REGIONS

- Defining the problem
- Low-level algorithms
- Introducing semantics

REGION SEGMENTATION

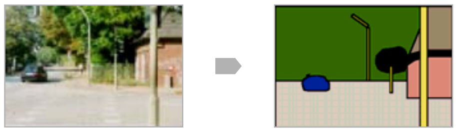

Ideal region: Set of pixels with the same statistical properties and corresponding to the same object.

IN PRACTICE

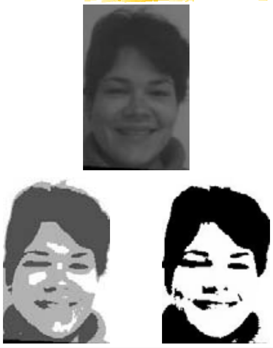

COLOR SEGMENTATION

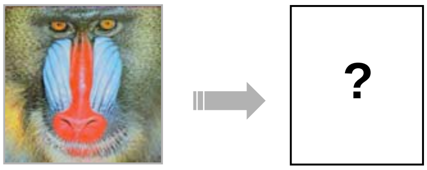

In general, regions can be formed from the original image data or from 'derived' images:

- color images from R, G, B
- textural images
- displacement images from motion analysis
- 3D depth images

CONTEXT IS ESSENTIAL

IN THEORY

Look for an image partition such that:

$$I = \bigcup_{i=1}^{m} S_i$$

$$S_i \bigcap S_j = 0, \forall i \neq j$$

$$H(S_i) = True, \forall i$$

$$H(S_i \bigcup S_j) = False, \text{ if } S_i \text{ and } S_j \text{ are adjacent.}$$

where H measures homogeneity.

GENERIC METHODS

Split:
- Start with a partition that satisfies Eq. 4.
- Split regions until they all satisfy Eq. 3.

Merge:
- Start with a partition that satisfies Eq. 3.
- Satisfy Eq. 4 by merging regions.

Homogeneity:
- Uniform gray-level or color statistics.
- Regions to which a parametric surface can be fitted.

LOW-LEVEL TECHNIQUES

- Histogram splitting.
- Region grouping.
- K-Means.
- Spectral methods.
- Relaxation.

HISTOGRAMS

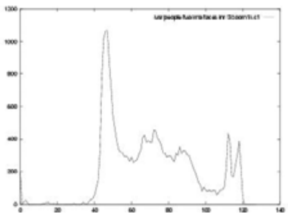

- Noisy histograms.
- Ill-defined boundaries.
- → High-level knowledge required.

HISTOGRAM SPLITTING

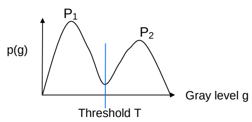

1. Compute image histogram.
2. Smooth histogram.
3. Look for peaks separated by deep valleys.
4. Group pixels into connex regions.
5. Smooth these regions
6. Iterate

RESULTS

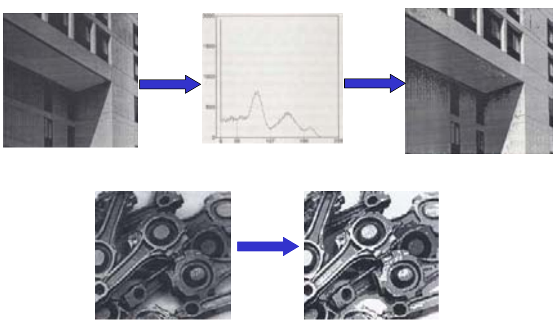

THRESHOLDS

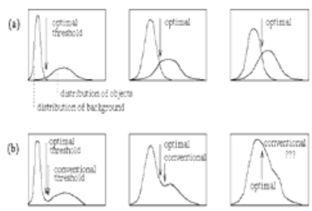

Probability distributions

Corresponding histograms

Choosing optimal thresholds is a difficult optimization problem.

HIERARCHICAL APPROACH

1. A first threshold is used to segment the dark pixels.
2. This yields two regions, the top half of the picture and the dark rectangle at the bottom.
3. The top half of the picture can now be easily segmented into two regions.

-> Decisions can be deferred until enough information becomes available.

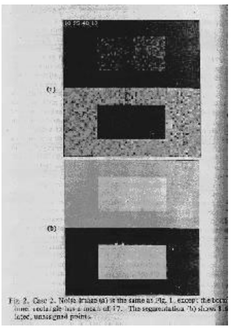

WEAKNESSES

- Histograms do not account for neighborhood relationships.
- Thresholds are hard to find.
- Some edges can have gray levels on both sides that belong to the same histogram peak.

NO VALID THRESHOLD!

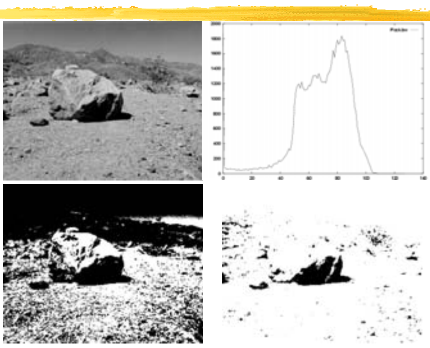

ILLUMINATION PROBLEMS

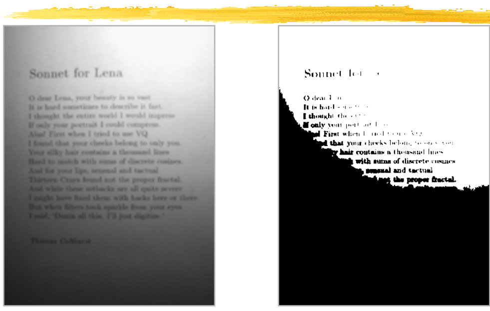

No global threshold -> Local ones are required.

LOCAL THRESHOLDING

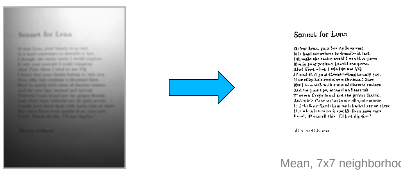

Mean, 7x7 neighborhood

- Examine statistically pixel values in local neighborhood around pixel to be thresholded.
- Use local statistic as threshold.
- Possibilities include mean, median, or mean of max and min value.

IMPROVED RESULTS

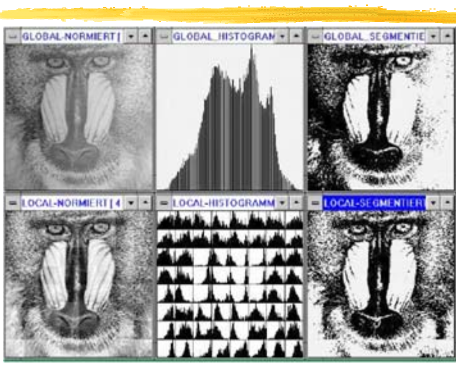

HOMOGENEOUS OR NOT?

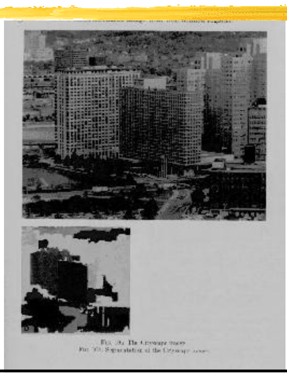

USING COLOR

RGB CHROMATICITY DIAGRAM

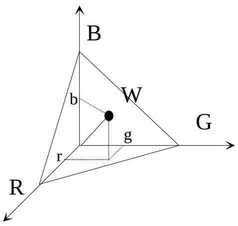

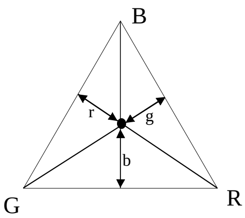

The Maxwell triangle involves projecting the colors in RGB space onto R+G+B=1 plane.

→ Chromaticity independently of luminance.

COLOR SPACE

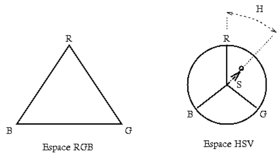

Hue 

Saturation 

Value 

HSV SPACE

Normalized Colors:

$$r = R/(R+G+B)$$

$$g = G/(R+G+B)$$

$$b = B/(R+G+B)$$

Hue/Saturation/Value:

$$I = R+G+B$$

$$S = 1 - \frac{3\min(R,G,B)}{I}$$

$$H = \arccos\left(\frac{0.5(2R-G-B)}{\sqrt{(R-G)^2+(R-B)(G-B)}}\right), \text{if } B < G$$

HSV IMAGES

IMPROVED HISTOGRAM

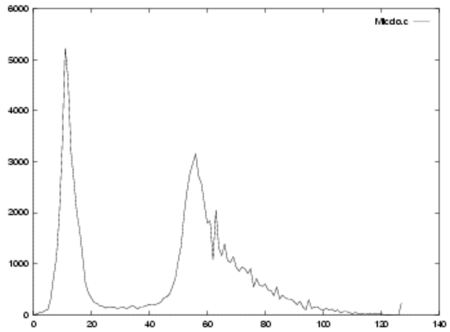

CIE LAB COLOR SPACE

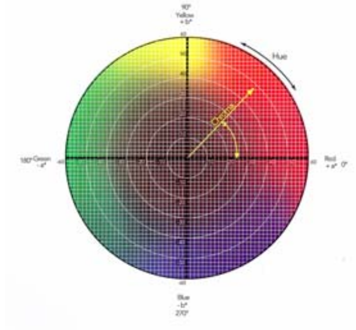

MULTIPLE HISTOGRAMS

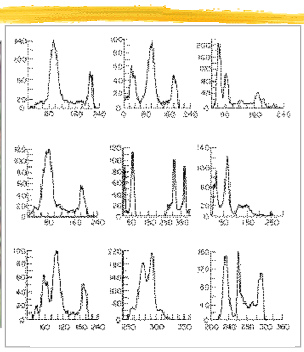

ALGORITHM

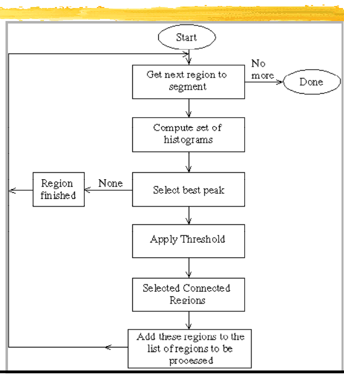

HIEARCHICAL SEGMENTATION

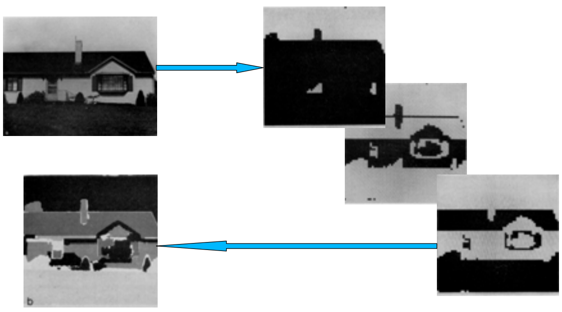

OHLANDER-PRICE

MERGING REGIONS

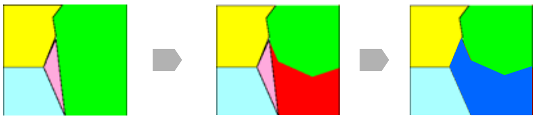

Split, merge, and split again on the basis of an homogeneity criterion.

RECURSIVE MERGING

1. Create an image partition.
2. Compute an adjacency graph.
3. For each image region:
   - Test its similarity with it neighbors.
   - Group the most similar ones.
4. Iterate until no more regions can be grouped.

STATISTICAL HOMOGENEITY

For each pair of neighboring regions:

- H0 hypothesis: The two regions correspond to the same object. --> Intensities are effectively described in terms of a single Gaussian distribution.
- H1 hypothesis: The two regions correspond to two different regions. Intensities must be described in terms of two distinct Gaussian distributions.
- Evaluate the probabilities of H0 and H1 and pick the best.

FISHER'S CRITERION

Discrimination between regions of different means and standard deviations can be done using

$$\frac{\mu_1 - \mu_2}{\sqrt{v_1^2 - v_2^2}} > \lambda$$

where λ is a threshold. If two regions have good separation in the means and low variance, then we can separate them.

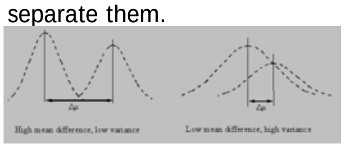

UNIFORMITY

The merging threshold for the mean intensity for two adjacent regions, λ, should vary depending on the expected uniformity of the merged region

- Less uniform regions will require a lower threshold to prevent under merging.
- Define uniformity as $1 - \nu / \mu$
- Take λ to be $(1 - \nu / \mu) \lambda_0$

→ User need supply only one threshold λ_0

PARAMETRIC HOMOGENEITY

Fit a 3—D surface to the region gray levels and evaluate the quality of this description.

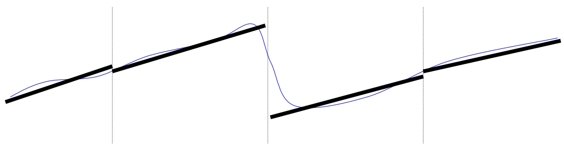

→ Merge if parametric surfaces are compatible.

RESULT

K-MEANS CLUSTERING

Assuming there are k regions and each one is described by a vector x_j, define:

- An objective function that measures the compactness of these regions.
- An optimization method.

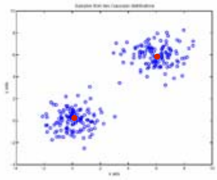

For a set of points in space, x_j is a coordinate vector. For black and white images, there is only one. For color images there are three.

OBJECTIVE FUNCTION

$$\Phi(\text{clusters}, \text{data}) = \sum_{i \in \text{clusters}} \left\{ \sum_{j \in i\text{'th cluster}} (\boldsymbol{x}_j - \boldsymbol{c}_i)^T (\boldsymbol{x}_j - \boldsymbol{c}_i) \right\}$$

If the allocation of points to clusters were known, we could compute the best centers easily.

But there are far too many combinations for an exhaustive search.

→ Define an algorithm that alternates
- Assume centers are known, allocate points
- Assume allocation is known, compute centers

ALGORITHM

1. Choose, for example randomly, k points that will serve as cluster centers.
2. Until their positions stabilize:
   - Associate each pixel to the cluster whose center is closest;
   - Recompute the centers by averaging the elements of each cluster.
3. Extract connected components.

K-MEANS CLUSTERING

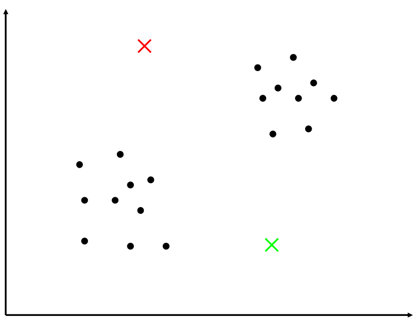

K-MEANS CLUSTERING

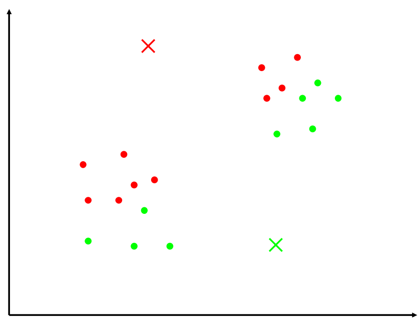

K-MEANS CLUSTERING

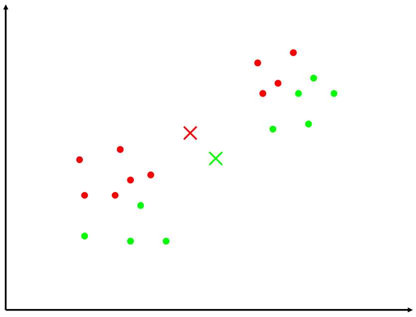

K-MEANS CLUSTERING

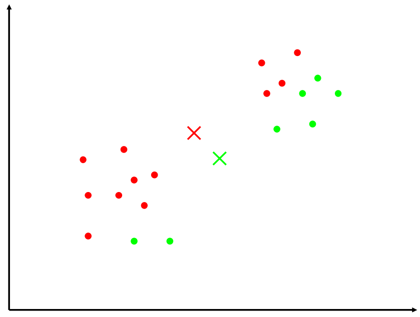

K-MEANS CLUSTERING

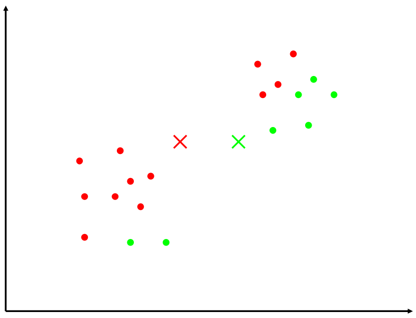

K-MEANS CLUSTERING

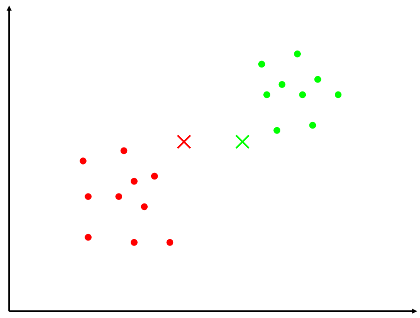

K-MEANS CLUSTERING

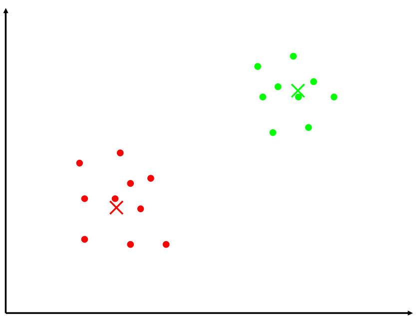

EXAMPLE (K=5)

Original Image | k-means on intensity | k-means on color

Each pixel is represented by the value of the cluster to which it belongs.

EXAMPLE (K=11)

Image

Clusters on color

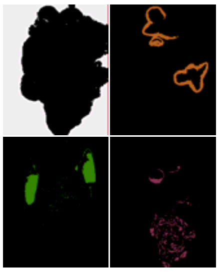

GRAPH THEORETIC CLUSTERING

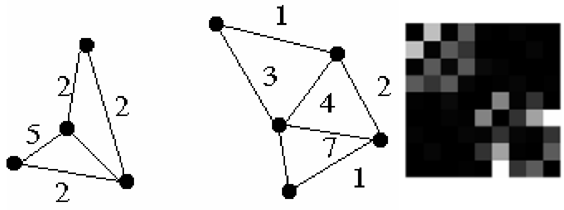

- Represent pixels as a weighted graph.
- Cut up this graph to get subgraphs with strong interior links

MEASURING AFFINITY

Intensity

$$aff(x,y)=\exp\left\{-\left(\frac{1}{2\sigma_i^2}\right)\left(\|I(x)-I(y)\|^2\right)\right\}$$

Distance

$$aff(x,y)=\exp\left\{-\left(\frac{1}{2\sigma_d^2}\right)\left(\|x-y\|^2\right)\right\}$$

Texture

$$aff(x,y)=\exp\left\{-\left(\frac{1}{2\sigma_t^2}\right)\left(\|c(x)-c(y)\|^2\right)\right\}$$

SCALE AFFECTS AFFINITY

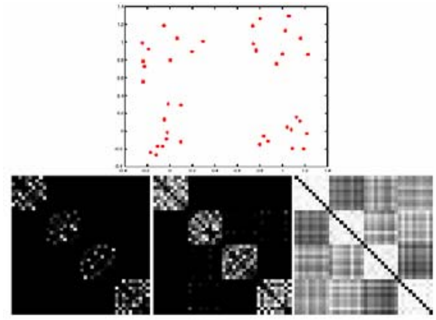

EIGENVECTORS OF THE AFFINITY MATRIX

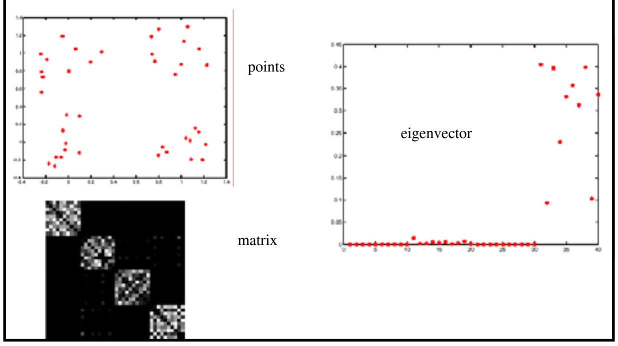

FINDING ONE GOOD CLUSTER

We want

A cluster in which elements have strong affinity with one another

A vector **a** that gives the association between each element and the cluster

We maximize $\mathbf{a}^t\mathbf{A}\mathbf{a}$ subject to $\mathbf{a}^t\mathbf{a}=1$

Solution is the eigenvector with the largest eigenvalue.

MORE THAN TWO SEGMENTS

Two options

- Recursively split each side to get a tree, continuing until the eigenvalues are too small.
- Use the other eigenvectors.

→Criterion evaluates within cluster similarity, but not across cluster difference.

NORMALIZED CUTS

Write graph as V, one cluster as A and the other as B and maximize

$$\left(\frac{assoc(A,A)}{assoc(A,V)}\right)+\left(\frac{assoc(B,B)}{assoc(B,V)}\right)$$

- Construct A, B such that their within cluster similarity is high compared to their association with the rest of the graph
- Maximize the within cluster similarity compared to the across cluster difference

NORMALIZED CUTS

1. Write a vector y whose elements are 1 if item is in A, -b if it's in B.
2. Write the matrix of the graph as W, and the matrix which has the row sums of W on its diagonal as D, 1 is the vector with all ones.

-> Criterion becomes $\min_y \left(\frac{y^T(D-W)y}{y^T Dy}\right)$

under the constraint $y^T D1 = 0$

->This is hard to do, because y's values are quantized

NORMALIZED CUTS

Instead, solve the generalized eigenvalue problem

$$\max_y \left(y^T (D-W)y\right) \text{ subject to } \left(y^T Dy = 1\right)$$

which gives $(D-W)y = \lambda Dy$

Now look for a quantization threshold that maximizes the criterion, that is, all components of y above that threshold go to one, all below go to -b

NORMALIZED CUTS RESULTS

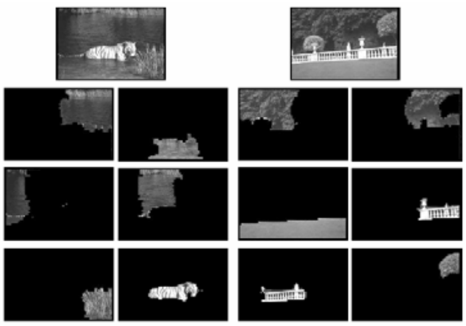

Shi and Malik, 1998

NORMALIZED CUTS RESULTS

Shi and Malik, 2000

RELAXATION

Assign label list
- Candidate list
- Confidence rating

Compare
- Inhibit or reinforce

Works locally to produce a global result

SIMULATED ANNEALING

Stochastic Relaxation

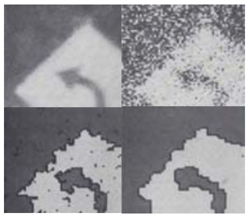

High Temperature
- Distributions are uniform

Low Temperature
- Focus on decreasing objective function

SEGMENTATION

Purely statistical methods give extract useful information but are inherently limited. One must also take into account:

- Region outlines.
- Region Shape.
- Context.

-> Yet another open problem!

SUPERVISED CLASSIFICATION

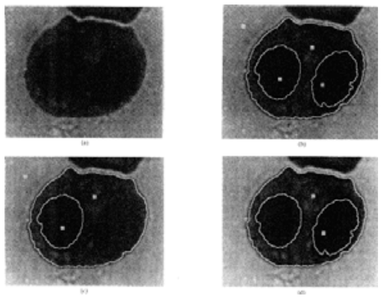

Interactive Segmentation of a Cell

REGION GROWING

<u>While</u> SSL is not empty <u>do</u>
Remove first pixel y from SSL.
Test neighbors of this point:
<u>If</u> all already labeled neighbors of y (other than boundary pixels) have the same label
<u>then</u>
Set y to this label.
Update running mean of corresponding region.
Add neighbors of y which are neither already set nor already in the SSL to the SSL according to their value of delta
<u>else</u>
Flag y as a boundary pixel.
fi

LIMITATIONS

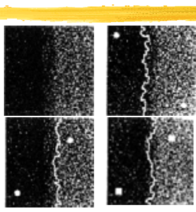

- In general, the result depends on the order in which the pixels are taken into consideration.
- La homogeneity measure is noise sensitive.

INTRODUCING SEMANTICS

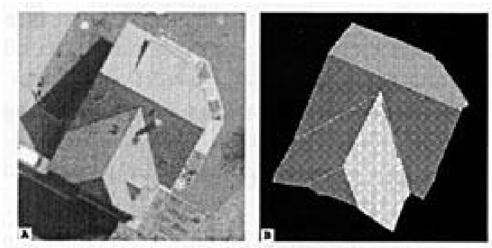

Segmentation of a complex roof with many sides

FROM ROOF EDGES TO ROOF PARTS

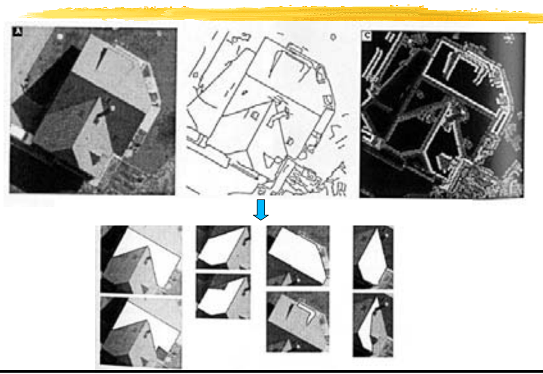

AMBIGUOUS INTERPRETATIONS

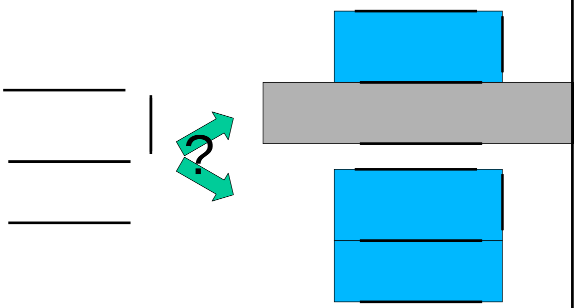

SEGMENTATION AS A GRAPH SEARCH PROBLEM

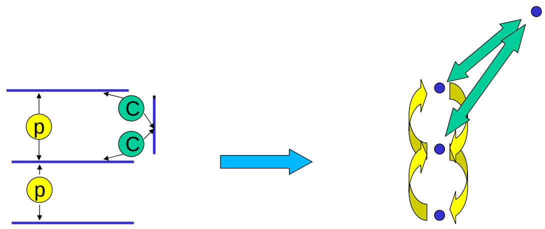

Finding candidate regions amounts to finding cycles in the graph → Can use graph-search techniques to handle the combinatorics.

COMBINING EDGE AND REGION INFORMATION

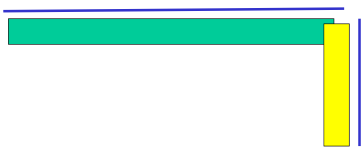

Check photometric consistency before associating edges to prune the graph.

ALGORITHM

1. Edge extraction
2. Photometric and chromatic analysis
3. Grouping on an homogeneity basis
4. Stereo
5. Coplanar grouping
6. Selection of compatible hypothesis

→ Ok for single houses but combinatorial explosion in a dense urban environment.