CSC420: Intro to Image Understanding
Introduction

Sanja Fidler

September 11, 2014

UNIVERSITY OF TORONTO

The Team

- **Instructor**:

Sanja Fidler (fidler@cs.toronto.edu)

- **Office**: 283B in Pratt
- **Office hours**: Tuesday 1.20-2.50pm, or by appointment
- **TAs**:

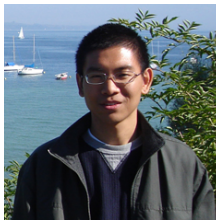

Tom Lee (tshlee@cs.toronto.edu)

Kaustav Kundu (kkundu@cs.toronto.edu)

- **Office hours**: TBA

Course Information

- **Class time**: Tuesday and Thursday at 3-4pm
- **Location**: BA2185
- **Tutorials**: demos and Q&A, we'll do it on demand
- **Class Website**:

  http://www.cs.utoronto.ca/~fidler/CSC420.html
- The class will use Piazza for **announcements** and **discussions**:

  https://piazza.com/utoronto.ca/fall2014/csc420
- Your grade will **not depend on your participation on Piazza**. It's just a good way for asking questions, discussing with your instructor, TAs and your peers

Course Information

- **Textbook**: We won't directly follow any book, but extra reading in this textbook will be useful:

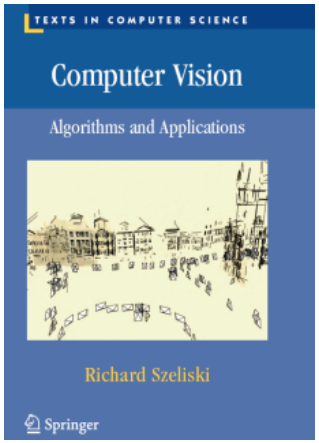

Rick Szeliski

Computer Vision: Algorithms and Applications

available free online:

http://szeliski.org/Book/

- Links to other material (papers, code, etc) will be posted on the class webpage

Course Prerequisites

Course Prerequisites:

- Data structures
- Linear Algebra
- Vector calculus

Without this you'll need some serious catching up to do!

Knowing some basics in this is a plus:

- Matlab (most programming assignments will be in Matlab)
- C++
- Machine Learning
- Solving assignments sooner rather than later

Requirements and Grading

- Each student expected to complete 5 assignments and a project
- **Grading**
  - **Assignments**: 50% (10% each)
  - **Project**: 50%
- **Assignments:**
  - Short **theoretical questions** and **programming exercises**
  - Will be given every **two weeks** (starting with second week of class)
  - You will have **a week to hand in the solution** to each assignment
  - You need to solve the assignment **alone**
- **Project:**
  - You will be able to choose from a list of projects or come up with your own project (discussed prior with your instructor)
  - Need to hand in a **report** and do an oral **presentation**
  - Can work **individually** or in **pairs**

Term Work Dates

Term Work	Post Date	Due Date	% of grade
Assignment 1	Sept 18	Sept 27	10%
Assignment 2	Oct 2	Oct 11	10%
Assignment 3	Oct 16	Oct 25	10%
Assignment 4	Oct 30	Nov 8	10%
Assignment 5	Nov 13	Nov 22	10%
Project Report		Dec 7	30%
Project Presentation		Dec 16	20%

- All dates are for 2014. ;)

Programming Language?

- Your assignments / project can be in Matlab, Python, C++
- As long as it compiles, runs, and you know how to defend it, we're happy
- HOWEVER, most code and examples we will provide during the class will be in Matlab
- Most code provided online by computer vision researchers is in Matlab
- Choose wisely

Lateness

Deadline The solutions to the assignments / project should be submitted **by 11.59pm on the date they are due**. Anything from 1 minute late to 24 hours will count as **one late day**.

Lateness Each student will be given a total of **3 free late days**. This means that you can hand in three of the assignments one day late, or one assignment three days late. It is up to the you to make a good planning of your work. **After you have used the 3 day budget, the late assignments will not be accepted.**

Syllabus

Tentative syllabus

Week nb.	Date	Topic
1	Sept 11	Intro
2	Sept 16 & Sept 18	Linear filters, edges
3	Sept 24 & Sept 25	Image features
4	Sept 30 & Oct 2	Keypoint detection
5	Oct 7 & Oct 9	Matching
6	Oct 14 & Oct 16	Segmentation
7	Oct 21 & Oct 23	Grouping
8	Oct 28 & Oct 30	Object, face recognition
9	Nov 4 & Nov 6	Object detection
10	Nov 11 & Nov 13	Stereo, multi-view
11	? & Nov 20	Recognition in 3D
12	Nov 25 & Nov 27	Motion, video

Let's begin!

Introduction to Intro to Image Understanding

- What is Computer Vision?
- Why study Computer Vision?
- Which cool applications can we do with it?
- Is vision a hard problem?
- What's an image?

What is Computer Vision?

What is Computer Vision?

- A field trying to develop automatic algorithms that would **"see"**

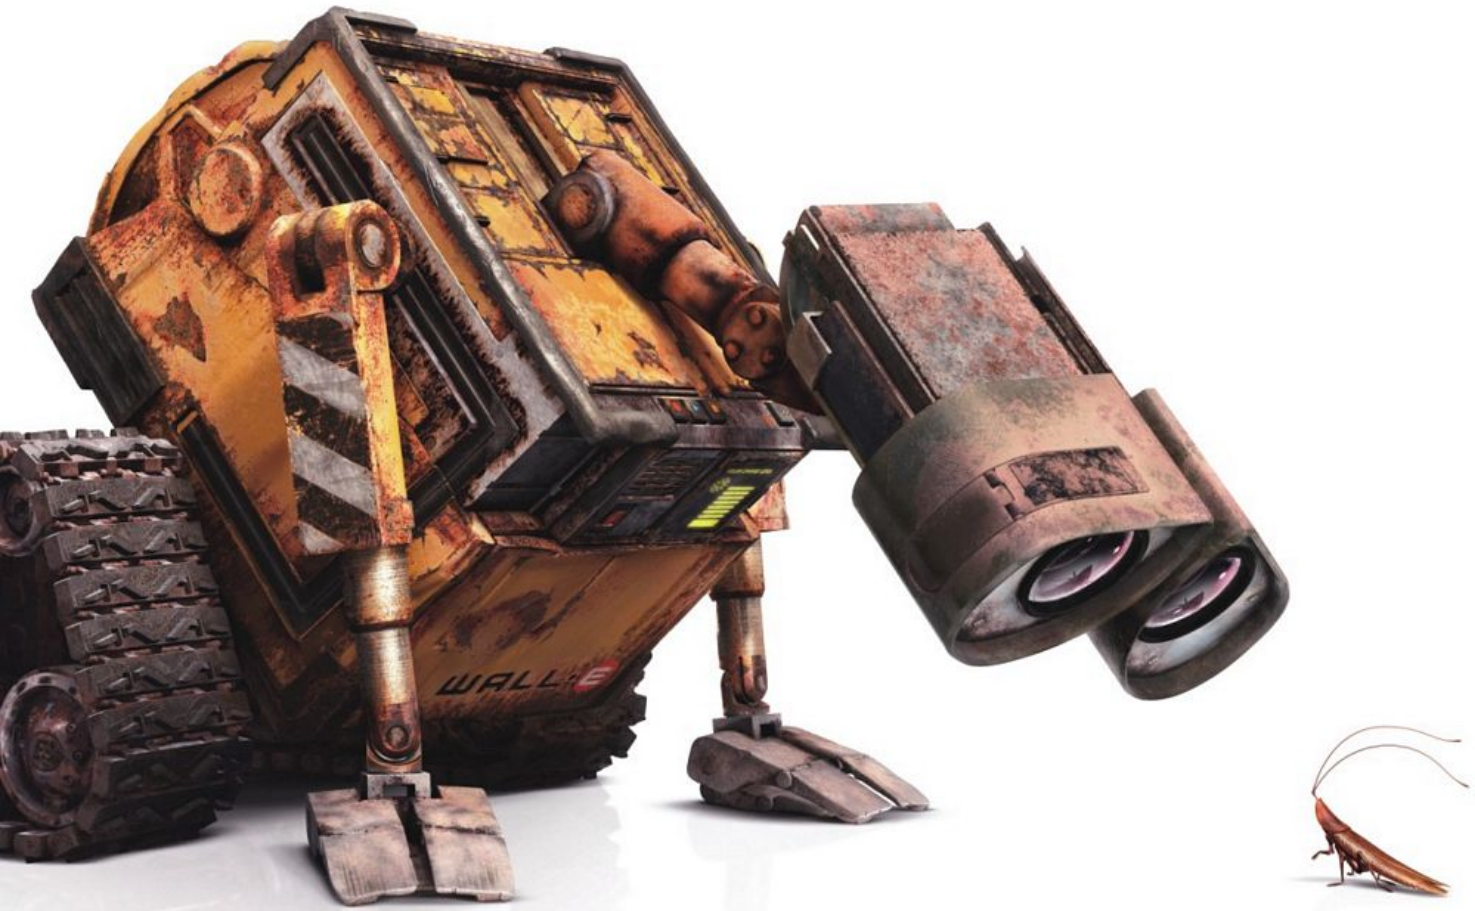

Disney · PIXAR
WALL·E

What is Computer Vision?

[text adopted from A. Torralba]

- What does it mean to see?
  - To know what is where by looking – Marr, 1982

What is Computer Vision?

[text adopted from A. Torralba]

- What does it mean to see?
  - To know what is where by looking – Marr, 1982
  - Understand where things are in the world

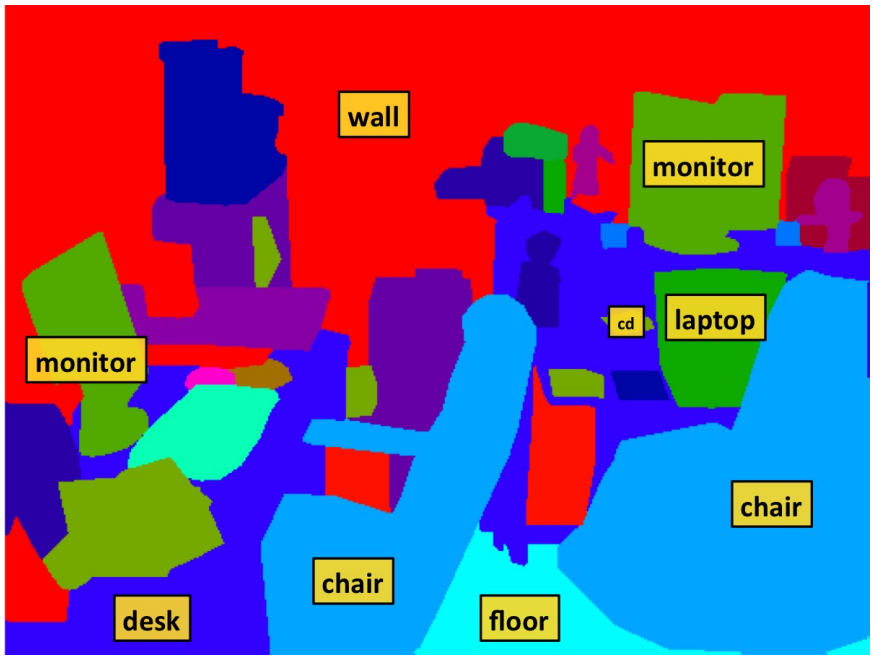

What is Computer Vision?

[text adopted from A. Torralba]

- What does it mean to see?
  - To know what is where by looking – Marr, 1982
  - Understand where things are in the world
  - What are their 3D properties?

image

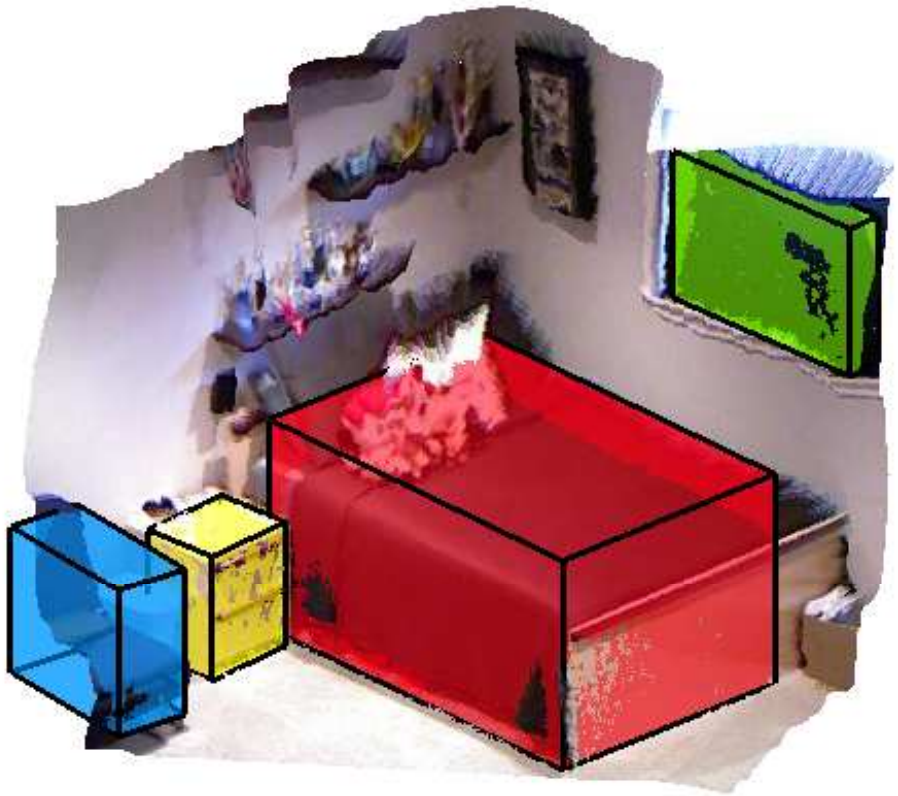

What is Computer Vision?

- What does it mean to see? [text adopted from A. Torralba]
  - To know what is where by looking – Marr, 1982
  - Understand where things are in the world
  - What are their 3D properties?
  - Wh

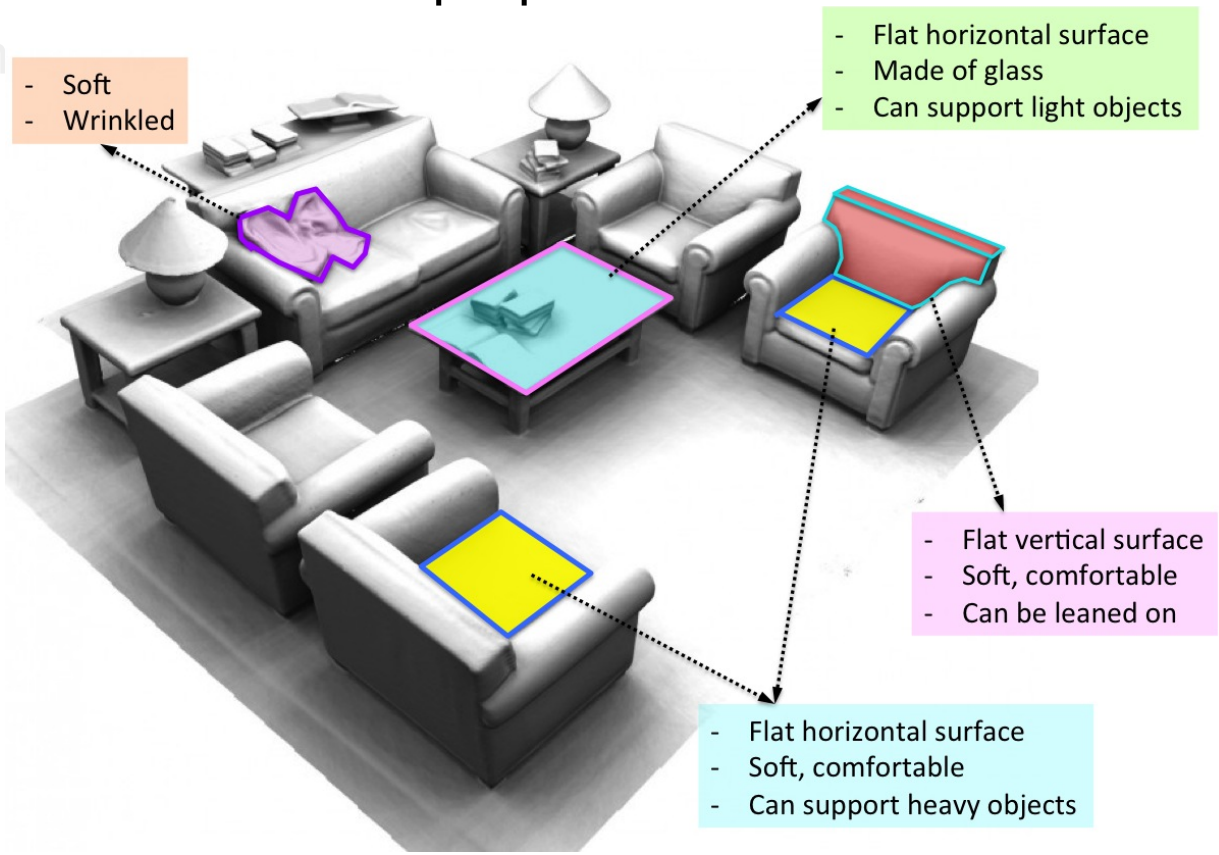

Depth pic from `http://vladlen.info`

What is Computer Vision?

[text adopted from A. Torralba]

- What does it mean to see?
  - To know what is where by looking – Marr, 1982
  - Understand where things are in the world
  - What are their 3D properties?
  - What actions are taking place?

Pic from www.cobblehillpuzzles.com

Why study Computer Vision?

Why study Computer Vision?

- Because it is challenging and fun

Jialiang Wang's (4th undergraduate year, UofT) video about his summer research in computer vision (click on the pic to see video – you'll need internet connection)

Why study Computer Vision?

- You are curious how to one day make the robot walk your dog

(click on the pic to see video)

Why study Computer Vision?

- ... and fold your laundry

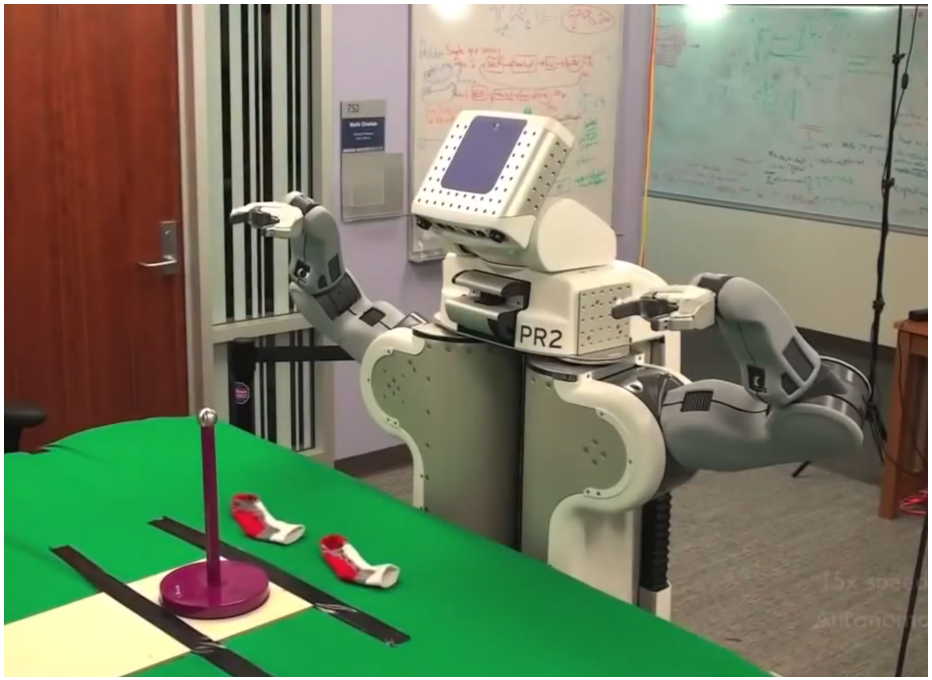

(click on each pic to see videos)

Why study Computer Vision?

- ... and drive you to work (video)

Amnon Shashua's Mobileye autonomous driving system

Why study Computer Vision?

- Allows you to manipulate your images

Scene Completion using Millions of Photographs, Hays & Efros, SIGGRAPH 2007

Why study Computer Vision?

- Allows you to manipulate your images

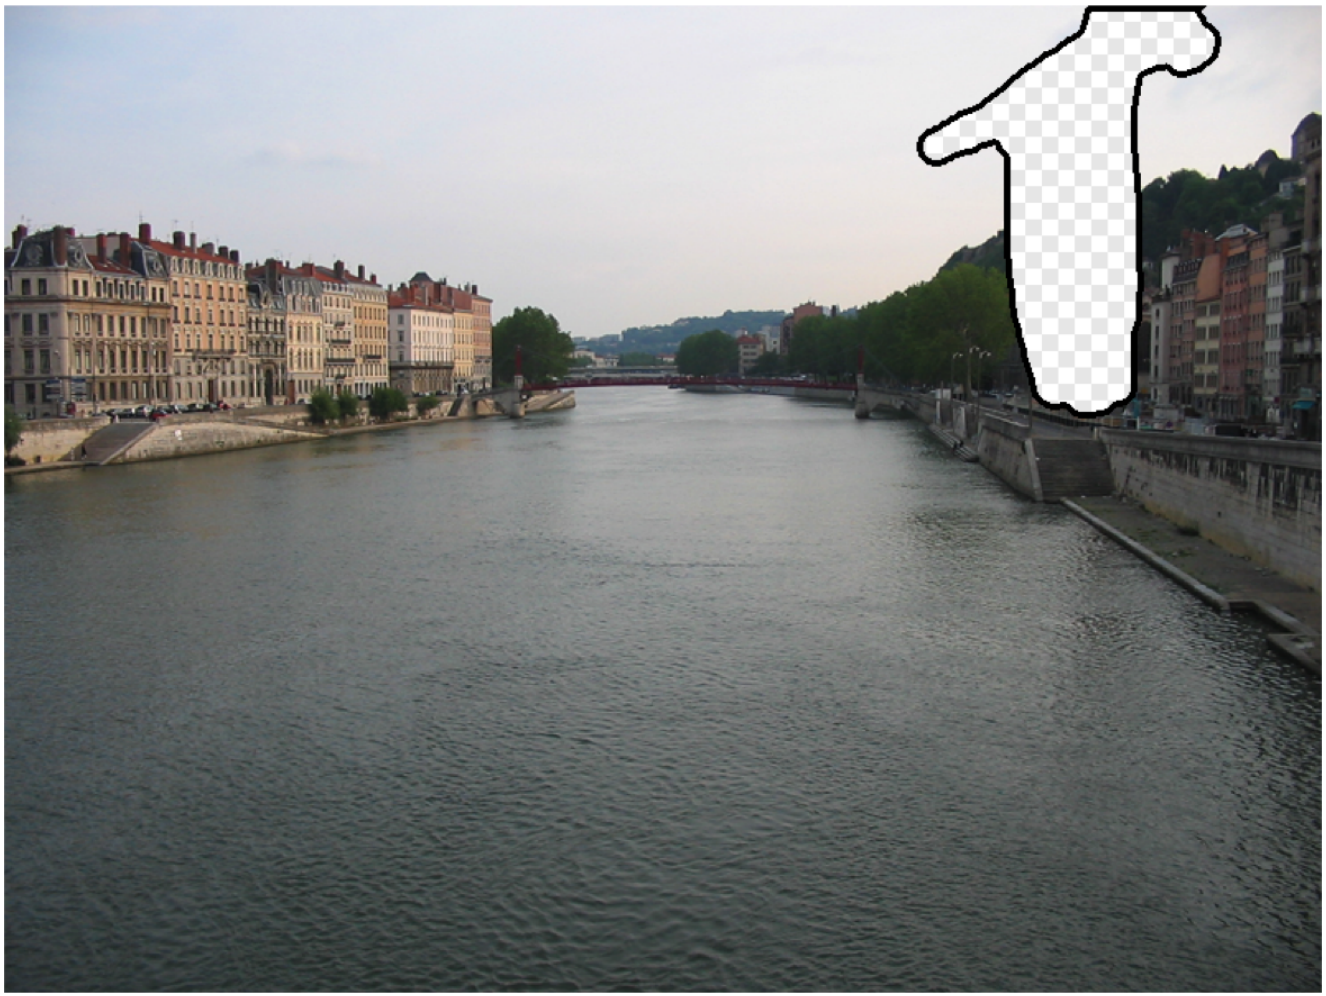

Scene Completion using Millions of Photographs, Hays & Efros, SIGGRAPH 2007

Why study Computer Vision?

- Allows you to manipulate your images

Scene Completion using Millions of Photographs, Hays & Efros, SIGGRAPH 2007

Why study Computer Vision?

- Allows you to manipulate your images

Scene Completion using Millions of Photographs, Hays & Efros, SIGGRAPH 2007

Why study Computer Vision?

- Allows you to manipulate your images

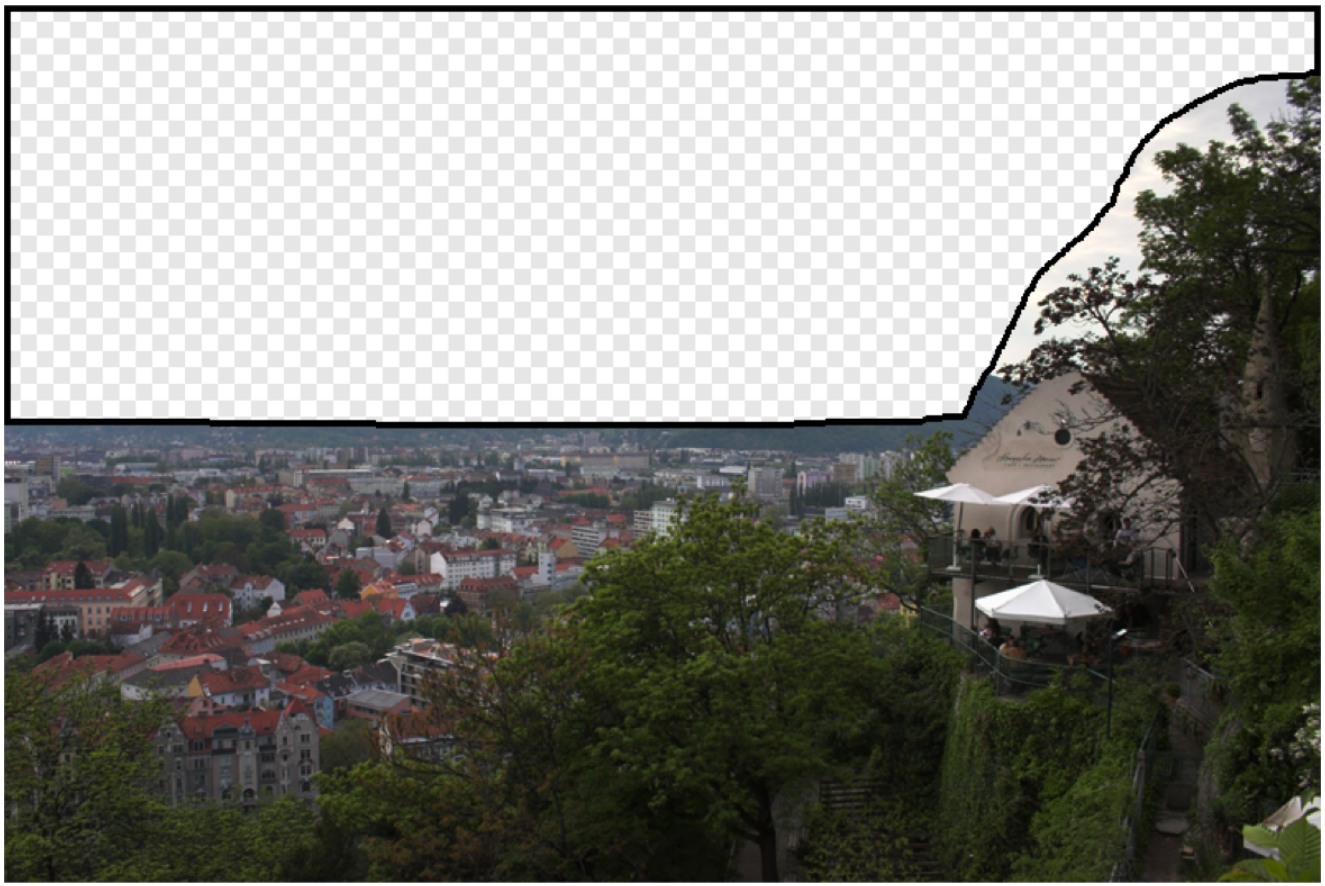

Scene Completion using Millions of Photographs, Hays & Efros, SIGGRAPH 2007

Why study Computer Vision?

- Allows you to manipulate your images

Scene Completion using Millions of Photographs, Hays & Efros, SIGGRAPH 2007

Why study Computer Vision?

- ... and make cool videos using a single image

3D Object Manipulation in a Single Photograph using Stock 3D Models, Kholgade, Simon, Efros, Sheikh, SIGGRAPH 2014

Why study Computer Vision?

- Fancy visualization and game analysis in sports

Why study Computer Vision?

- Fancy visualization and special effects in movies

[Source: `http://cvfxbook.com` and `http://vimeo.com/100095868`]

Why study Computer Vision?

- Reconstruct the world in 3D from online photos!
  (click on each pic to see videos)

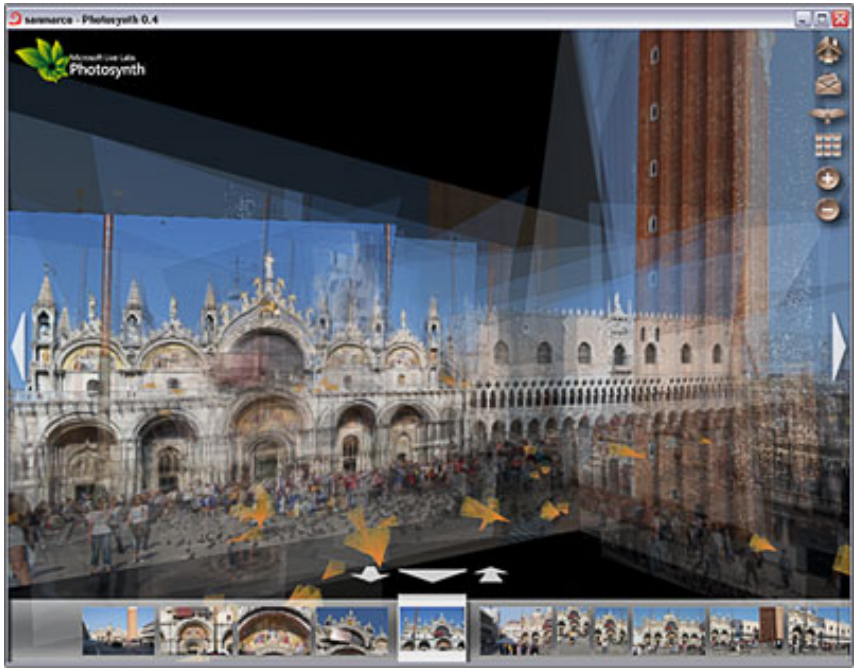

Photosynth, `https://photosynth.net/` (try it!)

Why study Computer Vision?

- Figure out what people are wearing

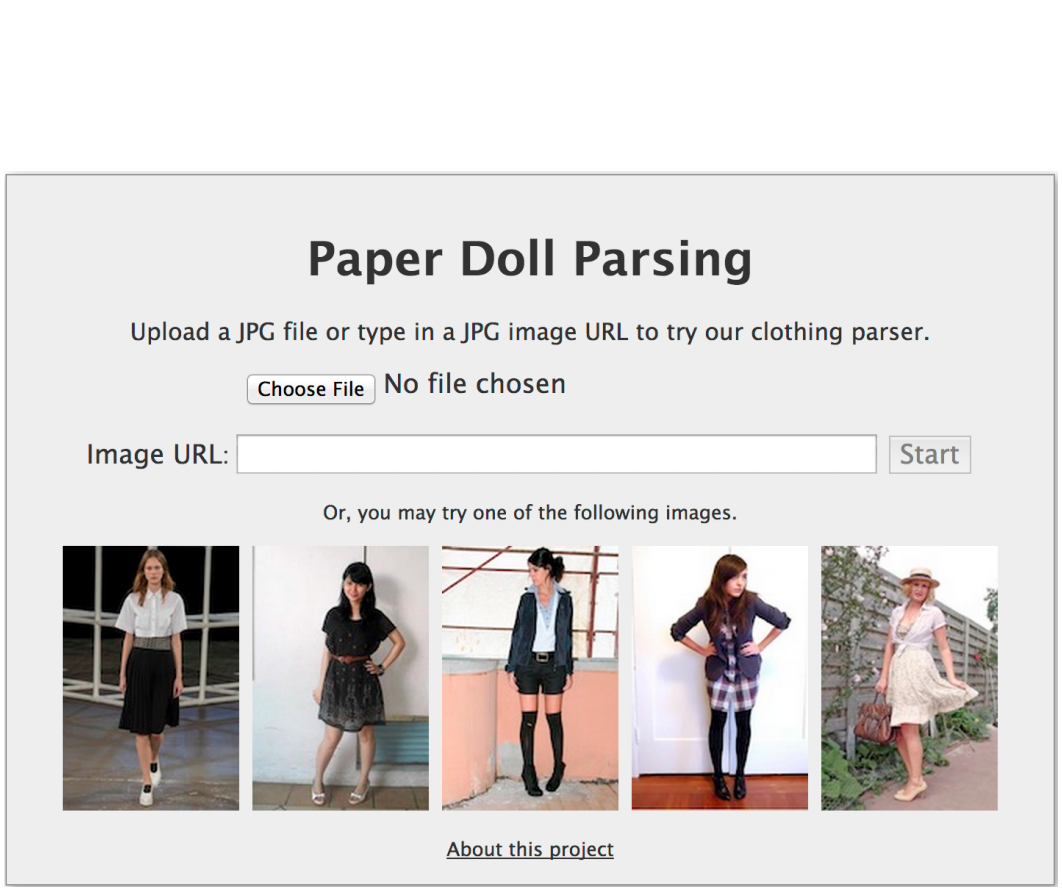

http://clothingparsing.com (try it!)

Why study Computer Vision?

- Detect and analyze faces

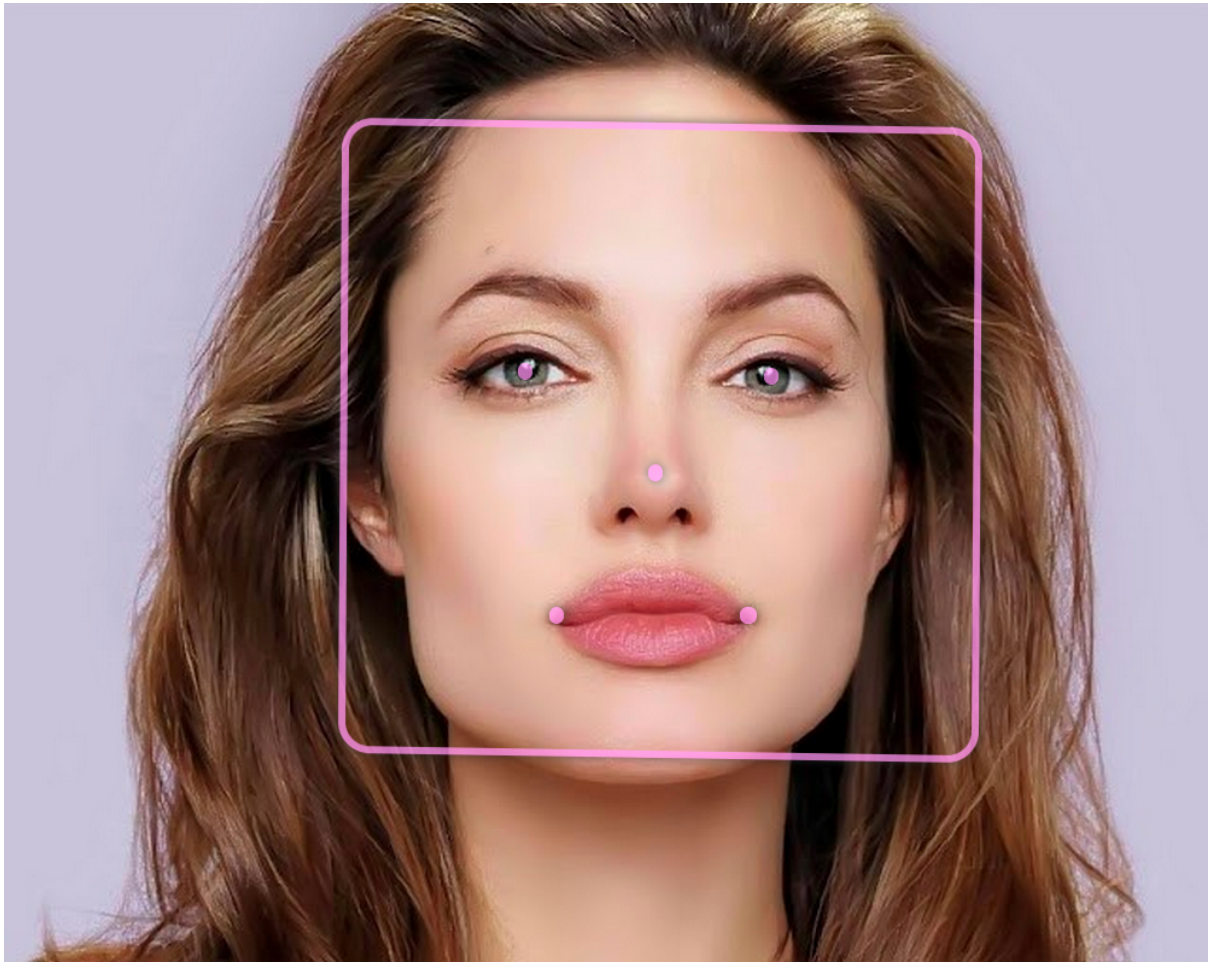

http://www.rekognition.com (try it!)

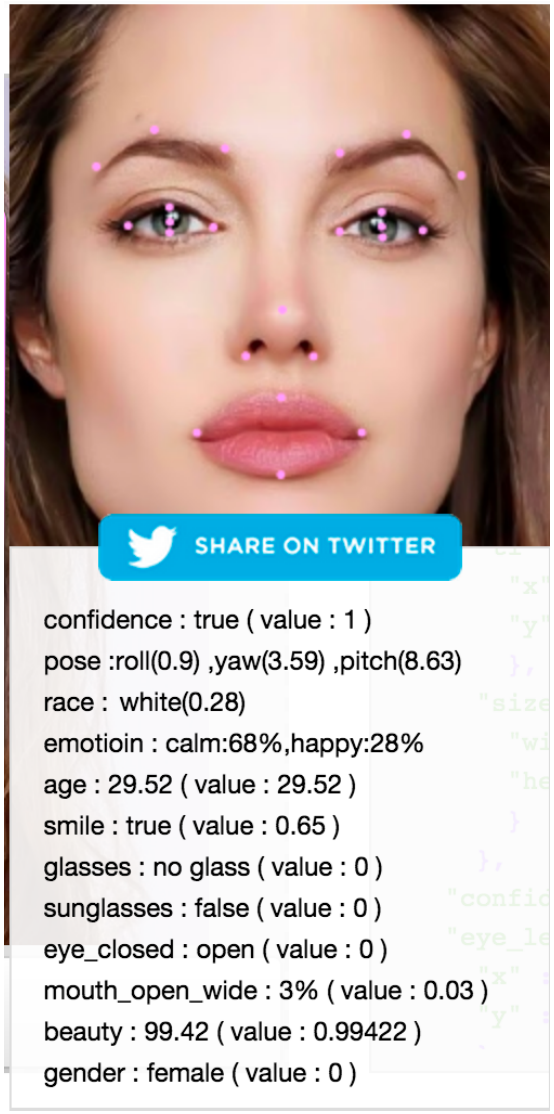

Why study Computer Vision?

- Detect and analyze faces

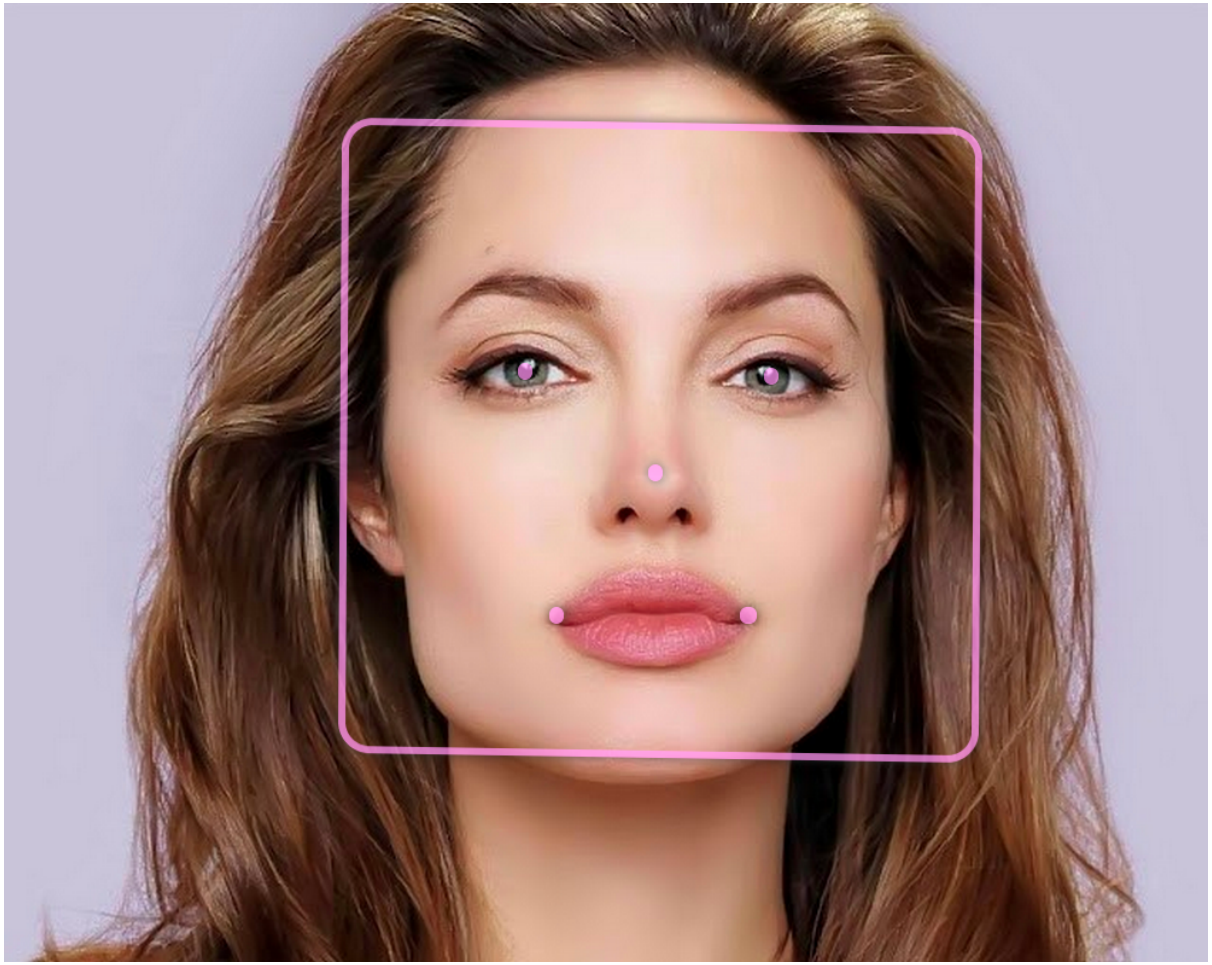

http://www.rekognition.com (try it!)

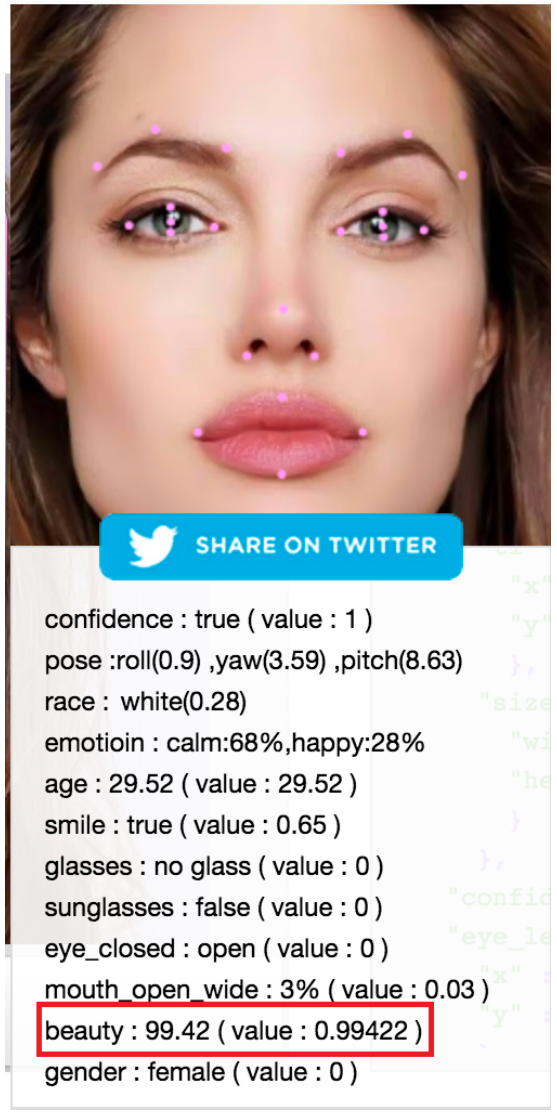

Why study Computer Vision?

- Detect and analyze faces

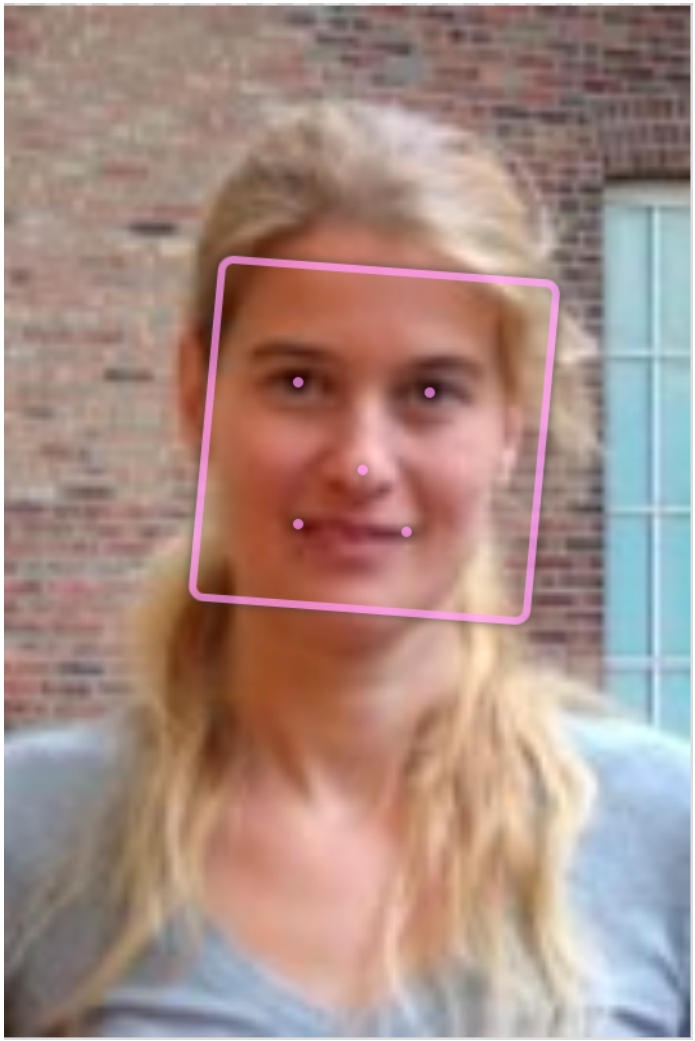

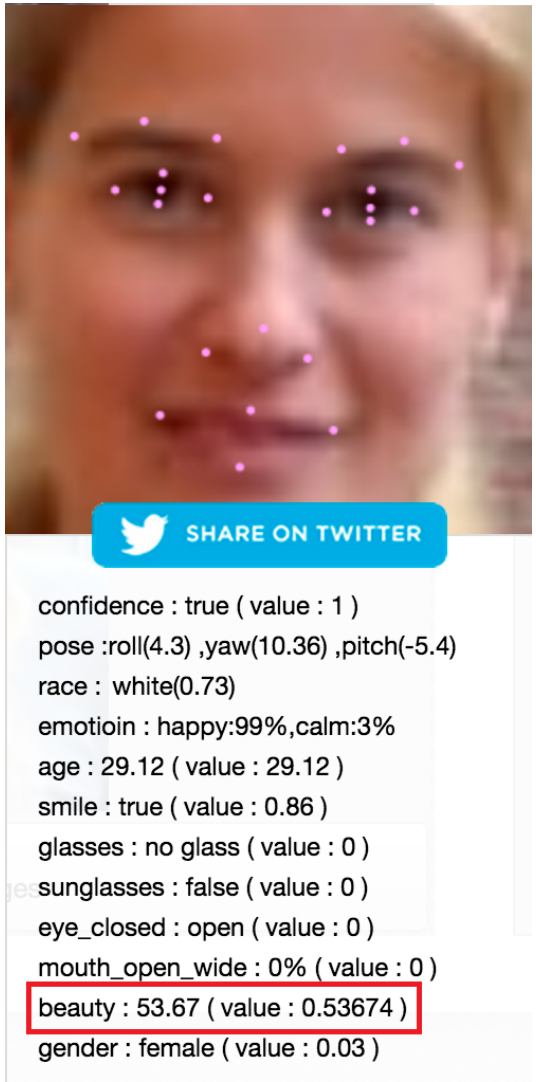

`http://www.rekognition.com` (try it!)

Why study Computer Vision?

- Detect and analyze faces

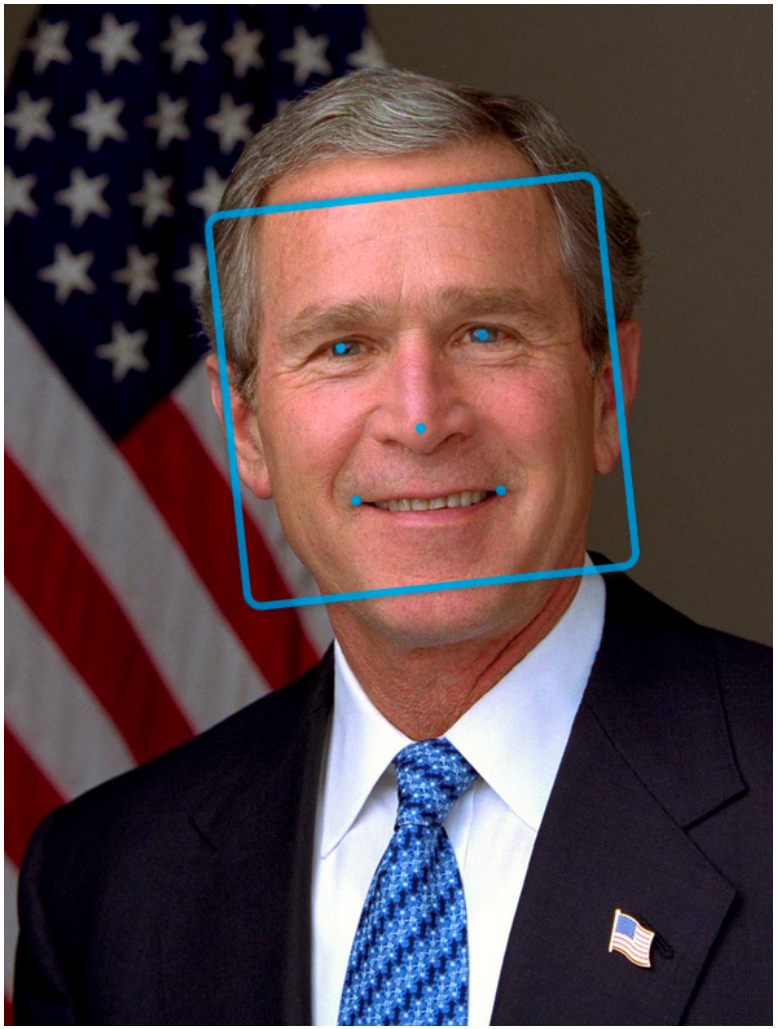

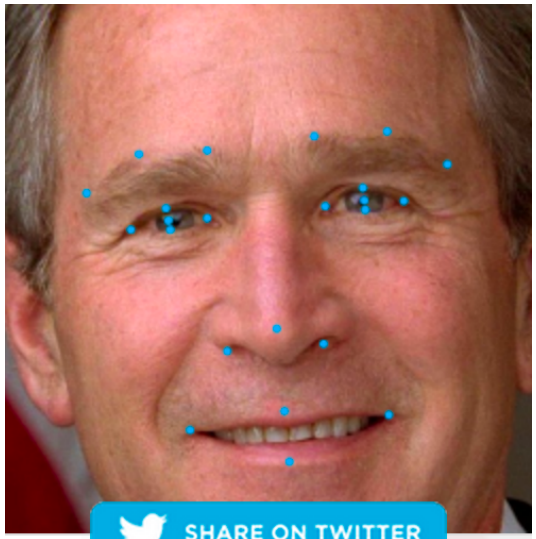

`http://www.rekognition.com` (try it!)

Why study Computer Vision?

- You can make yourself look better (and competitors worse)

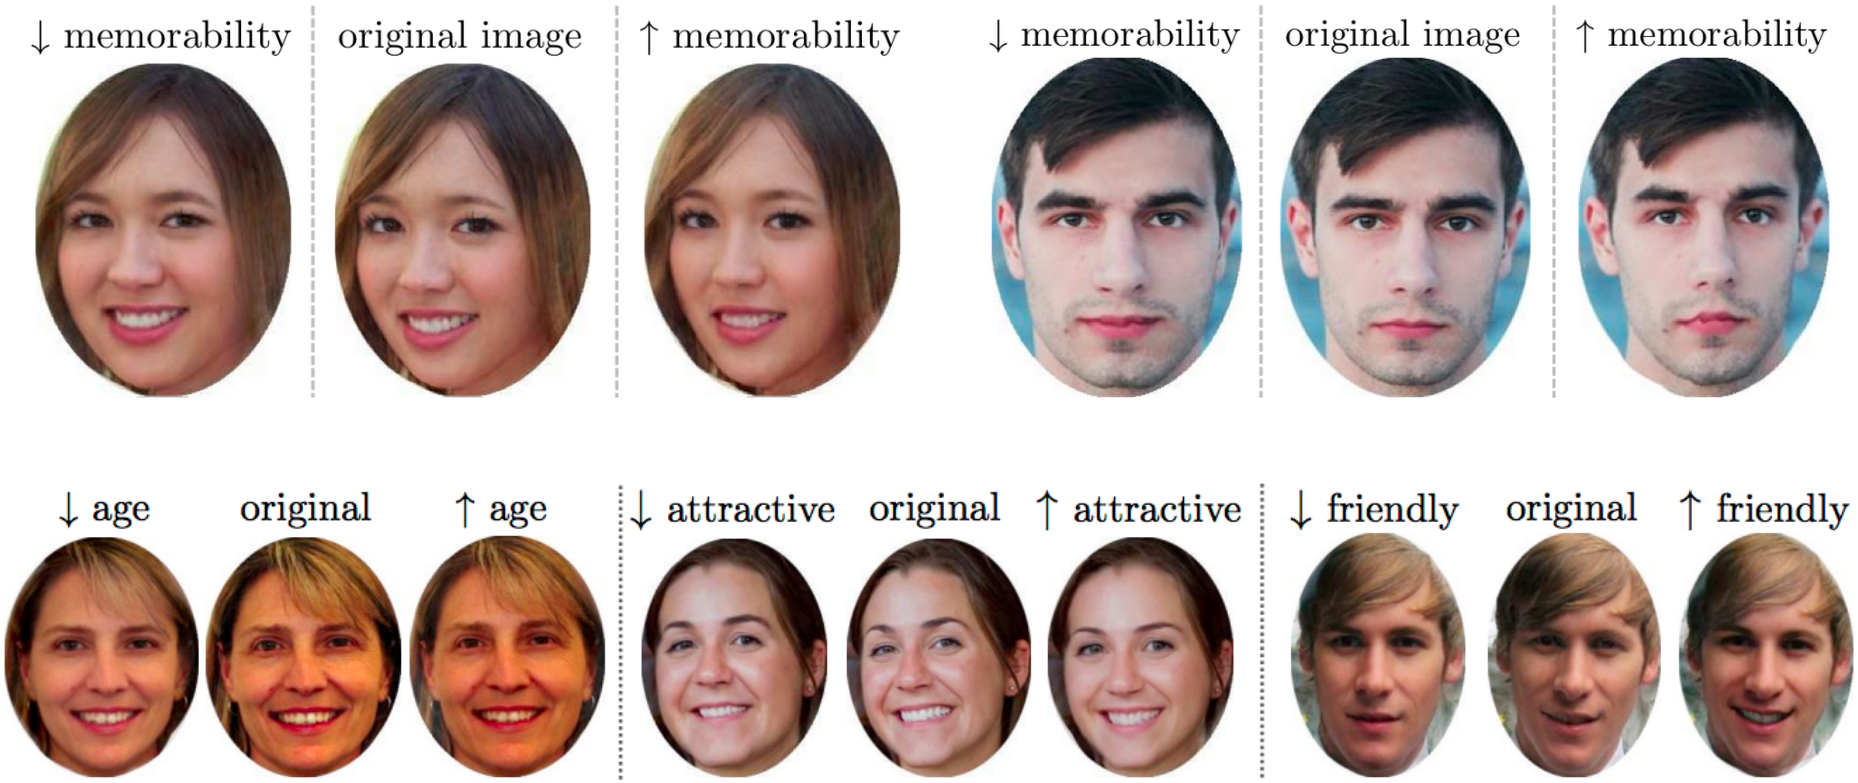

[Khosla, Bainbridge, Oliva, Torralba, Modifying the Memorability of Face Photographs, ICCV 2013]

Why study Computer Vision?

- Fingerprint recognition

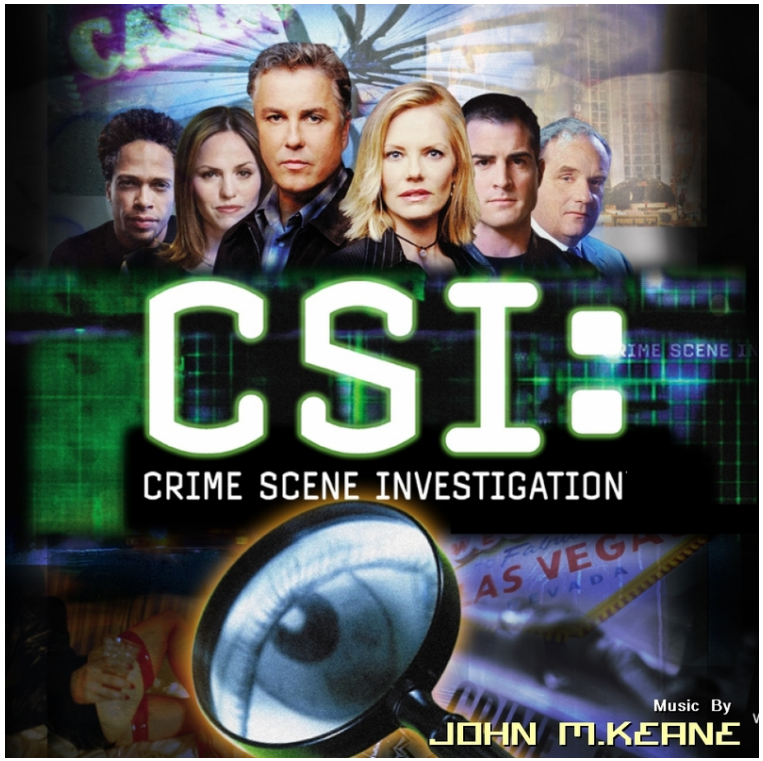

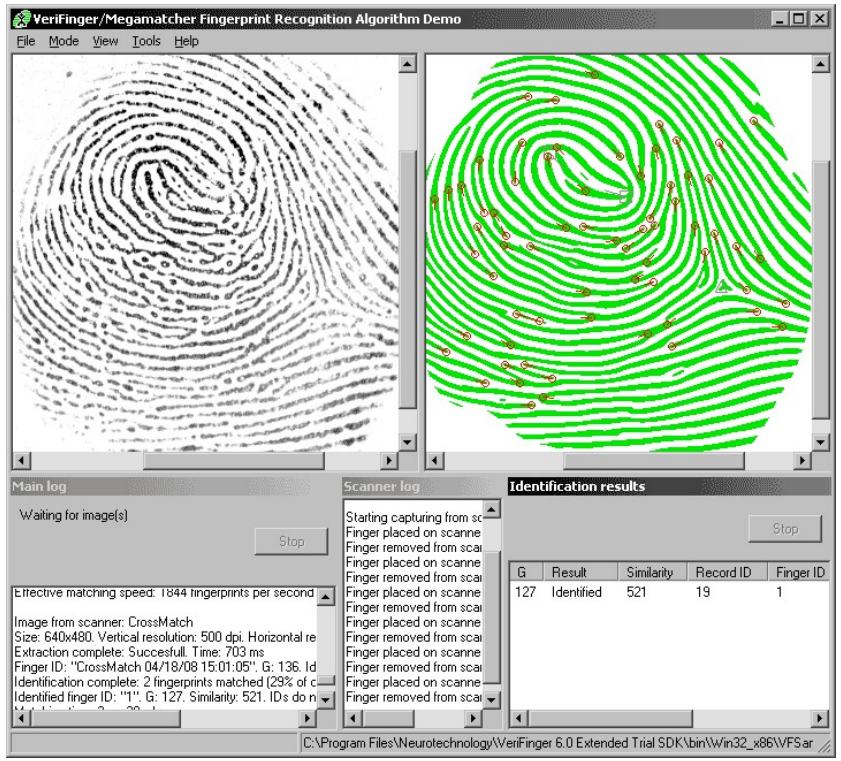

[Source: S. Lazebnik]

Why study Computer Vision?

- You can do some movie-like Forensics

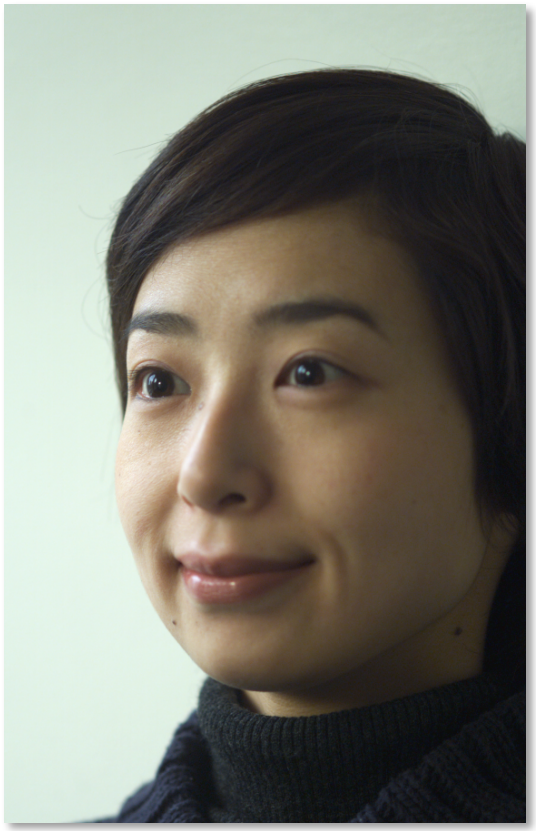

Figure: Source: Nayar and Nishino, Eyes for Relighting

[Source: N. Snavely]

Why study Computer Vision?

[Source: N. Snavely]

Why study Computer Vision?

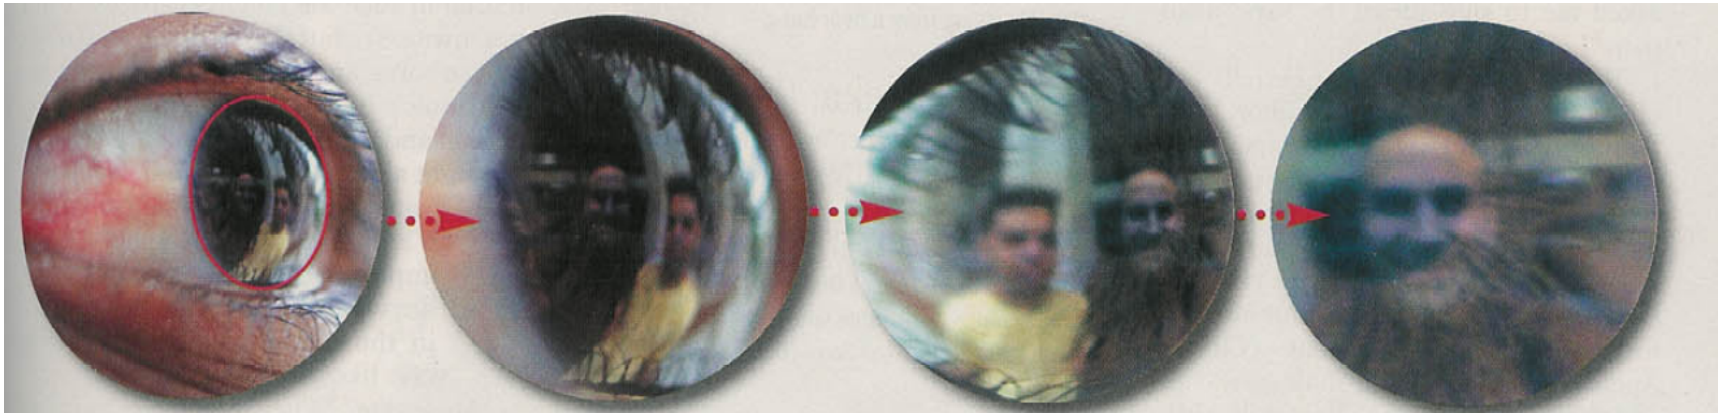

Figure: Source: Nayar and Nishino, Eyes for Relighting

[Source: N. Snavely]

Why study Computer Vision?

- Some more CSI

- Can you see something on the wall?

Torralba & Freeman, CVPR'12

Why study Computer Vision?

- Some more CSI

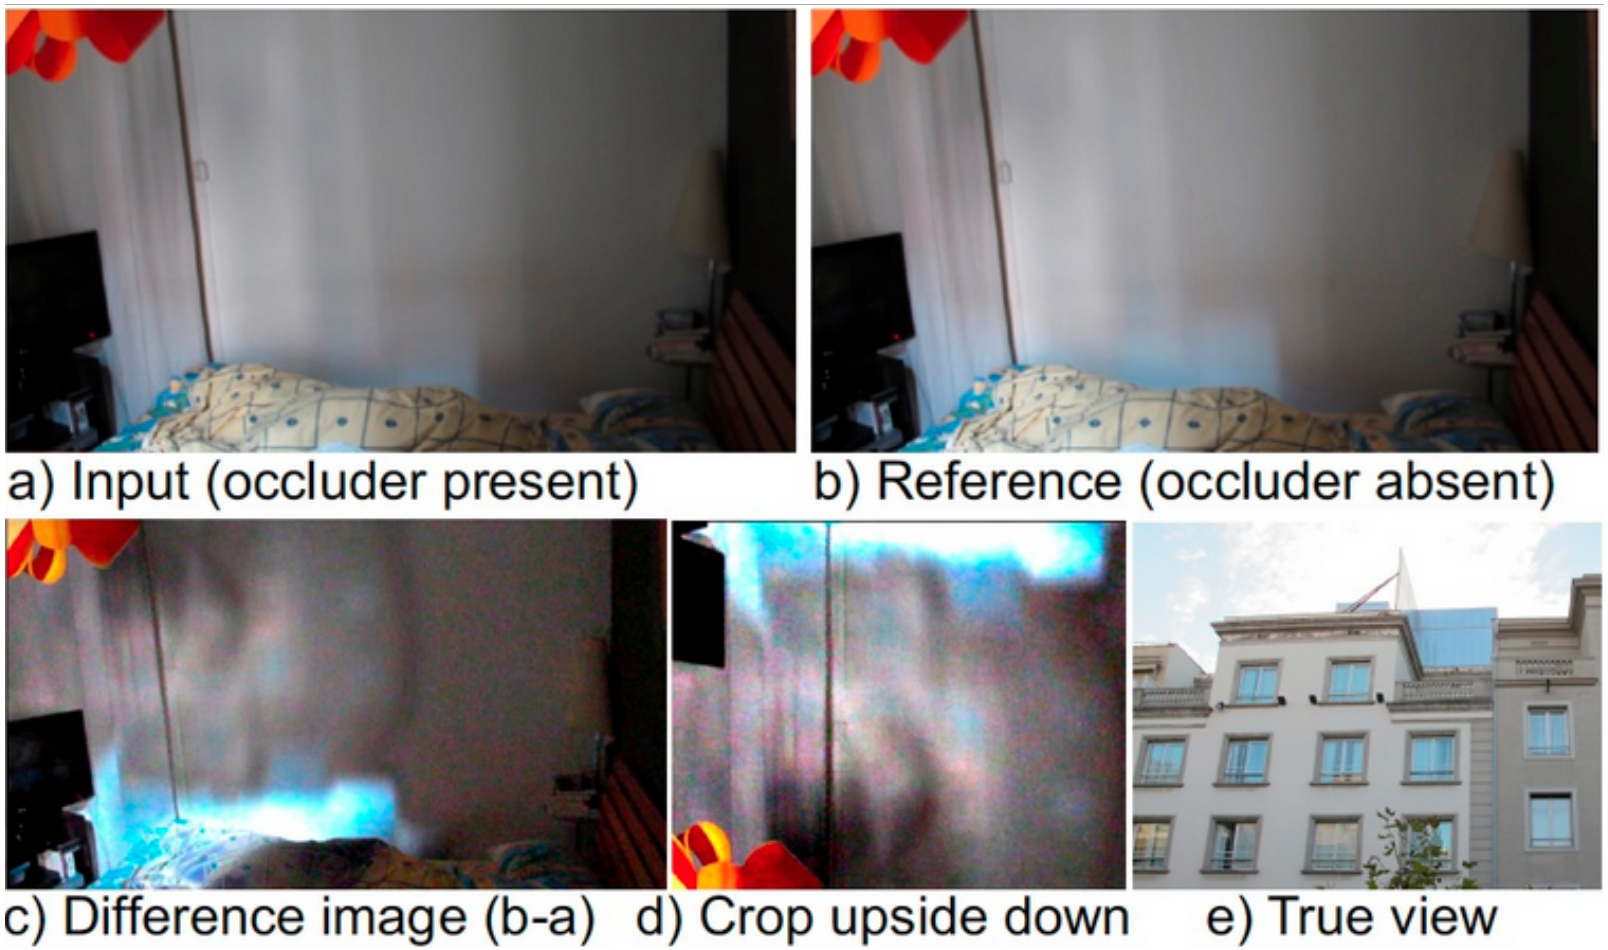

Why study Computer Vision?

- Object recognition (in mobile phones)

[Source: S. Seitz]

Why study Computer Vision?

- Recognizing movie posters (in mobile phones)

iPhone Apps: kooaba (www.kooaba.com)

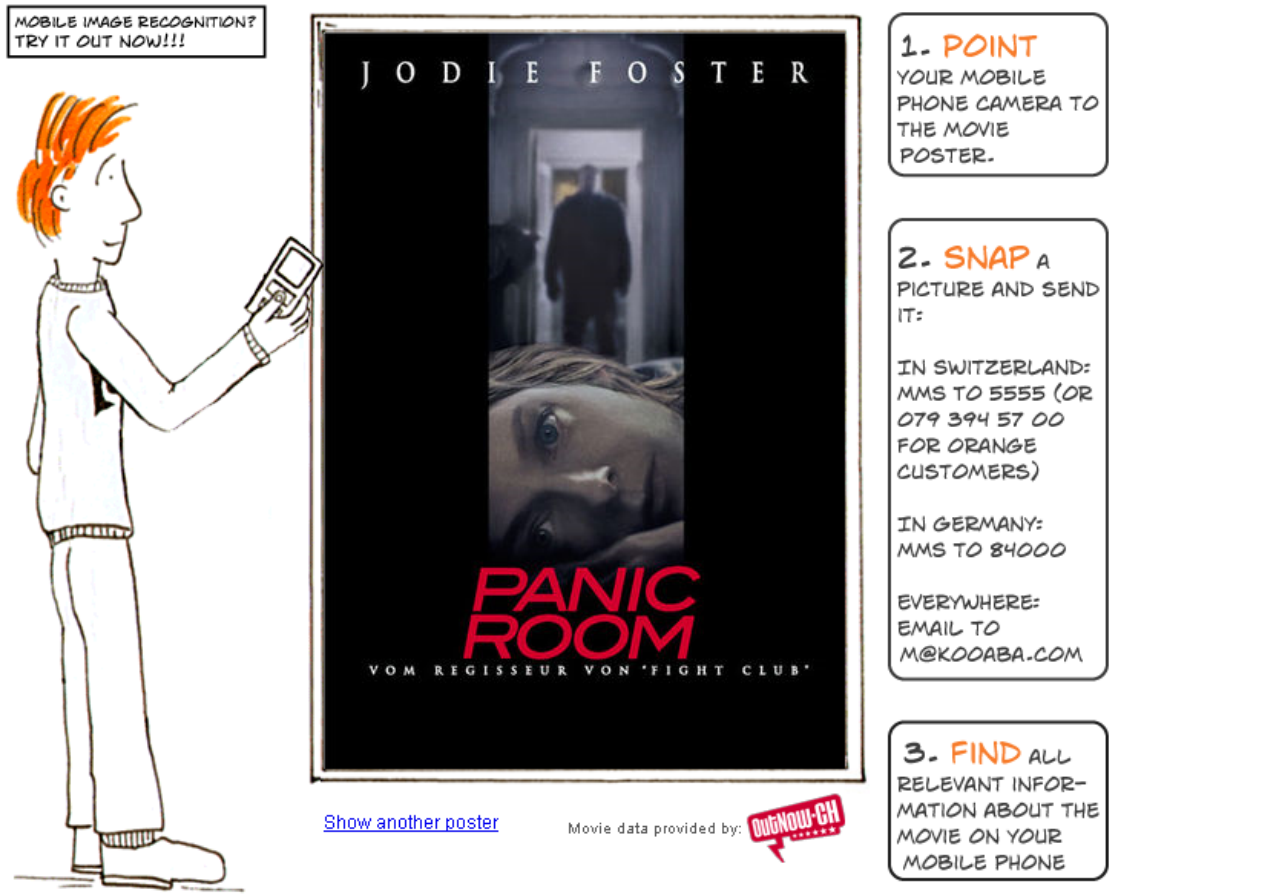

Source: S. Lazebnik

Why study Computer Vision?

- Games, games & games: 3D Pose Estimation with Depth Sensors

[Source: Microsoft Kinect]

How It All Began...

How It All Began...

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

PROJECT MAC

Artificial Intelligence Group — July 7, 1966
Vision Memo. No. 100.

THE SUMMER VISION PROJECT

Seymour Papert

The summer vision project is an attempt to use our summer workers effectively in the construction of a significant part of a visual system. The particular task was chosen partly because it can be segmented into sub-problems which will allow individuals to work independently and yet participate in the construction of a system complex enough to be a real landmark in the development of "pattern recognition".

[Slide credit: A. Torralba]

50 years and thousands of PhDs later...

Popular benchmarks:

The KITTI Vision Benchmark Suite
A project of Karlsruhe Institute of Technology and Toyota Technological Institute at Chicago

PASCAL2
Pattern Analysis, Statistical Modelling and Computational Learning

Car

Rank	Method	Setting	Code	Moderate	Easy	Hard	Runtime	Environment	Compare
1	SubCat			66.32 %	**81.94 %**	51.10 %	0.3 s	6 cores @ 2.5 Ghz (Matlab + C/C++)	
E. Ohn-Bar and M. Trivedi: Fast and Robust Object Detection Using Visual Subcategories. Computer Vision and Pattern Recognition Workshops Mobile Vision 2014.									
2	AOG		code	67.03 %	80.26 %	**55.60 %**	3 s	4 cores @ 2.5 Ghz (Matlab)	
B. Li, T. Wu and S. Zhu: Integrating Context and Occlusion for Car Detection by Hierarchical And-Or Model. ECCV 2014.									
3	SubCat-NoOcc			58.91 %	79.90 %	44.81 %	0.3 s	6 cores @ 2.5 Ghz (Matlab + C/C++)	
Anonymous submission									

Cyclist

Rank	Method	Setting	Code	Moderate	Easy	Hard	Runtime	Environment	Compare
1	pAUC			**38.03 %**	**51.62 %**	**33.38 %**	60 s	1 core @ 2.5 Ghz (Matlab + C/C++)	
S. Paisitkriangkrai, C. Shen and A. Hengel: Efficient pedestrian detection by directly optimizing the partial area under the ROC curve. ICCV 2013.									
2	DPM-C8B1			29.04 %	43.49 %	26.20 %	15 s	4 cores @ 2.5 Ghz (Matlab + C/C++)	
Anonymous submission									
3	LSVM-MDPM-us		code	29.88 %	38.84 %	27.31 %	10 s	4 cores @ 3.0 Ghz (C/C++)	
P. Felzenszwalb, R. Girshick, D. McAllester and D. Ramanan: Object Detection with Discriminatively Trained Part-Based Models. PAMI 2010.									

	mean	aero plane	bicycle	bird	boat	bottle	bus	car	cat	chair	cow	dining table	dog	horse	motor bike	person	potted plant	sheep	sofa	train	tv/ monitor	submission date
Feature Edit	56.4	74.8	69.2	55.7	41.9	36.1	64.7	62.3	69.5	31.3	53.3	43.7	69.9	64.0	71.8	60.5	32.7	63.0	44.1	63.6	56.6	2014-Sep-04
R-CNN (bbox reg)	53.7	71.8	65.8	53.0	36.8	35.9	59.7	60.0	69.9	27.9	50.6	41.4	70.0	62.0	69.0	58.1	29.5	59.4	39.3	61.2	52.4	2014-Mar-13
R-CNN	50.2	67.1	64.1	46.7	32.0	30.5	56.4	57.2	65.9	27.0	47.3	40.9	66.6	57.8	65.9	53.6	26.7	56.5	38.1	52.8	50.2	2014-Jan-30

50 years and thousands of PhDs later...

- Algorithms work **pretty** well
- Still some embarrassing mistakes...
- The general vision problem is not yet solved

Where pink means “person”

Why is vision hard?

Why is vision hard?

- Half of the cerebral cortex in primates is devoted to processing visual information. This is a lot. Means that vision has to be pretty hard!

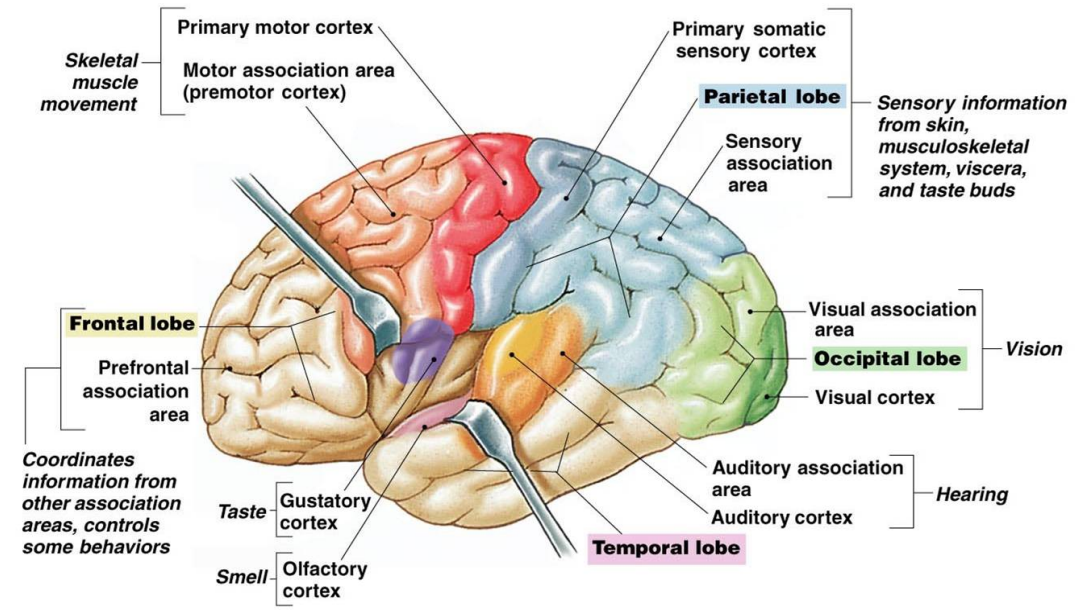

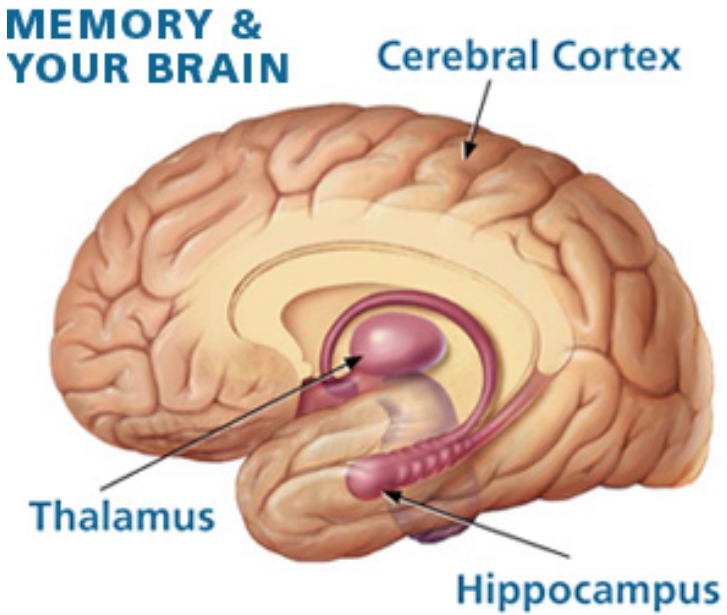

Why is vision hard?

Lots of data to process:

- Thousands to millions of pixels in an image
- 100 hours of video added to YouTube per minute [source: YouTube]
- Over 6 billion hours of video are watched each month on YouTube – almost an hour for every person on Earth [source: YouTube]

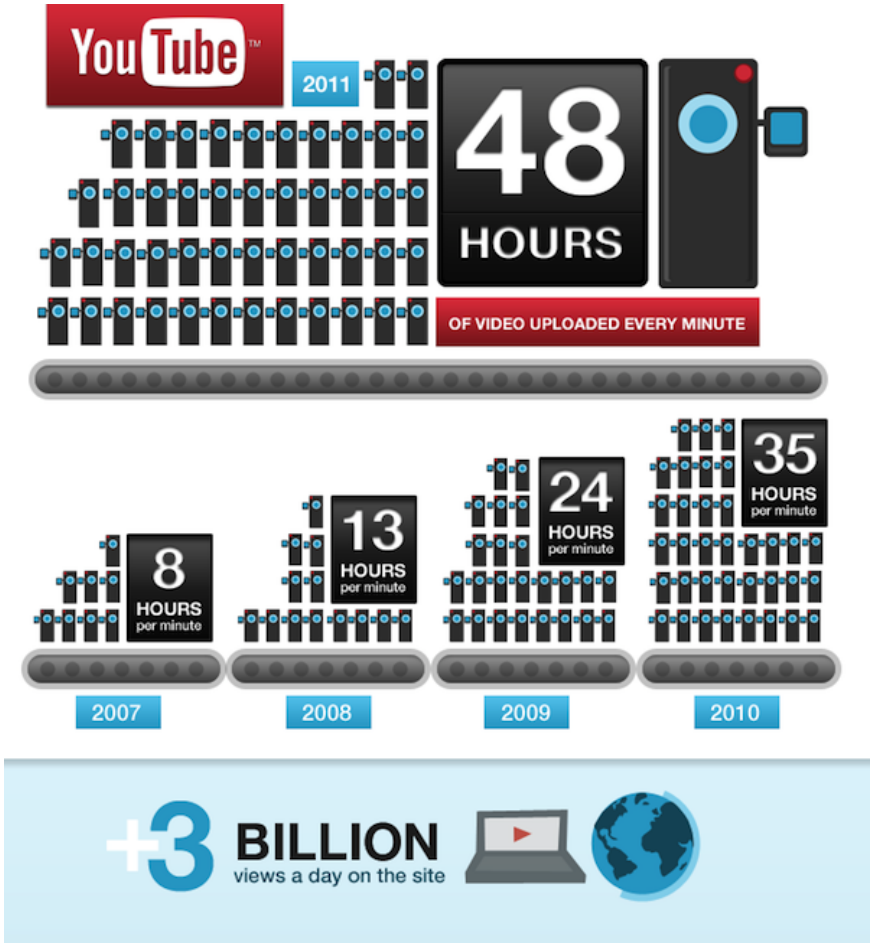

Why is vision hard?

Lots of data to process:

- $\sim$ 5000 new tagged photos added to Flickr per minute (7M per day)
- $\sim$ 60M photos uploaded to Instagram every day [source: Instagram]

How many photos are uploaded to Flickr every day, month, year?

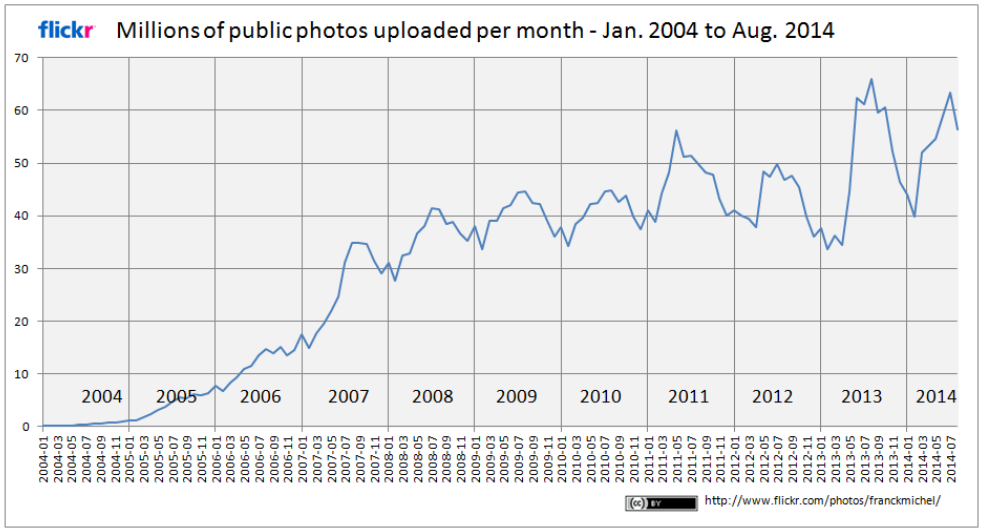

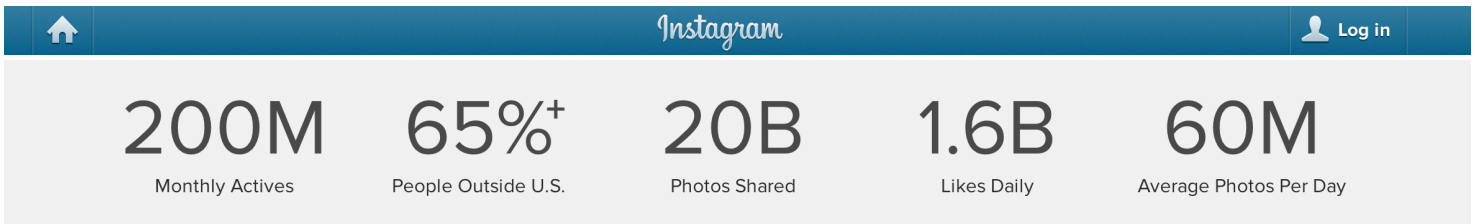

Why is vision hard?

All this is dog...

[slide adopted from: R. Urtasun]

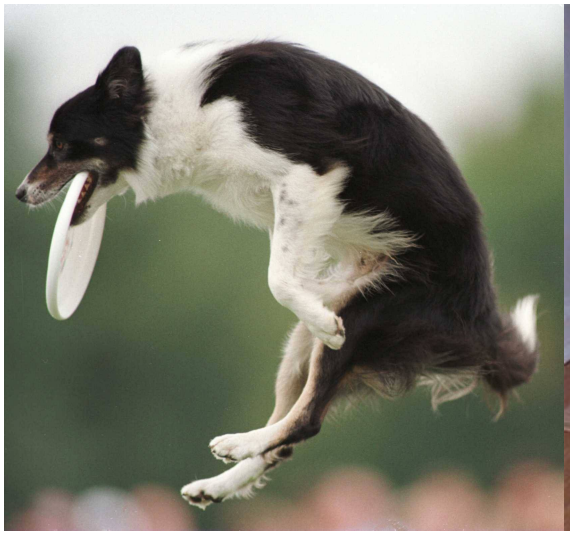

Why is vision hard?

Biederman, 1987

[slide credit: R. Urtasun]

Why is vision hard?

- Human vision seems to work quite well.
- How well does it really work?
- Let's play some games!

How good are humans?

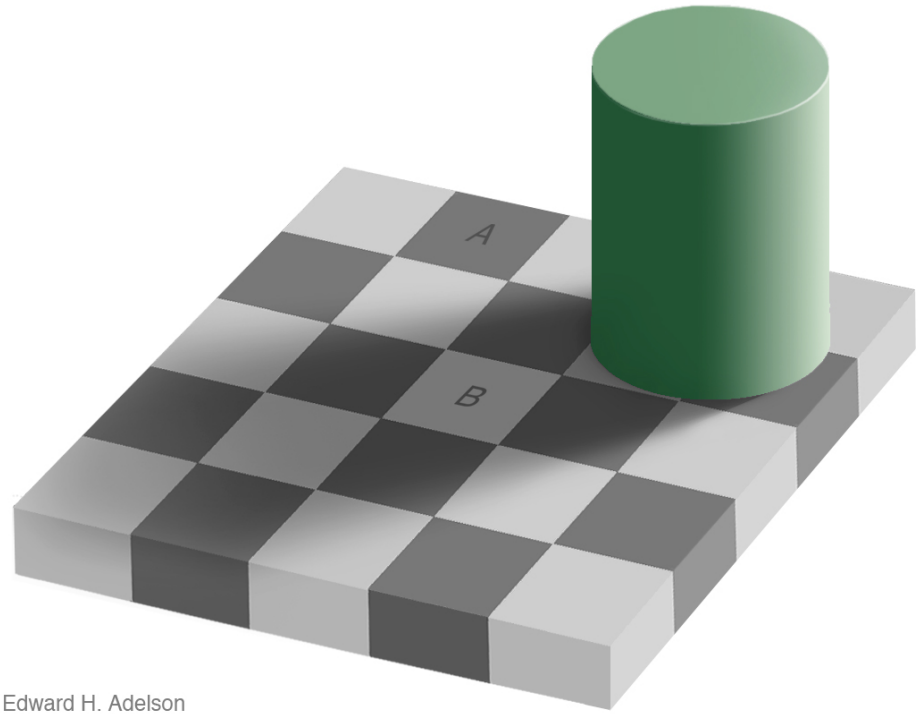

Edward H. Adelson

- Which square is lighter, A or B?

[Slide credit: A. Torralba]

How good are humans?

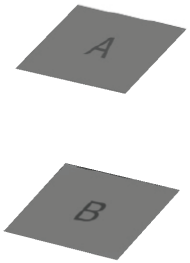

Edward H. Adelson

- Which square is lighter, A or B?

[Slide credit: A. Torralba]

How good are humans?

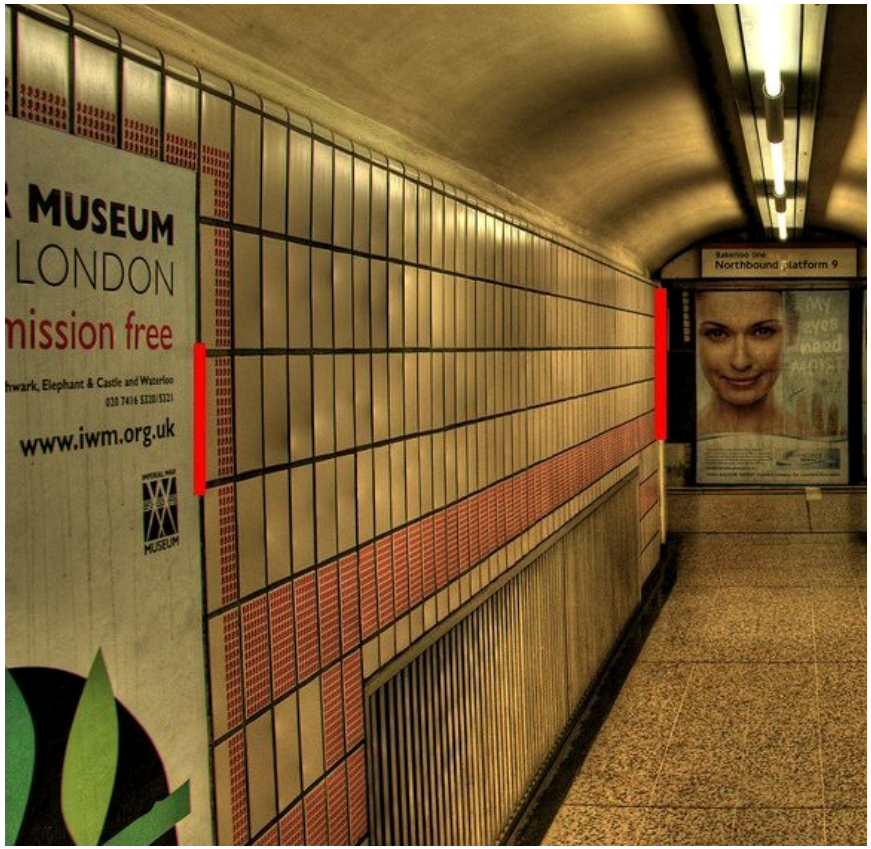

Figure: 2006 Walt Anthony

- Which red line is longer?

[Slide credit: A. Torralba]

How good are humans?

Figure: 2006 Walt Anthony

- Which red line is longer?

[Slide credit: A. Torralba]

How good are humans?

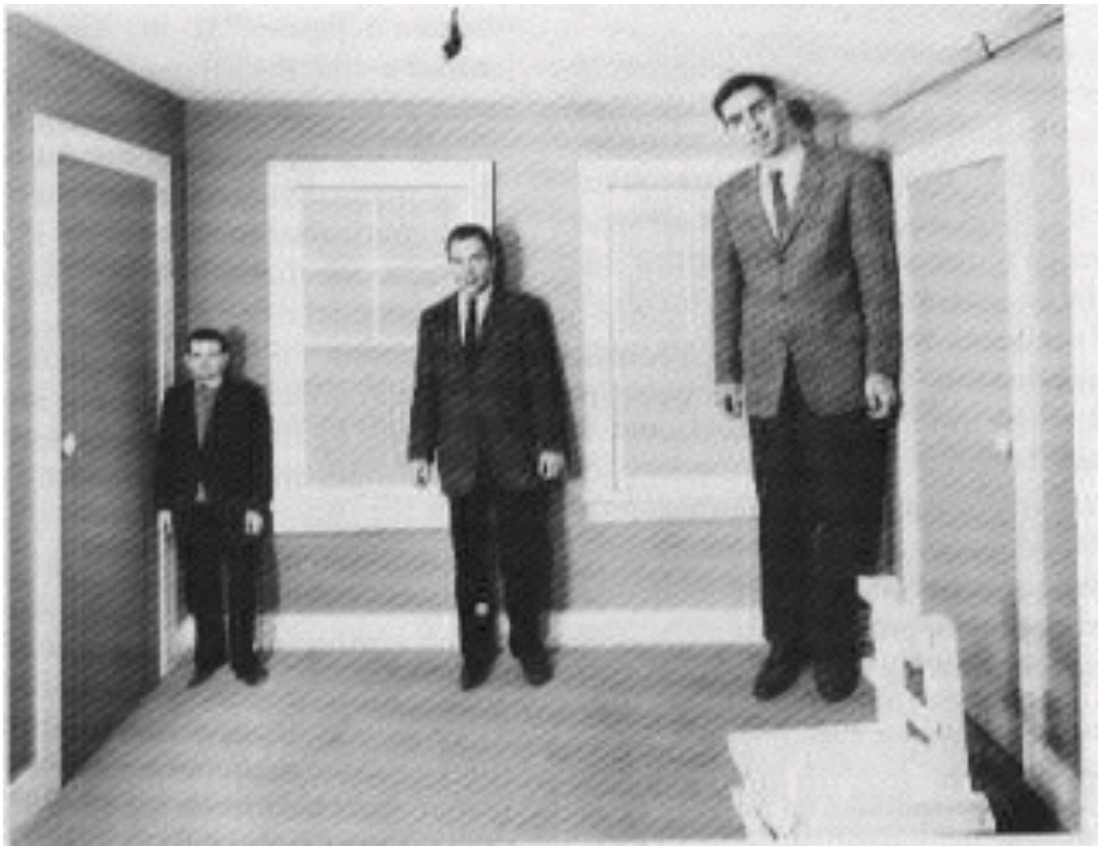

Figure: Ames room

- Assumptions can be wrong

[Slide credit: A. Torralba]

How good are humans?

Figure: Chabris & Simons

- Count the number of times the white team pass the ball
- Concentrate, it's difficult!

How good are humans?

Figure: Simons et al. (more videos here: `http://www.perceptionweb.com/misc.cgi?id=p3104`)

- Is something happening in the picture?

How good are humans?

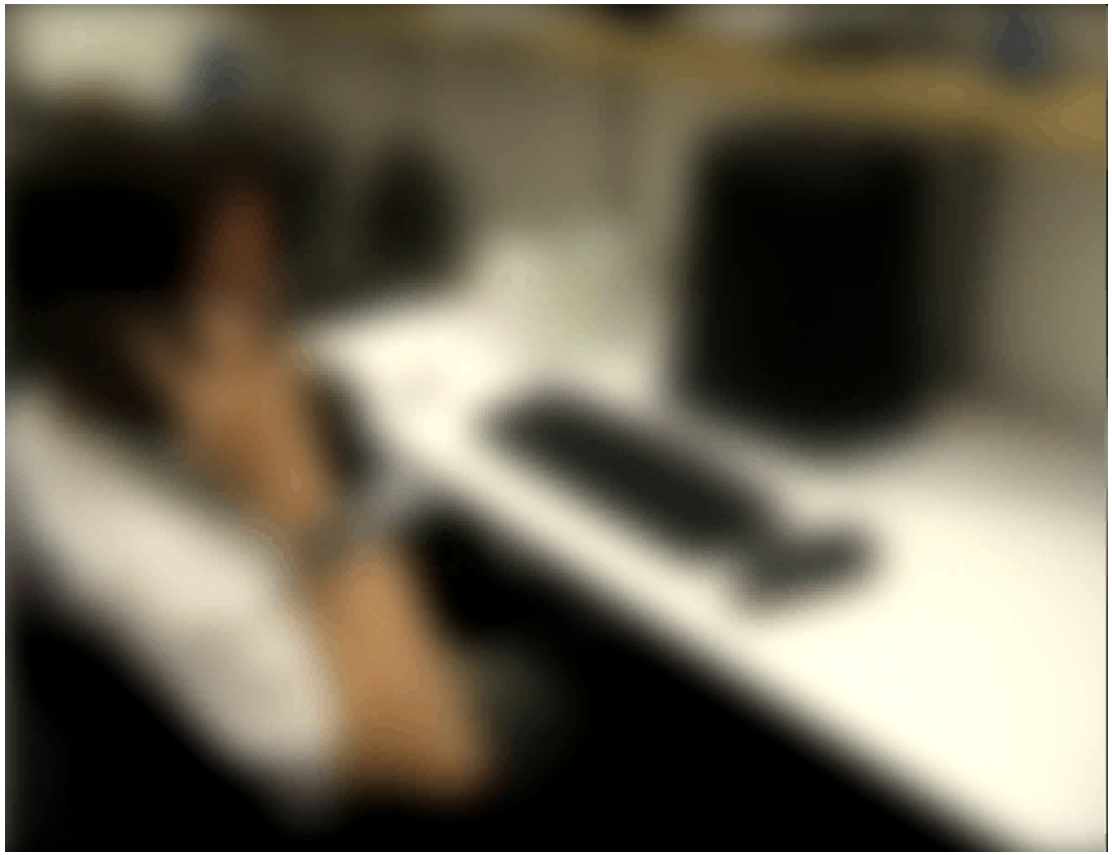

Figure: Torralba et al.

- Can you describe what's going on in the video?

How good are humans?

Figure: Torralba et al.

- Can you describe what's going on in the video?

What do I need...

What do I need to become a good Computer Vision researcher?

- Some math knowledge
- Good programming skills
- Imagination
- Even better intuition
- Lots of persistence
- Some luck always helps