Computer Vision, CS766

Staff

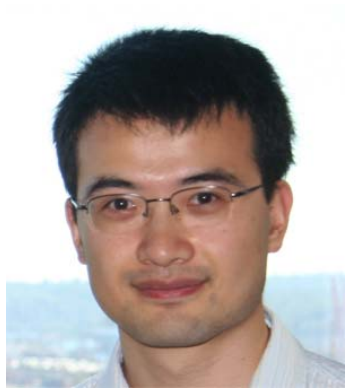

Instructor: Li Zhang
lizhang@cs.wisc.edu

TA: Yu-Chi Lai
yu-chi@cs.wisc.edu

Today

Introduction

Administrative Stuff

Overview of the Course

About Me

- Li Zhang (张力)
  - Last name pronounced as Jung
- New Faculty
  - PhD 2005, U of Washington
  - Research Scientist 06-07, Columbia U
- Research
  - Vision and Graphics
- Teaching
  - CS766 Computer Visoin
  - CS559 Computer Graphics

Previous Research Focus

- 3D shape reconstruction

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

Previous Research Focus

- 3D shape reconstruction
- Application

Licensed by SONY
for Games

Used by VA Hospital
for Prosthetics

Please tell me about you

Prerequisites

- Prerequisites—*these are essential*!
  - Data structures
  - A good working knowledge of C and C++ programming
    - (or willingness/time to pick it up quickly!)
  - Linear algebra
  - Vector calculus

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

Administrative Stuff

- **4 programming projects**
  - 15%, 2-3 weeks each
- **1 final project**
  - 40%, 5 weeks, open ended of your choosing, but needs
  - project proposal after 1 week
  - progress report after 3 weeks
  - Final presentation after 5 weeks
- **Computer account**:
  - Everyone registered in this class will get a Computer Systems Lab account to do project assignments.
- **Email list**:
  - compsci766-1-f07@lists.wisc.edu

Questions?

Every picture tells a story

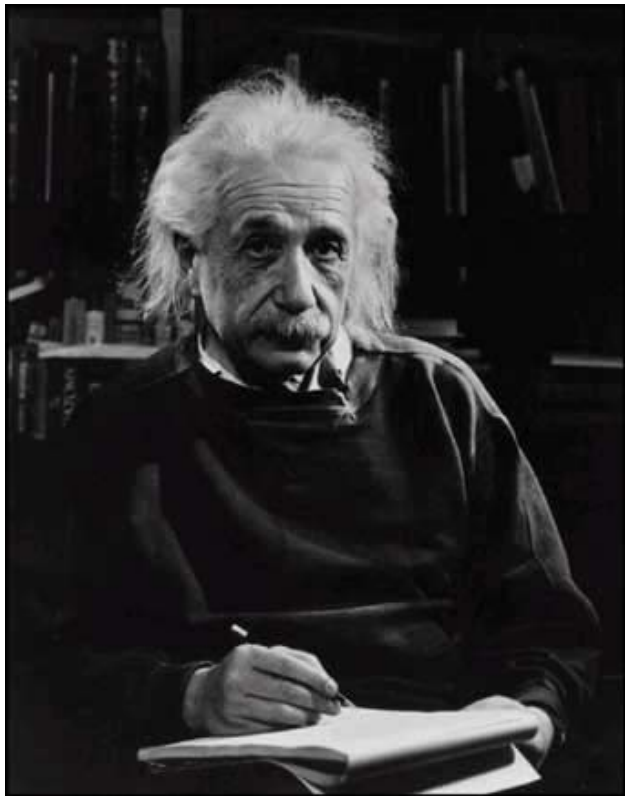

Goal of computer vision is to write computer programs that can interpret images

Can computer match human perception?

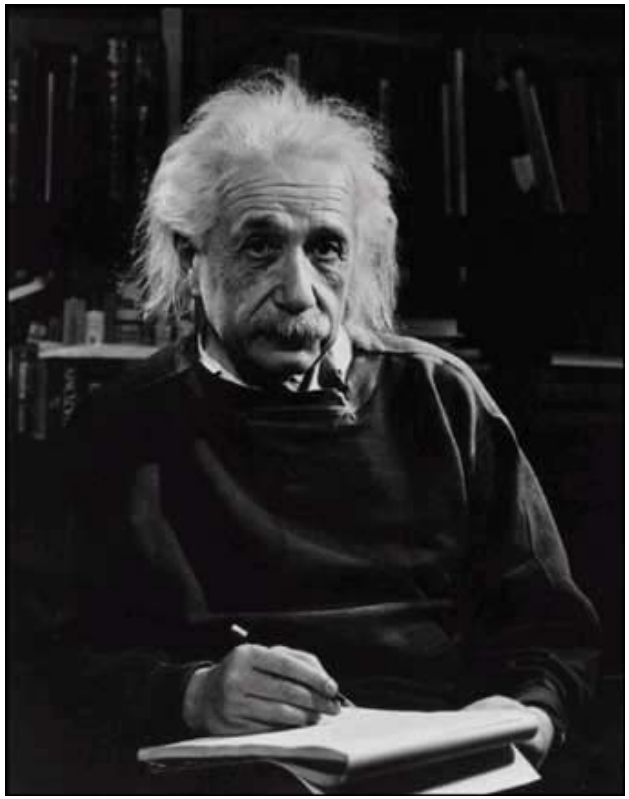

- Yes and no (but mostly no!)
  - computers can be better at “easy” things

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
  - Ambiguous description of the world
  - Doesn’t care about accuracy of world

Srinivasa Narasimhan’s slide

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

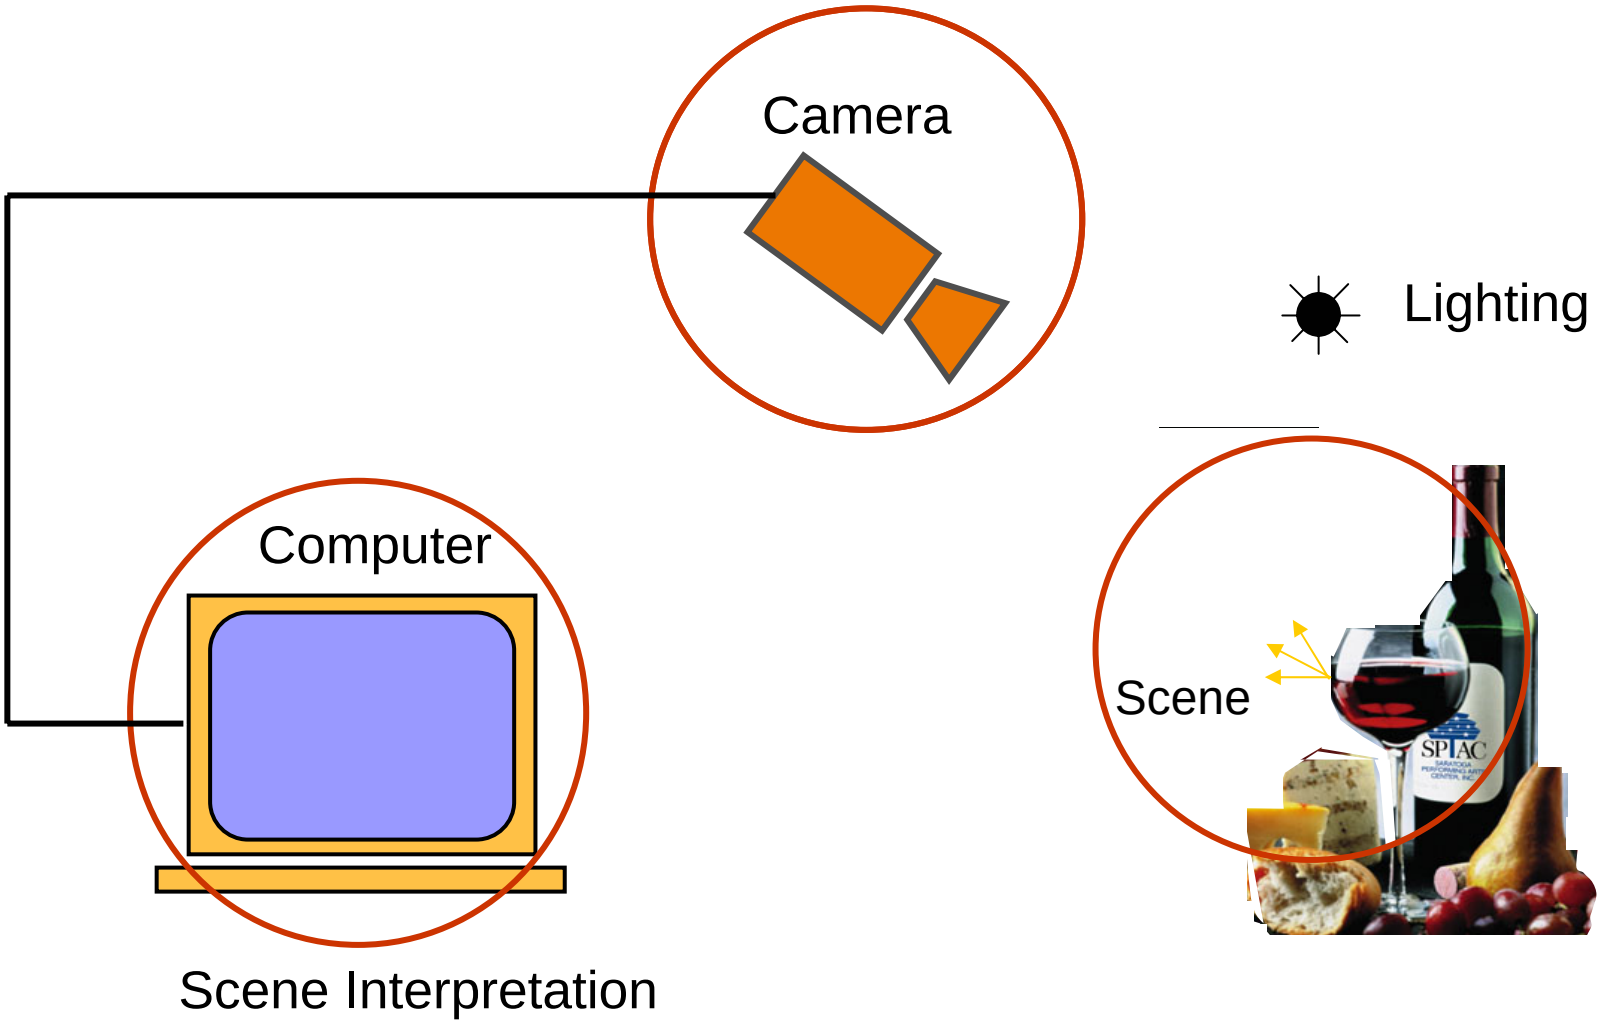

Srinivasa Narasimhan's slide

Topics Covered

Cameras and their optics

Today's Digital Cameras

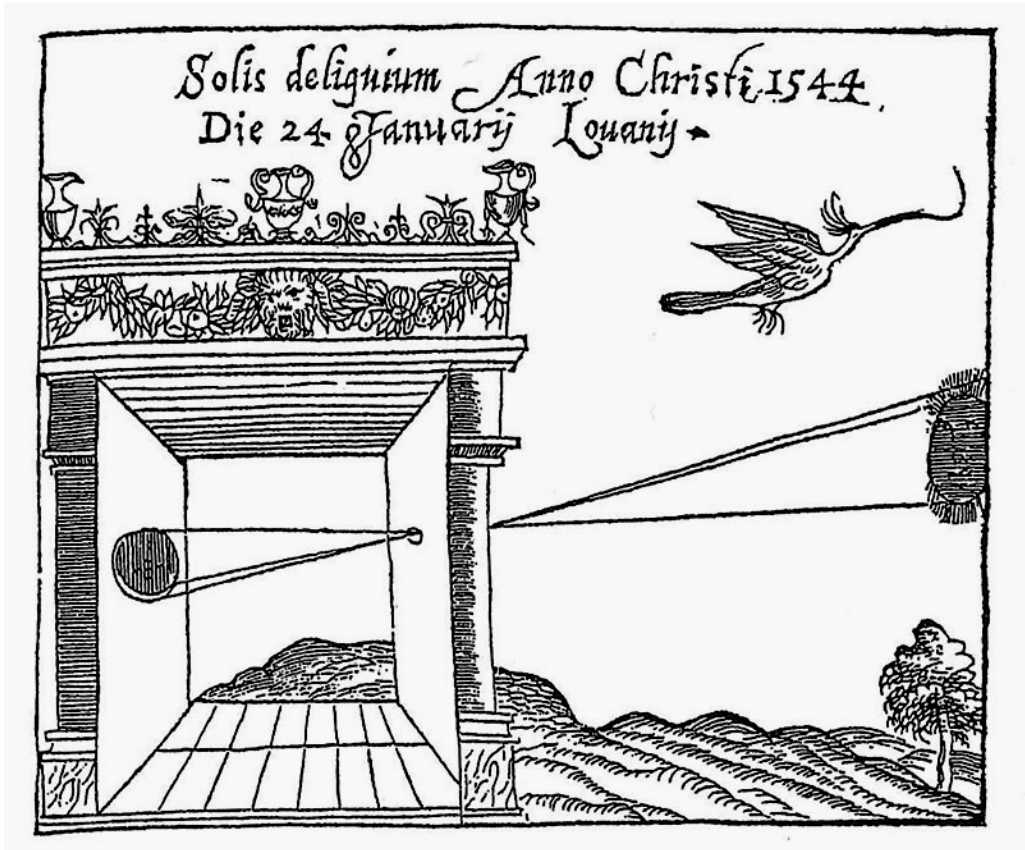

The *Camera Obscura*

Srinivasa Narasimhan's slide

Biological vision

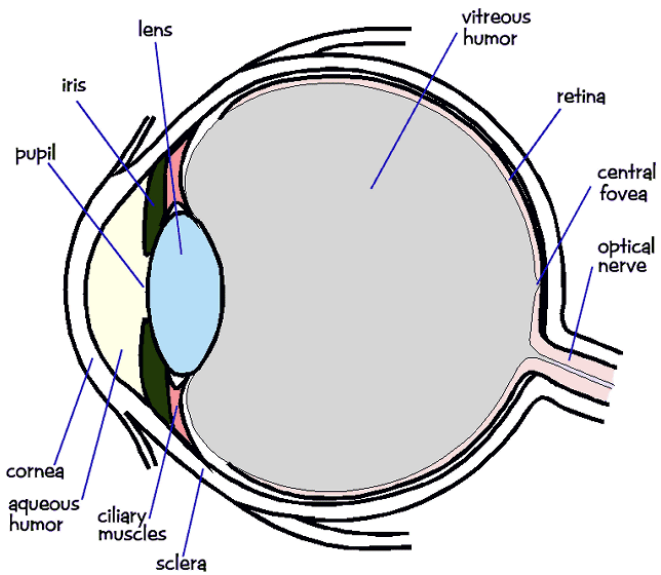

Human Eye

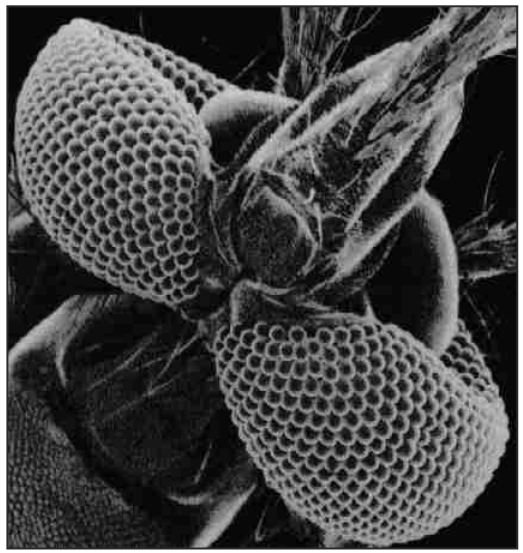

Mosquito Eye

Srinivasa Narasimhan's slide

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

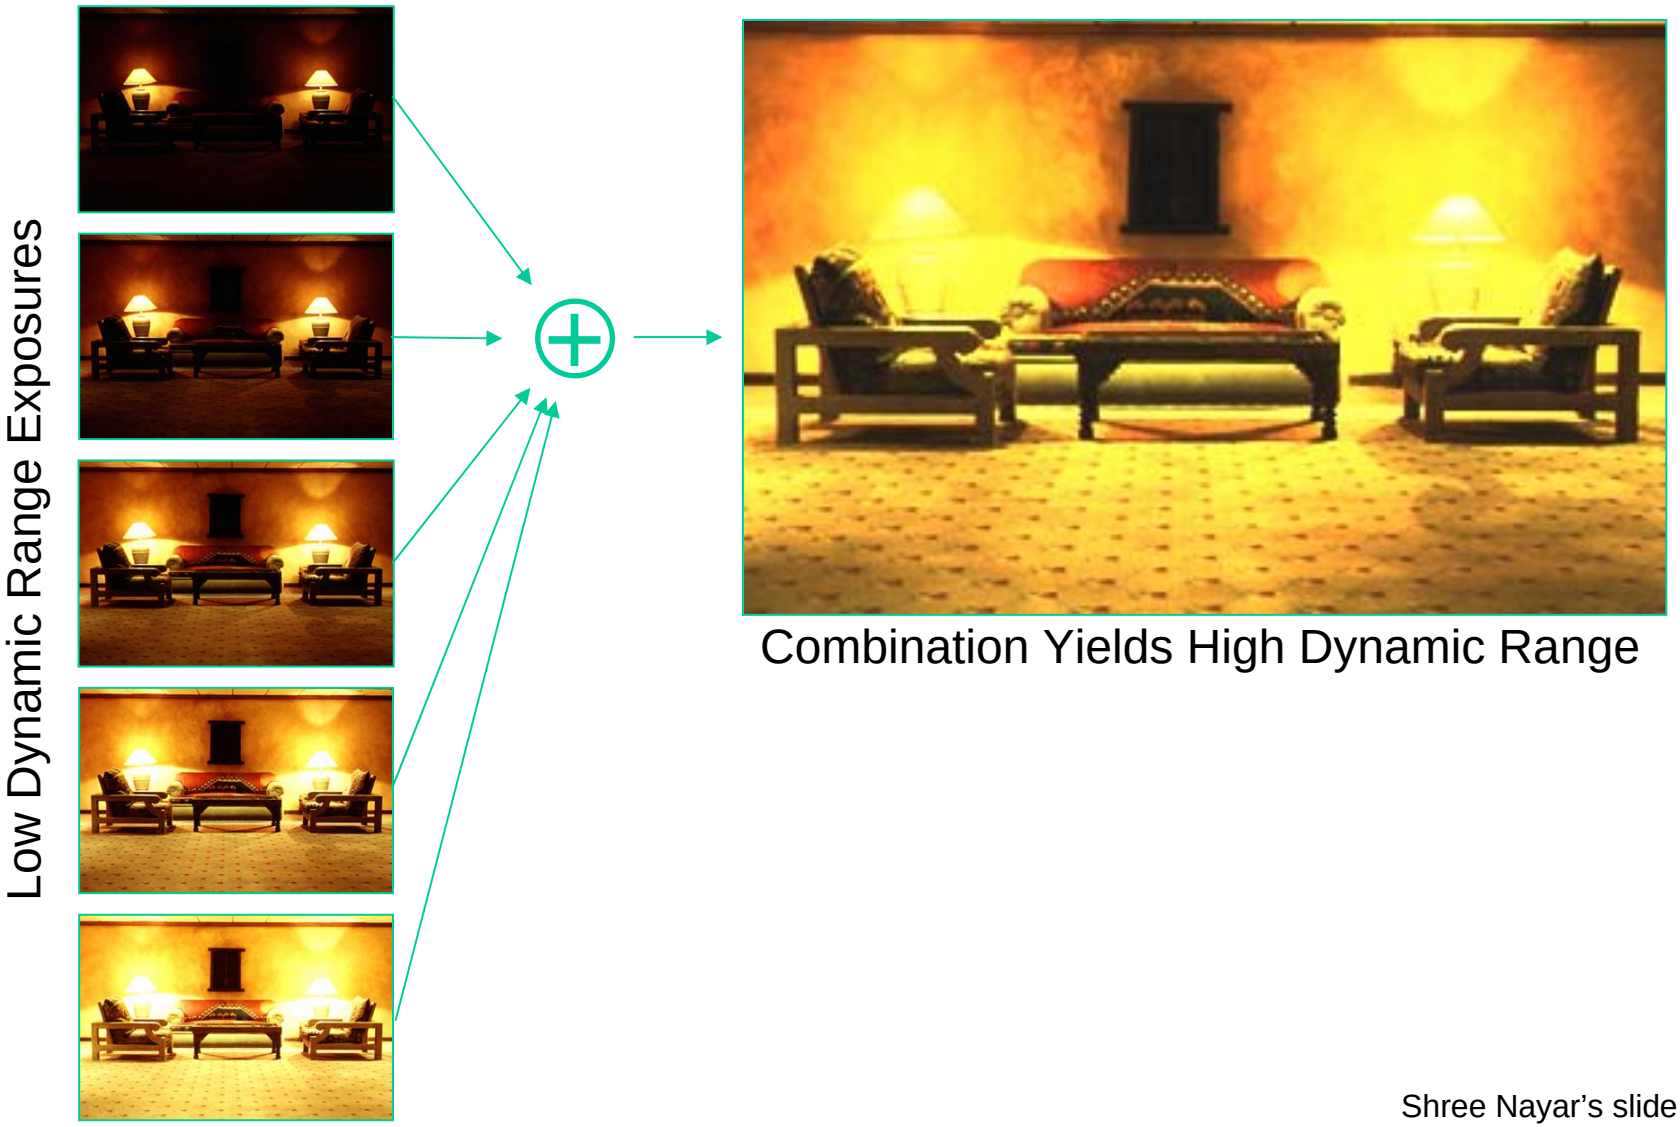

Shree Nayar's slide

Image Processing

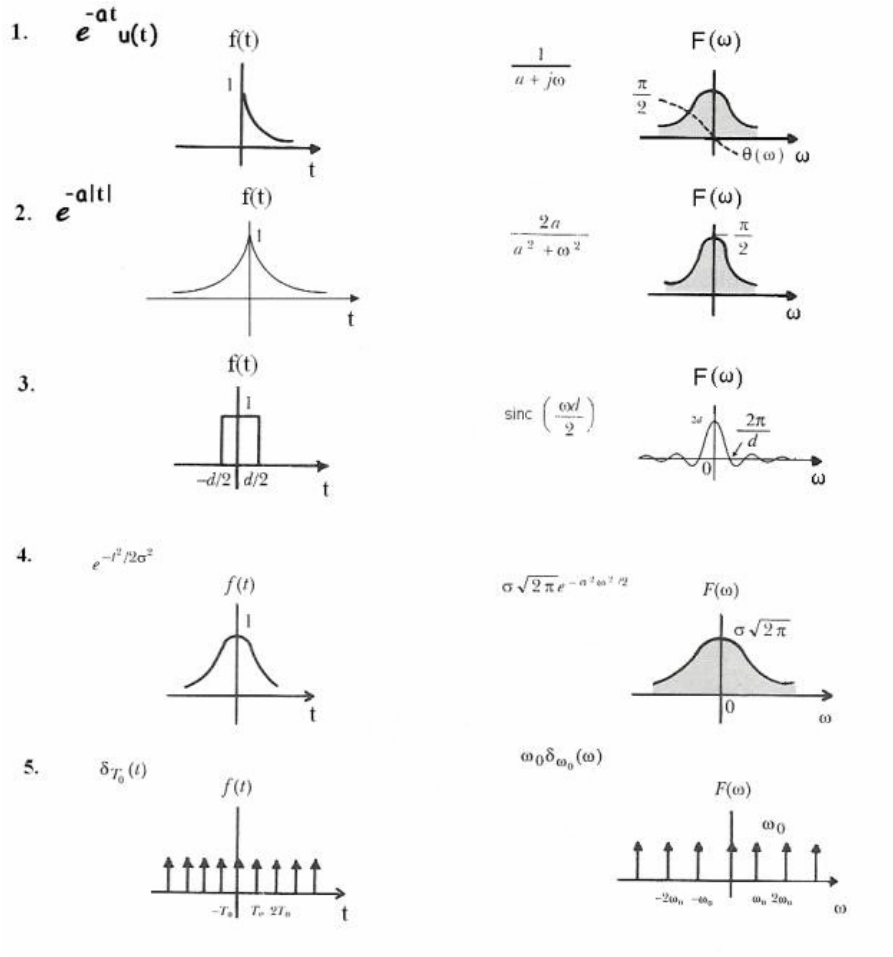

Fourier Transform
Sampling, Convolution

Image enhancement
Feature detection

Srinivasa Narasimhan's slide

Camera Projection

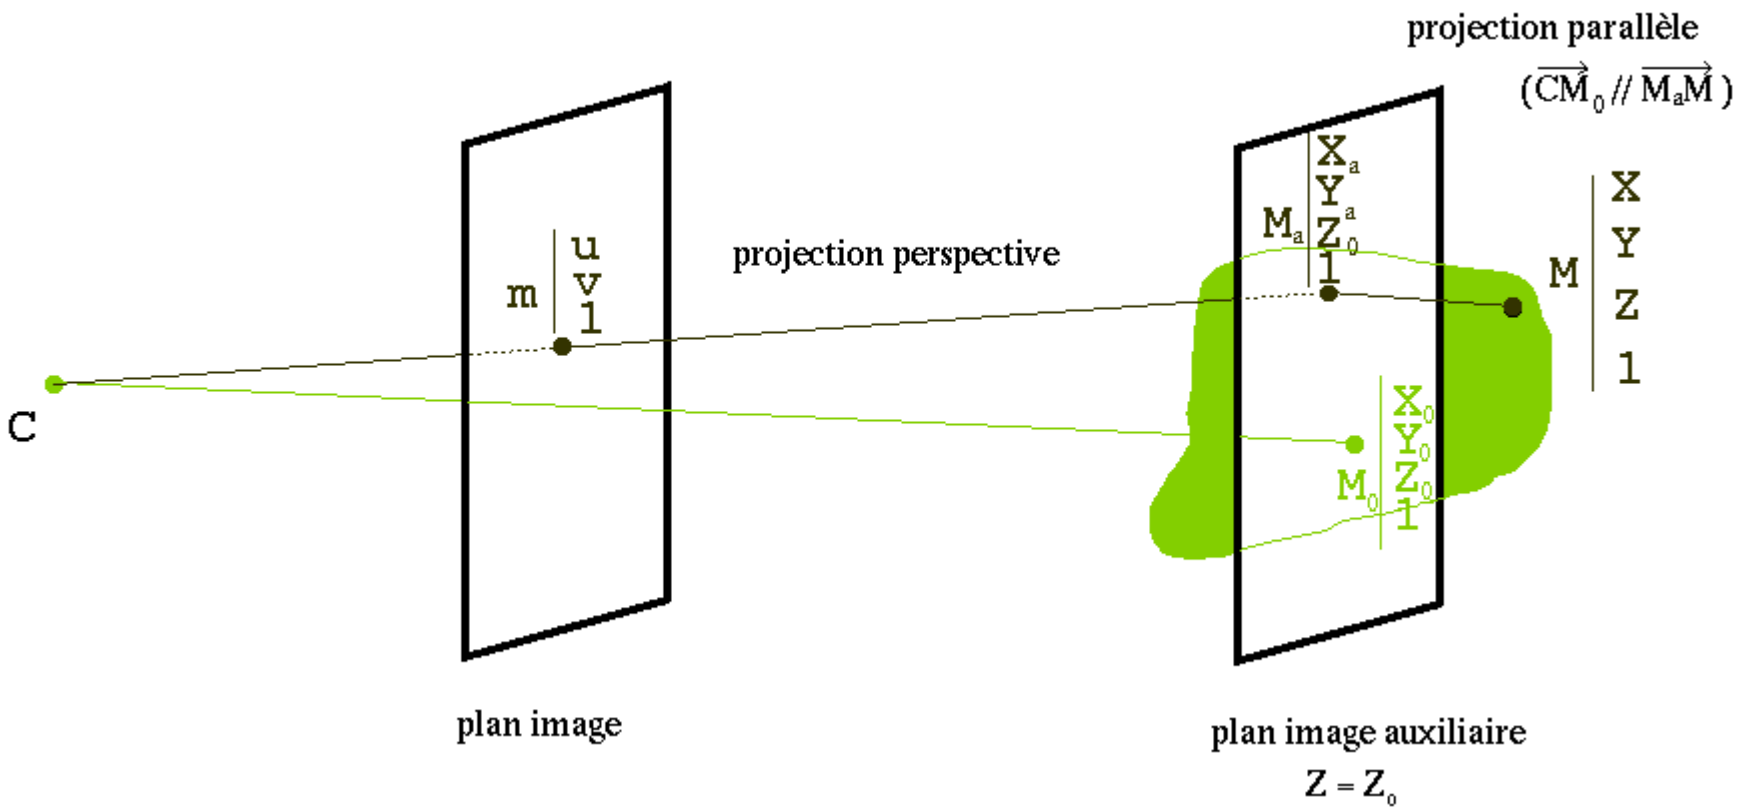

Image Transformation

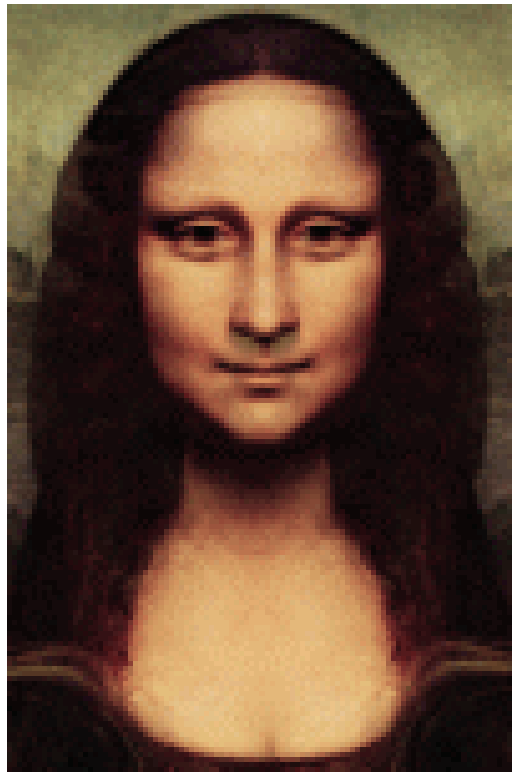

Steve Seitz and Chuck Dyer, View Morphing, SIGGRAPH 1996

Project 2: Panoramic Imaging

Input images:

Output Image:

Steve Seitz's slide

Projective Geometry

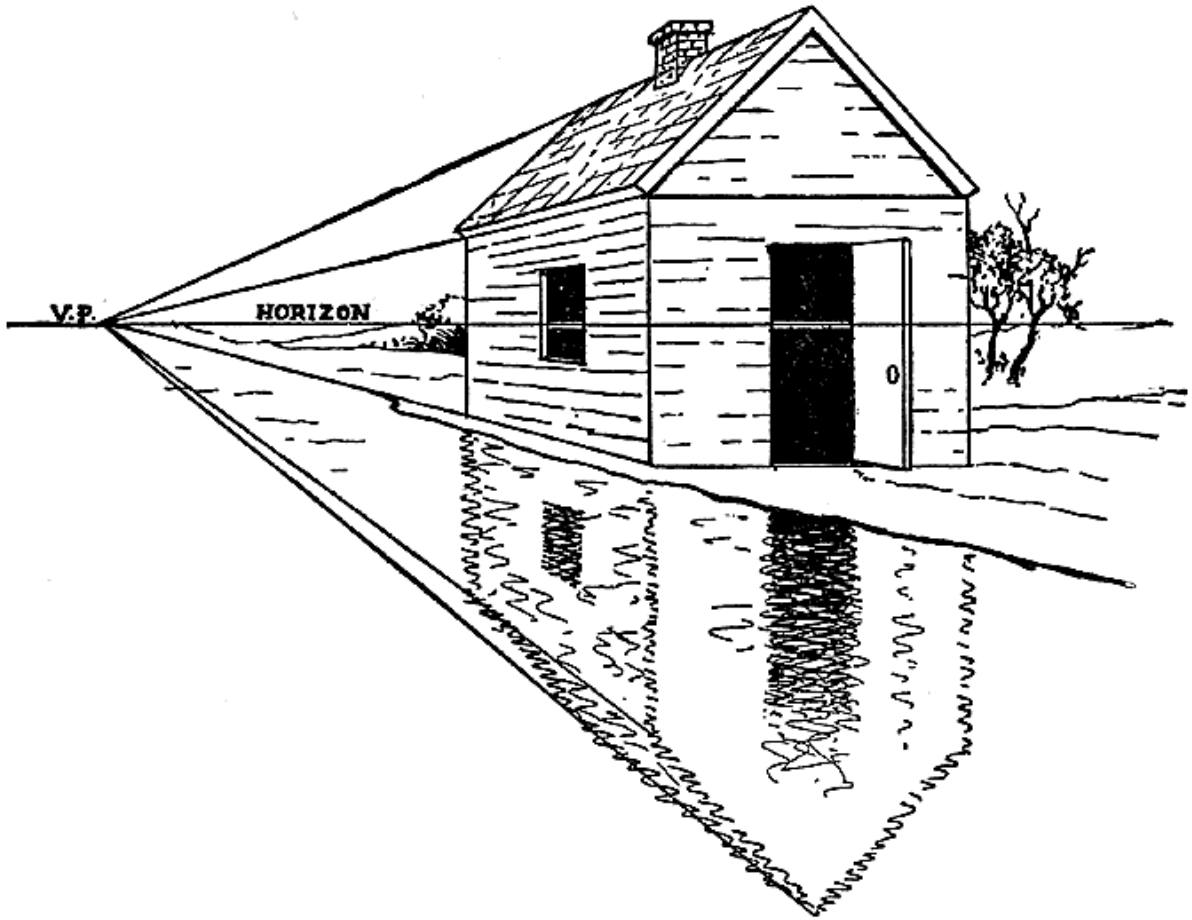

Single View Metrology

- https://research.microsoft.com/vision/cambridge/3d/3dart.htm

Single View Metrology

- https://research.microsoft.com/vision/cambridge/3d/3dart.htm

Shading and Photometric Stereo

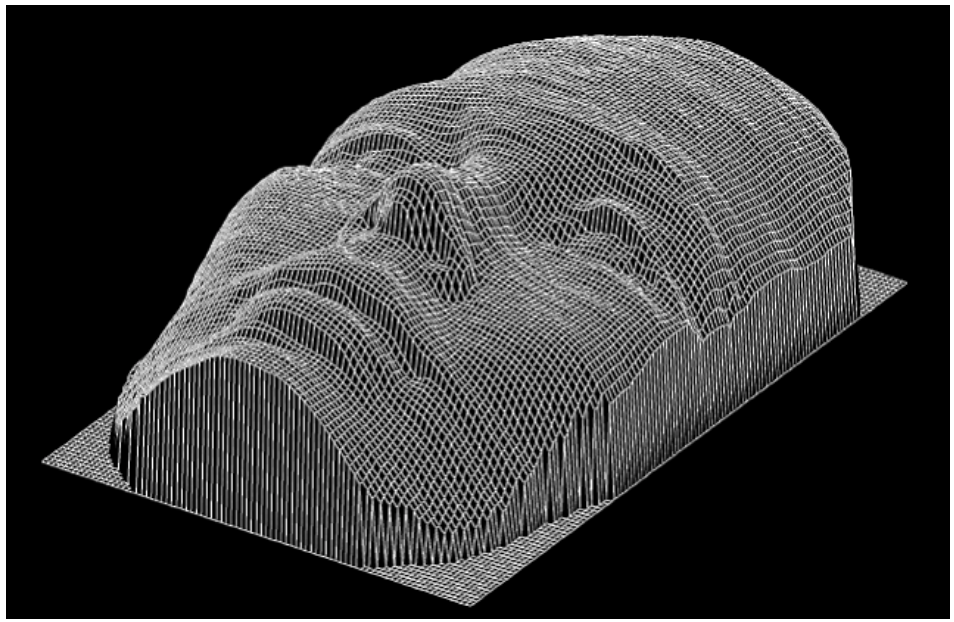

http://www.eecs.harvard.edu/~zickler/helmholtz.html

Texture Modeling

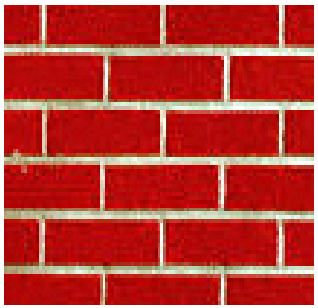

repeated

radishes rocks yogurt

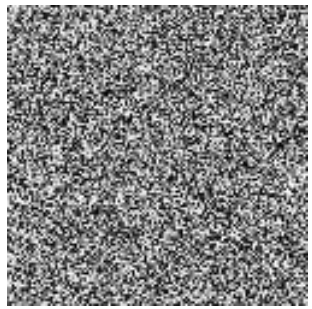

stochastic

“Semi-stochastic” structures

Alexei Efros’ slide

Project 3: Texture Synthesis

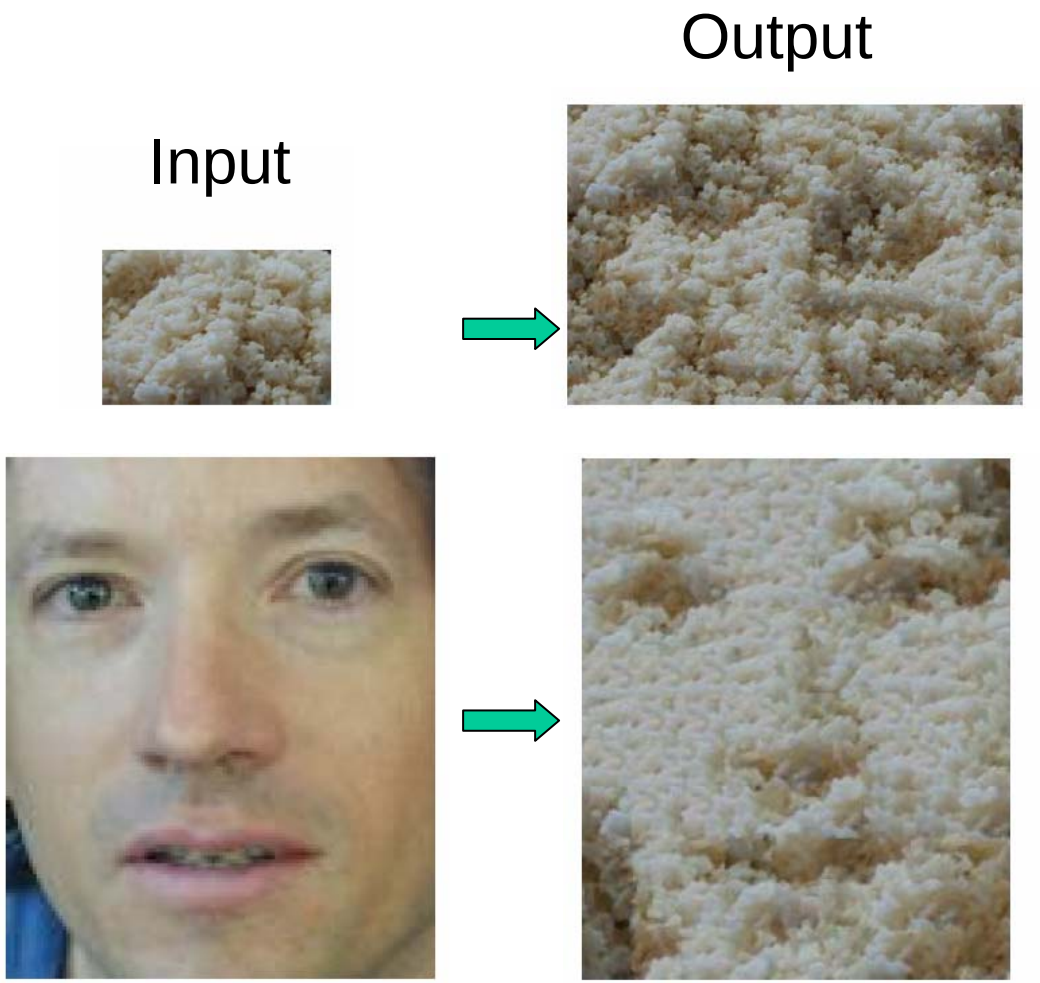

Image Quilting, Efros and Freeman., SIGGRAPH 2002.

Project 3: 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)

Face Detection and Recognition

Project 4: EigenFaces

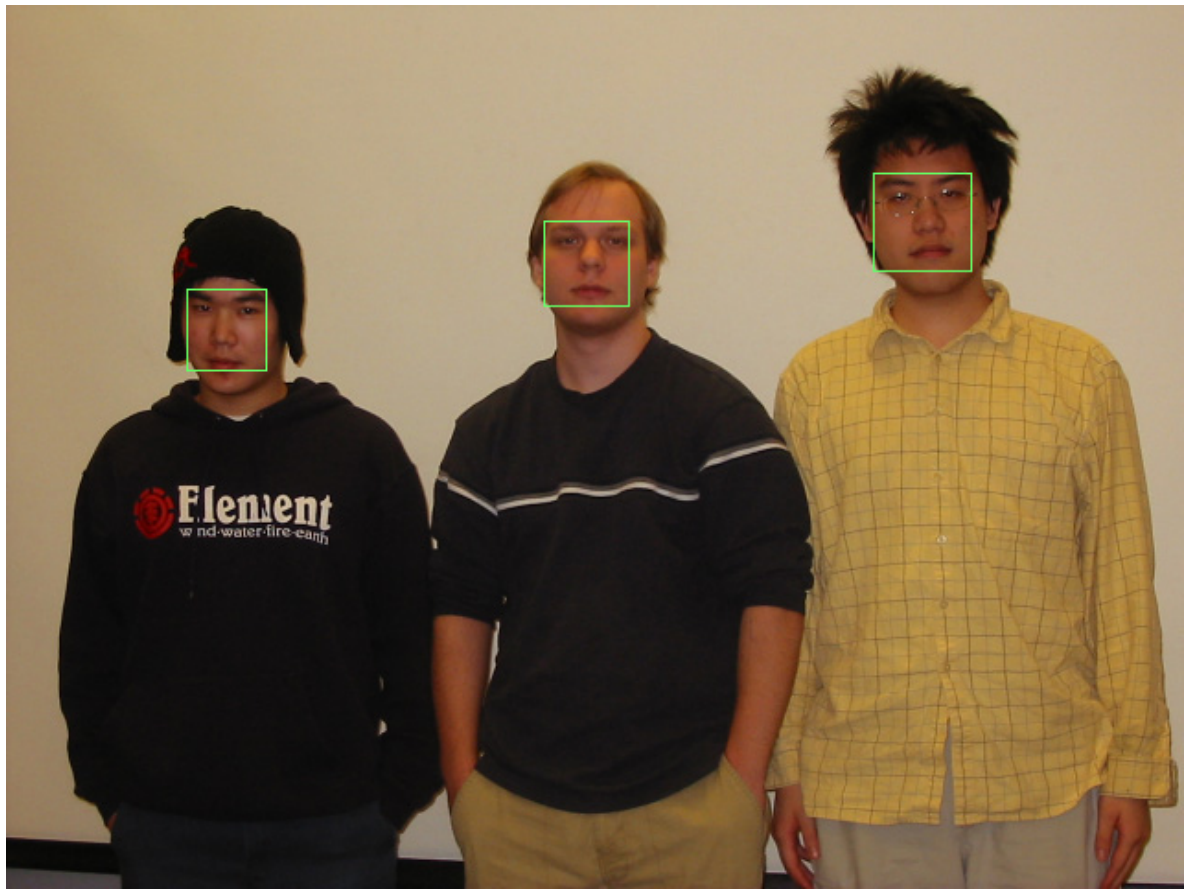

Face detection and recognition

Motion Estimation

Hidden Dragon Crouching Tiger

Motion Estimation

Application

Andy Serkis, Gollum, Lord of the Rings

Segmentation

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

Segmentation

Application

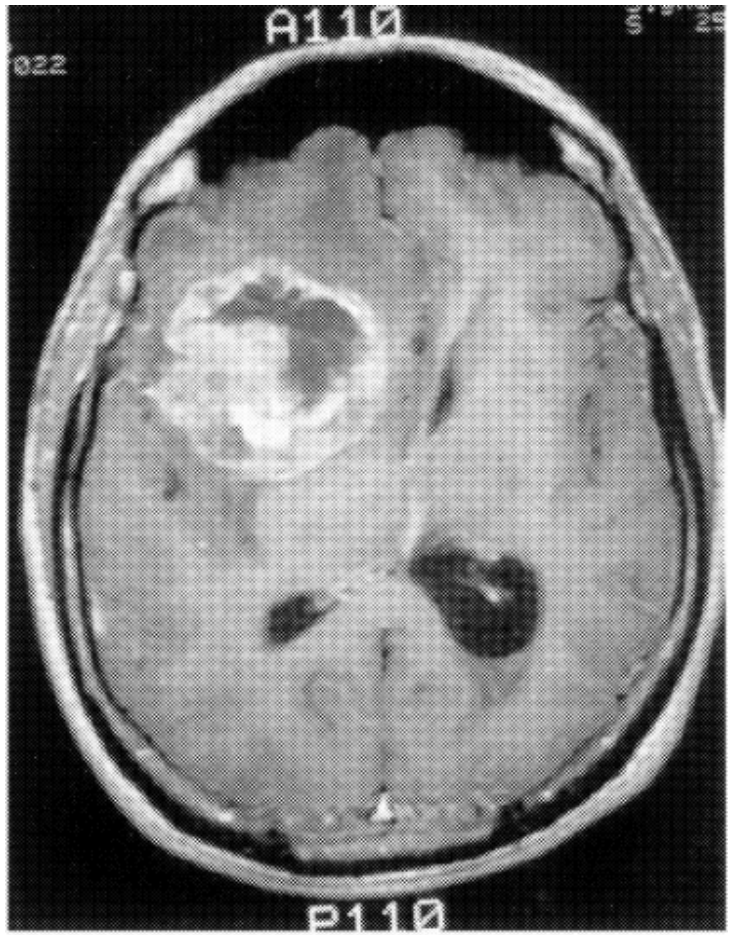

Medical Image Processing

Matting

Input

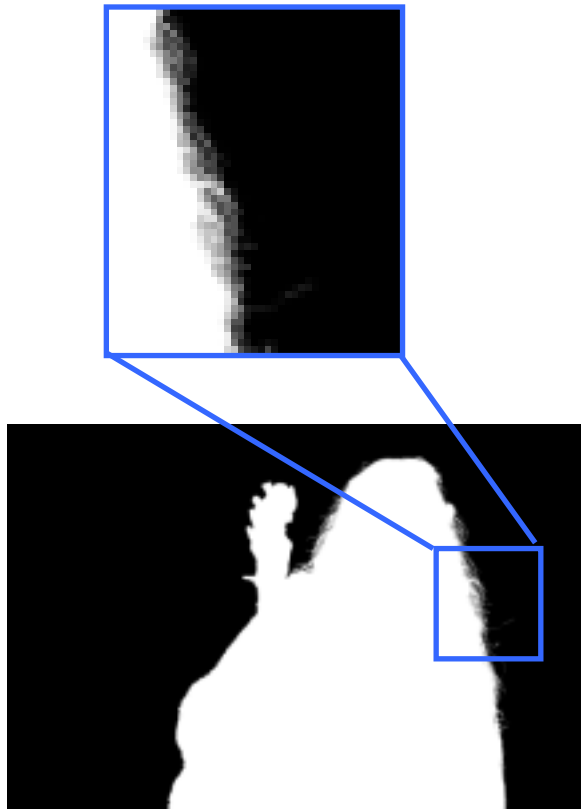

Matting

Composition

Light, Color, and Reflection

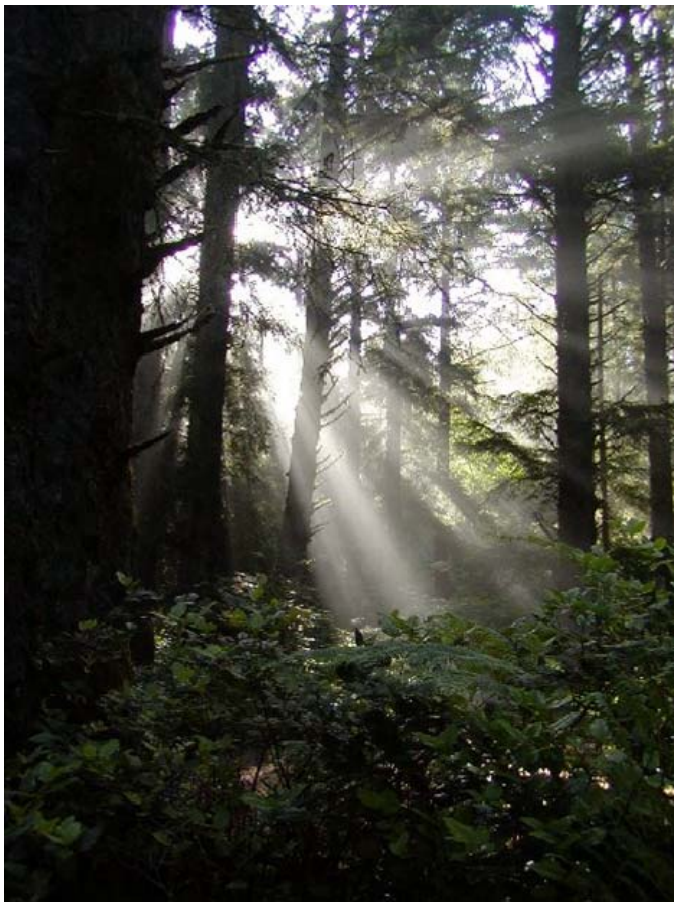

Capturing Light Field

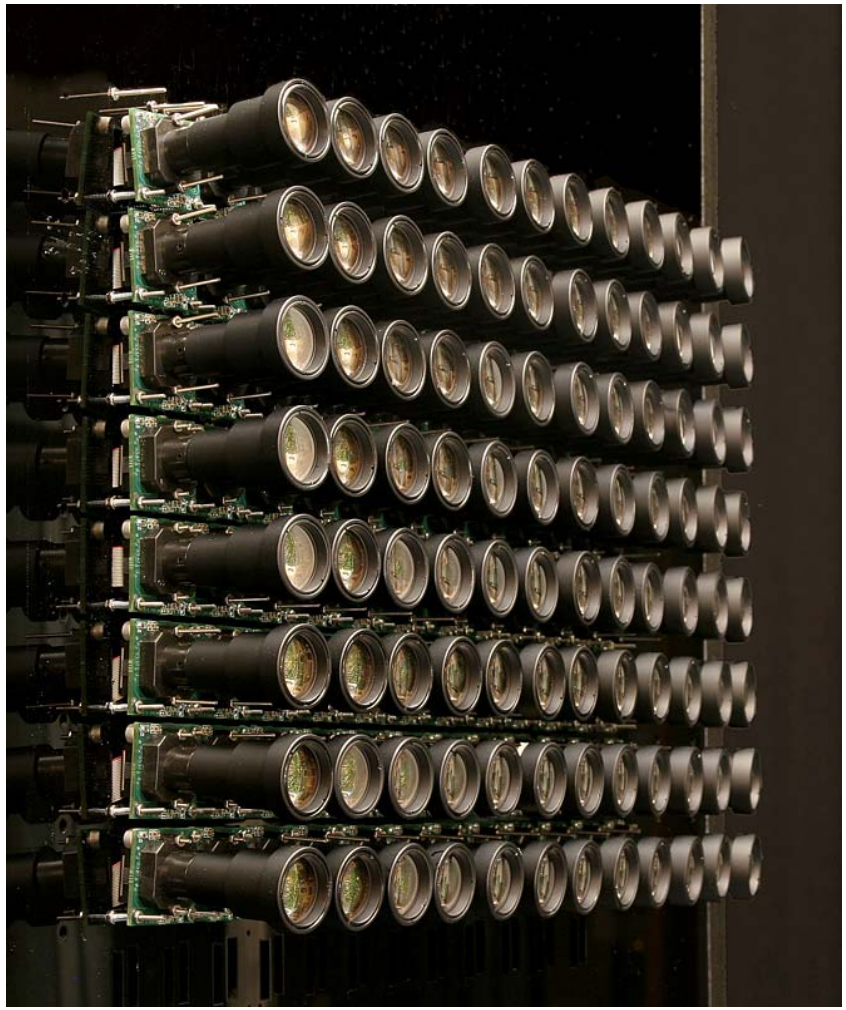

Camera Arrays, Graphics Lab, Stanford University

Capturing Light Field

Applications

Structured Light and Ranging Scanning

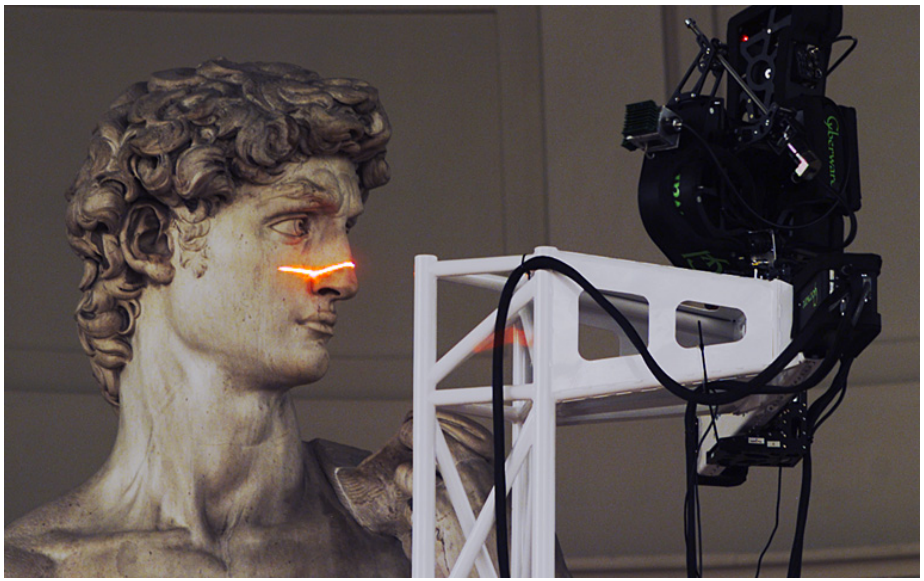

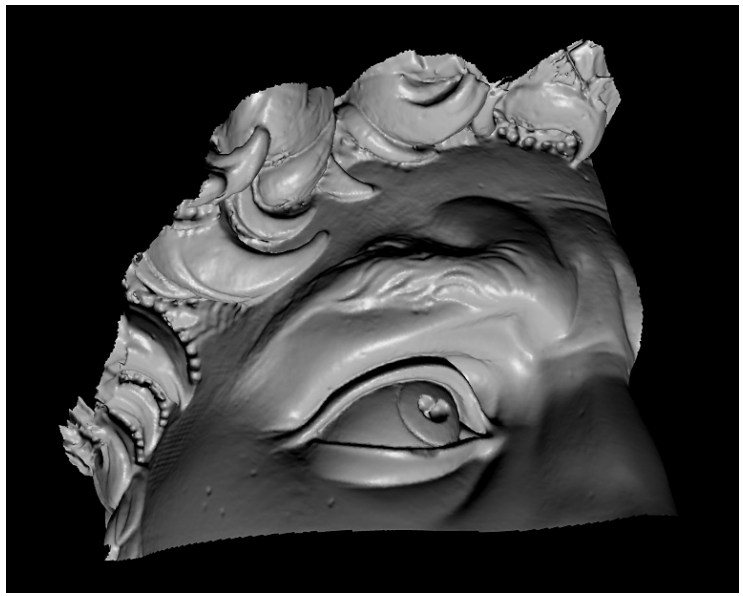

http://graphics.stanford.edu/projects/mich/

Structured Light and Ranging Scanning

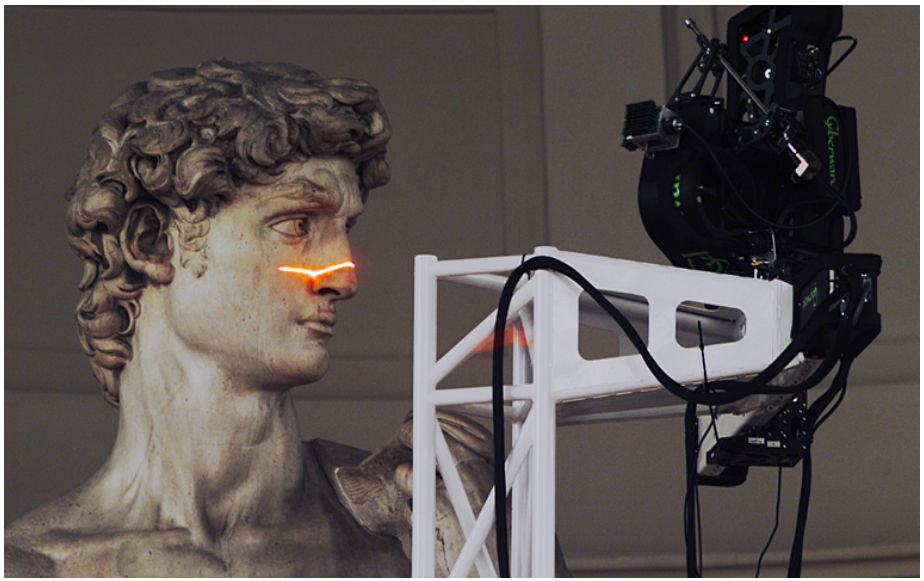

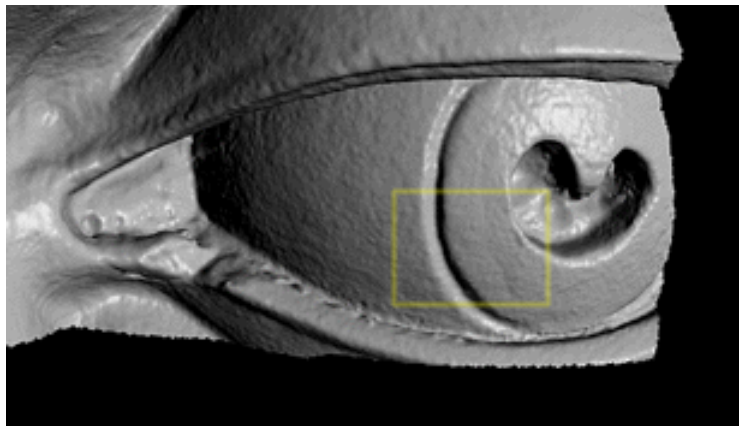

http://graphics.stanford.edu/projects/mich/

Structured Light and Ranging Scanning

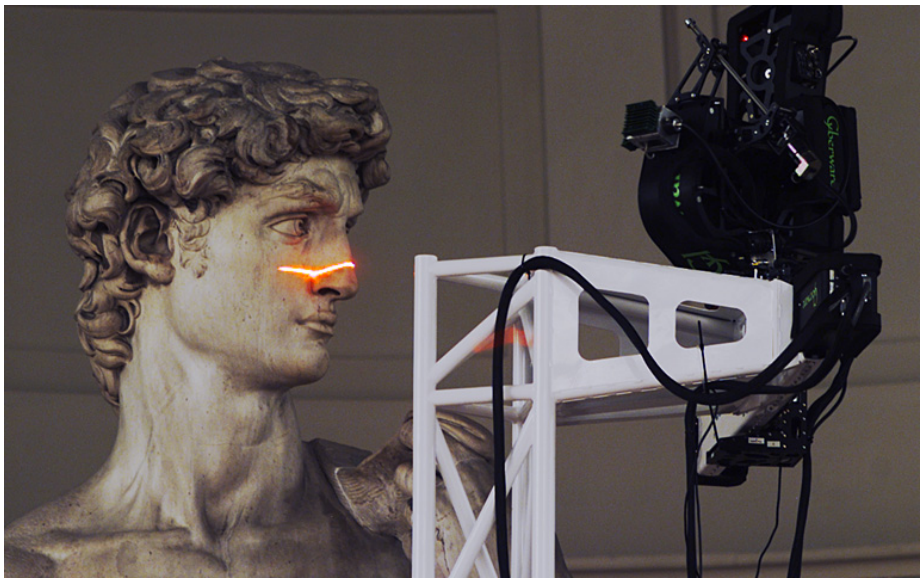

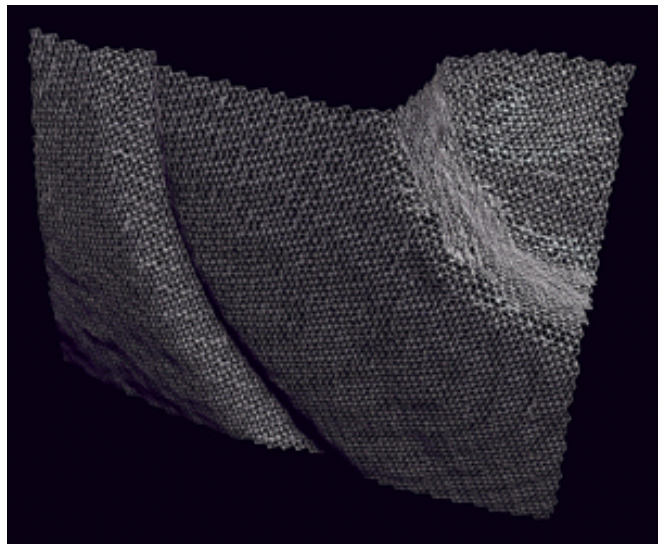

http://graphics.stanford.edu/projects/mich/

Novel Cameras and Displays

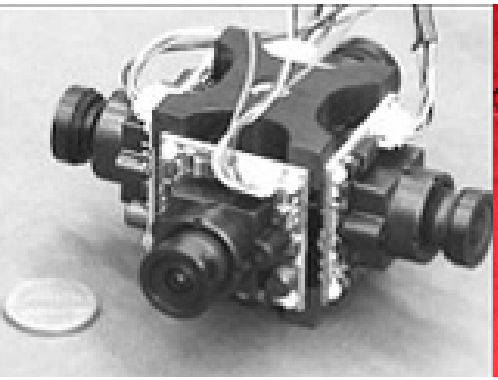

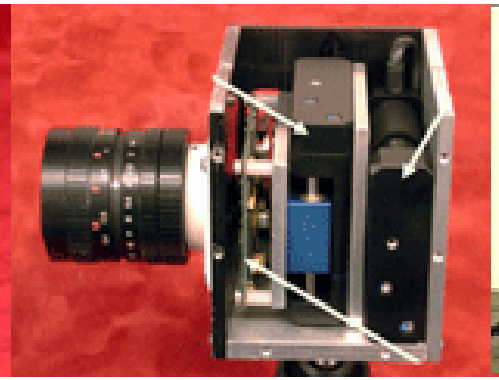
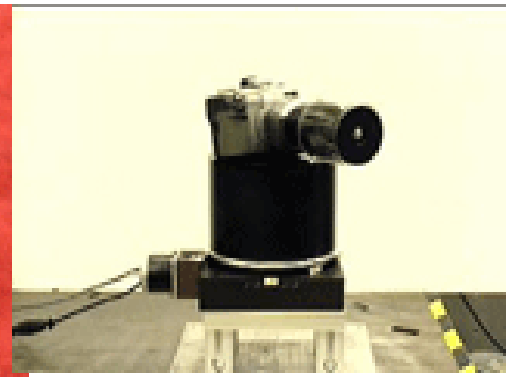
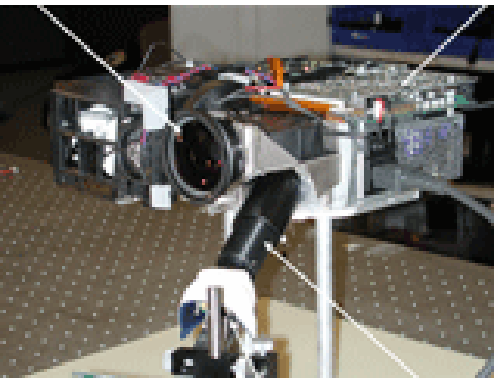


http://www1.cs.columbia.edu/CAVE/projects/cc.htm

Course Info

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