CS4670/5670: Intro to Computer Vision

Noah Snavely

Lecture 25: Introduction to Recognition

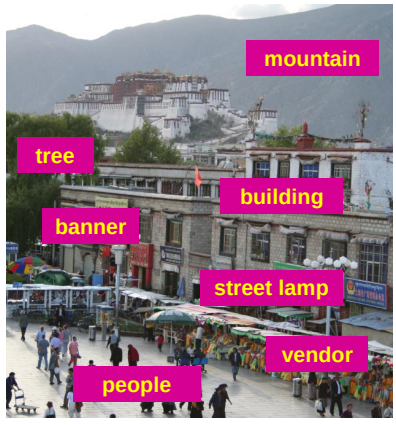

What do we mean by “object recognition”?

Next 15 slides adapted from Li, Fergus, & Torralba’s excellent short course on category and object recognition

Verification: is that a lamp?

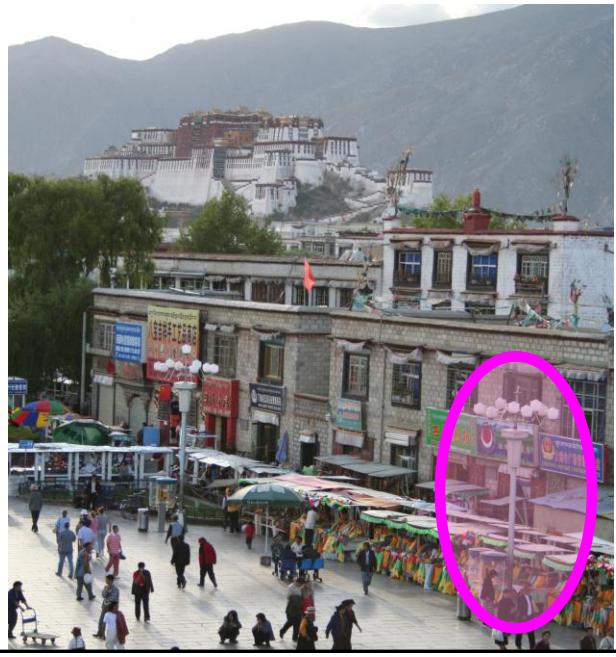

Detection: are there people?

Identification: is that Potala Palace?

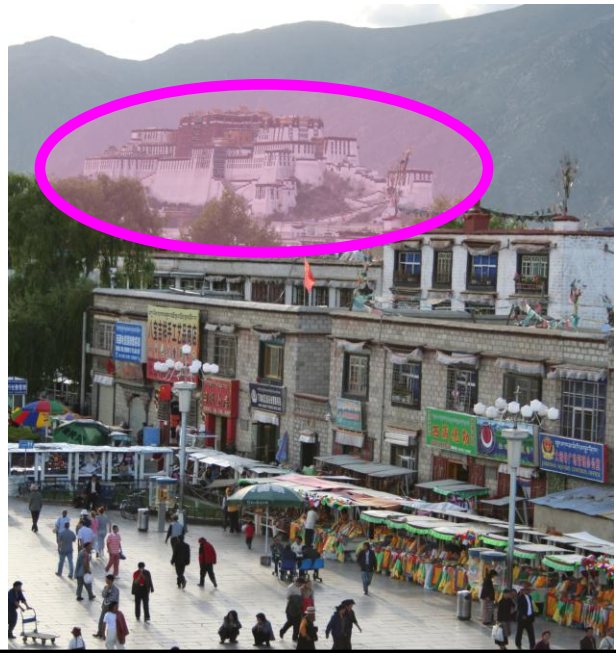

Object categorization

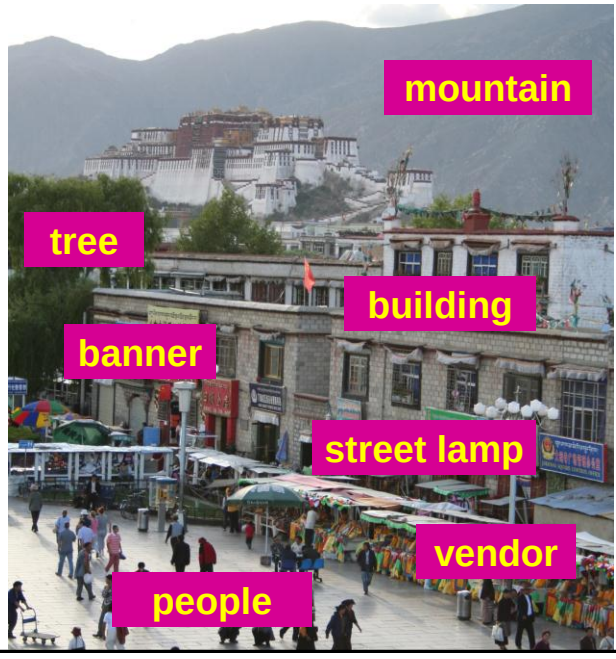

Scene and context categorization

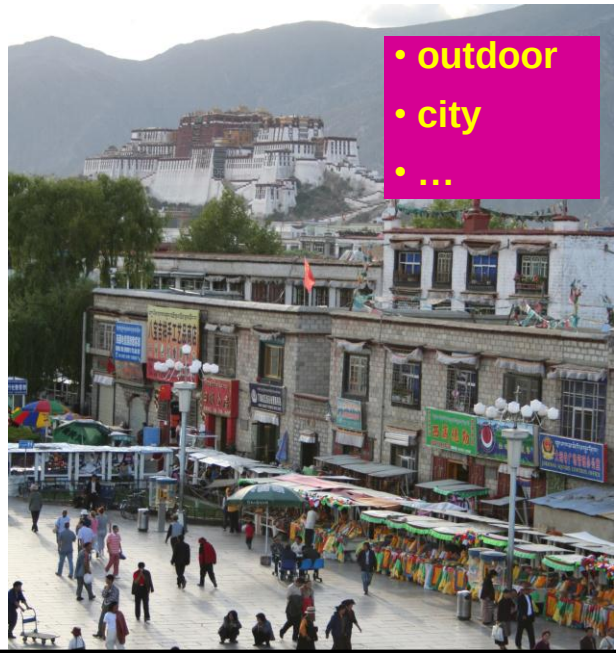

Object recognition
Is it really so hard?

This is a chair

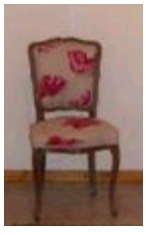

Find the chair in this image

Output of normalized correlation

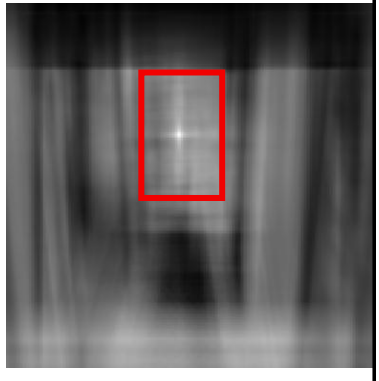

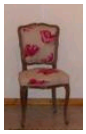

Object recognition
Is it really so hard?

Find the chair in this image

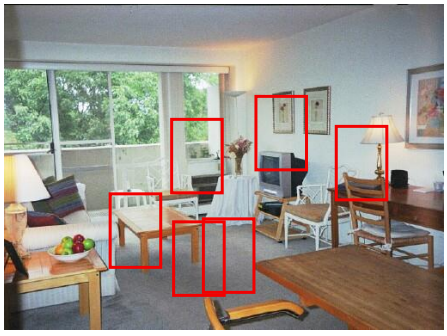

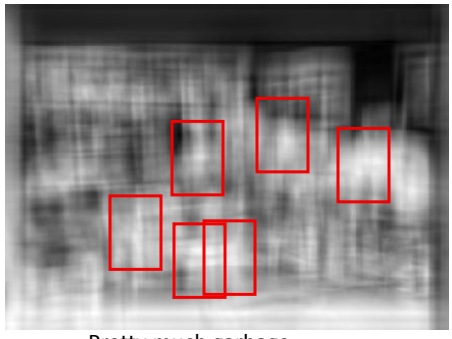

Pretty much garbage
Simple template matching is not going to make it

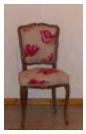

Object recognition
Is it really so hard?

Find the chair in this image

A “popular method is that of template matching, by point to point correlation of a model pattern with the image pattern. These techniques are inadequate for three-dimensional scene analysis for many reasons, such as occlusion, changes in viewing angle, and articulation of parts.” Nivatia & Binford, 1977.

Why not use SIFT matching for everything?

- Works well for object *instances*

- Not great for generic object *categories*

Applications: Computational photography

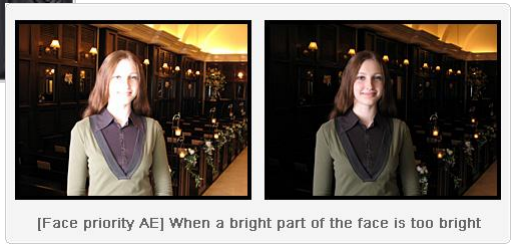

[Face priority AE] When a bright part of the face is too bright

Applications: Assisted driving

Pedestrian and car detection

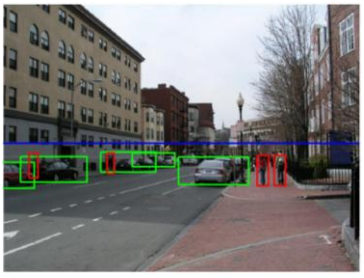

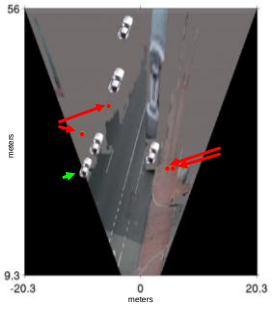

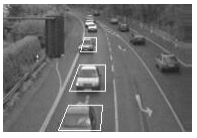

Lane detection

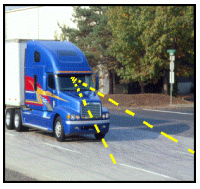

- Collision warning systems with adaptive cruise control,
- Lane departure warning systems,
- Rear object detection systems,

Applications: image search

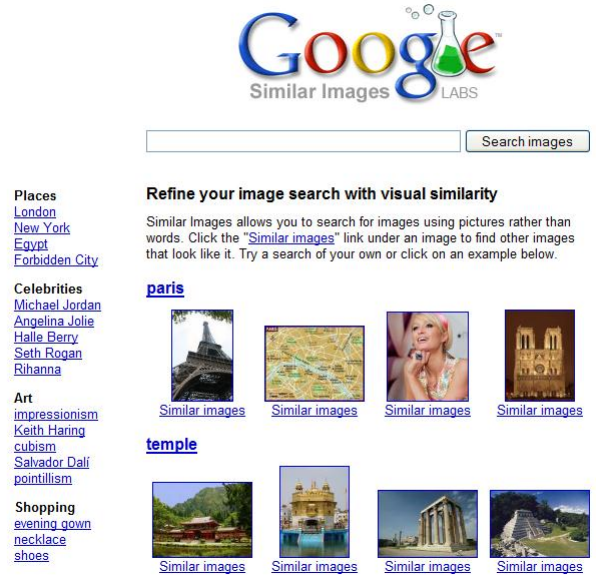

How do human do recognition?

- We don't completely know yet
- But we have some experimental observations.

Observation 1

- We can recognize familiar faces even in low-resolution images

Observation 2:

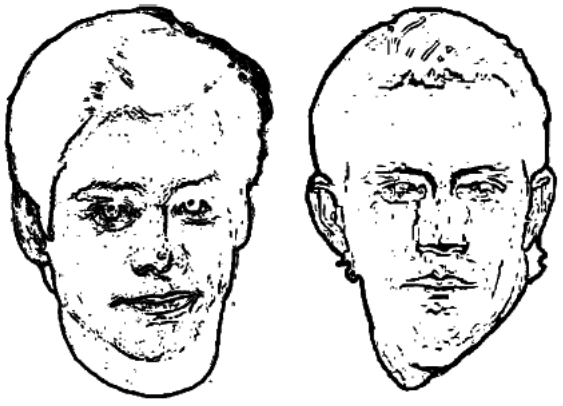

Jim Carrey Kevin Costner

- High frequency information is not enough

What is the single most important facial features for recognition?

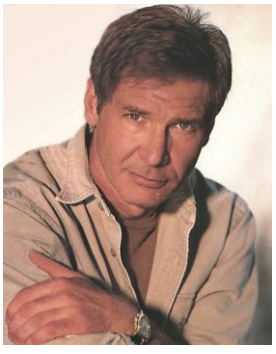

Observation 4:

- Image Warping is OK

The list goes on

Face Recognition by Humans: Nineteen Results All Computer Vision Researchers Should Know About

- http://web.mit.edu/bcs/sinha/papers/19results_sinha_etal.pdf

Let's start simple

- Today
  - skin detection
  - eigenfaces

Face detection

- Do these images contain faces? Where?

One simple method: skin detection

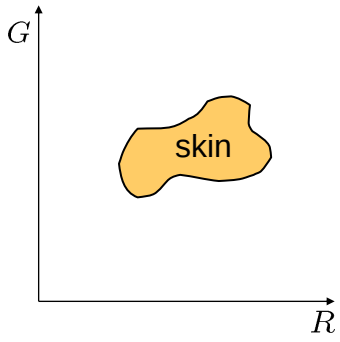

Skin pixels have a distinctive range of colors

- Corresponds to region(s) in RGB color space
  - for visualization, only R and G components are shown above

Skin classifier

- A pixel X = (R,G,B) is skin if it is in the skin region
- But how to find this region?