EECS 442

Computer Vision

Justin Johnson & David Fouhey

Winter 2021

https://web.eecs.umich.edu/~justincj/teaching/eecs442/WI2021/

Lecture 1: Introduction

Goals of Computer Vision

Get a computer to understand

Goals of Computer Vision

Goal: Naming

Goals of Computer Vision

Goal: Naming

Goals of Computer Vision

Goal: Naming

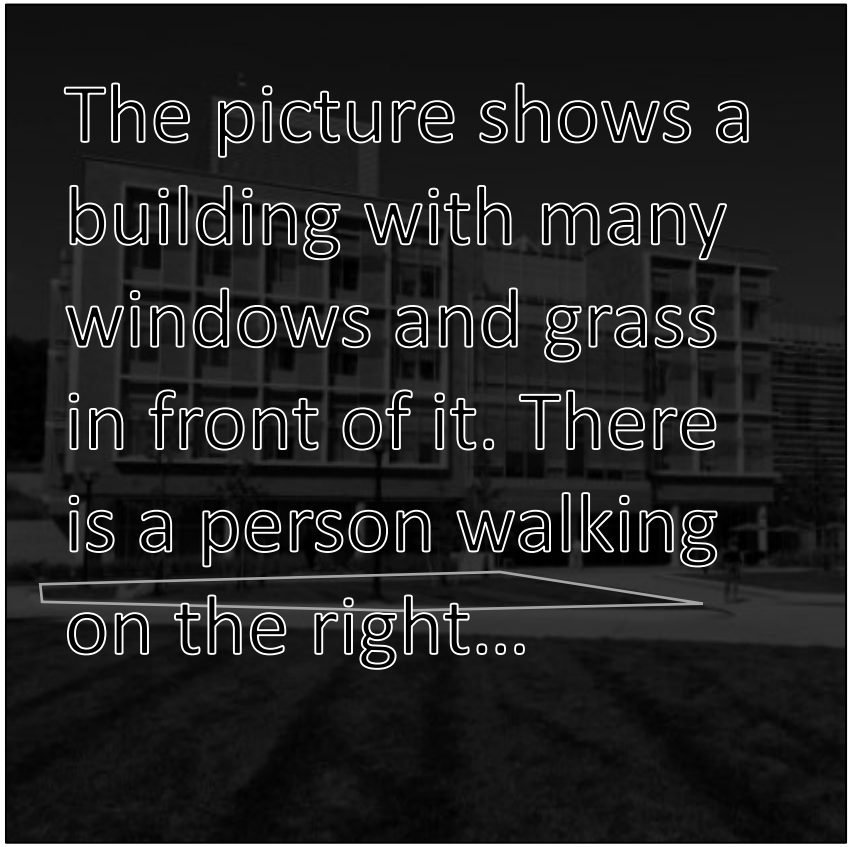

Goals of Computer Vision

Goal: 3D Structure

Goals of Computer Vision

Goal: Actions

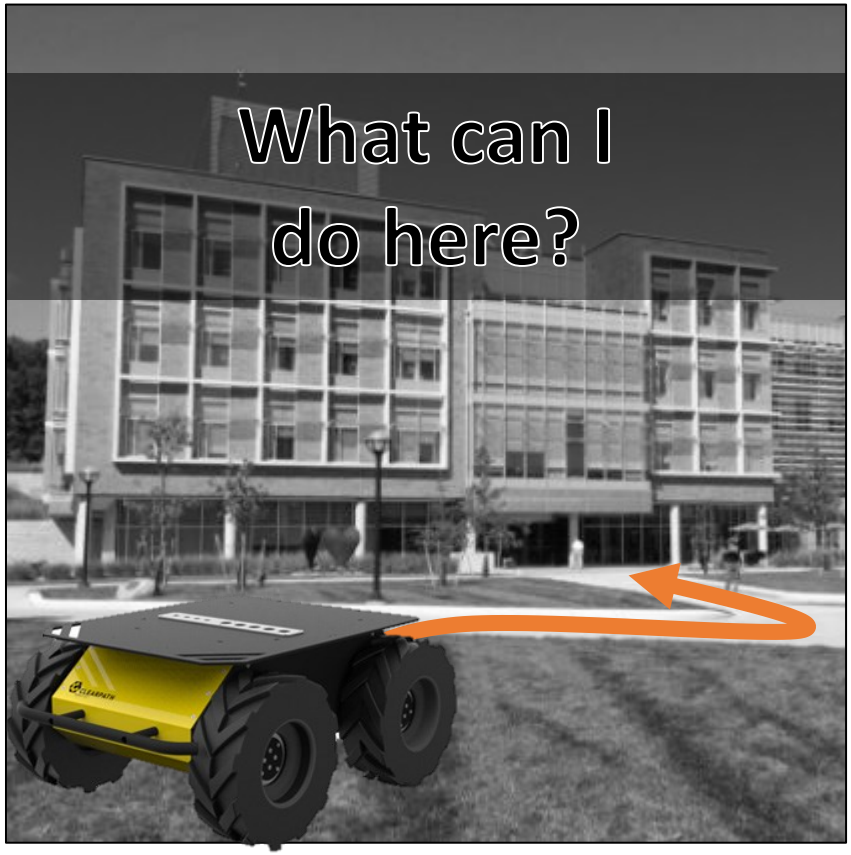

Goals of Computer Vision

Goal: Matching

Something I’m excited about…

Generating Images

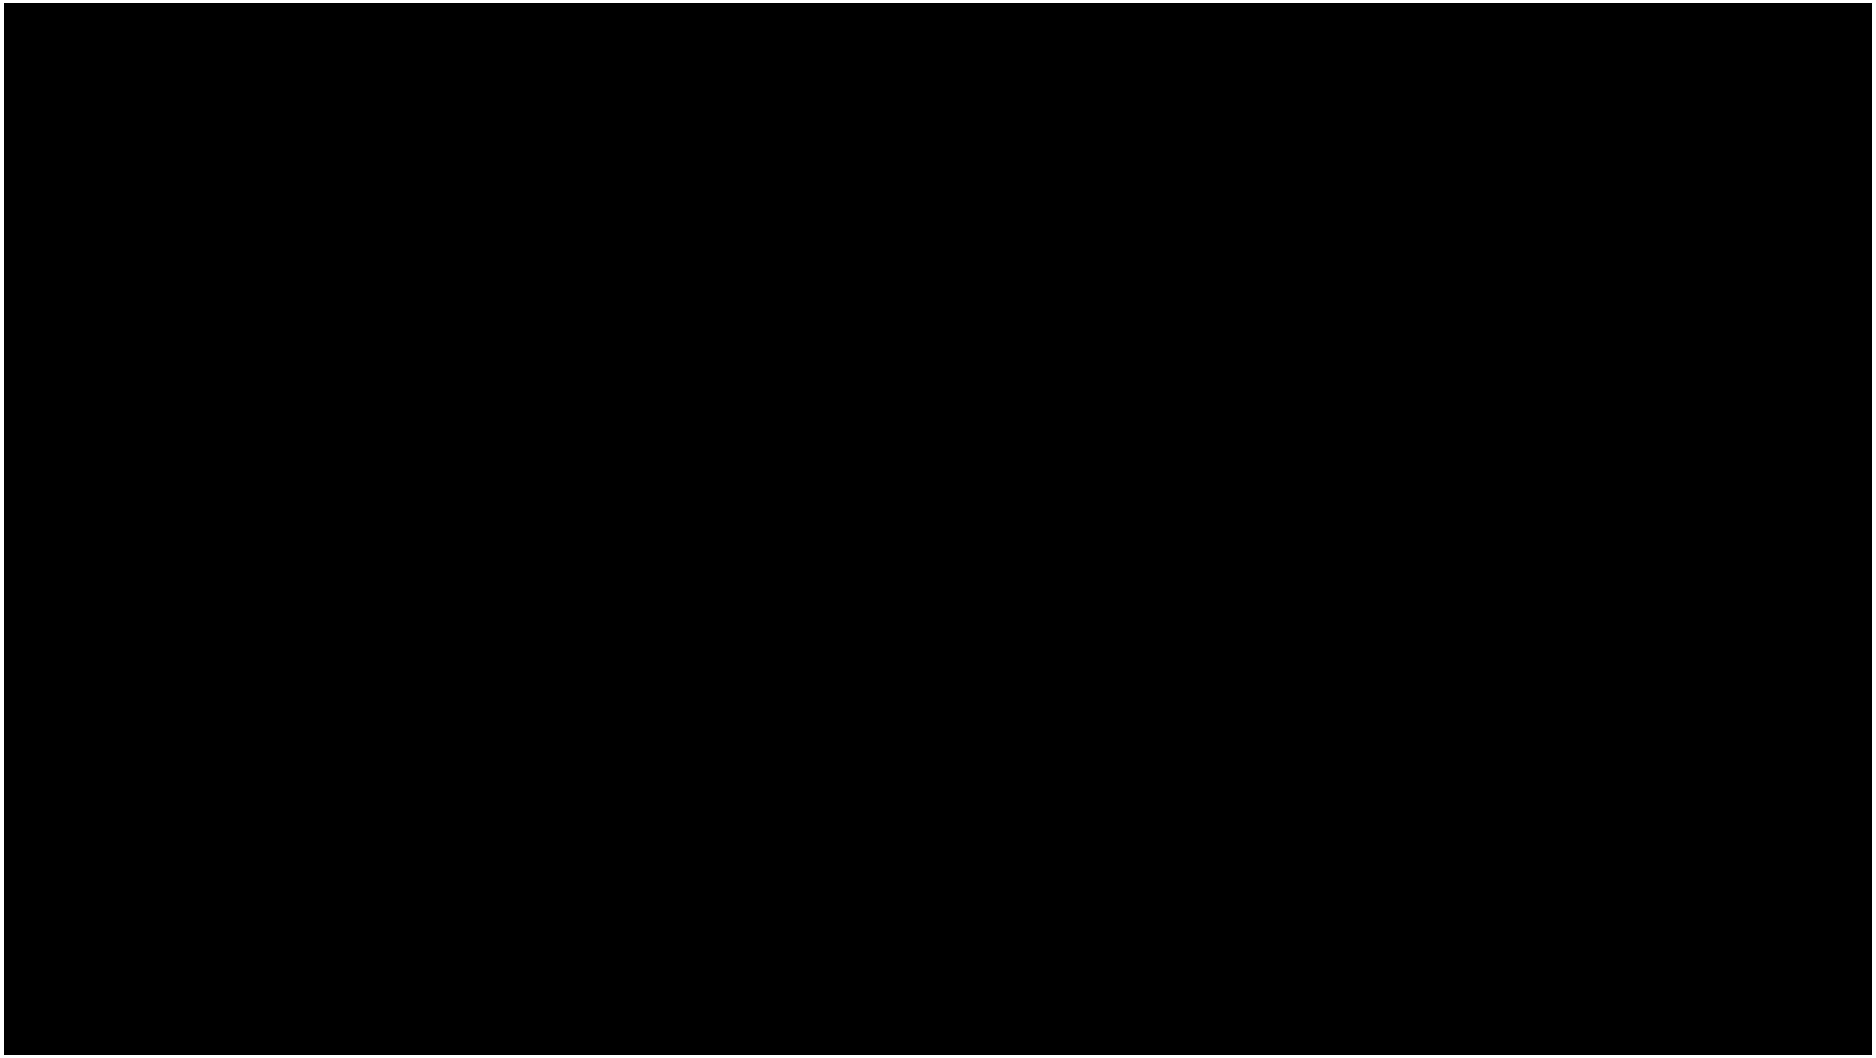

Karras et al, “Progressive Growing of GANs for Improved Quality, Stability, and Variation”, ICLR 2018

Generating Images from Language

Ramesh et al, "DALL·E: Creating Images from Text", 2021. https://openai.com/blog/dall-e/

Generating Images from Language

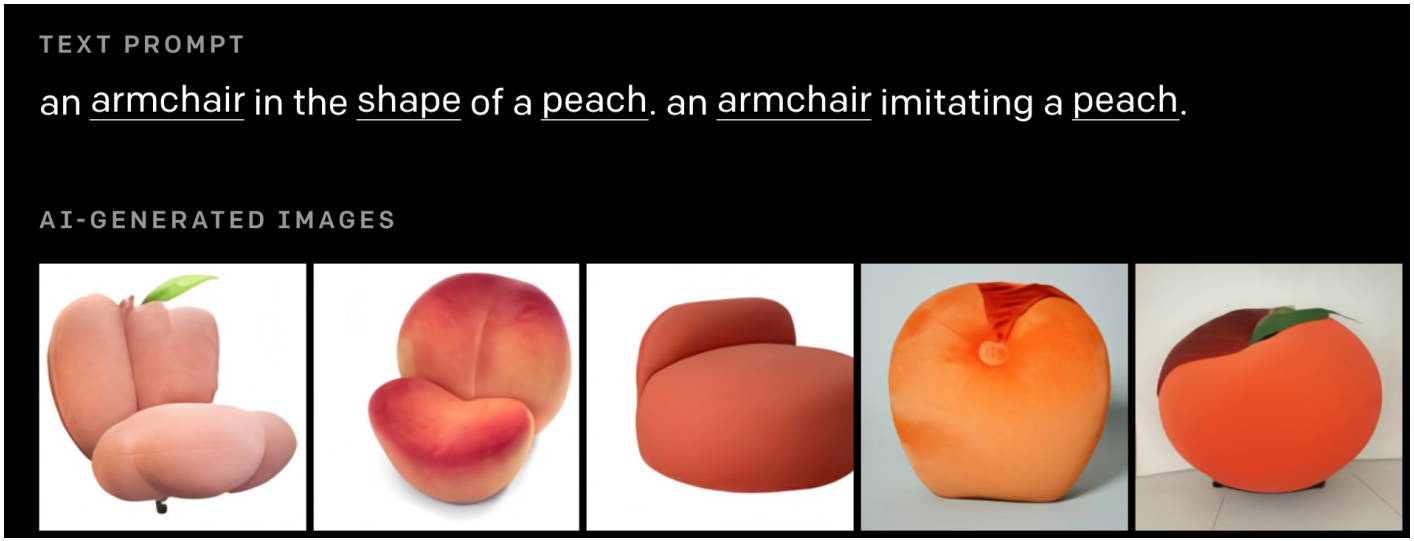

Ramesh et al, "DALL·E: Creating Images from Text", 2021. https://openai.com/blog/dall-e/

Reconstructing the 3D World

Giving People Tools

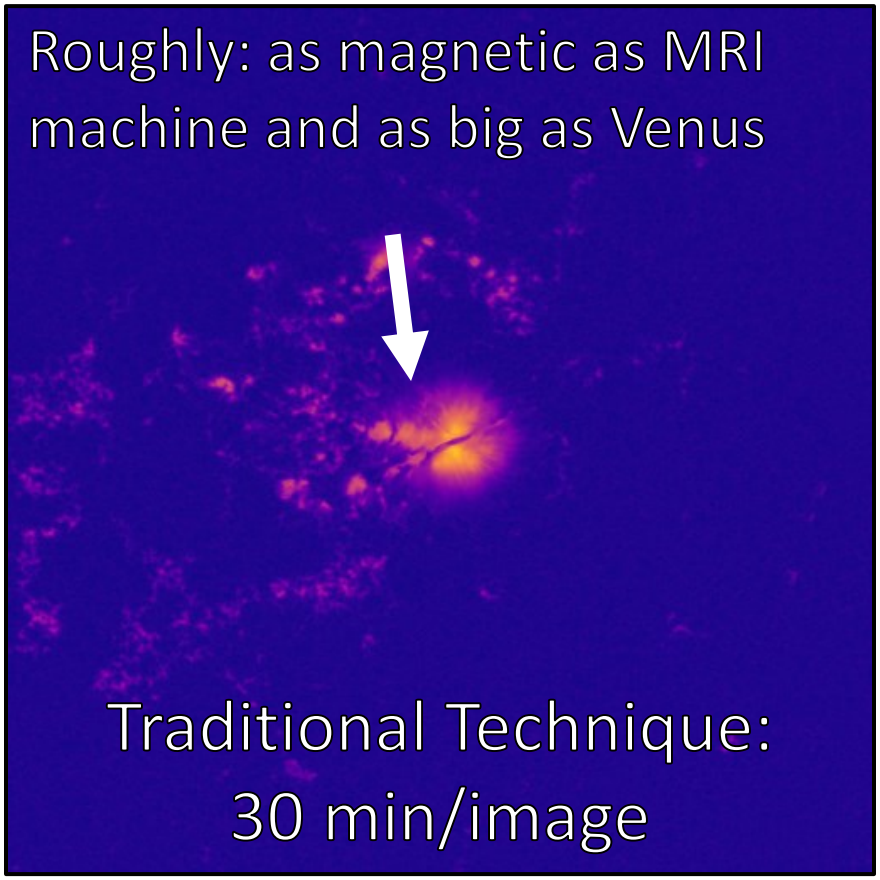

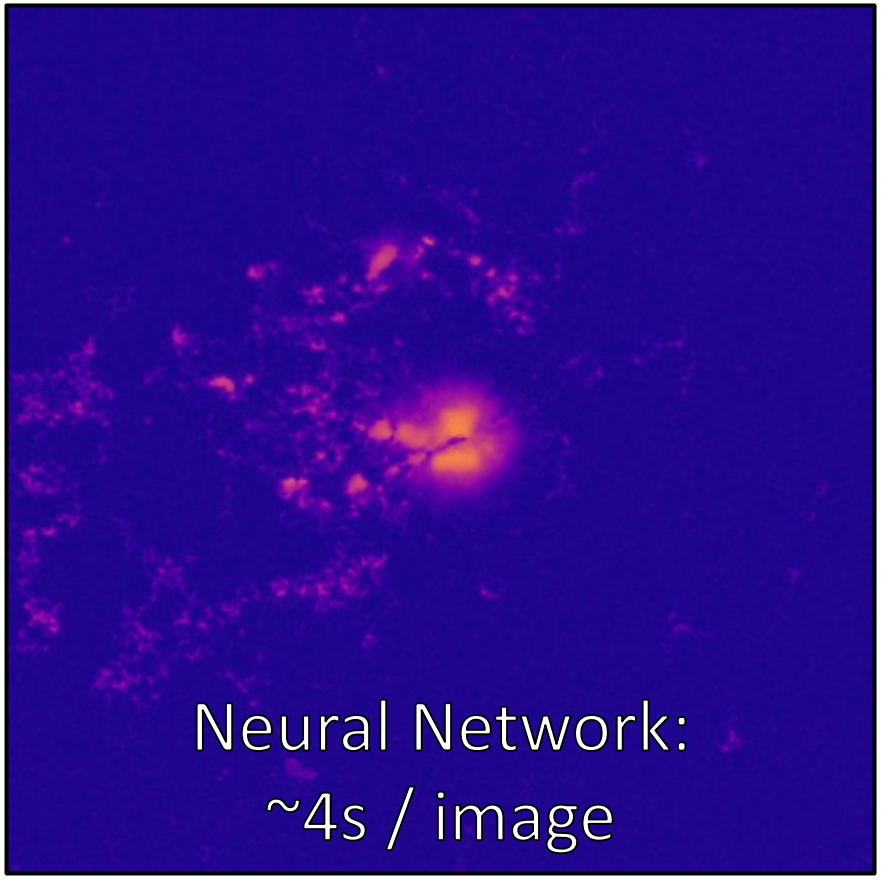

Why is this hard?

All of this stuff seems obvious and effortless!

Why do we need a whole subfield of CS for this?

Key Concept: We see with both our eyes and our brain

Why is this hard?

Why is this hard?

Why is this hard?

Problem: Semantic Gap

097	097	097	097	097	097	097	097	096	097	097	096	096	096
100	100	100	100	100	100	101	101	102	101	100	100	100	099
105	105	105	105	105	105	105	103	102	102	101	103	104	105
109	109	109	109	109	110	107	118	145	132	120	112	106	103
113	113	113	112	112	113	110	129	160	160	164	162	157	151
118	117	118	123	119	118	112	125	142	134	135	139	139	175
123	121	125	162	166	157	149	153	160	151	150	146	137	168
127	127	125	168	147	117	139	135	126	147	147	149	156	160
133	130	150	179	145	132	160	134	150	150	111	145	126	121
138	134	179	185	141	090	166	117	120	153	111	153	114	126
144	151	188	178	159	154	172	147	159	170	147	185	105	122
152	157	184	183	142	127	141	133	137	141	131	147	144	147
130	147	185	180	139	131	154	121	140	147	107	147	120	128
035	102	194	175	149	140	179	128	146	168	096	163	101	125

Why is it important?

Understand Web Data

Instagram: 100 million photos and videos uploaded per day

YouTube: 300 hours of video uploaded every minute

Looking at all content created in one day would take >50 years

Why is it important?

Autonomous Vehicles

Why is it important?

Robotics

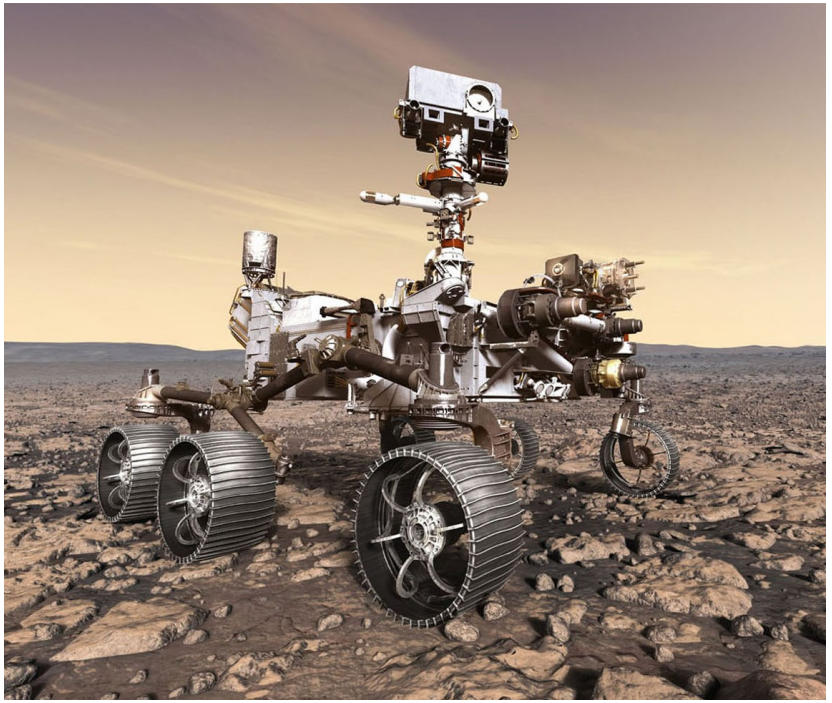

Image source

Why is it important?

Augmented / Virtual Reality

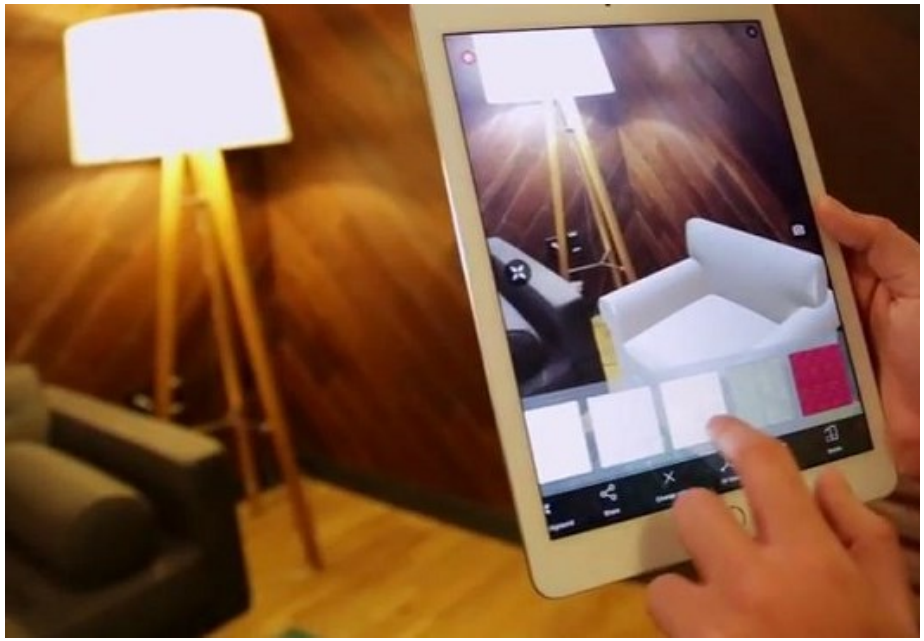

Why is it important?

Science and Medicine

Medical Imaging

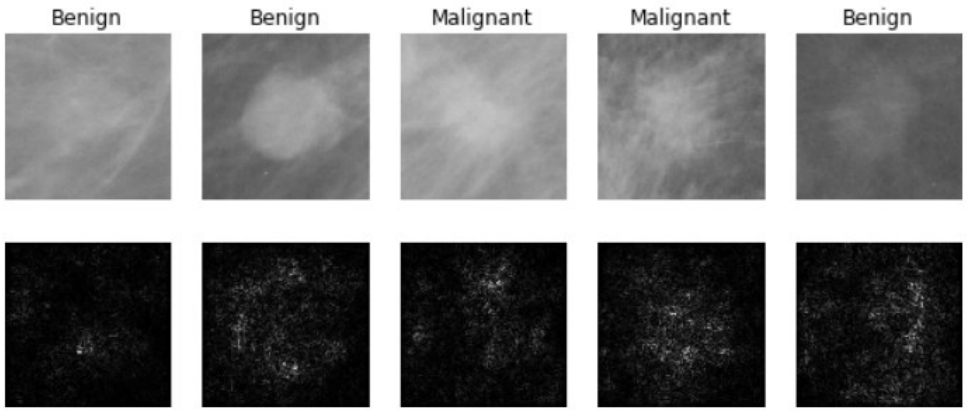

Levy et al, 2016 Figure reproduced with permission

Galaxy Classification

Dieleman et al, 2014

From left to right: public domain by NASA, usage permitted by ESA/Hubble, public domain by NASA, and public domain.

Whale recognition

Kaggle Challenge

This image by Christin Khan is in the public domain and originally came from the U.S. NOAA.

We have made great progress

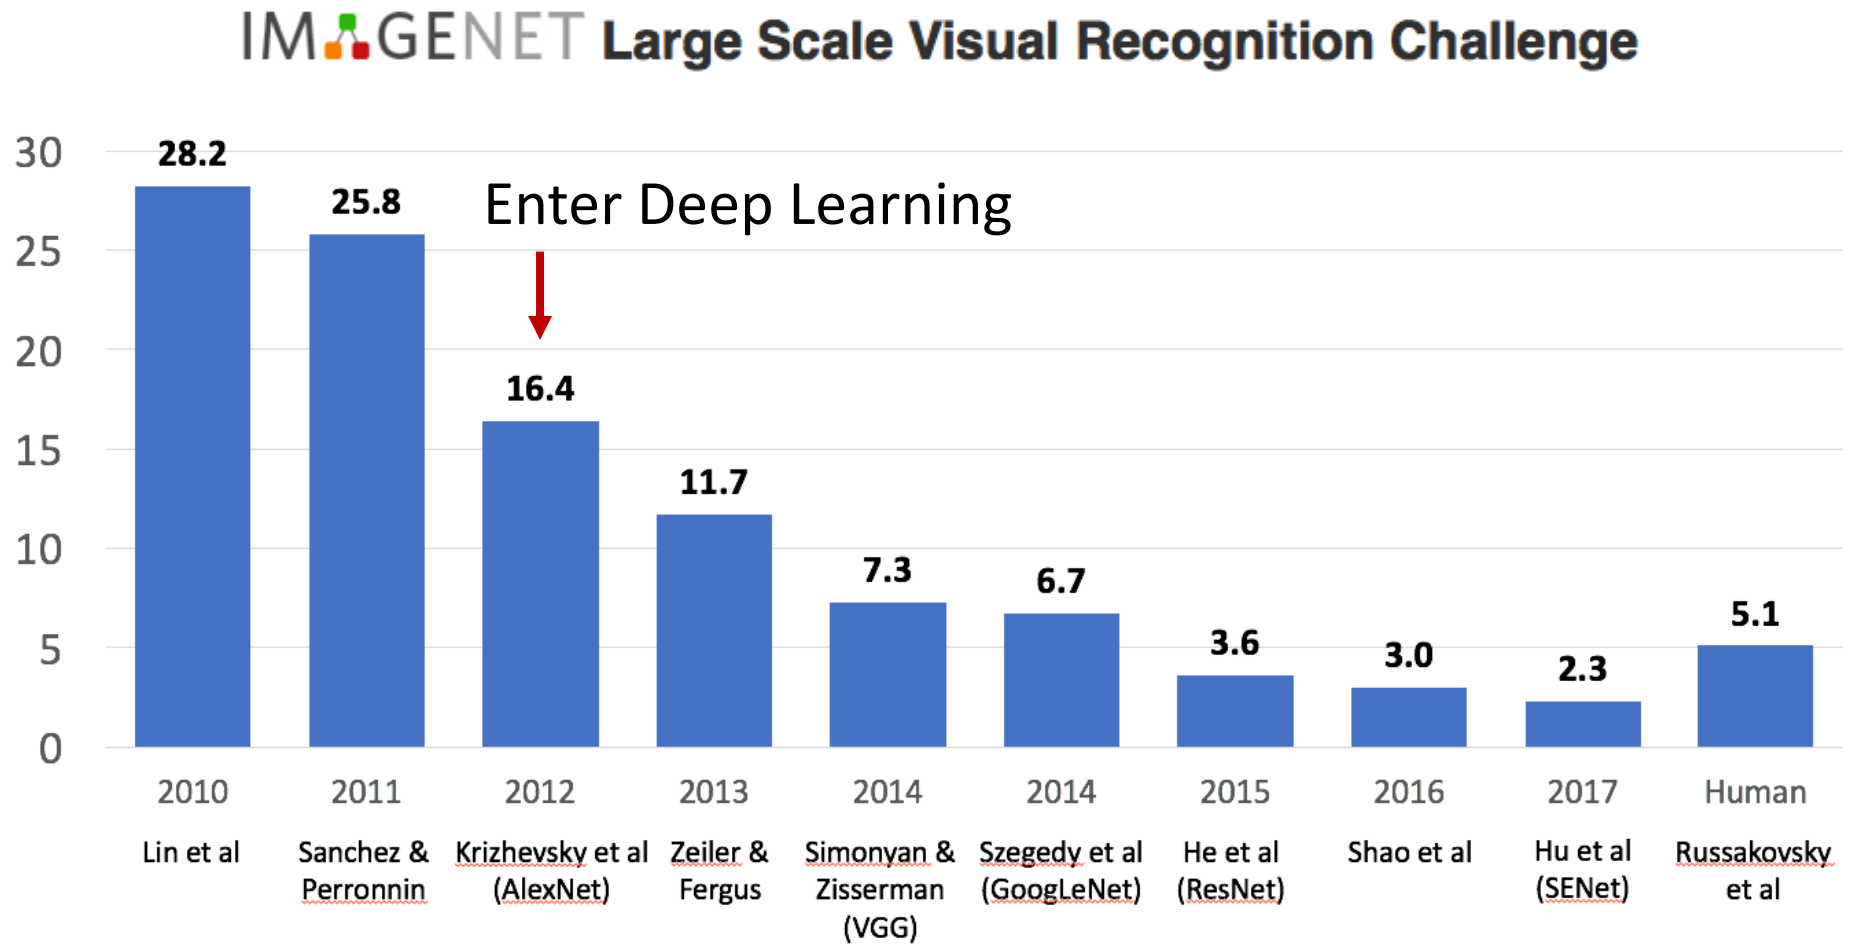

Explosion of Computer Vision

Publications at top Computer Vision conference:
Number of submitted and accepted papers at CVPR

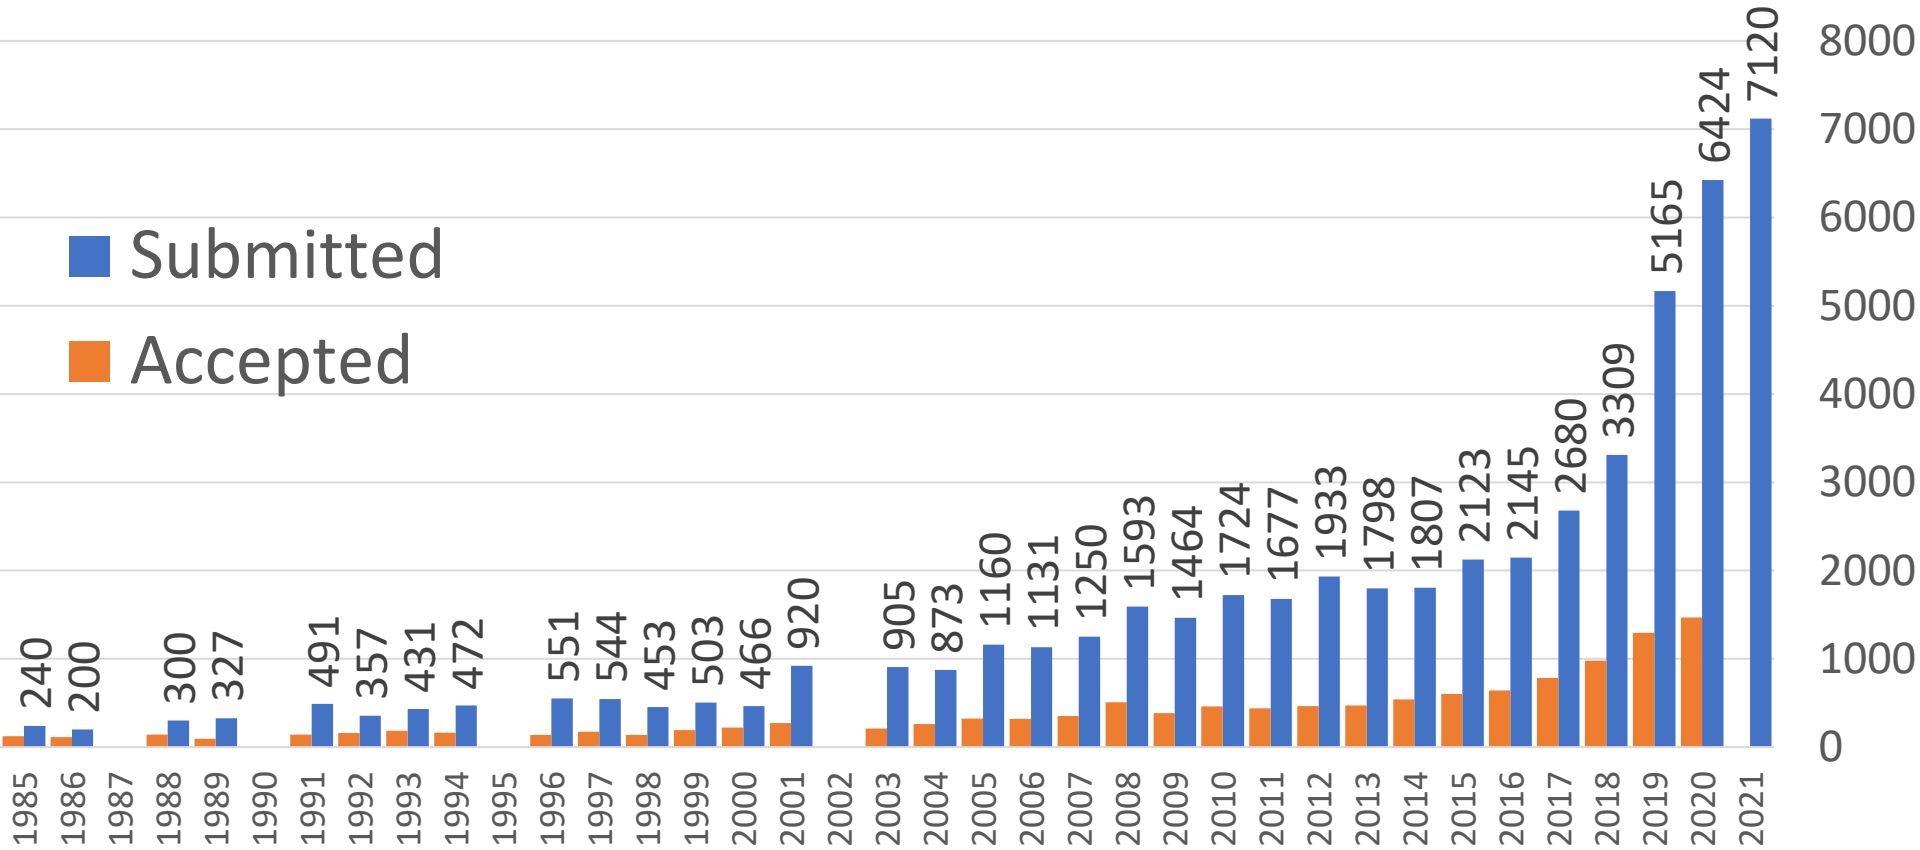

Despite our success, computer vision still has a long way to go.

Computer vision is far from solved!

“The Elephant in the Room”

Modern object detectors seem to work well!

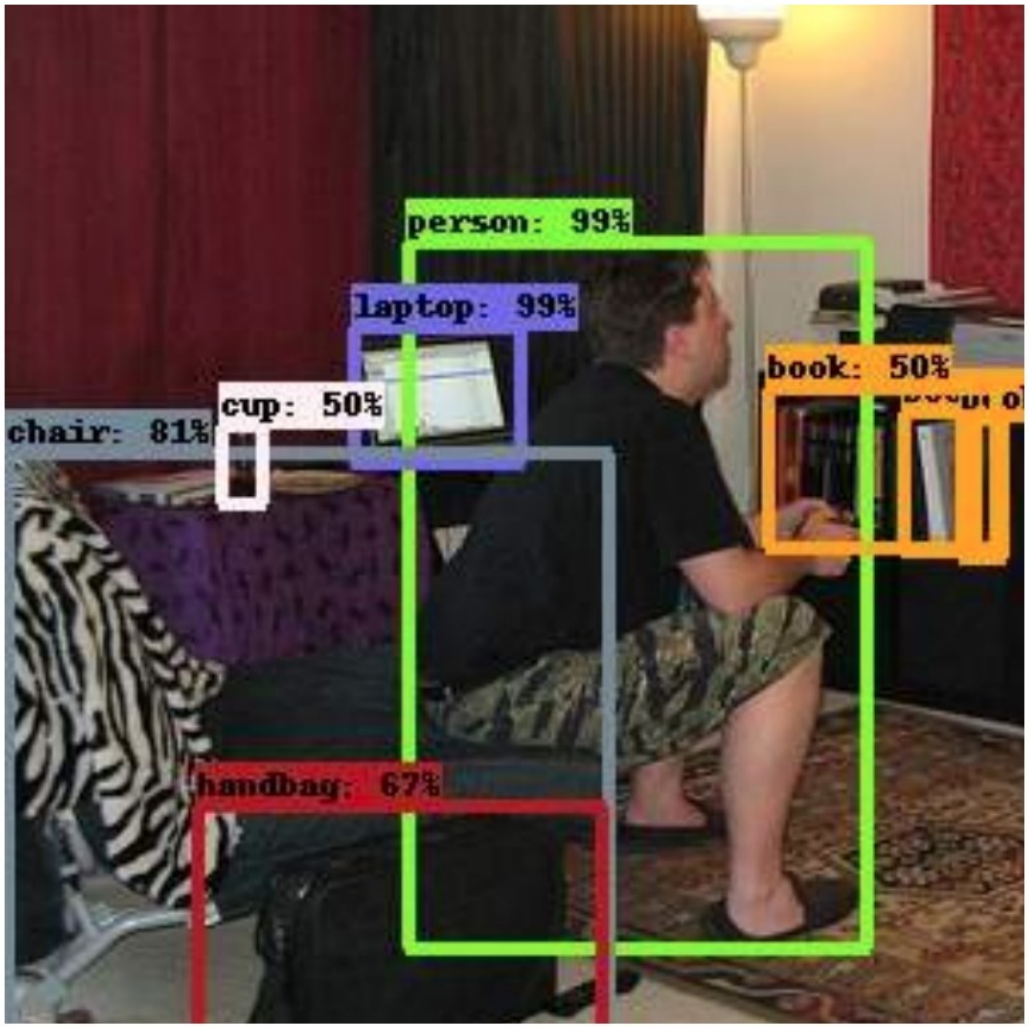

Rosenfeld et al, “The Elephant in the Room”, arXiv 2018

“The Elephant in the Room”

”Unusual” objects are often missed!

Computer Vision systems ”see” very differently than we do

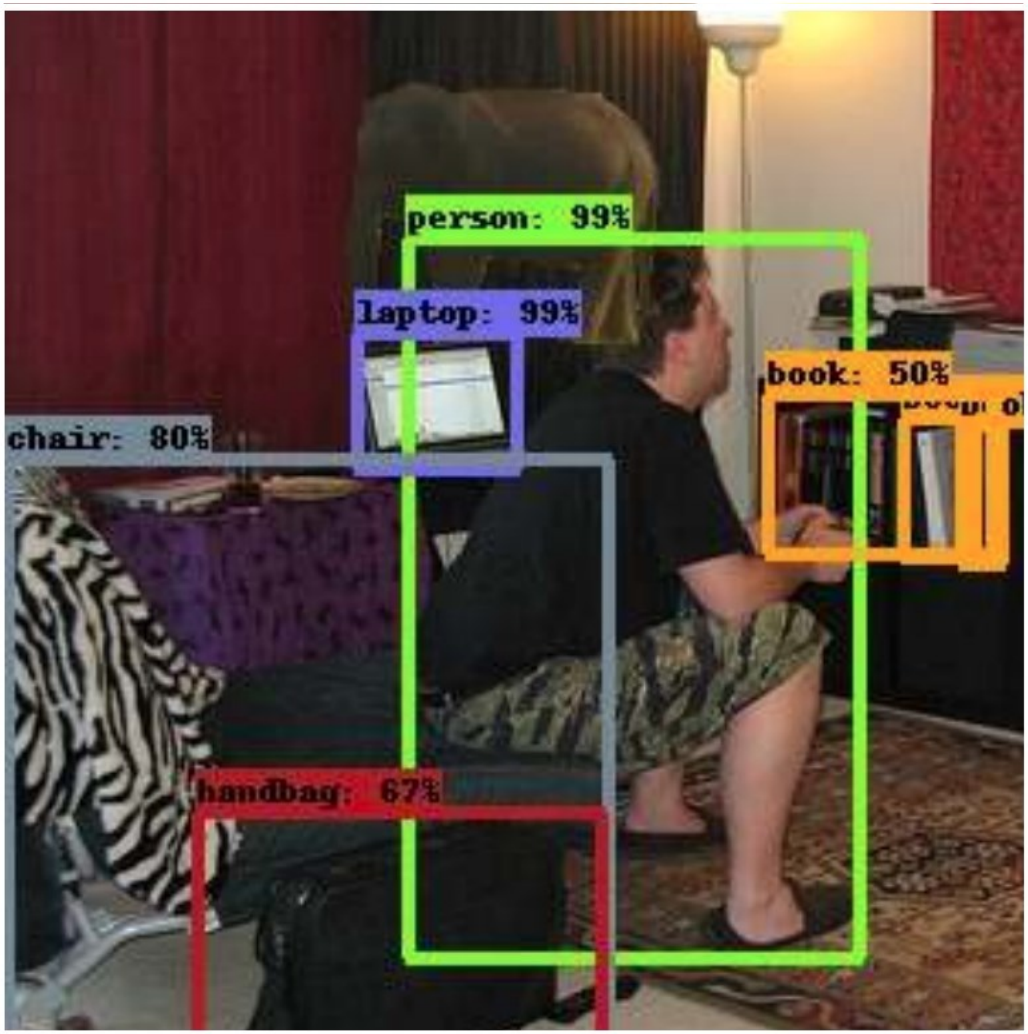

Rosenfeld et al, “The Elephant in the Room”, arXiv 2018

Example credit: Andrej Karpathy

Course Logistics

Meeting Times

- Lecture:
  - Section 1: Tue / Thu 10:30am – 12:00pm on Zoom
  - Section 2: Tue / Thu 12:00pm – 1:30pm on Zoom
- Office Hours
  - Lots of office hours! Schedule TBD
  - **GSI Office Hours**: Questions about homework, code
  - **Faculty Office Hours**: Prefer conceptual questions

Prerequisites

You **absolutely** need:

- **Programming**: EECS 281
- **Linear Algebra** (new): Math 214, 217, 296 or 417
  - Reach out to us if you have equivalent experience
  - We will have a one-lecture math refresher

You'll have to learn: Numpy + PyTorch, a little tiny bit of continuous optimization

Prerequisites

Suppose **K** in R^{3x3}, **x** in R^3 .Should know:

- How do I calculate **Kx?**
- When is **K** invertible?
- What is **x** if **Kx** = λ**x** for some λ?
- What's the set {**y:** $\mathbf{x}^T\mathbf{y}$ = 0} geometrically?

You should also be able to remember some notion of a derivative

Websites

- Course website: https://web.eecs.umich.edu/~justincj/teaching/eecs442/WI2021/
- Piazza: https://piazza.com/class/umich/winter2021/eecs442
- Canvas / Autograder for code submission
- Gradescope for writeup submission

Piazza

- Please ask questions on Piazza so we can answer the question once, officially, and quickly
- We will monitor Piazza in a systematic way, but we do not guarantee instant response times
- Please don't ask questions about course material, homework, or logistics over email – those should all go to Piazza
- If you have a sensitive or private issue, then email me / David directly

Textbooks

No required textbook.

Szeliski, *Computer Vision: Algorithms and Applications*, is a good reference and available online.

http://szeliski.org/Book/

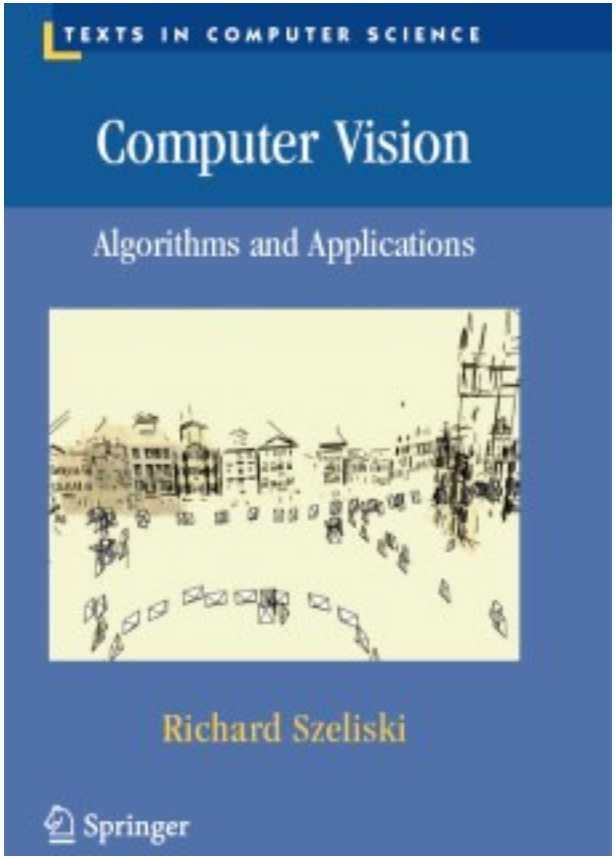

Evaluation

- Homework (76%)
  - Programming assignments in Python / numpy / Pytorch
  - HW0 is a numpy crash course (6%)
  - HW1-6 are guided mini-projects
  - HW1-5 are 12% each
  - HW6 is 10%
- Project (24%)
  - Project Proposal (2%)
  - Virtual Project Showcase (8%)
  - Project Report (14%)

Homework Late Policy

- Penalty: 1% per hour, rounded to nearest
- Example:
  - Due: Midnight Mon. (1s after 11:59:59pm Mon)
  - Submitted at 12:15am Tue: No penalty!
  - Submitted at 6:50am Tue: 7% penalty (specifically 90% -> 83%)
- Exceptions only for exceptional circumstances.
- Everyone gets **72 free late hours**, applied automatically

Copying: There are better options

- Read the syllabus
- Copying is usually *painfully* obvious and I don't have many options
- Submit it late (*that's why we have late days*), half-working (*that's why we have partial credit*), or take the zero on the homework – I guarantee you won't care about one bad homework in a year
- If you're overwhelmed, talk to us

Course Project

- Work in a team of 3-5 to do *something cool*
- There will be a piazza thread for pairing up
- Could be:
  - Applying vision to a problem you care about
  - Independent re-implementation of a paper
  - Trying to build and extend an approach
- Should be 2 homeworks worth of work per person

First Homework: Numpy warmup

First homework is meant as an introduction / refresher to array programming with numpy

Will be released tomorrow, due in one week on Wednesday 1/27 11:59pm

And now to David: Camera Models