

Segmentation & Grouping

Kristen Grauman, UT Austin

Last time

- Fitting an arbitrary shape with “active” deformable contours

Review questions


- How does Hough fitting and deformable contour fitting differ? How are they alike?
- What is the influence of the number of the vertices in an active contour?
- What is the influence in the number of candidate states (m) when fitting the active contour with DP?

Last time: deformable contours

- Representation of the contours
- Defining the energy functions
 - External
 - Internal
- Minimizing the energy function
- Extensions:
 - Tracking
 - Interactive segmentation


Tracking via deformable contours

1. Use final contour/model extracted at frame t as an initial solution for frame $t+1$
2. Evolve initial contour to fit exact object boundary at frame $t+1$
3. Repeat, initializing with most recent frame.



Tracking Heart Ventricles
(multiple frames)

3D active contours



Jørgen Ahlberg
<http://www.cvl.iui.uio.se/OutMaster/Papers/EXT1708.pdf>

Limitations

- May over-smooth the boundary
- Cannot follow topological changes of objects

Limitations

- External energy: snake does not really "see" object boundaries in the image unless it gets very close to it.

image gradients ∇I are large only directly on the boundary

Distance transform

- External image can instead be taken from the **distance transform** of the edge image.

Value at (x,y) tells how far that position is from the nearest edge point (or other binary image structure)

>> `help bwdist`

Slide credit: Kristen Grauman

Interactive forces

How can we implement such an *interactive* force with deformable contours?

Slide credit: Kristen Grauman

Interactive forces

- An energy function can be altered online based on user input – use the cursor to push or pull the initial snake away from a point.
- Modify external energy term to include:

$$E_{push} = \sum_{i=0}^{n-1} \frac{r^2}{|v_i - p|^2}$$

Nearby points get pushed hardest

Intelligent scissors

Another form of interactive segmentation:

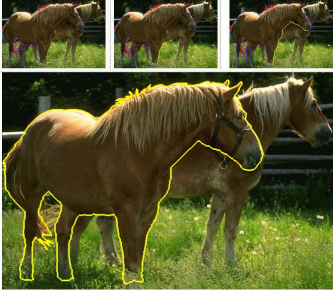
Compute optimal paths from **every point to the seed** based on edge-related costs.

VIDEO

Figure 2: Image demonstrating how the live-wire segment adapts and snaps to an object boundary as the free point moves (via cursor movement). The path of the free point is shown in white. Live-wire segments from previous free point positions (p_1 and p_2) are shown in green.

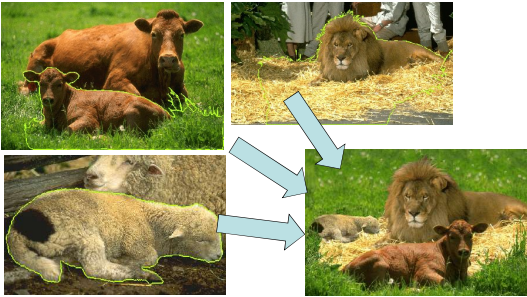
[Mortensen & Barrett, SIGGRAPH 1995, CVPR 1999]

Intelligent scissors



• <http://ivit.cs.byu.edu/Eric/Eric.html>

Intelligent scissors



• <http://ivit.cs.byu.edu/Eric/Eric.html>

Deformable contours: pros and cons

Pros:

- Useful to track and fit non-rigid shapes
- Contour remains connected
- Possible to fill in "subjective" contours
- Flexibility in how energy function is defined, weighted.

Cons:

- Must have decent initialization near true boundary, may get stuck in local minimum
- Parameters of energy function must be set well based on prior information

Recap: deformable contours

- Deformable shapes and active contours are useful for
 - Segmentation: fit or "snap" to boundary in image
 - Tracking: previous frame's estimate serves to initialize the next
- Fitting active contours:
 - Define terms to encourage certain shapes, smoothness, low curvature, push/pulls, ...
 - Use weights to control relative influence of each component cost
 - Can optimize 2d snakes with Viterbi algorithm.
- Image structure (esp. gradients) can act as attraction force for *interactive* segmentation methods.

Slide credit: Kristen Grauman

Outline

- What are grouping problems in vision?
- Inspiration from human perception
 - Gestalt properties
- Bottom-up segmentation via clustering
 - Algorithms:
 - Mode finding and mean shift: k-means, mean-shift
 - Graph-based: normalized cuts
 - Features: color, texture, ...
 - Quantization for texture summaries

Grouping in vision

- Goals:
 - Gather features that belong together
 - Obtain an intermediate representation that compactly describes key image or video parts

Examples of grouping in vision

Determine image regions

Group video frames into shots

Figure-ground

Object-level grouping

Slide credit: Kristen Grauman

Grouping in vision

- Goals:
 - Gather features that belong together
 - Obtain an intermediate representation that compactly describes key image (video) parts
- Top down vs. bottom up **segmentation**
 - Top down: pixels belong together because they are from the same object
 - Bottom up: pixels belong together because they look similar
- Hard to measure success
 - What is interesting depends on the app.

What are meta-cues for grouping?

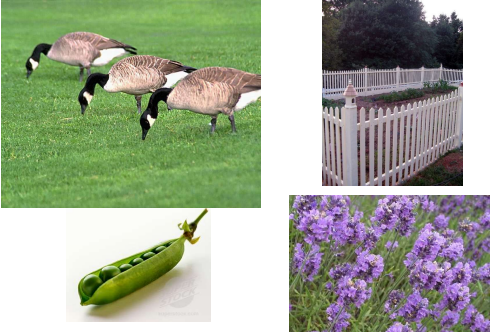
Muller-Lyer illusion

What things should be grouped?
 What cues indicate groups?

Gestalt

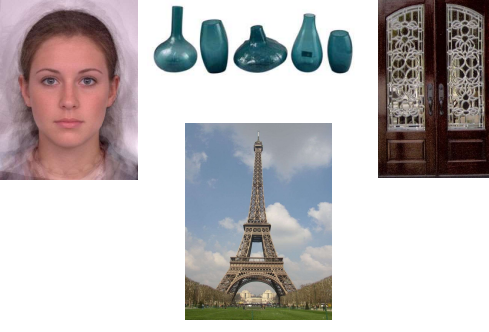
- Gestalt: whole or group
 - Whole is greater than sum of its parts
 - Relationships among parts can yield new properties/features
- Psychologists identified series of factors that predispose set of elements to be grouped (by human visual system)

Similarity



Slide credit: Kristen Grauman CC-BY-SA, <http://www.dailymotion.com/video/x2231532?view=Comp&paperDoc=15328-031.jpg>

Symmetry



Slide credit: Kristen Grauman in the processingfm site

Common fate





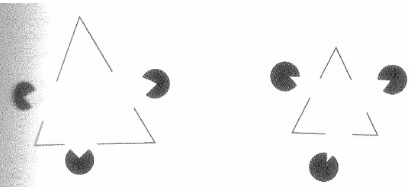
Image credit: Arthus-Bertrand (via F. Durand)

Proximity



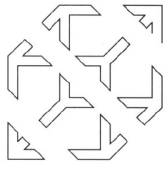
Slide credit: Kristen Grauman 035.jpg


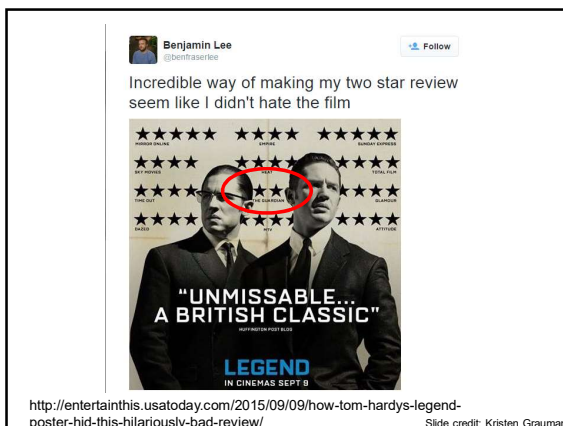
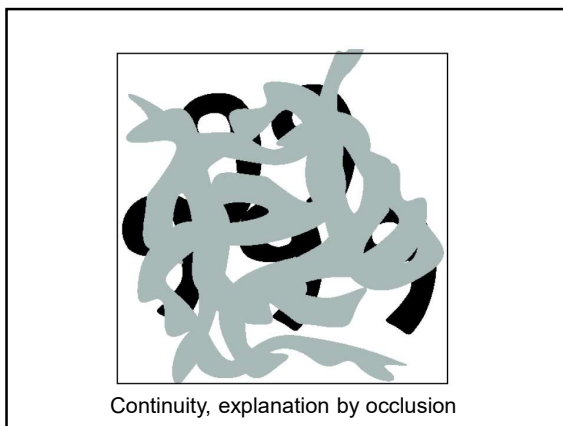

Illusory/subjective contours

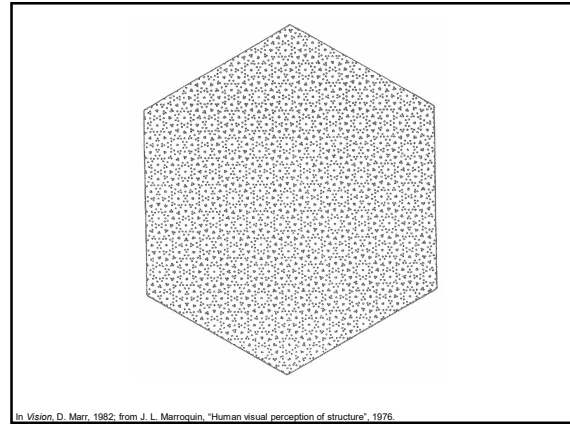



Interesting tendency to explain by occlusion

In Vision, D. Marr, 1982







Grouping phenomena in real life

Forsyth & Ponce, Figure 14.7

Grouping phenomena in real life

Forsyth & Ponce, Figure 14.7

Gestalt


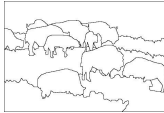

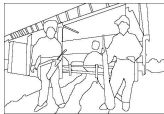
- Gestalt: whole or group
 - Whole is greater than sum of its parts
 - Relationships among parts can yield new properties/features
- Psychologists identified series of factors that predispose set of elements to be grouped (by human visual system)
- Inspiring observations/explanations; challenge remains how to best map to algorithms.

Outline

- What are grouping problems in vision?
- Inspiration from human perception
 - Gestalt properties
- Bottom-up segmentation via clustering
 - Algorithms:
 - Mode finding and mean shift: k-means, EM, mean-shift
 - Graph-based: normalized cuts
 - Features: color, texture, ...
 - Quantization for texture summaries

The goals of segmentation

Separate image into coherent "objects"

image	human segmentation
	
	

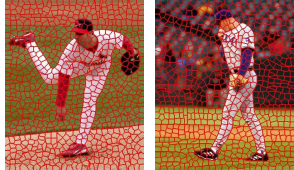
Source: Lana Lazebnik

The goals of segmentation

Separate image into coherent "objects"

Group together similar-looking pixels for efficiency of further processing

"superpixels"




X. Ren and J. Malik. [Learning a classification model for segmentation](#). ICCV 2003.

Source: Lana Lazebnik

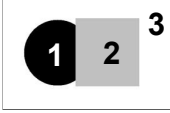
Clustering

- Clustering algorithms:
 - Unsupervised learning
 - Detect patterns in unlabeled data
 - E.g. group emails or search results
 - E.g. find categories of customers
 - E.g. group pixels into regions
- Useful when don't know what you're looking for
- Requires data, but no labels

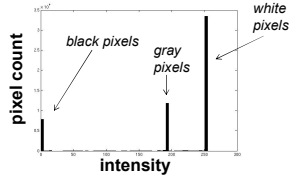


Slide credit: Dan Klein

Image segmentation: toy example



input image

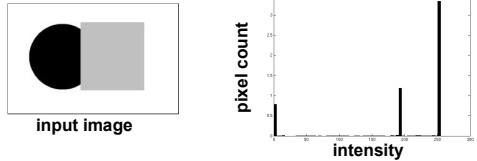


pixel count

intensity

- These intensities define the three groups.
- We could label every pixel in the image according to which of these primary intensities it is.
 - i.e., *segment* the image based on the intensity feature.
- What if the image isn't quite so simple?

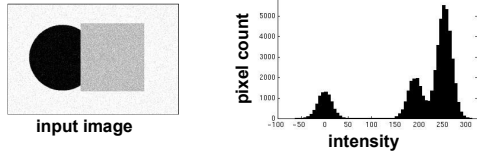
Slide credit: Kristen Grauman



input image

pixel count

intensity

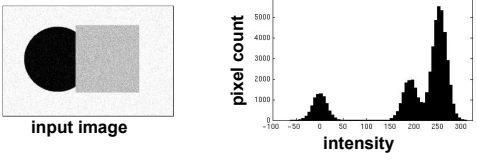


input image

pixel count

intensity

Slide credit: Kristen Grauman



input image

pixel count

intensity

- Now how to determine the three main intensities that define our groups?
- We need to *cluster*.

Slide credit: Kristen Grauman

0 intensity 190 255

- Goal: choose three “centers” as the **representative** intensities, and label every pixel according to which of these centers it is nearest to.
- Best cluster centers are those that minimize SSD between all points and their nearest cluster center c_i :

$$\sum_{\text{clusters } i} \sum_{\text{points } p \text{ in cluster } i} \|p - c_i\|^2$$

Slide credit: Kristen Grauman

Clustering

- With this objective, it is a “chicken and egg” problem:
 - If we knew the **cluster centers**, we could allocate points to groups by assigning each to its closest center.
 - If we knew the **group memberships**, we could get the centers by computing the mean per group.

Slide credit: Kristen Grauman

K-means clustering

- Basic idea: randomly initialize the k cluster centers, and iterate between the two steps we just saw.
 1. Randomly initialize the cluster centers, c_1, \dots, c_k
 2. Given cluster centers, determine points in each cluster
 - For each point p , find the closest c_i . Put p into cluster i
 3. Given points in each cluster, solve for c_i
 - Set c_i to be the mean of points in cluster i
 4. If c_i have changed, repeat Step 2

Properties

- Will always converge to *some* solution
- Can be a “local minimum”
 - does not always find the global minimum of objective function:

$$\sum_{\text{clusters } i} \sum_{\text{points } p \text{ in cluster } i} \|p - c_i\|^2$$

Source: Steve Seitz

K-Means

- An iterative clustering algorithm
 - Pick K random points as cluster centers (means)
 - Alternate:
 - Assign data instances to closest mean
 - Assign each mean to the average of its assigned points
 - Stop when no points' assignments change

K-means

1. Ask user how many clusters they'd like. (e.g. $k=5$)

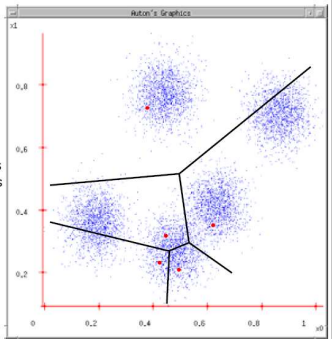
Slide credit Andrew Moore

K-means

1. Ask user how many clusters they'd like. (e.g. $k=5$)
2. Randomly guess k cluster center locations

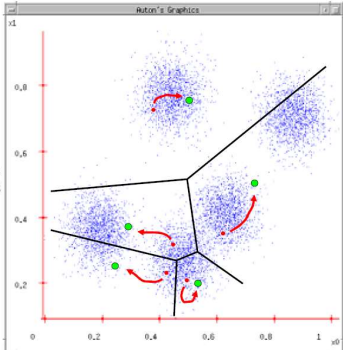
K-means

1. Ask user how many clusters they'd like. (e.g. $k=5$)
2. Randomly guess k cluster Center locations
3. Each datapoint finds out which Center it's closest to. (Thus each Center "owns" a set of datapoints)



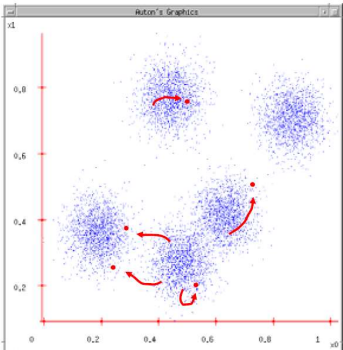
K-means

1. Ask user how many clusters they'd like. (e.g. $k=5$)
2. Randomly guess k cluster Center locations
3. Each datapoint finds out which Center it's closest to.
4. Each Center finds the centroid of the points it owns



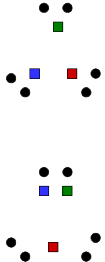
K-means

1. Ask user how many clusters they'd like. (e.g. $k=5$)
2. Randomly guess k cluster Center locations
3. Each datapoint finds out which Center it's closest to.
4. Each Center finds the centroid of the points it owns...
5. ...and jumps there
6. ...Repeat until terminated!



Initialization


- K-means is non-deterministic
 - Requires initial means
 - It does matter what you pick!
- What can go wrong?
- Various schemes for preventing this kind of thing



Slide credit: Dan Klein

K-Means Getting Stuck

- A local optimum:



Slide credit: Dan Klein

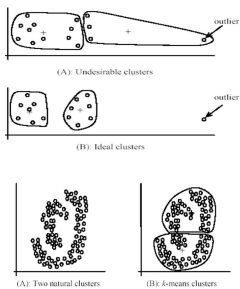
K-means: pros and cons

Pros

- Simple, fast to compute
- Converges to local minimum of within-cluster squared error

Cons/issues

- Setting k ?
- Sensitive to initial centers
- Sensitive to outliers
- Detects spherical clusters
- Assuming means can be computed



K-Means Questions

- Will K-means converge?
 - To a global optimum?
- Will it always find the true patterns in the data?
 - If the patterns are very very clear?
- Will it find something interesting?
- How many clusters to pick?
- Do people ever use it?

Slide credit: Dan Klein

Probabilistic clustering

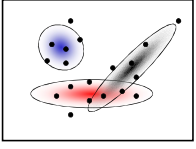
Basic questions

- what's the probability that a point x is in cluster m ?
- what's the shape of each cluster?

K-means doesn't answer these questions

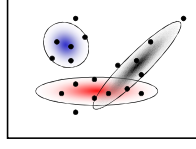
Probabilistic clustering (basic idea)

- Treat each cluster as a Gaussian density function



Slide credit: Steve Seitz

Expectation Maximization (EM)



A probabilistic variant of K-means:

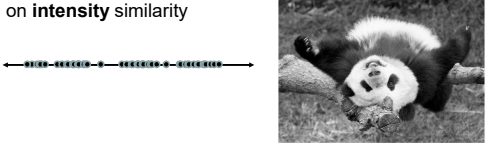
- E step: "soft assignment" of points to clusters
 - estimate probability that a point is in a cluster
- M step: update cluster parameters
 - mean and variance info (covariance matrix)
- maximizes the likelihood of the points given the clusters

Slide credit: Steve Seitz

Segmentation as clustering

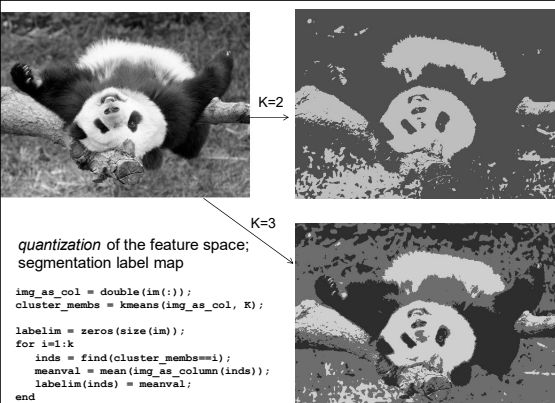
Depending on what we choose as the *feature space*, we can group pixels in different ways.

Grouping pixels based on **intensity** similarity



Feature space: intensity value (1-d)

Slide credit: Kristen Grauman



quantization of the feature space; segmentation label map

```

img_as_col = double(im(:));
cluster_membs = kmeans(img_as_col, K);

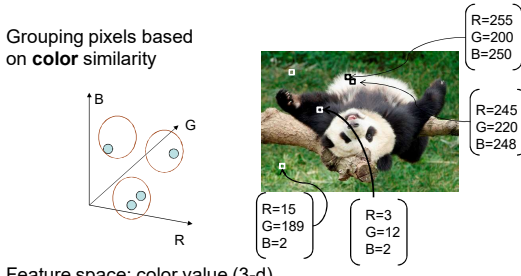
labelim = zeros(size(im));
for i=1:k
    inds = find(cluster_membs==i);
    meanval = mean(img_as_column(inds));
    labelim(inds) = meanval;
end
    
```

Slide credit: Kristen Grauman

Segmentation as clustering

Depending on what we choose as the *feature space*, we can group pixels in different ways.

Grouping pixels based on **color** similarity





Feature space: color value (3-d)

Slide credit: Kristen Grauman

Segmentation as clustering

Depending on what we choose as the *feature space*, we can group pixels in different ways.

Grouping pixels based on **intensity** similarity

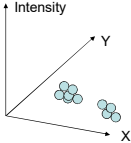
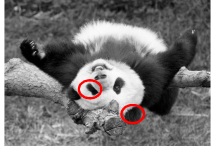
Clusters based on intensity similarity don't have to be spatially coherent.

Slide credit: Kristen Grauman

Segmentation as clustering

Depending on what we choose as the *feature space*, we can group pixels in different ways.

Grouping pixels based on **intensity+position** similarity





Both regions are black, but if we also include **position (x,y)**, then we could group the two into distinct segments; way to encode both similarity & proximity.

Slide credit: Kristen Grauman

Segmentation as clustering

- Color, brightness, position alone are not enough to distinguish all regions...

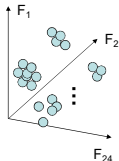



Slide credit: Kristen Grauman

Segmentation as clustering

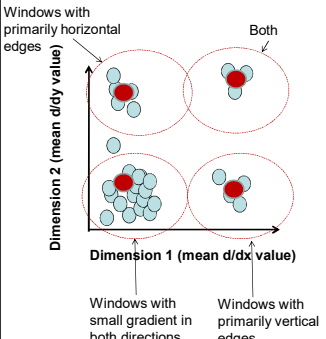
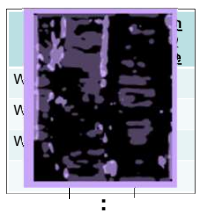
Depending on what we choose as the *feature space*, we can group pixels in different ways.

Grouping pixels based on **texture** similarity

Feature space: filter bank responses (e.g., 24-d)

Recall: texture representation example

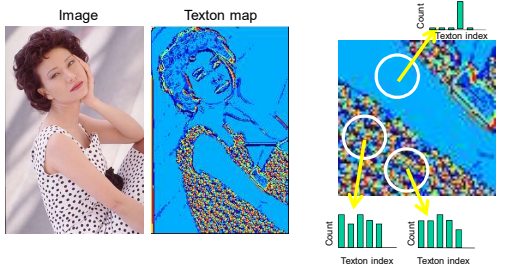



statistics to summarize patterns in small windows

Slide credit: Kristen Grauman

Segmentation with texture features

- Find "textons" by **clustering** vectors of filter bank outputs
- Describe texture in a window based on *texton histogram*



Malik, Belongie, Leung and Shi. IJCV 2001. Adapted from Lana Lazebnik

Image segmentation example

Texture-based regions

Color-based regions

Slide credit: Kristen Grauman

Pixel properties vs. neighborhood properties

query

query

These look very similar in terms of their color distributions (histograms).

How would their *texture* distributions compare?

Slide credit: Kristen Grauman

Outline

- What are grouping problems in vision?
- Inspiration from human perception
 - Gestalt properties
- Bottom-up segmentation via clustering
 - Algorithms:
 - Mode finding and mean shift: k-means, mean-shift
 - Graph-based: normalized cuts
 - Features: color, texture, ...
 - Quantization for texture summaries

Recall: K-means pros and cons

Pros

- Simple, fast to compute
- Converges to local minimum of within-cluster squared error

Cons/issues

- Setting k ?
- Sensitive to initial centers
- Sensitive to outliers
- Detects spherical clusters
- Assuming means can be computed

(A) Undesirable clusters

(B) Ideal clusters

(A) Two natural clusters

(B) k-means clusters

Mean shift clustering and segmentation

- An advanced and versatile technique for clustering-based segmentation

Segmented "landscape 1"

Segmented "landscape 2"

<http://www.caip.rutgers.edu/~comanici/MSPAMI/msPamiResults.html>



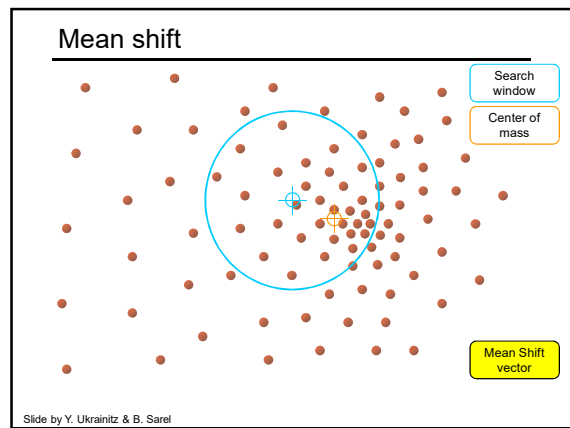
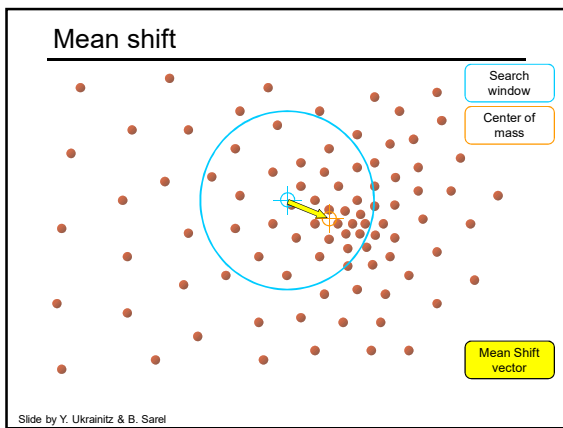
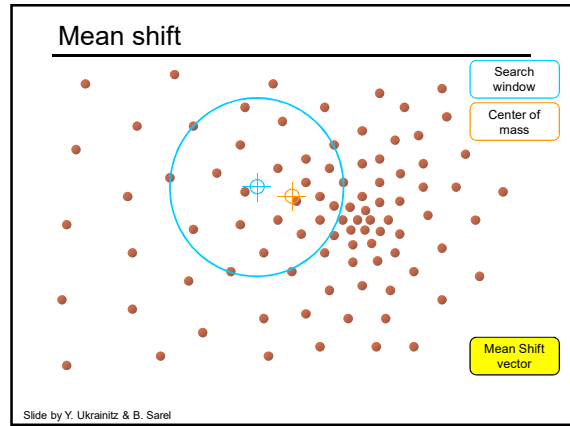
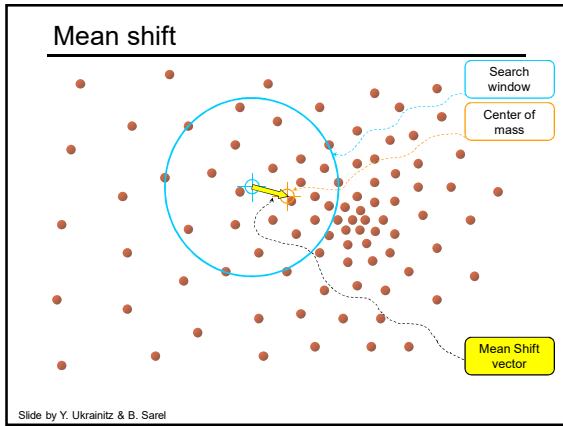
D. Comaniciu and P. Meer, [Mean Shift: A Robust Approach toward Feature Space Analysis](#), PAMI 2002.

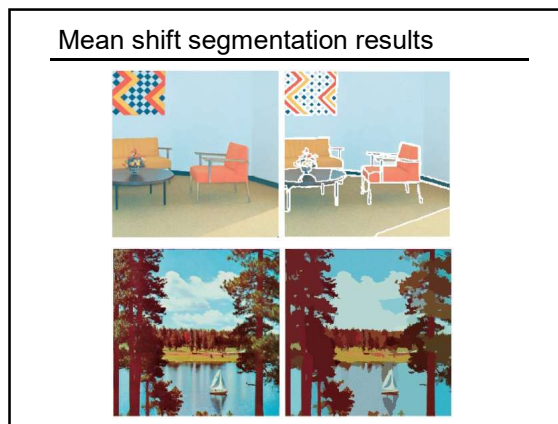
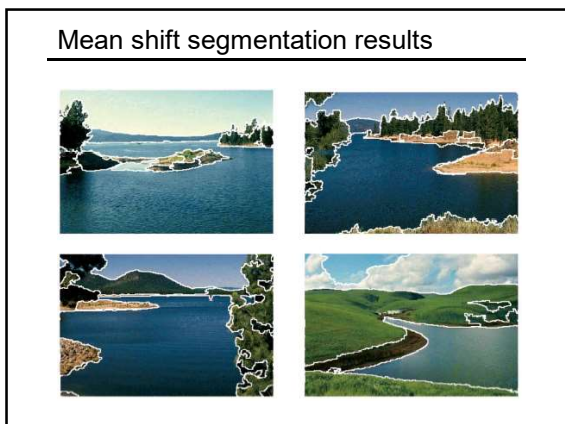
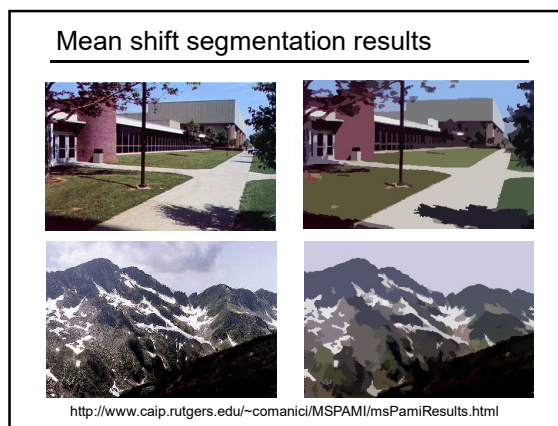
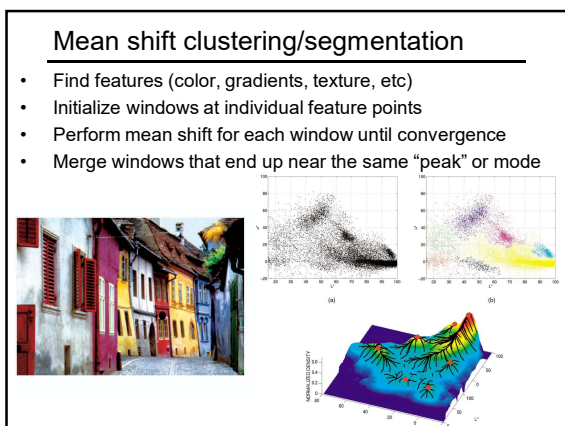
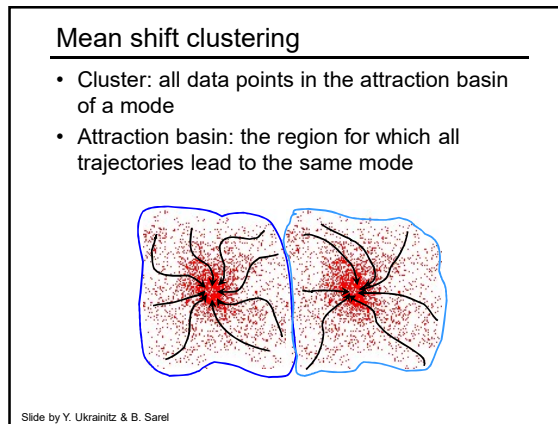

Mean shift algorithm

- The mean shift algorithm seeks *modes* or local maxima of density in the feature space

image

Feature space ($L^*u^*v^*$ color values)





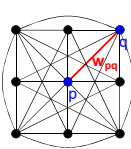
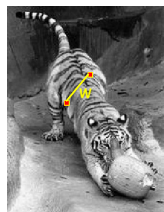
Mean shift

- **Pros:**
 - Does not assume shape on clusters
 - One parameter choice (window size, aka "bandwidth")
 - Generic technique
 - Find multiple modes
- **Cons:**
 - Selection of window size
 - Does not scale well with dimension of feature space

Outline

- What are grouping problems in vision?
- Inspiration from human perception
 - Gestalt properties
- Bottom-up segmentation via clustering
 - Algorithms:
 - Mode finding and mean shift: k-means, mean-shift
 - Graph-based: normalized cuts
 - Features: color, texture, ...
 - Quantization for texture summaries

Images as graphs

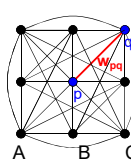




Fully-connected graph

- node (vertex) for every pixel
- link between every pair of pixels, p, q
- affinity weight w_{pq} for each link (edge)
 - w_{pq} measures *similarity*
 - » similarity is *inversely proportional* to difference (in color and position...)

Source: Steve Seitz

Segmentation by Graph Cuts

Break Graph into Segments

- Want to delete links that cross **between** segments
- Easiest to break links that have low similarity (low weight)
 - similar pixels should be in the same segments
 - dissimilar pixels should be in different segments

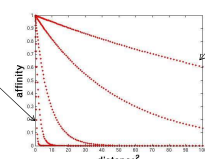
Source: Steve Seitz

Measuring affinity

- One possibility:

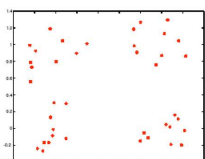
$$aff(x, y) = \exp\left\{-\left(\frac{1}{2\sigma_d^2}\right)\left(\|x - y\|^2\right)\right\}$$

Small sigma:
group only nearby points

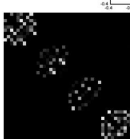
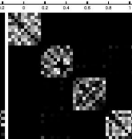
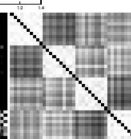


Large sigma:
group distant points

Measuring affinity



Data points $\sigma=2$

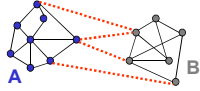




Affinity matrices

$\sigma=.1$ $\sigma=.2$ $\sigma=1$

Slide credit: Kristen Grauman

Cuts in a graph: Min cut



Link Cut

- set of links whose removal makes a graph disconnected
- cost of a cut: $cut(A, B) = \sum_{p \in A, q \in B} w_{p,q}$

Find minimum cut

- gives you a segmentation
- fast algorithms exist for doing this

Source: Steve Seitz

Minimum cut

- Problem with minimum cut:
Weight of cut proportional to number of edges in the cut; tends to produce small, isolated components.

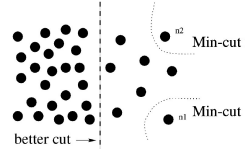
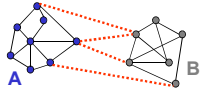


Fig. 1. A case where minimum cut gives a bad partition.

[Shi & Malik, 2000 PAMI]

Cuts in a graph: Normalized cut



Normalized Cut

- fix bias of Min Cut by **normalizing** for size of segments:

$$Ncut(A, B) = \frac{cut(A, B)}{assoc(A, V)} + \frac{cut(A, B)}{assoc(B, V)}$$


$assoc(A, V)$ = sum of weights of all edges that touch A

- Ncut value small when we get two clusters with many edges with high weights, and few edges of low weight between them
- Approximate solution for minimizing the Ncut value : generalized eigenvalue problem.


L. Shi and J. Malik, Normalized Cuts and Image Segmentation, CVPR, 2002

Source: Steve Seitz

Example results



Results: Berkeley Segmentation Engine



<http://www.cs.berkeley.edu/~fowlkes/BSE/>

Normalized cuts: pros and cons

Pros:

- Generic framework, flexible to choice of function that computes weights ("affinities") between nodes
- Does not require model of the data distribution

Cons:

- Time complexity can be high
 - Dense, highly connected graphs → many affinity computations
 - Solving eigenvalue problem
- Preference for balanced partitions

Summary

- Segmentation to find object boundaries or mid-level regions, tokens.
- Bottom-up segmentation via clustering
 - General choices -- features, affinity functions, and clustering algorithms
- Grouping also useful for quantization, can create new feature summaries
 - Texton histograms for texture within local region
- Example clustering methods
 - K-means
 - Mean shift
 - Graph cut, normalized cuts

Segments as primitives for recognition

Multiple segmentations

B. Russell et al., "Using Multiple Segmentations to Discover Objects and their Extent in Image Collections," CVPR 2006 Slide credit: Lana Lazebnik

Top-down segmentation

E. Borenstein and S. Ullman, "Class-specific, top-down segmentation," ECCV 2002
 A. Levin and Y. Weiss, "Learning to Combine Bottom-Up and Top-Down Segmentation," ECCV 2006. Slide credit: Lana Lazebnik

Top-down segmentation

E. Borenstein and S. Ullman, "Class-specific, top-down segmentation," ECCV 2002
 A. Levin and Y. Weiss, "Learning to Combine Bottom-Up and Top-Down Segmentation," ECCV 2006. Slide credit: Lana Lazebnik

Motion segmentation

A. Barbu, S.C. Zhu. Generalizing Swendsen-Wang to sampling arbitrary posterior probabilities, *IEEE Trans. PAMI*, August 2005.

Image grouping

K. Grauman & T. Darrell, Unsupervised Learning of Categories from Sets of Partially Matching Image Features, CVPR 2006.