Slide Credit: Kristen Grauman

CS376 Computer Vision

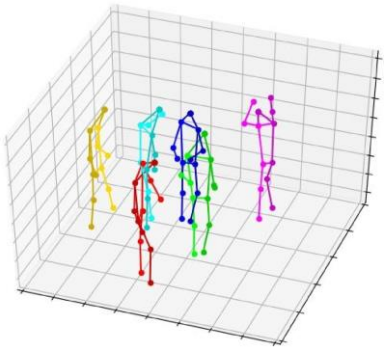

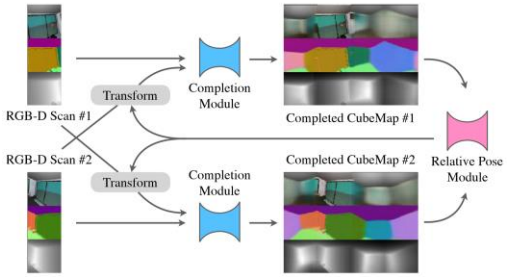

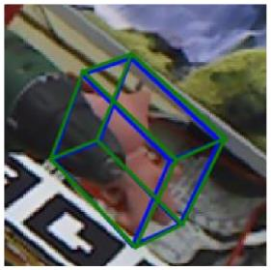

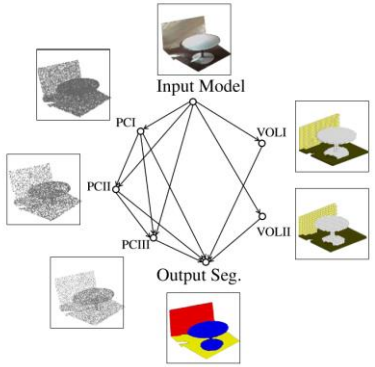

Qixing Huang

January 23th 2019

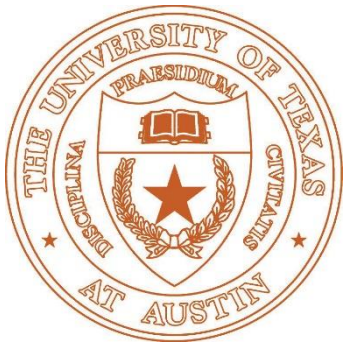

Introductions

- Instructor
  - Prof. Qixing Huang

- TA:
  - Zhenpei Yang

Today

- Course overview
- Requirements, logistics

What is computer vision

Done?

Computer Vision

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

1. Vision for measurement (Sensing)

Real-time stereo

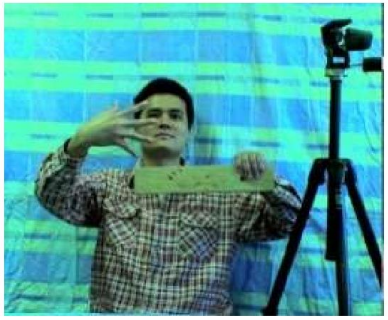

Wang et al.

Computer Vision

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

2. Vision for perception, interpretation

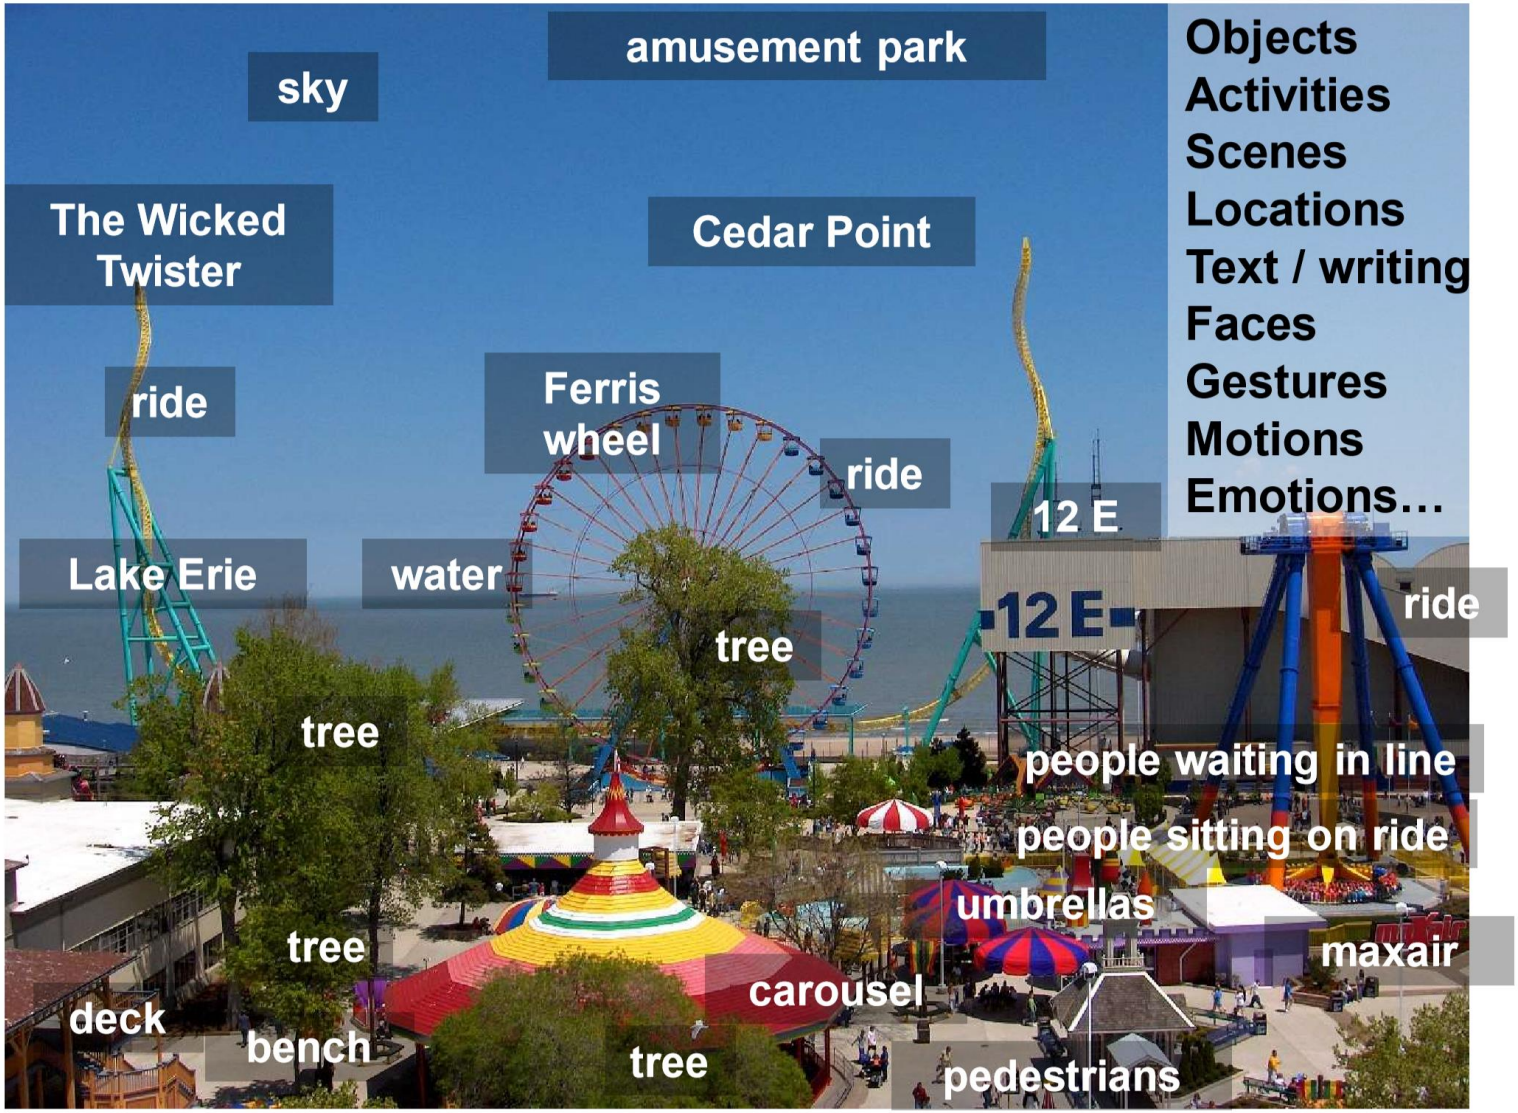

Computer Vision

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

3. Visual search, organization

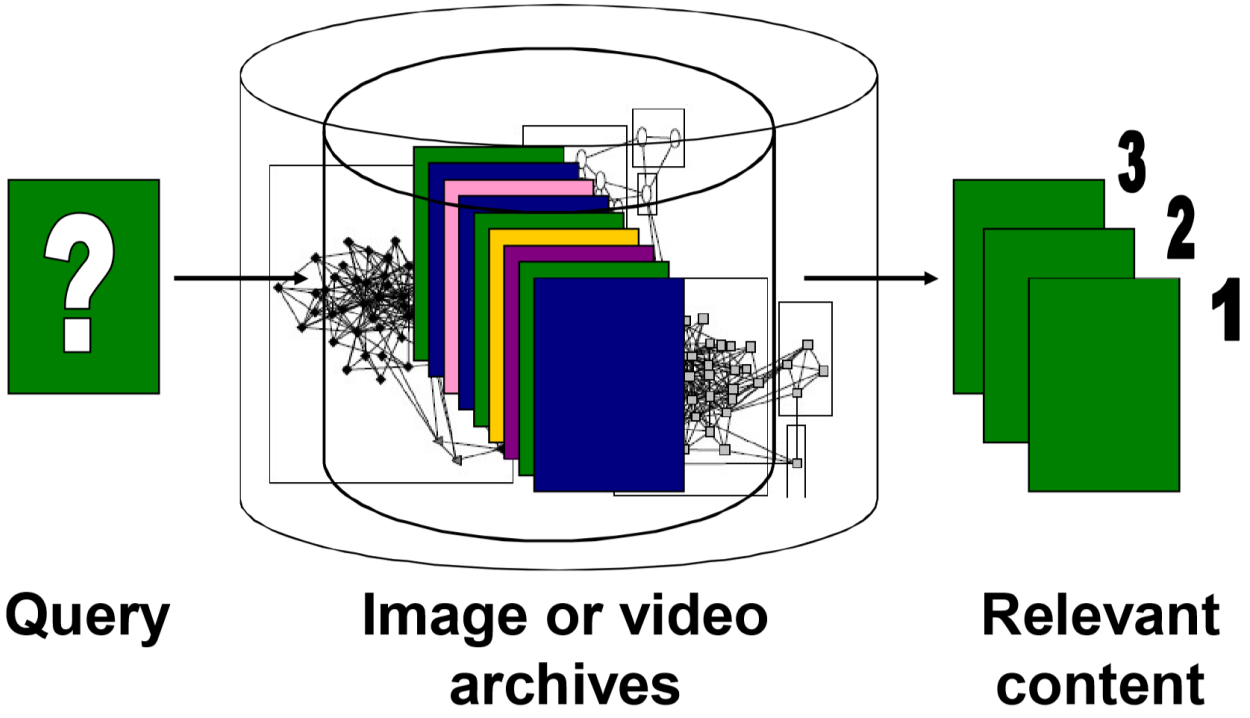

Different representations and algorithms involved

Computer Vision

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

Related disciplines

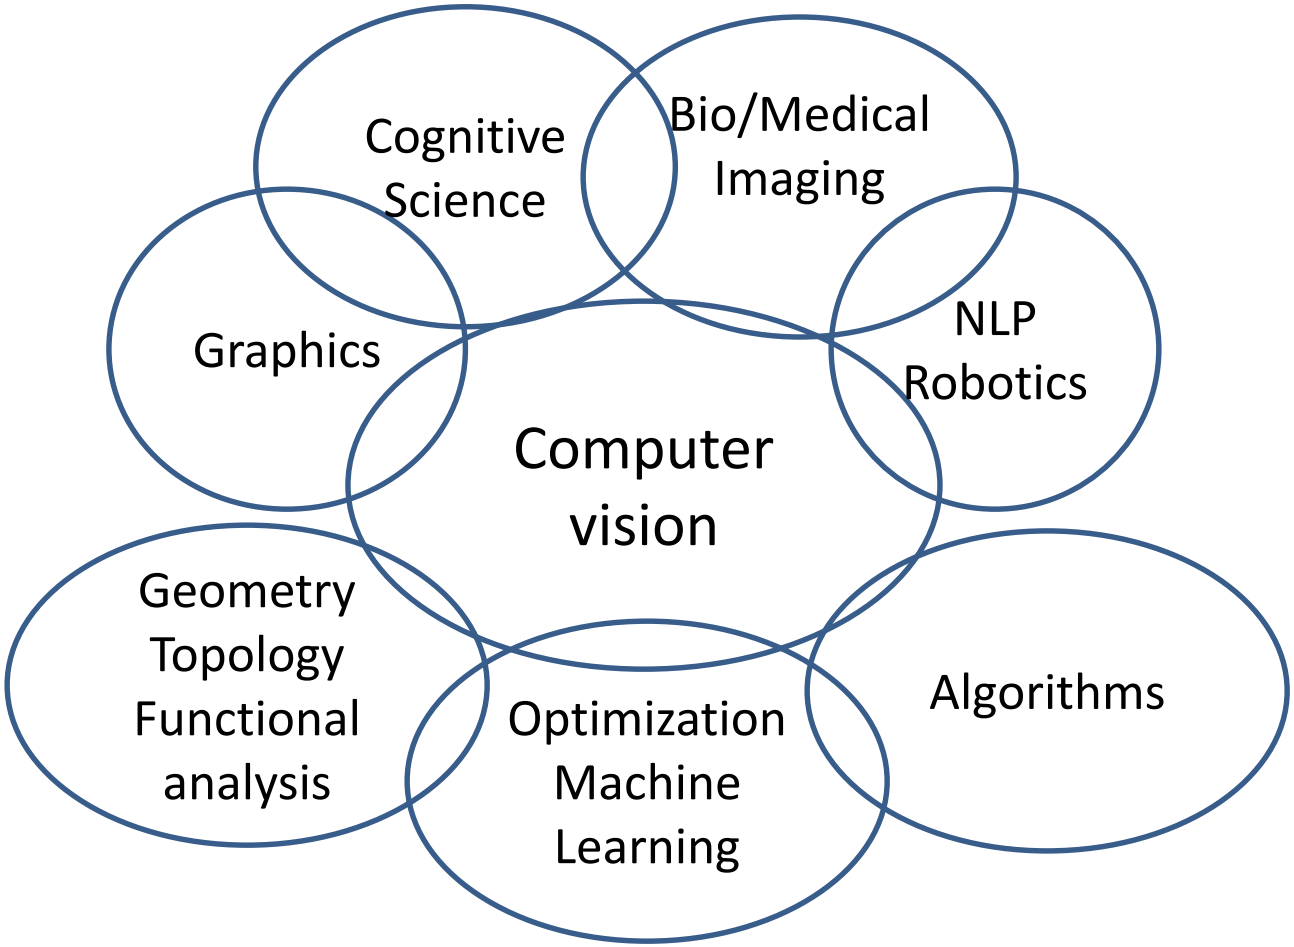

Vision and graphics

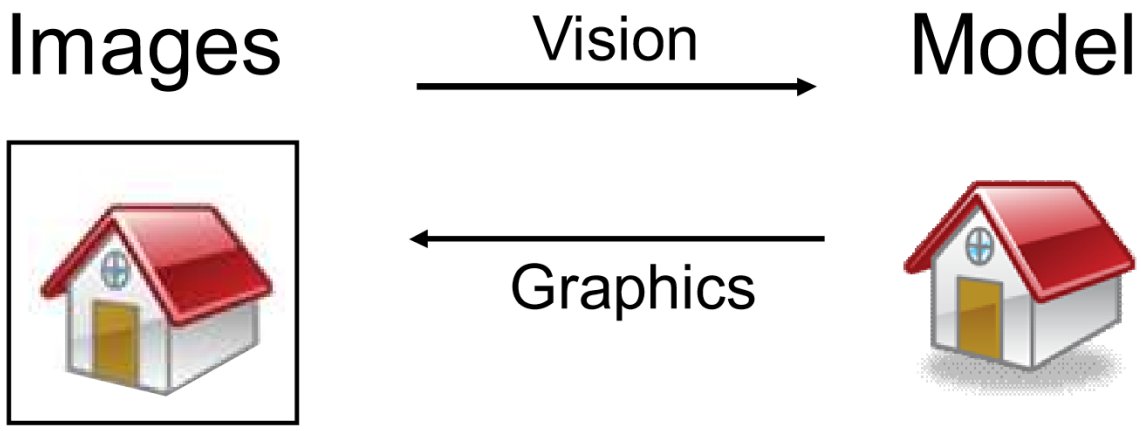

Inverse problems: analysis and synthesis

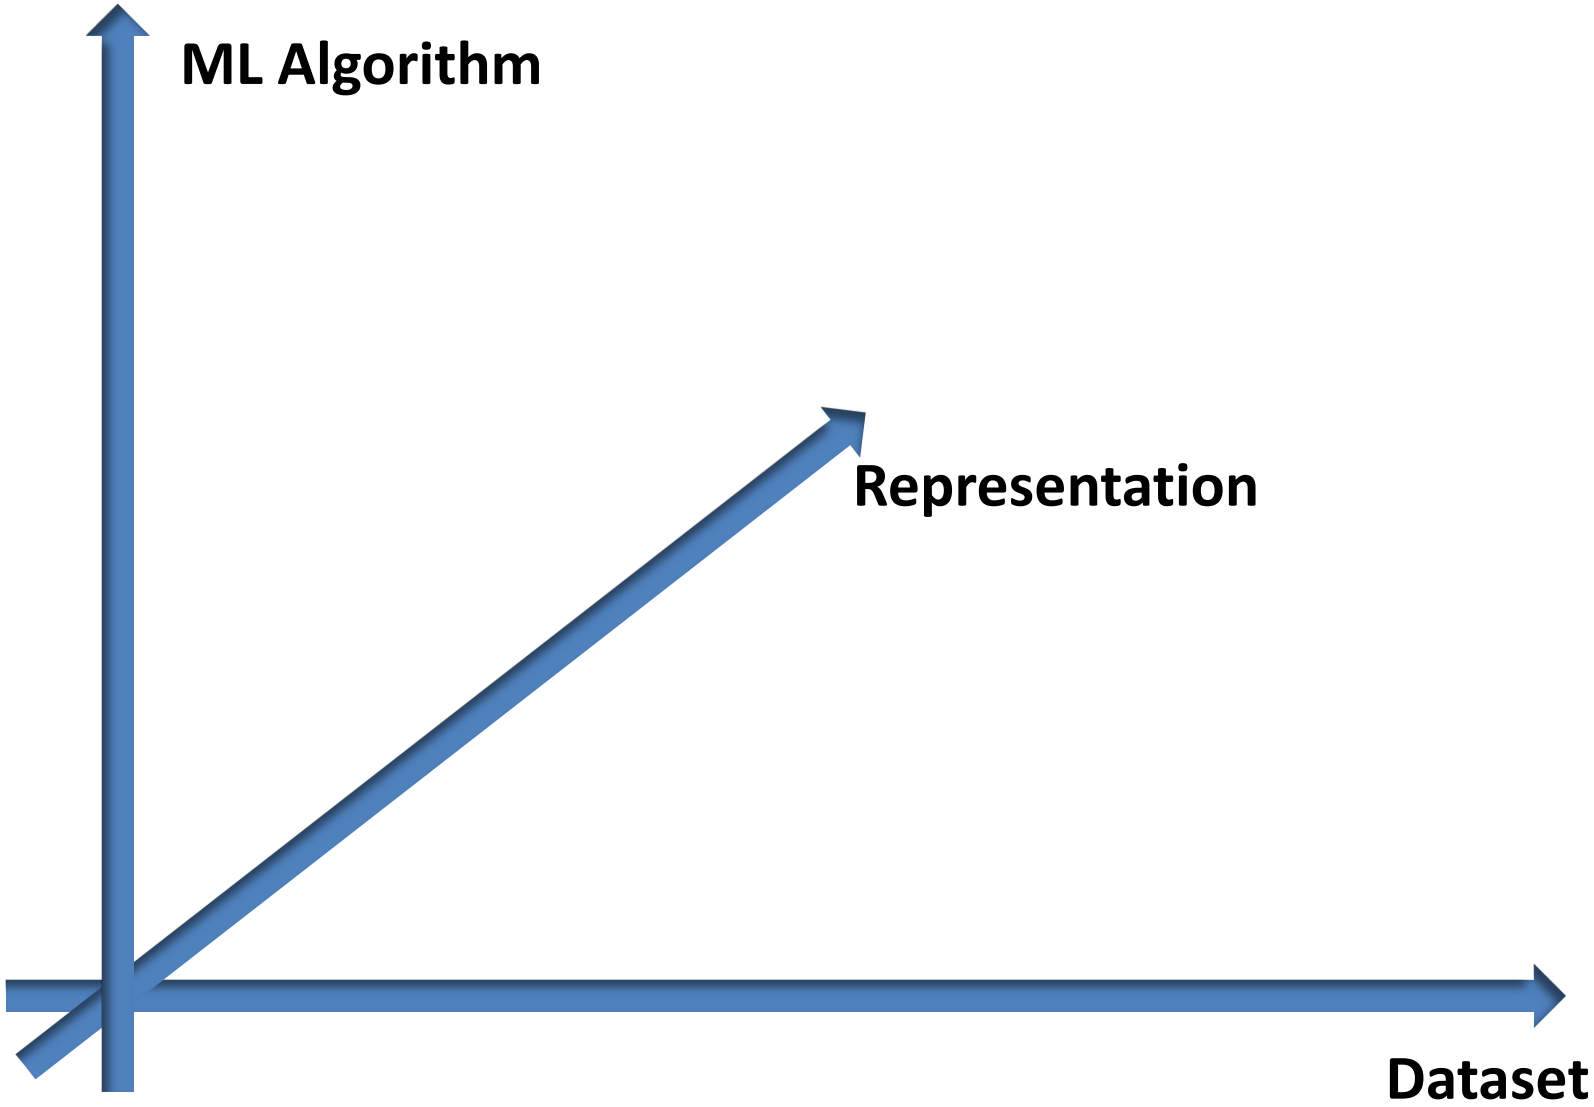
ML Algorithm
Representation
Dataset

Progress charted by datasets

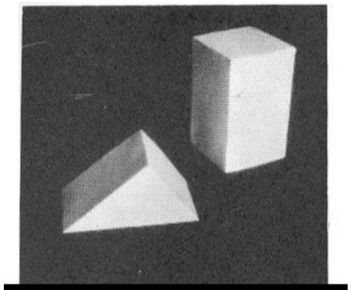

Roberts 1963

COIL

Progress charted by datasets

MSRC 21 Objects

Caltech-101

Caltech-256

Progress charted by datasets

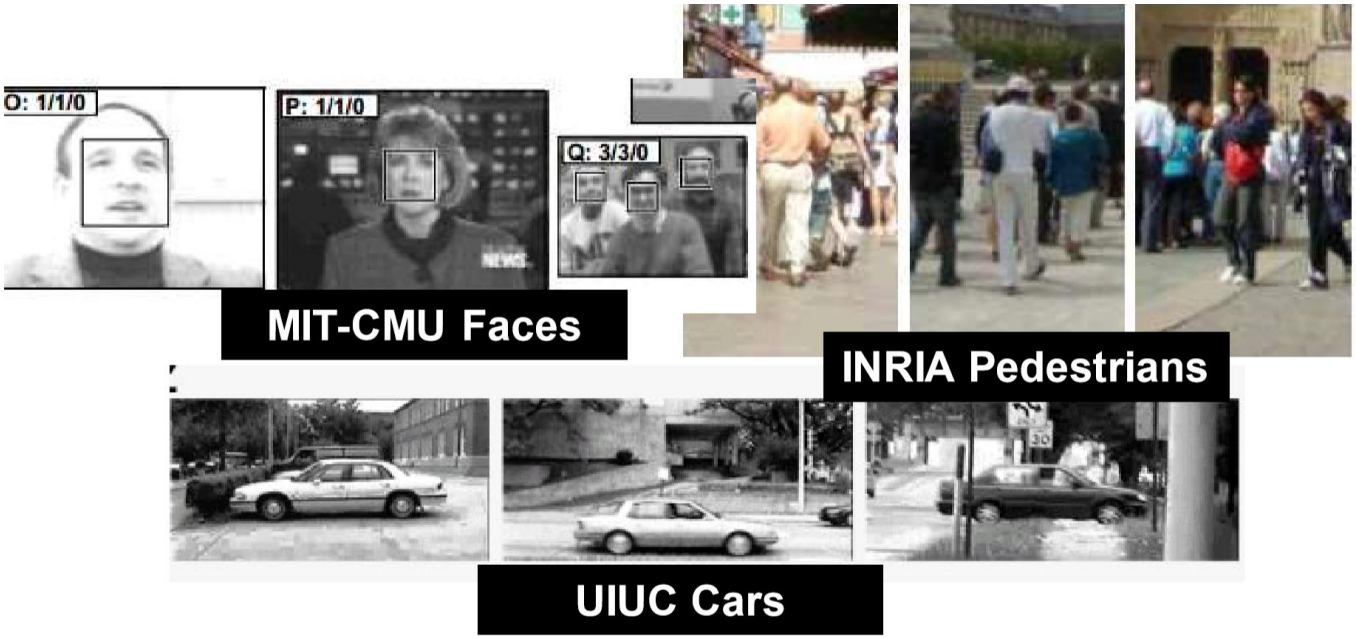

Progress charted by datasets

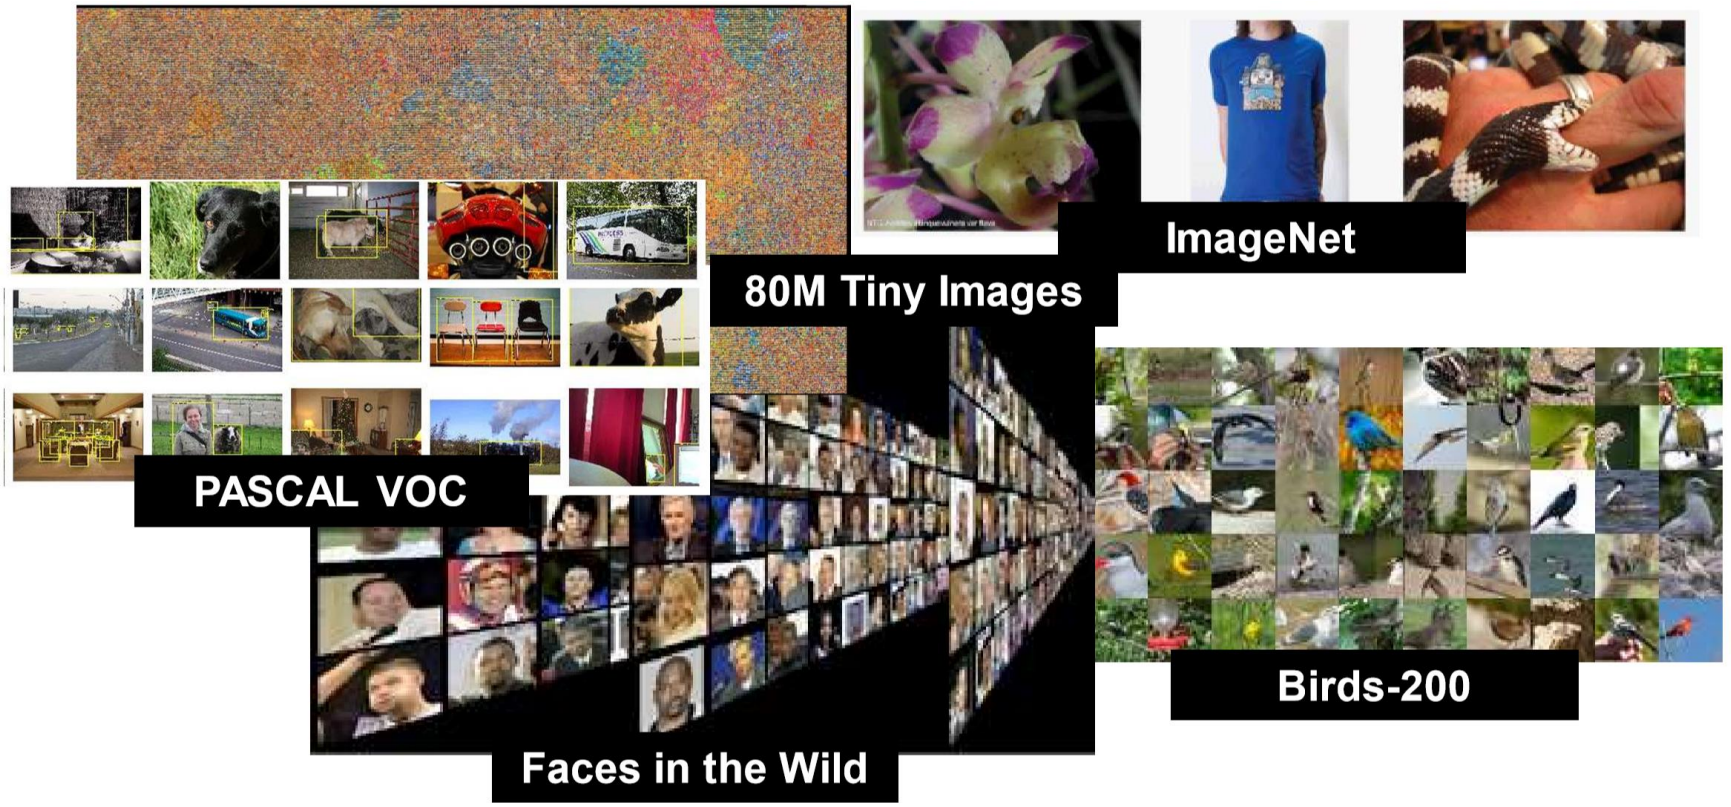

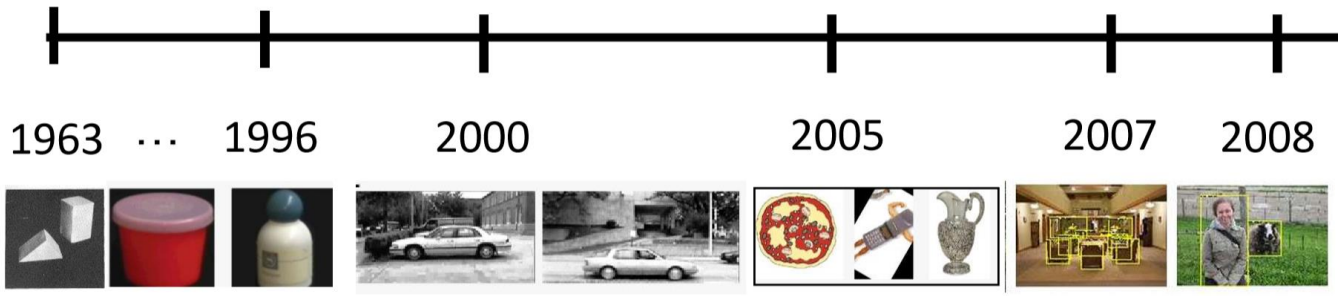

Data Representations

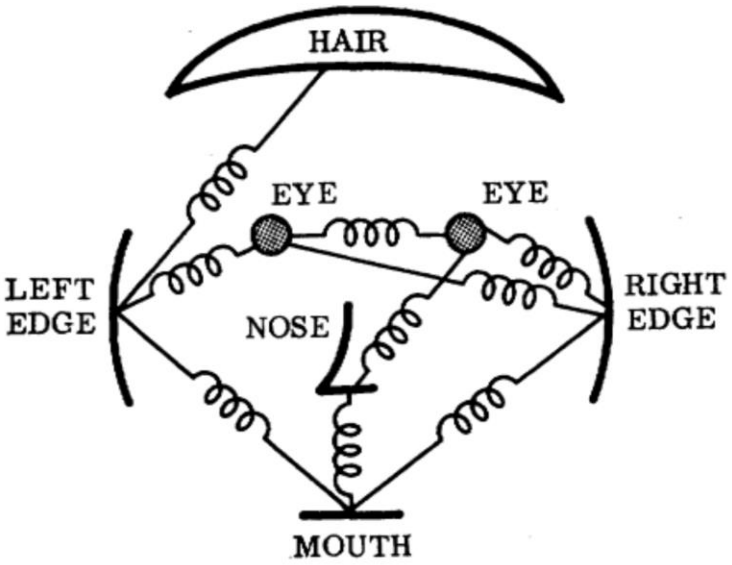

Pictorial Structures
(Fischler et al. 73)

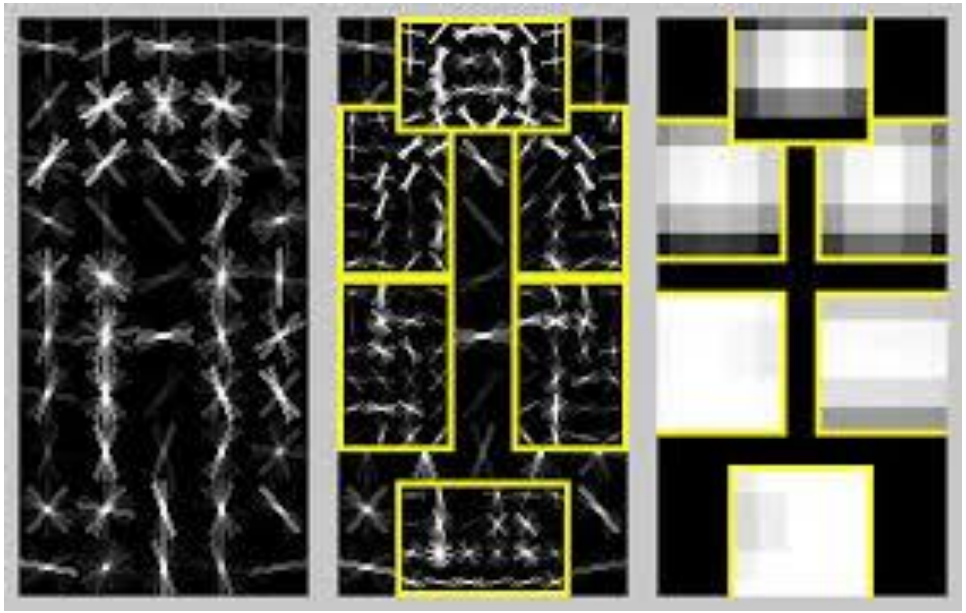

Deformable-Part-Model
(Felzenszwalb et al. 10)

1973

2010

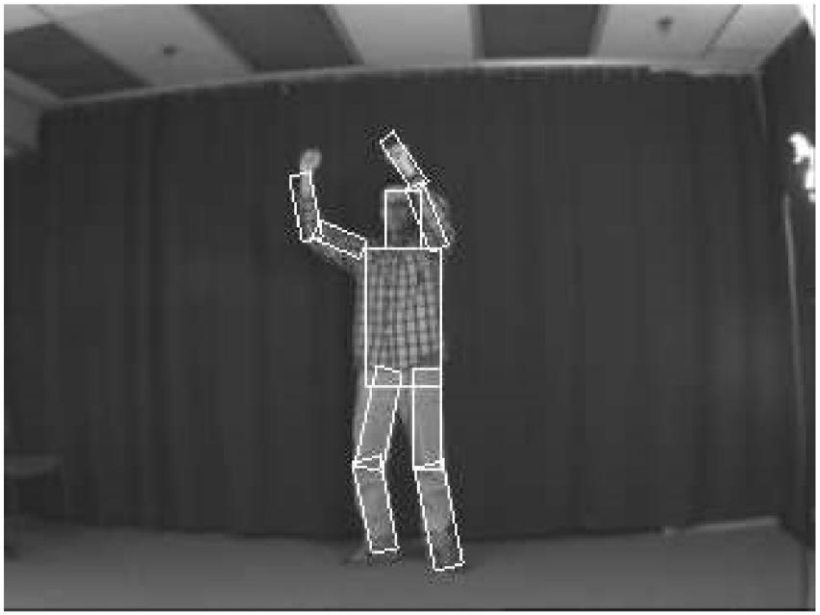

Pictorial Structures
(Fischler et al. 73)

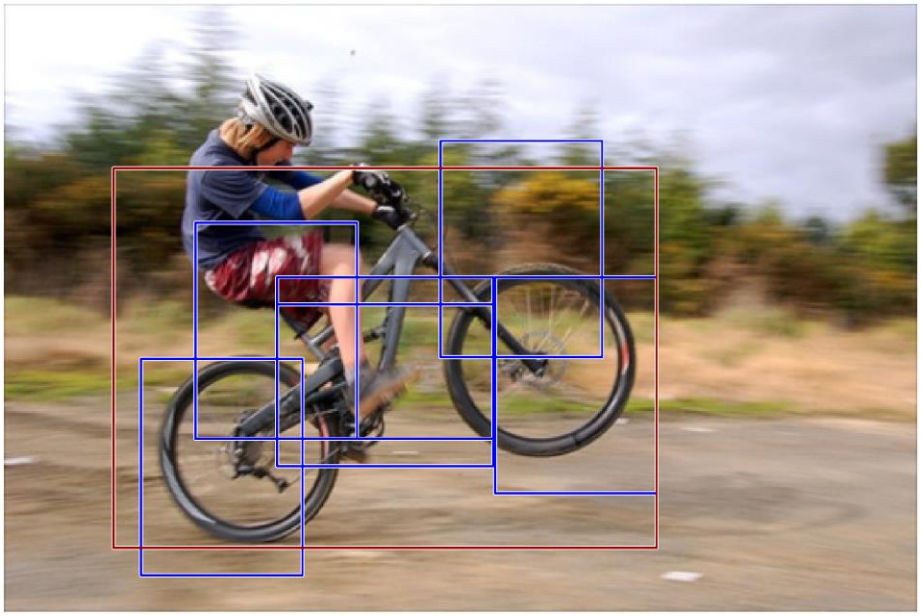

Deformable-Part-Model
(Felzenszwalb et al. 10)

1973

2010

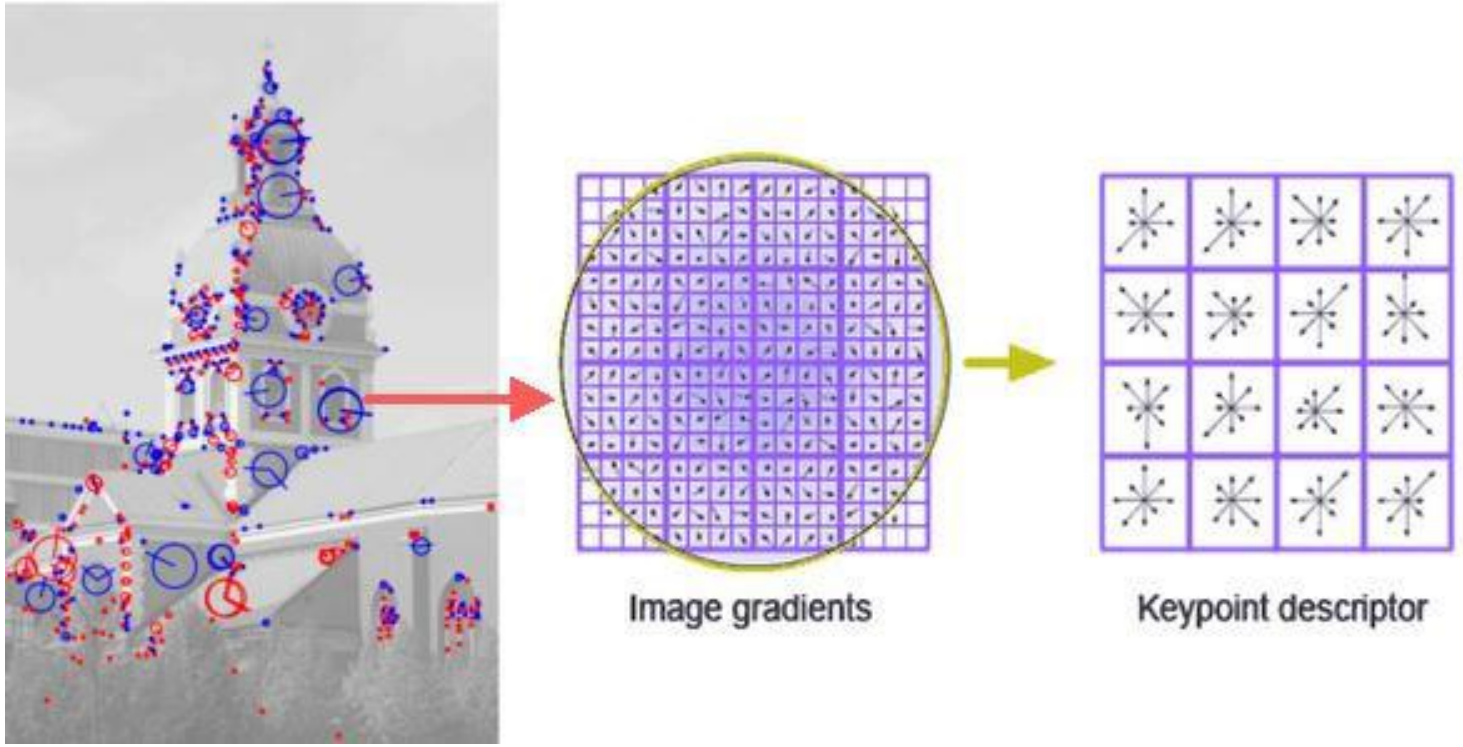

SIFT (Lowe 04)

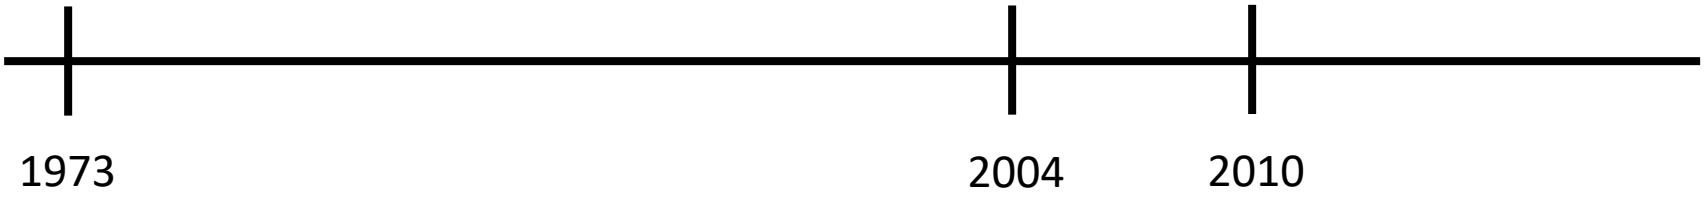

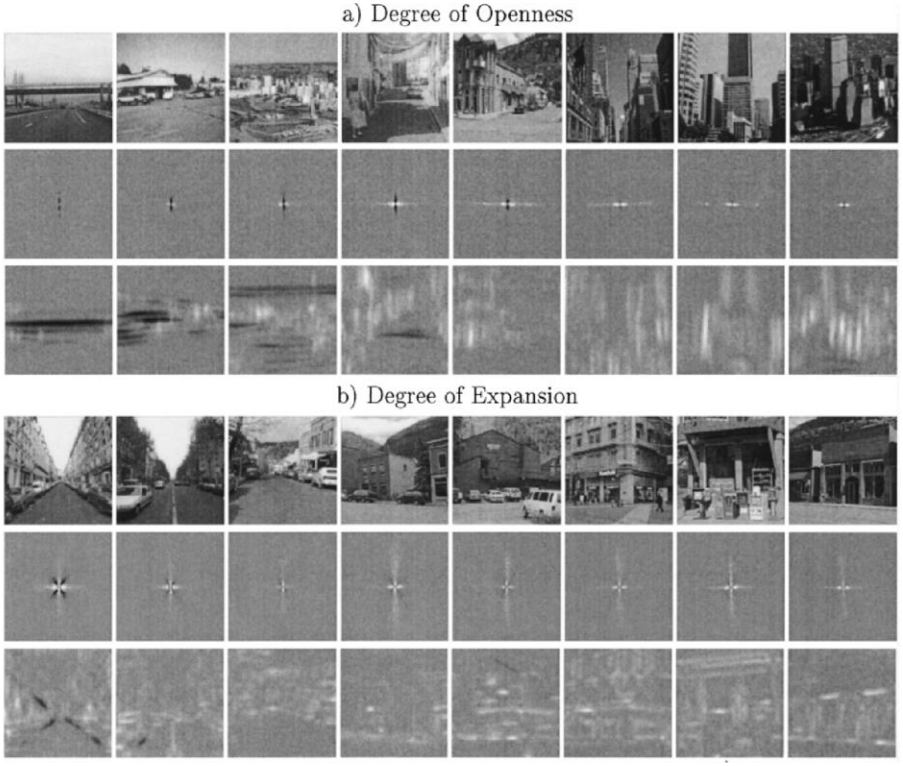

GIST (Oliva and Torralba 01)

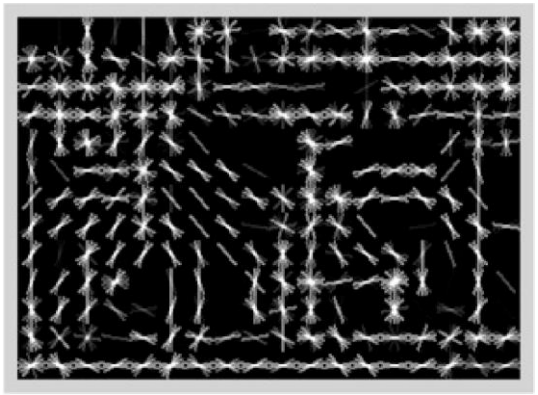

HOG (Dalal and Triggs 05)

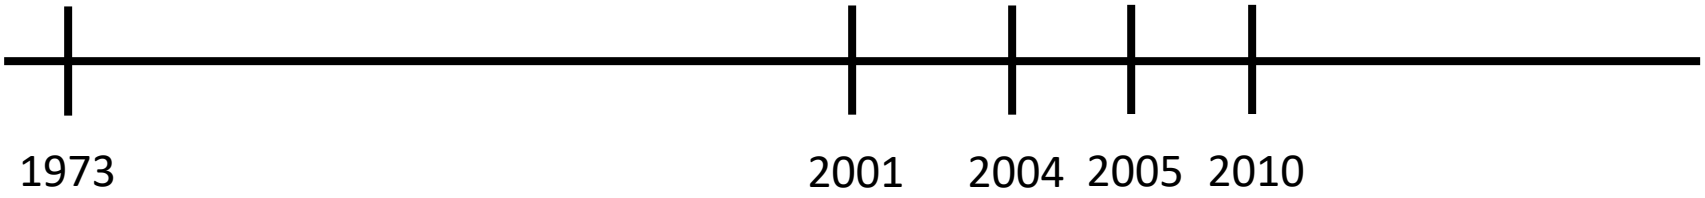

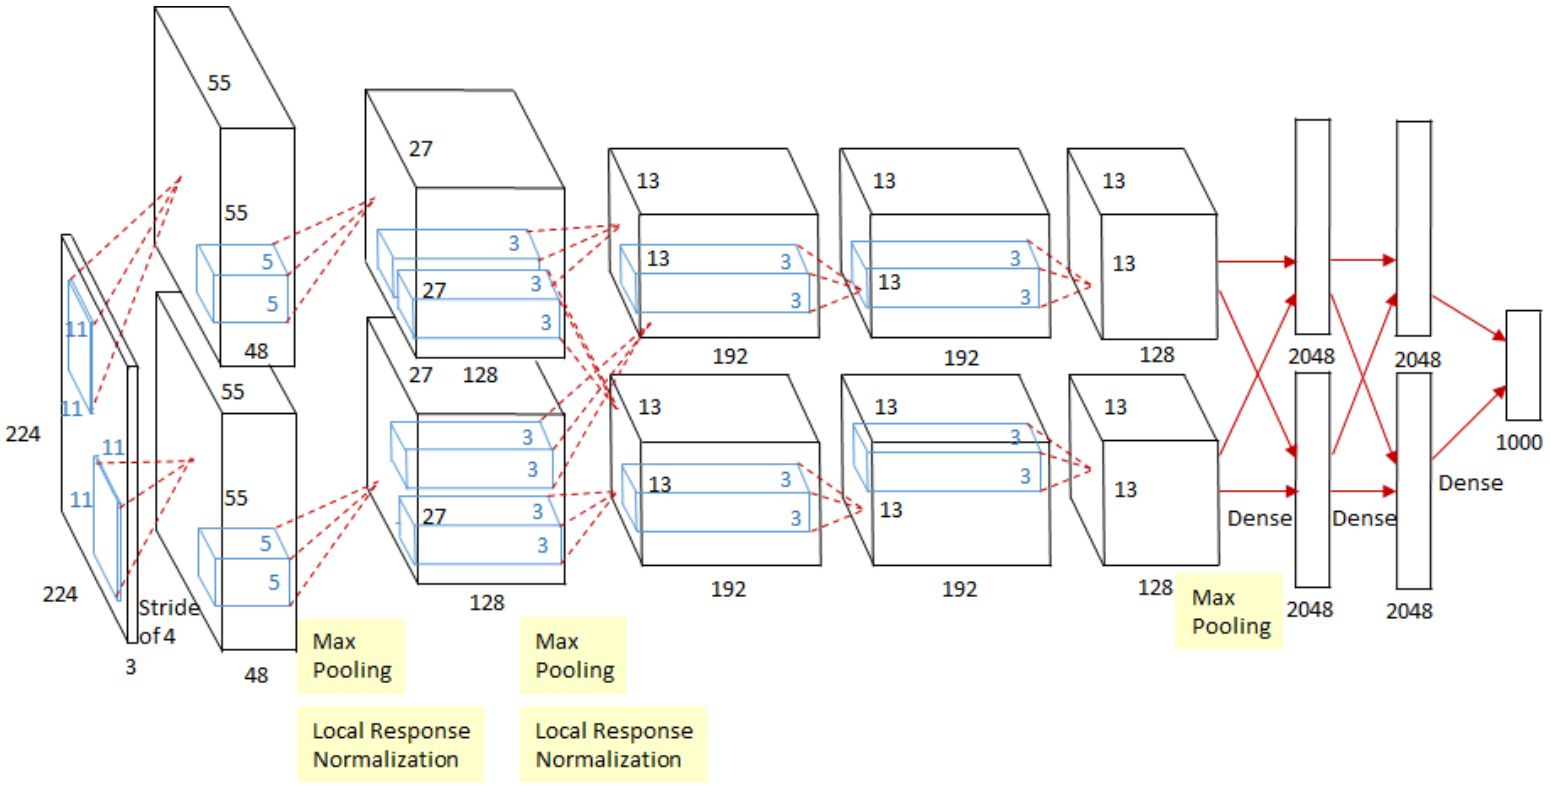

AlexNet (Krizhevsky et al. 12)

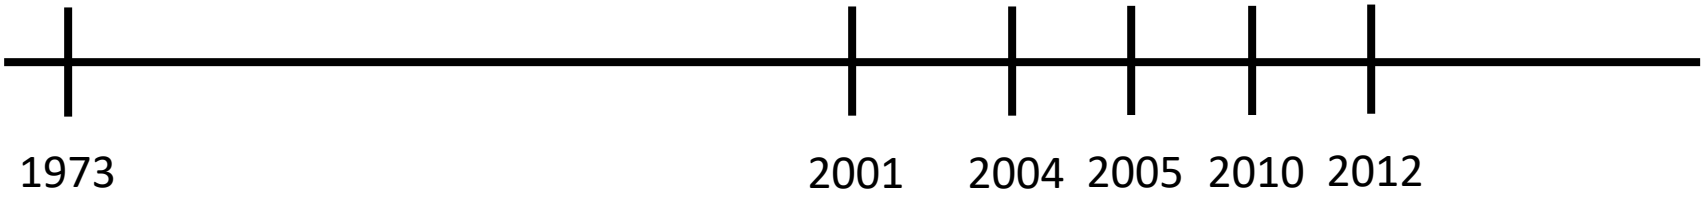

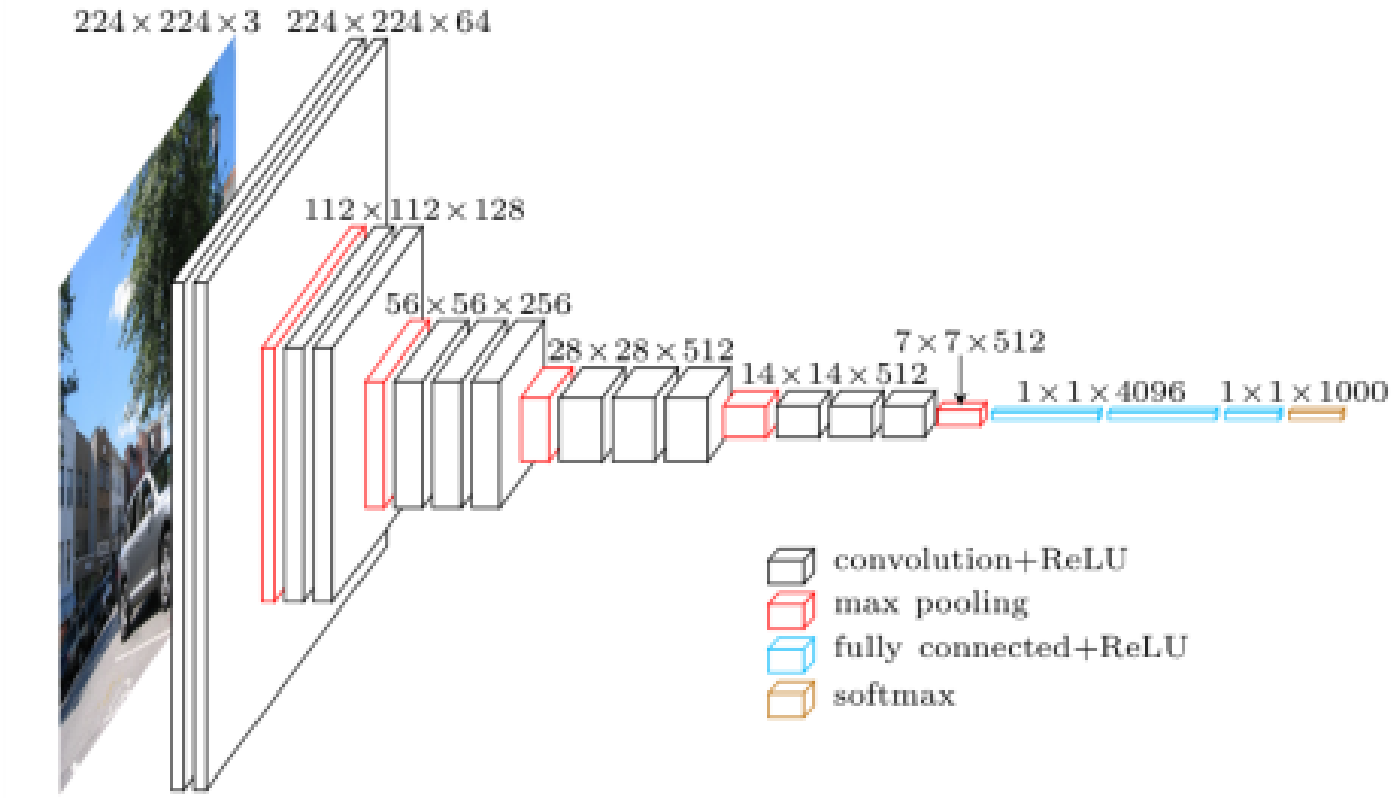

VGG19 (Simonyan and Zisserman 14)

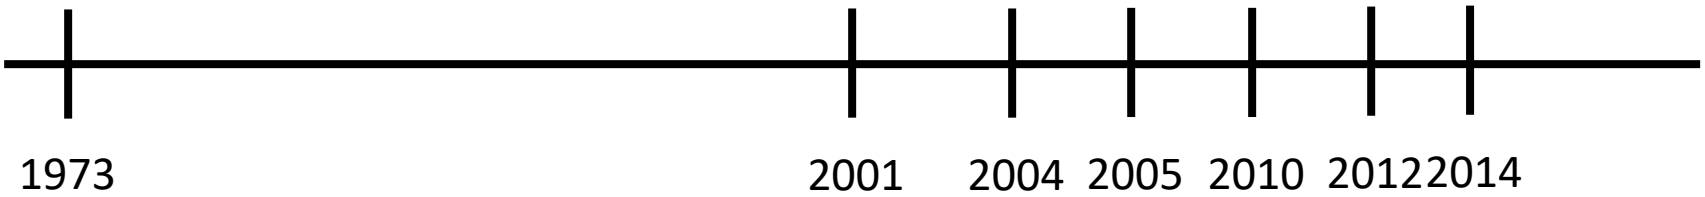

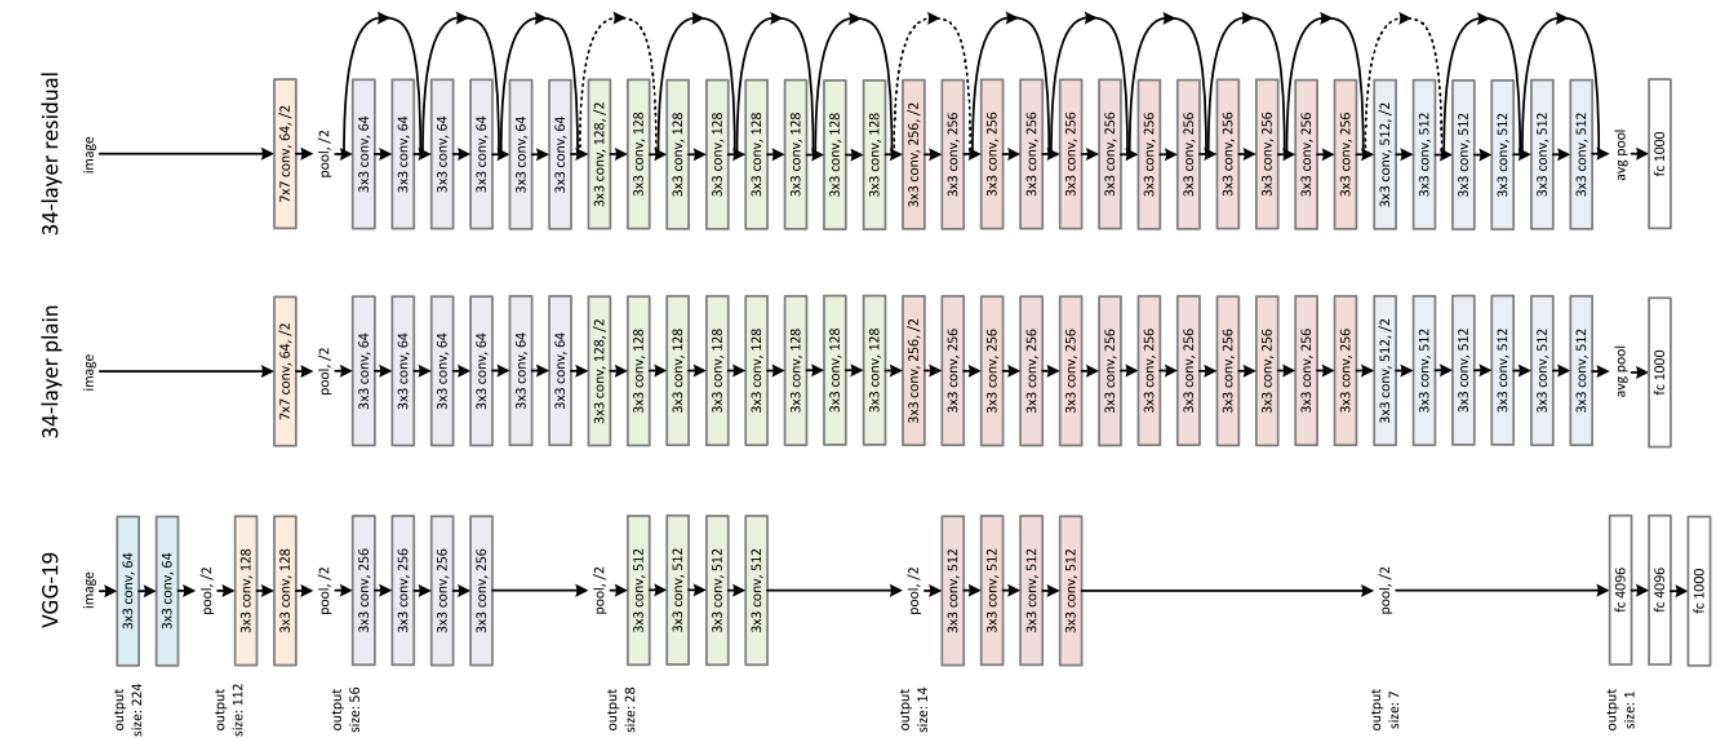

ResNet (He et al. 16)

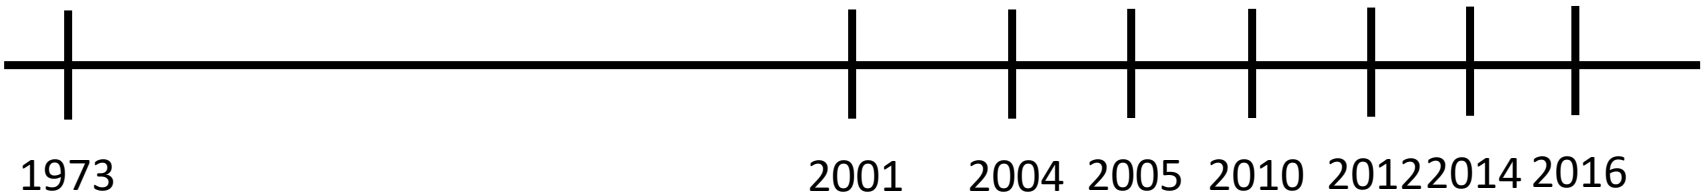

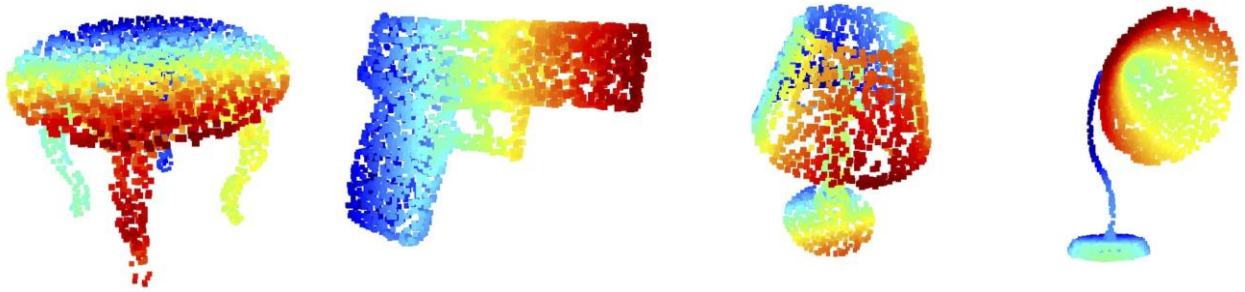

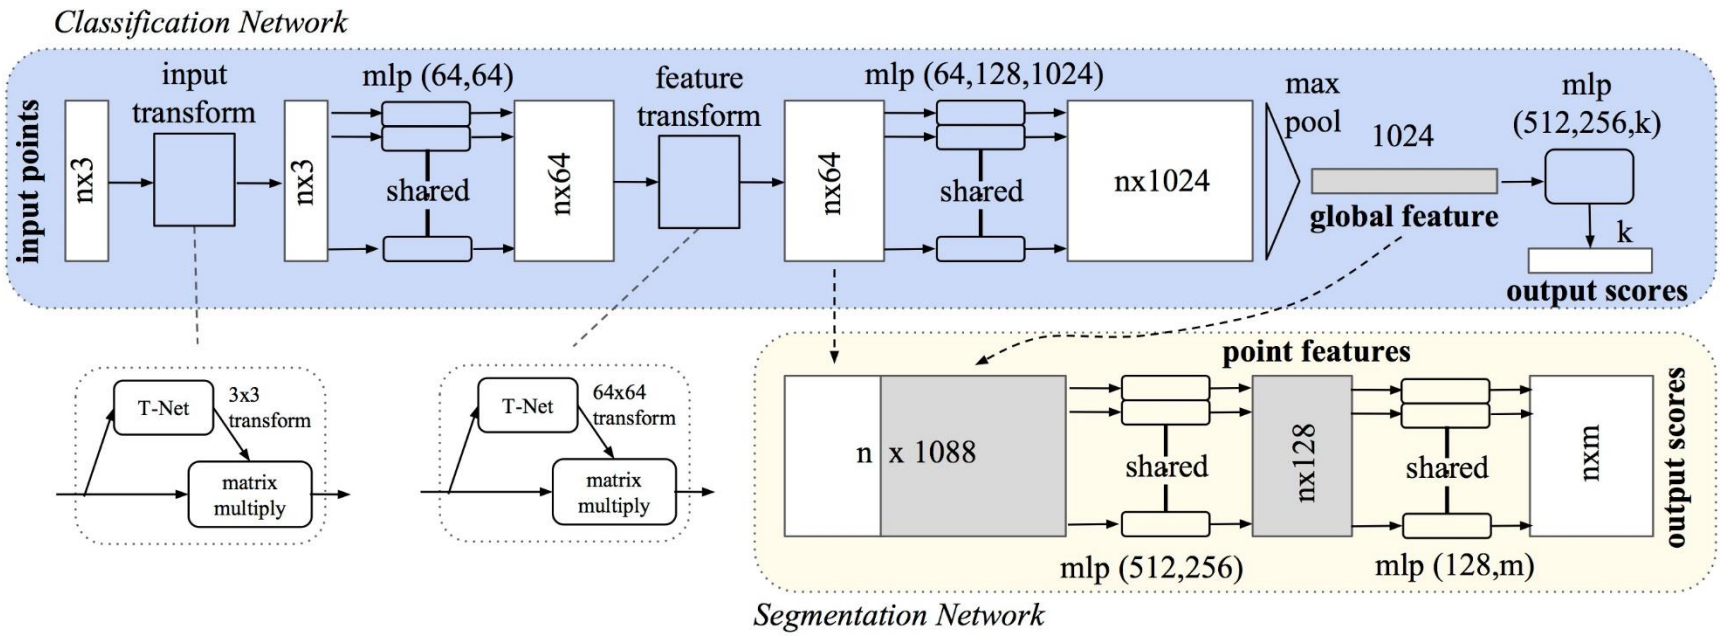

PointNet (R. Qi and Su et al. 17)

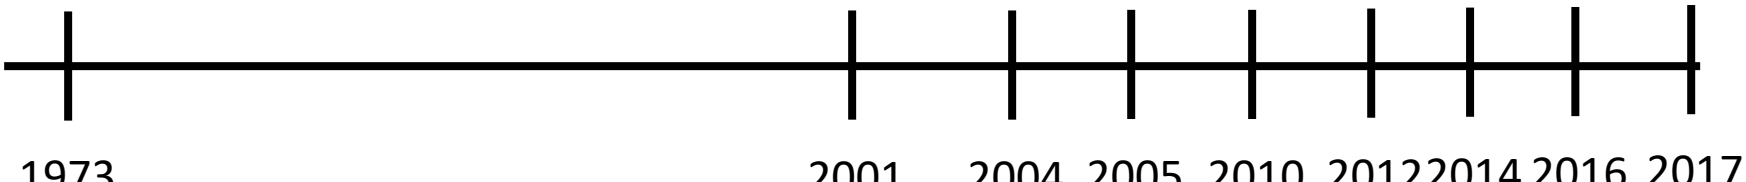

Machine Learning Algorithms

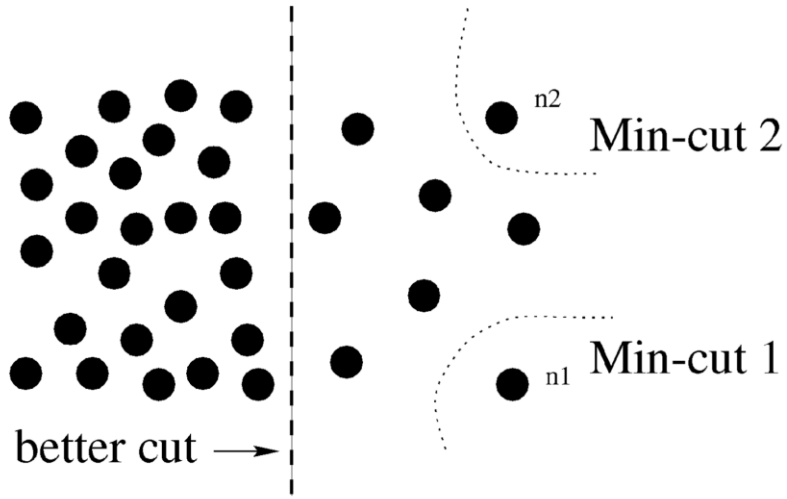

$$Ncut(A,B) = \frac{cut(A,B)}{assoc(A,V)} + \frac{cut(A,B)}{assoc(B,V)}$$

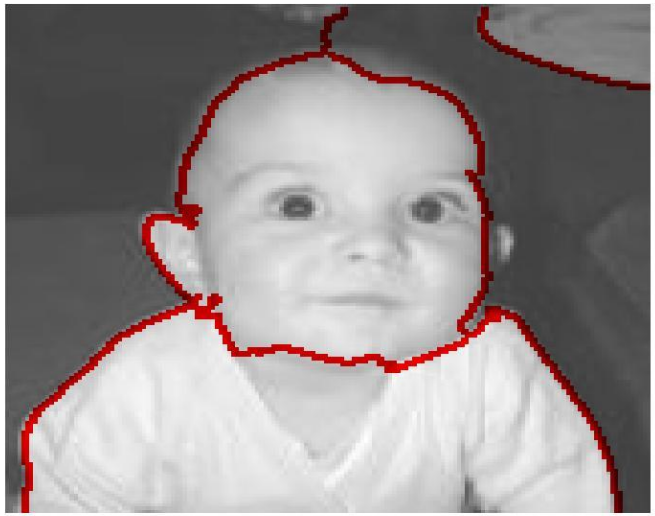

Normalized Cut (Shi and Malik 97)

1997

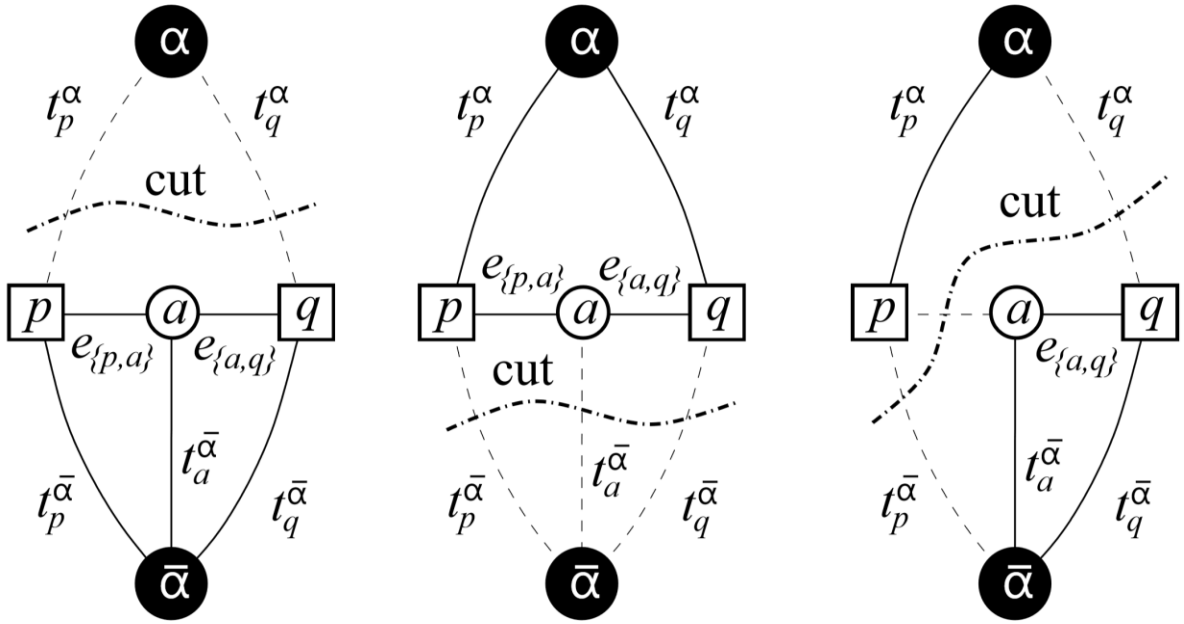

Graph Cut (Boykov et al. 99)

1997 1999

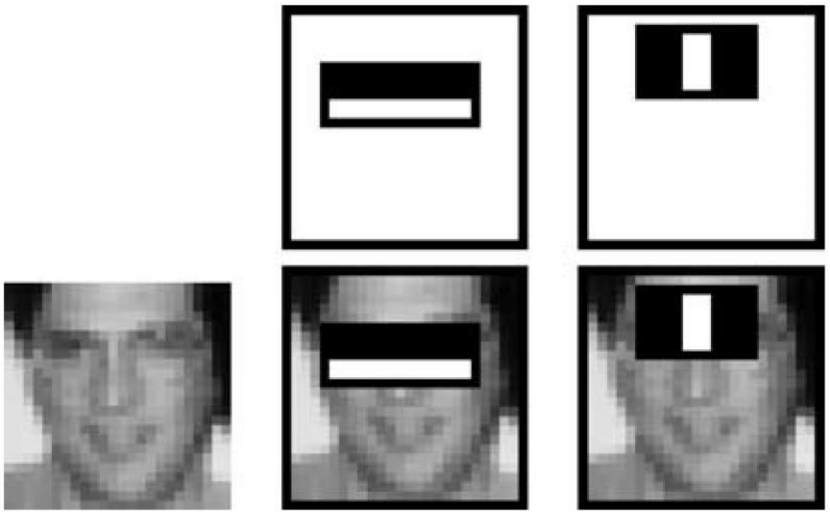

AdaBoosting for face detection
(Viola and Jones 04)

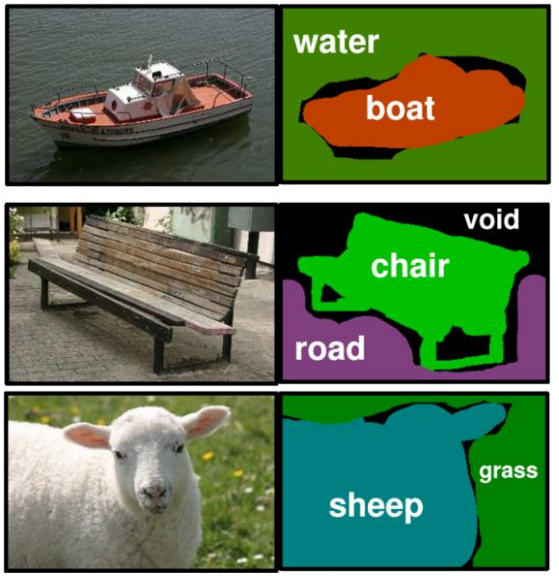

TextonBoost for segmentation
(Shotton et al. 06)

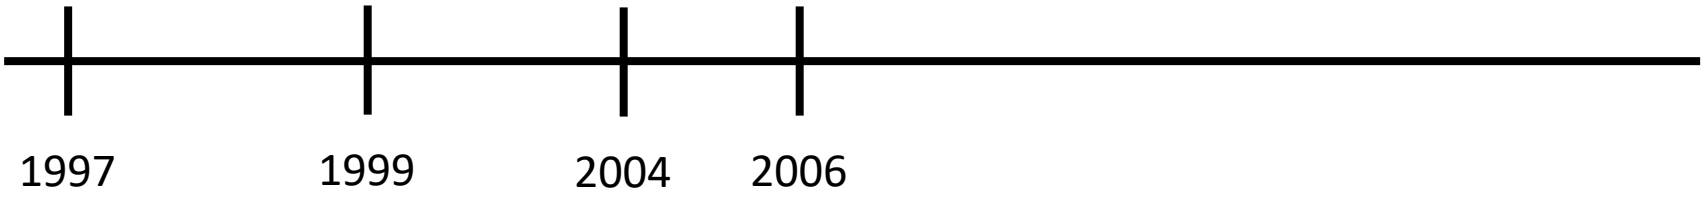

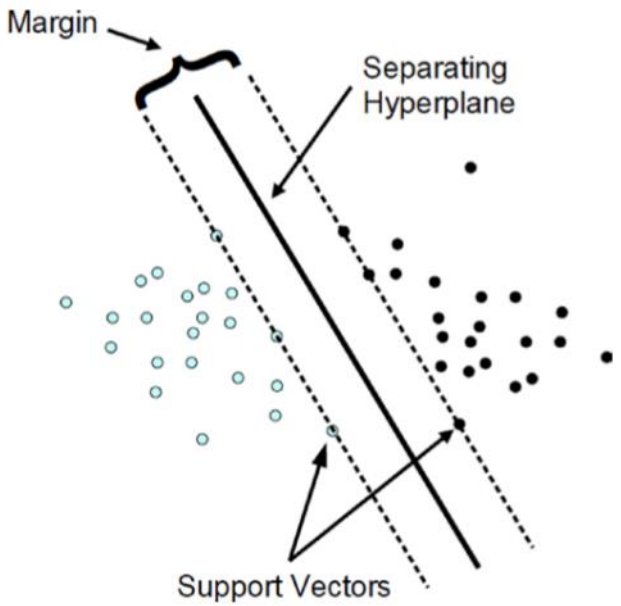

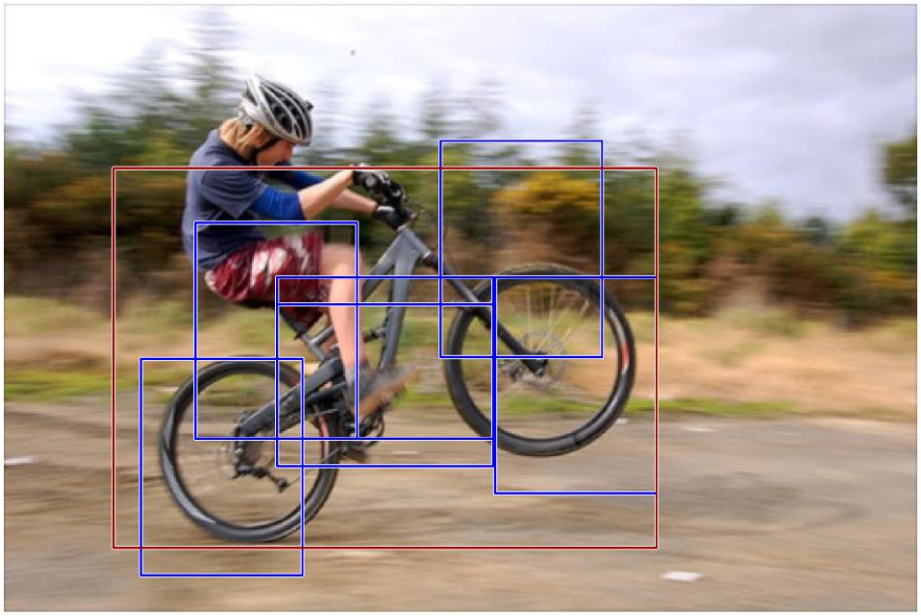

Support Vector Machine in Deformable Part Model
(Felzenszwalb et al. 10)

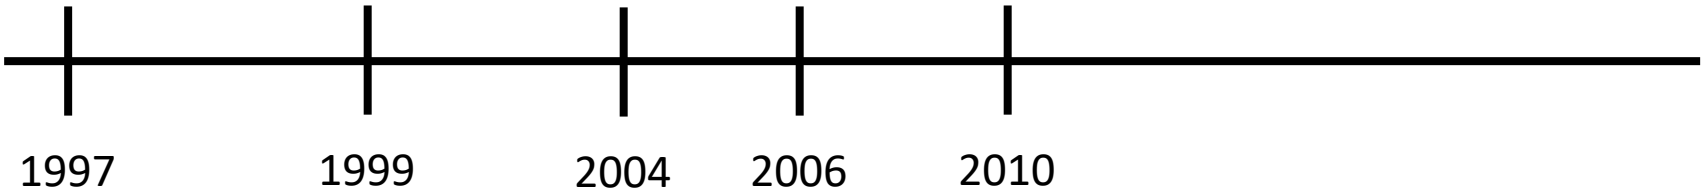

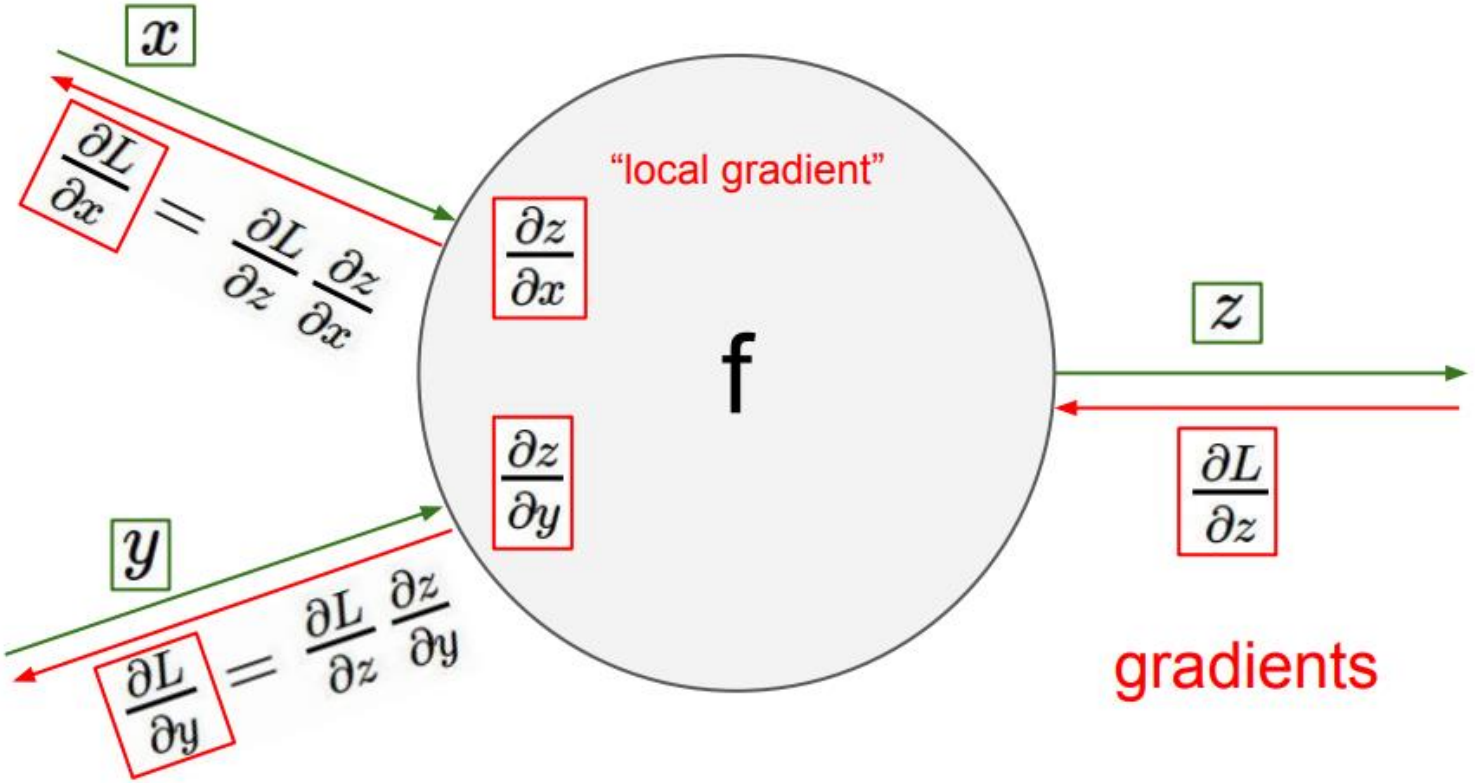

Back-propagation in neural network training/implementation
(Rumelhart et al. 86, LeCun et al. 98, Abadi et al. 16)

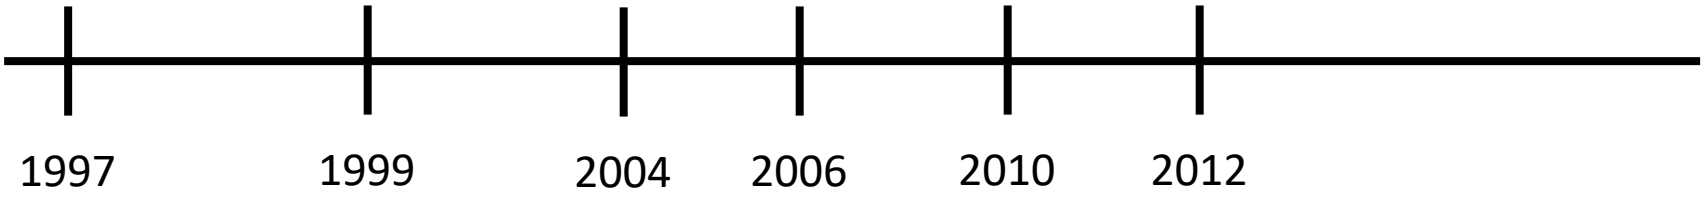

Optimization

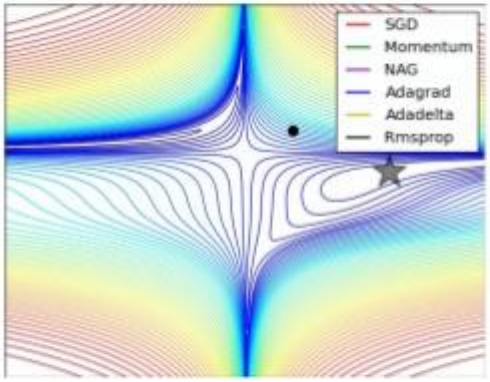

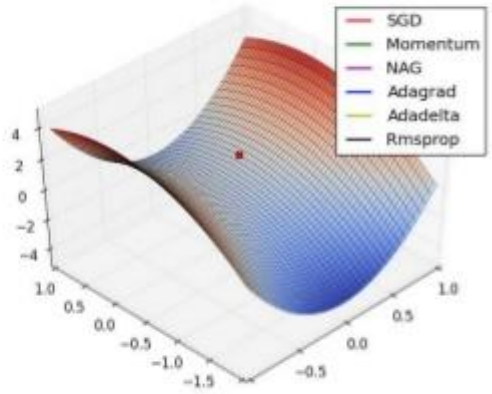

Adam: A method for stochastic optimization
(Kingma and Ba 14)

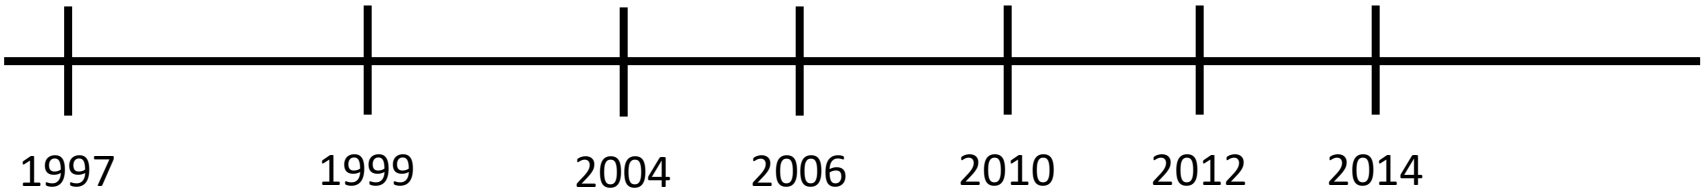

Goals of this course

- Upper division undergrad course
- Introduction to primary topics
  - Fundamentals of computer vision- image processing, grouping, multiple views
  - Recognition – emphasis on machine learning (the most active research area in computer vision)
- Hands-on experience with algorithms
- Basic knowledge of the research field of computer vision

Topics overview

- Features & filters
- Grouping & fitting
- Multiple views
- Recognition

Features and filters

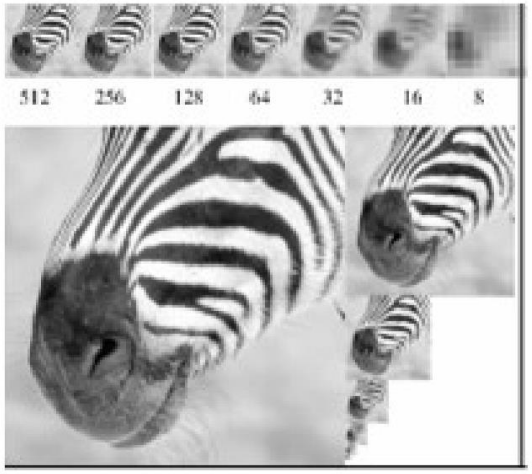

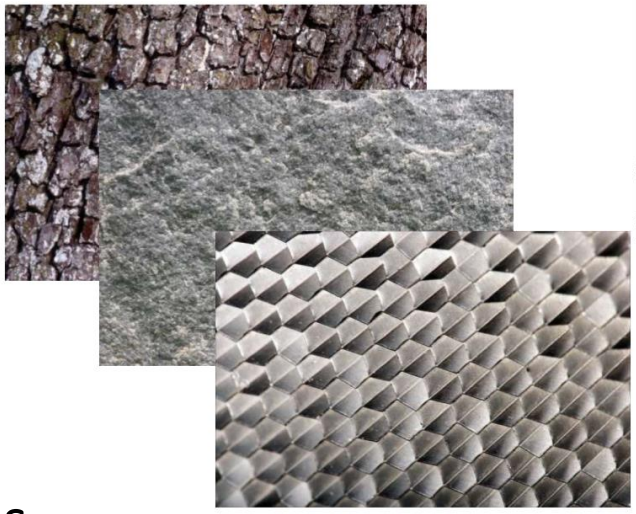

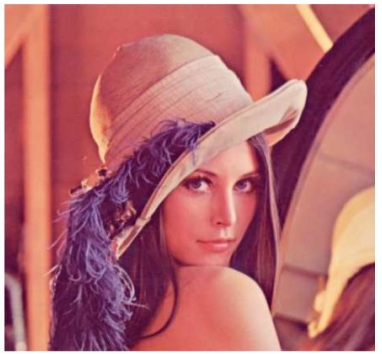

Transforming and describing images; Textures, colors, edges

Building blocks for neural networks

Grouping & fitting

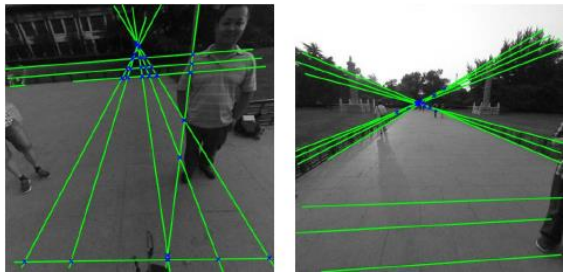

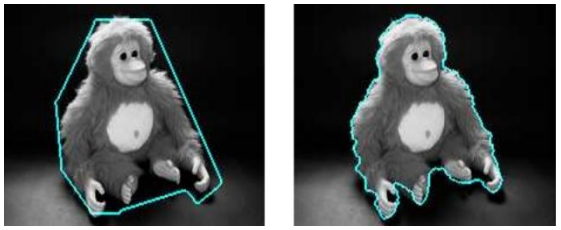

Shi et al.

Clustering, Segmentation, fitting; what parts belong together?

Multiple Views

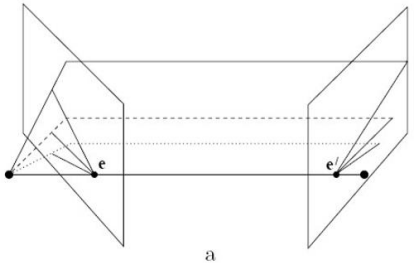

a

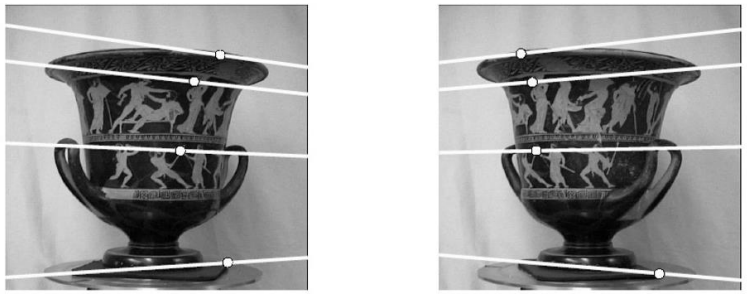

Hartley and Zisserman

Invariant features, matching
Epipolar geometry
Structure-from-motion, stereo

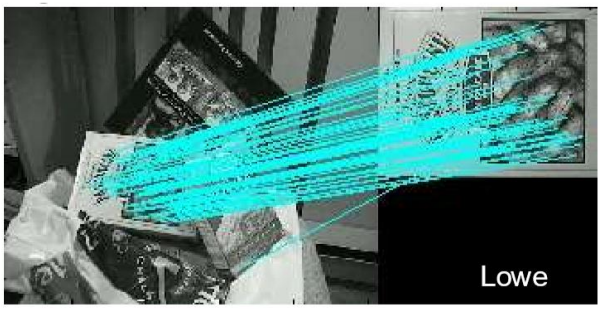

Recognition and learning

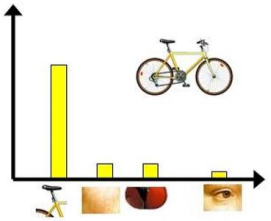

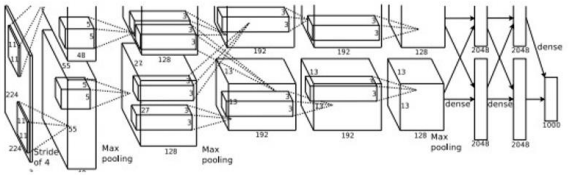

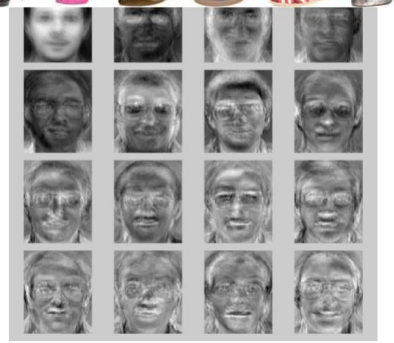

Recognizing categories (objects, scenes, activities, attributes…), learning techniques

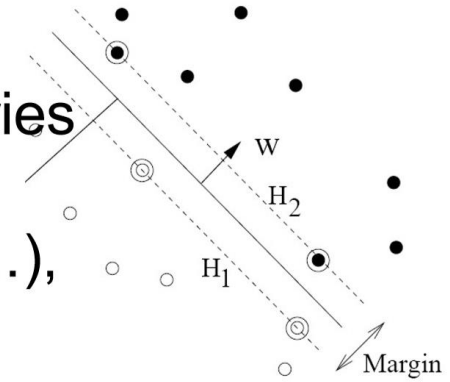

Data representation (vectorized) -> machine learning techniques

Textbooks

- Recommend book:
  - Computer Vision: Algorithms and Applications
  - By Rick Szeliski
  - http://szeliski.org/Book
  - Online

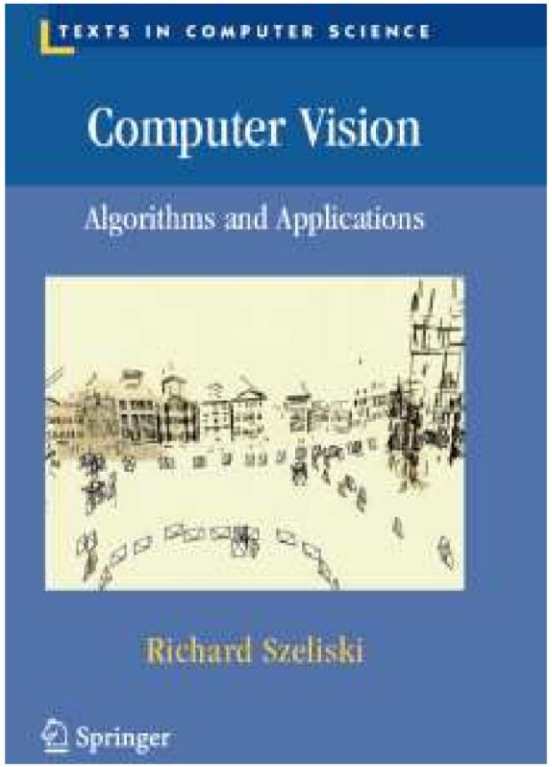

Assignments

- Majority - Programming problem
  - Implementation
  - Explanation, results
- Code in Matlab – available on CS Unix machines (see course page)
- Optional Latex templates
- Most of these assignments take significant time to do. We recommend starting early

Matlab

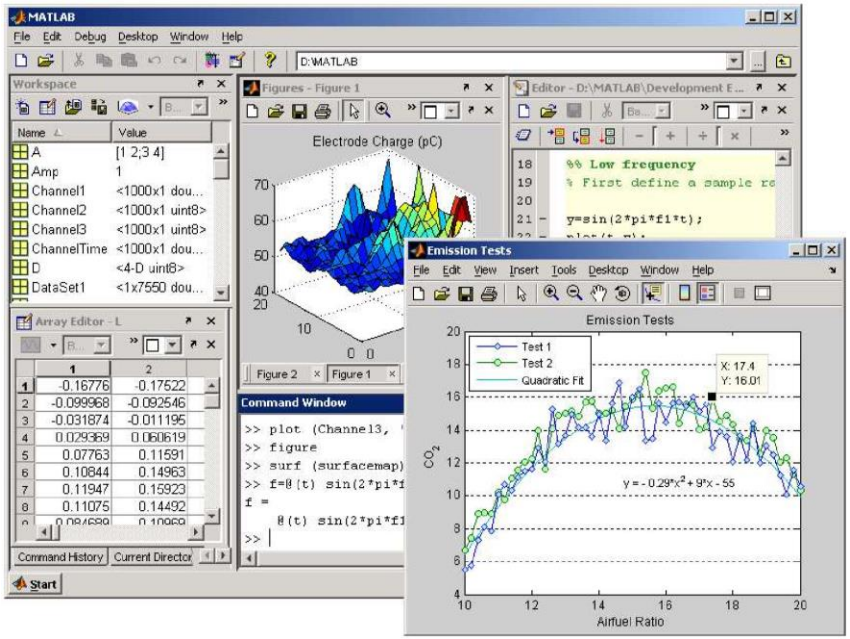

- Built-in toolboxes for low-level image processing, visualization
- Compact programs
- Intuitive interactive debugging
- Widely used in engineering
- I use it a lot!

Assignment 0

- A0: Matlab warmup + basic image manipulation
- Out today, due Tues Jan 29

- Verify CS account and Matlab access
- Look at the tutorial online

Preview of assignments

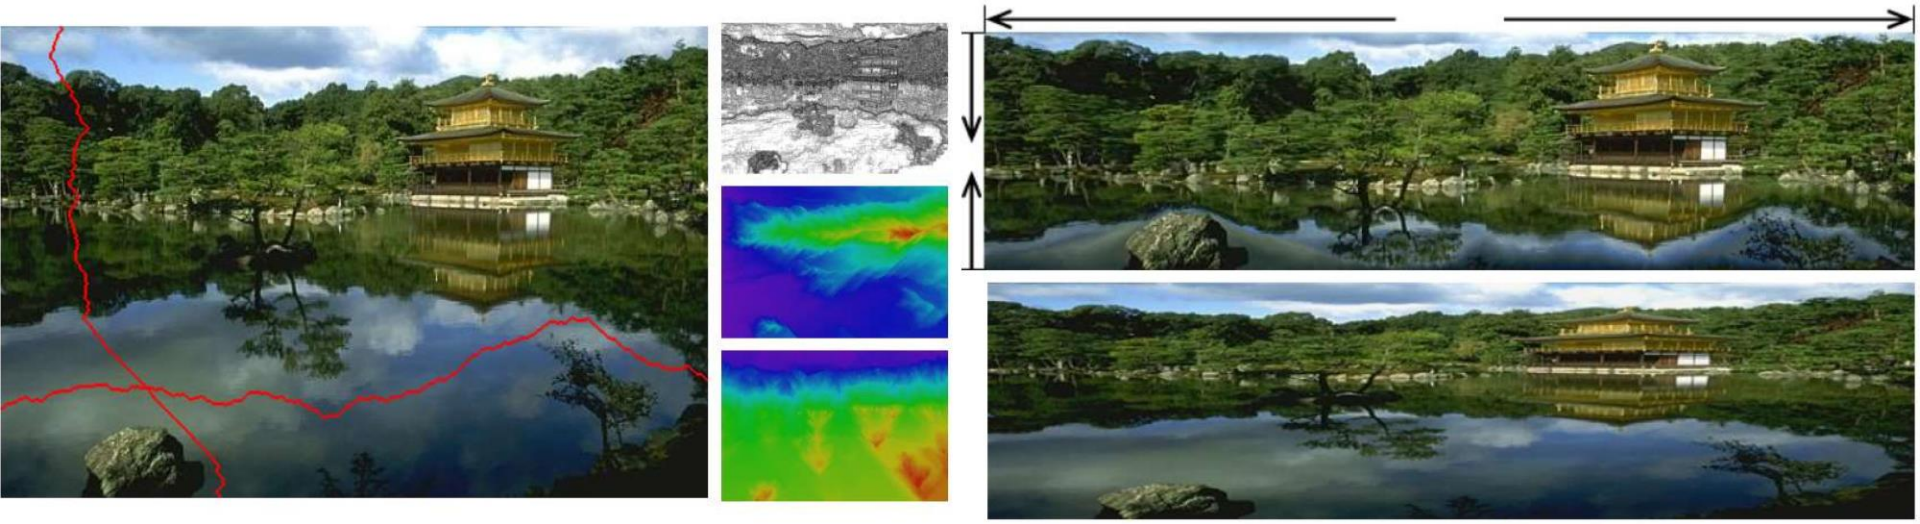

Seam carving

Preview of assignments

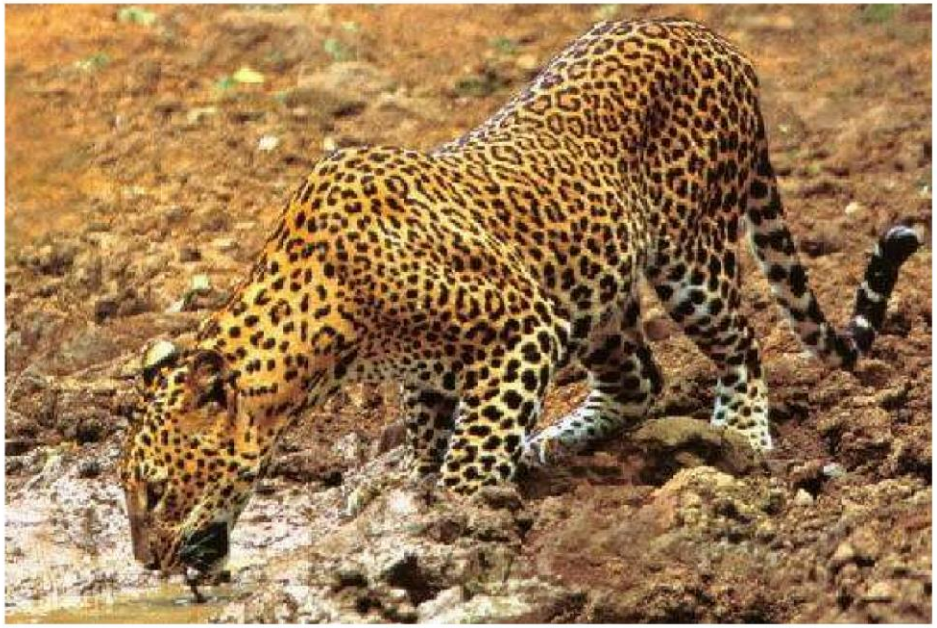

Grouping for segmentation

Preview of assignments

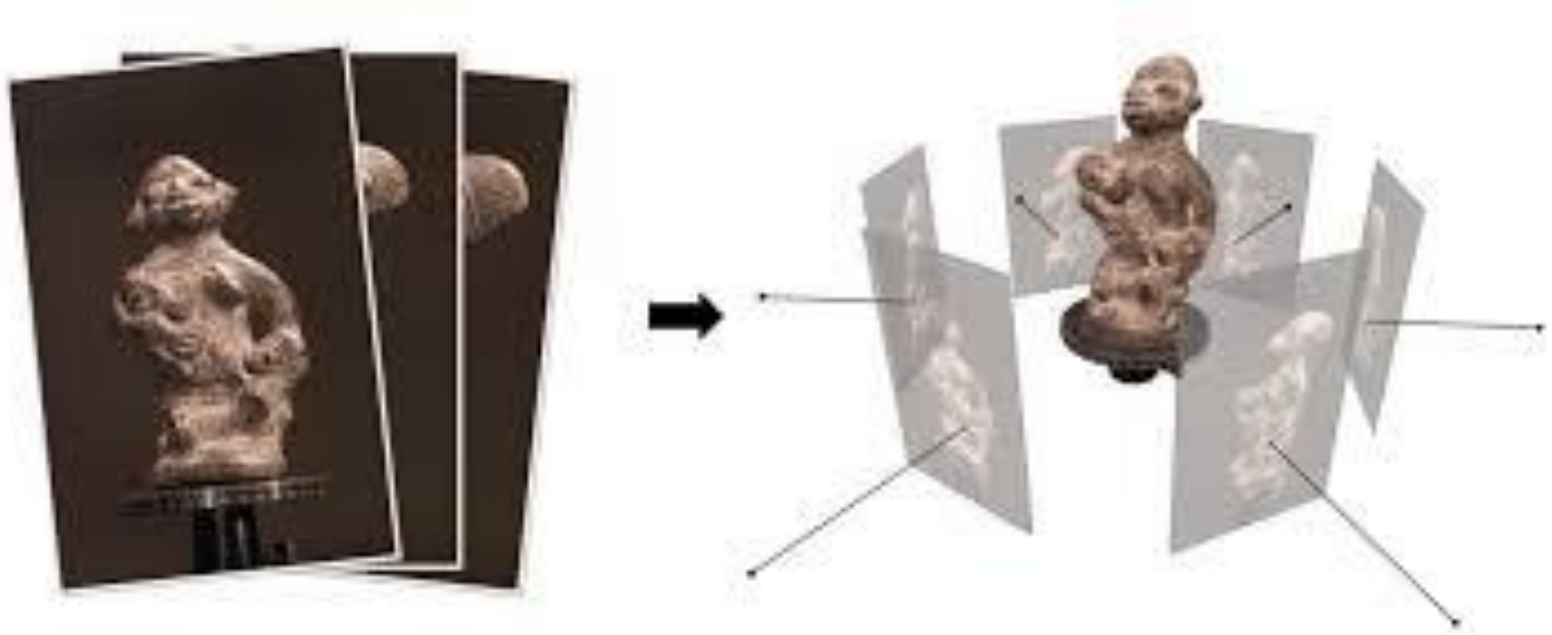

Structure-from-motion and stereo

Preview of assignments

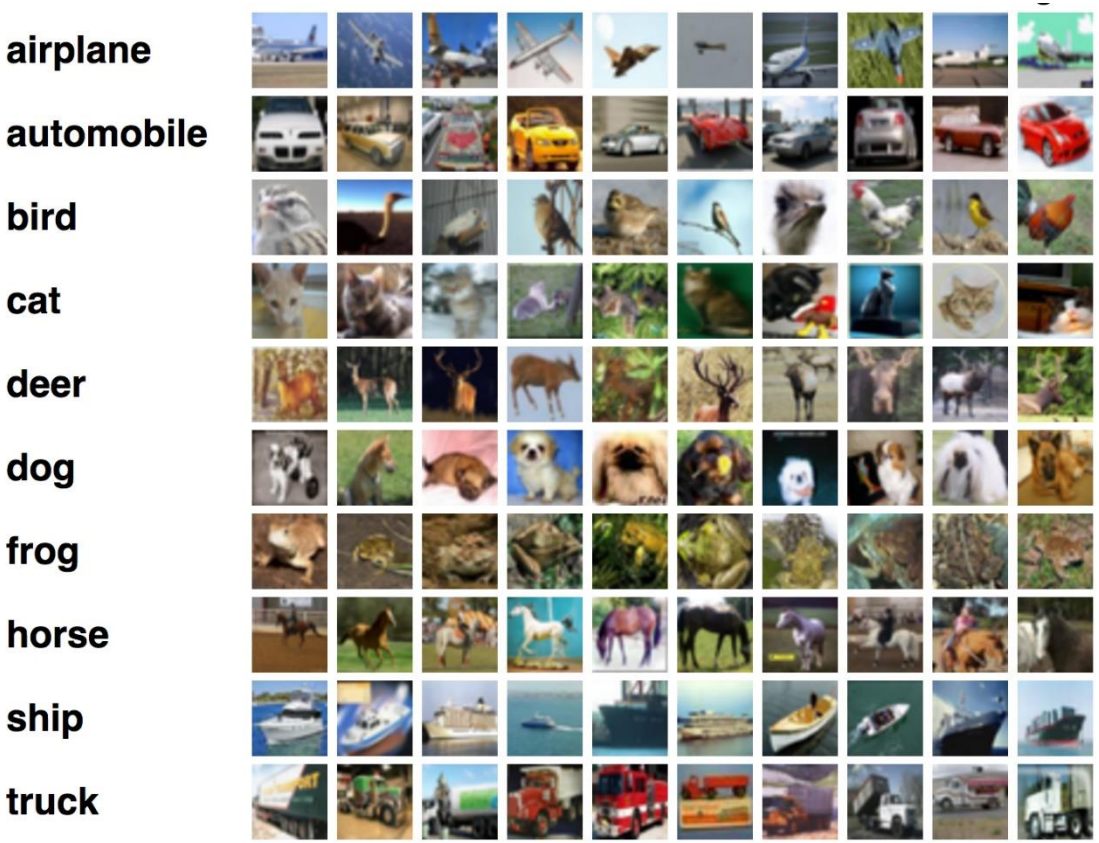

Image classification

Preview of assignments

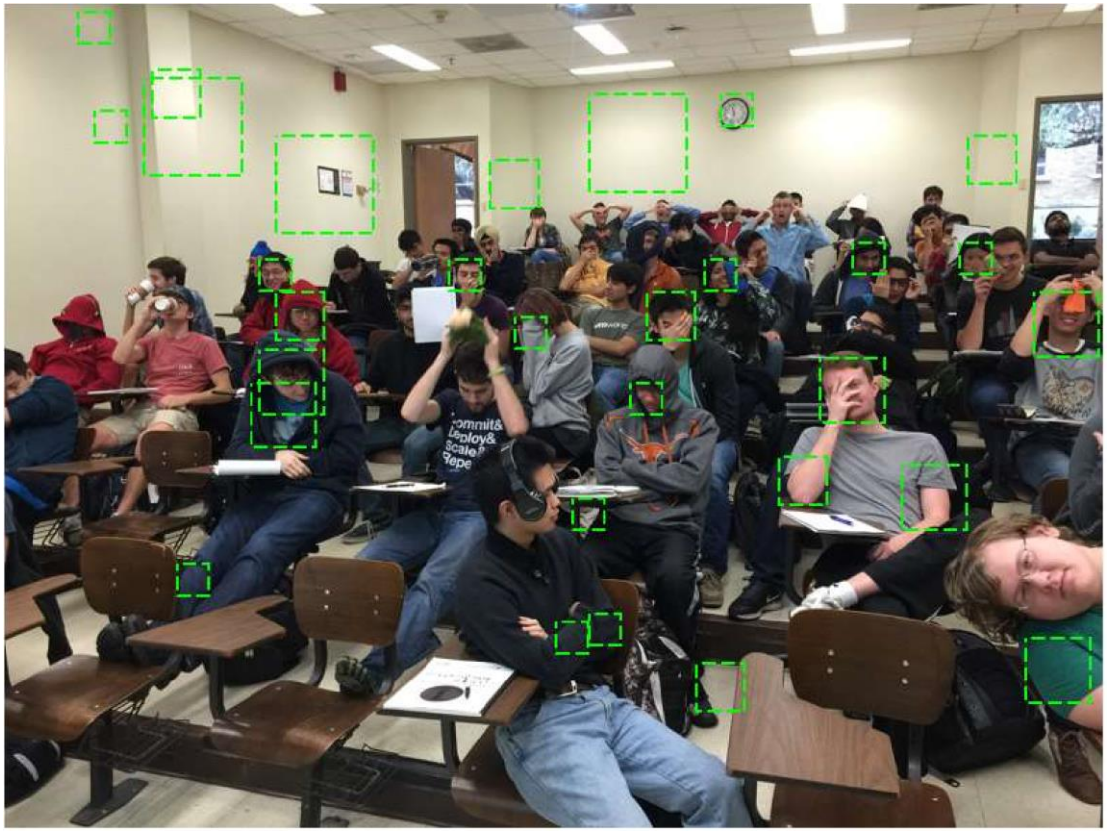

Object detection

Collaboration policy

- All responses and code must be written individually unless otherwise specified

- Students submitting answers or code found to be identical or substantially similar (due to inappropriate collaboration) risk failing the course.

Assignment deadlines

- Due about every two weeks
  - Tentative deadlines posted online but could slightly shift depending on lecture pace
- Assignments in by 11:59 pm on due date
  - Submit on Canvas, following submission instructions given in assignment
  - Deadlines are firm. We will use timestamp

Miscellaneous

- Slides, announcements via class website
- No laptops, phones, tablets, etc. open in class please
- Please use the front rows

Coming up

- Now: check out Matlab tutorial online
- A0 due Tuesday Jan. 29

- Textbook reading posted for next week