



# Master

# Reinforcement Learning 2022

## Lecture B:

# A Summary of Machine Learning

Aske Plaat



# What is Intelligence?

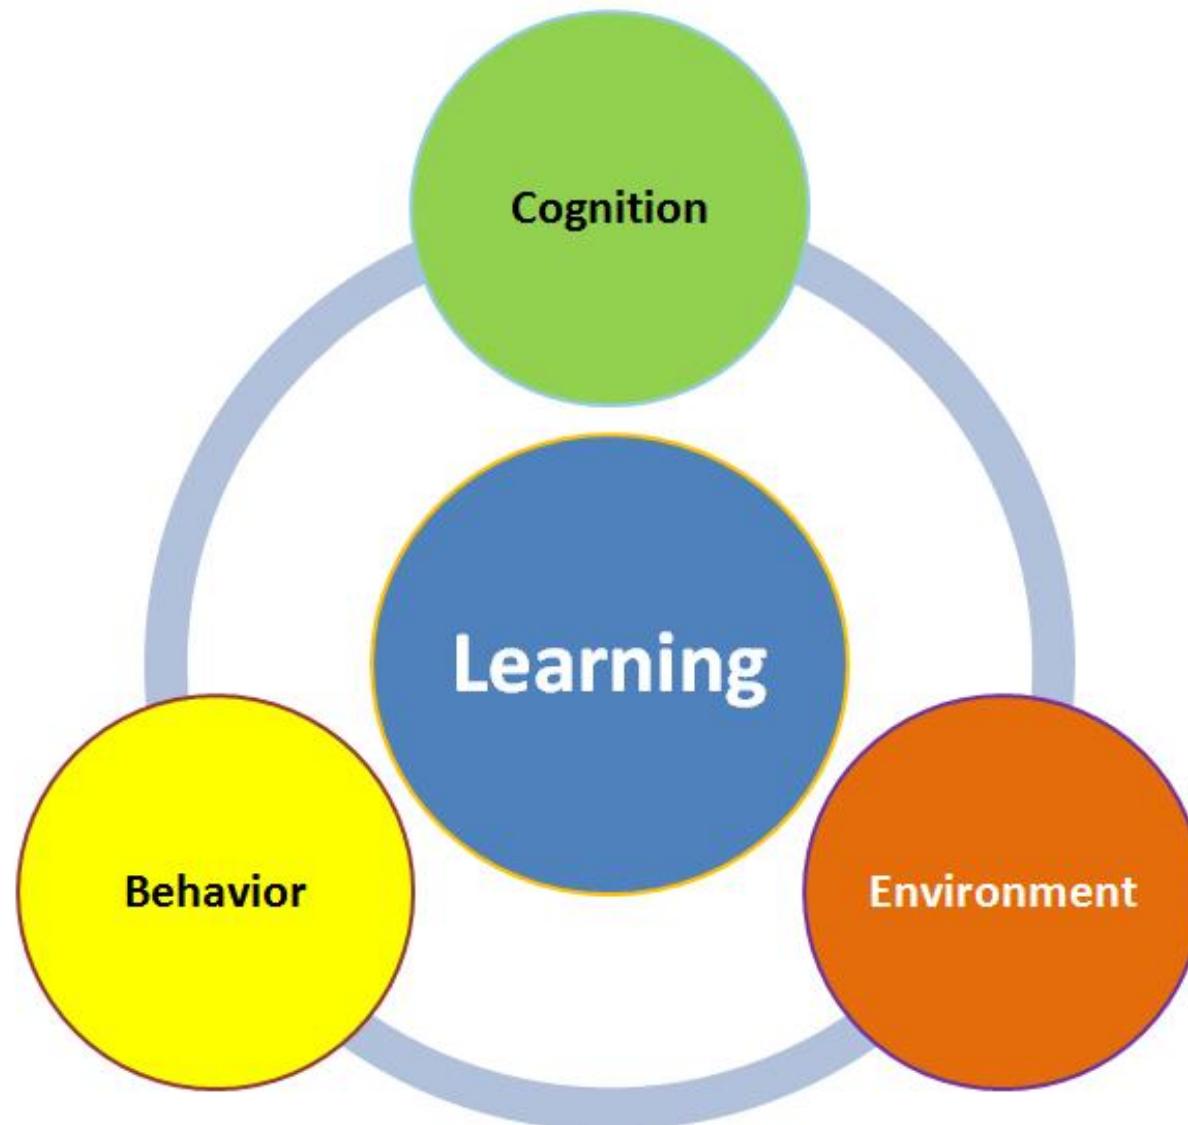
## Cognition

- Recognition
- Memory
- Logical Reasoning
- Learning
- Adaptive Behavior

## Behavior

- Creativity
- Intuition
- Free Will
- Self Awareness
- Consciousness

# What is Learning?



# A very short summary of Machine Learning and Deep Learning

## Traditional Programming



## Machine Learning

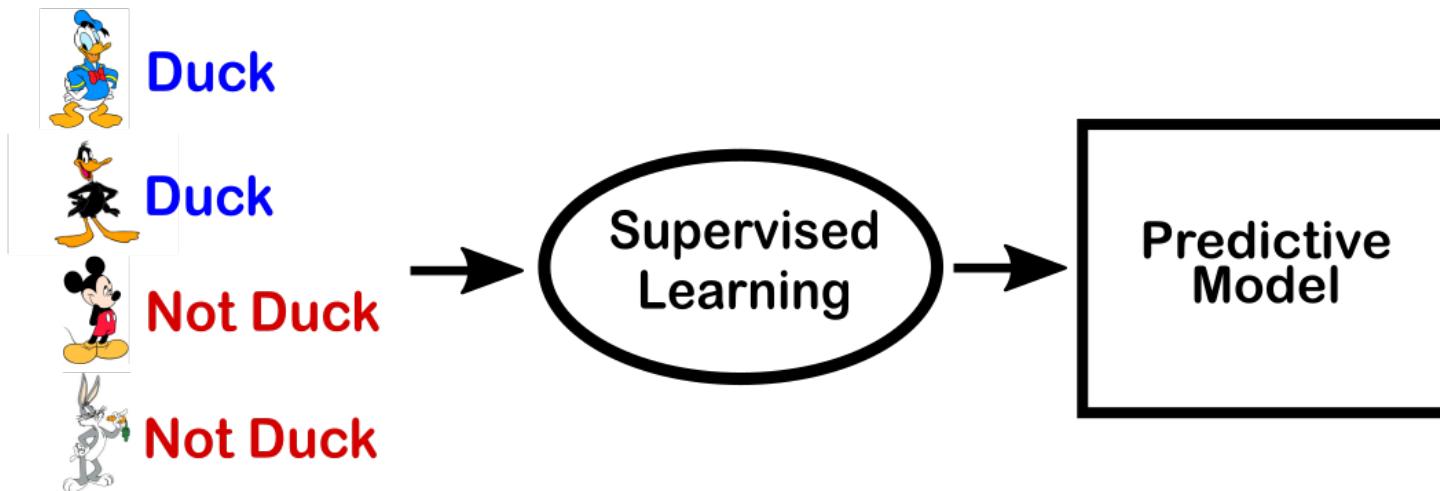


# Machine Learning

- Supervised Learning - learning from example & label
- Reinforcement Learning - learning from interaction number
- Unsupervised Learning - learning inherent relations

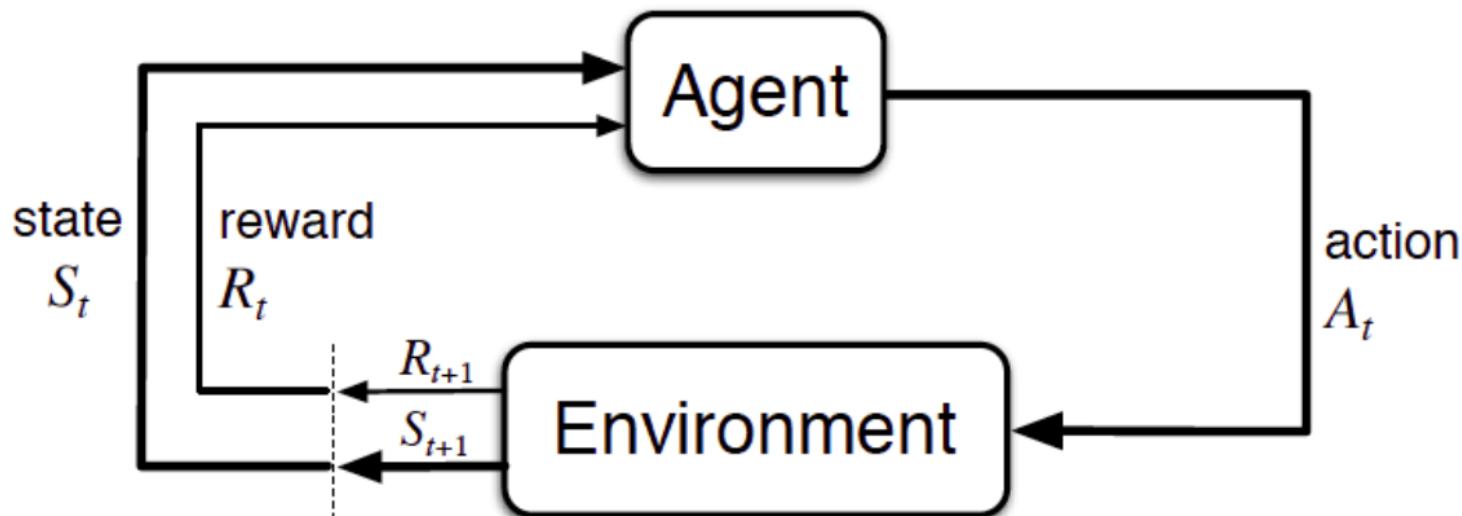
# Supervised Learning

- Is learning by example
- Database learning
- (Image, label) -> correct?



# Reinforcement Learning

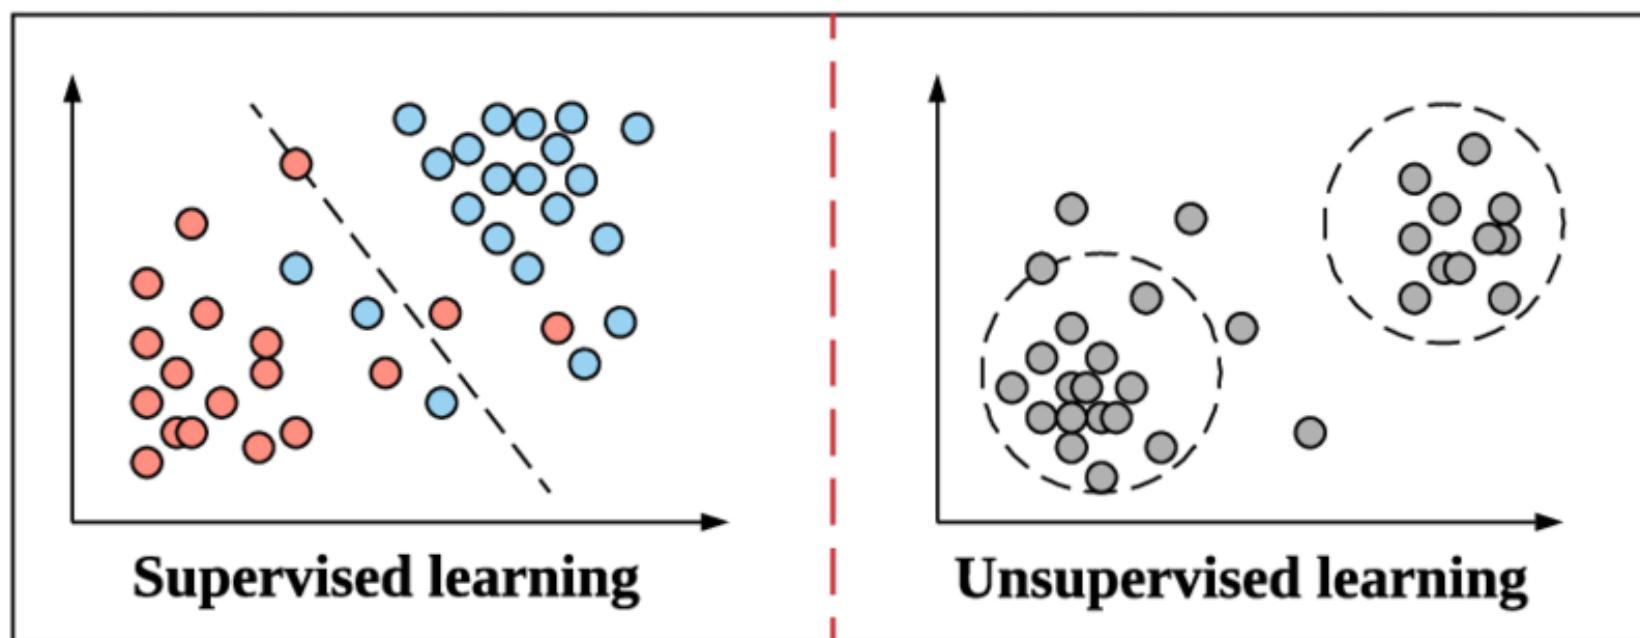
- is learning by interaction
- (state, action) -> reward number

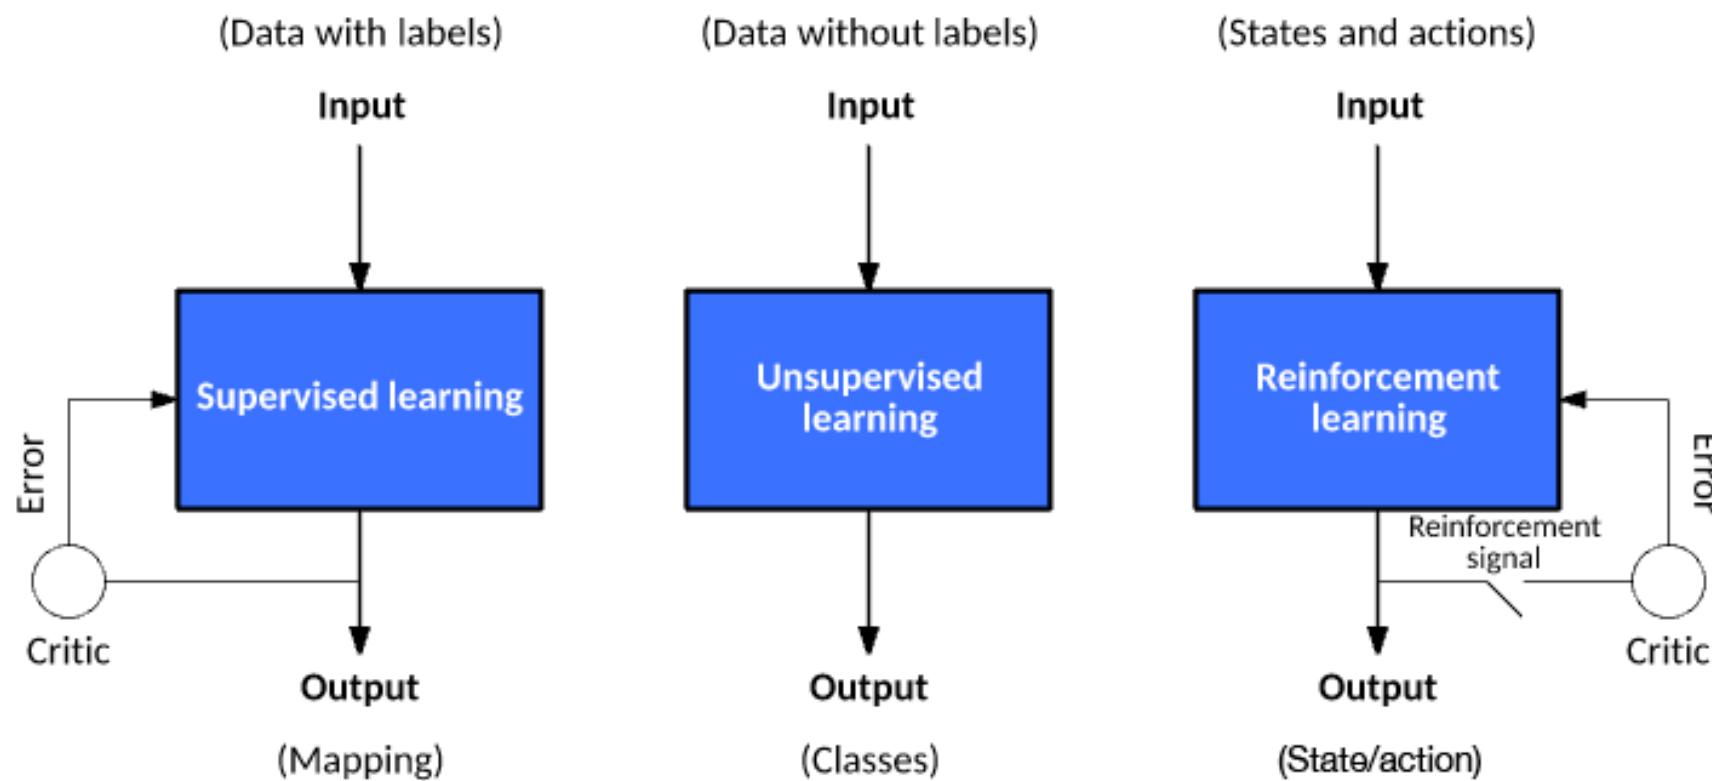


The agent-environment interaction in reinforcement learning. (Source: Sutton and Barto, 2017)

# Unsupervised Learning

- Is learning without examples, from inherent measures
- Database learning
- Clustering, data compression, dimensionality reduction





# Classical Machine Learning

Task Driven

## Supervised Learning

( Pre Categorized Data )



### Classification

( Divide the socks by Color )

Eg. Identity Fraud Detection



### Regression

( Divide the Ties by Length )

Eg. Market Forecasting

Data Driven

## Unsupervised Learning

( Unlabelled Data )



### Clustering

( Divide by Similarity )

Eg. Targeted Marketing



### Association

( Identify Sequences )

Eg. Customer Recommendation



### Dimensionality Reduction

( Wider Dependencies )

Eg. Big Data Visualization

Obj: Predictions & Predictive Models

Pattern/ Structure Recognition



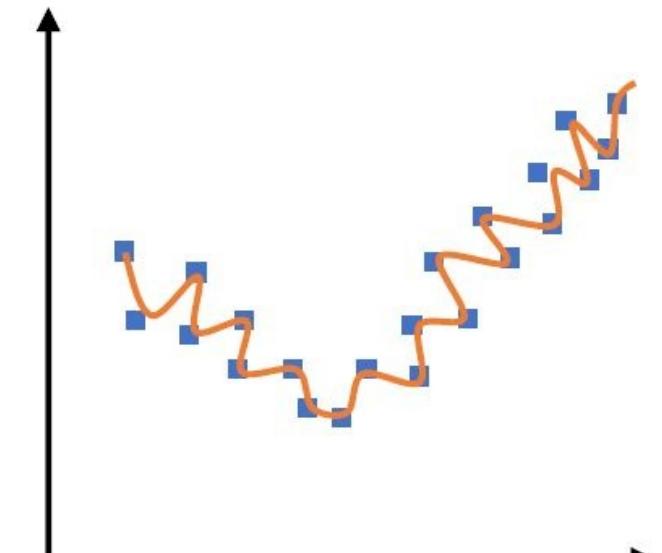
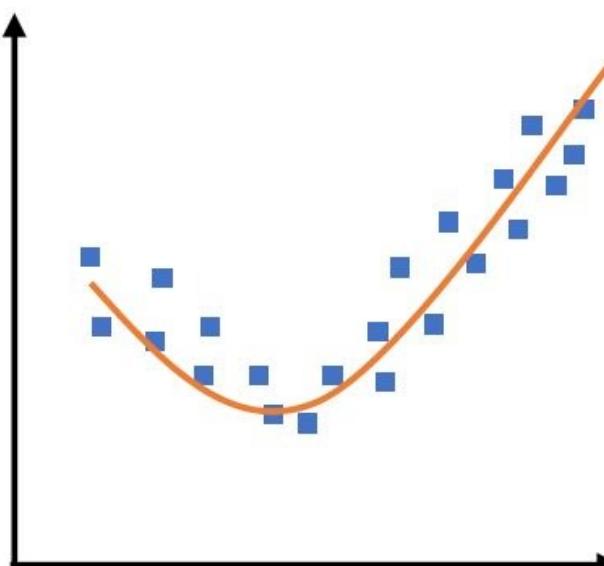
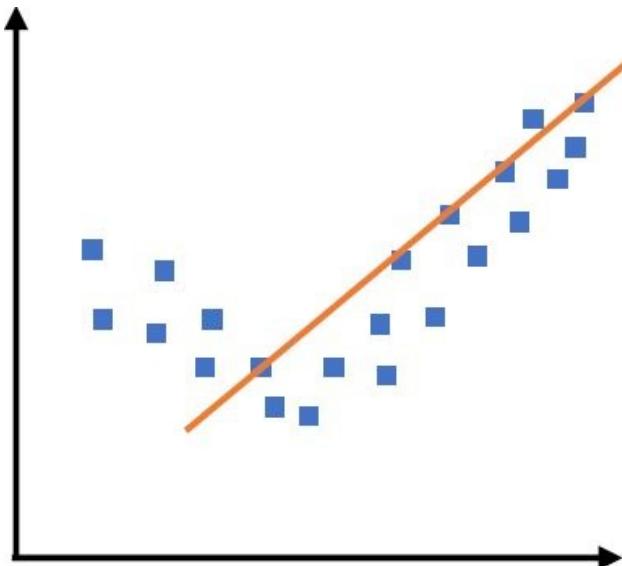
# Machine learning models cheat sheet

Supervised learning	Unsupervised learning	Semi-supervised learning	Reinforcement learning
<p>Data scientists provide input, output and feedback to build model (as the definition)</p> <p><b>EXAMPLE ALGORITHMS:</b></p> <p><b>Linear regressions</b></p> <ul style="list-style-type: none"><li>sales forecasting</li><li>risk assessment</li></ul> <p><b>Support vector machines</b></p> <ul style="list-style-type: none"><li>image classification</li><li>financial performance comparison</li></ul> <p><b>Decision tree</b></p> <ul style="list-style-type: none"><li>predictive analytics</li><li>pricing</li></ul>	<p>Use deep learning to arrive at conclusions and patterns through unlabeled training data.</p> <p><b>EXAMPLE ALGORITHMS:</b></p> <p><b>Apriori</b></p> <ul style="list-style-type: none"><li>sales functions</li><li>word associations</li><li>searcher</li></ul> <p><b>K-means clustering</b></p> <ul style="list-style-type: none"><li>performance monitoring</li><li>searcher intent</li></ul>	<p>Builds a model through a mix of labeled and unlabeled data, a set of categories, suggestions and exampled labels.</p> <p><b>EXAMPLE ALGORITHMS:</b></p> <p><b>Generative adversarial networks</b></p> <ul style="list-style-type: none"><li>audio and video manipulation</li><li>data creation</li></ul> <p><b>Self-trained Naïve Bayes classifier</b></p> <ul style="list-style-type: none"><li>natural language processing</li></ul>	<p>Self-interpreting but based on a system of rewards and punishments learned through trial and error, seeking maximum reward.</p> <p><b>EXAMPLE ALGORITHMS:</b></p> <p><b>Q-learning</b></p> <ul style="list-style-type: none"><li>policy creation</li><li>consumption reduction</li></ul> <p><b>Model-based value estimation</b></p> <ul style="list-style-type: none"><li>linear tasks</li><li>estimating parameters</li></ul>

# Machine Learning

## Learning a Function

- Fit a function to data (x, y) items
- Iterative minimization process of error between learned function  $f$  and data labels
- regression loss:  $L = \text{avg}(y - f)^2$
- categorization loss:  $L = -\text{sum}(y \log f)$



# Accuracy

- True Positives = # classified correctly as in class
- False Positives = # classified incorrectly as in class
- True Negatives = # classified correctly as not in class
- False Negatives = # classified incorrectly as not in class
- Accuracy = Trues/All =  $(TP+TN) / (TP+TN+FP+FN)$   
How well correct elements are predicted

# Confusion Matrix

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

# Function Approximation

- Small or Large data set?
- Exact Memoization: Small
- Feature Learning/Generalization: Large

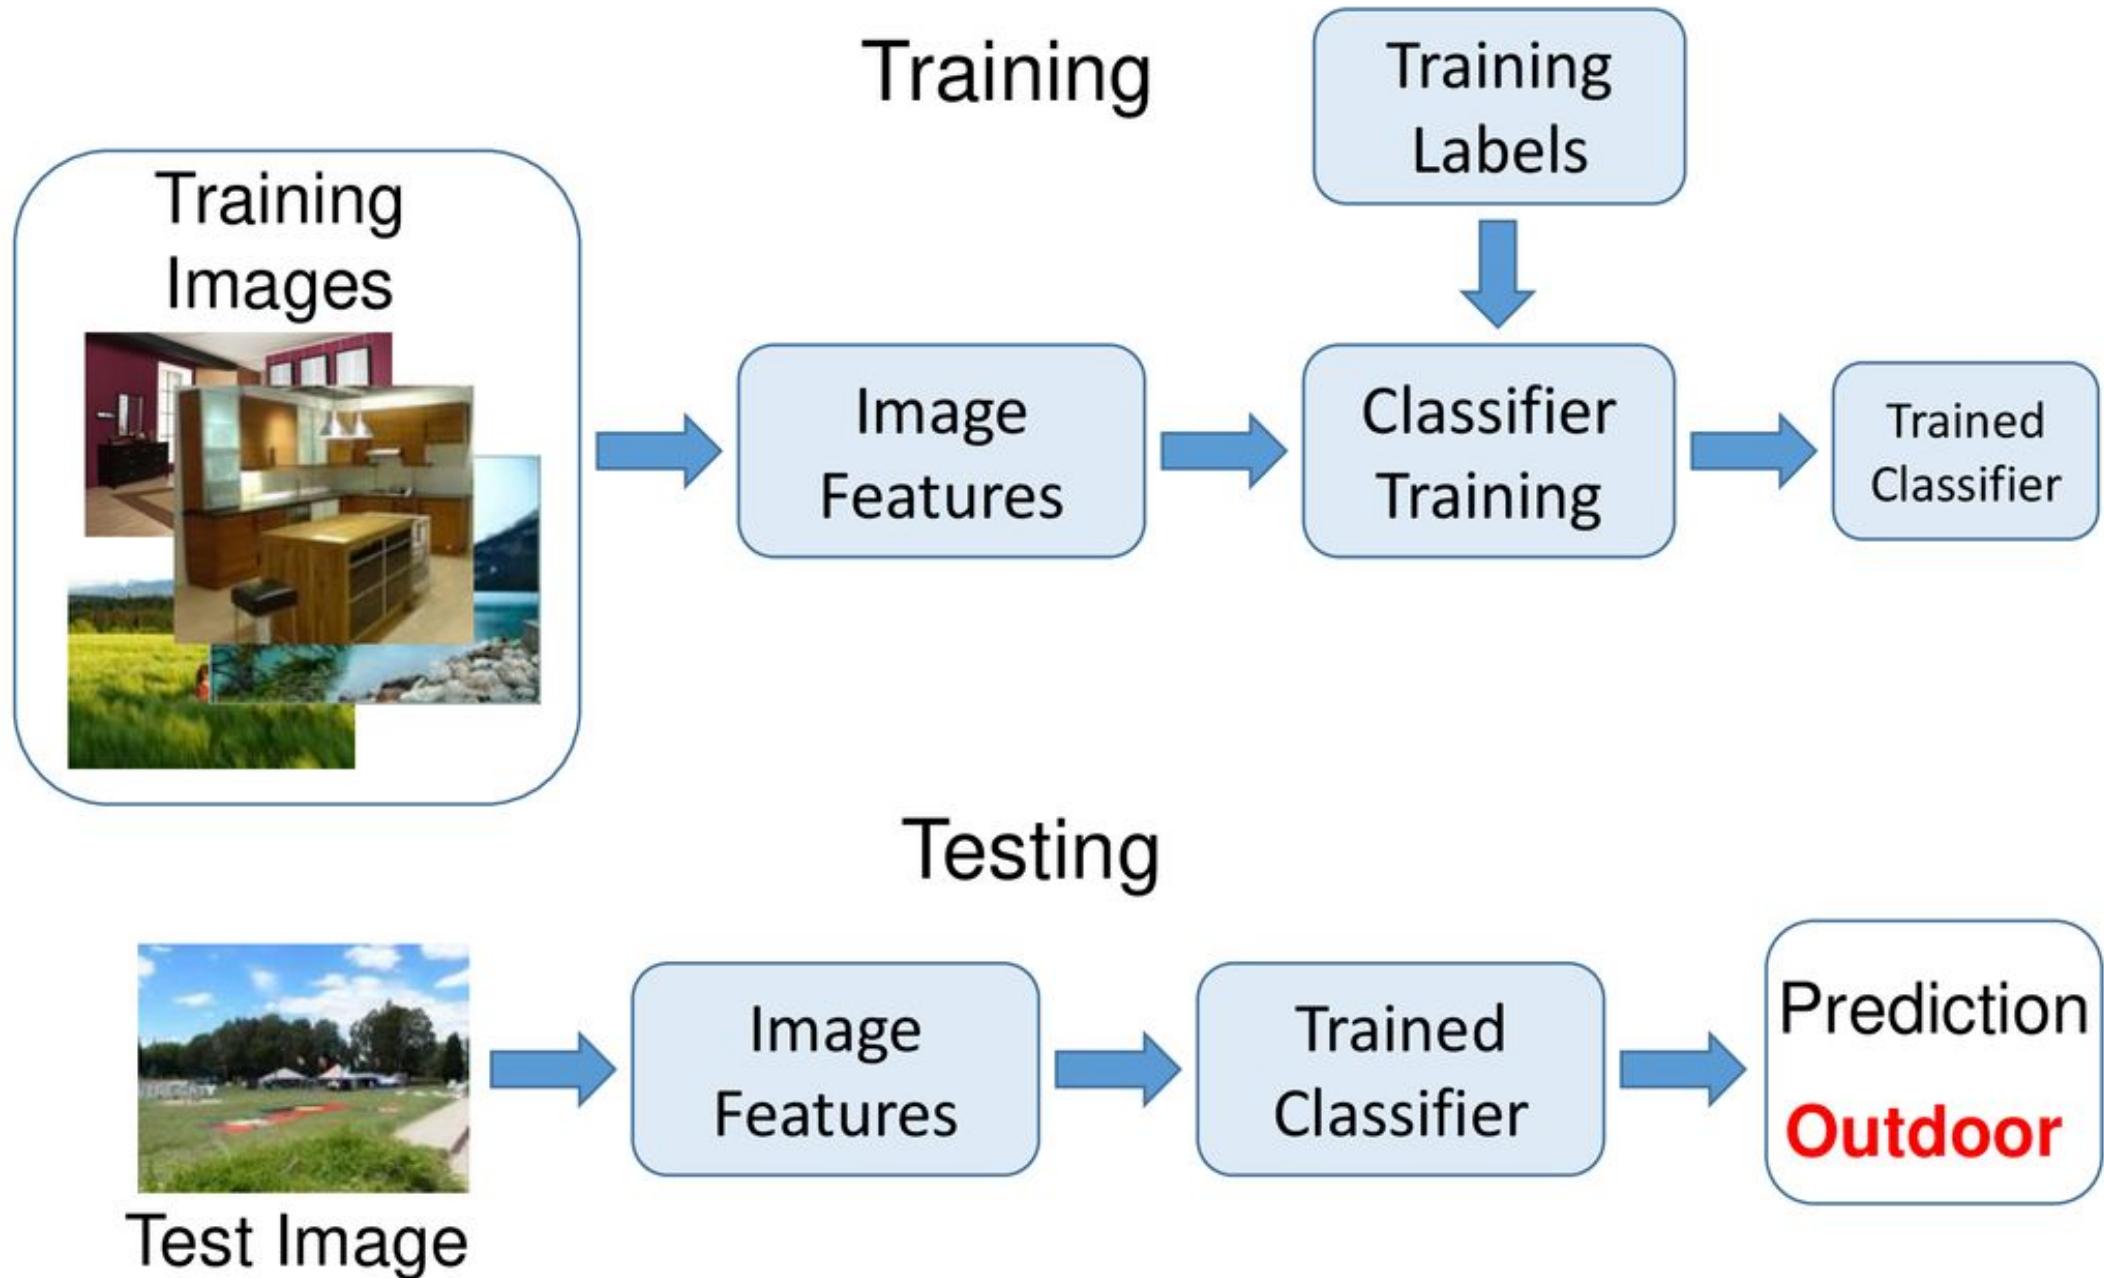
# What is a generalization?

- A generalization is a broad statement about a group of people, things, or ideas.
- It states something they have in common.

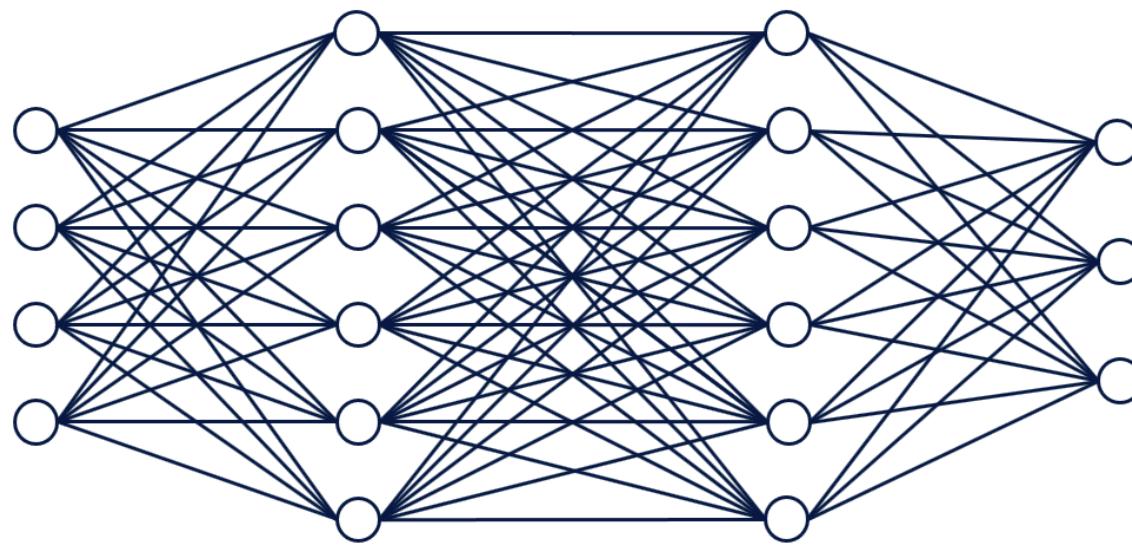
# Generalization

- Split data set D into Training set and Test set
- Train on Training set
- Test on Test set
- An algorithm generalizes well from training to test if the accuracy at test time is about as good as at training time

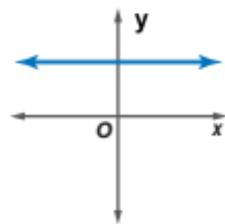
# Image Categorization



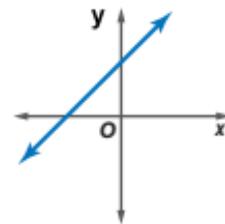
# Model Complexity



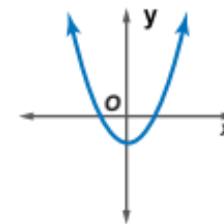
Constant function  
Degree 0



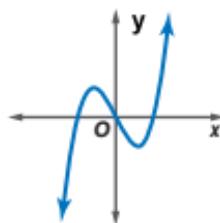
Linear function  
Degree 1



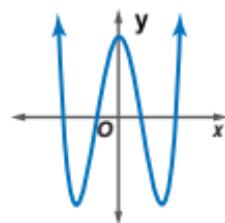
Quadratic function  
Degree 2



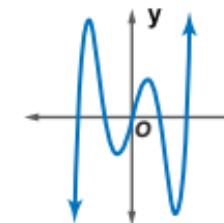
Cubic function  
Degree 3



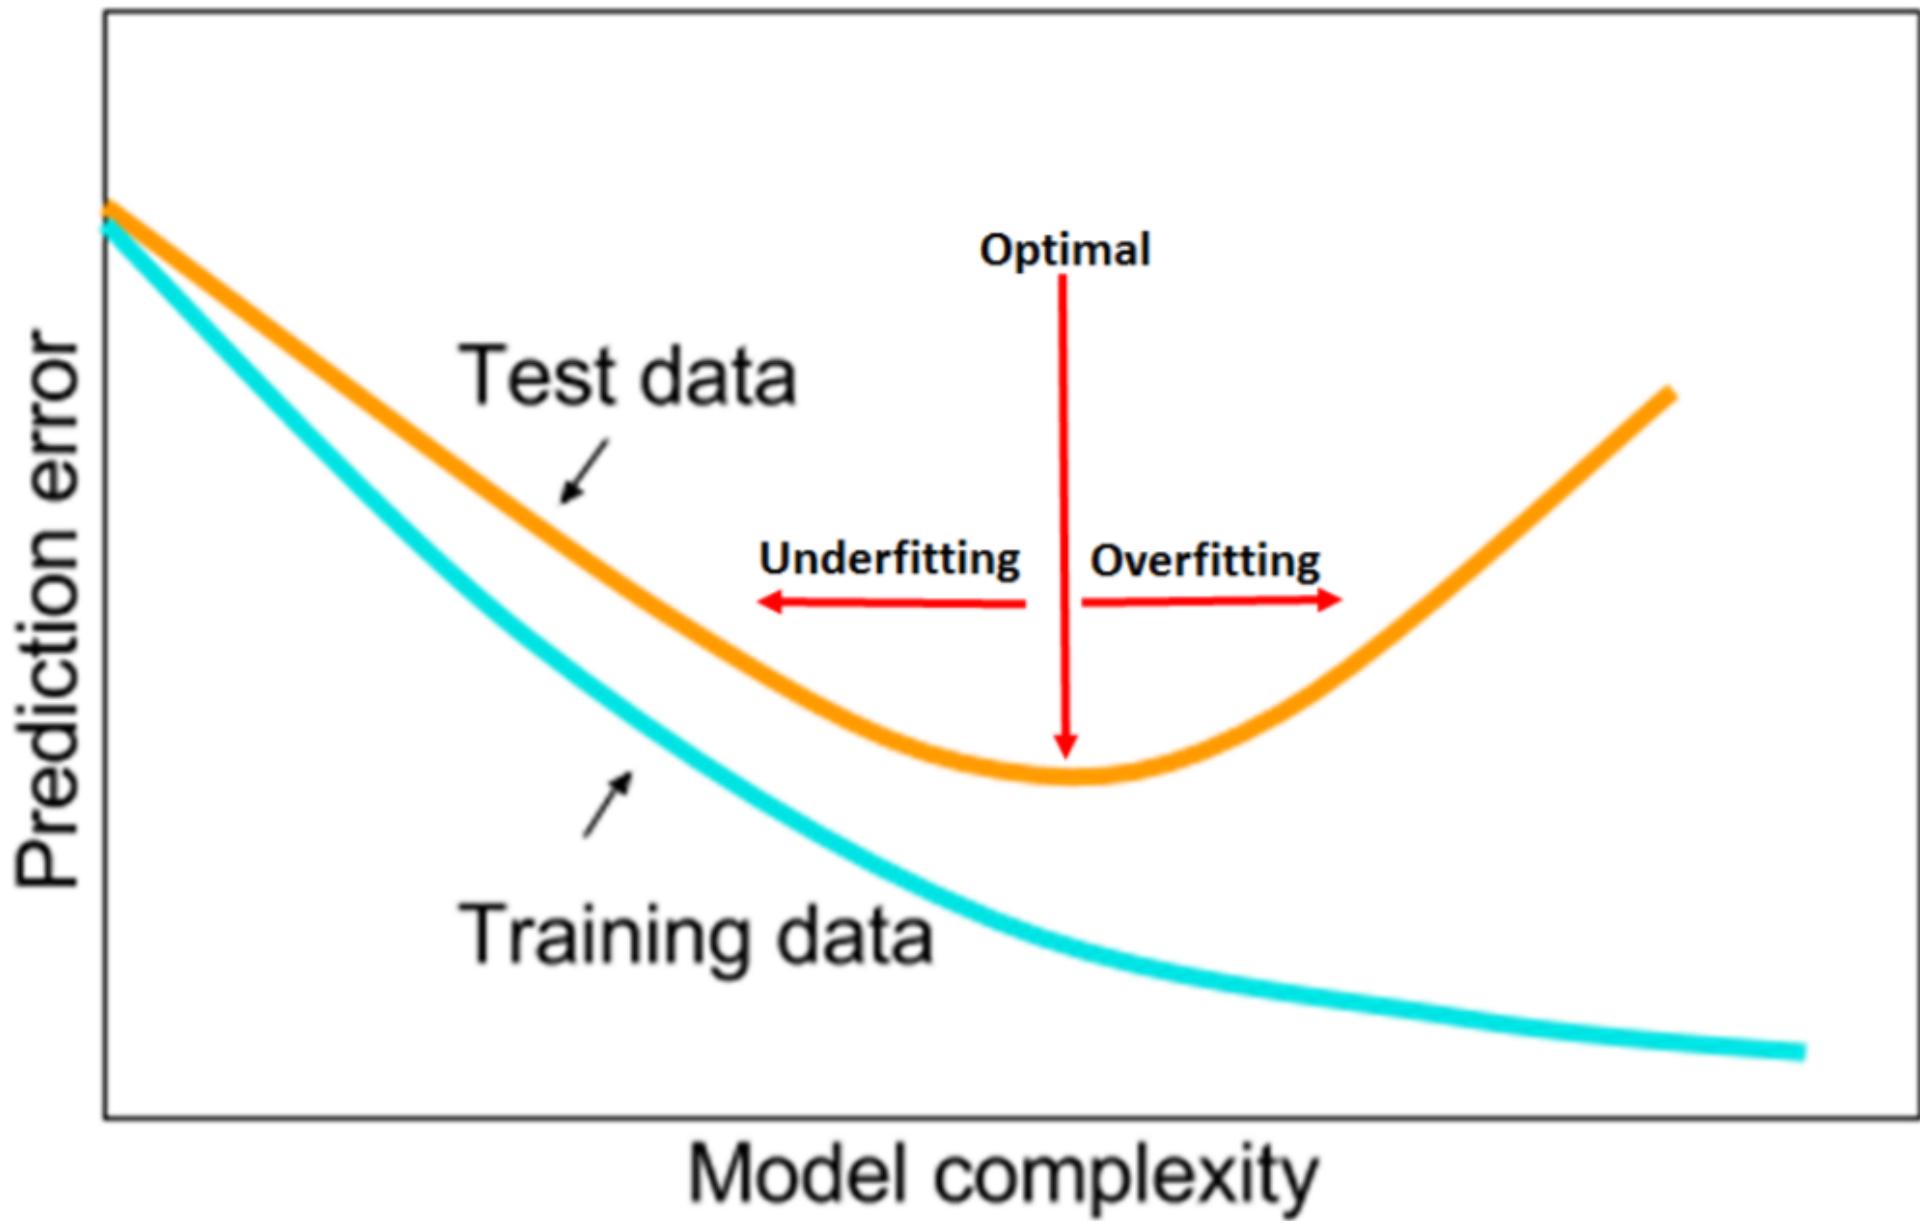
Quartic function  
Degree 4



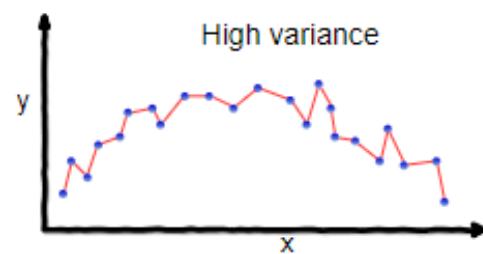
Quintic function  
Degree 5



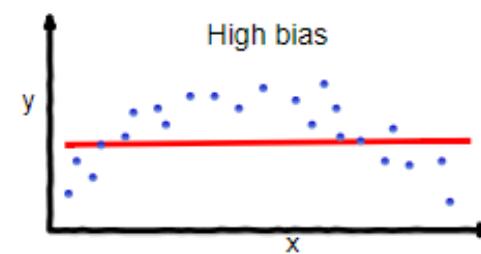
# Capacity



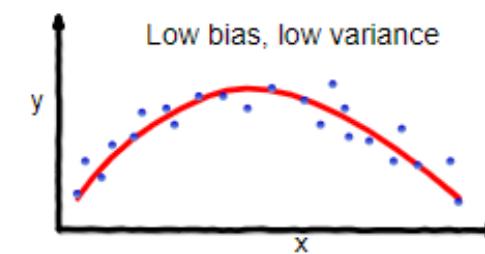
# Bias-Variance



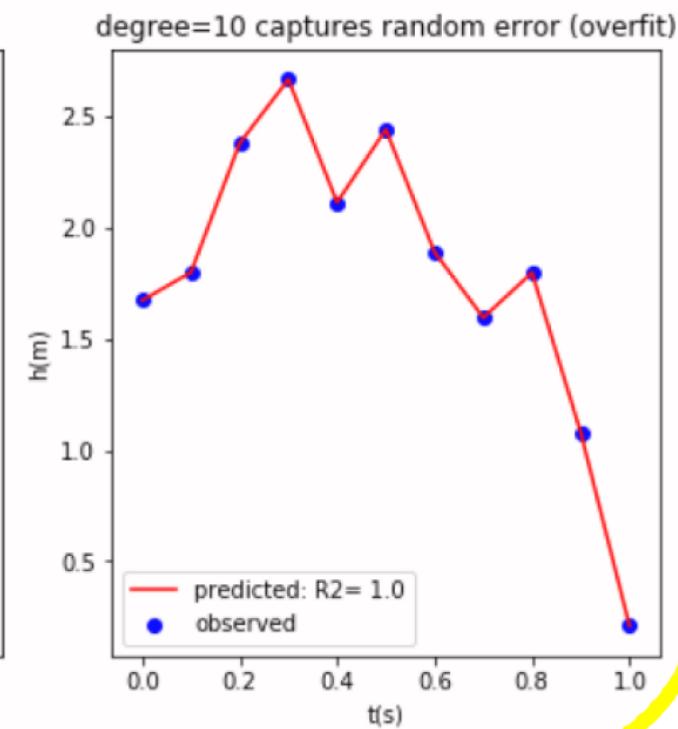
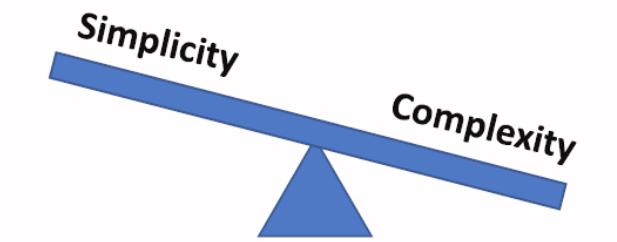
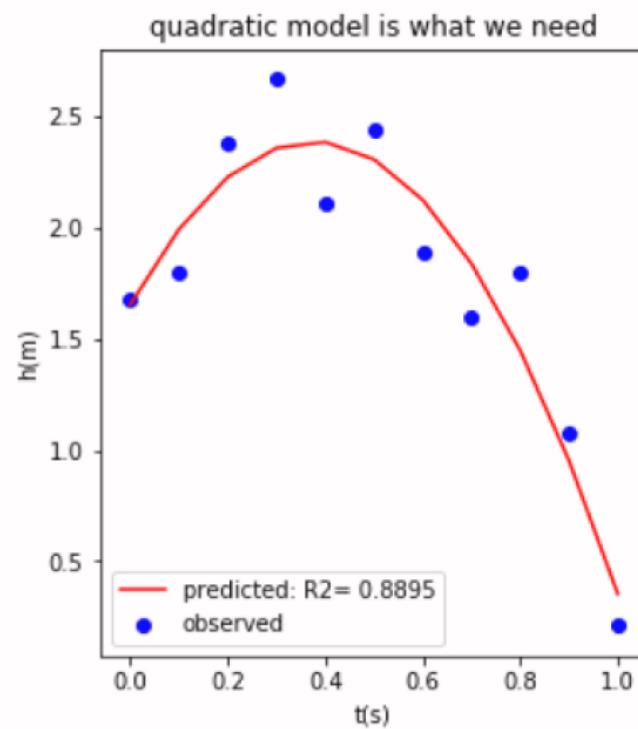
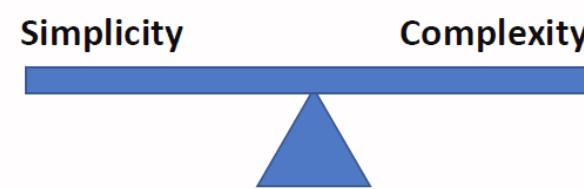
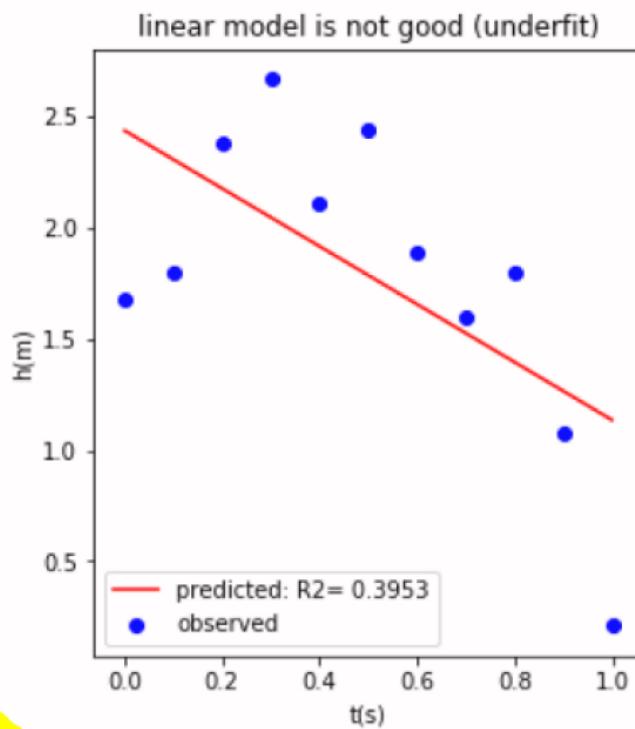
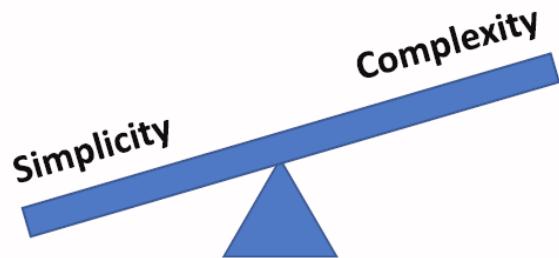
overfitting



underfitting

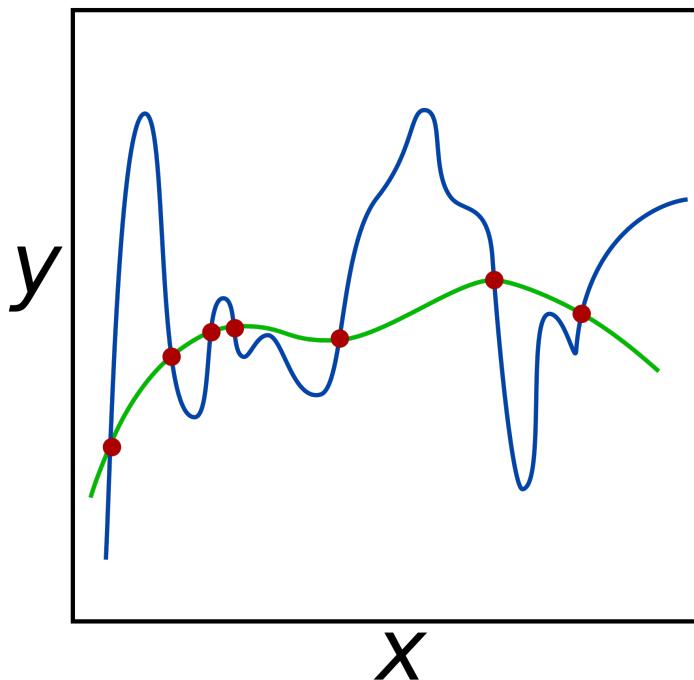


Good balance



# Regularization

- The world is smooth
- Penalize large weights

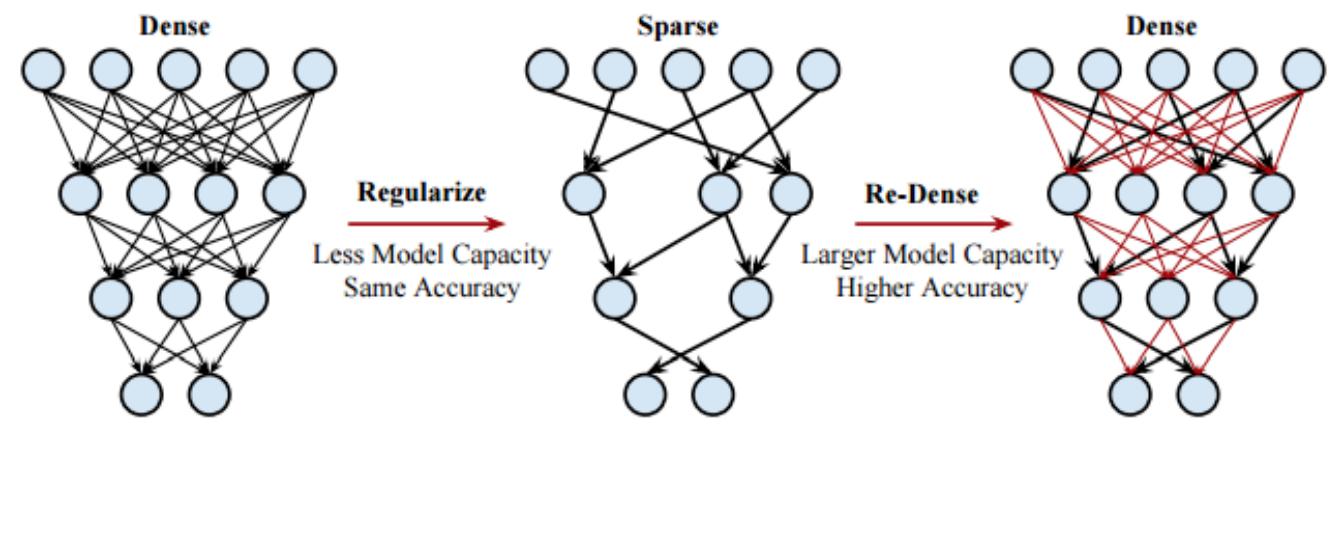


L1 regularization on least squares:

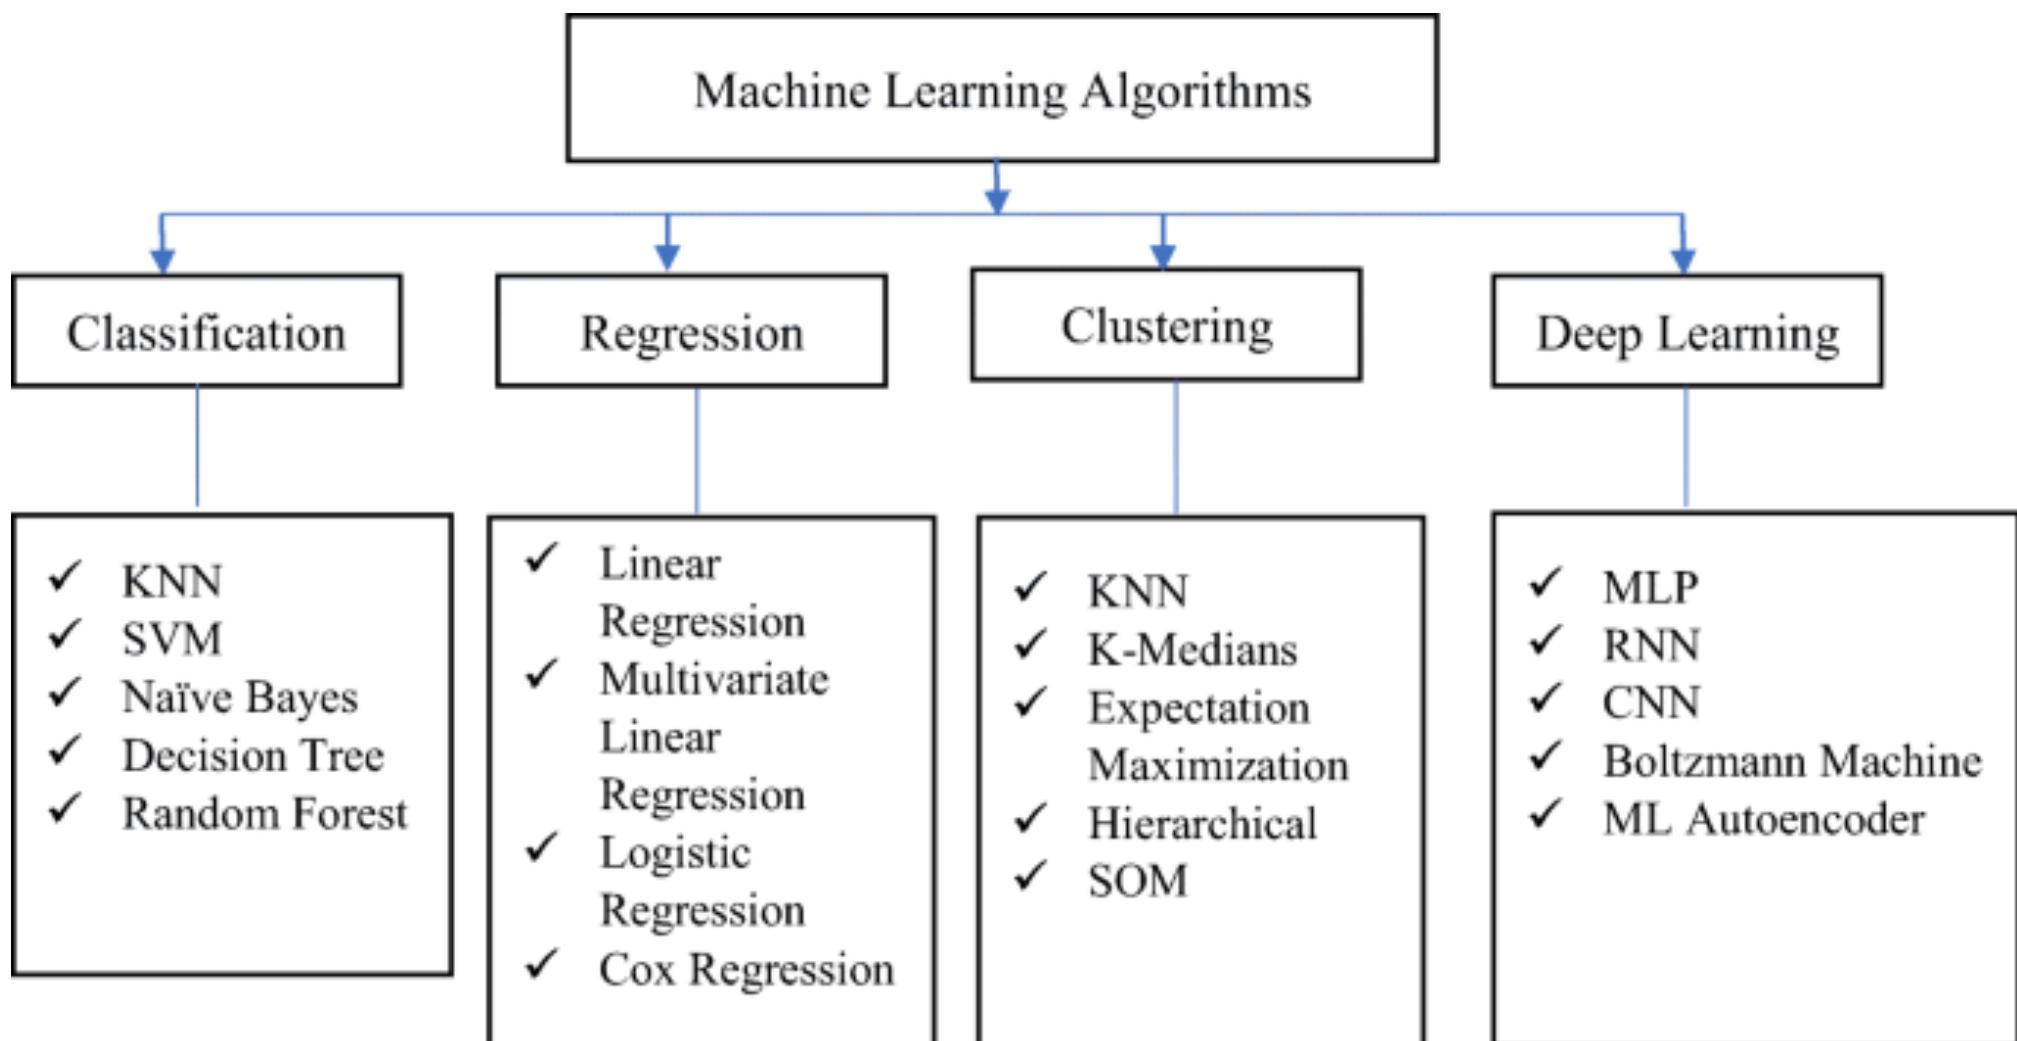
$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \sum_j \left( t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2 + \lambda \sum_{i=1}^k |w_i|$$

L2 regularization on least squares:

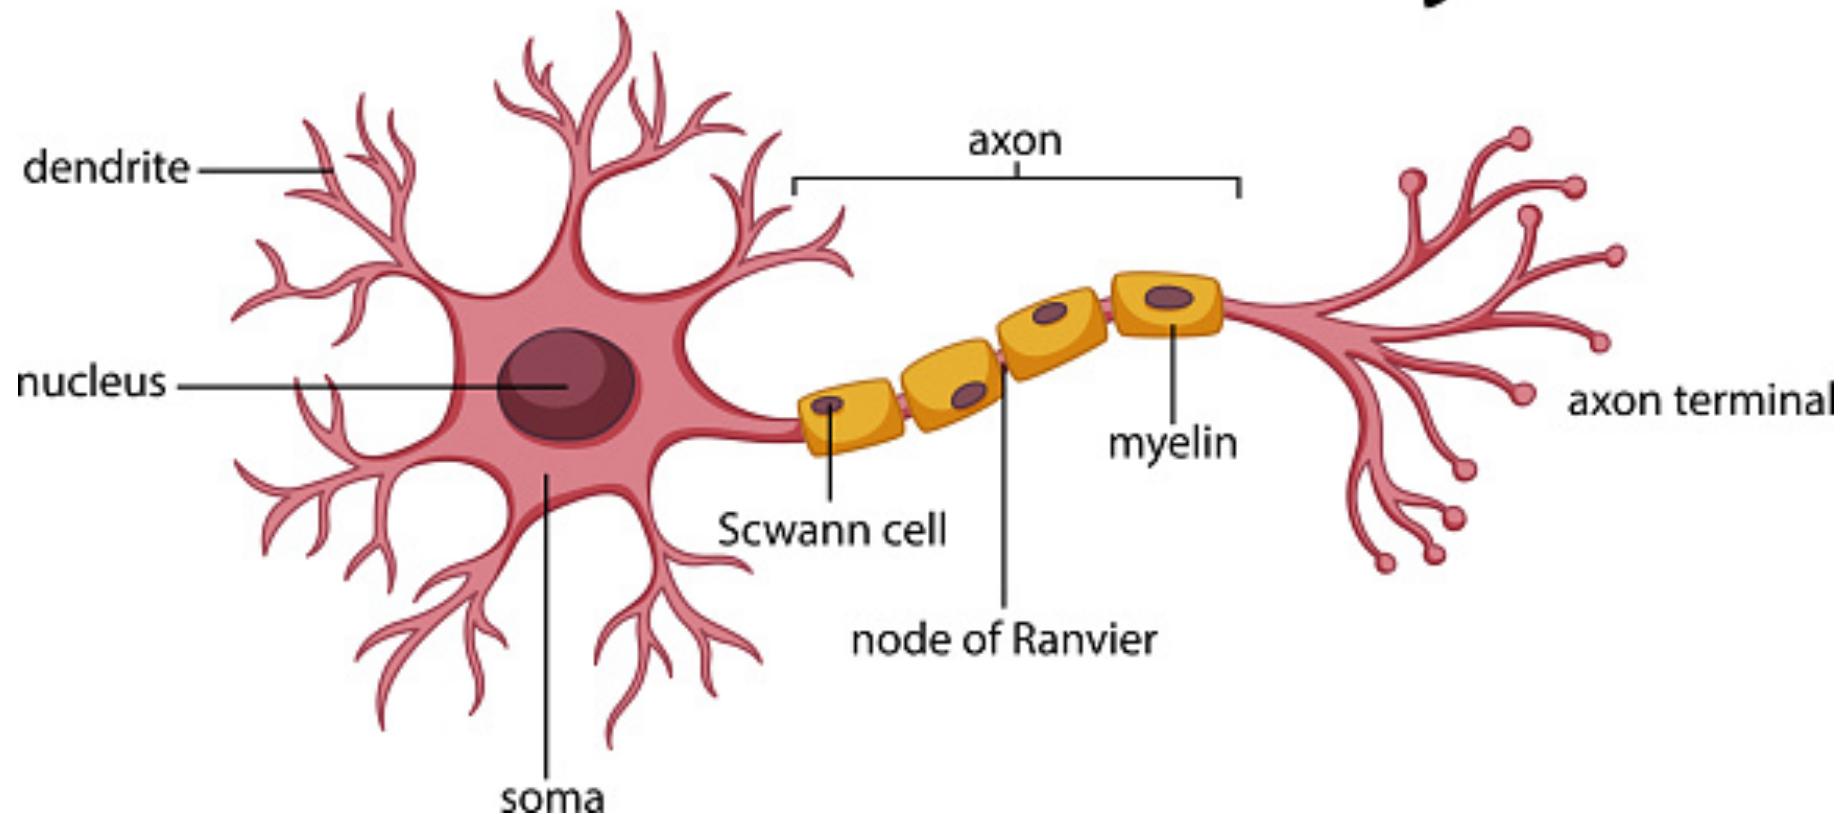
$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \sum_j \left( t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2 + \lambda \sum_{i=1}^k w_i^2$$



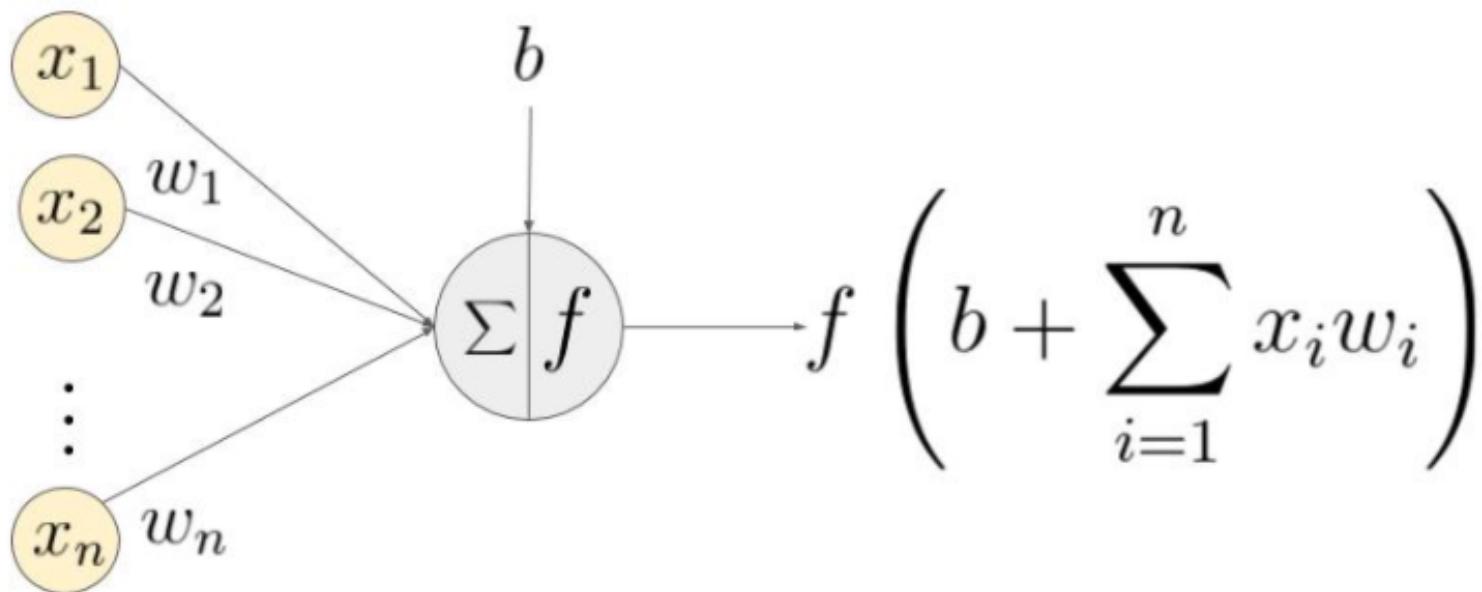
# Neural Networks



# Neuron Anatomy

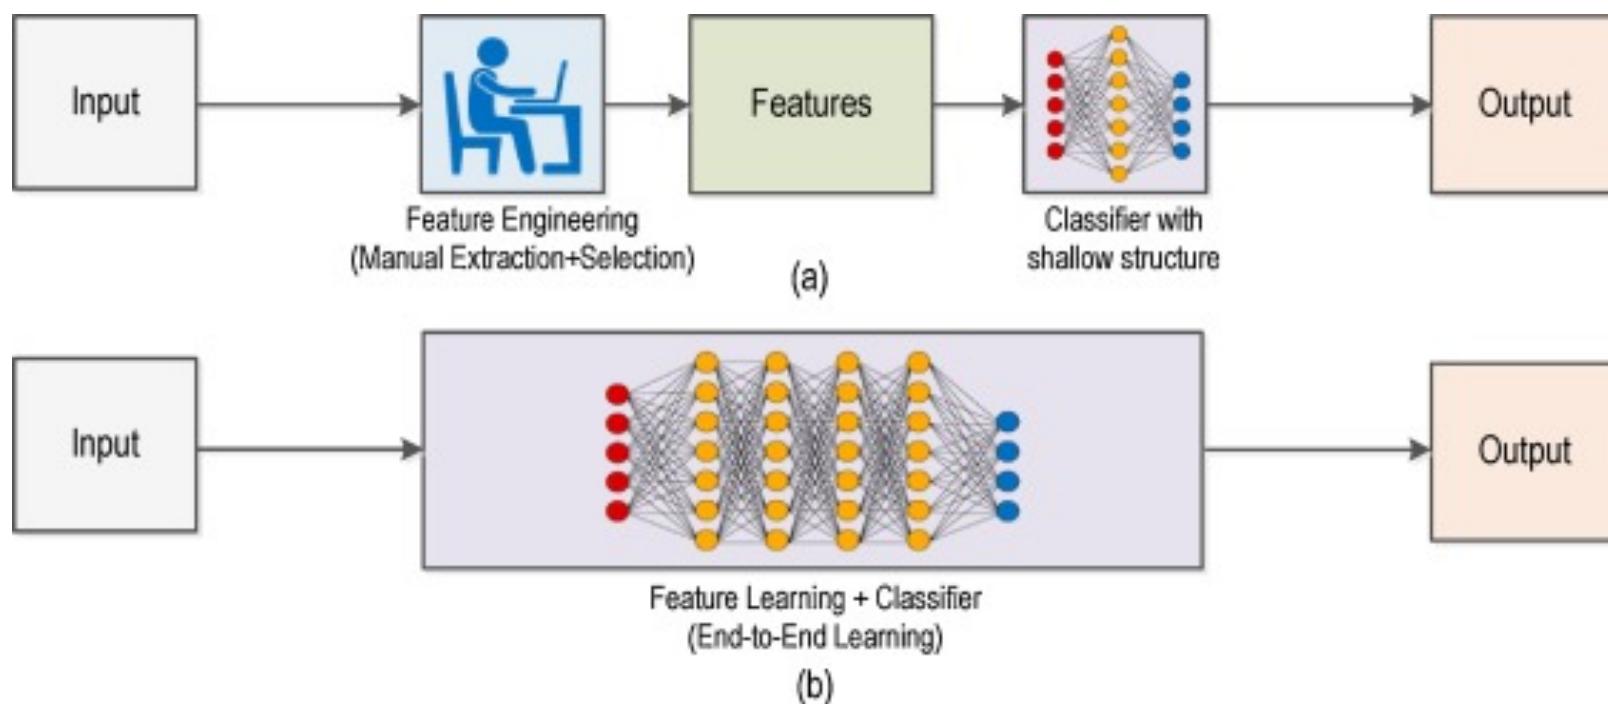


# Artificial Neuron



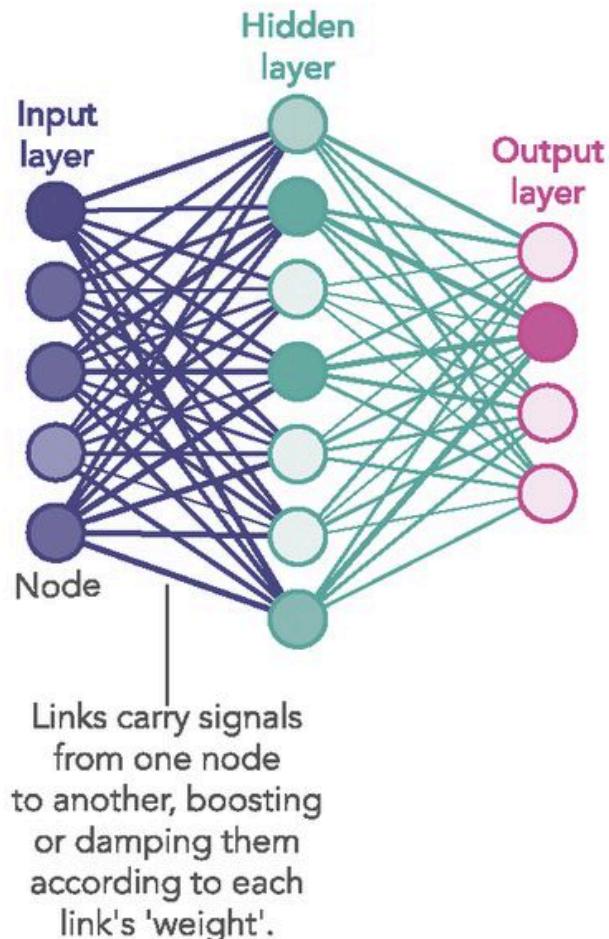
An example of a neuron showing the input ( $x_1 - x_n$ ), their corresponding weights ( $w_1 - w_n$ ), a bias ( $b$ ) and the activation function  $f$  applied to the weighted sum of the inputs.

# End-to-end Learning

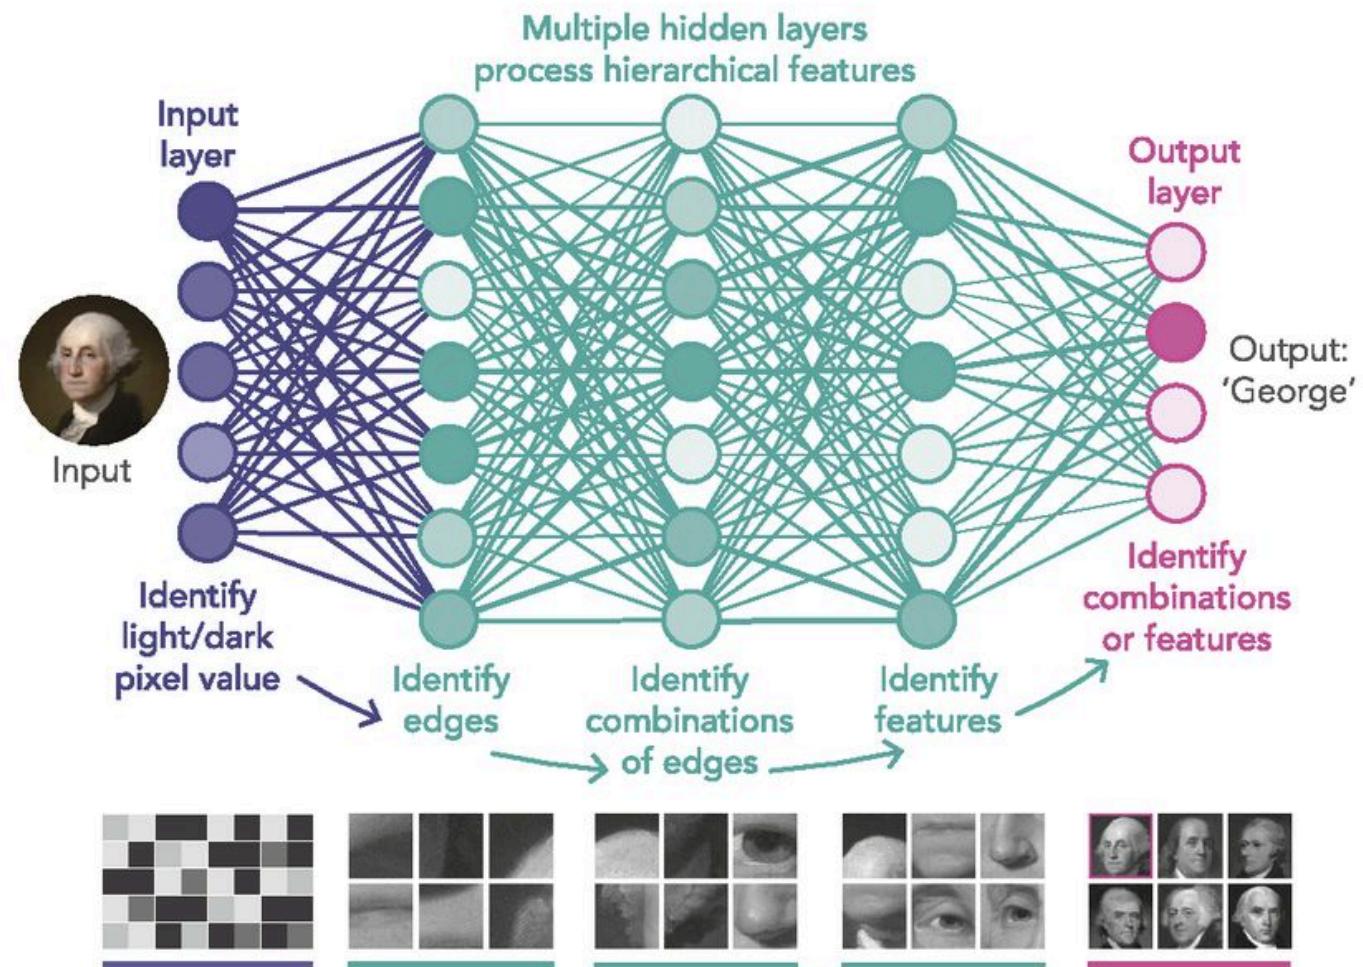


# Deep Learning

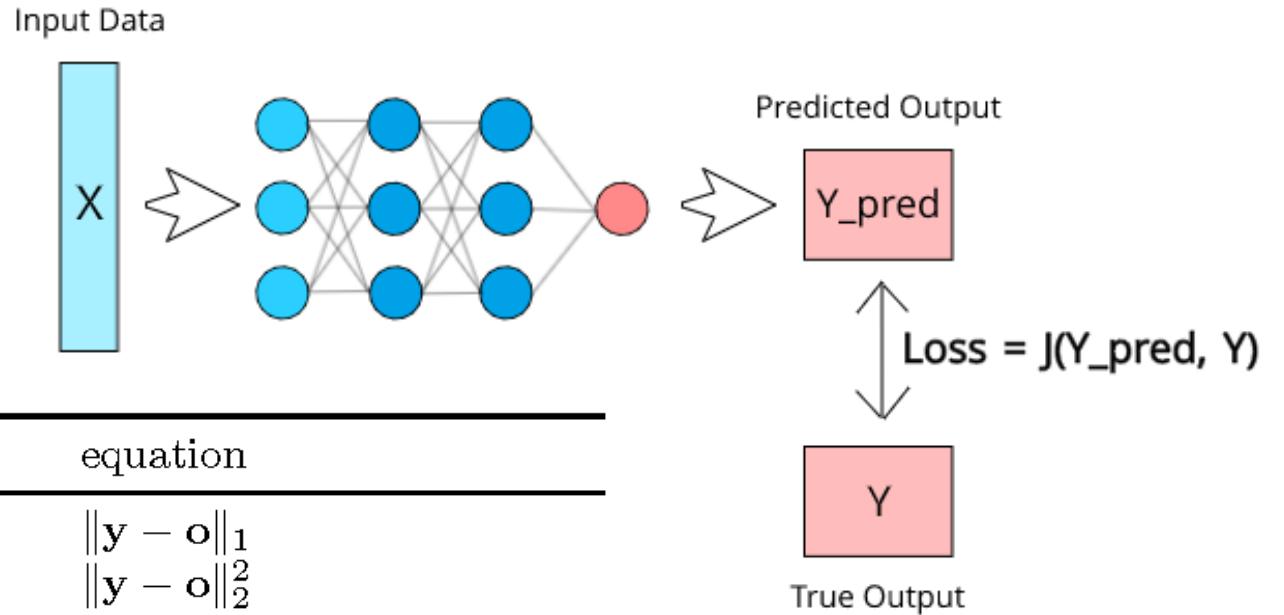
1980S-ERA NEURAL NETWORK



DEEP LEARNING NEURAL NETWORK



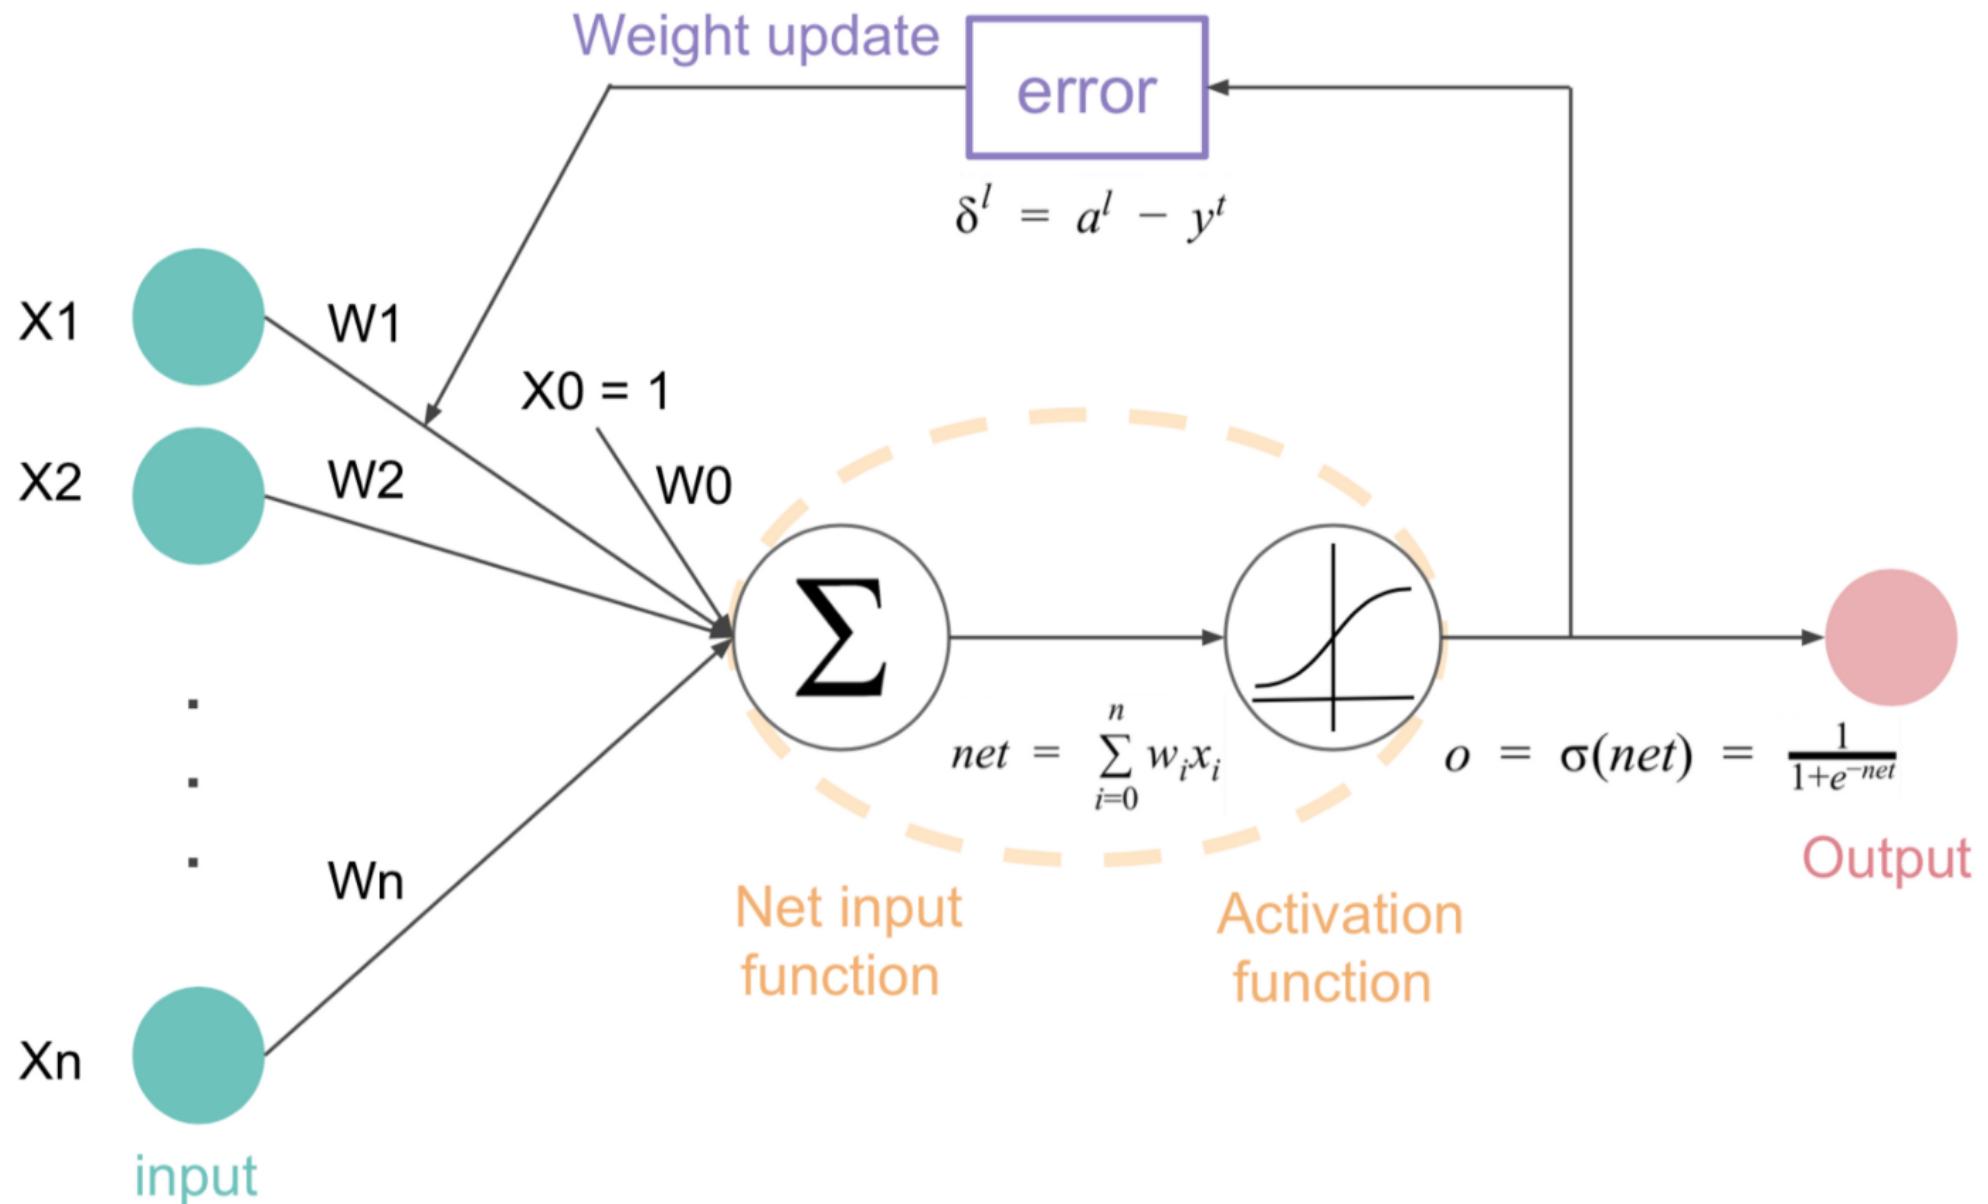
# LOSS



imate.

symbol	name	equation
$\mathcal{L}_1$	$L_1$ loss	$\ \mathbf{y} - \mathbf{o}\ _1$
$\mathcal{L}_2$	$L_2$ loss	$\ \mathbf{y} - \mathbf{o}\ _2^2$
$\mathcal{L}_1 \circ \sigma$	expectation loss	$\ \mathbf{y} - \sigma(\mathbf{o})\ _1$
$\mathcal{L}_2 \circ \sigma$	regularised expectation loss <sup>1</sup>	$\ \mathbf{y} - \sigma(\mathbf{o})\ _2^2$
$\mathcal{L}_\infty \circ \sigma$	Chebyshev loss	$\max_j  \sigma(\mathbf{o})^{(j)} - \mathbf{y}^{(j)} $
hinge	hinge [13] (margin) loss	$\sum_j \max(0, \frac{1}{2} - \hat{\mathbf{y}}^{(j)} \mathbf{o}^{(j)})$
hinge <sup>2</sup>	squared hinge (margin) loss	$\sum_j \max(0, \frac{1}{2} - \hat{\mathbf{y}}^{(j)} \mathbf{o}^{(j)})^2$
hinge <sup>3</sup>	cubed hinge (margin) loss	$\sum_j \max(0, \frac{1}{2} - \hat{\mathbf{y}}^{(j)} \mathbf{o}^{(j)})^3$
log	log (cross entropy) loss	$-\sum_j \mathbf{y}^{(j)} \log \sigma(\mathbf{o})^{(j)}$
log <sup>2</sup>	squared log loss	$-\sum_j [\mathbf{y}^{(j)} \log \sigma(\mathbf{o})^{(j)}]^2$
tan	Tanimoto loss	$-\sum_j \sigma(\mathbf{o})^{(j)} \mathbf{y}^{(j)} / \ \sigma(\mathbf{o})\ _2^2 + \ \mathbf{y}\ _2^2 - \sum_j \sigma(\mathbf{o})^{(j)} \mathbf{y}^{(j)}$
D <sub>CS</sub>	Cauchy-Schwarz Divergence [3]	$-\log \frac{\sum_j \sigma(\mathbf{o})^{(j)} \mathbf{y}^{(j)}}{\ \sigma(\mathbf{o})\ _2 \ \mathbf{y}\ _2}$

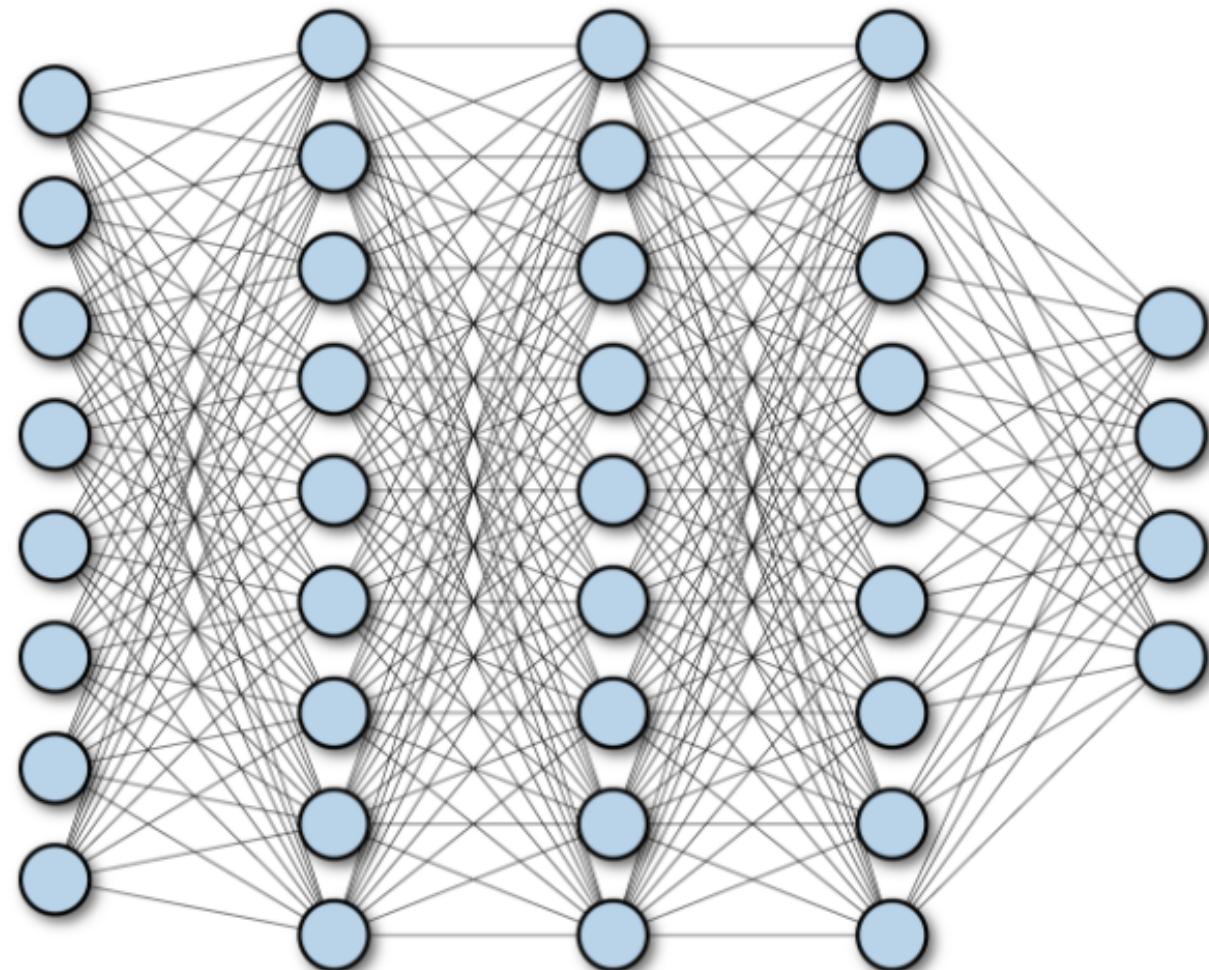
# Backpropagation



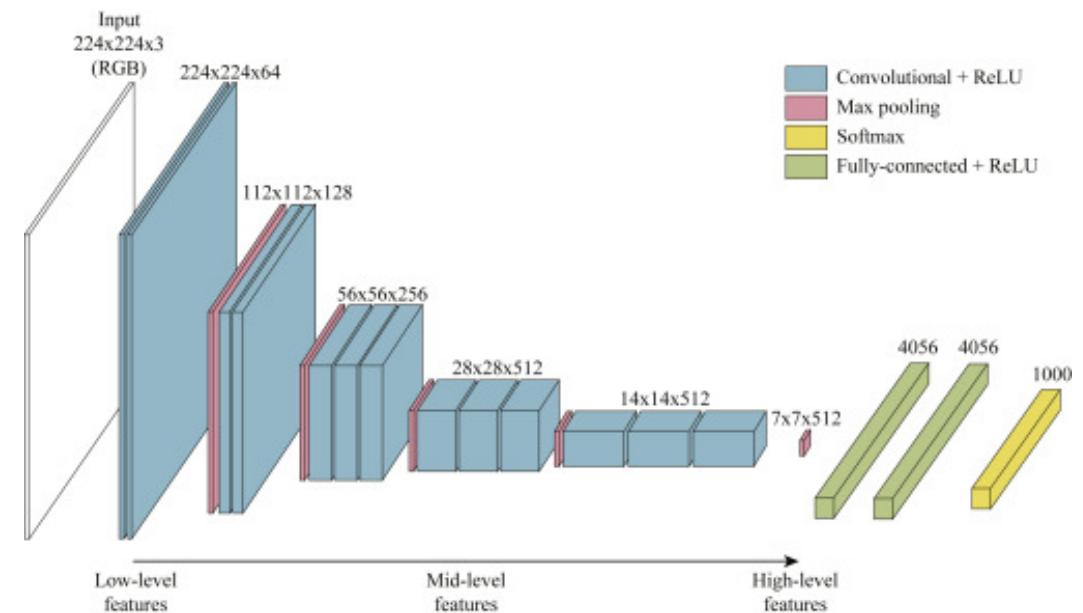
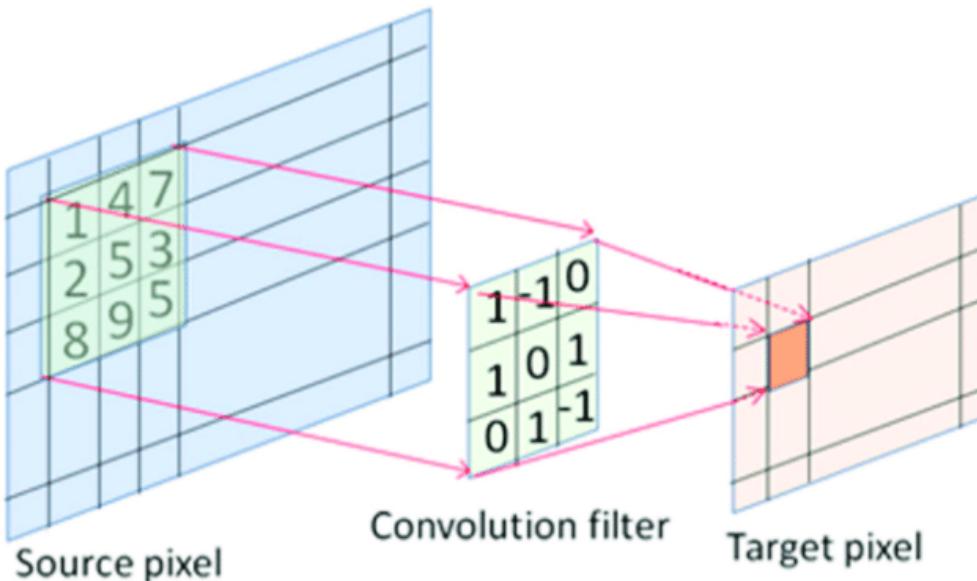
# Network Architectures

# FCN

- Computational Complexity
- Overfitting

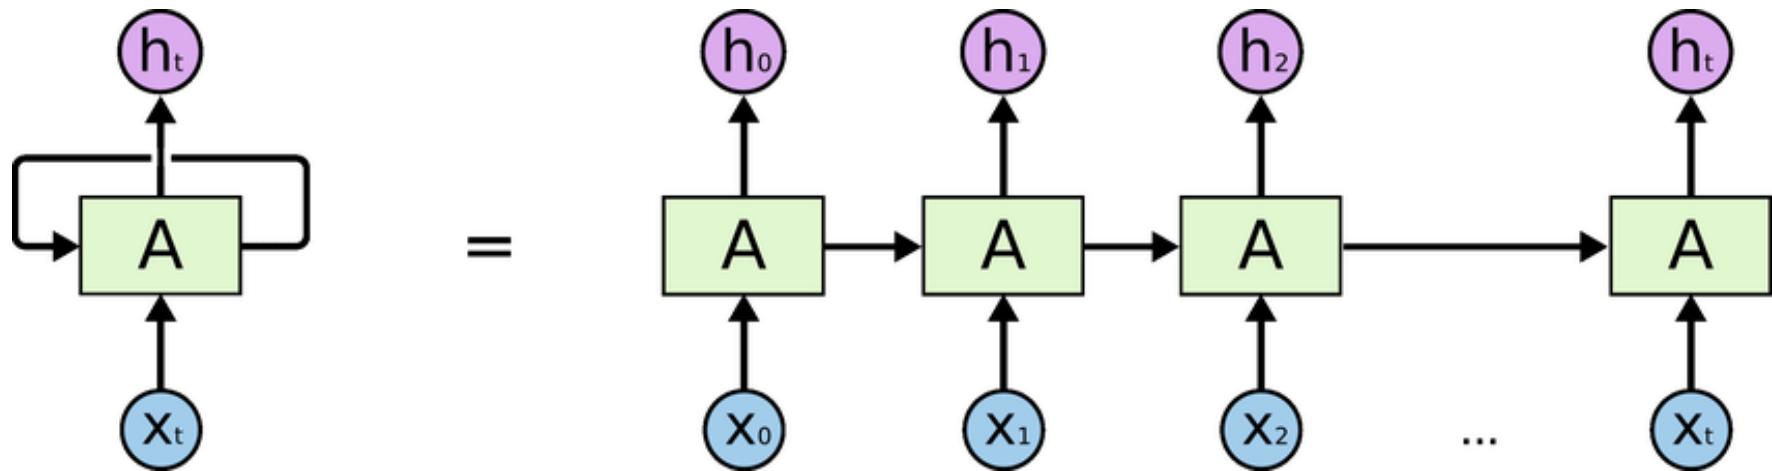


# CNN - spatial



- spatial hierarchy
- sparsely connected layers (filters share weights)
- pooling layers reduce dimensionality even more
- reduce overfitting

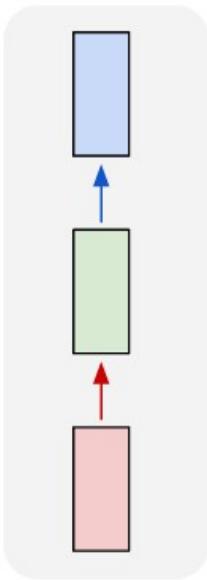
# RNN - sequential



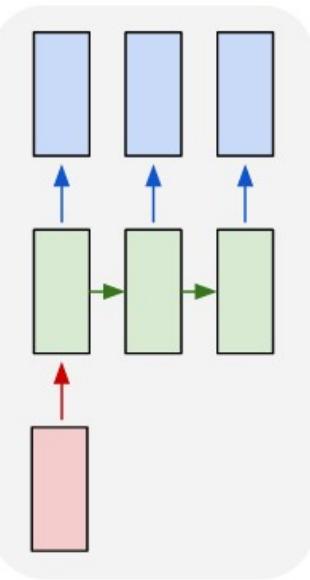
- state
- unfolding in deep layered network
- vanishing gradients - ReLU

# RNN - sequential

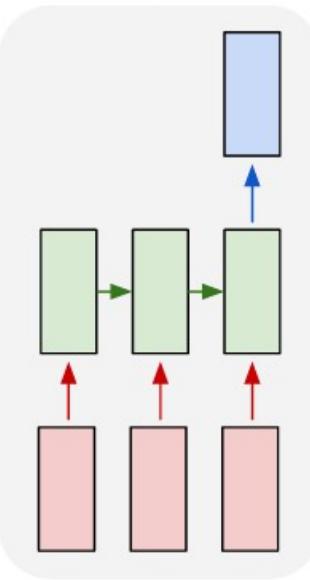
one to one



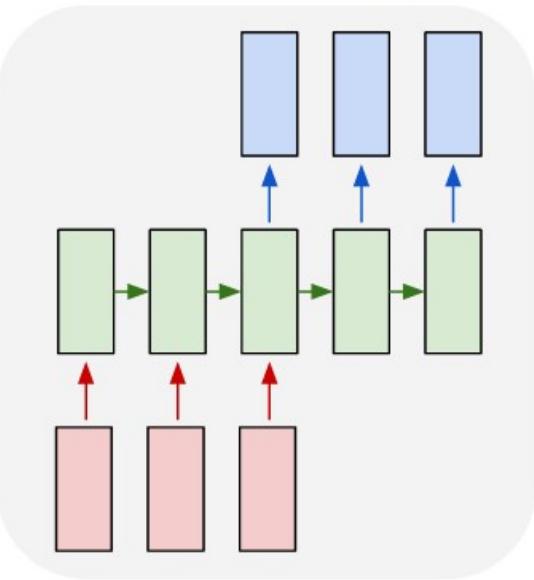
one to many



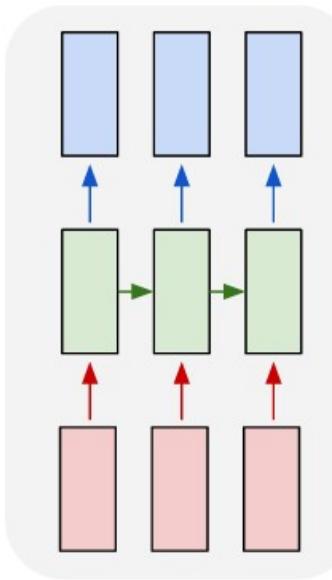
many to one



many to many

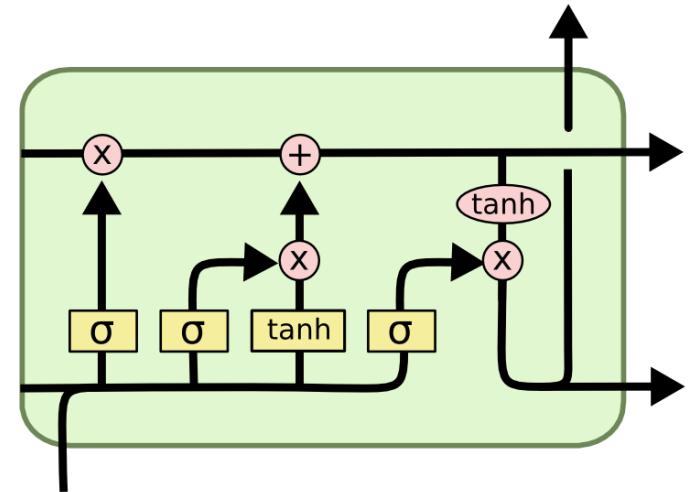


many to many



- 1-1: fixed to fixed: image classification
- 1-n. fixed to sequence. image to sentence of words
- n-1. sequence to fixed. sentiment analysis (sentence to class)
- n-n. sequence to sequence. machine translation (sentence to sentence)
- n-n synchronized. video classification

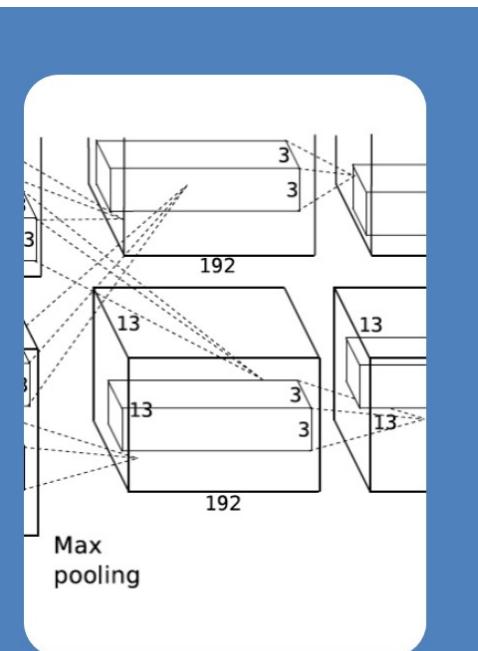
# LSTM - state



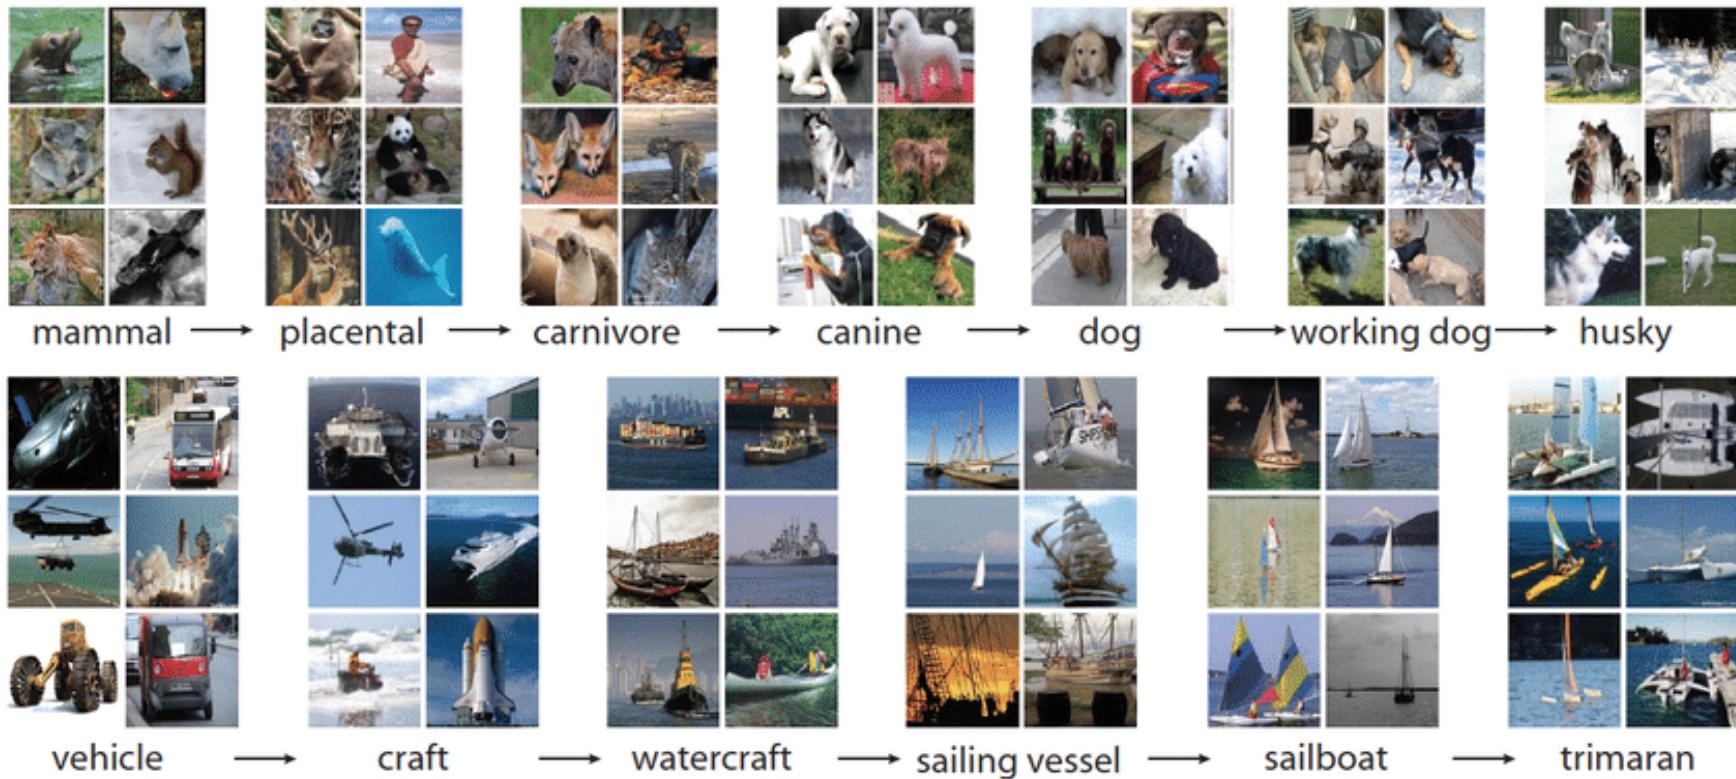
- Long Short Term Memory
- Explicit state, suffers less from vanishing gradient
- Good building block for large recurrent networks

# ML Revolution

- Why did ML take off in 2012?
- **Algorithms** – deep learning, CNN
- **Labeled Data sets** (ImageNet)
- Compute Power (**GPU**)
- Alex Krizhevsky [2012]



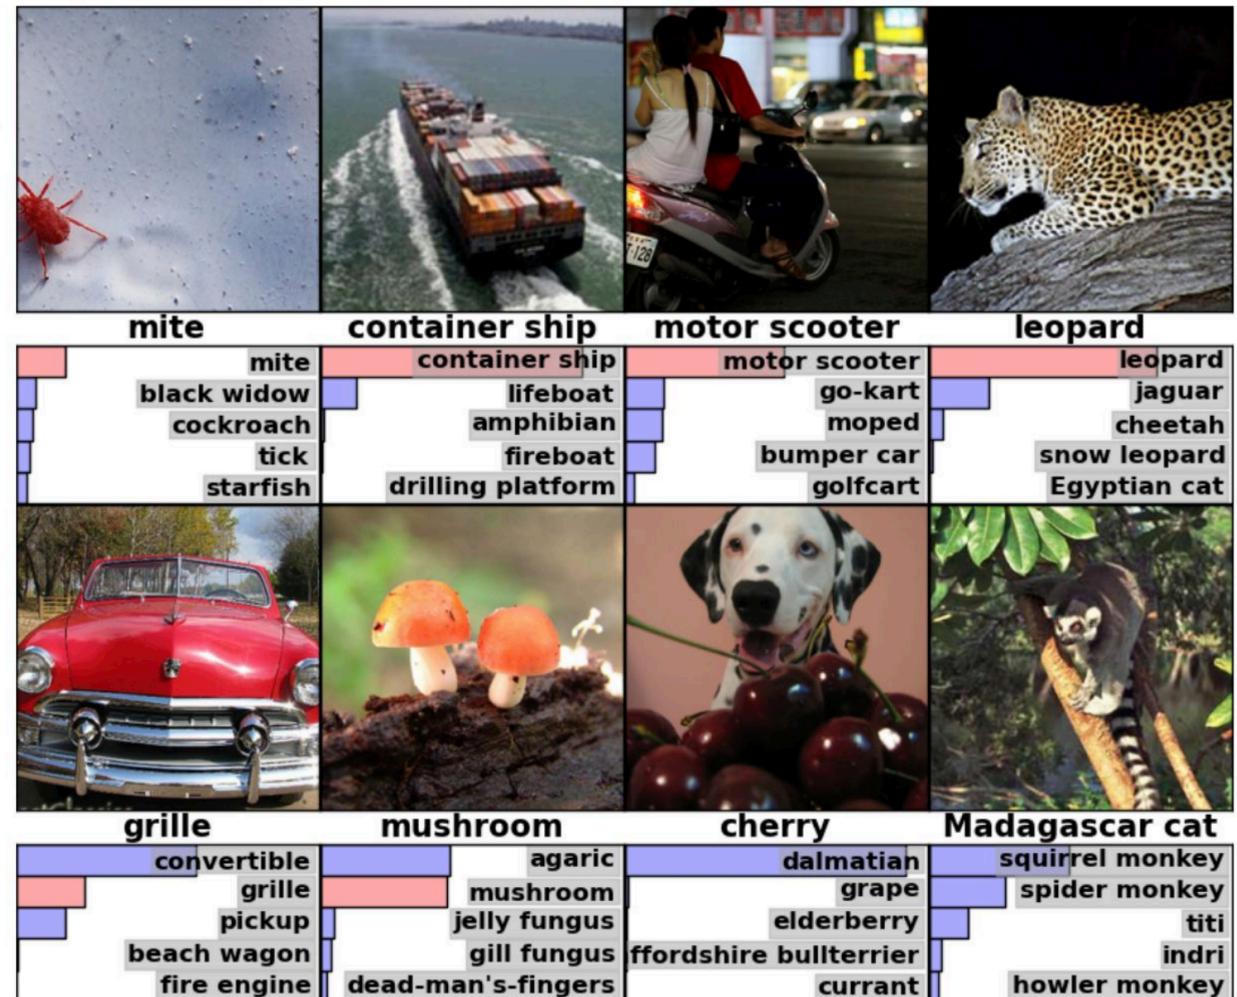
# ImageNet



# ImageNet Challenge



- 1,000 object classes (categories).
- Images:
  - 1.2 M train
  - 100k test.



# PyTorch & TensorFlow

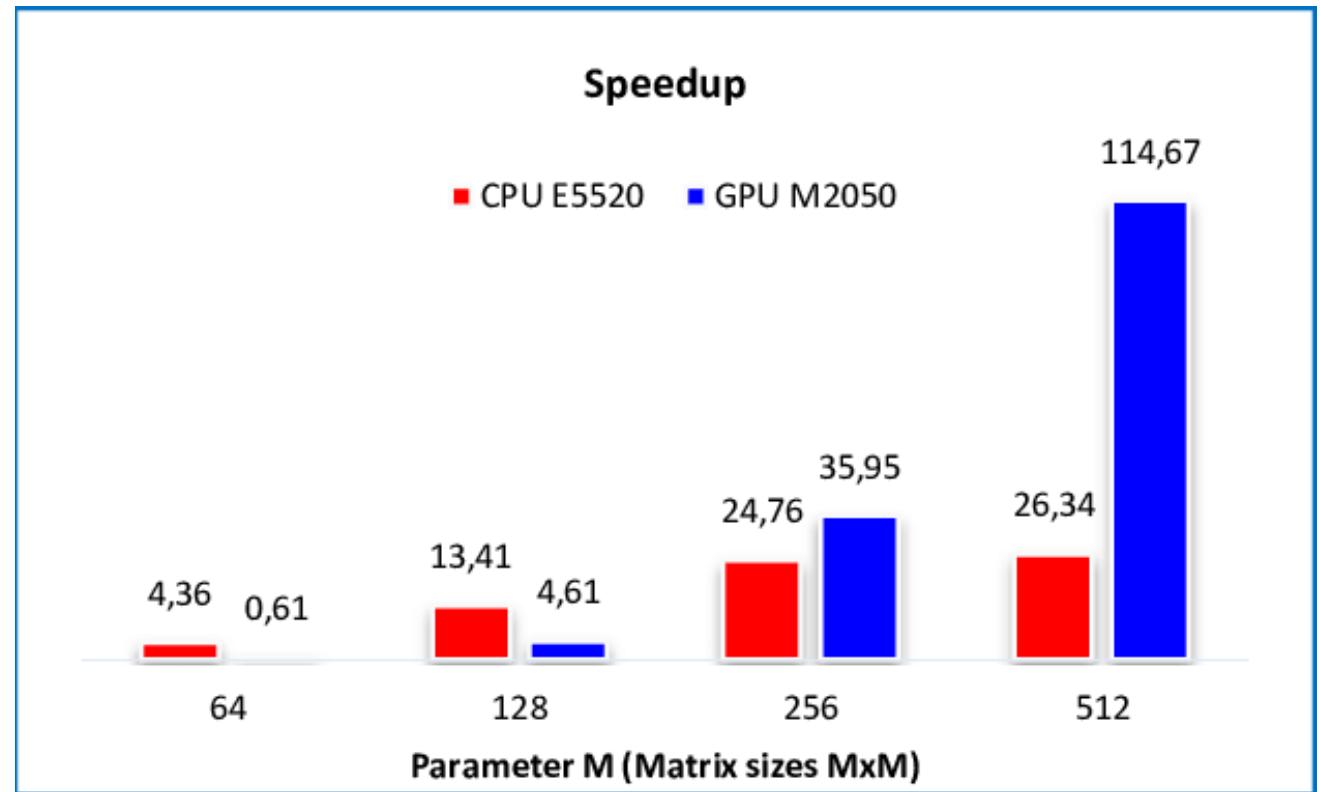
```
import torch.nn as nn
import torch.nn.functional as F

class NeuralNet(nn.Module):
    def __init__(self):
        super(NeuralNet, self).__init__()
        self.conv1 = nn.Conv2d(3, 6, 5)
        self.pool = nn.MaxPool2d(2, 2)
        self.conv2 = nn.Conv2d(6, 16, 5)
        self.fc1 = nn.Linear(16 * 5 * 5, 120)
        self.fc2 = nn.Linear(120, 84)
        self.fc3 = nn.Linear(84, 10)

    def forward(self, x):
        x = self.pool(F.relu(self.conv1(x)))
        x = self.pool(F.relu(self.conv2(x)))
        x = x.view(-1, 16 * 5 * 5)
        x = F.relu(self.fc1(x))
        x = F.relu(self.fc2(x))
        x = self.fc3(x)
        return x
```

# GPU

- GPU parallel Matrix Multiply large speedup
- Use GPU versions of TensorFlow and PyTorch



# DataScience Lab

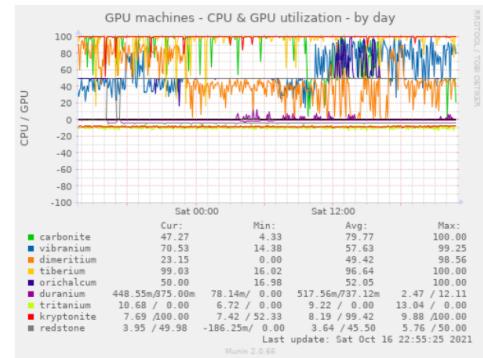
- DataScience Lab
- ALICE
- Google Colab



## GPU

The following machines contain high-end GPUs and can be used for CUDA-supported computations, such as training neural networks using e.g. TensorFlow and PyTorch.

Hostname	RAM	CPU/GPU	Local Storage	Available Status	to?
<a href="#">duranium.liacs.nl</a>	256GB	20 Intel Xeon E5-2650v3 cores @ 2.30GHz (40 threads), 6 NVIDIA GTX 980 Ti (6GB memory each), 2 NVIDIA Titan X (12GB memory each)	3TB under <a href="#">/local</a>	OK	Students
<a href="#">tritium.liacs.nl</a>	1TB	20 Intel Xeon E5-2650v3 cores @ 2.30GHz (40 threads), 16 NVIDIA Tesla K80 (11.5GB memory each) These are 8 dual-GPU boards	3TB under <a href="#">/local</a>	OK	Staff
<a href="#">kryptonite.liacs.nl</a>	64GB	16 Intel Xeon E5-2650 cores @ 2.00GHz (32 threads), 2 NVIDIA GeForce RTX 2080 Ti (11GB memory each)	7.3TB under <a href="#">/local</a>	OK	Staff
<a href="#">phlogiston.liacs.nl</a>	64GB	16 Intel Xeon E5-2650 cores @ 2.00GHz (32 threads), 2 NVIDIA GeForce RTX 2080 Ti (11GB memory each)	8.6TB under <a href="#">/local</a>	OK	Students
<a href="#">redstone.liacs.nl</a>	64GB	16 Intel Xeon E5-2650 cores @ 2.00GHz (32 threads), 2 NVIDIA GeForce RTX 2080 Ti (11GB memory each)	8.6TB under <a href="#">/local</a>	OK	Staff
<a href="#">runite.liacs.nl</a>	64GB	16 Intel Xeon E5-2650 cores @ 2.00GHz (32 threads), 2 NVIDIA GeForce RTX 2080 Ti (11GB memory each)	8.6TB under <a href="#">/local</a>	OK	Staff
<a href="#">carbonite.liacs.nl</a>	256GB	24 Intel Xeon Silver 4214 cores @ 2.20GHz (48 threads), 2 NVIDIA GeForce RTX 3090 Ti (24GB memory each)	3.4TB under <a href="#">/local</a>	OK	Staff
<a href="#">dimeritium.liacs.nl</a>	256GB	24 Intel Xeon Silver 4214 cores @ 2.20GHz (48 threads), 2 NVIDIA GeForce RTX 3090 (24GB memory each)	3.4TB under <a href="#">/local</a>	OK	Staff
<a href="#">orichalcum.liacs.nl</a>	256GB	24 Intel Xeon Silver 4214 cores @ 2.20GHz (48 threads),	3.4TB under <a href="#">/local</a>	OK	Staff



# About ALICE

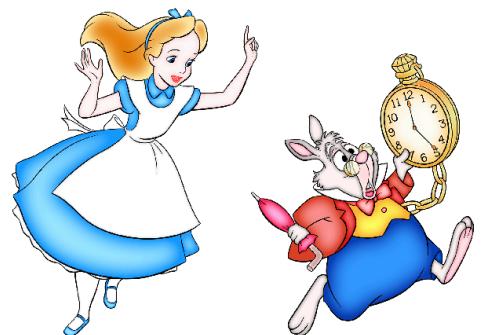
 Actions

*From ALICE Documentation*

[Alice Documentation](#) > [About ALICE](#)

ALICE (Academic Leiden Interdisciplinary Cluster Environment) is the high-performance computing (HPC) facility of the partnership between Leiden University and Leiden University Medical Center (LUMC). It is available to any researcher from both partners. Leiden University and LUMC aim to help deliver cutting edge research using innovative technology within the broad area of data-centric HPC. Both partners are responsible for the hosting, system support, scientific support and service delivery of several large super-computing and research data storage resources for the Leiden research community.

This wiki is the main source of documentation about the ALICE cluster.



Off to research computing Wonderland

## Contents

[\[hide\]](#)

- [1 Research Acknowledgement](#)
- [2 Why ALICE](#)
- [3 Overview of the cluster](#)
- [4 Future plans](#)
- [5 How to get involved with ALICE](#)
- [6 Costs overview](#)
- [7 Current Status Overview](#)

# Questions?