



## Announcements

- A2 goes out Thursday, due in 2 weeks
- Late submissions on Canvas
- Final exam dates now posted by registrar.
  - Ours is Dec 9, 2-5 pm.
- Check in on pace

## Review questions

- When describing texture, why do we collect filter response statistics within a window?
- How could we integrate rotation invariance into a filter-bank based texture representation?
- What is the Markov assumption?
  - And why is it relevant for the texture synthesis technique of Efros & Leung?
- What are key assumptions for computing optical flow based on image gradients?

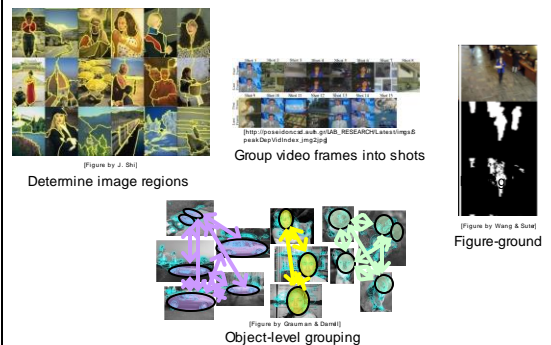
## Outline

- 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





## Common fate



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



Slide credit: Kristen Grauman


## Proximity



<http://www.capital.edu/Recreation/graphics/stock/035.jpg>



Slide credit: Kristen Grauman

## Illusory/subjective contours




Interesting tendency to explain by occlusion

In Vision, D. Marr, 1982



Continuity, explanation by occlusion



D. Forsyth



Continuity, explanation by occlusion




Slide credit: Kristen Grauman



<http://entertainthis.usatoday.com/2015/09/09/how-tom-hardys-legend-poster-hid-this-hilariously-bad-review/> Slide credit: Kristen Grauman

## Figure-ground



## Gestalt

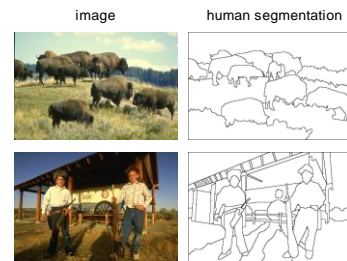
- Gestalt: w hole or group
  - Whole is greater than sum of its parts
  - Relationships among parts can yield new properties/features
- Psychologists identified series of factors that predispose set of elements to be grouped (by human visual system)
- Inspiring observations/explanations; challenge remains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




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

## Image segmentation: toy example





- 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




input image



Slide credit: Kristen Grauman



input image



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

Slide credit: Kristen Grauman


## Clustering

### Clustering systems:

- **Unsupervised learning**
- **Detect patterns** in unlabeled data
  - E.g. group emails or search results
  - E.g. find categories of customers
  - E.g. detect anomalous program executions
- Useful when don't know what you're looking for
- Requires data, but no labels
- Often get gibberish



Slide credit: Dan Klein



- 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


### 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.  $k=5$ )



K-means slides by 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!



## 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


## 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



Slide credit: Kristen Grauman

## 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

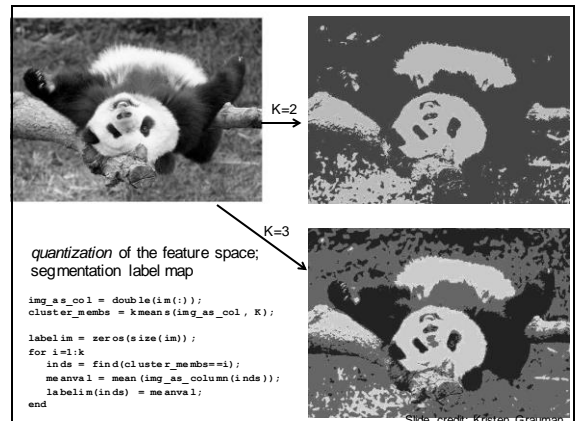
Depending on what we choose as the *feature space*, we can group pixels in different ways.

Grouping pixels based on **intensity** similarity



Feature space: intensity value (1-d)

Slide credit: Kristen Grauman





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...



## 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)

Slide credit: Kristen Grauman


## Recall: texture representation example

Windows with primarily horizontal edges



Windows with small gradient in both directions

Windows with primarily vertical edges




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



Slide credit: Kristen Grauman

## Pixel properties vs. neighborhood properties



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

How would their *texture* distributions compare?

Slide credit: Kristen Grauman

## Material classification example

For an image of a single texture, we can classify it according to its global (image-wide) texton histogram.





Figure from Varma & Zisserman, IJCV 2005

## Material classification example

*Nearest neighbor* classification: label the input according to the nearest known example's label.



$$\chi^2(h_i, h_j) = \frac{1}{2} \sum_{k=1}^K \frac{[h_i(k) - h_j(k)]^2}{h_i(k) + h_j(k)}$$

Manik Varma  
<http://www.robots.ox.ac.uk/~vgg/research/textclass/with.html>

## 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



## Mean shift clustering and segmentation

- An advanced and versatile technique for clustering-based segmentation




<http://www.caip.nyu.edu/~comanici/MSPAM/msPamResults.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




## Mean shift




Slide by Y. Ukrainitz & B. Sarel

## Mean shift




Slide by Y. Ukrainitz & B. Sarel

## Mean shift




Slide by Y. Ukrainitz & B. Sarel

### Mean shift




Slide by Y. Ukrainitz & B. Sarel

### Mean shift




Slide by Y. Ukrainitz & B. Sarel

### Mean shift



Slide by Y. Ukrainitz & B. Sarel


### Mean shift



Slide by Y. Ukrainitz & B. Sarel

### Mean shift clustering


- Cluster: all data points in the attraction basin of a mode
- Attraction basin: the region for which all trajectories lead to the same mode




Slide by Y. Ukrainitz & B. Sarel

### Mean shift clustering/segmentation

- Find features (color, gradients, texture, etc)
- Initialize windows at individual feature points
- Perform mean shift for each window until convergence
- Merge windows that end up near the same "peak" or mode



### Mean shift segmentation results



<http://www.caip.rutgers.edu/~comarid/MSPAM/msPamResults.html>

### Mean shift segmentation results



### Mean shift segmentation results




### 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

Slide credit: Kristen Grauman

## Measuring affinity



## 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 &amp; 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, 1997

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

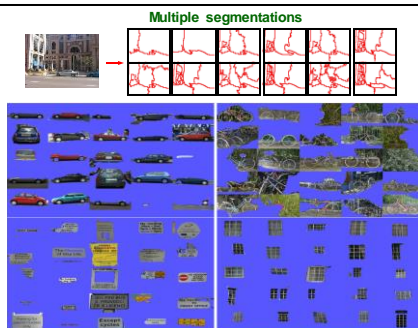
#### Cons:

- 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



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



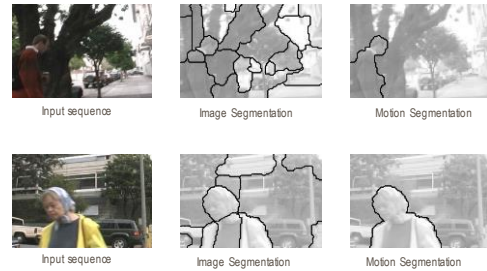
## Top-down segmentation



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.

## Coming up

- Fitting

