

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



































Eigenvectors and multiple cuts

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













Fitting lines

- Given points that belong to a line, what is the line?
- How many lines are there?
- Which points belong to which lines?



















































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: $A(x_c, y_c) = 0 \forall (x_c, y_c)$
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











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 k-means