GOOD MORNING

I am Pratik

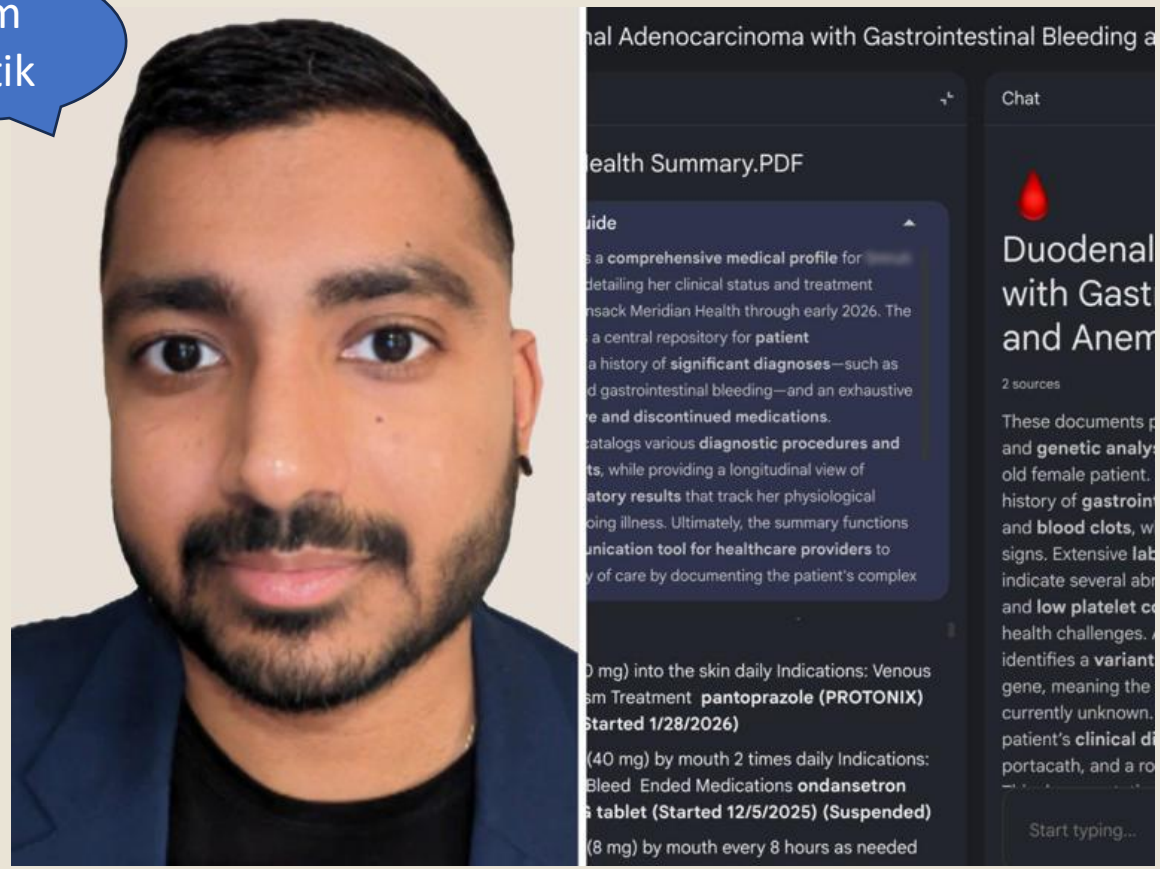

***4/5/2026**: Pratik's mother was diagnosed with an advanced cancer affecting the small intestine.*

Pratik: "I vibe coded an AI tool to help my mom fight stage 4 cancer. Now my friends use it to manage their parents' care" -link

"I wanted to keep things as simple as possible, but as her medical records grew to **1,600 pages long**, not including any of the images or scans, **we couldn't load it into a single thread anymore**. Eventually it was just way too much data, and I ran out of context.

"There were no less than three times that, guided by this workflow, we caught emergency situations and I believe saved her life"

Hyderabad: AstraZeneca Pharma India, a subsidiary of British-Swedish pharma giant AstraZeneca Plc, has signed a memorandum of understanding with the Telangana govt to introduce AI-powered lung cancer screening across 20 public health facilities.

India **4/7/2026**: AstraZeneca, Telangana govt to roll out AI-powered lung cancer screening in public hospitals. link

Introduce AI-powered lung cancer screening across 20 public health facilities

If they can do it, we can do it too!

Introduction to AI and Machine Learning

Amr Hilal
Tennessee Tech University

4/8/2026

What is AI?

- **What's AI but not ML?**
- Possible combinations of systems **Thinking/Acting Humanly/Rationally**:
  - Systems that act humanly.
  - Systems that think humanly.
  - Systems that think rationally.
  - Systems that act rationally. 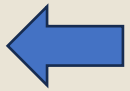

What's Machine Learning?

- People wanted to build **intelligent machines** but relatively few basic things can be **programmed**.
- What about more interesting things?
  - Almost impossible to programmatically describe.
  - So many cases and exception
- The only way is to have **a machine learn to do it** by itself.
  - New Capacity for computers to **learn like humans learn by experience**.

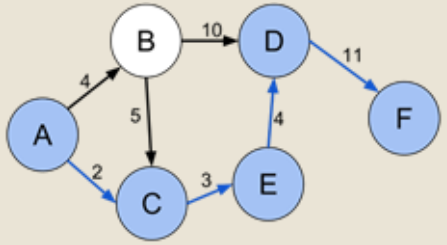

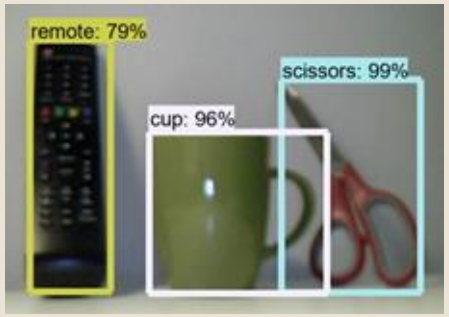

How Machine Learning came to exist?

- Term coined by Arthur Samuel in 1959 developing a **checkers game**
- **Samuel thought why not let the computer come up with its own rules instead of us telling it the rules.**
- Arthur wasn't a very good checkers player
- Had to program maybe tens of thousands of games to teach the computer good and bad positions
- Experience → Computer better checkers player.

World's first first successful **self-learning program**

Picture: https://www.ibm.com/ibm/history/ibm100/us/en/icons/ibm700series/impacts/

A formal definition of Machine Learning

- **Arthur Samuel (1959)**, older informal definition "the field of study that gives computers the ability to learn without being explicitly programmed."
- **Tom Mitchell (1998):** a more modern definition: "A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E."

  Example: playing checkers.

  E = the experience of playing many games of checkers

  T = the task of winning checkers game.

  P = the probability that the program will win the next game.

Machine Learning in the Science Process

- **Some phenomena** are **too complex to explain** in a **causal** form
- Machine Learning **makes up for the lack of complete understanding** of how a phenomenon works by **learning from many examples**
- **Data is essential for machine learning**
- We seek to make ML model **explainable** to reclaim the benefit of classical causal explanations

Induction vs Deduction

- **Deduction:** Going from general rules to a specific set of specific conclusions.
- **Induction (opposite):** Going from a specific set of observations to a general rule.
- **The Difference is:**
  - The **deductive** conclusions **are guaranteed** to be **correct if the premises are correct.**
  - The **inductive** conclusions **may** be **incorrect** (because we generalize).
    - That's a cost of learning via examples.

Machine Learning vs Traditional Programming

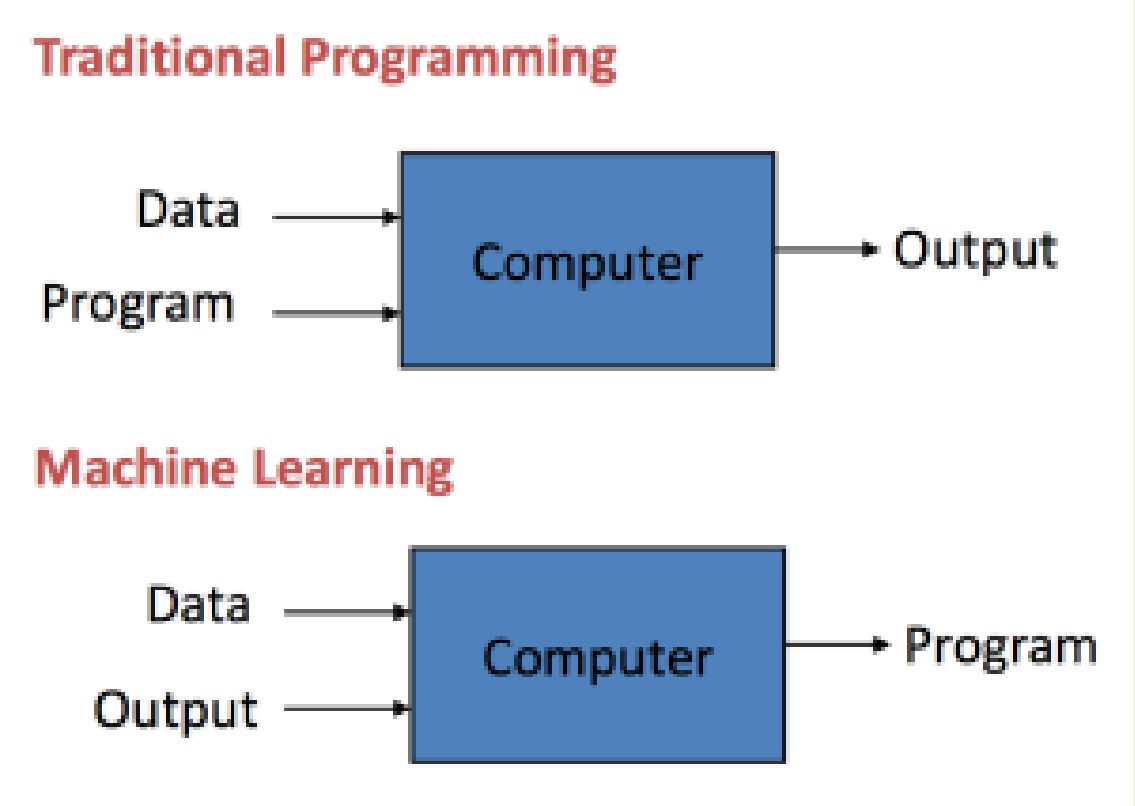

AI vs ML

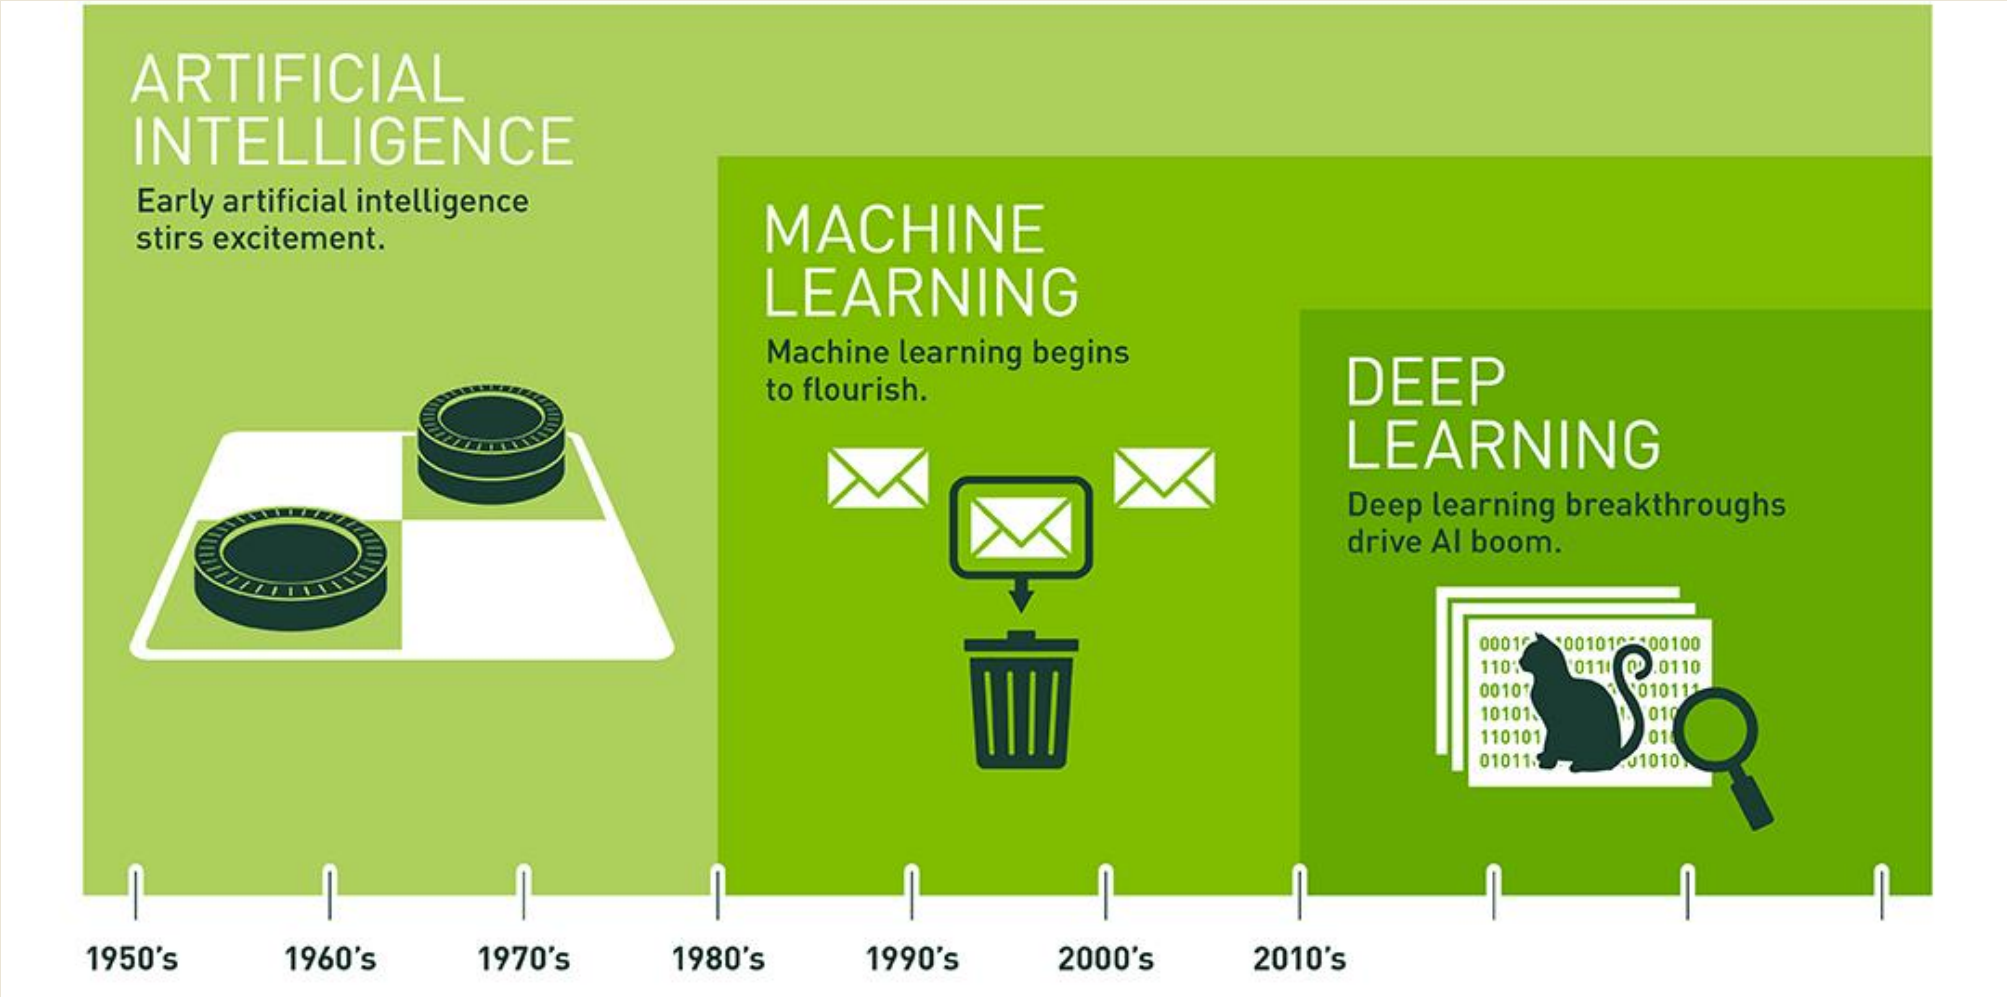
ARTIFICIAL INTELLIGENCE
Early artificial intelligence stirs excitement.
MACHINE LEARNING
Machine learning begins to flourish.
DEEP LEARNING
Deep learning breakthroughs drive AI boom.
1950's
1960's
1970's
1980's
1990's
2000's
2010's

Major Types of Machine Learning

Supervised Learning

- **Training data** are **labeled** (i.e., input and output)
- Two categories of problems
  - Regression (continuous output)
    - Housing prices, what is the input here?
    - Age given a picture
  - Classification (discrete output)
    - Benign vs. Malignant tumor, think about the input!!

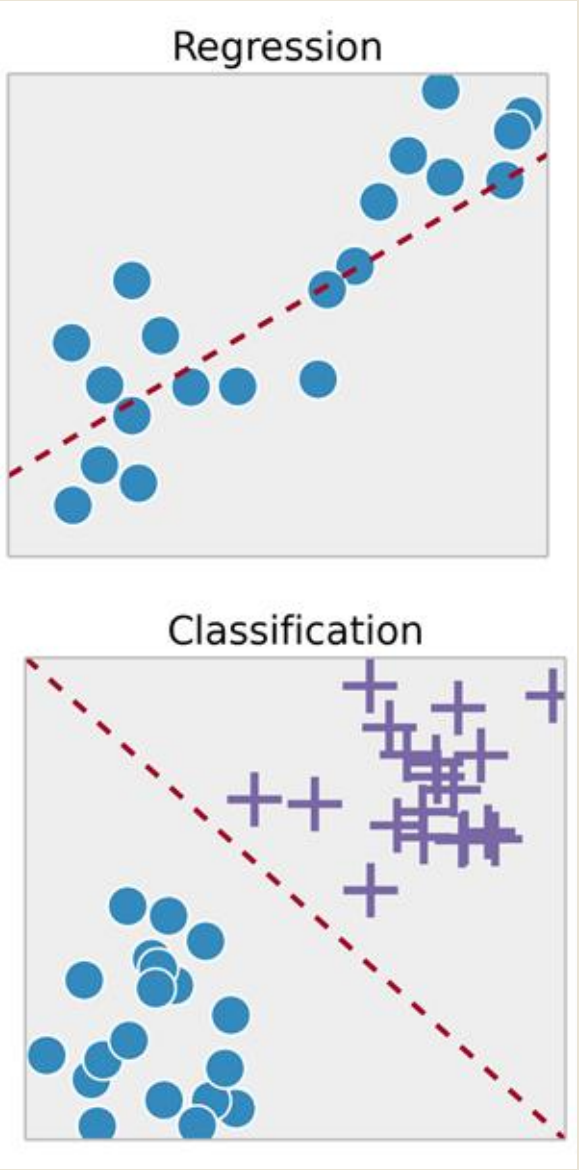

Unsupervised Learning

- The goal is to **find interesting patterns**
  - Knowledge Discovery
- Training data are **not labeled**
- Put data into buckets of similar features

  “Clustering”
- Examples
  - Fraud detection
  - Customer segmentation

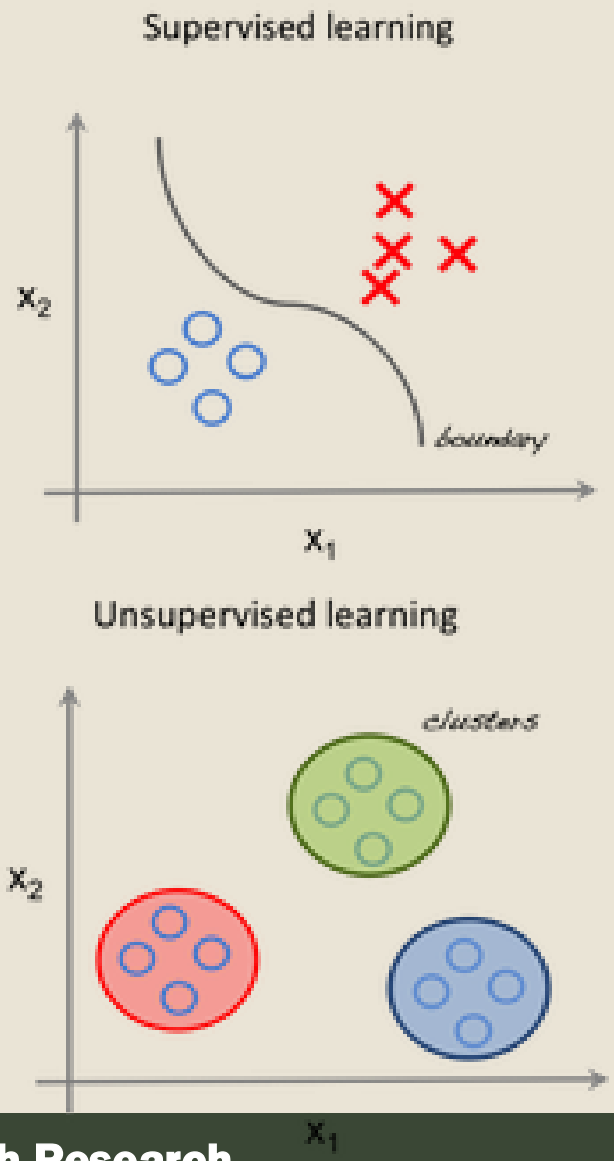

Reinforcement Learning

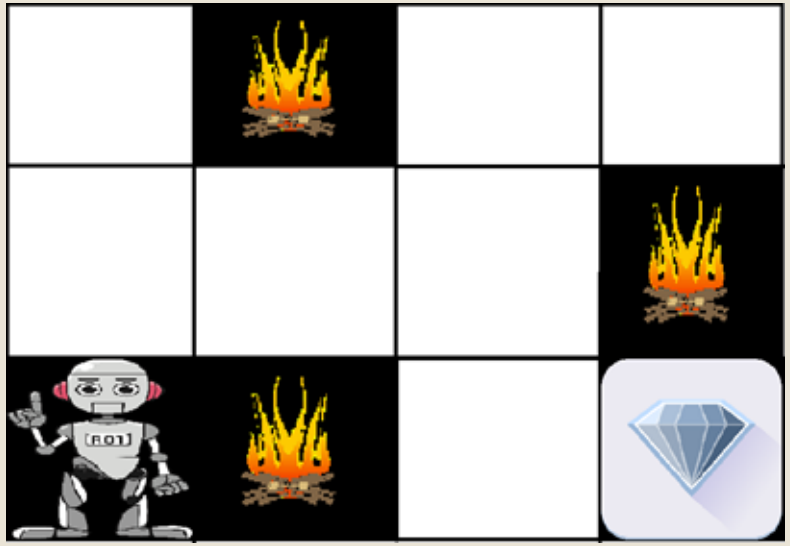

- Learns from a **sequential decision making** process
- An algorithm, learns by interacting with its environment **(Agent produces its own data, not ready before)**
- Rewards and penalties **signal** good and bad behavior
- Learns by **maximizing its reward and minimizing its penalty.**
- Games, real time decisions.

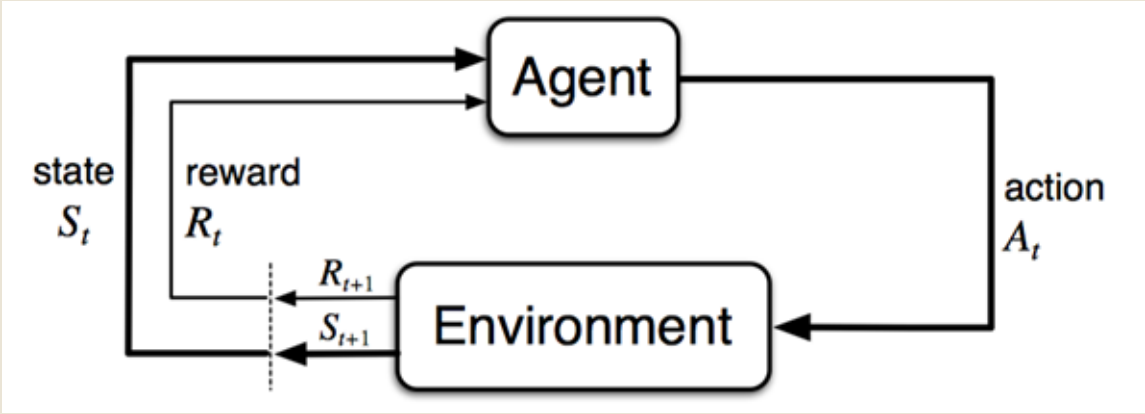

AI vs Agentic AI

- Traditional AI
  - Performs **specific tasks** (e.g., classification, prediction)
  - Does not maintain long-term goals or plans
  - Examples: image classifiers, recommendation systems
- Agentic AI
  - Acts as an **autonomous agent** with goals
  - Can **plan, decide, and take actions** over multiple steps
  - Often interacts with tools, environments, or APIs
  - Maintains **context and memory** across tasks
  - Example: Build a website, shopping, humanoid robots

How Machine Learning Works

(in a nutshell)

General Machine Learning Process

- Learning/Training (AKA model fitting)
  - Feature Extraction/Engineering
  - Data Preprocessing
  - Divide data into buckets
    - Training [+ Validation]
    - Test
  - Learn Iteratively
  - Eval, Error Analysis and Tuning
- Prediction

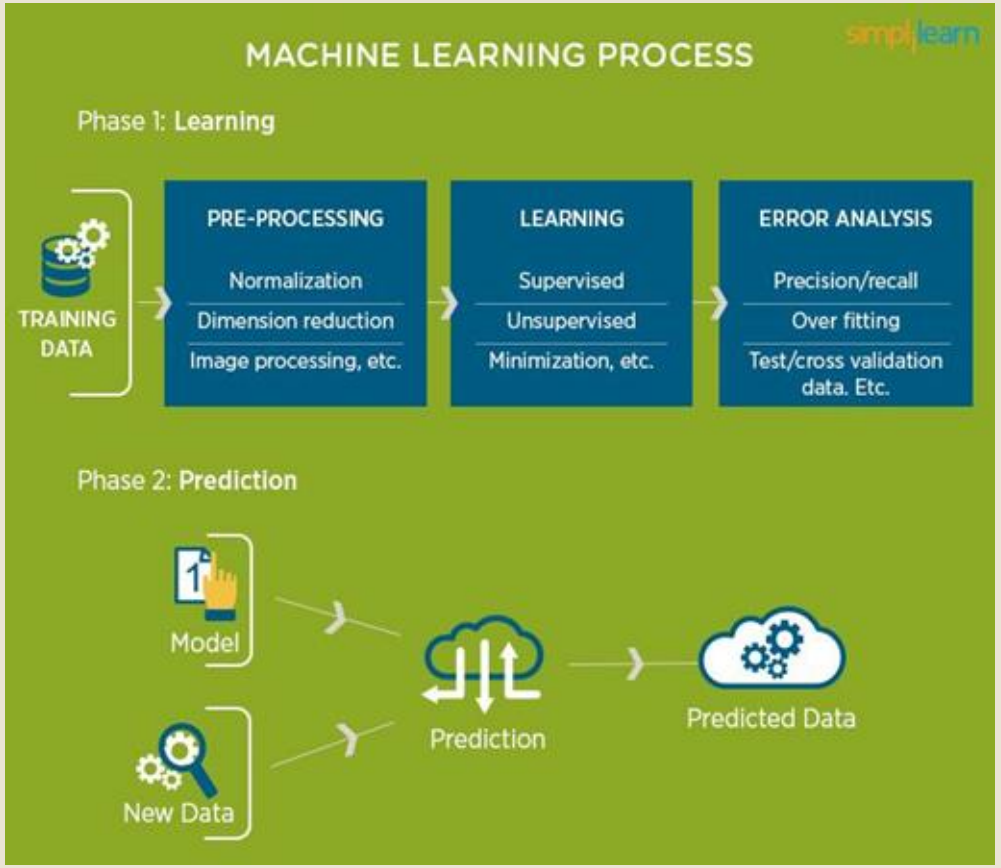

Decision Trees

- Powerful
- Interpretable
- Easy to build

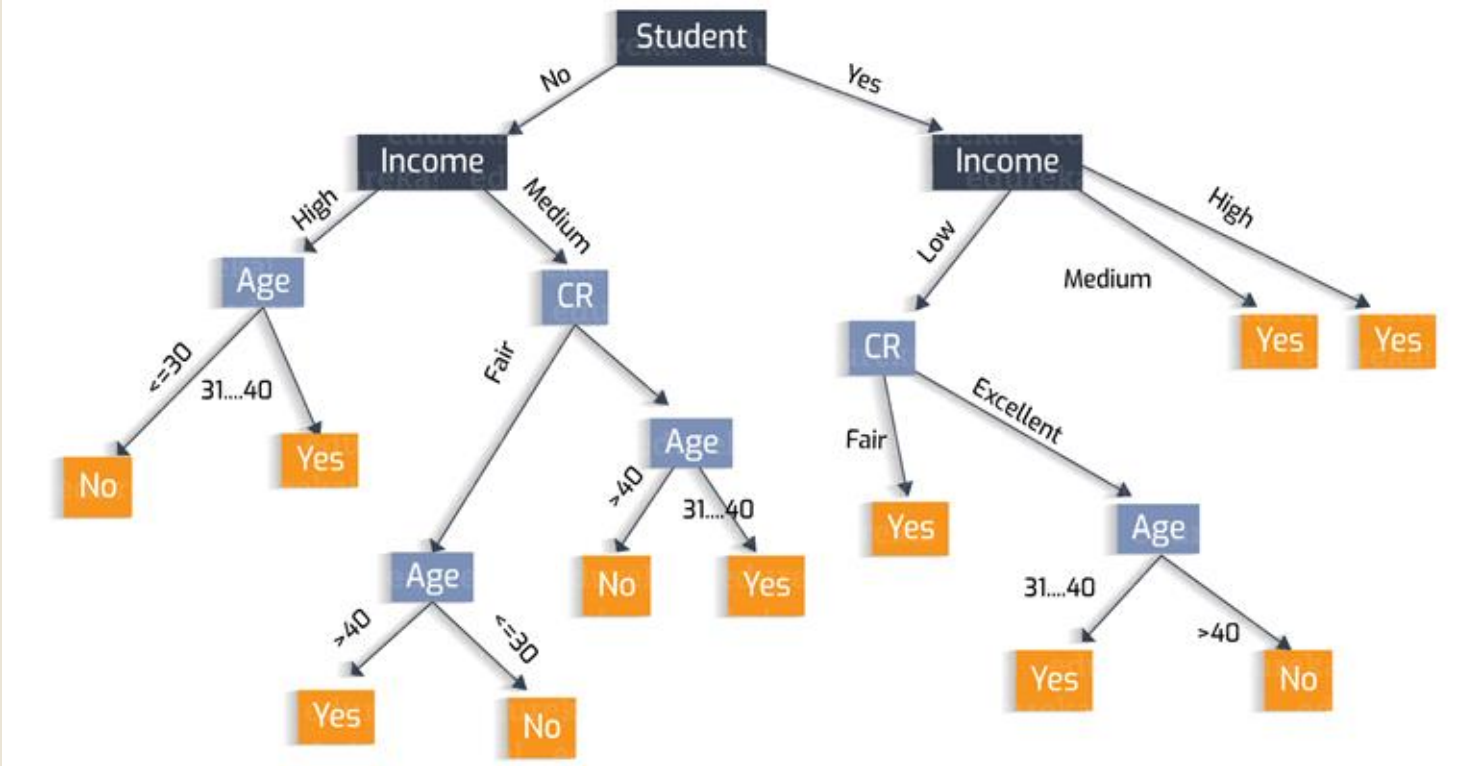

Picture: https://heartbeat.fritz.ai/understanding-the-mathematics-behind-decision-trees-22d86d55906

Linear Machine Learning Models

- **Knowns**: x's (feature), y (label)
- **Required**: w's and b
- Incremental process
- Iterative feedback system
  - **Optimize** to reduce **loss function** $f(y - \hat{y})$
- Less explainable

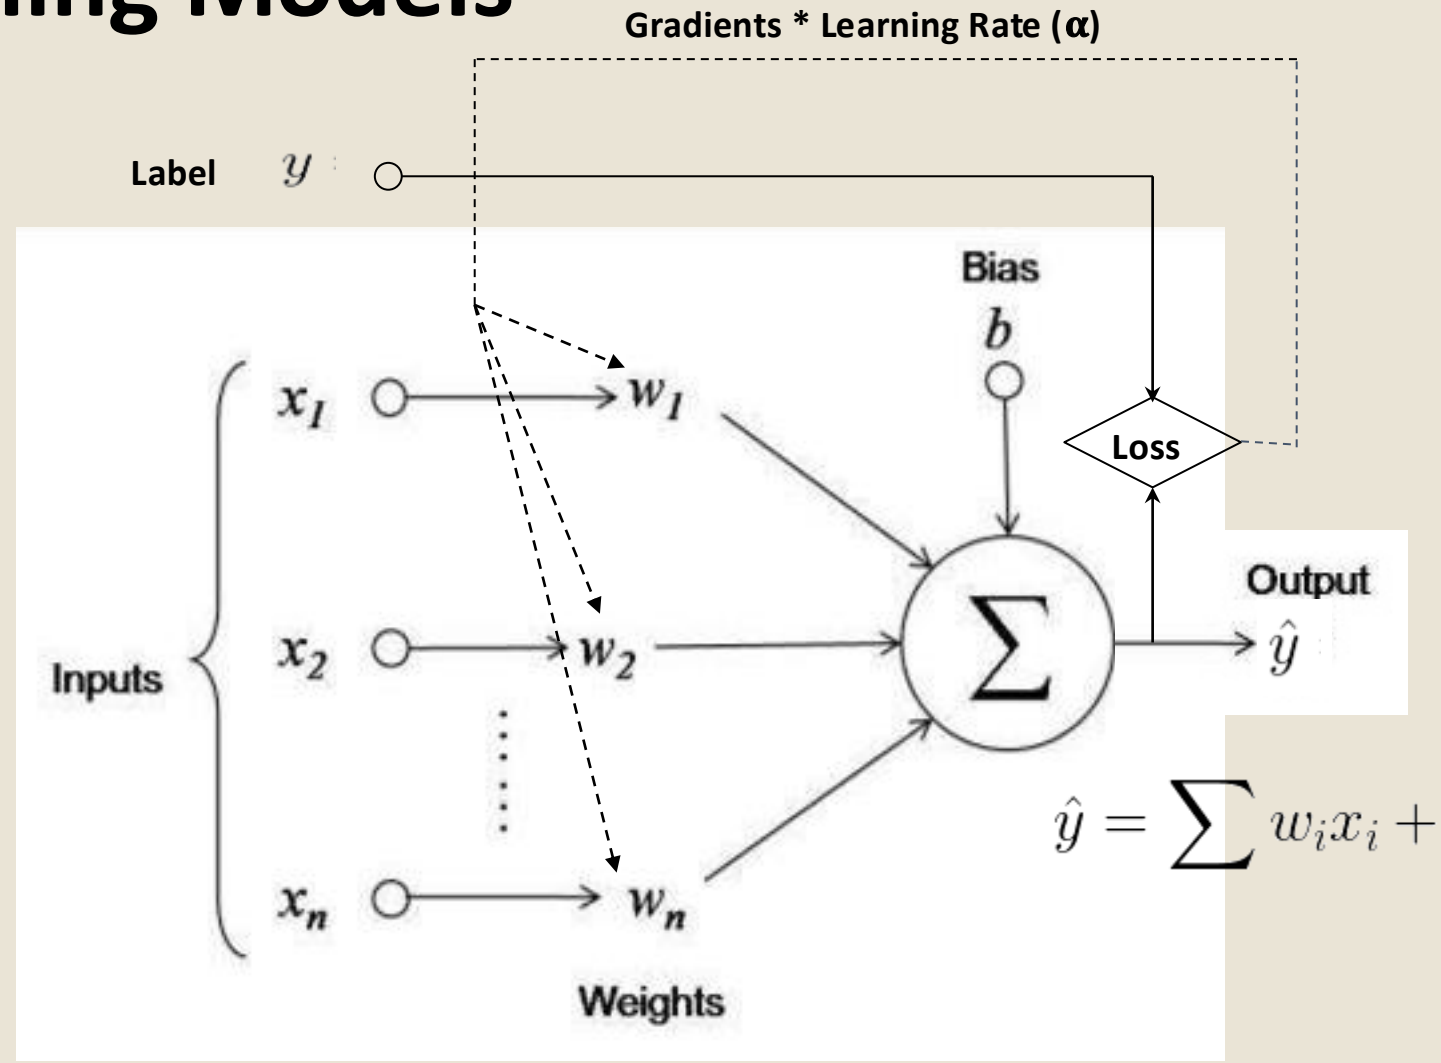

Optimizing the Loss Function

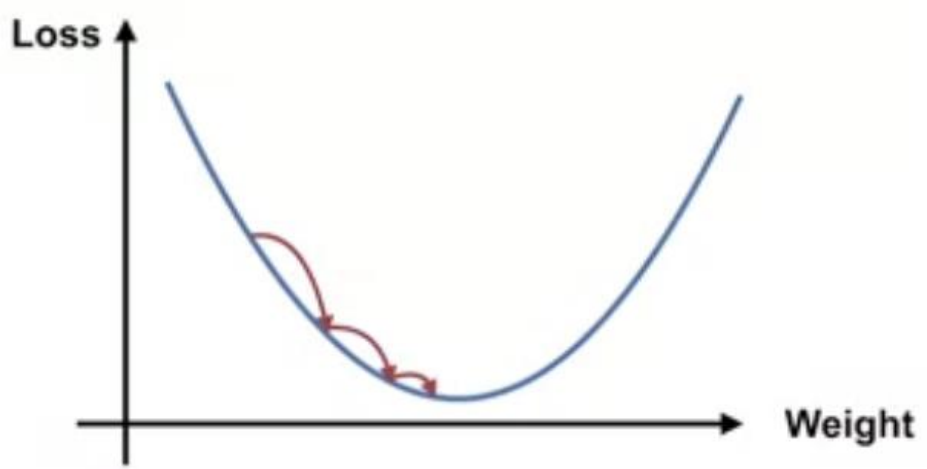

Simpler Linear Models

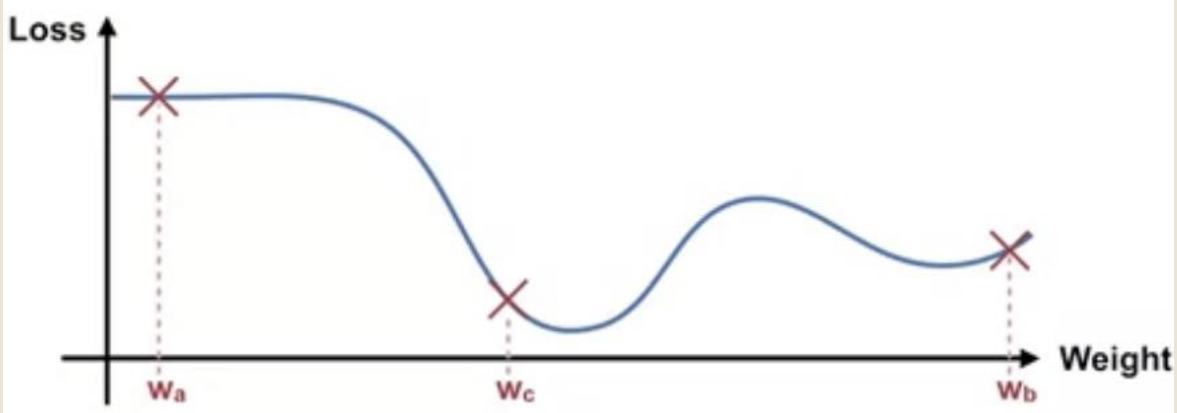

Neural Networks (more complex)

Neural Networks

- Humans learn by experience.
- Neurons connect in our brains to form chunks that represent acquired knowledge forming **Natural Neural Networks**
- More connected neurons → more cognitive capabilities

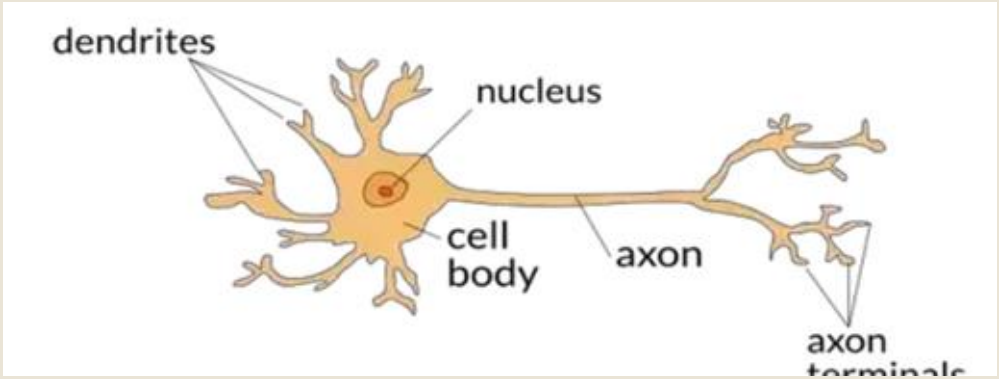

Neural Networks

- An artificial neural network tries to emulate our brains
- A neuron is similar to a Traditional ML unit
- Hooking units of Traditional ML forms a NN.
- **Deep Neural Networks:** Neural Networks with many layers
- An **Artificial Neural Network** is capable of **learning more complex problems**

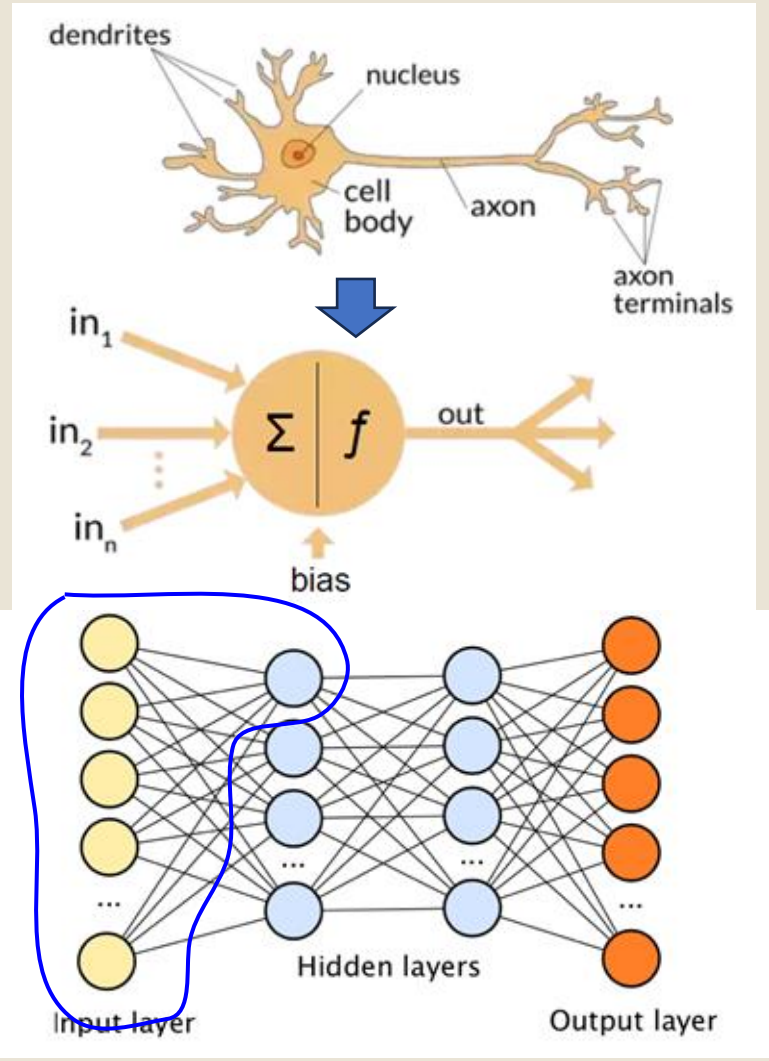

Parameters vs Hyper-parameters

- **Model Parameters**: the weights and biases that are learned during training
  - W's and b's
- **Model Hyper-parameters**: configuration **knobs** that optimize model performance
  - Regularization parameters, learning rate
  - Number of hidden layers, neurons per layer

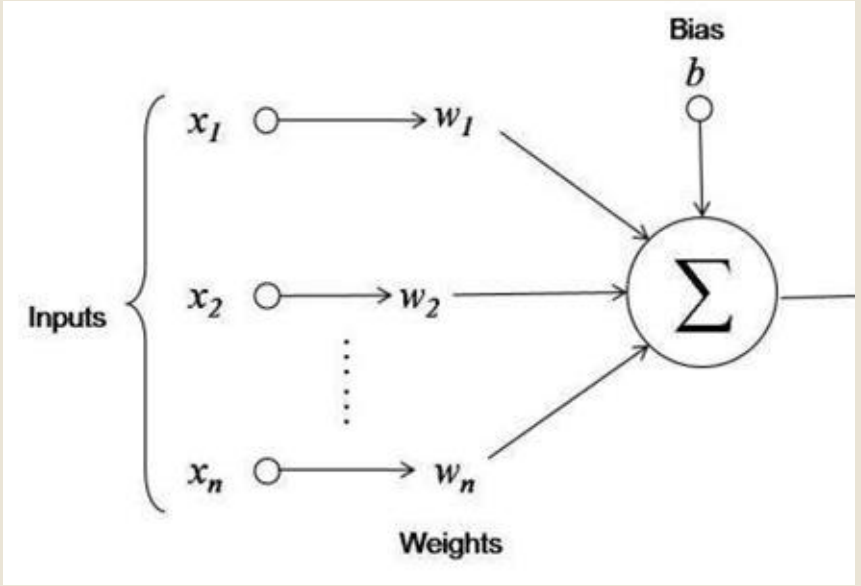

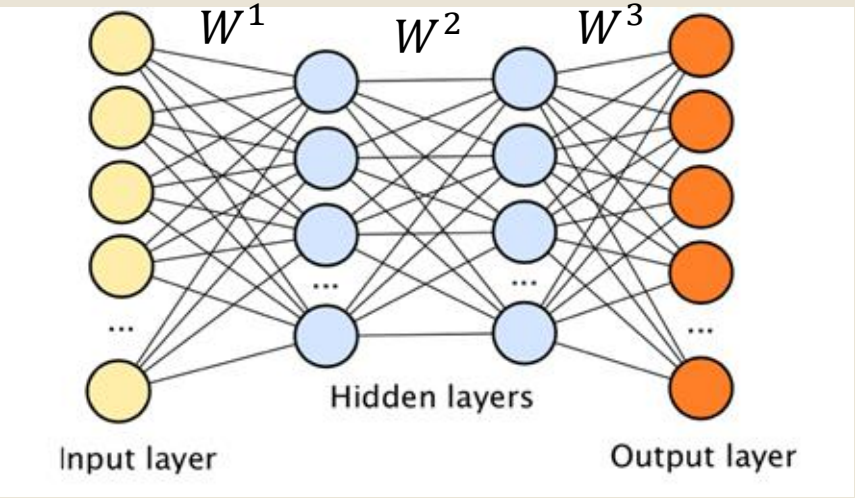

Machine Learning vs. Deep Learning

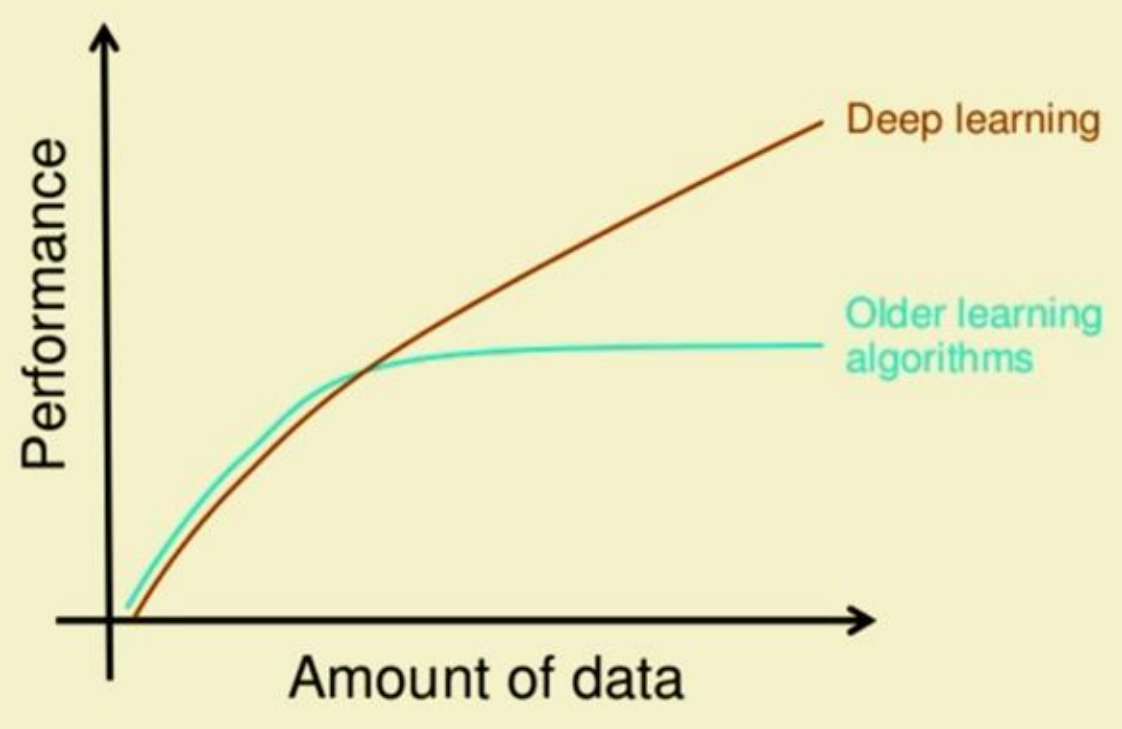

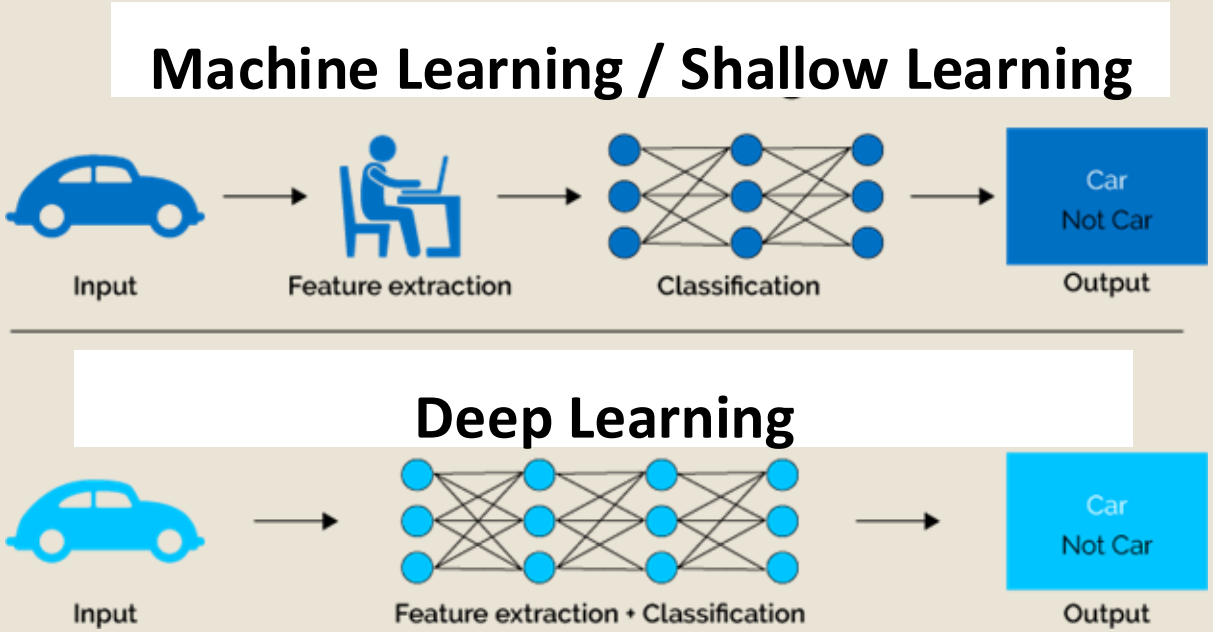

Model Evaluation, Error Analysis, and Tuning

Data Split, Model Selection, and Hyper Parameter Optimization

- **The ultimate goal of ML training is to generalize well to unseen data**
- We split data into Train and Test
  - Use test data to test model generalization on unseen data
- What should we do to find the best model settings then?
  - Use val dataset
  - Model Selection (Hyper Parameters Opt.)
- **Train model on train, evaluate on val, and pick the best model using test**

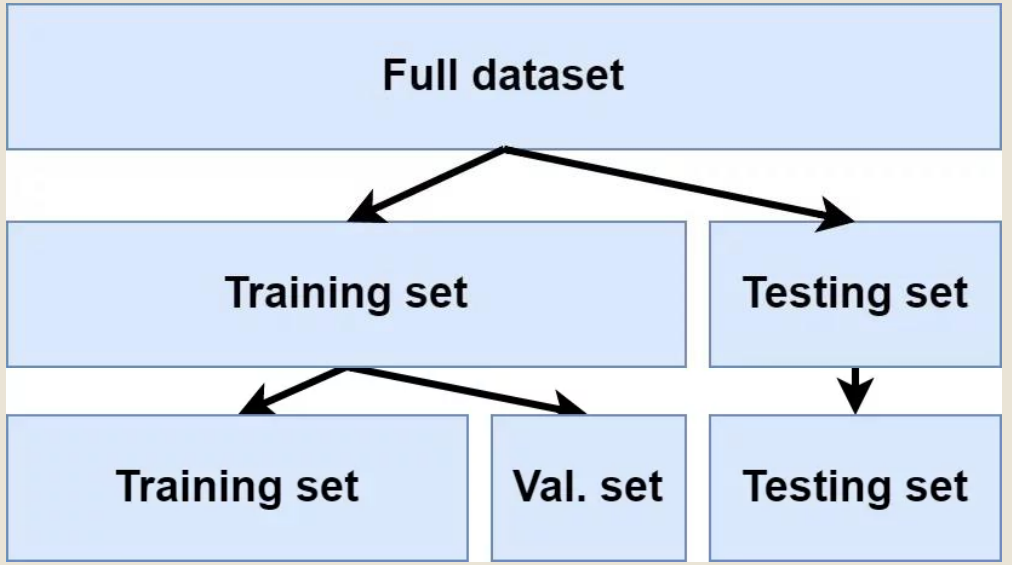

Evaluation Metrics: Accuracy

		ACTUAL CLASS	
		Class=Yes	Class=No
PREDICTED	Class=Yes	a True Pos	b False Pos
	Class=No	c False Neg	d True Neg

- Most widely-used metric used extensively when classes are **balanced**:

$$\text{Accuracy} = \frac{a+d}{a+b+c+d} = \frac{TP+TN}{TP+TN+FP+FN}$$

The Accuracy Paradox

- Consider a tumor classification problem
  - Number of benign cases = 990
  - Number of malignant cases = 10
- Detecting the rare class is usually more interesting
  - e.g., frauds, intrusions, defects, etc.

		Actual	
		Malignant	Benign
Predicted	Malignant	0	0
	Benign	10	990

- If a model predicts everything to be benign, accuracy is?
  - 990/1000 = 99 %
  - This is **misleading** because this trivial **model does not detect any class YES example**
- **We need a better measures than Accuracy**

Precision and Recall

$$precision = \frac{TP}{TP + FP}$$

$$recall = \frac{TP}{TP + FN}$$

$$F1 = \frac{2 \times precision \times recall}{precision + recall}$$

$$accuracy = \frac{TP + TN}{TP + FN + TN + FP}$$

$$specificity = \frac{TN}{TN + FP}$$

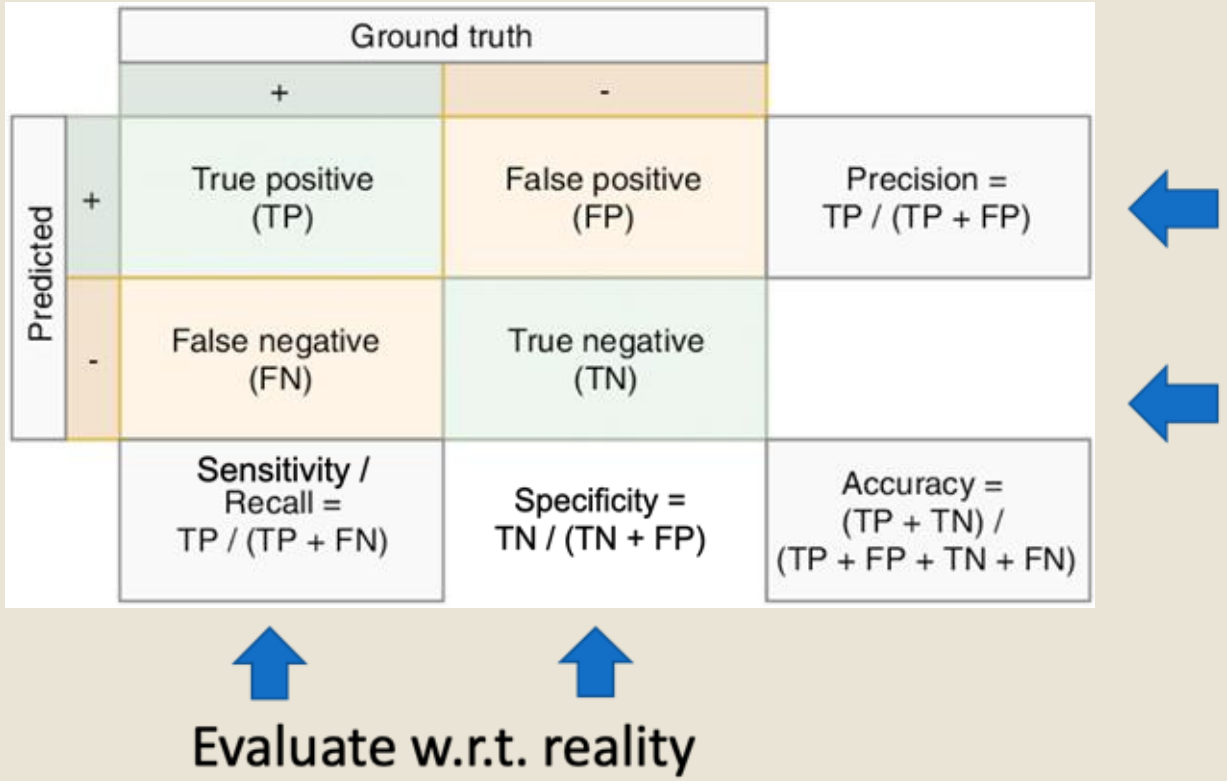

Confusion Matrix

		Ground truth +	Ground truth -	
Predicted	+	True positive (TP)	False positive (FP)	Precision = TP / (TP + FP)
Predicted	-	False negative (FN)	True negative (TN)	
		Recall = TP / (TP + FN)		Accuracy = (TP + TN) / (TP + FP + TN + FN)

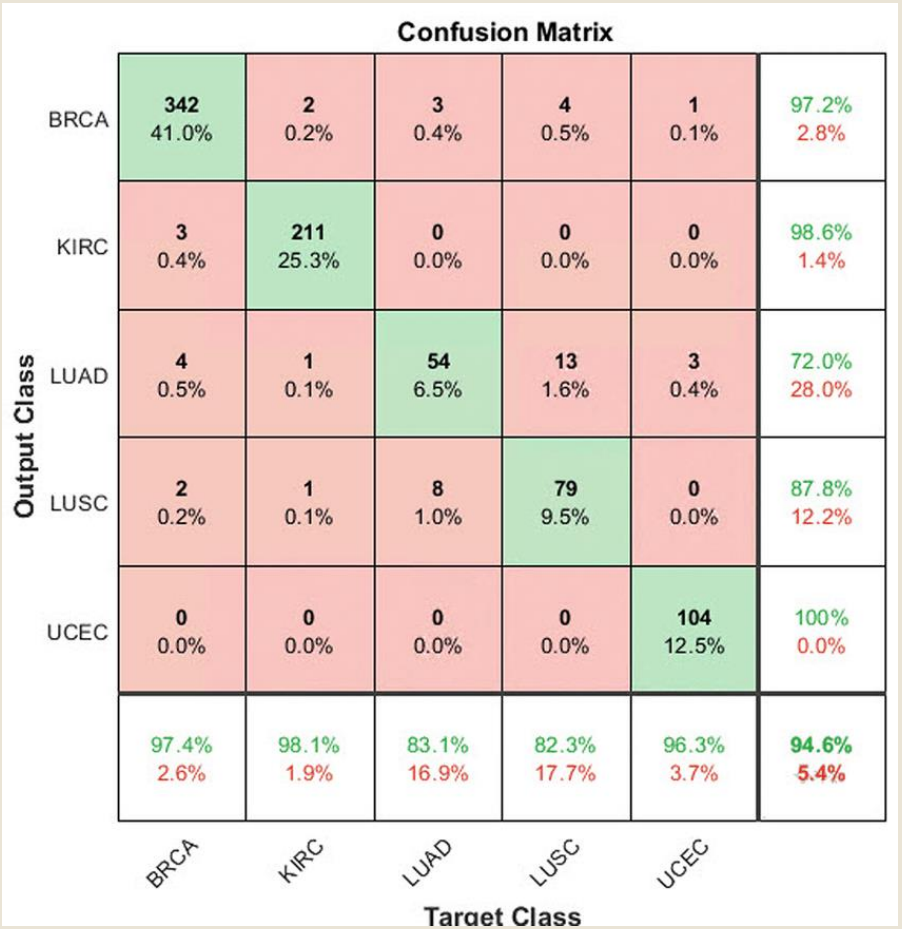

Generalization

- The ultimate goal of ML training is to **generalize well to unseen data**
- The training process need to **avoid overfitting and underfitting**

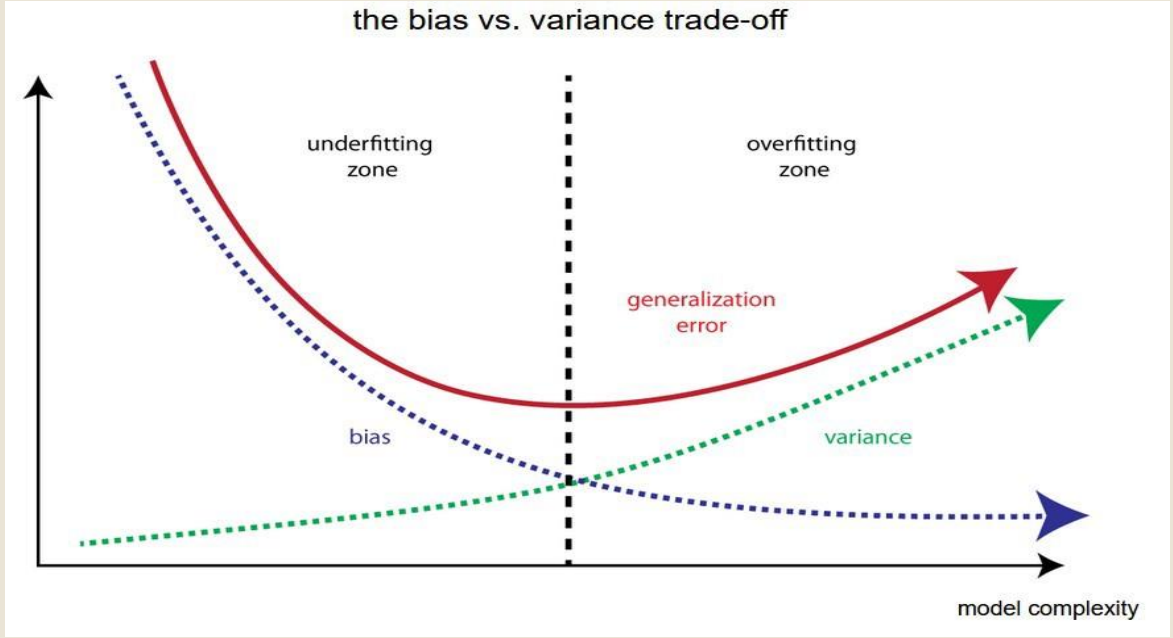

Underfitting vs Overfitting

- **Underfitting** (bias): Model is influenced by prior conception that limits its ability to learn from the training data.
- Reasons include:
  - Too simple model
  - Too little training

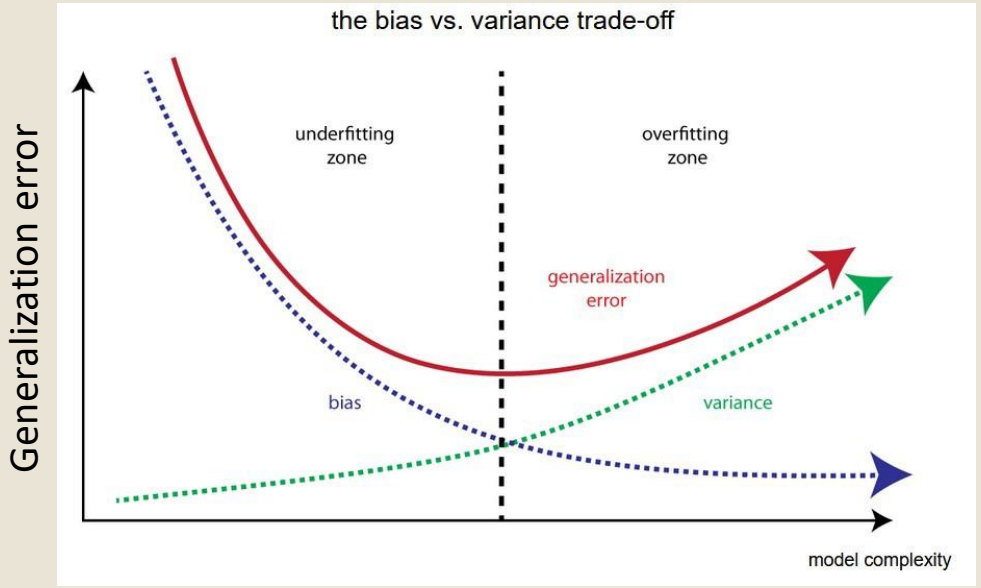

Underfitting vs Overfitting

- **Overfitting** (Variance): The model **attends too** much to the **details** of the **training data** which **limit its ability to generalize well.**
- Reasons include:
  - Too much training
  - Too complex model
  - Too noisy data

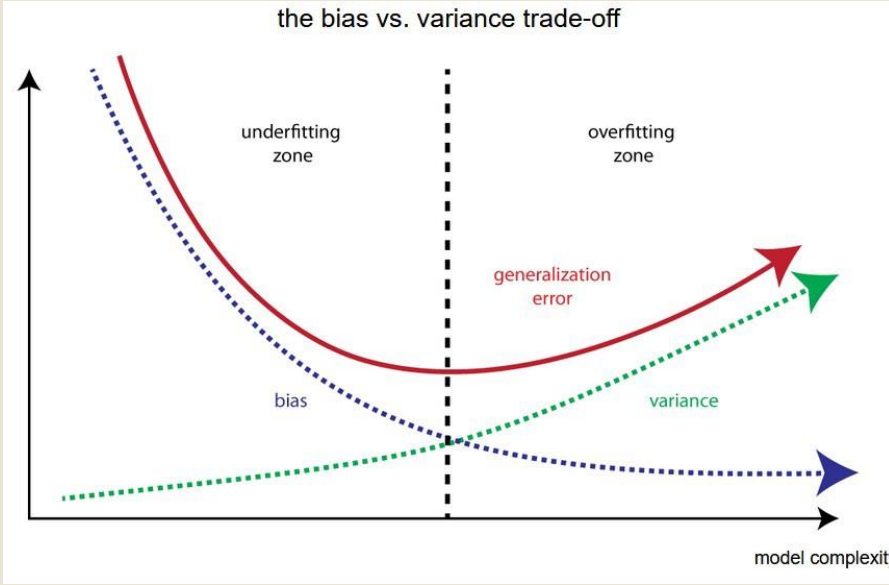

Underfitting vs Overfitting

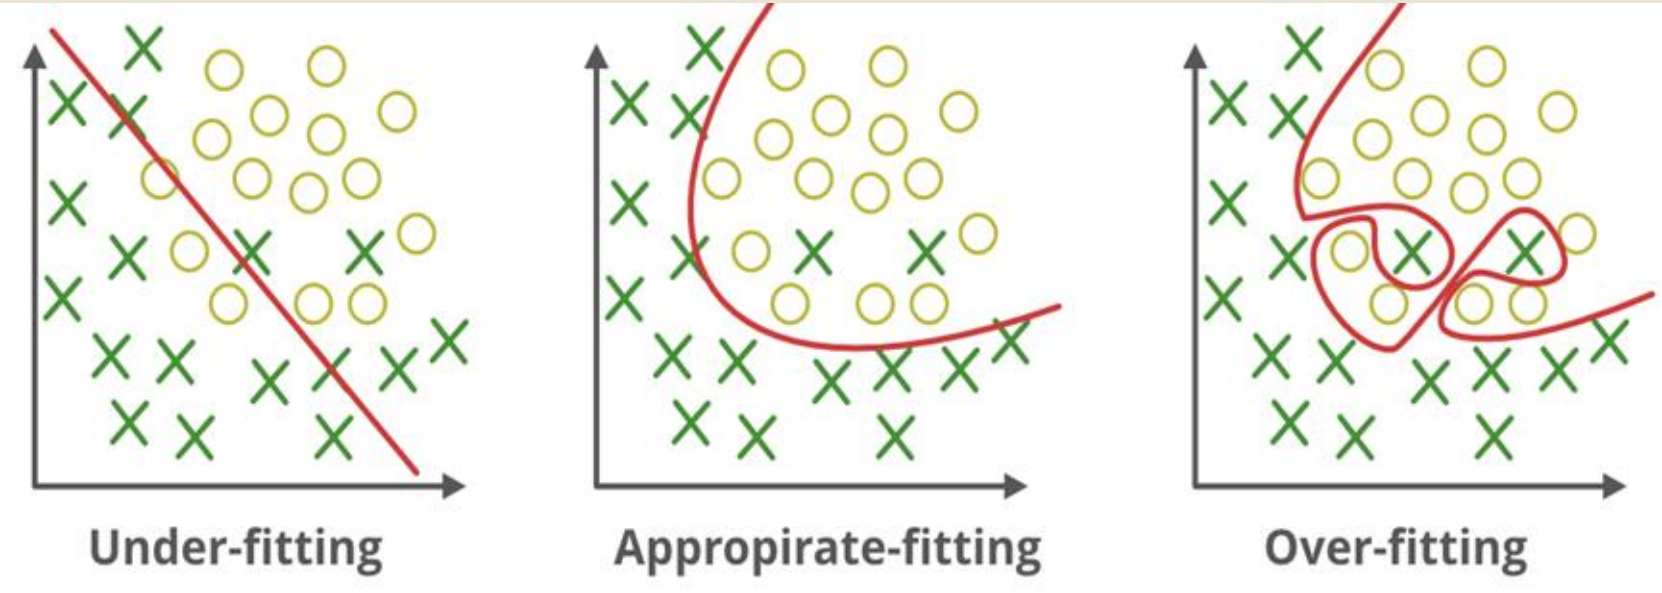

Learning Curves

- Analyze your model while training to detect and **prevent overfitting.**
  - e.g., use early stopping

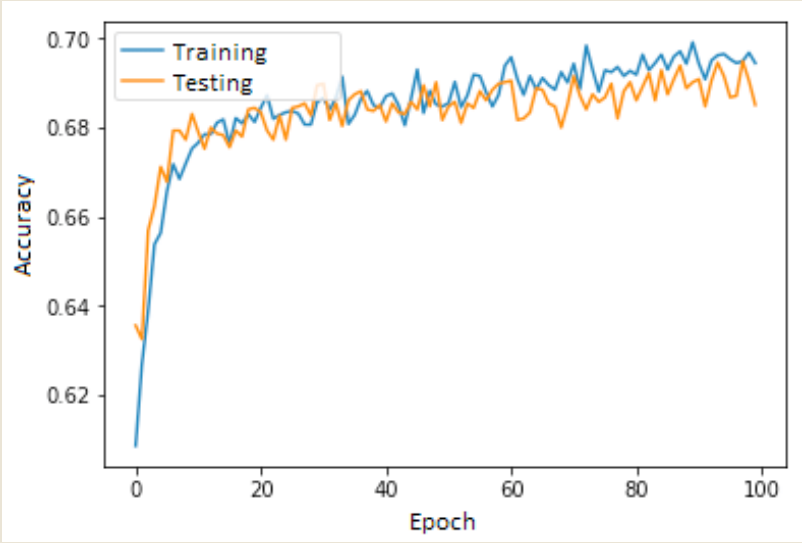

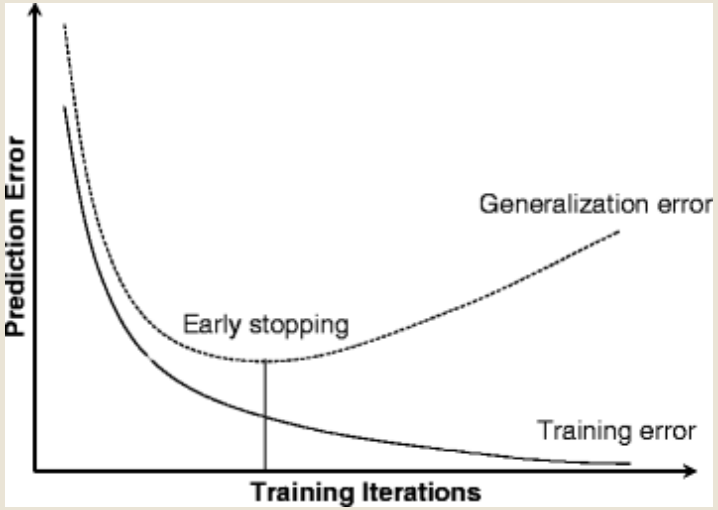

Model Tuning

- Goal: Speed up training and improve accuracy
- Learning rate (α)
  - Optimizes learning speed
- Regularization
  - Overcome overfitting and underfitting

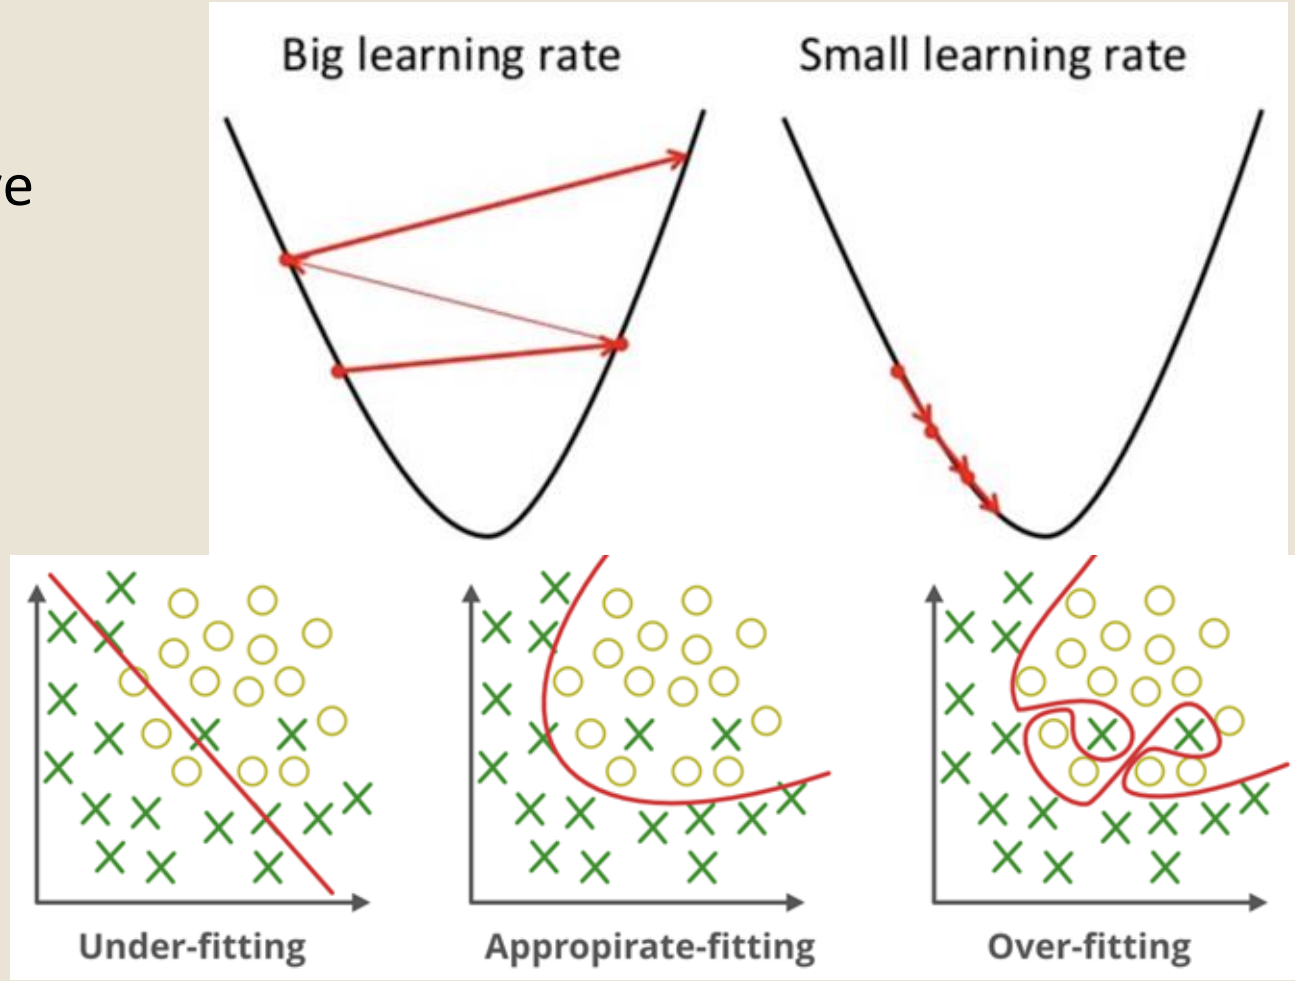

Best Practices and Common Mistakes

Problem Formulation

- Does your problem really need AI/ML?
  - May be some parts and others not
- Mapping your problem into the model components
  - E.g., features, states, performance measure, reward, transitions
- Is input available in real scenarios in the form you assume?
  - Adversary models

How to select your ML model

- Start simple. Don't start complex
  - Simple solutions are the hardest
  - Some simple models can provide the best results
    - SVMs, Random Forests variants like XGBoost, etc.
- Look for existing models
  - Image classifiers (VGG, resent, inception) and LLMs
- Adjust existing models
  - Transfer learning

Transfer Learning

- You don’t have to start from scratch
- Someone may have solved a problem similar to you
  - Borrow some of what it learned
  - Relearn the rest

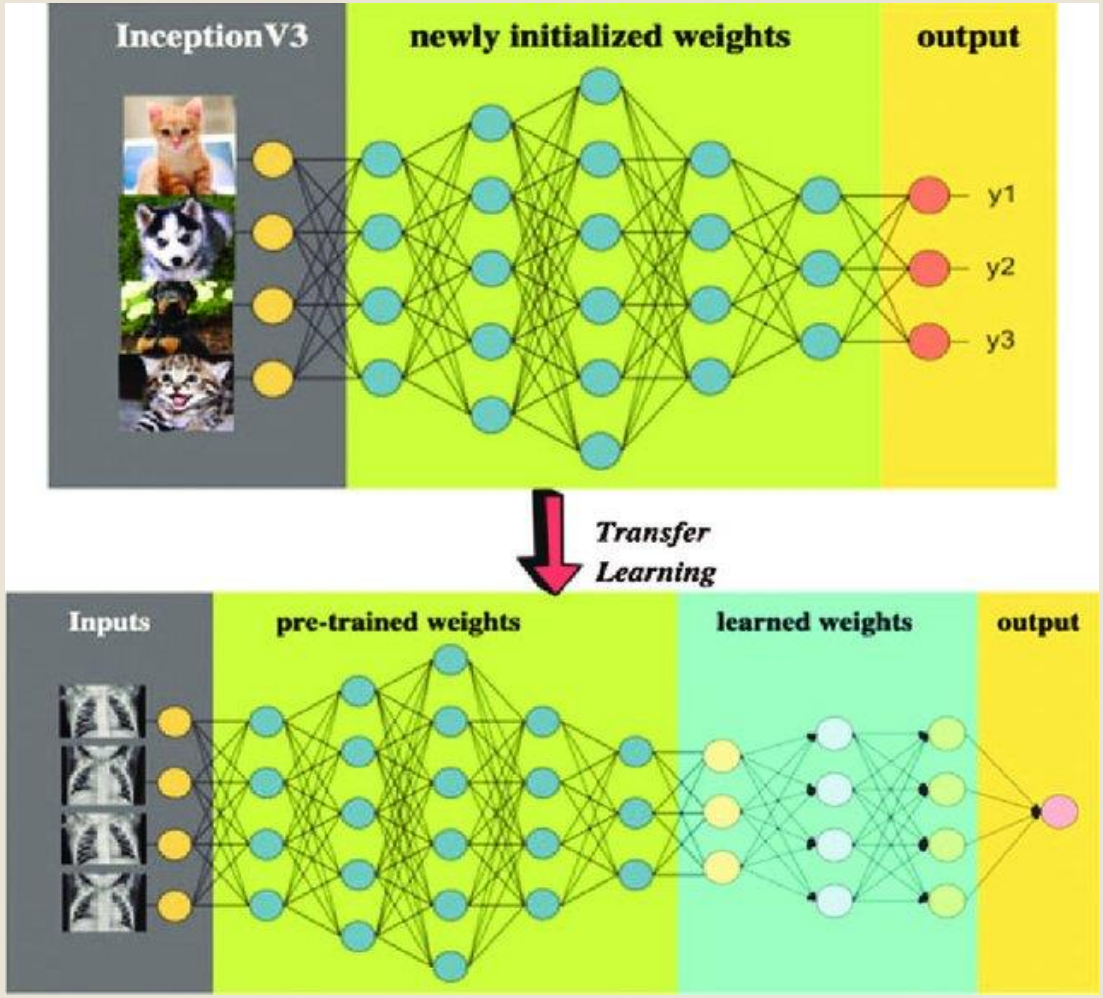

Have Data?

- Decides model to use
  - DL needs more data
- Online sources
  - Image repositories (Kaggle)
- Data augmentation (synthesized data)
  - Introduce distortions
  - Change color, orientation, size, etc.
  - Cropping

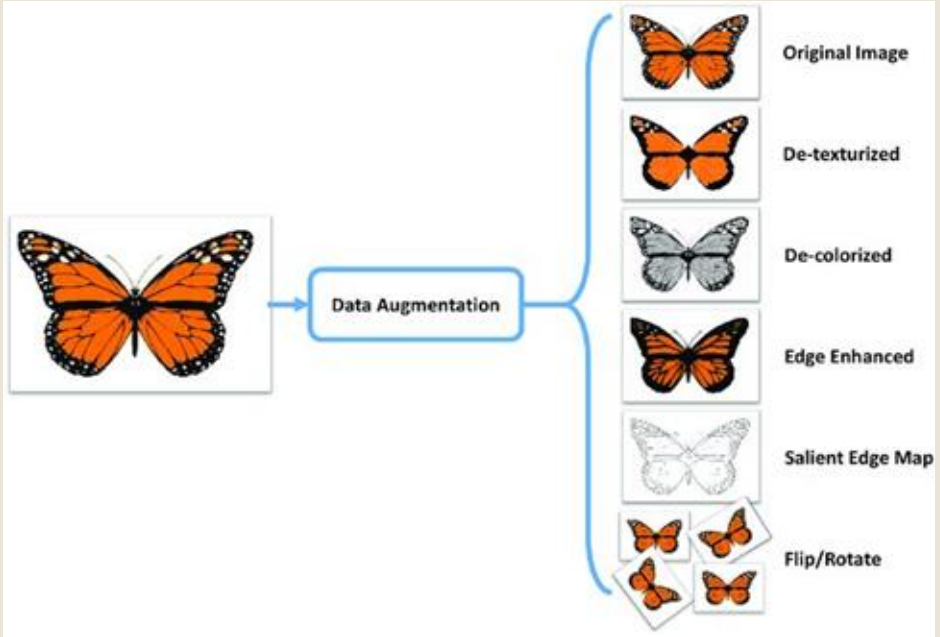

Data Contamination/Leakage

- Occurs when information from validation or test data leaks into training
- Common causes: improper data splitting (e.g., duplicates), preprocessing **before** splitting, or tuning on test data
- Leads to **overly optimistic** performance estimates
- Model may **fail to generalize** to truly unseen data
- Violates the principle of a **fair**, **unbiased evaluation**
- Prevent by maintaining strict separation of train, validation, and test pipelines

Machine Learning is Science and Art

- Tuning NNs is an iterative process
- There are common practices in HPO
- Try to use ONE main evaluation metric:
  - F1, accuracy, precision, recall
  - Others for discussion

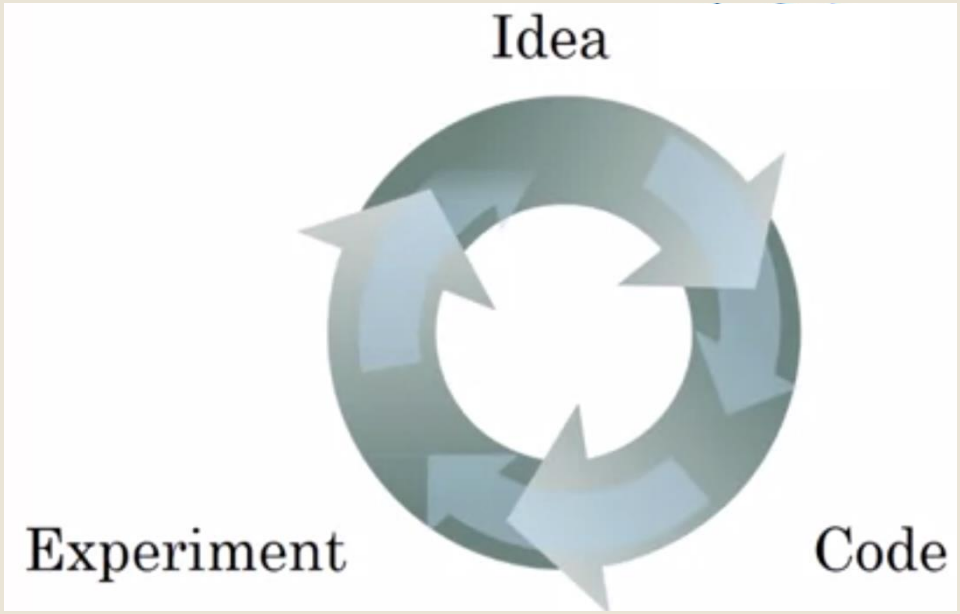

AI Risks and Ethics

Generation of toxic drugs using AI

- Nature 2022: [Dual use of artificial-intelligence-powered drug discovery]
- “in **less than 6 hours** after starting on our in-house server, our model generated **40,000 molecules** that scored within our desired threshold. In the process, the AI designed not only VX, but also many other known **chemical warfare agents** that we identified through visual confirmation with structures in public chemistry databases.”

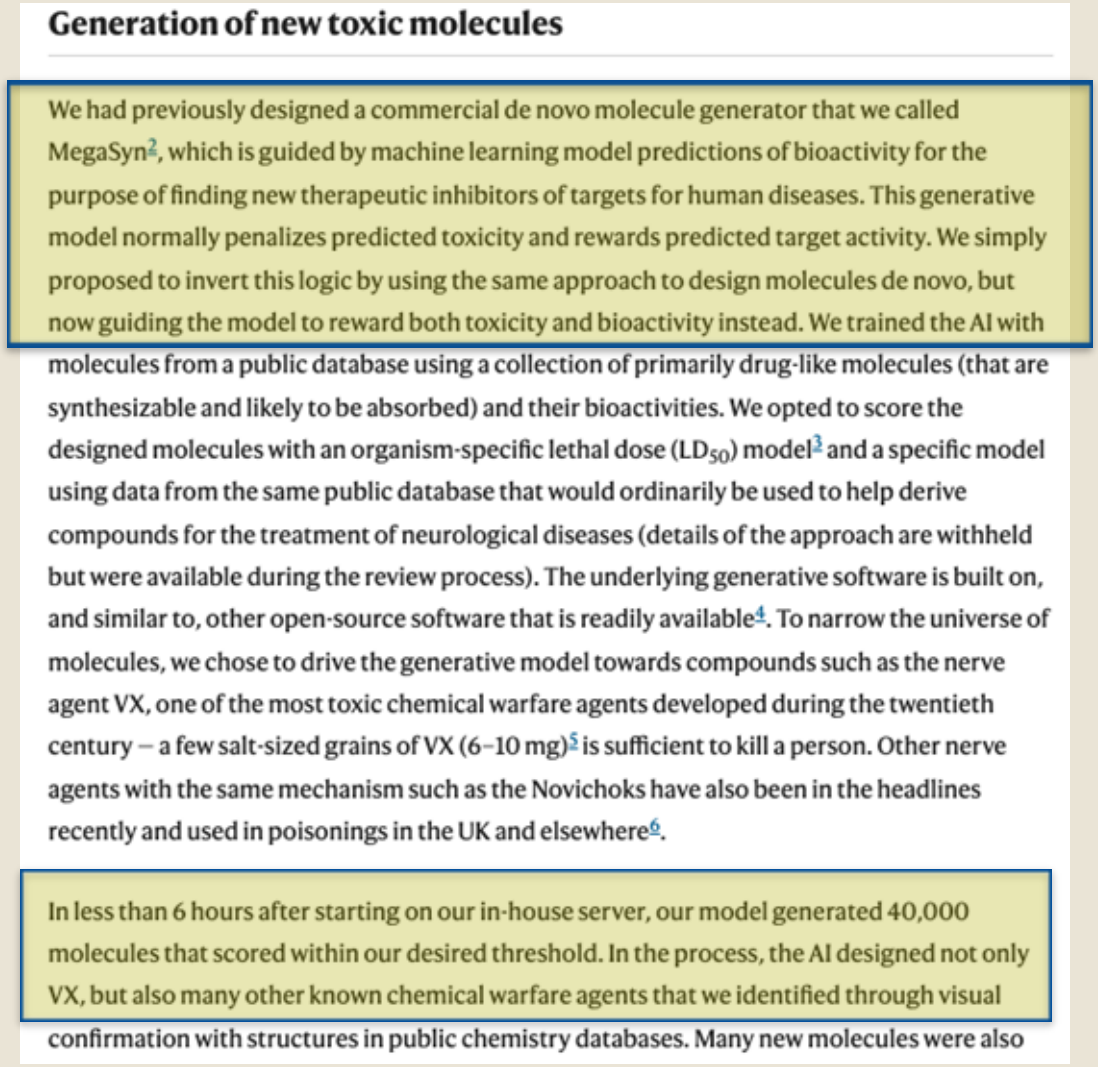

Generation of new toxic molecules

We had previously designed a commercial de novo molecule generator that we called MegaSyn[2], which is guided by machine learning model predictions of bioactivity for the purpose of finding new therapeutic inhibitors of targets for human diseases. This generative model normally penalizes predicted toxicity and rewards predicted target activity. We simply proposed to invert this logic by using the same approach to design molecules de novo, but now guiding the model to reward both toxicity and bioactivity instead. We trained the AI with molecules from a public database using a collection of primarily drug-like molecules (that are synthesizable and likely to be absorbed) and their bioactivities. We opted to score the designed molecules with an organism-specific lethal dose (LD_{50}) model[3] and a specific model using data from the same public database that would ordinarily be used to help derive compounds for the treatment of neurological diseases (details of the approach are withheld but were available during the review process). The underlying generative software is built on, and similar to, other open-source software that is readily available[4]. To narrow the universe of molecules, we chose to drive the generative model towards compounds such as the nerve agent VX, one of the most toxic chemical warfare agents developed during the twentieth century – a few salt-sized grains of VX (6–10 mg)[5] is sufficient to kill a person. Other nerve agents with the same mechanism such as the Novichoks have also been in the headlines recently and used in poisonings in the UK and elsewhere[6].

In less than 6 hours after starting on our in-house server, our model generated 40,000 molecules that scored within our desired threshold. In the process, the AI designed not only VX, but also many other known chemical warfare agents that we identified through visual confirmation with structures in public chemistry databases. Many new molecules were also