Introduction

Computer Vision
Fall 2018
Columbia University

Cameras everywhere

Also scary times

The New York Times

Facial Recognition Is Accurate, if You're a White Guy

By Steve Lohr

Feb. 9, 2018

Facial recognition technology is improving by leaps and bounds. Some commercial software can now tell the gender of a person in a photograph.

When the person in the photo is a white man, the software is right 99 percent of the time.

But the darker the skin, the more errors arise — up to nearly 35 percent for images of darker skinned women, according to a new study that breaks fresh ground by measuring how the technology works on people of different races and gender.

What is vision?

“What does it mean, to see? The plain man's answer (and Aristotle's, too) would be, to know what is where by looking.”

— David Marr, 1982

1945 - 1980 (35 years old)

Computational Photography

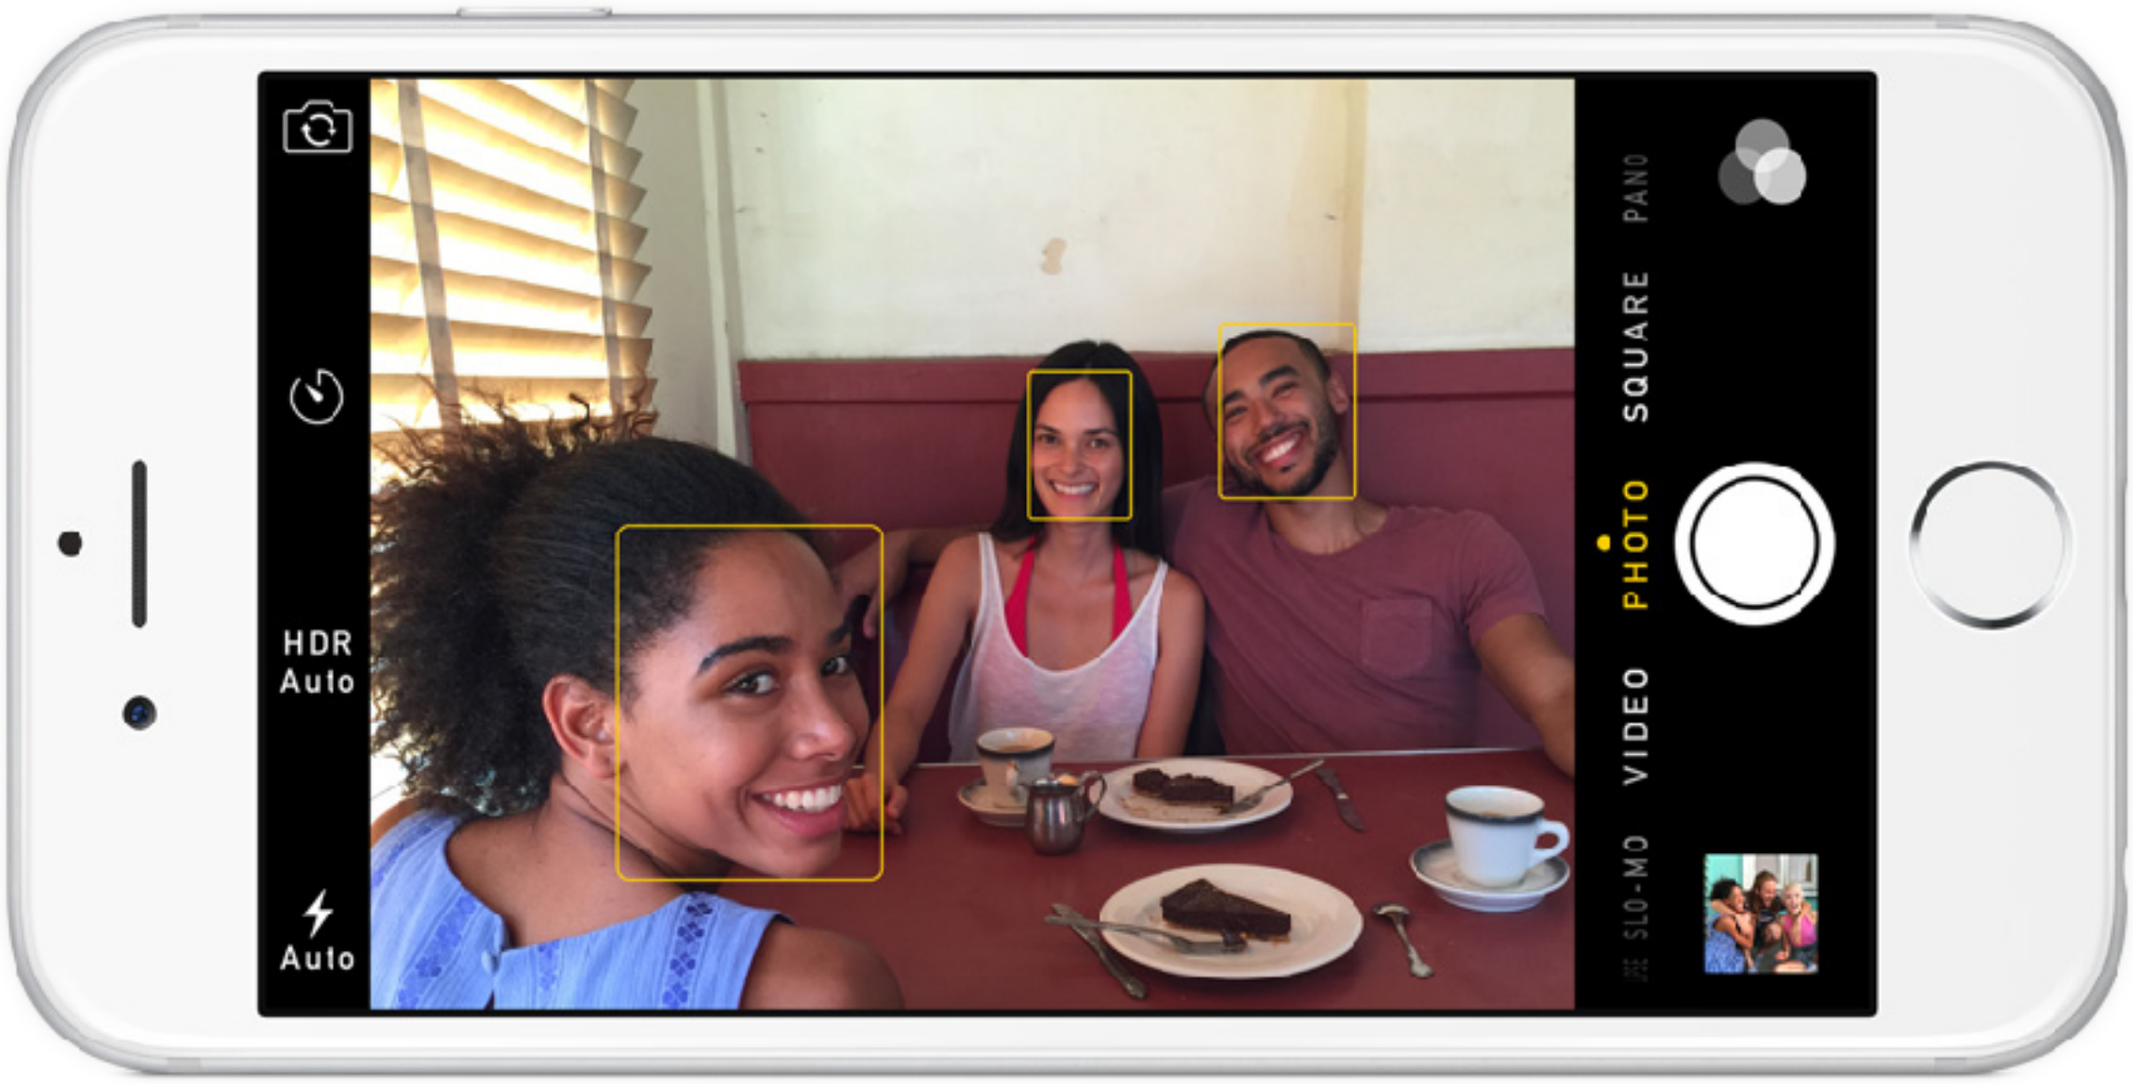

Biometrics

1984

- "the most recognized photograph" in the history of the *National Geographic* magazine

- No one knew her identity…

Biometrics

1984 2002

Optical Character Recognition

Security and Tracking

"The work was painstaking and mind-numbing: One agent watched the same segment of video 400 times. The goal was to construct a timeline of images, following possible suspects as they moved along the sidewalks. It took a couple of days" Washington Post

Health

Gaming

Shopping

Special Effects

Visual Search

Self-driving Cars

Space Exploration

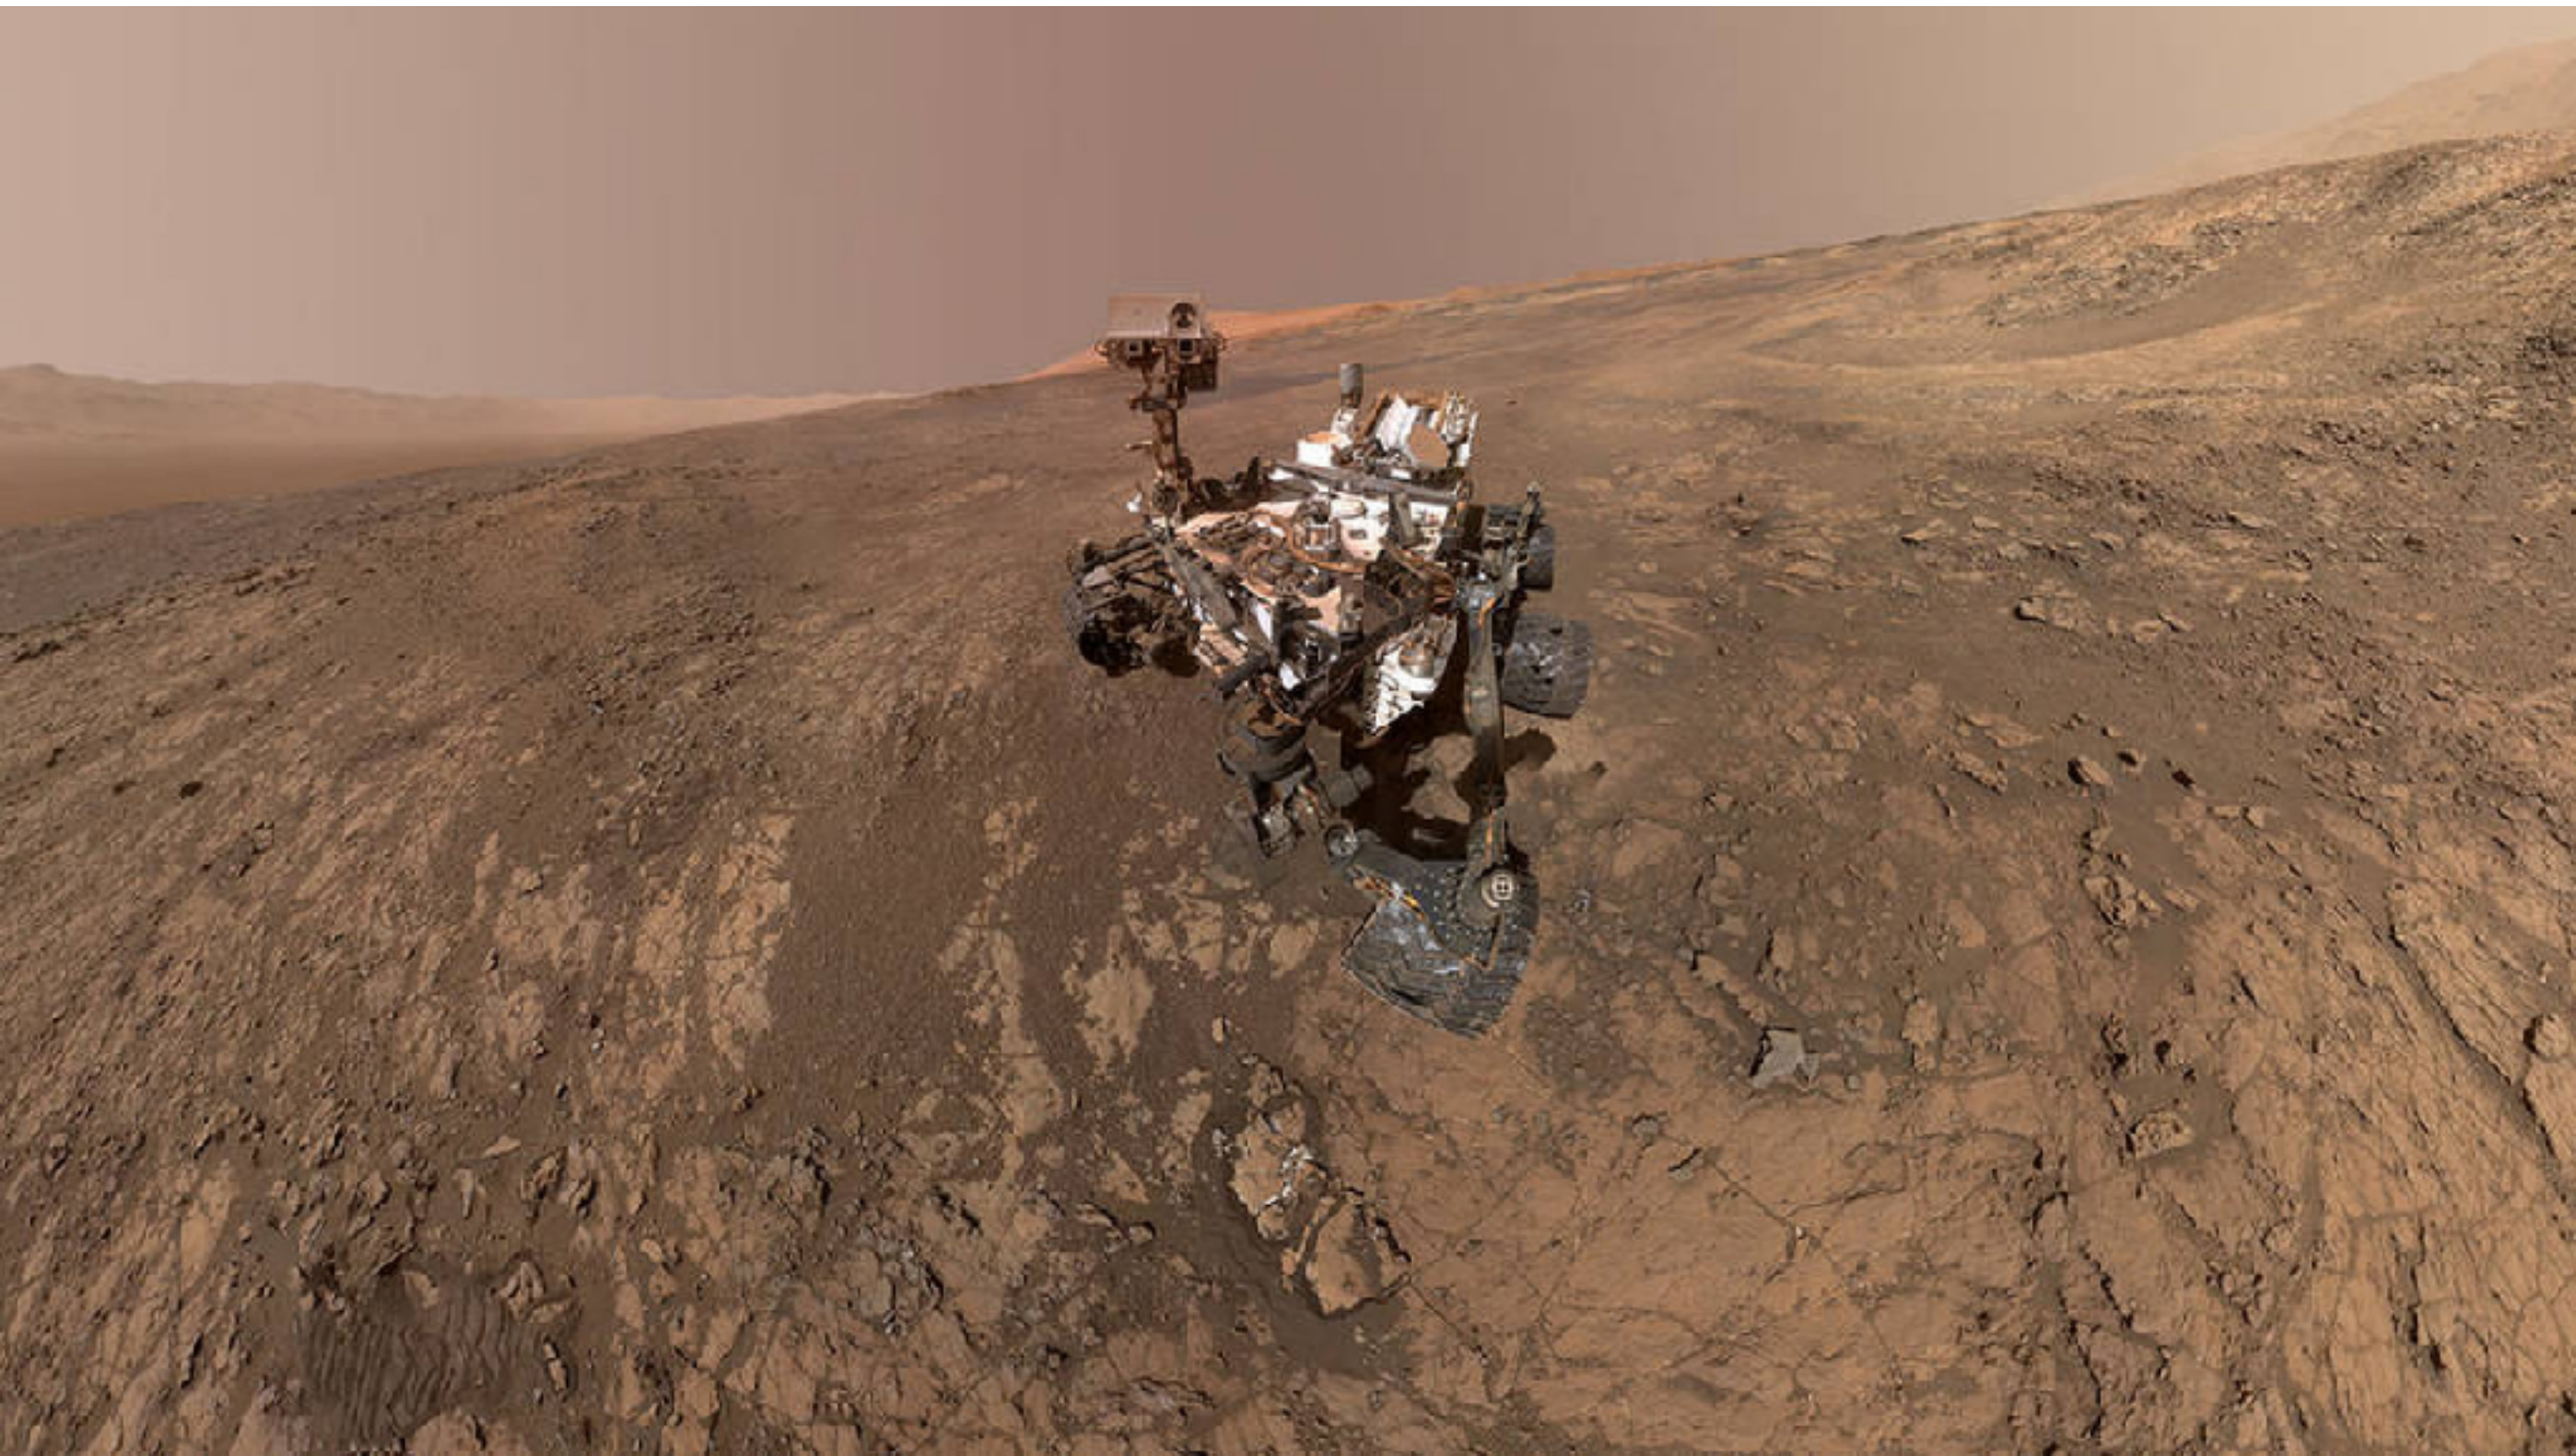

Augmented Reality

Worldwide Insight

Walmart in Wichita, Kansas

What is vision?

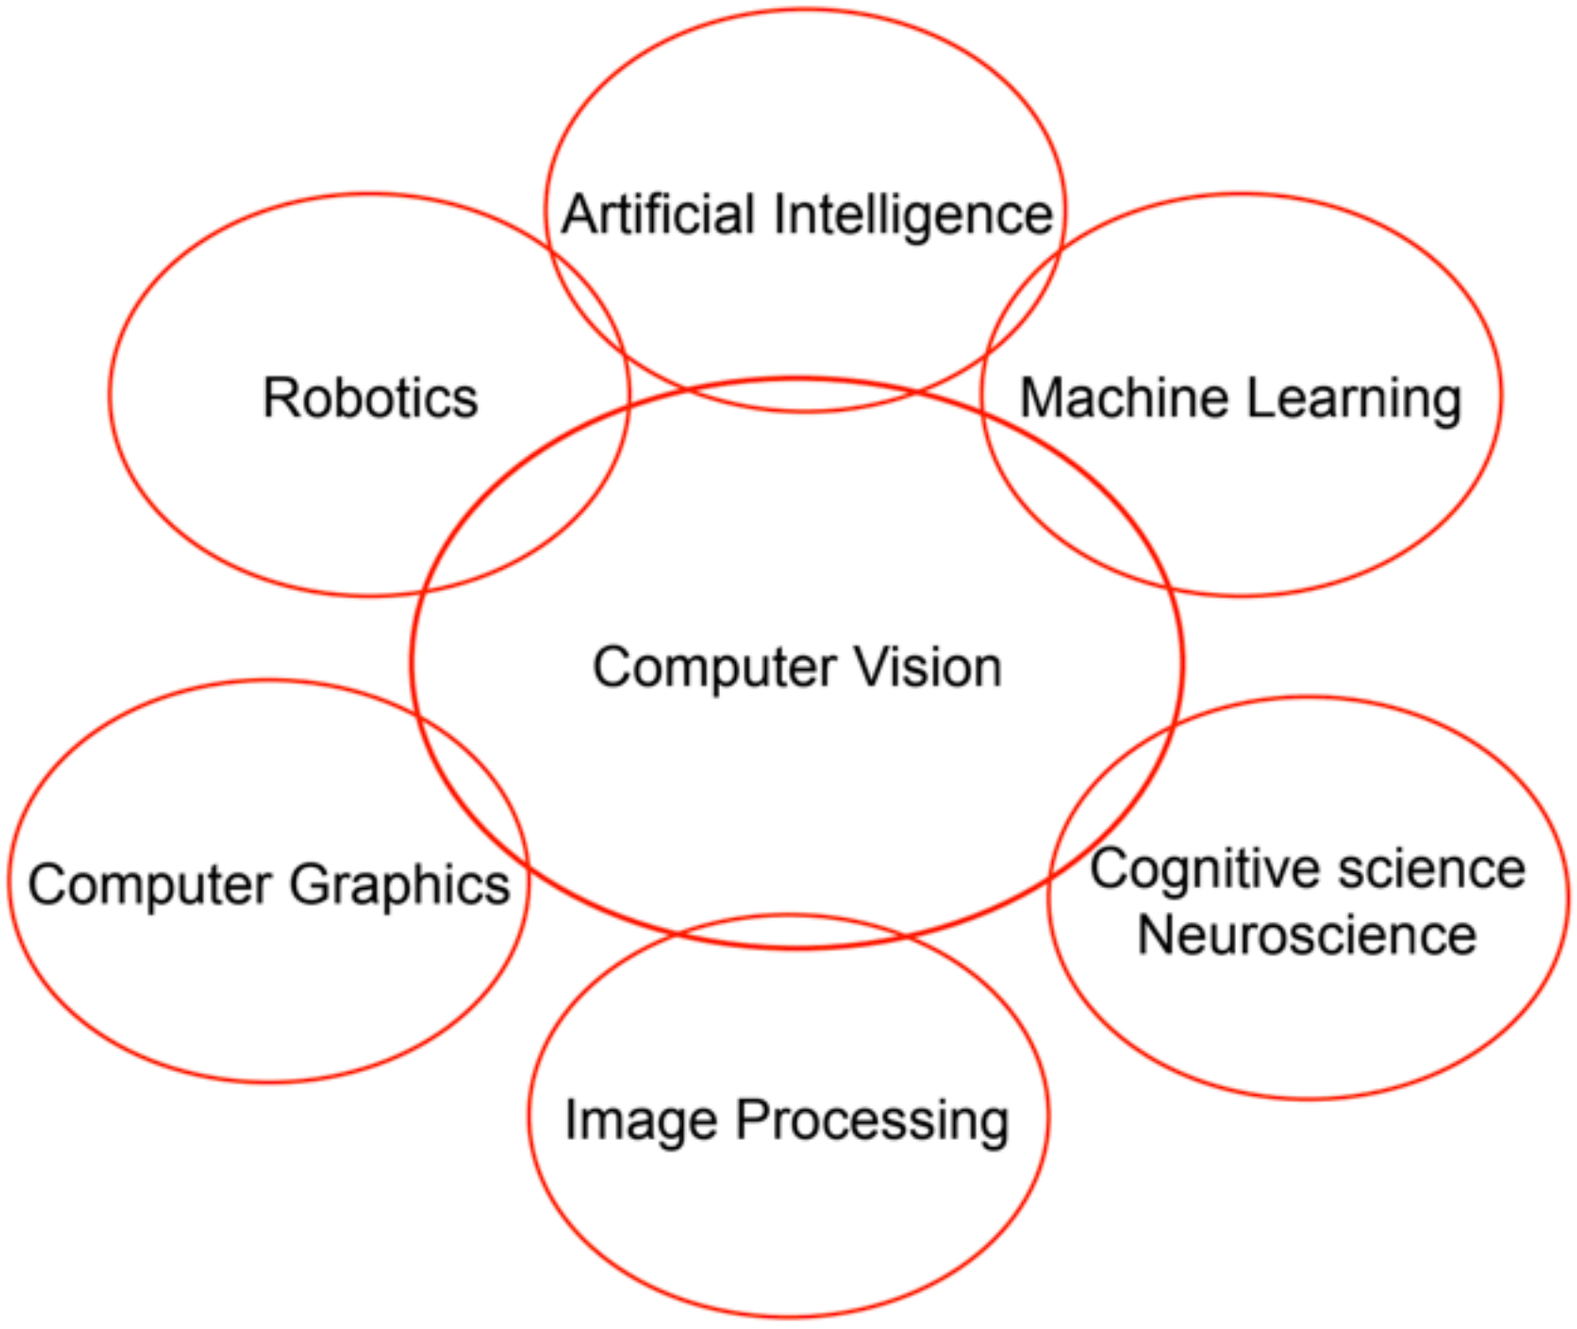

Slide credit: Kristen Grauman

Image Formation

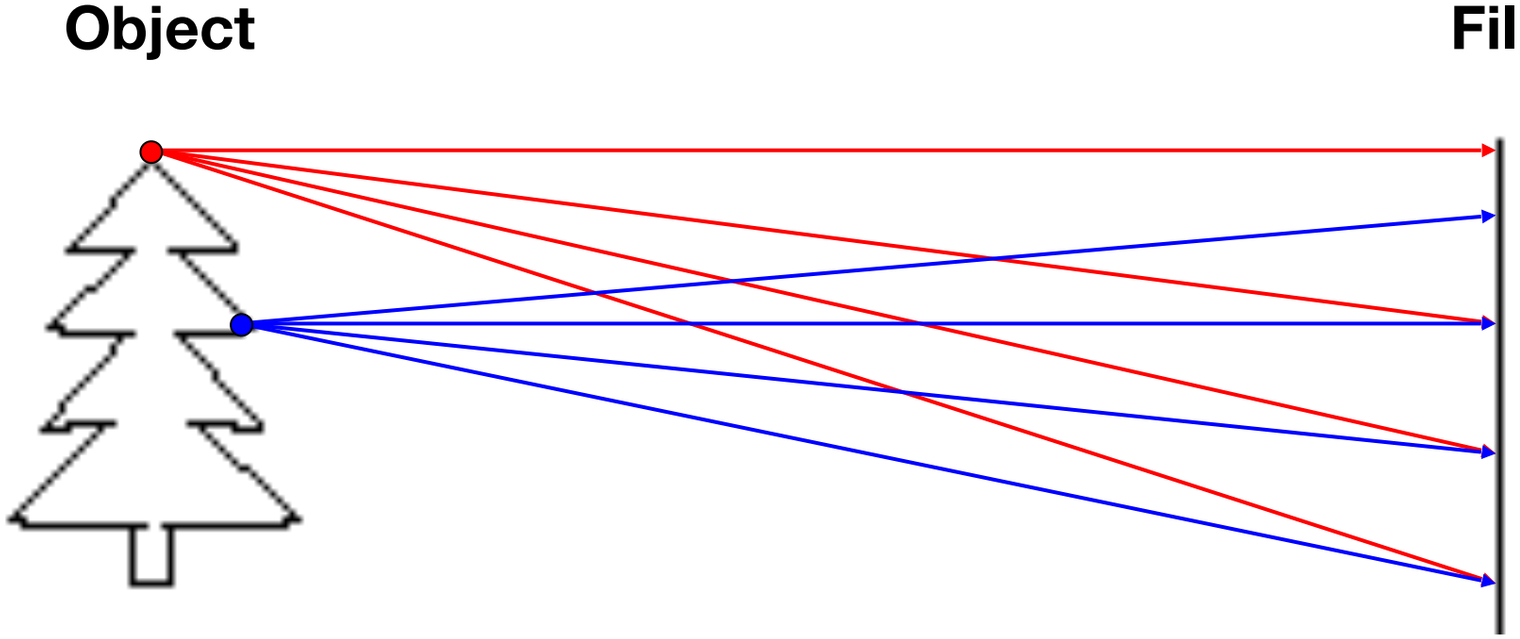

Slide credit: Steve Seitz

Image Formation

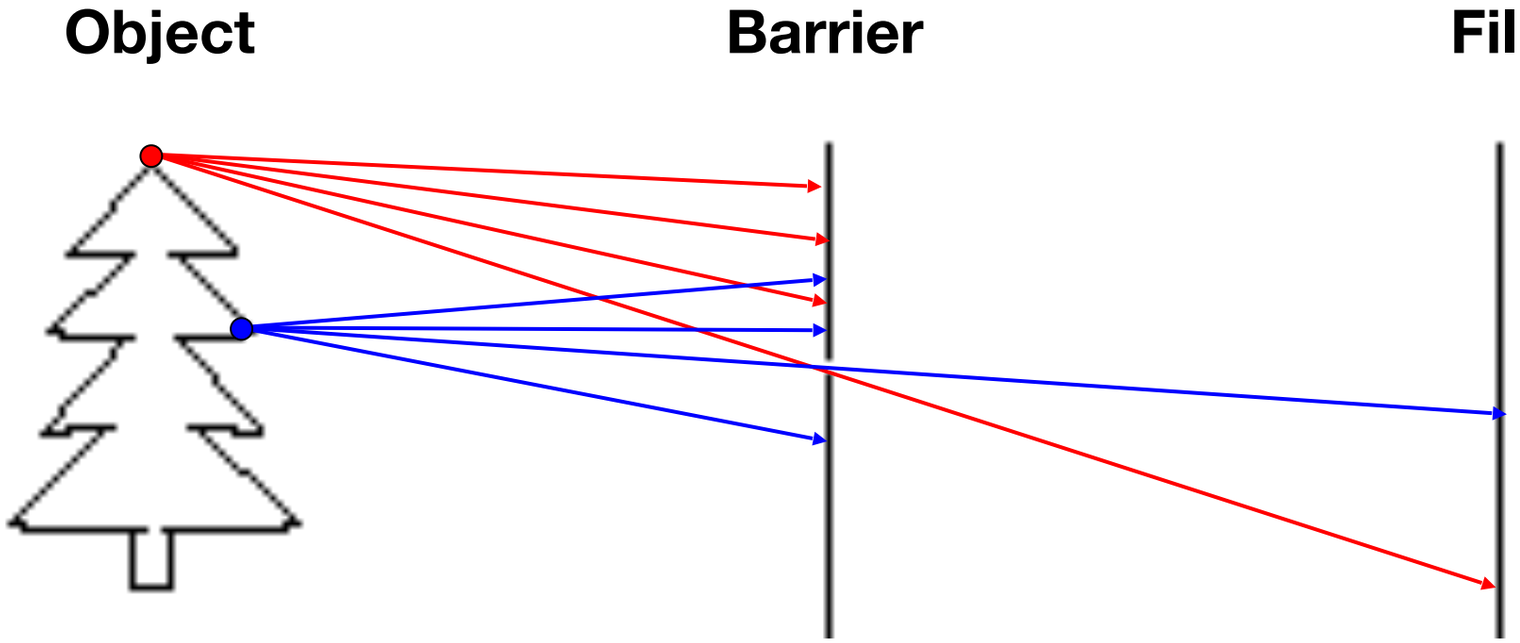

Add a barrier to block off most of the rays

Slide credit: Steve Seitz

Representing Digital Images

234	7	89	7	98	98	7	9	7	5
43	7	0	123	4	13	454	23	5	87
67	5	76	4	3	56	67	87	65	45
97	0	6	3	6	25	7	3	587	8
78	5	54	7	876	71	54	76	9	75
45	81	67	78	78	5	4	75	86	8
5	4	3	35	8	256	6	4	3	36
7	6	64	3	4	7	77	76	4	54
64	35	46	46	64	56	7	56	4	7
75	464	576	75	75	75	57	64	75	75

Slide credit: Deva Ramanan

Representing Digital Images

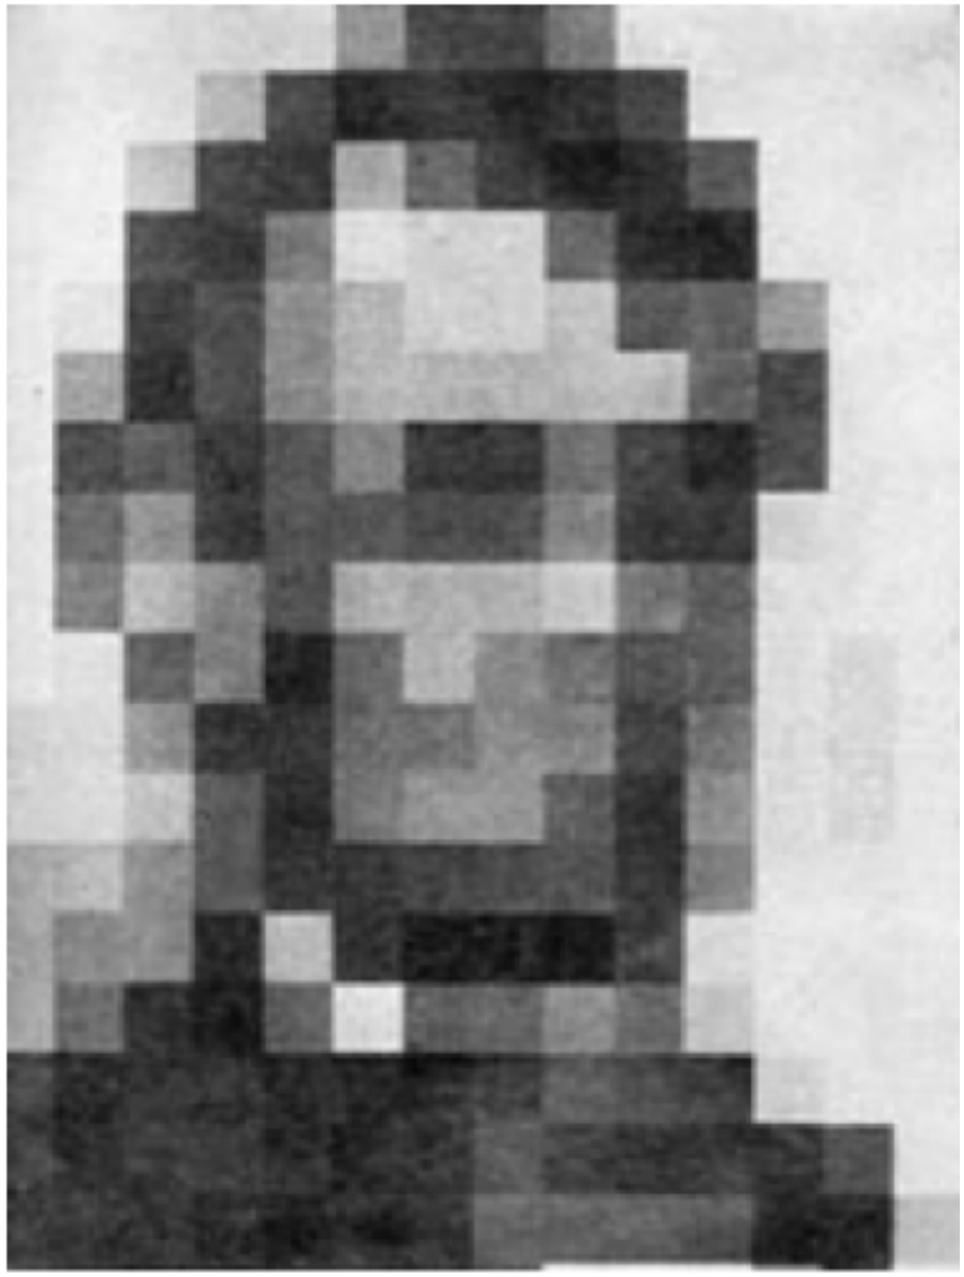

Slide credit: Deva Ramanan

Representing Color Images

Color images, RGB color space

R

G

B

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".

Illumination

“Neither Autopilot nor the driver noticed the white side of the tractor trailer against a brightly lit sky, so the brake was not applied.” — Tesla Company Blog

Slide credit: S. Ullman

Occlusion

René Magritte, 1957

Class Variation

Slide credit: Antonio Torralba

Clutter and Camouflage

Color

Motion

Slide credit: S. Lazebnik

Ill-posed Problem

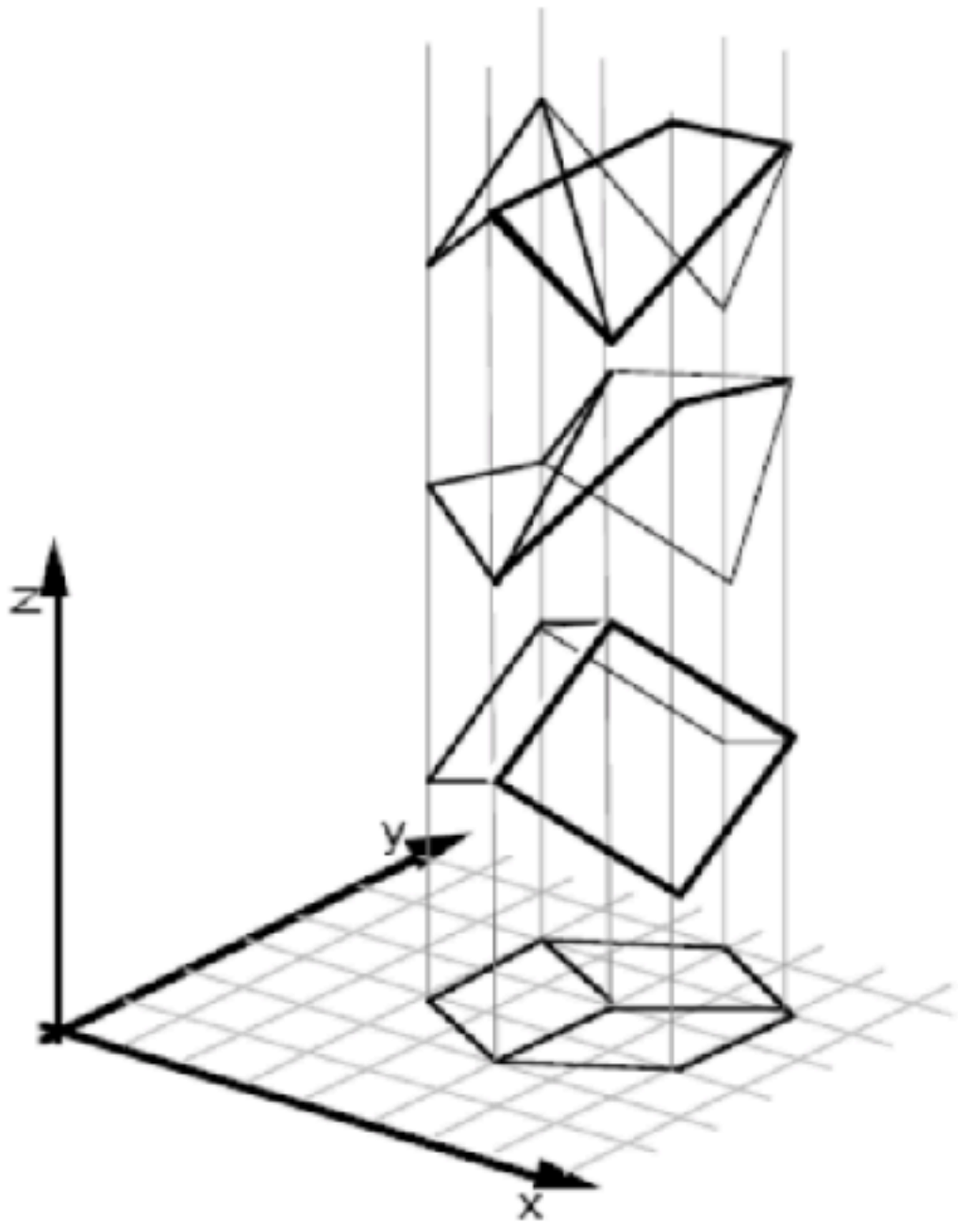

[Sinha and Adelson 1993]

Ill-posed Problem

Ill-posed Problem

Cambrian Explosion

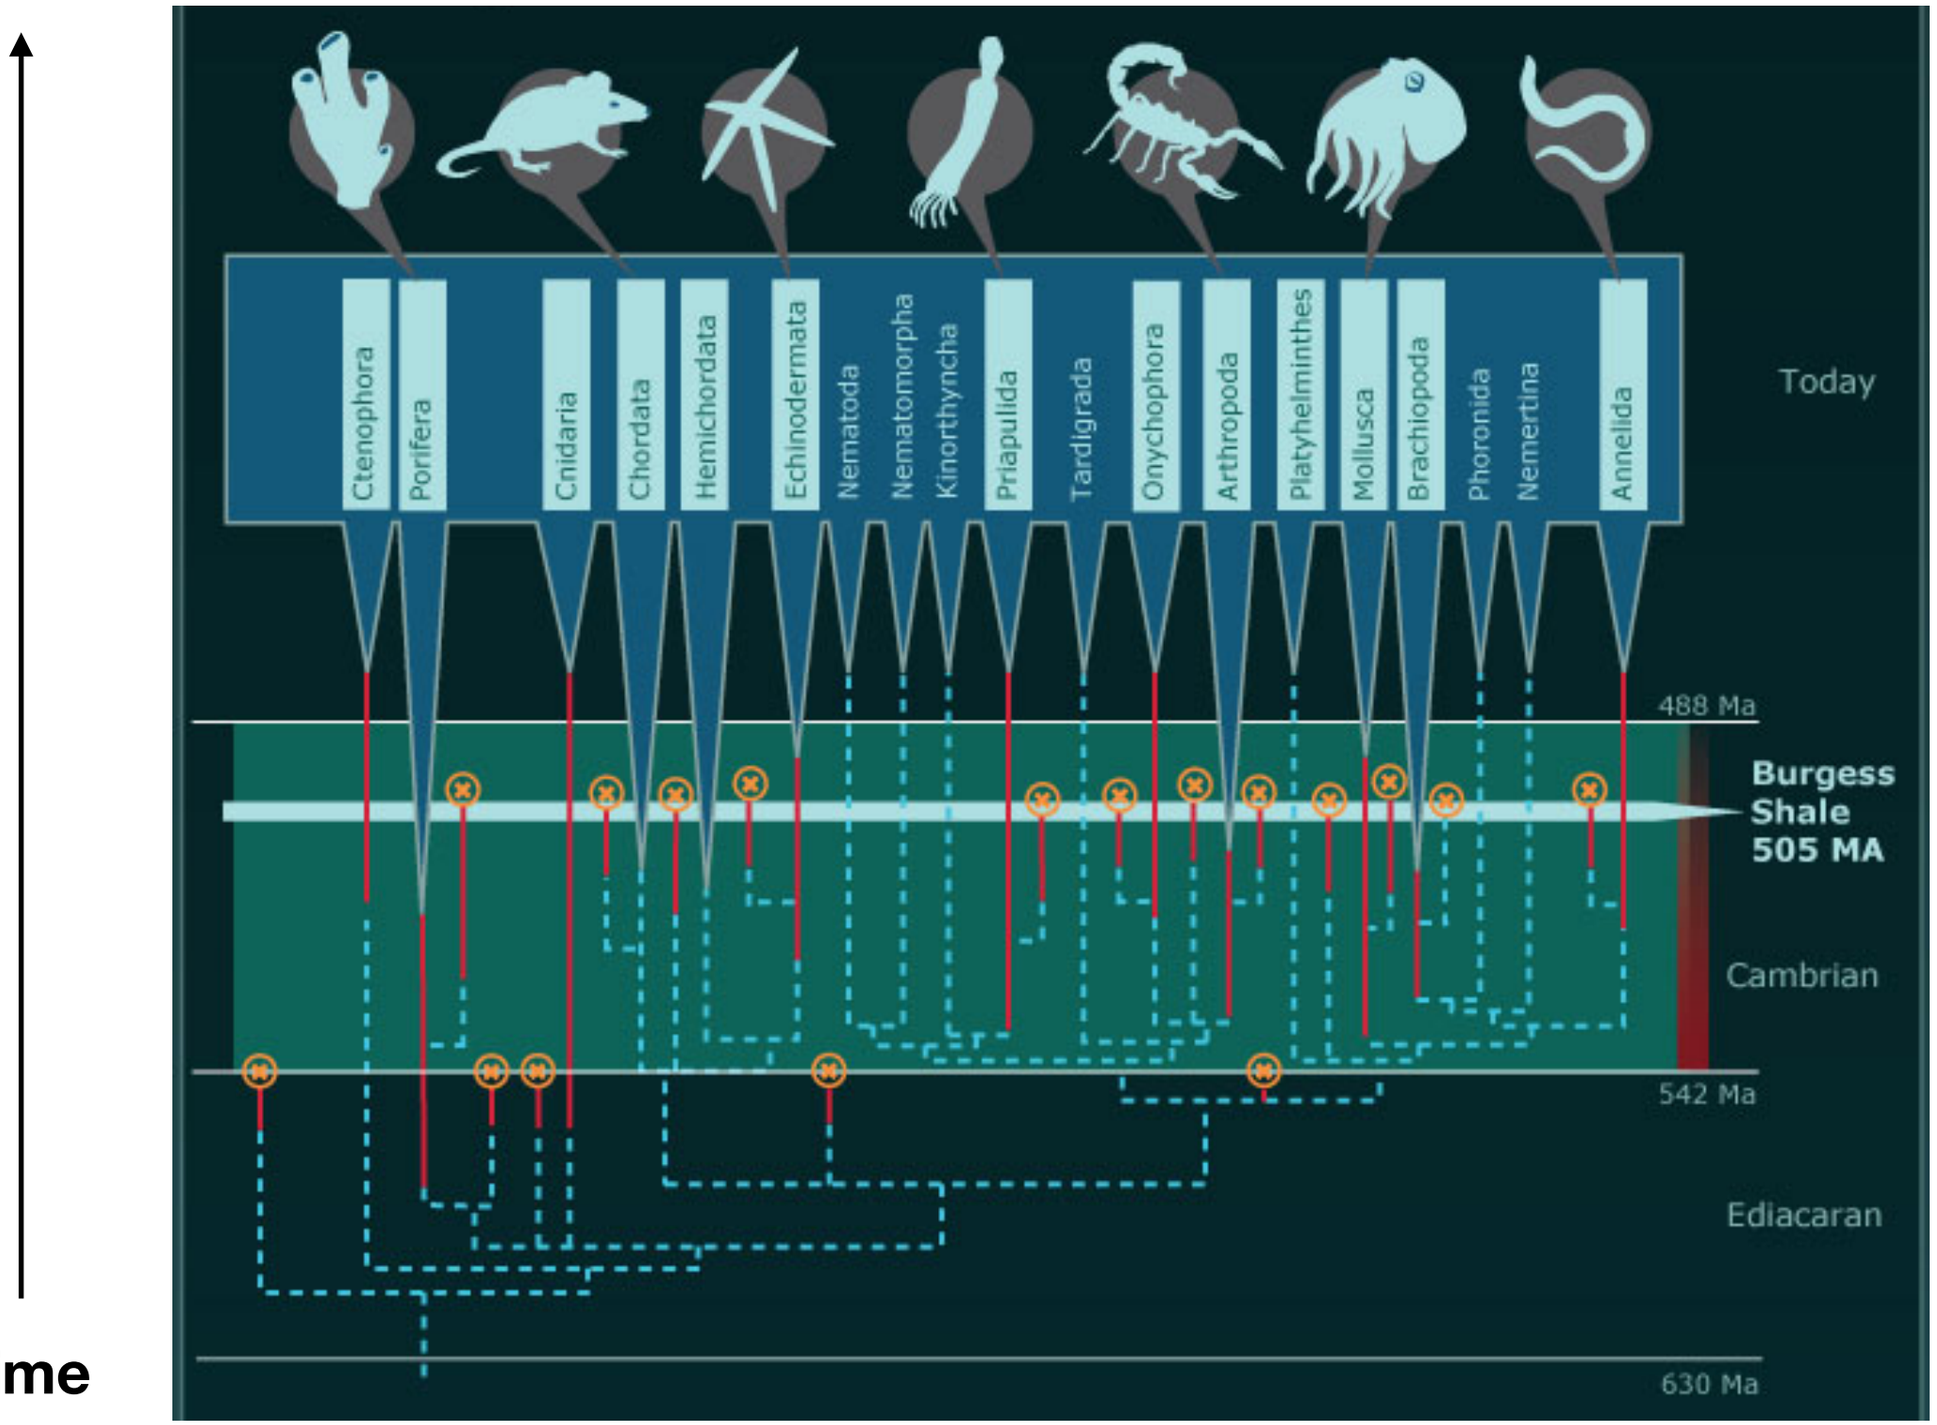

Cambrian Explosion

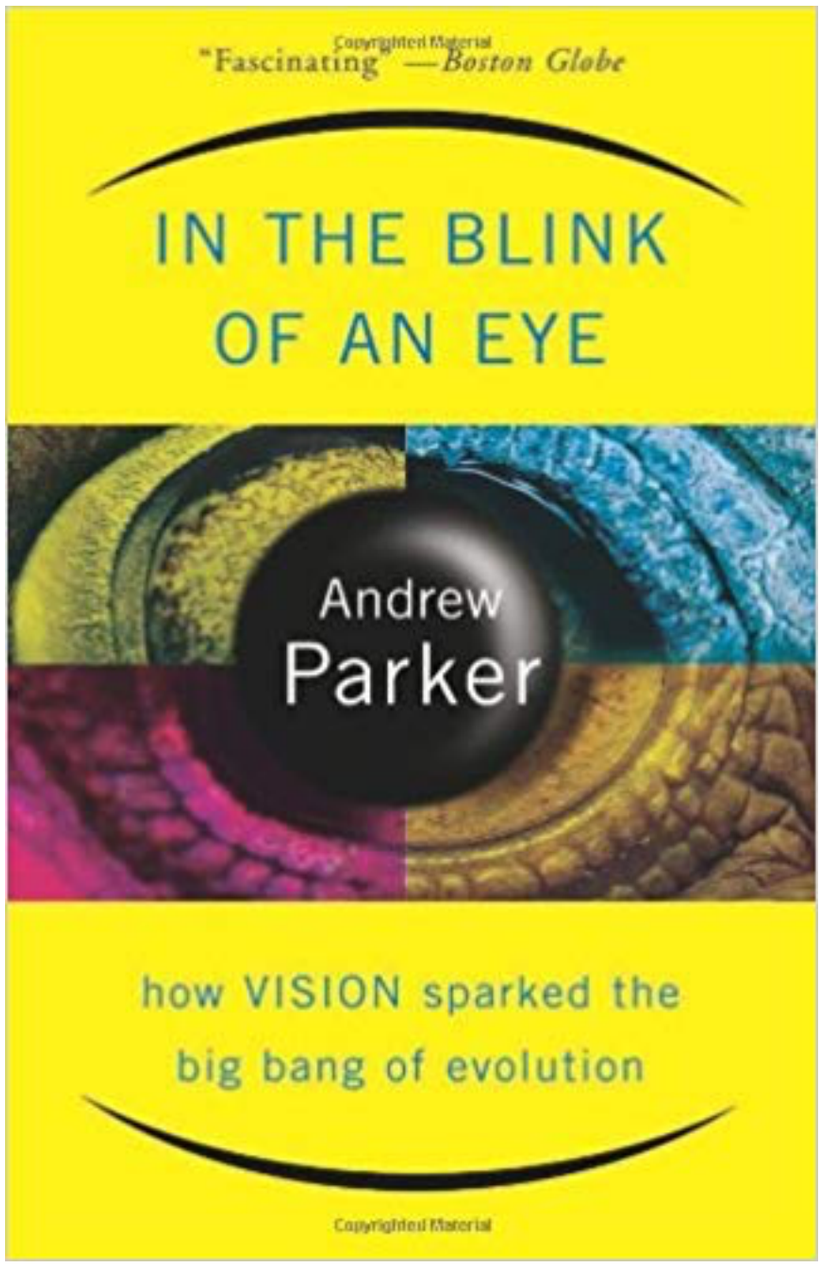

"The Cambrian Explosion is triggered by the sudden evolution of vision," which set off an evolutionary arms race where animals either evolved or died.
— Andrew Parker

Slide credit: Fei-Fei Li

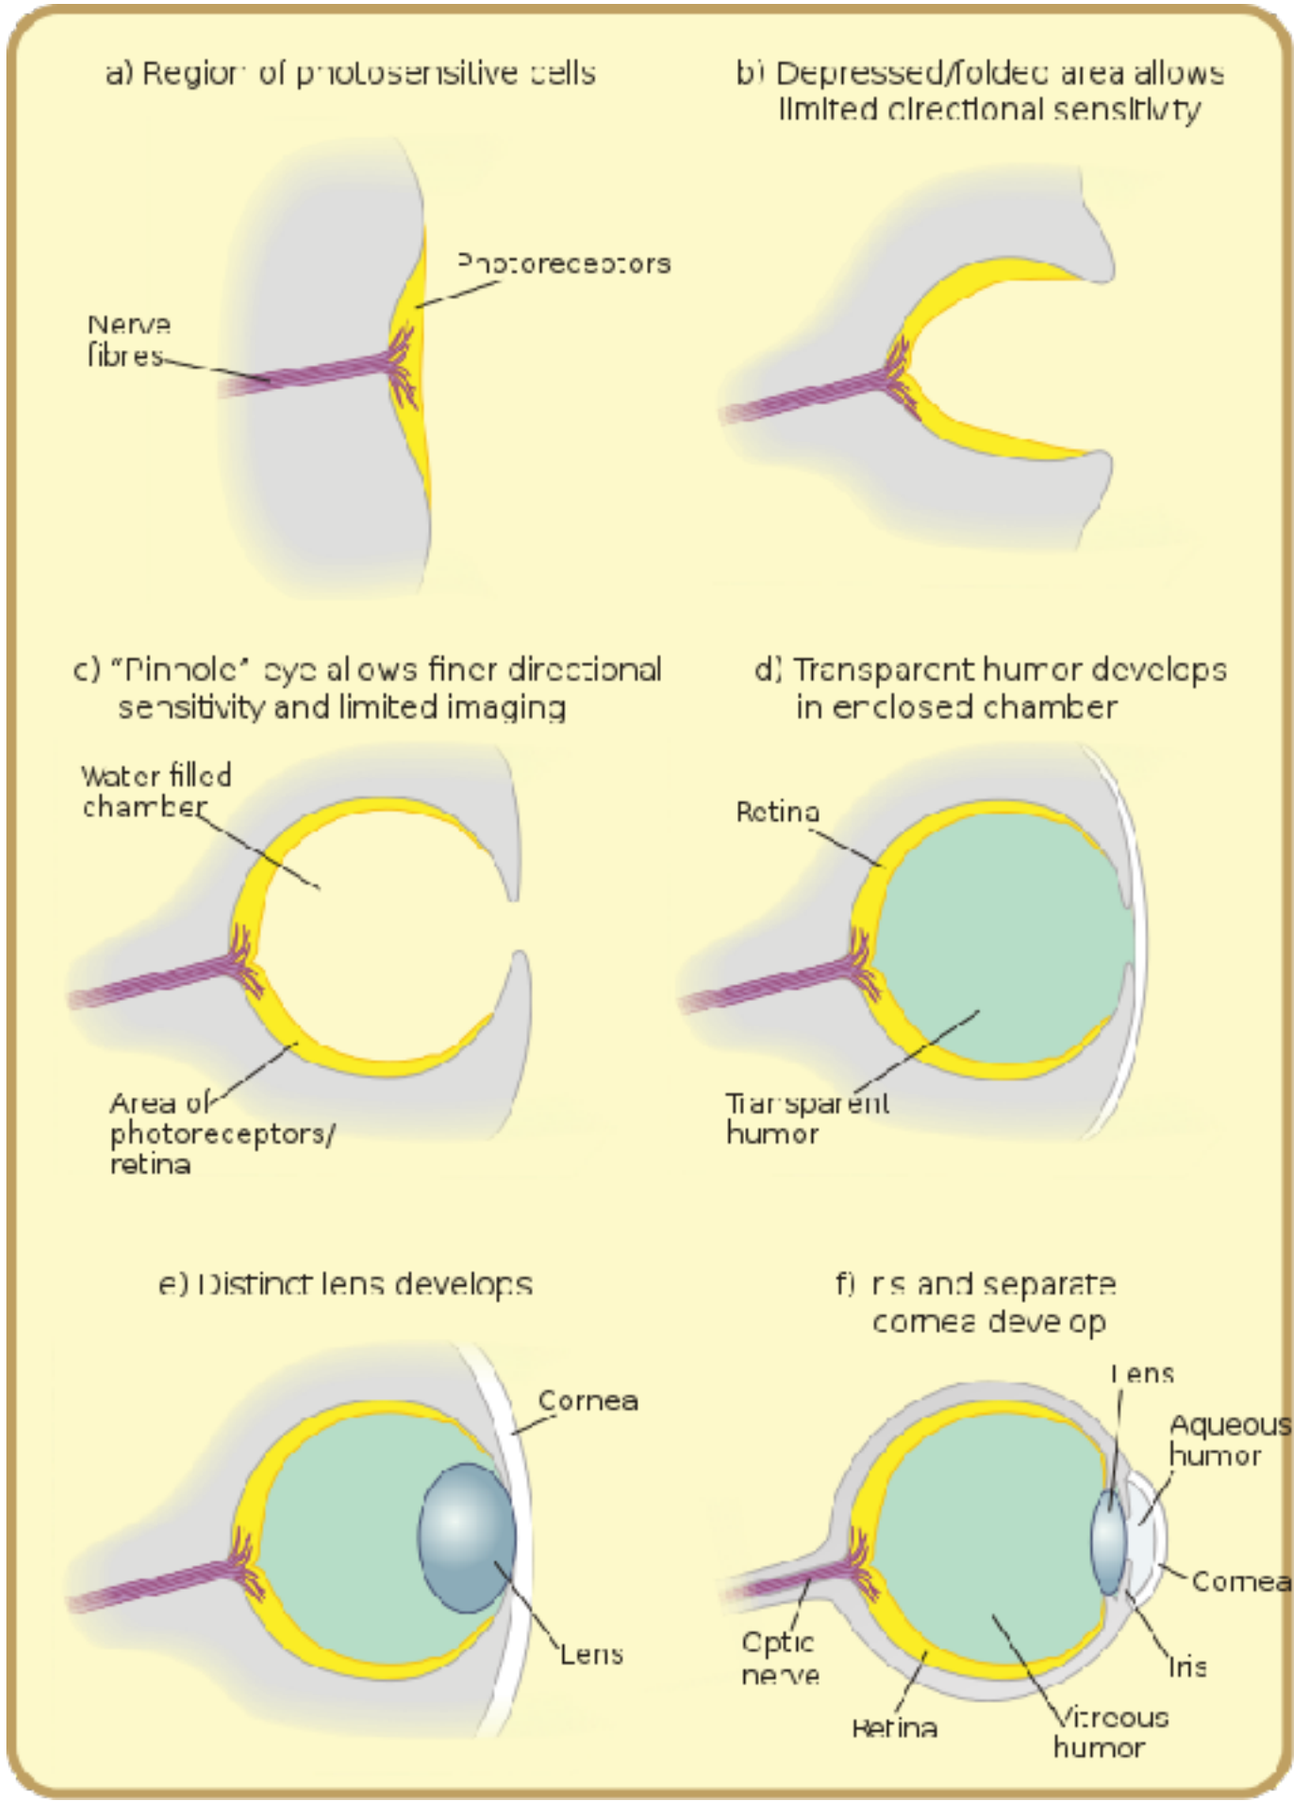

Evolution of Biological Eye

A quick experiment
Animals or Not?

Slide credit: Jia Deng


3aokg Al5ail
photographer ©


10636

Slide credit: Jia Deng

Thorpe, et al. *Nature*, 1996

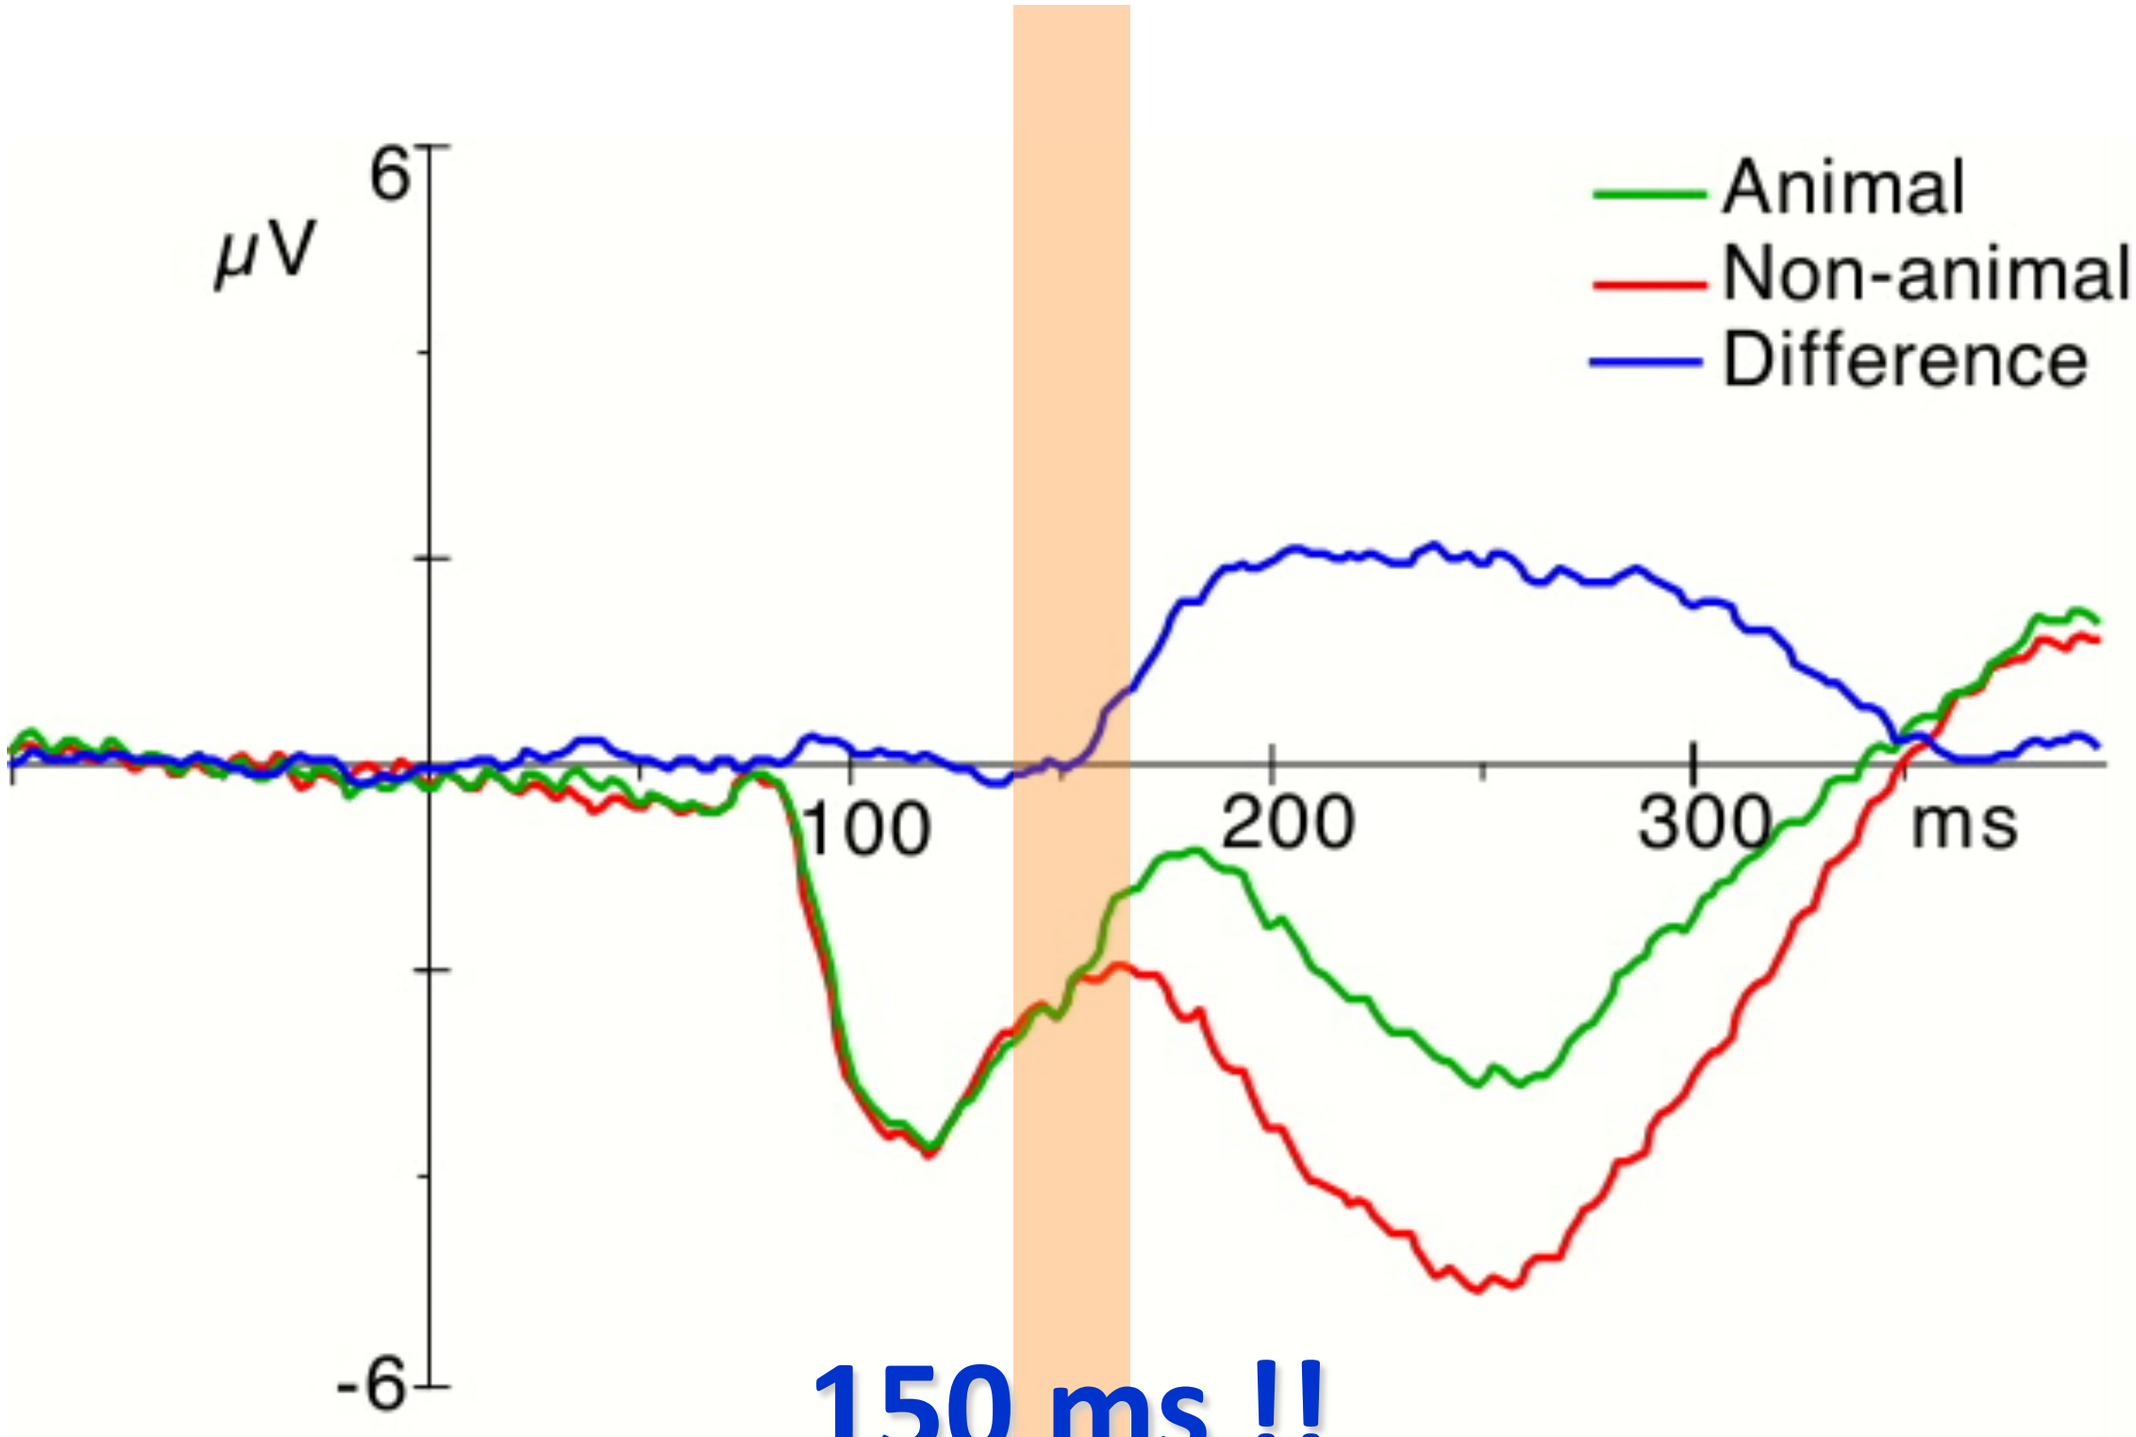

150 ms !!

Thorpe, et al. *Nature,* 1996

Why not build a brain?

About 1/3rd of the brain is devoted to visual processing

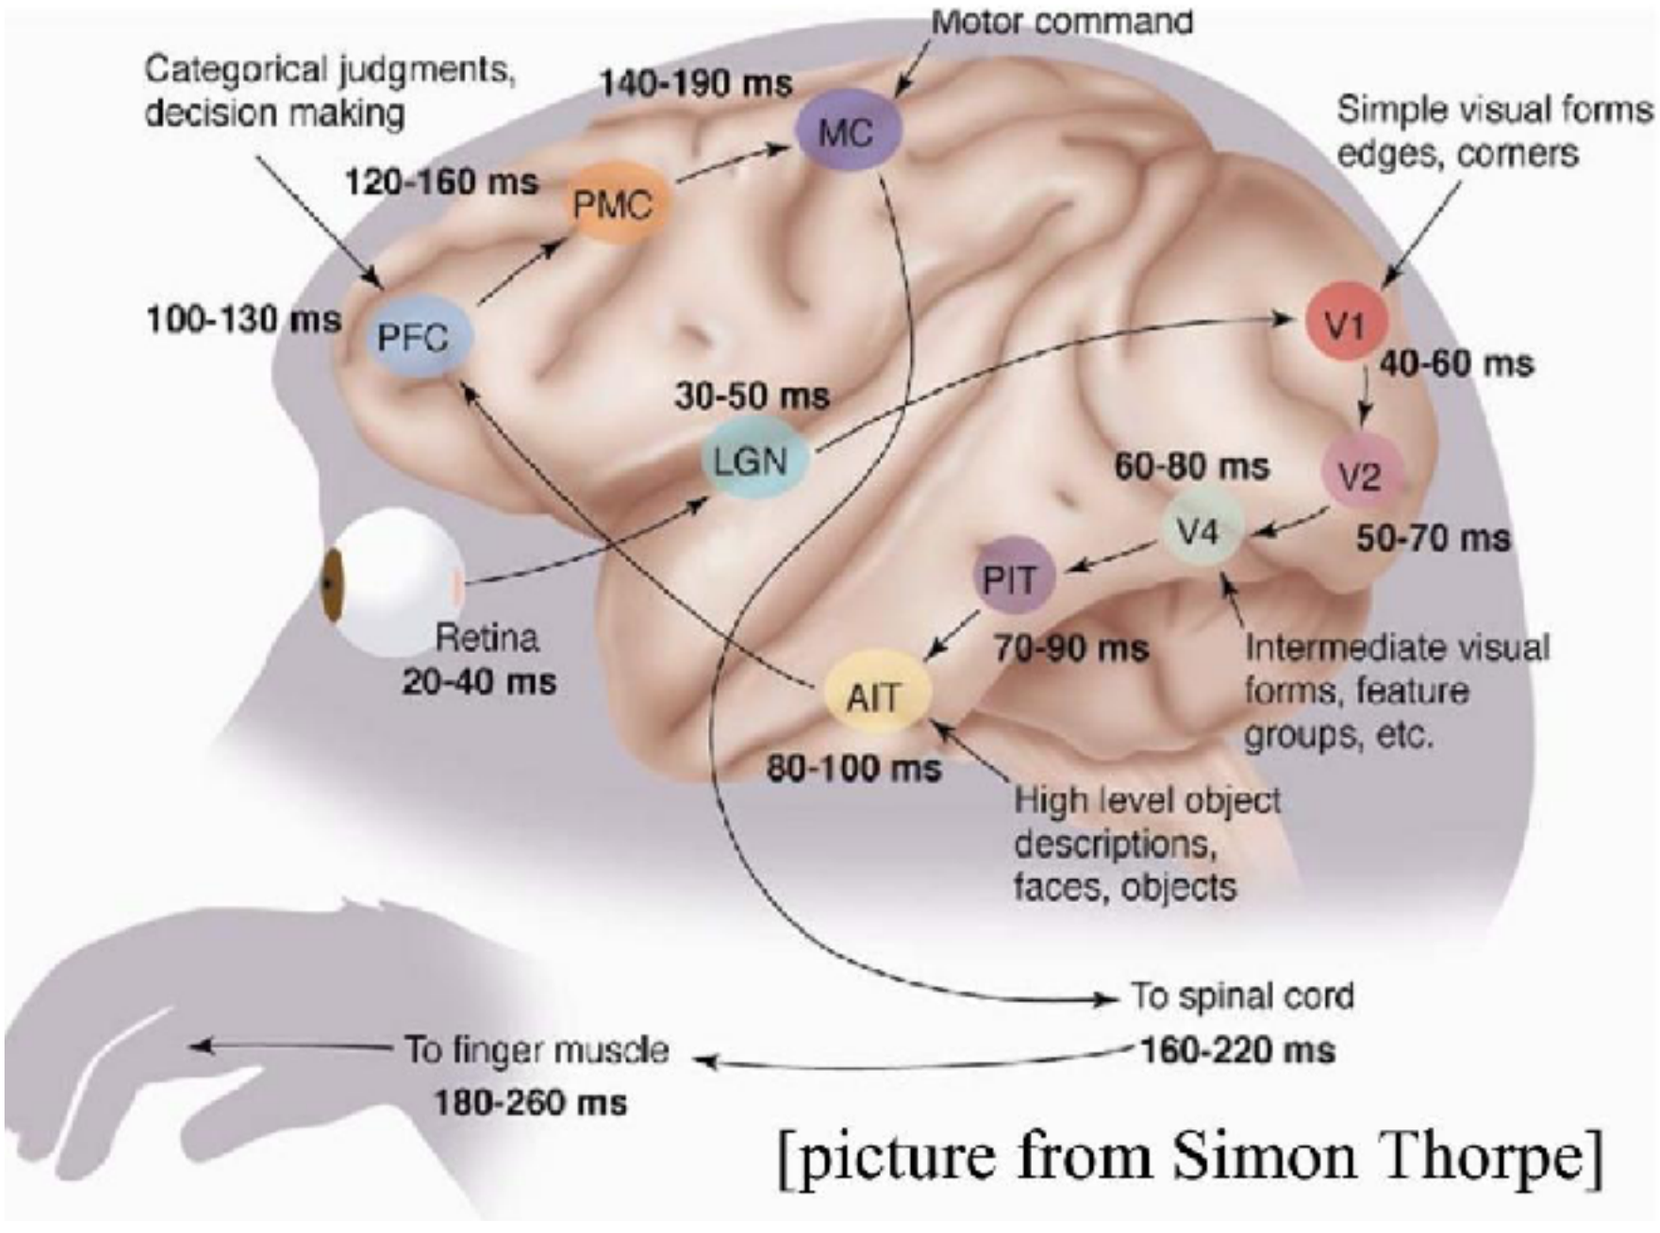

[picture from Simon Thorpe]

Do we have the hardware?

10^{11} parallel neurons

10^{8} serial transistors

We don’t know the software

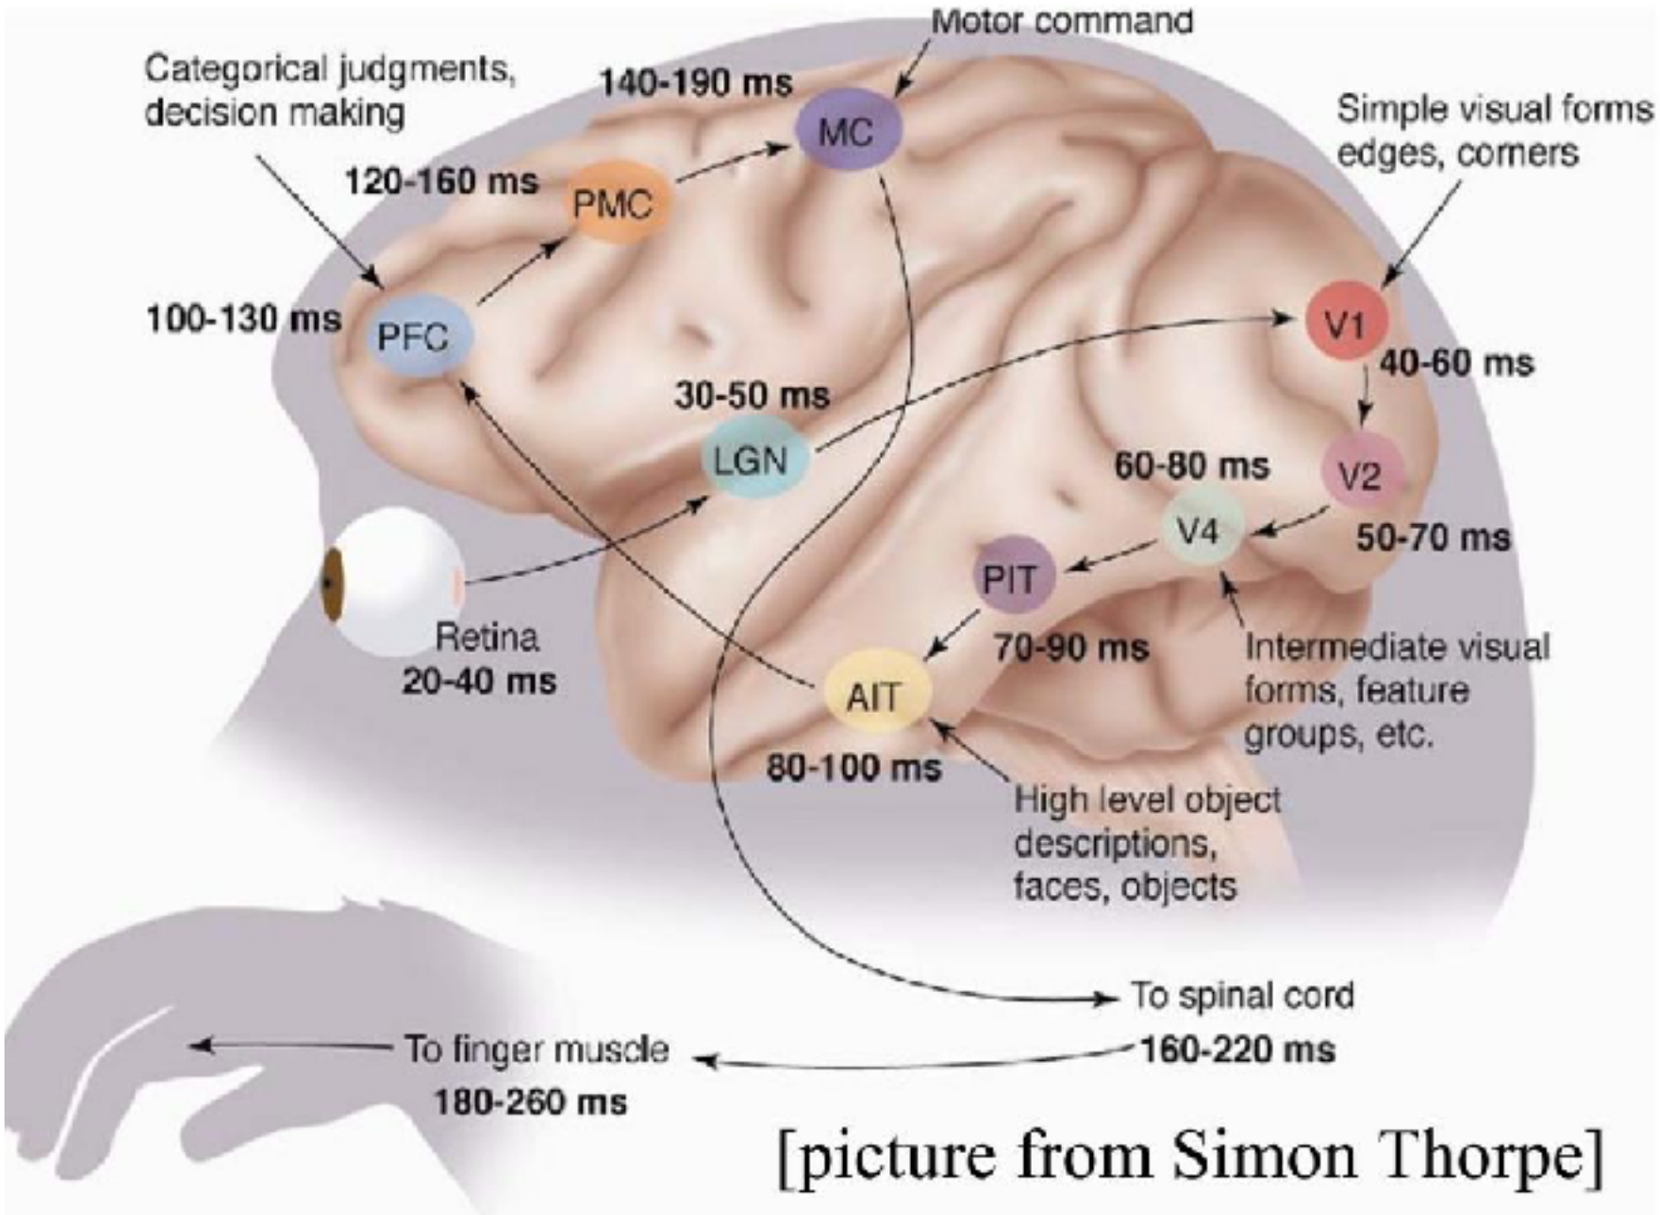

[picture from Simon Thorpe]

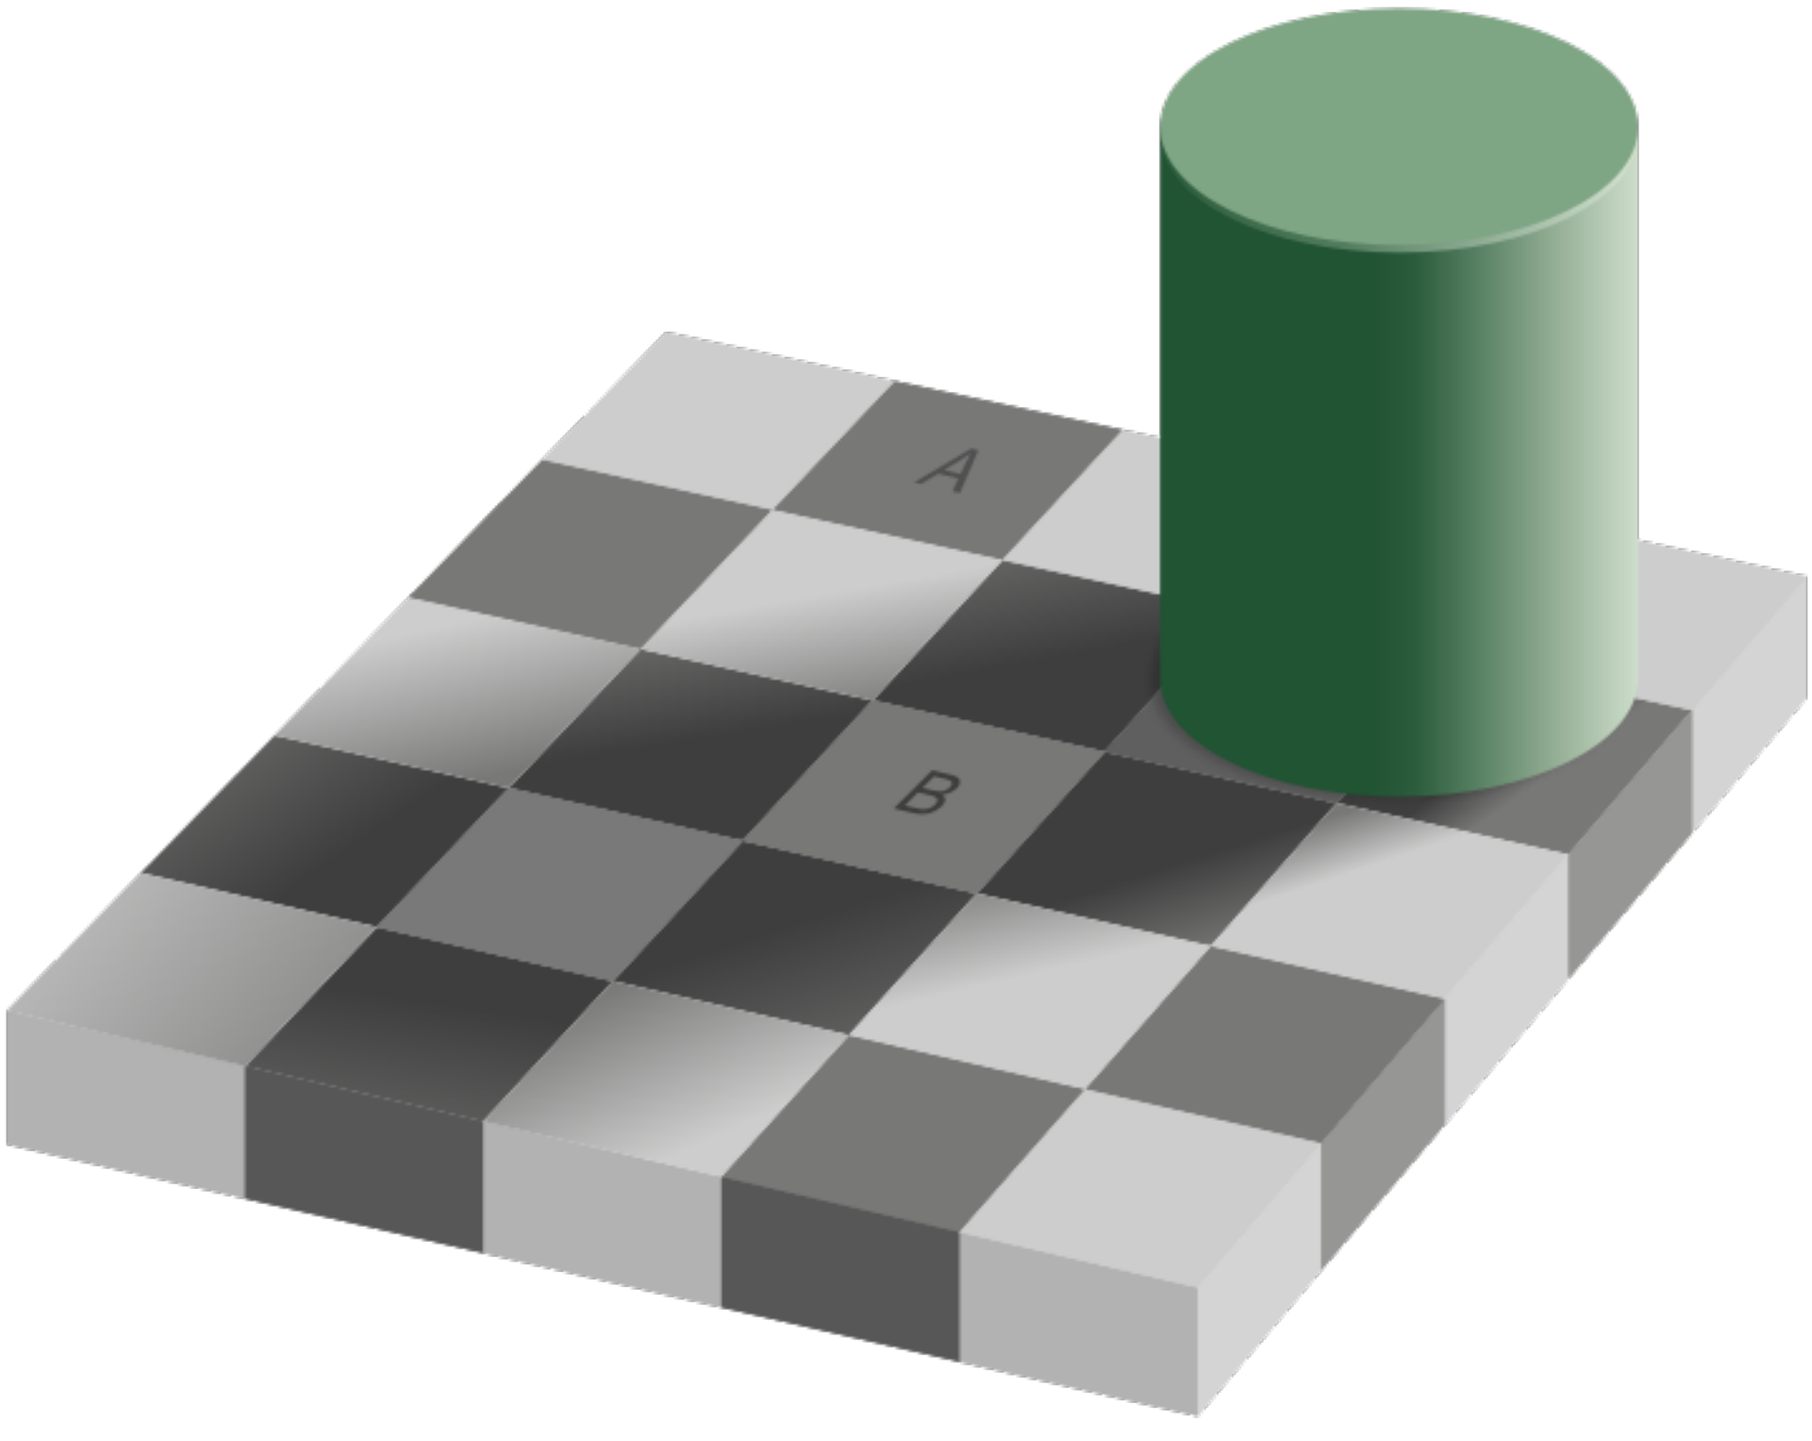
A
B

Adelson Illusion

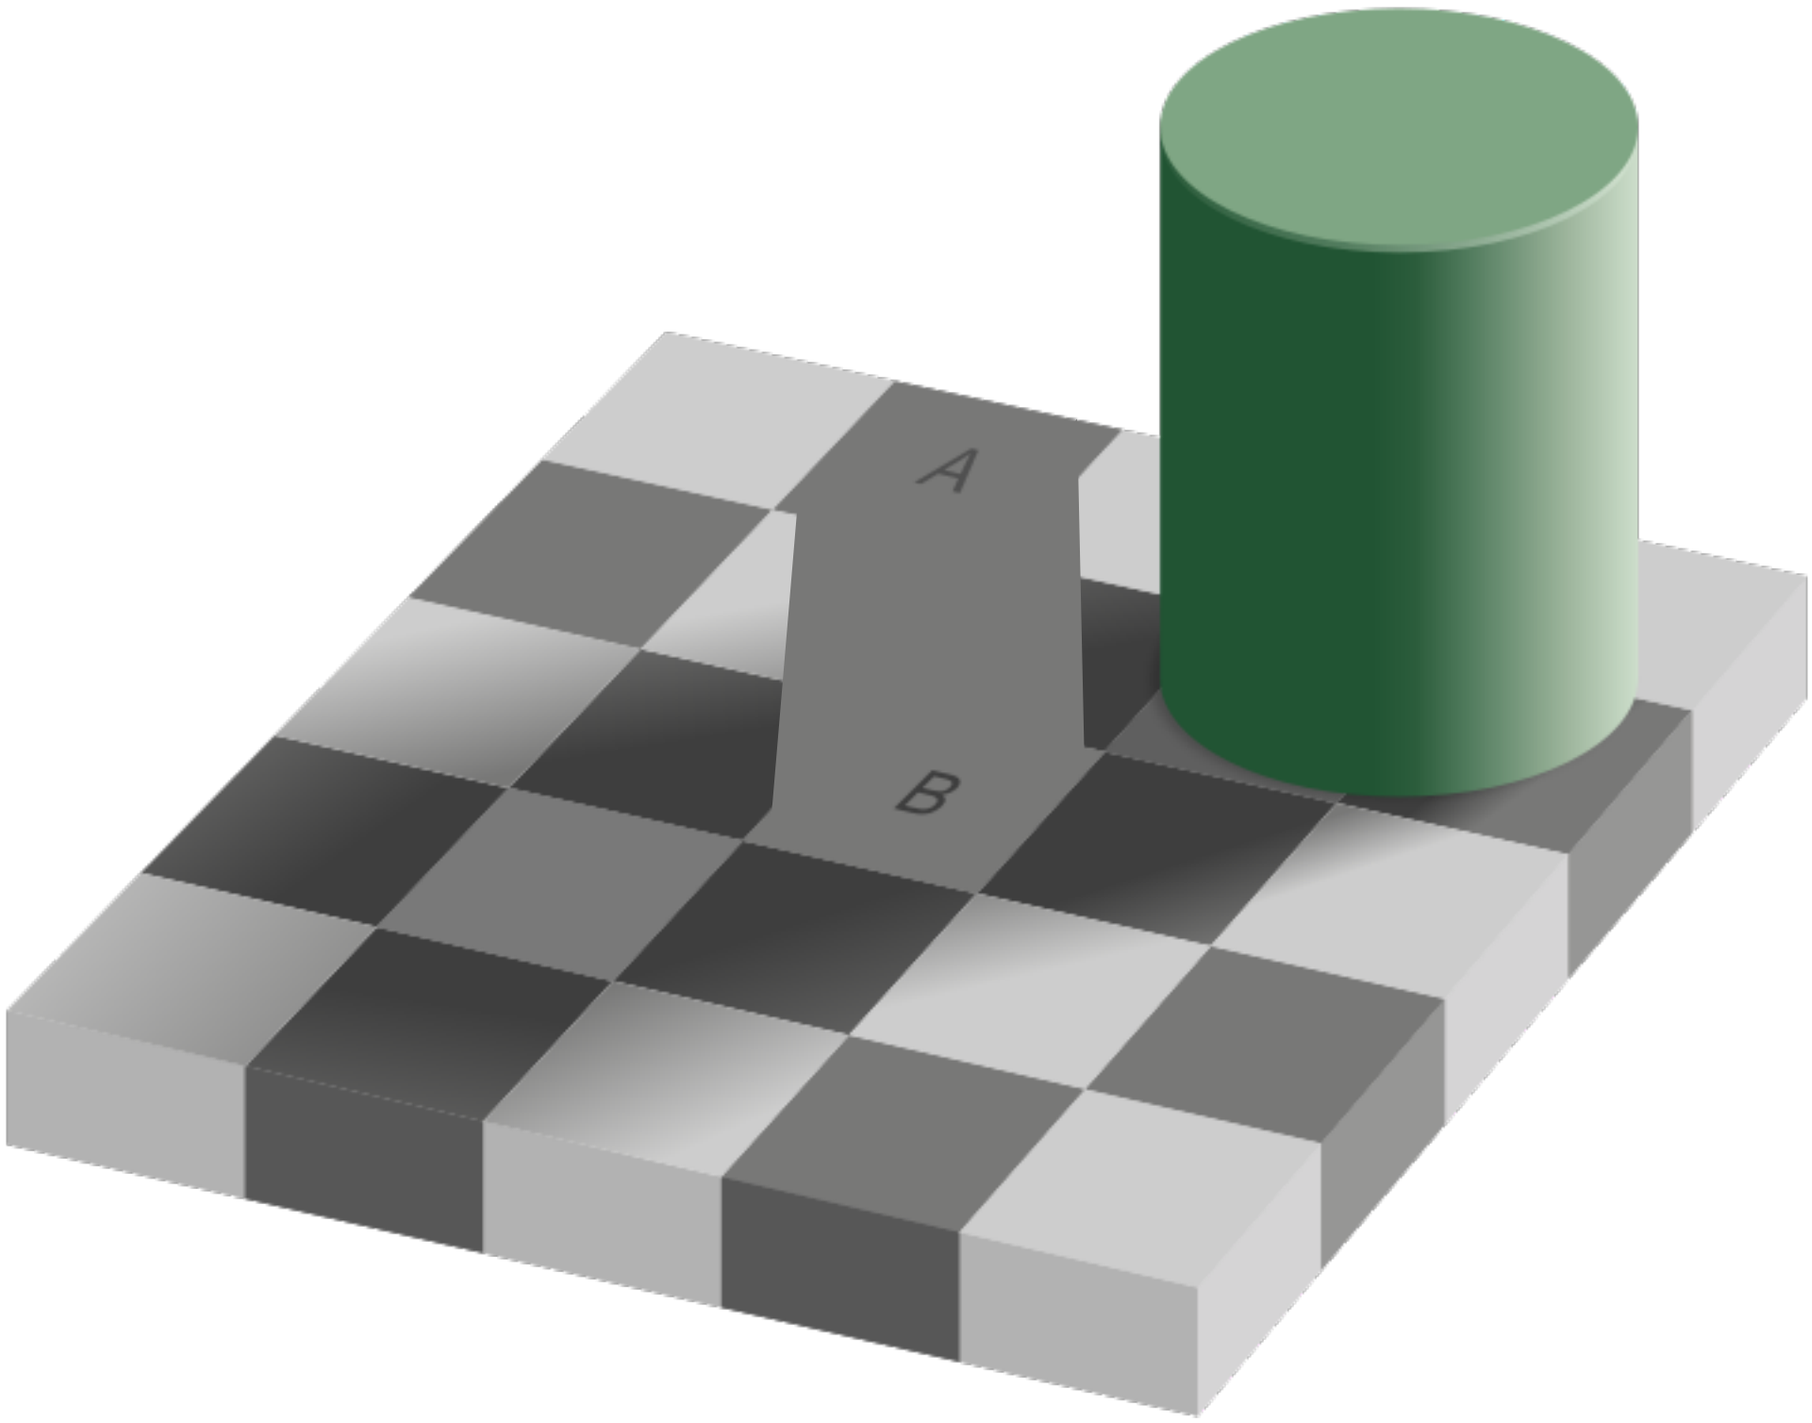

Illusionary Motion

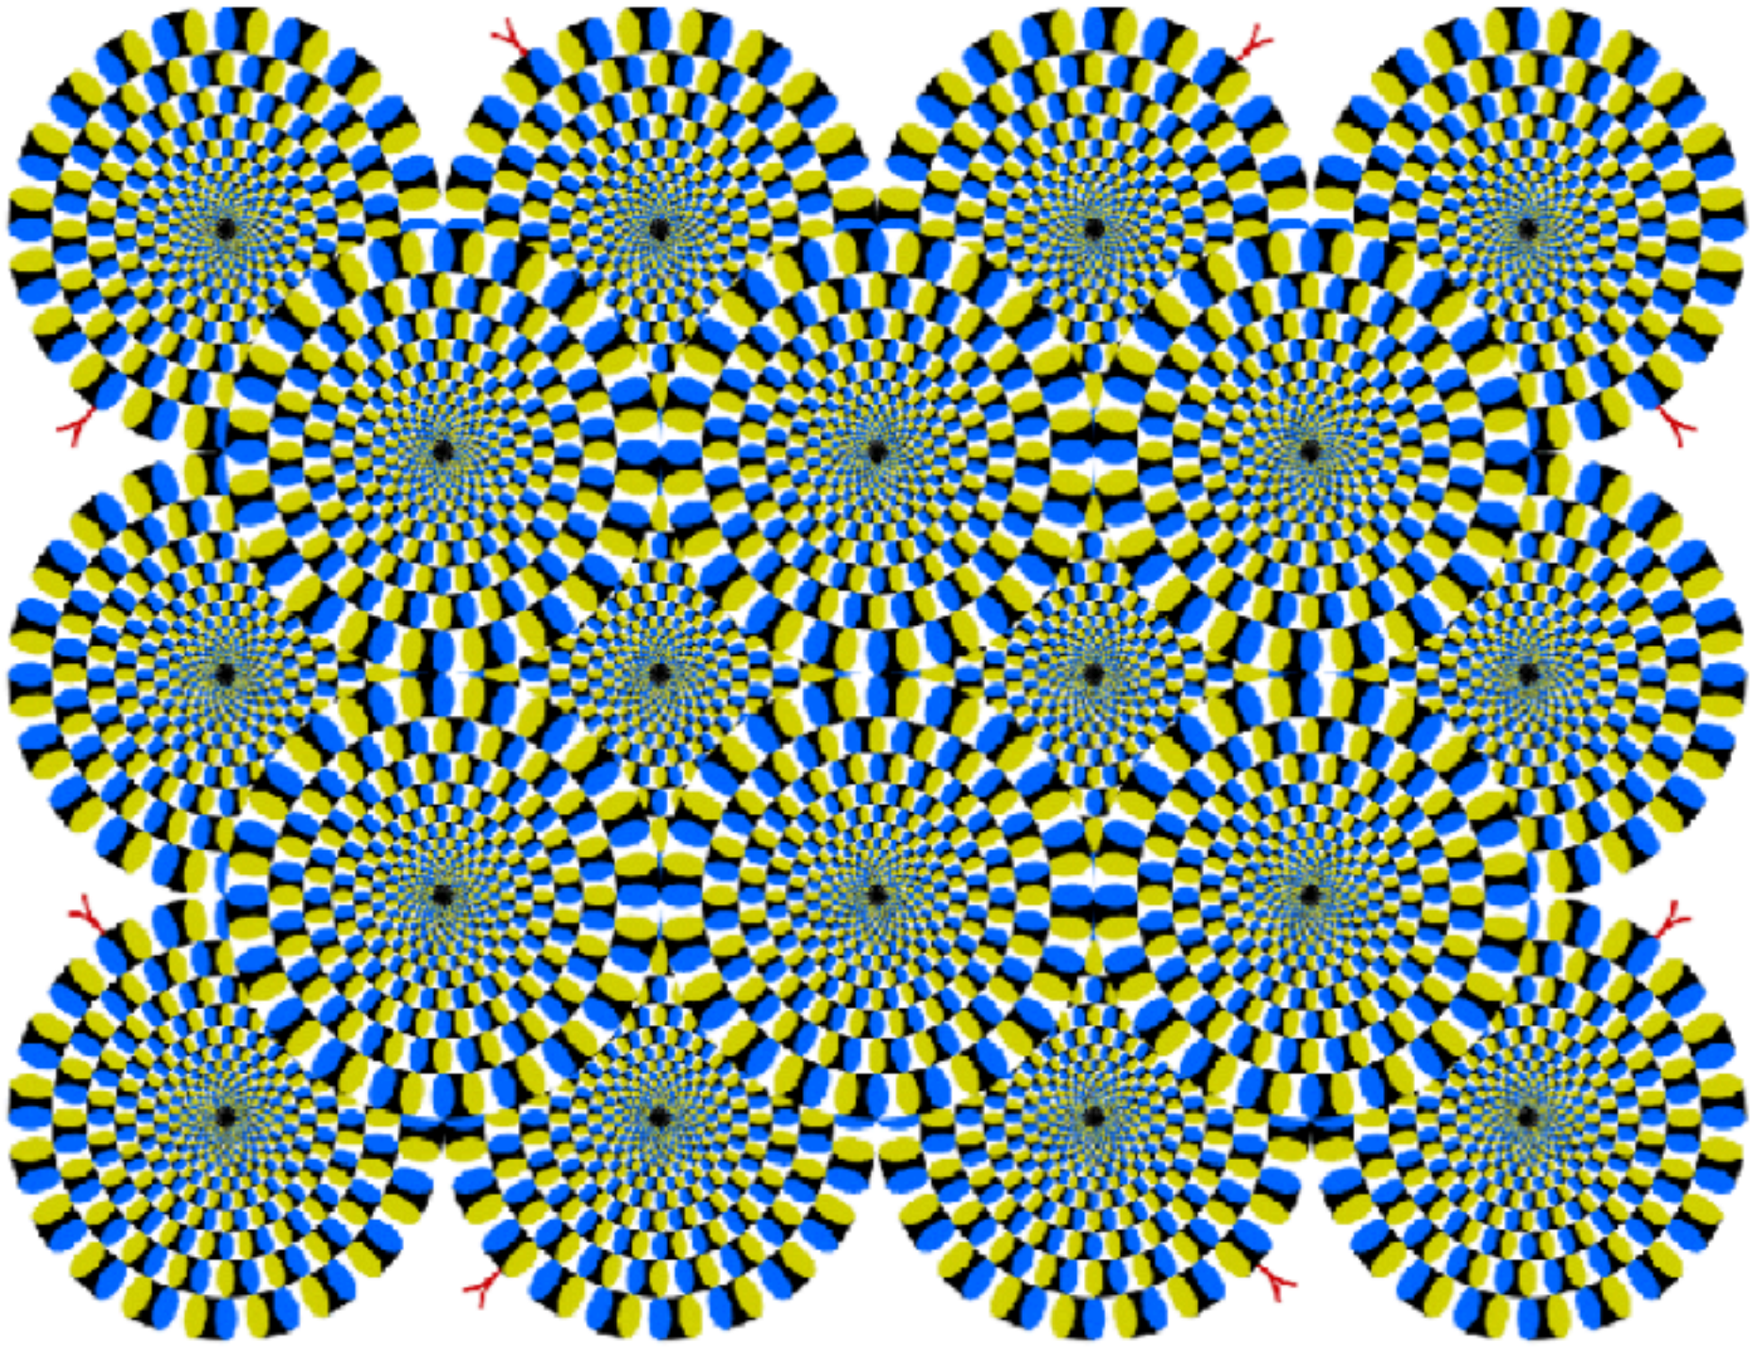

Scale Ambiguity

The Ames Room

The Ames Room

(Effect used in
Lord of the Rings)

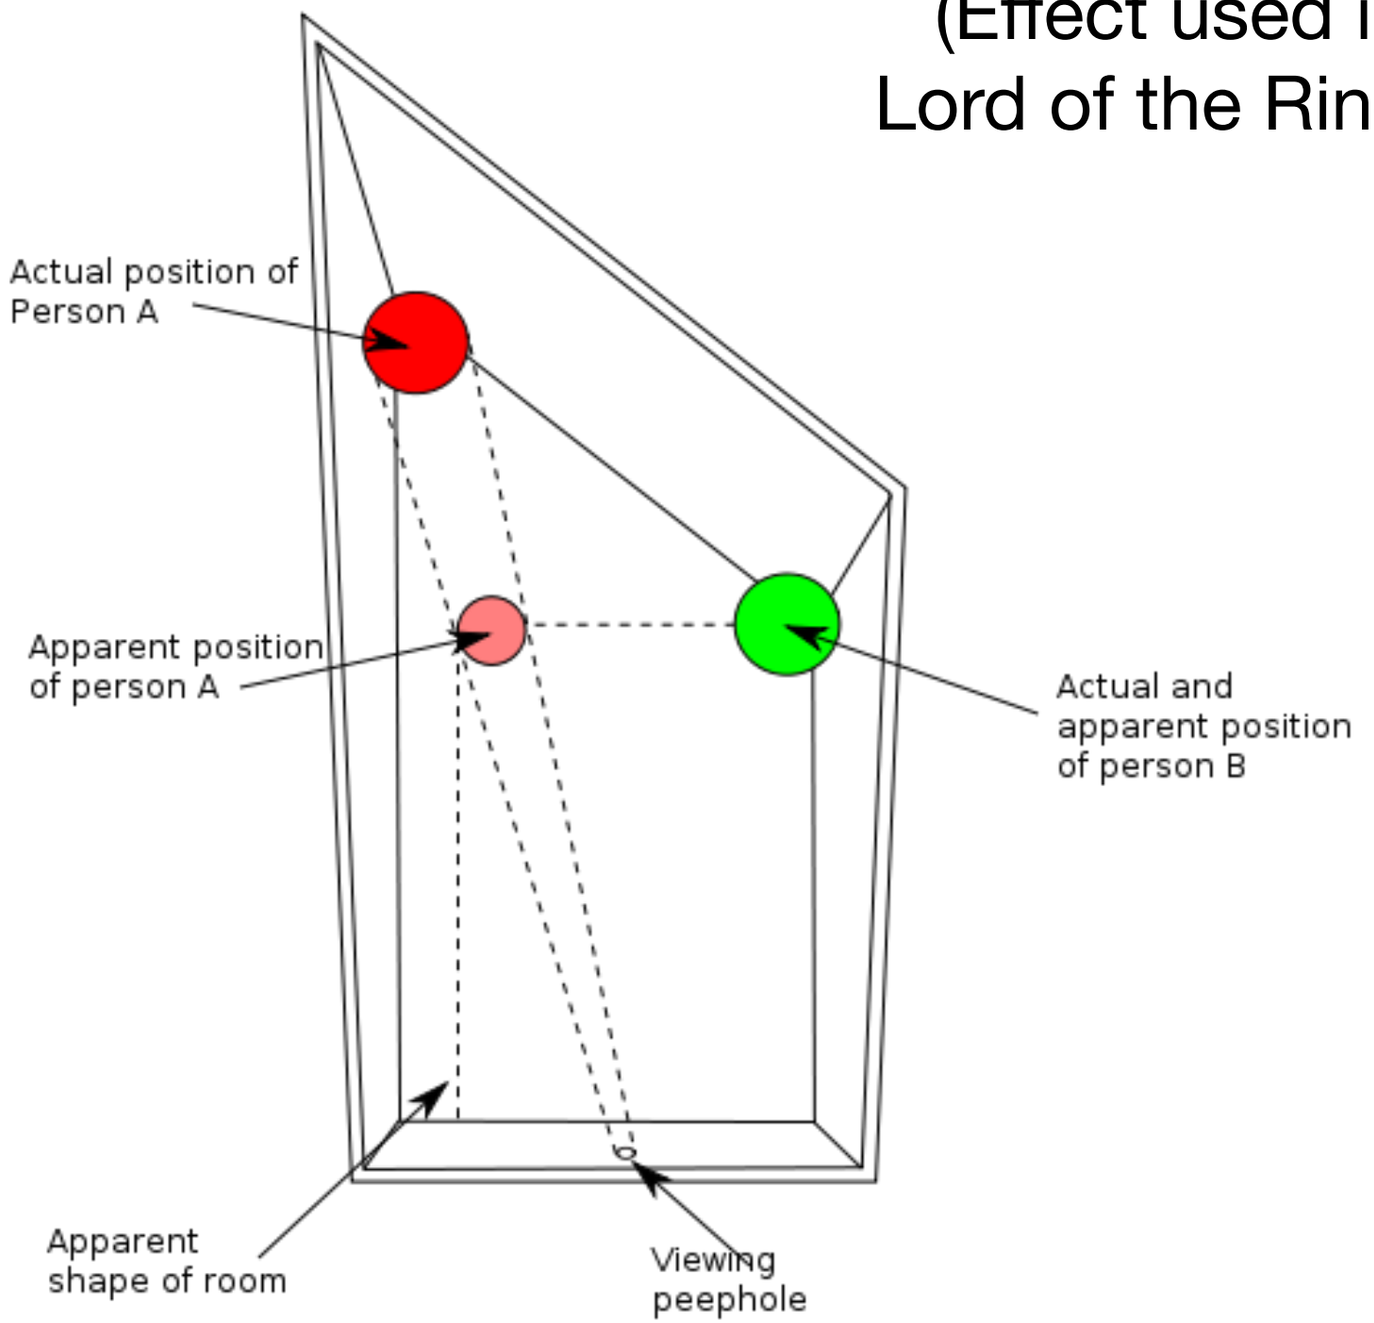

Heider-Simmel Illusion

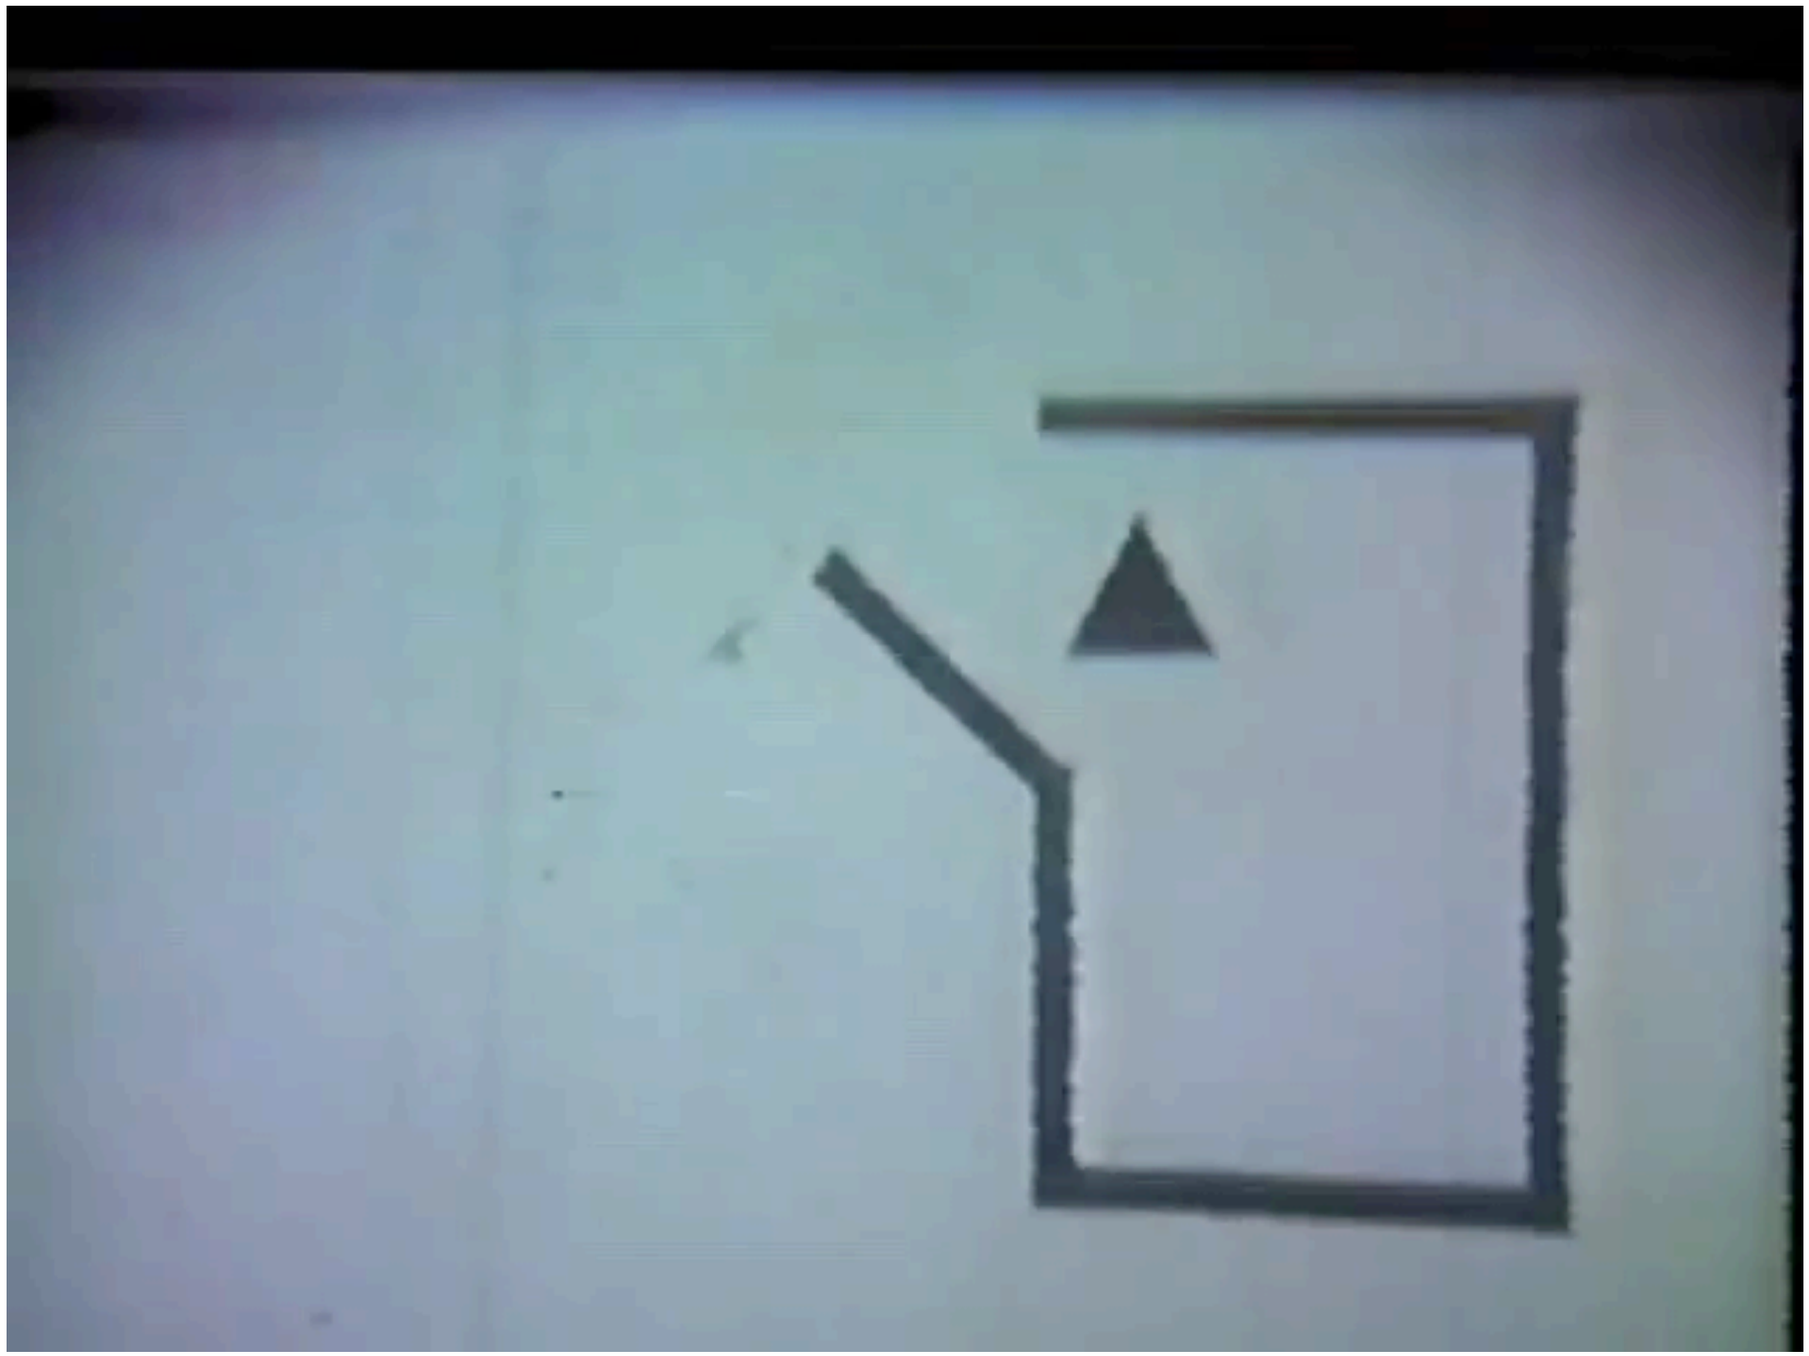

What objects are here?

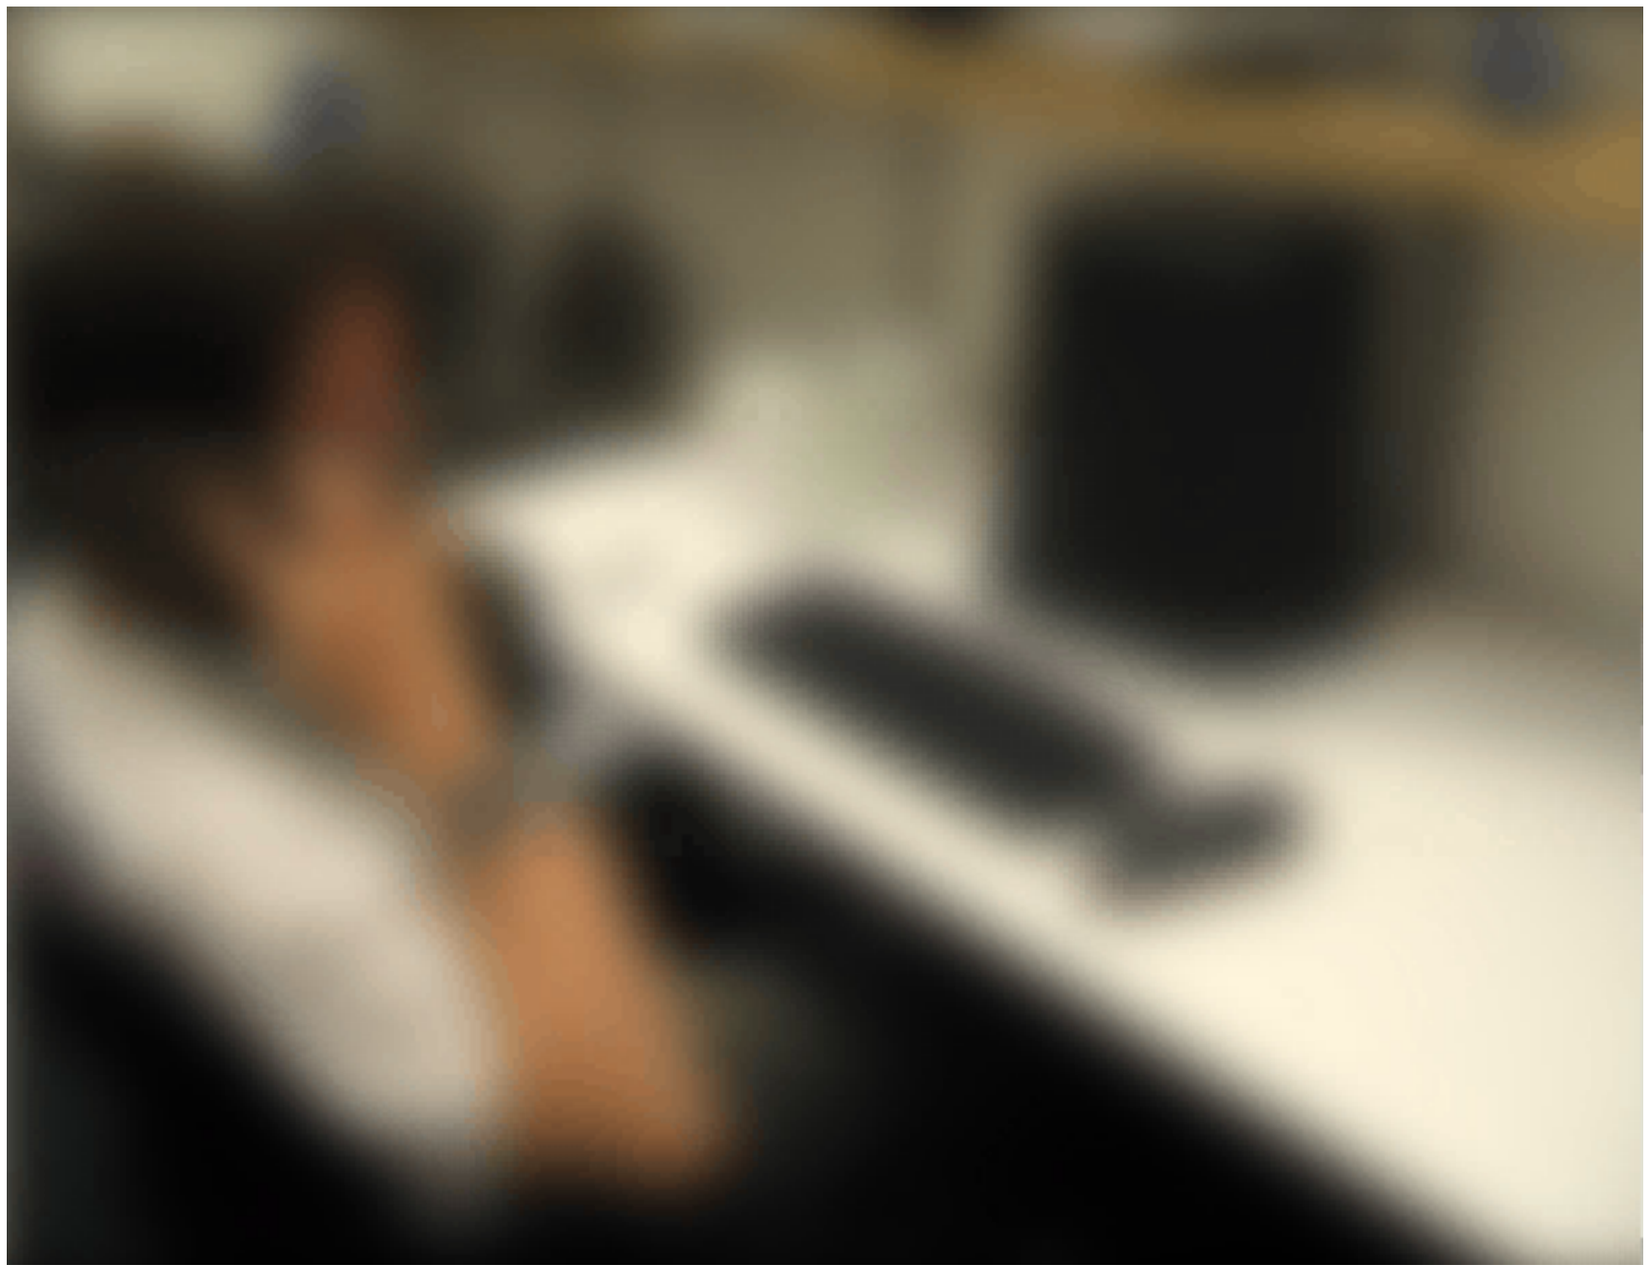

Slide credit: Rob Fergus and Antonio Torralba

Context

Slide credit: Rob Fergus and Antonio Torralba

Context

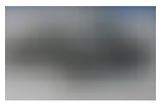

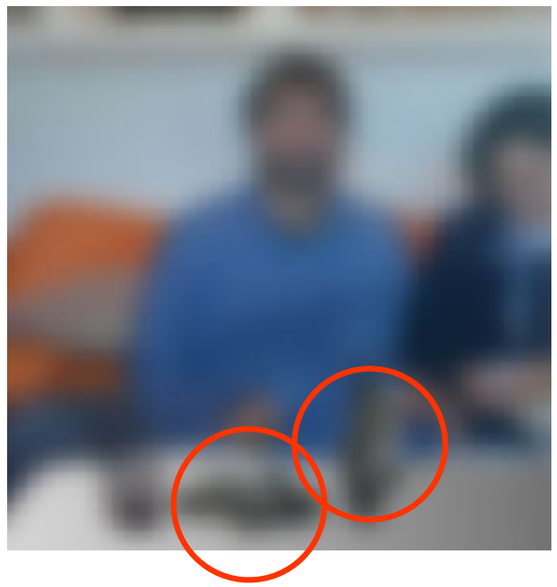

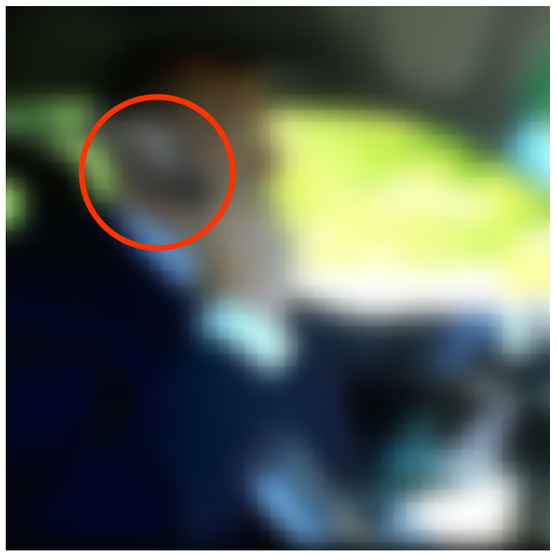

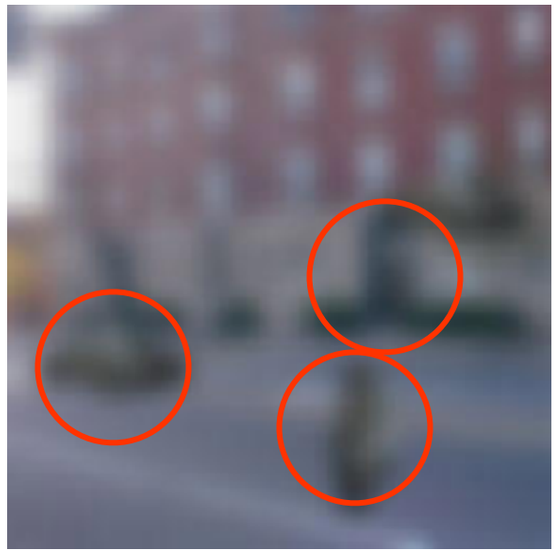

Slide credit: Fei-Fei Li, Rob Fergus and Antonio Torralba

Tool 1:
Physics and Geometry

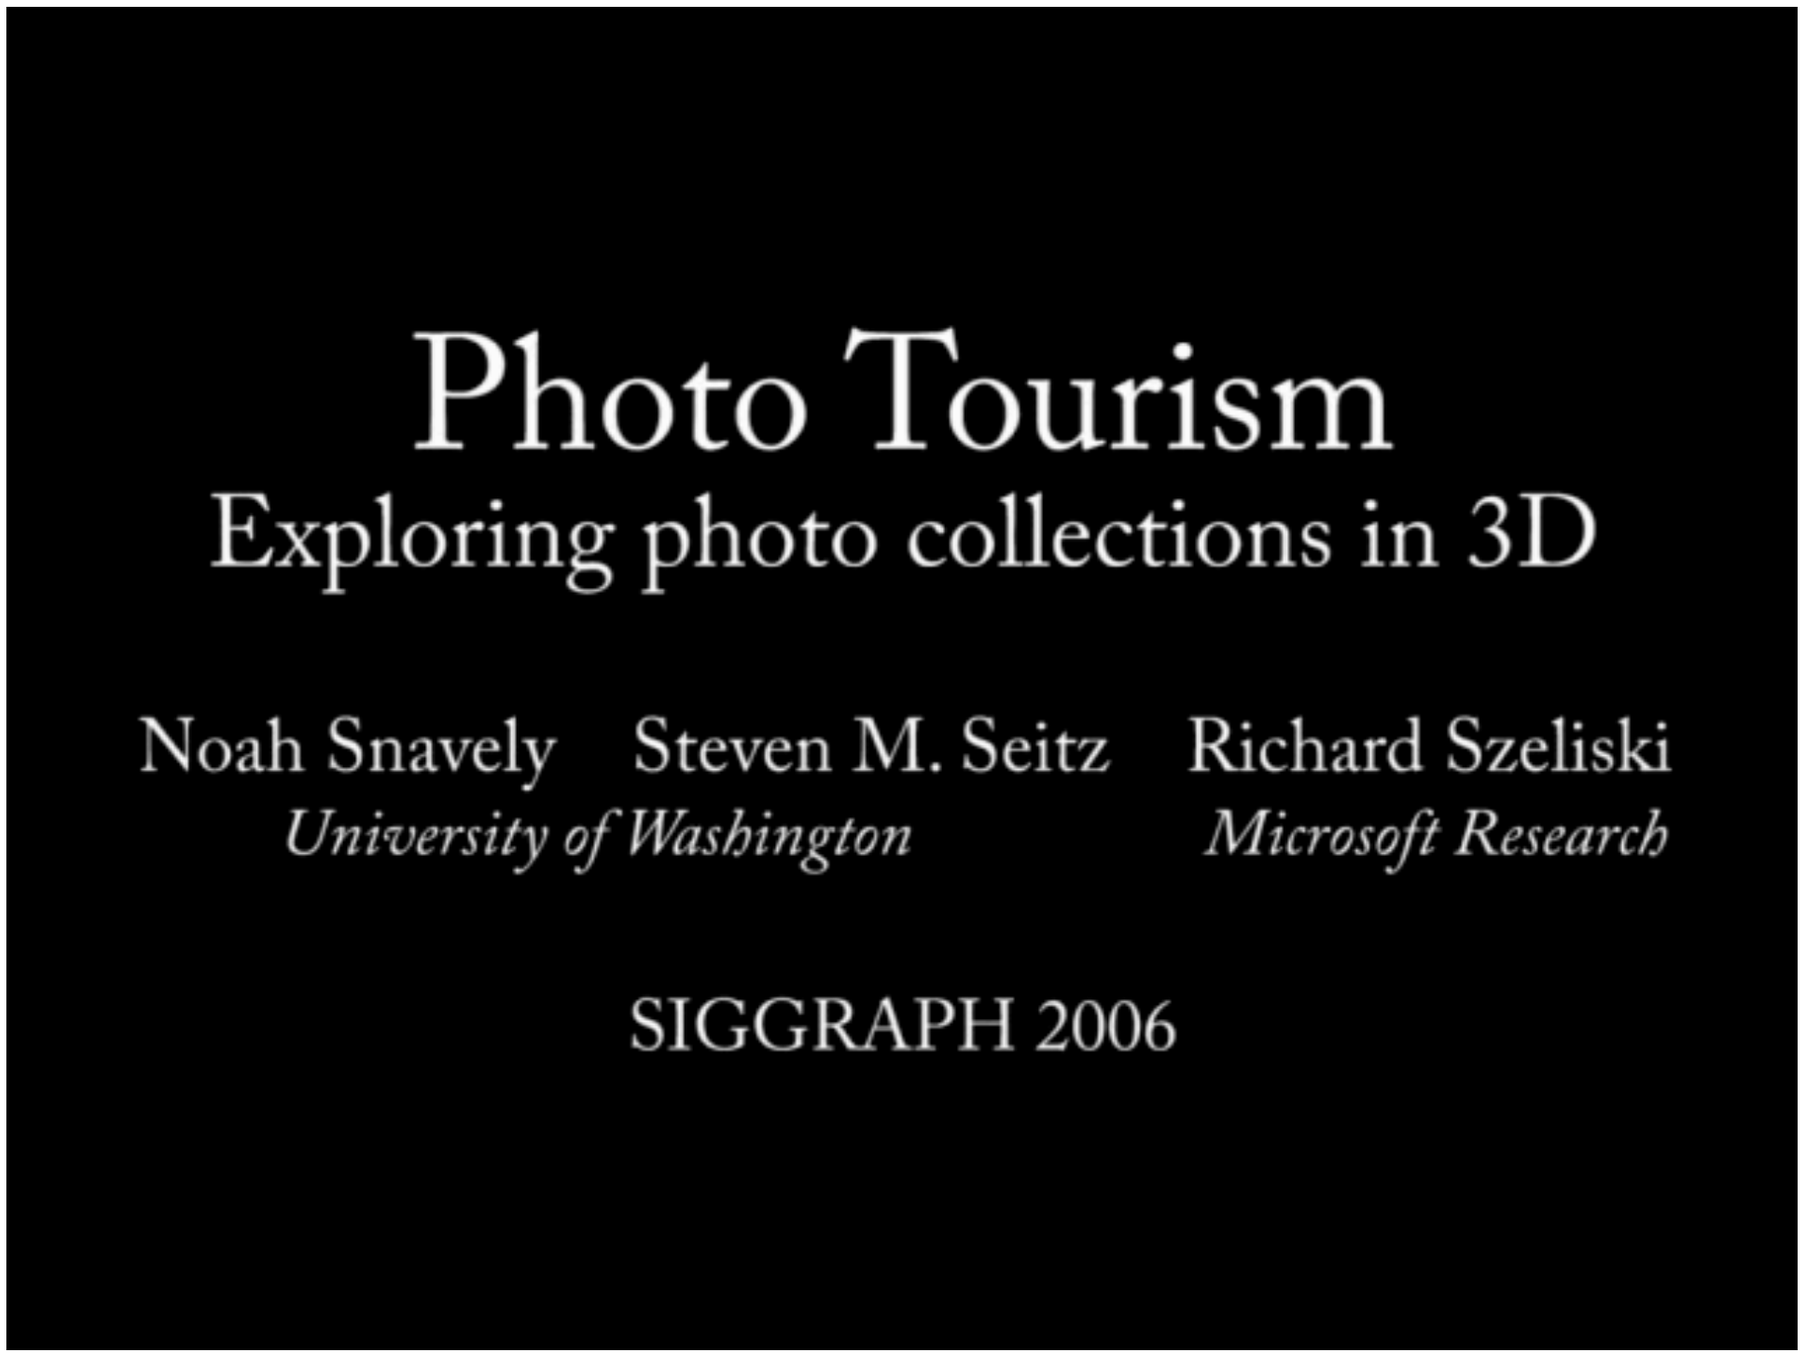

Tool 2:
Data and Learning

Two Extremes of Vision

Extrapolation problem
Generalization
Diagnostic features

Interpolation problem
Correspondence
Finding the differences

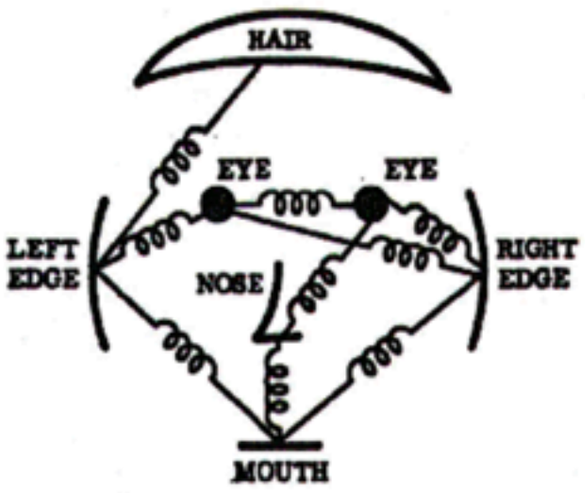

Slide credit: Aude Oliva

Evolution of Vision Datasets

Created here in 1996

COIL-20

Caltech 101

MNIST (1998)

IMAGENET (2009)

2 year old kid

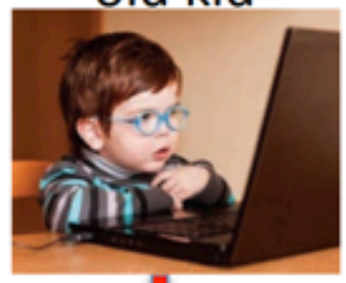

10^3 10^4 10^5 10^6 10^7 10^8 10^9 # images

Caltech-4 (2003)

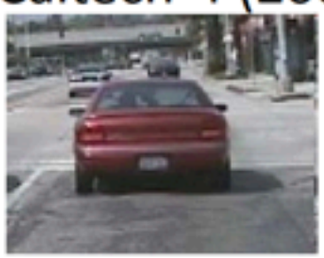

PASCAL (2005)

Slide credit: Aude Oliva

Course Information

Computer Vision
Fall 2018
Columbia University

About Me


SK
QC
ON
NB
WA
ME
MT
ND
VT
NH
OR
MN
ID
SD
WI
MI
NY
MA
CT
WY
PA
NJ
NE
IA
IL
IN
OH
MD
DE
NV
UT
United States
CA
CO
KS
MO
WV
VA
KY
NC
TN
AZ
OK
AR
SC
NM
MS
AL
GA
TX
LA
FL
UC Irvine
Gulf of
Mexico

About Me

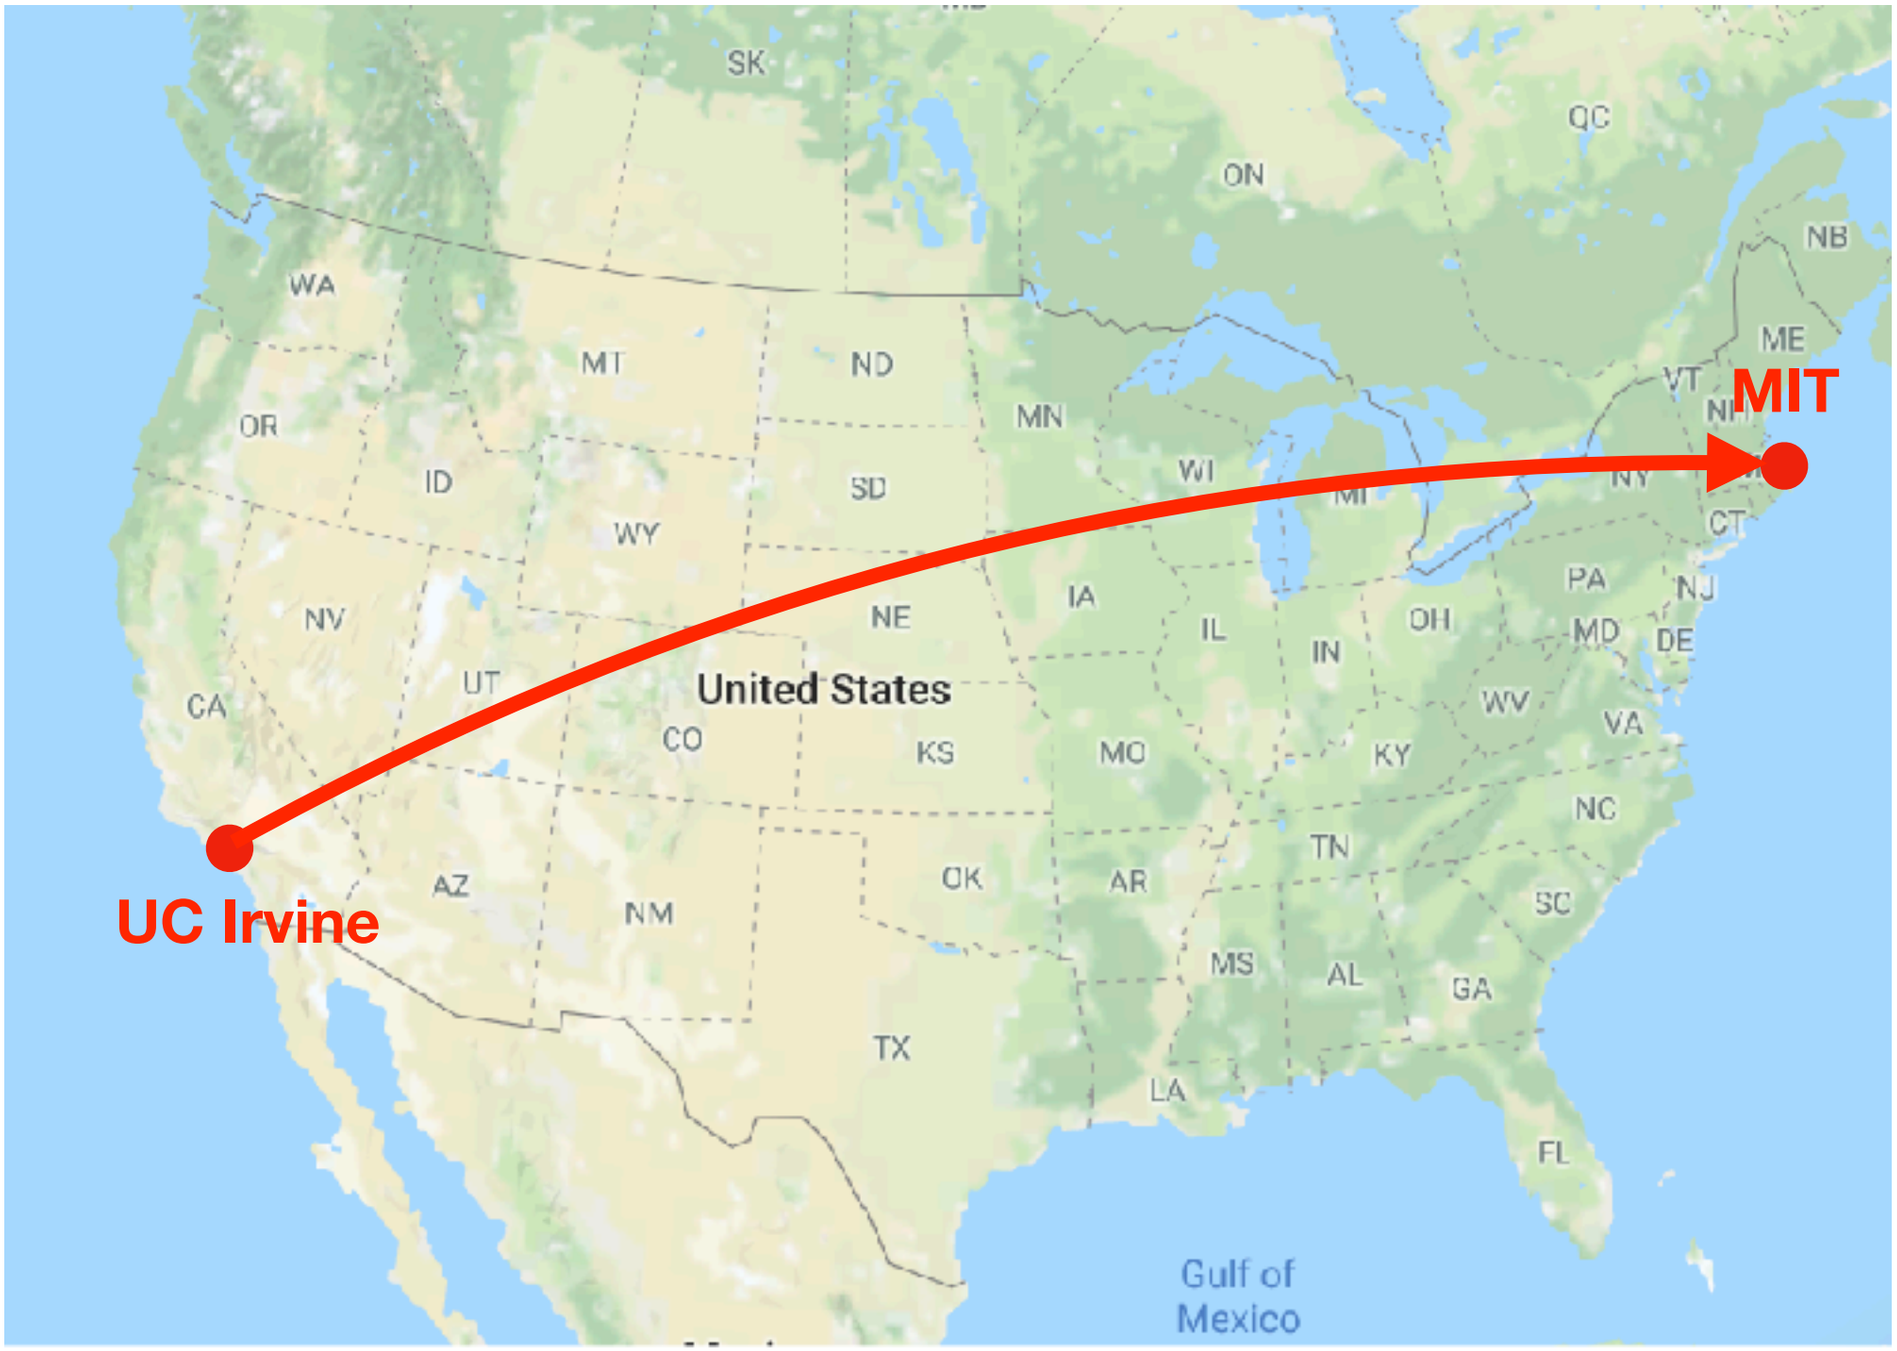
MIT
UC Irvine
United States
Gulf of Mexico
SK
ON
QC
NB
WA
MT
ND
MN
OR
ID
SD
WI
MI
NY
VT
ME
NH
CT
WY
PA
NJ
NV
NE
IA
IL
IN
OH
MD
DE
UT
CA
CO
KS
MO
KY
WV
VA
NC
TN
AZ
NM
OK
AR
SC
MS
AL
GA
TX
LA
FL

About Me

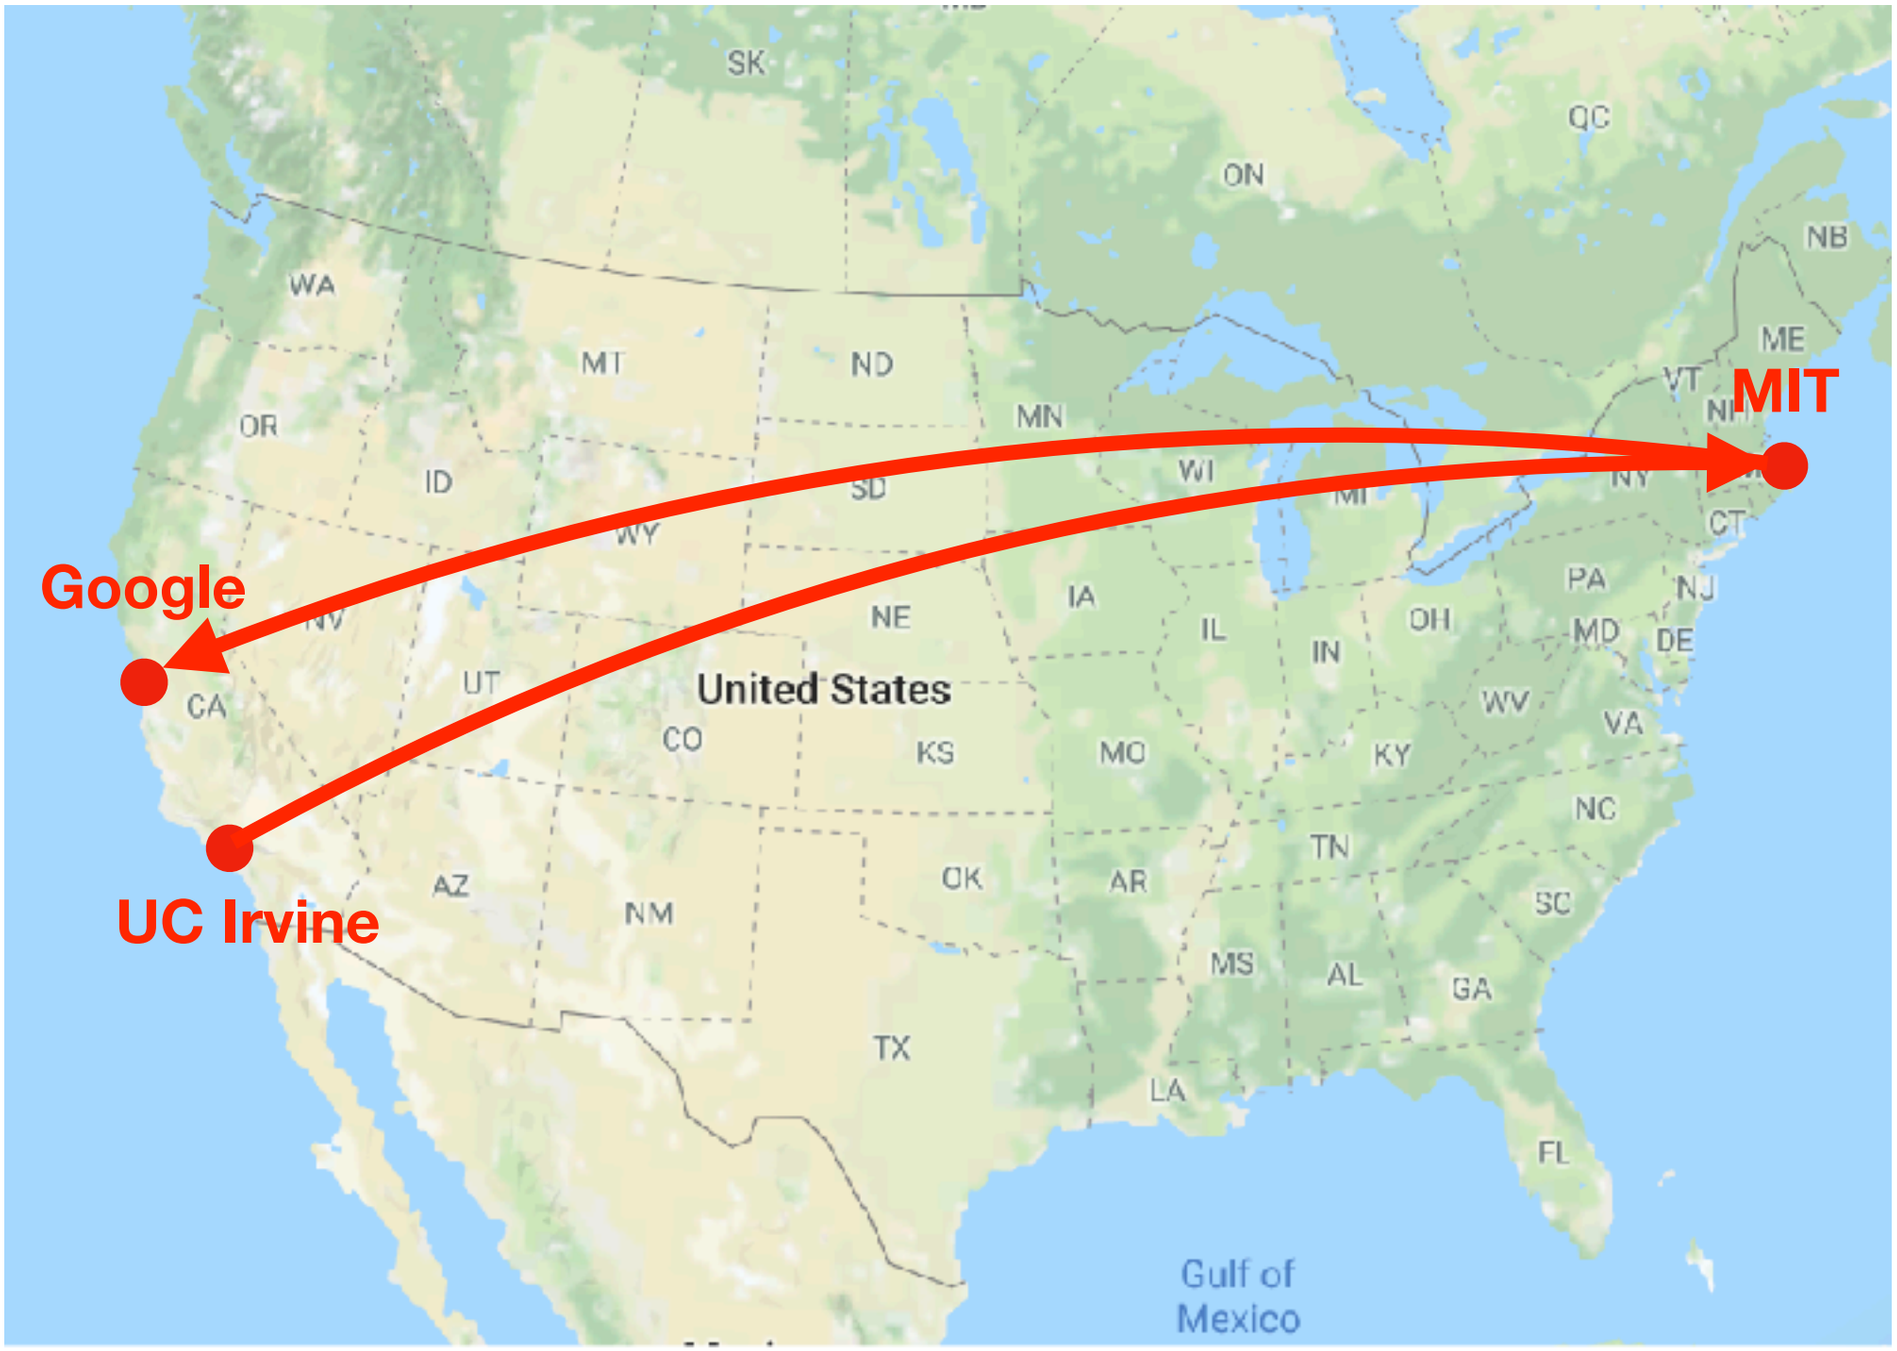
MIT
Google
UC Irvine
United States
Gulf of Mexico
SK
ON
QC
NB
WA
MT
ND
MN
OR
ID
SD
WI
MI
NY
VT
ME
NH
CT
WY
PA
NJ
NE
IA
IL
IN
OH
MD
DE
NV
UT
CA
CO
KS
MO
WV
VA
KY
NC
TN
AZ
NM
OK
AR
SC
MS
AL
GA
TX
LA
FL

About Me

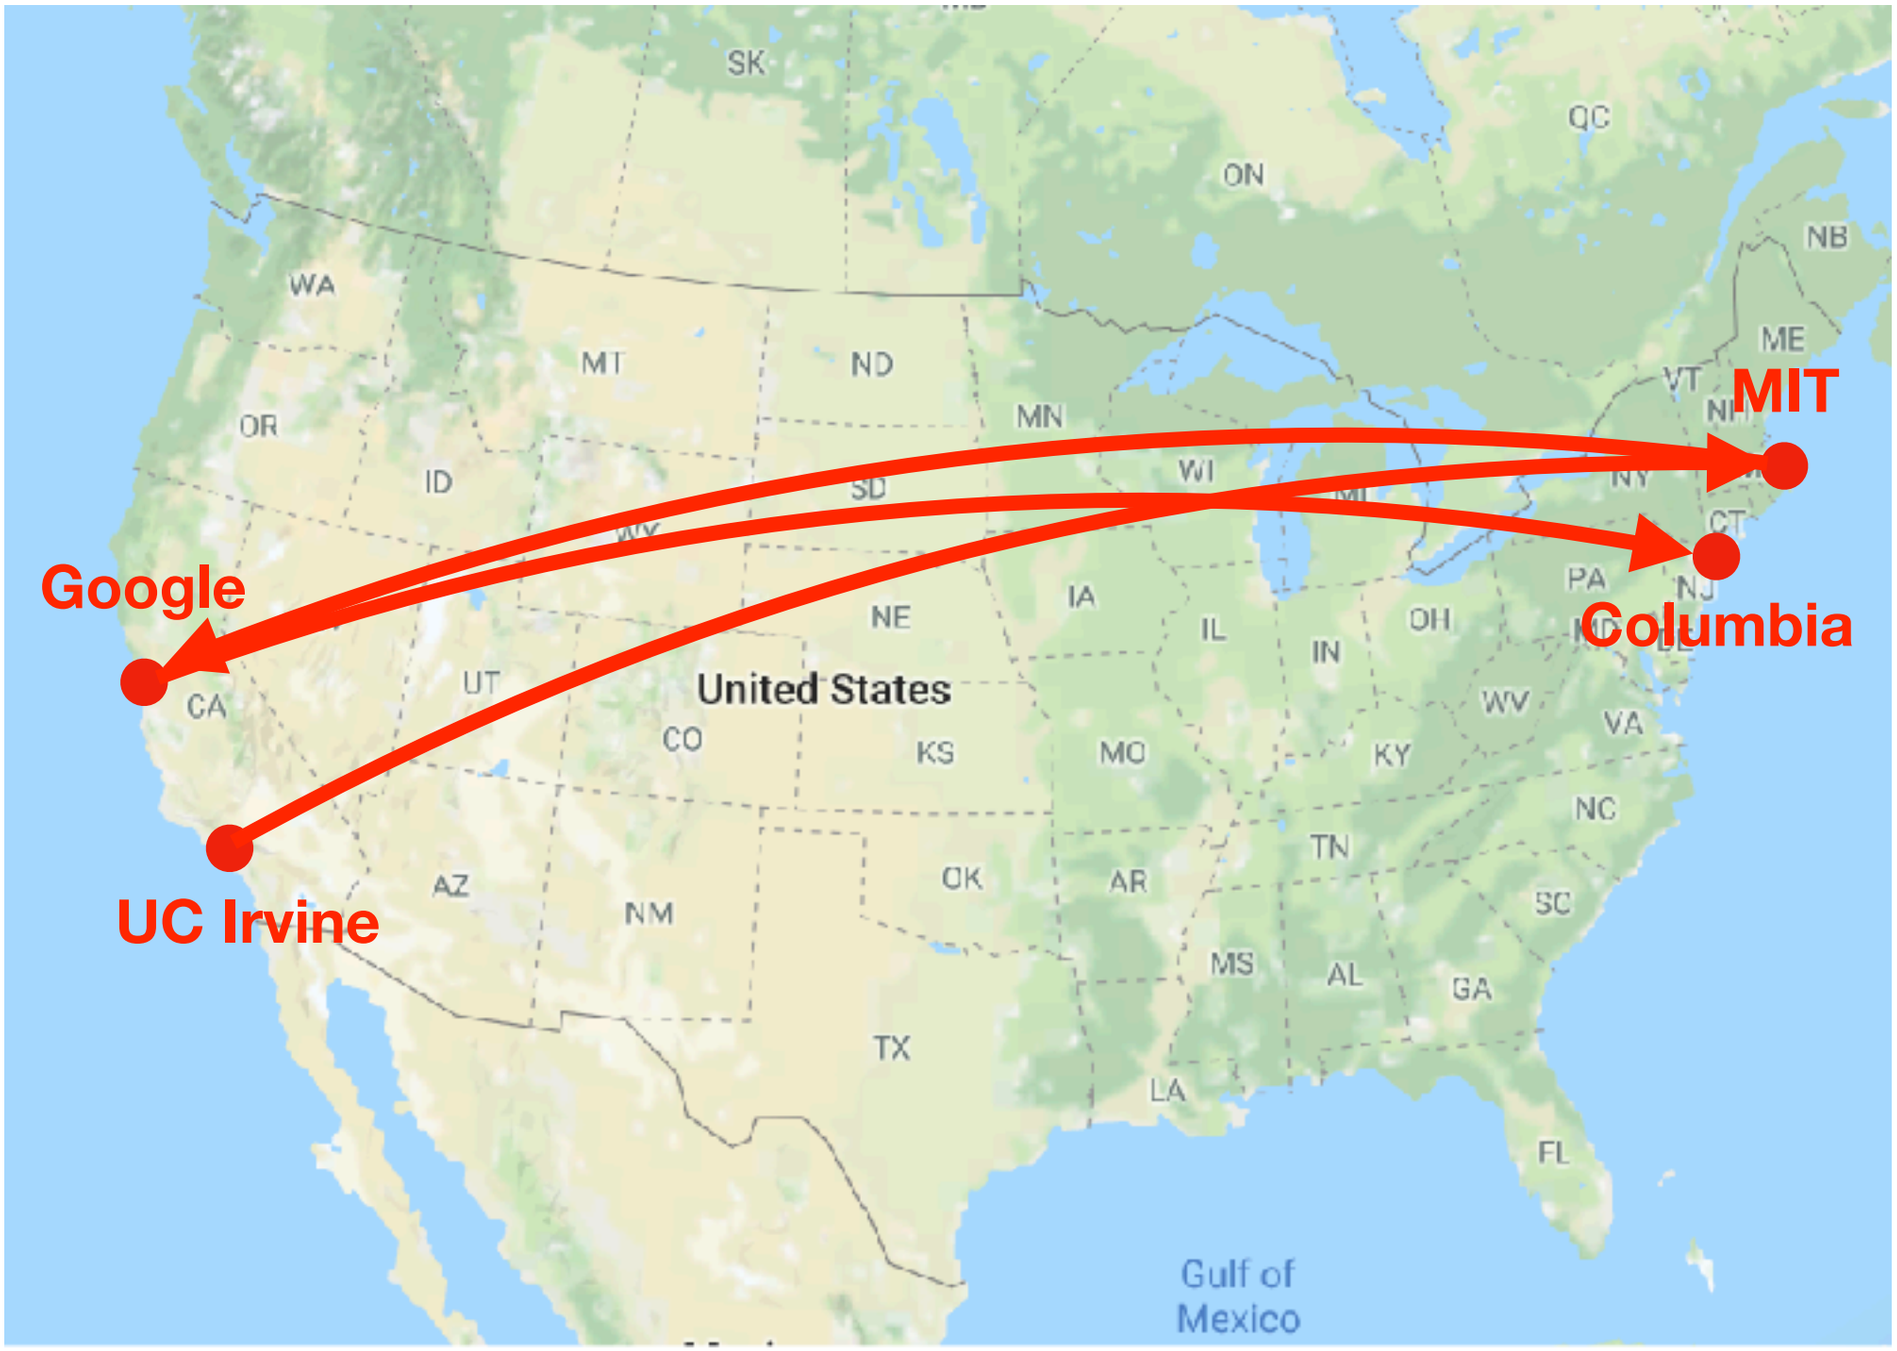
MIT
Google
Columbia
UC Irvine
United States
Gulf of Mexico
SK
ON
QC
NB
WA
MT
ND
ME
VT
NH
MN
OR
ID
SD
WI
NY
CT
WY
PA
NJ
IA
NE
IL
IN
OH
MD
DE
UT
CA
CO
KS
MO
WV
VA
KY
NC
TN
AZ
NM
OK
AR
SC
MS
AL
GA
TX
LA
FL

What about you?

- Major?
- Year?
- Research area?

Staff and Office Hours

- **Carl Vondrick**
  Office Hours: Monday 4:30pm to 5:30pm
  CSB 502 (temporary)

- TAs:
  - **Oscar:** TBA
  - **Xiaoning:** Monday, 5-6pm, CS TA Room
  - **Bo:** Tuesday, 3-4pm, CS TA Room
  - **James:** TBA
  - **Luc:** TBA

FAQ: Can you add me?

- We’re at capacity: 110 people enrolled
- 200 people on wait list
- If you don’t plan to take class, please drop soon

FAQ: Do I need to know C?

- No. The problem sets will use **Python**.

- Familiarity with linear algebra and calculus will be helpful but not required.

FAQ: How to contact you?

- No emails — please use **Piazza**
- You can send private messages on Piazza
- Course staff goes offline 7pm to 10am and weekends

#	Date	Topic	Reading	Assignments
1	Sep 5	Introduction		
Image Processing				
2	Sep 10	Image Processing I		HW1 Out
3	Sep 12	Image Processing II		
Cameras and Physics				
4	Sep 17	Image Formation		
5	Sep 19	Image Sensing		
6	Sep 24	Radiometry and Reflectance		HW1 Due, HW2 Out
Recognition and Matching				
7	Sep 26	Grouping: Edges and Boundaries		
8	Oct 1	Image Alignment and Stitching		
9	Oct 3	2D Recognition using SIFT		
10	Oct 8	Learning-based Vision I		HW2 Due, HW3 Out
11	Oct 10	Learning-based Vision II		
12	Oct 15	Object and Face Detection		
13	Oct 17	Image Segmentation		

3D Vision			
14	Oct 22	Photometric Stereo	HW3 Due, HW4 Out
15	Oct 24	Shape from Shading	
16	Oct 29	Depth from Focus/Defocus	
17	Oct 31	Camera Calibration and Simple Stereo	Project Proposal Due
-	Nov 5	Academic Holiday	
18	Nov 7	Uncalibrated Stereo	HW4 Due, HW5 Out
Video			
19	Nov 12	Motion and Optical Flow	
20	Nov 14	Object Tracking	
21	Nov 19	Structure from Motion	HW5 Due
-	Nov 21	Academic Holiday (Thanksgiving)	
Frontiers			
22	Nov 26	Vision, Sound, and Touch	
23	Nov 28	Unsupervised and Transfer Learning	
24	Dec 3	Bias and Ethics	
25	Dec 5	Project Presentations	
26	Dec 10	Project Presentations	

Grading

- 60% Problem Sets
- 40% Final Project
- No exams or quizzes

Problem Sets

- 5 problem sets, equally weighted
- Turn in via **CourseWorks** before class starts. Submit both PDF writeup and code online.
- One problem set may be a week late. No other extensions.
- Solutions available during TA office hours.
- Done individually, but you can have high-level discussion in pairs. Write up assignments individually

Final Project

- Individually or pairs (recommended)
- Final poster presentations: **Dec 5 and Dec 10**
- 4 page report in CVPR format
- Suggested projects and grading rubric to be announced

Academic Honesty

- Academic dishonesty may result in…
  - You fail course.
  - We refer your case to the Dean’s office.

Readings (Optional)

http://szeliski.org/Book/

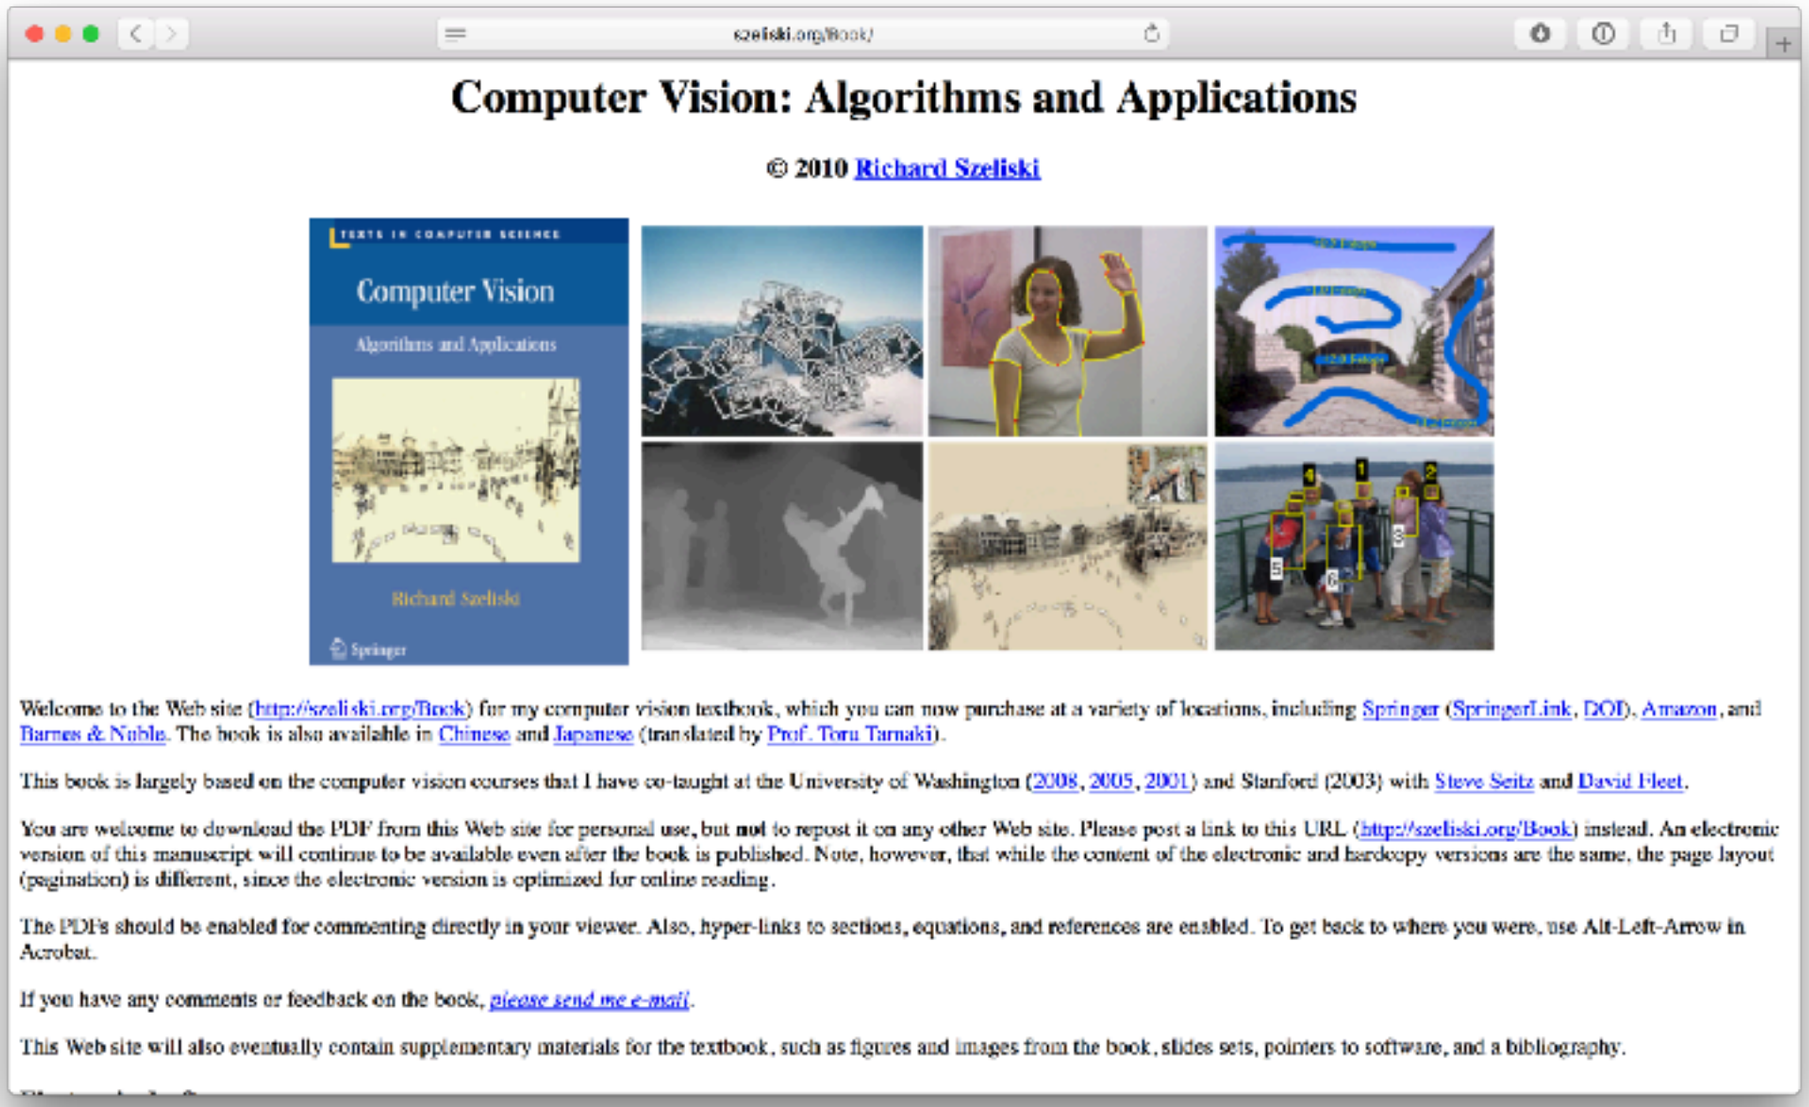

New Course

- Feedback appreciated.

- Please let us know if something works or not!

CS@CU

Distinguished Lecture Series 2018

Deep Learning to Learn

Sept. 10, 2018 (Monday)

11:30AM–12:30PM

Davis Auditorium

(412 CEPSR)

Pieter Abbeel (UC Berkeley)

Next Class: Linear Filters