COS 429: Computer Vision

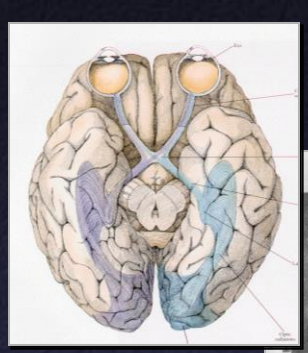
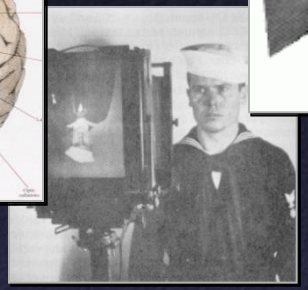

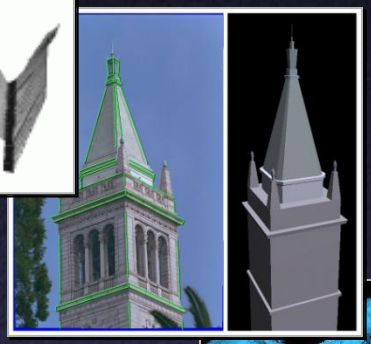
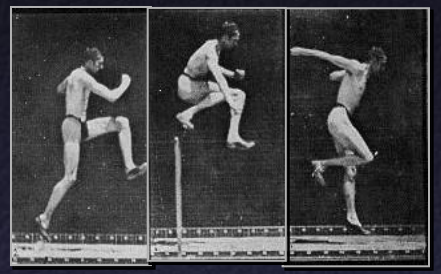
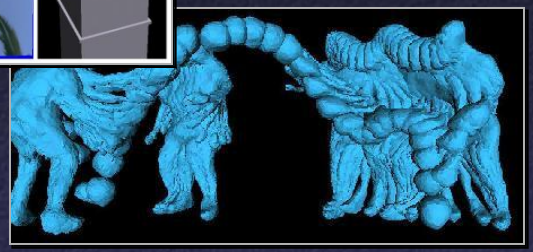
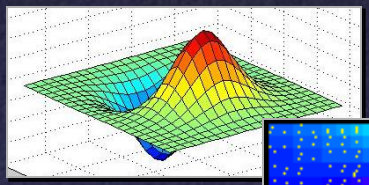
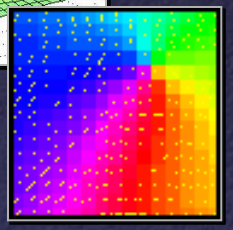
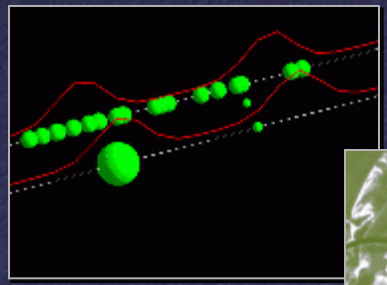


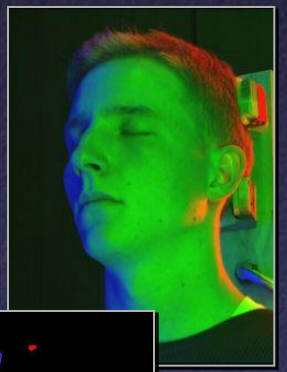
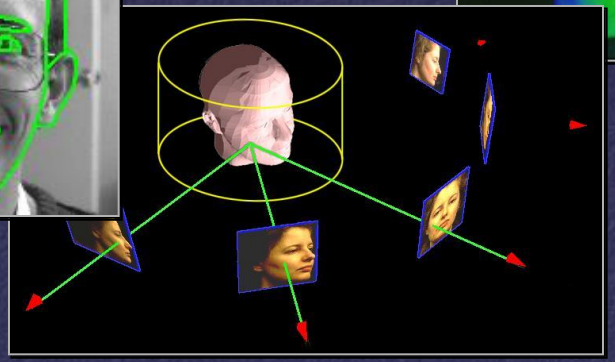

Thanks to Chris Bregler

COS 429: Computer Vision

Instructor: Thomas Funkhouser

funk@cs.princeton.edu

Preceptors: Ohad Fried, Xinyi Fan

{ohad,xinyi}@cs.princeton.edu

Web page:

http://www.cs.princeton.edu/courses/archive/fall13/cos429/

What is Computer Vision?

What is Computer Vision?

- Input: images
- Output: information about the world

What is Computer Vision?

Example:

- What is in this image?
- Who is in this image?
- Where are they?
- What are they doing?

What is Computer Vision?

Other questions:

- What camera settings were used?
- Which pixels go with which objects?
- What is the scene description in 3D?

Why is it Important?

- Applications of computer vision?

Why is it Important?

- Applications of computer vision?
  - In photography: automatic focus, etc.
  - In social networking: tagging people and places
  - In science: observing natural phenomena
  - In medicine: processing xrays, CAT scans, etc.
  - In surveillance: recognizing people and actions
  - In engineering: creating models of the world
  - In human-computer interfaces: recognizing gestures
  - In robotics: localization, mapping, scene understanding, etc.

Why is it Timely?

1) Lots of image data is being collected

Why is it Timely?

1a) Lots of image data is being collected now

Why is it Timely?

1b) Lots more image data will be collected soon

Why is it Timely?

1c) Image data can be collected continuously

Why is it Timely?

Why is it Timely?

2) There is a lot of information in images

Why is it Timely?

3) Computer vision is starting to work … kinda

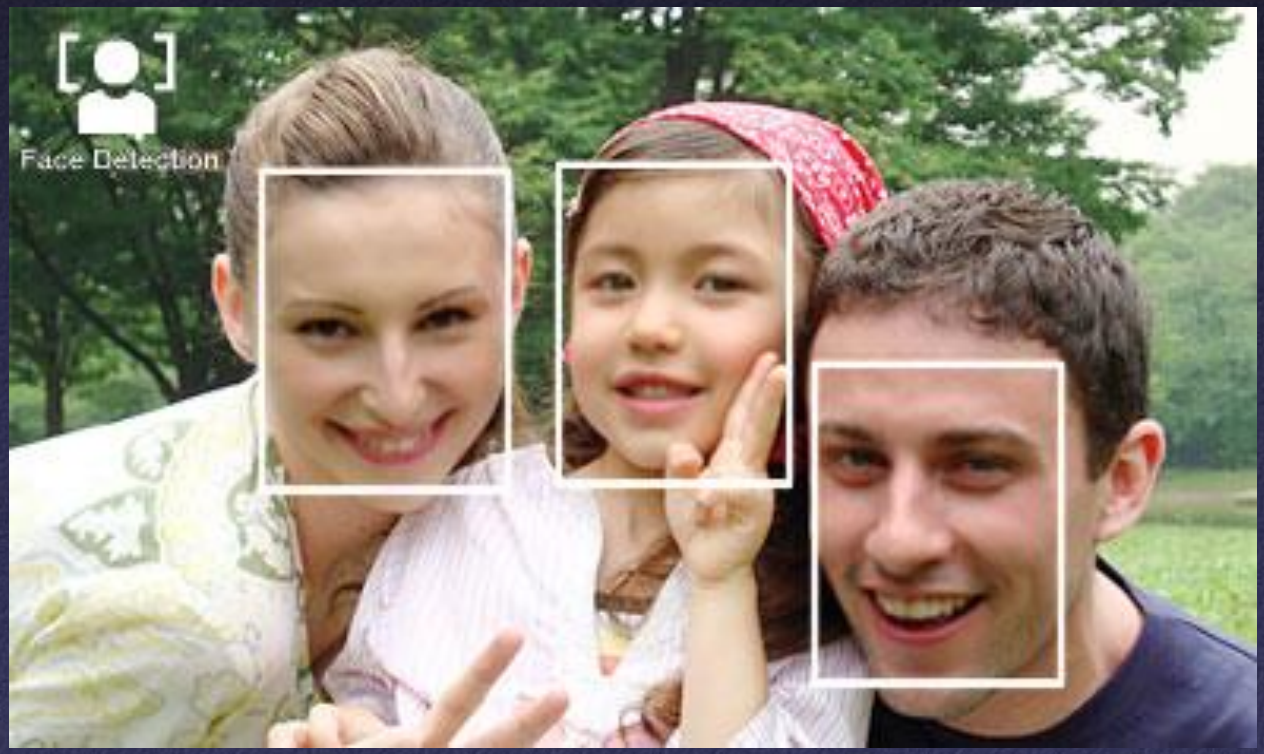

http://www.sony-asia.com/article/271940/section/product/product/dsc-wx1

How Does it Work?

How Does it Work?

- Example: face & eye detection

http://sourceforge.net/apps/mediawiki/pyvision/index.php?title=FaceL:_Facile_Face_Labeling

How Does it Work?

- Example: face & eye detection

How Does it Work?

- Example: face & eye detection

How Does it Work?

- Example: face & eye detection

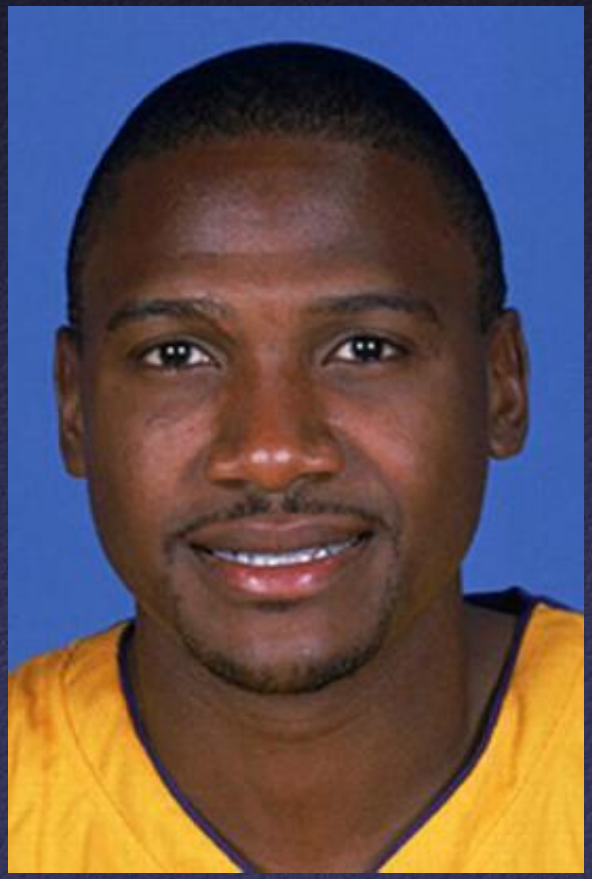

How Does it Work?

- Example: face & eye detection

How Does it Work?

- Example: face & eye detection

How Does it Work?

- Example: face & eye detection

How Does it Work?

- Example: face & eye detection

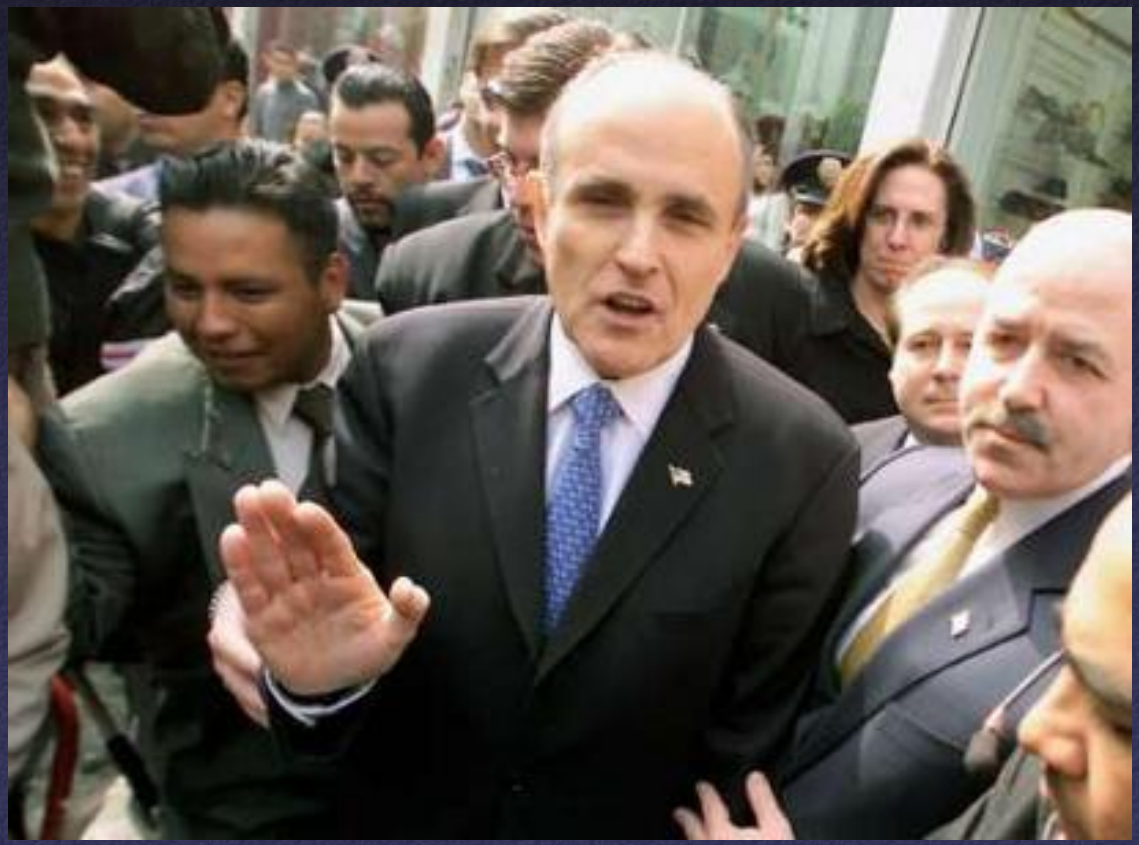

Why is it Hard?

Why is it Hard?

Example: inferring 3D scene from a single image

www.high-resolution-wallpapers.com

Why is it Hard?

1) It's an inverse problem

Image is 2D

Scene is 3D

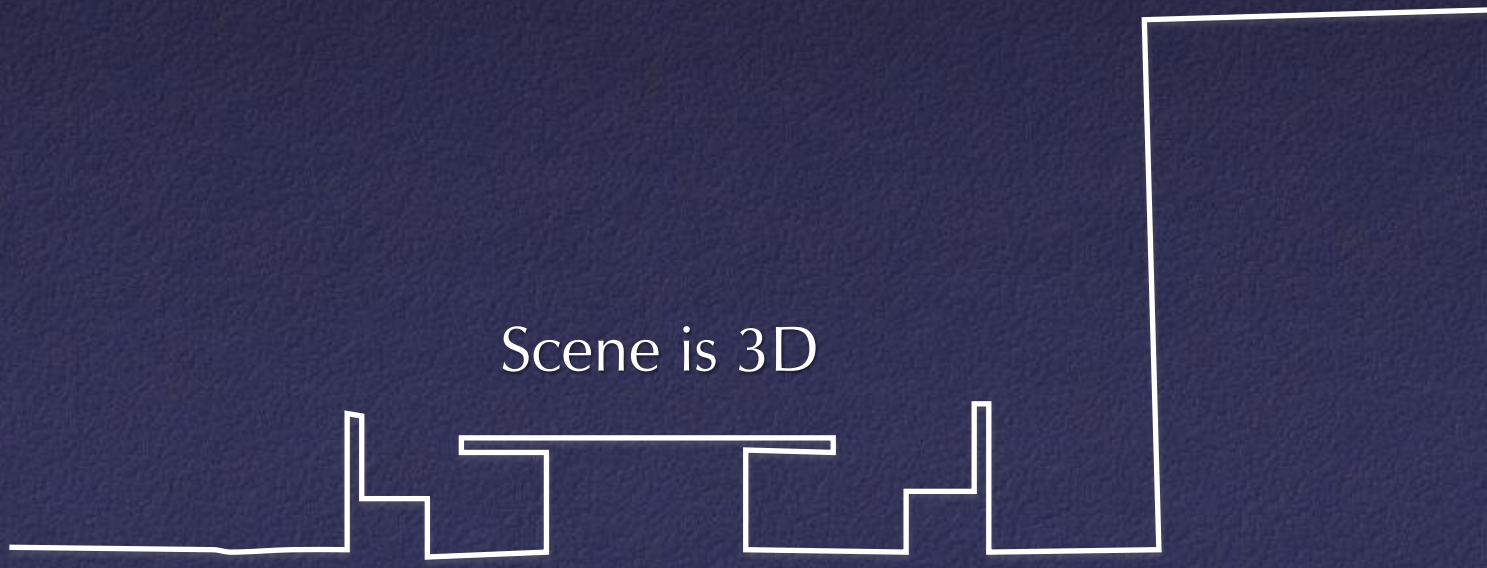

Why is it Hard?

1) It's an inverse problem

Image depends on viewpoint, …

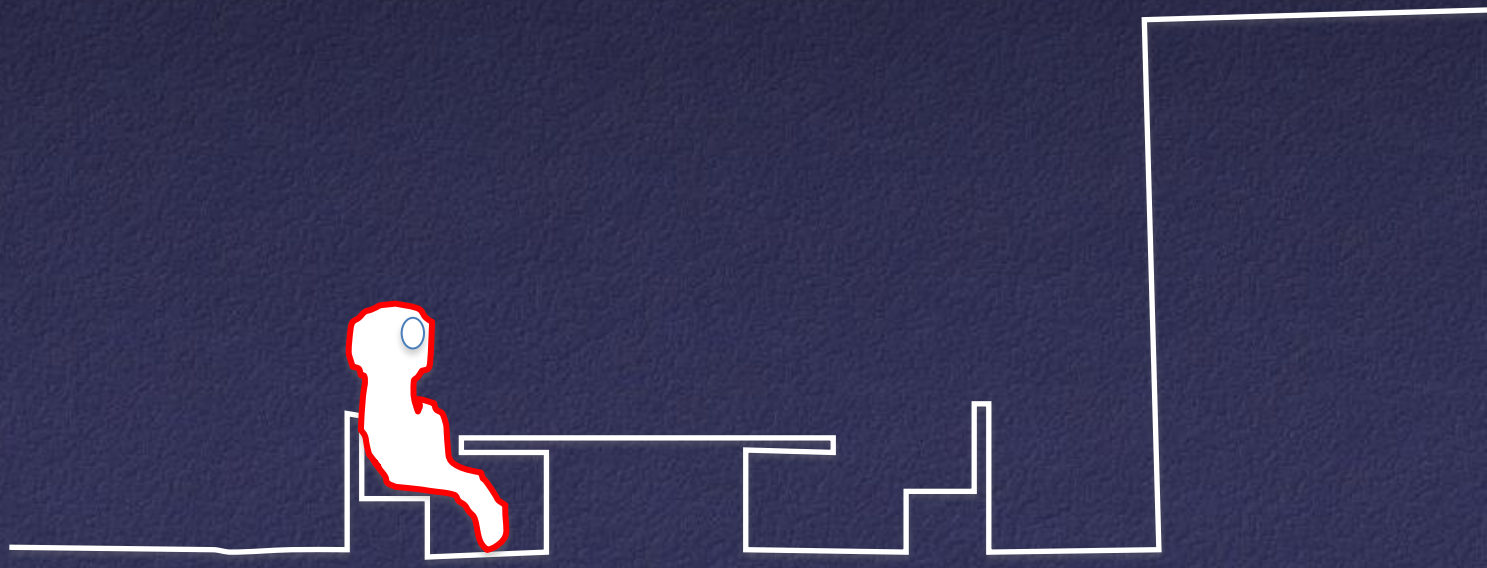

Why is it Hard?

1) It's an inverse problem

… and lighting, …

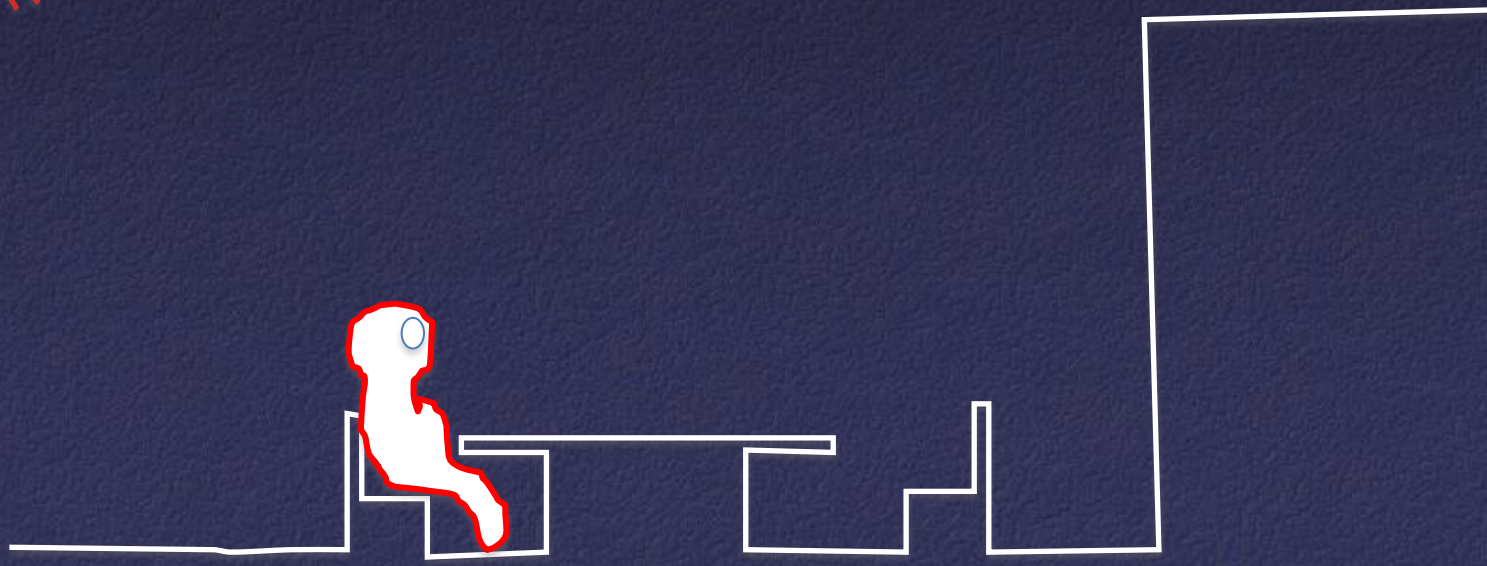

Why is it Hard?

1) It's an inverse problem

… and changes to light
as it travels through scene, …

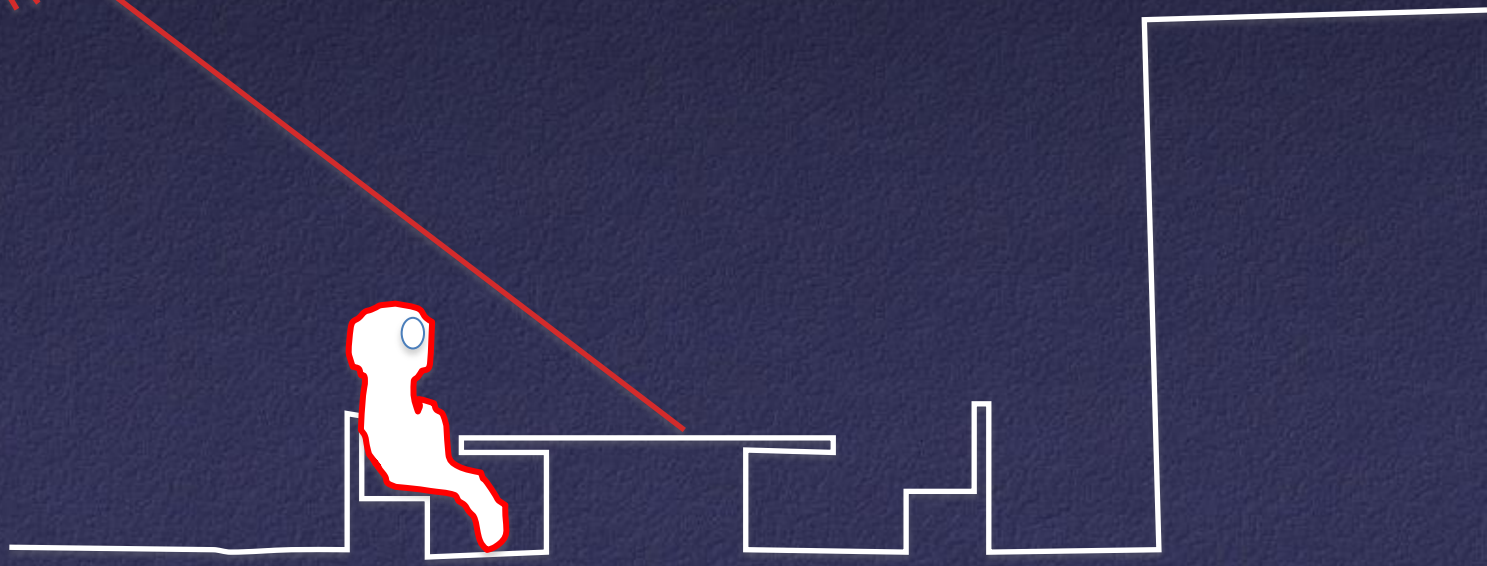

Why is it Hard?

1) It's an inverse problem

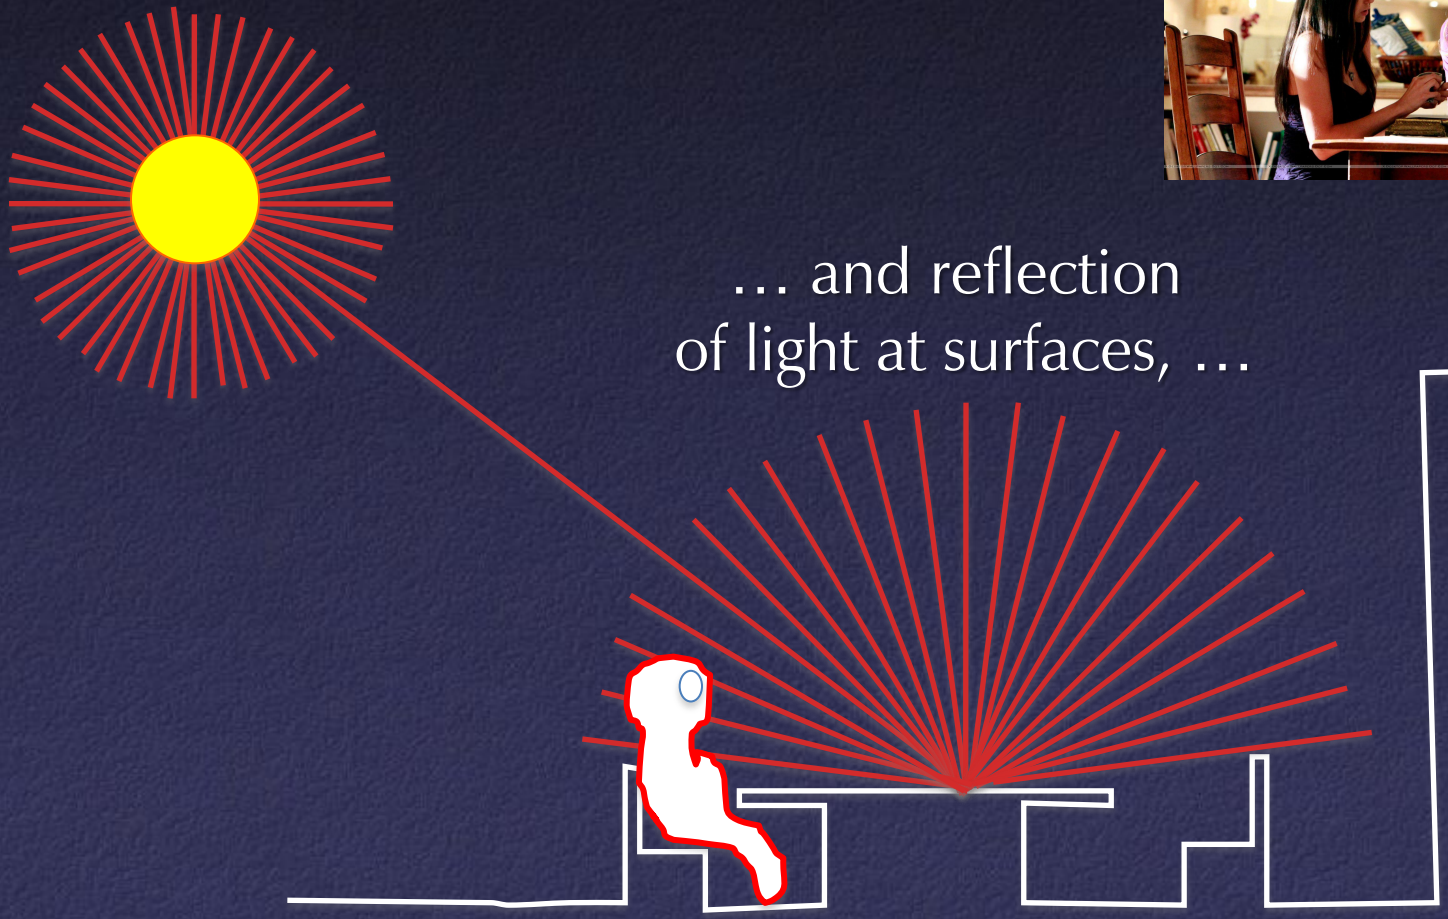

Why is it Hard?

1) It's an inverse problem

Why is it Hard?

1) It's an inverse problem

… and surface visibility, …

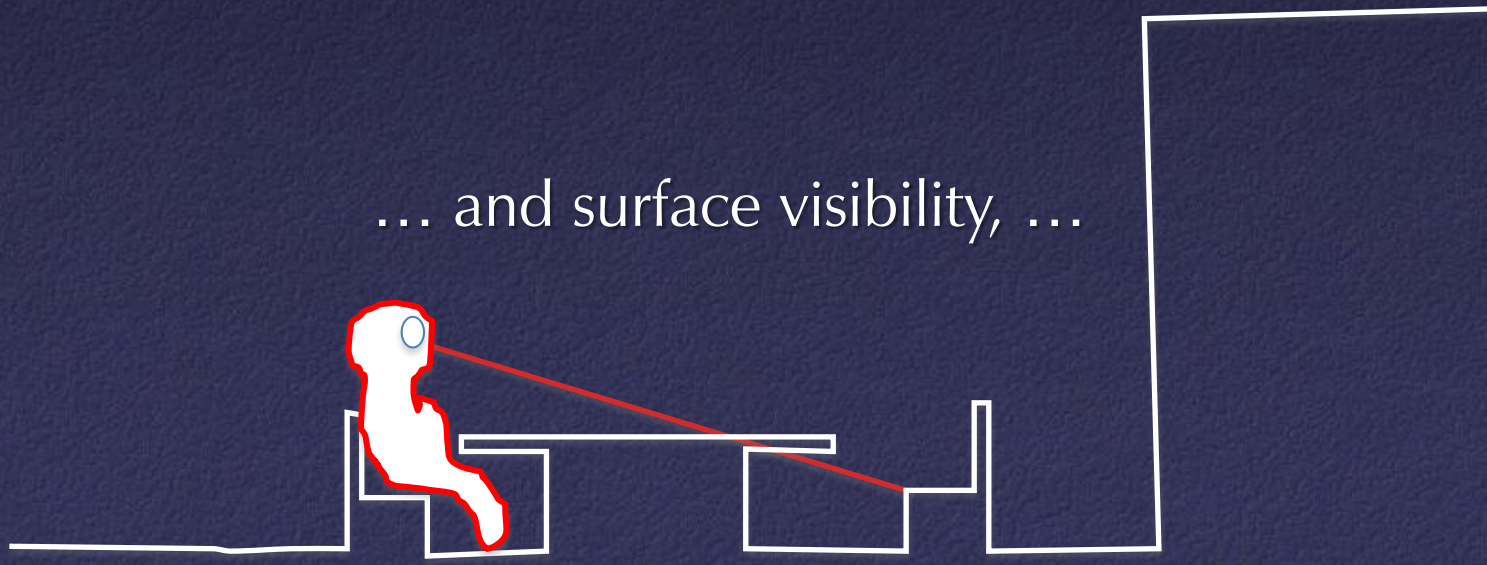

Why is it Hard?

1) It's an inverse problem

and camera properties
(more on that next time)

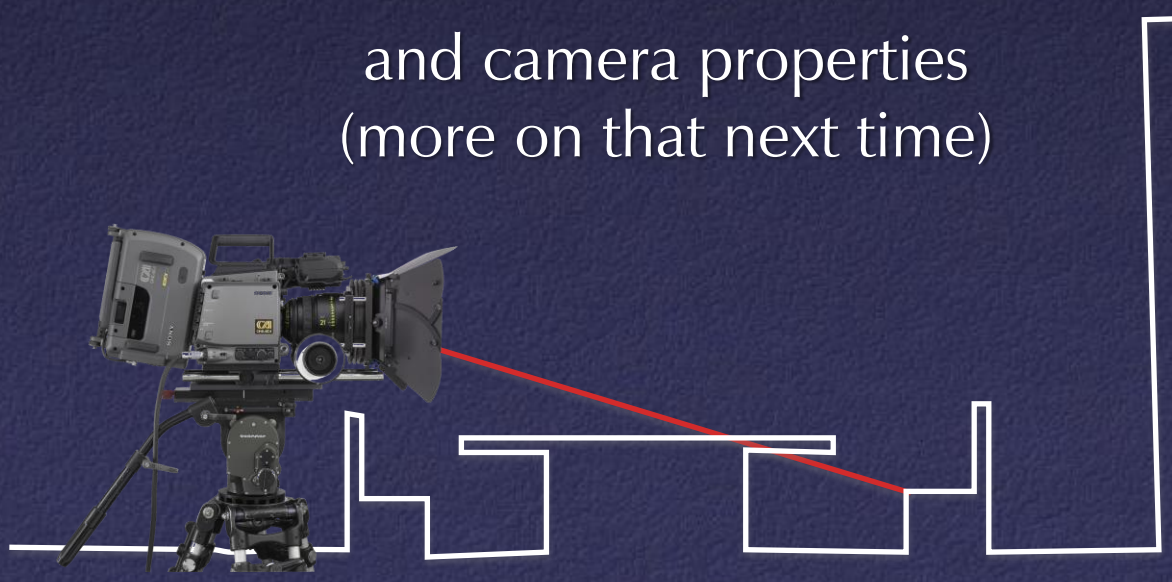

Why is it Hard?

1) It's an inverse problem

Computer graphics = scene → image

- Well-posed forward problem
- Can be solved by simulation

Computer vision = image → scene

- Ill-posed inverse problem
- Must be solved by inference

Why is it Hard?

1) Good inference requires good model of world

www.high-resolution-wallpapers.com

Why is it Hard?

2) Vision requires model of world

Ideally, high-level semantic model derived from real-world knowledge

- Wood tables are usually brown
- Wood tables are often in kitchens
- Cabinets are usually white
- Cabinets are usually in kitchens
- People often sit on chairs at tables
- People talk when they sit at tables
- etc.

Why is it Hard?

2) Vision requires model of world

In practice, low-level statistical model trained on examples

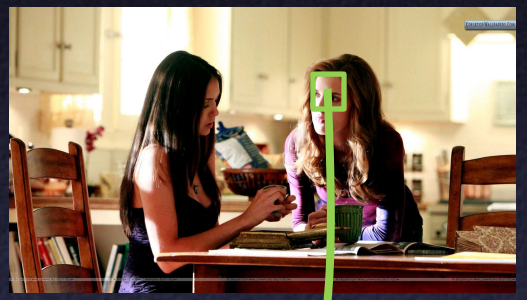

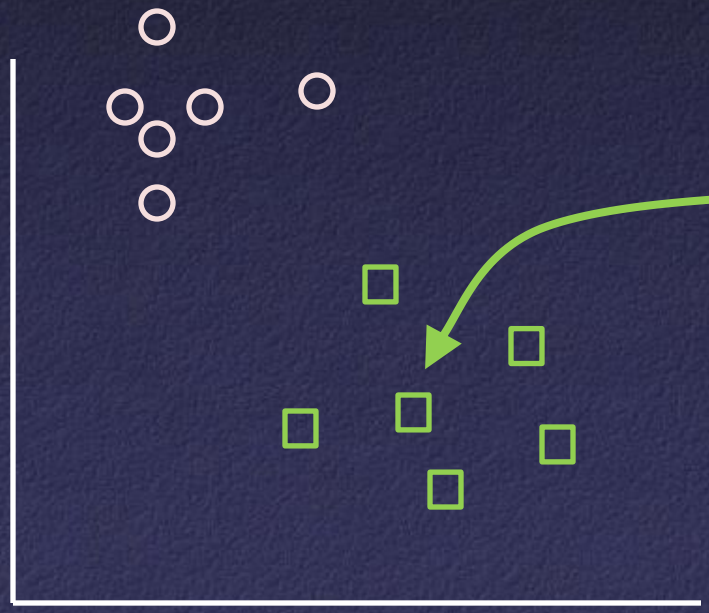

Taxonomy of Problems

- Low-level vision
- Mid-level vision
- High-level vision

Low-Level Vision

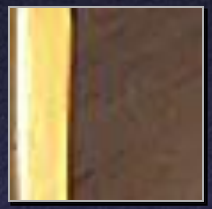

- Considers local properties of an image

"There's an edge!"

Mid-Level Vision

- Grouping and segmentation

"There's an object and a background!"

High-Level Vision

- Recognition
- Classification

"It's a chair! It's in a room!"

Course Outline

- Analyzing single image
  - Image formation and capture
  - Feature detection
  - Texture analysis
  - Segmentation
- Analyzing multiple images
  - Feature correspondences
  - Object tracking & optical flow
  - 3D reconstruction (shape from X)
- Classification and recognition
  - Object classification
  - Object detection
  - Place recognition
  - Scene understanding

Course Mechanics

- Lectures
  - Tues and Thurs 1:30-2:50PM, Friend 006
- Precept-like meetings
  - Not every week, optional
    0) Fri, Sep 13, 3PM, CS 105 ← Tomorrow
    1) Wed, Sep 25, 7:30PM, CS 105
    2) Wed, Oct 9, 7:30PM, CS 105
    3) Wed, Nov 6, 7:30PM, CS 105
    4) Wed, Nov 20, 7:30PM, CS 105

Course Mechanics

- Readings
  - Book: *Computer Vision: Algorithms and Applications* Richard Szeliski (available online)
  - Papers: seminal/relevant pubs (usually optional)
- Coursework
  - Four assignments, plus warmup
  - Final project

Grading

- 65%: Assignments
  - Short written answer
  - Programming exercise
  - Experimental analysis
- 30%: Final project
  - Small groups – 2-3 people
  - Presentation / demo in January
  - Writeup due on Dean's date
- 5%: Class participation

MATLAB

- Programming assignments in MATLAB
- Ohad and Xinyi will run a MATLAB tutorial "precept" tomorrow (Fri) at 3:00PM in CS 105 (the small auditorium).
- School of Engineering is running a short course 6:30-9:00PM on Mon Sep 16 – Tues Sep 17
  - Must register online

Assignment 0

- Warmup to get used to …
  - Thinking about computer vision
  - Programming in MATLAB
  - Analyzing computer vision results

Assignment 0

- Part 1: Thought exercise
  - Write one paragraph

Assignment 0

- Part 2: Programming exercise
  - Get used to programming in MATLAB

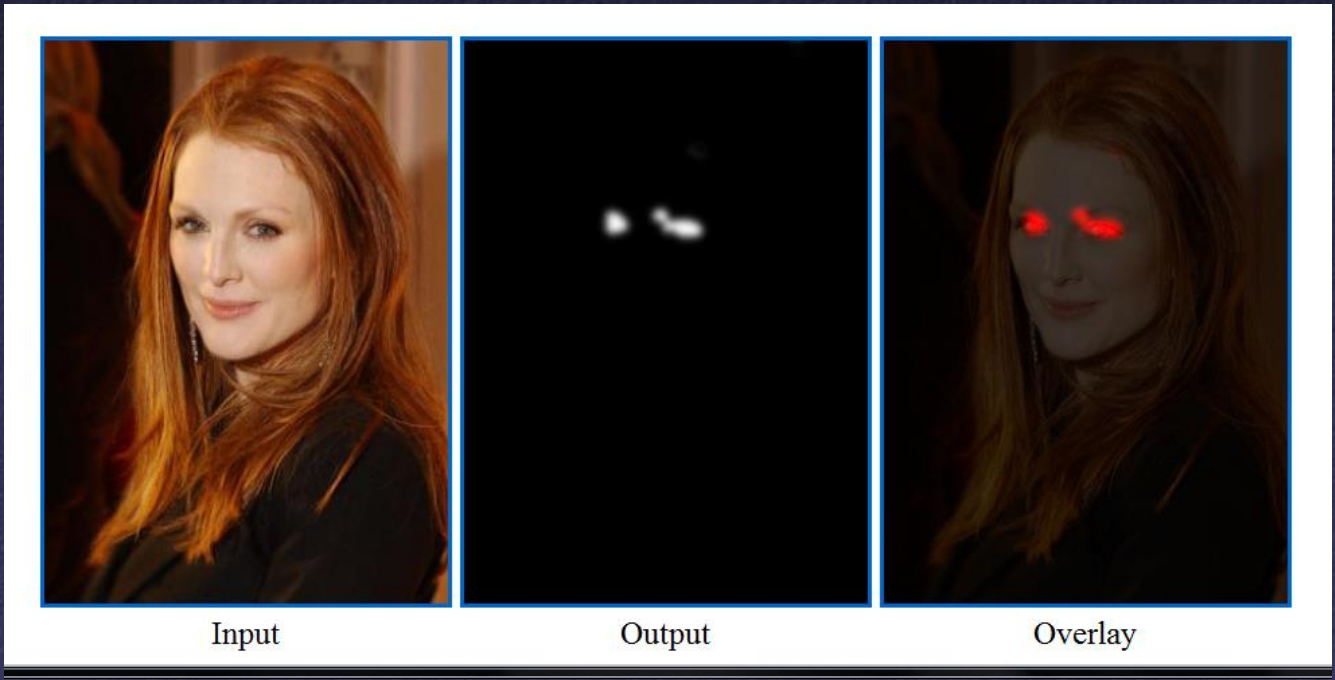

Assignment 0

- Part 3: Results and Analysis
  - How good is the result?

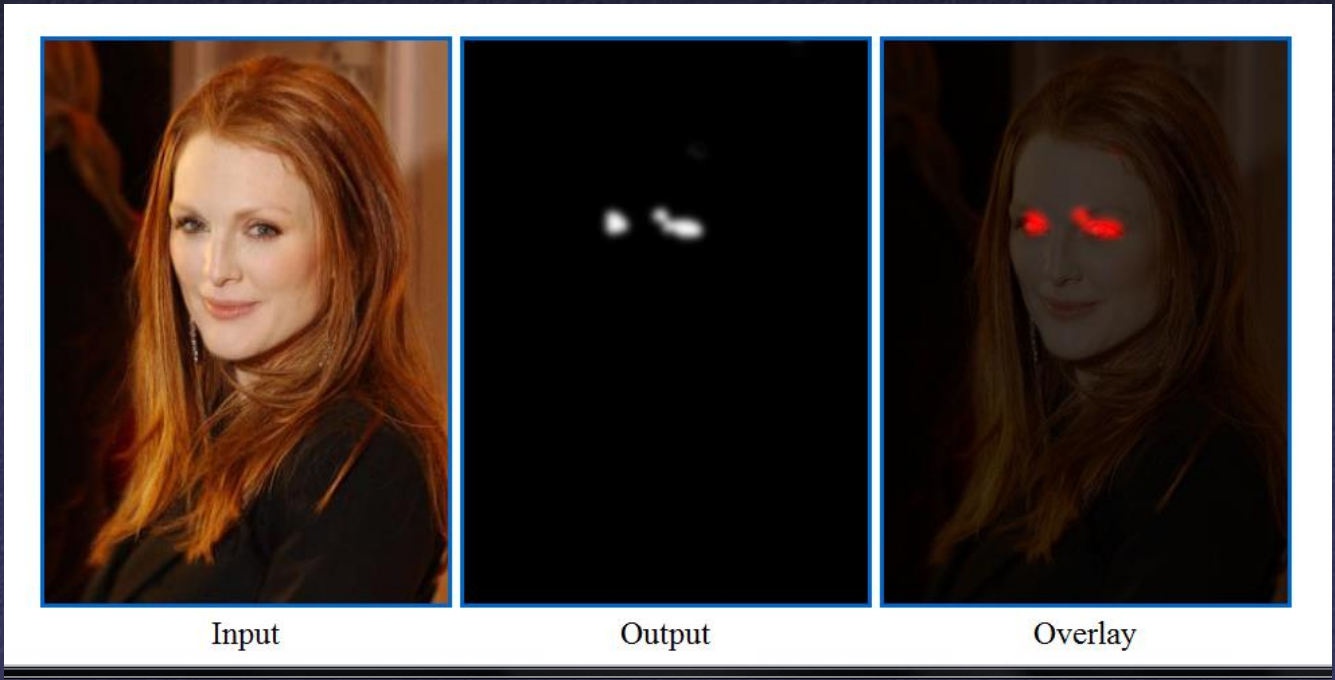

Assignment 0

- Part 3: Results and Analysis
  - How good is the result?
  - Which method gives a better result?
  - How do differences in input affect the results?
  - How do different parameters affect the results?

Collaboration Policy

- Assignments
  - All answers must be your own
  - May only use resources provided for this course
  - Do not Google for similar solutions
  - May ask other students for clarifications on assignment specification, lecture slides, readings, etc.
  - Nobody except prof and preceptors may see your code or help with detailed steps, ever

Q&A

- We will use piazza for Q&A.
- Please post non-private questions
  - E.g., MATLAB help, lecture clarification, assignment specification, bug reports, etc.
  - Please answer each others' questions (we will monitor and endorse students' answers)
- Please send specific questions about your implementation privately to preceptor
- Please visit office hours for longer discussions

The End

See you next time!

http://www.cs.princeton.edu/courses/archive/fall13/cos429/