

Medical Informatics

Lecture 13: Artificial Intelligence in Medicine

Dr Areti Manataki



Nanjing Medical University

Acknowledgements

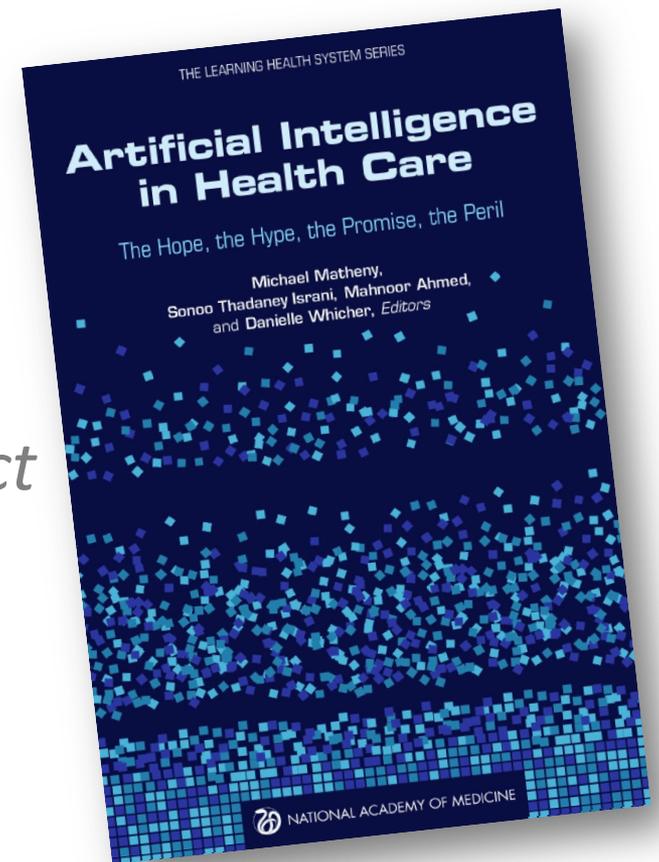
Some of the content of this lecture is based on the Data Science in Stratified Healthcare and Precision Medicine course on Coursera, which was developed by The University of Edinburgh under a Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0) license.

These lecture slides are also licensed under a CC BY-SA 4.0 license.

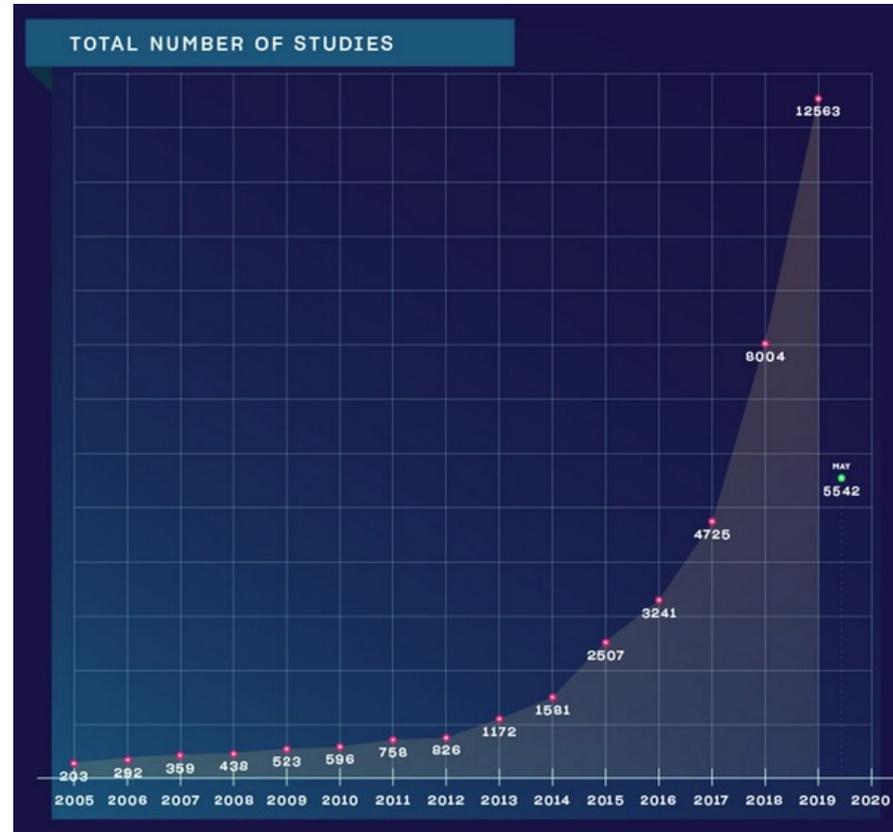


AI in Medicine: the big promise

*“The emergence of artificial intelligence in health care offers **unprecedented opportunities** to improve patient and clinical team outcomes, reduce costs, and impact population health.”*



Growing volume of medical studies using Machine Learning / AI

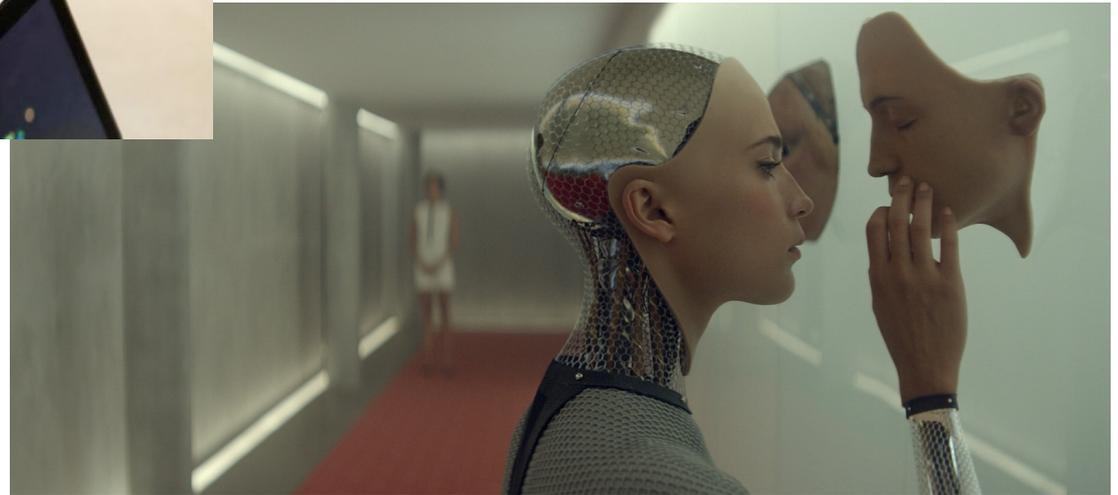
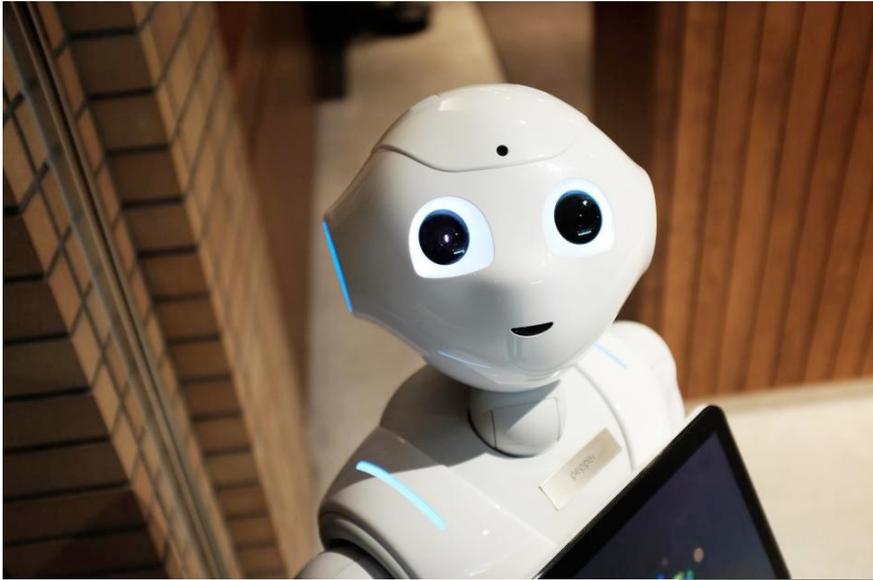


Number of medical AI studies on PubMed from 2010 to 2020

Image adapted from Meskó and Görög (2020) <https://doi.org/10.1038/s41746-020-00333-z>

What is AI?

What is Artificial Intelligence?



[Image by BagoGames on flickr, licensed as CC BY 2.0]

What is Artificial Intelligence?

- Systems that think like humans.
- Systems that act like humans.
- Systems that think rationally.
- Systems that act rationally.

*“Artificial Intelligence: A Modern Approach”
by Russel and Norvig*

What is Artificial Intelligence?

- Systems that think like humans.
- Systems that act like humans.
- Systems that think rationally.
- Systems that act rationally.

*“Artificial Intelligence: A Modern Approach”
by Russel and Norvig*

Two schools of thought:

- Rule-based
- Machine Learning

Machine Learning



BBC Sign in News Sport Weather iPlayer Sounds

NEWS

Home UK World Business Politics Tech Science Health Family & Education

Technology

Google AI defeats human Go champion

© 25 May 2017

f t e Share



FINANCIAL TIMES

UK COMPANIES TECH MARKETS GRAPHICS OPINION WORK & CAREERS LIFE & ARTS HOW TO SPEND IT

Driverless vehicles + Add to myFT

Self-driving cars raise fears over 'weaponisation'

Baidu warns global carmakers need to address security concerns

Machine Learning

- Machine learning algorithms allow computers to **identify patterns** in data, build models that **explain** the world and **make predictions** *without having explicit pre-programmed rules and models.*

Input x

2

3

5

10

8

Output y

5

7

11

21

?

Machine Learning

- Machine learning algorithms allow computers to **identify patterns** in data, build models that **explain** the world and **make predictions** *without having explicit pre-programmed rules and models.*

Input x

2

3

5

10

8

Output y

5

7

11

21

17

Machine Learning

- Machine learning algorithms allow computers to **identify patterns** in data, build models that **explain** the world and **make predictions** *without having explicit pre-programmed rules and models*.

<u>Input x</u>		<u>Output y</u>
2	$y = 2x + 1$	5
3		7
5		11
10		21
8		17

Machine Learning

- Machine learning algorithms allow computers to **identify patterns** in data, build models that **explain** the world and **make predictions** *without having explicit pre-programmed rules and models*.

<u>Input x</u>		<u>Output y</u>
2	$y = 2x + 1$	5
3		7
5		11
10		21
8		17

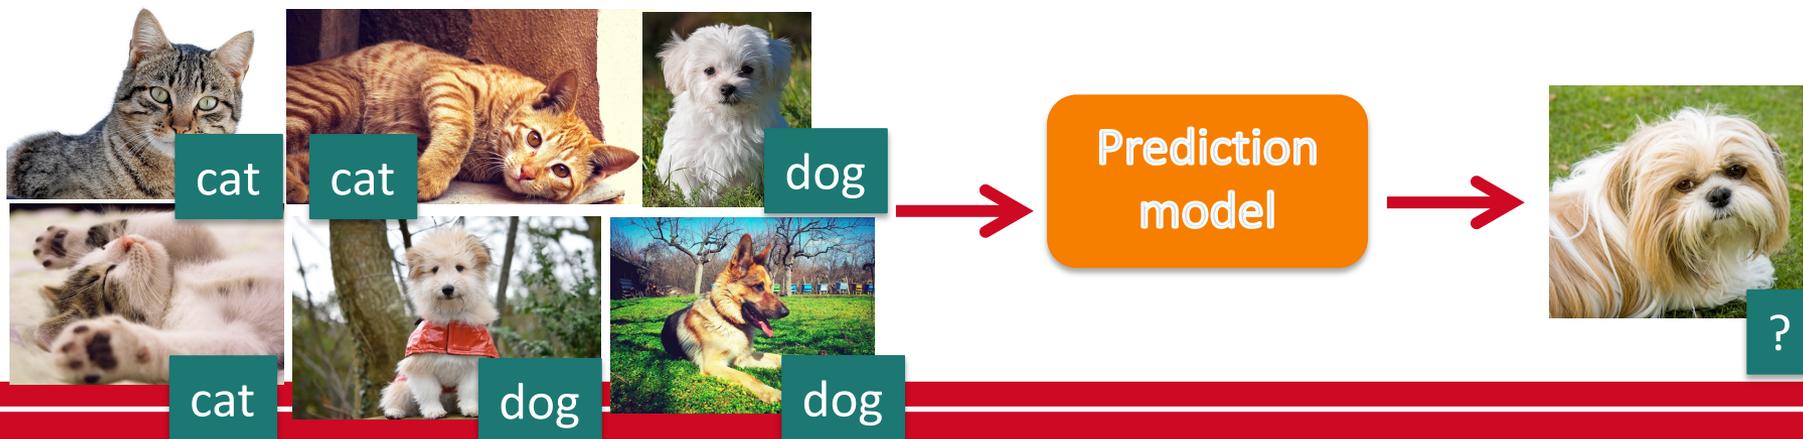
Machine Learning
can discover this

Supervised vs. Unsupervised Learning

- **Supervised learning:** We have a set of labelled training data, where we have pairs of inputs and outputs. The goal here is to build a prediction model, so that we can predict the output even for inputs never seen before.
 - Example: classifying tumours

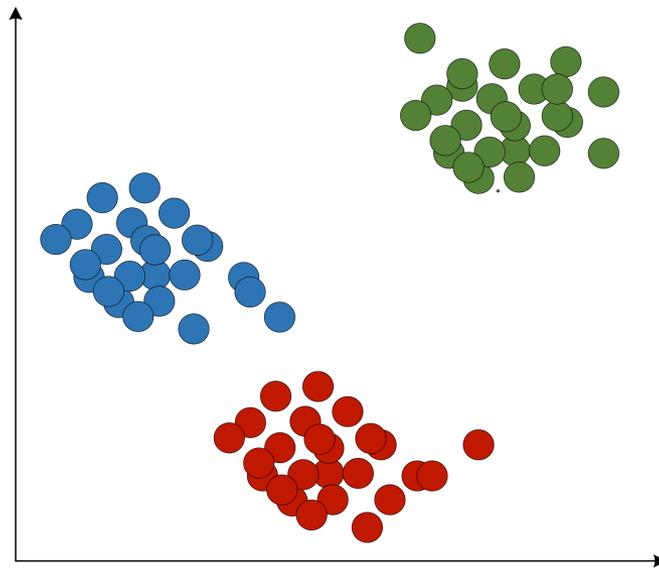
Supervised vs. Unsupervised Learning

- **Supervised learning:** We have a set of labelled training data, where we have pairs of inputs and outputs. The goal here is to build a prediction model, so that we can predict the output even for inputs never seen before.
 - Example: classifying tumours



Supervised vs. Unsupervised Learning

- **Unsupervised learning:** The dataset is unlabelled, and the objective is to extract knowledge from the data, for instance by discovering hidden patterns.
 - Example: identifying different clusters of cancer patients



*[Image from DataSciMed
MOOC]*

Two types of supervised learning

- **Classification**: we predict results in a discrete output (i.e. one or more classes)
 - Examples: predicting whether a tumour is benign or malignant, predicting cardiology diagnosis
- **Regression**: we predict results within a continuous numerical output
 - Examples: predicting a patient's life expectancy, predicting the waiting time at the A&E department

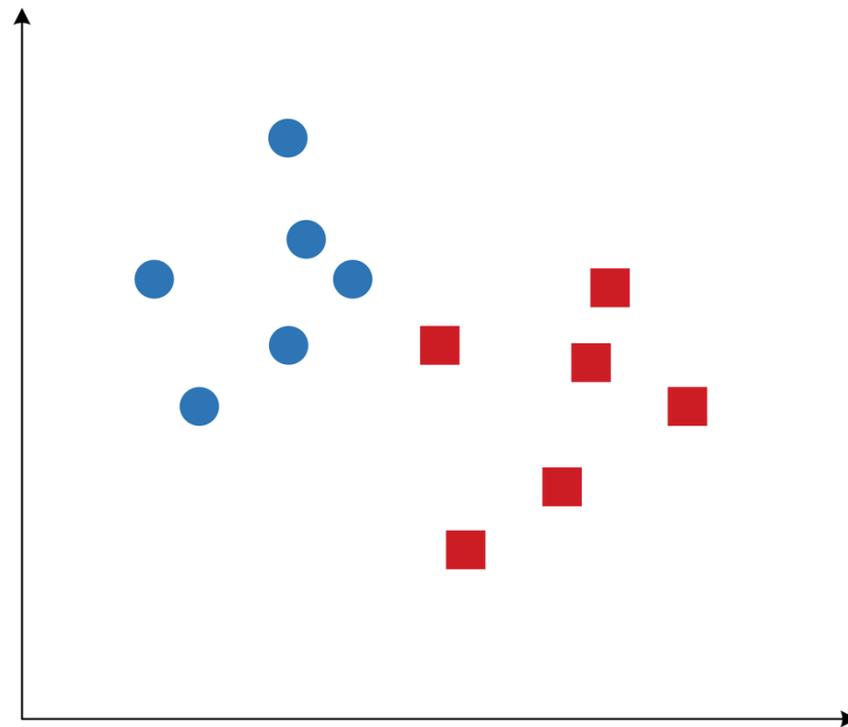
Machine Learning techniques

Machine learning techniques

- K-nearest neighbours
- Decision trees
- Random forests
- Neural networks
- Deep neural networks
- Support vector machines
- Linear regression
- Logistic regression

K-nearest neighbours

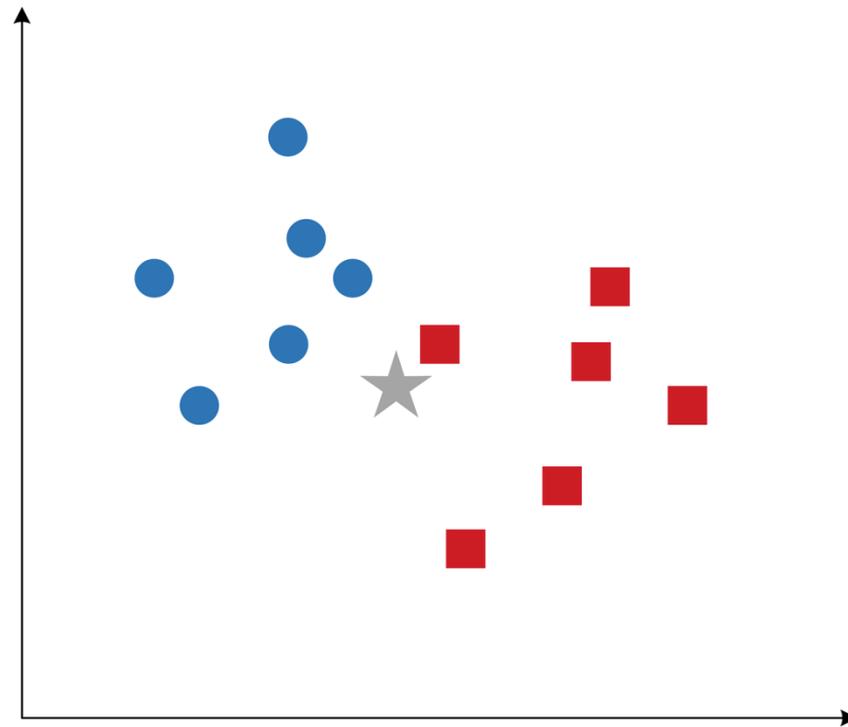
- Main idea: the prediction for a new datapoint is the mode or mean of the values of its k closest datapoints (i.e. its k nearest neighbours)



*[Image from DataSciMed
MOOC]*

K-nearest neighbours

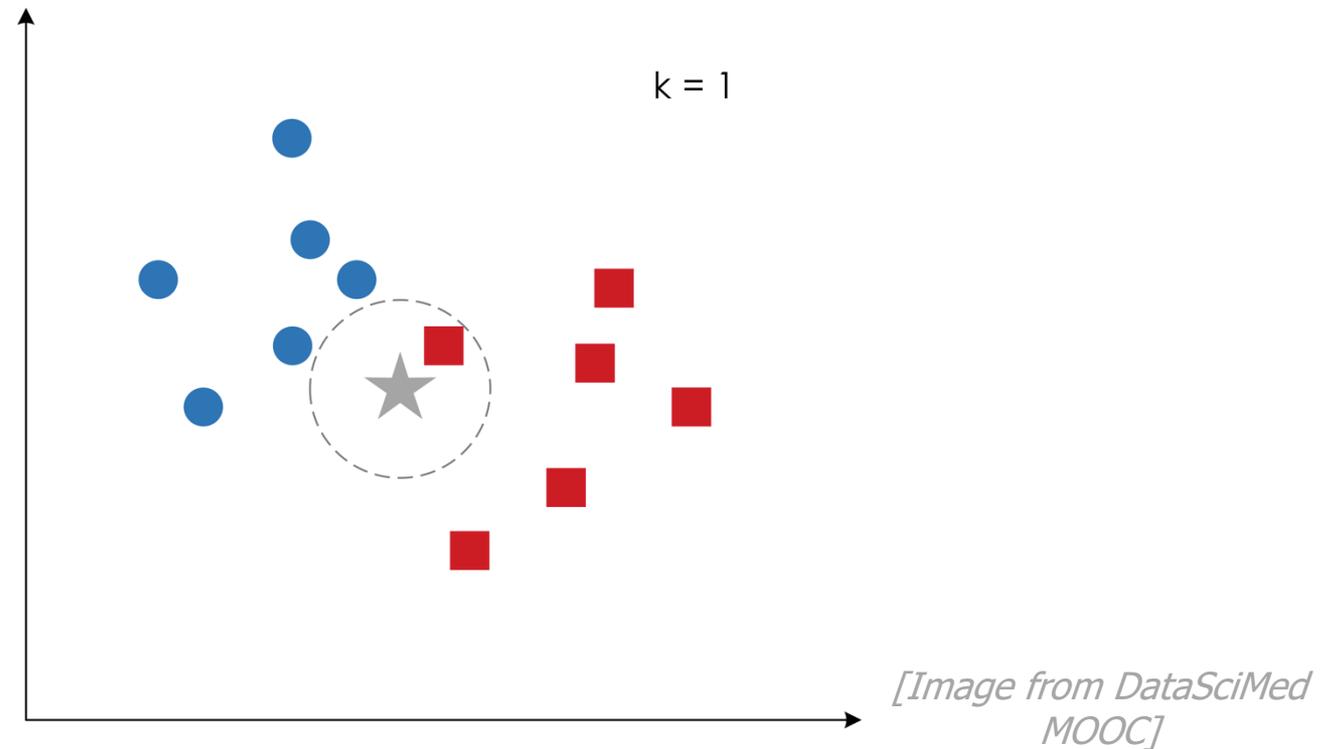
- Main idea: the prediction for a new datapoint is the mode or mean of the values of its k closest datapoints (i.e. its k nearest neighbours)



*[Image from DataSciMed
MOOC]*

K-nearest neighbours

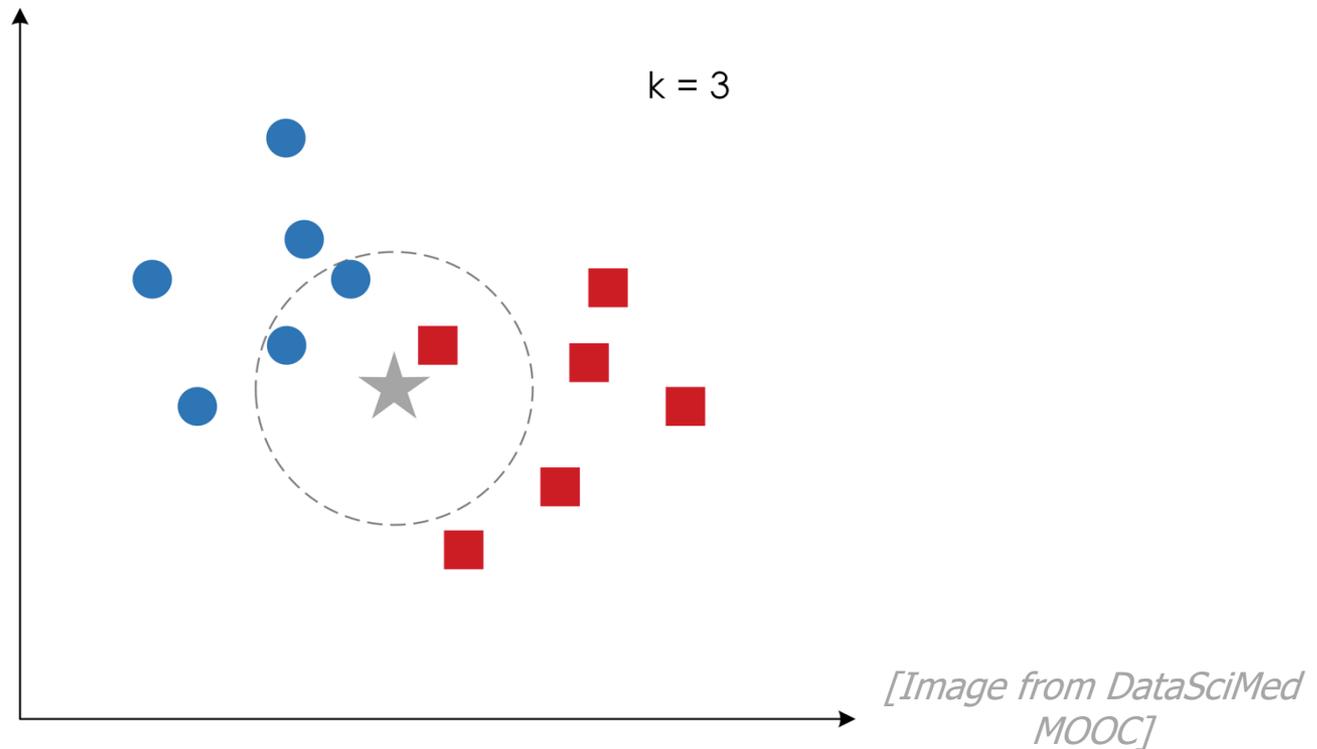
- Main idea: the prediction for a new datapoint is the mode or mean of the values of its k closest datapoints (i.e. its k nearest neighbours)



*[Image from DataSciMed
MOOC]*

K-nearest neighbours

- Main idea: the prediction for a new datapoint is the mode or mean of the values of its k closest datapoints (i.e. its k nearest neighbours)



K-nearest neighbours

- Main idea: the prediction for a new datapoint is the mode or mean of the values of its k closest datapoints (i.e. its k nearest neighbours)
- The algorithm is easy to understand and often gives reasonable performance without a lot of adjustments.
- Building the model is fast, but when a large number of features is involved, prediction can be slow.

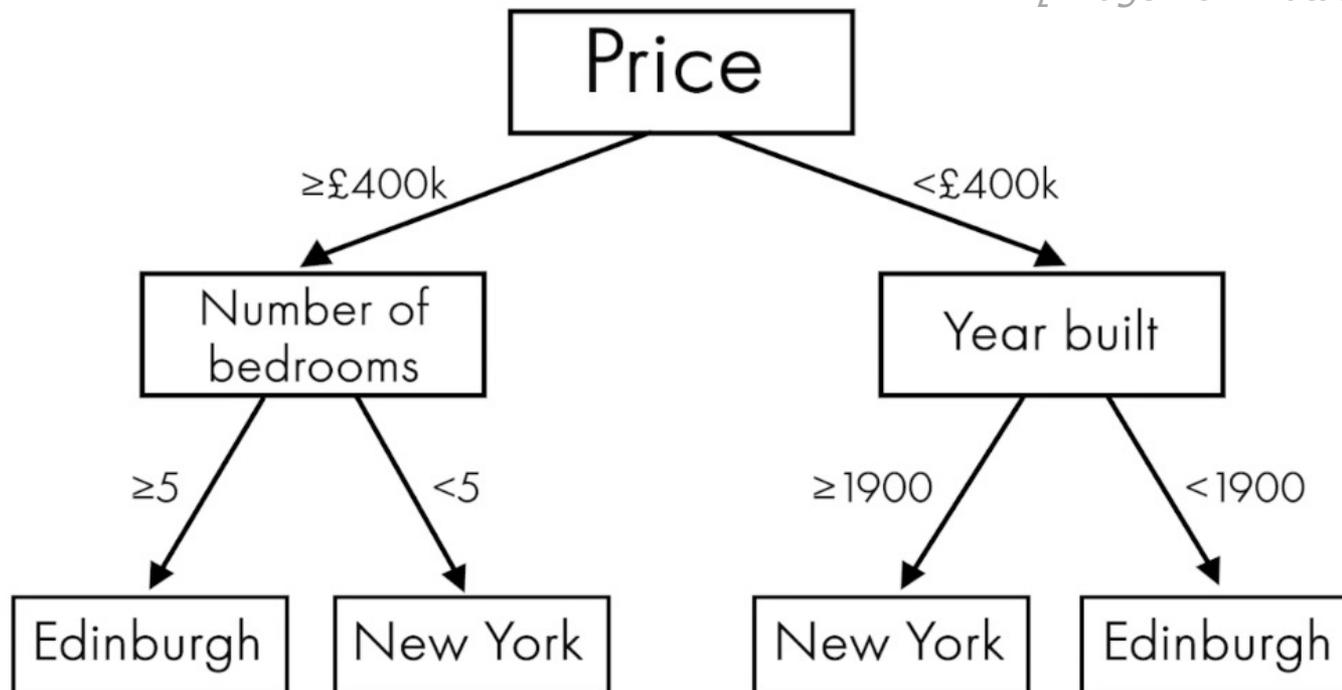
Decision trees

- Decision trees allow us to predict the value of a target variable based on a set of input variables.
- They have a **tree-like structure**, consisting of nodes and edges. Each node is associated with one of the input variables, while the edges coming from that node are the total possible values of that node.

Decision trees

- Decision trees allow us to predict the value of a target variable based on a set of input variables.

[Image from DataSciMed MOOC]



Decision trees

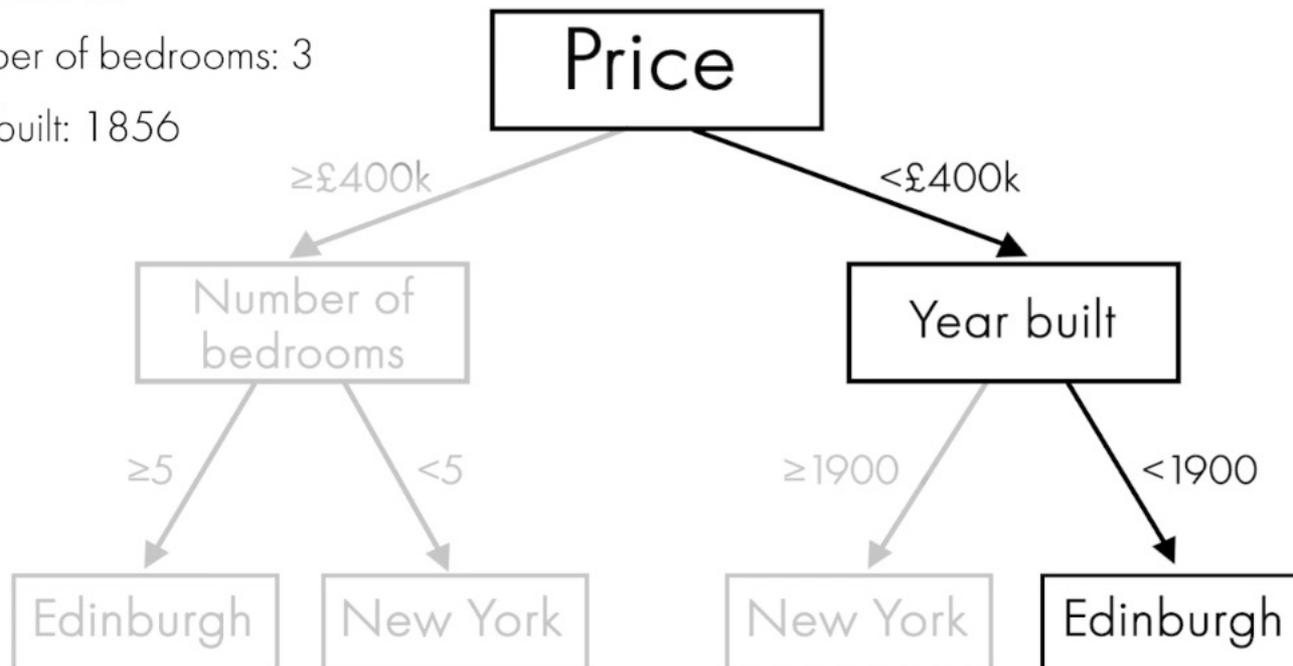
- Decision trees allow us to predict the value of a target variable based on a set of input variables.

[Image from DataSciMed MOOC]

Price: £280k

Number of bedrooms: 3

Year built: 1856

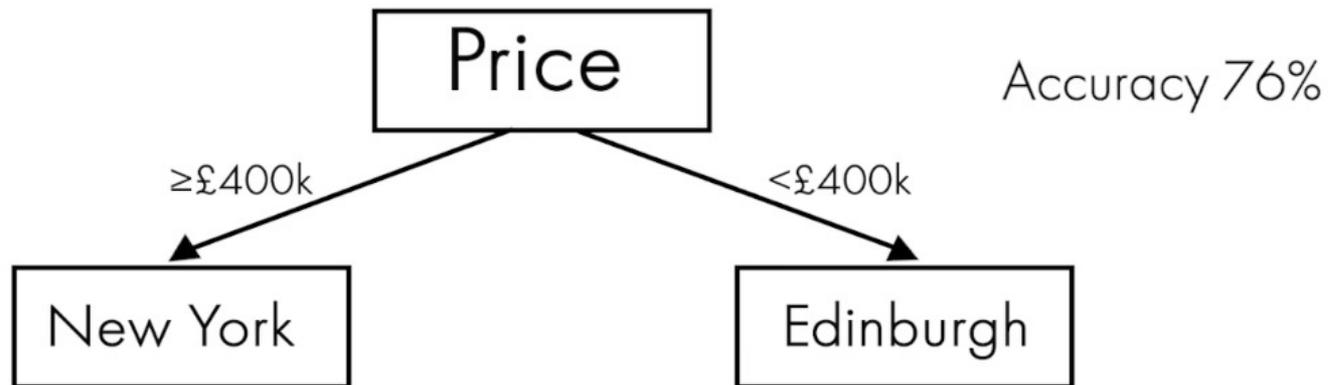


Decision trees

- Learning a decision tree is about learning the sequence of questions (i.e. sequence of variables and their split of possible values) that gets us to the correct answer more quickly.
 - Year built
 - Elevation
 - Square Metres
 - Price
 - Airconditioning
 - Number of bedrooms

Decision trees

- Learning a decision tree is about learning the sequence of questions (i.e. sequence of variables and their split of possible values) that gets us to the correct answer more quickly.

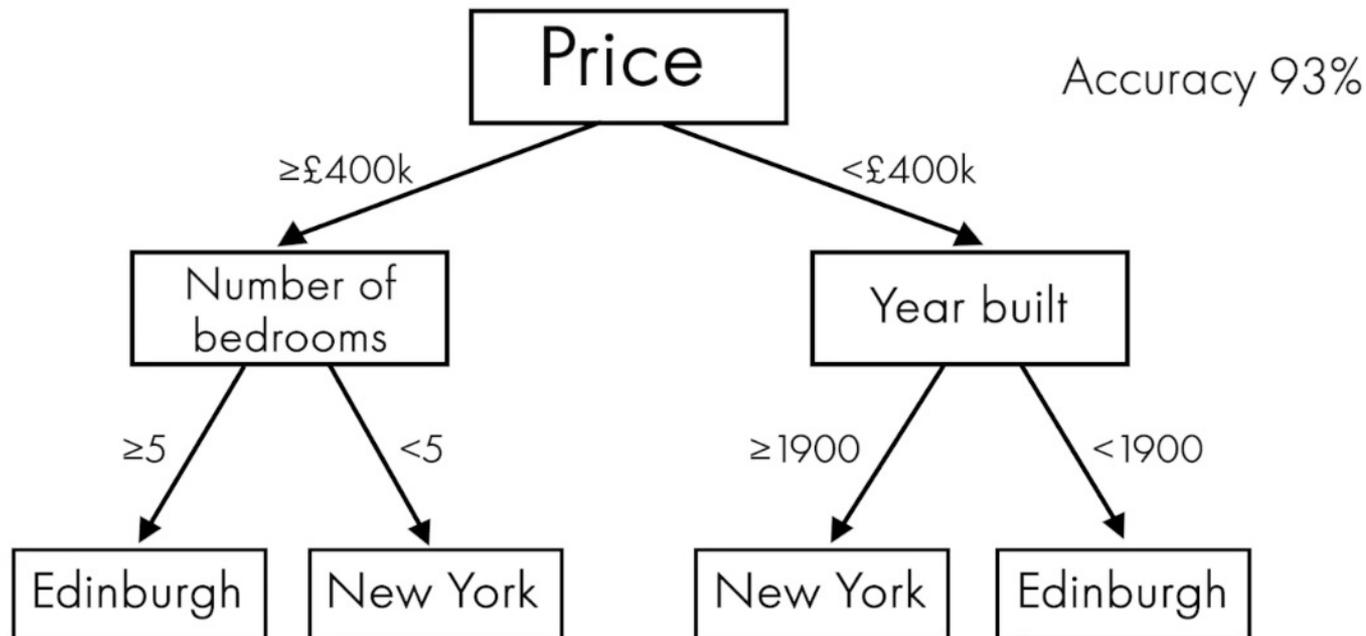


[Image from DataSciMed MOOC]

Decision trees

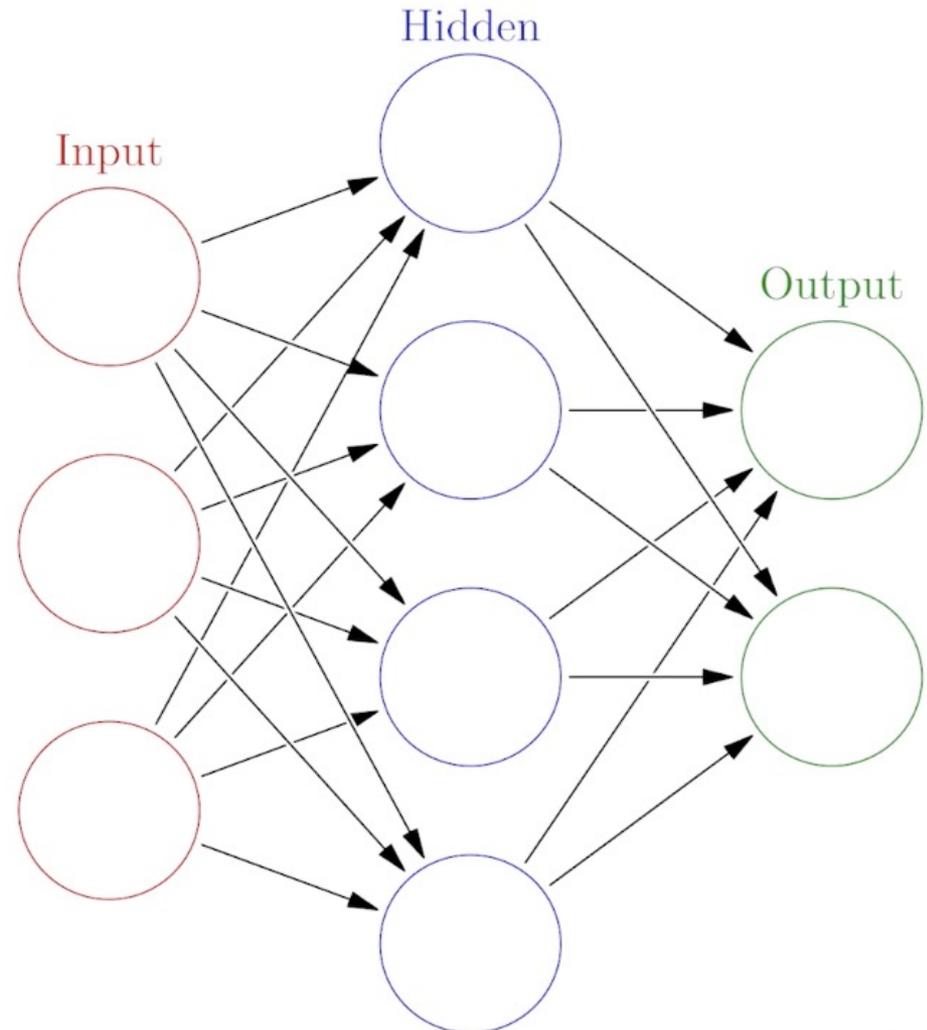
- Learning a decision tree is about learning the sequence of questions (i.e. sequence of variables and their split of possible values) that gets us to the correct answer more quickly.

[Image from DataSciMed MOOC]



Neural networks

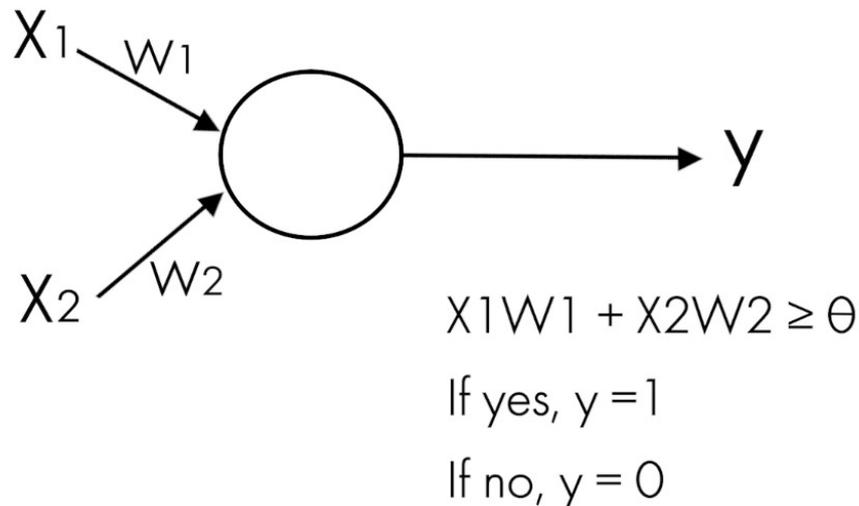
- A neural network is an interconnected group of nodes, called neurons.
- Neurons transmit signals to one another.
- Organized in layers: an input layer, an output layer, and one or more hidden layers



[Image by Glosser.ca on Wikimedia, licensed as CC BY –SA 3.0]

Neural networks

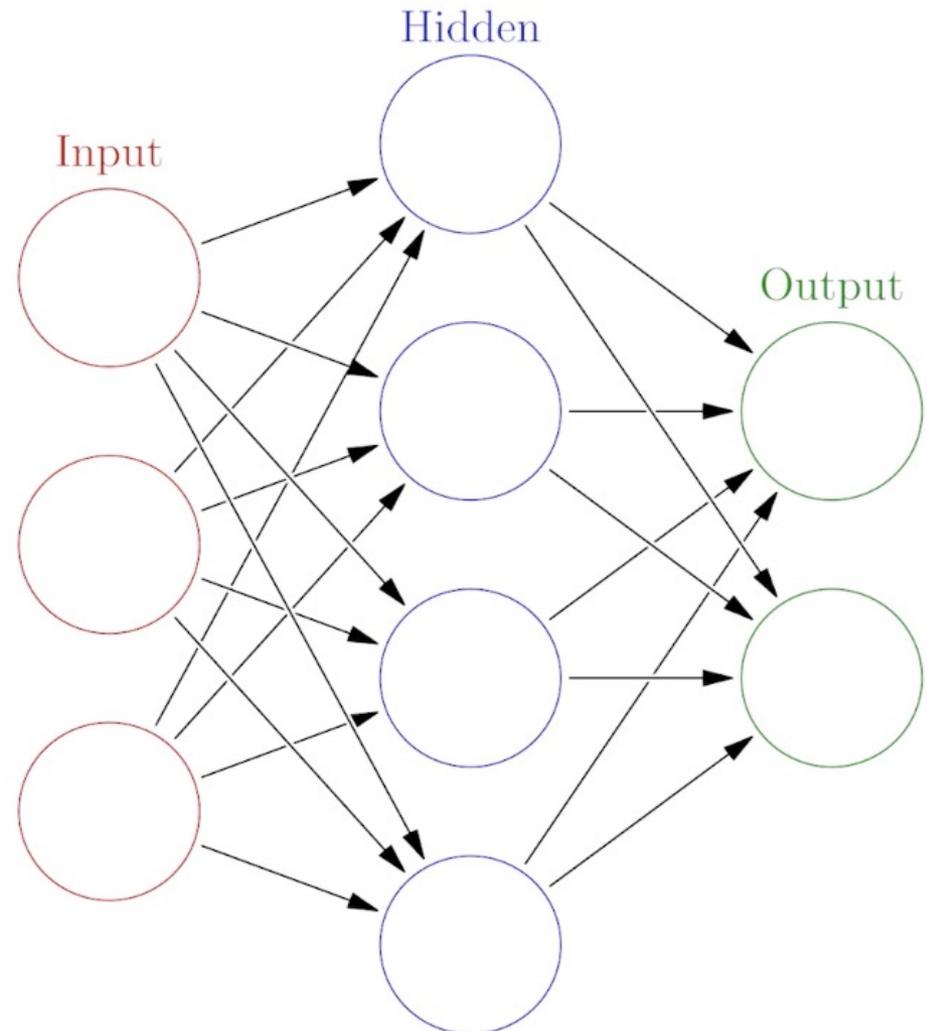
- Single neuron processing: A neuron calculates a weighted sum of its inputs, and then applies a threshold or activation function to determine whether the signal will be sent or not.



[Image from DataSciMed MOOC]

Neural networks

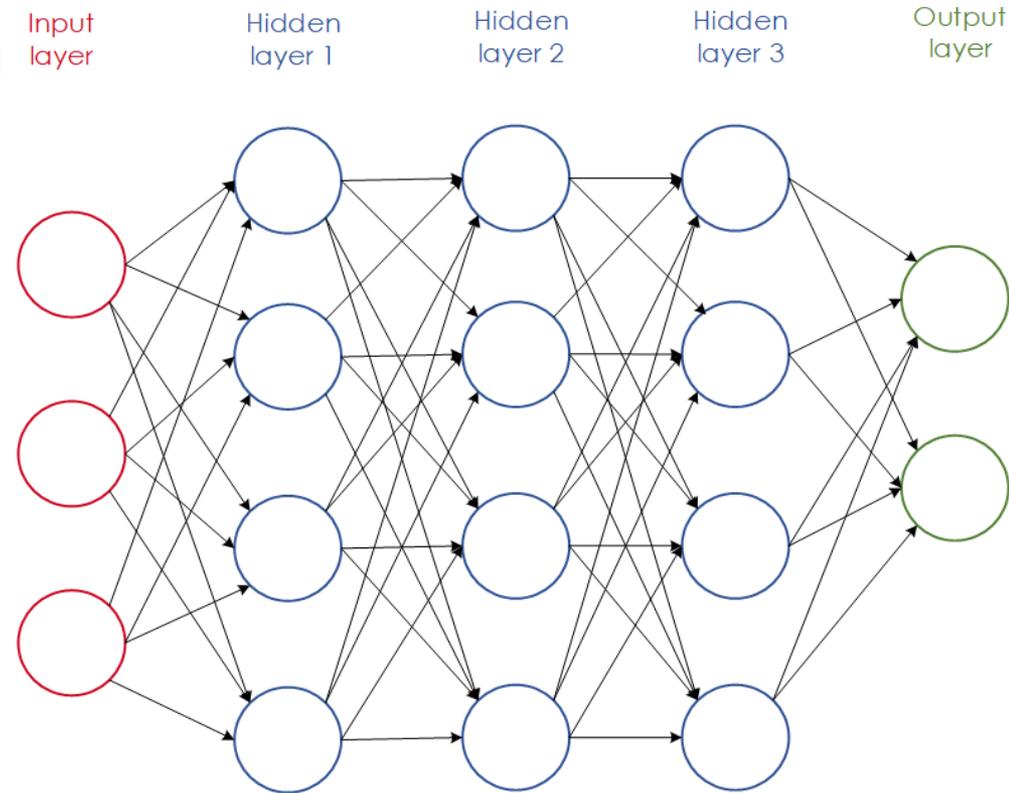
- In a neural network, this process is repeated several times: compute hidden units, and combine them to yield the final result.
- Purpose of neural networks in ML: learn the weights and the threshold functions with the use of existing data.



[Image by Glosser.ca on Wikimedia, licensed as CC BY -SA 3.0]

Deep neural networks

- Deep neural networks consist of several hidden layers.
- Deep learning has brought significant advances in several challenging domains, such as image analysis and natural language processing.
- Interpretability issues.

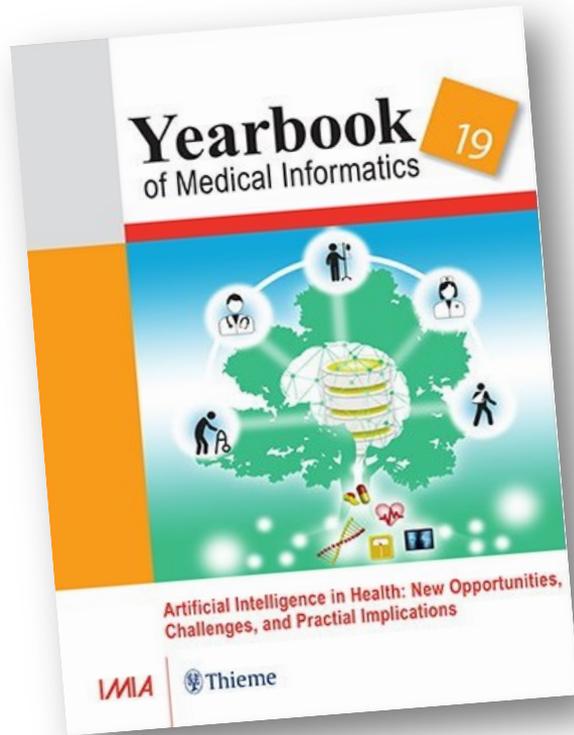


[Image from DataSciMed MOOC]

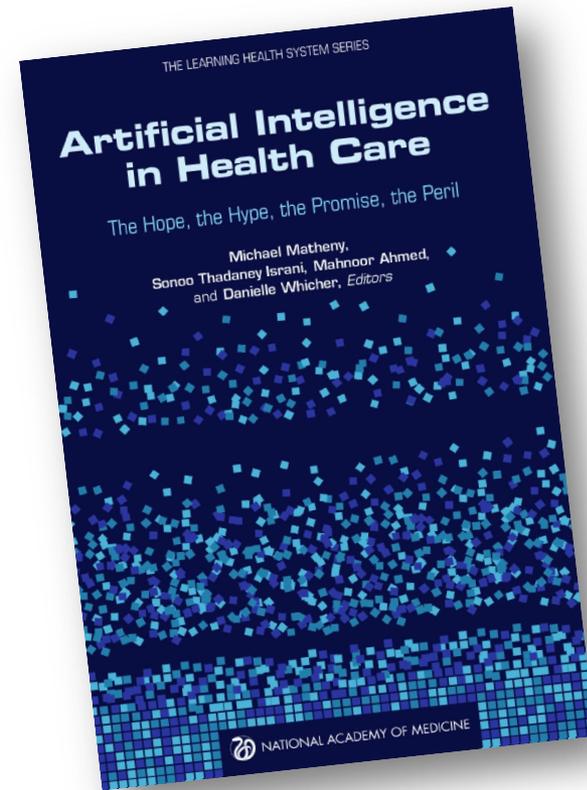
Reflecting on ML for AI

- Machine learning has been transforming the medical field in recent years.
 - “Clinically applicable deep learning for diagnosis and referral in retinal disease” by De Fauw et al.
 - “Highly accurate protein structure prediction with AlphaFold” by Jumper et al.
- Hype, unrealistic expectations
- Junk in, junk out
- Black box

Recommended readings



<https://imia-medinfo.org/wp/imia-yearbook-of-medical-informatics/>



<https://nam.edu/artificial-intelligence-special-publication/>

Conclusions

- ML is a subset of AI. There are other approaches too.
- Supervised vs. unsupervised learning
- There is a wide range of ML techniques.
- ML brings great opportunities to biomedicine and healthcare.
- Already some successful stories...but expectations can sometimes be unrealistically high, especially by the media.
- Nevertheless, an exciting area of research and innovation, with lots to offer.