CS 4476-B:
Computer Vision

Instructor: James Hays
TAs: **Ben Wilson** (head TA), Bharat Mamidibathula, Gunhyun Park, Jonathan Leo, Otis Smith, Pranav Khorana, Sukriti Bhardwaj, Tony Zhang, Xueqing Li, Yash Kothari, Yoonwoo Kim

Image by kirkh.deviantart.com

Today's Class

- Covid safety
- Who am I?
- What is Computer Vision?
- Specifics of this course
- Geometry of Image Formation
- Questions

New reported cases

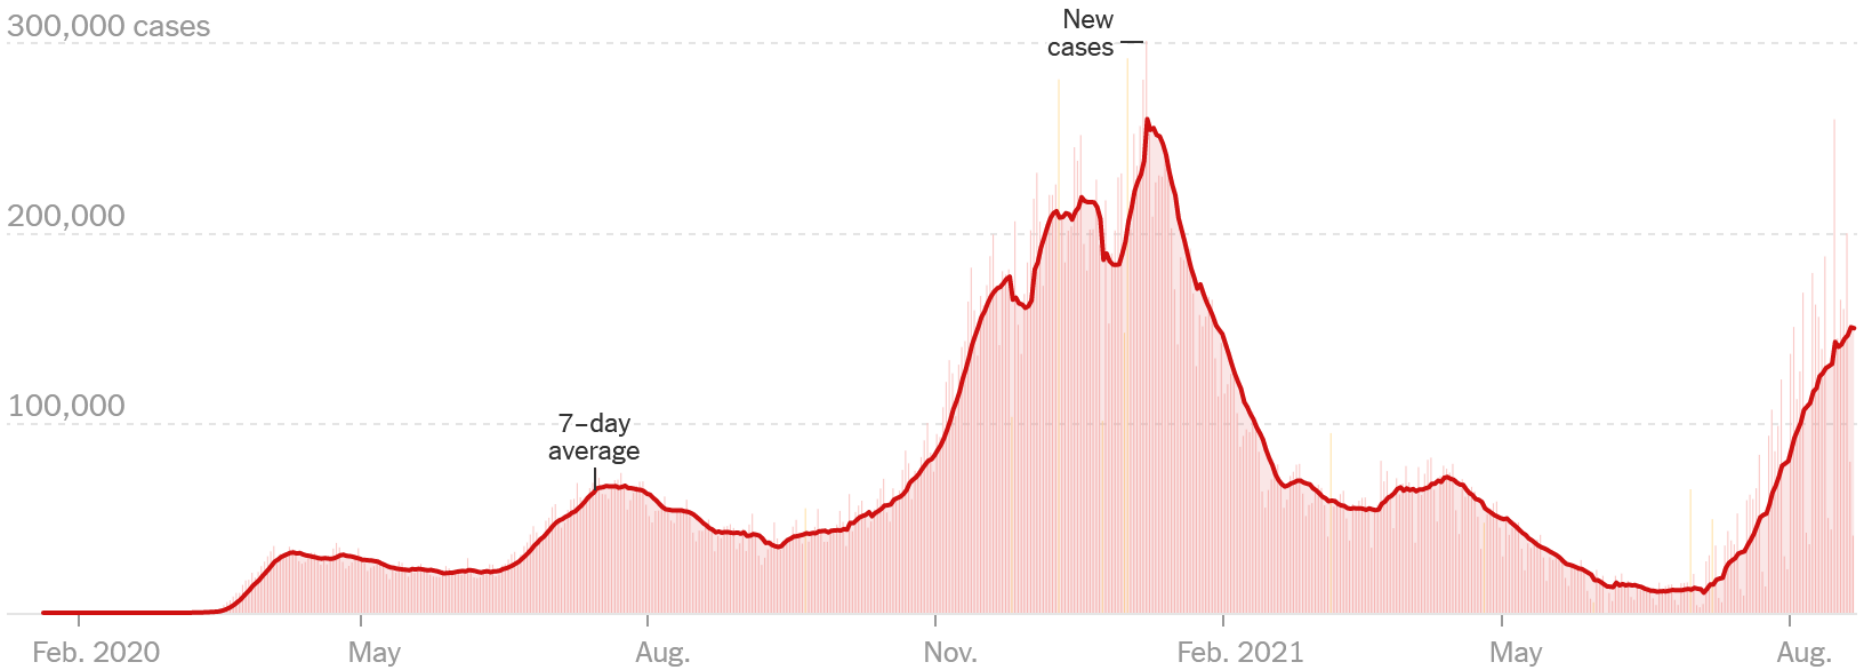

These are days with a reporting anomaly. Read more here.

Tests

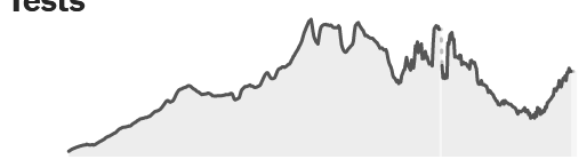

Hospitalized

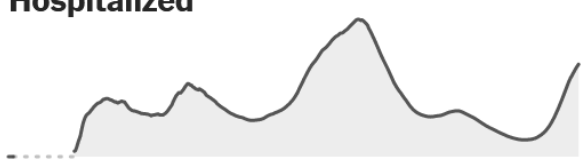

Deaths

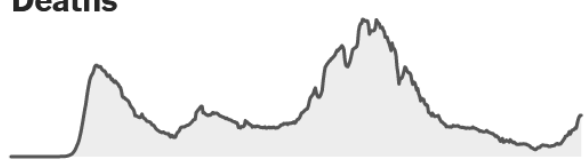

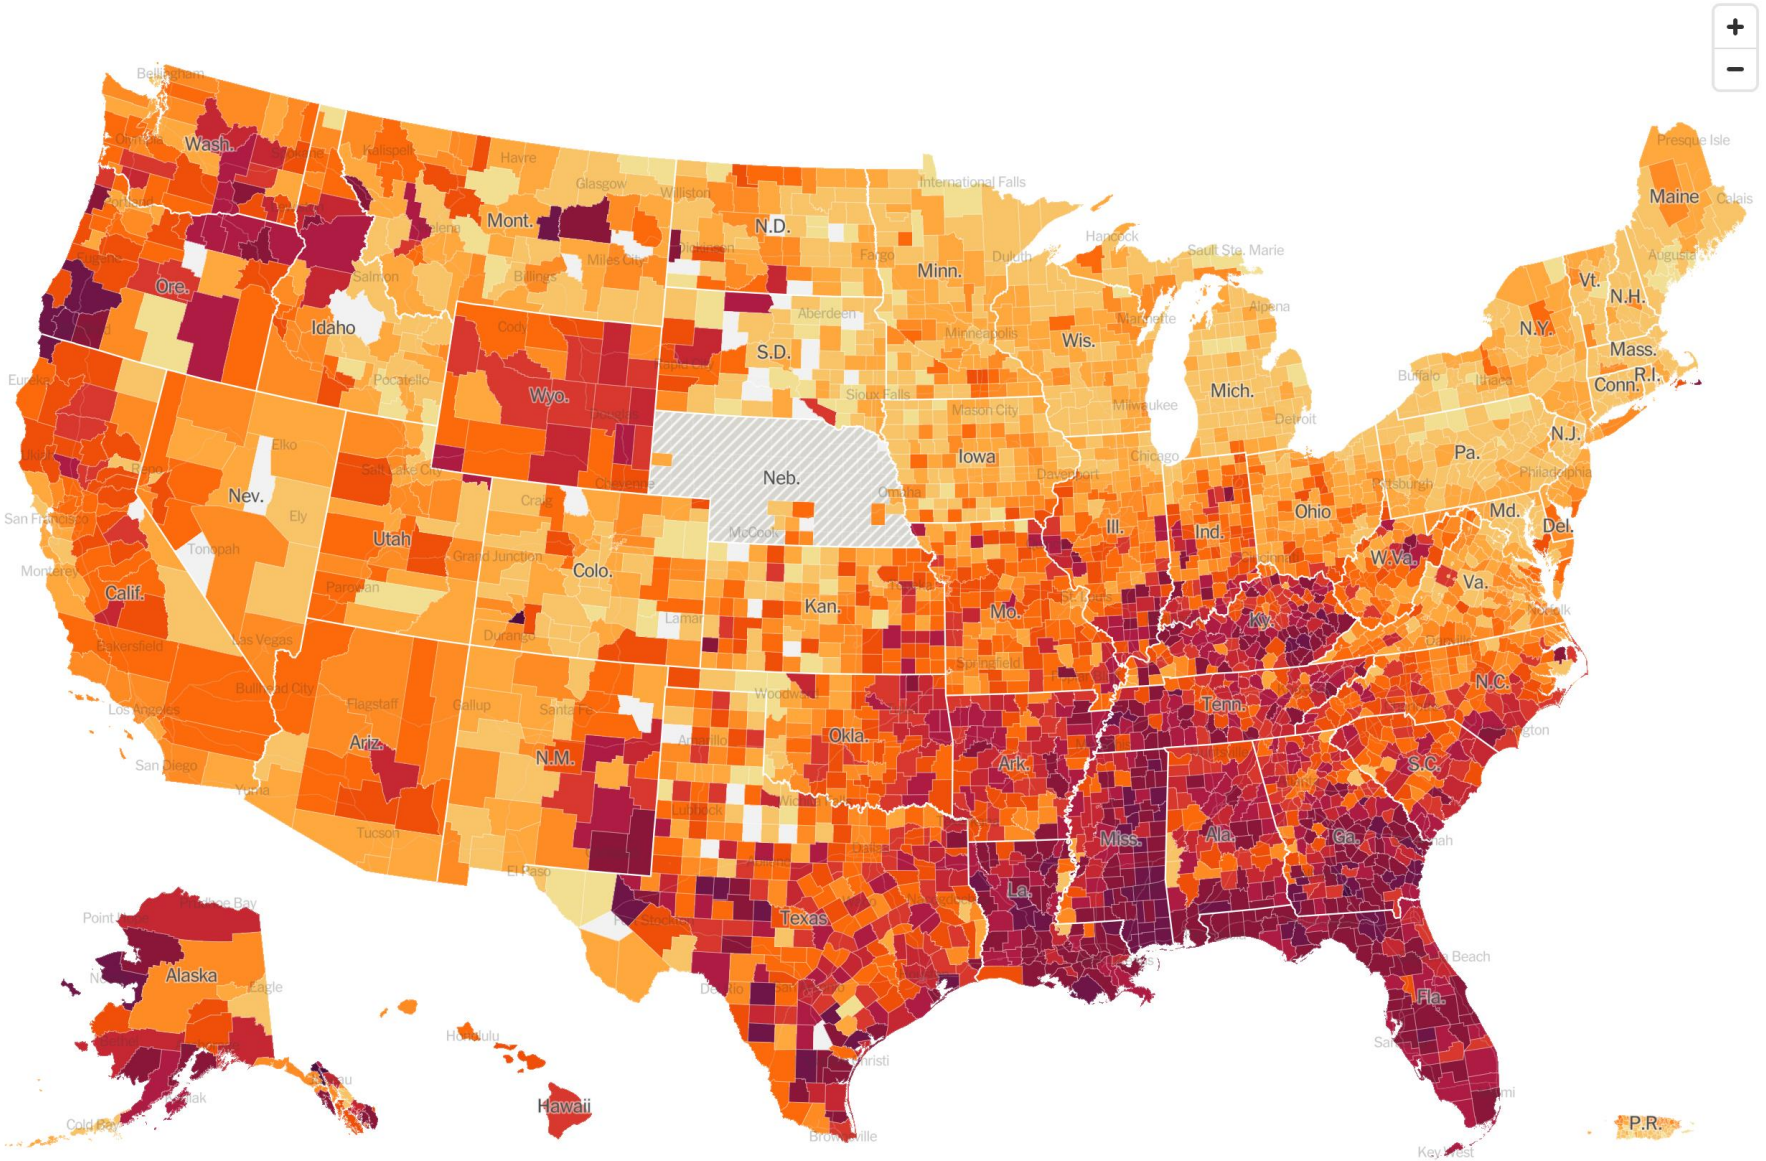

About this data

ON CAMPUS FALL 2021

Some Facts for Students

By Deven R. Desai, Scheller College

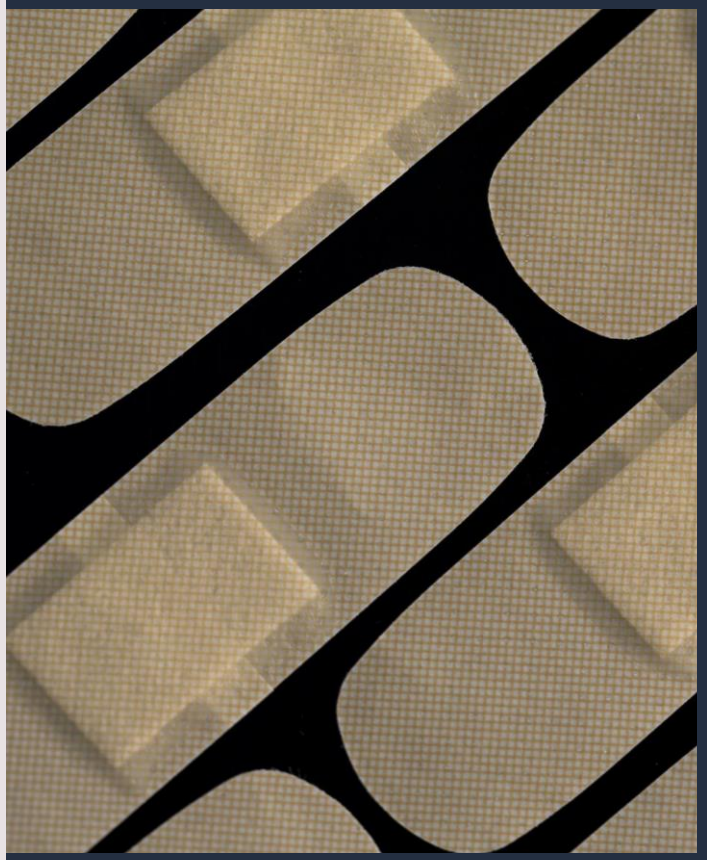

Layered Protection is the strongest protection

- Getting vaccinated is vital!
- Masking protects everyone!
- Getting tested weekly helps keep GA Tech open!
- Doing all three is the highest level of protection for you and the GA Tech community.
- Now, let’s dig in a bit.

Delta Is Different

The Delta Variant "is about a thousand times more infectious than the original strains of the virus"

Céline Gounder, a clinical assistant professor of medicine and infectious disease at NYU's Grossman School of Medicine, STATNews

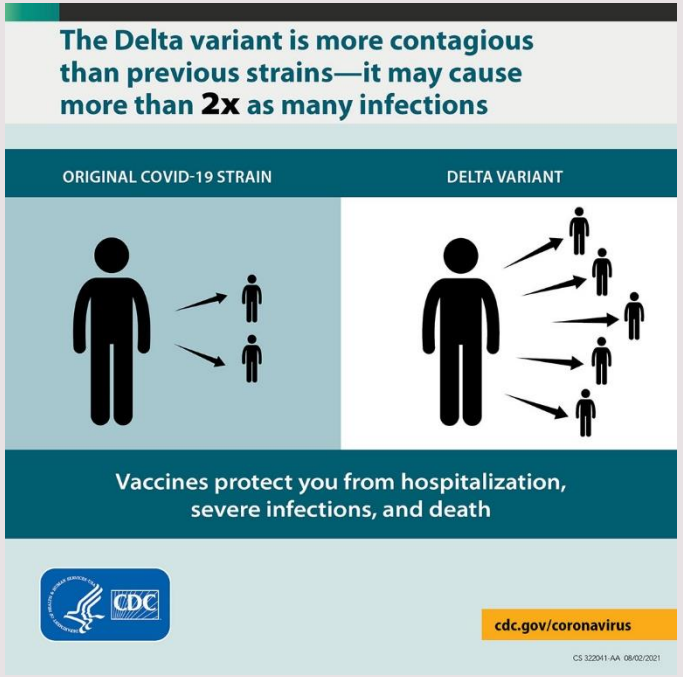

Getting vaccinated is vital!

"We know that the vaccines work," Kemp said. Fox5-Atlanta

“My family, myself, and other state leaders have all rolled up their sleeves and gotten their shot,” Kemp tweeted. “I encourage all Georgians who have concerns or questions to talk to a medical provider and get vaccinated as quickly as possible.” U.S. News

Getting vaccinated is vital!

Myth - Vaccination means no illness.

Reality: Vaccination prevents severe illness and reduces death.

Example: Provincetown outbreak - Of at least 965 positive cases that were traced to heavily vaccinated Provincetown, where around **60,000 people had gathered** for the holiday weekend, ***not a single death was reported and just seven people were hospitalized***. NY Times

Getting vaccinated is vital!

Vaccinated versus Unvaccinated outcomes

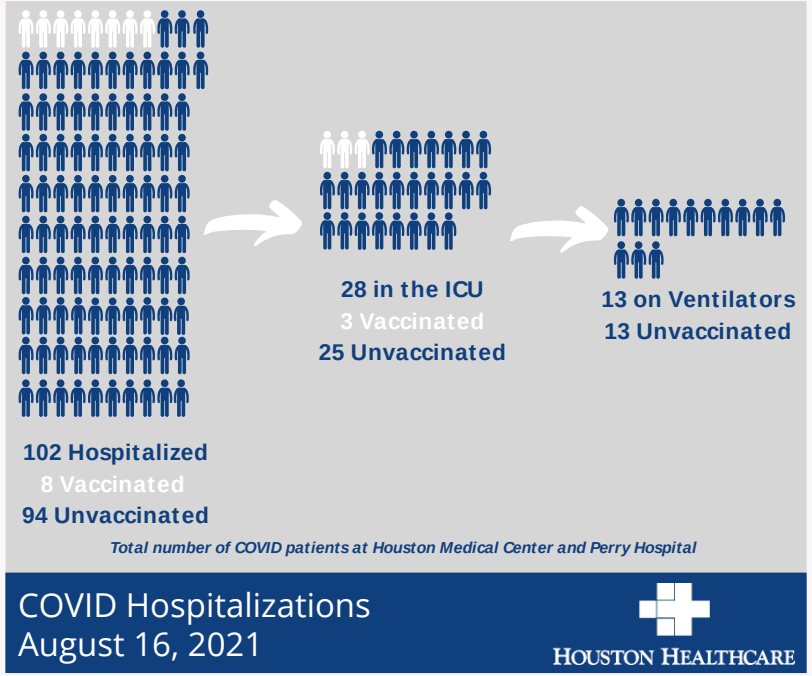

Houston Healthcare

Vaccination Protects Others and Reduces the Chance of Spreading Covid.

A recent study has shown that those who are fully vaccinated may carry the virus, and therefore be contagious, for fewer days than their unvaccinated counterparts. That suggests an even bigger overall difference in transmission between places with high and low vaccination rates. -- NY Times

In short: When you are vaccinated, you ensure that you are less likely to carry Covid to your family, friends, and those you love.

Masking protects everyone!

Q: WHY WEAR A MASK? AFTER ALL, IT’S NOT REQUIRED.

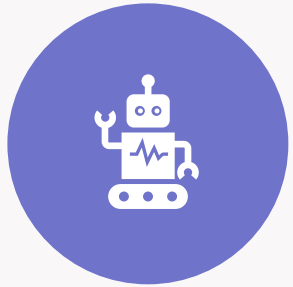

A: FACTS AND LOGIC.

Masking protects everyone! Facts

Exposure Matters: “The way to think about your exposure is dose times time. So your dose is a reflection of how much virus the person is carrying, but it’s also diluted in the air around them.”

Céline Gounder, a clinical assistant professor of medicine and infectious disease at NYU’s Grossman School of Medicine, STATNews

Breathing – “Airborne transmission arises through the inhalation of aerosol droplets exhaled by an infected person and is now thought to be the primary transmission route of COVID-19.” Bazant and Bush, PNAS, 2021

Masking protects everyone! Facts

1. You breathe through your nose and mouth.
2. “There's a lot of virus in my nose and throat, therefore, there's a lot of virus in the air that I cough out or breathe out.” Michael Marks, associate professor at the London School of Hygiene & Tropical Medicine, Gothamist.
3. To work, your mask **MUST** cover your NOSE and MOUTH.

Masking Protects You Inside a Room Facts – Nerd out GA Tech Style!

"**Our theoretical model quantifies the extent to which transmission risk is reduced in large rooms with high air exchange rates,** increased for more vigorous respiratory activities, and **dramatically reduced by the use of face masks**. Consideration of a number of outbreaks yields self-consistent estimates for the infectiousness of the new coronavirus." – Bazant and Bush, PNAS, 2021

Huh?

In short, the more vaccinated people in a room who are also masked, the lower the exposure risk.

Masking protects everyone! Logic

1. **The idea that you can not be masked while inside rested on *everyone being vaccinated*.**
2. **Problem: One cannot know whether everyone in the room is vaccinated.**
3. **Solution: When one cannot know the status of other folks, wearing your mask is the best move to**
    A. Decrease the chance of becoming ill and especially severely ill (as in hospitalized)
    B. Decrease the chance you might spread Covid to family, friends, and loved ones.

Conclusion

GA Tech continues to work to protect our community, but we all need to do our part. The simplest, strongest things you can do to protect our community, you, and our ability to stay open are

1. Get vaccinated,
2. Wear a mask, and
3. Get tested weekly.

United States | 18 - 29 Years

131,261

Total Admissions

Aug 01, 2020 - Aug 20, 2021

726

Current 7-Day Average

Aug 14, 2021 - Aug 20, 2021

715

Prior 7-Day Average

Aug 07, 2021 - Aug 13, 2021

744

Peak 7-Day Average

Aug 12, 2021 - Aug 18, 2021

+1.5%

Percent change from prior 7-day avg. of Aug 07, 2021 - Aug 13, 2021

-2.4%

Percent change from peak 7-day avg. of Aug 12, 2021 - Aug 18, 2021

New Admissions of Patients with Confirmed COVID-19 per 100,000 Population by Age Group, United States

Aug 01, 2020 - Aug 20, 2021

HHS Regions

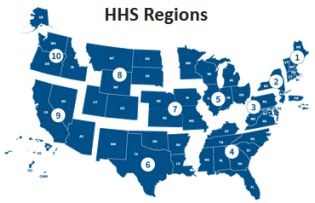

By Jurisdiction and Age Group | By Jurisdiction

Select a Jurisdiction: United States

Select an Age Group: 18 - 29 Years

United States | 18 - 29 Years

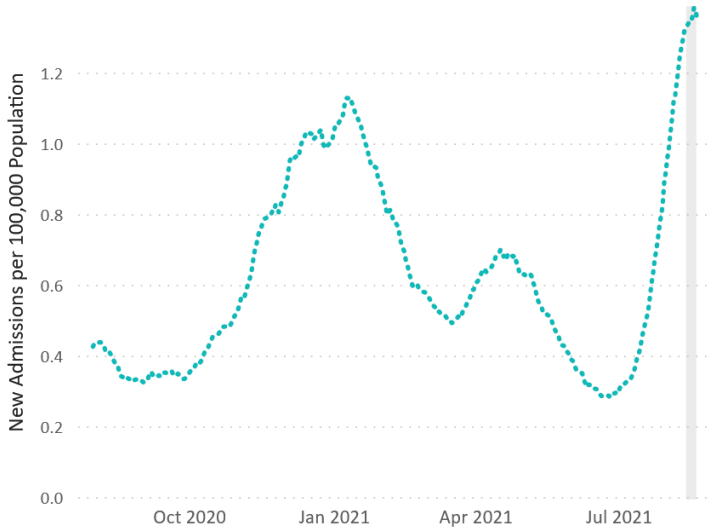

United States | All Age Groups

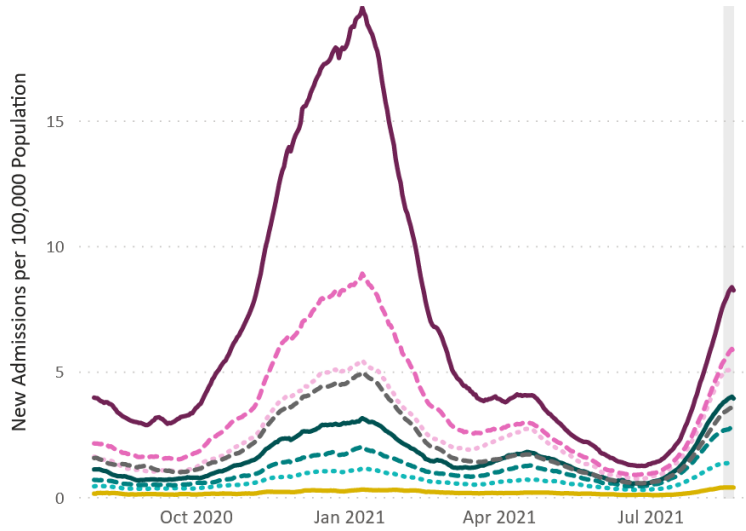

Age Group — 0 - 17 Years — 18 - 29 Years — 30 - 39 Years — 40 - 49 Years — 50 - 59 Years — 60 - 69 Years — 70+ Years — All Ages

Based on reporting from all hospitals (N=5,251). Due to potential reporting delays, data reported in the most recent 7 days (as represented by the shaded bar) should be interpreted with caution.
Small shifts in historic data may occur due to changes in the CMS Provider of Services file, which is used to identify the cohort of included hospitals. Data since December 1, 2020 have had error correction methodology applied. Data prior to this date may have anomalies that are still being resolved. Note that the above graphs are often shown on different scales. Data prior to August 1, 2020 are unavailable.

Last Updated: Aug 22, 2021

Unified Hospital Dataset, White House COVID-19 Team, Data Strategy and Execution Workgroup

U.S. COVID-19 vaccines by age

This chart shows the percentage of the U.S. population that has received a vaccination, broken down by age.

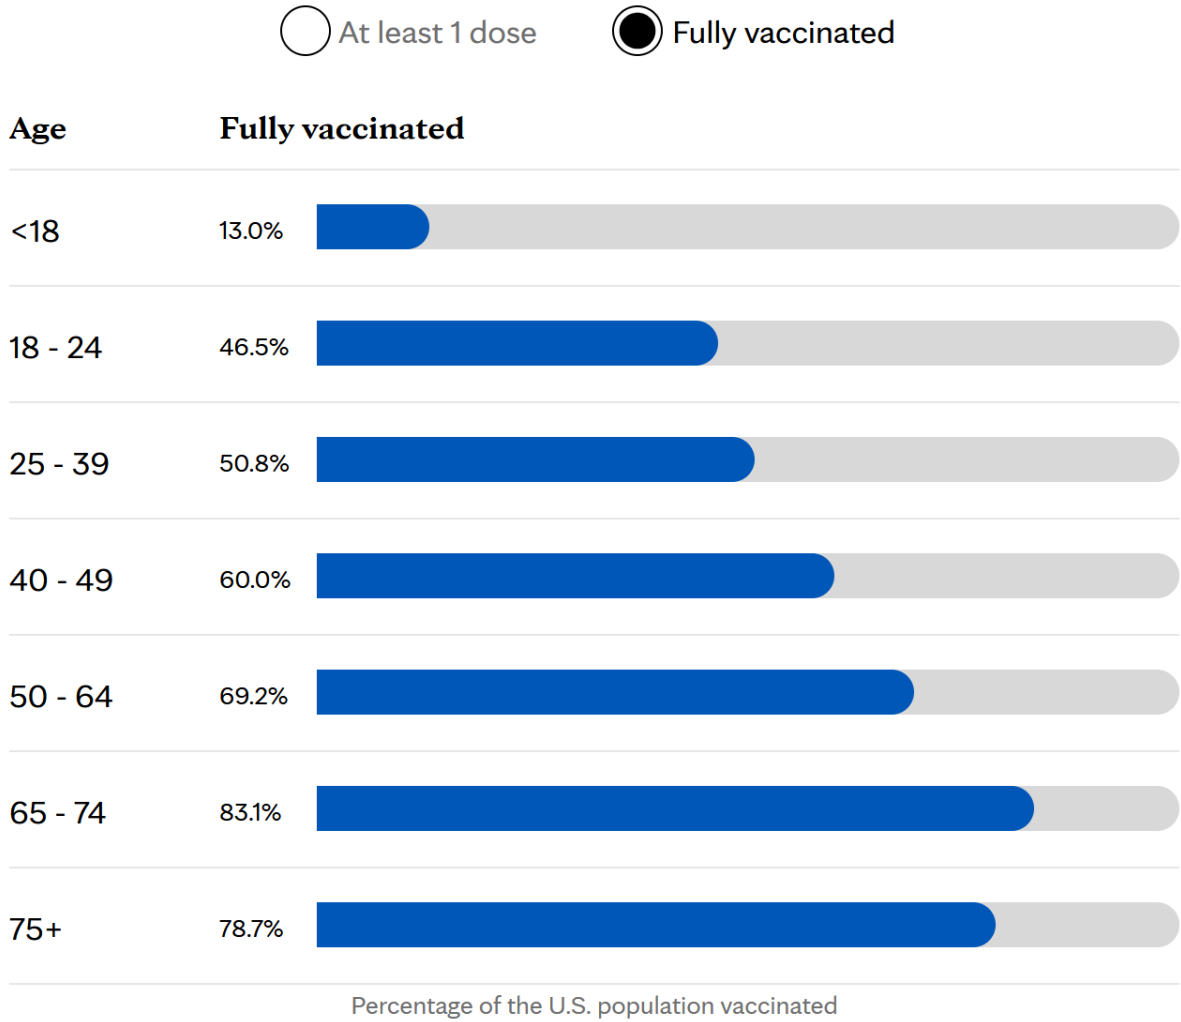

Carnegie Mellon University
Pittsburgh, PA
#1 in **Computer Science** (tie)
Save

University of Washington
Seattle, WA
#6 in **Computer Science** (tie)
Save

Columbia University
New York, NY
#13 in **Computer Science** (tie)
Save

University of Pennsylvania
Philadelphia, PA
#19 in **Computer Science**
Save

Massachusetts Institute of Technol
Cambridge, MA
#1 in **Computer Science** (tie)
Save

Georgia Institute of Technology
Atlanta, GA
#8 in **Computer Science** (tie)
Save

University of California--Los Angele
Los Angeles, CA
#13 in **Computer Science** (tie)
Save

Purdue University--West Lafayette
West Lafayette, IN
#20 in **Computer Science** (tie)
Save

Stanford University
Stanford, CA
#1 in **Computer Science** (tie)
Save

Princeton University
Princeton, NJ
#8 in **Computer Science** (tie)
Save

University of Wisconsin--Madison
Madison, WI
#13 in **Computer Science** (tie)
Save

Rice University
Houston, TX
#20 in **Computer Science** (tie)
Save

University of California--Berkeley
Berkeley, CA
#1 in **Computer Science** (tie)
Save

University of Texas--Austin
Austin, TX
#10 in **Computer Science**
Save

Harvard University
Cambridge, MA
#16 in **Computer Science** (tie)
Save

University of Massachusetts--Amherst
Amherst, MA
#20 in **Computer Science** (tie)
Save

University of Illinois--Urbana-Cham
Urbana, IL
#5 in **Computer Science**
Save

California Institute of Technology
Pasadena, CA
#11 in **Computer Science** (tie)
Save

University of California--San Diego
La Jolla, CA
#16 in **Computer Science** (tie)
Save

University of Southern California
Los Angeles, CA
#20 in **Computer Science** (tie)
Save

Cornell University
Ithaca, NY
#6 in **Computer Science** (tie)
Save

University of Michigan--Ann Arbor
Ann Arbor, MI
#11 in **Computer Science** (tie)
Save

University of Maryland--College Pa
College Park, MD
#16 in **Computer Science** (tie)
Save

Yale University
New Haven, CT
#20 in **Computer Science** (tie)
Save

Carnegie Mellon University
Pittsburgh, PA
#1 in **Computer Science** (tie)
Save

University of Washington
Seattle, WA
#6 in **Computer Science** (tie)
Save

Columbia University
New York, NY
#13 in **Computer Science** (tie)
Save

University of Pennsylvania
Philadelphia, PA
#19 in **Computer Science**
Save

Massachusetts Institute of Technol
Cambridge, MA
#1 in **Computer Science** (tie)
Save

Georgia Institute of Technology
Atlanta, GA
#8 in **Computer Science** (tie)
Save

University of California--Los Angele
Los Angeles, CA
#13 in **Computer Science** (tie)
Save

Purdue University--West Lafayette
West Lafayette, IN
#20 in **Computer Science** (tie)
Save

Stanford University
Stanford, CA
#1 in **Computer Science** (tie)
Save

Princeton University
Princeton, NJ
#8 in **Computer Science** (tie)
Save

University of Wisconsin--Madison
Madison, WI
#13 in **Computer Science** (tie)
Save

Rice University
Houston, TX
#20 in **Computer Science** (tie)
Save

University of California--Berkeley
Berkeley, CA
#1 in **Computer Science** (tie)
Save

University of Texas--Austin
Austin, TX
#10 in **Computer Science**
Save

Harvard University
Cambridge, MA
#16 in **Computer Science** (tie)
Save

University of Massachusetts--Amherst
Amherst, MA
#20 in **Computer Science** (tie)
Save

University of Illinois--Urbana-Cham
Urbana, IL
#5 in **Computer Science**
Save

California Institute of Technology
Pasadena, CA
#11 in **Computer Science** (tie)
Save

University of California--San Diego
La Jolla, CA
#16 in **Computer Science** (tie)
Save

University of Southern California
Los Angeles, CA
#20 in **Computer Science** (tie)
Save

Cornell University
Ithaca, NY
#6 in **Computer Science** (tie)
Save

University of Michigan--Ann Arbor
Ann Arbor, MI
#11 in **Computer Science** (tie)
Save

University of Maryland--College Par
College Park, MD
#16 in **Computer Science** (tie)
Save

Yale University
New Haven, CT
#20 in **Computer Science** (tie)
Save

For all of these reasons, in-person attendance is **discouraged** unless you are

- Masking
- Vaccinated
- Symptom free

Even doing all of those things, there is risk, so I am happy for students to attend lecture through Bluejeans

https://gatech.instructure.com/courses/217720/external_tools/717

Account
Dashboard
Courses
Calendar
Inbox
History
Commons
Help

Pages
Files
Syllabus
Outcomes
Rubrics
Quizzes
Modules
BigBlueButton (Formerly Conferences)
Collaborations
Attendance
Chat
Piazza
My Media
Media Gallery
BlueJeans
GaTech Roster
GradePub
CIOS
Badges
Well-Being Connect
GT Apps Labs
Mental Health Resources

in person for students who do not want to attend in person or who are under quarantine.

^ Every 1st Week; from August 23rd 2021 on Monday, Wednesday, at 12:30 pm for 1 hours 15 minutes till December 8th 2021

BlueJeans Meetings Currently In Progress

Name	Description	Ends At	Link

Recorded BlueJeans Meetings

Note: Recording links will not work until the file has been fully processed.

Name	Description	Recording Time	Actions

poll @6

Office hours poll

How would you prefer to attend TA office hours?

- [ ] In person, indoors
- [ ] In person, outdoors
- [ ] via Bluejeans

Submit

You have **not yet** voted.

Revoting is **allowed**. Select your vote and click submit to register your vote.

Your name will **not be visible to anyone.**

logistics

edit · good poll | 0

Office hours poll closes in 2 day(s)

A total of 0 votes in **0** hours

Today’s Class

~~Covid safety~~

Who am I?

What is Computer Vision?

Specifics of this course

Geometry of Image Formation

Questions

A bit about me

ARGO AI


ROTARY LIFT
ROTARY LIFT
ROTARY LIFT
RH
ing Systems Development

What type of stuff do I work on?

Understanding Lidar

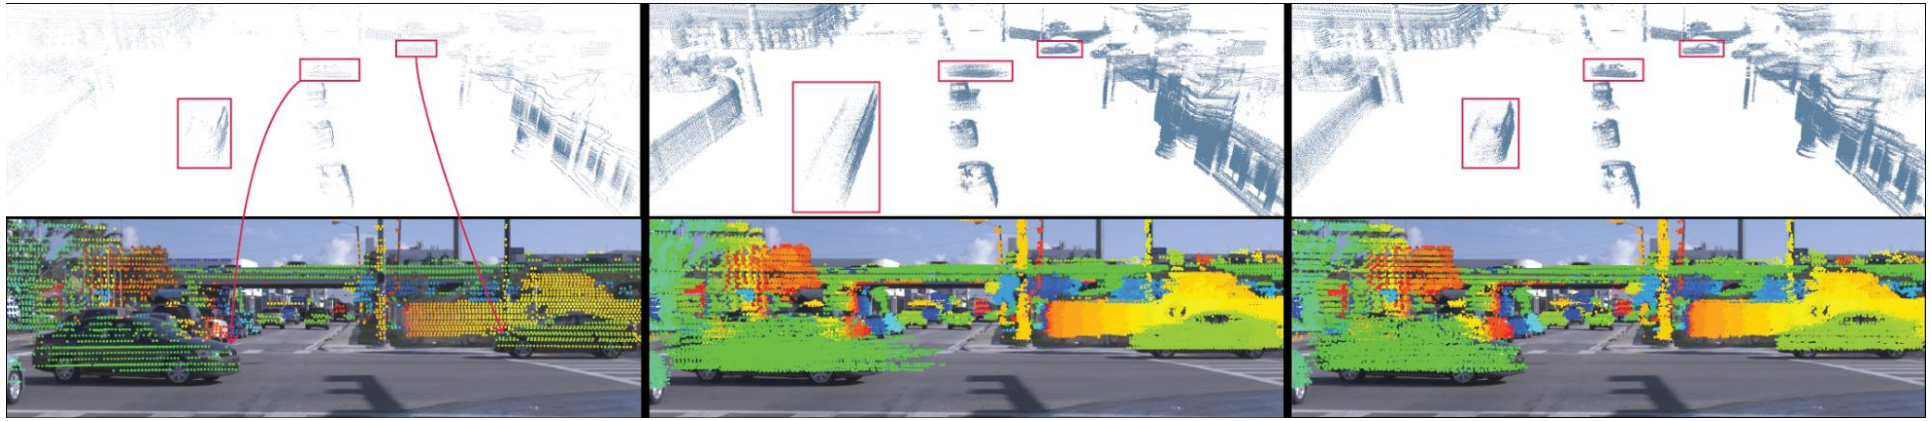

Scene Flow from Point Clouds with or without Learning
Jhony Kaesemodel Pontes, James Hays, Simon Lucey
https://jhonykaesemodel.com/publication/sceneflow-3dv2020/

Understanding Lidar

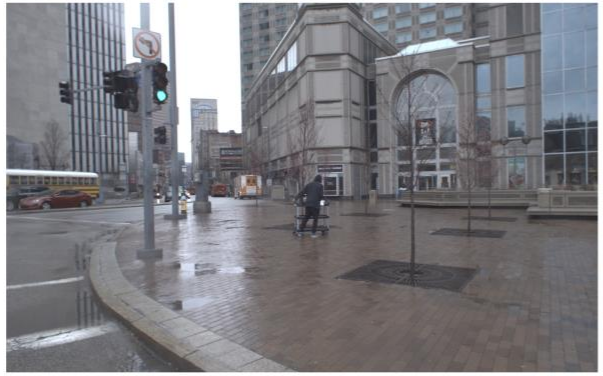

(a) Original camera image

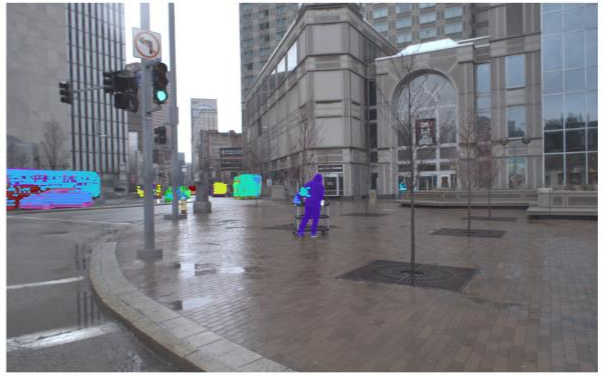

(b) Frustum proposals

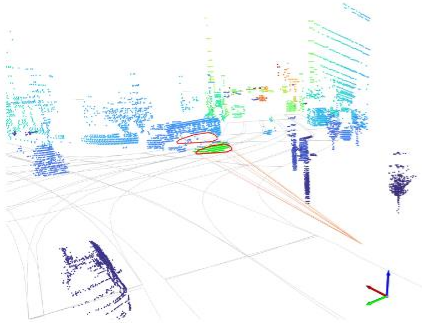

(c) Object frustum proposal

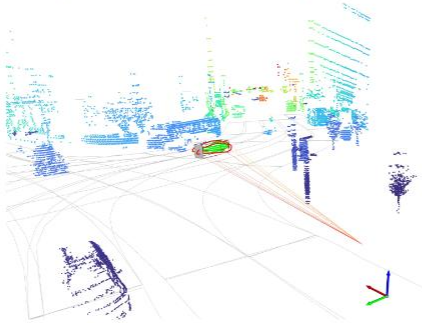

(d) LiDAR instance segmentation

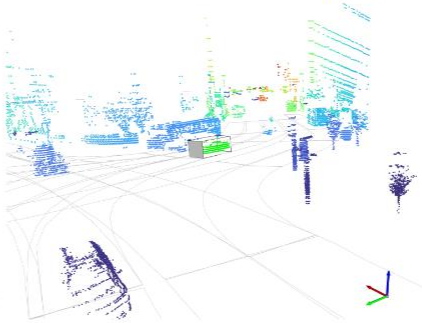

(e) Amodal completion

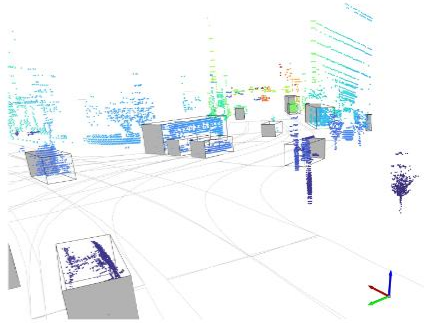

(f) Final cuboids

3D for Free: Crossmodal Transfer Learning using HD Maps
Benjamin Wilson, Zsolt Kira, James Hays
https://arxiv.org/abs/2008.10592

Exploring new data sources

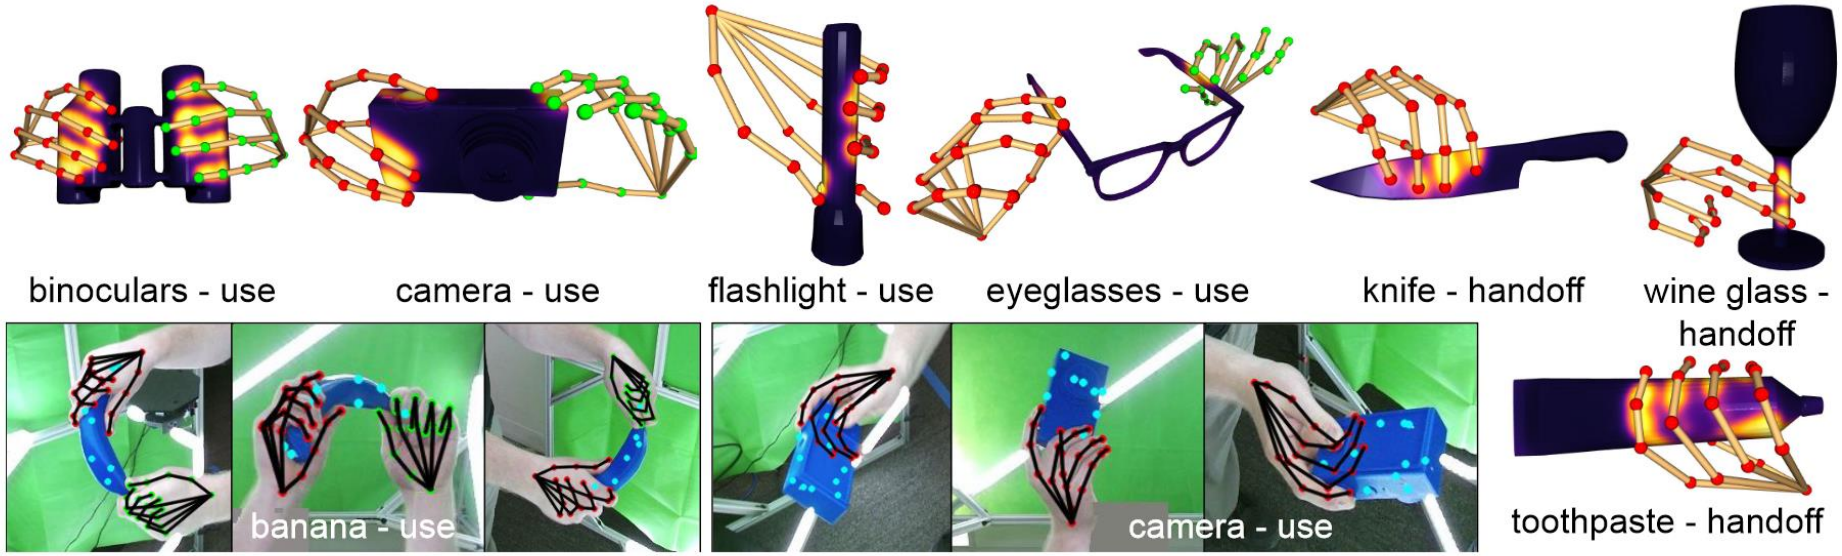

ContactPose: A Dataset of Grasps with Object Contact and Hand Pose
Samarth Brahmbhatt, Chengcheng Tang, Christopher D. Twigg, Charles C. Kemp, James Hays ECCV 2020

Exploring new data sources

LiDAR

- 2 roof-mounted LiDAR sensors
- Overlapping 40° vertical field of view
- Range of 200m
- On average, our LiDAR sensors produce a point cloud with ~ 107,000 points at 10 Hz

Cameras

- Seven high-resolution ring cameras (1920 x 1200) recording at 30 Hz with a combined 360° field of view
- Two front-view facing stereo cameras (2056 x 2464) sampled at 5 Hz

Localization

We use a city-specific coordinate system for vehicle localization. We include 6-DOF localization for each timestamp, from a combination of GPS-based and sensor-based localization methods.

Calibration

Sensor measurements for each driving session are stored in "logs." For each log, we provide intrinsic and extrinsic calibration data for LiDAR and all nine cameras.

https://www.argoverse.org/

Today's Class

- ~~Covid safety~~
- ~~Who am I?~~
- What is Computer Vision?
- Specifics of this course
- Geometry of Image Formation
- Questions

What is Computer Vision?

Derogatory summary of computer vision:
Machine learning applied to visual data

Computer Vision

- Automatic understanding of images and video
  1. Computing properties of the 3D world from visual data *(measurement)*

Slide credit: Kristen Grauman

1. Vision for measurement

Real-time stereo

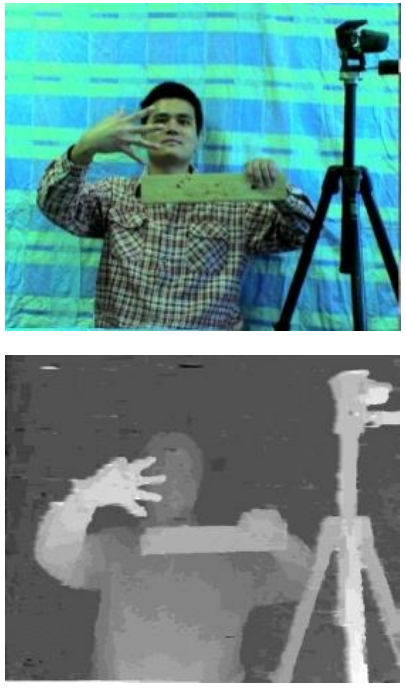

Wang et al.

Structure from motion

Snavely et al.

Tracking

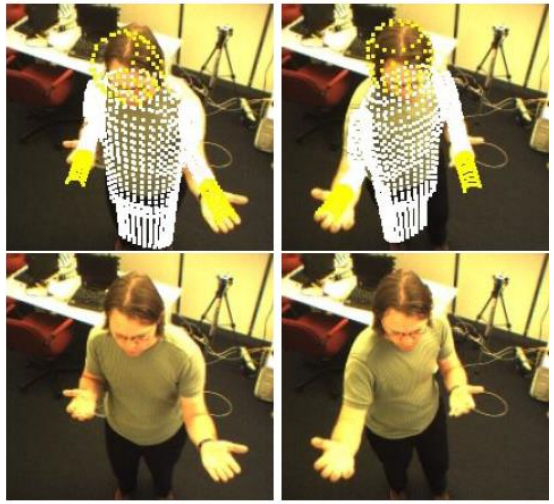

Demirdjian et al.

Slide credit: Kristen Grauman

Computer Vision

- Automatic understanding of images and video
  1. Computing properties of the 3D world from visual data *(measurement)*
  2. Algorithms and representations to allow a machine to recognize objects, people, scenes, and activities. *(perception and interpretation)*

Slide credit: Kristen Grauman

2. Vision for perception, interpretation

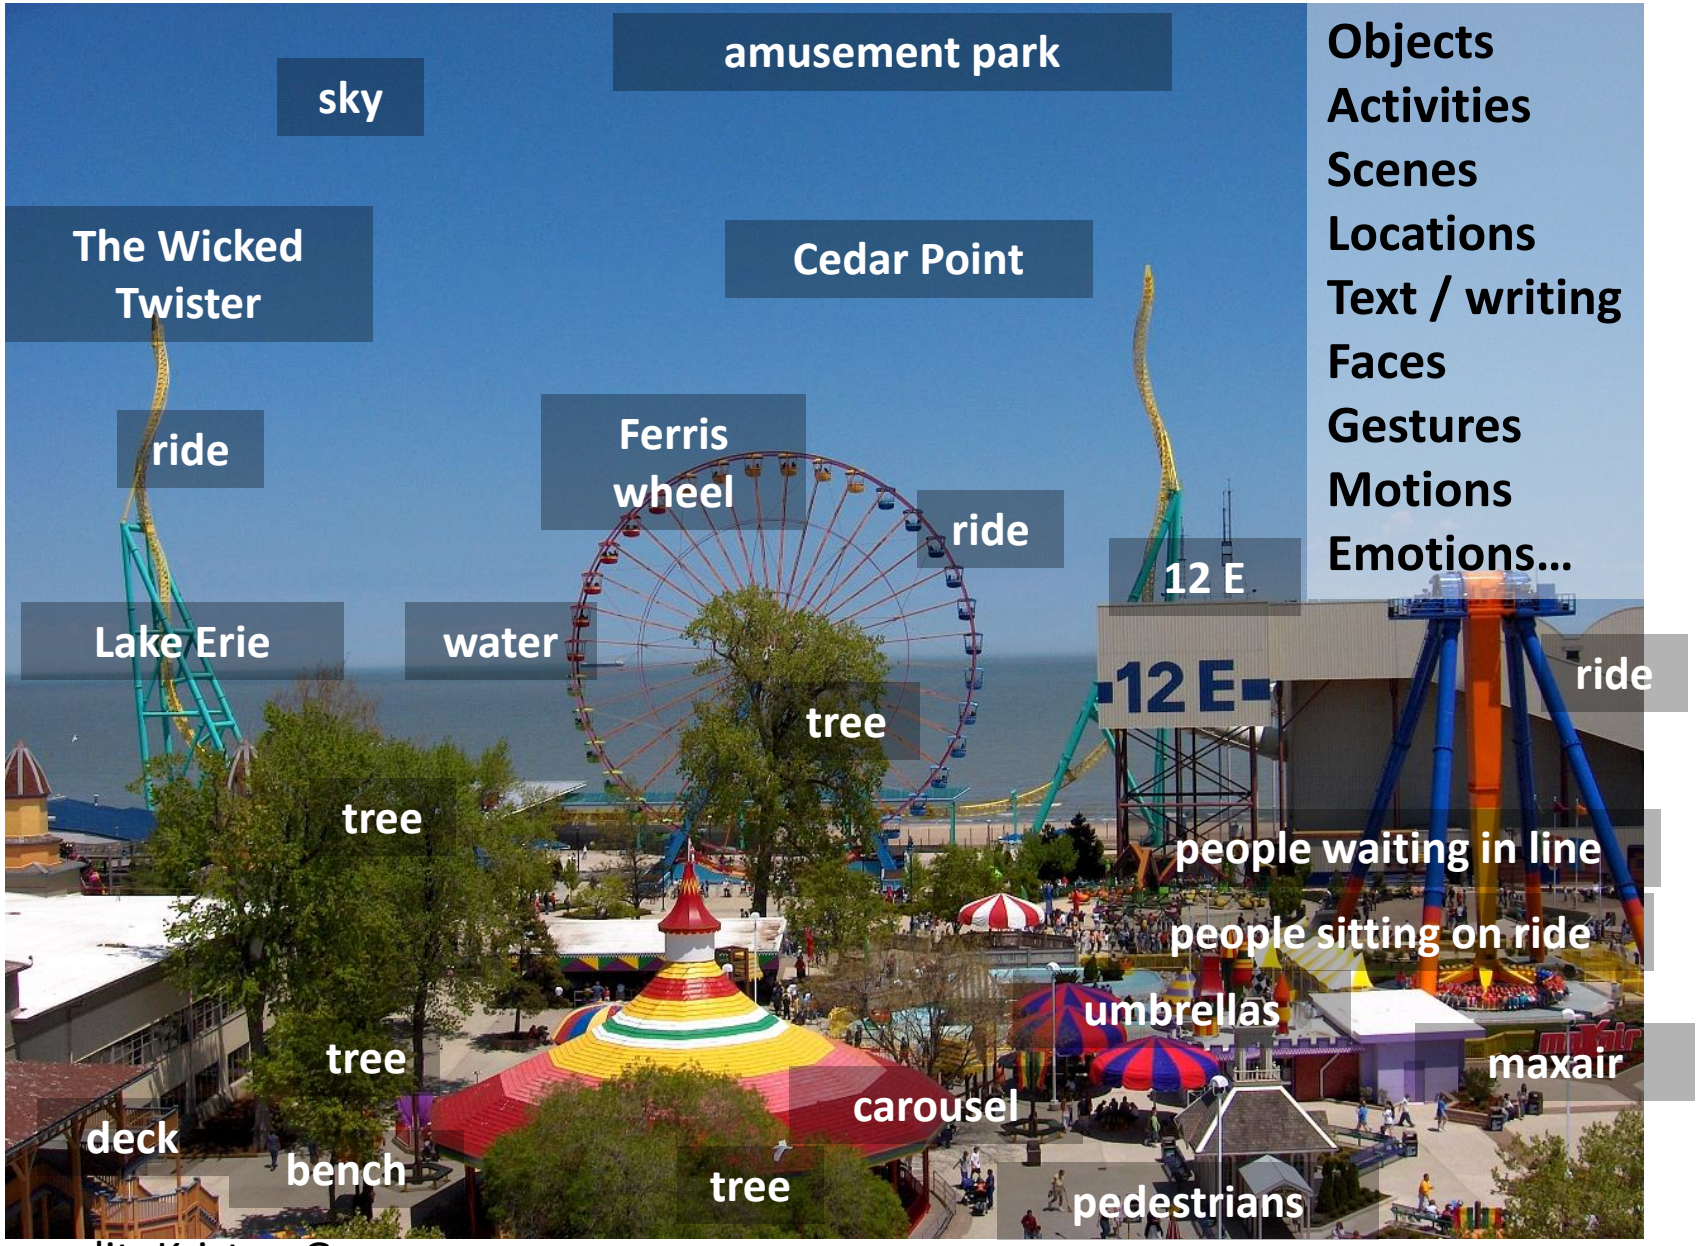

Slide credit: Kristen Grauman

Computer Vision

- Automatic understanding of images and video
  1. Computing properties of the 3D world from visual data *(measurement)*
  2. Algorithms and representations to allow a machine to recognize objects, people, scenes, and activities. *(perception and interpretation)*
  3. Algorithms to mine, search, and interact with visual data (*search and organization*)

Slide credit: Kristen Grauman

3. Visual search, organization

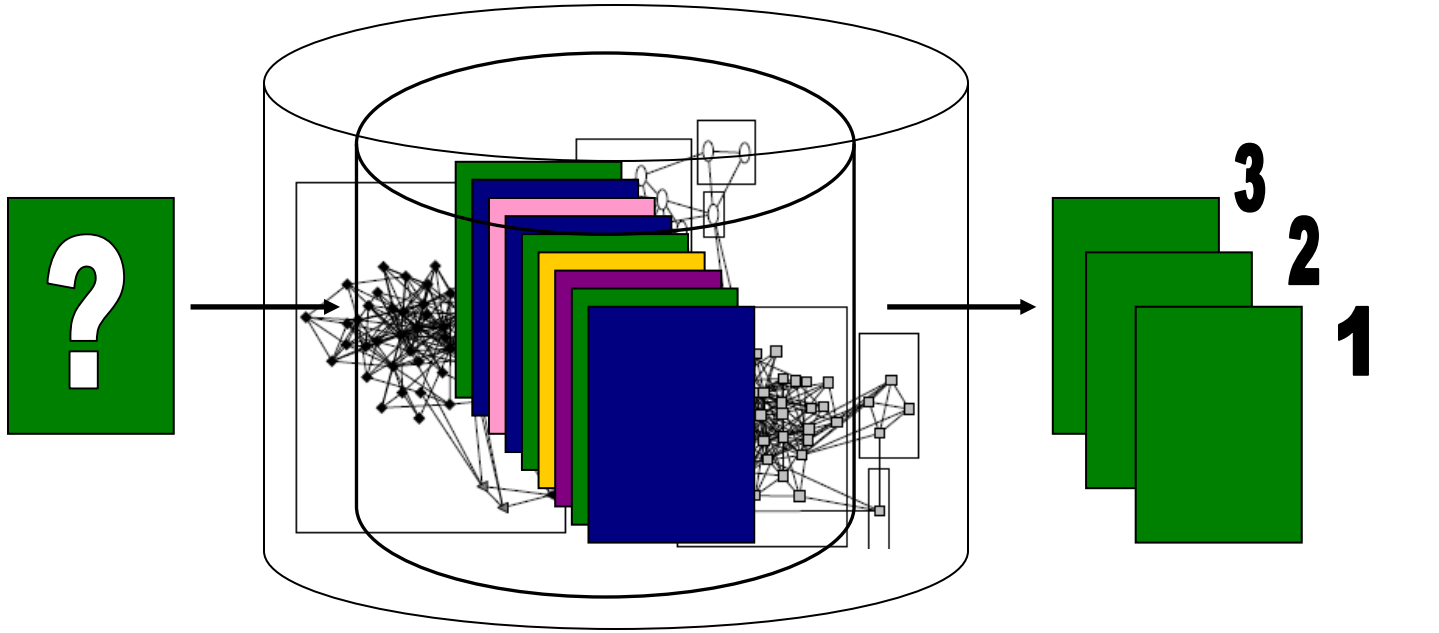

Slide credit: Kristen Grauman

Related disciplines

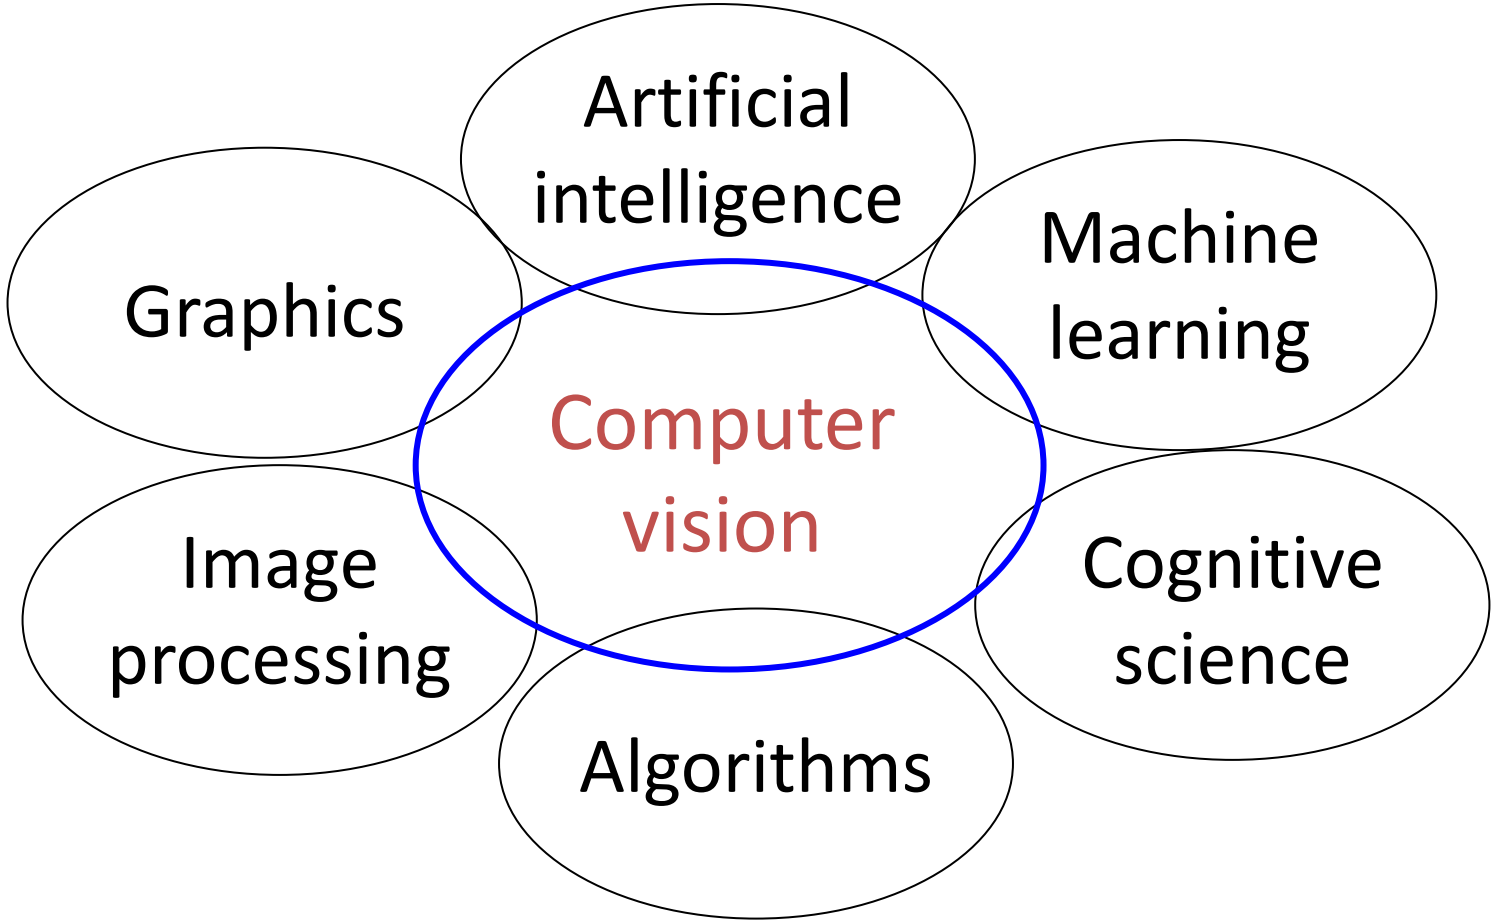

Slide credit: Kristen Grauman

Vision and graphics

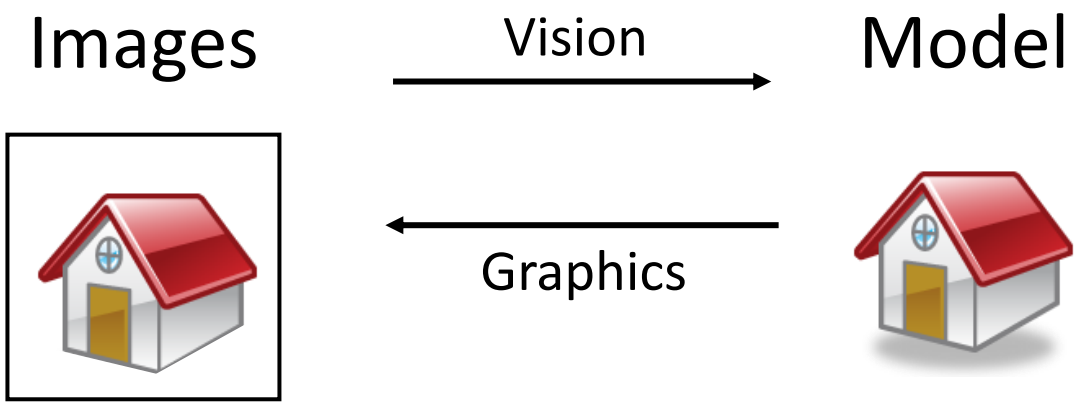

Inverse problems: analysis and synthesis.

Slide credit: Kristen Grauman

What humans see

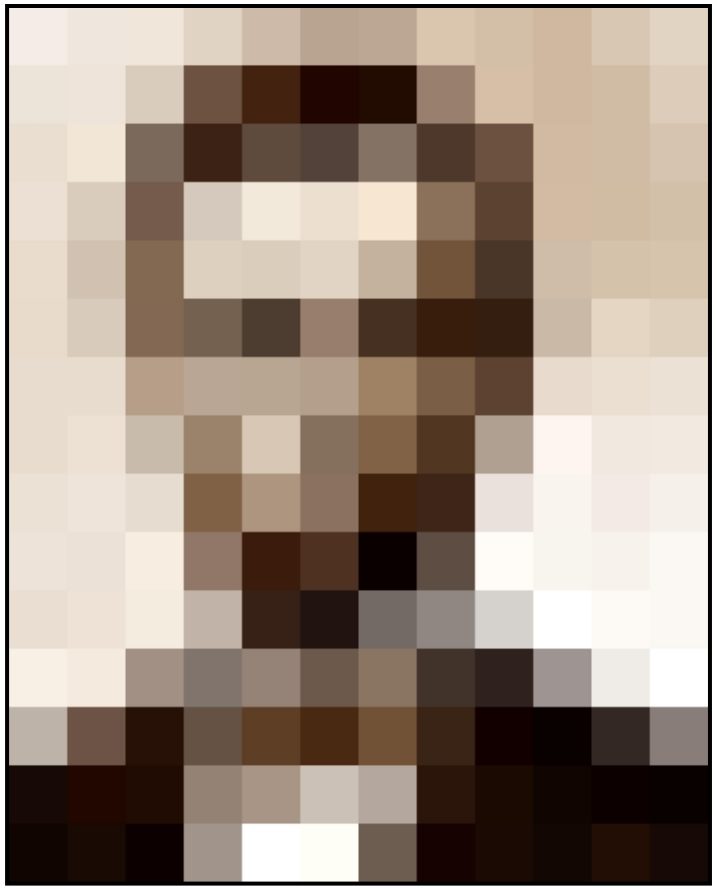

Slide credit: Larry Zitnick

What computers see

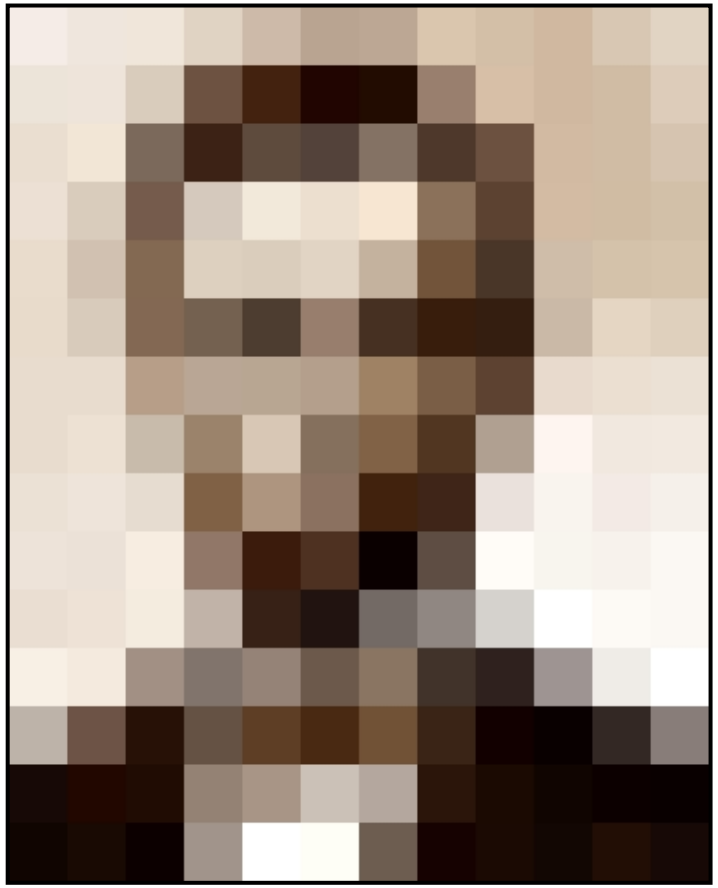

Slide credit: Larry Zitnick

What do humans see?

Slide credit: Larry Zitnick

Vision is really hard

- Vision is an amazing feat of natural intelligence
  - Visual cortex occupies about 50% of Macaque brain
  - One third of human brain devoted to vision (more than anything else)

Ridiculously brief history of computer vision

- 1966: Minsky assigns computer vision as an undergrad summer project
- 1960's: interpretation of synthetic worlds
- 1970's: some progress on interpreting selected images
- 1980's: ANNs come and go; shift toward geometry and increased mathematical rigor
- 1990's: face recognition; statistical analysis in vogue
- 2000's: broader recognition; large annotated datasets available; video processing starts
- 2010's: Deep learning with ConvNets
- 2020's: Widespread autonomous vehicles?
- 2030's: robot uprising?

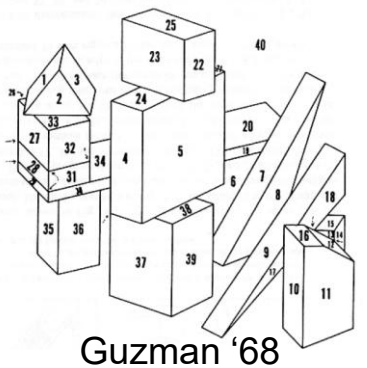

Guzman '68

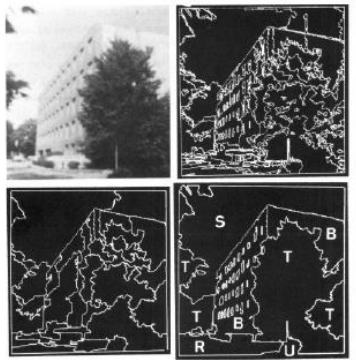

Ohta Kanade '78

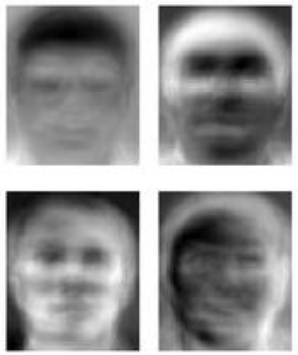

Turk and Pentland '91

How vision is used now

- Examples of real-world applications

Some of the following slides by Steve Seitz

Optical character recognition (OCR)

Technology to convert scanned docs to text

- If you have a scanner, it probably came with OCR software

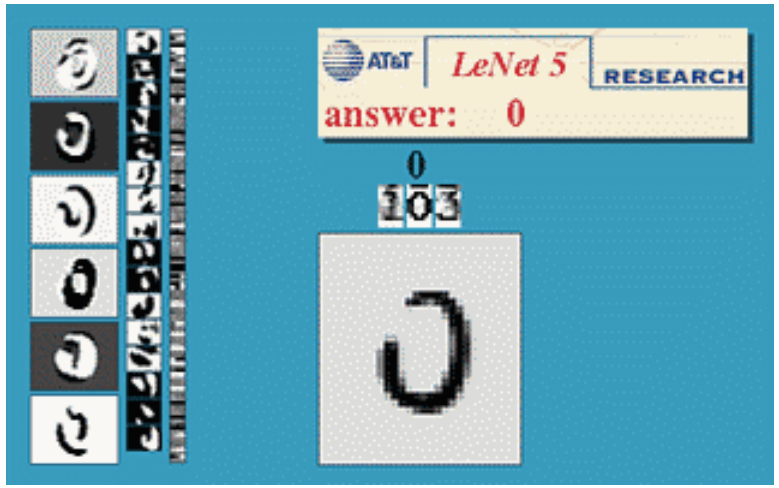

Digit recognition, AT&T labs
http://www.research.att.com/~yann/

License plate readers
http://en.wikipedia.org/wiki/Automatic_number_plate_recognition

Face detection

- Digital cameras detect faces

Vision in space

NASA'S Mars Exploration Rover Spirit captured this westward view from atop a low plateau where Spirit spent the closing months of 2007.

Vision systems (JPL) used for several tasks

- Panorama stitching
- 3D terrain modeling
- Obstacle detection, position tracking
- For more, read "Computer Vision on Mars" by Matthies et al.

iNaturalist

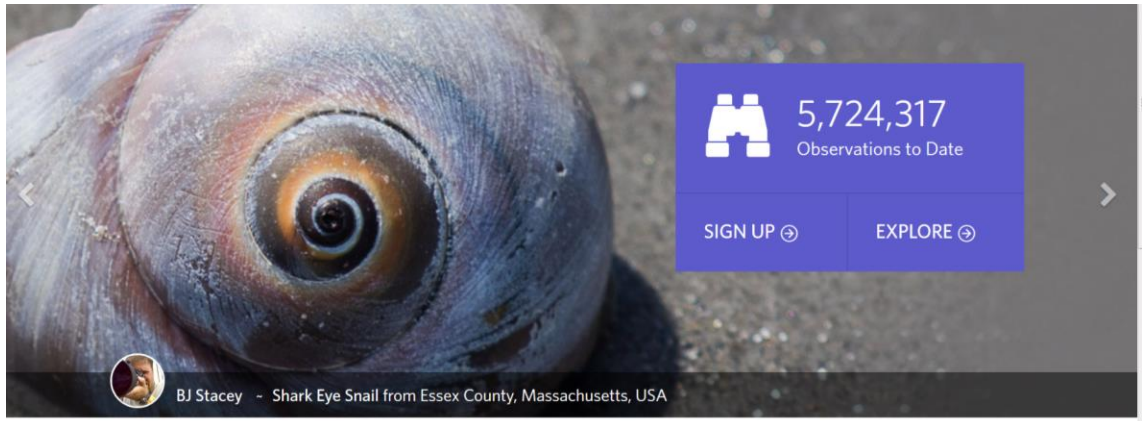

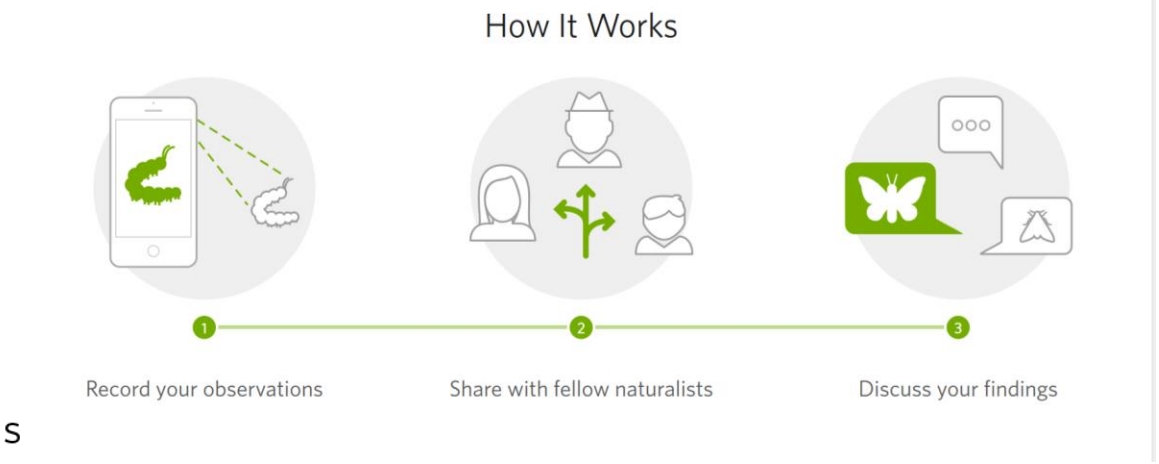

https://www.inaturalist.org/pages/computer_vision_demo

Skydio

https://www.skydio.com/

Zoox Computer Vision Demo

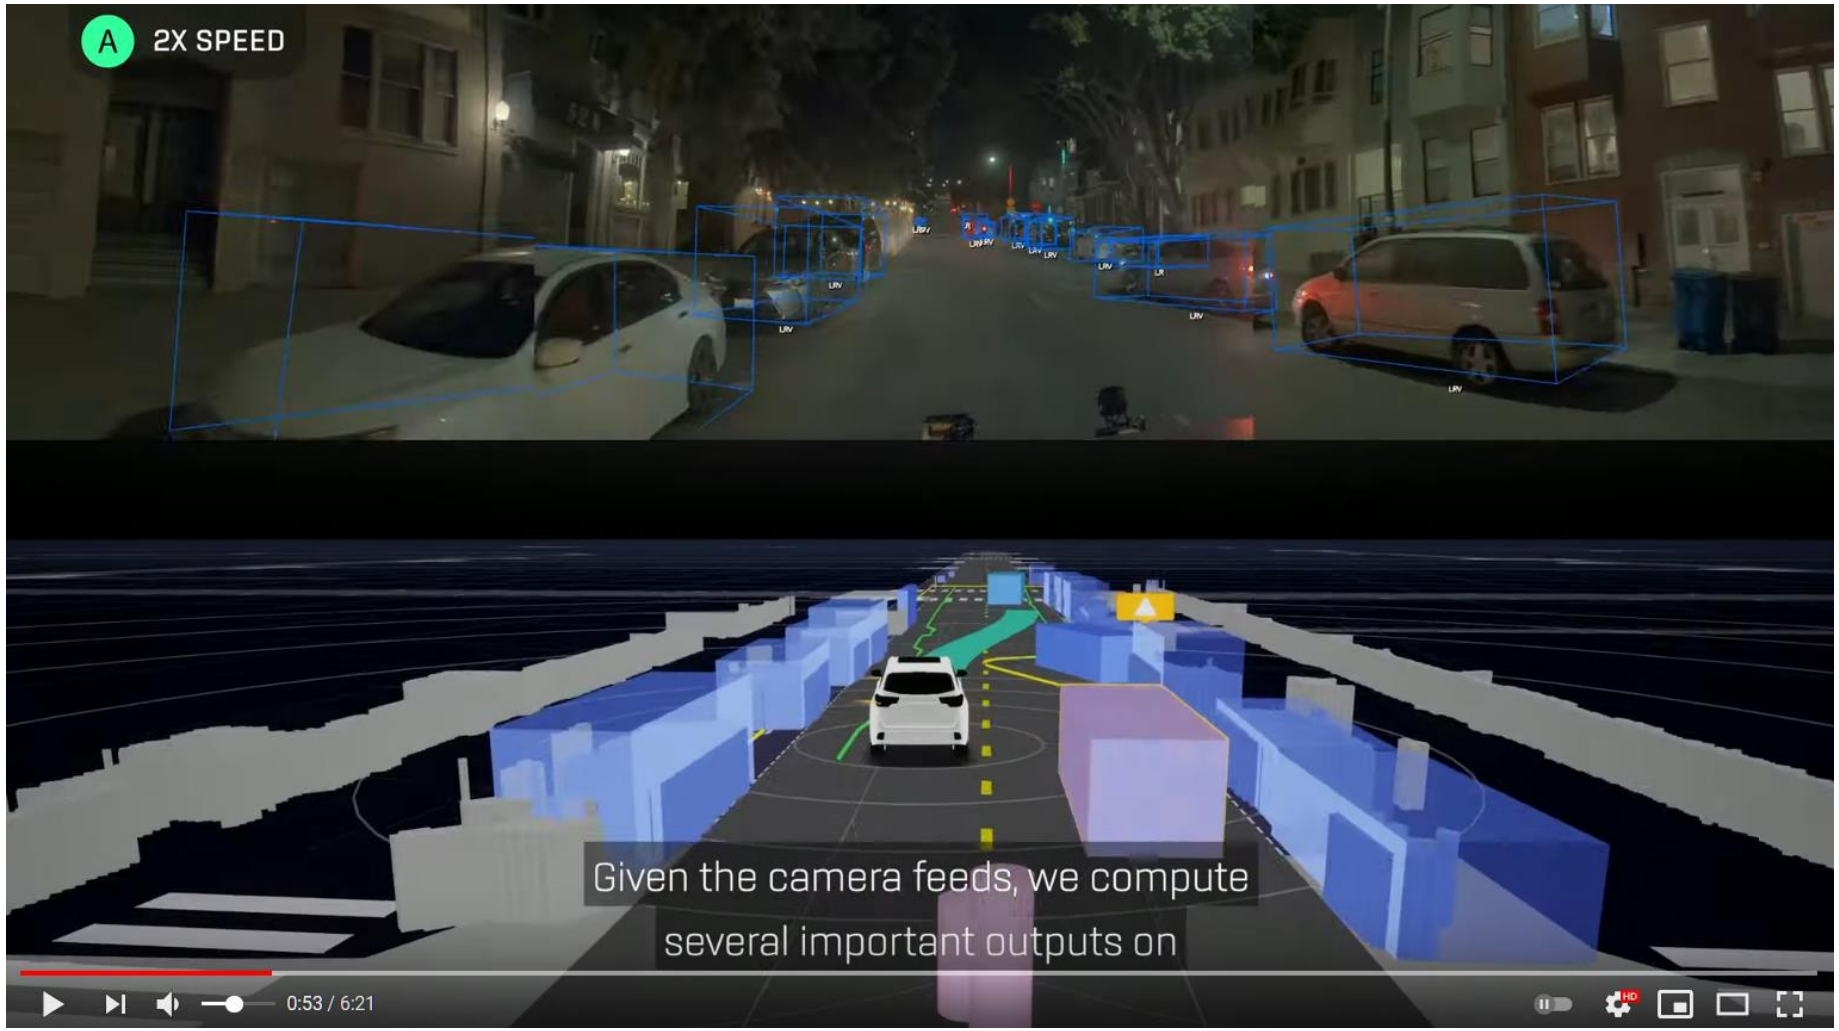

https://www.youtube.com/watch?v=BVRMh9NO9Cs

State of the art today?

With enough training data, computer vision ~~nearly~~ matches human vision at most recognition tasks

Deep learning has been an enormous disruption to the field. More and more techniques are being “deepified”.

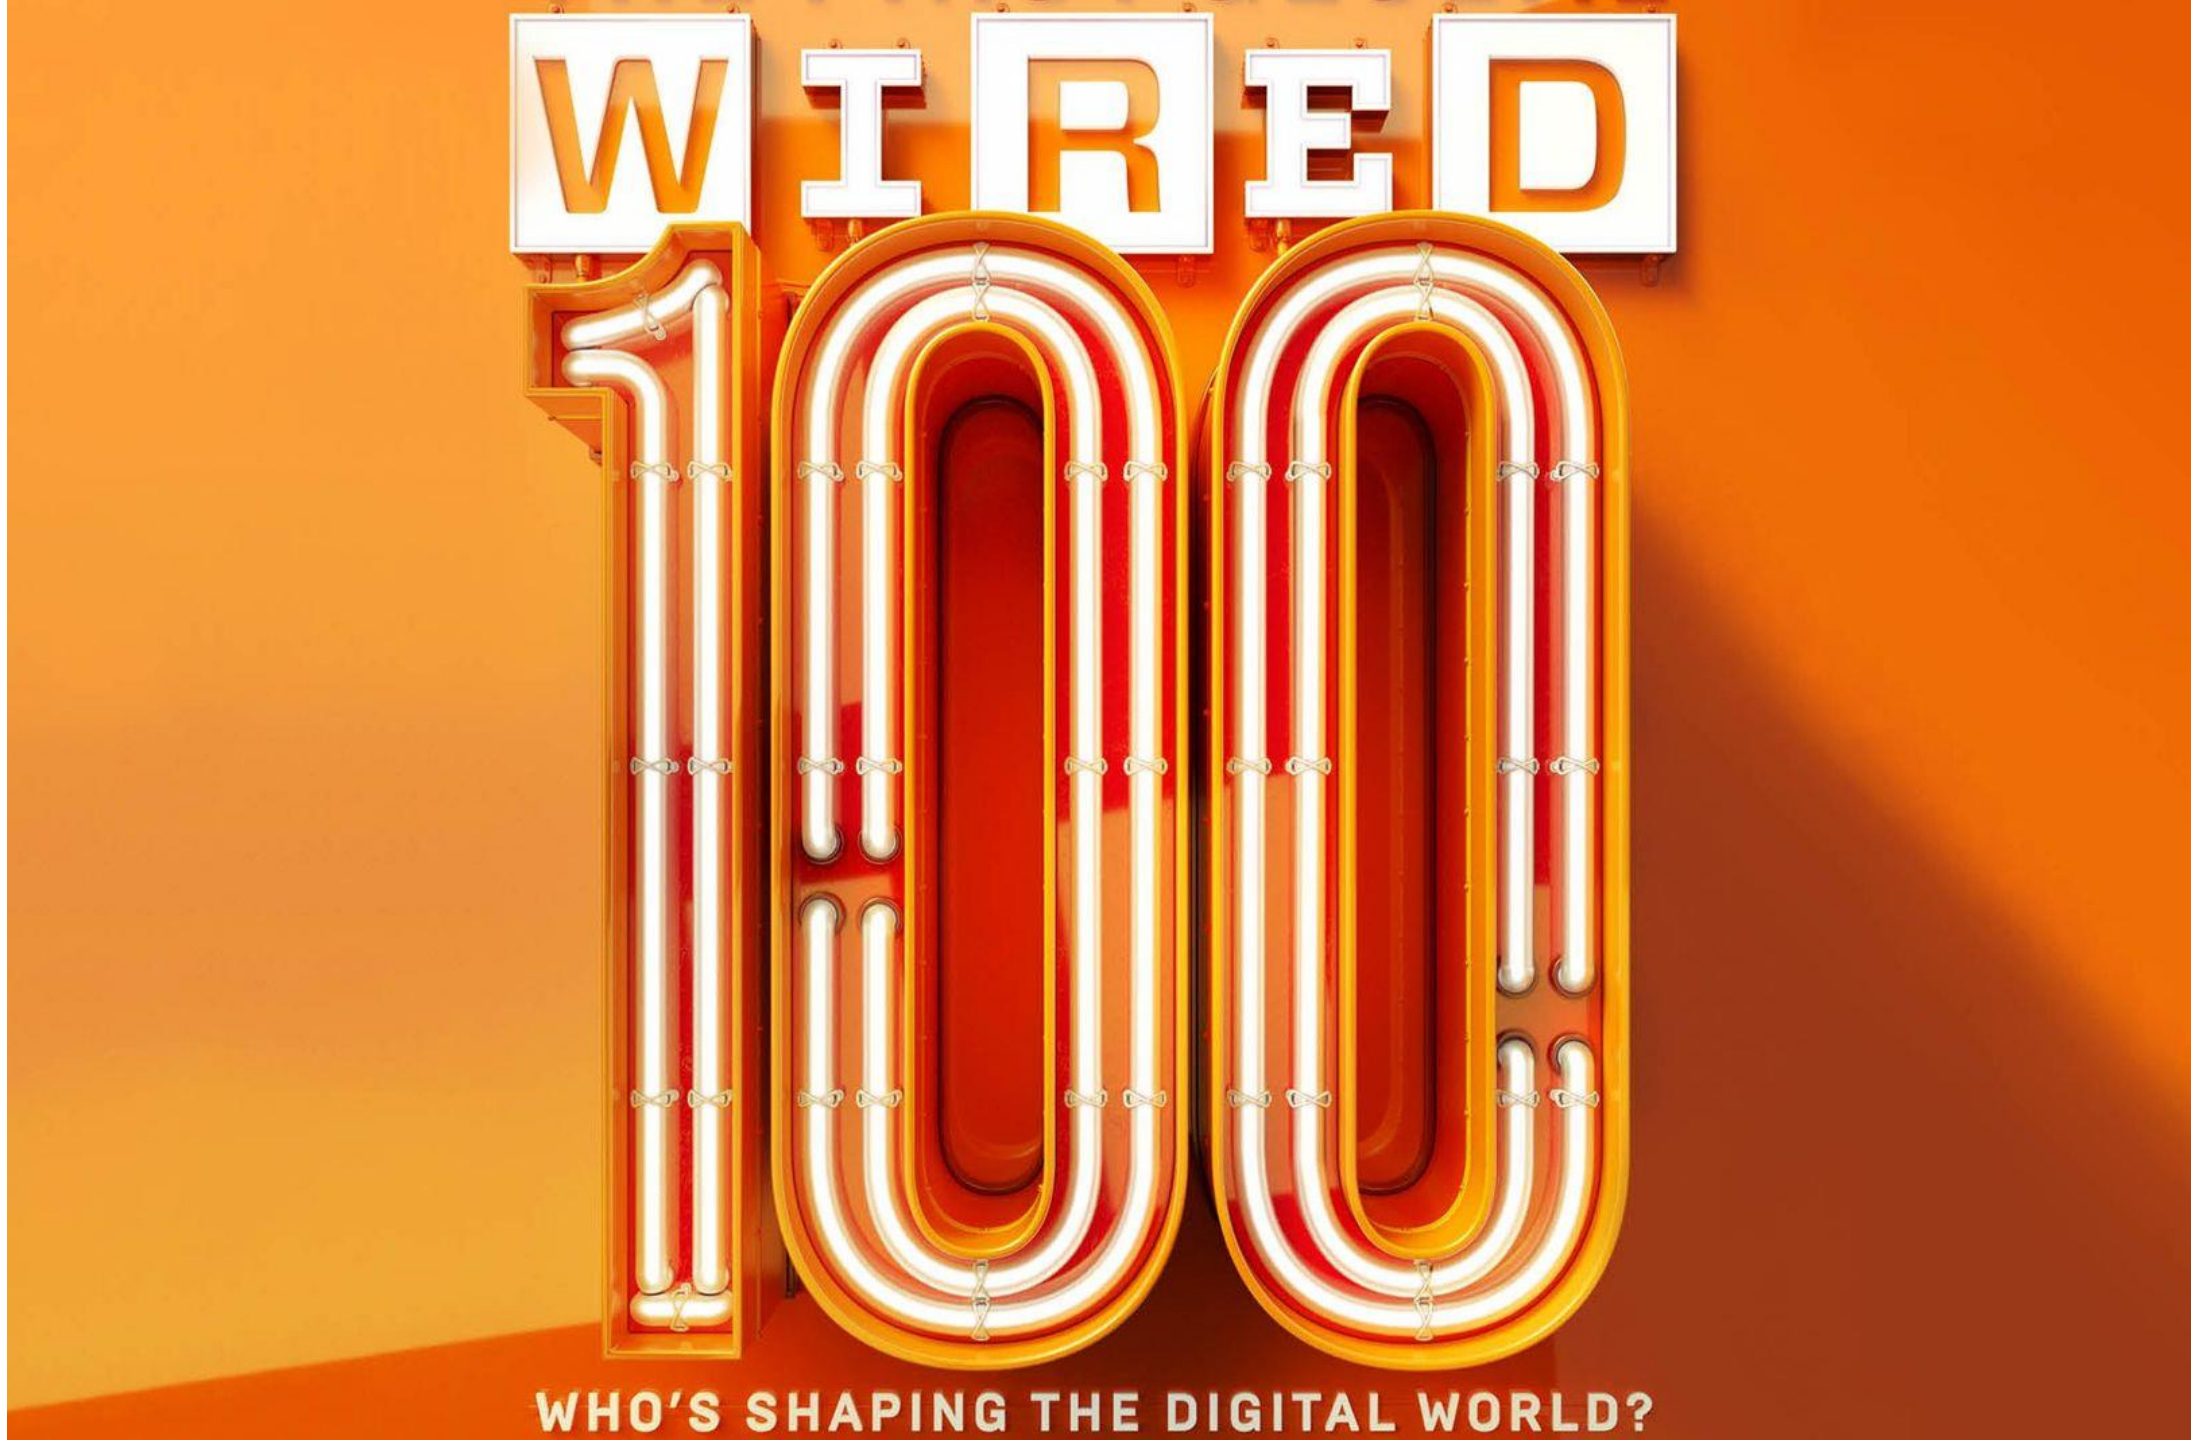
WIRED
100
WHO'S SHAPING THE DIGITAL WORLD?

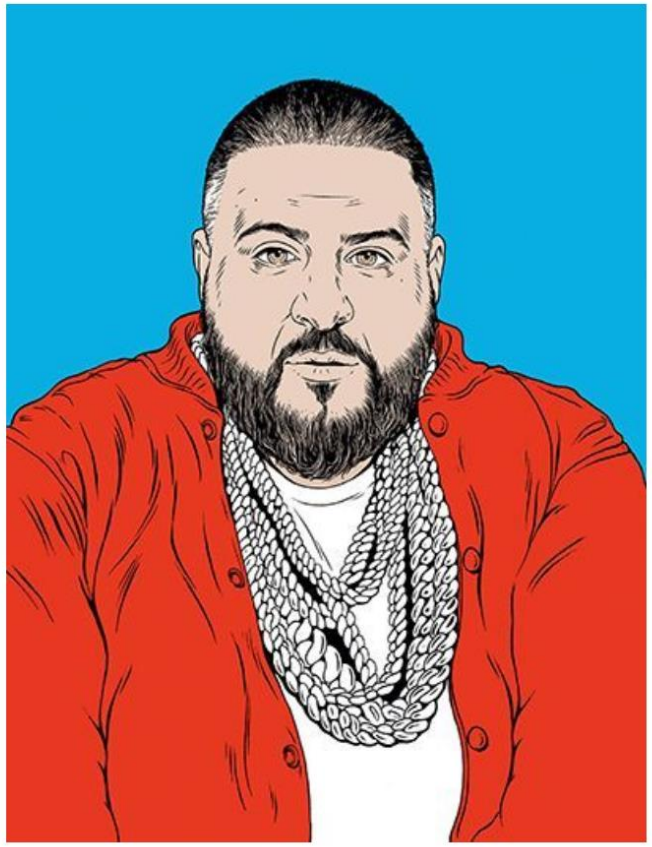

DJ Khaled

Credit **Louise Zergaeng Pomeroy**

73. DJ Khaled

Snapchat icon; DJ and producer

Louisiana-born Khaled Mohamed Khaled, aka DJ Khaled, cut his musical chops in the early 00s as a host for Miami urban music radio WEDR. He proceeded to build a solid if not dazzling career as a mixtape DJ and music producer (he founded his label We The Best Music Group in 2008, and was appointed president of Def Jam South in 2009).

69. Geoffrey Hinton

Psychologist, computer scientist; researcher, Google Toronto

British-born Hinton has been dubbed the "godfather of deep learning". The Cambridge-educated cognitive psychologist and computer scientist started being an ardent believer in the potential of neural networks and deep learning in the 80s, when those technologies enjoyed little support in the wider AI community.

But he soldiered on: in 2004, with support from the Canadian Institute for Advanced Research, he launched a University of Toronto programme in neural computation and adaptive perception, where, with a group of researchers, he carried on investigating how to create computers that could behave like brains.

Hinton's work – in particular his algorithms that train multilayered neural networks – caught the attention of tech giants in Silicon Valley, which realised how deep learning could be applied to voice recognition, predictive search and machine vision.

The spike in interest prompted him to launch a free course on neural networks on e-learning platform Coursera in 2012. Today, 68-year-old Hinton is chair of machine learning at the University of Toronto and moonlights at Google, where he has been using deep learning to help build internet tools since 2013.

63. Yann Lecun

Director of AI research, Facebook, Menlo Park

LeCun is a leading expert in deep learning and heads up what, for Facebook, could be a hugely significant source of revenue: understanding its user's intentions.

62. Richard Branson

Founder, Virgin Group, London

Branson saw his personal fortune grow £550 million when Alaska Air bought Virgin America for $2.6 billion in April. He is pressing on with civilian space travel with Virgin Galactic.

61. Taylor Swift

Entertainer, Los Angeles

Today’s Class

- ~~Covid safety~~
- ~~Who am I?~~
- ~~What is Computer Vision?~~
- Specifics of this course
- Geometry of Image Formation
- Questions

Grading

- 80% programming projects (5 total)
- 20% Quizzes or Problem sets

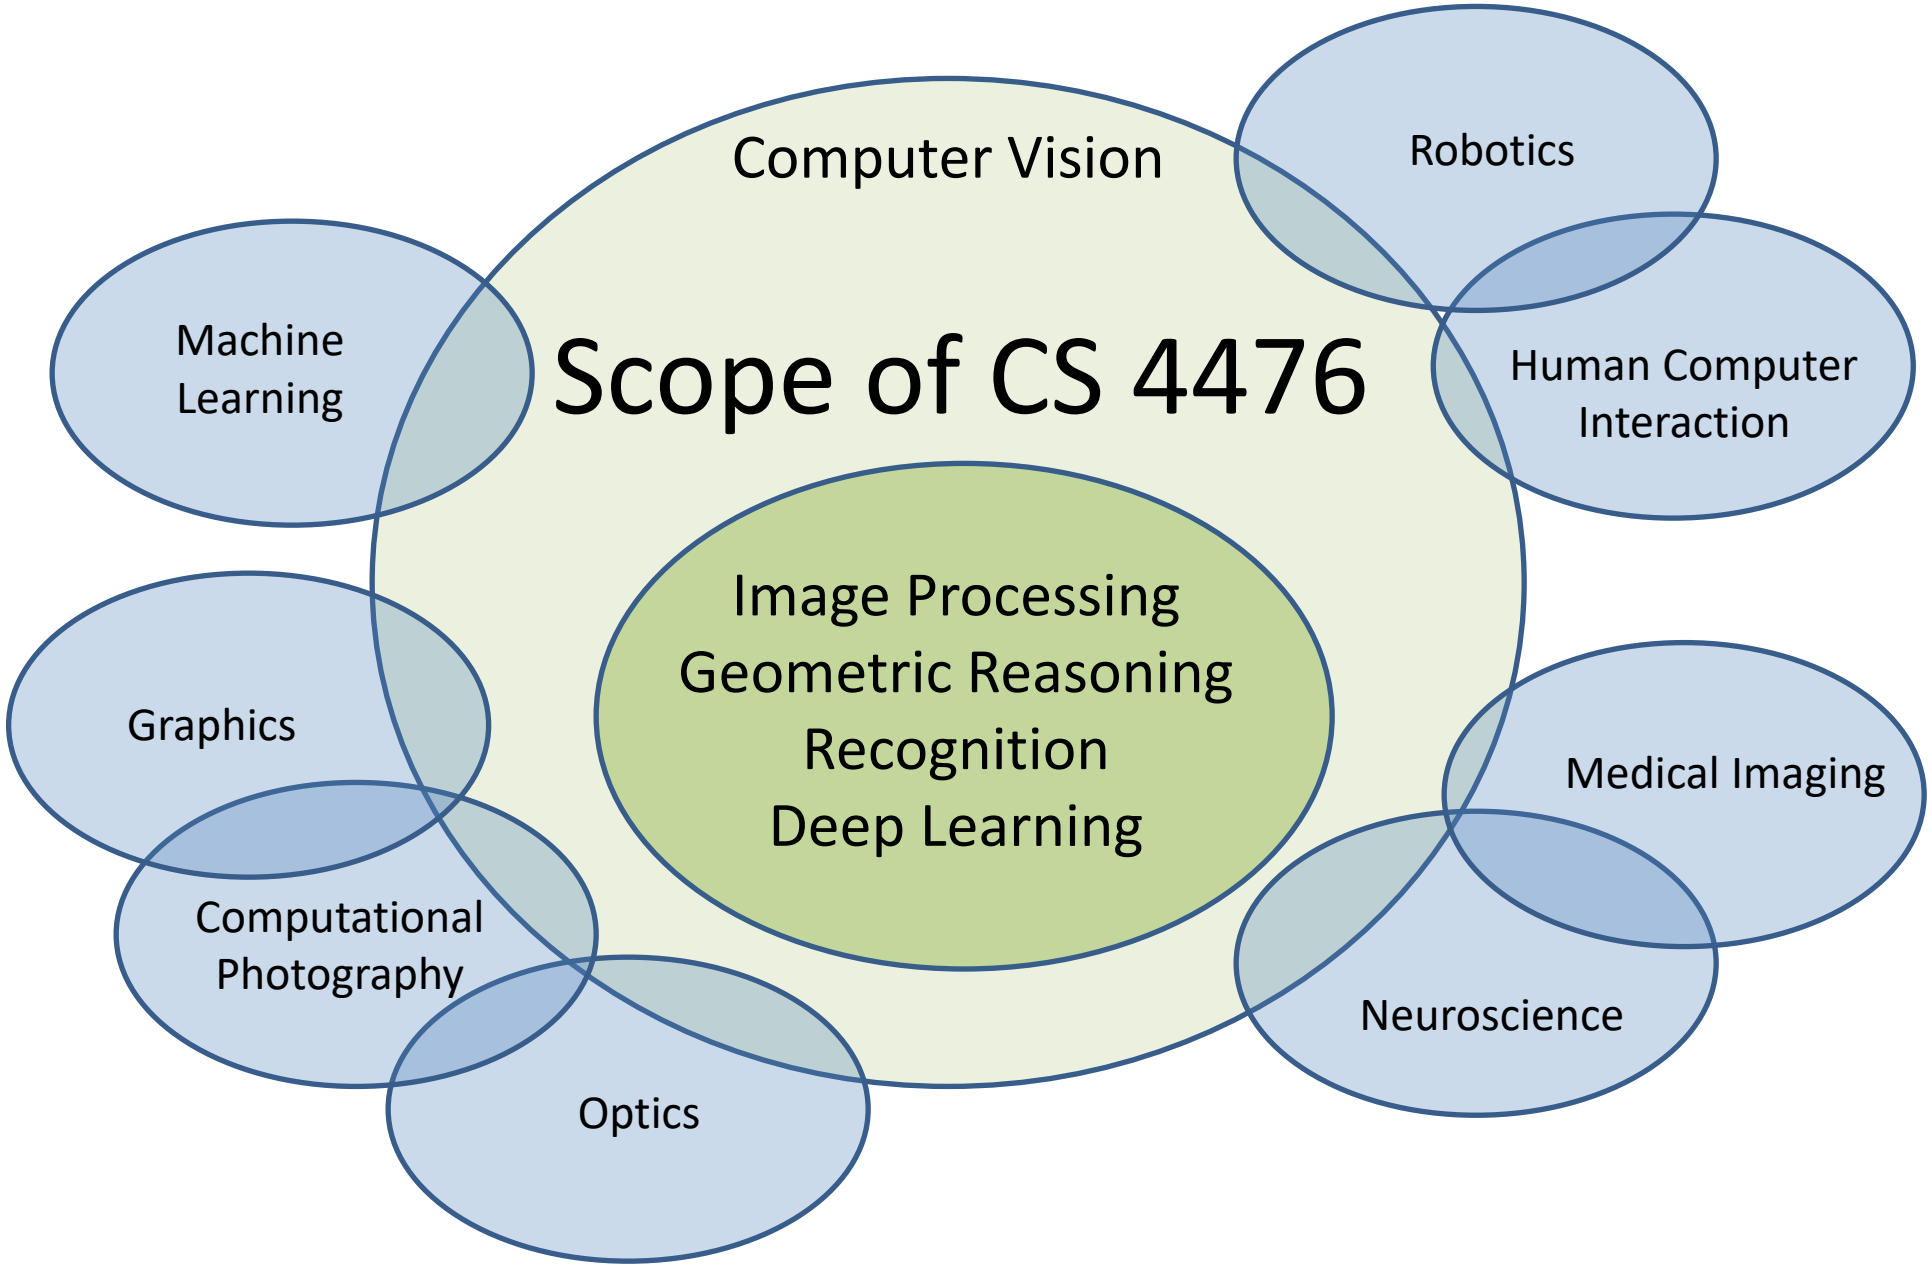
Computer Vision
Scope of CS 4476
Image Processing
Geometric Reasoning
Recognition
Deep Learning
Robotics
Human Computer Interaction
Machine Learning
Graphics
Computational Photography
Optics
Medical Imaging
Neuroscience

Textbook

Computer Vision: Algorithms and Applications, 2nd ed.

http://szeliski.org/Book/

Prerequisites

- **Linear algebra**, basic calculus, and probability
- Experience with image processing will help but is not necessary
- Experience with Python or Python-like languages will help

Projects

- Image Filtering and Hybrid Images
- Local Feature Matching
- Camera Calibration and Fundamental Matrix Estimation with RANSAC
- Image Classification with Deep Learning
- Semantic Segmentation with Deep Learning

Proj1: Image Filtering and Hybrid Images

- Implement image filtering to separate high and low frequencies
- Combine high frequencies and low frequencies from different images to create an image with scale-dependent interpretation

Proj2: Local Feature Matching

- Implement interest point detector, SIFT-like local feature descriptor, and simple matching algorithm.

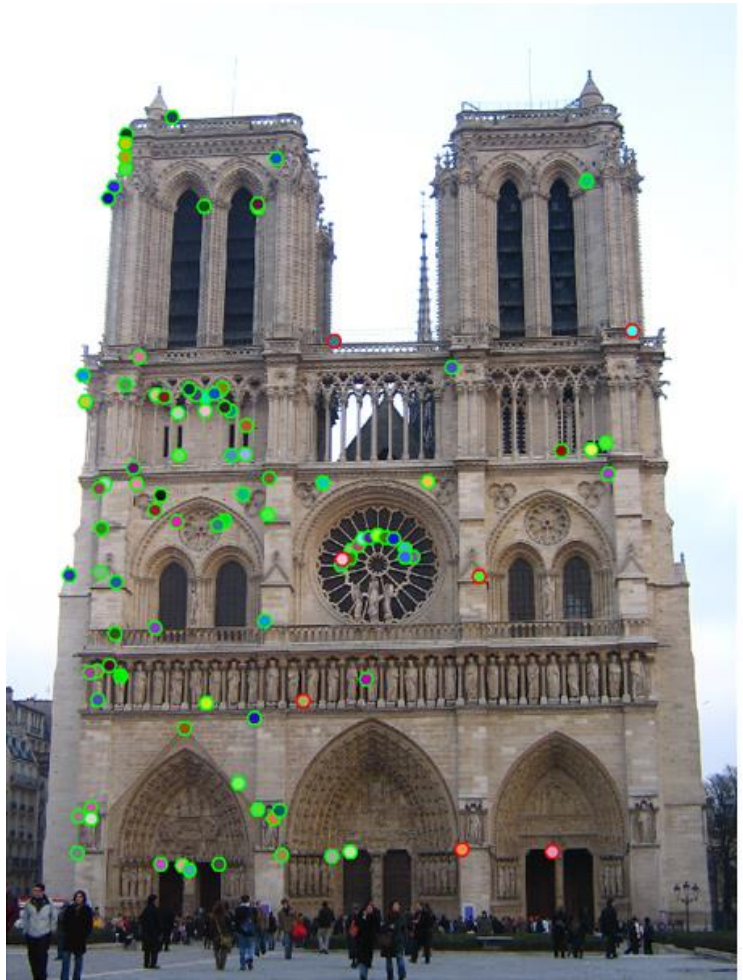

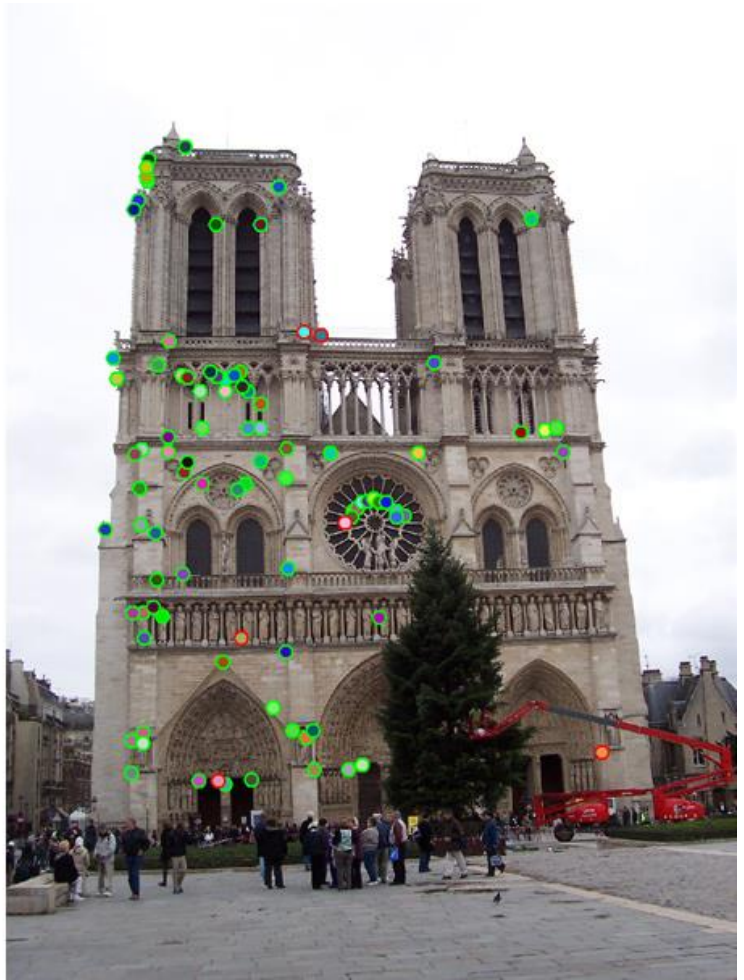

Course Syllabus (tentative)

http://www.cc.gatech.edu/~hays/compvision

Code of Conduct

Your work must be your own. We'll look for cheating. Don't talk at the level of code with other students.

Today's Class

- ~~Covid safety~~
- ~~Who am I?~~
- ~~What is Computer Vision?~~
- ~~Specifics of this course~~
- Geometry of Image Formation
- Questions


WÓDKA
Cuisinart
SNOW CRYSTALS
FOTOGRAFIE UND DAS UNSICHTBARE
DER KLEINE STOWASSER
MATHEMATICS
MAN AND SPACE
DRUGS
STACK-ON


WÓDKA
VODKA
Cuisinart
SNOW CRYSTALS
MATHEMATICS
MAN AND SPACE
DRUGS
STACK-ON

The Geometry of Image Formation

Mapping between image and world coordinates

- Pinhole camera model
- Projective geometry
  - Vanishing points and lines
- Projection matrix

What do you need to make a camera from scratch?

Image formation

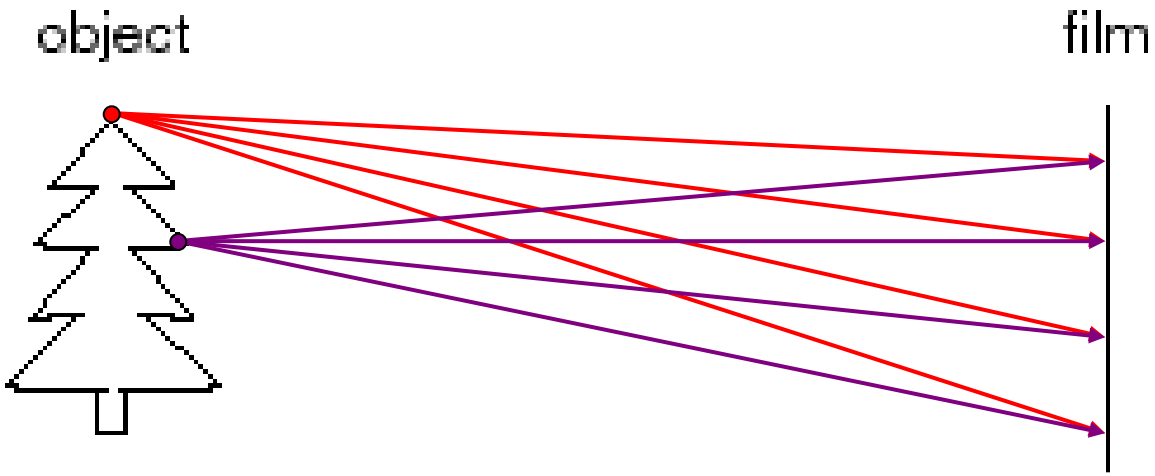

Let’s design a camera

- Idea 1: put a piece of film in front of an object
- Do we get a reasonable image?

Slide source: Seitz

Pinhole camera

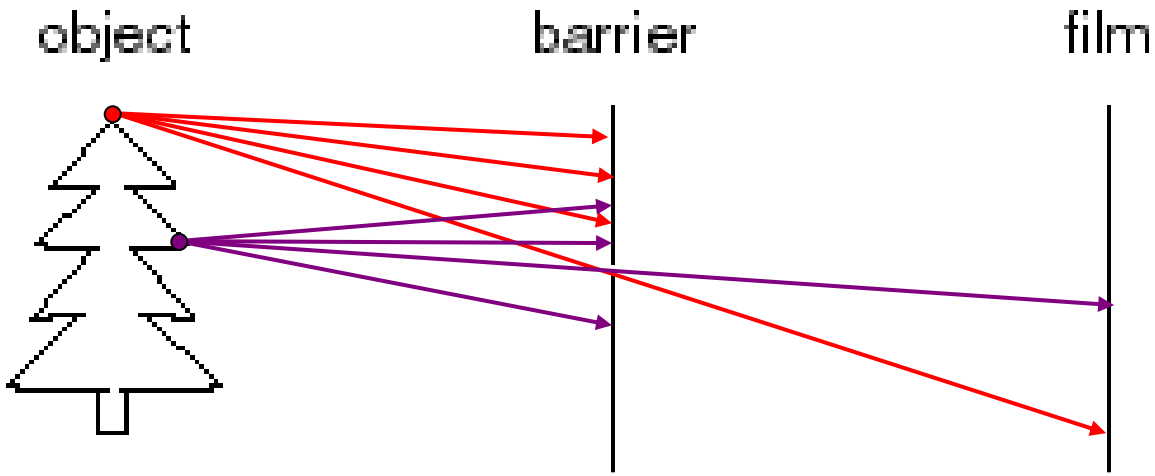

Idea 2: add a barrier to block off most of the rays

- This reduces blurring
- The opening known as the **aperture**

Slide source: Seitz

Pinhole camera

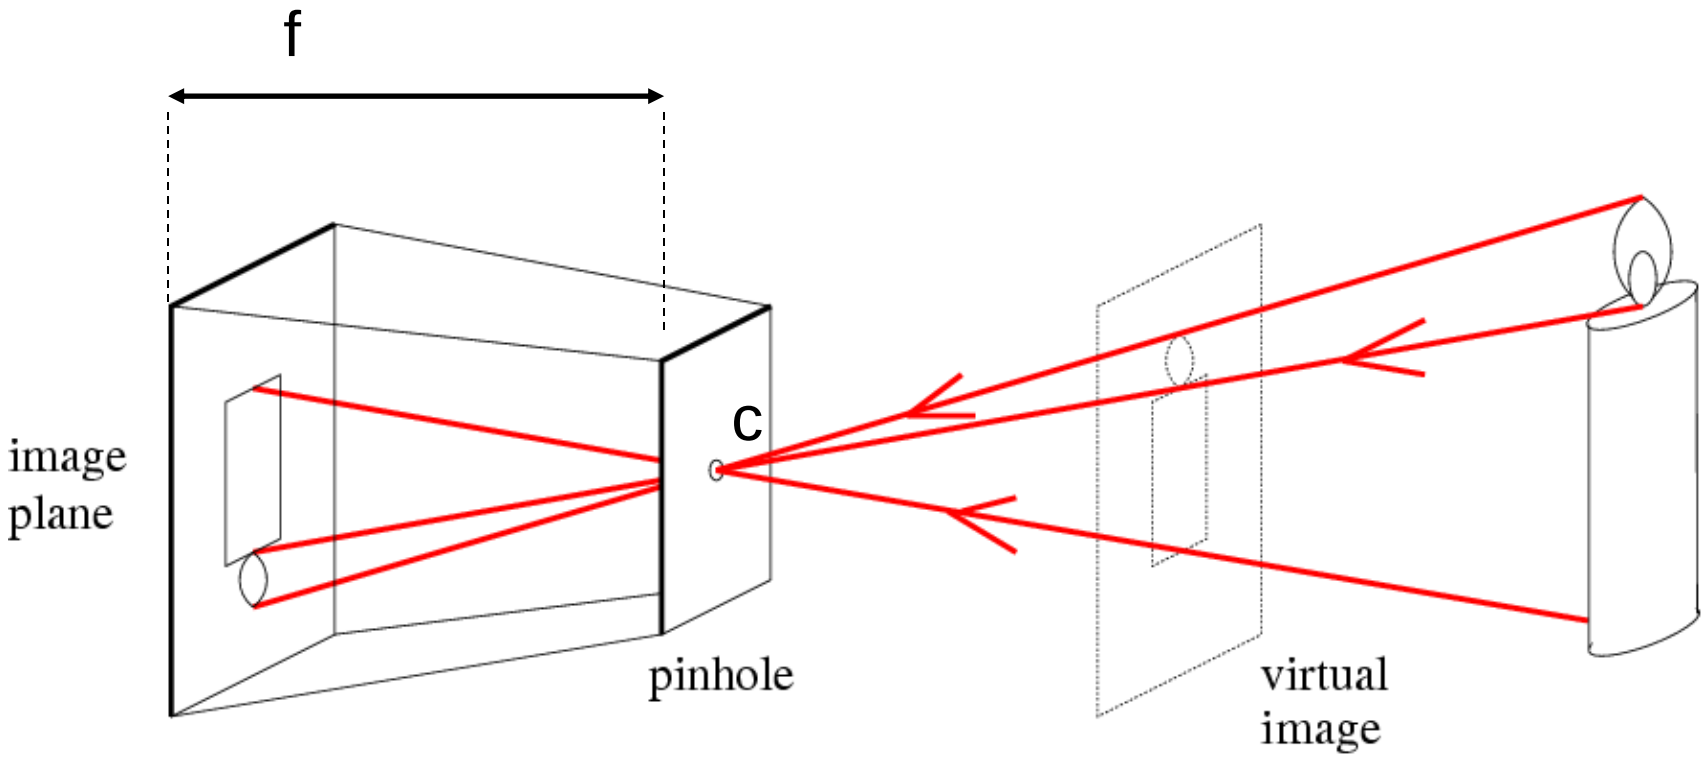

f = focal length
c = center of the camera

Figure from Forsyth

Camera obscura: the pre-camera

- Known during classical period in China and Greece (e.g. Mo-Ti, China, 470BC to 390BC)

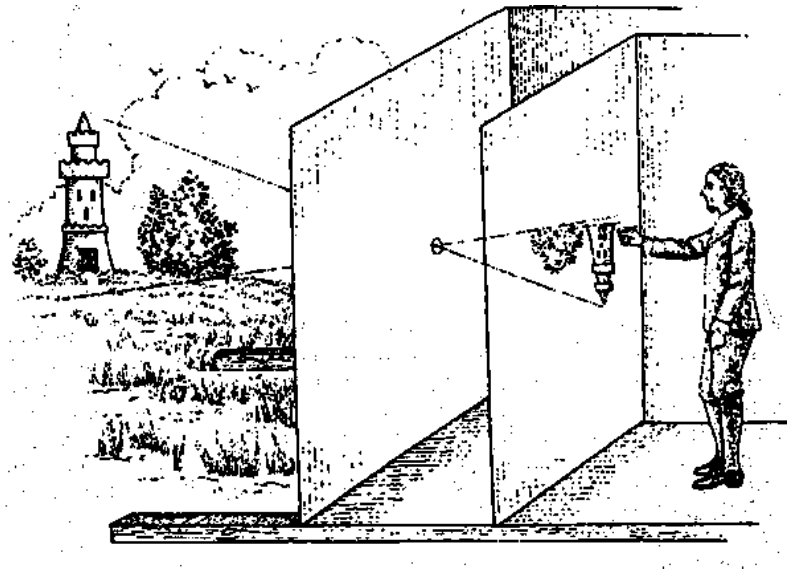

Illustration of Camera Obscura

Freestanding camera obscura at UNC Chapel Hill

Photo by Seth Ilys

Camera Obscura used for Tracing

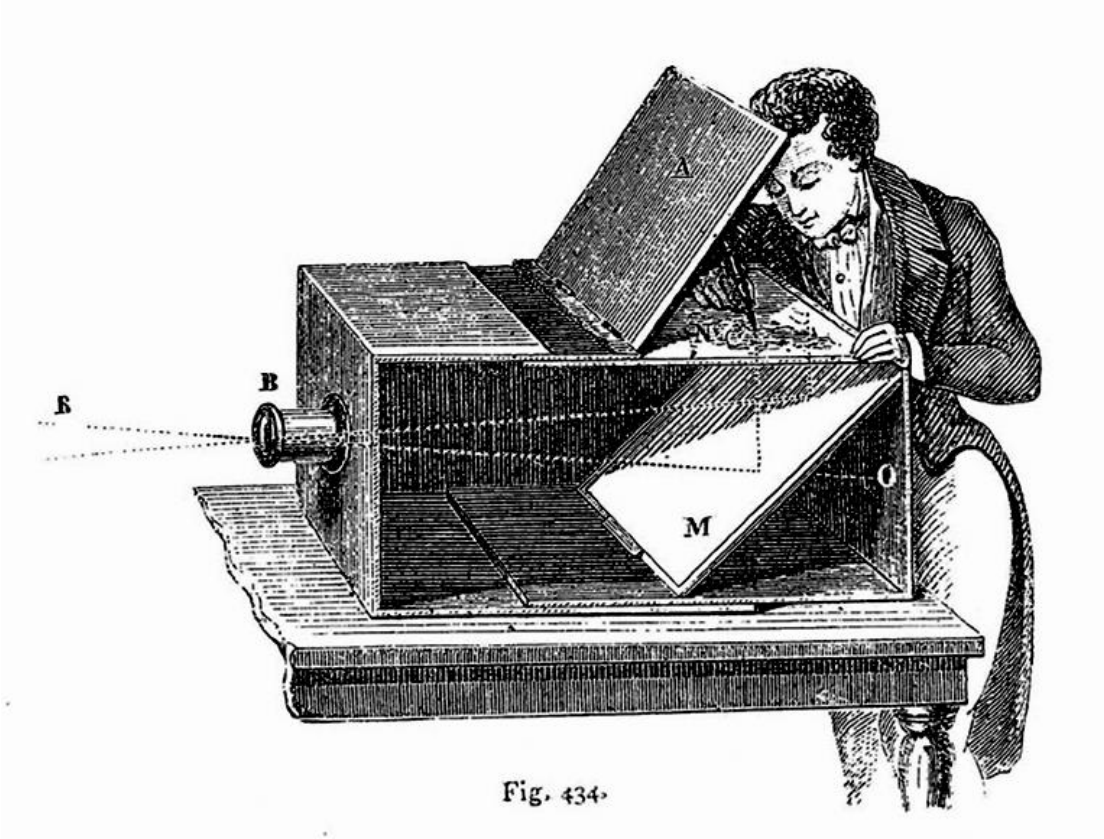

Lens Based Camera Obscura, 1568

Accidental Cameras

Accidental Pinhole and Pinspeck Cameras
Revealing the scene outside the picture.
Antonio Torralba, William T. Freeman

Accidental Cameras

a) Input (occluder present)

b) Reference (occluder absent)

c) Difference image (b-a)

d) Crop upside down

e) True view

First Photograph

Oldest surviving photograph

- Took 8 hours on pewter plate

Joseph Niepce, 1826

Photograph of the first photograph

Stored at UT Austin

Niepce later teamed up with Daguerre, who eventually created Daguerrotypes

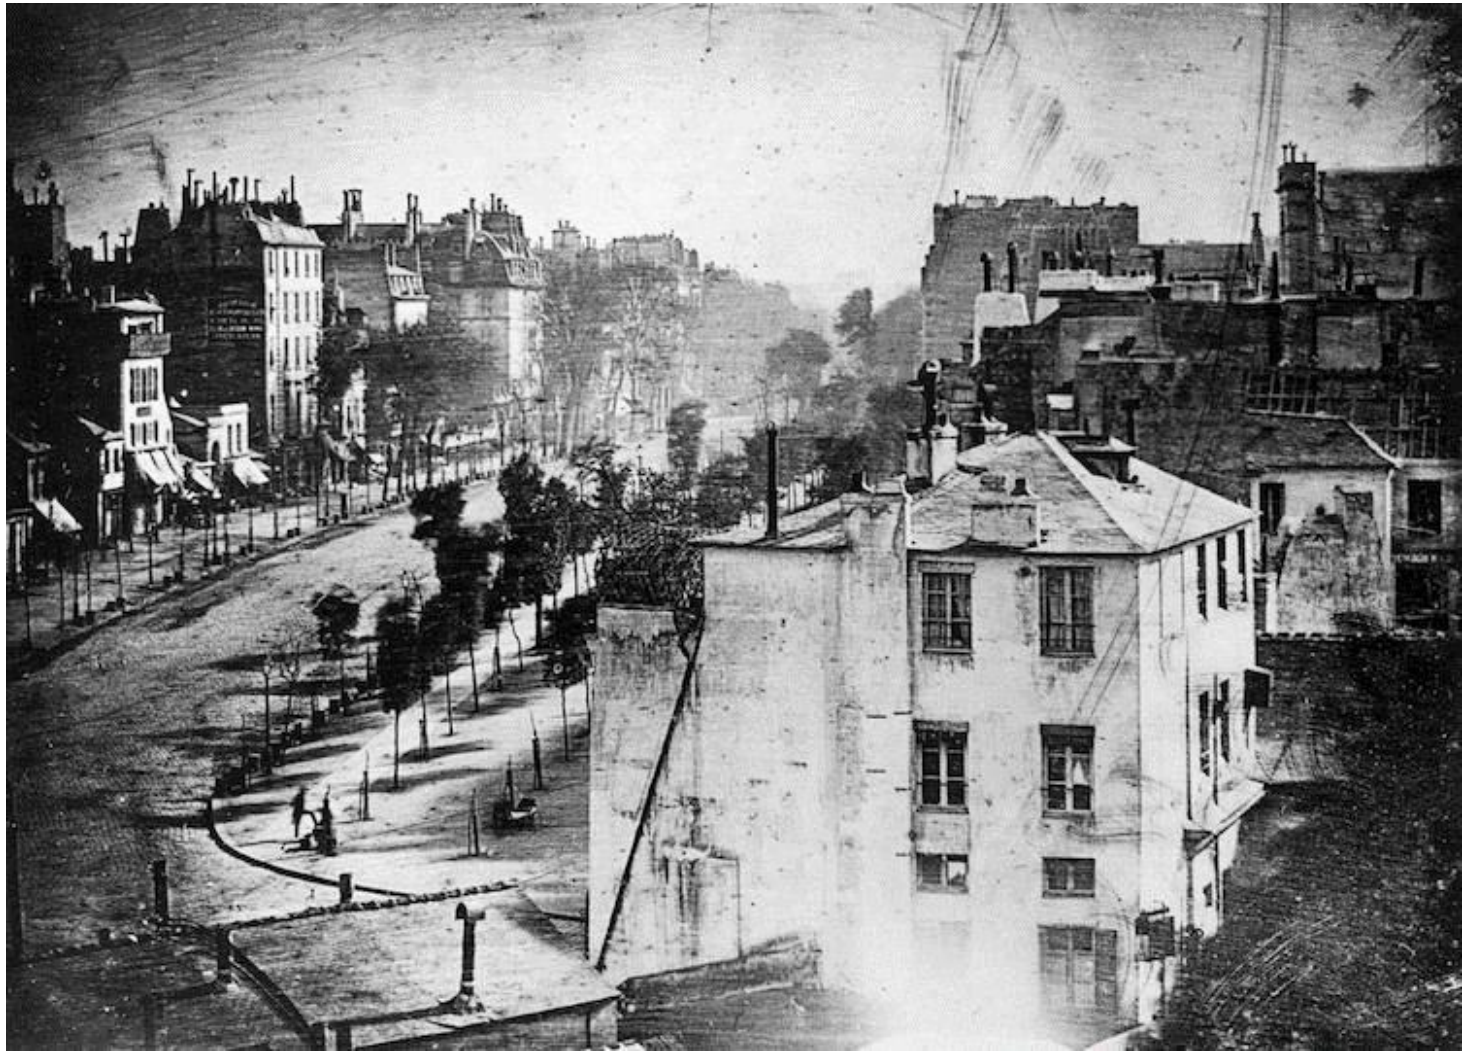

“Louis Daguerre—the inventor of daguerreotype—shot what is not only the world's oldest photograph of Paris, but also the first photo with humans. The 10-minute long exposure was taken in 1839 in Place de la République and it's just possible to make out two blurry figures in the left-hand corner.”

Source

Camera and World Geometry

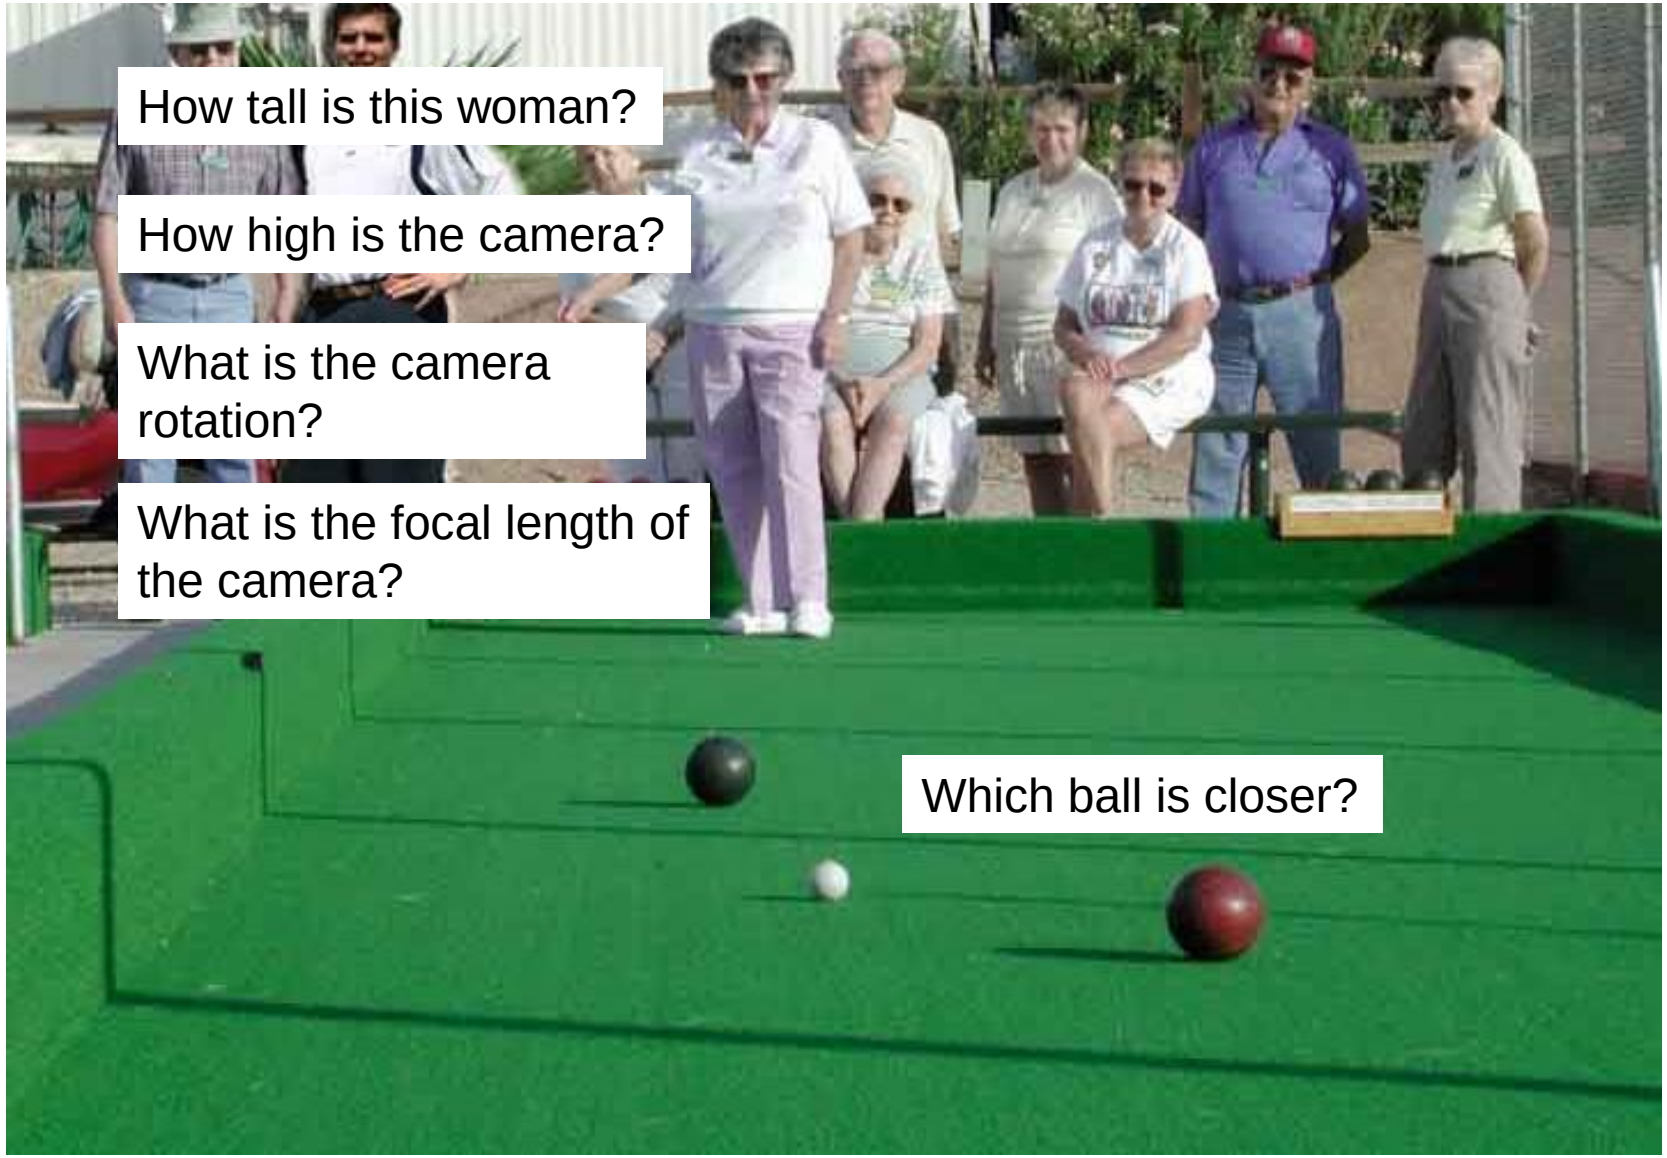

Dimensionality Reduction Machine (3D to 2D)

3D world

2D image

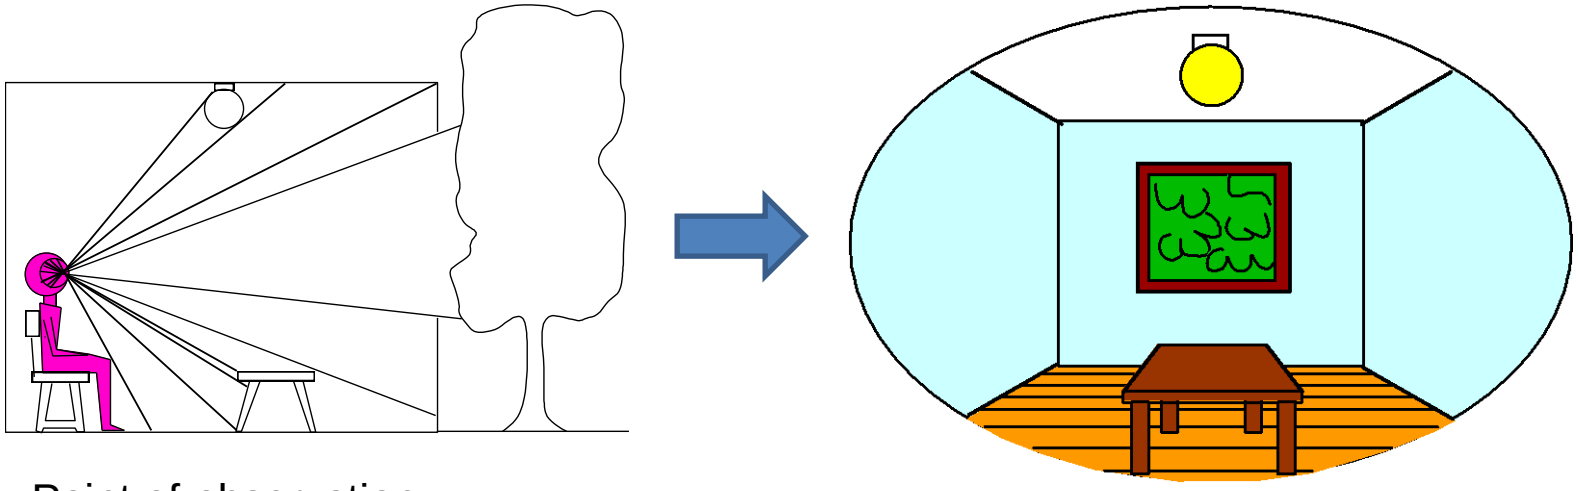

Point of observation

Slide source: Seitz

Projection can be tricky…

Slide source: Seitz

Projection can be tricky…

Projective Geometry

What is lost?

- Length

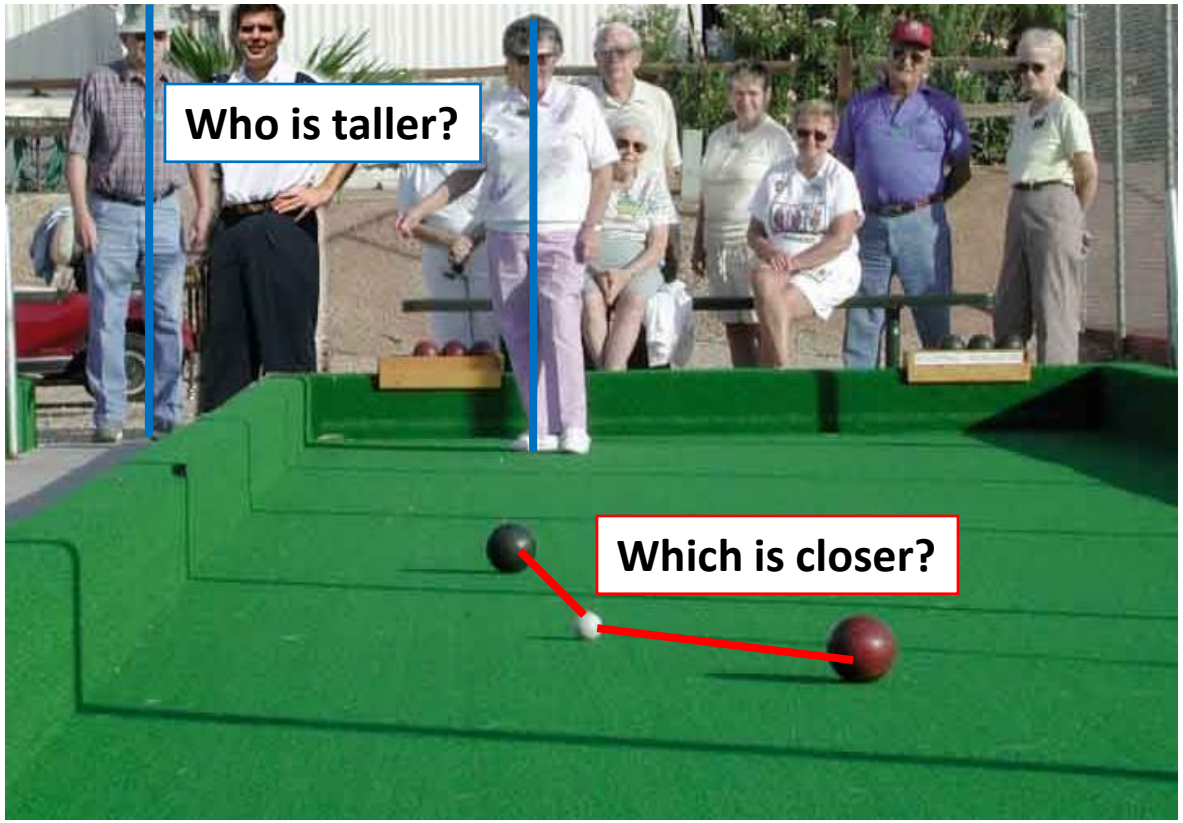

Length and area are not preserved

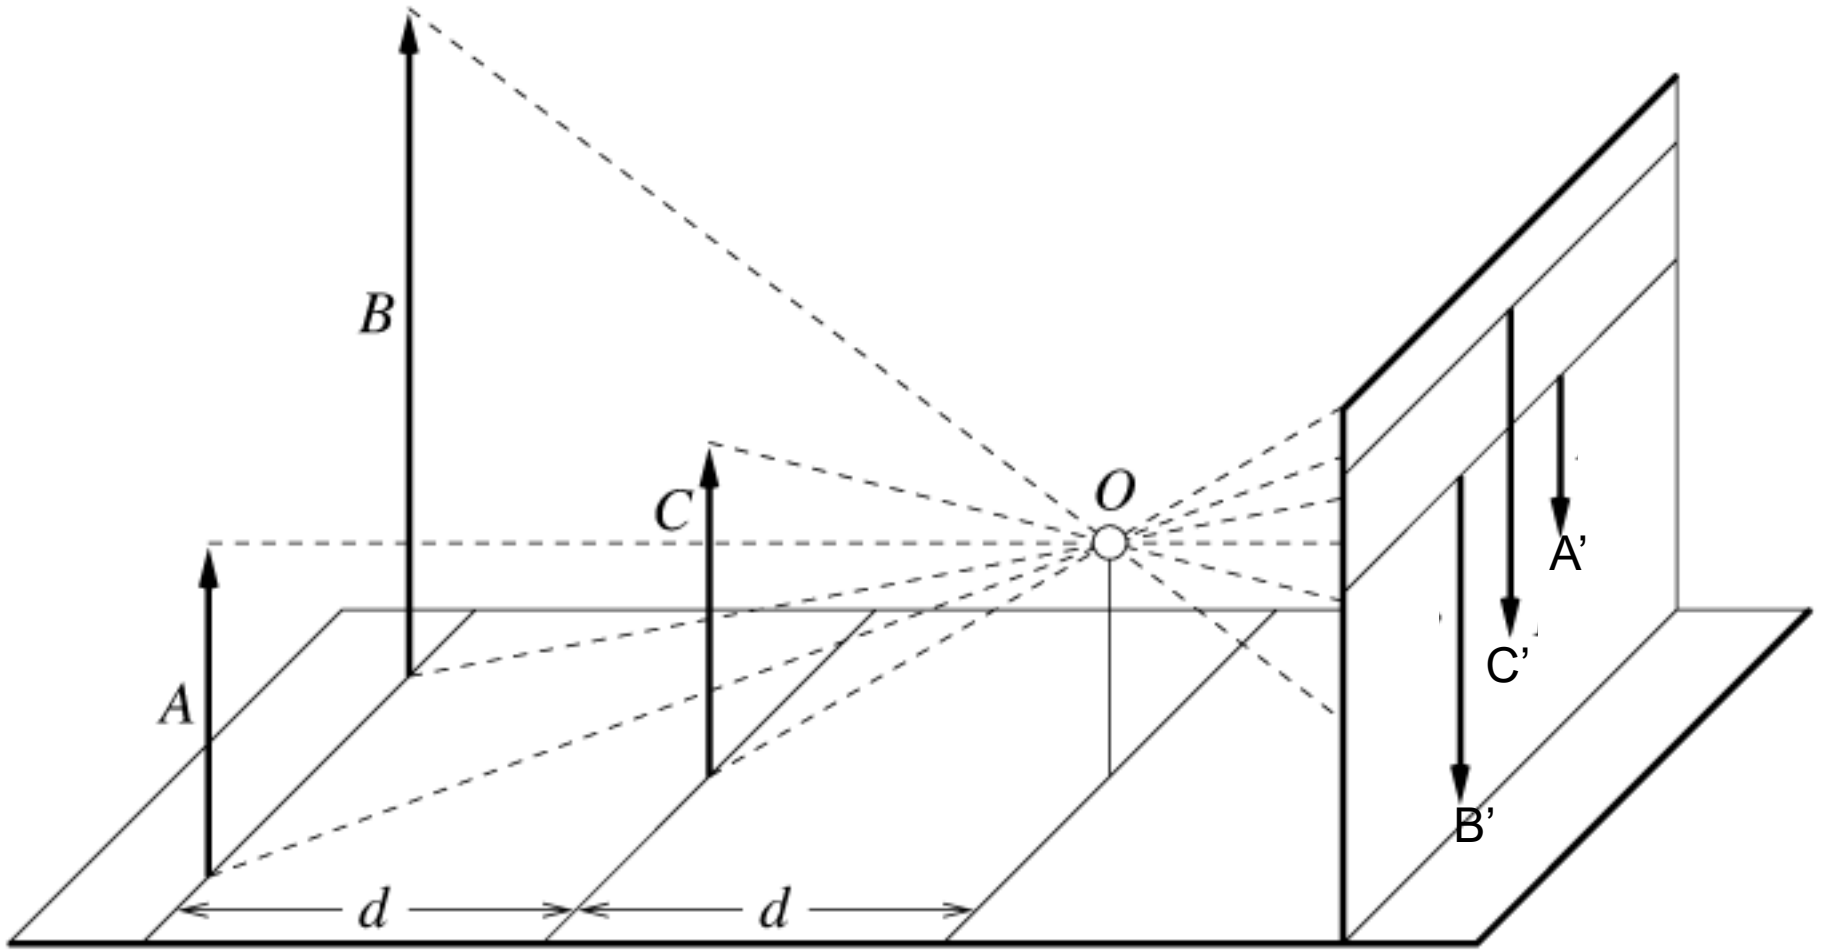

Figure by David Forsyth

Projective Geometry

What is lost?

- Length
- Angles

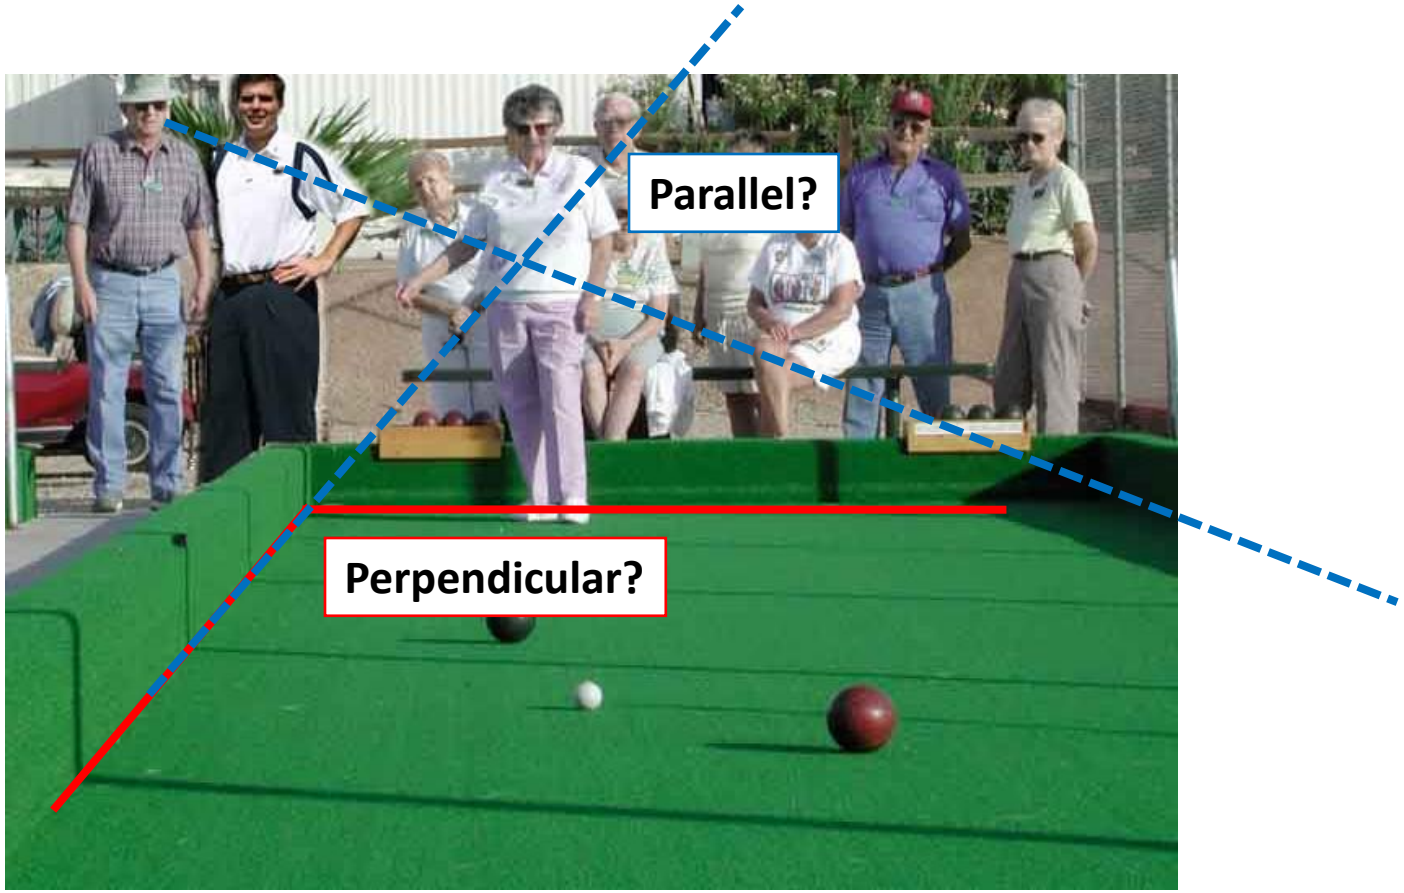

Projective Geometry

What is preserved?

- Straight lines are still straight

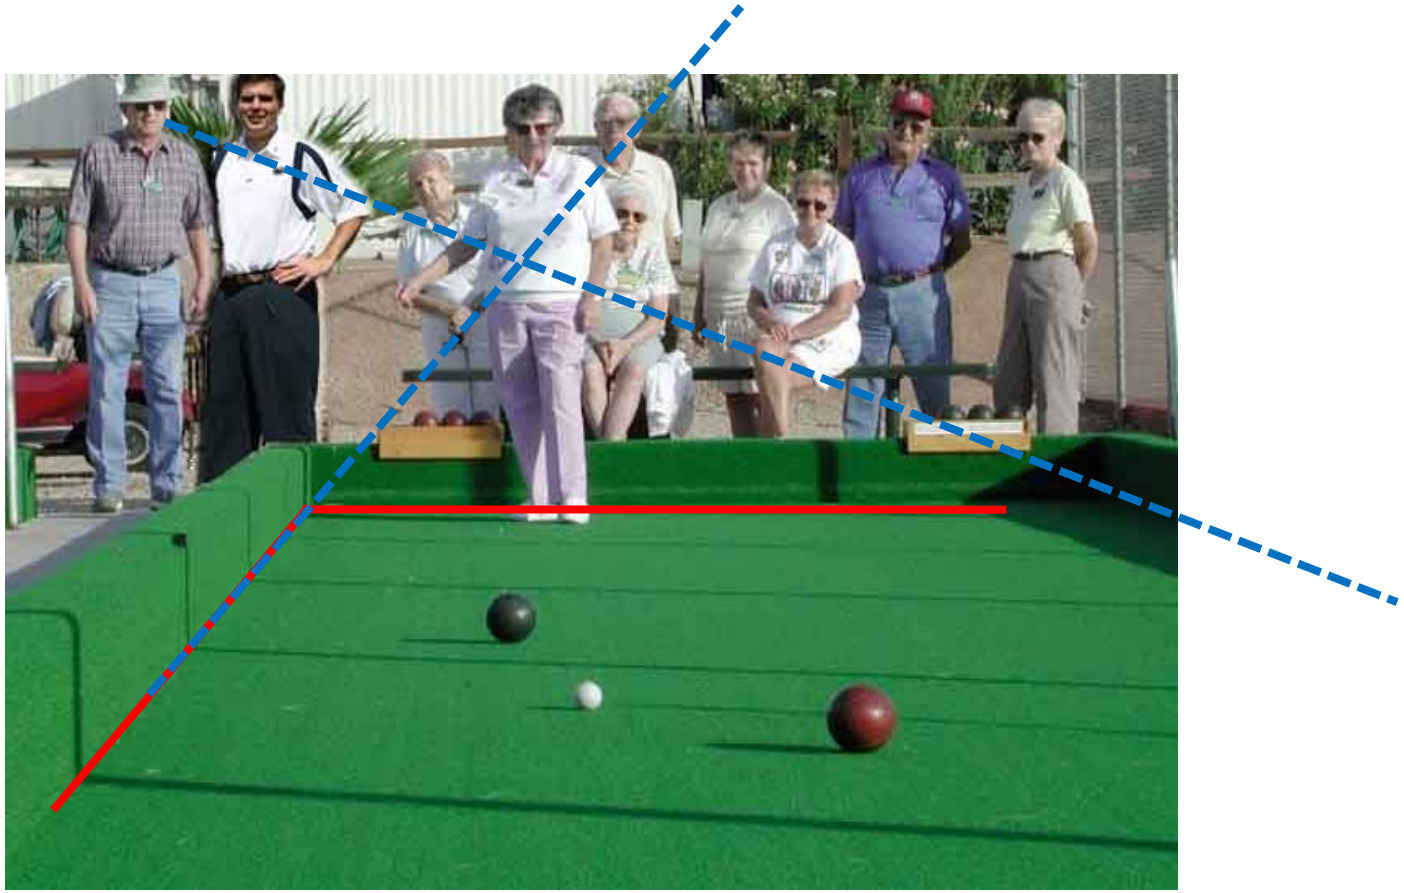

Vanishing points and lines

Parallel lines in the world intersect in the image at a “vanishing point”

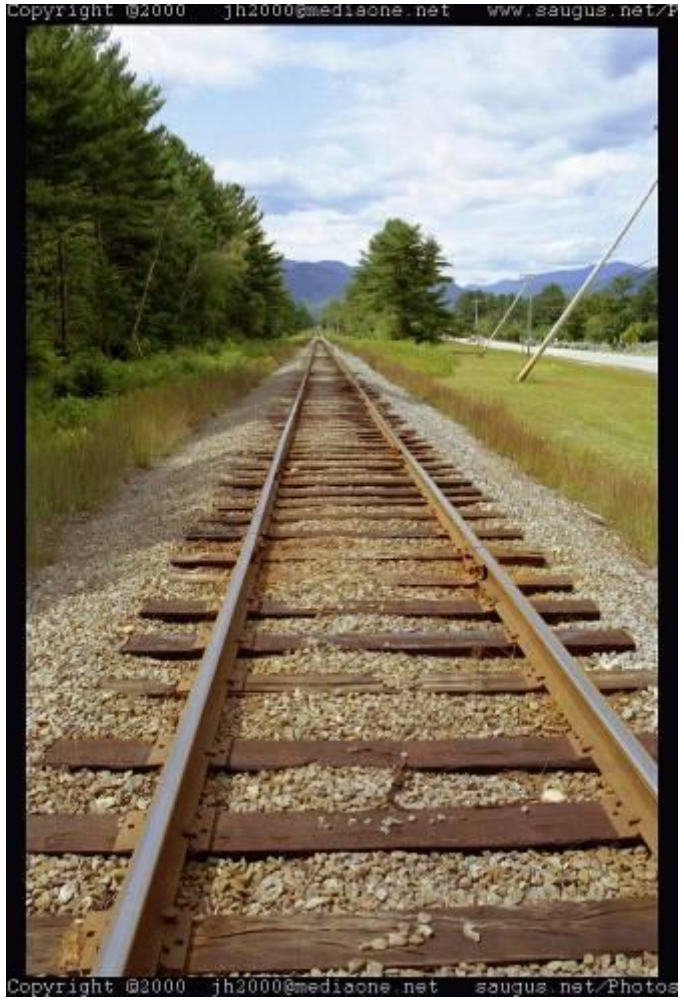

Vanishing points and lines

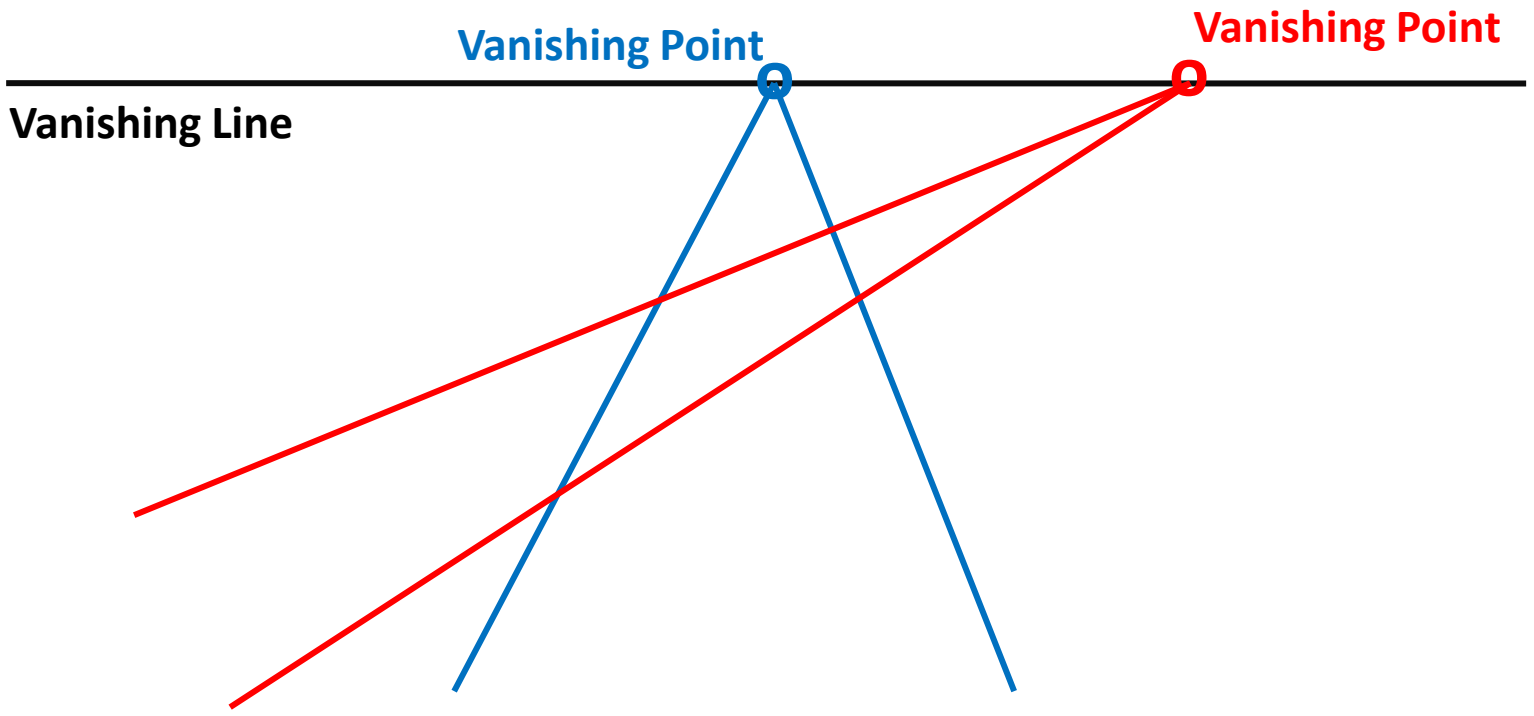

Vanishing points and lines

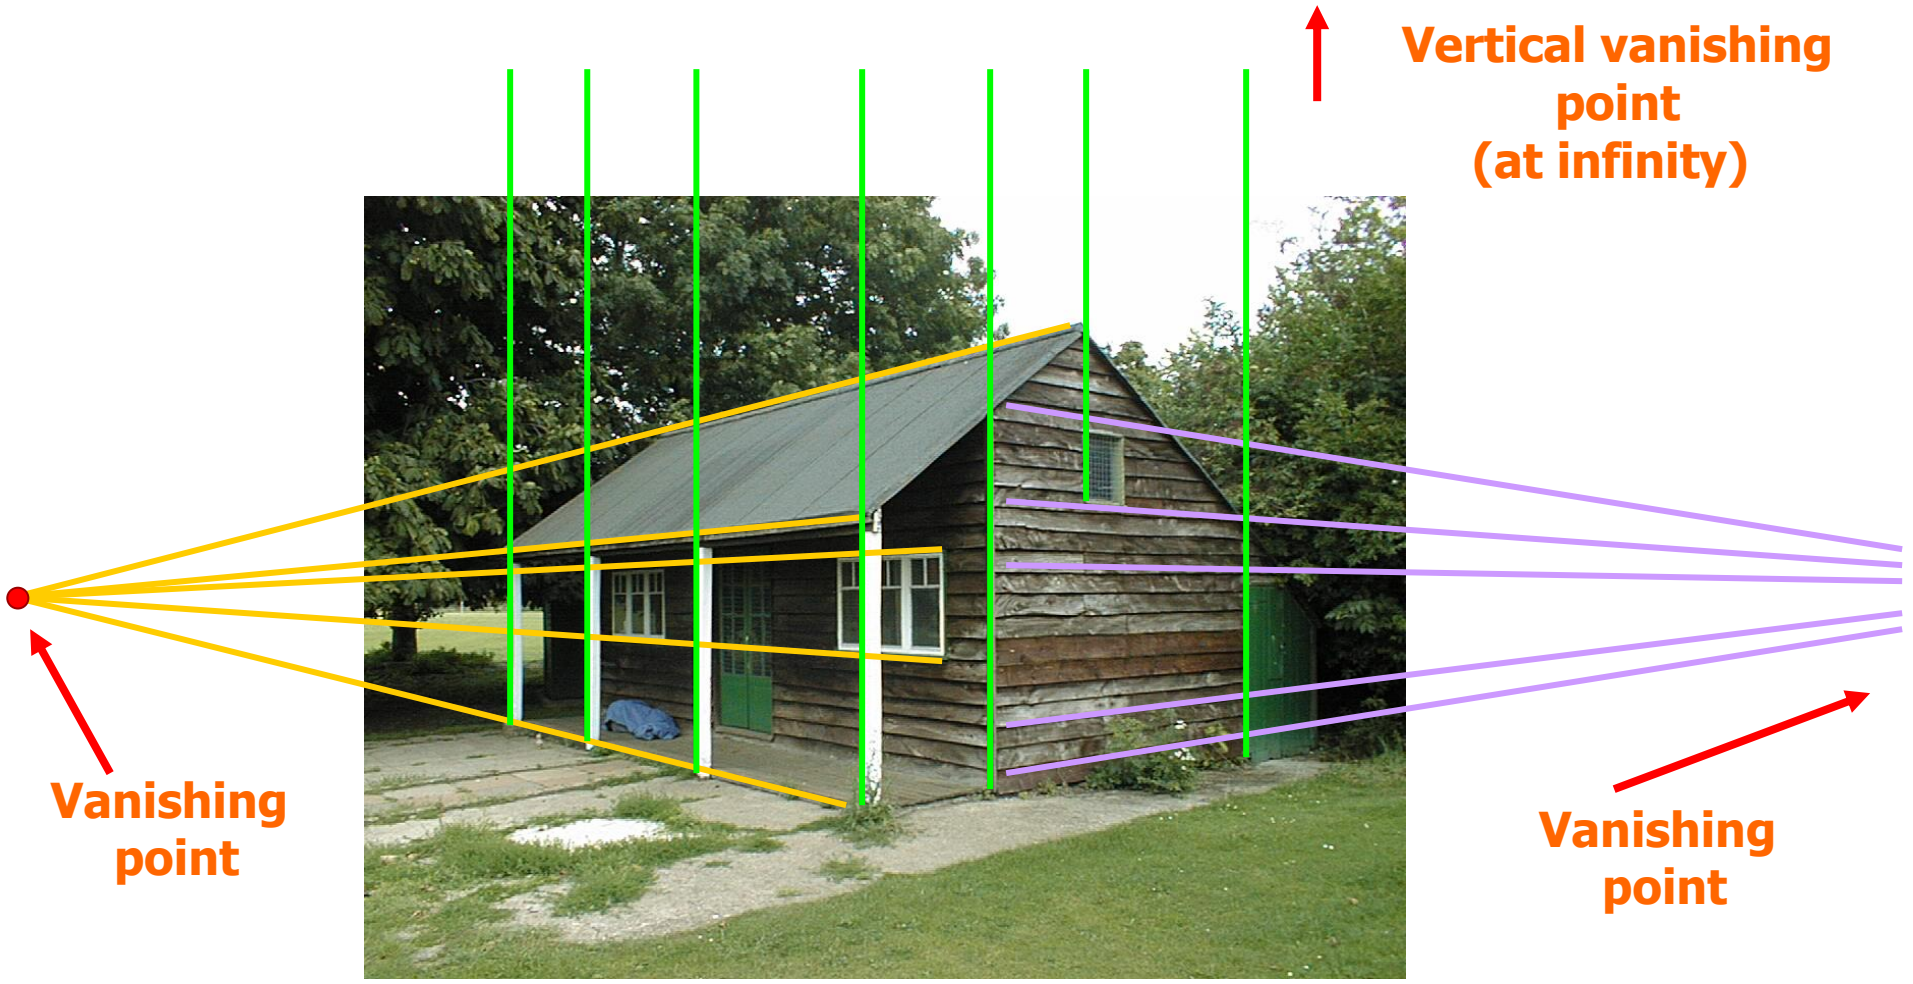

Slide from Efros, Photo from Criminisi

- Project 1 will be out soon
- Read Szeliski 2.1, especially 2.1.4
- Image projection
- Filtering

Chapter 2

Image formation

Projection: world coordinates→image coordinates

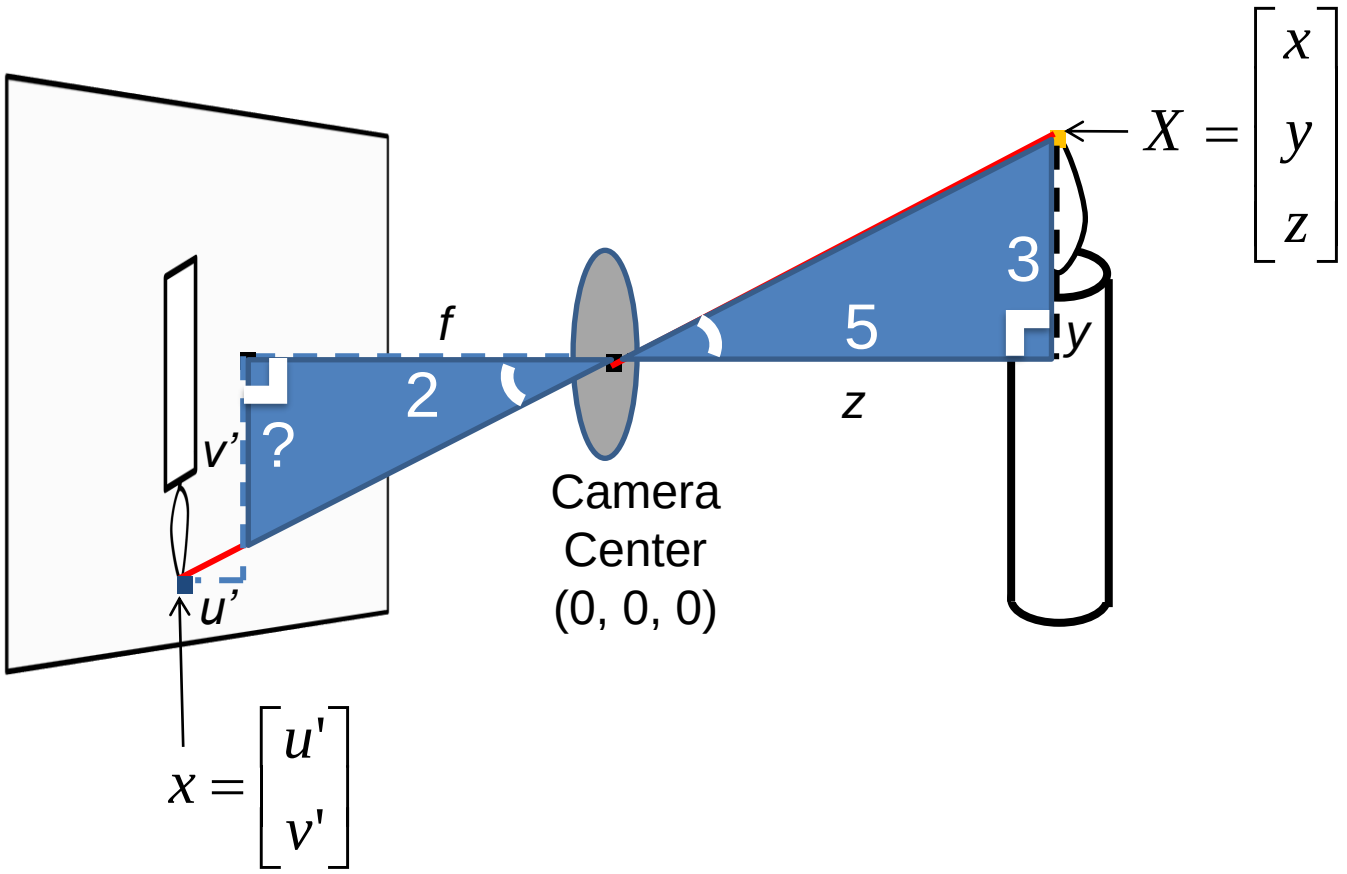

If X = 2, Y = 3,
Z = 5, and f = 2
What are U and V?

$$\frac{v'}{-f} = \frac{y}{z}$$

$$u' = -x * \frac{f}{z} \qquad u' = -2 * \frac{2}{5}$$

$$v' = -y * \frac{f}{z} \qquad v' = -3 * \frac{2}{5}$$