Australian Government

ANSTO

AI in medical imaging

Status, opportunities, pitfalls and challenges

Dr Mitra Safavi-Naeini

Principal Physicist - ANSTO

Science. Ingenuity. Sustainability.

ANSTO | We are the national nuclear science and technology agency

Operating safely for over 60 years

Home of Australia's landmark research infrastructure

Lucas Heights | NSW

Main campus

Clayton | VIC

Australian Synchrotron

Camperdown | NSW

Cyclotron

ANSTO

Artificial intelligence: what is intelligence?

ANSTO

Big Data

- Three Vs:
  - Volume: Exponential increase
  - Velocity: Rate at which it is produced
  - Variety: Different formats with no or little structure

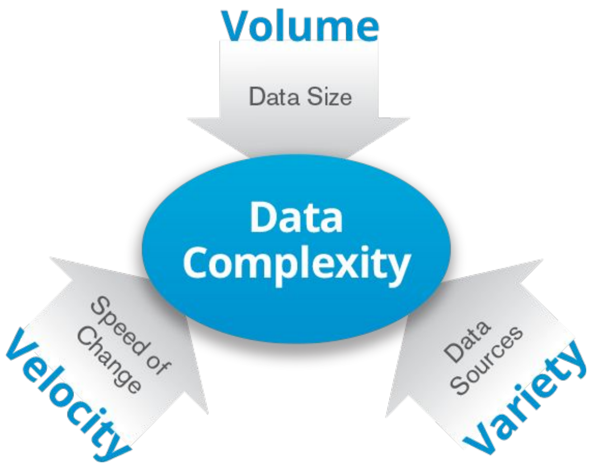

- There is an additional one that matches medical data:
  - Veracity: Data might be incomplete, noisy, not meaningful, uncertain, ambiguous and inconsistent.

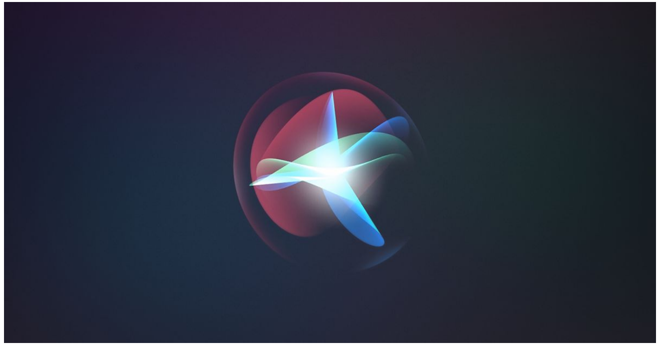

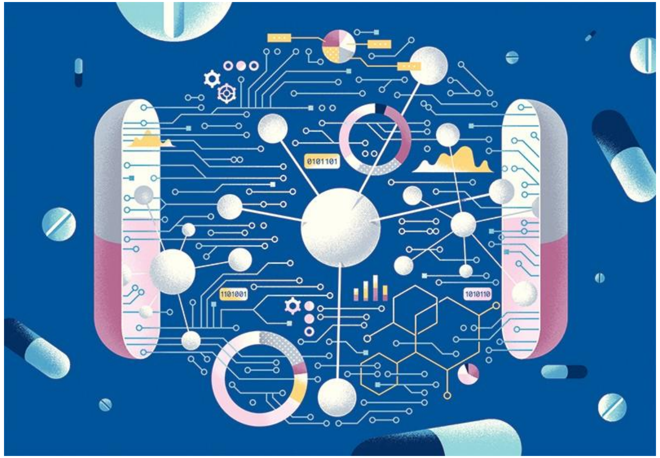

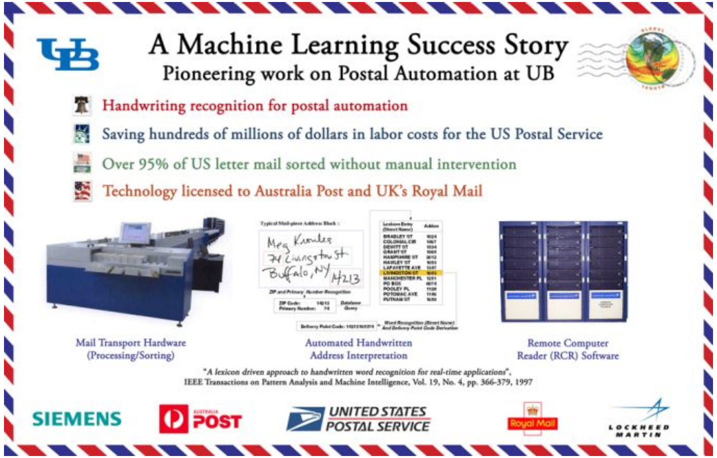
A Machine Learning Success Story
Pioneering work on Postal Automation at UB
Handwriting recognition for postal automation
Saving hundreds of millions of dollars in labor costs for the US Postal Service
Over 95% of US letter mail sorted without manual intervention
Technology licensed to Australia Post and UK's Royal Mail
Mail Transport Hardware (Processing/Sorting)
Automated Handwritten Address Interpretation
Remote Computer Reader (RCR) Software
"A lexicon driven approach to handwritten word recognition for real-time applications",
IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 19, No. 4, pp. 366-379, 1997
SIEMENS
POST
UNITED STATES POSTAL SERVICE
Royal Mail
LOCKHEED MARTIN

Is it AI?

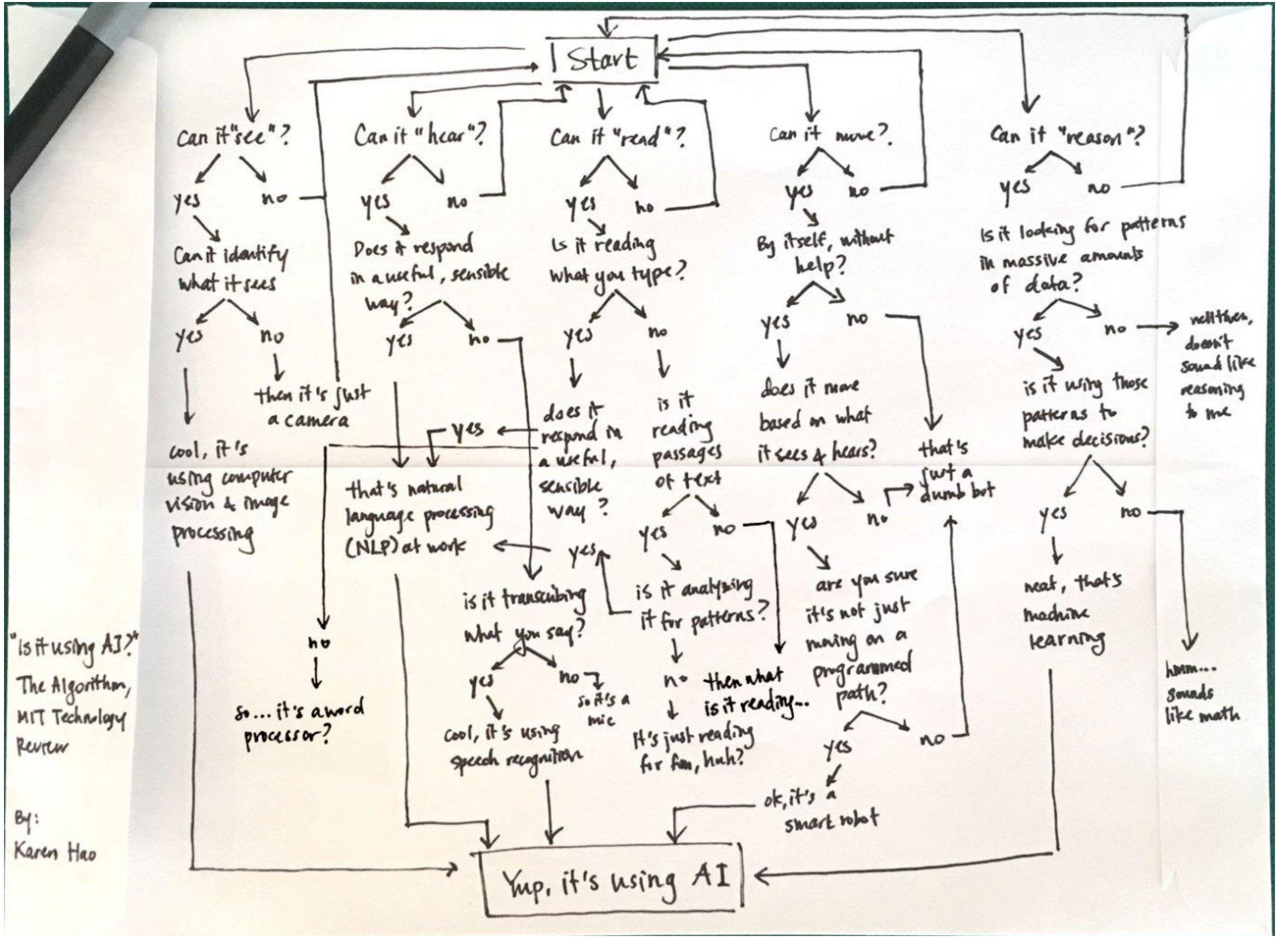

https://www.technologyreview.com/2018/11/10/139137/is-this-ai-we-drew-you-a-flowchart-to-work-it-out/

What is machine learning?

- All useful programs "learn" something
- Learning is improving at a task (set of skills) with experience
    - Experience = Data
    - The more data the better
    - However, more data requires more computational power and storage

"Field of study that gives computers the ability to learn without being explicitly programmed." ~Arthur Samuel (1959)

ANSTO

What is machine learning?

Conventional programming

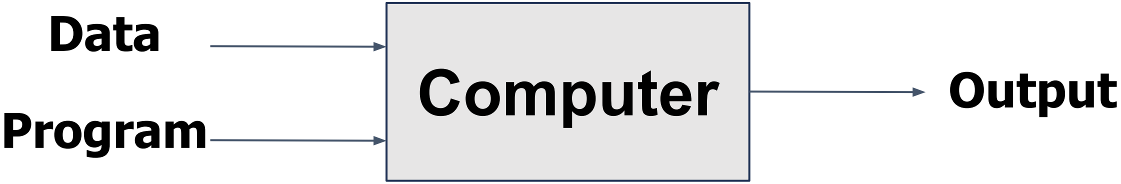

Example: Newton-Raphson method

$$x_1 = x_0 - f(x_0)/f'(x_0)$$

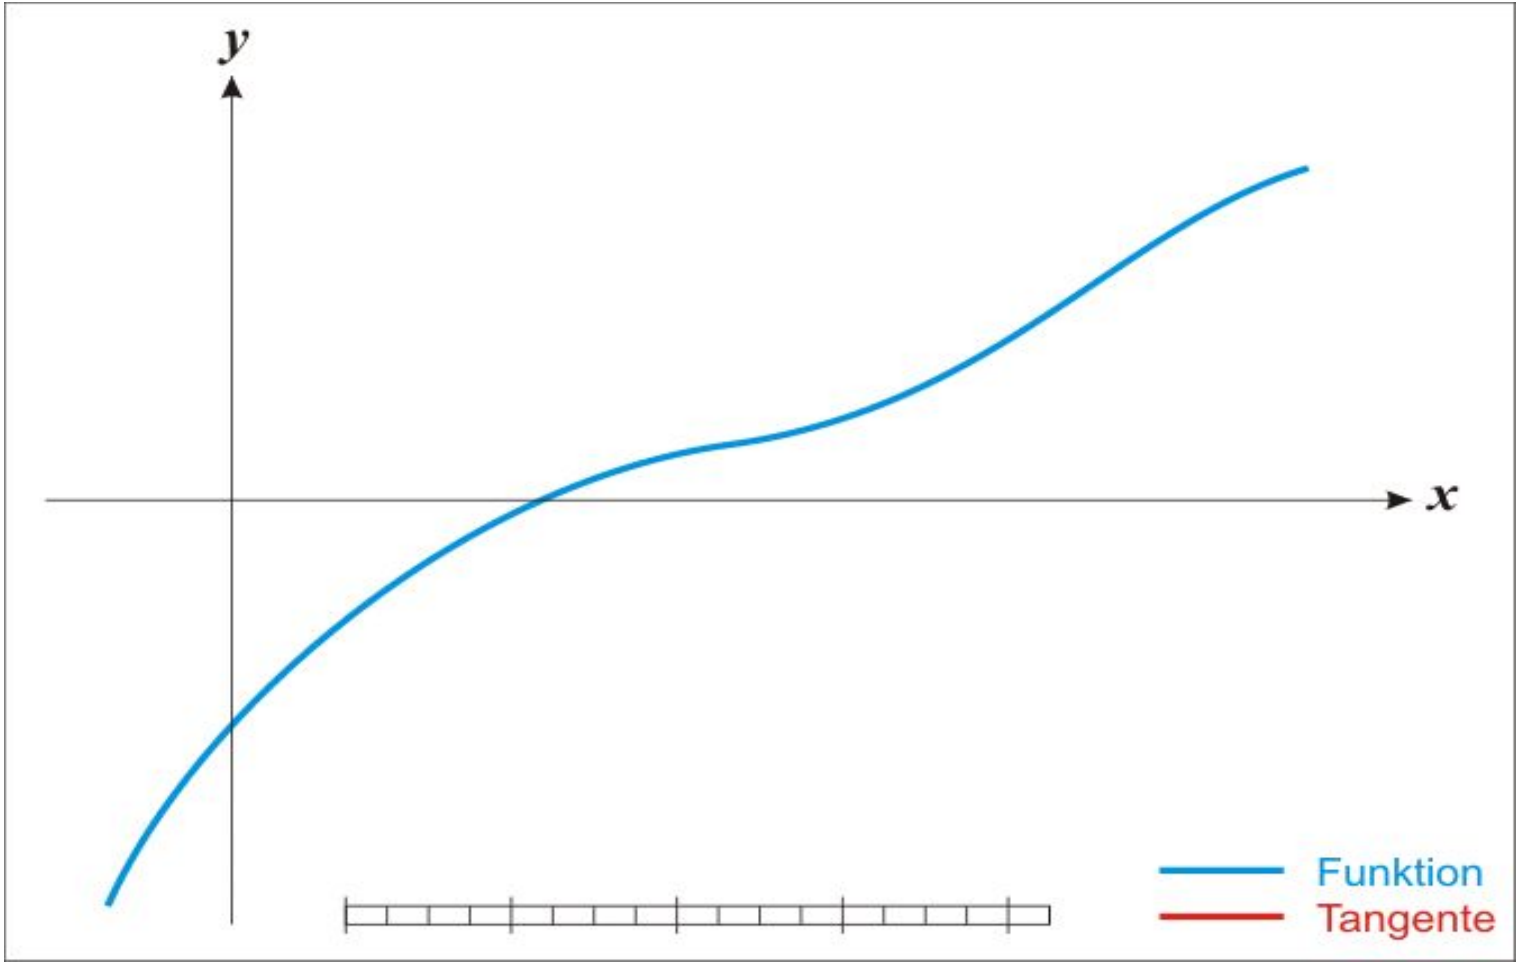
y
x
Funktion
Tangente

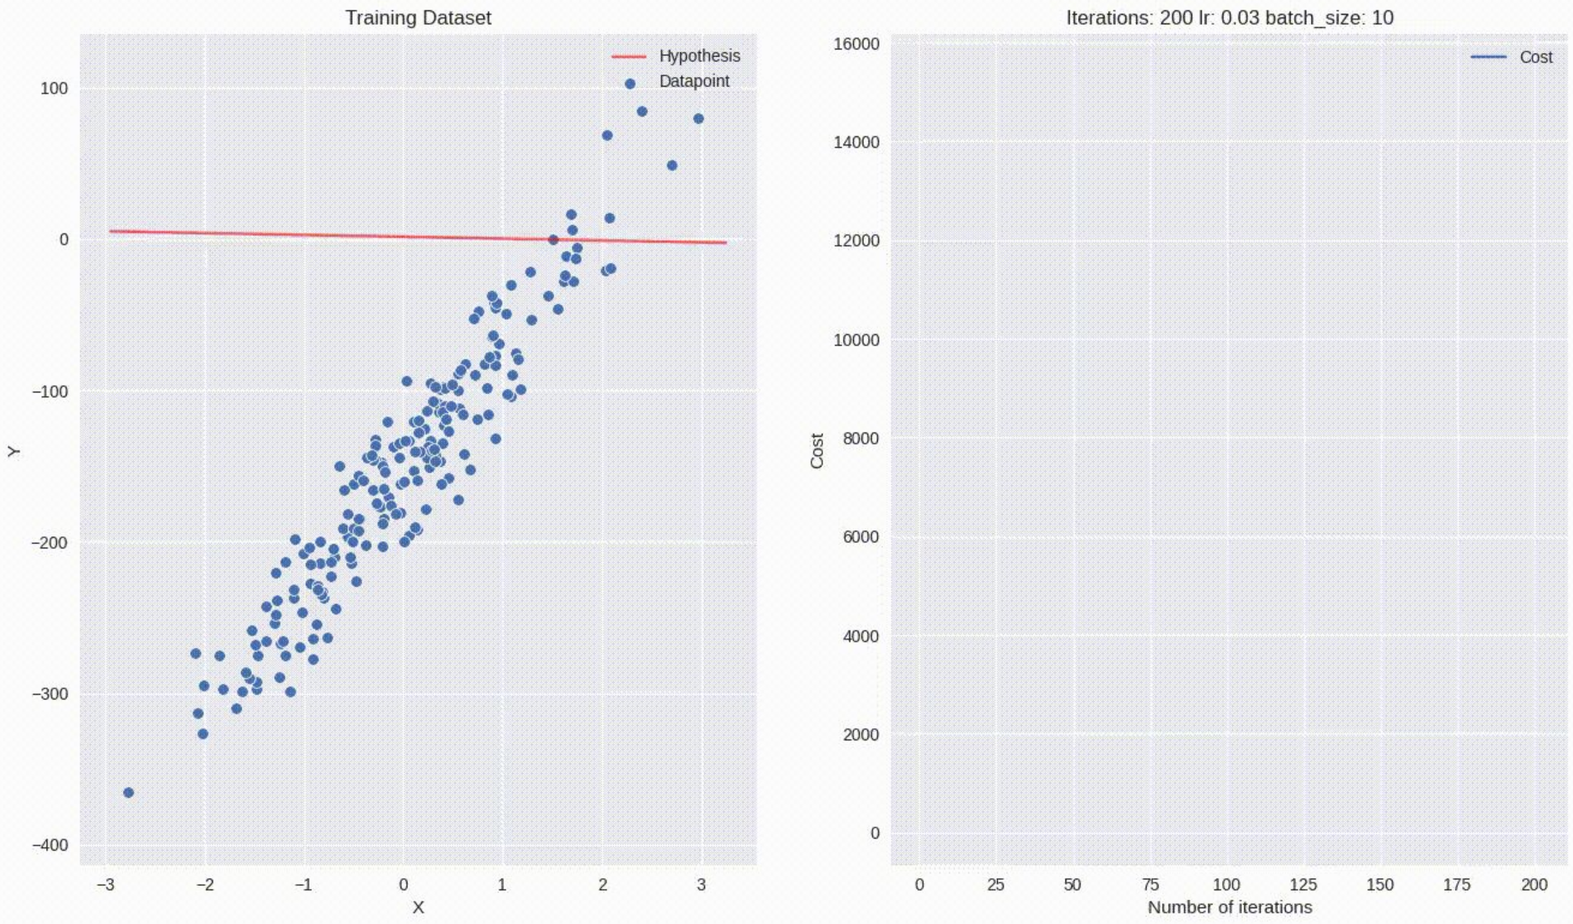

Source: https://github.com/Gautam-J/Machine-Learning

What is machine learning?

Conventional programming

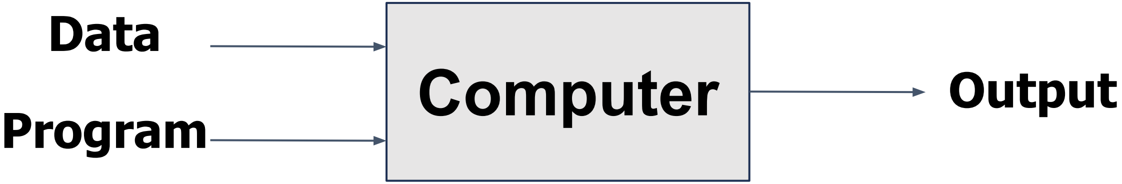

Machine Learning

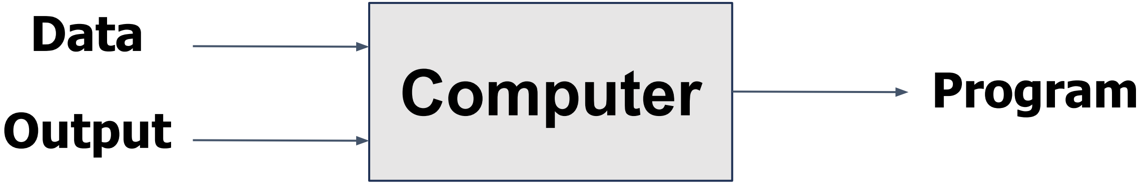

What is machine learning?

Conventional programming

Data

Program

Computer

Output

Machine Learning

Data

Output

Computer

Program

Basic paradigm

- Observe a set of examples:
  training data
- Infer something about the underlying model:
  inference engine
- Use inference to make predictions about unseen data:
  test data

Training dataset → Inference Engine → Test dataset

Supervised vs unsupervised learning

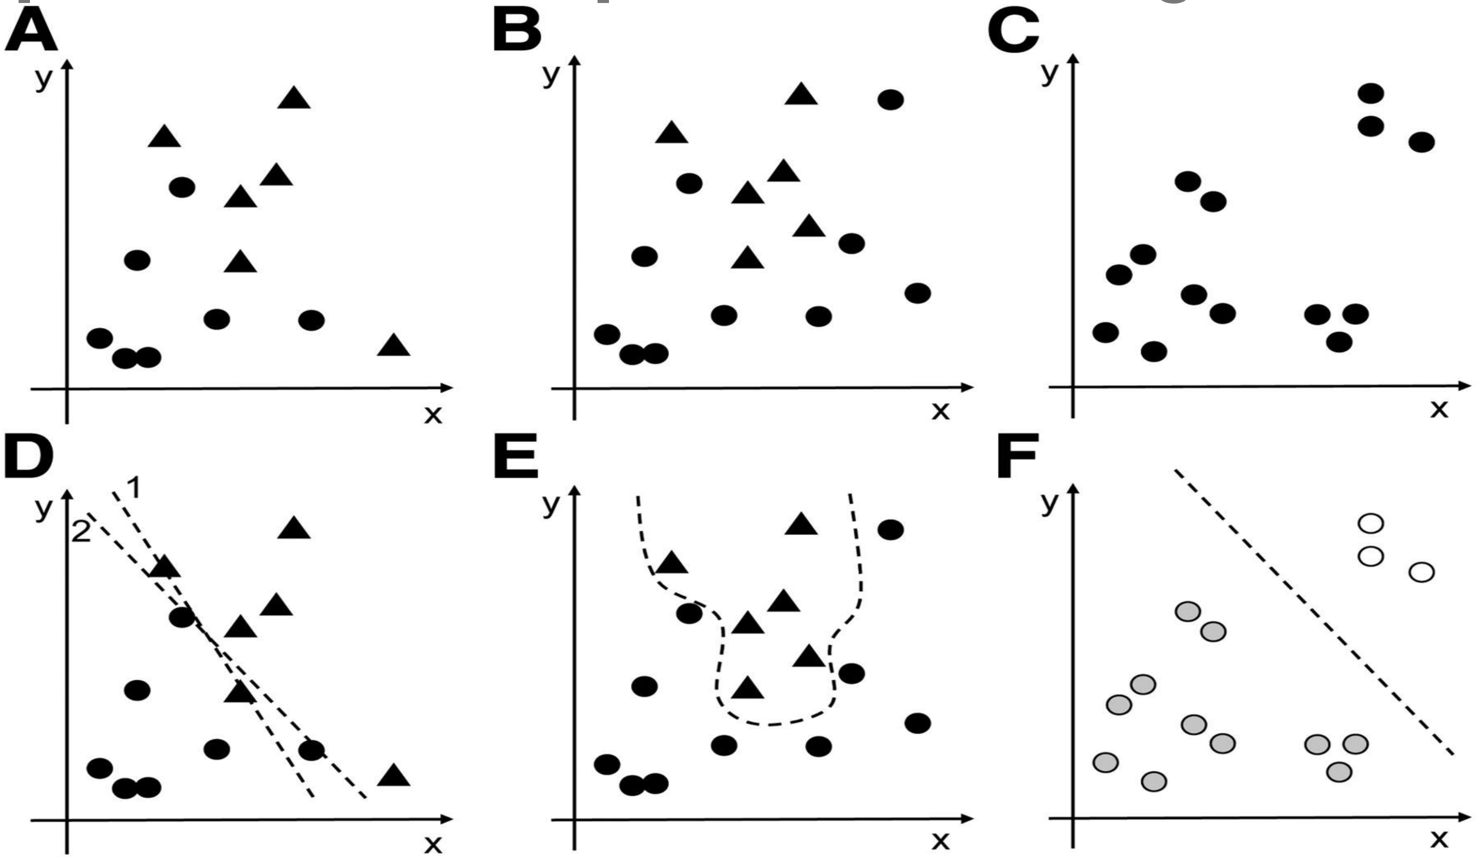

Uribe et al., 2019 "Machine Learning in Nuclear Medicine: Part 1—Introduction" JNM, 60 (4) 451-458;
DOI: *10.2967/jnumed.118.223495*

Supervised, semi-supervised, unsupervised learning

E. Montagnon et al., 2020 "Deep learning workflow in radiology: a primer". Insights Imaging 11, 22 (2020).
https://doi.org/10.1186/s13244-019-0832-5

Example: segmentation and quantification

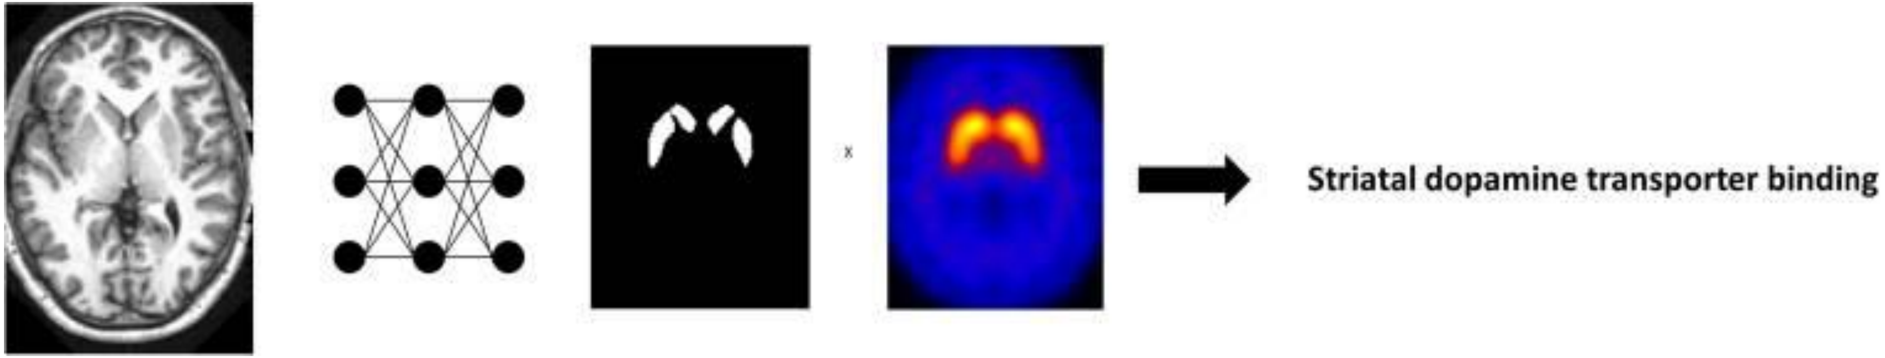

E. Montagnon et al., 2020 “Deep learning workflow in radiology: a primer”. Insights Imaging 11, 22 (2020).
https://doi.org/10.1186/s13244-019-0832-5

Status: what works?

- New Image reconstruction methods
- Quality assessment algorithms
- Triaging of cases to find similar patient data and diagnostic classification
- Lesion and disease detection and classification
- Segmentation, identification and feature extraction (radiomics) and quantification

What are the challenges? What is on the horizon?

- Diversification of learning for prognosis and predictive information.
- Building trust and reliability in the models
- Labelling the unlabelled image-data using unsupervised learning.
- Fairness and bias mitigation and understanding the limits of our algorithms.

Garbage in - garbage out!

Learning

- Declarative knowledge (memorising)
  - Accumulation of individual facts
  - Limitations: Storage capacity and observation time

- Imperative knowledge (generalisation)
  - Deduce new facts from old facts
  - Limitation: accuracy of deduction process (assuming that past can predict the future)

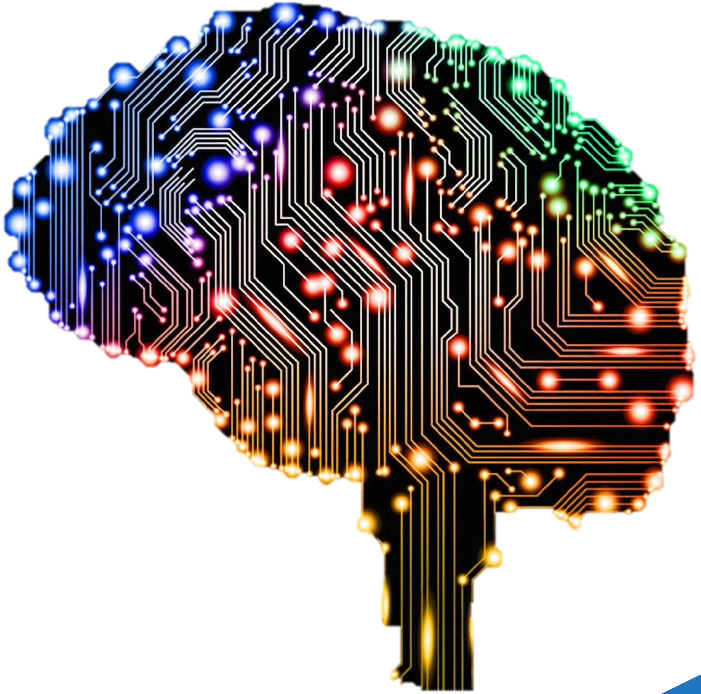

Demonstration: Brain tumour segmentation

- This is an example from Matlab, using the Deep Learning toolbox.
- The first and most important part of any ML solution is training data
- This example uses segmented MRI brain images from the BraTS data set from the Medical Segmentation decathlon (http://medicaldecathlon.com/)

ANSTO

Training data

- The data set includes 750 MRI images of brains with tumours, 484 of which have been manually labelled as **tumour** or **background**. These are what we will use.
- Image dimensions: **240 (h) x 240 (w) x 155 (d) x 4** (scan modality)
- Labelled images are divided into three subsets:
  - **Training** (400) - to train the neural network
  - **Validation** (29) - to monitor over/underfitting during training
  - **Testing** (55) - to evaluate the performance of the trained network

MRI image sample

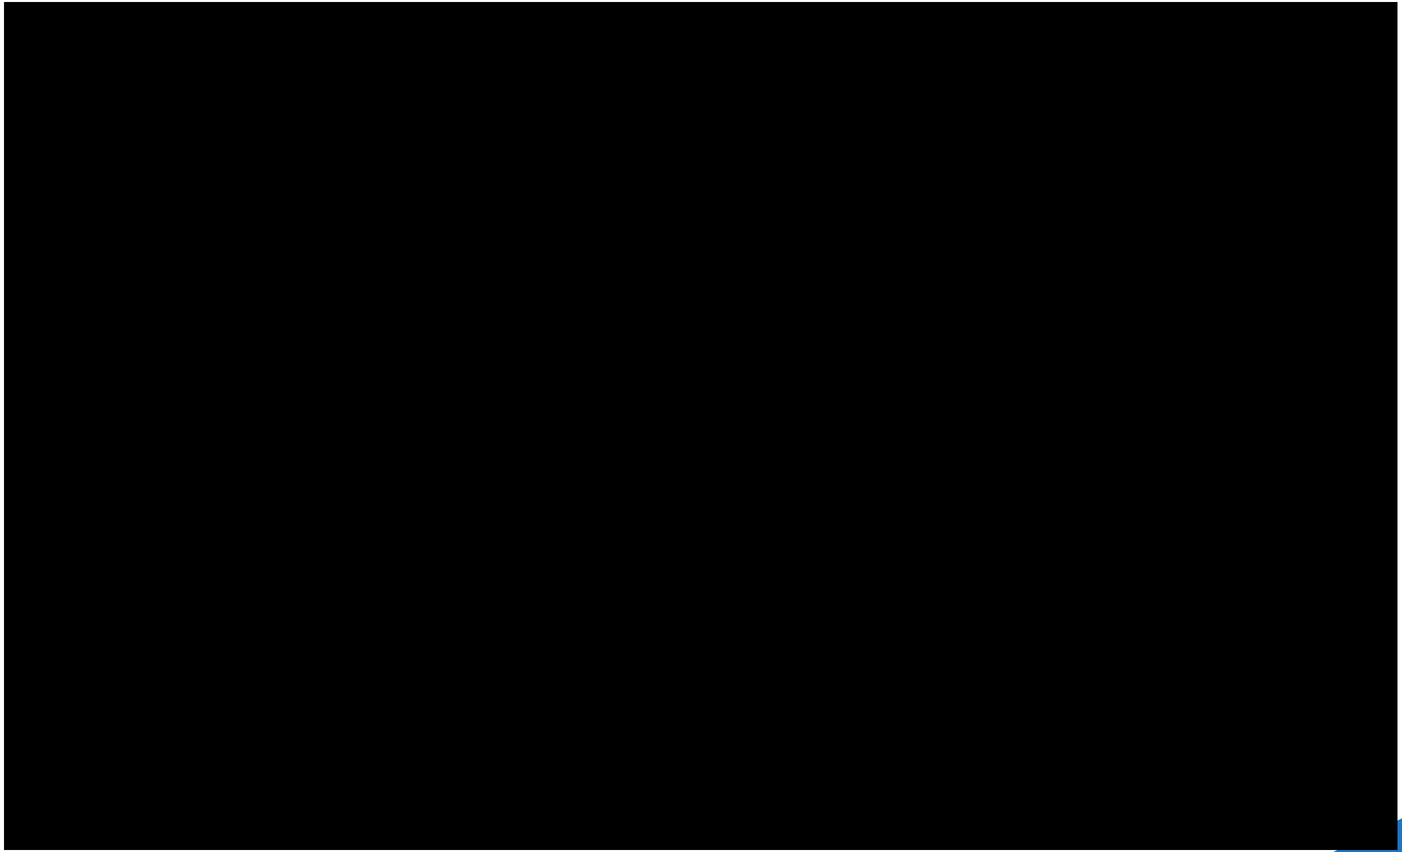

Preprocessing and preparation of training data

- Volume and label images are converted from NIFTI to Matlab binary format
- To reduce the memory and processing burden, the images are cropped to remove non-brain regions
- 16 132x132x132 subsets (‘patches’) of each image are randomly selected as training samples
- Patches are **normalised** - so we don’t have to do this later during training (faster to just do it once)

ANSTO

Constructing the network

- We will use U-Net, an encoder/decoder convolutional neural network (CNN) which is widely used for image segmentation
- A default 3D U-Net network is constructed; input dimension is the patch size (132x132x132x4) and output dimension is 44x44x44x2 (due to the use of 'valid' style padding on each patch - no zero-padding is used)
- The input and output layers are modified from defaults
    - Input layer does not need to do normalisation (already done during preprocessing)
    - Output layer is changed to a dicePixelClassificationLayer for better segmentation of small tumour regions

Training hyperparameters

- The main training hyperparameters of interest (and suggested values):
  - Use the **adam** optimiser (popular choice for segmentation)
  - Specify which data you are using for validation, and how often to validate (e.g. every 400 iterations)
  - Maximum number of training epochs (passes through all training data - e.g. 50)
  - Initial learning rate (0.0005), mode of decrease (‘piecewise’) and rate of decrease (x0.95 every 5 epochs)
  - MiniBatchSize - make larger if you have more GPU memory (but you might be better off with larger patch size)

ANSTO

Training

- This is the easy part - one line of code!

```
[net, info] = trainNetwork (dsTrain, lgraph, options);
```

- **lgraph** is the network, **dsTrain** is our training datastore and **options** are our training options
- Important: save the resulting network **net** for future use - it will be about 60 MB
- Takes roughly 30 hours to train on a powerful GPU/CPU system - *much* slower without a GPU

Applying the trained network

- Since the input size of the network is smaller than the dimension of the training data, we step through the input image size with a stride equal to the output patch size, and extract patches equal to the input patch size.
- Each patch is segmented:

```
patchSeg = semanticseg (patch, net);
```

- The resulting segmented patches are then stitched back together and the final image is cropped to the original size

ANSTO

Ground truth (L) and segmented (R)

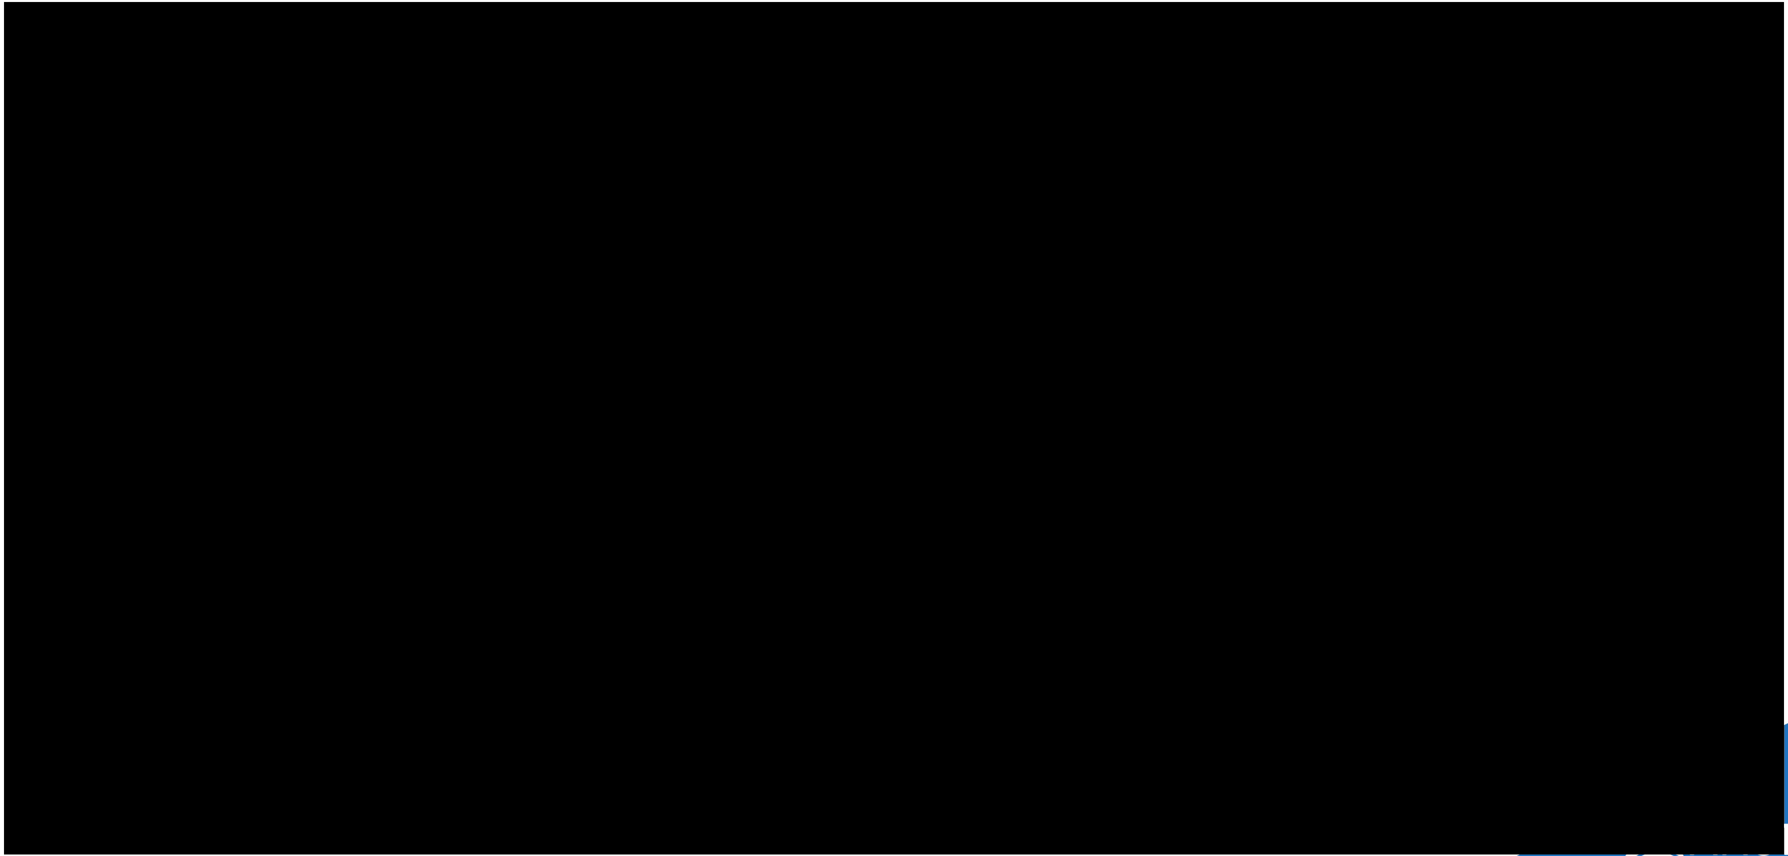

Performance

- The network performs very well. Results for testing data show that:
  - ~98% of the network-classified voxels agree with the ground truth (true positive + true negative)
  - Jaccard score is > 95% for all evaluated images
  - Dice similarity score > 99.9% for background and > 95% for tumour
- The approach can be generalised to other image segmentation problems

Code

Original example:

https://au.mathworks.com/help/vision/ug/segment-3d-brain-tumor-using-deep-learning.html

Commented code:

Thank you

Email Address: mitras@ansto.gov.au

ANSTO