Announcements

- PS0 grades are up on Canvas
- PS0 stats:
  - Mean: 92.13
  - Standard Dev: 9.71

2

Last time: Grouping

- Bottom-up segmentation via clustering
  - To find mid-level regions, tokens
  - General choices -- features, affinity functions, and clustering algorithms
  - Example clustering algorithms
    - Mean shift and mode finding: K-means, Mean shift
    - Graph theoretic: Graph cut, normalized cuts
- Grouping also useful for quantization
  - Texton histograms for texture within local region

3

Slide credit: Kristen Grauman

Recall: Images as graphs

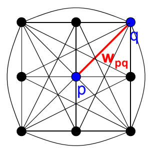

Fully-connected graph

- node for every pixel
- link between *every* pair of pixels, **p**,**q**
- similarity w_{pq} for each link
  - » similarity is *inversely proportional* to difference in color and position

4

Slide by Steve Seitz

Last time: Measuring affinity

40 data points

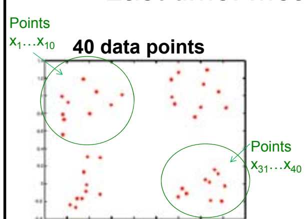

40 x 40 affinity matrix *A*

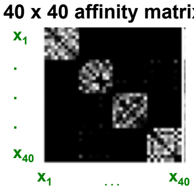

$$A(i,j) = \exp\{-(1/_{2\sigma^2})\|\boldsymbol{x}_i - \boldsymbol{x}_j\|^2\}$$

1. What do the **blocks** signify?
2. What does the **symmetry** of the matrix signify?
3. How would the matrix change with **larger value of σ**?

Slide credit: Kristen Grauman

5

Example: weighted graphs

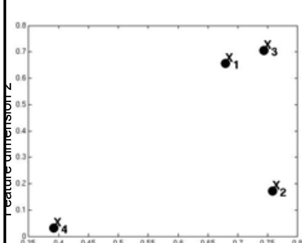

- Suppose we have a 4-pixel image (i.e., a 2 x 2 matrix)
- Each pixel described by 2 features

Dimension of data points : d = 2
Number of data points : N = 4

6

Kristen Grauman

Example: weighted graphs

Computing the distance matrix:

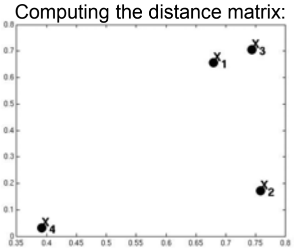

D(:,1)=

(1,:)=	(0)	0.24	0.01	0.47

for i=1:N
    for j=1:N
        $D(i,j) = ||x_i - x_j||^2$
    end
end

7

Kristen Grauman

Example: weighted graphs

Computing the distance matrix:

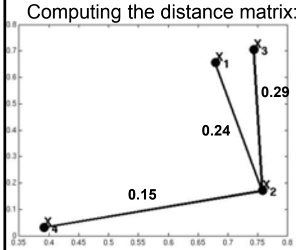

D(:,1)=

D(1,:)=	(0)	0.24	0.01	0.47

for i=1:N
    for j=1:N
        $D(i,j) = ||x_i - x_j||^2$
    end
end

8

Kristen Grauman

Example: weighted graphs

Computing the distance matrix:

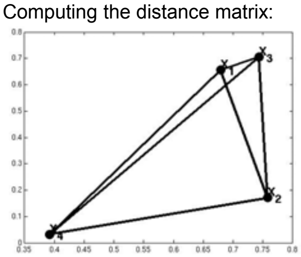

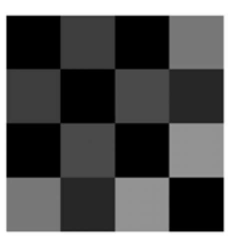

N x N matrix

for i=1:N
    for j=1:N
        $D(i,j) = ||x_i - x_j||^2$
    end
end

9

Kristen Grauman

Example: weighted graphs

D Distances→affinities **A**

σ = 0.5

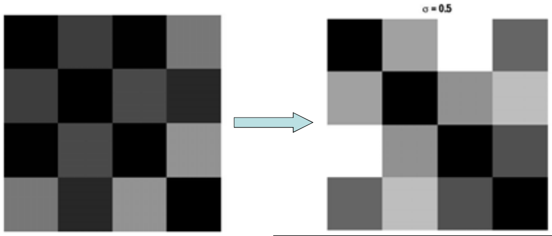

```
for i=1:N
  for j=1:N
    D(i,j) = ||x_i- x_j||^2
  end
end
```

```
for i=1:N
  for j=i+1:N
    A(i,j) = exp(-1/(2*σ^2)*||x_i- x_j||^2);
    A(j,i) = A(i,j);
  end
end
```

10

Kristen Grauman

Scale parameter σ affects affinity

Distance matrix **D=** 

Affinity matrix with increasing σ:

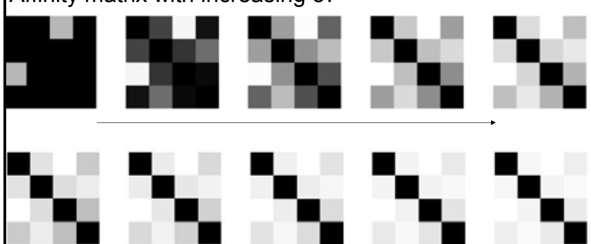

11

Kristen Grauman

Visualizing a shuffled affinity matrix

If we permute the order of the vertices as they are referred to in the affinity matrix, we see different patterns:

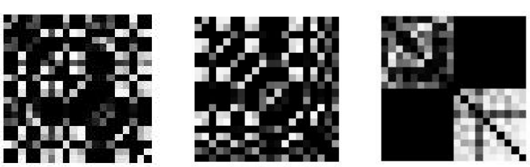

12

Kristen Grauman

Putting these two aspects together

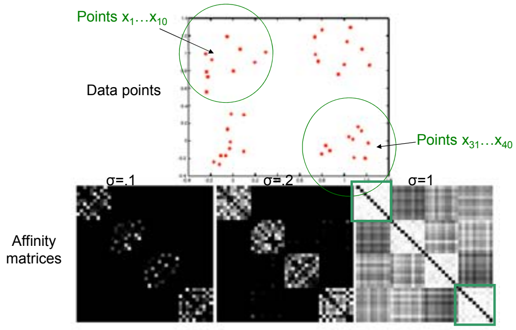

$$A(i,j) = \exp\{-(1/_{2\sigma^2})\|\boldsymbol{x}_i - \boldsymbol{x}_j\|^2\}$$

13

Kristen Grauman

Goal: Segmentation by Graph Cuts

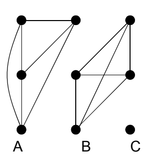

Break graph into segments

- Delete links that cross between segments
  - Easiest to break links that have low similarity
    - similar pixels should be in the same segments
    - dissimilar pixels should be in different segments

14

Slide credit: Kristen Grauman

Now: Fitting

- Want to associate a model with multiple observed features

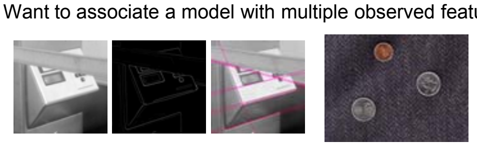

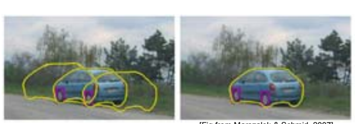

[Fig from Marszalek & Schmid, 2007]

For example, the model could be a line, a circle, or an arbitrary shape.

15

Fitting: Main idea

- Choose a parametric model that best represents a set of features
- Membership criterion is not local
  - Can't tell whether a point belongs to a given model just by looking at that point
- Three main questions:
  - What model represents this set of features best?
  - Which of several model instances gets which feature?
  - How many model instances are there?
- Computational complexity is important
  - It is infeasible to examine every possible set of parameters and every possible combination of features

16

Slide credit: L. Lazebnik

Example: Line fitting

- Why fit lines?
  Many objects characterized by presence of straight lines

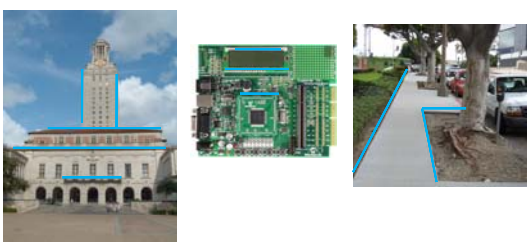

- Wait, why aren't we done just by running edge detection?

17

Kristen Grauman

Difficulty of line fitting

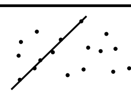

- **Extra** edge points (clutter), multiple models:
  - which points go with which line, if any?
- Only some parts of each line detected, and some parts are **missing:**
  - how to find a line that bridges missing evidence?
- **Noise** in measured edge points, orientations:
  - how to detect true underlying parameters?

18

Kristen Grauman

Voting

- It's not feasible to check all combinations of features by fitting a model to each possible subset.
- **Voting** is a general technique where we let each feature *vote for all models that are compatible with it.*
  - Cycle through features, cast votes for model parameters.
  - Look for model parameters that receive a lot of votes.
- Noise & clutter features will cast votes too, *but* typically their votes should be inconsistent with the majority of "good" features.

19
Kristen Grauman

Fitting lines: Hough transform

- Given points that belong to a line, what is the line?
- How many lines are there?
- Which points belong to which lines?
- **Hough Transform** is a voting technique that can be used to answer all of these questions.

Main idea:

1. Record vote for each possible line on which each edge point lies.
2. Look for lines that get many votes.

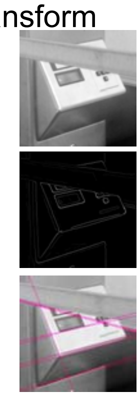

20
Kristen Grauman

Finding lines in an image: Hough space

Equation of a line?
y = mx + b

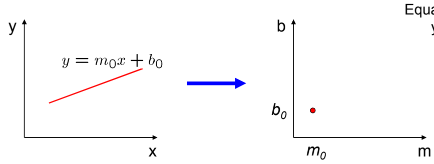

image space

Hough (parameter) space

Connection between image (x,y) and Hough (m,b) spaces

- A line in the image corresponds to a point in Hough space
- To go from image space to Hough space:
  - given a set of points (x,y), find all (m,b) such that y = mx + b

21
Slide credit: Steve Seitz

Finding lines in an image: Hough space

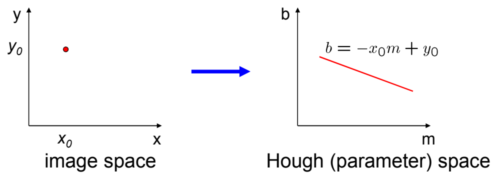

Connection between image (x,y) and Hough (m,b) spaces

- A line in the image corresponds to a point in Hough space
- To go from image space to Hough space:
  - given a set of points (x,y), find all (m,b) such that y = mx + b
- What does a point (x_0, y_0) in the image space map to?
  - Answer: the solutions of $b = -x_0 m + y_0$
  - this is a line in Hough space

22

Slide credit: Steve Seitz

Finding lines in an image: Hough space

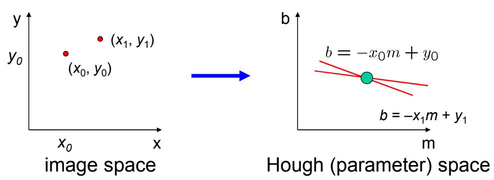

What are the line parameters for the line that contains both (x_0, y_0) and (x_1, y_1)?

- It is the intersection of the lines $b = -x_0 m + y_0$ and $b = -x_1 m + y_1$

23

Slide credit: Kristen Grauman

Finding lines in an image: Hough algorithm

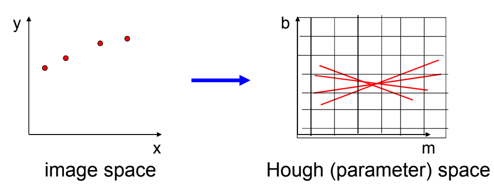

How can we use this to find the most likely parameters (m,b) for the most prominent line in the image space?

- Let each edge point in image space *vote* for a set of possible parameters in Hough space
- Accumulate votes in discrete set of bins; parameters with the most votes indicate line in image space.

24

Slide credit: Kristen Grauman

Polar representation for lines

Issues with usual (*m,b*) parameter space: can take on infinite values, undefined for vertical lines.

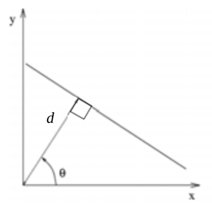

d : perpendicular distance from line to origin

θ : angle the perpendicular makes with the x-axis

$$x\cos\theta + y\sin\theta = d$$

Point in image space → sinusoid segment in Hough space

25

Adapted from Kristen Grauman

- Hough line demo

26

Hough transform algorithm

Using the polar parameterization:

$$x\cos\theta + y\sin\theta = d$$

Basic Hough transform algorithm

1. Initialize H[d, θ]=0
2. for each edge point I[x,y] in the image
   for θ = [θ_{min} to θ_{max}] // some quantization
   $d = x\cos\theta + y\sin\theta$
   H[d, θ] += 1
3. Find the value(s) of (d, θ) where H[d, θ] is maximum
4. The detected line in the image is given by $d = x\cos\theta + y\sin\theta$

H: accumulator array (votes)

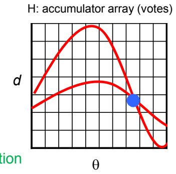

Time complexity (in terms of number of votes per pt)?

27

Source: Steve Seitz

1. Image → Canny

Derek Hoiem

2. Canny → Hough votes

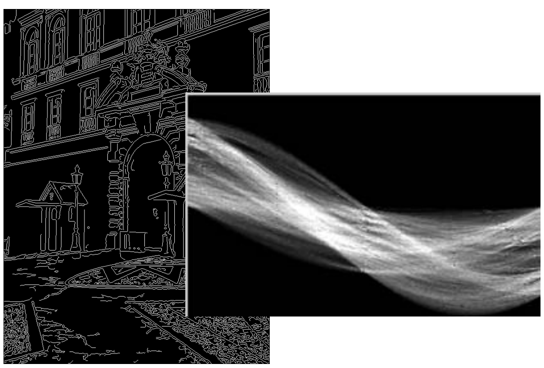

Derek Hoiem

3. Hough votes → Edges

Find peaks

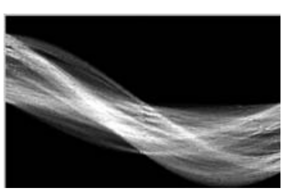

Derek Hoiem

Hough transform example

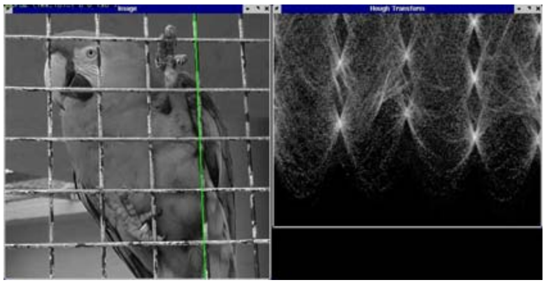

Derek Hoiem http://ostatic.com/files/images/ss_hough.jpg

Original image

Canny edges

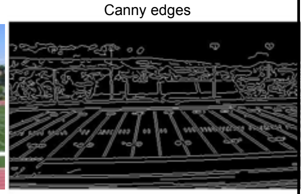

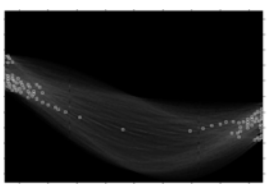

Vote space and top peaks

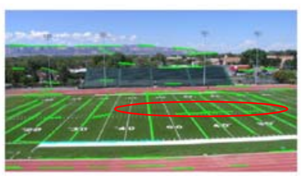

Showing longest segments found

32

Kristen Grauman

Impact of noise on Hough

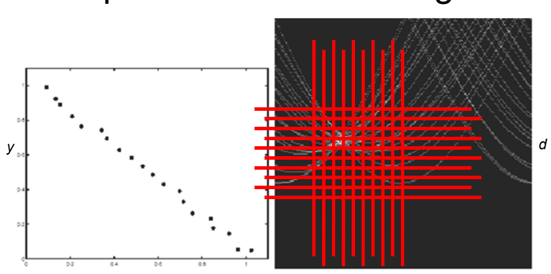

Image space edge coordinates

Votes

What difficulty does this present for an implementation?

Impact of noise on Hough

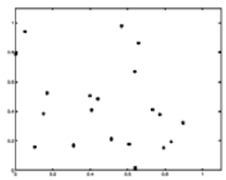

Image space edge coordinates

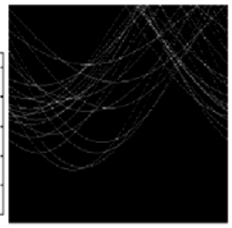

Votes

Here, everything appears to be "noise", or random edge points, but we still see peaks in the vote space. 34

Slide credit: Kristen Grauman

Extensions

Recall: when we detect an edge point, we also know its gradient direction

Extension 1: Use the image gradient
1. same
2. for each edge point I[x,y] in the image
   θ = gradient at (x,y)
   $d = x\cos\theta + y\sin\theta$
   H[d, θ] += 1
3. same
4. same

(Reduces degrees of freedom)

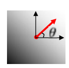

$\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right]$

$\theta = \tan^{-1}\left(\frac{\partial f}{\partial y} / \frac{\partial f}{\partial x}\right)$

35

Slide credit: Kristen Grauman

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…

36

Source: Steve Seitz

Hough transform for circles

- Circle: center (a,b) and radius r

$$(x_i - a)^2 + (y_i - b)^2 = r^2$$

- For a fixed radius r

Equation of circle?

Equation of set of circles that all pass through a point?

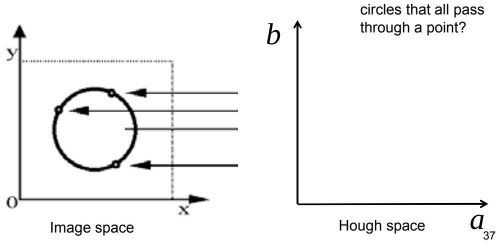

37

Adapted by Devi Parikh from: Kristen Grauman

Hough transform for circles

- Circle: center (a,b) and radius r

$$(x_i - a)^2 + (y_i - b)^2 = r^2$$

- For a fixed radius r

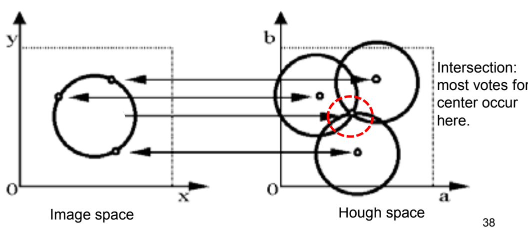

38

Kristen Grauman

Hough transform for circles

- Circle: center (a,b) and radius r

$$(x_i - a)^2 + (y_i - b)^2 = r^2$$

- For an unknown radius r

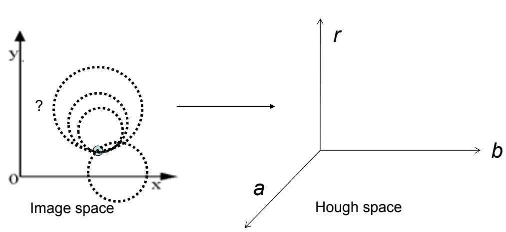

39

Kristen Grauman

Hough transform for circles

- Circle: center (a,b) and radius r

$$(x_i - a)^2 + (y_i - b)^2 = r^2$$

- For an unknown radius r

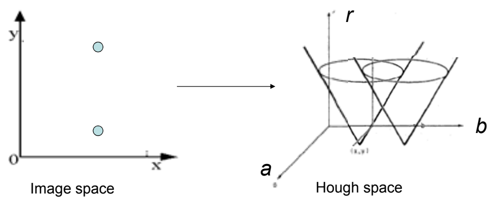

40

Kristen Grauman

Hough transform for circles

- Circle: center (a,b) and radius r

$$(x_i - a)^2 + (y_i - b)^2 = r^2$$

- For an unknown radius r, **known** gradient direction

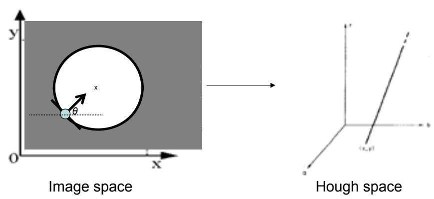

41

Kristen Grauman

Hough transform for circles

For every edge pixel (x,y) :
  For each possible radius value r:
    For each possible gradient direction θ:
    // or use estimated gradient at (x,y)
      $a = x - r\cos(\theta)$ // column
      $b = y + r\sin(\theta)$ // row
      $H[a,b,r]$ += 1
  end
end

Time complexity per edge pixel?

- Check out online demo : http://www.markschulze.net/java/hough/

42

Kristen Grauman

Example: detecting circles with Hough

Original | Edges | Votes: Penny

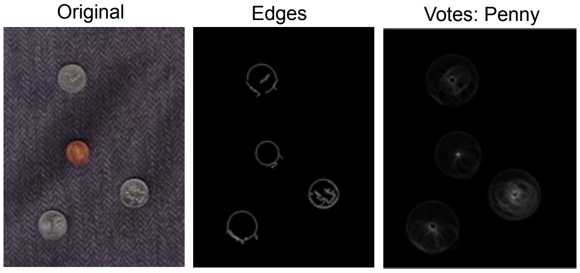

Note: a different Hough transform (with separate accumulators) was used for each circle radius (quarters vs. penny).

43

Coin finding sample images from: Vivek Kwatra

Example: detecting circles with Hough

Combined detections / Original | Edges | Votes: Quarter

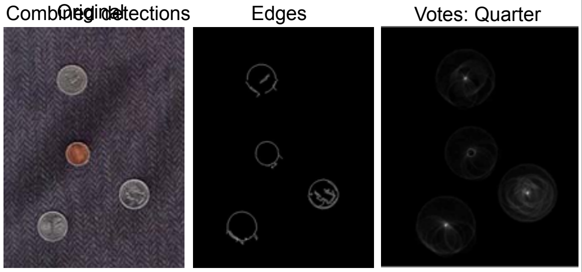

44

Coin finding sample images from: Vivek Kwatra

Example: iris detection

Gradient+threshold | Hough space (fixed radius) | Max detections

- Hemerson Pistori and Eduardo Rocha Costa http://rsbweb.nih.gov/ij/plugins/hough-circles.html

45

Kristen Grauman

Example: iris detection

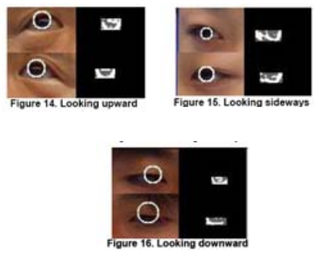

- An Iris Detection Method Using the Hough Transform and Its Evaluation for Facial and Eye Movement, by Hideki Kashima, Hitoshi Hongo, Kunihito Kato, Kazuhiko Yamamoto, ACCV 2002.

46

Voting: practical tips

- Minimize irrelevant tokens first
- Choose a good grid / discretization

  Too fine ? Too coarse

- Vote for neighbors, also (smoothing in accumulator array)
- Use direction of edge to reduce parameters by 1

47

Kristen Grauman

Hough transform: pros and cons

Pros

- All points are processed independently, so can cope with occlusion, gaps
- 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

Cons

- Complexity of search time increases exponentially with the number of model parameters
- Non-target shapes can produce spurious peaks in parameter space
- Quantization: can be tricky to pick a good grid size

48

Kristen Grauman

Generalized Hough Transform

- What if we want to detect arbitrary shapes?

Intuition:

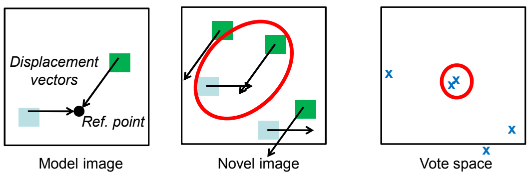

Now suppose those colors encode gradient directions…

49

Kristen Grauman

Generalized Hough Transform

- Define a model shape by its boundary points and a reference point.

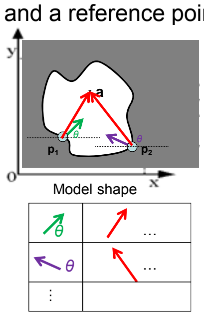

Offline procedure:

At each boundary point, compute displacement vector: $\mathbf{r} = \mathbf{a} - \mathbf{p}_i$.

Store these vectors in a table indexed by gradient orientation θ.

50

[Dana H. Ballard, Generalizing the Hough Transform to Detect Arbitrary Shapes, 1980]

Generalized Hough Transform

Detection procedure:

For each edge point:

- Use its gradient orientation θ to index into stored table
- Use retrieved **r** vectors to vote for reference point

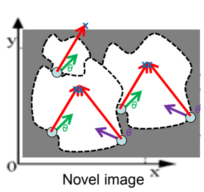

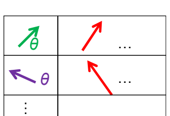

Assuming translation is the only transformation here, i.e., orientation and scale are fixed.

51

Generalized Hough for object detection

- Instead of indexing displacements by gradient orientation, index by matched local patterns.

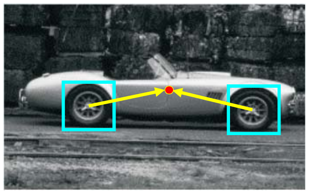

training image

"visual codeword" with displacement vectors

B. Leibe, A. Leonardis, and B. Schiele, Combined Object Categorization and Segmentation with an Implicit Shape Model, ECCV Workshop on Statistical Learning in Computer Vision 2004

52

Source: L. Lazebnik

Generalized Hough for object detection

- Instead of indexing displacements by gradient orientation, index by "visual codeword"

test image

B. Leibe, A. Leonardis, and B. Schiele, Combined Object Categorization and Segmentation with an Implicit Shape Model, ECCV Workshop on Statistical Learning in Computer Vision 2004

53

Source: L. Lazebnik

Summary

- **Fitting** problems require finding any supporting evidence for a model, even within clutter and missing features
  - associate features with an explicit model
- **Voting** approaches, such as the **Hough transform**, make it possible to find likely model parameters without searching all combinations of features
  - Hough transform approach for lines, circles, …, arbitrary shapes defined by a set of boundary points, recognition from patches

54

Questions?

See you Tuesday!

55