Light field sensing

Marc Levoy

Computer Science Department
Stanford University

The scalar light field
(in geometrical optics)

Radiance as a function of position and direction in a static scene with fixed illumination

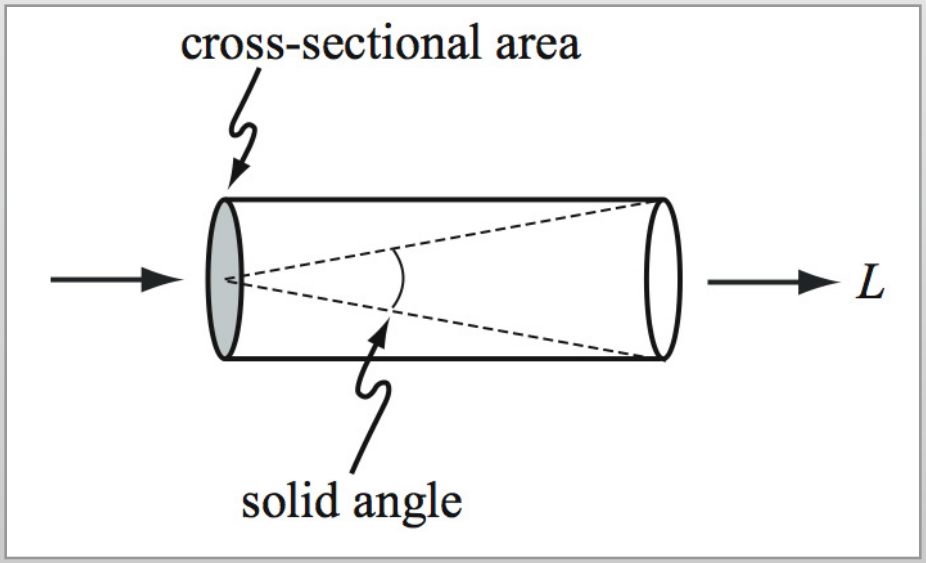

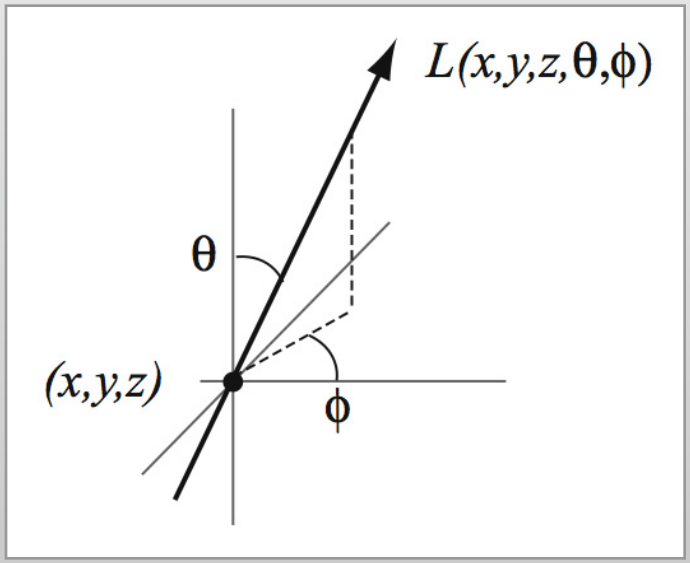

L is radiance in watts / (m^2 steradians)

5-dimensional function

The vector light field
[Gershun 1936]

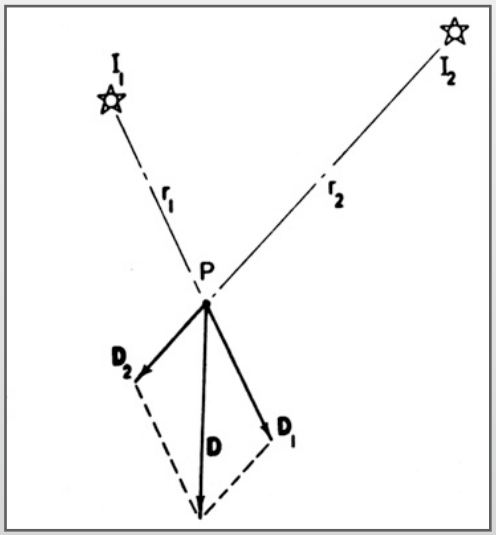

adding two light vectors

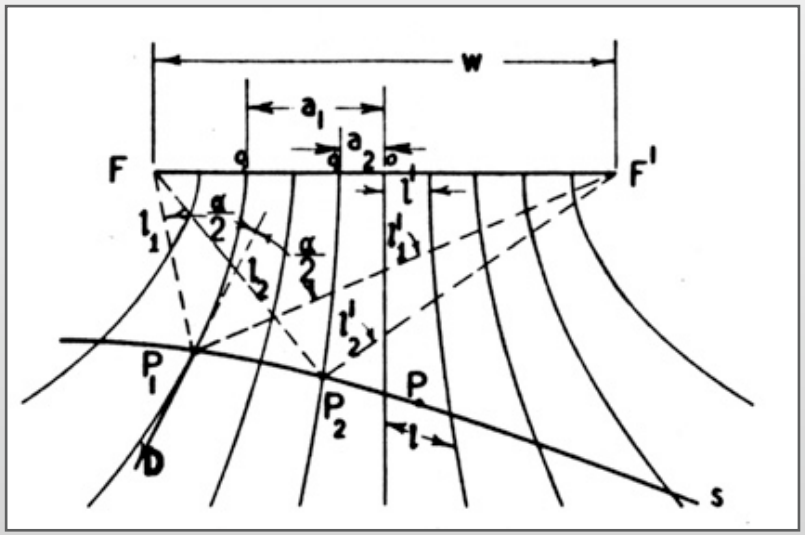

the vector light field
produced by a luminous strip

- amplitude gives irradiance at that point
- direction tells which way to orient a surface for maximum brightness under uniform illumination

Visualizing the vector irradiance field

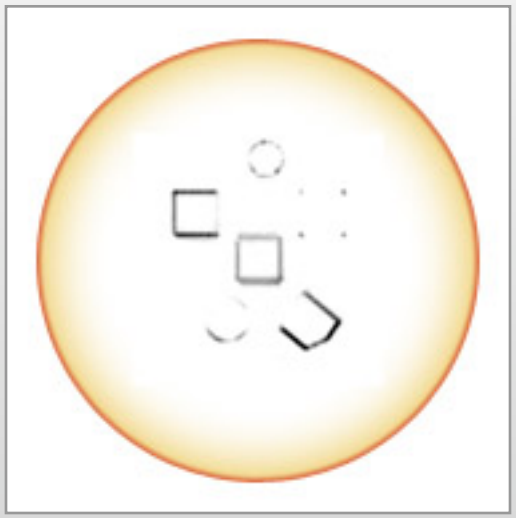

flatland scene with
partially opaque blockers
under uniform illumination

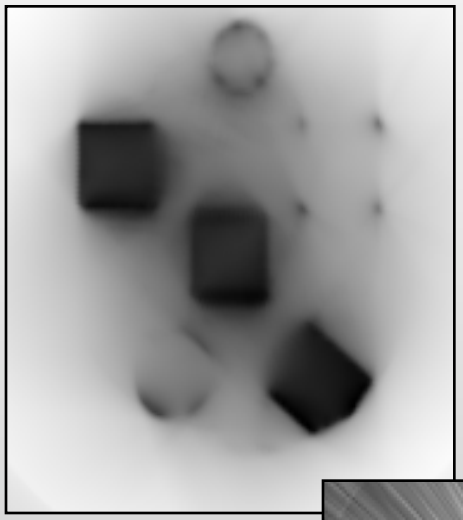

scalar irradiance
at each point

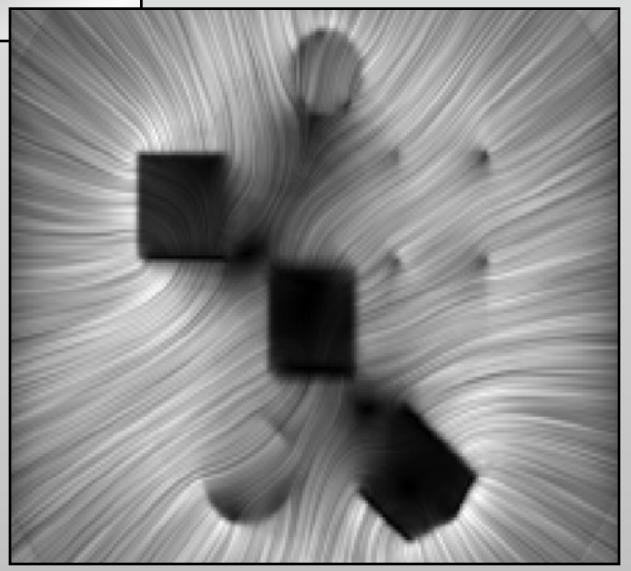

vector directions,visualized using
line integral convolution (LIC) [Cabral 1993]

Dimensionality of the scalar light field

- for general scenes

  $\Rightarrow$ 5D function

  "plenoptic function"

  $L (x, y, z, \theta, \phi)$

- in free space

  $\Rightarrow$ 4D function

  " the (scalar) light field"

  $L (?)$

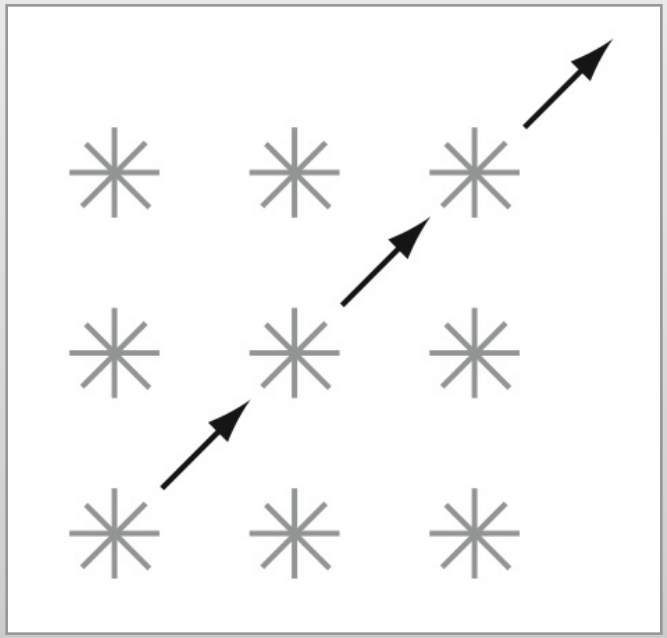

Some candidate parameterizations for the 4D light field

Point-on-plane + direction
(or point-on-surface + direction)

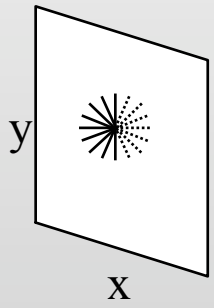

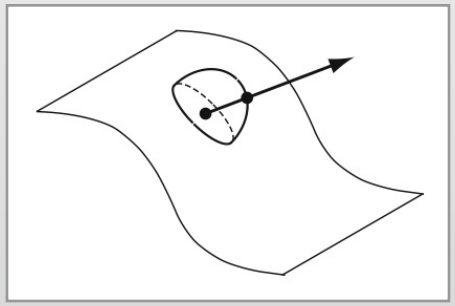

$L(x, y, \theta, \phi)$ or $L(u, v, \theta, \phi)$

- convenient for measuring BRDFs
- restriction to line gives looming field

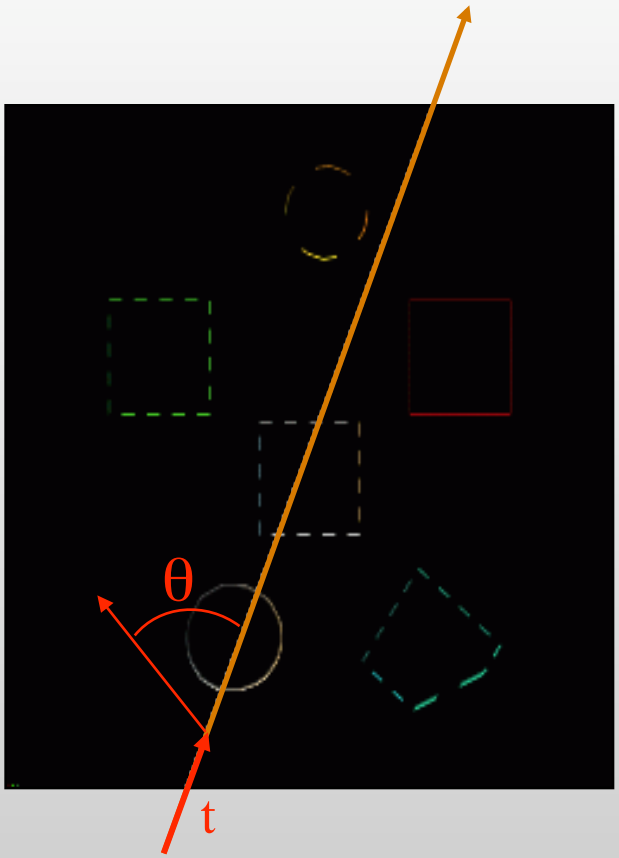

flight path through
a flatland scene

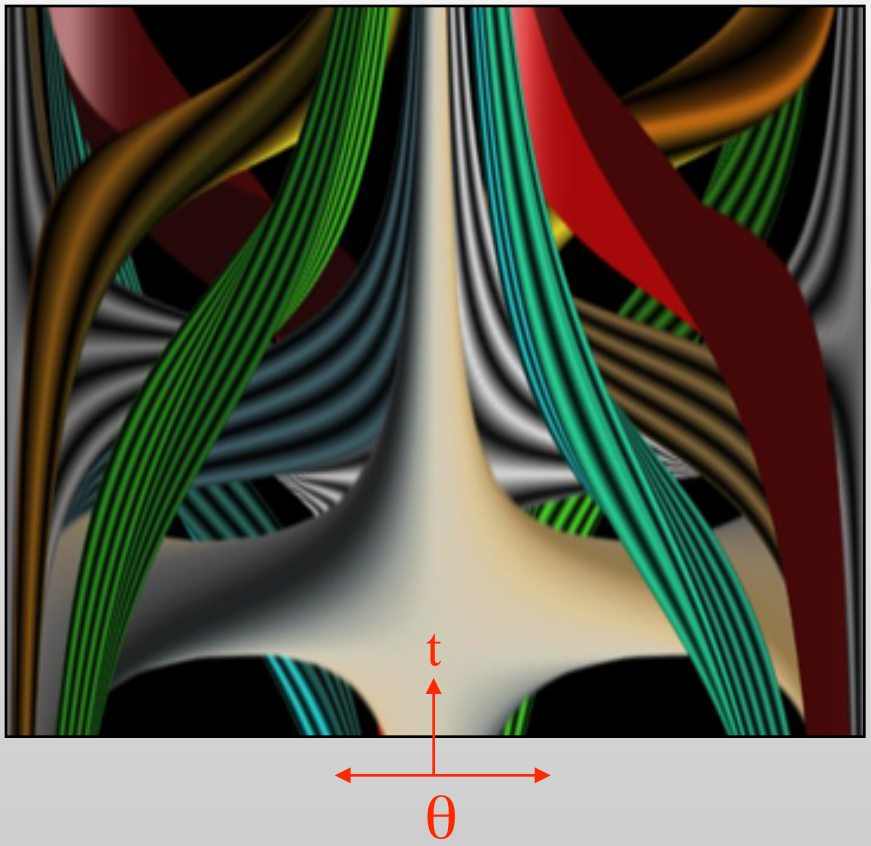

corresponding looming light field
(see also [Hasinoff 2006])

The looming field

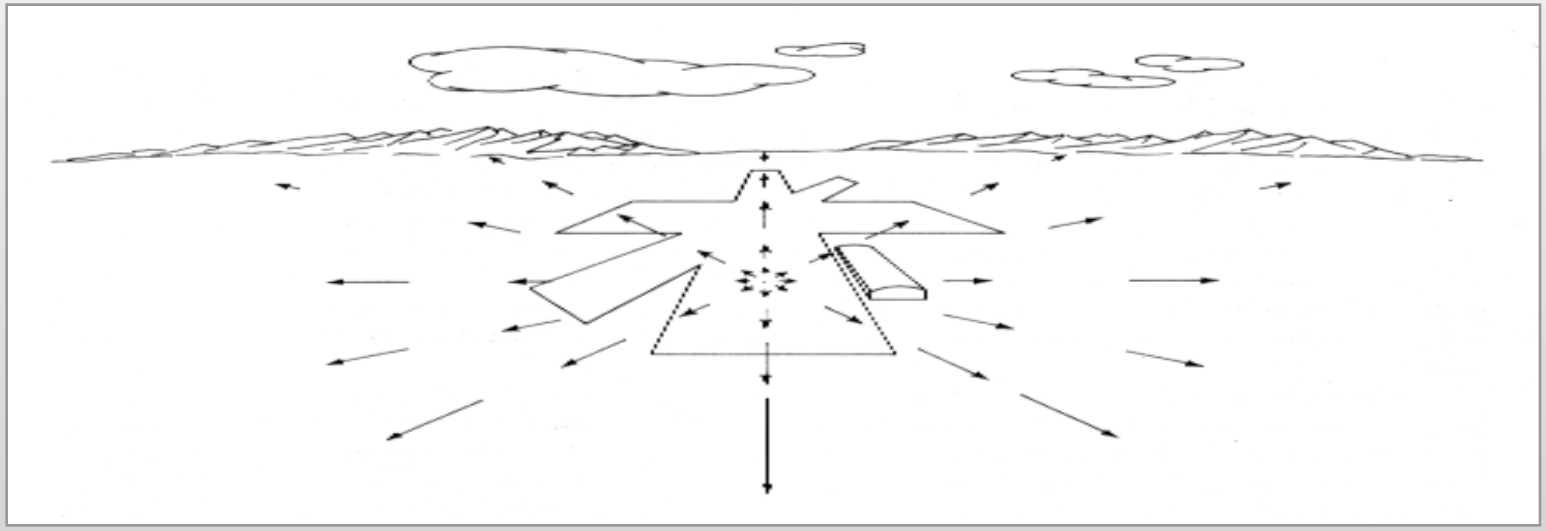

[Gibson]

More parameterizations

Chords of a sphere

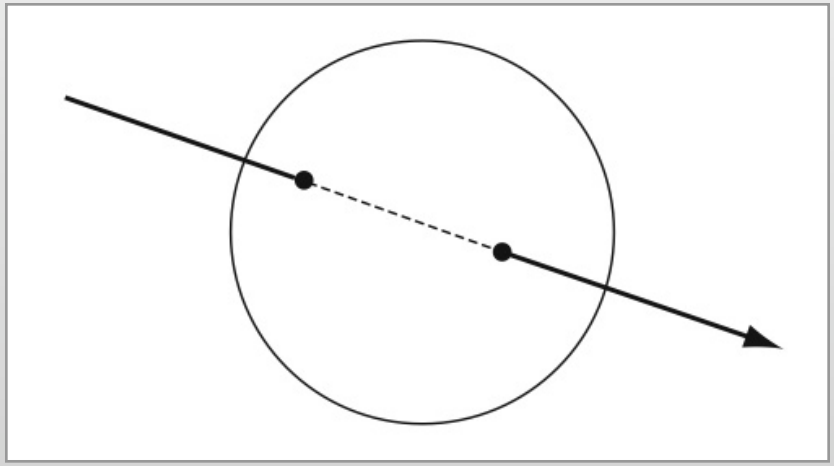

$L (\theta_1, \phi_1, \theta_2, \phi_2)$

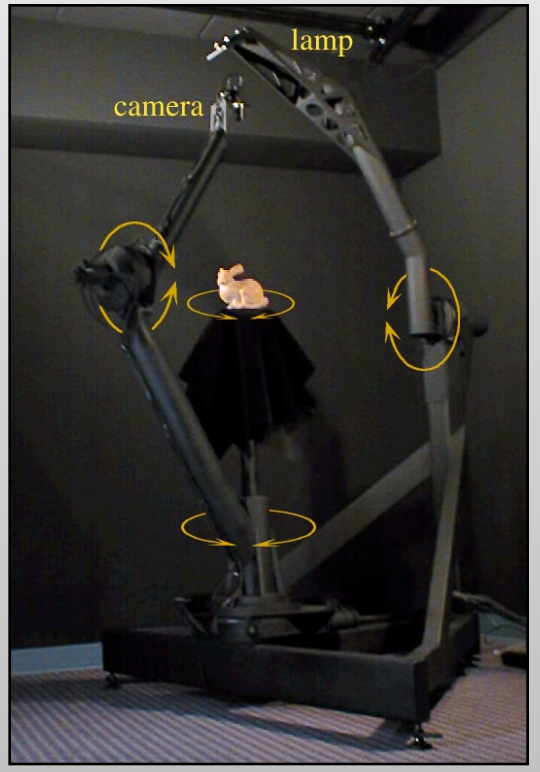

- convenient for spherical gantry
- facilitates uniform sampling

Two planes ("light slab")

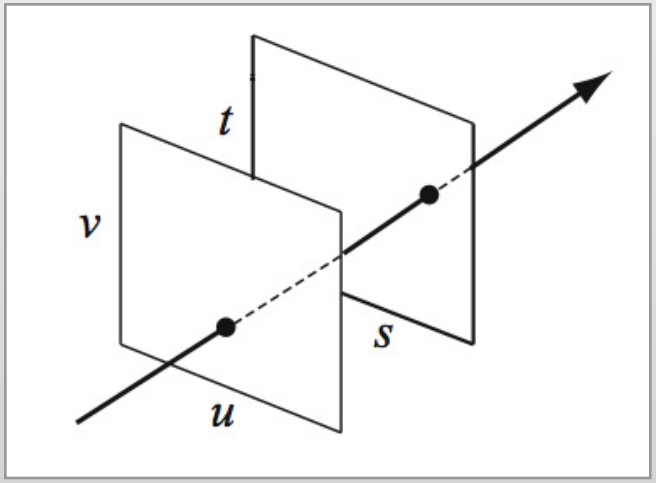

L (u, v, s, t)

- uses projective geometry
  - one plane at infinity $\Rightarrow$ array of orthographic images
  - fast incremental display algorithms

The free-space assumption

- Where can you use free-space light fields?
  - the 3D space around a compact object
  - the 3D space inside an uncluttered environment

- stitching together light fields

  [Chen, Levoy, Hanrahan (unpublished)]

  - partition scene into disjoint cells
  - links between cells are light fields
  - hierarchy of cells, links, light fields
  - # of light fields is linear in # of cells

Light field rendering

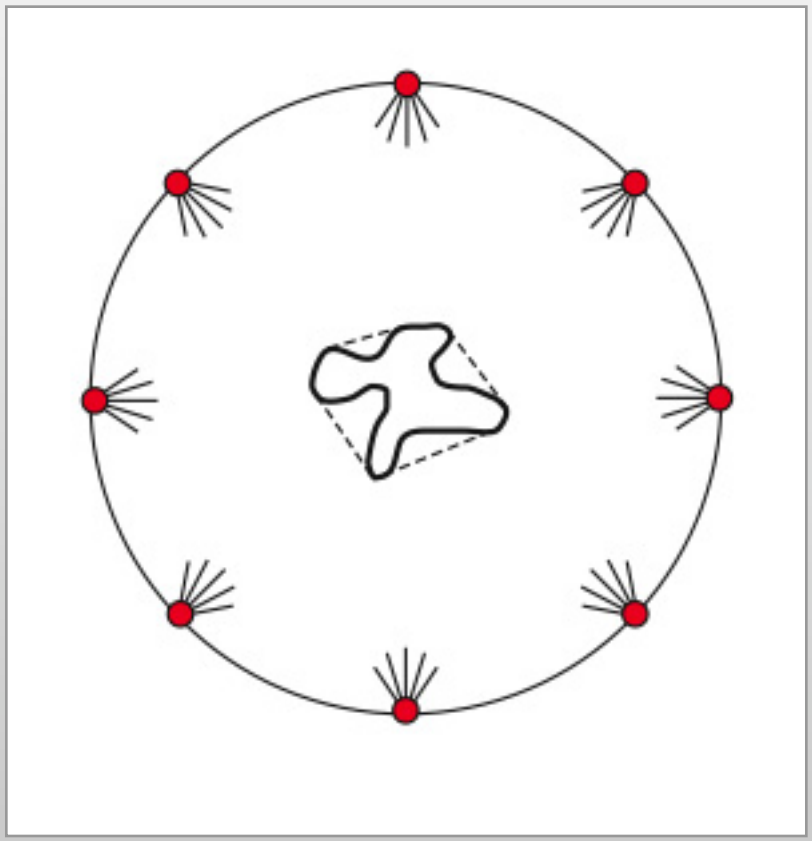

flipbook animation
(QuickTime VR)

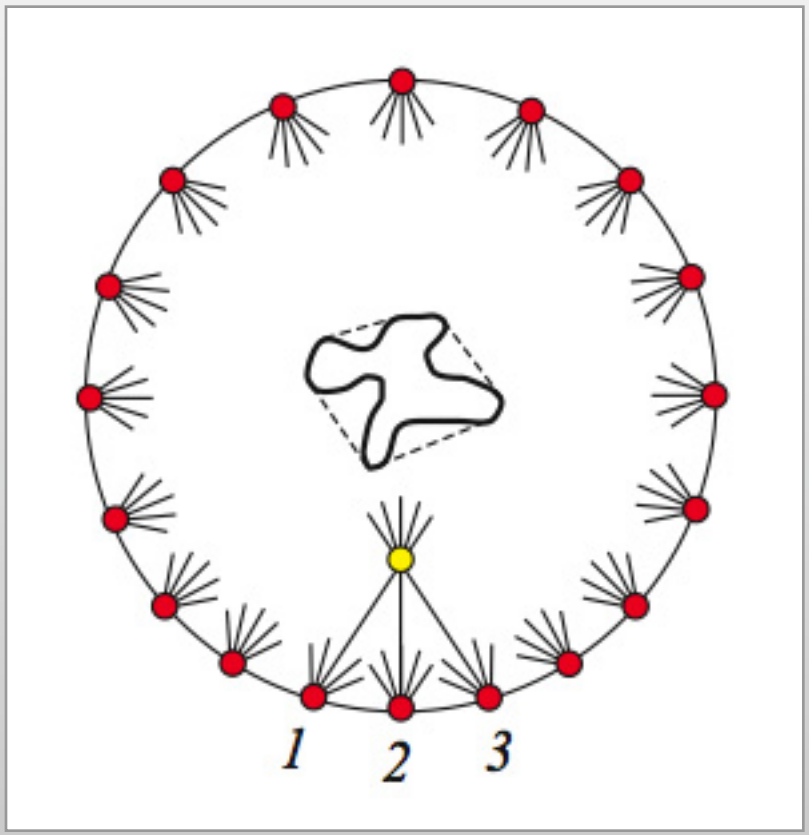

rebinning the rays
to create new views

(movie is available at http://graphics.stanford.edu/papers/light)

Alternative parameterizations for the 5D plenoptic function

- Two-plane ray field

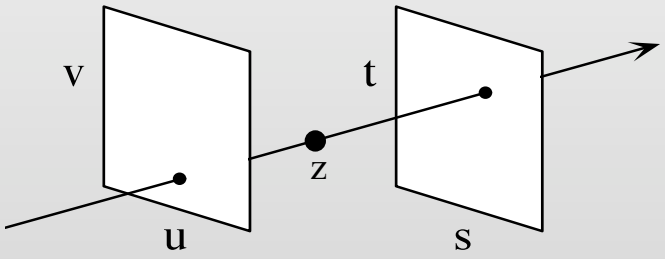

L (u, v, s, t, z)

- allows multiple colors, in sequence, along one line
- alternative to L (x, y, z, θ, ϕ)
- inspired by Salesin's ZZ-buffer [1989]

Devices for recording light fields

big scenes

small scenes

- handheld camera [Buehler 2001]
- array of cameras [Wilburn 2005]
- plenoptic camera [Ng 2005]
- light field microscope [Levoy 2006]

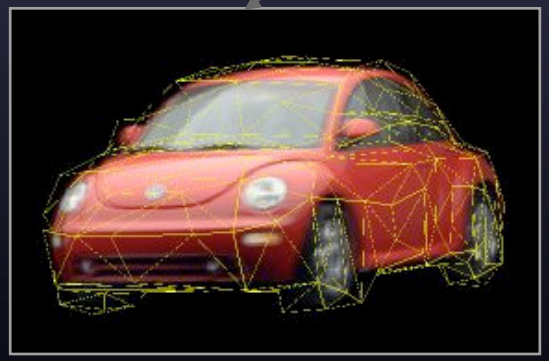

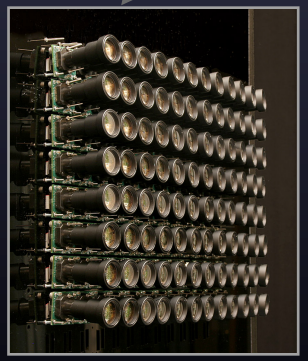

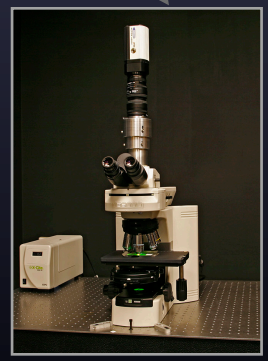

Stanford Multi-Camera Array

[Wilburn SIGGRAPH 2005]

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

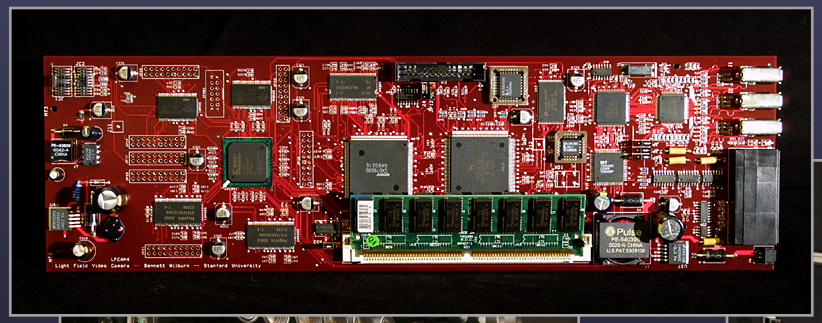

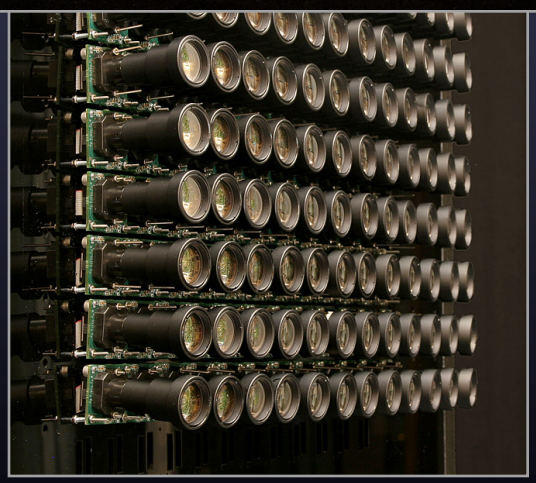

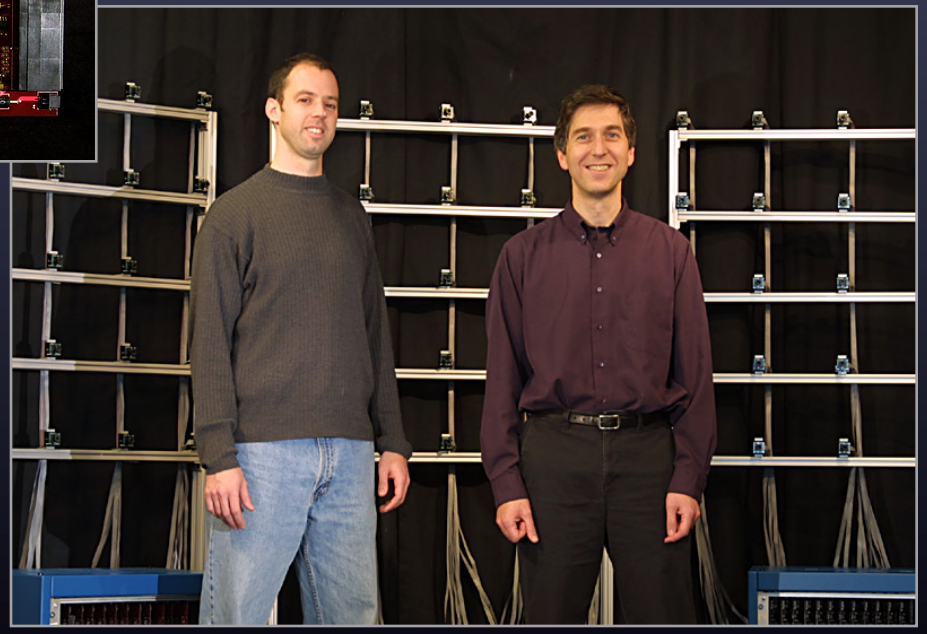

Ways to use large camera arrays

- widely spaced → light field capture

Manex's bullet time array

(movie is available at http://graphics.stanford.edu/projects/array)

Ways to use large camera arrays

- widely spaced → light field capture
- tightly packed → high-performance imaging

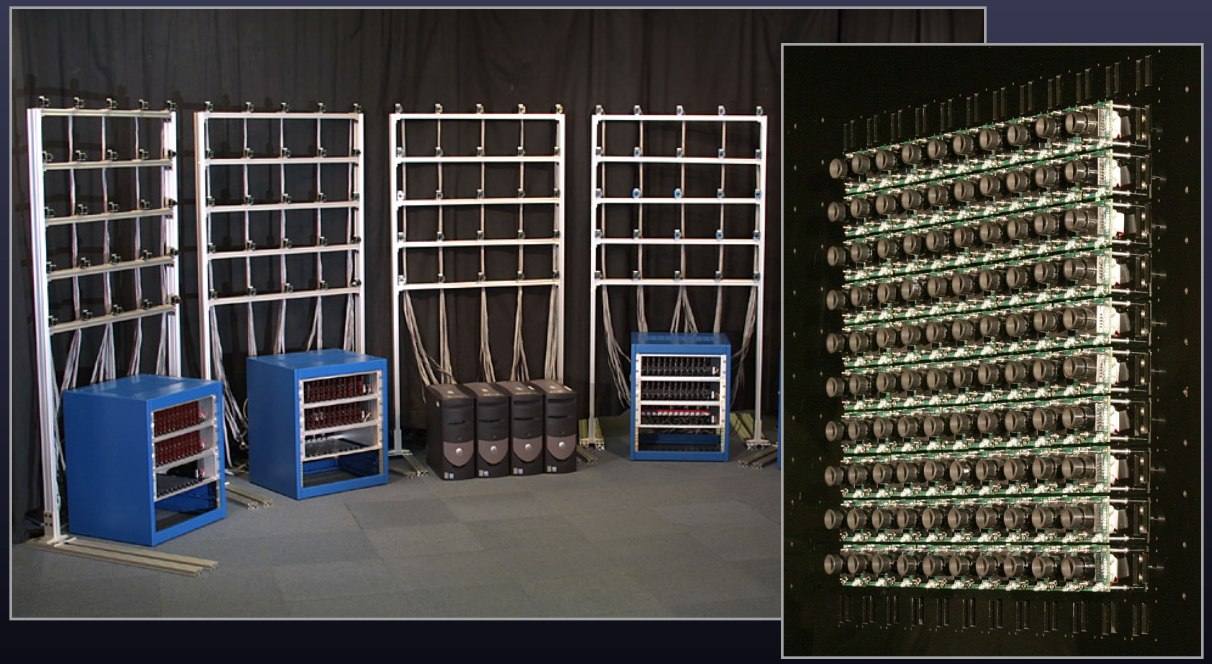

Ways to use large camera arrays

- widely spaced → light field capture
- tightly packed → high-performance imaging
- intermediate spacing → synthetic aperture photography

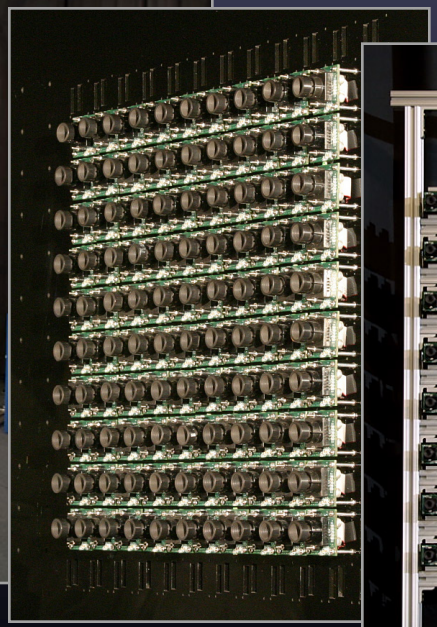

Synthetic aperture photography

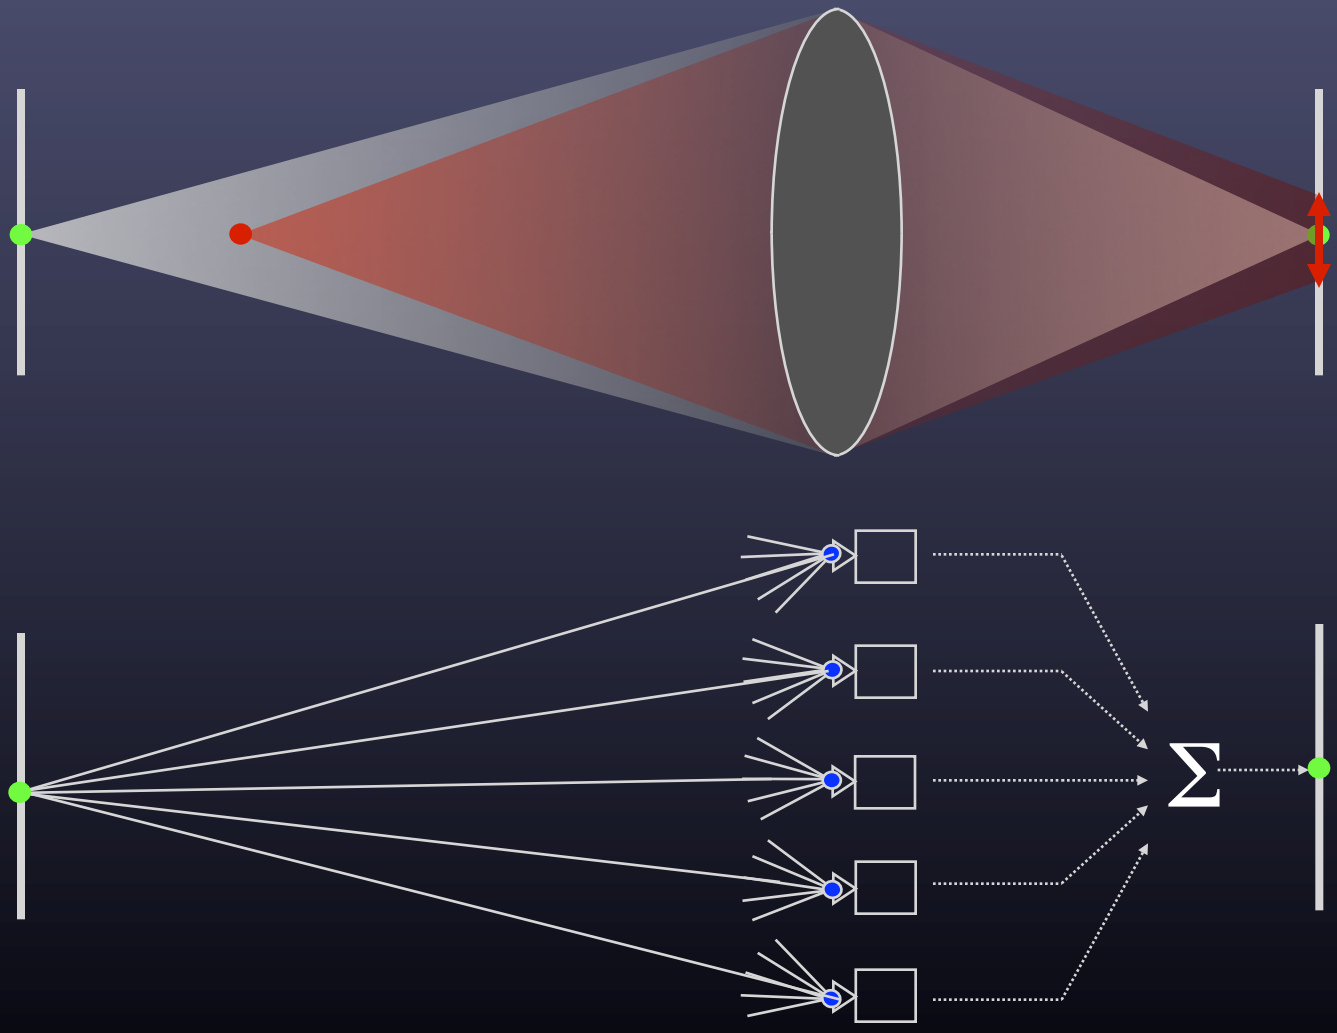

Example using 45 cameras
[Vaish CVPR 2004]

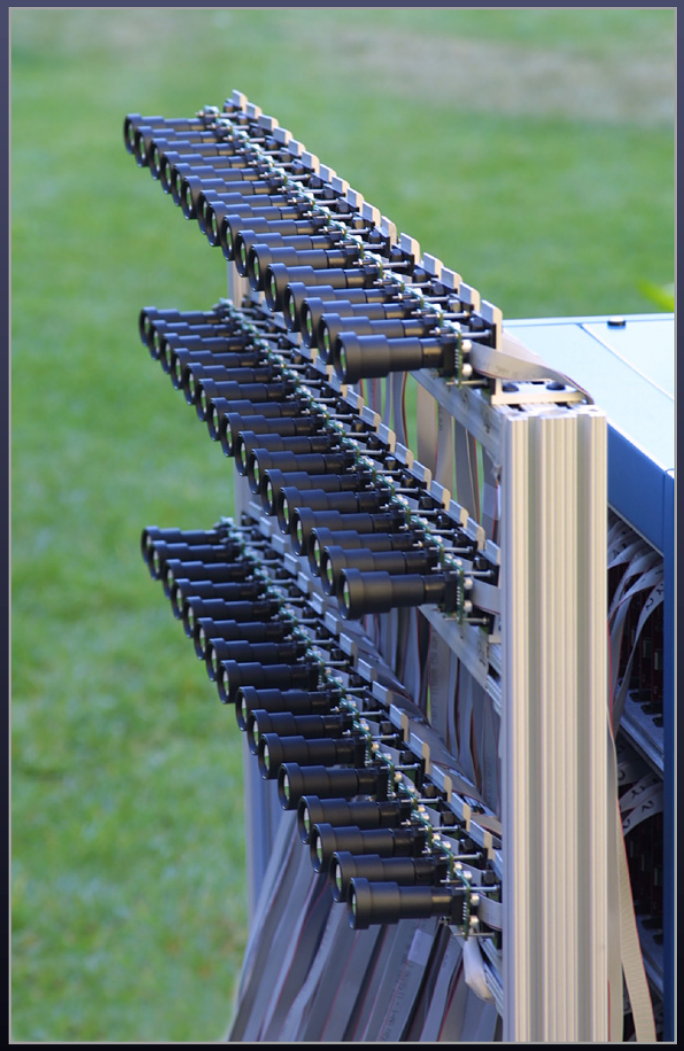

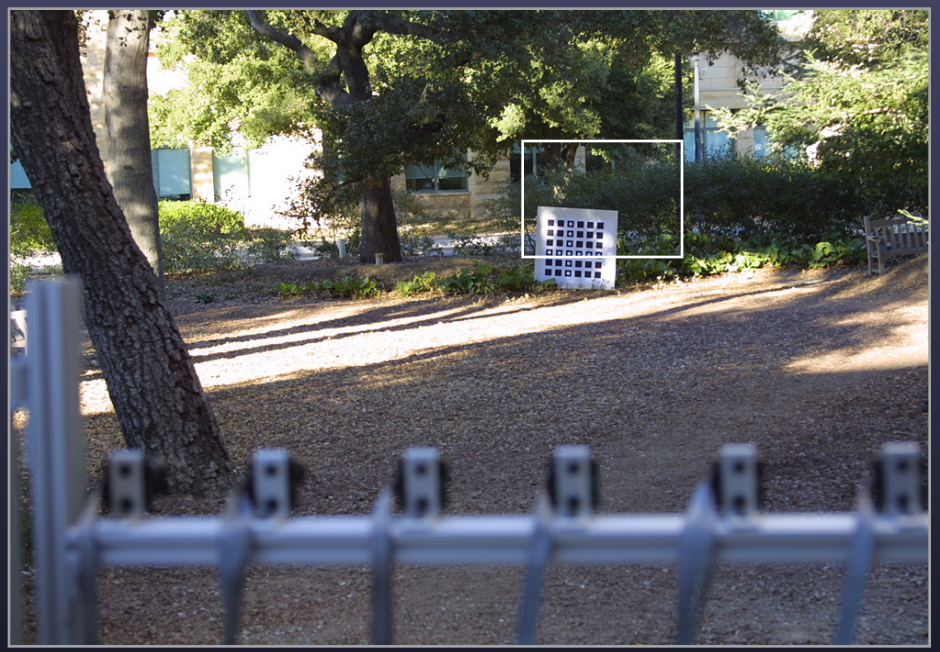

one camera’s view

synthetic aperture view

(movie is available at http://graphics.stanford.edu/projects/array)

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)

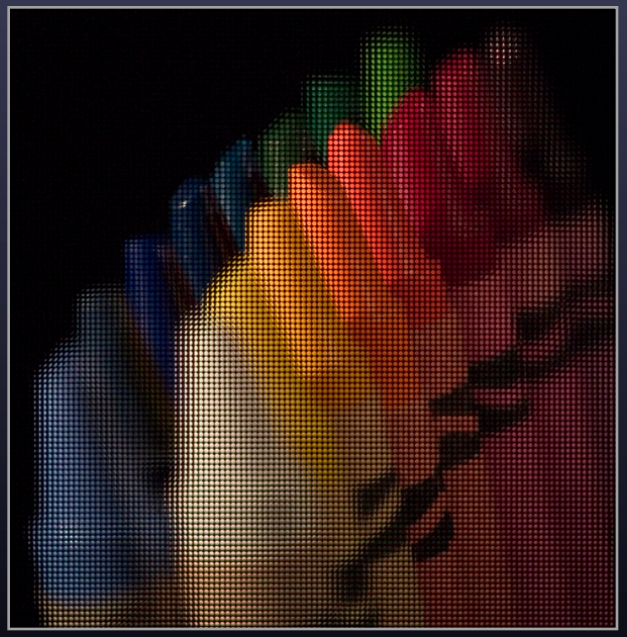

Conventional versus light field camera

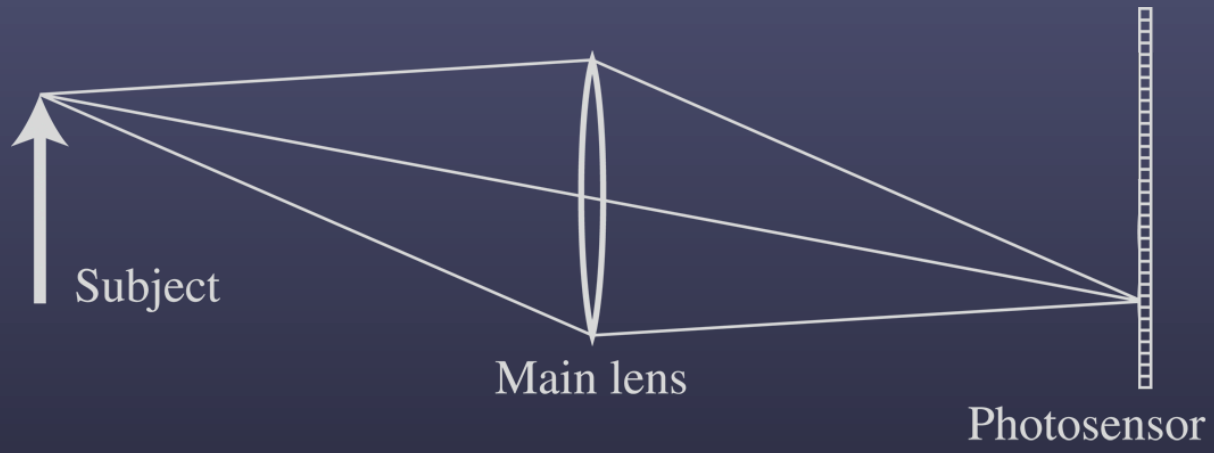

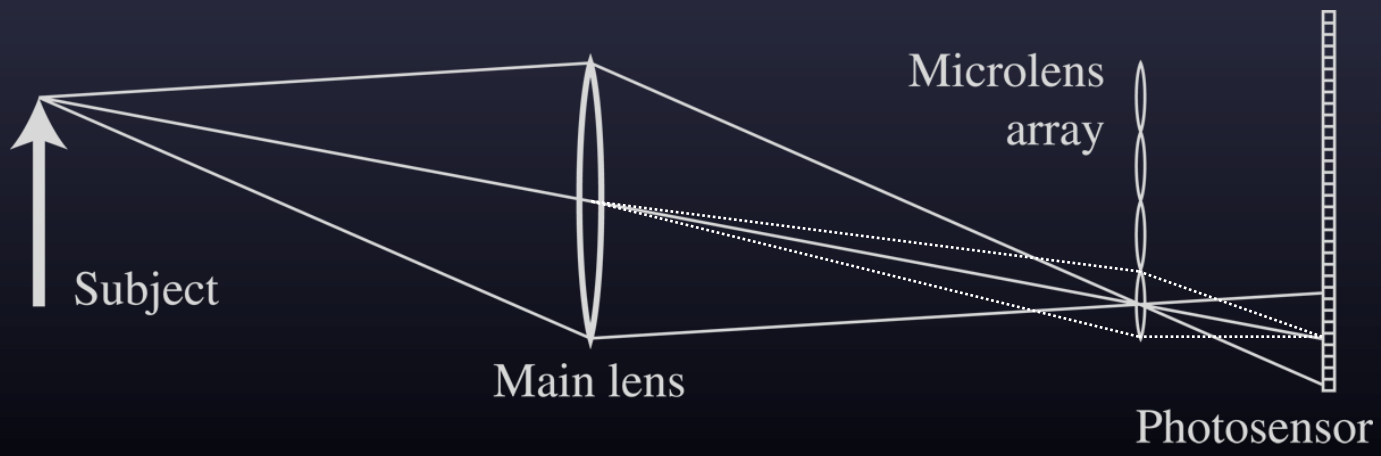

Conventional versus light field camera

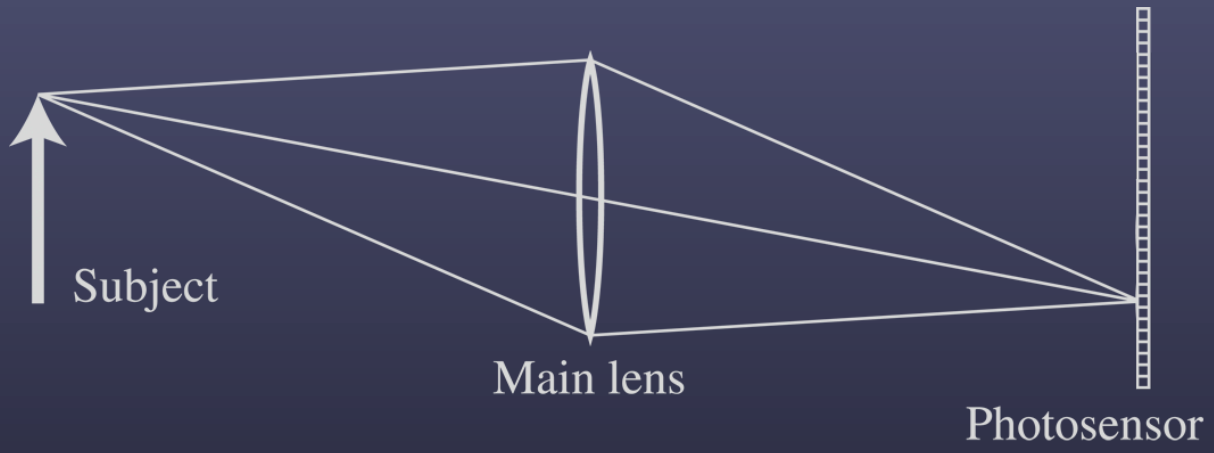

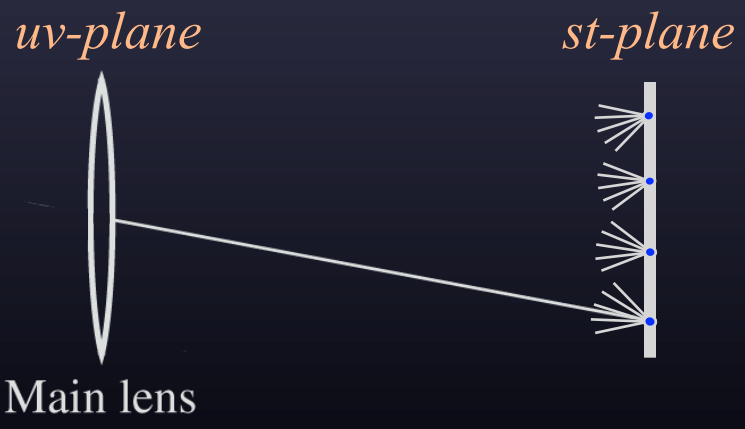

Prototype camera

Contax medium format camera

Kodak 16-megapixel sensor

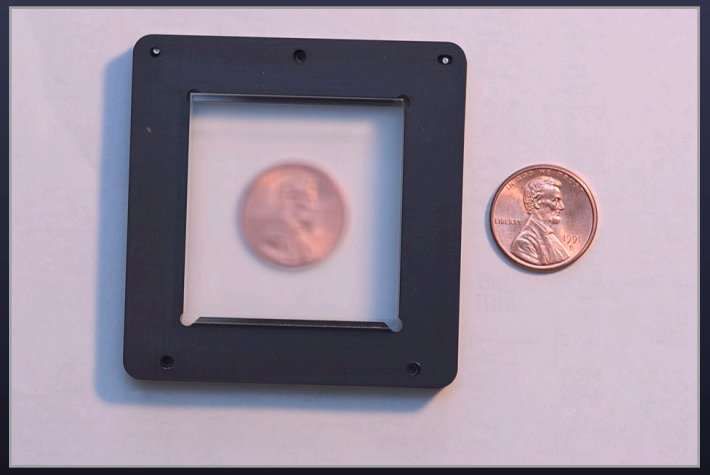

Adaptive Optics microlens array

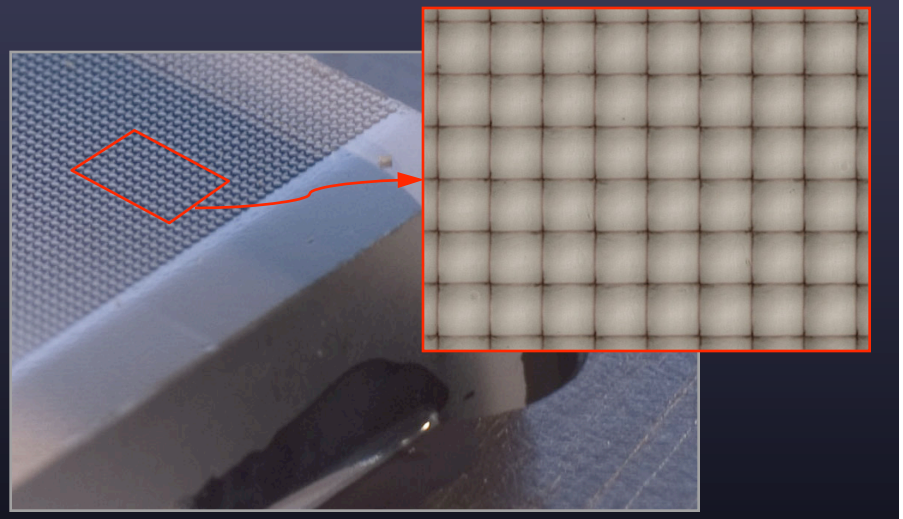

125µ square-sided microlenses

4000 × 4000 pixels ÷ 292 × 292 lenses = 14 × 14 pixels per lens

Typical image captured by camera (show here at low res)

Digital refocusing

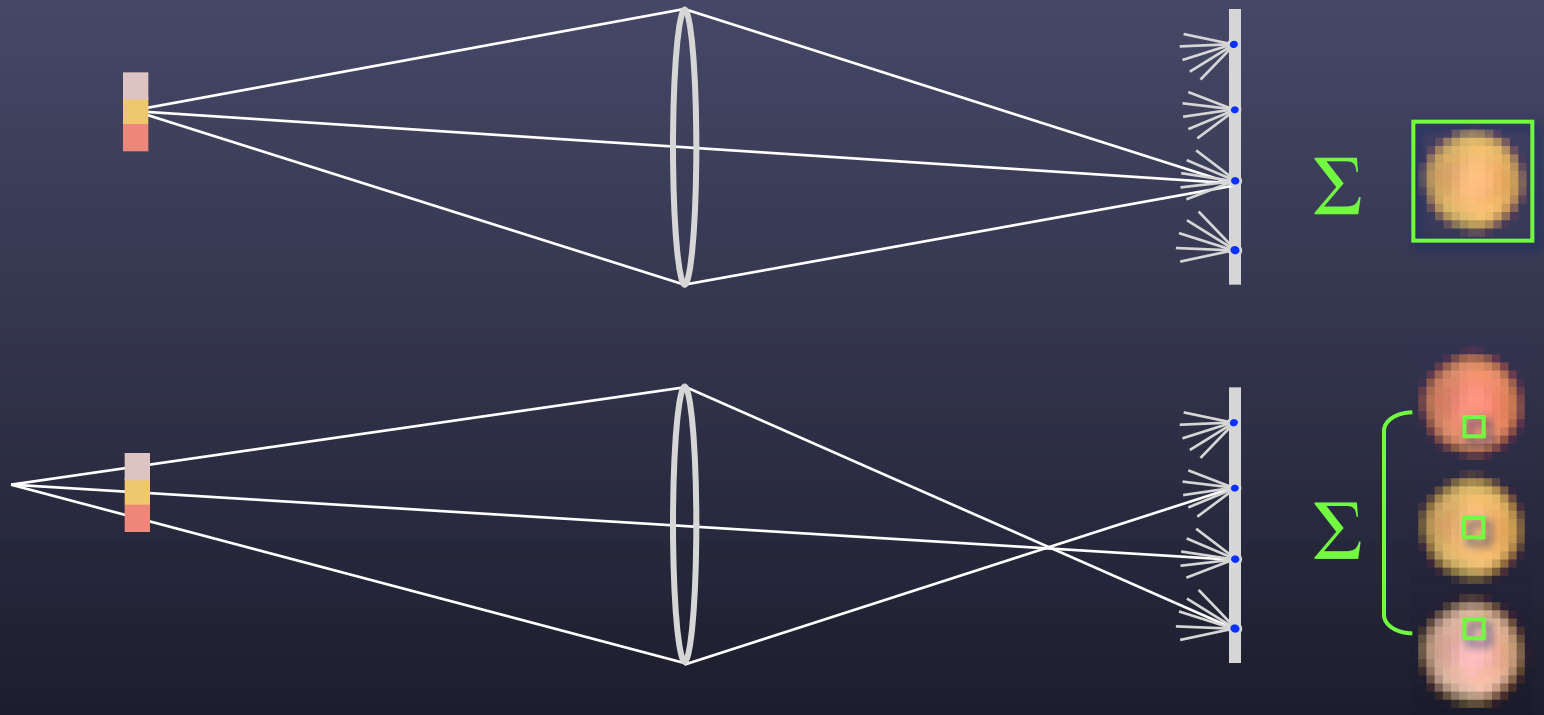

- refocusing = summing windows extracted from several microlenses

A digital refocusing theorem

- an f / N light field camera, with P × P pixels under each microlens, can produce views as sharp as an f / (N × P) conventional camera

– or –

- it can produce views with a shallow depth of field (f / N) focused anywhere within the depth of field of an f / (N × P) camera

Example of digital refocusing

Example of digital refocusing

Example of digital refocusing

Example of digital refocusing

Example of digital refocusing

Refocusing portraits

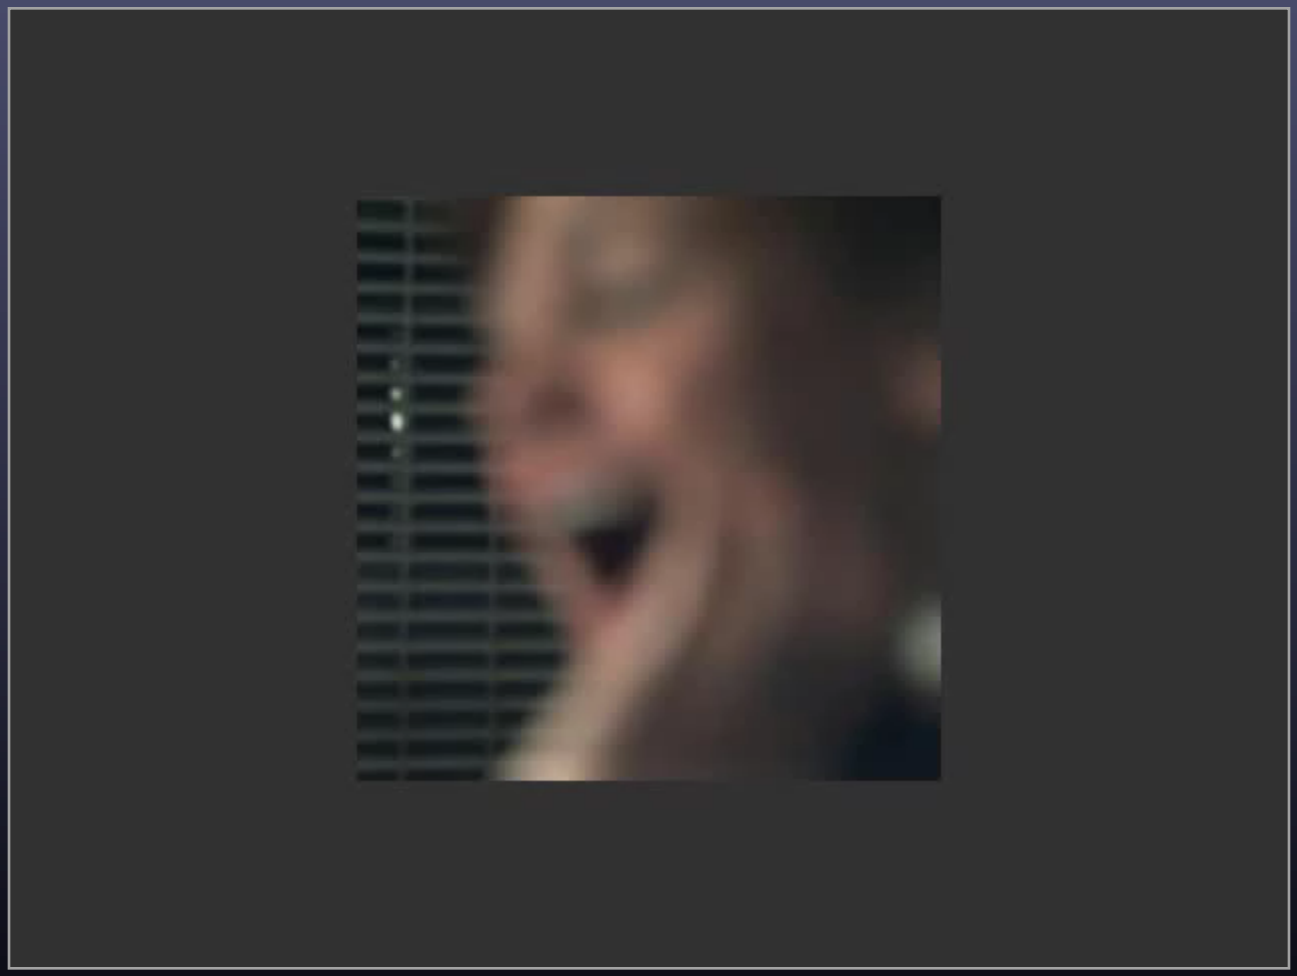

(movie is available at http://refocusimaging.com)

Extending the depth of field

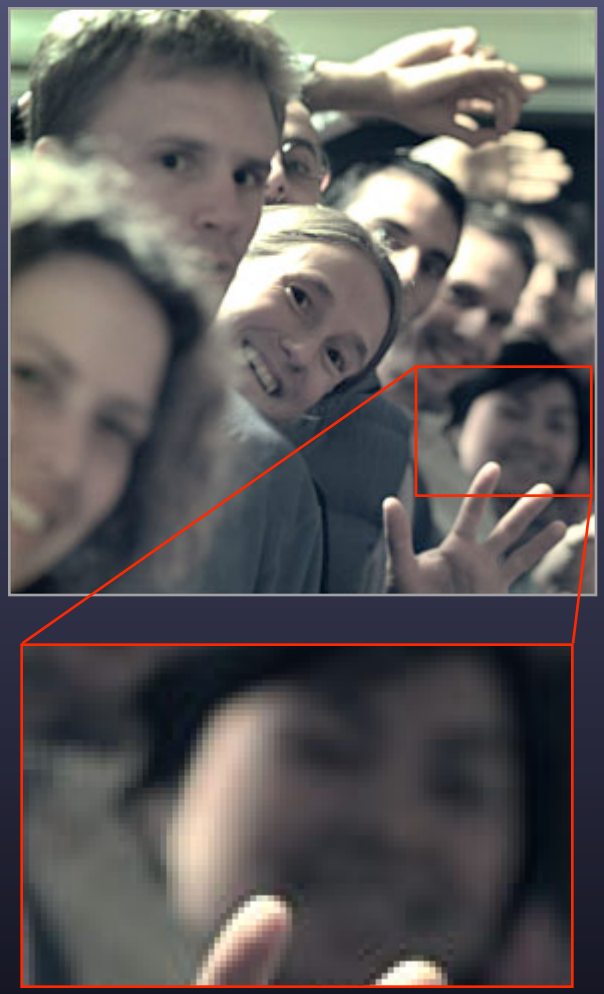

conventional photograph, main lens at f / 4

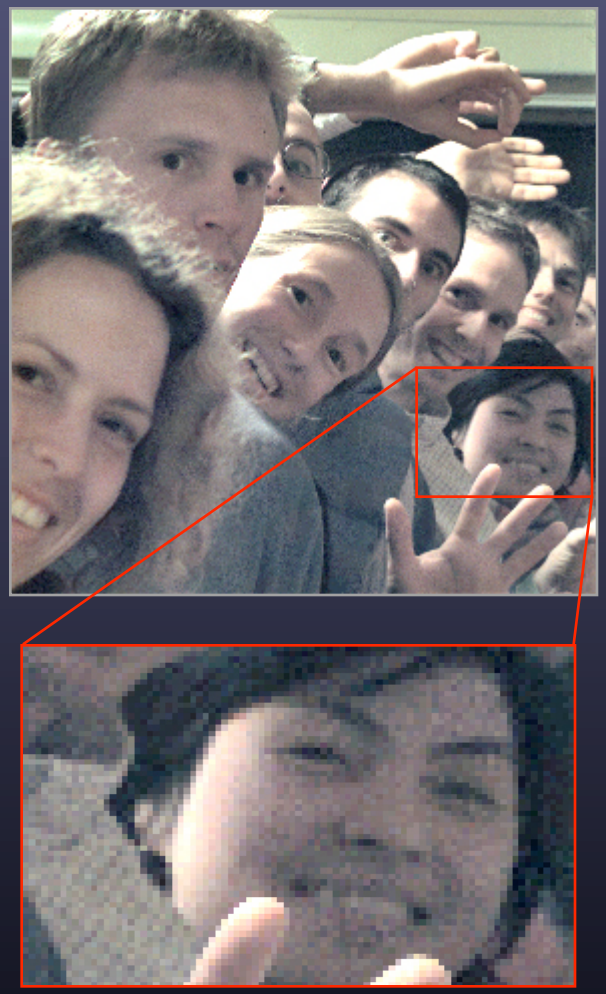

conventional photograph, main lens at f / 22

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

Digitally moving the observer

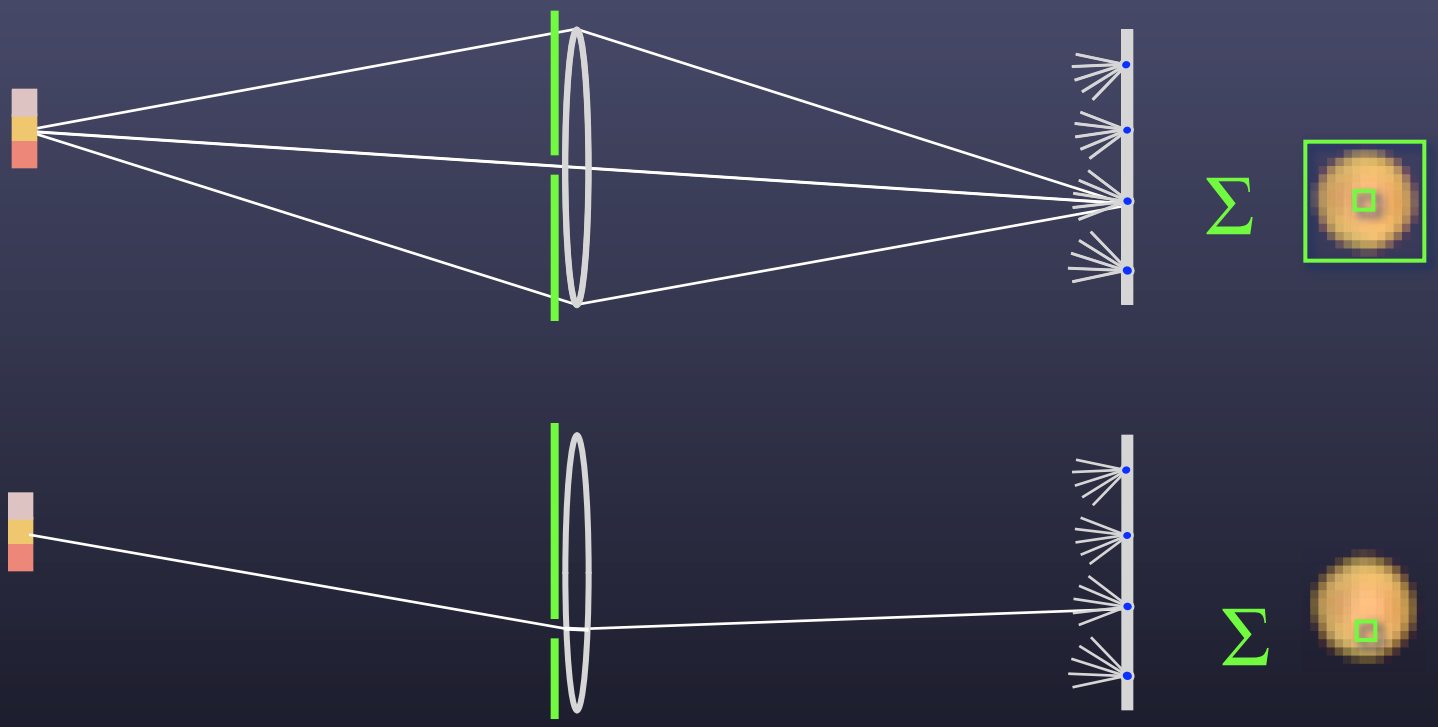

- moving the observer = moving the window we extract from the microlenses

Example of moving the observer

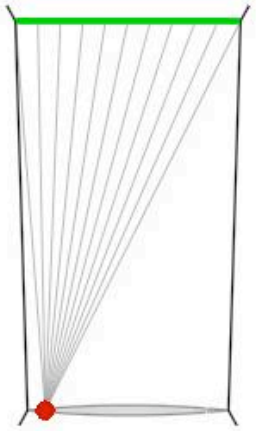

Example of moving the observer

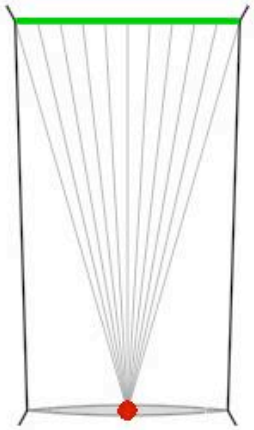

Example of moving the observer

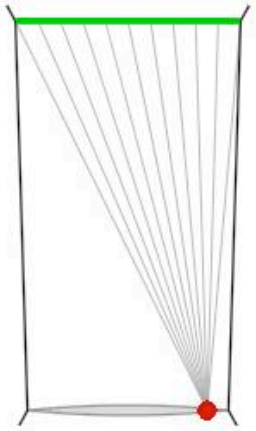

Moving backward and forward

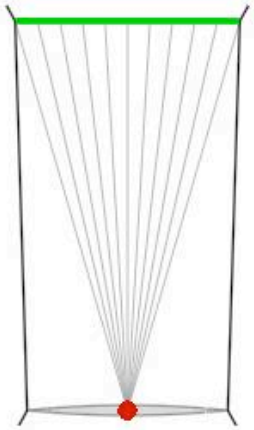

Moving backward and forward

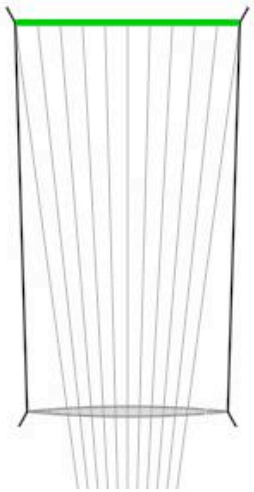

Moving backward and forward

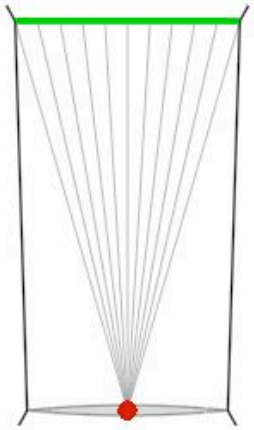

Moving backward and forward

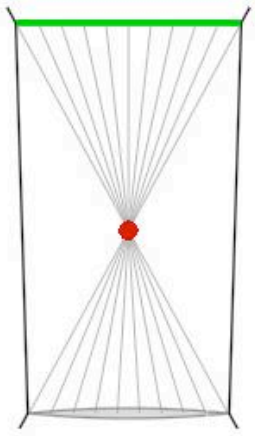

Lego gantry for capturing light fields
(built by Andrew Adams)

Flash-based viewer for light fields
(written by Andrew Adams)

(light field can be viewed at http://lightfield.stanford.edu/lfs.html)

Implications / commercialization
(see refocusimaging.com)

- 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 Mpix

  20K × 13K pixels

  4000 × 2666 pixels × 20 × 20 rays per pixel

or

  2000 × 1500 pixels × 3 × 3 rays per pixel = 27 Mpix

Light Field Microscopy

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

(Proc. SIGGRAPH 2006)

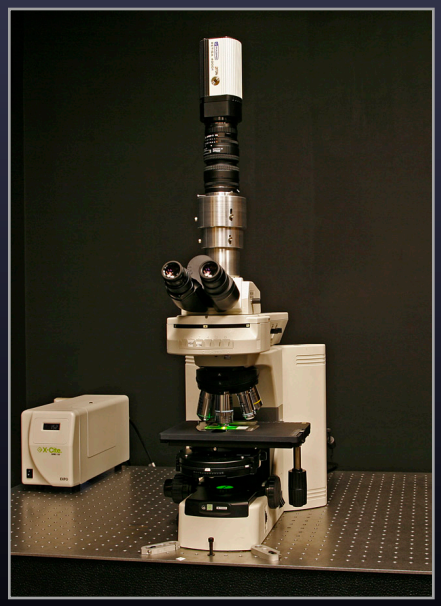

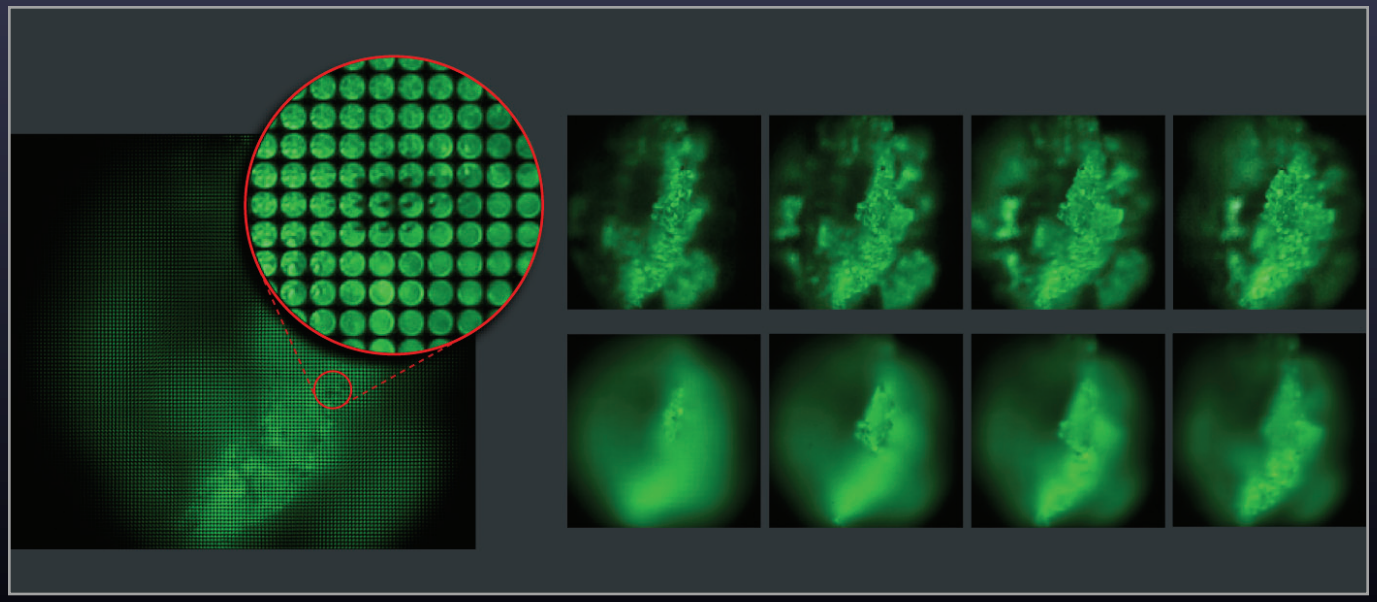

A traditional microscope

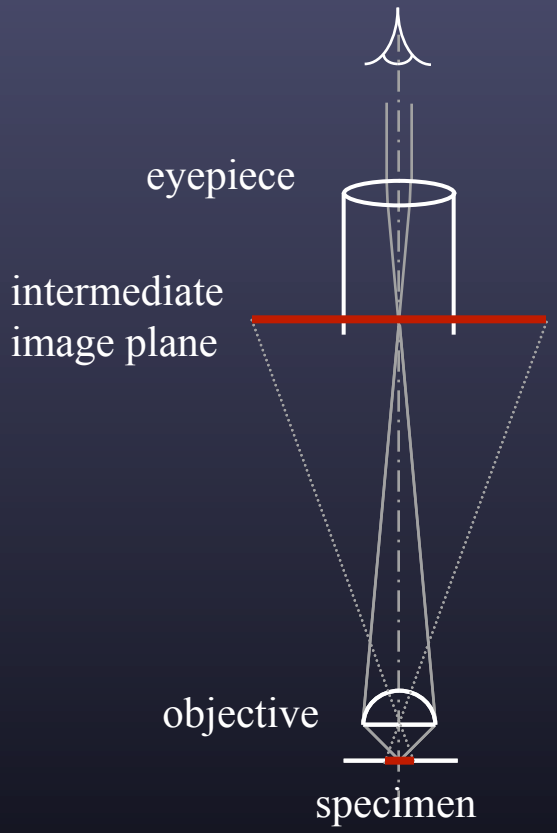

A light field microscope (LFM)

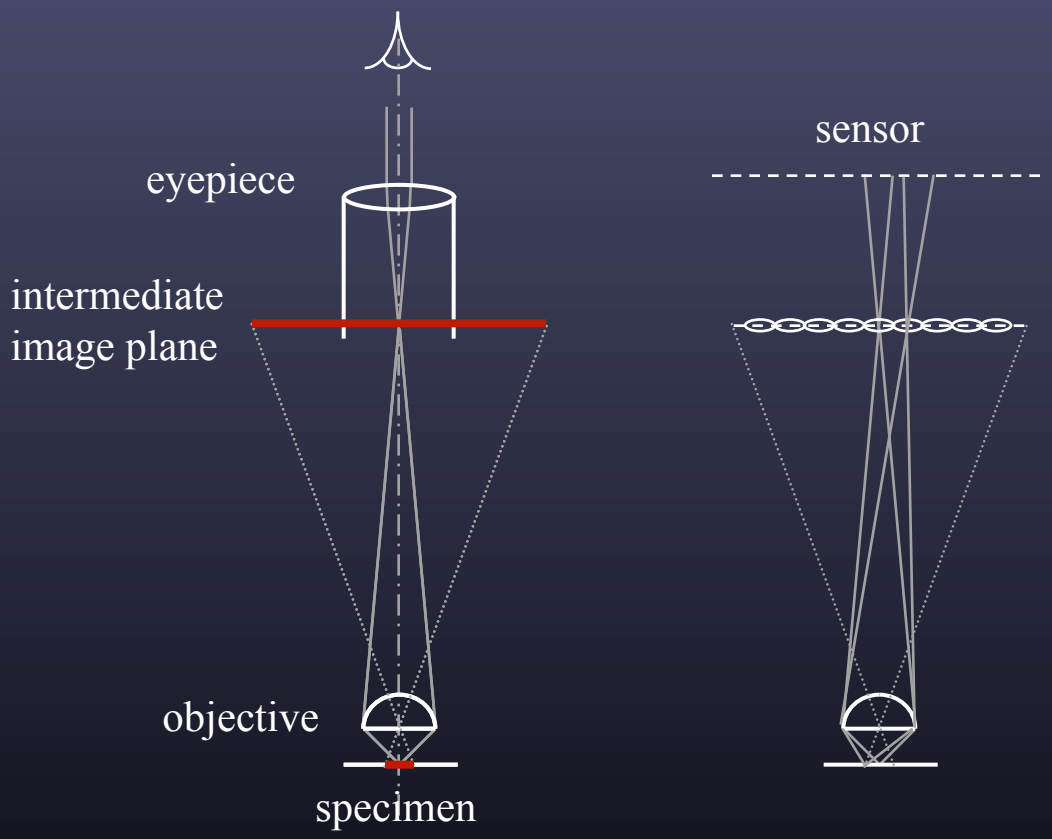

- 40x / 0.95NA objective

↓

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

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

A light field microscope (LFM)

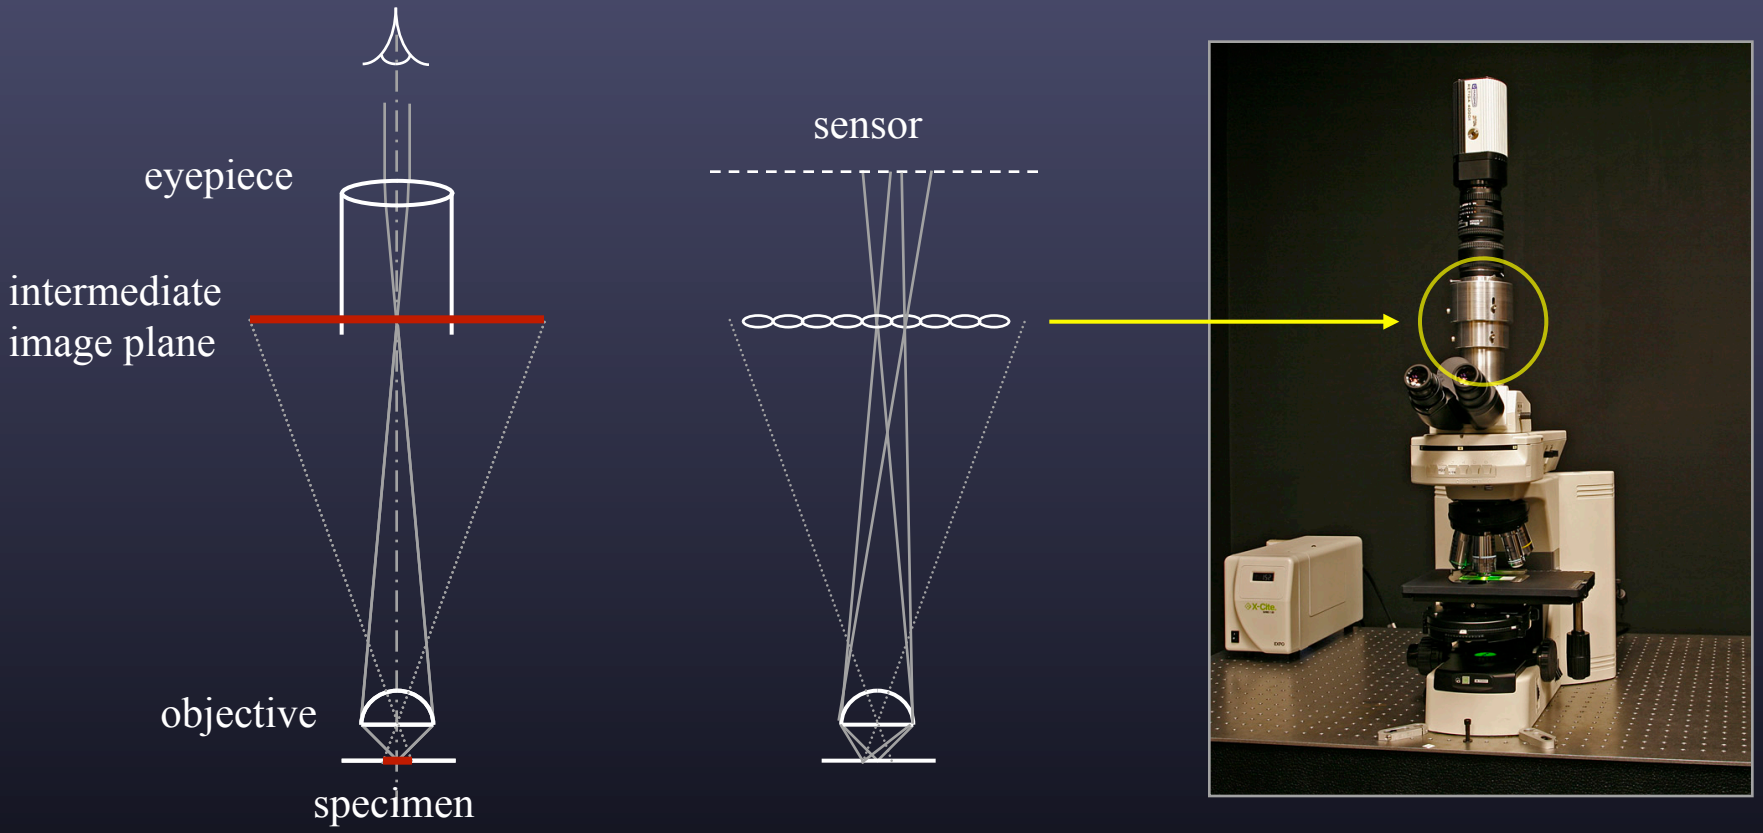

Example light field micrograph

- 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

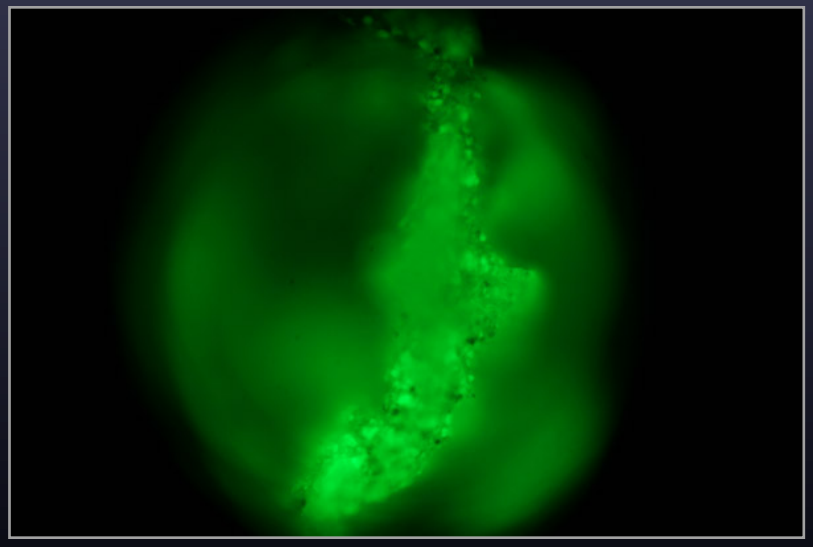

ordinary microscope

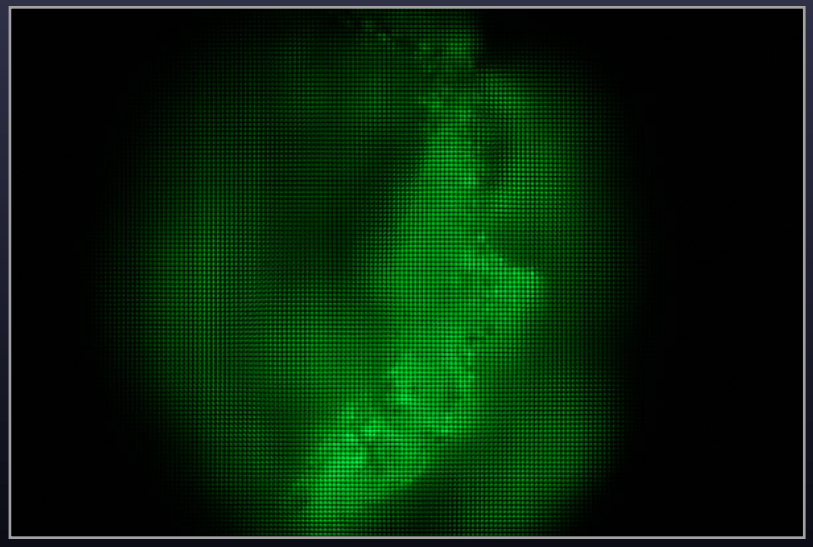

light field microscope

The geometry of the light field in a microscope

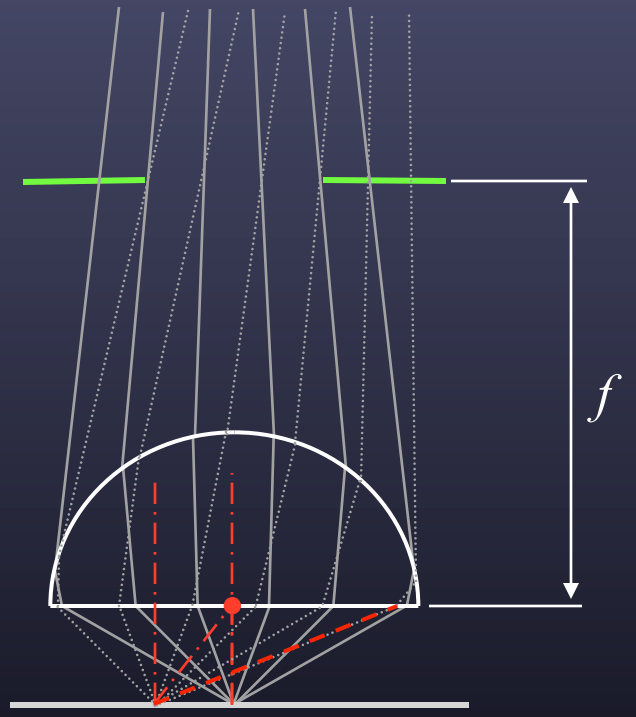

objective lenses are telecentric

- 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)

Example light field micrograph

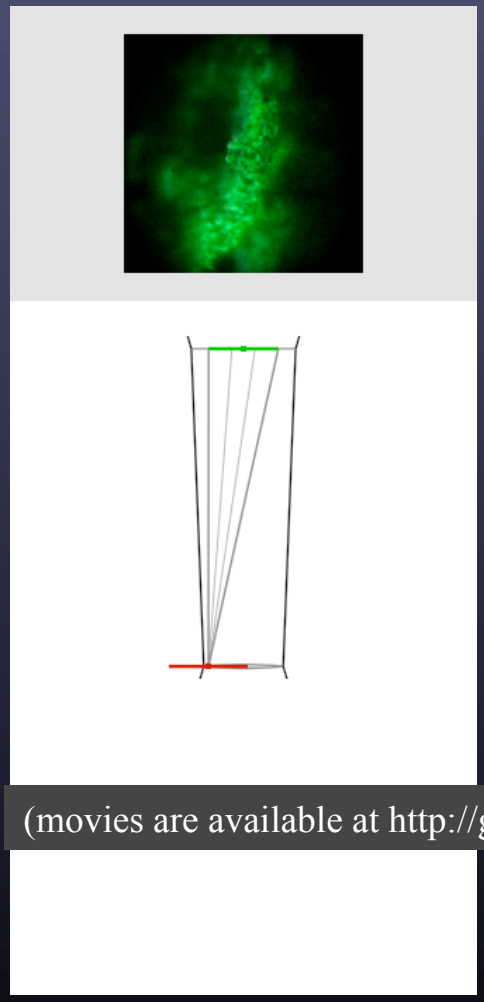

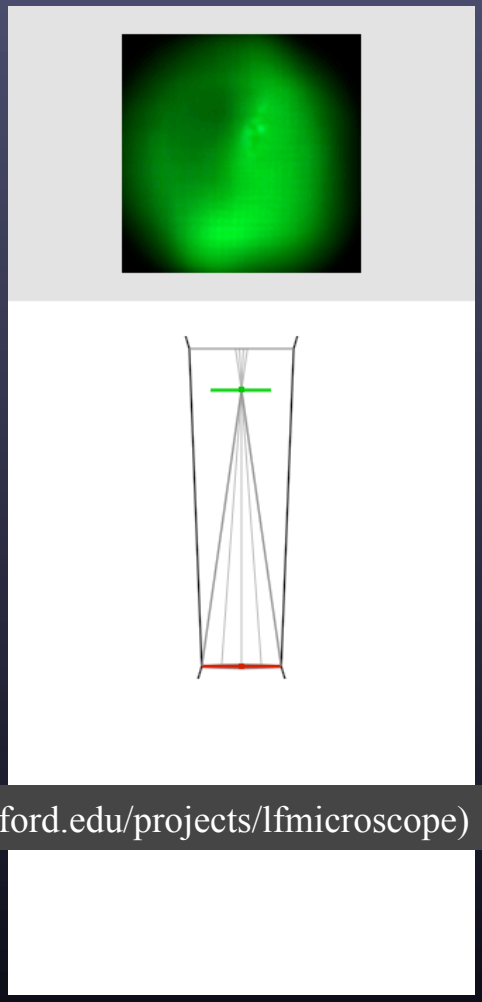

(movies are available at http://graphics.stanford.edu/projects/lfmicroscope)

panning sequence

focal stack

Real-time viewer

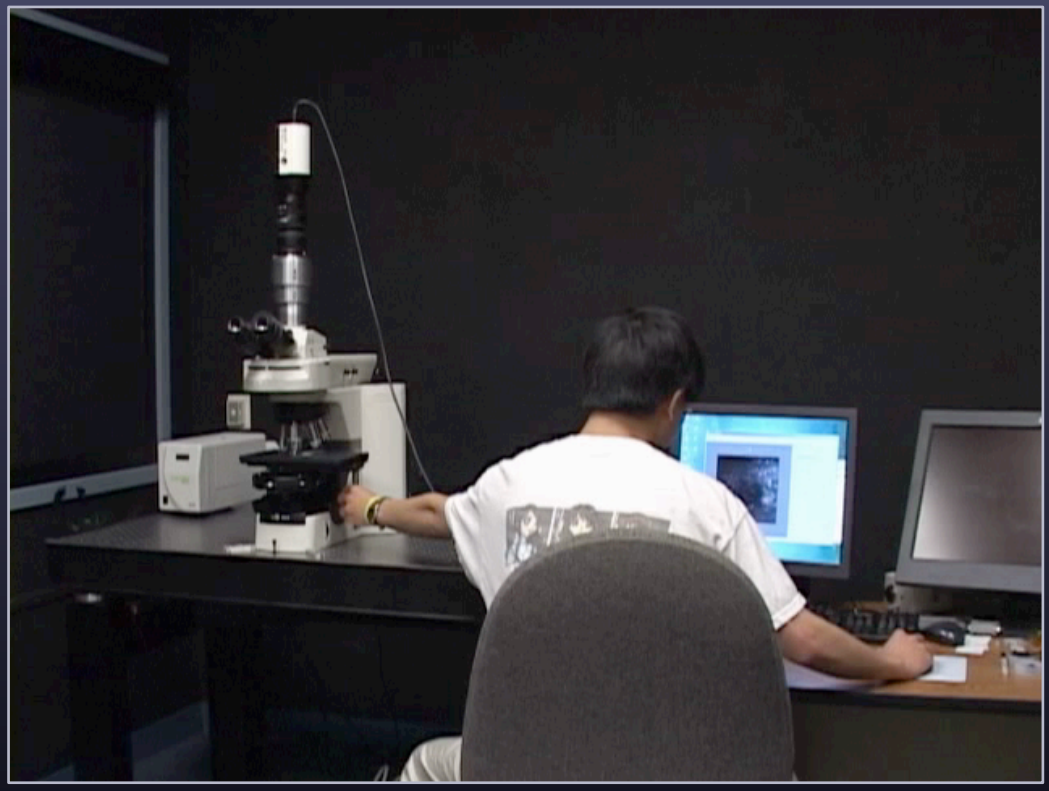

(movie is available at http://graphics.stanford.edu/projects/lfmicroscope/2007.html)

Other examples

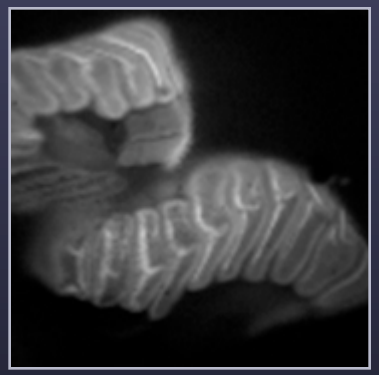

fern spore
(60x, autofluorescence)

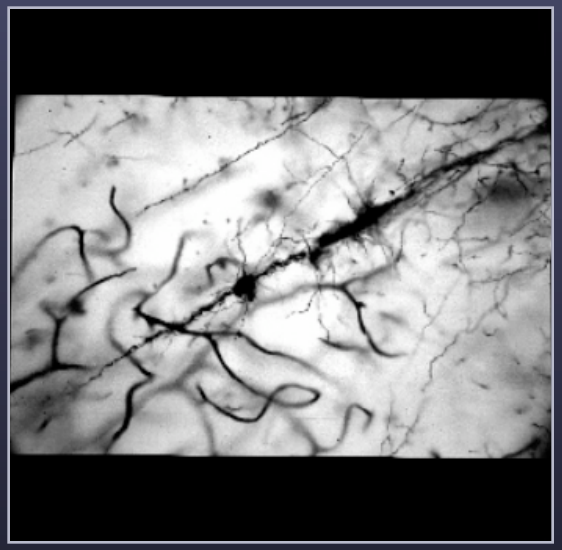

Golgi-stained neurons
(40x, transmitted light)

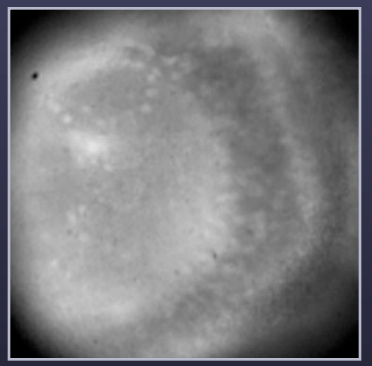

zebrafish optic tectum
(calcium imaging
of neural activity)

(movies are available at http://graphics.stanford.edu/projects/lfmicroscope)

3D reconstruction

- 4D light field → *digital refocusing* →
  3D focal stack → *deconvolution microscopy* →
  3D volume data

(DeltaVision)

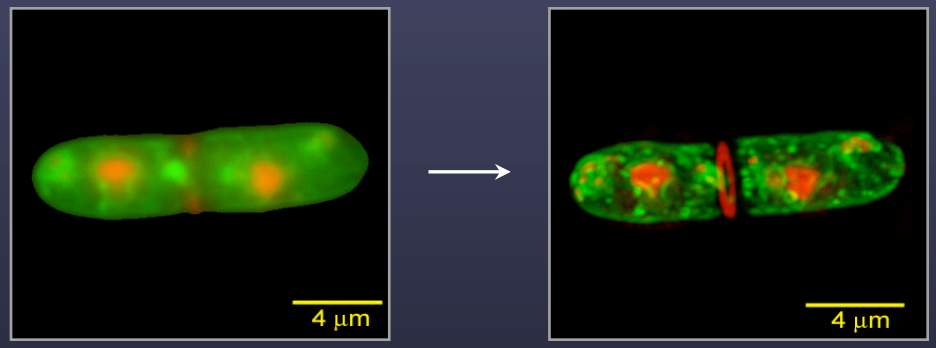

- 4D light field → *tomographic reconstruction* →
  3D volume data

(from Kak & Slaney)

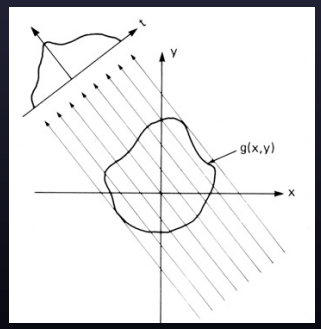

Silkworm mouth

(40x / 1.3NA oil immersion)

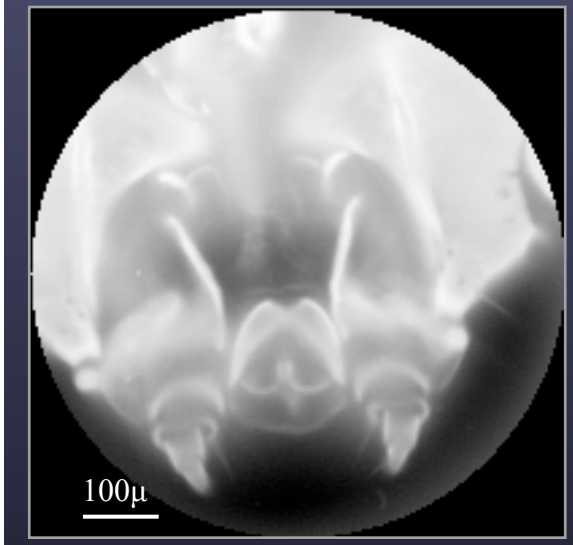

slice of focal stack

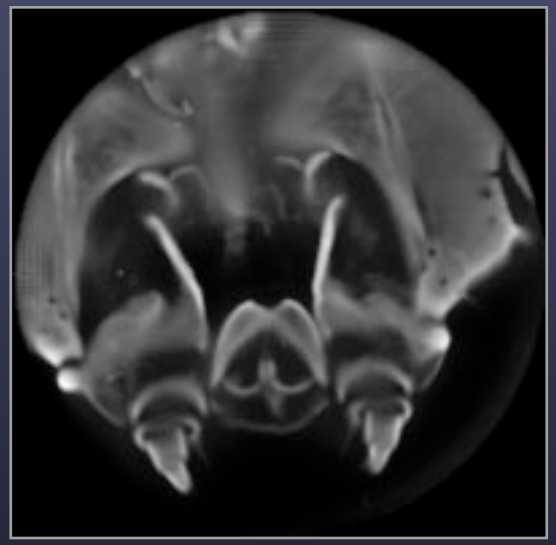

slice of volume

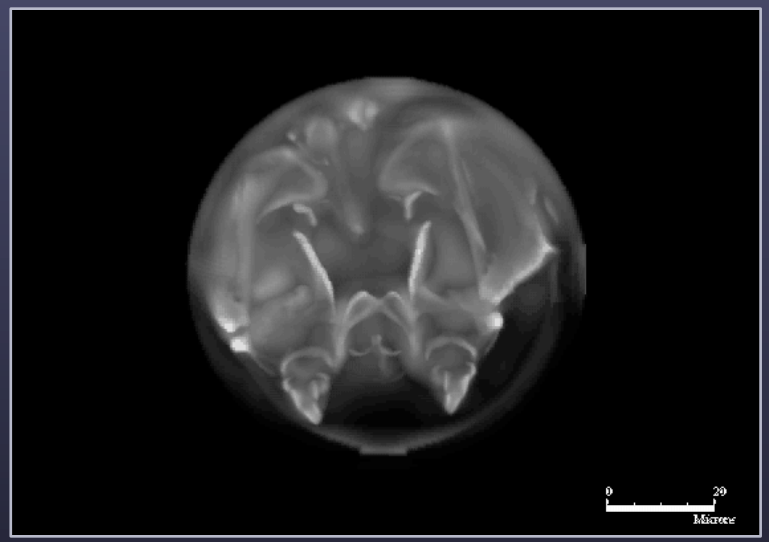

volume rendering

GFP-labeled zebrafish neurons

(40x / 0.8NA water immersion)

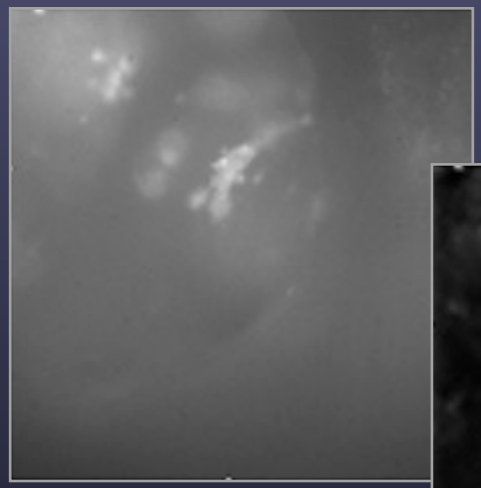

focal stack

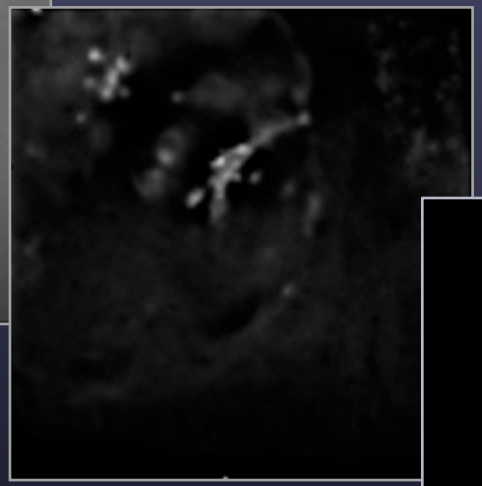

deconvolved

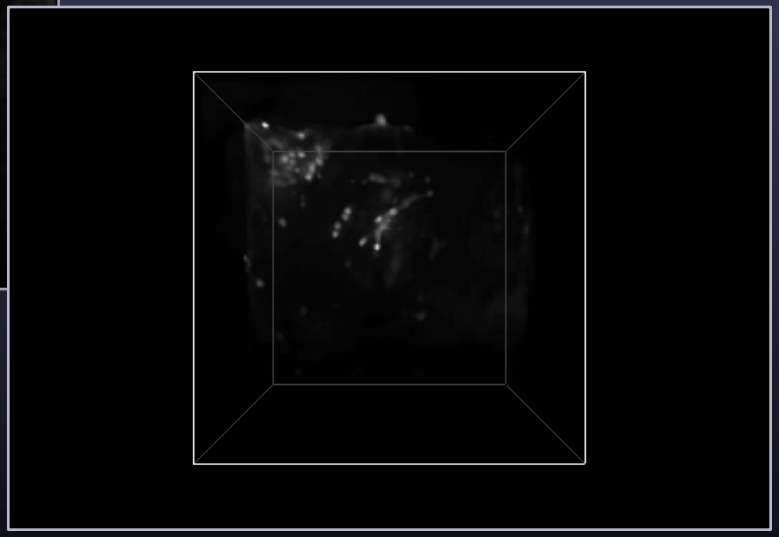

volume rendering

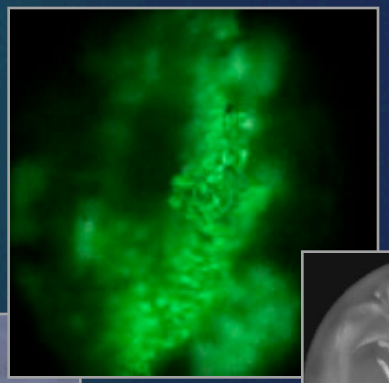

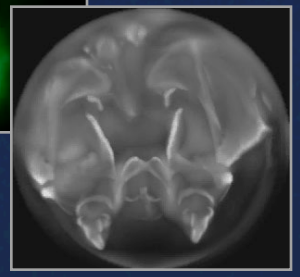

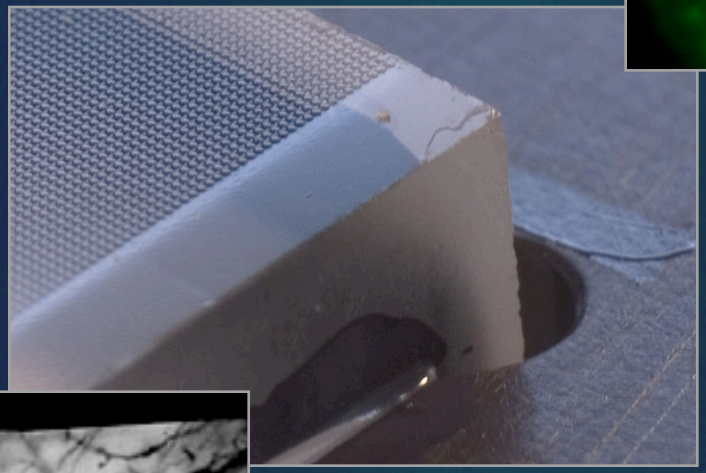

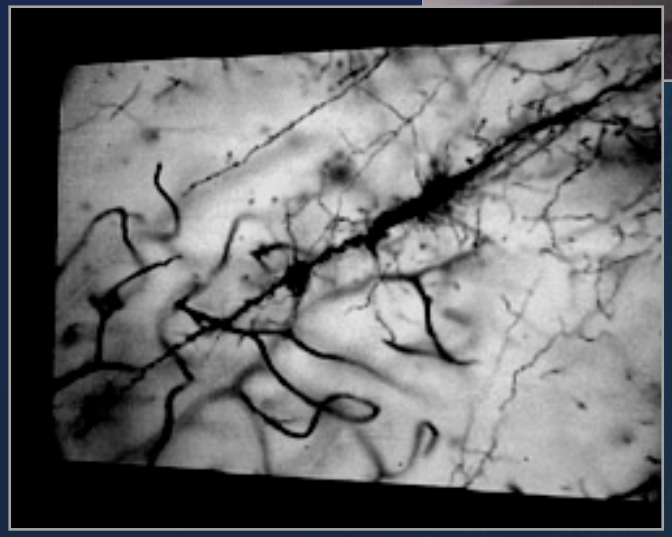

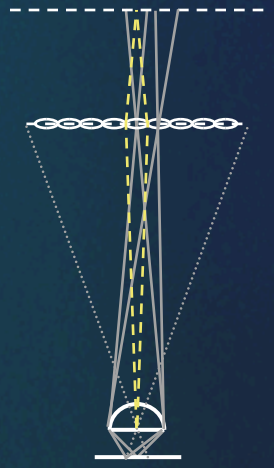

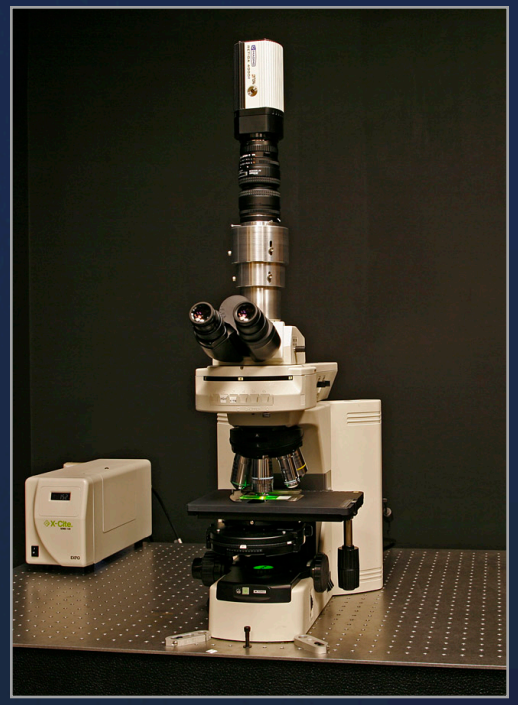

http://graphics.stanford.edu