Advanced Computer Vision

Jia-Bin Huang
Electrical and Computer Engineering
Virginia Tech

Today’s class

- A little about me
- A little about you
- Course logistics
- Topic overview

About me

7 months old!

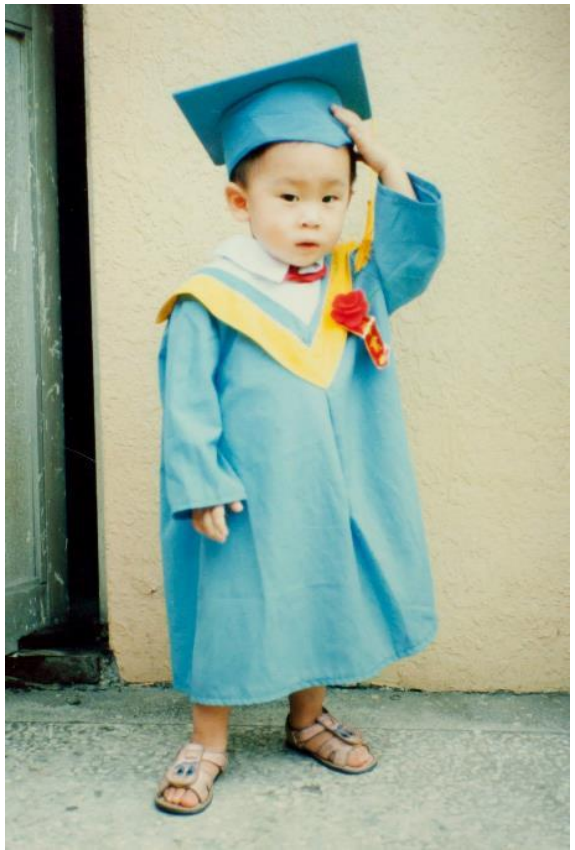

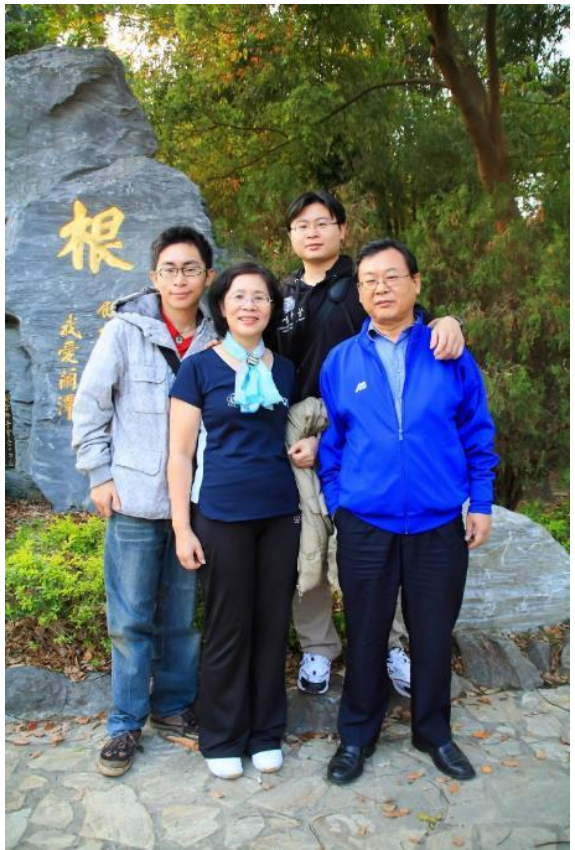

National Chiao-Tung University
B.S. in EE

Microsoft Research
Research Intern
2012, 2013

IIS, Academia Sinica
Research Assistant

Disney Research
Research Intern 2014

UC, Merced
Visiting Student

UIUC
Ph.D. in ECE 2016




Image Completion [SIGGRAPH14]

- Revealing unseen pixels




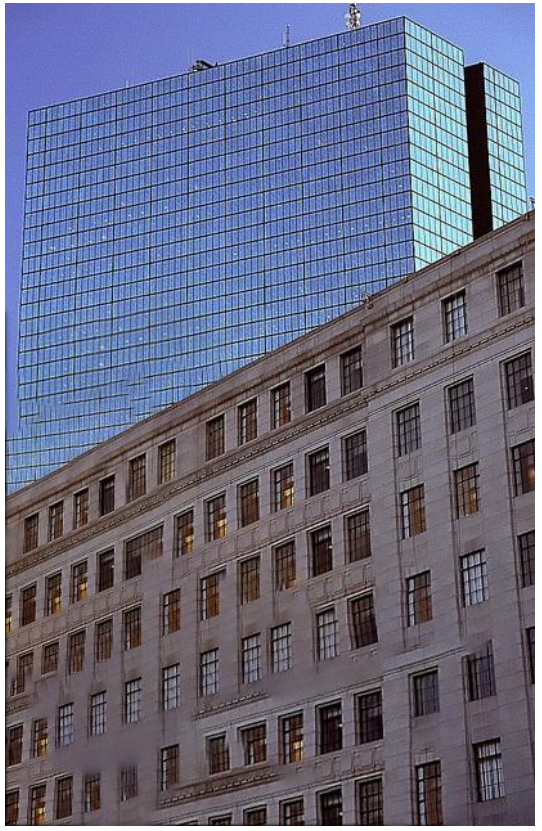



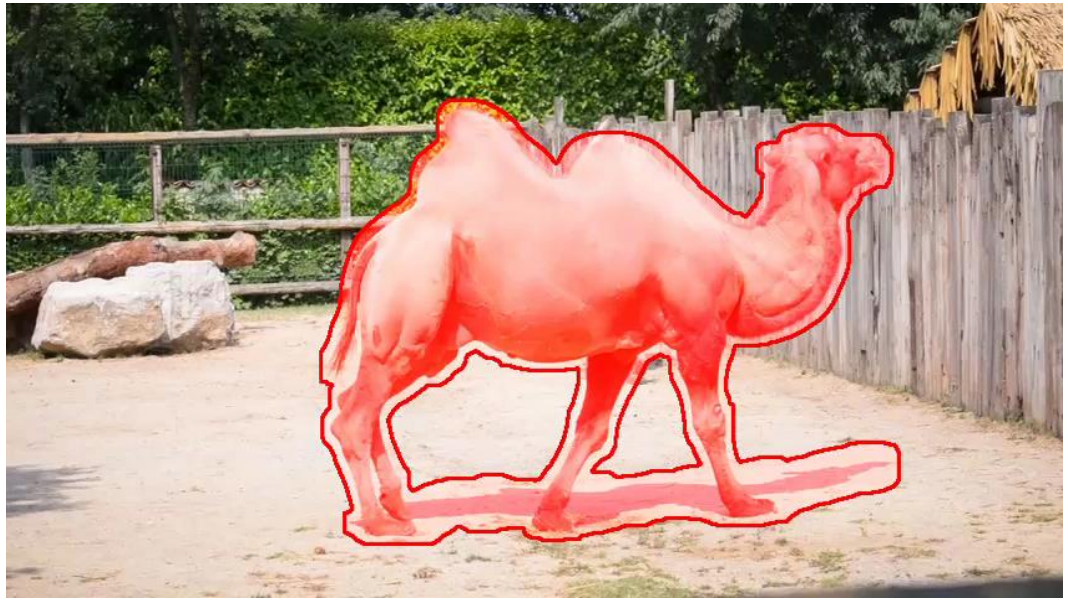


Video Completion [SIGGRAPH Asia16]

- Revealing temporally coherent pixels

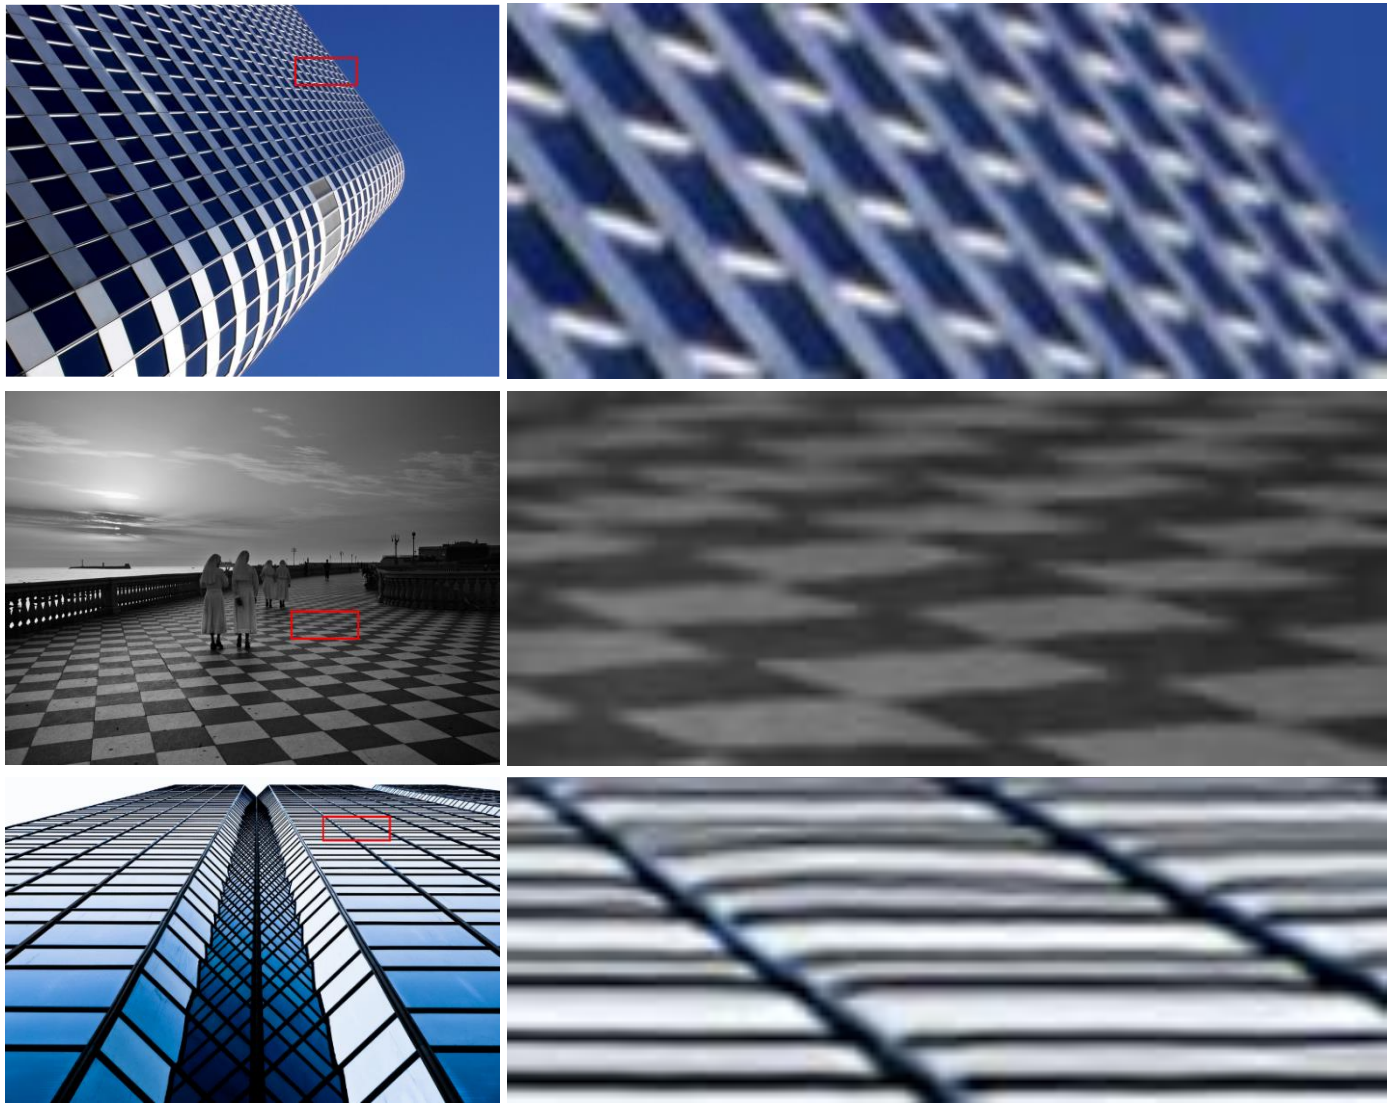

Image super-resolution [CVPR15]

- Revealing unseen high frequency details

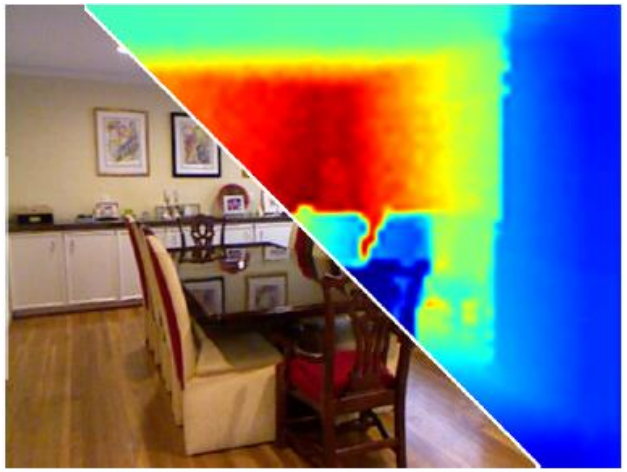



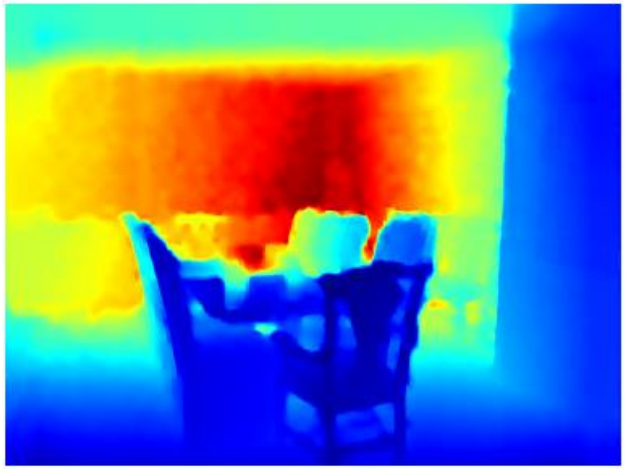


Depth upsampling | Noise reduction | Inverse halftoning | Texture removal

Deep Joint Image Filtering [ECCV16]

- Transferring structural details

Object tracking [ICCV15]

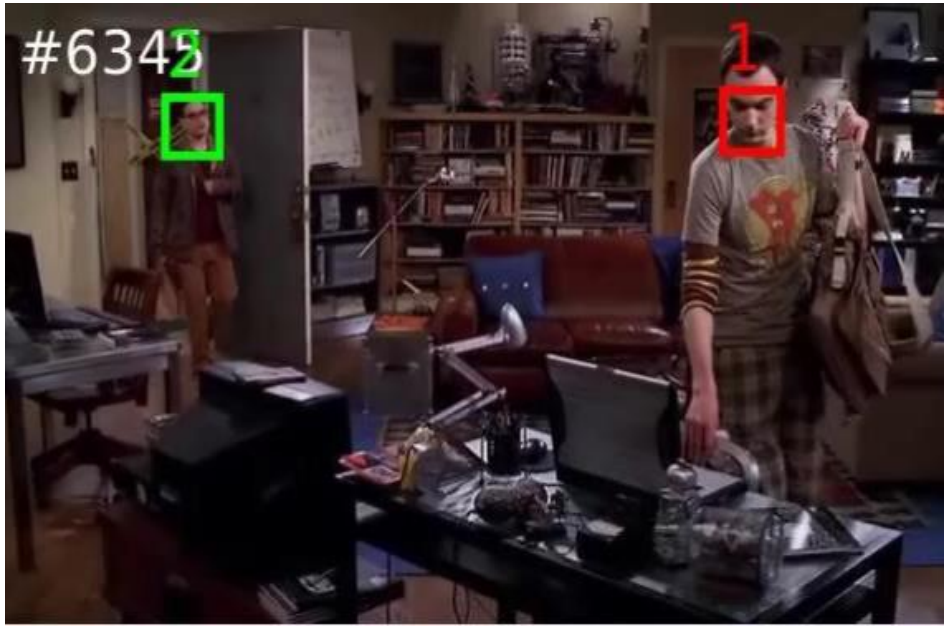

Multi-face tracking [ECCV16]

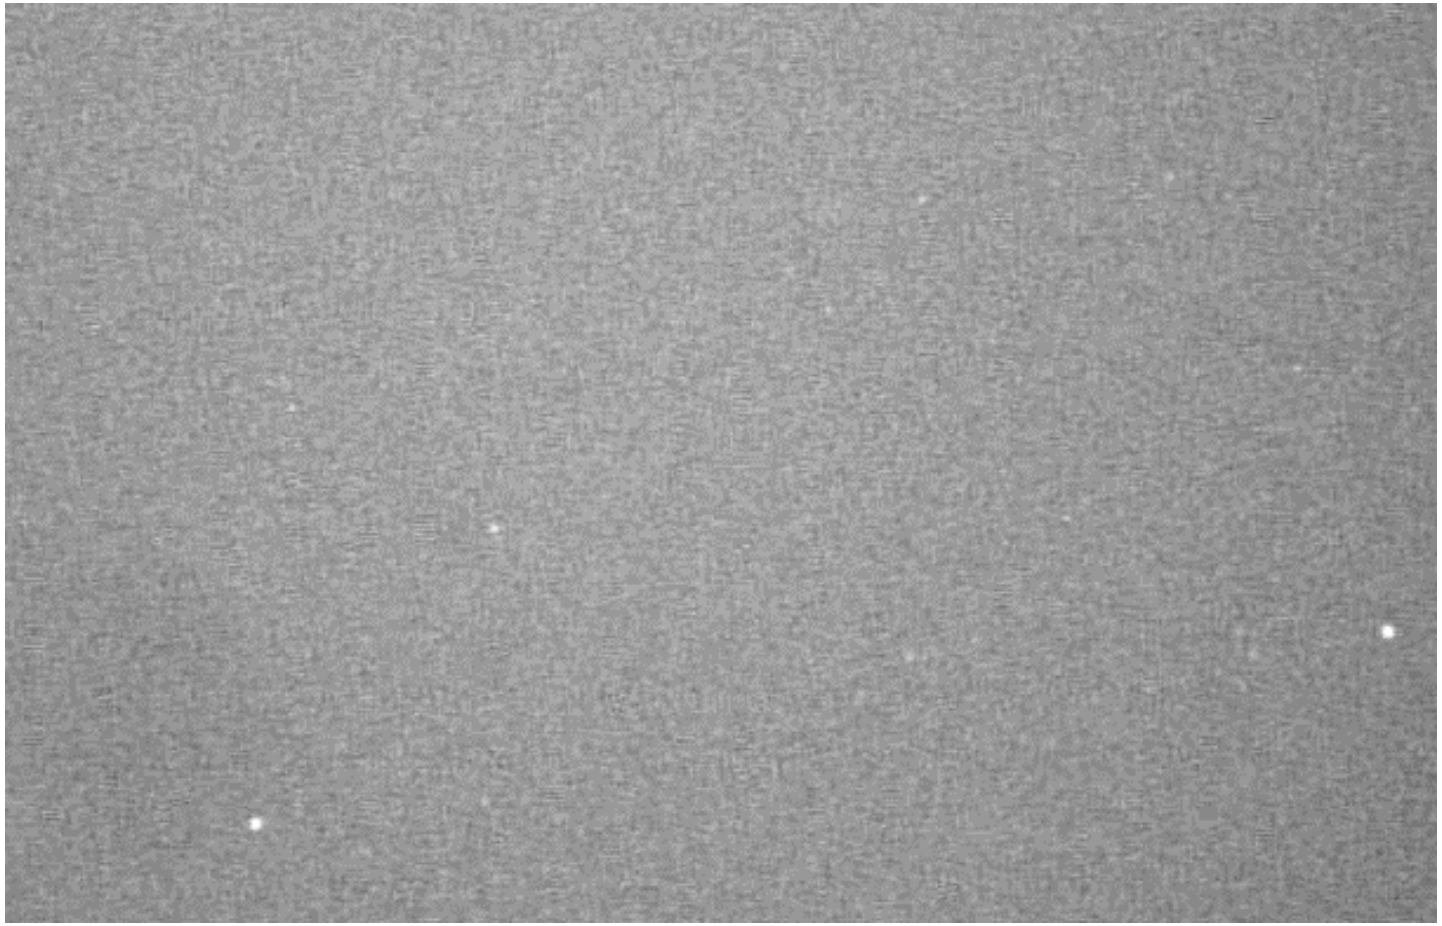

Detecting migrating birds [CVPR16]

Visual Tracking

- Locating moving objects across video frames

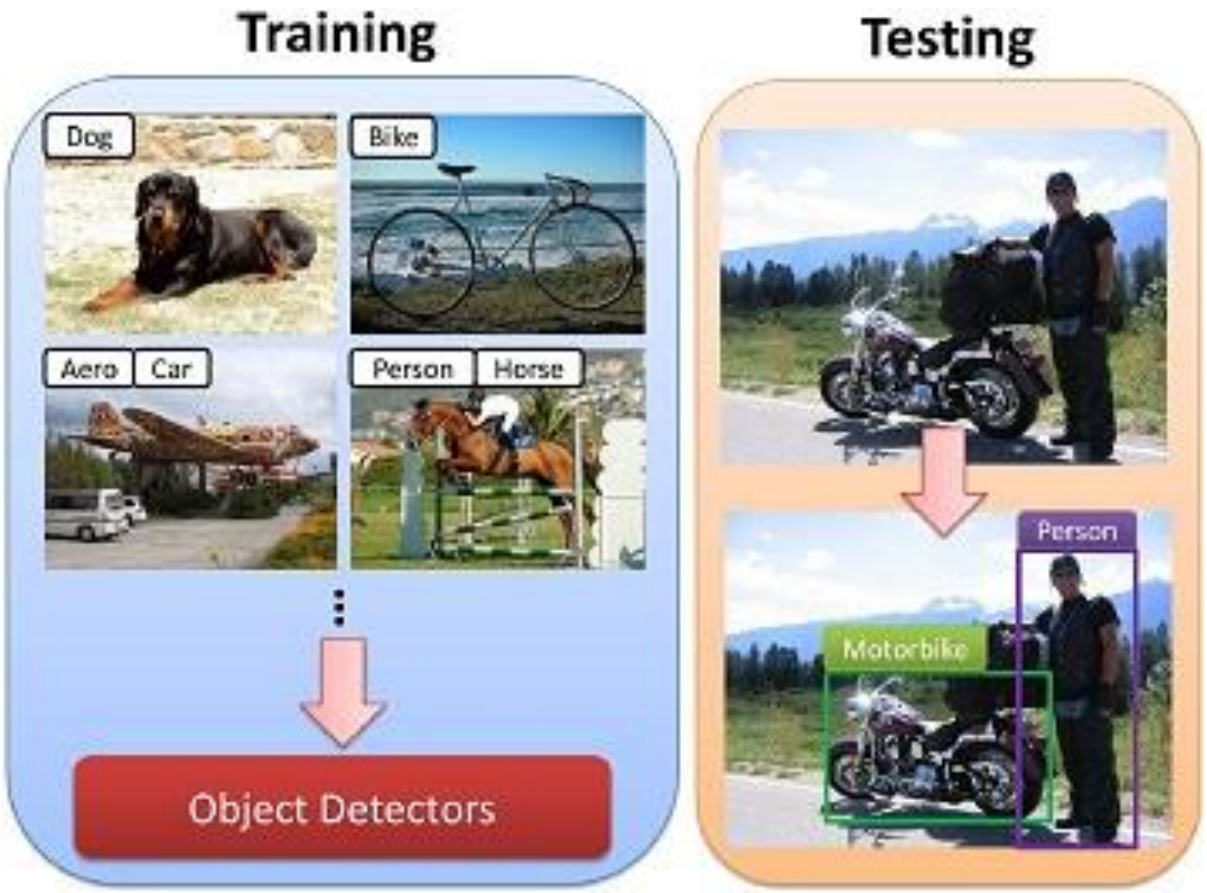

Weakly supervised localization [CVPR16]

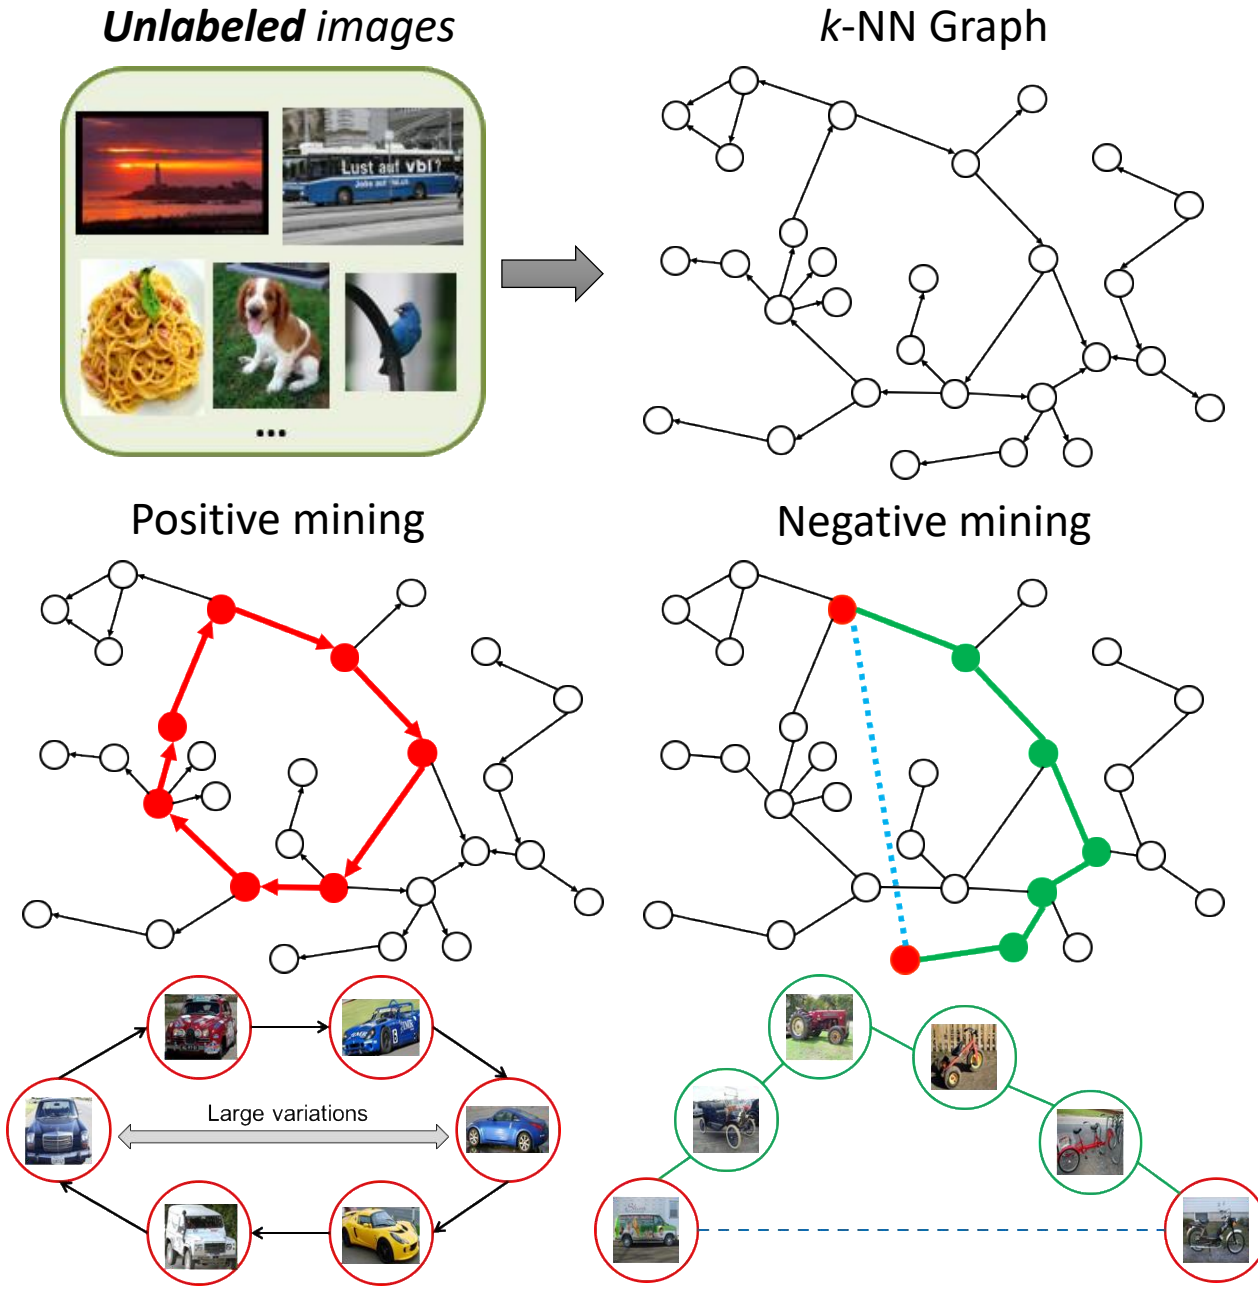

Unsupervised feature learning [ECCV16]

Learning with weak labels

A little about you

- Find a partner.
- Introduce yourself
  - Name?
  - Department?
  - Why are you taking this course?
  - One interesting fact?
- 3 mins

- **Introduce your partner!**

Course Overview

- ECE 6554
  - Tuesday and Thursday 2:00 pm to 3:15 pm
  - Randolph Hall 220

- Office hours (Jia-Bin)
  - TBD (will post a doodle link on piazza)

- Course webpage: http://bit.ly/vt-computer-vision-spring-2017

- Piazza discussion forum: http://piazza.com/vt/spring2017/ece6554

This course

- Focus on more advanced techniques and ideas in computer vision
- Presented in research papers
- High-level recognition problems, innovative applications.

Goals

- Understand state-of-the-art approaches
- Analyze and critique current approaches
- Identify interesting open questions
- Present clearly and methodically

Expectations

- [5%] Class participation
- [25%] Paper reviews
- [10%] Leading discussions
- [20%] Topic presentation
- [10%] Experiment presentation
- [30%] Final project

Requirements – Class participation

- Read the assigned papers before each class
- Actively participate in discussions in class.
- If you are unable to attend a specific class, please let me know ahead of time via email (and have a good excuse).
- No laptops, cell phone or other distractions in class please.

Requirements – Paper reviews

- One page review of the selected paper
- Write in your own words
- Due date
  - 12:00 PM (noon) the day of the class (i.e. on Tuesdays and Thursdays).
- Submission via piazza
- Skip the review if
  - You are presenting a paper on that day
  - You are a discussion lead on that day

Paper reviews – Suggested structure

- Short summary of the paper
- Main contributions
- Strengths and weaknesses?
- Are the experiments convincing?
- How could the work be extended?
- Additional comments, including unclear points, open research questions, and applications.

Requirements – Leading discussions

- ~ One of you will be assigned to argue for the paper
- ~ One of you will be assigned to argue against the paper
- Come prepared with 5 points

Requirements – Topic presentation

- 30 minutes in-class presentation
- Meet with me 3 days prior the talk (i.e. Mon/Friday) with a complete set of slides for a dry run
- IMPORTANT: Don’t present papers – present the topic!

Topic presentation - Structure

- High-level topic overview
- Main motivation
- Clear statement of the problem
- Overview of the technical approach
- Strengths/weaknesses of the approach
- Overview of the experimental evaluation
- Strengths/weaknesses of evaluation
- Discussion: future direction, links to other work

Requirements – Experiment presentation

- 15 minutes in-class presentation
- Send the draft to me 3 days prior your talk
- Implement/download code for a main idea in the paper and evaluate it:
  - Experiment with different types of training/testing data sets
  - Evaluate sensitivity to important parameter settings
  - Show an example to analyze a strength/weakness of the approach
  - Show qualitative and quantitative results

Requirements – Final project

Possibilities:

- Extension of a technique studied in class
- Analysis and empirical evaluation of an existing technique
- Comparison between two approaches
- Design and evaluate a novel approach
- Be creative!

Work individually or in pair

Computational resources: http://www.arc.vt.edu/

Requirements – Project tentative timeline

- Project proposals (1 page) **[10%]**
  - March 2nd
- Final project presentations **[40%]**
  - 10-mins in-class presentation
  - May 4 to 8
- Final project report **[50%]**
  - 6-8 pages research paper
  - May 8

Each Lecture

- ~ 15 minute discussion on paper we read
  - Led by two students: “for” and “against”
- ~ 30 minute presentation on topic
- ~ 15 minute presentation on experiment
- ~ 15 minutes for questions, interruptions, unplanned discussions

Computer Vision Fundamentals

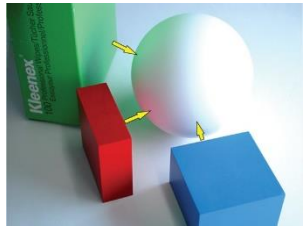

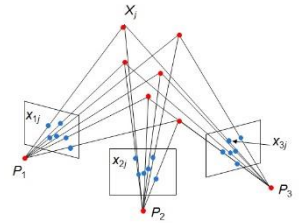

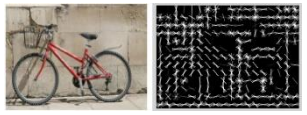

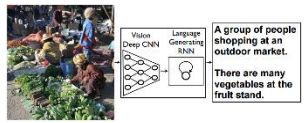

- Interpreting Intensities
  - What determines the brightness and color of a pixel?
  - How can we use image filters to extract meaningful information from the image?
- Correspondence and Alignment
  - How can we find corresponding points in objects or scenes?
  - How can we estimate the transformation between them?
- Perspective and 3D Geometry
  - How can we map between the 3D world and the 2D image?
  - How can we recover 3D coordinates from images or video?
- Grouping and Segmentation
  - How can we group pixels into meaningful regions?
- Categorization and Object Recognition
  - How can we represent images and categorize them?
  - How can we recognize categories of objects?
- Advanced Topics
  - Action recognition, 3D scenes and context, CNNs, …

Major topics in this class

- **Visual Recognition**
  - Instance/category recognition, ConvNets, detection, segmentation, pose
- **Representation Learning**
  - Attributes, self-supervised learning, generative models, image styles
- **Activity and Event**
  - Action recognition, active perception, object relationships, first-person
- **Multi-modality**
  - Language, sketches, sounds
- **Applications and Data**
  - Robotics, graphics, big data, human-in-the-loop, crowdsourcing

R-CNN: *Regions with CNN features*

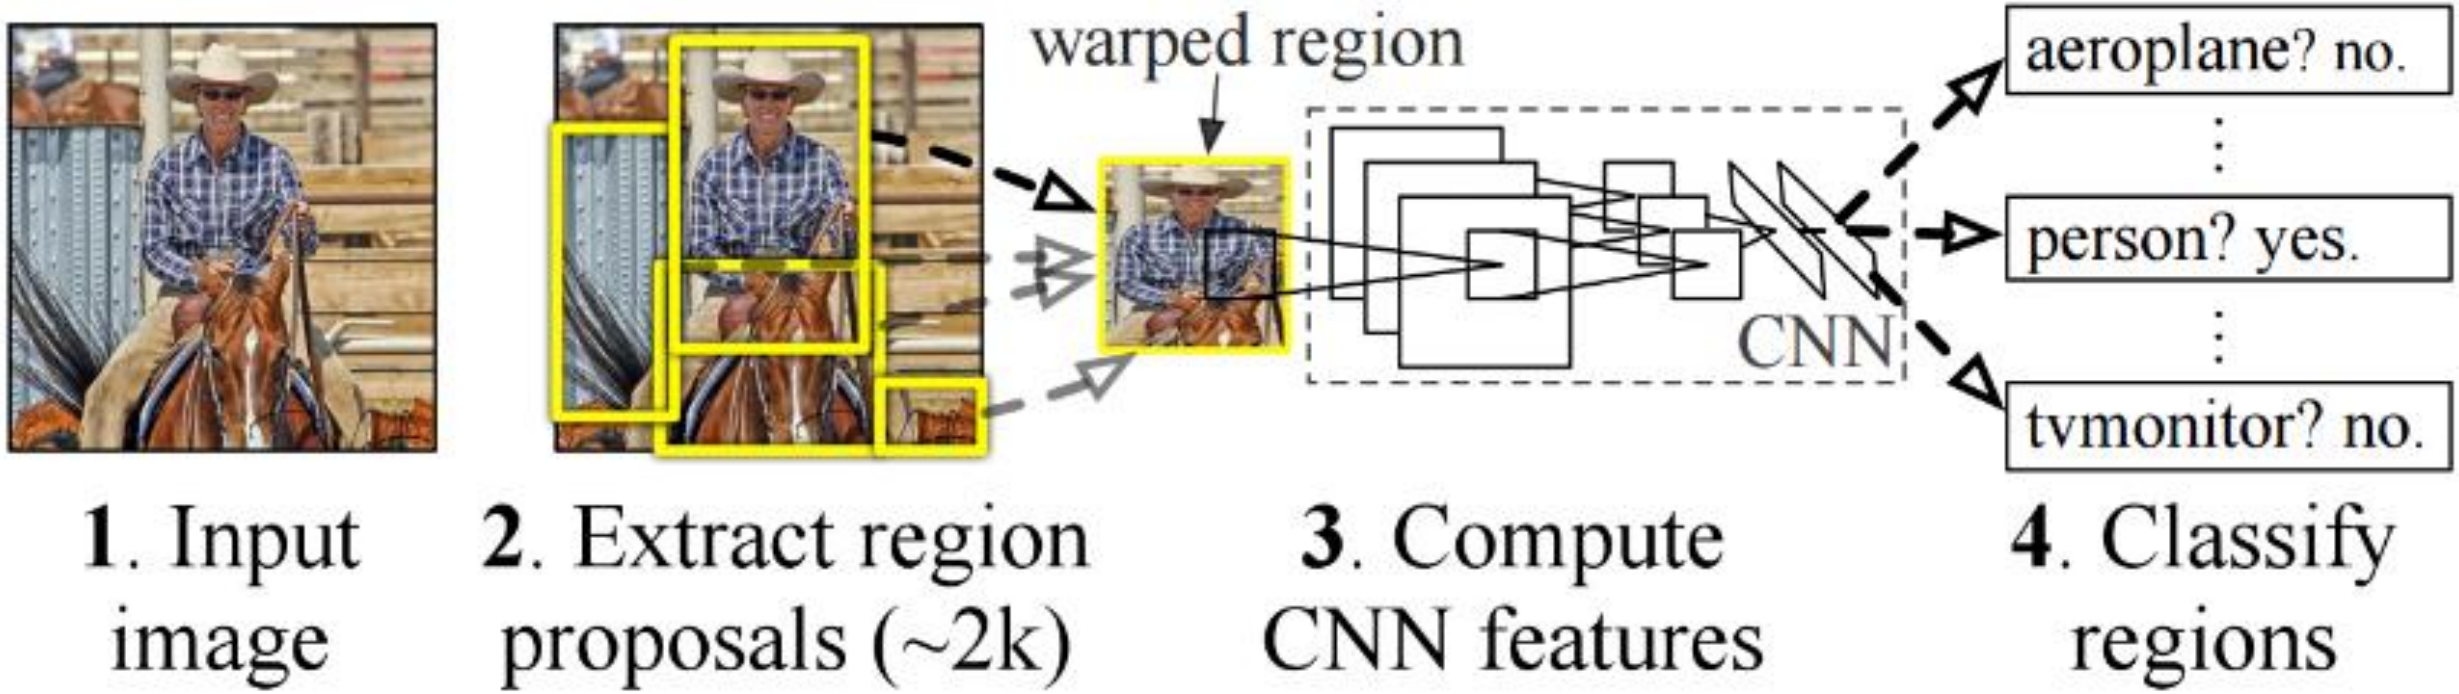

RCNN [Girshick et al. CVPR 2014]

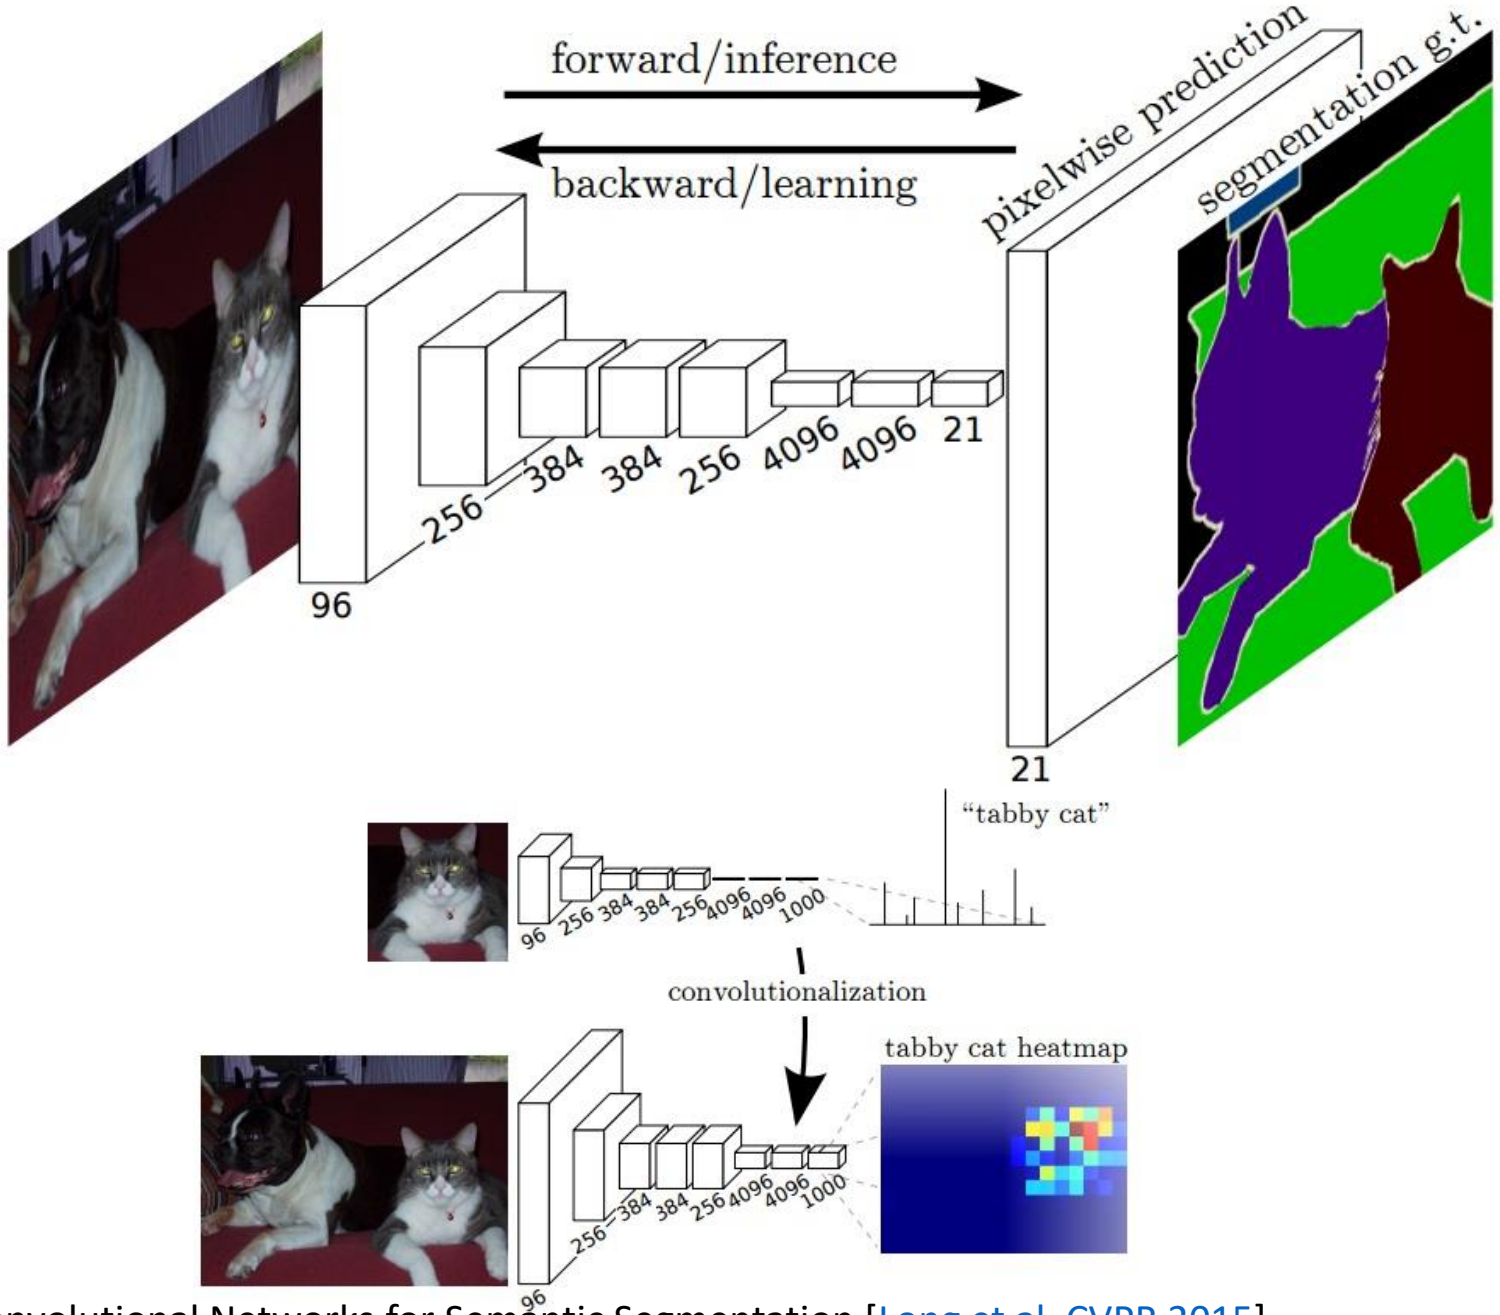

Fully Convolutional Networks for Semantic Segmentation [Long et al. CVPR 2015]

CNN as a Similarity Measure for Matching

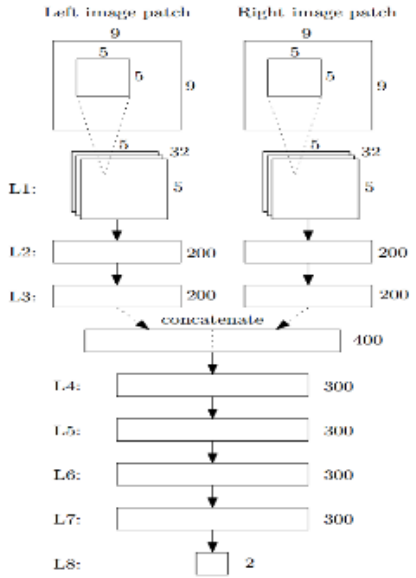

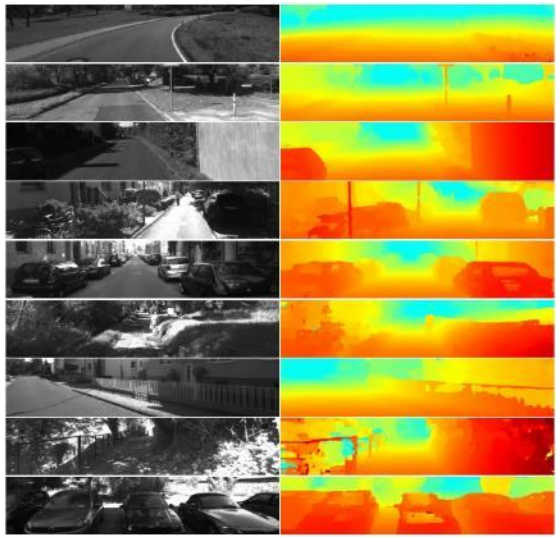

Stereo matching [Zbontar and LeCun CVPR 2015]
Compare patch [Zagoruyko and Komodakis 2015]

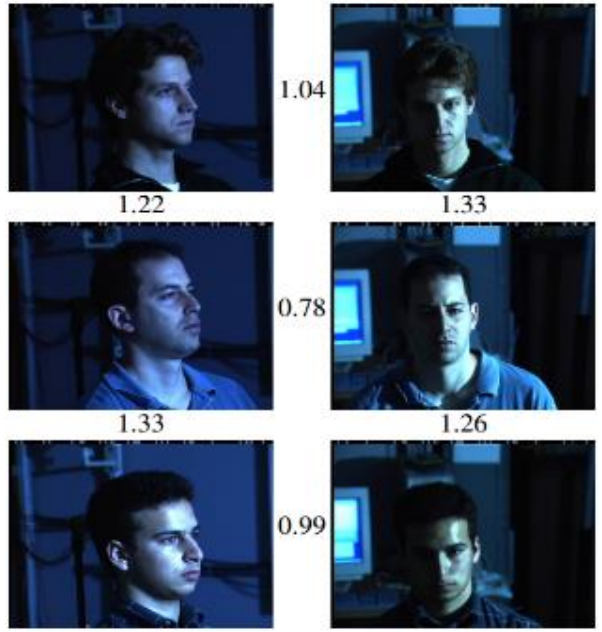

FaceNet [Schroff et al. 2015]

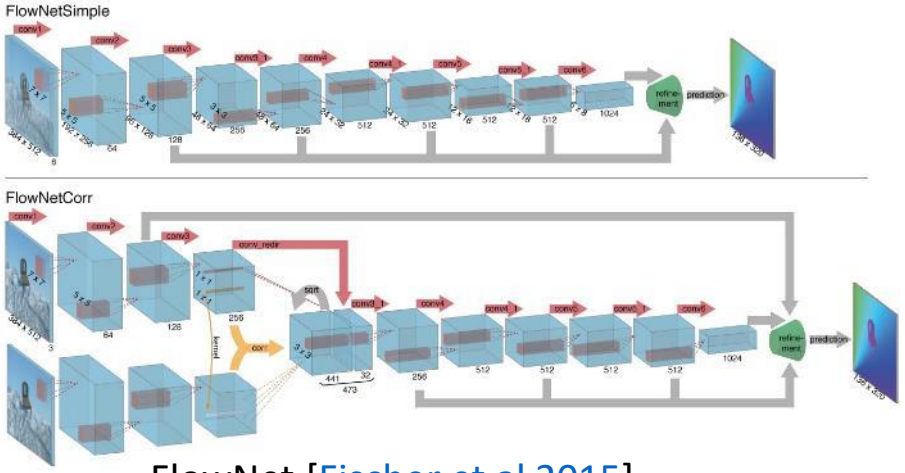

FlowNet [Fischer et al 2015]

Match ground and aerial images
[Lin et al. CVPR 2015]

DeepPose [Toshev and Szegedy CVPR 2014]

otter

black:	yes
white:	no
brown:	yes
stripes:	no
water:	yes
eats fish:	yes

polar bear

black:	no
white:	yes
brown:	no
stripes:	no
water:	yes
eats fish:	yes

zebra

black:	yes
white:	yes
brown:	no
stripes:	yes
water:	no
eats fish:	no

Learning To Detect Unseen Object Classes by Between-Class Attribute Transfer. [Lampert et al. CVPR 2009]

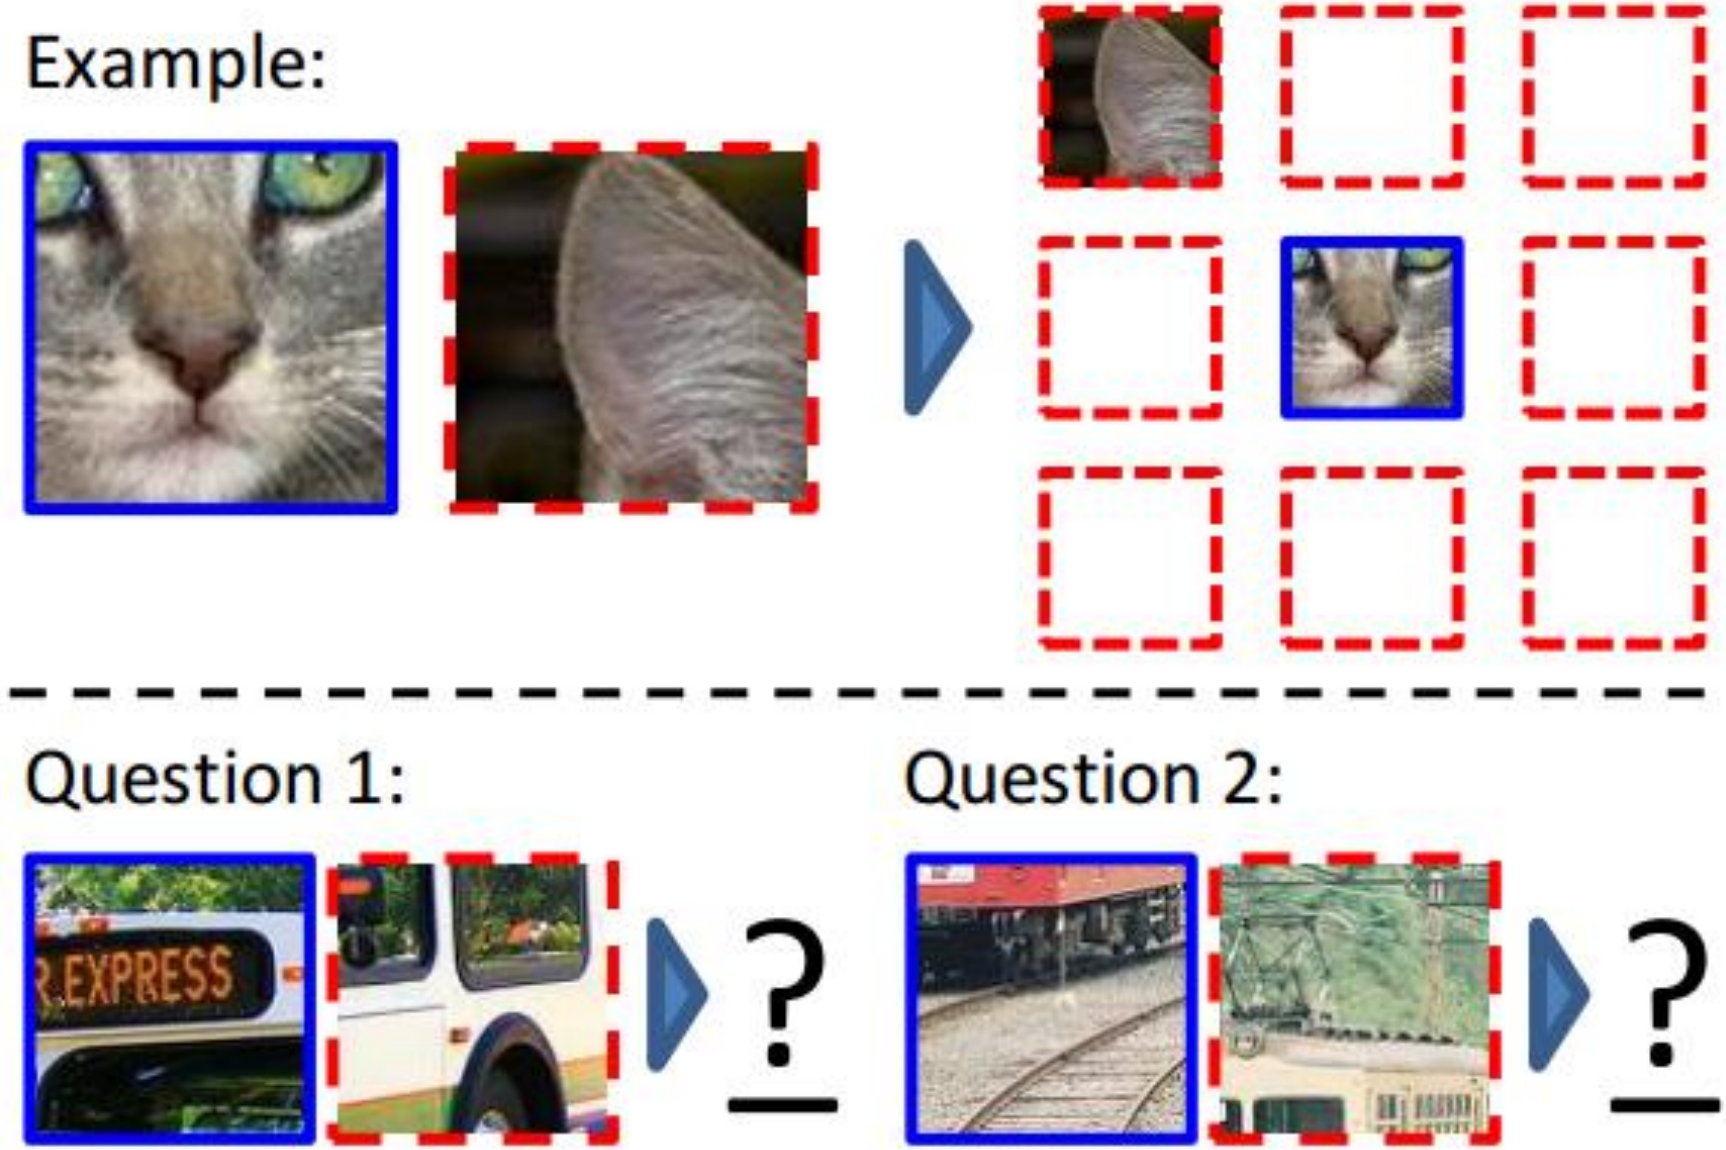

Context Prediction [Doersch et al. ICCV 2015]

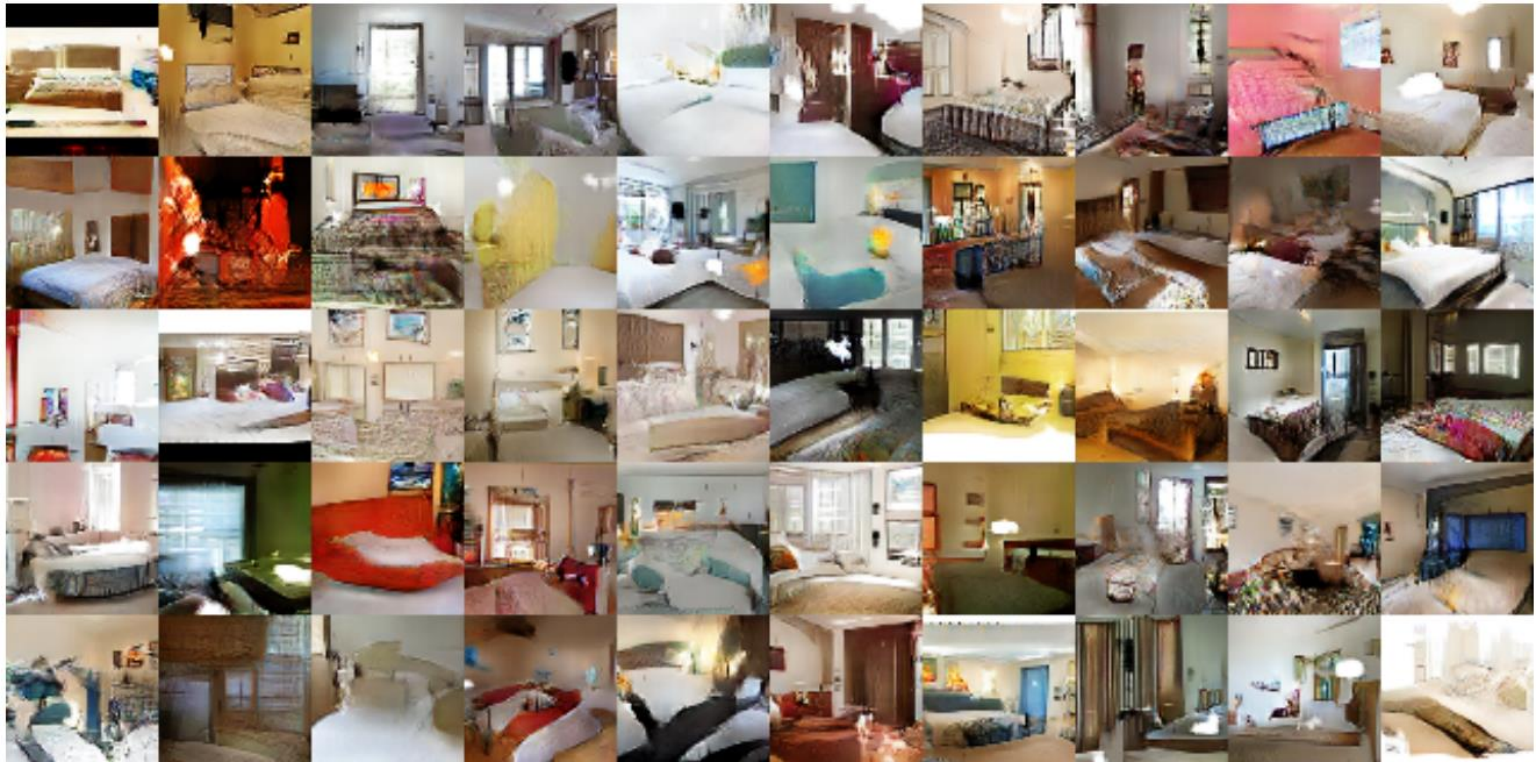

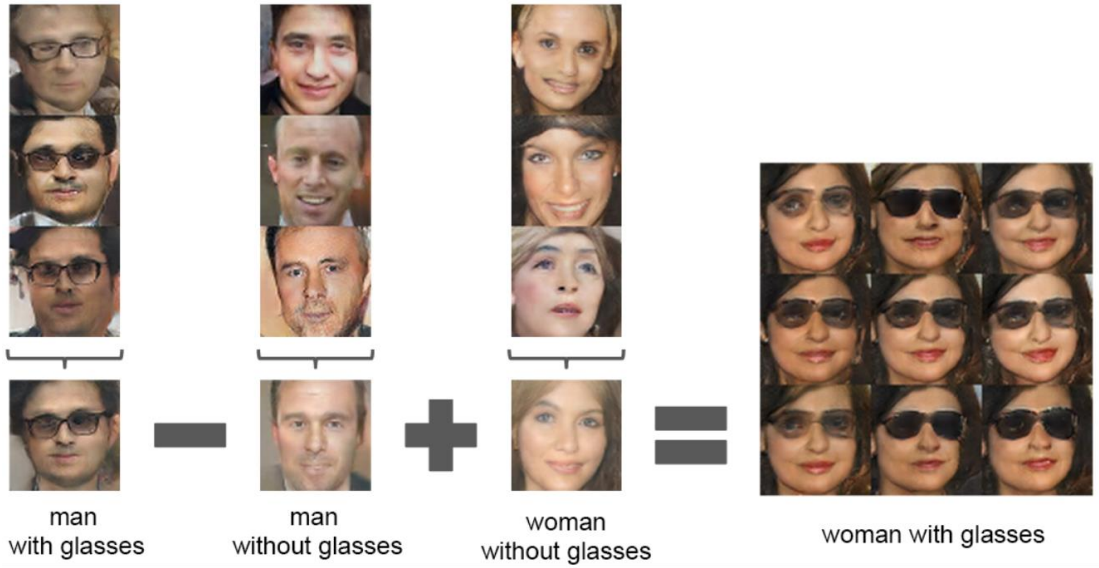

[Radford et al. ICLR 2016]

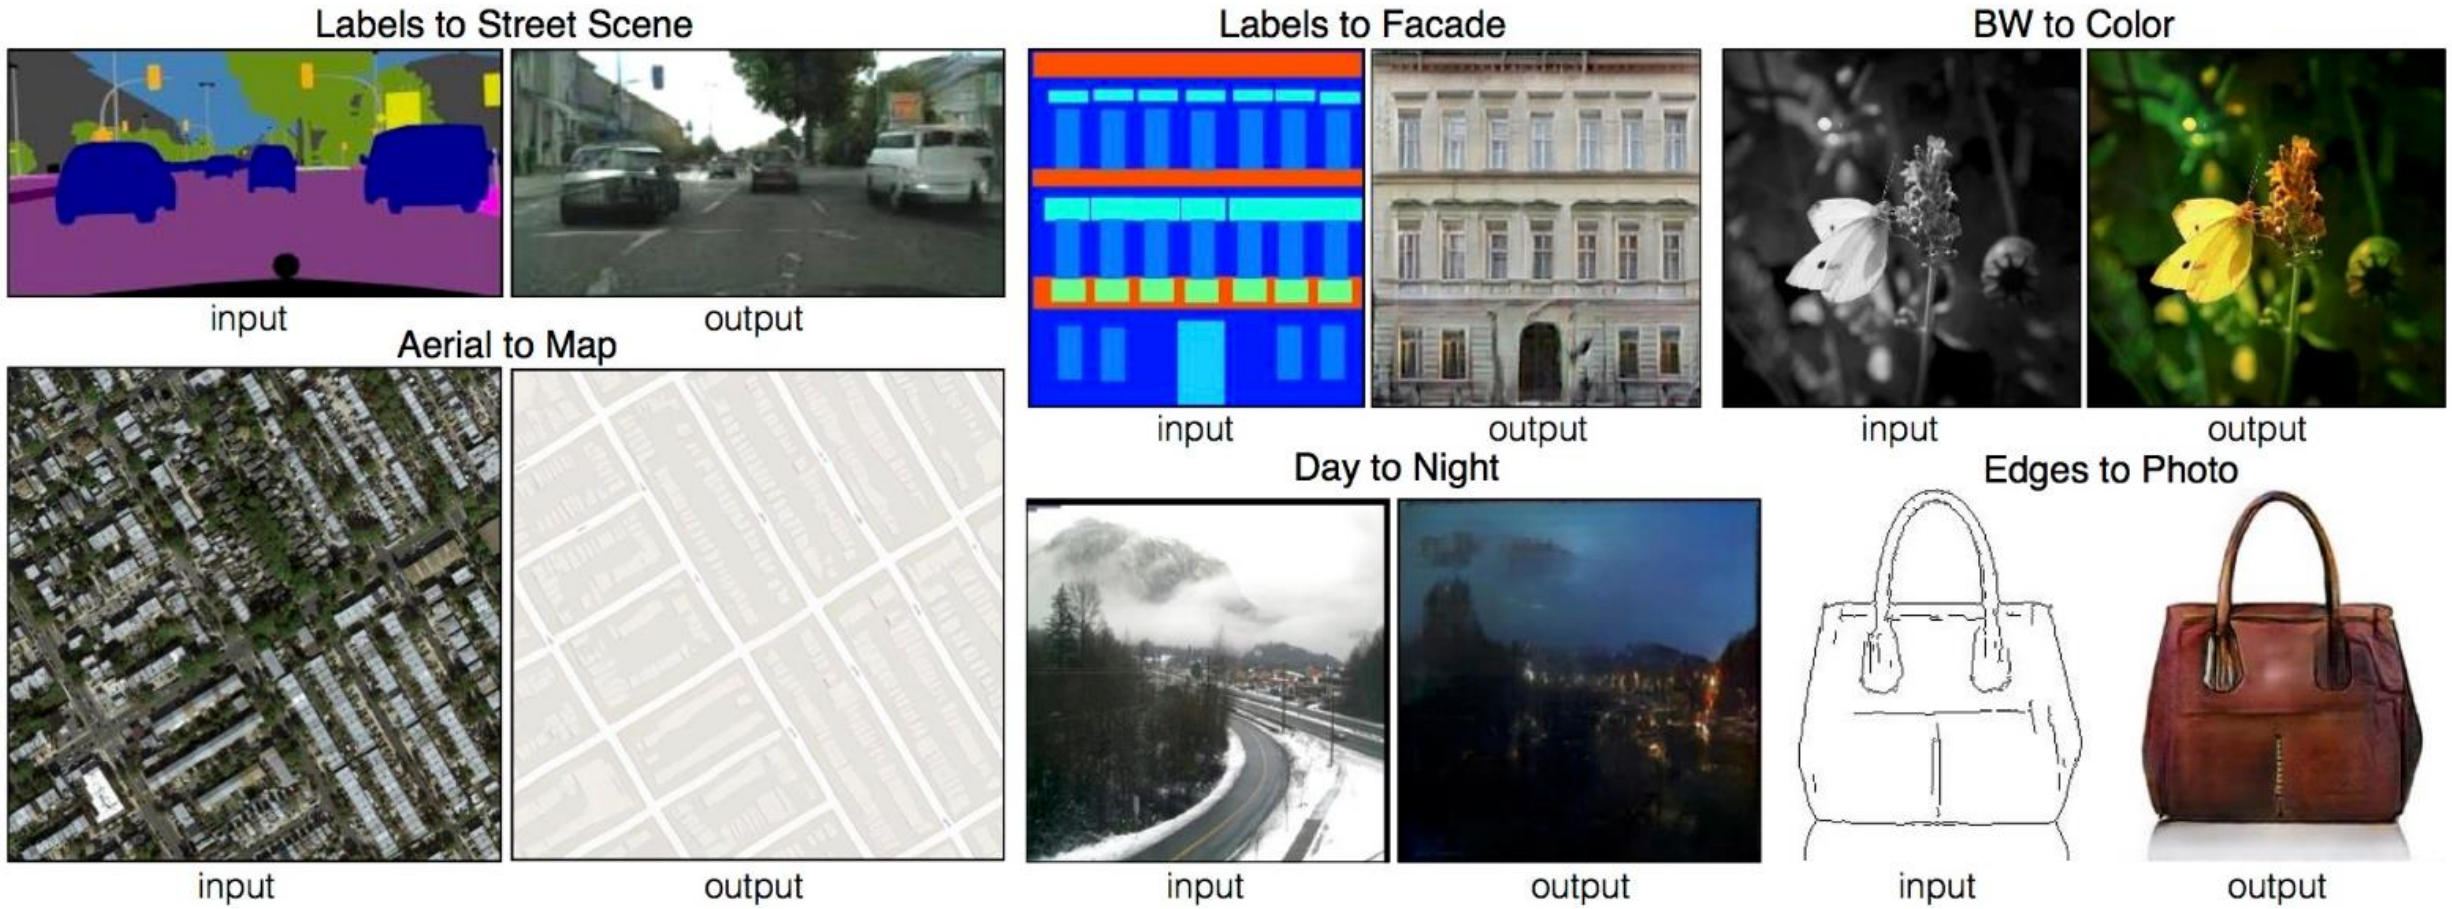

[Isola et al. 2016]

[Gatys et al. CVPR 2016]

Action recognition

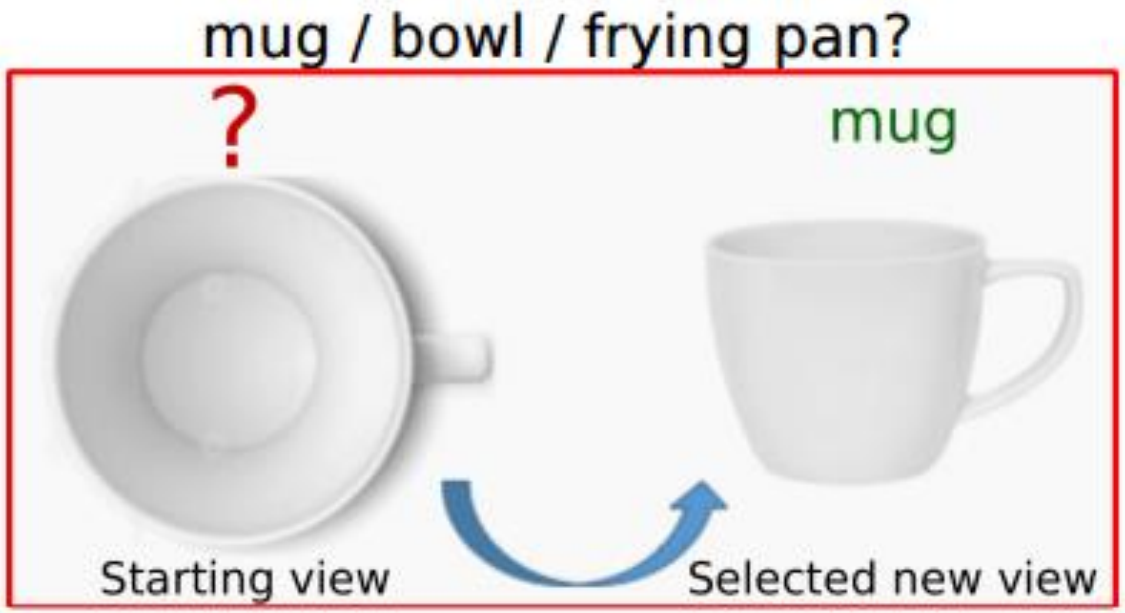

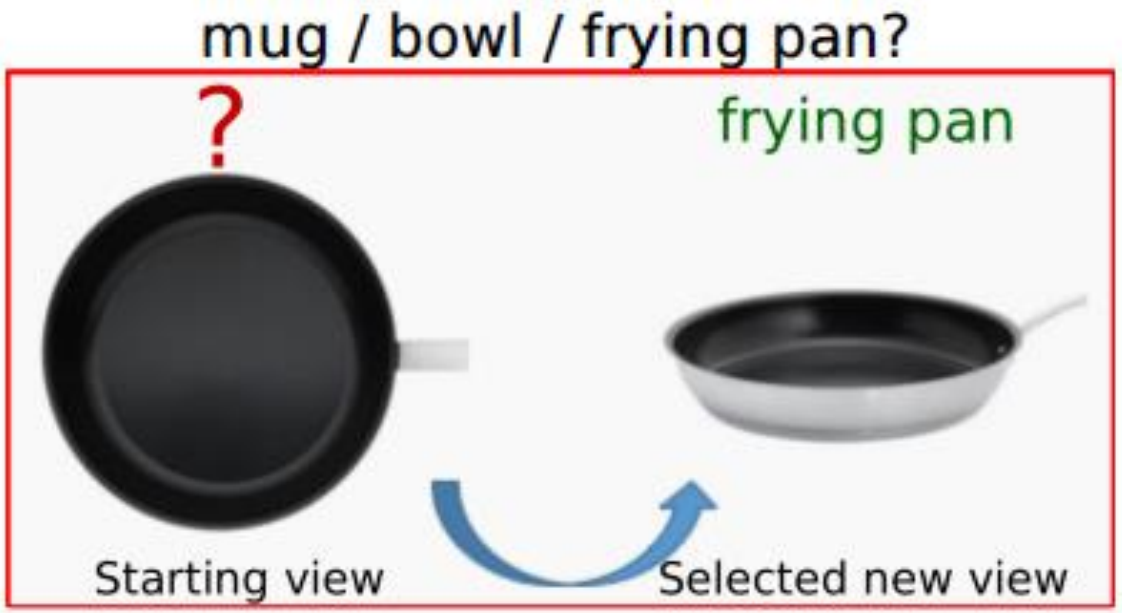

Active Perception

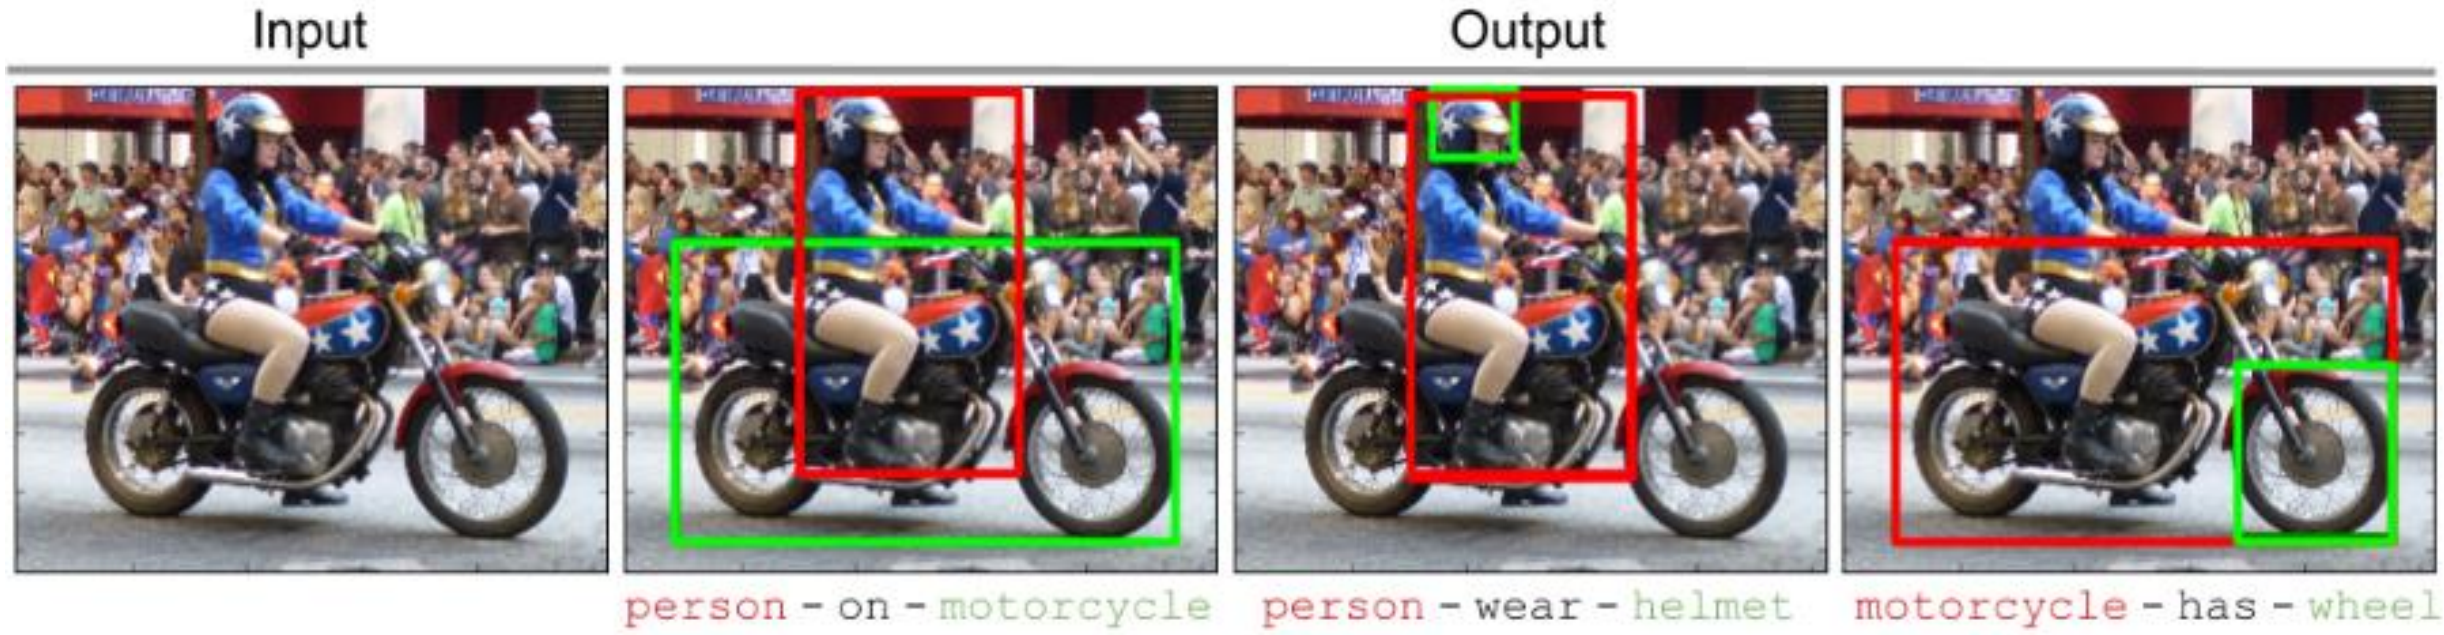

Groups of Objects

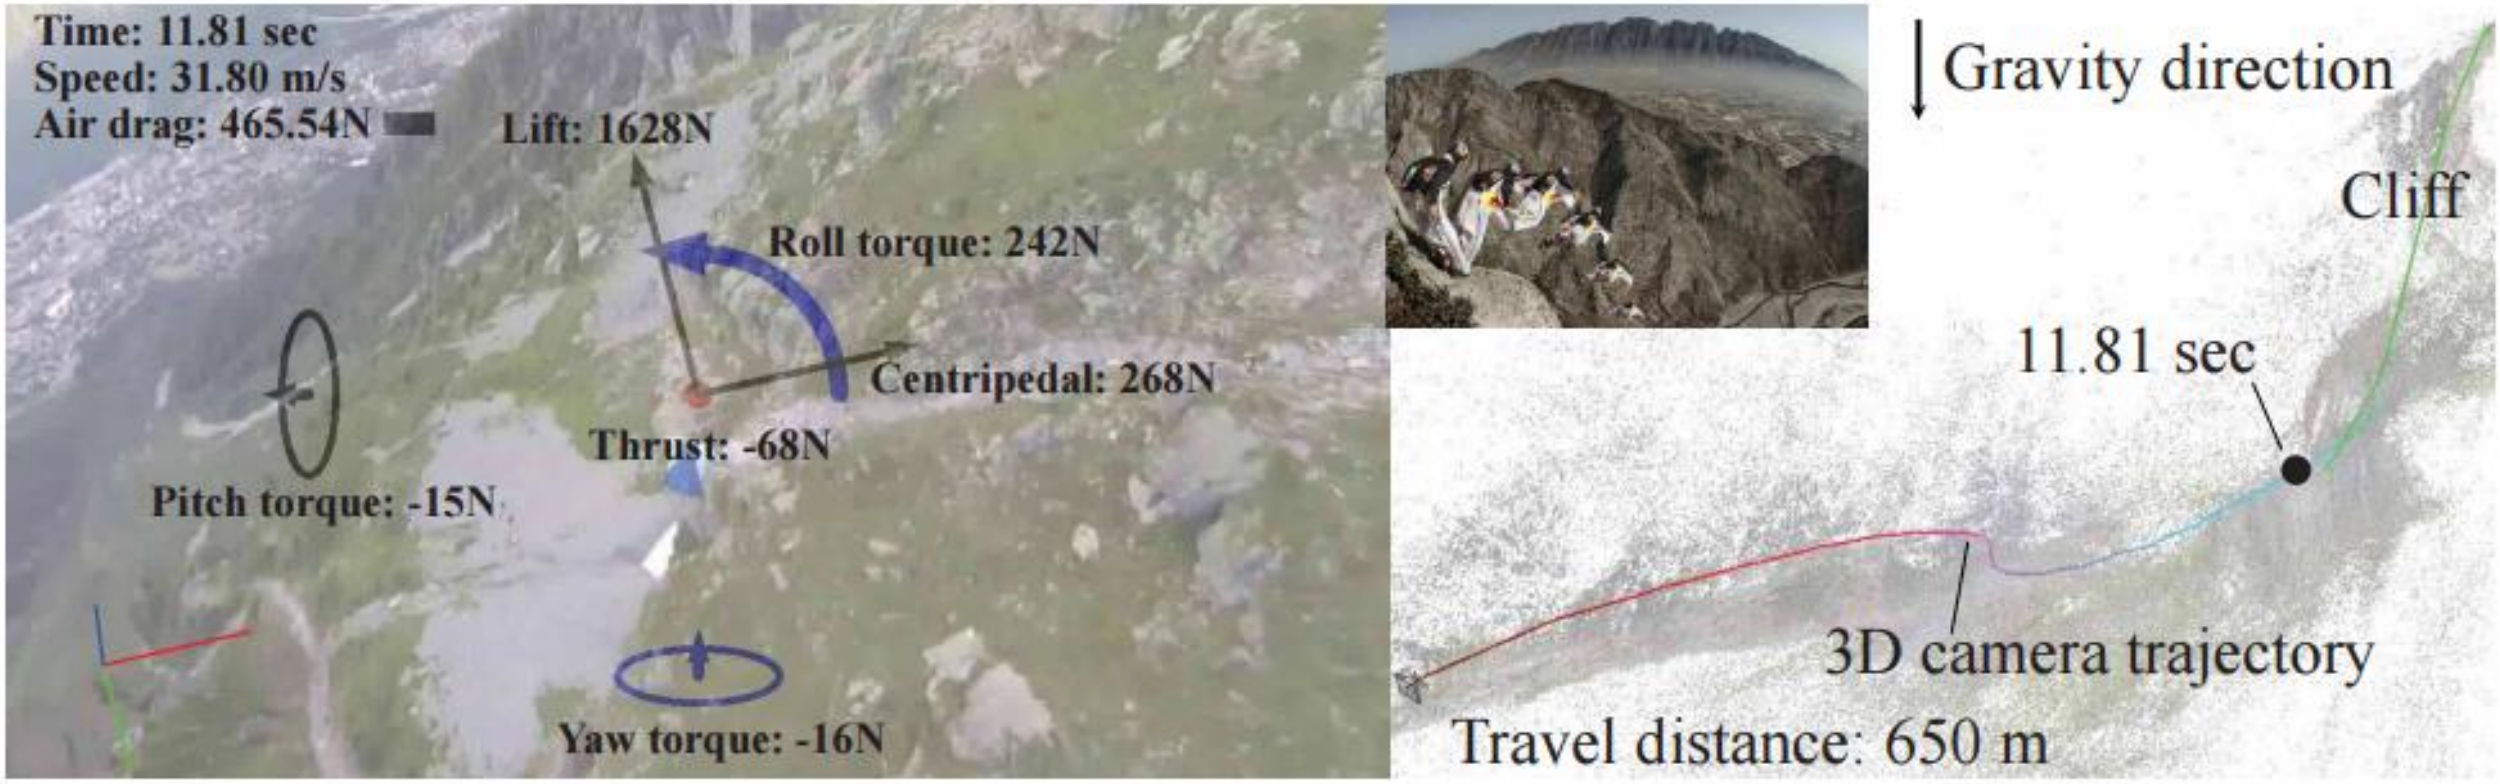

First-person vision

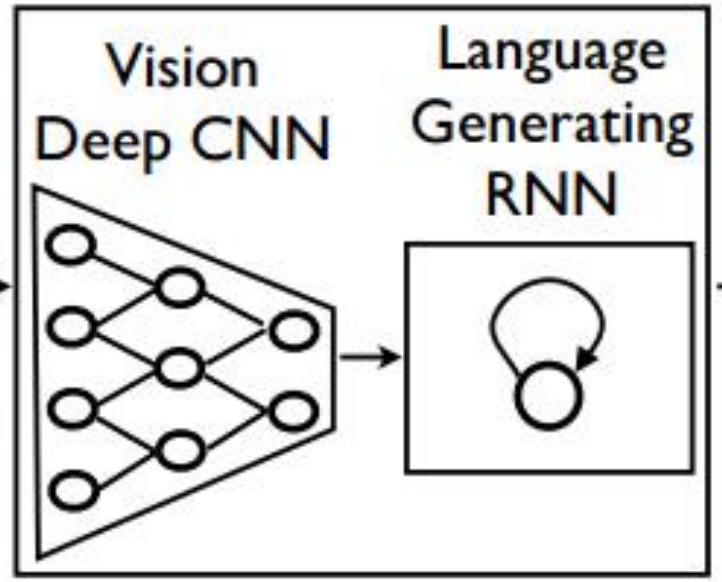

A group of people shopping at an outdoor market.

There are many vegetables at the fruit stand.

Language and vision

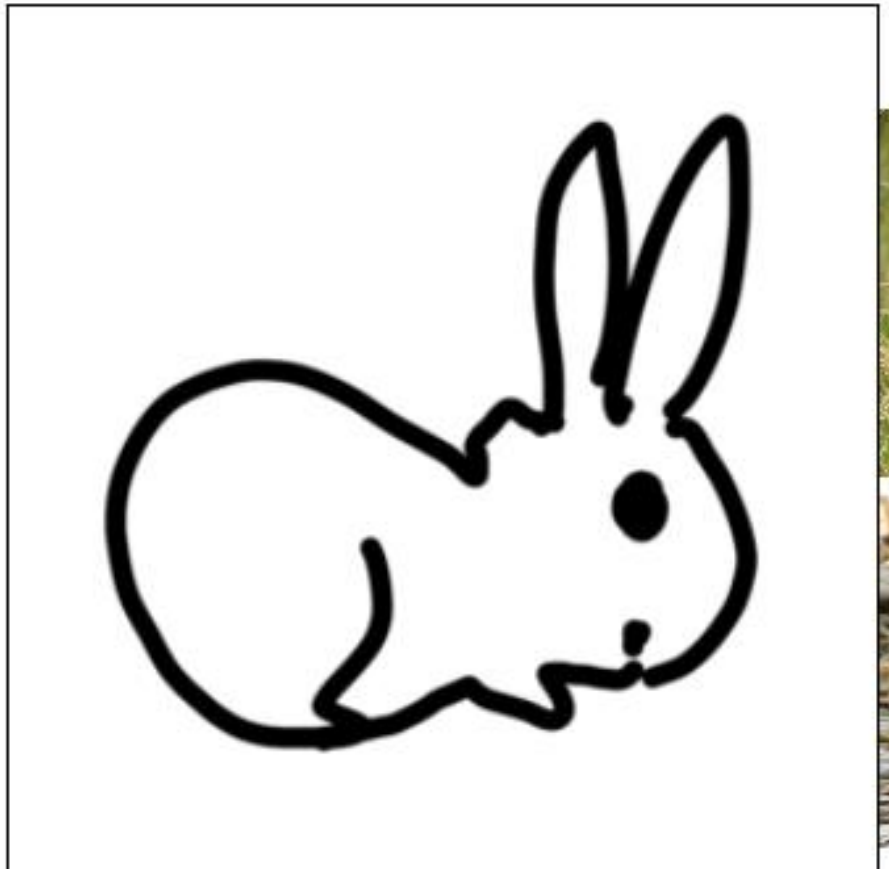

Sketches

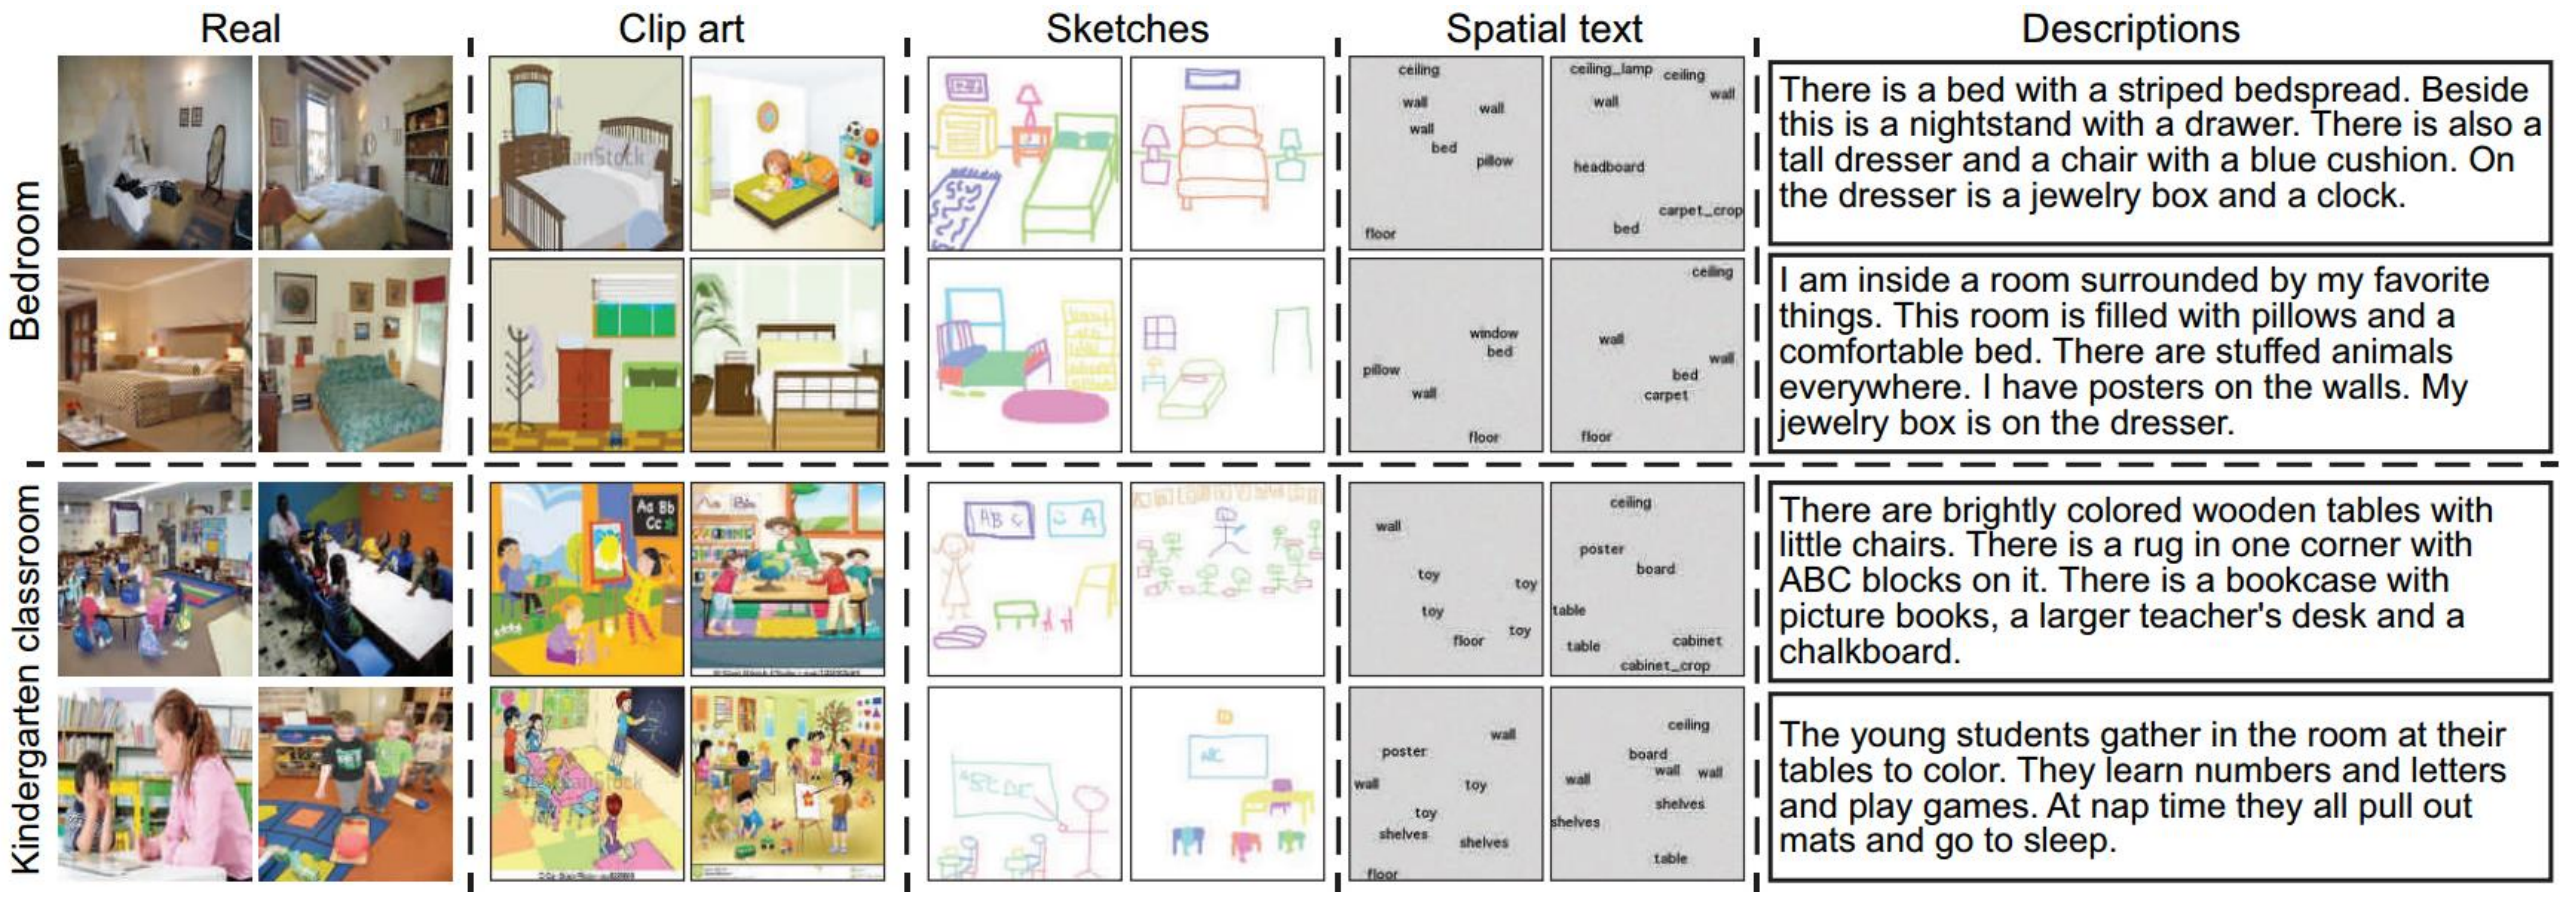

Cross-modal learning

Getting help outside of class

Discussion Board:
http://piazza.com/vt/spring2017/ece6554

Readings/Textbook:

http://szeliski.org/Book/

Lecture notes: will be posted online

Use Office Hours / After class

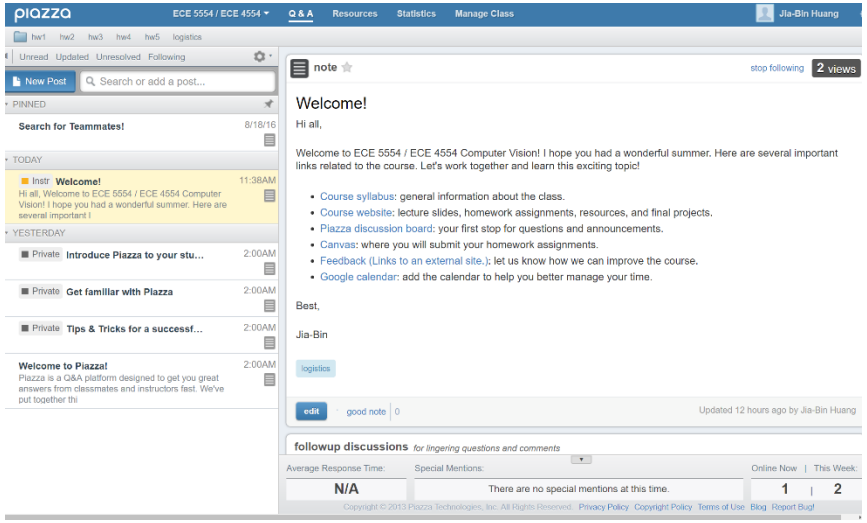

Things to remember

- To-Do
  - Sign up piazza discussion board
  - Sign up topics you would like to present/discuss/experiment
  - Submit your paper review via piazza
- Next class: Instance recognition
- Questions?