Module 1: Crash course in AI

INFO901

Marija Slavkovik 2022

What is machine learning

- How we humans understand learning: an agent is learning if it improves its performance on future tasks after making observations about the world.
- What is learning in AI:
    - supervised learning
    - unsupervised learning
    - (combinations of the above two)
    - reinforcement learning

Step back to knowledge representation (KR)

A computer

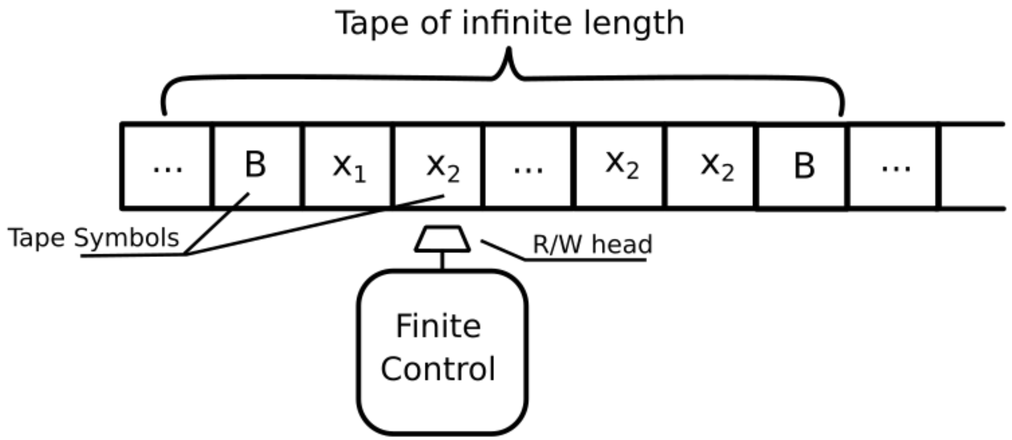

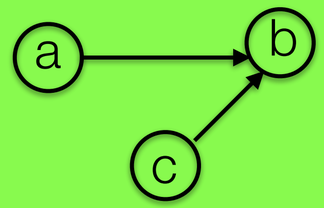

	a	b	c
a	0	1	0
b	0	0	0
c	0	1	0

Example	Input Attributes										Goal
	Alt	*Bar*	*Fri*	*Hun*	*Pat*	*Price*	*Rain*	*Res*	*Type*	*Est*	*WillWait*
$\mathbf{x}_1$	*Yes*	*No*	*No*	*Yes*	*Some*	*$$$*	*No*	*Yes*	*French*	*0–10*	$y_1 = Yes$
$\mathbf{x}_2$	*Yes*	*No*	*No*	*Yes*	*Full*	*$*	*No*	*No*	*Thai*	*30–60*	$y_2 = No$
$\mathbf{x}_3$	*No*	*Yes*	*No*	*No*	*Some*	*$*	*No*	*No*	*Burger*	*0–10*	$y_3 = Yes$
$\mathbf{x}_4$	*Yes*	*No*	*Yes*	*Yes*	*Full*	*$*	*Yes*	*No*	*Thai*	*10–30*	$y_4 = Yes$
$\mathbf{x}_5$	*Yes*	*No*	*Yes*	*No*	*Full*	*$$$*	*No*	*Yes*	*French*	*>60*	$y_5 = No$
$\mathbf{x}_6$	*No*	*Yes*	*No*	*Yes*	*Some*	*$$*	*Yes*	*Yes*	*Italian*	*0–10*	$y_6 = Yes$
$\mathbf{x}_7$	*No*	*Yes*	*No*	*No*	*None*	*$*	*Yes*	*No*	*Burger*	*0–10*	$y_7 = No$
$\mathbf{x}_8$	*No*	*No*	*No*	*Yes*	*Some*	*$$*	*Yes*	*Yes*	*Thai*	*0–10*	$y_8 = Yes$
$\mathbf{x}_9$	*No*	*Yes*	*Yes*	*No*	*Full*	*$*	*Yes*	*No*	*Burger*	*>60*	$y_9 = No$
$\mathbf{x}_{10}$	*Yes*	*Yes*	*Yes*	*Yes*	*Full*	*$$$*	*No*	*Yes*	*Italian*	*10–30*	$y_{10} = No$
$\mathbf{x}_{11}$	*No*	*No*	*No*	*No*	*None*	*$*	*No*	*No*	*Thai*	*0–10*	$y_{11} = No$
$\mathbf{x}_{12}$	*Yes*	*Yes*	*Yes*	*Yes*	*Full*	*$*	*No*	*No*	*Burger*	*30–60*	$y_{12} = Yes$

Figure 18.3 Examples for the restaurant domain.

$(p \vee q) \leftrightarrow r$

p	q	
0	1	1
1	0	1
1	1	1
0	0	0

What is data?

Data

data is an individual fact, statistics, or item of information

digital data is data that can be created by a computer

structured data is data that is organised in a way that makes it easy for processing

Factored representation

real world to data

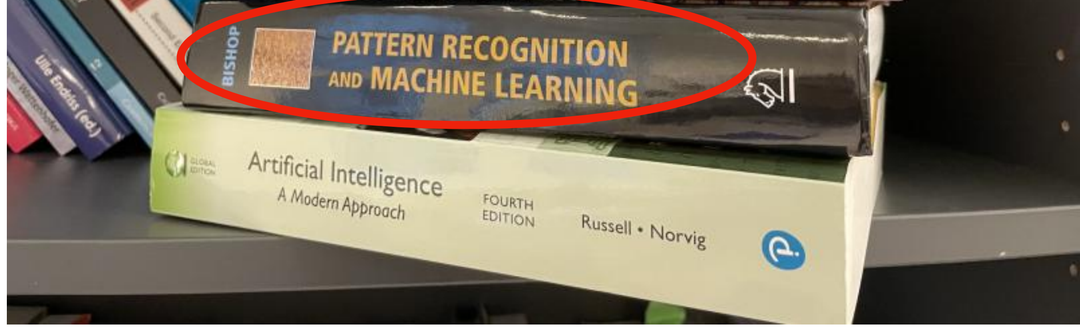

https://en.wikipedia.org/wiki/Iris_flower_data_set

(Colour
Width
Length
Animate
Scent
(2,1,1,0,5)

“In machine learning, a feature is an individual measurable property or characteristic of a phenomenon being observed” Bishop

Factored representation

real world to data

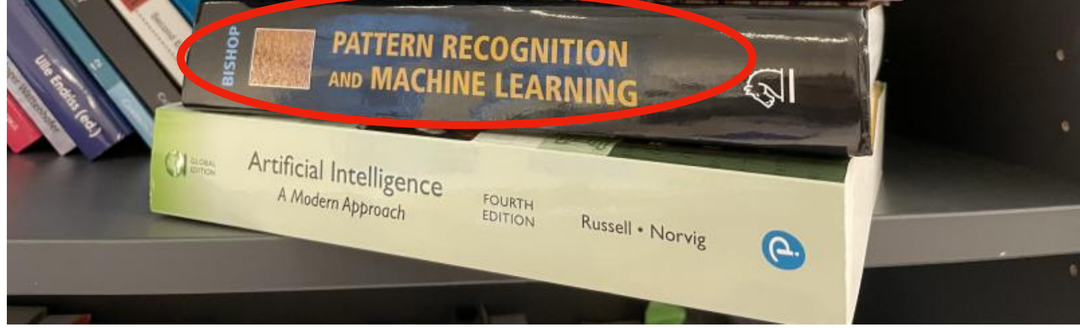

https://en.wikipedia.org/wiki/Iris_flower_data_set

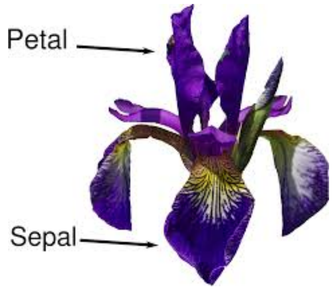


```
First five columns of data:
[[ 5.1  3.5  1.4  0.2]
 [ 4.9  3.   1.4  0.2]
 [ 4.7  3.2  1.3  0.2]
 [ 4.6  3.1  1.5  0.2]
 [ 5.   3.6  1.4  0.2]]
```

Colour
Width
Length
Animate
Scent

Features

(2,1,1,0,5)

“In machine learning, a feature is an individual measurable property or characteristic of a phenomenon being observed” Bishop

Factored representation

real world to data

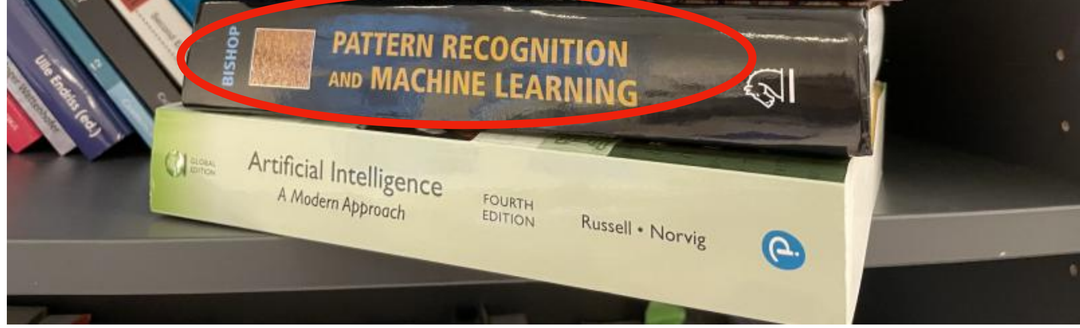

https://en.wikipedia.org/wiki/Iris_flower_data_set

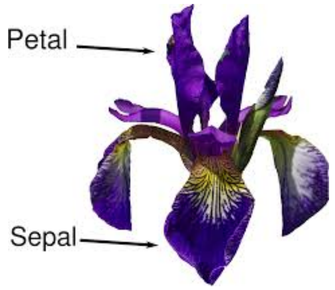


```
First five columns of data:
[[ 5.1  3.5  1.4  0.2]
 [ 4.9  3.   1.4  0.2]
 [ 4.7  3.2  1.3  0.2]
 [ 4.6  3.1  1.5  0.2]
 [ 5.   3.6  1.4  0.2]]
```

Colour
Width
Length
Animate
Scent

Features

(2,1,1,0,5)

“In machine learning, a feature is an individual measurable property or characteristic of a phenomenon being observed” Bishop

Factored representation

real world to data

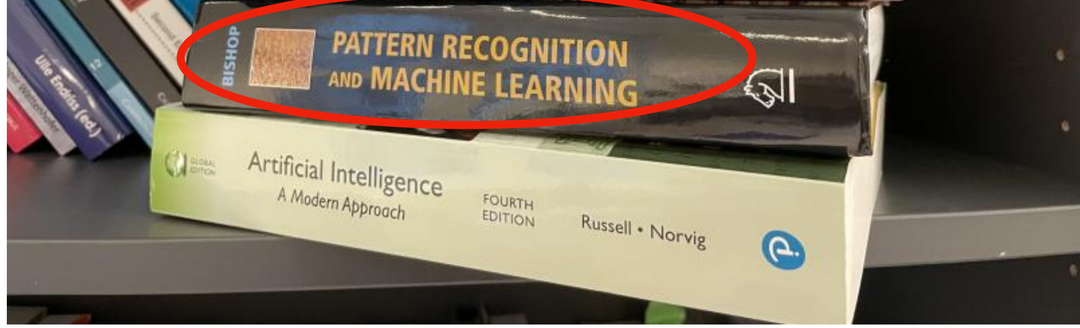

https://en.wikipedia.org/wiki/Iris_flower_data_set

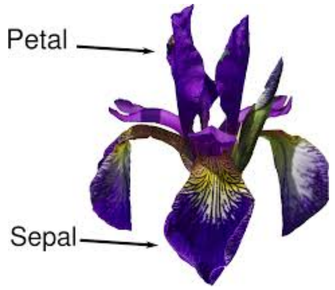

```
First five columns of data:
[[ 5.1  3.5  1.4  0.2]
 [ 4.9  3.   1.4  0.2]
 [ 4.7  3.2  1.3  0.2]
 [ 4.6  3.1  1.5  0.2]
 [ 5.   3.6  1.4  0.2]]
```

Colour
Width
Length
Animate
Scent

Features

(2,1,1,0,5)

“In machine learning, a feature is an individual measurable property or characteristic of a phenomenon being observed” Bishop

Factored representation

real world to data

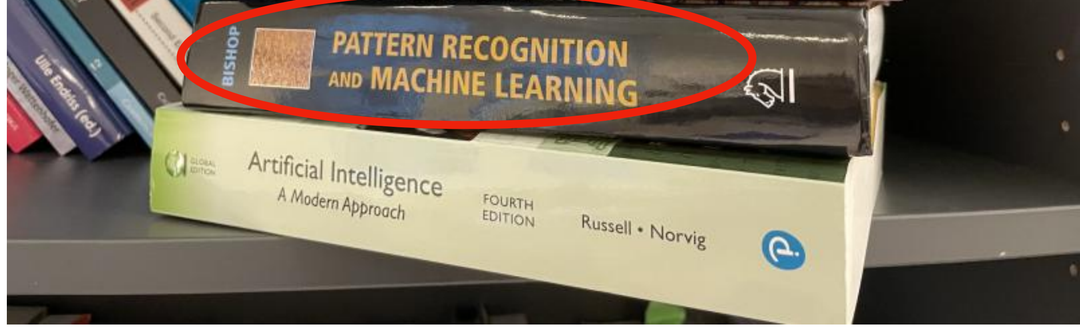

https://en.wikipedia.org/wiki/Iris_flower_data_set

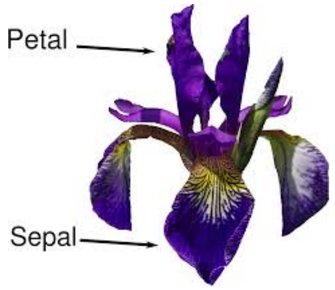

```
First five columns of data:
[[ 5.1  3.5  1.4  0.2]
 [ 4.7  3.2  1.3  0.2]
 [ 5.   3.6  1.4  0.2]]
```

Colour
Width
Length
Animate
Scent

Features

(2,1,1,0,5)

“In machine learning, a feature is an individual measurable property or characteristic of a phenomenon being observed” Bishop

Factored representation

real world to data

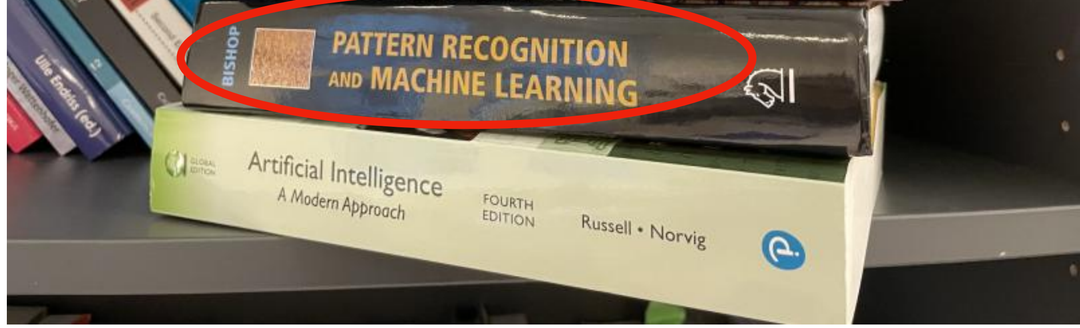

https://en.wikipedia.org/wiki/Iris_flower_data_set

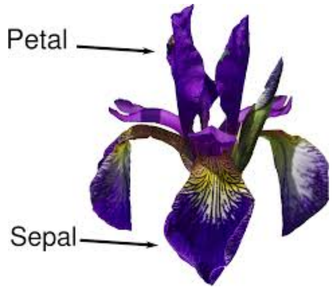


```
First five columns of data:
[[ 5.1  3.5  1.4  0.2]
 [ 4.9  3.   1.4  0.2]
 [ 4.7  3.2  1.3  0.2]
 [ 4.6  3.1  1.5  0.2]
 [ 5.   3.6  1.4  0.2]]
```

Colour
Width
Length
Animate
Scent

Features

(2,1,1,0,5)

“In machine learning, a feature is an individual measurable property or characteristic of a phenomenon being observed” Bishop

Factored representation

real world to data

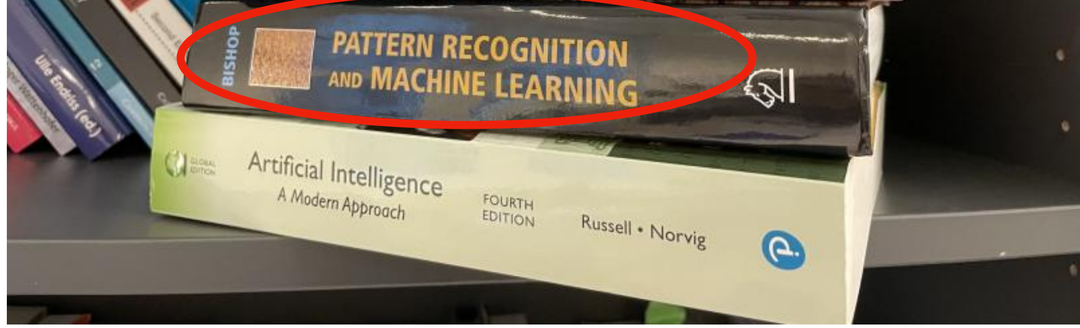

https://en.wikipedia.org/wiki/Iris_flower_data_set

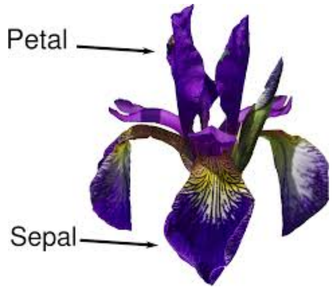


```
First five columns of data:
[[ 5.1  3.5  1.4  0.2]
 [ 4.9  3.   1.4  0.2]
 [ 4.7  3.2  1.3  0.2]
 [ 4.6  3.1  1.5  0.2]
 [ 5.   3.6  1.4  0.2]]
```

Colour
Width
Length
Animate
Scent

Features

(2,1,1,0,5)

"In machine learning, a feature is an individual measurable property or characteristic of a phenomenon being observed" Bishop

Factored representation

real world to data

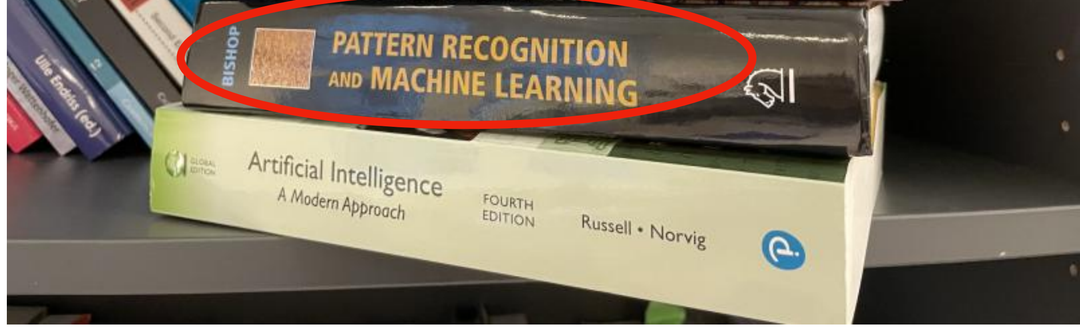

https://en.wikipedia.org/wiki/Iris_flower_data_set

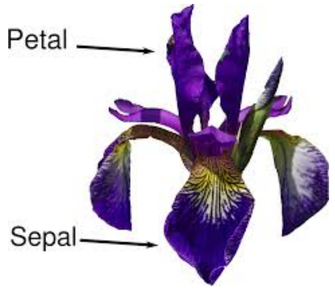

sepal
length

```
First five columns of data:
[ 5.1  3.5  1.4  0.2]
  4.9  3.   1.4  0.2]
  4.7  3.2  1.3  0.2]
  4.6  3.1  1.5  0.2]
  5.   3.6  1.4  0.2]]
```

Colour
Width
Length
Animate
Scent

Features

(2,1,1,0,5)

“In machine learning, a feature is an individual measurable property or characteristic of a phenomenon being observed” Bishop

Factored representation

real world to data

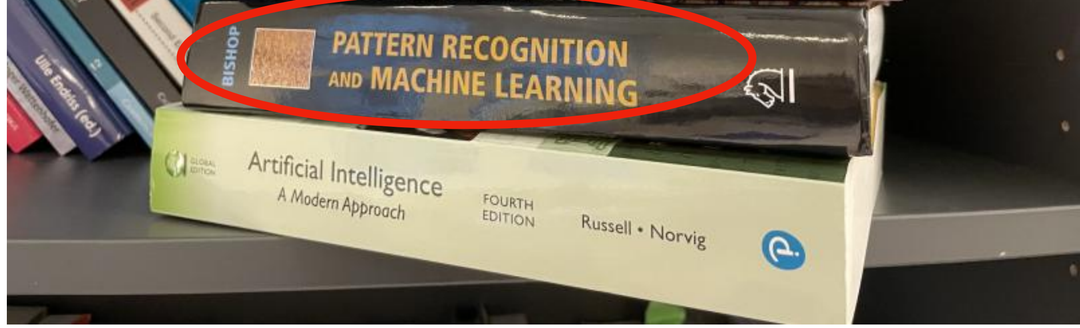

https://en.wikipedia.org/wiki/Iris_flower_data_set

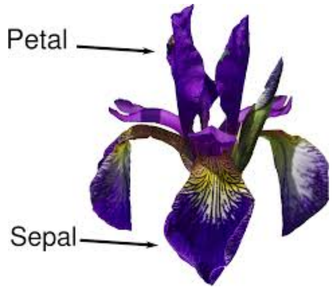

sepal
width

```
First five columns of data:
[[ 5.1  3.5  1.4  0.2]
 [ 4.9  3.   1.4  0.2]
 [ 4.7  3.2  1.3  0.2]
 [ 4.6  3.1  1.5  0.2]
 [ 5.   3.6  1.4  0.2]]
```

Colour
Width
Length
Animate
Scent

Features

(2,1,1,0,5)

“In machine learning, a feature is an individual measurable property or characteristic of a phenomenon being observed” Bishop

Factored representation

real world to data

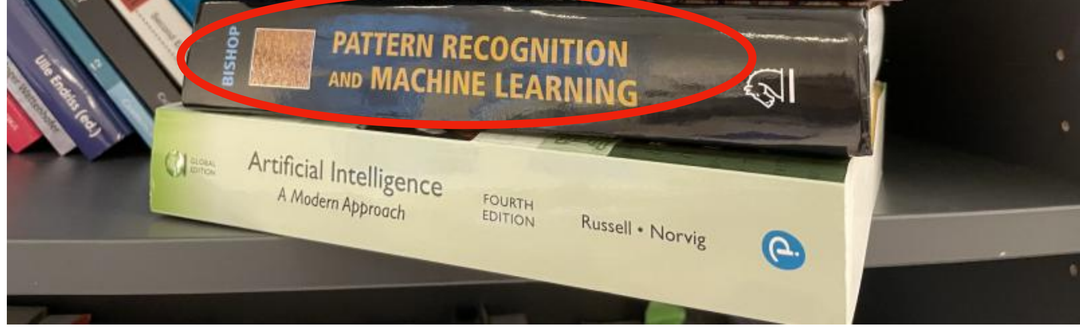

https://en.wikipedia.org/wiki/Iris_flower_data_set

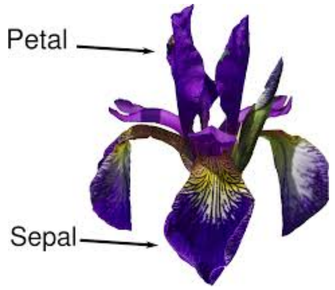

petal
length

```
First five columns of data:
[[ 5.1  3.5  1.4  0.2]
 [ 4.9  3.   1.4  0.2]
 [ 4.7  3.2  1.3  0.2]
 [ 4.6  3.1  1.5  0.2]
 [ 5.   3.6  1.4  0.2]]
```

Colour
Width
Length
Animate
Scent

Features

(2,1,1,0,5)

“In machine learning, a feature is an individual measurable property or characteristic of a phenomenon being observed” Bishop

Factored representation

real world to data

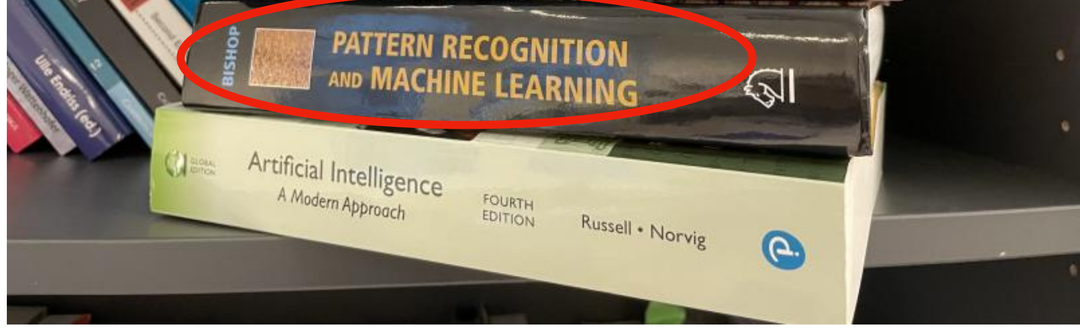

https://en.wikipedia.org/wiki/Iris_flower_data_set

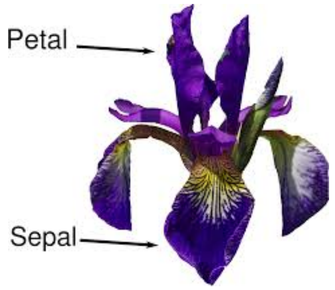

petal
width

```
First five columns of data:
[[ 5.1  3.5  1.4  0.2]
 [ 4.9  3.   1.4  0.2]
 [ 4.7  3.2  1.3  0.2]
 [ 4.6  3.1  1.5  0.2]
 [ 5.   3.6  1.4  0.2]]
```

Colour
Width
Length
Animate
Scent

Features

(2,1,1,0,5)

“In machine learning, a feature is an individual measurable property or characteristic of a phenomenon being observed” Bishop

Even behind data there (should be) humans

- **Feature engineering** is the process of using domain knowledge of the data to create **features** that make machine learning algorithms work. Also: how to represent your data best for a particular application
- In machine learning more is not better, sometimes it is just more
- **Feature extraction** is the process of transforming the initial measured values into features. Eg: remove redundancy, transform data type, bundle together
- What works best in ML: continuous features. In the wild: categorical (discrete) features

Supervised learning

- Given: a set of input-output pairs
- Learned: predicting output for given input

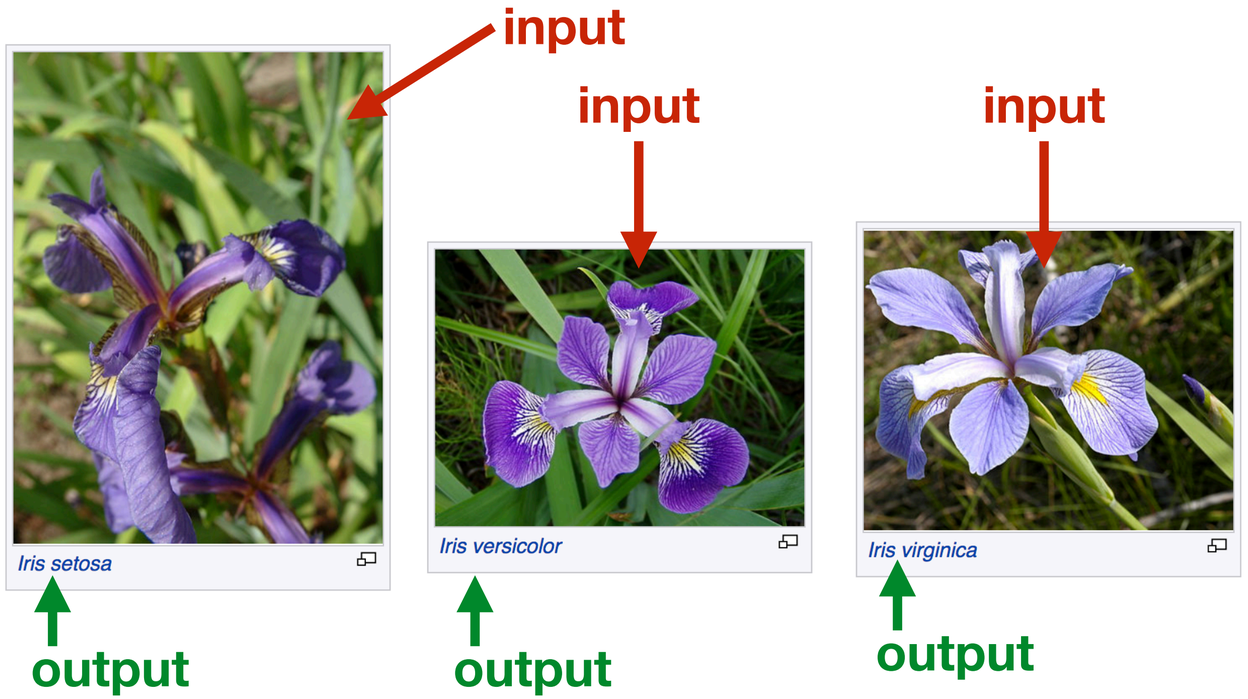

- Supervised learning involves “observing” several examples of a random vector **x** and an associated value or vector **y**, then “learning” to predict y from x:
  output
  - by estimating p(**y**|**x**)
  - by estimating F(**x**) = **y**
- To understand the above a background is needed in linear algebra and probability theory

Linear algebra

- Algebra - the study of (mathematical) structures and the rules for manipulating these structures
- Linear - special types of structure

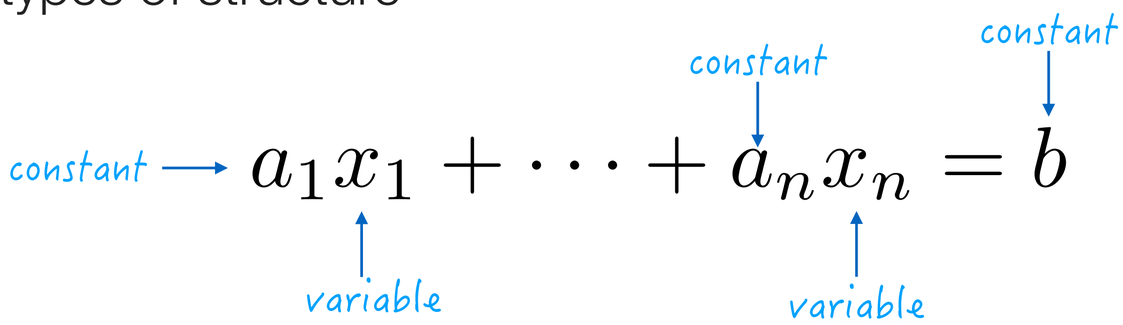

no power for variables if x^2 would be allowed at least once we call it quadratic

https://www.desmos.com/calculator

Linear algebra

$$2x=4$$

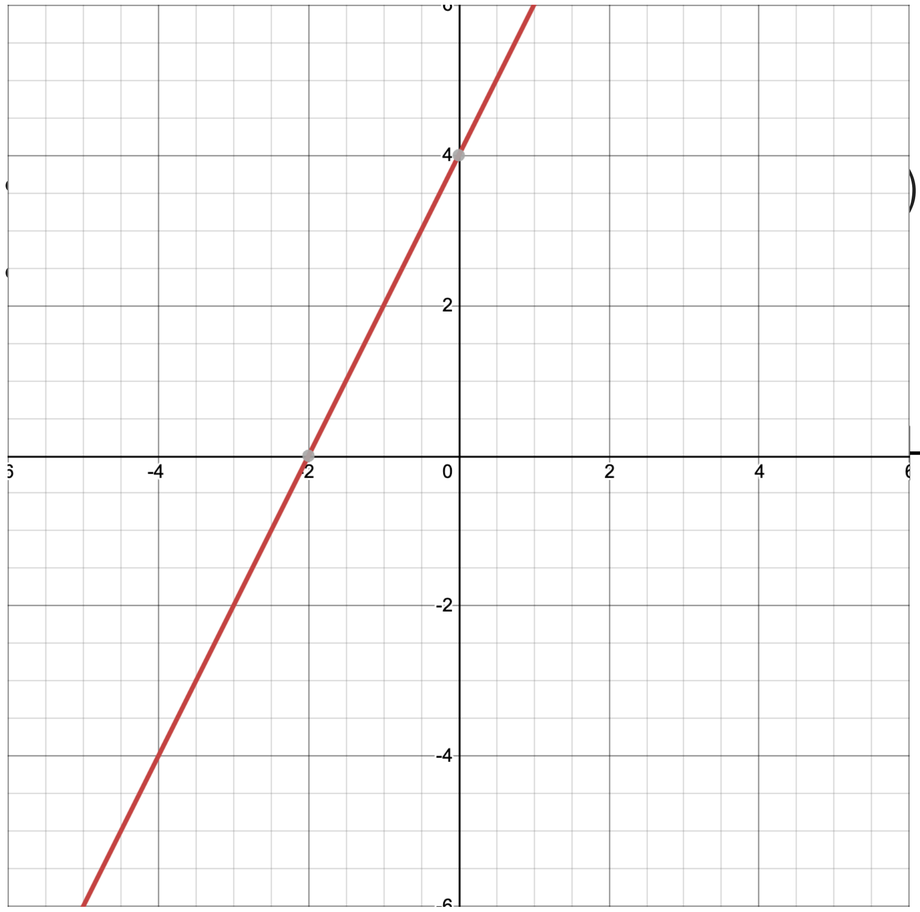

) structures and the rules for manipulating these structures

$$+ \cdots + a_n x_n = b$$

constant

constant

variable

no power for variables if x^2 would be allowed at least once we call it quadratic

https://www.desmos.com/calculator

Linear algebra

$$2x=4$$

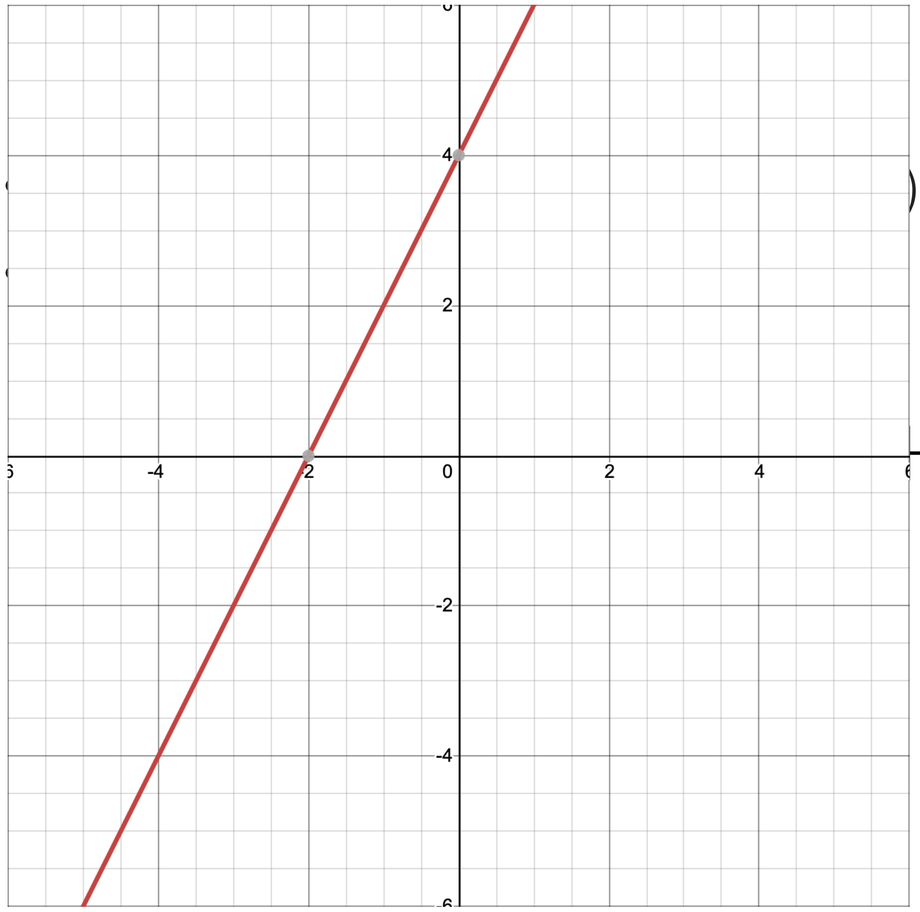

) structures and the rules for manipulating these structures

$$+ \cdots + a_n x_n = b$$

constant

bias

variable

no power for variables if x^2 would be allowed at least once we call it quadratic

https://www.desmos.com/calculator

Linear algebra

- Algebra - the study of (mathematical) structures and the rules for manipulating these structures
- Linear - special types of structure

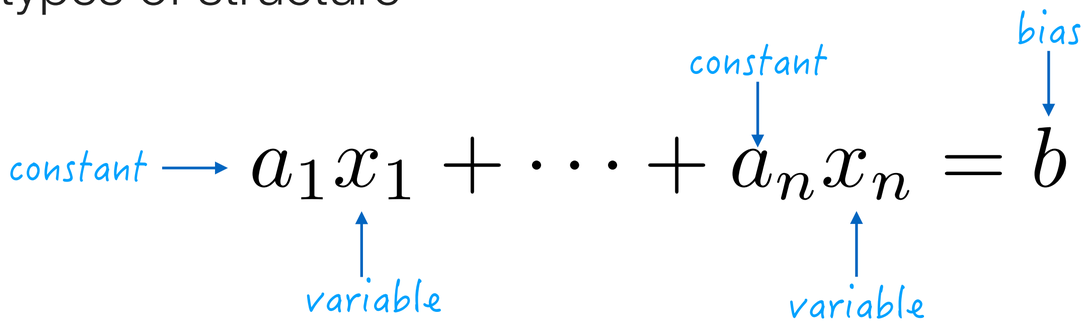

no power for variables if x^2 would be allowed at least once we call it quadratic

- Scalar - single number
- Vector - array of numbers.
- An **array** is a container object that holds a fixed number of values of a single type.

https://www.desmos.com/calculator

Vector, matrix, tensor

Vector

- In a vector the scalars are arranged in an order.
- We can identify each number by its index.
- Vector notation - bold small case. Eg. **x**

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$$

Matrix

- Matrix is a two dimensional array of numbers.
- We can identify each number in a matrix by **two** indices.
- Matrix notation - bold uppers case. Eg. **A**

$$\mathbf{A} = \begin{bmatrix} A_{1,1}, A_{1,2}, \cdots, A_{1_m} \\ A_{2,1}, A_{2,2}, \cdots, A_{2_m} \\ \vdots \\ A_{n,1}, A_{n,2}, \cdots, A_{n_m} \end{bmatrix}$$

Tensor

..you know..like in TensorFlow

- Tenstor is an array of numbers arranged on a regular grid with a varian number of axes.
- An *n*-ranked tensor has *n* indeces.
- Used to describe a matrix of values together with how those values are transformed by some function

Tensor

..you know..like in TensorFlow

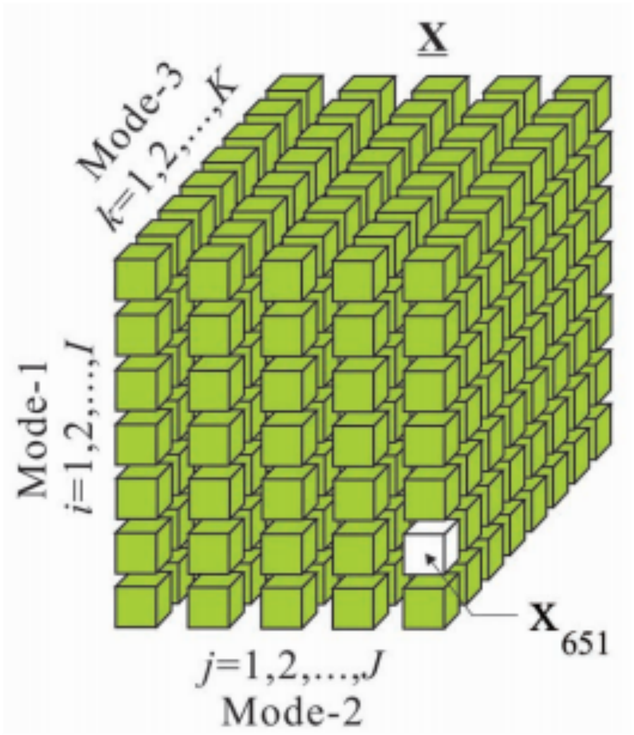

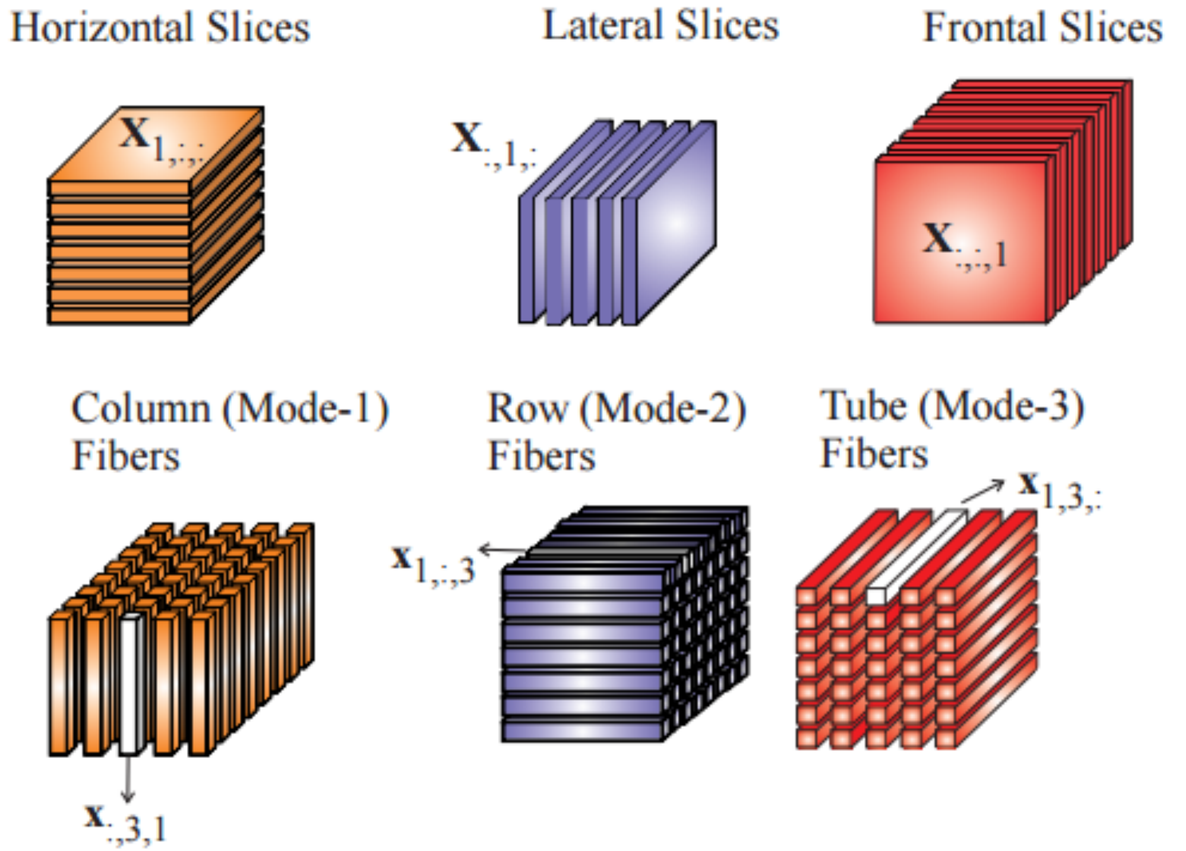

Figure 2: A 3rd-order tensor $\underline{\mathbf{X}} \in \mathbb{R}^{I \times J \times K}$, with entries $x_{i,j,k} =$

Conditional probability

- Conditional probability is the probability of some event y, given that some other event x has happened p(y|x).

$$p(y \mid x) = \frac{p(y,x)}{p(x)}$$

Back to supervised learning

- Given a training set of N example input-output pairs (x_1,y_1), $(x_2,y_2),\ldots,(x_N,y_N)$ where each y_j was generated by an unknown function $y=f(x)$, discover a function h that approximates the true function f.

hypothesis

- Why discover a function h that approximates the true function f. Why not f directly?

Unsupervised learning

no output available, only input

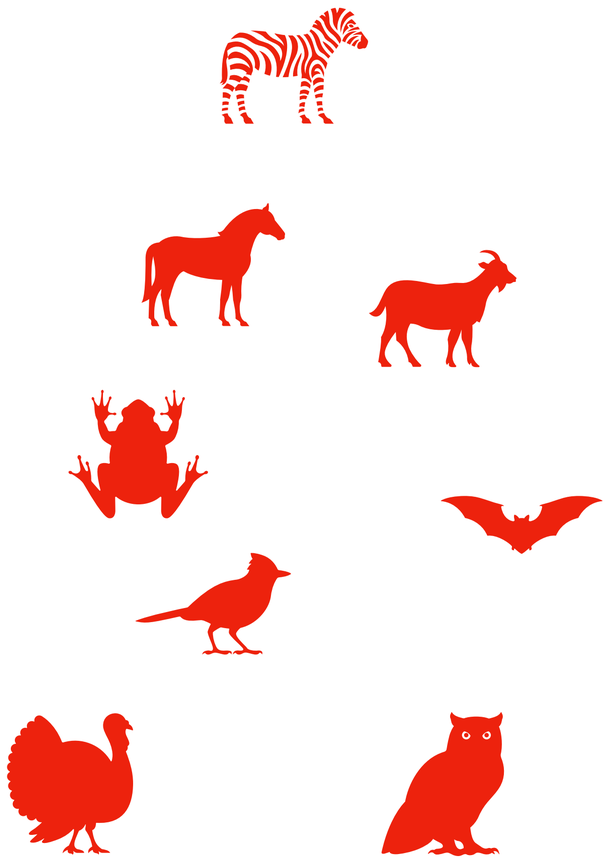

- Find patterns in the input (although no explicit feedback is supplied = no correct answer)
- Types: clustering, data transformations (eg visualisation), density estimation

Unsupervised learning

no output available, only input

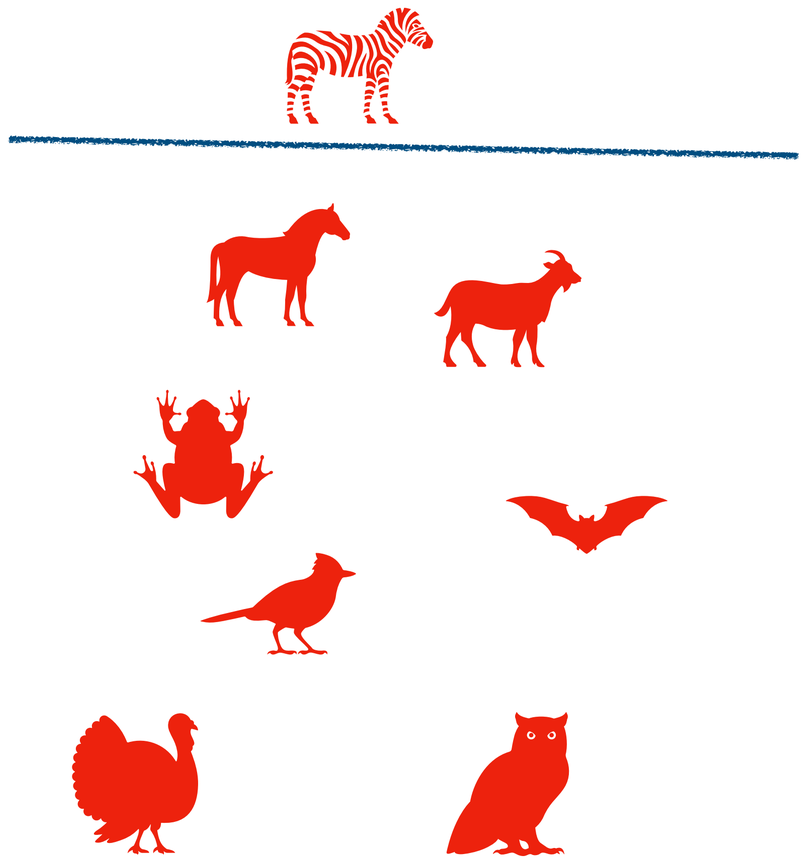

- Find patterns in the input (although no explicit feedback is supplied = no correct answer)
- Types: clustering, data transformations (eg visualisation), density estimation

Unsupervised learning

no output available, only input

- Find patterns in the input (although no explicit feedback is supplied = no correct answer)
- Types: clustering, data transformations (eg visualisation), density estimation

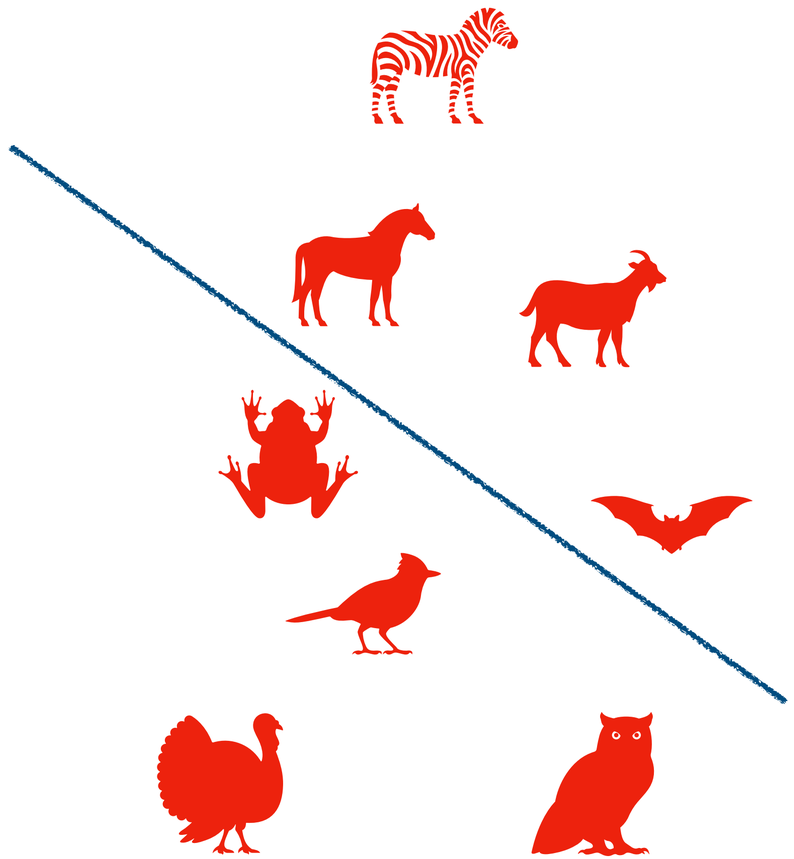

Unsupervised learning

no output available, only input

- Find patterns in the input (although no explicit feedback is supplied = no correct answer)
- Types: clustering, data transformations (eg visualisation), density estimation

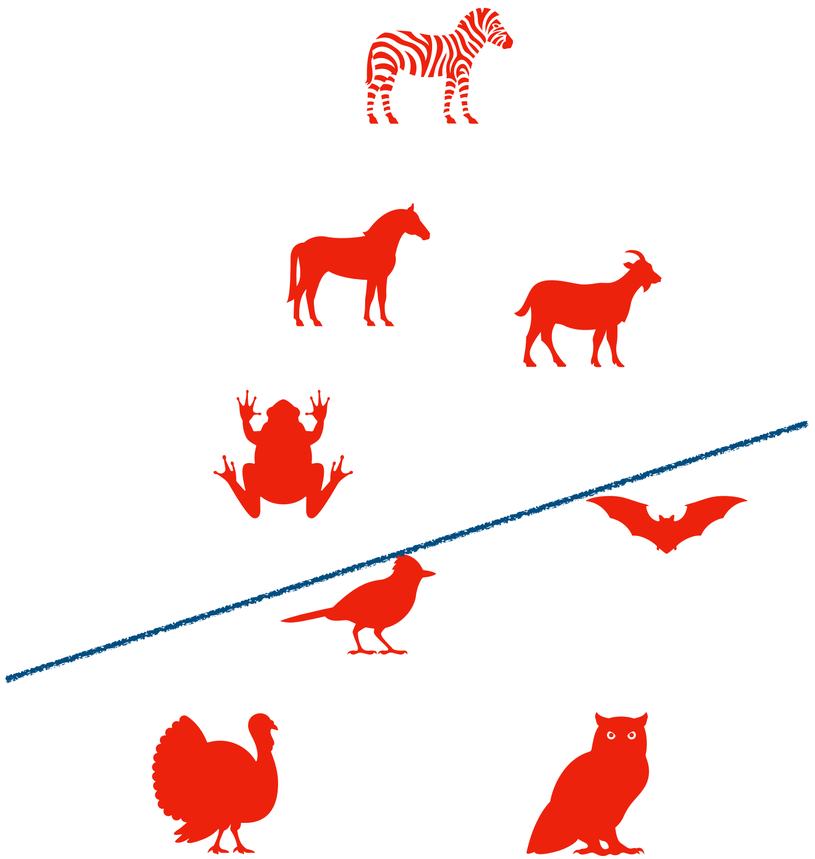

Unsupervised learning

no output available, only input

- Find patterns in the input (although no explicit feedback is supplied = no correct answer)
- Types: clustering, data transformations (eg visualisation), density estimation

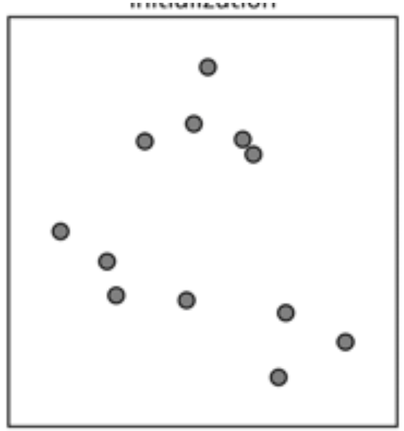

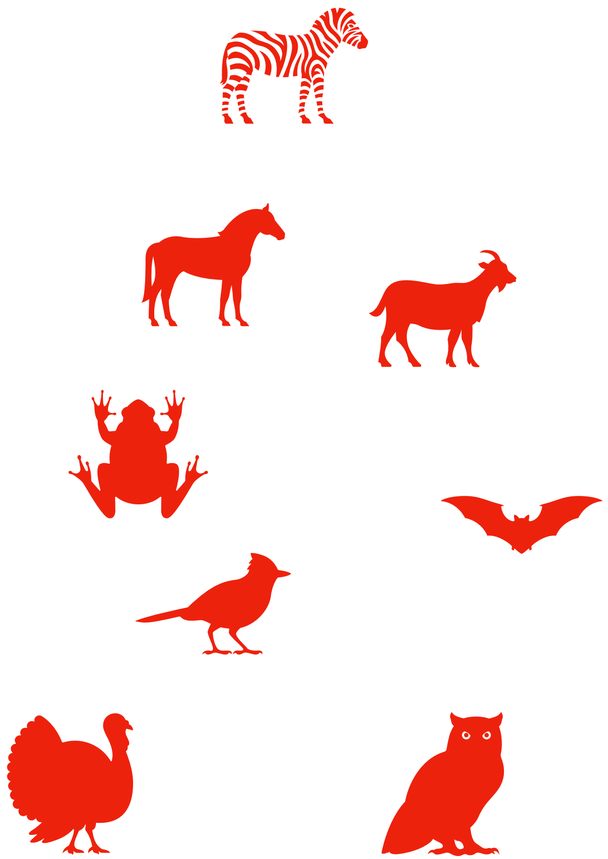

Reinforcement learning

- An agent takes actions in an environment and observes the rewards (positive utility) and punishments (costs = negative utility). Identify a strategy (which action to take in which situation, also called a policy) that maximises a reward function how to maximise an objective from a series of reinforcements - rewards and punishments
- Relies on Markov Decision Processes as a formalism

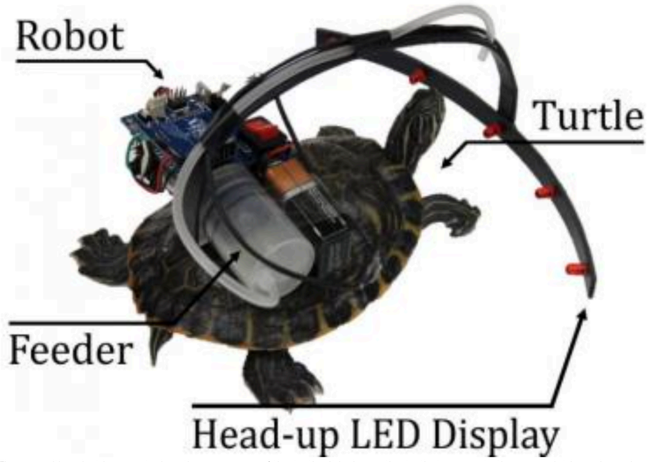

The parasitic robot is mounted on the carapace of the turtle. It induces the turtle to move to the waypoint by using a heads-up LED display as well as rewarding the turtle. Source: KAIST


SCORE
3470
9
3

Reinforcement learning

why (relatively) limited to games?

- The world is huge and unpredictable
- Thousands of trials and errors are needed to be run for a good strategy to be found. Not all of them safe
- Finding a good objective function (what to reward and what to punish) is difficult.

Reinfo

why (relat

- The world is
- Thousands o o be found. Not a
- Finding a goo difficult.

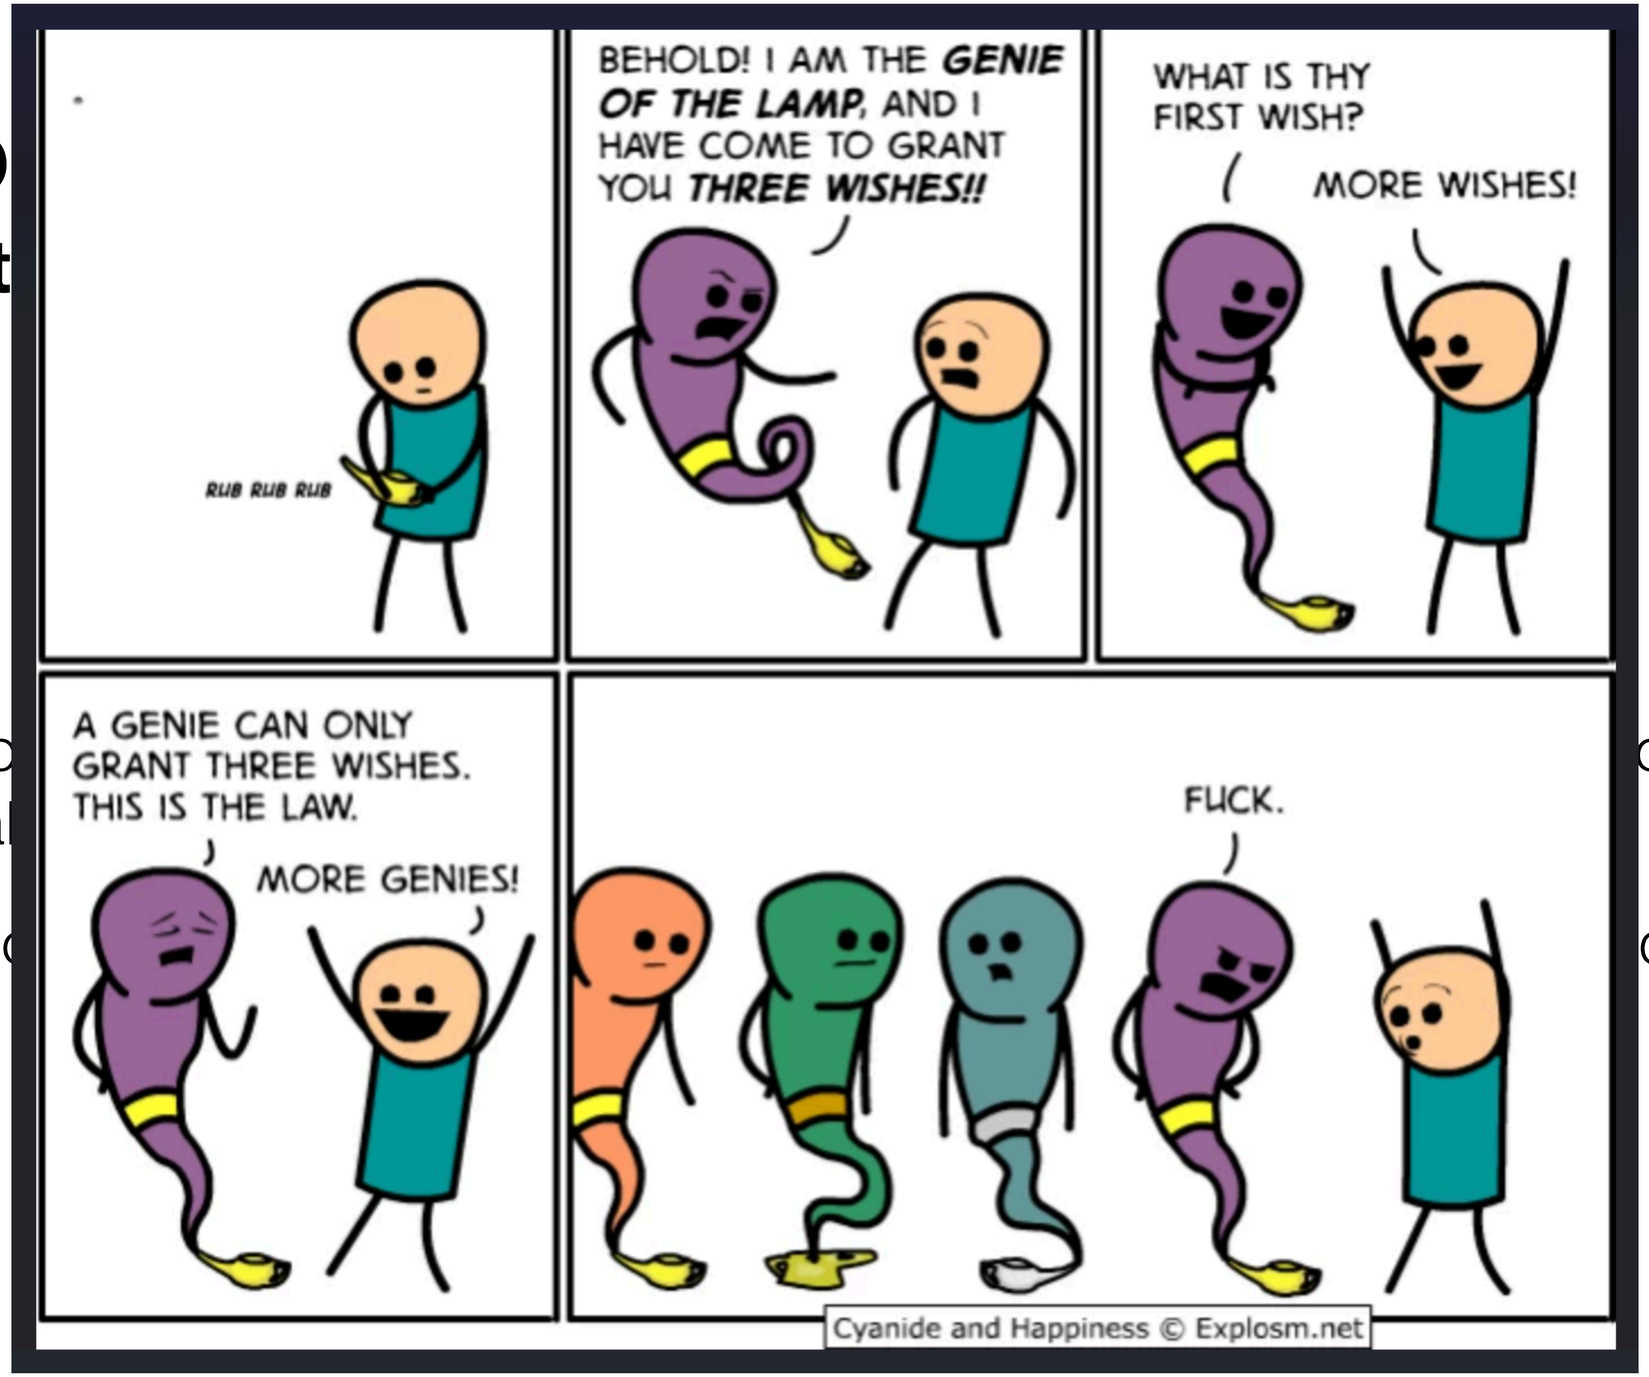

Agents in AI

An entity that acts in an environment.

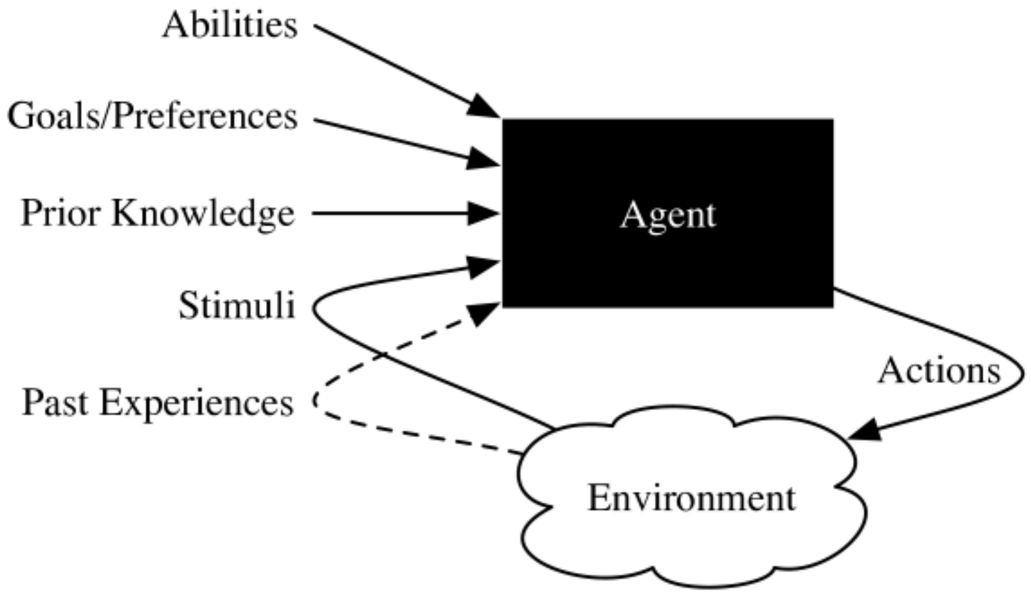

An agent as an input-output system

The agent's actions depend on:

- **prior knowledge** about the agent and the environment
- **history of interaction** with the environment, which is composed of
  - – stimuli received from the current environment, which can include observations about the environment, as well as actions that the environment imposes on the agent and
  - – past experiences of previous actions and stimuli, or other data, from which it can learn
- **goals** that it must try to achieve or preferences over states of the world
- **abilities**, the primitive actions the agent is capable of carrying out.

Types of agents

- Reactive agent: when triggered by a particular stimuli, executes a specified action
- Rational agent: choices are made to advance a goal
- Intelligent agent: is rational and can adapt to an environment = can learn to do better
- Computational agent: uses computation to agent
- Embodied agent: another word for robot
- Software agent: no body, just code
- Autonomous agent: can operate without human oversight and control for some period of time

Multi Agent Systems

- How to solve problems in an environment, alone or in team
- How to reason when your choices and outcomes can depend on other agents?
- Analysis of events when multiple agents interact (by simulation)