# Light fields and plenoptic cameras

CS 448A, Winter 2010

Marc Levoy
Computer Science Department
Stanford University

# Outline

- scalar and vector light fields
- light field capture and rendering
  - –parameterization
  - –creation and display
  - –devices for capturing light fields
  - –sampling issues
- 3D reconstruction from light fields
- applications of light fields

# The scalar light field

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

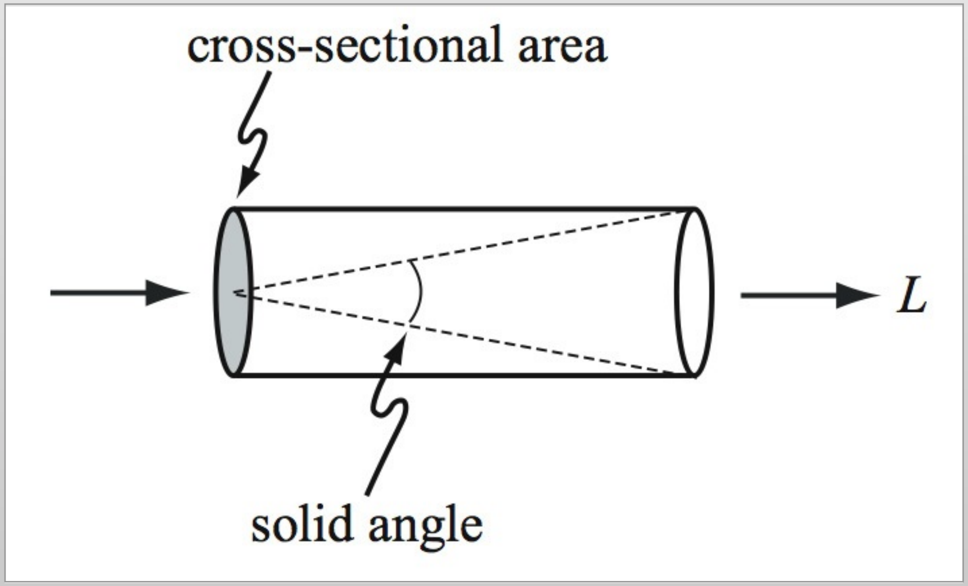


L is radiance in watts / ($m^2$ steradians)

# The vector light field
## [Gershun 1936]

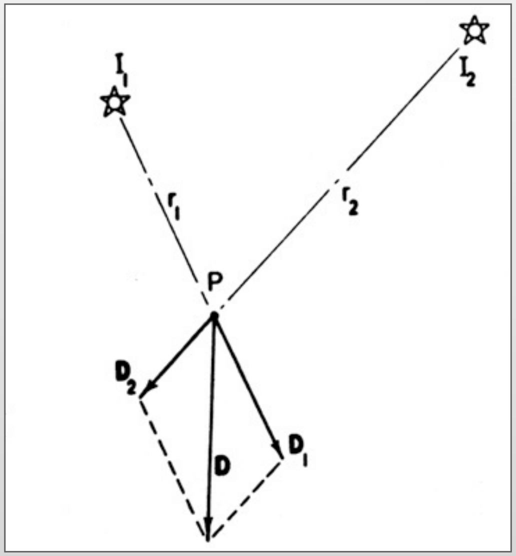


adding two light vectors

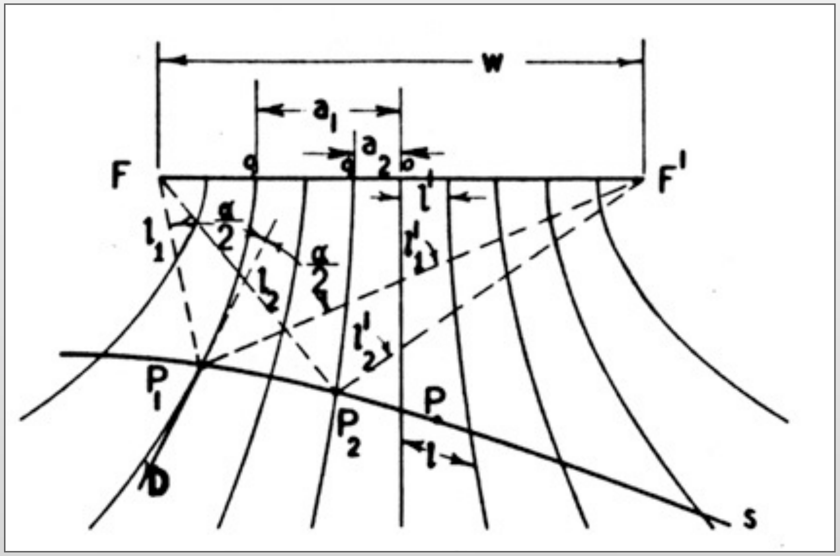


the vector light field
produced by a luminous strip

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

# Visualizing the vector light field

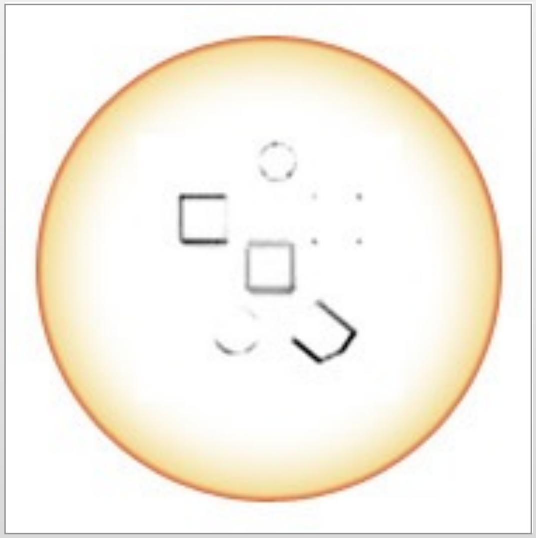

flatland scene with
partially opaque blockers
under uniform illumination

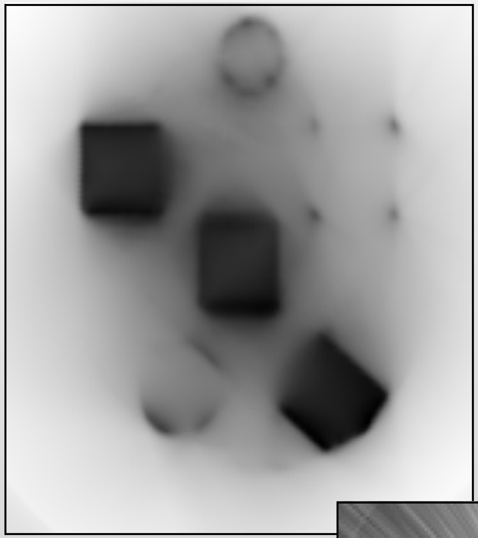

scalar irradiance
at each point

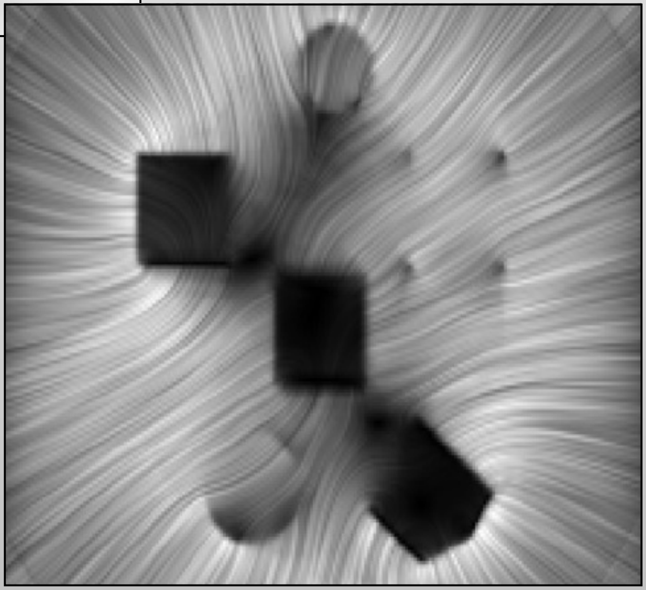

vector directions,visualized using
line integral convolution (LIC) [Cabral 1993]

# Dimensionality of the scalar light field

- for general scenes

  ⇒ 5D function

  "plenoptic function"
  [Adelson 1991]

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

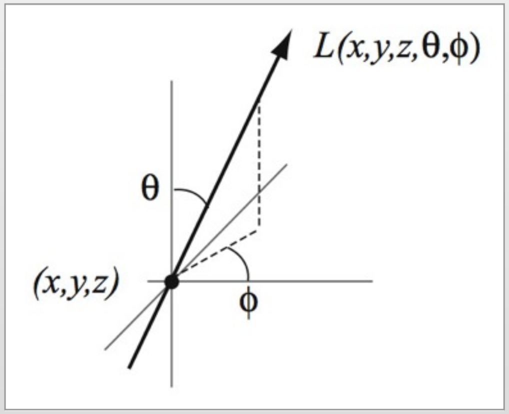


- in free space

  ⇒ 4D function

  " the (scalar) light field"
  [Moon 1981]

  $L(\,?\,)$

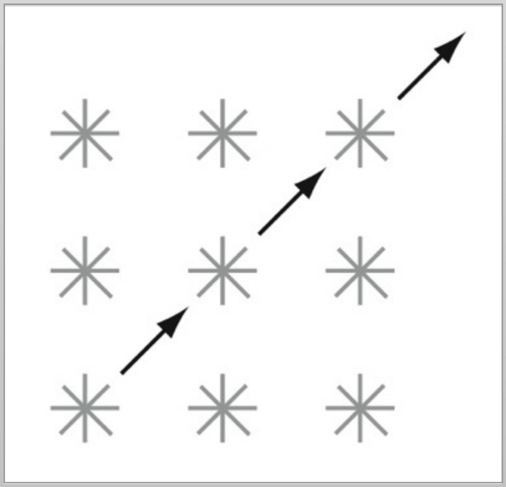

# The free-space assumption

- the 3D space around a compact object
- the 3D space inside an uncluttered environment

# Light field rendering

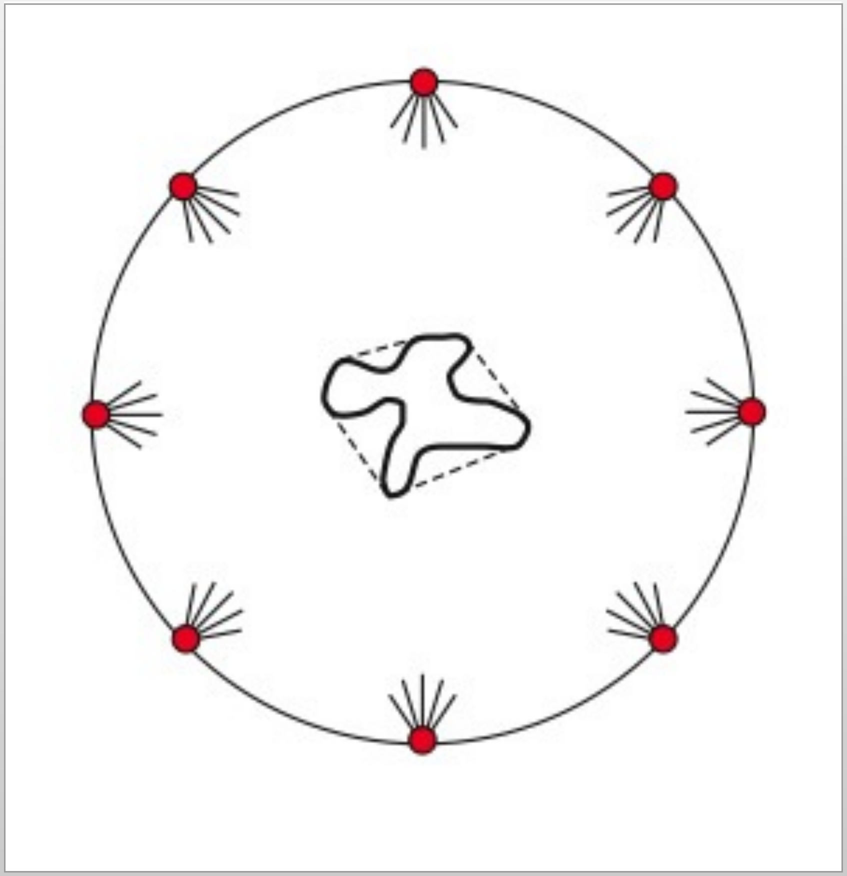

flipbook animation
(QuickTime VR)
[Chen 1995]

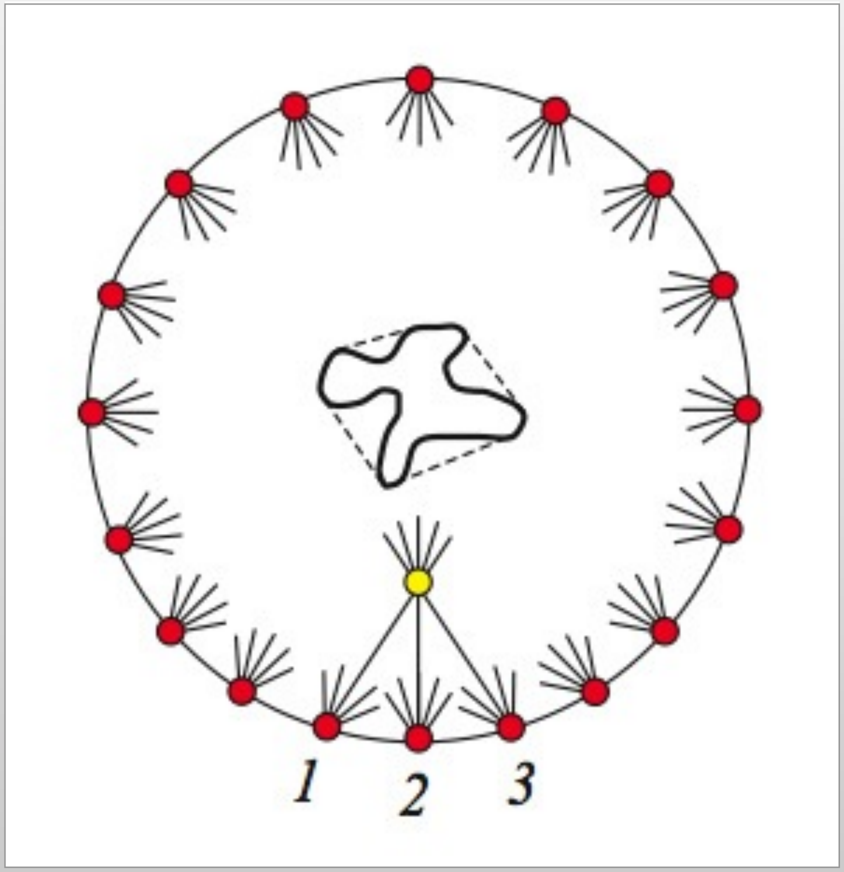


rebinning the rays
to create new views
[Levoy 1996]

# Outline

- scalar and vector light fields
- ☞ light field capture and rendering
  - parameterization
  - creation and display
  - devices for capturing light fields
  - sampling issues
- 3D reconstruction from light fields
- applications of light fields

# Some candidate parameterizations

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

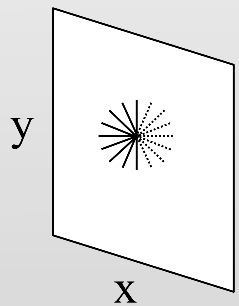


L ( x, y, θ, ϕ )

or

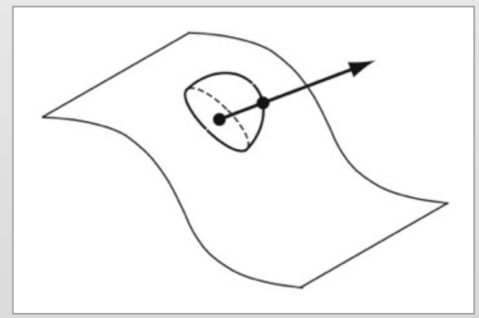

L ( u, v, θ, ϕ )

- convenient for representing the light field at a surface
  [Wood 2000]

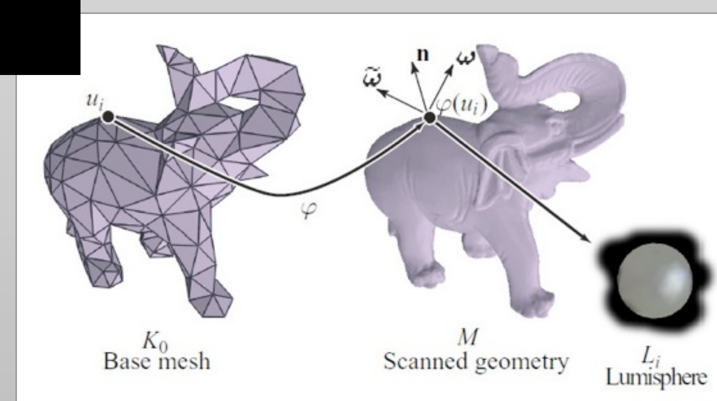

# Chords of a sphere

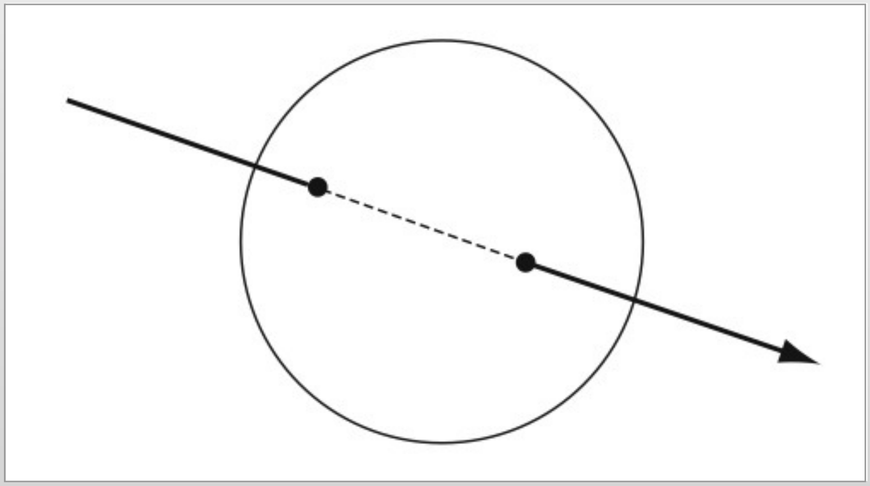

$L(\theta_1, \phi_1, \theta_2, \phi_2)$

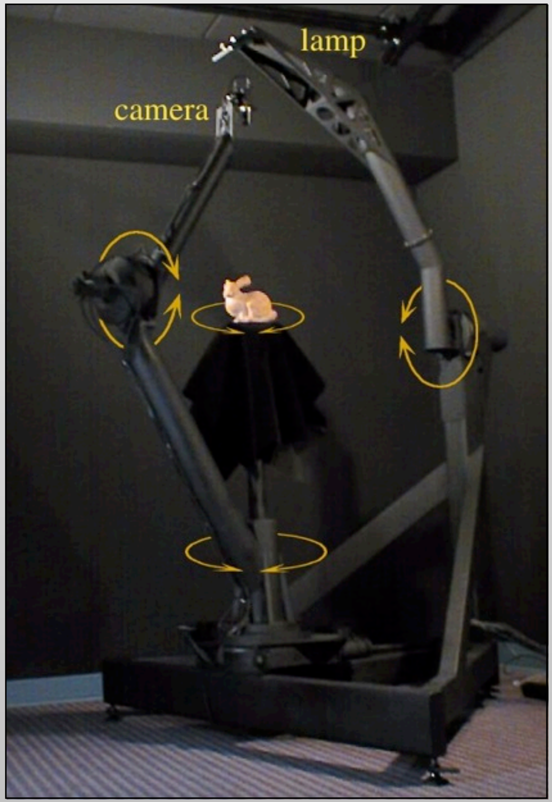


- if points on sphere are chosen at random, sampling of light field will be uniform
- useful for spherical gantries

# Two planes ("light slab")

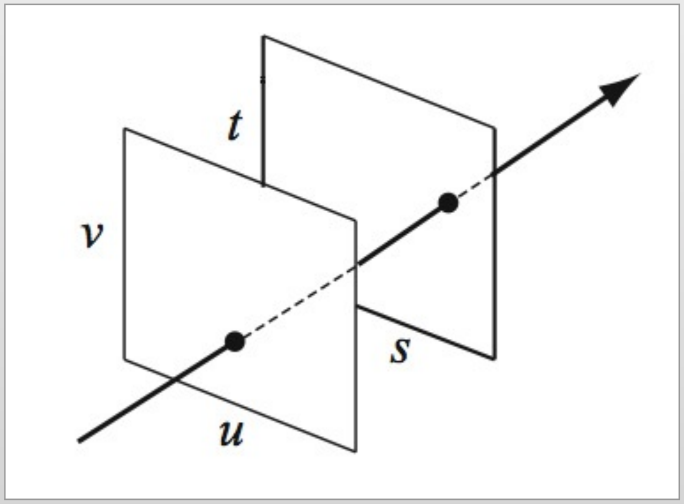


L ( u, v, s, t )

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

# Alternative parameterization for the <u>5D</u> plenoptic function

- Two-plane ray field

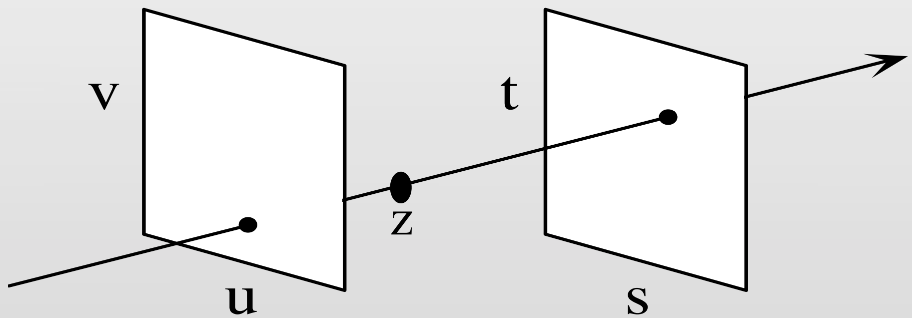


L ( u, v, s, t, z )

- allows multiple colors, in sequence, along one line
- alternative to L ( x, y, z, θ, φ )
- inspired by Salesin's ZZ-buffer [1990]

# A light field is an array of images

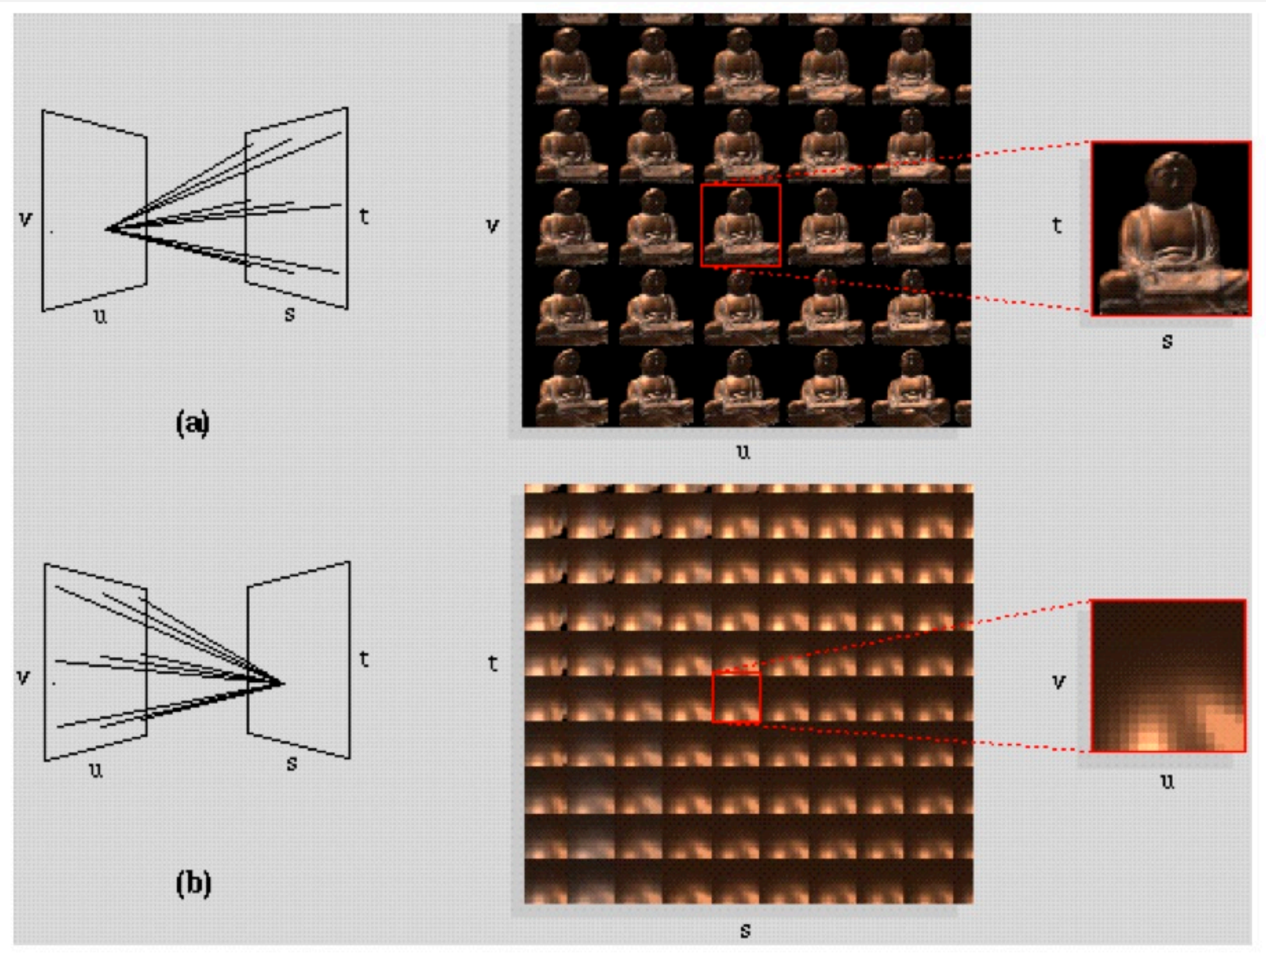


- an array of perspective views
- or an array of reflectance maps

...depending on where the object is relative to the two defining planes

# Creating a (synthetic) light field using the 2-plane parameterization

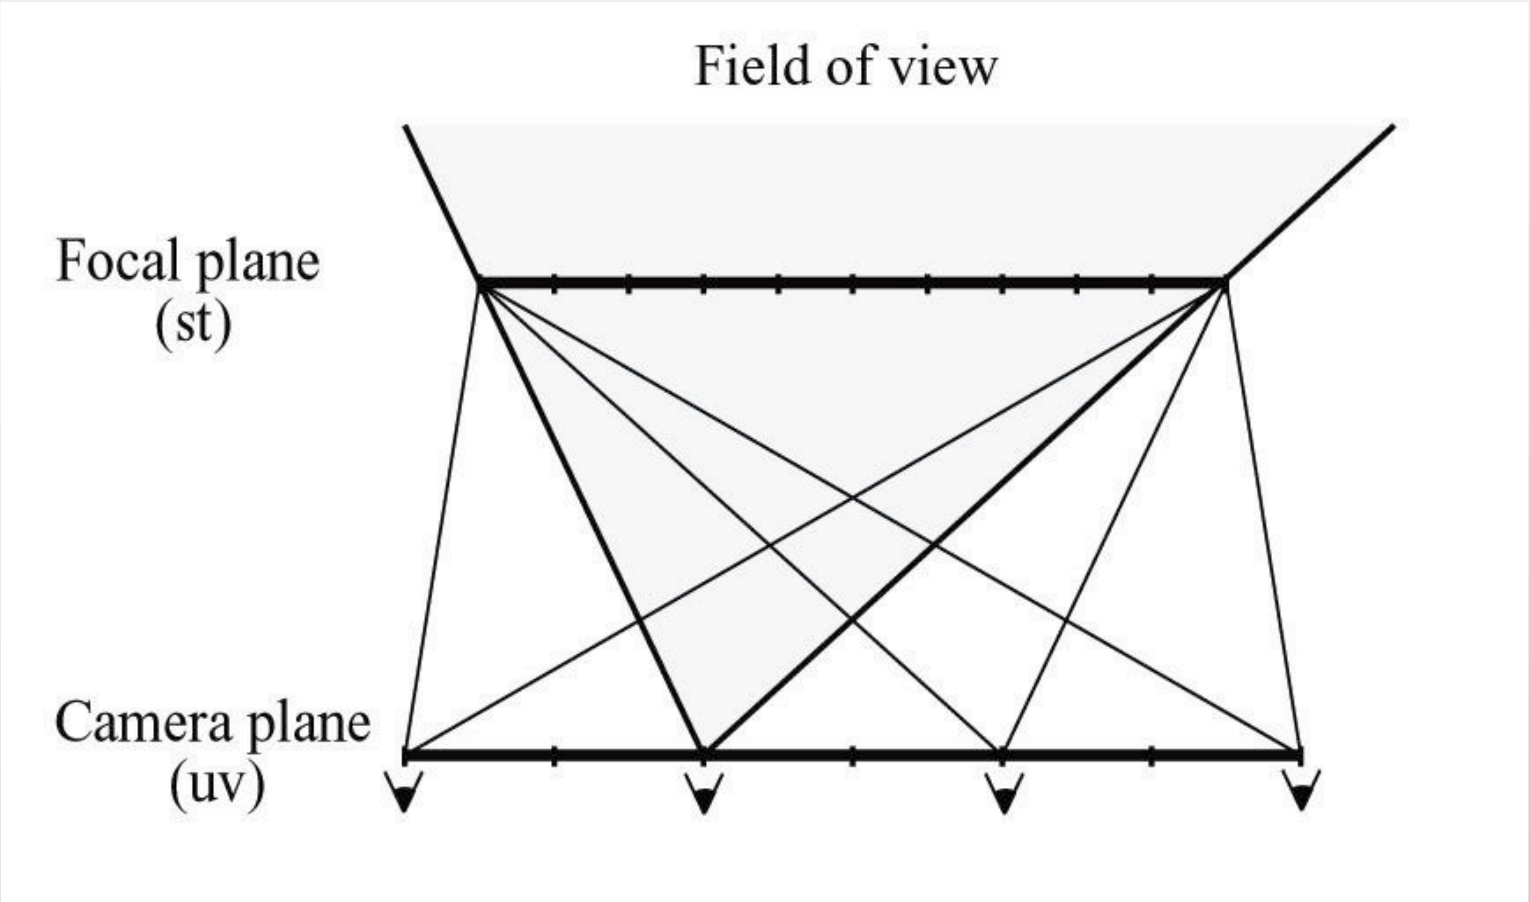


- off-axis (sheared) perspective views

# Displaying a light field

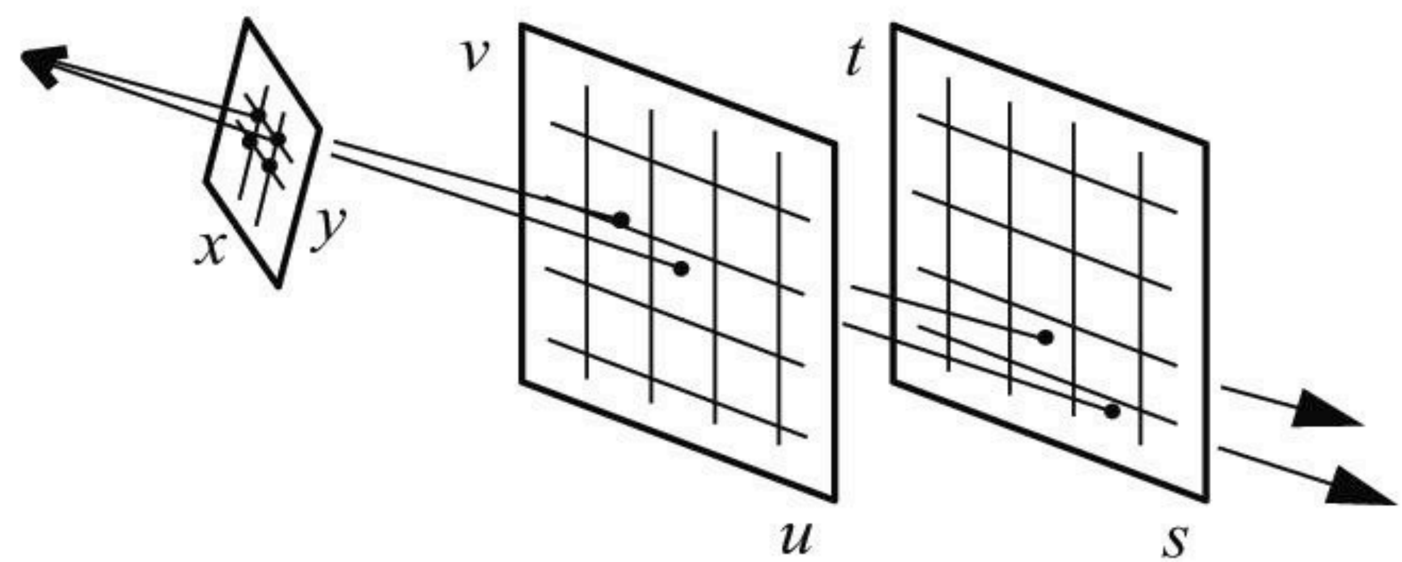


```
foreach x,y
     compute u,v,s,t
     I(x,y) = L(u,v,s,t)
```

(lightfield.rm)

# Pre-filtering during creation

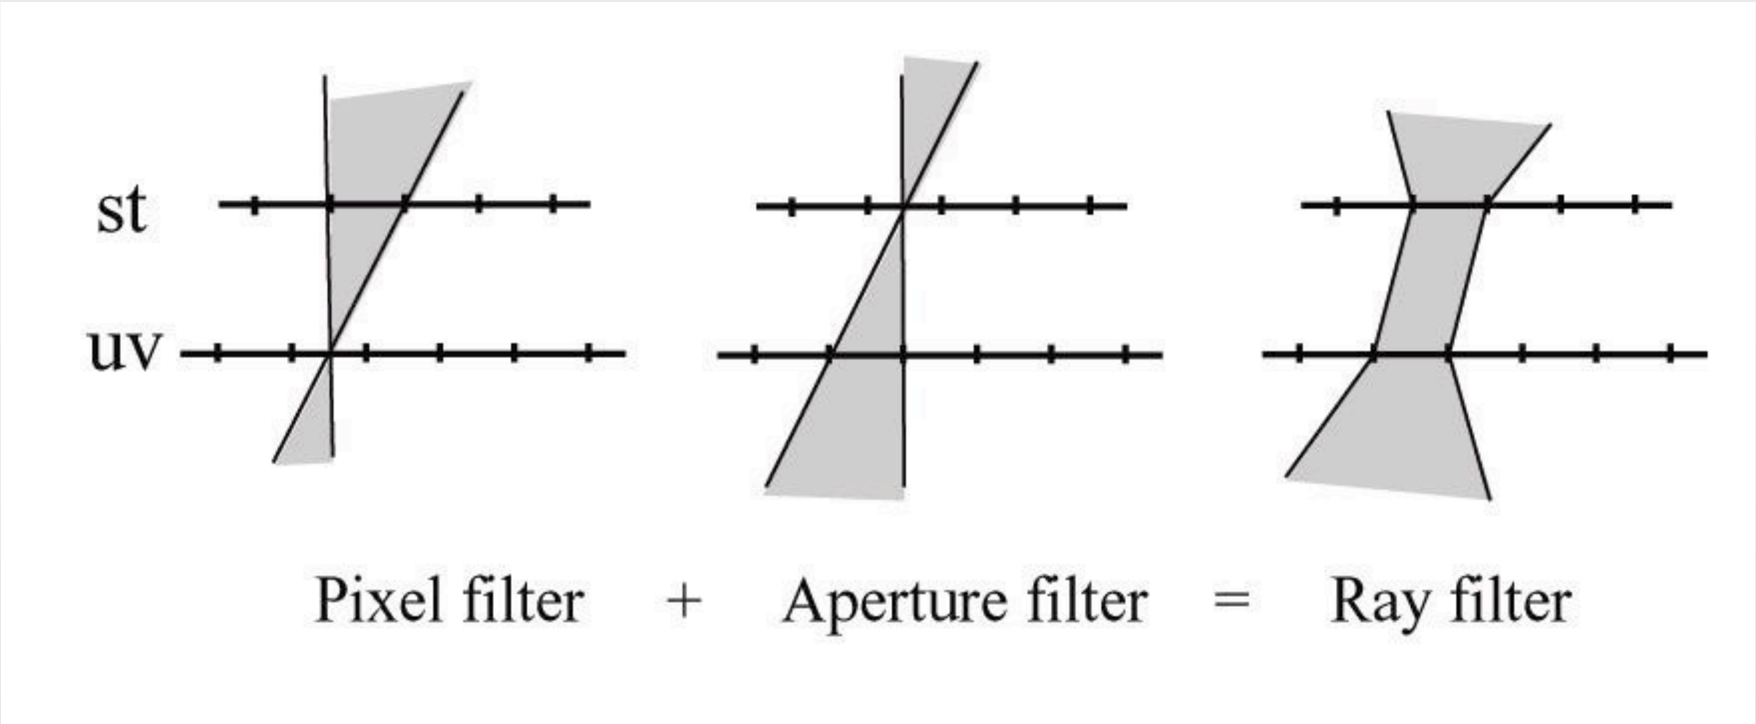

# The light field captured by a (single) (ordinary) camera

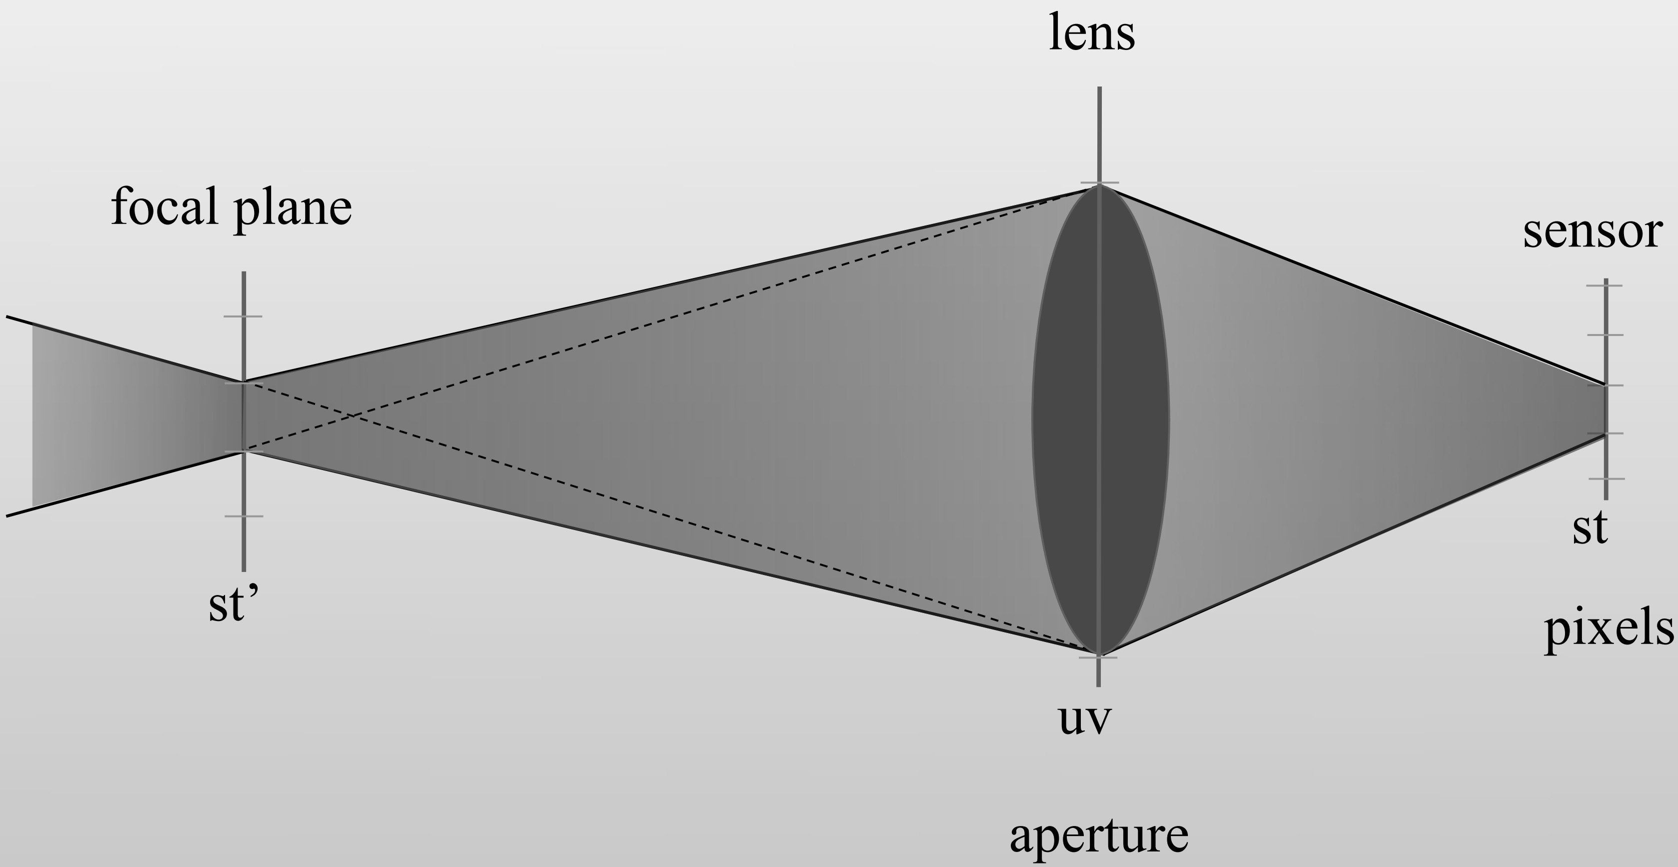

# Interpolation for display

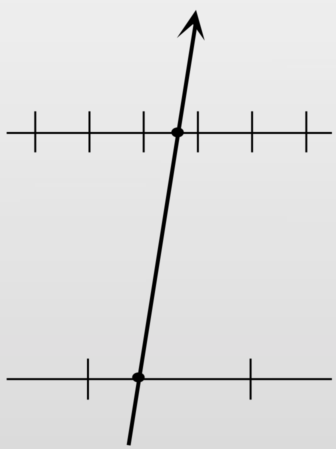

point sample

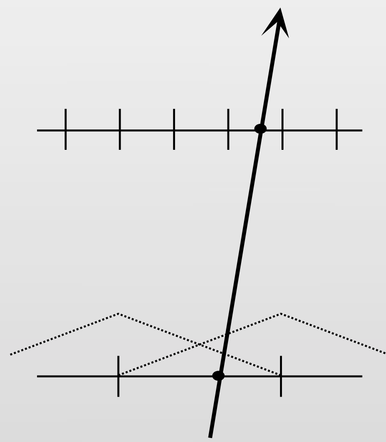

uv bilerp

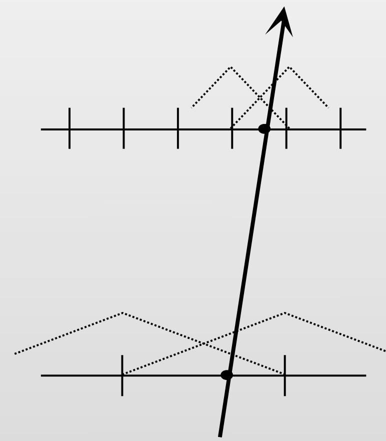

uv and st bilerp

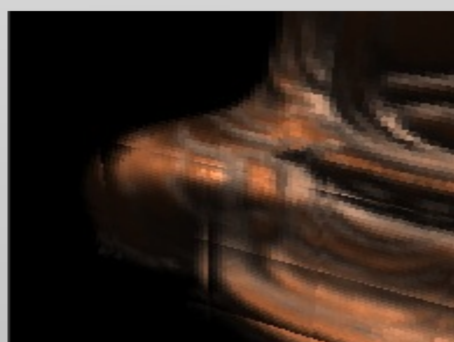

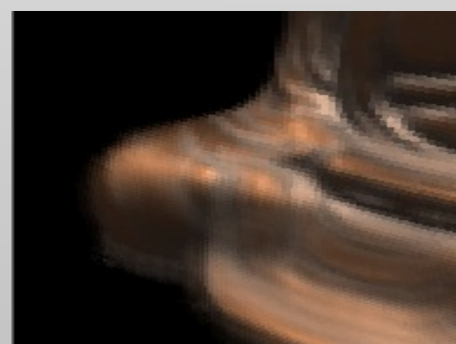

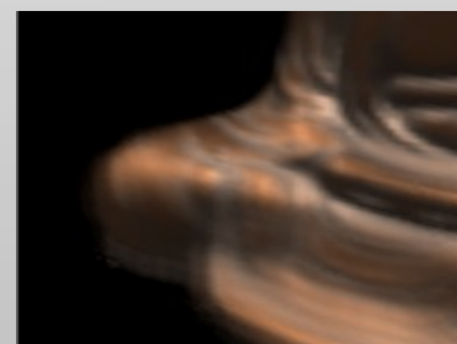

(alldragon.lif)

# Our planned light field of the Medici Chapel

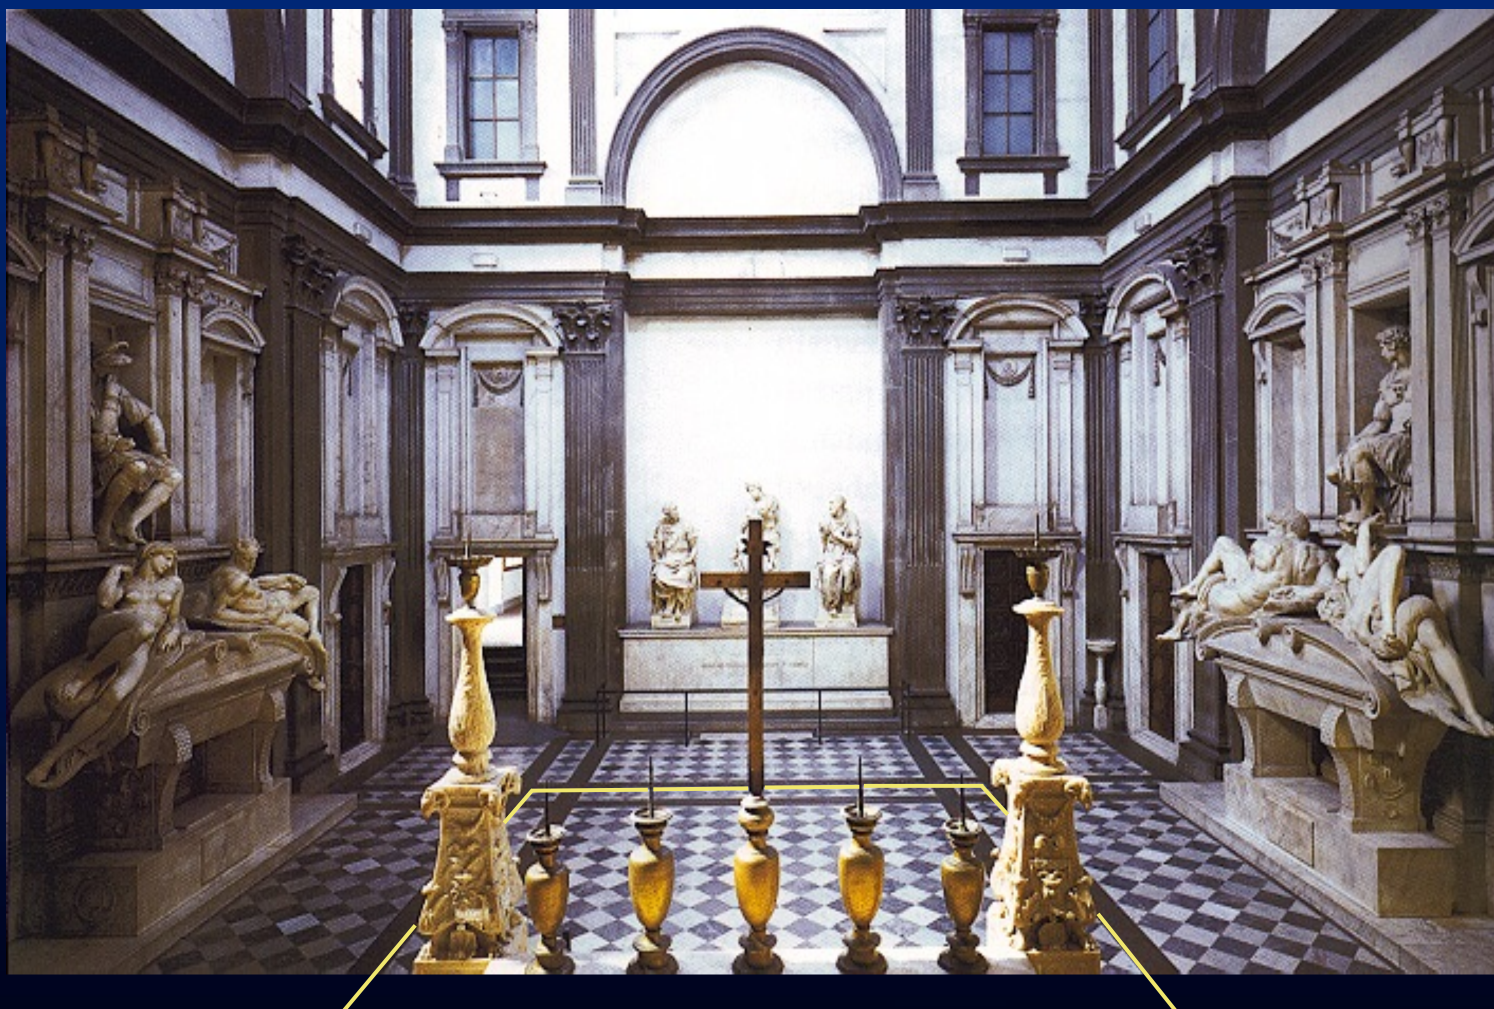

# What got in the way of this plan

# An optically complex statue

Night (Medici Chapel)

# Capturing the light field

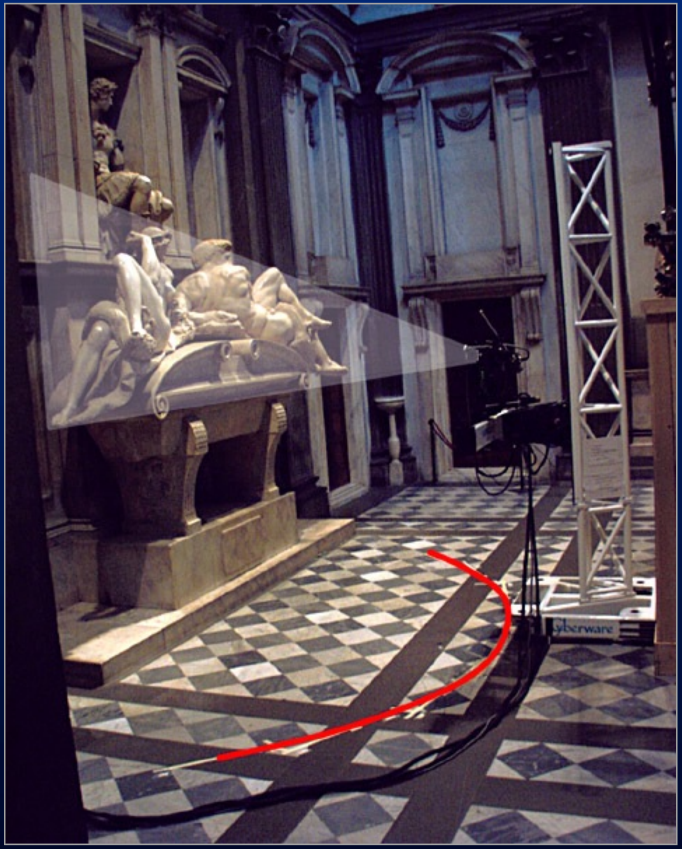


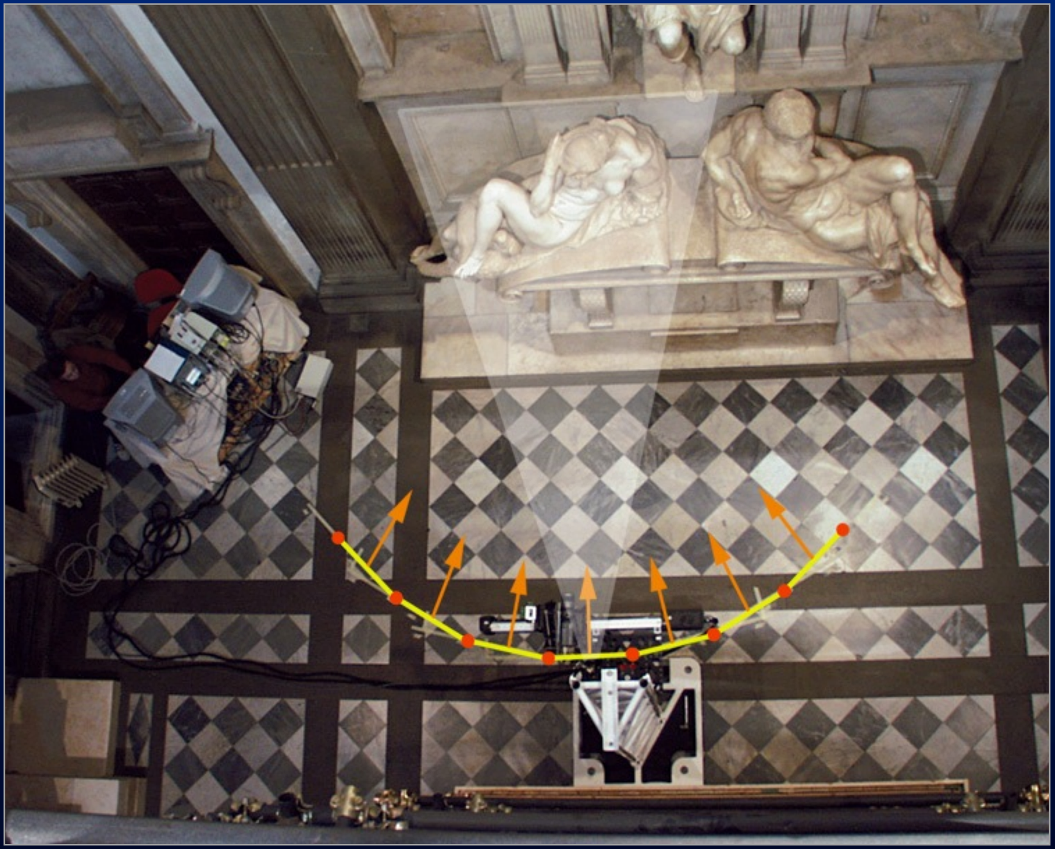

7 light slabs,
each 70cm x 70cm

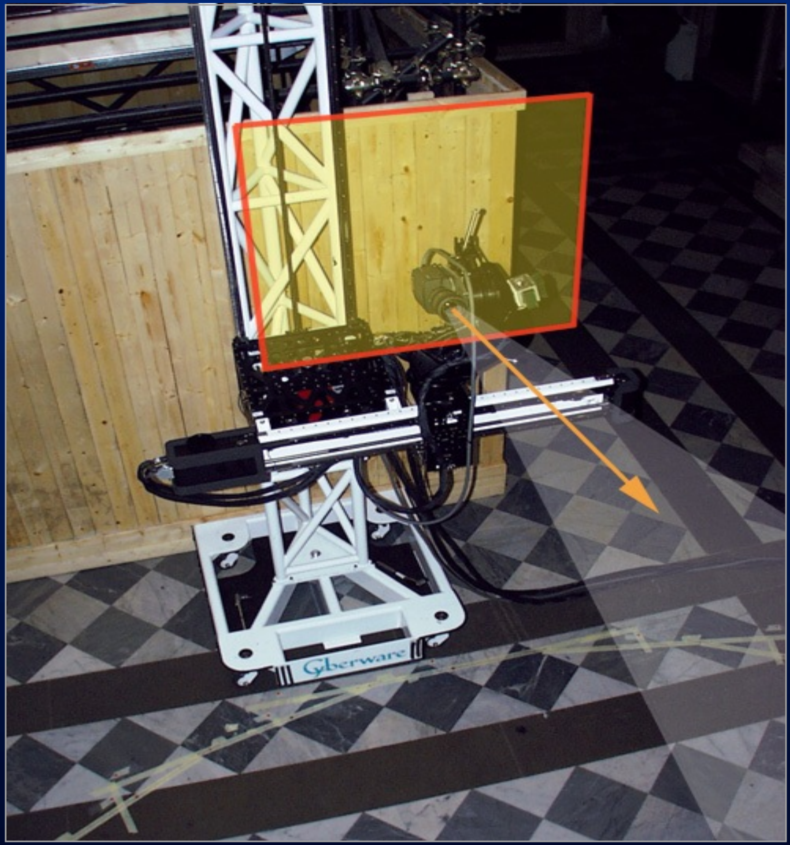


each slab contained 56 x 56
images spaced 12.5mm apart

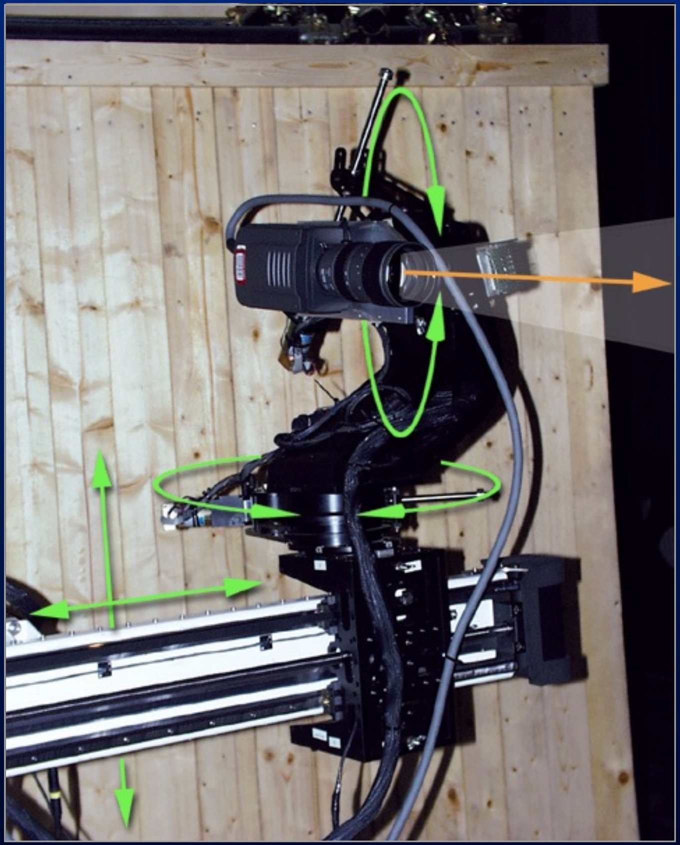

the camera was always aimed
at the center of the statue

# Statistics about the light field

- 392 x 56 images
- 1300 x 1000 pixels each
- 96 gigabytes (uncompressed)
- 35 hours of shooting  (over 4 nights)
- also acquired a 0.29 mm 3D model of statue
- data still hasn’t been calibrated and aligned!

Single image from the light field

# One row of one light field slab

# Other devices for capturing light fields

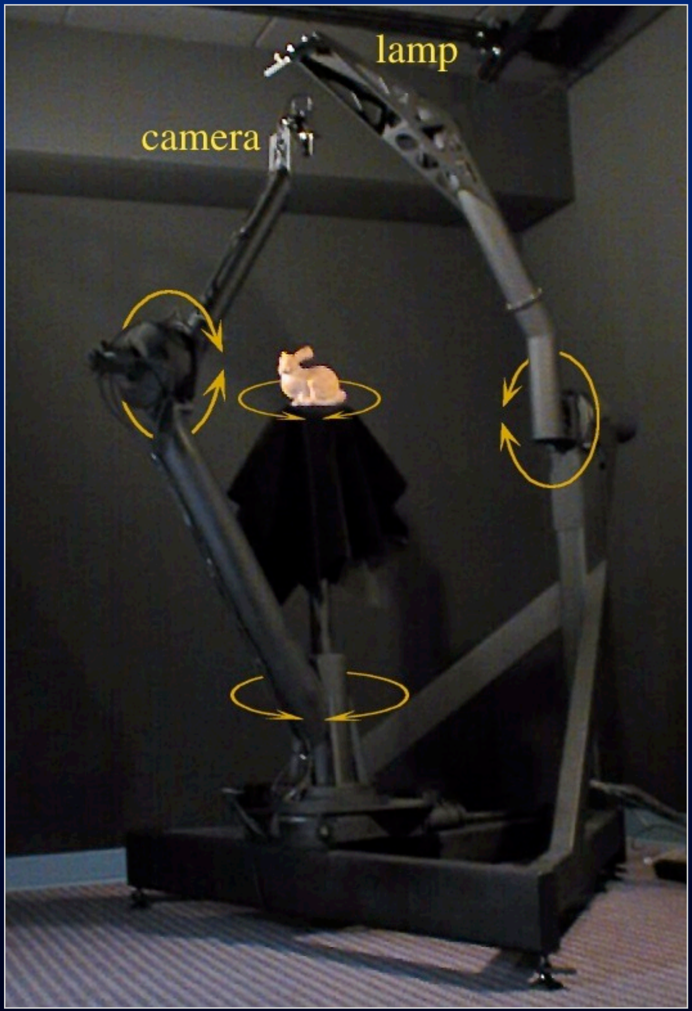


Stanford Spherical Gantry

- Stanford spherical gantry
  [Levoy 2002]
- MIT camera array
  [Yang 2002]
- CMU camera array
  [Zhang 2004]
- MSR/China concentric mosaics
  [Shum 2000]
- Stanford camera array
  [Wilburn 2005]
- Ren Ng's plenoptic camera
  [Ng 2005]

# Handheld plenoptic camera
## [Ng 2005]

- array of microlenses <u>behind</u> the main lens
  - –requires modifying the camera

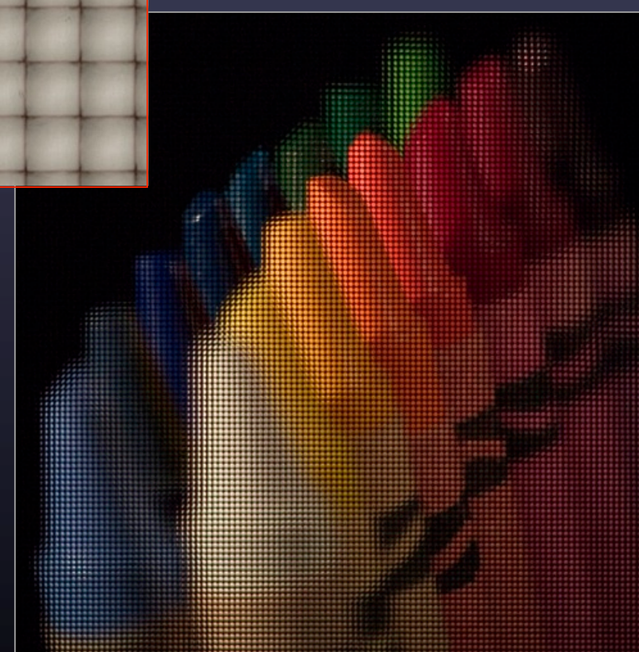

# Adobe light field camera
## [Georgeiv 2006]

- array of lenslets <u>outside</u> the main lens
  - –each lenslet must be well-corrected

# Light field microscope
[Levoy 2006]

- array of microlenses behind a microscope objective
  - allows oblique views, refocusing, 3D reconstruction

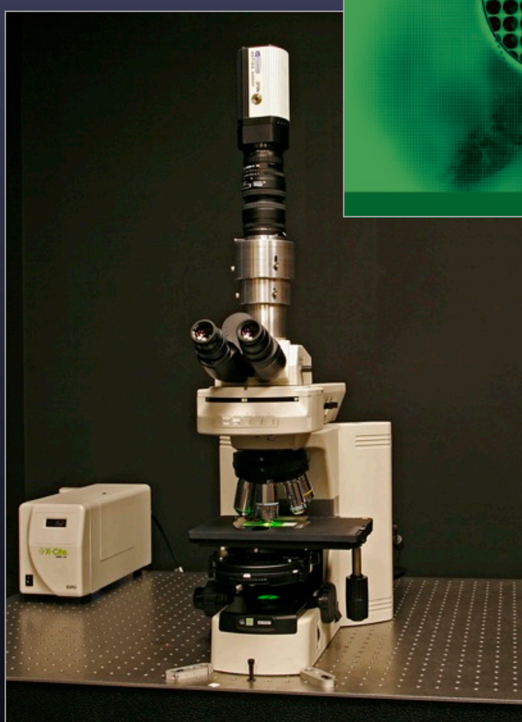

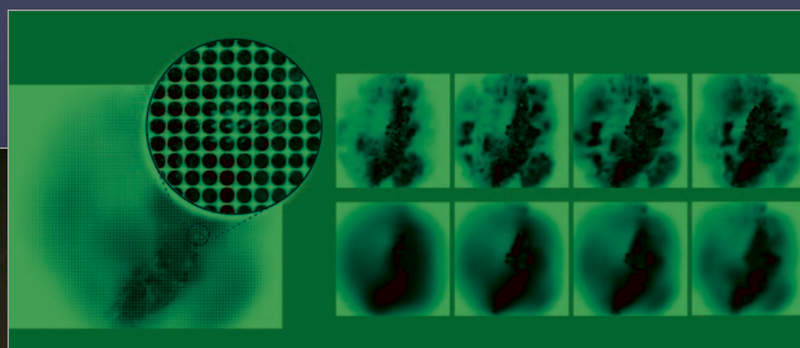

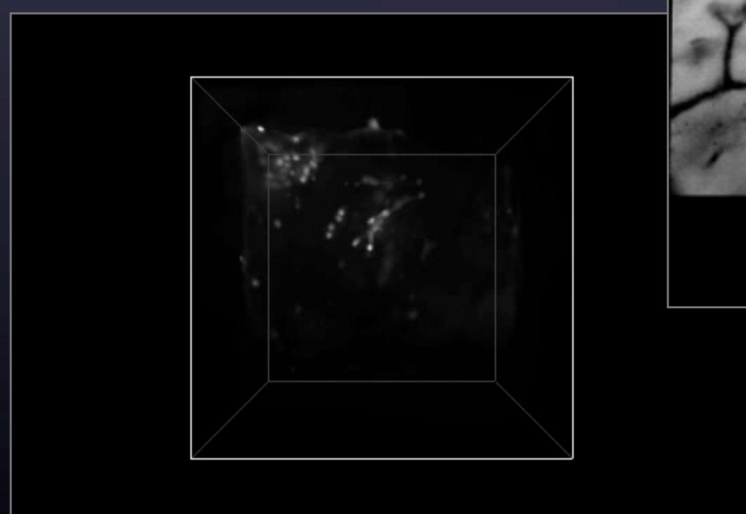

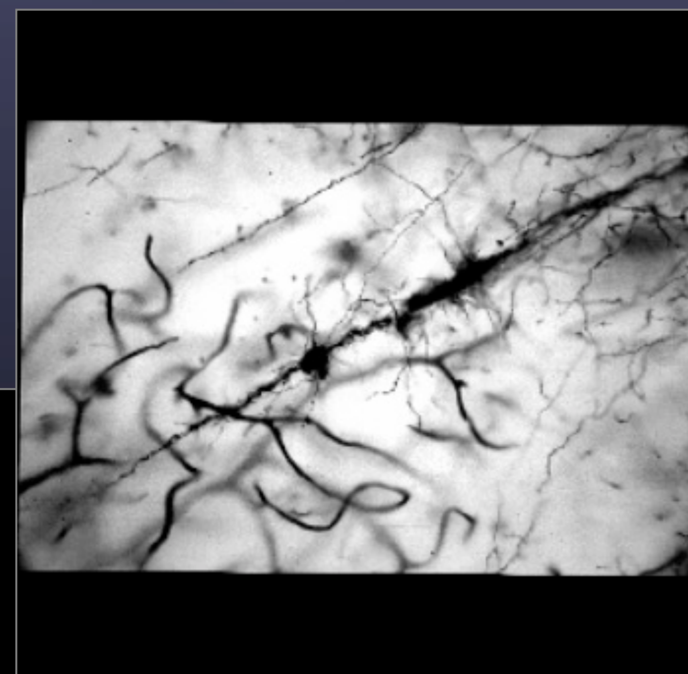

# Capturing unstructured light fields using a handheld video camera

- video camera and calibration target
  [Gortler 1996]

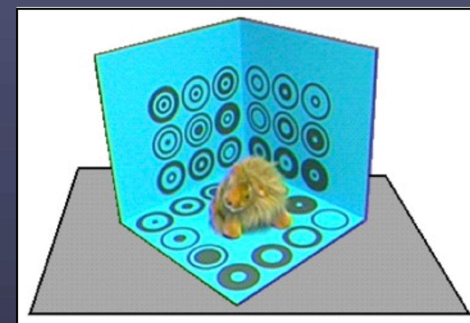

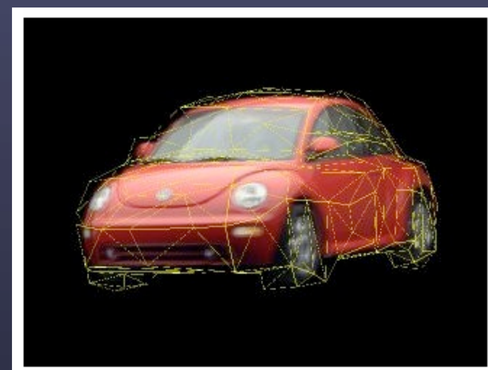

- markerless capture
  [Buehler 2001]

- interactive capture
  (Abe Davis)

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

calibration point

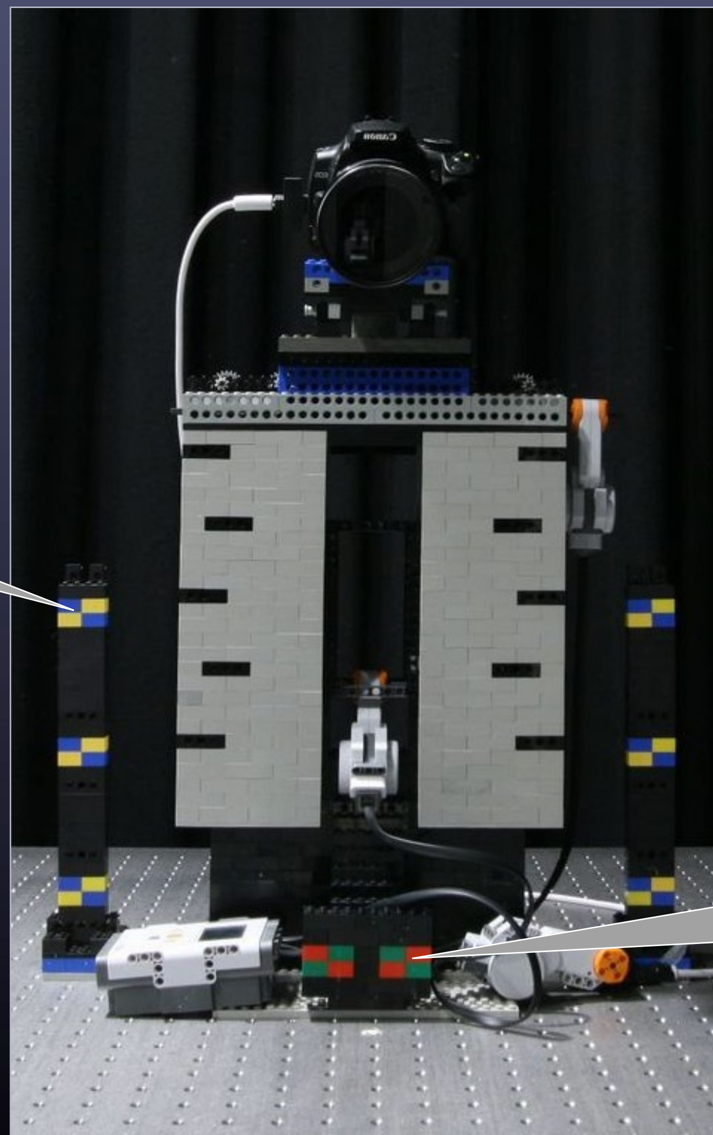

plane + parallax [Vaish 2004]

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


Try it yourself at http://lightfield.stanford.edu/

# The Lego gantry captures a light field of itself

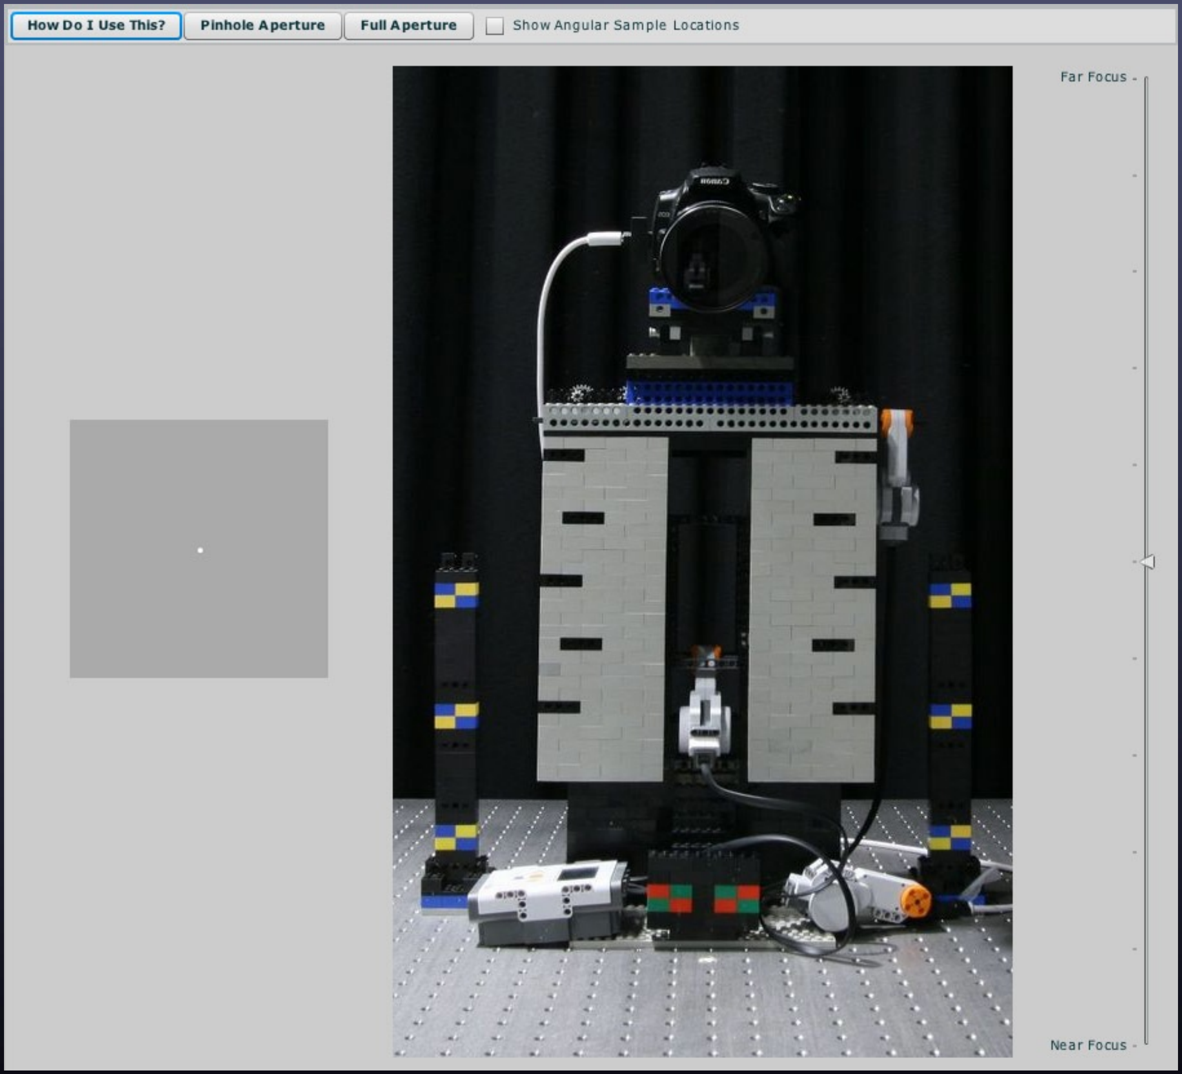

# The BRDF kaleidoscope

[Han 2003]

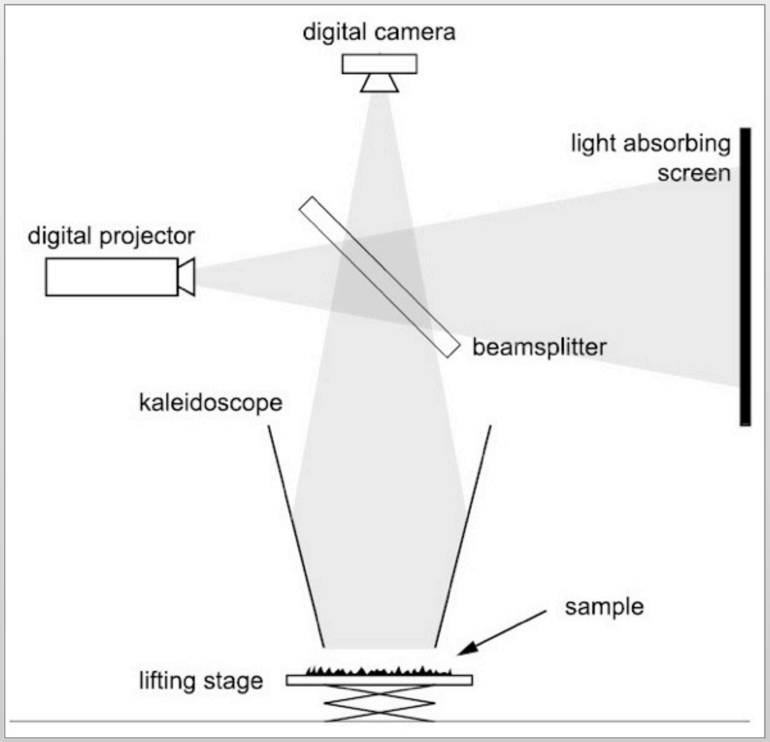


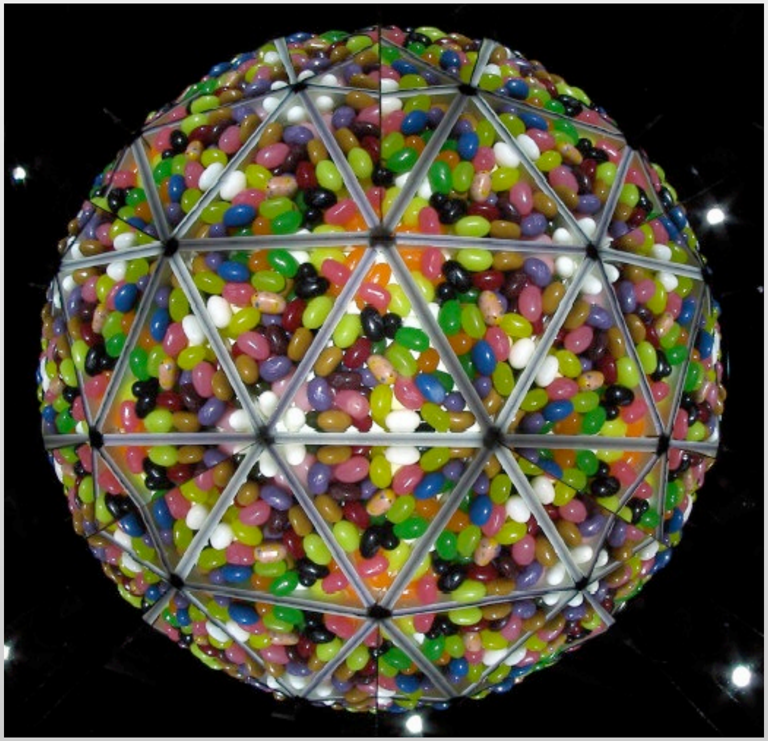

# Photographing through mirrors


(Loreco)

(Nayar)

(Nayar)

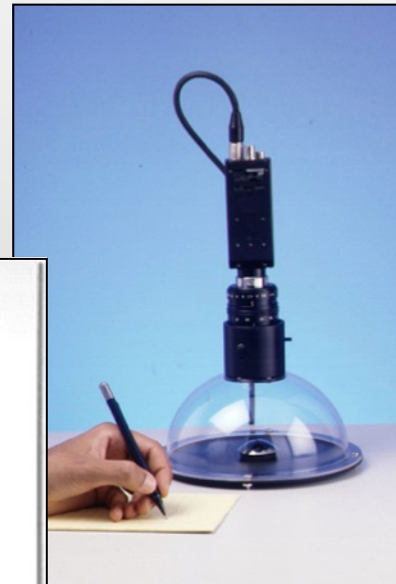

(http://www.math.toronto.edu/~drorbn/Gallery/Symmetry/Tilings/S333/Kaleidoscope-2.html)

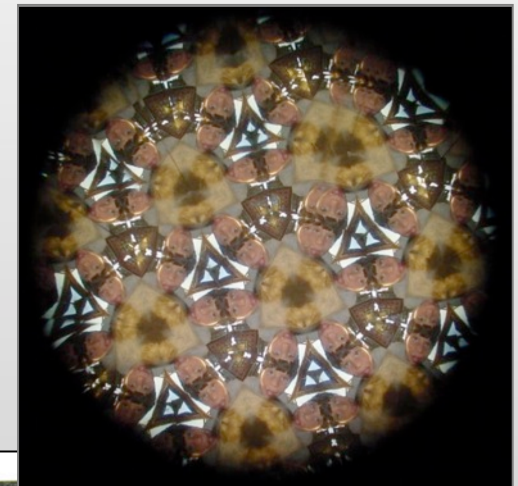

# Outline

- scalar and vector light fields
- light field capture and rendering
  - –parameterization
  - –creation and display
  - –devices for capturing light fields
  - –sampling issues
- 3D reconstruction from light fields
- applications of light fields

# The coverage of a light field

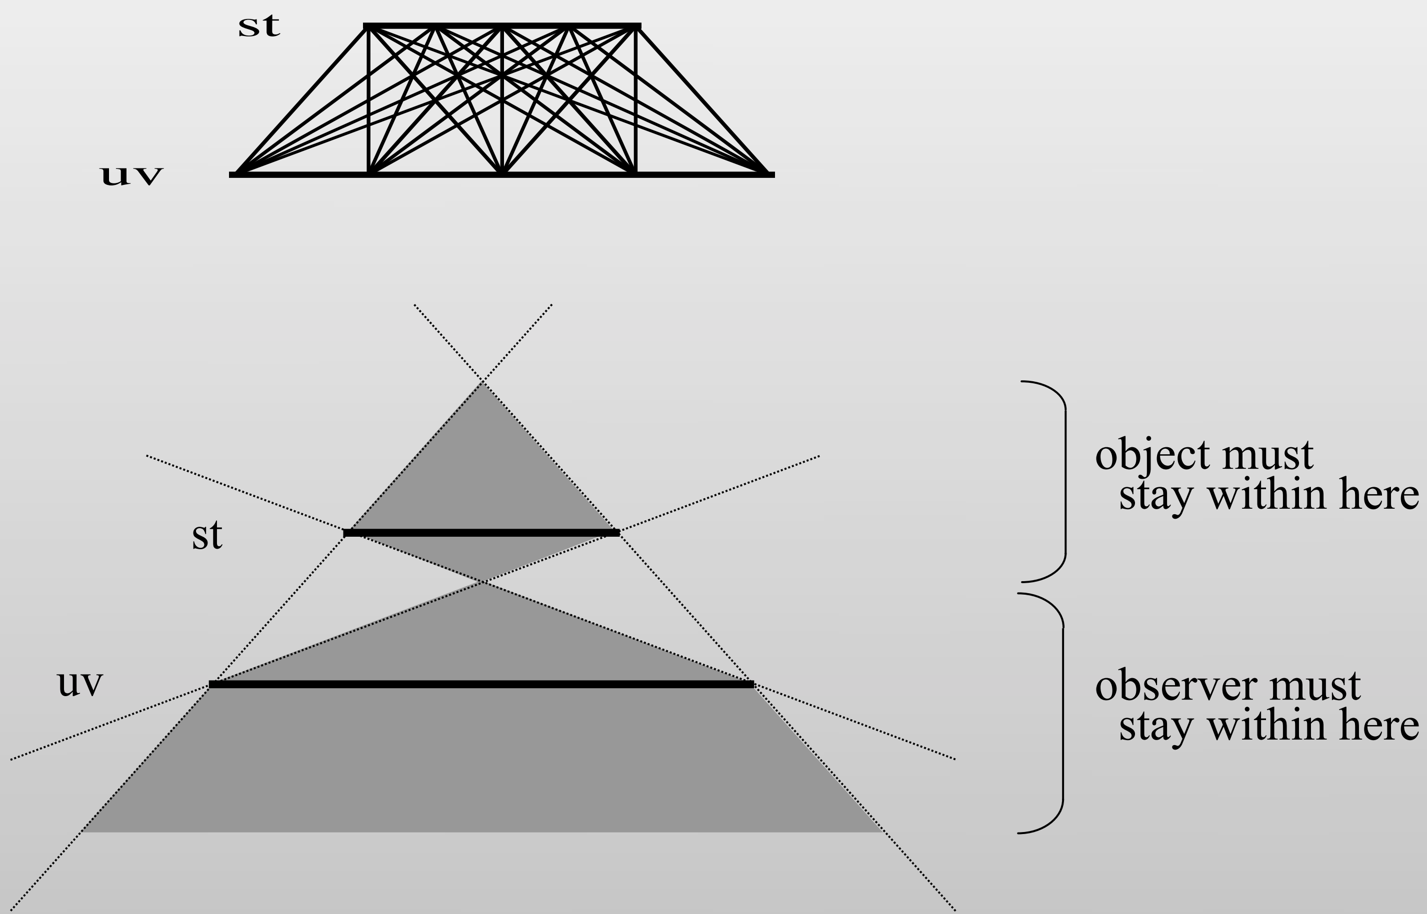

# An omni-directional light field

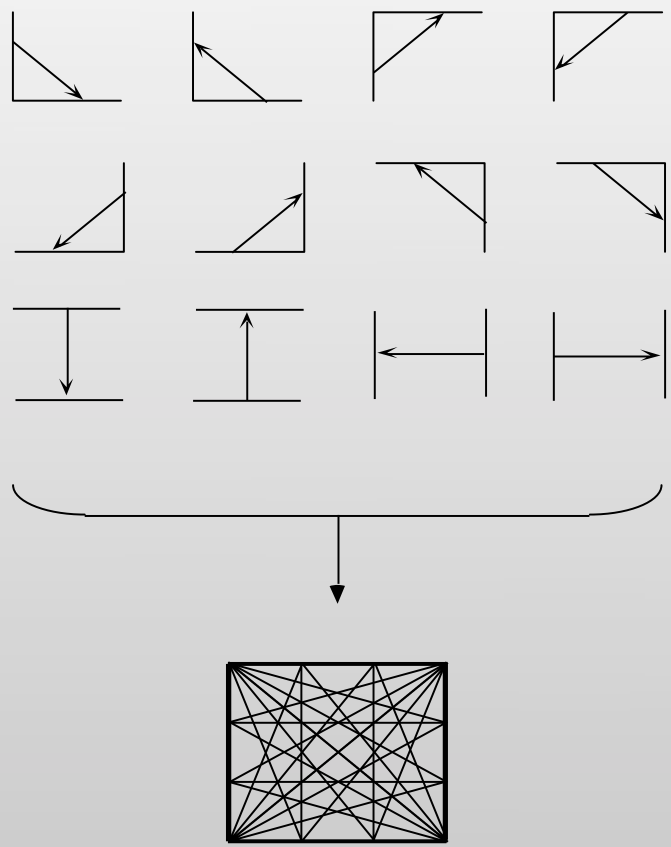

- object can fill square
- observer can stand anywhere outside square

# Line space interpretation

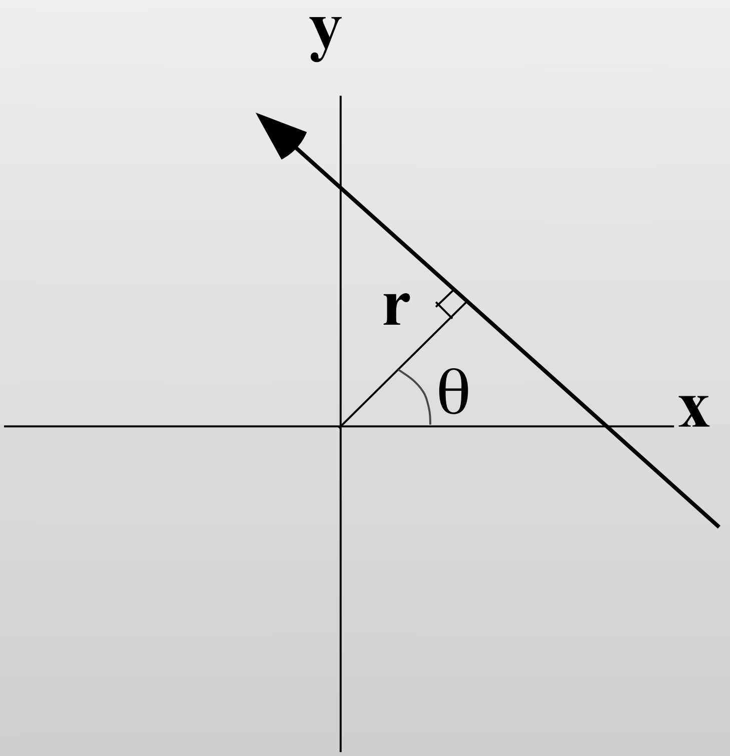


**point space**

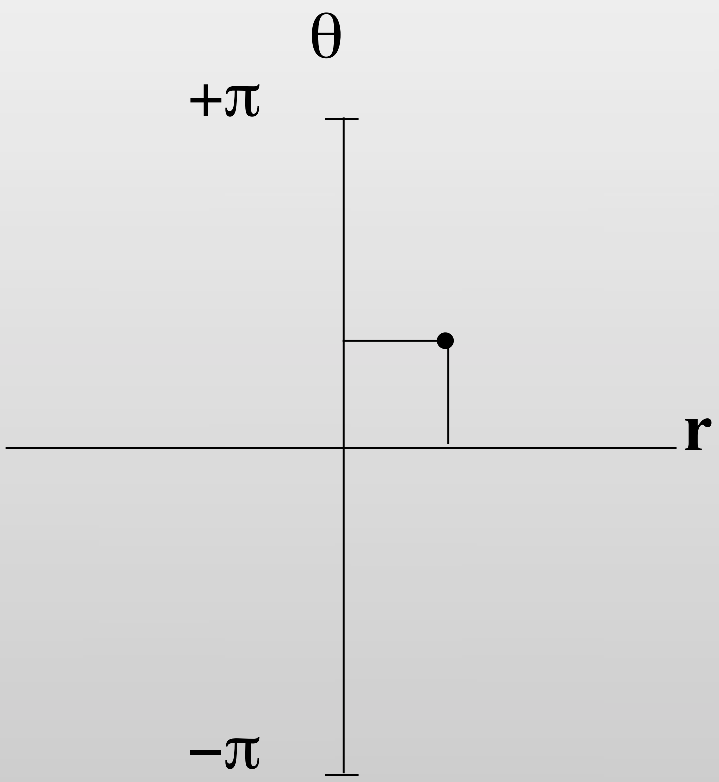


**line space**

# Using line space to visualize ray coverage

**point space** | **line space**

**one slab**

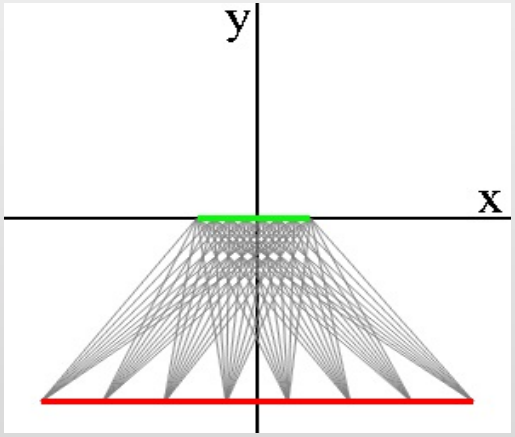


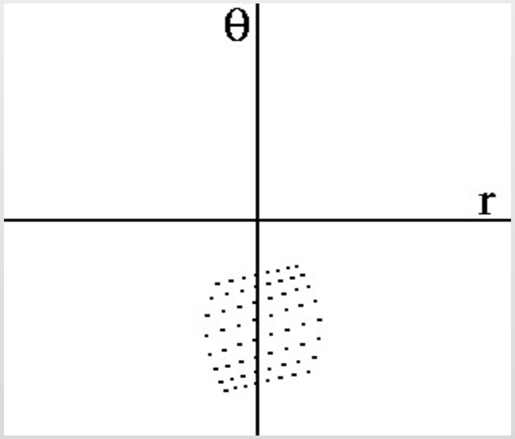


(alldragon.lif)

**four slabs**

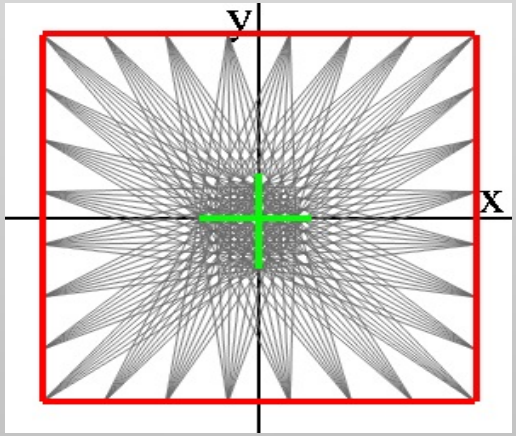


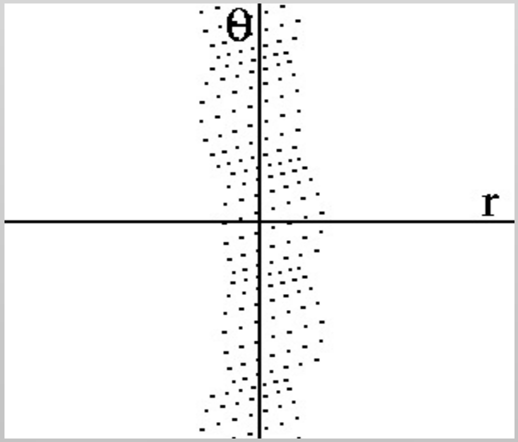


(lion.lif)

# Using line space to visualize sampling uniformity

**point space** | **line space**

**corner**

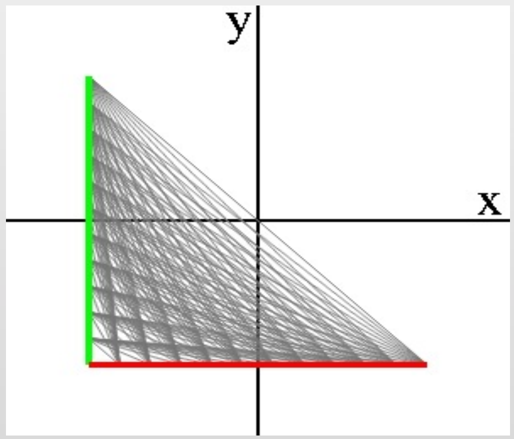


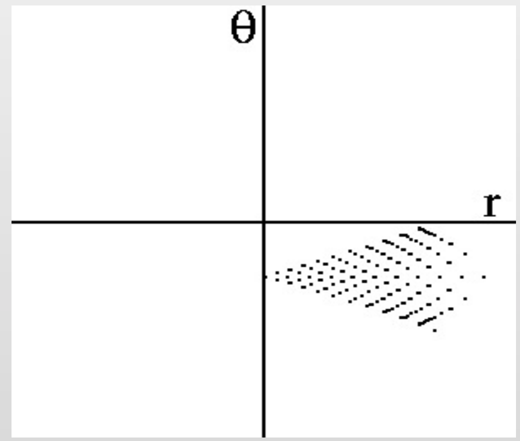


**orthographic**

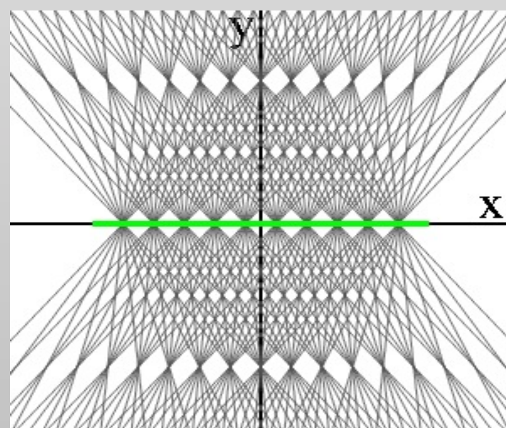


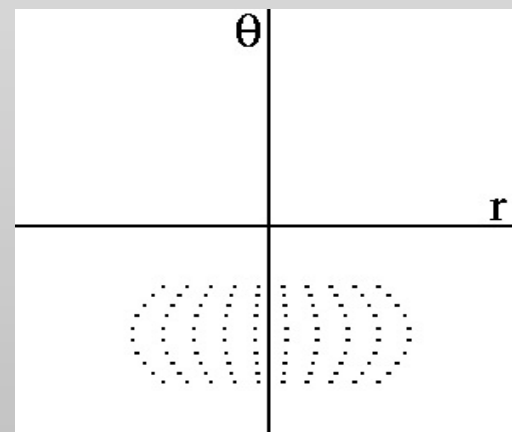

# Disparity artifacts in light fields

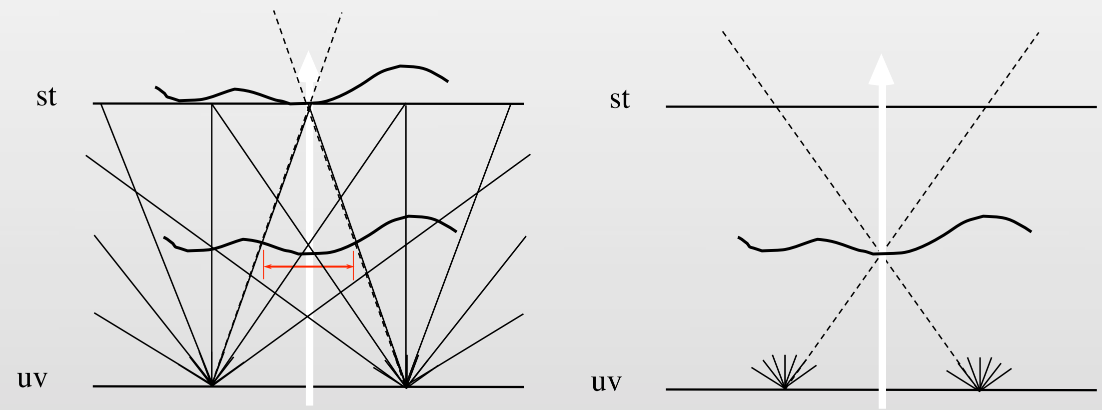


- disparity depends on density of samples in the light field and range of depths in the scene
  - [Levoy 1996]

- if the depth of a surface is known, a better set of rays can be extracted, reducing disparity
  - [Gortler 1996]

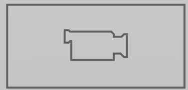
(alldragon.lif)

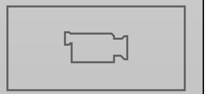
(lion.lif)

# Iso-disparity curves

[Chai et al., Siggraph 2000]

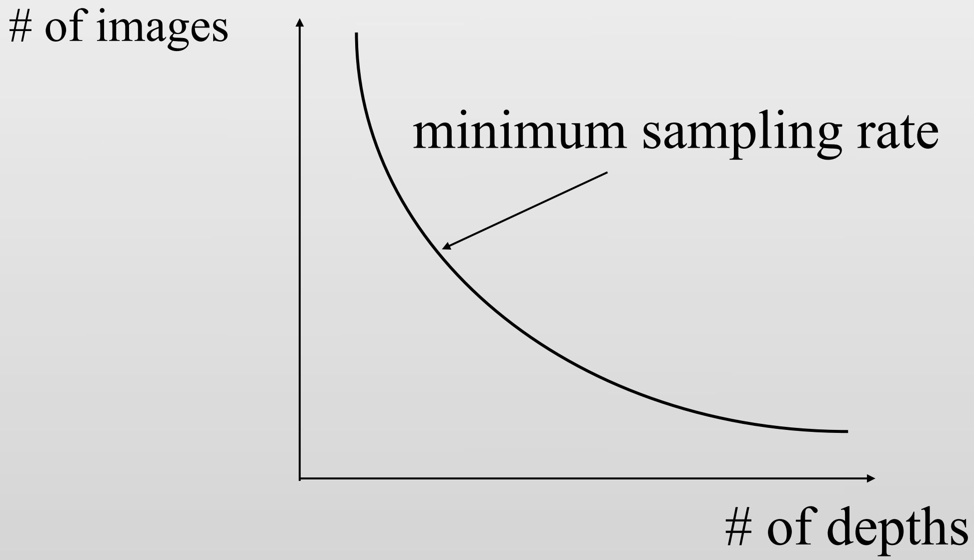


- for a given minimum acceptable disparity (in pixels), any combination of # of images and # of depths (in bits) falling on the curve will (barely) satisfy it

# Iso-disparity curves

[Chai et al., Siggraph 2000]

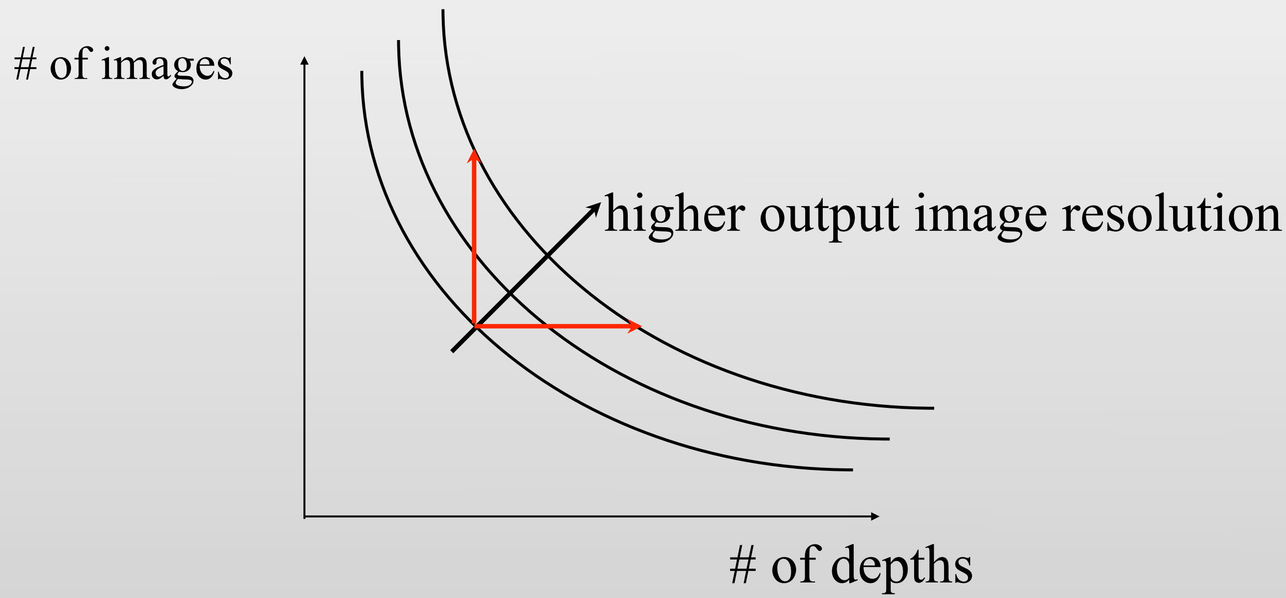


- higher output image resolution requires some combination of more images or more depths (to obtain a fixed circle of confusion (in pixels))

# Iso-disparity curves
[Chai et al., Siggraph 2000]

48 x 48 images, no geometry

(Chai)

16 x 16 images, 8 depths

# Outline

- scalar and vector light fields
- light field capture and rendering
  - –parameterization
  - –creation and display
  - –devices for capturing light fields
  - –sampling issues
- ☞ 3D reconstruction from light fields
- applications of light fields

# 3D reonstruction from light fields

light field

synthetic focal sequence

# Transpose of the light field

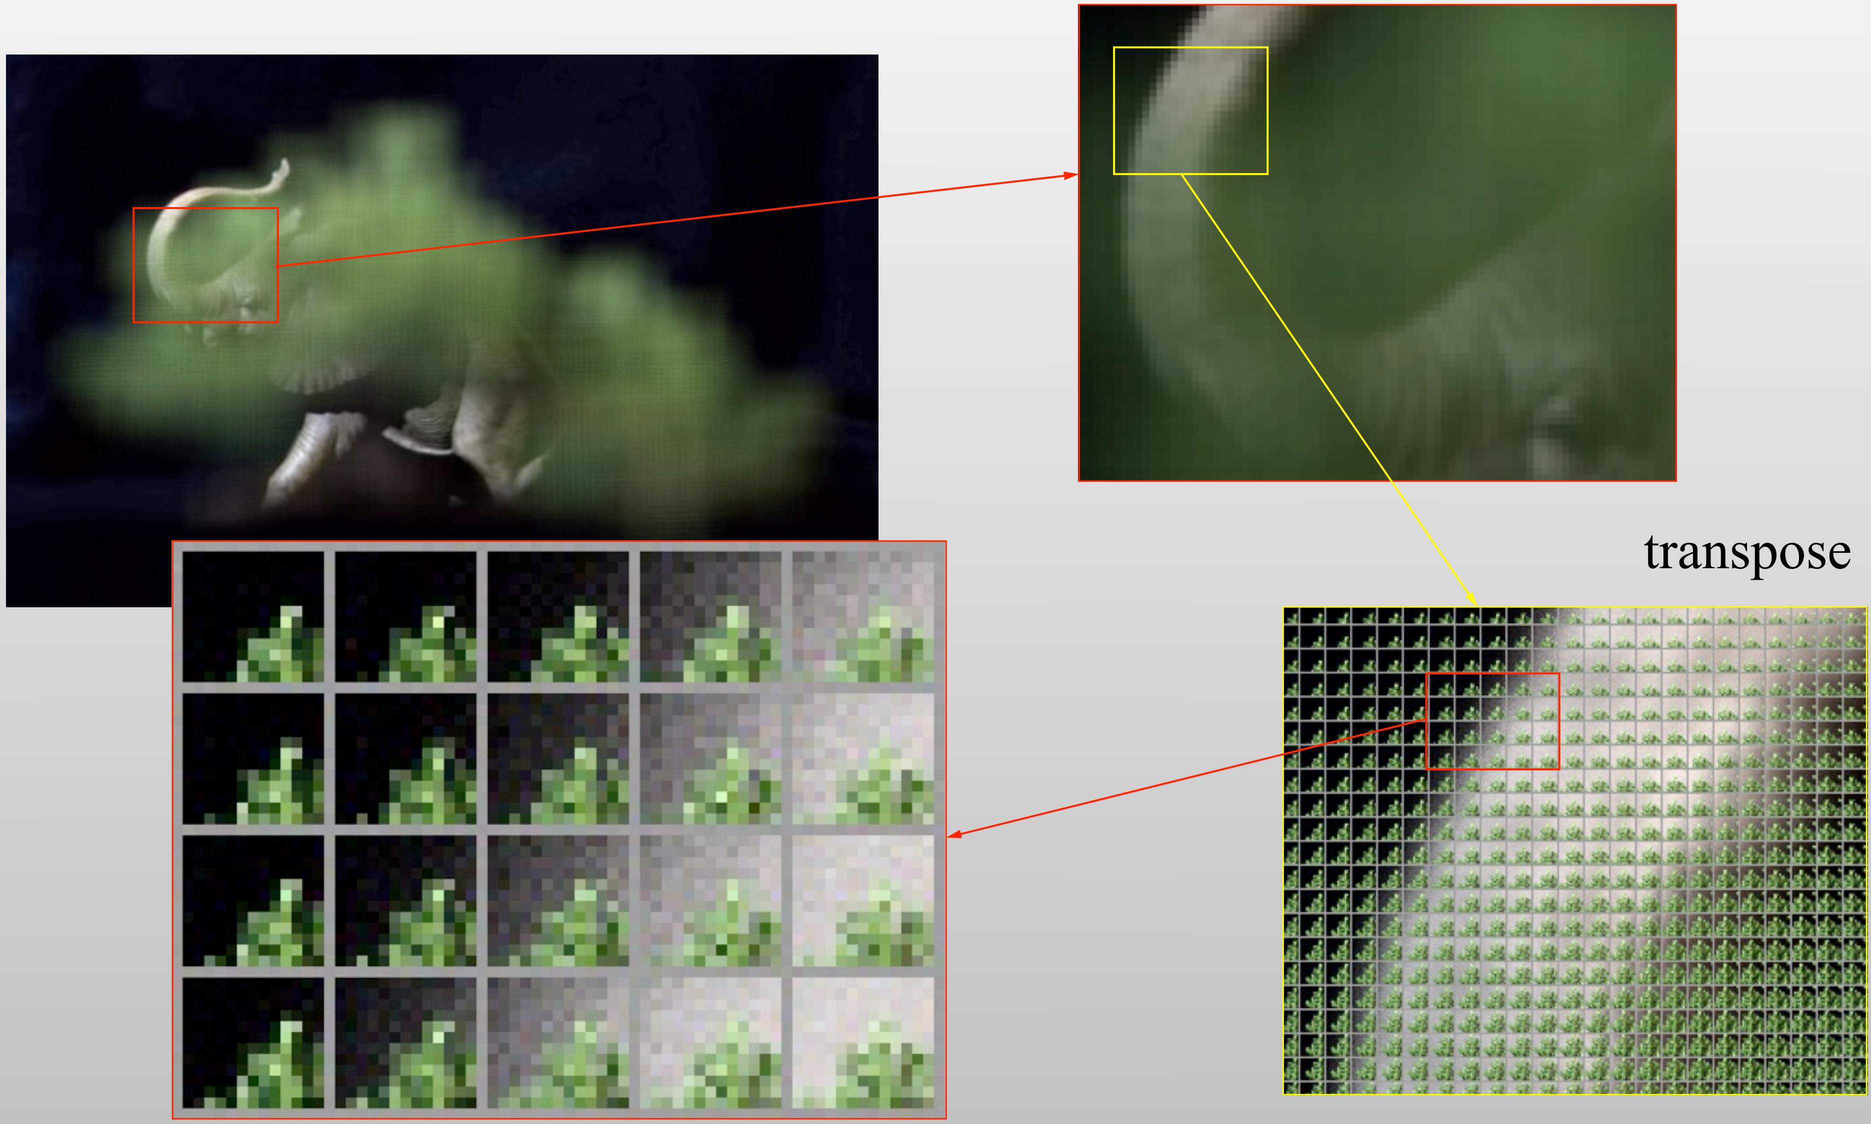

# Vision algorithms interpreted in line space

- flatland scene

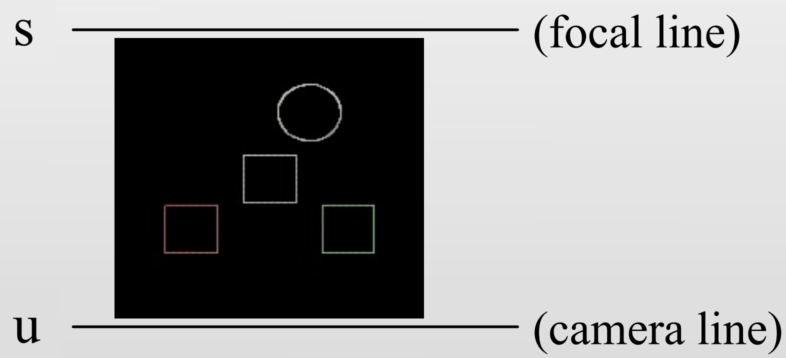


- flatland light field (a.k.a. epipolar image) [Bolles 1987]

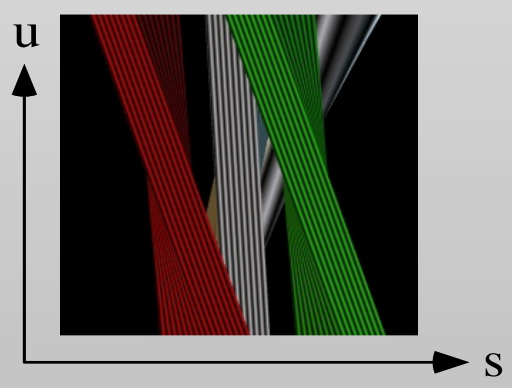

# Line space dualities

- a line

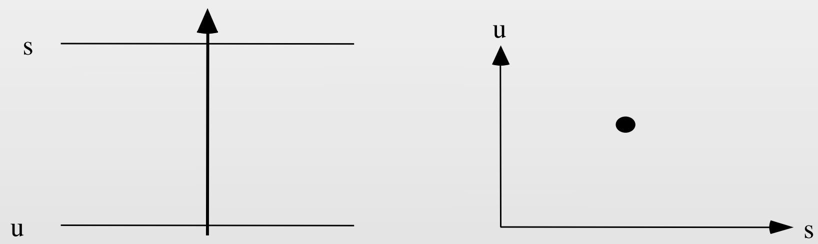


- lines through a point

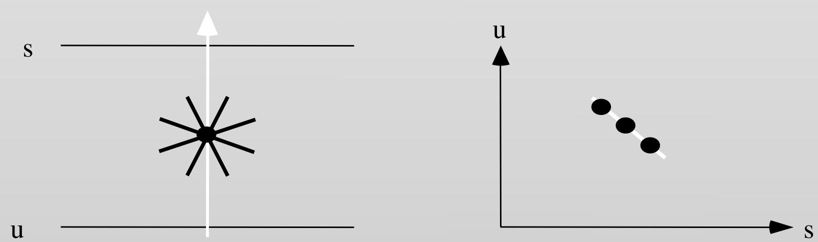


- points on a line

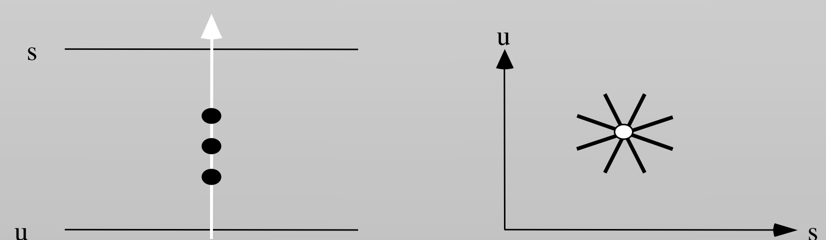

# Shape from stereo versus shape from focus

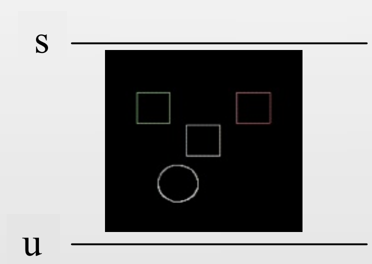


scene

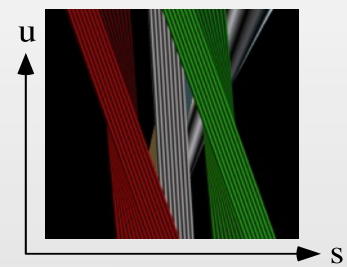


light field

- shape from stereo

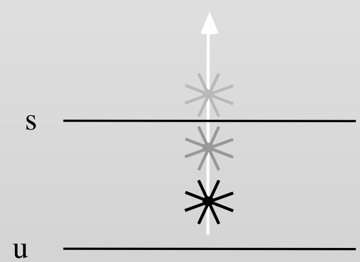


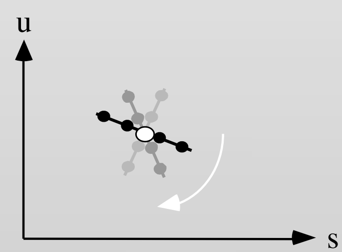


- shape from focus

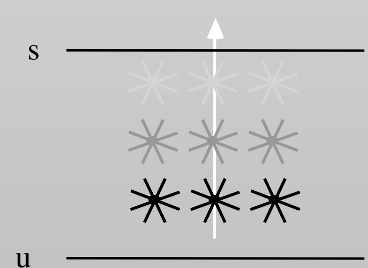


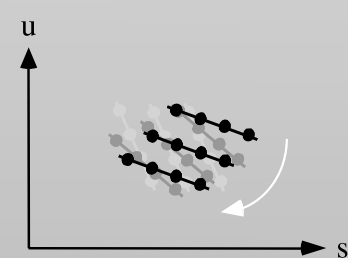

# Shape from focus example

rectified camera images

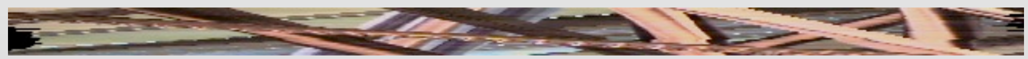

slice of epipolar volume at scanline 119

synthetic focus sequence

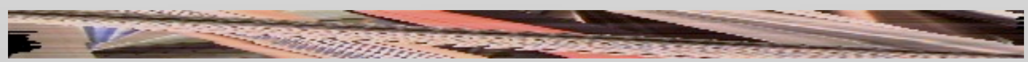

slice at scanline 261

# Shape from focus example

rectified camera images

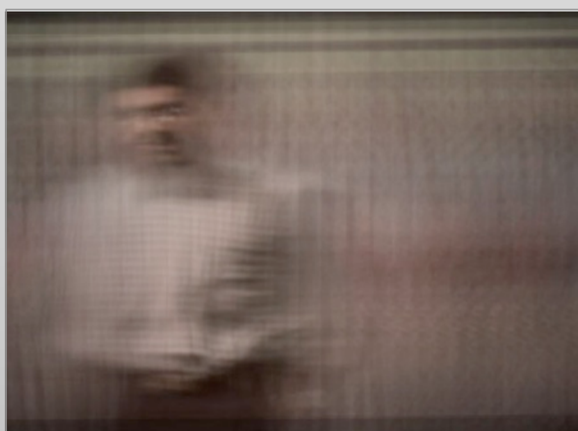

synthetic focus sequence

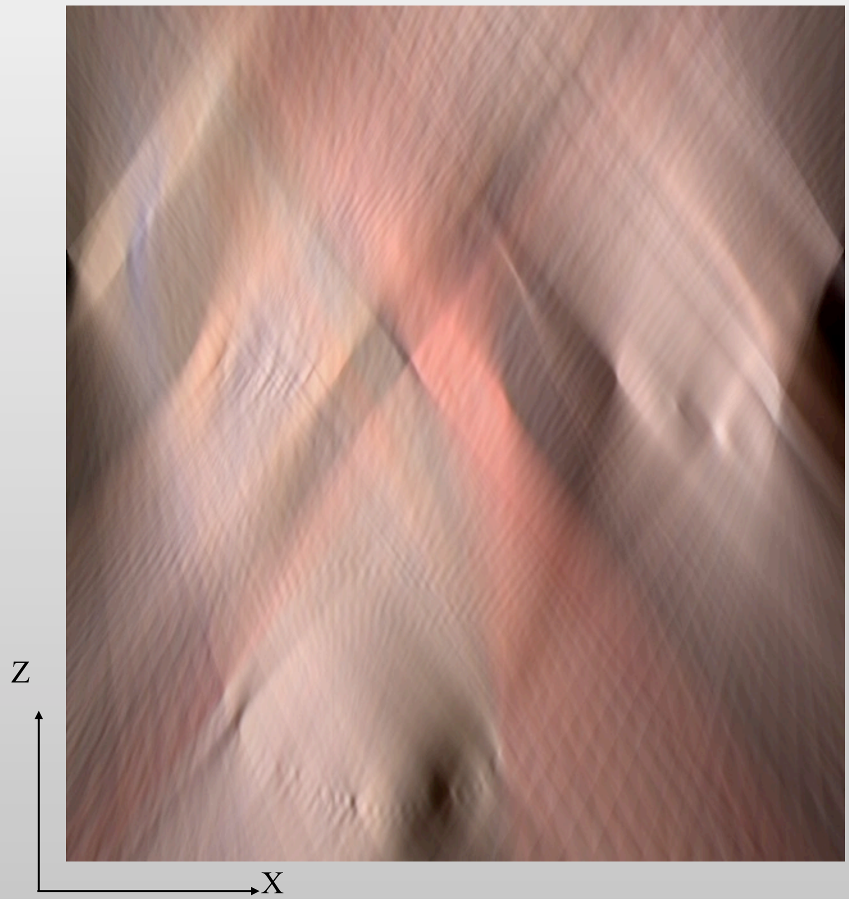


one scanline with different focal distances,
i.e. one slice from a focal stack

# Shape from focus example

rectified camera images

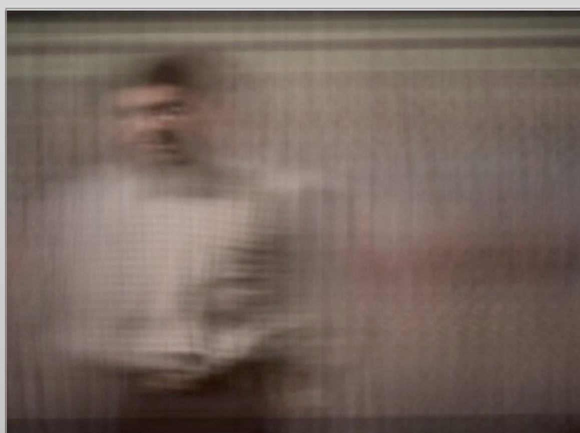

synthetic focus sequence

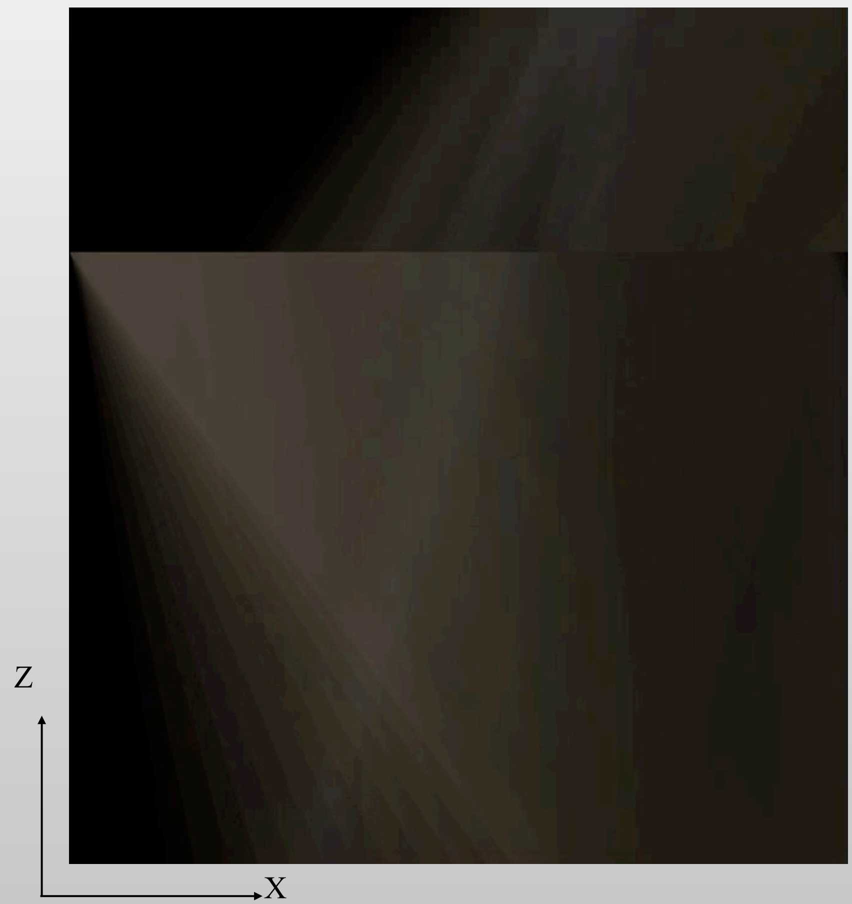


one scanline with different focal distances,
i.e. one slice from a focal stack

# Shape from focus example

rectified camera images

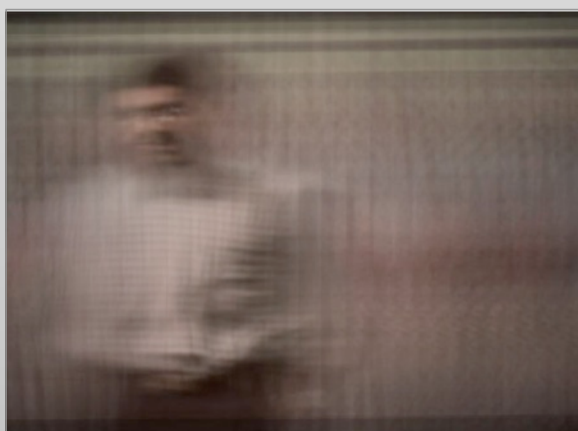

synthetic focus sequence

after applying x-sharpness operator

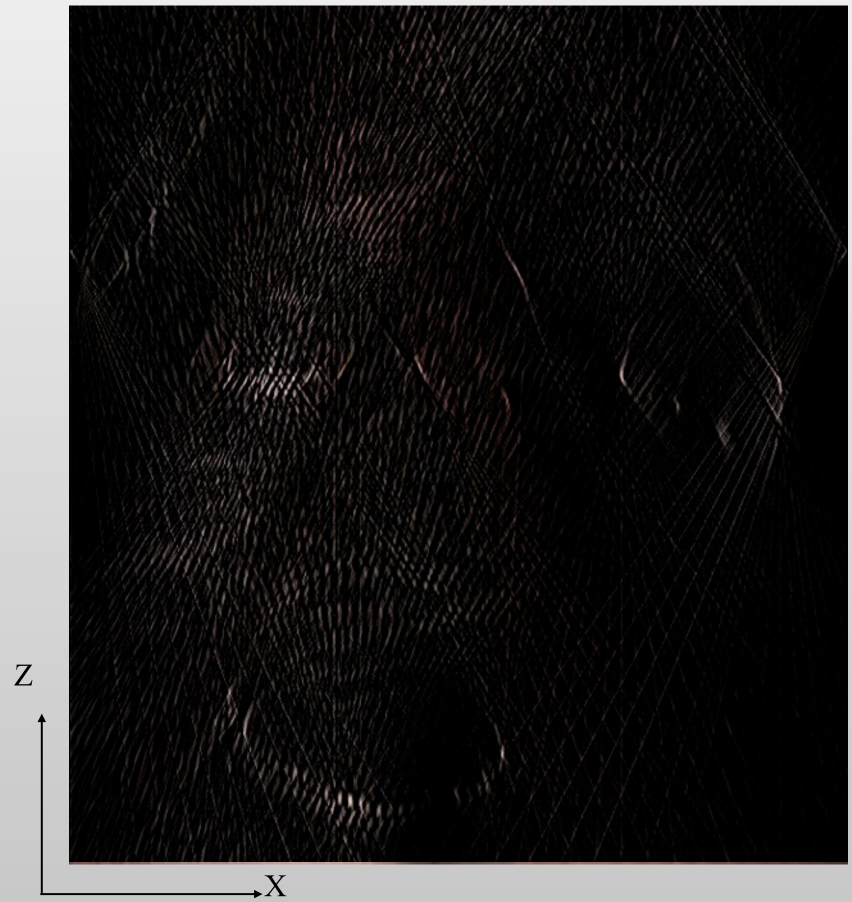


one scanline with different focal distances, i.e. one slice from a focal stack

# Which is better: stereo or focus?

- stereo
  - fails on heavily occluded scenes
- focus
  - fails on surfaces with linear ramp shading
- hybrid?

# Outline

- scalar and vector light fields
- light field capture and rendering
  - –parameterization
  - –creation and display
  - –devices for capturing light fields
  - –sampling issues
- 3D reconstruction from light fields
- ☞ applications of light fields

# Applications of light fields

- perspective flyarounds (light field rendering)
- digital refocusing
- 3D reconstruction
- 4D texture synthesis
- light field editing
- light field morphing
- autostereoscopic display of light fields

# Plenoptic image editing

[Seitz 1998]

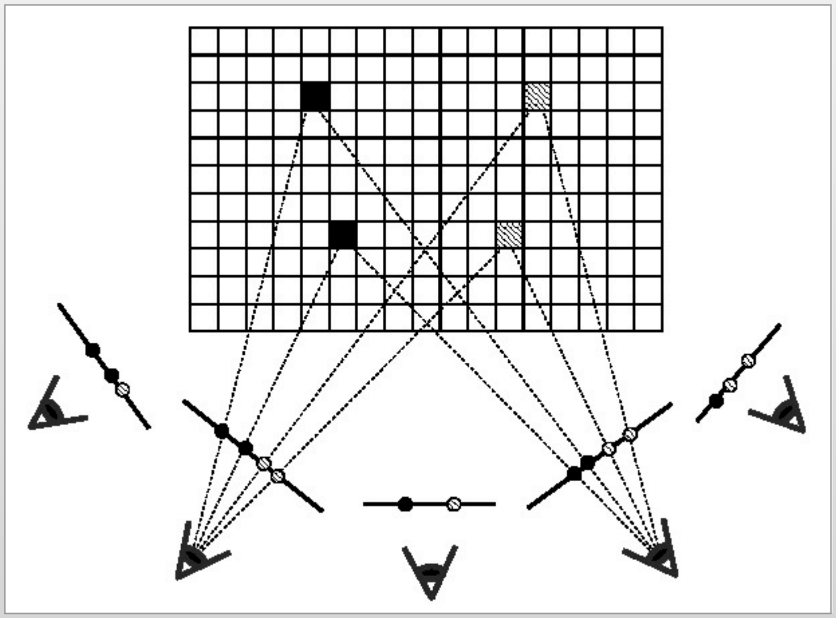

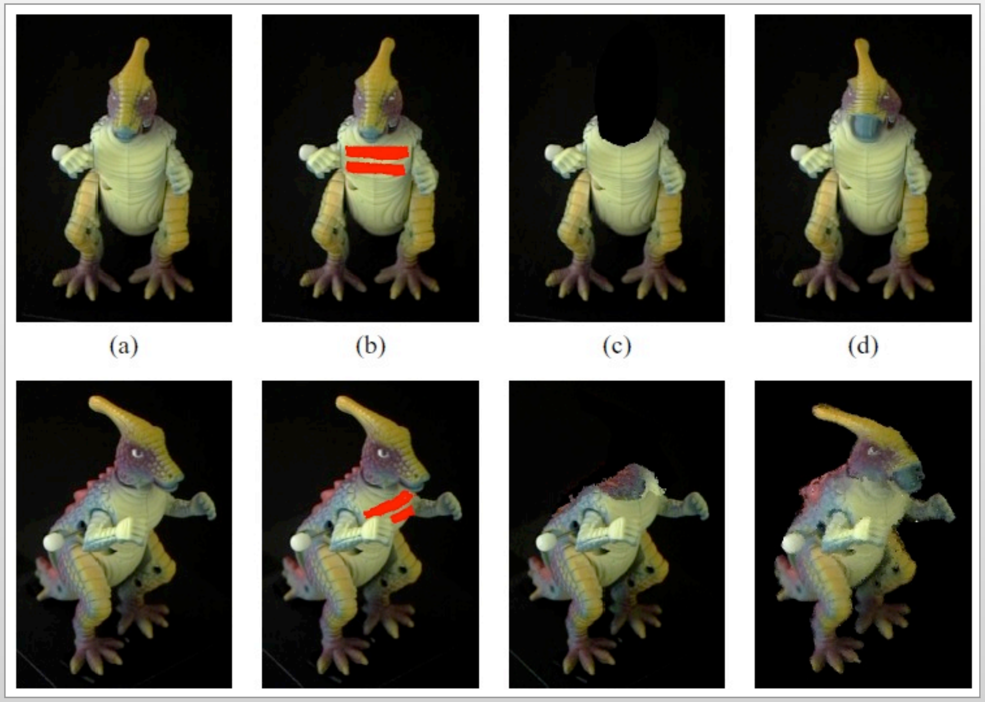


- use color consistency constraints [Seitz 1997] to obtain voxel model
- store mapping between pixels in images and voxels in model
- for each pixel changed during editing
  - change corresponding voxel
  - change corresponding pixels in all other views

# Light field morphing
[Zhang et al., SIGGRAPH 2002]

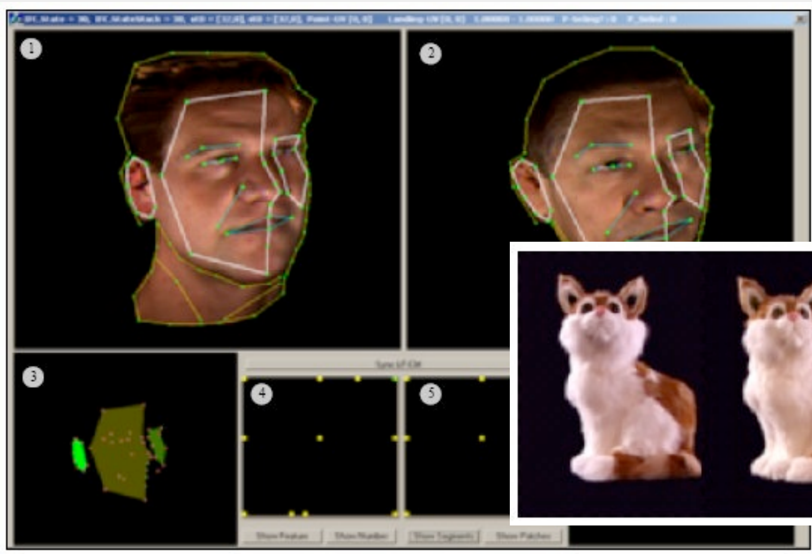

UI for specifying feature polygons
and their correspondences

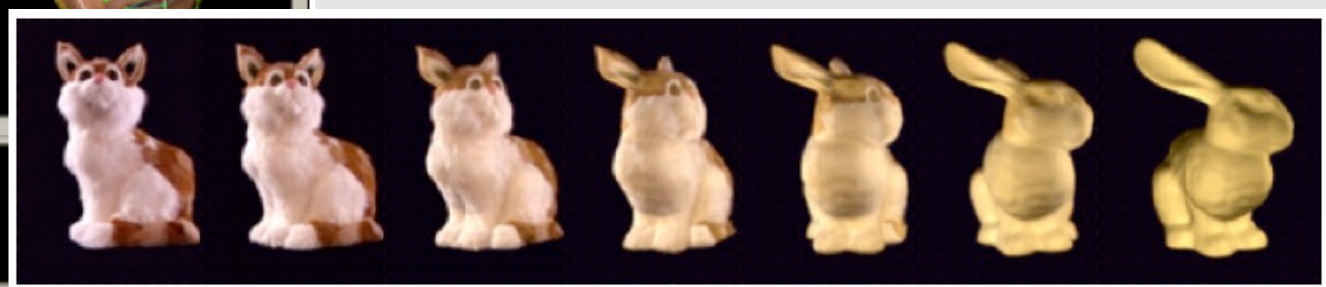

sample morph

- feature correspondences = 3D model

# Autostereoscopic display of light fields

[Isaksen 2000]

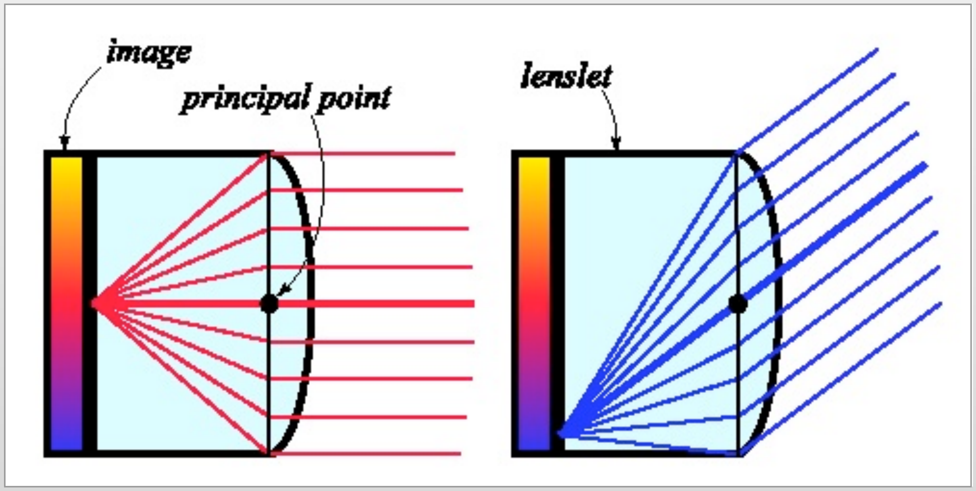


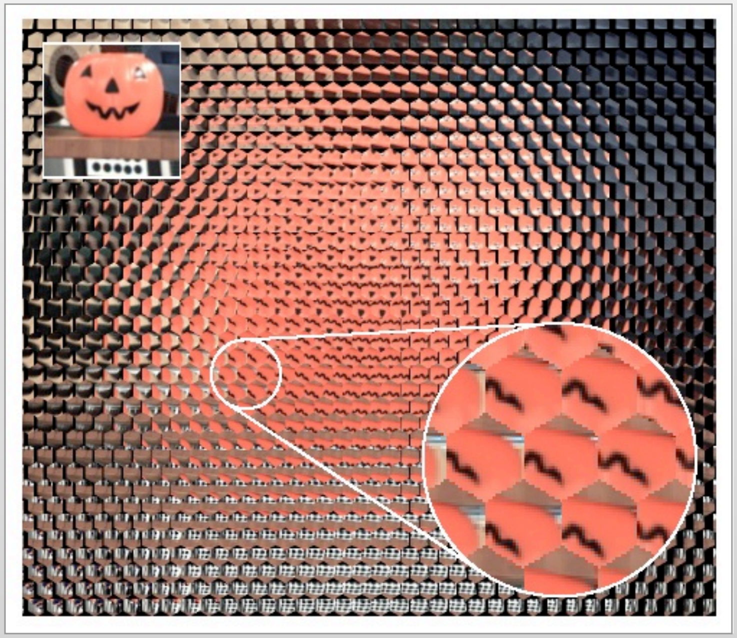

- image is at focal distance of lenslet ⇒ collimated rays
- spatial resolution ~ # of lenslets in the array
- angular resolution ~ # of pixels behind each lenslet
- each eye sees a different sets of pixels ⇒ stereo

# Autostereoscopic display of light fields
[Matusik 2004]

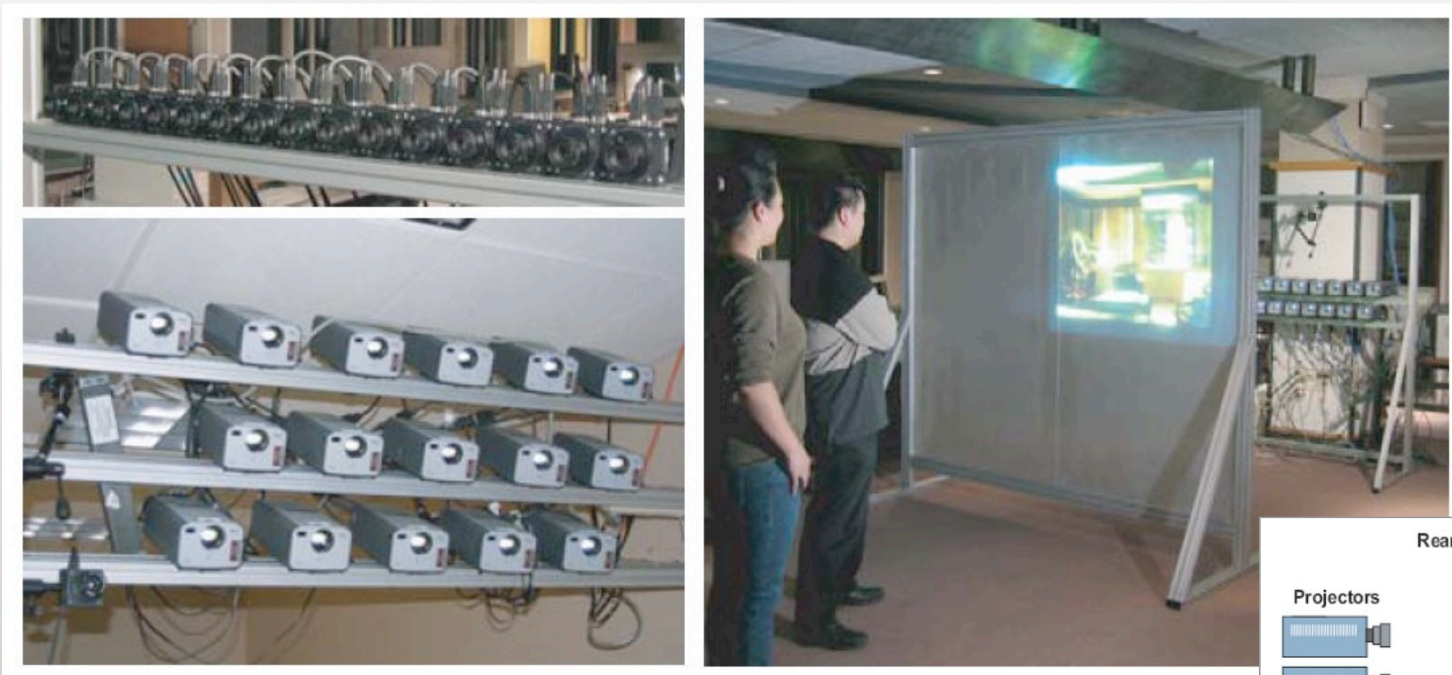

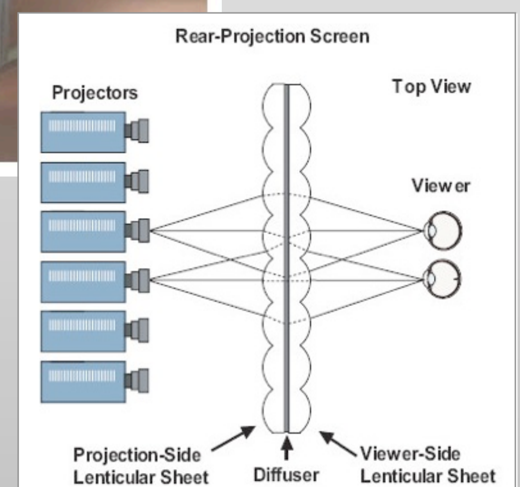


- 16 cameras, 16 projectors
- spatial resolution ~ # of pixels in projector
- angular resolution ~ # of projectors
- # of lenslets is unimportant

# Slide credits


- Gershun, A., "The Light Field", Moscow, 1936. Translated by P. Moon and G. Timoshenko in *Journal of Mathematics and Physics*, Vol. XVIII, MIT, 1939, pp. 51–151.
- Adelson, E.H., Bergen, J.R., “The plenoptic function and the elements of early vision,” In *Computation Models of Visual Processing*, MIT Press, 1991.
- Moon, P., Spencer, D.E., *The Photic Field*, MIT Press, 1981.
- Chen, S.E., Williams, L., “View interpolation for image synthesis,” *Proc. SIGGRAPH 1995.*
- Levoy, M., Hanrahan, P., “Light Field Rendering,” *Proc. SIGGRAPH 1996.*
- Wood, D.N. et al., “Surface Light Fields for 3D Photography,” *Proc. SIGGRAPH 2000.*
- Salesin, D., Stolfi, J., “Rendering CSG Models with a ZZ-Buffer,” *Proc. SIGGRAPH 1990.*
- Levoy, M., Stanford Spherical Gantry, 2002, http://graphics.stanford.edu/projects/gantry/.
- Yang, J.C., Everett, M., Buehler, C., McMillan, L., “A Real-Time Distributed Light Field Camera,” *Proc. Eurographics Rendering Workshop 2002.*
- Zhang, C., Chen, T., “A Self-Reconfigurable Camera Array,” *Proc. Eurographics Rendering Workshop 2004*.
- Shum, H., He, L.-W., “Rendering with Concentric Mosaics,” *Proc. SIGGRAPH 1999*.
- Wilburn, B., et al., “High Performance Imaging Using Large Camera Arrays,” *Proc. SIGGRAPH 2005.*
- Ng, R., Levoy, M., Brédif, M., Duval, G., Horowitz, M., Hanrahan, P., “Light Field Photography with a Hand-Held Plenoptic Camera,” *Stanford Tech Report CTSR 2005-02,* April, 2005.

- Georgiev, T. et al., “Spatio-Angular Resolution Tradeoff in Integral Photography,” *Proc. Eurographics Symposium on Rendering 2006.*
- Gortler, S.J., Grzeszczuk, R., Szeliski, R., Cohen, M.F., “The Lumigraph,” *Proc. SIGGRAPH 1996.*
- Buehler, C. et al., “Unstructured Lumigraph rendering,” *Proc. SIGGRAPH 2001.*
- Vaish, V., Wilburn, B., Joshi, N., Levoy, M., “Using Plane + Parallax for Calibrating Dense Camera Arrays,” *Proc. CVPR 2004* .
- Levoy, M., Ng, R., Adams, A., Footer, M., Horowitz, M., “Light field microscopy,” *Proc. SIGGRAPH 2006.*
- Han, J.Y., Perlin, K., “Measuring Bidirectional Texture Reflectance With a Kaleidoscope,” *Proc. SIGGRAPH 2003.*
- Bolles, R., Baker, H., Marimont, D., “Epipolar-Plane Image Analysis: An Approach to Determining Structure from Motion”, *IJCV* **1**(1), 1987.
- Seitz, S.M., Dyer, C.R., “Photorealistic Scene Reconstruction by Voxel Coloring,” *Proc. CVPR 1997.*
- Seitz, S.M., Kutulakos, K.N., “Plenoptic Image Editing,” *Proc. ICCV 1998.*
- Zhang, Z., Wang, L., Guo, B., Shum, H.-Y., “Feature-based light field morphing,” *Proc. SIGGRAPH 2002.*
- Isaksen, A., McMillan, L., Gortler, S.J., “Dynamically reparameterized light fields,” *Proc. SIGGRAPH 2000.*
- Matusik, W. et al., “3D TV: A Scalable System for Real-Time Acquisition, Transmission, and Autostereoscopic Display of Dynamic Scenes,” *Proc. SIGGRAPH 2004.*