Lecture 3: The Shape Context

Wesley Snyder, Ph.D.

UWA, CSSE

NCSU, ECE

Giant Quokka

The Shape Context

- Need for invariance
  - Translation, rotation, scale
- How to achieve invariance
- Definition of the Shape Context
- Computing the Shape Context
- The correspondence problem
- Linear assignment problem
- Implementation
- Performance

The Shape Recognition Problem

- (For now) assume 2-D shapes
- Allow arbitrary in-viewing plane
  - rotation
  - translation
  - scale change
- Consider partial occlusion
- Consider only boundary information

Invariance

Translation: Choose an object-centered point of reference (center of gravity is good), and subtract those coordinates from all the points on the boundary

Invariance

Scale: Options:

- find distance from CG to boundary points
  - Average, median, maximum
- find distance between pairs of boundary points
  - Average, median, maximum

Median of pairs seems to be most robust to outliers

Invariance

Rotation: lots of ways. Let's just do one: Make all directions relative to the tangent line at each point. Which means, of course, that we have to find the tangent line.

Finding the tangent to a curve

Let $p_i = (x_i, y_i)$ be a point on the boundary, and let p_j denote its neighbors along the boundary

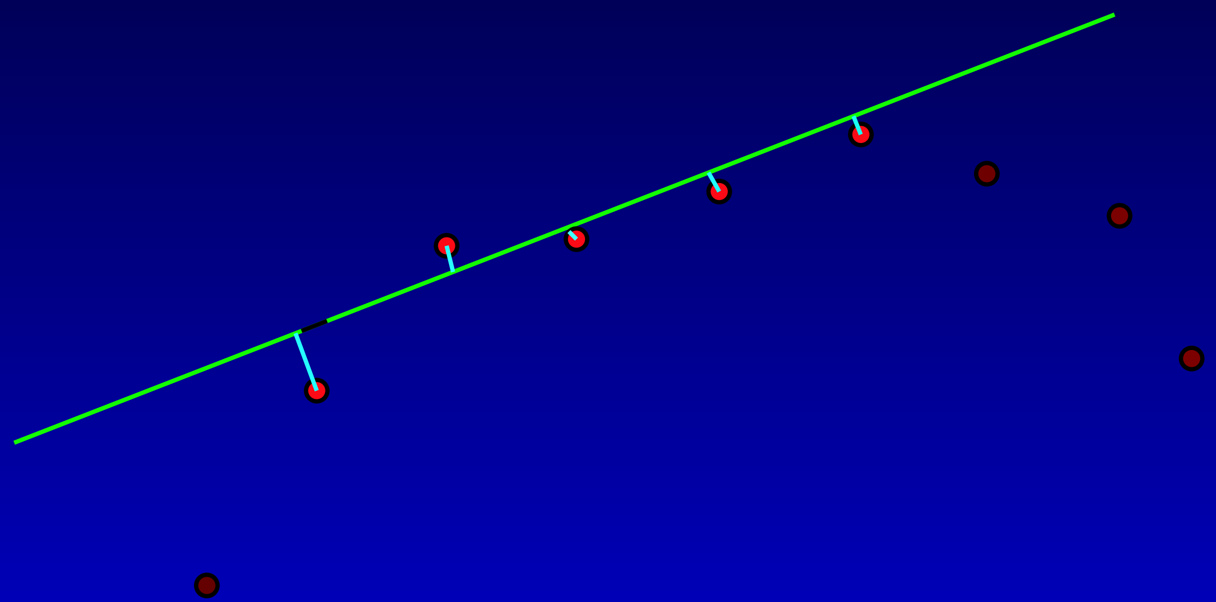

We seek a line (shown here in green) that minimizes the sum of squared perpendicular distances in a local neighborhood (shown in red) of the point of interest.

Finding the Tangent to a Curve at a Point

Find the normal which minimizes the squared projections of points onto that normal. This requires the points all be shifted to be relative to their center of gravity. $p_j \leftarrow p_j - m$, where $m = \frac{1}{N}\sum_j p_j$.

$$\sum_j (n^T p_j)^2$$

Finding the Tangent to a Curve at a Point

Unfortunately, this is minimized by $n = 0$, an interesting, but totally useless result. Fix it[a] by constraining n to be a unit vector. Use a Lagrange multiplier:

$$\sum_j (n^T p_j)^2 + \lambda n^T n = n^T \left(\sum_j p_j p_j^T \right) n + \lambda n^T n$$

[a] There is a typo on this slide. Can you find it?

Determine n

Differentiate and set the result to zero (using the derivative of a quadratic form twice:

$$2\sum_j p_j p_j^T n = 2\lambda n$$

An eigenvalue problem! n is one of the eigenvectors of the scatter matrix. But which one? The one corresponding to the larger eigenvalue (but which one?). Find the vector from the center point (x_i) to the next point (x_{i+1}),

$$\delta = \frac{x_{i+1} - x_i}{||x_{i+1} - x_i||}$$

and project the eigenvector onto δ.

Projecting onto δ

Project the principal eigenvector on this vector named δ. $p = \delta \cdot n$. If p is negative, flip the sign of the eigenvector. Let this (possibly) new eigenvector (which is the tangent vector) also be denoted n.

Measure angles in the tangent frame

Shape Context requires we find the angle (in the tangent frame) of the vector to each other point. The angle is just the inner product of the tangent vector (which we just found) with the vector to the other point.

$$\theta = \cos^{-1} n \cdot \left(\frac{p_j}{||p_j||} \right)$$

Shape Context for a Point

From a point, P, measure distance and angle to all other points. Histogram it. That histogram is the shape context for that point.

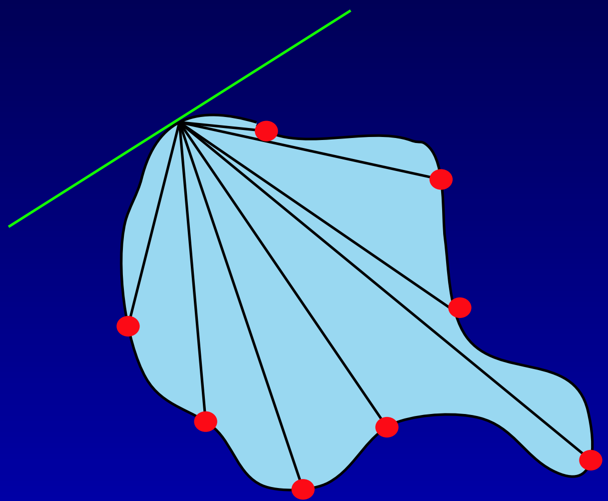

No, of course it's more complicated than that. The angle is relative to the local tangent. And the measurements are logs of distance, but that's the gist of it.

Shape Context –entire shape

Quantize angle into 12 bins, log-distance into 5 bins. Now, we have a histogram with 60 elements, which is the shape context for a single point. The difference between the SC of two points can be written

$$\gamma_{PQ} = \frac{1}{2} \sum_{k=i}^{K} \frac{(h(P,k) - h(Q,k))^2}{h(P,k) + h(Q,k)}$$

Shape Context –entire shape

All these assignment costs may be calculated and entered into a matrix.

$$\Gamma_{ij} = [\gamma_{PQ}] \quad P \in C_i, Q \in C_j \tag{1}$$

The cost of matching shapes. For a scalar measure of this cost, find the function $f : C_i \longrightarrow C_j$, a bijection (1:1 and onto), minimizes the total cost. That is, to the P^{th} element of C_i we must find exactly one element of C_j to assign it to.

The Linear Assignment Problem

The linear assignment problem is a special case of the transportation problem, which is a special case of the minimum cost flow problem, which is a special case of a linear program. Simplex algorithm– exponential cost in worst case

Special algorithms for the special cases

LAP(2)

Suppose we wish to assign n workers to n tasks, (or establish a correspondence between n points on one boundary and n on another), and assume we can compute the cost of assigning each worker to each task. Could describe those costs as a matrix:

$$\begin{bmatrix} a1 & a2 & a3 & a4 \\ b1 & b2 & b3 & b4 \\ c1 & c2 & c3 & c4 \\ d1 & d2 & d3 & d4 \end{bmatrix}$$

For example, $a3$ is the cost of assigning worker a to perform task 3.

LAP(3)

On each row, subtract the smallest member. That will leave (at least) one zero. Rename each of the other elements using primes; for example:

$$\begin{bmatrix} 0 & a2' & 0 & a4' \\ b1' & b2' & b3' & 0 \\ 0 & c2' & c3' & c4' \\ d1' & 0 & d3' & d4' \end{bmatrix}$$

LAP(4)

Now, since it is possible to choose zeros in such a way that each row and column has exactly one zero, we are done. The assignments are:

$$\begin{bmatrix} 0 & a2\prime & \mathbf{0} & a4\prime \\ b1\prime & b2\prime & b3\prime & \mathbf{0} \\ \mathbf{0} & c2\prime & c3\prime & c4\prime \\ d1\prime & \mathbf{0} & d3\prime & d4\prime \end{bmatrix}$$

LAP(5)

Of course, life isn't always so nice. Sometimes you can't just assign like last slide.

$$\begin{bmatrix} 0 & a2' & \mathbf{0} & a4' \\ b1' & b2' & b3' & \mathbf{0} \\ \mathbf{0} & c2' & c3' & c4' \\ d1' & \mathbf{0} & d3' & d4' \end{bmatrix}$$

This is too hard. Let's do an example: