EECS 442

Computer Vision

Prof. David Fouhey

Winter 2019, University of Michigan

http://web.eecs.umich.edu/~fouhey/teaching/EECS442_W19/

Goals of Computer Vision

Get a computer to understand

Goal: Naming

Goal: Naming

The picture shows a building with many windows and grass in front of it. There is a person walking on the right…

Goal: 3D

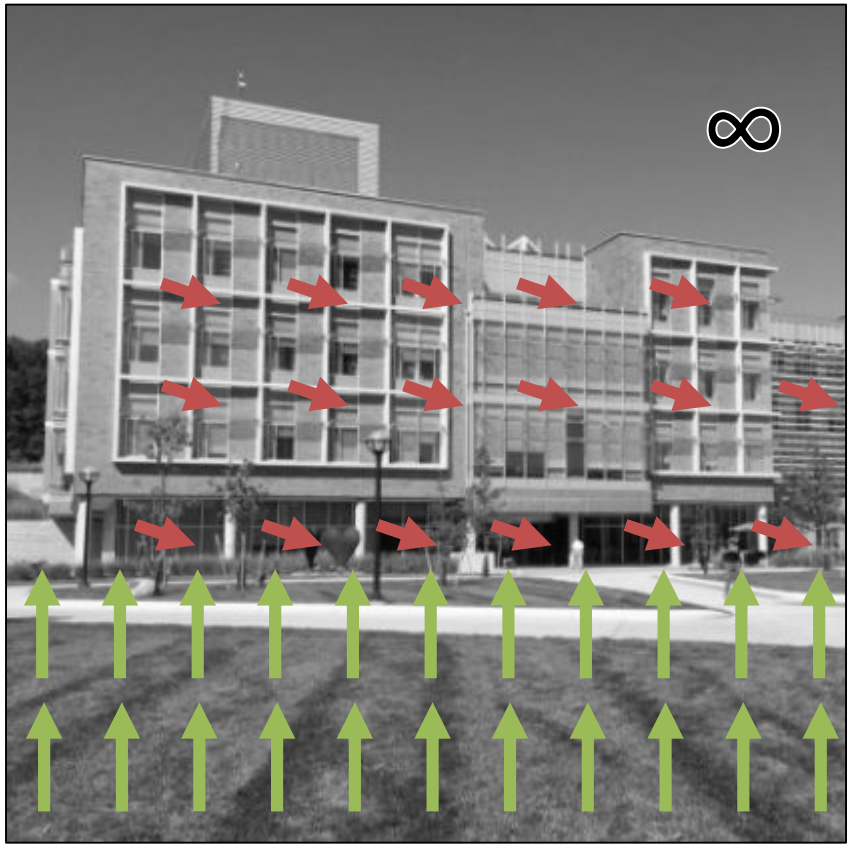

Goal: Actions

Seems Obvious, Right?

- **Key concept to keep in mind throughout the course**: you see with both your eyes **and** your brain.

Why is it Hard?

Why is it Hard?

Goal of computer vision

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

Despite This, We've Made Progress

- Few of these problems are **solved** *(and there are lots of dangers to pretending things are solved when they aren't)*
- But we do have systems with performance ranging from non-embarrassing to super-human (with the right caveats)

Look at Your Phone

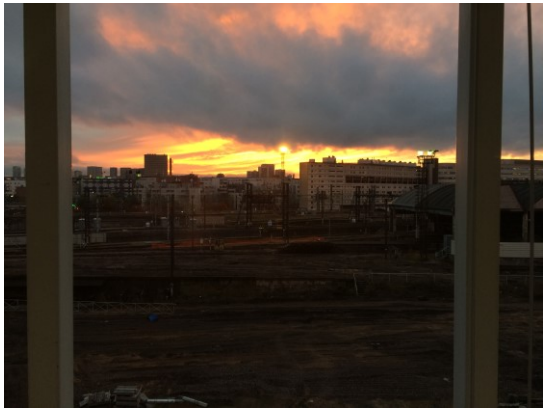

Iphone Image Credit: Wikipedia

Graphics

https://affinelayer.com/pixsrv/

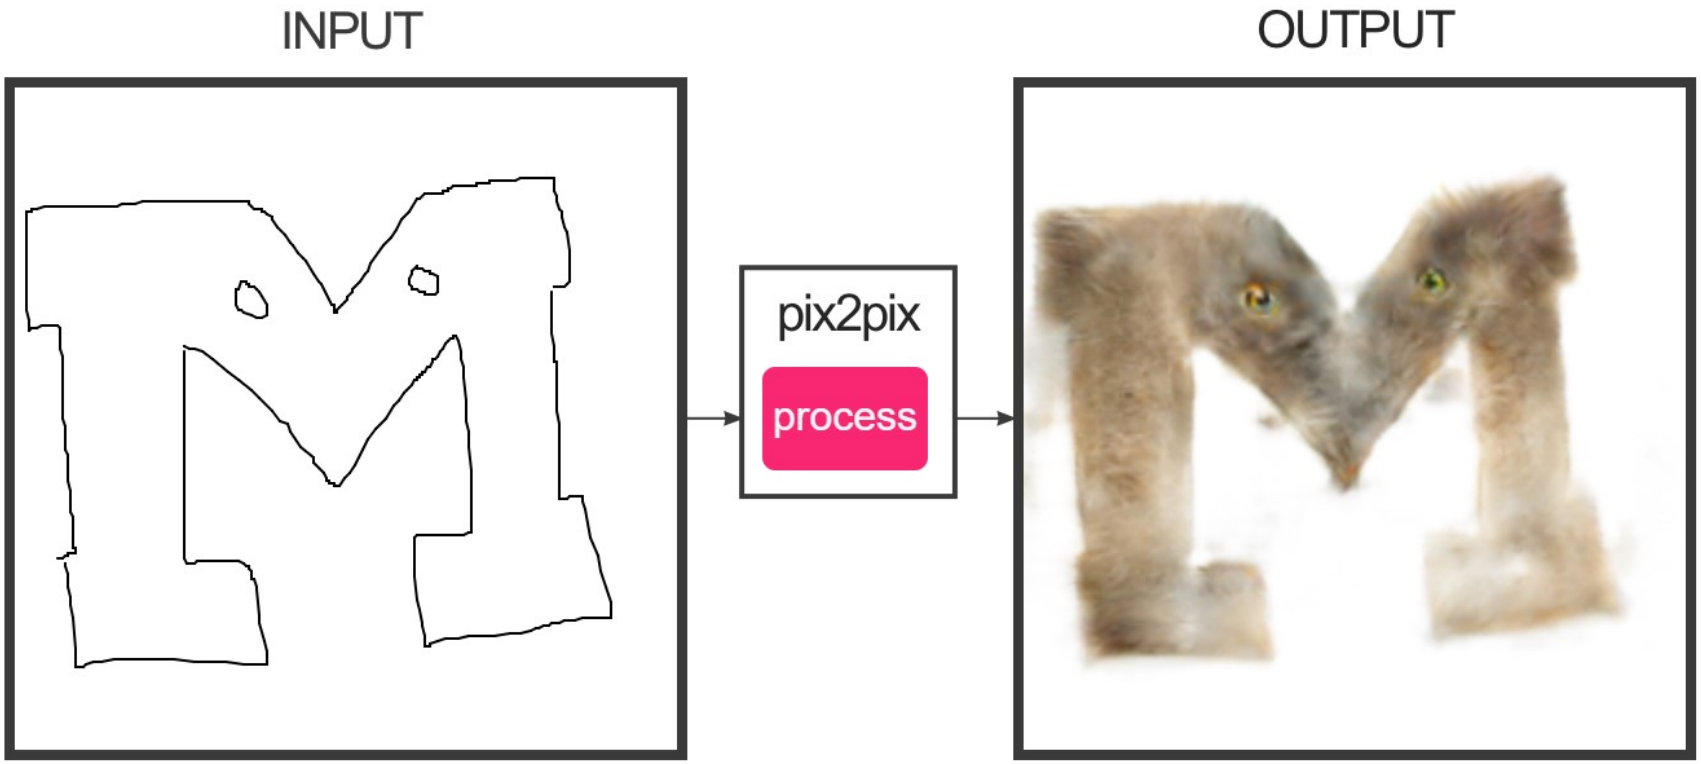

Isola et al. *Image-to-Image Translation with Conditional Adversarial Networks*. CVPR 2017

Graphics

Slide Credit: S. Seitz

Faces

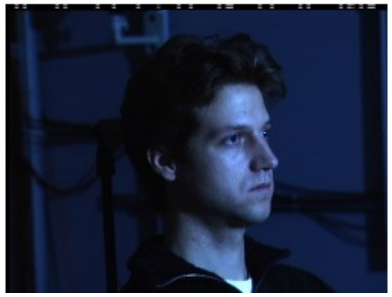

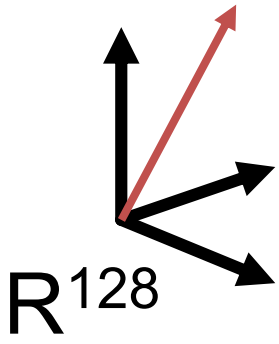

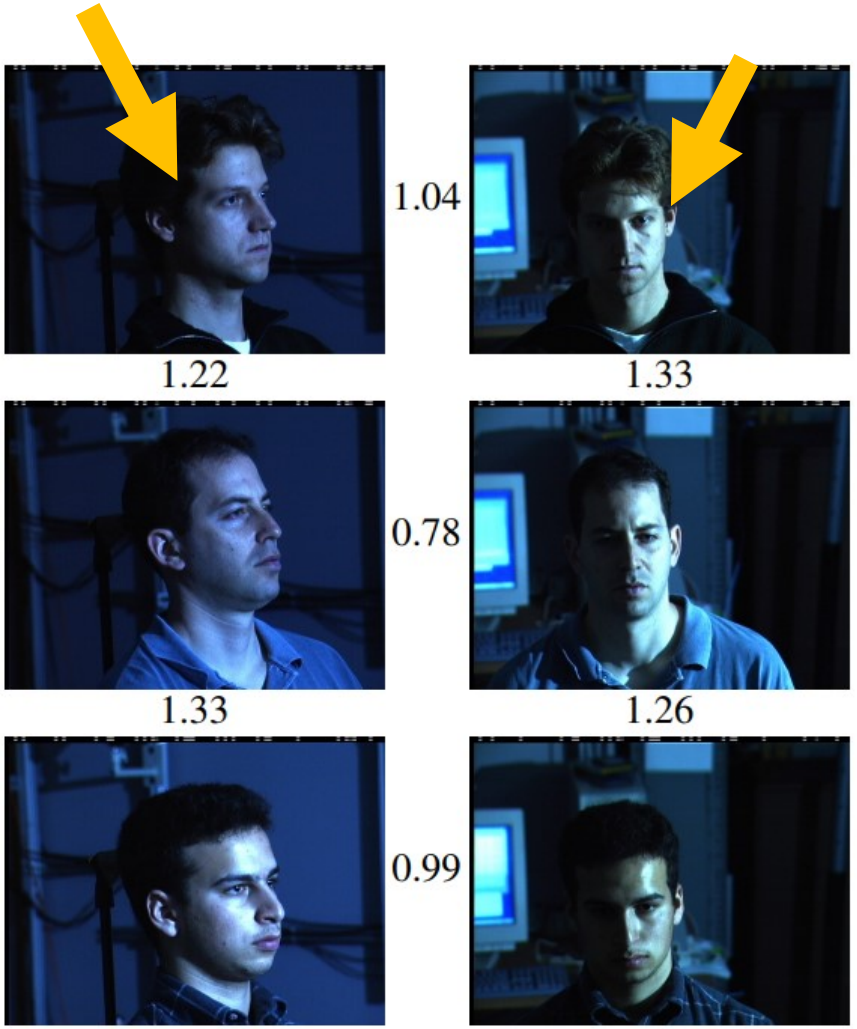

Schroff et al. *FaceNet: A Unified Embedding for Face Recognition and Clustering*. CVPR 2015

Humans

Real-time Multi-Person 2D Pose Estimation Using Part Affinity Fields

Zhe Cao, Tomas Simon, Shih-En Wei, Yaser Sheikh
Carnegie Mellon University

Cao et al. *Realtime Multi-person 2D Pose Estimation using Part Affinity Fields*. CVPR 2017

Recognition

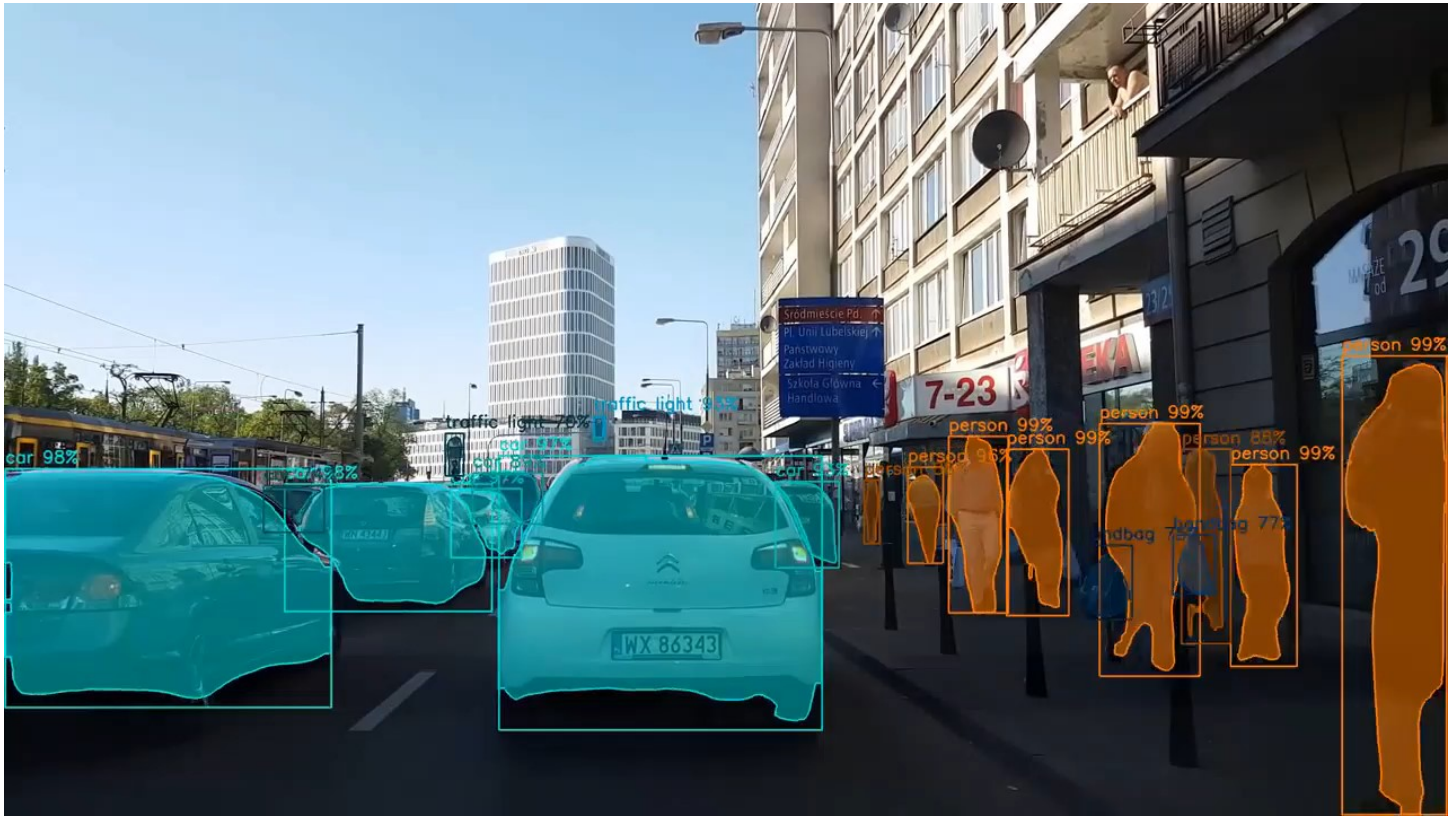

He et al. *Mask RCNN*. ICCV 2017.
Video Credit: Karol Majek (https://www.youtube.com/watch?v=OOT3UIXZztE)

3D

Agarwal et al. *Building Rome In A Day*. ICCV 2009.

3D

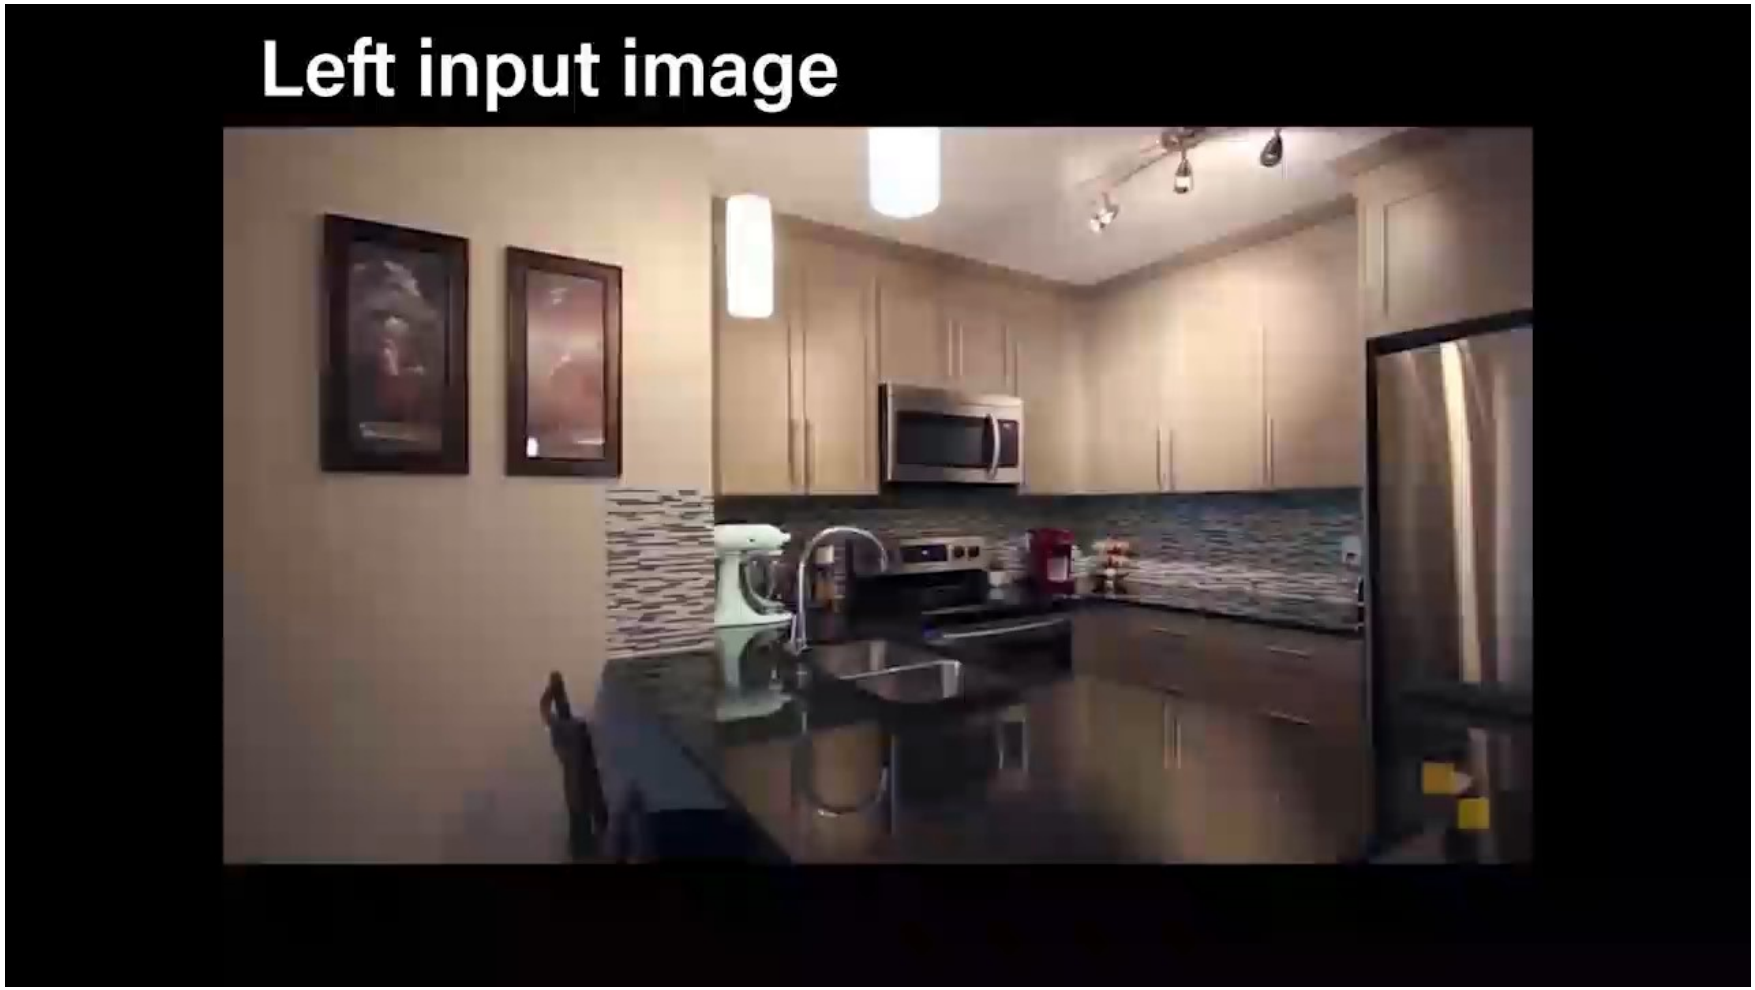

Zhou et al. Stereo Magnification: Learning View Synthesis using Multiplane Images. SIGGRAPH 2018.

Vision Assisting Things

Owens et al. *Audio-Visual Scene Analysis with Self-Supervised Multisensory Features .* ECCV 2018

Why is Computer Vision Difficult?

Viewpoint Variation



Slide Credit: L. Lazebnik

Illumination Variation

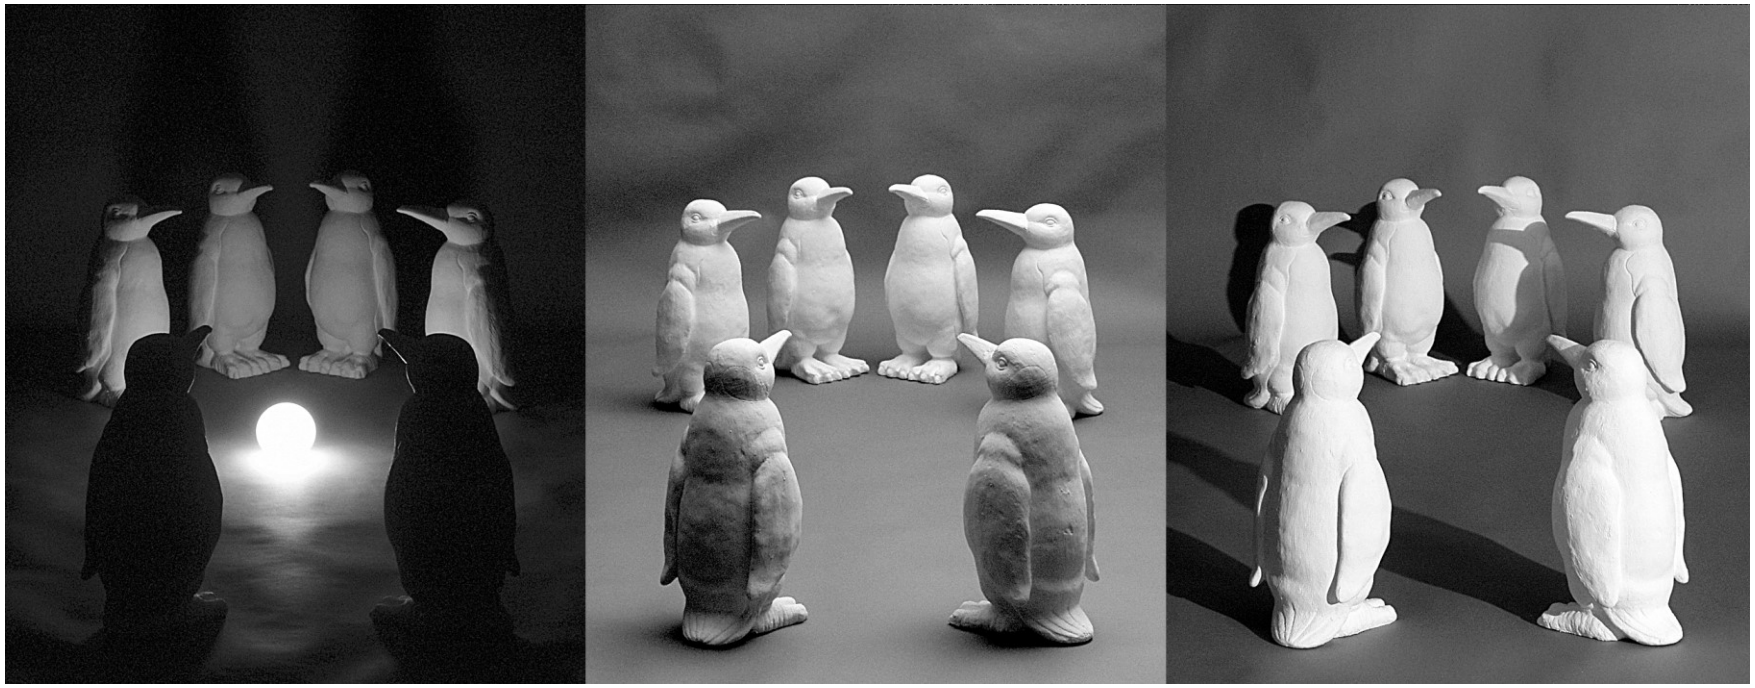

Image Credit: J. Koenderink

Scale Variation

Slide Credit: L. Fei-Fei, Fergus & Torralba

Deformation

Image Credit: Peng et al., SIGGRAPH ASIA 2018

Intra-Object Class Variation

Slide Credit: Fei-Fei, Fergus & Torralba

Occlusion, Clutter

Image Credit: Wikipedia

Ambiguity

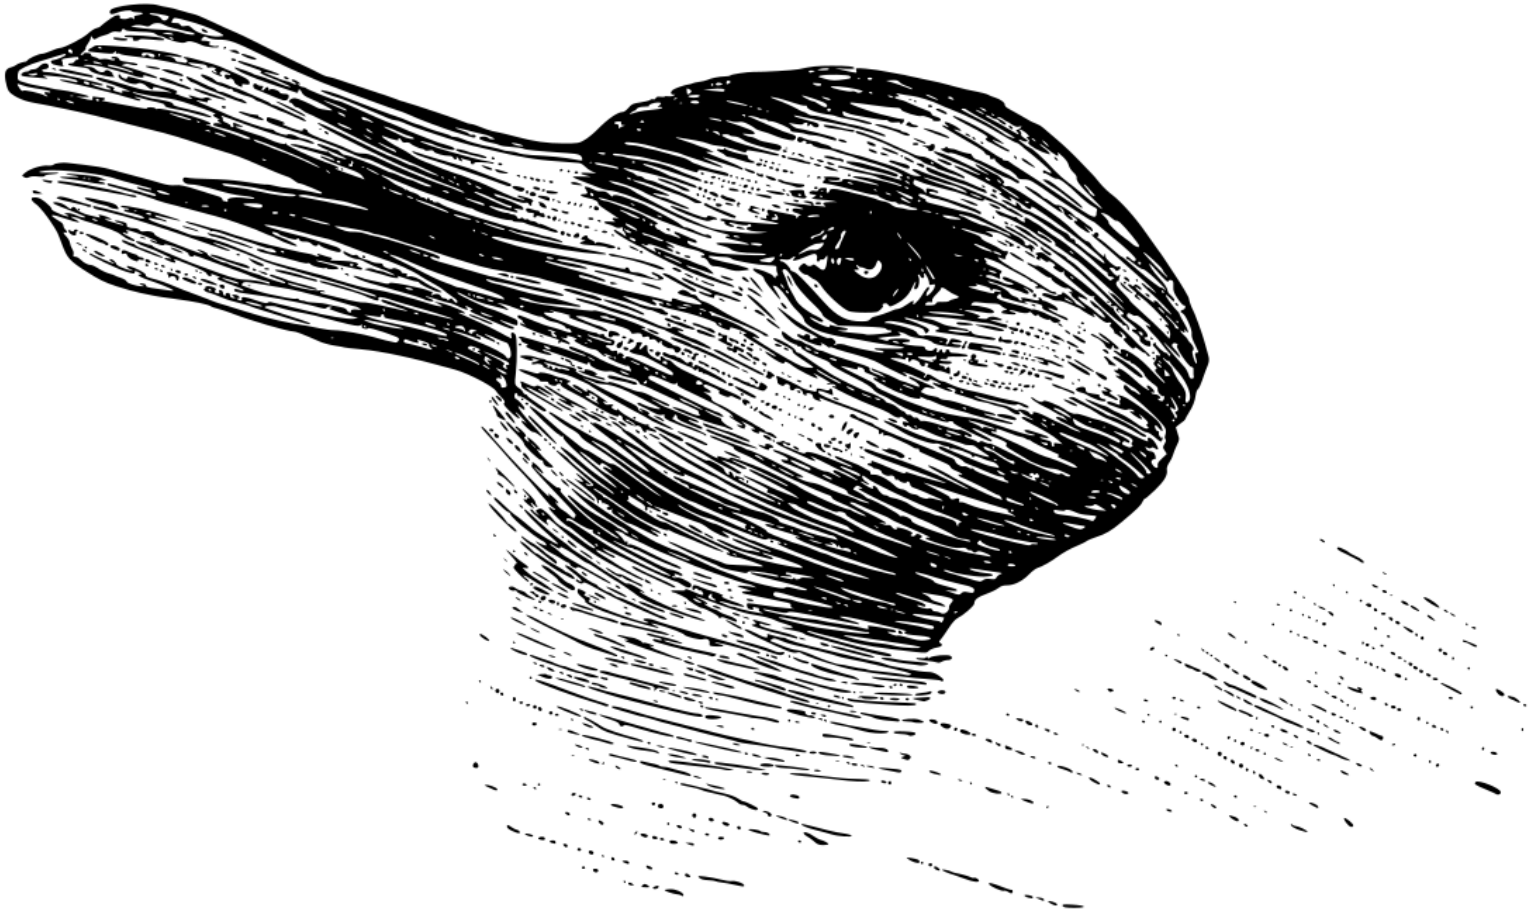

Slide Credit: Fliegende Blätter

Ambiguity

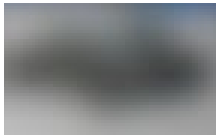

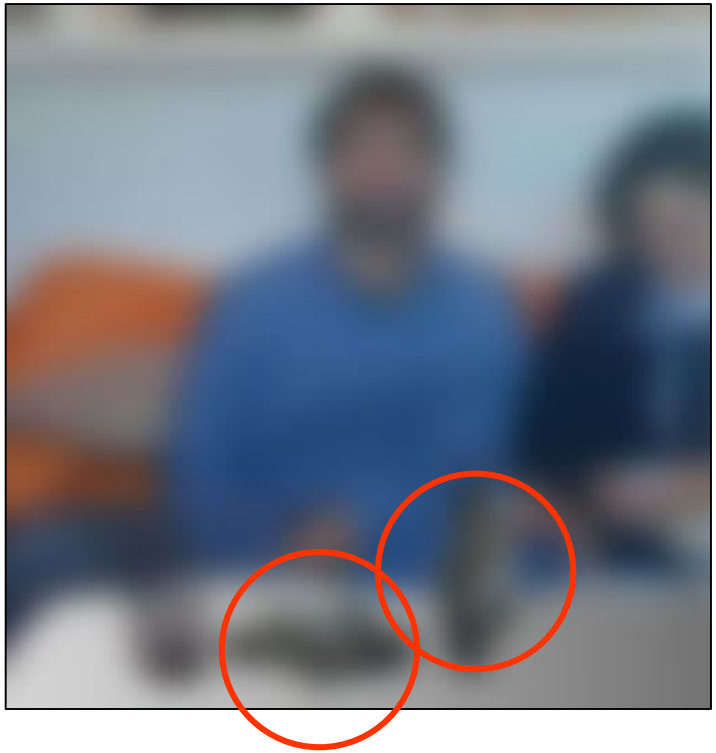

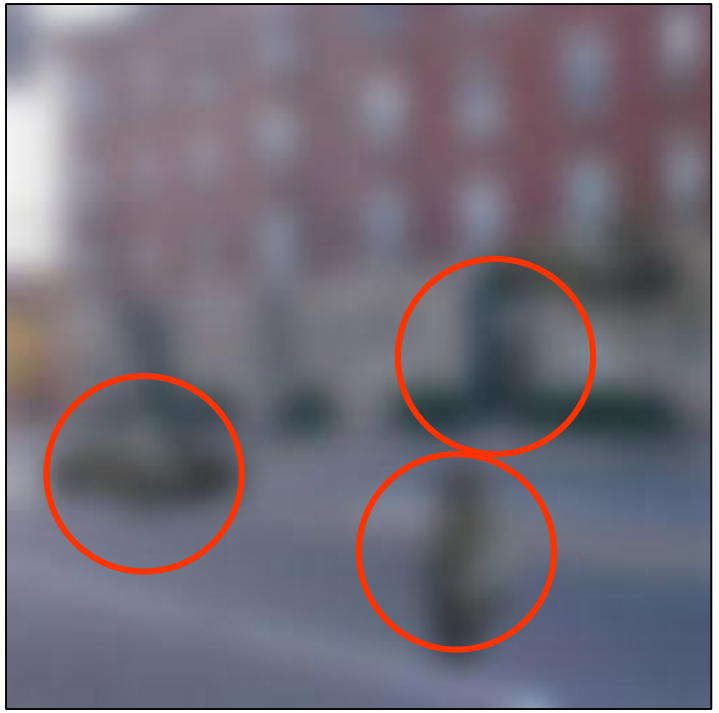

Slide Credit: L. Fei-Fei, Fergus & Torralba

Ambiguity

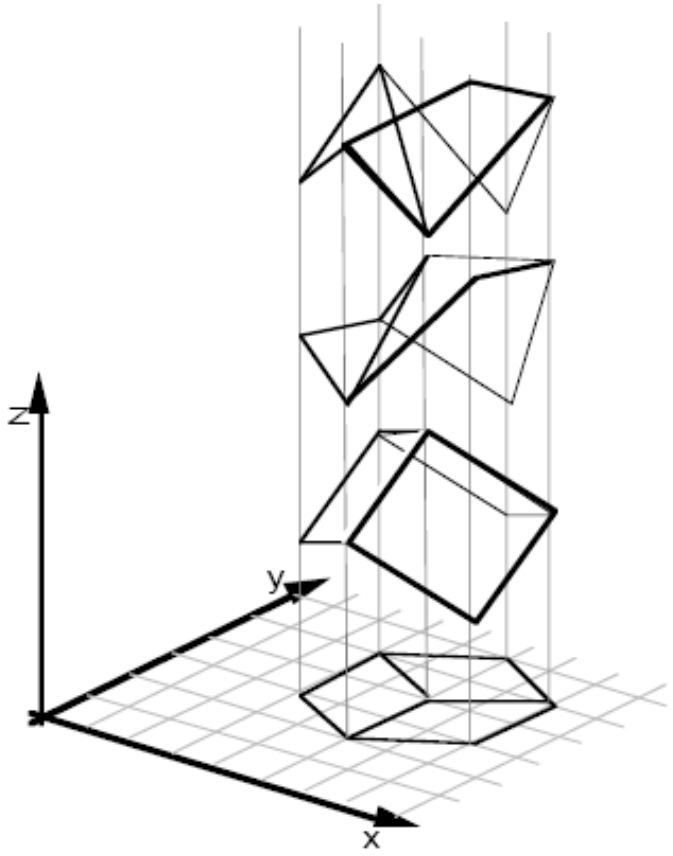

Slide Credit: Sinha and Adelson 1993

Why is it Possible?

Has regularity

Has rules

Has rules and regularity

The World

Imaging

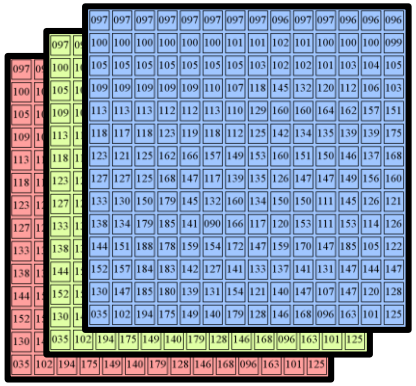

Our Job

Sift through: evidence (the image) and past experience (knowledge) to interpret the image correctly.

Slide Credit: J. Deng

Cues: Perspective

Cues: Shading

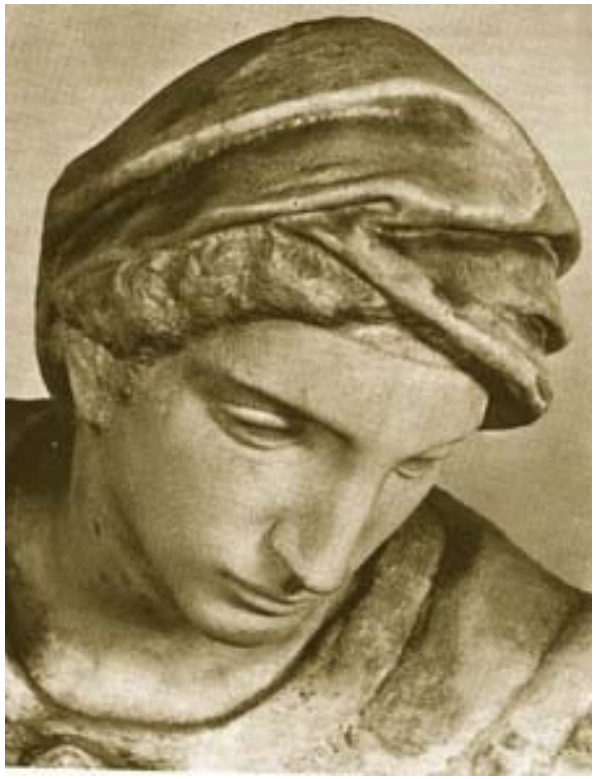

Slide Credit: L. Lazebnik, L. Fei-Fei, Fergus & Torralba

Cues: Texture Gradient

Slide Credit: J. Deng

Cues: Common Fate

Image Credit: Pathak et al. *Learning Features by Watching Objects Move.* CVPR 2017.

Course overview

1. Image formation and processing
2. Learning and deep learning
3. Transformations and motion
4. 3D reconstruction
5. Advanced topics

Part 1: Formation and Processing

Camera Models

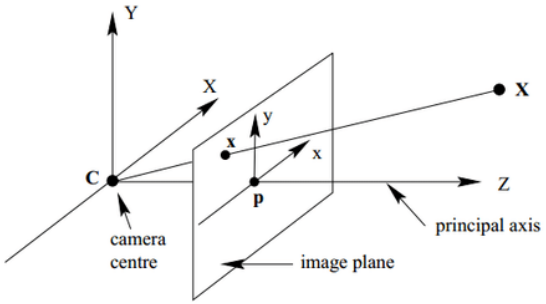

Linear Filtering

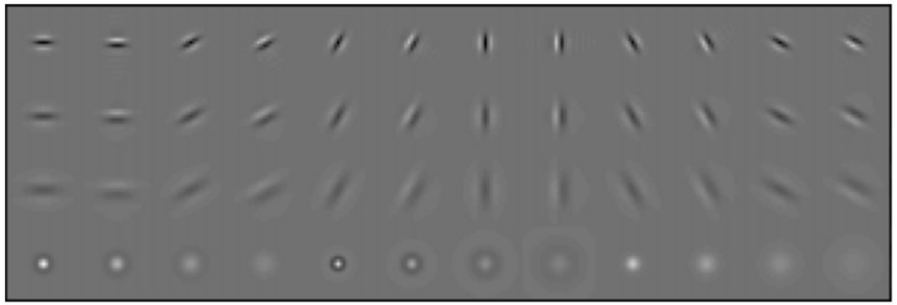

Feature Detection

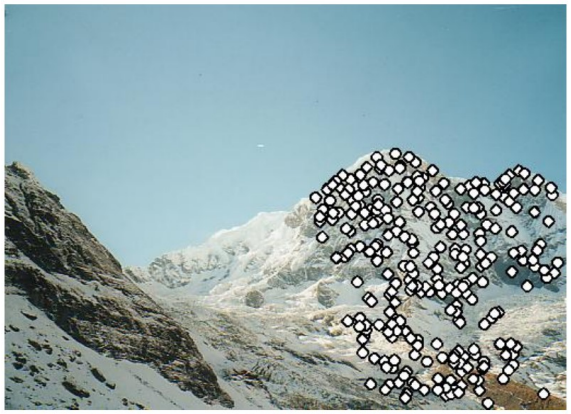

Image Credit: Hartley and Zisserman 04, Leung and Malik IJCV 01, Brown and Lowe ICCV 03,

Part 2: Transformations and Fitting

Transformations

Robust Fitting

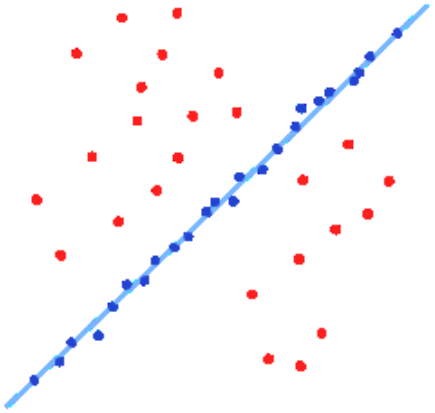

Image Credit: Wikipedia

Part 3: Learning and Deep Learning

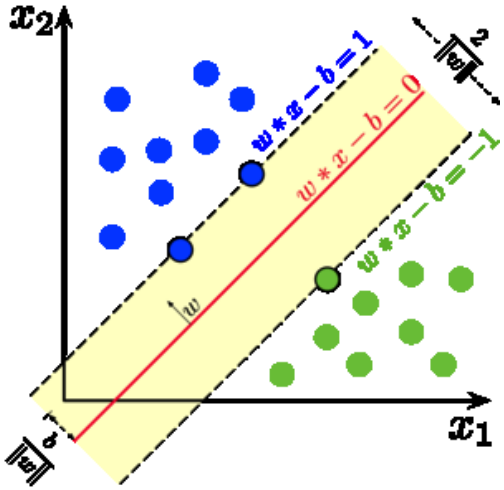

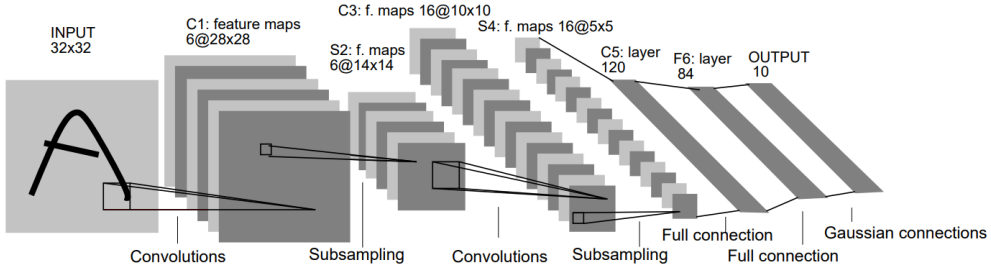

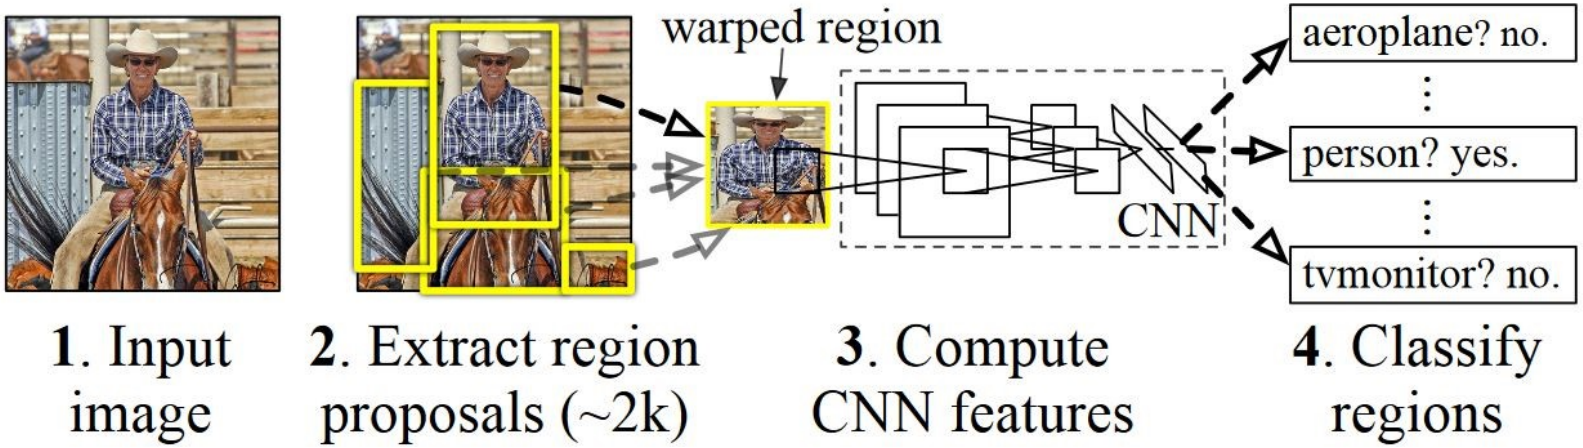

Image Credit: Wikipedia, LeCun et al. Proc IEEE 01, Girshick et al. CVPR14

Part 4: 3D Reconstruction

Stereo Vision

Multiview Stereo and Structure From Motion

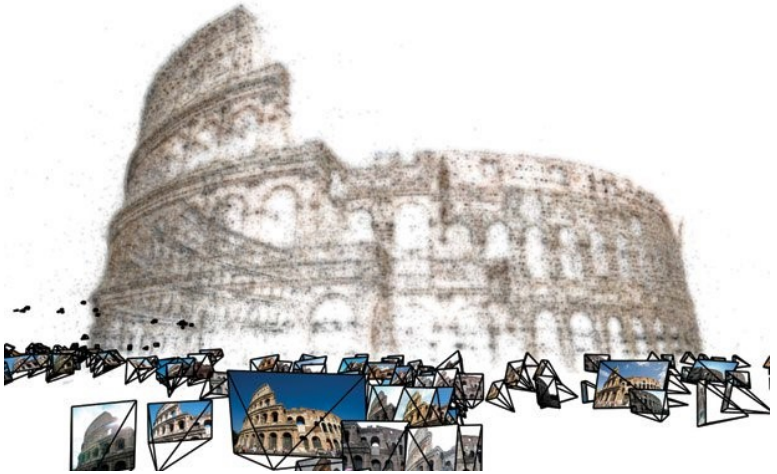

Part 5: Advanced Topics

Vision & Language

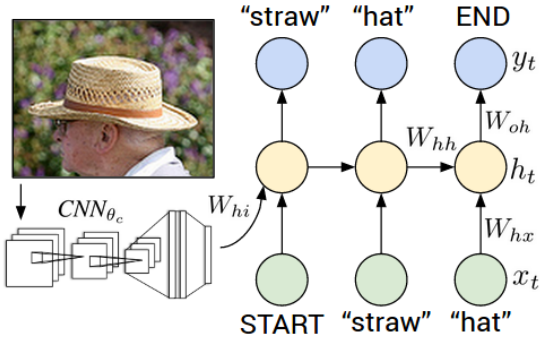

Video

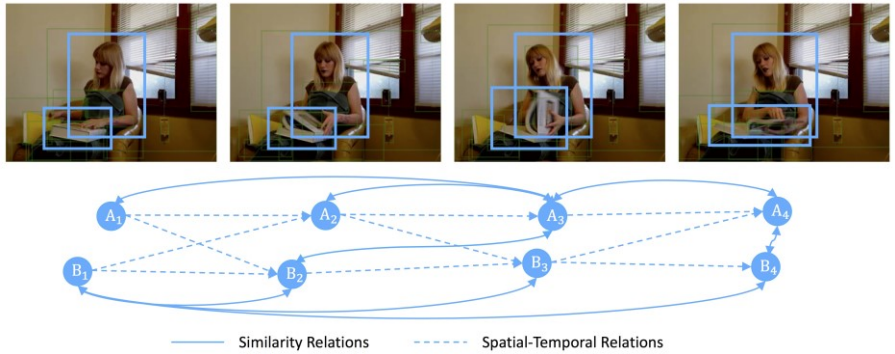

Learning and Geometry

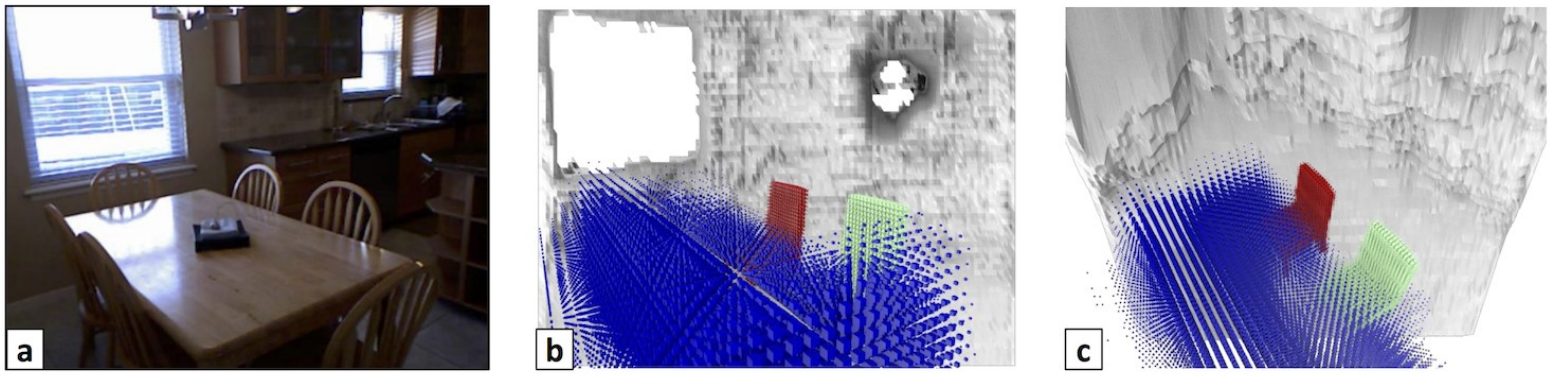

Image Credit: Karpathy et al. CVPR 2015, Wang et al. ECCV 2018, Tulsiani et al. CVPR 2018

Textbooks

No textbook, but Szeliski, *Computer Vision: Algorithms and Applications*, is a good reference and available online.

http://szeliski.org/Book/

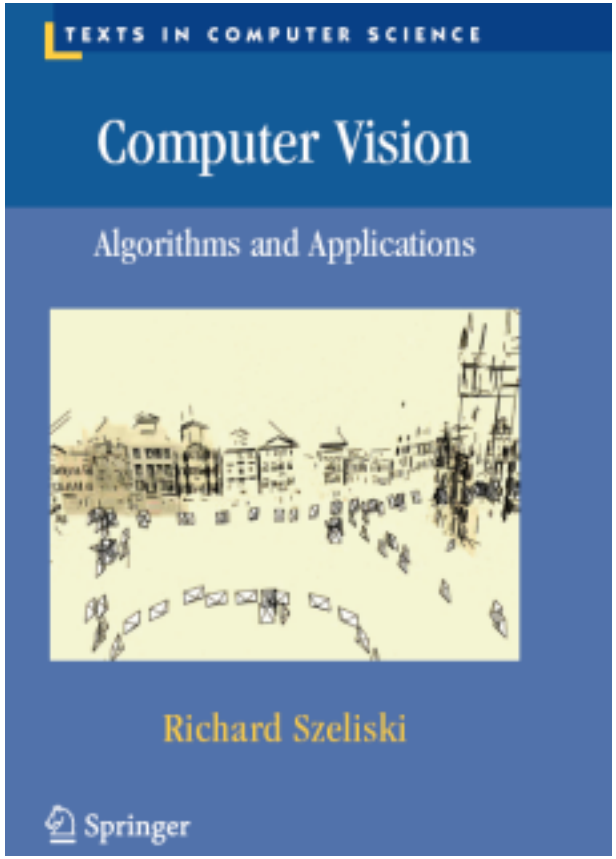

Administrivia

- Websites / Staff
- Prerequisites
- Waitlist etc.
- Evaluation
- Classes/Discussions/Piazza/Office Hours

Websites

- Course website: http://web.eecs.umich.edu/~fouhey/teaching/EECS442_W19/
- Piazza: You should have access via canvas
- We'll use Piazza to make announcements/discussions, and things like homework will appear on the website.

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 cannot guarantee instant response times
- Same goes for email

Staff

- Professor: (me) David Fouhey
- GSIs / IAs:
  - Linyi Jin,
  - Richard Higgins
  - Shengyi Qian
  - Yi Wen

Prerequisites

You **absolutely** need: EECS 281 and corresponding programming ability.

You will **struggle continuously** without: Basic knowledge of linear algebra, calculus.

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: x**^{T}**y** = 0} geometrically?

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

Waitlist Policies

1. Waitlist right now is huge
2. I will move as many people off as possible
3. I will not reorder the waitlist
4. If you are dropping, please drop quickly so others can be added quickly

Evaluation

- Mid-term Exam: 15%
- Homeworks: 5 x 10%
- Project: 35%

Evaluation: Mid-term

- 15% of grade
- Thursday before Spring Break (2/28) in class
  - **Please do not schedule things.**
- Will cover:
  - Images and image processing
  - Fitting and matching
  - Basics of Learning

Evaluation: Homework

- 5 Homeworks, 10% Each
- Submit a tiny project (code) + write-up (pdf)
- **You should** discuss, but your implementations should be your own.
- **No:** copying off the Internet or your classmates, asking reddit / stackoverflow, over-the-shoulder debugging
- **Overall:** should not know the code for how others solved it.

Evaluation: Homework Late Days

- 3 late days in *The Ann Arbor Bank of Late HW*
- Spend these as you choose. No loans!
- No need to announce you're taking a late day – we'll just deduct it automatically.

Evaluation: Homework Late Policy

- Penalty: 1% per hour, round to nearest hour
- Example:
  - Due: Midnight Mon. (1s after 11:59:59pm Mon)
  - Submitted at 12:15am Tue: No penalty!
  - Submitted at 6:50am Tue: 7% penalty
- Exceptions only for exceptional circumstances (talk to us)
- **Questions?**

Evaluation: Homework Advice

- Start early: vision often takes a while to run. Think of both computer time and your time. They're different.
- Vision code often "works" a little, but poorly, with bugs. Build in time for two full screwups
- Make things modular: visualize and test on smaller data.

All three interact – bugs are expensive since they may require lengthy reruns

Evaluation: Term Project

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

Evaluation: Term Project

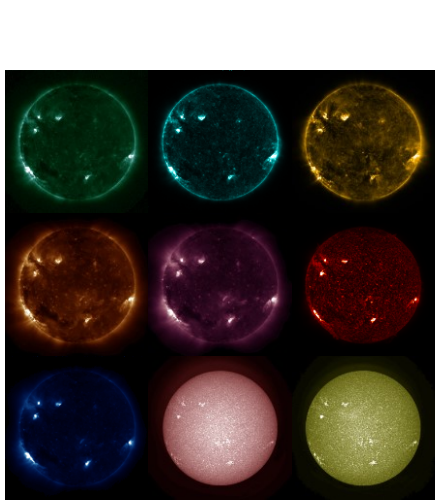

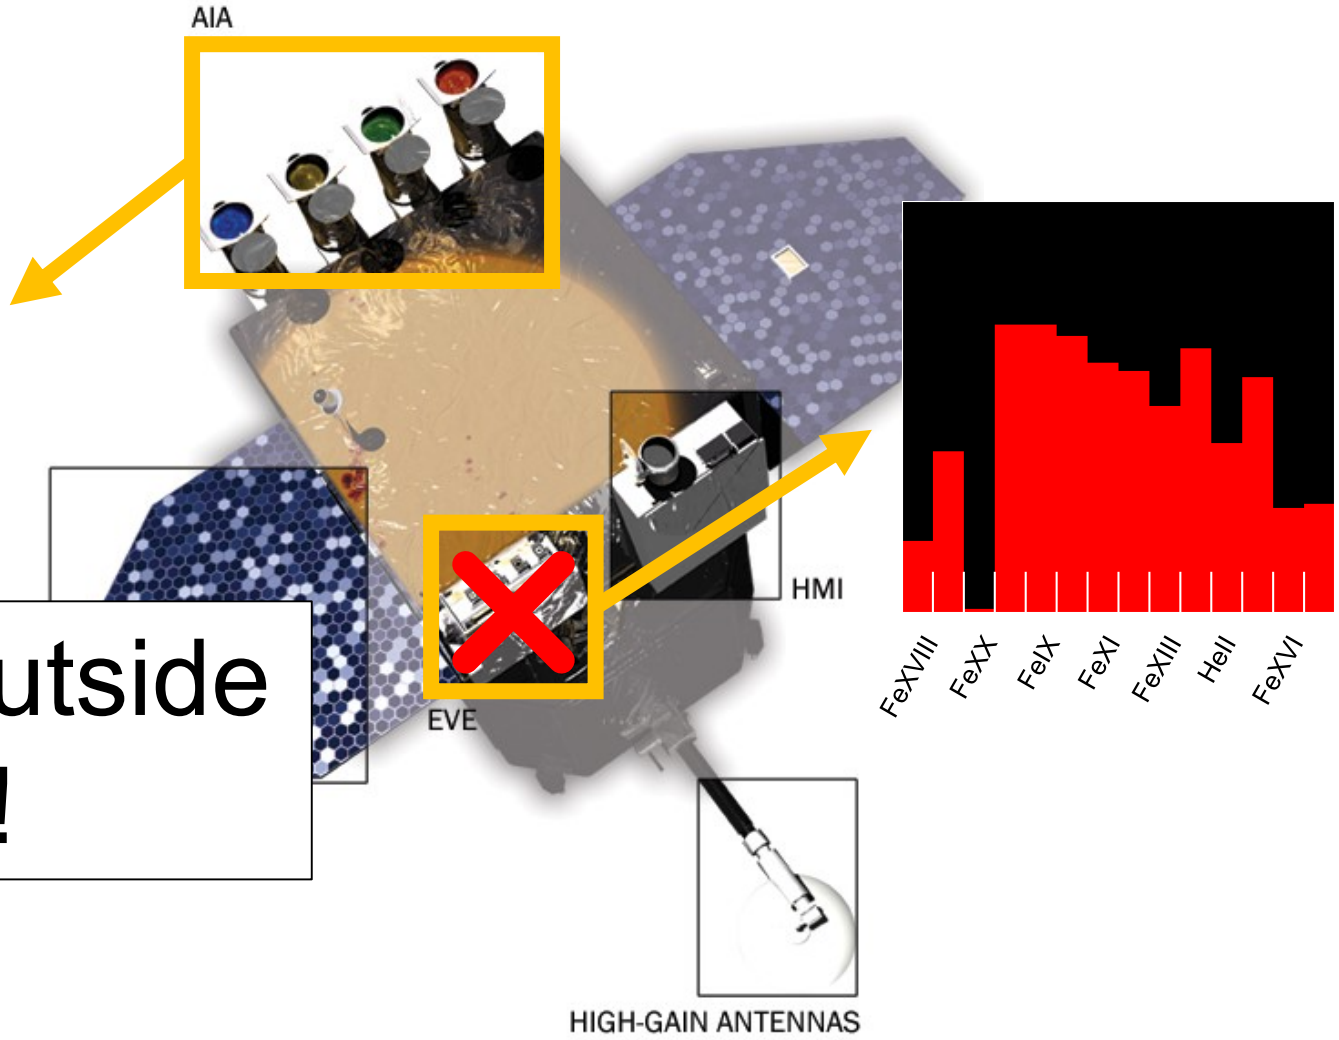

Think outside the box!

Image Credit: Wikipedia

Evaluation: Term Project

- Proposal due between Feb 14 – March 19. We will provide some inspiration. You can turn it in at any point and we will give you feedback quickly.
- Progress Report due April 4: what have you done, what is left?
- Final Project (code + report) due April 23 at the earliest (may give an extension).
- Poster Session during Exams.
- **Questions?**

Meetings

- Class:
  - Tue/Thu 10:30am – Noon, 1571 GGBL
- Discussion Section
  - Wed 5PM-6PM, 1571 GGBL
  - Mon 12:30PM – 1:30PM, 1200 EECS
- Office Hours
  - Professor: 10:30am-Noon Fridays (*BBB 3777*)
  - GSI/IAs: 3:00-4:30pm Tuesday, 2:30-4:00pm Thursday (*BBB Learning Center*)

Meetings

Mon	**Tue**	**Wed**	**Thu**	**Fri**
Discussion 12:30pm-1:30pm	*Class* 10:30am-12:00pm	*Discussion* 5:00pm-6:00pm	*Class* 10:30am-12:00pm	*Office Hours* 10:30am-12:00pm
	Office Hours 3:00pm-4:30pm		*Office Hours* *2:30pm-4pm*	

Questions?

Slide Credit: L. Lazebnik