



Image Analysis

Rasmus R. Paulsen

Tim B. Dyrby

DTU Compute

<http://compute.dtu.dk/courses/02502>



Week 1 - today

8:00 – 10:00

Exercises







10:00 – 12:00

Introduction and practical matters

Lecture – An introduction to image analysis

Lecture – A tutorial on Principal Component Analysis (PCA)

Rasmus R. Paulsen



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

Tim B. Dyrby



- Associate Professor at DTU Compute and Danish Research Centre for Magnetic Resonance (DRCMR)



Teaching Assistants

- Mathias Micheelsen Lowes
- Bjørn Marius Schreblowski
- Thina Lundsgaard Thøgersen
- Andreas With Aspe



Practical matters

- 13 days over the DTU 13 week semester
- Flipped class room
 - 8-10 Computer exercises (also on MS Teams)
 - 10-12 Lecture with quizzes
- Lectures are streamed, recorded and made available
 - Links to the stream will be posted on DTU Learn before the lecture
 - Links to video on the homepage (under schedule)
 - [Courses.compute.dtu.dk/02502](https://courses.compute.dtu.dk/02502)



From 2021 on onwards

- The exercises are now much more related to the exam
- Learning objectives stated in all exercises
- You will be examined in these learning objectives
- You will also be examined in the more theoretical learning objectives from the lectures
- We expect that you can run Python during the exam!

Very Important I: Do the exercises!

Very Important II:
We can not help you after the course period!



Materials

■ Book:

- Rasmus R. Paulsen and Thomas B. Moeslund: *Introduction to Medical Image Analysis (MIA)*
- Polyteknisk boghandel
- <http://mediabook.compute.dtu.dk>

■ Notes

- Notes will be provided during the course

■ At the end of the course a complete reading list will be published



DTU Learn and the homepage

- Homepage : The main entry to the course
 - <http://courses.compute.dtu.dk/02502>
 - Schedule / Exercises / Data
 - Updates happen!
- Course messages will be given through DTU Learn



#	Date	Topic	Video	Material	Exercise
1	4/2	Introduction to image analysis (Rasmus) Introduction to Principal Component Analysis (PCA) (Rasmus)	Stream	MIA 1, 2, app. A. PCA Note (except Section VI (SVD) and App. A)	1
2	11/2	Cameras, lenses, image compression, image storage and change detection in videos (Rasmus)	Stream	MIA 2, 3 + CDV Note	1 + 1b
3	18/2	Pixelwise operations, Colour images. PCA Analysis on images (Claes)	Stream	MIA 4, 8 Eigenfaces article (only sections marked with yellow)	2 + 2b
4	25/2	Neighborhood Processing (Filtering) and Morphology (Tim)	Stream Recordings Fall 2023	MIA 5, 6	3 + 8
5	4/3	Blob analysis and object classification (Rasmus)	Stream	MIA 7	4 + 4b
6	11/3	Pixel classification and advanced classification (Tim)	Stream Recordings Fall 2023	MIA 9 + LDA note on Learn	5
7	18/3	Industry presentations:		none	Exercise catch-up
8	25/3	Geometric transformations and landmark based registration (Tim)	Stream Recordings Fall 2023	MIA 10, 11	6
9	1/4	Boundary Tracing (Hough Transformation and Dynamic Programming) (Tim)	Stream Recordings Fall 2023	MIA 12	6b
10	8/4	Advanced registration (Tim)	Stream Recordings Fall 2023	Elastix manual (5.2.0) chapter 2.	7
11	22/4	Real time face detection using the Viola Jones method (Rasmus)	Stream	Rapid Object Detection using a Boosted Cascade of Simple Features	9
12	29/4	Statistical models of shape and appearance and active shape models (Rasmus)	Stream	Statistical Models of Appearance for Computer Vision (p. 12 - 20 and p. 29 - 43)	Exercise catchup
13	6/5	Advanced topics (Claes Nøhr Ladefoged)	Stream	none	Digital test exam and exercise catch-up



Learning Objectives (Læringsmål)

- A list of learning objectives for each lecture and exercise
- A learning objective describes what you can do after the lecture/exercise
- 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

- Four hours multiple-choice exam
- Please see details here:
 - <http://courses.compute.dtu.dk/02502/exam.html>
- Previous exam sets are also available
 - Most relevant is from Spring 2021 and onwards



AI assisted tools during the course and the exam

- You are allowed to use AI tools like ChatGPT and Copilot
 - Both during the course and at the exam
- It is your responsibility to
 - install and keep your tools up to date
 - Verify if the output of the tools are correct
- The exam can ALSO be solved without the use of AI assisted tools



AI related learning objectives

- General 02502 course learning objective:
 - Estimate the correctness of the answer given by an AI assisted tool like ChatGPT and Copilot
- We are gradually adding AI tools related learning objectives to the exercises



PollEverywhere quizzes

<https://pollev.com/rasmuspaulse538>




What programming language are you most comfortable with?



Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app

What is your practical experience with Python programming



What is your experience with Jupyter Notebooks or JupyterLab

I am a complete beginner - I need instructions on software installation and a crash course in Jupyter notebooks.

0%

I have tried it a few times but I need a refresher

0%

I am comfortable with Jupyter notebooks but would not mind a quick refresher

0%

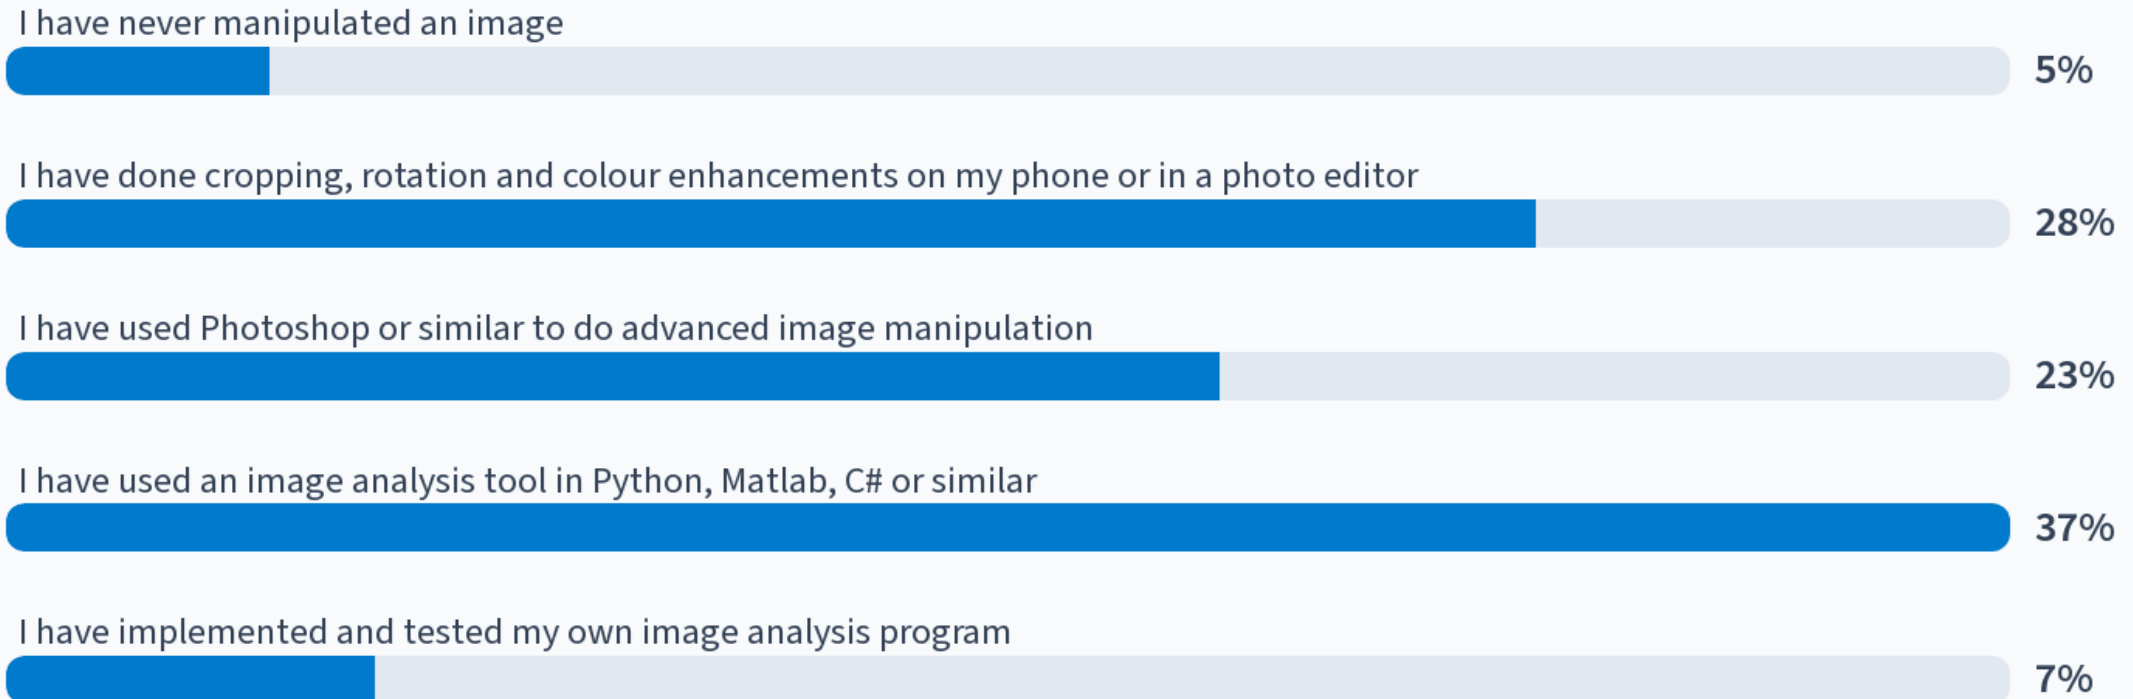
I am all ready - hit me!

0%

I am an expert Jupyter notebook user and have made my own

0%


What is your experience with image manipulation, image processing and image analysis?



To what extent are you currently using AI assisted tools like ChatGPT and Copilot

105

They are invaluable tool in my learning and university life



17%

I use them actively and often



45%

I am an occasional user




32%

I rarely use the tools



3%

I have never or very rarely used these tools



3%

Why are my slides black?



With a white background,
the strongest visual
contrast is here




No – not the (only)
reason

Norwegian Black Metal

?

Why are my slides black?



With a dark background, the strongest visual contrast is here
(which I find more important)



What is image analysis

- Automatic extraction of information from images
- A sub-topic within
 - Pattern recognition
 - Machine learning
 - Deep learning





What is image processing

- Changing the information in images – but not necessarily getting any knowledge
 - Photoshopping
 - Changing the visual appearance of photos
 - Cropping / rotating
 - Filters / effects



Face tracking – all features including eyes

- For digital cameras / phones
 - Automatic focus on the face + face beautification
- Tracking and manipulation for apps
 - Messenger / WhatsApp / SnapChat ...
- Awareness tracking for car drivers
 - Warning if you fall a sleep



A 100 million \$ industry



- This image is worth 100 of millions of dollars!
- Well – perhaps not that exact photo.
- The ability to track faces fast and accurate
 - Including estimates of 3D structure
 - App developers pays buckets of money for that
- It all started in 2001 with:
P. Viola and M. Jones. "Rapid object detection using a boosted cascade of simple features.". CVPR 2001
- Suddenly you could track faces fast and relatively accurate
- Now a lot of focus on deep learning


Self driving cars

■ Modern self driving cars rely on many sensors

- Lidar – radar system
- GPS
- Accelerometers, gyroscopes, magnetometers etc.
- Stereo cameras or multiple cameras
- Lots of advanced image analysis – sensor fusion



Sports tracking – human body tracking




- Huge commercial impact
- Lots of research in human body tracking
- Personal trainers
- TV player tracking and smart overlays



What is medical image analysis?




- Extraction of information from digital images
- Find unknown connections between diseases and what can be seen in images
- Can enhance the signs of diseases
 - Tumours / heart diseases / brain diseases / bone fractures
- Reproduce expert diagnostics
 - More accurate
 - Variation between doctors opinions removed
- Computer aided diagnostics



Automatic localization of the heart and its major substructures




Medical image analysis examples



Relevance



1980
Magnetic
resonance
prototype





Now – PET/MR

- Images is an important tool in
 - Diagnosis
 - Treatment
 - Follow-up
- Very high-tech!
- New imaging technologies are developed all the time.

Digital Images – Learning Objectives

- Describe the fundamental properties of a digital image
- Describe and use the commonly used image coordinate systems
- Describe pixel types
- Describe the binary, the color, the label, the multispectral, the floating point, and the 16-bit image



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?
 - $(2^8 = 256)$
- $00000001 = 1$
- $00000010 = 2$
- $00000100 = 4$
- $00001010 = 10$
- $00001111 = 15$
- $11111111 = 255$

128	64	32	16	8	4	2	1
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>



What is decimal 67 as a binary number?

1010 1010

0100 0011

0101 0001

0110 0101

0100 0101

I do not know

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app



What is decimal 67 as a binary number?





What is decimal 67 as a binary number?


1010 1010	0%
✓ 0100 0011	99%
0101 0001	0%
0110 0101	0%
0100 0101	1%
I do not know	0%

A digital image


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

- between 0 and 255.
- The pure image data takes up 16 bytes of computer memory

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
- Bone is white and air is black
 - The more radiation the darker
- What are they used for?
 - Fractures
 - Arthritis
 - 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

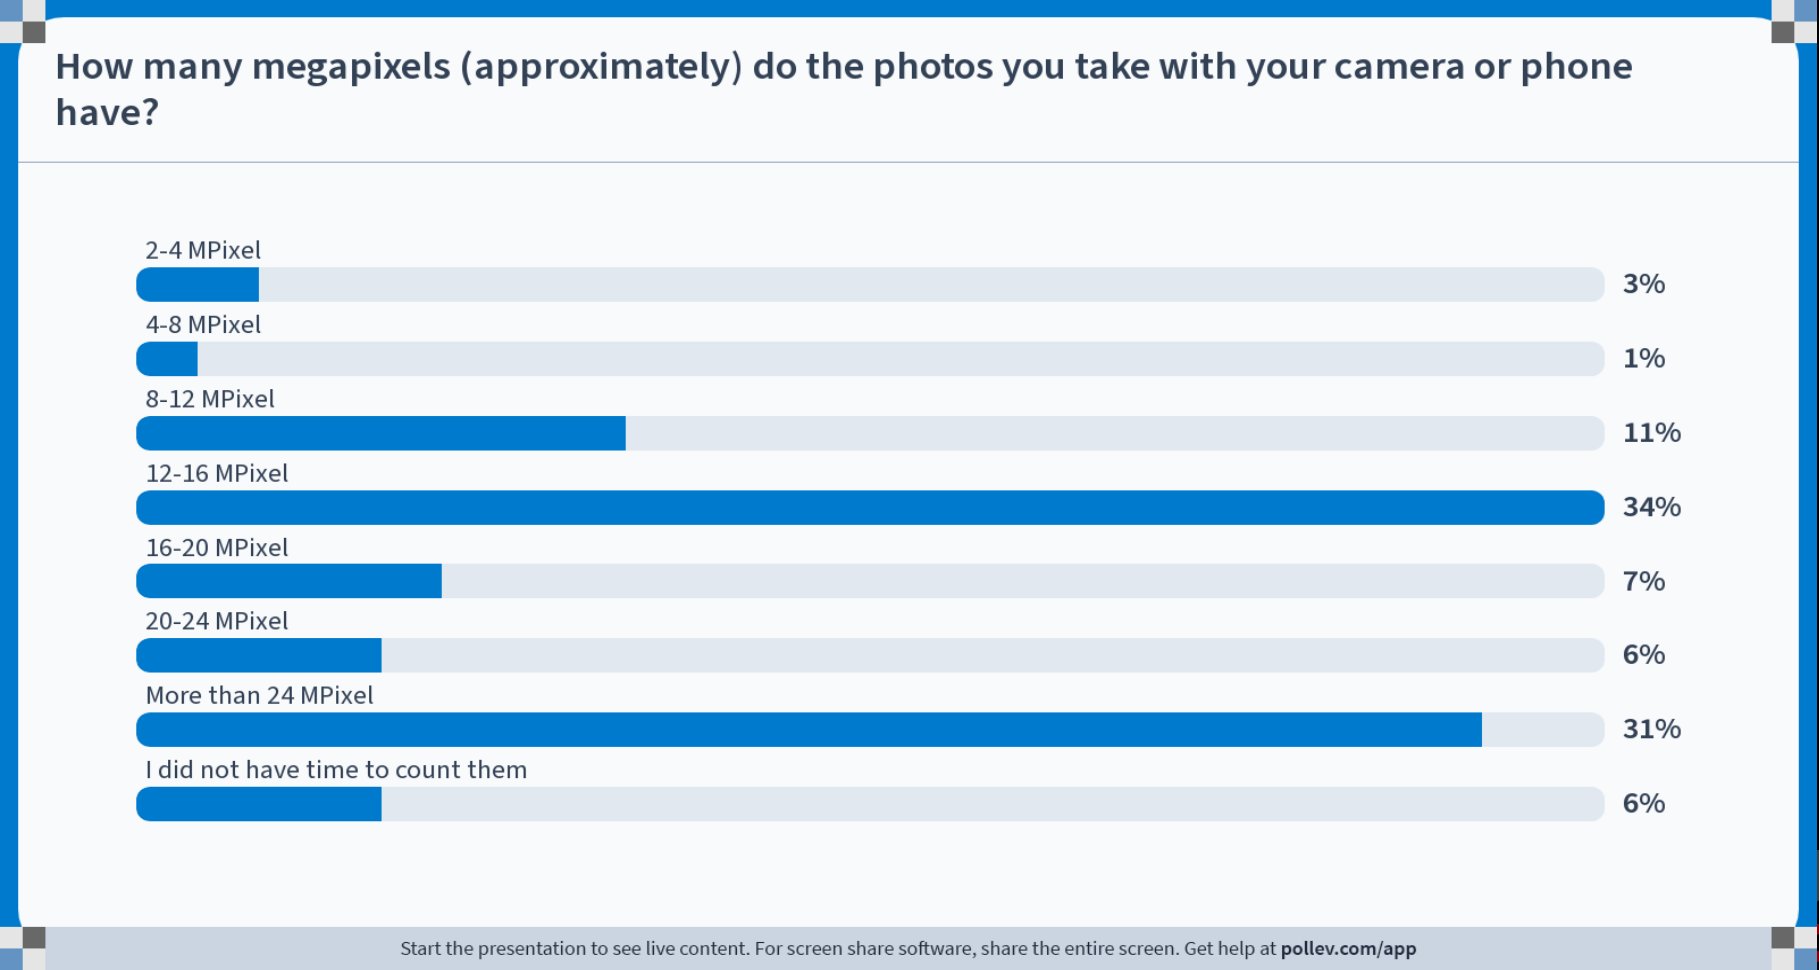
■ Danish

- Rumlig – jeg har en god rumlig forståelse

Spatial resolution

The number of pixels used to represent the image







How many pixels?

Width	Height	Pixels	Mega-pixels	Camera
320	240	10.000	0.01	Prototype 1975
1600	1200	1.920.000	2	Nikon Coolpix 950
4032	3024	12.192.768	12	Samsung Galaxy S7 edge
6240	4160	26.000.000	26	Canon EOS 6D M2
8984	6732	60.480.288	60.5	Phase One P65+




Grey level resolution

The number of grey levels in an image



An image as a matrix



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

- An image is stored in the computer memory as a 2 dimensional matrix
- 4 rows and 4 columns
- Can also be seen as a discrete function $f(r, c)$
- In Python a pixel can be stored as an `uint8`
- `uint8` = Unsigned 8-bit integer = 1 byte


Pixel types and their ranges

Data type	Range
uint8	0 to 255
uint16	0 to 65535
uint32	0 to $2^{32} - 1$
float	-1 to 1 or 0 to 1
int8	-128 to 127
int16	-32768 to 32767
int32	-2^{31} to $2^{31} - 1$

- A pixel can be processed and stored as different *types*
- The uint8 is the most common type
- For processing a pixel is often transformed to a float
- When processing speed and memory space is an issue you should be careful about the pixel type – more about that later in the course.

https://scikit-image.org/docs/stable/user_guide/data_types.html


Pixel coordinates – Python matrix



	0	1	2	3
0	23	216	120	55
1	4	89	158	130
2	65	76	189	34
3	19	234	7	45

- Origin is in upper left corner
- 0-based
- (row, column) system
 - Vertical axis is the first axis
- M rows and N columns
- Row range $[0, M-1]$
- Column range $[0, N-1]$


Pixel coordinates – Photoshop etc.



	0	1	2	3
0	23	216	120	55
1	4	89	158	130
2	65	76	189	34
3	19	234	7	45

- Used in many graphics programs
- Origin in upper left corner
- 0-based
- (X,Y) system
 - Horizontal axis is the first coordinate
- Often width (W) and height (H) are used to denote image dimensions
- X range $[0, W-1]$
- Y range $[0, H-1]$

Plot coordinates




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

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

3




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


Where are the flower leaves in the histogram?



- Range 1
- Range 2
- Range 3
- Range 4
- Range 5
- Range 6
- Oh bugger - my flower allergy has been trigge...




Where are the flower leaves in the histogram?



Range 1	<div></div>	1%
Range 2	<div></div>	0%
Range 3	<div></div>	1%
Range 4	<div></div>	5%
✓ Range 5	<div></div>	1%
✓ Range 6	<div></div>	89%
Oh bugger - my flower allergy has been triggered	<div></div>	1%




Where are the flower leaves in the histogram?



Range 1	1%
Range 2	0%
Range 3	1%
Range 4	5%
✓ Range 5	1%
✓ Range 6	89%
Oh bugger - my flower allergy has been triggered	1%

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? $P(v=3)$
- $h(3) = 7$
- $Np = 36$
- $P(v=3) = 7/36 * 100\%$
- The histogram divided by the total pixel count can be seen as a probability density function

A random pixel is chosen in the image. What is the probability that the value of the pixel is 3?

6%

28%

39%

51%



72%

I do not know

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app




A random pixel is chosen in the image. What is the probability that the value of the pixel is 3?





A random pixel is chosen in the image. What is the probability that the value of the pixel is 3?

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





Other Image Types


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

Colour images



- 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




RGB = (0,0,0)




RGB = (255,255,255)


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




(255,0,0)




(0,255,0)



(0,0,255)



(255,255,0)




(0,255,255)




(255,0,255)

Additive color mixing



<http://hyperphysics.phy-astr.gsu.edu/hbase/vision/addcol.html>

Subtractive color mixing



Wikipedia



Processing RGB images

- Each pixel in a colour image contains 3 values
- Equal to a “vector function” in mathematics
- More complicated to analyse
- Medical images are typically grayscale
 - Why?
- Often images are converted 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?

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



CT Scanning

→
Threshold



"Bone Image"

Binary image – why?





Image from 3D laboratory

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


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?

- Segment images into regions
- **Example:** Recognize the major structures of the human heart as seen in a computed tomography image. Construct a 3D model of a given patient heart. Use the 3D model for diagnostics and surgery planning.




Label image from semantic segmentation

- Scene understanding for self navigating vehicles




<https://towardsdatascience.com/semantic-segmentation-of-150-classes-of-objects-with-5-lines-of-code-7f244fa96b6c>

Deep learning for semantic segmentation




Badrinarayanan, Vijay, Alex Kendall, and Roberto Cipolla. "Segnet: A deep convolutional encoder-decoder architecture for image segmentation." IEEE transactions on pattern analysis and machine intelligence 39.12 (2017): 2481-2495.

Label images in this course



Pixel Classification

BLOB analysis and 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 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 units
 - 0 water
 - -1000 air
 - -120 fat
 - 400+ bone




Floating point images

- The pixel type is often changed when applying image processing functions
- For example when scaling an image, the output will be a floating point image:

```
image_rescaled = rescale(im_org, 0.25, anti_aliasing=True, channel_axis=2)
print(image_rescaled.shape)
print(image_rescaled.dtype)
```



Python scripts vs. Jupyter Notebooks

- In this course, you can do the exercises and the exam in both Jupyter Notebooks or as Python scripts
- Strengths and benefits of both approaches





What is your experience with Github?





PCA Analysis

Next week:

Image acquisition, digital cameras, compression and storage and real-time image analysis

