

Edges and Lines

Readings: Chapter 10: 10.2.3-10.3


- better edge detectors
- line finding
- circle finding

Lines and Arcs Segmentation

In some image sets, lines, curves, and circular arcs are more useful than regions or helpful in addition to regions.


Lines and arcs are often used in

- object recognition
- stereo matching
- document analysis



Zero Crossing Operators

Motivation: The zero crossings of the second derivative of the image function are more precise than the peaks of the first derivative.



step edge
smoothed

1st derivative

2nd derivative

How do we estimate the Second Derivative?

- Laplacian Filter: $\nabla^2 f = \partial^2 f / \partial x^2 + \partial^2 f / \partial y^2$

0	1	0
1	-4	1
0	1	0

- Standard mask implementation
- Derivation: In 1D, the first derivative can be computed with mask [-1 0 1]
- The 1D second derivative is [1 -2 1]
- The Laplacian mask estimates the 2D second derivative.

How did you get those masks?

1D function $f(x)$

$[f(-1) \quad f(0) \quad f(1)]$ pixel values

$f(0)-f(-1) \quad f(1)-f(0)$ first difference

$(f(1)-f(0))-(f(0)-f(-1))$ second difference

$1f(-1)-2f(0)+1f(1)$ simplify

$[\quad 1 \quad -2 \quad 1]$ mask

and in 2D

	$f(0,1)$	
$f(-1,0)$	$f(0,0)$	$f(1,0)$
	$f(0,-1)$	

$$\begin{aligned}\partial f / \partial x(1/2) &= f(1,0) - f(0,0) \\ \partial f / \partial x(-1/2) &= f(0,0) - f(-1,0)\end{aligned}$$

$$\partial f / \partial x^2 = f(1,0) - 2f(0,0) + f(-1,0)$$

$$\partial f / \partial y^2 = f(0,1) - 2f(0,0) + f(0,-1)$$

$$\begin{aligned}\nabla^2 f &= \partial f / \partial x^2 + \partial f / \partial y^2 \\ &= 1f(1,0) - 4f(0,0) + 1f(-1,0) + 1f(0,1) + 1f(0,-1)\end{aligned}$$

0	1	0
1	-4	1
0	1	0

Properties of Derivative Masks

- Coordinates of derivative masks have opposite signs in order to obtain a high response in regions of high contrast.
- The sum of coordinates of derivative masks is zero, so that a zero response is obtained on constant regions.
- First derivative masks produce high absolute values at points of high contrast.
- Second derivative masks produce zero-crossings at points of high contrast.

History: Marr/Hildreth Operator

- First **smooth** the image via a Gaussian convolution.
- Apply a **Laplacian filter** (estimate 2nd derivative).
- Find **zero crossings** of the Laplacian of the Gaussian.

This can be done at multiple resolutions.


History: Haralick Operator

- Fit the gray-tone intensity surface to a piecewise **cubic polynomial approximation**.
- Use the approximation to **find zero crossings** of the second directional derivative in the direction that maximizes the first directional derivative.

The derivatives here are calculated from **direct mathematical expressions** wrt the cubic polynomial.

Reality: Canny Edge Detector

- Smooth the image with a Gaussian filter with spread σ .
- Compute gradient **magnitude and direction** at each pixel of the smoothed image.
- **Zero out** any pixel response \leq the two neighboring pixels on either side of it, along the direction of the gradient. This is called **nonmaximum suppression**.



- **Track high-magnitude contours.**
- **Keep only pixels along these contours**, so weak little segments go away.

Canny on Kidney



Canny Characteristics

- The Canny operator gives single-pixel-wide images with good continuation between adjacent pixels
- It is the most widely used edge operator today; no one has done better since it came out in the late 80s. Many implementations are available.
- It is very sensitive to its parameters, which need to be adjusted for different application domains.


Finding Lines

- Parameter Estimation Methods
- Tracking Methods

1. Parameter Estimation Methods


Hough Transform

- The Hough transform is a method for detecting lines or curves specified by a **parametric function**.
- If the parameters are p_1, p_2, \dots, p_n , then the Hough procedure uses an **n-dimensional accumulator** array in which it accumulates votes for the correct parameters of the lines or curves found on the image.



Finding Straight Line Segments

- $y = mx + b$ is not suitable (why?)
- The equation generally used is: $d = r \sin \theta + c \cos \theta$



d is the distance from the line to origin

θ is the angle the perpendicular makes with the column axis

Procedure to Accumulate Lines

- Set accumulator array A to all zero.
Set point list array $PTLIST$ to all NIL.
- For each pixel (R,C) in the image {
 - compute gradient magnitude $GMAG$
 - if $GMAG > \text{gradient_threshold}$ {
 - compute quantized tangent angle $THETAQ$
 - compute quantized distance to origin DQ
 - increment $A(DQ, THETAQ)$
 - update $PTLIST(DQ, THETAQ)$ } }

Example

gray-tone image

0	0	0	100	100
0	0	0	100	100
0	0	0	100	100
100	100	100	100	100
100	100	100	100	100

DQ

-	-	3	3	-
-	-	3	3	-
3	3	3	3	-
3	3	3	3	-
-	-	-	-	-

THETAQ

-	-	0	0	-
-	-	0	0	-
90	90	40	20	-
90	90	90	40	-
-	-	-	-	-


Accumulator A

360	-	-	-	-	-	-
.	-	-	-	-	-	-
6	-	-	-	-	-	-
3	4	-	1	-	2	5
0	-	-	-	-	-	-
distance	0	10	20	30	40	...90
angle						


PTLIST

360	-	-	-	-	-	-
.	-	-	-	-	-	-
6	-	-	-	-	-	-
3	*	-	*	-	*	-
0	-	-	-	-	-	-
	(1,3)	(1,4)	(2,3)	(2,4)		
					(3,1)	(3,2)
					(4,1)	(4,2)
					(4,3)	


Examples




Original Image



Detected Edges



Detected lines



The vote histogram with the detected lines marked with 'o'

Examples cont

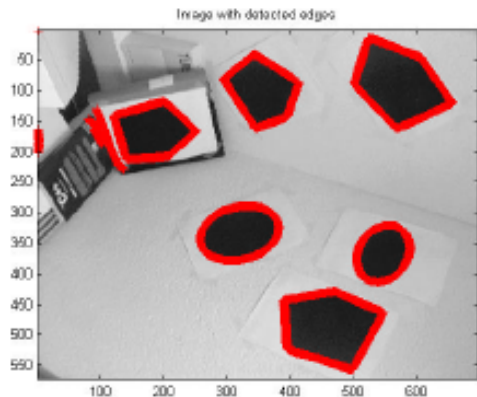


Image with detected edges

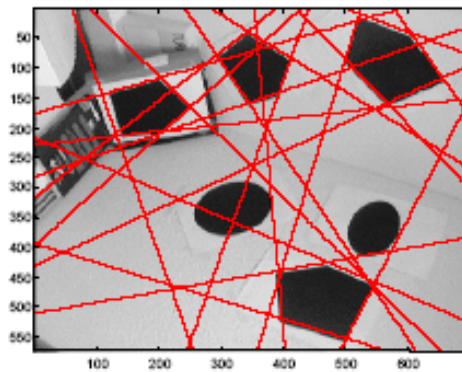



Image with the detected lines



The vote histogram with the selected lines

How do you extract the line segments from the accumulators?

pick the bin of A with highest value V

while $V > \text{value_threshold}$ {

- order the corresponding pointlist from PTLIST
- merge in high gradient neighbors within 10 degrees
- create line segment from final point list
- zero out that bin of A
- pick the bin of A with highest value V }

Line segments from Hough Transform





Fig.7. Puppet scenes 211, 212, 214, 225 and the edges recovered by the algorithm.

A Nice Hough Variant



The Burns Line Finder





1. Compute gradient magnitude and direction at each pixel.
2. For high gradient magnitude points, assign **direction labels** to **two symbolic images** for **two different quantizations**.
3. Find connected components of each symbolic image.

- Each pixel belongs to 2 components, one for each symbolic image.
- Each pixel votes for its longer component.
- Each component receives a count of pixels who voted for it.
- The components that receive majority support are selected.

Burns Example 1



Burns Example 2




2. Tracking Methods

Mask-based Approach

Use masks to identify the following events:

1. start of a new segment
2. interior point continuing a segment
3. end of a segment
4. junction between multiple segments
5. corner that breaks a segment into two



Edge Tracking Procedure

for each edge pixel P {
 classify its pixel type using masks

 case

1. isolated point : ignore it
2. start point : make a new segment
3. interior point : add to current segment
4. end point : add to current segment and finish it
5. junction or corner : add to incoming segment
 finish incoming segment
 make new outgoing segment(s)

A Good Tracking Package: the ORT toolkit

- Part of the C software available on the class web page
- Updated versions are available
- How does it work?

How ORT finds segments

(Communicated by Ata Etemadi who designed it; this is really what he said.)


- The algorithm is called Strider and is like a **spider striding** along pixel chains of an image.
- The spider is looking for **local symmetries**.
- When it is moving along a **straight or curved segment** with no interruptions, its legs are symmetric about its body.
- When it encounters an obstacle (ie. a **corner or junction**) its legs are no longer symmetric.
- If the **obstacle is small** (compared to the spider), it soon becomes symmetrical.
- If the **obstacle is large**, it will take longer.

Strider

- Strider tracks along a pixel chain, looking for junctions and corners.
- It identifies them by a measure of asymmetry.
- The accuracy depends on the length of the spider and the size of its stride.
- The larger they are, the less sensitive it becomes.

Strider

The measure of asymmetry is the angle between two line segments.




The angle must be $\leq \arctan(2/\text{length}(L2))$


Longer spiders allow less of an angle.

Strider

- The parameters are the length of the spider and the number of pixels per step.
- These parameters can be changed to allow for less sensitivity, so that we get longer line segments.
- The algorithm has a final phase in which adjacent segments whose angle differs by less than a given angle are joined.

Ort finds line segments for building detection





(410,1): 0-RGB(255,255,255)

Advantages of Strider

- works on pixel chains of arbitrary complexity
- can be implemented in parallel
- no assumptions and the effects of the parameters are well understood


Hough Transform for Finding Circles

Equations:


$$\begin{aligned} r &= r_0 + d \sin \theta \\ c &= c_0 - d \cos \theta \end{aligned}$$

r, c, d are parameters

Main idea: The gradient vector at an edge pixel points to the center of the circle.




Why it works



Filled Circle:

Outer points of circle have gradient direction pointing to center.



Circular Ring:


Outer points gradient towards center.

Inner points gradient away from center.

The points in the away direction don't accumulate in one bin!

Procedure to Accumulate Circles

- Set accumulator array A to all zero.
Set point list array PTLIST to all NIL.
- For each pixel (R,C) in the image {
For each possible value of D {
 - compute gradient magnitude GMAG
 - if GMAG > gradient_threshold {
 - . Compute THETA(R,C,D)
 - . $R0 := R - D * \sin(\text{THETA})$
 - . $C0 := C + D * \cos(\text{THETA})$
 - . increment A(R0,C0,D)
 - . update PTLIST(R0,C0,D) }}



Finding lung nodules (Kimme & Ballard)




Fig. 4.7 Using the Hough technique for circular shapes. (a) Radiograph. (b) Window. (c) Accumulator array for $r = 3$. (d) Results of maxima detection.

Summary

- The Canny edge operators is still the best one.
- The Hough transform and its variants can be used to find line segments or circles.
- It has also been generalized to find other shapes
- The original Hough method does not work well for line segments, but works very well for circles.
- The Burns method improves the Hough for line segments and is easy to code.
- The Sruider algorithm in the ORT package gives excellent line and curve segments by tracking.