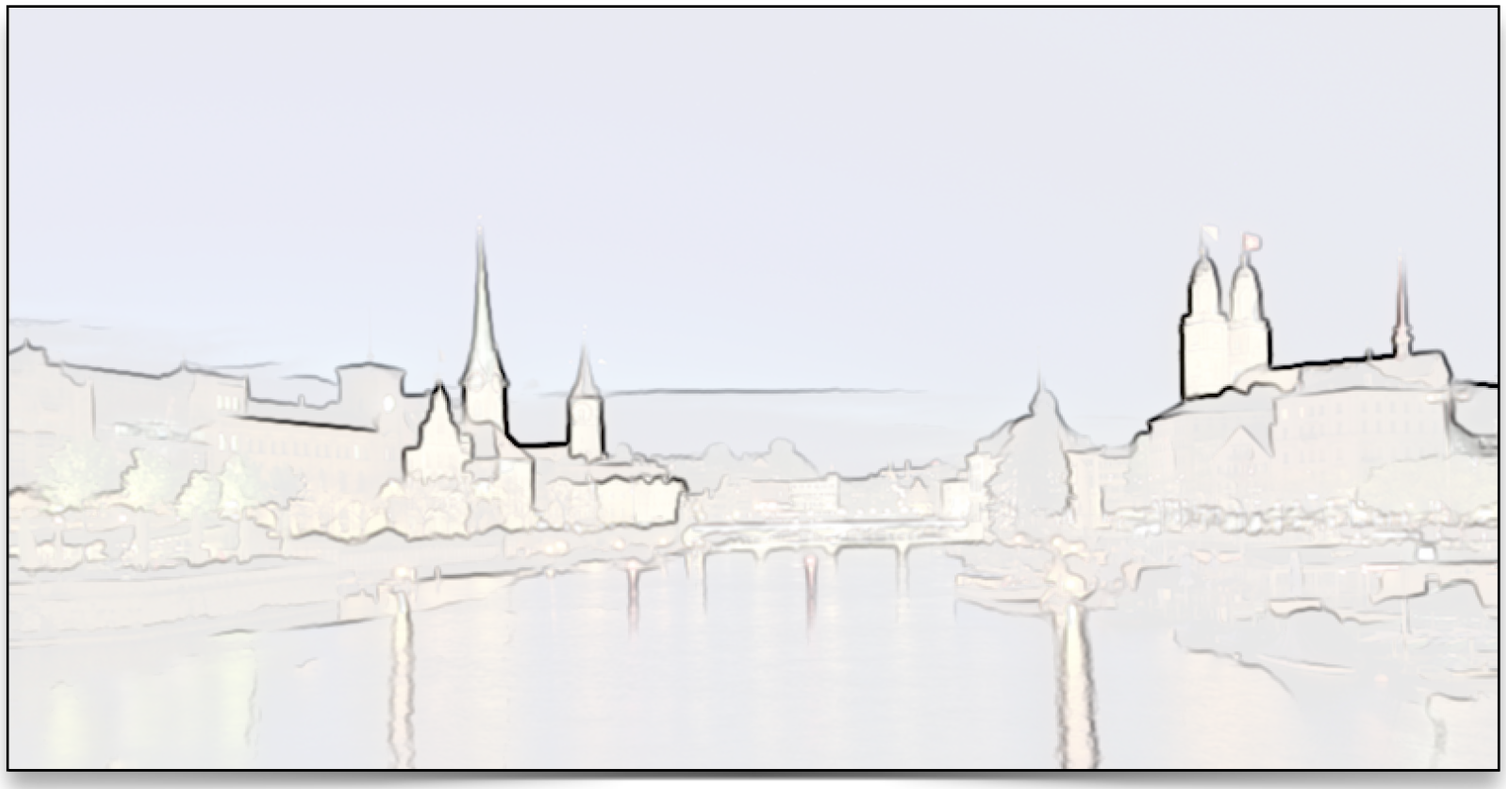

Crisp Boundary Detection Using Pointwise Mutual Information

Phillip Isola, Daniel Zoran, Dilip Krishnan, Edward Adelson
MIT

Original image

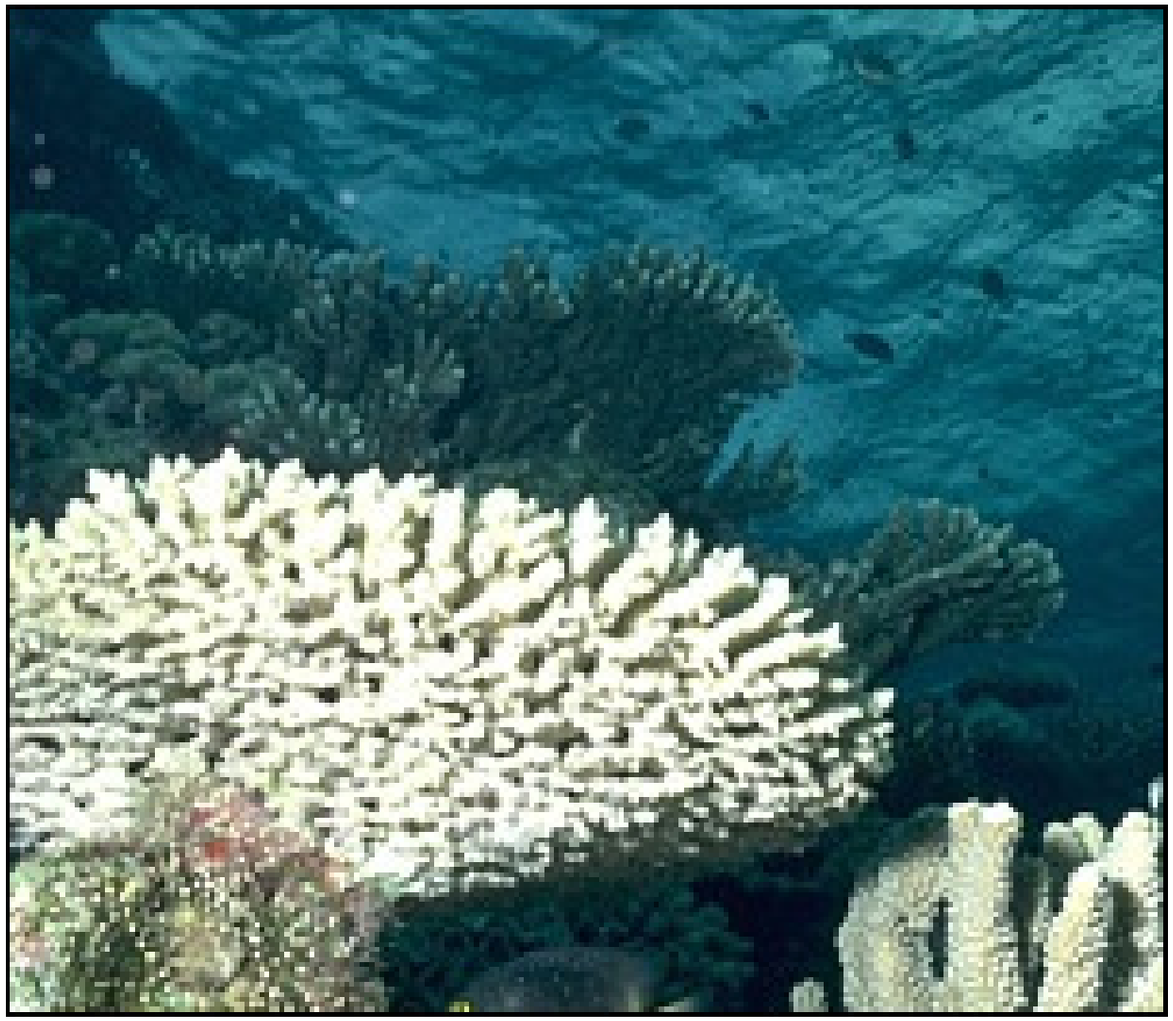

Human drawn contours
(Martin et al. 2004)

Goal: Find boundaries between image regions in a way that mimics human performance.

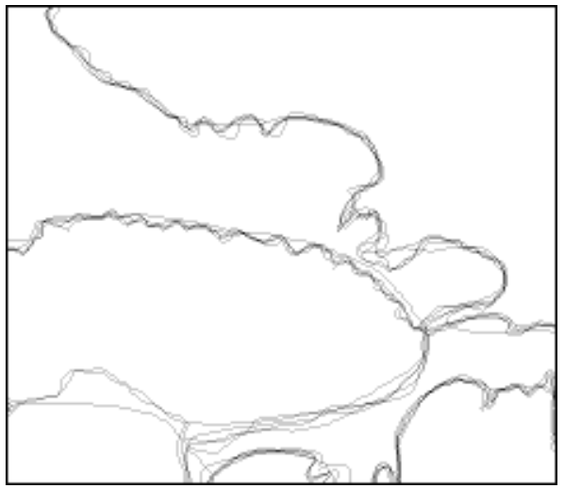

How do you find a boundary?

Sobel & Feldman
1968

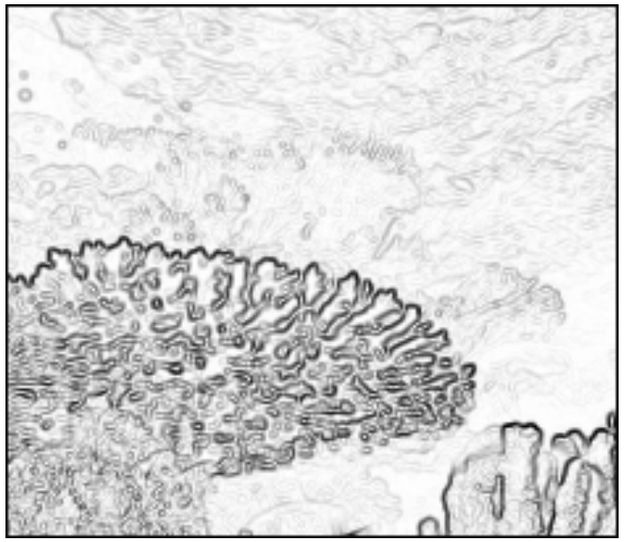

Find a change in luminance

Arbeláez et al.
2011 (gPb)

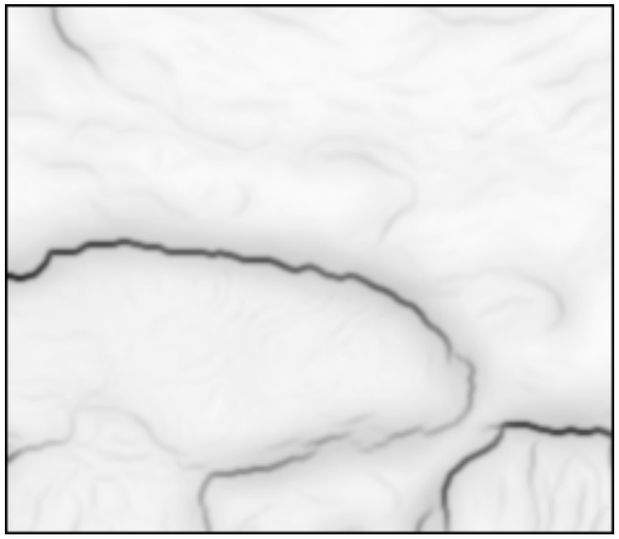

Find a change color/texture

Dollár & Zitnick
2014 (Structured Edges)

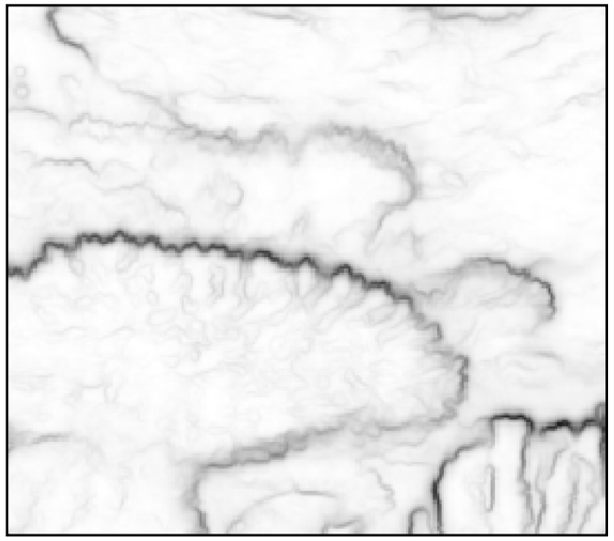

Recognize familiar patches

Pointwise mutual information

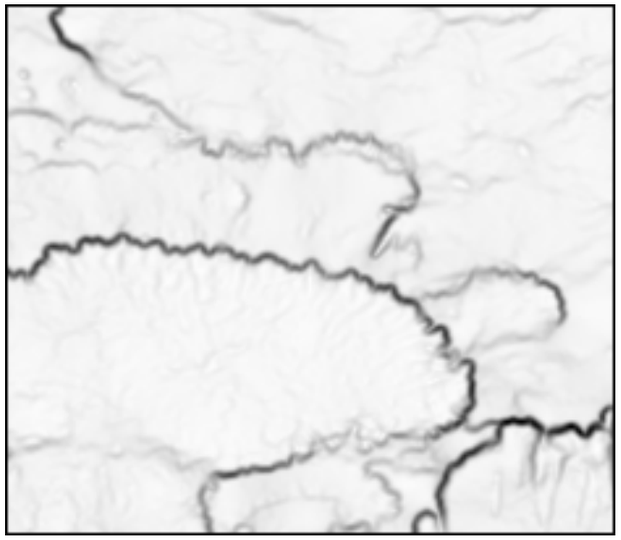

Look for a statistical dissociation

Key observation: *Pixels belonging to the same object have higher statistical association than pixels belonging to different objects.*

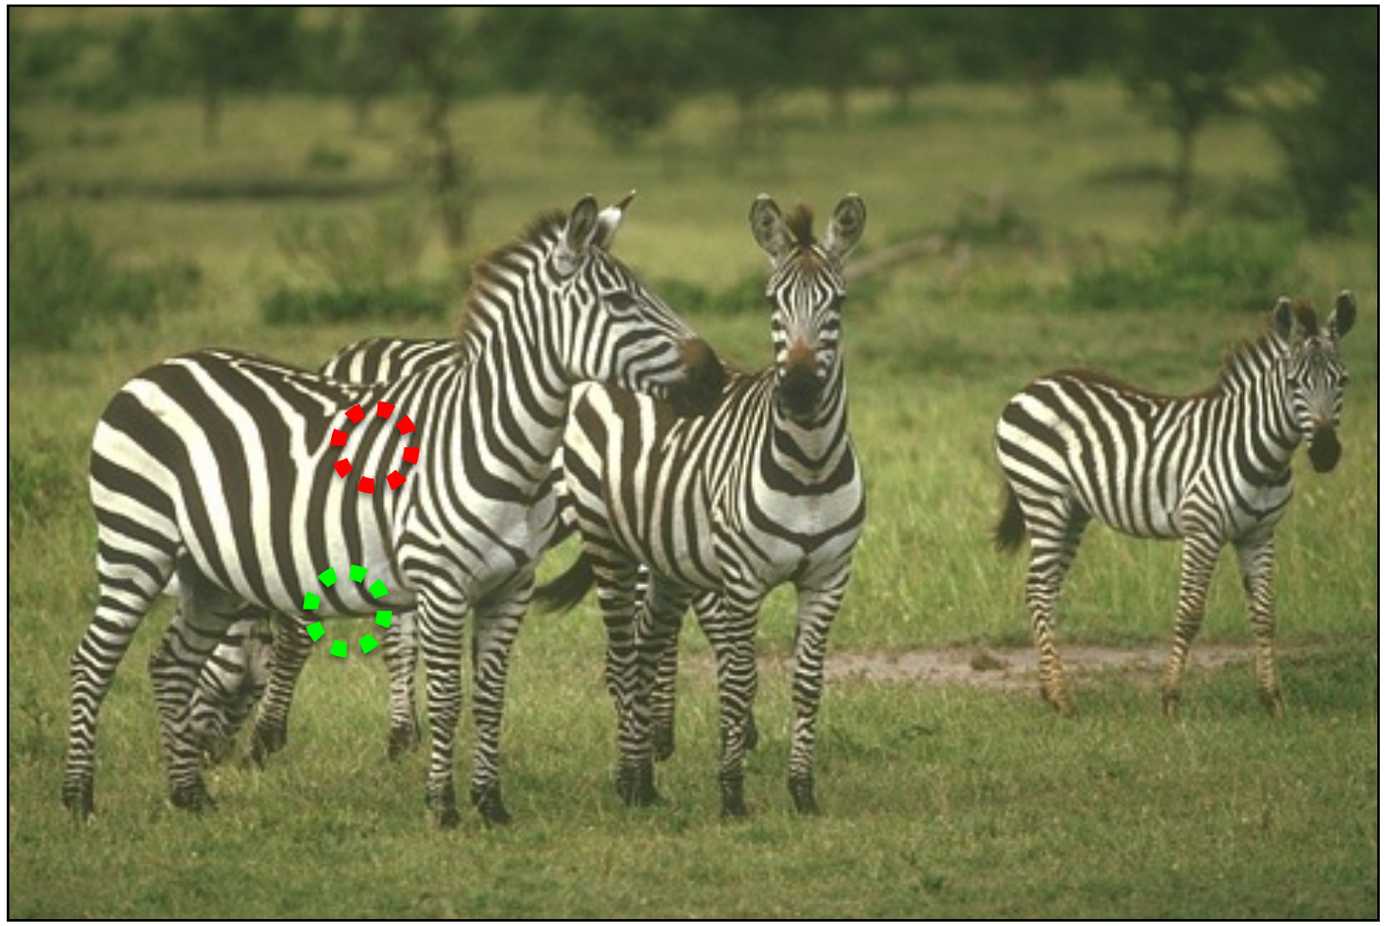

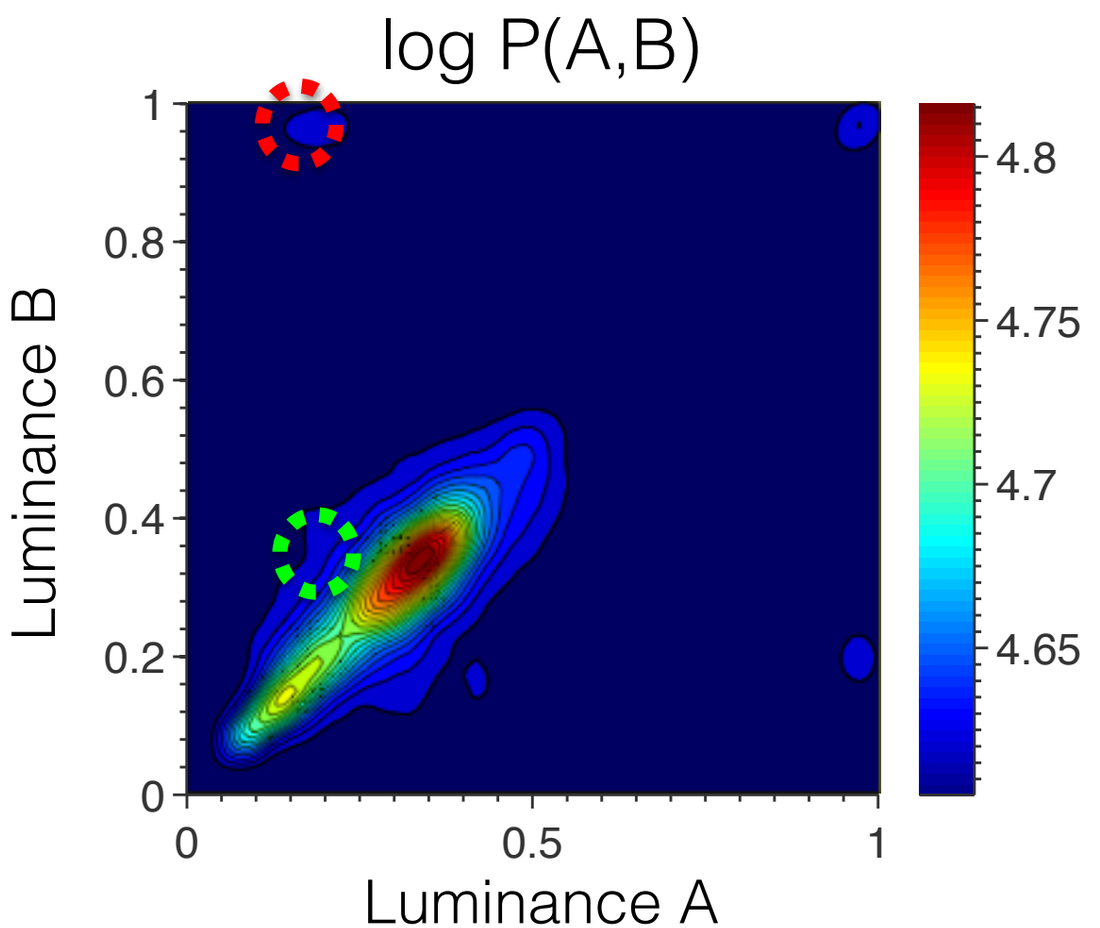

$P(A, B) =$ how often each color A occurs next to each color B *within this image*.

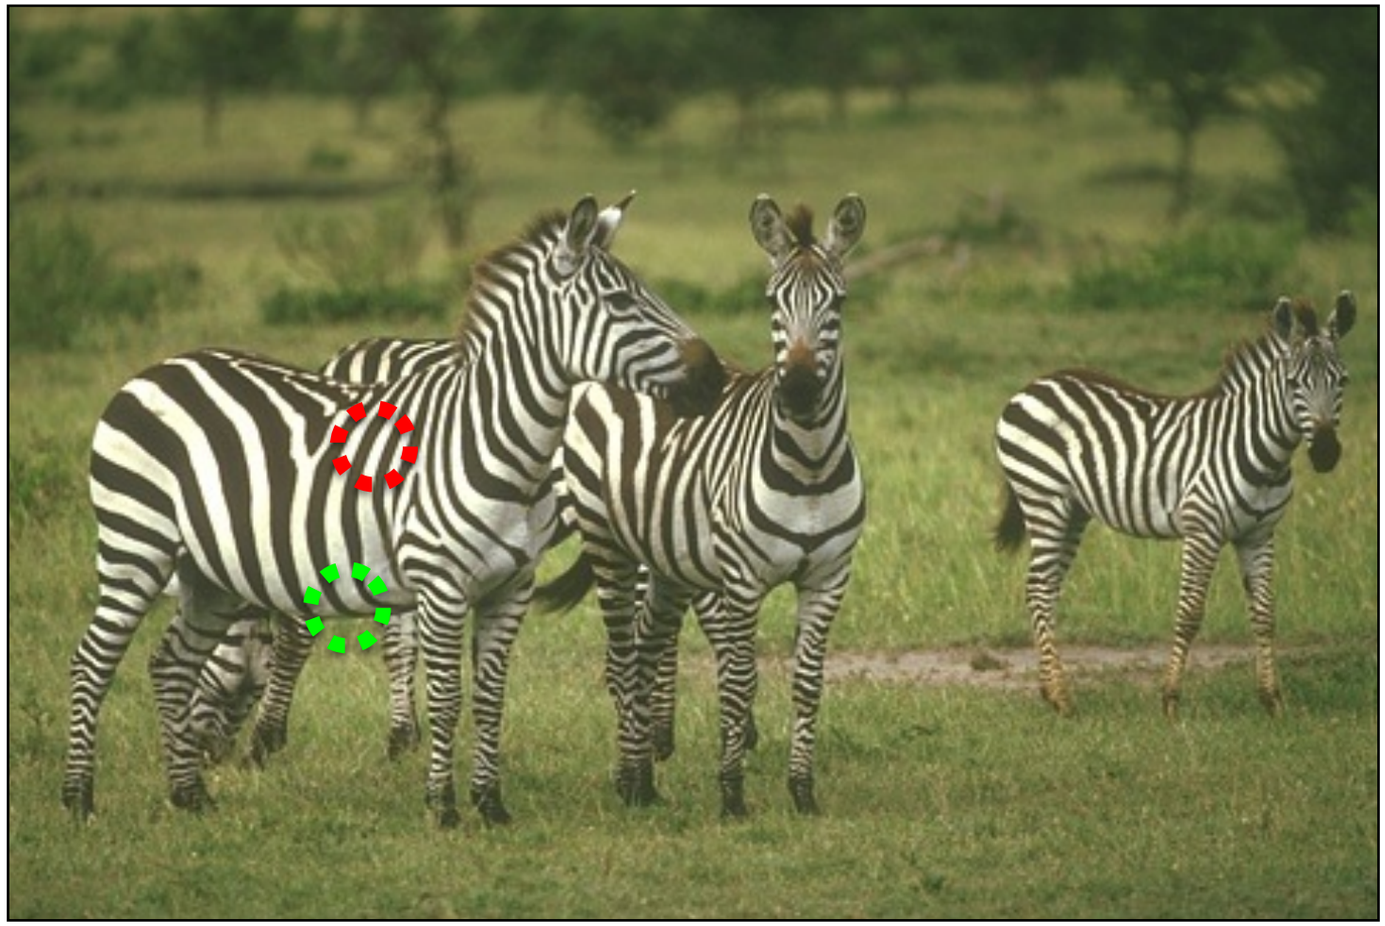

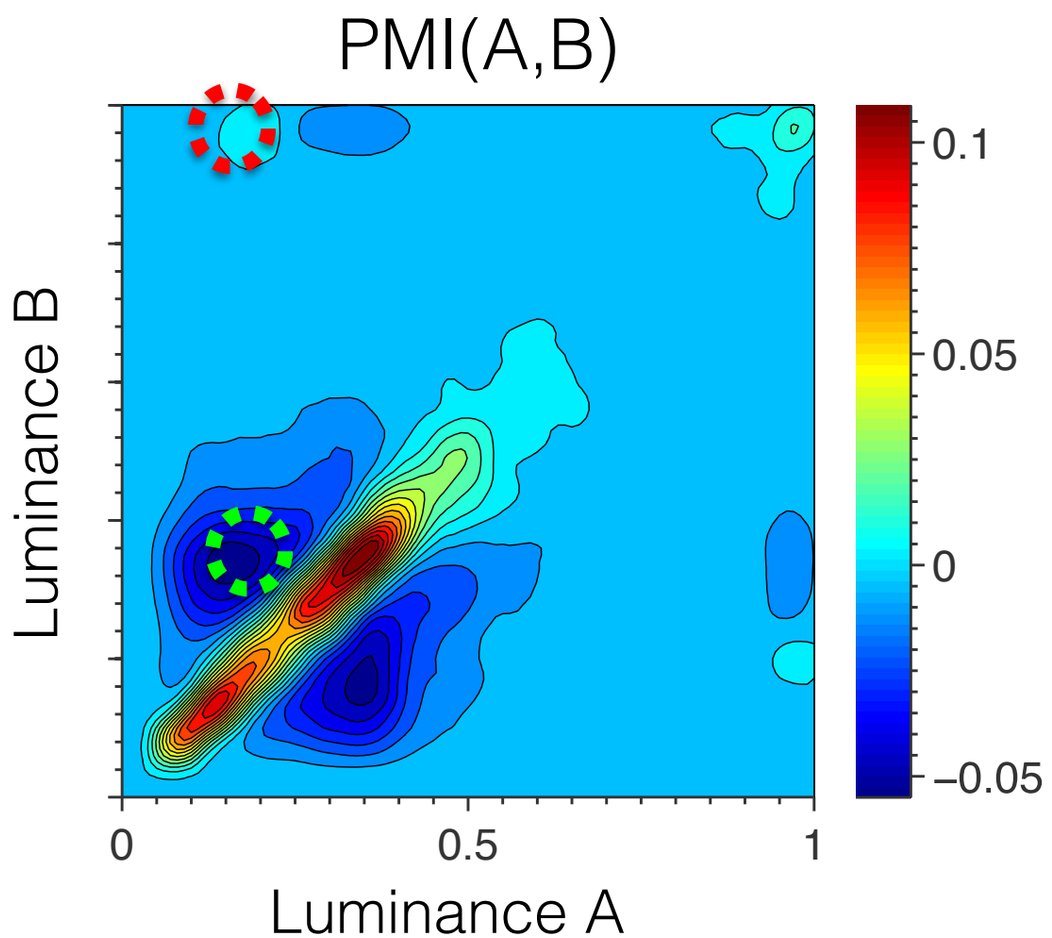

Pointwise mutual information (PMI)

$$\mathrm{PMI}_\rho(A,B) = \log \frac{P(A,B)^\rho}{P(A)P(B)}$$

Use PMI as affinity measure for affinity-based pixel grouping.

Is PMI informative about object boundaries?

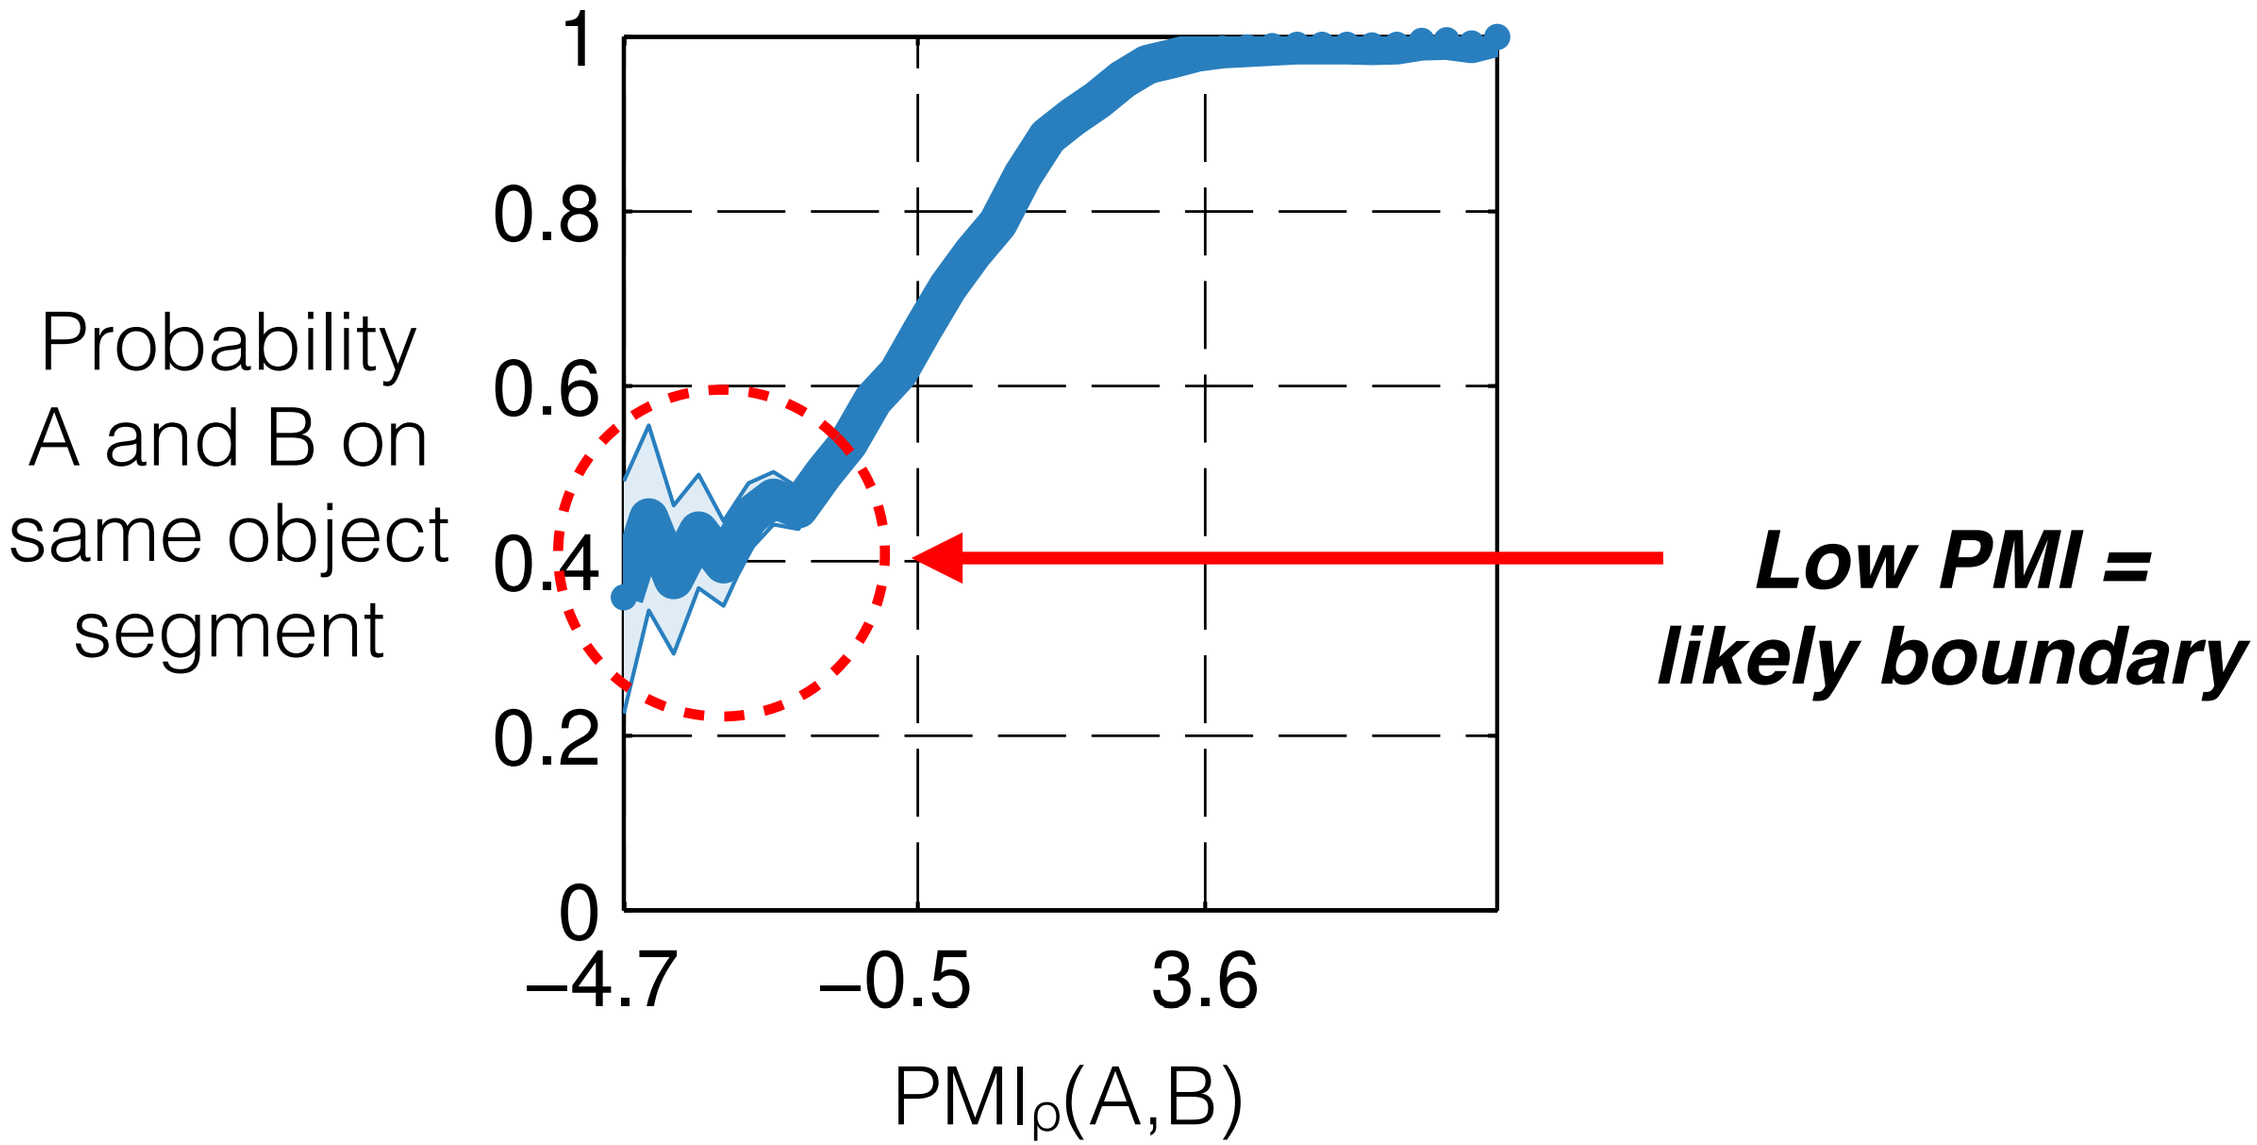

Is PMI informative about object boundaries?

Yes!

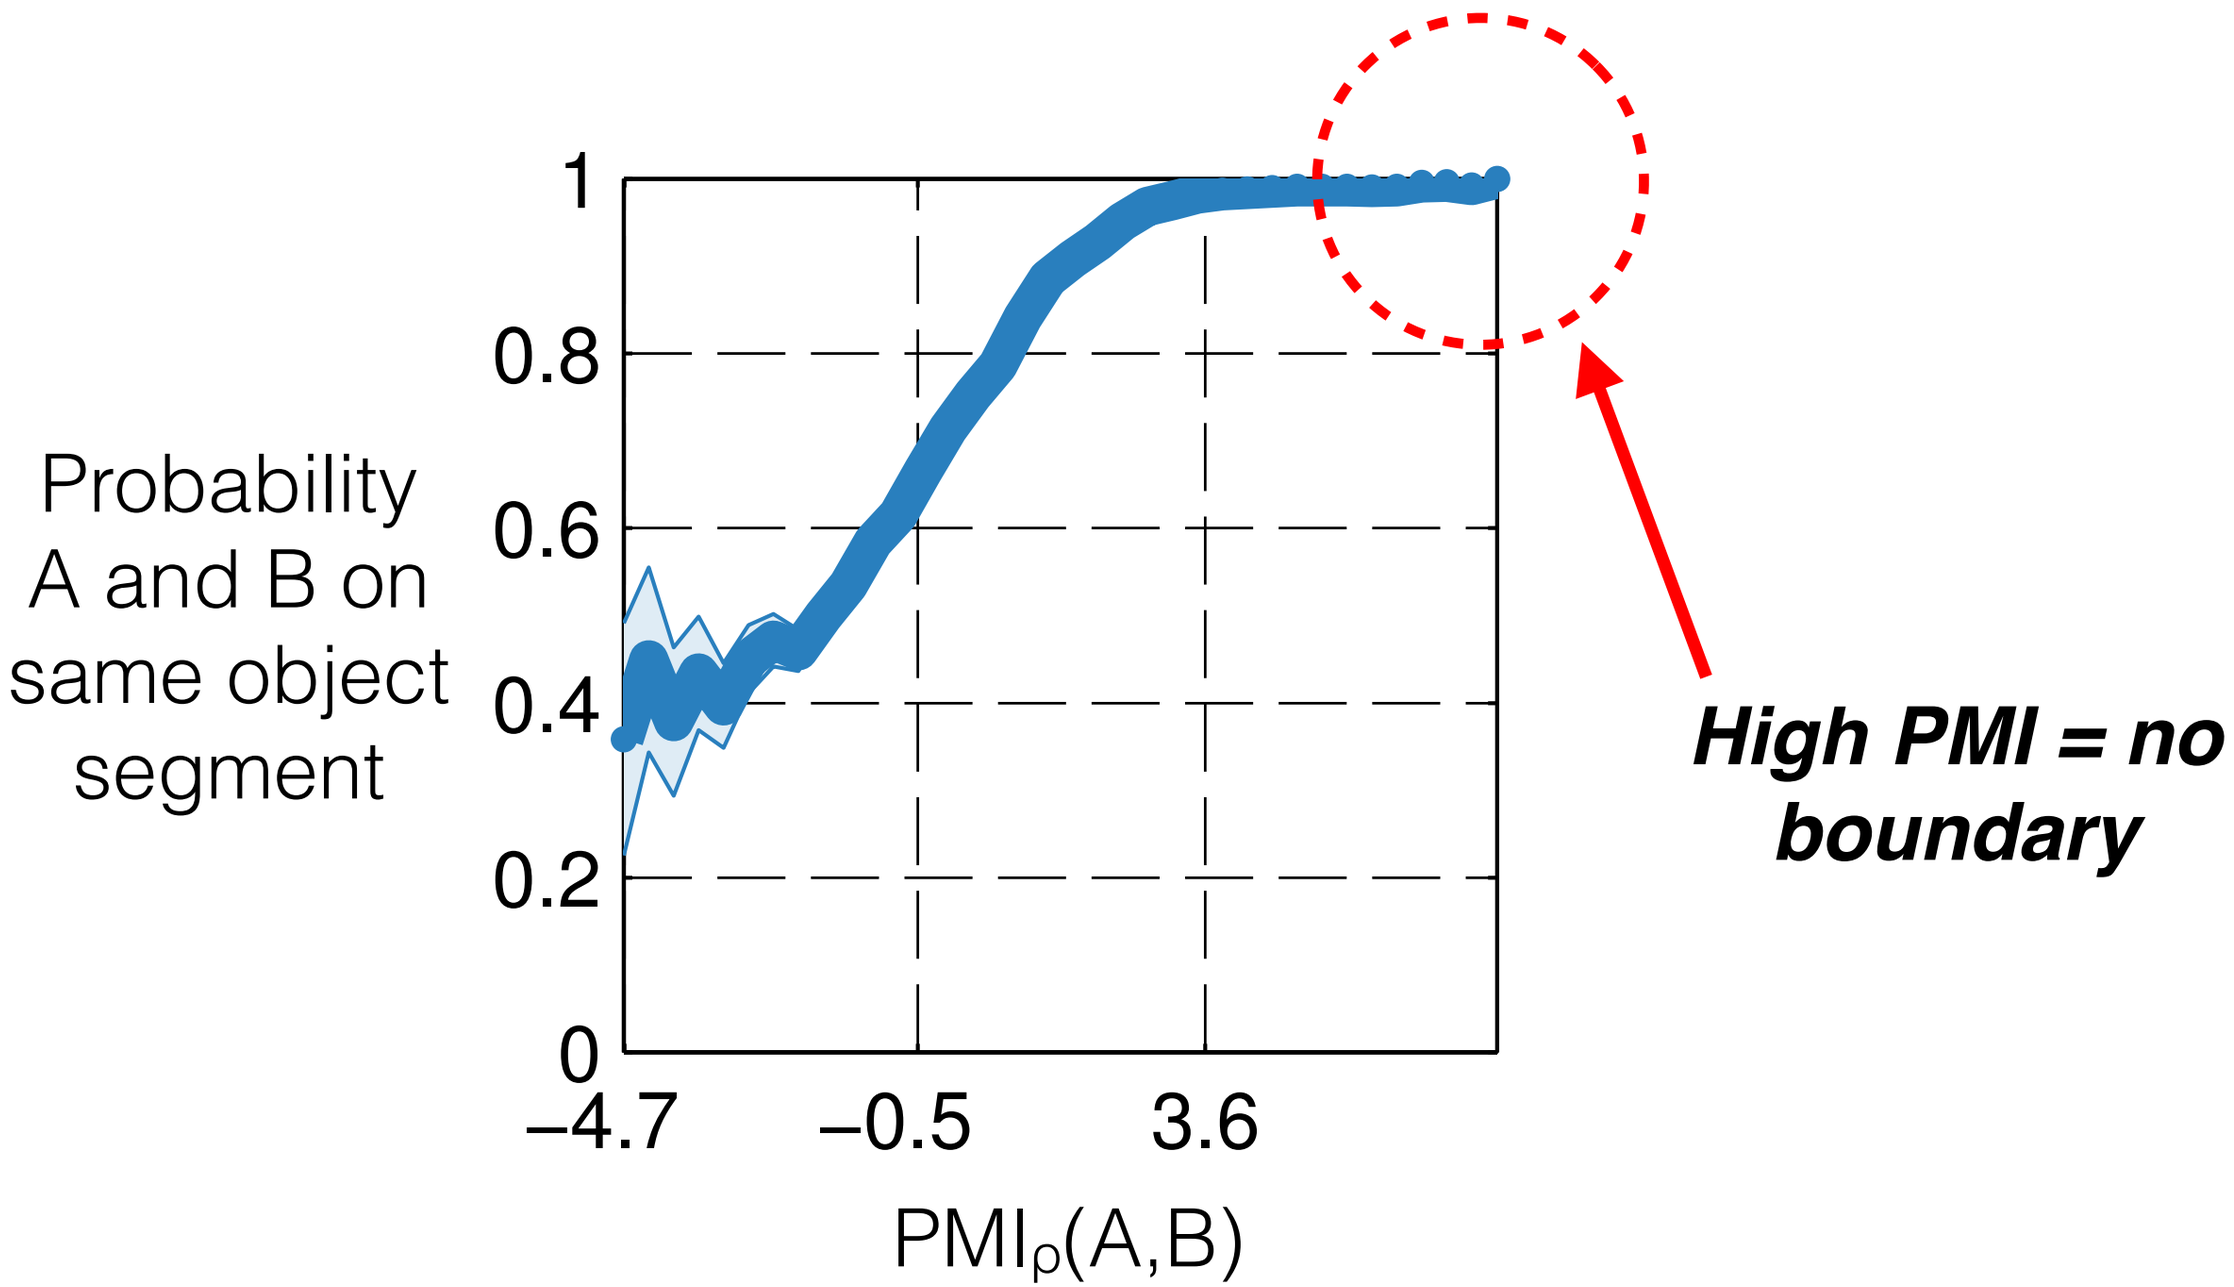

PMI leverages internal image statistics

Internal statistics

External statistics

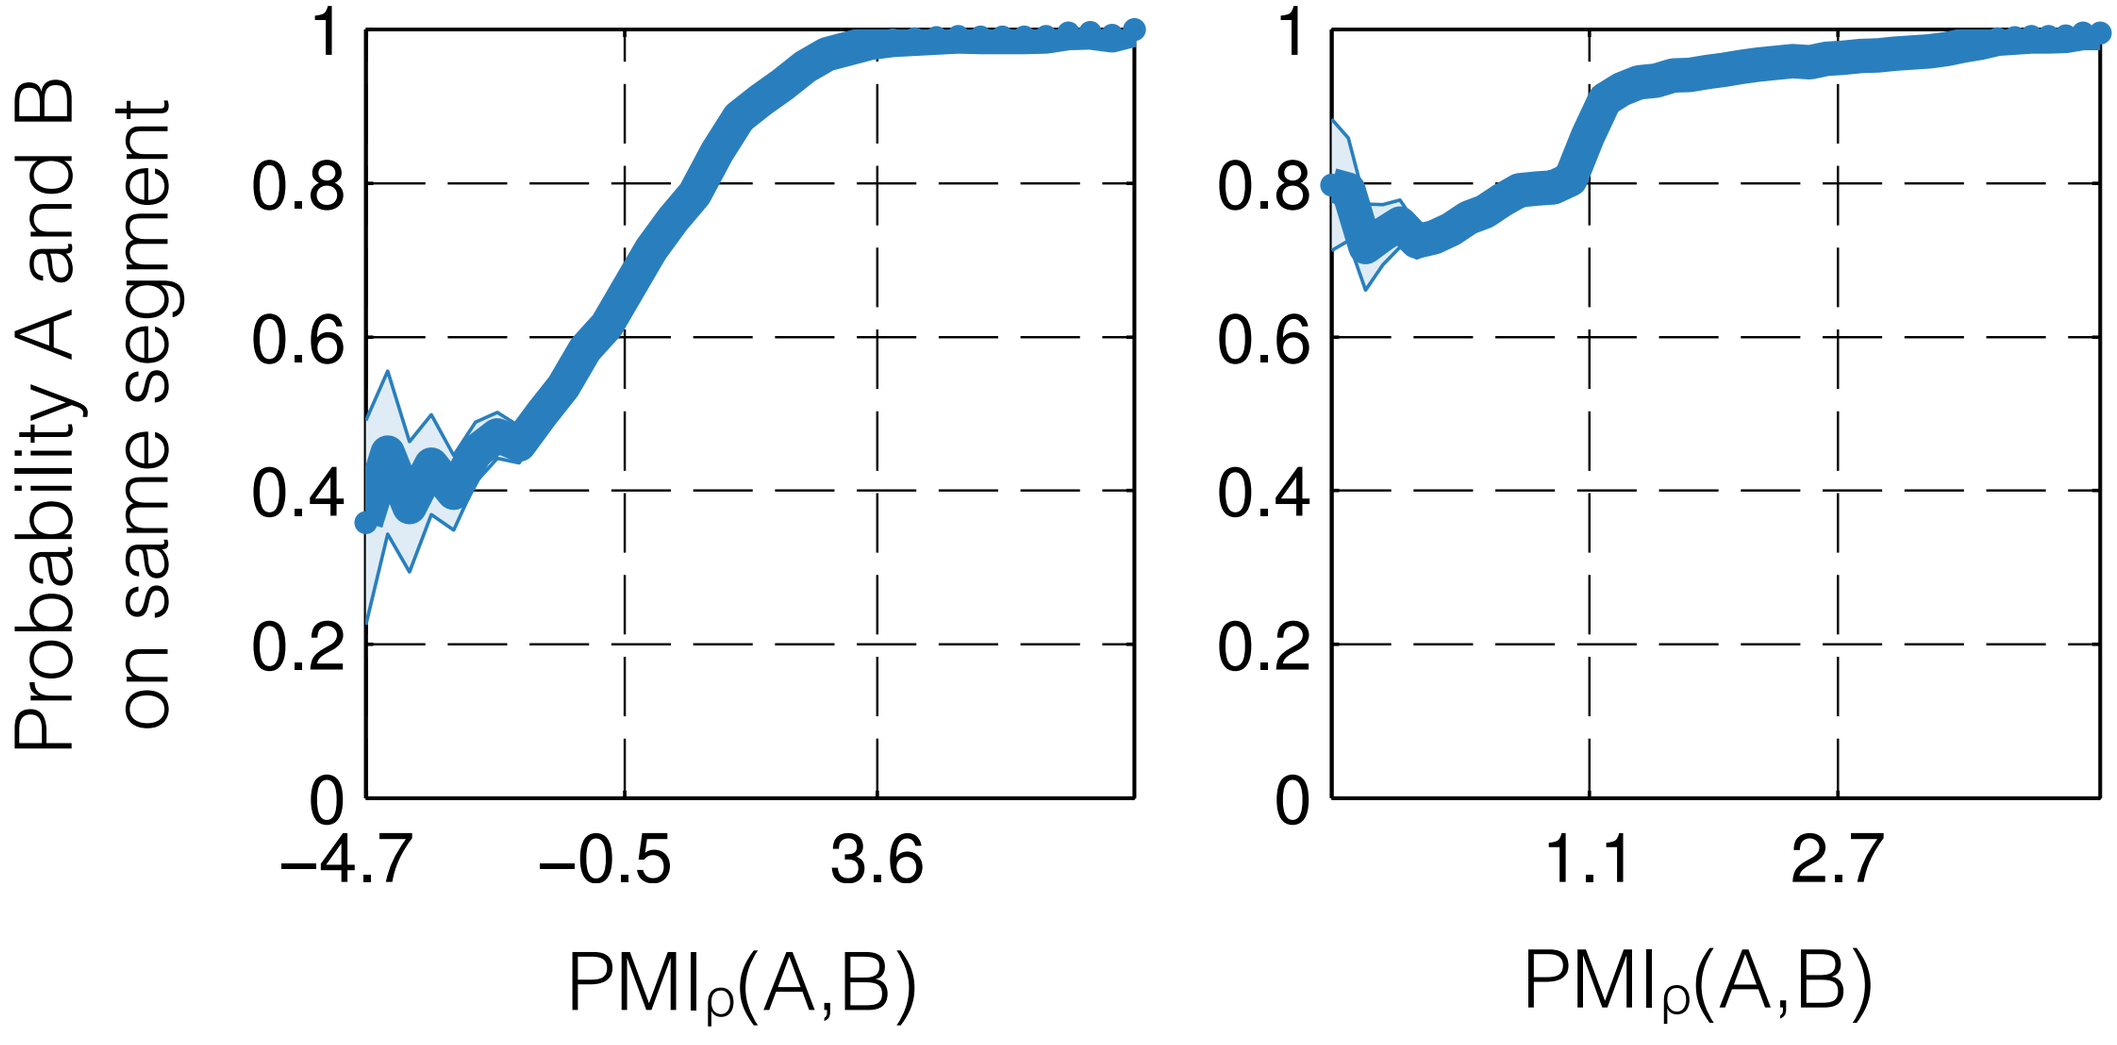

Algorithm

Algorithm outline

1. Get affinity between pixel pairs using PMI — our contribution
2. Apply affinity-based boundary detection — standard techniques
   (Arbeláez et al. 2011, gPb)

Step 1: Estimate feature co-occurrence distribution $P(A, B)$

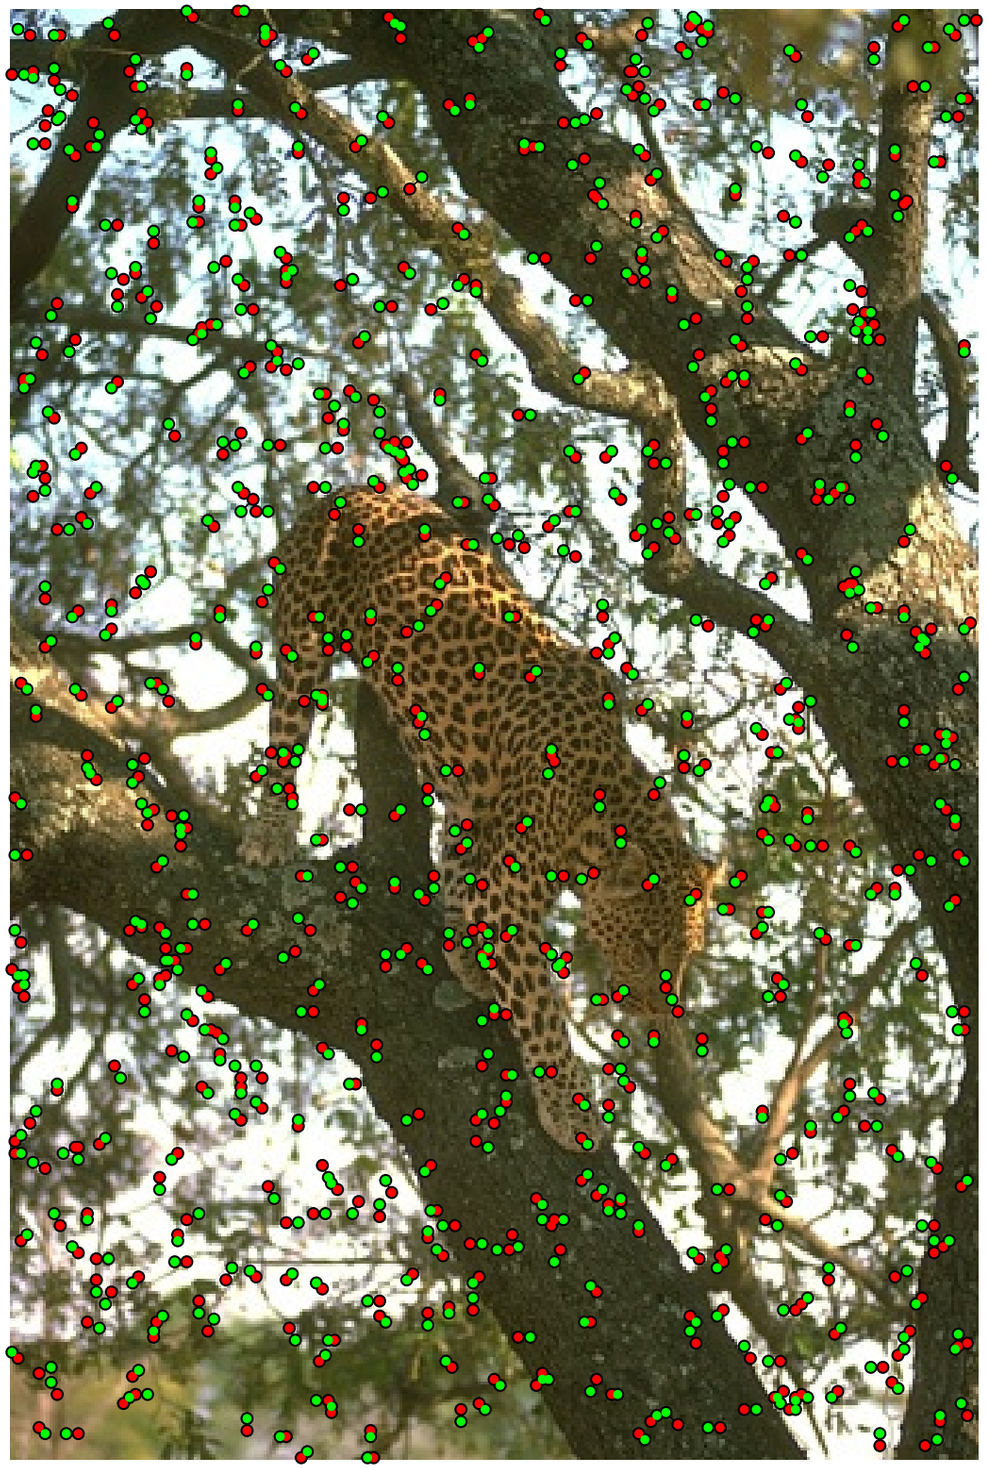

Samples

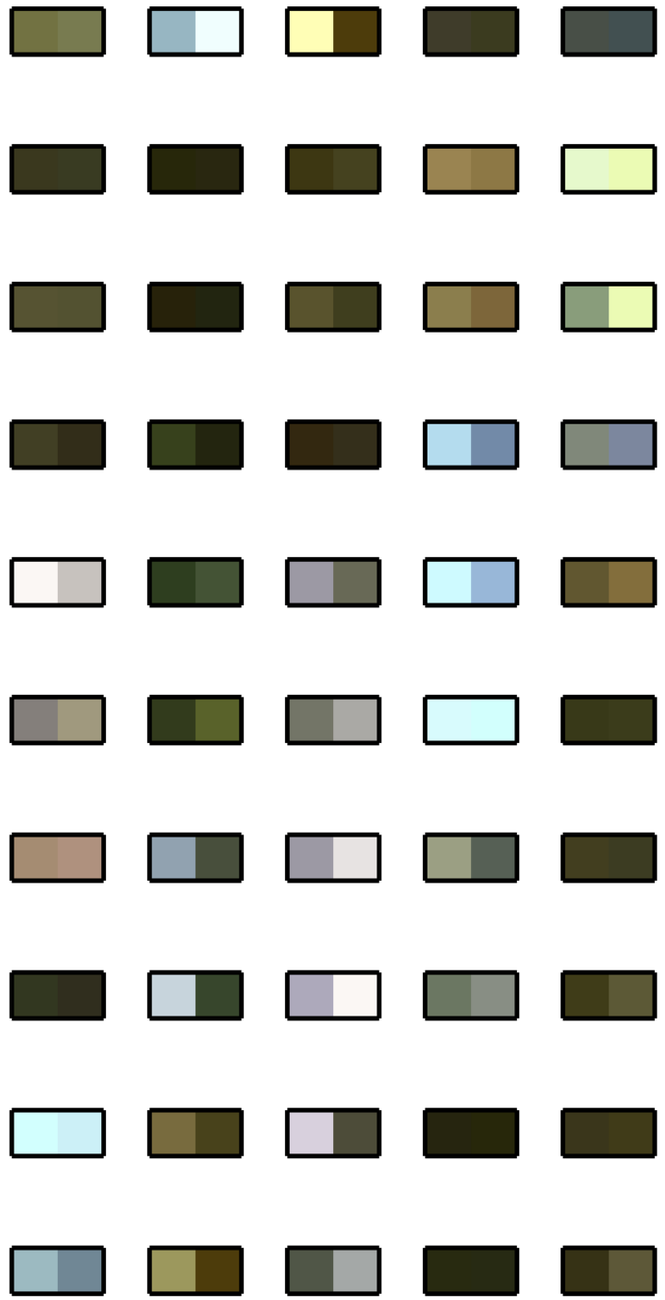

Step 2: Derive PMI(A,B) from feature co-occurrence distribution

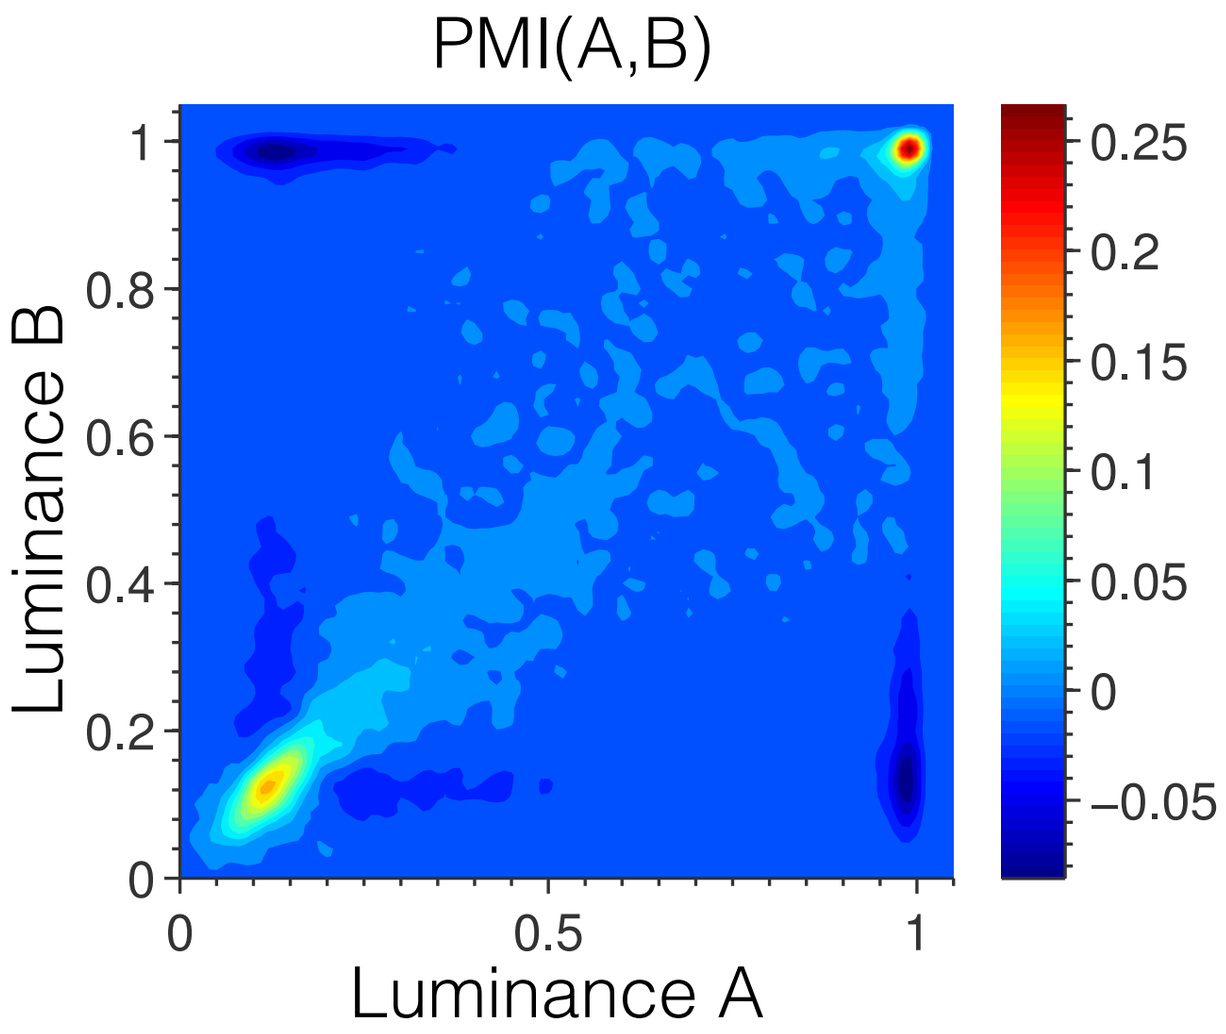

$$\mathrm{PMI}_\rho(A,B) = \log \frac{P(A,B)^\rho}{P(A)P(B)}$$

Step 3: Use PMI as affinity between each pair of nearby pixels

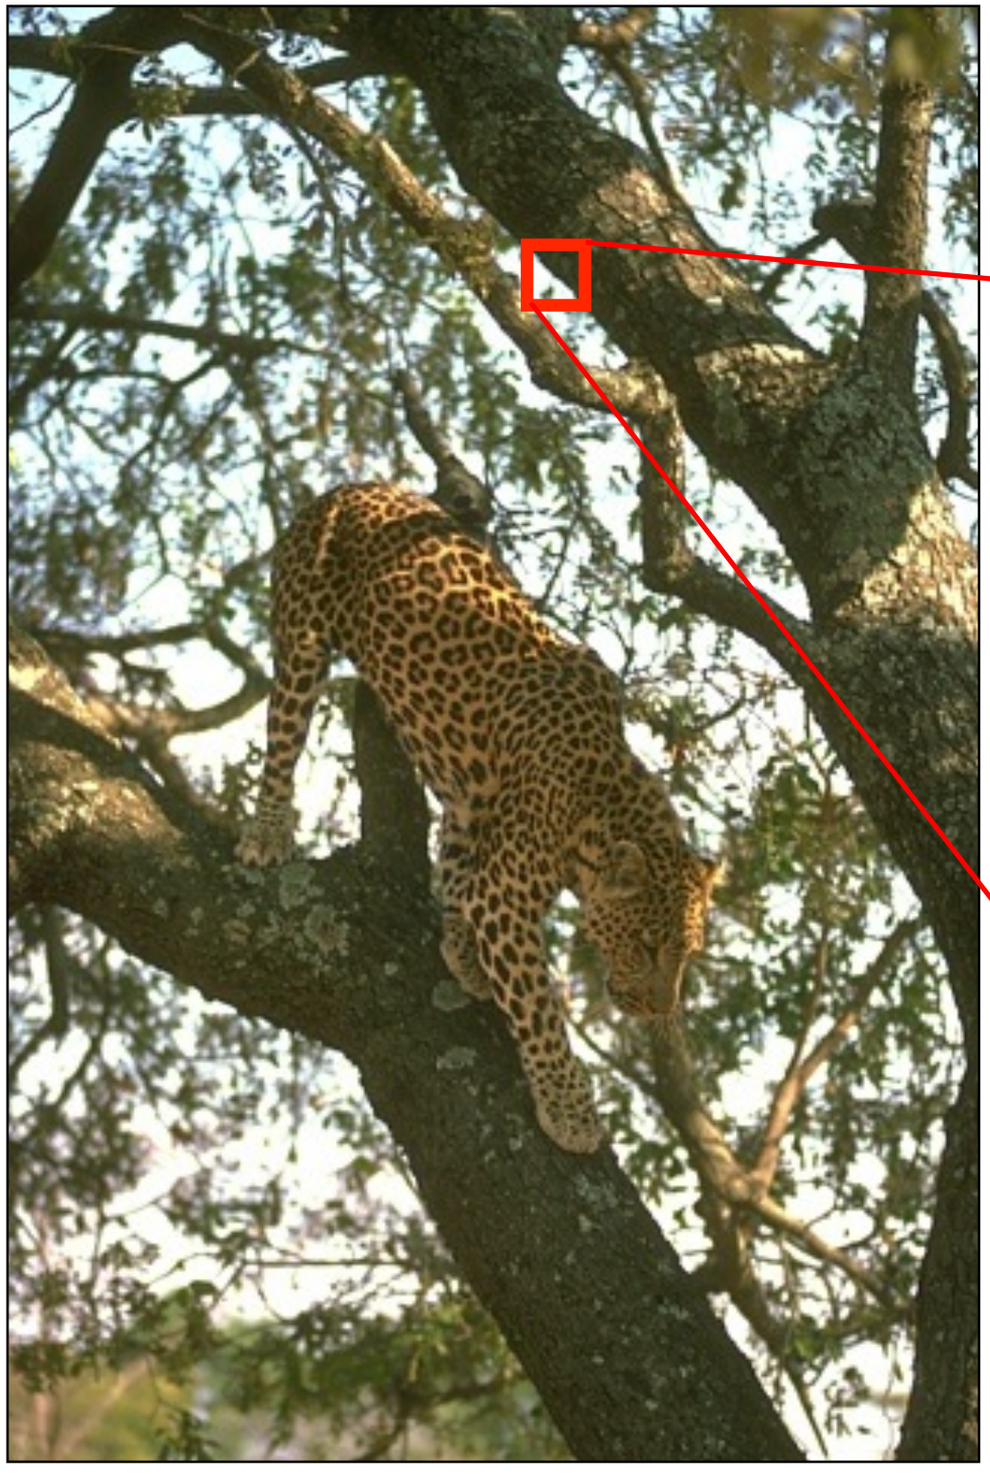

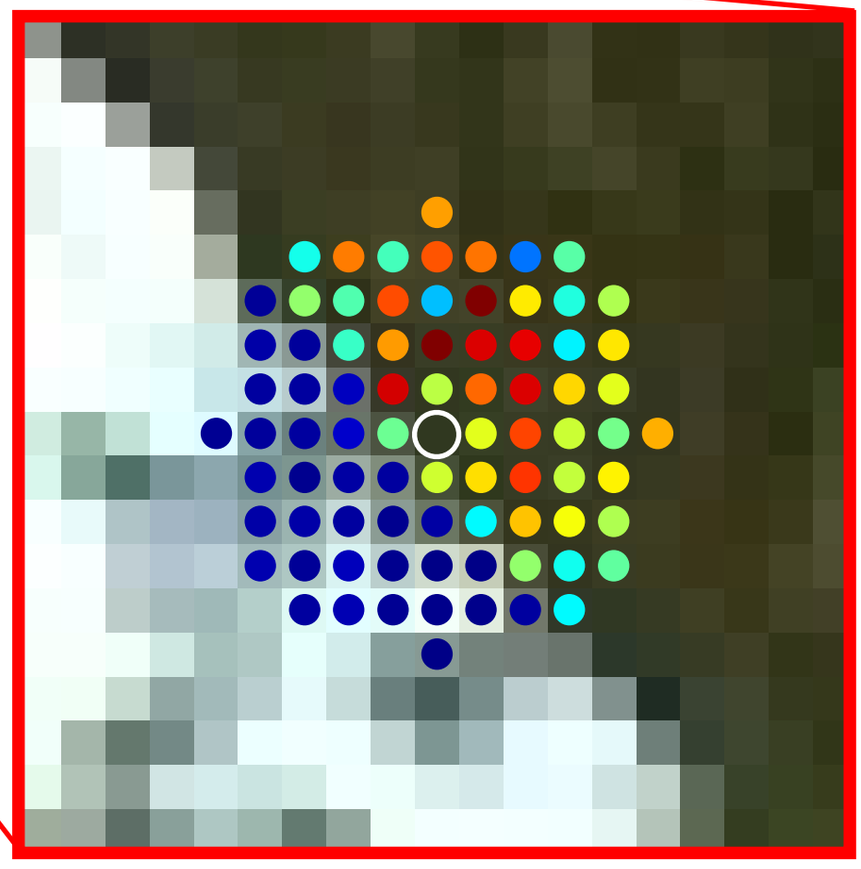

Step 3: Use PMI as affinity between each pair of nearby pixels

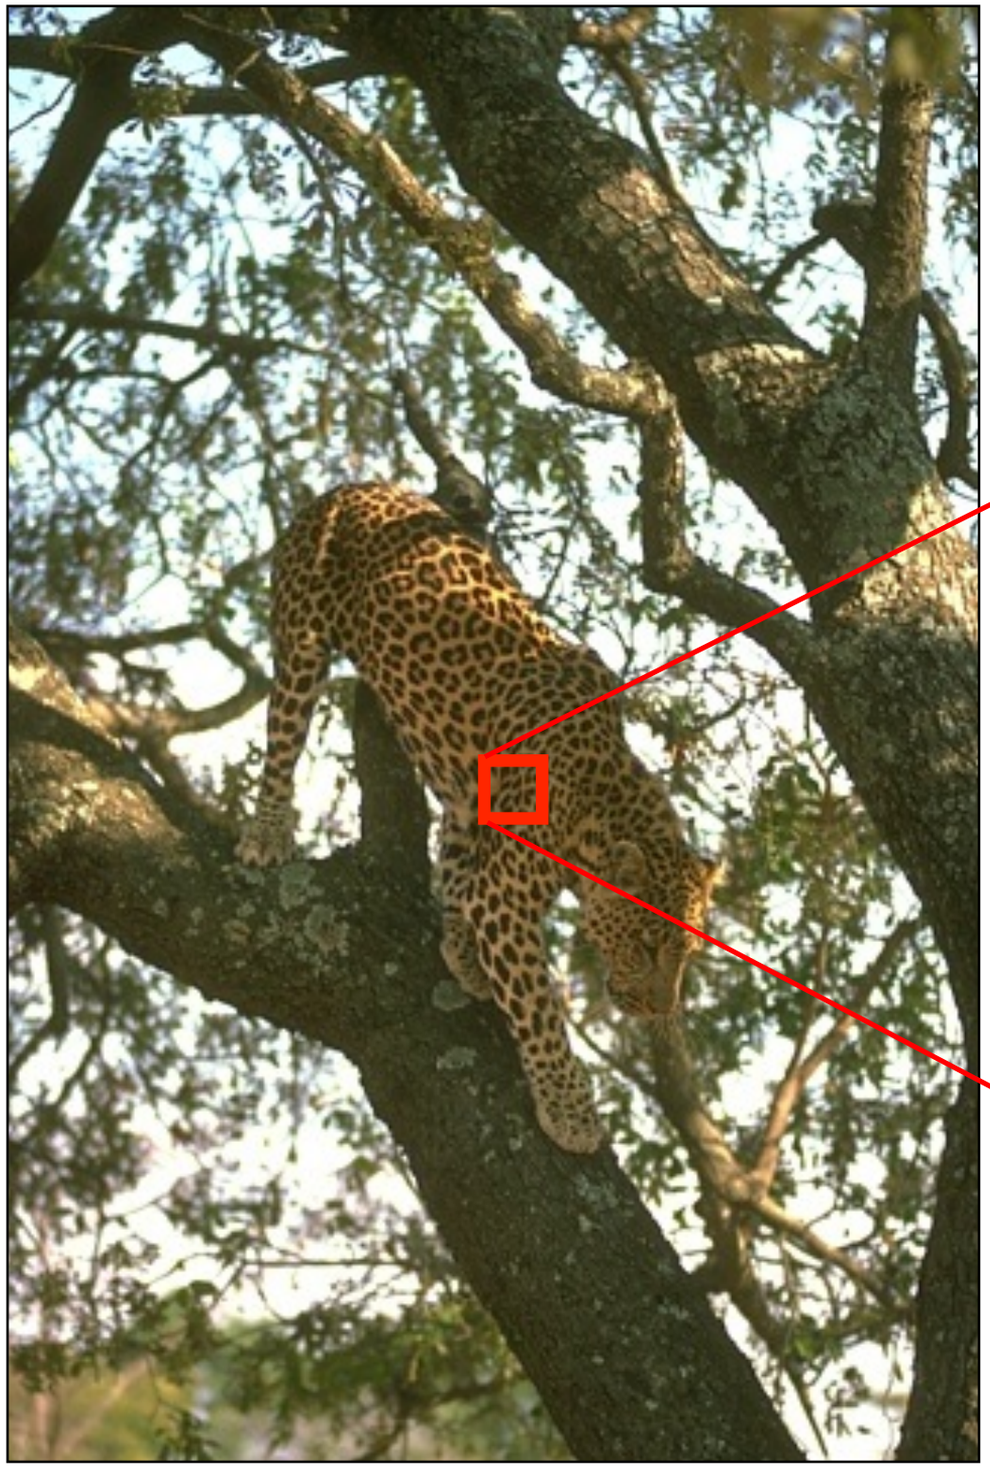

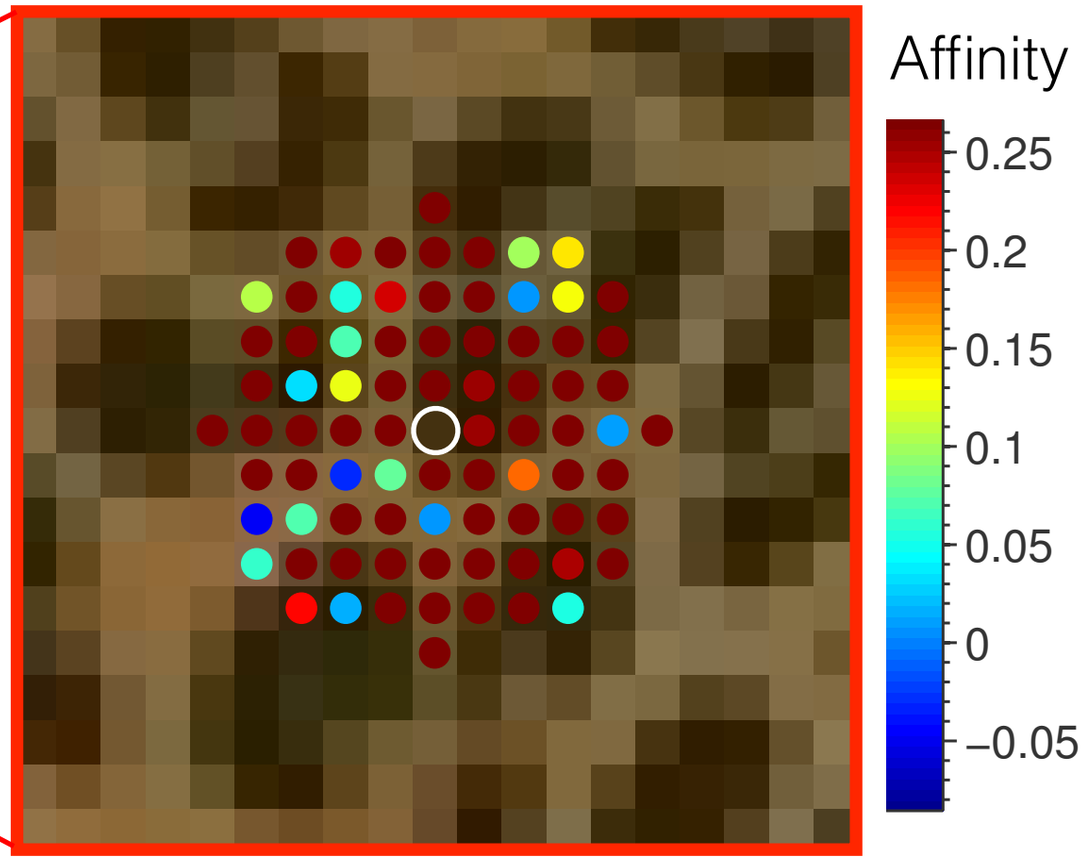

Step 4: Group pixels based on affinity (spectral clustering)

Input

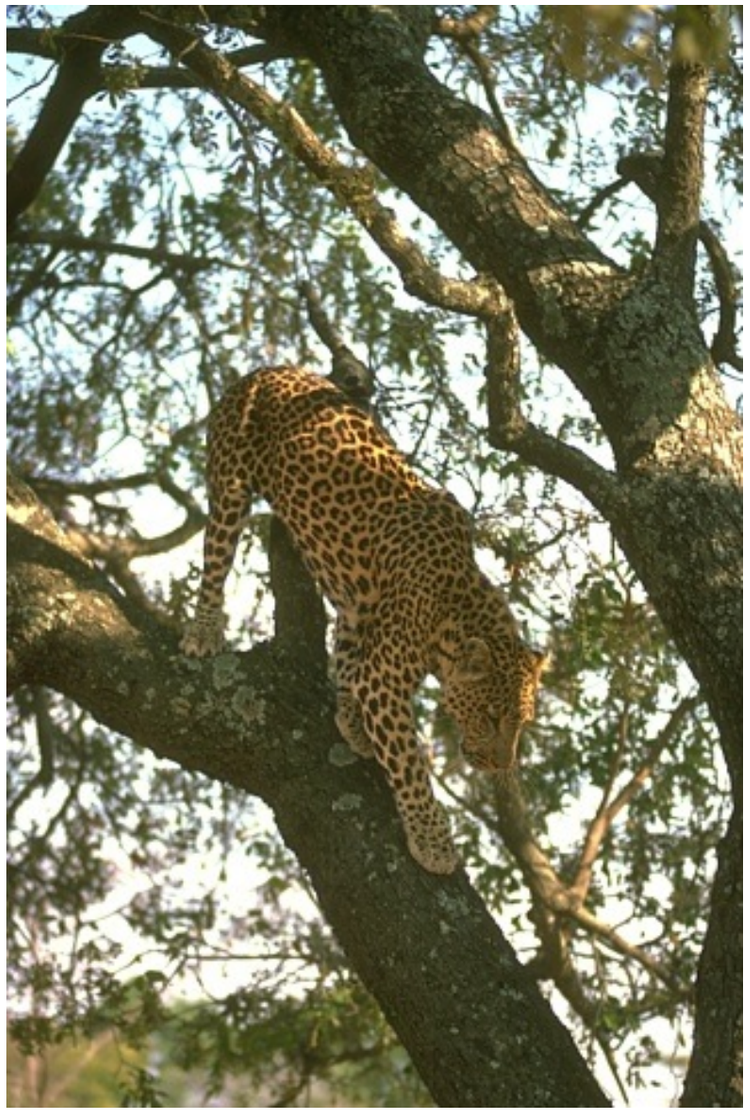

Boundaries

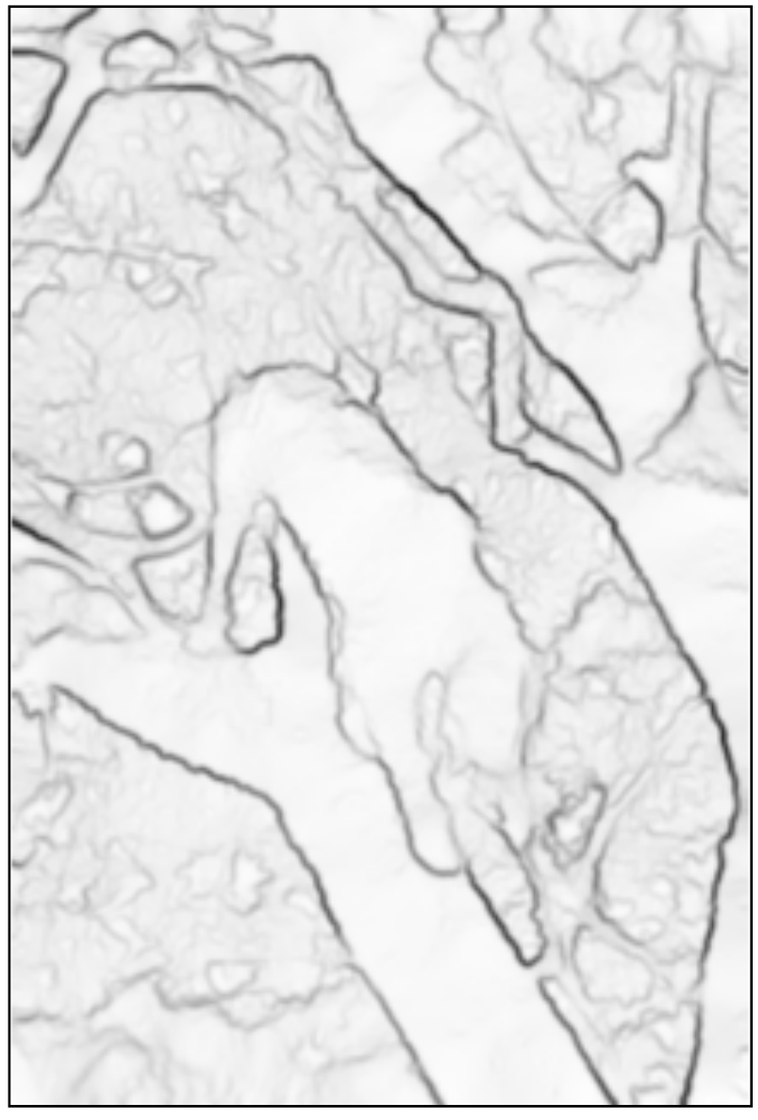

Segments

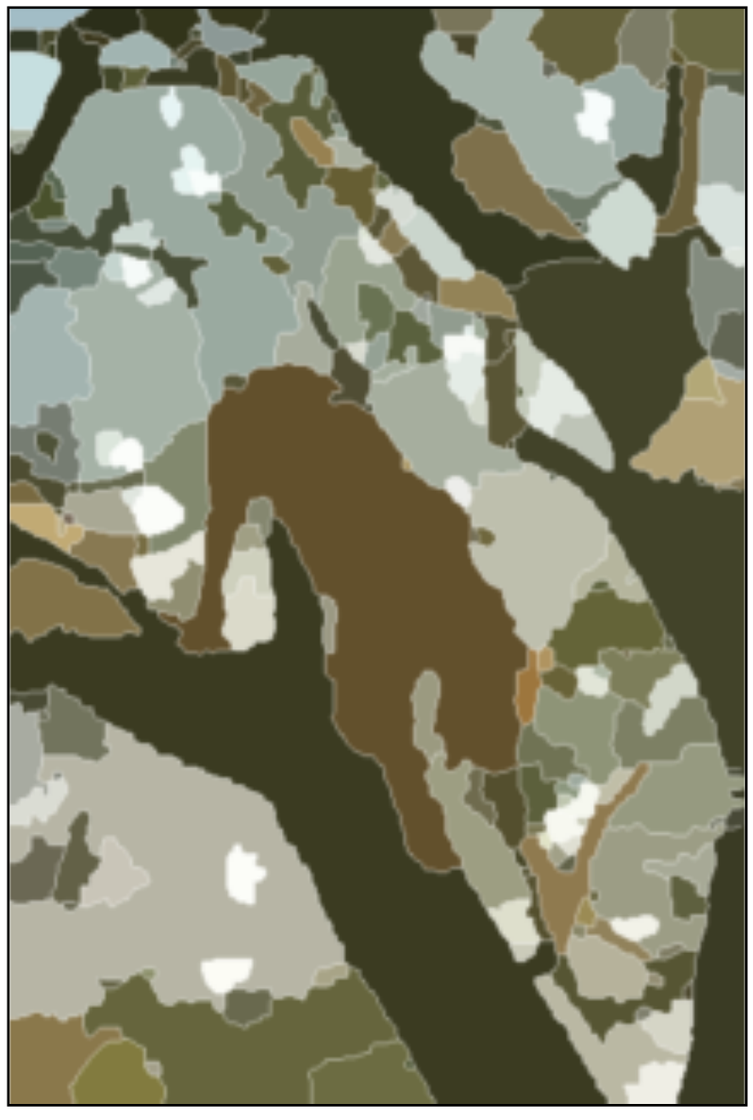

Arbeláez et al. 2011

Results

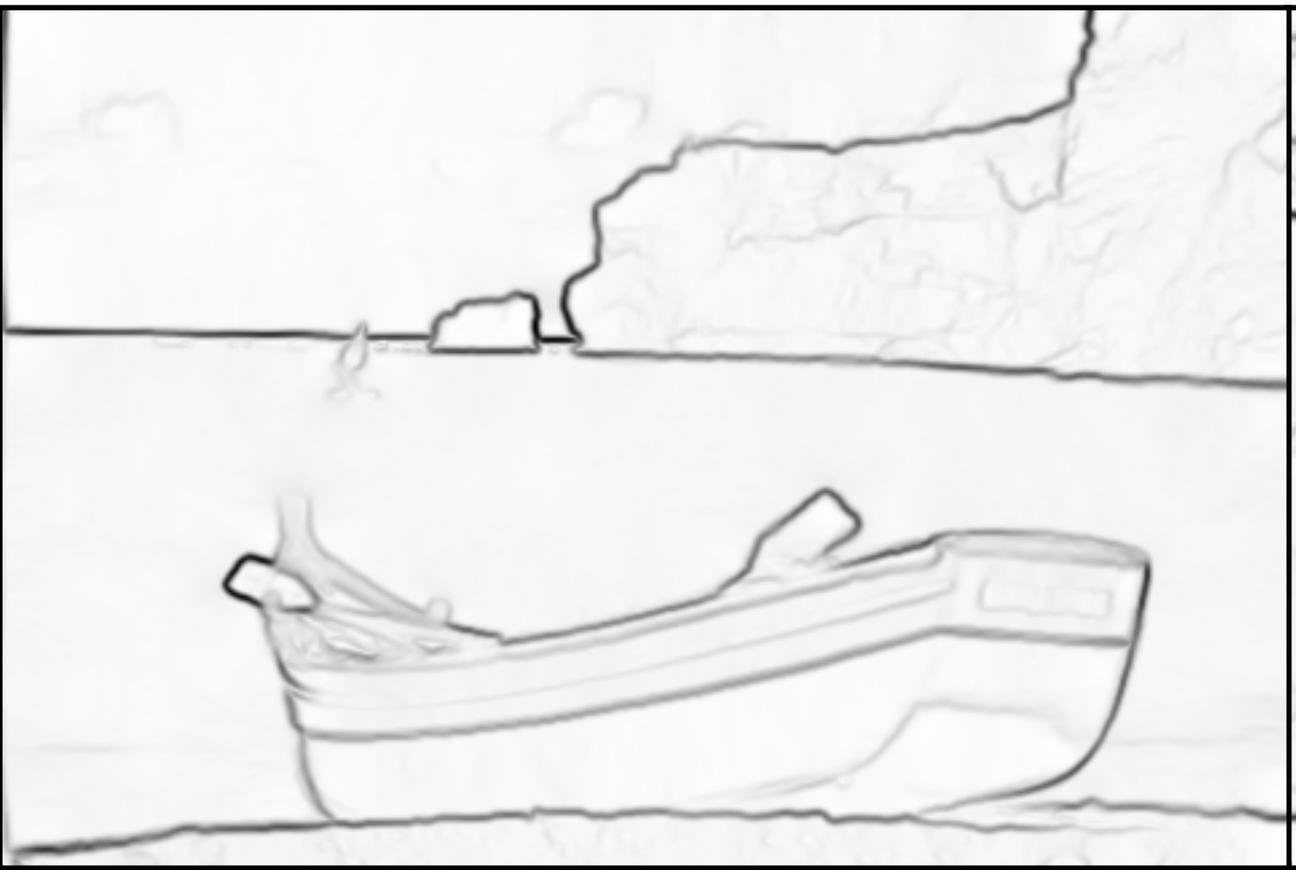

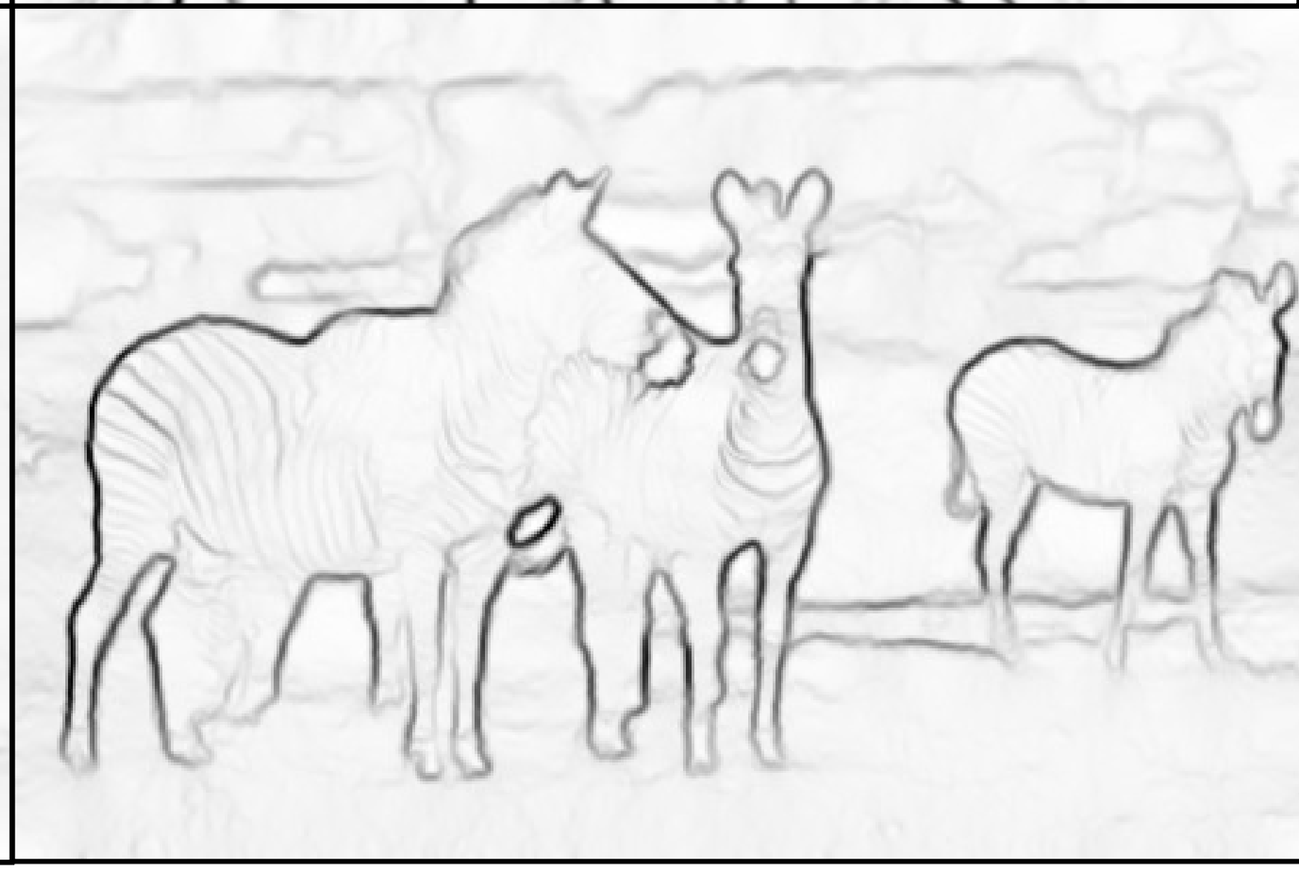

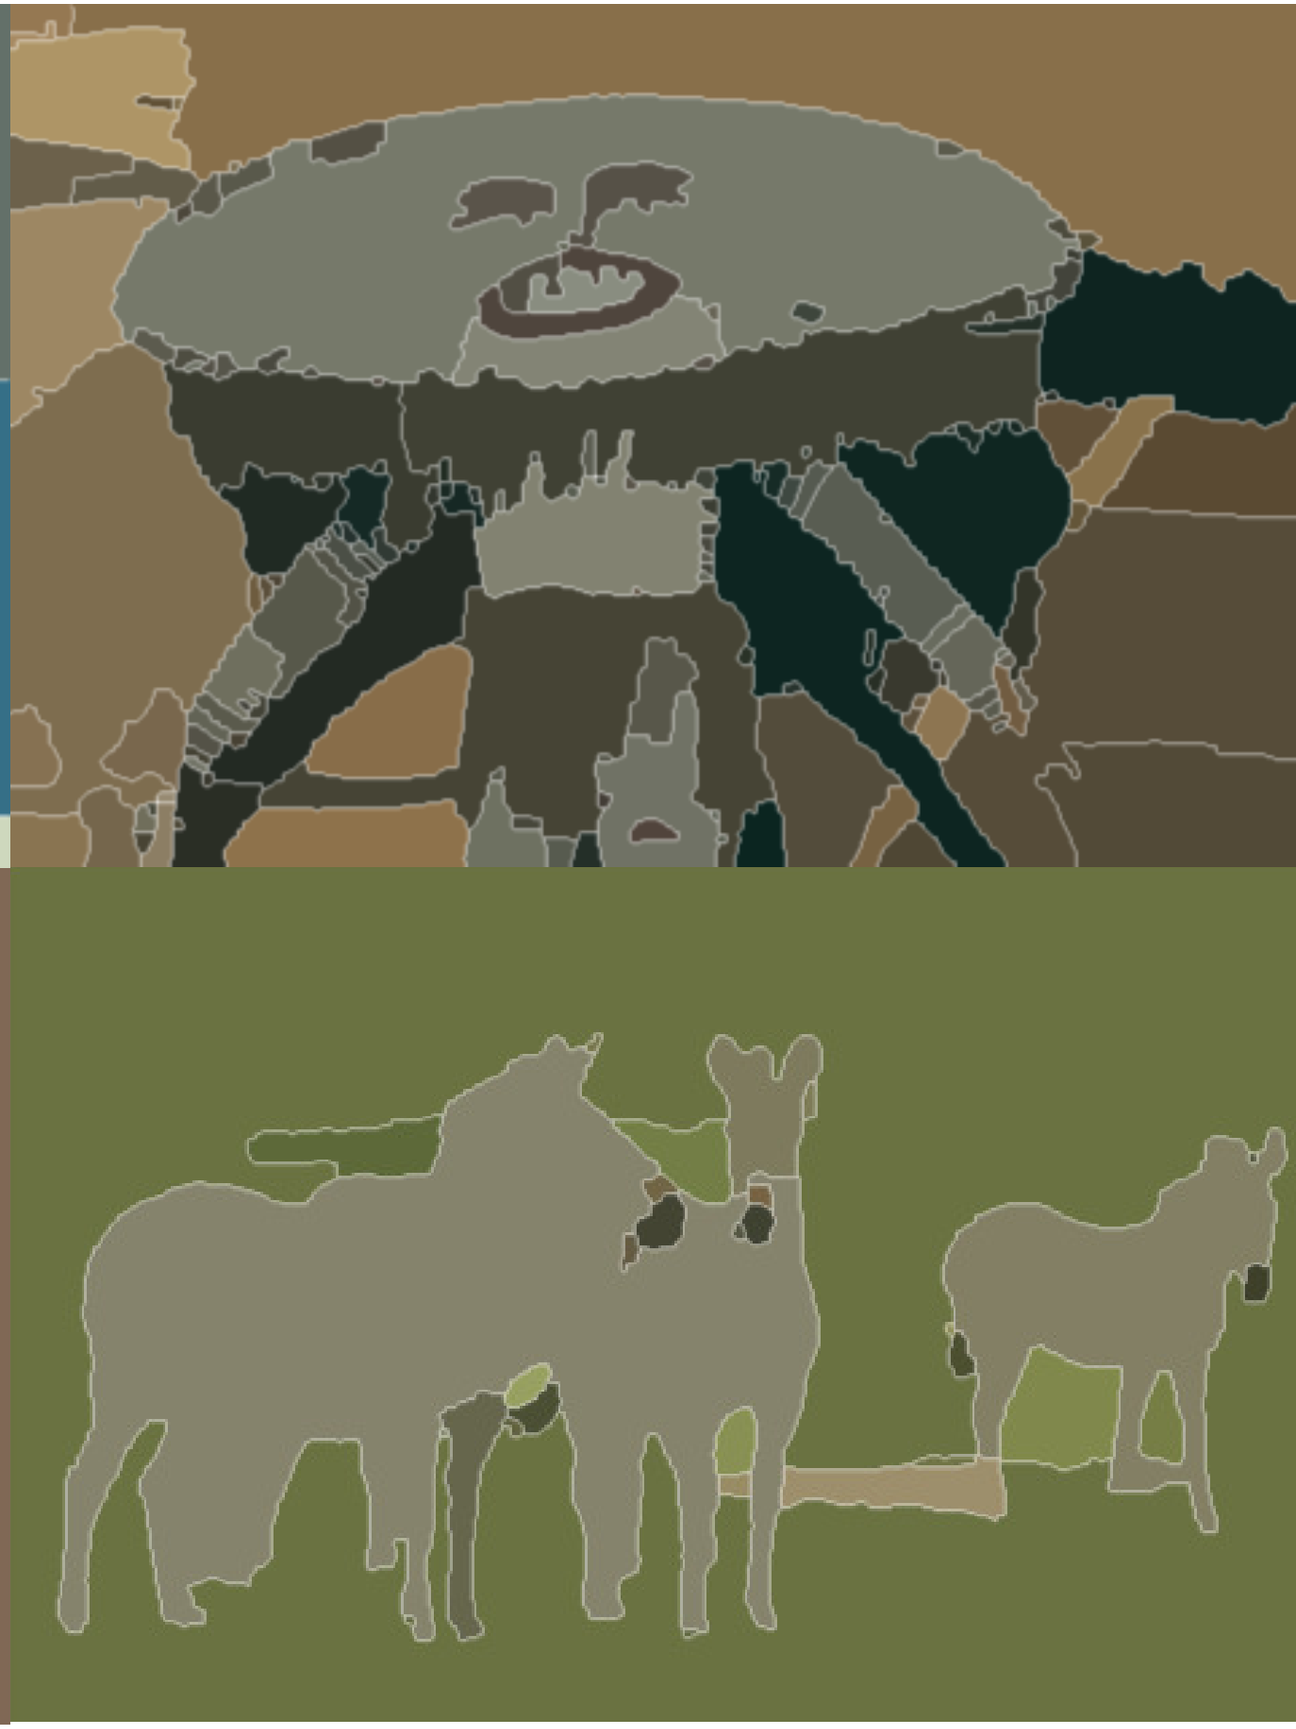

Performance on BSDS500

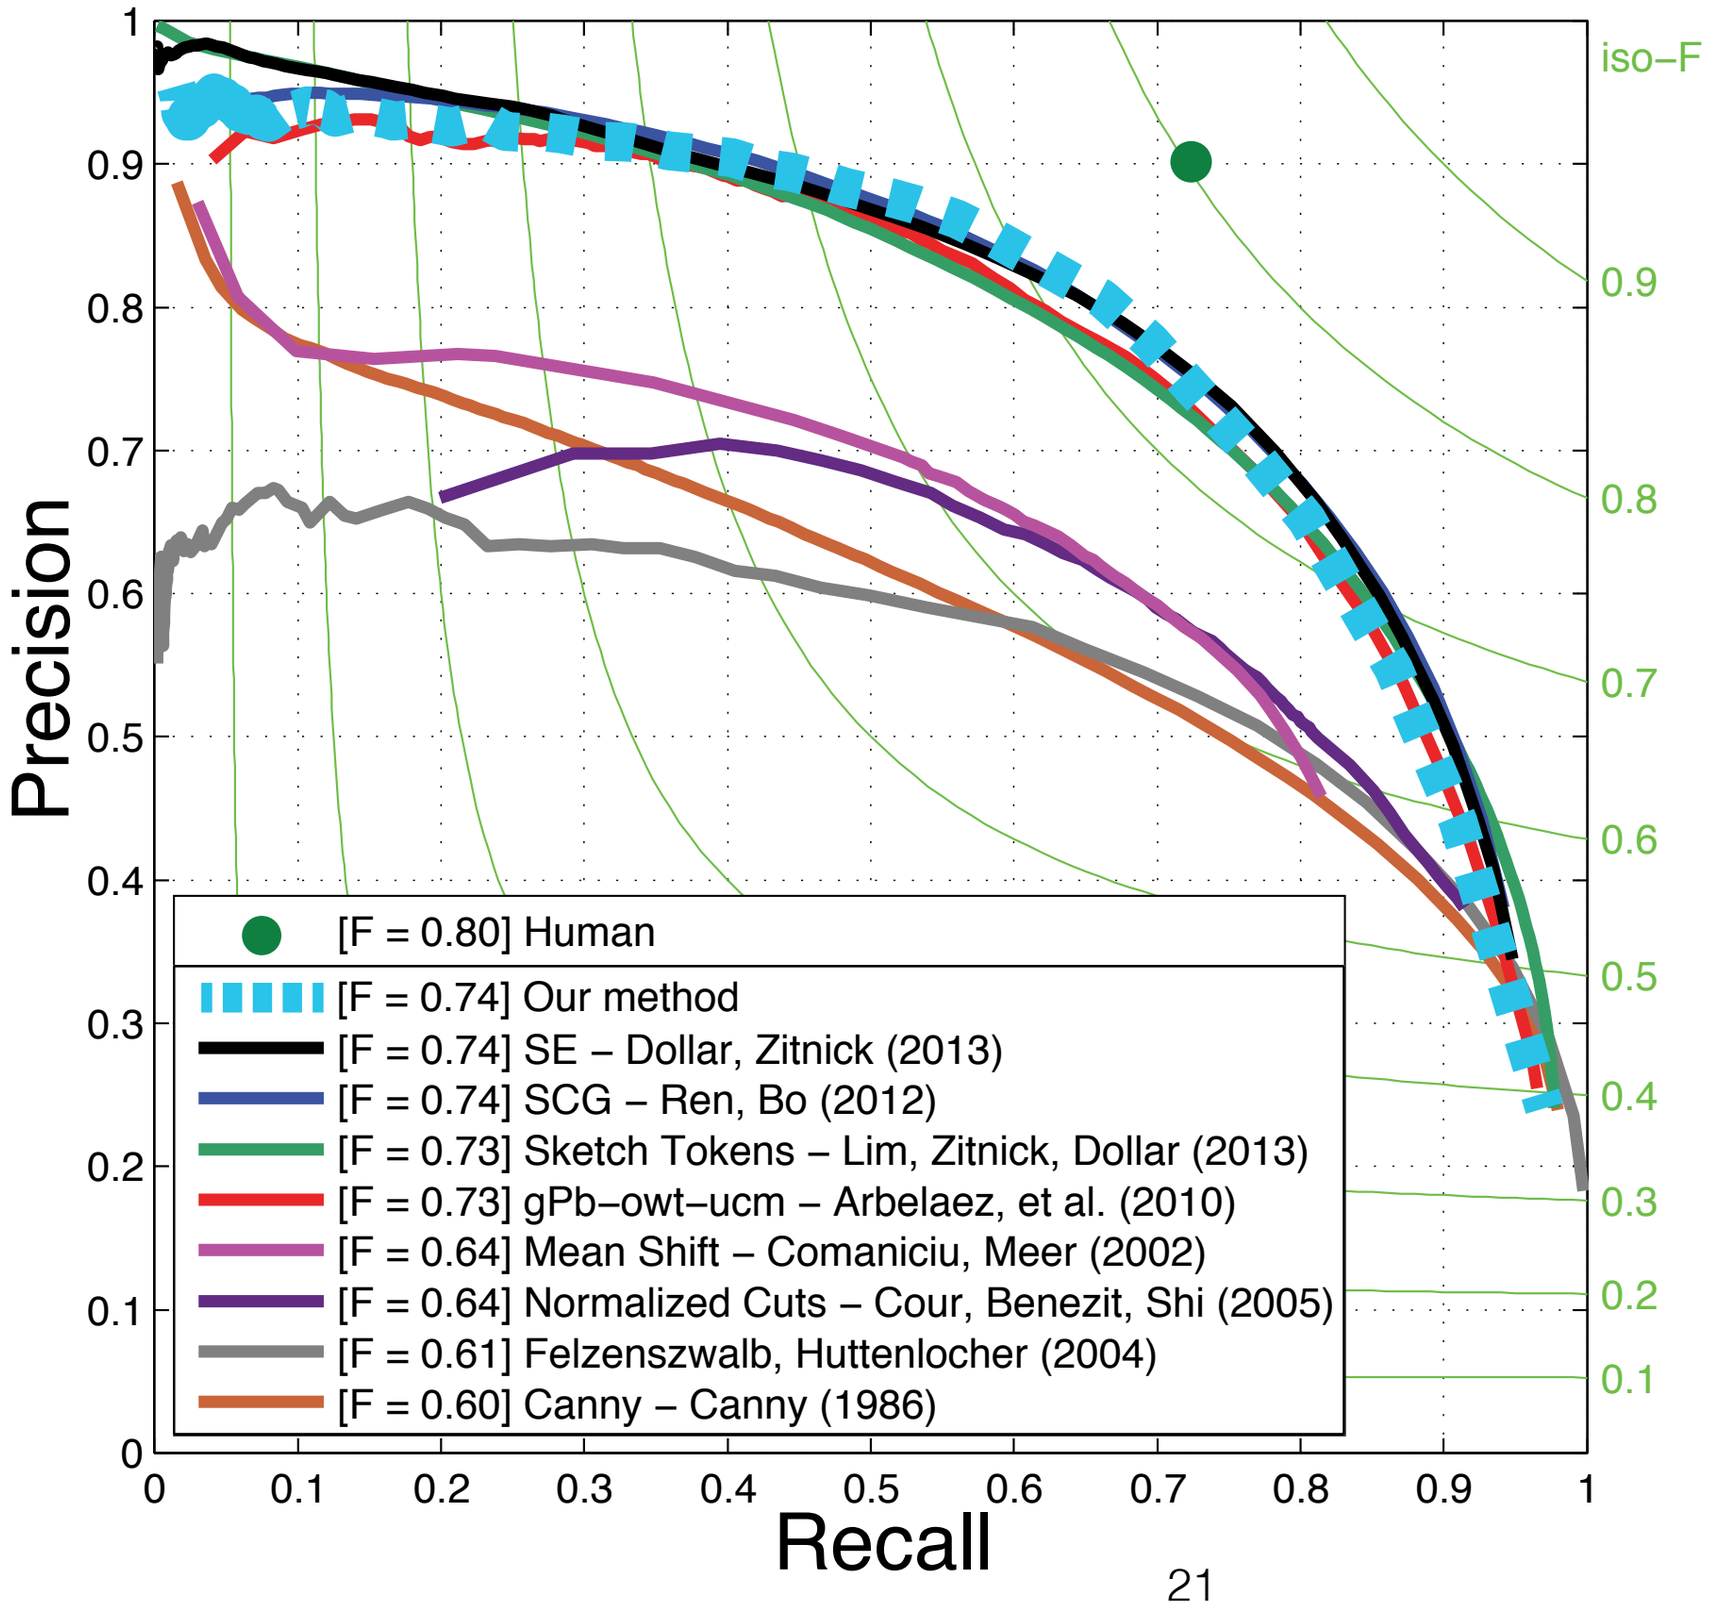

ODS: **0.74**
OIS: **0.77**
AP: **0.80**

Works on diverse stimuli

Cellphone photo

Art

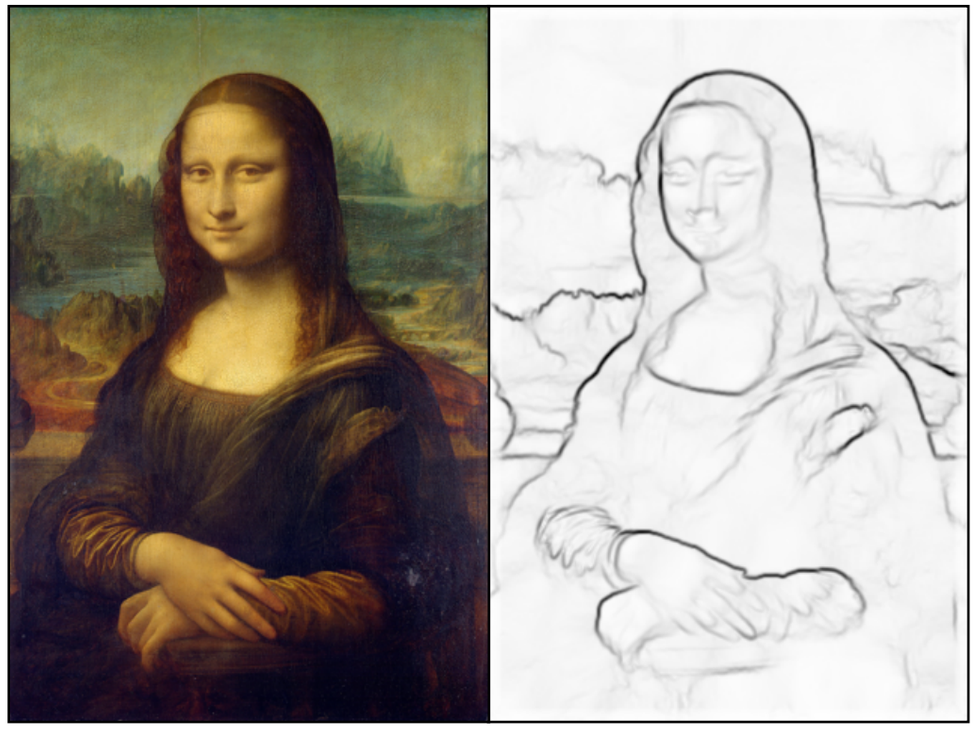

Satellite imagery

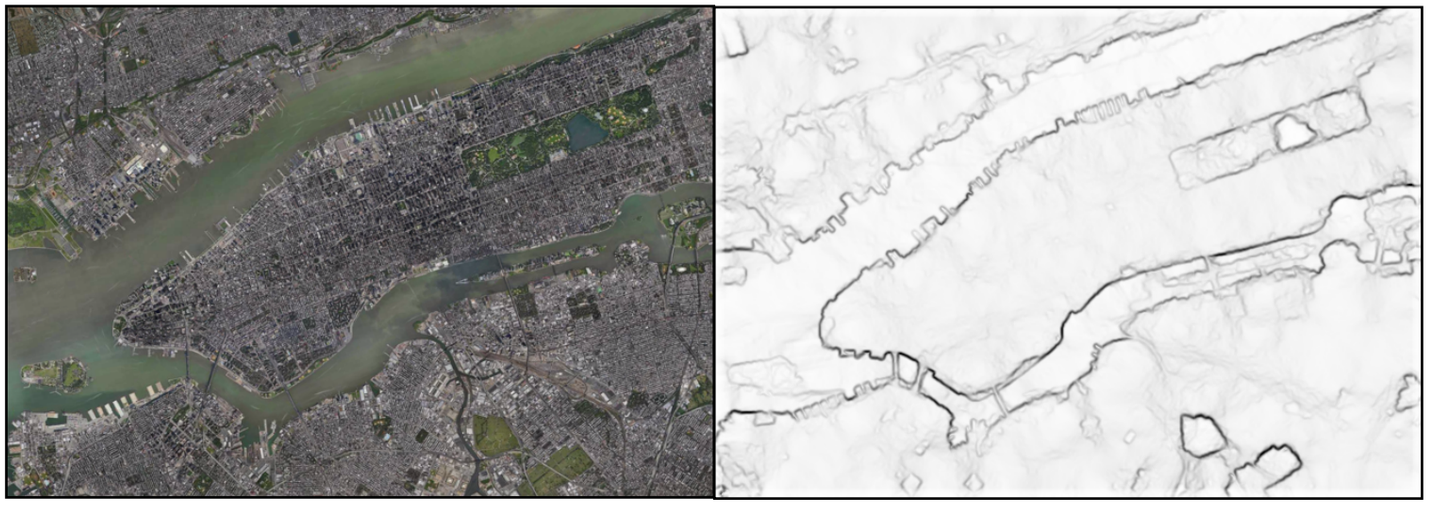

Boundary detection in XYT

Summary

Pointwise mutual information is a powerful affinity measure, with applications to boundary detection and segmentation.

It is unsupervised and relies entirely on simple internal image statistics.

Code available:

```
boundaries = findBoundaries(I);
```

mit.edu/pmi-boundaries

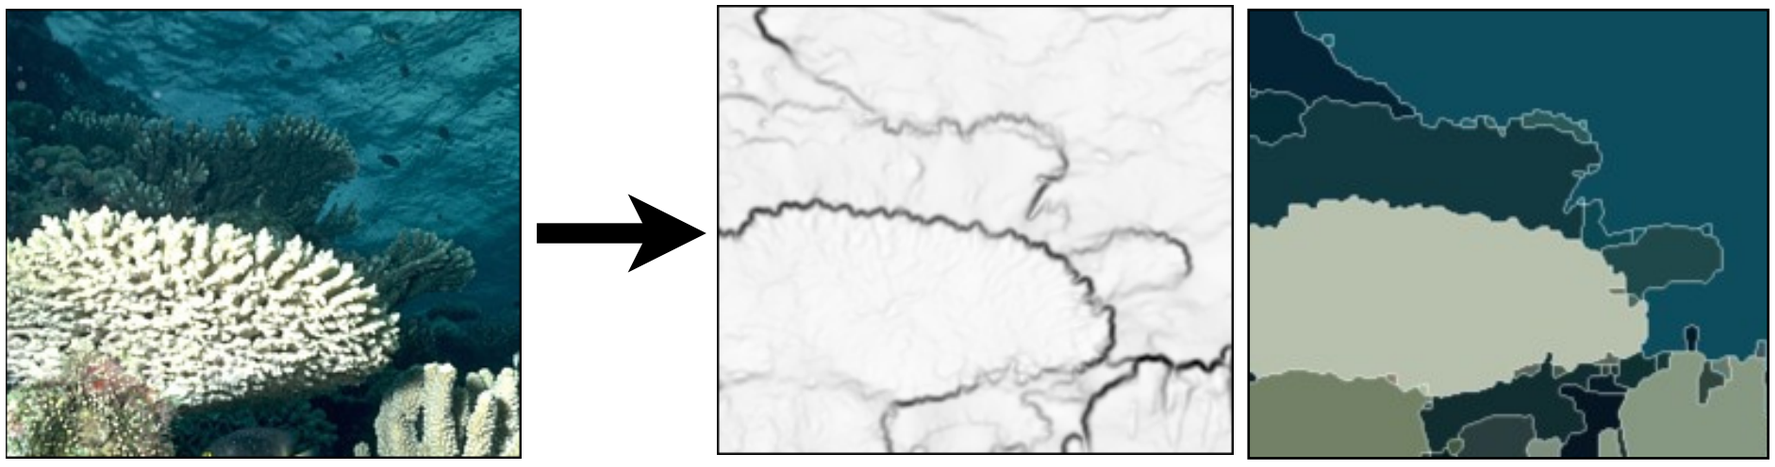

Thanks!