Image-based rendering

Michael F. Cohen
Microsoft Research

Computer Graphics

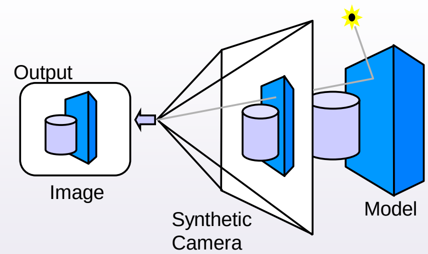

Computer Vision

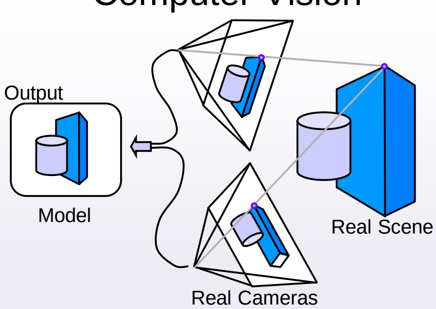

Combined

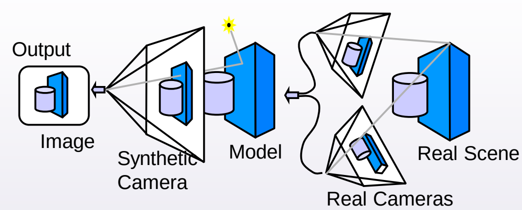

But, vision technology falls short

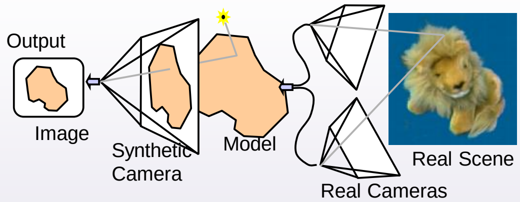

… and so does graphics.

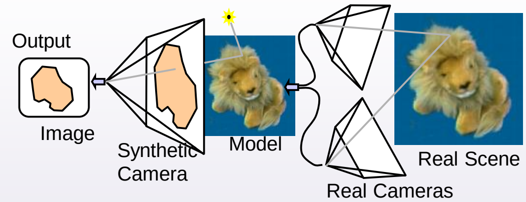

Image Based Rendering

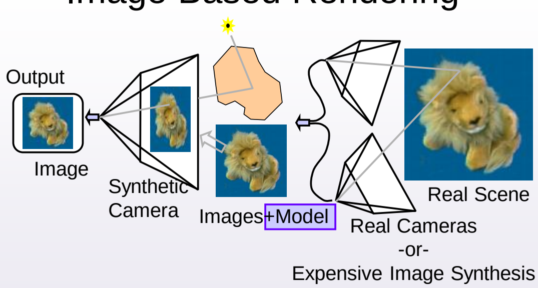

Ray

- Constant radiance
  - time is fixed

- 5D
  - 3D position
  - 2D direction

All Rays

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

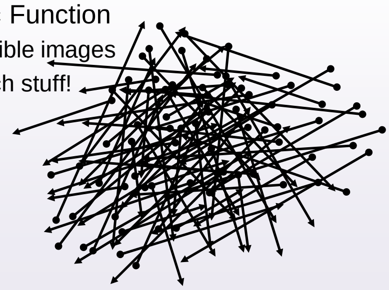

Line

- Infinite line

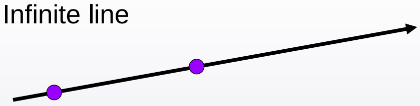

- 4D
  - 2D direction
  - 2D position

Ray

- Discretize

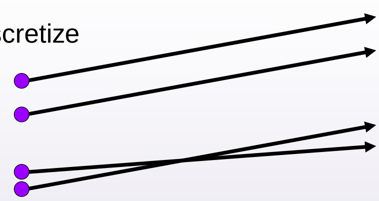

- Distance between 2 rays
  - Which is closer together?

Image

- What is an image?

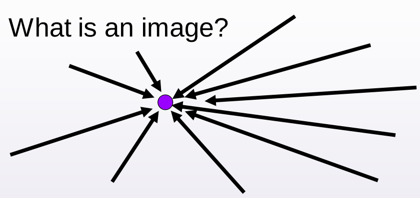

- All rays through a point
  - Panorama?

Image

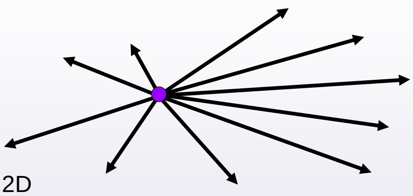

- 2D
  - position of rays has been fixed
  - direction remains

Image

- Image plane

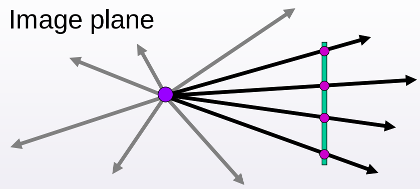

- 2D
  - position

Image

- Image plane

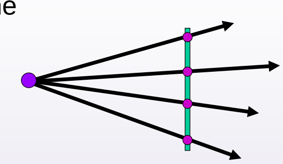

- 2D
  - position

Object

- Light leaving towards “eye”

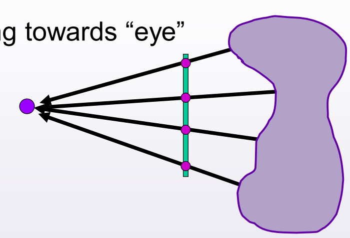

- 2D
  - just dual of image

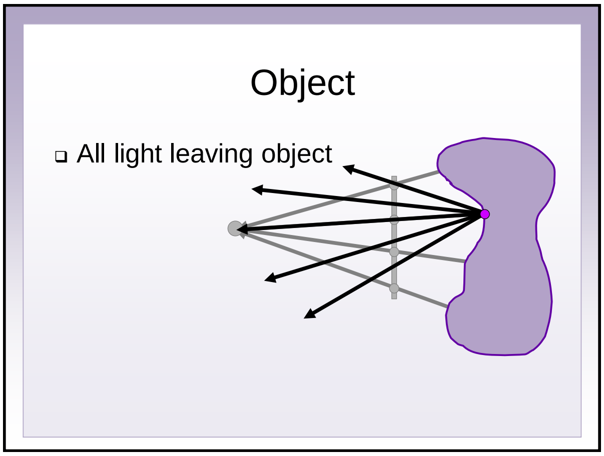
Object
All light leaving object

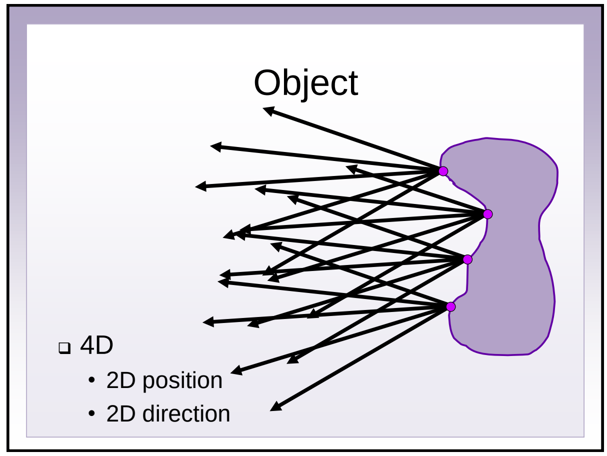
Object
4D
2D position
2D direction

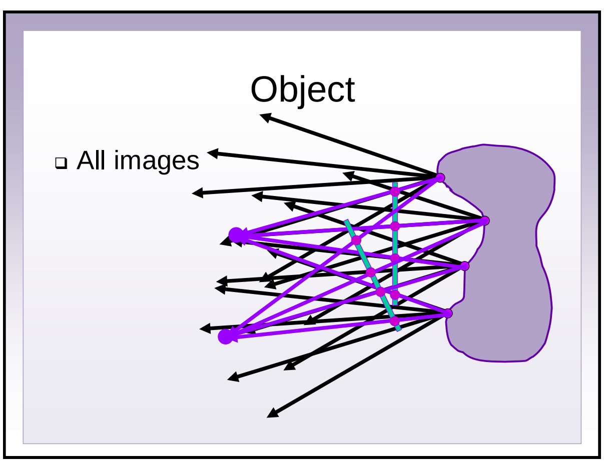
Object
All images

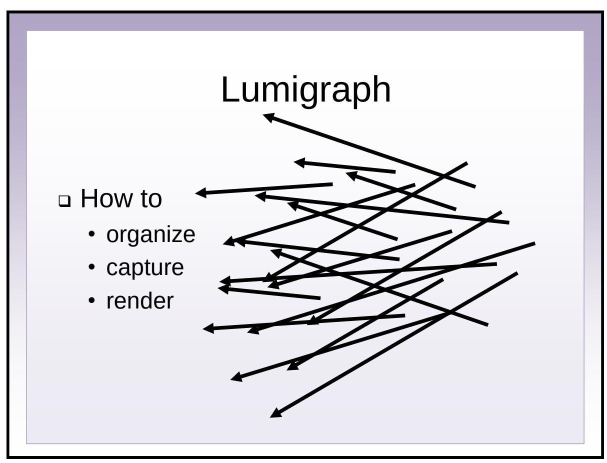
Lumigraph
How to
organize
capture
render

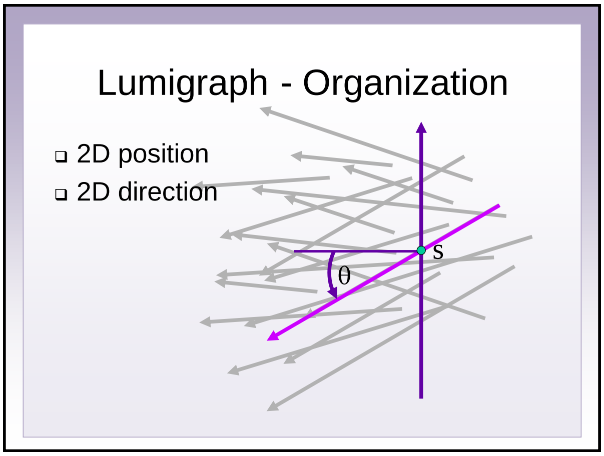
Lumigraph - Organization
2D position
2D direction
s
θ

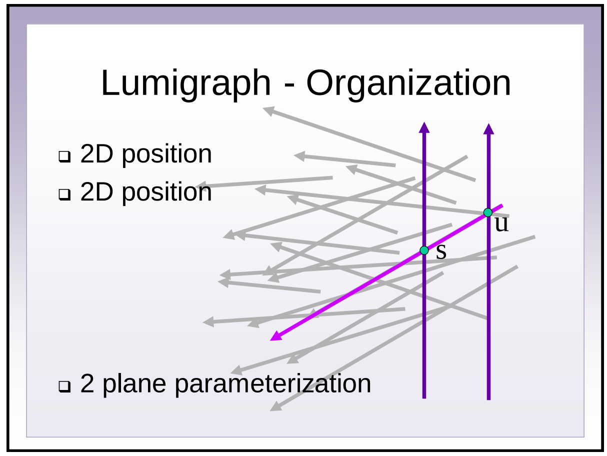
Lumigraph - Organization
2D position
2D position
u
s
2 plane parameterization

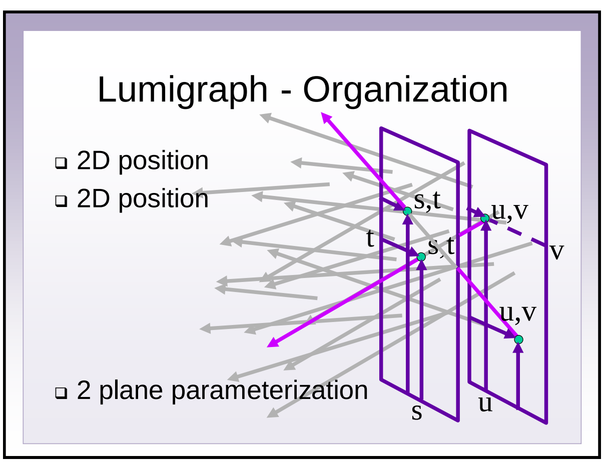
Lumigraph - Organization
2D position
2D position
s,t
u,v
t
s,t
v
u,v
2 plane parameterization
s
u

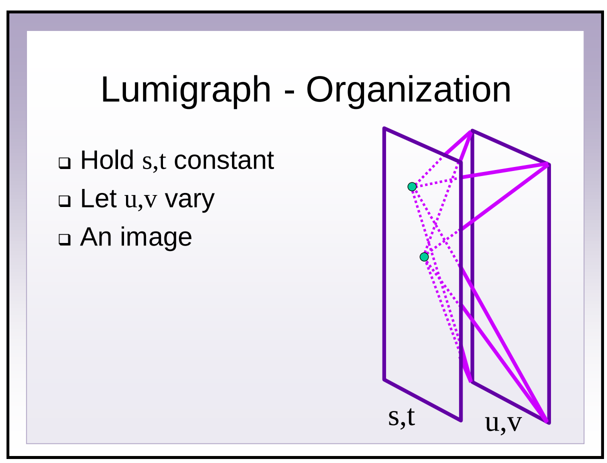
Lumigraph - Organization
Hold s,t constant
Let u,v vary
An image
s,t
u,v

Lumigraph - Organization

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

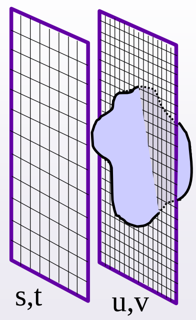

Lumigraph - Capture

- Idea 1
  - Move camera carefully over s,t plane
  - Gantry
    - see Lightfield paper

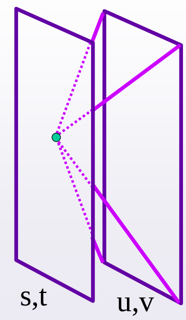

Lumigraph - Capture

- Idea 2
  - Move camera anywhere
  - Rebinning
    - see Lumigraph paper

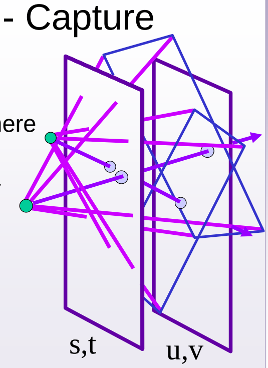

Lumigraph - Rendering

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

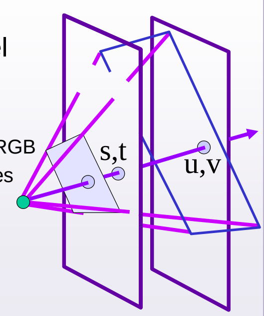

Lumigraph - Rendering

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

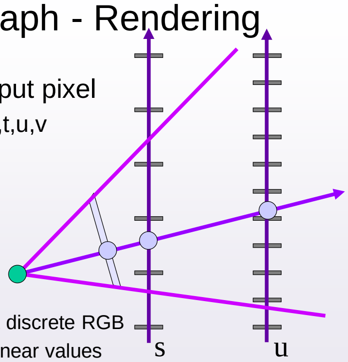

Lumigraph - Rendering

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

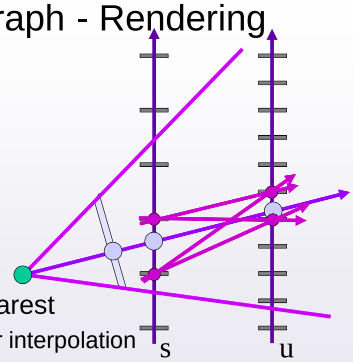

High-Quality Video View Interpolation Using a Layered Representation

Larry Zitnick
Sing Bing Kang
Matt Uyttendaele
Simon Winder
Rick Szeliski

Interactive Visual Media Group
Microsoft Research

Current practice

free viewpoint video

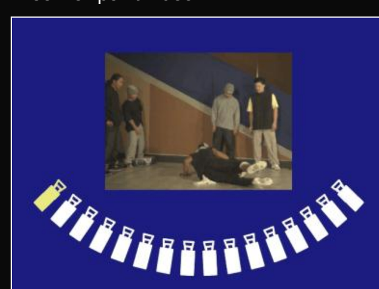

Many cameras

vs.

Motion Jitter

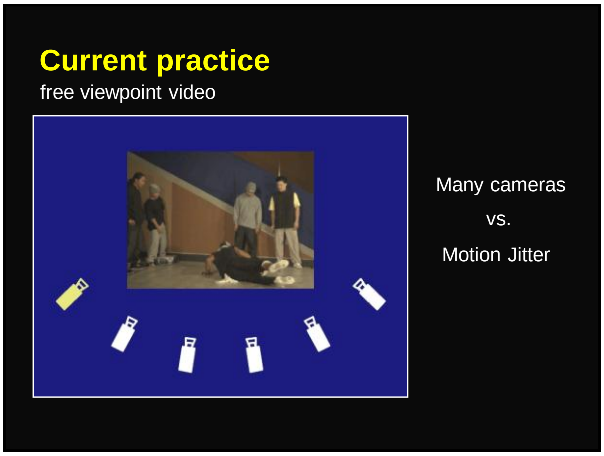
Current practice
free viewpoint video
Many cameras
vs.
Motion Jitter

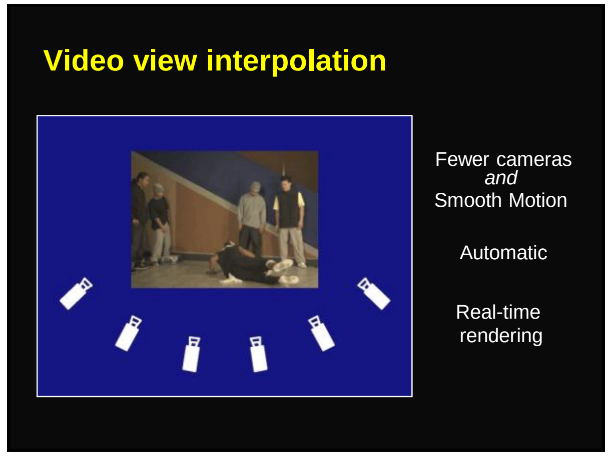
Video view interpolation
Fewer cameras
and
Smooth Motion
Automatic
Real-time
rendering

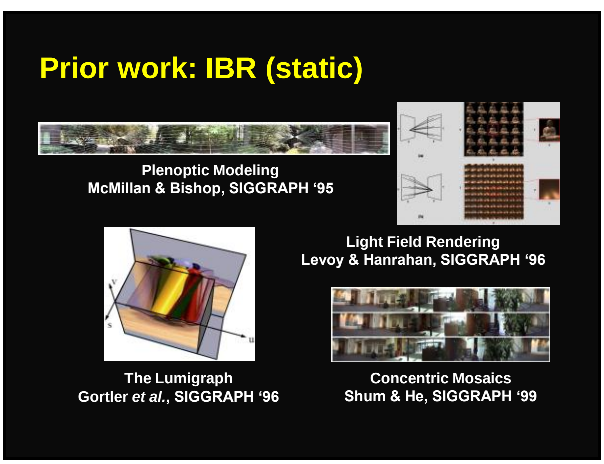
Prior work: IBR (static)
Plenoptic Modeling
McMillan & Bishop, SIGGRAPH '95
Light Field Rendering
Levoy & Hanrahan, SIGGRAPH '96
The Lumigraph
Gortler et al., SIGGRAPH '96
Concentric Mosaics
Shum & He, SIGGRAPH '99

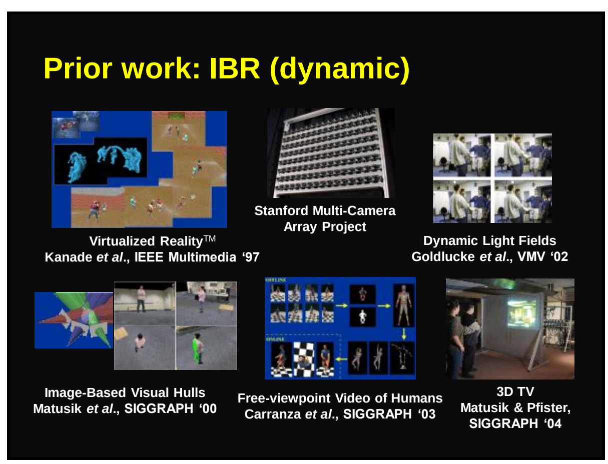
Prior work: IBR (dynamic)
Virtualized Reality™
Kanade et al., IEEE Multimedia '97
Stanford Multi-Camera
Array Project
Dynamic Light Fields
Goldlucke et al., VMV '02
Image-Based Visual Hulls
Matusik et al., SIGGRAPH '00
Free-viewpoint Video of Humans
Carranza et al., SIGGRAPH '03
3D TV
Matusik & Pfister,
SIGGRAPH '04

System overview

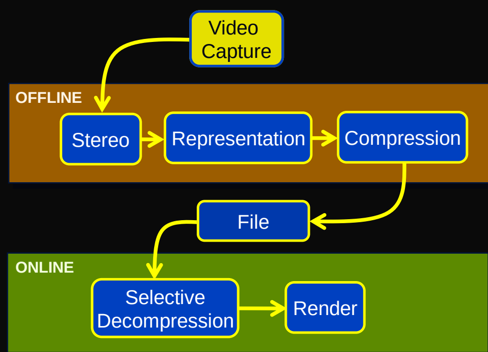

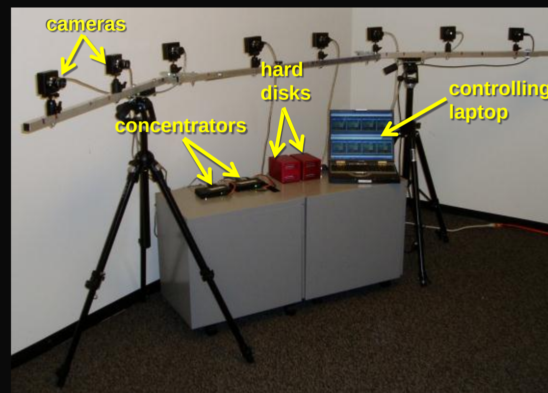

Calibration

Zhengyou Zhang, 2000

Input videos

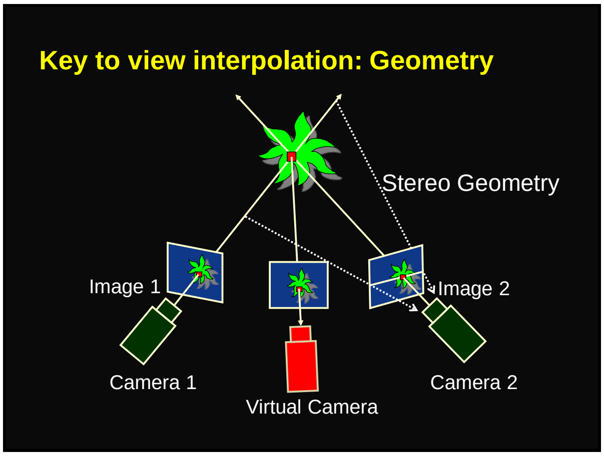
Key to view interpolation: Geometry
Stereo Geometry
Image 1
Image 2
Camera 1
Virtual Camera
Camera 2

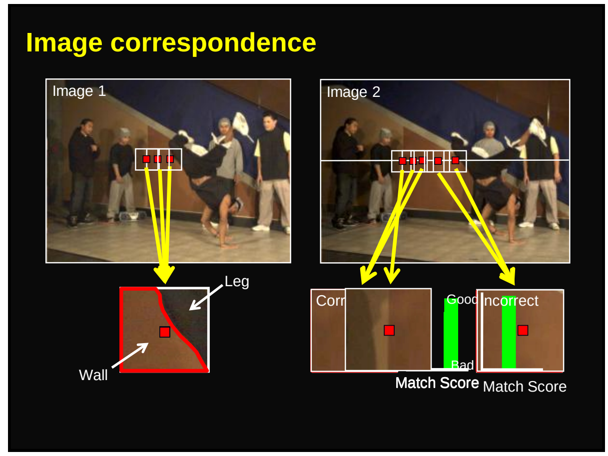
Image correspondence
Image 1
Image 2
Leg
Wall
Incorrect
Match Score
Match Score

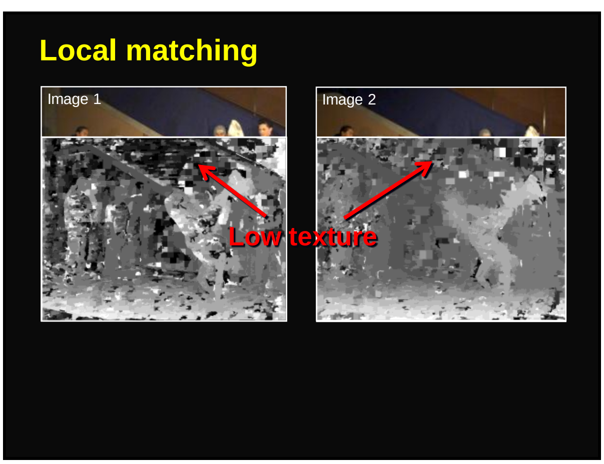
Local matching
Image 1
Image 2
Low texture

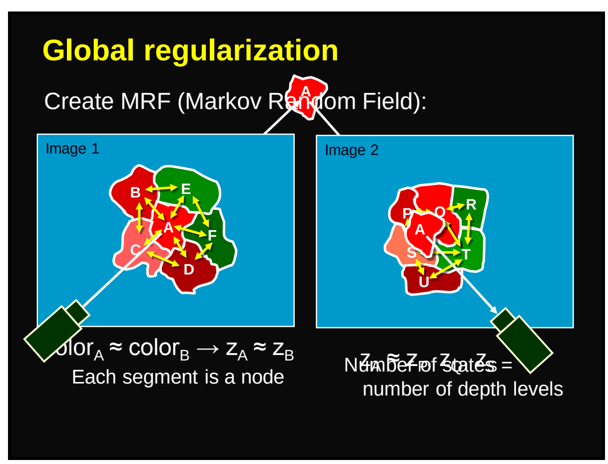
Global regularization
Create MRF (Markov Random Field):
Image 1
Image 2
Each segment is a node
number of depth levels

Iteratively solve MRF

Depth through time

Matting

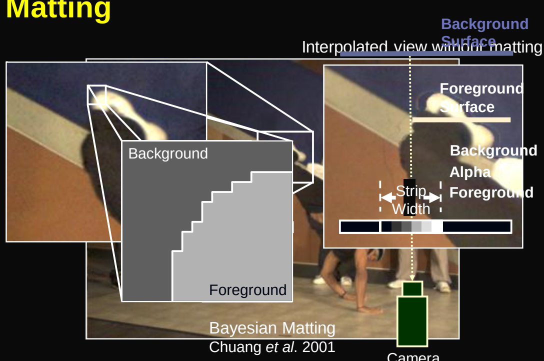

Rendering with matting

No Matting

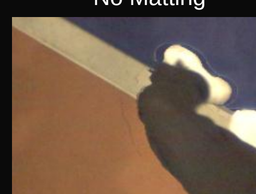

Matting

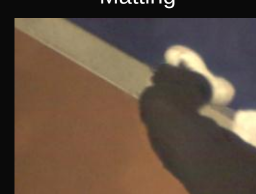

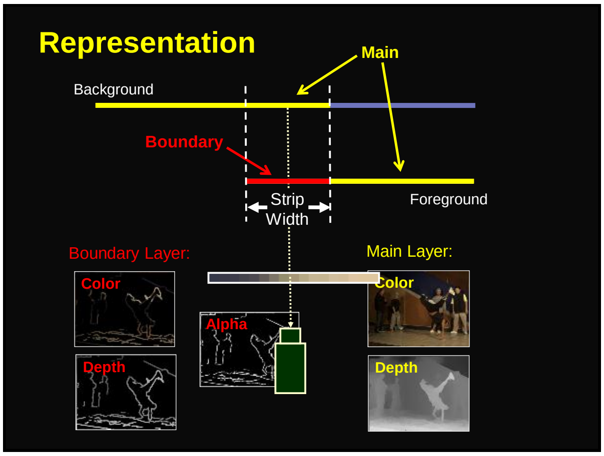
Representation
Main
Background
Boundary
Strip
Width
Foreground
Boundary Layer:
Main Layer:
Color
Alpha
Color
Depth
Depth


Massive Arabesque