Image Segmentation

Presented by: Hanbyul Joo, Xinwu Yang

Image Segmentation

Presented by: Hanbyul Joo, Xinwu Yang

Image Segmentation

Objective

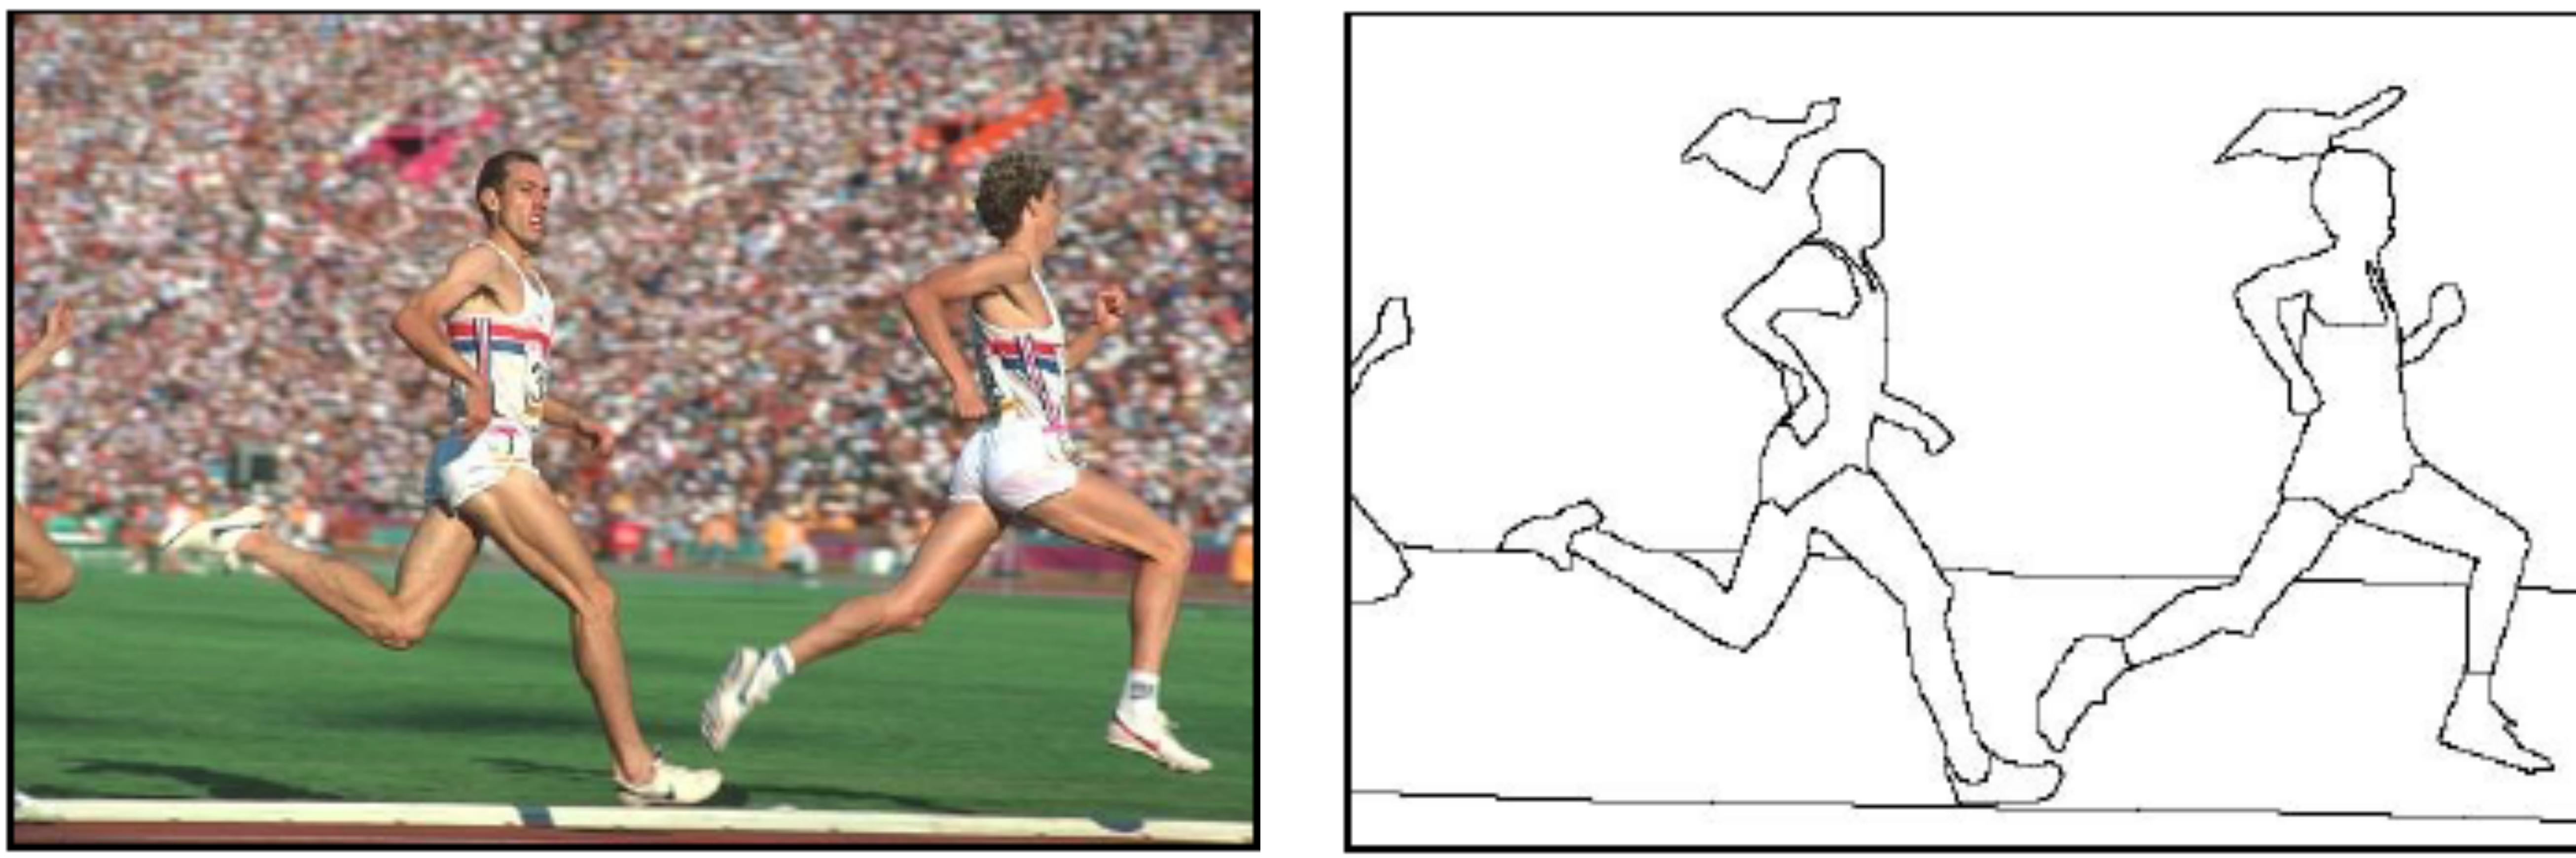

Image Segmentation

Which one is correct?

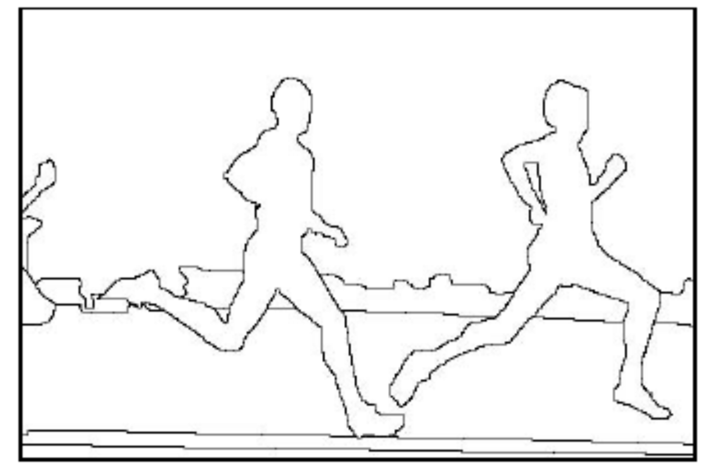

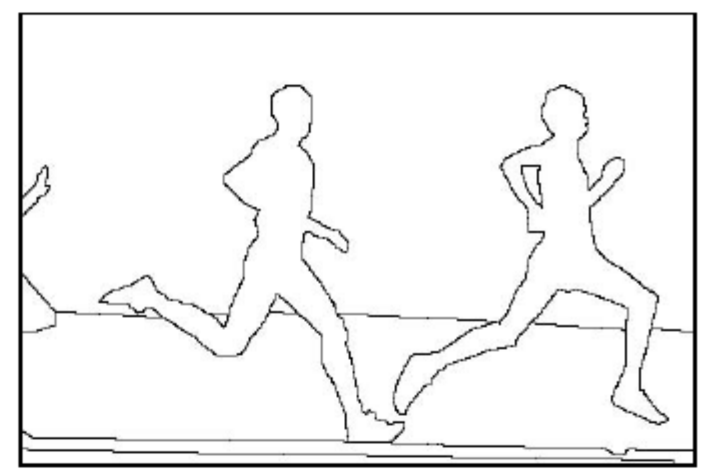

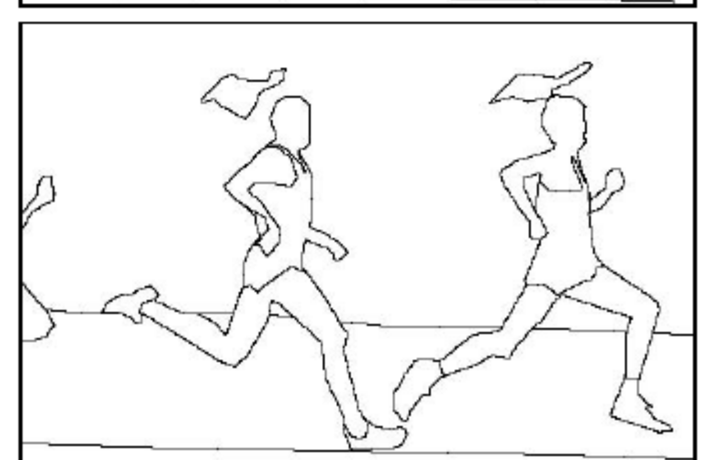

Image Segmentation

Top down cue

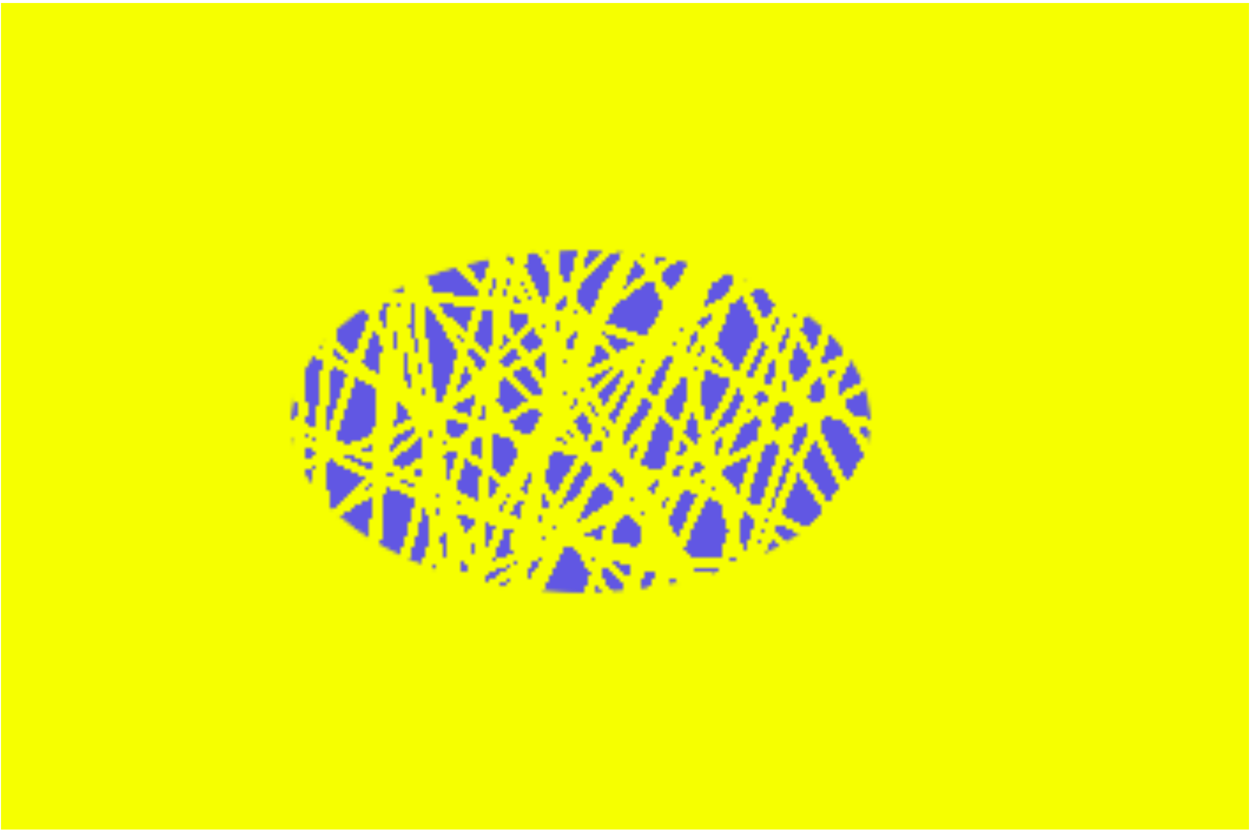

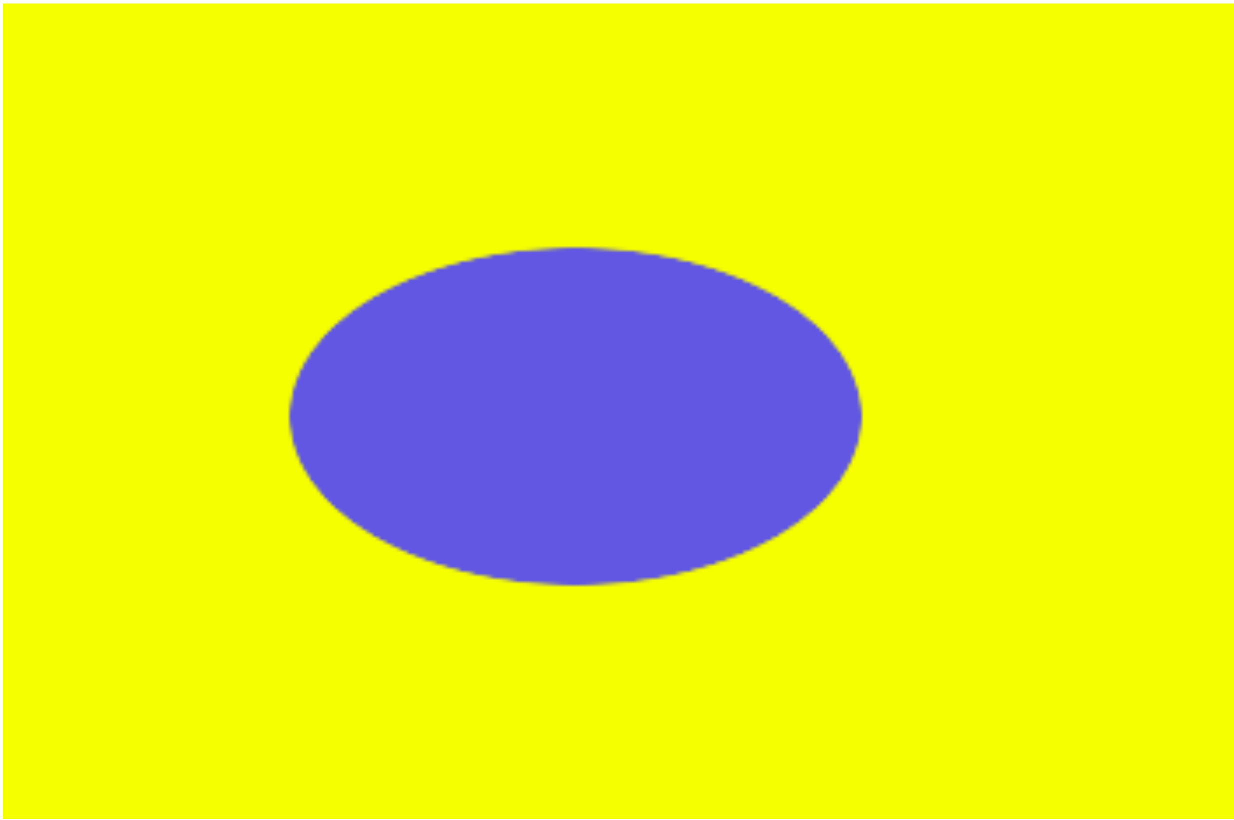

Images by Mike McCann & Jackie Chen

BSDS
(The berkeley Segmentation Dataset)

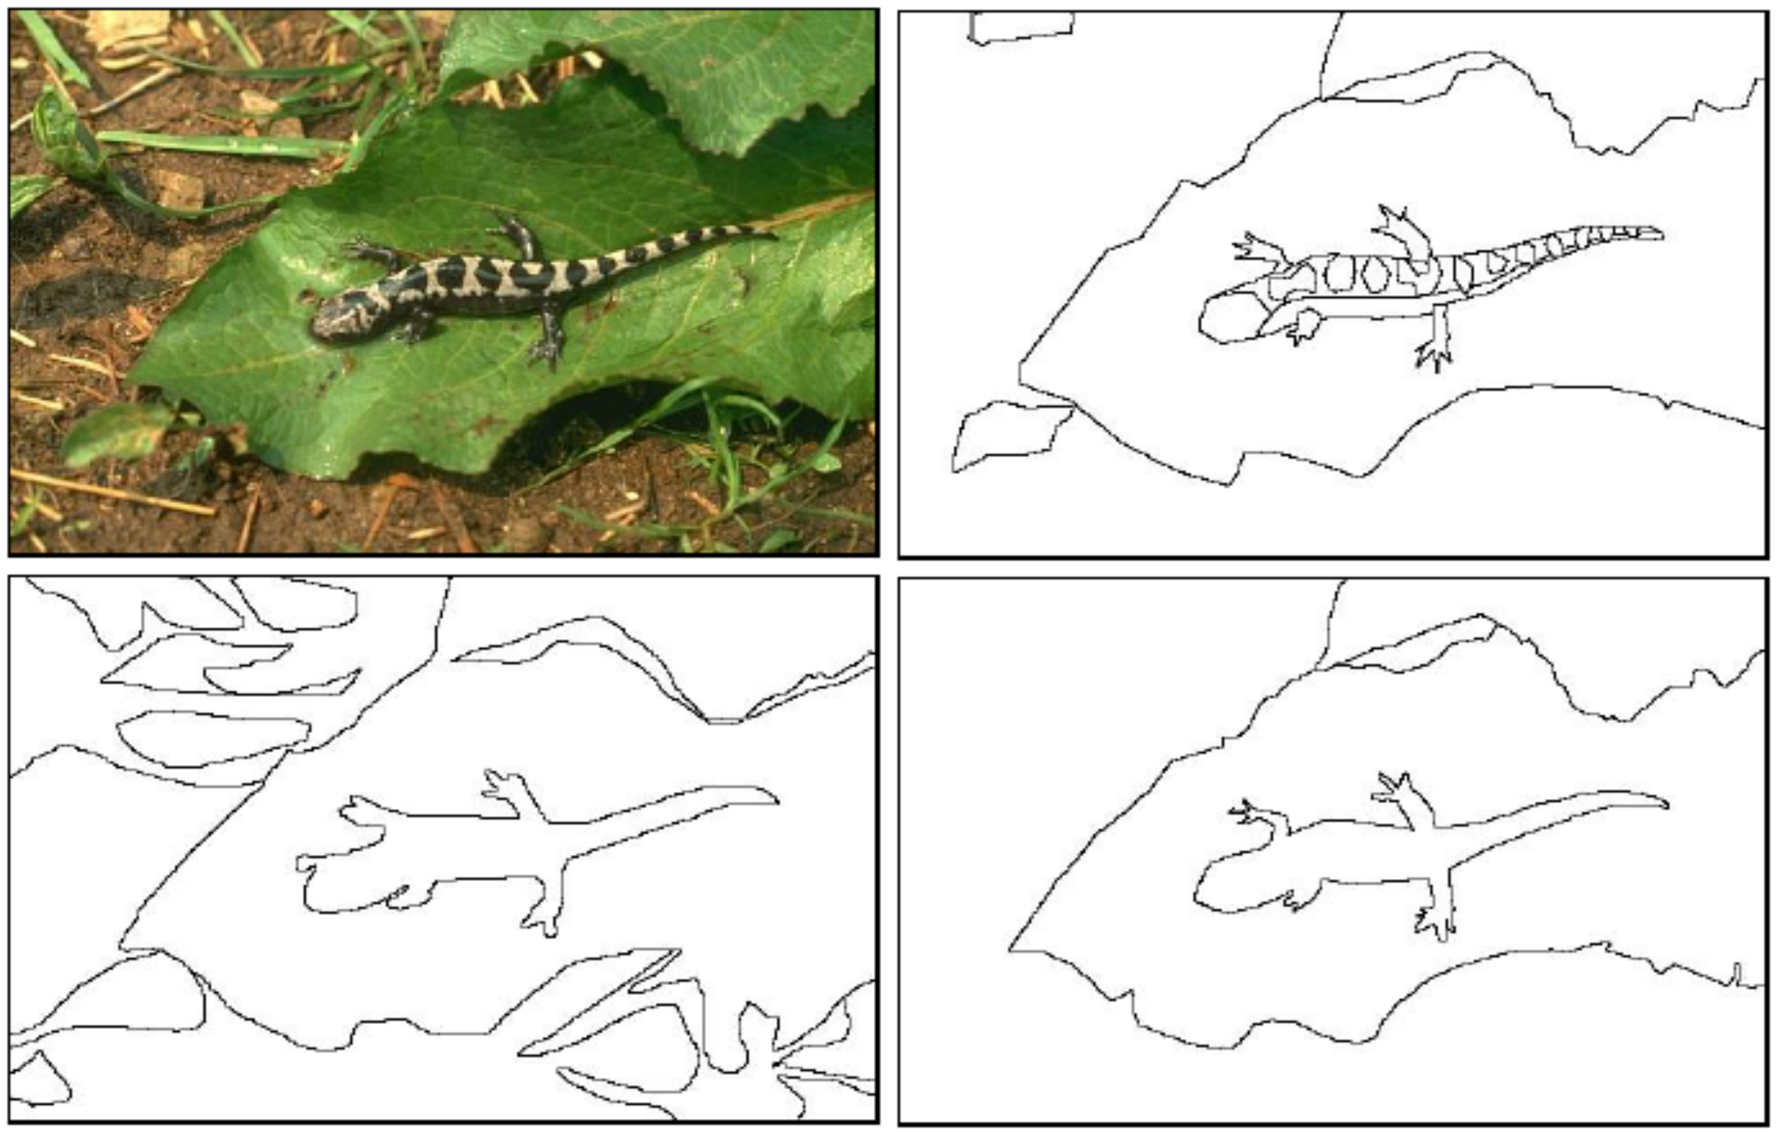

http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/

Weizmann Segmentation Dataset

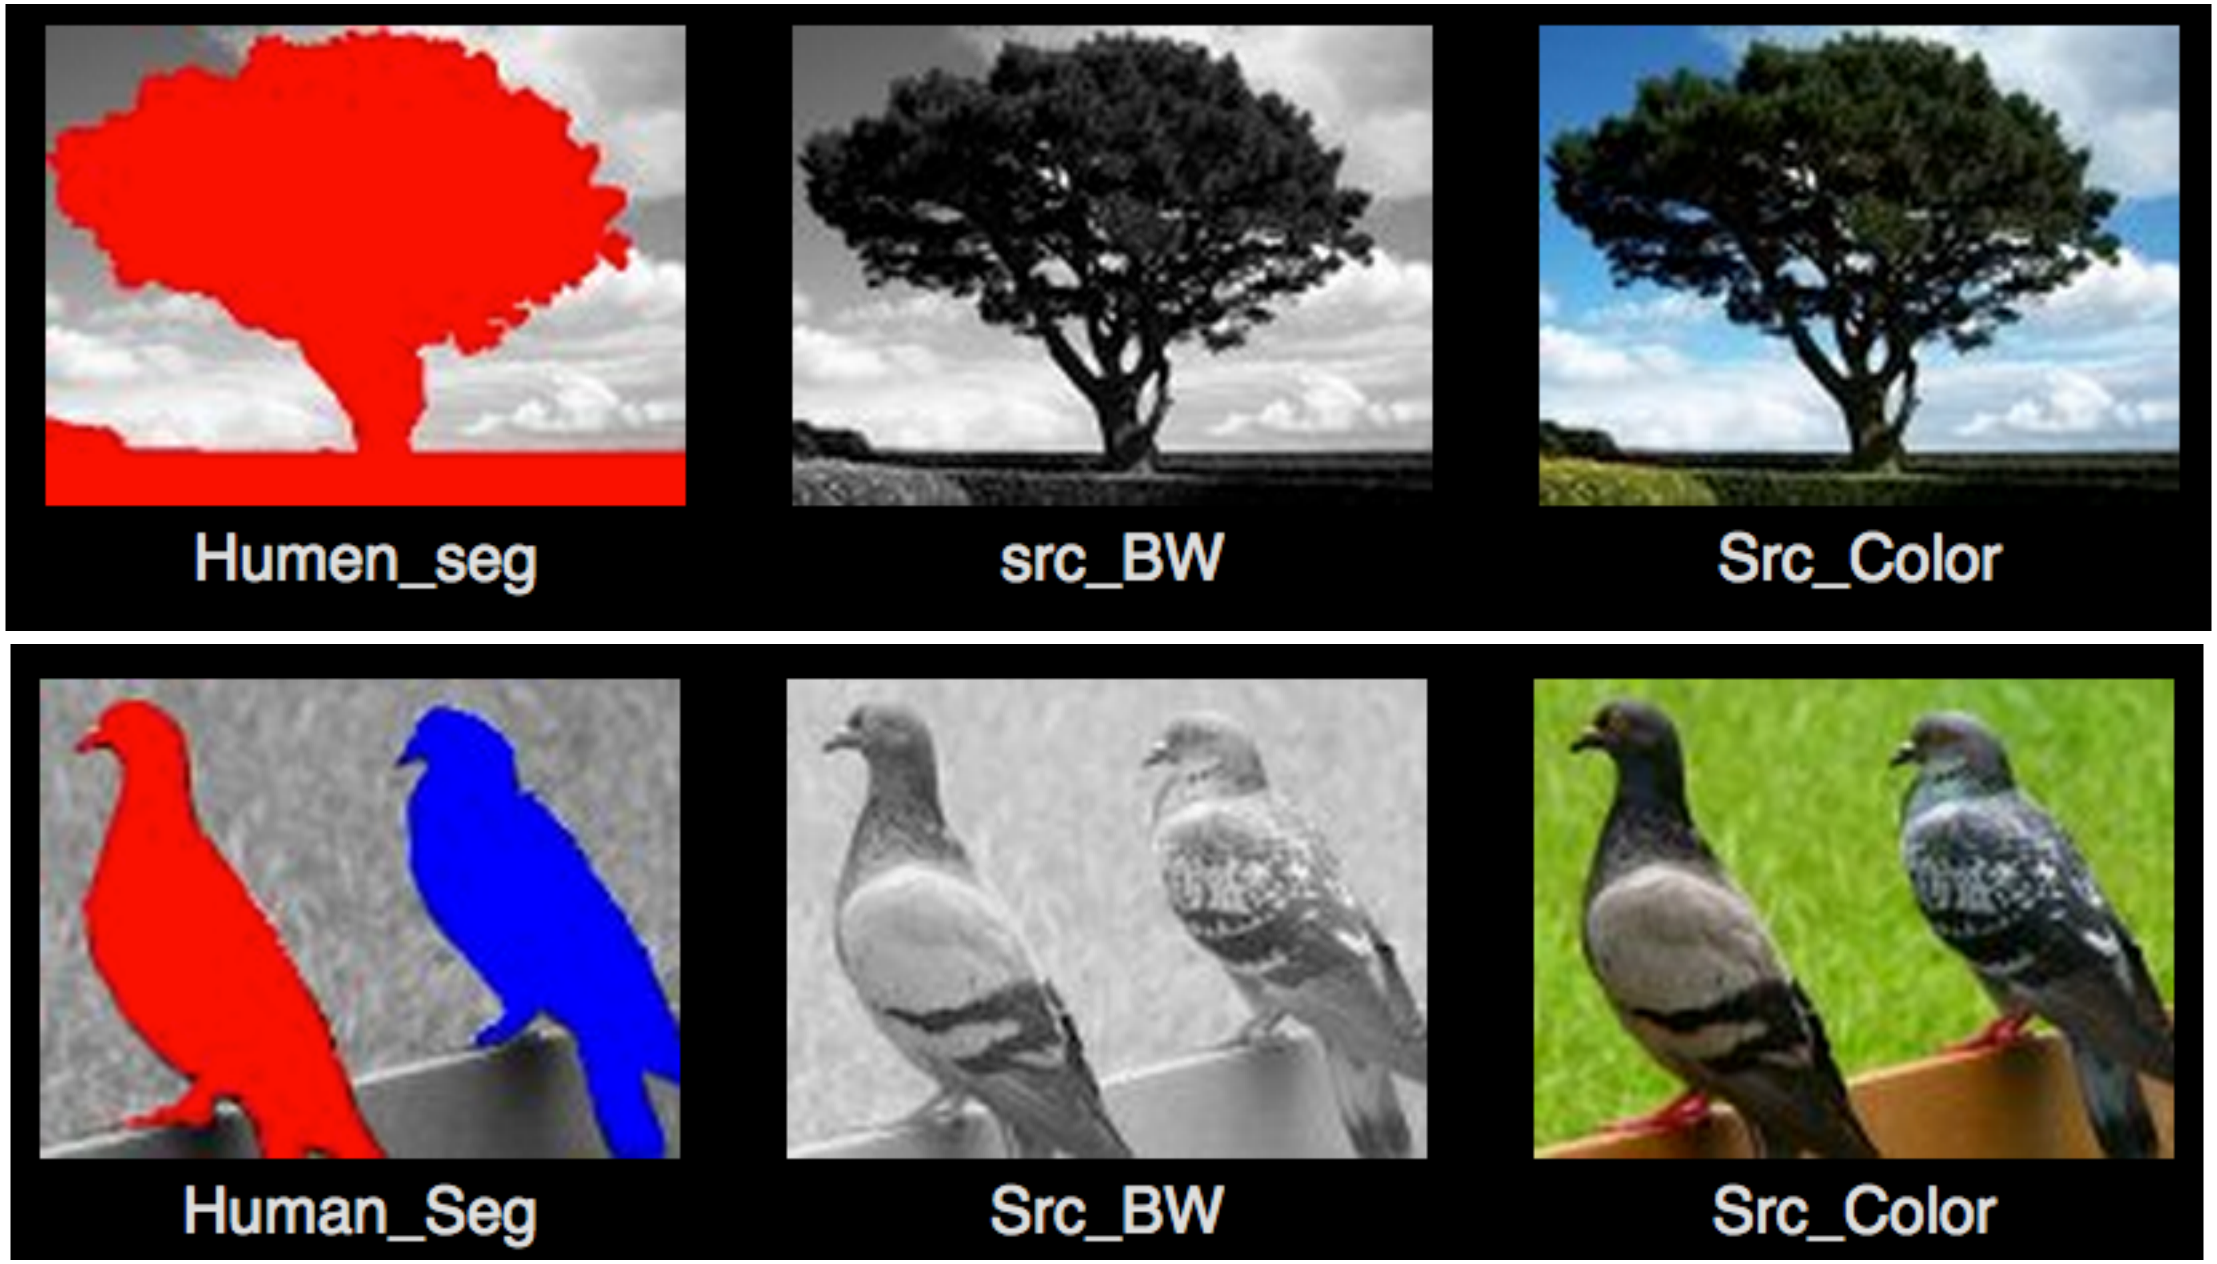

http://www.wisdom.weizmann.ac.il/~vision/Seg_Evaluation_DB/index.html

Pascal (VOC)

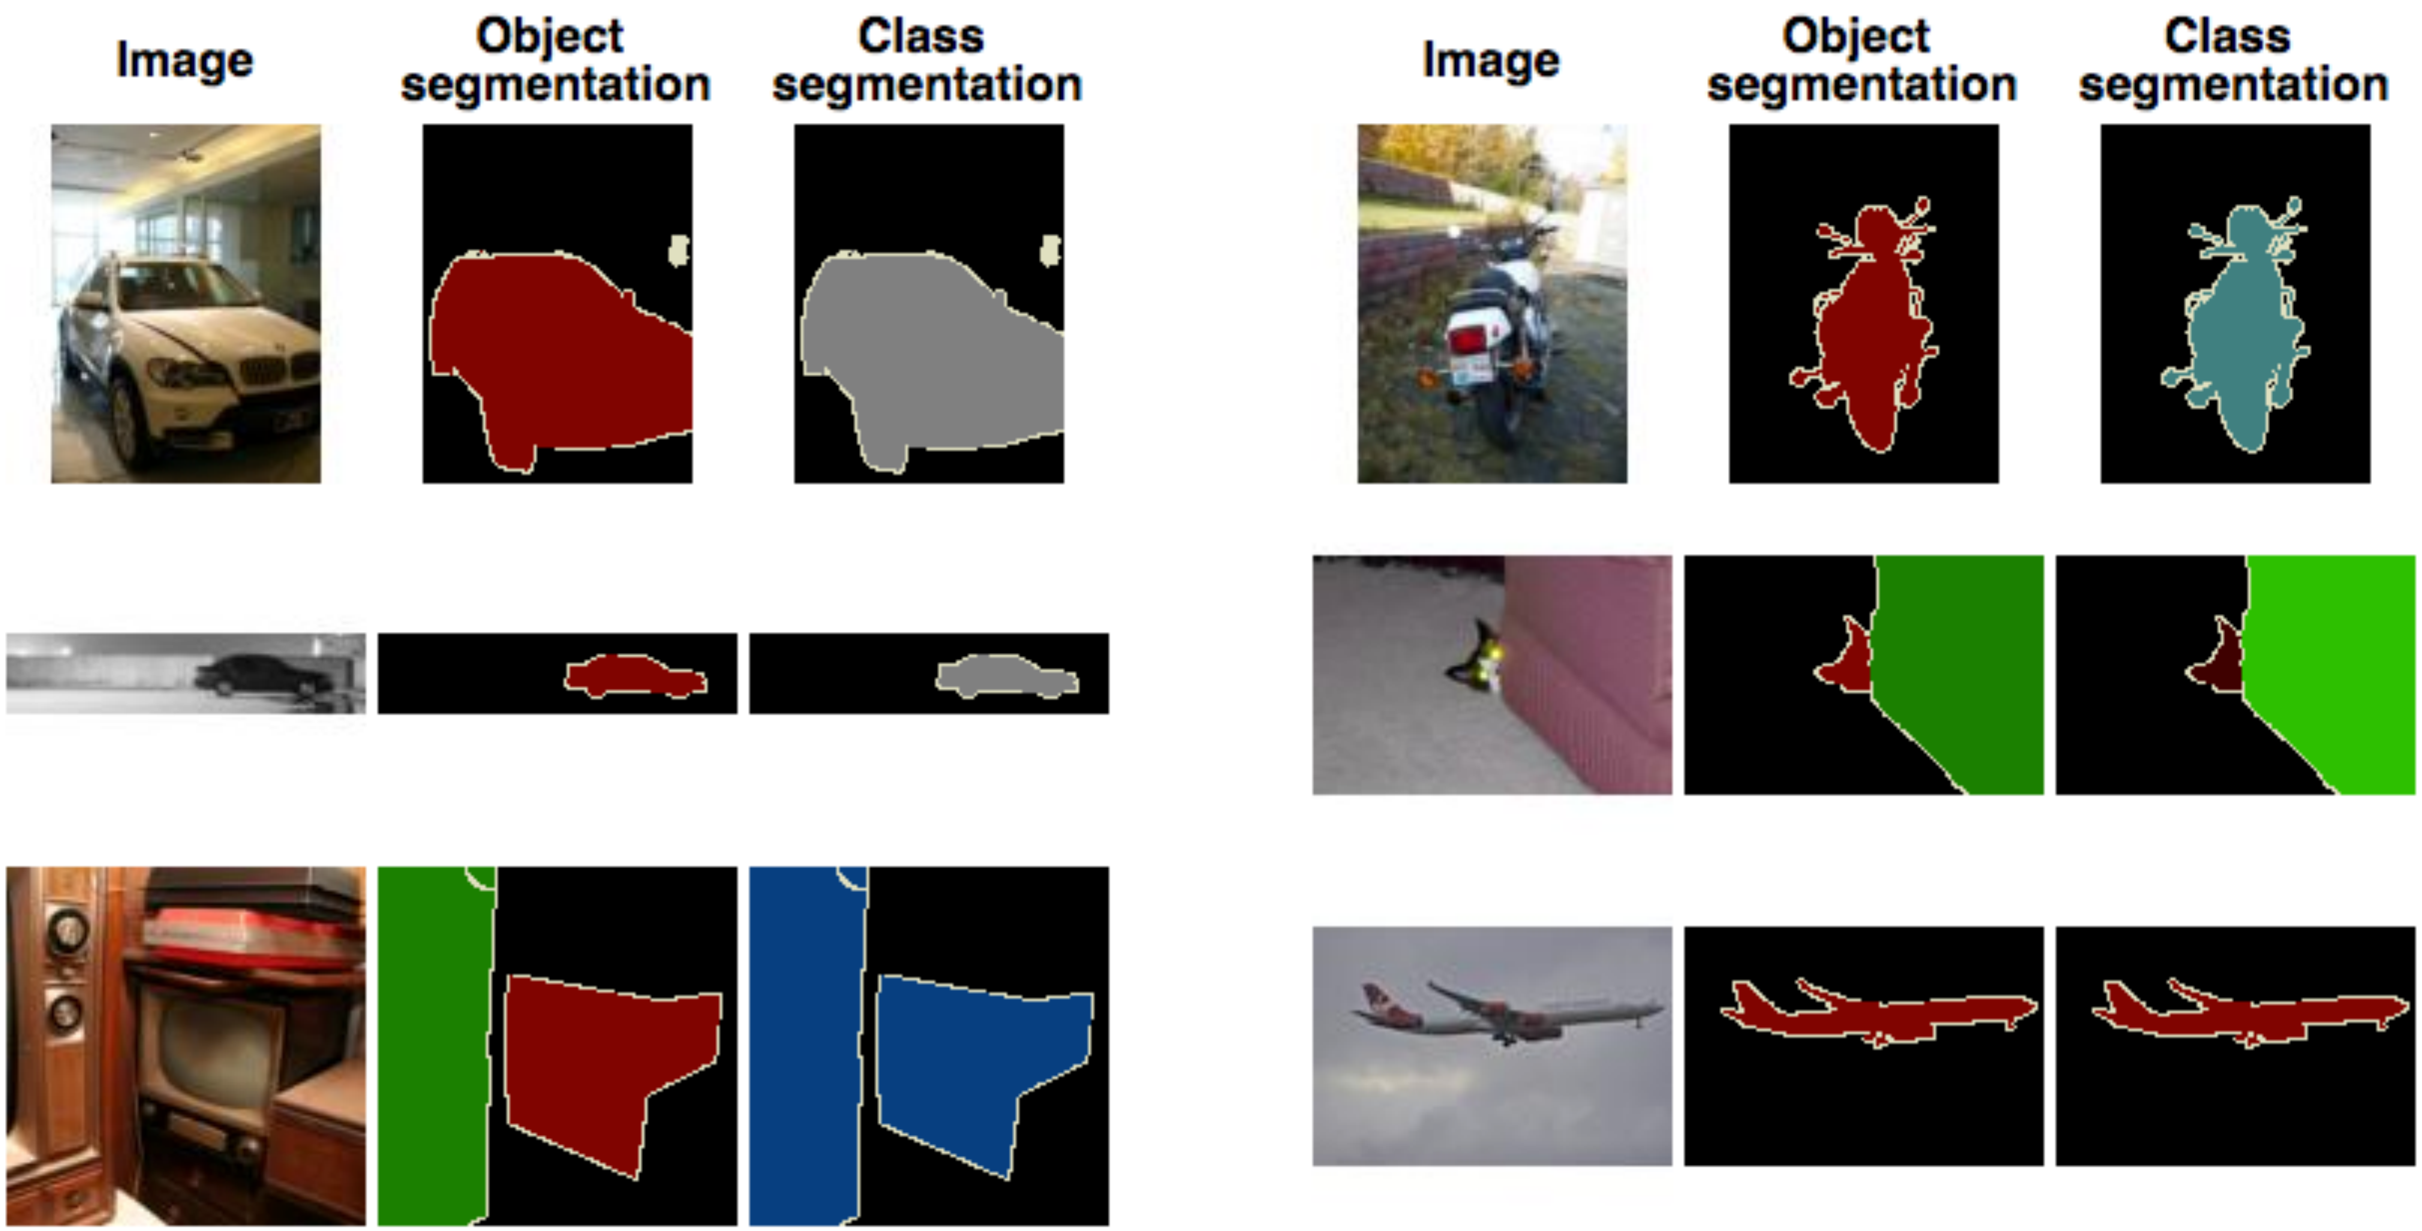

http://pascallin.ecs.soton.ac.uk/challenges/VOC/voc2012/segexamples/index.html

Today

P. Arbelaez, M. Maire, C. Fowlkes and J. Malik. **Contour detection and hierarchical image segmentation**, TPAMI 33, no. 5 (2011): 898-916.

J. Carreira and C. Sminchisescu. **Constrained parametric min-cuts for automatic object segmentation,** CVPR, 2010.

J. Long, E. Shelhamer and T. Darrell. **Fully convolutional networks for semantic segmentation**, arXiv preprint arXiv:1411.4038 (2014).

Today

P. Arbelaez, M. Maire, C. Fowlkes and J. Malik. **Contour detection and hierarchical image segmentation**, TPAMI 33, no. 5 (2011): 898-916.

J. Carreira and C. Sminchisescu. **Constrained parametric min-cuts for automatic object segmentation,** CVPR, 2010.

J. Long, E. Shelhamer and T. Darrell. **Fully convolutional networks for semantic segmentation**, arXiv preprint arXiv:1411.4038 (2014).

Introduction

Contour Detection

Introduction

Image Segmentation

Obvious: Segmentation -> Contour
Not obvious: Contour -> Segmentations

Introduction
Main Contributions

- Contour detection:
  Multiple local cues (multi-scale)
  + globalization framework (spectral clustering)

- Segmentation:
  Transforming contour results to segmentation

- Through evaluation on BSDS

Contour Detection [2]
Oriented Gradient Signal

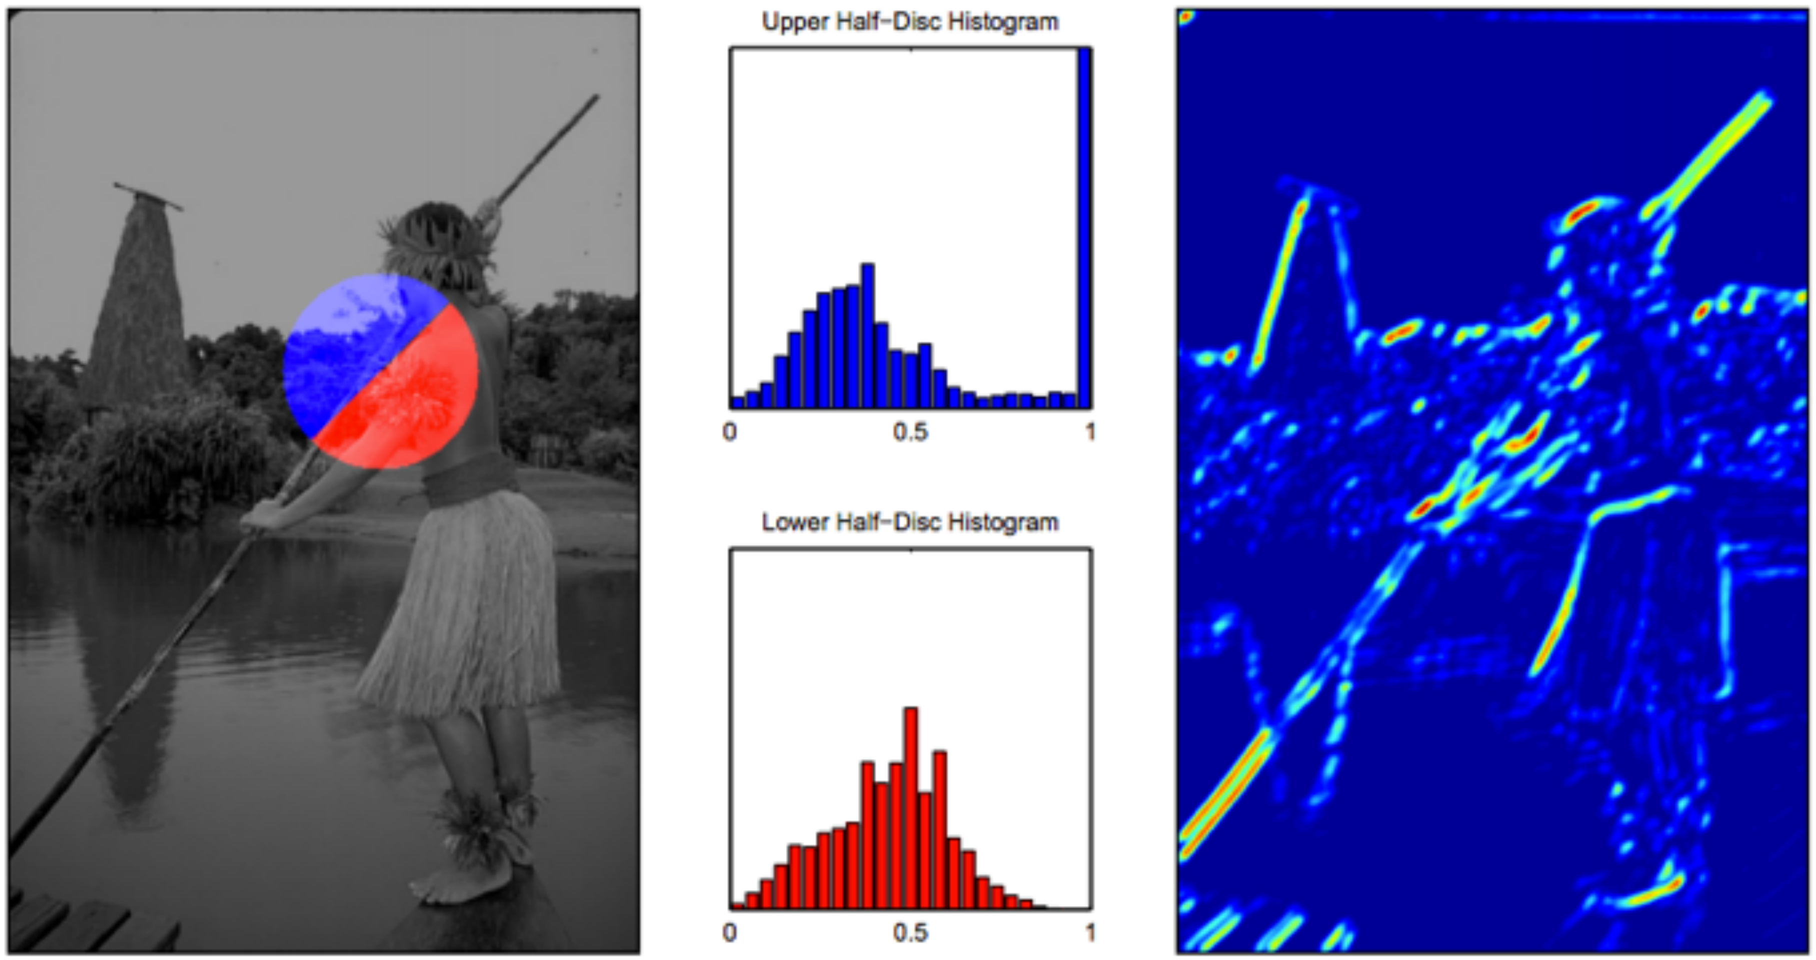

- Brightness
- Color
- Texture

$$G(x, y, \theta) = \chi^2(g, h) = \frac{1}{2} \sum_i \frac{(g(i) - h(i))^2}{g(i) + h(i)}$$

8 orientations

Contour Detection [2]

Four Separate Features

Brightness

Color

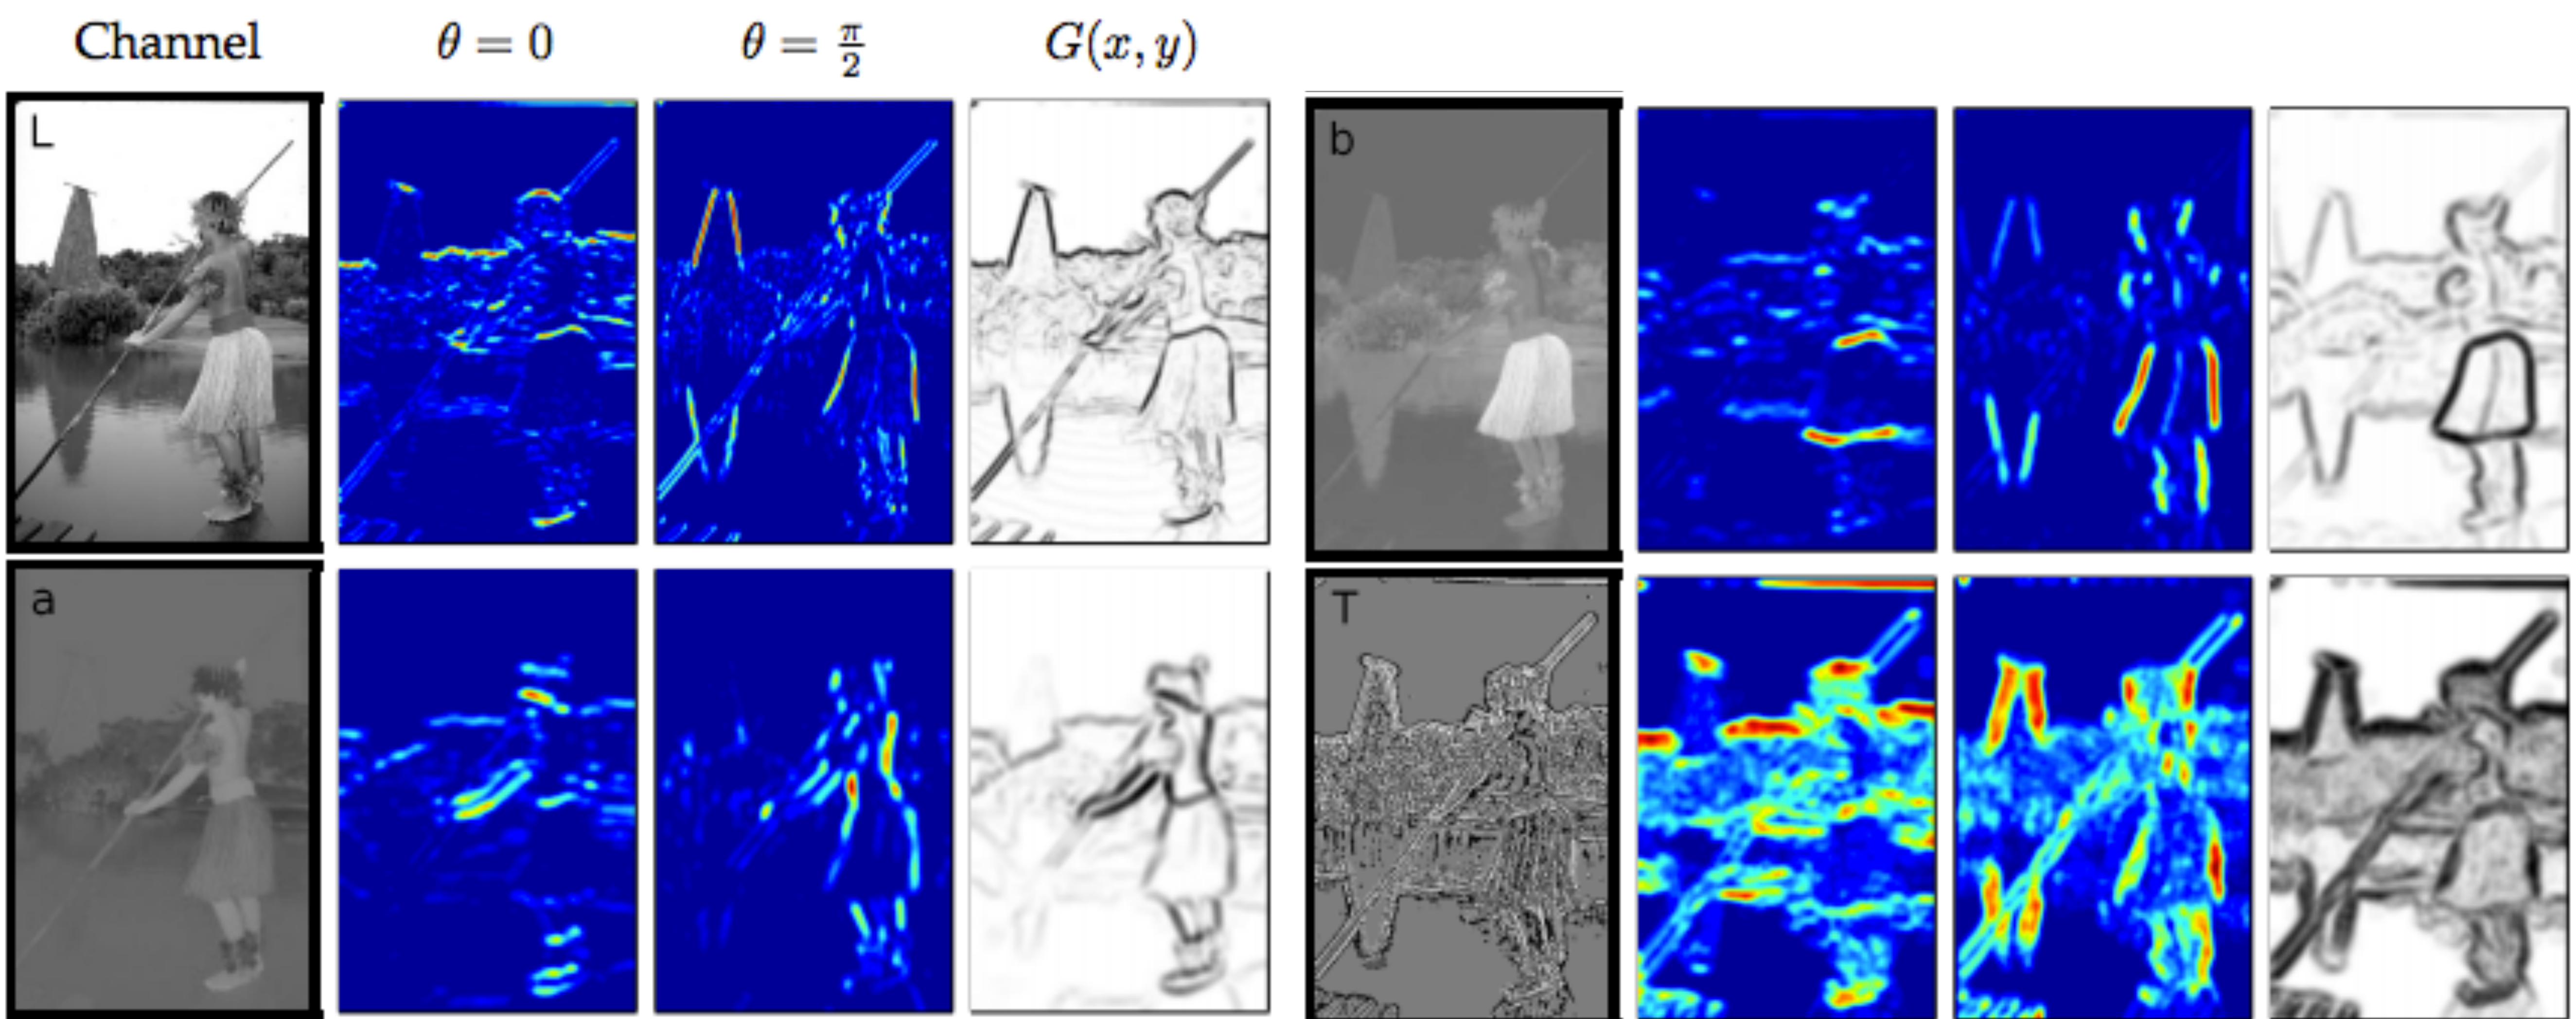

Color

Maximum response among all 8 orientation

Texture (texton id)

Contour Detection

Multiscale Cue Combination (Proposed)

$[\frac{\sigma}{2}, \sigma, 2\sigma]$

Three scales

From training data

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

$$mPb(x,y) = \max_{\theta}\{mPb(x,y,\theta)\}$$

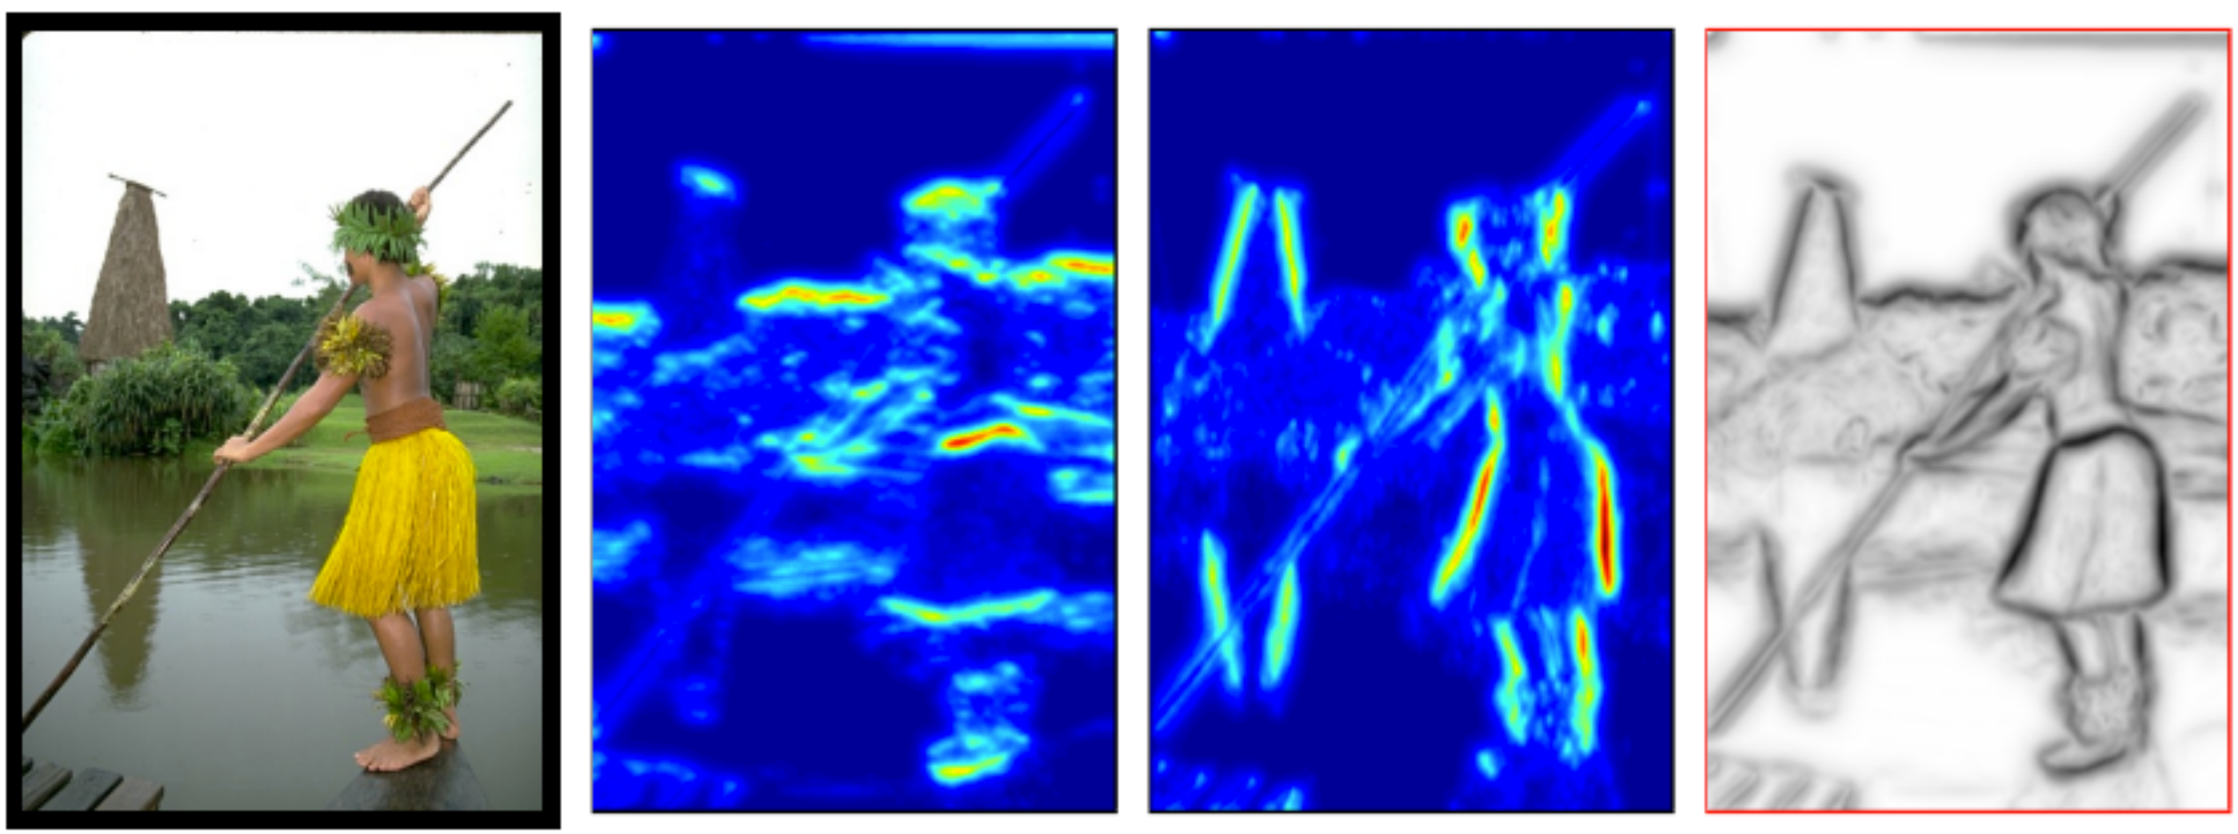

$\theta = 0$ $\theta = \frac{\pi}{2}$ $mPb(x,y)$

Contour Detection

Global Cues (Affinity)

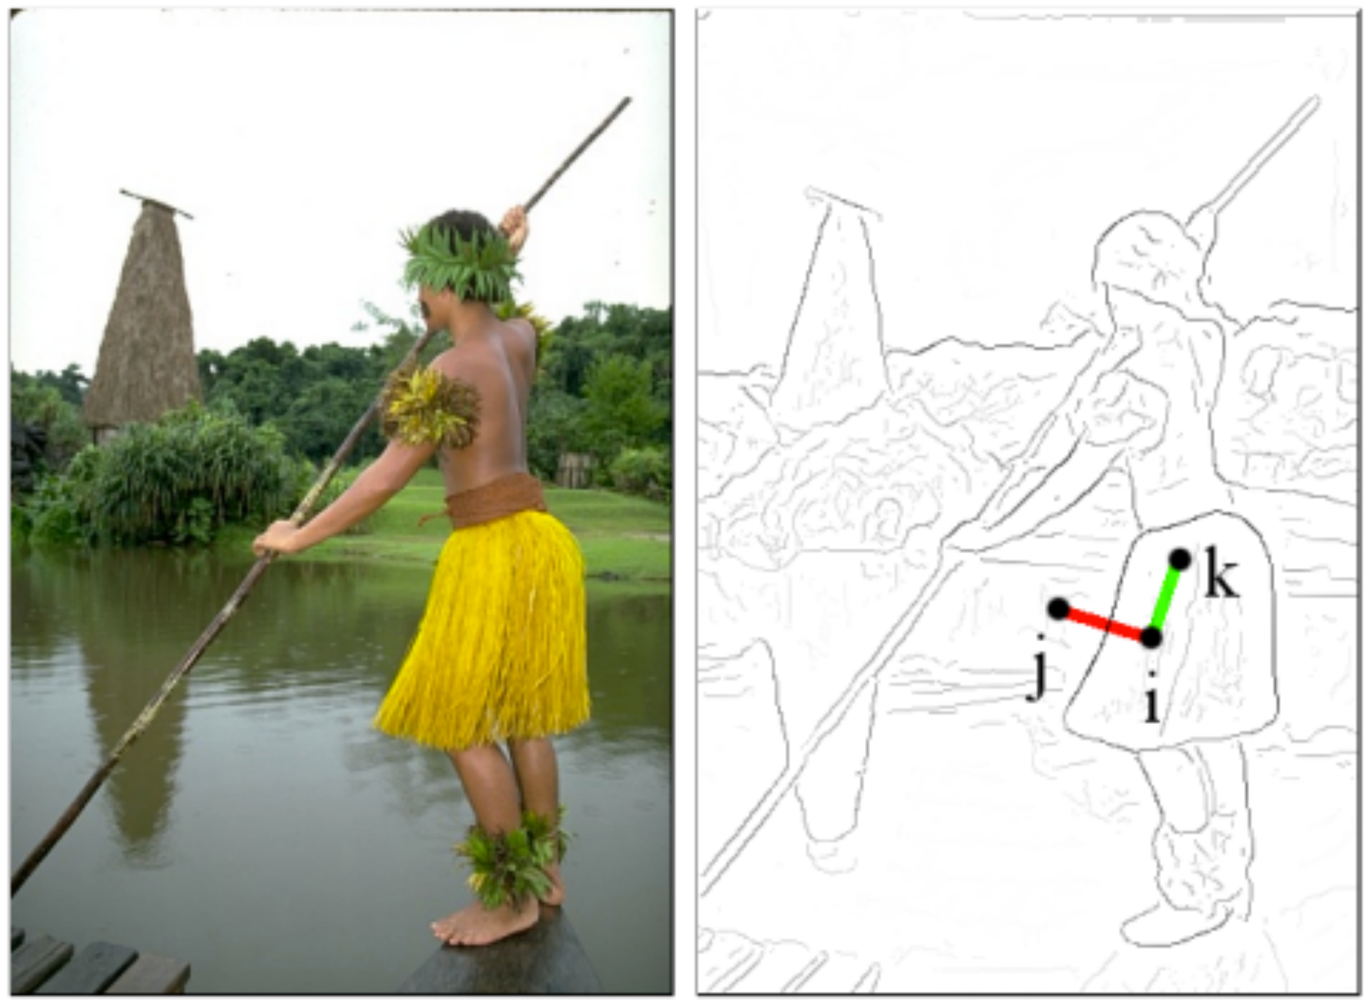

$$W_{ij} = \exp\left(-\max_{p \in \overline{ij}}\{mPb(p)\}/\rho\right)$$

Contour Detection

Global Cues (Problem)

Typical problem of NCut

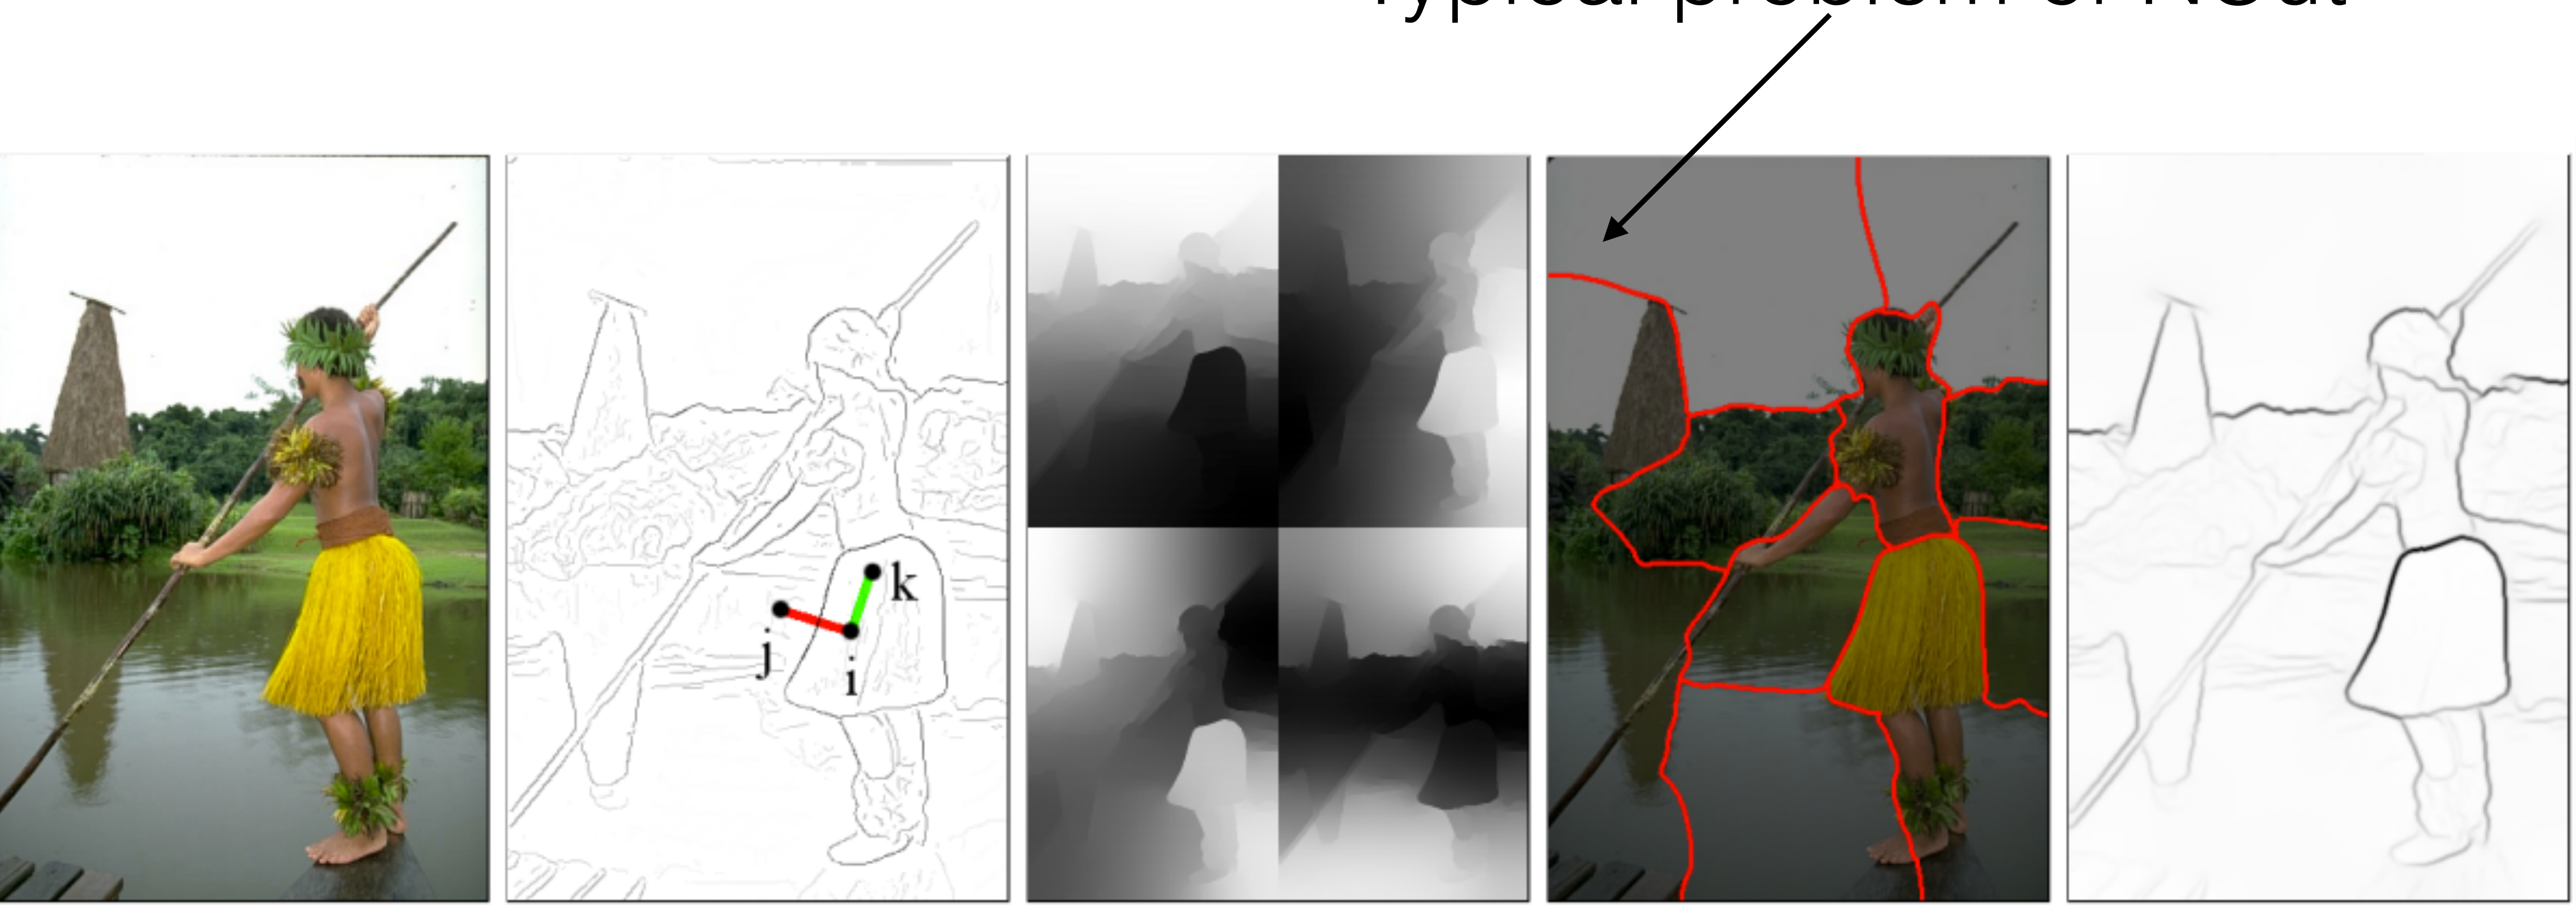

Contour Detection

Global Cues ("Spectral Component")

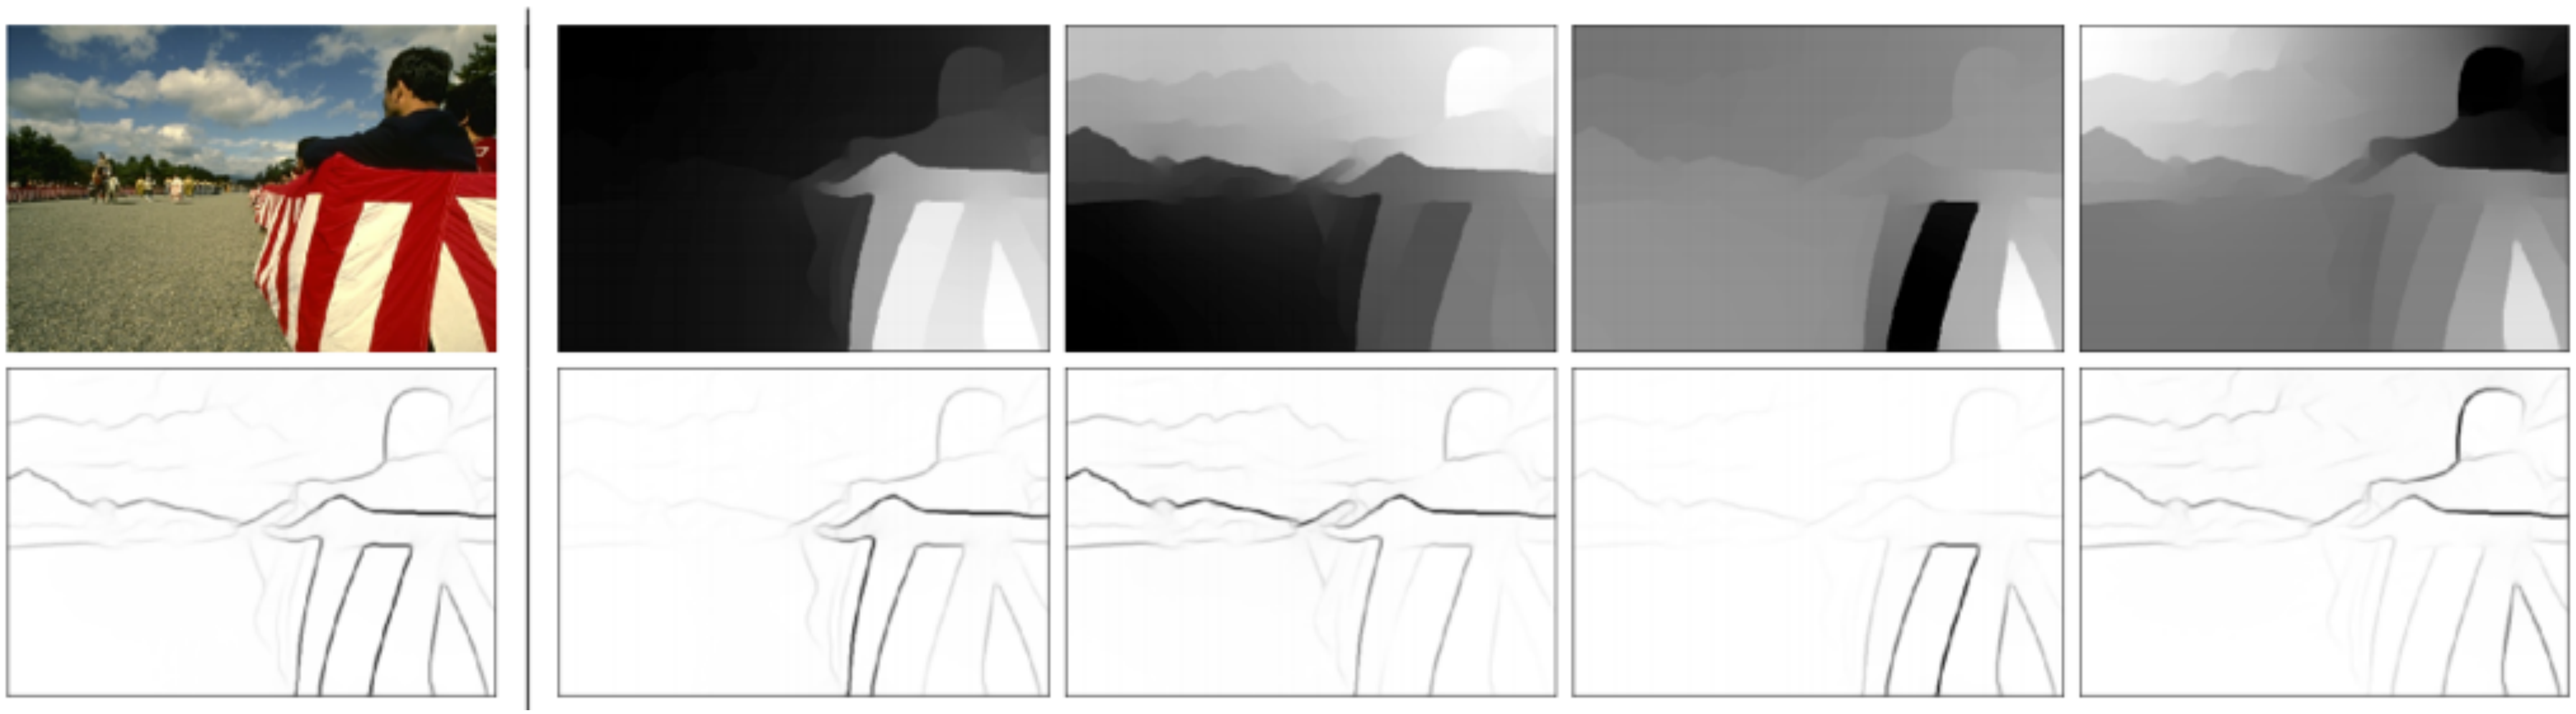

Maximum sPb over orientations

$$sPb(x, y, \theta) = \sum_{k=1}^{n} \frac{1}{\sqrt{\lambda_k}} \cdot \nabla_\theta \mathbf{v_k}(x, y)$$

Gaussian directional derivative

Contour Detection

Final Probability

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

All the edges

Salient edges

Contour Detection

Results

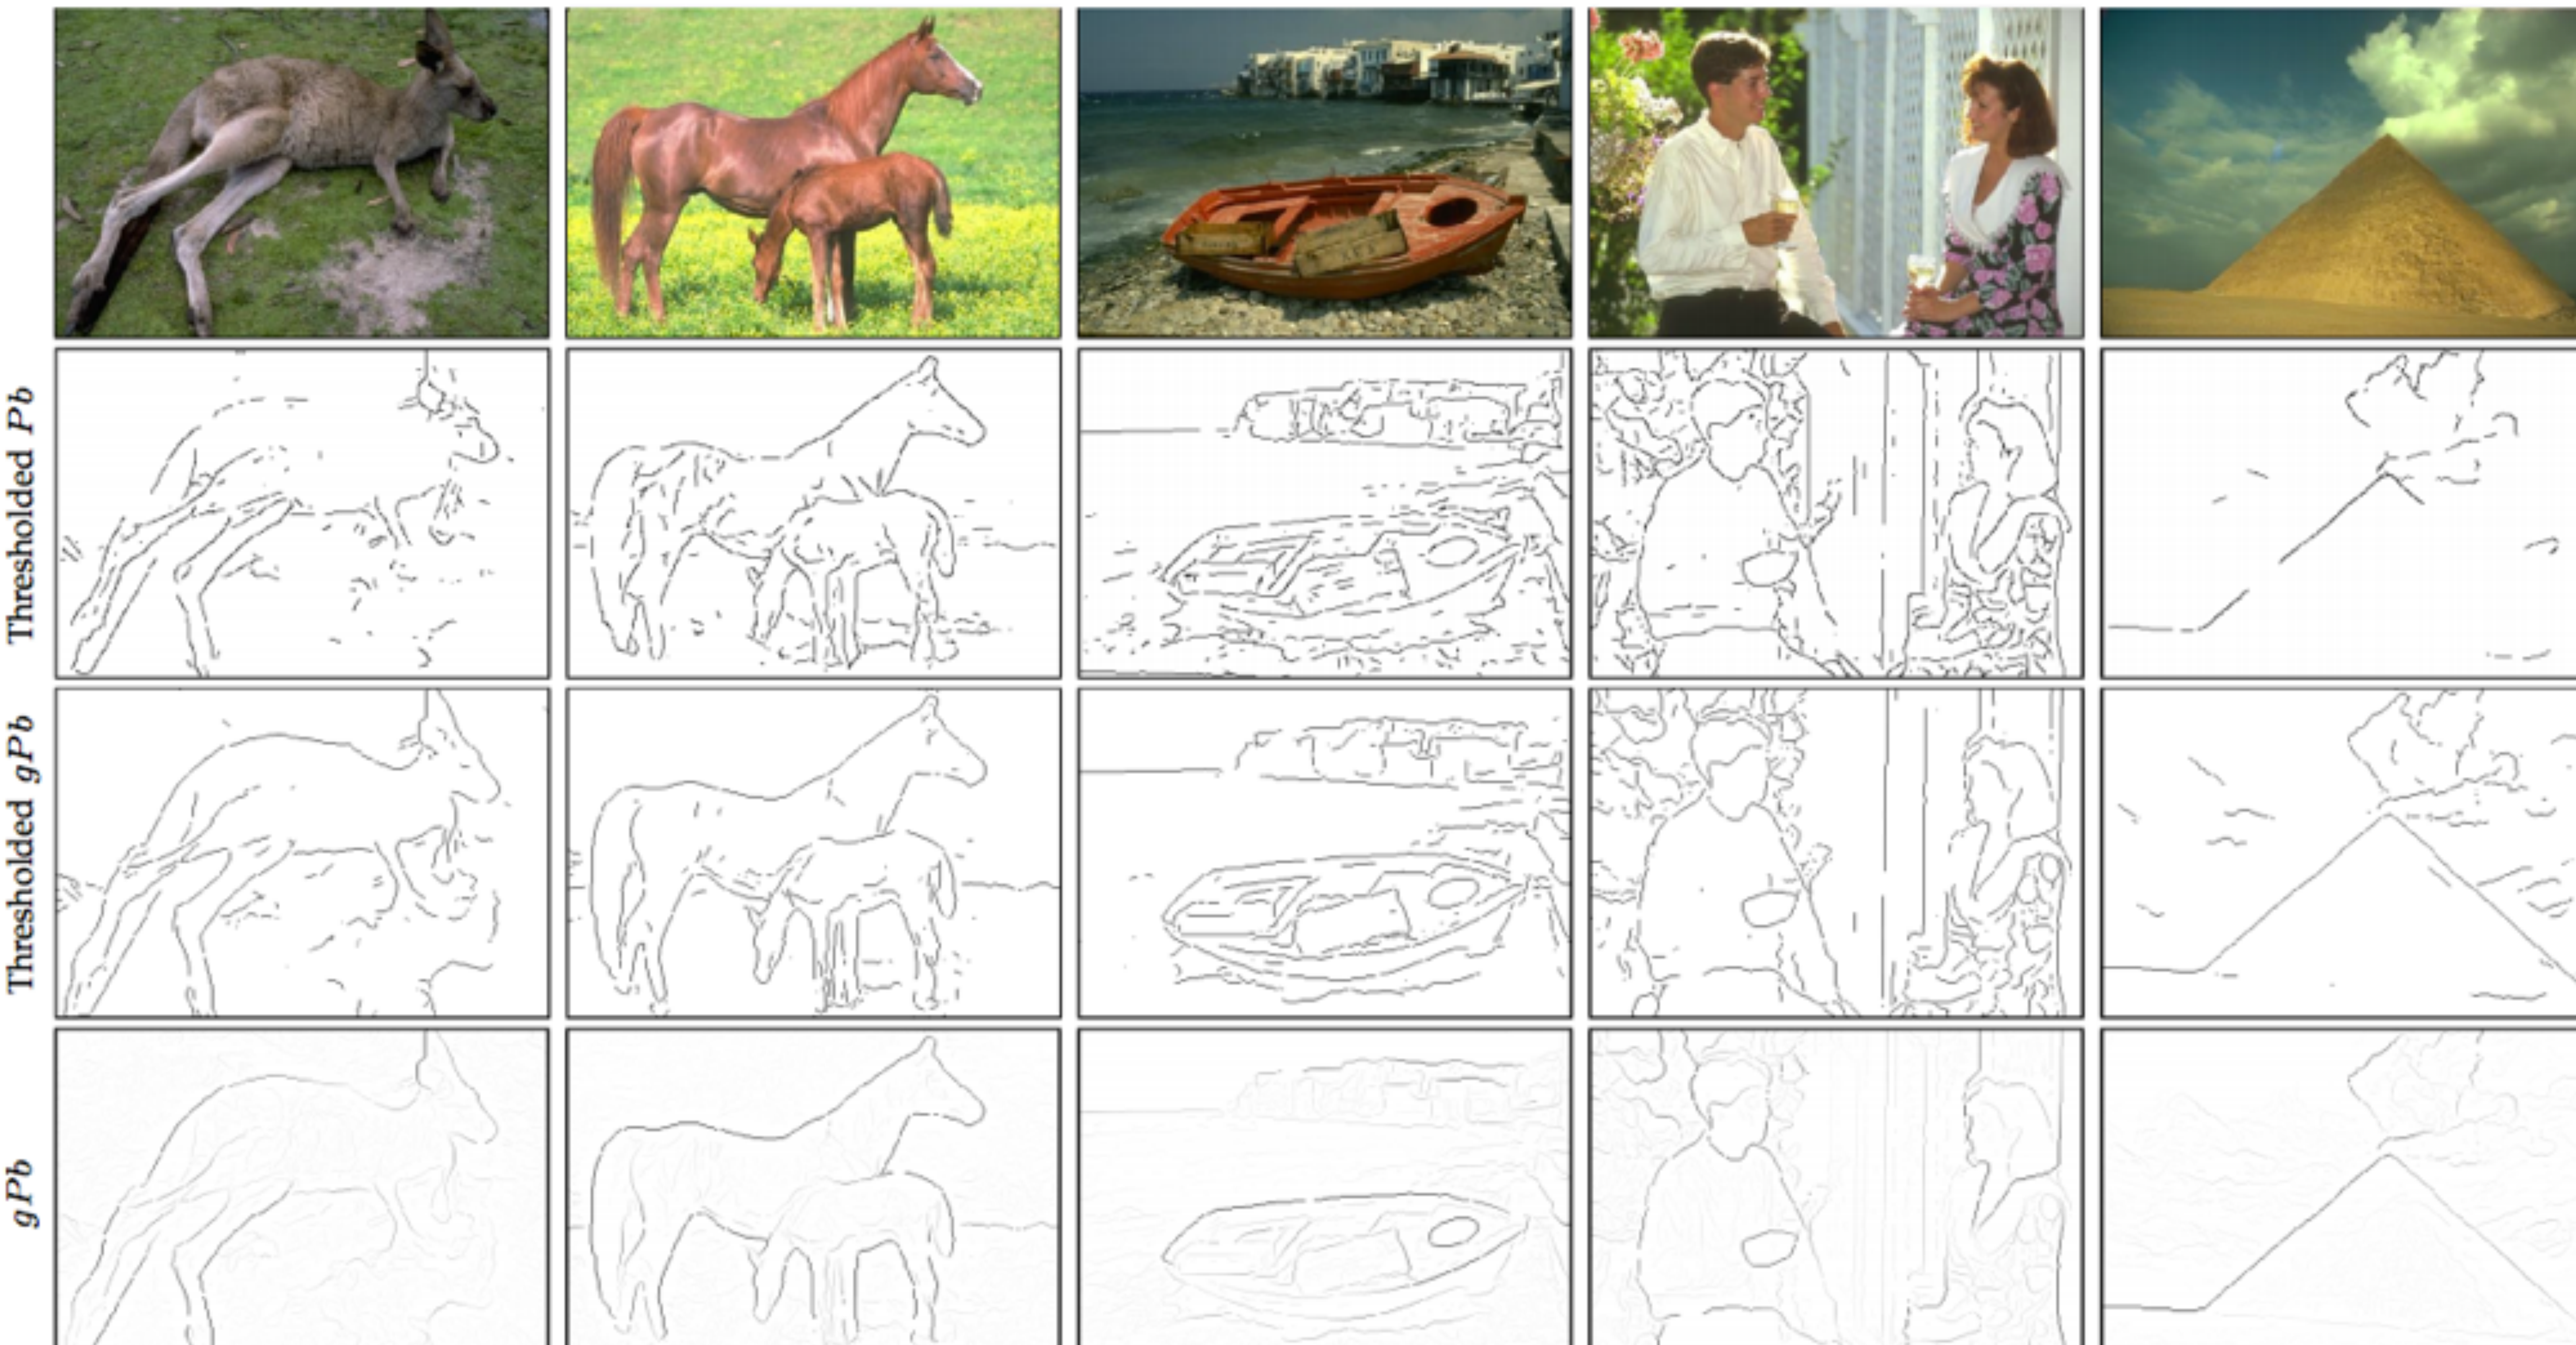

Contour Detection

Results (benefit of globalization)

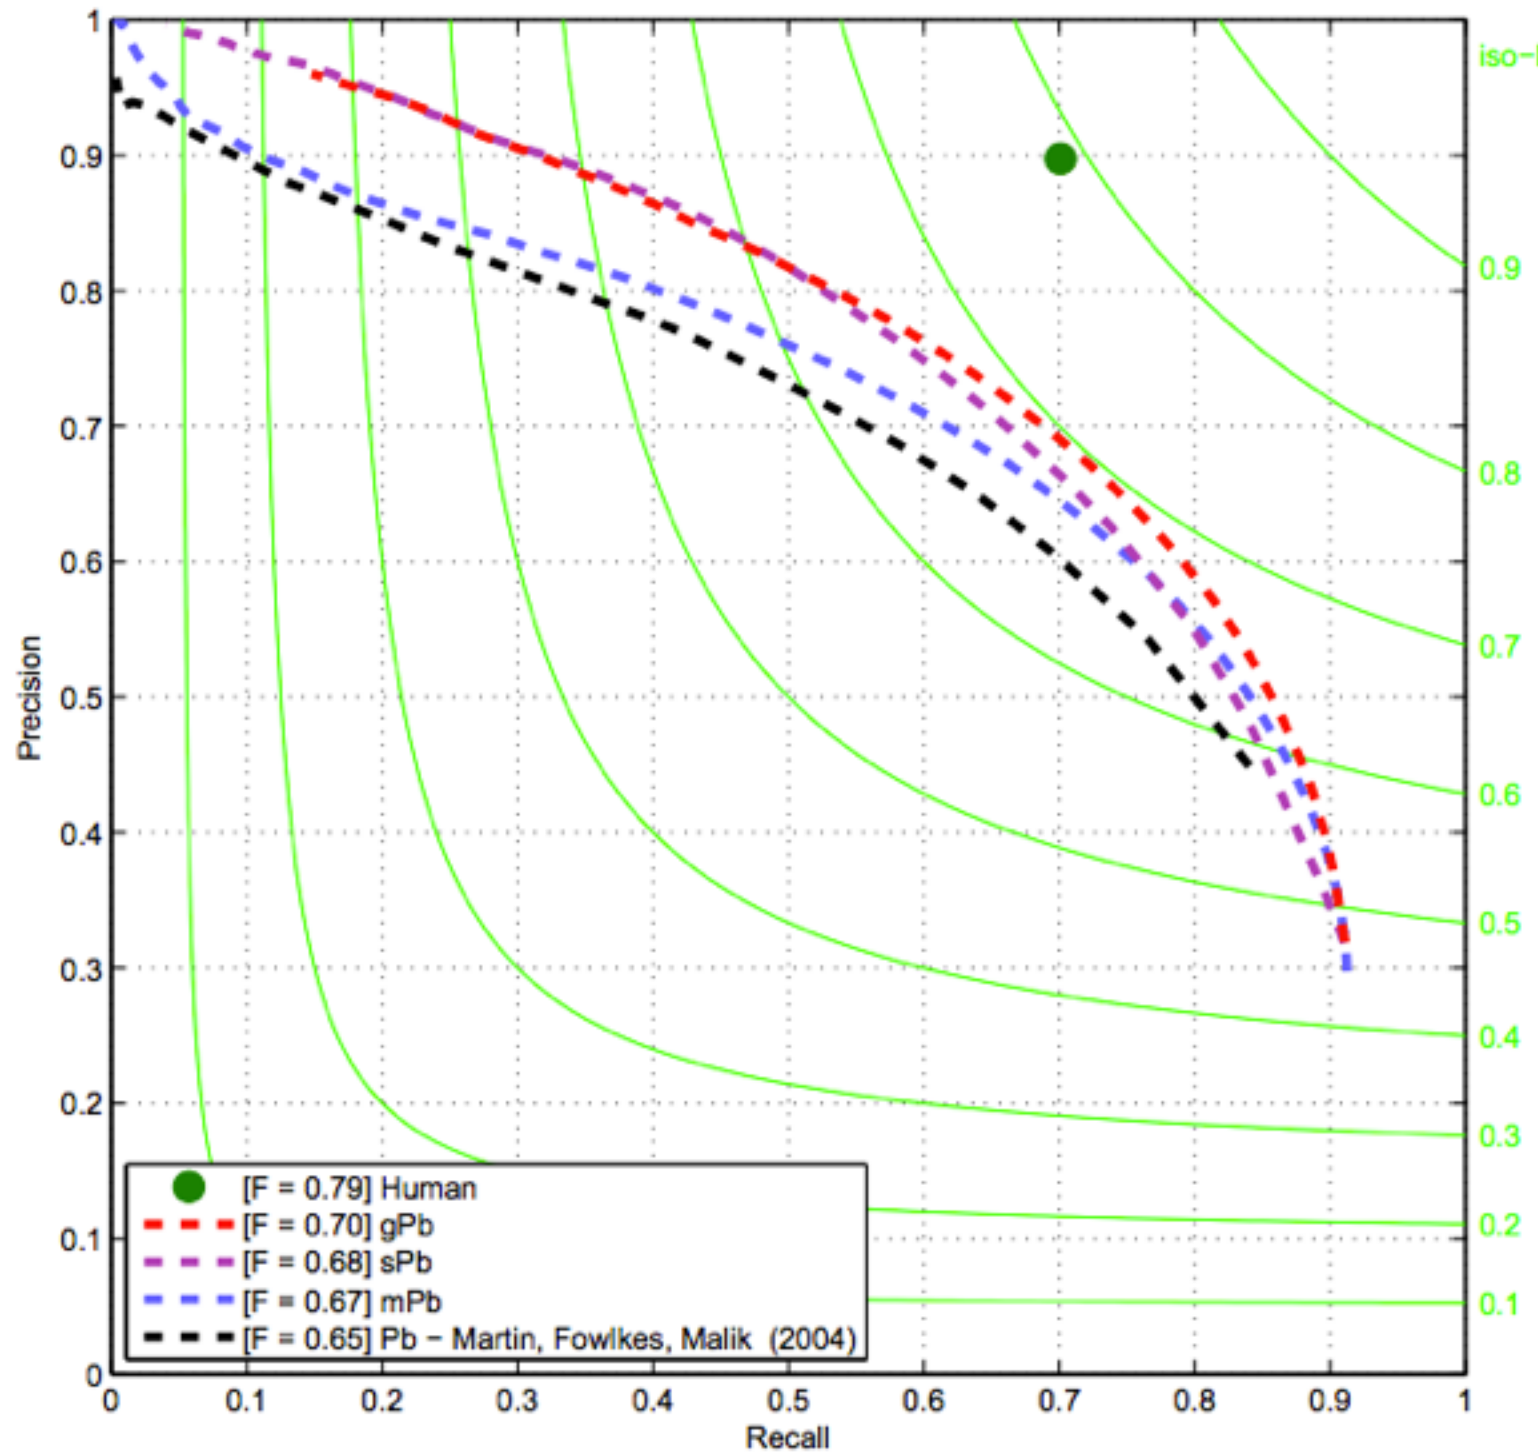

Segmentation

Transform contour to closed contour

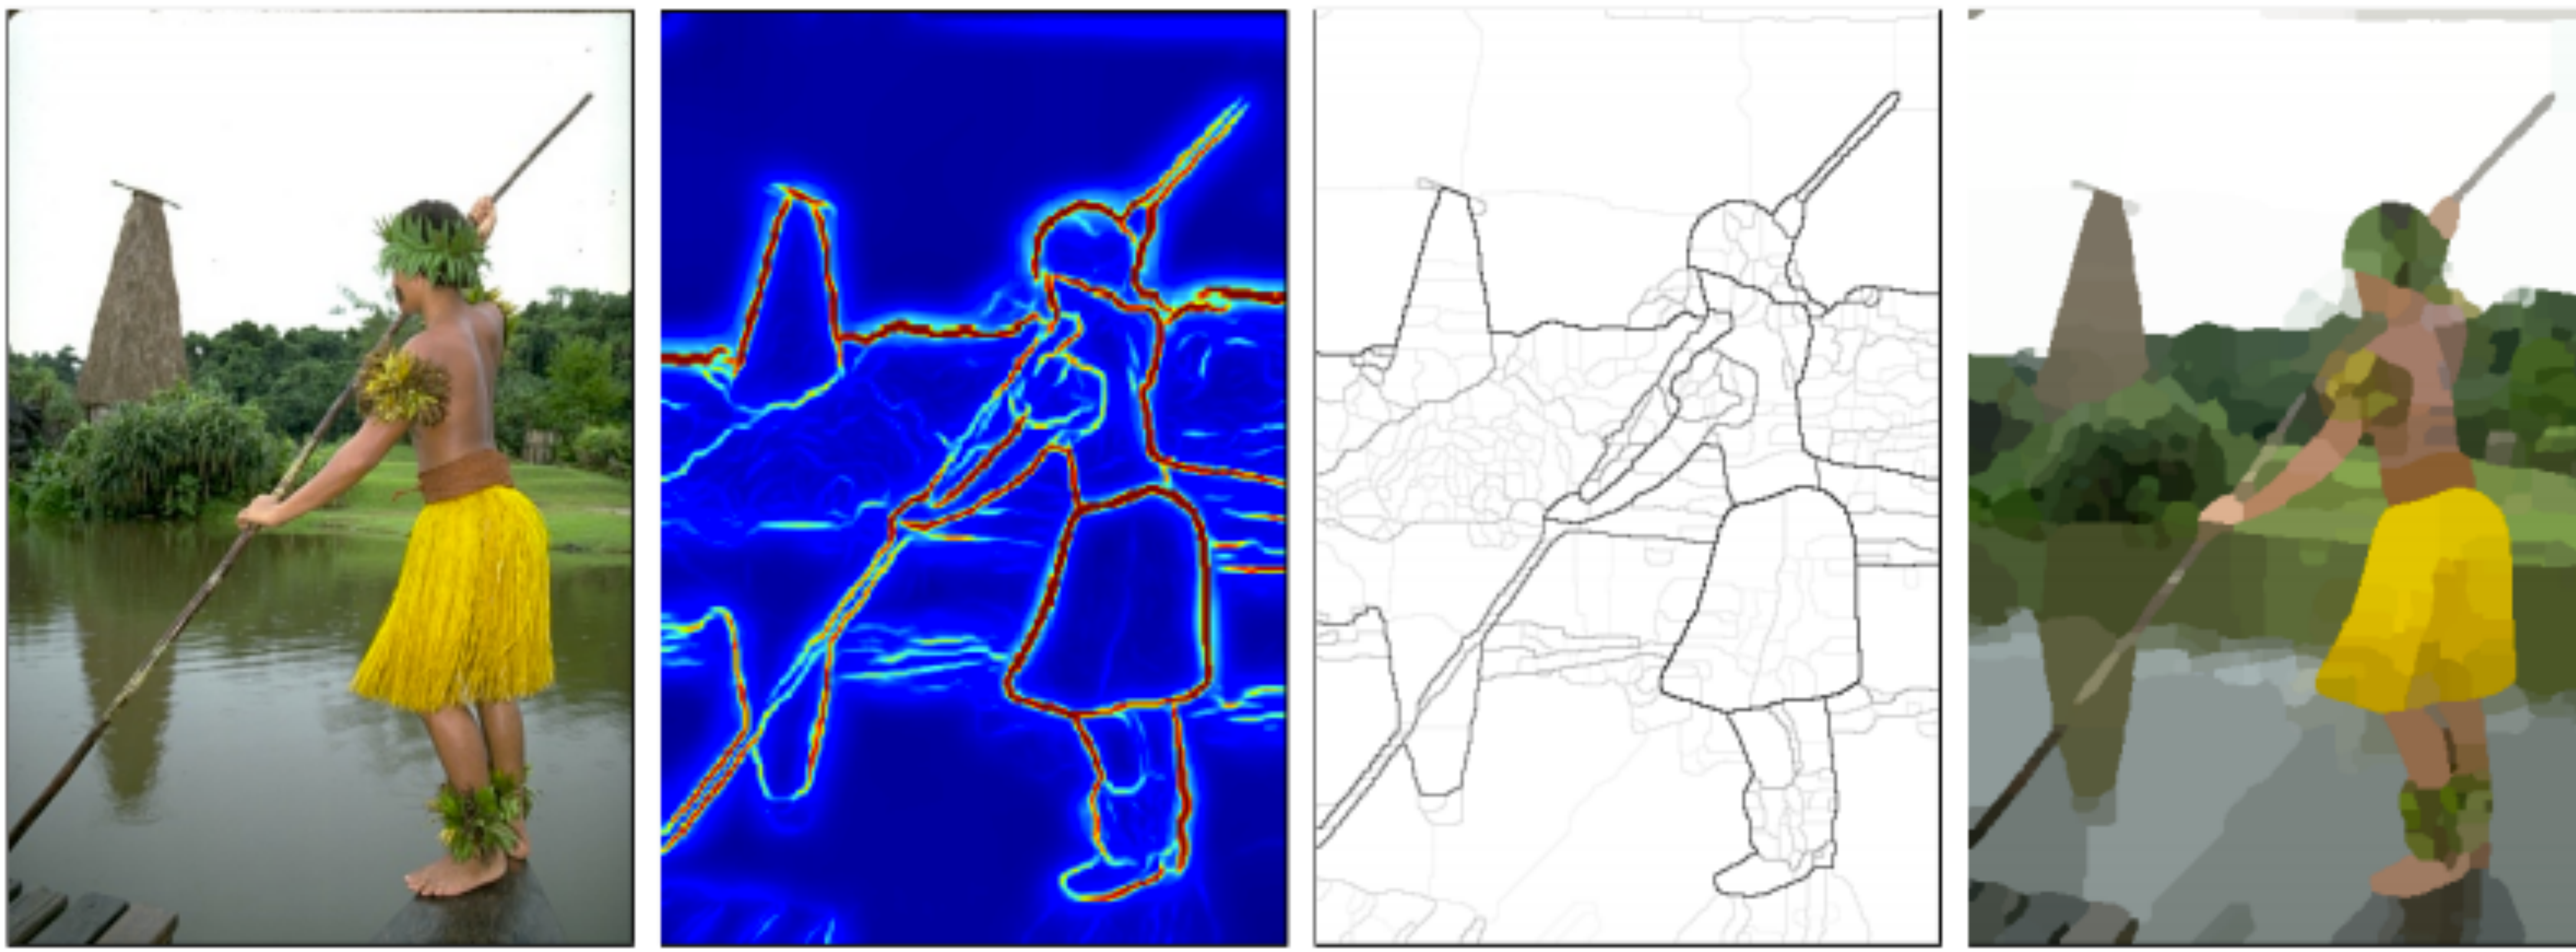

Why?

- Regions provide scale information
- Natural domain for computing features for recognition
- Super pixels

Segmentation

Watershed Transform

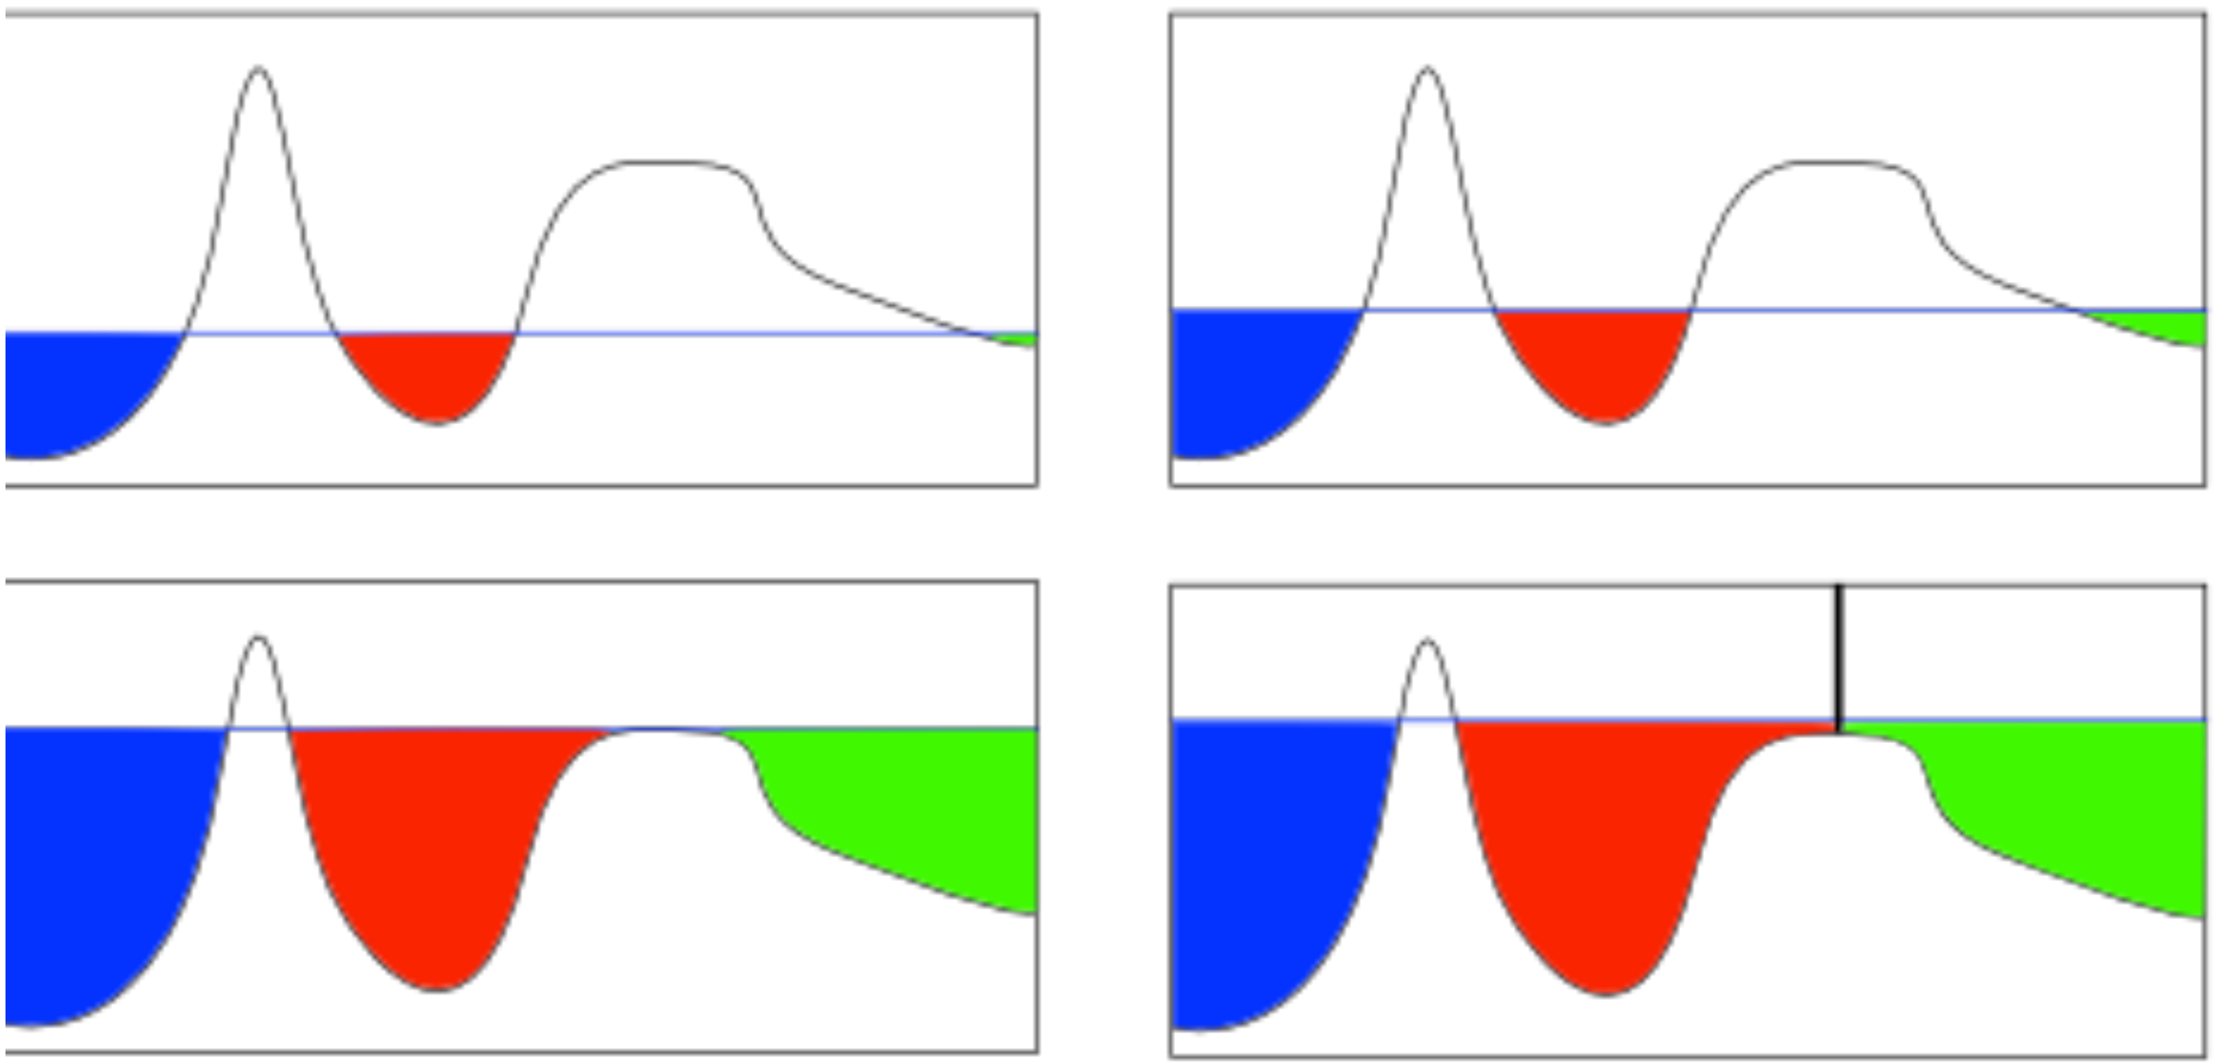

Image from
http://adessowiki.fee.unicamp.br/adesso/wiki/MICLab/view/

Segmentation

Watershed Transform

Regional minima

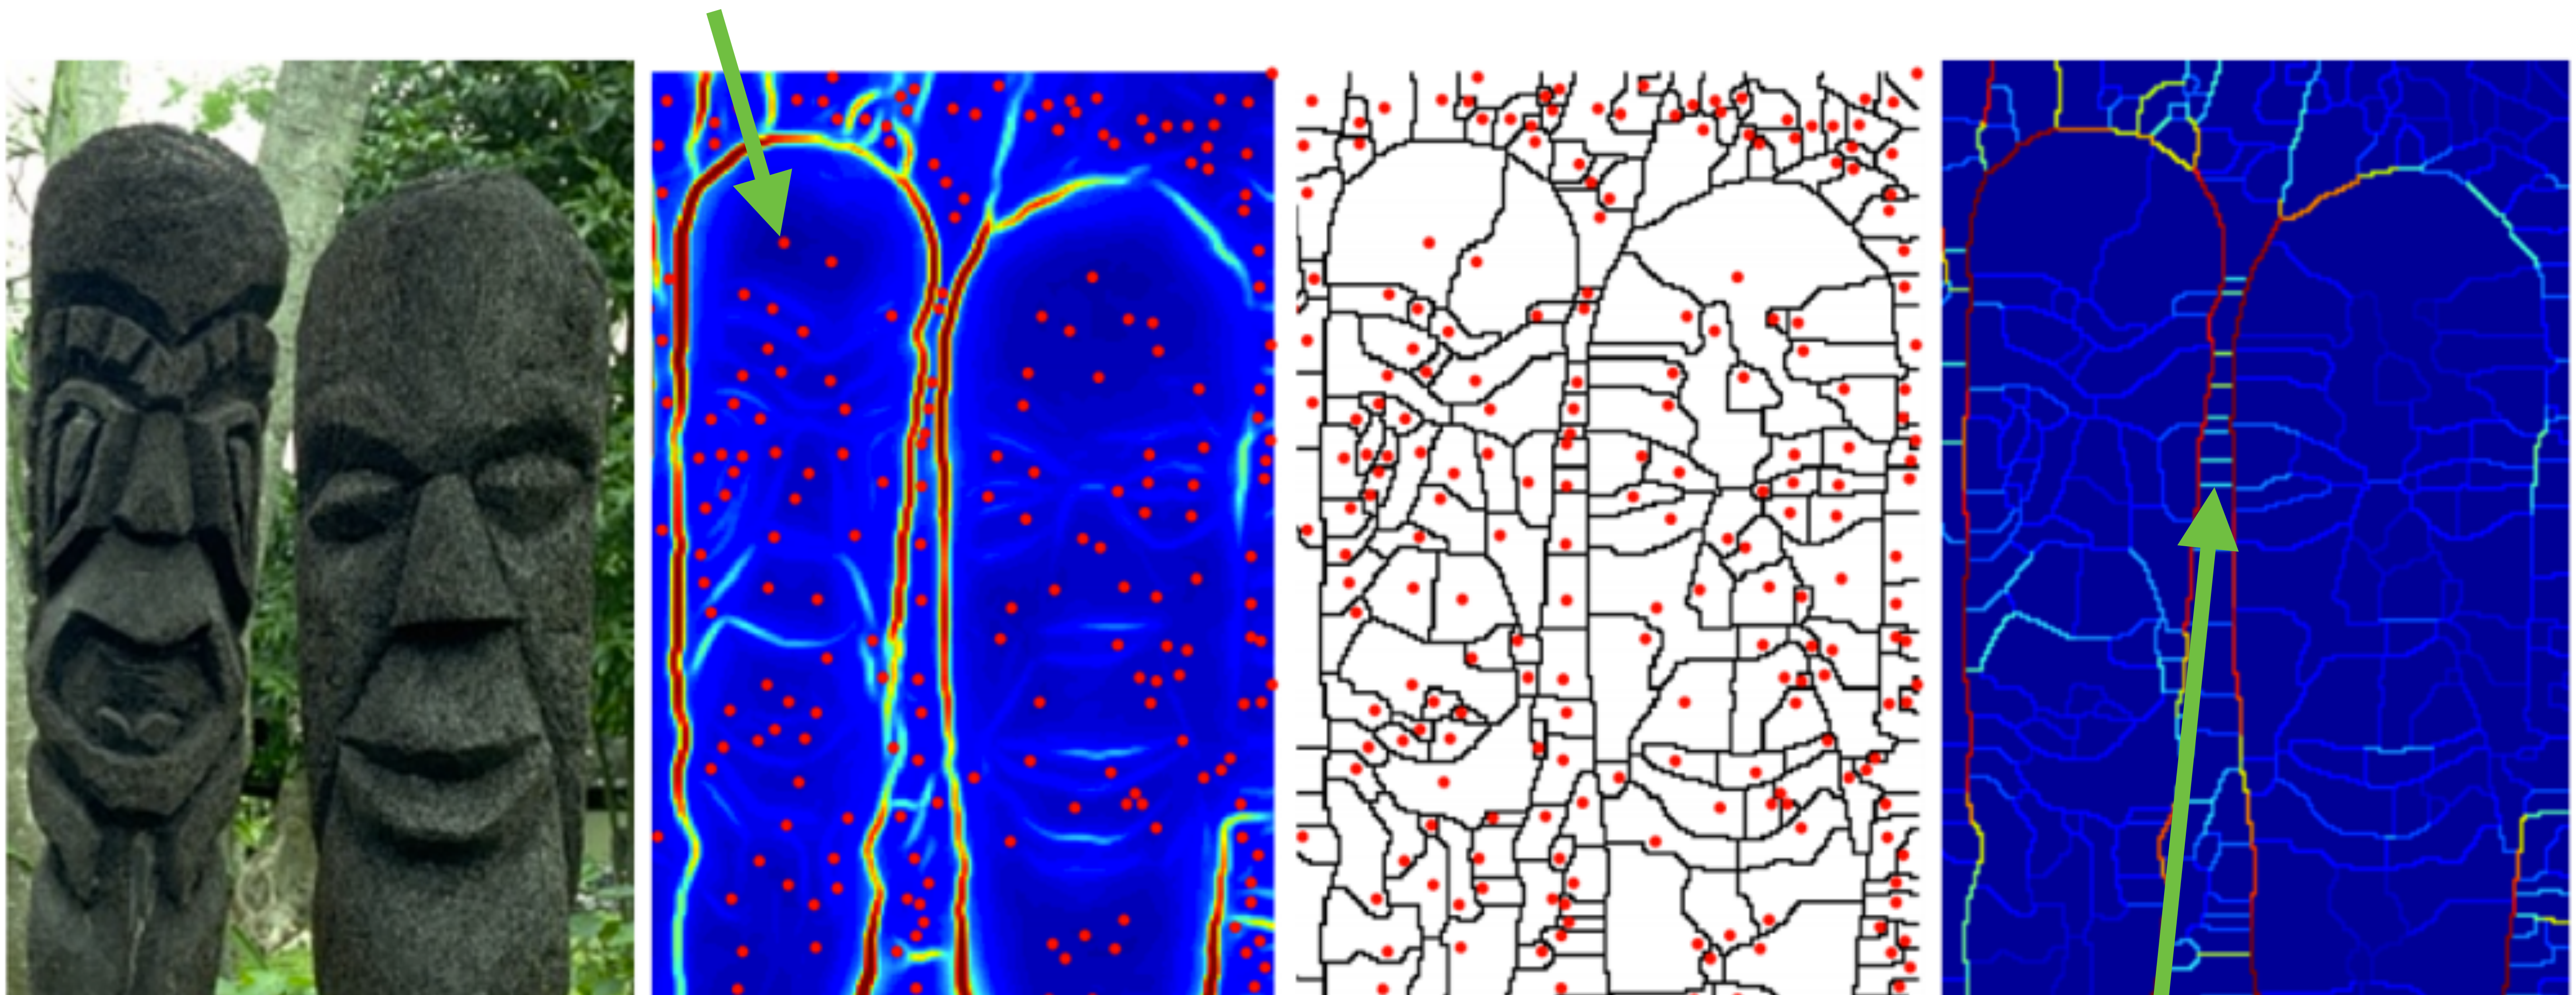

$$E(x, y) = \max_\theta E(x, y, \theta)$$

Artifacts

Weight of each arc = mean of E(x,y) for the pixels on the arc

Segmentation

“Oriented” Watershed Transform

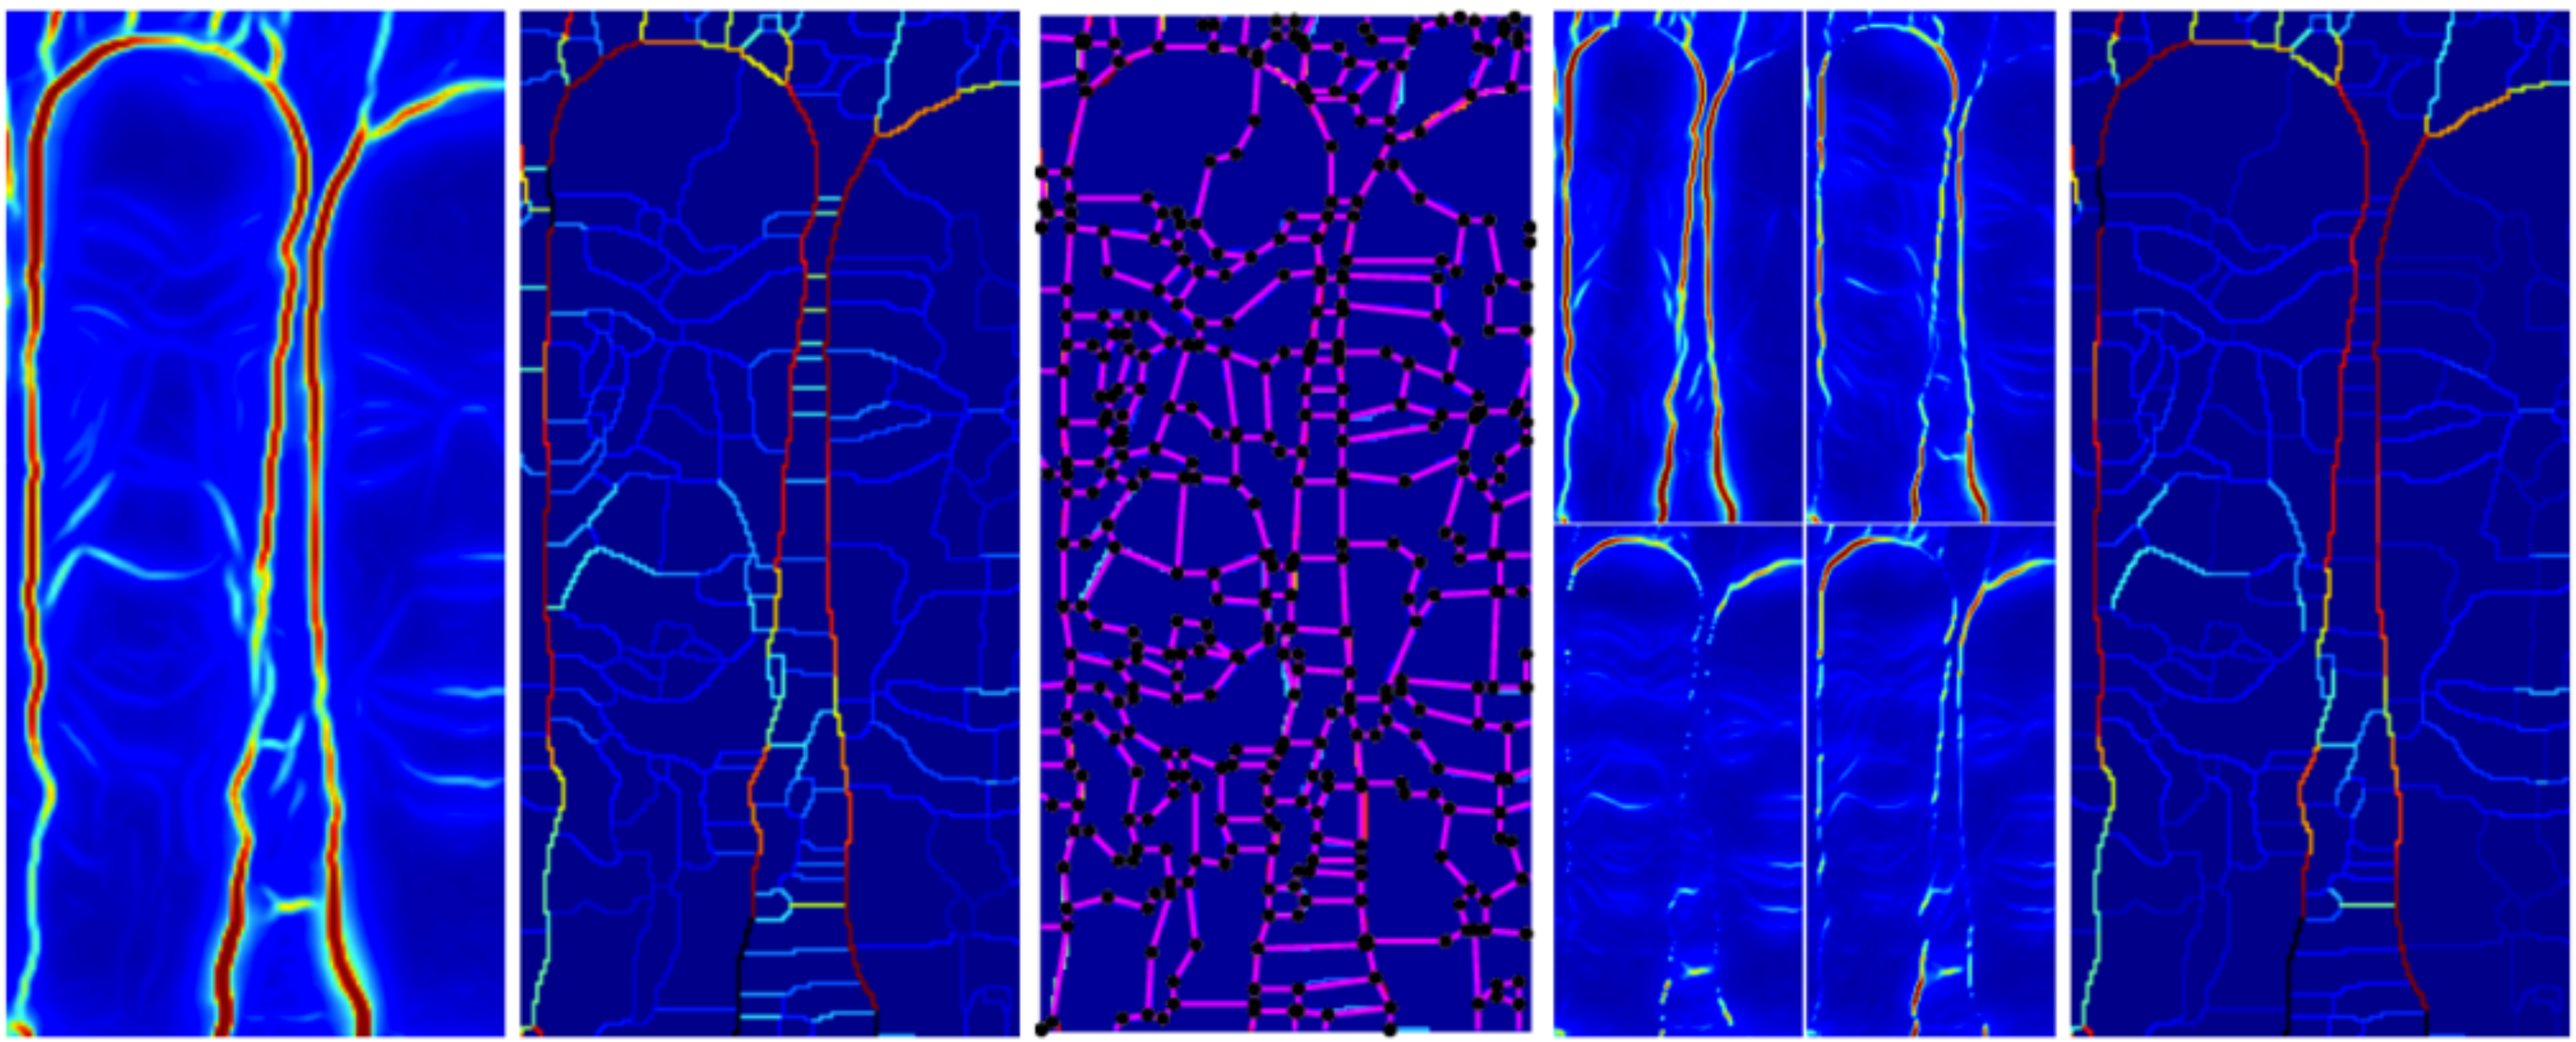

$\max_\theta E(x, y, \theta)$

$E(x, y, o(x, y))$

Enforce consistency in orientation

Hierarchical Segmentation

Represent Uncertainty of a Segmentation

3 Region Image

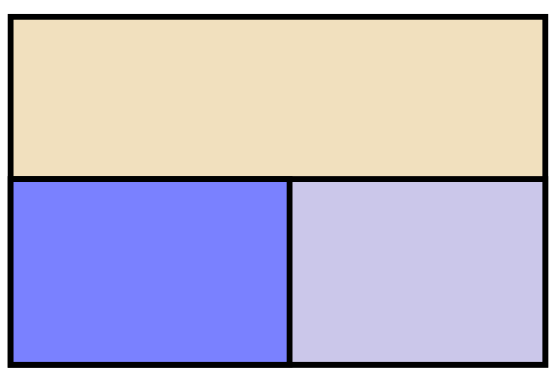

Dendrogram

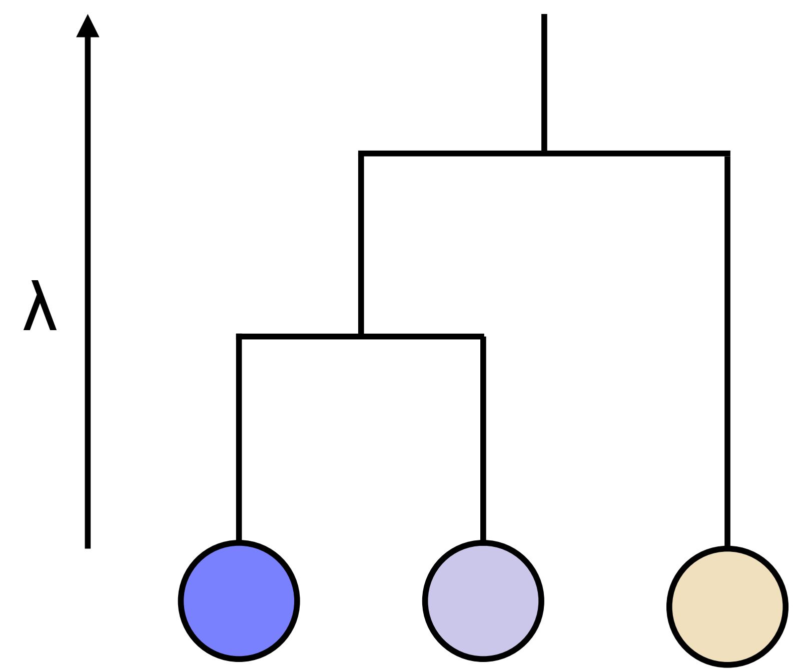

Contour Image

Slide by Derek Hoeim

Segmentation
Ultrametric Contour Map

- Ultrametric
  - Definition: D(x,y) <= max{ D(x,z), D(z,y) }

The union R_{12} of two regions R_1 and R_2 must have >= distance to adjacent region R_3 than either R_1 or R_2

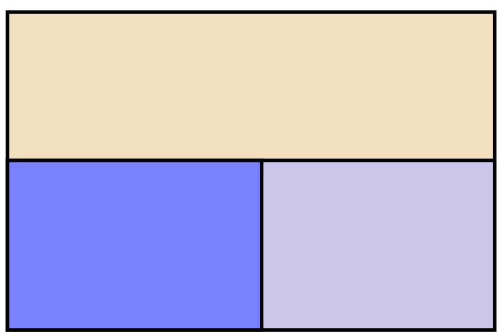

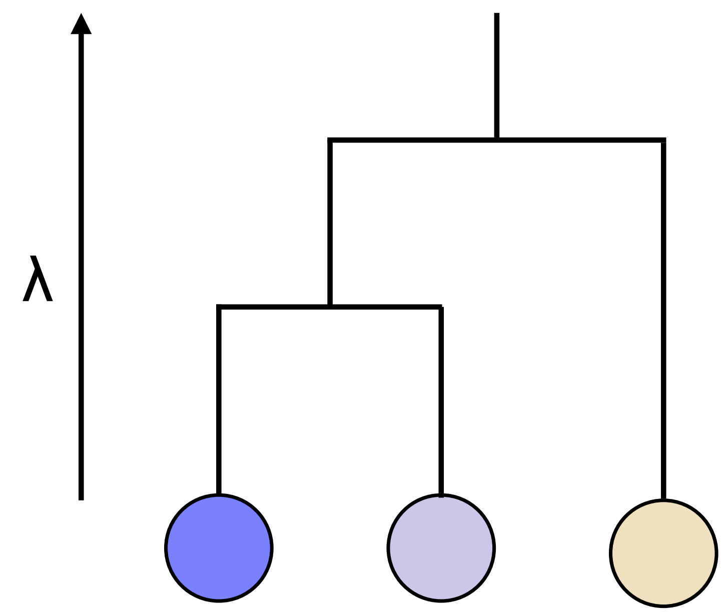

Slide by Derek Hoiem

Segmentation
Ultrametric Contour Map

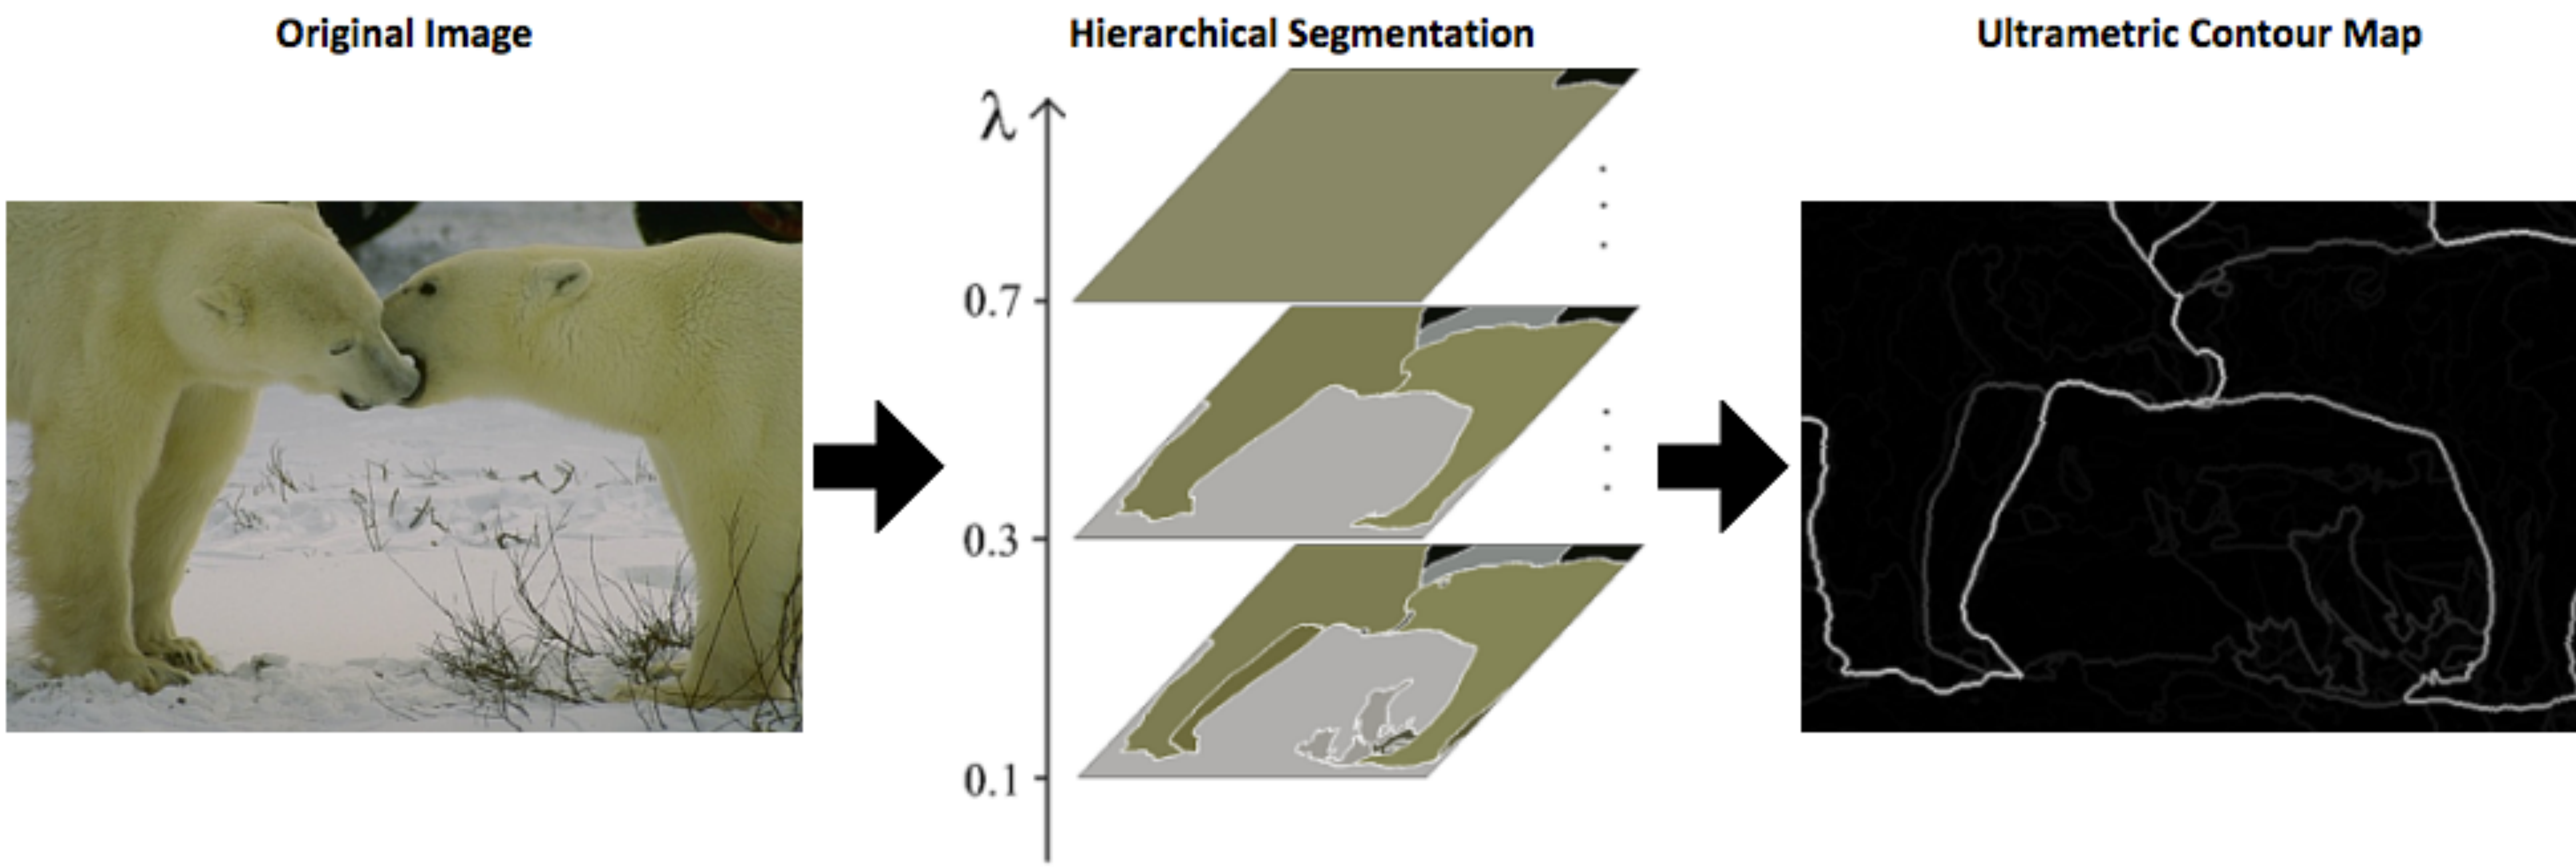

Segmentation

Results

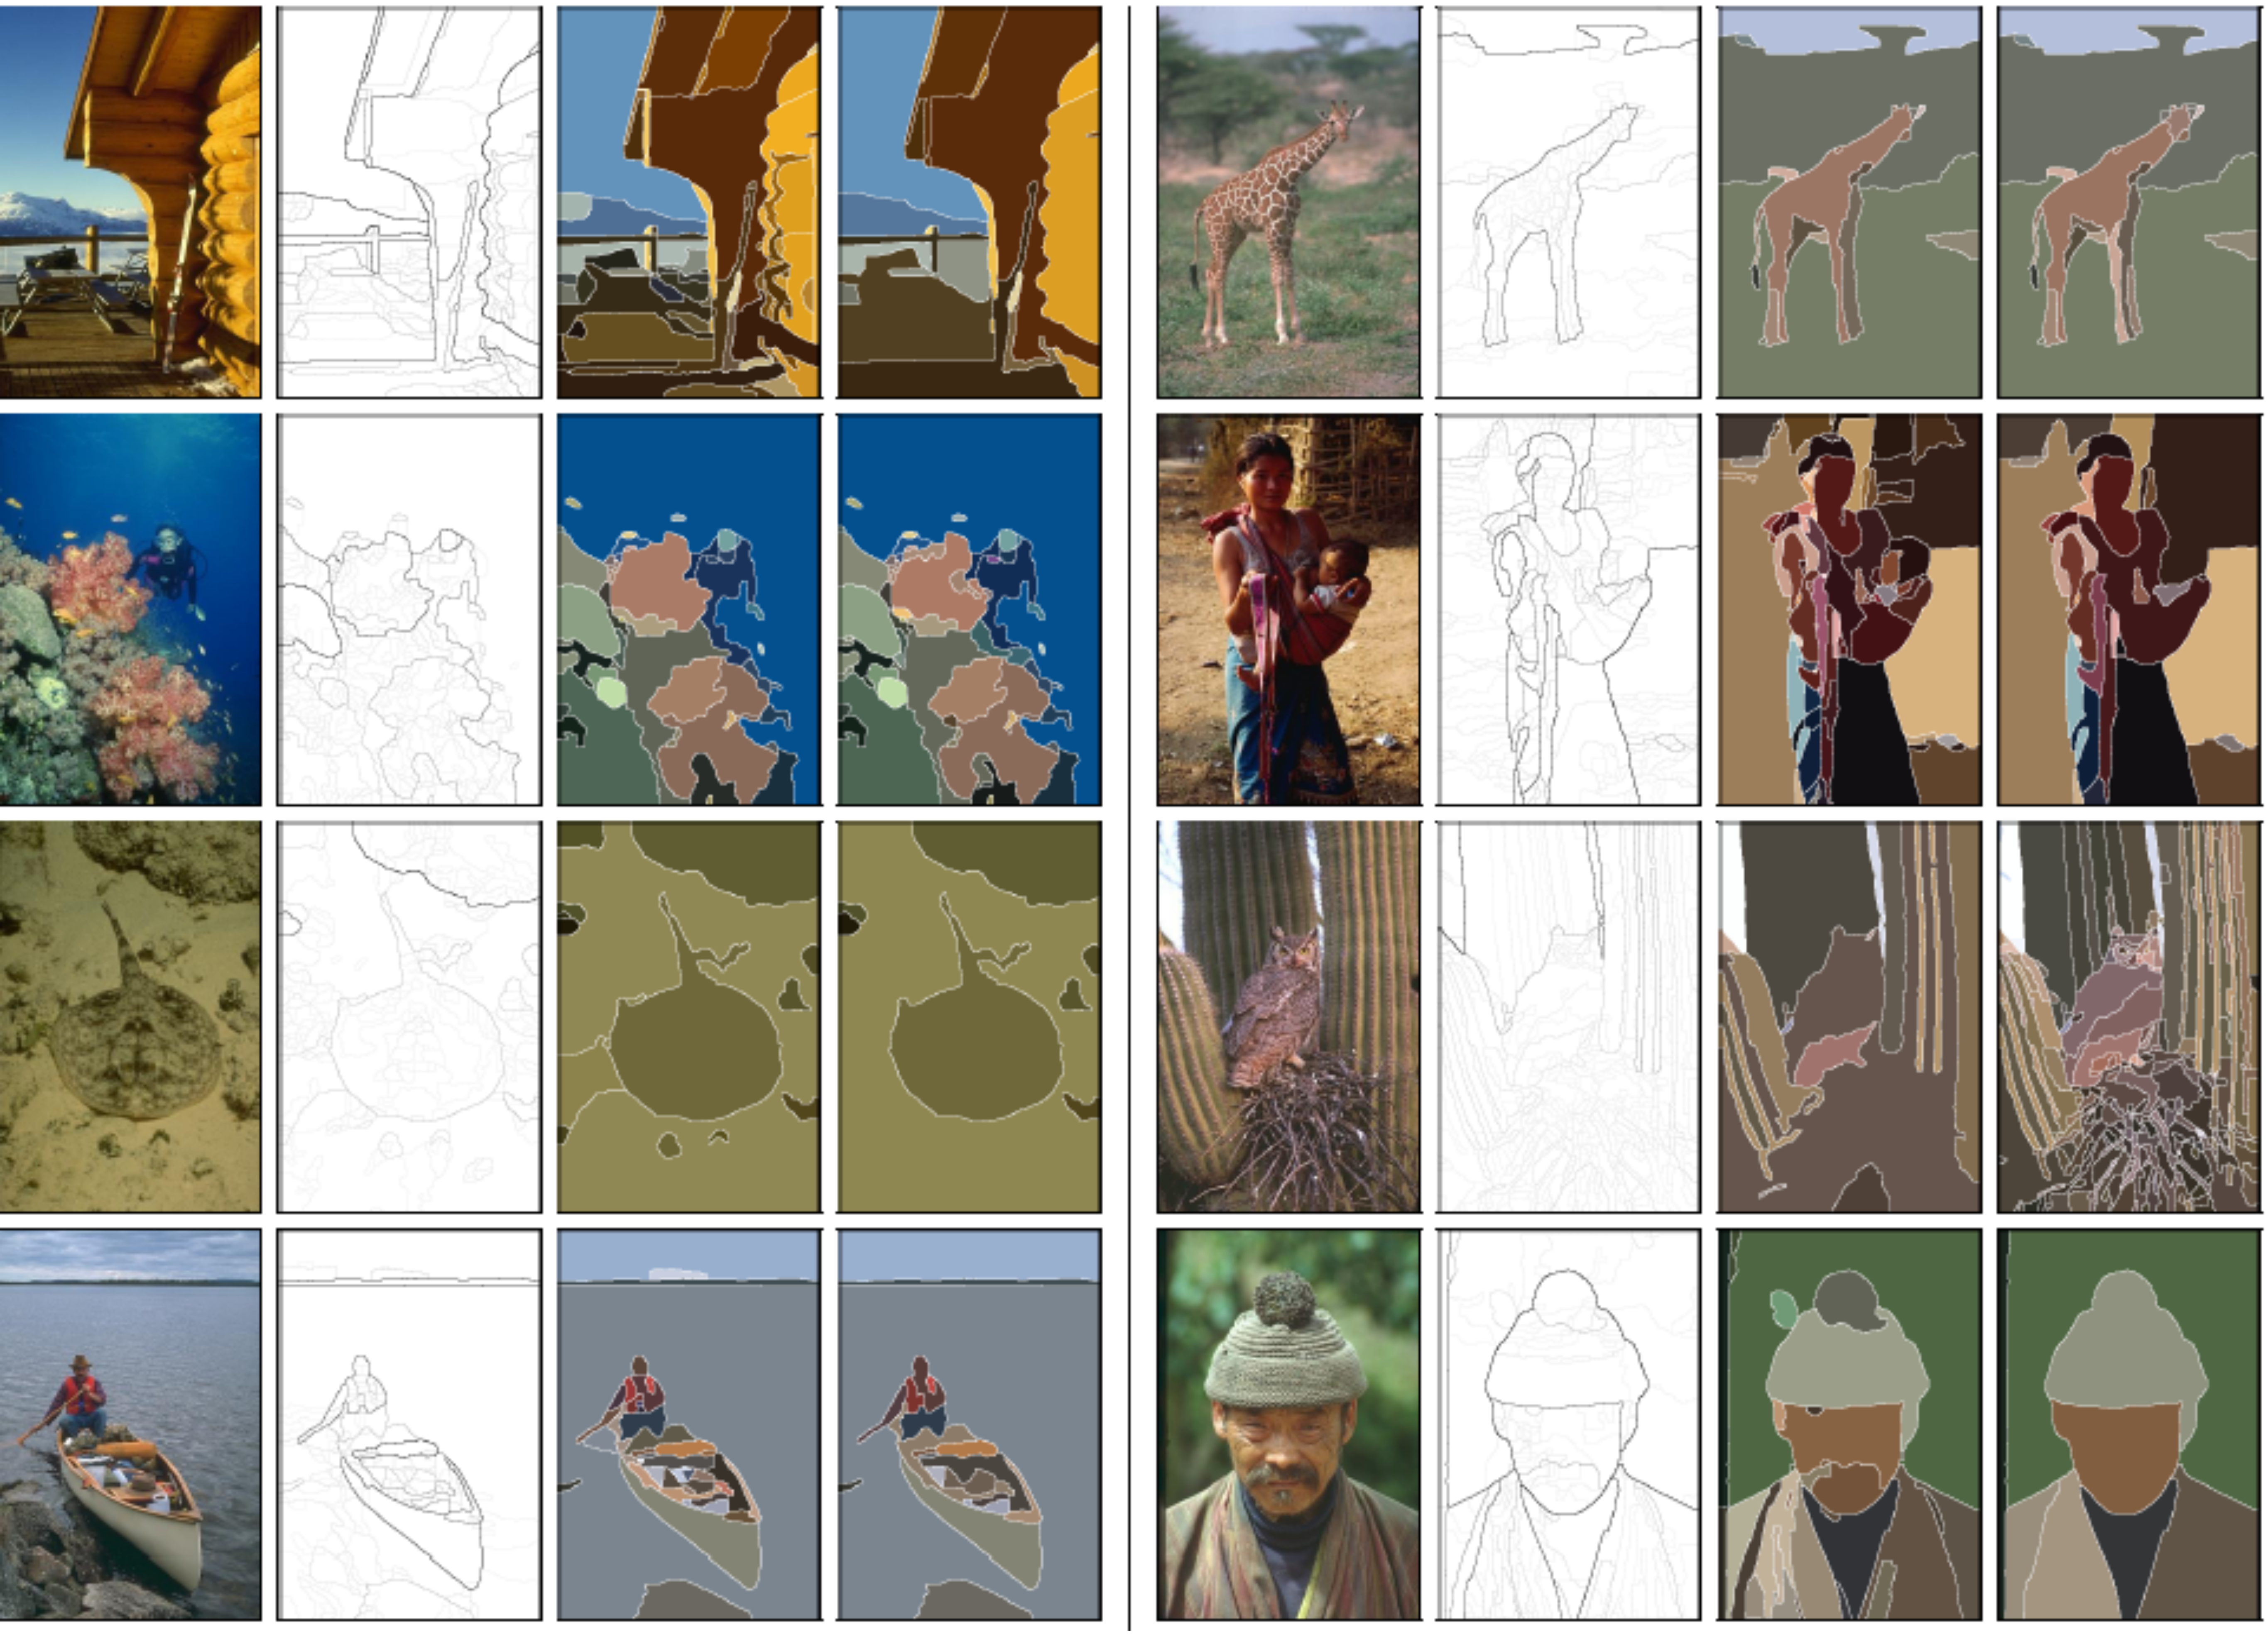

Evaluation on BSDS

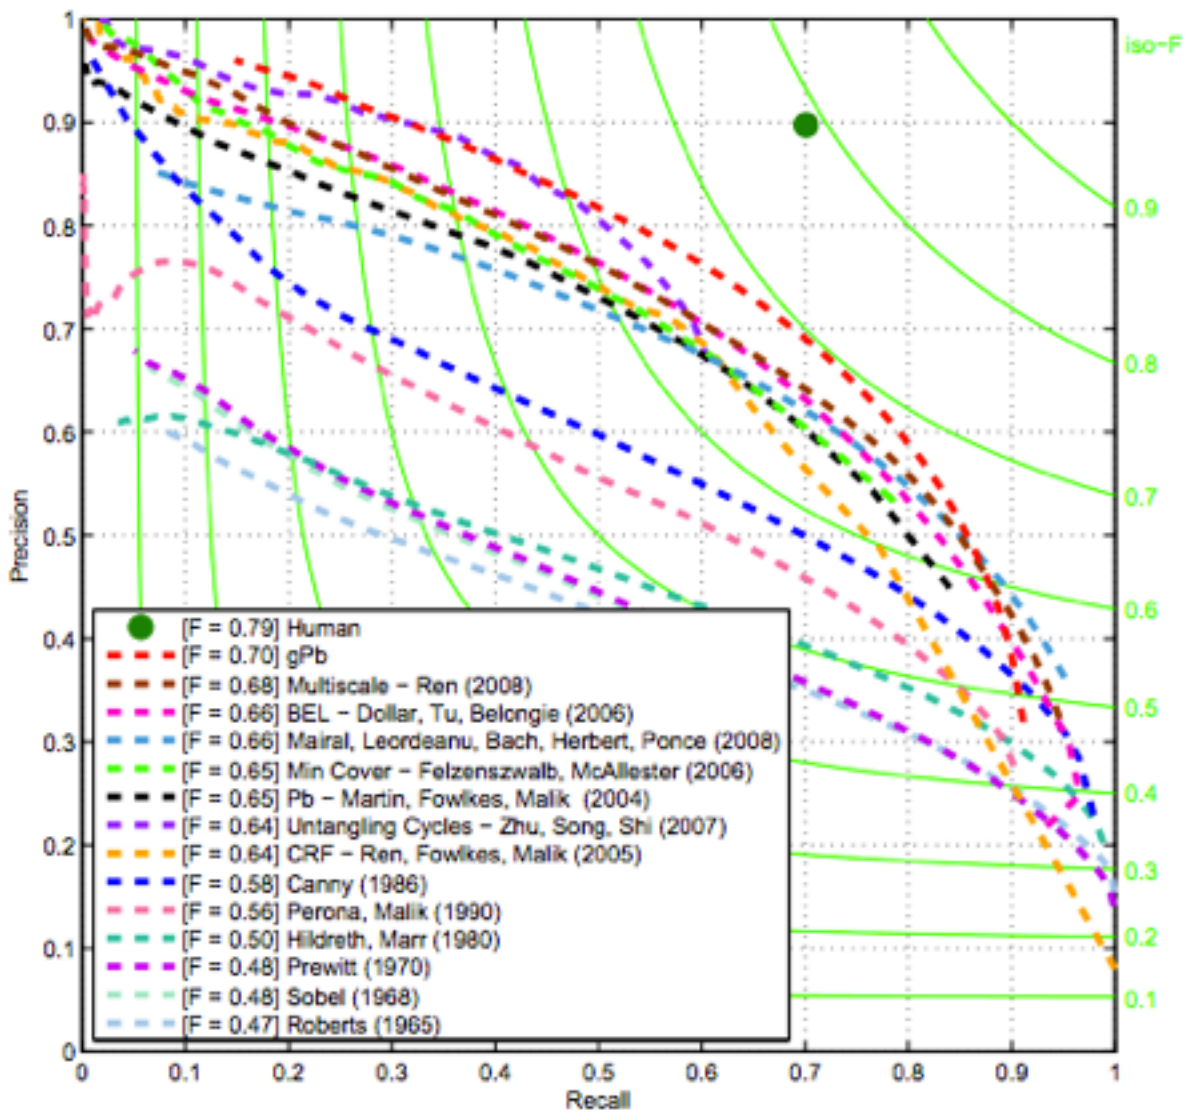

Boundary(Contour) detector

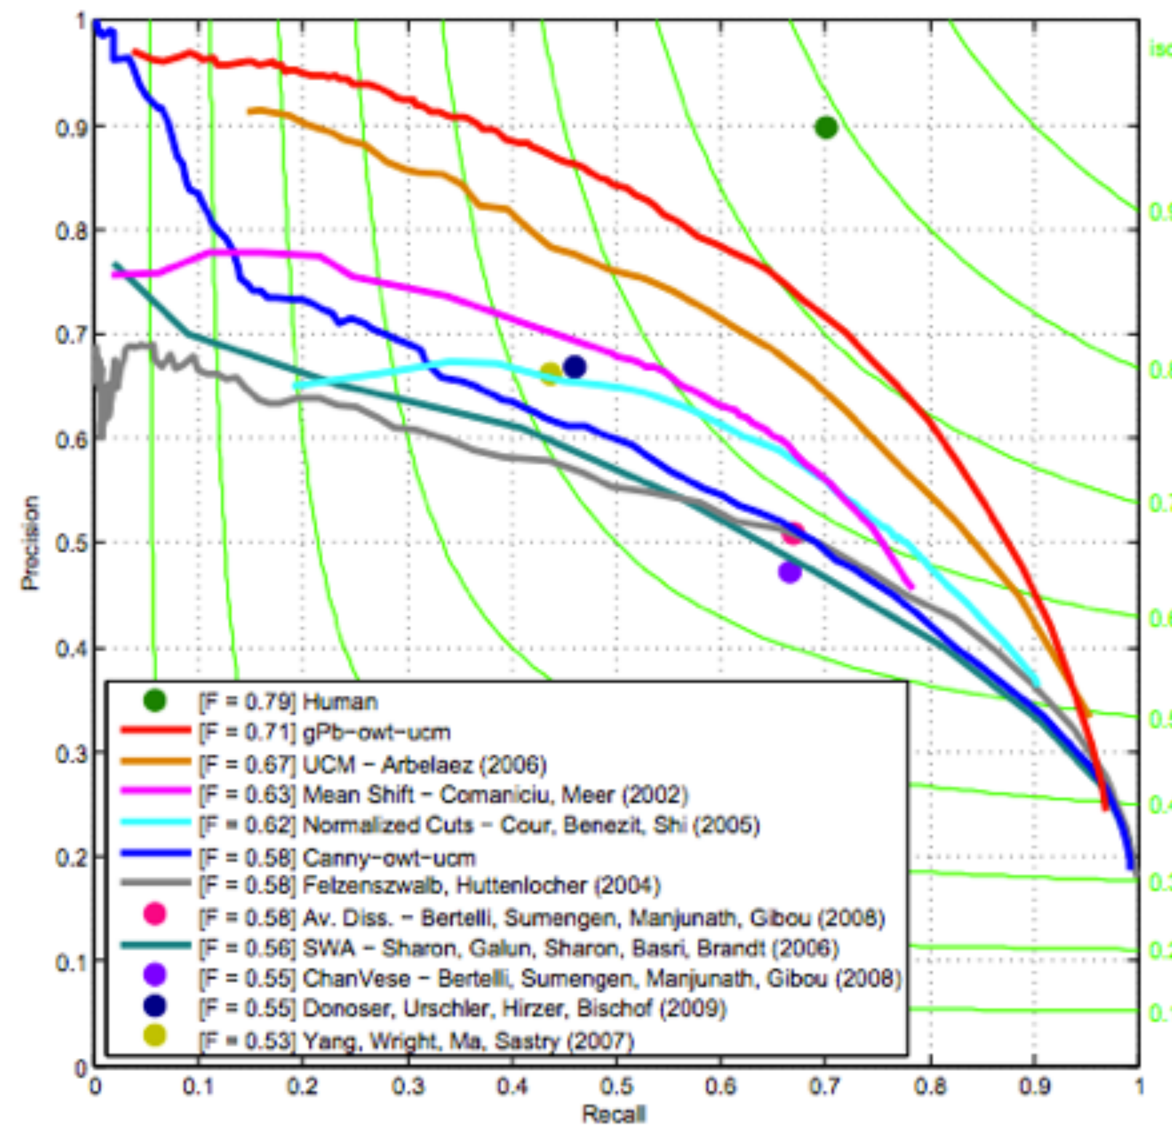

Segmentation

Evaluation on BSDS

Boundary benchmarks

	BSDS300			BSDS500		
	ODS	OIS	AP	ODS	OIS	AP
Human	0.79	0.79	—	0.80	0.80	—
gPb-owt-ucm	**0.71**	**0.74**	**0.73**	**0.73**	**0.76**	**0.73**
[34] Mean Shift	0.63	0.66	0.54	0.64	0.68	0.56
[33] NCuts	0.62	0.66	0.43	0.64	0.68	0.45
Canny-owt-ucm	0.58	0.63	0.58	0.60	0.64	0.58
[32] Felz-Hutt	0.58	0.62	0.53	0.61	0.64	0.56
[31] SWA	0.56	0.59	0.54	—	—	—
Quad-Tree	0.37	0.39	0.26	0.38	0.39	0.26
gPb	0.70	0.72	0.66	0.71	0.74	0.65
Canny	0.58	0.62	0.58	0.60	0.63	0.58

$$F = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

Evaluation on BSDS

Region benchmarks

	BSDS300						
	Covering			PRI		VI	
	ODS	OIS	Best	ODS	OIS	ODS	OIS
Human	0.73	0.73	–	0.87	0.87	1.16	1.16
gPb-owt-ucm	**0.59**	**0.65**	**0.75**	**0.81**	**0.85**	**1.65**	**1.47**
[34] Mean Shift	0.54	0.58	0.66	0.78	0.80	1.83	1.63
[32] Felz-Hutt	0.51	0.58	0.68	0.77	0.82	2.15	1.79
Canny-owt-ucm	0.48	0.56	0.66	0.77	0.82	2.11	1.81
[33] NCuts	0.44	0.53	0.66	0.75	0.79	2.18	1.84
[31] SWA	0.47	0.55	0.66	0.75	0.80	2.06	1.75
[29] Total Var.	0.57	–	–	0.78	–	1.81	–
[70] T+B Encode	0.54	–	–	0.78	–	1.86	–
[30] Av. Diss.	0.47	–	–	0.76	–	2.62	–
[30] ChanVese	0.49	–	–	0.75	–	2.54	–
Quad-Tree	0.33	0.39	0.47	0.71	0.75	2.34	2.22

Evaluation on BSDS

Region benchmarks

	BSDS500						
	Covering			PRI		VI	
	ODS	OIS	Best	ODS	OIS	ODS	OIS
Human	0.72	0.72	–	0.88	0.88	1.17	1.17
gPb-owt-ucm	**0.59**	**0.65**	**0.74**	**0.83**	**0.86**	**1.69**	**1.48**
[34] Mean Shift	0.54	0.58	0.66	0.79	0.81	1.85	1.64
[32] Felz-Hutt	0.52	0.57	0.69	0.80	0.82	2.21	1.87
Canny-owt-ucm	0.49	0.55	0.66	0.79	0.83	2.19	1.89
[33] NCuts	0.45	0.53	0.67	0.78	0.80	2.23	1.89
Quad-Tree	0.32	0.37	0.46	0.73	0.74	2.46	2.32

Discussions

Any better local cues?

F-measure based gradient decent?

Why Normalized cut capture global cue

Intuition of the Ultrametric?

Greedy(watershed) vs Graph-cut

Can handle bad illumination?

Hierarchies vs. Multiple Segmentations

References

- P. Arbelaez, M. Maire, C. Fowlkes and J. Malik, Contour detection and hierarchical image segmentation, TPAMI, 2011.
- P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik, From contours to regions: An empirical evaluation, CVPR, 2009.
- M. Maire, P. Arbeláez, C. Fowlkes, and J. Malik, Using contours to detect and localize junctions in natural images. CVPR, 2008.
- P. Arbelaez, Boundary Extraction in Natural Images Using Ultrametric Contour Maps, POCV, 2006
- Derek Hoiem's slides
- Elad Liebman, http://www.cs.utexas.edu/~cv-fall2012/slides/elad-expt.pdf

Backup Slides

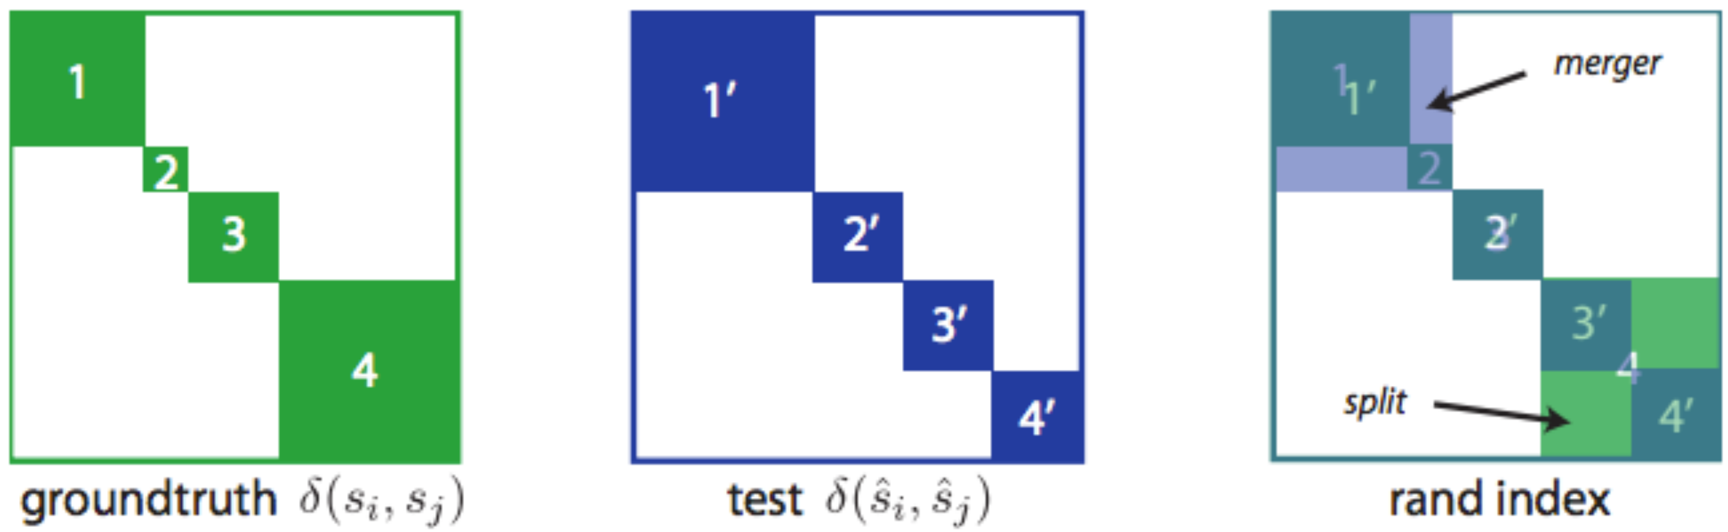

Figure 2: **The Rand index quantifies segmentation performance by comparing the difference in pixel pair connectivity between the groundtruth and test segmentations.** Pixel pair connectivities can be visualized as symmetric binary block-diagonal matrices $\delta(s_i, s_j)$. Each diagonal block corresponds to connected pixel pairs belonging to one of the image segments. The Rand index incurs penalties when pixels pairs that must not be connected are connected or vice versa. This corresponds to locations where the two matrices disagree. An erroneous merger of two groundtruth segments incurs a penalty proportional to the product of the sizes of the two segments. Split errors are similarly penalized.

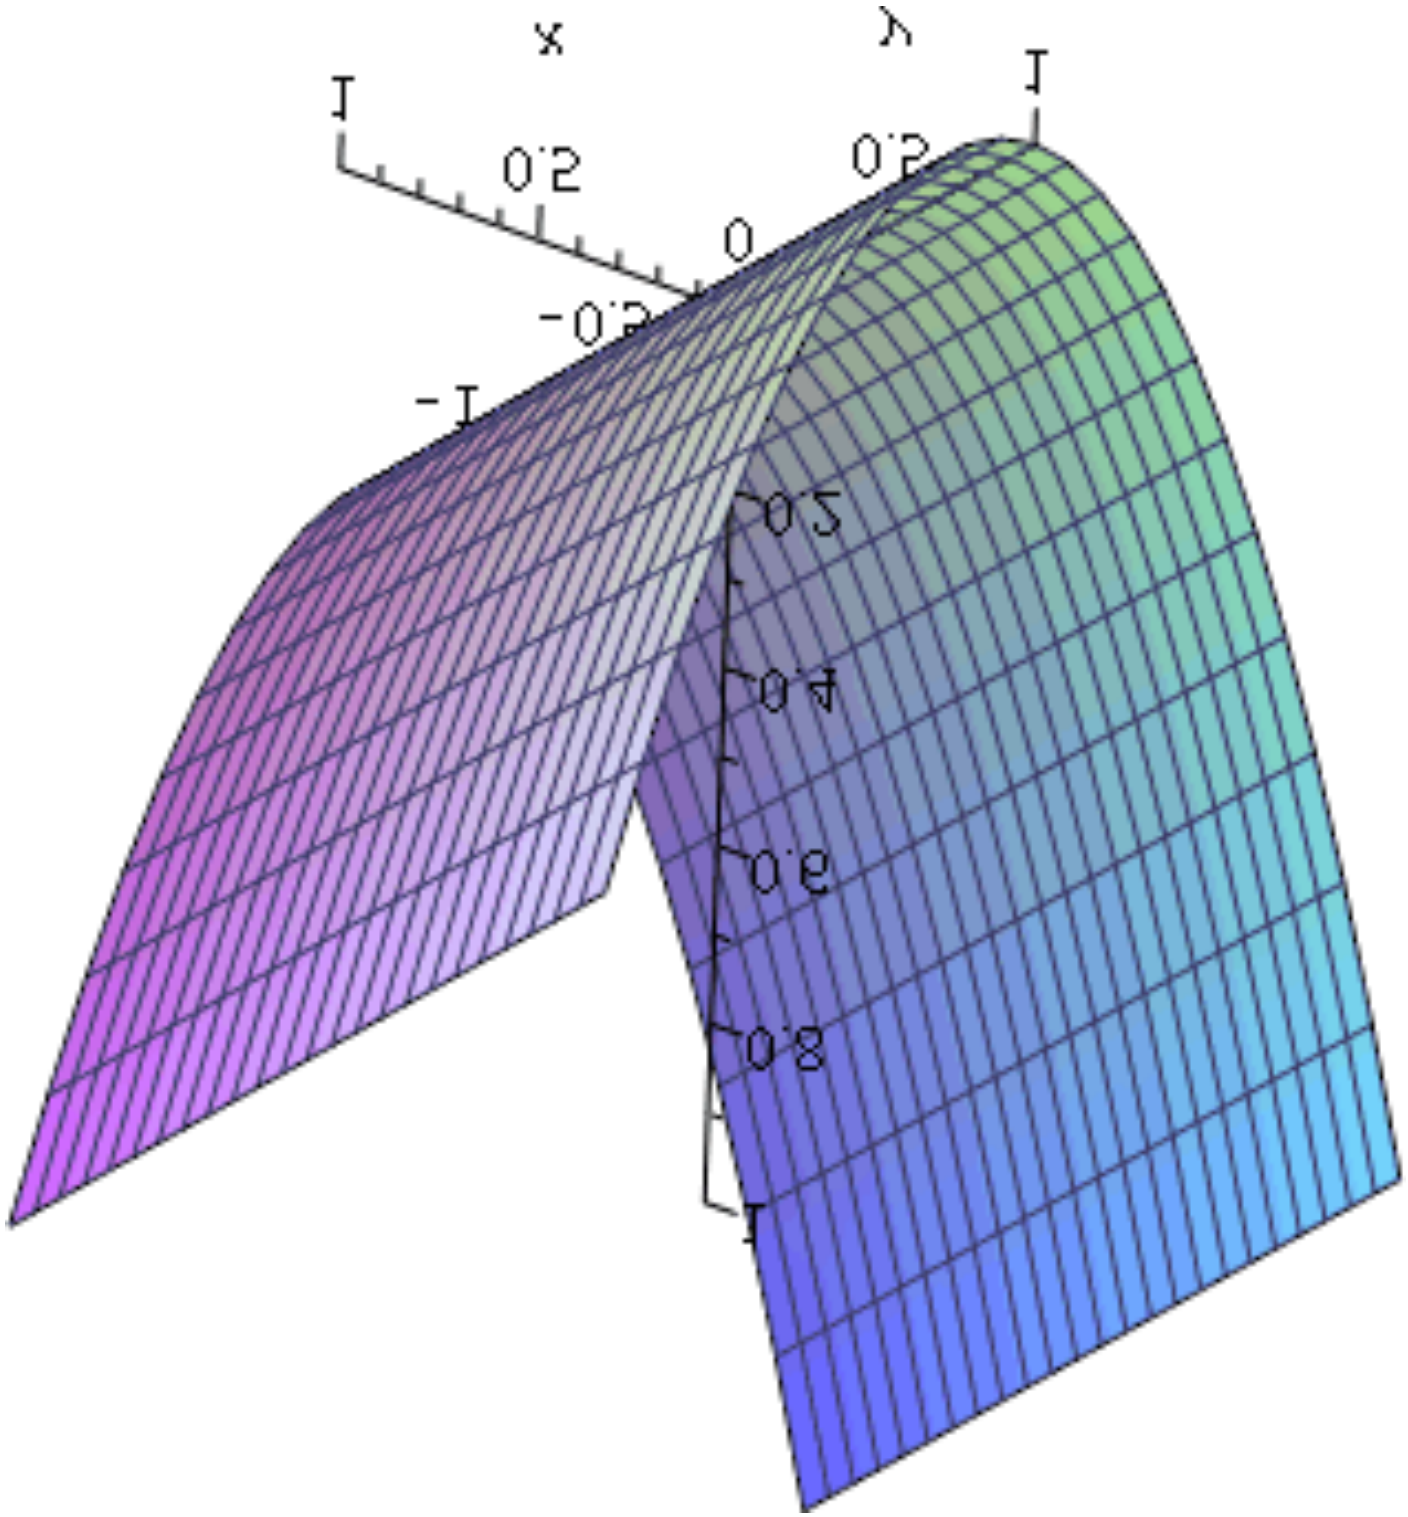
x
y
1
0.5
0
-0.5
-1
0.2
0.4
0.6
0.8

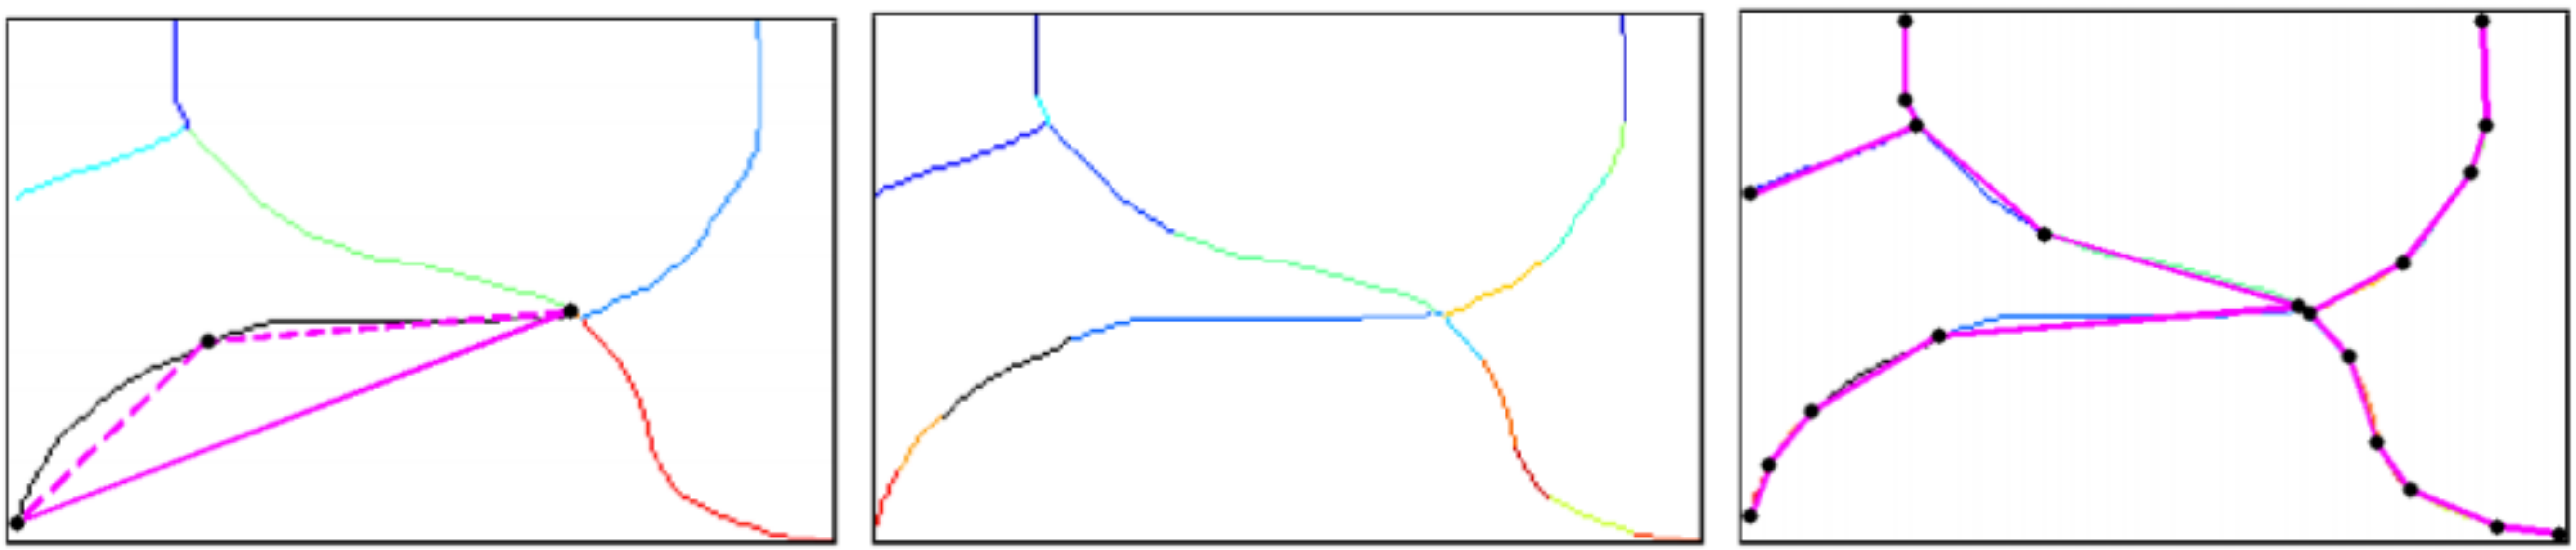

Benchmarks

Region-based Metrics

- Variation of Information

$$VI(S, S') = H(S) + H(S') - 2I(S, S')$$

Entropy **Mutual information**

- Rand Index **Sum of labeling error**

$$PRI(S, \{G_k\}) = \frac{1}{T} \sum_{i<j} [c_{ij} p_{ij} + (1 - c_{ij})(1 - p_{ij})]$$

i,j same label

- Segmentation Covering

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

Overlap **Covering**

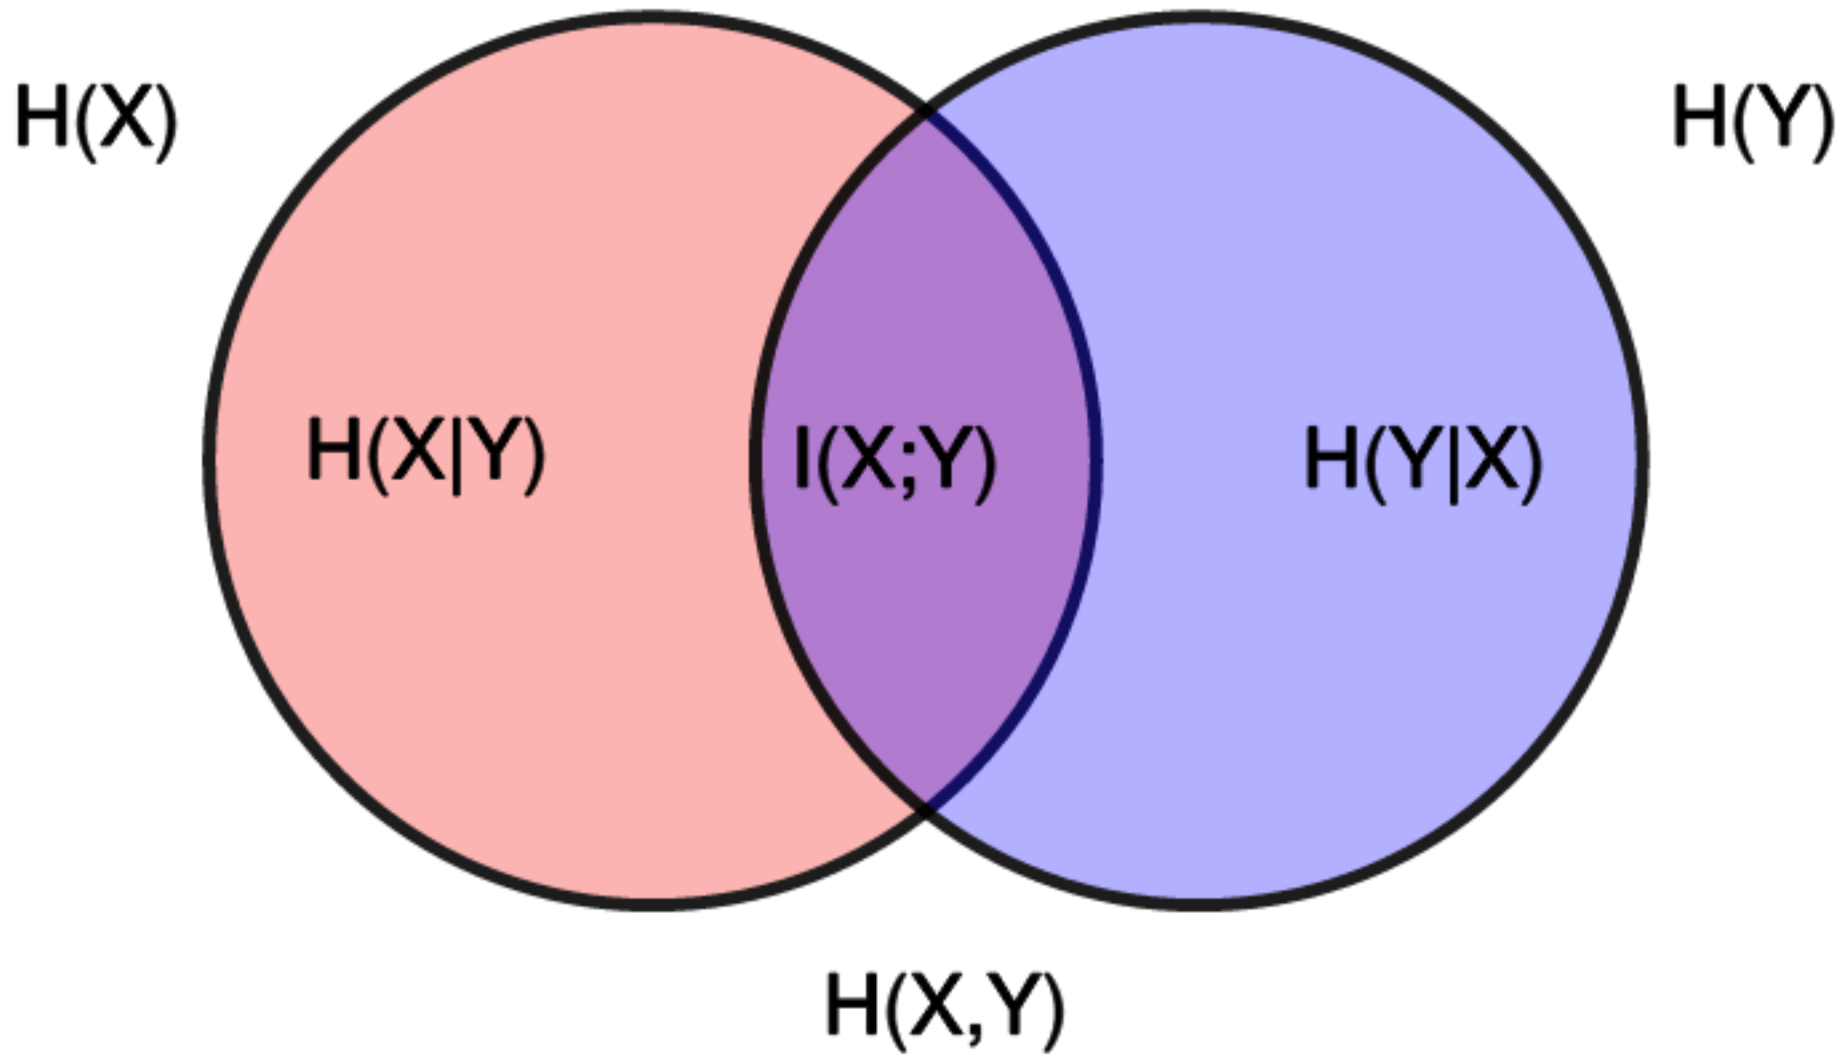
H(X)
H(Y)
H(X|Y)
I(X;Y)
H(Y|X)
H(X,Y)

Previous Work

- Contours
  - Roberts, Sobel, Prewitt
  - Canny
  - etc

- Regions
  - Felz-Hutt
  - Normalized Cuts
  - etc

+1	0
0	-1

Gx

0	+1
-1	0

Gy

-1	0	+1
-2	0	+2
-1	0	+1

Gx

+1	+2	+1
0	0	0
-1	-2	-1

Gy

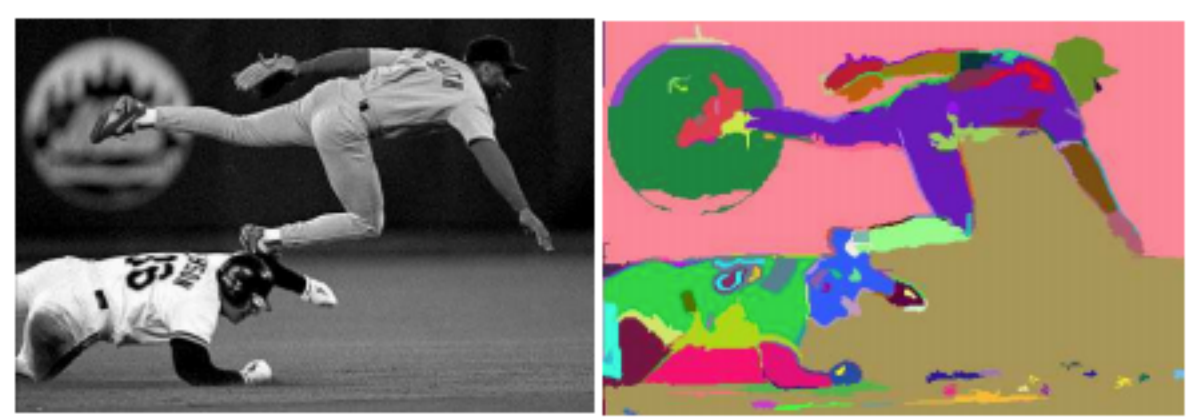

Segmentation
Ultrametric Contour Map

- Hierarchal region representation
- Greedy graph-based region margin algorithm

1) Select minimum weight contour:
   $C^* = \text{argmin}_{C \in \mathcal{K}_0} W(C)$.
2) Let $R_1, R_2 \in \mathcal{P}_0$ be the regions separated by C^*.
3) Set $R = R_1 \cup R_2$, and update:
   $\mathcal{P}_0 \leftarrow \mathcal{P}_0 \backslash \{R_1, R_2\} \cup \{R\}$ and $\mathcal{K}_0 \leftarrow \mathcal{K}_0 \backslash \{C^*\}$.
4) Stop if $\mathcal{K}_0$ is empty.
   Otherwise, update weights $W(\mathcal{K}_0)$ and repeat.

Interactive Segmentation

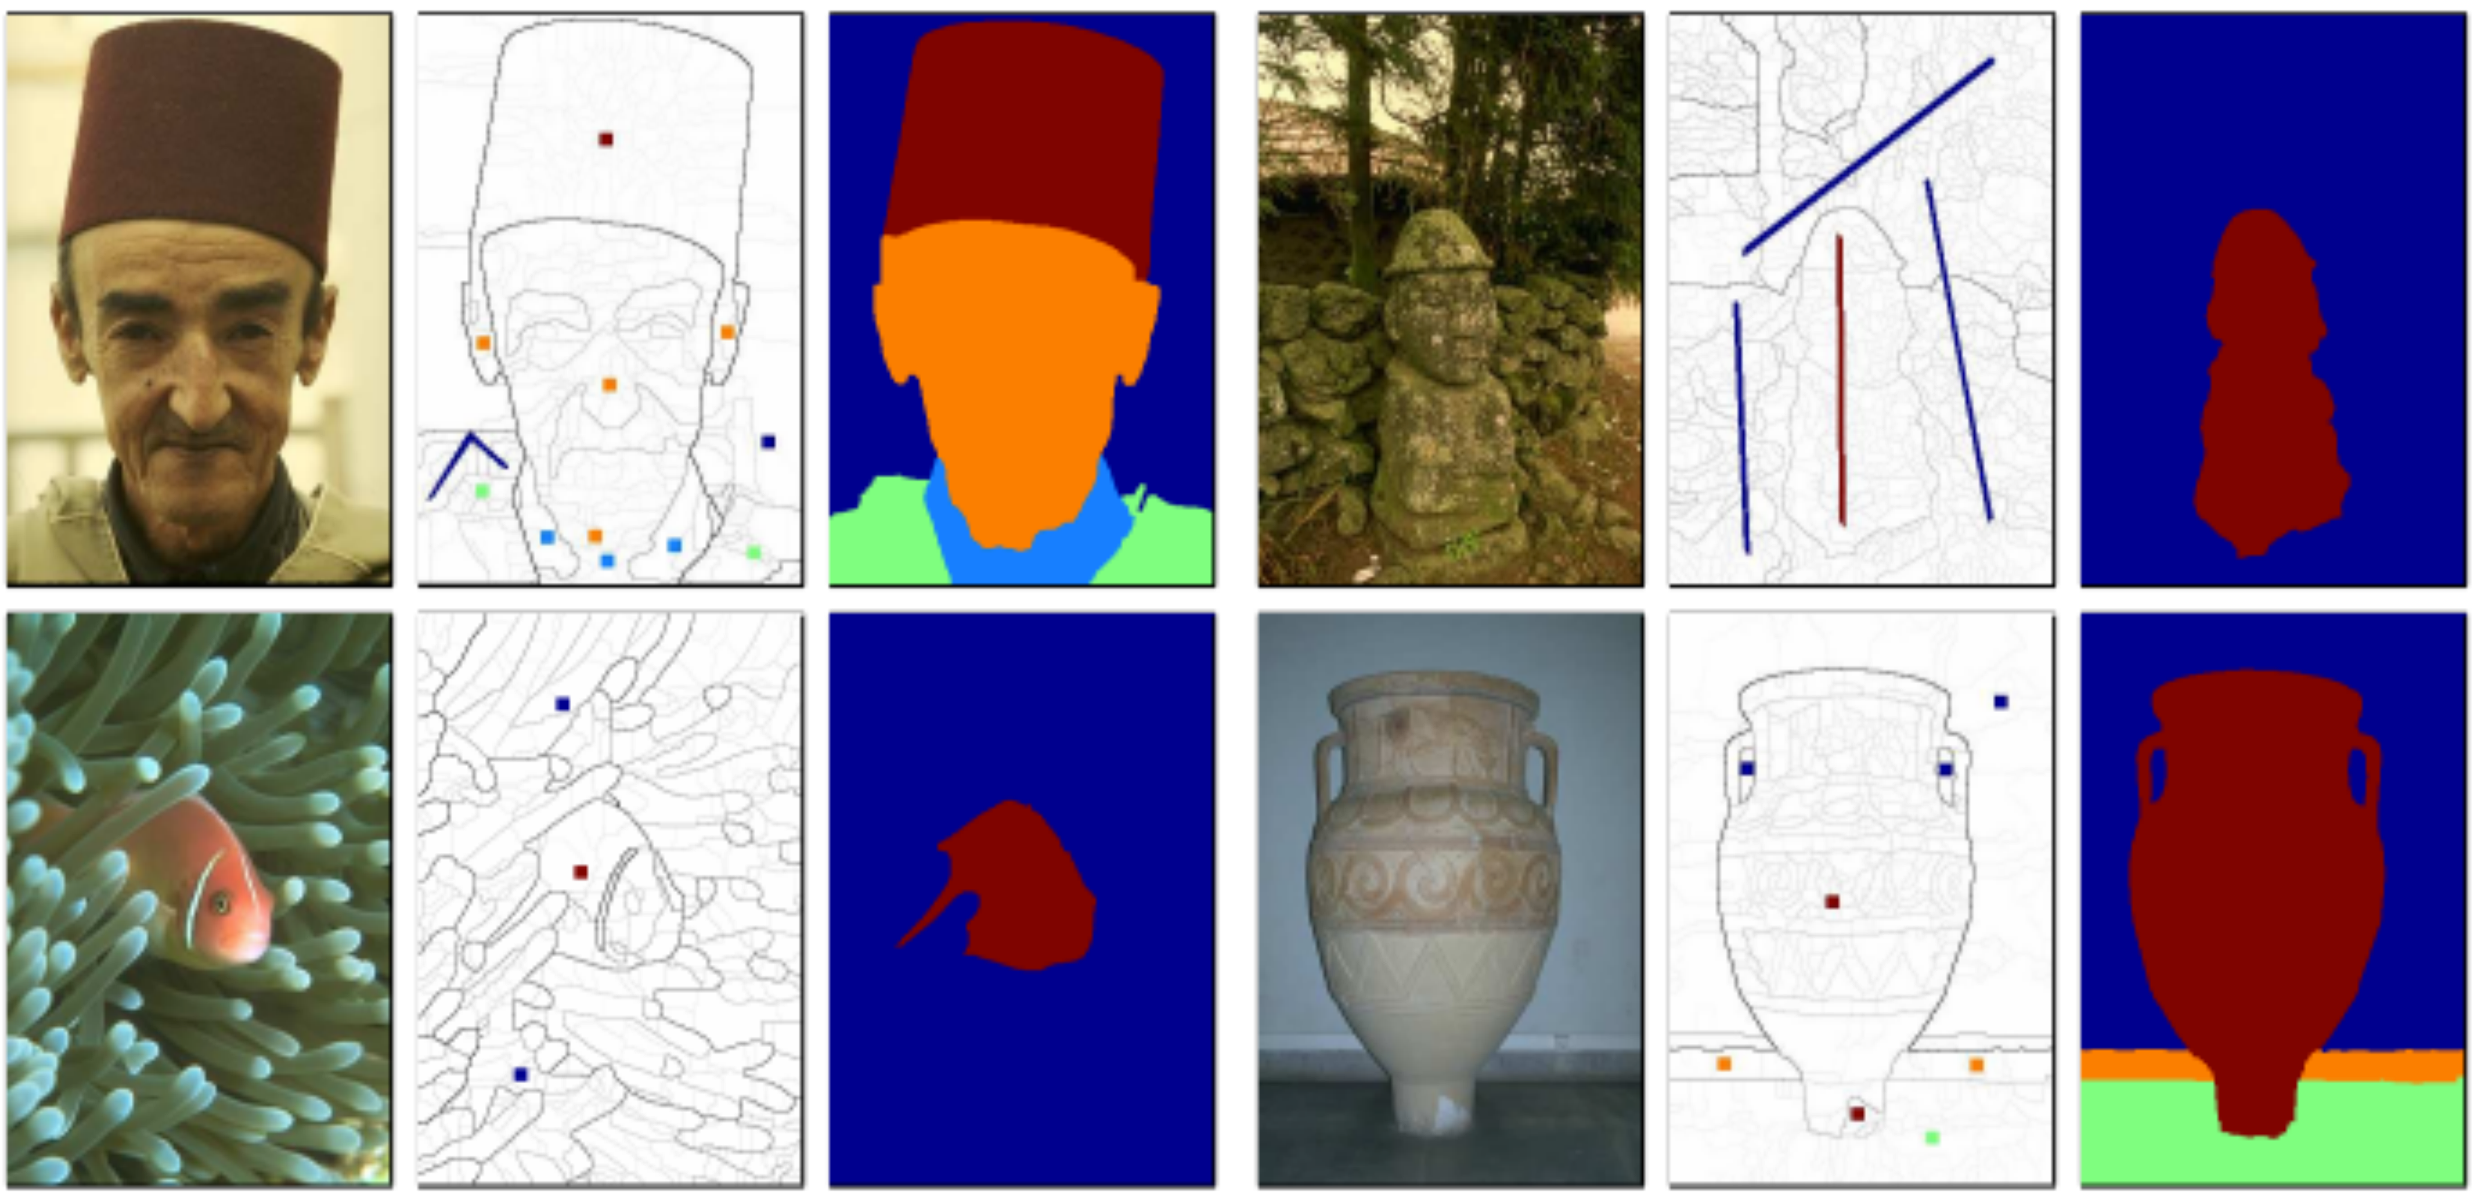

Multiscale Segmentation for Object Detection

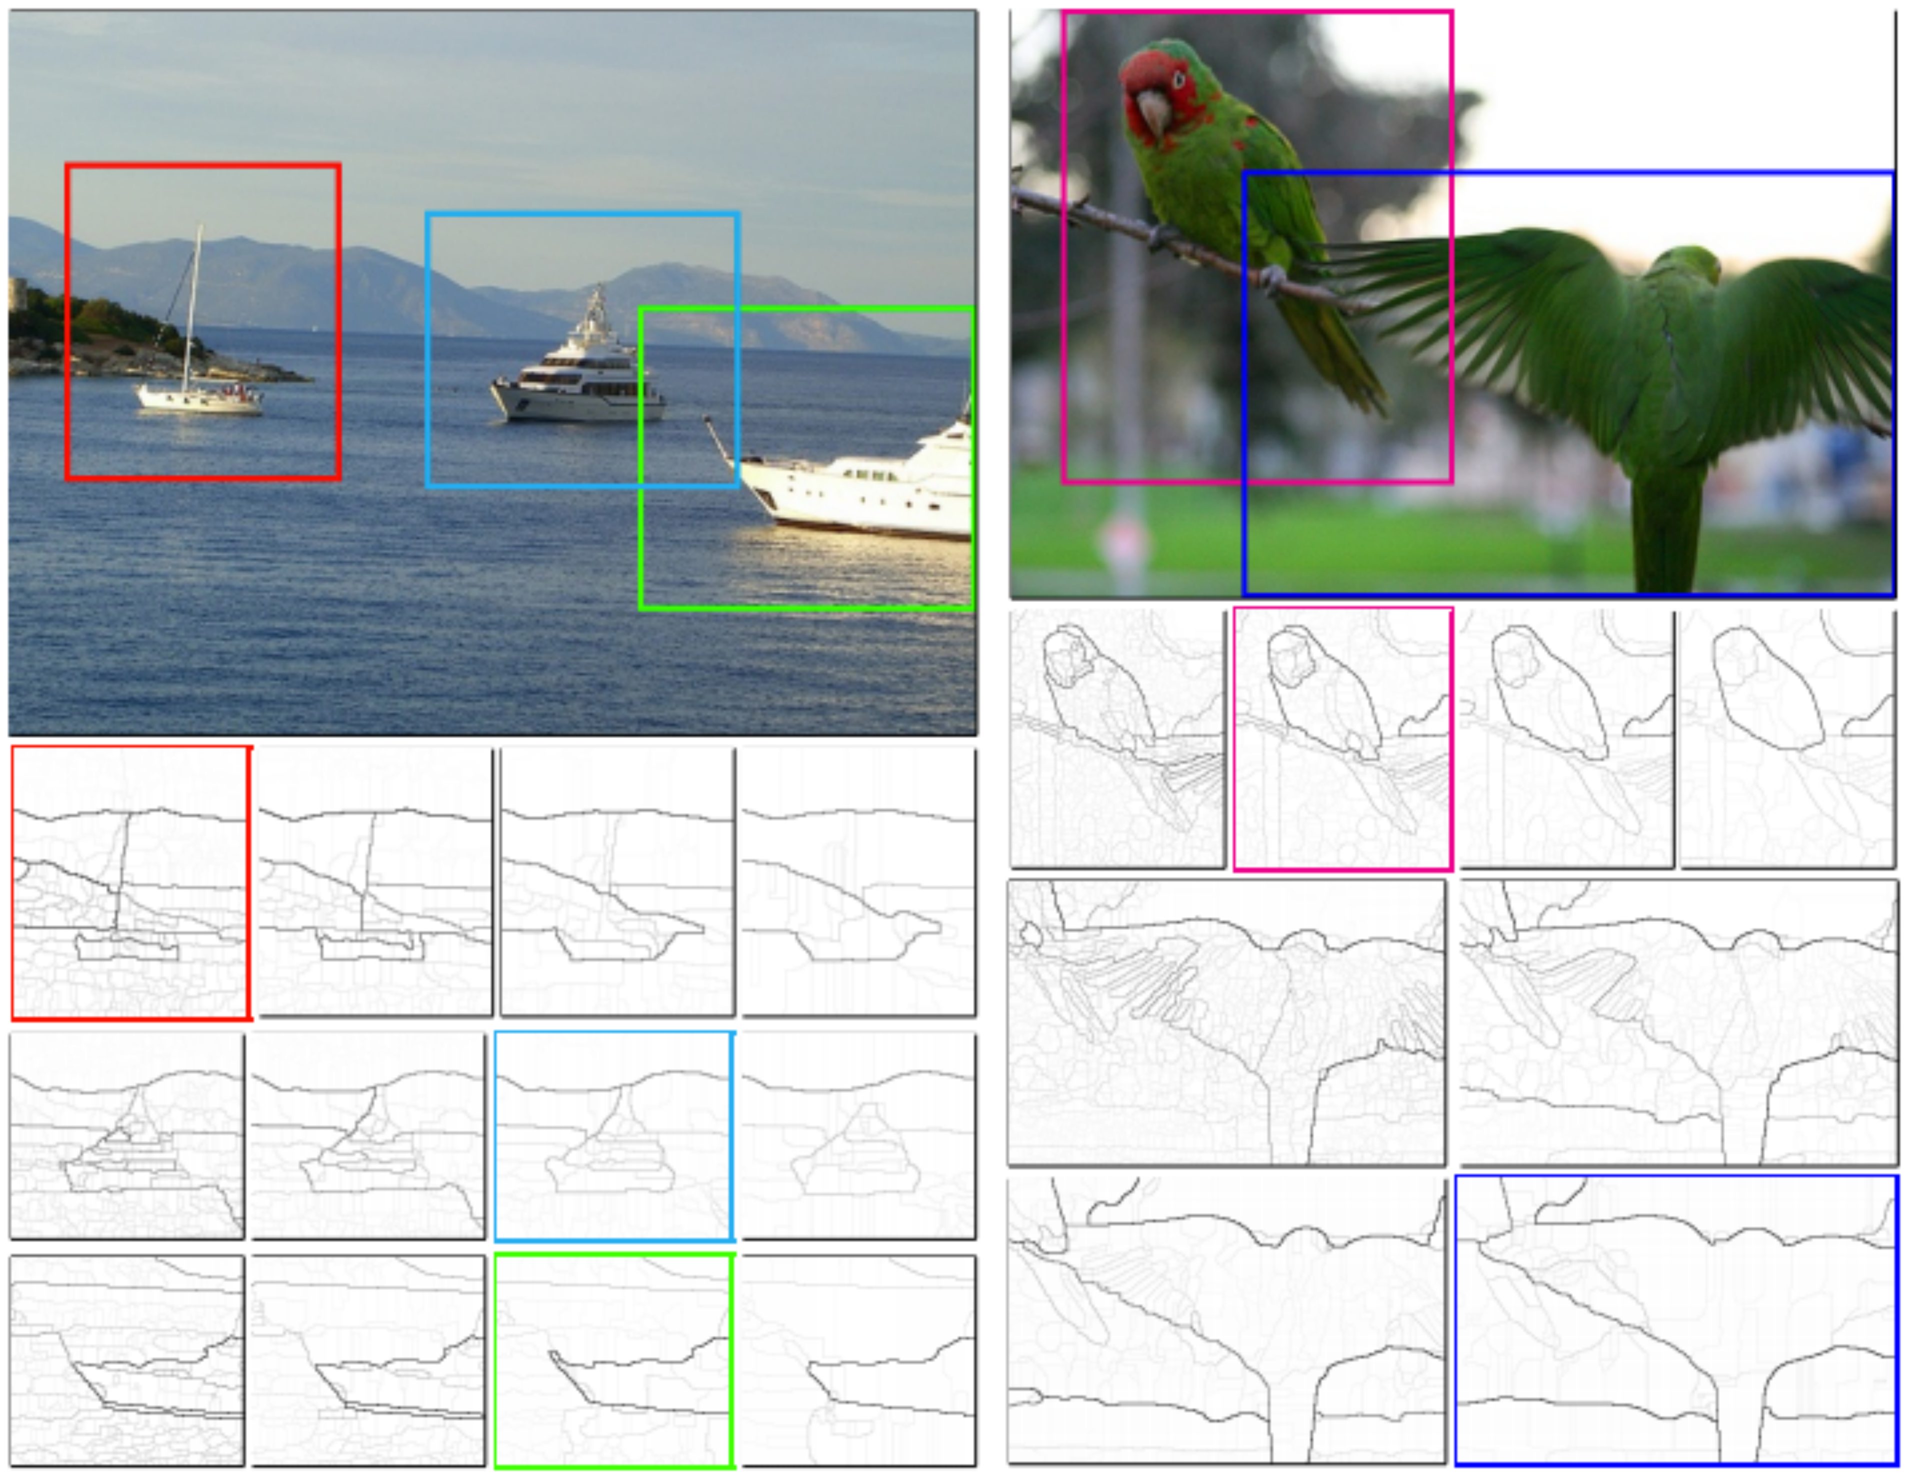

Pairwise comparison of Segmentation Algorithms