Modern Boundary Detection II

Szeliski 4.2

Computer Vision

CS 143, Brown

James Hays

Many slides Michael Maire, Jitendra Malek

Today’s lecture

- pB boundary detector, from “local” to “global”
- Finding Junctions
- Sketch Tokens

How good are humans locally?

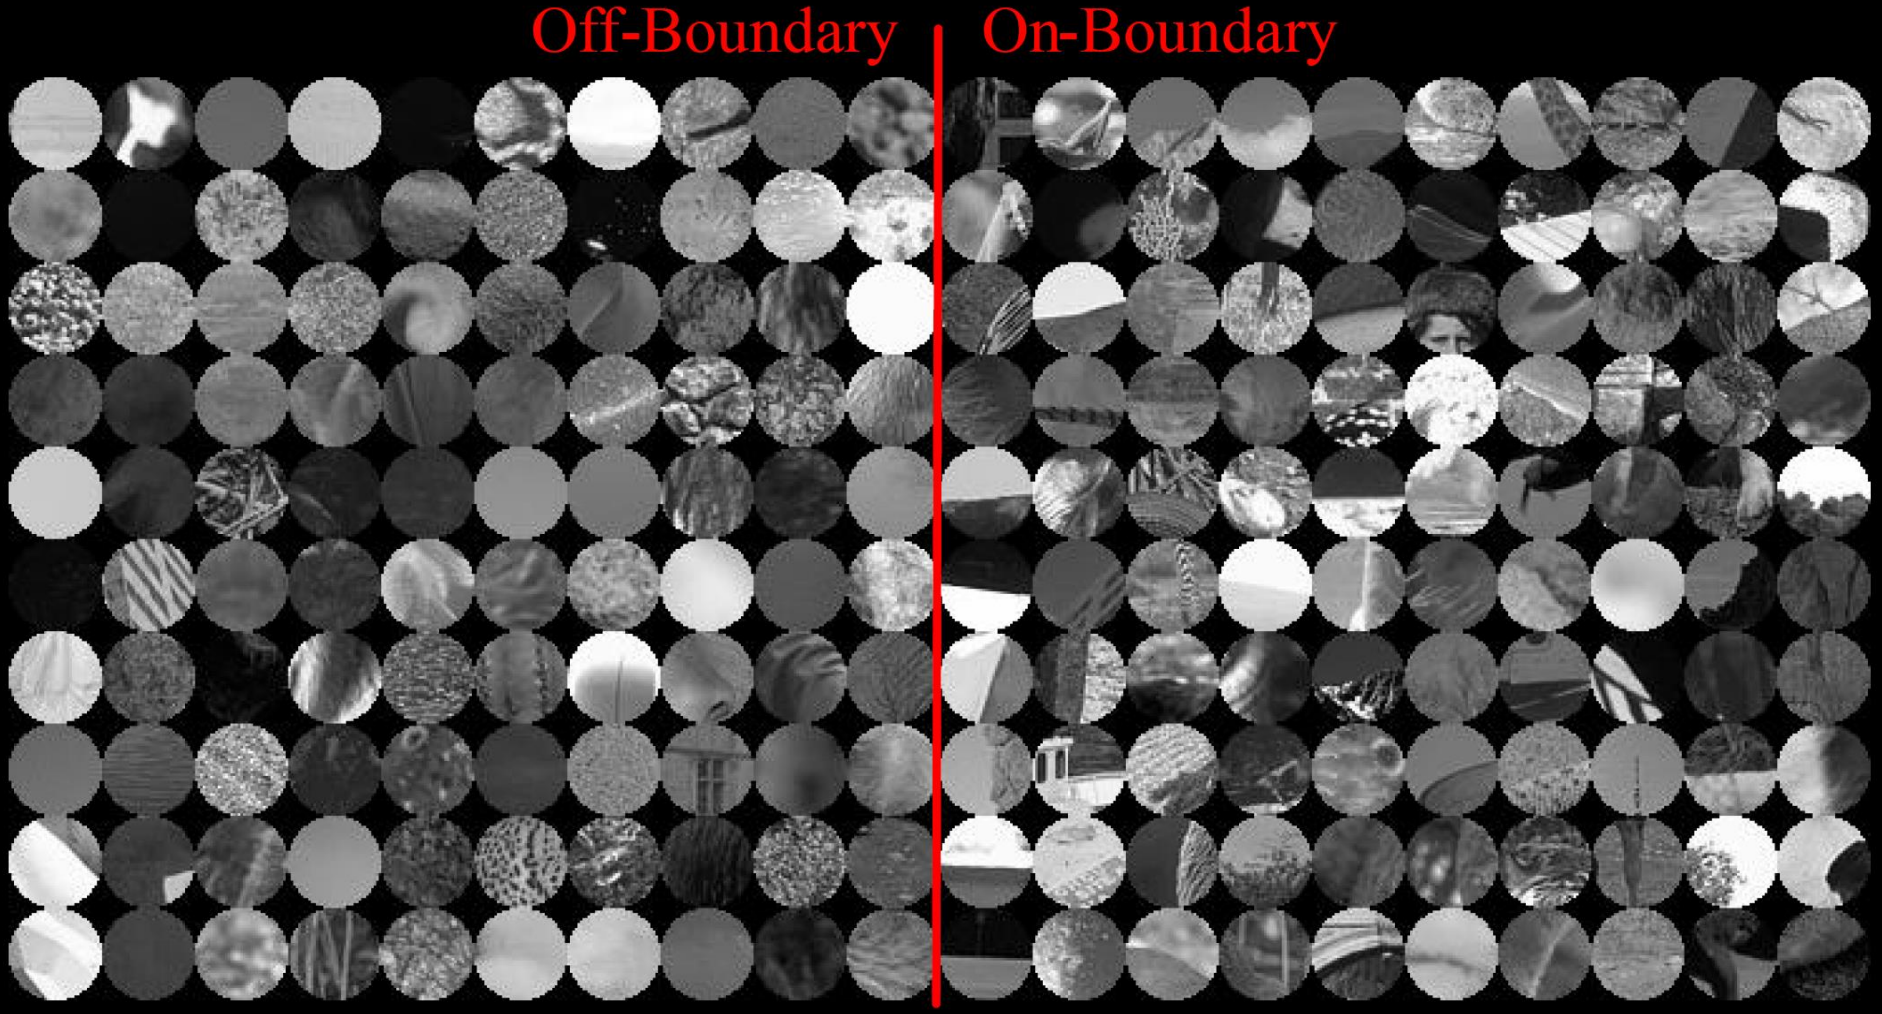

- Algorithm: r = 9, Humans: r = {5,9,18}
- Fixation(2s) -> Patch(200ms) -> Mask(1s)

Man versus Machine:

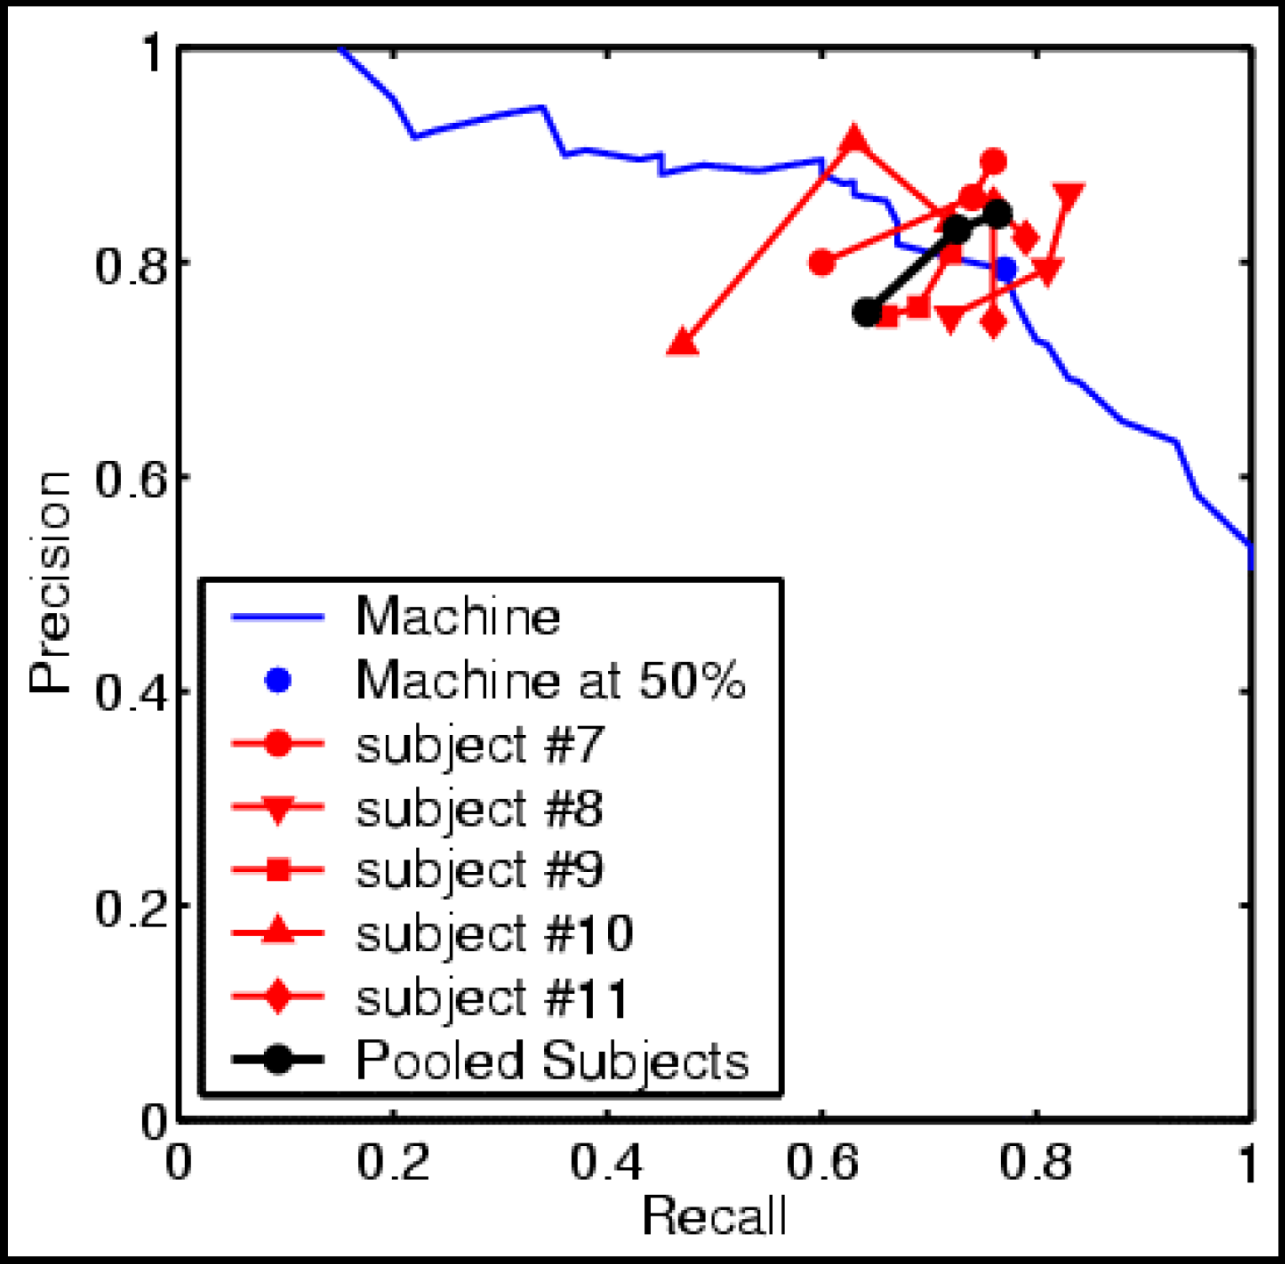

Sketch Tokens: A Learned Mid-level Representation for Contour and Object Detection

Joseph Lim, C Lawrence Zitnick, Piotr Dollar

CVPR 2013

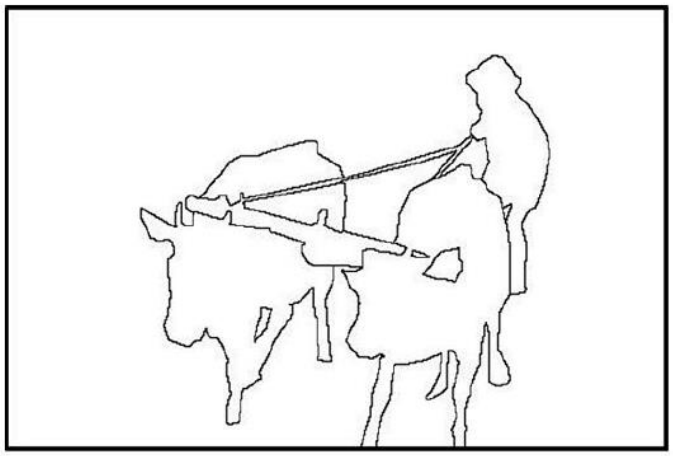

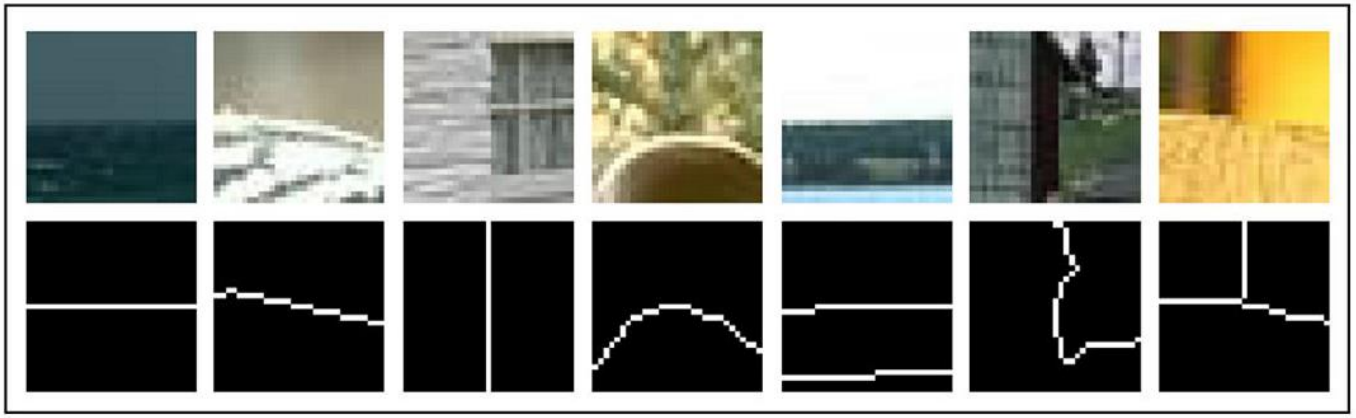

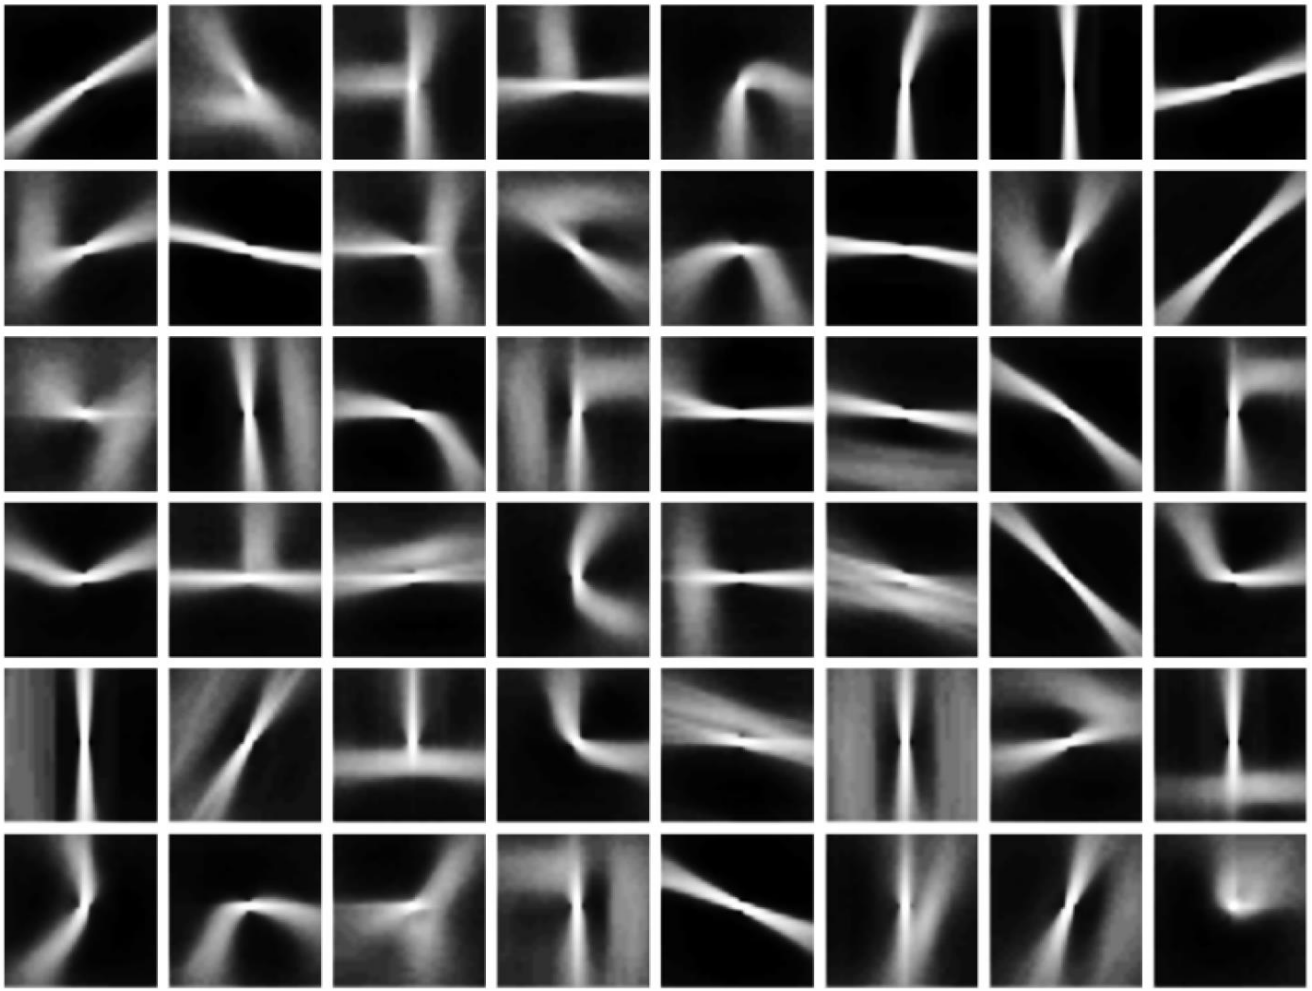

Image Features – 21350 dimensions!

- 35x35 patches centered at every pixel
- 35x35 “channels” of many types:
  - Color (3 channels)
  - Gradients (3 unoriented + 8 oriented channels)
    - Sigma = 0, T heta = 0, pi/2, pi, 3pi/2
    - Sigma = 1.5, Theta = 0, pi/2, pi, 3pi/2
    - Sigma = 5
  - Self Similarity
    - 5x5 maps of self similarity within the above channels for a particular anchor point.

Self-similarity features

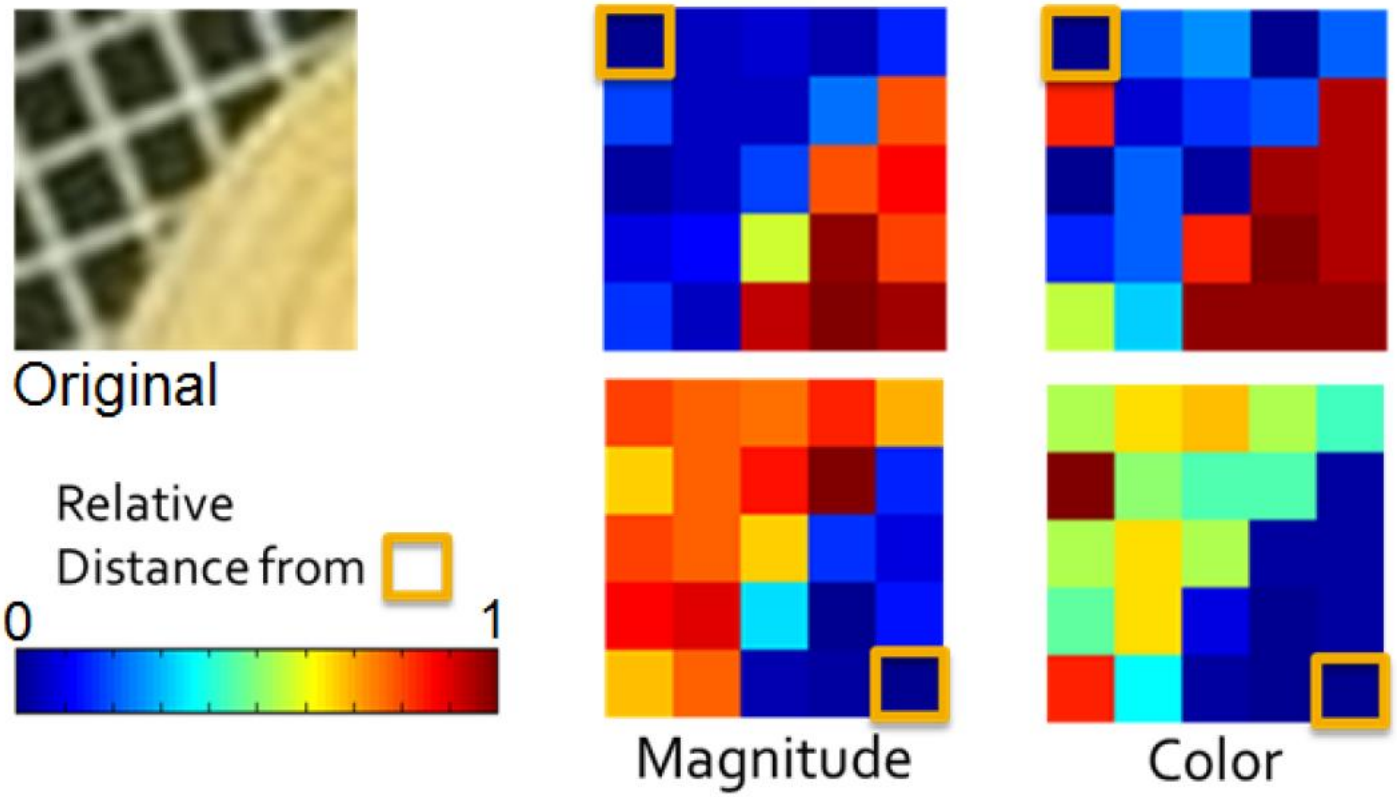

Self-similarity features: The L1 distance from the anchor cell (yellow box) to the other 5 x 5 cells are shown for color and gradient magnitude channels. The original patch is shown to the left.

Learning

- Random Forest Classifiers, one for each sketch token + background, trained 1-vs-all
- Advantages:
  - Fast at test time, especially for a non-linear classifier.
  - Don’t have to explicitly compute independent descriptors for every patch. Just look up what the decision tree wants to know at each branch.

Learning

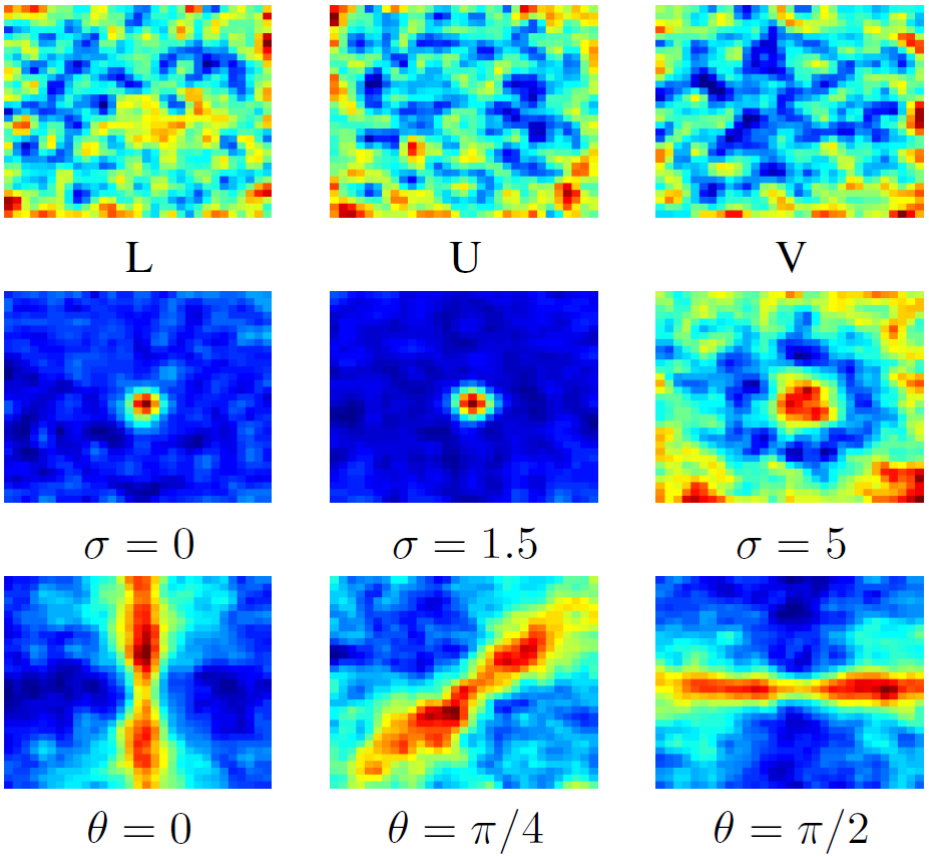

Frequency of example features being selected by the random forest: (first row) color channels, (second row) gradient magnitude channels, (third row) selected orientation channels.

Detections of individual sketch tokens

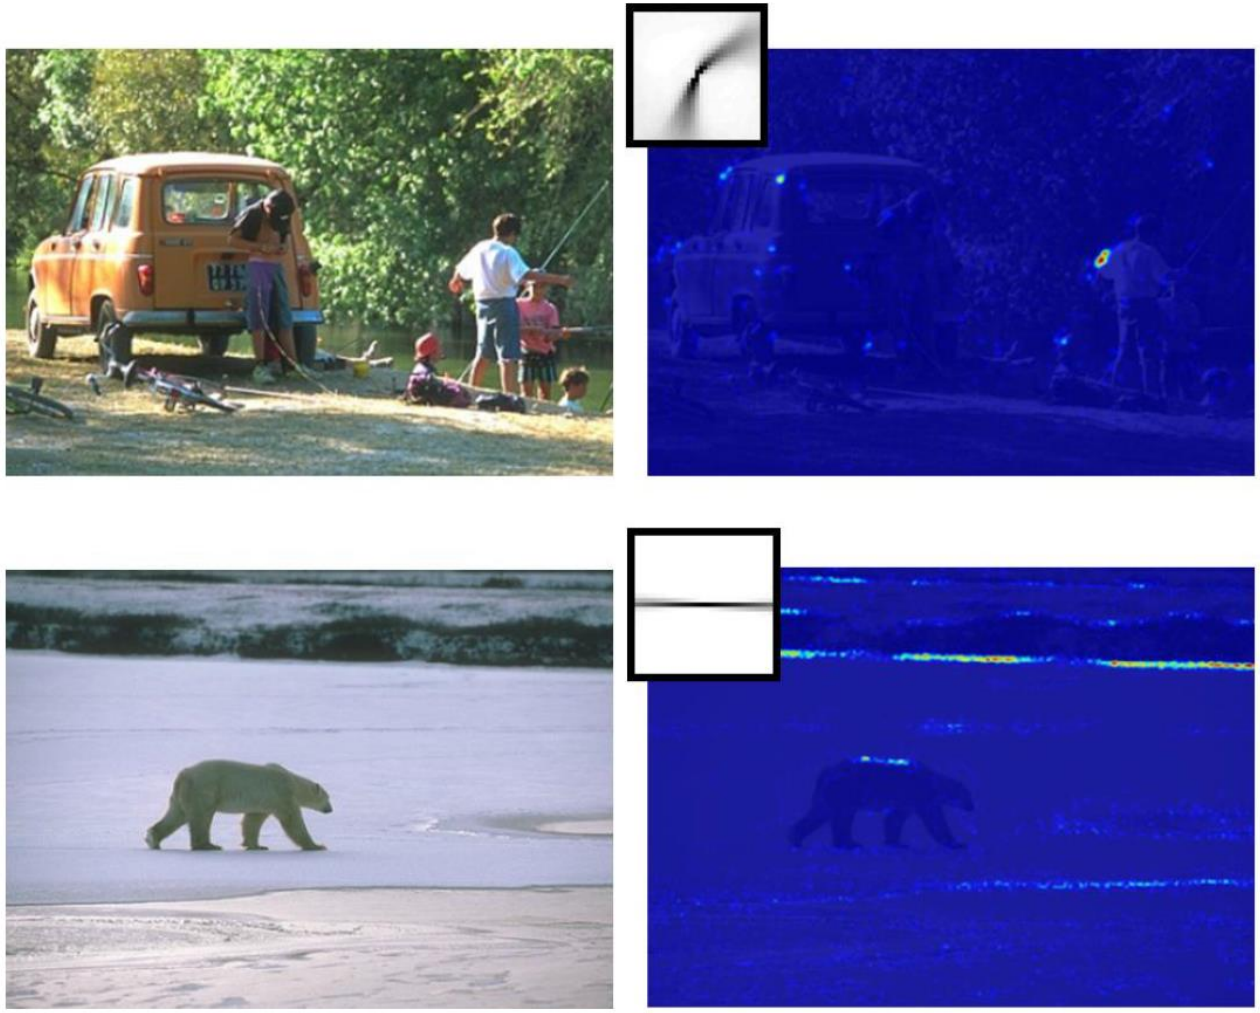

Detections of individual sketch tokens

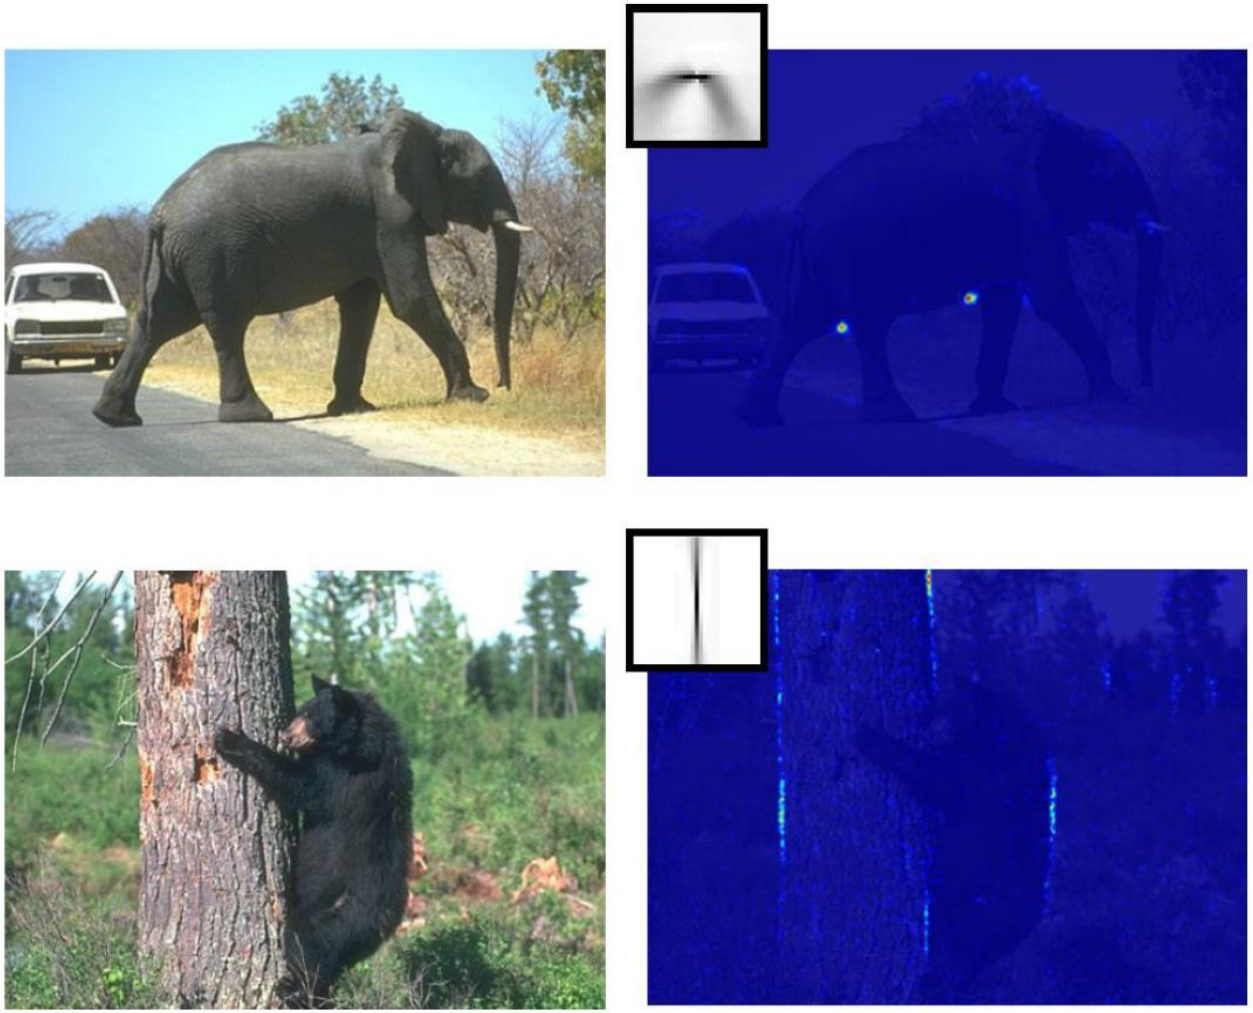

Combining sketch token detections

- Simply add the probability of all non-background sketch tokens

- Free parameter: number of sketch tokens
  - k = 1 works poorly, k = 16 and above work OK.


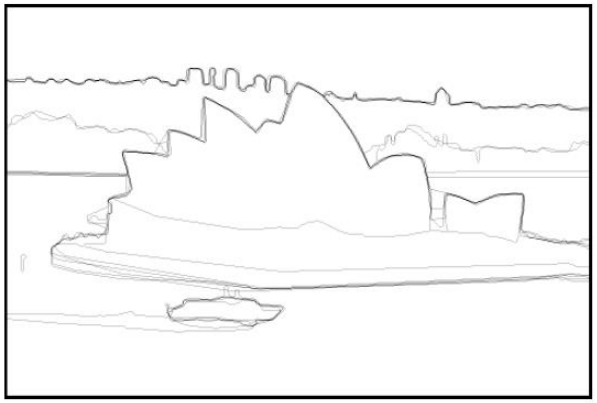
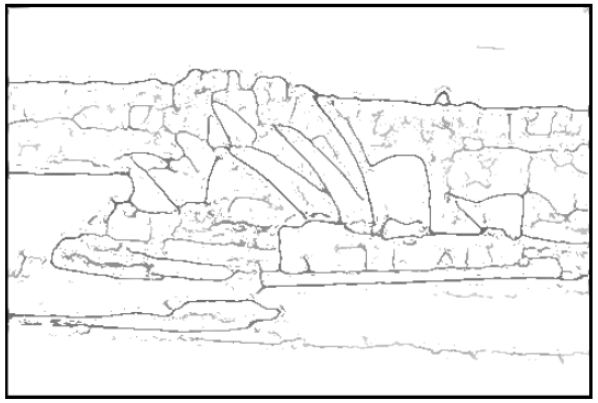

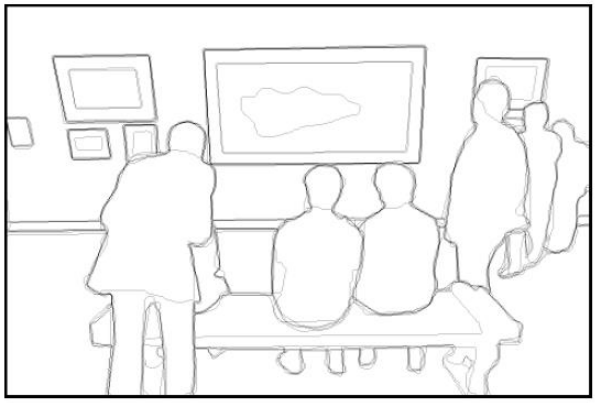
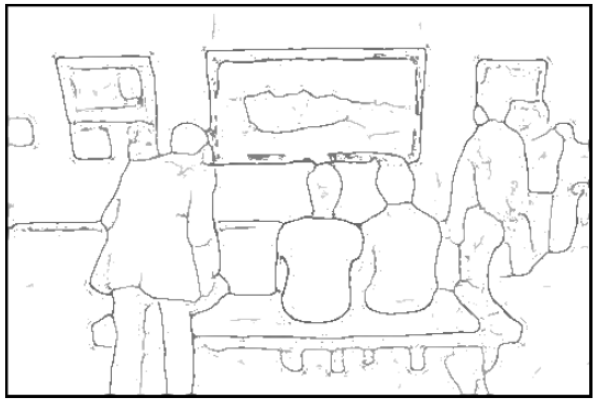

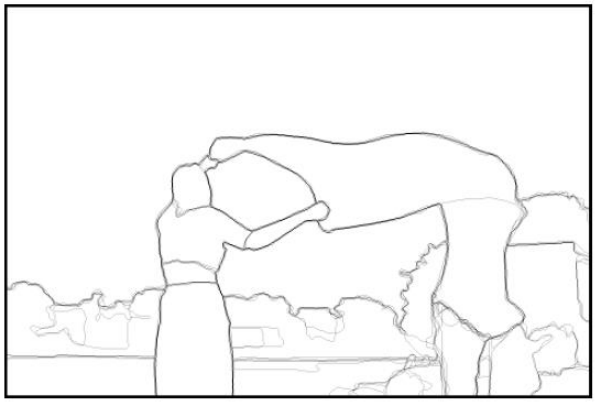
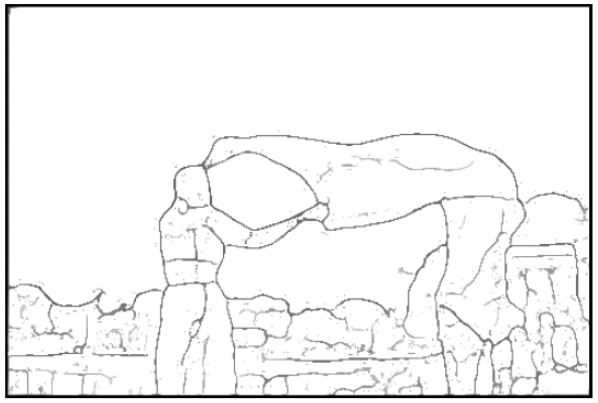

Input Image Ground Truth Sketch Tokens

Evaluation on BSDS

Method	ODS	OIS	AP	Speed
Human	.80	.80	-	-
Canny	.60	.64	.58	1/15 s
Felz-Hutt [12]	.61	.64	.56	1/10 s
gPb (local) [1]	.71	.74	.65	60 s
SCG (local) [24]	.72	.74	.75	100 s
Sketch tokens	**.73**	**.75**	**.78**	**1 s**
gPb (global) [1]	.73	.76	.73	240 s
SCG (global) [24]	.74	.76	.77	280 s

Evaluation on BSDS

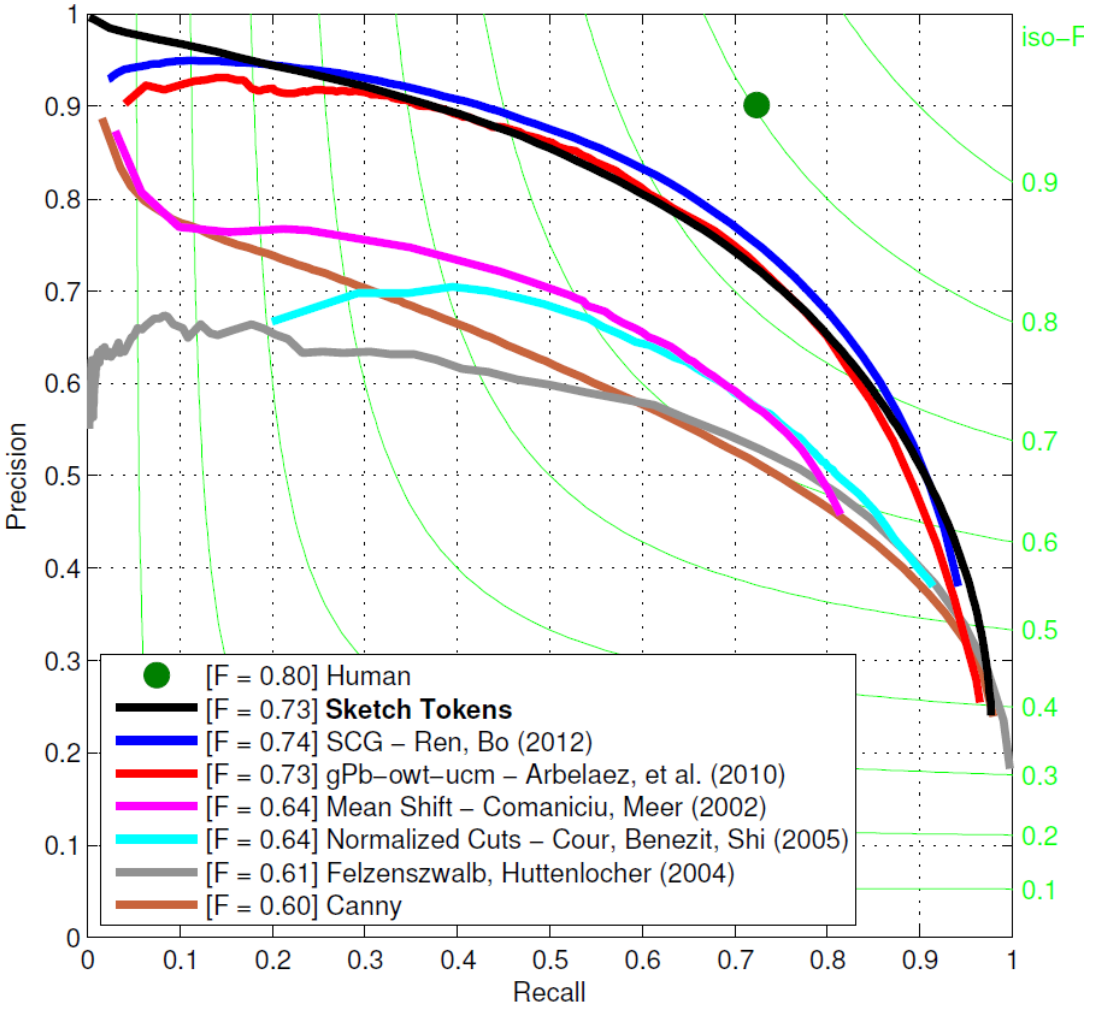

Summary

- Distinct from previous work, cluster the *human annotations* to discover the mid-level structures that you want to detect.
- Train a classifier for every sketch token.
- Is as accurate as any other method while being 200 times faster and using no global information.