Edge Detection

Computer Vision P. Schrater Spring 2003

$$\log \frac{P(\phi(I(x)) \mid \text{on-edge})}{P(\phi(I(x)) \mid \text{off-edge})} > T,$$

Simplest Model: (Canny)

Edge(x) = a U(x) + n(x)



Convolve image with U and find points with high magnitude. Choose value by comparing with a threshold determined by the noise model.



Probability of a filter response off an edge Probability of a filter response on an edge

Fundamental limits on edge detection



Figure 2.5 Empirical distributions for a gradient filter response on boundaries (left) and off boundaries (right).



Figure 2.6 The edge estimated on the glove image by ML with the filter at scale 0 (left), filter at scale 1 (centre), and filter with scales 0, 1, 2, 4 (right). Observe that ML significantly overestimates the number of edges in this image.



Need to take into account base edge rate.

Smoothing and Differentiation

- Issue: noise
 - smooth before differentiation
 - two convolutions to smooth, then differentiate?
 - actually, no we can use a derivative of Gaussian filter
 - because differentiation is convolution, and convolution is associative







There are three major issues:

- 1) The gradient magnitude at different scales is different; which should we choose?
- 2) The gradient magnitude is large along thick trail; how do we identify the significant points?
- 3) How do we link the relevant points up into curves?



1 pixel

3 pixels

7 pixels

The scale of the smoothing filter affects derivative estimates, and also the semantics of the edges recovered.

Computing Edges via the gradient



We wish to mark points along the curve where the magnitude is biggest. We can do this by looking for a maximum along a slice normal to the curve (non-maximum suppression). These points should form a curve. There are then two algorithmic issues: at which point is the maximum, and where is the next one?



Non-maximum suppression

At q, we have a maximum if the value is larger than those at both p and at r. Interpolate to get these values.







Predicting the next edge point

Assume the marked point is an edge point. Then we construct the tangent to the edge curve (which is normal to the gradient at that point) and use this to predict the next points (here either r or s).



Remaining issues

- Check that maximum value of gradient value is sufficiently large
 - drop-outs? use hysteresis
 - use a high threshold to start edge curves and a low threshold to continue them.

Edge Detection Algorithm

- Define Edge filter: $G_{uv} = \exp(-\frac{1}{2}\left[\frac{(x_u N/2)^2}{pixelwidth^2} + \frac{(y_v M/2)^2}{pixelwidth^2}\right])$ $G = \begin{bmatrix} D_{ij} = [-1 & 0 & 1] \\ D_x = G \otimes D \\ Dy = G \otimes D^T \end{bmatrix}$
- Next Compute Gradient

Gradient Magnitude

$$I_{x} = Dx \otimes I$$

$$I_{y} = Dy \otimes I$$

$$\|\nabla I(i,j)\| = \sqrt{I_{x}(i,j)^{2} + I_{y}(i,j)^{2}}$$

$$\nabla I(i,j) = \begin{bmatrix} I_{x}(i,j) \\ I_{y}(i,j) \end{bmatrix}$$

$$\angle \nabla I(i,j) = \tan^{-1} \left(\frac{I_{y}(i,j)}{I_{x}(i,j)} \right)$$

Gradient Direction

Edge Detection Algorithm

- Find an initial point (i,j) such that: $\|\nabla I(i,j)\| > c$ Where c is relatively large
- Find nearby points in the direction of (i,j)'s gradient $\vec{n}(i,j) = \begin{bmatrix} I_x(i,j) / \|\nabla I(i,j)\| \\ I_y(i,j) / \|\nabla I(i,j)\| \end{bmatrix}$ Find nearby points $\begin{bmatrix} u \\ v \end{bmatrix}$ such that $n(u,v)^T \vec{n}(i,j) < \varepsilon$ Choose $\begin{bmatrix} u \\ v \end{bmatrix}$ with maximal gradient magnitude

Edge Detection Algorithm

- While set of (i,j) with $\|\nabla I(i,j)\|$ **bot** visited
 - Find a start point, erasing points that have been checked.
 - While possible, expand a chain through the current point by:
 - 1) predicting a next set of points (not visited) using the direction perp. to gradient
 - 2) Finding which of set (if any) is a local maximum
 - -3) Test if grad mag for local max > k
 - 4) Recording all the set as visited
 - record point and set to current point.
 - end
- end

Notice

- Something nasty is happening at corners
- Scale affects contrast
- Edges aren't bounding contours





fine scale high threshold







Zero mean image, -1:1 scale





Positive responses

Zero mean image, -max:max scale





Zero mean image, -1:1 scale





Positive responses

Zero mean image, -max:max scale





Figure from "Computer Vision for Interactive Computer Graphics," W.Freeman et al, IEEE Computer Graphics and Applications, 1998 copyright 1998, IEEE

The Laplacian of Gaussian

- Another way to detect an extremal first derivative is to look for a zero second derivative
- Bad idea to apply a Laplacian without smoothing
 - smooth with Gaussian, apply Laplacian
 - this is the same as filtering with a _ Laplacian of Gaussian filter
- Now mark the zero points *where there* is a sufficiently large derivative, and enough contrast

• In 2D:





contrast=1

LOG zero crossings contrast=4





We still have unfortunate behaviour at corners



Orientation representations

- The gradient magnitude is affected by illumination changes
 - but it's direction isn't
- We can describe image patches by the swing of the gradient orientation

- Important types:
 - constant window
 - small gradient mags
 - edge window
 - few large gradient mags in one direction
 - flow window
 - many large gradient mags in one direction
 - corner window
 - large gradient mags that swing

Representing Windows

- Types
 - constant
 - small eigenvalues
 - Edge
 - one medium, one small
 - Flow
 - one large, one small
 - corner
 - two large eigenvalues

$$H = \sum_{window} (\nabla I) (\nabla I)^T$$





9	9	×.	÷	2	1	×.	2	E.	×.	1	×.	×.	N	\$	0	ø	0	-	-	e.	0	Ø	N	ъ	N	•	٩.	ι.
٥	¢	R.	F.	н.	P	а,	ы		•	1	1	Z.	-		-	P	æ,	-1	2	2	1	0	D	L.	٥	•	•	8
	κ.	۰.	1	0	ų,	0	N	•	•	0	0	7	1	•	х.	2	•	•	1	10	1	7	0	ũ.	9	٩.	6	٥
-	•	•		-	0	Ø	•	•		1	Z	0	Z	1	0	æ,	o	•	2	2			a.	ō	٥	•	D.	0
		•	•	•	z	r	•	х.	х.	•	¥.	1	1	Ó	0	Б.	•	×.	٥	z	1		e.	o	Ø.	۵		-
•			•		×,	2	z		•		10	1	1	o	0	11	•		9.	И	2	1	×.	1	a.	٥	•	
1	31	N	10	- 1	-	ъ	2	Z.		•			ъ	9	a.	1				e	0	1	2	х.	1	e.	•	
	11	E.	11	4	8	0	0	2	k.	1			11	а.	11	н.	2	0	8	o	5	0	ø					
a.	1	4	8	Δ.	ñ,	Ô.	0	•	9	•		т.		•	Т	х.		н.	7	ø	ъ.	•	0	1	•	•		
0	6		5	0	6	ð	ø									1	κ.						0	0	0			
	0	8	ς.	5		0							4		9	•				×.	σ.	-0	μ.	0	0			1
	B		0	6	b.	2	-		τ.						5	D						-	2	_				
		,	7	Ĭ.	i.	Ľ,			ė.	a.	0		N	1		0			9			0	h	Ц			Ξ.	
12	4	2	8		÷.							E	ě	÷.	10	n					1		Ā					
F.	Ľ.	U.	÷.	Ū.	Ū.					u	R	F	Я	Ă.	Š.	÷.		9	Ì.	ù,		Ě.	2			1		
14						÷.	1	1		7	1	ų,	2	2	U.	ų.	÷.		÷.	Ū.		ŭ.		1	2	2	2	
		01			ŝ.	λ.				Å.	2	1	Č.	Ľ.	Ξ.	1	Ĵ.	1	0		П.	÷.	0	2	8	0	3	2
1	5		2	1	d.	2	5	2	×.	Ľ	Ľ.	Ű.	0	1.	5	Ű.	Ű.			÷.	U.	а.	Ĵ.	2	×.	2		ų,
1	4		2	1	0		5	2	X.	2	2	1	1		۰.	0		0	1	1		0	ľ.	9	2	Ľ.	2	<u>H.</u>
			5	2	3	0	P*	9	Ч,	¥.	"	•	1		1	٢.	11	1	N	1		9	ø	2	•	۲	~	
		•	4	К	2	6	10	0	9	ų.	0	2	۰.	2			1	1	۰.						2	٩.		
1	p	8	1	$\boldsymbol{\nu}$	P.	6	P	1	0	U	0	2	1		•	•	•					1	÷.			1	-	
9	0	1	2	•	D	φ	0	Q	2	0	9					1	8	E,		1	10	10	۰.					
Ð	Ø	0			ę	•	0	0	9	2	1	0	2	•	•		•			8	-	н.	2	П	4			
т	1	•		•	1	Ø	Q	ц,	η.	ŝ.	0	Q		P.	•	0	8	1	2		0				0	4	•	
0		2	9	Ð	ł,	٩	α		Q	0	0	٩.	×.	9	Q	Q			2	0	Ø	0	1	1	н.	•		•
0	•	b	¢,	0	۰	-	-	0	Ū	0	0	8	8	Ð	Q	1	0		Ξ	o	0	0	-		2	1		-
0	4	N	٦.	•	٠	•	•	•	Q	ο	ю	0	-	Ε.	٠	•	•	2	-	e	P	-	*			ø		
			ß	0	8	٠	9	0	4	0	0	0	-	•	٠	•	•		-	•	-	-	•	-	đ	ē	~	-
	1	5	0	Ð	6	0		Q	0	0	0	1	1	•	•		-	1		0	•>	-	-	•	4	ь		-00
							8		8																			



Filters are templates

- Applying a filter at some point can be seen as taking a dotproduct between the image and some vector
- Filtering the image is a ot products

- Insight
 - filters look like the effects they are intended to find
 - filters find effects they look like



Normalized correlation

- Think of filters of a dot product
 - now measure the angle
 - i.e normalised
 correlation output is
 filter output, divided
 by root sum of squares
 of values over which
 filter lies

- Tricks:
 - ensure that filter has a zero response to a constant region (helps reduce response to irrelevant background)
 - subtract image average when computing the normalizing constant (i.e. subtract the image mean in the neighbourhood)
 - absolute value deals with contrast reversal