Computer Vision

Thursday, Aug 28

Introductions

- **Instructor**: Prof. Kristen Grauman
  grauman @ cs
  CSA 114, Tues/Thurs 2-3 pm
- **TA**: Harshdeep Singh
  harshd @ cs
  TAY basement CS Lab, Time TBD

Today

- Course overview
- Requirements, logistics

What is computer vision?

Done?

What is computer vision?

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

Vision for measurement

Real-time stereo

Pollefeys et al.

Structure from motion

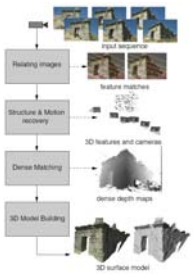

Multi-view stereo for community photo collections

Goesele et al.

Slide credit: L. Lazebnik

Vision for perception, interpretation

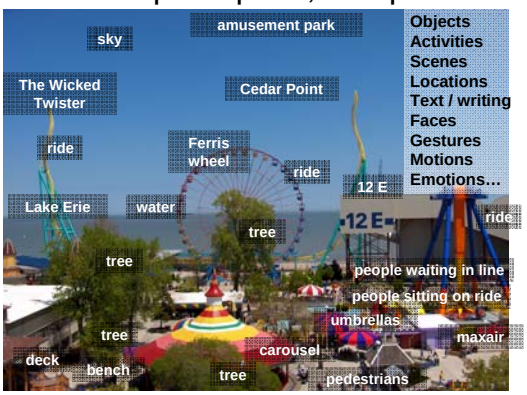

Related disciplines

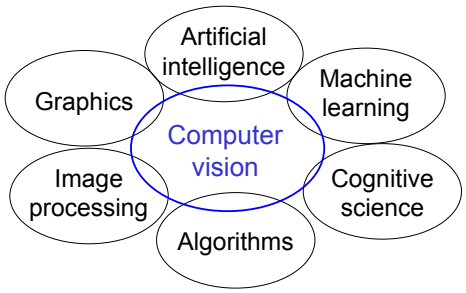

Vision and graphics

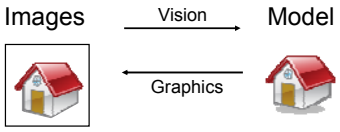

Inverse problems: analysis and synthesis.

Why vision?

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

Why vision?

- Images and video are everywhere!

Slide credit: L. Lazebnik

Faces and digital cameras

Camera waits for everyone to smile to take a photo [Canon]

Setting camera focus via face detection

Linking to info with a mobile device

Situated search
Yeh et al., MIT

kooaba

MSR Lincoln

Video-based interfaces

Human joystick
NewsBreaker Live

Assistive technology systems
Camera Mouse
Boston College

Vision for medical & neuroimages

Image guided surgery
MIT AI Vision Group

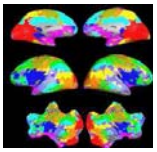

fMRI data
Golland et al.

Special visual effects

The Matrix

What Dreams May Come

Mocap for *Pirates of the Carribean*, Industrial Light and Magic
Source: S. Seitz

Safety & security

Navigation,
driver safety

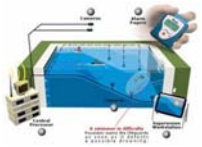

Monitoring pool
(Poseidon)

Pedestrian detection
MERL, Viola et al.

Surveillance

Obstacles?

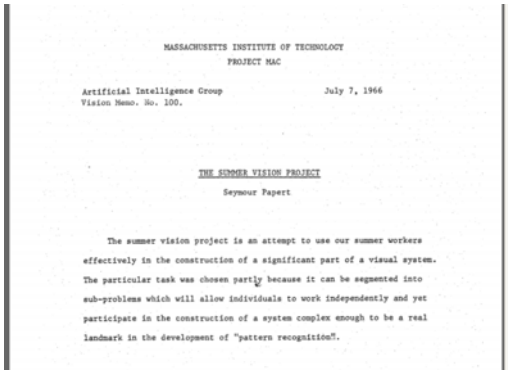

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
PROJECT MAC

Artificial Intelligence Group
Vision Memo. No. 100.

July 7, 1966

THE SUMMER VISION PROJECT

Seymour Papert

The summer vision project is an attempt to use our summer workers effectively in the construction of a significant part of a visual system. The particular task was chosen partly because it can be segmented into sub-problems which will allow individuals to work independently and yet participate in the construction of a system complex enough to be a real landmark in the development of "pattern recognition".

What the computer gets

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

Challenges: many nuisance parameters

Illumination

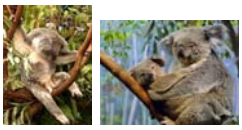

Object pose

Clutter

Occlusions

Intra-class appearance

Viewpoint

Challenges: intra-class variation

slide credit: Fei-Fei, Fergus & Torralba

Challenges: importance of context

slide credit: Fei-Fei, Fergus & Torralba

Challenges: complexity

- Thousands to millions of pixels in an image
- 3,000-30,000 human recognizable object categories
- 30+ degrees of freedom in the pose of articulated objects (humans)
- Billions of images indexed by Google Image Search
- 18 billion+ prints produced from digital camera images in 2004
- 295.5 million camera phones sold in 2005
- About half of the cerebral cortex in primates is devoted to processing visual information [Felleman and van Essen 1991]

- Ok, clearly the vision problem is deep and challenging…time to give up?
- Active research area with exciting progress!

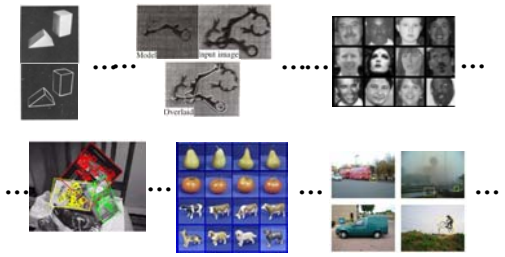

Brainstorm

1. What functionality should the system have?
2. Intuitively, what are the technical components or sub-problems that must be solved?

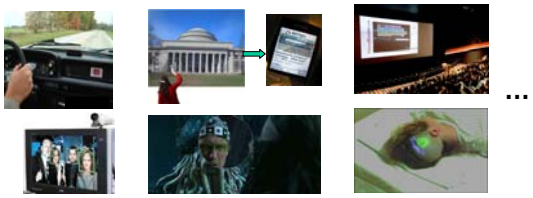

Goals of this course

- Upper division undergrad course

- Introduction to primary topics
- Hands-on experience with algorithms
- Views of vision as a research area

Topics overview

- Image formation
- Features
- Grouping & fitting
- Multi-view geometry
- Recognition & learning
- Motion & tracking

Focus is on algorithms, rather than specific systems.

Image formation

- How does light in 3d world project to form 2d images?

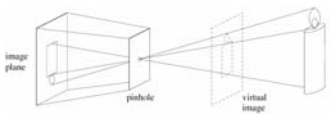

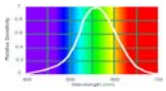

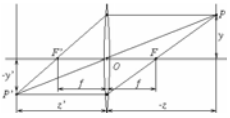

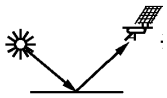

Features and filters

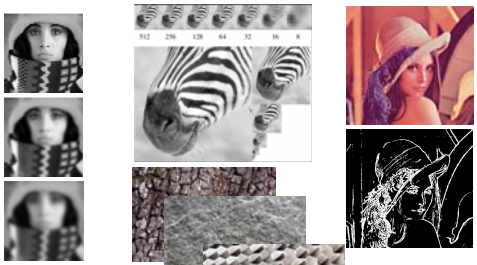

Transforming and describing images; textures, colors, edges

Grouping & fitting

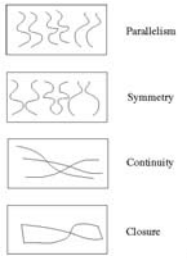

[fig from Shi et al]

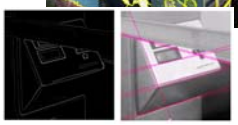

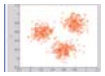

Clustering, segmentation, fitting; what parts belong together?

Multiple views

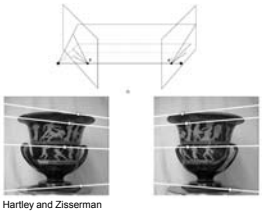

Hartley and Zisserman

Multi-view geometry, matching, invariant features, stereo vision

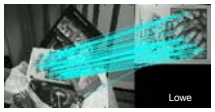

Lowe

Fei-Fei Li

Recognition and learning

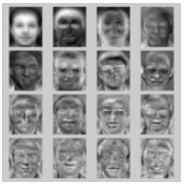

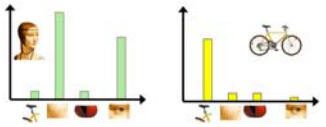

Recognizing objects and categories, learning techniques

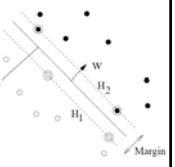

Motion and tracking

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

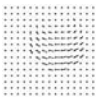

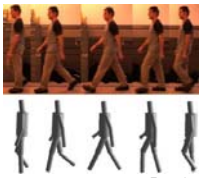

Tomas Izo

Course web page

http://www.cs.utexas.edu/~grauman/courses/fall2008/main.htm

Fall 2008: CS 378 Computer Vision

Computer Vision
Fall 2008

Tues/Thurs 12:30 – 2:00 pm
Parlin Hall 1 (PAR 1)

CS 378, Unique # 55770

Instructor: Prof. Kristen Grauman
Email: last name @ cs.utexas.edu
Office hours: Tues/Thurs 2:00-3:00 pm in CSA 114

TA : Harshdeep Singh
Email: first name @ gmail.com
Office hours : TBD

The TA station is in the basement of ENS inside room 31NR.
Directions to the TA stations are posted right outside the basement elevator, and also outside room 31NR.

Schedule | Blackboard | Overview | Requirements | Links | Syllabus outline

Announcements

Pset 0 is due Thurs Sept 4: getting to know Matlab.

Schedule

CS 378 / 395T Computer Vision : Schedule

Computer Vision
Fall 2008

Please note - specifics of this schedule are subject to change.

F&P = Forsyth & Ponce
S&S = Shapiro & Stockman
T&V = Trucco & Verri

Dates	Topic	Reading and references	Of related interest	Lectures	Assignments
8/28 9/2	Image formation Matlab tutorial	F&P Chapter 1 Matlab intro	Who Invented Ray Tracing? By G. Hofmann Building a camera with a Pringles can		Pset 0 Pset 0 images
9/4	Color	F&P Chapter 6 The foundations of color measurement and color perception by Brian A. Wandell (optional)			**Pset 0 due 9/4**
9/9 9/11 9/16	Features and texture	F&P Chapters 7, 9 [T&V Chapter 4] [S&S Chapter 3] [S&S Chapter 5]			
9/18 9/23 9/25	Grouping and fitting				**Pset 1 due 9/18**

Textbooks

- Forsyth & Ponce is recommended book
- This and others are on reserve at PCL
- For some topics, we'll post small sections from external sources on Blackboard.

Requirements / Grading

- Problem sets (55%)
- Midterm exam (15%)
- Final exam (20%)
- Class participation, including attendance (10%)

Problem sets

- Some short answer concept questions
- Programming problem
  - Implementation
  - Explanation, results
- Code in Matlab – available on CS Unix machines

Preview of problem sets

- Pset 0: Matlab warmup

Preview of problem sets

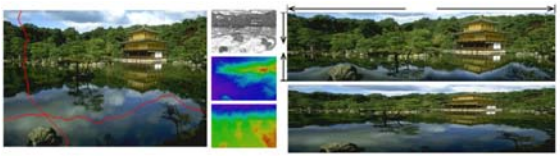

Pset 1: Content-aware image resizing

Shai Avidan, Ariel Shamir
Seam Carving for Content-Aware Image Resizing
ACM Transactions on Graphics, Volume 26, Number 3, SIGGRAPH 2007

Preview of problem sets

Pset 2: Image mosaics

Image from Fei-Fei Li

Preview of problem sets

Pset 3: Video search

Preview of problem sets

Pset 4: Face recognition

Preview of problem sets

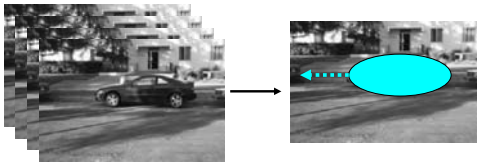

Pset 5: Tracking

Collaboration policy

All responses and code must be written individually.

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

CSA

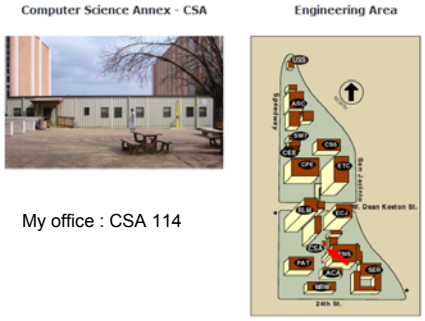

My office : CSA 114

TAY

Harshdeep's office hours in TAY CS Lab, basement

Due dates

- Assignments in by11:59 PM on due date
  - Follow submission instructions given in assignment regarding hardcopy/electronic
- Lose half of possible remaining credit each day late
- Three free late days, total
  - Notify TA **by due date** if using

Sharing results (optional)

- We'll review results for some problems afterwards
- Share your results in class for extra credit
  - Email beforehand
  - Brief description (≈2 minutes), visual
  - Up to 6 points extra credit on midterm

Current events (optional)

- Any vision-related piece of news; may revolve around policy, editorial, technology, new product, …
- Brief overview to the class (≈5 minutes)
- Must be current

- Email the relevant links or information beforehand
- Extra credit: up to 6 points on midterm score

Miscellaneous

- Check class website regularly
- Make sure you are on class mailing list
- No laptops in class please

- Feedback welcome and useful

Next

Problem set 0 due Sept 4 (next Thursday)
- Matlab warmup: download from class page

Tuesday: in-class tutorial on Matlab

Thursday: Image formation
- Read F&P Chapter 1,6