Computational Photography
CS445

Instructor: Yuxiong Wang

Some Slides adopted from Derek Hoiem

People - Instructor

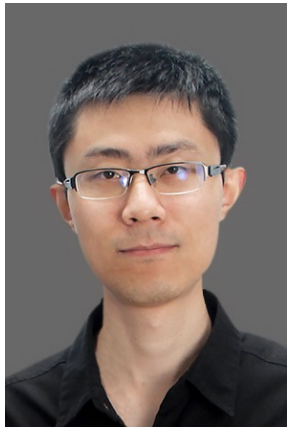

Yuxiong Wang

- Ph.D. & Postdoctoral Fellow
  - CMU, Robotics Institute
  - Thesis: Learning to learn for small sample visual recognition
- Assistant Professor, 2020-
  - UIUC, CS
- Research Interests
  - Meta-Learning, Computer Vision, Robotics

Hiring! Ph.D., MS., Undergraduates

People - Teaching Assistant

Kai Yan

- Intro
  - 3rd year PhD student
  - Co-advised by Prof. Yu-Xiong Wang and Prof. Alexander Schwing
  - 1st time as a TA for CS445
- Research Topic – Reinforcement / Imitation Learning
  - Learn to make decisions by historical data and trial-and-error
- What's fun about CS445?
  - Understand and DIY your PhotoShop/FaceApp!

People - Teaching Assistant

Hang Yu

- Intro
- 2nd year MSCS student
- 1st time as CS445 TA
- Love freediving

- Research Topic - Long-Tailed Recognition
- solve the hideous data imbalance issue

- What’s fun about CS445?
- Understanding what’s happening in Photoshop and Premiere Pro
- Know your camera better!

People - Teaching Assistant

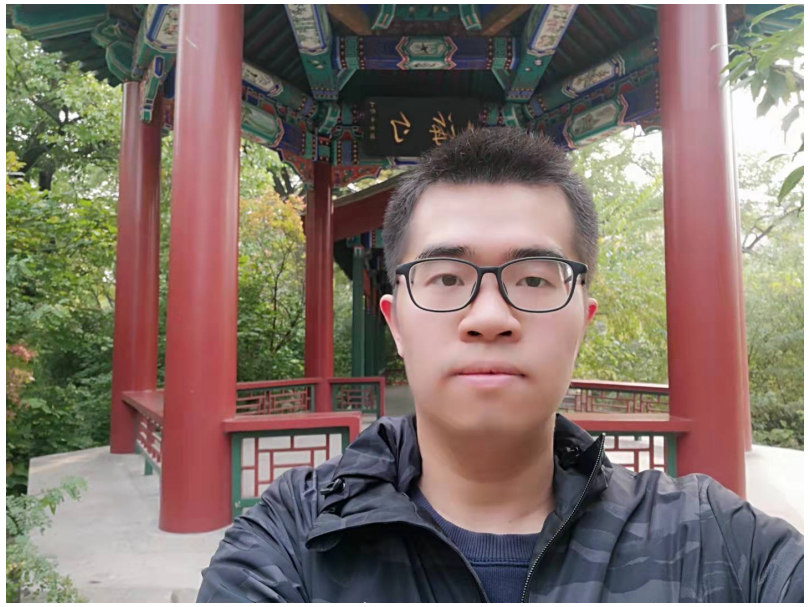

Shuhong Zheng

- Intro
  - 2nd year MS student
  - Advised by Prof. Yu-Xiong Wang
  - 2nd time as a TA for CS445
- Research Topic – Bridging Discriminative and Generative Learning
  - Generative models like NeRF and diffusion models to perform visual perception tasks
- What's fun About CS445?
  - The experiences of handling with your own images! (e.g., Generate fakes and get rid of something on the images that you dislike!)

Today’s Class

- A little about my research
- Intro to Computational Photography
- Course outline and logistics

About My Research

Pottopod

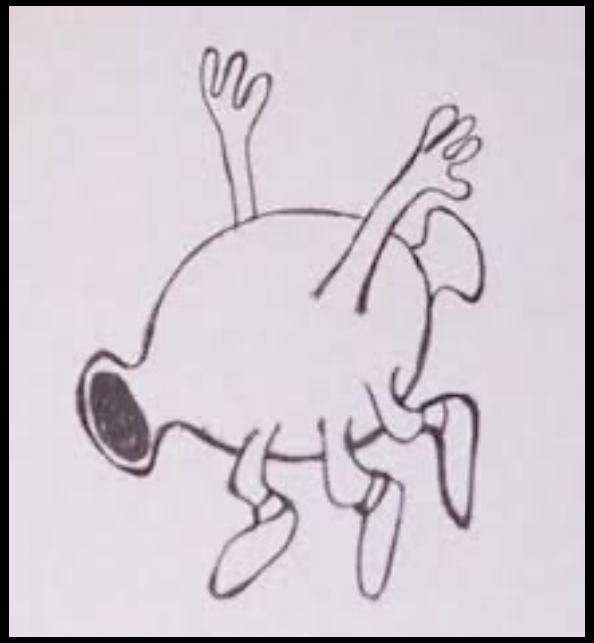

Image credit: S. Savarese, P. Perona

Find the Pottopod

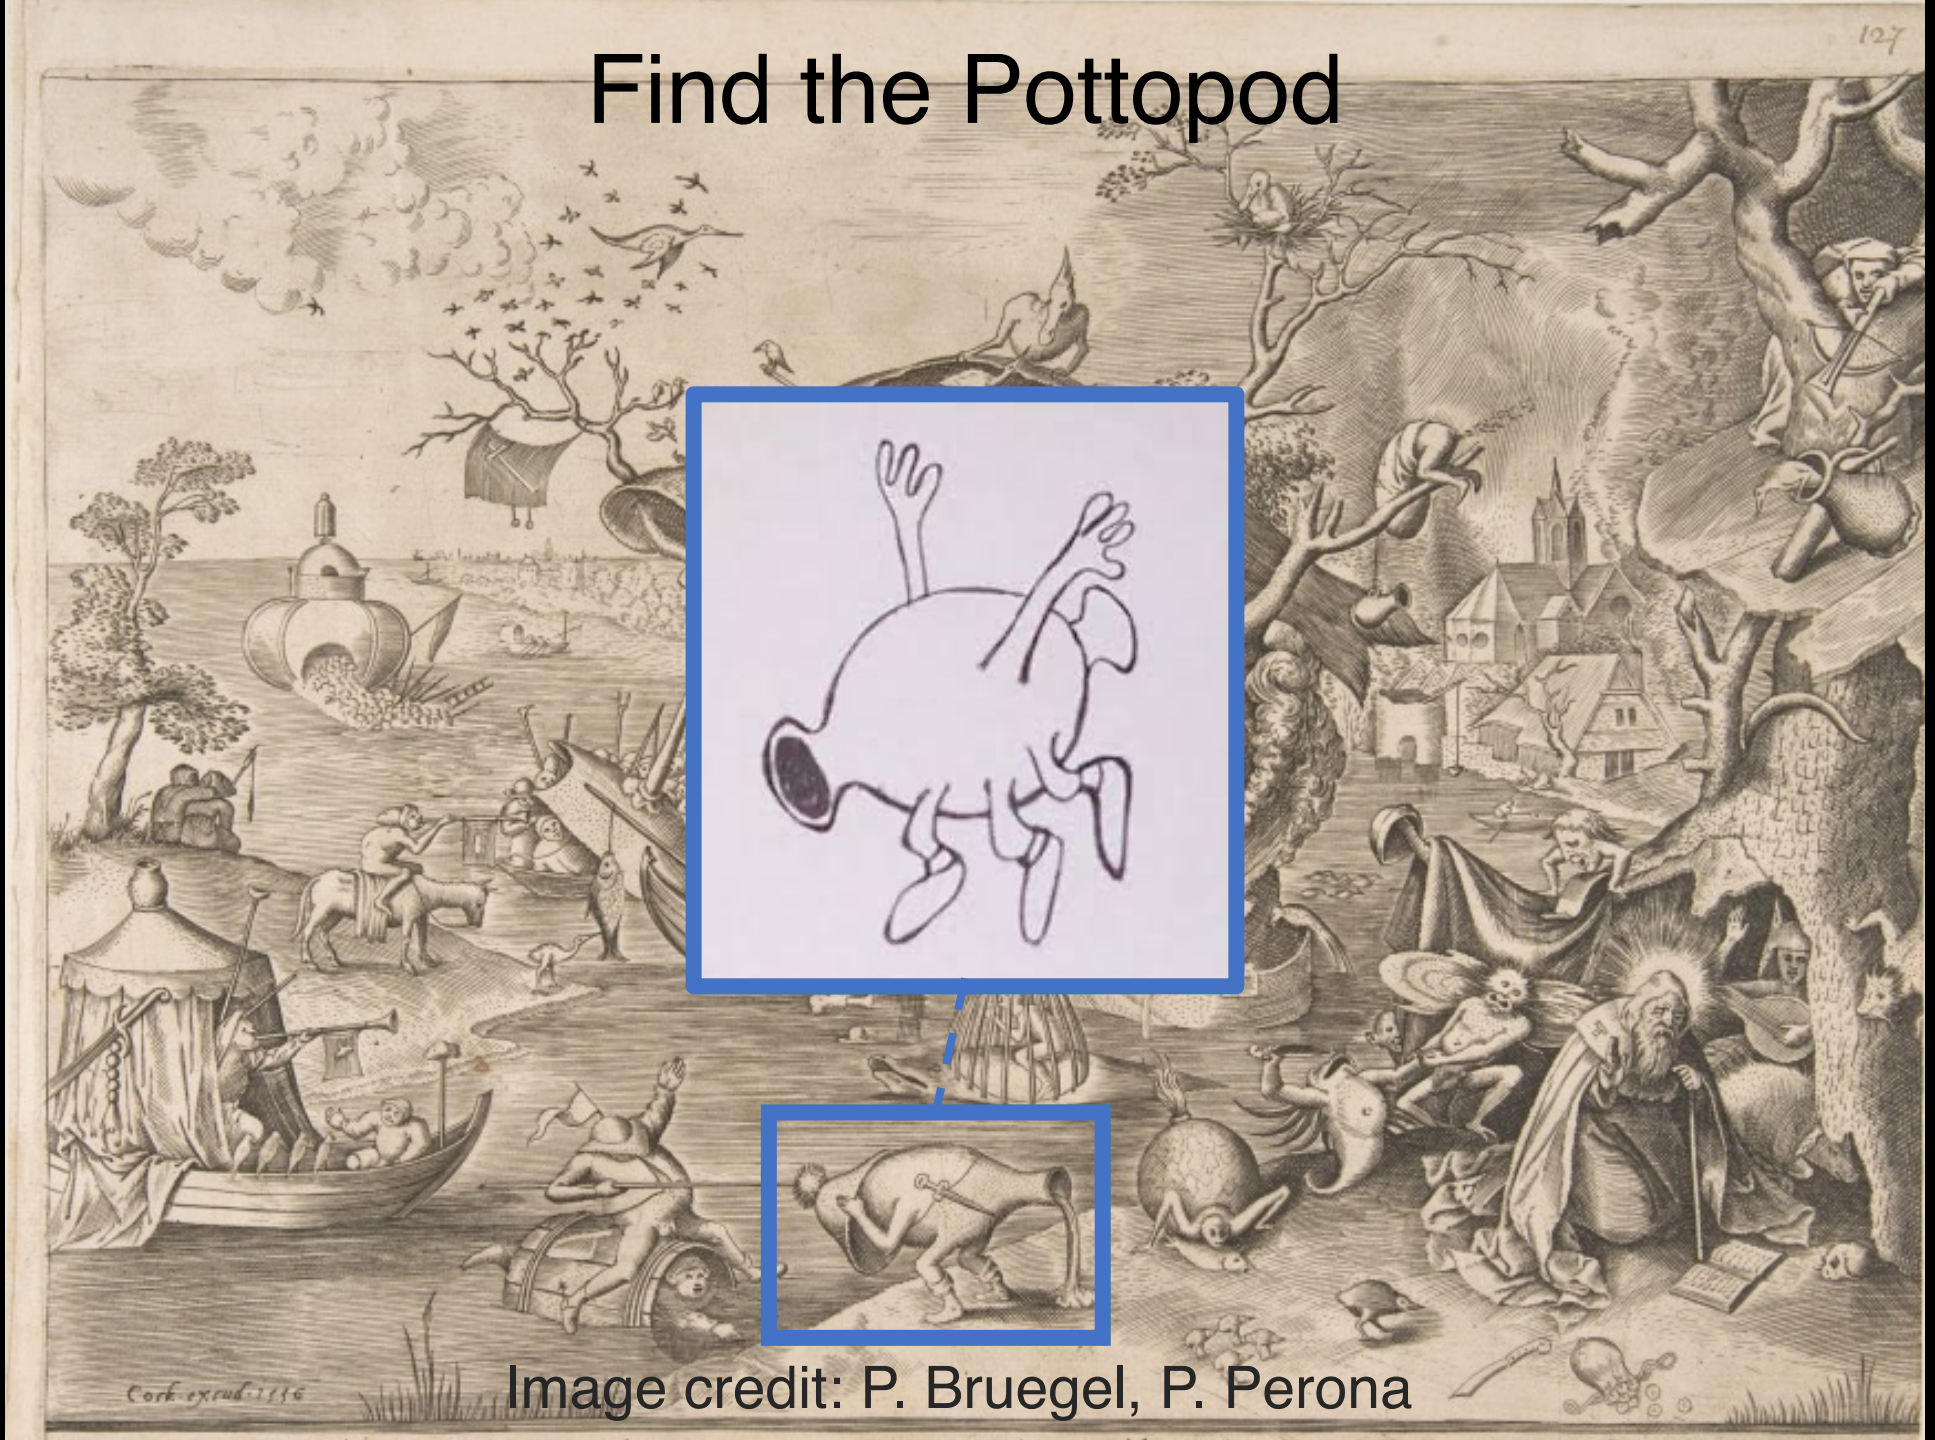

Image credit: P. Bruegel, P. Perona

Machine Recognition

supervised learning

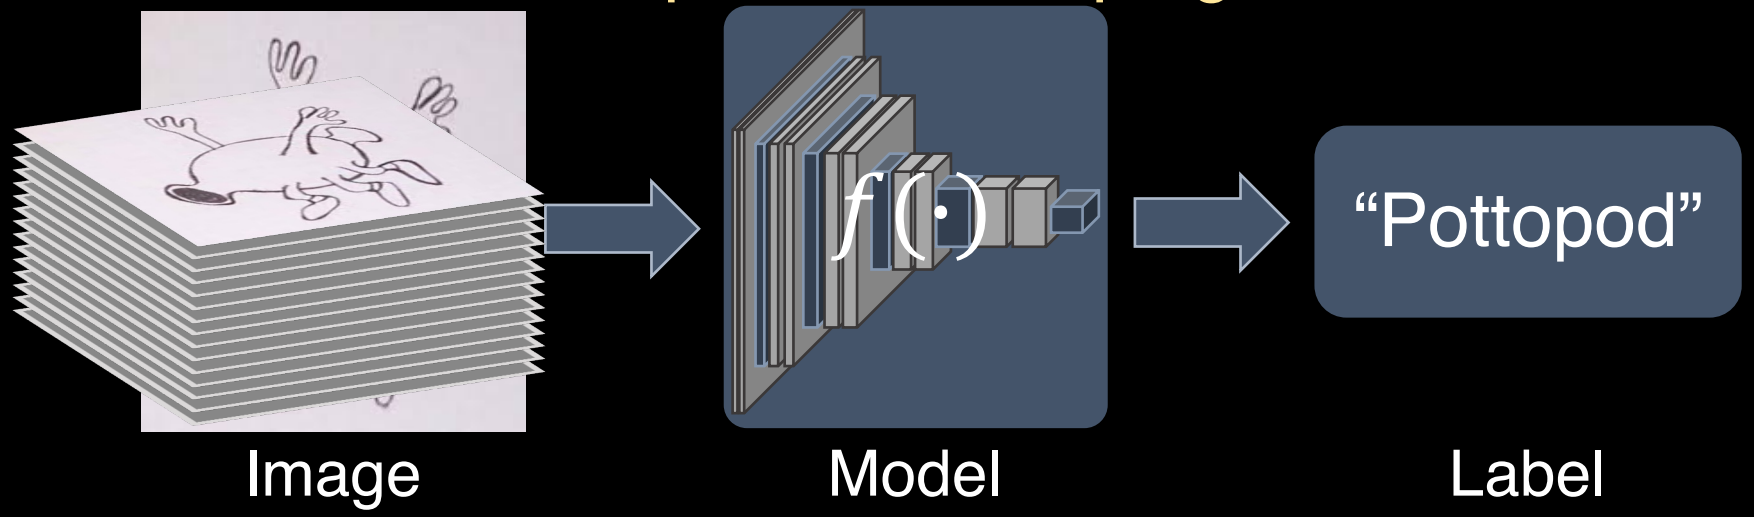

Why Not Simply Collect More Data?

- The real world is long-tailed
  - Unscalable: Recollect huge data for each new
  - Expensive
  - Difficult

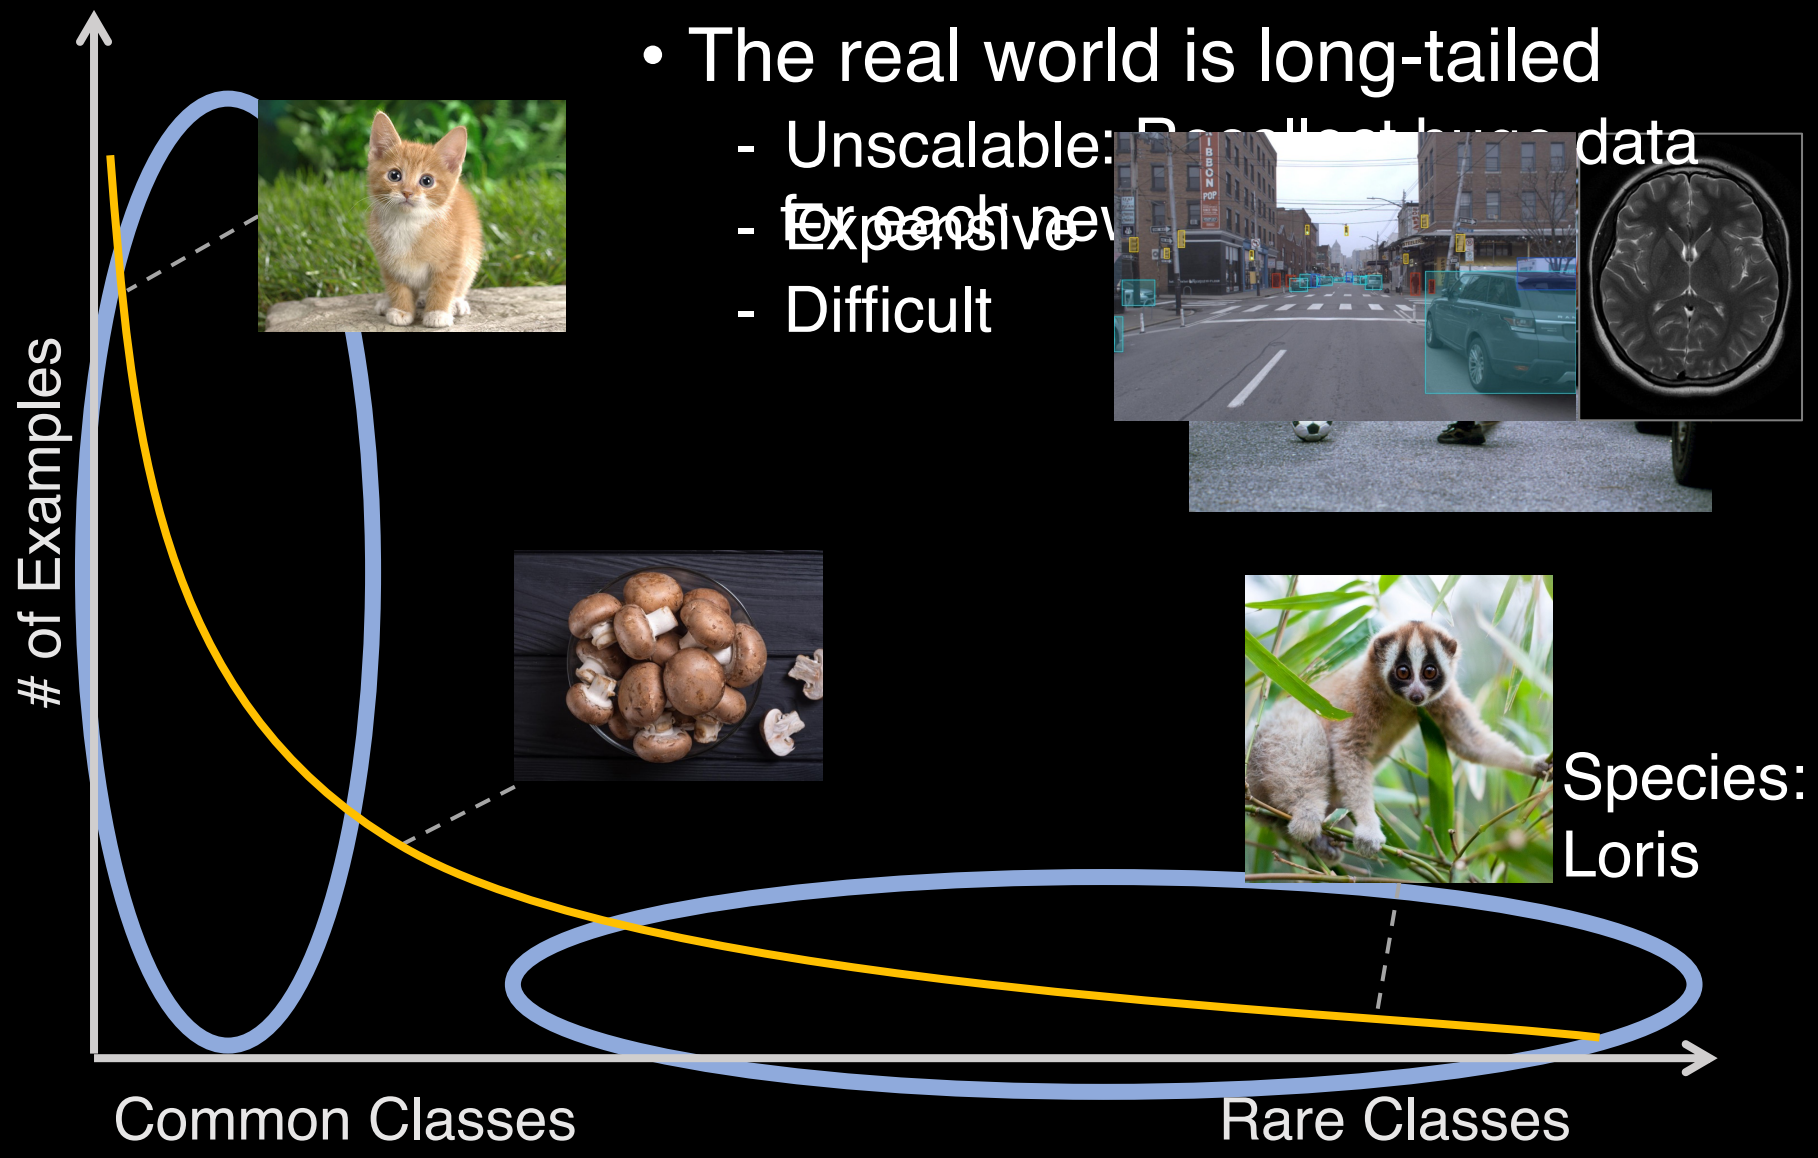

Our Problem: Few-Shot Learning

Few-Shot Loris Classification

Novel (Rare) Class

f_{Loris}

Model

One-shot
(single image)

Miller, 2000. Fink, 2005. Bart and Ullman, 2005. Fei-Fei et al., 2006.

Key Insight: Meta-Learning

Learning how to learn on
Base (Common) Classes

Learning from
Experience

Few-Shot Loris Classification

Novel (Rare) Class

f_{Loris}

Model

One-shot
(single image)

Address Rich Space of Few-Shot Tasks

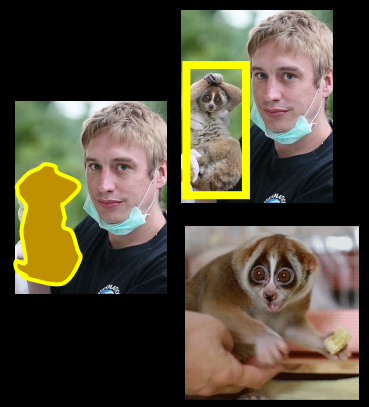

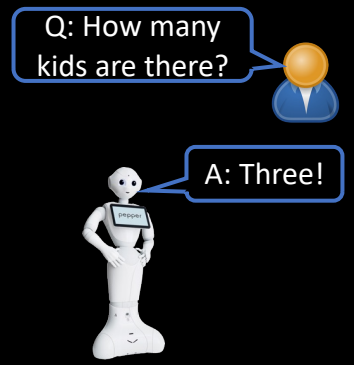

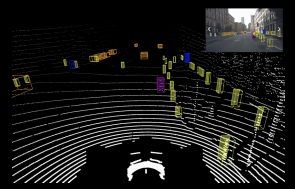

Beyond: NLP, RL…

Task Complexity

Recognition → Synthesis → Interaction → Other Modality

My Work:
A Principled Meta-Learning Framework

- Address *rich* space of few-shot tasks
- Explore *various* experience & knowledge

Learning to Generate Data

[CVPR18] [CVPR19]
[MM19] [ICLR21]

Learning to Transform Model

[ECCV16] [NeurIPS17]
[ECCV18] [ICCV19]

Unsupervised Learning

[CVPR15] [AAAI16]
[NeurIPS16] [ICCV17]

Compositional Learning

[ICCV19]

Experience

Developmental Learning

[CVPR17]

Predictive Learning

[ECCV18] [ECCV18]
[IROS18]

Some background to computational photography and …

The Pursuit of Realism

Several of following slides from Alyosha Efros

Depicting Our World: The Beginning

Prehistoric Painting, Lascaux Cave, France
~ 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

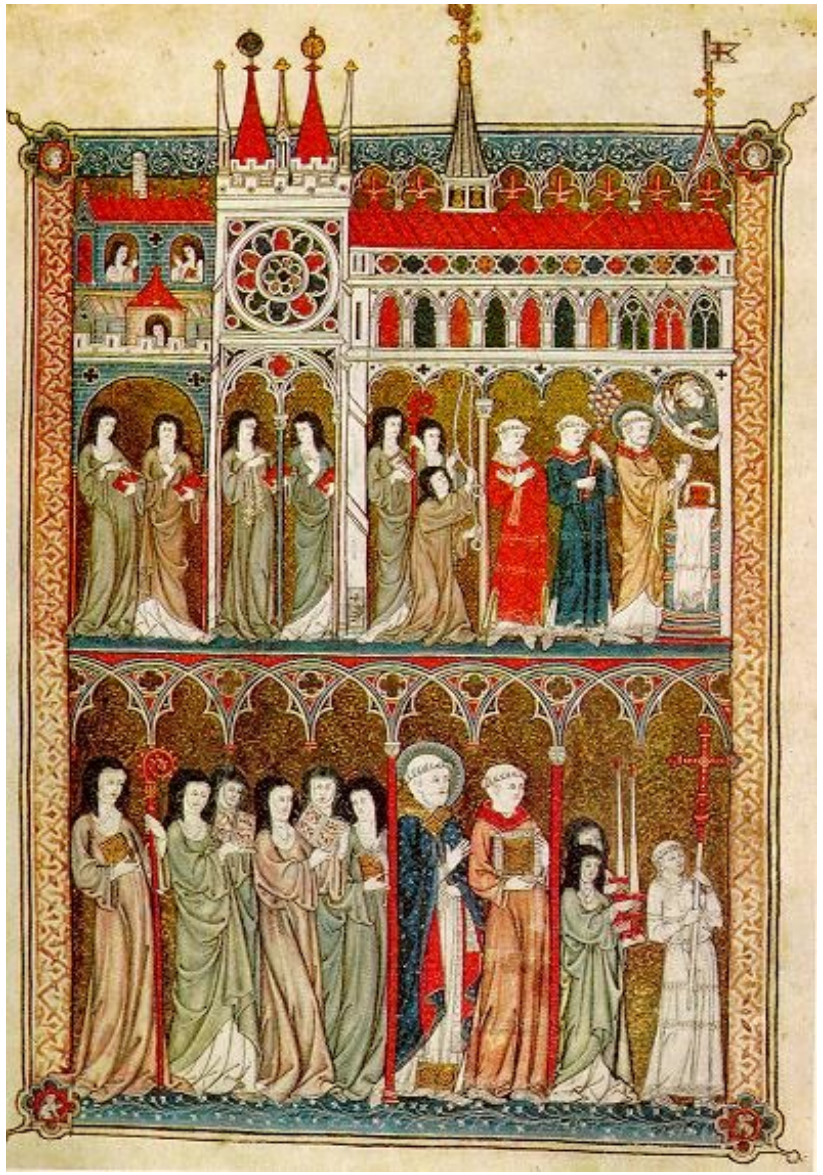

Nuns in Procession. French ms. ca. 1300.

Depicting Our World: Renaissance

North Doors (1424)

Lorenzo Ghiberti (1378-1455)

East Doors (1452)

Depicting Our World: Renaissance

Paolo Uccello,
Miracle of the Profaned Host (c.1467-9)

Depicting Our World: Toward Perfection

Jan van Eyck, *The Arnolfini Portrait (1426-1434)*

Depicting Our World: Toward Perfection

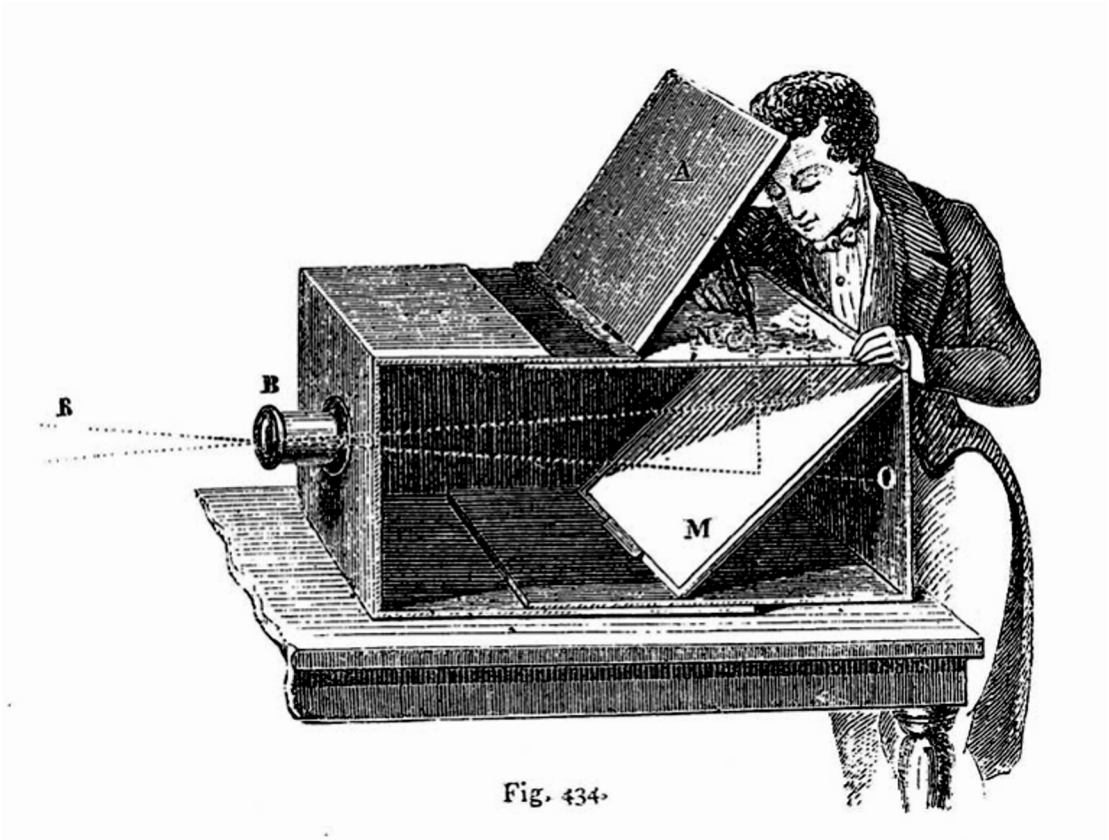

Lens Based Camera Obscura, 1568

Depicting Our World: Perfection!

Still Life, Louis Jaques Mande Daguerre, 1837

Is reality what we want?

http://www.freephotoediting.com/samples/photo-glamourisation/020_celebrity-skin-retouching.htm

Enter Computer Graphics...

Traditional Computer Graphics

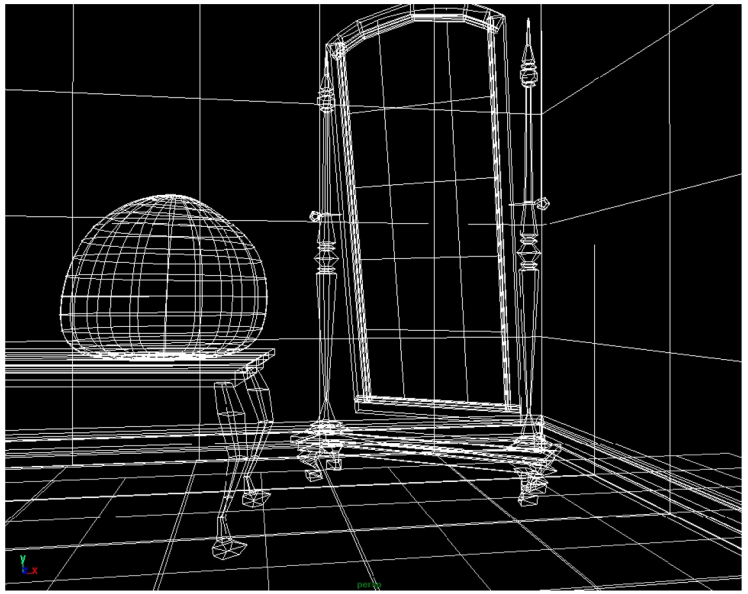

3D geometry

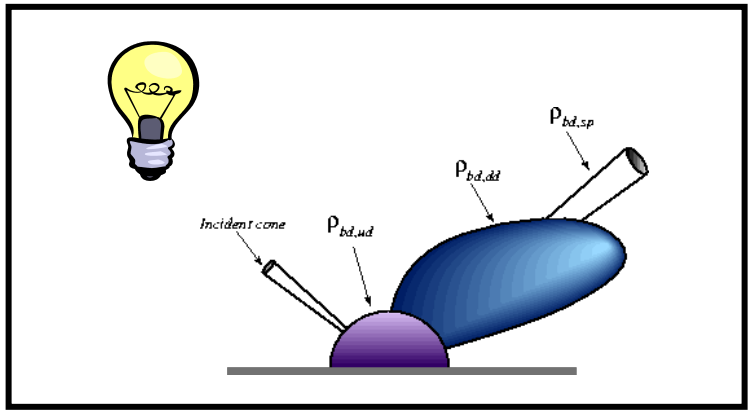

physics

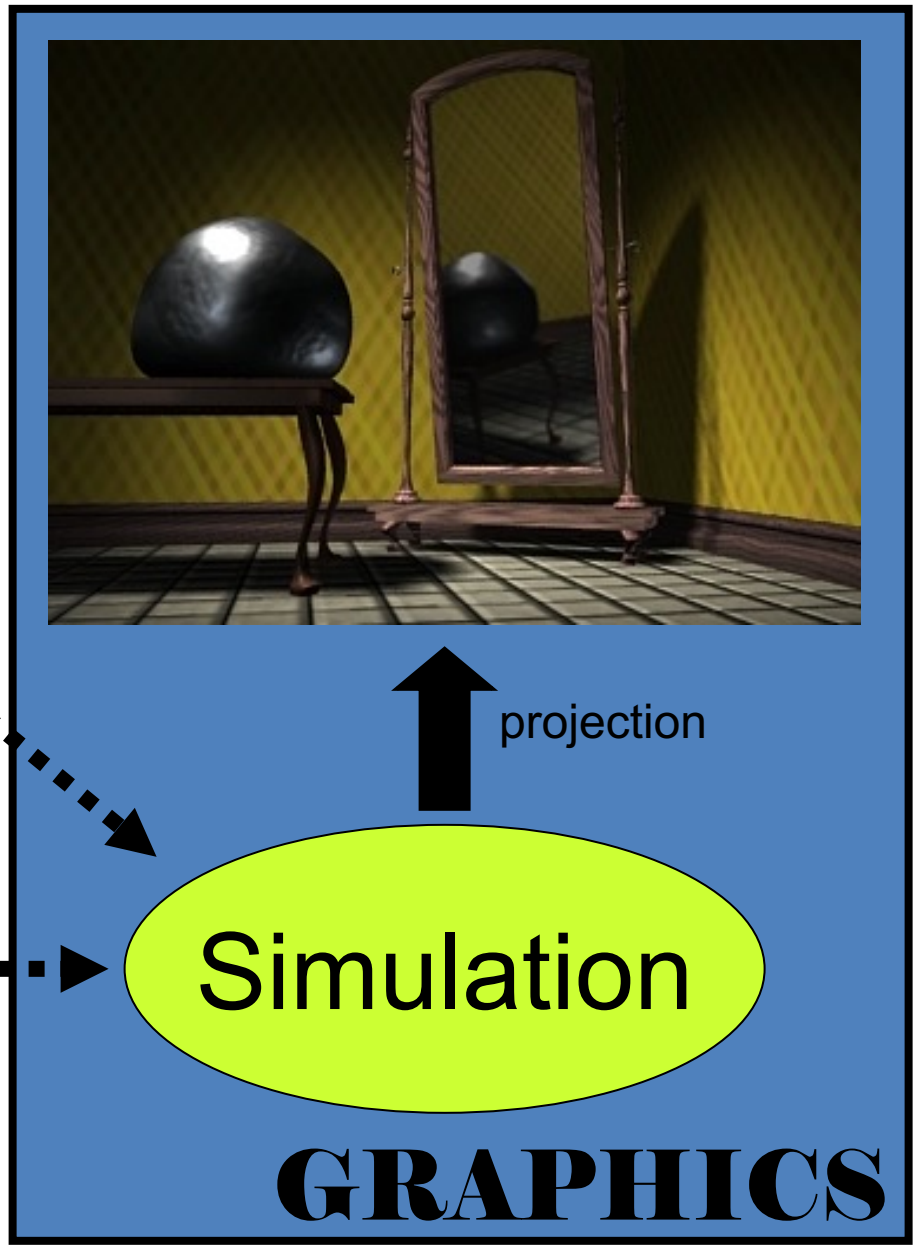

Computer graphics

What’s wrong?

The richness of our everyday world

Photo by Svetlana Lazebnik

Which parts are hard to model?

Photo by Svetlana Lazebnik

People

From “Final Fantasy”

Alyosha Efros - On the Tube, London

Faces / Hair

From “Final Fantasy”

Photo by Joaquin Rosales Gomez

Urban Scenes

Virtual LA (SGI)

Photo of I LA

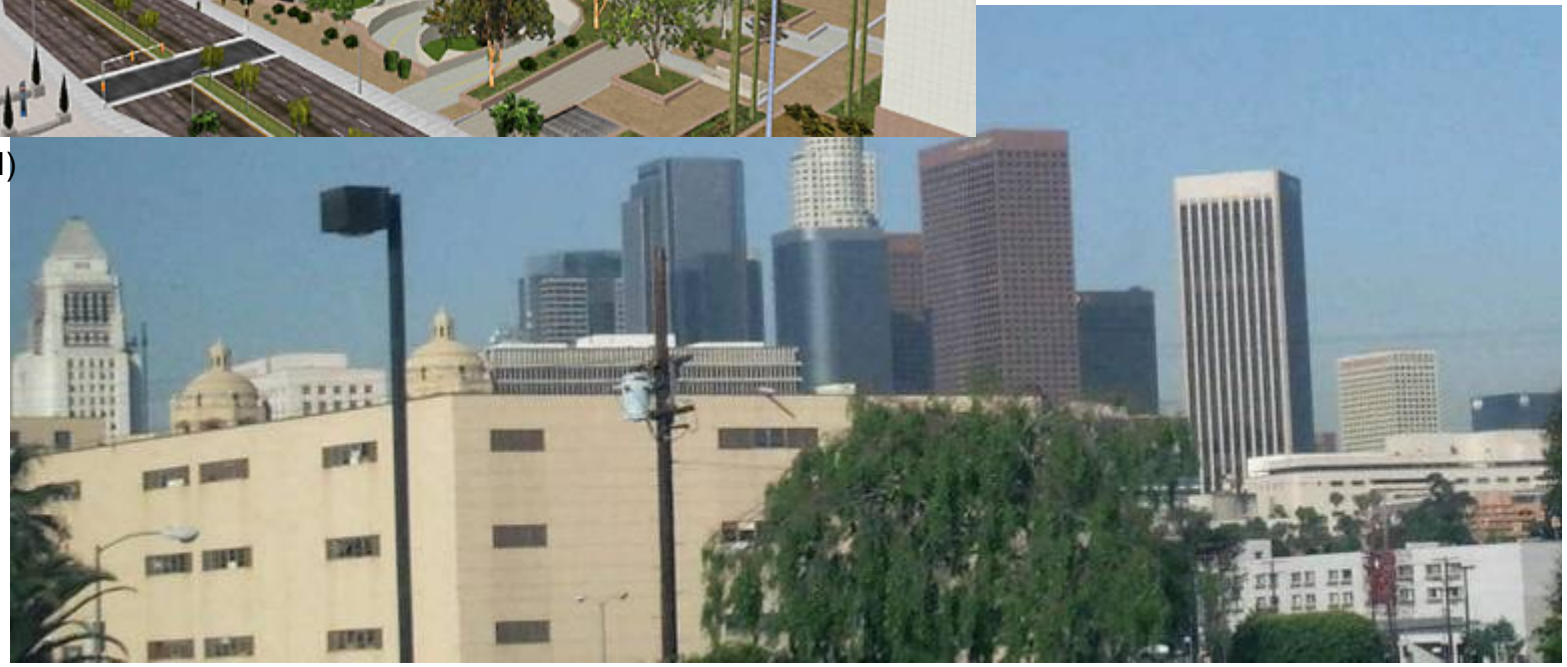

Nature

River Cherwell, Oxford

The Realism Spectrum

Computer Graphics

+ easy to create new worlds

+ easy to manipulate objects/viewpoint

- very hard to look realistic

Computational Photography

Realism
Manipulation
Ease of capture

Photography

+ instantly realistic

+ easy to aquire

- very hard to manipulate objects/viewpoint

Computational Photography

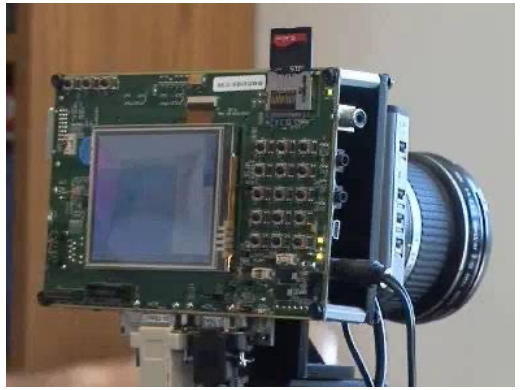

How can I use computational techniques to capture light in new ways?

How can I use computational techniques to breathe new life into the photograph?

How can I use computational techniques to synthesize and organize photo collections?

Virtual Real World

Campanile Movie (1997)

http://www.debevec.org/Campanile/

Going beyond reality…

Benjamin Button (2008)

http://www.digitaldomain.com/work/the-curious-case-of-benjamin-button/
https://vimeo.com/72240367

Another example of blending reality with fantasy

Samsung Galaxy S6 regular and “beauty” selfie

FaceApp

Course objectives

1. You will have new abilities for visual creation.

Graphic by James Hays

Course objectives

2. You will get a foundation in computer vision.

Safety

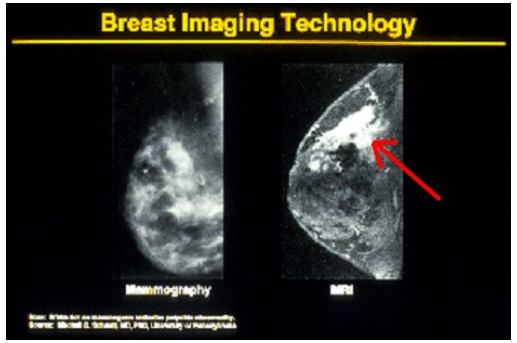

Health

Security

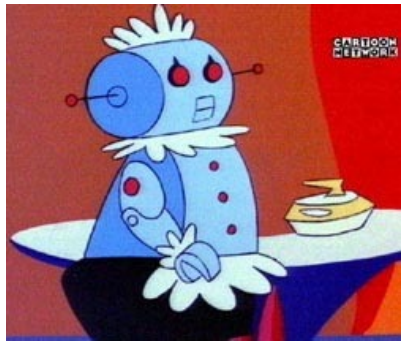

Comfort

Fun

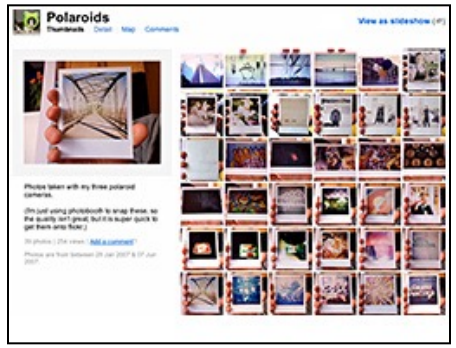

Access

Got job?

- Google, Facebook, Microsoft, Sony, iRobot, Amazon, Snapchat, Ebay, tons of startups, etc.

- http://www.cs.ubc.ca/~lowe/vision.html

Course objectives

3. You'll better appreciate your own visual ability.

Course objectives

4. You’ll have fun doing cool stuff!

Projects

Project 1: Hybrid Images

Project 2: Image Quilting for Texture Synthesis and Transfer

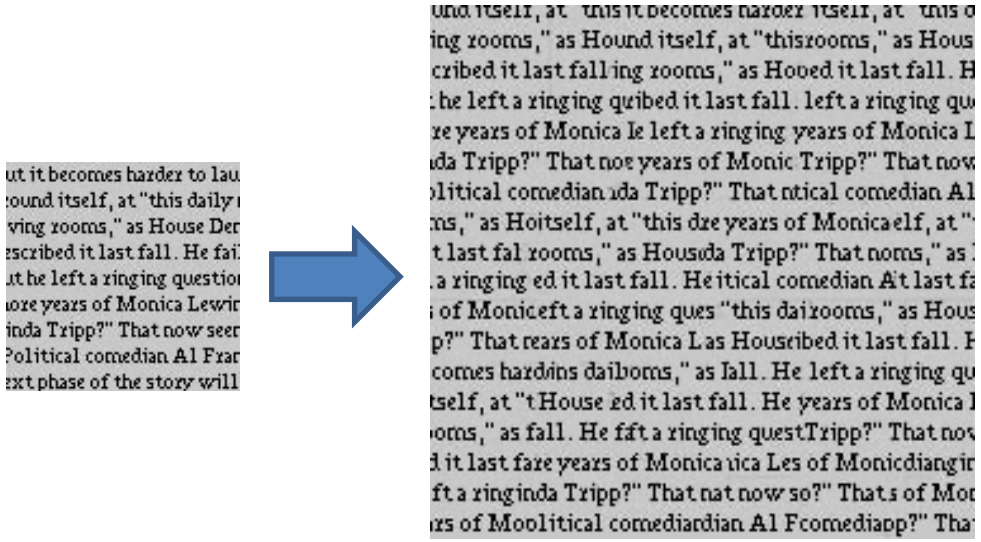

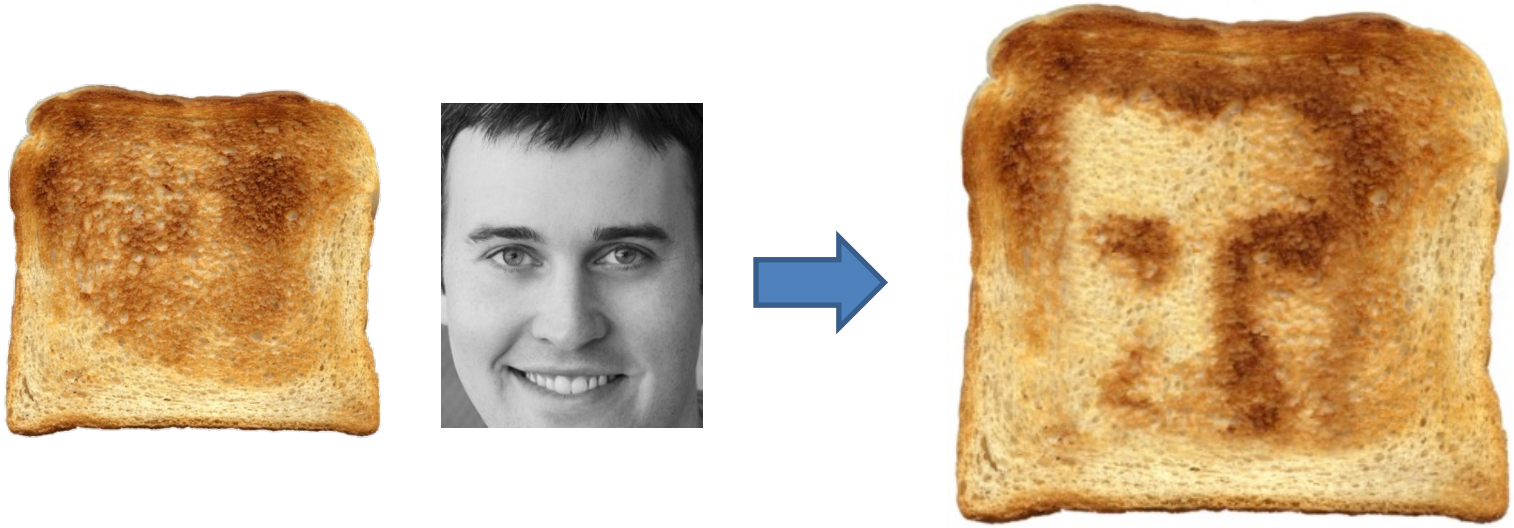

Project 3: Gradient-Domain Fusion (Poisson Editing)

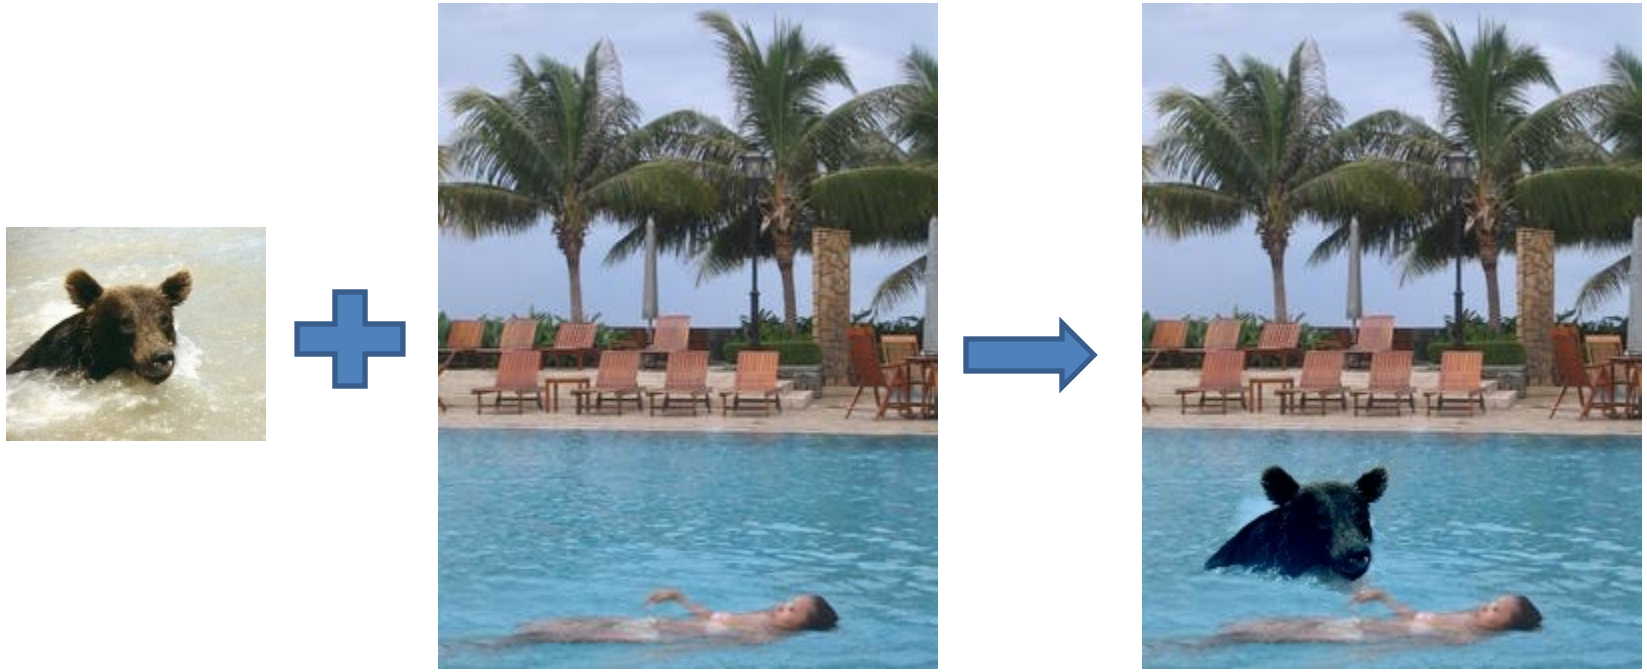

Photos from James Hays

Project 3: Gradient-Domain Fusion (Poisson Editing)

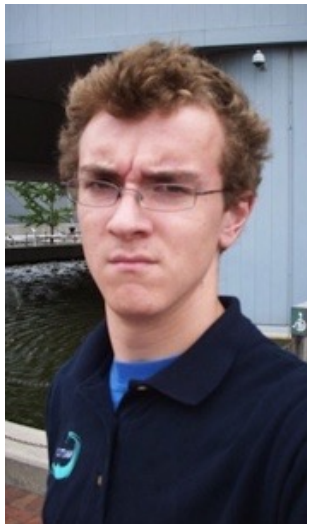

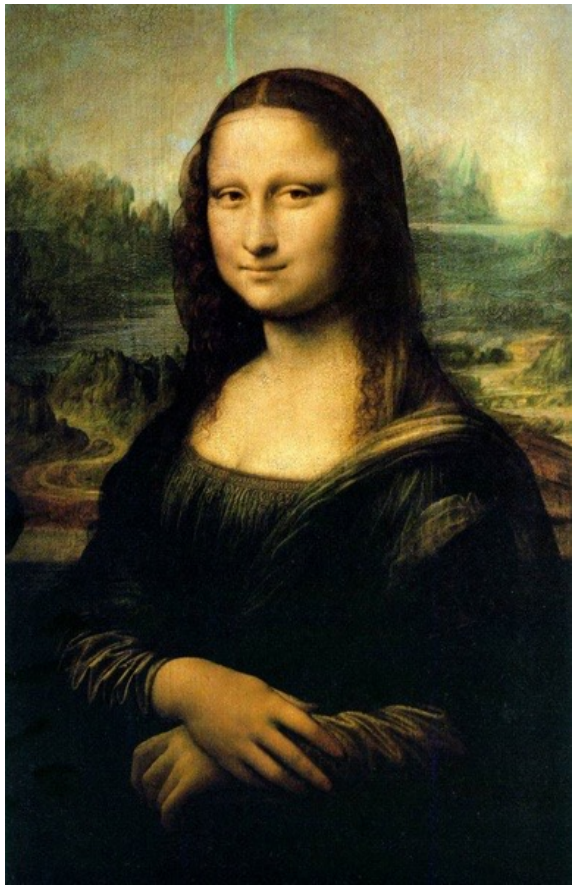

Photos from Evan Wallace

Project 4: Image-Based Lighting

Project 5: Video alignment, stitching, and editing

Final Project

Something cool!

Course outline

Prof: Yuxiong Wang yxw@illinois.edu

TAs

- Kai Yan, kaiyan3@illinois.edu
- Hang Yu, hangy6@illinois.edu
- Shuhong Zheng, szheng36@illinois.edu

Learning resources

Lectures

- Lecture modules on Coursera
- Recorded lectures: mainly from Prof. Derek Hoiem

Website

- https://yxw.cs.illinois.edu/course/CS445/S24/

Slides

- On Coursera & Website

Office hours

- On Coursera & Website

Campuswire: for announcements & discussion

- https://campuswire.com/p/G5A5FBEF7
- Code: 5425

Readings/textbook: for depth and details not covered in lecture

Grades

- Projects (55%)
  - 5 projects: each with 100 core points with more optional “bells and whistles”
  - 3 credit (ugrad): graded out of 425 points
  - 4 credit (grad/MCS): graded out of 500 points
- Exams (30%)
  - Midterm 15%: covers first half
  - Final 15%: covers entire semester
- Final Project (15%)
  - 1% for proposal, 14% for final submission
  - Work in a group of two or up to four
  - 2-4 page short report

Late policy

- Up to ten free days total – use them wisely!
- 5 point penalty per day after that
- Project must be submitted within two weeks of due date to receive any points

Project details

- Implement stuff from scratch and apply it to your own photos

- Submit PDF, Jupyter notebook, and Python code

Academic Integrity

These are OK

- Discuss projects with classmates (don’t show each other code)
- Use Stack Overflow to learn how to use a Python module
- Get images from online (make sure to attribute the source)

Not OK

- Copying or looking at project-specific code (i.e. so that you claim credit for part of an assignment based on code that you didn’t write)
- Using external resources (code, images) without acknowledging them

Remember

- Ask if you’re not sure if it’s ok
- You are safe as long as you acknowledge all of your sources of inspiration, images, code, etc. in your write-up

Other comments

Prerequisites

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

Equipment

- Your own camera, but a smartphone is probably good enough
- A mirrored sphere for project 4 (12 cm or bigger) e.g. https://www.amazon.com/Stainless-Mirror-Polished-Sphere-Ornament/dp/B01ING7L4U

Feedback is welcome