Image-based Rendering

© Michal Havlik

15-462: Graphics I
Alexei Efros, CMU, Fall 2007

What is light?

Electromagnetic radiation (EMR) moving along rays in space

- R(λ) is EMR, measured in units of power (watts)
  - λ is wavelength

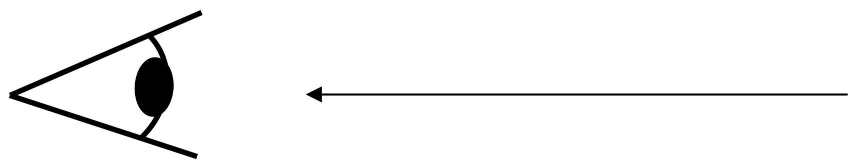

Useful things:

- Light travels in straight lines
- In vacuum, radiance emitted = radiance arriving
  - i.e. there is no transmission loss

What do we see?

3D world

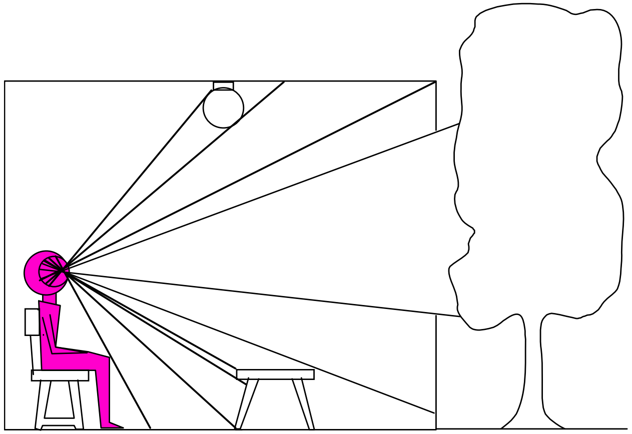

Point of observation

2D image

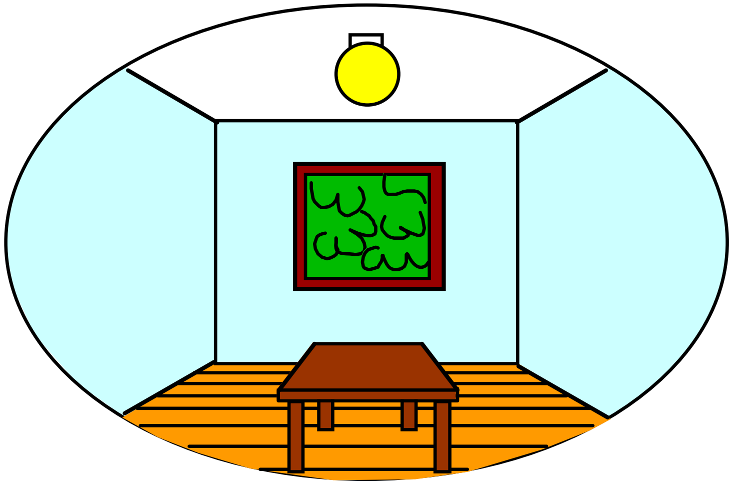

What do we see?

3D world

2D image

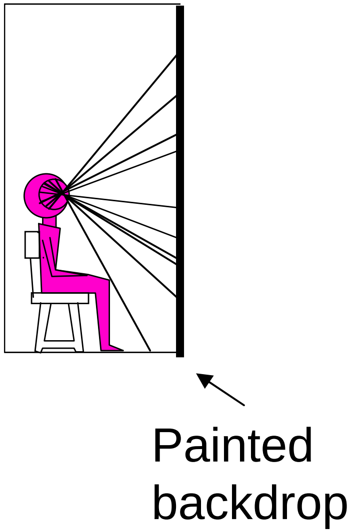

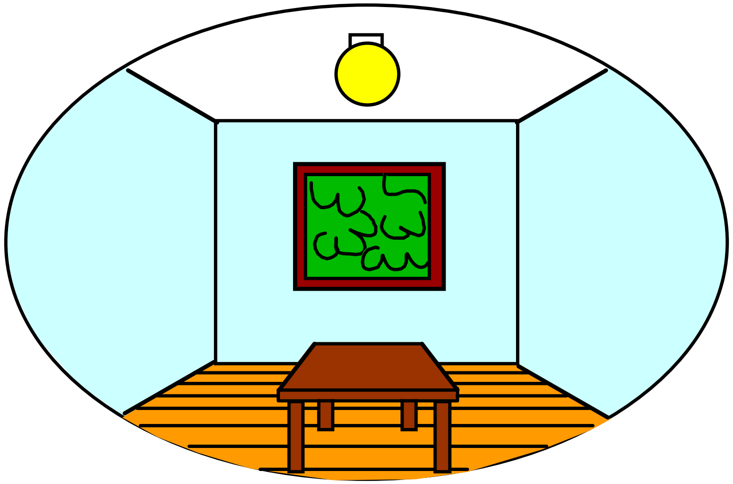

On Simulating the Visual Experience

Just feed the eyes the right data

- No one will know the difference!

Philosophy:

- Ancient question: “Does the world really exist?”

Science fiction:

- Many, many, many books on the subject, e.g. *slowglass*
- Latest take: *The Matrix*

Physics:

- <u>*Slowglass*</u> might be possible?

Computer Science:

- Virtual Reality

To simulate we need to know:

What does a person see?

The Plenoptic Function

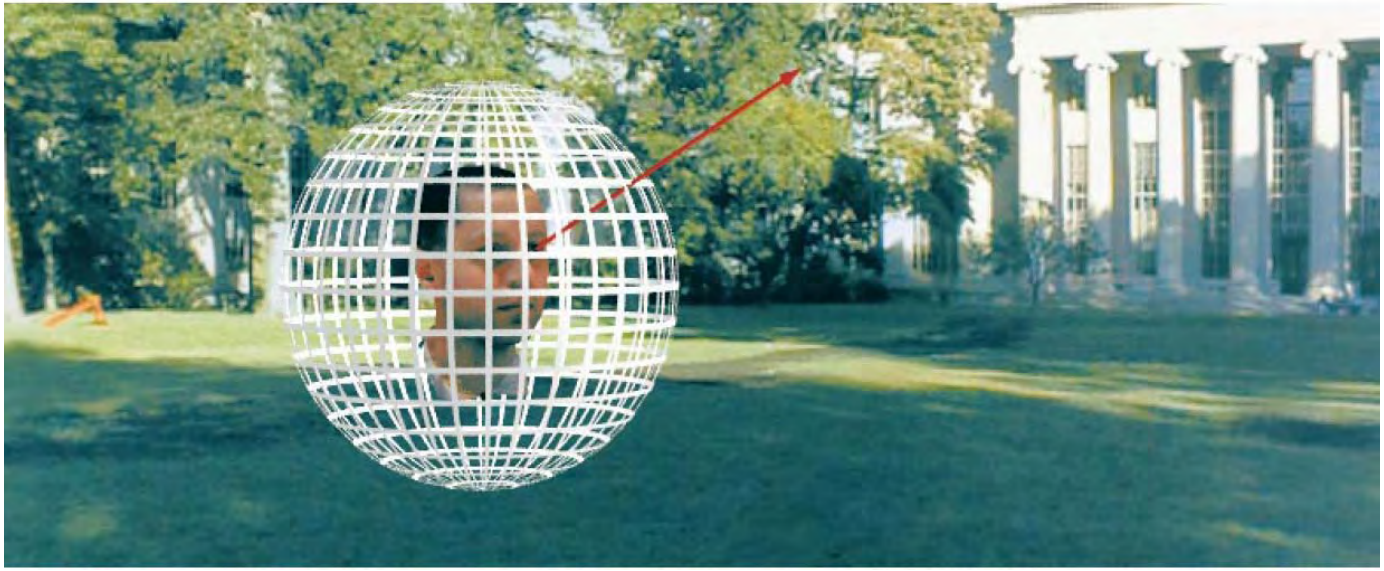

Figure by Leonard McMillan

Q: What is the set of all things that we can ever see?

A: The Plenoptic Function (Adelson & Bergen)

Let's start with a stationary person and try to parameterize <u>everything</u> that he can see…

Grayscale snapshot

$$P(\theta,\phi)$$

is intensity of light

- Seen from a single view point
- At a single time
- Averaged over the wavelengths of the visible spectrum

(can also do *P(x,y)*, but spherical coordinate are nicer)

Color snapshot

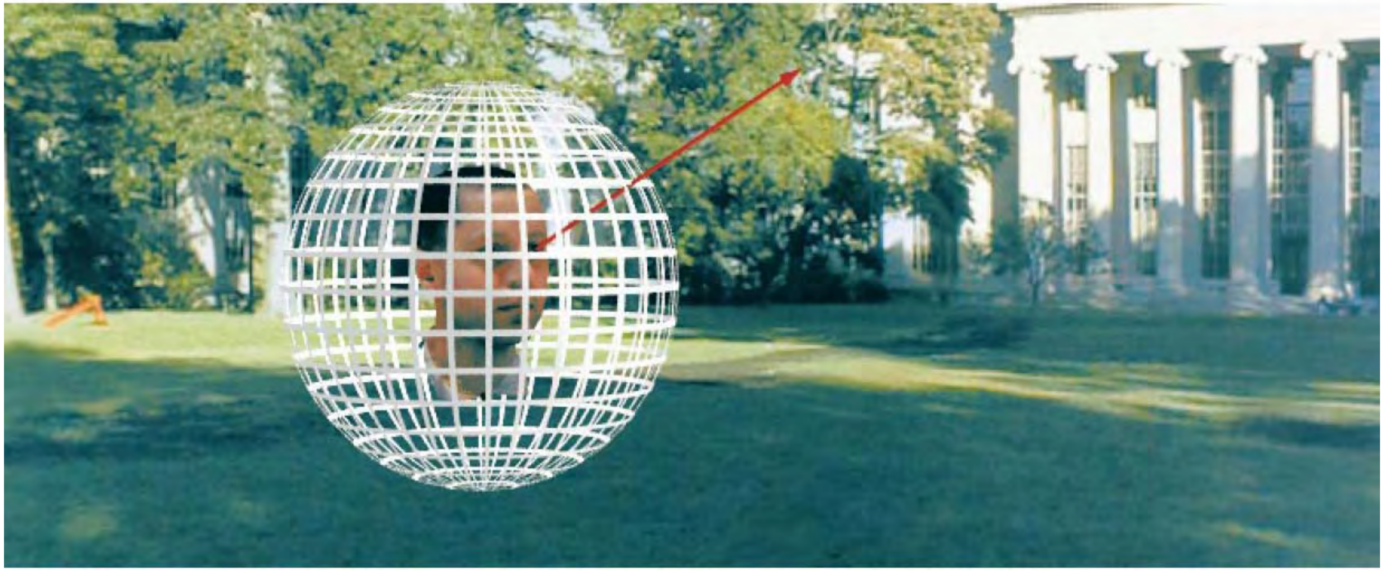

$$P(\theta,\phi,\lambda)$$

is intensity of light

- Seen from a single view point
- At a single time
- As a function of wavelength

A movie

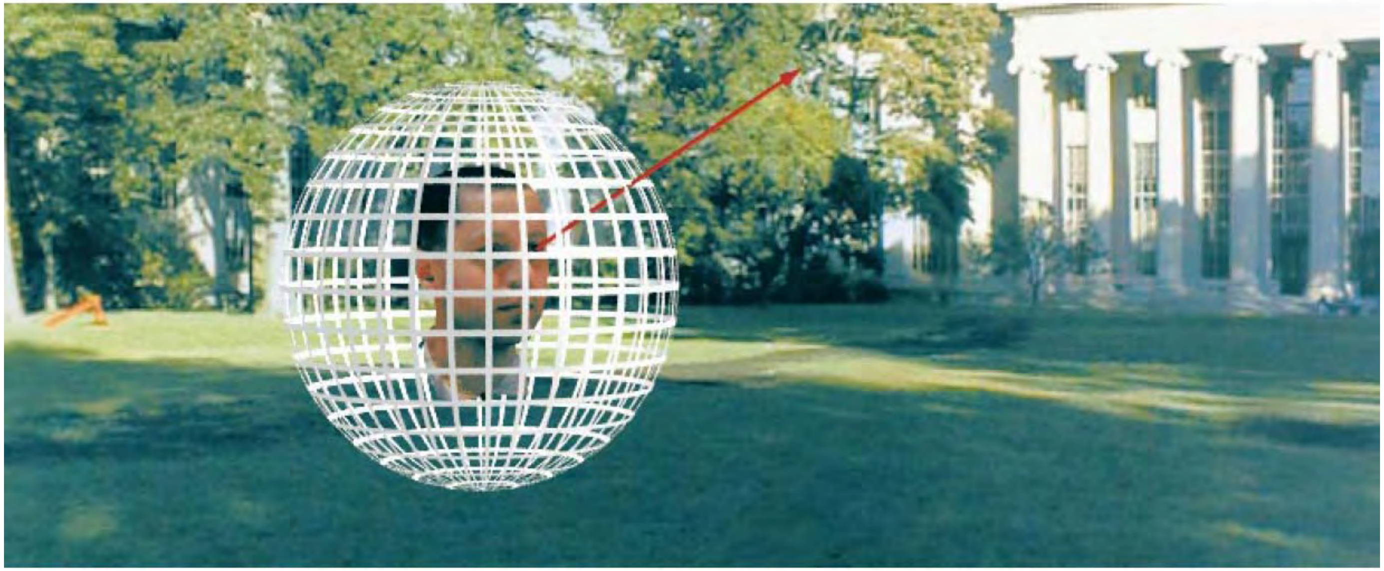

$$P(\theta,\phi,\lambda,t)$$

is intensity of light

- Seen from a single view point
- Over time
- As a function of wavelength

Holographic movie

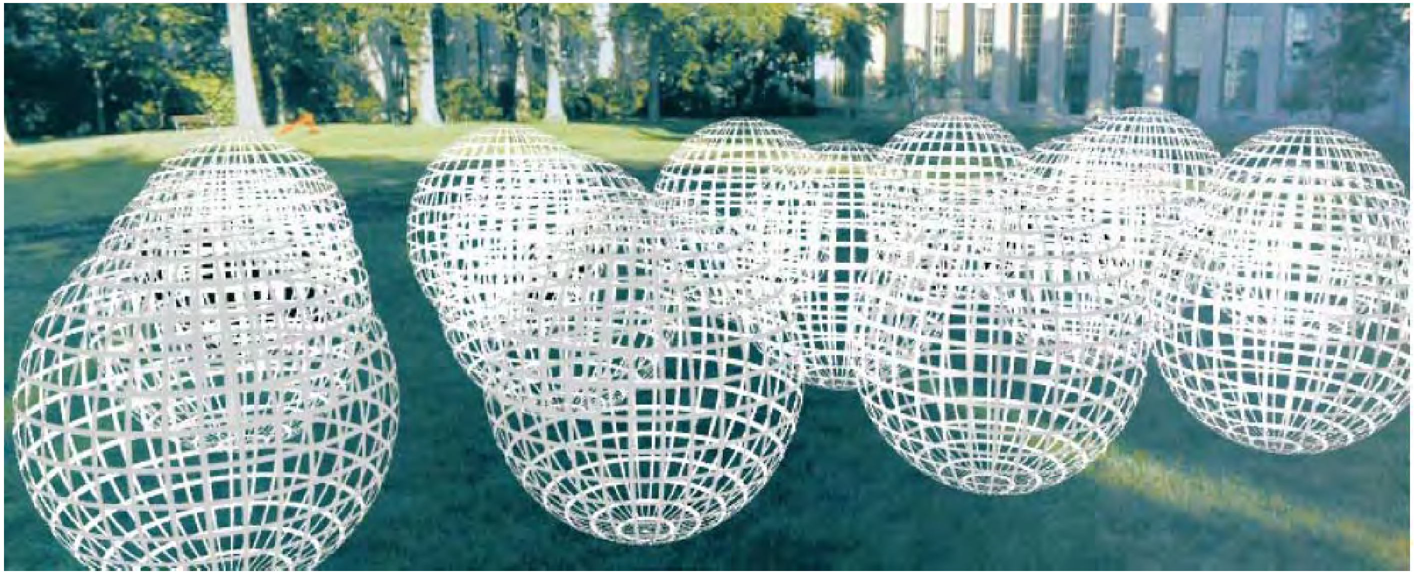

$$P(\theta,\phi,\lambda,t,V_X,V_Y,V_Z)$$

is intensity of light

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

The Plenoptic Function

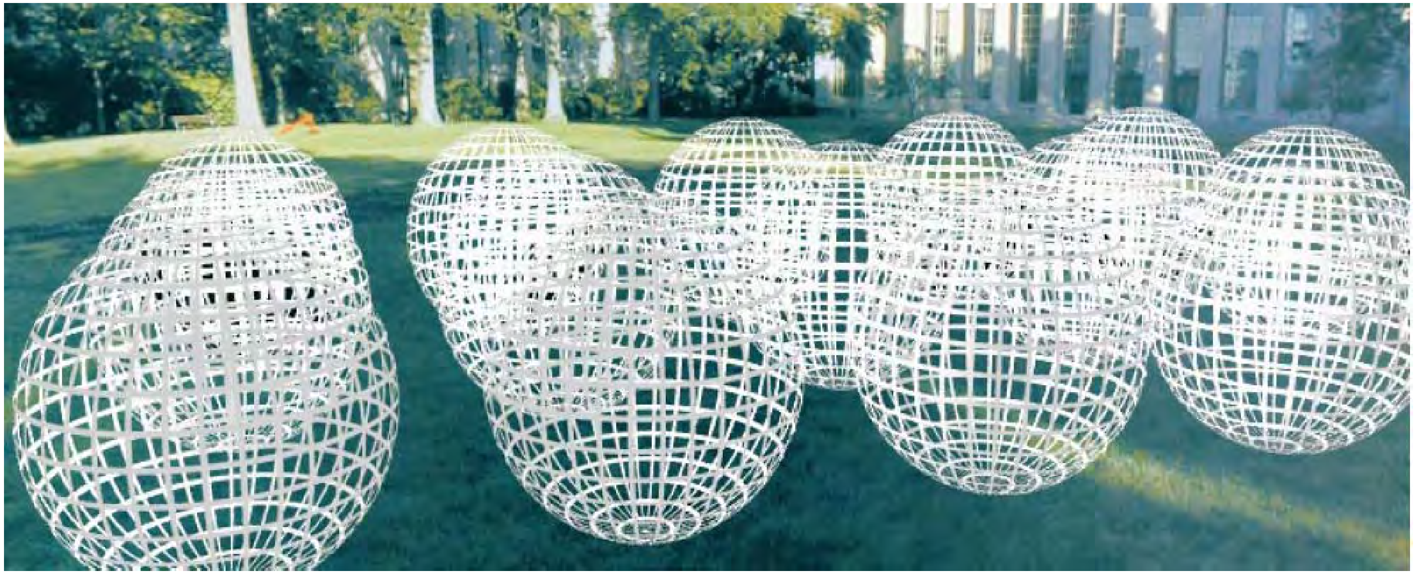

$$P(\theta,\phi,\lambda,t,V_X,V_Y,V_Z)$$

- 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! it completely captures our visual reality!   Not bad for a function…

Sampling Plenoptic Function (top view)

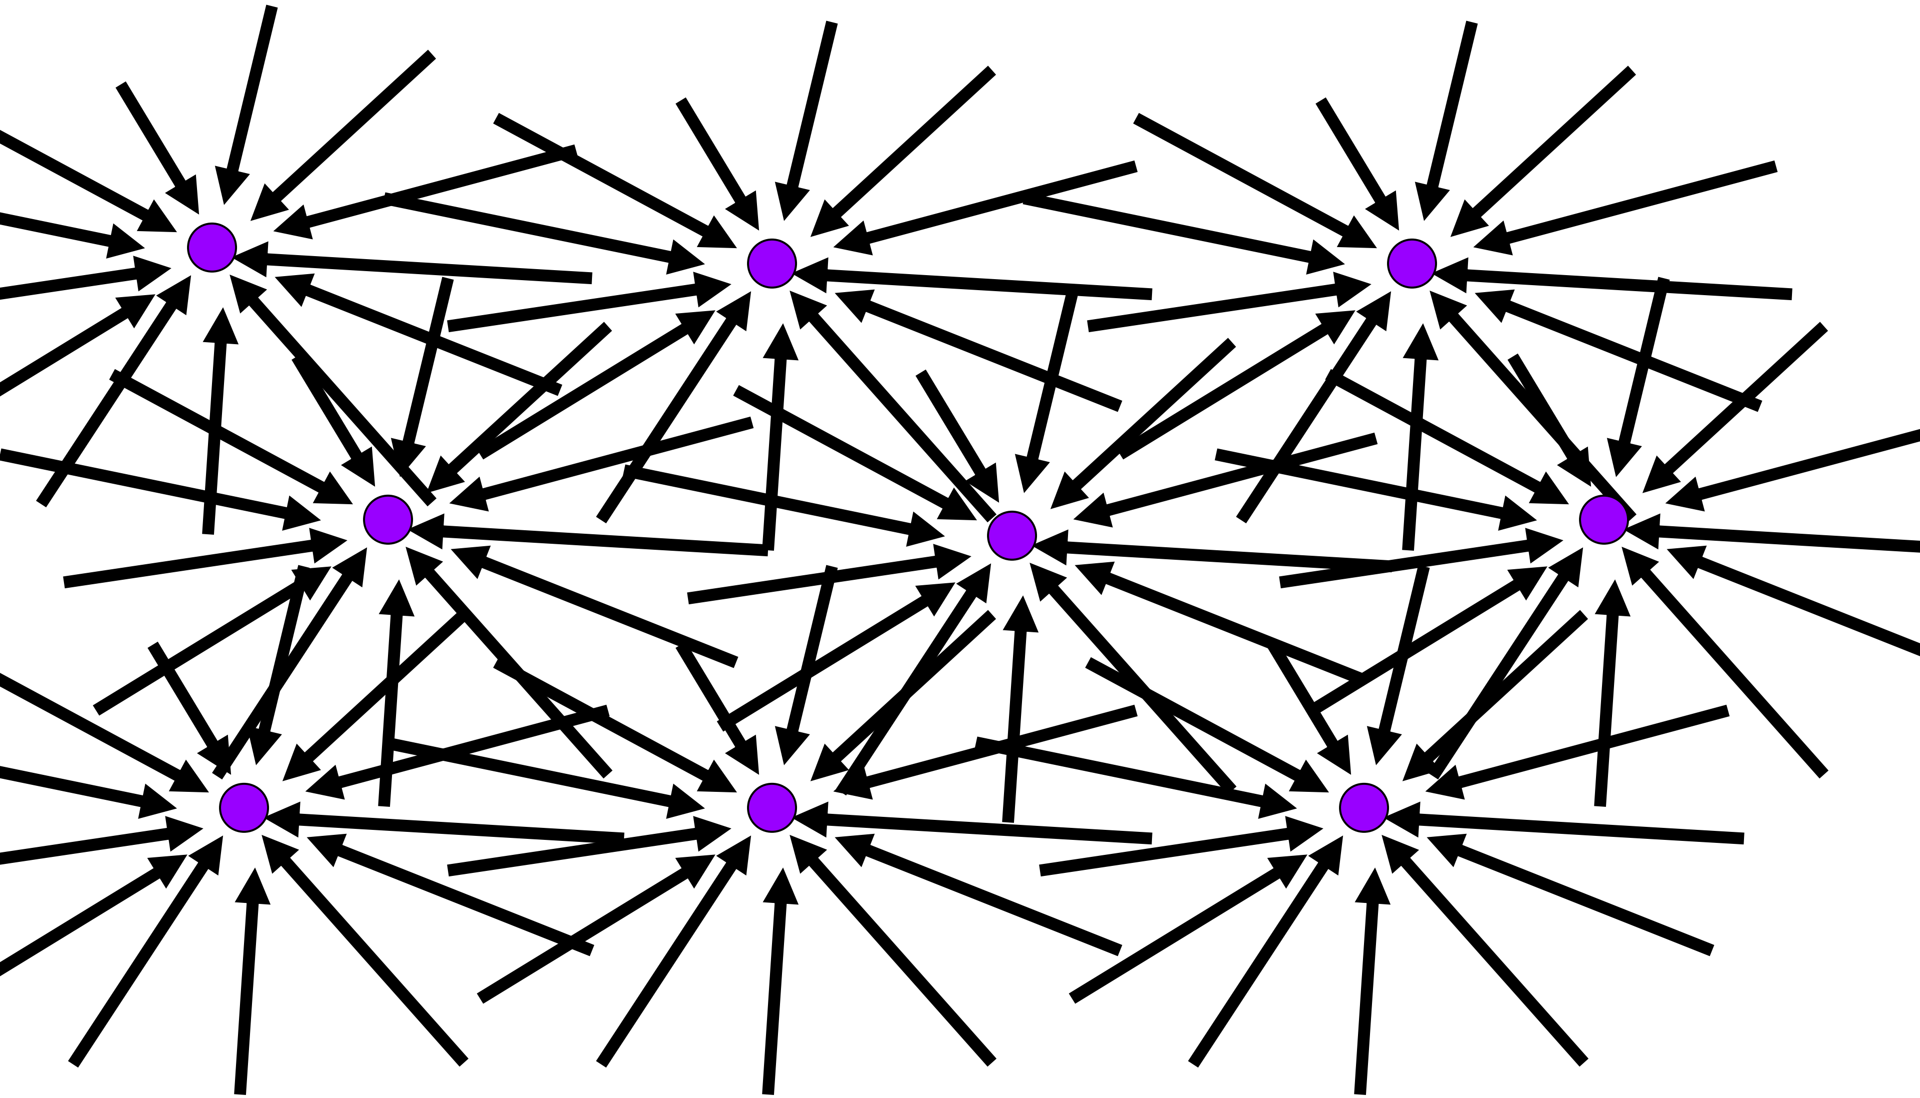

Just lookup -- Quicktime VR

Ray

Let’s not worry about time and color:

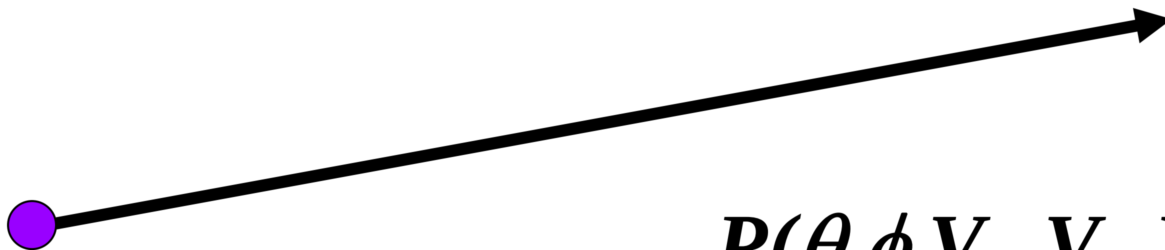

$$P(\theta,\phi,V_X,V_Y,V_Z)$$

5D

- 3D position
- 2D direction

Slide by Rick Szeliski and Michael Cohen

How can we use this?

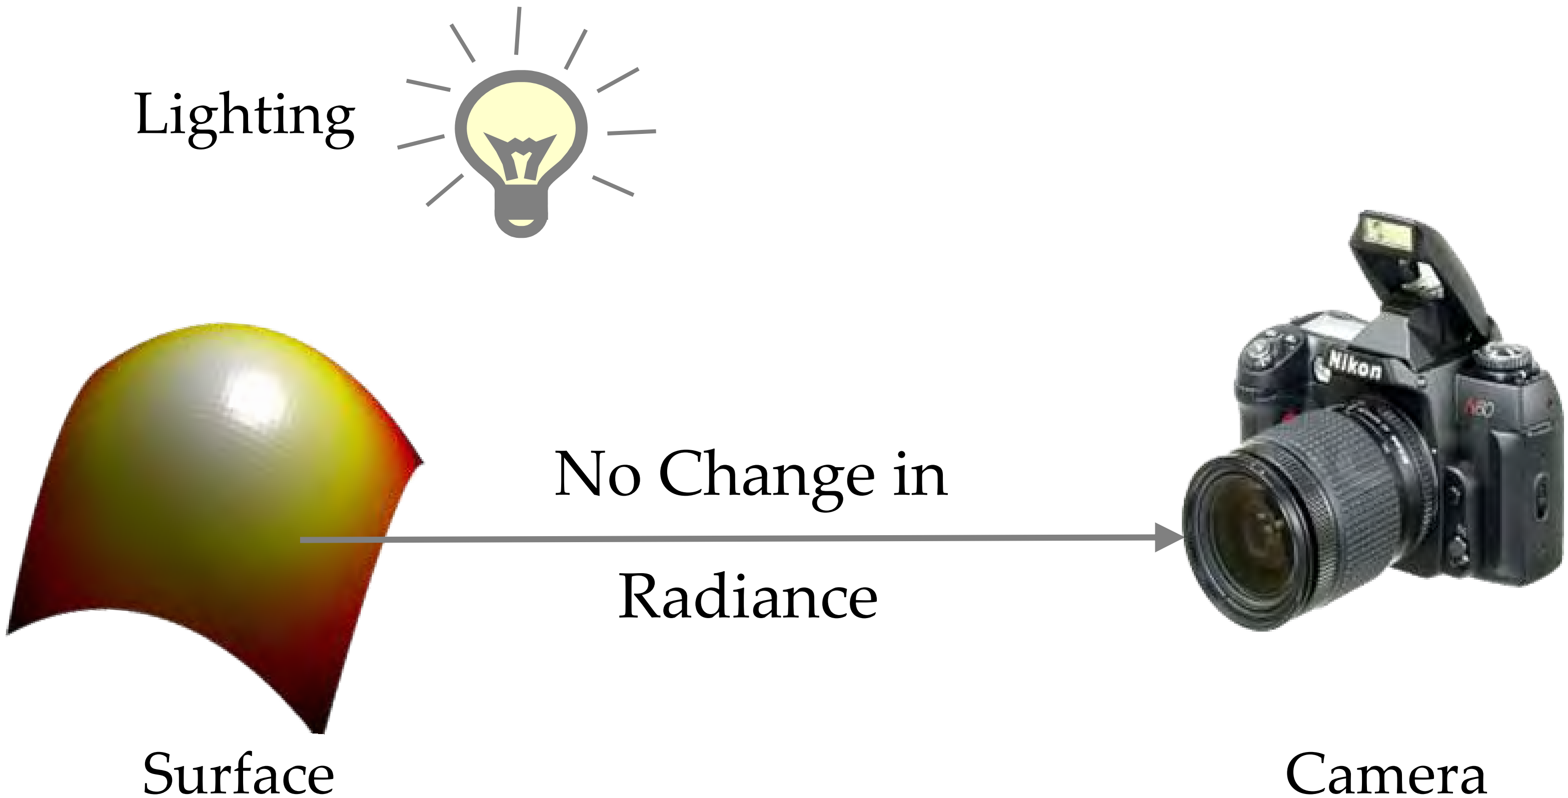

Ray Reuse

Infinite line

- Assume light is constant (vacuum)

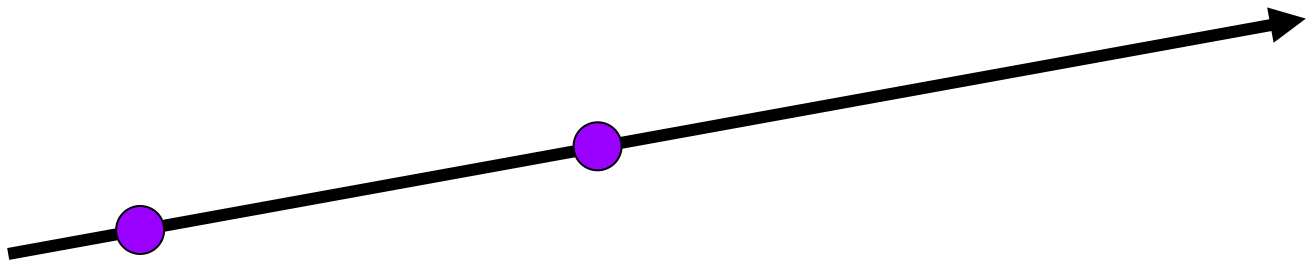

4D

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

Slide by Rick Szeliski and Michael Cohen

Only need plenoptic surface

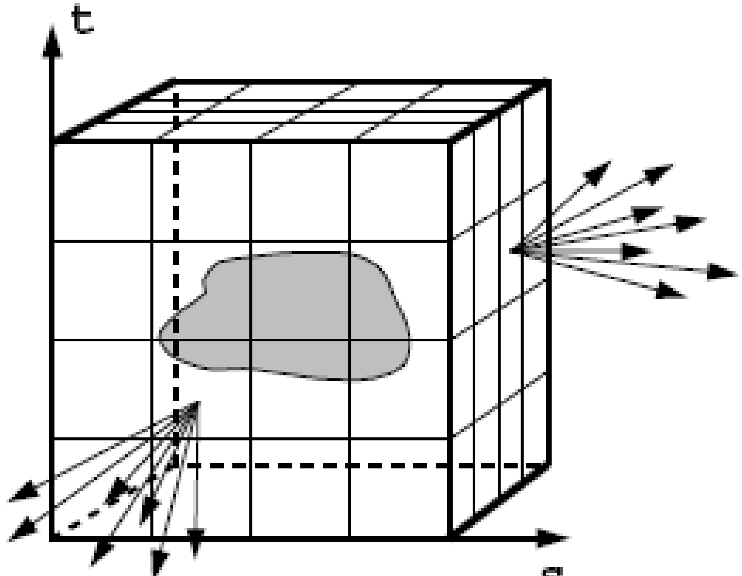

Figure 1: The surface of a cube holds all the radiance information due to the enclosed object.

Synthesizing novel views

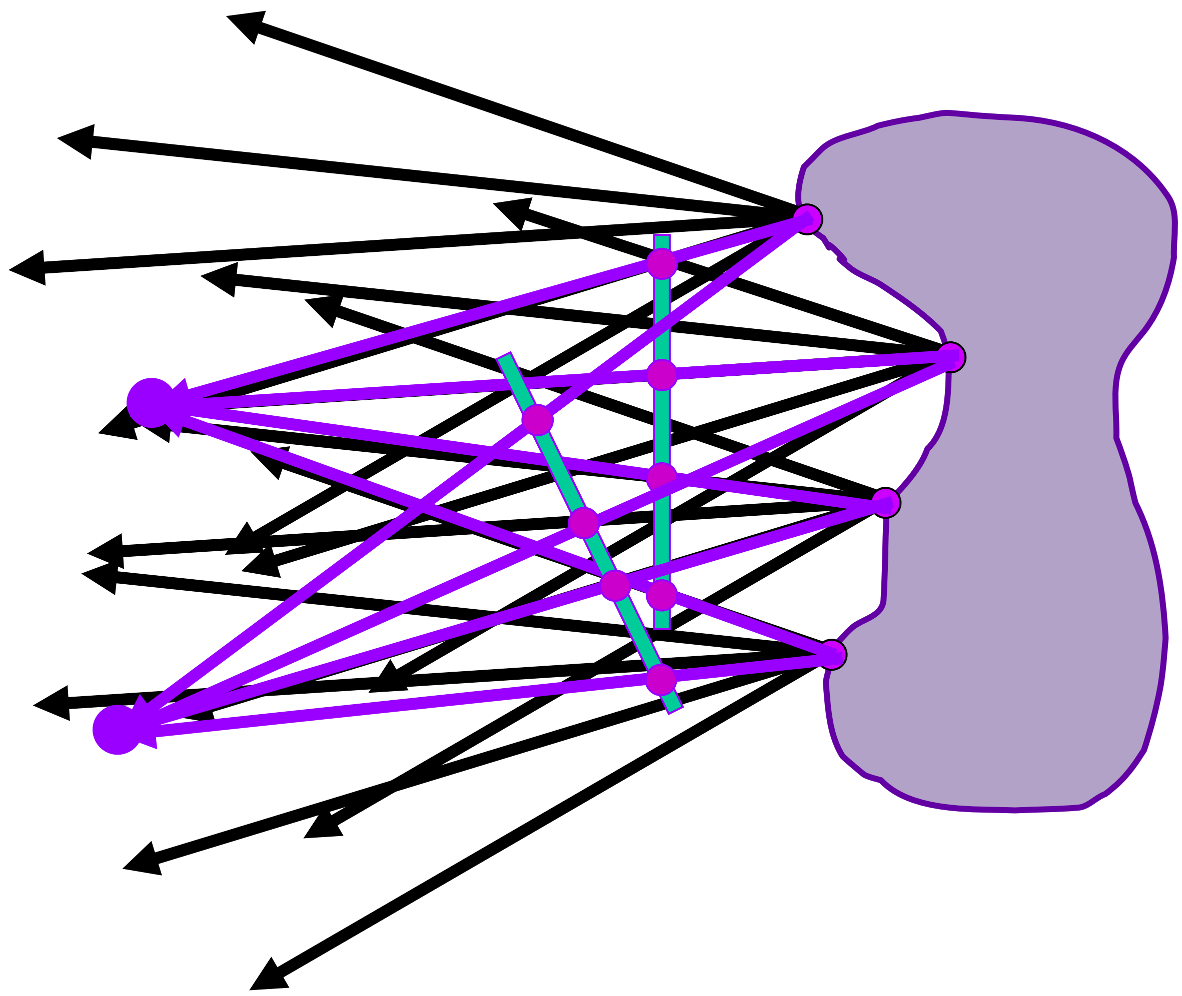

Slide by Rick Szeliski and Michael Cohen

Lumigraph / Lightfield

Outside convex space

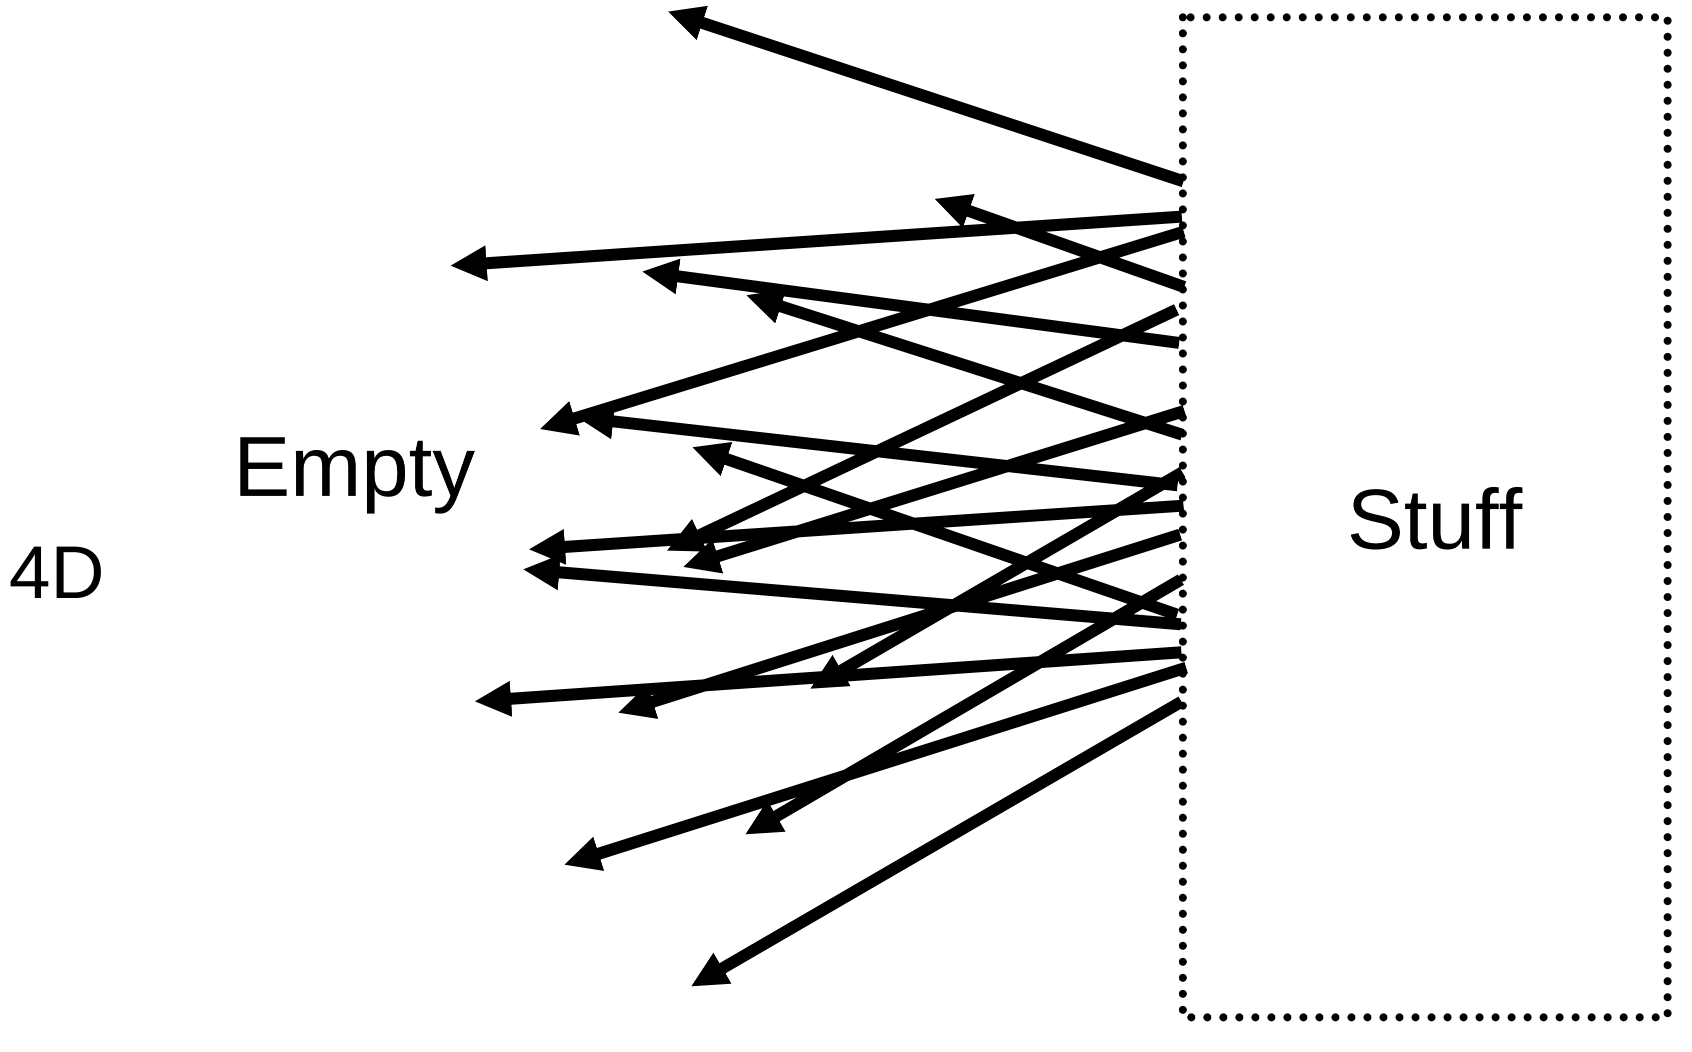

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Organization

2D position

2D direction

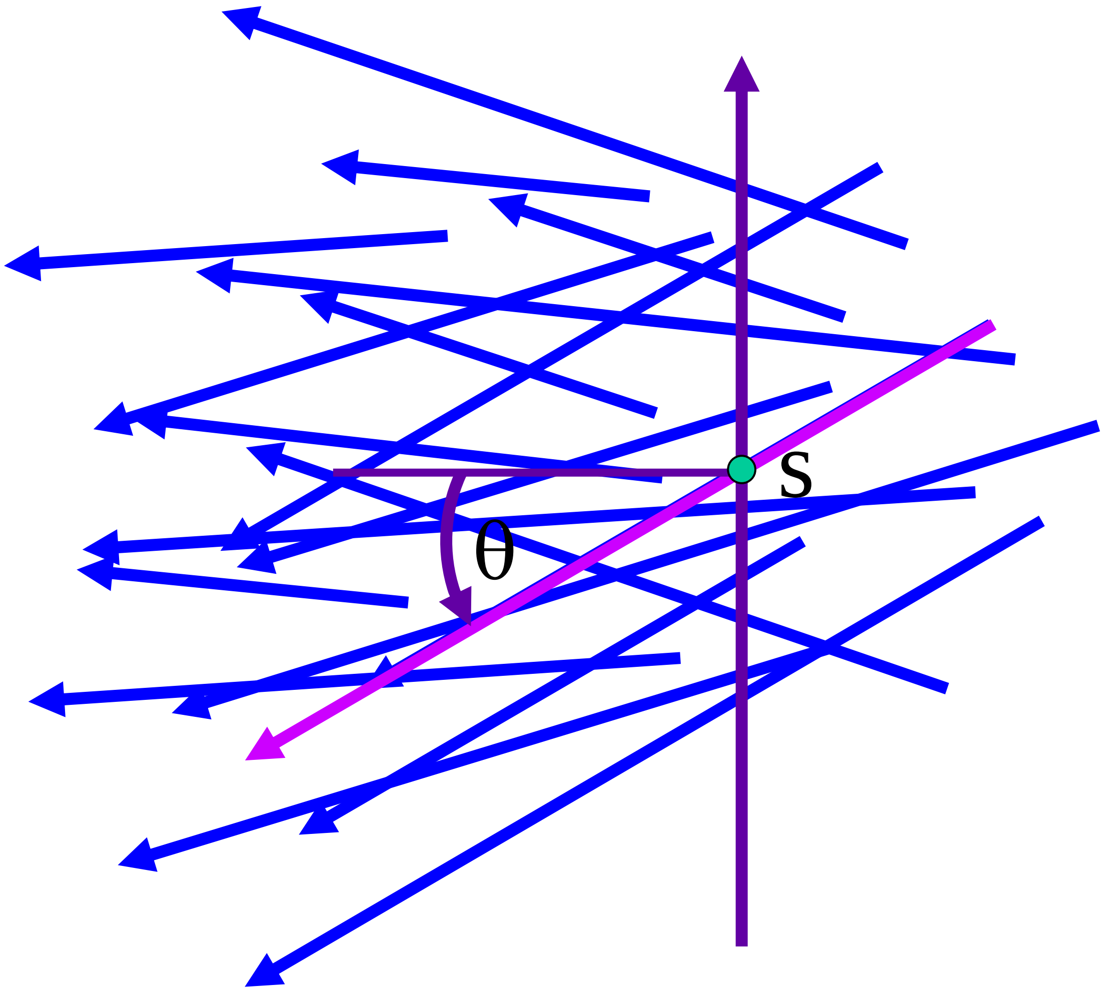

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Organization

2D position

2D position

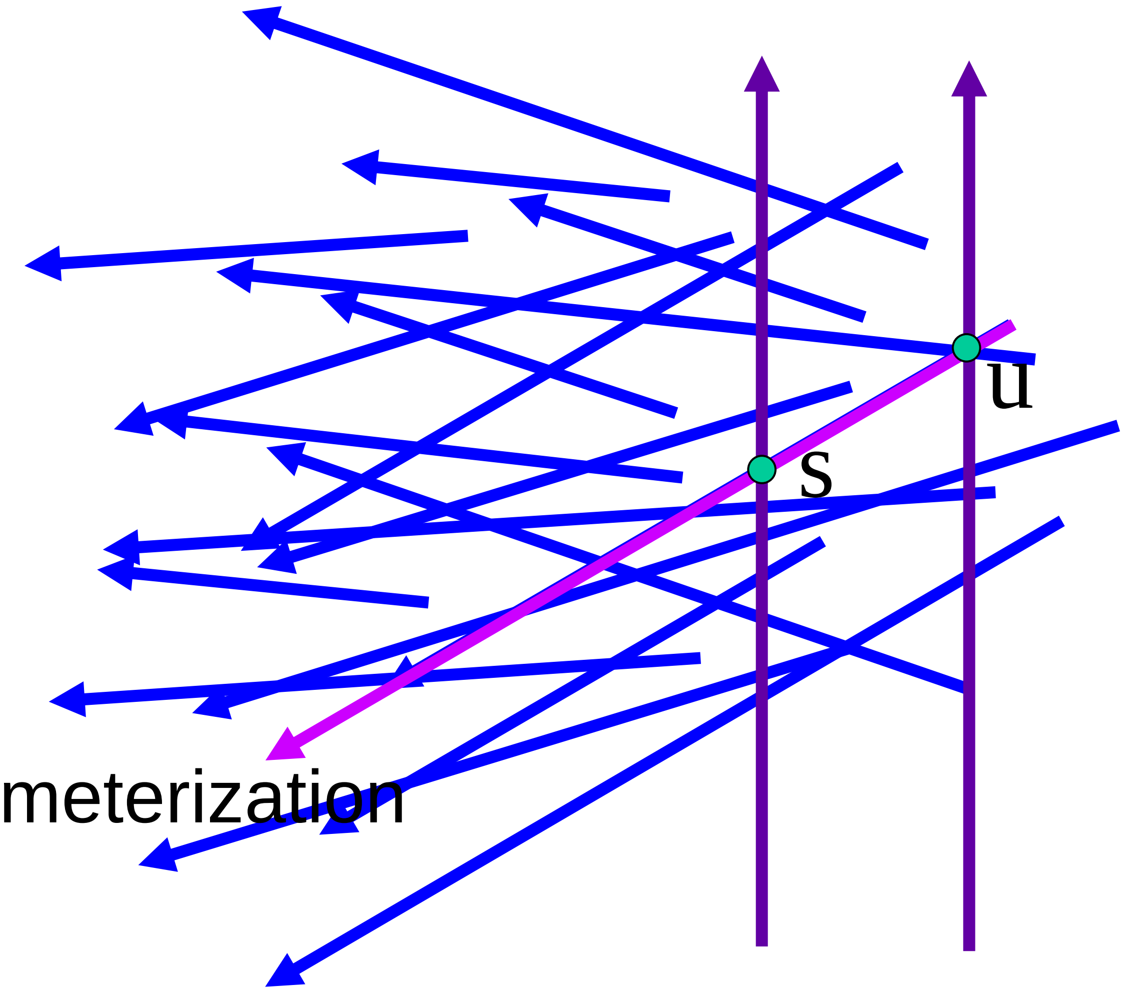

2 plane parameterization

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Organization

2D position

2D position

2 plane parameterization

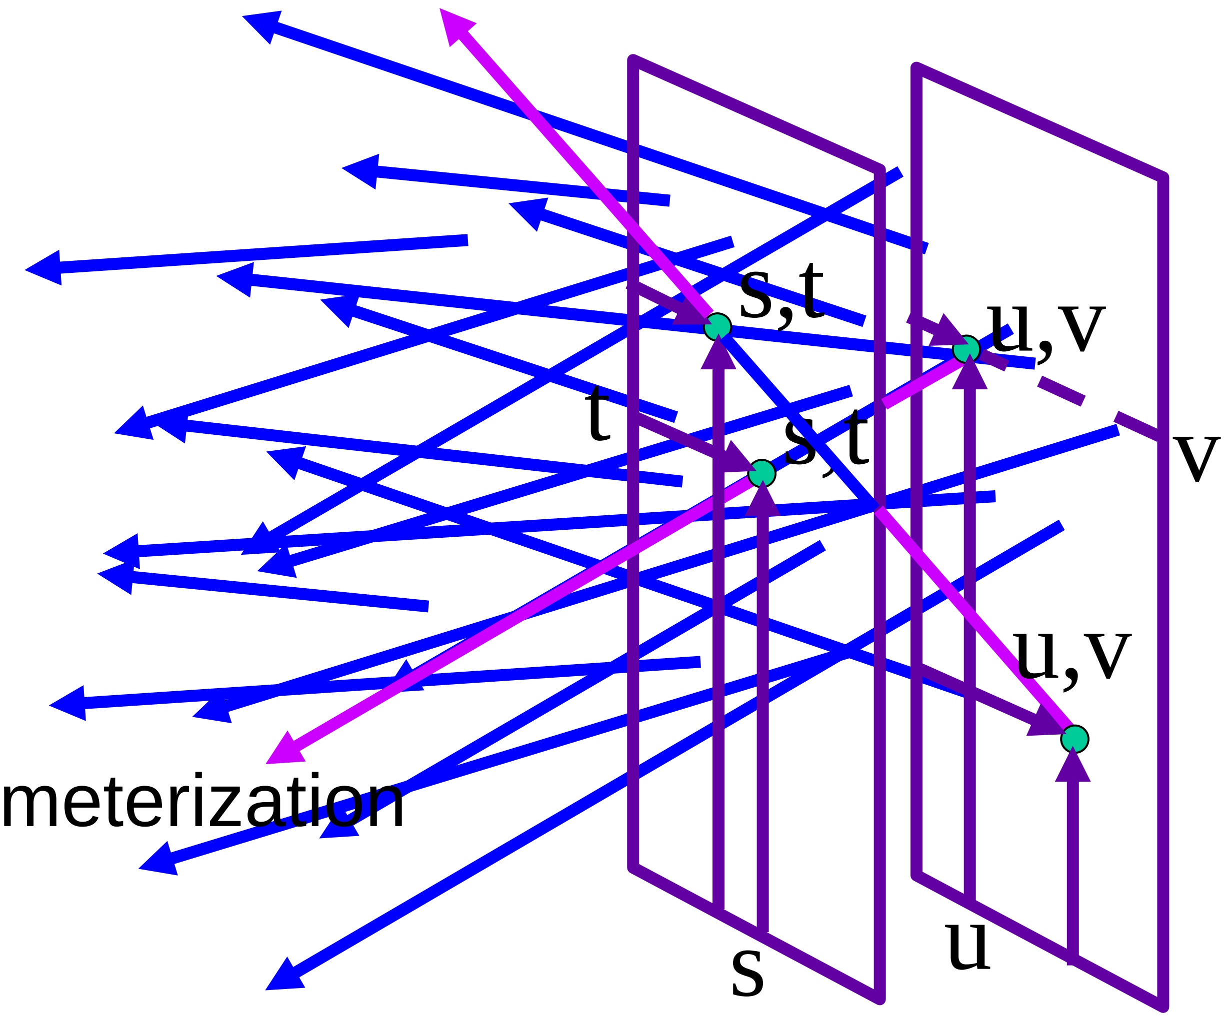

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Organization

Hold s,t constant

Let u,v vary

An image

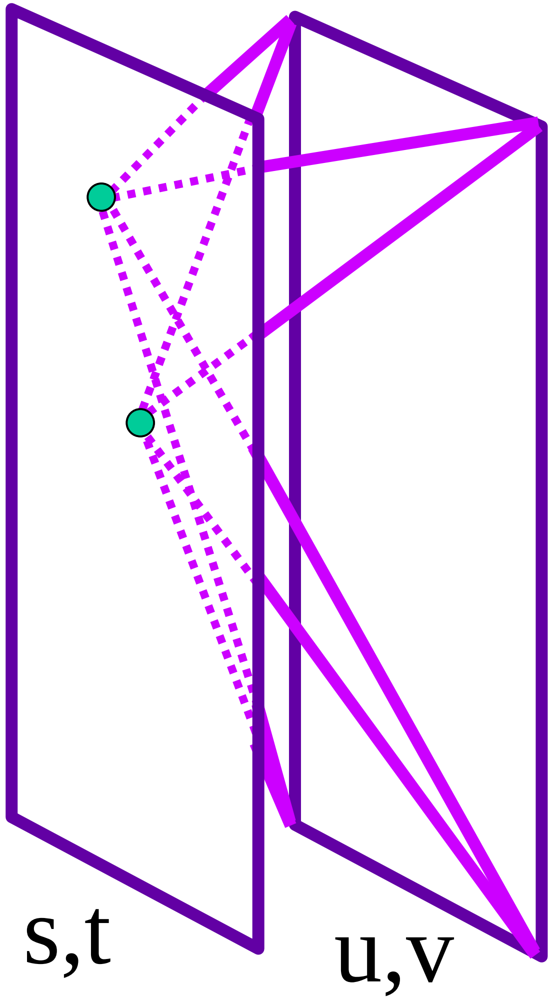

Slide by Rick Szeliski and Michael Cohen

Lumigraph / Lightfield

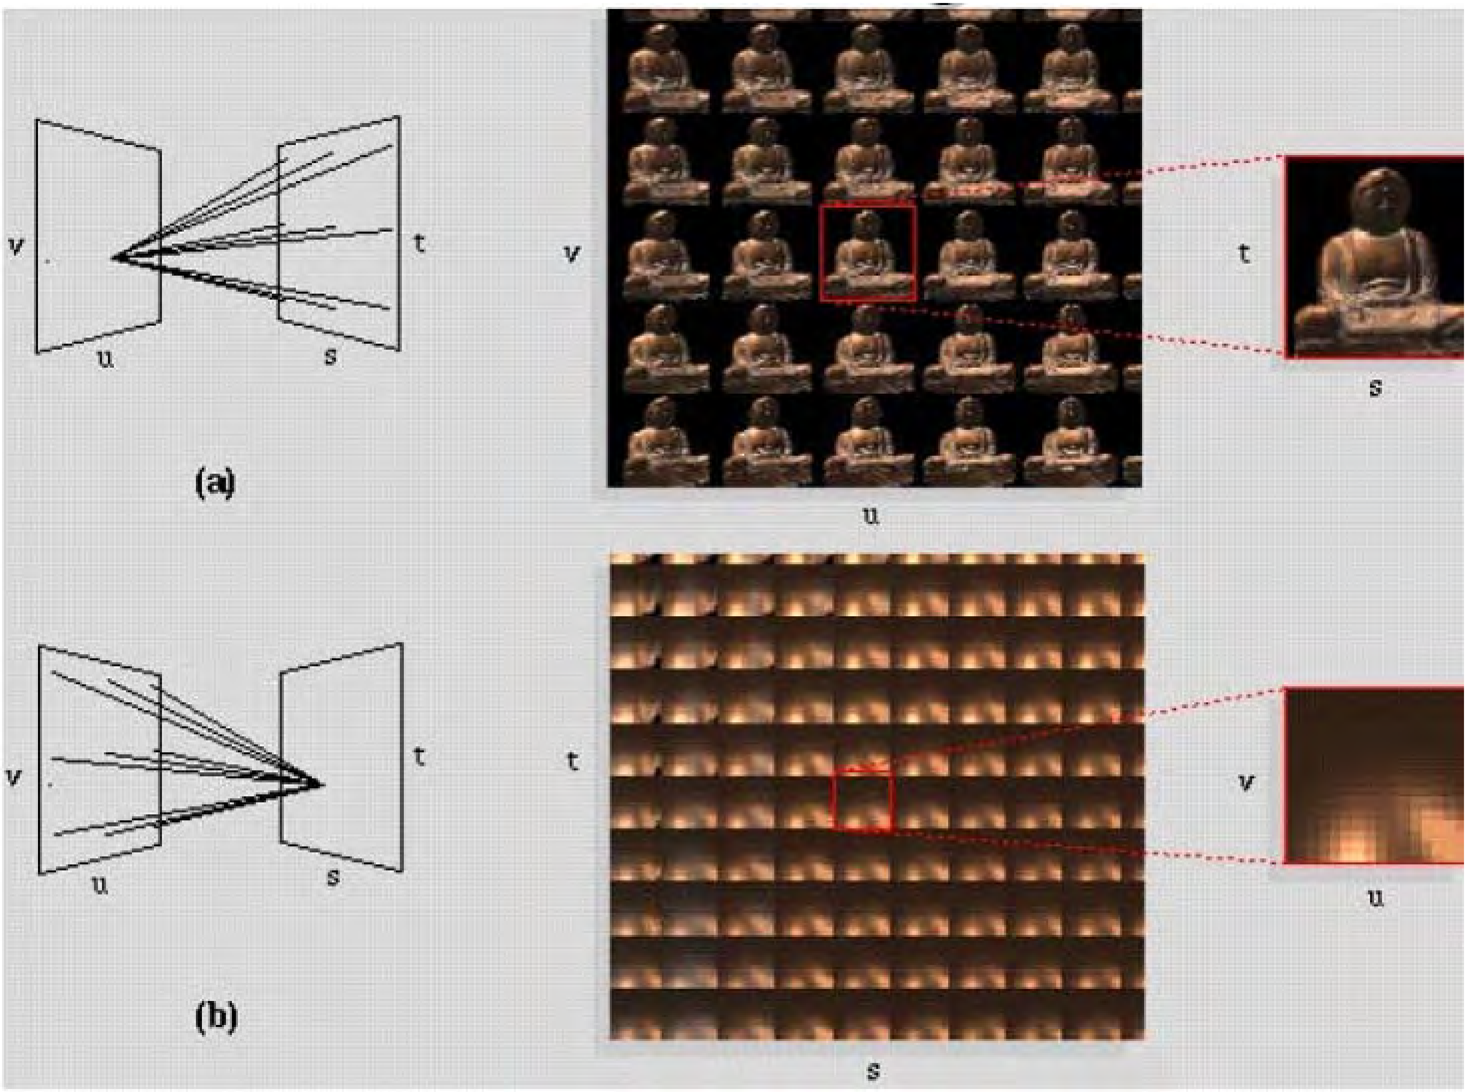

Lumigraph - Capture

Idea 1

- Move camera carefully over s,t plane
- Gantry
  – see Lightfield paper

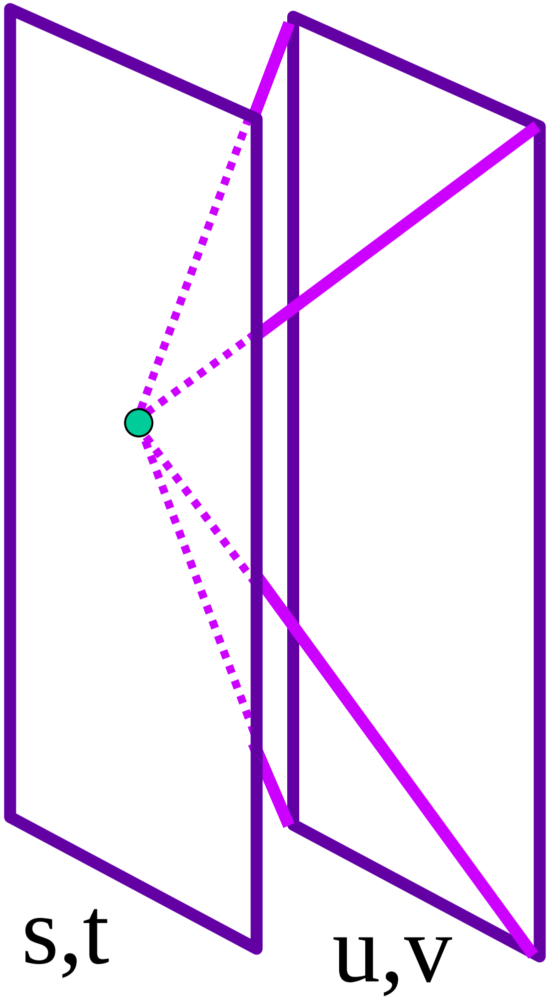

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Capture

Idea 2

- Move camera anywhere
- Rebinning
  - see Lumigraph paper

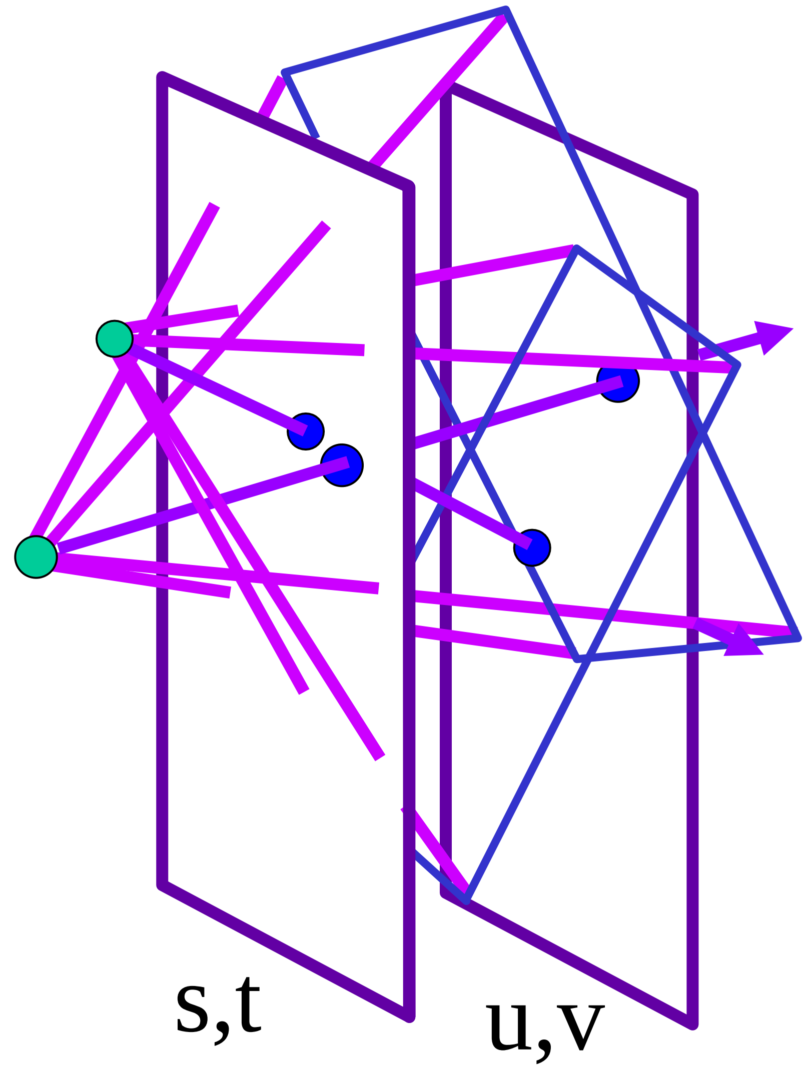

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Rendering

- For each output pixel
  - determine s,t,u,v
  - either
    - use closest discrete RGB
    - interpolate near values

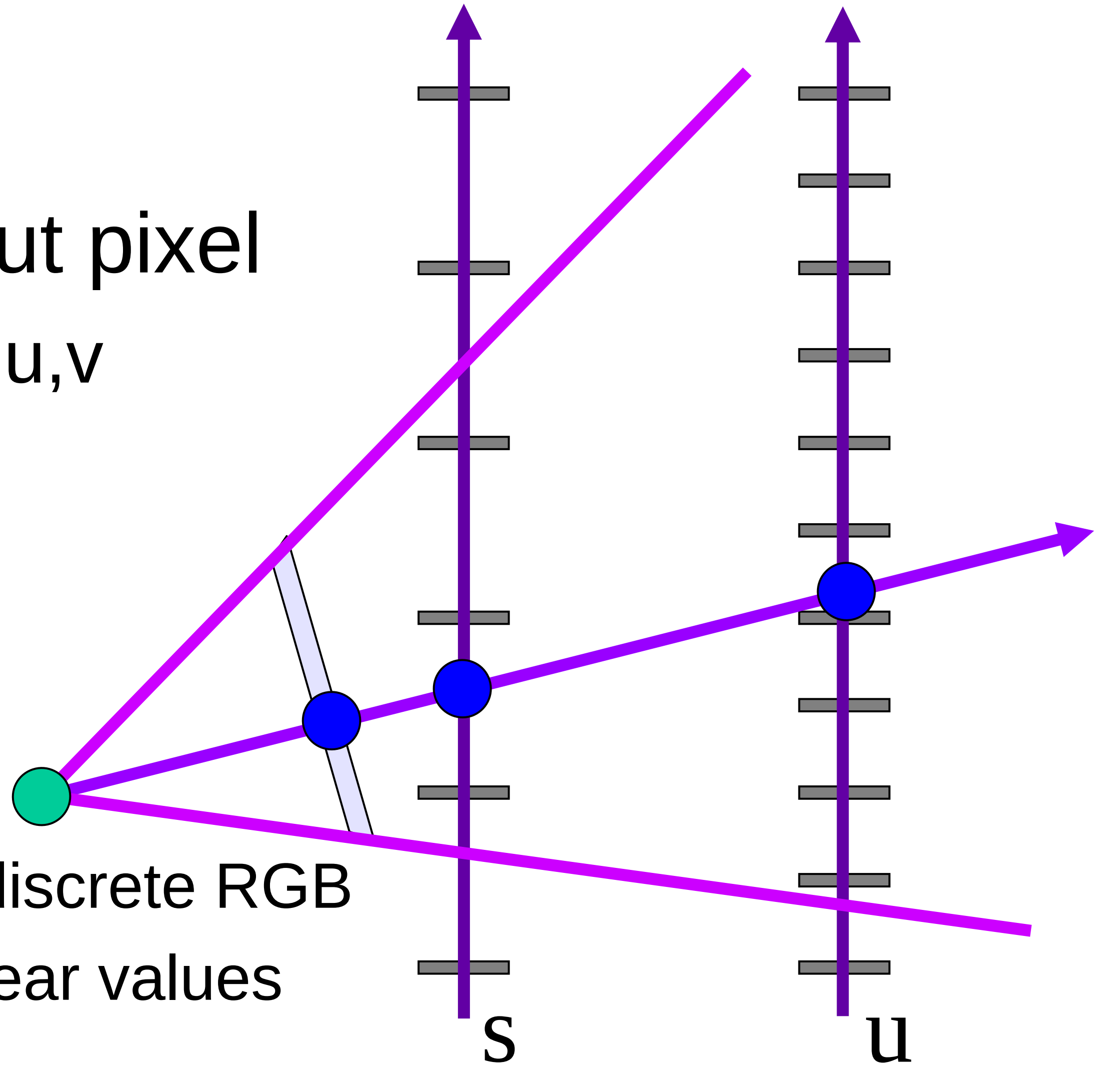

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Rendering

Nearest

- closest s
- closest u
- draw it

Blend 16 nearest

- quadrilinear interpolation

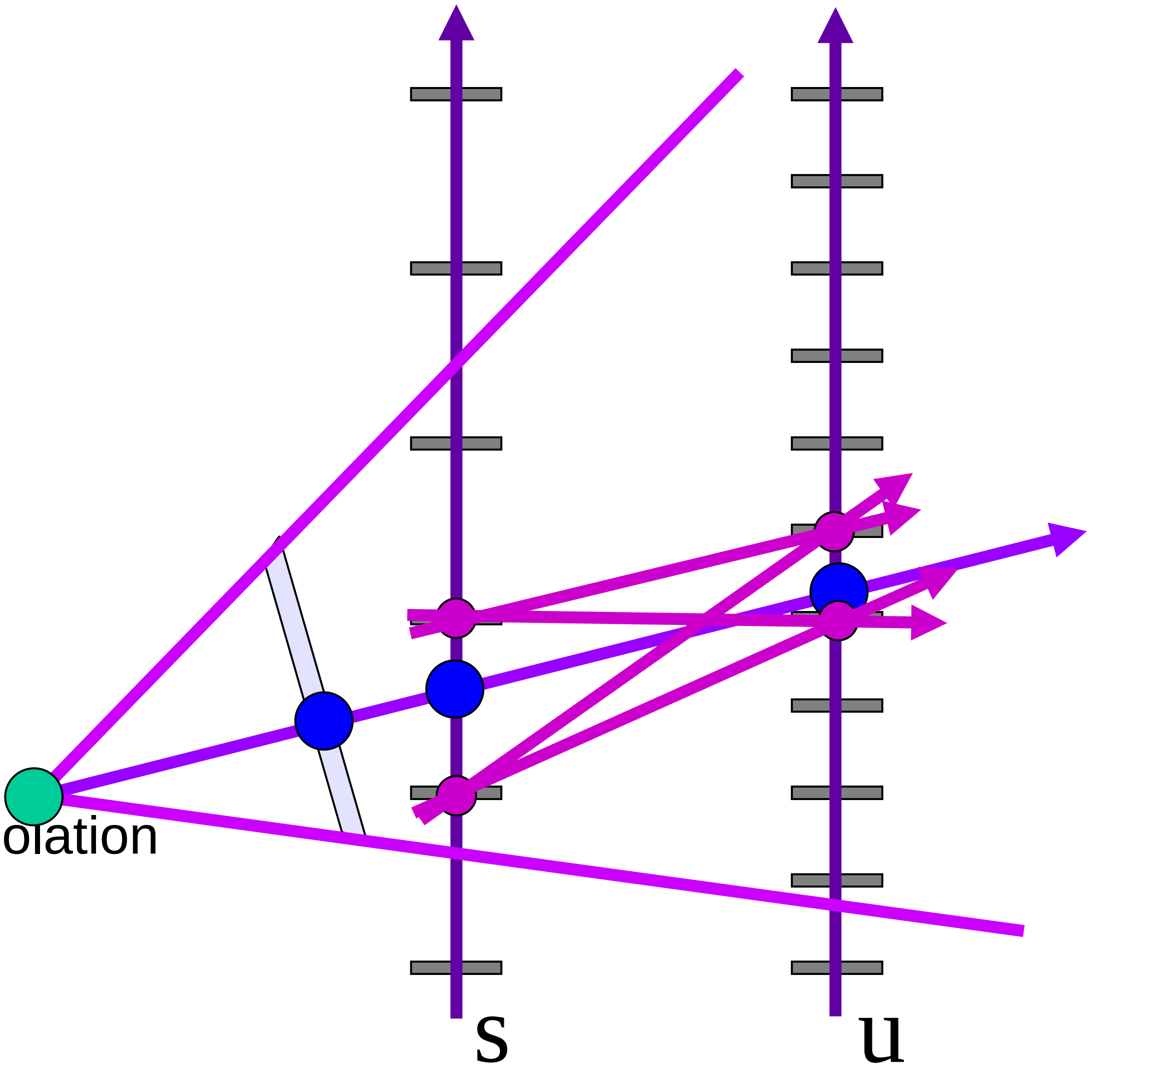

Slide by Rick Szeliski and Michael Cohen

Stanford multi-camera array

- 640 × 480 pixels × 30 fps × 128 cameras

- synchronized timing
- continuous streaming
- flexible arrangement

Light field photography using a handheld plenoptic camera

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

Conventional versus light field camera

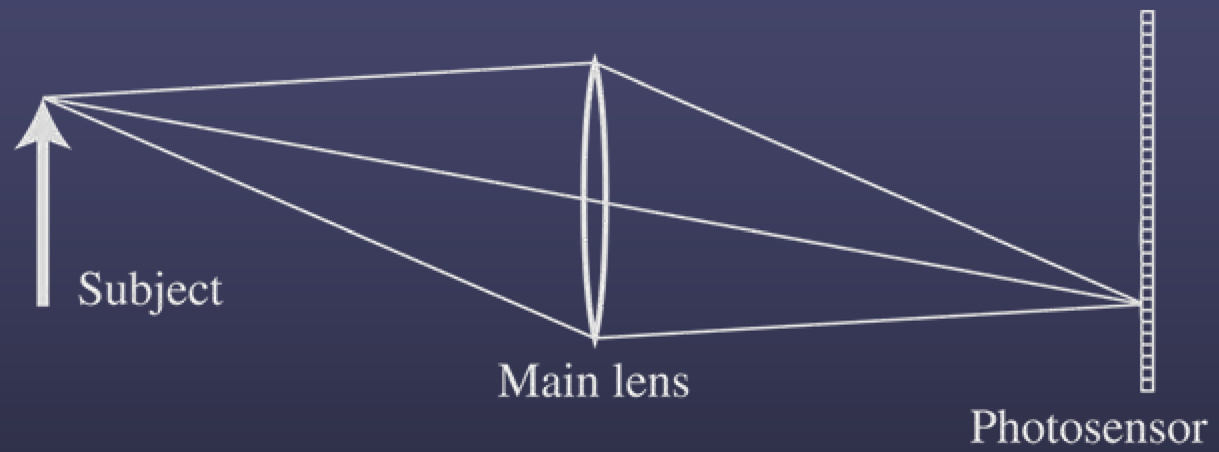

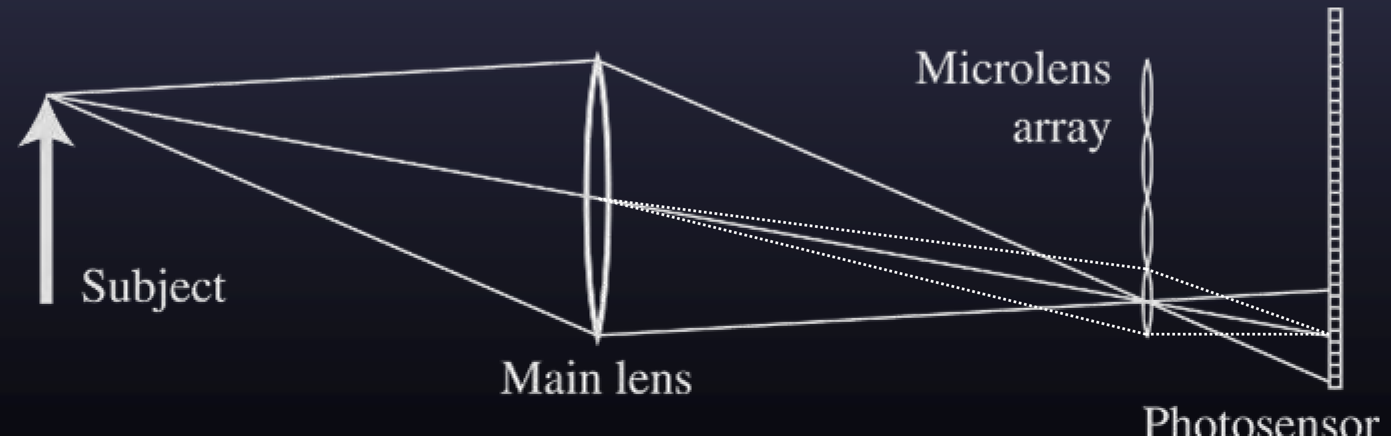

Conventional versus light field camera

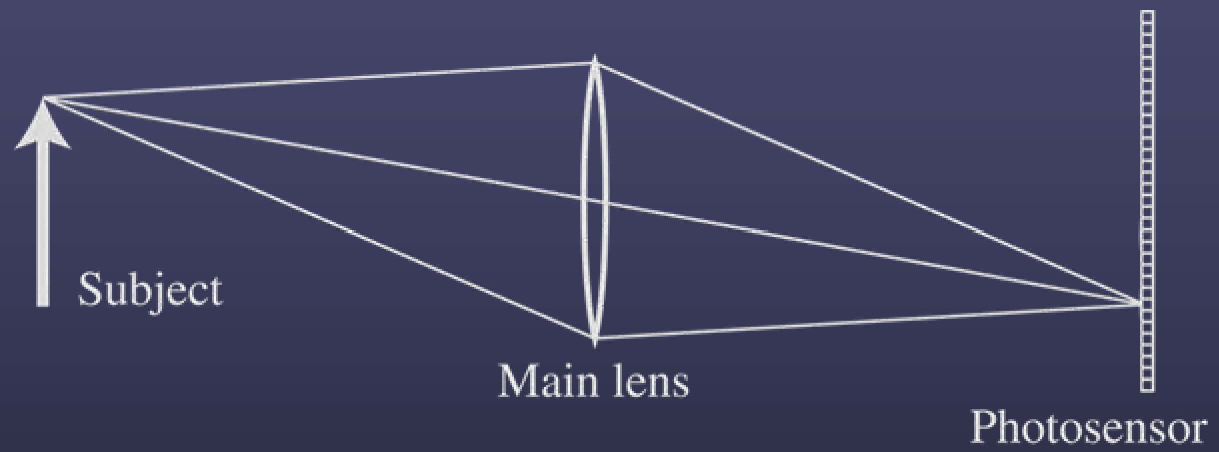

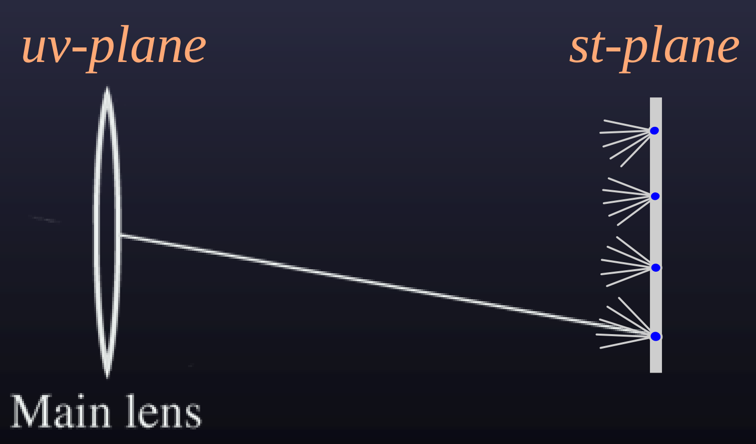

Prototype camera

Contax medium format camera

Kodak 16-megapixel sensor

Adaptive Optics microlens array

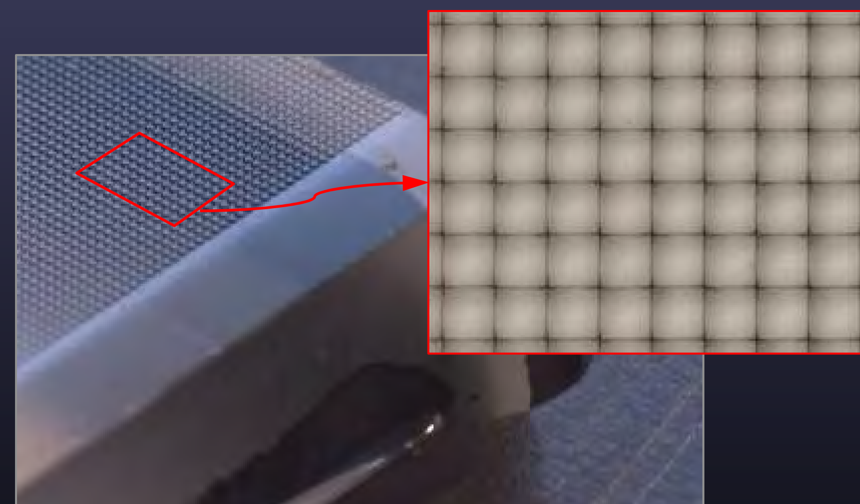

125μ square-sided microlenses

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

Digitally stopping-down

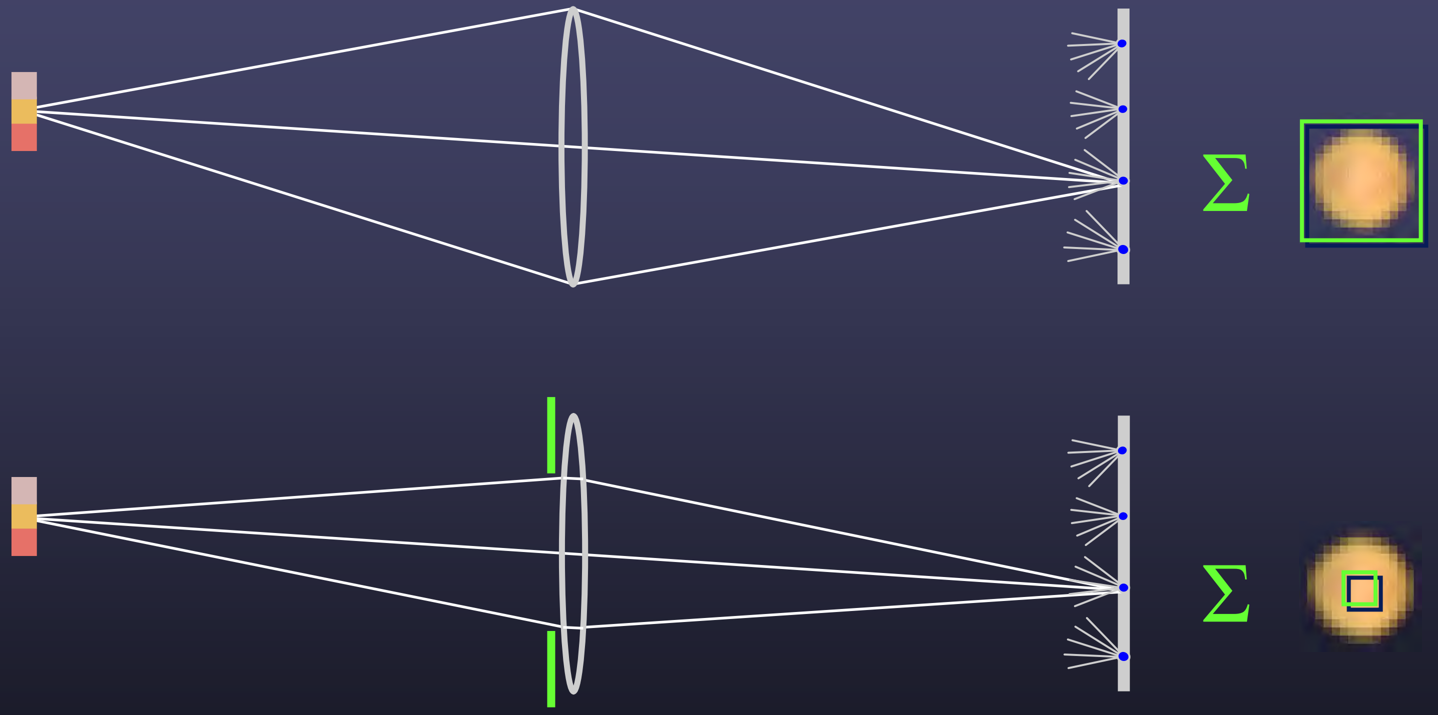

- stopping down = summing only the central portion of each microlens

Digital refocusing

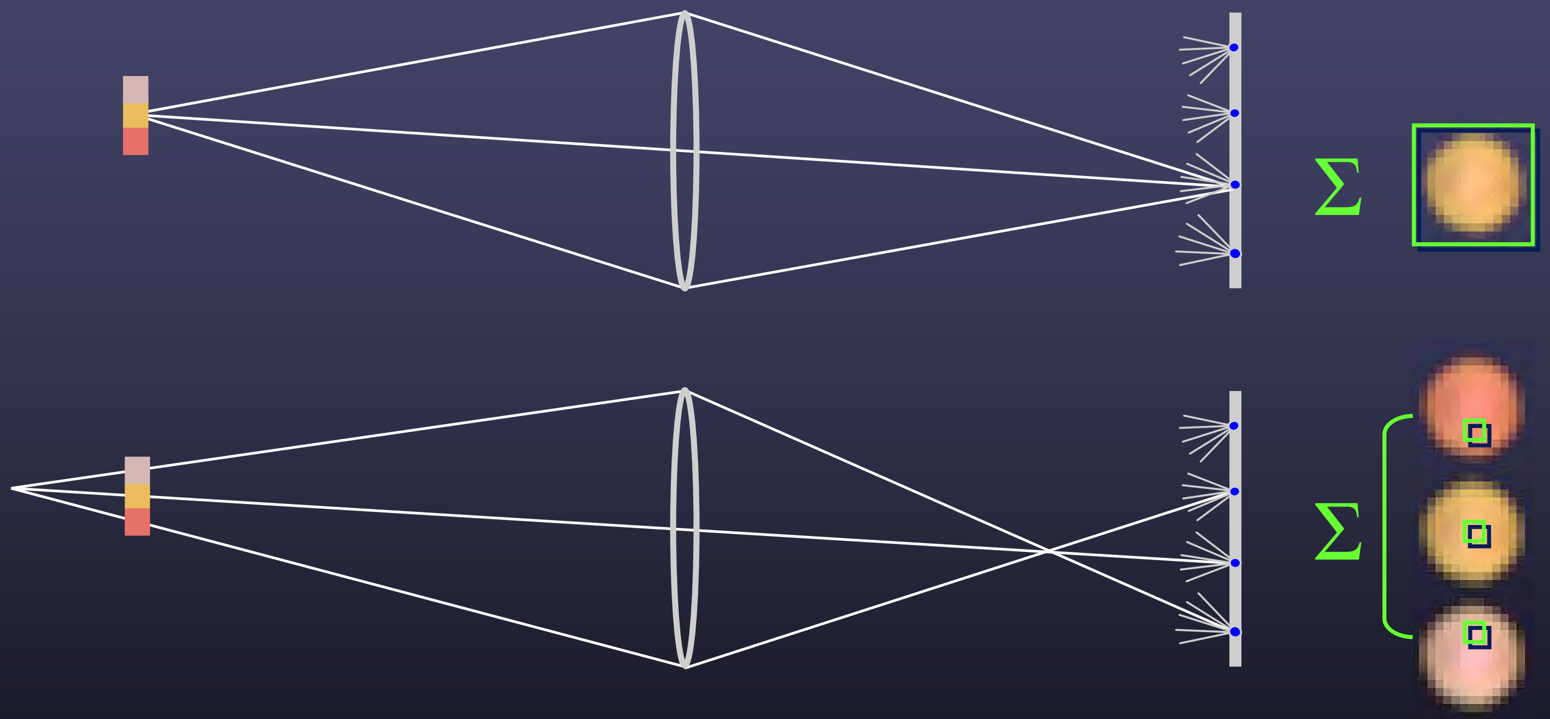

- refocusing = summing windows extracted from several microlenses

Example of digital refocusing

Digitally moving the observer

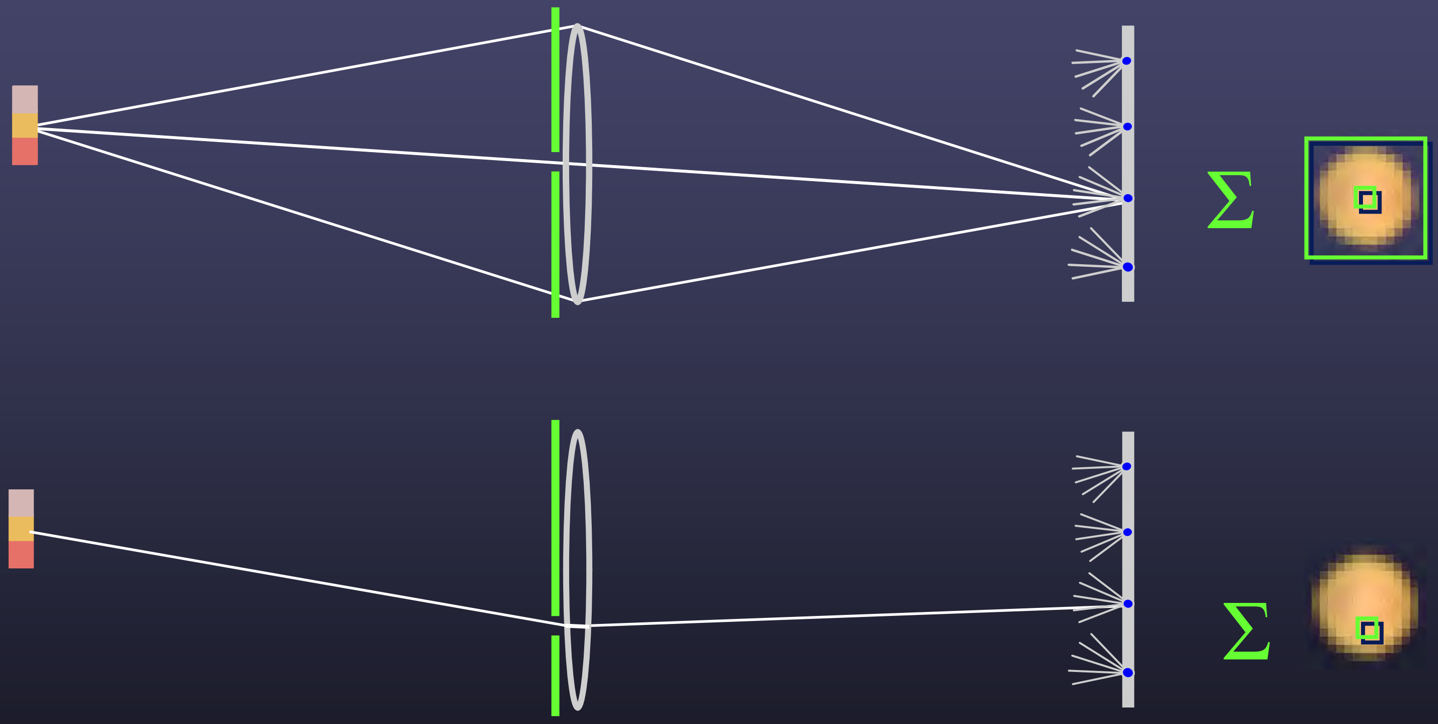

- moving the observer = moving the window we extract from the microlenses

Example of moving the observer

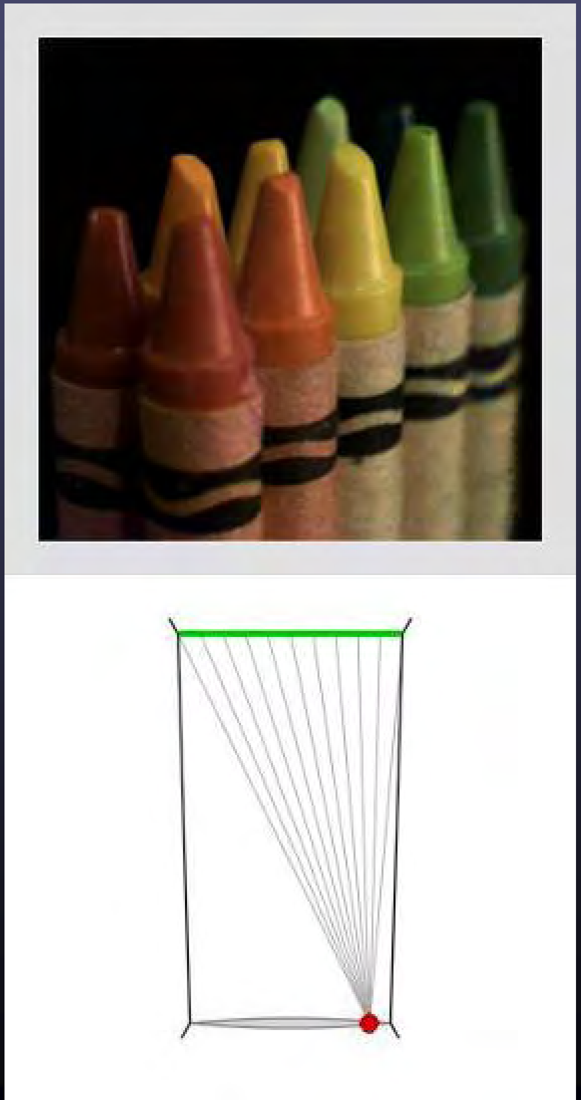

Moving backward and forward

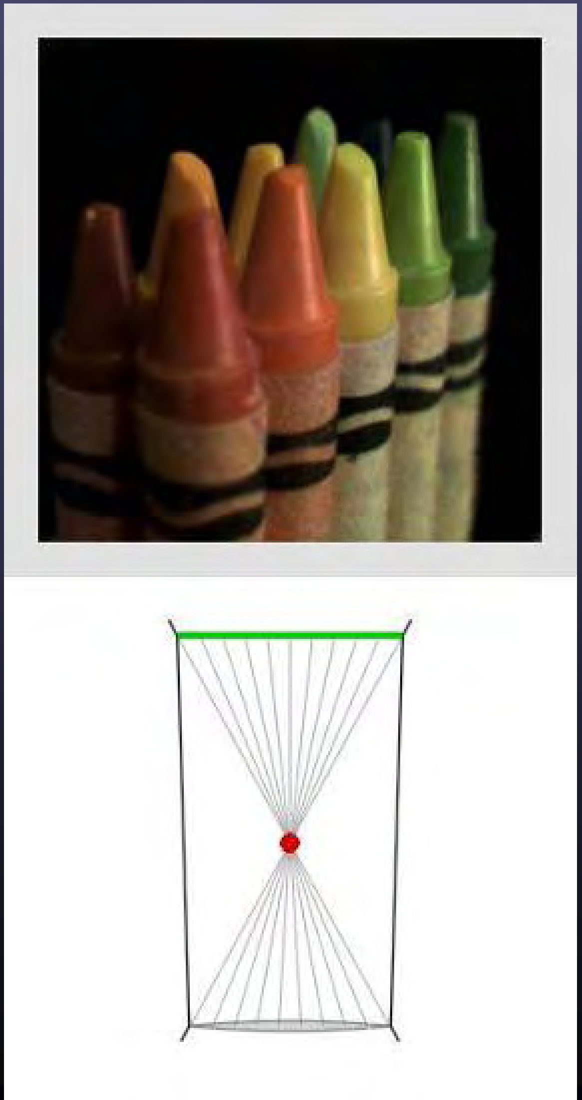

3D Lumigraph

One row of s,t plane

- i.e., hold t constant

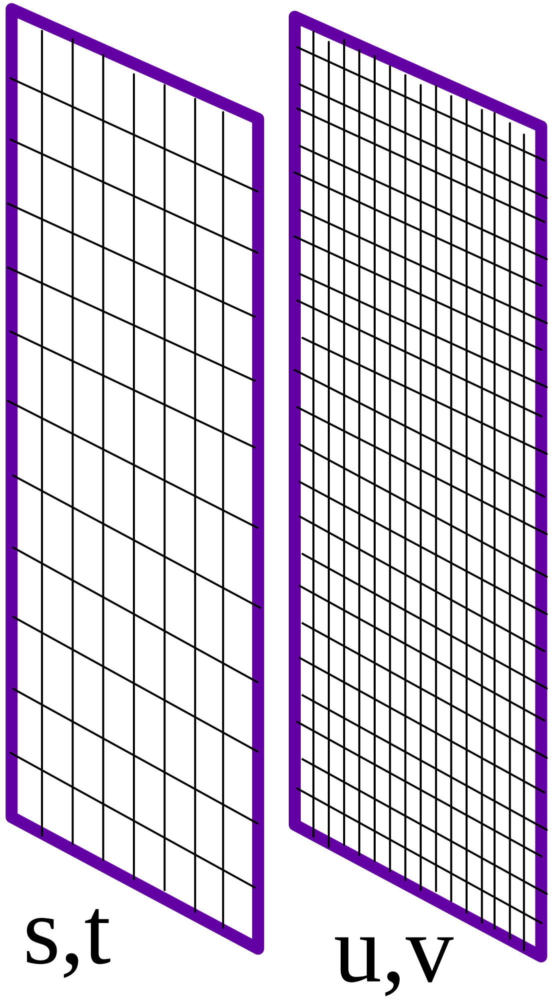

3D Lumigraph

One row of s,t plane

- i.e., hold t constant
- thus s,u,v
- a “row of images”

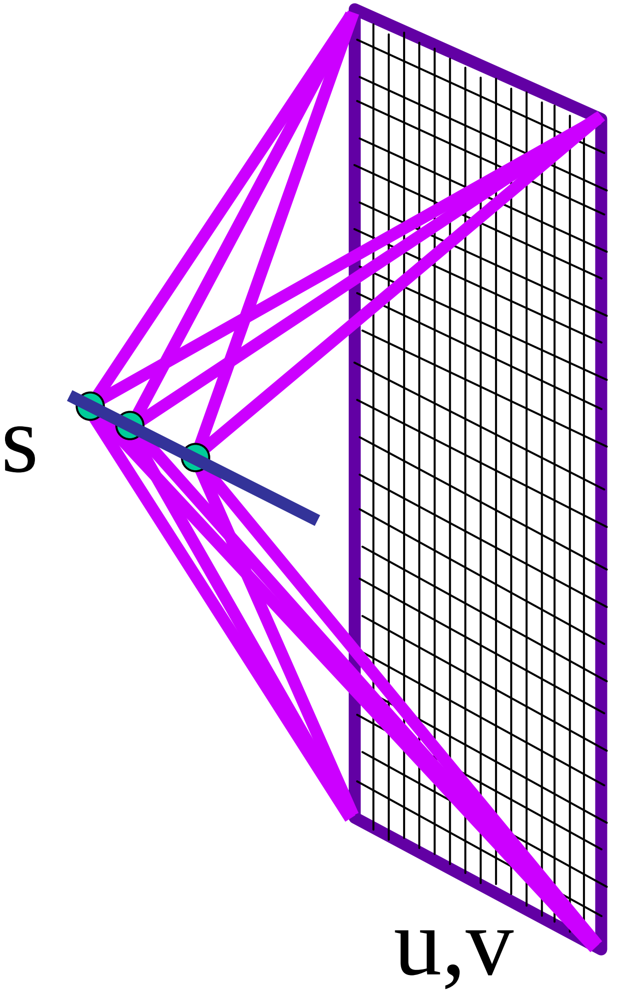

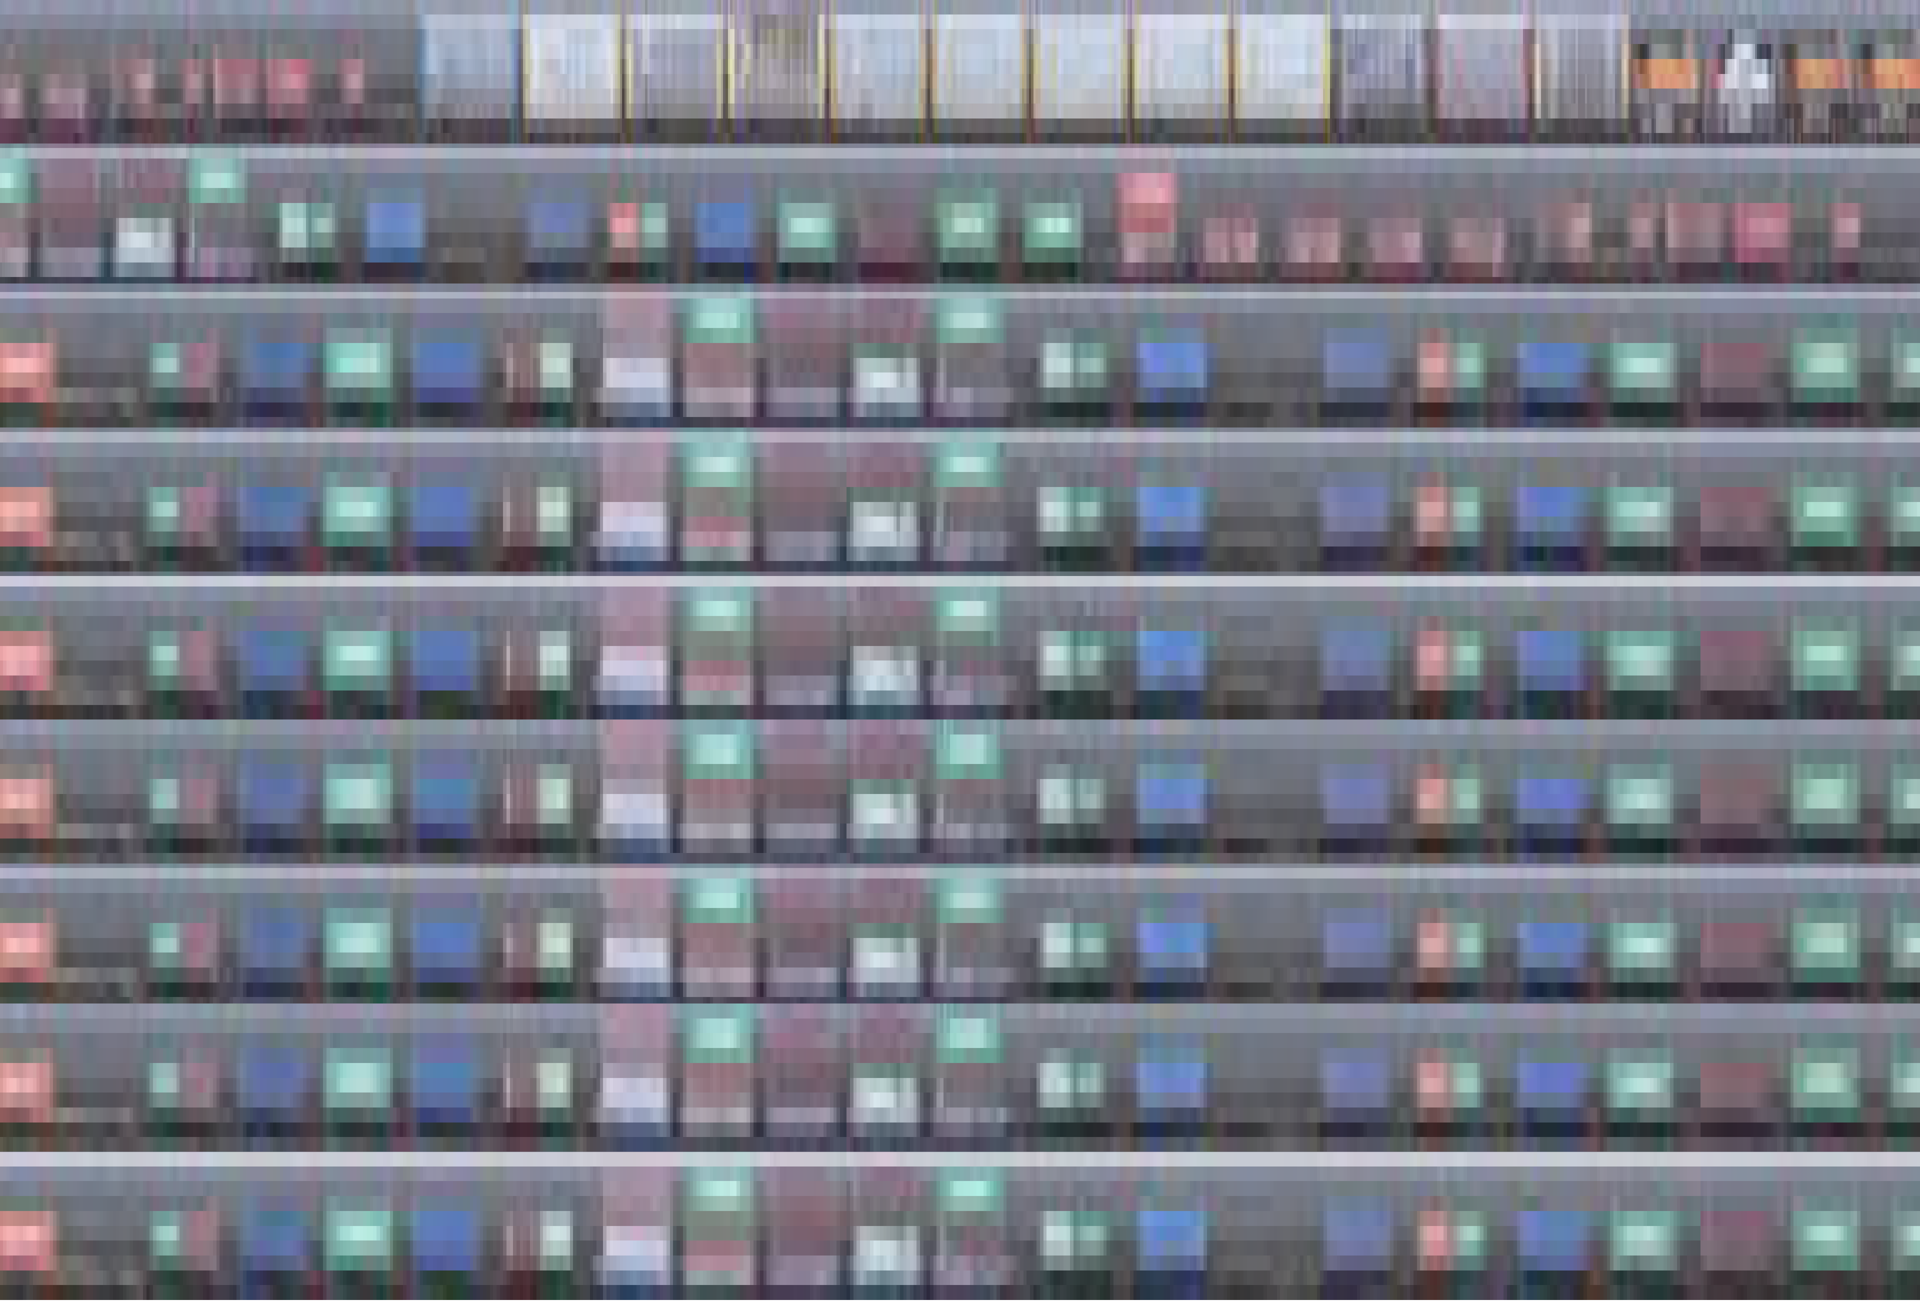

P(x,t) by David Dewey

2D: Image

What is an image?

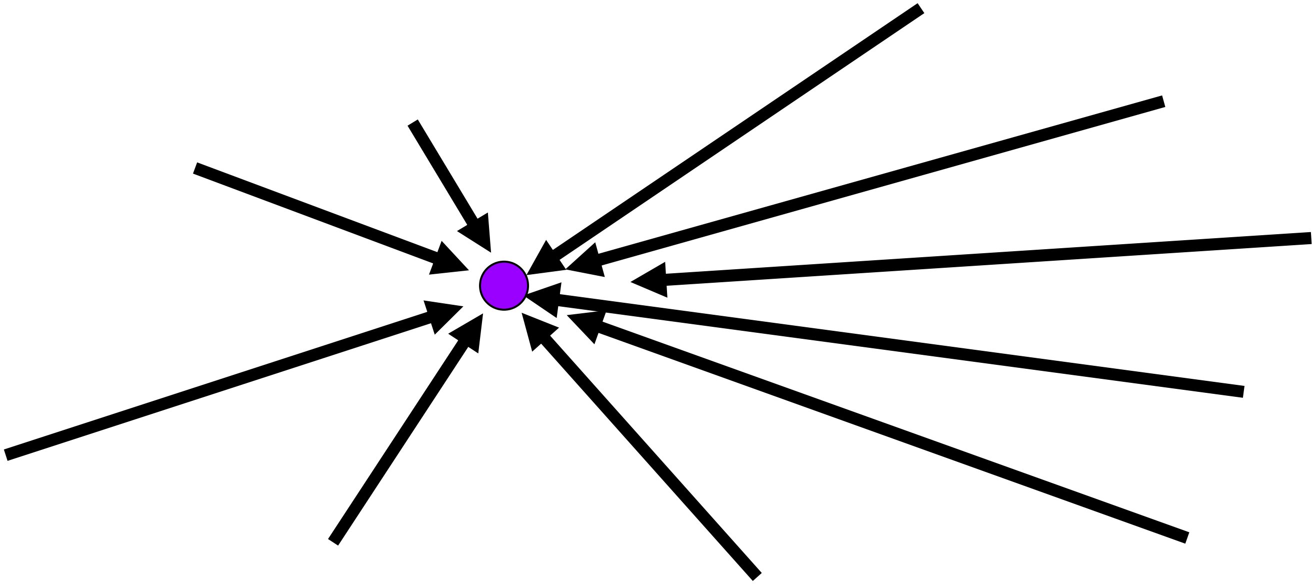

All rays through a point

- Panorama?

Image

Image plane

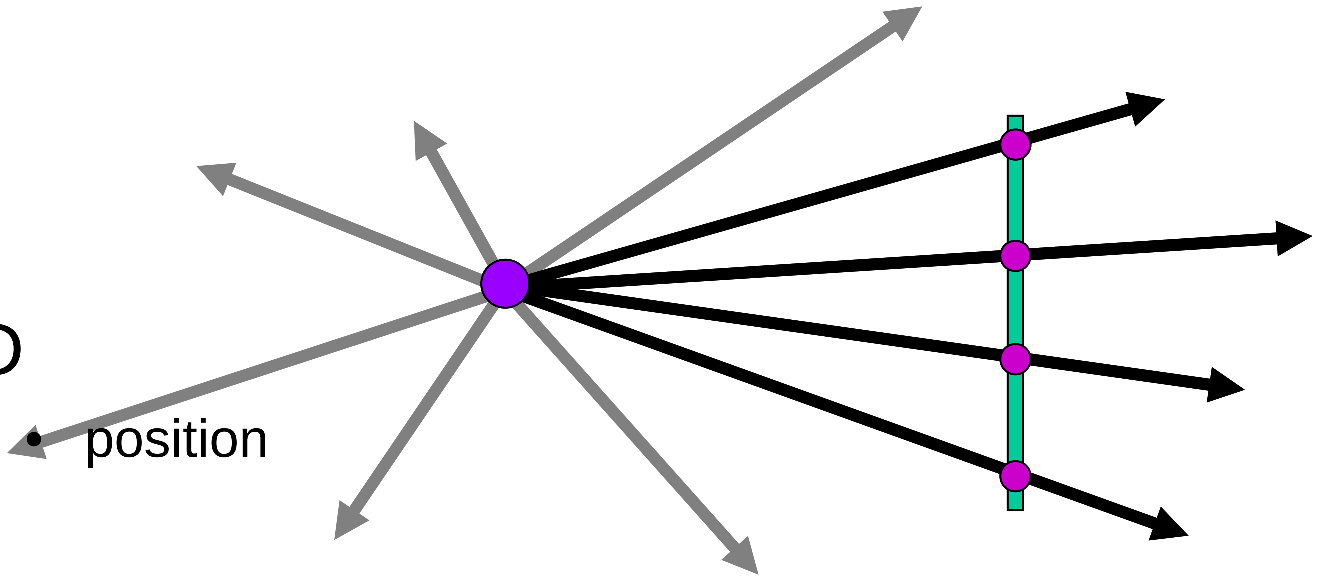

2D

- position

Spherical Panorama

See also: 2003 New Years Eve

http://www.panoramas.dk/fullscreen3/f1.html

All light rays through a point form a ponorama

Totally captured in a 2D array -- $P(\theta,\phi)$

Where is the geometry???

Other ways to sample Plenoptic Function

Moving in time:

- Spatio-temporal volume: $P(\theta,\phi,t)$
- Useful to study temporal changes
- Long an interest of artists:



Claude Monet, Haystacks studies

Space-time images

Other ways to slice the plenoptic function…

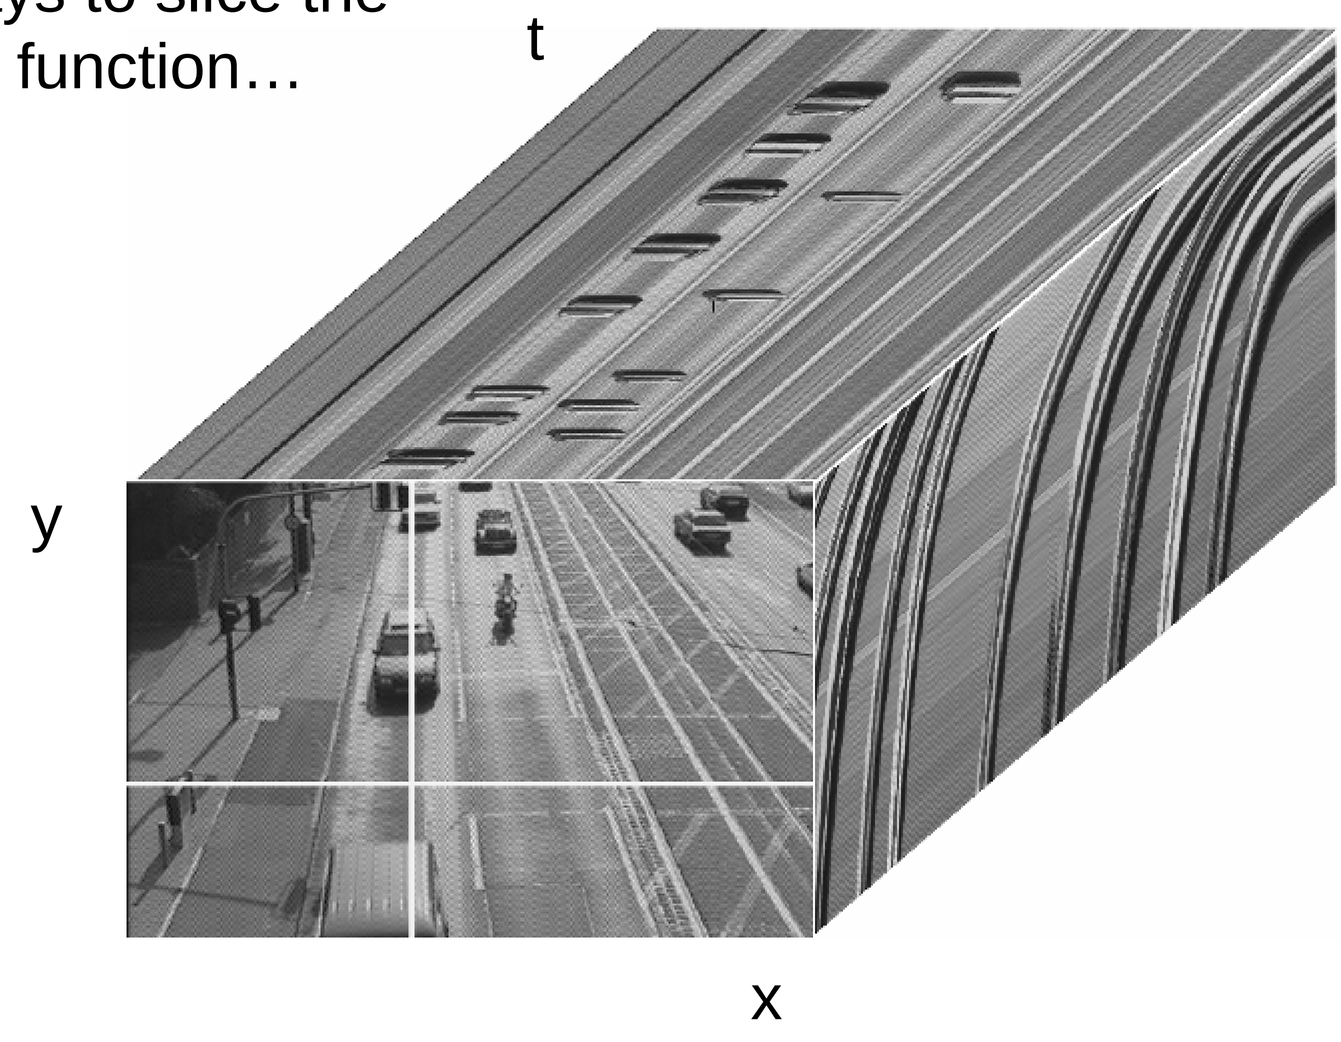

The “Theatre Workshop” Metaphor

(Adelson & Pentland,1996)

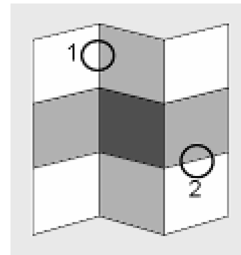

desired image

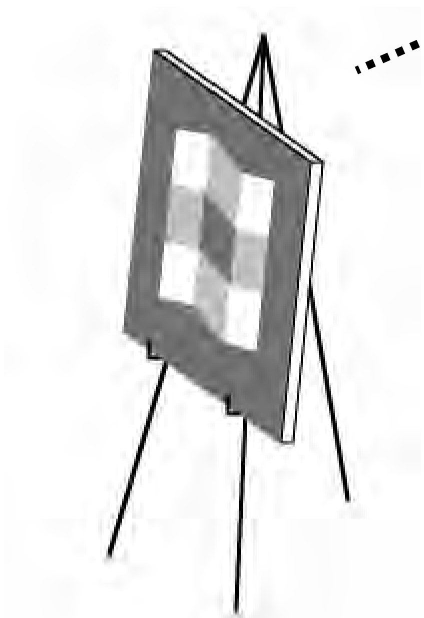

Painter

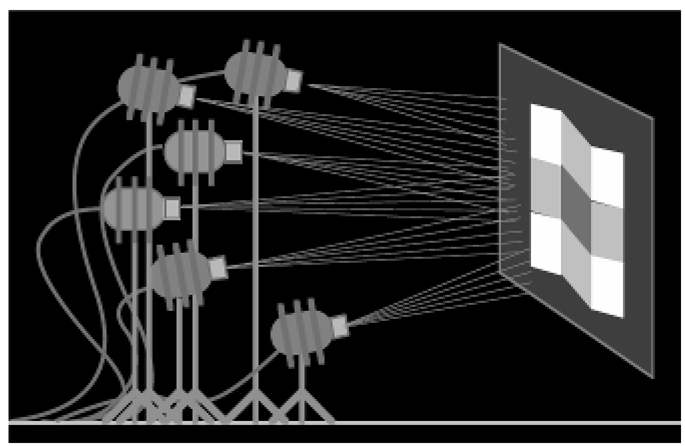

Lighting Designer

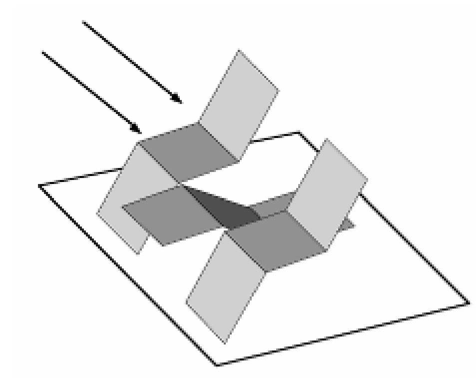

Sheet-metal worker

Painter (texture)

Lighting Designer (environment maps)

Show Naimark SF MOMA video

http://www.debevec.org/Naimark/naimark-displacements.mov

Sheet-metal Worker (geometry)

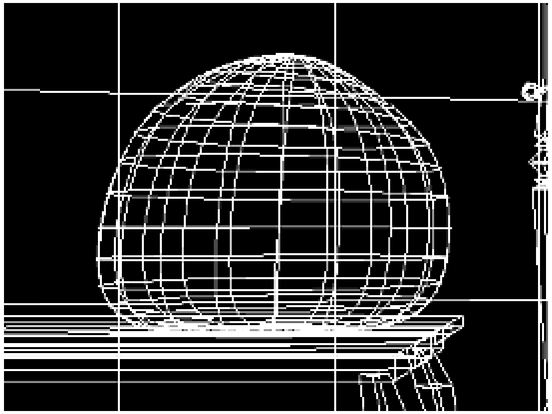

Let surface normals do all the work!

… working together

clever Italians

Want to minimize cost

Each one does what’s easiest for him

- Geometry – big things
- Images – detail
- Lighting – illumination effects

Three approaches

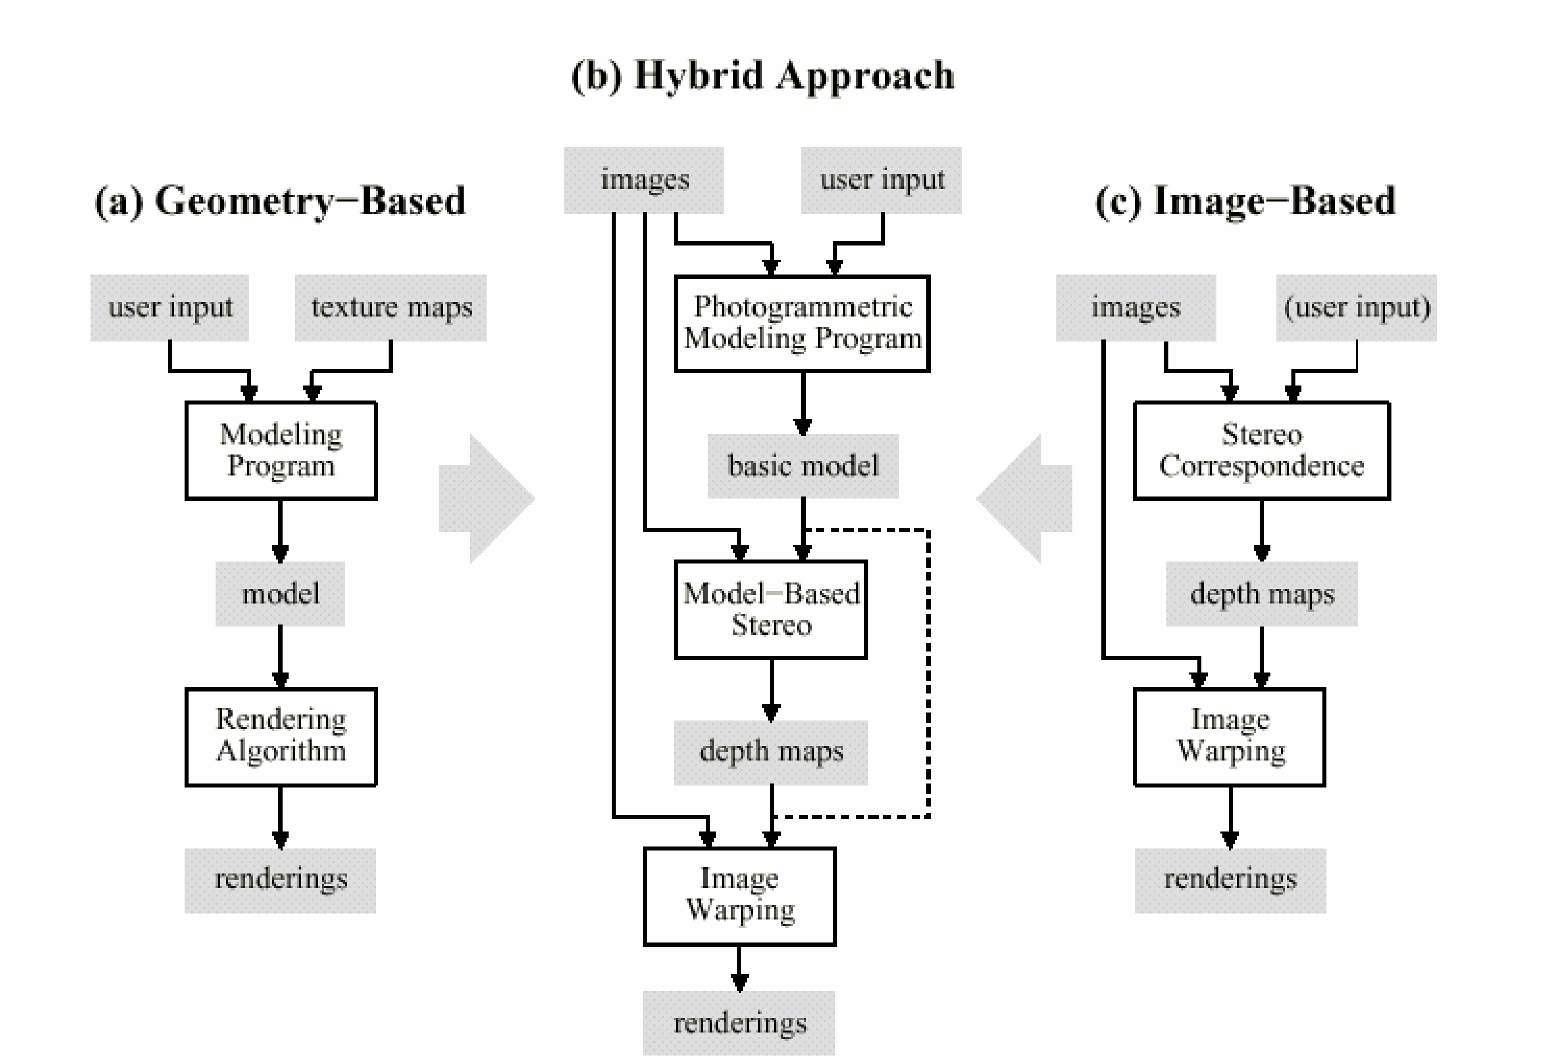

Outline of a simple algorithm (1)

- Based on constraints
- Input to the algorithm (1): two images

Outline of a simple algorithm (2)

- Input to the algorithm (2):
  User select edges and corners

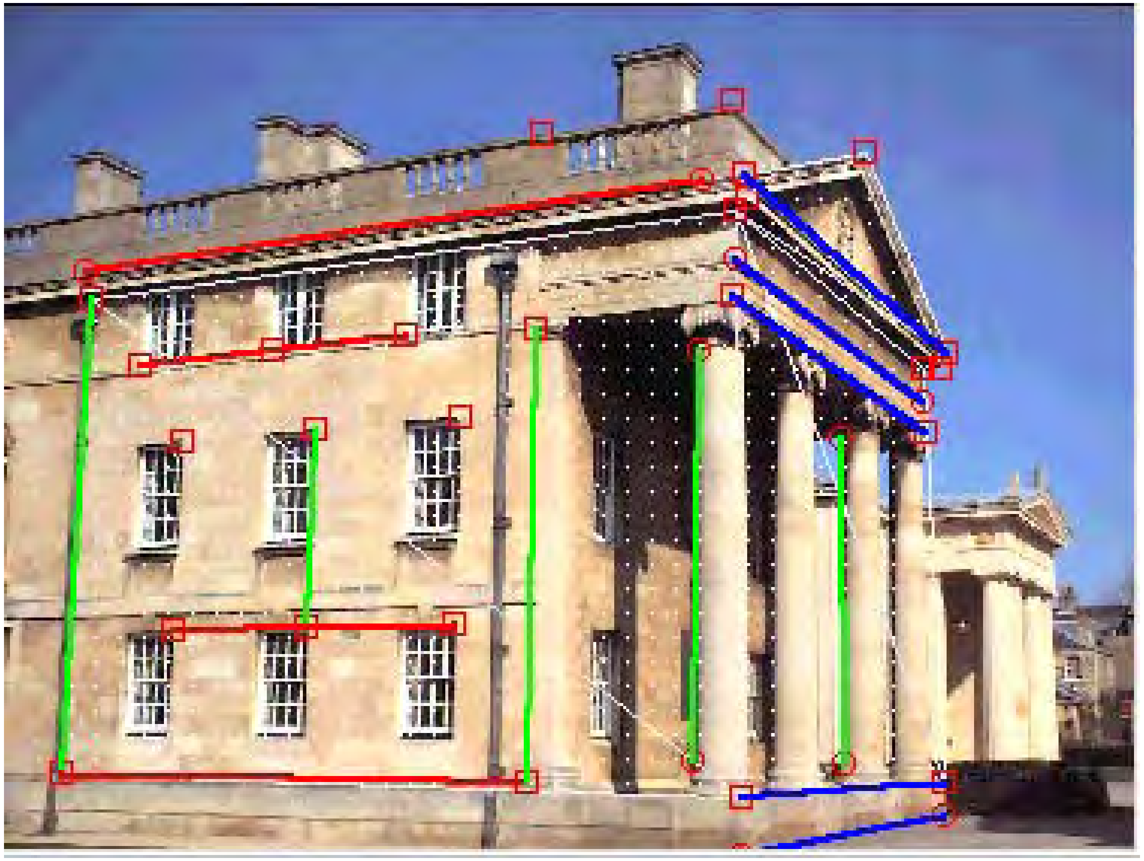

Outline of a simple algorithm (3)

- Camera Position and Orientation

  Determine the position and orientation of camera

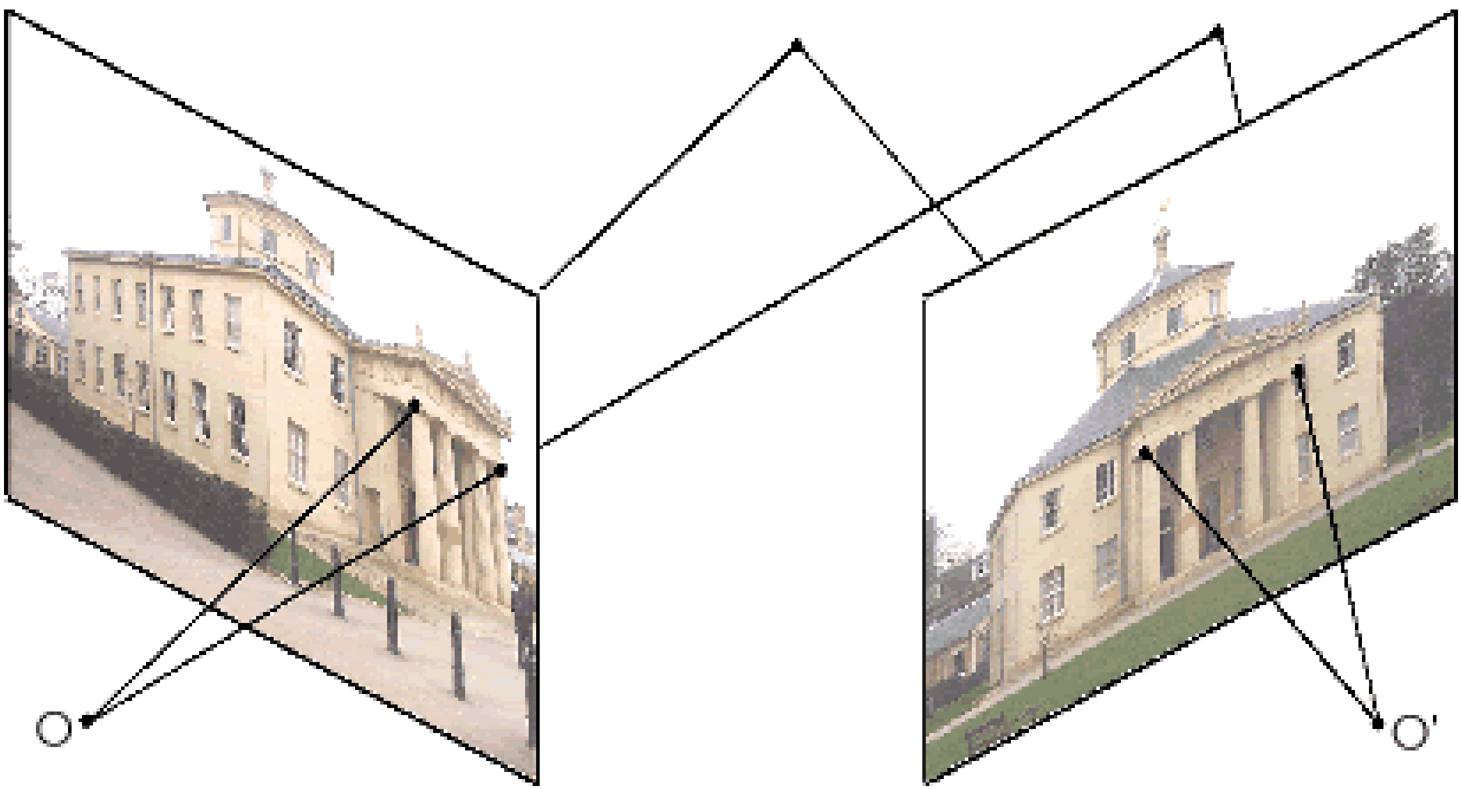

Outline of a simple algorithm (4)

- Computing projection matrix and Reconstruction

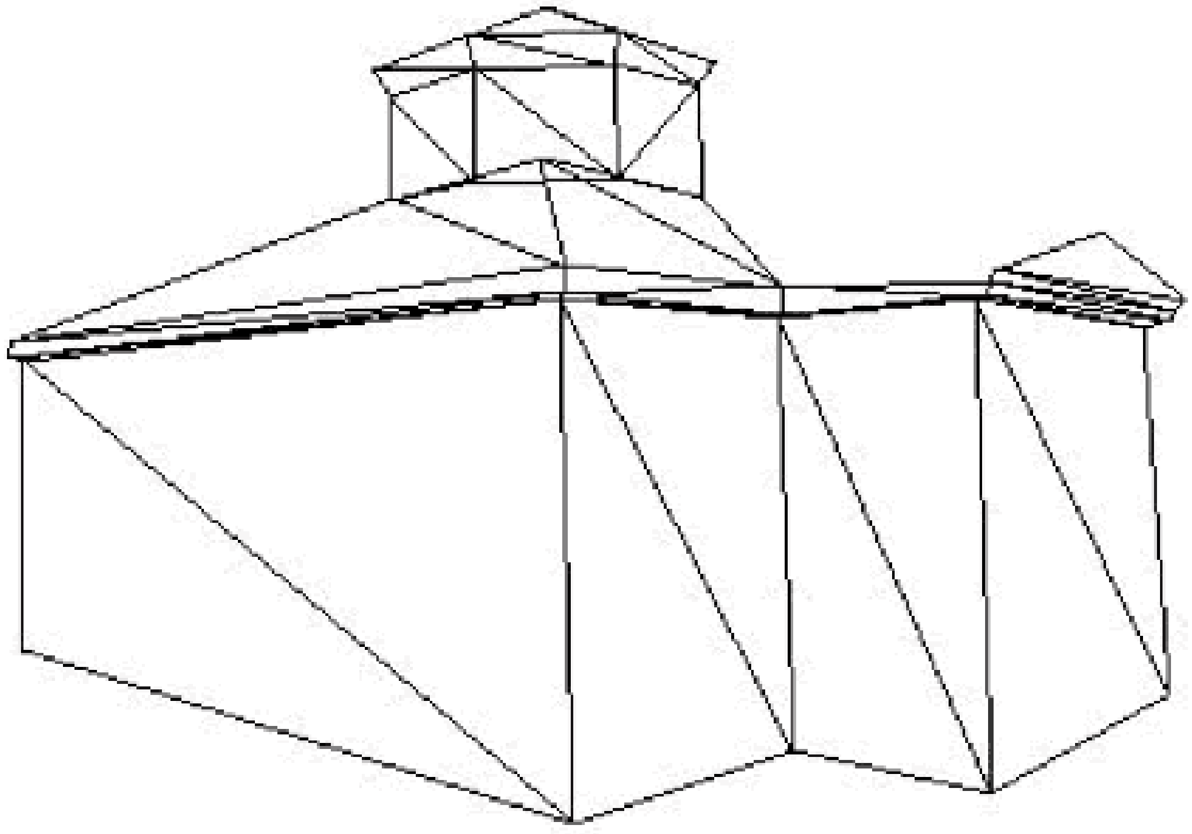

Outline of a simple algorithm (5)

- Compute 3D textured triangles

Facade

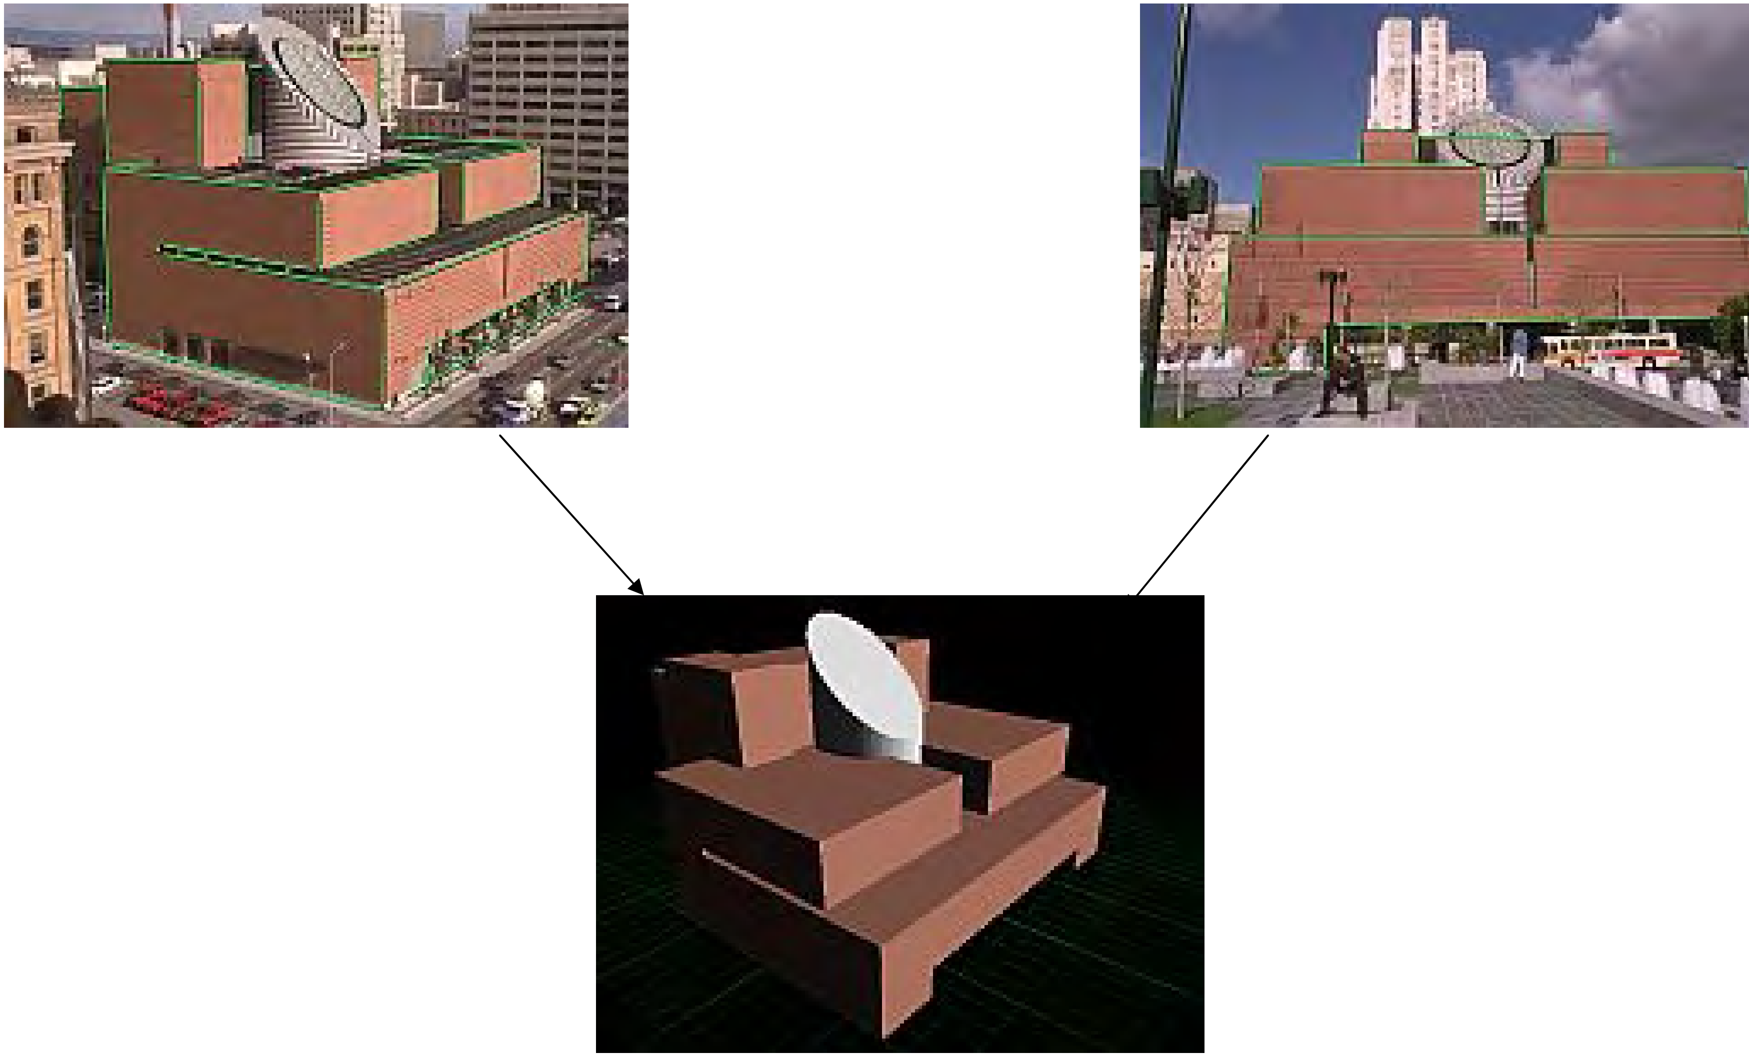

SFMOMA (San Francisco Museum of Modern Art) by Yizhou Yu,

Façade (Debevec et al) inputs

Façade (Debevec et al)

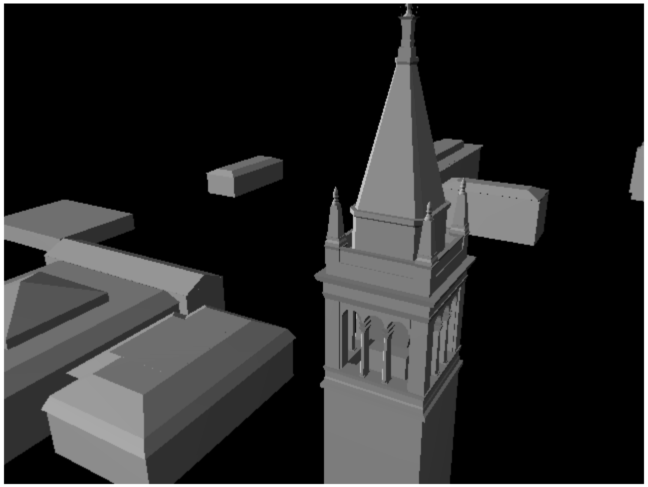

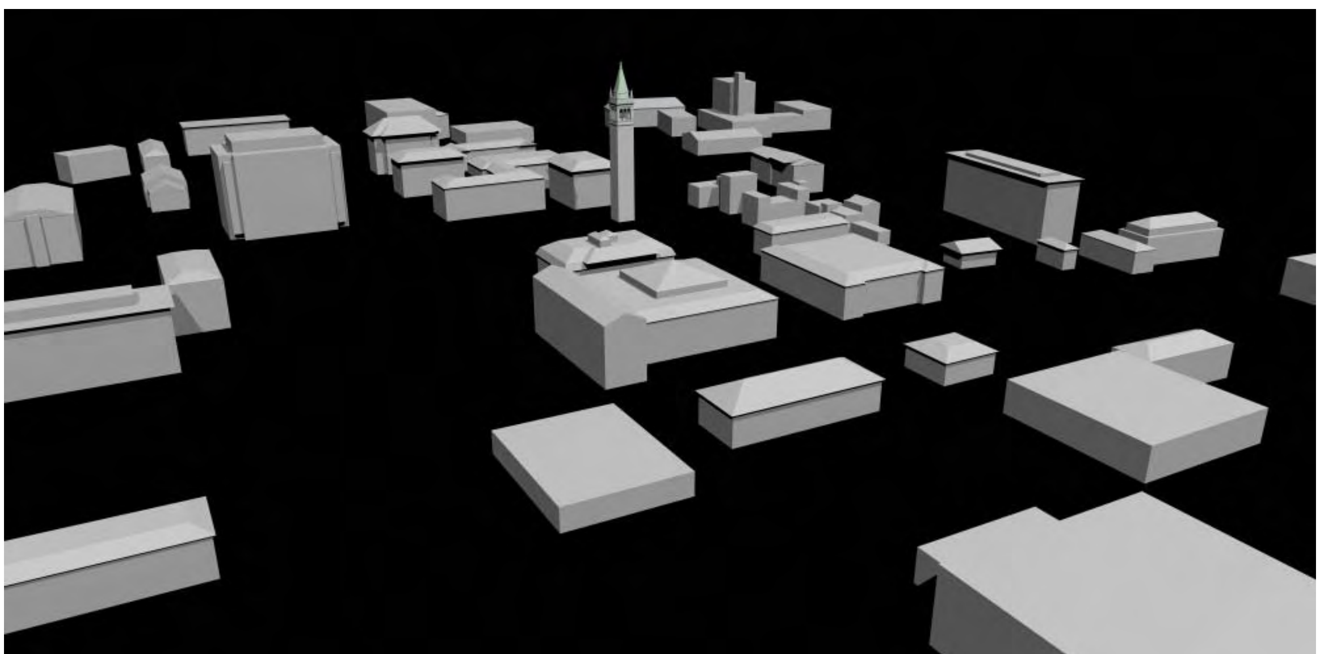

Campanile movie

www.debevec.org