L'intelligenza artificiale: cos'è e come possiamo accoglierla

Prof. **Daniela Besozzi**

Dept. Informatics, Systems and Communication
University of Milano-Bicocca

UNIVERSITÀ DEGLI STUDI DI MILANO BICOCCA

DIPARTIMENTO DI INFORMATICA SISTEMISTICA E COMUNICAZIONE

Outline

- **Artificial intelligence** vs **machine learning**
- Artificial intelligence in **ophtalmology**
- **What is (NOT)** artificial intelligence
- What it means (for a machine) to **learn**
- The **three main categories** of machine learning
- **Data, features, labels,** evaluation **metrics**
- **Bias** and **interpretability** issues
- To conclude

Si dichiara **assenza di conflitto di interessi** commerciali nel campo sanitario

Artificial intelligence (or machine learning?)

Natural Language Processing
Fuzzy Inference Systems
Evolutionary Computation

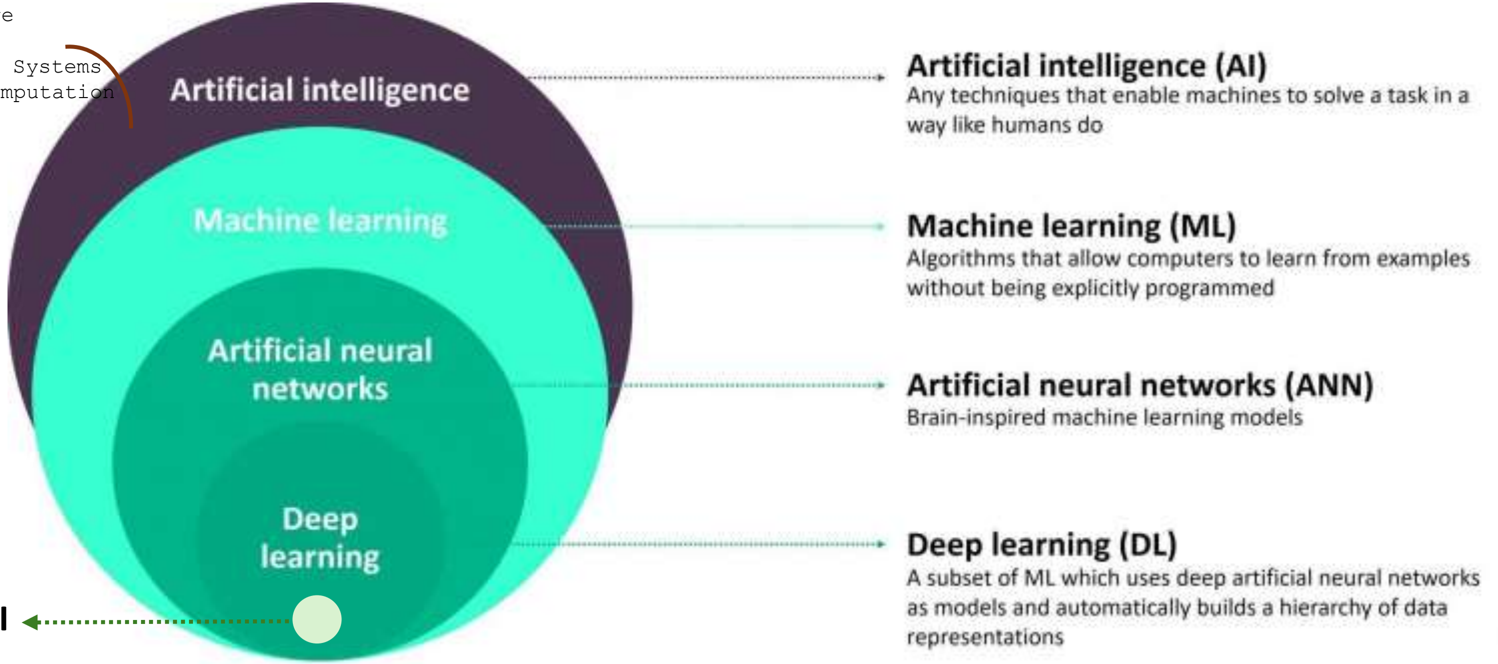

Generative AI

Artificial intelligence in ophthalmology

REVIEW

Curr Opin Ophthalmol 2026, 37:000–000
DOI:10.1097/ICU.0000000000001225

CURRENT OPINION

The evolving role of artificial intelligence in ophthalmology: basic science, translation, and clinical integration

Meghan N. Miller[a], David Hsu[b], Prashant D. Tailor[b] and Matthew R. Starr[b]

Purpose of review
Ophthalmology has emerged as the proving ground for medical artificial intelligence (AI) because of its imaging-centric workflows, standardized data acquisition, and well defined clinical endpoints. Recently, the field has started transitioning from narrowly trained, task-specific algorithms toward multimodal models that are trained on large datasets and can be adapted to many tasks. These models may integrate information from images, text, and clinical data and are increasingly trialed in real-world settings. This review examines recent advances and ongoing controversies in the use of AI as a basic science, translational, and clinical research tool in ophthalmology.

Multi-modal models and **large-scale integration** of images, clinical and omics data

REVIEW

Curr Opin Ophthalmol 2026, 37:236–243
DOI:10.1097/ICU.0000000000001209

CURRENT OPINION

Opportunities and challenges in leveraging multiomics and biobanks for vision science

Dolly Ann Padovani-Claudio[a,b] and Lisa Bastarache[b]

Purpose of review
Emerging biobank resources allow large-scale integration of eye-specific phenotypes with clinical, genomic, and multiomic data. This convergence enables unprecedented opportunities to systematically dissect the genetic architecture, epidemiology, and mechanistic pathways of both rare monogenic and common polygenic diseases. The review aims to critically examine how contemporary data extraction, multiomics, and analytic methodologies are reshaping disease classification, genetic discovery, and translational research in ophthalmology, while highlighting the associated challenges in leveraging these advanced tools.

Artificial intelligence in ophtalmology

TECHNIQUES

OPEN

Agentic Artificial Intelligence in Eye Banking: A Proposed Workflow

Ranit Karmakar, PhD, Elizabeth Kassahun Kiros, MD, MPH,† Siddharth Nath, MD, PhD,† Pearse A. Keane, MD, MSc,‡ Krishna Kumar,§ John Lohmeier, CEBT,§ Shameema Sikder, MD,† Eric Meinecke, CEBT,¶ Chris Hanna, MBA, CEBT,‖ and Allen O. Eghrari, MD, MPH†*

Purpose: To develop an agentic artificial intelligence (AI) framework that streamlines and standardizes eye bank operations by automating donor screening, image analysis, and tissue suitability assessment under expert supervision.

transplantation. It carries implications for both limited resource settings and the broader context of transplantation medicine.

Key Words: artificial intelligence, AI, eye banking, specular microscopy

Standardization of **eye bank operations** by automating **donors' screening**, **image analysis**, and **tissue suitability assessment**

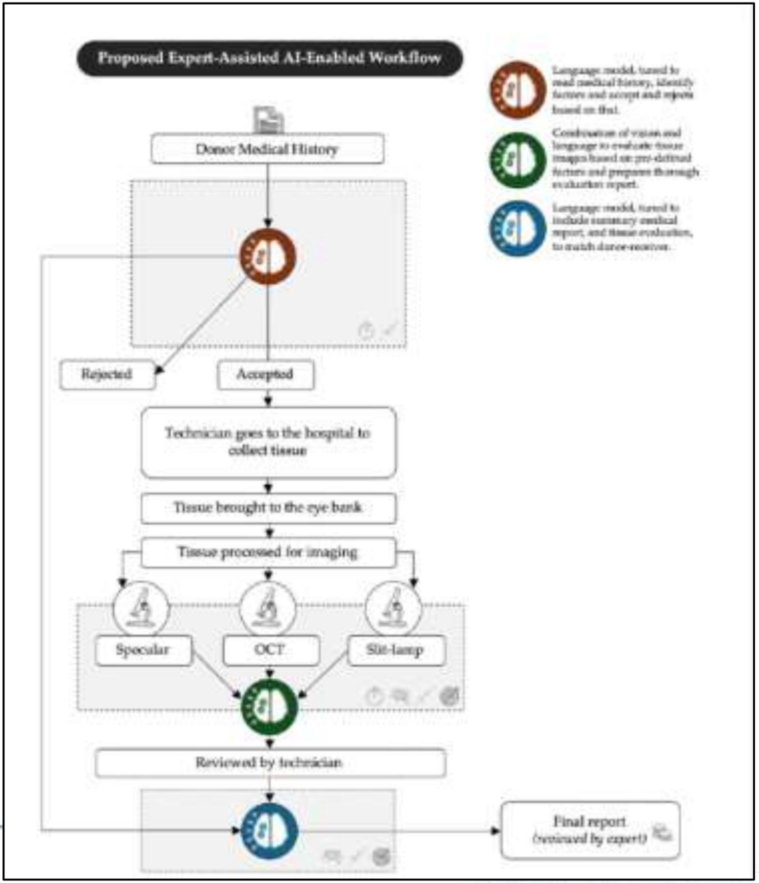

Artificial intelligence in ophtalmology

Article | Open access | Published: 13 September 2023

A foundation model for generalizable disease detection from retinal images

Yukun Zhou, Mark A. Chia, Siegfried K. Wagner, Murat S. Ayhan, Dominic J. Williamson, Robbert R. Struyven, Timing Liu, Moucheng Xu, Mateo G. Lozano, Peter Woodward-Court, Yuka Kihara, UK Biobank Eye & Vision Consortium, Andre Altmann, Aaron Y. Lee, Eric J. Topol, Alastair K. Denniston, Daniel C. Alexander & Pearse A. Keane

Nature **622**, 156–163 (2023) | Cite this article

Article | Open access | Published: 25 September 2025

Multimodal foundation model and benchmark for comprehensive retinal OCT image analysis

José Morano, Botond Fazekas, Emese Sükei, Ronald Fecso, Taha Emre, Markus Gumpinger, Georg Faustmann, Marzieh Oghbaie, Ursula Schmidt-Erfurth & Hrvoje Bogunović

npj Digital Medicine **8**, Article number: 576 (2025) | Cite this article

Article | Open access | Published: 29 August 2025

An artificial intelligence cloud platform for OCT-based retinal anomalies screening system in real clinical environments

Xinjian Chen, Jingtao Wang, Tianwei Qian, Jingcheng Wang, Yiming Ding, Su Zhang, Jingjing Liao, Qian Cheng, Ting Yang, Muhammad Mateen, Yu Fan, Zongming Song, Jili Chen, Suyan Li, Jianmin Hu, Wentao Yan, Haoyu Chen, Wencan Wu, Jing Huang, Tien Yin Wong & Xun Xu

npj Digital Medicine **8**, Article number: 559 (2025) | Cite this article

ORIGINAL ARTICLE

Development and Validation of a Multimodal Multitask Vision Foundation Model for Generalist Ophthalmic Artificial Intelligence

Authors: Jianing Qiu, Ph.D., Jian Wu, M.D., Ph.D., Hao Wei, M.Sc., Peilun Shi, M.Res., Minqing Zhang, M.Sc., Yunyun Sun, M.D., Ph.D., Lin Li, Ph.D., +35, and Wu Yuan, Ph.D. Author Info & Affiliations

Published November 27, 2024 | NEJM AI 2024;1(12) | DOI: 10.1056/AIoa2300221 | VOL. 1 NO. 12

Artificial intelligence in ophtalmology

Cureus
Part of SPRINGER NATURE

Open Access Original Article

Assessment of Generative Artificial Intelligence

Review began 08/12/2024
Review ended 08/22/2024
Published 08/26/2024

DOI: 10.7759/cureus.67833

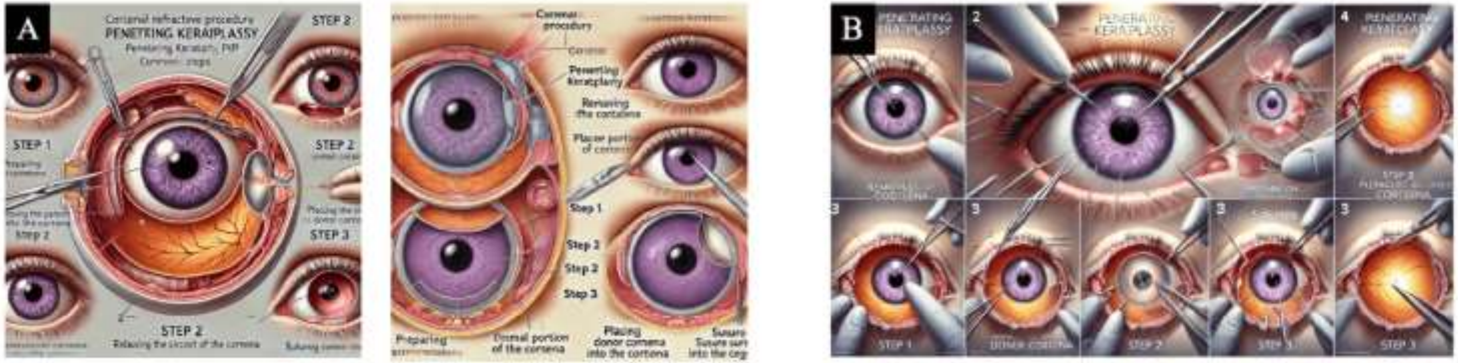

FIGURE 6: PKP

PKP medical illustrations created by the generative AI models after inputting final prompt number 6 (Table 1). Panel (a) final image generated by DALL-E 3. Panel (b) final image generated by MIM

Be aware of large language models

- **ChatGPT** is a chatbot based on a **generative AI architecture** named **transformer**
  - GPT (Generative Pretrained Transformer) is a **large language model** trained with contents available online
  - The name alludes the purpose, that is, it transforms things (e.g. translation): **(input) text → (output) text**
- **ChatGPT does not know what it is talking about:**
  - It concatenates the most likely sequence of words to correctly answer the prompt, but:
    it **does not understand** the prompt and it **does not know** the context for the answer
  - It is **not designed to create logically coherent, consistent, nor truth-based sentences**
    (there is no mechanism for verifying truth and correctness)
- Since GPT is not based on truth or on reliable sources, it **can create** (formally flawless but) **completely FALSE texts**
  - It can create false documents, invent facts and statements, cite non-

What is NOT artificial intelligence

What artificial intelligence is, actually..

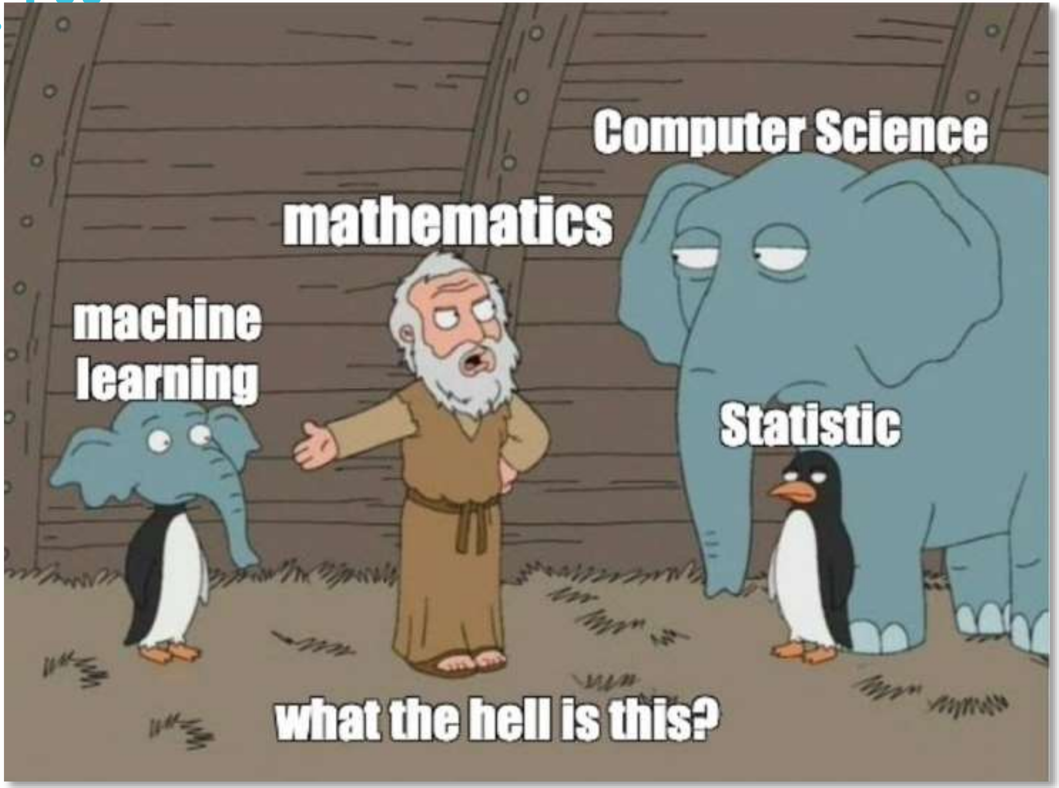

algorithm

What ~~artificial intelligence~~ is, actually:

just a set of instructions

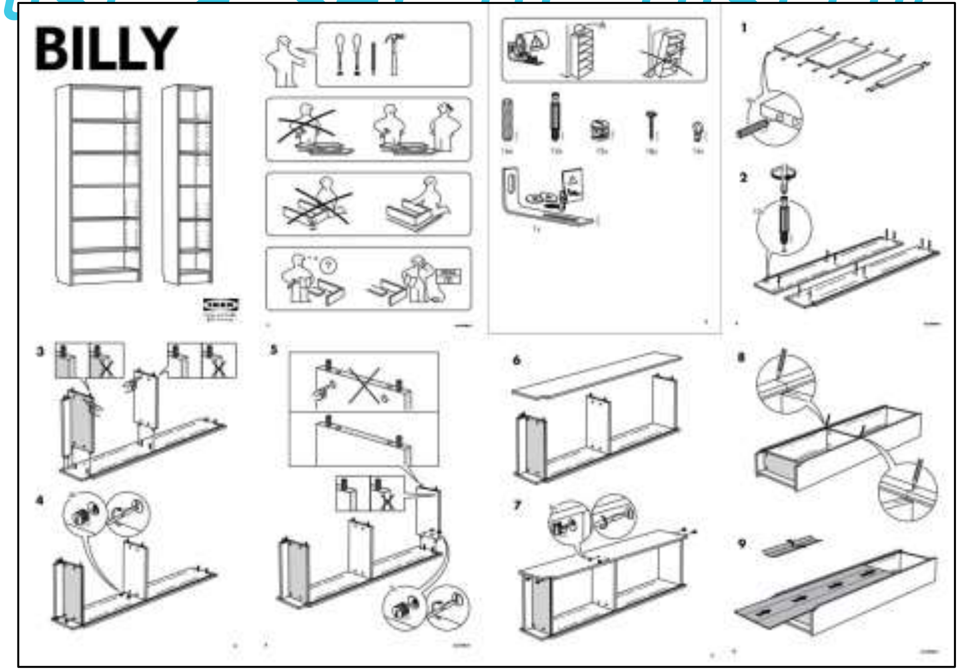

What does it mean «to learn»?

- In order to learn, machines need to be **trained with data**
- Machine learning is a **computational/statistical tool** that can be used to **solve problems** and **make inference** on a wide range of problems
  - Thanks to their **huge processing power**, machines can quickly **highlight** or **find patterns in big data** that would have otherwise been missed by human beings
- The **fundamental goal** of is to **generalize beyond the training samples**, that is, to successfully interpret data that the algorithm has never "seen" before
- There are different approaches to getting machines to learn, depending on:

The **task** you are trying to accomplish, i.e. the **scientific question**

The **type**, **quality** and **amount of data** that you **have in hand**

The three main categories of machin

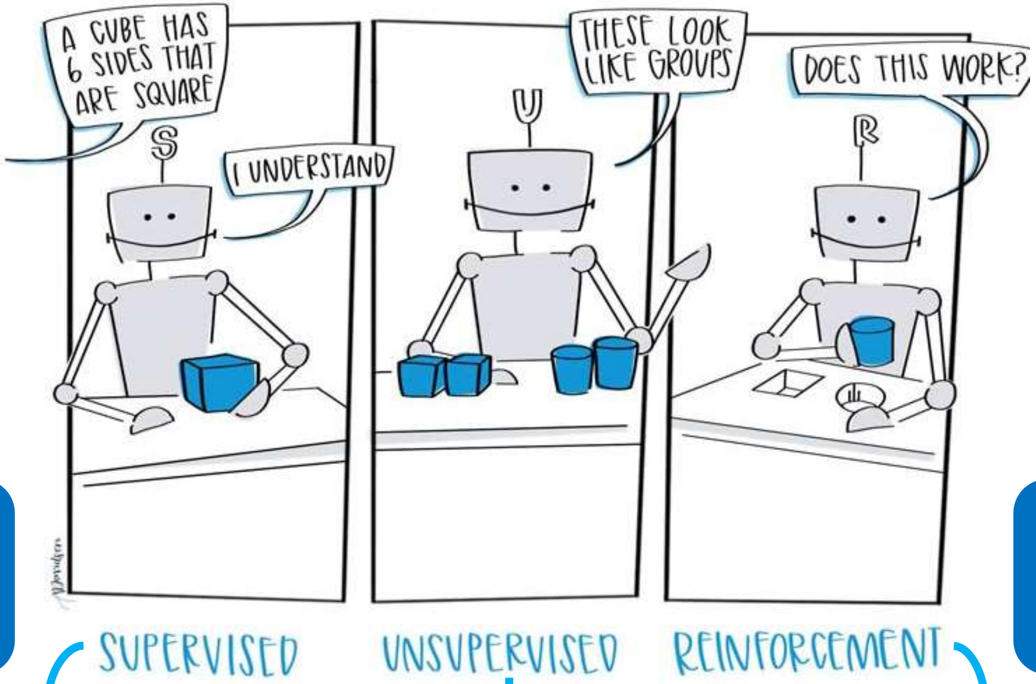

General rules from examples

Identify similarities

Make decisions (trial and error)

LABELS !

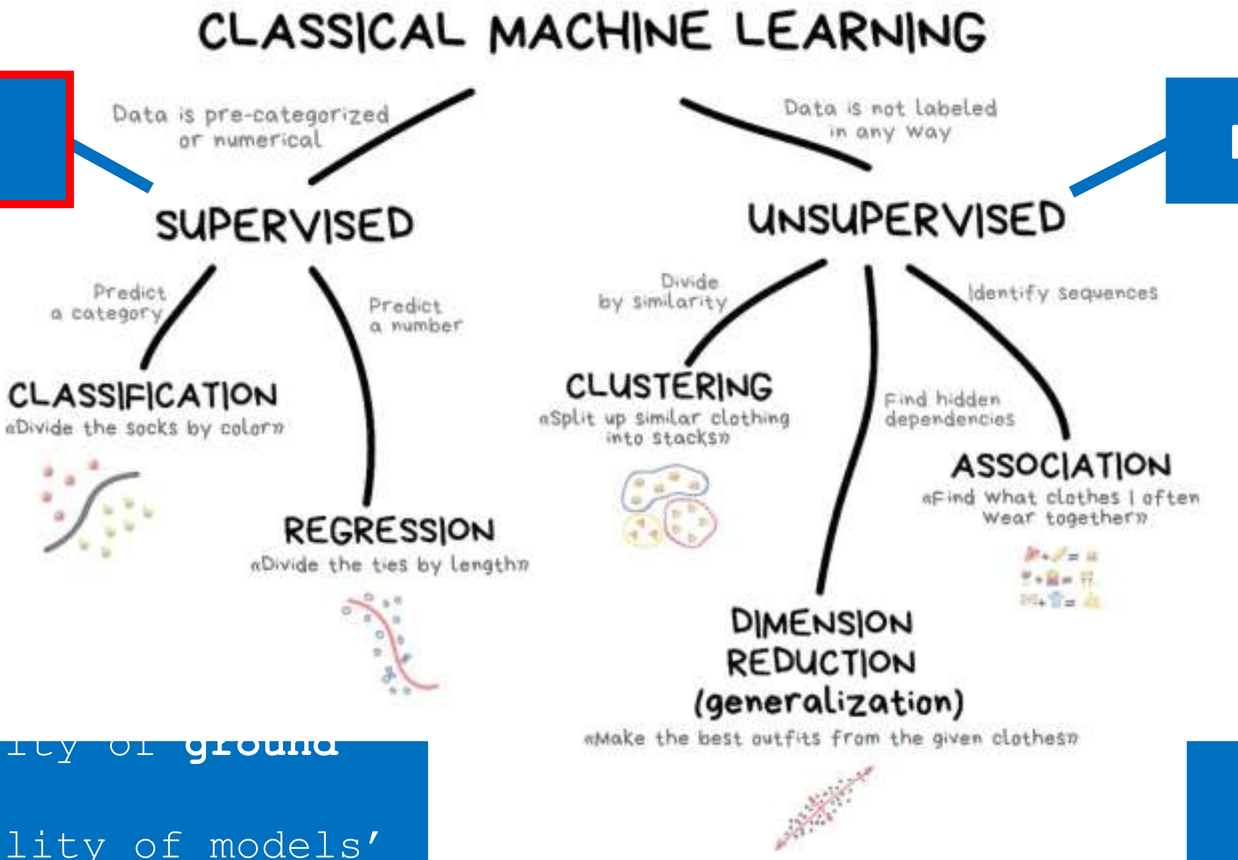

NO LABELS

- Availability of **ground truth data**
- Measurability of models' **accuracy**

...

Data, features, labels

- A **feature** is a characteristic or a measurable property of the data
  - Features can be **quantitative** (numeric values) or **qualitative** (string labels)
- For instance, a quantitative **feature** in OCT could

From: Evaluation of thickness of individual macular retinal layers in diabetic eyes from optical coherence tomography

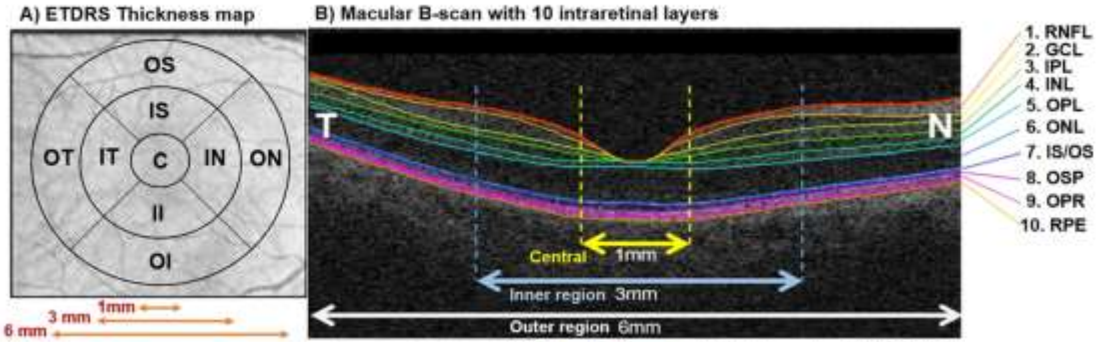

Scientific Reports 14, Article number: 17909 (2024)

The **label** is a **chosen feature**, used for training and for model evaluation.

For instance, **health status**:

- Patient with multiple sclerosis
- Healthy control

The importance of labels…

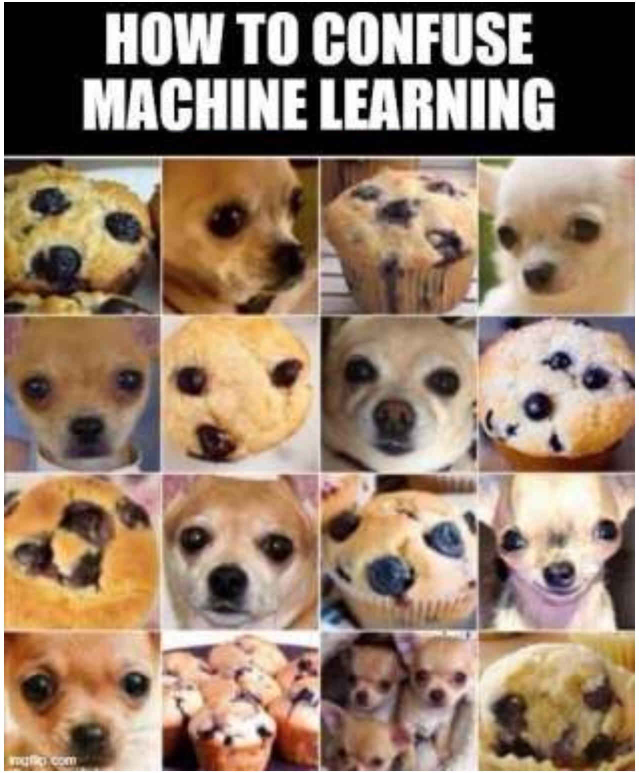

How supervised learning works: classification task

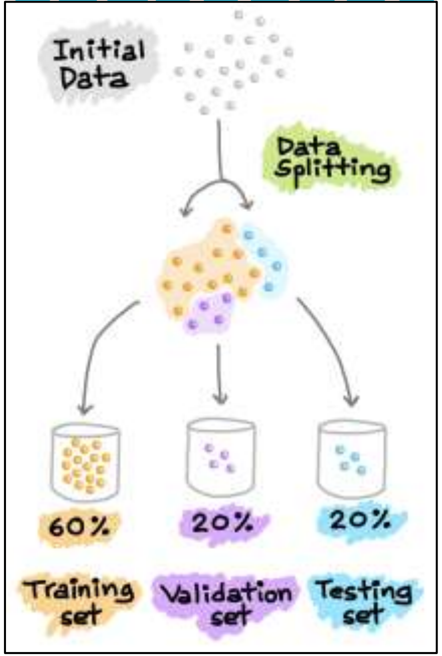

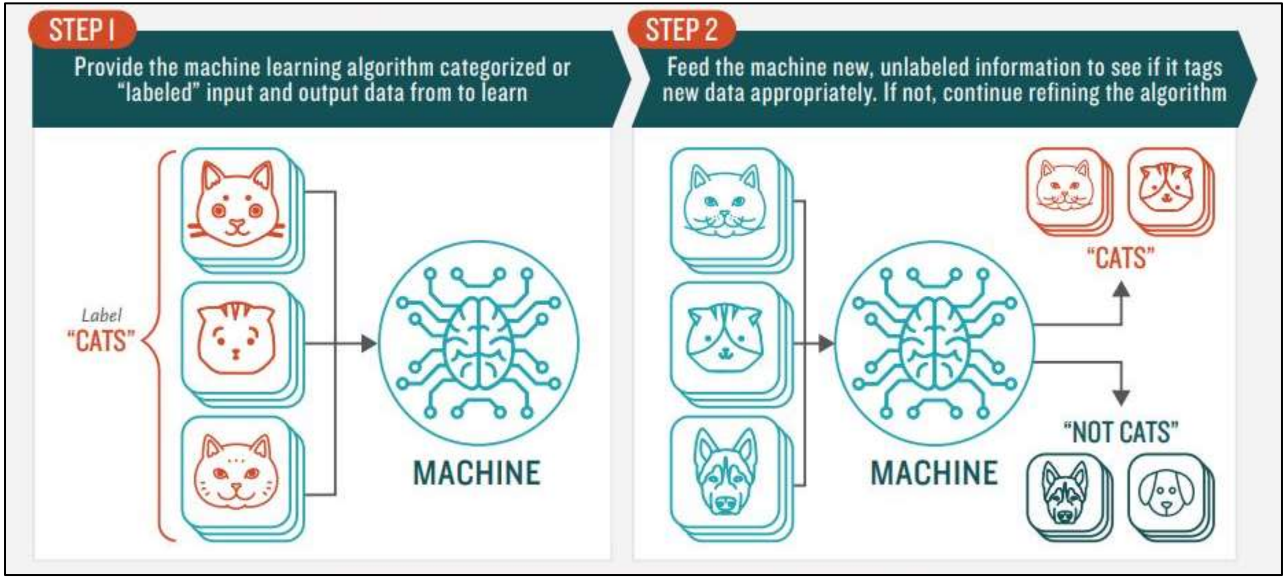

Supervised learning – Neural networks

- A **neural network** consists of layers of interconnected **neurons** that cooperate to solve a **classification** task

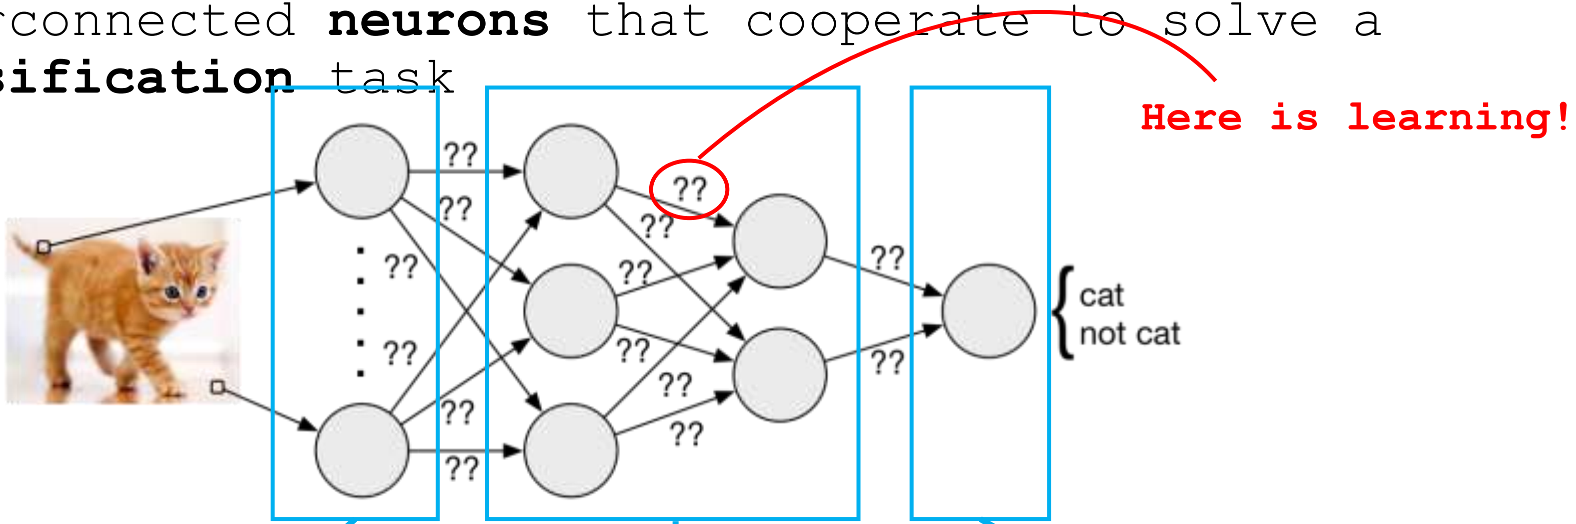

Input layer: no computation is done here,
the input neurons just pass the information
to the next layer

Hidden layer(s): hidden neurons
transfer the weights from the input
layer to the following layer

Output layer: an "activation function" maps the network results to the output class

Supervised learning - Deep neural networks

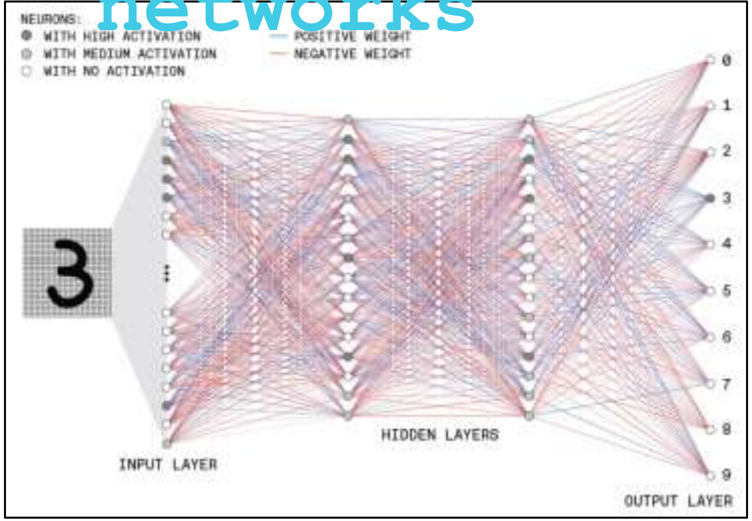

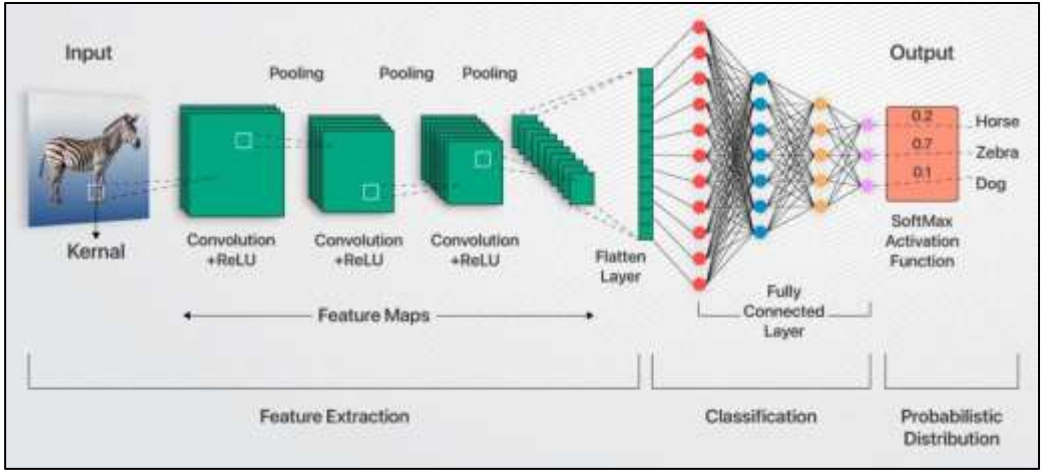

What *might go wrong* in neural network learning?

- Which neural network **archite**
  - How many layers? How many neuro
- Which **activation function**?
- Which **loss function**?
- Which **optimization algorith**
- Which data?
  - How many **samples**?
  - How many **features**?
  - Are the **labels** reliable/accurate
- …

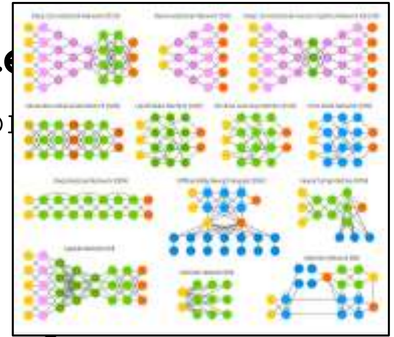

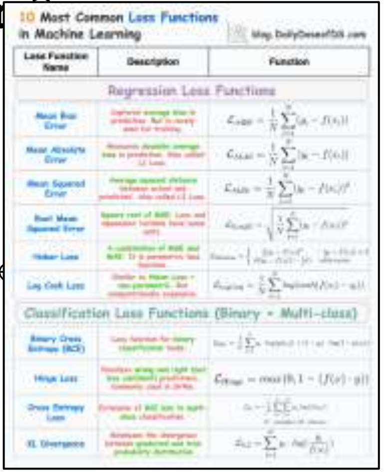

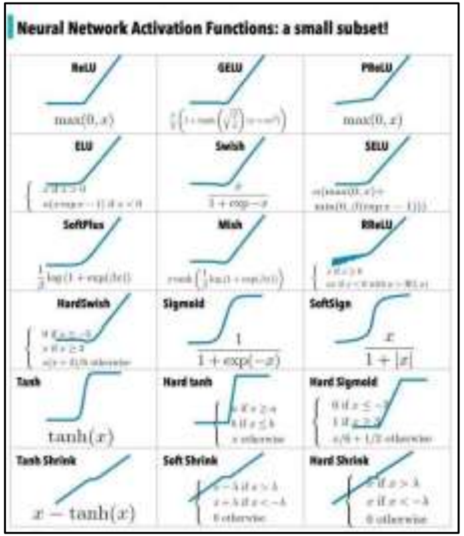

How to evaluate a neural network model?

Thanks to labels!

- Given a classification of a specific data set, we can obtain:
  true positives and **negatives** (TP, TN), **false positives** and **negatives** (FP, FN)

- M … **goodness of a** … m

Metric	Formula	Interpretation
Accuracy	$\frac{TP + TN}{TP + TN + FP + FN}$	Overall performance of model
Precision	$\frac{TP}{TP + FP}$	How accurate the positive predictions are
Recall Sensitivity	$\frac{TP}{TP + FN}$	Coverage of actual positive sample
Specificity	$\frac{TN}{TN + FP}$	Coverage of actual negative sample
F1 score	$\frac{2TP}{2TP + FP + FN}$	Hybrid metric useful for unbalanced classes

Accuracy measures the proportion of correct predictions out of all predictions

Sensitivity measures the proportion of actual positives that are correctly identified as such

Specificity measures the proportion of actual negatives that are correctly identified as such

Which algorithm should you use?

- **The size/quality/nature of data, and the availability of labels**
- **What you want to do with the d**
  - Prediction? Description? Data explorati
- **The required computational ti**
  - Speed/performance on new data?
- **Accuracy and ease of use**
  - Easy to implement? Interpretability?
- **Presence of any type of bias**
  - Class imbalance, annotation bias, aggregation bias, algorithmic bias, demographic bias, distributional shift,

Which algorithm should you use?

- Choosing the right algorithm requires **trading off one benefit against another,** including model speed, accuracy, and complexity
  - **The "best method ever" <u>does not exist</u>**
- **More than one algorithm** should always be tested and compared

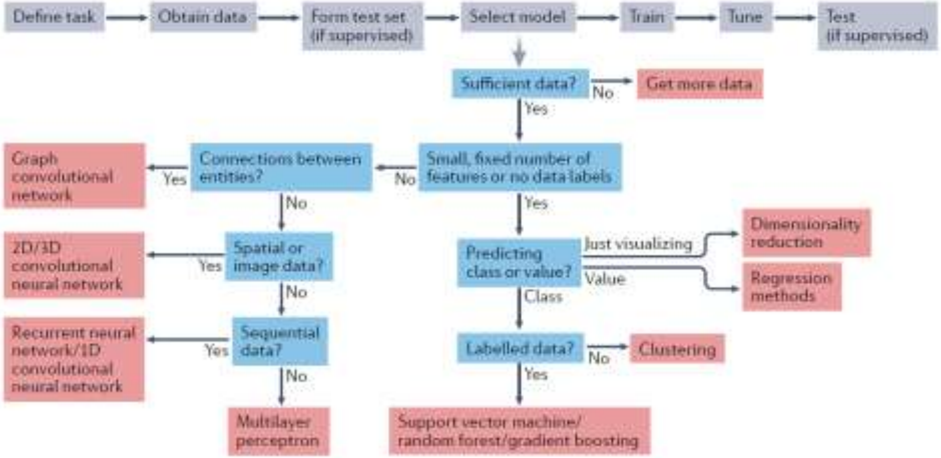

The black box metaphor and interpretability

- AI sometimes appears as a kind of magic trick or a **black box**, that is, a system for which we can only observe the inputs and outputs but not the internal workings
  - Thus, it is difficult to figure out **why and how** they behave the way they do
  - The lack on **interpretability** **y** is a major issue i

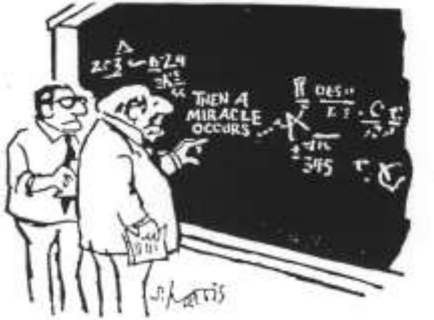

I think you should be a little more specific, here in Step 2

Interpretability vs accuracy

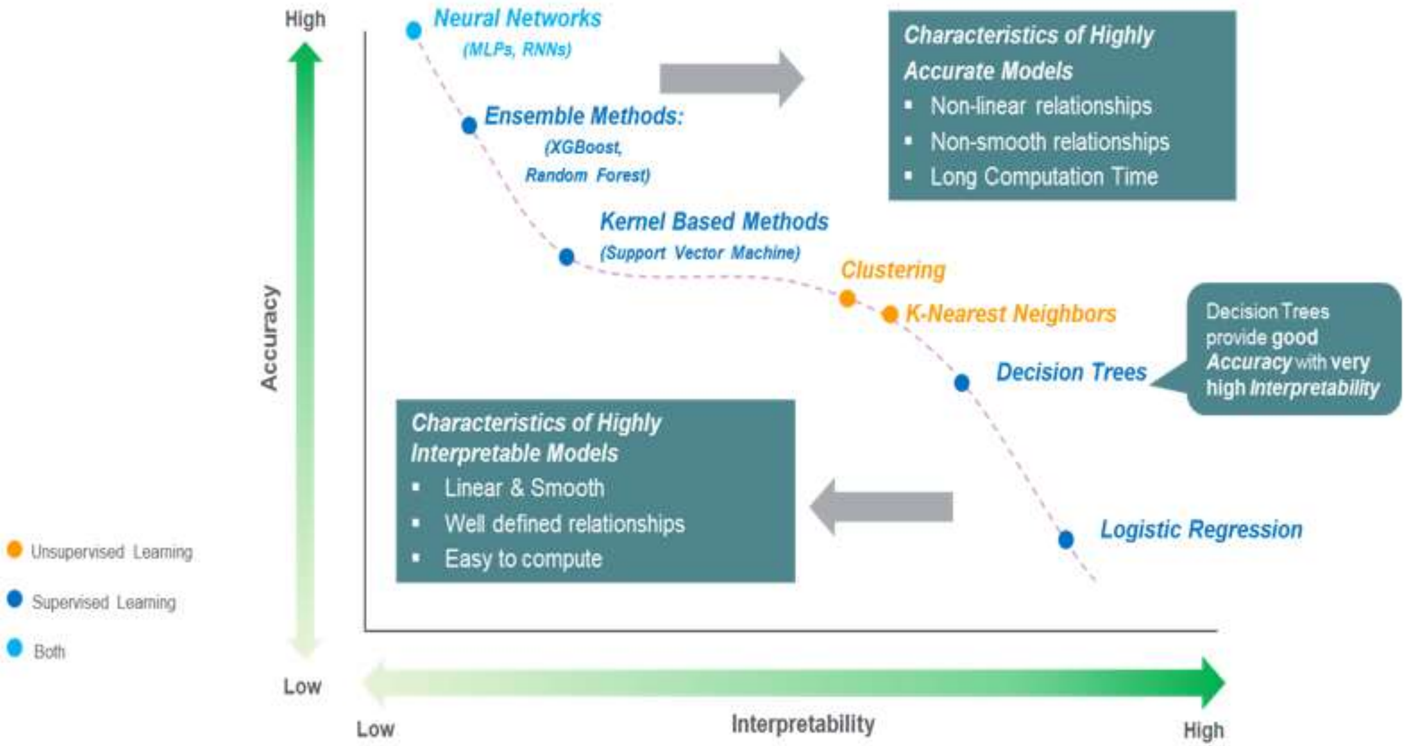

To conclude

- Artificial intelligence is **transforming the way we process data**, build models, and develop new tools in medical disciplines

- Artificial intelligence applied to **healthcare** requires:
  - issues **awareness**
  - **caution**
  - **transparency**
  - **responsibility**
  - **effective**, yet **understandable** and **interpretable** decision-making tools
  - **concrete impact** on healthcare
  - open and clear **communication with society**

Thanks for your attention!