

Accurate IBR

- Is there a "photograph" that gives us all the information we need to view a scene correctly from any viewpoint?

- What dimension is the "image?"
- How can we represent the "image?"

- Answers

- Plenoptic Function
- Light Field
- Lumigraph
- Layered Depth Image




Images as a Collection of Rays



An image is a subset of the rays seen from a given point
- this "space" of rays occupies two dimensions

The Plenoptic Function

- ✓ The set of rays seen from all points ...




$$p = P(\theta, \phi, x, y, z, \lambda, t)$$

Where to Begin?


- ✓ Pinhole camera model

- Defines a mapping from image points to rays in space



Mapping from Rays to Points


- ✓ Simple Derivation



$$P = \begin{bmatrix} u_x & v_x & o_x \\ u_y & v_y & o_y \\ u_z & v_z & o_z \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\vec{X} = \vec{C} + t P \vec{x}$$

Correspondence



$$\begin{aligned} \vec{C}_2 + t_2 P_2 \vec{x}_2 &= \vec{C}_1 + t_1 P_1 \vec{x}_1 \\ t_2 P_2 \vec{x}_2 &= \vec{C}_1 - \vec{C}_2 + t_1 P_1 \vec{x}_1 \\ t_2 \vec{x}_2 &= P_2^{-1} (\vec{C}_1 - \vec{C}_2) + t_1 P_2^{-1} P_1 \vec{x}_1 \\ \frac{t_2}{t_1} \vec{x}_2 &= \frac{1}{t_1} P_2^{-1} (\vec{C}_1 - \vec{C}_2) + P_2^{-1} P_1 \vec{x}_1 \\ \vec{x}_2 &= \frac{1}{t_1} \frac{P_2^{-1} (\vec{C}_1 - \vec{C}_2)}{e_{21}} + \frac{P_2^{-1} P_1 \vec{x}_1}{h_{21}} \end{aligned}$$

Resulting Warping Function

✓ A perturbed planar warp ...

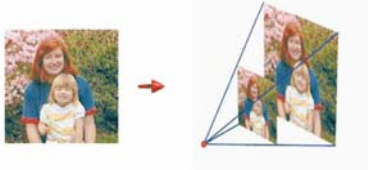
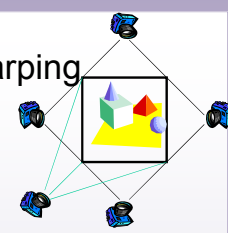
$$\bar{x}_i = \delta \bar{e}_i + H_i \bar{x}_i$$



Image Warping



- Warping allows us to replace geometric detail with textures
 - Textures created from photographs
 - Mapped to coarse planar model
- Warping problems
 - Warping incorrect for non-planes
 - Depth warping creates "holes"
- Image warping alone not enough to correctly reconstruct arbitrary scene

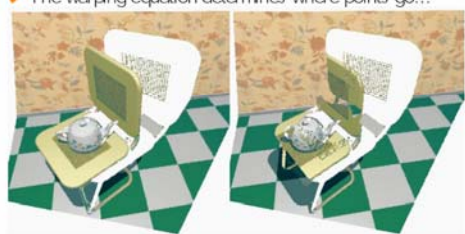
Warping in Action

✓ A 3D Warp



Visibility

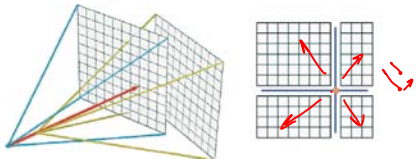
✓ The warping equation determines where points go...



... but that is not sufficient

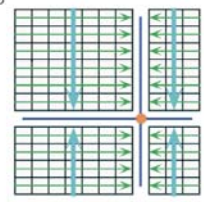
Partition Reference Image

✓ Project the desired center-of-projection onto the reference image




Enumeration

- ✓ Drawing toward the projected point guarantees an *occlusion compatible* ordering
- ✓ Ordering is consistent with a painter's algorithm
- ✓ Independent of the scene's contents
- ✓ Easily generalized to other viewing surfaces
- ✓ No auxiliary information required




Panoramic Cameras



- ✓ Warping equation can be easily adapted
- ✓ Visibility algorithm works
- ✓ Nonlinear mapping functions

Examples

✓ Cylindrical camera




Constructing Panoramas

✓ Images are related by a projective transforms

$$\vec{x}_2 = H_{21} \vec{x}_1$$


✓ Optimization problem

- maximize normalized correlation
- minimize sum of squared error



Initial Guesses and Constraints

- ✓ Sum of angles is 2π
 - constrains focal length
- ✓ Skew of camera is near 0
- ✓ Aspect ratio near 1




More General Case

- Constant radiance
 - time is fixed

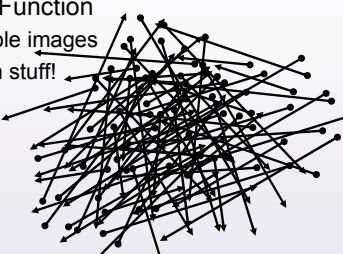
Ray

- Constant radiance
 - time is fixed
- 5D
 - 3D position
 - 2D direction

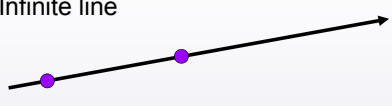


All Rays

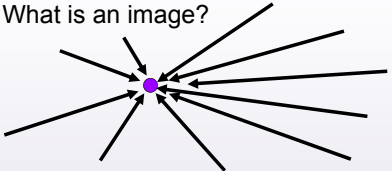
- Plenoptic Function
 - all possible images
 - too much stuff!



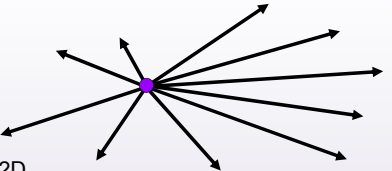
Line

- Infinite line
 
- 4D
 - 2D direction
 - 2D position

Image

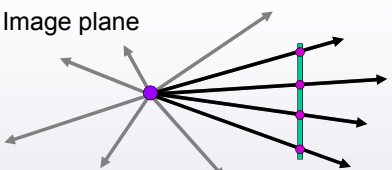
- What is an image?
 
- All rays through a point
 - Panorama?

Image

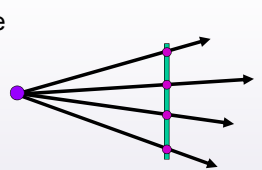



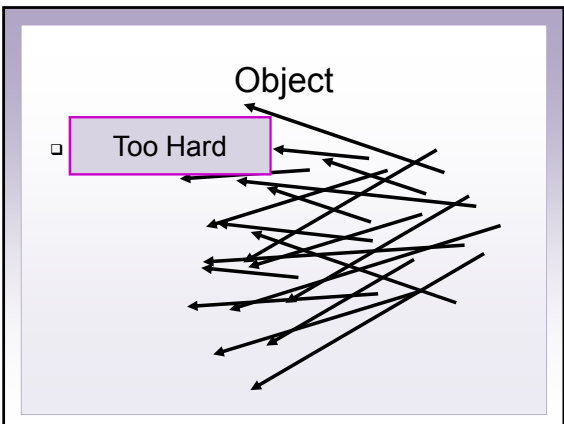




- 2D
 - position of rays has been fixed
 - direction remains


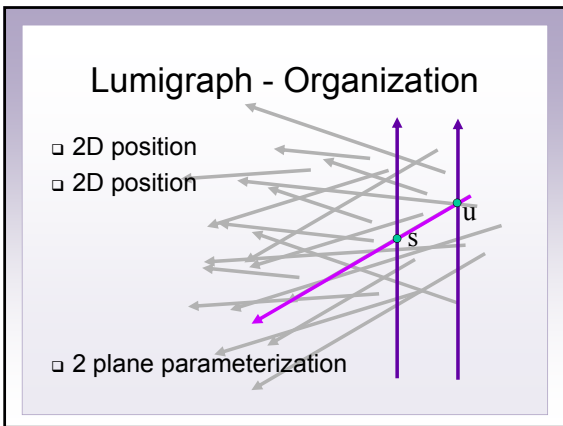

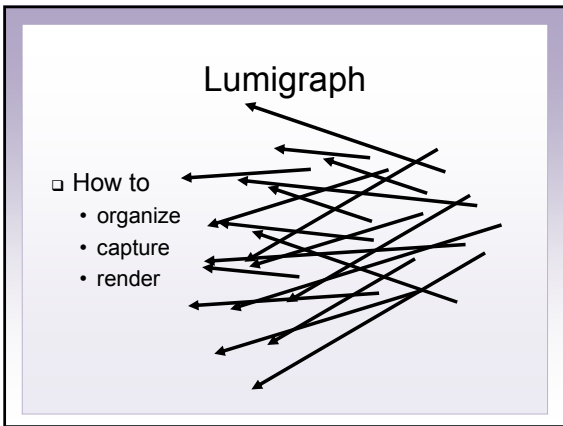


Image

- Image plane
 
- 2D
 - position

Image


- Image plane
 
- 2D
 - position



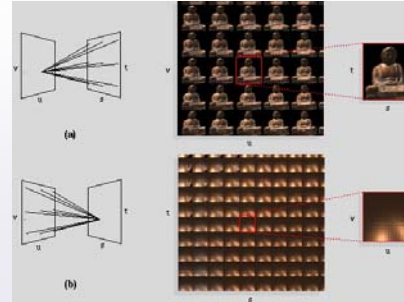


Lumigraph - Organization

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




Lightfield - Organization




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




Stanford Gantry



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
 - quadrilinear 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

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

Lumigraph - Example

Parallel Lumigraph

Lightfield - Examples


Lightfield Video

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

Modeling and Rendering Architecture from Photographs

Debevec, Taylor, and Malik 1996



Facade Quicktime

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

Plenoptic Cameras

Thanks

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