Plan for today

- Topic overview
- Introductions
- Course overview:
  - Logistics and requirements

2

What is Computer Vision?

3

Computer Vision

Enable machines to "see" the visual world as we do

4

Computer Vision

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

5

1. Vision for measurement

Real-time stereo

Wang et al.

Structure from motion

Snavely et al.

Tracking

Demirdjian et al.

Slide credit: Kristen Grauman

6

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)*

7

2. Vision for perception, interpretation

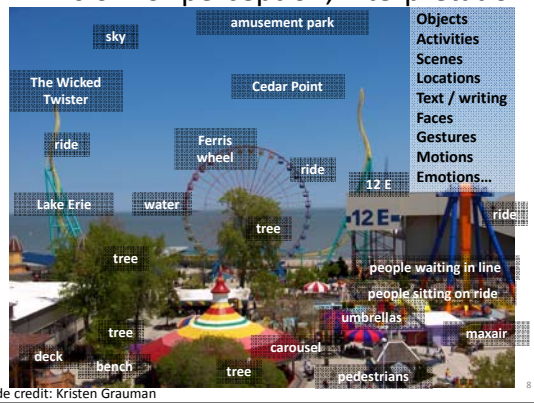

Slide credit: Kristen Grauman

8

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*)

9

3. Visual search, organization

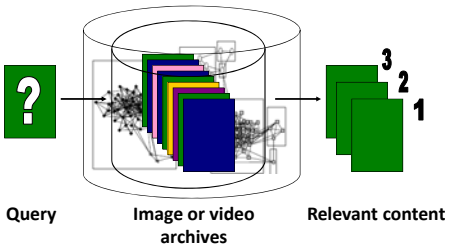

Slide credit: Kristen Grauman

10

Related disciplines

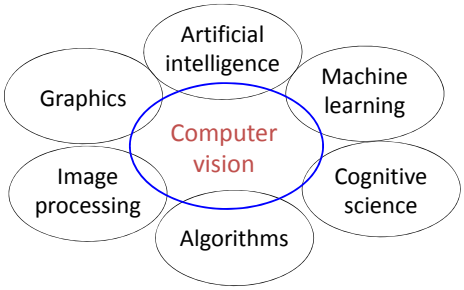

Slide credit: Kristen Grauman

11

Vision and graphics

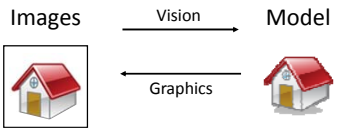

Inverse problems: analysis and synthesis

Slide credit: Kristen Grauman

12

Why is vision difficult?

13

What humans see

14

What computers see

Slide credit: Larry Zitnick

15

Why is vision difficult?

- Ill-posed problem: real world much more complex than what we can measure in images
  - 3D → 2D
- Impossible to literally "invert" image formation process

Slide credit: Kristen Grauman

16

Challenges: ambiguity

- Many different 3D scenes could have given rise to a particular 2D picture

Slide credit: Svetlana Lazebnik

17

Challenges: many nuisance parameters

Illumination

Object pose

Clutter

Occlusions

Intra-class appearance

Viewpoint

Slide credit: Kristen Grauman

18

Challenges: scale

slide credit: Fei-Fei, Fergus, Torralba

19

Challenges: Motion

slide credit: Svetlana Lazebnik

20

Challenges: occlusion, clutter

slide credit: Svetlana Lazebnik

21

Challenges: object intra-class variation

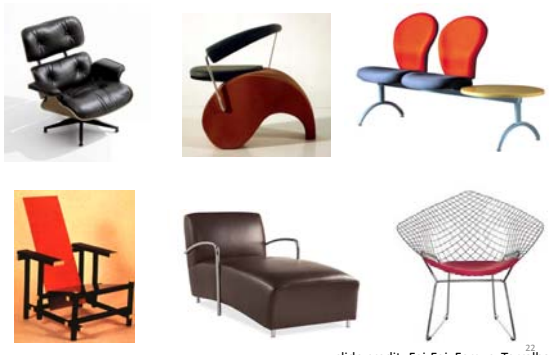

slide credit: Fei-Fei, Fergus, Torralba

22

Challenges: context and human experience

Slide credit: Fei-Fei, Fergus, Torralba

23

Challenges: context and human experience

Fei Fei Li, Rob Fergus, Antonio Torralba

24

Challenges: context and human experience

Fei Fei Li, Rob Fergus, Antonio Torralba

25

Challenges: complexity

How many object categories are there?

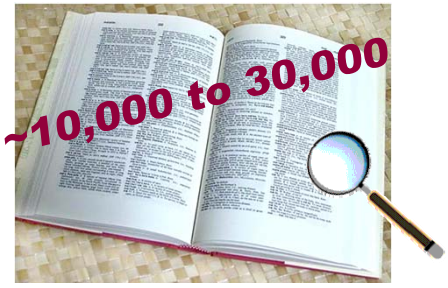

Biederman 1987

Slide credit: Fei-Fei, Fergus, Torralba

26

Challenges: complexity

flickr
10 billion images

facebook
250 billion images

imgur
1 billion images served daily

YouTube
100 hours uploaded per minute

photobucket
10 billion images

From cisco:

Almost 90% of web traffic is visual!

27

Challenges: complexity

- Thousands to millions of pixels in an image
- 30+ degrees of freedom in the pose of articulated objects (humans)
- About half of the cerebral cortex in primates is devoted to processing visual information [Felleman and van Essen 1991]

Slide credit: Kristen Grauman

28

What works well today?

29

Optical character recognition (OCR)

Digit recognition
yann.lecun.com

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

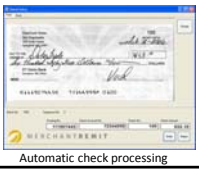

Automatic check processing

Sudoku grabber
http://sudokugrab.blogspot.com/

30

Source: S. Seitz, N. Snavely

Image classification and object detection

Image classification

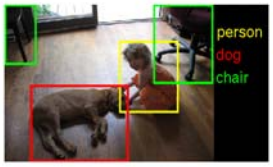

Object detection

31

Biometrics

Fingerprint scanners

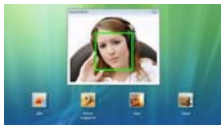

Face recognition systems

32

Face detection

Source: S. Seitz

33

Face detection for privacy protection

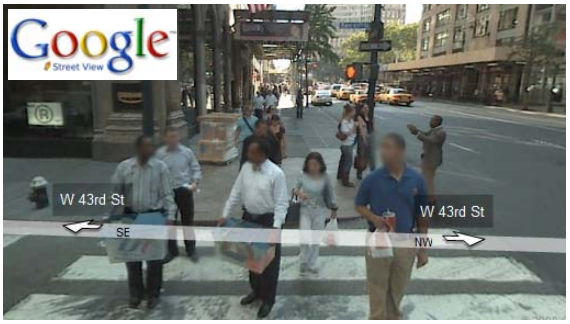

34

slide credit: Svetlana Lazebnik

Technology gone wild…

35

slide credit: Svetlana Lazebnik

Face recognition

Slide credit: Devi Parikh

36

Face anonymization

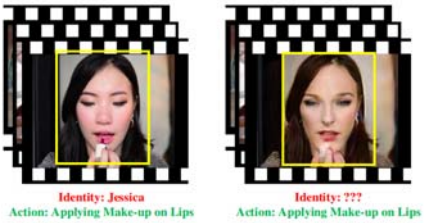

Zhongzheng Ren, Yong Jae Lee, and Michael Ryoo.
Learning to Anonymize Faces for Privacy Preserving Action Detection. arXiv 2018.

37

Interactive systems

Shotton et al.

38

Instance recognition

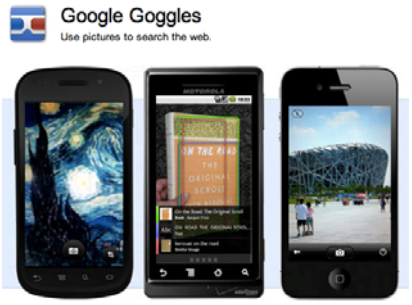

Slide credit: Devi Parikh

39

Pedestrian detection

Slide credit: Devi Parikh

40

Autonomous agents

Mars rover

Self-driving car

41

3D reconstruction from photo collections

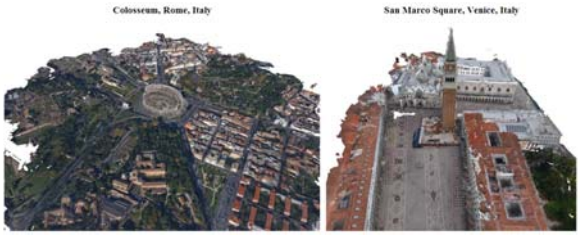

Q. Shan, R. Adams, B. Curless, Y. Furukawa, and S. Seitz, The Visual Turing Test for Scene Reconstruction, 3DV 2013

slide credit: Svetlana Lazebnik

42

Special effects: shape capture


The Matrix movies, ESC Entertainment, XYZRGB, NRC

Source: S. Seitz

Special effects: motion capture

Pirates of the Carribean, Industrial Light and Magic

Source: S. Seitz

Medical imaging

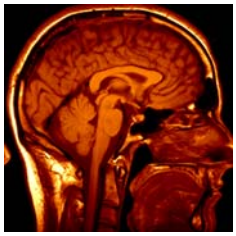

3D imaging
MRI, CT

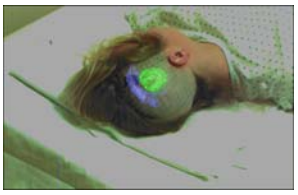

Image guided surgery
Grimson et al., MIT

Source: S. Seitz

Visual data in 1963

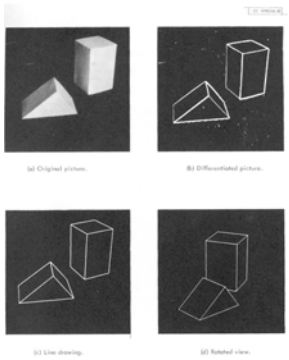

L. G. Roberts, *Machine Perception of Three Dimensional Solids*, Ph.D. thesis, MIT Department of Electrical Engineering, 1963.

Slide credit: Kristen Grauman

46

Visual data today

Svetlana Lazebnik

47

Why vision?

- As image sources multiply, so do applications
  - Relieve humans of boring, easy tasks
  - Enhance human abilities
  - Advance human-computer interaction, visualization
  - Perception for robotics / autonomous agents
  - Organize and give access to visual content

Slide credit: Kristen Grauman

48

Applications

- Law enforcement / Surveillance
- Robotics
- Autonomous driving
- Medical imaging
- Photo organization
- Image search
- E-commerce
- … cell phone cameras, social media, Google Glass, etc.

Slide adapted from Devi Parikh

49

Summary

- Computer Vision is useful, interesting, and difficult
- A growing and exciting field
- Lots of cool and important applications
- New teams in existing companies, startups, etc.

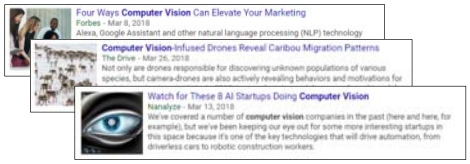

50

Introductions

- Instructor
  - Yong Jae Lee
  - yongjaelee@ucdavis.edu
  - Assistant Professor in CS, UC Davis since July 2014
  - Research areas: Computer vision and machine learning
    - Visual Recognition
    - Graphics Applications

51

Introductions

- TAs:
  - Chongruo Wu
  - crwu@ucdavis.edu
  - PhD student in CS

  - Maheen Rashid
  - mhnrashid@ucdavis.edu
  - PhD student in CS

52

This course

- ECS 174 (4-units)
- Lecture: Tues & Thurs 12:10-1:30 pm, Giedt Hall 1002
- Discussion section: Mon 5:10-6pm, Chem 179 (*location may change*)
- Office hours: Academic Surge 1044/2075
  - Maheen: Tues 10 am-noon (AS 1044)
  - Chongruo: Thurs 3-5 pm (AS 1044)
  - Yong Jae: Fri 3-5 pm (AS 2075)

53

This course

- Course webpage
  https://sites.google.com/a/ucdavis.edu/ecs-174-computer-vision---spring-2018/
- Canvas (assignment submission, grades)
  https://canvas.ucdavis.edu/courses/225624
- Piazza
  piazza.com/uc_davis/spring2018/ecs174/

54

Goals of this course

- Introduction to primary topics in Computer Vision
- Basics and fundamentals
- Practical experience through assignments
- Views of computer vision as a research area

55

Prerequisites

- Upper-division undergrad course
- Basic knowledge of probability and linear algebra
- Data structures, algorithms
- Programming experience
- Experience with image processing or Matlab will help but is not necessary

56

Topics overview

- Features and filters
- Grouping and fitting
- Recognition and learning

Focus is on algorithms, rather than specific systems

57

Features and filters

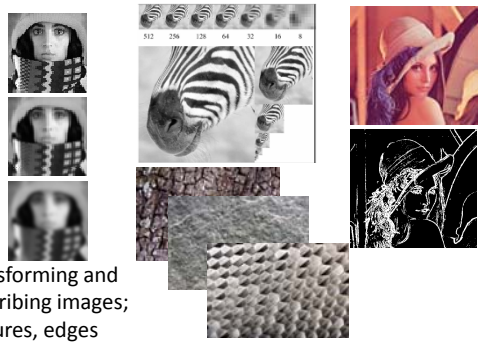

Transforming and describing images; textures, edges

Slide credit: Kristen Grauman

58

Grouping and fitting

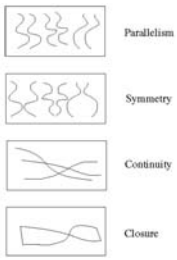

[fig from Shi et al]

Clustering, segmentation, fitting; what parts belong together?

Slide credit: Kristen Grauman

59

Recognition and learning

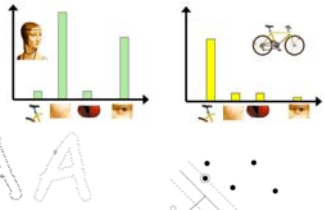

Recognizing objects and categories, learning techniques

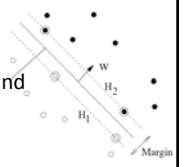

Slide credit: Kristen Grauman

60

Recognition and learning

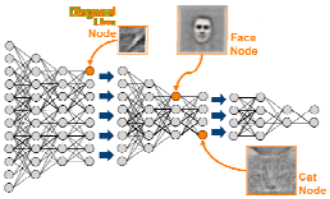

Deep learning

61

Not covered: Multiple views and motion

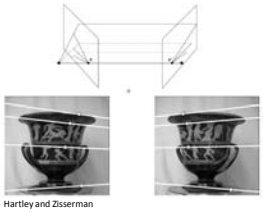

Hartley and Zisserman

Multi-view geometry, stereo vision

Fei-Fei Li

Slide credit: Kristen Grauman

62

Not covered: Video processing

Tracking objects, video analysis, low level motion, optical flow

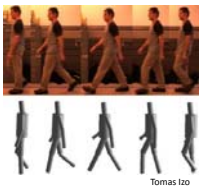

Tomas Izo

Slide credit: Kristen Grauman

63

Textbooks

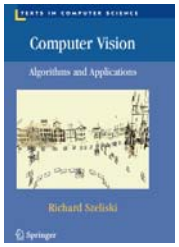

By Rick Szeliski
http://szeliski.org/Book/

By Kristen Grauman, Bastian Leibe
Visual Object Recognition

64

Requirements / Grading

- Problem sets (60%)
- Final exam (35%)
  - comprehensive (cover all topics learned in class)
- Class and Piazza participation, including attendance (5%)
  - Piazza: participation points for posting (sensible) questions and answers

65

Problem sets

- Some short answer concept questions
- Matlab programming problems
  - Implementation
  - Explanation, results
- Follow instructions; points will be deducted if we can't run your code out of the box
- Ask questions on Piazza first
- Submit to Canvas
- The assignments will take significant time to do
- Start early
- TA will go over problem set during discussion sections after release (others will be used as extra office hours)

66

Matlab

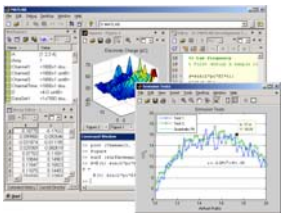

- Built-in toolboxes for low-level image processing, visualization
- Compact programs
- Intuitive interactive debugging
- Widely used in engineering

Slide credit: Kristen Grauman

67

Matlab

- CSIF labs 67, 71, 75 (pc1-pc60)
- Academic Surge 1044 and 1116
- Lab schedule (reservations) and remote access info found on class website
- Matlab available for free from campus software site

68

Problem Set 0

- Matlab warmup
- Basic image manipulation
- Out Thursday, due 4/13

69

Preview of some problem sets

resize: castle squished

crop: castle cropped

content aware resizing: seam carving

70

Preview of some problem sets

Grouping

71

Preview of some problem sets

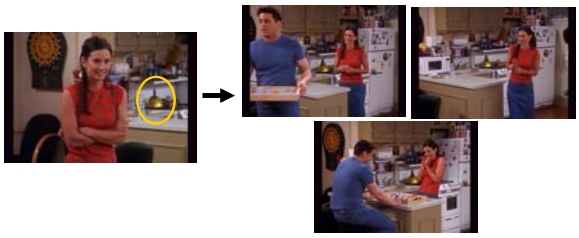

Object search and recognition

72

Problem set deadlines

- Problem sets due 11:59 PM
  - Follow submission instructions given in assignment
  - Submit to Canvas; no hard copy submissions
  - Deadlines are firm. We'll use Canvas timestamp. Even 1 minute late is late.
- Late submissions: 1 point deduction for every hour after the deadline up to 72 hours; after 72 hours, you will receive a 0
- If your program doesn't work, clean up the code, comment it well, explain what you have, and still submit. Draw our attention to this in your answer sheet.

73

Collaboration policy

- **All responses and code must be written individually or in pairs (a group of 2)**
- Students submitting answers or code found to be identical or substantially similar (due to inappropriate collaboration) risk failing the course
  - We will be using MOSS to check for cheating!
  - Copying online solutions also counts as cheating!
  - Please don't cheat… you are going to get caught!
- Read and follow UC Davis code of conduct

74

MOSS

75

Miscellaneous

- Check class website regularly for assignment files, notes, announcements, etc.
- Come to lecture on time
- No laptops, phones, tablets, etc. in class please
- Please interrupt with questions at any time

76

Coming up

- Read the class webpage carefully
- Next class (Thurs): lecture on linear filters
- PS0 out Thursday, due 4/13

77

Questions?

See you Thursday!

78