Image or Object?

Michael F. Cohen
Microsoft Research

Is this real?

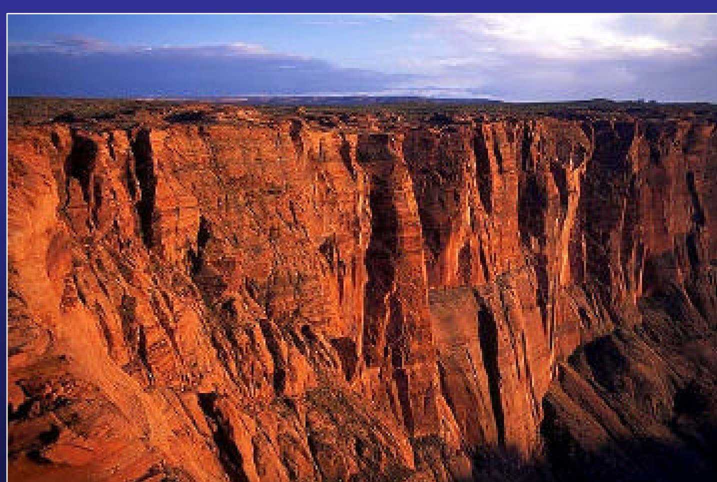

Photo by Patrick Jennings (patrick@synaptic.bc.ca), Copyright 1995, 96, 97 Whistler B. C. Canada

A mental model

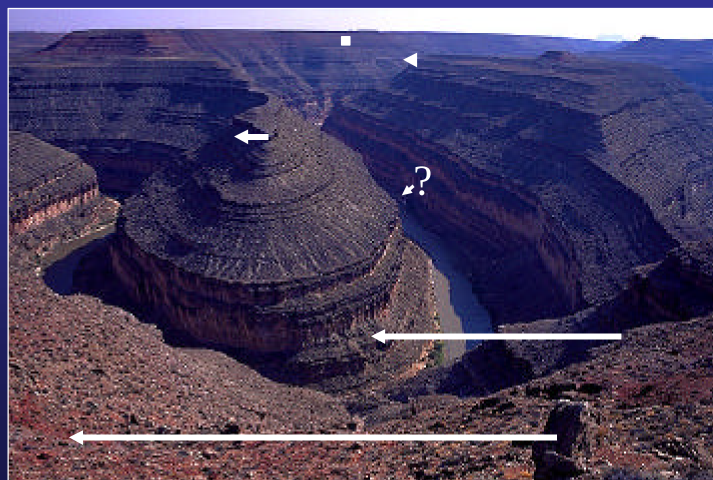

Photo by Patrick Jennings (patrick@synaptic.bc.ca), Copyright 1995, 96, 97 Whistler B. C. Canada

Sue Vision

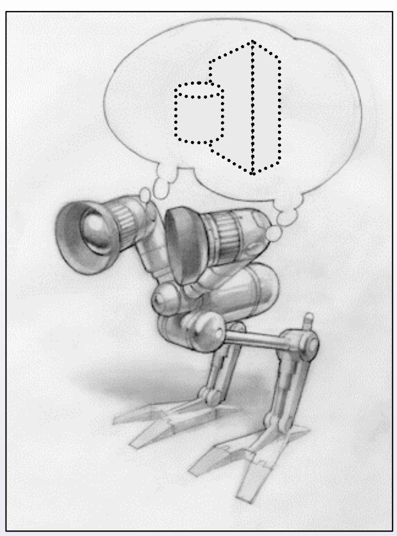

Is this real?

Michael F. Cohen, Shenchang Eric Chen, John R. Wallace, and Donald P. Greenberg

1988, *A Progressive Refinement Approach to Fast Radiosity Image Generation*.

Bob Graphics

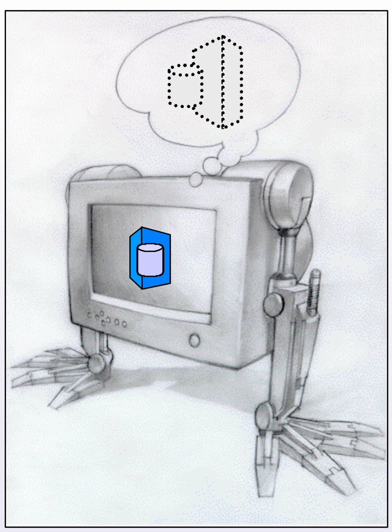

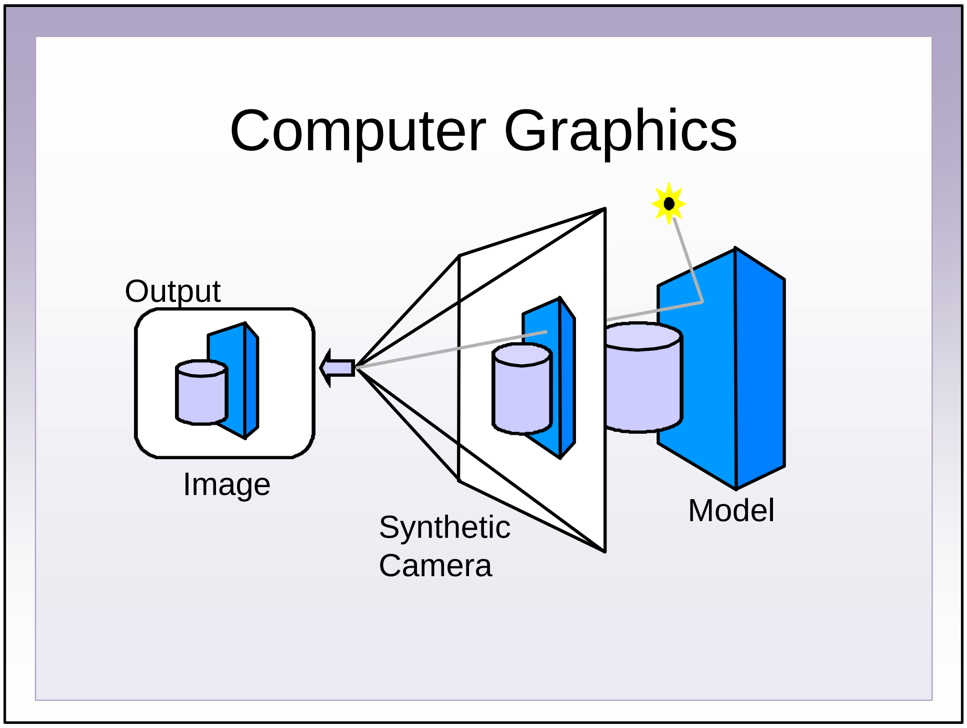
Computer Graphics
Output
Image
Synthetic
Camera
Model

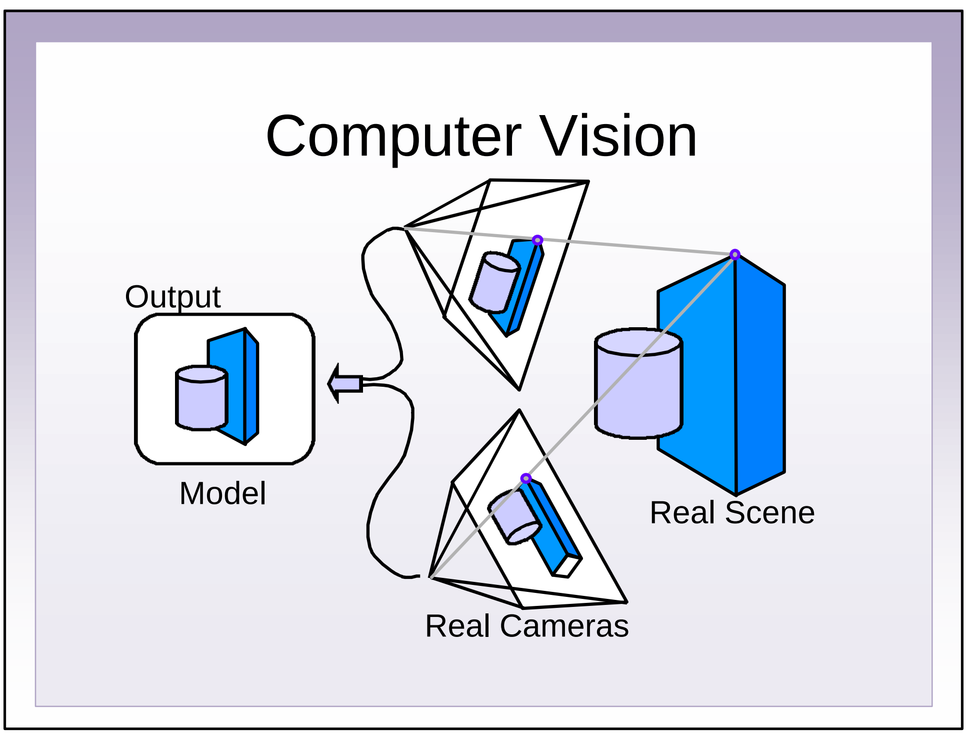
Computer Vision
Output
Model
Real Scene
Real Cameras

Computer Graphics Meets Computer Vision

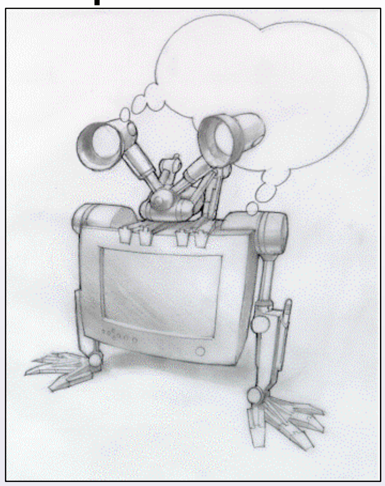

Computer Graphics Meets Computer Vision

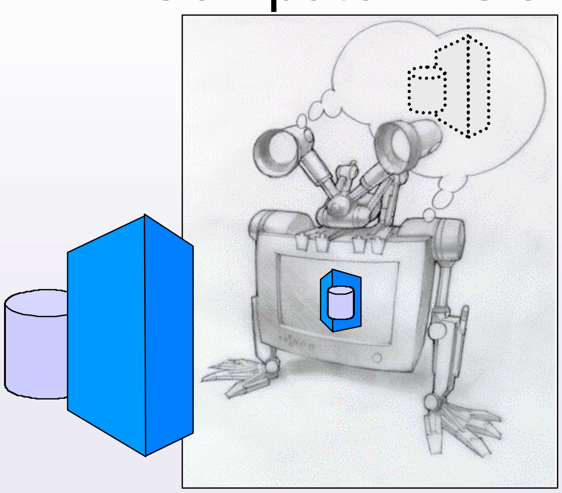

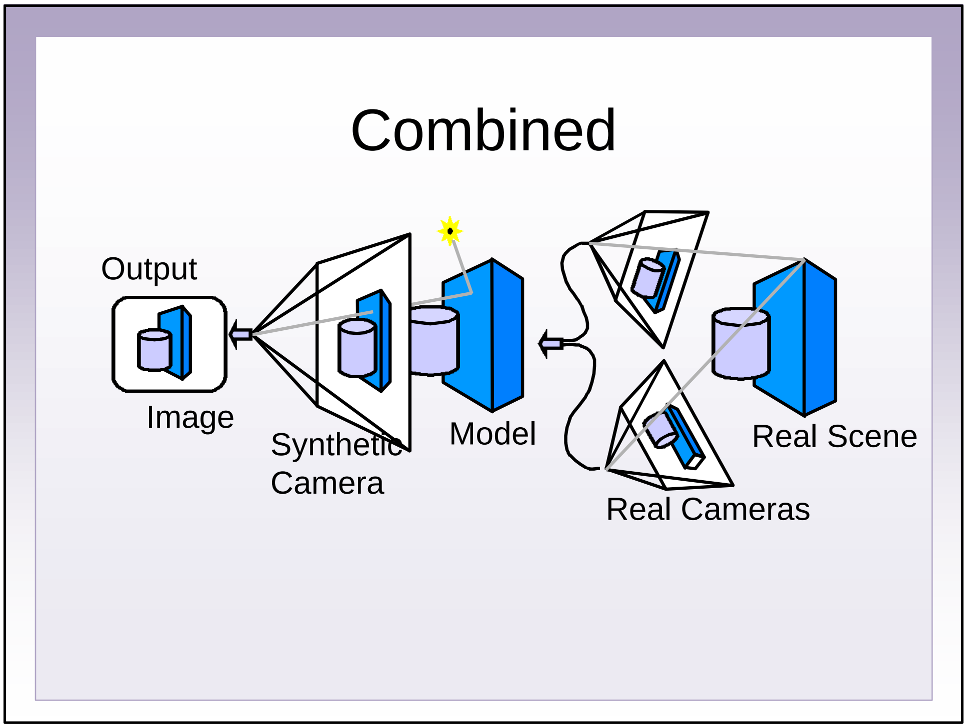
Combined
Output
Image
Synthetic
Camera
Model
Real Scene
Real Cameras

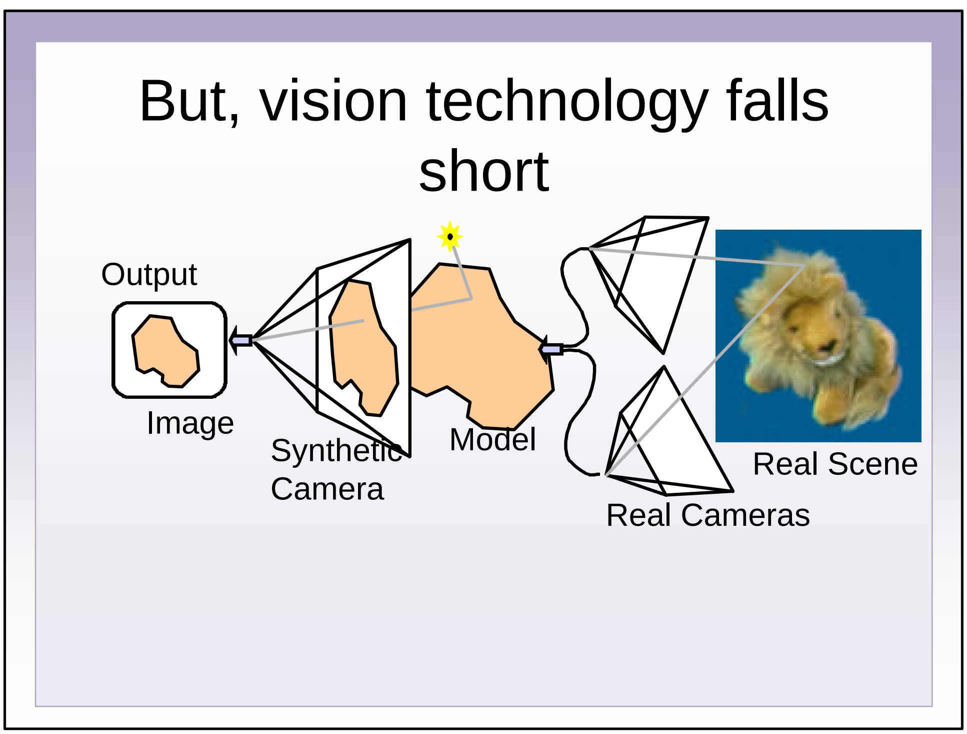
But, vision technology falls short
Output
Image
Synthetic
Camera
Model
Real Scene
Real Cameras

… and so does graphics.

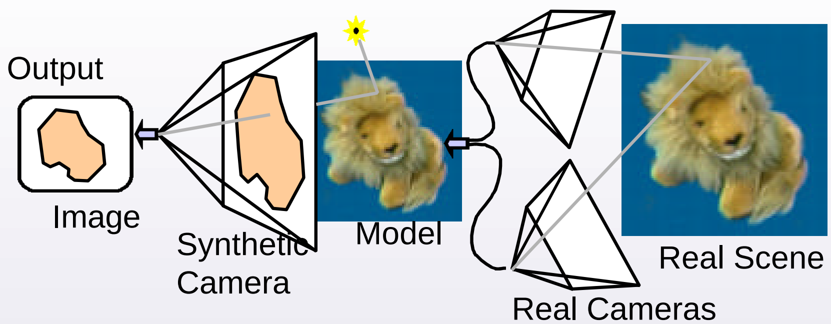

Image Based Rendering

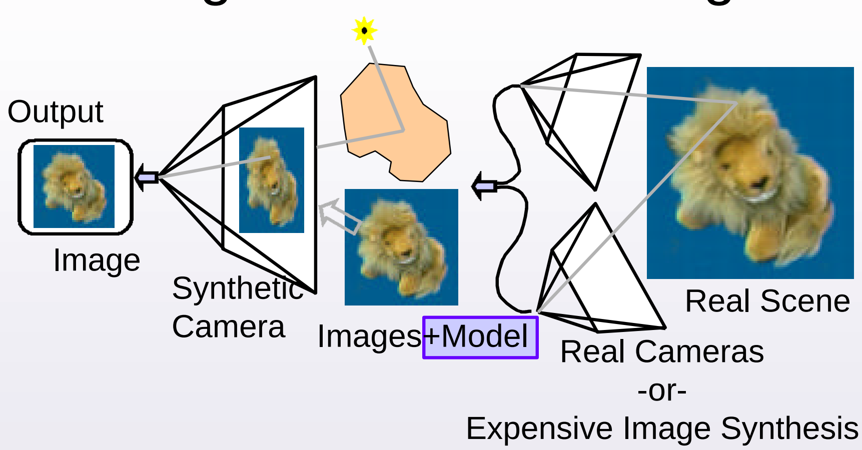

Ray

- Constant radiance
  - time is fixed

- 5D
  - 3D position
  - 2D direction

All Rays

- Plenoptic Function
  - all possible images
  - too much stuff!

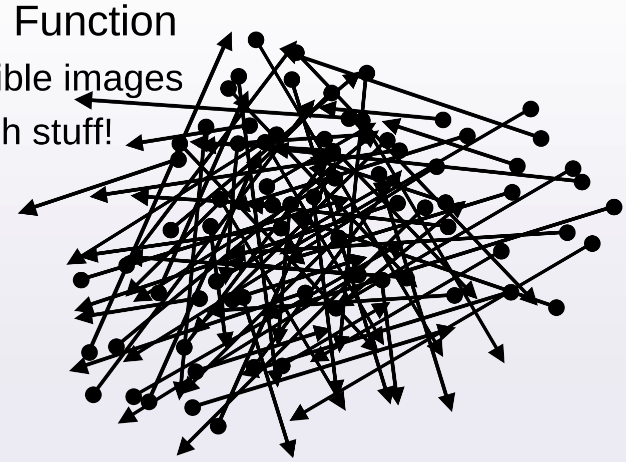

Line

- Infinite line

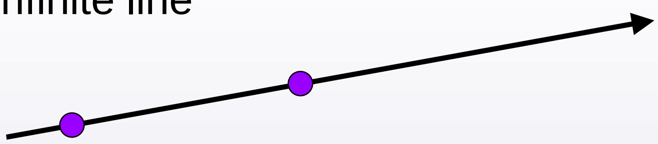

- 4D
  - 2D direction
  - 2D position

Ray

- Discretize

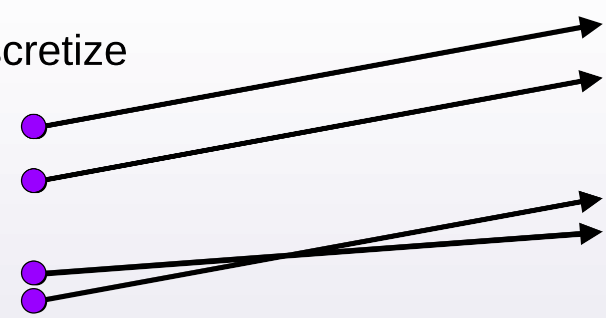

- Distance between 2 rays
  - Which is closer together?

Image

- What is an image?

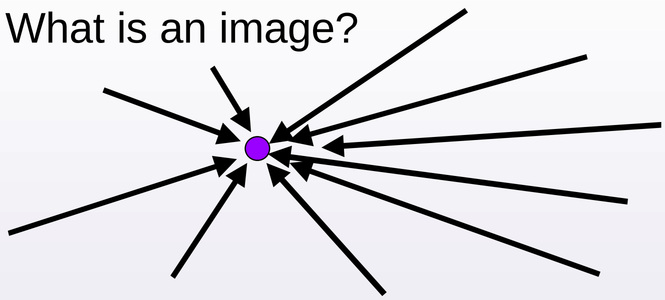

- All rays through a point
  - Panorama?

Image

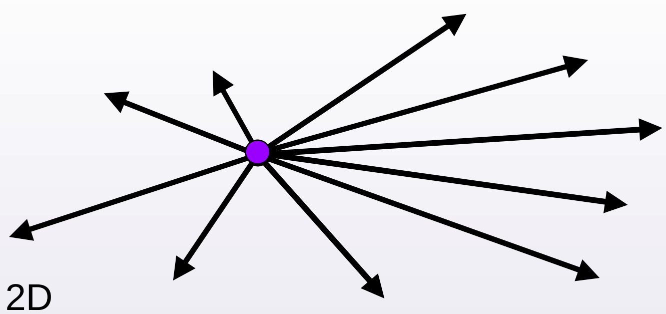

- 2D
  - position of rays has been fixed
  - direction remains

Image

- Image plane

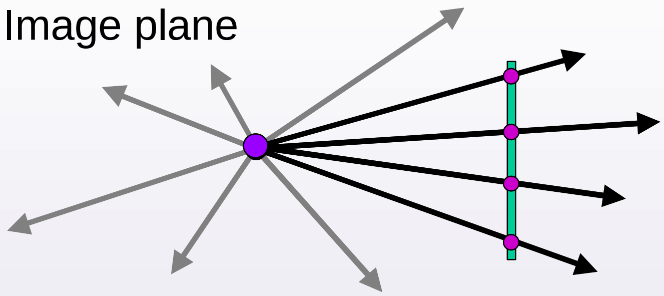

- 2D
  - position

Image

- Image plane

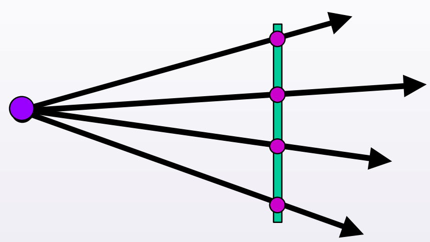

- 2D
  - position

Object

- Light leaving towards "eye"

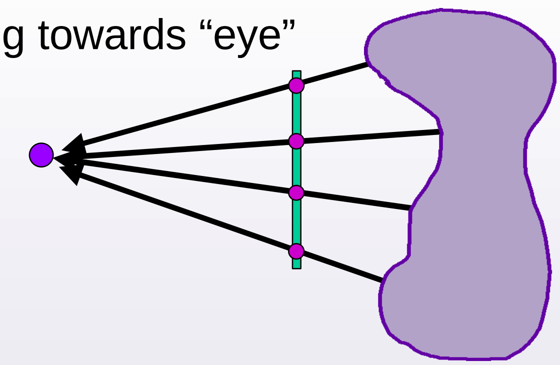

- 2D
  - just dual of image

Object

- All light leaving object

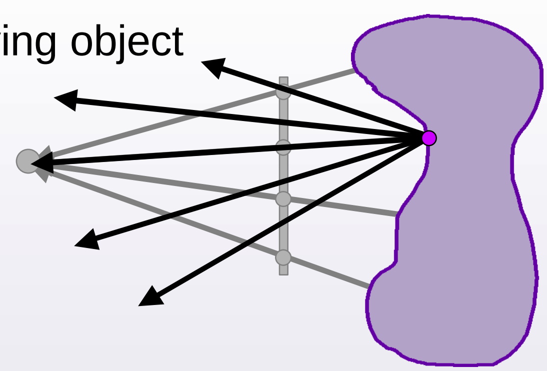

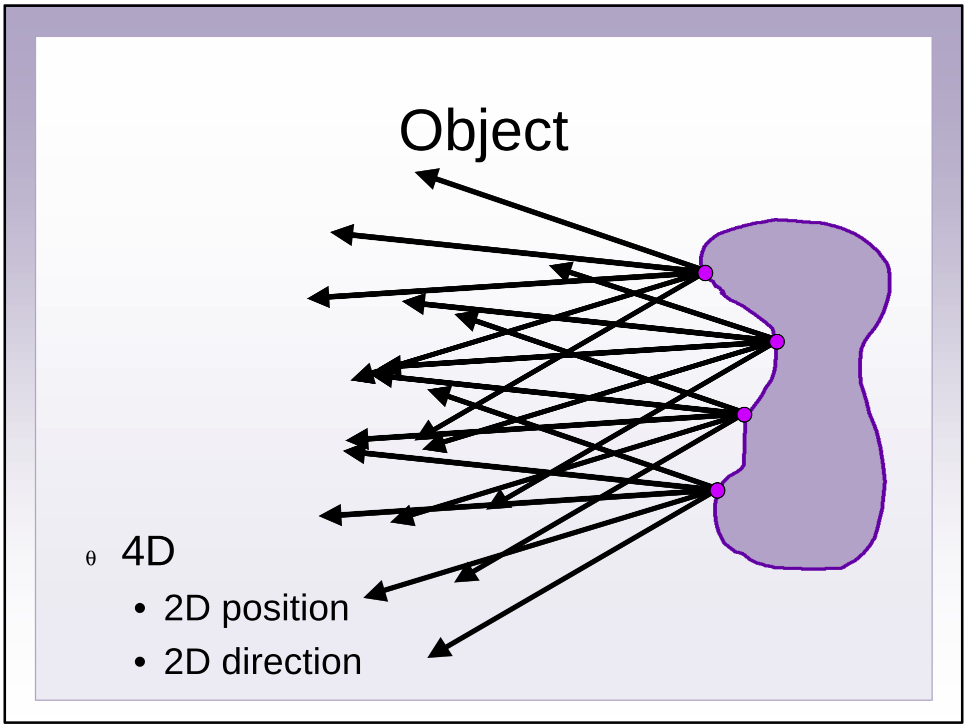
Object
θ 4D
• 2D position
• 2D direction

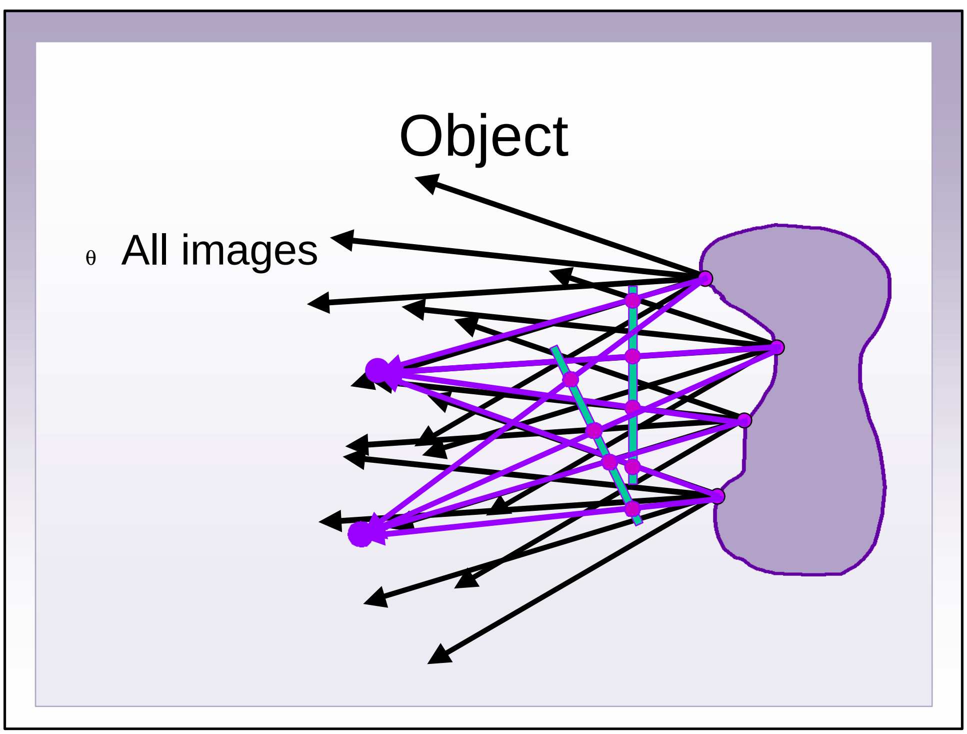
Object
θ All images

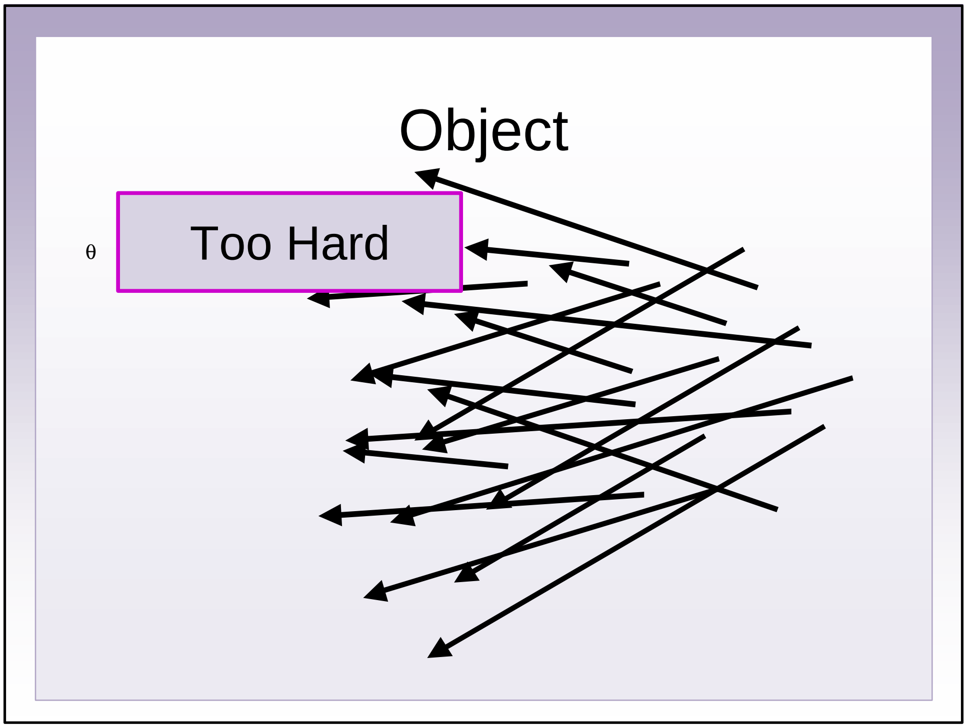
Object
Too Hard

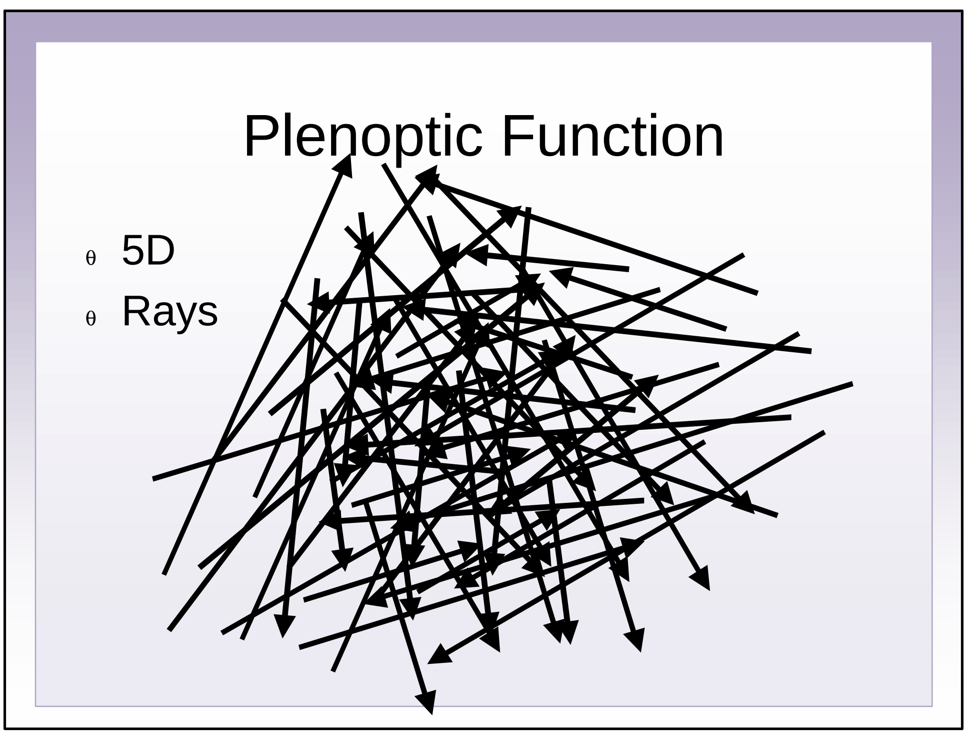
Plenoptic Function
5D
Rays

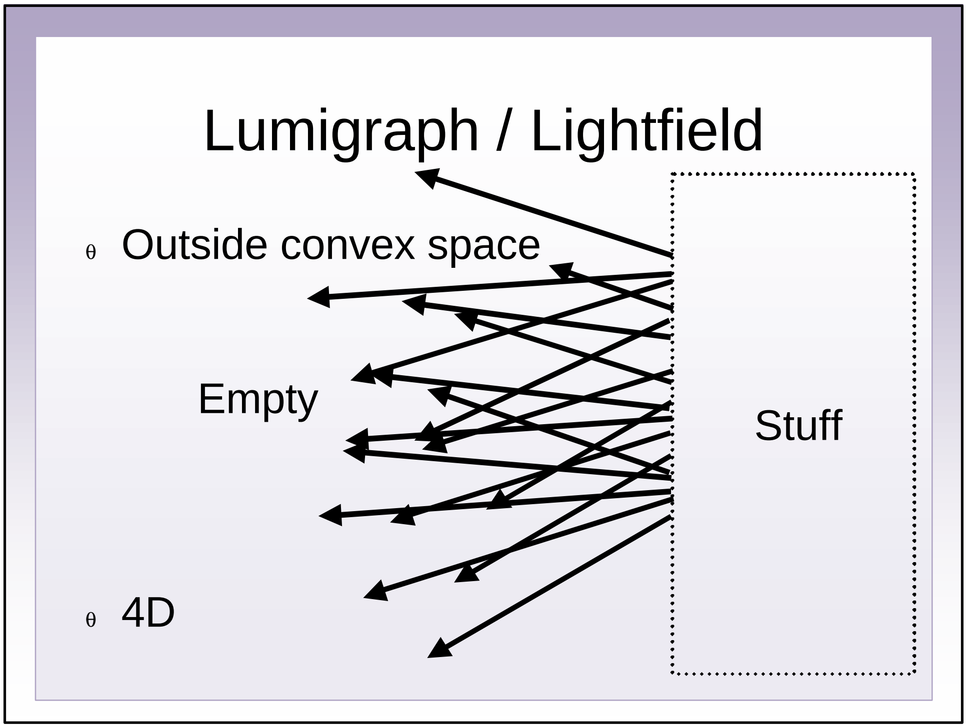
Lumigraph / Lightfield
θ Outside convex space
Empty
Stuff
θ 4D

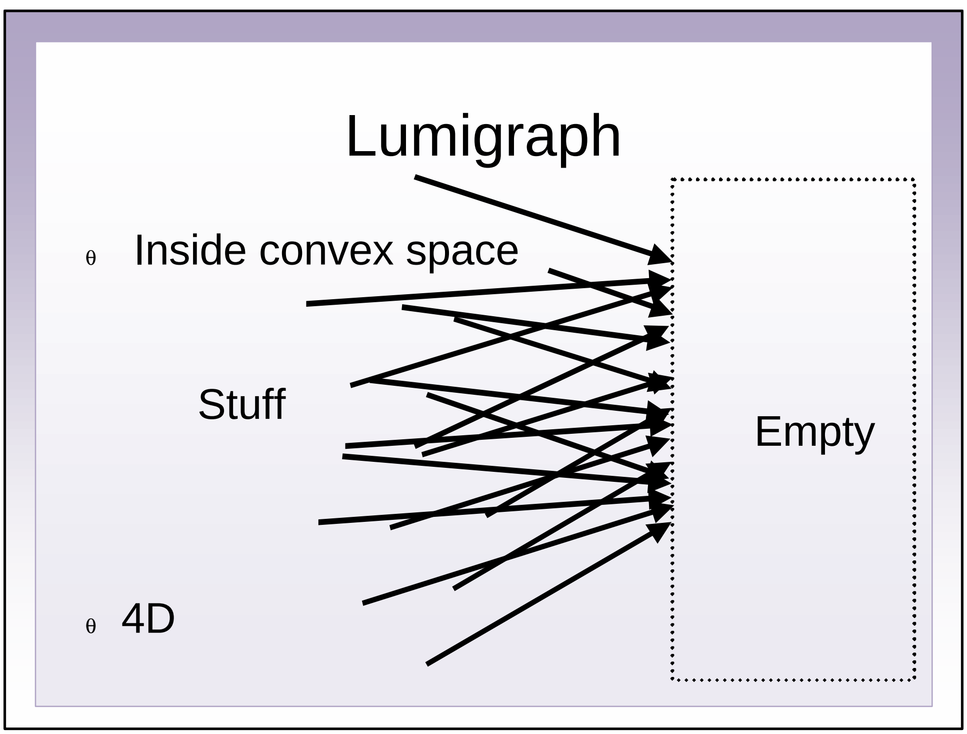
Lumigraph
θ Inside convex space
Stuff
Empty
θ 4D

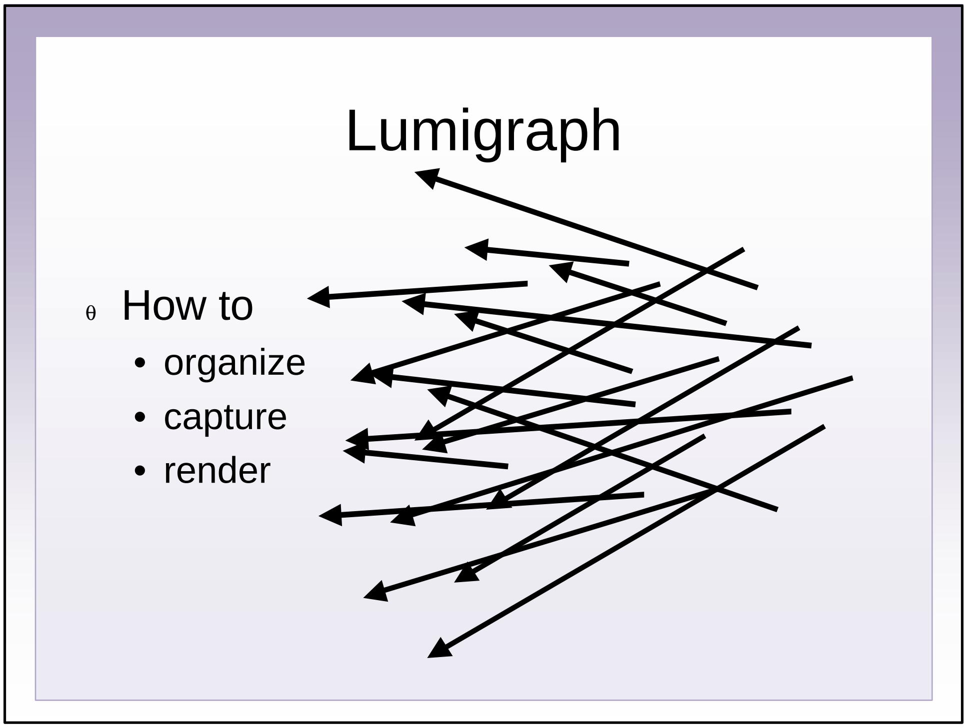
Lumigraph
θ How to
• organize
• capture
• render

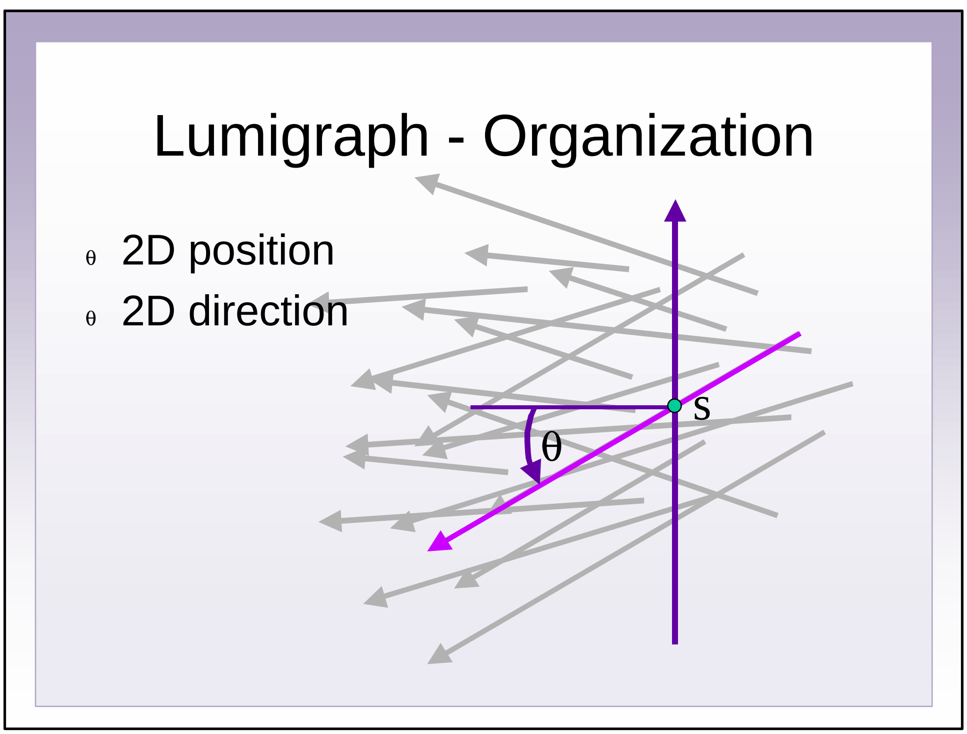
Lumigraph - Organization
θ 2D position
θ 2D direction
s
θ

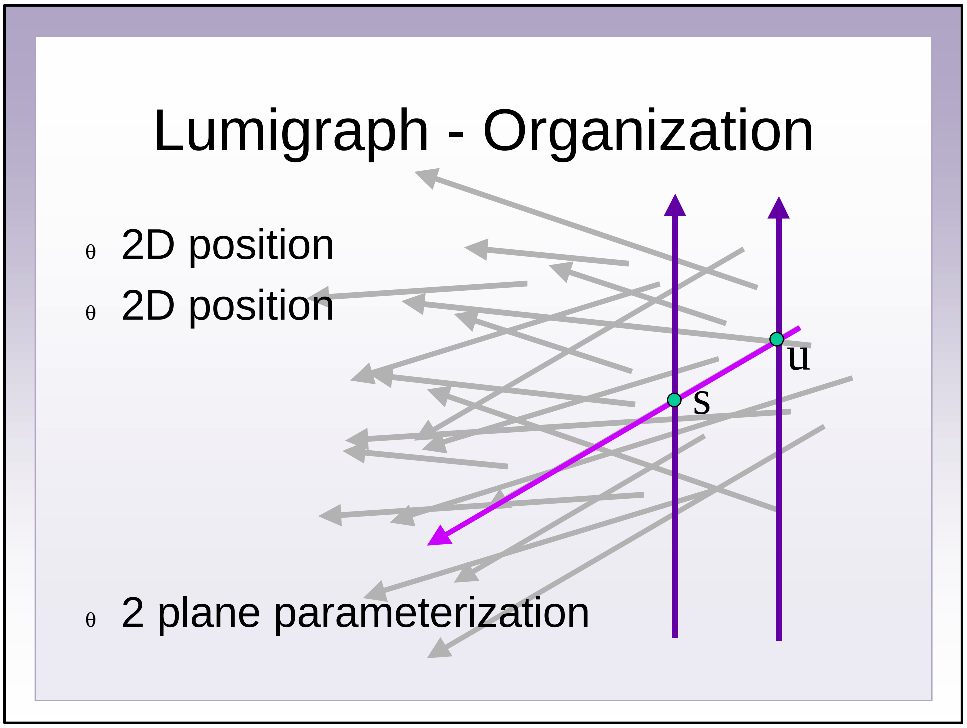
Lumigraph - Organization
2D position
2D position
2 plane parameterization
s
u

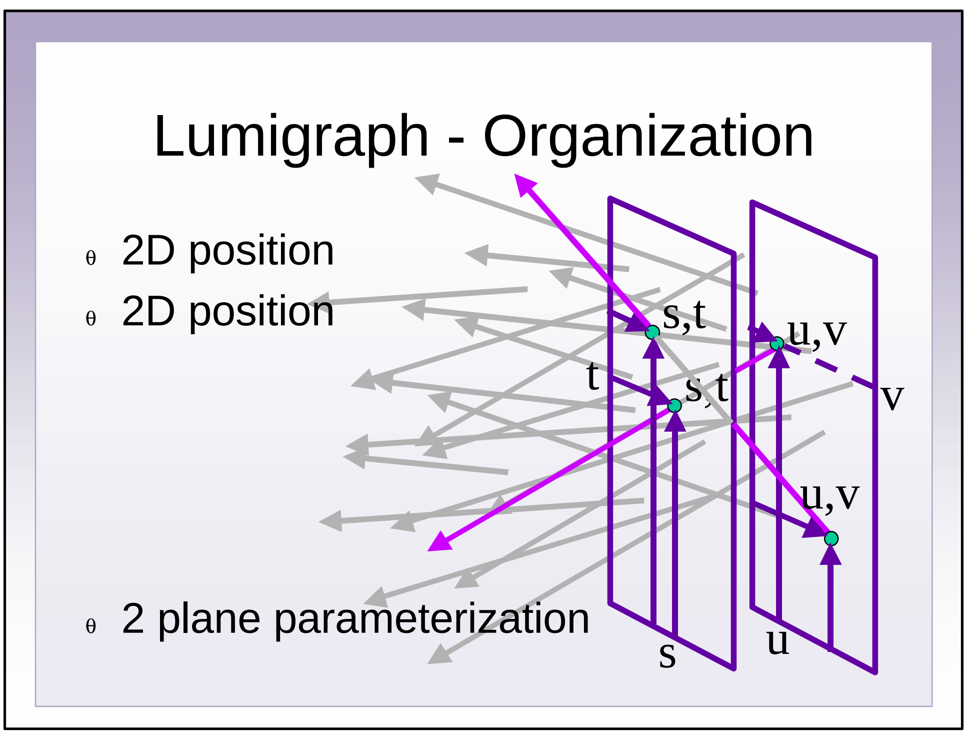
Lumigraph - Organization
2D position
2D position
2 plane parameterization
s,t
s,t
t
u,v
u,v
v
s
u

Lumigraph - Organization

- Hold s,t constant
- Let u,v vary
- An image

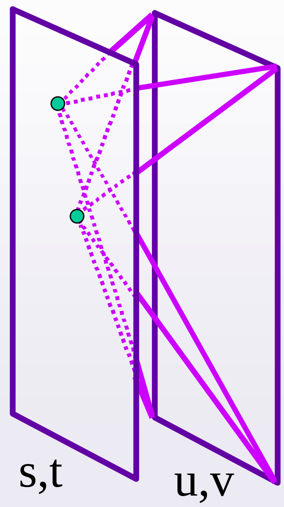

Lumigraph - Organization

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

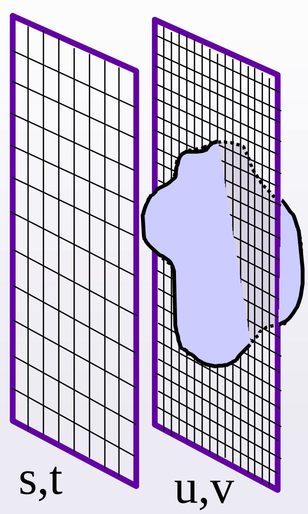

Lumigraph - Capture

θ Idea 1

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

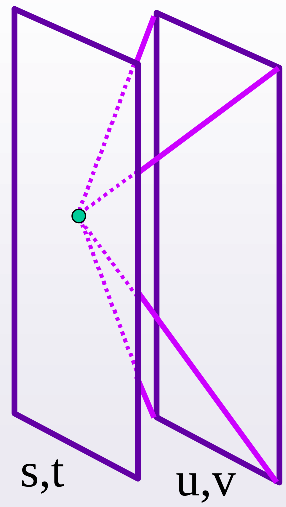

Lumigraph - Capture

θ Idea 2

- Move camera anywhere
- Rebinning
  - see Lumigraph paper

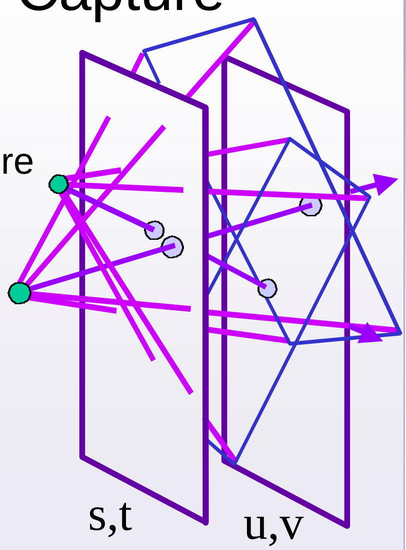

Lumigraph - Rendering

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

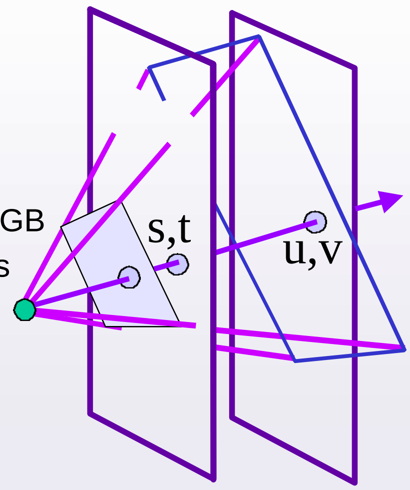

Lumigraph - Rendering

- For each output pixel
  - determine s,t,u,v

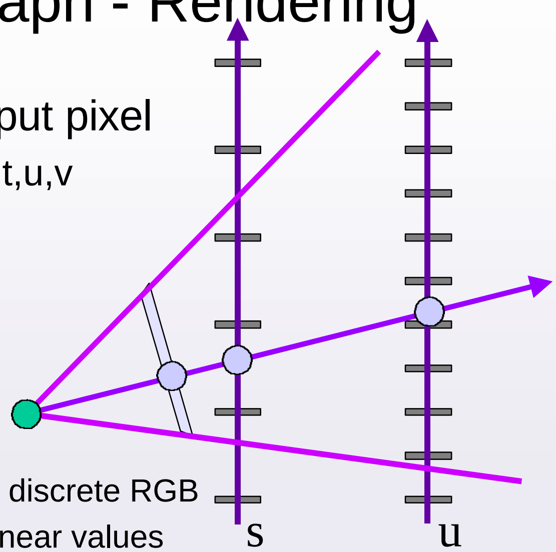

  - either
    - use closest discrete RGB
    - interpolate near values

Lumigraph - Rendering

- Nearest
  - closest s
  - closest u
  - draw it
- Blend 16 nearest
  - quadrilinear interpolation

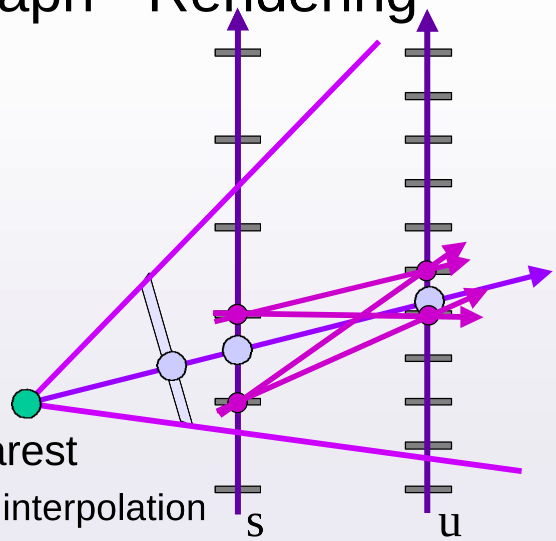

Lumigraph - Rendering

- Depth Correction
  - closest s
  - intersection with "object"
  - best u
  - closest u

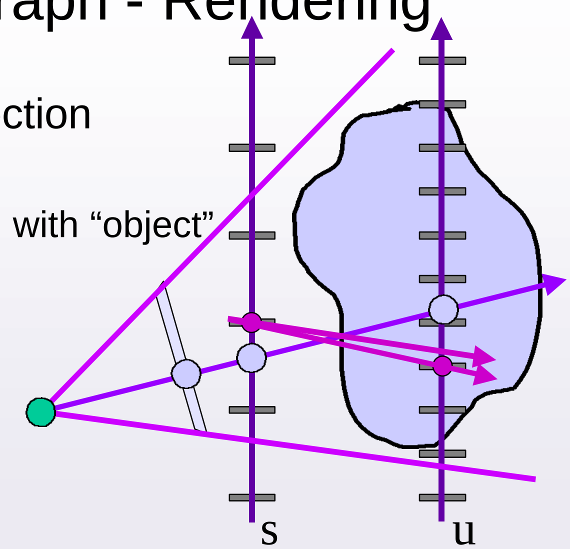

Lumigraph - Rendering

- Depth Correction
  - quadralinear interpolation
  - new “closest”
  - like focus

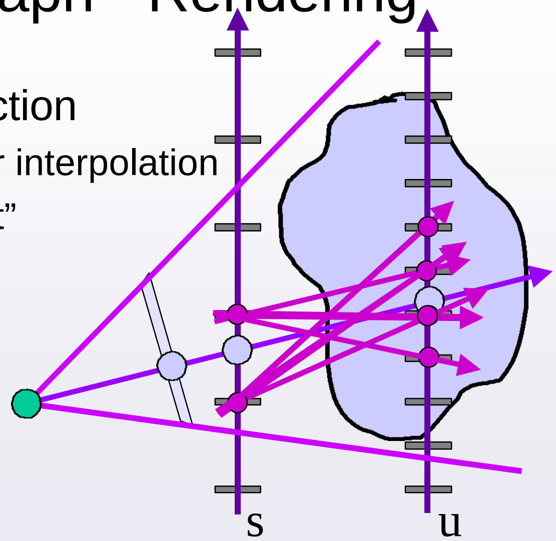

Lumigraph - Rendering

- Fast s,t,u,v finding
  - scanline interpolate
  - texture mapping
  - shear warp

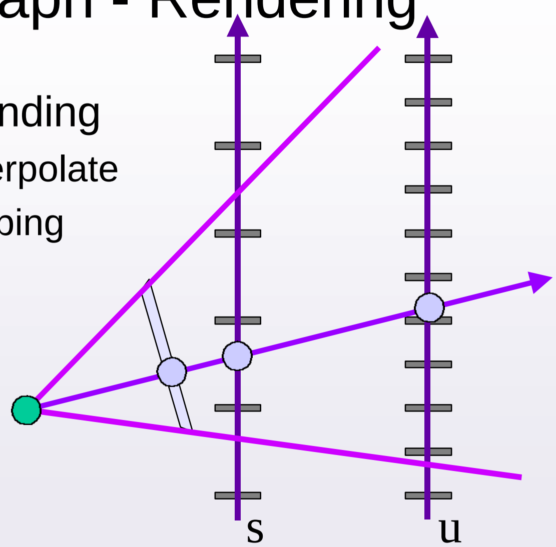

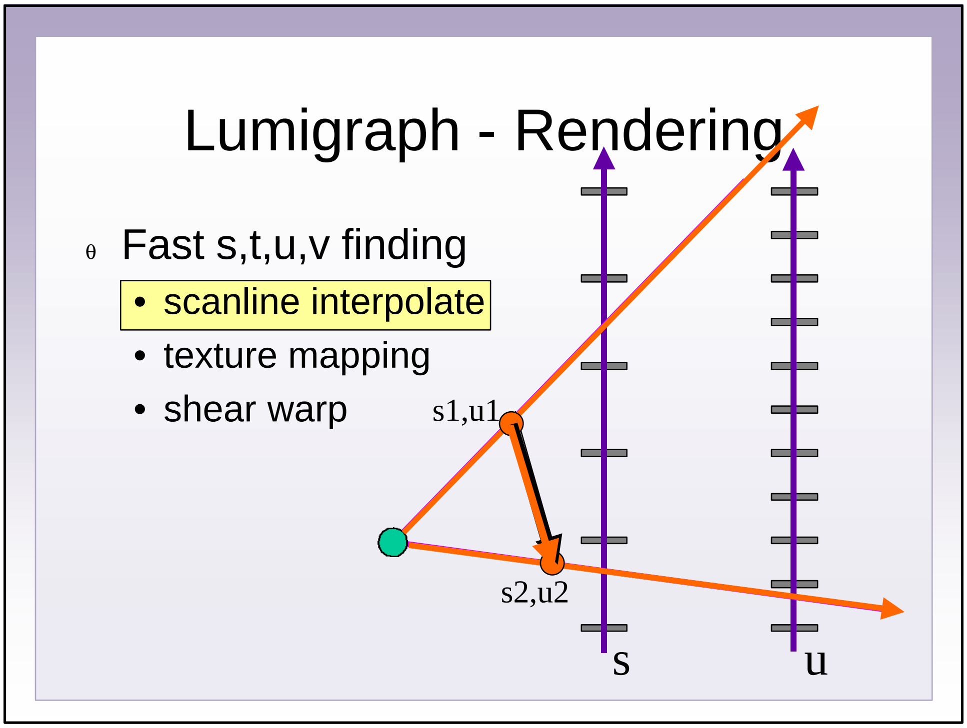
Lumigraph - Rendering
Fast s,t,u,v finding
• scanline interpolate
• texture mapping
• shear warp
s1,u1
s2,u2
s
u

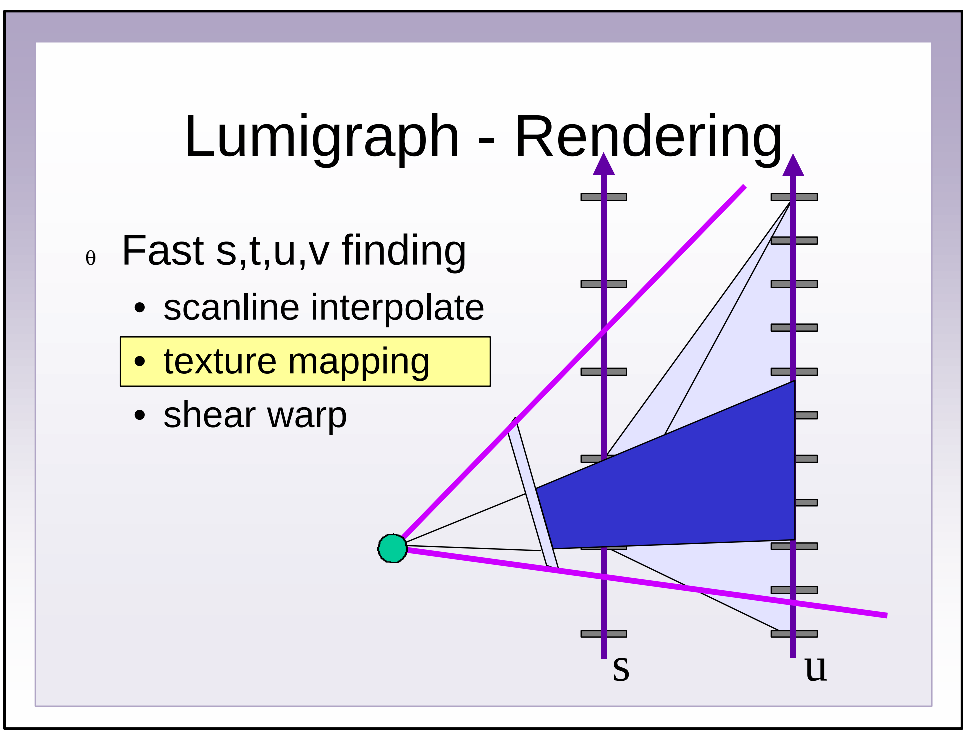
Lumigraph - Rendering
Fast s,t,u,v finding
• scanline interpolate
• texture mapping
• shear warp
s
u

Lumigraph - Rendering

- Fast s,t,u,v finding
  - scanline interpolate
  - texture mapping
  - shear warp

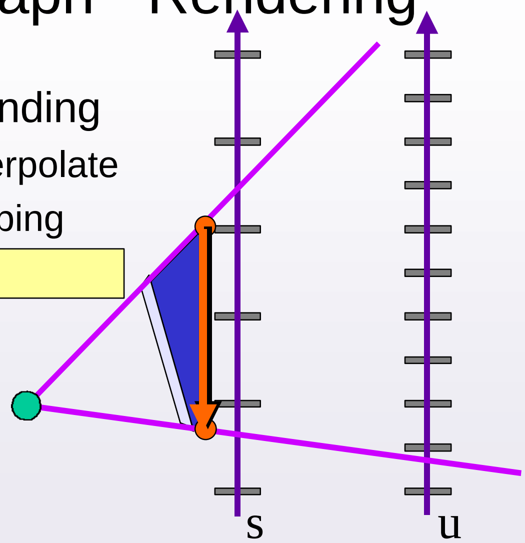

Lumigraph - Demo

- Lumigraph
  - Lion, Fruit Bowl, Visible Woman, Path Tracing

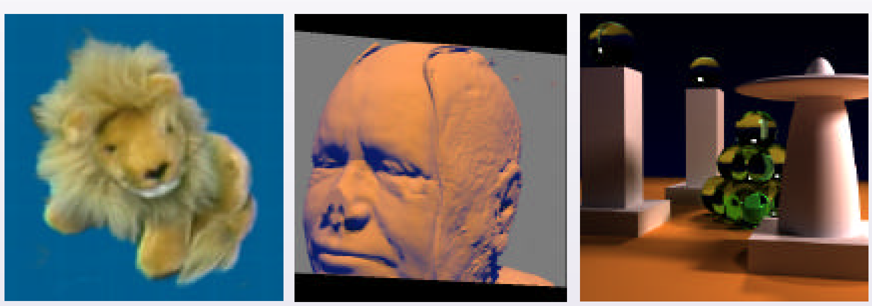

Lightfield - Demo

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

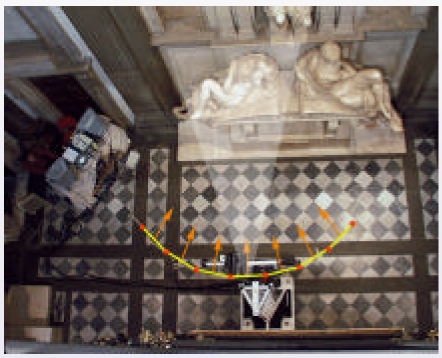

3D Representations

- 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

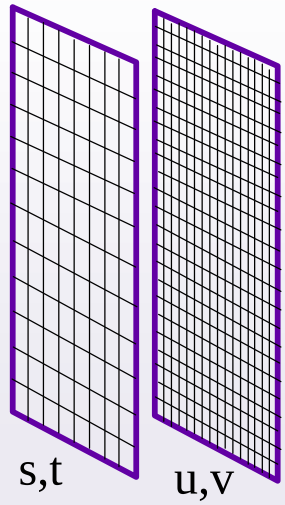

3D Lumigraph

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

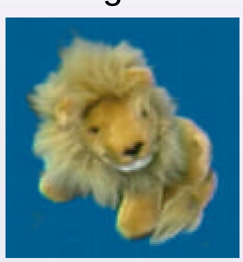

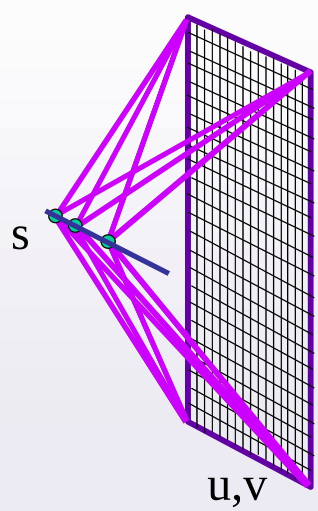

Concentric Mosaics

- Replace “row” with “circle” of images

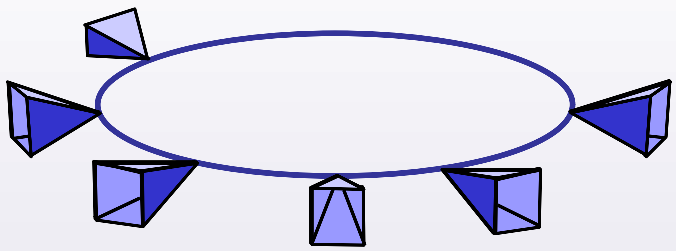

Concentric Mosaics

Concentric Mosaics

- From above

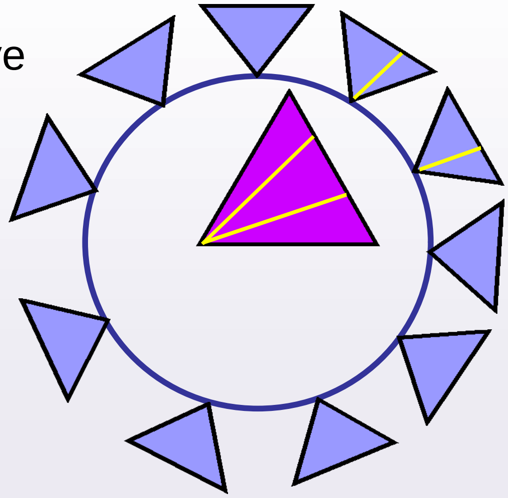

Concentric Mosaics

- From above

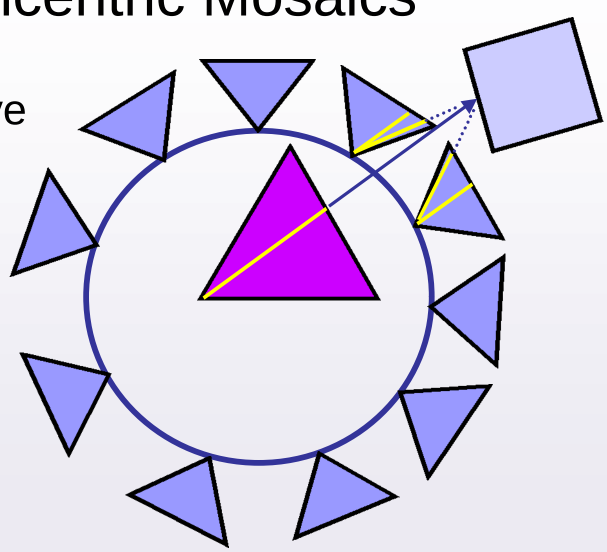

Concentric Mosaics

- Panorama

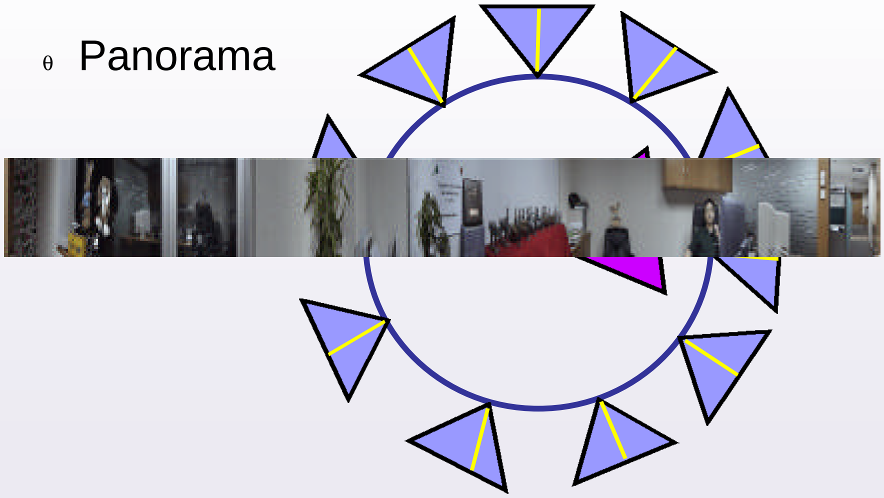

2.5D Representations

- 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 ?

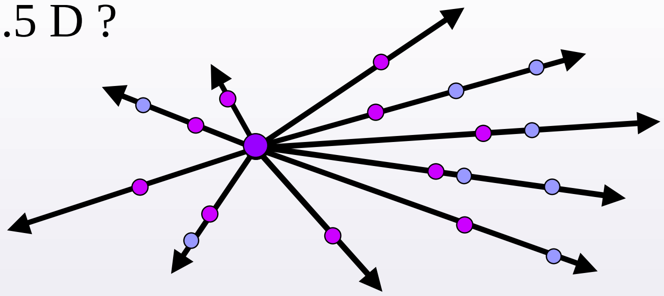

Layered Depth Image

Layered Depth Image

- Rendering from LDI

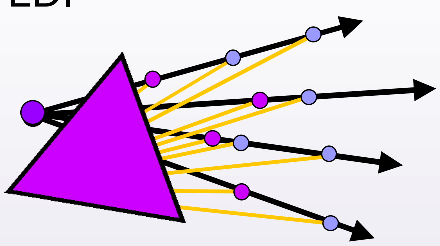

- Incremental in LDI X and Y
- Guaranteed to be in back-to-front order

Layered Depth Image

- Rendering from LDI

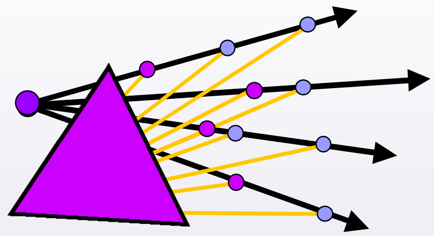

• Incremental in LDI X and Y
• Guaranteed to be in back-to-front order

Layered Depth Image

- Rendering from LDI

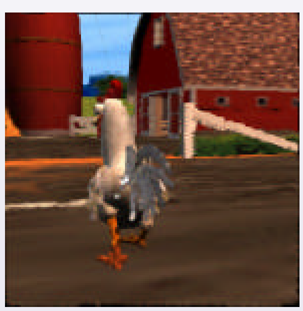

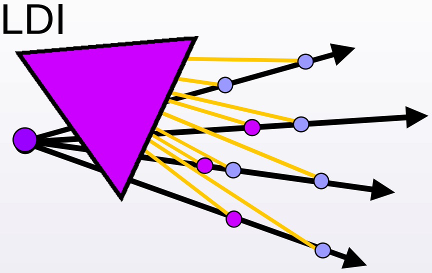

• Incremental in LDI X and Y
• Guaranteed to be in back-to-front order

Layered Depth Image

View Dependent Surfaces

- From multiple input images
  - determine geometry
    - with user's help
  - view dependent texture

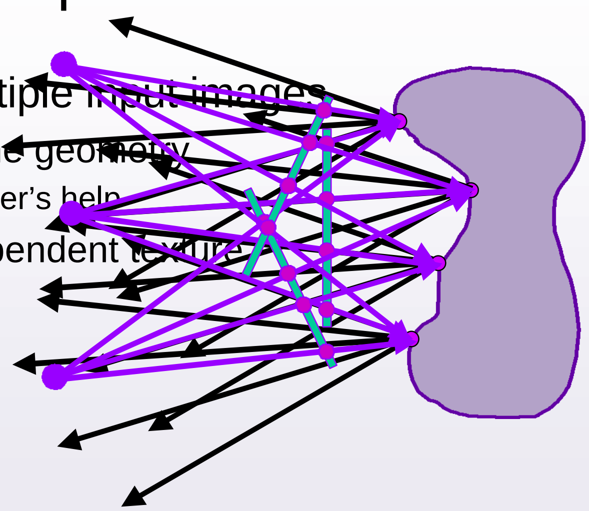

View Dependent Surfaces

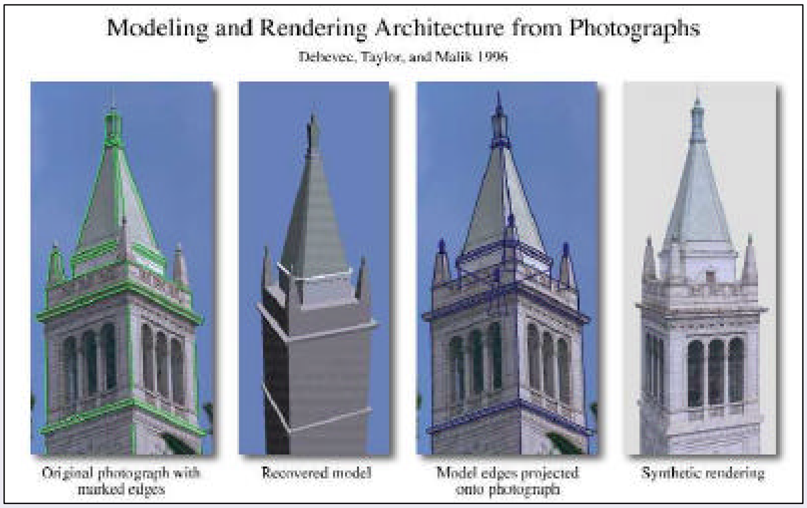

Summary

- 5D: Plenoptic Function (Ray)
- 4D: Lumigraph / Lightfield
- 3D: Lumigraph Subset
- 3D: Concentric Mosaics
- 2.5D: Layered Depth Image
- 2.5D: View Dependent Models
- 2D: Image

Thanks

- Peter-Pike Sloan (Lumigraph)
- Jonathan Shade (Lightfield, LDI)
- Marc Levoy, Stanford University
  - Michaelangelo data set