CS4670: Intro to Computer Vision

Noah Snavely

Lecture 25: Introduction to Recognition

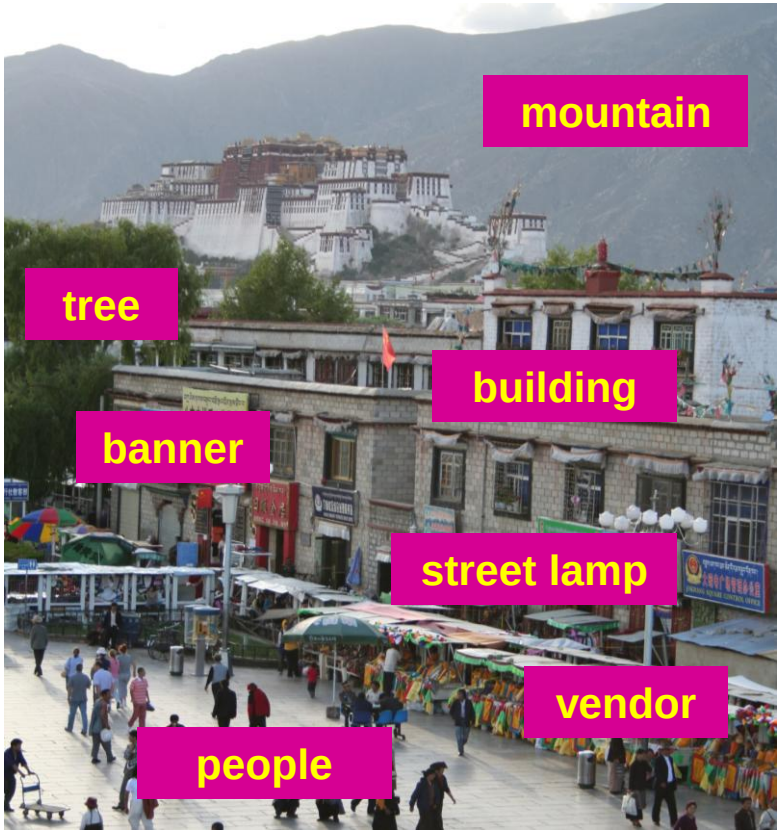

What do we mean by “object recognition”?

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?

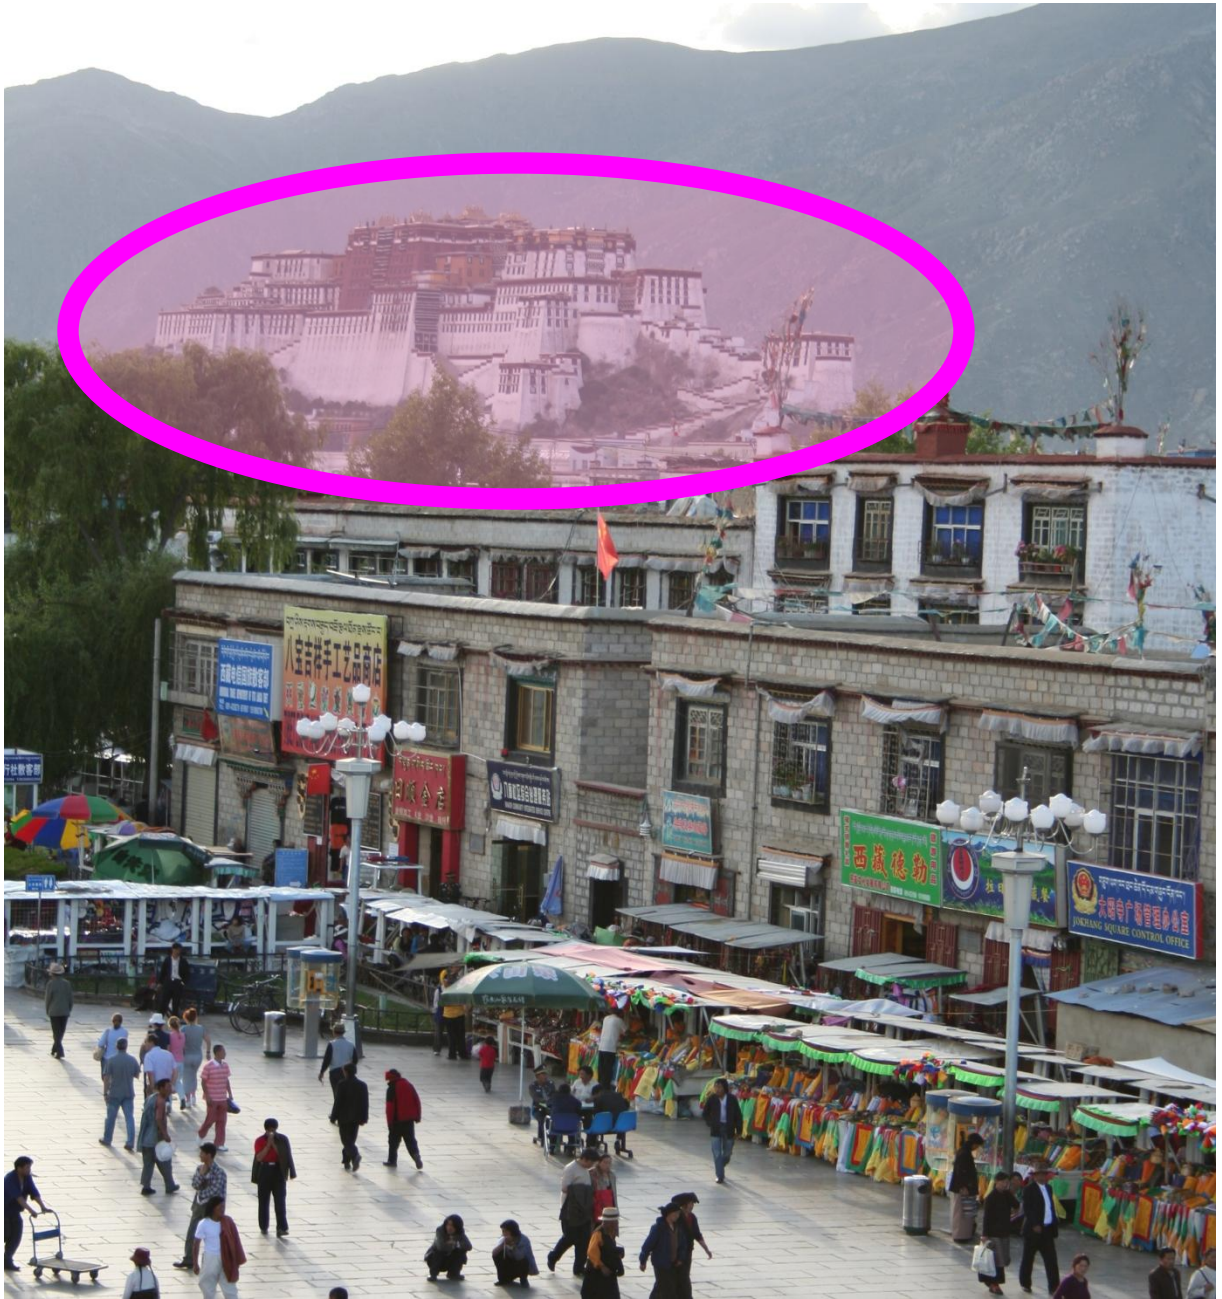

Object categorization

Scene and context categorization

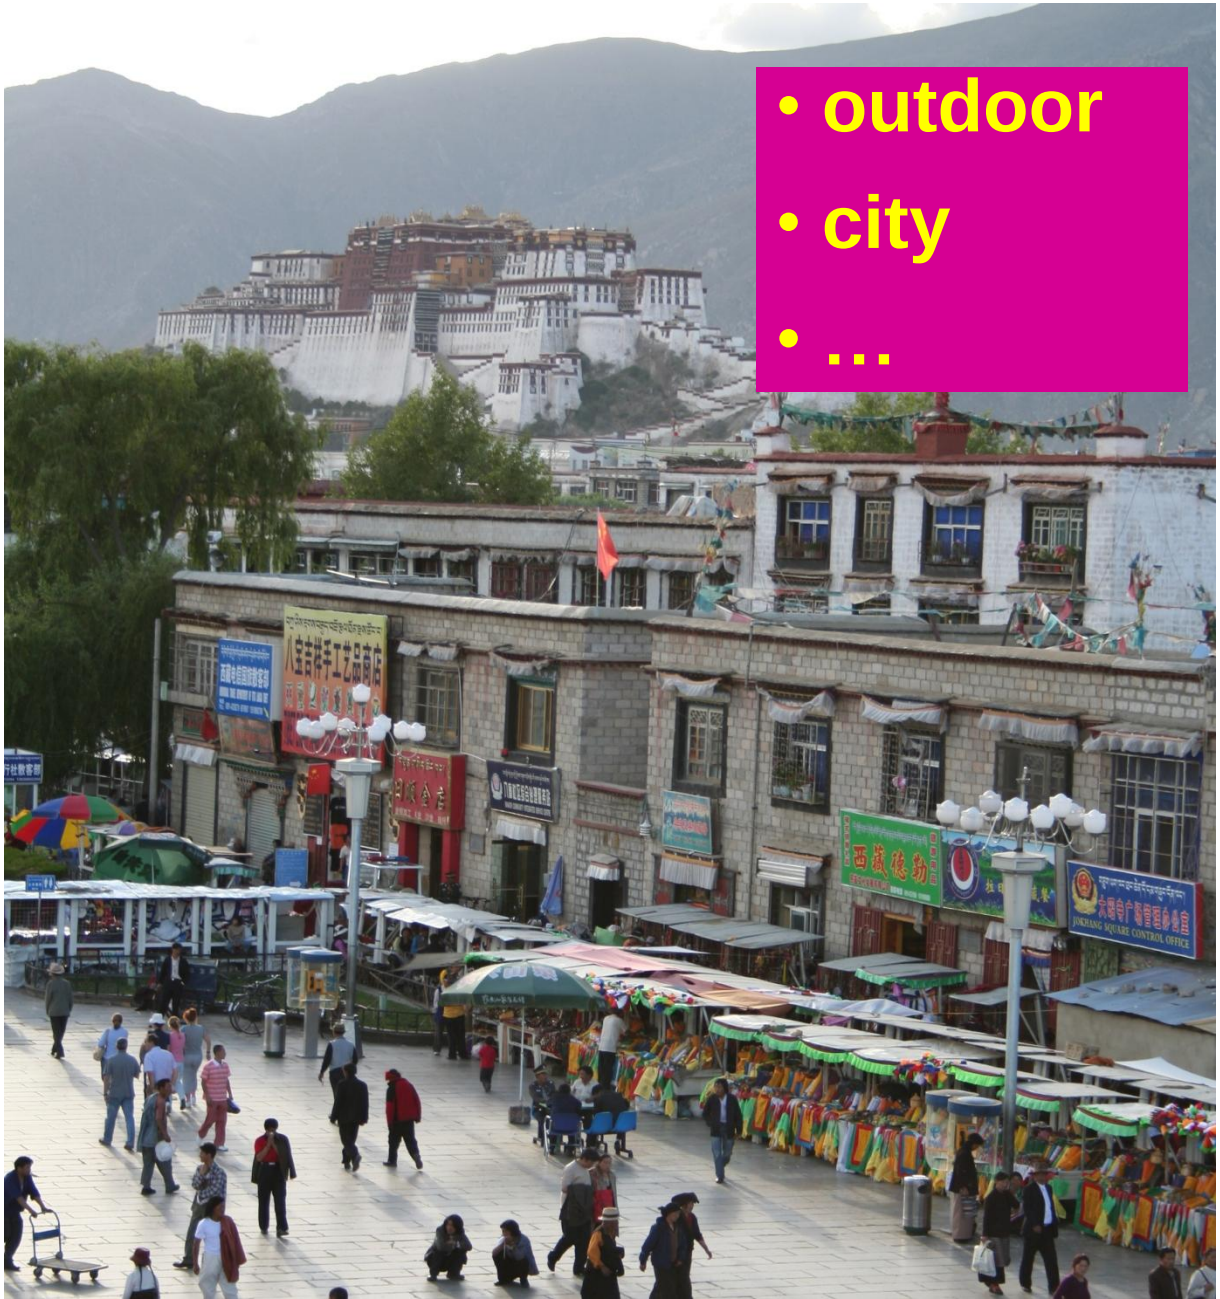

Object recognition
Is it really so hard?

This is a chair

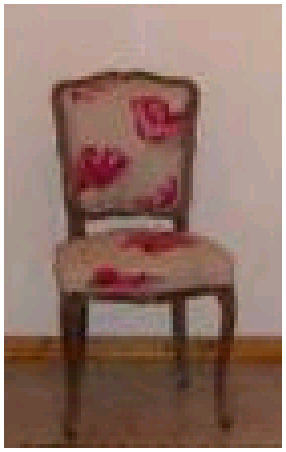

Find the chair in this image

Output of normalized correlation

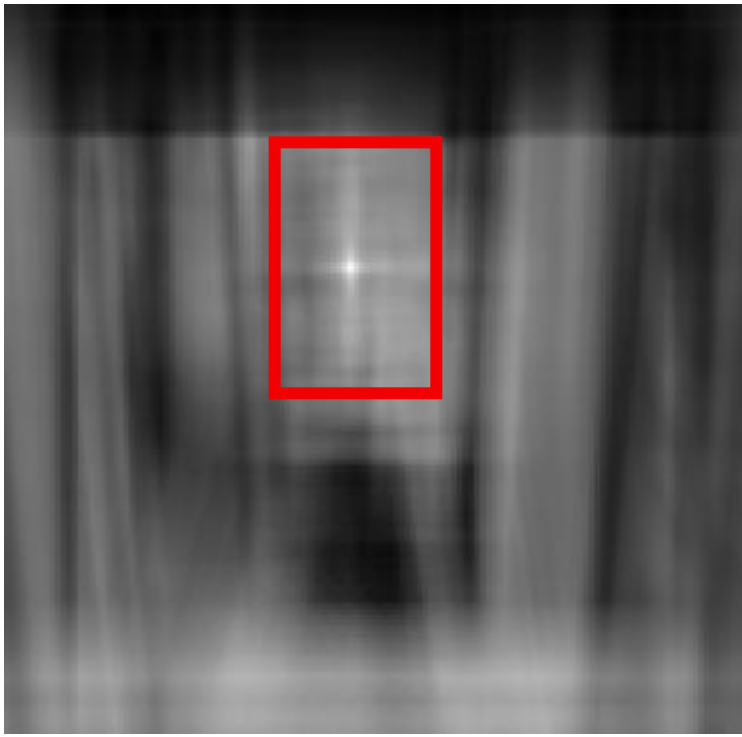

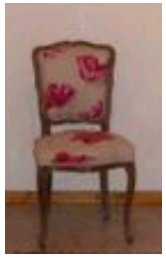

Object recognition
Is it really so hard?

Find the chair in this image

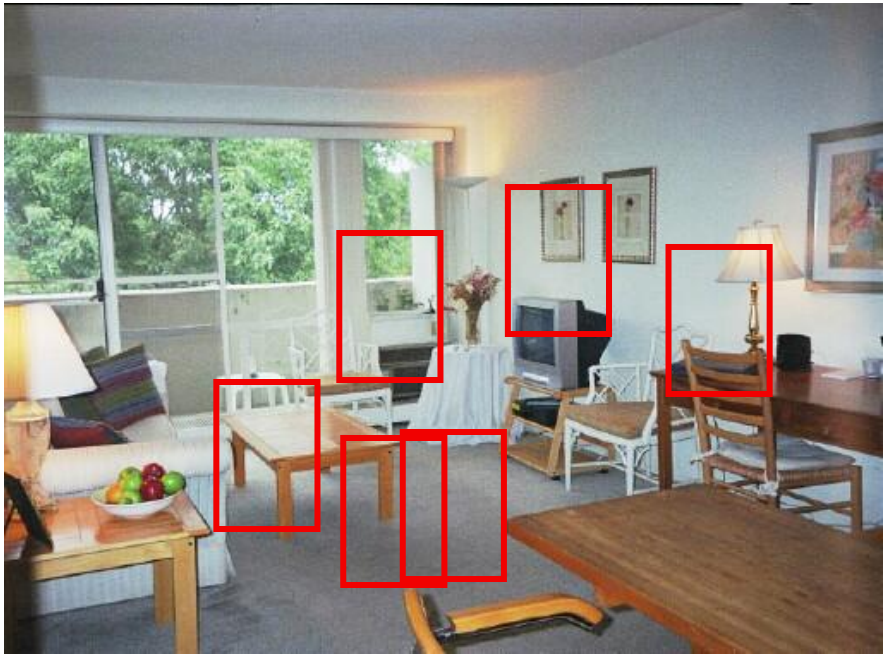

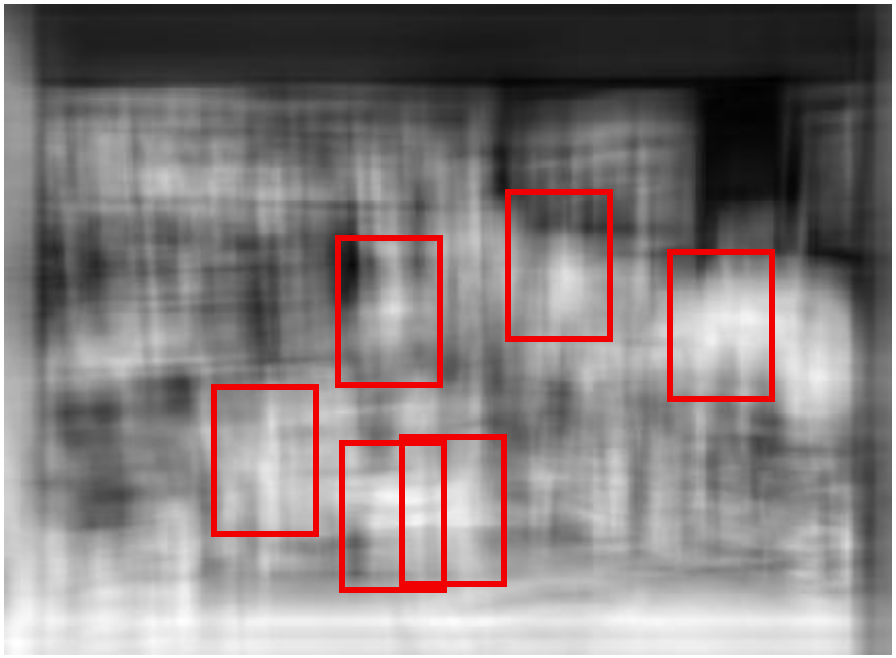

Pretty much garbage

Simple template matching is not going to make it

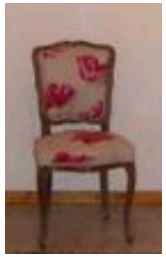

Object recognition
Is it really so hard?

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

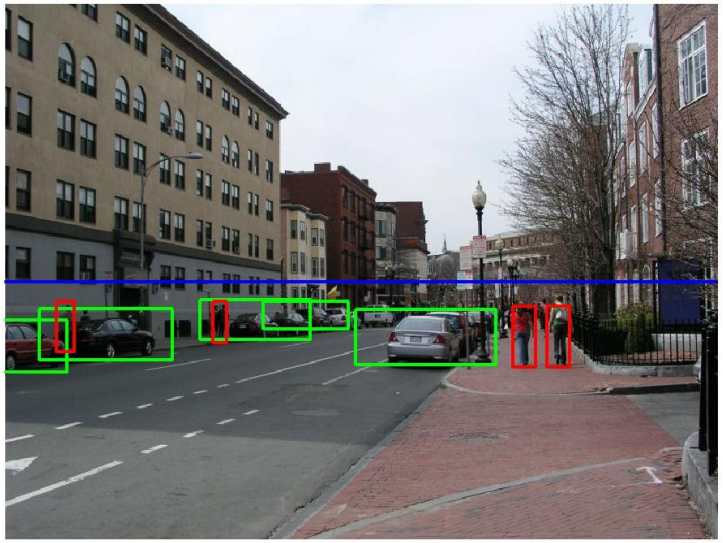

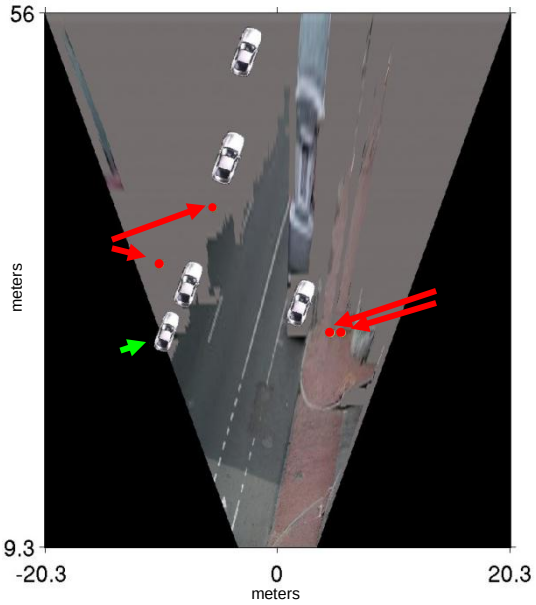

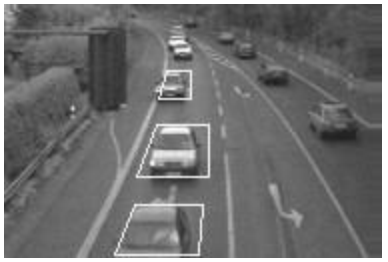

Lane detection

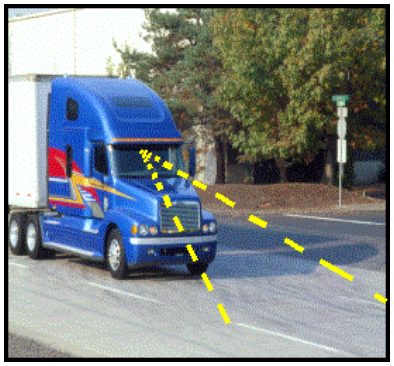

- Collision warning systems with adaptive cruise control,
- Lane departure warning systems,
- Rear object detection systems,

Applications: image search

Google Similar Images LABS

Search images

Places
London
New York
Egypt
Forbidden City

Celebrities
Michael Jordan
Angelina Jolie
Halle Berry
Seth Rogan
Rihanna

Art
impressionism
Keith Haring
cubism
Salvador Dalí
pointillism

Shopping
evening gown
necklace
shoes

Refine your image search with visual similarity

Similar Images allows you to search for images using pictures rather than words. Click the "Similar images" link under an image to find other images that look like it. Try a search of your own or click on an example below.

paris


Similar images


Similar images


Similar images


Similar images

temple


Similar images


Similar images


Similar images


Similar images

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?

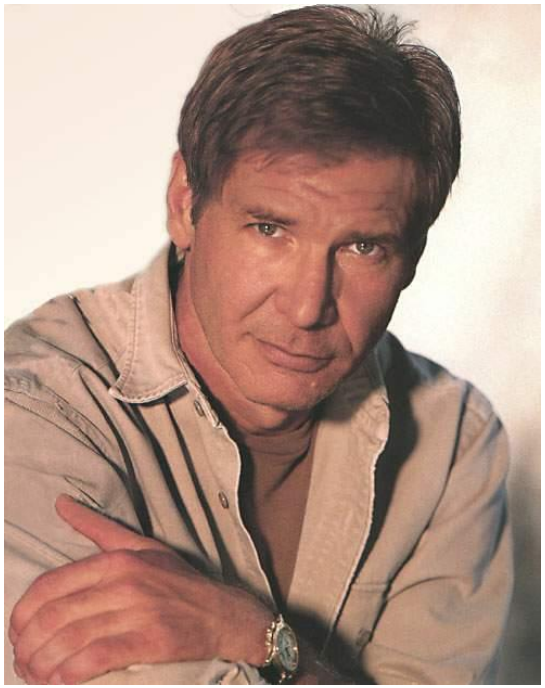

Observation 4:

- 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

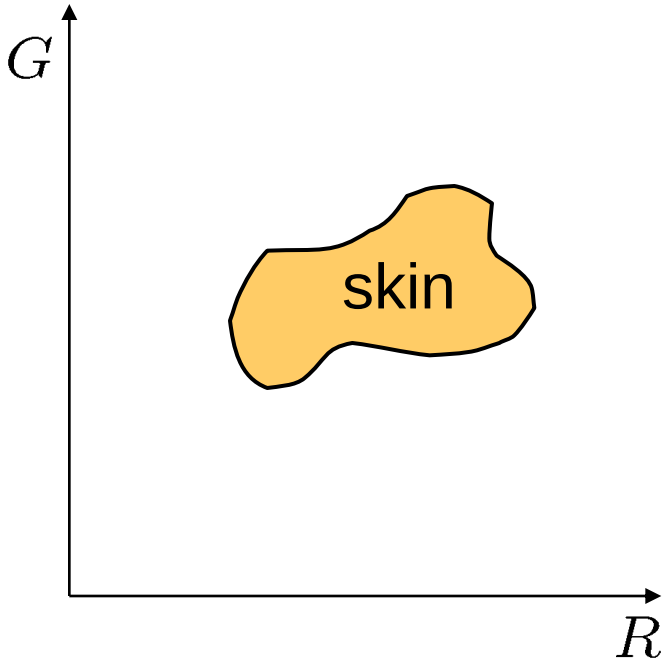

Skin pixels have a distinctive range of colors

- 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?

Skin detection

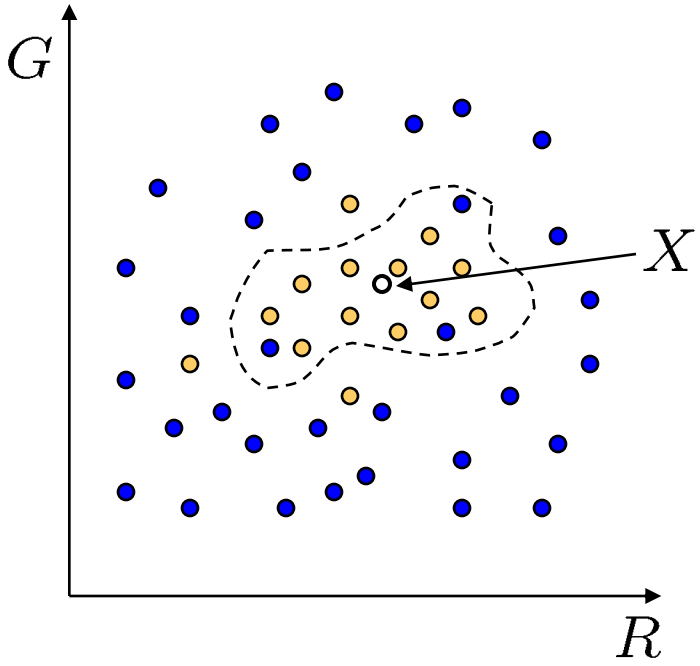

Learn the skin region from examples

- Manually label pixels in one or more “training images” as skin or not skin
- Plot the training data in RGB space
  - skin pixels shown in orange, non-skin pixels shown in blue
  - some skin pixels may be outside the region, non-skin pixels inside. Why?

Skin classifier

- Given X = (R,G,B): how to determine if it is skin or not?

Skin classification techniques

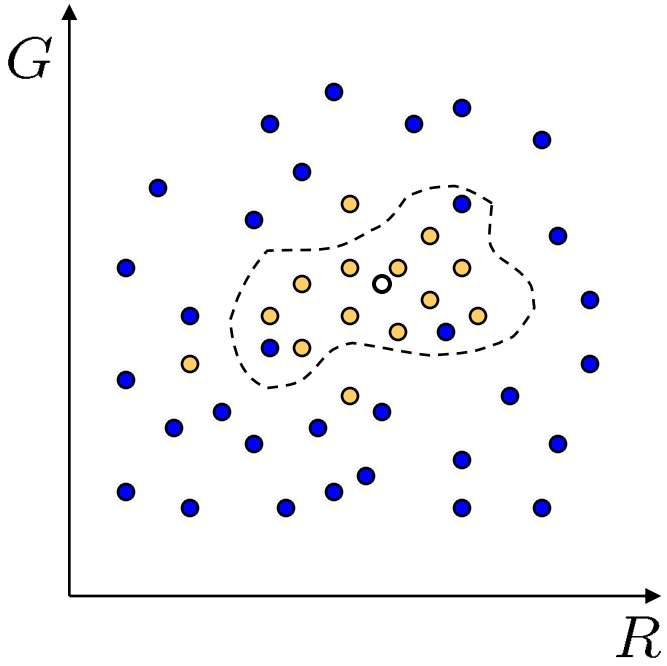

Skin classifier

- Given X = (R,G,B): how to determine if it is skin or not?
- Nearest neighbor
  - find labeled pixel closest to X
  - choose the label for that pixel
- Data modeling
  - fit a model (curve, surface, or volume) to each class
- Probabilistic data modeling
  - fit a probability model to each class

Probability

Basic probability

- **X** is a random variable
- **P(X)** is the probability that **X** achieves a certain value

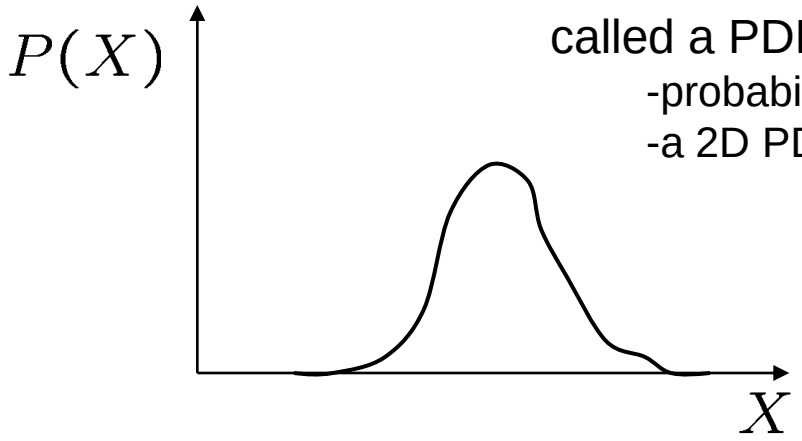

called a PDF

-probability distribution/density function

-a 2D PDF is a surface, 3D PDF is a volume

- $0 \leq P(X) \leq 1$
- $\int_{-\infty}^{\infty} P(X)dX = 1$ or $\sum P(X) = 1$

  continuous **X** discrete **X**

- Conditional probability: **P(X | Y)**
  - probability of **X** given that we already know **Y**

Probabilistic skin classification

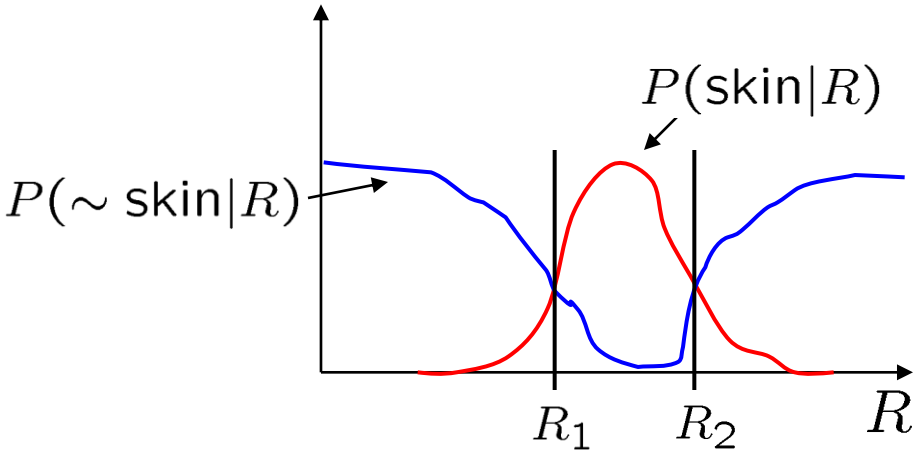

Now we can model uncertainty

- Each pixel has a probability of being skin or not skin
  - $P(\sim \text{skin}|R) = 1 - P(\text{skin}|R)$

Skin classifier

- Given X = (R,G,B): how to determine if it is skin or not?
- Choose interpretation of highest probability
  - set X to be a skin pixel if and only if $R_1 < X \leq R_2$

Where do we get $P(\text{skin}|R)$ and $P(\sim \text{skin}|R)$?