Lightfields and Lumigraphs

Vision for Graphics

CSE 590SS, Winter 2001

Richard Szeliski

(with lots of slides from Michael Cohen)

Modeling light

How do we generate new scenes and animations from existing ones?

Classic “3D Vision + Graphics”:

- take (lots of) pictures
- recover camera pose
- build 3D model
- extract texture maps / BRDFs
- synthesize new views

Computer Graphics

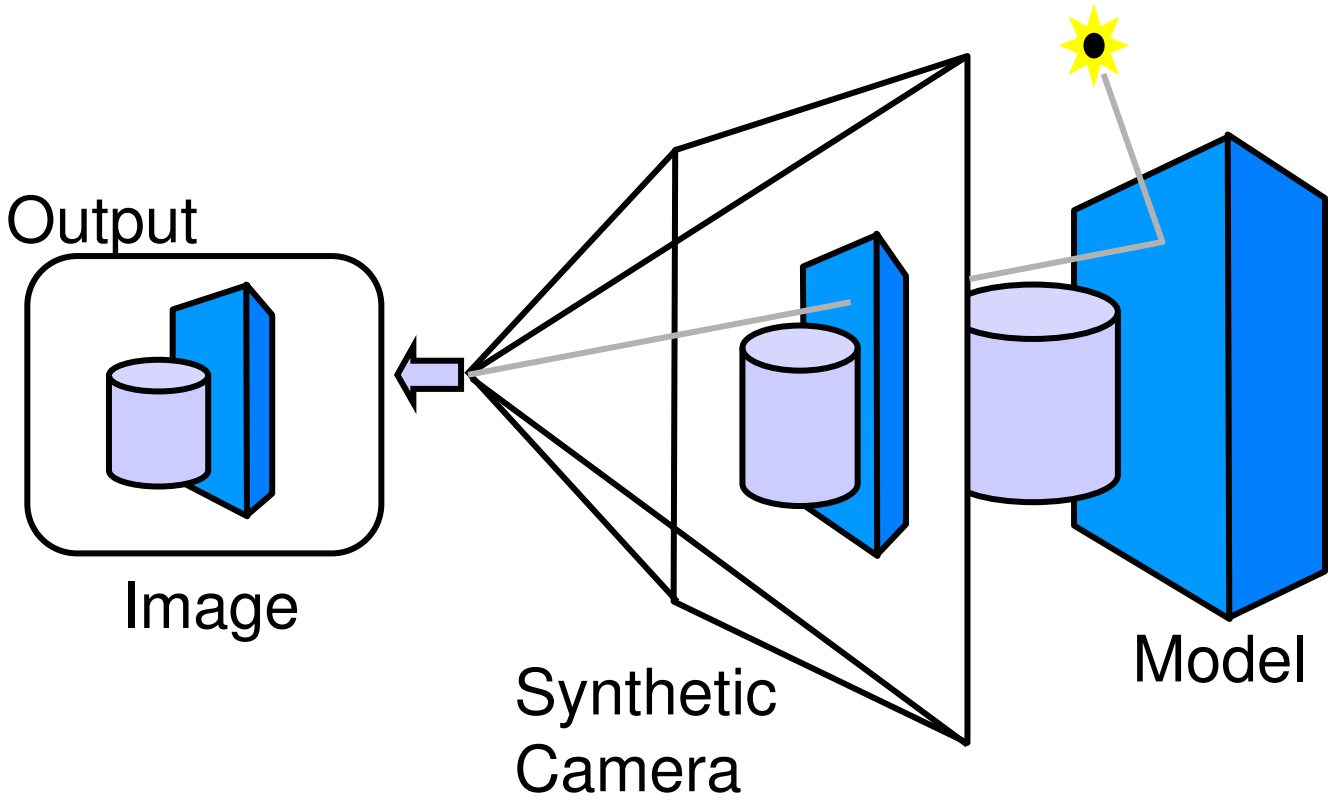

Computer Vision

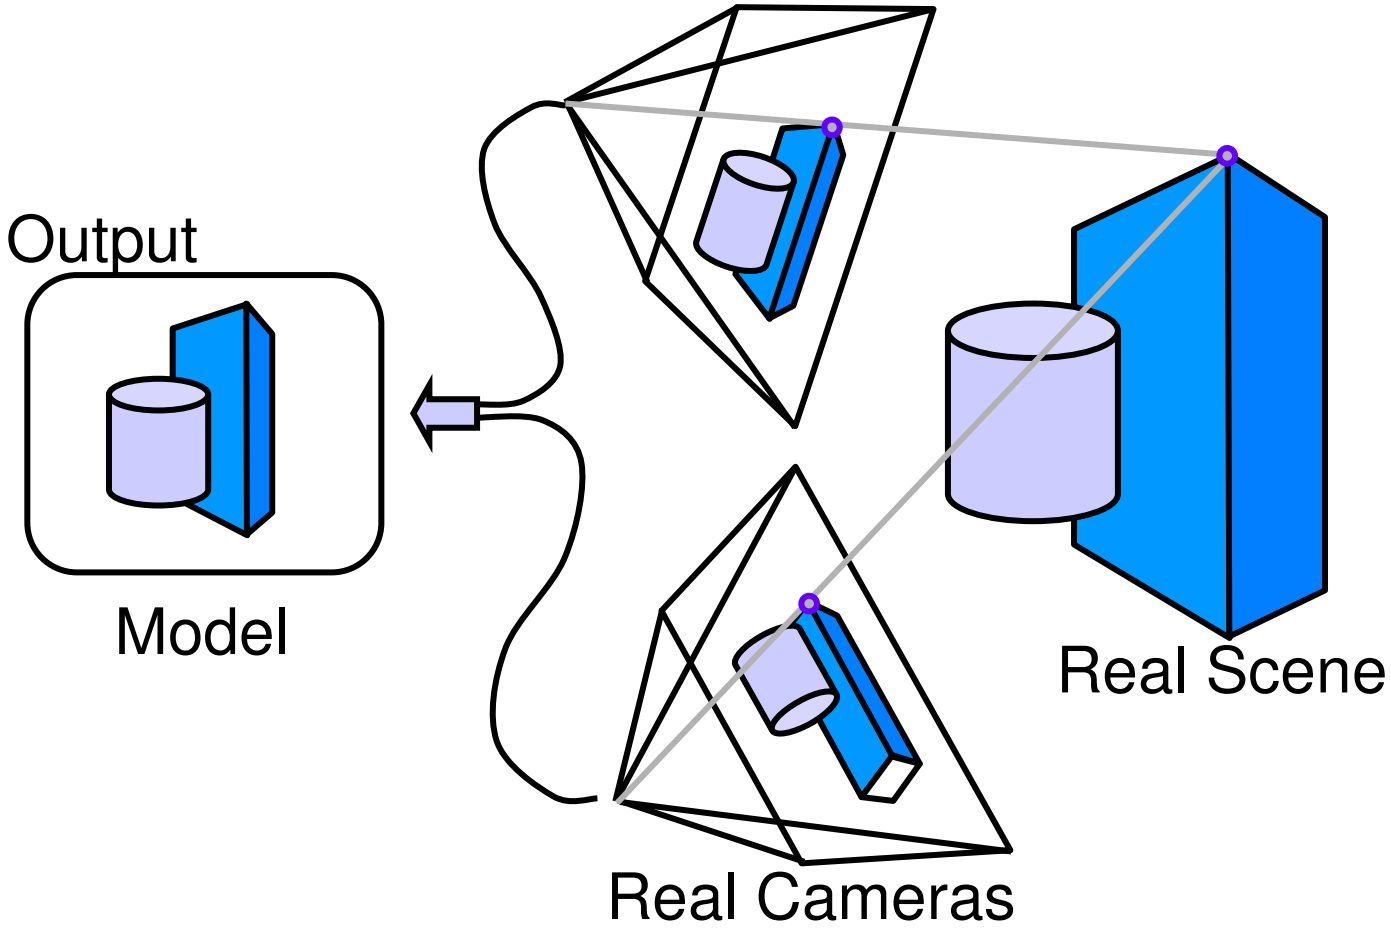

Combined

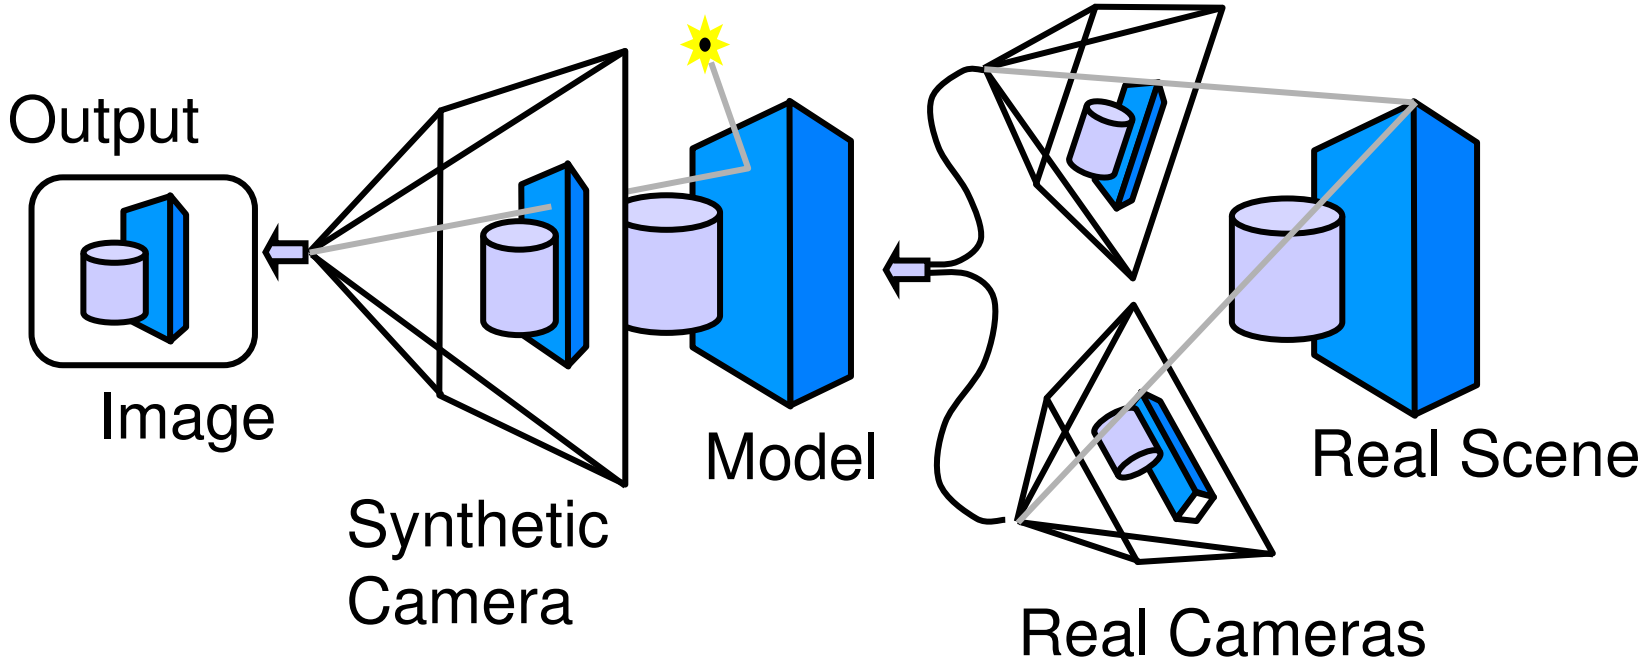
Output
Image
Synthetic
Camera
Model
Real Scene
Real Cameras

But, vision technology falls short

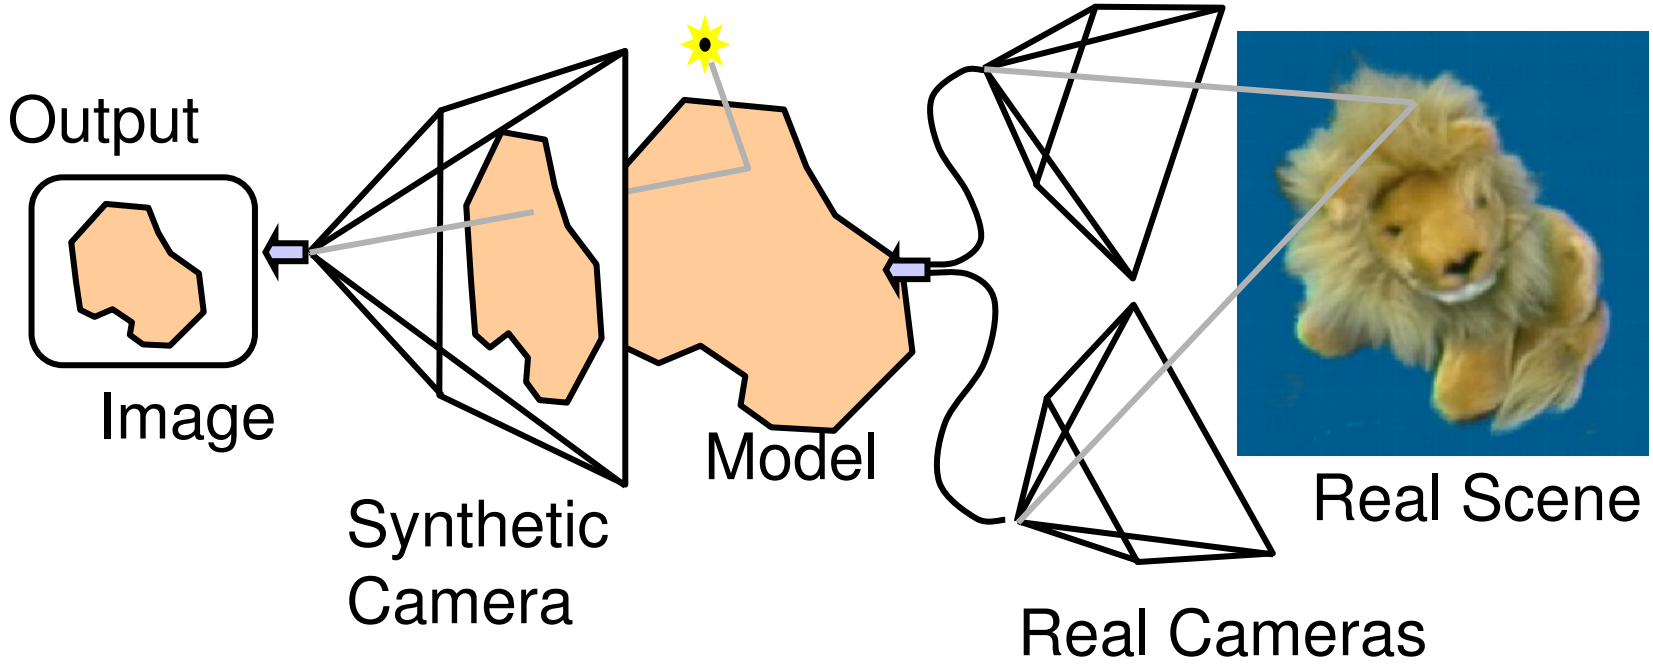

… and so does graphics.

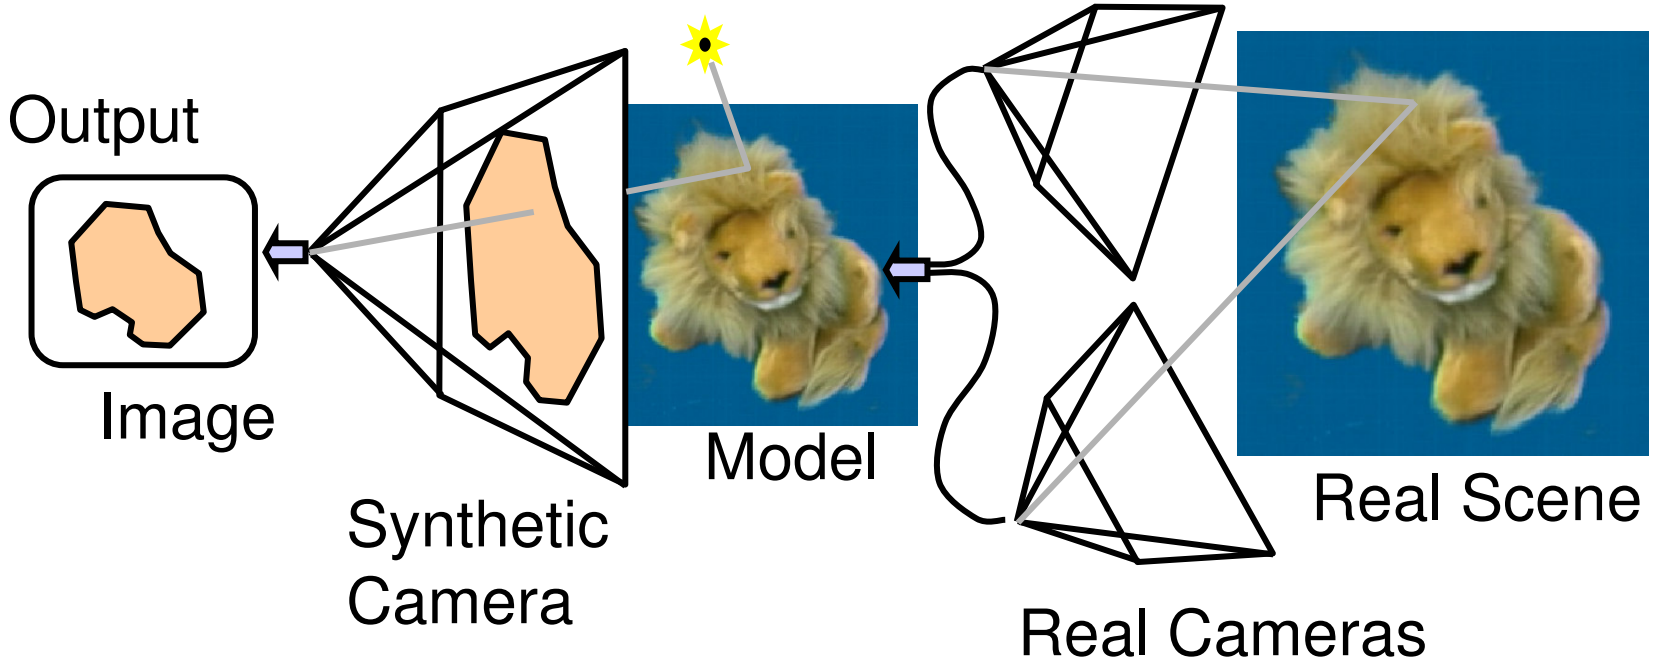

Image Based Rendering

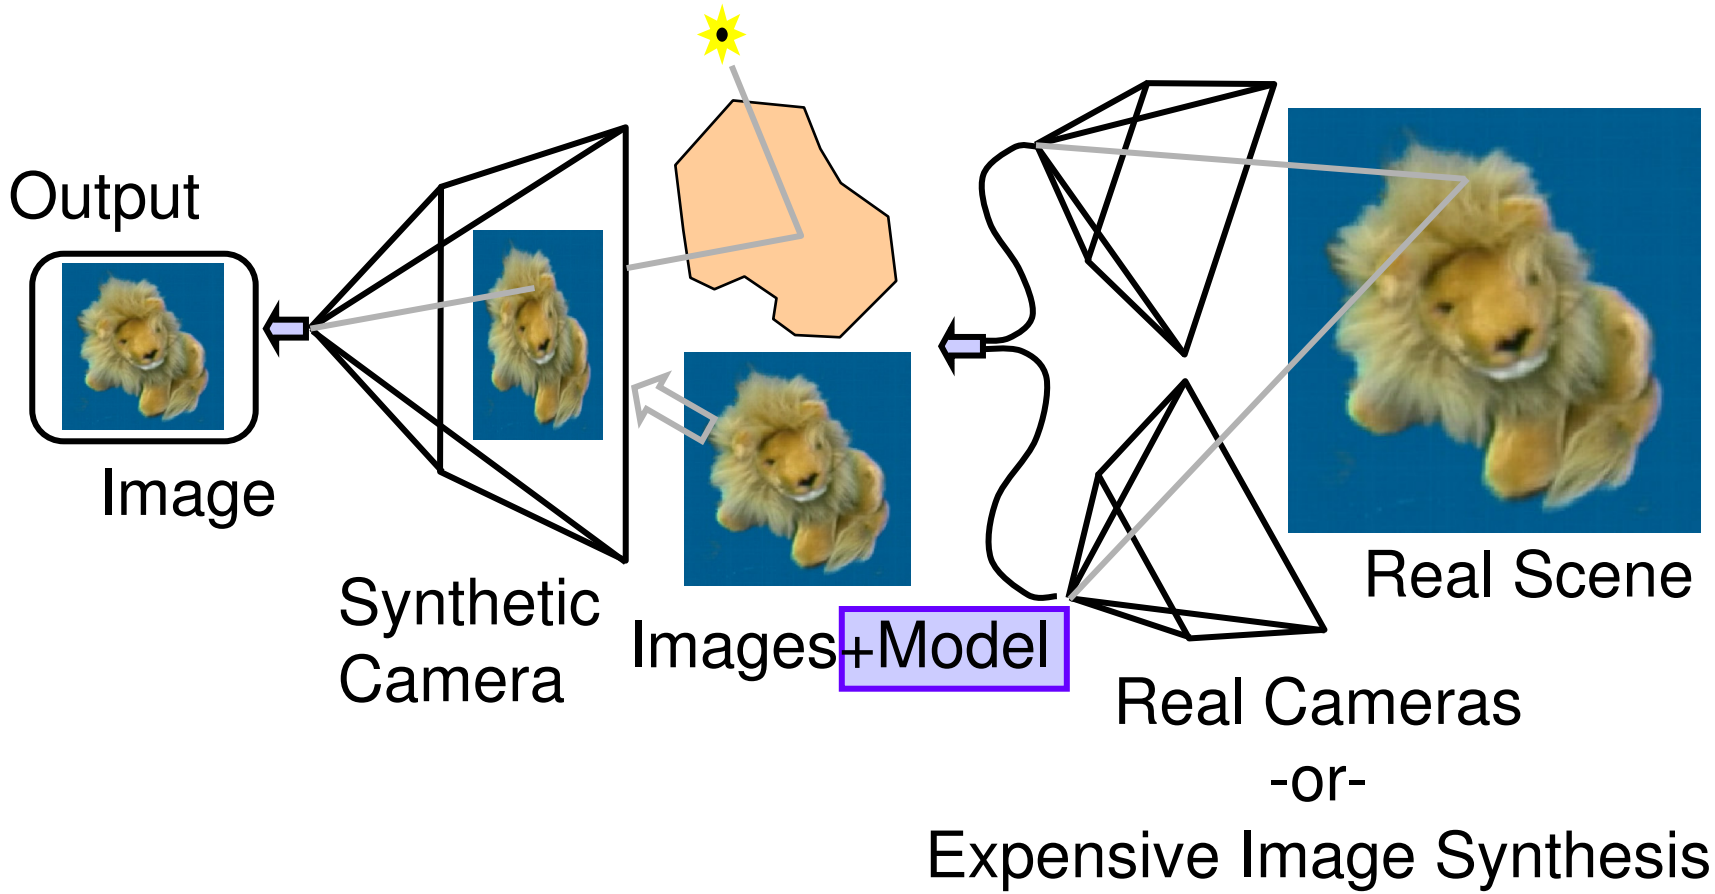

Ray

Constant radiance

- time is fixed

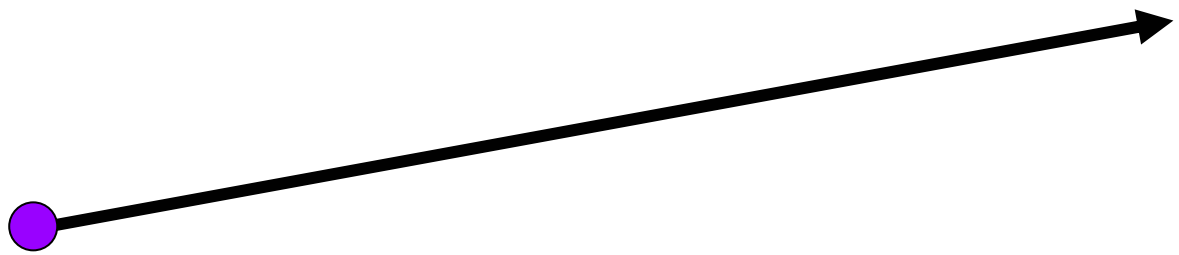

5D

- 3D position
- 2D direction

All Rays

Plenoptic Function:

- all possible images
- too much stuff!

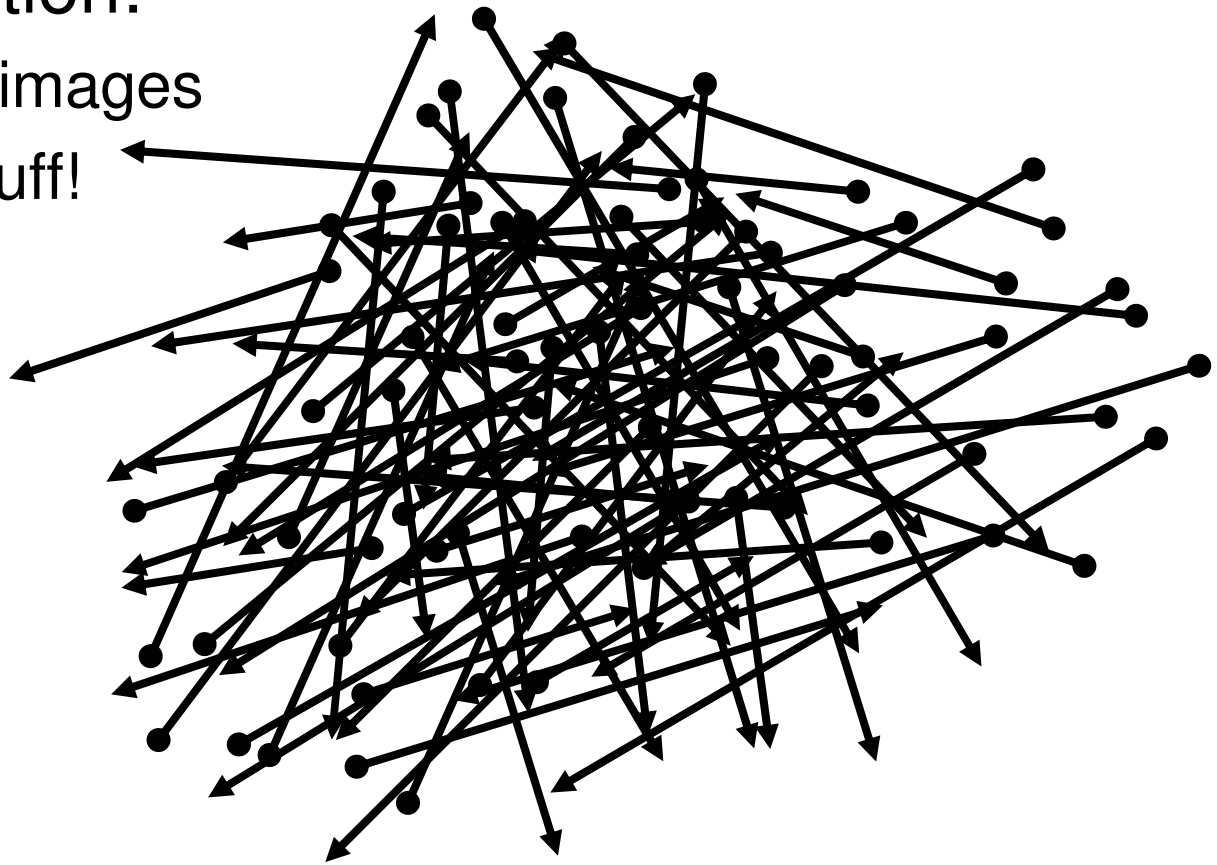

Line

Infinite line

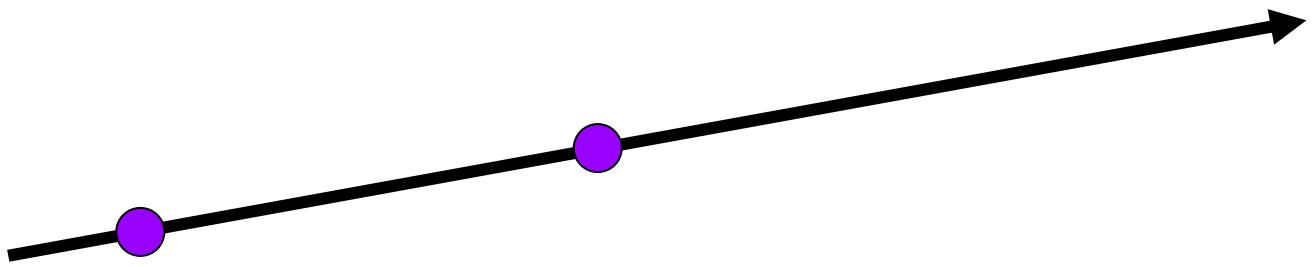

4D

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

Ray

Discretize, then interpolate

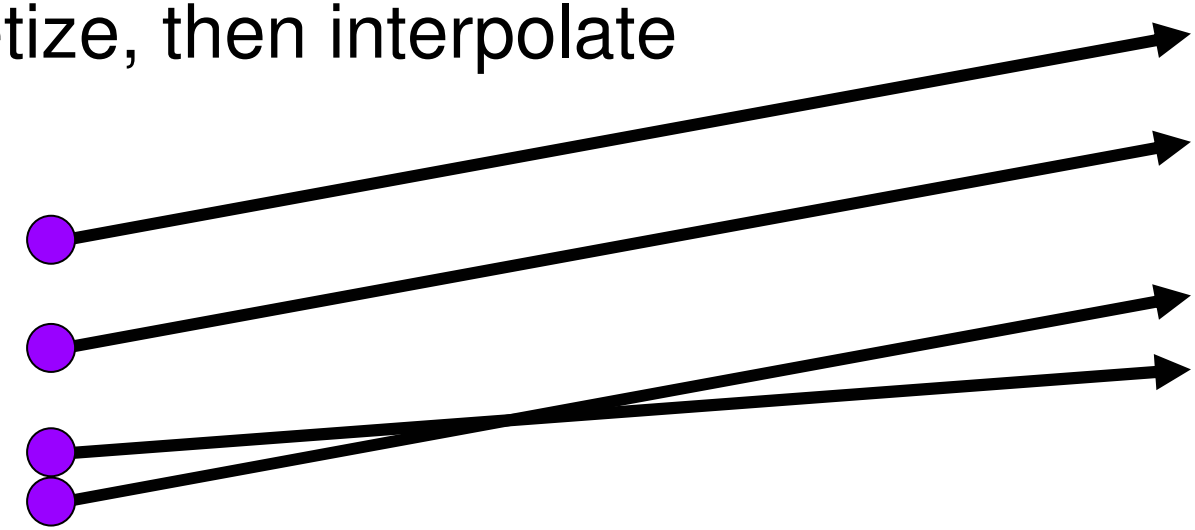

Distance between 2 rays

- Which is closer together?

Image

What is an image?

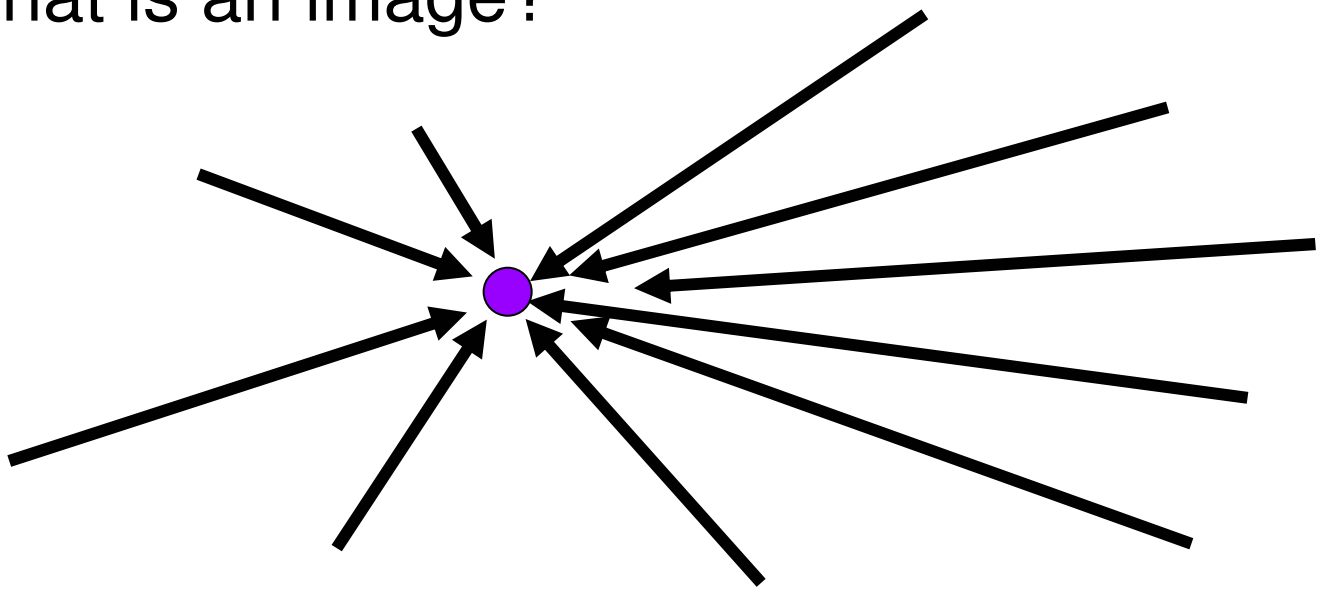

All rays through a point

- Panorama?

Image

Image plane

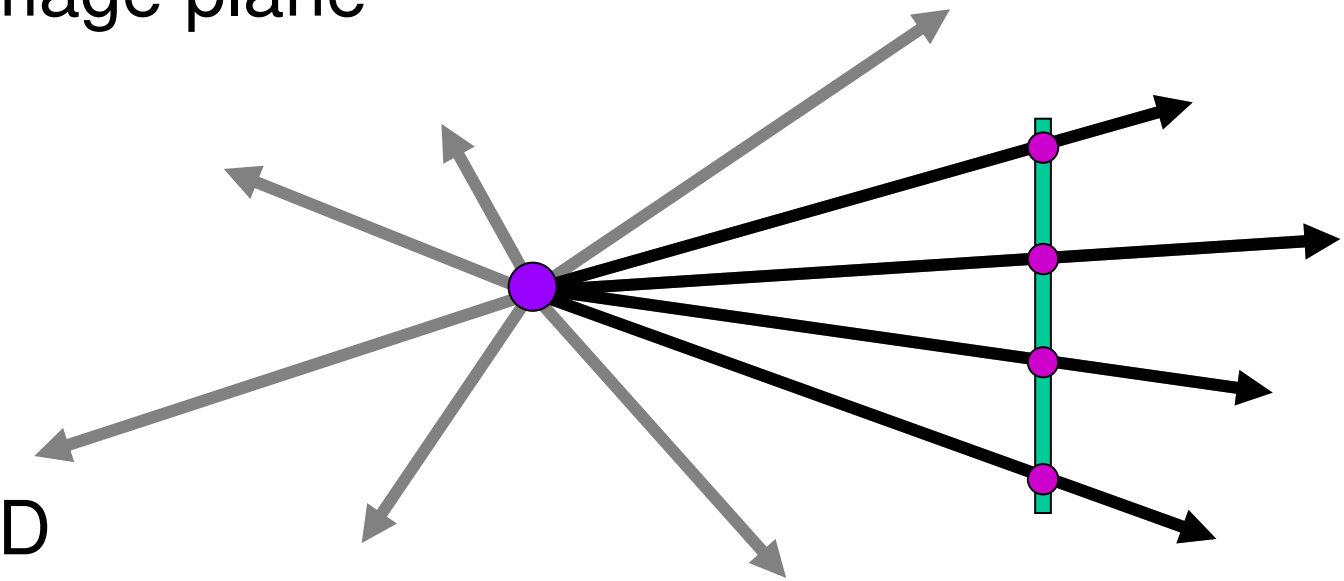

2D

- position

Object

Light leaving towards “eye”

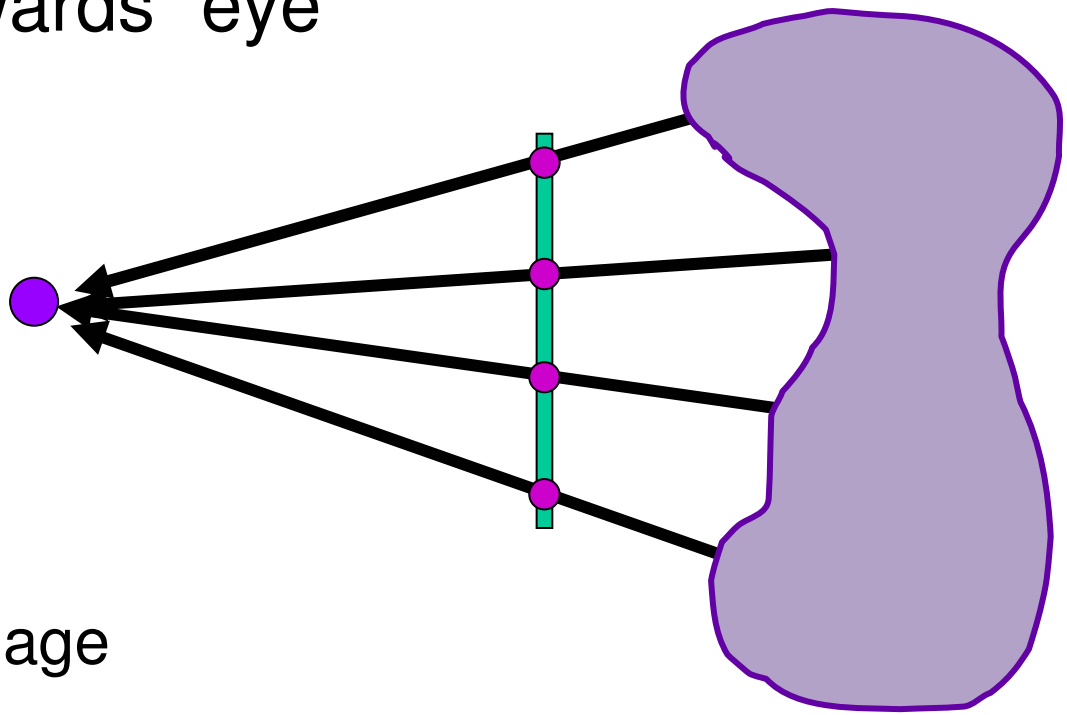

2D

- just dual of image

Object

All light leaving object

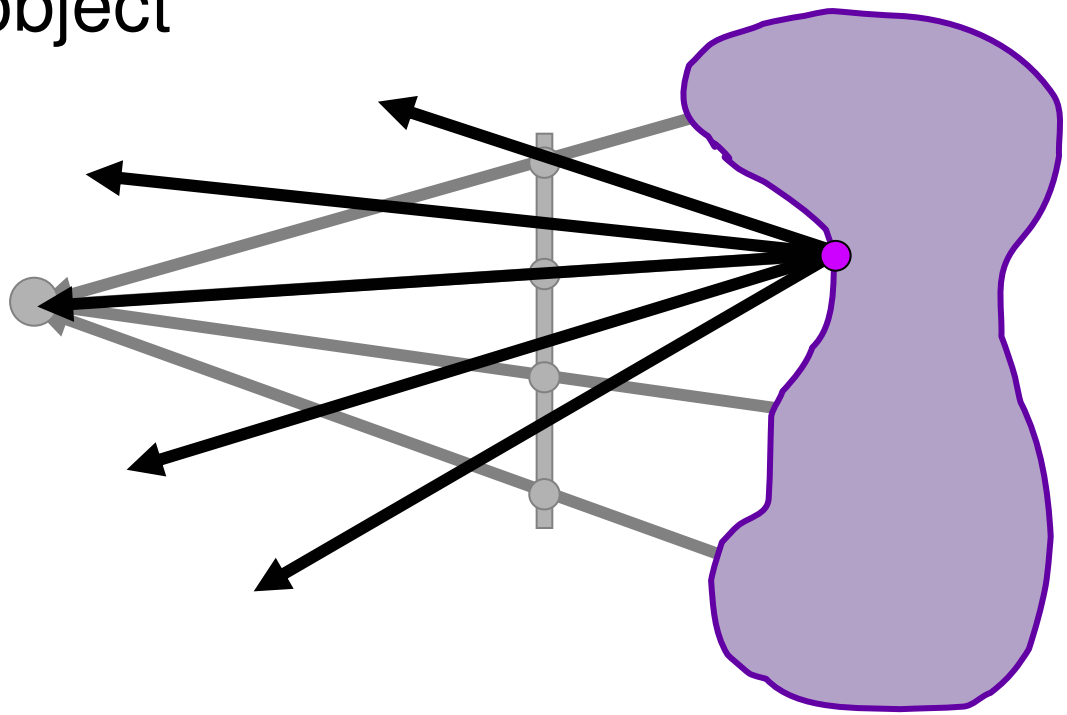

Object

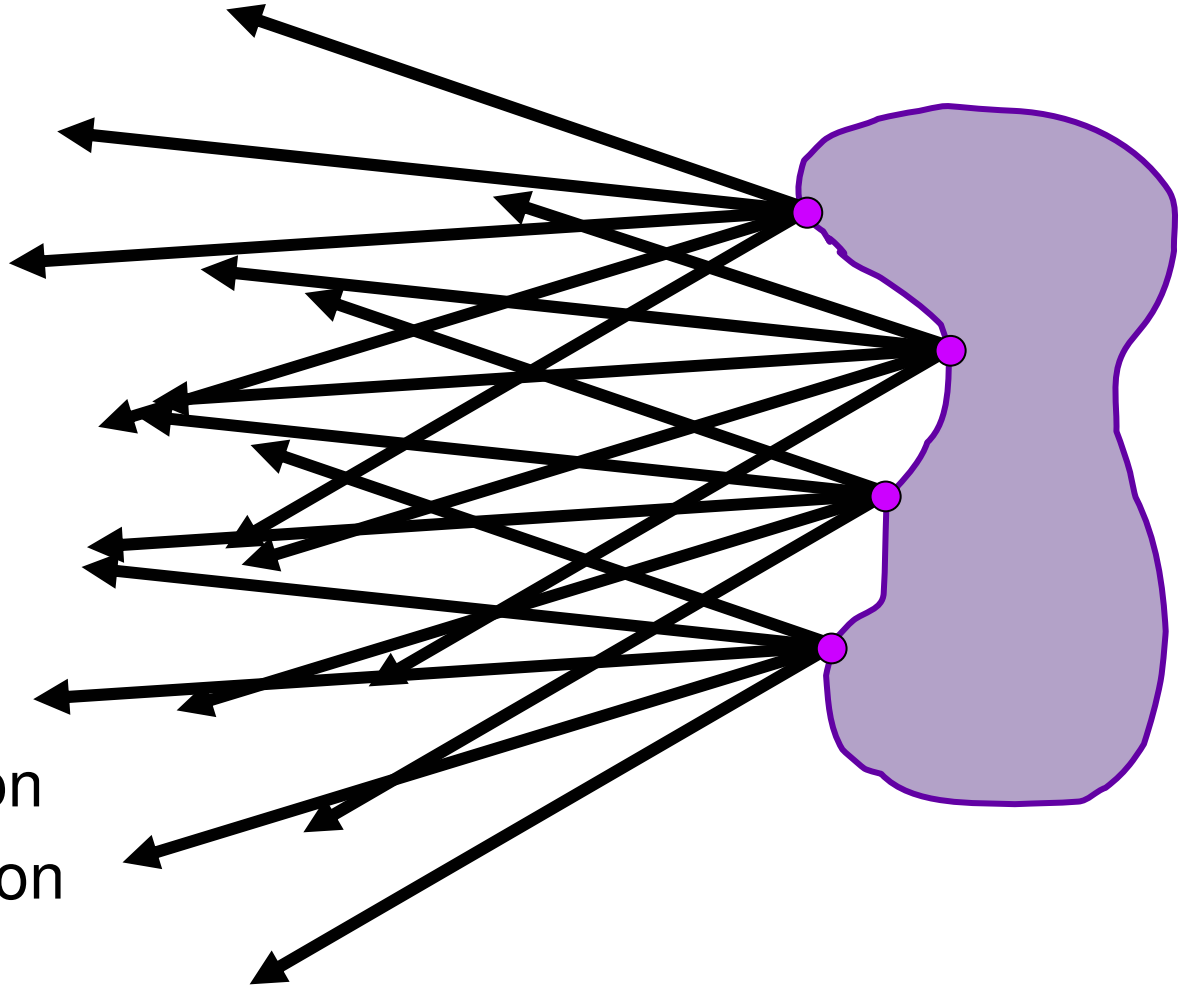

4D

- 2D position
- 2D direction

Object

All images

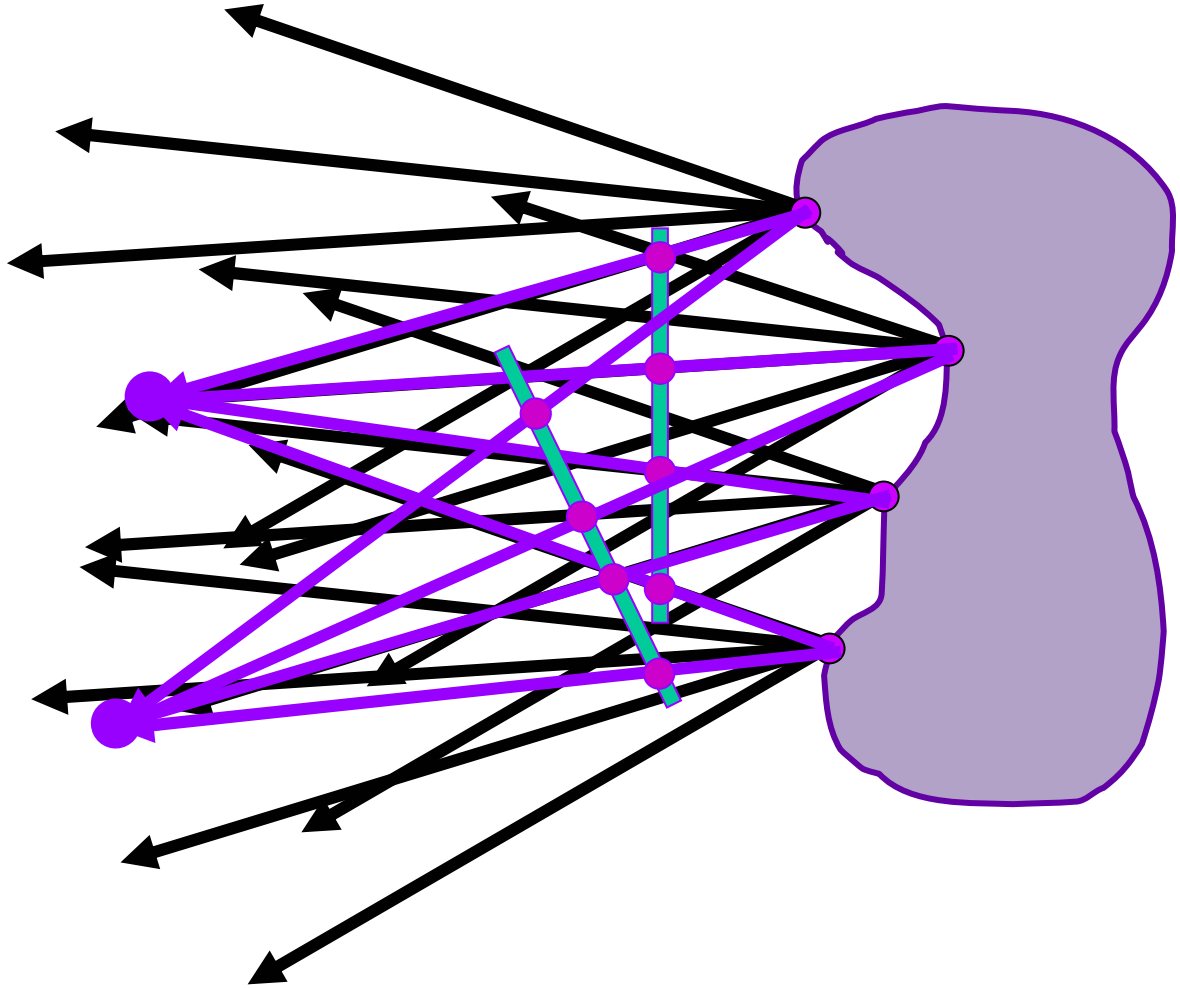

Lumigraph / Lightfield

Outside convex space

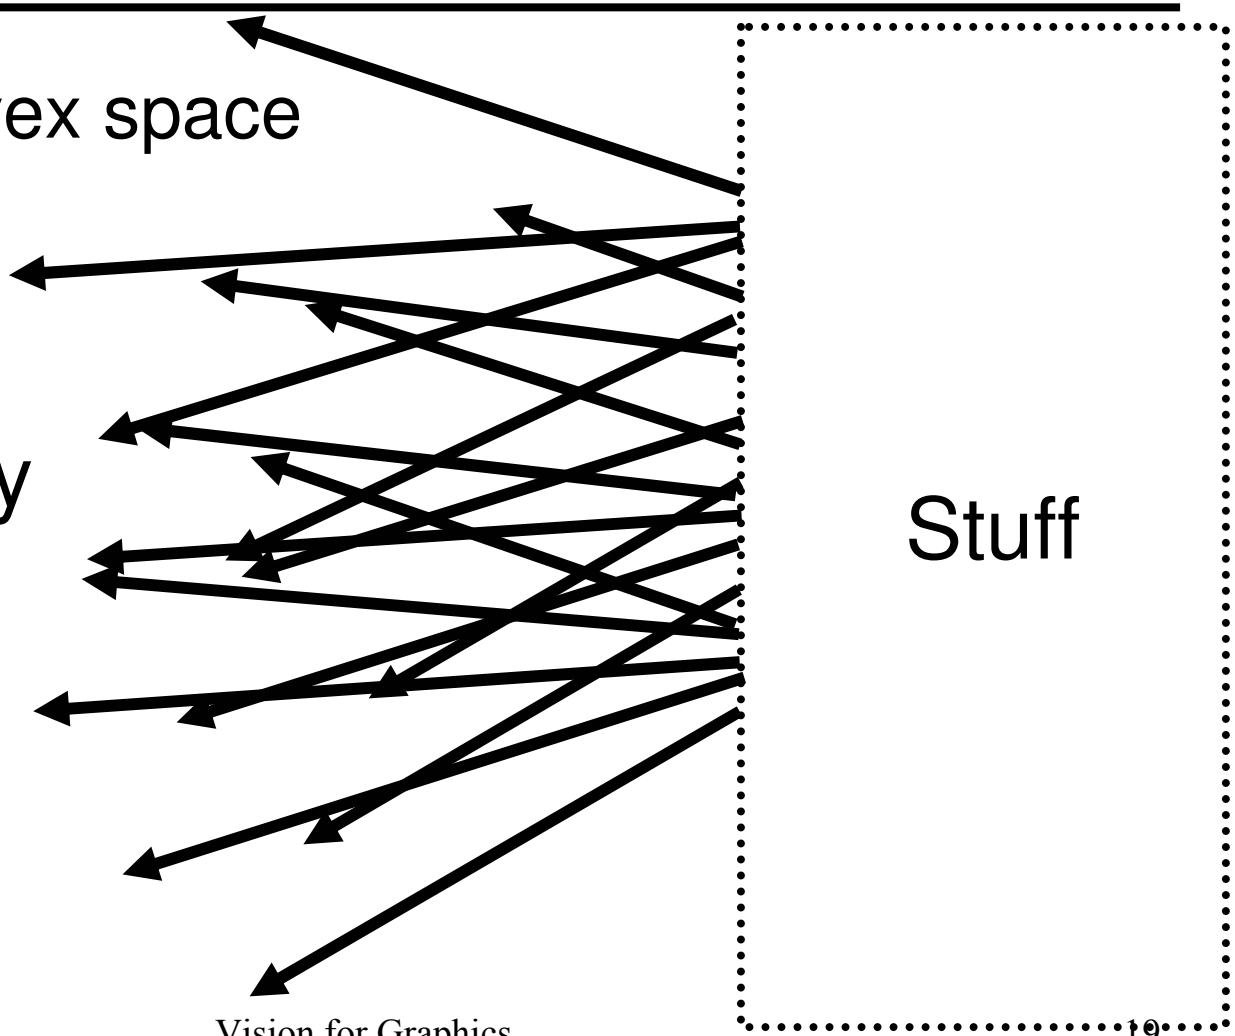

4D

Lumigraph

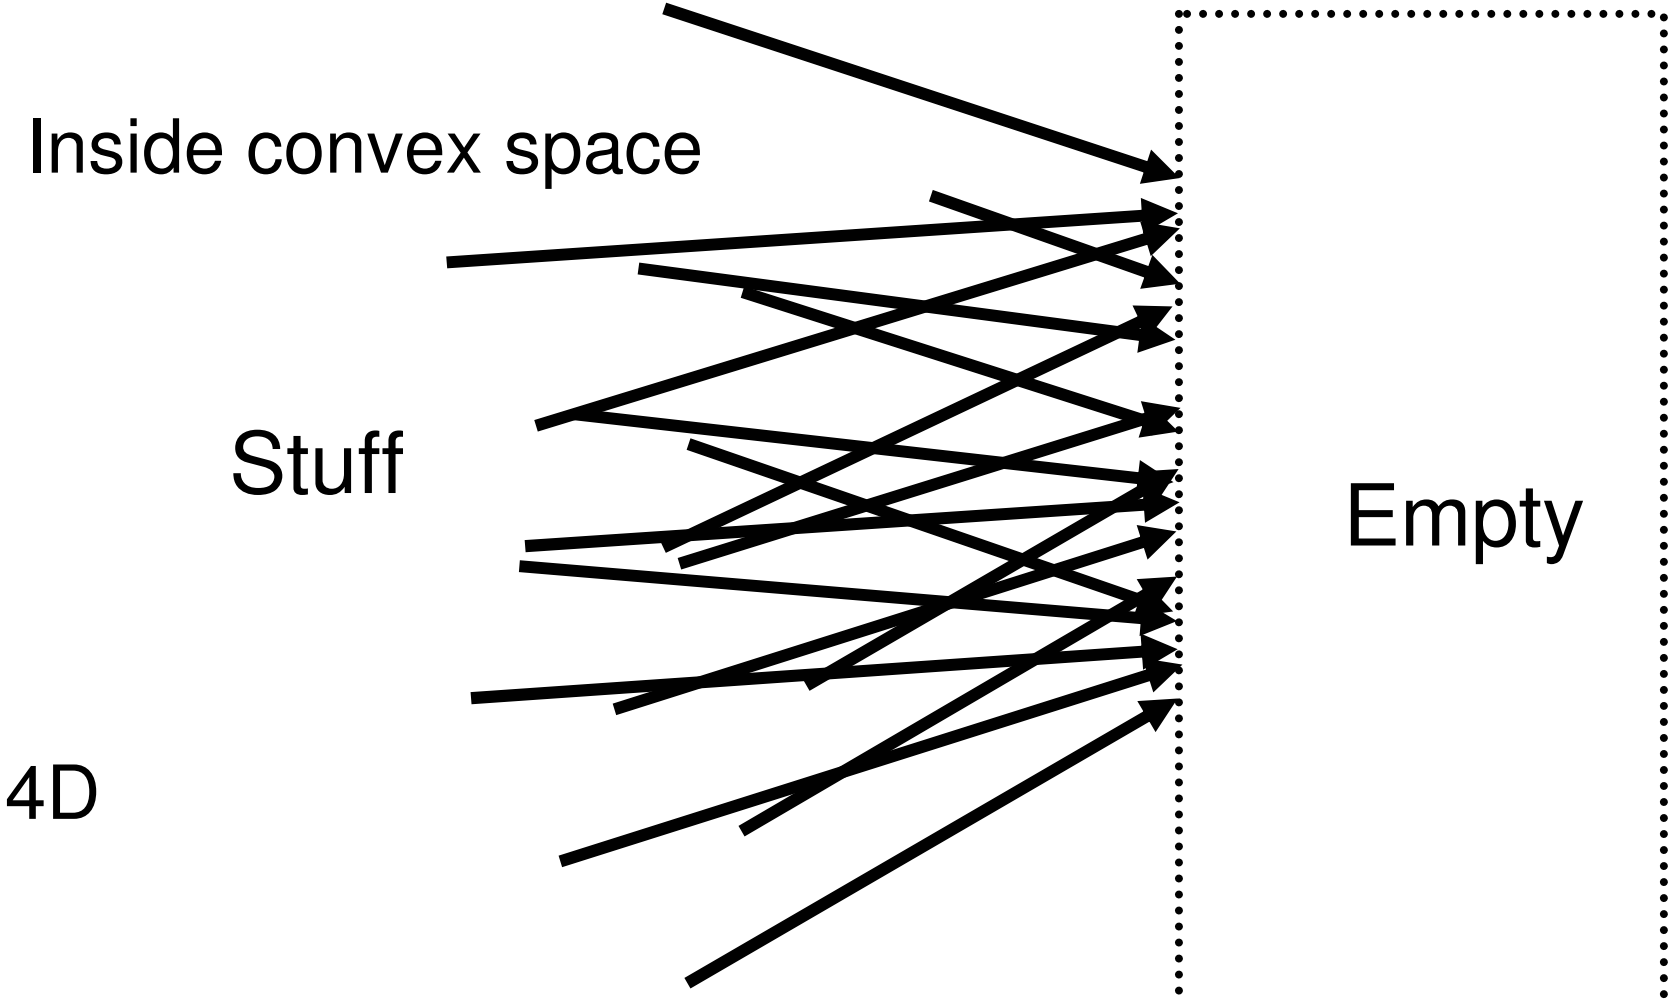

Lumigraph

How to

- organize
- capture
- render

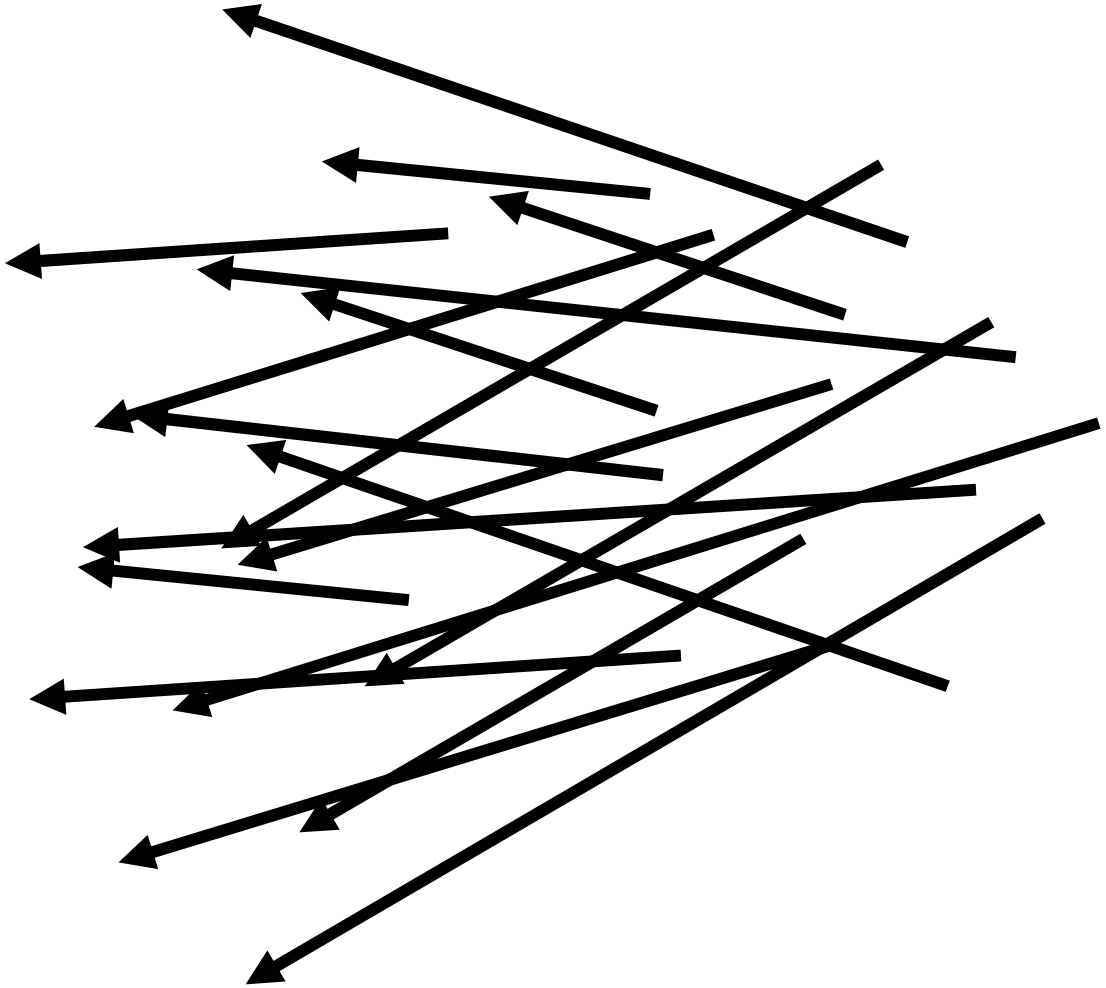

Lumigraph - Organization

2D position

2D direction

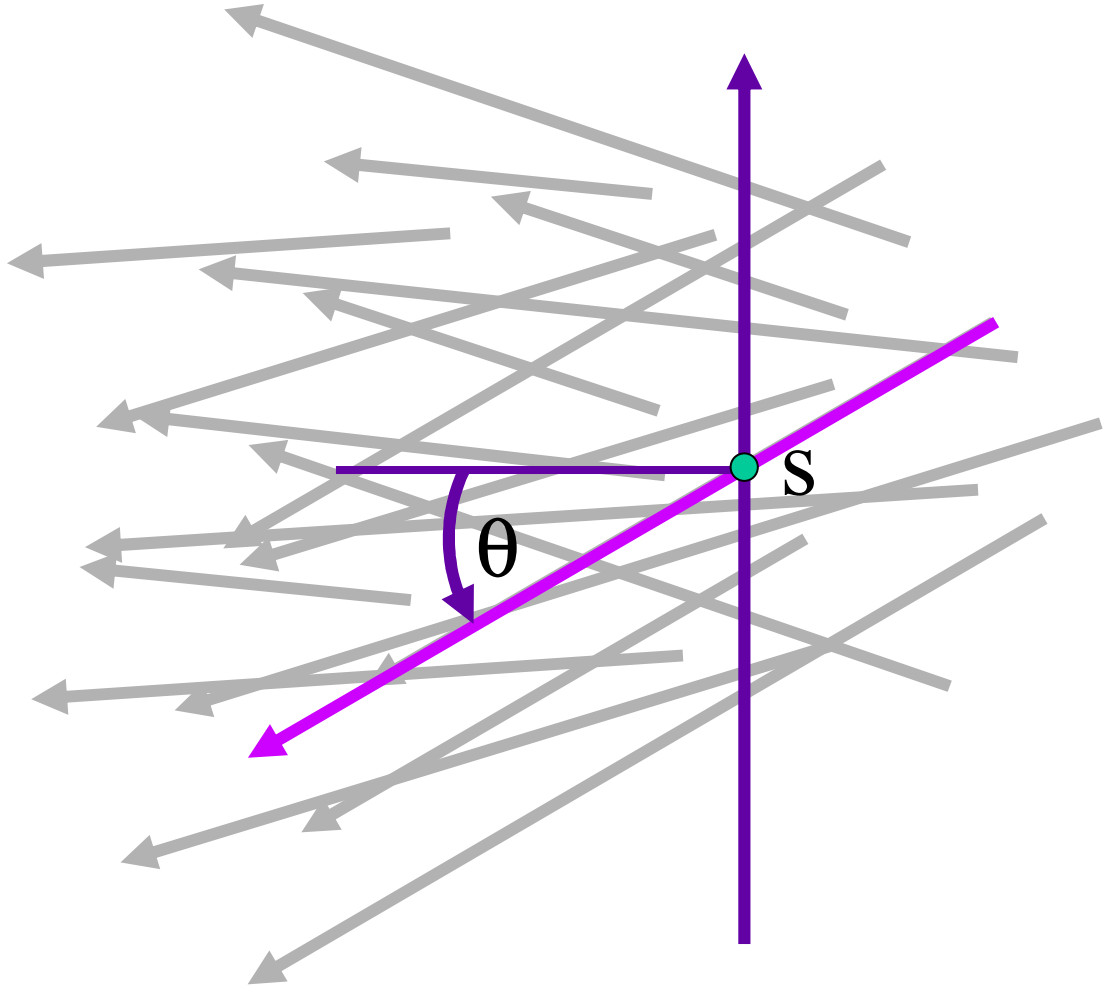

Lumigraph - Organization

2D position

2D position

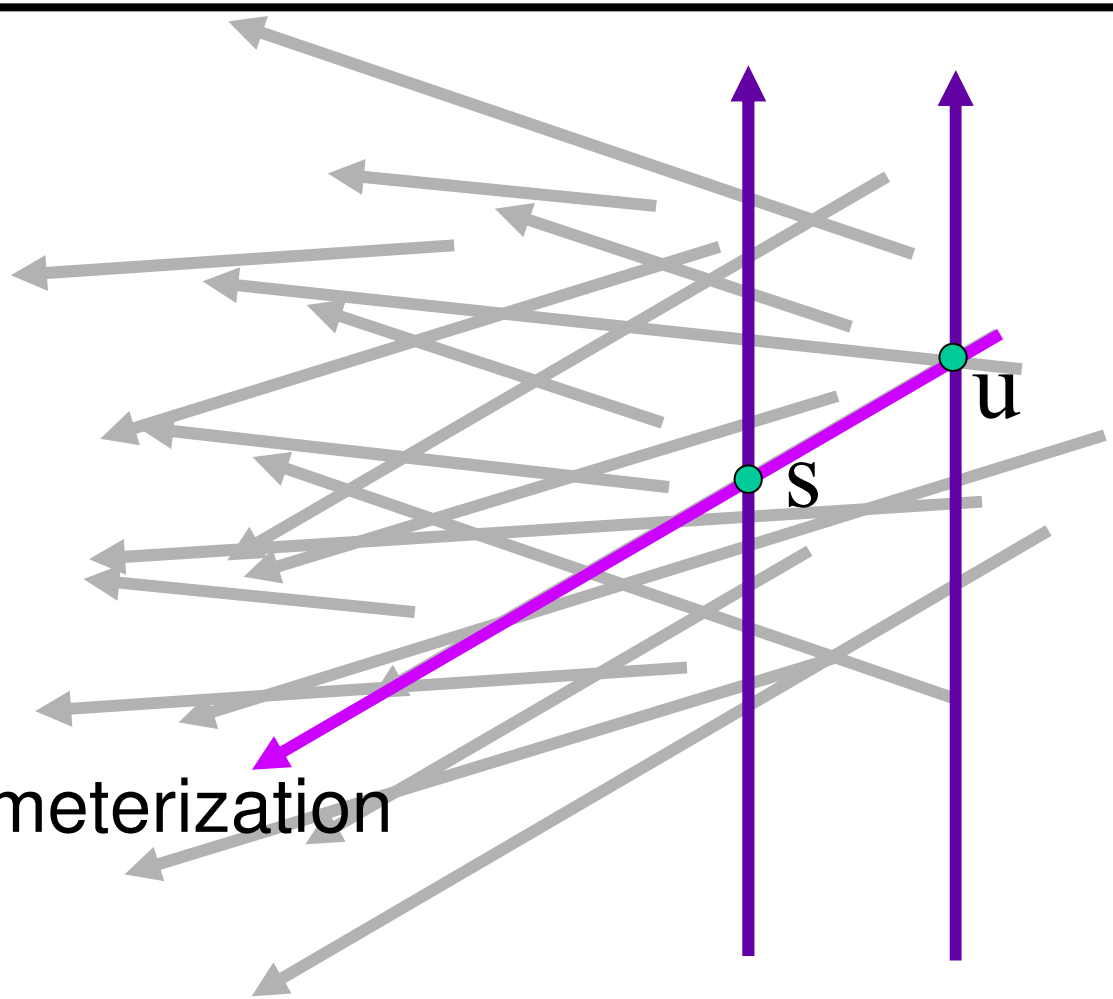

2 plane parameterization

Lumigraph - Organization

2D position

2D position

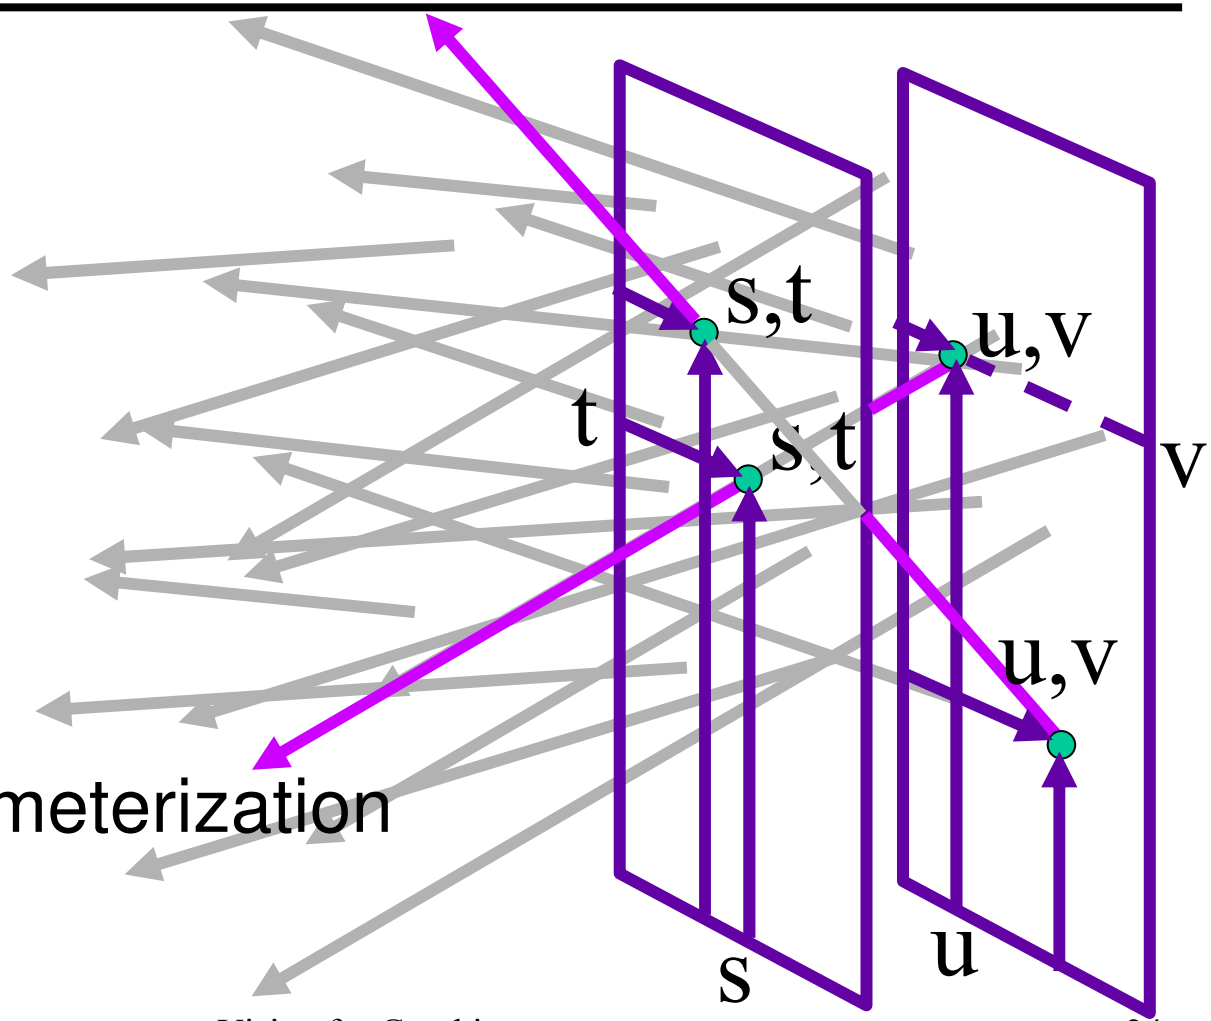

2 plane parameterization

Lumigraph - Organization

Hold s,t constant

Let u,v vary

An image

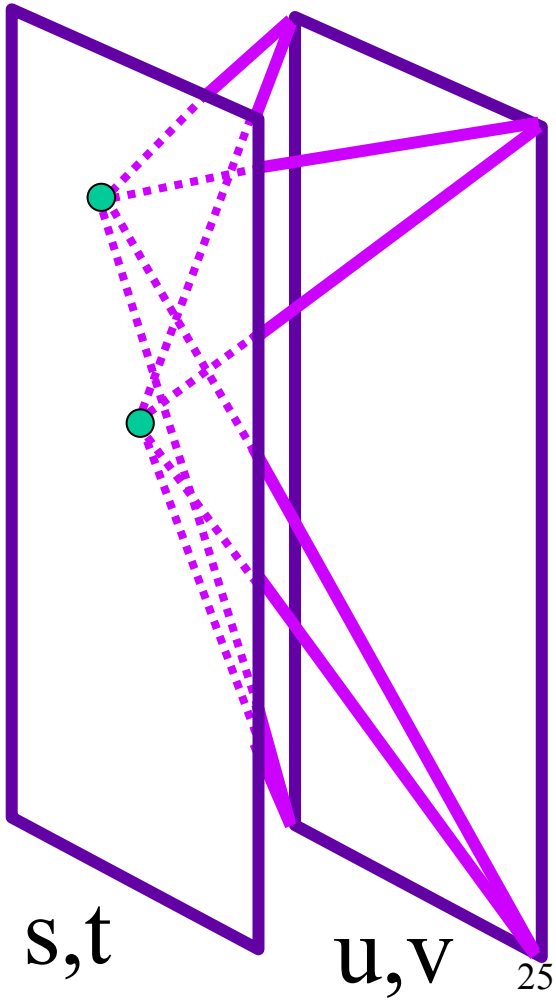

Lumigraph - Organization

Discretization

- higher res near object
  - if diffuse
  - captures texture
- lower res away
  - captures directions

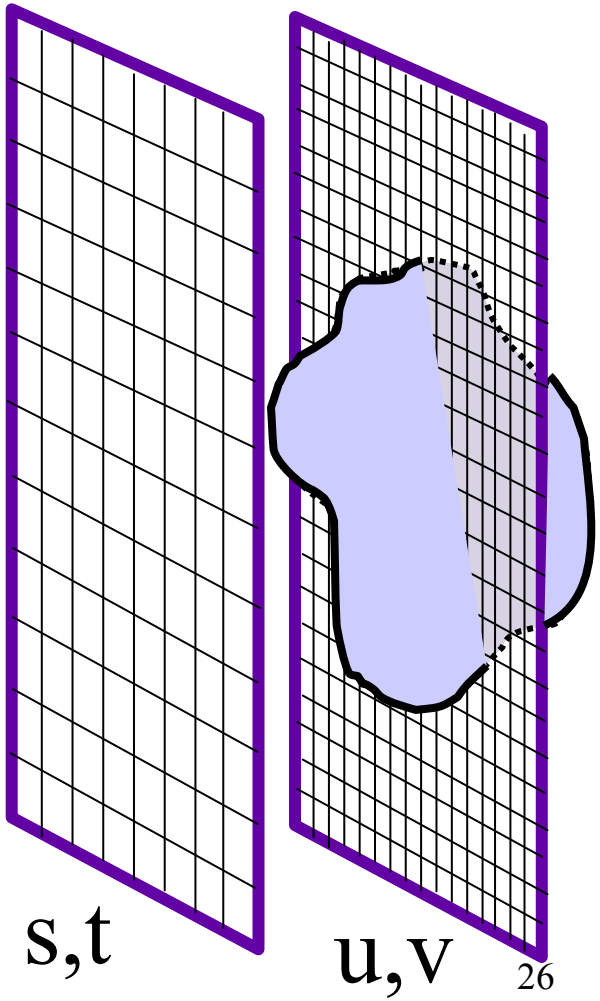

Lumigraph - Capture

Idea 1

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

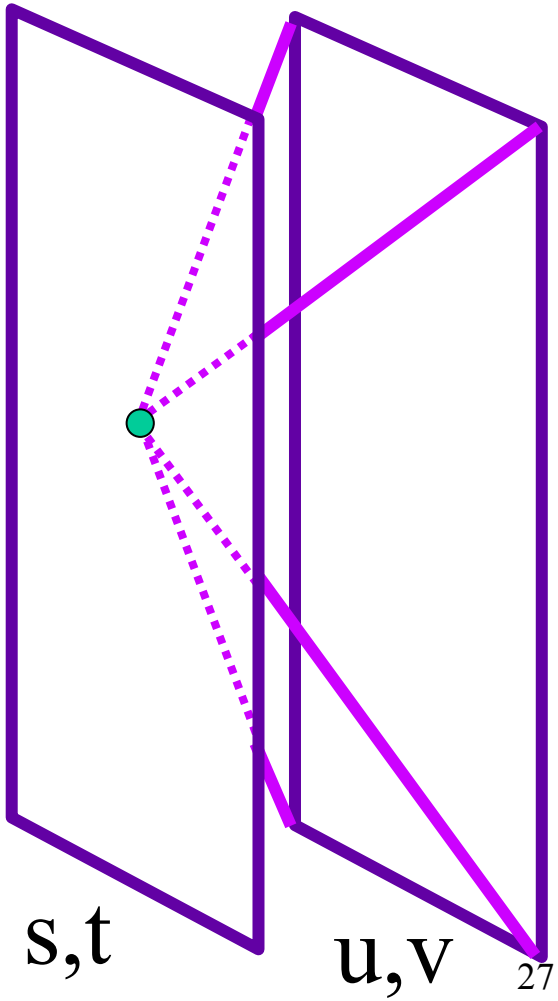

Lumigraph - Capture

Idea 2

- Move camera anywhere
- Rebinning
  - see Lumigraph paper

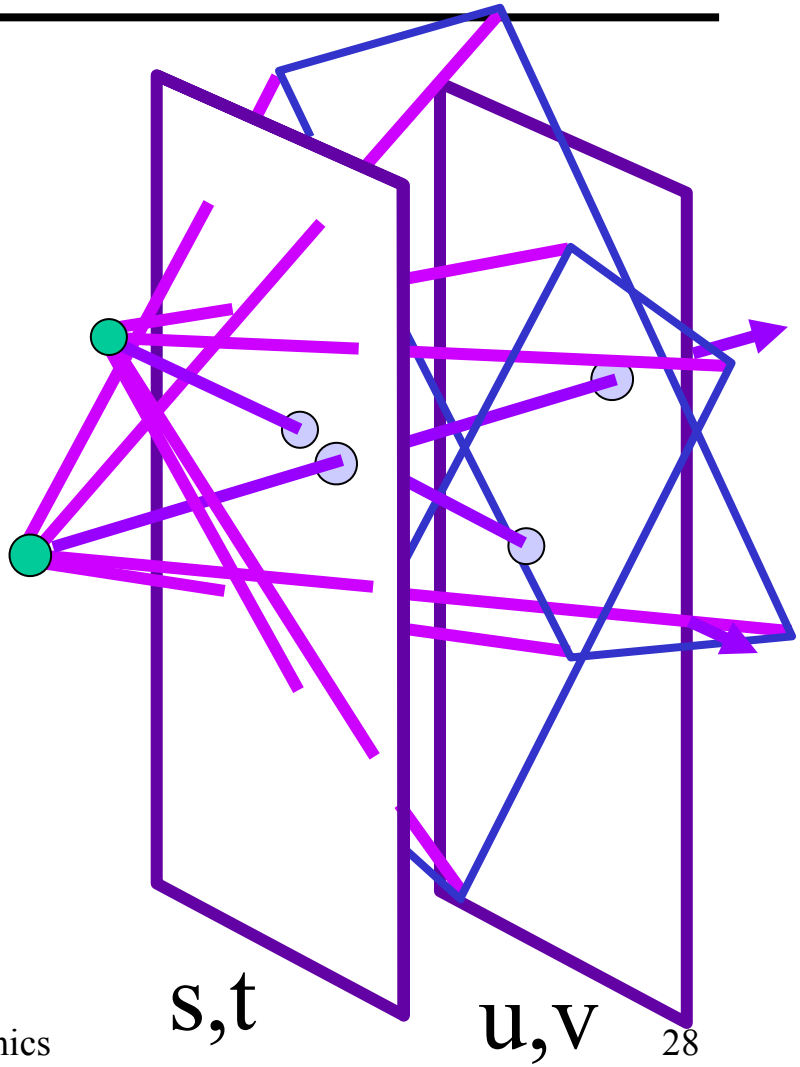

Lumigraph - Rendering

For each output pixel

- determine s,t,u,v
- either
  - find closest discrete RGB
  - interpolate near values

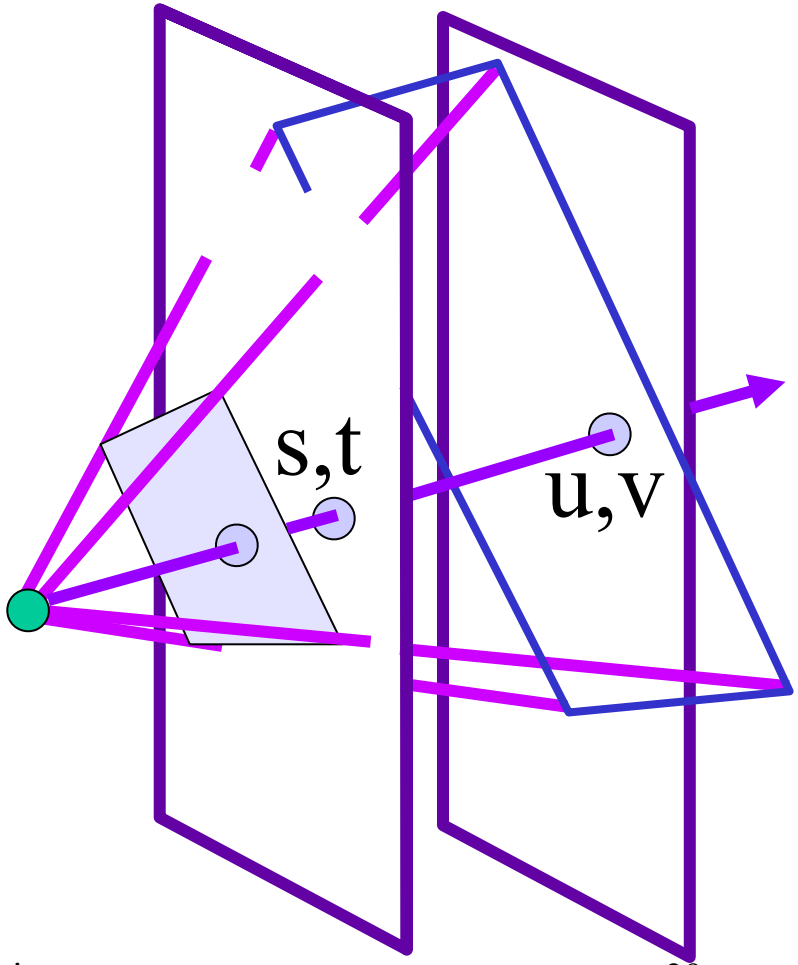

Lumigraph - Rendering

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

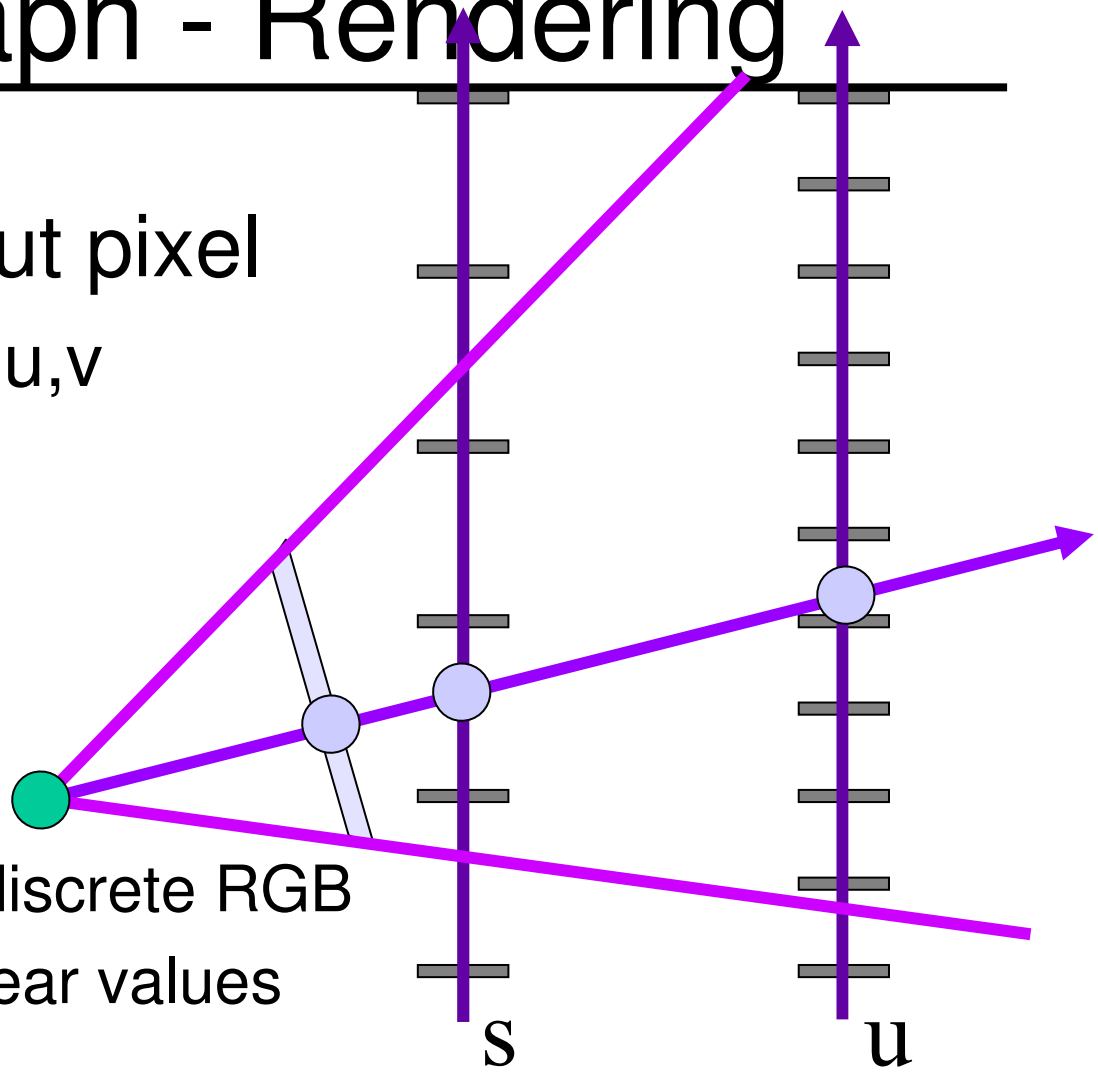

Lumigraph - Rendering

Nearest

- closest s
- closest u
- draw it

Blend 16 nearest

- quadrilinear interpolation

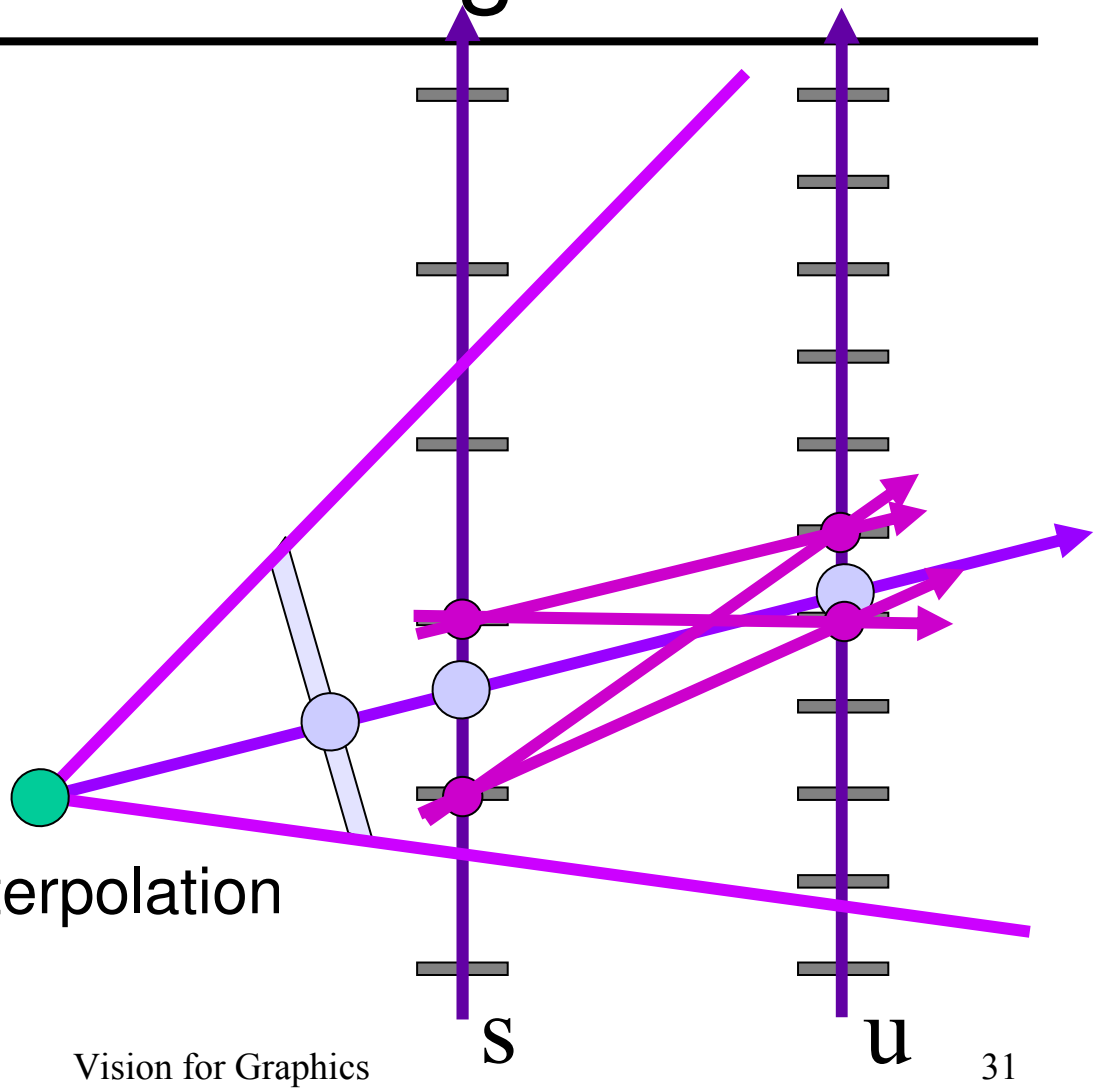

Lumigraph - Rendering

Depth Correction

- closest s
- intersection with “object”
- best u
- closest u

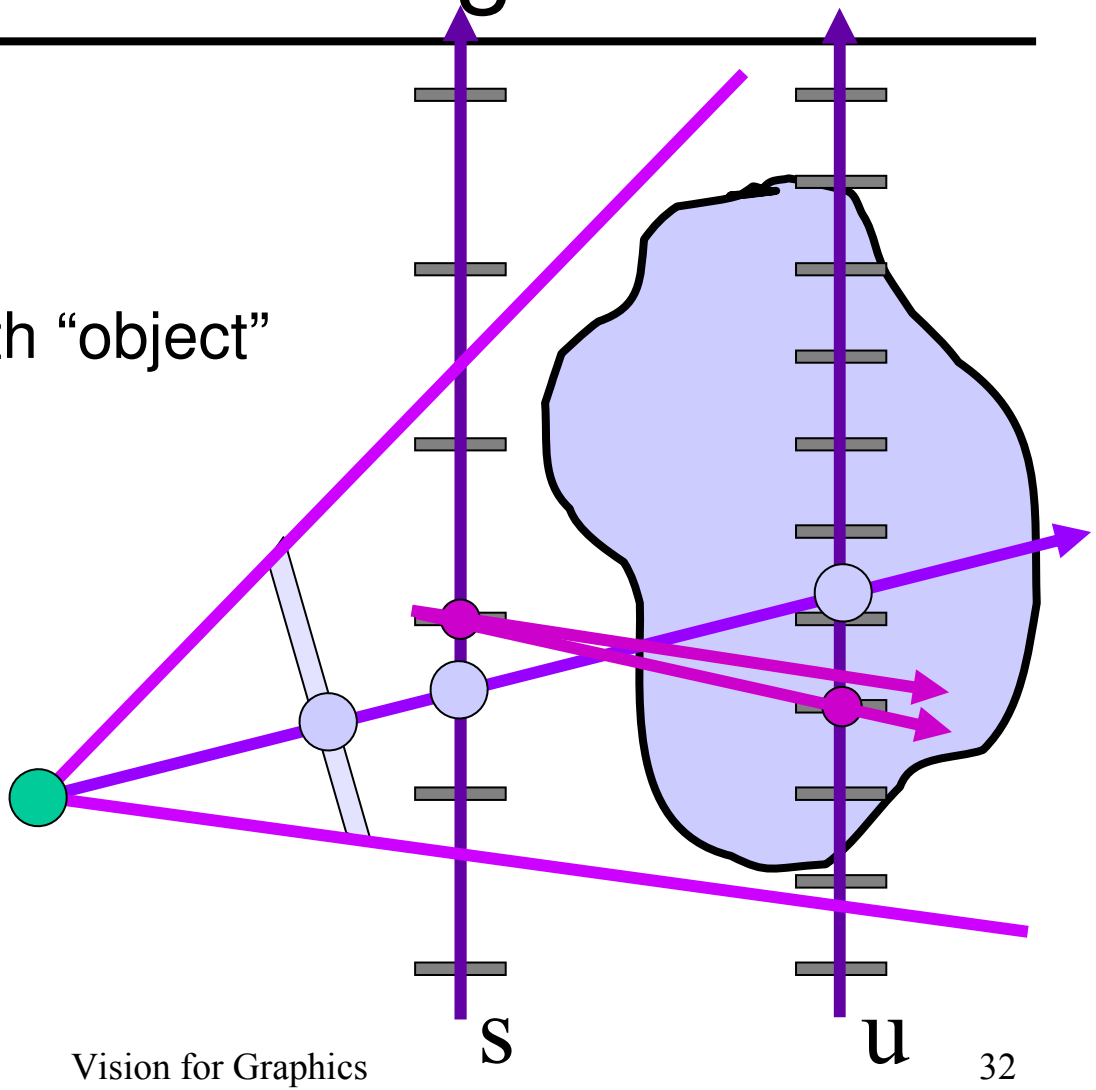

Lumigraph - Rendering

Depth Correction

- quadralinear interpolation
- new “closest”
- like focus

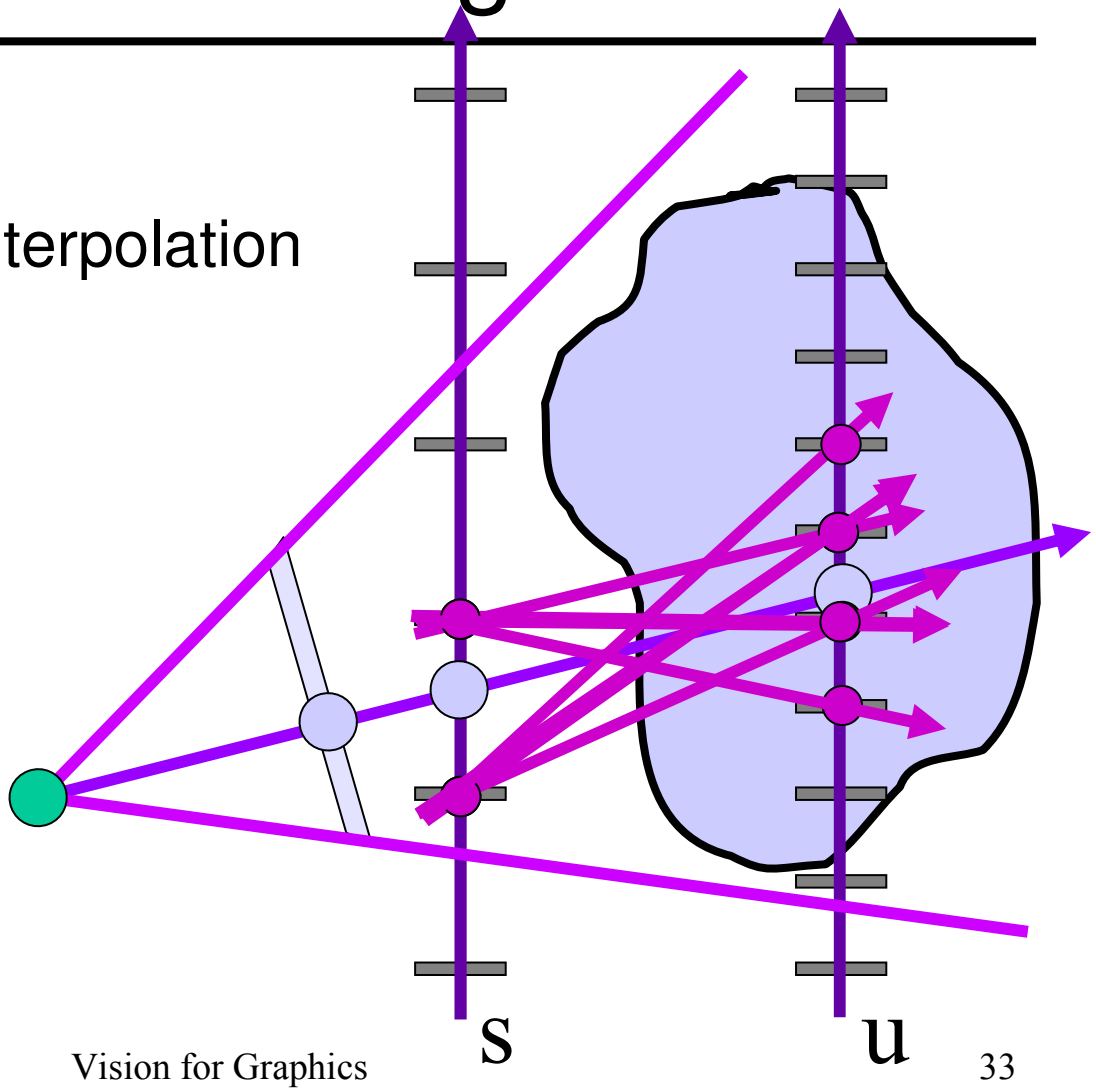

Lumigraph - Rendering

Fast s,t,u,v finding

- scanline interpolate
- texture mapping
- shear warp

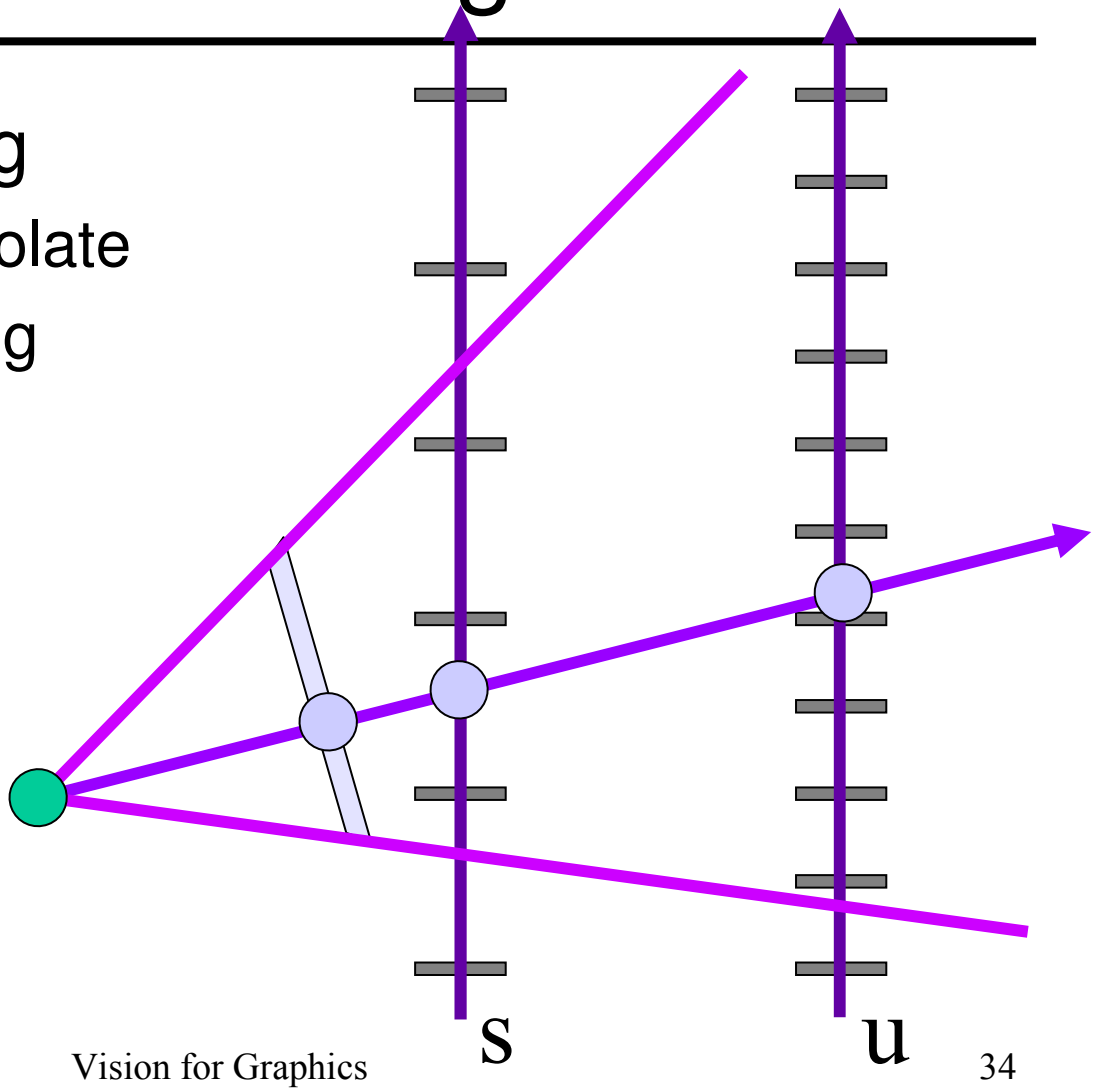

Lumigraph - Ray Space

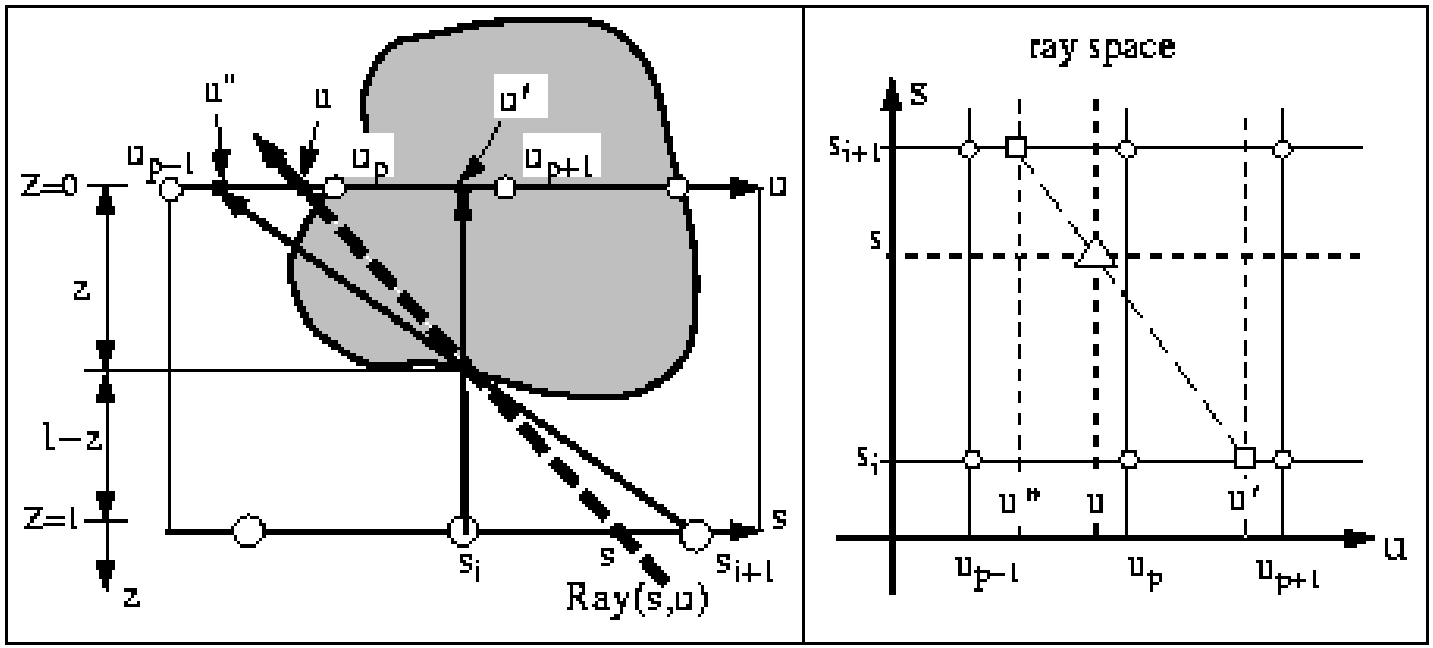

Lumigraph - Ray Space

Image effects:

- parallax
- occlusion
- transparency
- highlights

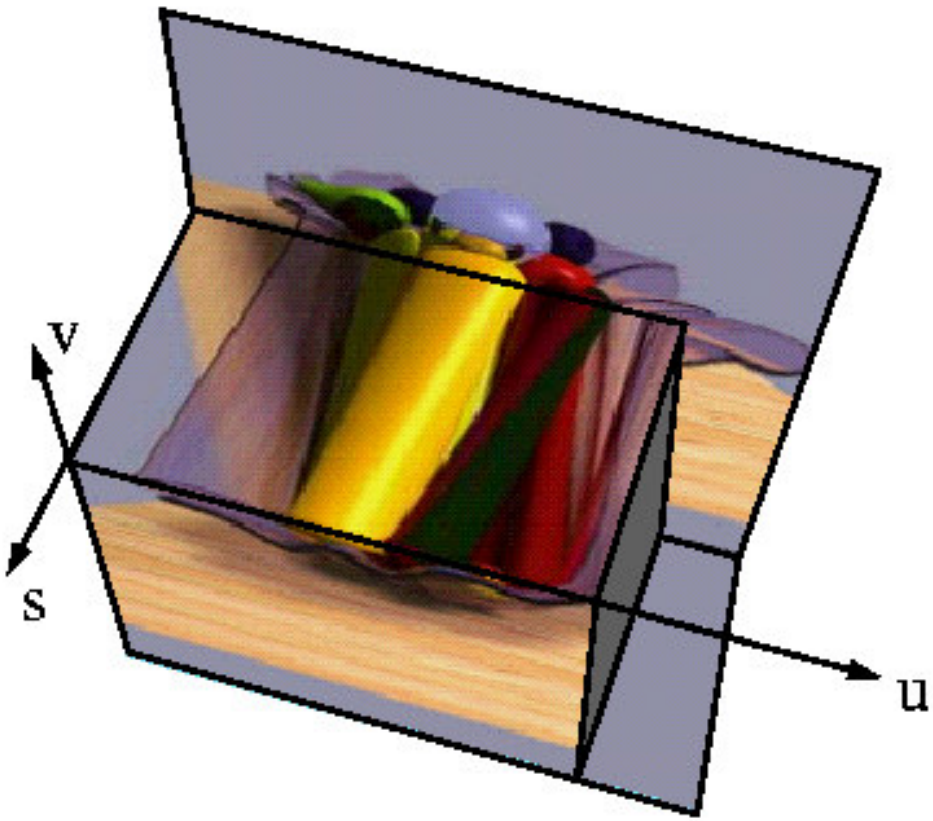

Lumigraph - Demo

Lumigraph

- Lion, Fruit Bowl, Visible Woman, Path Tracing


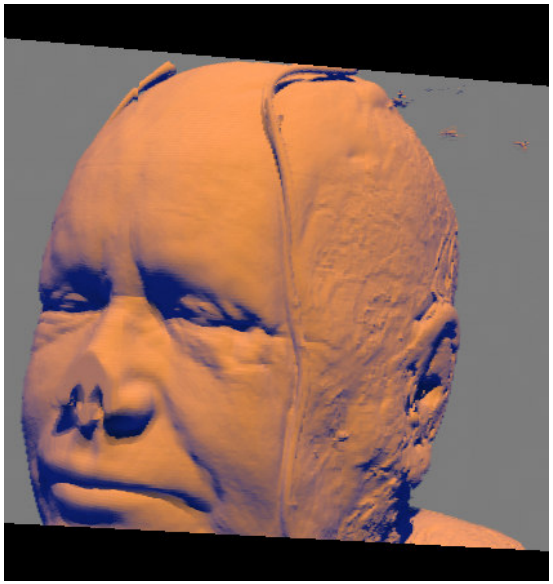
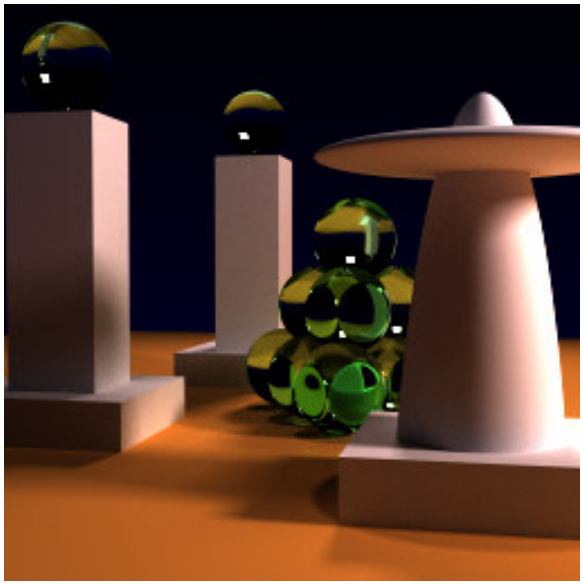

Lightfield - Demo

Digital Michelangelo Project

- Marc Levoy, Stanford University
- Lightfield (“night”) assembled by Jon Shade

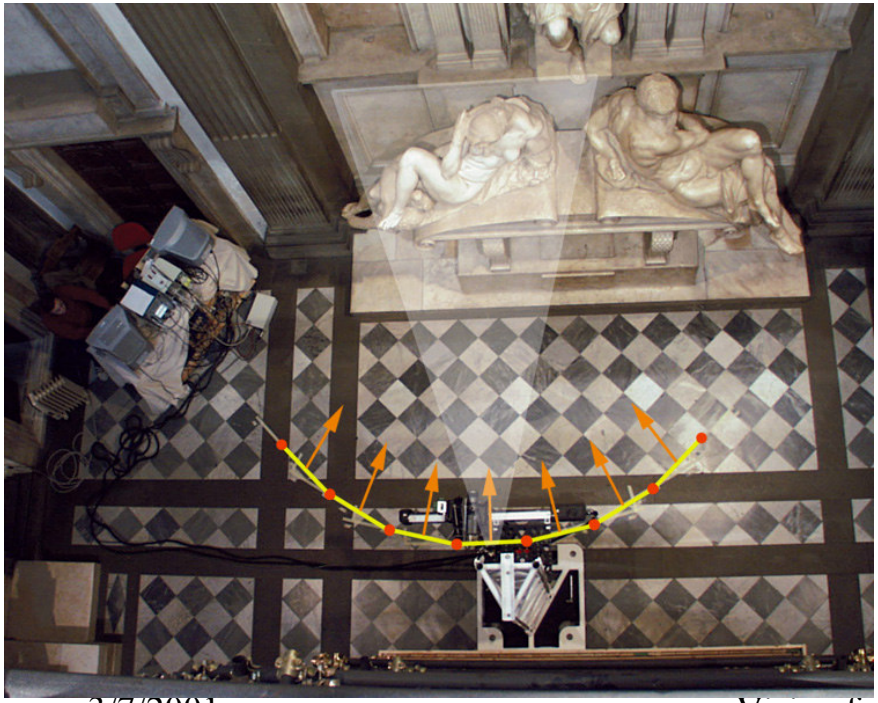

Surface Lightfields

Turn 4D parameterization around

Leverage coherence

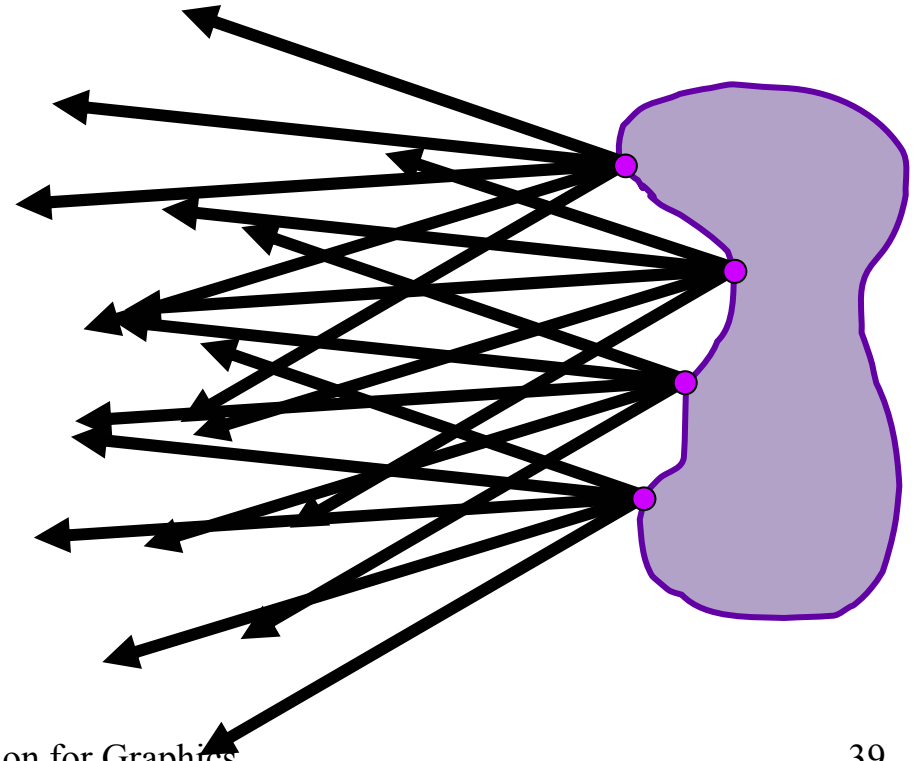

Surface Lightfields

Wood et al, SIGGRAPH 2000

3D Representations

Image is 2D

Lumigraph is 4D

What happened to 3D?

- 3D Lumigraph subset
- Concentric mosaics

3D Lumigraph

One row of s,t plane

- i.e., hold t constant

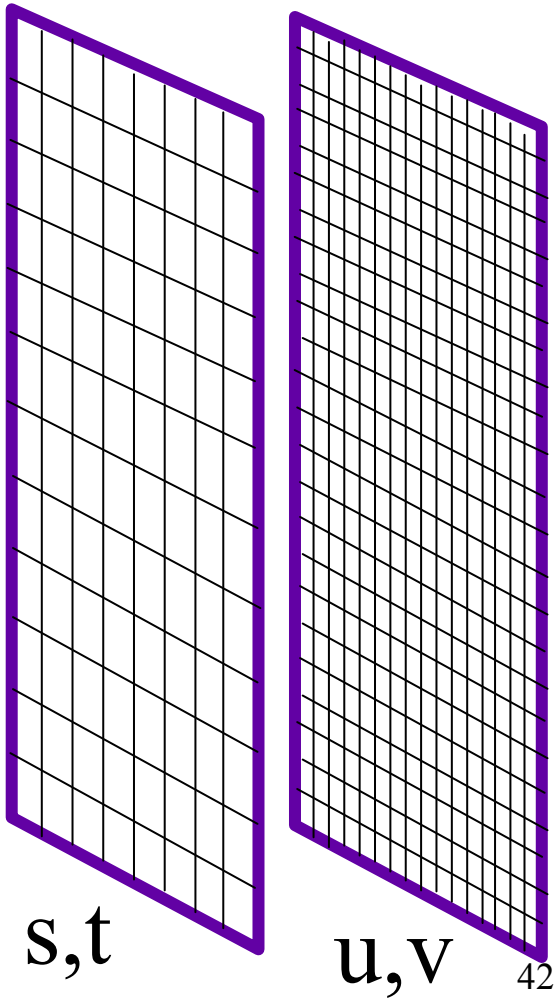

3D Lumigraph

One row of s,t plane

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

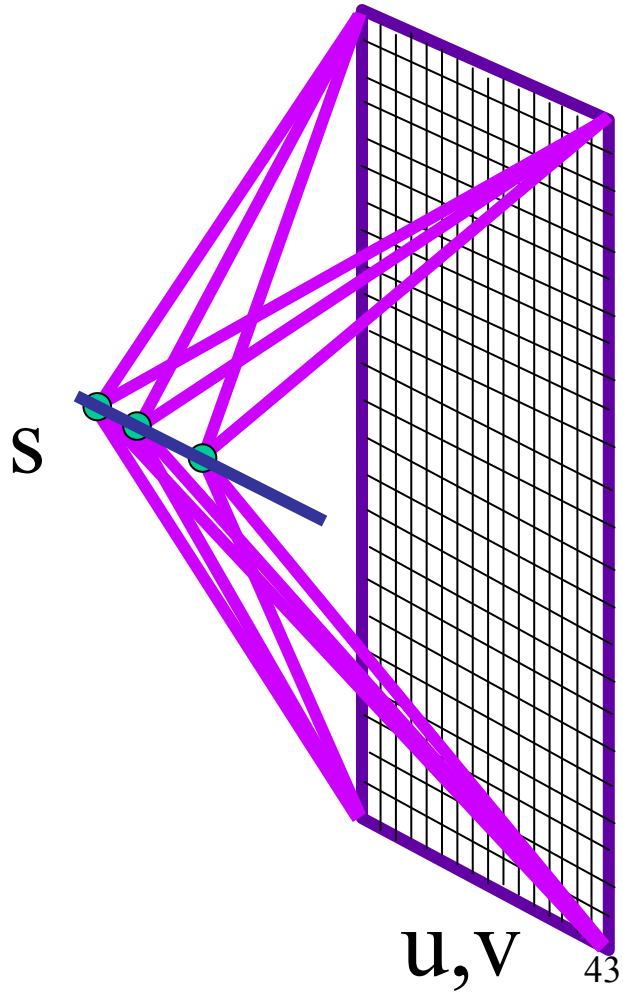

Concentric Mosaics

Replace “row” with “circle” of images

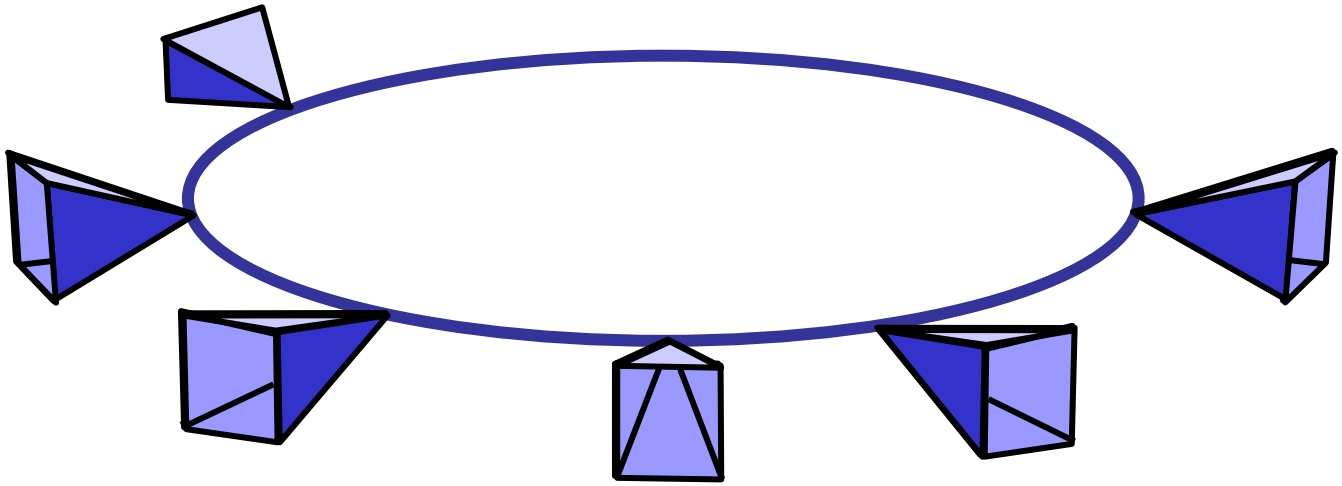

Concentric Mosaics

Concentric Mosaics

From above

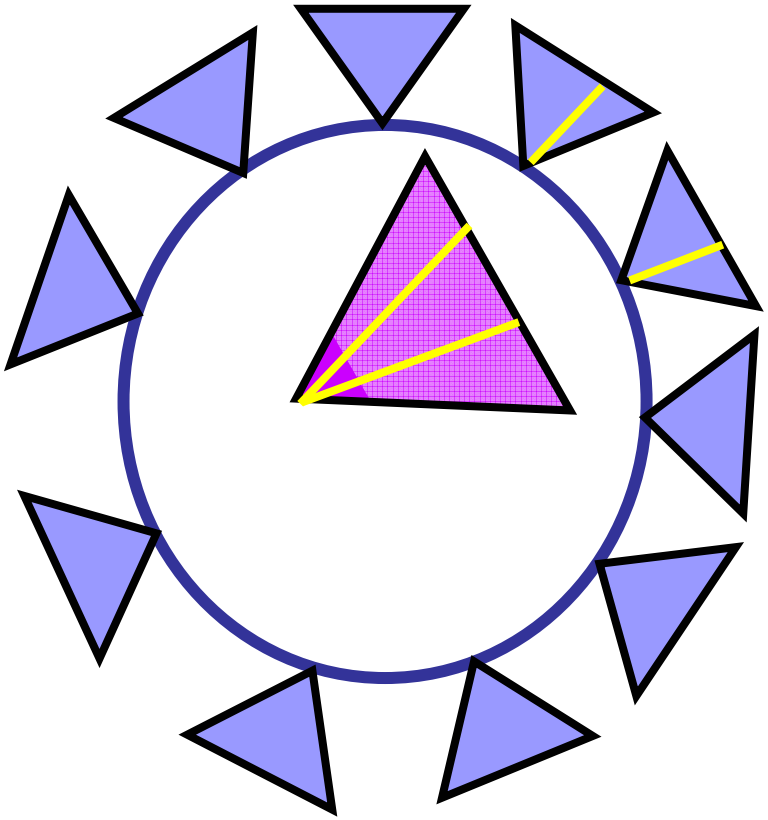

Concentric Mosaics

Depth correction

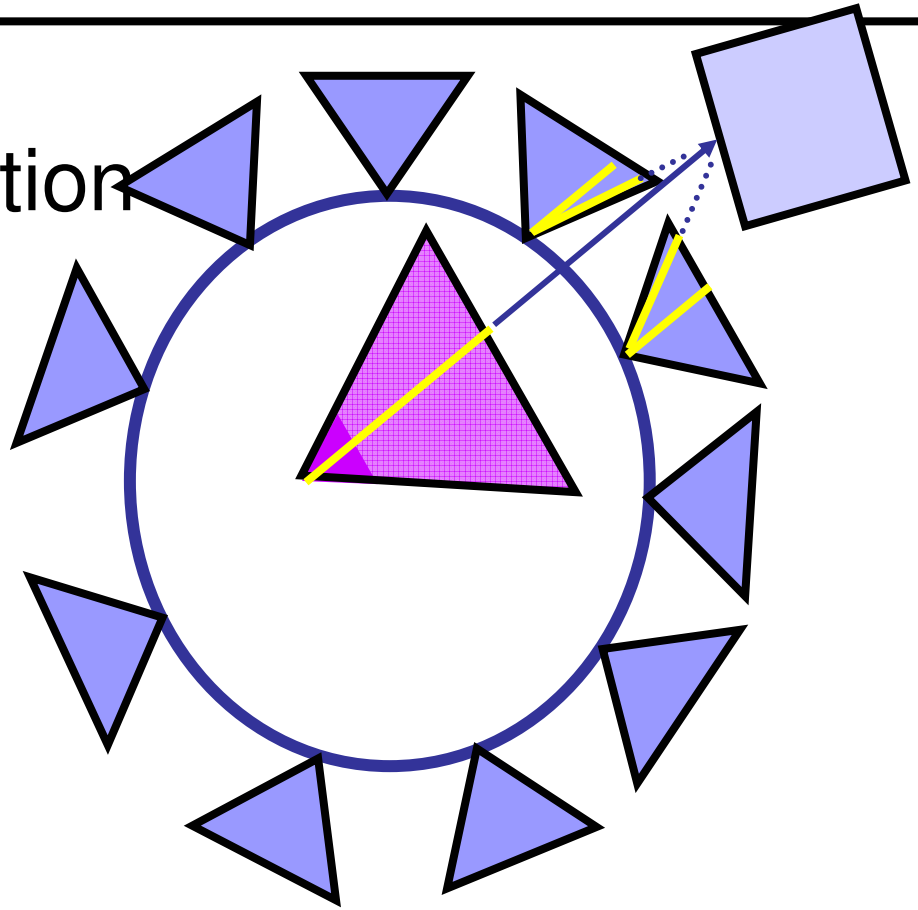

Concentric Mosaics

- Panorama

2.5D Representations

Image is 2D

Lumigraph is 4D

3D

- 3D Lumigraph subset
- Concentric mosaics

2.5D

- Layered Depth Images
- View Dependent Surfaces

Layered Depth Image

2.5 D ?

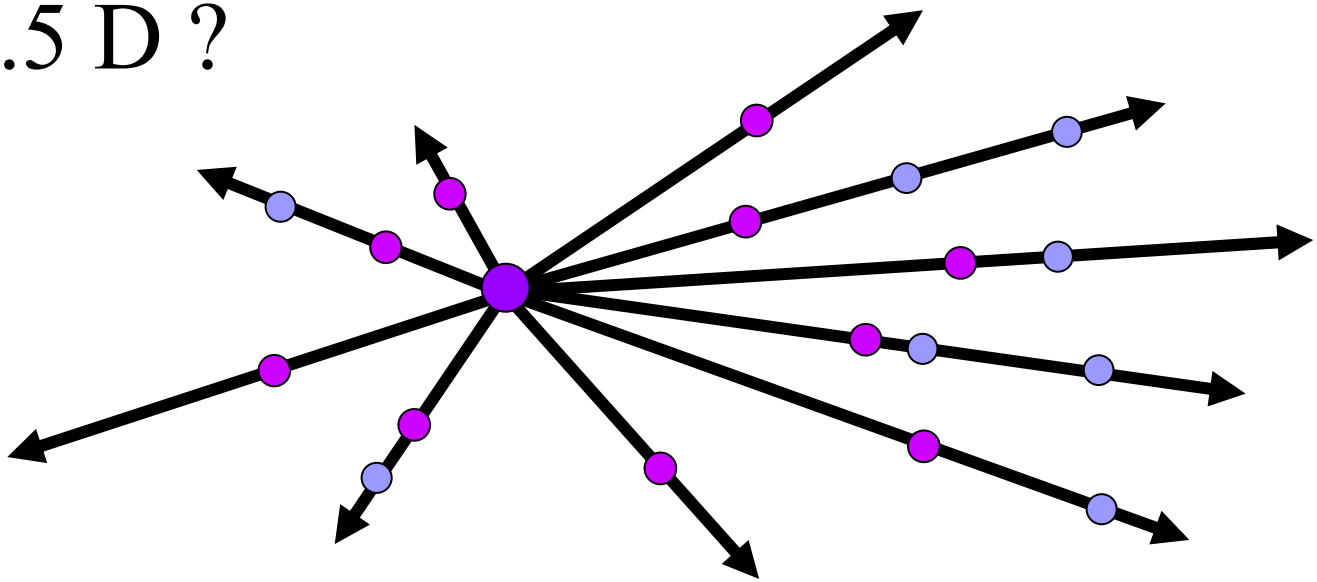

Layered Depth Image

Layered Depth Image

- Rendering from LDI

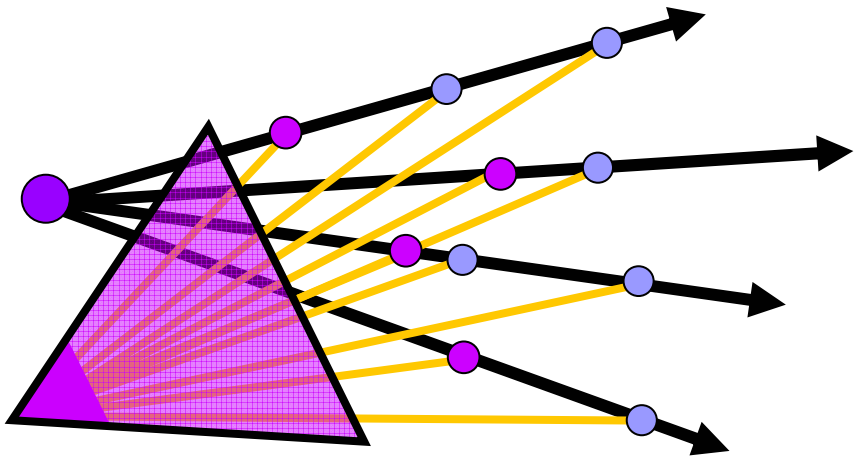

- Incremental in LDI X and Y
- Guaranteed to be in back-to-front order

Layered Depth Image

- Rendering from LDI

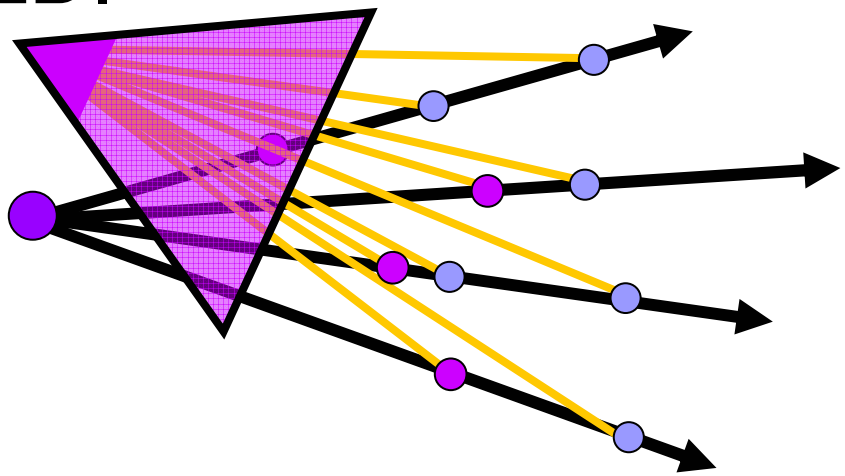

- Incremental in LDI X and Y
- Guaranteed to be in back-to-front order

Layered Depth Image (LDI)

Tiled LDIs

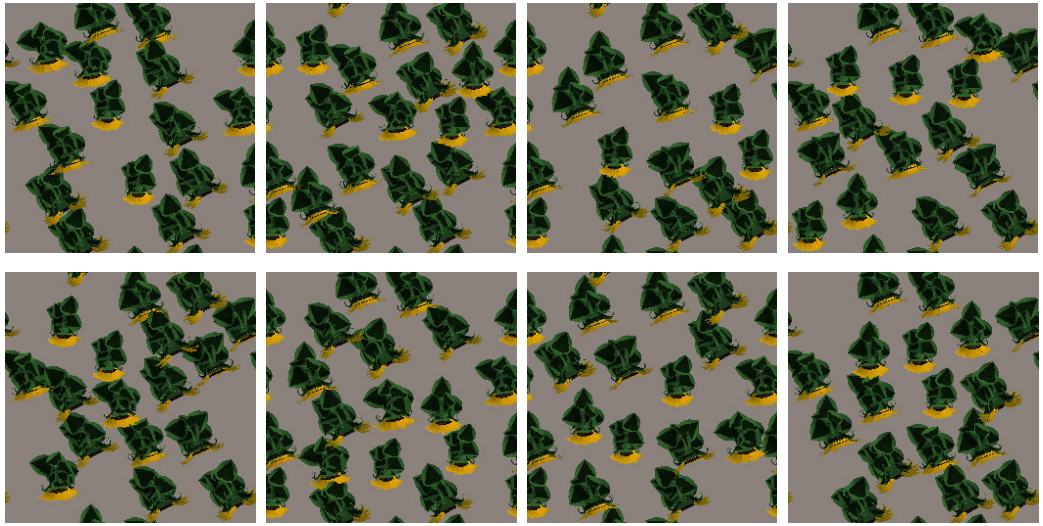

Tiled LDIs

Multiresolution

Directionally dependent

Image Based Models

Modeling and Rendering Architecture from Photographs

Debevec, Taylor, and Malik 1996

Original photograph with marked edges | Recovered model | Model edges projected onto photograph | Synthetic rendering

Image Based Faces

Estimate shape from images

Match metrics to shape

Project video onto shape

- Texture map

Animate

Summary

5D: Plenoptic Function (Ray)

4D: Lumigraph / Lightfield

3D: Lumigraph Subset

3D: Concentric Mosaics

2.5D: Layered Depth Image

2.5D: Image Based Models

2D: Image