Lecture 8: Fitting

Tuesday, Sept 25



Announcements, schedule

- Grad student extensions
 - Due end of term
 - Data sets, suggestions
- Reminder: Midterm Tuesday 10/9
- Problem set 2 out Thursday, due 10/11

Outline

- Review from Thursday (affinity, cuts)
- Local scale and affinity computation
- Hough transform
- Generalized Hough transform
 - Shape matching applications
- Fitting lines
 - Least squares
 - Incremental fitting, k-means





Mean shift

- Labeling of data points: points visited by any window converging to same mode get the same label
- [Comaniciu & Meer, PAMI 2002] : If data point is visited by multiple diverging mean shift processes, "majority vote"



Slide by Steve Seitz

<image><section-header><section-header><section-header><section-header><section-header>









Measuring affinity

- One possibility: $aff(x,y) = \exp\left\{-\left(\frac{1}{2\sigma_d^2}\right)\left(\|x-y\|^2\right)\right\}$
- Map distances to similarities, use as edge weights on graph















- Want a vector w giving the "association" between each element and a cluster
- Want elements within this cluster to have strong affinity with one another

• Maximize
$$\mathbf{w}^T A \mathbf{w}$$

subject to the constraint $\mathbf{w}^T \mathbf{w} = 1$

... Aw = λw

Eigenvalue problem : choose the eigenvector of A with largest eigenvalue



- Given a symmetric matrix A, find a vector x such that
- $\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x}$ is maximum AND
- $\|\boldsymbol{x}\|^2 = 1$
- Find x such that $\frac{x^T A x}{x^T x}$ is maximum.

The solution to this problem is given by the following theorem:

• $\frac{x^T A x}{x^T x}$ reaches its absolute maximum when x is an eigenvector of A corresponding to the *largest* eigenvalue λ_{max} .



Eigenvectors and multiple cuts

- Use eigenvectors associated with k largest eigenvalues as cluster weights
- Or re-solve recursively







Figure 2: The effect of local scaling. (a) Input data points. A tight cluster resides within a background cluster. (b) The affinity between each point and its surrounding neighbors is indicated by the thickness of the line connecting them. The affinities across clusters are larger than the affinities within the background cluster. (c) The corresponding visualization of affinities after local scaling. The affinities across clusters are now significantly lower than the affinities within any single cluster.

[Self-Tuning Spectral Clustering, L. Zelnik-Manor and P. Perona, NIPS 2004]











Hough transform

- Maps model (pattern) detection problem to simple peak detection problem
- Record all the structures on which each point lies, then look for structures that get many votes
- Useful for line fitting






Polar representation for lines

- Issues with (*m*,*b*) parameter space:
 - Can take on infinite values
 - Undefined for vertical lines (x=constant)















Extensions

Extension 1: Use the image gradient

1. same

2. for each edge point I[x,y] in the image

compute unique (d, θ) based on image gradient at (x,y)

H[d, θ] += 1

3. same

4. same

(Reduces degrees of freedom)



Extensions

Extension 1: Use the image gradient

1. same

2. for each edge point I[x,y] in the image

compute unique (d, θ) based on image gradient at (x,y)

H[d, θ] += 1

3. same

4. same

(Reduces degrees of freedom)

Extension 2

• give more votes for stronger edges (use magnitude of gradient)

Extension 3

- change the sampling of (d, θ) to give more/less resolution

Extension 4

• The same procedure can be used with circles, squares, or any other shape...







Hough transform for circles

```
For every edge pixel (x,y):

For each possible radius value r:

\theta = gradient direction, from x,y to center

a = x - r \cos(\theta)

b = y + r \sin(\theta)

H[a,b,r] += 1

end

end
```







Finding Coins



Note: a different Hough transform (with separate accumulators) was used for each circle radius (quarters vs. penny).

Coin finding sample images from: Vivek Kwatra

Hough transform for parameterized curves

For any curve with analytic form f(**x**,**a**) = 0, where

x : edge pixel in image coordinates

a : vector of parameters

1. Initialize accumulator array: H[**a**] = 0

2. For each edge pixel:

determine each **a** such that f(**x**,**a**) = 0, and increment H[**a**]

3. Local maxima in H correspond to curves

Practical tips

- Minimize irrelevant tokens first (take edge points with significant gradient magnitude)
- Choose a good grid / discretization (trial and error)
- Vote for neighbors, also (smoothing in accumulator array)
- Utilize direction of edge to reduce free parameters by 1

Hough transform

- Pros
 - All points are processed independently, so can cope with occlusion
 - Some robustness to noise: noise points unlikely to contribute consistently to any single bin
 - Can detect multiple instances of a model in a single pass

Hough transform

- Cons
 - Complexity of search time increases exponentially with the number of model parameters
 - Non-target shapes can produce spurious peaks in parameter space
 - Quantization: hard to pick a good grid size
 - Too coarse → large votes obtained when too many different lines correspond to a single bucket
 - Too fine → miss lines because some points that are not exactly collinear cast votes for different buckets



Generalized Hough transform

- For model shape: construct a table storing these
 r vectors as function of gradient direction
- To detect model shape: for each edge point
 - Index into table with θ
 - Use indexed r vectors to vote for (x,y) position of reference point
- Peak in this Hough space is reference point with most supporting edges

Assuming translation is the only transformation here, i.e., orientation and scale are fixed.





Generalized Hough Transform

Find Object Center (x_c, y_c) given edges (x_i, y_i, ϕ_i)

Create Accumulator Array $A(x_c, y_c)$ Initialize: $A(x_c, y_c) = 0 \quad \forall (x_c, y_c)$

For each edge point (x_i, y_i, ϕ_i)



For each **r** vector entry indexed in table, compute:

$$x_c = x_i + r_k^i \cos \alpha_k^i$$
$$y_c = y_i + r_k^i \sin \alpha_k^i$$

Increment Accumulator: $A(x_c, y_c) = A(x_c, y_c) + 1$

Find peaks in $A(x_c, y_c)$

With modifications to table can generalize to add scale, orientation – increases size of accumulator array

Generalizing for scale, orientation

- To search for shapes at arbitrary scale and orientation
 - Add the parameters to the accumulator array (4d)
 - Update table

Example in recognition: implicit shape model

 Build 'codebook' of local appearance for each category using agglomerative clustering

[B. Leibe, A. Leonardis, and B. Schiele, Combined Object Categorization and Segmentation with an Implicit Shape Model, 2004]





Implicit shape model

- Given test image, extract patches, match to codebook entry (or entries)
- Cast votes for possible positions of object center
- Search for maxima in voting space using Mean-Shift
- (Extract weighted segmentation mask based on stored masks for the codebook occurrences)

[B. Leibe, A. Leonardis, and B. Schiele, Combined Object Categorization and Segmentation with an Implicit Shape Model, 2004]



Fig. 1. The recognition procedure. Image patches are extracted around interest points and compared to the codebook. Matching patches then cast probabilistic votes, which lead to object hypotheses that can later be refined. Based on the refined hypotheses, we compute a categoryspecific segmentation.

[B. Leibe, A. Leonardis, and B. Schiele, Combined Object Categorization and Segmentation with an Implicit Shape Model, 2004]

Grouping and fitting

- Grouping, segmentation: make a compact representation that merges similar features
 - Relevant algorithms: K-means, hierarchical clustering, Mean Shift, Graph cuts
- Fitting: fit a model to your observed features
 - Relevant algorithms: Hough transform for lines, circles (parameterized curves), generalized Hough transform for arbitrary boundaries; least squares; assigning points to lines incrementally or with kmeans