2D Image Analysis

2D Image Analysis

- Segmentation
- Localization
- Shape Analysis
- Classification & Categorization

Segmentation

- Partition images into *meaningful* entities
- The "*Holy-Grail*" problem in low-level computer vision
- A panacea to many high-level vision problems

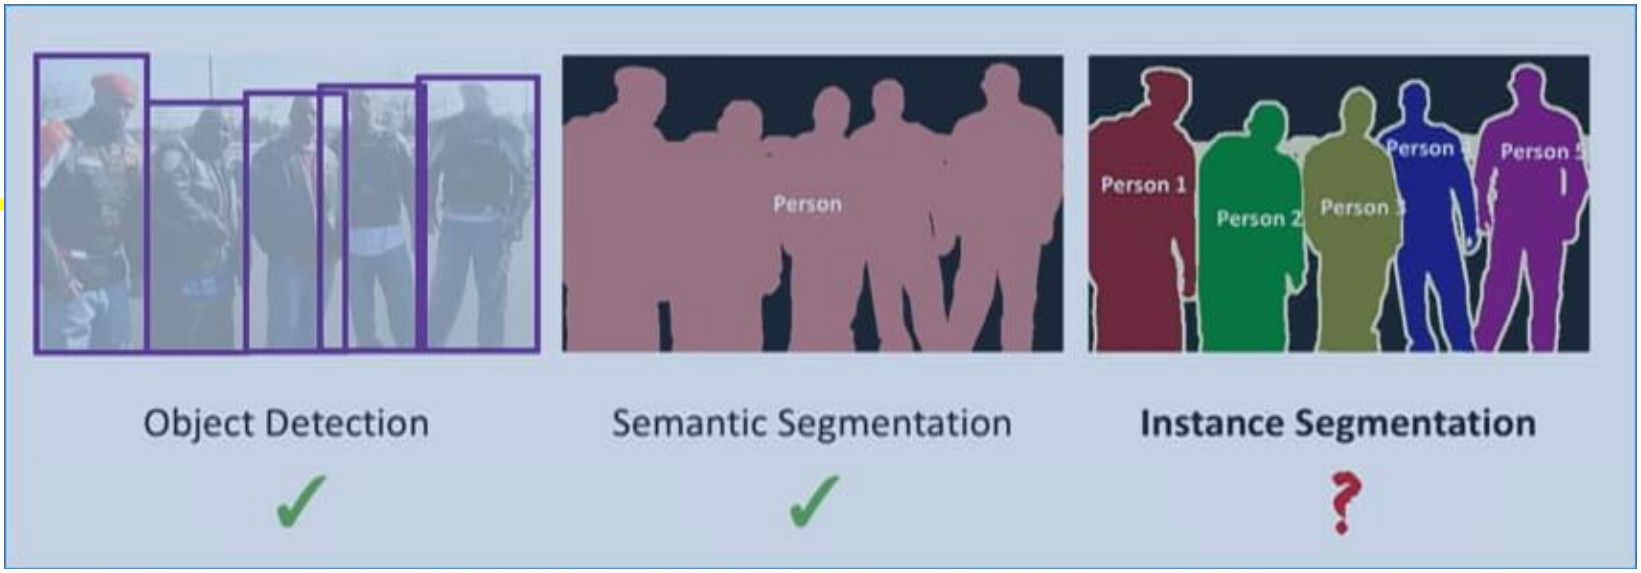

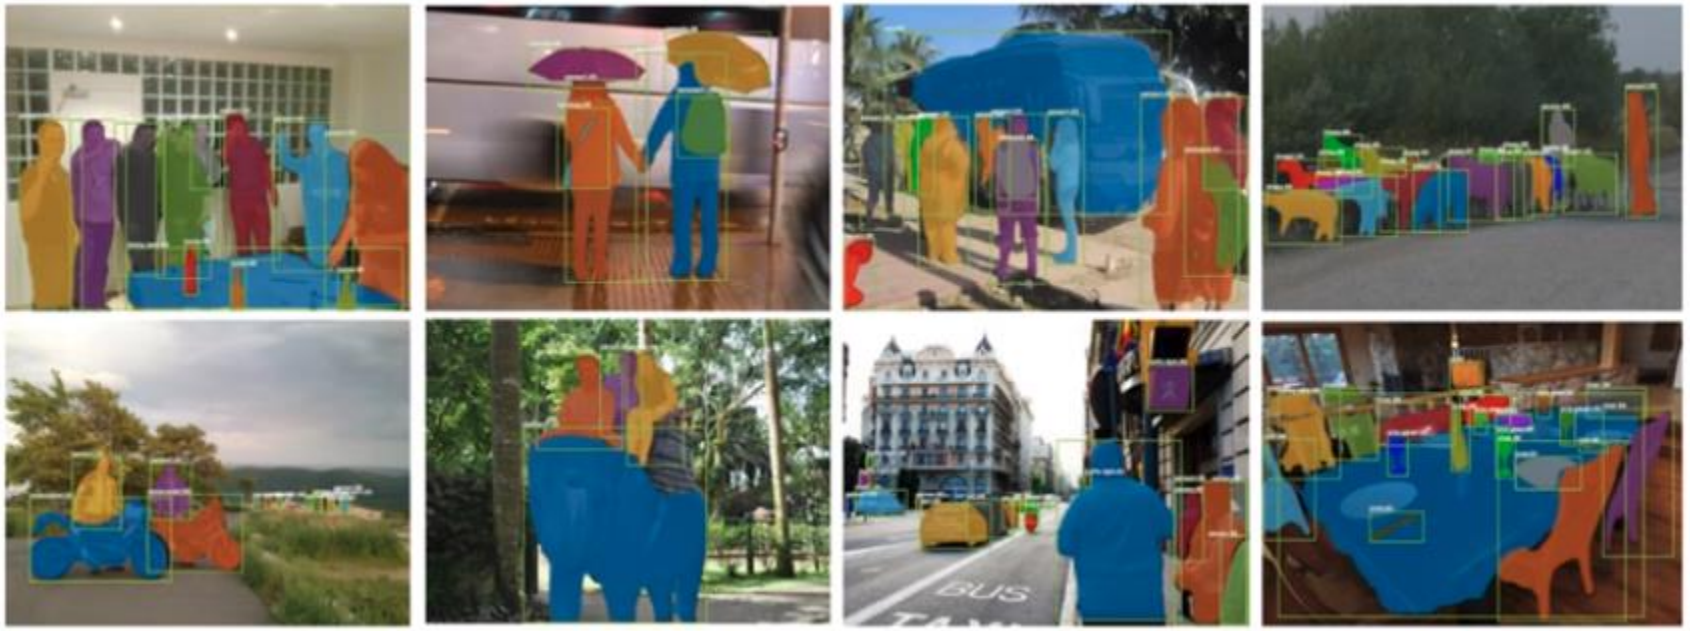

https://www.youtube.com/watch?v=OOT3UIXZztE

https://www.youtube.com/watch?v=g7z4mkfRjI4

University of California Santa Barbara

Segmentation (cont.)

- Even though appeared simple, the problem is extremely hard
  - Noise
  - Sensing and lighting conditions
  - Repetitive patterns
  - Syntactic vs. semantic grouping
  - Top down vs. bottom up approaches

What is segmentation?

- Isolating a specific region of interest ("find the star" or "bluish thing")

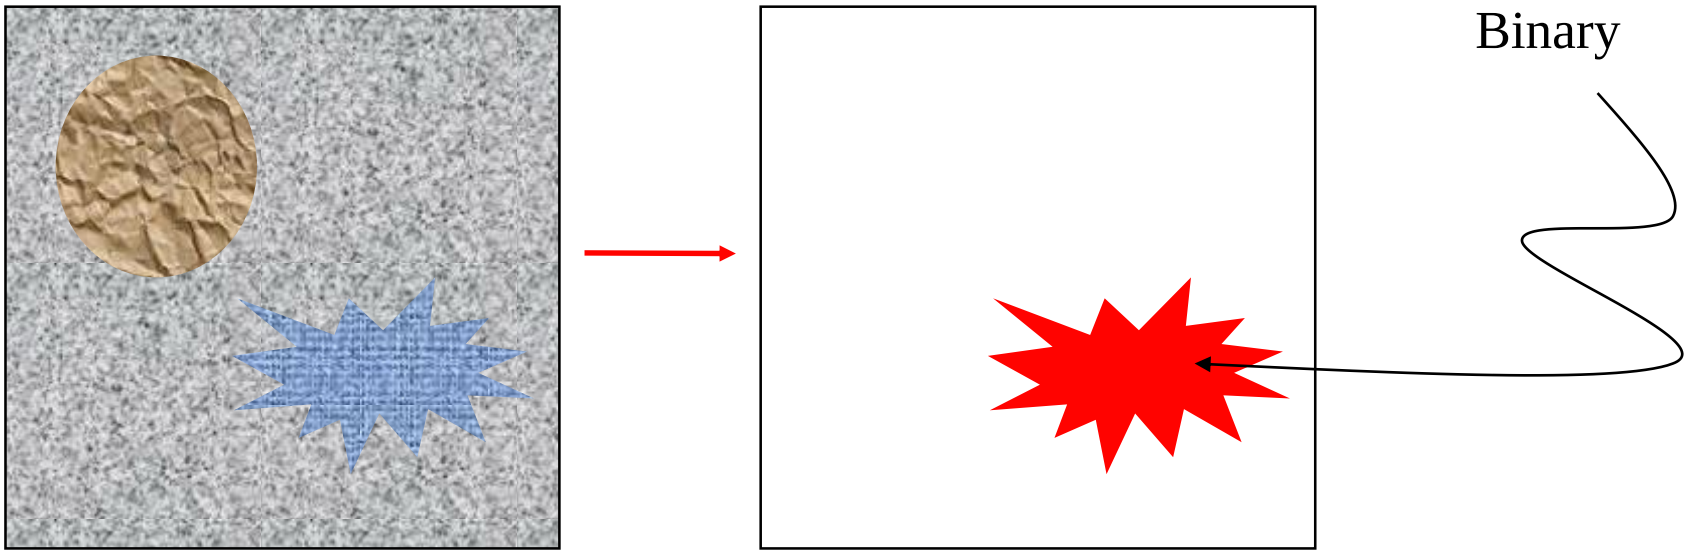

What is segmentation?

- Partitioning images/volumes into meaningful pieces

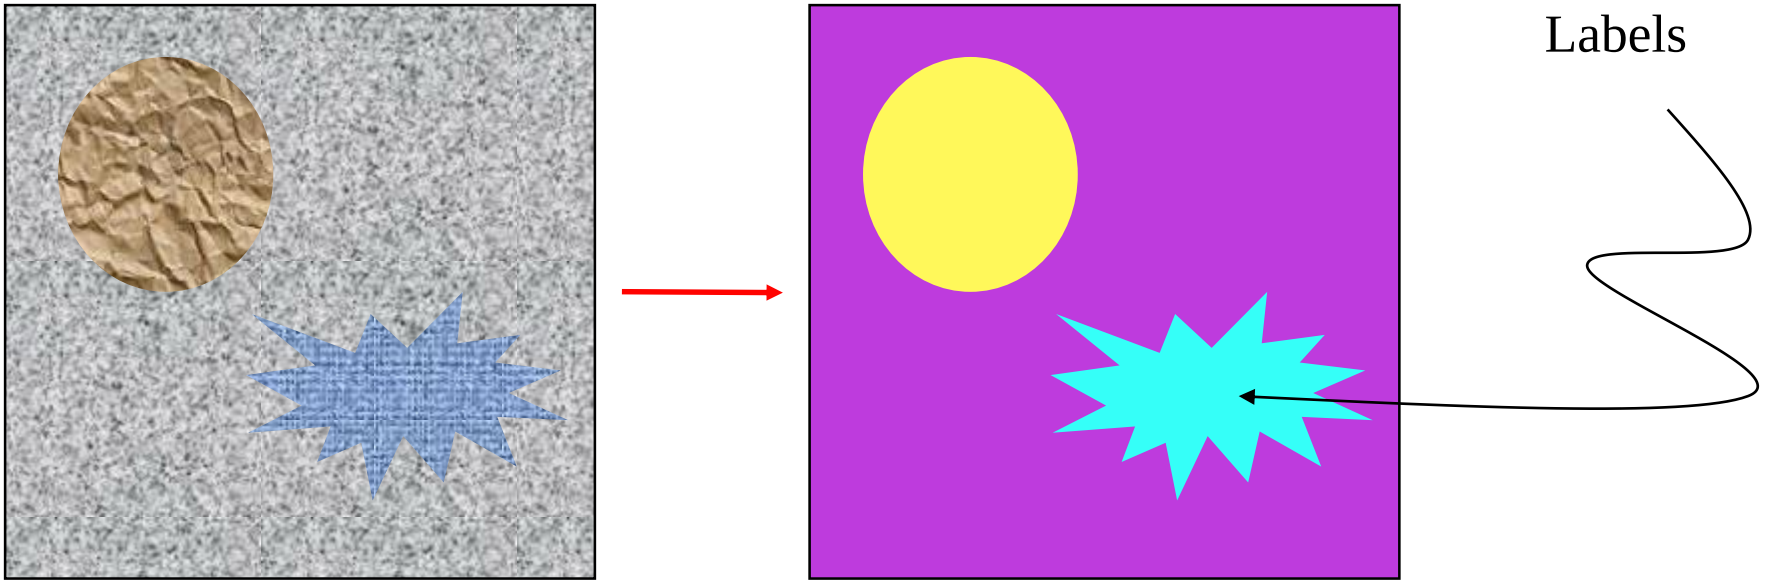

What is segmentation?

- Assigning each pixel a type (tissue or material)

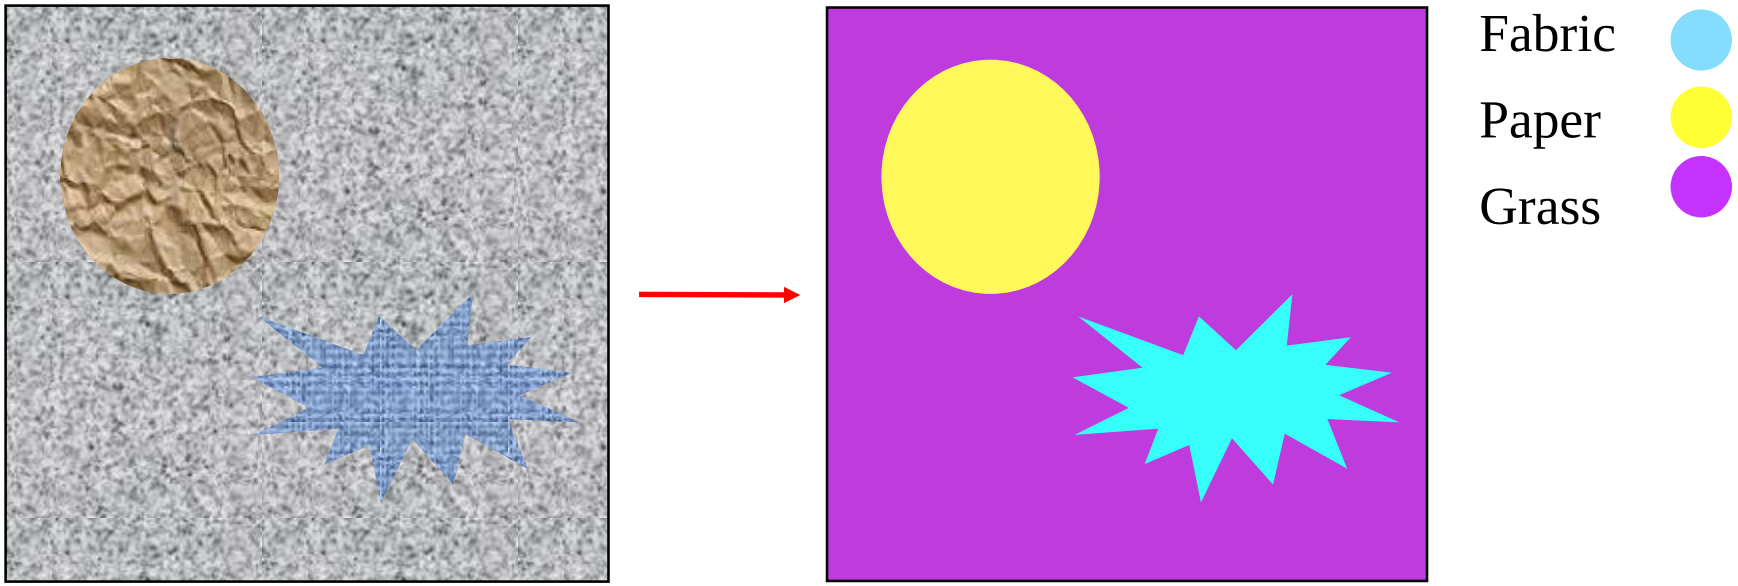

Examples

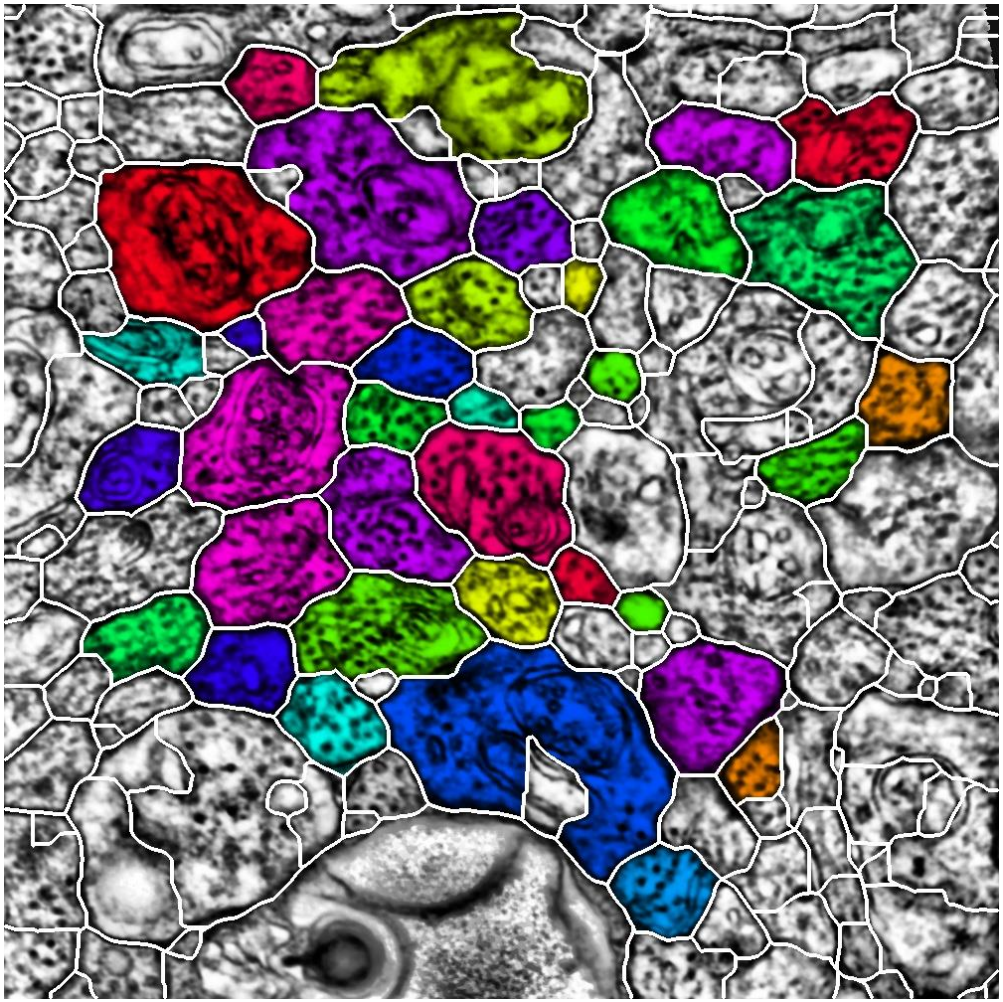

Example

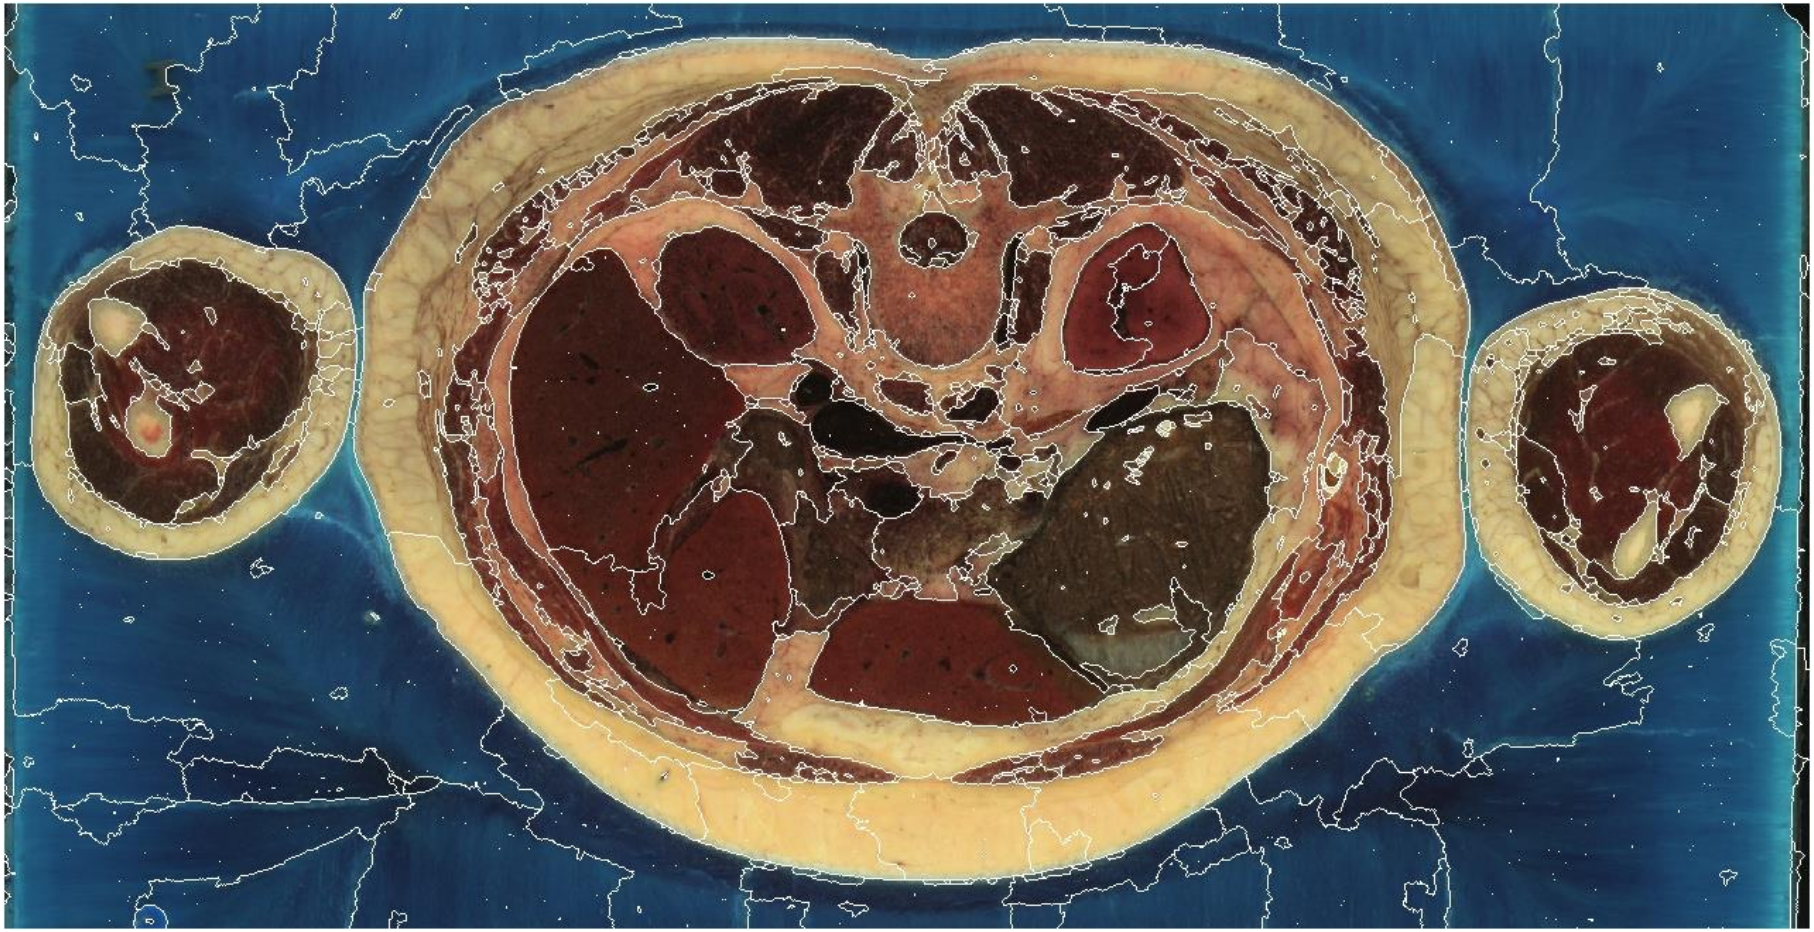

Examples

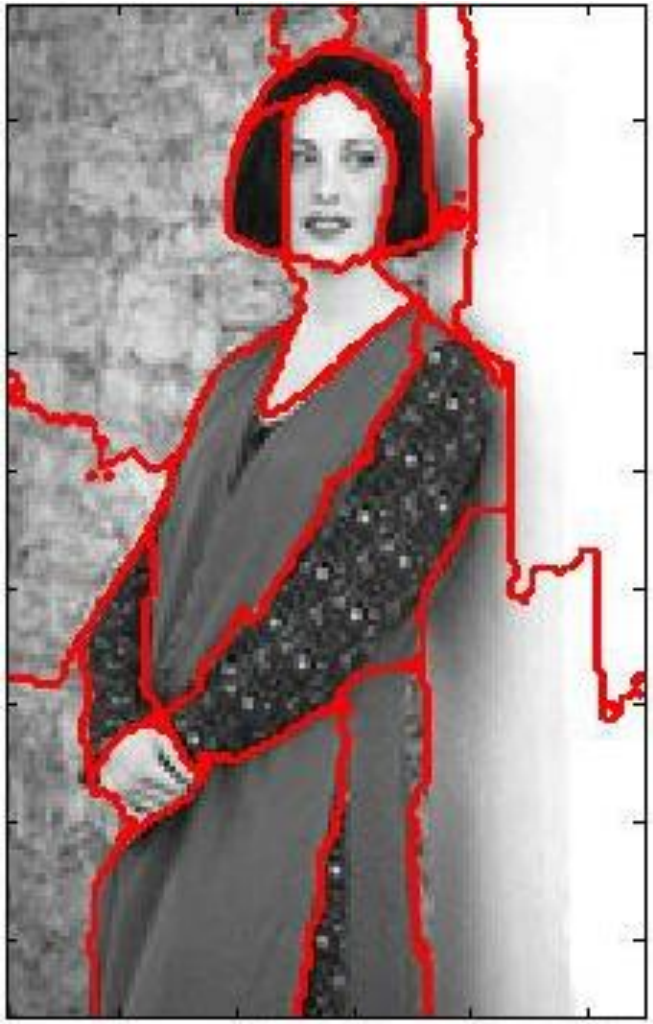

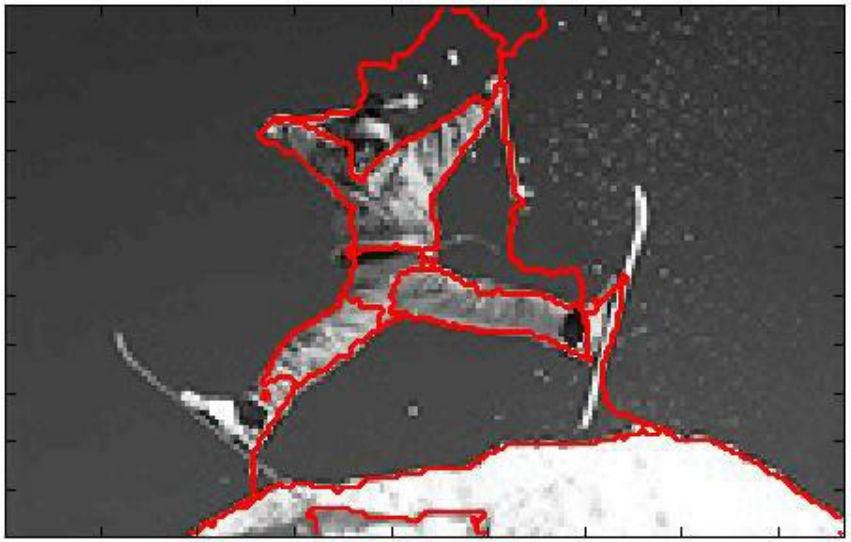

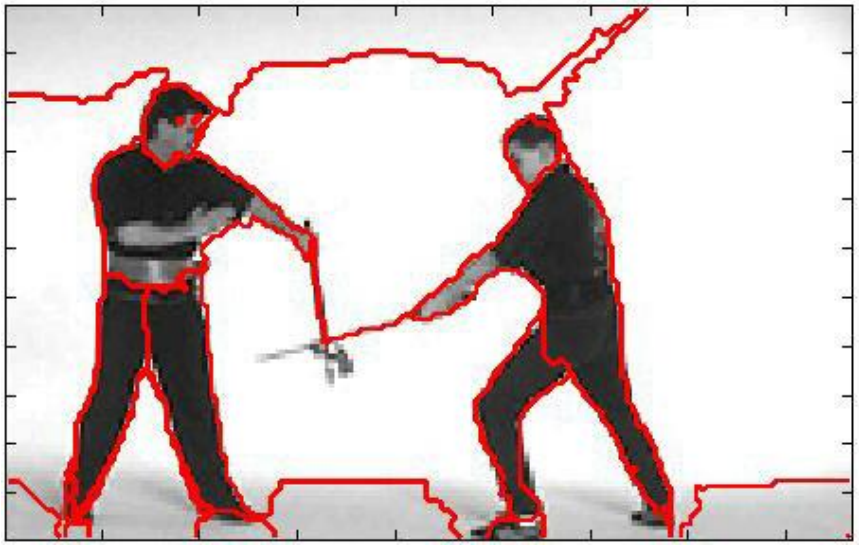

Examples

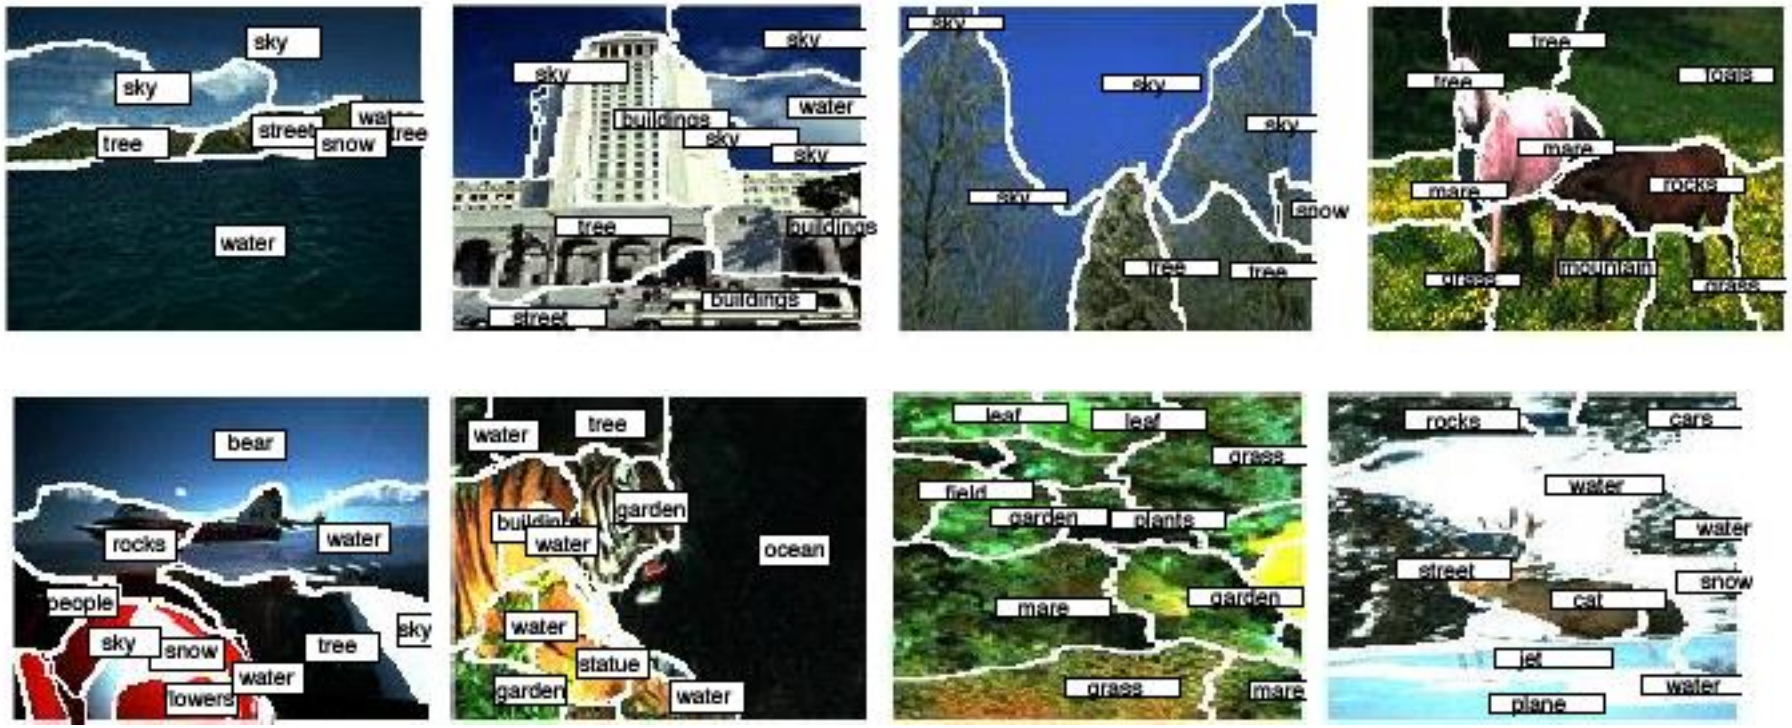

Query image: 108019

Query blobs

	feature importance:				
	overall	**color**	**texture**	**location**	**shape**
blob	very	very	somewhat	not	not
background	somewhat	very	not	not	not

Querying from 35000 images (2000 returned by the filter).

1: 108044 (score = 0.99)

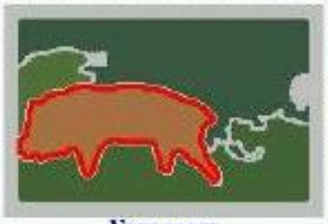

New query

2: 108023 (score = 0.98)

New query

3: 108006 (score = 0.98)

New query

4: 108029 (score = 0.98)

New query

5: 108051 (score = 0.98)

New query

6: 108084 (score = 0.97)

New query

7: 108037 (score = 0.97)

New query

8: 108004 (score = 0.97)

New query

Examples

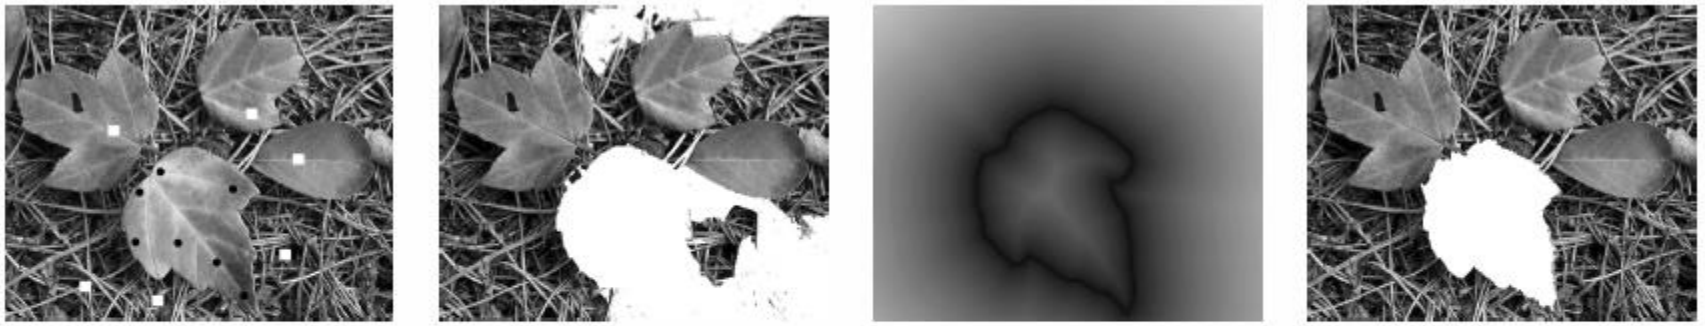

Figure 6. Segmenting a maple leaf. (a) The user input: circles indicate the object, squares indicate the background (colour is for visualization purposes only). (b) The result without shape priors – segmentation is shown in white. (c) The level-set of the shape template after transformation, $\bar{\phi}_{trans}$. (d) The result with shape priors.

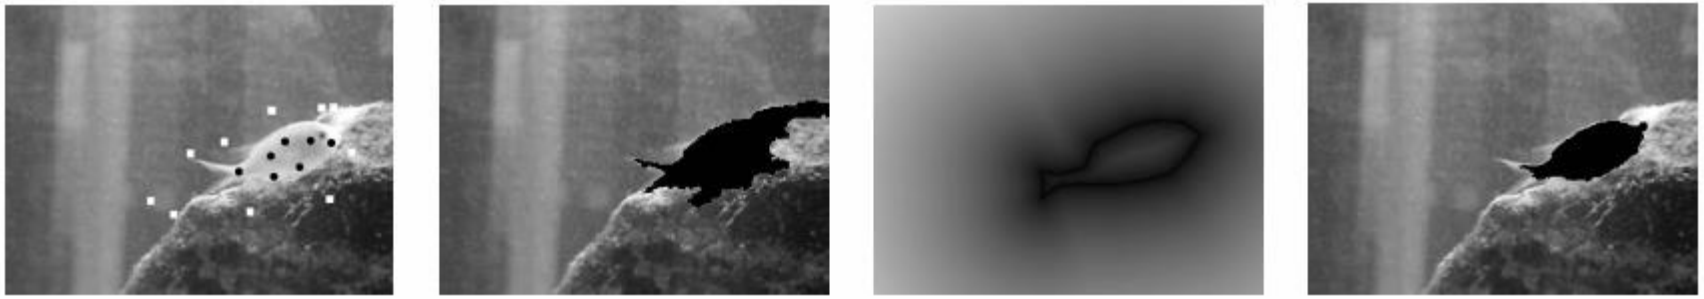

Figure 7. Segmenting a fish. (a) The user input: circles indicate the object, squares indicate the background (colour is for visualization purposes only). (b) The result without shape priors – segmentation is shown in black. (c) The level-set of the shape template after transformation, $\bar{\phi}_{trans}$. (d) The result with shape priors.

Examples

- Sony EyeToy
  - Background, motion, and color segmentation

Segmentation

- Spatial and temporal segmentation
  - And spatial-temporal (“spatiotemporal”) segmentation
- Segment images/video based on:
  - Grayscales
  - Color
  - Textures
  - Depth
  - Motion
  - Low-level features
  - Etc.

Gelstalt Examples

Not grouped

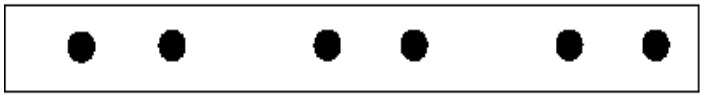

Proximity

Similarity

Similarity

Common Fate

Common Region

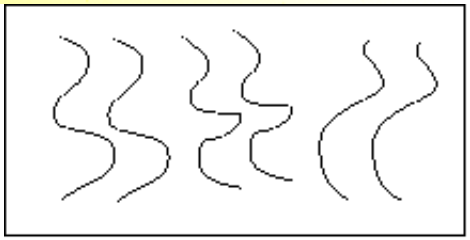

Parallelism

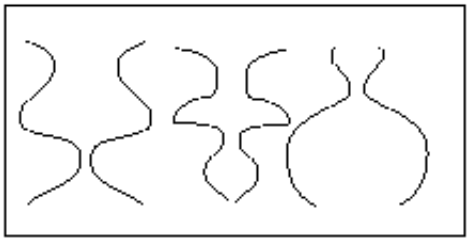

Symmetry

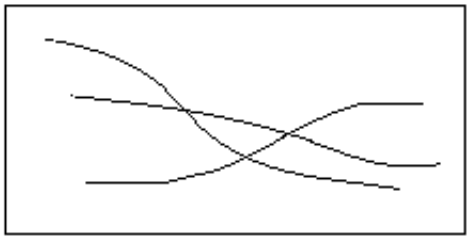

Continuity

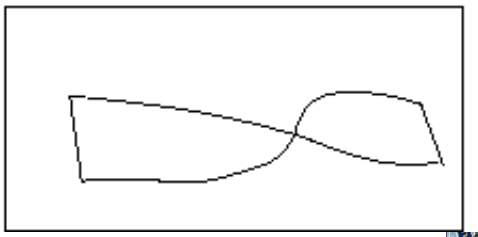

Closure

- How edges should be grouped?
- How regions should be defined?
- Semantic vs. syntactic

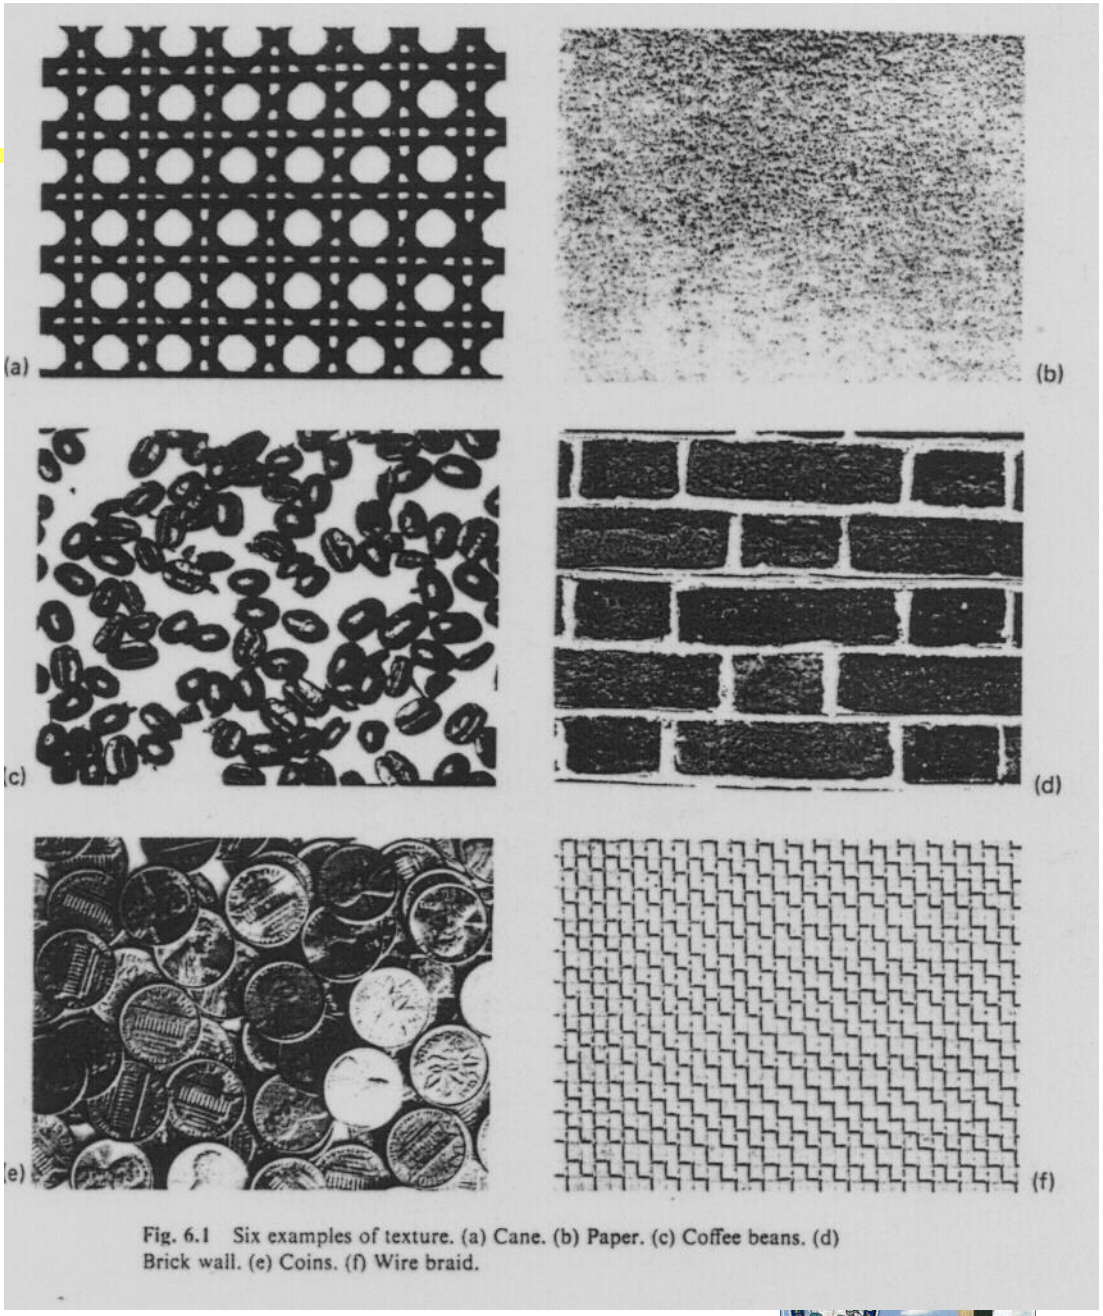

Fig. 6.1 Six examples of texture. (a) Cane. (b) Paper. (c) Coffee beans. (d) Brick wall. (e) Coins. (f) Wire braid.

Syntactic vs. Semantic

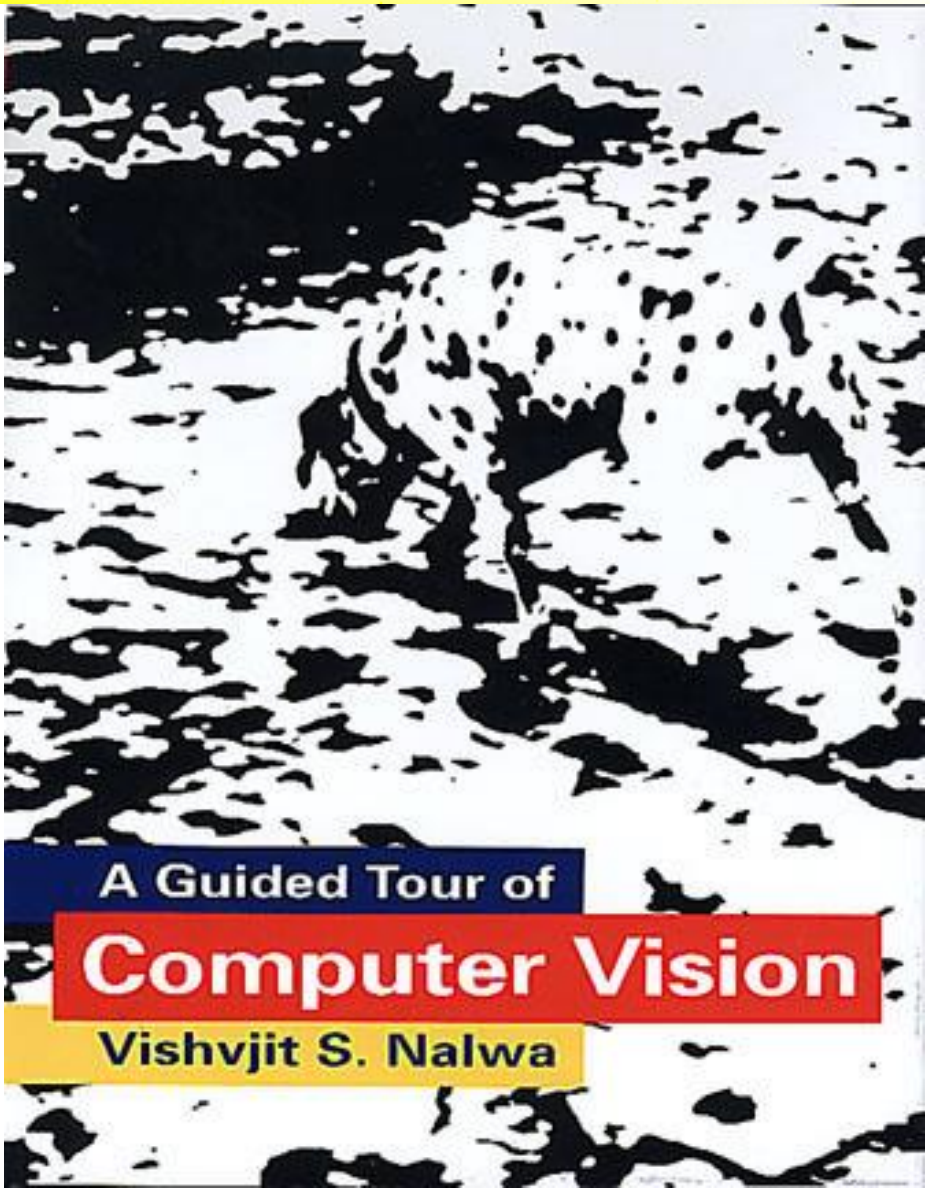

2D Image Analysis (cont.)

- Representation (*syntactic* level)
  - Describe the *shape* (*appearance*) of edges and regions
    - regions: size, location, orientation, etc.
    - edges: curvature, orientation, length, etc.
    - info can be extracted from images alone

2D Image Analysis (cont.)

- Interpretation (*semantic* analysis)
  - Describe the *identity* of image features
  - Regions: sky, water body, etc.
  - Edges: 3D orientation, occluding contours, road boundaries, etc.
  - Often need domain specific knowledge and contextual information

General purpose segmentation strategies

- Region-based methods
  - Regions are locally homogeneous (in some property)
  - Regions satisfy some property (to within an tolerance)
  - E.g., Flood fill

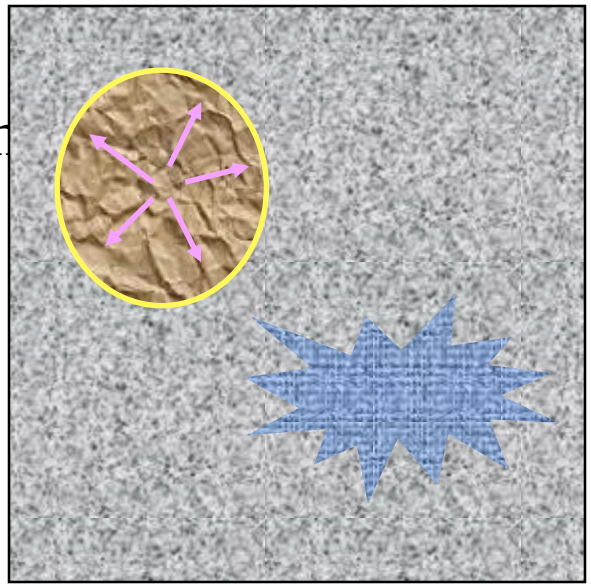

- Edge- or contour-based methods
  - Regions are bounded by features
  - Features → sharp contrast
  - E.g., Canny Edges

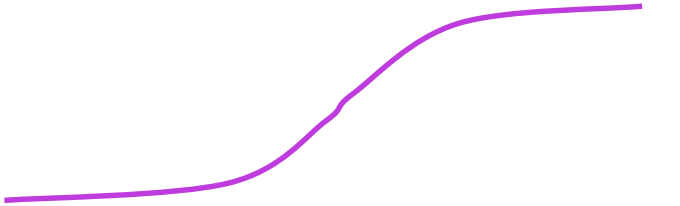

- Bottom up (from images to features and objects)
- Syntactic information

Segmentation via deformable models

- Active contours
  - Train models to learn certain shapes
- Snakes (polyline)
- Level sets

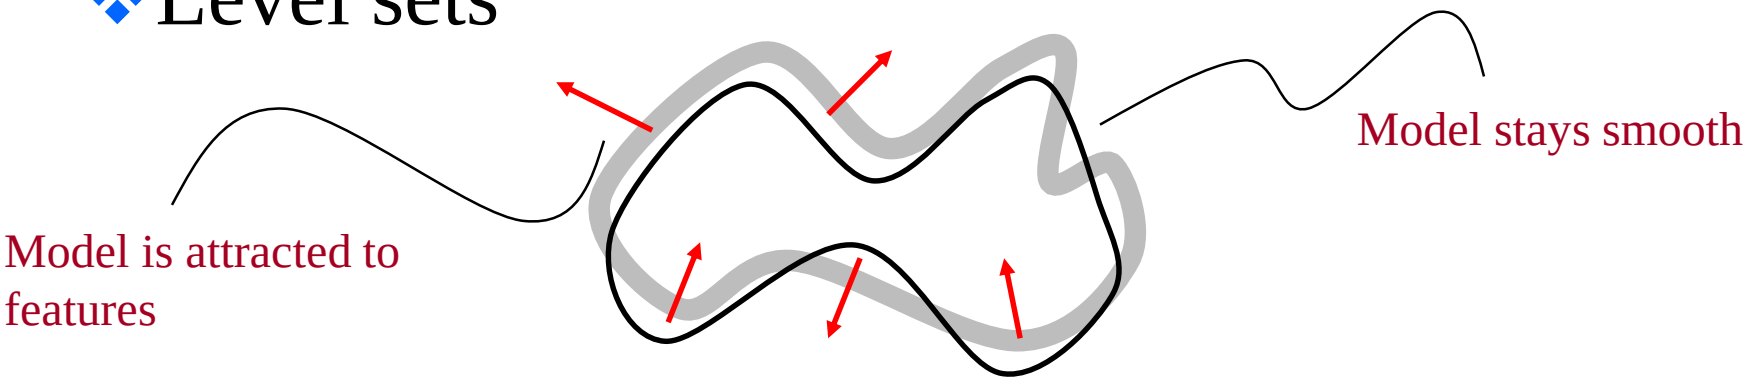

Grayscale-based segmentation (Thresholding)

- "Together" = similar grayscale values

Input image

Foreground segmentation

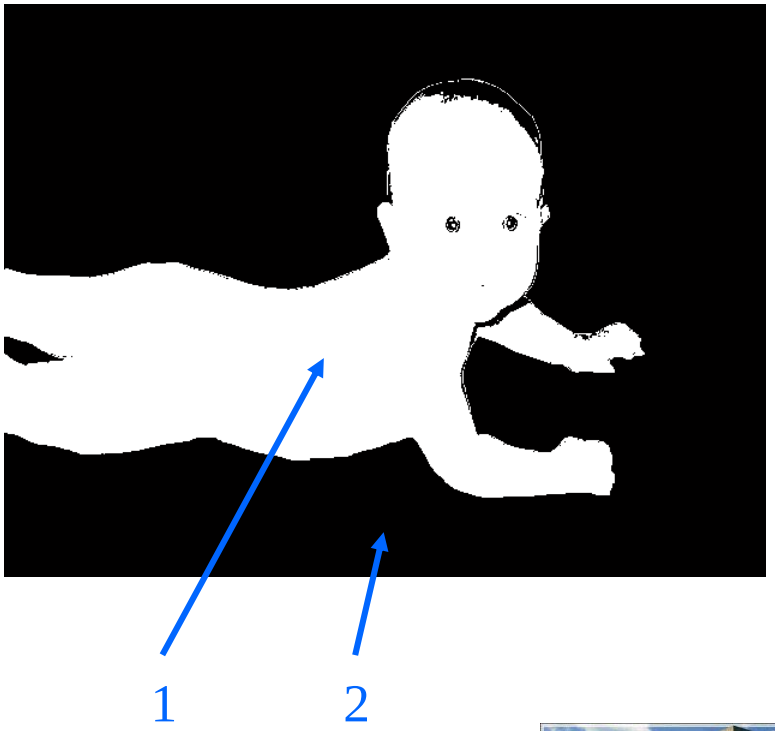

H

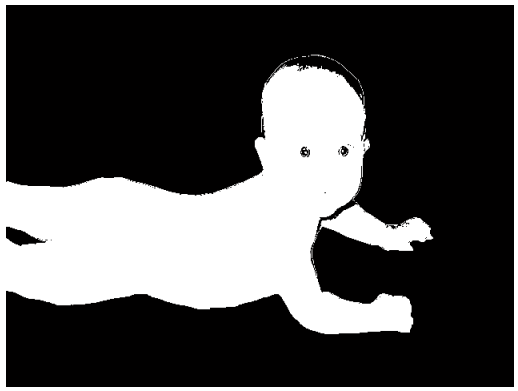

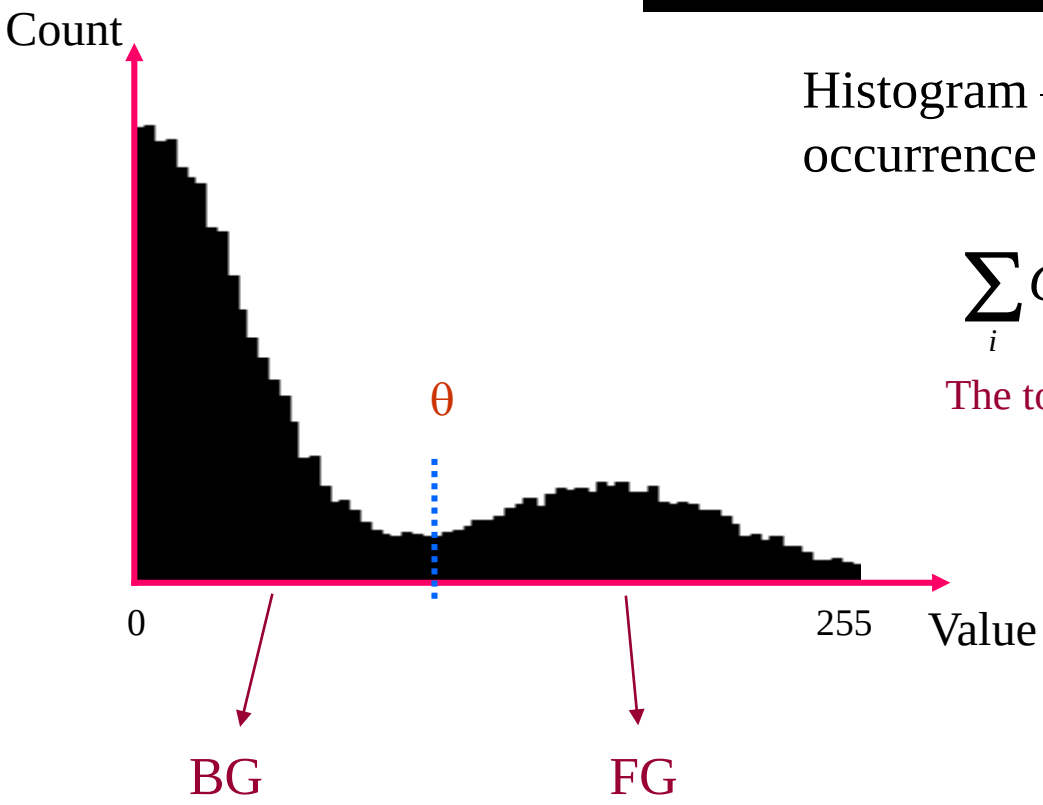

Histogram – a plot of frequency of occurrence (count) of pixel values

$$\sum_i Count(i) = ?$$

The total number of pixels in the image

FG/BG segmentation

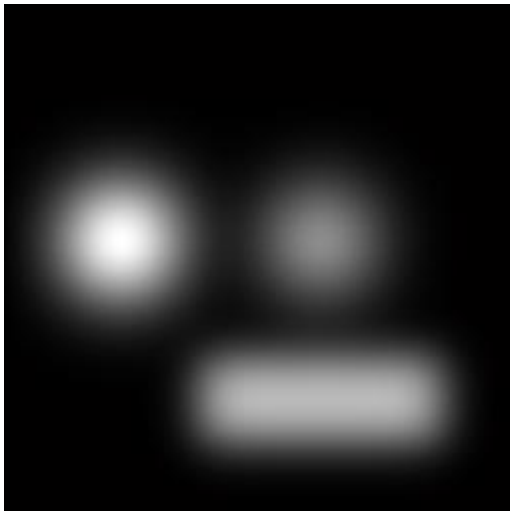

Original

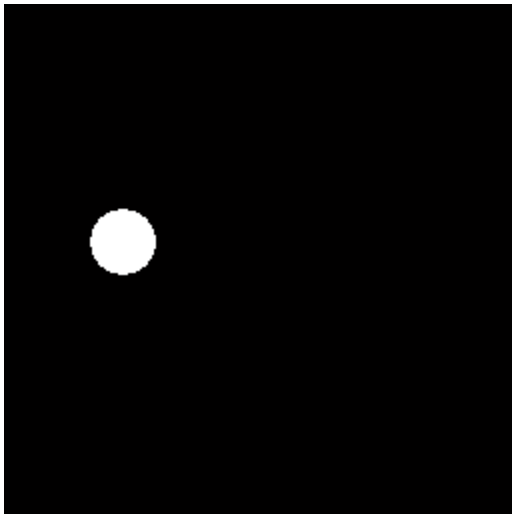

θ = 100
1 FG region

θ = 75
2 FG regions

θ = 50
3 FG regions

θ = 25
3 (large) FG regions

θ = 5
1 FG region

Example with noise

Noisy input image

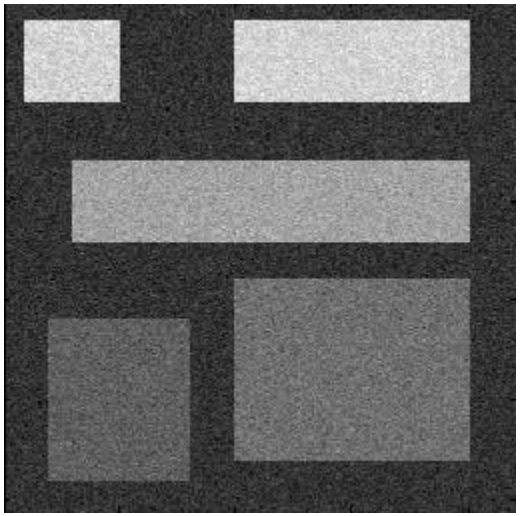

Desired segmentation

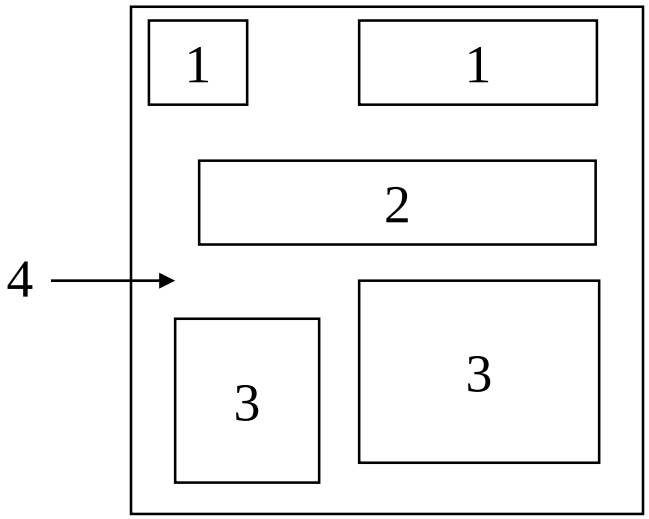

Where to threshold to get this segmentation?

Let's look at the histogram of the input image…

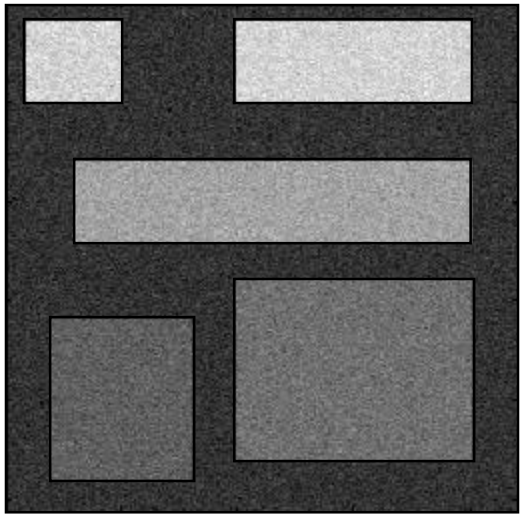

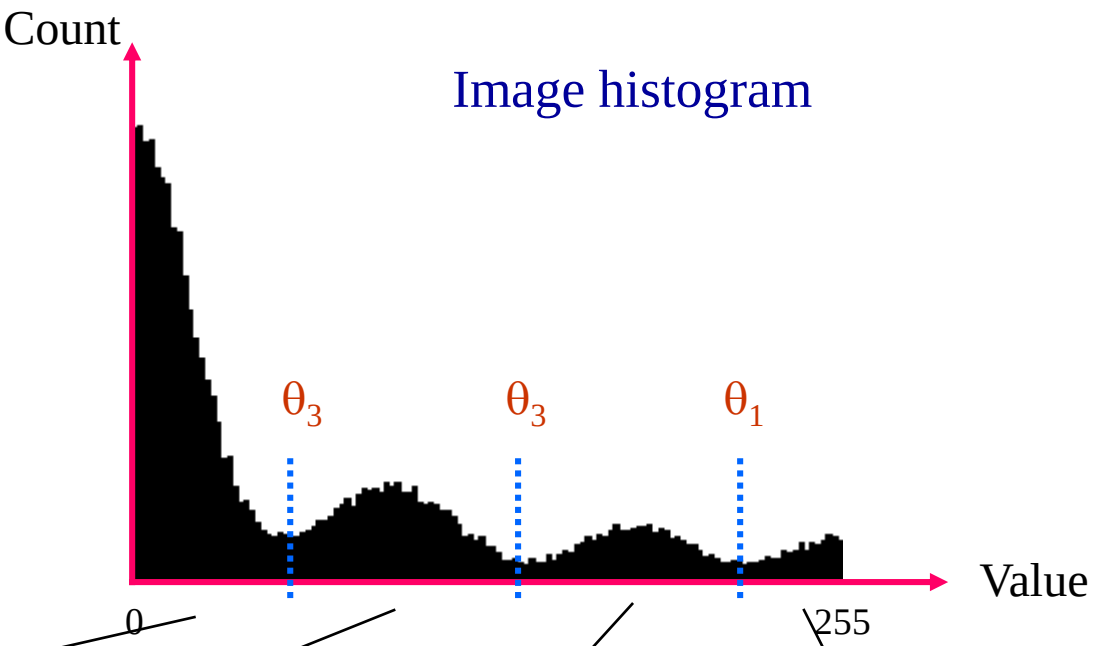

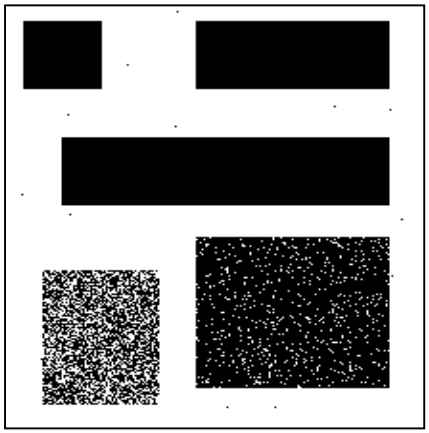

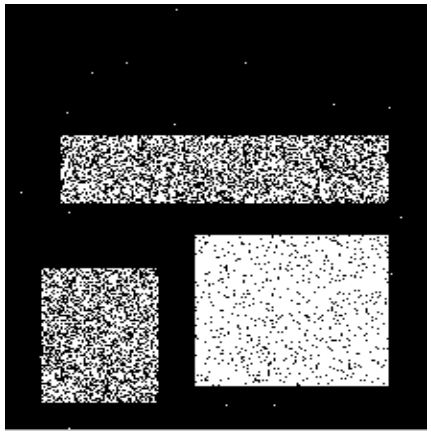

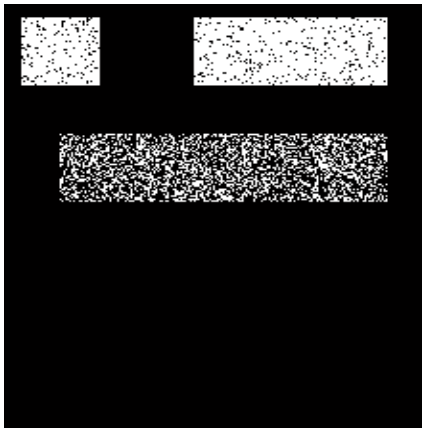

Color-based segmentation

- "Together" = similar color values
  - Color and intensity, or just color???
    - E.g., are "dark green" and "bright green" similar?

- Segment based on partitioning of color space
  - RGB, YUV, HSV, …?

- Several ways to model the color range of a region, including…

Color-based segmentation

- Color cube:
  - $r_{min} < R < r_{max}$ AND $g_{min} < G < g_{max}$ AND $b_{min} < B < b_{max}$

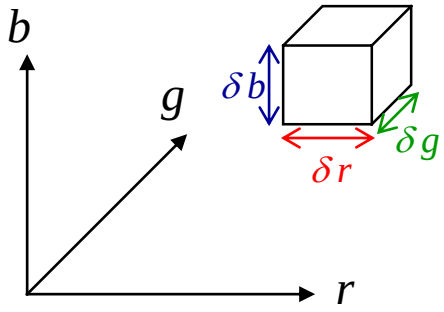

- Euclidian distance:
  - $d = \| (R,G,B) - (r_c, g_c, b_c) \|$
  - $d < d\theta$

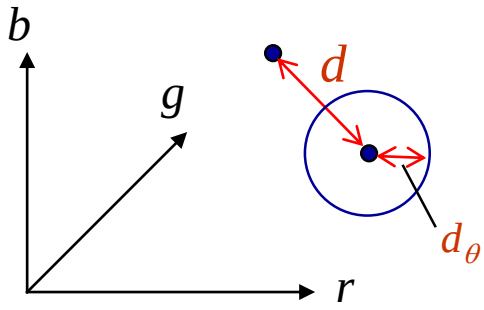

- Mahalanobis distance: Takes into account variance in all dimensions:
  - $d^2 = (x - x_m)^T C_x^{-1} (x - x_m)$
  - $d < d\theta$
    - x is the (R,G,B) vector
    - x_m is the mean of the class distribution
    - C_x is the covariance matrix of the distribution

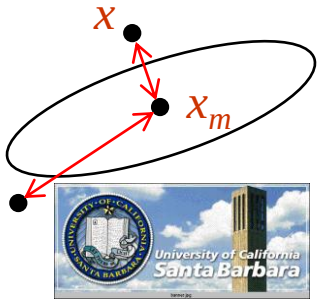

Color-based skin segmentation

Original

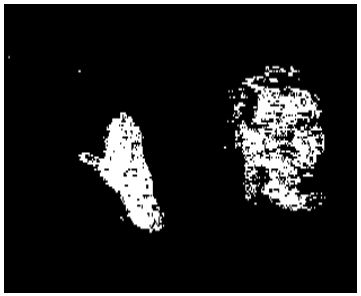

Threshold based on color

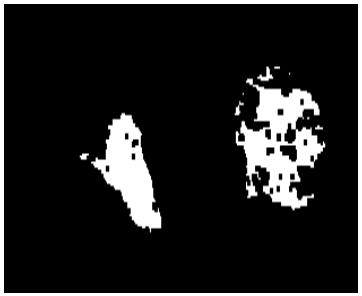

After morphological analysis

Original

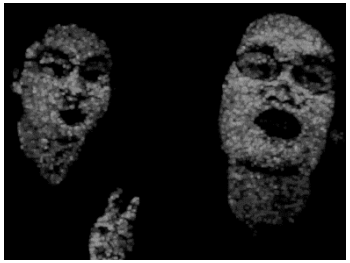

Threshold based on color

Texture painted back on face regions

May want to enforce **spatial** proximity as well as **color** proximity

Color segmentation examples

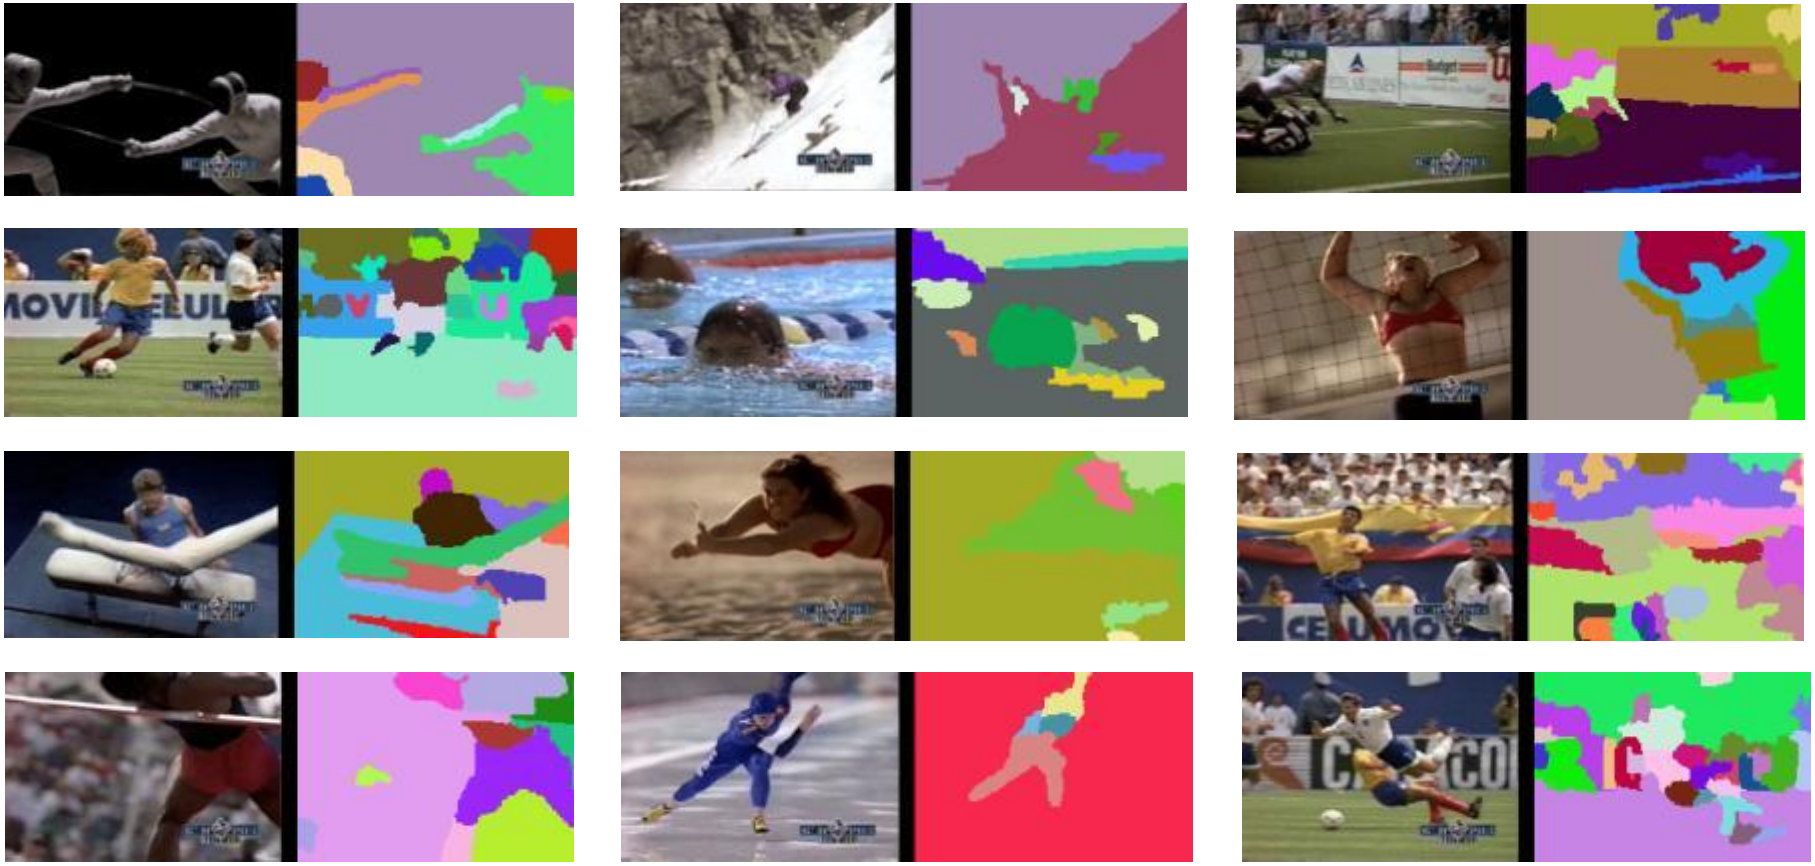

http://www.ee.columbia.edu/~dzhong/rtrack/demo.htm

Color segmentation examples

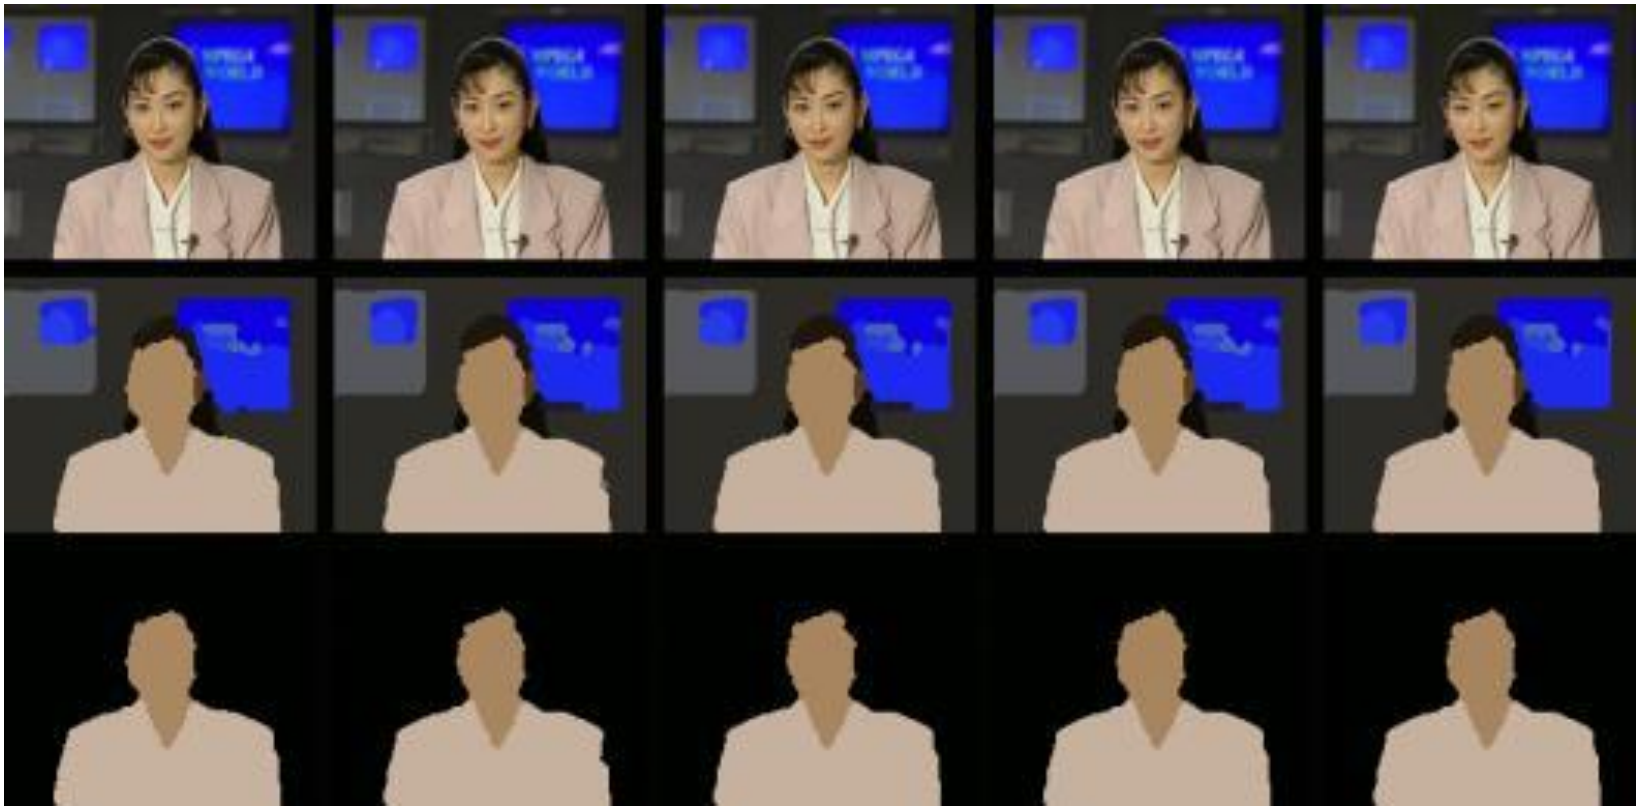

Texture-based segmentation

- "Together" = similar texture properties
  - Fundamentally an area-based measure, not a single pixel
- There is no single definition/measure of texture
  - Number of edge segments per unit area (magnitude, orientation…)
  - Frequency of occurrence of a pattern
  - Statistics, statistics, statistics
  - How large of an area defines a texture?

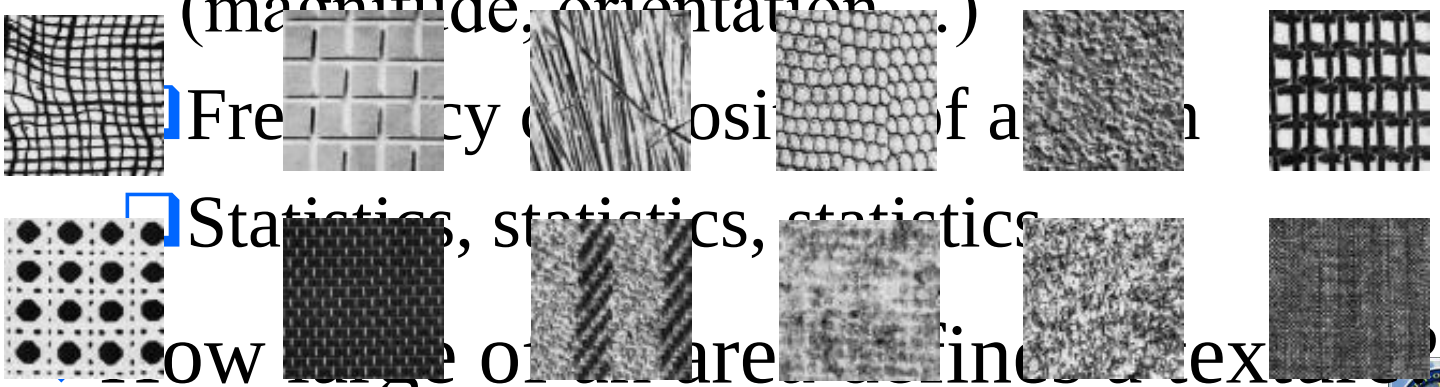

Brodatz textures

Texture-based segmentation

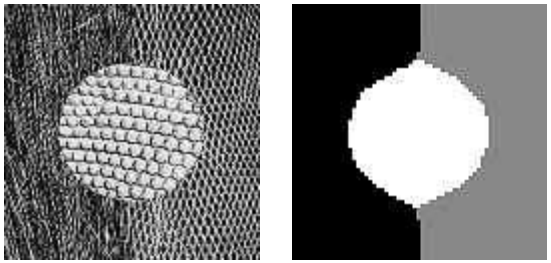

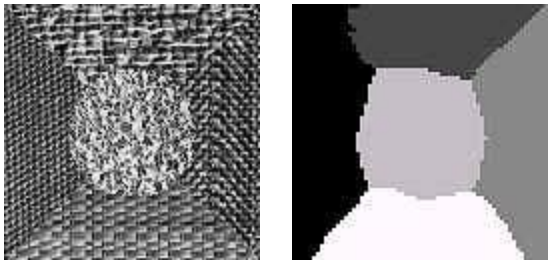

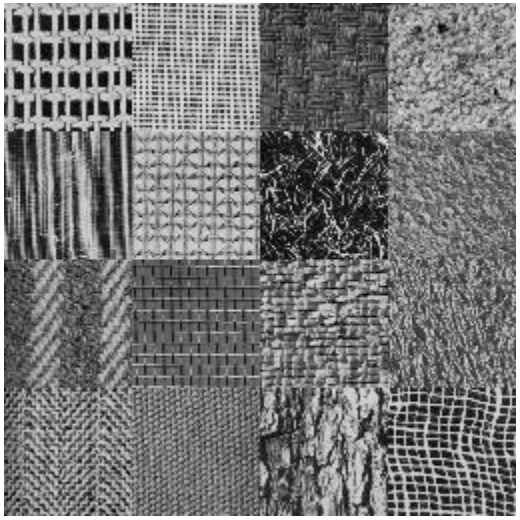

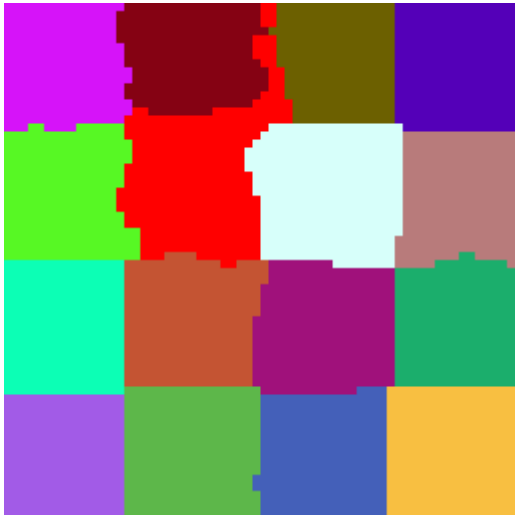

Texture-based segmentation

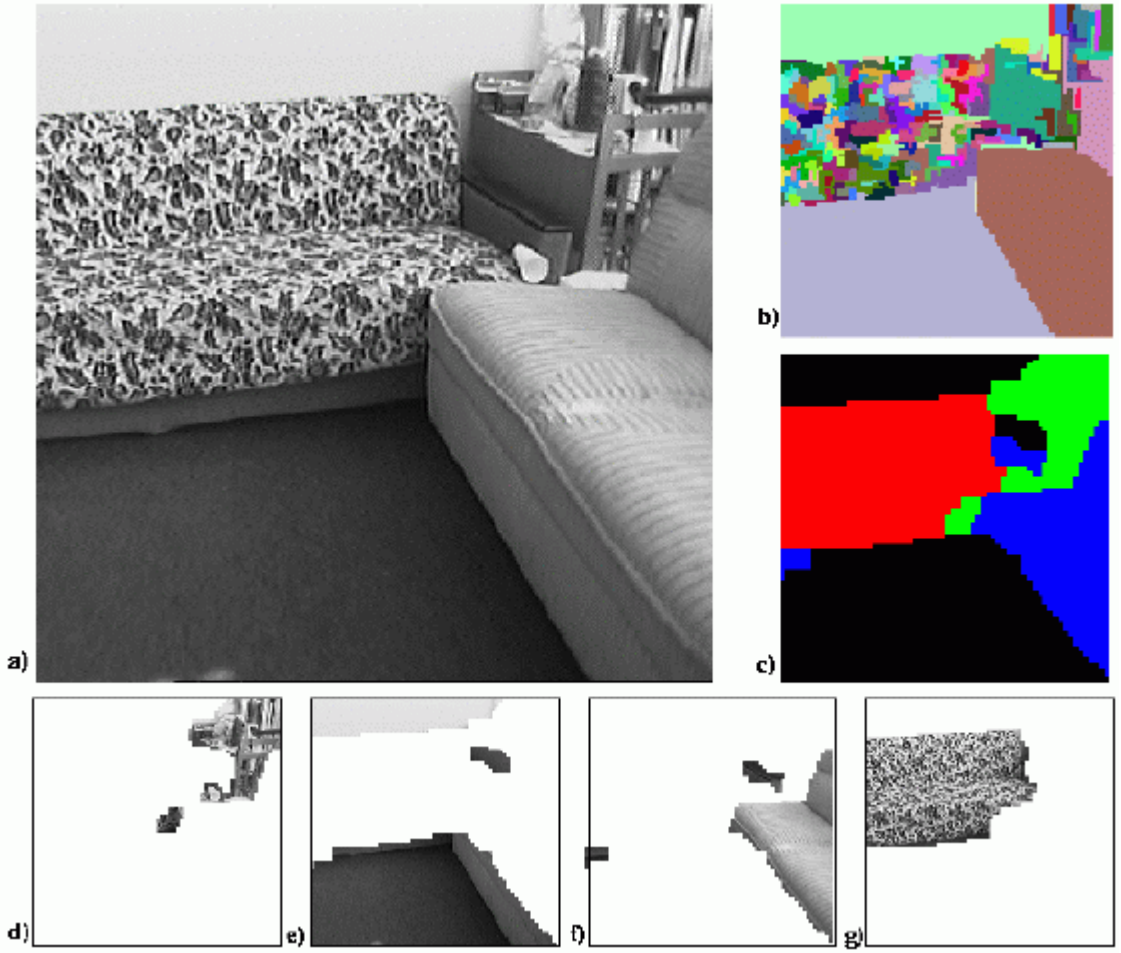

Texture segmentation

http://www-dbv.cs.uni-bonn.de/image/example6.html

Motion-based segmentation

- "Together" = similar motion
  - Rigid motion: all object points described by the same transformation
    - Pencils, coffee mugs, computer monitors, marbles, …
  - Non-rigid motion: articulated objects, bending objects, squishy objects…
    - Clouds, fluids, faces, hair, arms, scissors, …

Relative motion (depth)

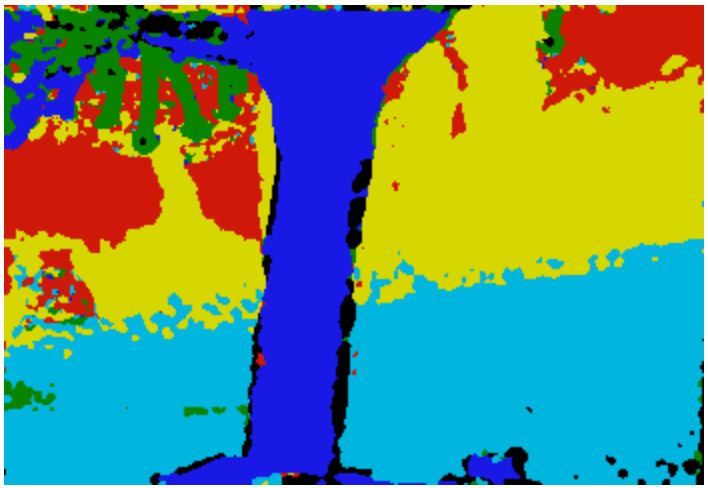

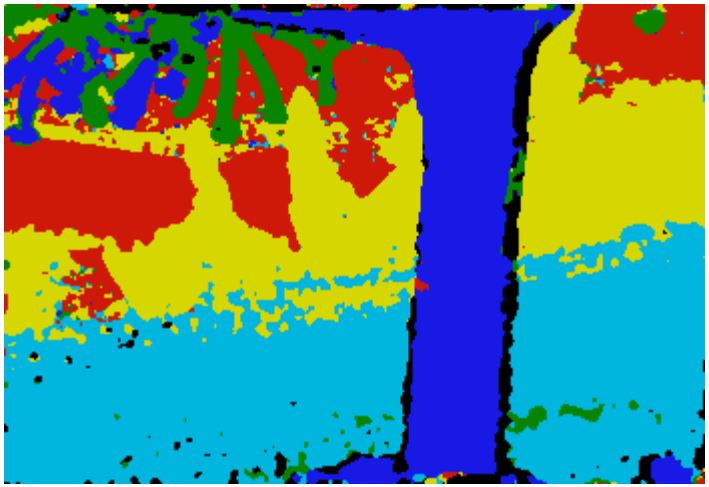

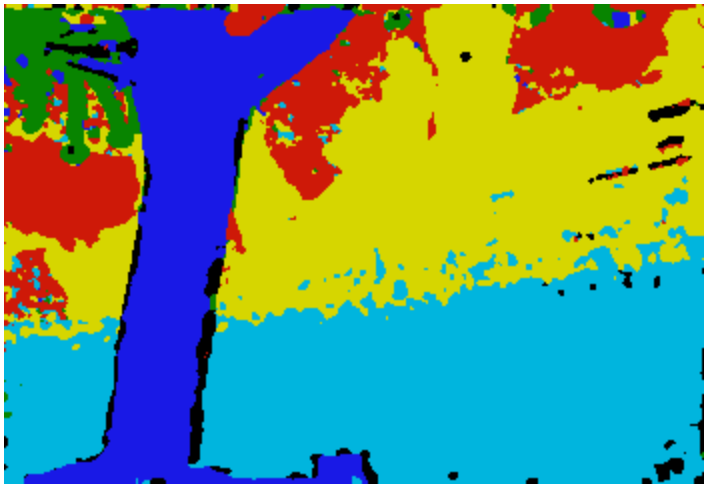

Segmenting two moving objects

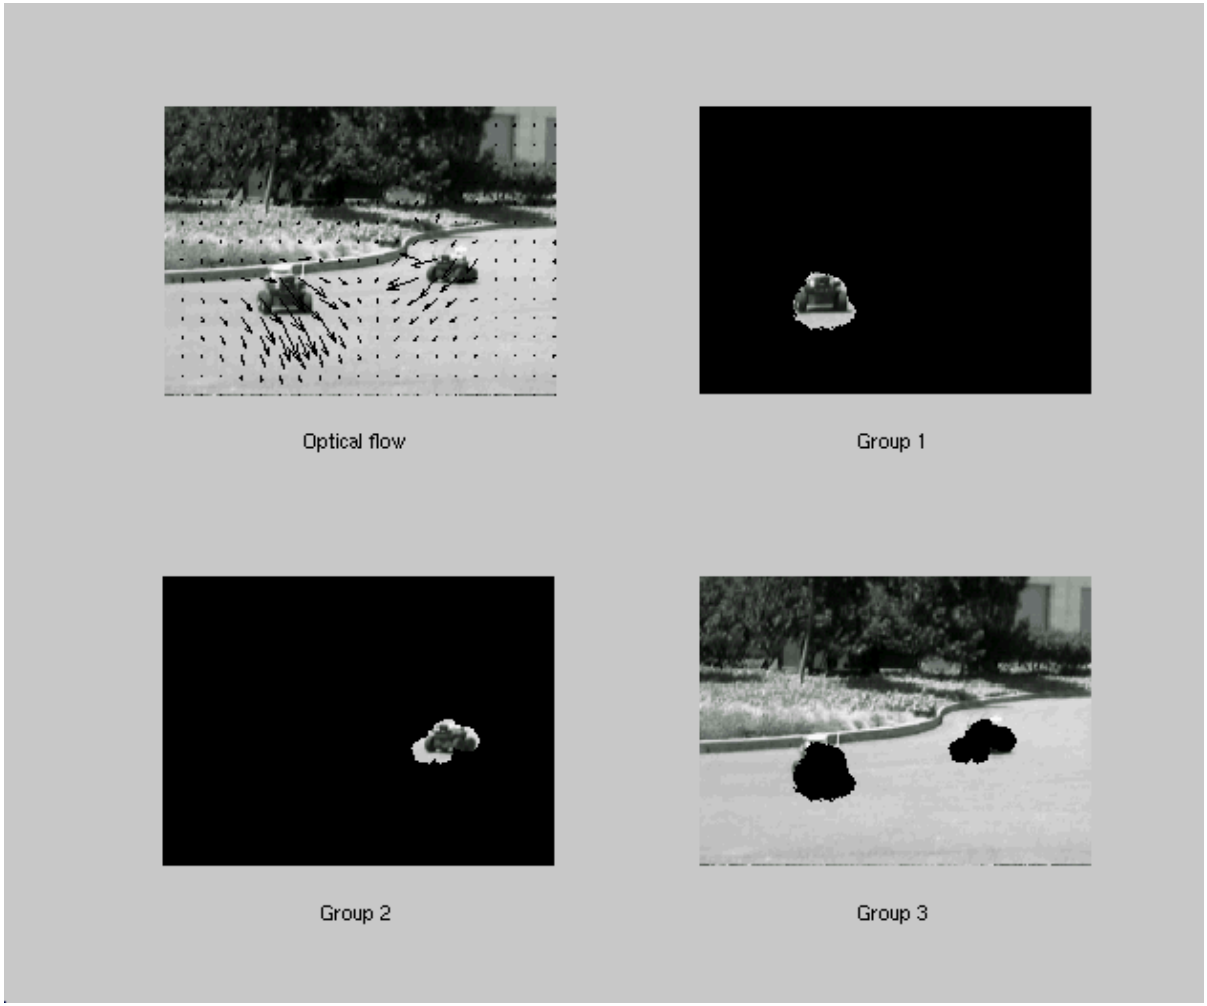

http://robotics.eecs.berkeley.edu/~rvidal/segment.html

Surveillance

(a)

(b) (c) (d)

Depth-based segmentation

- "Together" = similar depth (distance from reference)
  - How similar?
  - How to segment a large object? A wall?
- Or surface normal
- Or contiguous object

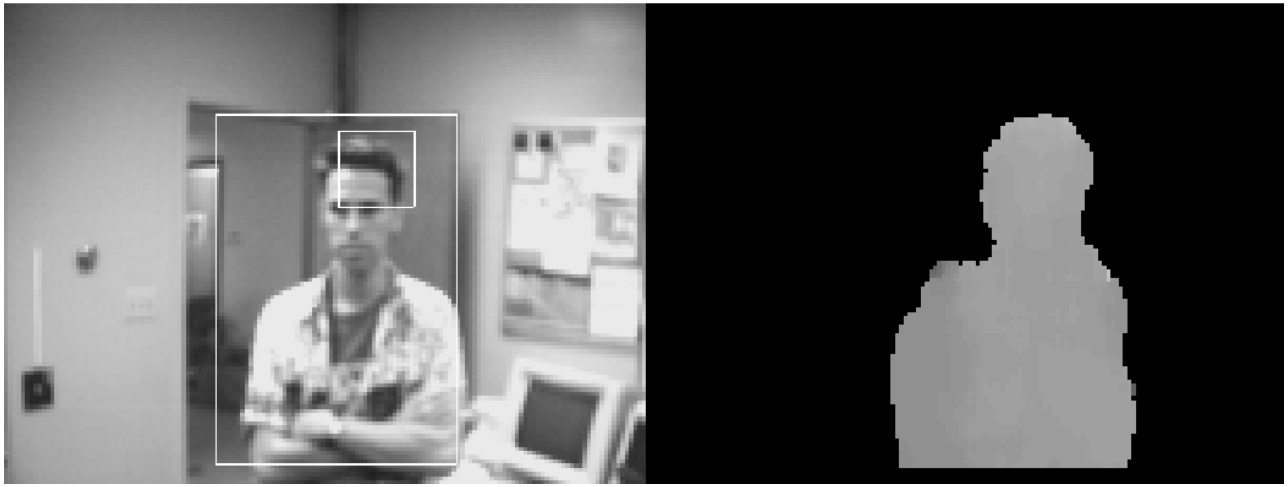

Segmentation: Background subtraction

- Goal: Separate the "foreground" from the "background" in the scene
  - Not necessarily related to depth

- Approach: Model the background, then detect significant changes from the model

Temporal segmentation of video

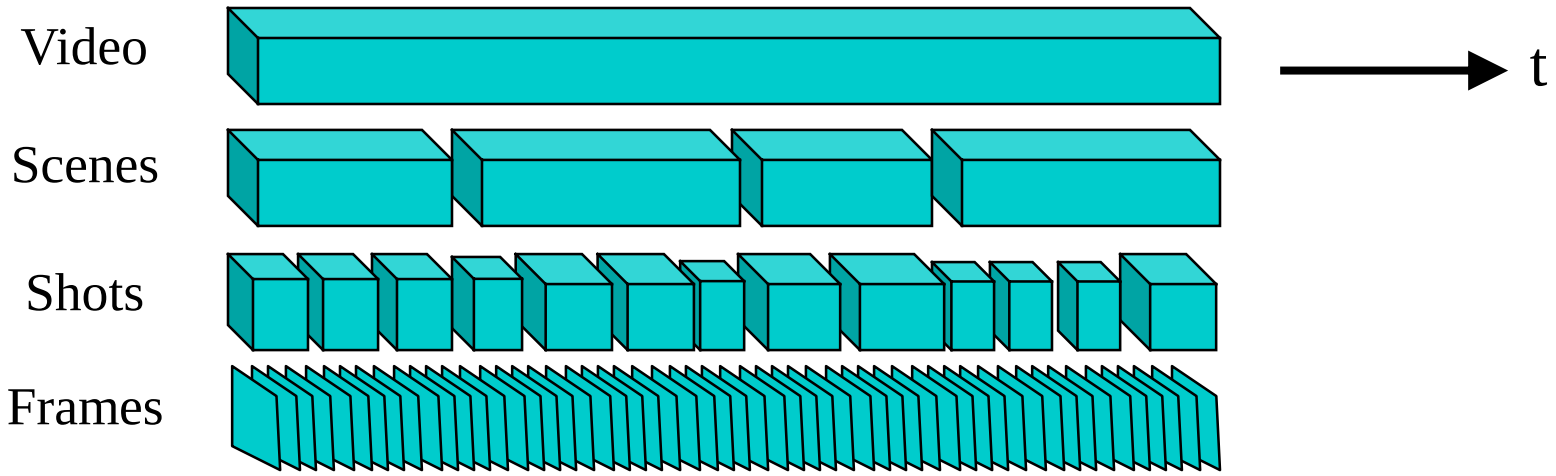

- Segment video into clips (shots) by looking for large changes
  - Overall frame-to-frame change (frame differencing)
    - Color, grayscale pixel values
  - Histogram change
    - Can be faster to compute