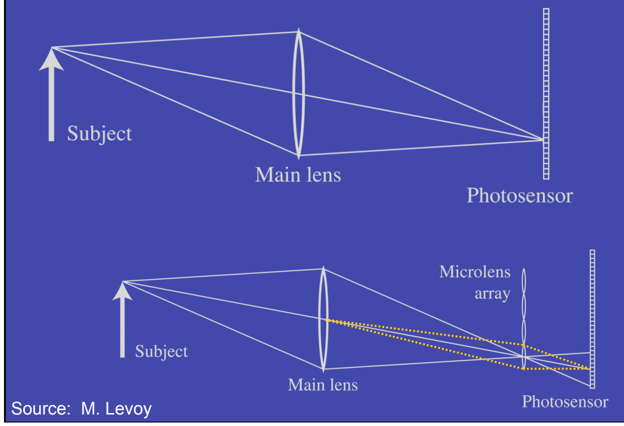

Source: M. Levoy

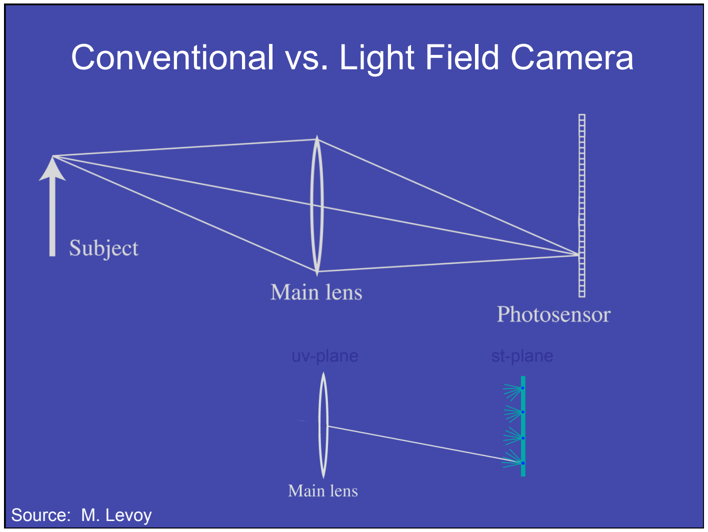
Conventional vs. Light Field Camera
Subject
Main lens
Photosensor
uv-plane
st-plane
Main lens
Source: M. Levoy

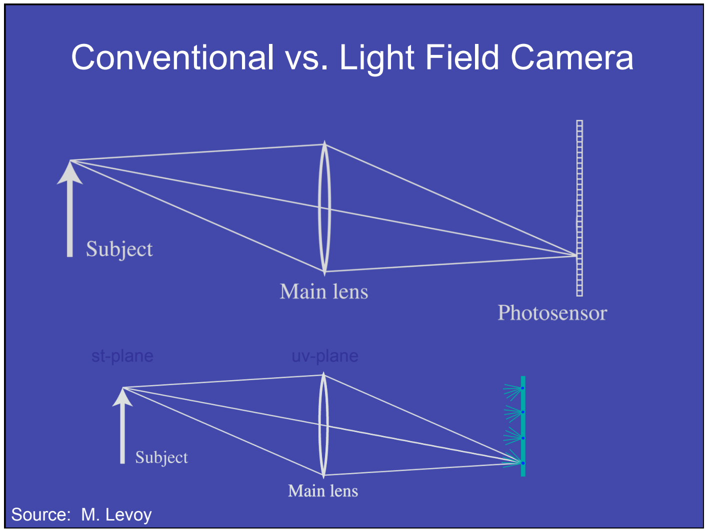
Conventional vs. Light Field Camera
Subject
Main lens
Photosensor
st-plane
uv-plane
Subject
Main lens
Source: M. Levoy

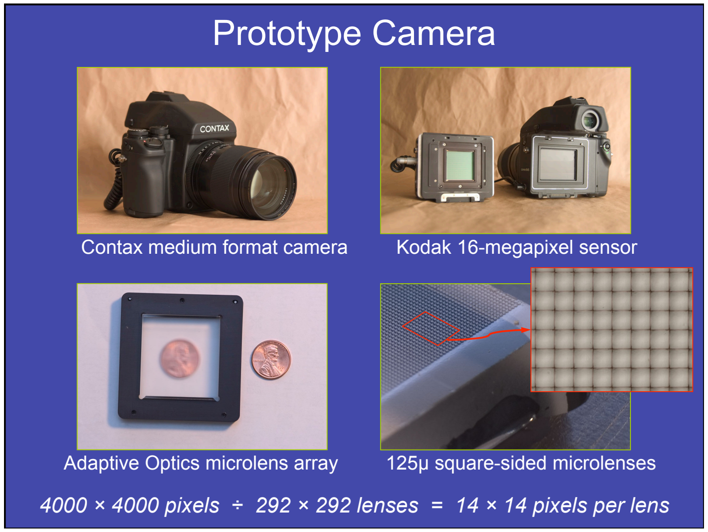
Prototype Camera
CONTAX
Contax medium format camera
Kodak 16-megapixel sensor
Adaptive Optics microlens array
125μ square-sided microlenses
4000 × 4000 pixels ÷ 292 × 292 lenses = 14 × 14 pixels per lens

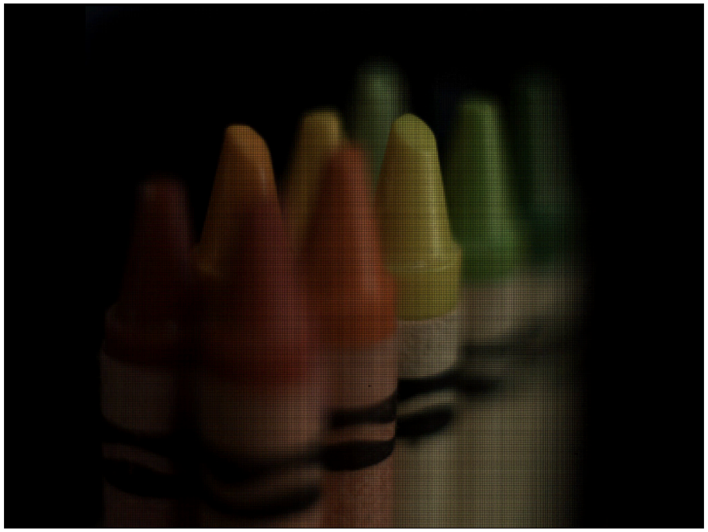

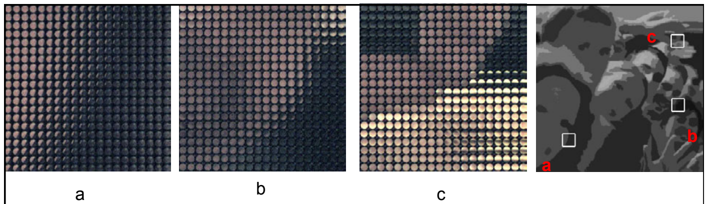

(a) illustrates microlenses at depths ***closer*** than the focal plane. In these *right-side up microlens images, the woman's cheek appears on the left, as it appears in the macroscopic image.* (b) *illustrates microlenses at depths* ***further*** than the focal plane. In these *inverted microlens images, the man's cheek appears on the right, opposite the macroscopic world. This effect is due to inversion* of the microlens' rays as they pass through the world focal plane before arriving at the main lens. (c) illustrates microlenses on edges ***at the focal plane*** (the fingers that are clasped together). The microlenses at this depth are constant in color because all the rays arriving at the microlens originate from the same point on the fingers, which reflect light diffusely.

Digitally Stopping-Down

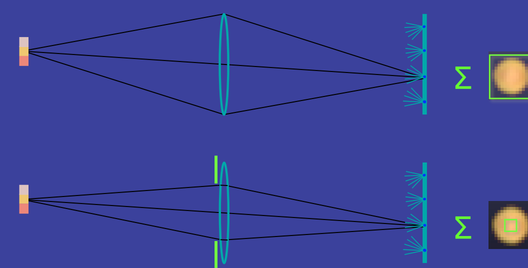

stopping down = summing only the central portion of each microlens

Source: M. Levoy

Digital Refocusing

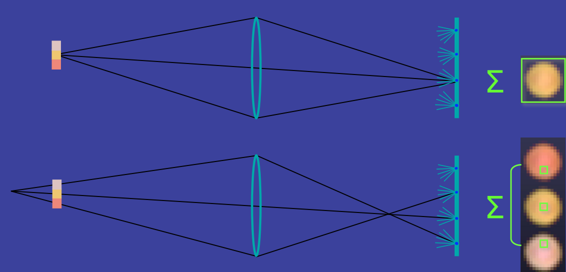

refocusing = summing windows extracted from several microlenses

Source: M. Levoy

Example of Digital Refocusing

Source: M. Levoy

Refocusing Portraits

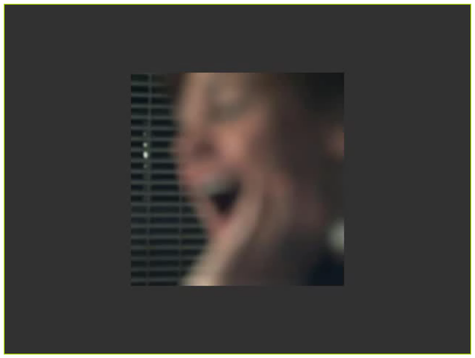

Extending the Depth of Field

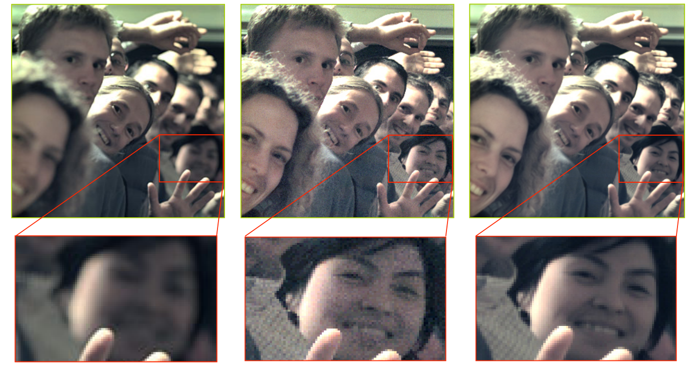

conventional photograph, main lens at *f* / 4

conventional photograph, main lens at *f* / 22

light field, main lens at *f* / 4, after all-focus algorithm [Agarwala 2004]

Digitally Moving the Observer

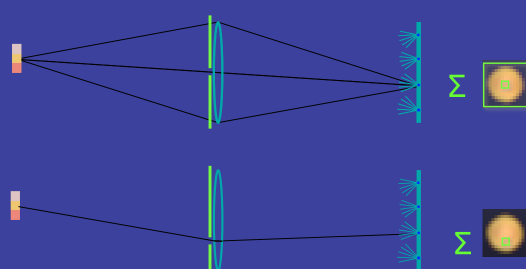

moving the observer = moving the window we extract from the microlenses

Source: M. Levoy

Example of Moving the Observer

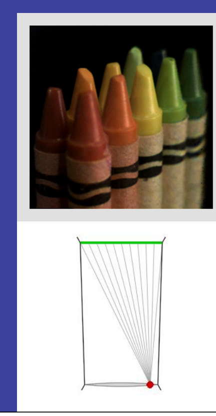

Source: M. Levoy

Moving Backward and Forward

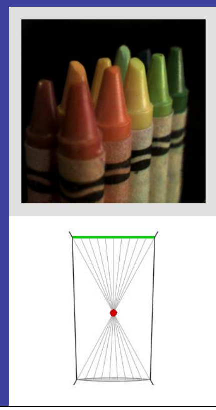

Source: M. Levoy

Implications

- Cuts the unwanted link between exposure (due to the aperture) and depth of field
- Trades off spatial resolution for ability to refocus and adjust the perspective
- Sensor pixels should be made even smaller, subject to the diffraction limit
  - 36mm × 24mm ÷ 2μ pixels = 216 megapixels
  - 18K × 12K pixels
  - 1800 × 1200 pixels × 10 × 10 rays per pixel

Source: M. Levoy

Other ways to Sample the Plenoptic Function

- Moving in **time**:
  - Spatio-temporal volume: $P(q, f, t)$
  - Useful to study temporal changes
  - Long an interest of artists

Claude Monet, Haystacks studies

Space-Time Images

Other ways to slice the plenoptic function: