Computer Vision, CS766

Staff

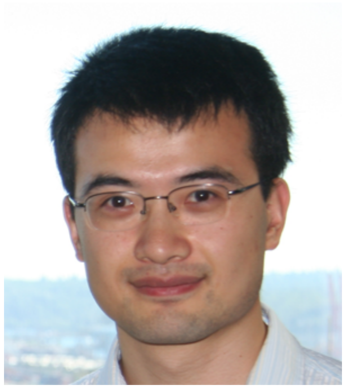

Instructor: Li Zhang
lizhang@cs.wisc.edu

TA:
Brandon Smith, bmsmith@cs.wisc.edu
Shengqi Zhu, sqzhu@cs.wisc.edu
Travis Portz, tportz@cs.wisc.edu
Yiqing Yang, breakds@cs.wisc.edu

Today

Introduction

Administrative Stuff

Overview of the Course

About Me

- Li Zhang
  - Last name pronounced as Jung
  - www.cs.wisc.edu/~lizhang
- Research
  - Computer Vision
  - Computer Graphics
- Teaching
  - CS766 Computer Vision
  - CS559 Computer Graphics

Example Research Projects

- 3D shape reconstruction

Scene

3D Model

Depth Map

Example Research Projects

- 3D shape reconstruction

Four examples of recovered 3D shapes
of a moving face from six video streams

Example Research Projects

- 3D shape reconstruction
- Application

Entertainment:
Games & Movies

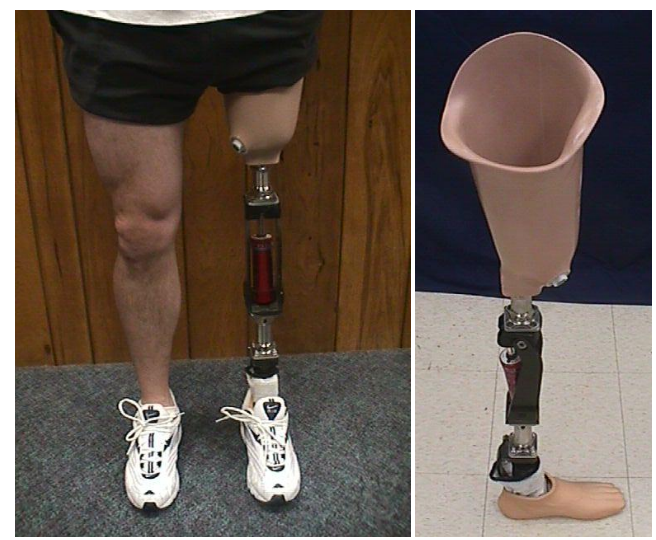

Medical Practice:
Prosthetics

Example Research Projects

- 3D shape reconstruction
- Image recognition/segmentation

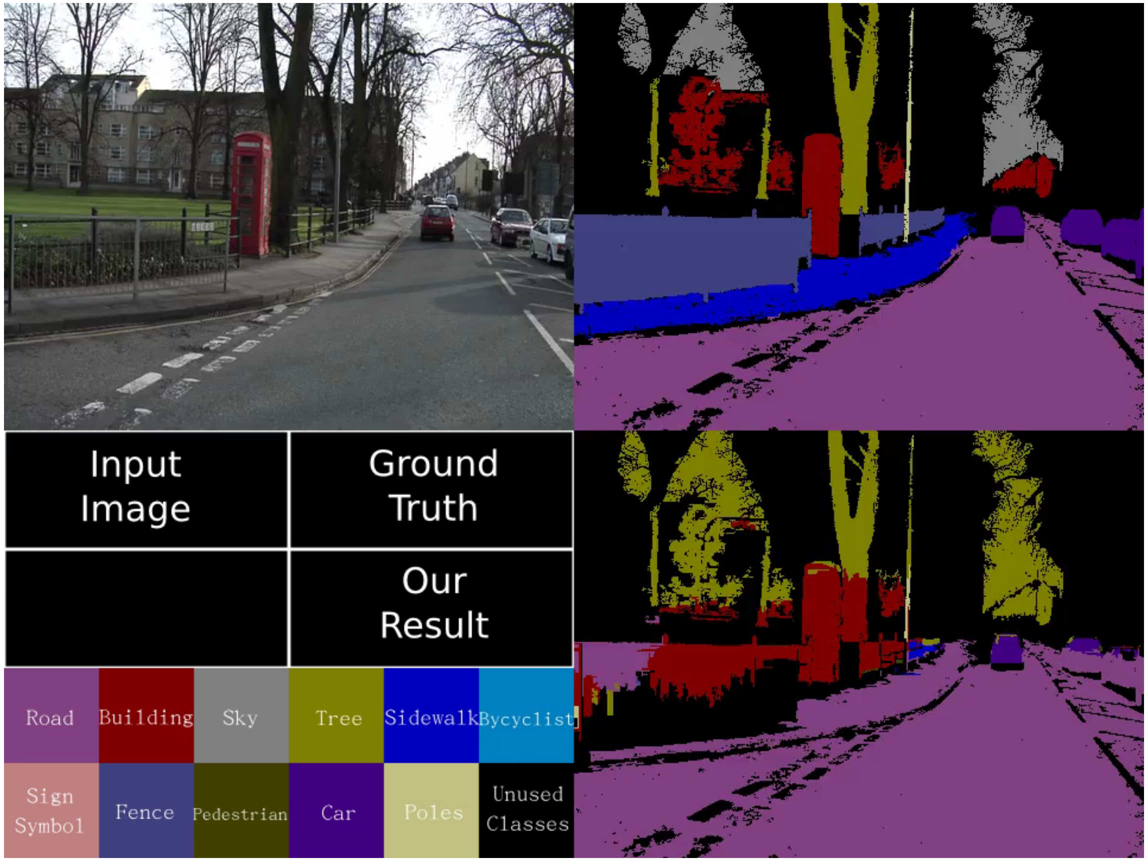

Example Research Projects

- 3D shape reconstruction
- Image recognition/segmentation
- Image video enhancement

A Conventional Photo

Automatic Gain Adjustment

Example Research Projects

- 3D shape reconstruction
- Image recognition/segmentation
- Image video enhancement

A Short Image Sequence

Our result 2009

Please tell me about you

Who you are?
Why you are taking this class?
What do you want to learn?

Prerequisites

- Prerequisites—*these are essential*!
  - Data structures & Algorithms
  - A good working knowledge of C++/Java programming
    - (or willingness/time to pick it up quickly!)
    - If you know Matlab, projects will be easier
  - Linear algebra
  - Calculus

- Course does ***not*** assume prior imaging experience
  - no image processing, graphics, etc.

Administrative Stuff

- **1 written assignment**
  - 5% (this week)
- **3 programming projects**
  - 15%, 2-3 weeks each
- **Paper presentation**
  - 15%, over a month
  - Students grade the presenters
- **1 final project**
  - 35%, 5 weeks, open ended of your choice, but needs
  - project proposal after 1 week
  - progress report after 3 weeks
  - Final presentation after 5 weeks

Administrative Stuff

- **Computer account**:
  - Everyone registered in this class will get a Computer Systems Lab account to do project assignments.
- **Email list**:
  - compsci766-1-f12@lists.wisc.edu

Questions?

Every picture tells a story

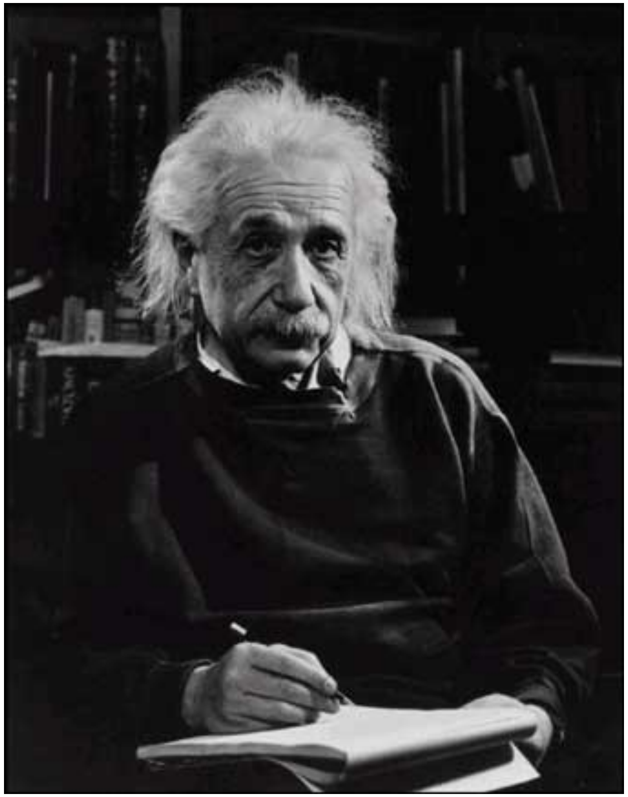

Goal of computer vision is to write computer programs that can interpret images

A little story

In 1966, Marvin Minsky asked his undergraduate student Gerald Jay Sussman to "spend the summer linking a camera to a computer and getting the computer to describe what it saw" (Boden 2006, p. 781). We now know that the problem is slightly more difficult than that.

(Szeliski 2010, Computer Vision)

Can computer match human perception?

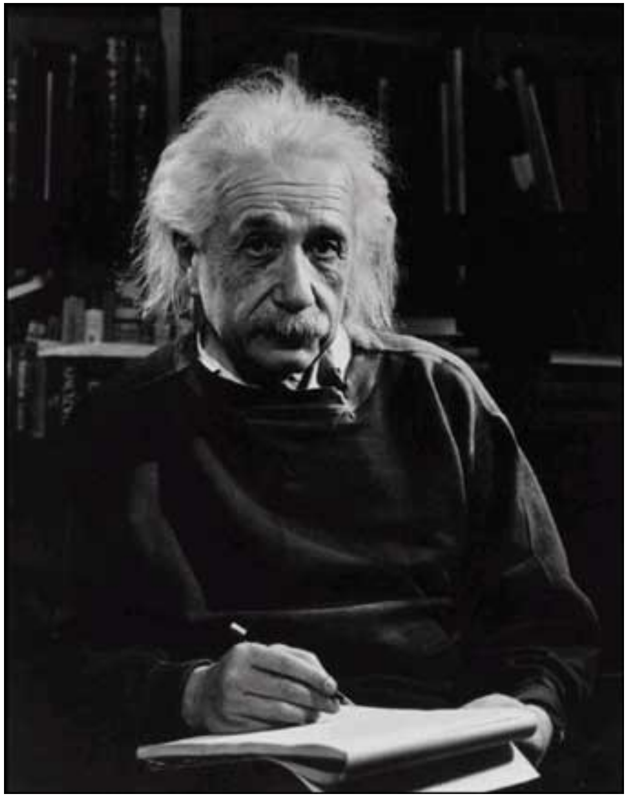

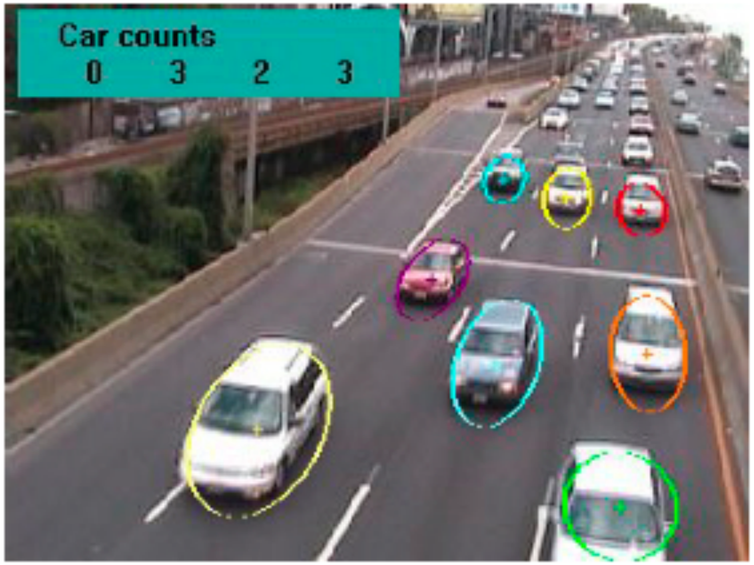

- Yes and no (but mostly no!)
  - computers can be better at “easy” things

Can computer match human perception?

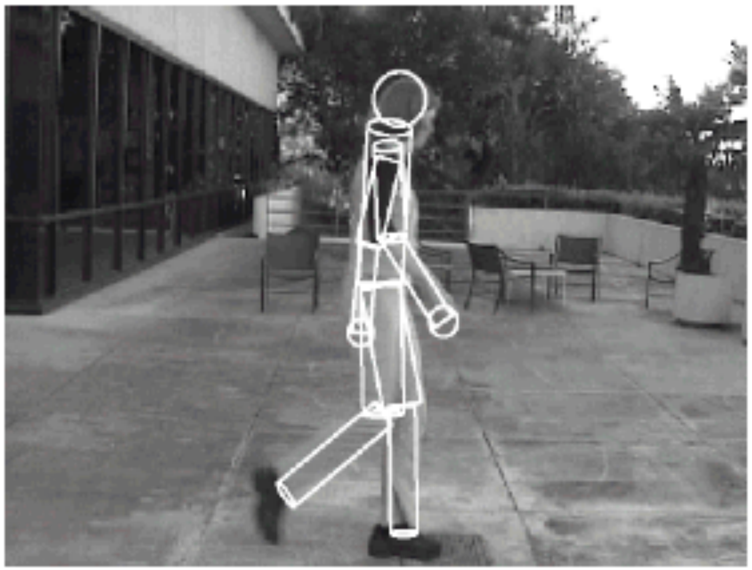

- Yes and no (but mostly no!)
  - computers can be better at “easy” things
  - humans are much better at “hard” things

Computer Vision vs Human Vision

- Can do amazing things like:
  - Recognize people and objects
  - Navigate through obstacles
  - Understand mood in the scene
  - Imagine stories

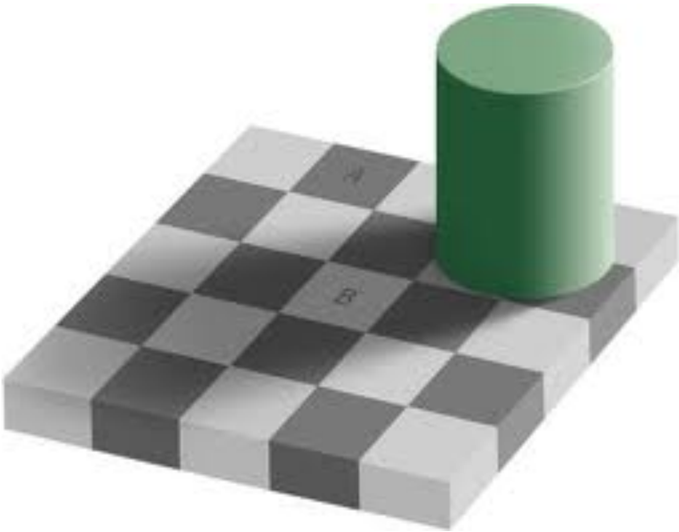

- But still is not perfect:
  - Suffers from Illusions
  - Ignores many details
  - Doesn’t care about accuracy of world

Srinivasa Narasimhan’s slide

Computer Vision vs Human Vision

- Can do amazing things like:
  - Recognize people and objects
  - Navigate through obstacles
  - Understand mood in the scene
  - Imagine stories

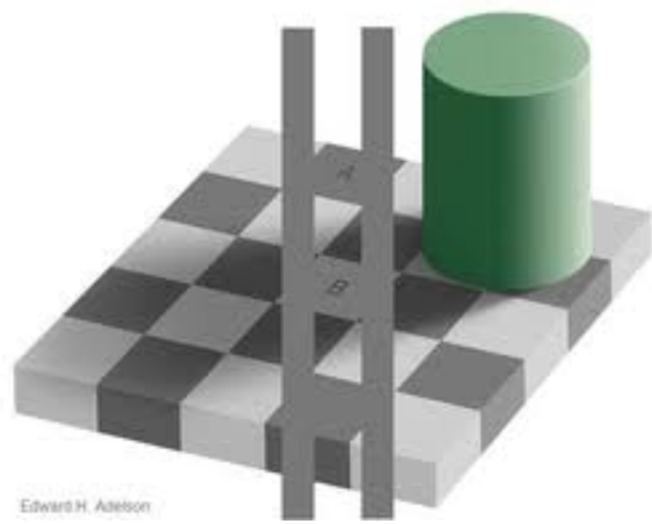

- But still is not perfect:
  - Suffers from Illusions
  - Ignores many details
  - Doesn’t care about accuracy of world

Computer vision vs Human Vision

What we see

0	3	2	5	4	7	6	9	8
3	0	1	2	3	4	5	6	7
2	1	0	3	2	5	4	7	6
5	2	3	0	1	2	3	4	5
4	3	2	1	0	3	2	5	4
7	4	5	2	3	0	1	2	3
6	5	4	3	2	1	0	3	2
9	6	7	4	5	2	3	0	1
8	7	6	5	4	3	2	1	0

What a computer sees

Srinivasa Narasimhan's slide

Components of a computer vision system

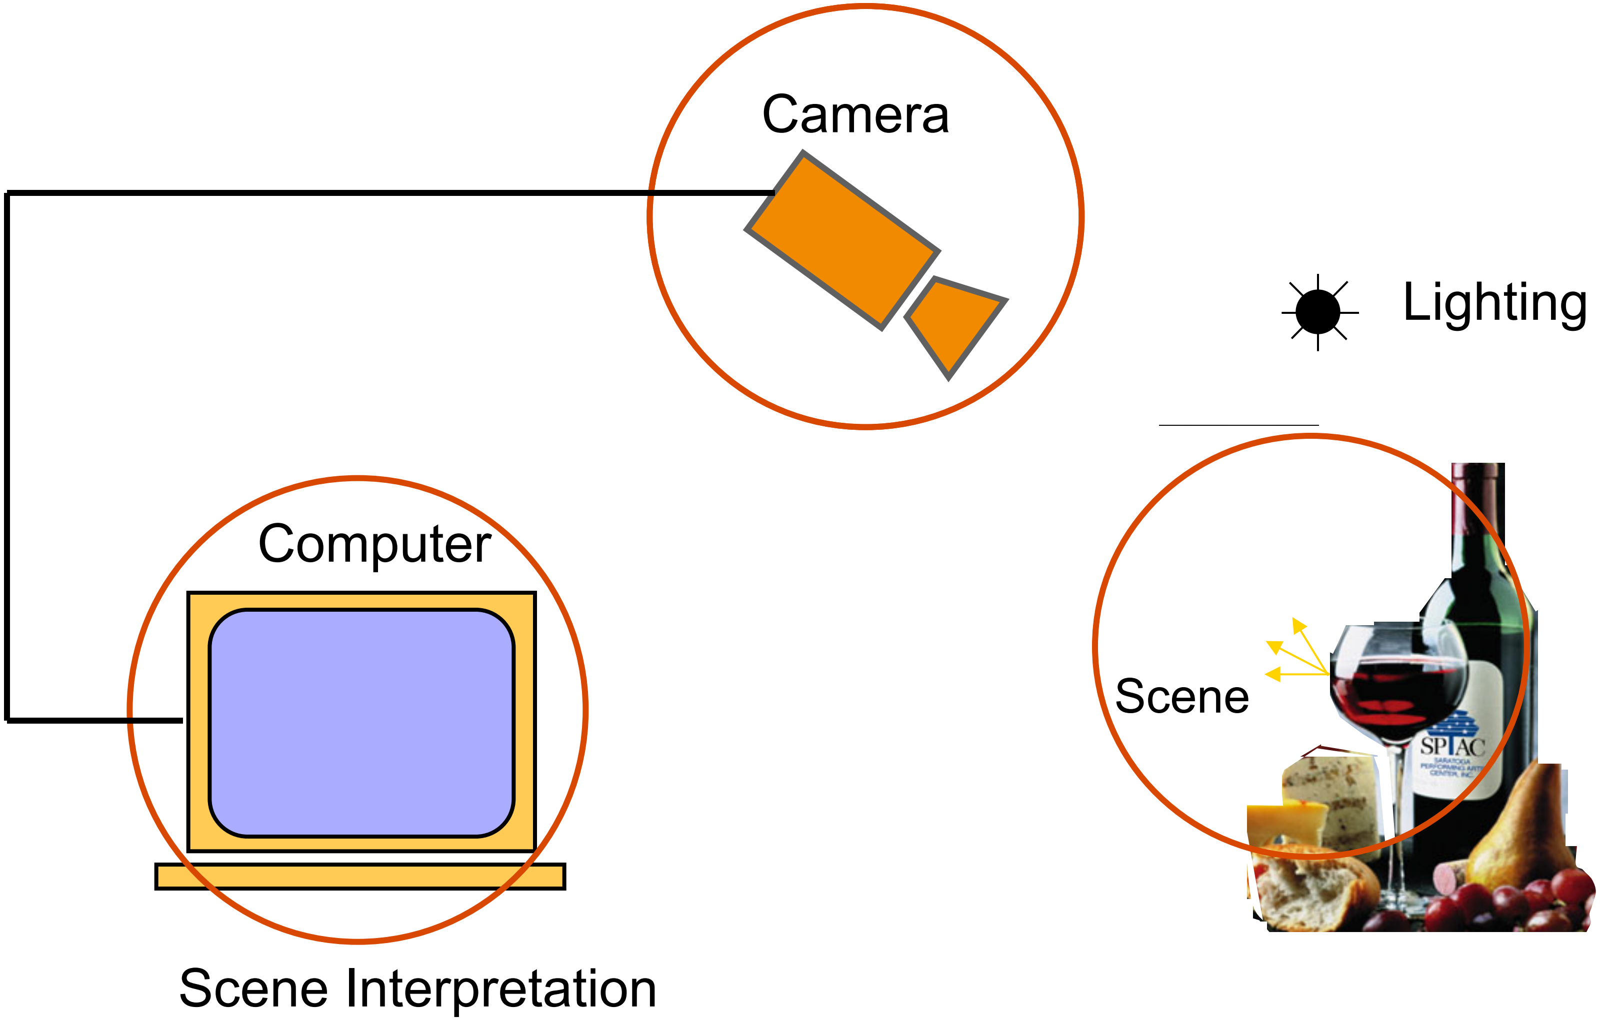

Srinivasa Narasimhan's slide

Topics Covered

Cameras and their optics

Today’s Digital Cameras

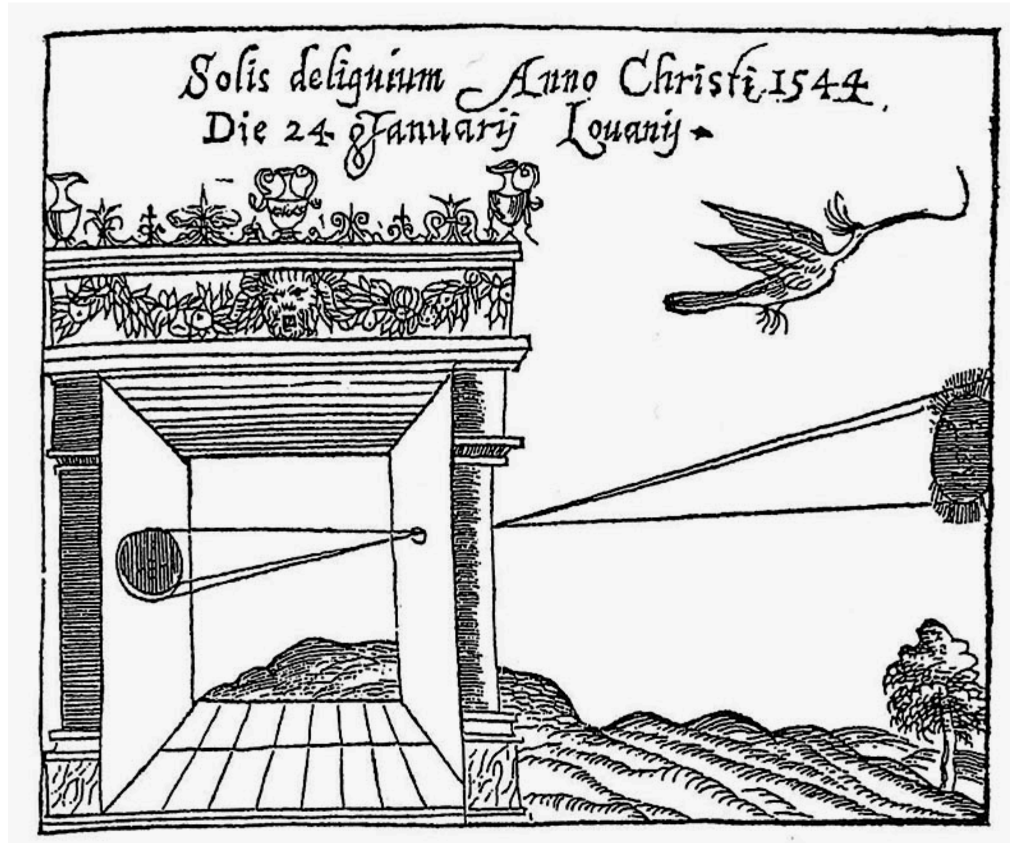

The *Camera Obscura*

Srinivasa Narasimhan’s slide

Biological vision

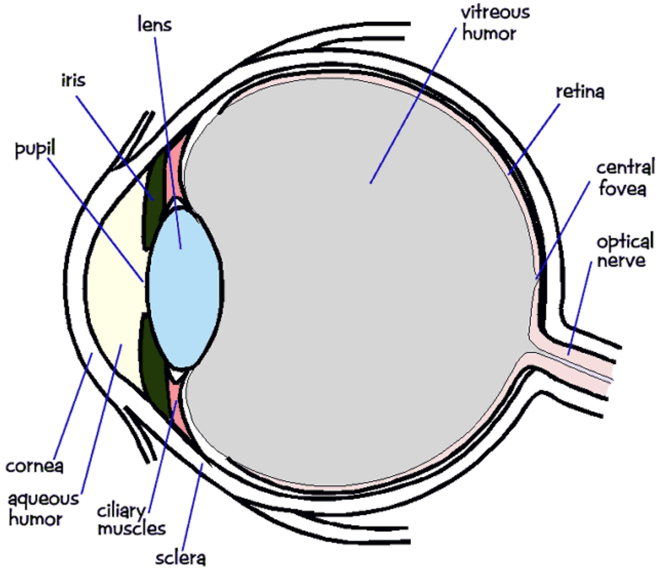

Human Eye

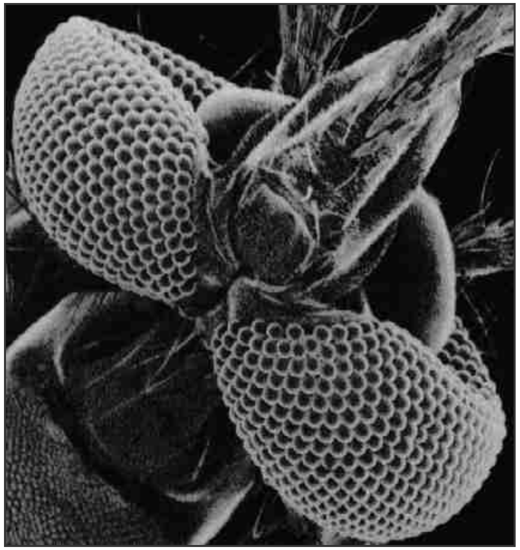

Mosquito Eye

Srinivasa Narasimhan's slide

A tiny camera

PHOTO: FRAUNHOFER INSTITUTE FOR BIOMEDICAL ENGINEERING

A tiny camera

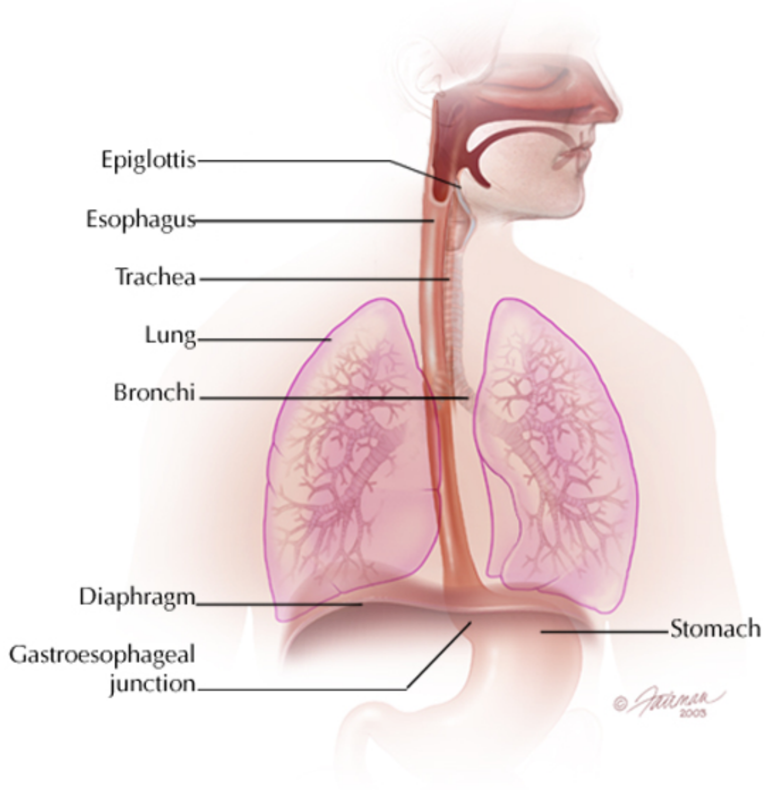

PHOTO: FRAUNHOFER INSTITUTE FOR BIOMEDICAL ENGINEERING

Project 1: High Dynamic Range Imaging

- Cameras have limited dynamic range

Short Exposure

Long Exposure

Desired Image

Shree Nayar's slide

Project 1: High Dynamic Range Imaging

iPhone 4

Project 1: High Dynamic Range Imaging

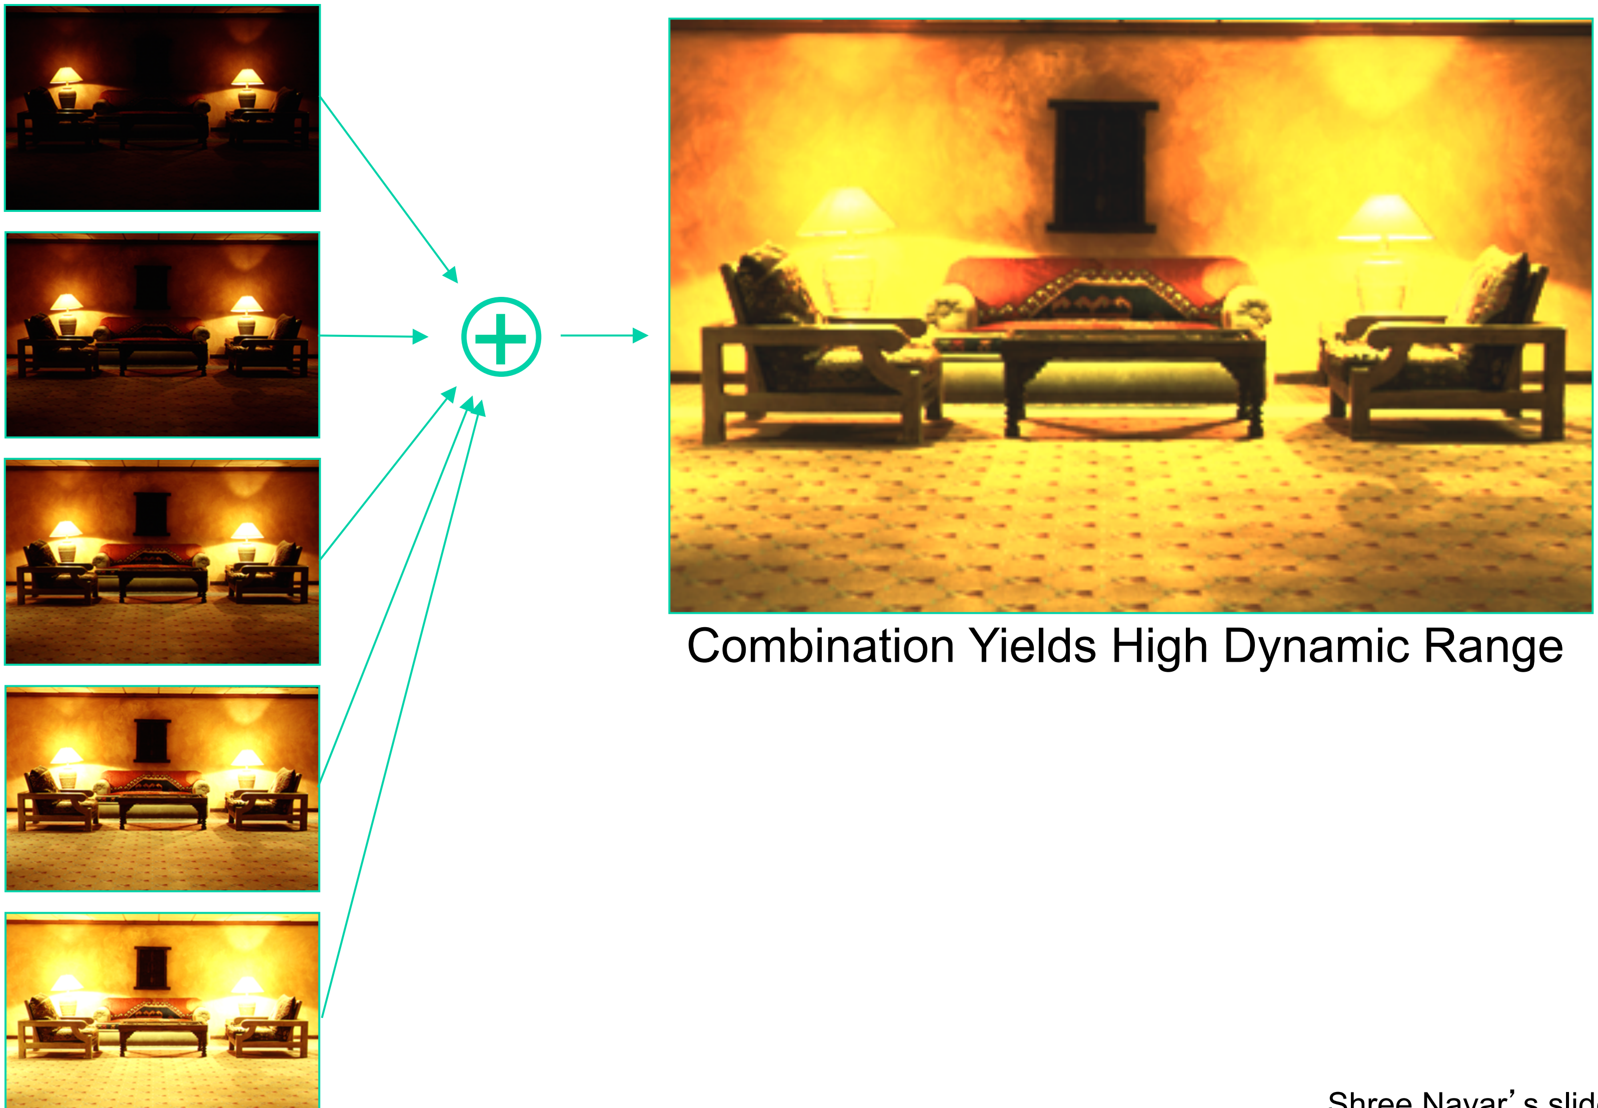

Low Dynamic Range Exposures

Combination Yields High Dynamic Range

Shree Nayar’s slide

Image Processing

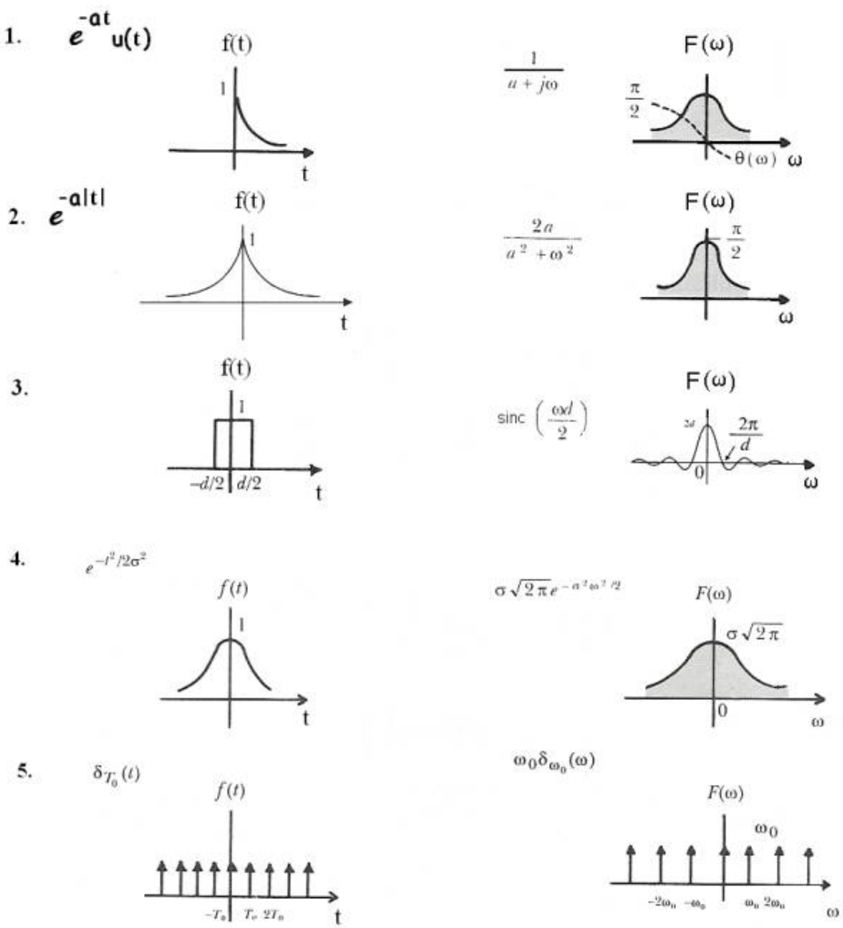

Fourier Transform
Sampling, Convolution

Image enhancement
Feature detection

Srinivasa Narasimhan's slide

Image Processing

Camera Projection

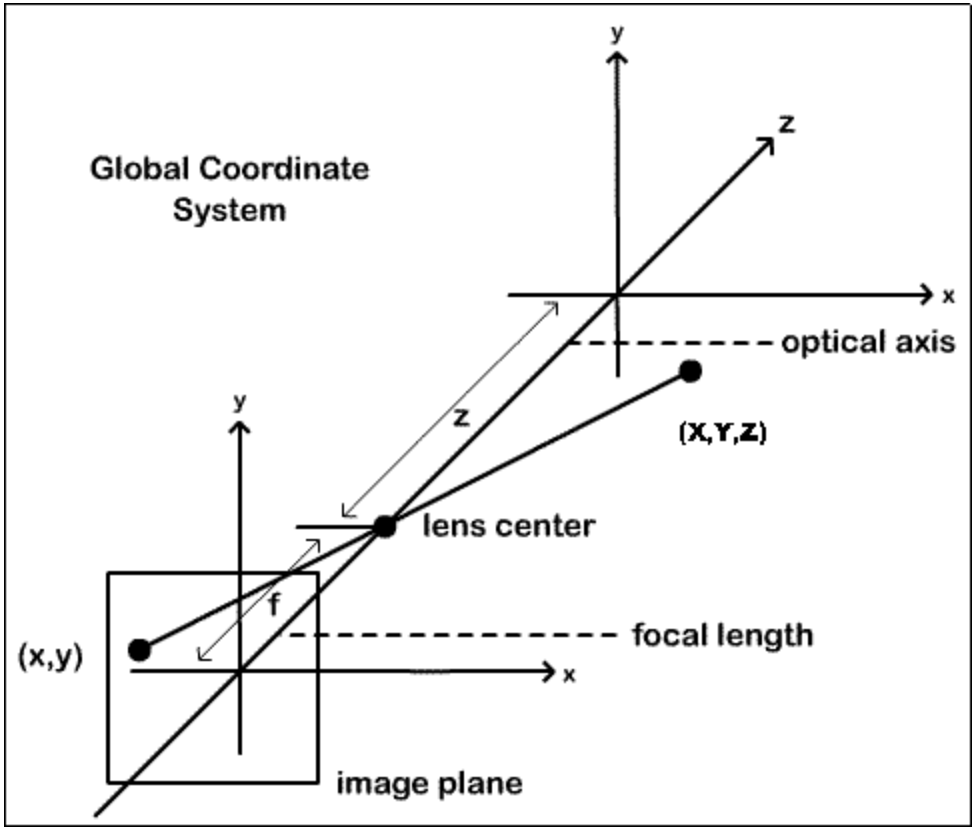

Image Transformation

Steve Seitz and Chuck Dyer, View Morphing, SIGGRAPH 1996

Project 2: Panoramic Imaging

Input images:

Output Image:

Steve Seitz's slide

Project 2: Panoramic Imaging

360-degree panorama, Curiosity Mars rover, 2012

Projective Geometry

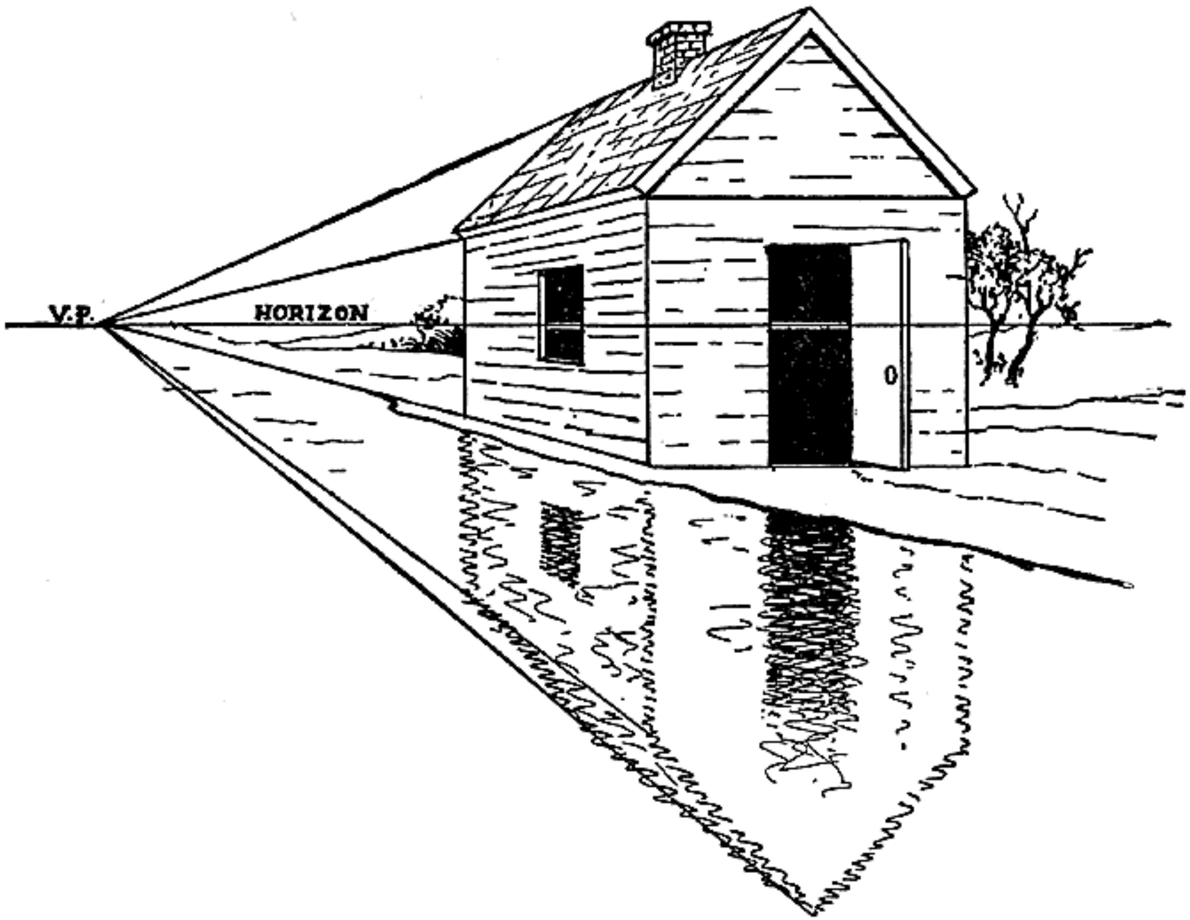

Single View Metrology

- http://research.microsoft.com/vision/cambridge/3d/default.htm

Single View Metrology

- http://research.microsoft.com/vision/cambridge/3d/default.htm

Shading and Photometric Stereo

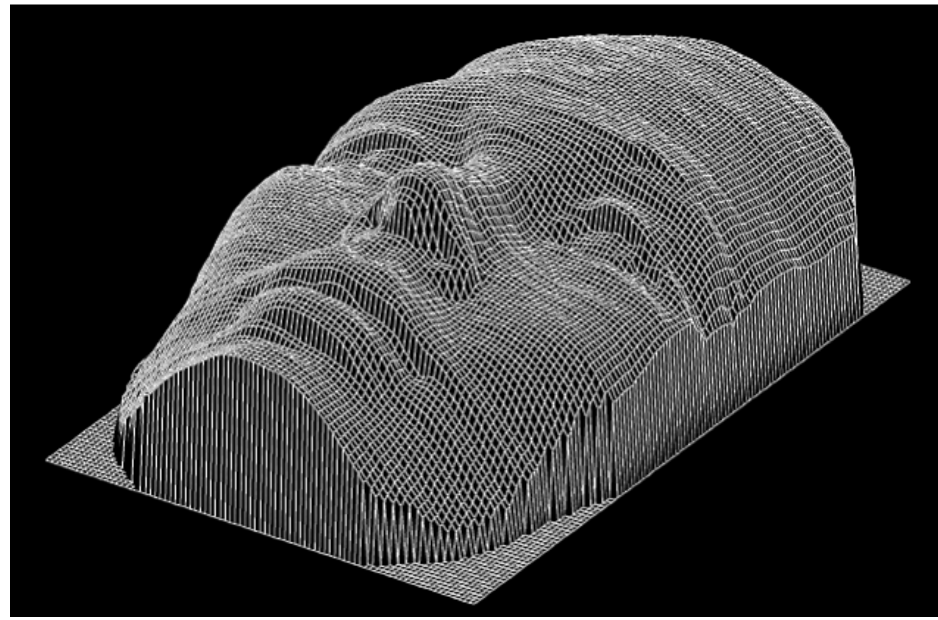

http://www.eecs.harvard.edu/~zickler/helmholtz.html

Project 3: photometric stereo

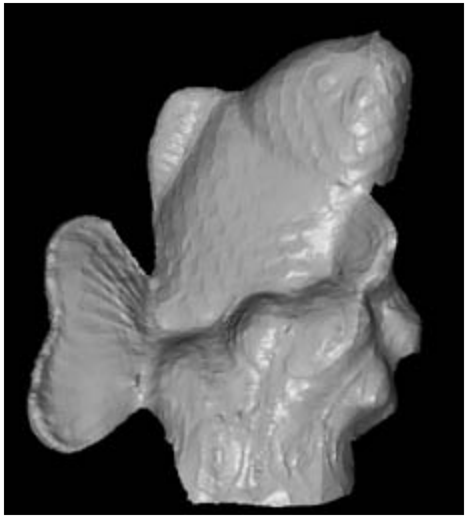

Texture Modeling

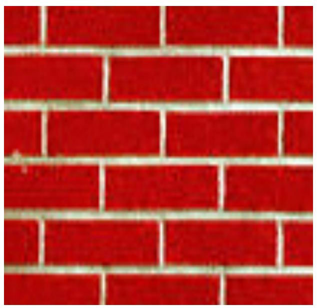

repeated

radishes

rocks

yogurt

“Semi-stochastic” structures

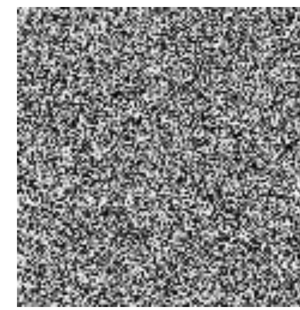

stochastic

Alexei Efros’ slide

Texture Synthesis

Output

Input

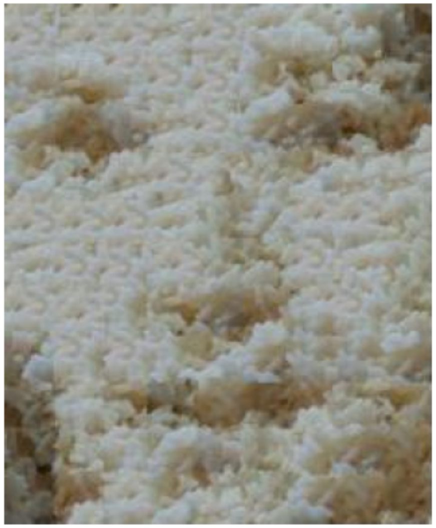

Image Quilting, Efros and Freeman., SIGGRAPH 2002.

Texture Synthesis

Input images:

Output Image:

Graphcut Textures, Kwatra et al., SIGGRAPH 2003.

Multi-view Geometry

http://phototour.cs.washington.edu/

- Binocular Stereo (2 classes)
- Multiview Stereo (1 class)
- Structure from Motion (2 classes)

Applications

- http://photosynth.net/default.aspx

Applications

- http://photosynth.net/default.aspx
- Visual odometry, curiosity Mars rover
  - http://www.csmonitor.com/Science/2012/0904/Curiosity-Mars-rover-shoots-spectacular-full-circle-panorama

Related Applications

- Apple 3D map, Google 3D,

Face Detection and Recognition

Face Detection and Recognition

Smart cameras: auto focus, red eye removal, auto color correction

Face Detection and Tracking

Face Detection and Tracking

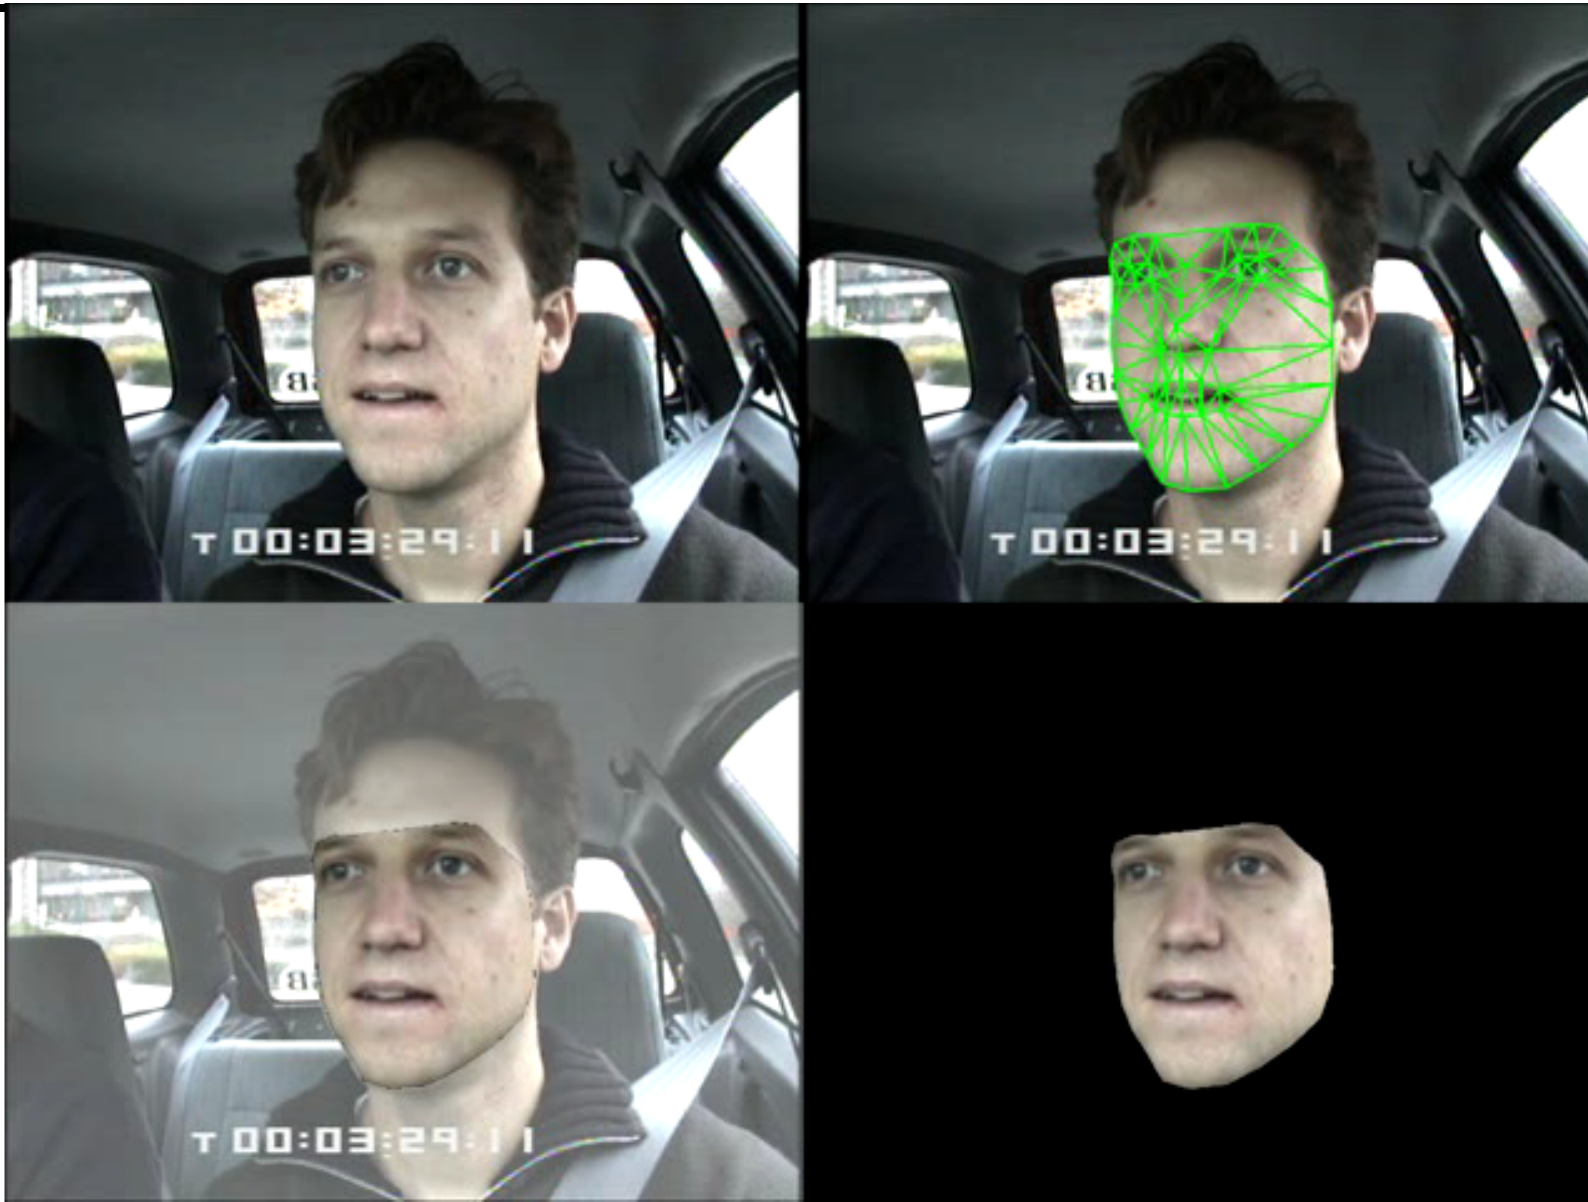

Face Detection and Tracking

Lexus LS600 Driver Monitor System

Motion Estimation

Hidden Dragon Crouching Tiger

Motion Estimation

Application

Andy Serkis, Gollum, Lord of the Rings

Motion Estimation

KINECT
for XBOX 360

Segmentation

http://www.eecs.berkeley.edu/Research/Projects/CS/vision/bsds/

Segmentation

Application

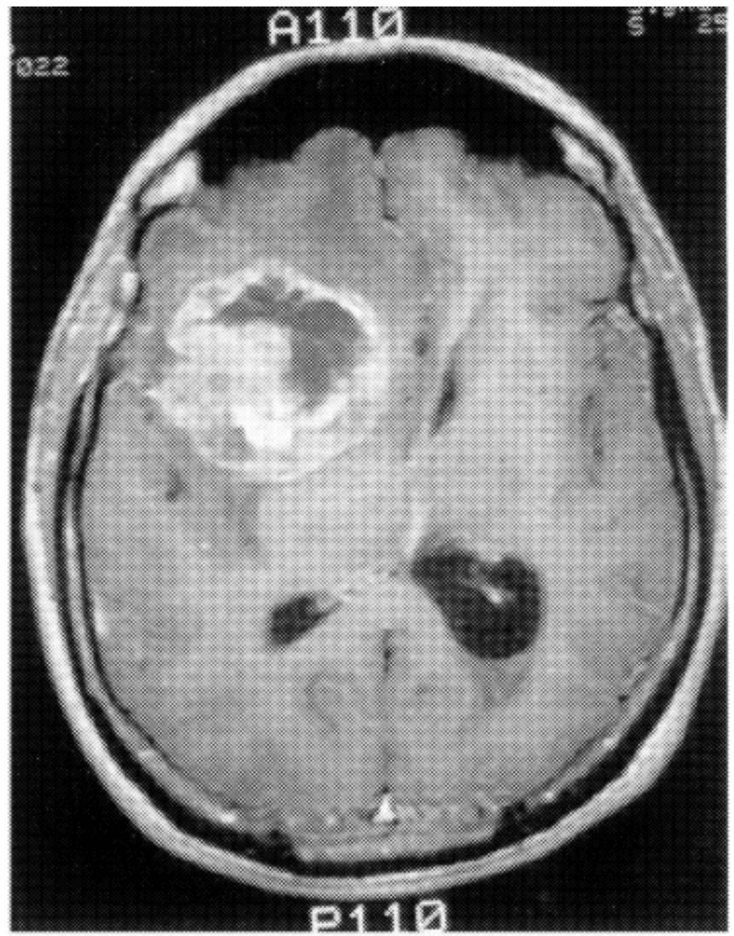

Medical Image Processing

Segmentation

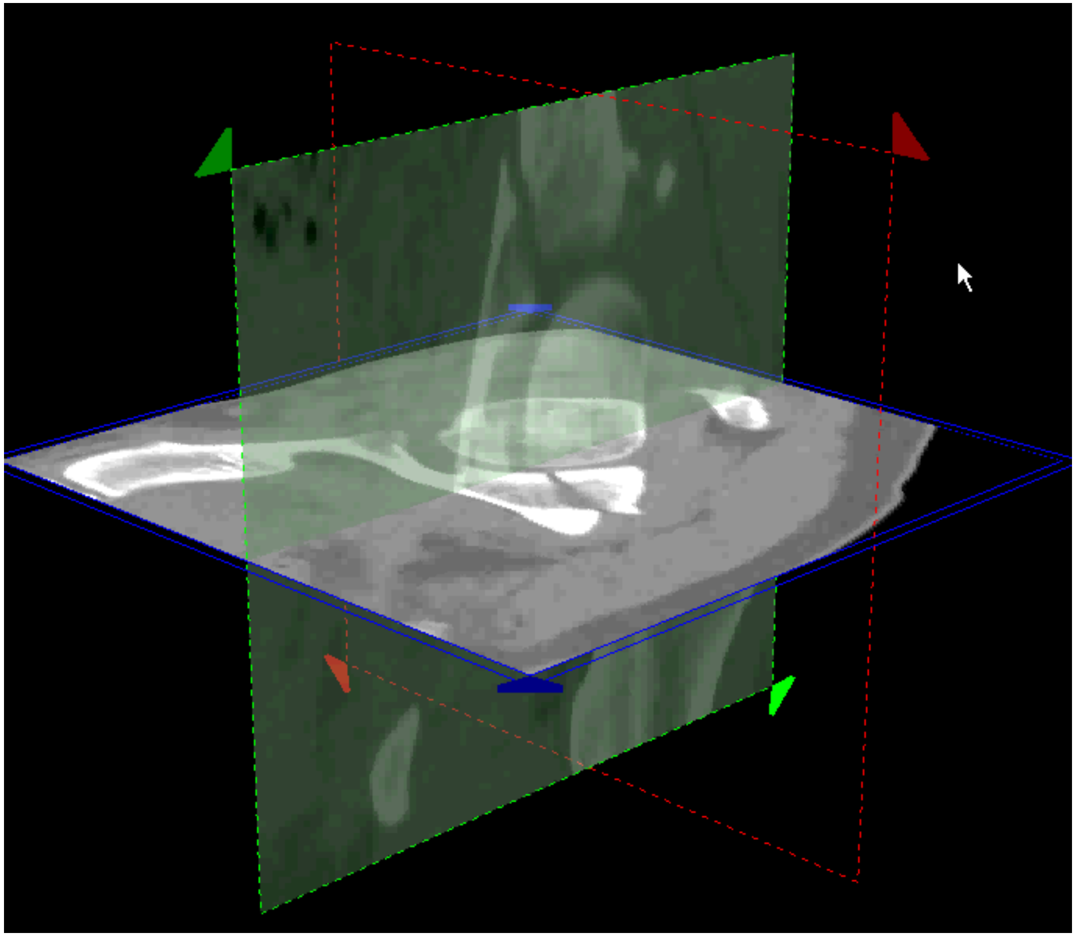

Matting

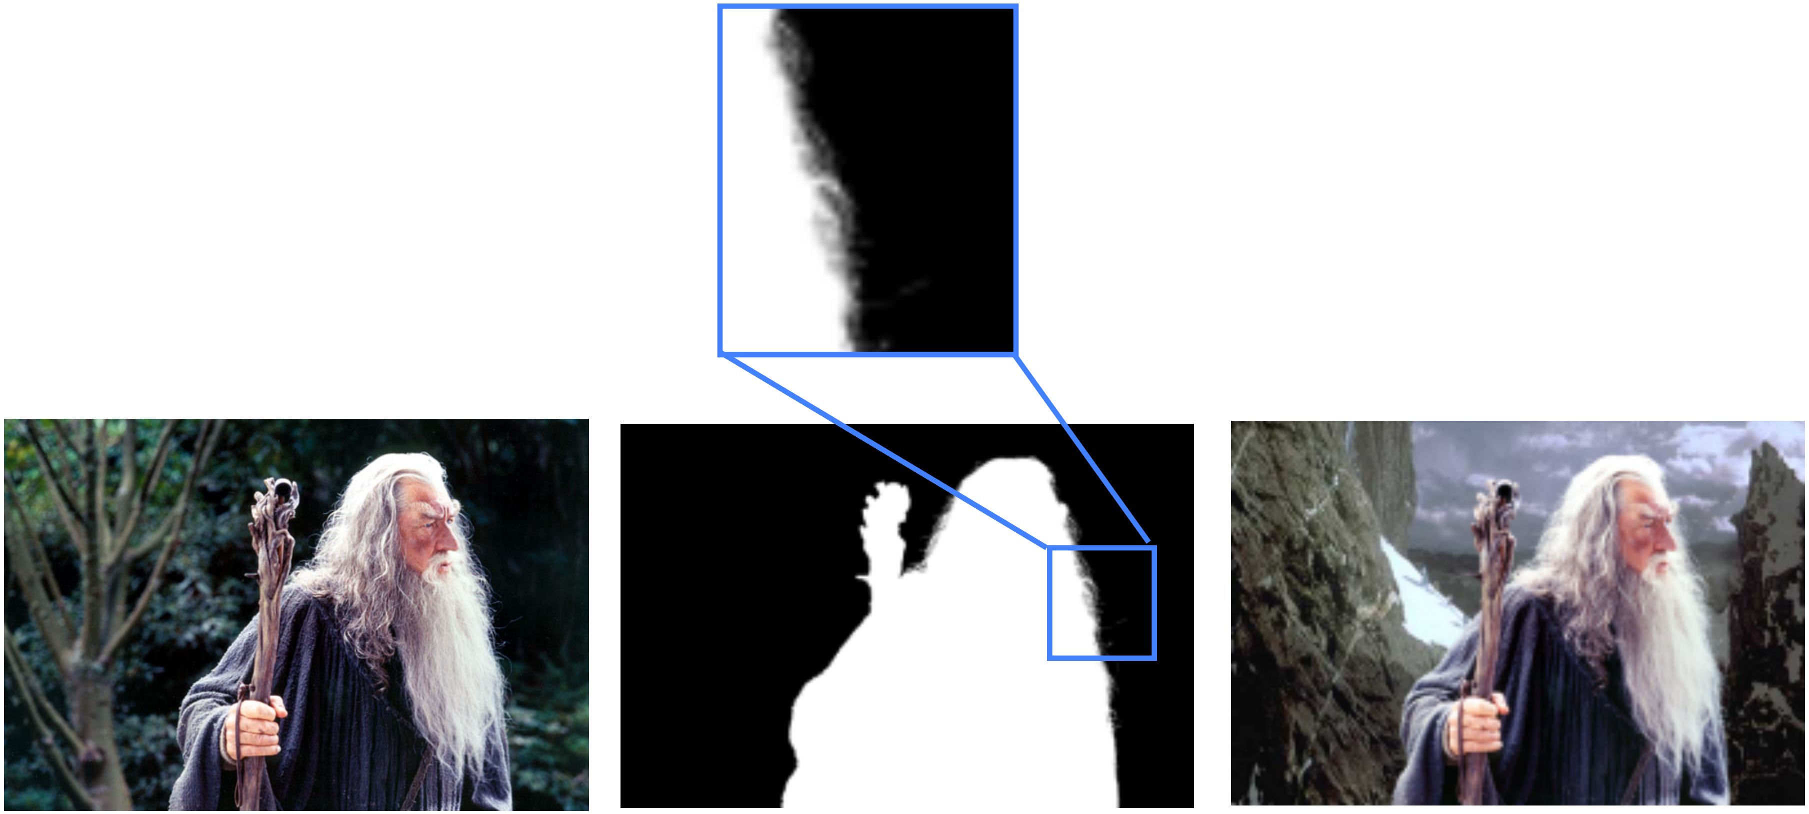

Input Matting Composition

Capturing Light Field

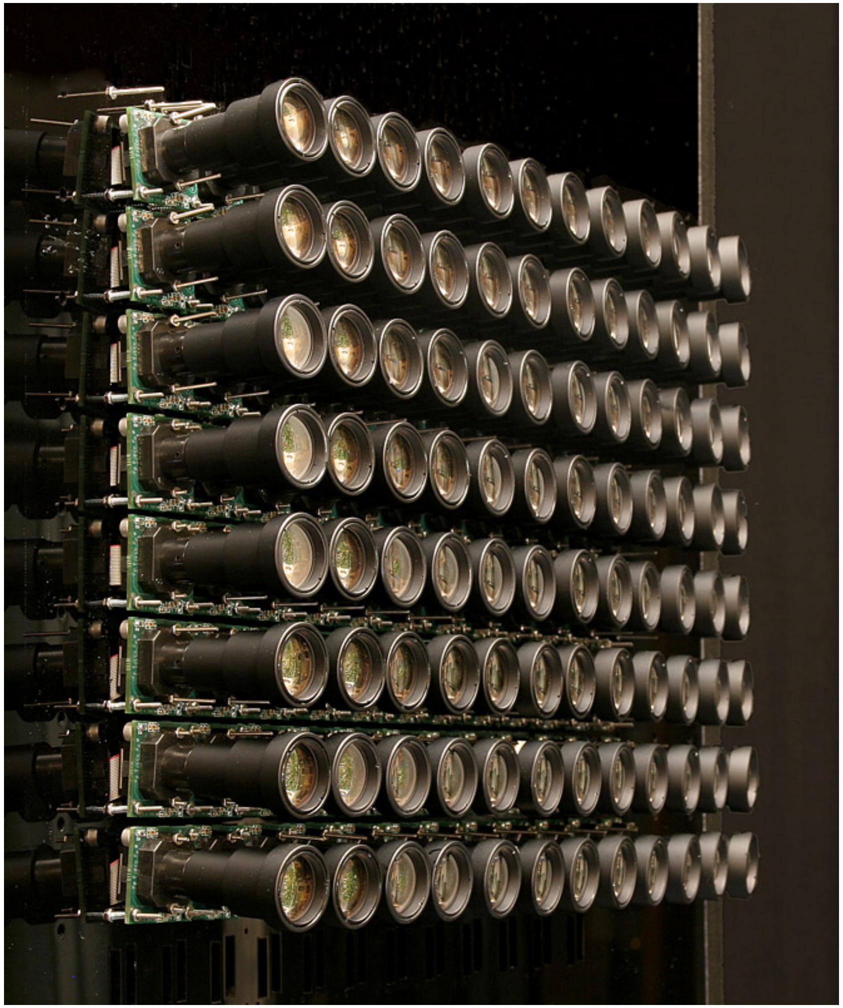

Camera Arrays, Graphics Lab, Stanford University

Capturing Light Field

Applications: synthetic aperture imaging

crowd0-parallax.mov

bike-sap.mov

Camera Arrays, Graphics Lab, Stanford University

Capturing Light Field

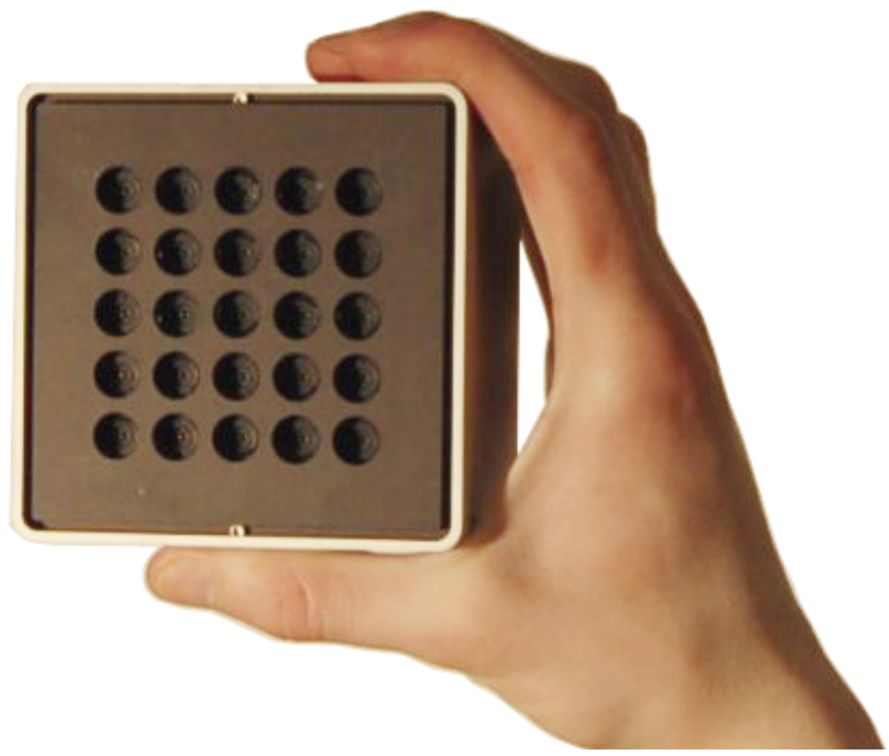

Portable Camera Arrays, U of Wisc

Dynamic Scene: Video Game

1 of 5 Input Views

Our Result

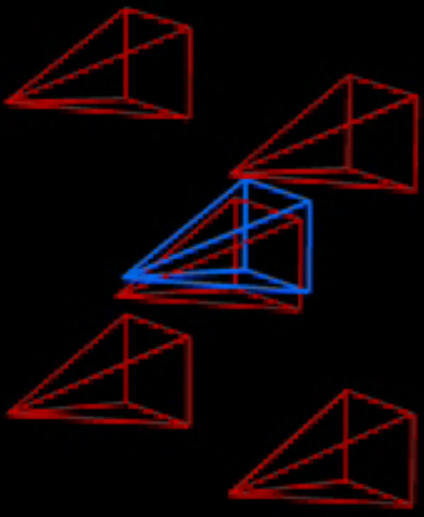

Camera Array vs Virtual Camera

Dynamic Video: Video Game

Our Result

iMovie

Dynamic Scene: Crowd

1 of 5 Input Views

Our Result

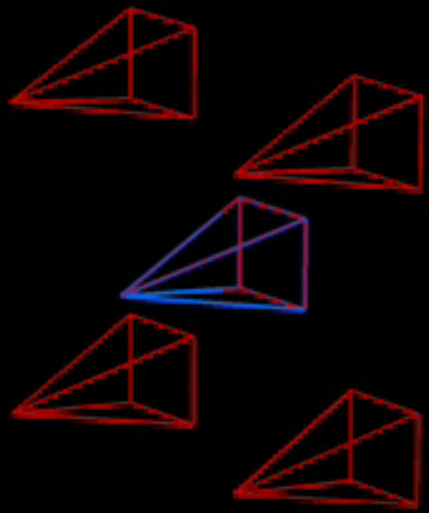

Dynamic Video: Crowd

Our Result

iMovie

Stable 3D Videos

Augmented 3D Videos

Structured Light and Ranging Scanning

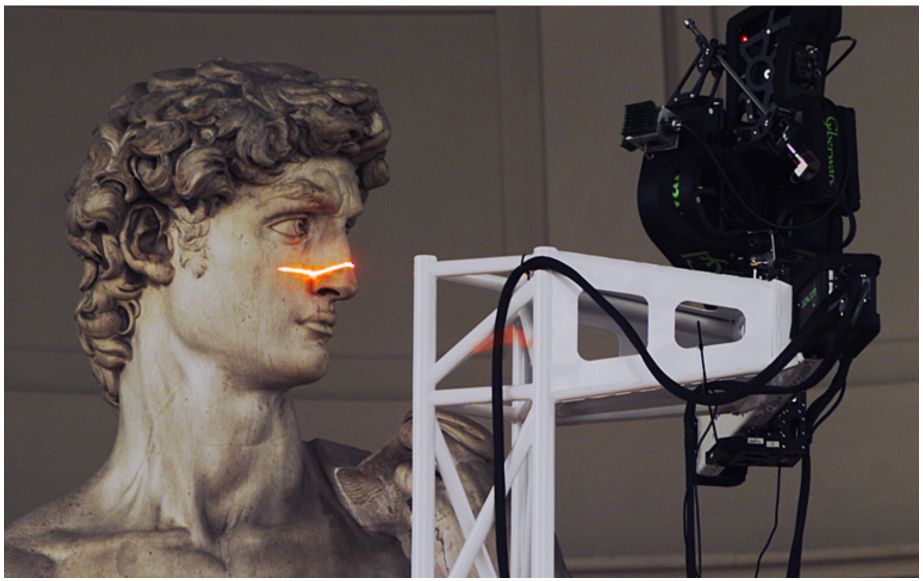

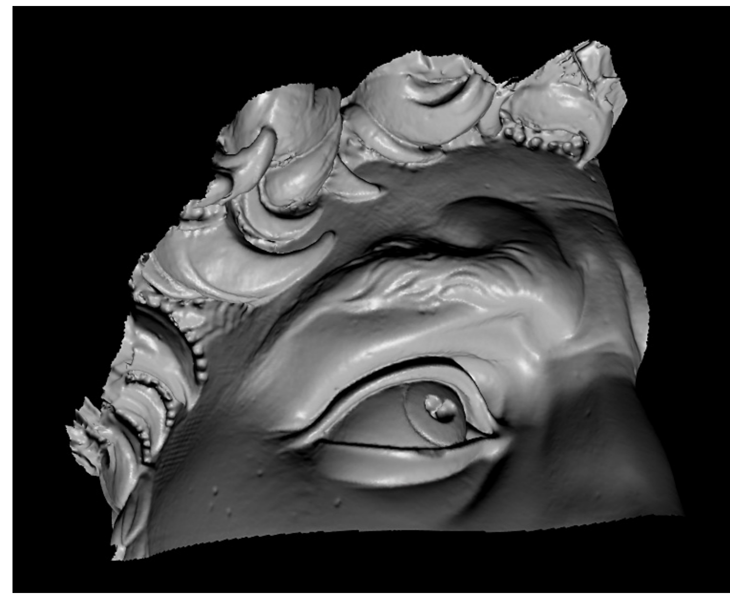

http://graphics.stanford.edu/projects/mich/

Structured Light and Ranging Scanning

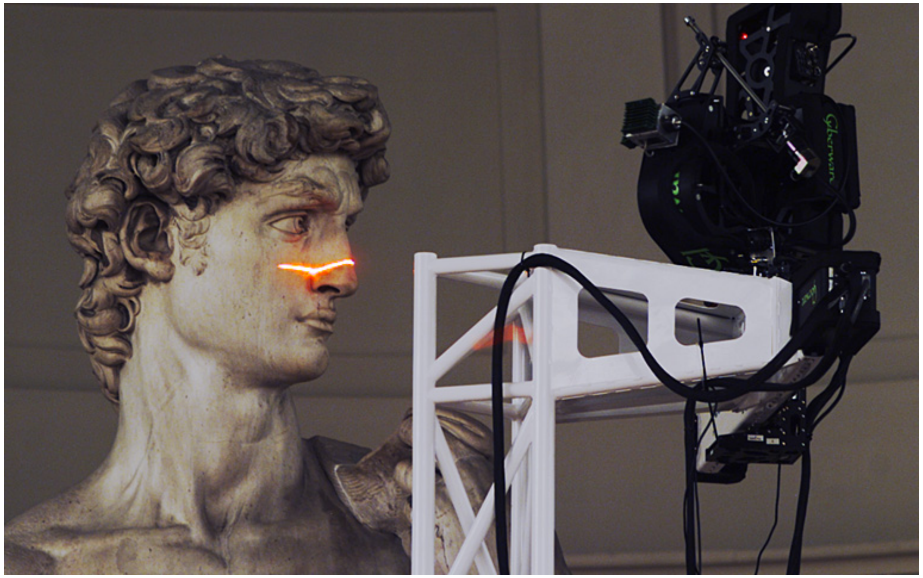

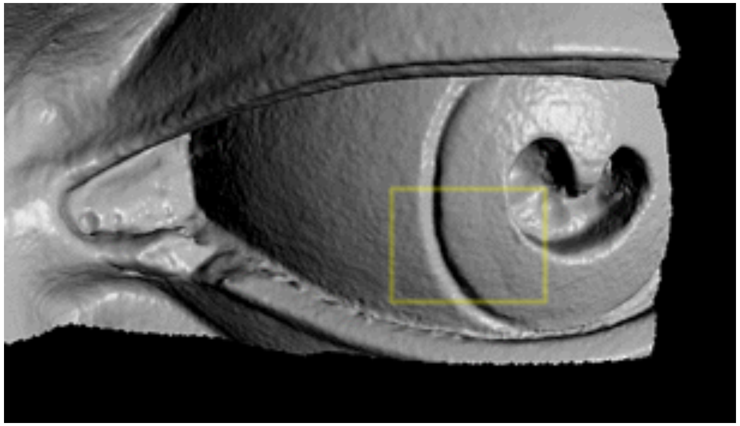

http://graphics.stanford.edu/projects/mich/

Structured Light and Ranging Scanning

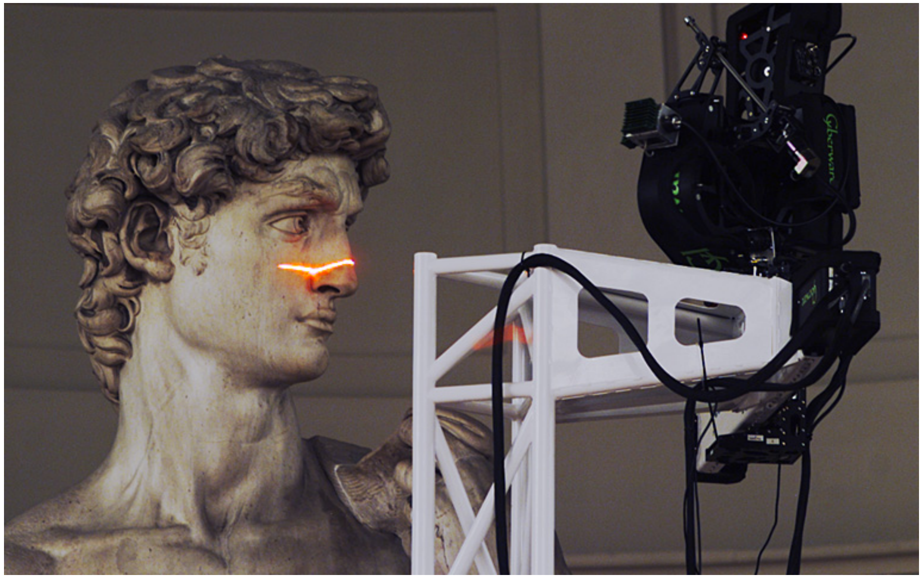

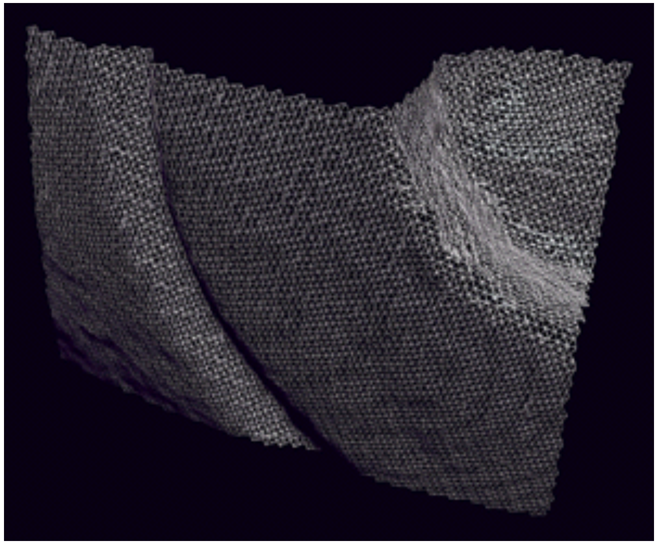

http://graphics.stanford.edu/projects/mich/

Novel Cameras and Displays

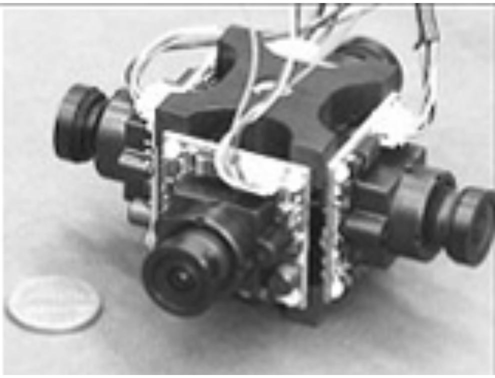

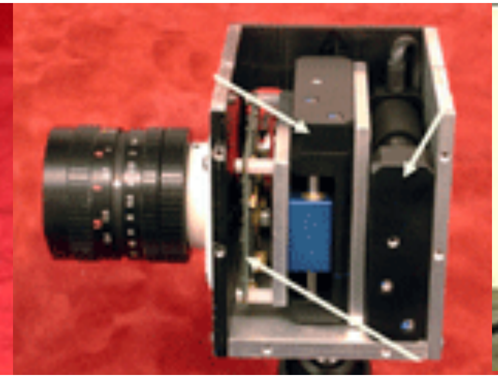
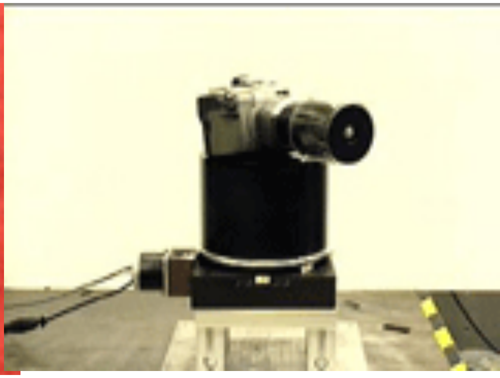
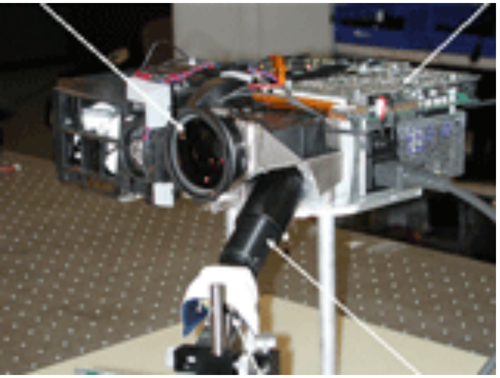


http://www1.cs.columbia.edu/CAVE/projects/cc.htm

Assignment 0, Imagination

- Due next Thursday
- Give THREE interesting things that you may wish to do with images
  - Better Image Capture
  - Making use of images
  - Design imaging systems
  - Combining vision methods in your research area

Course Info

http://www.cs.wisc.edu/~cs766-1/