



# Segmentation and Fitting

CS 543 D.A. Forsyth



# Graph theoretic clustering

- Represent tokens using a weighted graph.
  - affinity matrix
- Cut up this graph to get subgraphs with strong interior links

# Affinity matrix



# Good split



# 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 and cuts

- Simplest idea: we want a vector  $a$  giving the association between each element and a cluster
- We want elements within this cluster to, on the whole, have strong affinity with one another
- We could maximize  $a^T A a$
- But need the constraint  $a^T a = 1$ 
  - This is an eigenvalue problem - choose the eigenvector of  $A$  with largest eigenvalue


# Example eigenvector



points



matrix



eigenvector



# More than two segments

- Two options
  - Recursively split each side to get a tree, continuing till the eigenvalues are too small
  - Use the other eigenvectors

# Normalized cuts

- Current criterion evaluates within cluster similarity, but not across cluster difference
- Instead, we'd like to maximize the within cluster similarity compared to the across cluster difference
- Write graph as  $V$ , one cluster as  $A$  and the other as  $B$

Maximize

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

- i.e. construct  $A, B$  such that their within cluster similarity is high compared to their association with the rest of the graph

# Normalized cuts

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

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

- and we have a constraint
  - This is hard to do, because  $y$ 's values are quantized

$$y^T D \mathbf{1} = 0$$

# Normalized cuts

- Instead, solve the generalized eigenvalue problem

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

- which gives

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

- Now look for a quantization threshold that maximises the criterion --- i.e all components of  $y$  above that threshold go to one, all below go to -b

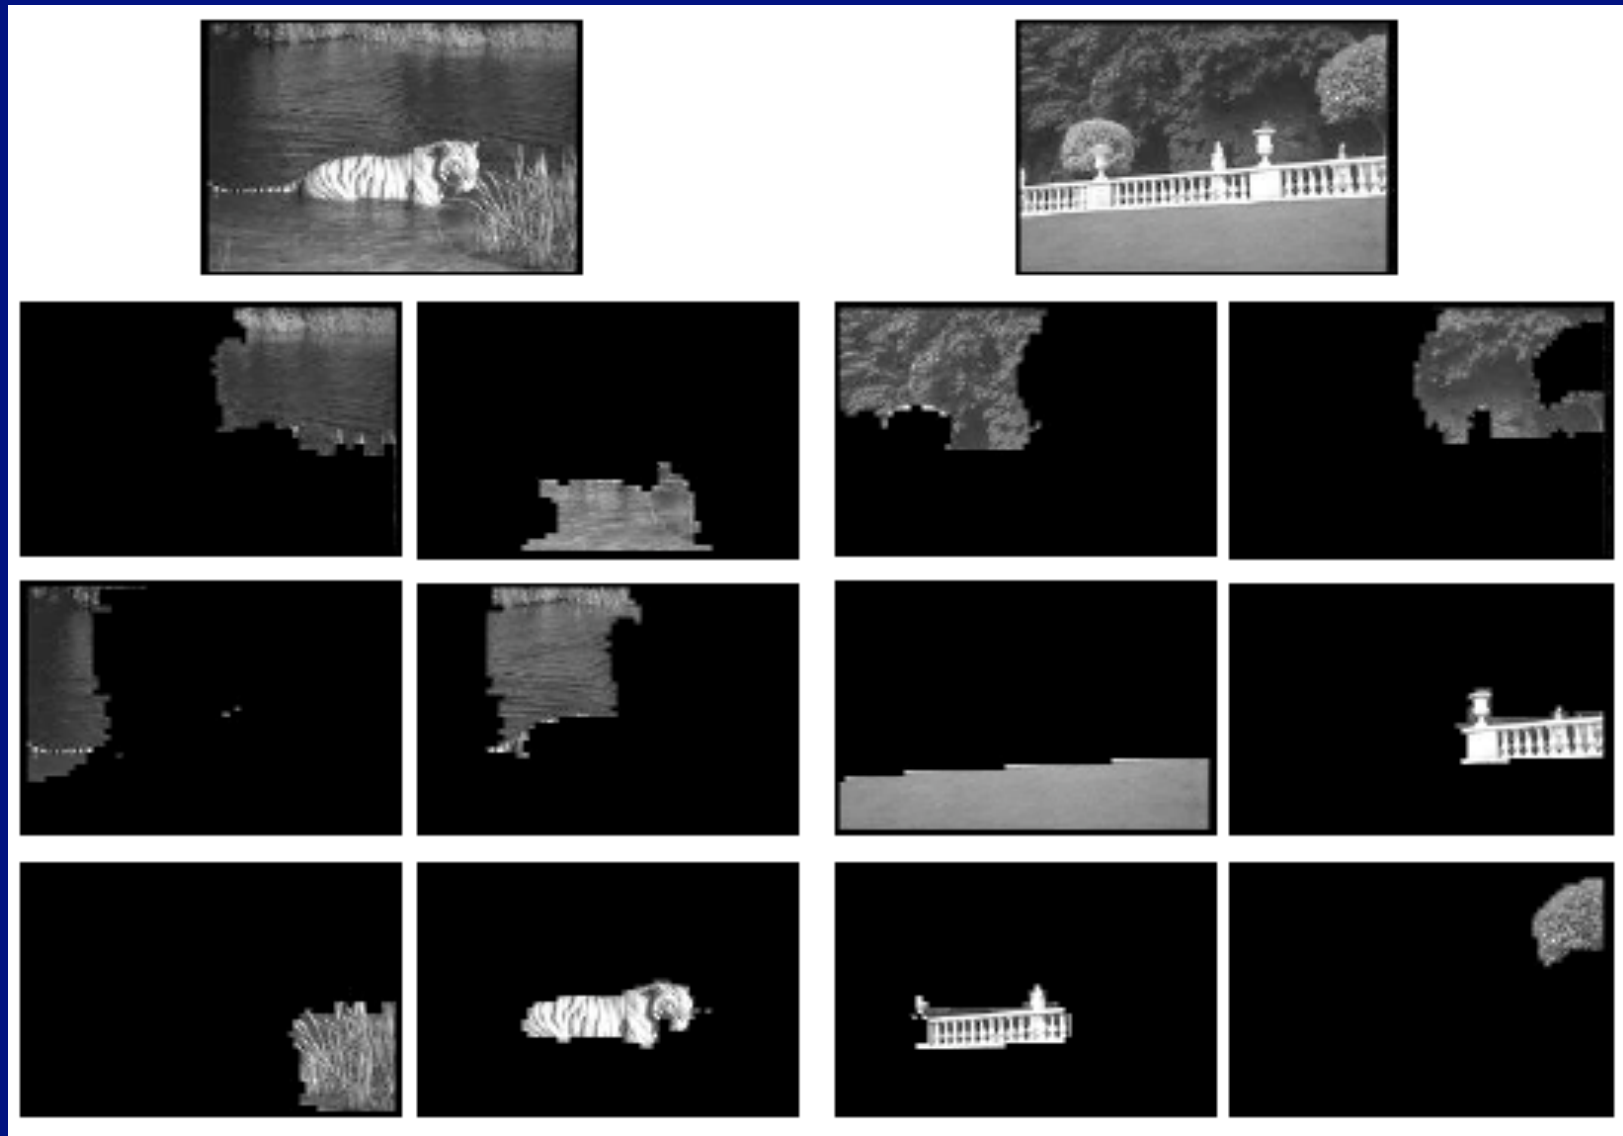


Figure from "Image and video segmentation: the normalised cut framework",  
by Shi and Malik, copyright IEEE, 1998

# Fitting

- Choose a parametric object/some objects to represent a set of tokens
- Most interesting case is when criterion is not local
  - can't tell whether a set of points lies on a line by looking only at each point and the next.
- Three main questions:
  - what object represents this set of tokens best?
  - which of several objects gets which token?
  - how many objects are there?

(you could read line for object here, or circle, or ellipse or...)


# Fitting and the Hough Transform

- Purports to answer all three questions
  - in practice, answer isn't usually all that much help
- We do for lines only
- A line is the set of points  $(x, y)$  such that


$$(\sin \theta)x + (\cos \theta)y + d = 0$$

- Different choices of  $\theta, d > 0$  give different lines
  - For any  $(x, y)$  there is a one parameter family of lines through this point, given by
- Each point gets to vote for each line in the family; if there is a line that has lots of votes, that should be the line passing through the points

$$(\sin \theta)x + (\cos \theta)y + d = 0$$



tokens




votes




# Mechanics of the Hough transform



- Construct an array representing  $\theta, d$
- For each point, render the curve  $(\theta, d)$  into this array, adding one at each cell
- Difficulties
  - how big should the cells be? (too big, and we cannot distinguish between quite different lines; too small, and noise causes lines to be missed)
  - How many lines?
    - count the peaks in the Hough array
  - Who belongs to which line?
    - tag the votes
- Hardly ever satisfactory in practice, because problems with noise and cell size defeat it




tokens




votes







Maximum number of votes



Number of noise points



Line fitting can be max. likelihood - but choice of model is important



# Who came from which line?

- Assume we know how many lines there are - but which lines are they?
  - easy, if we know who came from which line
- Three strategies
  - Incremental line fitting
  - K-means
  - Probabilistic (later!)

**Algorithm 15.1:** Incremental line fitting by walking along a curve, fitting a line to runs of pixels along the curve, and breaking the curve when the residual is too large

```
Put all points on curve list, in order along the curve
Empty the line point list
Empty the line list
Until there are too few points on the curve
  Transfer first few points on the curve to the line point list
  Fit line to line point list
  While fitted line is good enough
    Transfer the next point on the curve
      to the line point list and refit the line
  end
  Transfer last point(s) back to curve
  Refit line
  Attach line to line list
end
```



**Algorithm 15.2:** K-means line fitting by allocating points to the closest line and then refitting.

Hypothesize  $k$  lines (perhaps uniformly at random)

*or*

Hypothesize an assignment of lines to points  
and then fit lines using this assignment

Until convergence


    Allocate each point to the closest line


    Refit lines


end


# Robustness


- As we have seen, squared error can be a source of bias in the presence of noise points
  - One fix is EM - we'll do this shortly
  - Another is an M-estimator
    - Square nearby, threshold far away
  - A third is RANSAC
    - Search for good points


















Too small



Too large



# RANSAC

- Algorithm
  - Choose a small subset uniformly at random
  - Fit to that
  - Anything that is close to result is signal; all others are noise
  - Refit
  - Do this many times and choose the best
- Issues
  - How many times?
    - Often enough that we are likely to have a good line
  - How big a subset?
    - Smallest possible
  - What does close mean?
    - Depends on the problem
  - What is a good line?
    - The number of nearby points is so big it is unlikely to be all outliers

### Algorithm 15.4: RANSAC: fitting lines using random sample consensus

Determine:

$n$  — the smallest number of points required

$k$  — the number of iterations required

$t$  — the threshold used to identify a point that fits well

$d$  — the number of nearby points required  
to assert a model fits well

Until  $k$  iterations have occurred

Draw a sample of  $n$  points from the data  
uniformly and at random

Fit to that set of  $n$  points

For each data point outside the sample

Test the distance from the point to the line  
against  $t$ ; if the distance from the point to the line  
is less than  $t$ , the point is close

end

If there are  $d$  or more points close to the line  
then there is a good fit. Refit the line using all  
these points.

end

Use the best fit from this collection, using the  
fitting error as a criterion

# Fitting curves other than lines

- In principle, an easy generalisation
  - The probability of obtaining a point, given a curve, is given by a negative exponential of distance squared
- In practice, rather hard
  - It is generally difficult to compute the distance between a point and a curve

# Missing variable problems

- In many vision problems, if some variables were known the maximum likelihood inference problem would be easy
  - fitting; if we knew which line each token came from, it would be easy to determine line parameters
  - segmentation; if we knew the segment each pixel came from, it would be easy to determine the segment parameters
  - fundamental matrix estimation; if we knew which feature corresponded to which, it would be easy to determine the fundamental matrix
  - etc.
- This sort of thing happens in statistics, too

# Missing variable problems

- Strategy
  - estimate appropriate values for the missing variables
  - plug these in, now estimate parameters
  - re-estimate appropriate values for missing variables, continue
- eg
  - guess which line gets which point
  - now fit the lines
  - now reallocate points to lines, using our knowledge of the lines
  - now refit, etc.
- We've seen this line of thought before (k means)

# Missing variables - strategy

- We have a problem with parameters, missing variables
- This suggests:
- Iterate until convergence
  - replace missing variable with expected values, given fixed values of parameters
  - fix missing variables, choose parameters to maximise likelihood given fixed values of missing variable
- e.g., iterate till convergence
  - allocate each point to a line with a weight, which is the probability of the point given the line
  - refit lines to the weighted set of points
  - Converges to local extremum
  - Somewhat more general form is available



# Lines and robustness

- We have one line, and  $n$  points
- Some come from the line, some from “noise”
- This is a mixture model:

$$\begin{aligned} P(\text{point} \mid \text{line and noise params}) &= P(\text{point} \mid \text{line})P(\text{comes from line}) + \\ &\quad P(\text{point} \mid \text{noise})P(\text{comes from noise}) \\ &= P(\text{point} \mid \text{line})\lambda + P(\text{point} \mid \text{noise})(1 - \lambda) \end{aligned}$$

- We wish to determine
  - line parameters
  - $p(\text{comes from line})$

# Algorithm for line fitting


- Obtain some start point

$$\theta^{(0)} = (\phi^{(0)}, c^{(0)}, \lambda^{(0)})$$


- Now compute  $\delta$ 's using formula above
- Now compute maximum likelihood estimate of

$$\theta^{(1)}$$


- $\phi, c$  come from fitting to weighted points
- $\lambda$  comes by counting
- Iterate to convergence



The expected values of the deltas at the maximum  
(notice the one value close to zero).



Closeup of the fit



# Choosing parameters

- What about the noise parameter, and the sigma for the line?
  - several methods
    - from first principles knowledge of the problem (seldom really possible)
    - play around with a few examples and choose (usually quite effective, as precise choice doesn't matter much)
  - notice that if  $kn$  is large, this says that points very seldom come from noise, however far from the line they lie
    - usually biases the fit, by pushing outliers into the line
    - rule of thumb; its better to fit to the better fitting points, within reason; if this is hard to do, then the model could be a problem

# Other examples


- Fitting multiple lines
  - rather like fitting one line, except there are more hidden variables
  - easiest is to encode as an array of hidden variables
    - a table with a
      - one where the  $i$ 'th point comes from the  $j$ 'th line,
      - zeros otherwise
  - rest is on same lines as above

# Issues with EM


- Local maxima
  - can be a serious nuisance in some problems
  - no guarantee that we have reached the “right” maximum
- Starting
  - $k$  means to cluster the points is often a good idea




Local maximum




which is an excellent fit to some points




and the deltas for this maximum




A dataset that is well fitted by four lines




Result of EM fitting, with one line (or at least, one available local maximum).



Result of EM fitting, with two lines (or at least,  
one available local maximum).



Seven lines can produce a rather logical answer



# Other examples

- Segmentation
  - a segment is a gaussian that emits feature vectors (which could contain colour; or colour and position; or colour, texture and position).
  - segment parameters are mean and (perhaps) covariance
  - if we knew which segment each point belonged to, estimating these parameters would be easy
  - rest is on same lines as fitting line



## Segmentation with EM

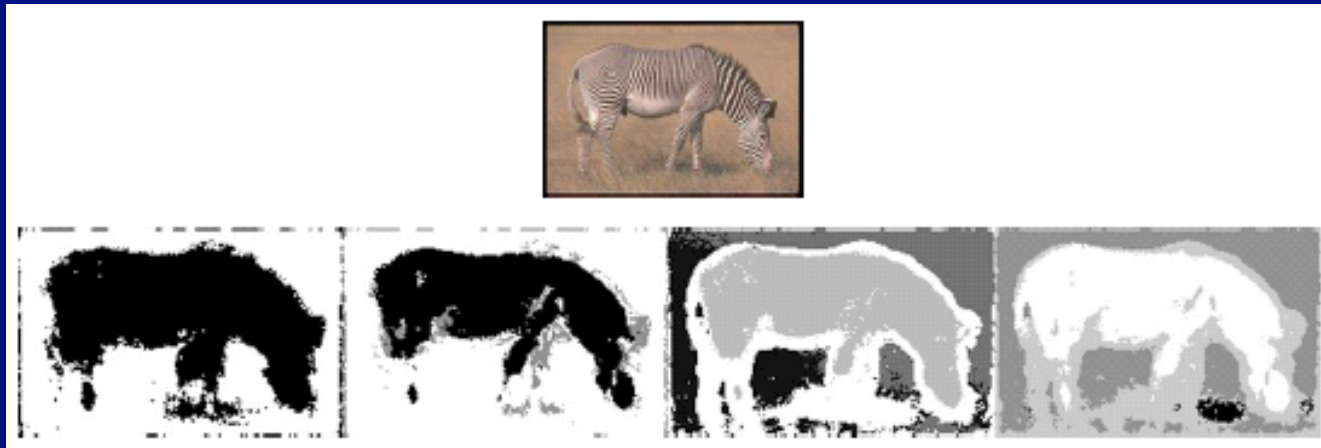


Figure from "Color and Texture Based Image Segmentation Using EM and Its Application to Content Based Image Retrieval", S.J. Belongie et al., Proc. Int. Conf. Computer Vision, 1998, c1998, IEEE

# Motion segmentation with EM

- Model image pair (or video sequence) as consisting of regions of parametric motion

- affine motion is popular

$$\begin{pmatrix} v_x \\ v_y \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} t_x \\ t_y \end{pmatrix}$$

- Now we need to
- determine which pixels belong to which region
- estimate parameters
- Likelihood

- assume

$$I(x, y, t) = I(x + v_x, y + v_y, t + 1) + \textit{noise}$$

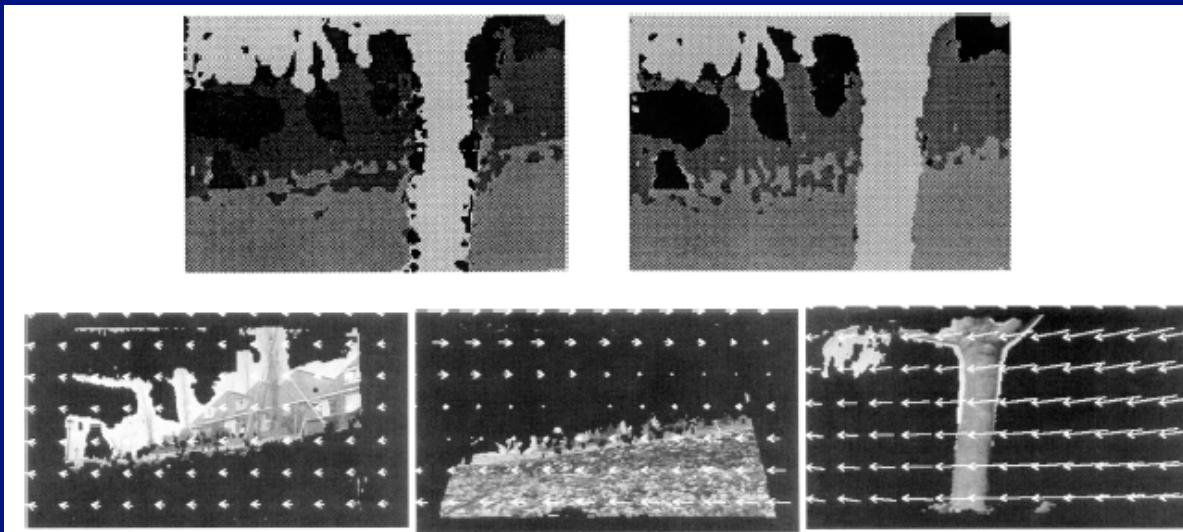
- Straightforward missing variable problem, rest is calculation



Three frames from the MPEG “flower garden” sequence

Figure from “Representing Images with layers,” by J. Wang and E.H. Adelson, IEEE  
Transactions on Image Processing, 1994, c 1994, IEEE

Grey level shows region no. with highest probability



Segments and motion fields associated with them

Figure from "Representing Images with layers," by J. Wang and E.H. Adelson, IEEE  
Transactions on Image Processing, 1994, c 1994, IEEE



If we use multiple frames to estimate the appearance of a segment, we can fill in occlusions; so we can re-render the sequence with some segments removed.

Figure from "Representing Images with layers," by J. Wang and E.H. Adelson, IEEE Transactions on Image Processing, 1994, c 1994, IEEE