Lecture 1: Introduction

COS 429: Computer Vision

http://cs.princeton.edu/~cos429/

Slide credits: Fei-Fei Li, Deva Ramanan, Szymon Rusinkiewicz

What does the future world look like?

Slide credit: Fei-Fei Li

Y.Z. Dai, 7 year old

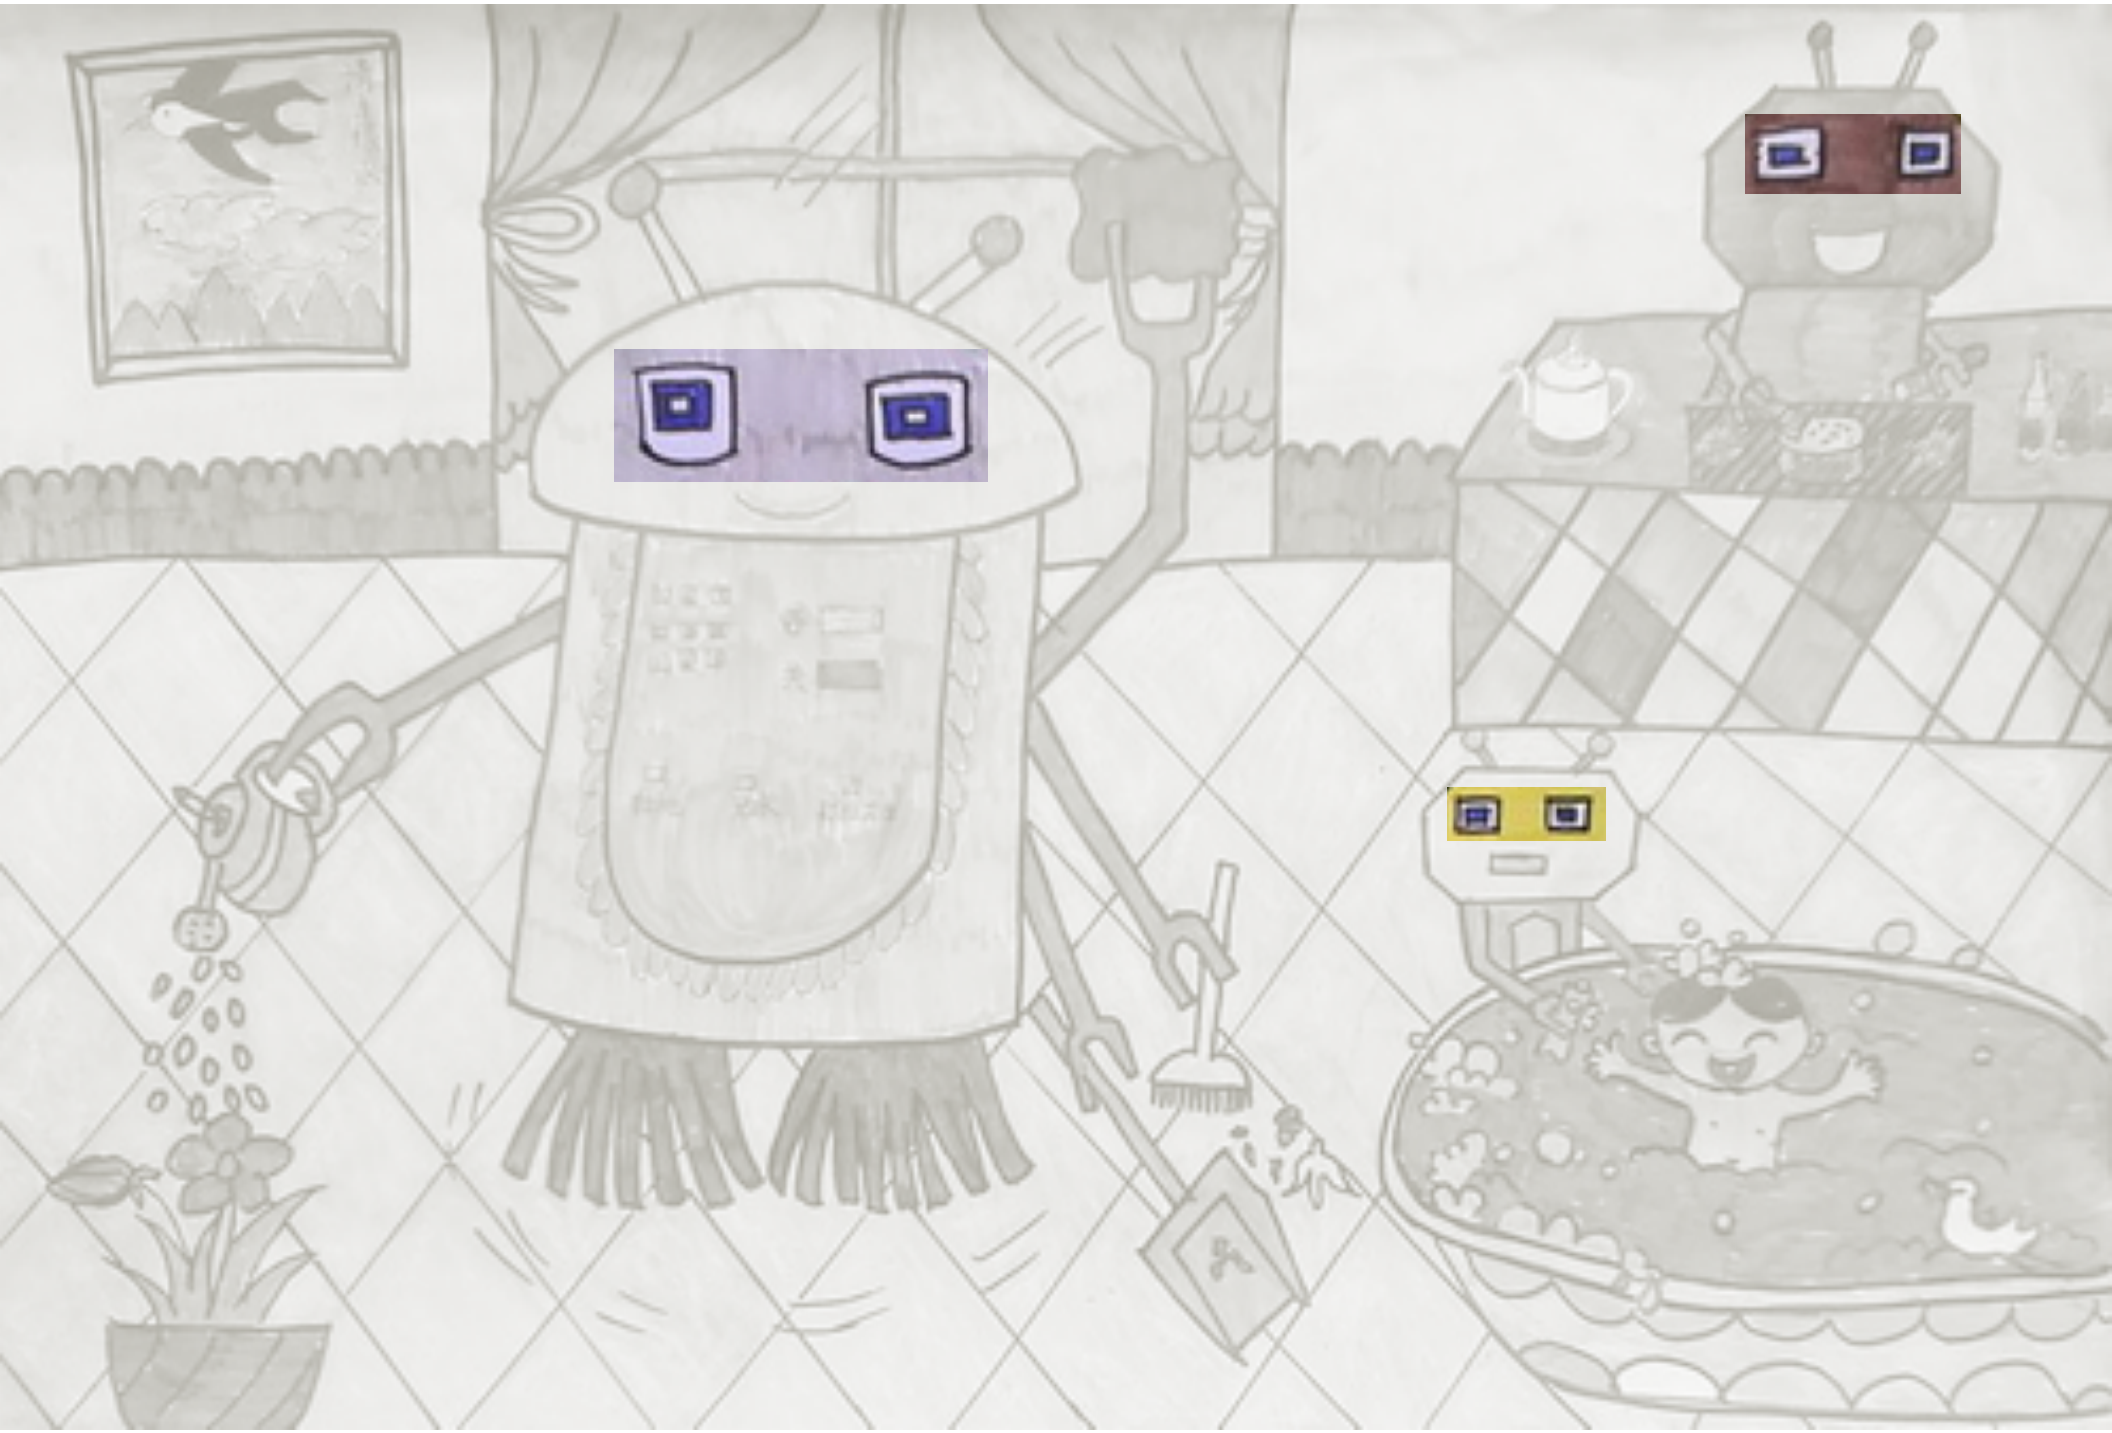

Slide credit: Fei-Fei Li

Y.Z. Dai, 7 year old

Slide credit: Fei-Fei Li

Y.Z. Dai, 7 year old

Slide credit: Fei-Fei Li

Slide credit: Fei-Fei Li

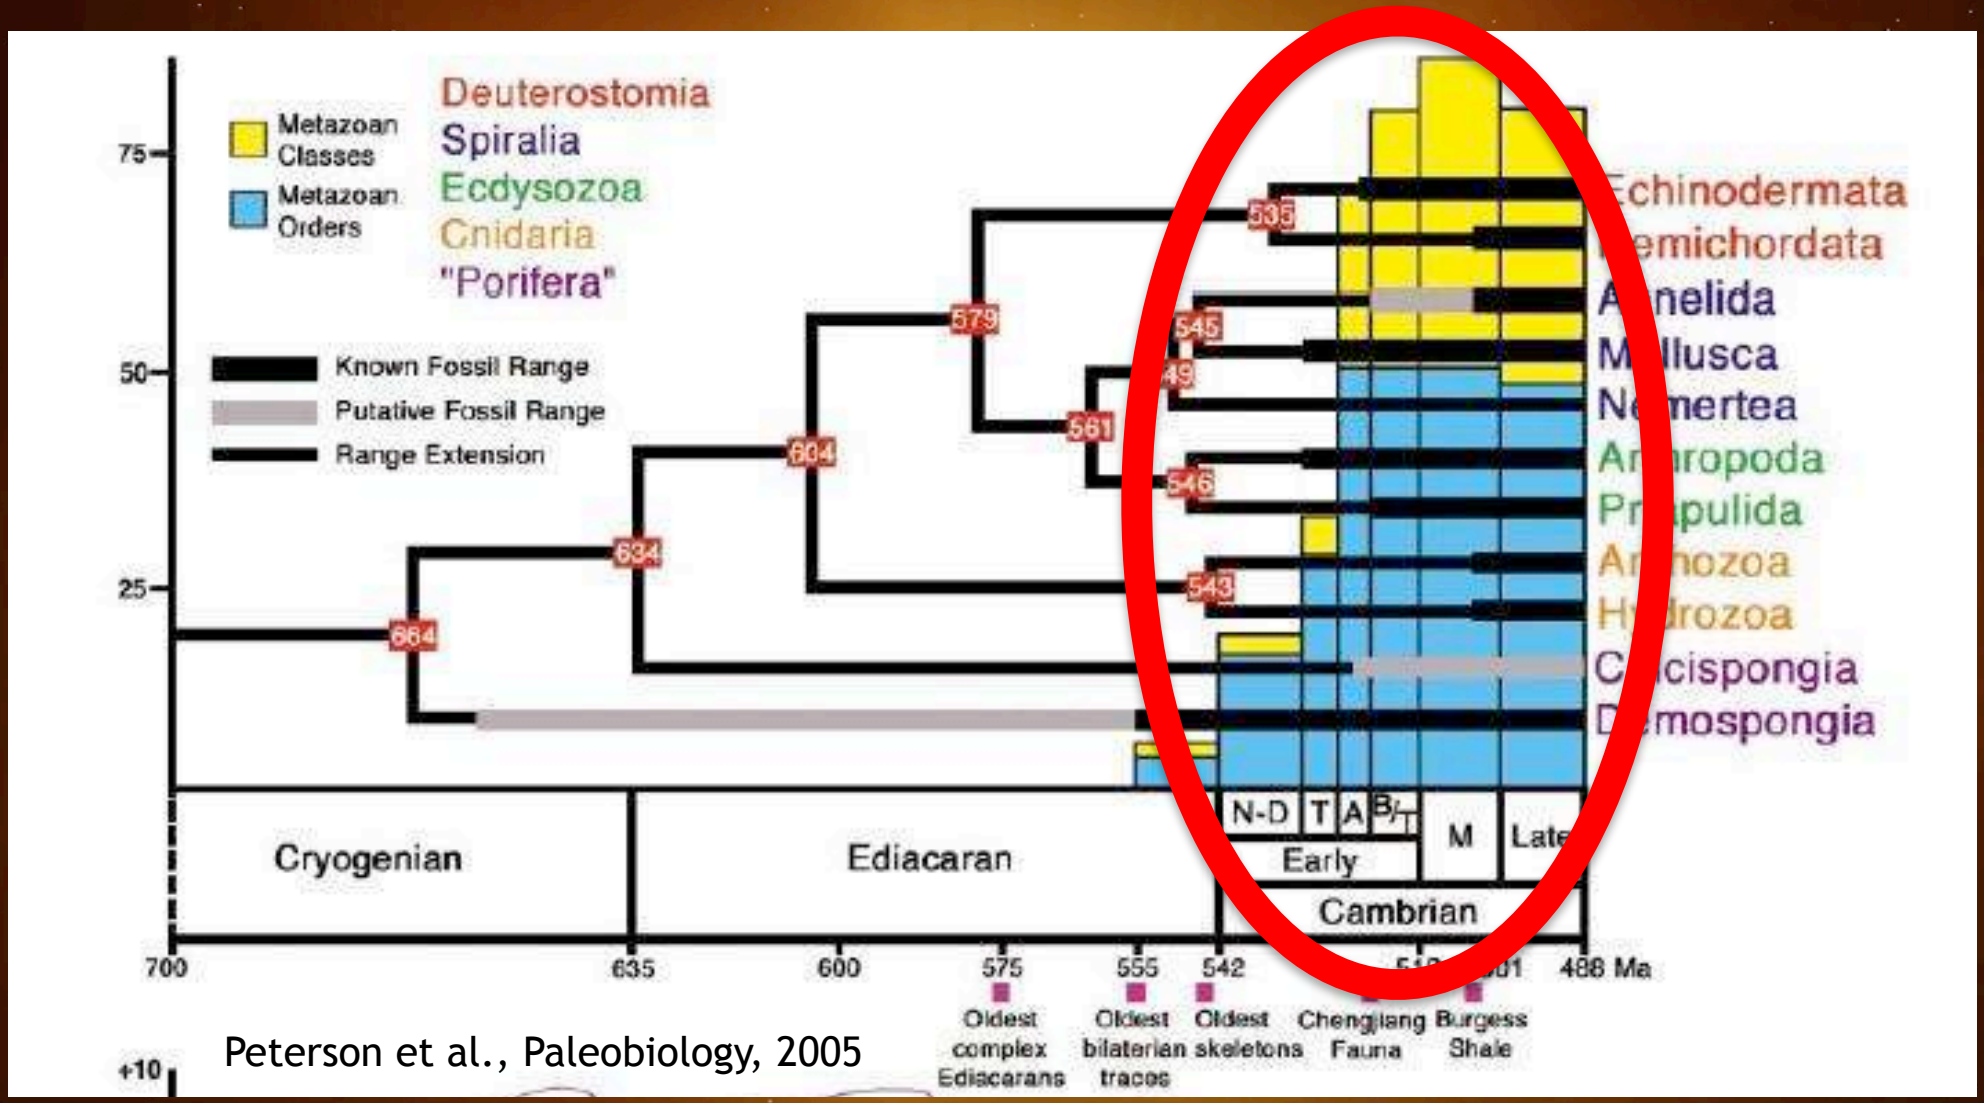

Peterson et al., Paleobiology, 2005

Slide credit: Fei-Fei Li

"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, zoologist

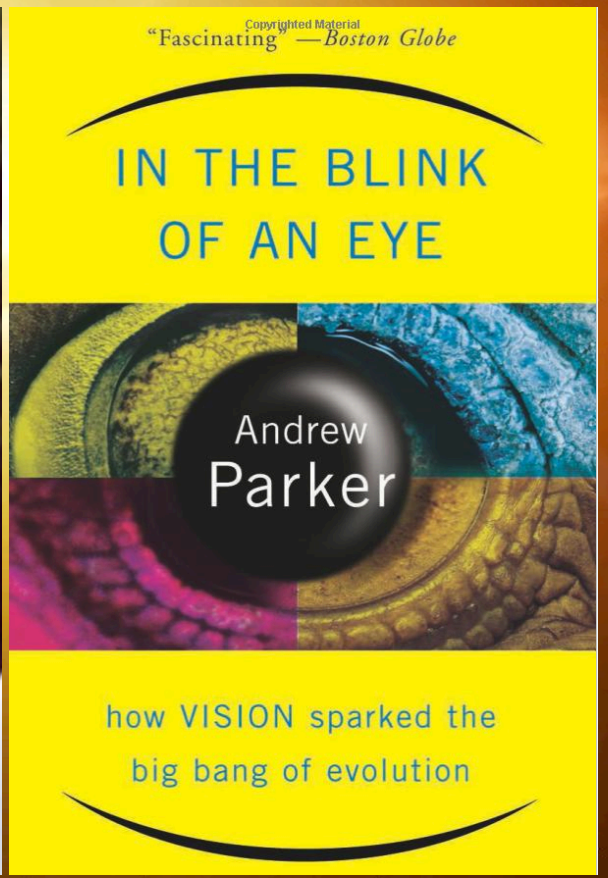

Slide credit: Fei-Fei Li

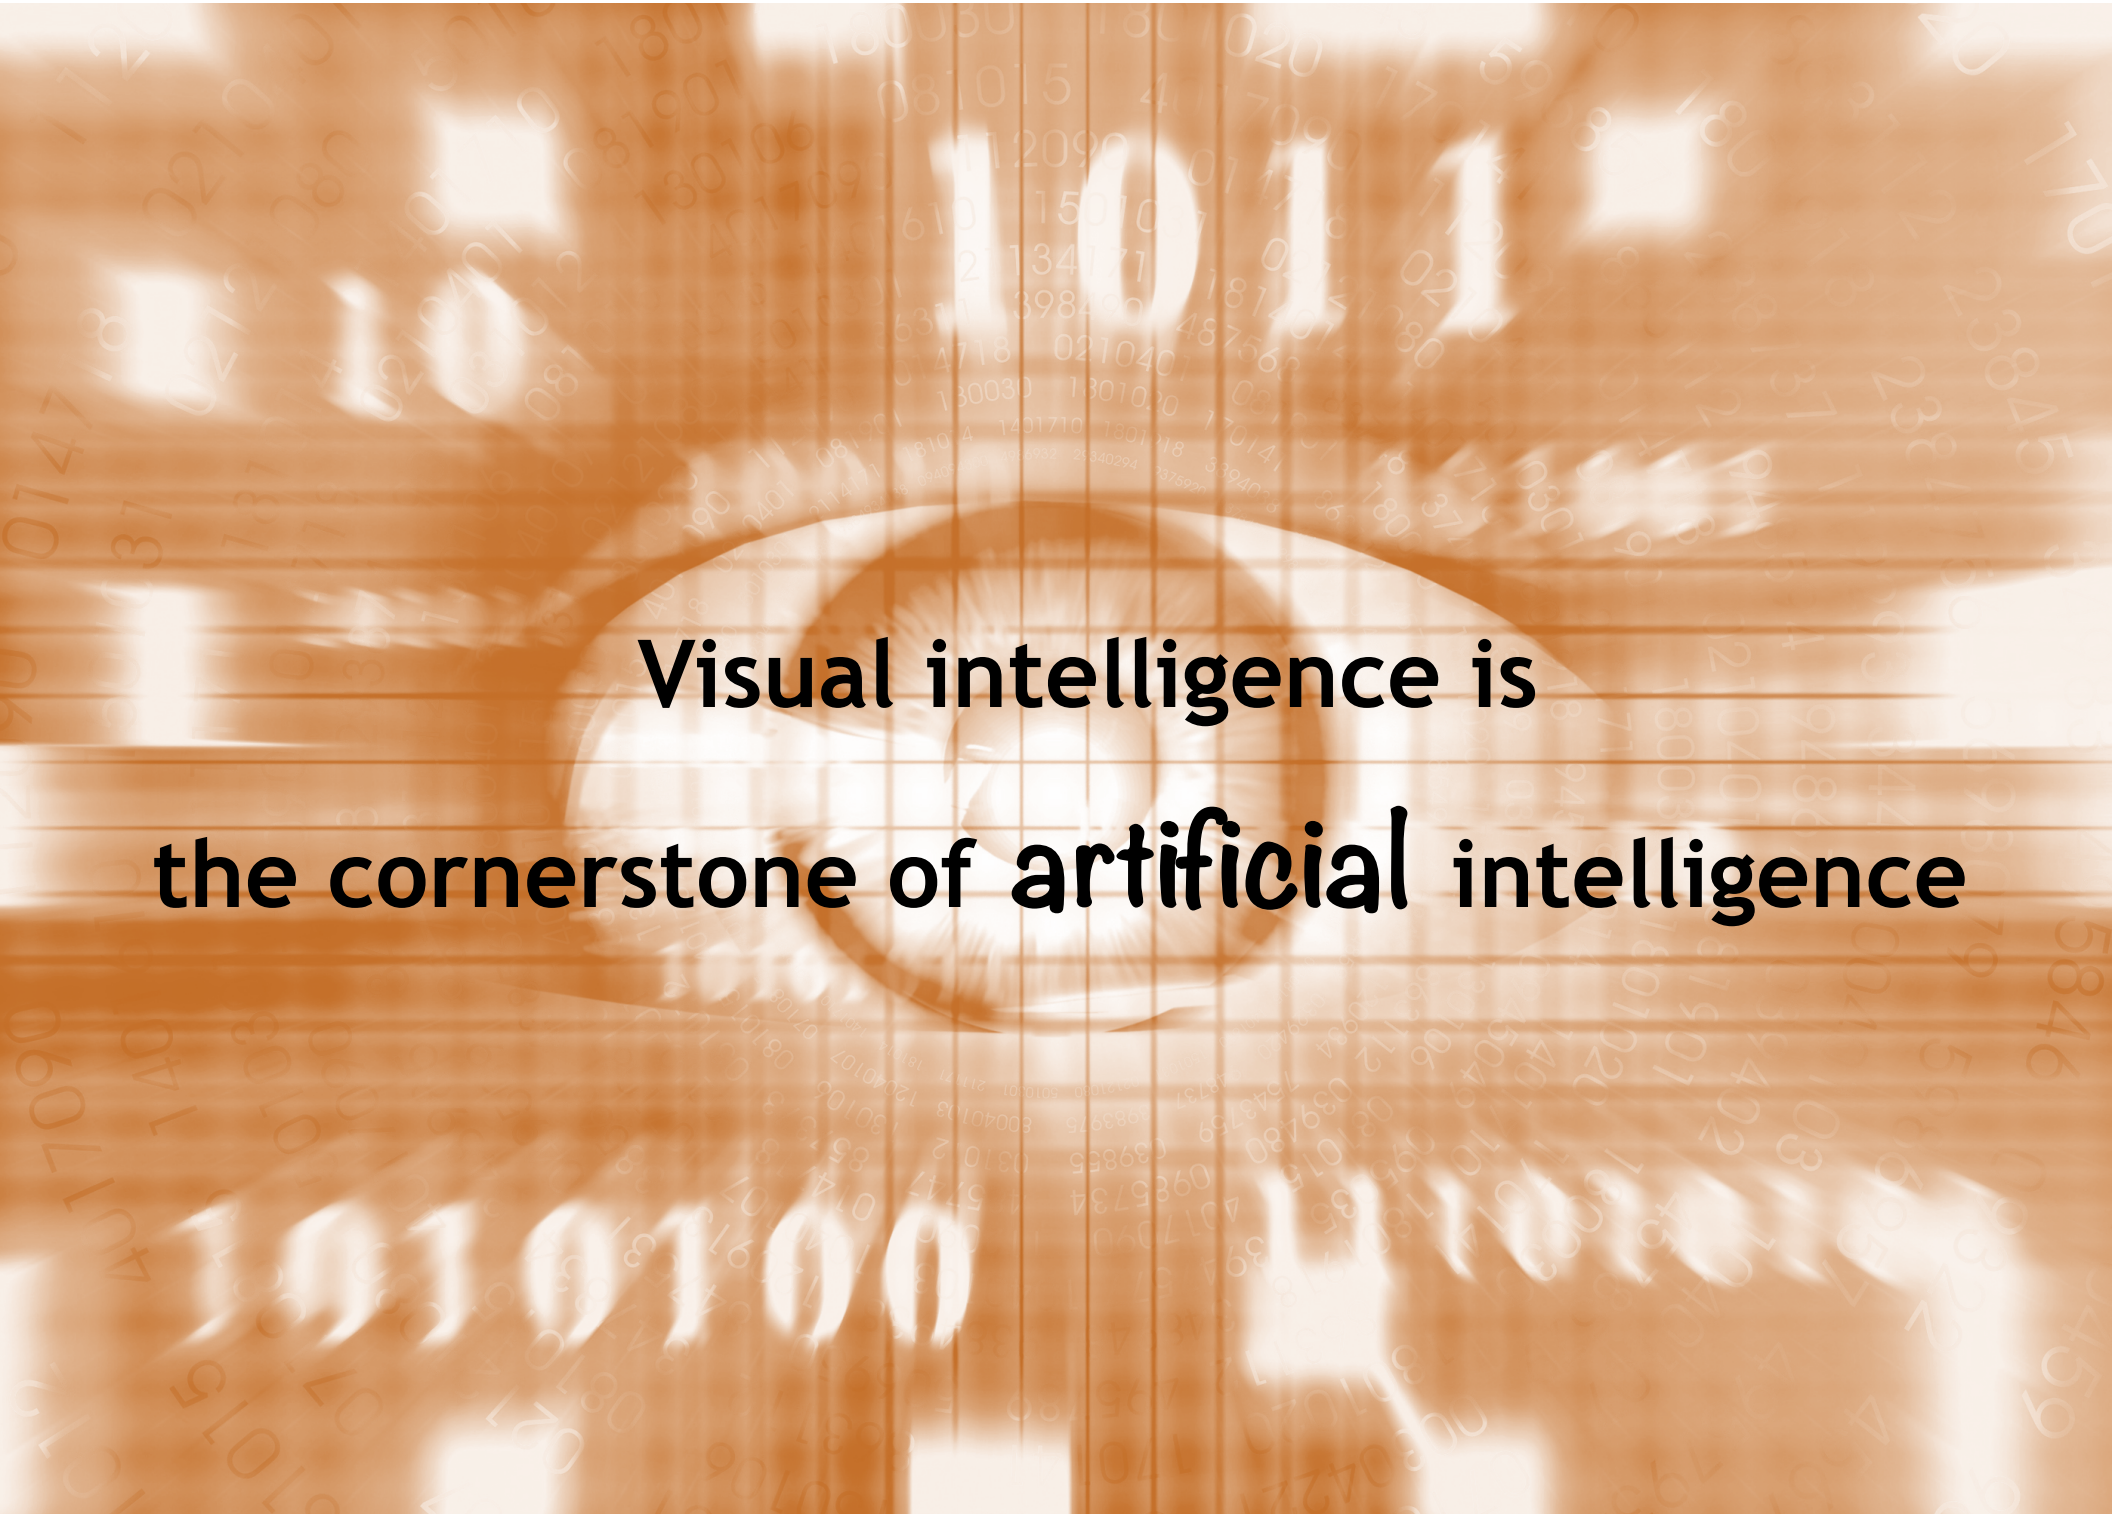

Slide credit: Fei-Fei Li






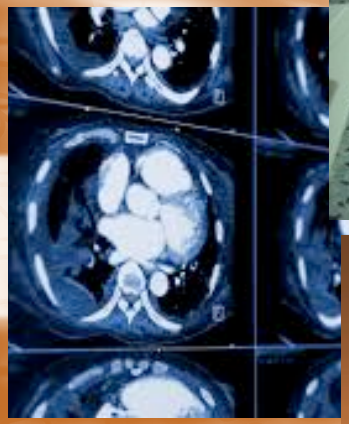





Slide credit: Fei-Fei Li

Seeing is challenging

Challenge: illumination

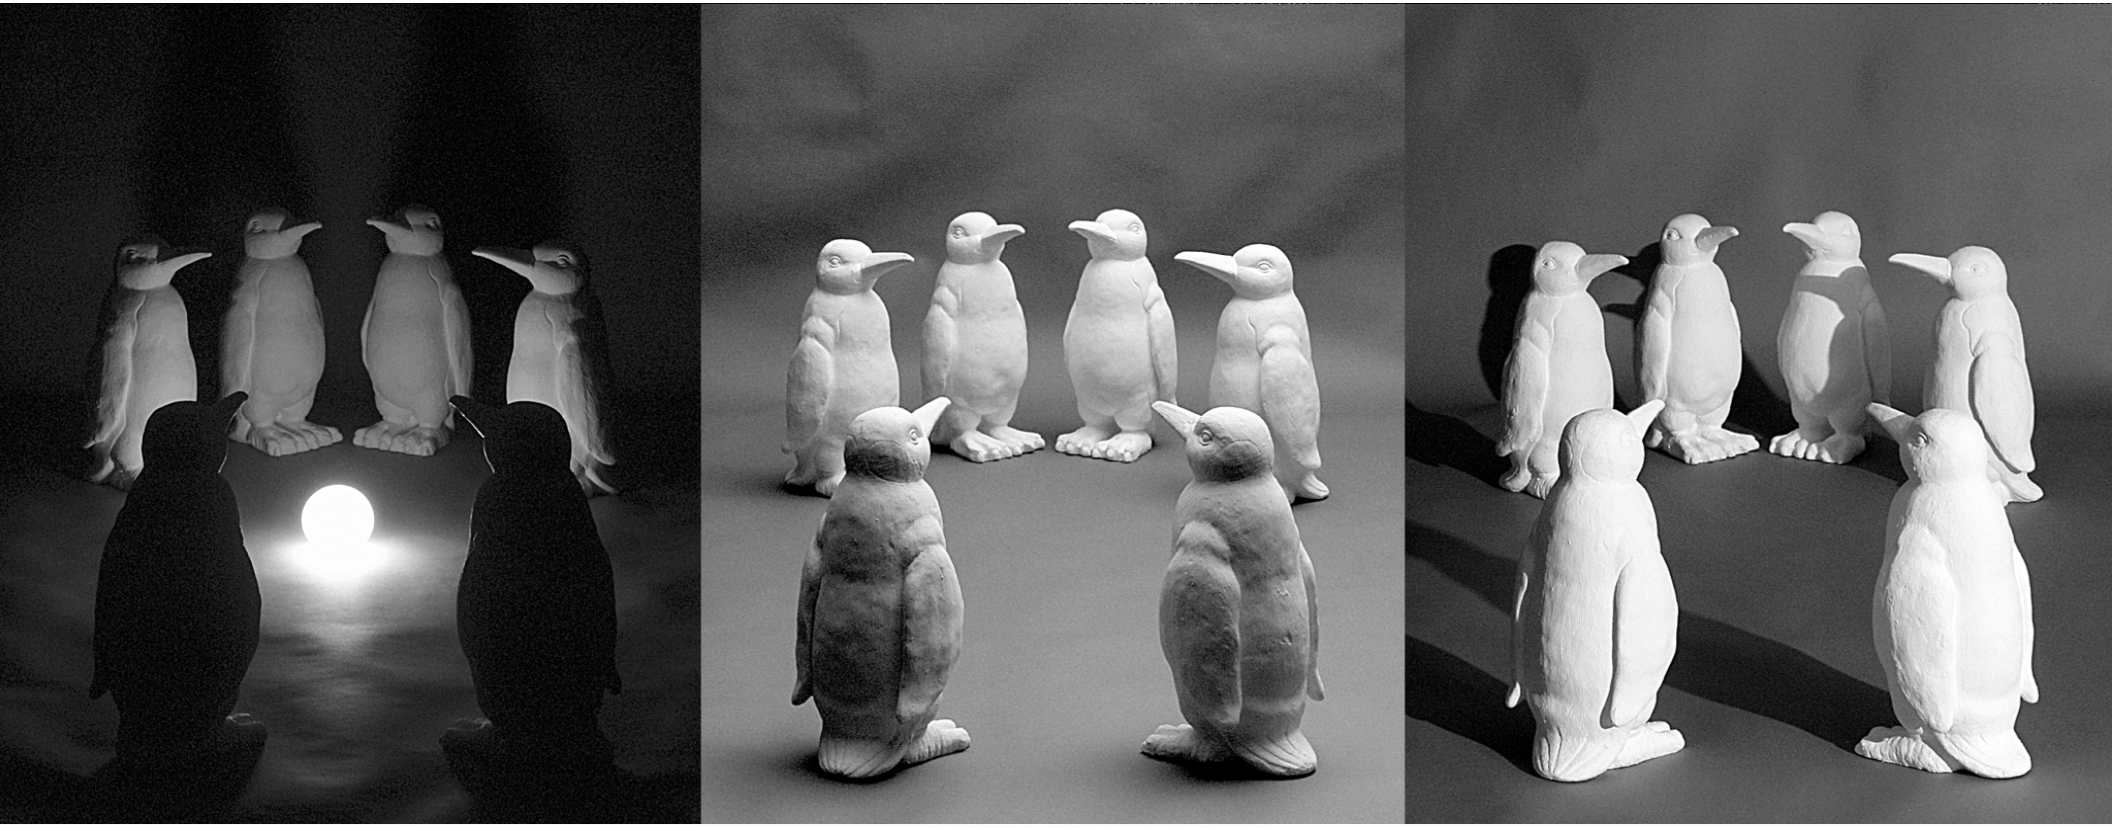

image credit: J. Koenderink

Challenge: light and shadow

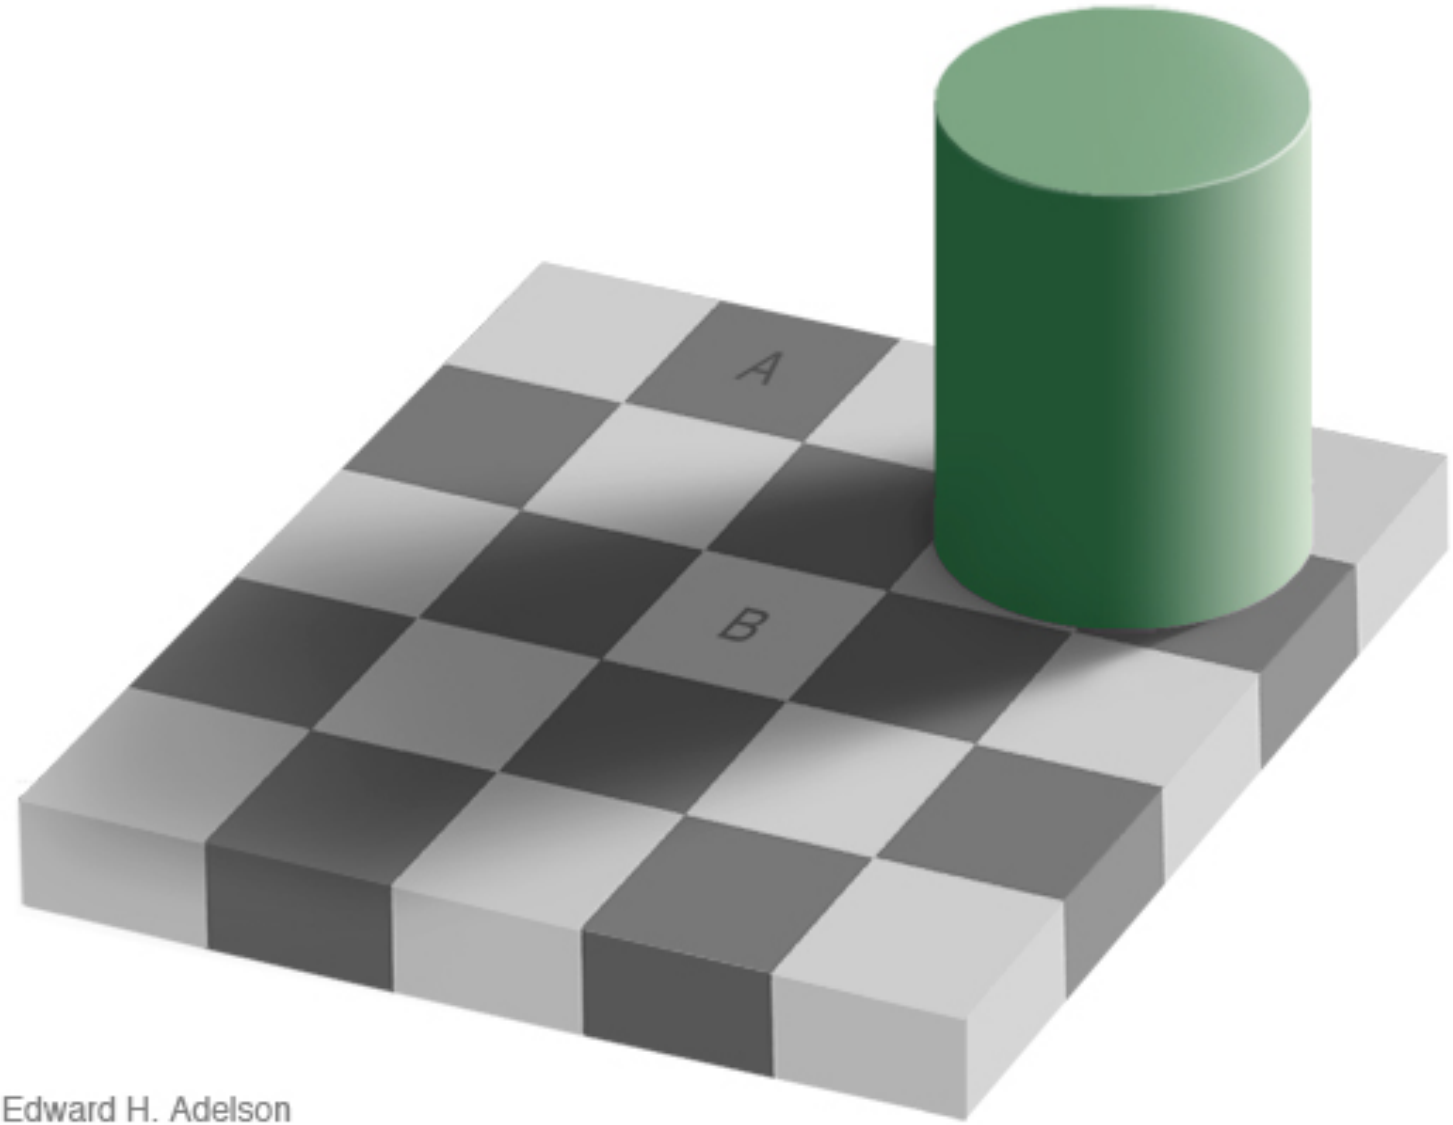

Edward H. Adelson

Challenge: light and shadow

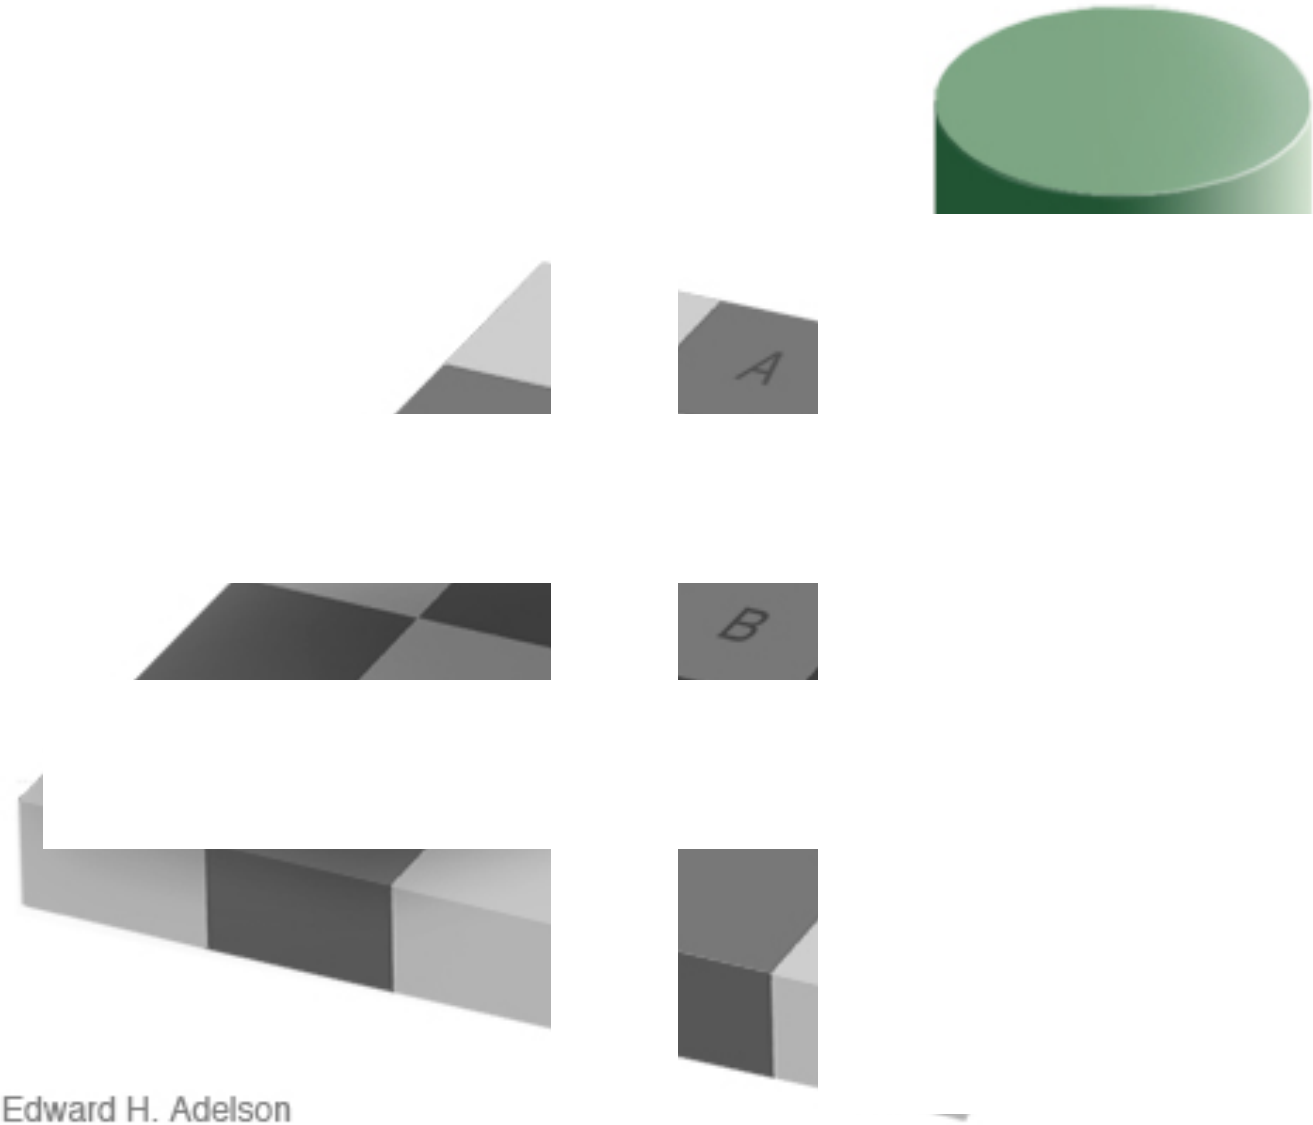

Edward H. Adelson

Challenge: light and shadow

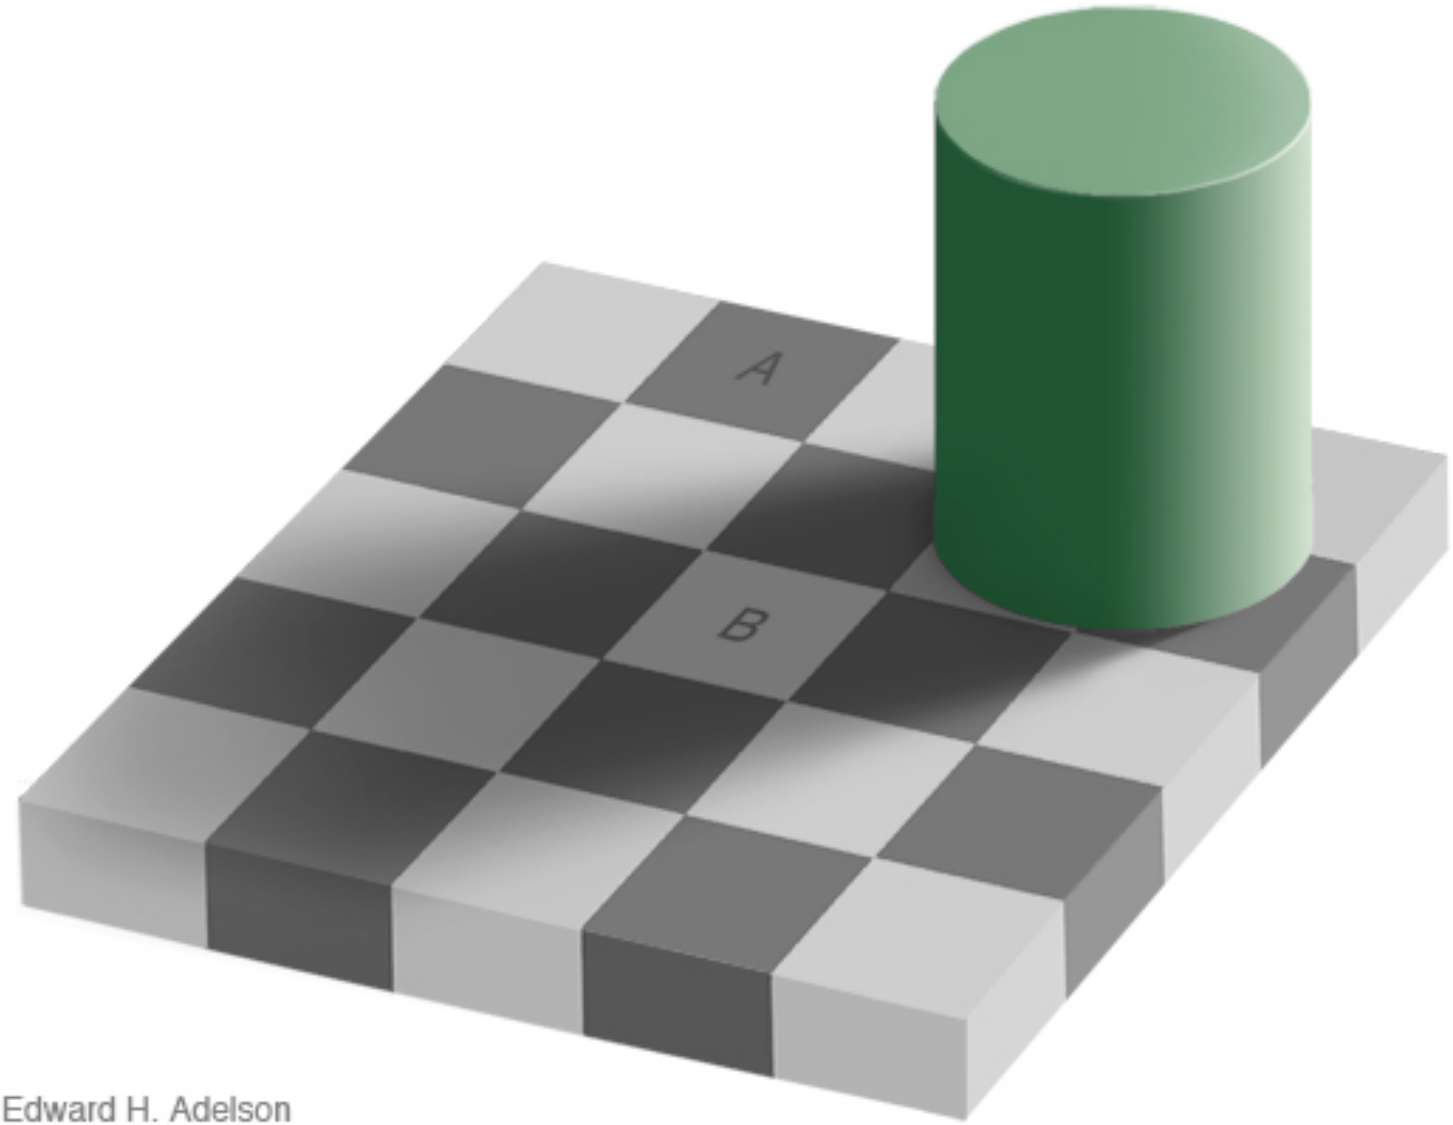

Edward H. Adelson

Challenge: scale

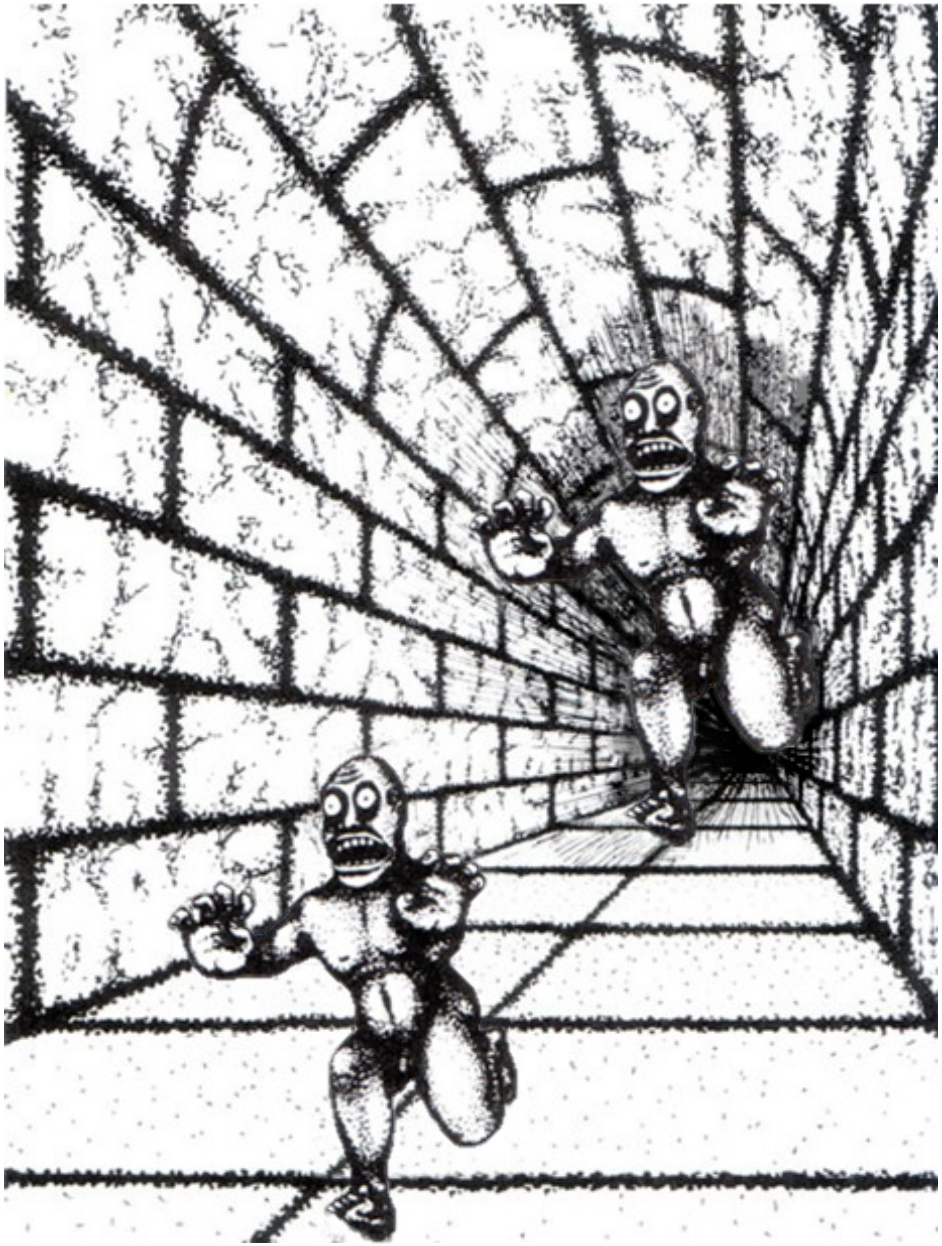

Challenge: scale

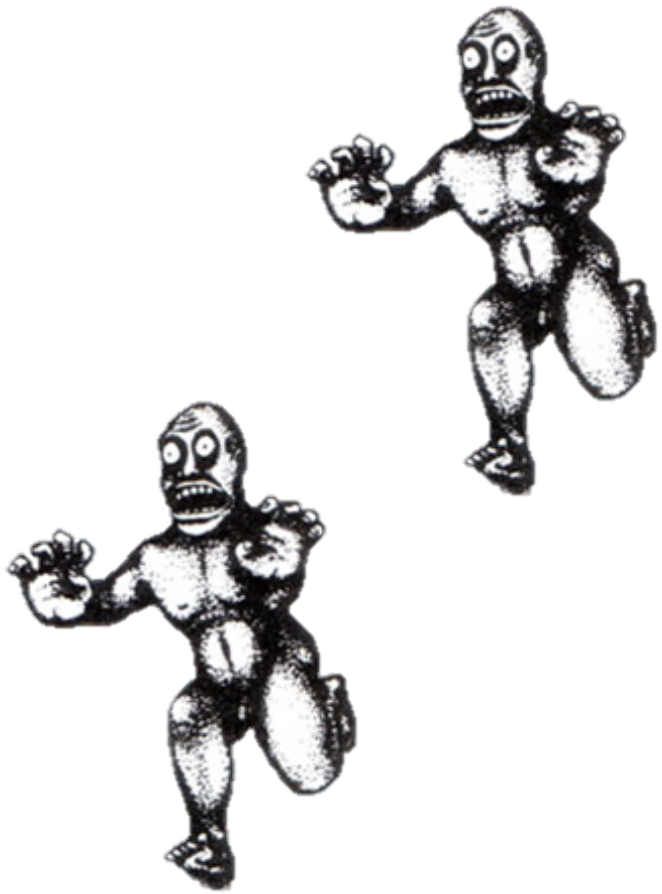

Challenge: scale

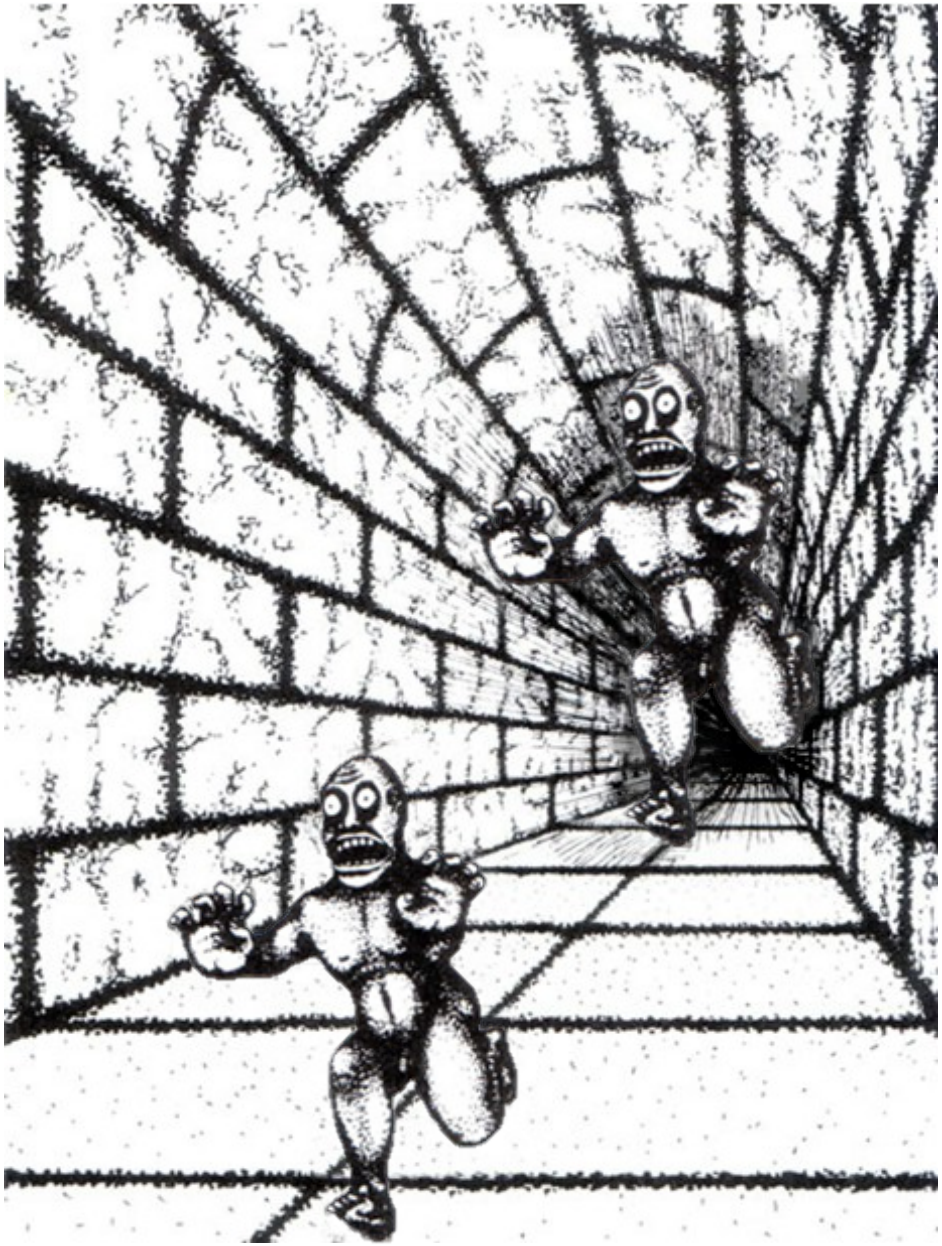

Challenge: occlusion and clutter

Image source: National Geographic

Challenge: intra-class variation

slide credit: Fei-Fei, Fergus & Torralba

Challenge: deformation

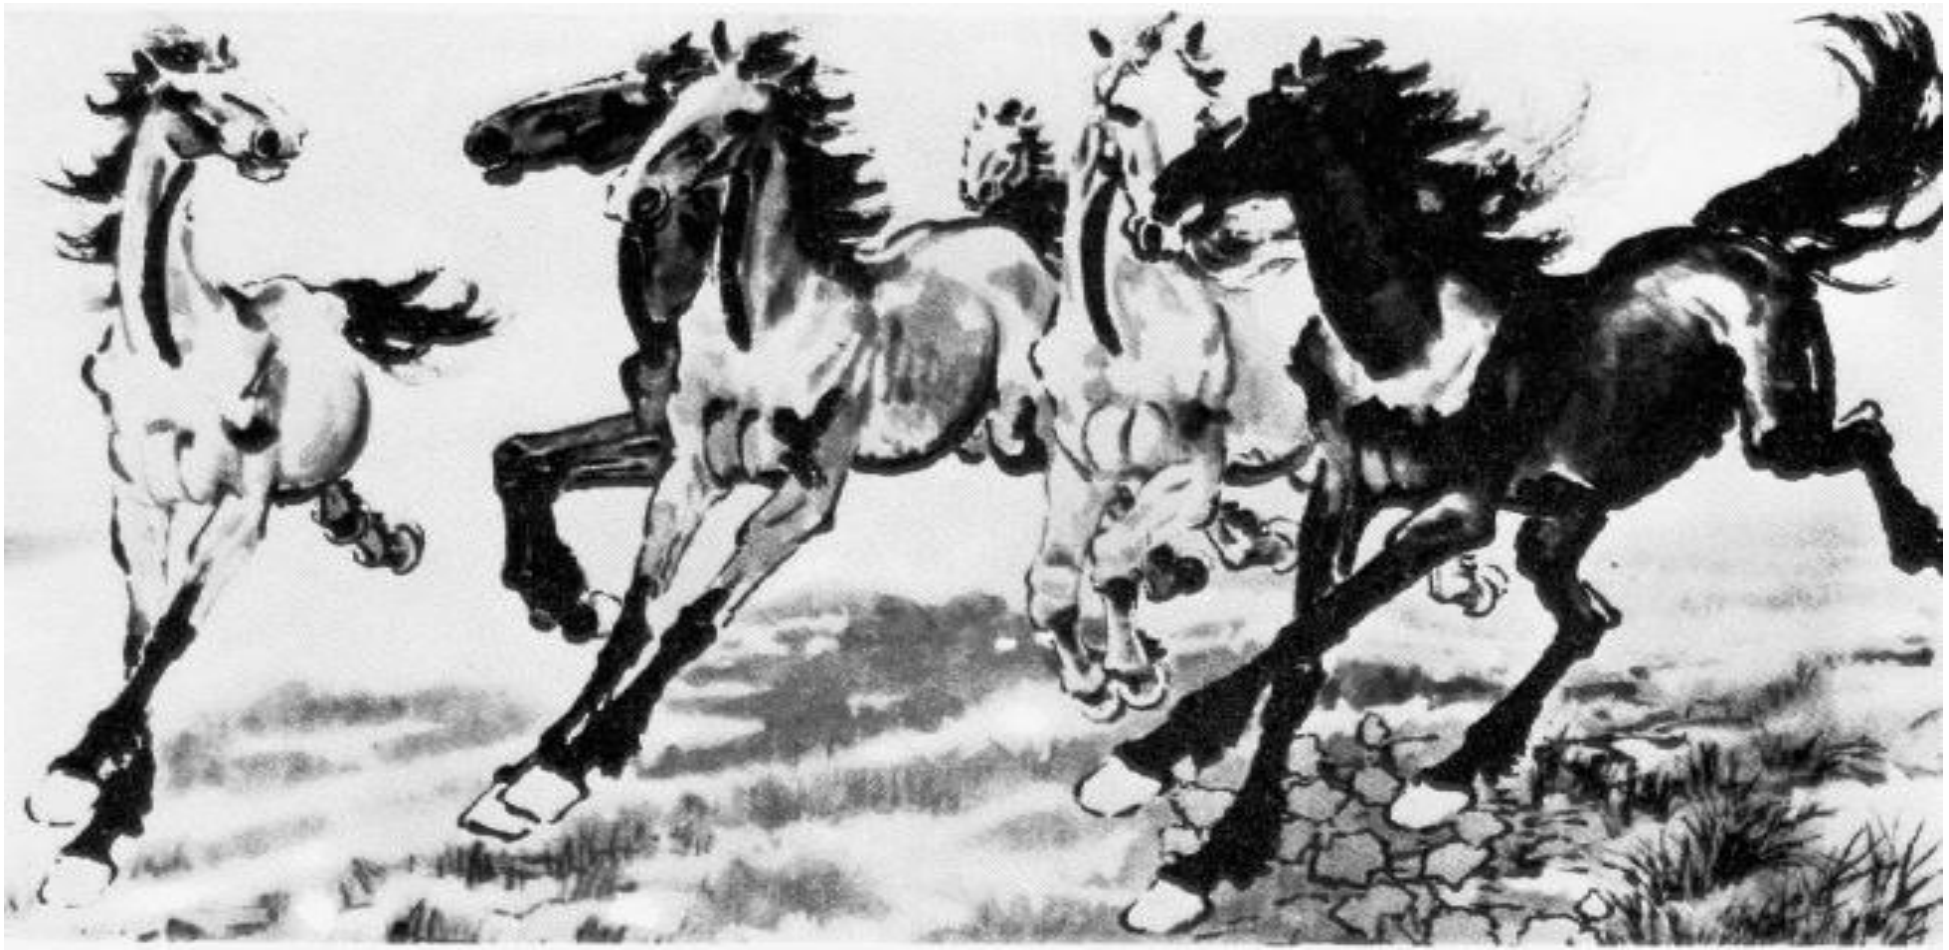

Xu, Beihong 1943

slide credit: Fei-Fei, Fergus & Torralba

Challenge: motion

Challenge: perspective

(Aside) Challenge: intentions and goals

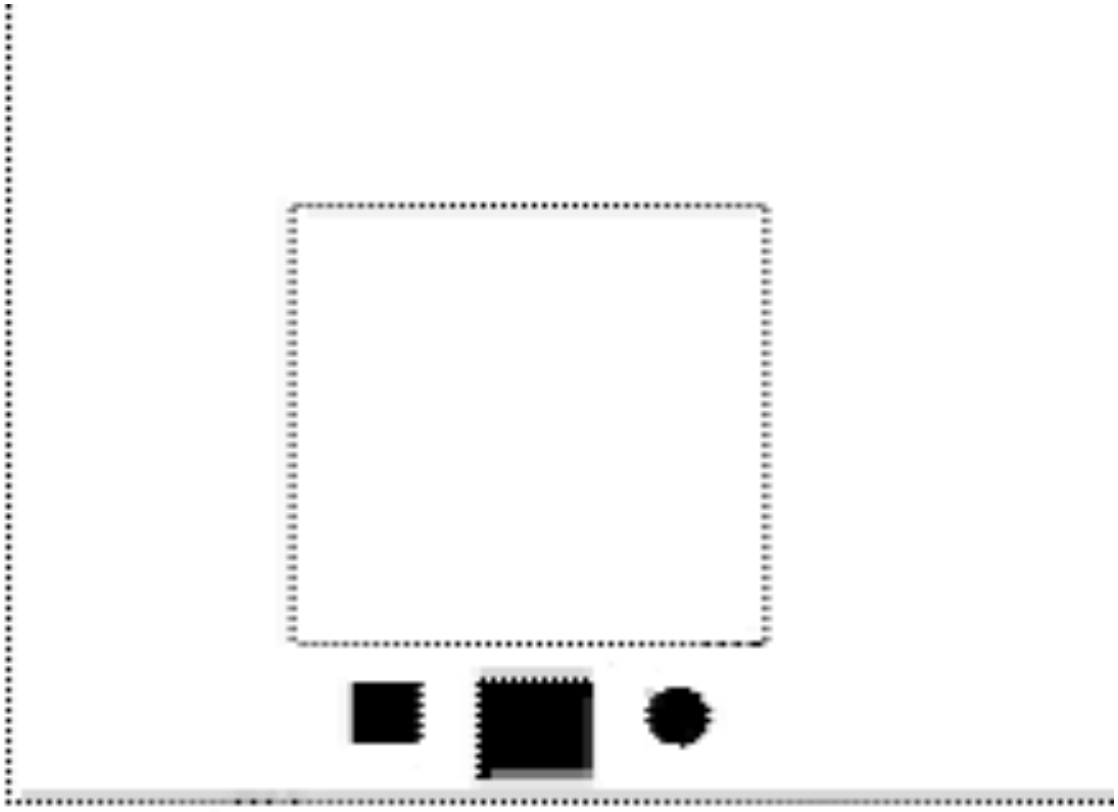

Slide credit: Deva Ramanan

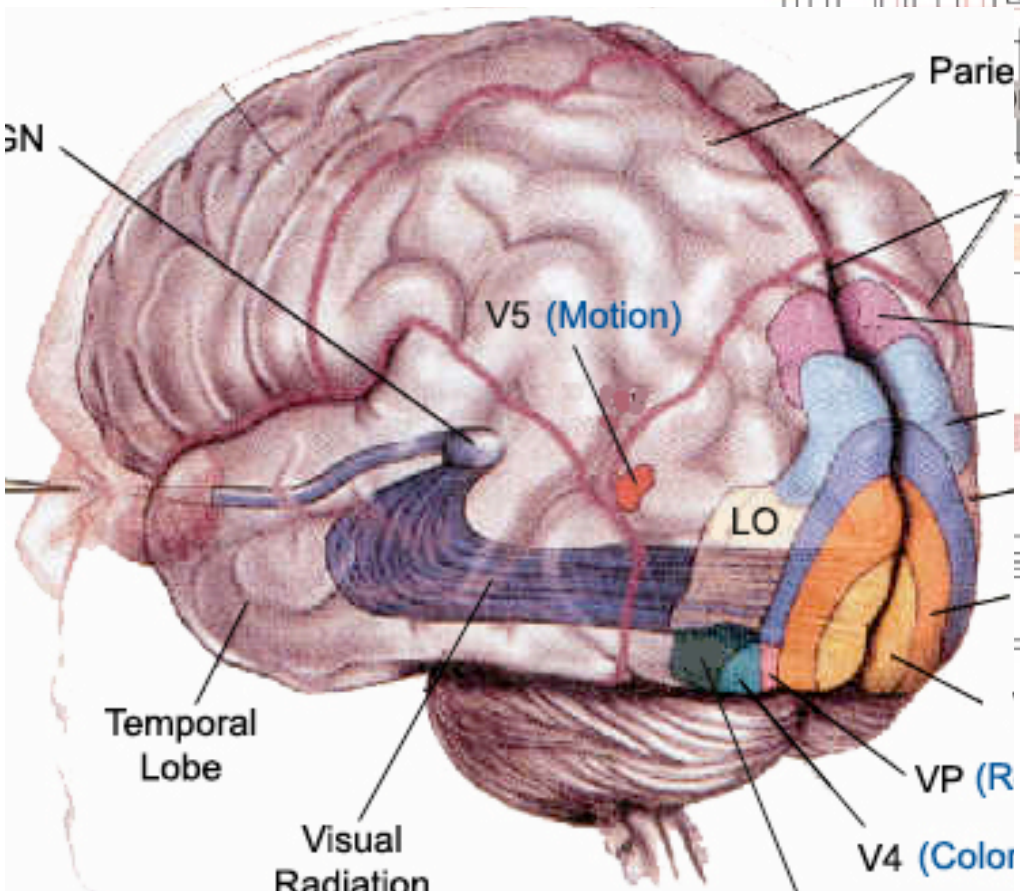

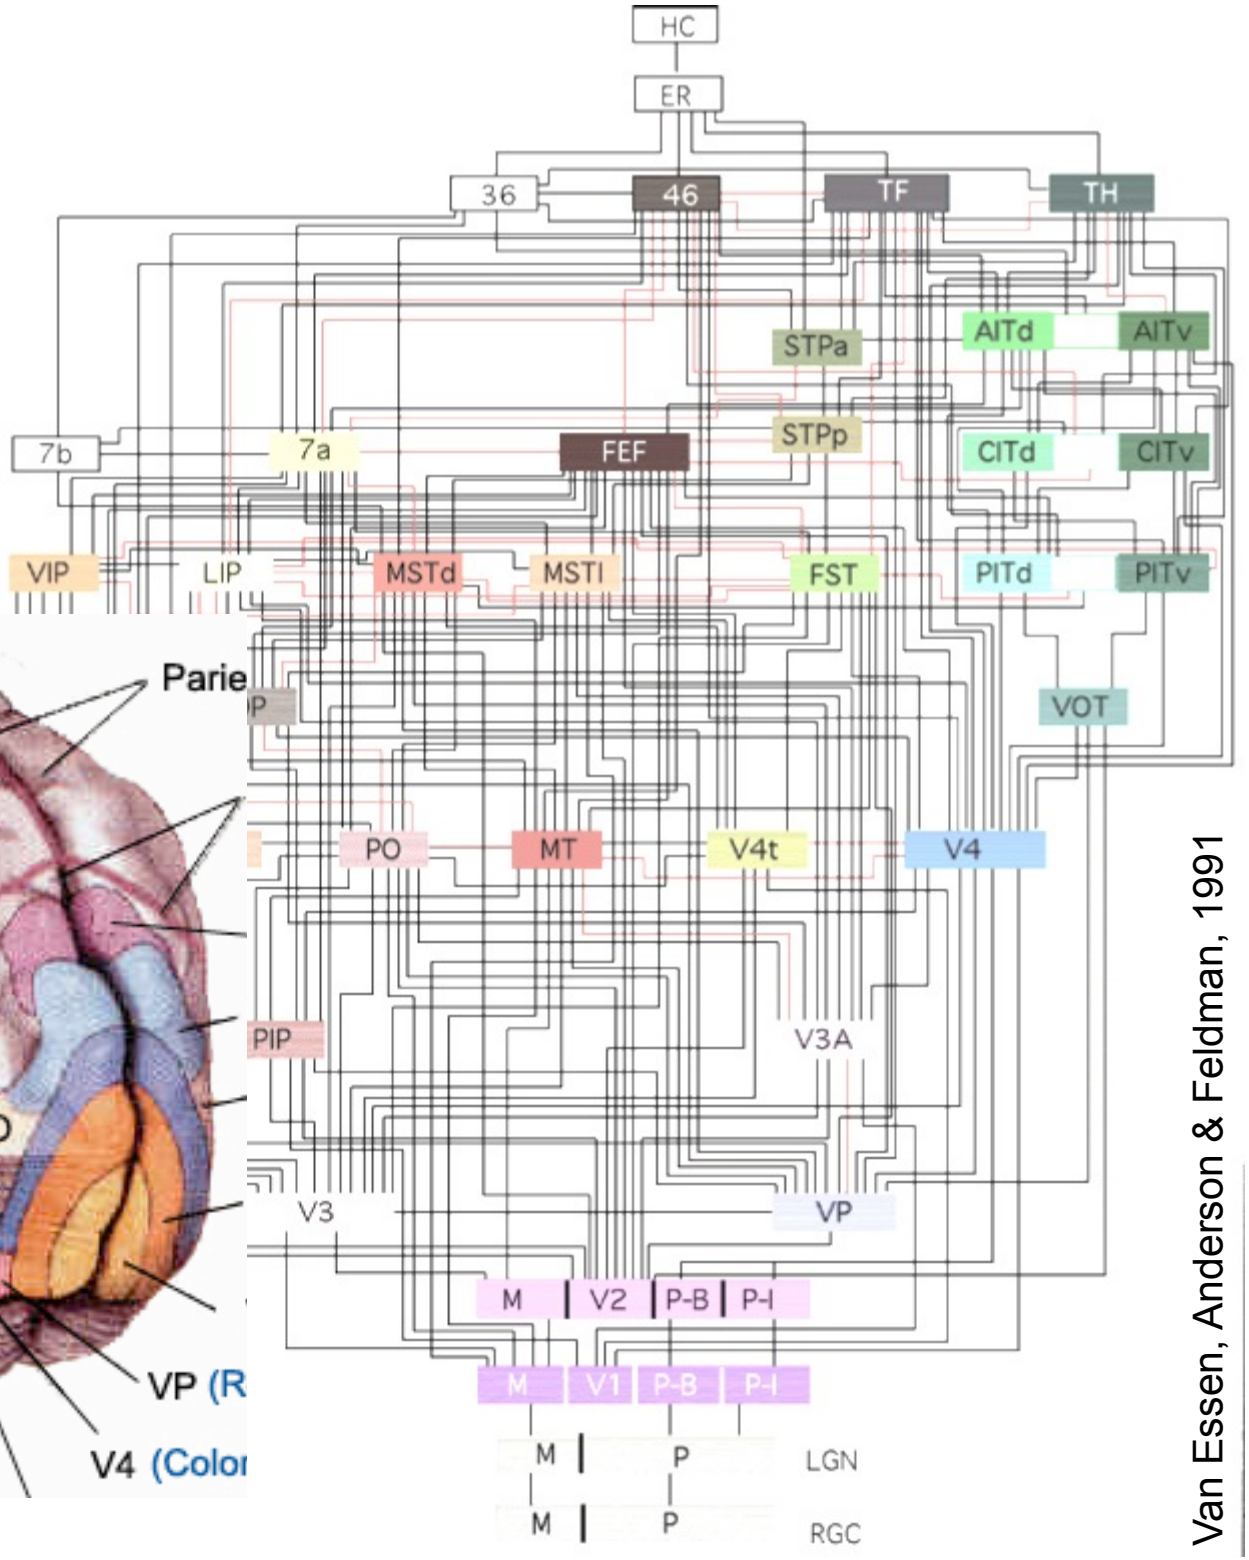

Van Essen, Anderson & Feldman, 1991

What is computer vision?

Computer vision

- Input: images or video
- Output: description of the world
  - Also: measuring, classifying, interpreting visual information

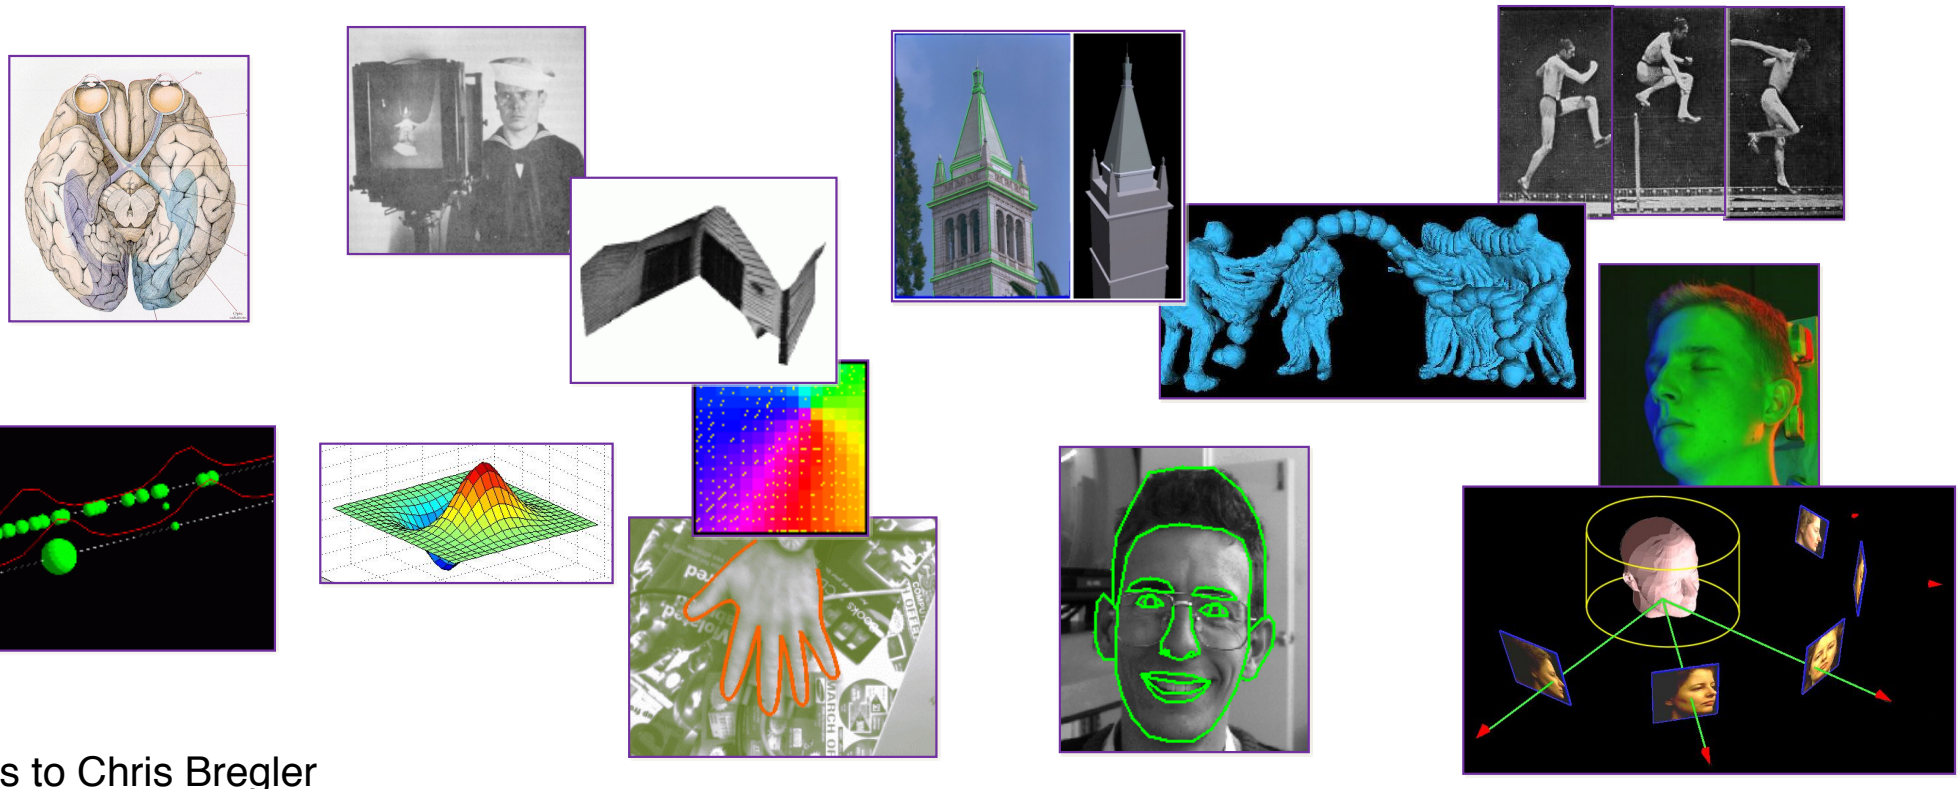

Thanks to Chris Bregler

One Pixel

- Amount of light recorded by a photoreceptor

"Is this the object's color?
Illumination? Noise?
I can't tell!"

Slide credit: Szymon Rusinkiewicz

Low-Level or “Early” Vision

- Local image/shape properties

“There’s an edge!”

Slide credit: Szymon Rusinkiewicz

Mid-Level Vision

"There's an object and a background!"

- Grouping and segmentation

Slide credit: Szymon Rusinkiewicz

High-Level Vision

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

- Recognition
- Classification

Slide credit: Szymon Rusinkiewicz

What We Would Like to Infer…

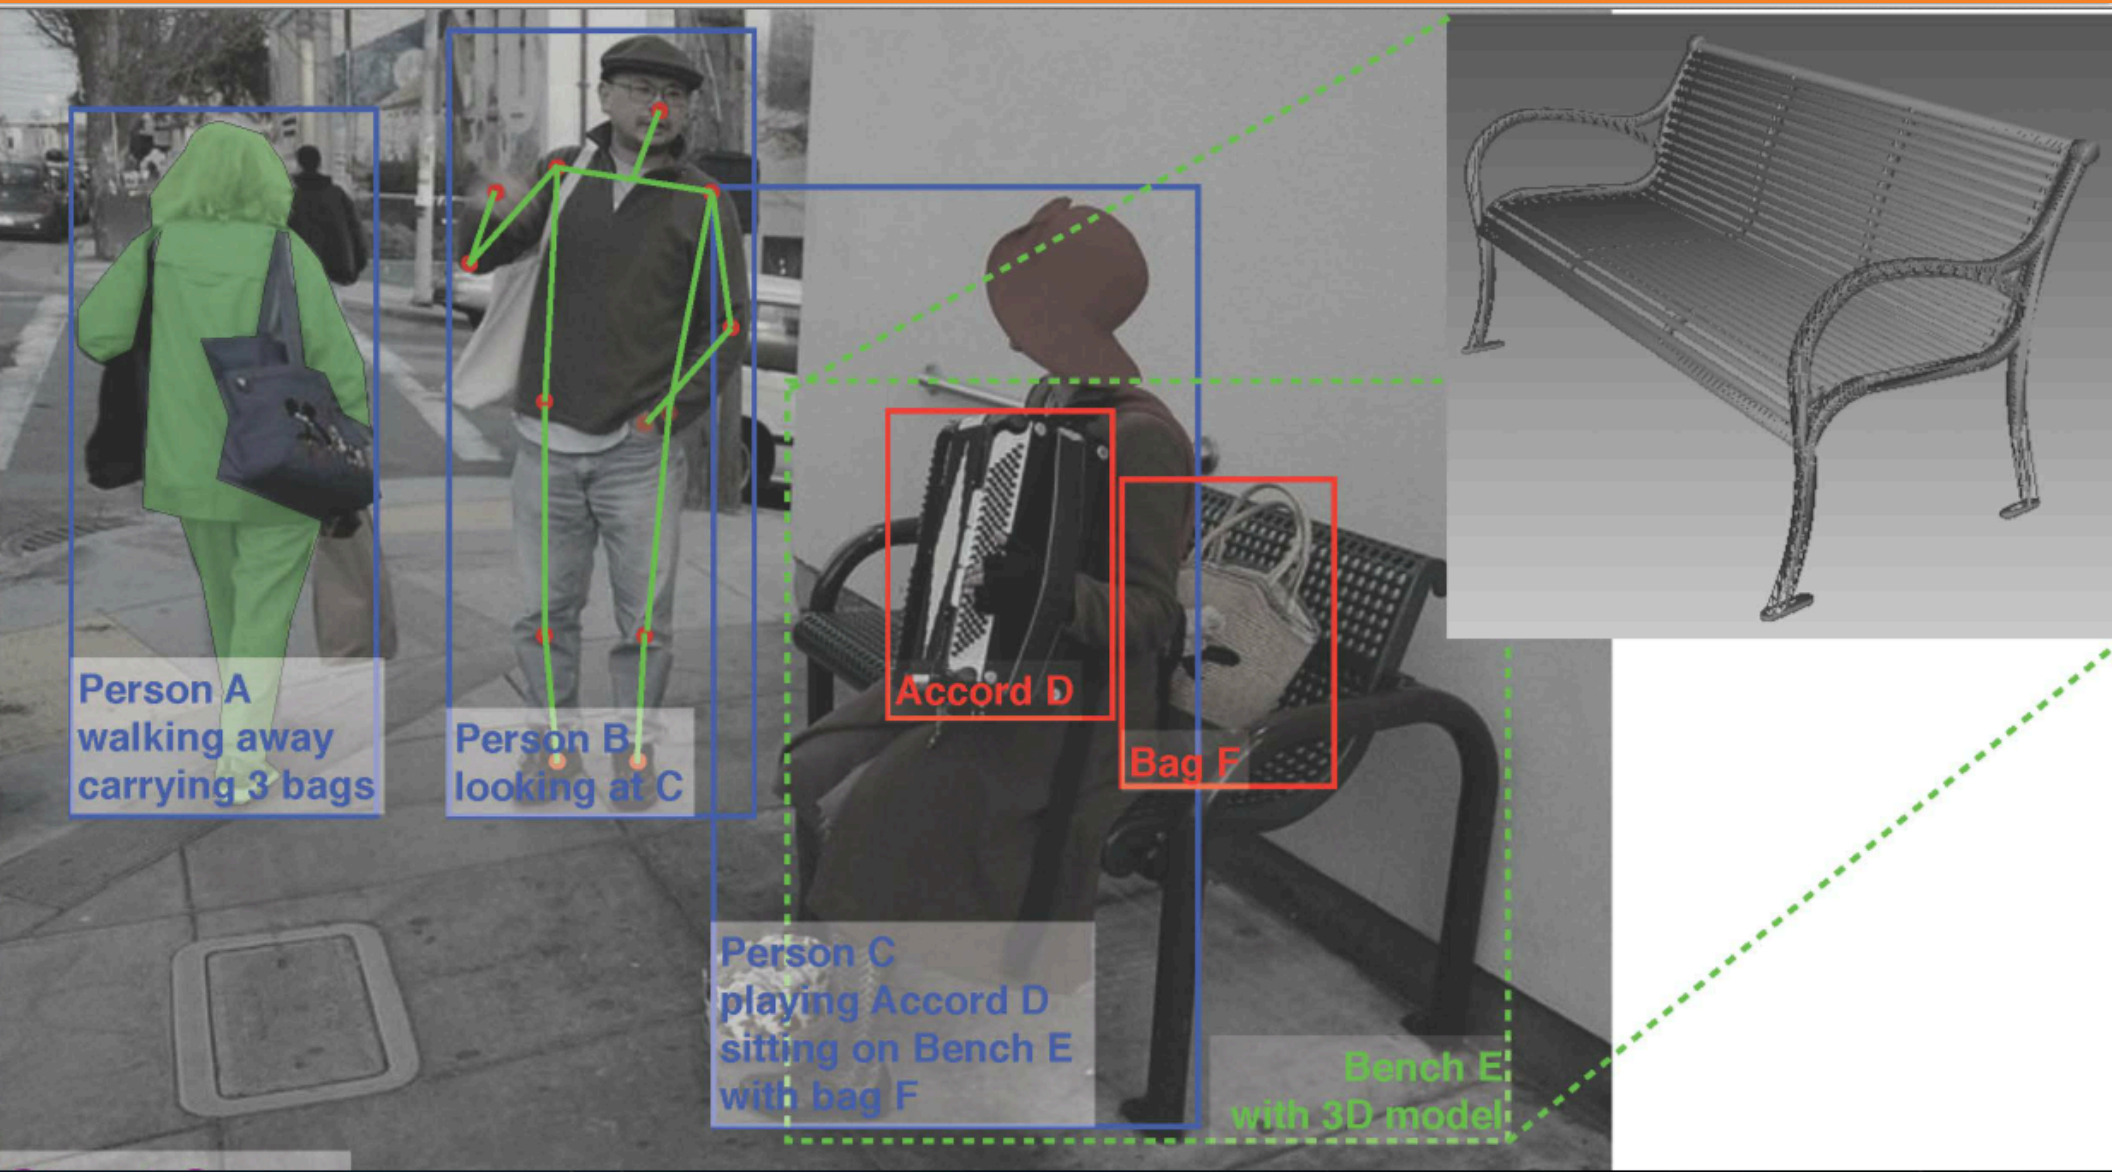

Will person B put some money into Person C's tip bag?

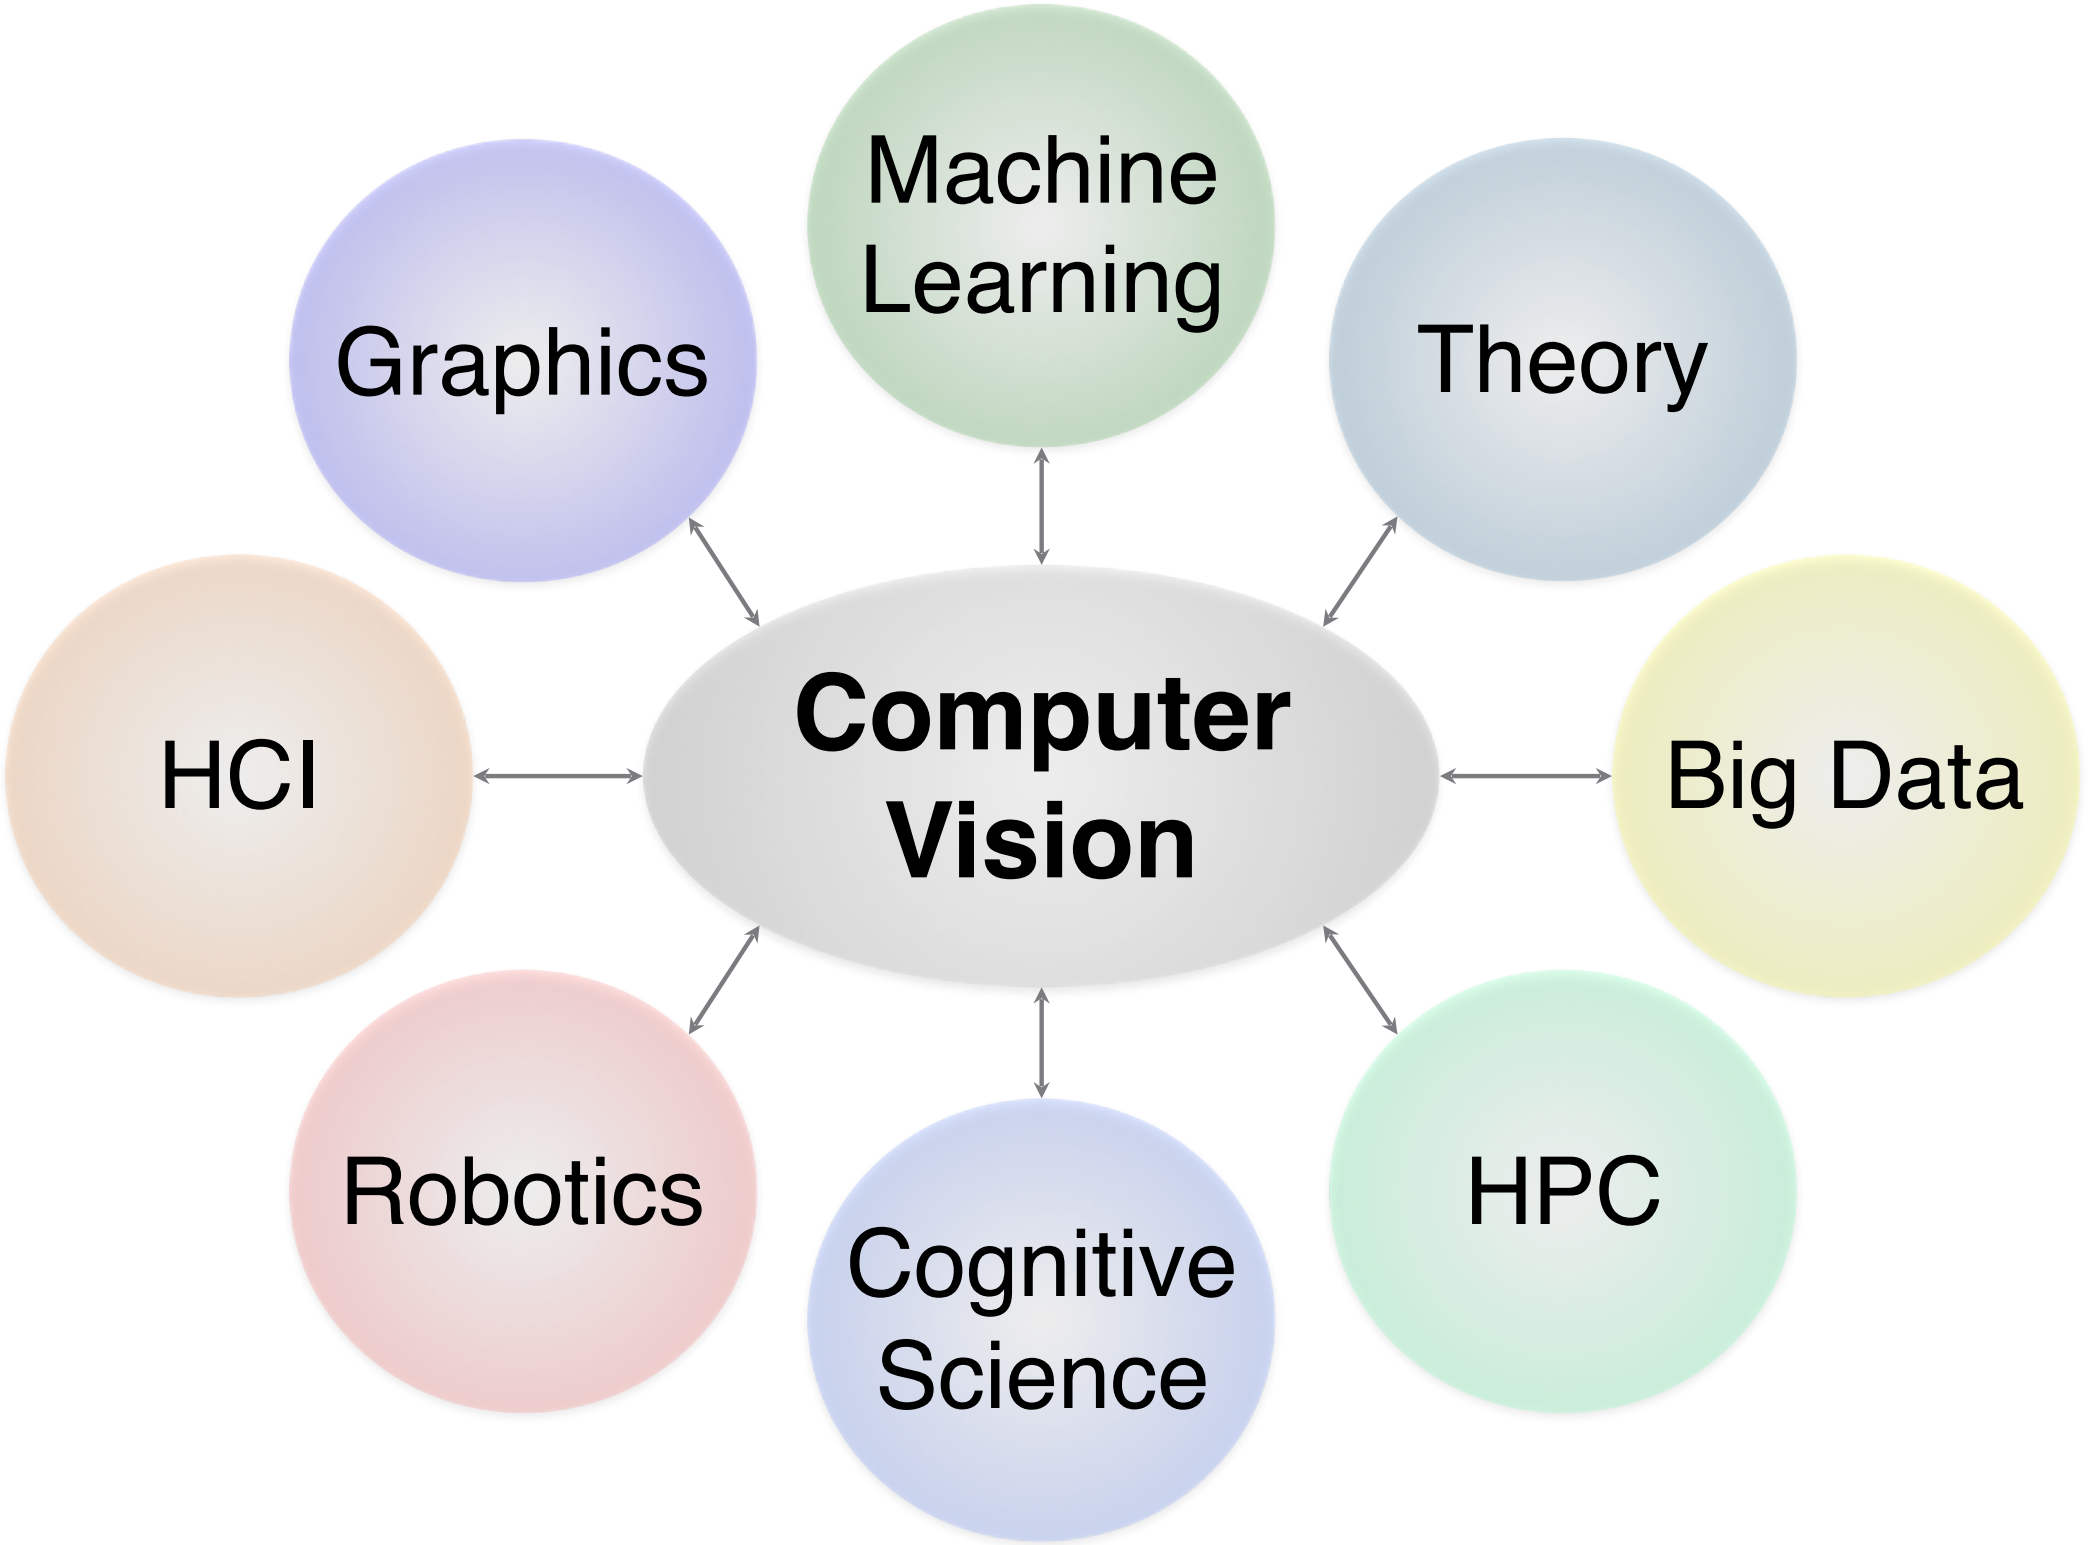
Machine Learning
Graphics
Theory
HCI
Computer Vision
Big Data
Robotics
Cognitive Science
HPC

Computer vision success stories

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".

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

PROJECT MAC

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

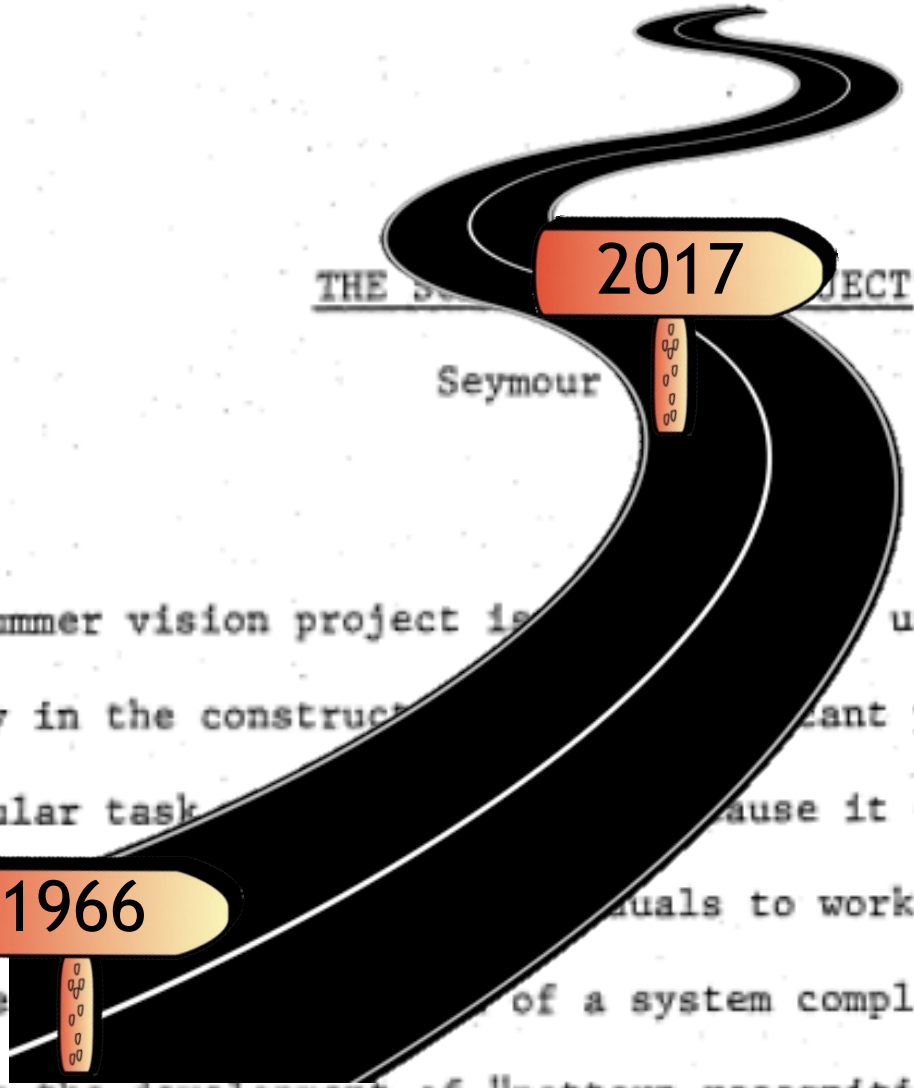

THE S ECT

Seymour

The summer vision project is use our summer workers effectively in the construct ant part of a visual system. The particular task ause it can be segmented into sub-prob uals to work independently and yet participate of a system complex enough to be a real landmark in the development of "pattern recognition".

A Brief History of Computer Vision

- 1966: Marvin Minsky assigns computer vision as an undergrad summer project
- 1960s: interpretation of synthetic worlds
- 1970s: interpretation of carefully selected images
- 1980s: NNs come and go; shift towards geometry and increased mathematical rigor
- 1990s: face recognition; statistical analysis
- 2000s: broader recognition; large annotated datasets available; video processing

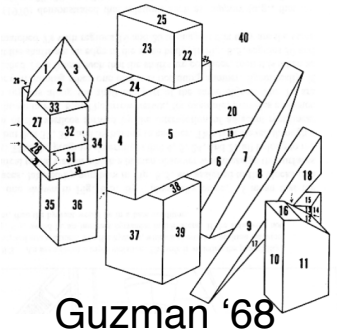

Guzman ‘68

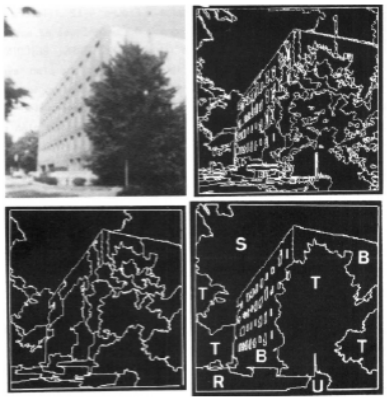

Ohta Kanade ‘78

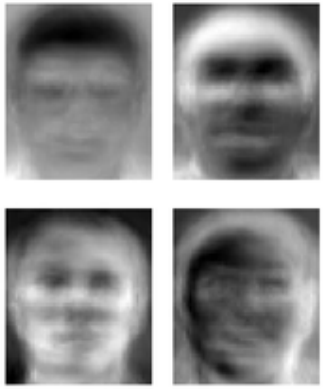

Turk and Pentland ‘91

Sorting our mail

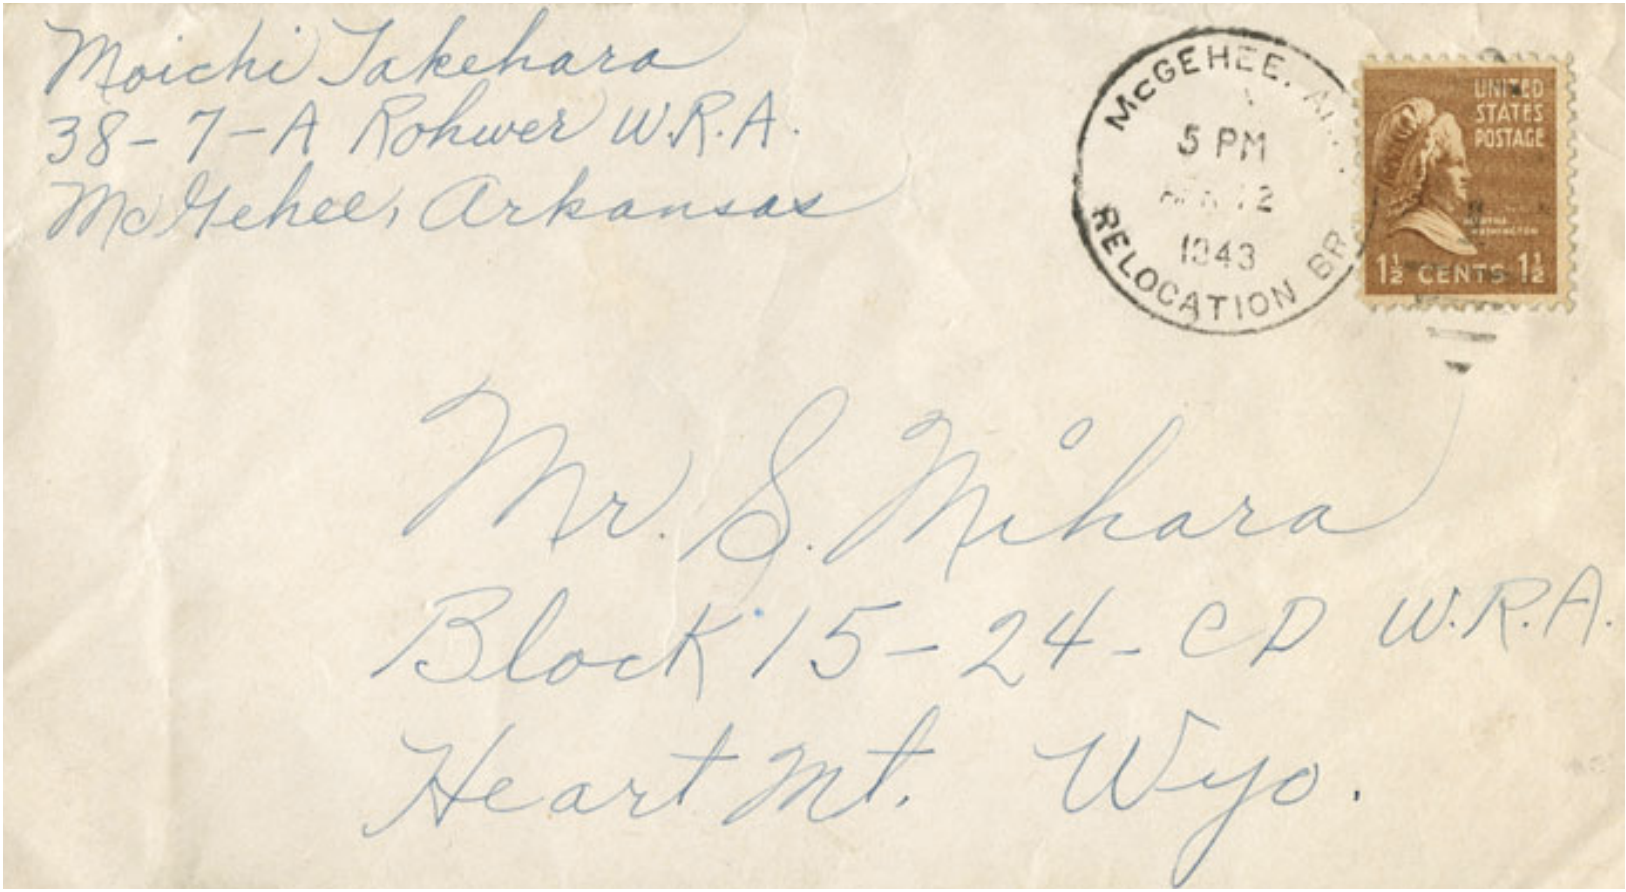

Depositing checks

Detecting (frontal) faces

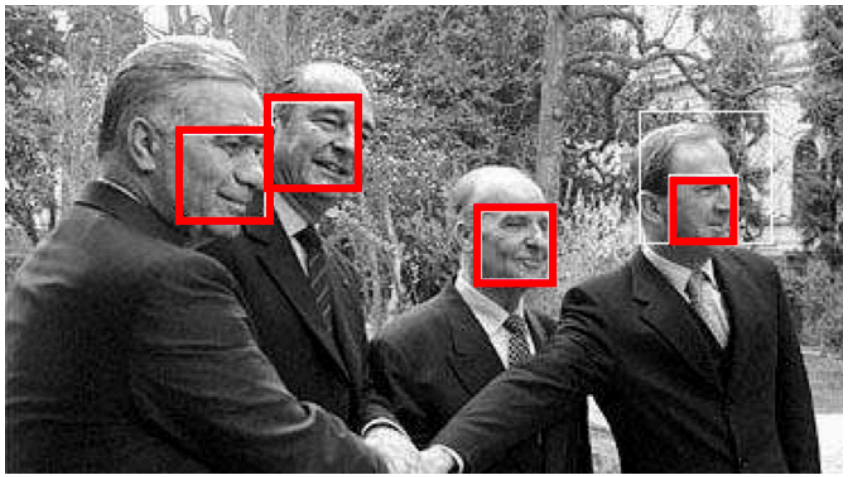

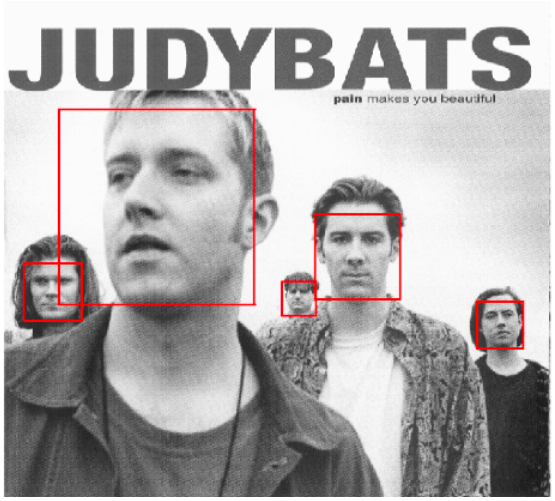

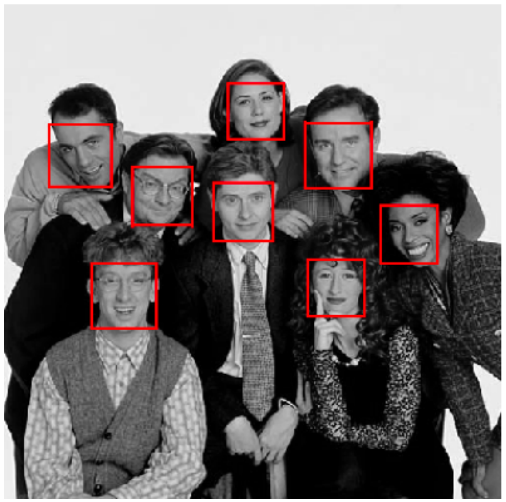

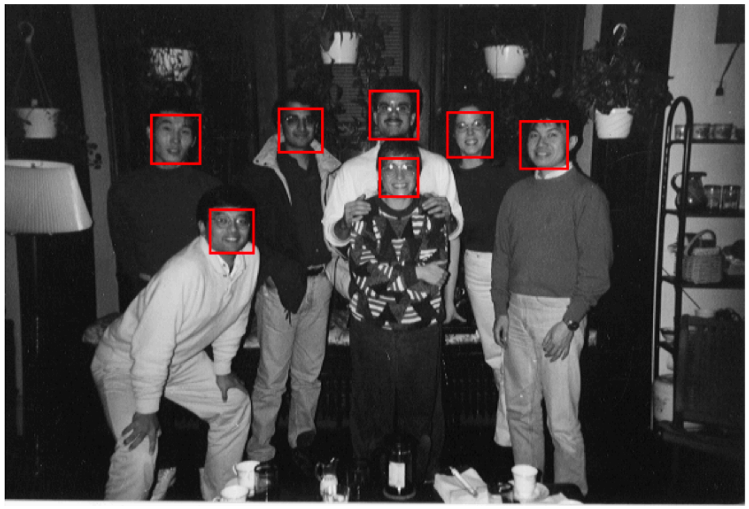

FinePix S6000fd, by Fujifilm, 2006

Viola & Jones, 2001

Face Recognition

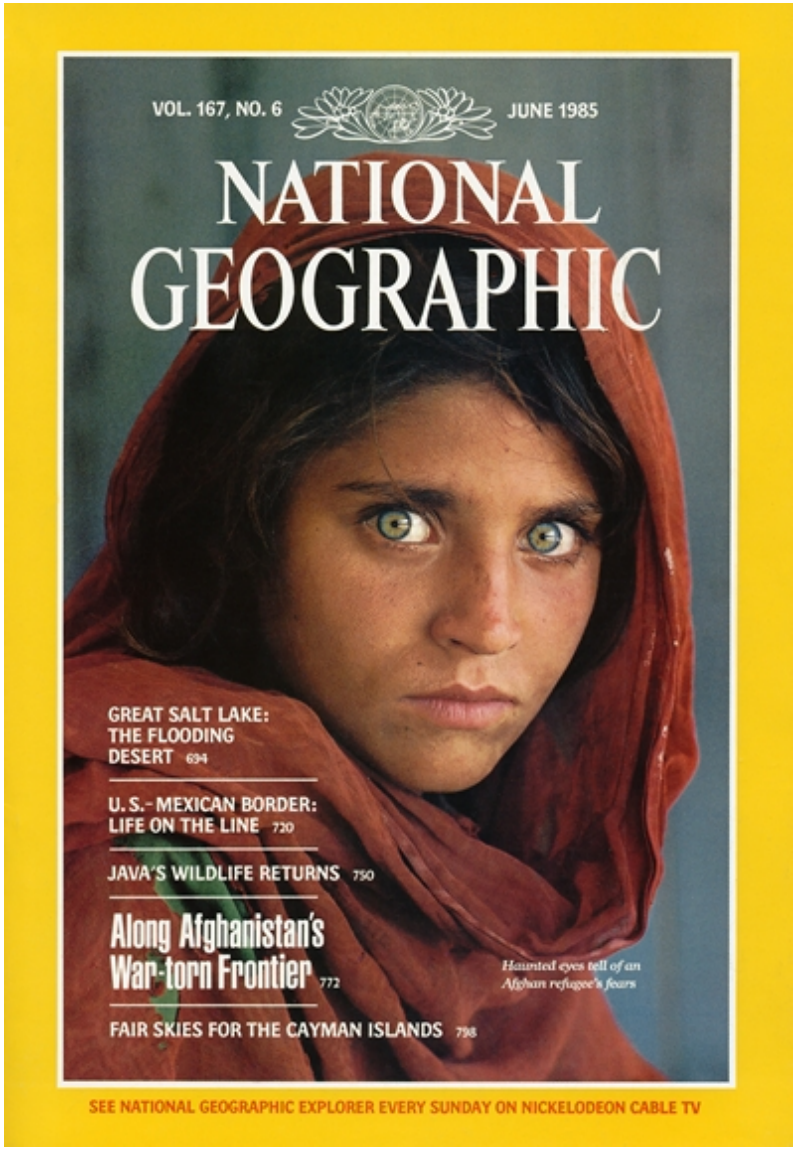

Who is she?

Slide credit: Deva Ramanan

Vision-Based Biometrics

"*How the Afghan Girl was Identified by Her Iris Patterns*" Read the story

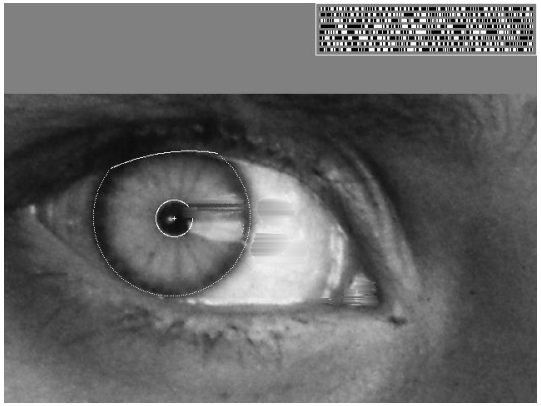

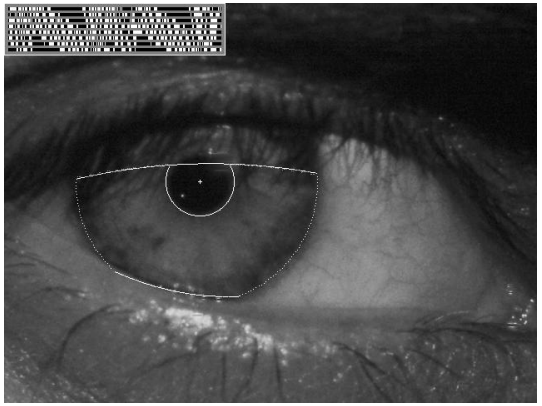

Slide credit: Deva Ramanan

Vision-based security

iPhone X

Secure Authentication

Your face is now your password. Face ID is a secure new way to unlock, authenticate, and pay.

Sports

Sportvision first down line
Nice explanation on www.howstuffworks.com

Slide credit: Deva Ramanan

3D Shape Capture for Special Effects

The Matrix movies, ESC Entertainment, XYZRGB, NRC

Slide credit: Deva Ramanan

3D Maps

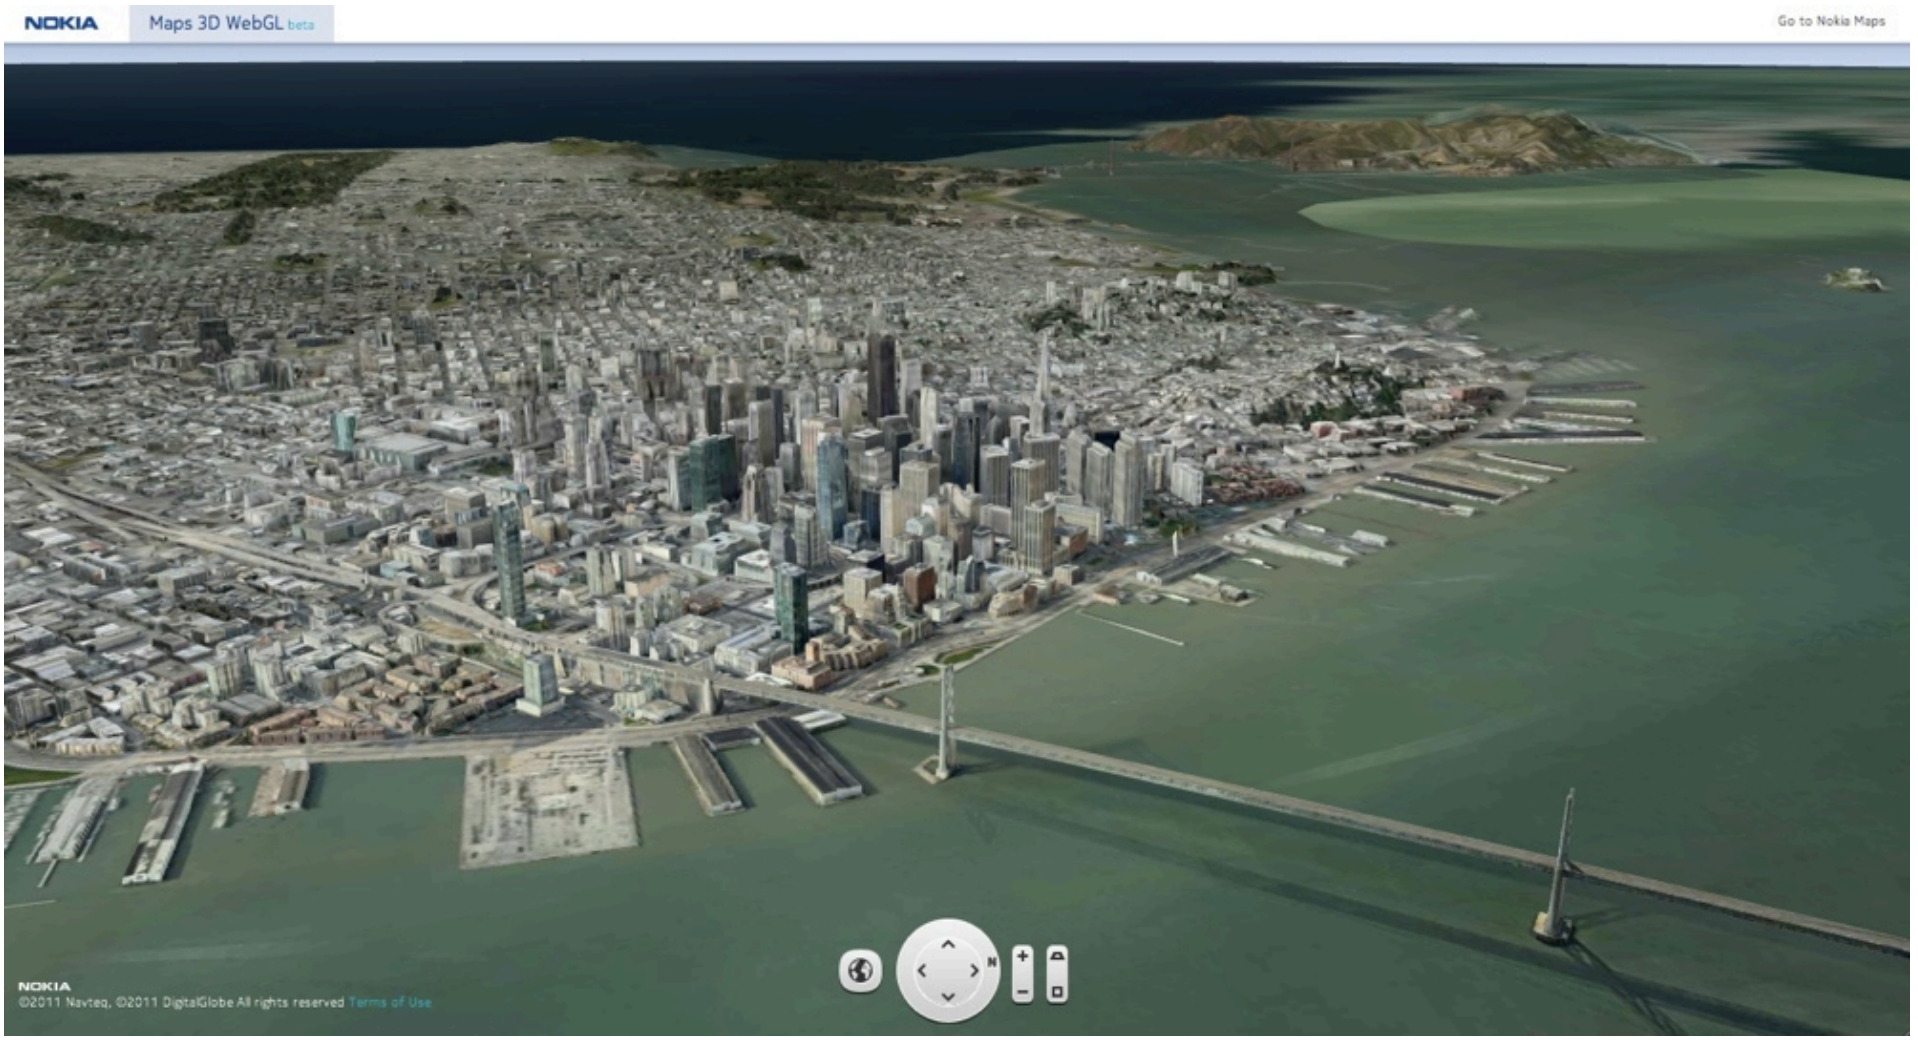

Image from Nokia's Maps 3D WebGL
(see also: Google Maps GL, Google Earth)

Slide credit: Deva Ramanan

Photo tourism

Reconstructing the 4D world
(UWashington/Microsoft)

Slide credit: Deva Ramanan

Understanding traffic patterns

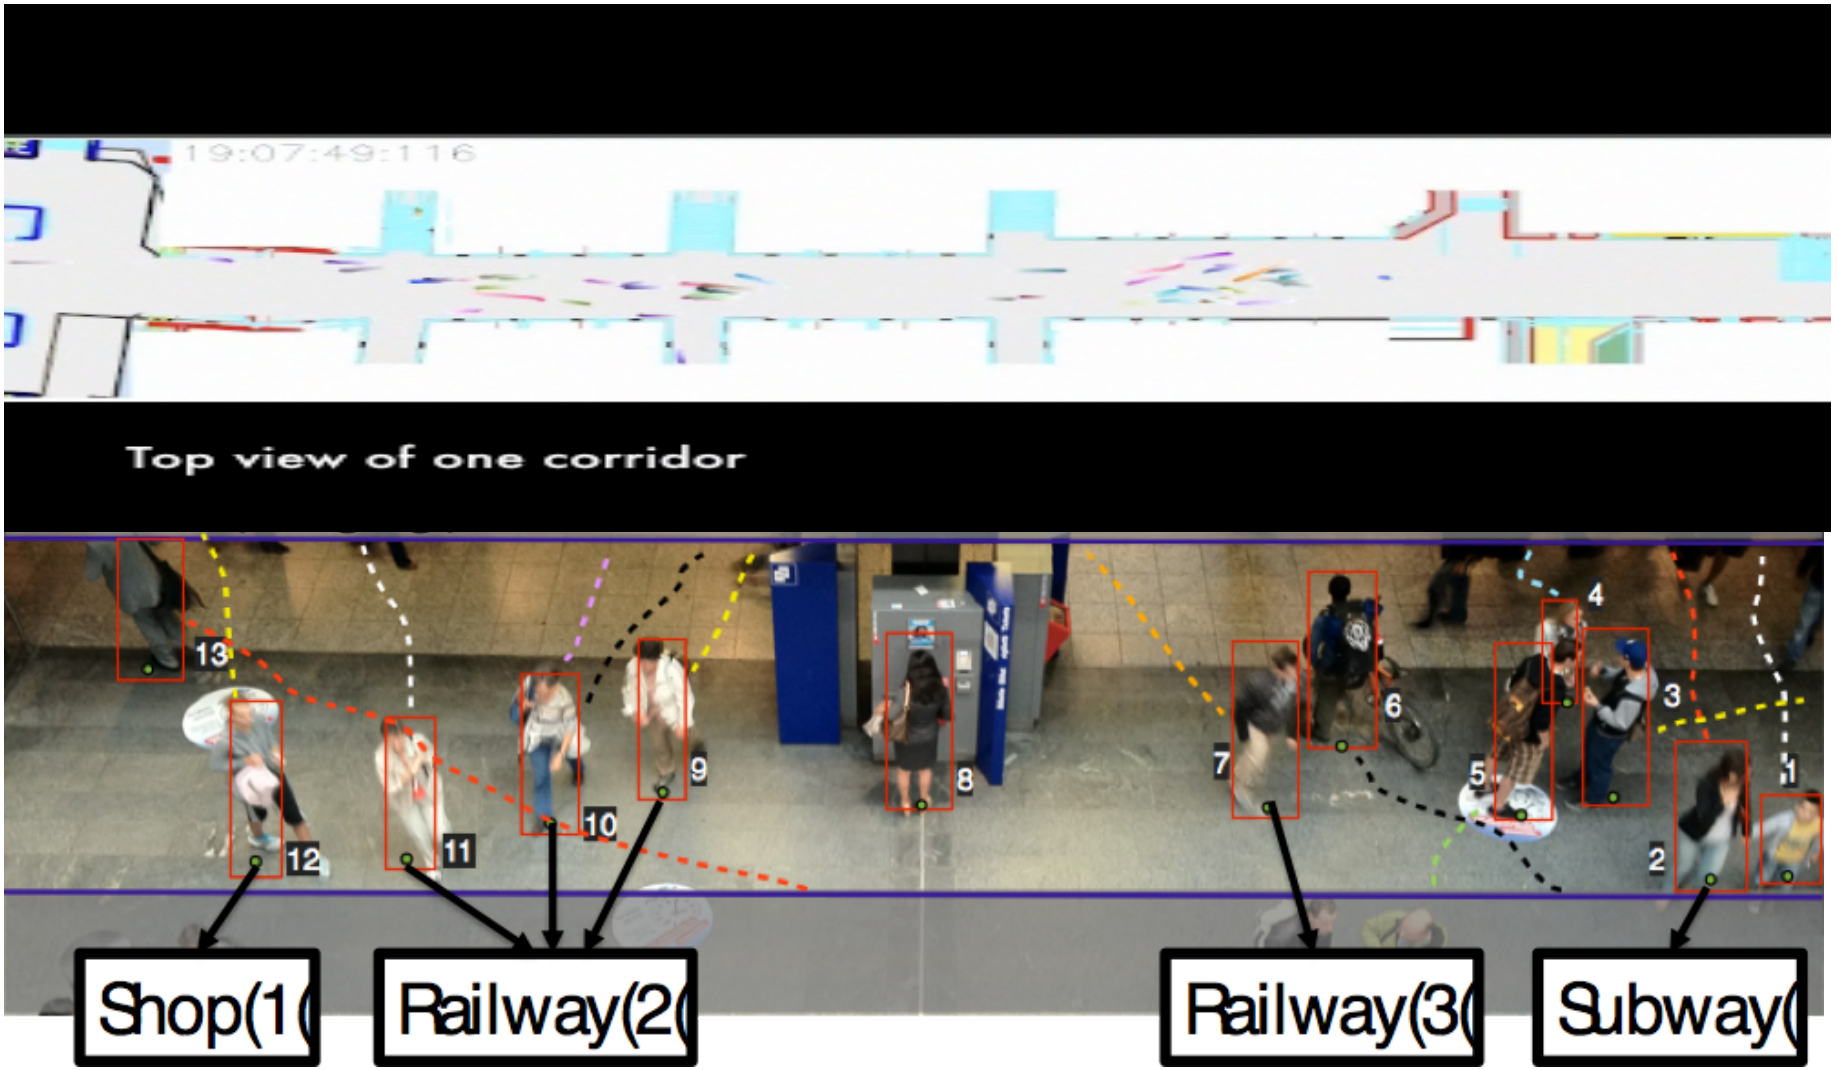

Alahi & Fei-Fei, 2014

Robotics

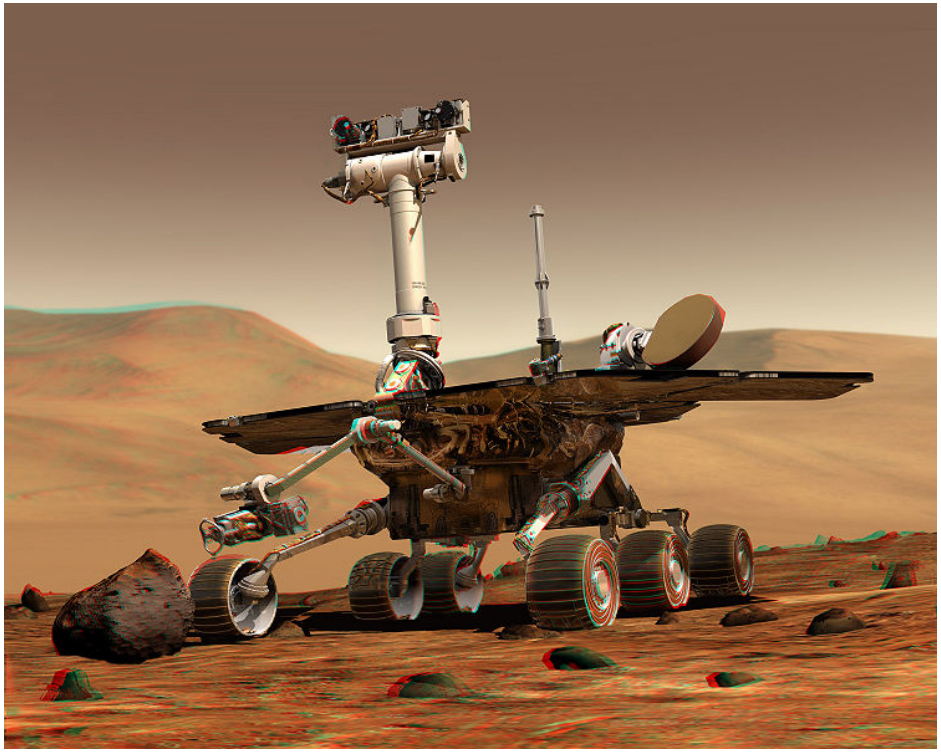

NASA's Mars Spirit Rover
http://en.wikipedia.org/wiki/Spirit_rover

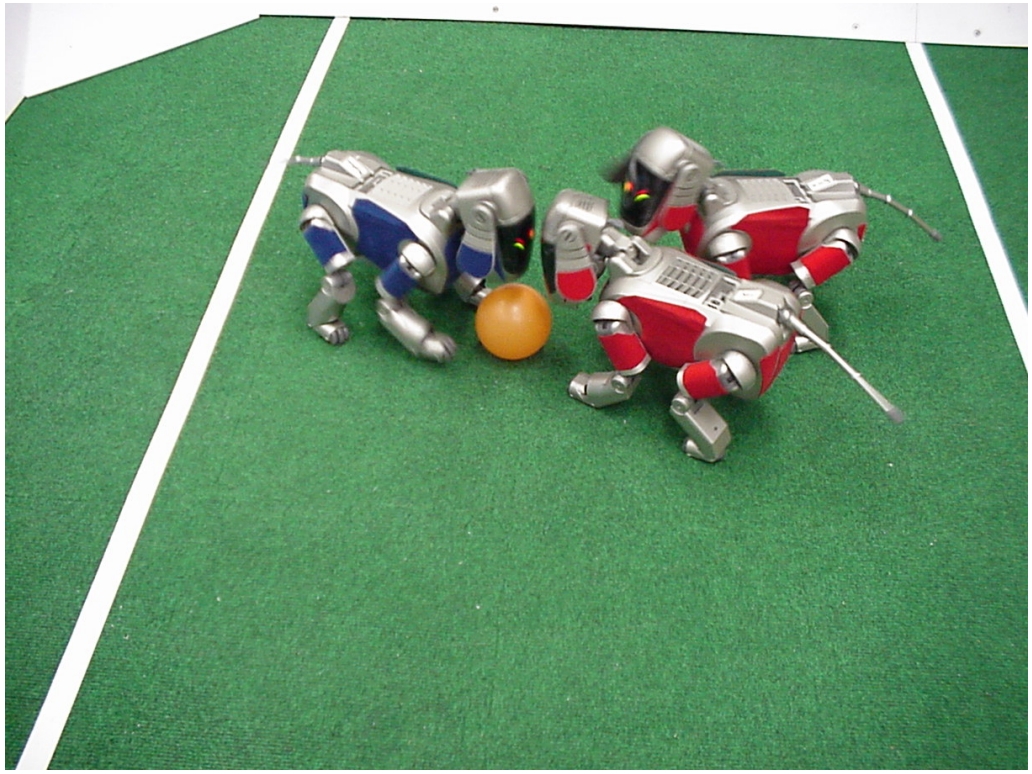

http://www.robocup.org/

Self-Driving Cars

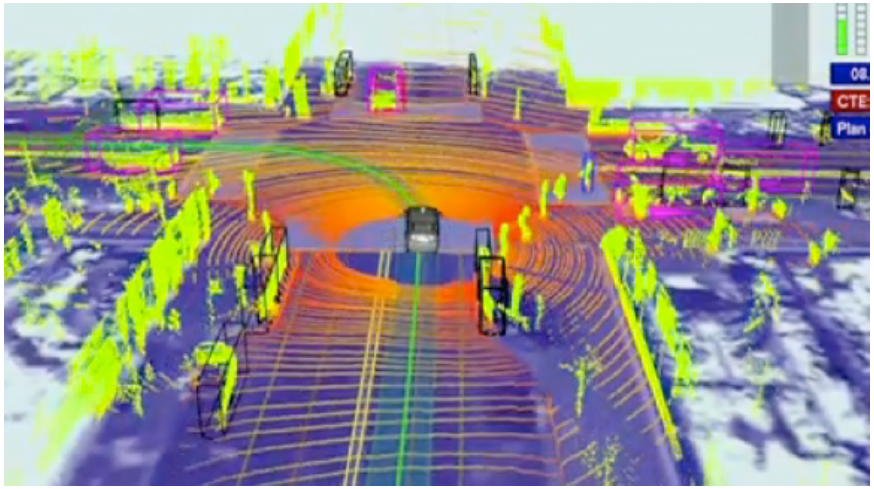

Object recognition: 1979

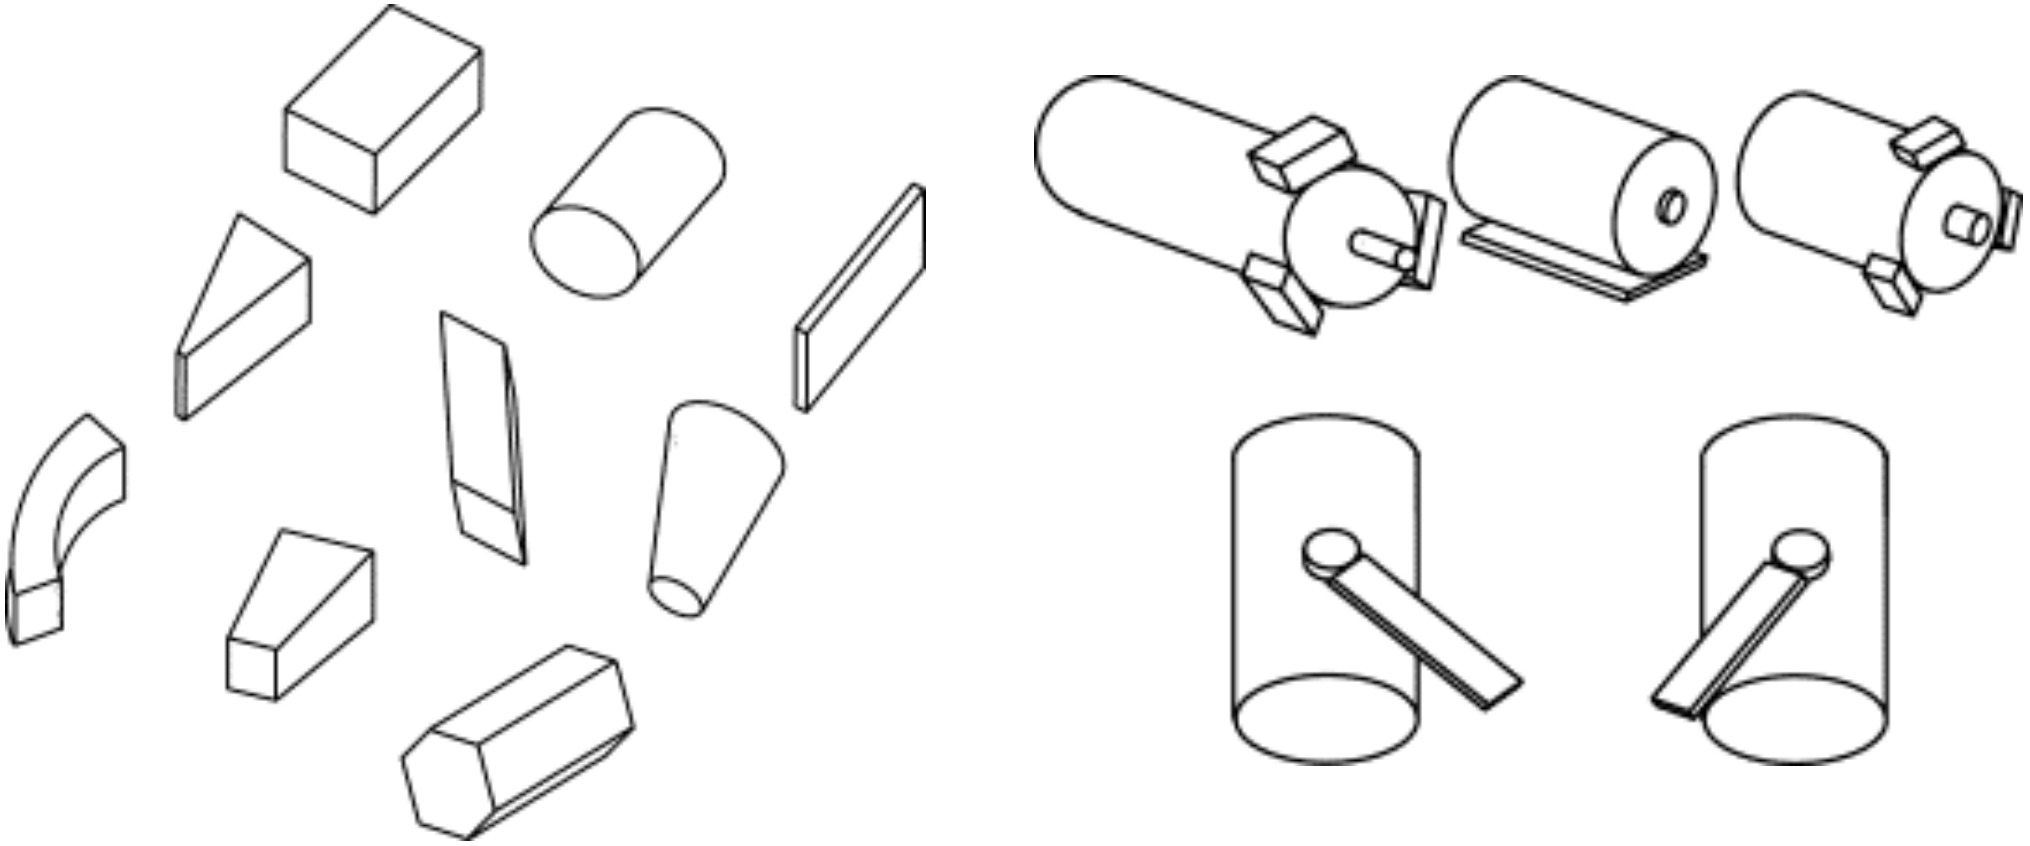

"Generalized Cylinders", by Brooks & Binford, 1979

Object recognition: 2010

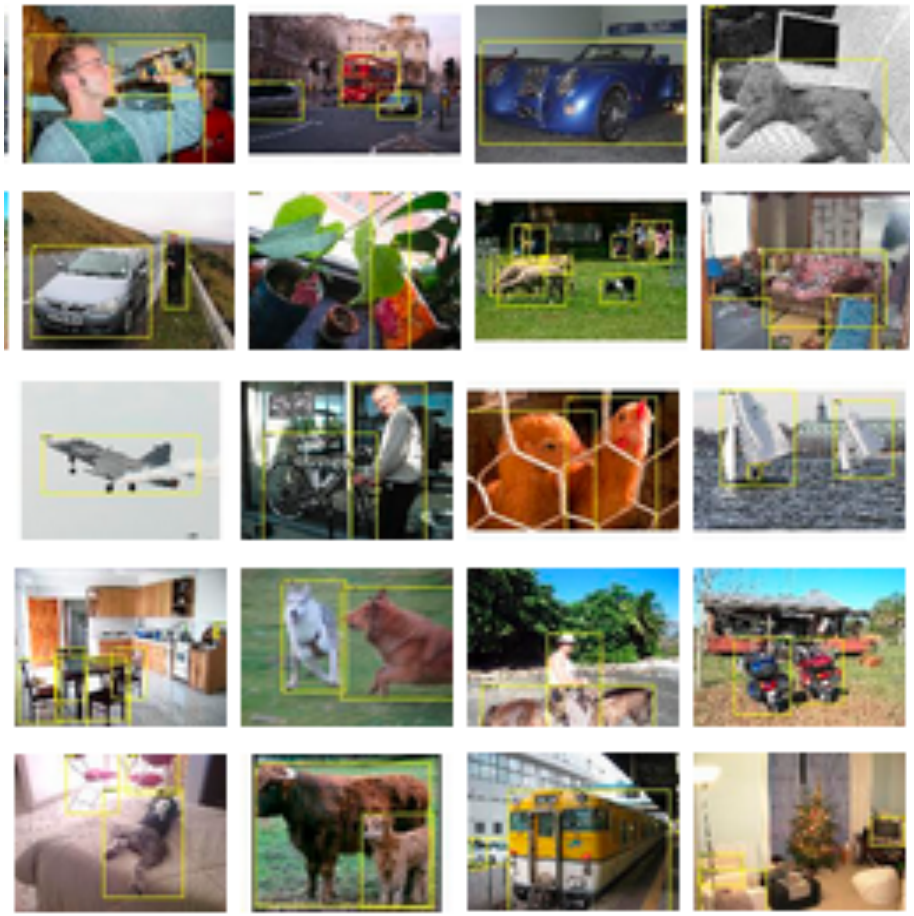

Classification: person, motorcycle

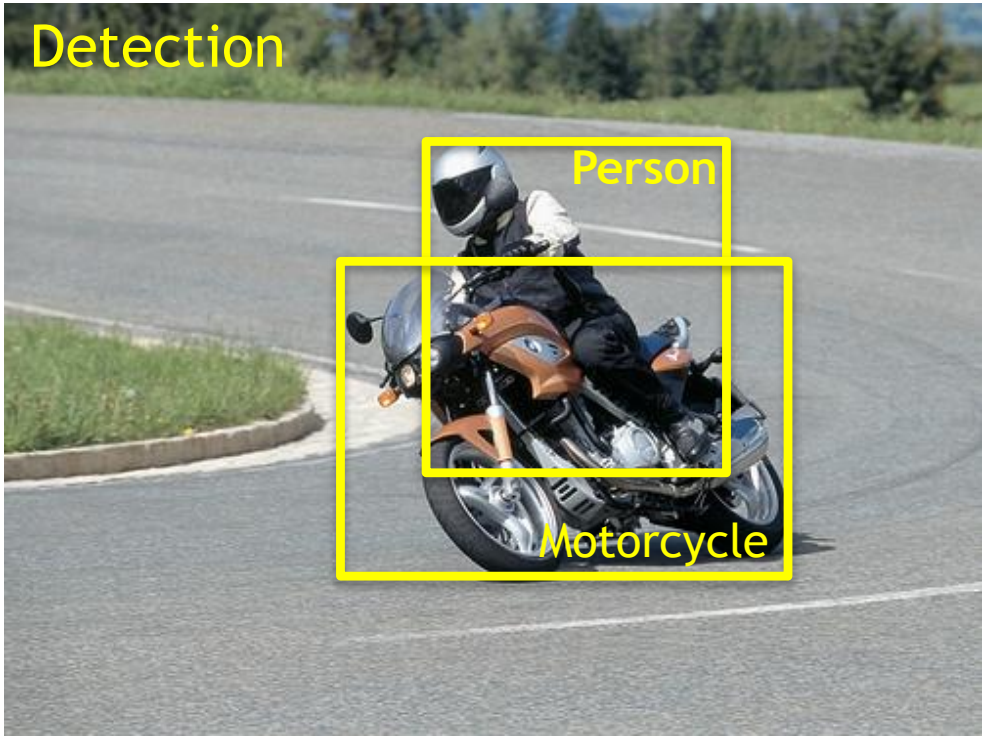

Everingham et al. IJCV 2010.

Object recognition: now

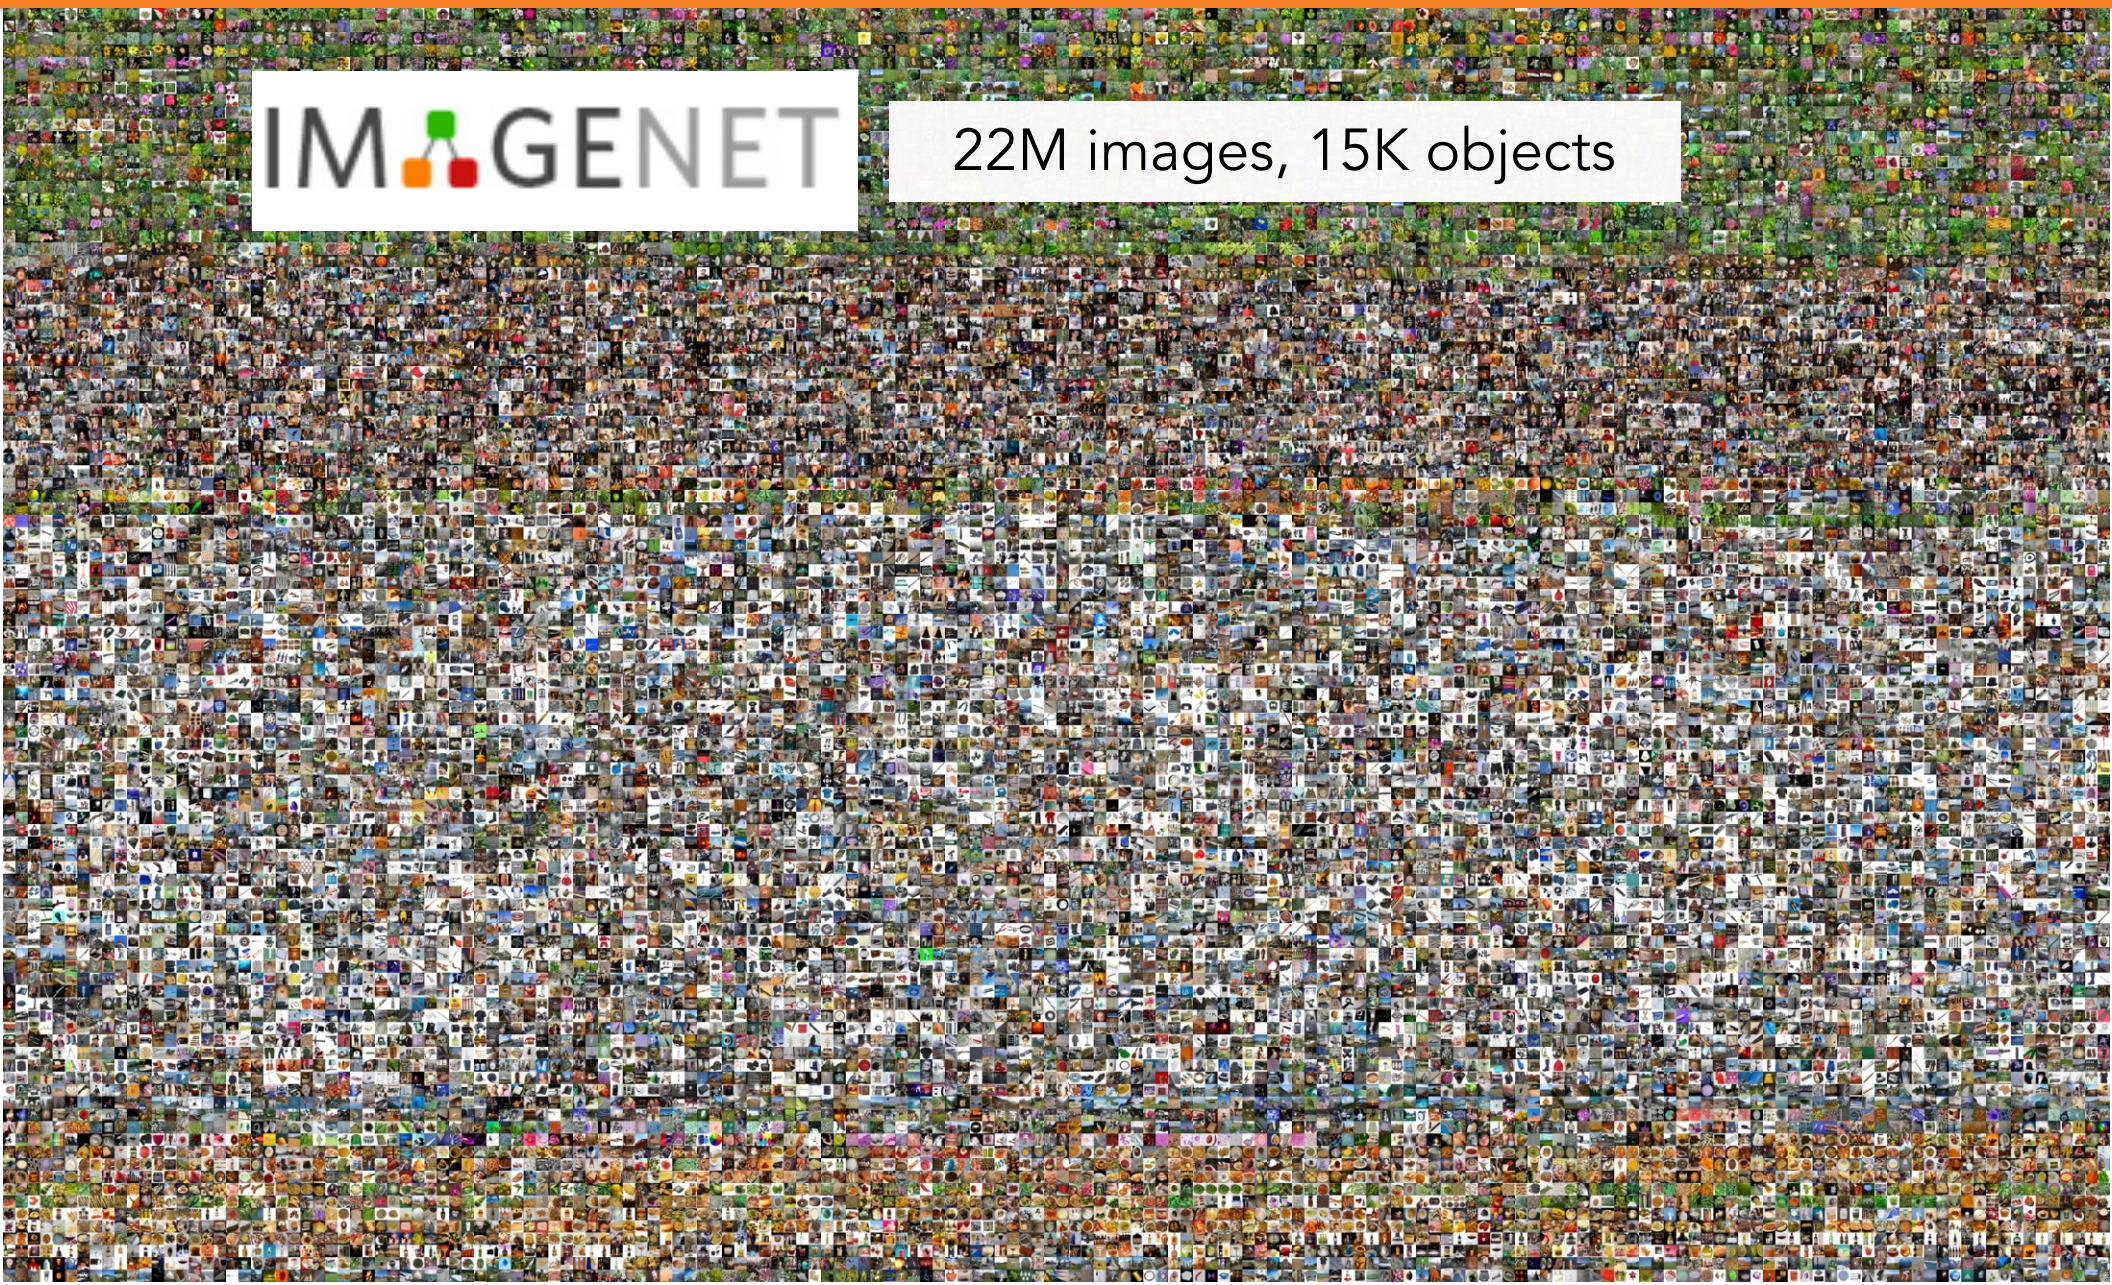

Object recognition: now

IMAGENET

(1.2M images, 1000 objects)

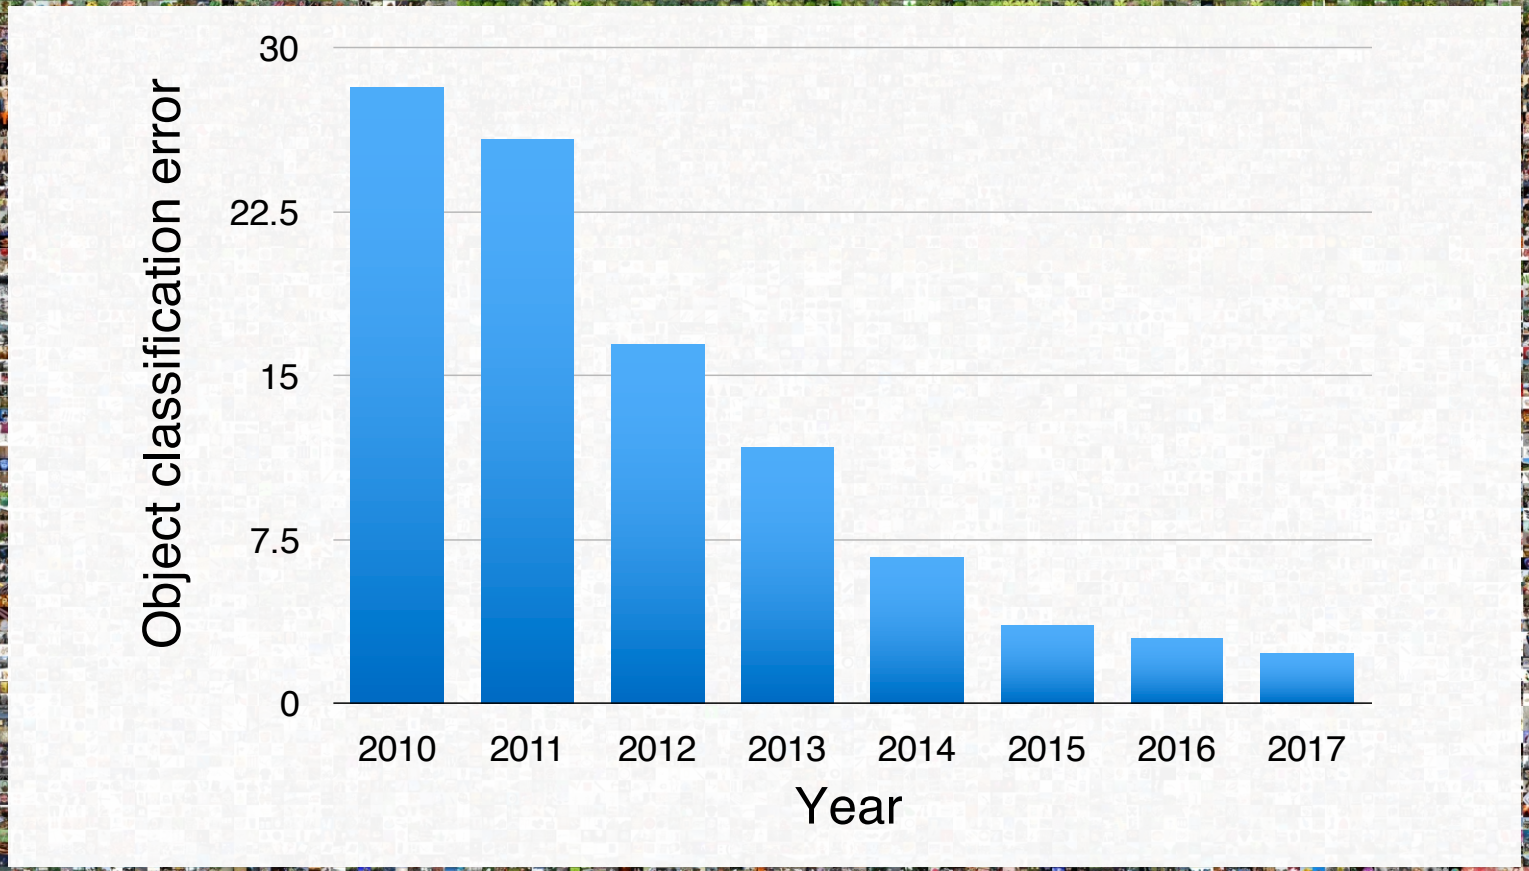

http://image-net.org

Object recognition: now

IMAGENET

(1.2M images, 1000 objects)

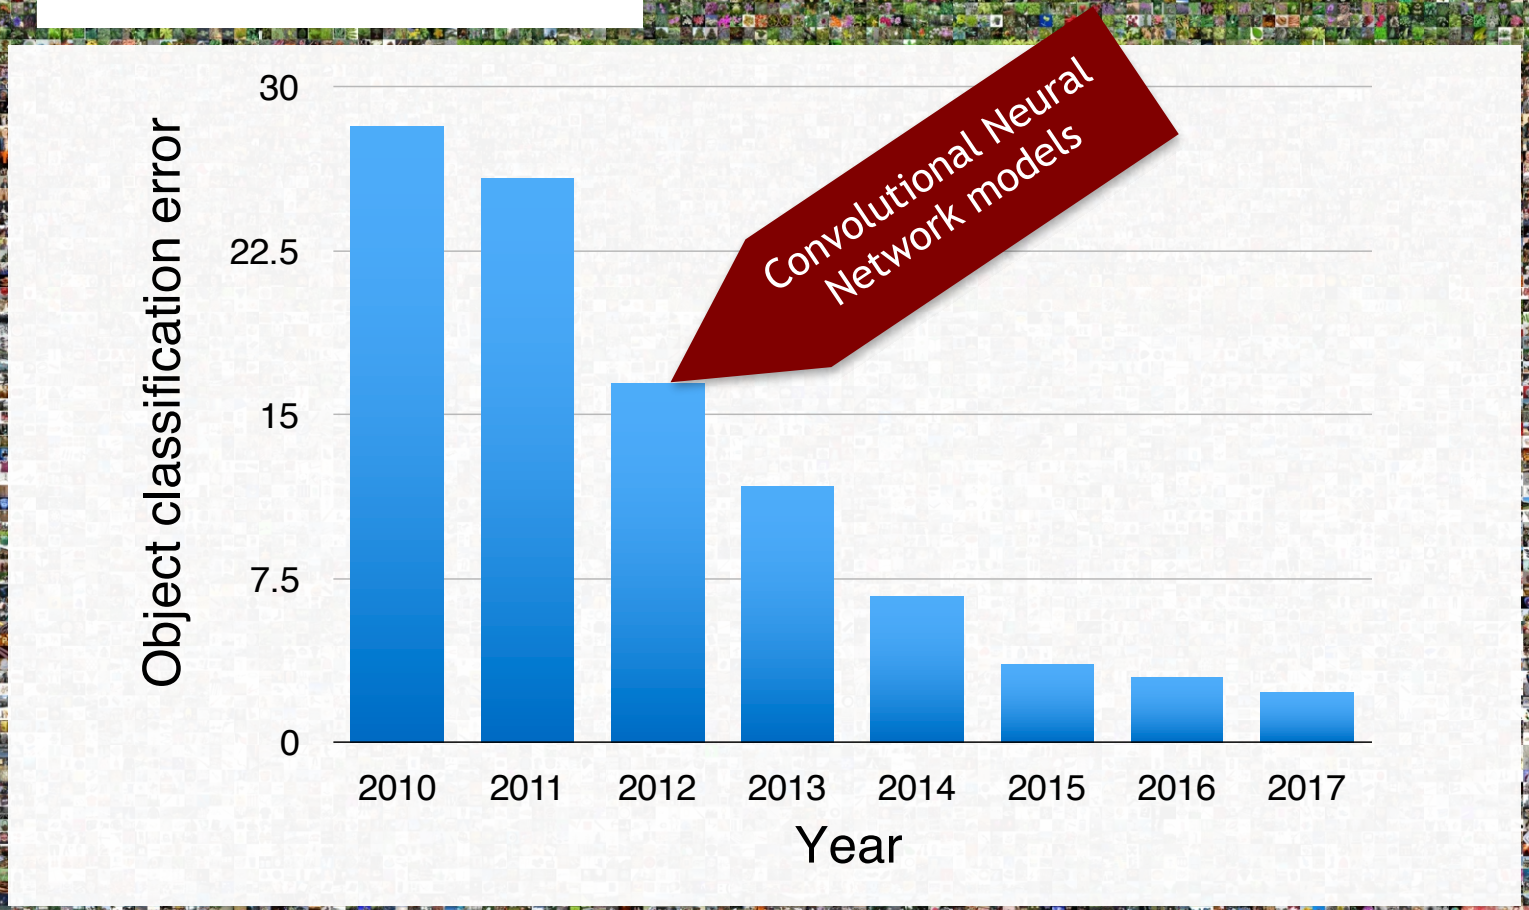

http://image-net.org

Although haven’t quite solved everything

Unconstrained environments

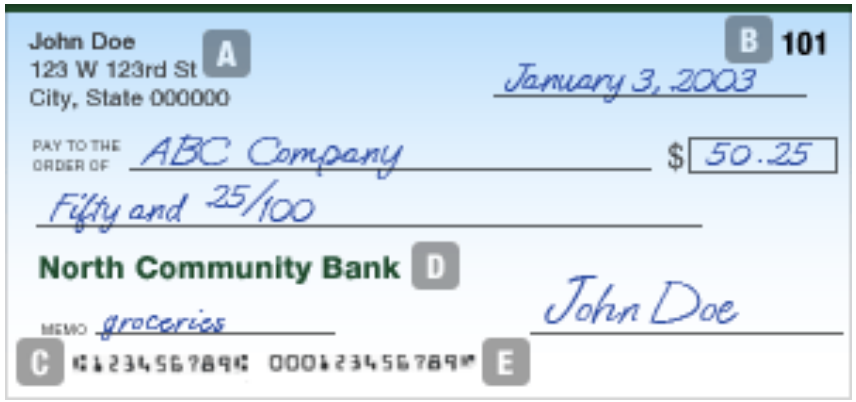

VS

Unconstrained environments

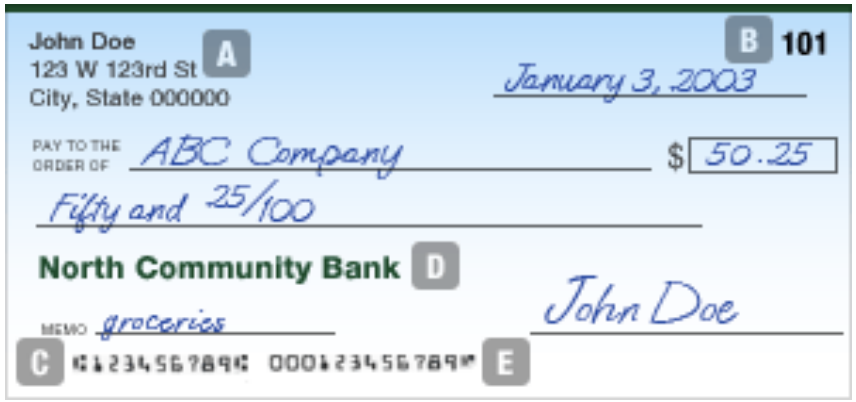

VS

VS

Pushing the limits with new tasks

Action recognition in videos

Approach	mAP	Approach	mAP
Random [43]	5.9	RGB++	15.6
C3D [54]	10.9	Two-Stream++	16.8
AlexNet [20]	11.3	Two-Stream+LSTM	17.8
IDT [57]	17.2	Two-Stream Extended	18.6
Two-Stream [44]	14.3	Ours (RGB Only)	[illegible]
		Ours	**22.4**

Table 1. Video classification results on Charades [43]. The left shows the published baselines from [43] and the right show additional new baselines. Our proposed approach outperforms all competing methods on this dataset.

[Sigurdsson CVPR 2017]

Pitfalls of looking at "overall accuracy"

Object Segmentation

Detection Leaderboard

BBOX: Dev Standard15 Chal15 Chal16 Chal17

SEGM: Dev Standard15 Chal15 Chal16 Chal17 Chal18

Small objects

Large objects

Copy to Clipboard | Export to CSV | Search:

	AP	AP^{50}	AP^{75}	AP^{S}	AP^{M}	AP^{L}	AR^{1}	AR^{10}	AR^{100}	AR^{S}	AR^{M}	AR^{L}	date
MMDet	0.486	0.730	0.530	0.339	0.52	0.602	0.368	0.593	0.632	0.464	0.665	0.777	2018-08-18
Megvii (Face++)	0.485	0.737	0.532	0.298	0.50	0.641	0.369	0.594	0.630	0.474	0.659	0.767	2018-08-18
FirstShot	0.463	0.681	0.508	0.258	0.48	0.636	0.359	0.580	0.622	0.445	0.655	0.776	2018-08-17
VISION DU	0.455	0.694	0.494	0.248	0.48	0.629	0.353	0.566	0.604	0.418	0.641	0.762	2018-08-17
DL-61	0.442	0.687	0.478	0.238	0.46	0.616	0.350	0.557	0.591	0.402	0.629	0.751	2018-08-18
BingBing	0.438	0.686	0.471	0.239	0.46	0.609	0.347	0.564	0.605	0.417	0.643	0.768	2018-08-18

http://cocodataset.org/#detections-leaderboard, Aug 30, 2019

Pitfalls of looking at "overall accuracy"

Visual Question Answering

What color are her eyes?
What is the mustache made of?

How many slices of pizza are there?
Is this a vegetarian pizza?

Rank	Participant Team	yes/no	number	other	overall	Last Submission at
1	**AIOZ**	87.96	54.99	63.28	72.61	9 months ago
2	**HDU-UCAS-USYD**	87.97	52.51	63.58	72.49	1 year ago
3	**MSRA-MSM**	87.17	55.19	62.56	71.96	7 months ago
4	**casia_iva**	86.98	51.05	62.31	71.31	1 year ago
5	**Tohoku CV Lab**	87.29	53.25	61.13	71.12	1 year ago
6	**MIL-UT**	87.00	51.65	61.62	71.06	1 year ago

https://evalai.cloudcv.org/featured-challenges/80/leaderboard/124, results as of 08/2019

Surprising failures

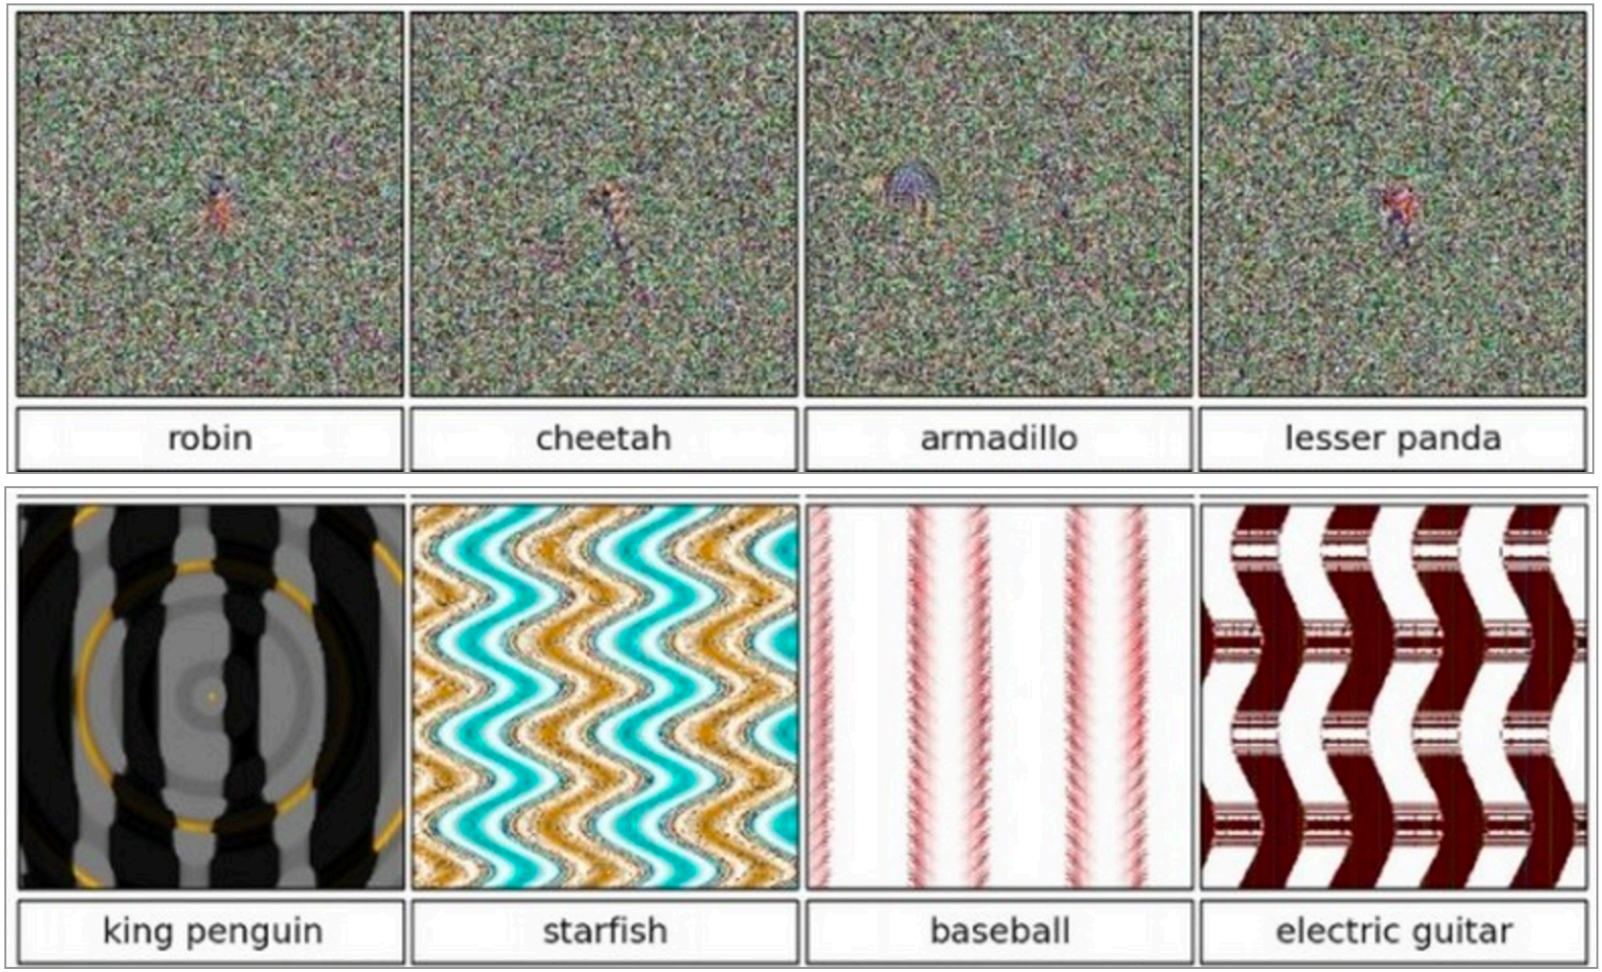

These images are classified with >99.6% confidence as the shown class by a Convolutional Network.

[Nguyen, Yosinski, Clune CVPR 2015]

http://karpathy.github.io/2015/03/30/breaking-convnets/

Surprising failures

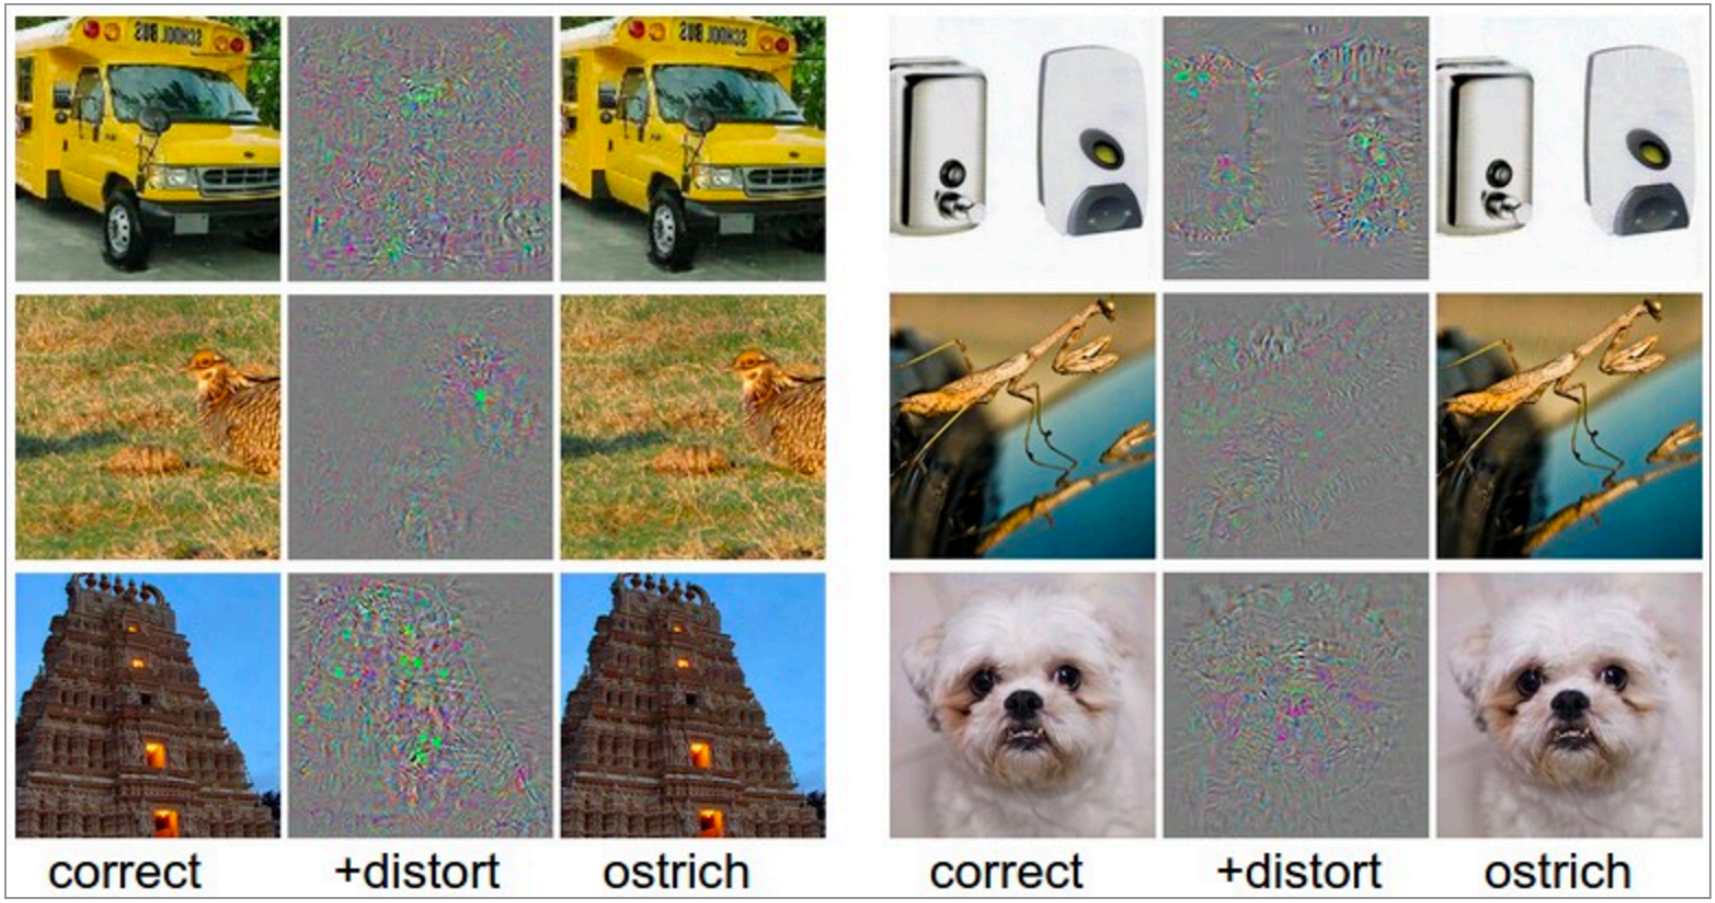

Take a correctly classified image (left image in both columns), and add a tiny distortion (middle) to fool the ConvNet with the resulting image (right).

[Szegedy et al. Intriguing properties of neural networks]

http://karpathy.github.io/2015/03/30/breaking-convnets/

(aside) Assignment 0

- See if you can fool a modern image classification model

Our plan in COS 429

Course Outline

- Image formation and capture

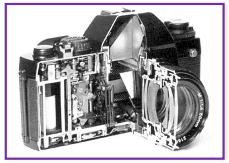

- Filtering and feature detection

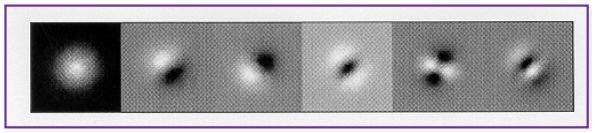

- Recognition and classification

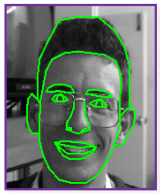

- Segmentation and clustering

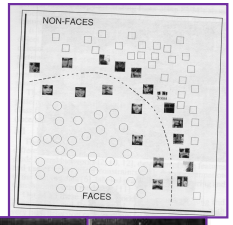

- Motion estimation and tracking

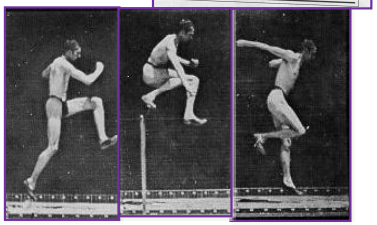

- 3D shape reconstruction

- Convolutional neural nets / deep learning

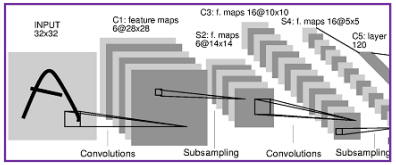

Course Staff

Instructor: Prof. Olga Russakovsky

TAs: Xingyuan Sun, Fangyin Wei, Jimmy Wu and Felix Yu

Q&A

- We will use piazza for Q&A. Please direct all non-private questions there. If emailing the Professor, consult the FAQ first.

- Feel free to answer each others' questions (we will monitor and endorse students' answers) but keep in mind the collaboration policy

Course Mechanics

- Recommended book:

Computer Vision: Algorithms and Applications

© 2010 Richard Szeliski, Microsoft Research

  - Also available online: http://szeliski.org/Book/

- Assigned papers / other readings

Course Mechanics

- 56%: 4 written / programming assignments
  - Individual or with a partner (but same partner on at most 2 assignments)
  - 4 free late days, at most 2 per assignment
- 20%: Midterm
  - Thu, Oct 24th, **no exceptions**
- 24%: Final project
  - Small groups of 1-3 people
  - Start working soon and do a cool project!
  - Milestone due Dec 13, poster session Mon, Jan 13 3-6pm
  - Writeup due on Dean's date Tue Jan 14

Python

- The assignments will be re-written to use the Python language instead of Matlab
  - Be sure to thank the TAs!

Next class: image formation and capture