

CMPT 882 – Recognition Problems in Computer Vision

Greg Mori

Outline

- Intro to class
- Administrative details

Overview

- This class is about visual “recognition”
 - **Objects:** cups, cars, horses, ... accordions to zebras
 - **Textures:** grass, leaves, dirt, water, ...
 - **Human figures:** faces; whole body; elbows, wrists, knees,...
 - **Human actions:** running, jumping, waving, ...
 - **Places:** office, city street, beach, jungle, ...
- Goal is to provide view of state-of-art for these problems

Objects

- What is “Object recognition?”
 - overloaded term



- Is there a car in this image?
 - Object/image categorization
 - Object category recognition
- Where is the car?
 - Object localization
 - Object detection
- Which car is it?
 - Object recognition
 - Object identification


Challenges in Recognition

- Intra-class variation
- Object pose variation
- Background clutter
- Occlusion
- Lighting



Object Recognition - Shape


- Template matching using shape



Berg et al. CVPR 05

Object Recognition – Appearance

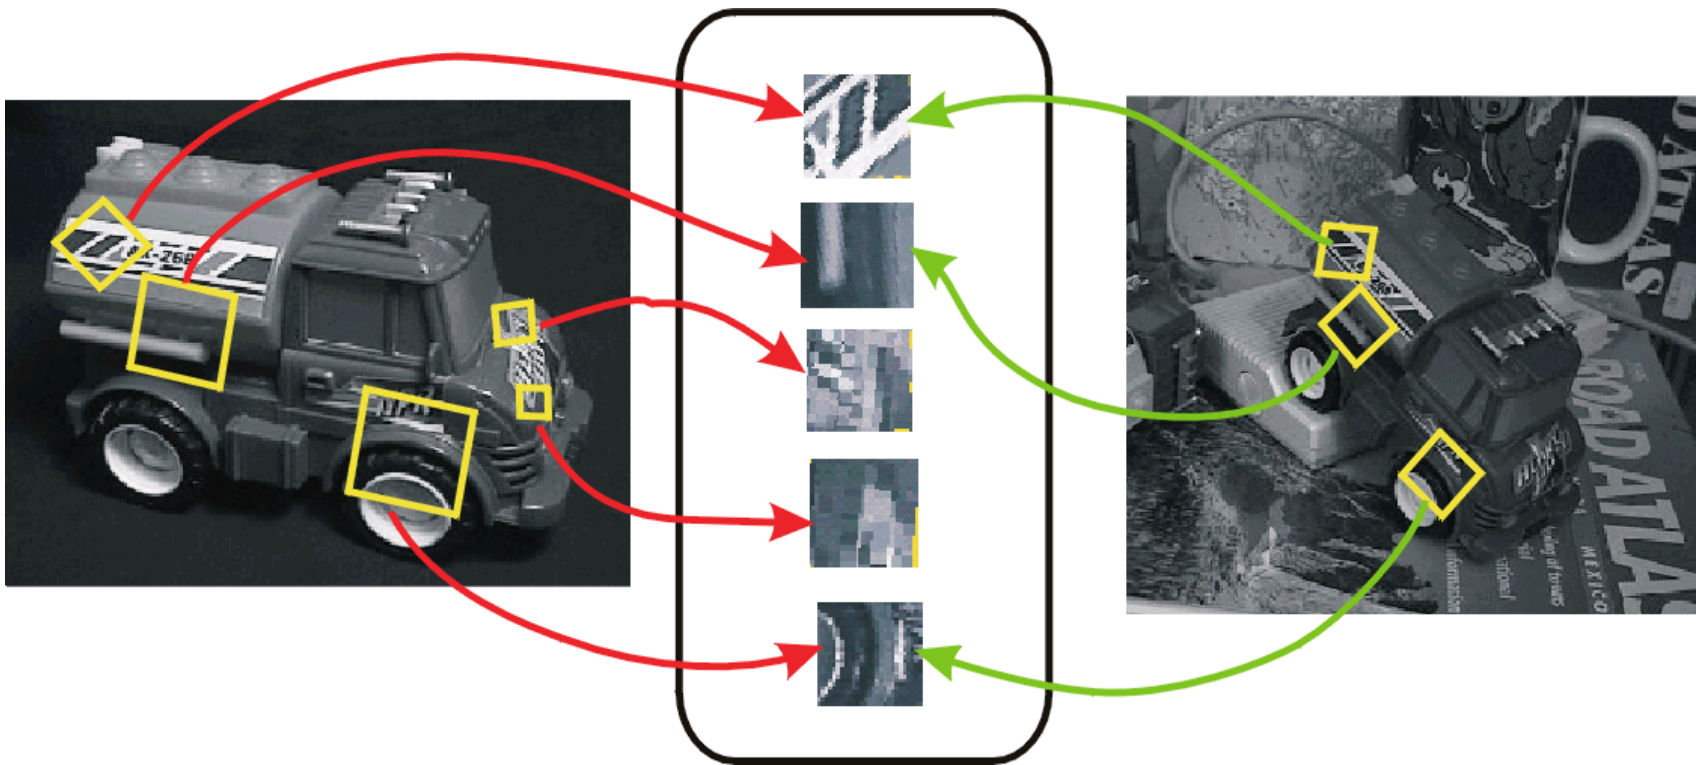
- Histograms of gradients



Dalal and Triggs CVPR 05


Object Recognition – Local Features

- D. Lowe SIFT (ICCV 99, IJCV 04)



Fast Object Retrieval


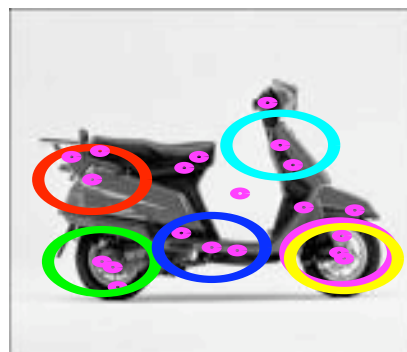
- Stewenius + Nister, CVPR 06
 - 50,000 images at 8Hz (laptop)



cf. SnapTell



Object Recognition – Part-based Models

- Constellation models



Fergus et al. CVPR 03

- Latent SVM




Felzenszwalb et al. CVPR 08

Photosynth

- Noah Snavely, Steven M. Seitz, Richard Szeliski, "Photo tourism: Exploring photo collections in 3D," SIGGRAPH06

[Photo tourism video](#)

Textures



Clothing Textures




Human Figures

- Faces (Viola + Jones CVPR 01)




Human Figures

- Implicit shape model



Leibe et al. CVPR 05




Leibe et al. CVPR 07

Human Figures – Pose Estimation



Human Actions





Shechtman and Irani CVPR 05

Real-time Gesture Recognition



Bayazit et al. MVA 09

Places

livingroom bedroom



highway



kitchen



office



ins. city



Using Context

We know there is a keyboard present in this scene even if we cannot see it clearly.



We know there is no keyboard present in this scene



... even if there is one indeed.

Slide: Torralba

Course Plan

- Read research papers
 - For each topic I present important papers
 - Students each present a recent paper
 - We discuss
- Do a project
 - Gain in-depth experience on a problem and algorithm

Introductions

Prerequisite

- No formal prerequisites
- You will need to do the usual things
 - Math (continuous), programming, reading, writing, presenting
- Ask me if you are concerned

Grading Scheme

- 10% Class participation
 - Participate in discussions about papers, ask/answer questions
- 10% Reading assignments
 - 1 or 2 papers each week; the ones I present
- 10% Paper presentation
 - List of recommended papers online
- 10% Assignment
 - Small programming assignment on edges and texture
- 60% Project
 - Individual or in small groups
 - Presentation, written report

Reading Assignments

- Similar to mini paper review
 - One paragraph summarizing paper
 - Critical discussion (what you like / don't like)
 - Questions you have (for me to explain)
- Due before start of lecture via email
- These details and list of papers are online

Paper Presentations

- Choose one recent paper from area that interests you
 - Recommended list online
- 20 minute presentation
 - 10+ minutes questions/discussion
 - Feel free to use slides provided by authors


Assignment

- Short programming assignment
 - Canny edge detection
 - Texture recognition
- Out next week, due 2 weeks later
- Choice of language yours
 - MATLAB recommended

Project

- Major component of course
- Recommended projects:
 - Object category recognition (Caltech 101)
 - Human action recognition (Weizmann)
- Implement existing technique
 - Or variant thereof
- Proposal, presentation, report

Caltech 101



- Object category recognition
 - 101 classes, ~50-100 examples of each

Weizmann Human Action Dataset



- 9 subjects, each performs 9* actions

- Wednesday
 - Edge detection basics
- Next week
 - Edge detection, texture