A lightning introduction to (modern) image processing

Pierre Weiss, CNRS, Université de Toulouse

M2R Bio-Santé

Introduction

Objectives of this lecture

- Quick introduction on the need for computers:
    - image acquisition
    - image restoration
    - image analysis
    - image generation
- Starting with the basics:
    - Python programming
    - Pixel operations
    - Linear filtering and convolutions
    - Morphological operations
    - Inverse problems
- Some basic principles of deep learning

These are very introductory lectures!

Watch this excellent contents for further information.

Computers and image acquisition

A modern view on image acquisition

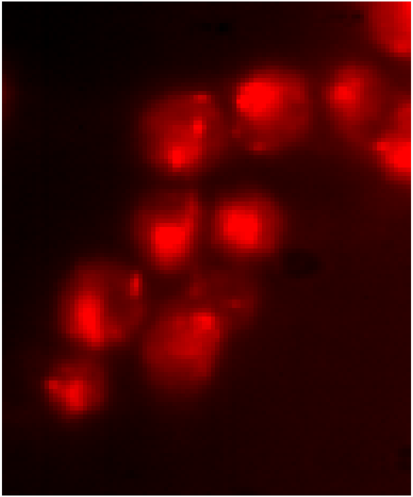

Baker's yeast with a wide-field microscope - Resolution $= 0.2\mu m$

A modern view on image acquisition

Nobel prize E. Betzig 2014 - resolution 50nm.
Image by S. Cantaloube, T. Mangeat

A modern view on image acquisition

Single Molecule Localization Microscopy, histone H3 (10^6 images).

Introduction: single molecule localization microscopy

From microscopy to nanoscopy

1. Standard microscope: 200-500 nanometers (1mm / 5000)
2. SMLM: 10-50 nanometers (1mm / 100,000)
3. Atom size: 0.1nm (1mm / 10, 000, 000)

Current limit is mostly computation (and thermal agitation)

A usual camera (WYSIWYG)

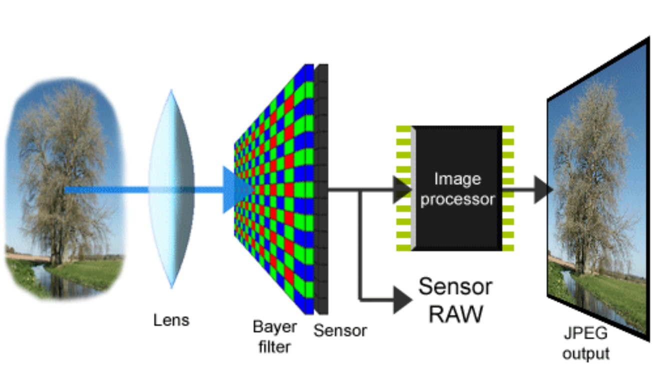

A modern view on image acquisition

Why is this not so simple?

- Impossible to measure the signal directly.
- Sub-optimal as well.

Examples

- Echography, oil prospection.
- Magnetic Resonance Imaging, radio-interferometry.
- X-ray Tomography.
- In fact nearly any device...

X-Ray Tomography

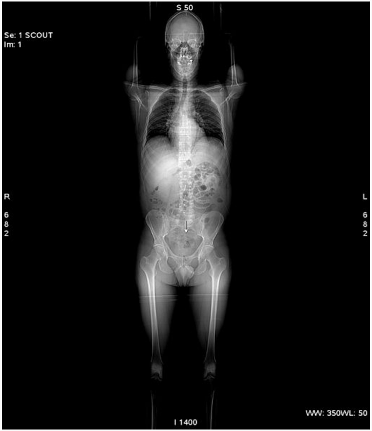

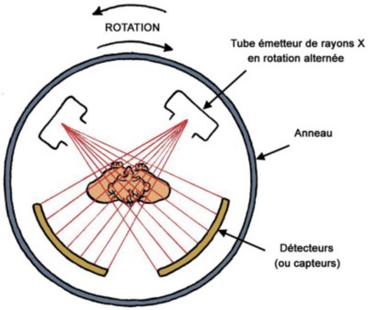

Magnetic Resonance Imaging

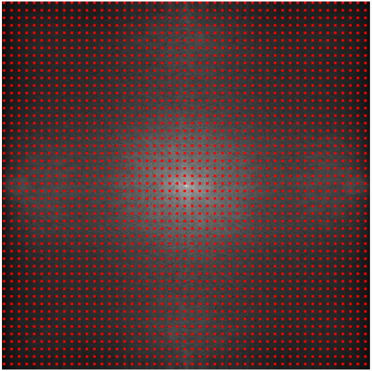 

Summary

$y_m = \langle$ 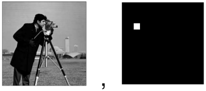 $\rangle$

Camera

$y_m = \langle$ 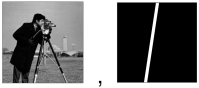 $\rangle$

Tomography

$y_m = \langle$ 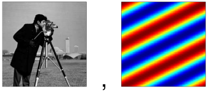 $\rangle$

MRI

$y_k = \langle$  , ? $\rangle$

Other choices?

How to reconstruct images?

1. First step: modelling the system

$$y = Ax + b, \tag{1}$$

   where

   - The image $x \in \mathcal{X}$ is a vector
   - The operator $A : \mathcal{X} \to \mathbb{R}^M$ describes the system
   - The perturbation b models noise

2. Second step: design a reconstruction algorithm

$$\hat{x} = \operatorname*{argmin}_{x \in \mathcal{X}} \frac{1}{2} \|Ax - y\|_2^2 + R(x), \tag{2}$$

   where R promotes "realistic looking images".

3. Third step: design algorithms to interpret/extract/quantify data

A modern view on image acquisition

What can be done with computational imaging?

Image super-resolution (20Mpix image)

What can be done with computational imaging?

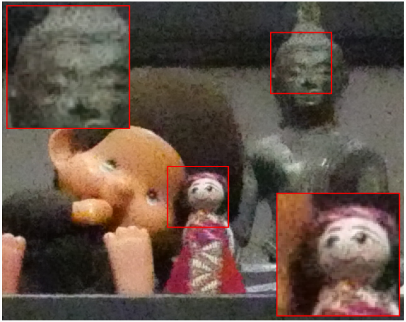

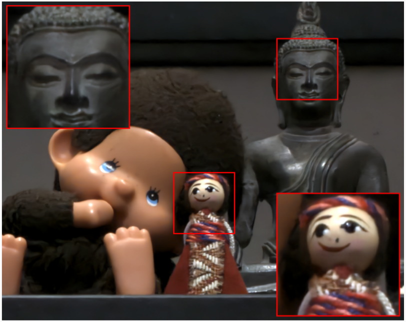

Image super-resolution ×4 with 30 burst images

A modern view on image acquisition

What can be done with computational imaging?

Image super-resolution ×4 with 30 burst images

A modern view on image acquisition

What can be done with computational imaging?

Image super-resolution ×4 with 30 burst images

A modern view on image acquisition

Computational imaging: augmented reality

Computers and image improvement

Image improvement

Improving the image quality

Most images suffer from problems

- Noise
- Blur
- Unwanted objects
- Too large size

Image improvement

Low signal-to-noise-ratio

Image denoising

Discarding structured noise

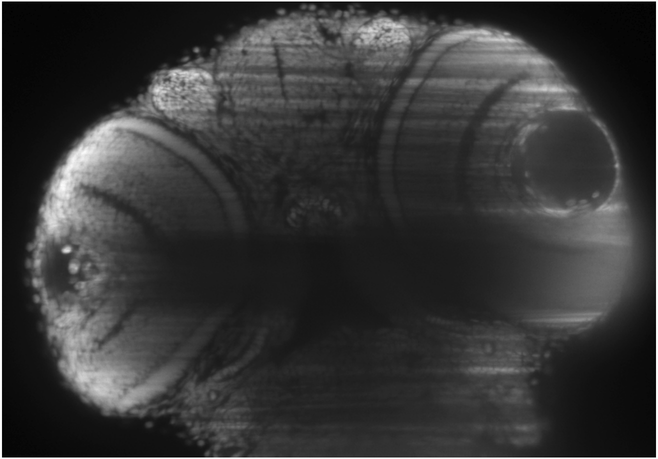

Image denoising (LSFM image of zebrafish)

Discarding structured noise

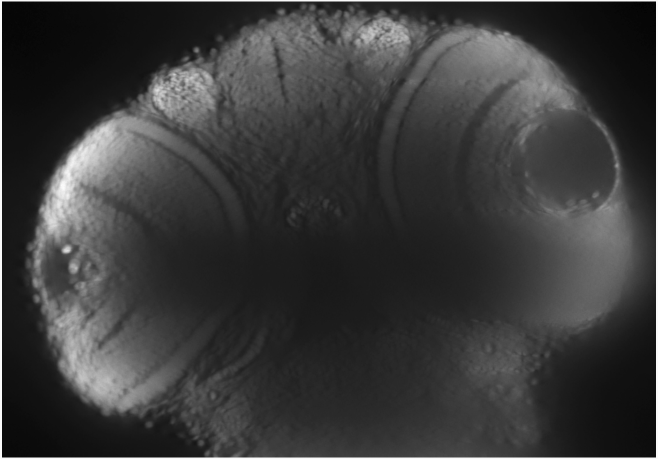

Image denoising (here with VSNR)

Sharpening images

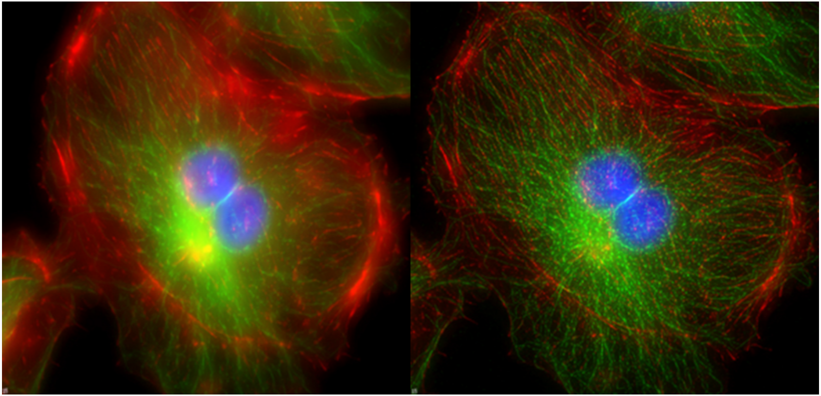

Image deblurring

Explain the difference between blind and non blind image deblurring

Image improvement

Discarding unwanted objects

Image inpainting

Discarding unwanted objects

Image inpainting

Image improvement

Reducing the image sizes

Original (226 KB)
PNG

(16 KB)
JPEG

(16 KB)
JPEG 2000

Image Compression

Think about it seriously (though noone usually cares):

- Buy a larger server to store your data, or...

Image improvement

Reducing the image sizes

Image Compression

Think about it seriously (though noone usually cares):

- Buy a larger server to store your data, or...
- Work a bit to keep only useful information?

Computers and image analysis

Image analysis

How to have a computer automatically:

- Segment objects?
- Detect objects?
- Interpret the image contents?

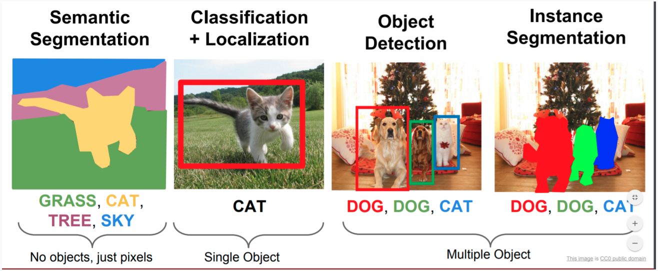

The different types of image segmentation

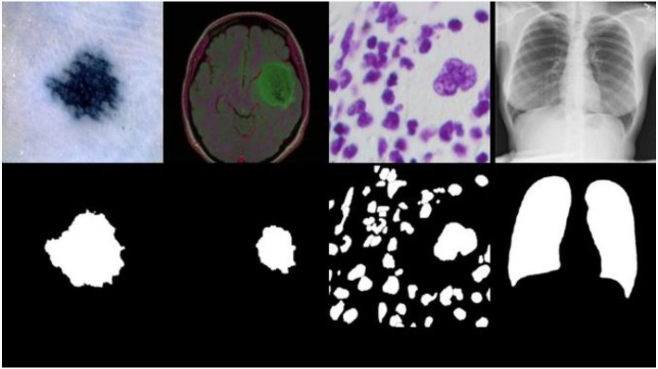

Image segmentation

Image analysis

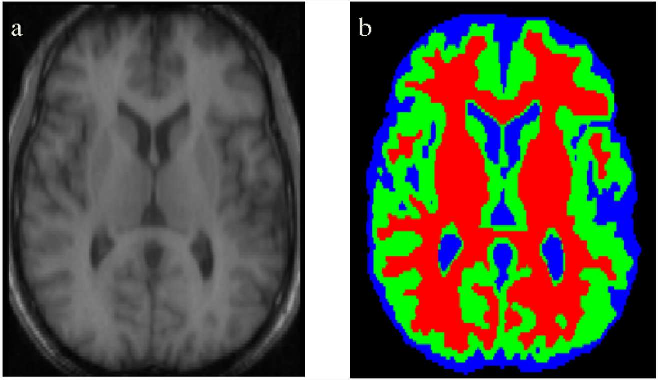

Semantic segmentation

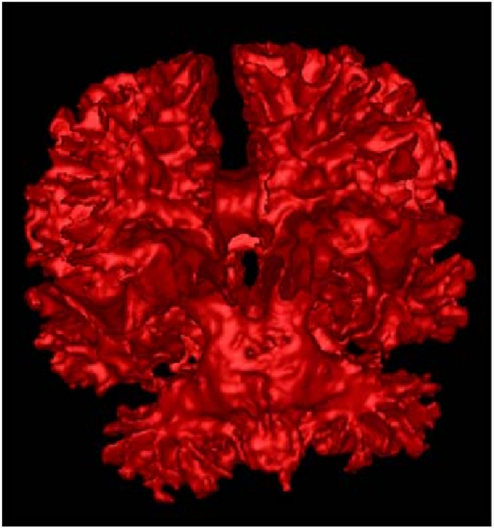

Semantic segmentation

Image analysis

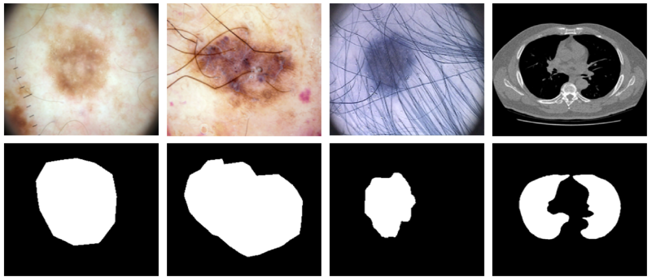

Semantic segmentation – harder cases

Why is semantic segmentation important?

- Assess volumes, boundaries, numbers of cells...
- Assist doctors, biologists...
- Reduce human subjectivity
- Some automatic algorithms now perform better than humans
- Treat large volumes of data for statistical analysis

Image analysis

Image classification

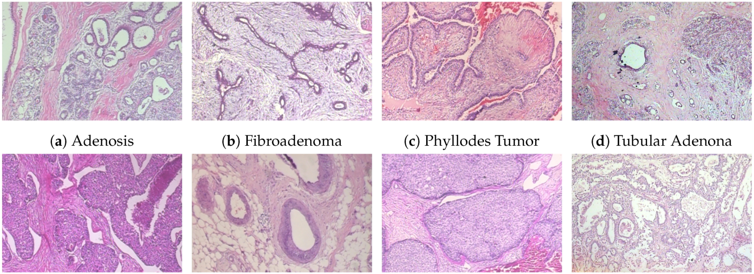

(**a**) Adenosis (**b**) Fibroadenoma (**c**) Phyllodes Tumor (**d**) Tubular Adenona

(**e**) Ductal Carcinoma (**f**) Lobular Carcinoma (**g**) Mucinous Carcinoma (**h**) Papillary Carcinoma

Image classification

We'll see this aspect in Part II.

Image interpretation

Image interpretation

An example of automatic scene interpretation

Image analysis

Image interpretation

An example of application for self-driving cars

Image analysis

Unexpected applications

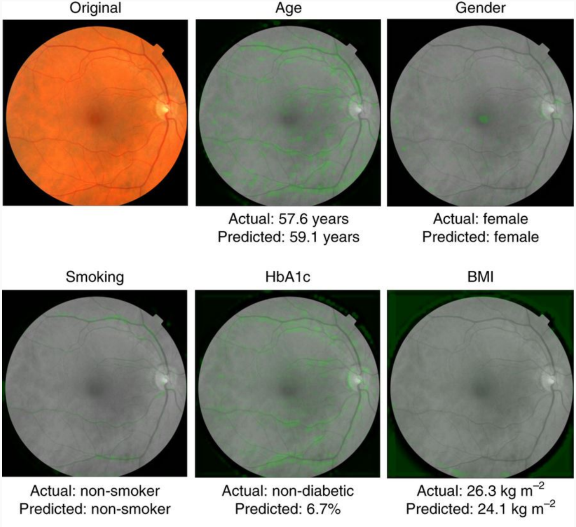

What can be deduced from fundus images?

Computers and image/model generation

Deep face

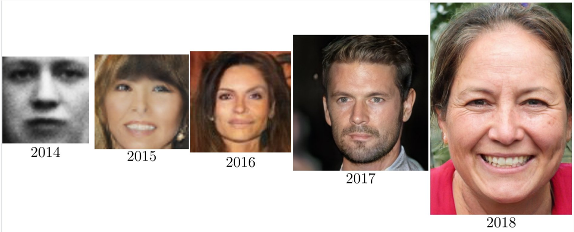

The rapid evolution of deepface
Now it is perfect enough ⇒ new applications

Image generation

Animated deepfaces = deepfakes

Diffusion models – Dall-E

Draw an astronaut riding a horse on the moon

Diffusion models – Dall-E

Draw a very muscled teapot

Diffusion models – Dall-E

Complete this painting of Vermeer (Girl with a ring)

Impact in biology?

- Image acquisition/analysis now relies on the same mechanisms
- Computers can learn complex patterns
- What if we could train Deepcells? Deeporganisms?
- These technologies have not yet strongly entered biology...
- Things may change rapidly... Stay tuned

Ending words

Why should you care about computers?

- Image acquisition (all modern digital devices rely on computing)
- Image improvement (deblurring, denoising, inpainting)
- Image analysis (segmentation, detection, classification, interpretation)
- Image generation (learning complex models from observations)

How to use computers?

- Acquisition/improvement:
  - Physics/Mathematics model of acquisition device (with manufacturers)
  - Invert the acquisition model (with optimization, neural networks)
- Analysis (the main focus of this course):
  - Python, thresholding, filtering, morphology, iterative methods (Part I)
  - Basics of deep learning (Part II)

Hands-on!

Many existing languages: C, C++, Java, Matlab, Fiji, Icy,...

Why Python?

- Python is getting dominant for imaging
- Plenty of libraries are currently developed/maintained
- Wide community with plenty of forums
- 99% of deep learning models
- Allows GitHub and ctrl+c ctrl+v programming
- Relatively easy and universal
- Basis of modern tools such as Napari

Images and Python

Importing image processing libraries

Plenty of powerful image processing libraries

- OpenCV
- Scikit image
- MatPlotLib
- Scikit learn
- Numpy
- Pillow...

Machine learning libraries

- Scikit-Learn
- PyTorch
- TensorFlow
- ...

Opening an image

- Import scikit image
- Open a 2D gray-scale image
- Open a 2D colour image
- Open a 3D gray-scale image
- Open a 3D colour image
- Show channels, slices
- Crop, flip

Representing an image

- Show a surface plot
- Show a contour plot
- Explain level lines
- Change lookup tables

Pointwise operations in Python

Pointwise operations

- Inverting contrasts
- Thresholding
- Change of contrast
- Sine (or other random grayscale changes)
- Quantification

Convolution/filtering

- Take a look here
- Average blur
- Gaussian blur
- Edge detection (x,y gradients + norm of gradient)
- Laplacian
- General convolution + implementation with for loops

Morphological operations in Python

Image morphology

- You can take a look here
- Invariance to contrast changes + level lines
- Median filtering (+ pepper noise, structuring elements)
- Dilation, erosion, opening, closing

Iterative methods

- Total variation deblurring
- Active contours

The limits of handcrafted methods

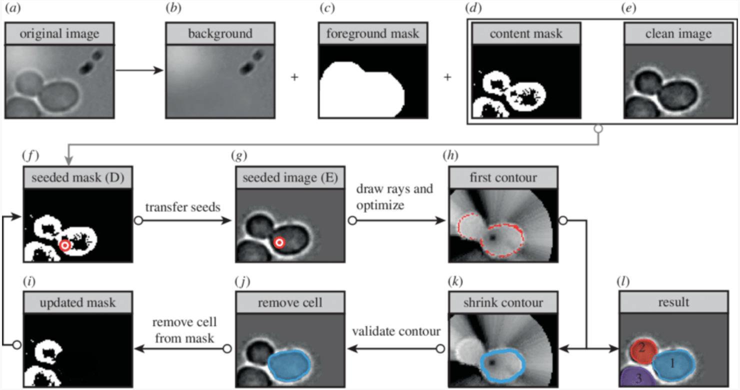

A typical pipeline

Lightning presentation of supervised learning

Supervised learning

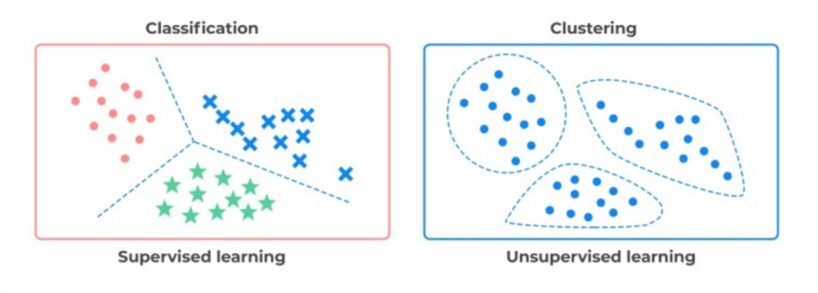

Supervised vs unsupervised learning

- Supervised: labeled data
- Unsupervised: unlabeled data (clustering)

Where to get your data

Many open datasets/challenges are being created

- Fast MRI (10000+ 3D volumes)
- ImageNet (10^6+ images)
- Data science bowl
- Go Pro deblur challenge

Supervised learning

Where to get your data

Many open datasets/challenges are being created

- Fast MRI (10000+ 3D volumes)
- ImageNet (10^6+ images)
- Data science bowl
- Go Pro deblur challenge

Create your own!

Really worth the investment on the long run!

Data + computing $\equiv$ sinews of war

Supervised learning

Data = pair (image, desired output)

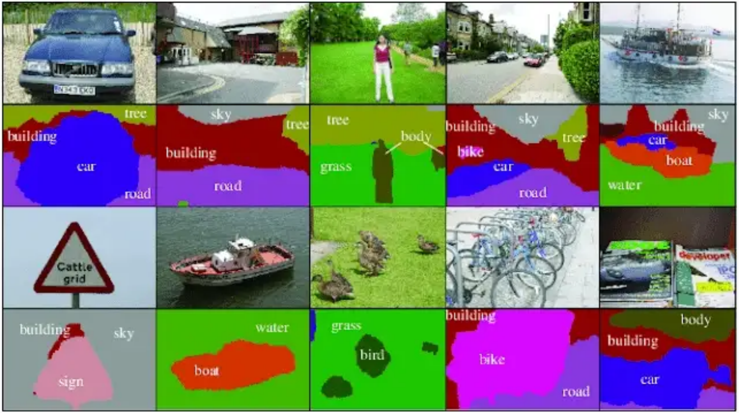

Labeling for image classification

Supervised learning

Data = pair (image, desired output)

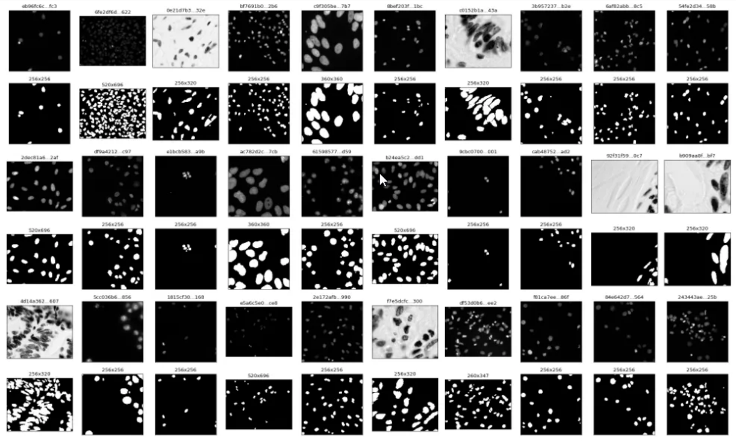

Labeling for image segmentation

Supervised learning

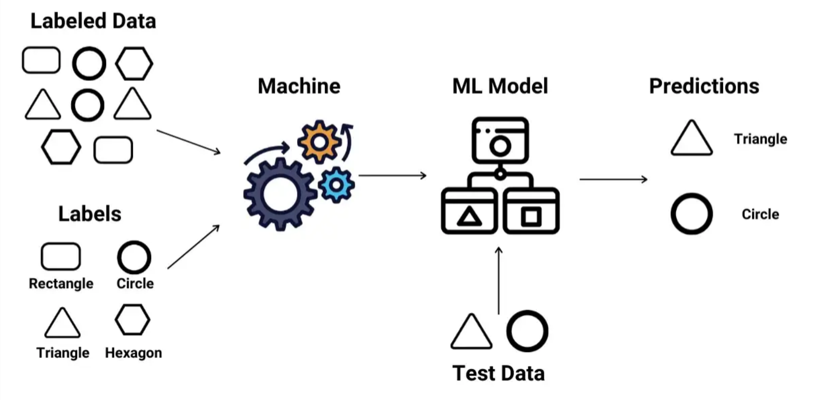

The basic principle...

Opening the black-box

The main existing tools:

- Linear regression
- Support vector machines
- Decision trees
- Random forests
- Neural networks

Decision trees

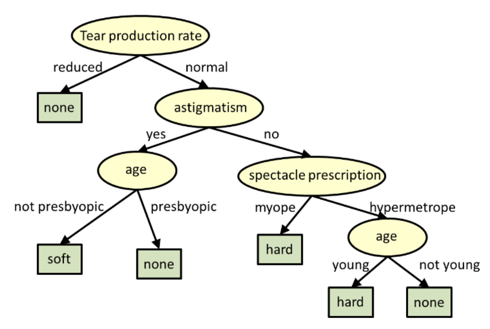

A decision tree to decide the type of lens

Tree = decision path based on features

- Easy to interpret ($\simeq$ human decision)
- Not so efficient

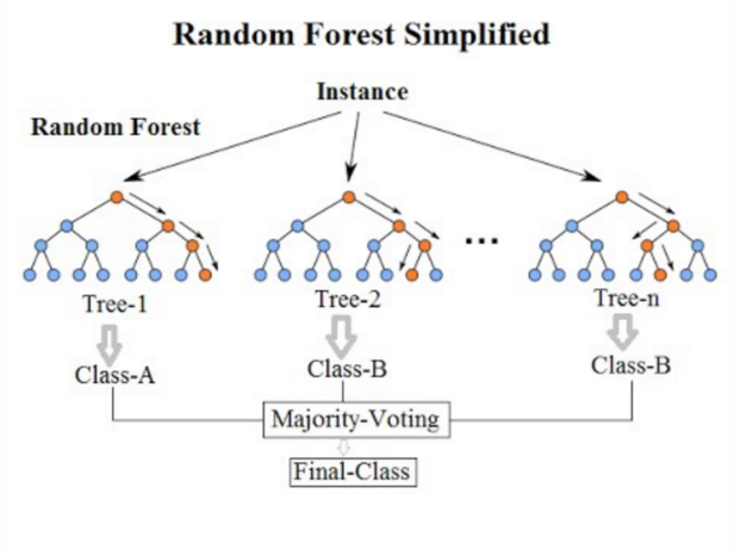

A random forest classifier – Hands on Ilastik!

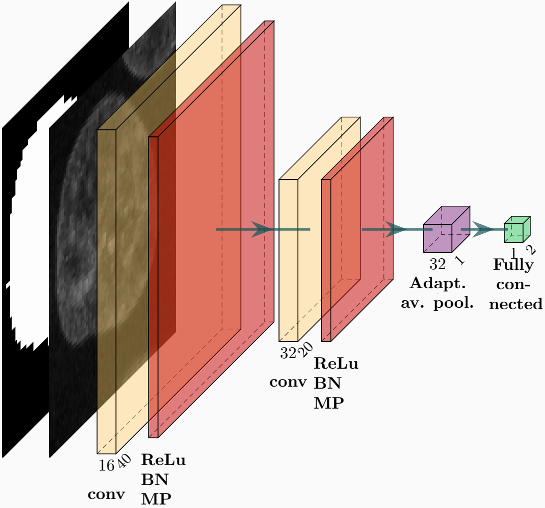

A Neural Network for Classification (depth = 2)

Supervised learning

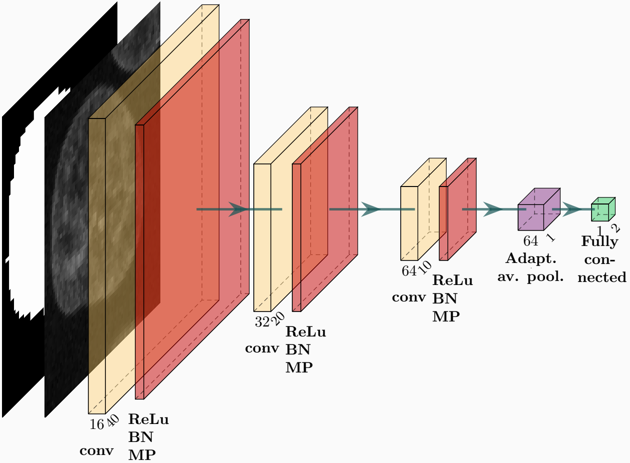

A Neural Network for Classification (depth = 3)

Supervised learning