

# Introduction to Medical Image Processing with Python

**Michal Kazmierski**



*michal.kazmierski@mail.utoronto.c  
a*

Haibe-Kains Lab

# Agenda


- **Introduction & brief overview of medical imaging**
- **Imaging basics**
  - What is an image?
  - Image formation
  - Images are objects in space
  - Digital imaging, sampling and quantization
- **Operations on images**
  - Resampling & interpolation
  - Thresholding
  - Noise
  - Convolution and filtering
- **Overview of machine learning in medical imaging**
- **Practical workshop in Python**

# What is medical imaging?

**Medical imaging:** non-invasive examination of **structure, physiology** and **pathology** inside the human body.

# What is medical imaging?


**Medical imaging:** non-invasive examination of **structure, physiology** and **pathology** inside the human body.




**Diagnose cancer:** computed tomography (CT), magnetic resonance imaging (MRI)

# What is medical imaging?

**Medical imaging:** non-invasive examination of **structure, physiology** and **pathology** inside the human body.




**Diagnose cancer: CT, MRI**




**Diagnose and monitor neurodegenerative diseases: positron emission tomography (PET), MRI**

# What is medical imaging?


**Medical imaging:** non-invasive examination of **structure, physiology** and **pathology** inside the human body.



**Diagnose cancer: CT, MRI**




**Diagnose and monitor neurodegenerative diseases: PET, MRI**




**Identify traumatic injuries: CT, radiography**

# What is medical imaging?


**Medical imaging:** non-invasive examination of **structure, physiology** and **pathology** inside the human body.



**Diagnose cancer: CT, MRI**



**Diagnose and monitor neurodegenerative diseases: PET, MRI**



**Identify traumatic injuries: CT, radiography**



**Examine foetal health *in utero*: ultrasound (US)**

*From left:*


*Ardila et al., Nat Med. (2019)*

*Okamura et al. Clin Transl Imaging (2018)*

*Mutch et al. Neurosurg Clin N Am. (2016)*

*GE Healthcare in Prince & Links, Medical Imaging Signals and Systems 2nd ed. (2015)*


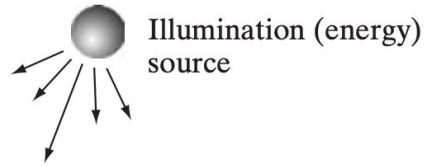
# Imaging basics



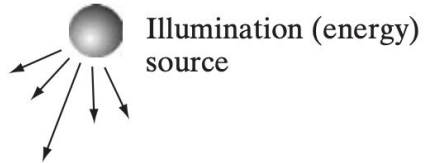
Object




# Imaging basics



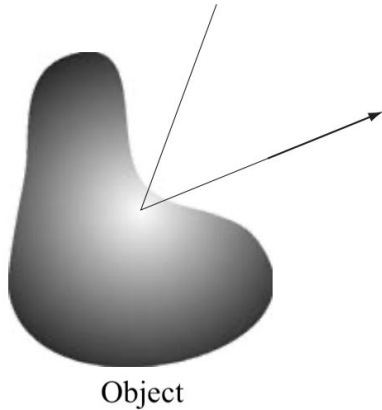
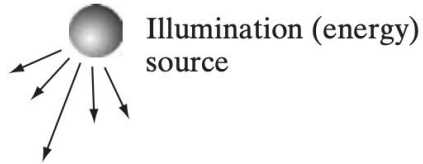
# Imaging basics



- **Visible light source (photography)**
- **X-ray tube (CT, radiography)**
- **Radioisotope injected into the patient (PET, SPECT)**
- **RF pulse (MR)**



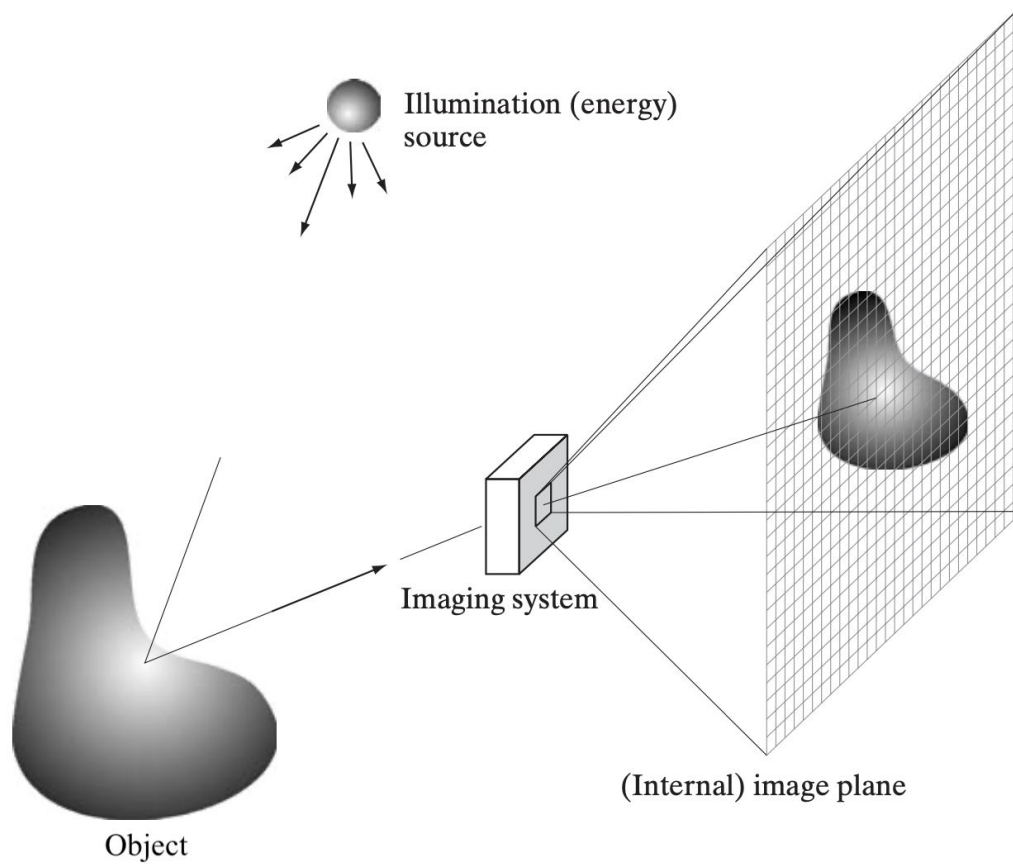
# Imaging basics



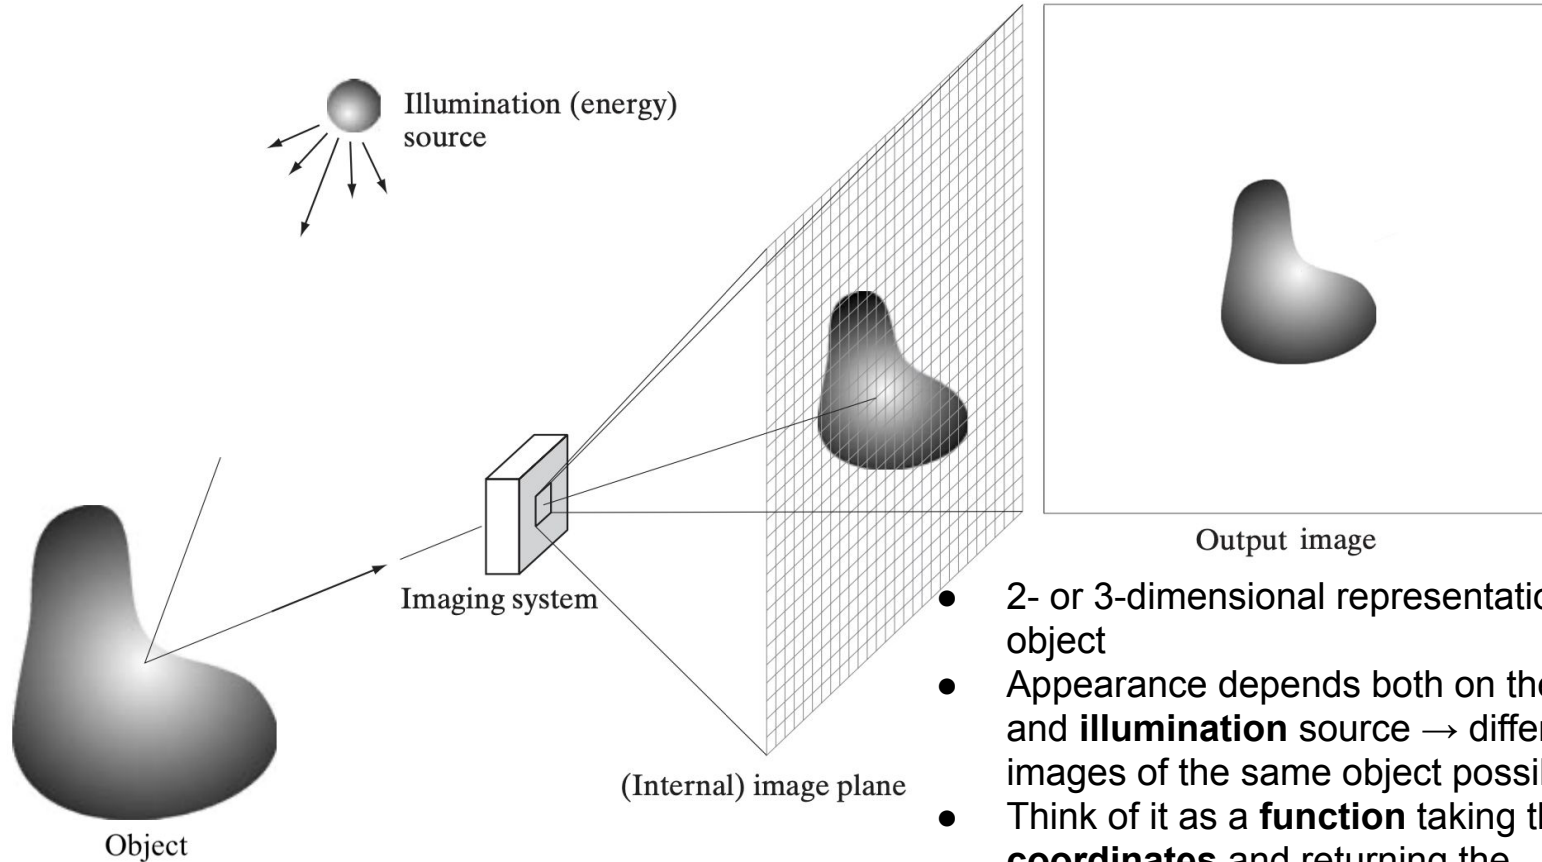
**Interaction** depends on illumination & object:

- Light reflected off the skin
- Change in X-ray energy depending on tissue type

# Imaging basics




# Imaging basics




- 2- or 3-dimensional representation of the object
- Appearance depends both on the **object** and **illumination** source → different images of the same object possible
- Think of it as a **function** taking the **spatial coordinates** and returning the **measurement at given location**

# Imaging basics



Image

=




Function of  
two/three  
arguments  $(x, y, z)$

# Imaging basics

## Example: projection radiography (“X-ray”):



- **Object:** the patient’s body
- **Energy source:** X-ray tube
- **Imaging system:** radiographic film/digital sensor
- **Image:** function of (x, y) coordinates giving the measured X-ray intensity at that point



# Imaging basics


## Example: computed tomography (CT):

- **Object:** the patient's body
- **Energy source:** X-ray tube
- **Imaging system:** CT scanner
- **Image:** function of (x, y, z) coordinates giving the measured X-ray intensity at that point (3D image made up of 2D slices)






# Digital imaging



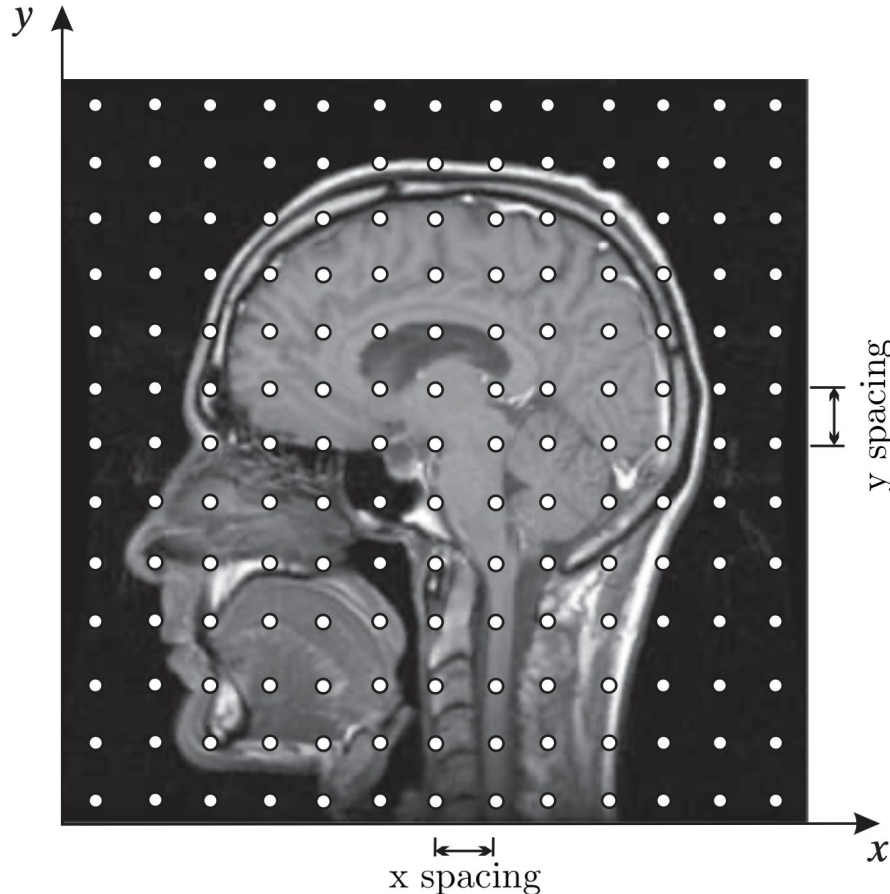
# Digital imaging

- **Analog signals:**  
take values in a **continuous range** and are defined on **continuous set of points**
- **Digital signals:**  
take values from a **discrete set** and are defined on **discrete points**
- Computers work with **digital signals**;  
any signal can be digitized in 2 steps:
  - **Sampling** (discretizing the coordinates)
  - **Quantization** (discretizing the values)




The **red** signal is obtained by digitizing the **grey** signal

# Sampling



- Taking measurements at points in a discrete set (e.g. on a rectangular grid)
- Point measurements are called **pixels** (*picture elements*) in 2D or **voxels** (*volume elements*) in 3D (they are *mathematical points*, **not** little squares/cubes!)
- **Pixel (voxel) spacing** is the distance between measurement points, determines the **spatial resolution**


# Quantization




- Similar to sampling, but on measured **intensity values**
- Example:
  - X-ray intensity can take any value  $> 0$
  - Restrict to “bins” every 10 units
  - E.g. [26.94, 35.98, 0.48, 15.13] → [20, 30, 0, 10]
- Number of quantization levels determines the **intensity resolution**

# Images are objects in space

## Image (Voxel) Coordinate System




## World Coordinate System




# The three components of image geometry


- **Origin:** the location of the (0, 0, 0) voxel in world coordinates (with respect to some known reference point in the scanner)
- **Spacing:** the distance between voxels in (x, y, z) directions in mm
- **Direction:** the 3D rotation of the coordinate axes (most commonly no rotation, used in certain CT and MR acquisition protocols)



**Origin:** translation




**Spacing:** scaling




**Direction:** rotation

# Why it matters?



Spacing: 1 px = 1 mm


# Why it matters?




Spacing: 1 px = 1 mm  
Image space size: 100 px



# Why it matters?




Spacing: 1 px = 1 mm  
Image space size: 100 px




Spacing: 1 px = 2 mm

# Why it matters?




Spacing: 1 px = 1 mm  
Image space size: 100 px




Spacing: 1 px = 2 mm  
Image space size: 50 px

# Why it matters?





**Spacing: 1 px = 1 mm**  
**Image space size: 100 px**  
**Real size: 100 mm**



**Spacing: 1 px = 2 mm**  
**Image space size: 50 px**  
**Real size: 100 mm**

# Why it matters?

- A **PET-CT** machine acquires images in 2 modalities at the same time
- Can easily match them knowing the position of the patient in scanner coordinates




# Resampling


**Resampling:** sampling a sampled image

- Change size (e.g. for faster processing with smaller images)
- Harmonize acquisition parameters (e.g. one scanner might use .5 mm spacing, another 1 mm spacing)
- Perform geometric transformations (rotations, translations, etc.)

# Resampling example




# Resampling example



**Spacing: (1 mm, 1 mm)**

# Resampling example




**Spacing: (1 mm, 1 mm)**


**New spacing: (1.5 mm, 1.5 mm)**



# Resampling example




**Spacing: (1 mm, 1 mm)**




**New spacing: (1.5 mm, 1.5 mm)**

# Resampling example




**Spacing: (1 mm, 1 mm)**  
**Size: (512 px, 512 px)**



**New spacing: (1.5 mm, 1.5 mm)**  
**New size: (341 px, 341 px)**


# Resampling: interpolation

- In general, the new sampling points will not align with the current sampling grid

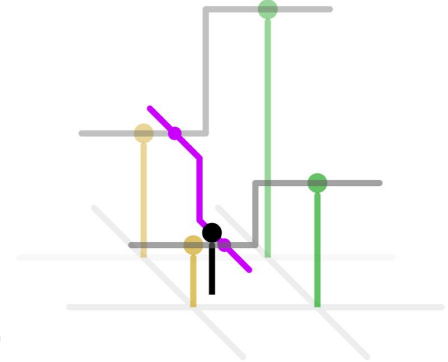


# Resampling: interpolation

- In general, the new sampling points will not align with the current sampling grid
- Simple solution: use the nearest known sample (**nearest neighbour interpolation**)



1D nearest-neighbour



2D nearest-neighbour

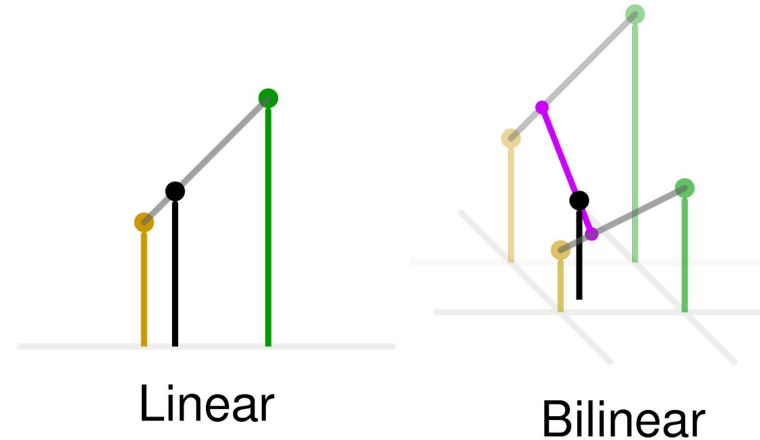
# Resampling: interpolation

- In general, the new sampling points will not align with the current sampling grid
- Simple solution: use the nearest known sample (**nearest neighbour interpolation**)



# Resampling: interpolation

- In general, the new sampling points will not align with the current sampling grid
- Simple solution: use the nearest known sample (**nearest neighbour interpolation**)
- Better solution: assume intensity changes approximately linearly between nearby values, fit 2 sets of lines (in x/y directions) and read out the value (**(bi)linear interpolation**)





# Resampling: interpolation

- In general, the new sampling points will not align with the current sampling grid
- Simple solution: use the nearest known sample (**nearest neighbour interpolation**)
- Better solution: assume intensity changes approximately linearly between nearby values, fit 2 sets of lines (in x/y directions) and read out the value (**(bi)linear interpolation**)





# Resampling: rotation





# Resampling: rotation

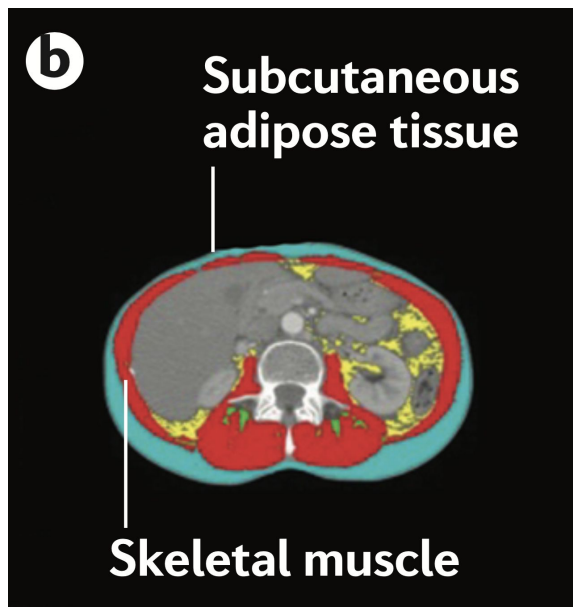


# Intensity transformations


- Most medical imaging modalities are **greyscale**: a single scalar value per pixel → pixel values are called **grey-level intensities**
- Some image modalities have defined (pseudo) units for grey level intensities, allowing absolute comparison between images:
  - e.g CT: Hounsfield units (HU), defined as 0 for water and -1000 for air
- Other modalities can have arbitrary grey-level values with no defined units, allowing only for relative comparison within an image
  - e.g. T2-weighted MR

# Intensity transformations: segmentation

- **Image segmentation:** separating one or more *regions of interest* from regions that do not contain information relevant for the task



Separate different types of normal tissue



Separate tumour from normal tissue

# Thresholding

- Simplest method of **image segmentation**: separating one or more *regions of interest* from regions that do not contain relevant information
- Useful for e.g. locating a specific tissue type in an image
- Algorithm:


for each pixel in image:

if (pixel  $\geq$  lower threshold) and (pixel  $\leq$  upper threshold):

set pixel to 1


otherwise:

set pixel to 0



# Noise


- **Noise** refers to variation in grey level intensity values that does not correspond to real features of the object



*Image corrupted with additive Gaussian noise.*

# Noise


- **Noise** refers to variation in grey level intensity values that does not correspond to real features of the object
  - Random variation due to stochastic nature of underlying physical processes
  - Errors during image reconstruction
  - Corruption during transmission/storage



*Image corrupted with additive Gaussian noise.*

# Noise


- **Noise** refers to variation in grey level intensity values that does not correspond to real features of the object
  - Random variation due to stochastic nature of underlying physical processes
  - Errors during image reconstruction
  - Corruption during transmission/storage
- The best way to remove noise is at the source; can be reduced using **filtering** methods



*Image corrupted with additive Gaussian noise.*

# Filtering



- Altering the intensity values in an image to achieve a particular effect
- Output determined by the specifics of the operator (**filter**)
- Basis of most intensity-altering operations






# Example: Gaussian blur

- Replaces each pixel in the image by the average of values in its neighbourhood weighted using the Gaussian function
- Gives more weight to nearby values, less to more distant neighbours
- Can be efficiently calculated using **convolution**



# Example: Gaussian blur



# Machine learning in medical imaging

- Most of the algorithms we've talked about so far require explicit programming (e.g. thresholding: need to specify the threshold bounds)

# Machine learning in medical imaging

- Most of the algorithms we've talked about so far require explicit programming (e.g. thresholding: need to specify the threshold bounds)
- **Machine learning** algorithms build a mathematical model of sample data in order to make predictions or decisions without being explicitly programmed to perform the task. (Wikipedia)


# Machine learning in medical imaging

- Most of the algorithms we've talked about so far require explicit programming (e.g. thresholding: need to specify the threshold bounds)
- **Machine learning** algorithms build a mathematical model of sample data in order to make predictions or decisions without being explicitly programmed to perform the task. (Wikipedia)
- ⇒ can use general purpose learning algorithms (currently trending: deep convolutional neural networks) to **learn to solve problems from large datasets**

# Machine learning in medical imaging

## Classification:

e.g. prognosis in head & neck cancer




**Output: image-level class labels**

# Machine learning in medical imaging

## Semantic segmentation:

e.g tumour & adjacent organ segmentation for radiotherapy planning




- Brain
- Brainstem
- Mandible
- Parotid Left / Right

**Output: pixel/voxel-level class labels**

# Machine learning in medical imaging

## De-noising & superresolution:

e.g. improving quality of cardiac MRI




**Output: higher quality image**



# Machine learning in medical imaging

**Object detection, localisation, recognition:**

e.g. detection & malignancy classification of lung nodules




**Output: object location/bounding box & class probability**

# Machine learning in medical imaging

## Image registration:

e.g. multi-modal image fusion



**Output: aligned images**

# Machine learning in medical imaging

- That doesn't mean the simple algorithms are useless or obsolete! They in fact serve as components of ML systems (e.g. in pre-processing) and can be used to solve certain problems without requiring a lot of training data

# Useful resources

- SimpleITK Notebooks (<http://insightsoftwareconsortium.github.io/SimpleITK-Notebooks/>) **A much more thorough overview of SimpleITK features, including advanced concepts like image registration.**
- SimpleITK documentation ([https://itk.org/SimpleITKDoxygen/html/namespaceitk\\_1\\_1simple.html](https://itk.org/SimpleITKDoxygen/html/namespaceitk_1_1simple.html)) **The ultimate SimpleITK reference, although at times difficult to navigate.**
- A. R. Smith, 'A Pixel Is Not A Little Square, A Pixel Is Not A Little Square, A Pixel Is Not A Little Square! (And a Voxel is Not a Little Cube)', p. 11, Jul. 1995. **A good explanation of image reconstruction in general, and why the “small square” model of pixels is not correct.**
- R. C. Gonzalez and R. E. Woods, Digital image processing, 3rd ed. Upper Saddle River, N.J: Prentice Hall, 2008. **Classic computer vision textbook, covers most aspects of traditional image processing.**
- M. A. Haidekker, Advanced biomedical image analysis. Hoboken, N.J: John Wiley & Sons, 2011. **Very good book with a lot of useful medical imaging-specific (and not only) algorithms, although some approaches have been superseded by machine learning.**
- J. L. Prince and J. M. Links, Medical imaging signals and systems, 2nd ed. Boston: Pearson, 2015. **Good introductory textbook focusing on the mathematics and physics of medical imaging.**
- PyDicom (<https://pydicom.github.io/pydicom/stable/>) and the DICOM standard (<http://dicom.nema.org/medical/dicom/current/output/html/part01.html>) **Very useful when working with clinical images, especially in oncology.**
- A. Hosny, C. Parmar, J. Quackenbush, L. H. Schwartz, and H. J. W. L. Aerts, 'Artificial intelligence in radiology', Nature Reviews Cancer, vol. 18, no. 8, pp. 500–510, Aug. 2018, doi: 10.1038/s41568-018-0016-5. **Interesting review of current applications and challenges for machine learning in radiology.**



# Python workshop

# Tools of the trade: SimpleITK

## Pros:

- Seamlessly handles 2D and 3D images
- Keeps track of image geometry (spacing, direction, world origin)
- Handles many common image storage formats (DICOM, NIfTI, NRRD)
- Many fast algorithms for image processing, tailored to medical images

## Cons:

- Python wrapper around a C++ library → API can get pretty ugly
- Integration with the rest of the Python scientific stack is not great
- Has its own way of doing things, often subtly incompatible with the rest of Python ecosystem

# Tools of the trade: Numpy, SciPy and Matplotlib

## Pros:

- The Python scientific computing stack is built on them
- Backed by very fast C/Fortran libraries
- Excellent implementations of many common algorithms
- Seamless integration with many ML libraries (scikit-learn, PyTorch, Tensorflow)

## Cons:

- The basic data structure (Numpy ndarray) not made for imaging
- Do not keep track of image geometry, need to do it manually
- Lack I/O for many medical image formats
- Do not include implementations of more specialized medical image processing algorithms

# Tools of the trade: Numpy, SciPy and Matplotlib

## Pros:

- The Python scientific computing stack is built on them
- Backed by very fast C/Fortran libraries
- Excellent implementations of many common algorithms
- Seamless integration with many ML libraries (scikit-learn, PyTorch, Tensorflow)

## Cons:

- The basic data structure (Numpy ndarray) not made for imaging
- Do not keep track of image geometry, need to do it manually
- Lack I/O for many medical image formats
- Do not include implementations of more specialized medical image processing algorithms

**General recommendation:** Load the image and do as much processing as possible with SimpleITK, convert to Numpy array for visualization/some algorithm not implemented in SimpleITK/fancy deep learning stuff.