Lightfields

Lecture 13

Admin

- Projects due by the end of today
  - Email me source code, result images and short report

Overview

- Lightfield representation of a scene
  - Unified representation of all rays
- Lightfield hardware
  - Clever cameras that can capture a lightfield
- Other types of exotic cameras

Overview

- Lightfield representation of a scene
  - Unified representation of all rays
- Lightfield hardware
  - Clever cameras that can capture a lightfield
- Other types of exotic cameras

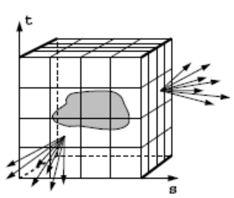

Figure 1: The surface of a cube holds all the radiance information due to the enclosed object.

The Lumigraph, Gortler et al. 1996

Idea

- **Reduce to outside the convex hull of a scene**
- **For every line in space**
- **Store RGB radiance**
- **Then rendering is just a lookup**
- **Two major publication in 1996:**
  - Light field rendering [Levoy & Hanrahan]
    - http://graphics.stanford.edu/papers/light/
  - The Lumigraph [Gortler et al.]
    - Adds some depth information
    - http://cs.harvard.edu/~sjg/papers/lumigraph.pdf

How many dimensions for 3D lines ?

- 4: e.g. 2 for direction, 2 for intersection with plane

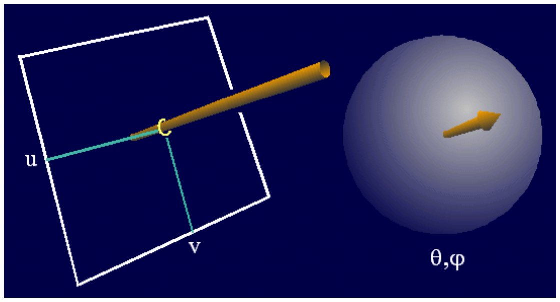

4-D lightfield

- Alternate names: Lumigraph, Plenoptic function

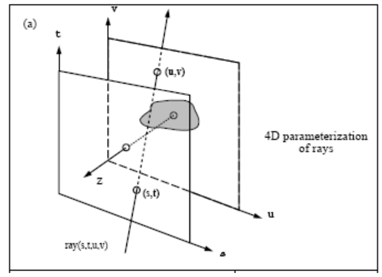

Figure from Gortler et al. SIGGRAPH 1996

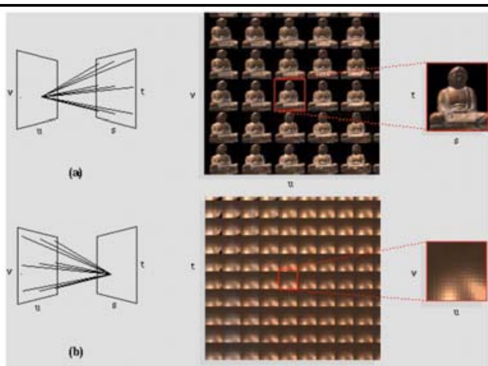

Figure 6: Two visualizations of a light field. (a) Each image in the array represents the rays arriving at one point on the uv plane from all points on the st plane, as shown at left. (b) Each image represents the rays leaving one point on the st plane bound for all points on the uv plane. The images in (a) are off-axis (i.e. sheared) perspective views of the scene, while the images in (b) look like reflectance maps. The latter occurs because the object has been placed astride the focal plane, making sets of rays leaving points on the focal plane similar in character to sets of rays leaving points on the object.

Cool visualization

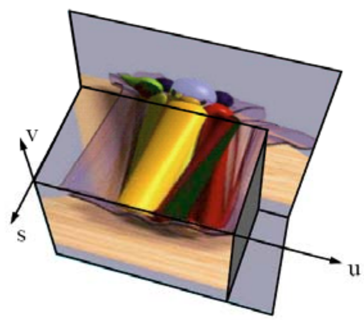

Figure 7: An (s, u, v) slice of a Lumigraph

From Gortler et al.

View = 2D plane in 4D

- With various resampling issues

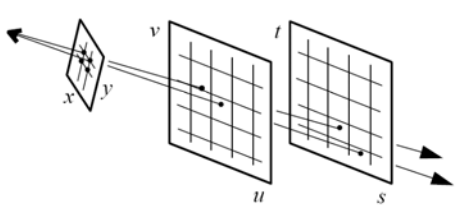

Figure 12: The process of resampling a light slab during display.

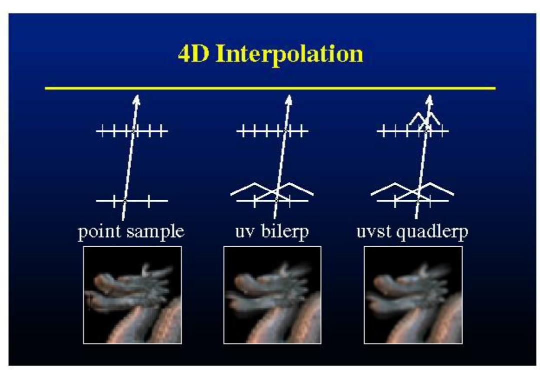

Slide by Marc Levoy

Plenoptic function

Adelson & Bergen '91

- 4-D → Lightfield (transparent medium)
- 5-D → Lightfield + attenuation along rays
- 6-D → Time-varying lightfield w/attenuation
- 7-D → 6-D + spectral information

Complete representation

- Captures exterior of convex hull

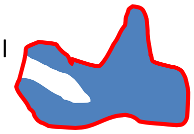

- Depth
  - Corresponds to slope
- BRDF
  - Bidirectional reflectance function

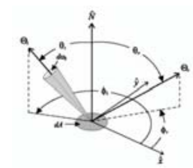

http://math.nist.gov/~FHunt/appearance/brdf.html

Demo Demo

Fourier Slice Photography

- Ren Ng Stanford University

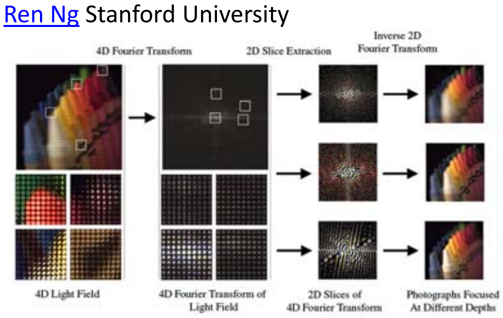

Fourier slice photography

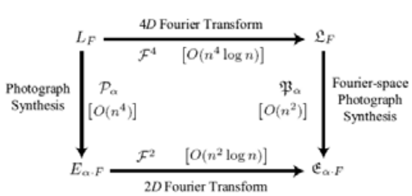

Figure 5: Fourier Slice Photography Theorem (Eq. 9). Transform relationships between the 4D light field L_F, a lens-formed 2D photograph $E_{\alpha \cdot F}$, and their respective Fourier spectra, $\mathfrak{L}_F$ and $\mathfrak{E}_{\alpha \cdot F}$. n is the number of samples in each dimension.

Overview

- Lightfield representation of a scene
  - Unified representation of all rays
- Lightfield hardware
  - Clever cameras that can capture a lightfield
- Other types of exotic cameras

Light field photography and videography

Marc Levoy

Computer Science Department
Stanford University

High performance imaging using large camera arrays

Bennett Wilburn, Neel Joshi, Vaibhav Vaish, Eino-Ville Talvala, Emilio Antunez, Adam Barth, Andrew Adams, Mark Horowitz, Marc Levoy

(Proc. SIGGRAPH 2005)

Stanford multi-camera array

- 640 × 480 pixels × 30 fps × 128 cameras
- synchronized timing

Figure 2: Our camera tiles contain an Omnivision 8610 image sensor, passive electronics, and a lens mount. The ribbon cables carry video data, synchronization signals, control signals, and power between the tile and the processing board. To keep costs low, we use fixed-focus, fixed-aperture lenses.

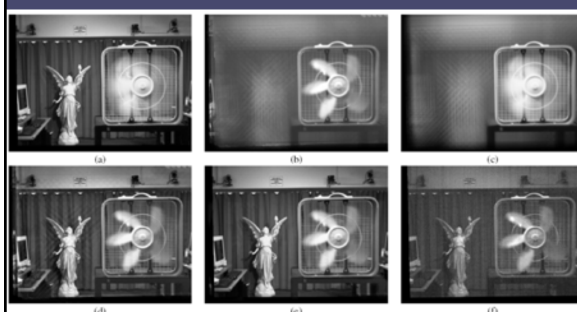

Figure 12: Hybrid synthetic aperture photography for combining high depth of field and low motion blur. (a-c) Images captured of a scene simultaneously through three different apertures: a single camera with a long exposure time (a), a large synthetic aperture with short exposure time (b), and a large synthetic aperture with a long exposure time. Computing (a+b-c) yields image (d), which has aliasing artifacts because the synthetic apertures are sampled sparsely from slightly different locations. Masking pixels not in focus in the synthetic aperture images before computing the difference (a + b - c) removes the aliasing (e). For comparison, image (f) shows the image taken with an aperture that is narrow in both space and time. The entire scene is in focus and the fan motion is frozen, but the image is much noisier.

Robotic Camera

Image Leonard McMillan

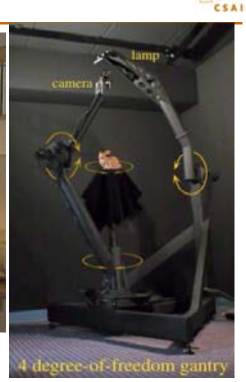

Image Levoy et al.

Plenoptic camera

- **For depth extraction**
- **Adelson & Wang 92**
  http://www-bcs.mit.edu/people/jyawang/demos/plenoptic/plenoptic.html

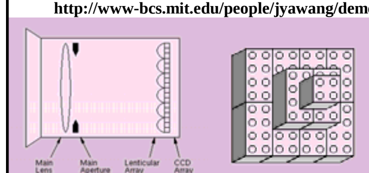

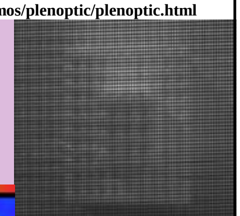

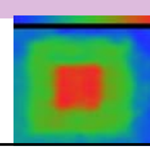

Light field photography using a handheld plenoptic camera

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

(Proc. SIGGRAPH 2005
and TR 2005-02)

Conventional Photograph

Slide by Ren Ng.

Light Field Photography

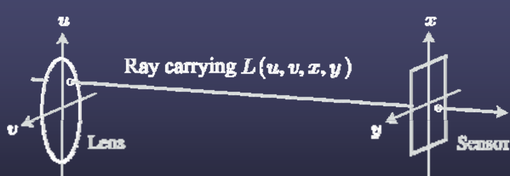

- Capture the light field inside the camera body

Slide by Ren Ng.

Hand-Held Light Field Camera

Medium format digital camera

Camera in-use

16 megapixel sensor

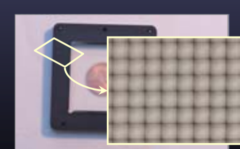

Microlens array

Slide by Ren Ng.

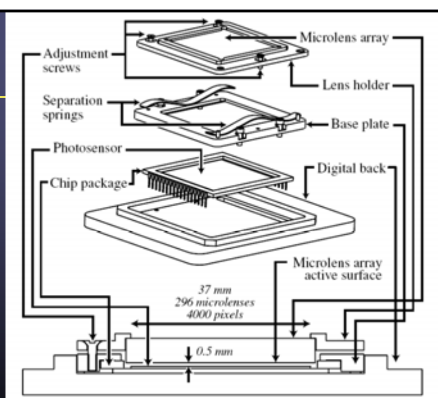

Figure 8: *Top:* Exploded view of assembly for attaching the microlens array to the digital back. *Bottom:* Cross-section through assembled parts.

Light Field in a Single Exposure

Slide by Ren Ng.

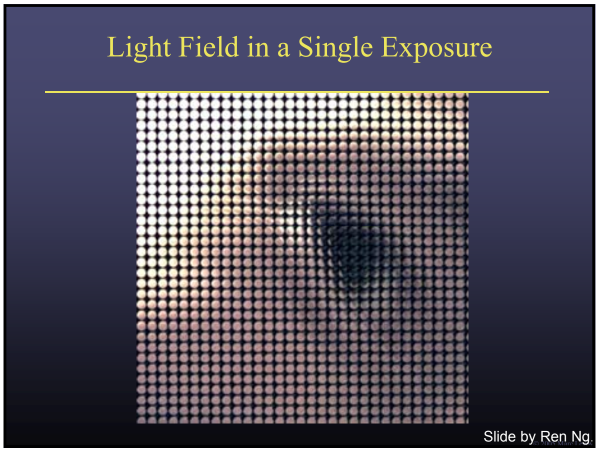
Light Field in a Single Exposure
Slide by Ren Ng.

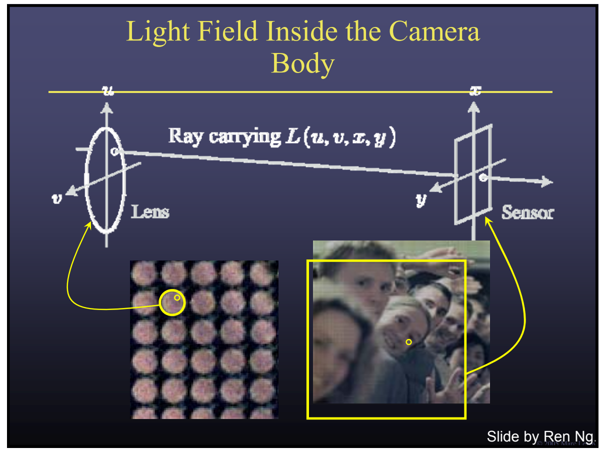
Light Field Inside the Camera Body
u
x
Ray carrying $L(u, v, x, y)$
v
Lens
y
Sensor
Slide by Ren Ng.


Digital Refocusing
Slide by Ren Ng.


Digital Refocusing
Slide by Ren Ng.

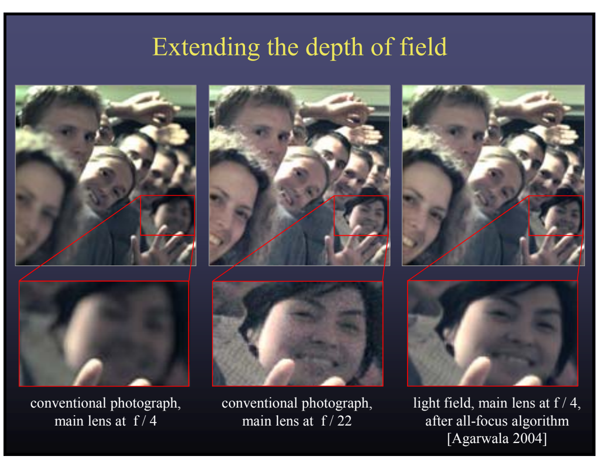
Extending the depth of field
conventional photograph, main lens at f / 4
conventional photograph, main lens at f / 22
light field, main lens at f / 4, after all-focus algorithm [Agarwala 2004]

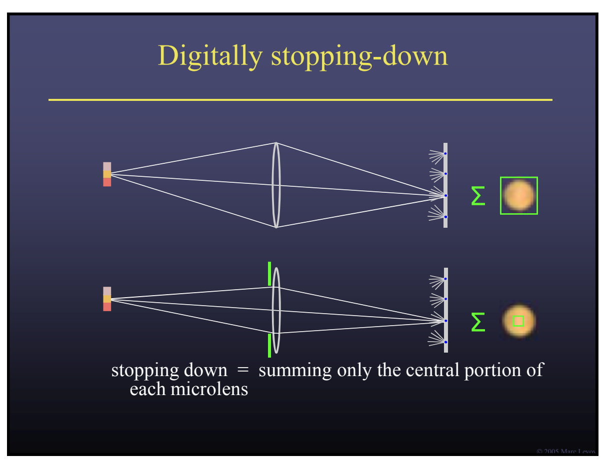
Digitally stopping-down
Σ
Σ
stopping down = summing only the central portion of each microlens

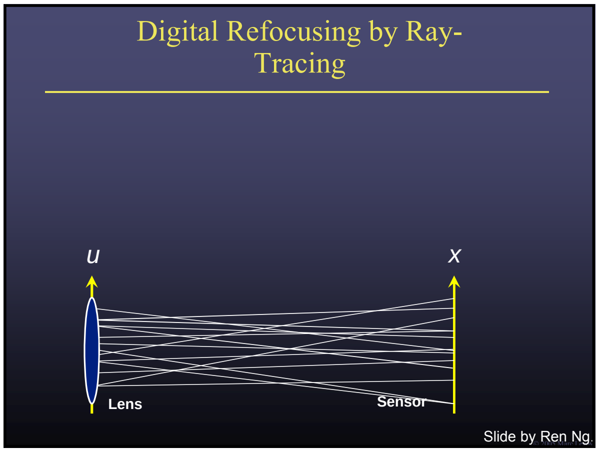
Digital Refocusing by Ray-Tracing
u
x
Lens
Sensor
Slide by Ren Ng.

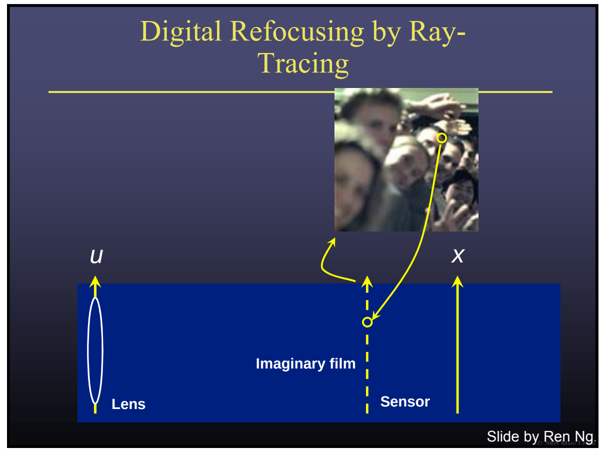
Digital Refocusing by Ray-Tracing
u
x
Imaginary film
Lens
Sensor
Slide by Ren Ng.

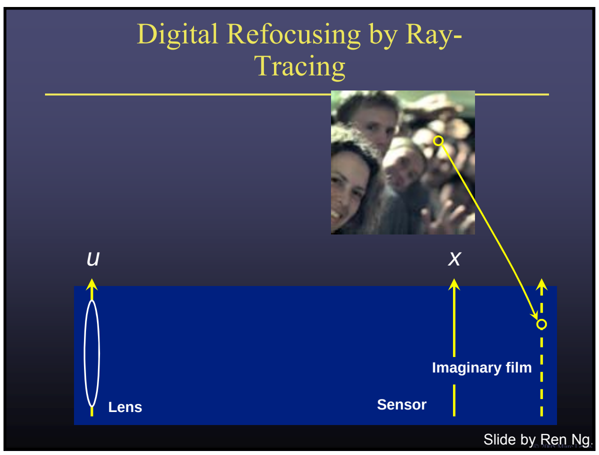
Digital Refocusing by Ray-Tracing
u
x
Imaginary film
Lens
Sensor
Slide by Ren Ng.

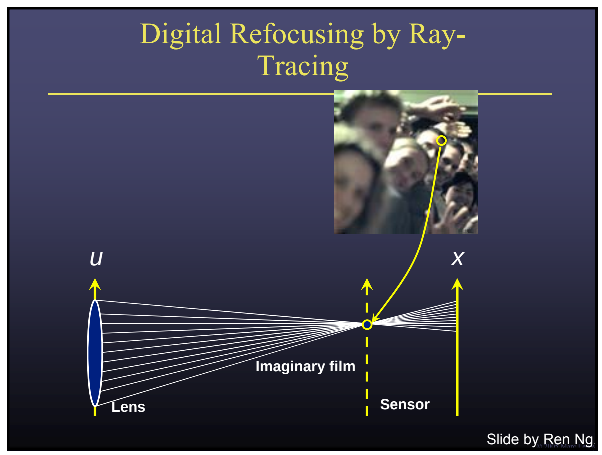
Digital Refocusing by Ray-Tracing
u
x
Imaginary film
Lens
Sensor
Slide by Ren Ng.

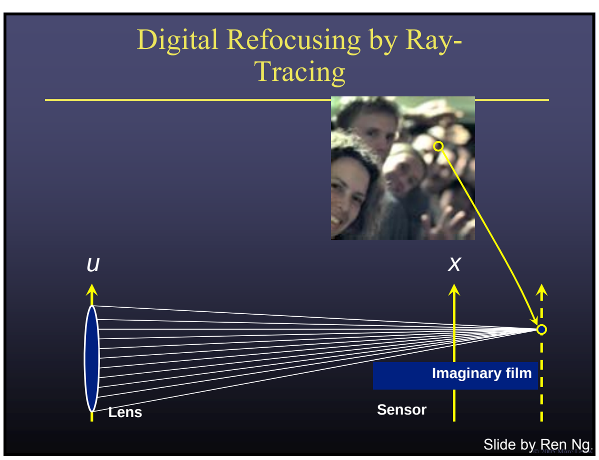
Digital Refocusing by Ray-Tracing
u
x
Imaginary film
Lens
Sensor
Slide by Ren Ng.

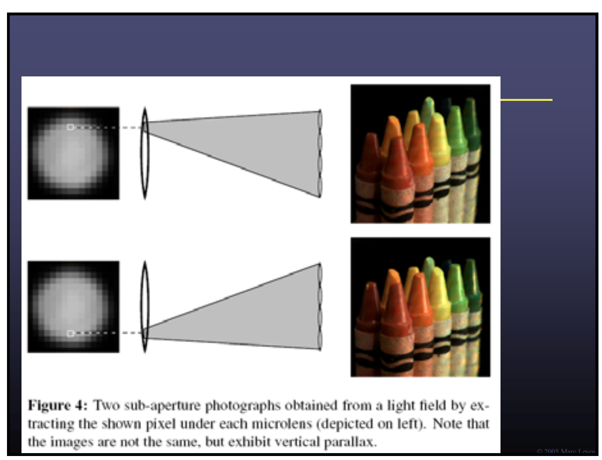
Figure 4: Two sub-aperture photographs obtained from a light field by extracting the shown pixel under each microlens (depicted on left). Note that the images are not the same, but exhibit vertical parallax.

Digitally moving the observer

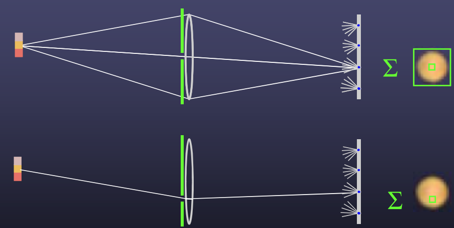

- moving the observer = moving the window we extract from the microlenses

Example of moving the observer

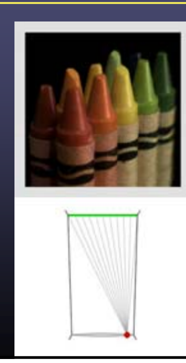

Moving backward and forward

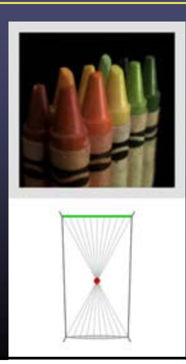

Results of Band-Limited Analysis

- Assume a light field camera with
  - An *f* /*A* lens
  - *N* x *N* pixels under each microlens
- From its light fields we can
  - Refocus *exactly* within depth of field of an *f* /(*A* * *N*) lens
- In our prototype camera
  - Lens is *f* /4
  - 12 x 12 pixels under each microlens
- Theoretically refocus within depth of field of an *f*/48 lens

Slide by Ren Ng.

Implications

- cuts the unwanted link between exposure (due to the aperture) and depth of field
- trades off (excess) spatial resolution for ability to refocus and adjust the perspective
- sensor pixels should be made even smaller, subject to the diffraction limit

  36mm × 24mm ÷ 2.5μ pixels = 266 megapixels

  20K × 13K pixels

  4000 × 2666 pixels × 20 × 20 rays per pixel

Light Field Microscopy

Marc Levoy, Ren Ng, Andrew Adams, Matthew Footer, and Mark Horowitz

(Proc. SIGGRAPH 2006)

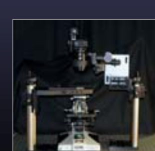

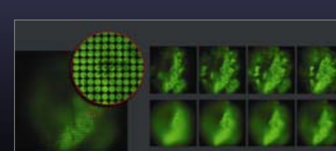

A traditional microscope

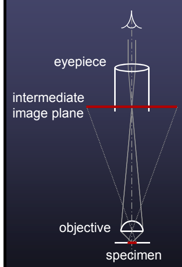

A light field microscope (LFM)

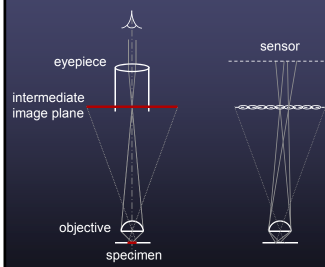

→ reduced lateral resolution on specimen = 0.26μ × 12 spots = 3.1μ

- 40x / 0.95NA objective
  ↓
  0.26μ spot on specimen × 40x = 10.4μ on sensor
  ↓
  2400 spots over 25mm field
- 125^2-micron microlenses
  ↓
  200 × 200 microlenses with 12 × 12 spots per microlens

A light field microscope (LFM)

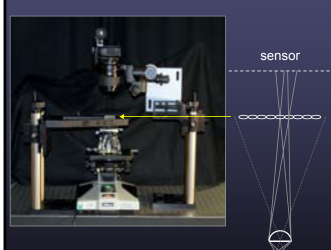

Example light field micrograph

- orange fluorescent crayon
- mercury-arc source + blue dichroic filter
- 16x / 0.5NA (dry) objective
- f/20 microlens array
- 65mm f/2.8 macro lens at 1:1
- Canon 20D digital camera

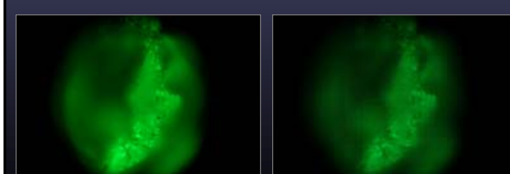

ordinary microscope | light field microscope

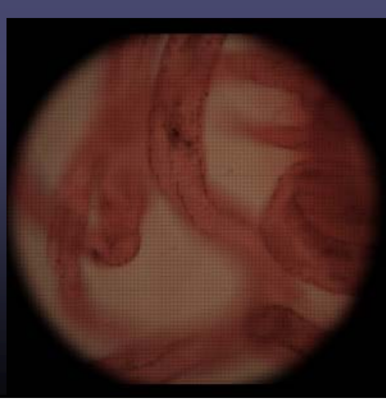

Show video

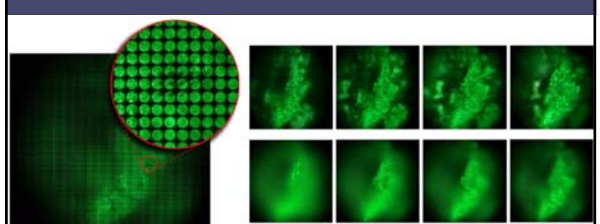

Figure 1: At left is a light field captured by photographing a speck of fluorescent crayon wax through a microscope objective and microlens array. The objective magnification is 16×, and the field of view is 1.3mm wide. The image consists of 170^2 subimages, one per microlens, each depicting a different part of the specimen. An individual subimage contains 20^2 pixels, each representing a different point on the objective lens and hence a unique direction of view. By extracting one pixel from each subimage, we can produce perspective views of the specimen, a sequence of which is shown at top-right. Alternatively, by summing the pixels in each subimage, we can produce orthographic views with a shallow depth of field, like an ordinary microscope but of lower spatial resolution. Shearing the light field before summing, we can focus at different depths, as shown in the sequence at bottom-right. These images were computed in real-time on a PC.

The geometry of the light field in a microscope

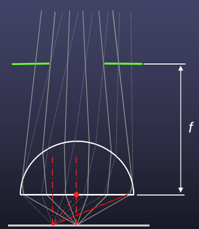

objective lenses are telecentric

- microscopes make orthographic views
- translating the stage in X or Y provides no parallax on the specimen
- out-of-plane features don't shift position when they come into focus
- front lens element size = aperture width + field width
- PSF for 3D deconvolution microscopy is shift-invariant (i.e. doesn't change across the field of view)

Panning and focusing

panning sequence

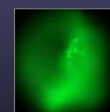

focal stack

Real-time viewer

Other examples

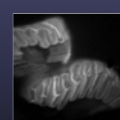

fern spore
(60x, autofluorescence)

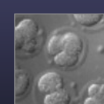

mouse oocyte
(40x, DIC)

Golgi-stained neurons
(40x, transmitted light)

Extensions

- digital correction of aberrations
  - by capturing and resampling the light field

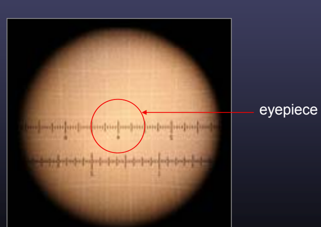

Nikon 40x 0.95NA (dry) Plan-Apo

Extensions

- digital correction of aberrations
  - by capturing and resampling the light field

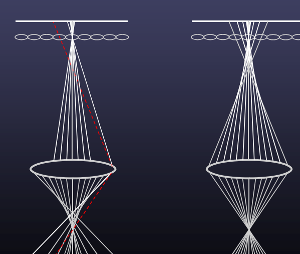

Extensions

- digital correction of aberrations
  - by capturing and resampling the light field

correcting for aberrations caused by imaging through thick specimens whose index of refraction doesn't match that of the immersion medium

Extensions

- digital correction of aberrations
  - by capturing and resampling the light field
- multiplexing of variables other than angle
  - by placing gradient filters at the aperture plane, such as neutral density, spectral, or polarization

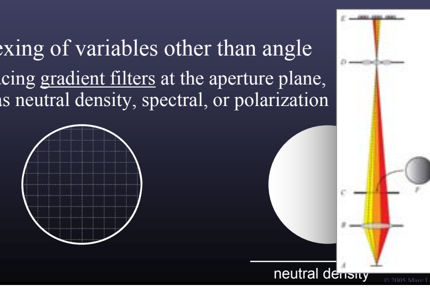

Extensions

- digital correction of aberrations
  - by capturing and resampling the light field
- multiplexing of variables other than angle
  - by placing gradient filters at the aperture plane, such as neutral density, spectral, or polarization

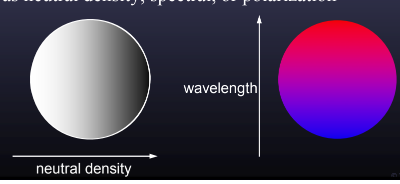

Extensions

- digital correction of aberrations
  - by capturing and resampling the light field
- multiplexing of variables other than angle
  - by placing gradient filters at the aperture plane, such as neutral density, spectral, or polarization

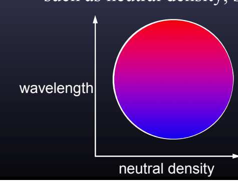

... or polarization direction
... or ???

- gives up digital refocusing?

Extensions

- digital correction of aberrations
  - by capturing and resampling the light field
- multiplexing of variables other than angle
  - by placing gradient filters at the aperture plane, such as neutral density, spectral, or polarization
- microscope scatterometry
  - by controlling the incident light field using a micromirror array + microlens array

Programmable incident light field

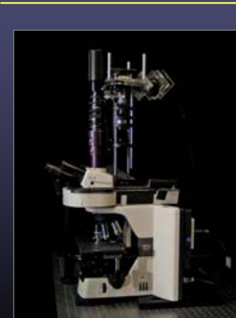

- light source + micromirror array + microlens array
- 800 × 800 pixels = 40 × 40 tiles × 20 × 20 directions
- driven by image from PC graphics card

Other applications of light field illumination: "4D designer lighting"

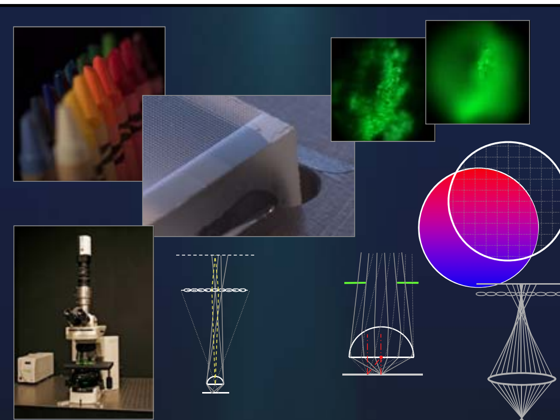

http://graphics.stanford.edu

Overview

- Lightfield representation of a scene
  - Unified representation of all rays
- Lightfield hardware
  - Clever cameras that can capture a lightfield
- Other types of exotic cameras

Computational Imaging: Recent Advances in Optics

Shree K. Nayar

Computer Science
Columbia University

The Eye's Lens

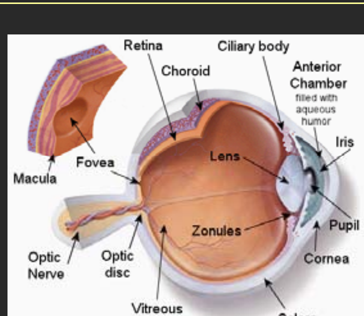

Varioptic Liquid Lens: Electrowetting

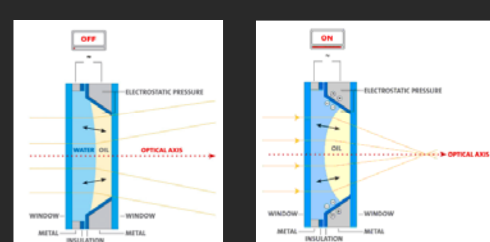

Varioptic, Inc., 2007

Varioptic Liquid Lens

(Courtesy Varioptic Inc.)

Captured Video

(Courtesy Varioptic Inc.)

Conventional Compound Lens

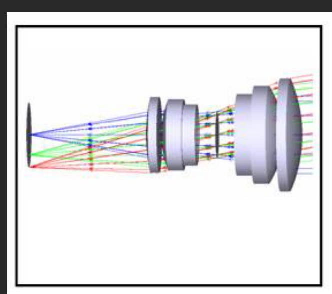

"Origami Lens": Thin Folded Optics (2007)

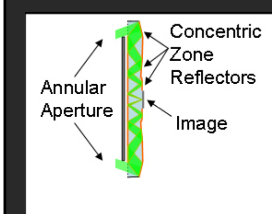

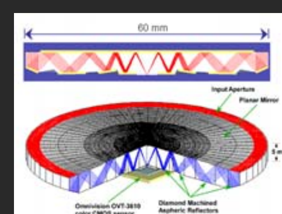

"Ultrathin Cameras Using Annular Folded Optics,"
E. J. **Tremblay**, R. A. Stack, R. L. Morrison, J. E. **Ford**
Applied Optics, 2007 - OSA

Origami Lens

Optical Performance

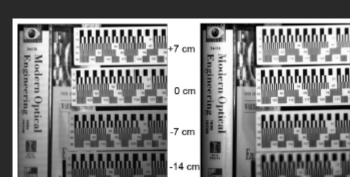

Compound Lens of Dragonfly

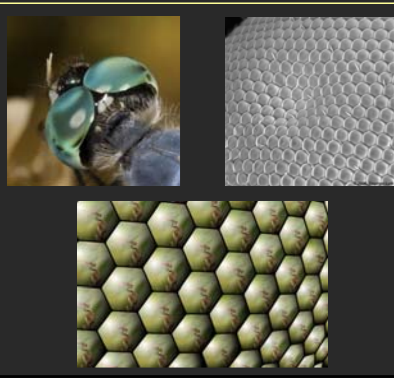

TOMBO: Thin Camera (2001)

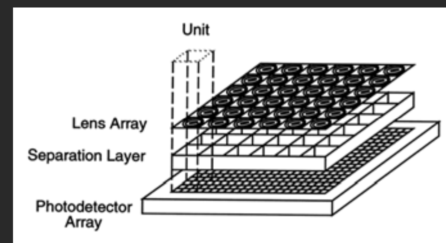

"Thin observation module by bound optics (TOMBO),"
J. Tanida, T. Kumagai, K. Yamada, S. Miyatake
Applied Optics, 2001

TOMBO: Thin Camera

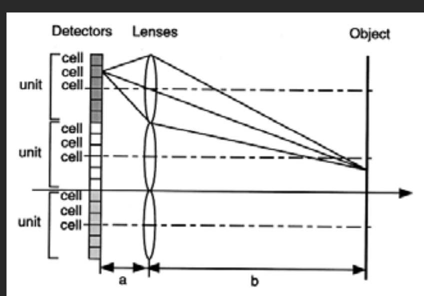

Captured Image

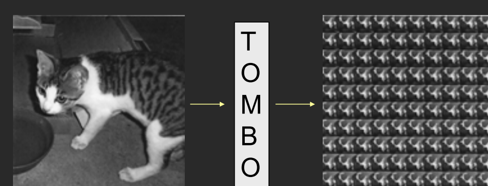

Scene

Captured Image

(Multiple low-resolution copies of the scene)

$$\mathbf{g} = \mathbf{H}\mathbf{f}$$

Image = Optics . Scene

Reconstructed Image

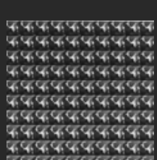

$$\mathbf{f} = \mathbf{H}^{+}\mathbf{g}$$

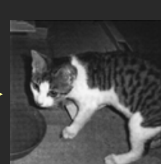

Conventional Lens: Limited Depth of Field

Open Aperture

Smaller Aperture

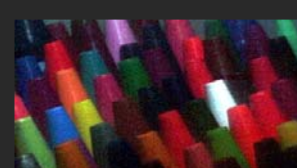

Wavefront Coding using Cubic Phase Plate

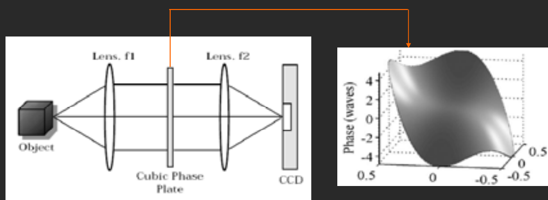

"*Wavefront Coding: jointly optimized optical and digital imaging systems*",
E. Dowski, R. H. Cormack and S. D. Sarama ,
Aerosense Conference, April 25, 2000

Wavefront coding

- Insert special element into lens
- All-depths blurred equally
- Single deconvolution yields all-focus image
- Parabolic lightfield integration path
- Parabola is only shape that is invariant under shear

Photography as Integration

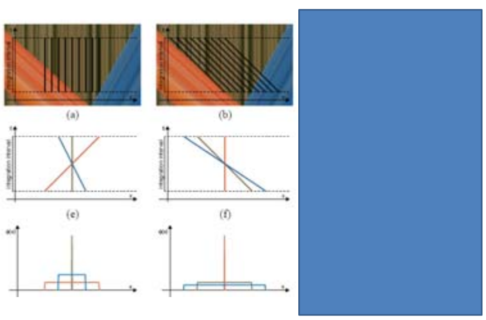

Figure from Levin et al. 2008

Parabola is shear invariant

- $f(t) = a_0 t^2$

$$f_s(t) = f(t) - st = a_0(t - \frac{s}{2a_0})^2 - \frac{s^2}{4a_0}$$

Depth Invariant Blur

Conventional System | Wavefront Coded System

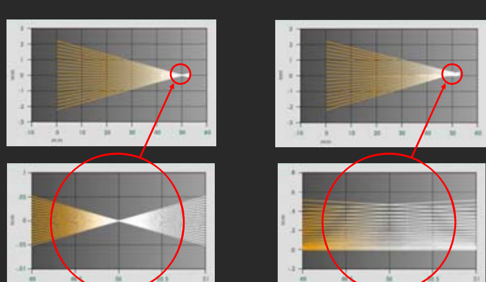

Point Spread Functions

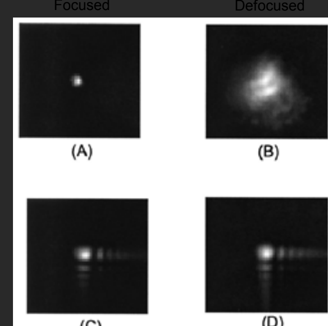

Example

Conventional System	Wavefront Coded System
Open Aperture	Captured Image
	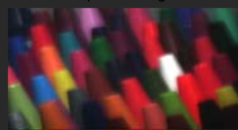
Stopped Down	After Processing
	

Compressed Imaging

Sparsity of Image: $X = \Psi\ \theta$ (sparse basis, coefficients)

Measurements: $Y = \Phi\ X$ (measurement basis)

"A New Compressive Imaging Camera Architecture"
D. Takhar et al., Proc. SPIE Symp. on Electronic Imaging, Vol. 6065, 2006.

Single Pixel Camera

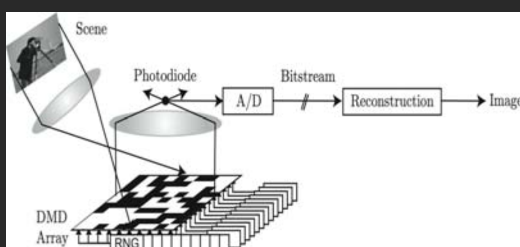

Single Pixel Camera

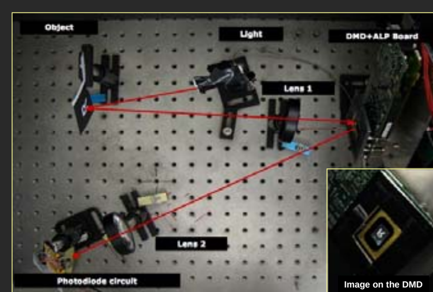

Example

Original	Compressed Imaging	
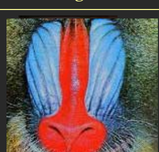	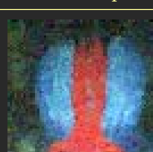	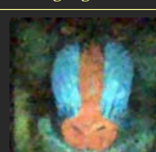
	4096 Pixels 1600 Measurements (40%)	65536 Pixels 6600 Measurements (10%)

Example

Original	Compressed Imaging	
	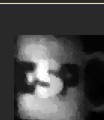	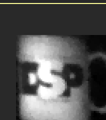
	4096 Pixels 800 Measurements (20%)	4096 Pixels 1600 Measurements (40%)