Mechatronic Engineering program

Basics of AI and Deep Learning:
Computer Vision - An Introduction

Krzysztof Holak
AGH University of Krakow

1

Schedule

- **Lecture 1: An Introduction**
- Image Preprocessing,
- Image Filtering
- Edge Detection

- Lecture 2: Image Segmentation
- Lecture 3: Image Features
- Lecture 4: Video Processing

2

What is an image?

Image as **2D function** $z = f(x,y)$

x, y – coordinates of points in the image
z – value of the points (intensity, color etc.)

This approach lets us **make** image processing algorithms in terms of mathematical approach (e.g. image gradient)

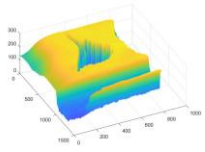

3

What is an image?

Image as **a matrix** of **pixels** A = A(i,j)

i, j – row and column of a pixel in a matrix
A(i,j) – value of the pixels (intensity, color etc.)

This approach lets us **implement** our image processing algorithms
(after all, all computer objects are discrete!)

4

Image representation

Let's think about our images as matrices A(i,j) of size MxN
Digital image is **discretized** and **quantized**

M,N – number of rows and columns in our matrix – **image spatial resolution**
Number of bits per pixel to represent color or value – **bit depth**

Grayscale image – 8 bits per pixel (256 gray levels)
Color image (True Color) – 8 bits per color channel, 24 bits per pixels
Binary image (Black and White) – 1 bit per pixel

Multispectral, Hyperspectral…

5

Color images

RGB Color Space

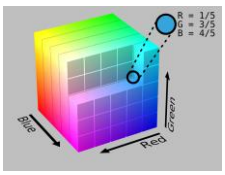

Default in Matlab
Non-intuitive
Channels correlated

HSV Color Space

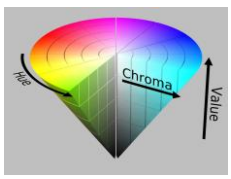

Intuitive
Decoupled color and intensity information

6

RGB Color images

R

G

B

7

HSV Color images

H

S

V

8

Image Histogram

Image histogram – a bar graph representing number of pixels of a given intensity value

Histogram characterizes the brightness, contrast and intensity value spread over the image – tells a lot how the image looks like!
Discrete version of probability density function

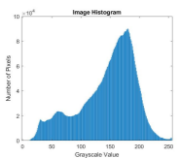

9

Image Histogram examples

A dark image

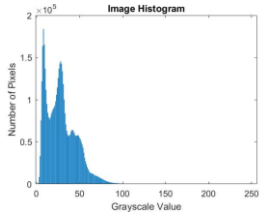

10

Image Histogram examples

High contrast image vs low contract image

11

Image Histogram examples

A high contrast image

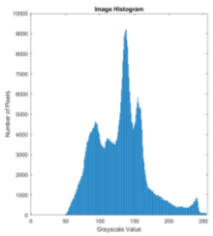

12

Image Histogram examples

A low contrast image

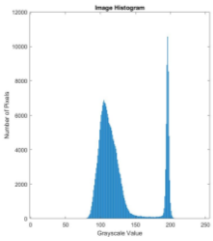

13

Image Histogram examples

Compared low contrast and high contrast noise images and histograms

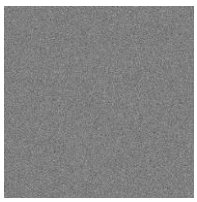

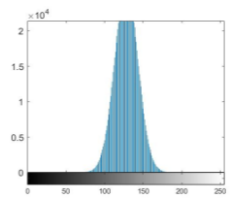

14

Image Histogram examples

Compared low contrast and high contrast noise images and histograms

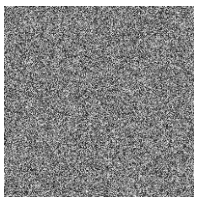

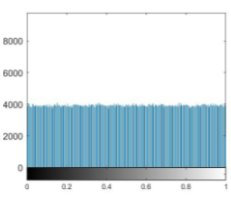

15

Image Equalization

Previous example gives us a way to improve a contrast of an image by histogram stretching – **histogram equalization**

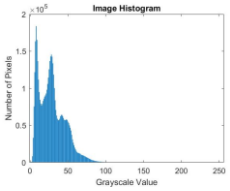

16

Image Equalization

Proof from **probability theory** – the probability densities for monotonic functions of random variables

17

Histograms for color images

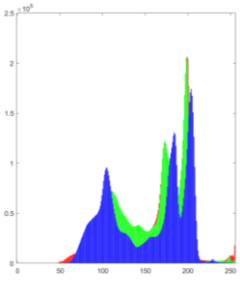

18

Histograms equalization for color images

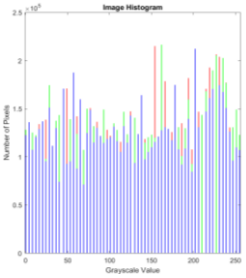

19

Histograms in practice

Use histogram to find a threshold an binarize an image!

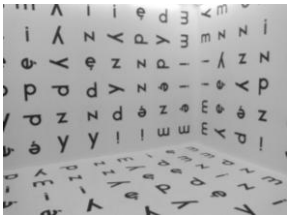

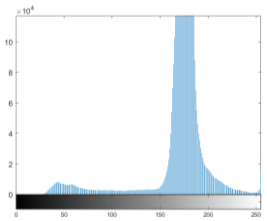

20

Histograms in practice

Use histogram to find a threshold an binarize an image!

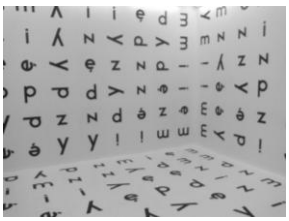

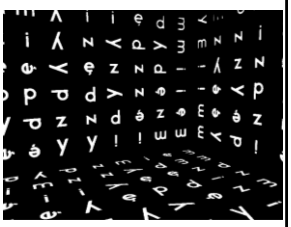

21

Basic Operations on Images

- **Point Processing**

  The output value of a pixel is a function of input pixel's value
  $y = f(x)$

- **Local Processing**

  The output value of a pixel is a function of input pixel's value and value of its neighbors
  $y = f(x,h(x))$

- **Global Processing**

  The output value of a pixel is a function of all pixels' values in the image
  $y = f(x,I)$

22

Basic Operations on Images

Classification of operations on images

- **Point Processing**

  Change of brightness
  Gamma correction
  Arithmetic operations
  Logic operations
  Thresholding

- Local Processing
- Global Processing

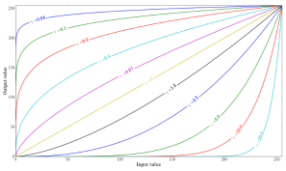

Source: OpenCV

23

Basic Operations on Images

- Point Processing
- **Local Processing**

  Linear filters
  Nonlinear filters
  Morphological filters
  Edge detection
  Image correlation

- Global Processing

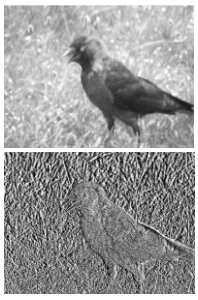

24

Basic Operations on Images

- Point Processing
- Local Processing
- **Global Processing**

  Image Transforms:
  - Fourier Transform
  - Cosine Transform
  - Wavelet Transform

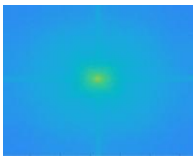

25

Image Filters

- To understand the concept of filtering, let us return to the idea that an image is 2D function - or even better, a **2D signal**!
- Image, as any signal has frequency representation
  It contains **spatial frequencies (SF)!**
  Two sets of them (in **x and y directions**)

X direction

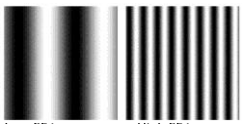

Low SF image High SF image

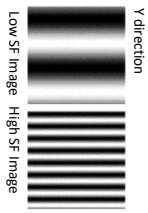

26

Image Filters

Low frequency conents

Average brightness (zero frequency, DC content)
Slowly varying brightnes across an image

High frequency content

Rapidly varying brightness across an image
Sharp Edges
Details
Noise

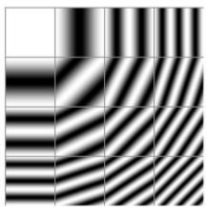

27

Linear image filters

Mathematically, they are realized as a **discrete convolution** of a filter mask with an image
Results depend on filter mask

$$h[m,n] = \sum_{k,l} g[k,l]\ f[m+k,n+l]$$

Example: g - a box filter of size 3x3
f- a simple binary image

$$\frac{1}{9}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

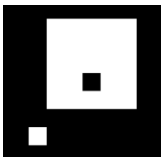

28

Discrete convolution

$$\frac{1}{9}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

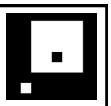

$$h[m,n] = \sum_{k,l} g[k,l]\ f[m+k,n+l]$$

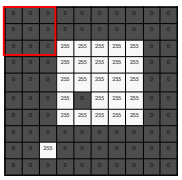

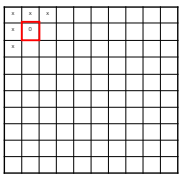

29

Discrete convolution

$$\frac{1}{9}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

$$h[m,n] = \sum_{k,l} g[k,l]\ f[m+k,n+l]$$

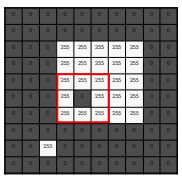

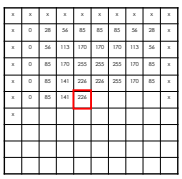

30

Filtering Results

$\frac{1}{9}\begin{bmatrix}1&1&1\\1&1&1\\1&1&1\end{bmatrix}$

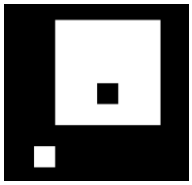

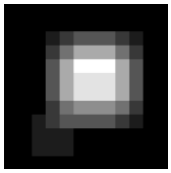

What do we observe?
Blurred edges, „noise" reduced, smaller image (boundary „cut")

MATLAB

31

Low Pass Filters

Masks have only positive coefficients
Sum of all coefficients is 1 (no change in brightness!)
Coefficients decide on the impact on neighbors on central pixel value

Examples: box filter, **Gaussian filter**

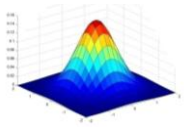

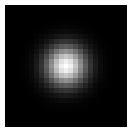

$$G_\sigma = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}}$$

32

Filtering Results

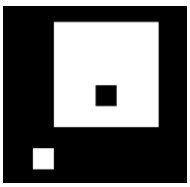

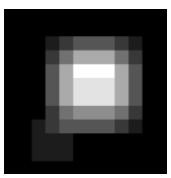

$\frac{1}{9}\begin{bmatrix}1&1&1\\1&1&1\\1&1&1\end{bmatrix}$

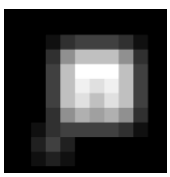

$\frac{1}{16}\begin{bmatrix}1&2&1\\2&4&2\\1&2&1\end{bmatrix}$

33

Filtering Results

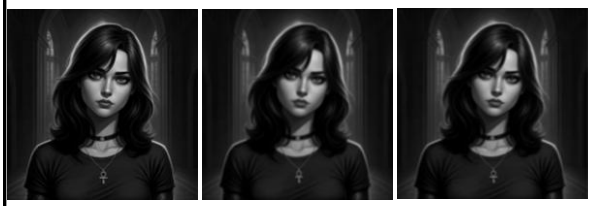

Box 13 x 13 Gauss 13 x 13, sigma 5

What is the effect of size of the filter on the results?
What about parameters of Gaussian filter?

34

Filtering Results

Box 13 x 13 Gauss 13 x 13, sigma 5

Look at details!

35

Low pass filters to remove noise

Salt and Pepper Noise: 0,3 Gauss 7 x 7, sigma 2

How good is it?

36

Low pass filters to remove noise

Salt and Pepper Noise: 0,5

Gauss 7 x 7, sigma 2

How good is it?

37

Nonlinear noise reduction

For some kind of noise, e.g. impulse noise (or „salt and pepper") we can do much better!

Median filter – nonlinear filter, cannot be realized by convolution operations!
It computes a median of pixel values under the mask at current position

$$\begin{matrix} 255 & 255 & 255 \\ 255 & 18 & 255 \\ 255 & 255 & 255 \end{matrix}$$

What would be the result of filtering?

Using 3x3 Linear LP box filter: 288 (and surronding fixels are changed as well?
Using 3x3 Median filter: 255 (and no change for surrounding pixels)

38

Filtering Results

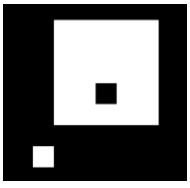

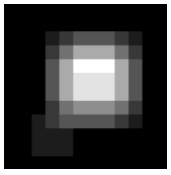

$\frac{1}{9}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$

Median Filter

39

Filtering Results

Salt and Pepper Noise: 0,5

Median Filter 7 x 7

How good is it?

40

High Pass Filters

We are back to linear filtering (convolution)

Let's use the following masks

$$S_x = \begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix} \qquad S_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$$

How would they work?

41

High Pass Filters Results

$$S_x = \begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix} \qquad S_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$$

42

High Pass Filters

They pass only high frequencies – edges, corners, noise etc.
They are **edge detectors!**
Called also **image derivative filters** or **image gradient filters**

$$\frac{\partial f(x,y)}{\partial x} = \lim_{\varepsilon \to 0} \frac{f(x+\varepsilon, y) - f(x,y)}{\varepsilon}$$

$$\frac{\partial f(x,y)}{\partial x} \approx \frac{f(x+1, y) - f(x,y)}{1}$$

Many masks used in practice – Sobel, Prewitt, Roberts for 1st image derivative

Laplacian for 2nd image derivative $\nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$

43

Edge Detection

An edge – a place in the image where there is a rapid change of pixel intensity
Sources of edges?
Are all edges in the image true edges from 3D world?

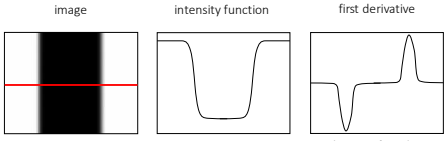

Edges are found as **maxima** of **image first derivatives**

44

What causes an edge?

45

Image Gradient

1) Choose directional first derivative image filters (Sobels masks)
2) Convolve the filter mask with an image to obtain image derivatives
3) Compute gradient at each pixel – its magnitude and direction

$\nabla f = \left[\frac{\partial f}{\partial x}, 0\right]$

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

$\nabla f = \left[0, \frac{\partial f}{\partial y}\right]$

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

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

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

46

Edge Detection

What are edges in real images?
Noise!

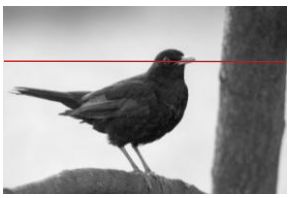

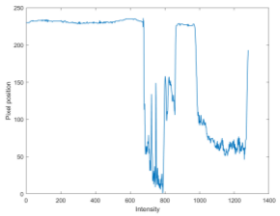

47

Edge Detection

Solution – first reduce noise by convolving an image with **LP filter** and then use **HP filter** to **detect edges**
If we use Laplace second derivative mask – well known LoG Filter

48

Laplacian and LoG Filtering Results

Laplacian Filter

Laplacian of Gaussian Filter

49

HP Filtering vs Edge Detection

LoG Filter

LoG Edge Detector

50

Canny Edge Detection

Goal: Get all real edges, remove all edges due to any kind of noise
Edge – it is 1D curve, should be one pixel wide

After HP: Multiple edges, „smeared" wide edges, because the gradient value is large across the edge

Edges may be broken into pieces

Solution: **Canny Edge detector** – one of the most known image processing algorithms

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

51

Canny Edge Detection

General Canny Edge detection algorithm

1) Compute image derivatives in x and y direction
2) Compute magnitude and orientation of the gradient
3) Thin „wide" edges to one-pixel width – **Non-maxima supression**
4) Use two thresholds and **hysteresis** for thresholding and edge linking

52

Non-maxima supression

Computer for pixel edges for each edge orientation

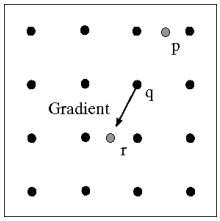

There is a maximum (Edge point) in pixel q, if the magnitude is greater than in points p and r, on both sides along gradient direction (values interpolated)

Source: D. Forsyth

53

Edge Linking

Start at edge pixels, link the edge following the edge's direction

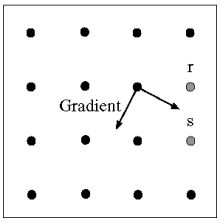

From an edge point, construct a tangent to the edge curve, then follow it to predict the position of the next edge pixel

Source: D. Forsyth

54

Closing gaps – two threshold hysteresis

Start at edge pixels, link the edge following the edge's direction

Use two threshold to find edges without gaps
Use higher threshold t_1 to start following an edge curve, then use the lower threshold t_2 to stop following it

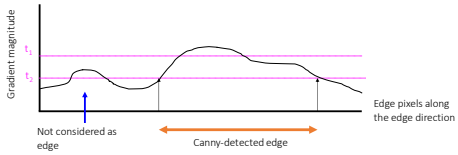

55

Canny Edge Detection

t1 = 0,25, t2 = 0,1 sigma =1

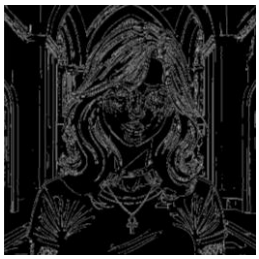

t1 = 0,05, t2 = 0,01 sigma =1

56

Canny Edge Detection

t1 = 0,25, t2 = 0,1 sigma =1

t1 = 0,25, t2 = 0,1 sigma =5

57

Morphological Image Filtering

They are usually used to proces binary images
Theoretical and detailed explanation next lecture
Today, let us say that they act as **minimum/maximum filters**
The mask can be of any shape (**structuring element**)

The ouput pixel value becomes lowest/highest pixel value of pixels directly under mask centered at that pixel

Basic **morphological operations**

Erosion
Dilation
Opening
Closing

58

Morphological Filtering

Erosion: disk element size = 9

Dilation: disk element size = 9

59

Morphological Filtering

Opening: disk element size = 9

Closing: disk element size = 9

60

Morphological Filtering

Compare details on the head of the bird

Erosion

Opening

Dilation

Closing

61

Image Restoration

62

Review material

- Describe a digital representation of grayscale and color images
- What is image histogram?
- What are characteristics of low/high contrast images' histograms?
- What is the effect of histogram equalization on image and its histogram?
- What do high and low spatial frequencies represent in images?
- How is linear filtering performed on 2D signals?
- How low/high pass filters change images?
- How does the median filtering work?
- How to compute image gradient?
- Describe linear filters which are used to detect edges on images
- What are the main steps of Canny edge detection algorithm?

63

Thank you for attention

64