Recap from Monday

- Visualizing Networks
- Caffe overview
- Slides are now online

Today

- Edges and Regions, GPB
- Fast Edge Detection Using Structured Forests
  - Zhihao Li
- Holistically-Nested Edge Detection
  - Yuxin Wu
- Selective Search for Object Recognition
  - Chun-Liang Li

Logistics

- Please read:
  - **Region-based Convolutional Networks for Accurate Object Detection and Semantic Segmentation**
- If you're up next, please meet us
- Project Proposals Due in < 1 week
  - If you have questions, ask to meet

Edges and Regions

David Fouhey

Task

"I stand at the window and see a house, trees, sky. Theoretically I might say there were 327 brightnesses and nuances of colour. Do I have "327"? No. I have sky, house, and trees."

-Max Wertheimer

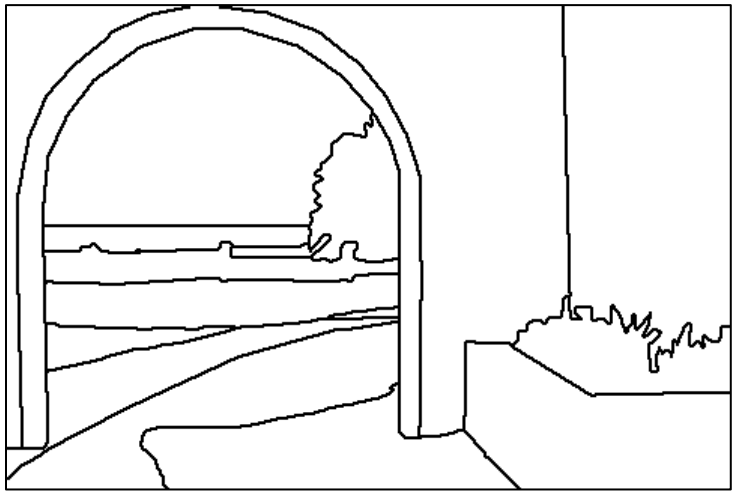

Quote from Jitendra Malik's page

Approaching the Task – Regions

Decomposing image into K connected regions
(Clustering task)

Approaching the Task – Edges

HxWx {0,1} classification problem

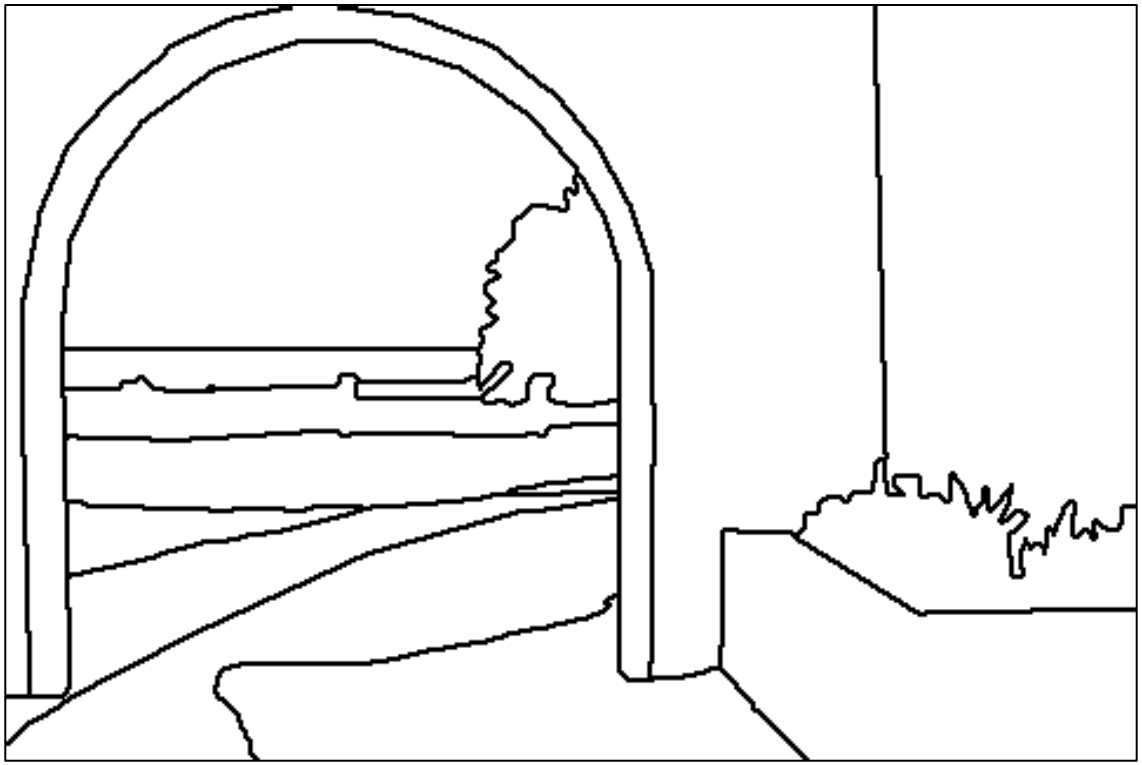

Are the Tasks Equivalent?

Segmentation

?

Boundaries

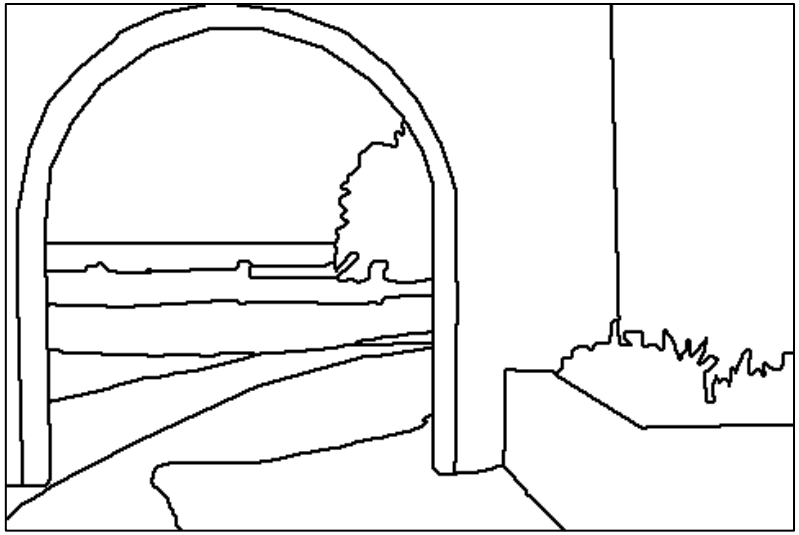

Are the Tasks Equivalent?

Boundaries

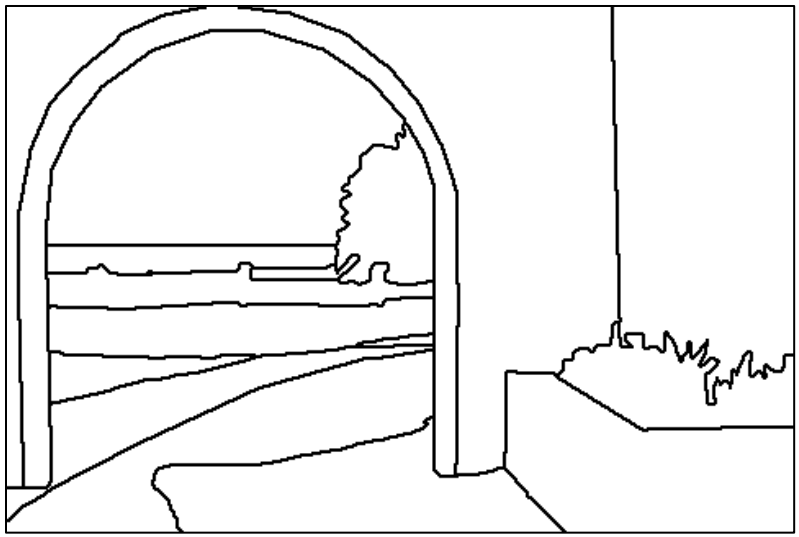

?

Segmentation

Are the Tasks Equivalent?

Boundaries

Segmentation

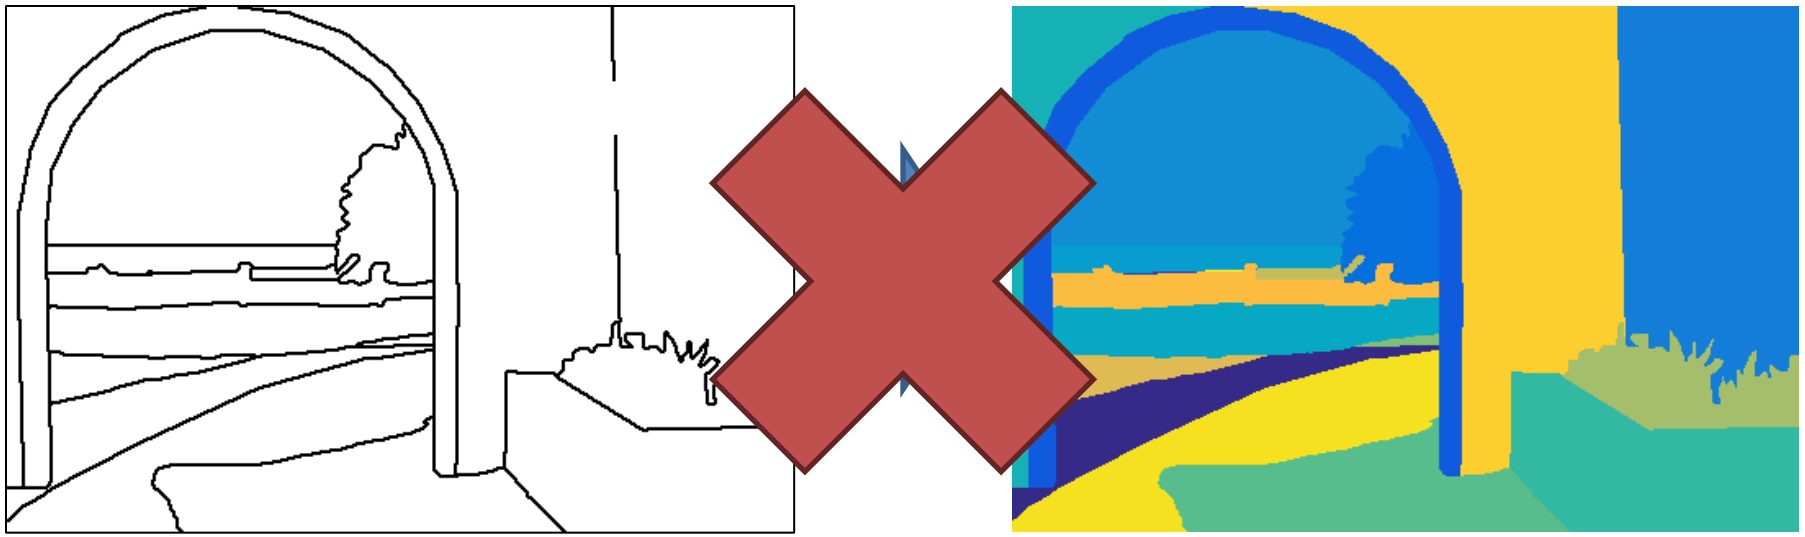

Contours have to be closed!

Does This Matter in the CNN Era?

HED – State of the Art

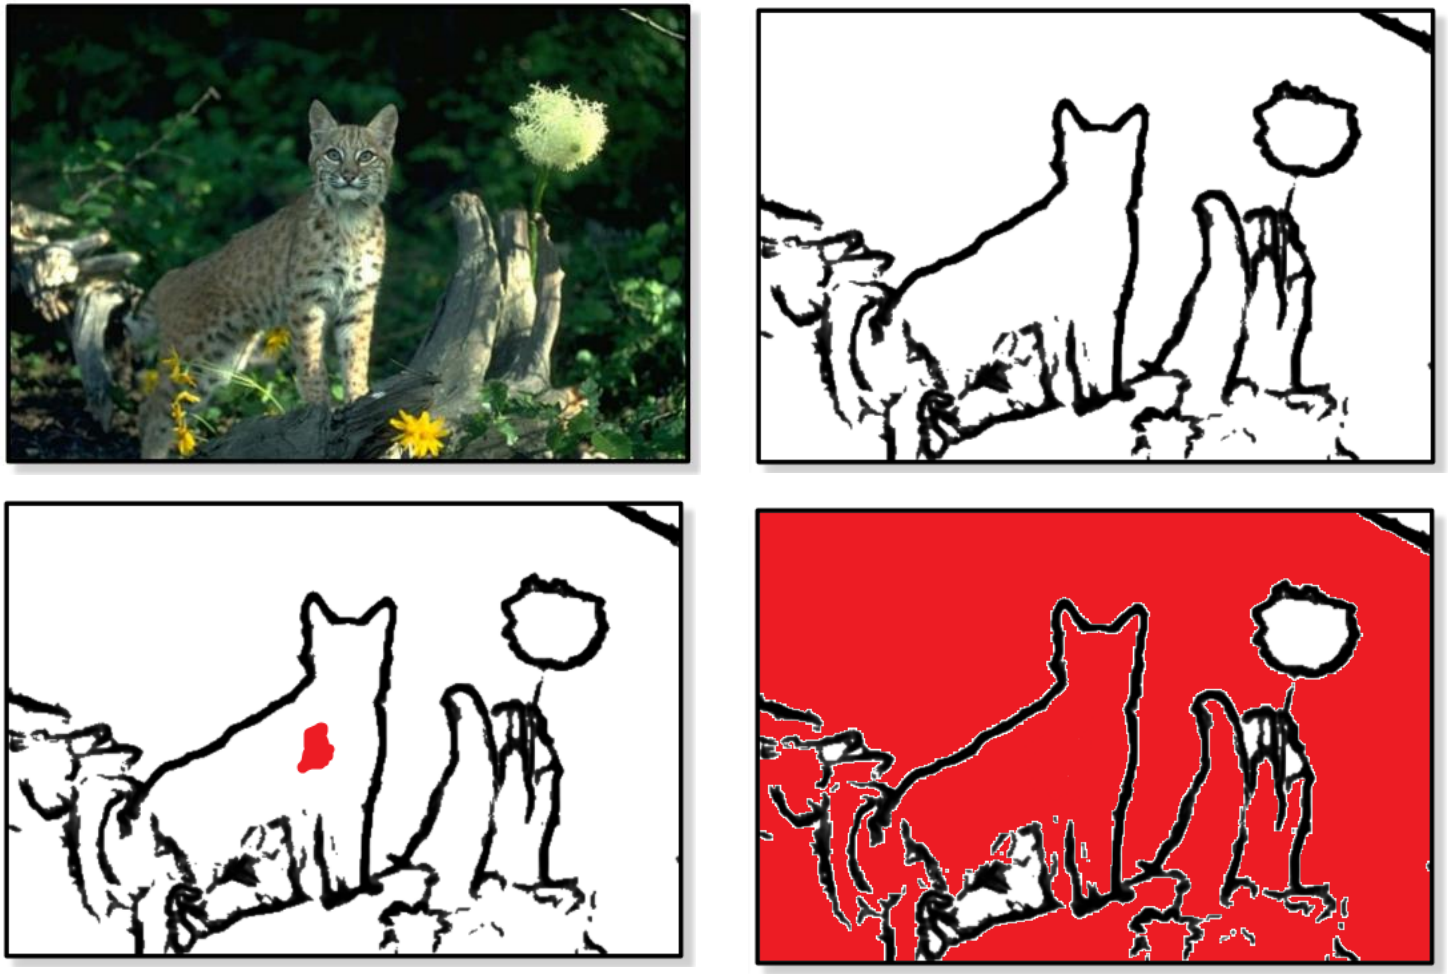

Are These Well-Defined Tasks?

Should blue and yellow go in the same segment?

Image credit: NYU depth dataset

Successes – Superpixels

Problem: >10^5 pixels intractable for reasoning

Solution: use bigger/super pixels that don't ruin any boundaries

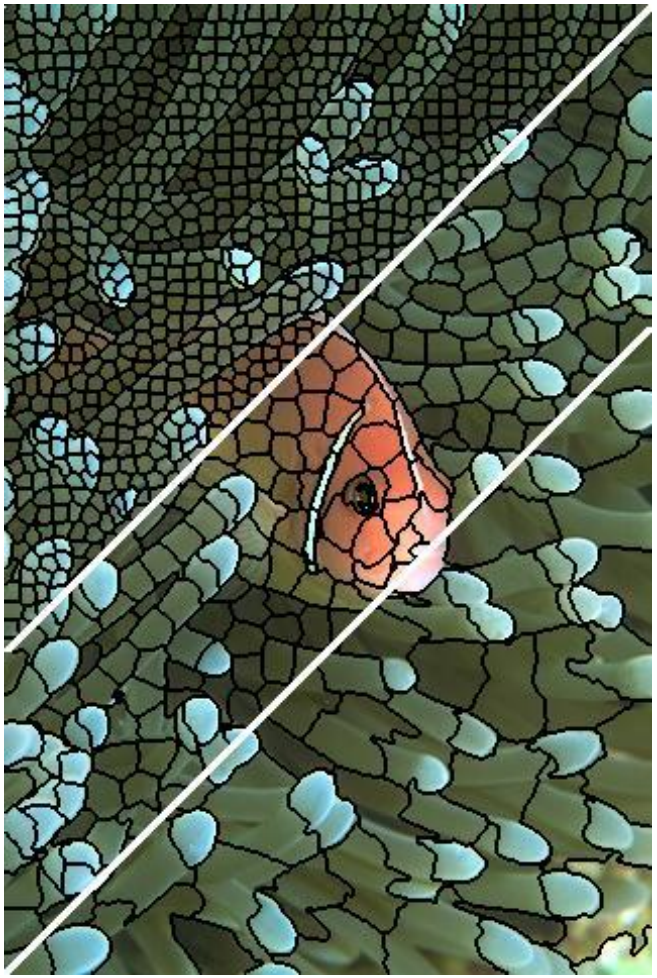

First from Ren et al. 2003, Fish image from Achanta et al. 2012

Successes – Multiple Segmentations

- **Problem**: No one segmentation is good
- **Solution**: Use many, figure it out later

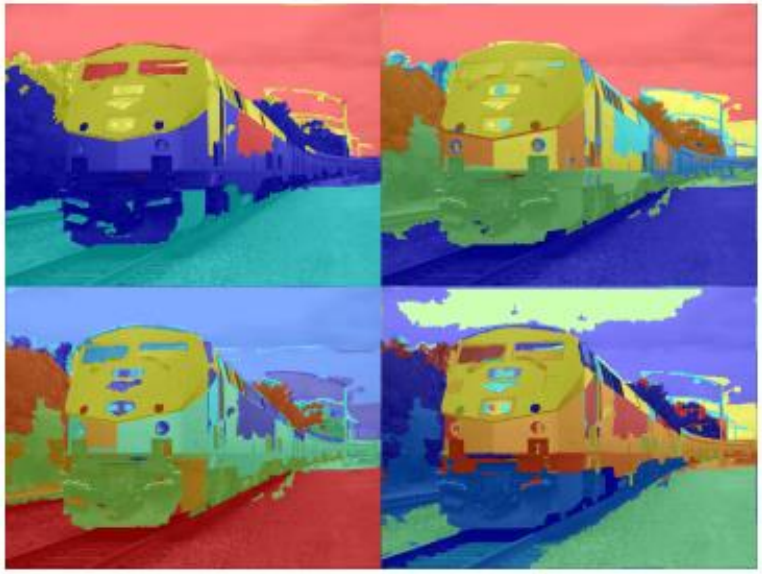

Hoiem et al. 2005

Contributions of Paper

- Merges the (edges + regions) approaches
- Introduces machinery used throughout vision
- Landmark paper in segmentation/boundary detection
- **Note:** the questions are often as important as the answers

Questions from Piazza

- Where's the learning?!
  - Great idea! Two papers next
- What's this useful for?
  - Great question! Last paper today, paper for Monday.

Dataset – BSDS 500

Images

- 500 Total
- 300 Training, 200 Testing

Annotation

- 5 annotators (CV students) per image
- Annotators annotate segment

Dataset – Instructions

Divide each image into pieces, where each piece represents a distinguished thing in the image. It is important that all of the pieces have approximately equal importance. The number of things in each image is up to you. Something between 2 and 20 should be reasonable for any of our images

Martin et al. “A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics.” ICCV 2001

Dataset – Image and Annotations

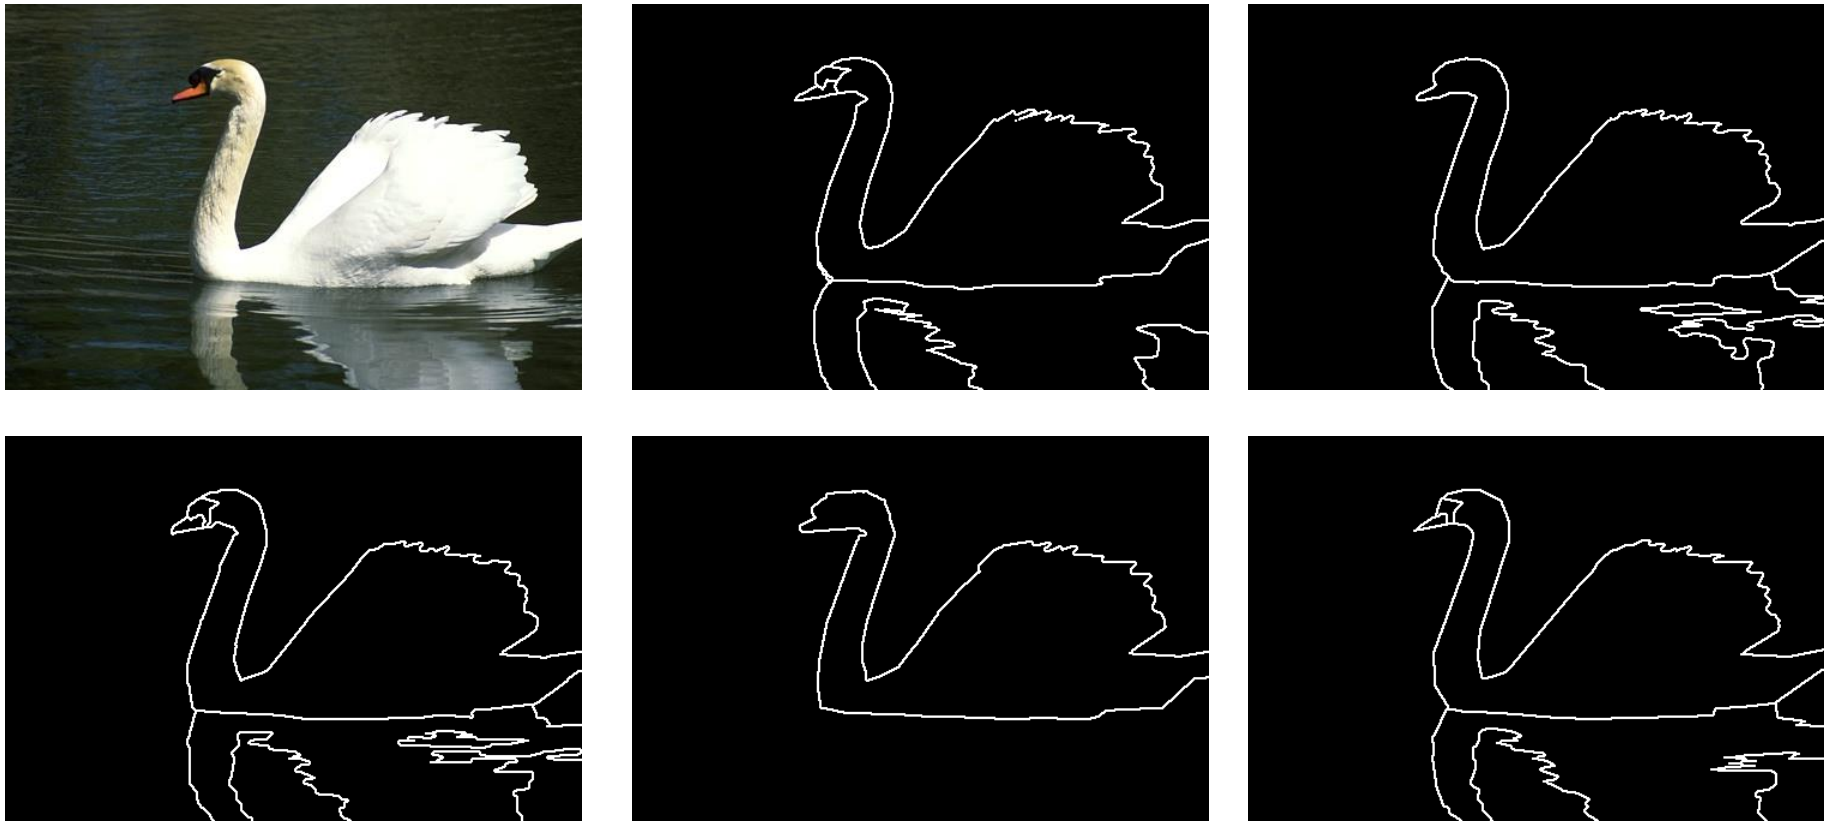

Evaluation Criteria – Boundaries

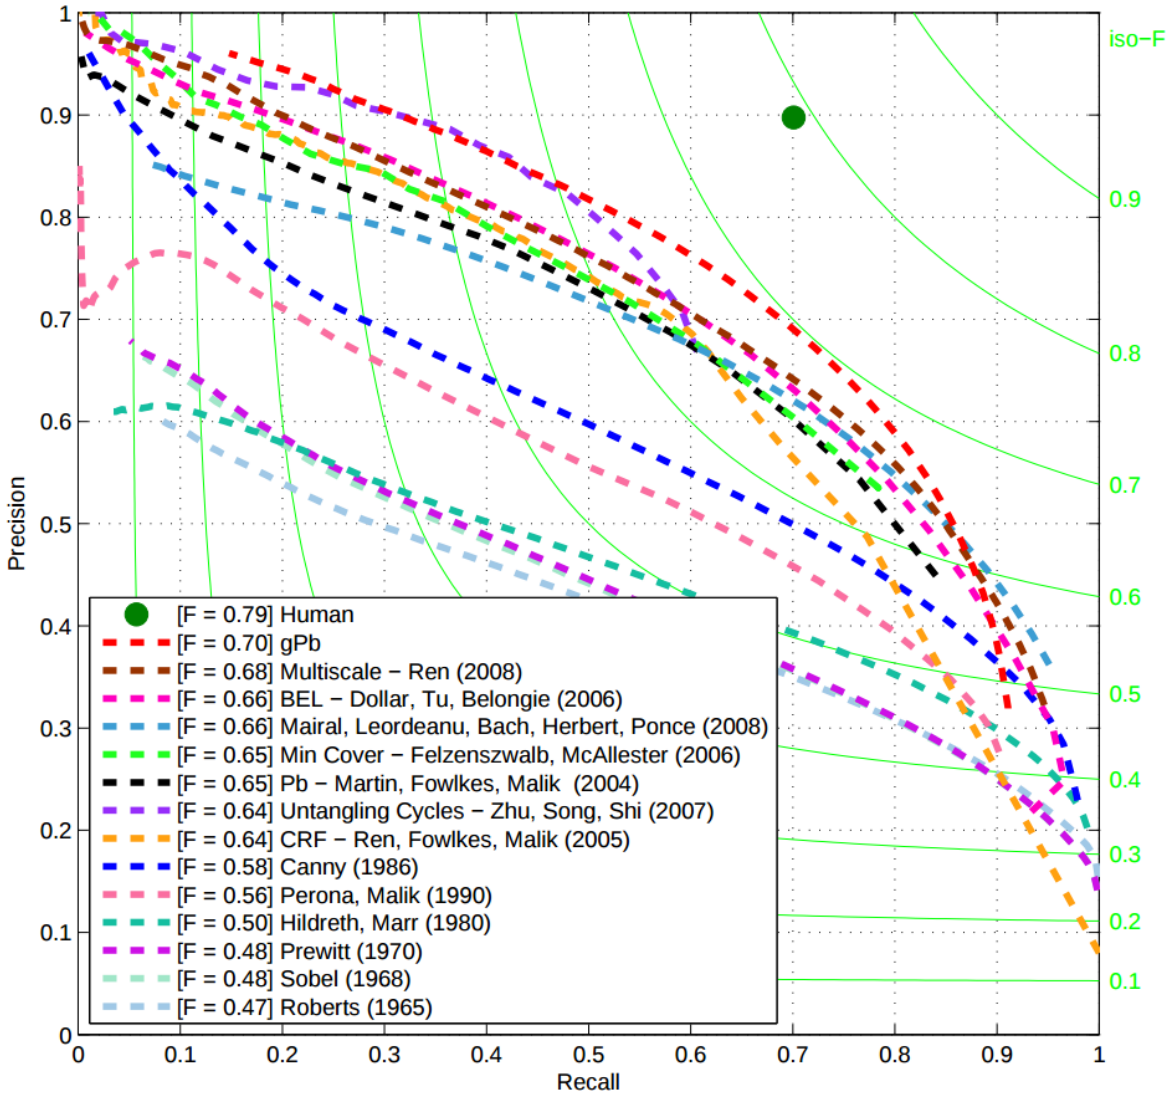

Precision
= TP / (TP+FP)
(fraction of predicted + results that are +)

Recall
= TP / (TP+FN)
(fraction of + results that are predicted +)

Evaluation Criteria – Segments

- In words: Average the intersection/union of the best predicted region for all GT regions, weighted by GT region size

$$\mathcal{C}(S' \to S) = \frac{1}{N} \sum_{R \in S} |R| \cdot \max_{R' \in S'} \mathcal{O}(R, R') \quad \mathcal{O}(R, R') = \frac{|R \cap R'|}{|R \cup R'|}$$

- Previous evaluation criteria don't clearly distinguish dumb baselines from algorithm outputs

GPB-OWT-UCM

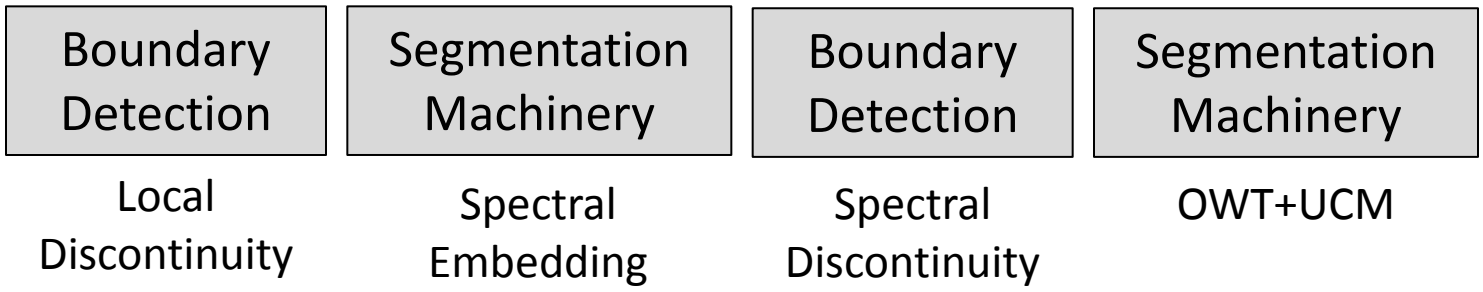
Boundary Detection
Segmentation Machinery
Boundary Detection
Segmentation Machinery
Local Discontinuity
Spectral Embedding
Spectral Discontinuity
OWT+UCM

Local Terms

- **Core Idea:** can compute histogram distances

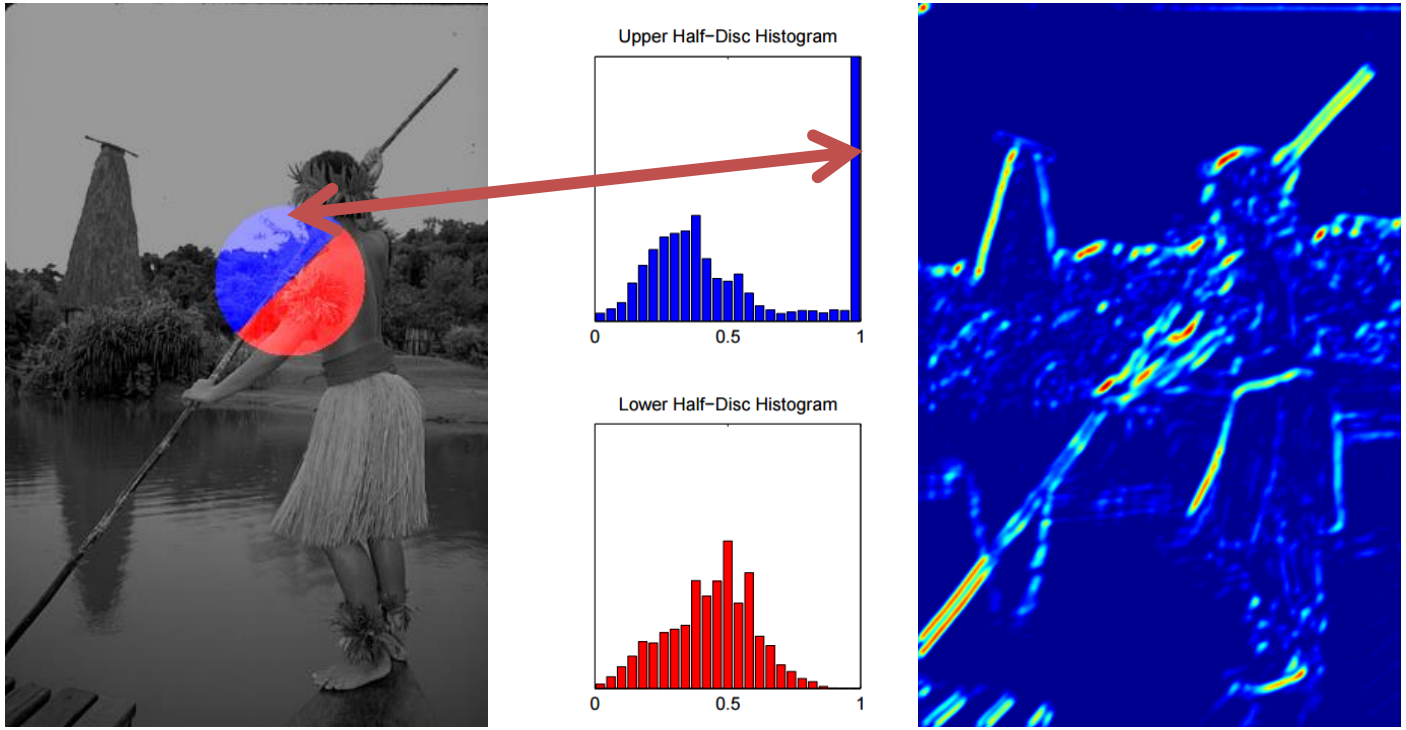

Local Terms

Luminance Image

Orientation 1

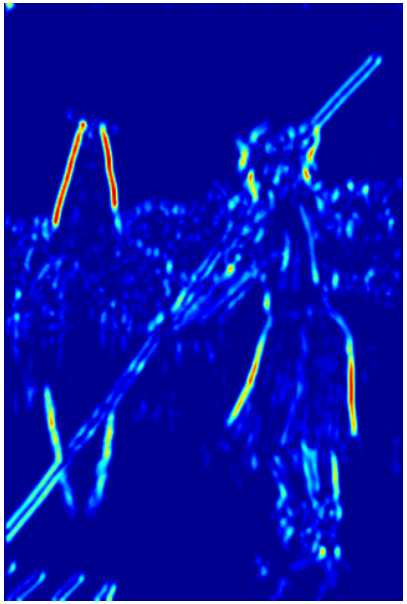

Orientation 2

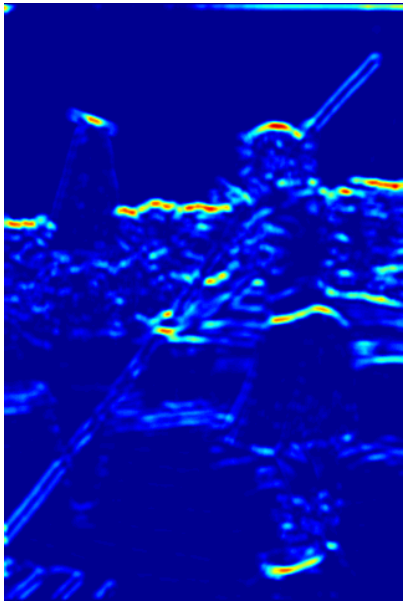

Max over Orientations

Local Terms

Luminance Image

Max over Orientations

Local Terms – Multiple Cues

Accumulate evidence per-orientation

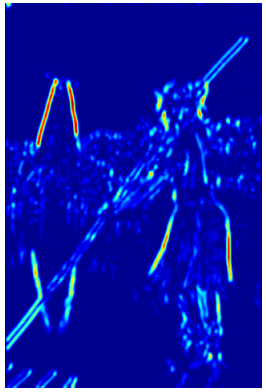

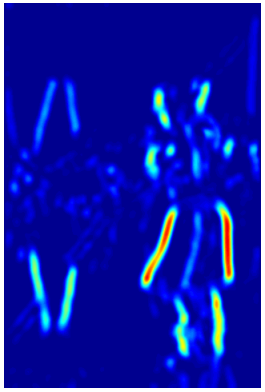

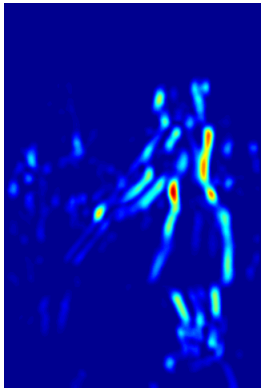

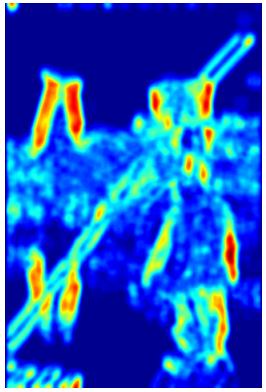

Weighted Sum of Predictions

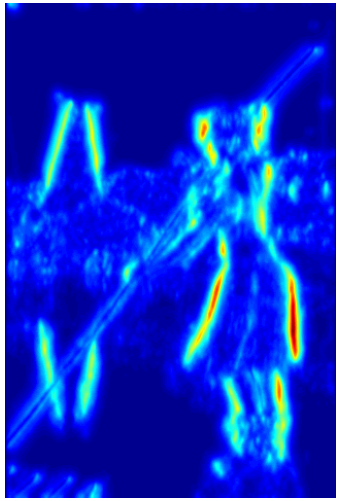

Learning

- Simple linear combinations = few parameters
- Gradient ascent in the reading
- Logistic regression in past

weights

Contour strength
in feature + scale

$$mPb(x, y, \theta) = \sum_s \sum_i \alpha_{i,s} G_{i,\sigma(i,s)}(x, y, \theta)$$

GPB-OWT-UCM

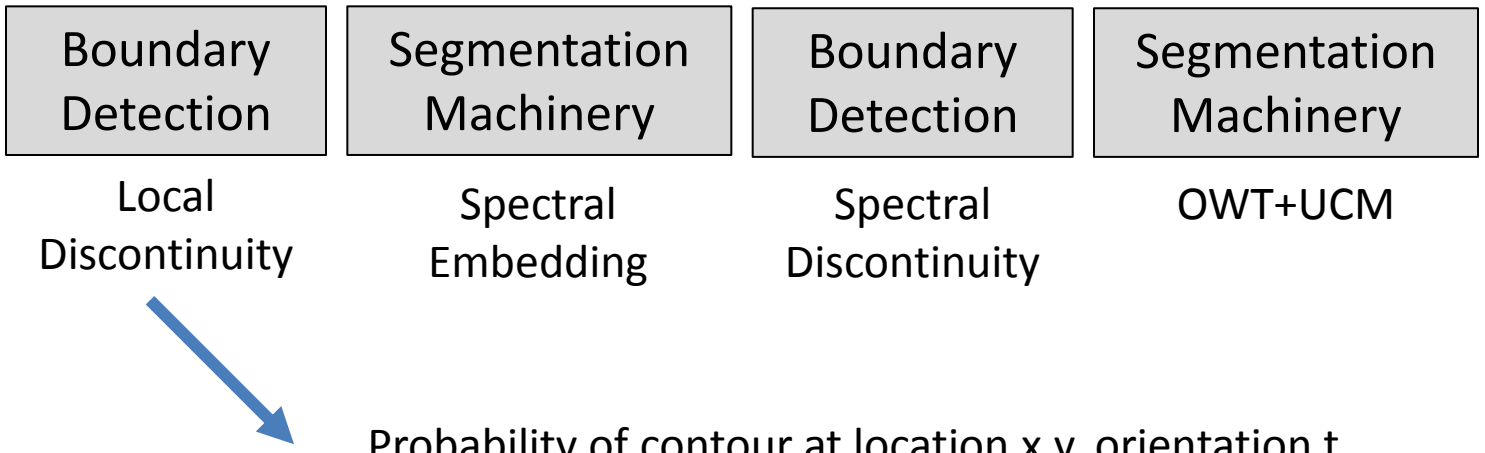
Boundary Detection
Segmentation Machinery
Boundary Detection
Segmentation Machinery
Local Discontinuity
Spectral Embedding
Spectral Discontinuity
OWT+UCM
Probability of contour at location x,y, orientation t

Globalization – Motivation

Local

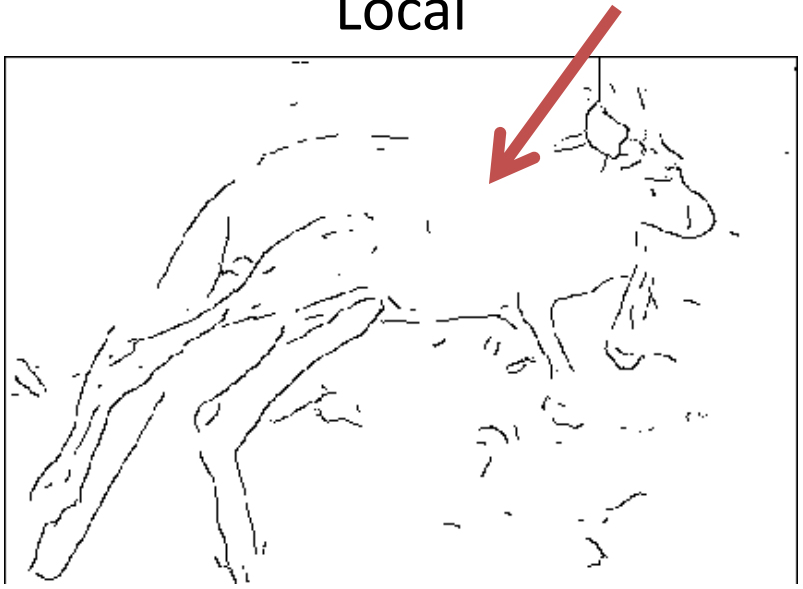

Globalized

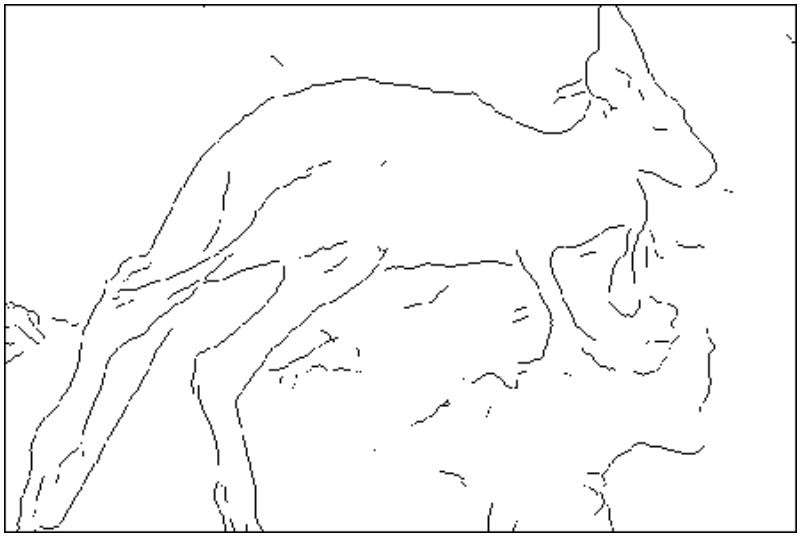

Globalization

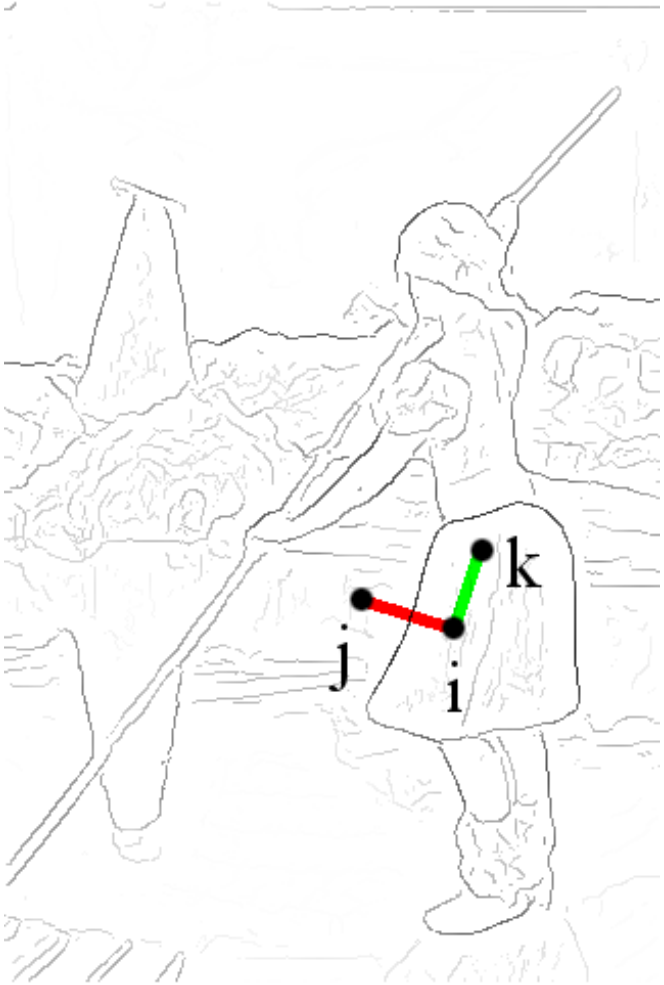

$$W \in R^{(HW)x(HW)}$$

Normal Spectral Clustering

1. Use W to produce embedding/space defined by eigenvectors of a system of equations. *See links on Piazza for why*
2. Cluster in induced space

This Paper

1. Use W to produce embedding/space defined by eigenvectors of a system of equations
2. Treat eigenvectors as images, compute gradient

Globalization

Input

Eigenvectors of Spectral System

Weighted
Sum of Gradients

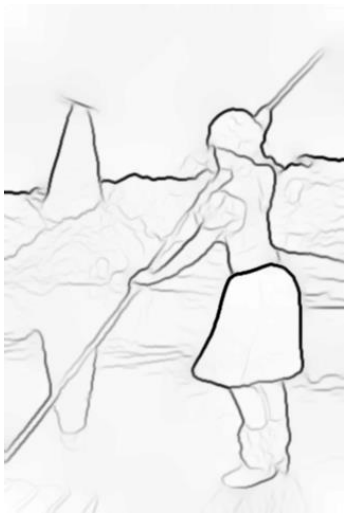

Combining Global + Local

- Linear weighting; weights learned with gradient ascent

$$gPb(x, y, \theta) = \sum_s \sum_i \beta_{i,s} G_{i,\sigma(i,s)}(x, y, \theta) + \gamma \cdot sPb(x, y, \theta)$$

Orientations processed **separately** throughout

Why is this important?

GPB-OWT-UCM

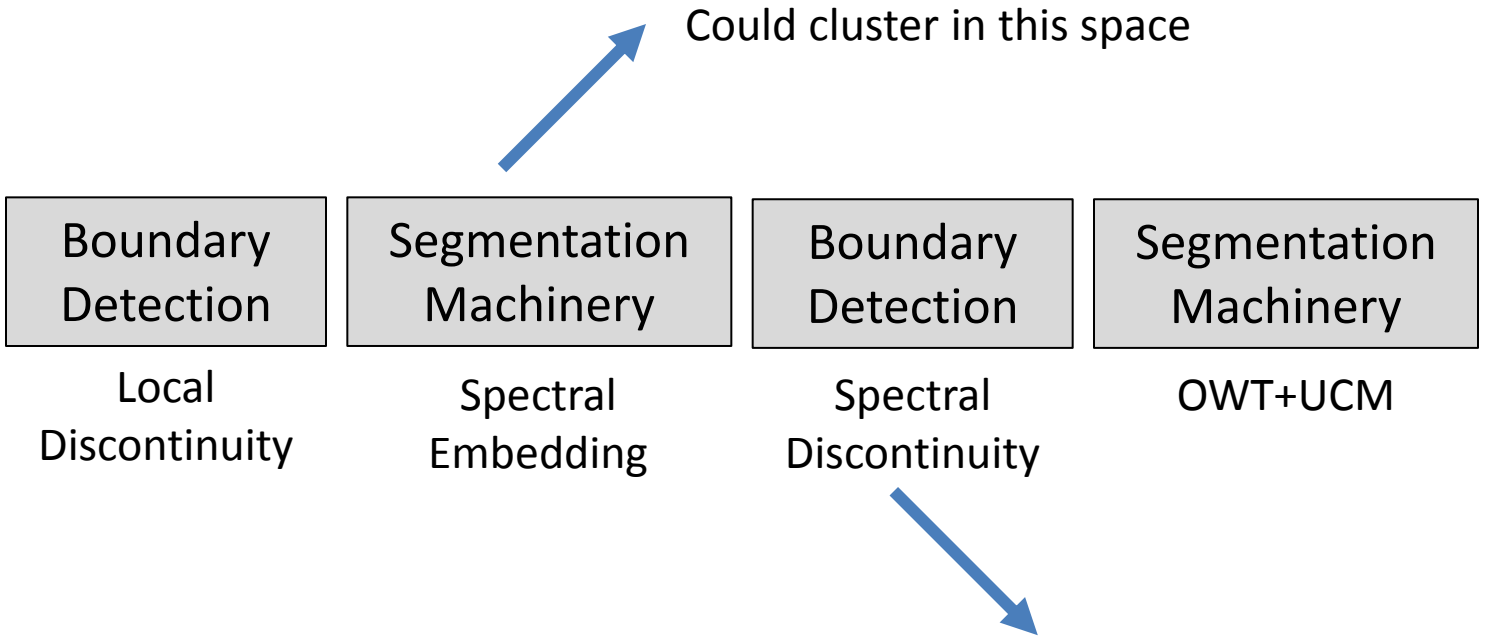
Could cluster in this space
Boundary Detection
Segmentation Machinery
Boundary Detection
Segmentation Machinery
Local Discontinuity
Spectral Embedding
Spectral Discontinuity
OWT+UCM
Probability of contour at location x,y, orientation t taking into consideration soft segmentations

Watershed Transform – 1D Version

- **Black region:** probability of boundary
- **Black lines:** watershed boundaries

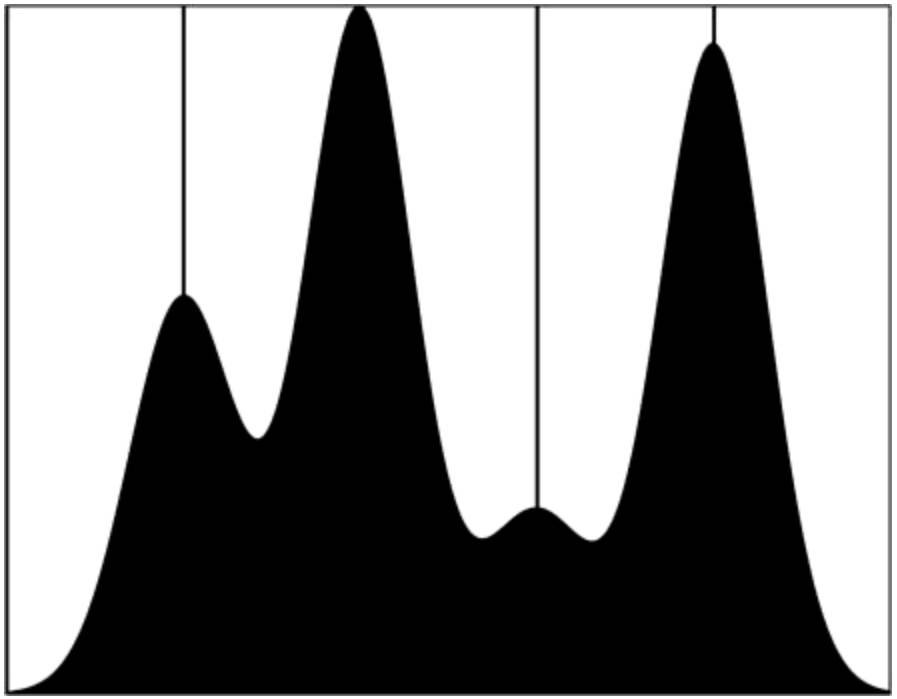

Orientation

Problem: probability of boundary is orientation-dependent

Solution: get probability of boundary in direction

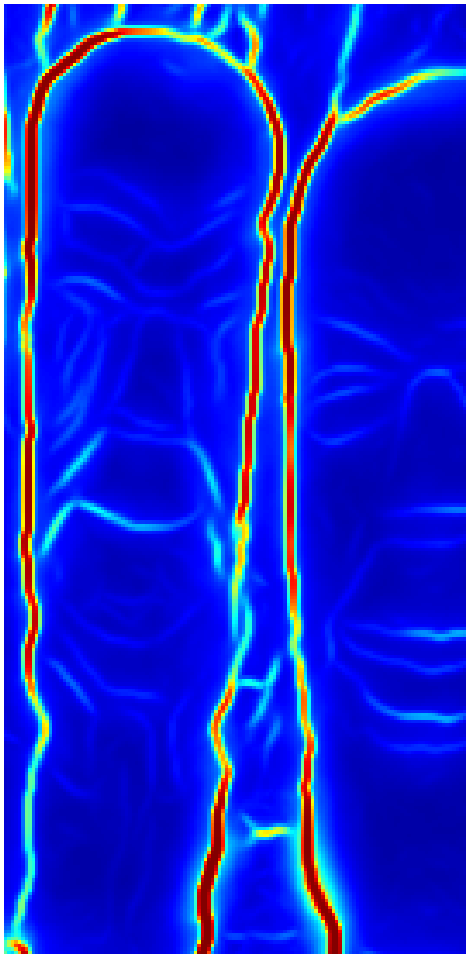

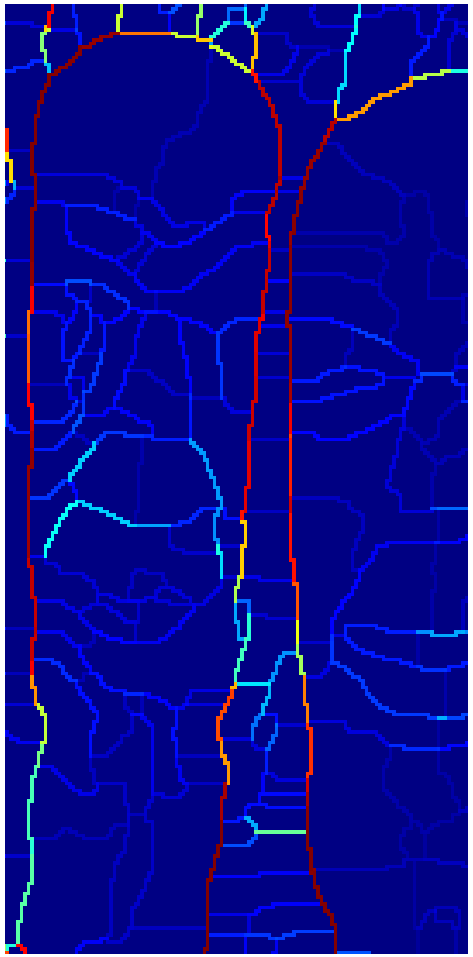

Output of Watershed Transform

"Oversegmentation" of image
with boundary strengths

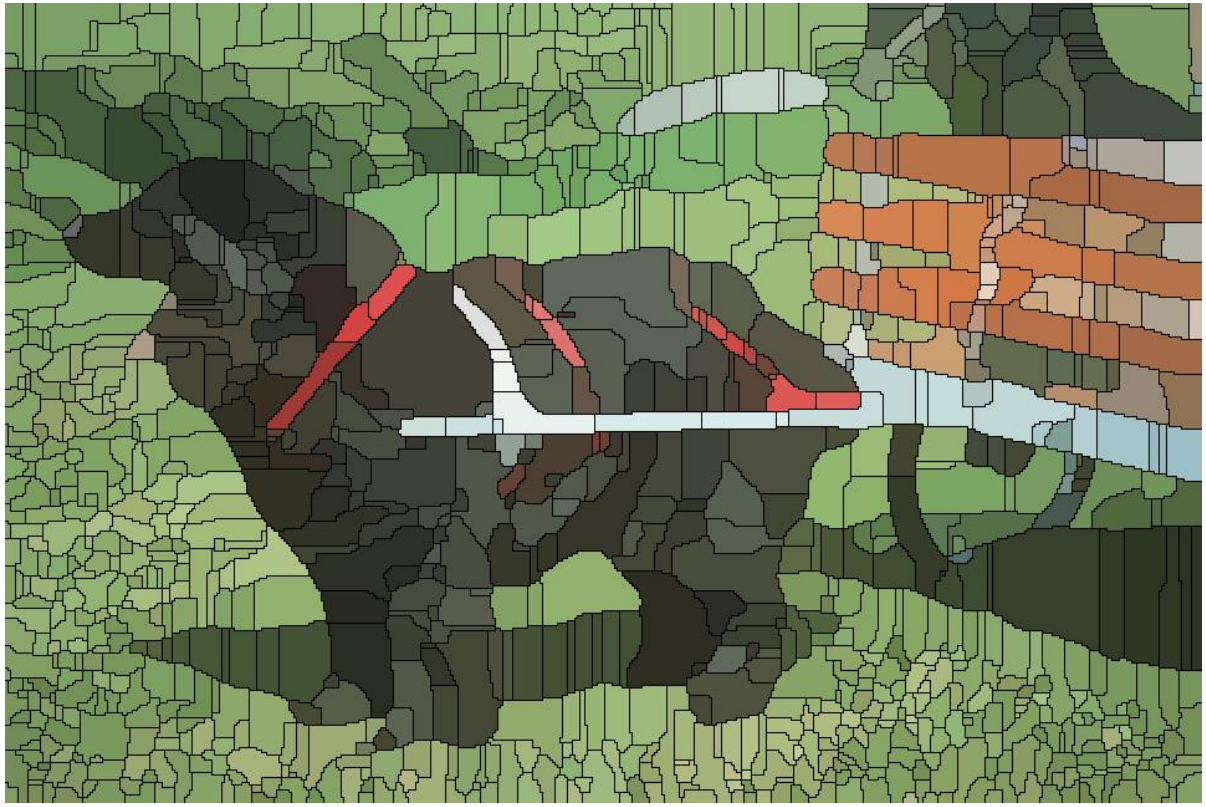

UCM

- Hierarchical merging; guarantees closed contours

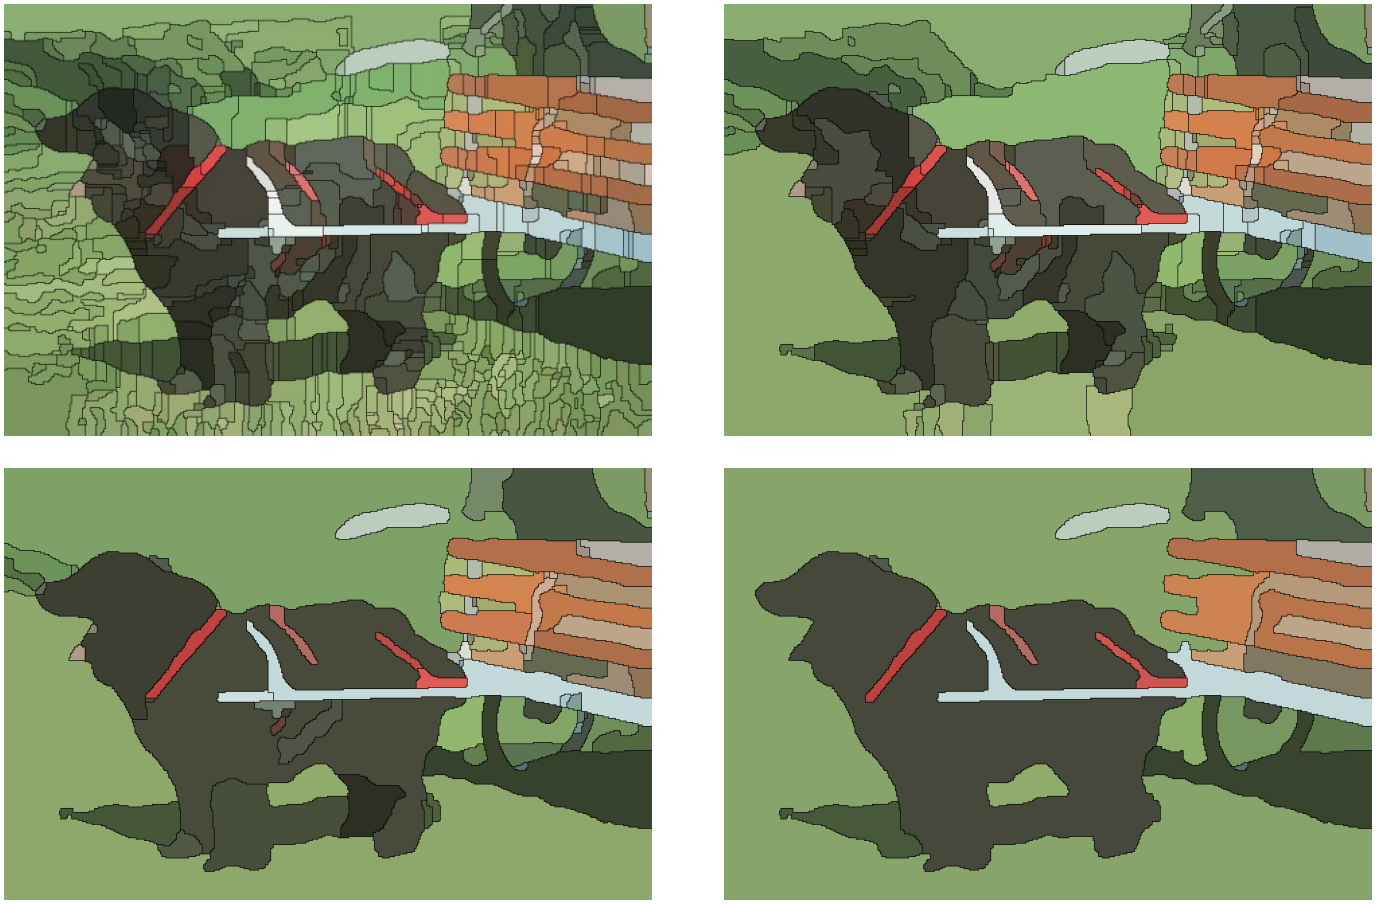

GPB-OWT-UCM

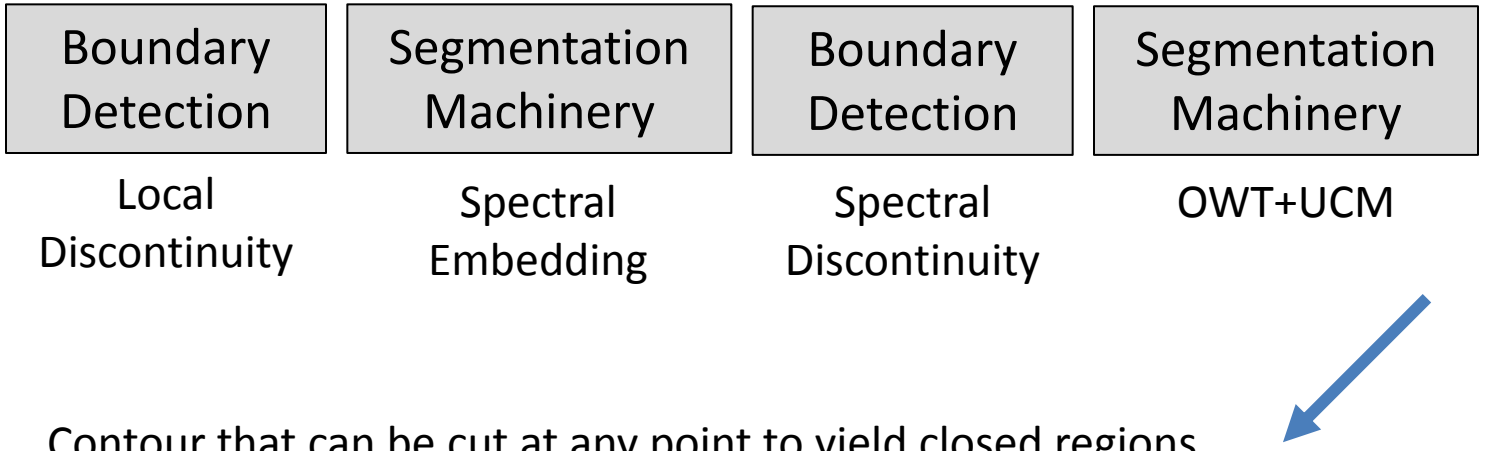
Boundary Detection
Segmentation Machinery
Boundary Detection
Segmentation Machinery
Local Discontinuity
Spectral Embedding
Spectral Discontinuity
OWT+UCM
Contour that can be cut at any point to yield closed regions

Results – State of the Art

This: 72.6

Current SOA: 78.2

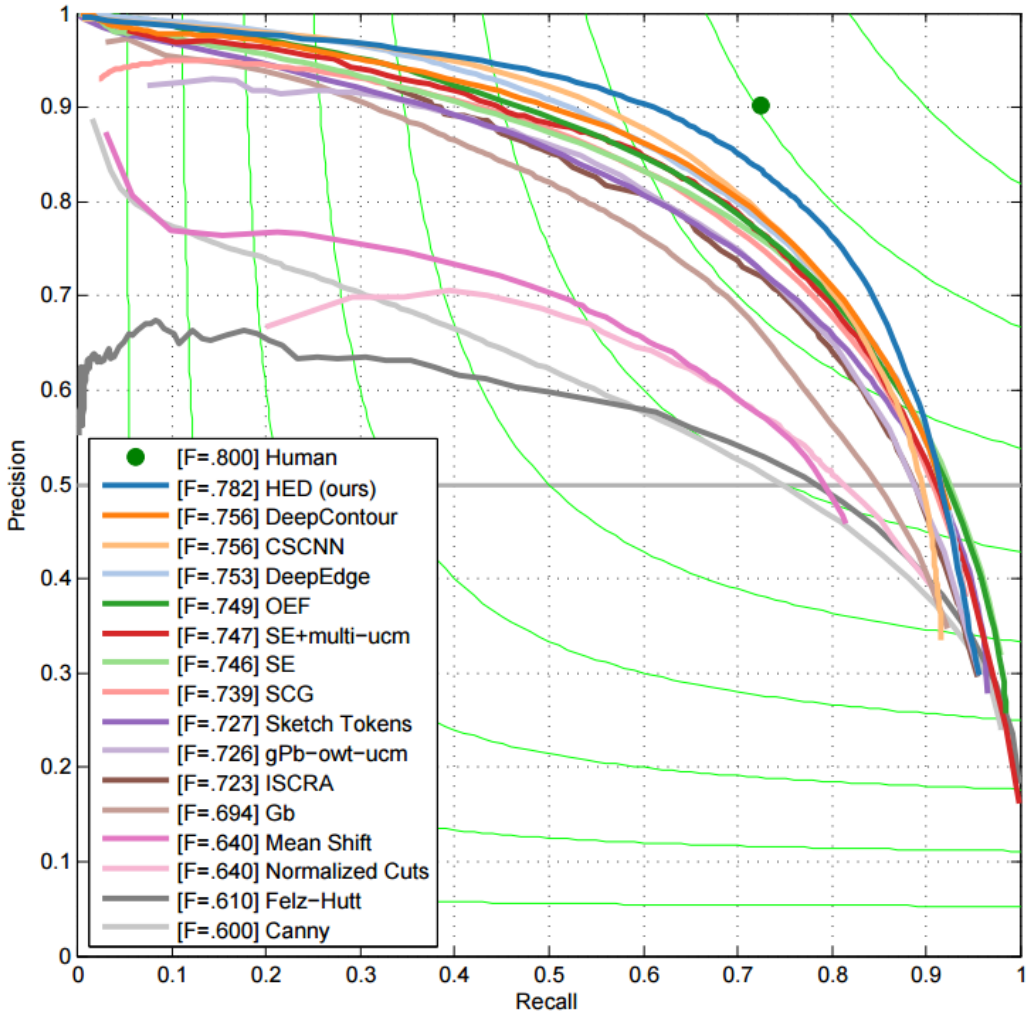

Results – Ablative Analysis

- Combining Local + Global helps
- **Why does local help in high-recall regime?**

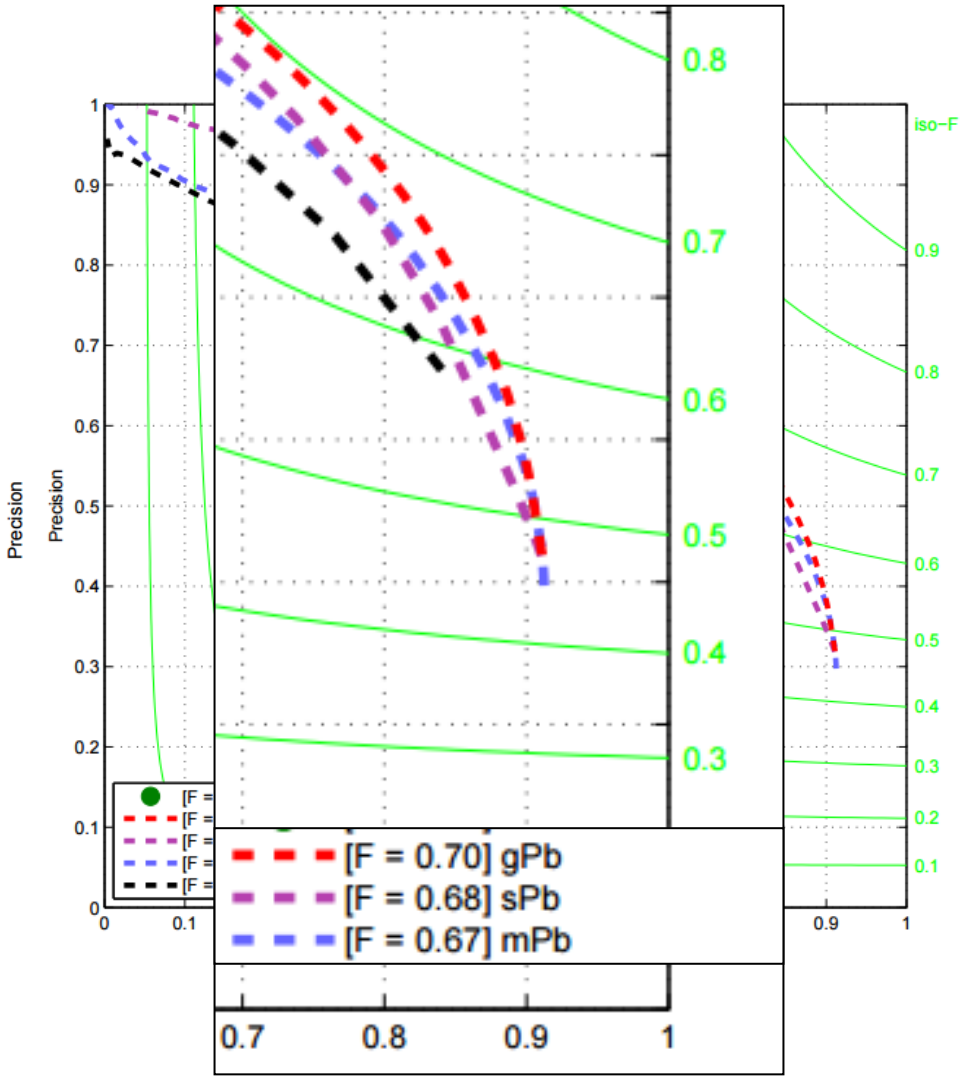

Results – Ablative Analysis

OWT/UCM:

- Ensures closed boundaries
- Helps a little

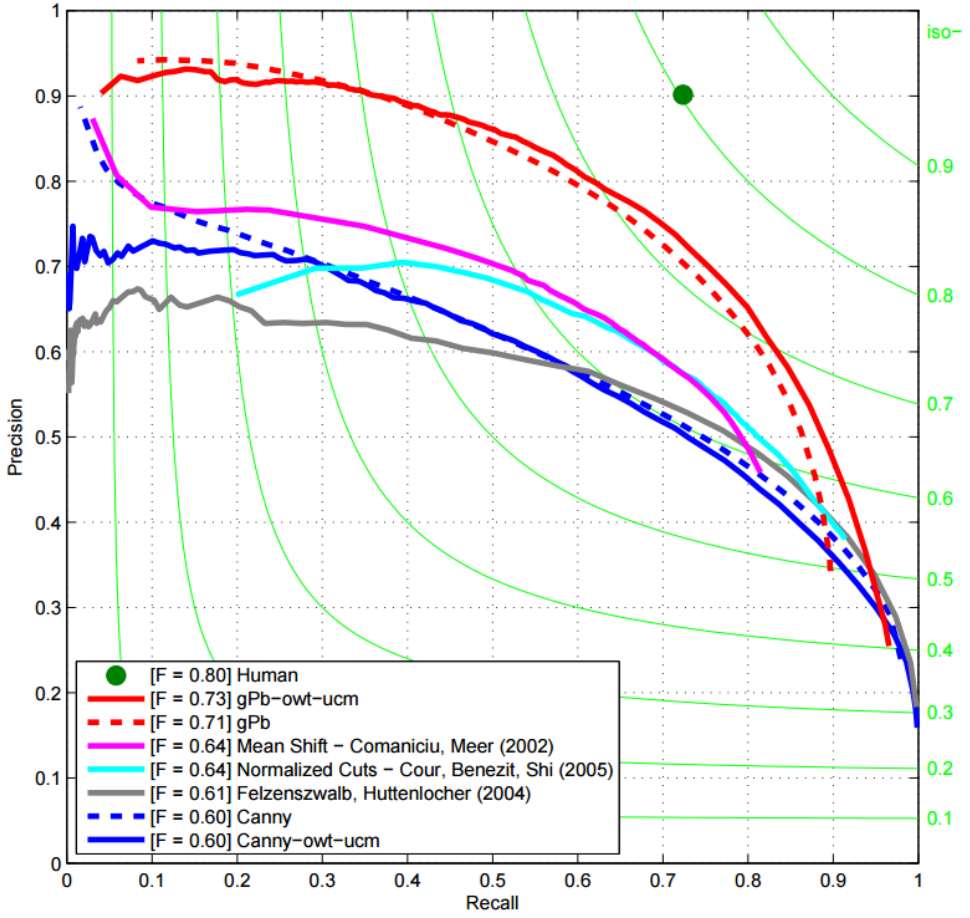

Next Up

- Fast Edge Detection Using Structured Forests
  - Zhihao Li
- Holistically-Nested Edge Detection
  - Yuxin Wu
- Selective Search for Object Recognition
  - Chun-Liang Li

Stash