

RWTH AACHEN
UNIVERSITY

Computer Vision - Lecture 4

Gradients & Edges

02.11.2016

Computer Vision WS 16/17

Bastian Leibe
RWTH Aachen
<http://www.vision.rwth-aachen.de>
leibe@vision.rwth-aachen.de

RWTH AACHEN
UNIVERSITY

Announcements

- Exercise sheet 2 is available
 - Thresholding, Morphology
 - Gaussian smoothing
 - Image gradients
 - Edge Detection
 ⇒ *Deadline: Sunday night, 13.11. (in two weeks).*
- Reminder
 - You're encouraged to form teams of up to 3 people!
 - Hints:
 - Turn in everything as a single zip archive.
 - Use the provided Matlab framework.
 - For each exercise, you need to implement the corresponding `apply` function. If the screen output matches the expected output (shown in class), you will know that your solution is correct.
 - Matlab helps you to find errors (red lines under your code)!

2
B. Leibe

RWTH AACHEN
UNIVERSITY

Course Outline

- Image Processing Basics
 - Image Formation
 - Binary Image Processing
 - Linear Filters
 - Edge & Structure Extraction
- Segmentation
- Local Features & Matching
- Object Recognition and Categorization
- 3D Reconstruction
- Motion and Tracking

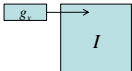
Computer Vision WS 16/17

4
B. Leibe

RWTH AACHEN
UNIVERSITY

Topics of This Lecture

- Recap: Linear Filters
- Multi-Scale representations
 - How to properly rescale an image?
- Filters as templates
 - Correlation as template matching
- Image gradients
 - Derivatives of Gaussian
- Edge detection
 - Canny edge detector



Computer Vision WS 16/17

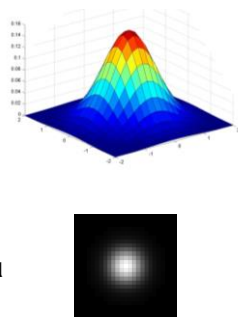
5
B. Leibe

RWTH AACHEN
UNIVERSITY

Recap: Gaussian Smoothing

- Gaussian kernel

$$G_\sigma = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}}$$
- Rotationally symmetric
- Weights nearby pixels more than distant ones
 - This makes sense as 'probabilistic' inference about the signal
- A Gaussian gives a good model of a fuzzy blob




Computer Vision WS 16/17

6
B. Leibe
Image Source: Forsyth & Ponce

RWTH AACHEN
UNIVERSITY

Recap: Smoothing with a Gaussian

- Parameter σ is the "scale" / "width" / "spread" of the Gaussian kernel and controls the amount of smoothing.



```

for sigma=1:3:10
    h = fspecial('gaussian', fsize, sigma);
    out = imfilter(im, h);
    imshow(out);
    pause;
end
  
```

Computer Vision WS 16/17

7
B. Leibe
Slide credit: Kristen Grauman

Recap: Effect of Filtering

- Noise introduces high frequencies. To remove them, we want to apply a "low-pass" filter.
- The ideal filter shape in the frequency domain would be a box. But this transfers to a spatial sinc, which has infinite spatial support.
- A compact spatial box filter transfers to a frequency sinc, which creates artifacts.
- A Gaussian has compact support in both domains. This makes it a convenient choice for a low-pass filter.

Computer Vision WS 16/17

B. Leibe

8

Recap: Low-Pass vs. High-Pass

Original image

Low-pass filtered

High-pass filtered

Computer Vision WS 16/17

B. Leibe

Image Source: S. Chens

9

Topics of This Lecture

- Recap: Linear Filters
- Multi-Scale representations**
 - How to properly rescale an image?
- Filters as templates
 - Correlation as template matching
- Image gradients
 - Derivatives of Gaussian
- Edge detection
 - Canny edge detector

Computer Vision WS 16/17

B. Leibe

10

Motivation: Fast Search Across Scales

Computer Vision WS 16/17

B. Leibe

Image Source: Irani & Basri

11

Recap: Sampling and Aliasing

Computer Vision WS 16/17

B. Leibe

Image Source: Forsyth & Ponce

12

Recap: Sampling and Aliasing

Computer Vision WS 16/17

B. Leibe

Image Source: Forsyth & Ponce

13

Computer Vision WS 16/17

Recap: Sampling and Aliasing

Signal → Fourier Transform → Magnitude Spectrum

Sample → Sampled Signal → Fourier Transform → Magnitude Spectrum

Inaccurately Reconstructed Signal ← Inverse Fourier Transform

Copy and Shift

Cut out by multiplication with box filter

B. Leibe
Image Source: Forsyth & Ponce

Computer Vision WS 16/17

Recap: Resampling with Prior Smoothing

256 × 256 128 × 128 64 × 64 32 × 32 16 × 16

Artifacts!

no smoothing

Gaussian $\sigma = 1$

Gaussian $\sigma = 2$

- Note: We cannot recover the high frequencies, but we can avoid artifacts by smoothing before resampling.

B. Leibe
Image Source: Forsyth & Ponce

Computer Vision WS 16/17

The Gaussian Pyramid

Low resolution

High resolution

$G_4 = (G_3 * \text{gaussian}) \downarrow 2$

$G_3 = (G_2 * \text{gaussian}) \downarrow 2$ down-sample

$G_2 = (G_1 * \text{gaussian}) \downarrow 2$ blur down-sample

$G_1 = (G_0 * \text{gaussian}) \downarrow 2$ blur down-sample

$G_0 = \text{Image}$ blur

B. Leibe
Source: Irani & Basri

Computer Vision WS 16/17

Gaussian Pyramid - Stored Information

All the extra levels add very little overhead for memory or computation!

B. Leibe
Source: Irani & Basri

Computer Vision WS 16/17

Summary: Gaussian Pyramid

- Construction: create each level from previous one
 - Smooth and sample
- Smooth with Gaussians, in part because
 - a Gaussian * Gaussian = another Gaussian
 - $G(\sigma_1) * G(\sigma_2) = G(\sqrt{\sigma_1^2 + \sigma_2^2})$
- Gaussians are low-pass filters, so the representation is redundant once smoothing has been performed.
 - ⇒ There is no need to store smoothed images at the full original resolution.

Slide credit: David Lowe
B. Leibe

Computer Vision WS 16/17

The Laplacian Pyramid

Gaussian Pyramid

$L_i = G_i - \text{expand}(G_{i+1})$

Laplacian Pyramid

$L_n = G_n$

G_2 expand

G_1 expand

G_0 expand

L_2

L_1

L_0

Why is this useful?

B. Leibe
Source: Irani & Basri

RWTH AACHEN UNIVERSITY

Laplacian ~ Difference of Gaussian

DoG = Difference of Gaussians
 Cheap approximation - no derivatives needed.

20

RWTH AACHEN UNIVERSITY

Topics of This Lecture

- Recap: Linear Filters
- Multi-Scale representations
 - How to properly rescale an image?
- Filters as templates
 - Correlation as template matching
- Image gradients
 - Derivatives of Gaussian
- Edge detection
 - Canny edge detector

21

RWTH AACHEN UNIVERSITY

Note: Filters are Templates

- Applying a filter at some point can be seen as taking a dot-product between the image and some vector.
- Filtering the image is a set of dot products.
- Insight
 - Filters look like the effects they are intended to find.
 - Filters find effects they look like.

22

RWTH AACHEN UNIVERSITY

Where's Waldo?

23

RWTH AACHEN UNIVERSITY

Where's Waldo?

24

RWTH AACHEN UNIVERSITY

Where's Waldo?


25

RWTH AACHEN UNIVERSITY

Correlation as Template Matching

- Think of filters as a dot product of the filter vector with the image region
 - Now measure the angle between the vectors

$$a \cdot b = |a| |b| \cos \theta \quad \cos \theta = \frac{a \cdot b}{|a| |b|}$$
 - Angle (similarity) between vectors can be measured by normalizing the length of each vector to 1 and taking the dot product.



Template


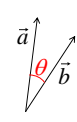


Image region





Vector interpretation

26

RWTH AACHEN UNIVERSITY

Topics of This Lecture

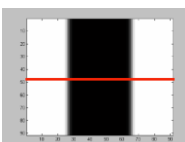
- Recap: Linear Filters
- Multi-Scale representations
 - How to properly rescale an image?
- Filters as templates
 - Correlation as template matching
- Image gradients
 - Derivatives of Gaussian
- Edge detection
 - Canny edge detector

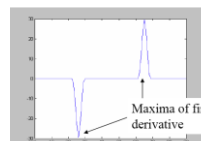



27

RWTH AACHEN UNIVERSITY

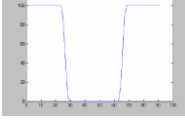
Derivatives and Edges...

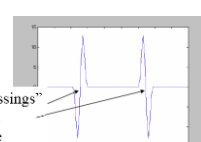




1st derivative

Maxima of first derivative





"zero crossings" of second derivative

28

RWTH AACHEN UNIVERSITY

Differentiation and Convolution

- For the 2D function $f(x,y)$, the partial derivative is:

$$\frac{\partial f(x,y)}{\partial x} = \lim_{\epsilon \rightarrow 0} \frac{f(x+\epsilon, y) - f(x, y)}{\epsilon}$$
- For discrete data, we can approximate this using finite differences:

$$\frac{\partial f(x,y)}{\partial x} \approx \frac{f(x+1, y) - f(x, y)}{1}$$
- To implement the above as convolution, what would be the associated filter?


1	-1
---	----

29

RWTH AACHEN UNIVERSITY

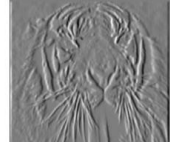
Partial Derivatives of an Image

$\frac{\partial f(x,y)}{\partial x}$



$\frac{\partial f(x,y)}{\partial y}$

-1	1
----	---



-1	?	1
1	or	-1

Which shows changes with respect to x?

30

RWTH AACHEN UNIVERSITY

Assorted Finite Difference Filters

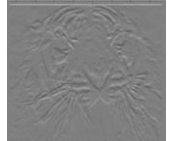
Prewitt: $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$

Sobel: $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$

Roberts: $M_x = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$

```

>> My = fspecial('sobel');
>> outim = imfilter(double(im), My);
>> imagesc(outim);
>> colormap gray;
    
```



31

Image Gradient

- The gradient of an image:

$$\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right]$$
- The gradient points in the direction of most rapid intensity change
- The gradient direction (orientation of edge normal) is given by:

$$\theta = \tan^{-1} \left(\frac{\partial f / \partial y}{\partial f / \partial x} \right)$$
- The edge strength is given by the gradient magnitude

$$\|\nabla f\| = \sqrt{\left(\frac{\partial f}{\partial x} \right)^2 + \left(\frac{\partial f}{\partial y} \right)^2}$$

Slide credit: Steve Seitz. B. Leibe

Effect of Noise

- Consider a single row or column of the image
 - Plotting intensity as a function of position gives a signal

Where is the edge?

Slide credit: Steve Seitz. B. Leibe

Solution: Smooth First

Where is the edge? Look for peaks in $\frac{\partial}{\partial x}(h * f)$

Slide credit: Steve Seitz. B. Leibe

Derivative Theorem of Convolution

$$\frac{\partial}{\partial x}(h * f) = \left(\frac{\partial}{\partial x} h \right) * f$$

- Differentiation property of convolution.

Slide credit: Steve Seitz. B. Leibe

Derivative of Gaussian Filter

$$g * (h * I) = (g * h) * I$$

$$\begin{bmatrix} 1 & -1 \end{bmatrix} * \begin{bmatrix} 0.0030 & 0.0133 & 0.0219 & 0.0133 & 0.0030 \\ 0.0133 & 0.0596 & 0.0983 & 0.0596 & 0.0133 \\ 0.0219 & 0.0983 & 0.1621 & 0.0983 & 0.0219 \\ 0.0133 & 0.0596 & 0.0983 & 0.0596 & 0.0133 \\ 0.0030 & 0.0133 & 0.0219 & 0.0133 & 0.0030 \end{bmatrix}$$

Why is this preferable?

Slide adapted from Kristen Grauman. B. Leibe

Derivative of Gaussian Filters

x-direction y-direction

Slide credit: Steve Seitz. B. Leibe. Source: Sreethana Lazebnik

RWTH AACHEN UNIVERSITY

Laplacian of Gaussian (LoG)

- Consider $\frac{\partial^2}{\partial x^2}(h * f)$

f

Signal

$\frac{\partial^2}{\partial x^2}h$

Kernel

$(\frac{\partial^2}{\partial x^2}h) * f$

Convolution

Sigma = 50

Where is the edge? Zero-crossings of bottom graph

Slide credit: Steve Seitz B. Leibe 38

RWTH AACHEN UNIVERSITY

Summary: 2D Edge Detection Filters

Gaussian

$$h_\sigma(u, v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}}$$

Derivative of Gaussian

$$\frac{\partial}{\partial x} h_\sigma(u, v)$$

Laplacian of Gaussian

$$\nabla^2 h_\sigma(u, v)$$

- ∇^2 is the Laplacian operator:

$$\nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$

Slide credit: Kristen Grauman B. Leibe 39

RWTH AACHEN UNIVERSITY

Topics of This Lecture

- Recap: Linear Filters
- Multi-Scale representations
 - How to properly rescale an image?
- Filters as templates
 - Correlation as template matching
- Image gradients
 - Derivatives of Gaussian
- Edge detection
 - Canny edge detector

Computer Vision WS 16/17 B. Leibe 40

RWTH AACHEN UNIVERSITY

Edge Detection

- Goal: map image from 2D array of pixels to a set of curves or line segments or contours.
- Why?

Figure from J. Shotton et al., PAMI 2007

- Main idea: look for strong gradients, post-process

Slide credit: Kristen Grauman, David Lowe B. Leibe 41

RWTH AACHEN UNIVERSITY

Designing an Edge Detector

- Criteria for an "optimal" edge detector:
 - Good detection:** the optimal detector should minimize the probability of false positives (detecting spurious edges caused by noise), as well as that of false negatives (missing real edges).
 - Good localization:** the edges detected should be as close as possible to the true edges.
 - Single response:** the detector should return one point only for each true edge point; that is, minimize the number of local maxima around the true edge.

True edge

Poor robustness to noise

Poor localization

Too many responses

Source: Li Fei-Fei B. Leibe 42

RWTH AACHEN UNIVERSITY

Gradients → Edges

Primary edge detection steps

- Smoothing: suppress noise
- Edge enhancement: filter for contrast
- Edge localization
 - Determine which local maxima from filter output are actually edges vs. noise
 - Thresholding, thinning

- Two issues
 - At what scale do we want to extract structures?
 - How sensitive should the edge extractor be?

adapted from Kristen Grauman B. Leibe 44

RWTH AACHEN UNIVERSITY

Scale: Effect of σ on Derivatives

$\sigma = 1 \text{ pixel}$ $\sigma = 3 \text{ pixels}$

- The apparent structures differ depending on Gaussian's scale parameter.

⇒ Larger values: larger-scale edges detected
 ⇒ Smaller values: finer features detected

Computer Vision WS 16/17 45

Slide credit: Kristen Grauman B. Leibe

RWTH AACHEN UNIVERSITY

Sensitivity: Recall Thresholding

- Choose a threshold t
- Set any pixels less than t to zero (off).
- Set any pixels greater than or equal t to one (on).

$$F_t[i, j] = \begin{cases} 1, & \text{if } F[i, j] \geq t \\ 0, & \text{otherwise} \end{cases}$$

Computer Vision WS 16/17 46

B. Leibe

RWTH AACHEN UNIVERSITY

Original Image

Computer Vision WS 16/17 47

Slide credit: Kristen Grauman B. Leibe

RWTH AACHEN UNIVERSITY

Gradient Magnitude Image

Computer Vision WS 16/17 48

Slide credit: Kristen Grauman B. Leibe

RWTH AACHEN UNIVERSITY

Thresholding with a Lower Threshold

Computer Vision WS 16/17 49

Slide credit: Kristen Grauman B. Leibe

RWTH AACHEN UNIVERSITY

Thresholding with a Higher Threshold

Computer Vision WS 16/17 50

Slide credit: Kristen Grauman B. Leibe

RWTH AACHEN UNIVERSITY

Canny Edge Detector

- Probably the most widely used edge detector in computer vision
- Theoretical model: step-edges corrupted by additive Gaussian noise
- Canny has shown that the first derivative of the Gaussian closely approximates the operator that optimizes the product of *signal-to-noise ratio* and localization.


J. Canny, *A Computational Approach To Edge Detection*, *IEEE Trans. Pattern Analysis and Machine Intelligence*, 8:679-714, 1986.

51

B. Leibe Source: Li Fei-Fei

RWTH AACHEN UNIVERSITY

The Canny Edge Detector




Original image (Lenna)

52

B. Leibe Slide credit: Kristen Grauman

RWTH AACHEN UNIVERSITY

The Canny Edge Detector



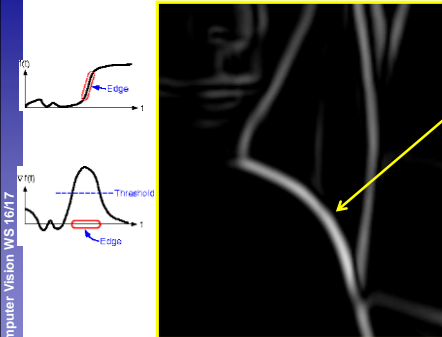
Gradient magnitude

53

B. Leibe Slide credit: Kristen Grauman

RWTH AACHEN UNIVERSITY

The Canny Edge Detector



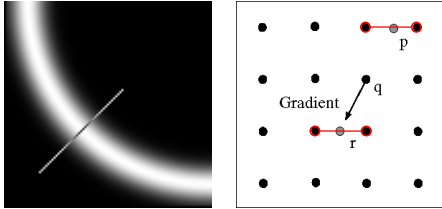
How to turn these thick regions of the gradient into curves?

55

B. Leibe Slide credit: Kristen Grauman

RWTH AACHEN UNIVERSITY

Non-Maximum Suppression



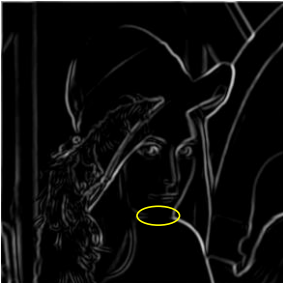
- Check if pixel is local maximum along gradient direction, select single max across width of the edge
 - Requires checking interpolated pixels p and r
 - ⇒ Linear interpolation based on gradient direction

56

B. Leibe Source: Forsyth & Ponce

RWTH AACHEN UNIVERSITY

The Canny Edge Detector



Thinning
(non-maximum suppression)

Problem: pixels along this edge didn't survive the thresholding.

57


B. Leibe Slide credit: Kristen Grauman

RWTH AACHEN UNIVERSITY

Solution: Hysteresis Thresholding

- Hysteresis: A lag or momentum factor
- Idea: Maintain two thresholds k_{high} and k_{low}
 - Use k_{high} to find strong edges to start edge chain
 - Use k_{low} to find weak edges which continue edge chain
- Typical ratio of thresholds is roughly


$$k_{high} / k_{low} = 2$$




58
Source: D. Lowe, S. Seitz

RWTH AACHEN UNIVERSITY


Hysteresis Thresholding




Original image



High threshold (strong edges)



Low threshold (weak edges)



Hysteresis threshold

59
Source: L. Fei-Fei

RWTH AACHEN UNIVERSITY

Summary: Canny Edge Detector

1. Filter image with derivative of Gaussian
2. Find magnitude and orientation of gradient
3. Non-maximum suppression:
 - Thin multi-pixel wide "ridges" down to single pixel width
4. Linking and thresholding (hysteresis):
 - Define two thresholds: low and high
 - Use the high threshold to start edge curves and the low threshold to continue them

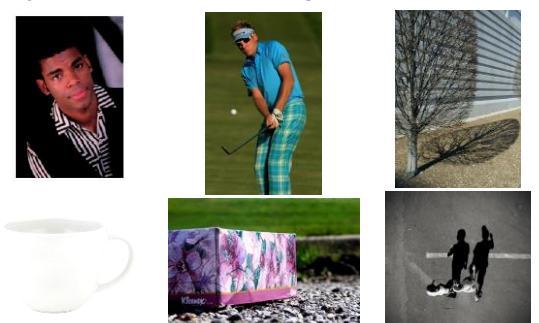
- MATLAB:


```
>> edge(image, 'canny');
>> help edge
```

60
Source: D. Lowe, L. Fei-Fei

RWTH AACHEN UNIVERSITY

Object Boundaries vs. Edges









Background Texture Shadows

61
Slide credit: Kristen Grauman

RWTH AACHEN UNIVERSITY

Edge Detection is Just the Beginning...

Image	Human segmentation	Gradient magnitude
		
		

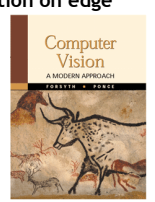
- Berkeley segmentation database:
 <http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/>

62
Source: L. Lazebnik

RWTH AACHEN UNIVERSITY

References and Further Reading

- Background information on linear filters and their connection with the Fourier transform can be found in Chapter 7 of F&P. Additional information on edge detection is available in Chapter 8.
 - D. Forsyth, J. Ponce, *Computer Vision - A Modern Approach*. Prentice Hall, 2003



63