Introduction to Medical Image Analysis

Rasmus R. Paulsen
DTU Compute

rapa@dtu.dk

http://www.compute.dtu.dk/courses/02511
http://www.compute.dtu.dk/courses/02512

Rasmus R. Paulsen


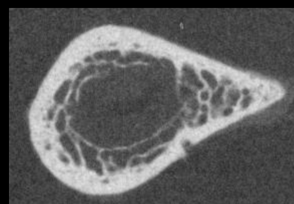
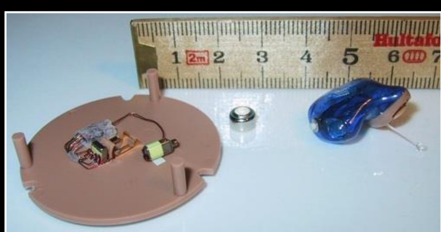


- Master of Science (Eng). DTU 1998
- Industrial PhD with Oticon A/S
- Research and development at Oticon A/S
- Associate Professor DTU Compute

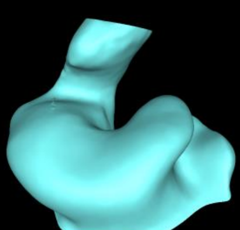

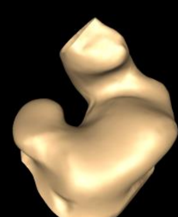
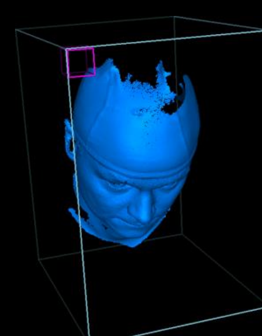
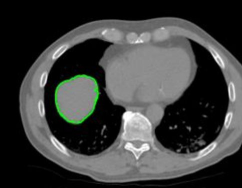
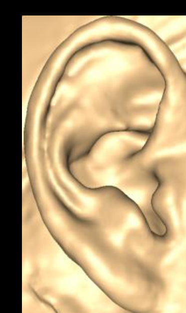

Teaching Assistants

- Gudmundur Einarsson
  Ph.D. student at DTU Compute
- Sigbjørn Hokland
  Medicine and Technology

A statement from me and the TA's

- The lectures and exercises are offers
  - We do not notice if you are here or not
- We want to lift you as far as possible
  - Make you understand our topic and the learning objectives
- We are grateful for critical and constructive feedback

- We do expect that you are here to learn
- We do expect that you are responsible for your own learning
- The TA WILL help the ones that have a hard time understanding a problem
- The TA WILL NOT help the ones that do not even try to understand the problem

Practical matters

- 13 days over the DTU 13 week semester
- Full day with lectures and exercises
- Lectures in auditorium 045 in building 303A
- Exercises in group rooms south 001 in building 308

Week 1 - today

9.00	Introduction and practical matters
	Lecture – Digital Images
	Exercises
12.00 – 13.00	Lunch break
13.00 - 17.00	Exercises

Materials

- Book:
  - Rasmus R. Paulsen and Thomas B. Moeslund: *Introduction to Medical Imaging.* **(MIA). 4th print. 2015**
  - **~250 kr**
  - Polyteknisk boghandel
  - http://people.compute.dtu.dk/rapa/MedIABook/
  - Errata for earlier versions here
- Notes
  - A few other hand outs
- At the end of the course a complete reading list will be published

CampusNet and homepage

- Get used to use CampusNet
  - Course messages will be given through CampusNet
  - Exercises delivered using CampusNet

- Homepage
  - http://www.compute.dtu.dk/courses/02511
  - Date / Exercises / Schedule
  - Updates happen!

#	Dato	Emne	Materiale	Øvelse
1	4/2	Introduction and digital images	MIA kap. 1 og 2, app. A	1
2	11/2	Image acquisition, and digital images, and file formats	MIA Kap. 2 og 3	2
3	18/2	Pixelwise operations	MIA kap. 4	3
4	25/2	Neighborhood Processing (Filtering)	MIA kap. 5	4
5	4/3	Morphology **Eksternt foredrag:** Projektleder Oline V. Olesen: Tracoline, fra ide til klinisk virkelighed.	MIA kap. 6	5
6	11/3	Blob analysis	MIA kap. 7	6
7	18/3	Pixel classification	MIA kap. 8 og 9	7
8	25/3	X-Ray imaging and CT scanning	Noter	8
9	8/4	Geometric Transformations **Eksternt foredrag:** Sales Manager Michel Strauss, Siemens Healthcare	MIA kap. 10	9
10	15/4	Image Registration	MIA kap. 11	10
11	22/4	Boundary Tracing (Hough Transformation and Dynamic Programming) **Eksternt foredrag:** CEO Michael Grunkin, Visiopharm	MIA kap. 12	11
12	29/4	Segmentation and Clustering	noter	12
13	6/5	3D medical image analysis	noter	Tidligere eksamensopgaver

- Nearly a chapter per lesson
- Few notes

Invited Speakers

- Relevant invited speakers from the industry/research sectors
  - CEO Michael Grunkin, Visiopharm
  - Sales Manager Michel Strauss, Siemens Healthcare
  - Projektleder Oline V. Olesen, Tracoline
  - DTU Compute PhD Students

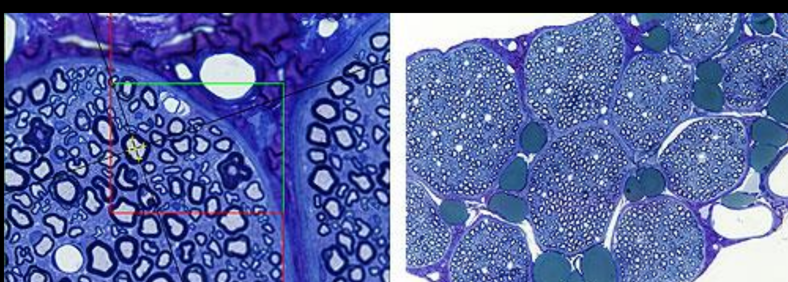

Learning Objectives (Læringsmål)

- A list of learning of objectives shown before each lecture
- A learning objective describes what you can do after the lecture
- If you fulfil all learning objectives you get 12
- Low-level learning objective
  - Apply the Prewitt edge filter to an image
- High-Level learning objective
  - Evaluate and compare the performance of a selection of image analysis algorithms

Exam

- 4 hour written exam
- Multiple choice
  - 25 questions
  - For each question there are 5 answers and a "I do not know"
  - Correct answer 5
  - Wrong answer -1
  - No answer 0
- Throughout the course example exam questions will be given
- Previous exams can be found on the homepage at the end of the course

Exercises

- 12 exercises in total
- You should deliver three exercise reports
  - Only for course 02511
- Count approximately 30% of the final grade
- Exercise 4 should be delivered at latest March 11!
- Can be delivered by 1 or 2 persons
  - You are responsible for forming the groups
  - Not allowed to deliver "duplicated reports"
- Other deadlines at the homepage

Borrowing stuff or cheating

- Cheating is generally punished hard at the university
- Reports are checked automatically when delivered
  - Offenders are reported to the study service
- It is ok to:
  - Use small pieces of text from other works if it is clearly written (cited) where it comes from
  - Discuss the problems with other people
- It is NOT ok to:
  - Copying parts of other peoples reports
  - Deliver text that you have not written
  - Directly translate text from for example Wikipedia

Matlab and computers

- No databar
- We assume that you can use your own portable computer with Matlab
- Try to arrange yourself into groups with at least one working Matlab installation

Exercise report

- When doing the exercises:
  - Create a Matlab file and keep all your Matlab code for the exercise in this one
  - Create a text document (word, Latex, OpenOffice etc) and paste the results here
- Writing the report
  - Evaluate the results
    - What do you see?
    - Why does the results look like they do?
    - Can it be improved? How?

Bachelor and master thesis

- Many of you should soon choose a bachelor project
  - Some even a master project
- Difficult to choose topic and supervisor
- Later in this course we will give you some tips and ideas

Questions!

- How many of you:
  - use image manipulation software
    - Photoshop
    - Paintshop
    - GIMP
    - Something else (what?)
  - Adjust images before putting them on Facebook or somewhere else?

Image Manipulation

- What operations do you do
  - Crop?
  - Resize?
  - Contrast?
  - Rotate?
  - Colors?

Camera

- Bring your own camera/smartphone to the exercises
- Learn to transfer photos from your camera/phone so you can use them on your computer

What is image analysis?

- Extraction of information from digital images
- Reproduce expert diagnostics
  - More accurate
  - Variation between doctors opinions removed
- Computer aided diagnostics – the doctor has the last word
- Can enhance the signs of diseases
  - Tumours
  - Bleedings

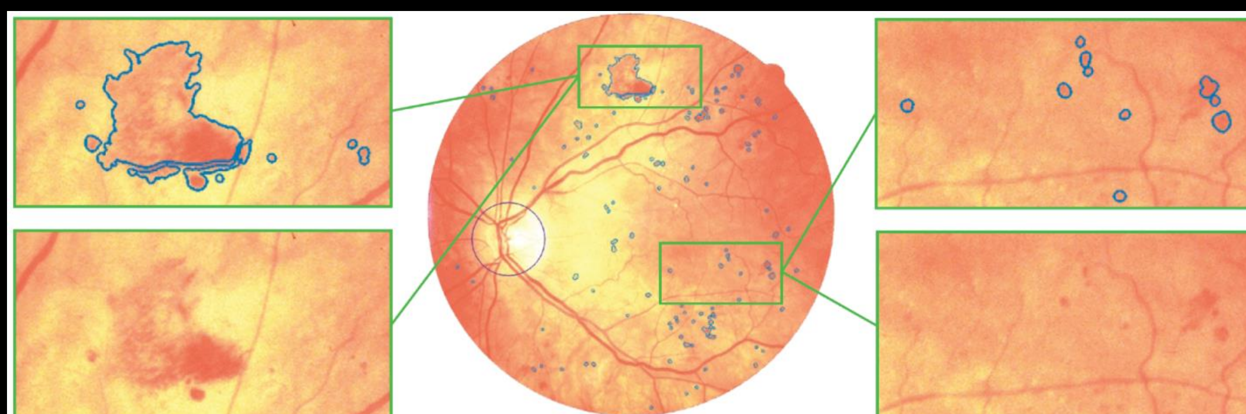

Automatically detected haemorrhages and micro aneurysms in digitized fundus images

Examples

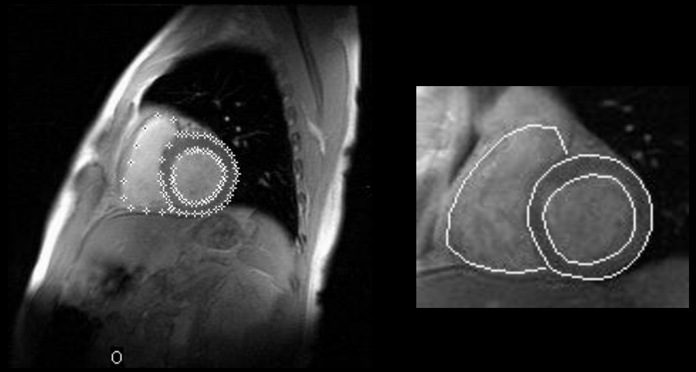

Recognise and track the heart

Face recognition and tracking

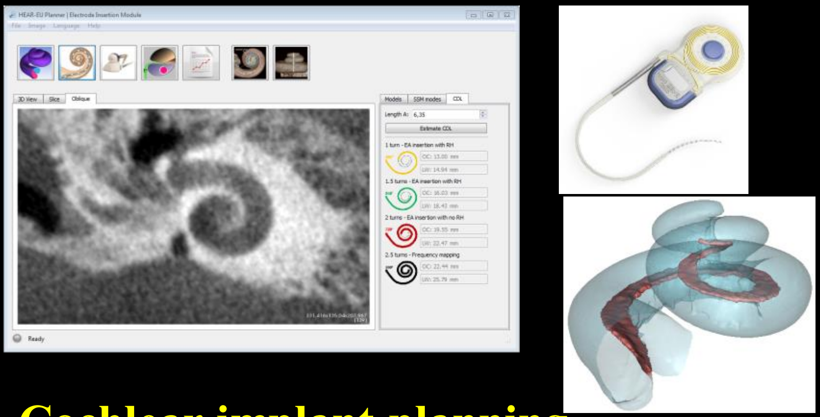

Cochlear implant planning

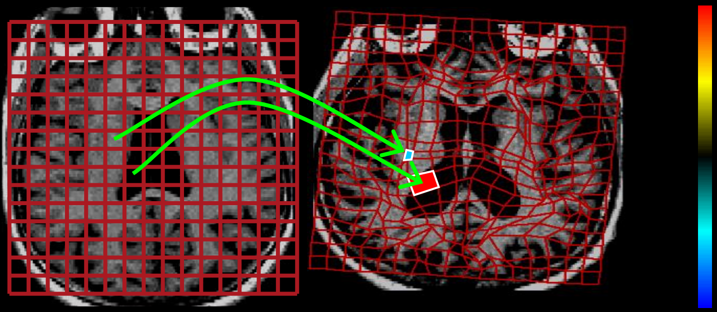

Shape changes in brain structures

Relevance

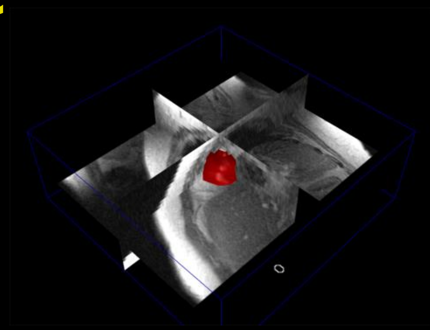

1980

2012

Magnetic resonance (MR) machine

- A lot of the data on hospitals are images
- Images are an important tool in
  - Diagnosis
  - Treatment
  - Follow-up
- Very high-tech!
- New imaging technologies are developed all the time.
- Connection between patient journals and images

Relevance

- Siemens PET/MR machine
- Installed at Rigshospitalet december 2011
- Extremely advanced
- New types of images and information

Relevance

- 7 tesla MR scanner installed at Hvidovre hospital
  - New anatomical details visible
  - More information in brain images

- Fast and accurate 3D face scanner installed at the Bloodbank at Glostrup Hospital
  - Creating a database of normal human faces
  - Can be used to identify facial features connected to non-normal growth

Digital Images – Learning Objectives

- Describe the fundamental properties of a digital image
- Read and show an image in Matlab
- Describe the commonly used image coordinate systems
- Describe the binary, the label, the multispectral, and the 16-bit image
- Show and manipulate your own images in Matlab

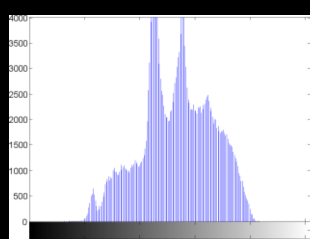

A digital image

23	216	120	55
4	89	158	130
65	76	189	34
19	234	7	45

- Consists of pixels (picture elements)
- Each pixel has a value between 0 and 255? Why?

Bits and Bytes!

- A bit is a tiny tiny little switch that can be either 0 or 1 – the “memory of a computer” consists of insanely many bits
- One byte is 8 bits together. It is the “basic” unit in a computer.
- With 8 bits how many possible values can be made?
  - (2^8 = 256)

- 00000001 = 1
- 00000010 = 2
- 00000100 = 4
- 00001010 = 10
- 00001111 = 15

128	64	32	16	8	4	2	1
□	□	□	□	□	□	□	□

Bit the Byte!

MSB							LSB

LSB = Least significant bit
MSB = Most significant bit

Binary numbers

- Decimal 6
  - 0000 0110
- Decimal 67
  - 0100 0011
- Decimal 126
  - 0111 1110
- Decimal 225
  - 1110 0001

A digital image

23	216	120	55
4	89	158	130
65	76	189	34
19	234	7	45

- between 0 and 255.
- How many bytes do our image take up in the computer memory?
  - 16
- How many different images can be made?
  - 256^16 = $(3.4\cdot10^{38})$

Grayscale digital images

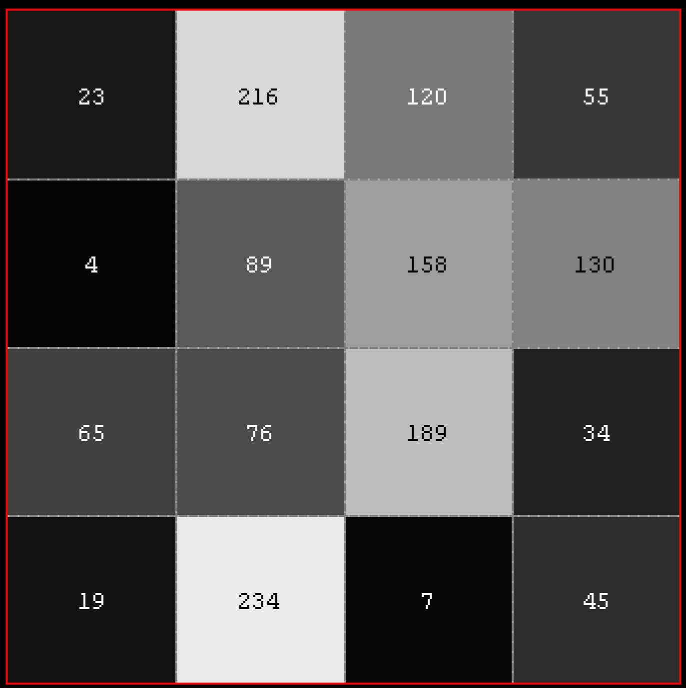

- 0 is black and 255 is white!
- The values in between are shown as shades of gray

Typical Grayscale image

- Traditional film X-ray
- Scanned on a flatbed scanner
- Do you know what an X-ray is?
- Bone is white and air is black
  - The more radiation the darker
- What are they used for?
  - Fractures
  - Arthrisis
  - Osteoporosis

Image Resolution

- Determines how much the image fills in the memory and on the hard disk
- Spatial resolution
- Gray level resolution

Spatial?

- Spatial
  - relating to the position, area and size of things
- Example:
  - This task is designed to test the child's *spatial* awareness
- Danish
  - Rumlig – barnet har en god rumlig forståelse

Spatial resolution

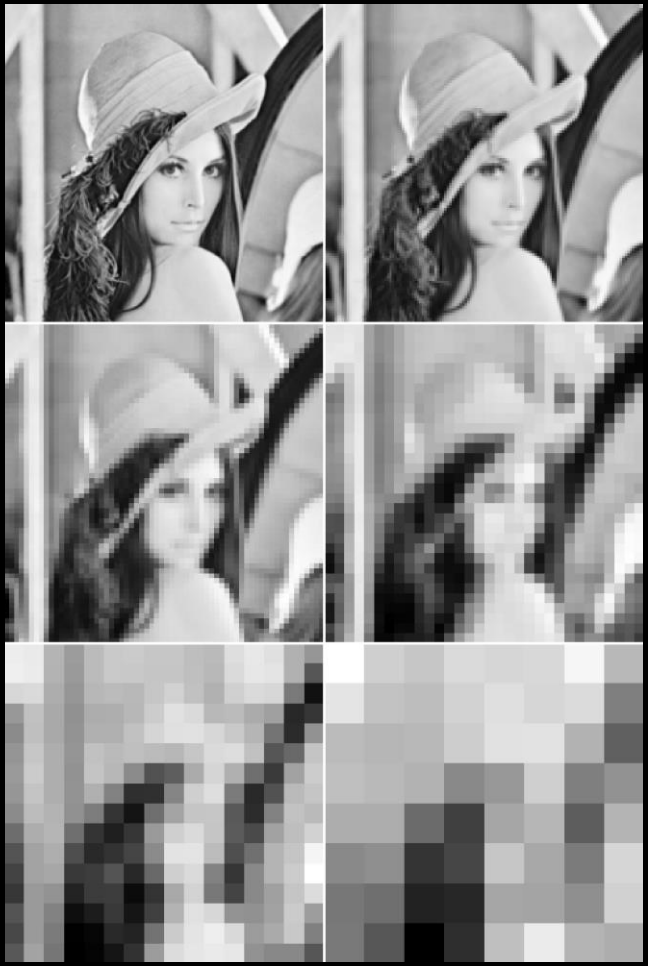

- The number of pixels used to represent the image
- 256 x 256
- 128 x 128
- 64 x 64
- 32 x 32
- 16 x 16
- 8 x 8

- How many pixels do you have in your camera?

How many pixels?

Width	Height	Pixels	Mega-pixels	Camera
320	240	10.000	0.01	Prototype 1975
640	480	307.200	0.03	Apple Quicktake 100
1600	1200	1.920.000	2	Nikon Coolpix 950
3520	2344	8.250.880	8.25	Canon 20D
8984	6732	60.480.288	60.5	Phase One P65+

Gray level resolution

- The number of gray levels in the image
- 256
- 64
- 16
- 8
- 4
- 2

Image as a matrix

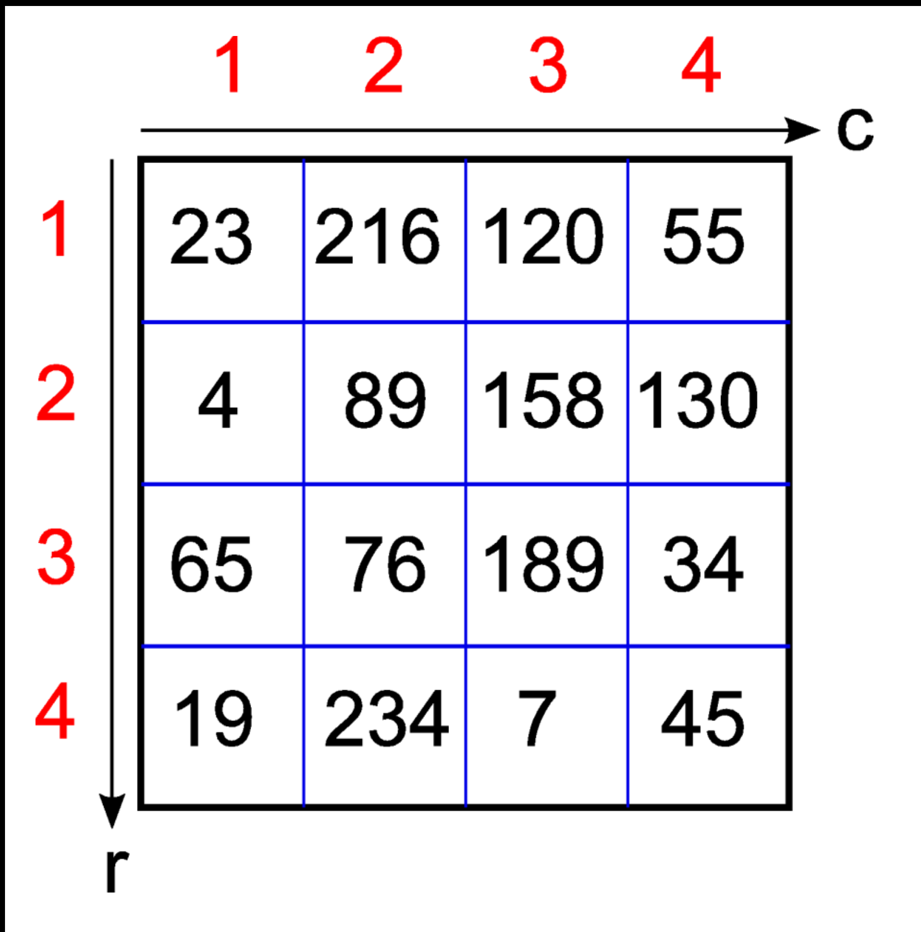

- An image is stored in the computer memory as a 2 dimensional matrix
- 4 rows and 4 columns
- Matlab image I – what is I(2,3) ?
- Can also be seen as a discrete function f(r, c)
- In Matlab a pixel is stored as an UINT8!
- UINT8 = Unsigned 8-bit integer = 1 byte

Pixel coordinates – Matlab matrix

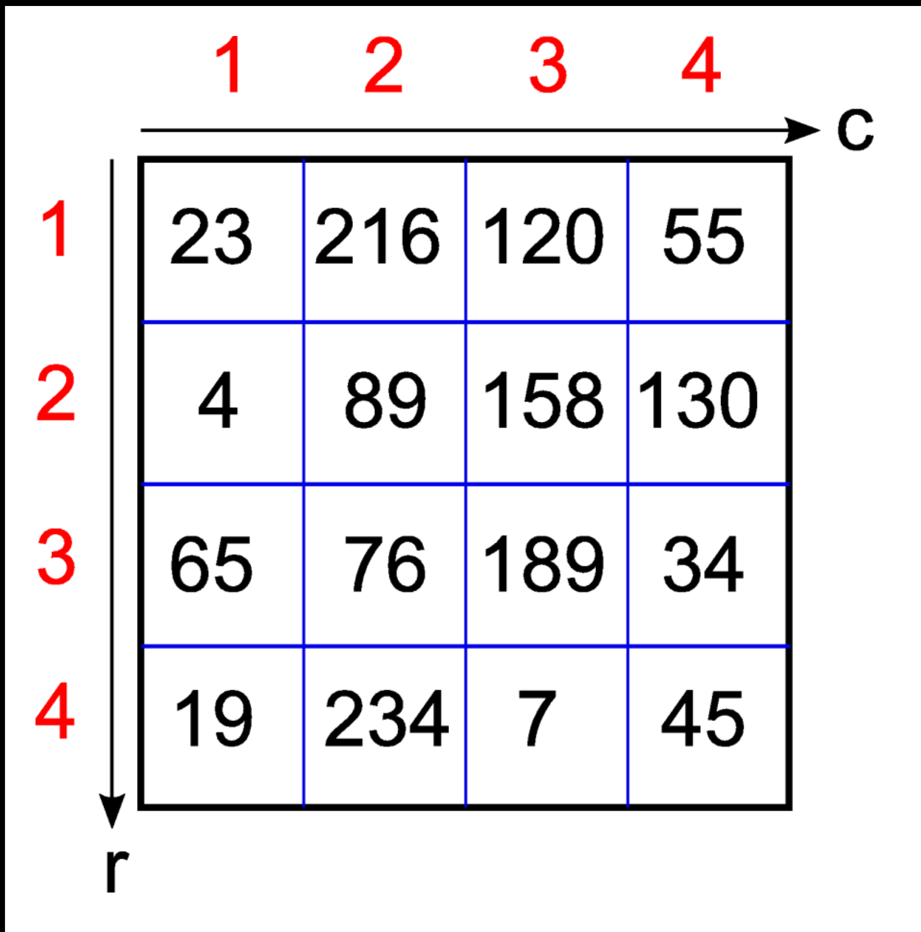

- Used in Matlab
- Origin is in upper left corner
- 1-based
- (row, column) system
- M rows and N columns

- What is the coordinates of the pixel with value 34?

Pixel coordinates – Photoshop etc.

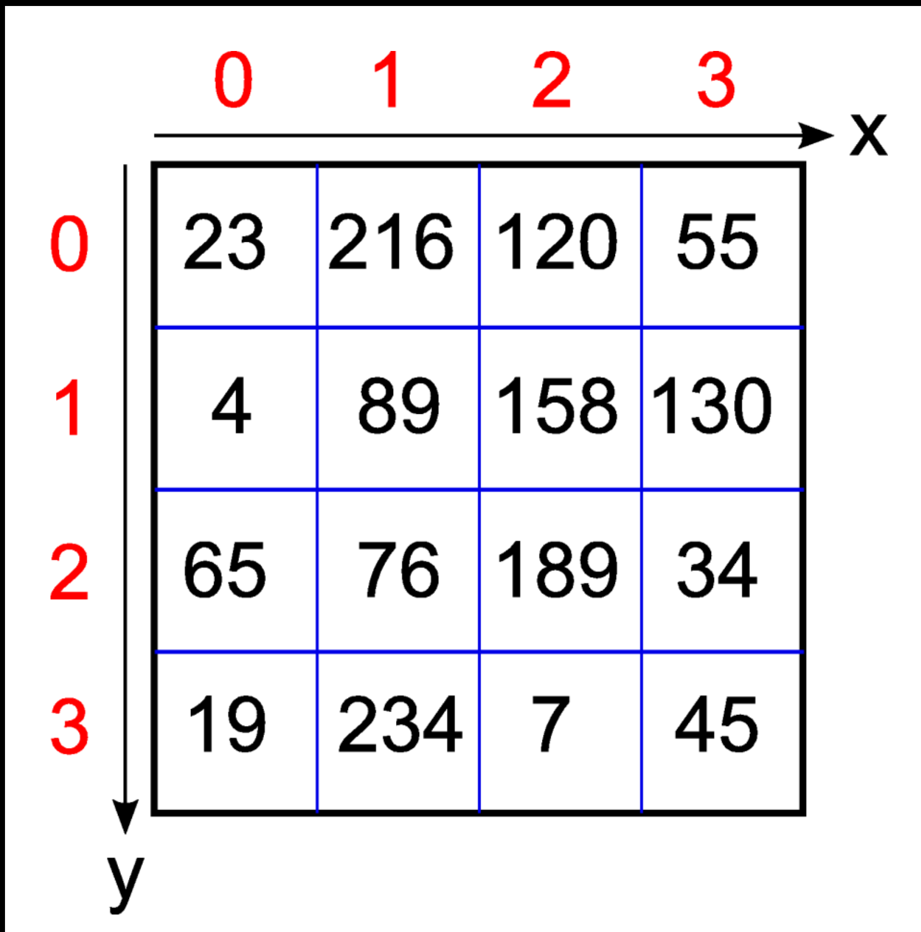

- Used in many graphics programs
- Origin in upper left corner
- 0-based
- (X,Y) system

- What is the coordinates of the pixel with value 34?

Pixel coordinates – Matlab plots

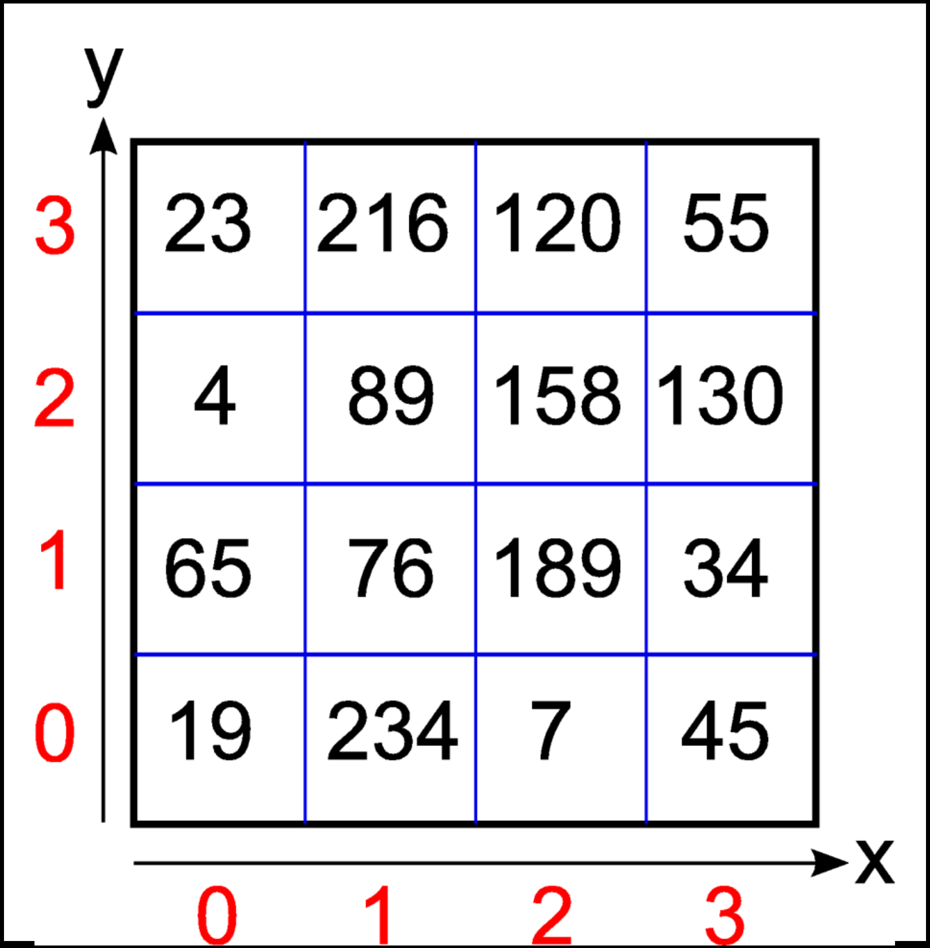

- Used when plotting – known from mathematics
- Origin in lower left corner
- 0-based
- (X,Y) system

- What is the coordinates of the pixel with value 34?

Why should I care?

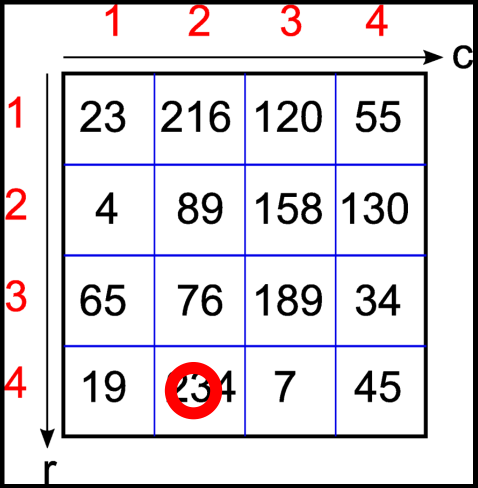

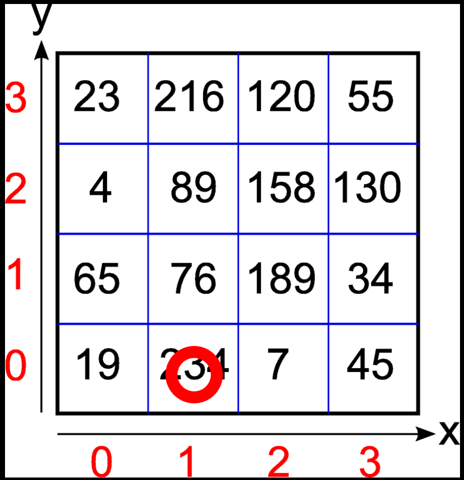

- You have a Matlab image in the matrix system
- Found the pixel with the maximum value
- Want to plot a red circle on top of it
- Plotting is done in the Matlab plot system
- How is this done in this image?
  - Max = 234 at (r,c) = (4,2)
  - Plot circle at (x,y) = (1,0)
- General conversion
  - x = c-1
  - y = M-r

The Image Histogram

- A histogram normally contains the same number of “bins” as the possible pixel values
- A bin stores the number of pixel with that value

0	2	6	6	3	3
1	4	3	4	4	4
3	2	5	1	5	2
1	4	2	1	3	1
2	5	3	0	2	0
4	2	5	6	3	1

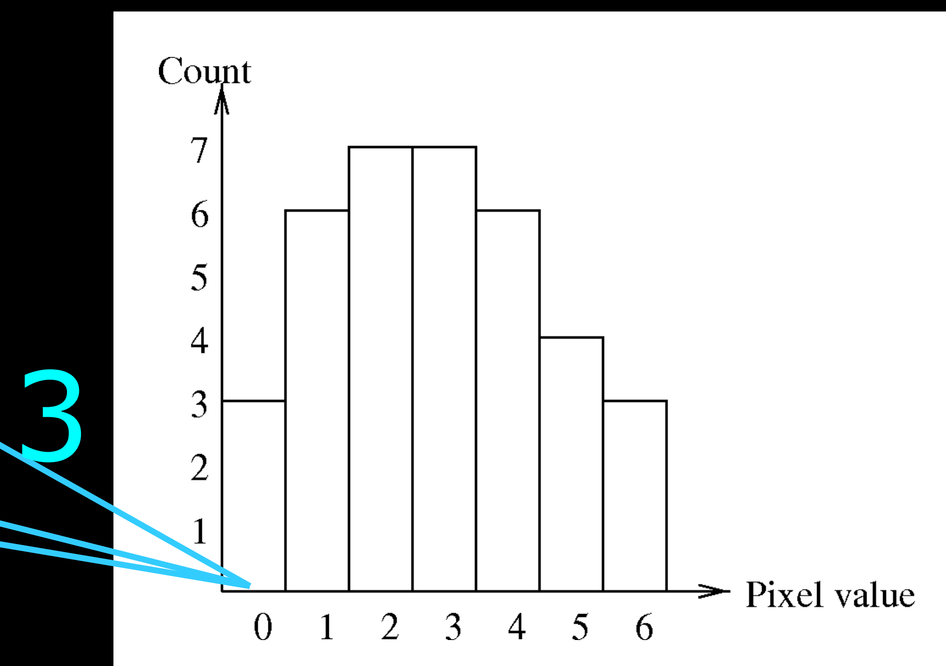

A real grayscale image histogram

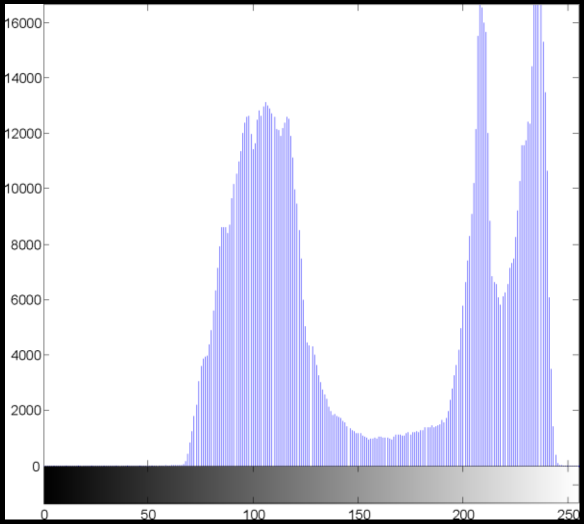

- 256 gray levels in the image = 256 bins in the histogram
- The shape of the histogram tells us something about the image
- Can you “recognise” the flower in the histogram?
- What “colors” are missing?

The histogram function

0	2	6	6	3	3
1	4	3	4	4	4
3	2	5	1	5	2
1	4	2	1	3	1
2	5	3	0	2	0
4	2	5	6	3	1

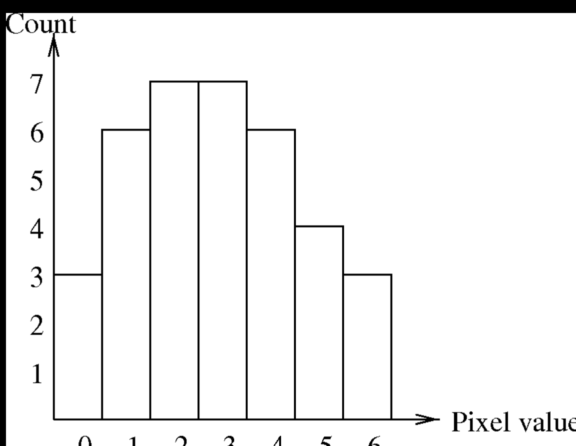

- Can be seen as a function h(v)
- v is the pixel value
- h(2) = 7
- h(5) = 4

- Total number of pixels is the sum of all h

Pixel value statistics

0	2	6	6	3	3
1	4	3	4	4	4
3	2	5	1	5	2
1	4	2	1	3	1
2	5	3	0	2	0
4	2	5	6	3	1

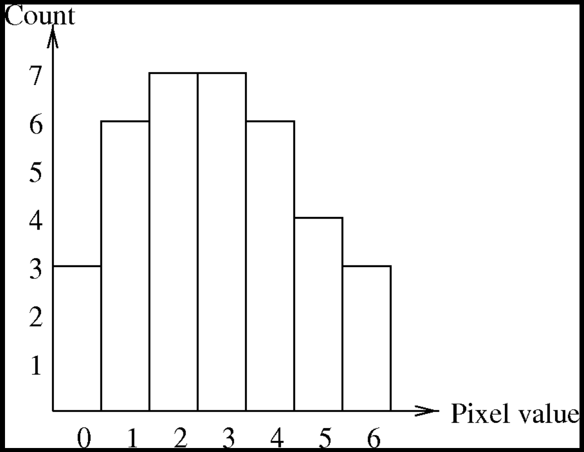

- Pick a random pixel in the image
- What is the probability of it having value 3? P(v=3)

- h(3) = 7
- Np = 36
- P(v=3) = 7/36 * 100%

Opgave t09.5

Der vælges en tilfældig pixel i billedet vist i Figur 4. Hvad er sandsynligheden for at pixelen har værdi 3?

1. 6%
2. 28%
3. 39%
4. 51%
5. 72%
6. Ved ikke

2	5	4	0	6	3
3	3	1	2	3	5
0	0	1	3	2	3
2	3	2	5	5	3
0	0	3	2	5	2
3	2	4	5	1	1

Figure 4: Et billede

Normalised histogram

- A normalised histogram is made by dividing each bin count with the total number of pixels
- H(v) is the normalised histogram function
- H(v) is the probability that a random pixel has value v

- Equal to a probability density function

Other Image Types

- Colour images
- Binary Images
- Label Images
- 16-bit images

Short intro

Will be covered in detail later in the course.

Colour images

- Anyone heard of RGB?
- RGB = Red, Green, and Blue
- Television, computers, digital cameras use the “RGB color space”
- Additive colours: Final colour is made by mixing red, green, and blue
- Typically the values of R, G, and B lie between 0 and 255 (total 3 bytes)!

RGB Colours

- When alle three “Lamps” are turned of we get black
- When all three “lamps” are on what do we get?

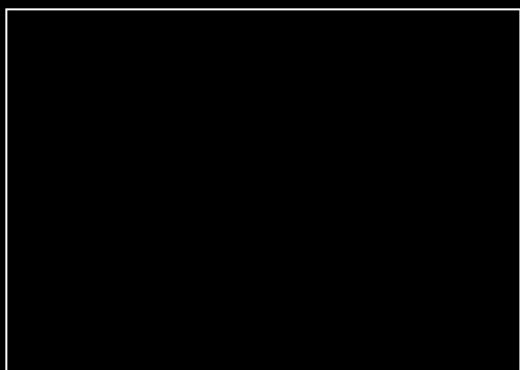

RGB = (0,0,0)

RGB = (255,255,255)

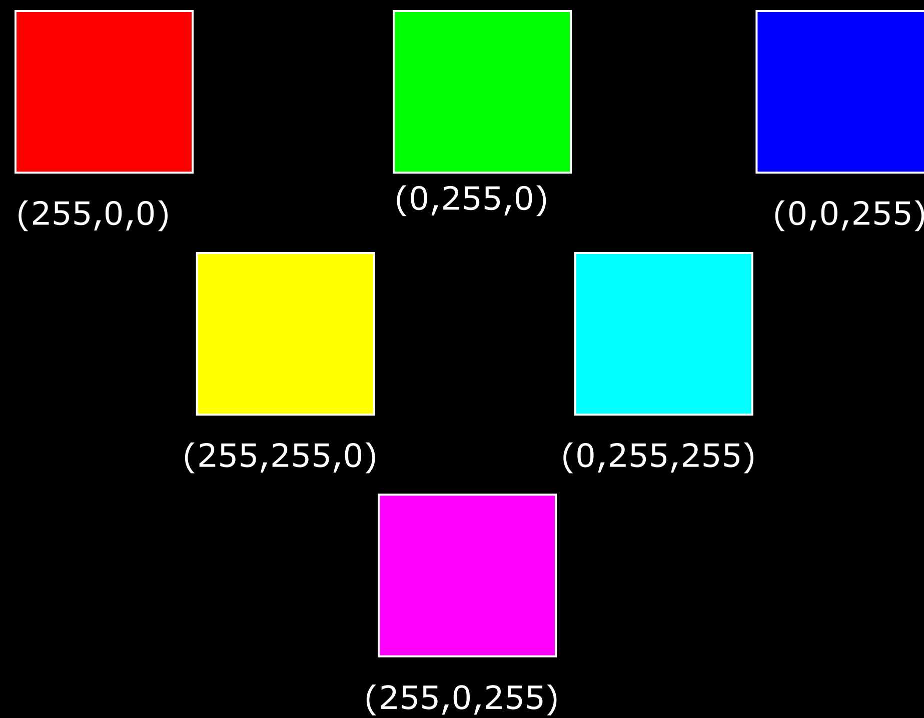

Processing RGB images

- Each pixel in a colour image contains 3 values
- Equal to a “vector function” in mathematics
- Much more complicated to analyse
- Medical images are typically grayscale
- Therefore we convert from colours to grayscale before the analysis

Converting colour to grayscale

$$v = 0.2989 * R + 0.5870 * G + 0.1140 * B$$

Is it possible to convert a grayscale image back to a color image?

Opgave t09.24

Et farvebillede konverteres til et gråtone billede. En farvepixel med RGB værdien (240, 120, 200) vil have have hvilken gråtone værdi efter konvertering?

1. 7
2. 57
3. 123
4. 165
5. 243
6. Ved ikke

Binary images

- Binary – means on or off
- Binary image – only two colors
- Background (0 = black)
- Foreground (1 = white)

- Simple representation of CT scanning of the head

Gray scale to Binary Image

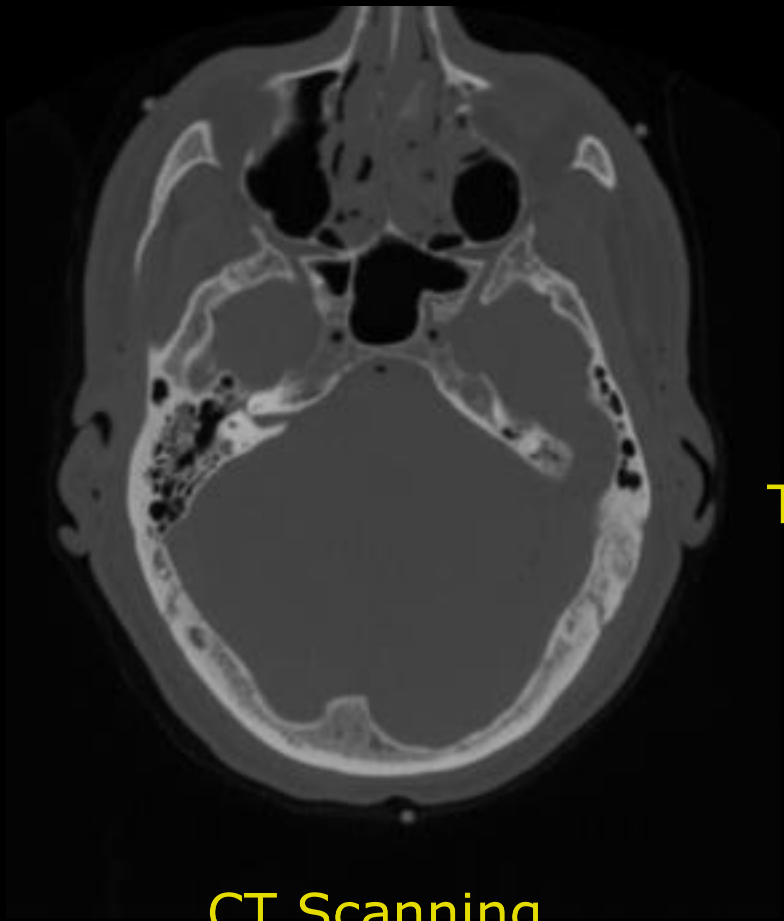

Threshold

CT Scanning

"Bone Image"

Binary image – why?

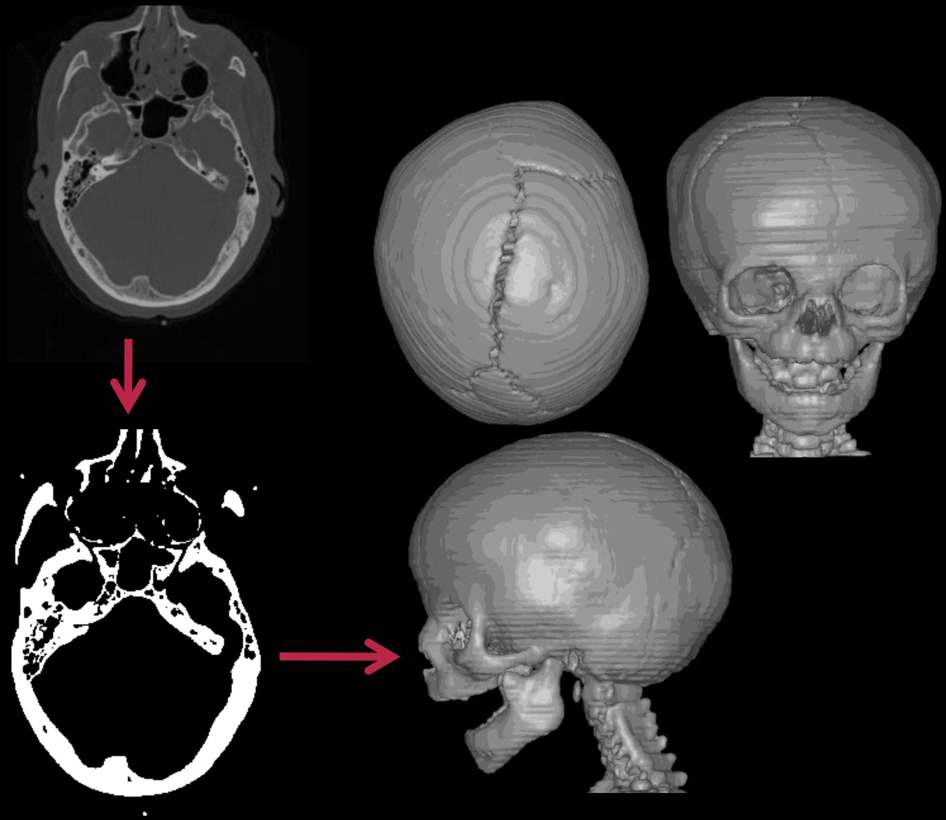

- Separating objects from background
- Count the number of the objects
- Measure the size and shape of objects
- Advanced 3D visualisations

Image from 3D laboratory

Label Image

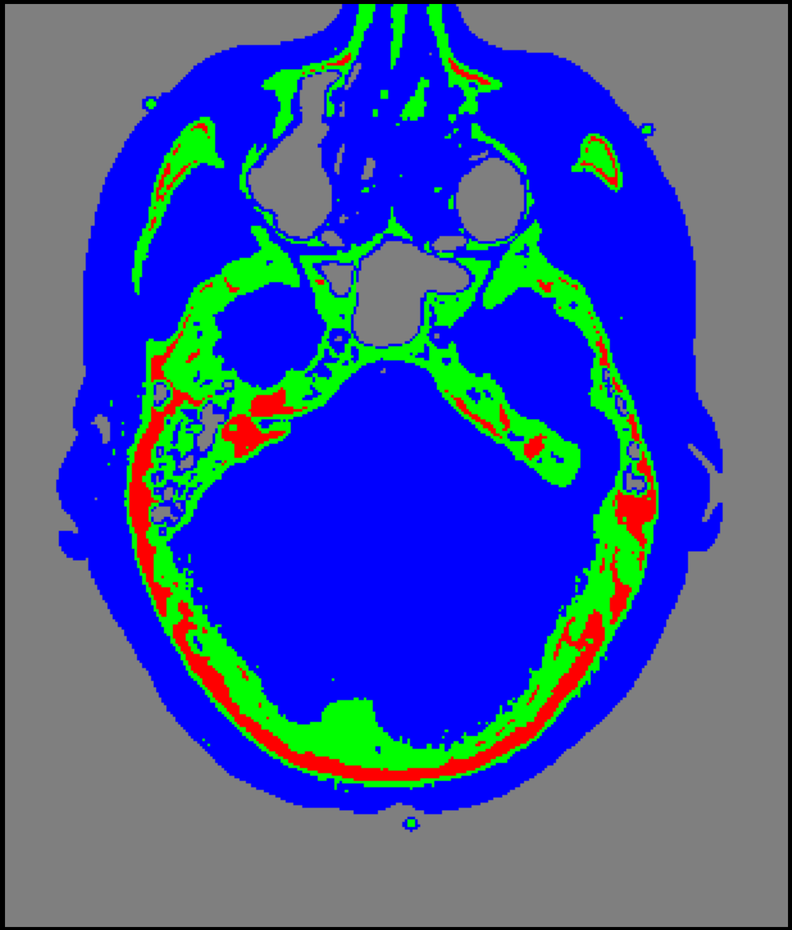

- The pixel value tells the *type* of the pixel
  - (0) Gray – background
  - (1) Blue – soft tissue
  - (2) Green – hard bone
  - (3) Red – spongy bone
- Only 4 different pixel values
- Colours made using a *look-up-table*

Label Image –why?

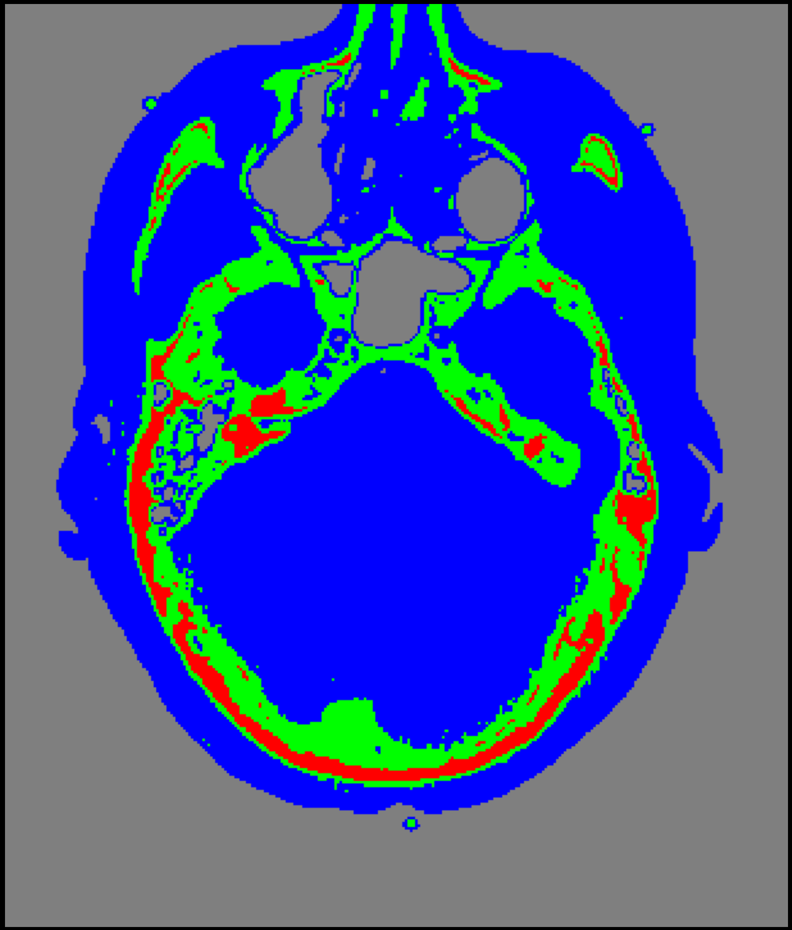

- How big is a tumour? (volume / percent)
- Bone density
- General anatomy recognition
  - Blood vessels
  - Calcifications

Label Image – how?

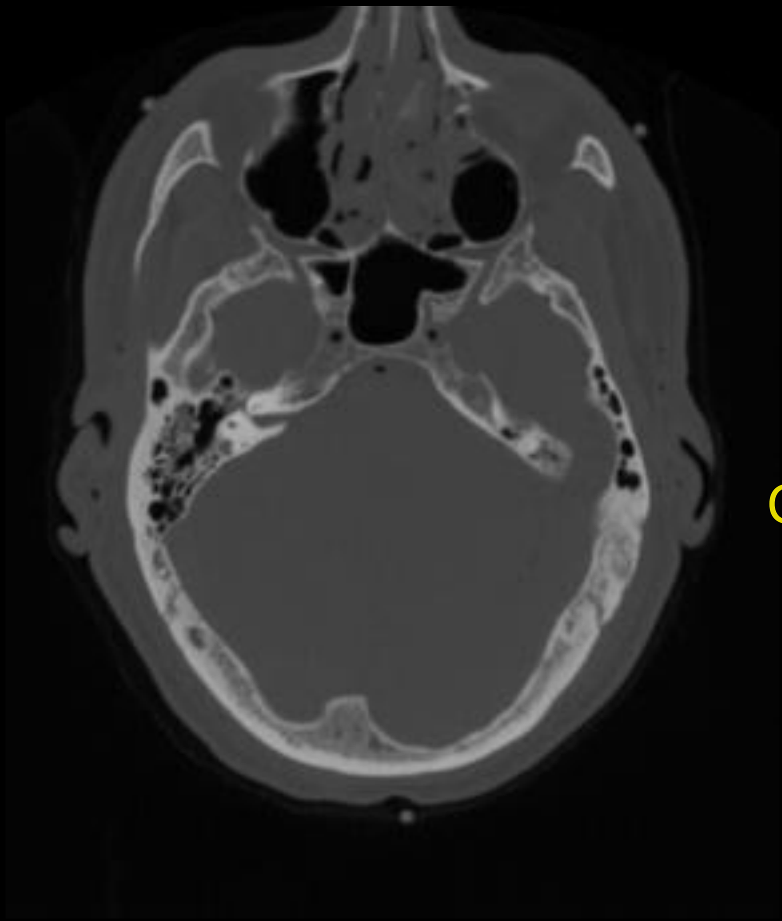

Classification

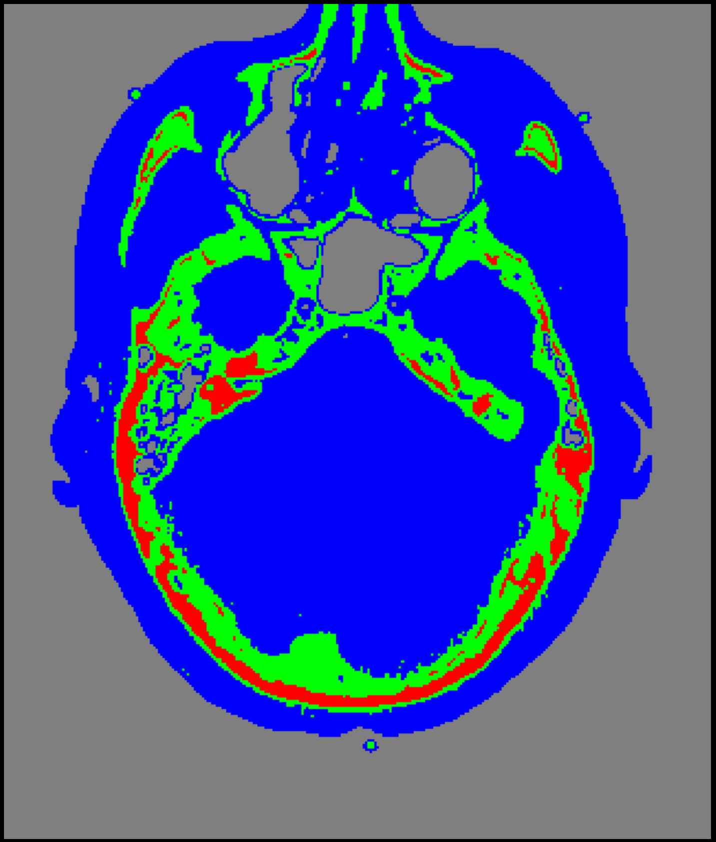

Multispectral images

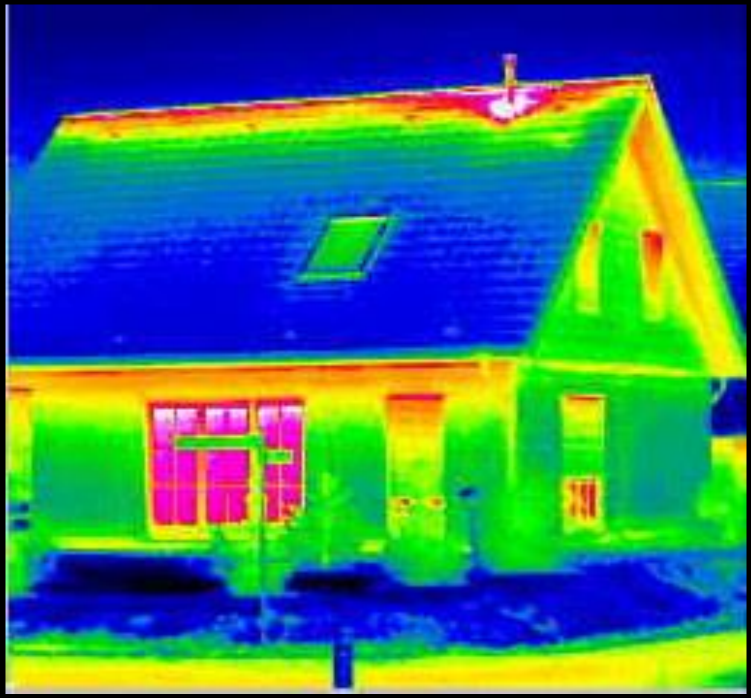

Infrared

- There are more visual information than what can be seen with the human eye
- Standard cameras captures the red, green, blue colours
- Capture systems that capture more bands and other frequencies exist
- Creates multispectral images
  - Each pixel contains perhaps 20 values from different spectral bands

Multispectral System - VideometerLab

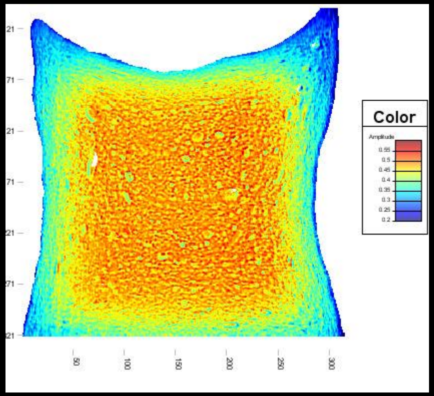

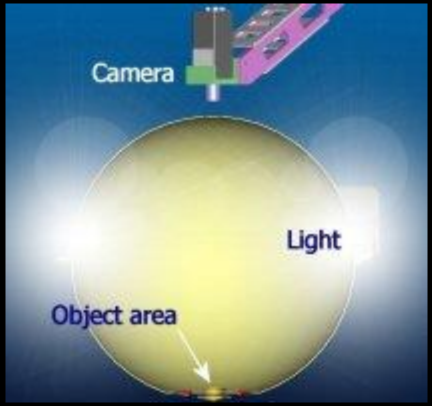

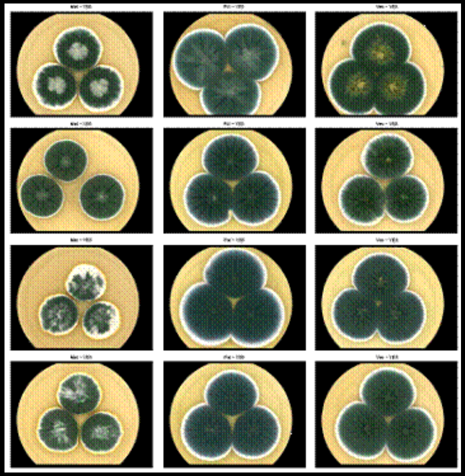

- Integrating sphere
- Light emitting diodes with different wavelengths
  - From near infrared to ultraviolet
- High resolution B/W camera
- Water in bread
- Classification of fungi
- Skin diseases

16-bit images

- 256 values fine for the human eye
- Pixel values not only for display
  - Physical meaning
- Computed Tomography
  - X-ray attenuation
- Hounsfield unit
  - 0 water
  - -1000 air
  - -120 fat
  - 400+ bone

Exercises