CS 180/280A: Intro to Computer Vision and Computational Photography

Instructors: Angjoo Kanazawa
Alexei Efros

GSIs: Konpat Preechakul
Justin Kerr
Brent Yi
Chung Min Kim

Tutors: Jorge Diaz Chao
Jingfeng Yang
Natalie Wei
Jameson Crate

UC Berkeley, Fall 2025

Today

Introductions

Why this Course?

Administrative stuff

- Project #0 out!

Brief History of Visual Data

Course Staff

Prof. Angjoo Kanazawa

Prof. Alexei Efros

Course Staff

GSIs

Justin Kerr

Konpat Preechakul

Brent Yi

Chung Min Kim

Tutors

Jameson Crate

Jorge Diaz Chao

Jingfeng Yang

Natalie Wei

3D/4D perception of people

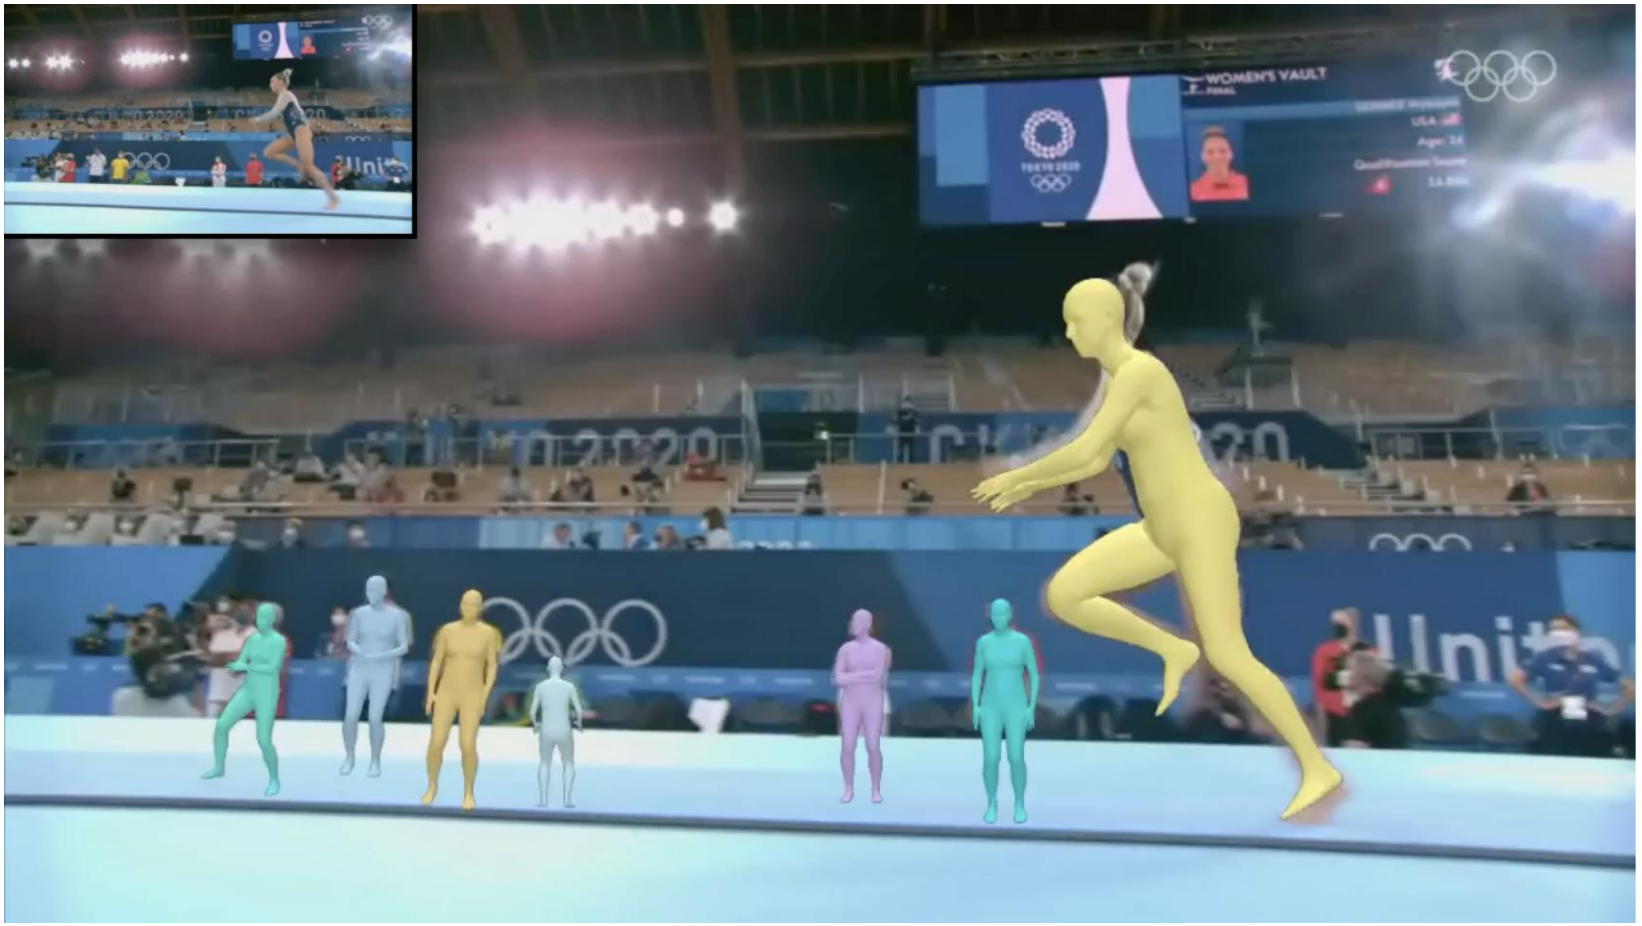

Goel, Pavlakos, Rajasegaran, Kanazawa*, Malik*, ICCV 2023

People + Scenes

HSfM: Reconstructing People, Places, and Cameras

Output: 3D point maps, camera poses, and 3D human meshes

Müller and Choi et al. CVPR 2025

Reconstructing immersive 3D world

From this video..

4D Ayuna

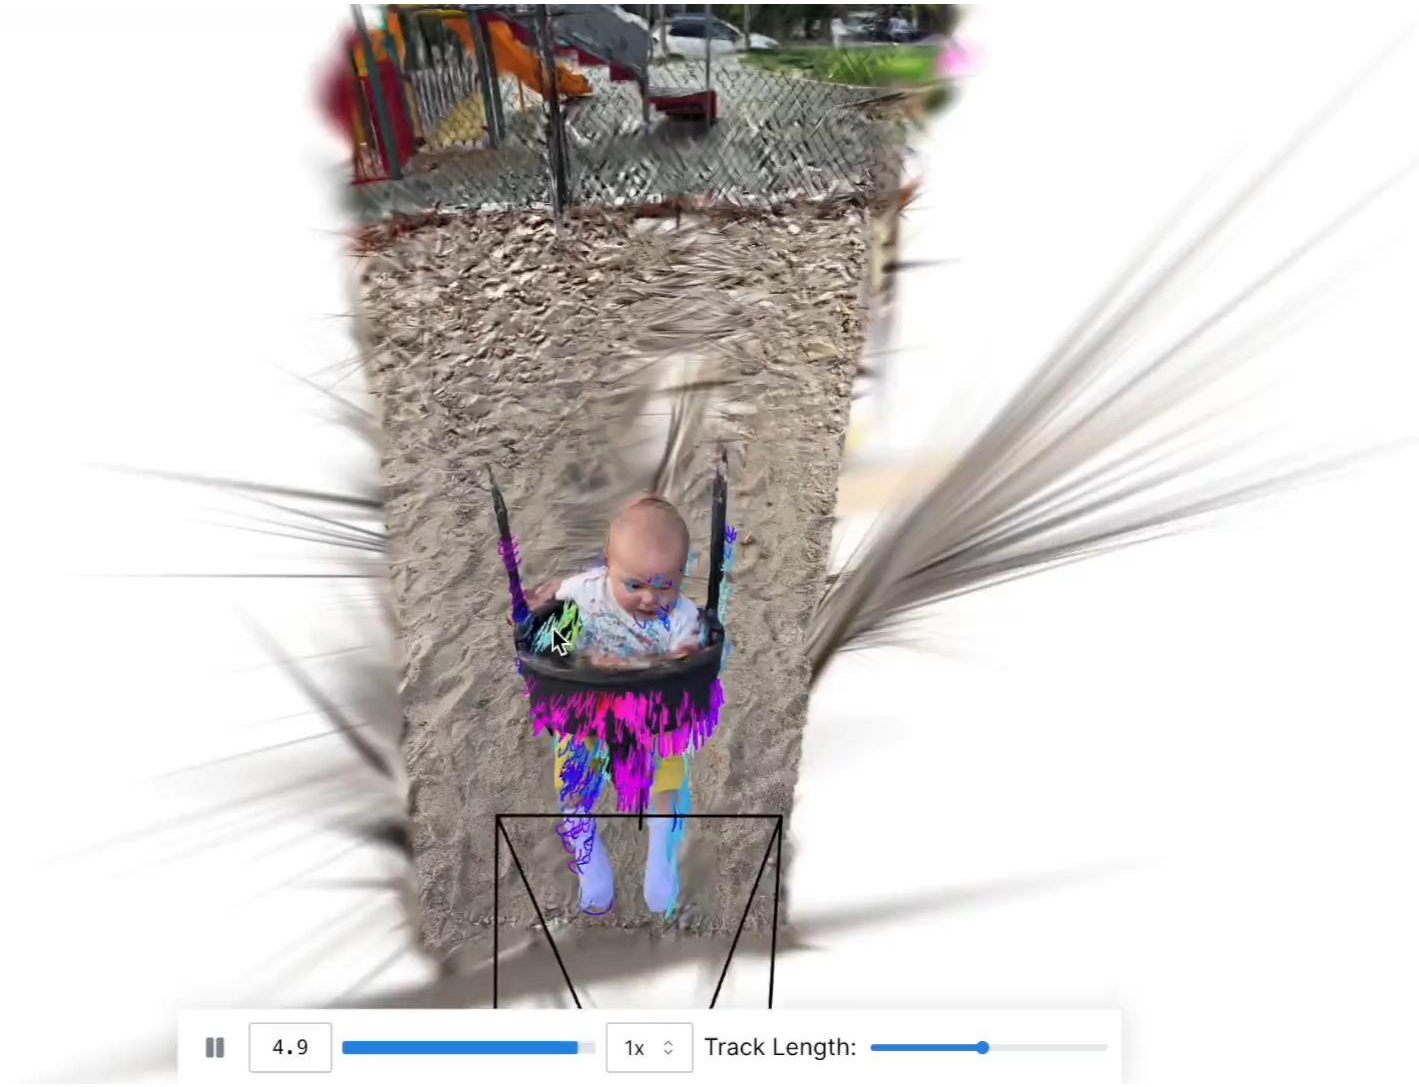

Angjoo’s research

Prof. Efros' research

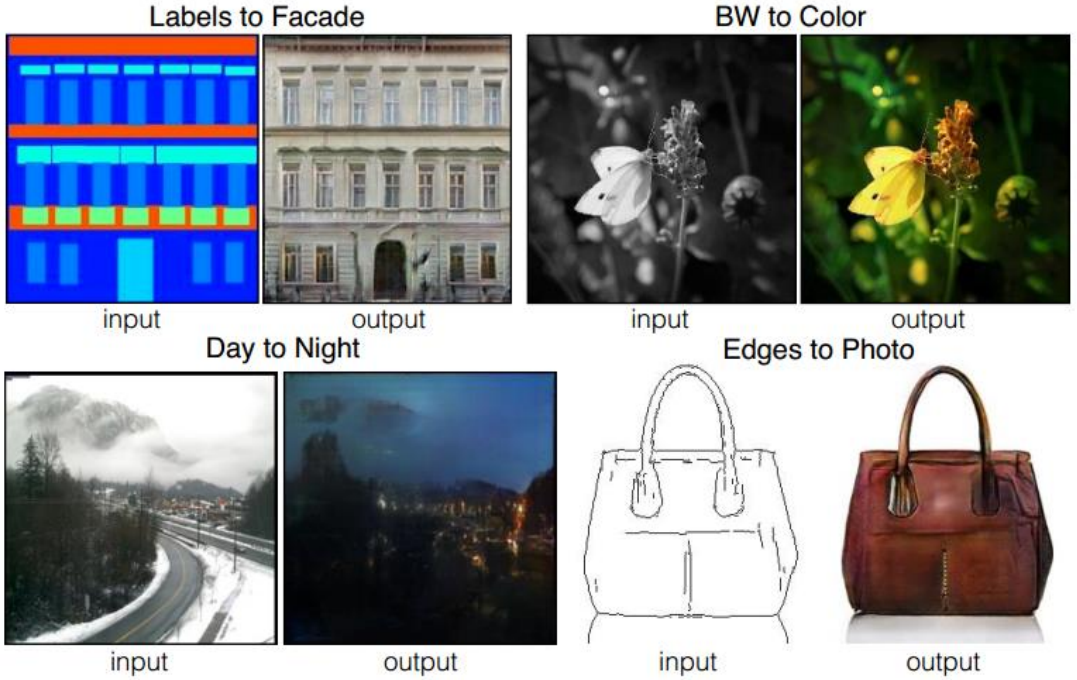

Pretty Pictures: Pix2Pix, cycleGAN, *instructPix2Pix*, etc.

Emergence via self-supervision

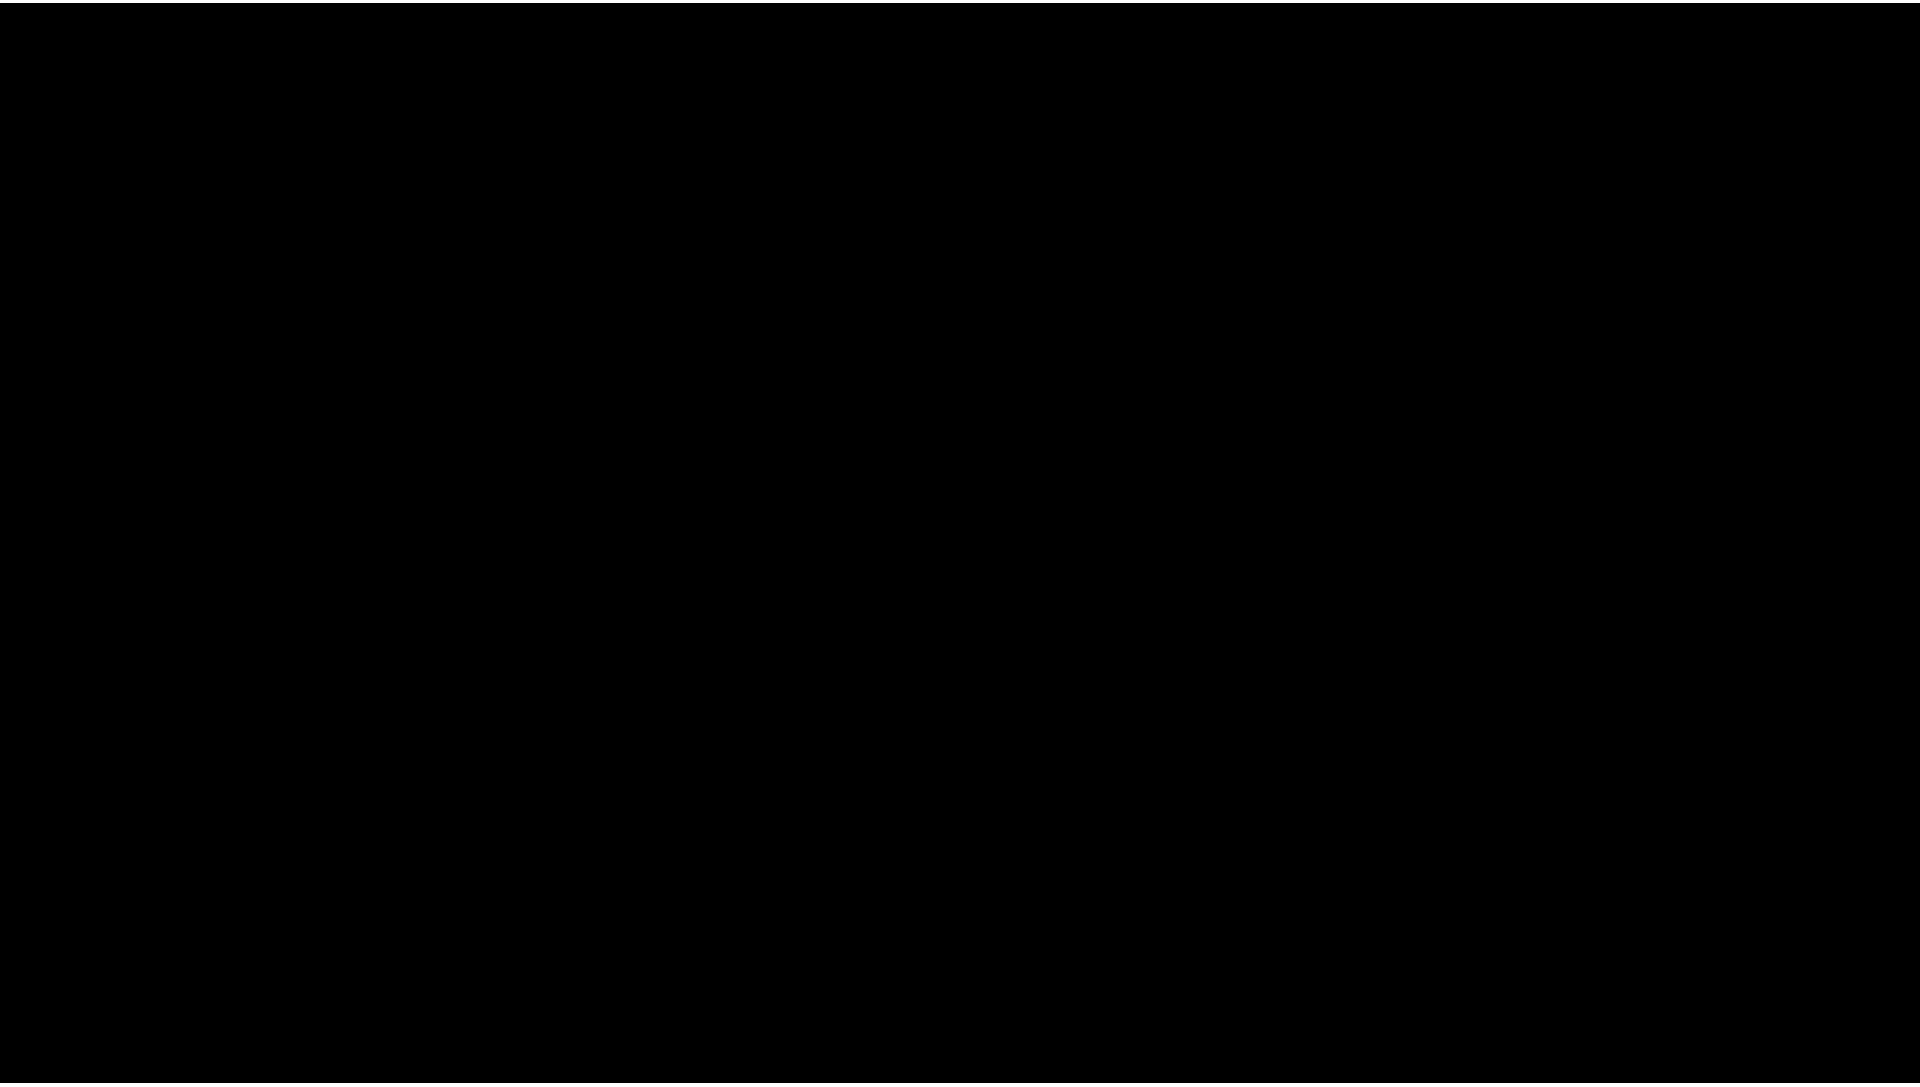

Why This Course?

Visual Computing in the old days…

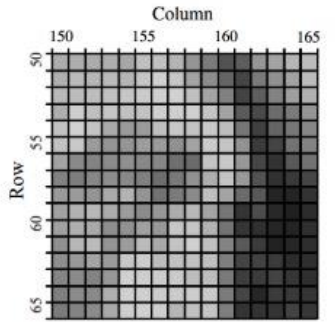

Column

Row	150	155	160	165
50	183 183 181 184 177 200 200 189 159 135 94 105 160 174 191 196			
	186 195 190 195 191 205 216 206 174 153 112 80 134 157 174 196			
	194 196 198 201 206 209 215 216 199 175 140 77 106 142 170 186			
	184 212 200 204 201 202 214 214 214 205 173 102 84 120 134 159			
55	202 215 203 179 165 165 199 207 202 208 197 129 73 112 131 146			
	203 208 166 159 160 168 166 157 174 211 204 158 69 79 127 143			
	174 149 143 151 156 148 146 123 118 203 208 162 81 58 101 125			
	143 137 147 153 150 140 121 133 157 184 203 164 94 56 66 80			
60	164 165 159 179 188 159 126 134 150 199 174 119 100 41 41 58			
	173 187 193 181 167 151 162 182 192 175 129 60 88 47 37 50			
	172 184 179 153 158 172 163 207 205 188 127 63 56 43 42 55			
	156 191 196 159 167 195 178 203 214 201 143 101 69 38 44 52			
65	154 163 175 165 207 211 197 201 201 199 138 79 76 67 51 53			
	144 150 143 162 215 212 211 209 197 198 133 71 69 77 63 53			
	140 151 150 185 215 214 210 210 211 209 135 80 45 69 66 60			
	135 143 151 179 213 216 214 191 201 205 138 61 59 61 77 63			

Image Processing
EECS 225B

Computer Graphics
CS 184

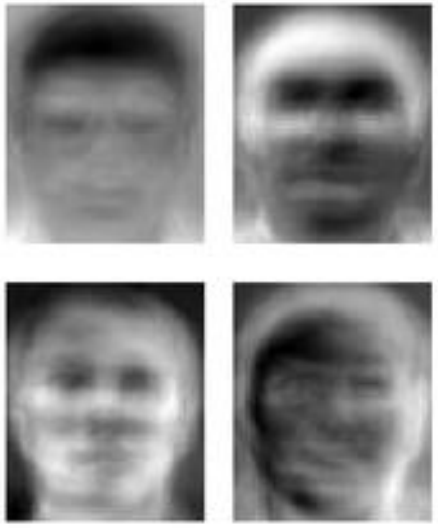

Computer Vision CS 280

Visual Computing gets interconnected

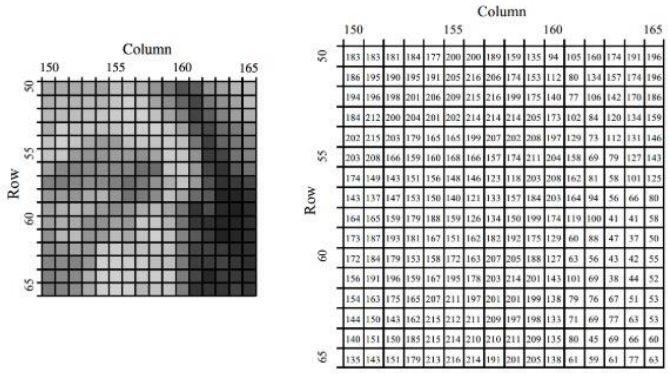

Image Processing
EECS 225B

Art History
ART 10

Computer Graphics
CS 184

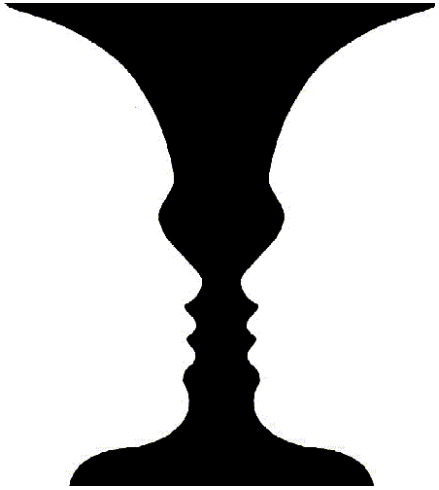

Visual Perception
PSYCH

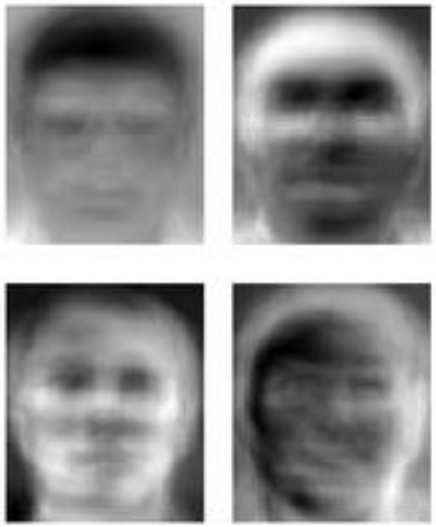

Computer Vision CS 280

Computational Photography

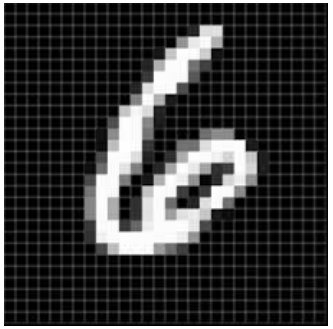

Machine Learning

Visual Computing gets interconnected

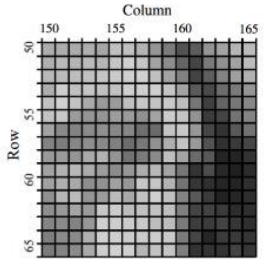

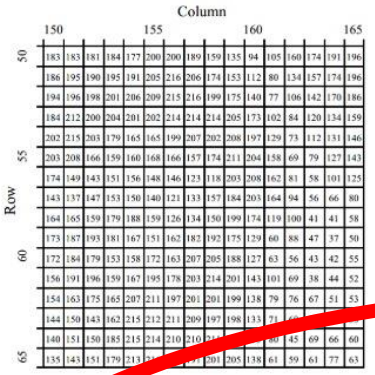

Image Processing
EECS 225B

Art History
ART 10

Computer Graphics
CS 184

CS180

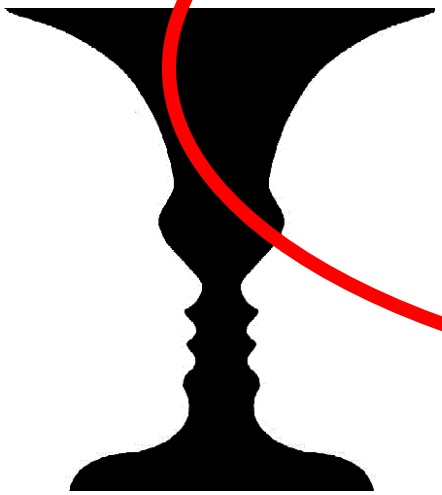

Visual Perception
PSYCH

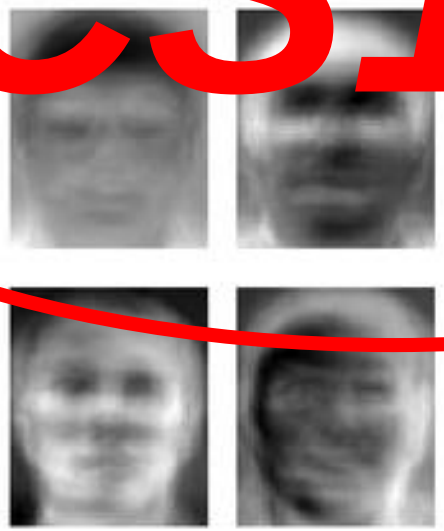

Computer Vision CS 280

Computational Photography

Machine Learning

CS180: Focus on Visual Data

The key objective of this class is **to become friends with every pixel!**

Course objectives

1. You will appreciate the fundamental difficulty of understanding and computing with **visual data**

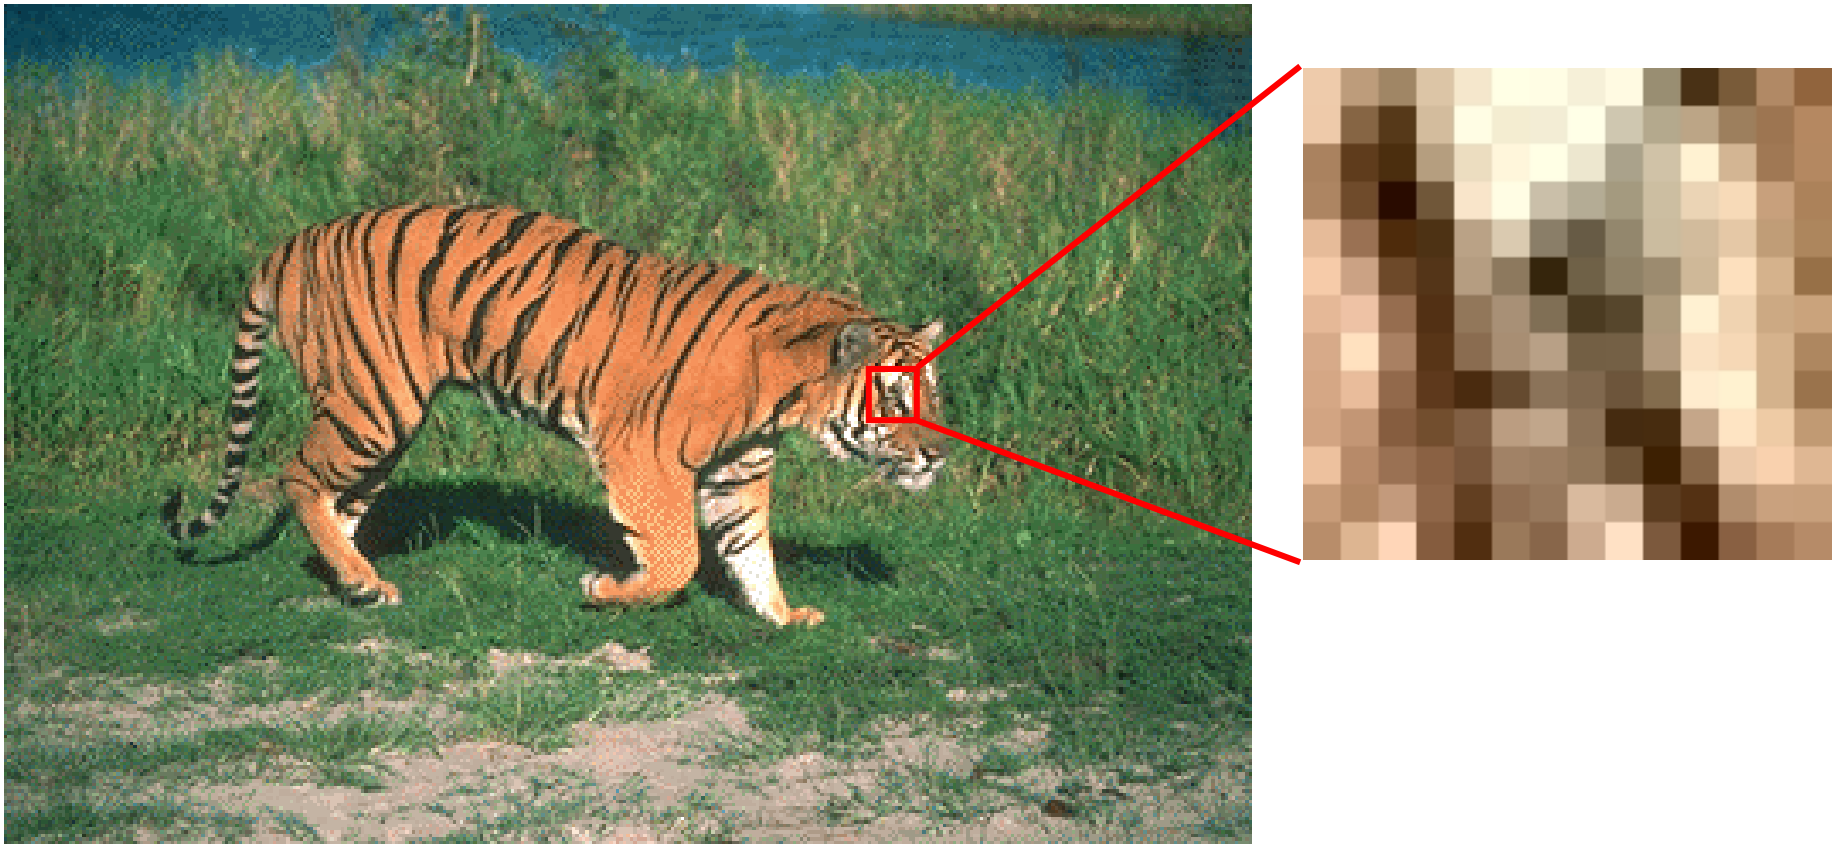

Course objectives

2. You will get a foundation in image processing and computer vision, **from the ground up**:

- Camera basics, image formation
- Convolutions, filtering
- Image and Video Processing (filtering, anti-aliasing, pyramids)
- Image Manipulation (warping, morphing, mosaicing, matting, compositing)
- Projection, 3D, stereo
- Data-driven methods
- Generative Models
- …

Course objectives

3. You will get a more intuitive understanding of important mathematical and computational concepts

- Signal Processing
- Gradients
- Change of basis
- interpolation, extrapolation
- Furrier Transforms
- PCA
- Deep Learning
- Auto-regressive Visual Models
- Diffusion Models
- NeRFs
- …

Course objectives

4. You will learn approaches for **visual synthesis**

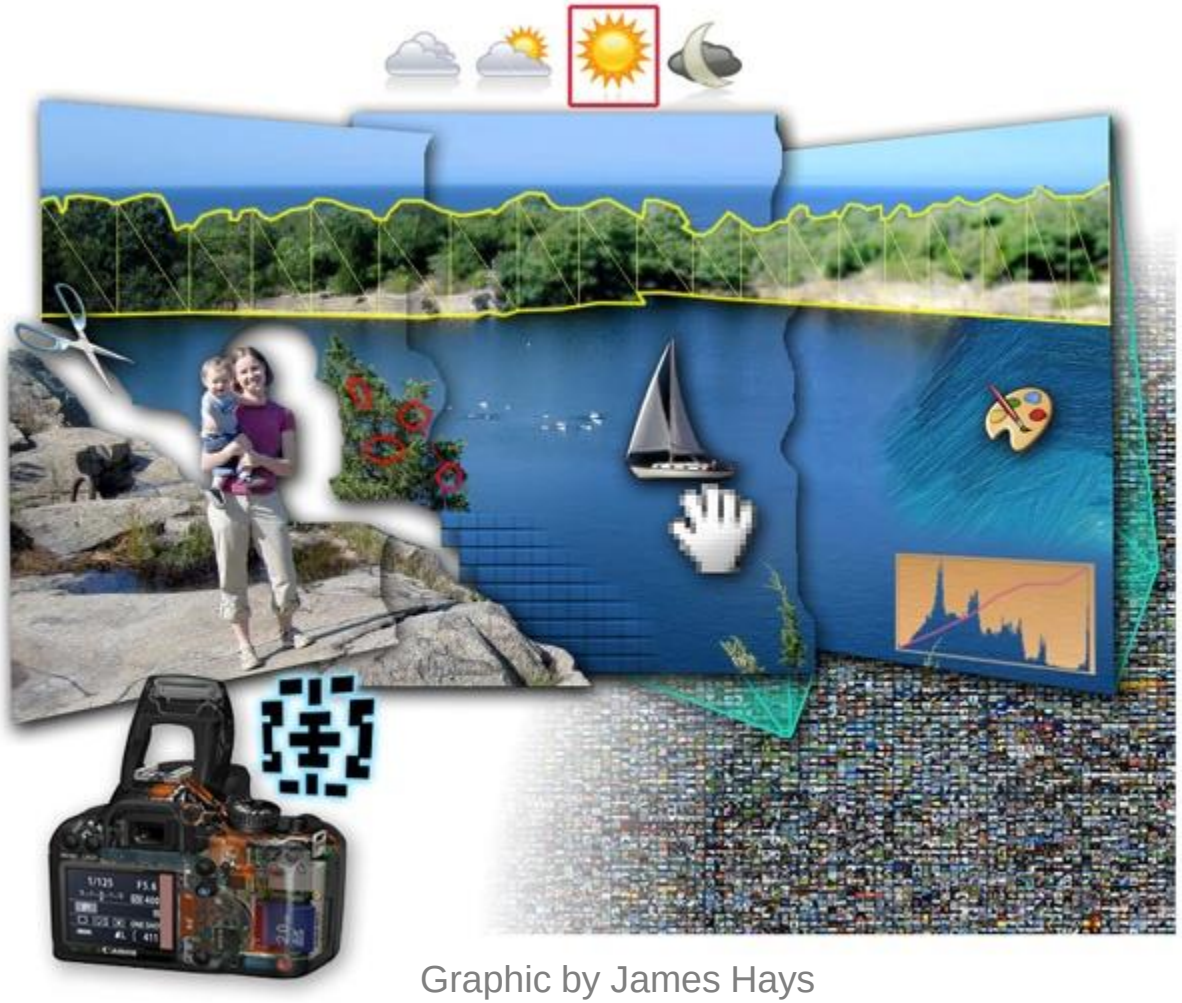

Graphic by James Hays

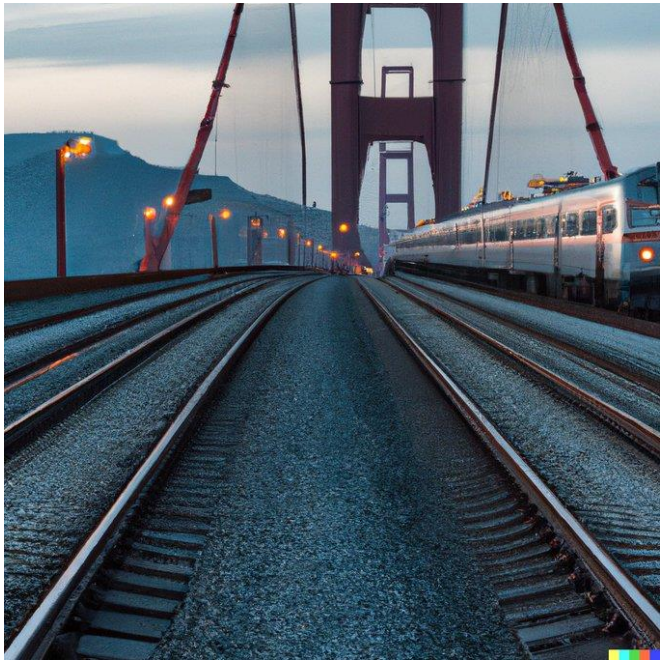

DALL-E + Danielle Baskin

Course objectives

4. You'll better appreciate human visual perception

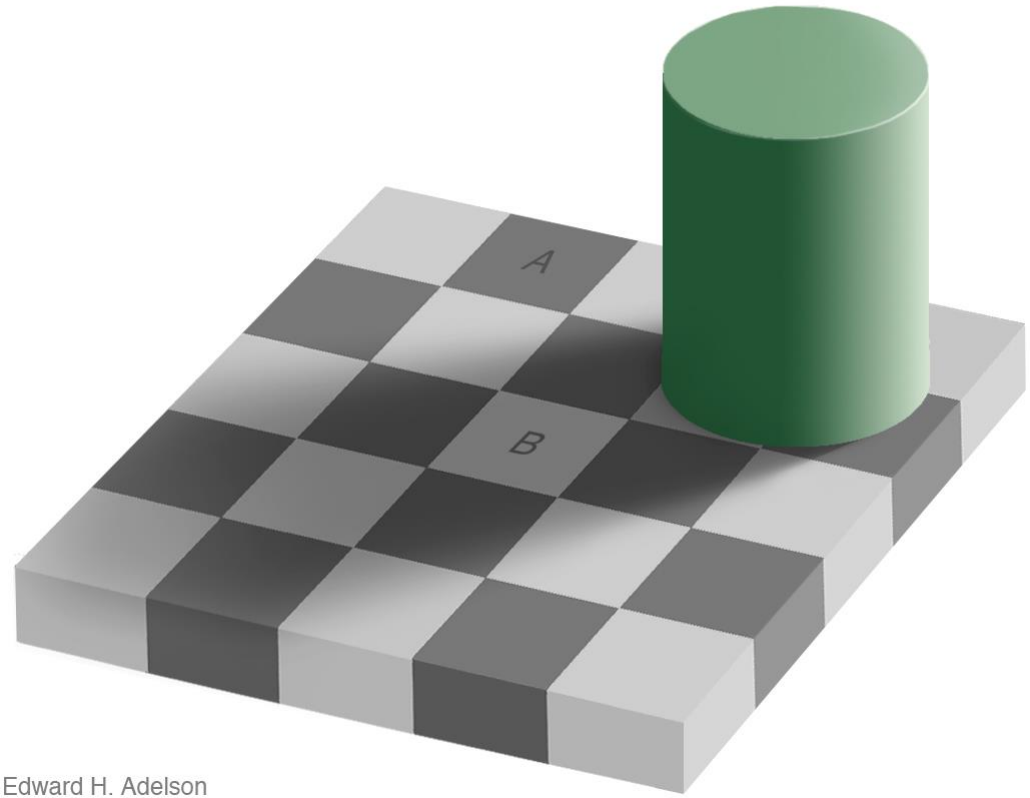

Edward H. Adelson

Course objectives

4. You'll better appreciate human visual perception

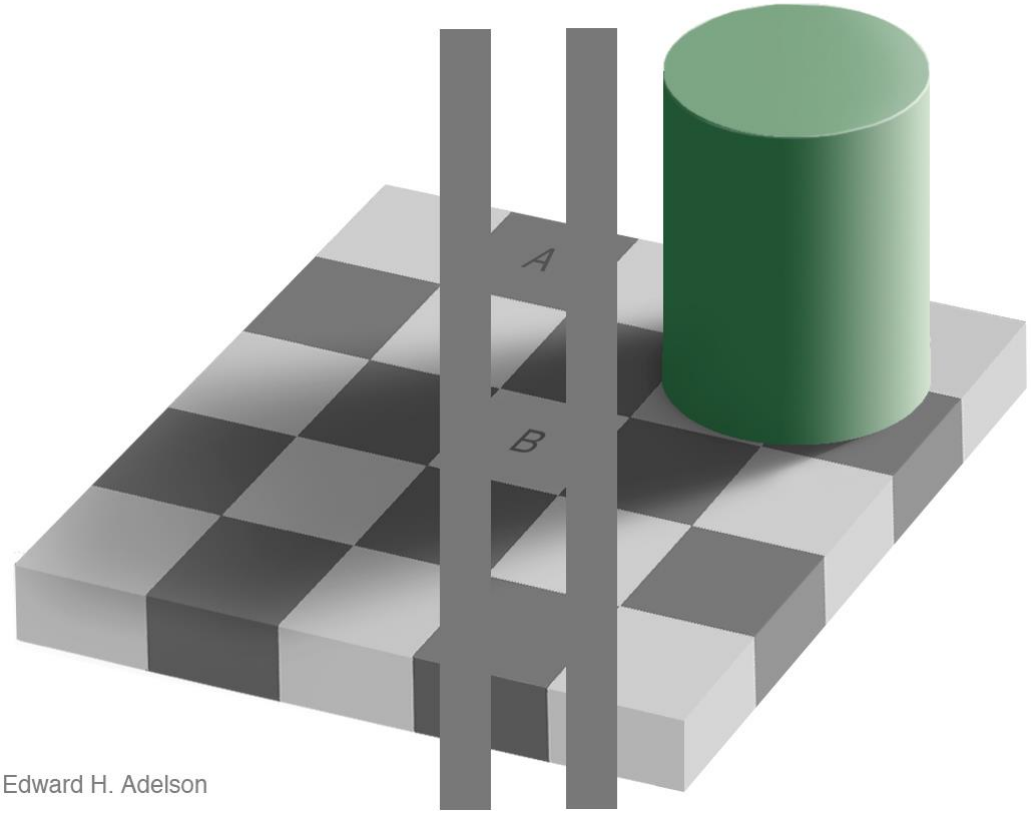

Edward H. Adelson

Different people see different things

https://en.wikipedia.org/wiki/The_dress

Course objectives

5. You will learn about the **history of ideas** in visual computing
   - Did you know Large Generative Models go back to 1940s?
   - Or that Deep Learning started with a Nobel Prize on Neuroscience of the Visual Cortex in the 1960s?
   - …

Course objectives

6. You'll have fun doing cool stuff, coding up a storm, largely from scratch

Programming Project #0

Becoming Friends with your camera:

1. Learn to take better selfies
2. Experiment with perspective compression
3. Create **a Dolly Zoom** shot

Project out TODAY, due Tuesday!

Programming Project #1

Prokudin-Gorskii's Color Photography (1907)

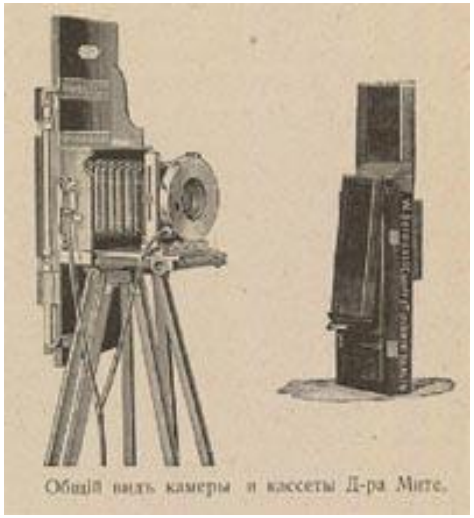

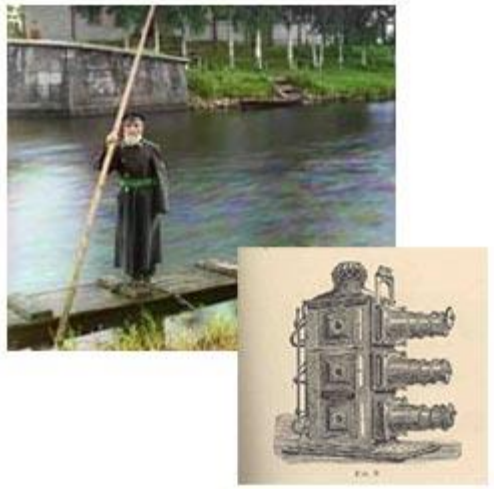

Programming Project #1

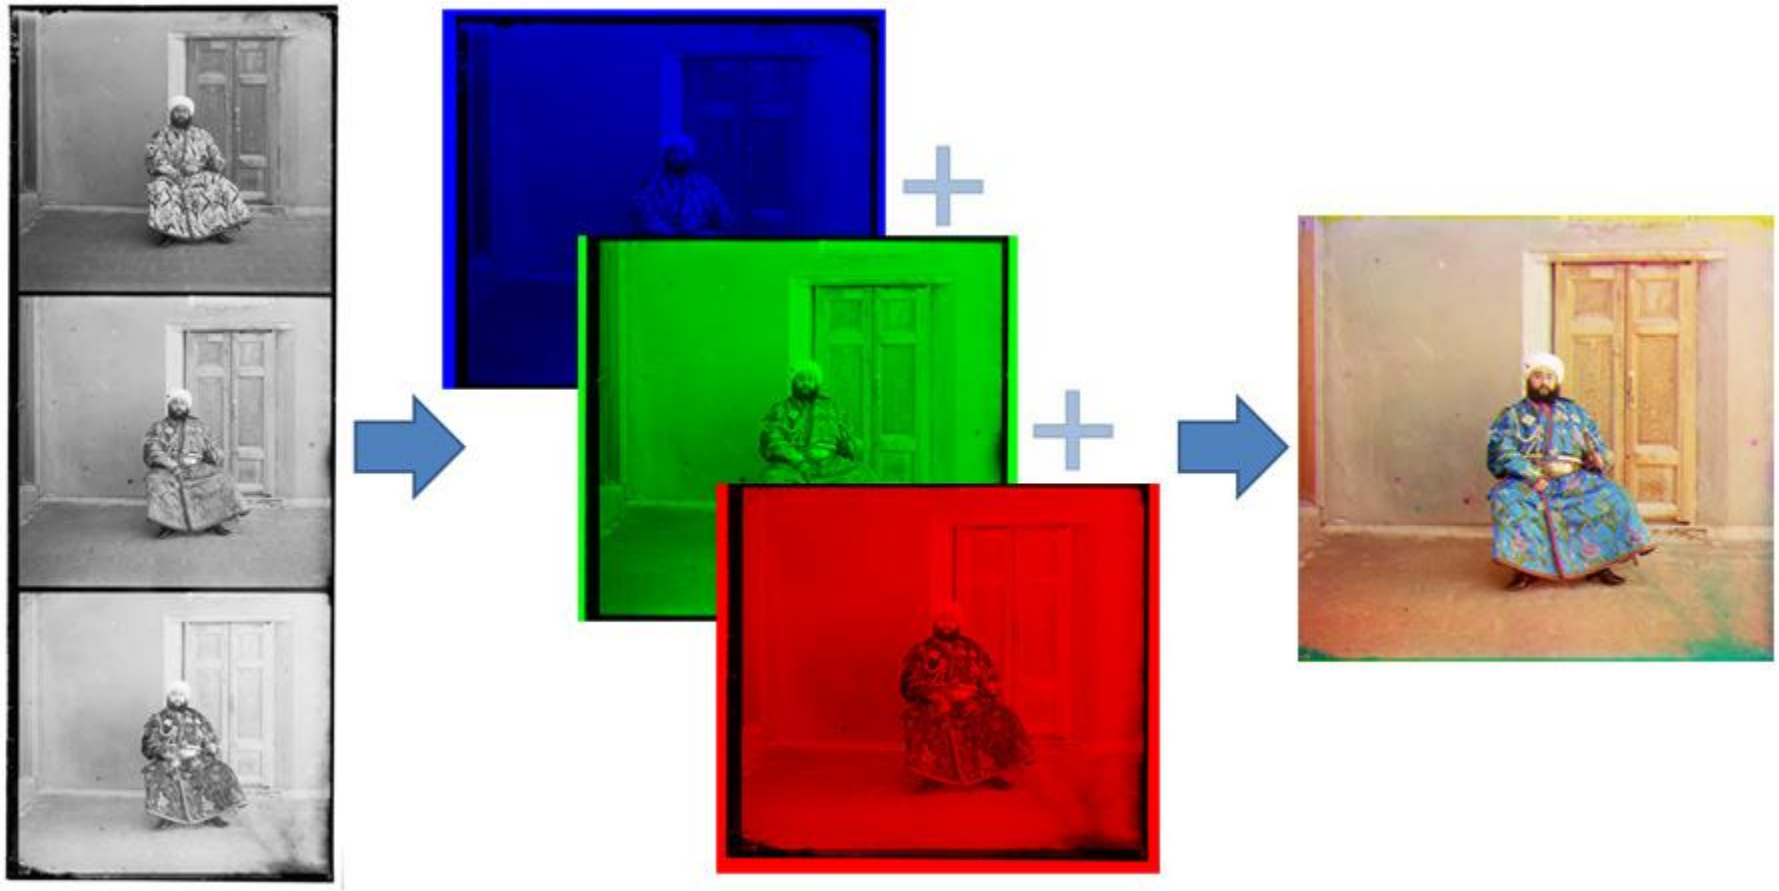

Project 2: Fun with frequencies

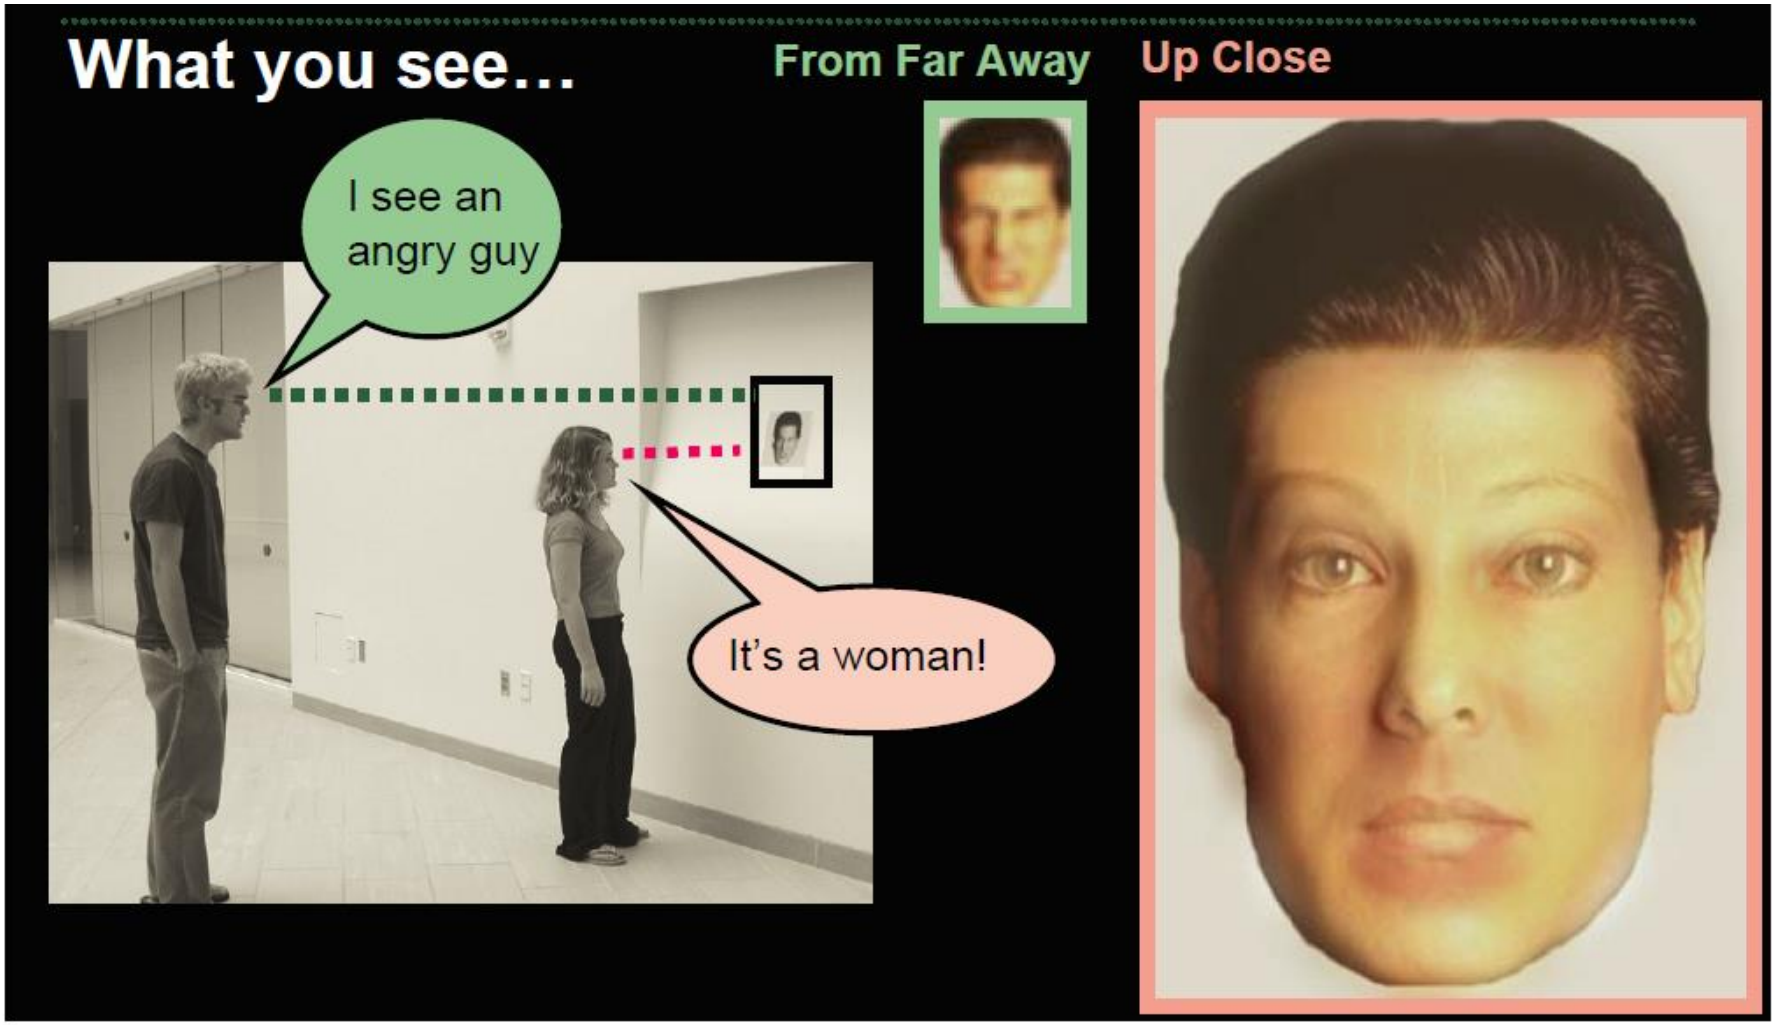

Project 2: Fun with frequencies

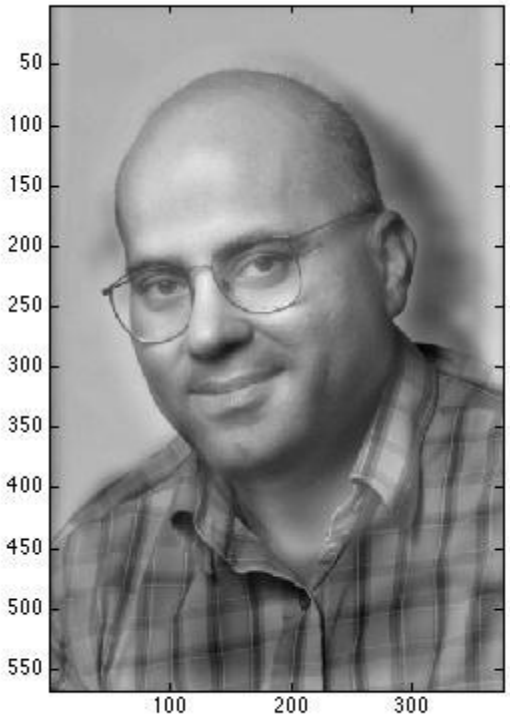

Prof. Christos Papadimalik

Project 2: Fun with Frequencies

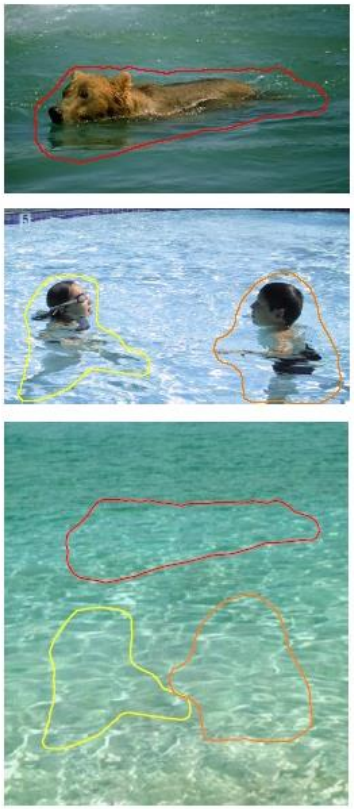

sources/destinations

cloning

seamless cloning

Project 3: Panorama Stitching

Photo Mosaics

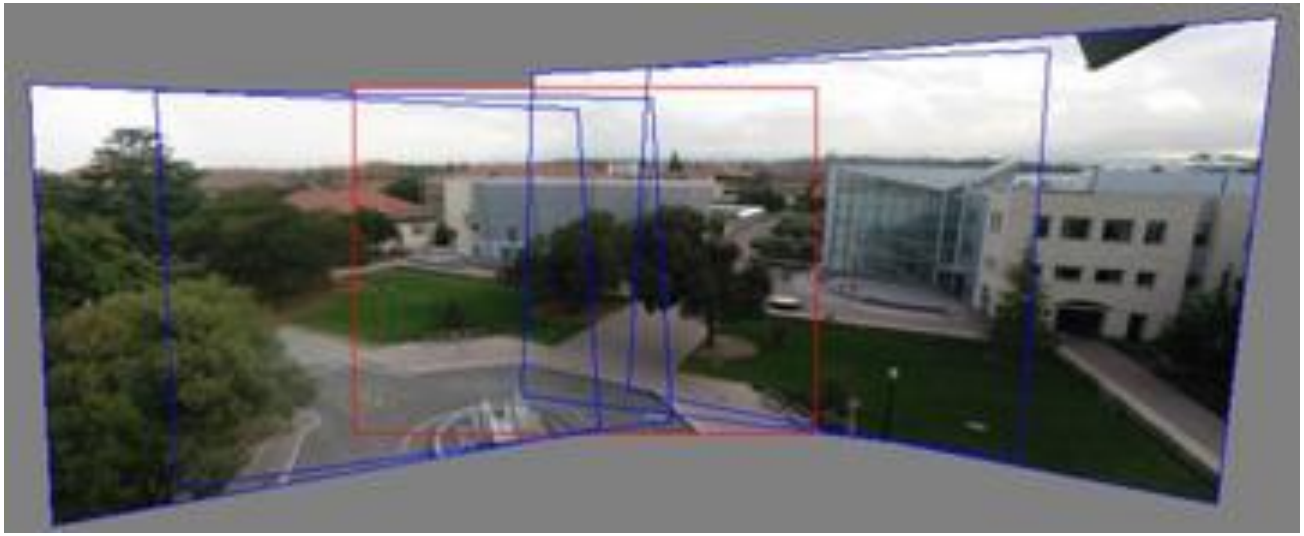

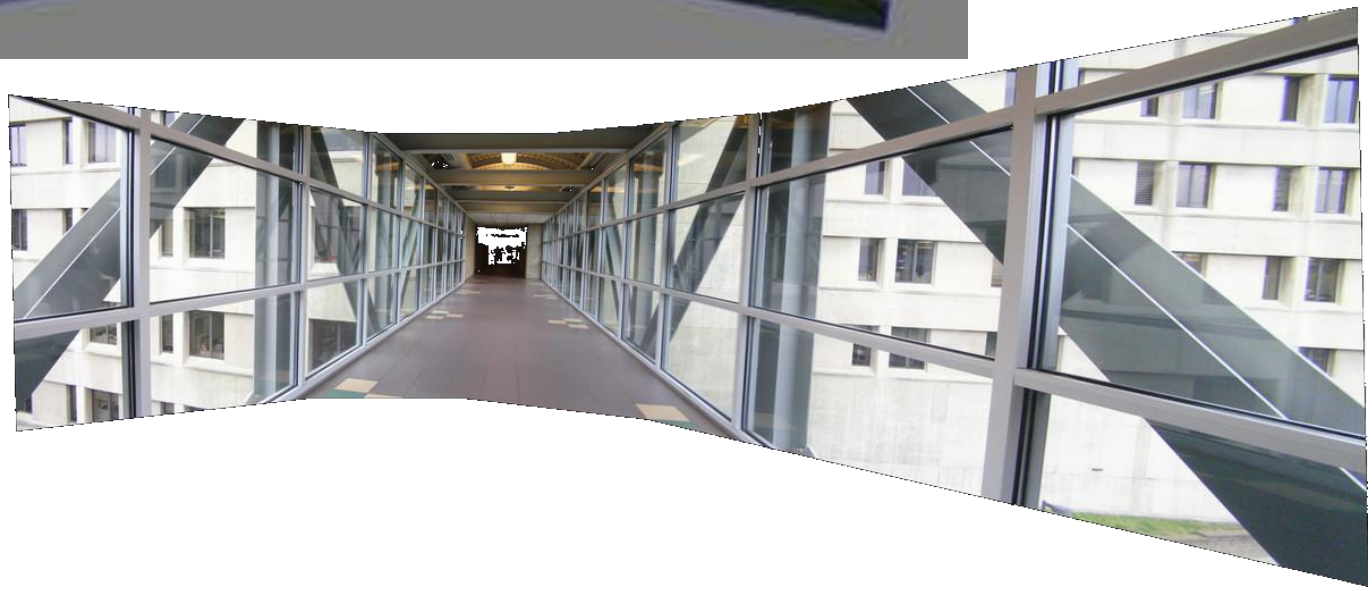

Project 4: Neural Radiance Fields

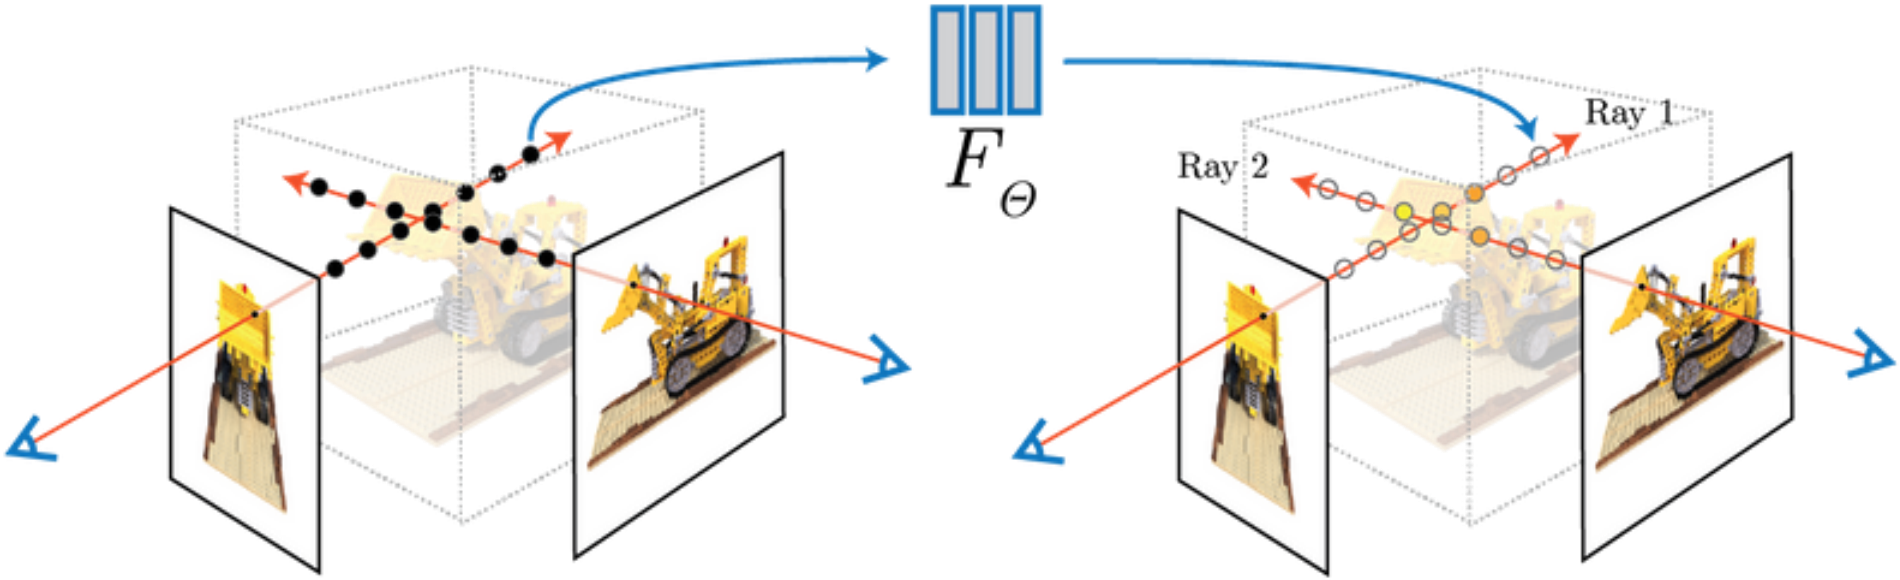

Project 5: Fun with Diffusion

a photo of an old woman

a painting of a deer

For each project:

Derive the **math**, implement stuff **from scratch**, and apply it to your **own** photos

Every person does their own project (no groups!)

Reporting via web page (plus submit code)

Class Organization / Administrivia

General

Prerequisites

- Linear algebra!!! (EE16A, Math 54, Math 56, or Math 110)
- Multivariable Calculus (Math 53)
- Good programming skills (at least CS61B)
- Deep Learning experience (e.g. CS189, CS182, may be concurrent)
- Creativity

Emphasis on programming projects!

- Building something from scratch

Graduate Version (CS280A):

- Need to do “bells & whistles” for each project
- different grading curve than CS180

Administrative Stuff

Grading

- 50% Programming Projects (5 total)
- 20% Midterm (**Oct 23, in class**)
- 25% Final Exam (**Dec 19**)
- 5% Pop Attendance Quizzes

Late Policy

- Five (5) **emergency** late days for semester**.** The expectation is you will never use them.
- 10% off each additional day afterwards

Rule # 1:
No lecture recordings

This is **an in-person class**. You are to come to the lecture and participate! Attendance is required.

Rule # 2:
Deadline is a deadline

In real life there are no extensions

This is a FUN but INTENSE class, projects come **one after another**

Slip days are for **emergencies.** Unless something terrible happens to you, you should have all your slip days left.

Projects are time consuming. Start early!!!

Rule # 3:
TA’s don’t debug code

TA’s don’t debug code for you.

Part of the skill is to learn how to ask questions to debug the issue without presenting the code

Visualize the results and send those to figure out what is wrong

Use the pixels – become friends with visual debugging

Getting help outside of class

Course Web Page

- https://cal-cs180.github.io/fa25/

Online forums:

- **Ed**
- **Gradescope** (add yourself via course webpage)

Discussion Sections:

- 6-9 sections each week (watch announcements)
- Do attend, **specially first one next week!**

Office hours

- For profs: after each lecture
- For others, see webpage

Textbooks

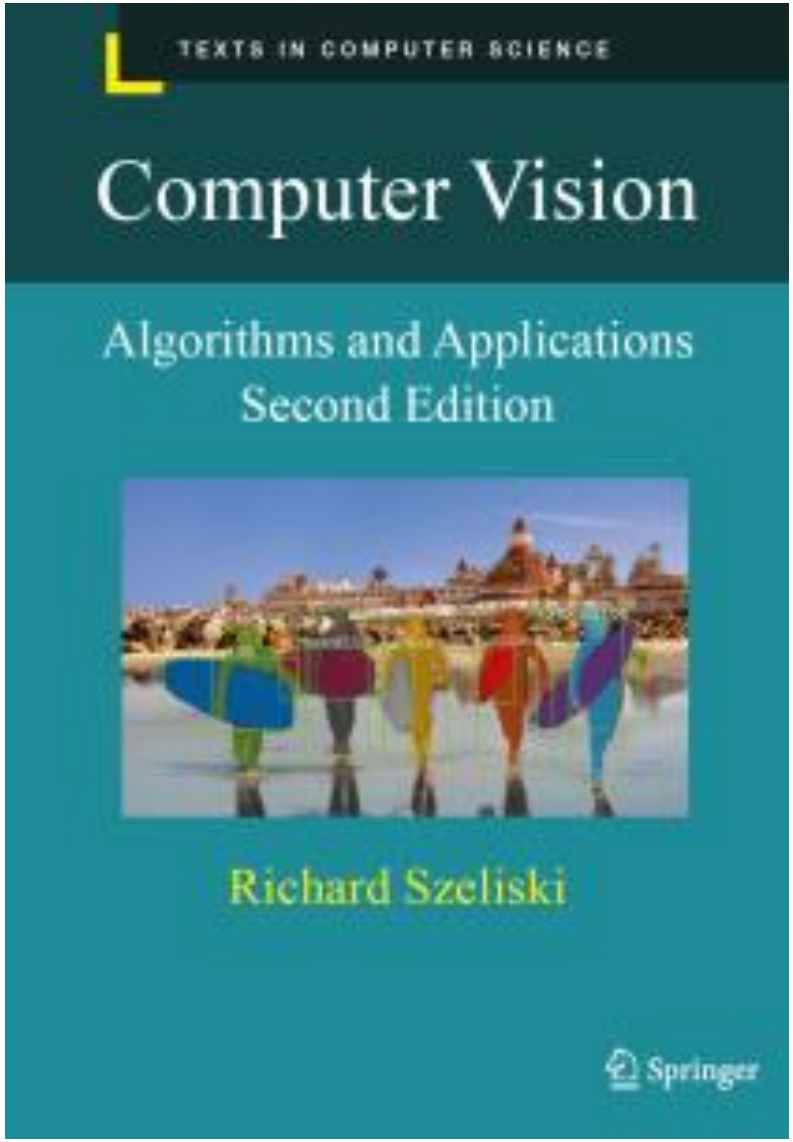

http://szeliski.org/Book/

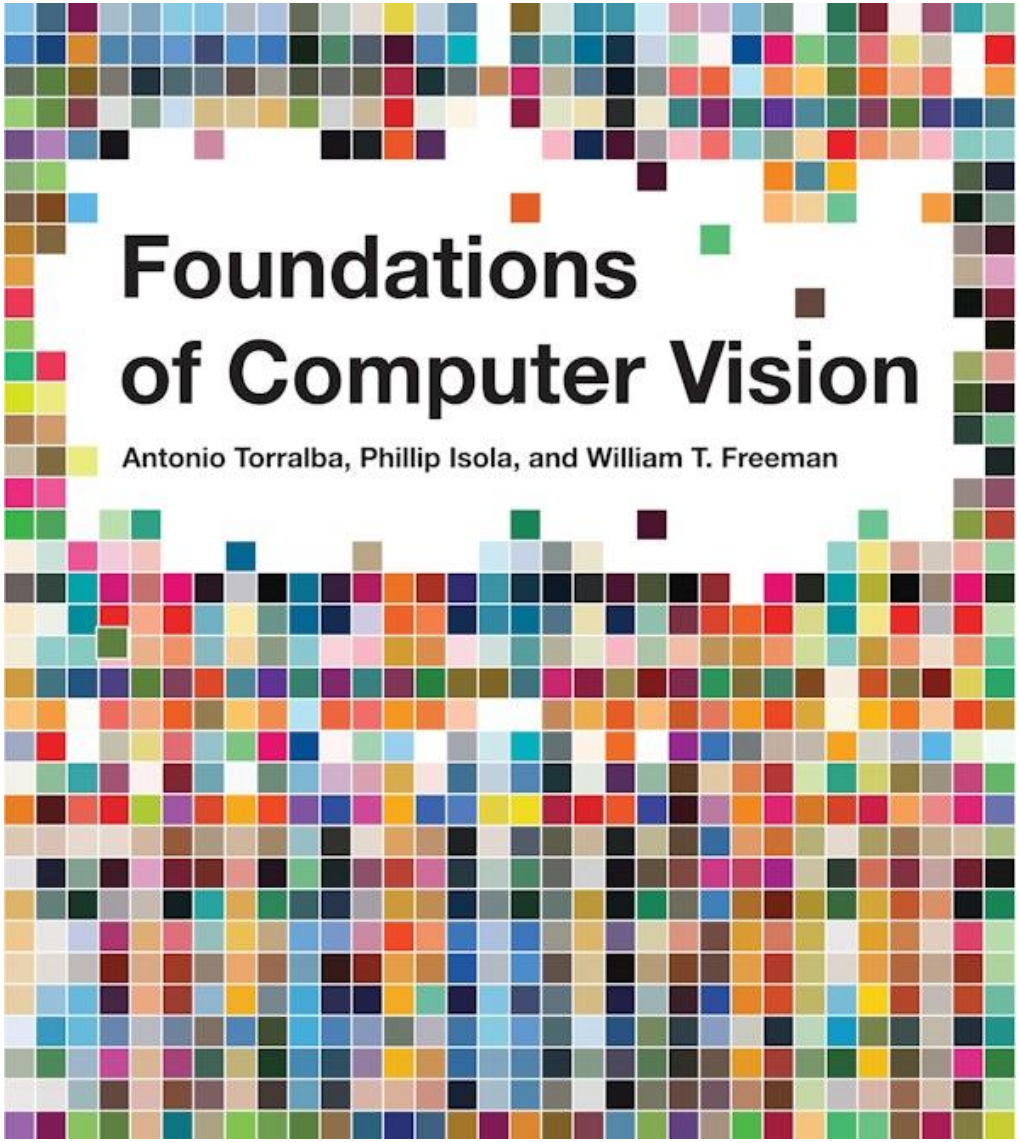

https://mitpress.mit.edu/9780262048972/foundations-of-computer-vision/

Academic Integrity

- You are encouraged to discuss projects, but **never share code**
  - You must type every character yourself
- Just like in any academic setting, cite any sources and inspirations

Our LLM policy

- GPT, Claude, etc are wonderful tools
  - And so is calculator, Wolfram Alpha, Wikipedia, Stack Overflow, etc.
  - but before you use a calculator, it's important to learn how to do long division by hand.
- We hope you are here to learn! In this course, we want you to do things **from scratch**.
- Letting LLMs to code the assignments for you will be considered **cheating**
- Exams will test your knowledge of the code you have written during the semester.
  - E.g. we might ask you to write code during exam.

Waitlists

- We are expecting to add 50-60 seats to cs180
- Also, we expect 50-70 people to drop after the first couple of projects ☺
- So, chances are good to get in, **but you need to start doing projects**

For CS280A waitlist
https://forms.gle/fdEPsXs4HAubxqKN8

For Concurrent Enrollment (CE)
https://forms.gle/JLs8Jvb1eWXTpdg4A

Warning:
historically high GPA of this course

- Survivor bias
- High class GPA != easy course
- This is a FUN but INTENSE class
- Many people will drop out, switch to pass/fail.

Why you should NOT take this class

- Project-based class
  - No canned problem sets
  - No clean rubrics
  - Open-ended by design
  - Coding from scratch
  - Will try to make sure everyone understands the basics super-well, before covering advanced topics
- Need time to think, not just hack
  - **Creativity** is a class requirement Aesthetics is part of the grade.
  - We already expect you to know Deep Learning!
- Lots of work…There are easier classes if
  - you just need some units
  - you care more about the grades than about learning stuff
- Not worth it if you don’t enjoy it

Now… reasons TO take this class

- It's your reward after 3 grueling years ☺
- You get to work with pictures, unleash your creative potential
- Gateway to CS280… grad school…

A Brief History of the Visual Data

Depicting Our World: The Beginning

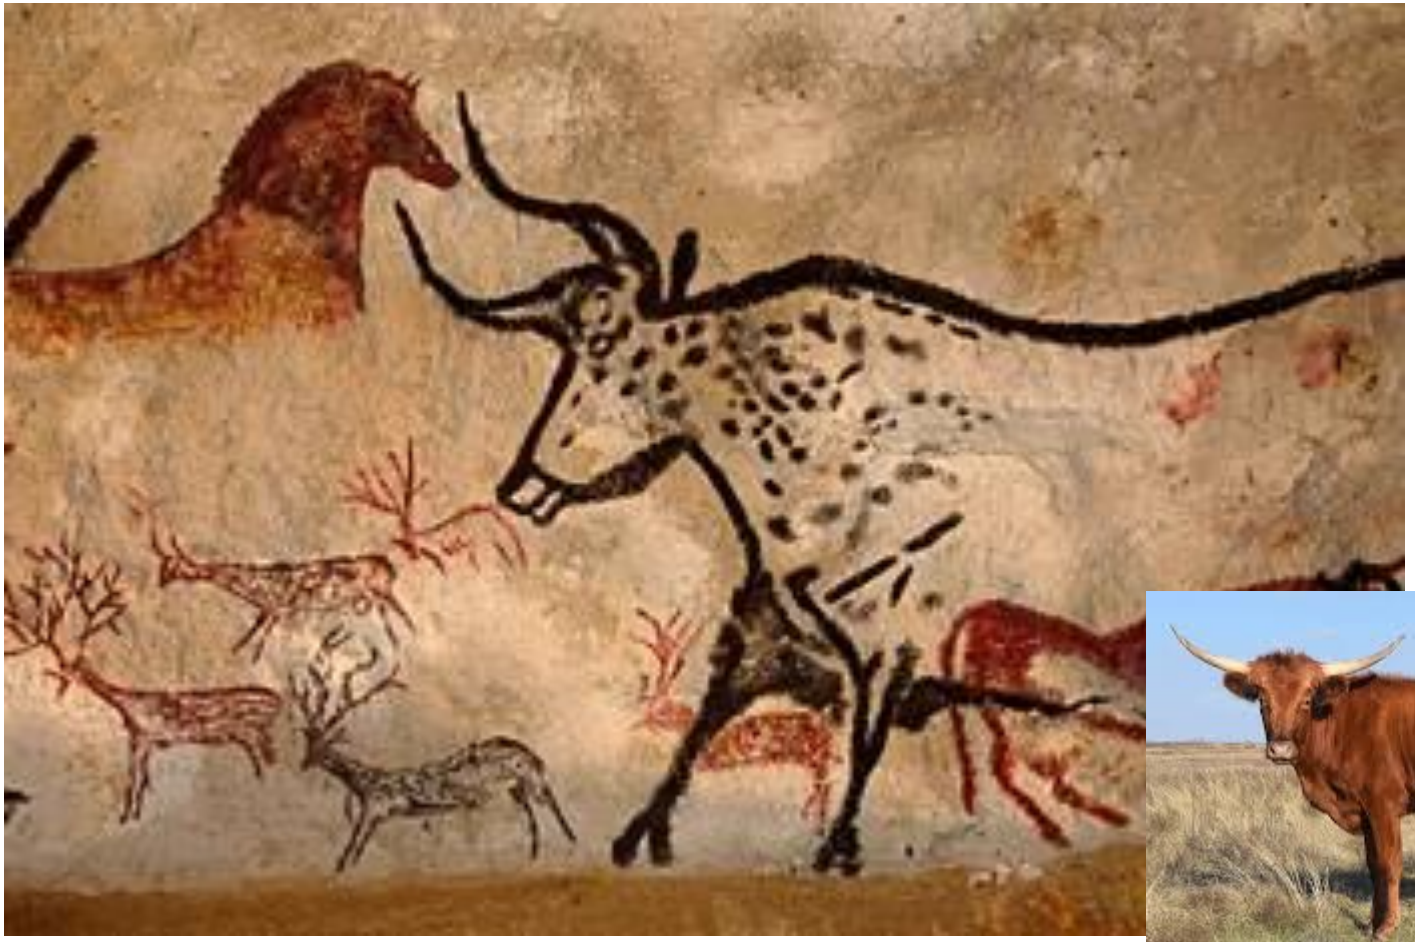

Prehistoric Painting, Lascaux Cave, France
~ 13,000 -- 15,000 B.C.

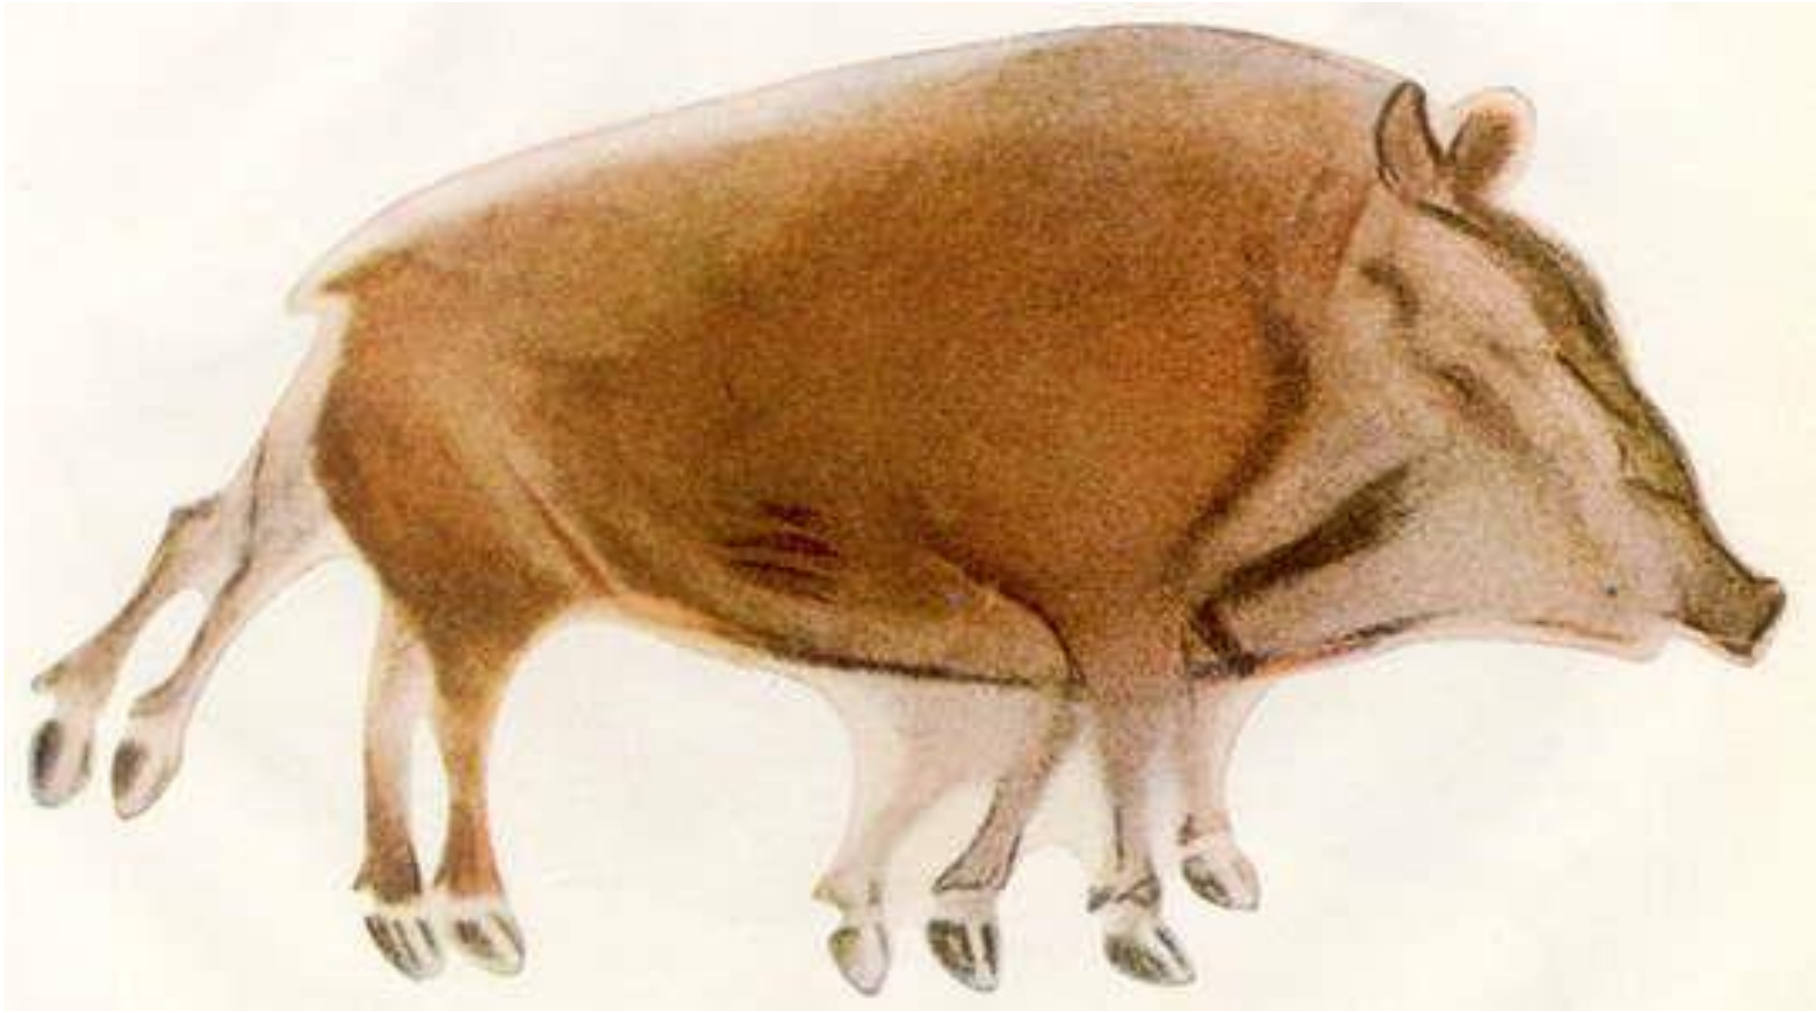

Prehistoric Cave Painting, Altamira
~ 20,000 – 15,000 B.C.

Depicting Our World: Middle Ages

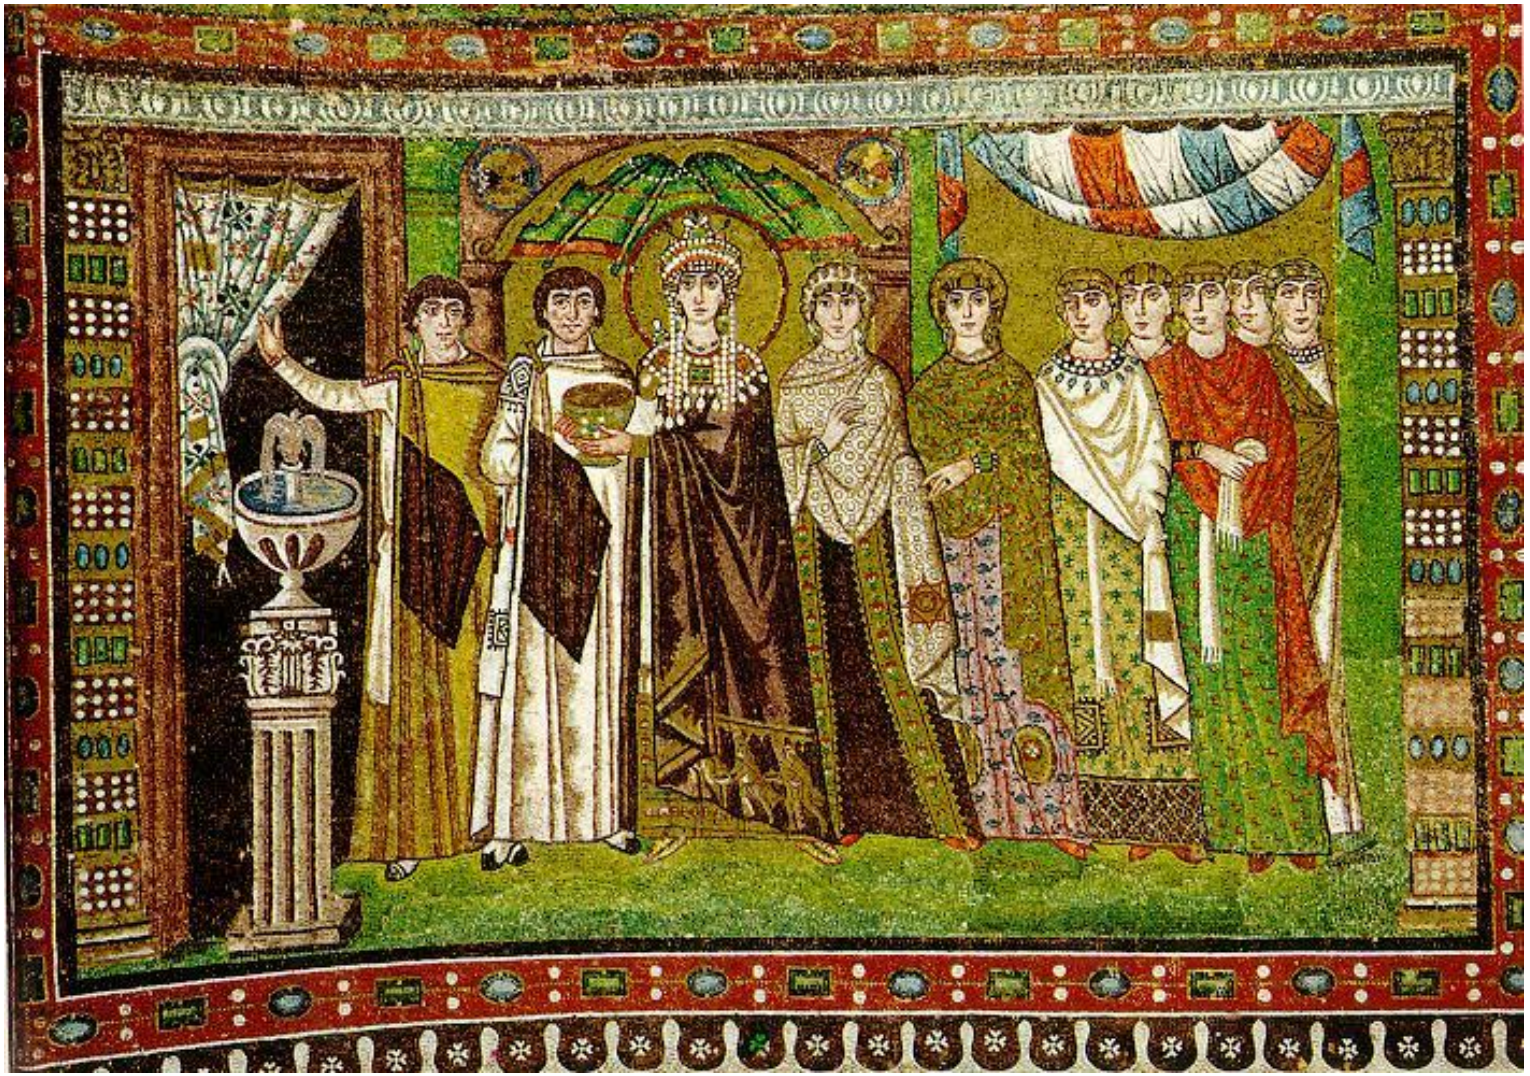

The Empress Theodora with her court.
Ravenna, St. Vitale 6th c.

Depicting Our World: Middle Ages

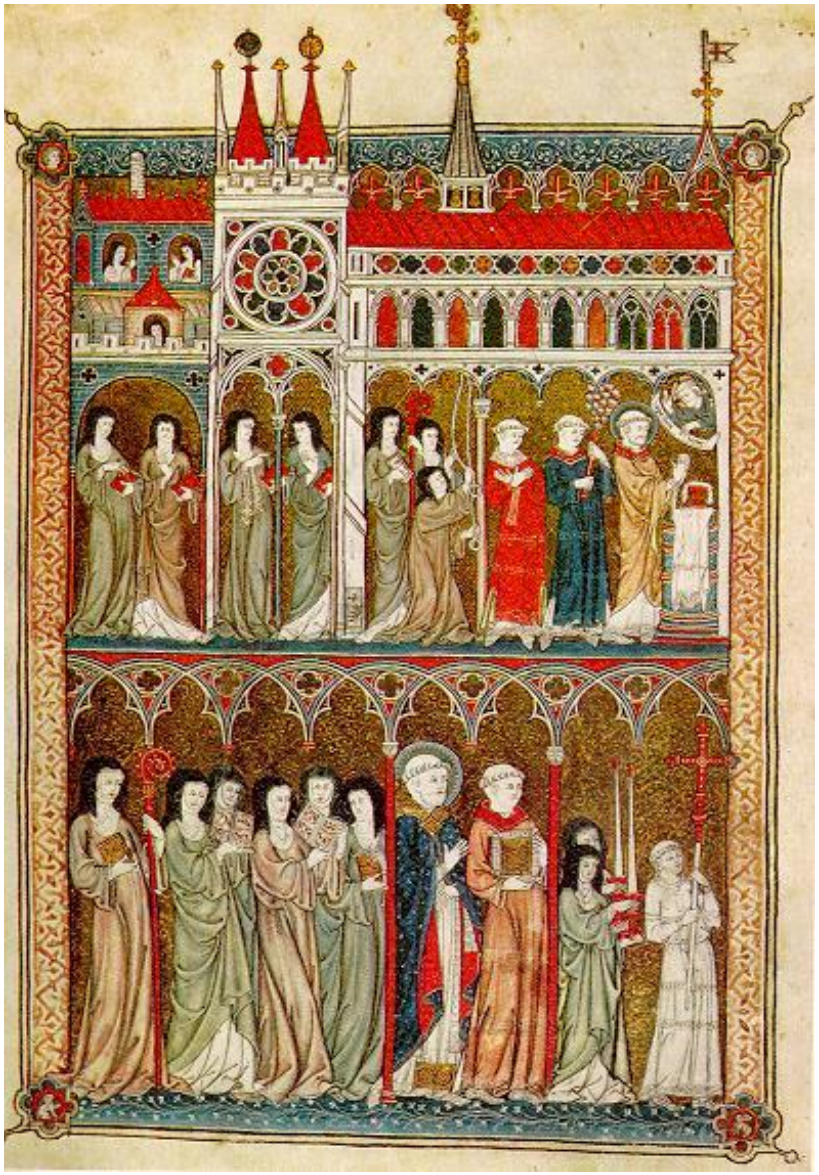

Nuns in Procession. French ms. ca. 1300.

Beginnings of the Renaissance

Giotto, *The Mourning of Christ*, c.1305

Depicting Our World: Renaissance

North Doors (1424)

Lorenzo Ghiberti (1378-1455)

East Doors (1452)

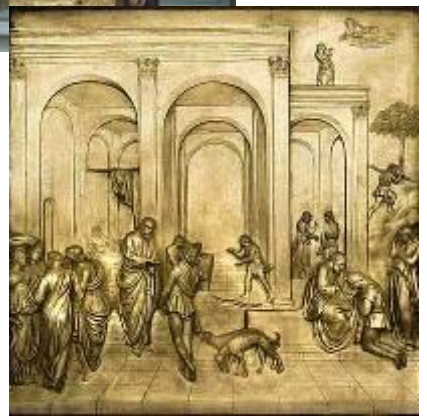

Depicting Our World: Renaissance

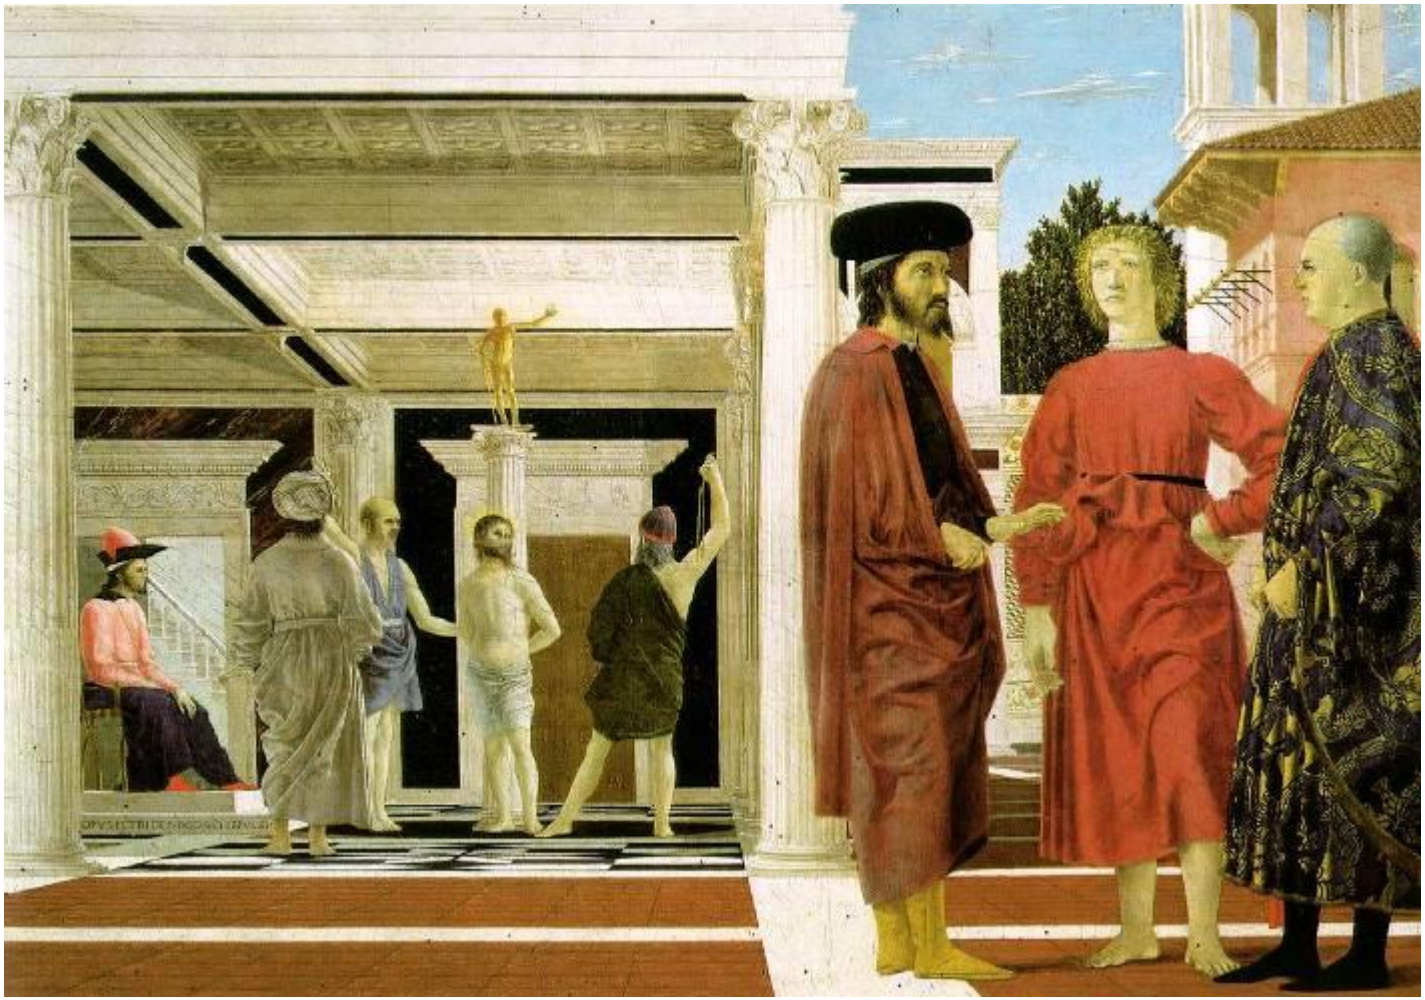

Piero della Francesca,
The Flagellation (c.1469)

Depicting Our World: Toward Perfection

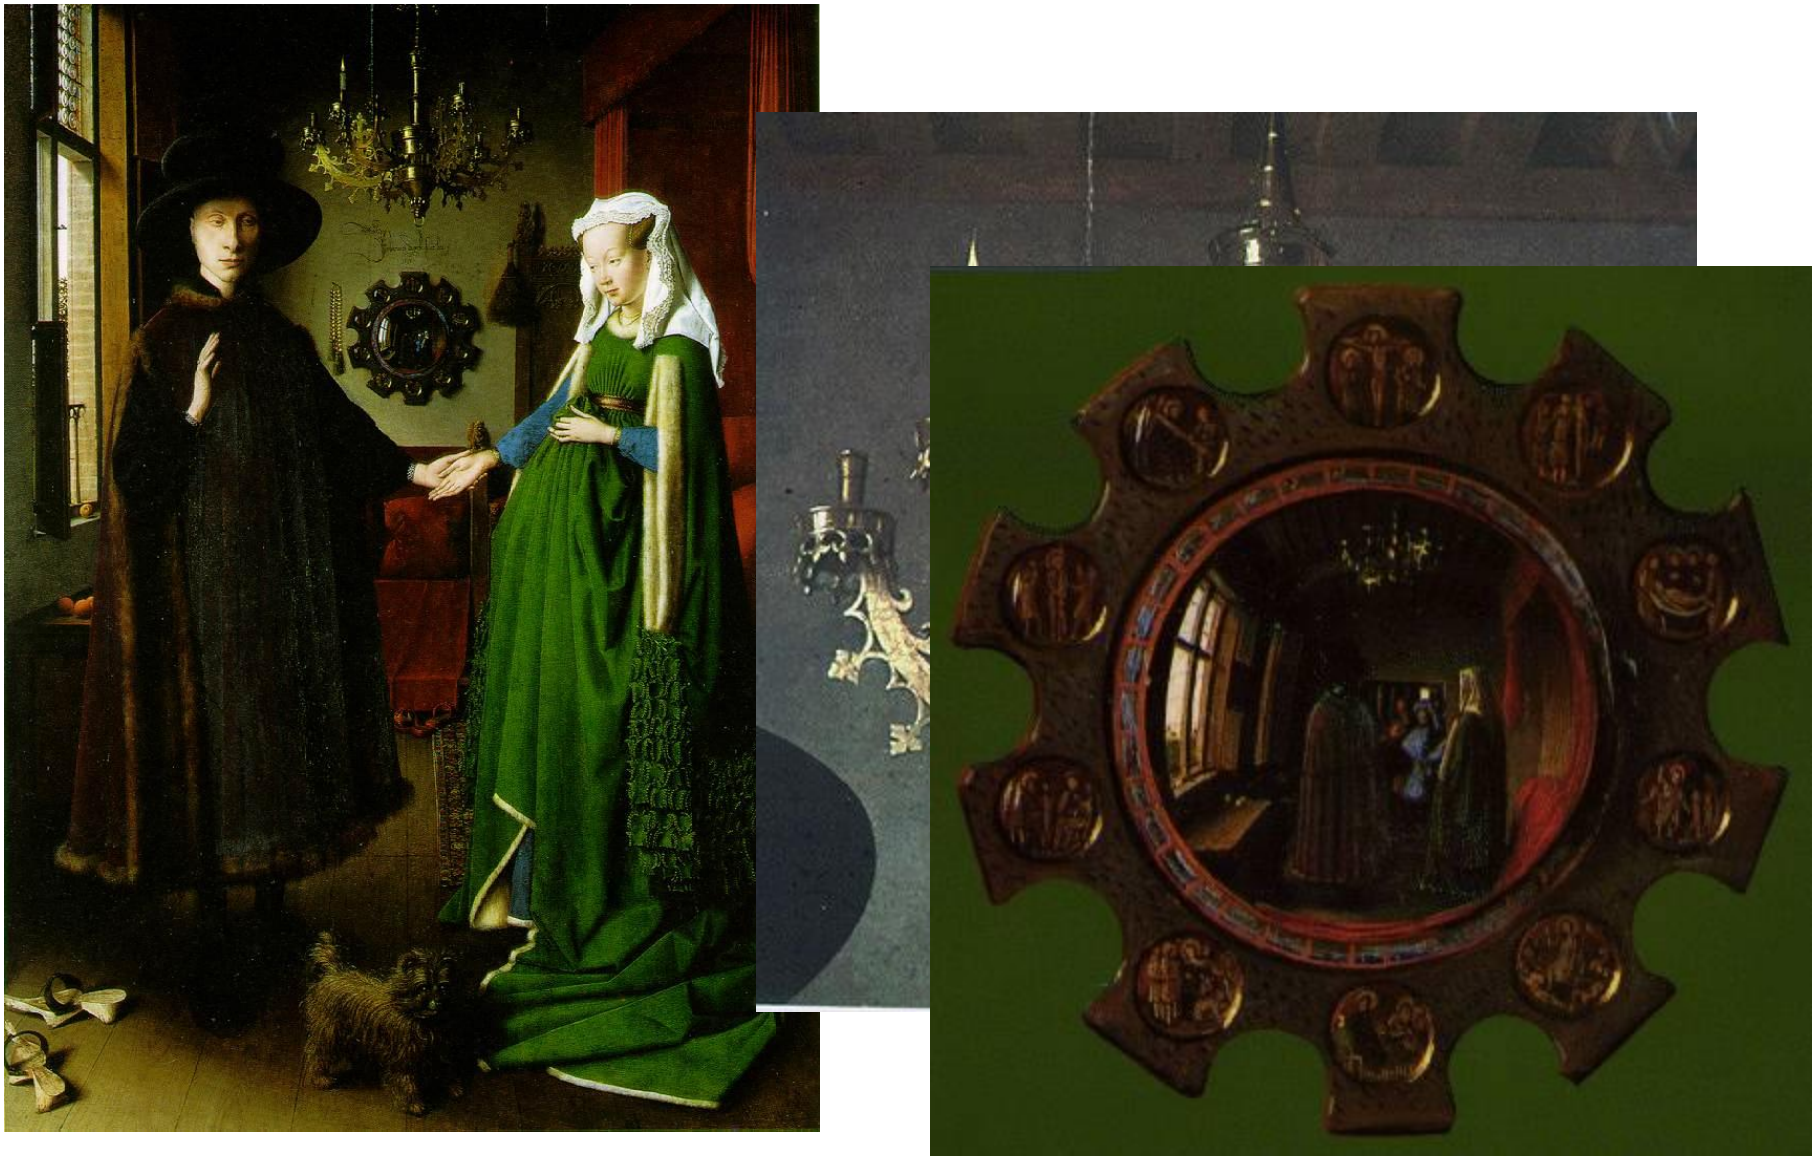

Jan van Eyck, *The Arnolfini Marriage (c.1434)*

Depicting Our World: Toward Perfection

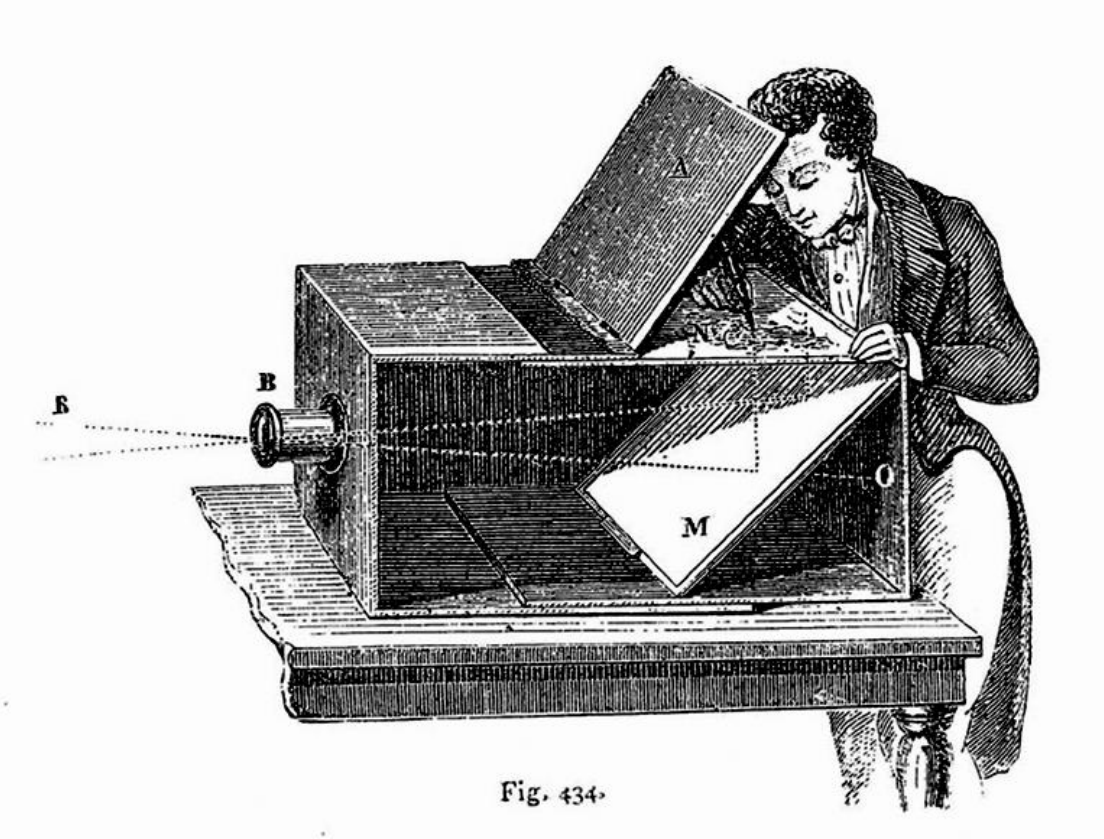

Lens Based Camera Obscura, 1568

Depicting Our World: Perfection!

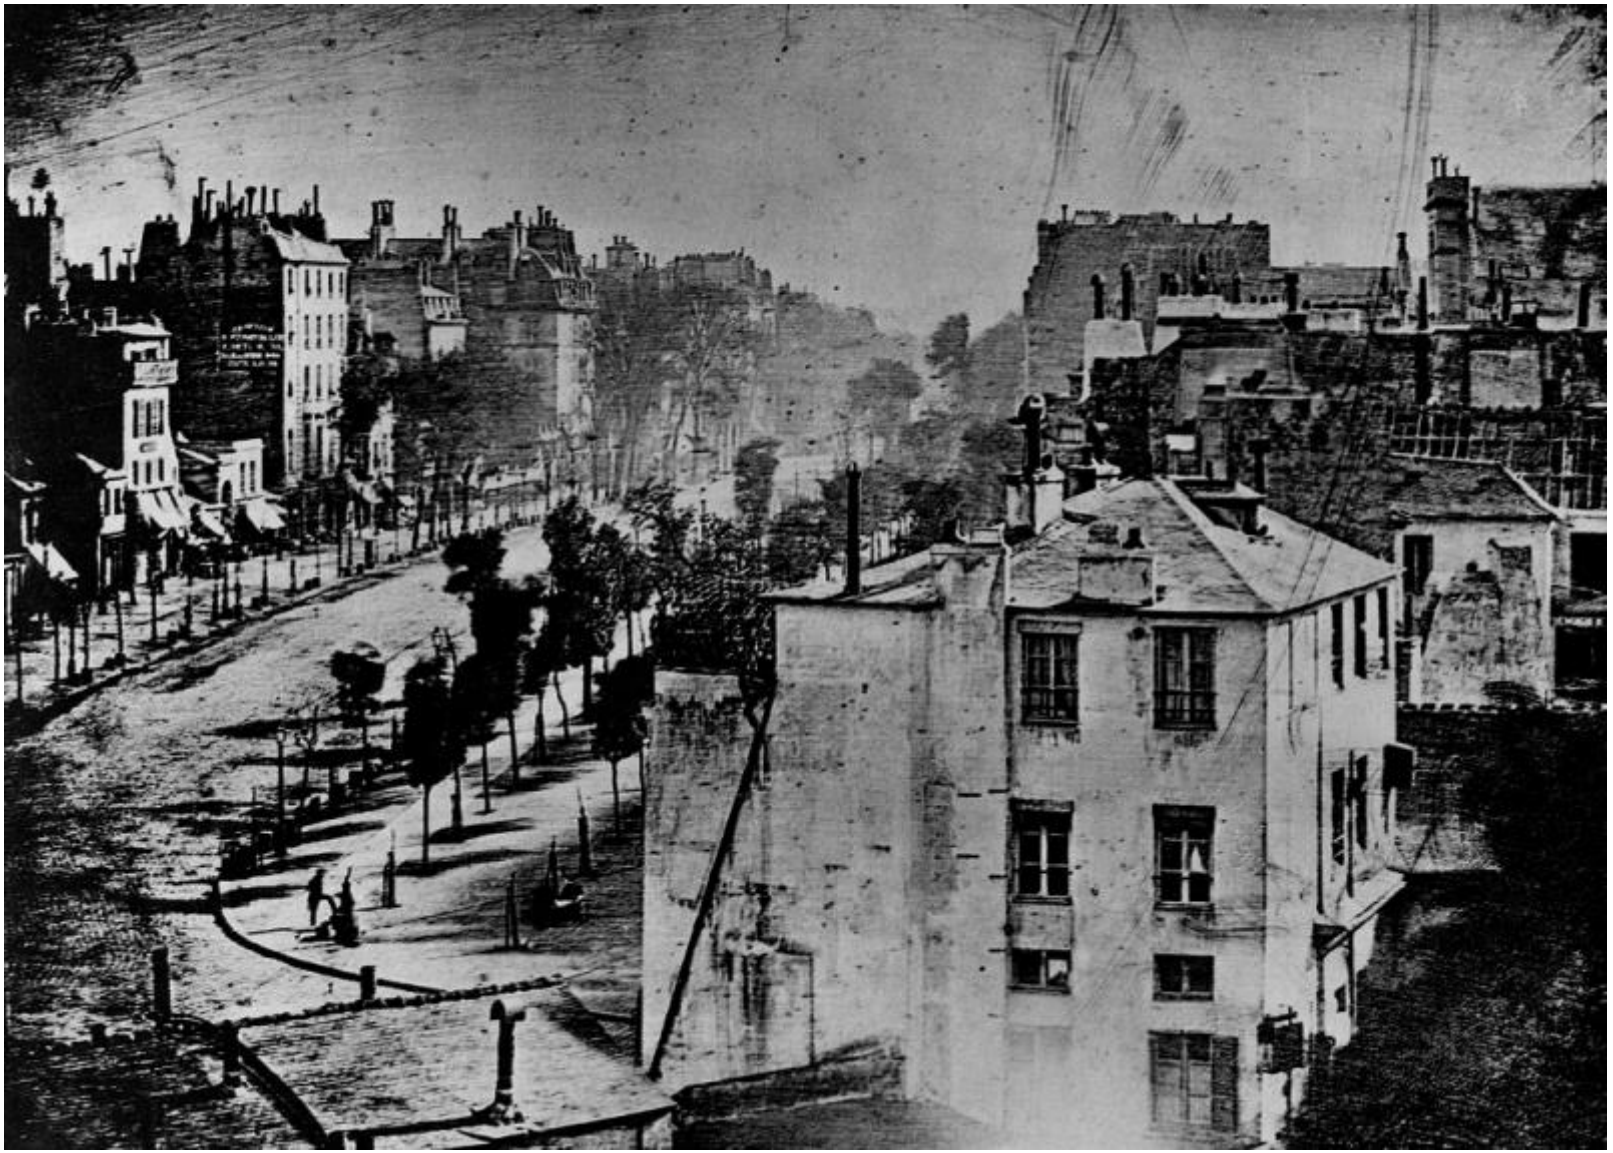

Boulevard du Temple, Louis Daguerre, 1838

Depicting Our World: Realism?

Paris, according to Flickr

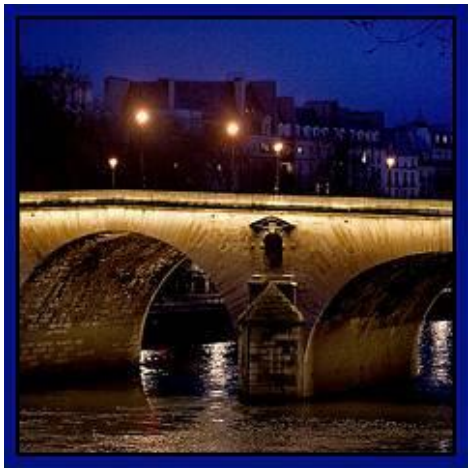

Paris, according to Google StreetView

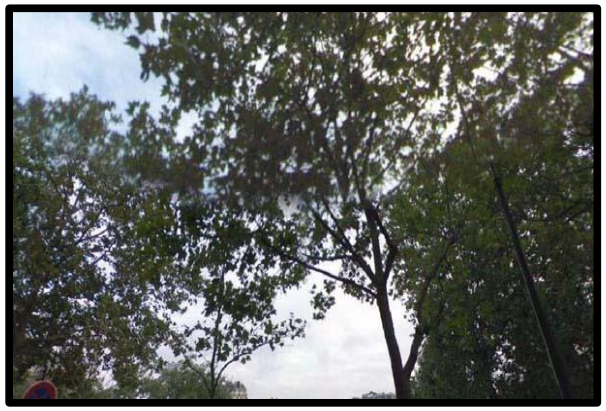
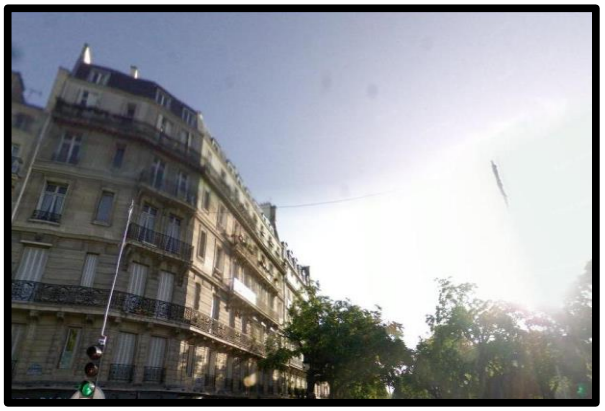


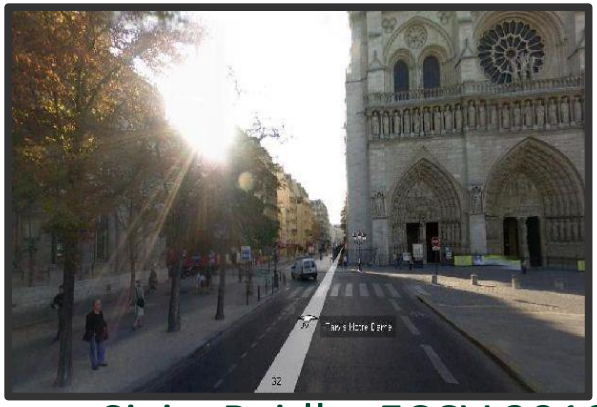

Knopp, Sivic, Pajdla, ECCV 2010

Paris, according to me





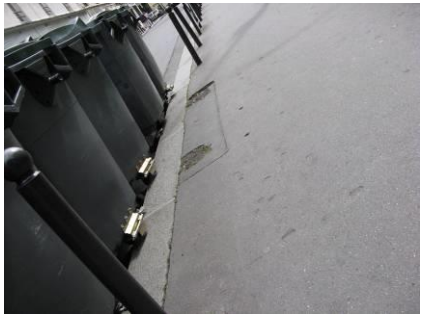




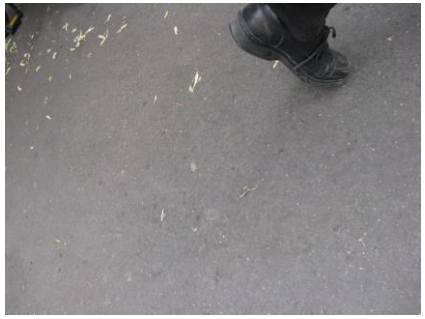

After realism…

Monet,
La rue Montorgueil

Depicting Our World: Ongoing Quest

Pablo Picasso

David Hockney

Better than realism?

David Hockney, Place Furstenberg (1985)

Which one is right?

Multiple viewpoints

David Hockney,
Place Furstenberg,
1985

Single viewpoint

Alyosha Efros
Place Furstenberg,
2009

Depicting Our World: Ongoing Quest

Antonio Torralba & Aude Oliva (2002)

Enter Computer Graphics...

Traditional Computer Graphics

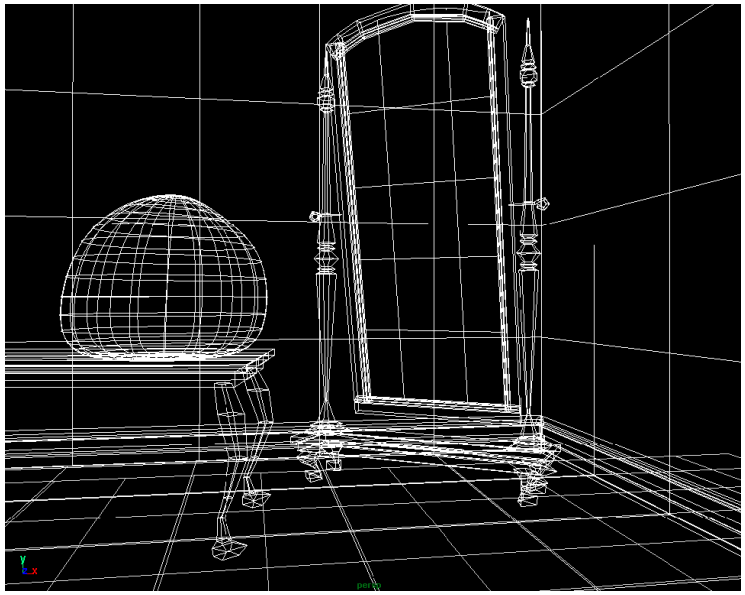

3D geometry

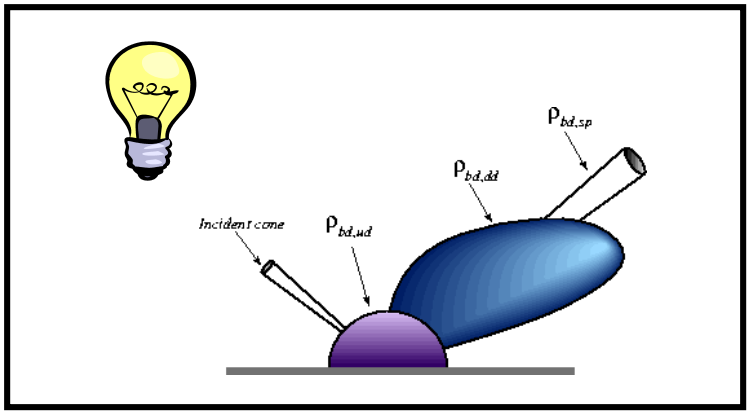

physics

projection

Simulation

GRAPHICS

Modern Computer Graphics

- Amazingly real
- But so sterile, lifeless, *futuristic (why?)*

The richness of our everyday world

Photo by Svetlana Lazebnik

Beauty in complexity

University Parks, Oxford

Which parts are hard to model?

Photo by Svetlana Lazebnik

Creating Realistic Imagery

Computer Graphics

+ great creative possibilities

+ easy to manipulate objects/viewpoint

-Tremendous expertise and effort to obtain realism

Computational Photography

Realism
Manipulation
Ease of capture

Photography

+ instantly realistic

+ easy to aquire

- very hard to manipulate objects/viewpoint

Pop Quiz!

- Don’t worry, this time, we grade only on participation
- Quiz designed to diagnose your level of preparation for CS180
- Should take about 10-15 minutes
  - If you are DSP, you can take up to 30 minutes

https://forms.gle/ByV8zdW1UbS8wDgH7