



# Artificial Intelligence (AI)

Ong Wee Hong

Assistant Professor

Robotics and Intelligent Systems Lab (Robolab), School of Digital Science

Universiti Brunei Darussalam

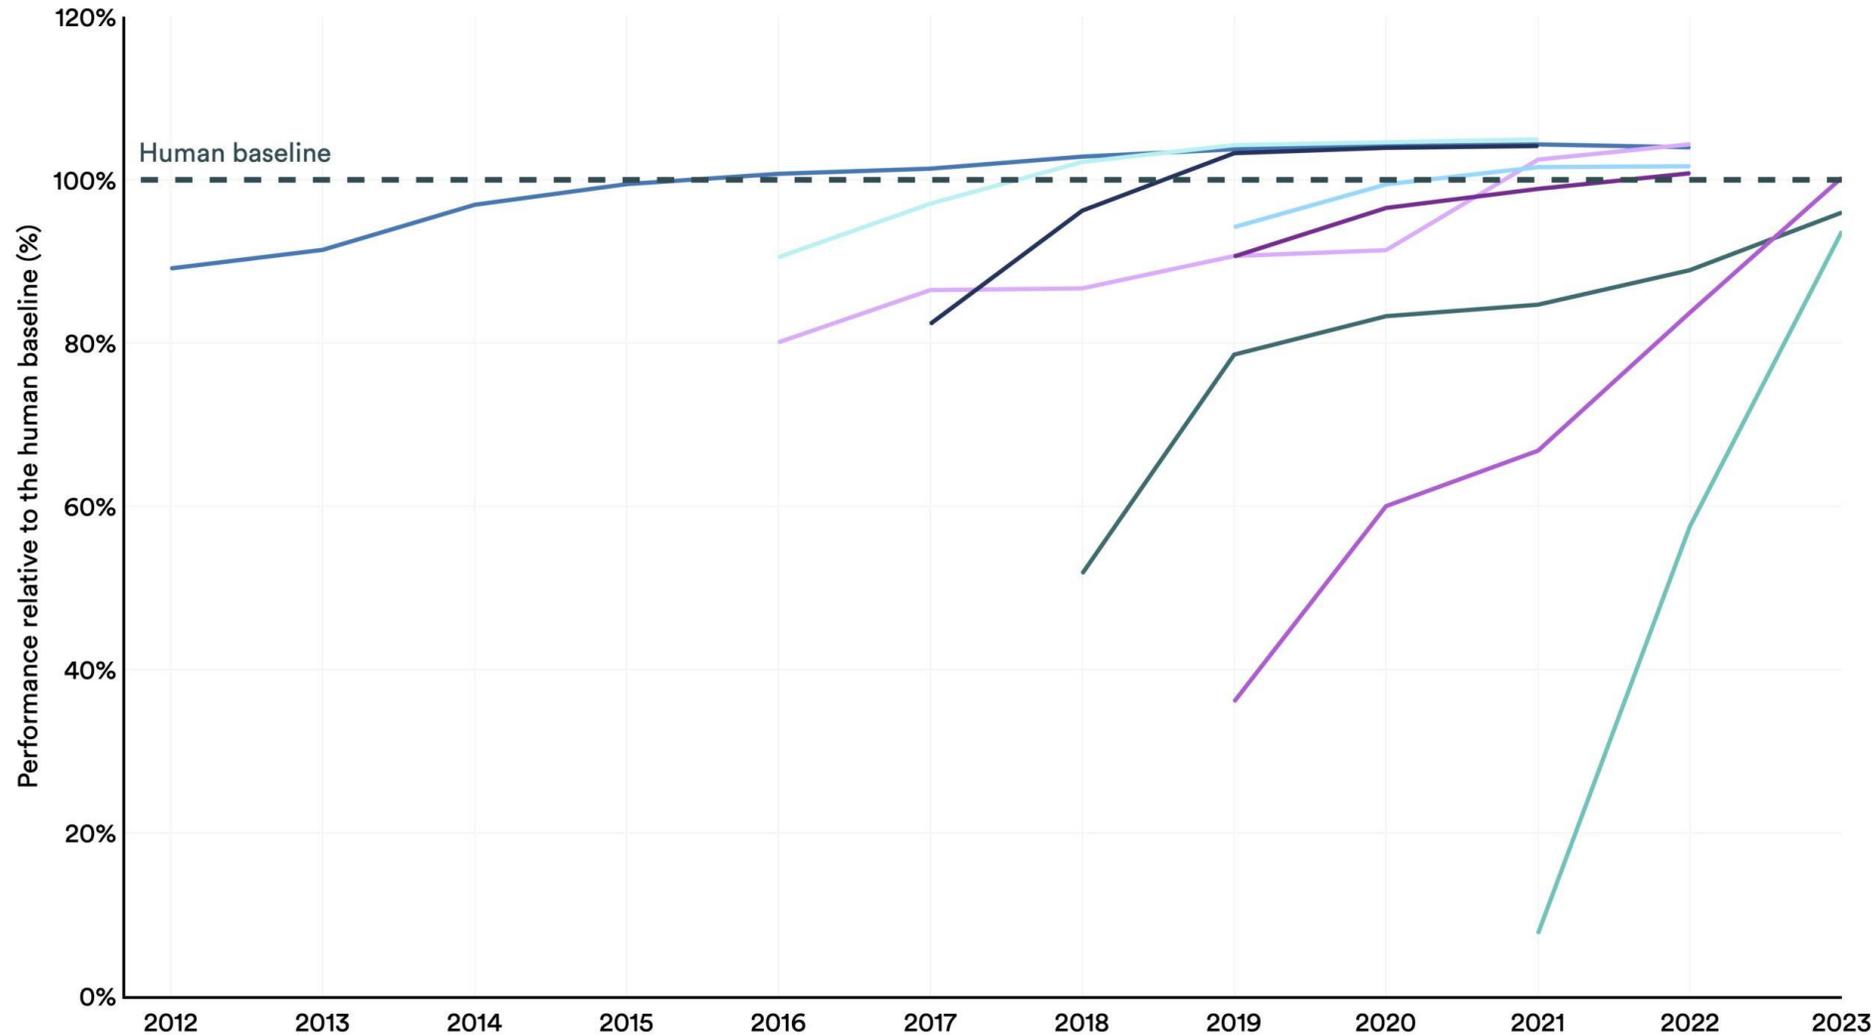
Everyone knows AI



**ChatGPT**

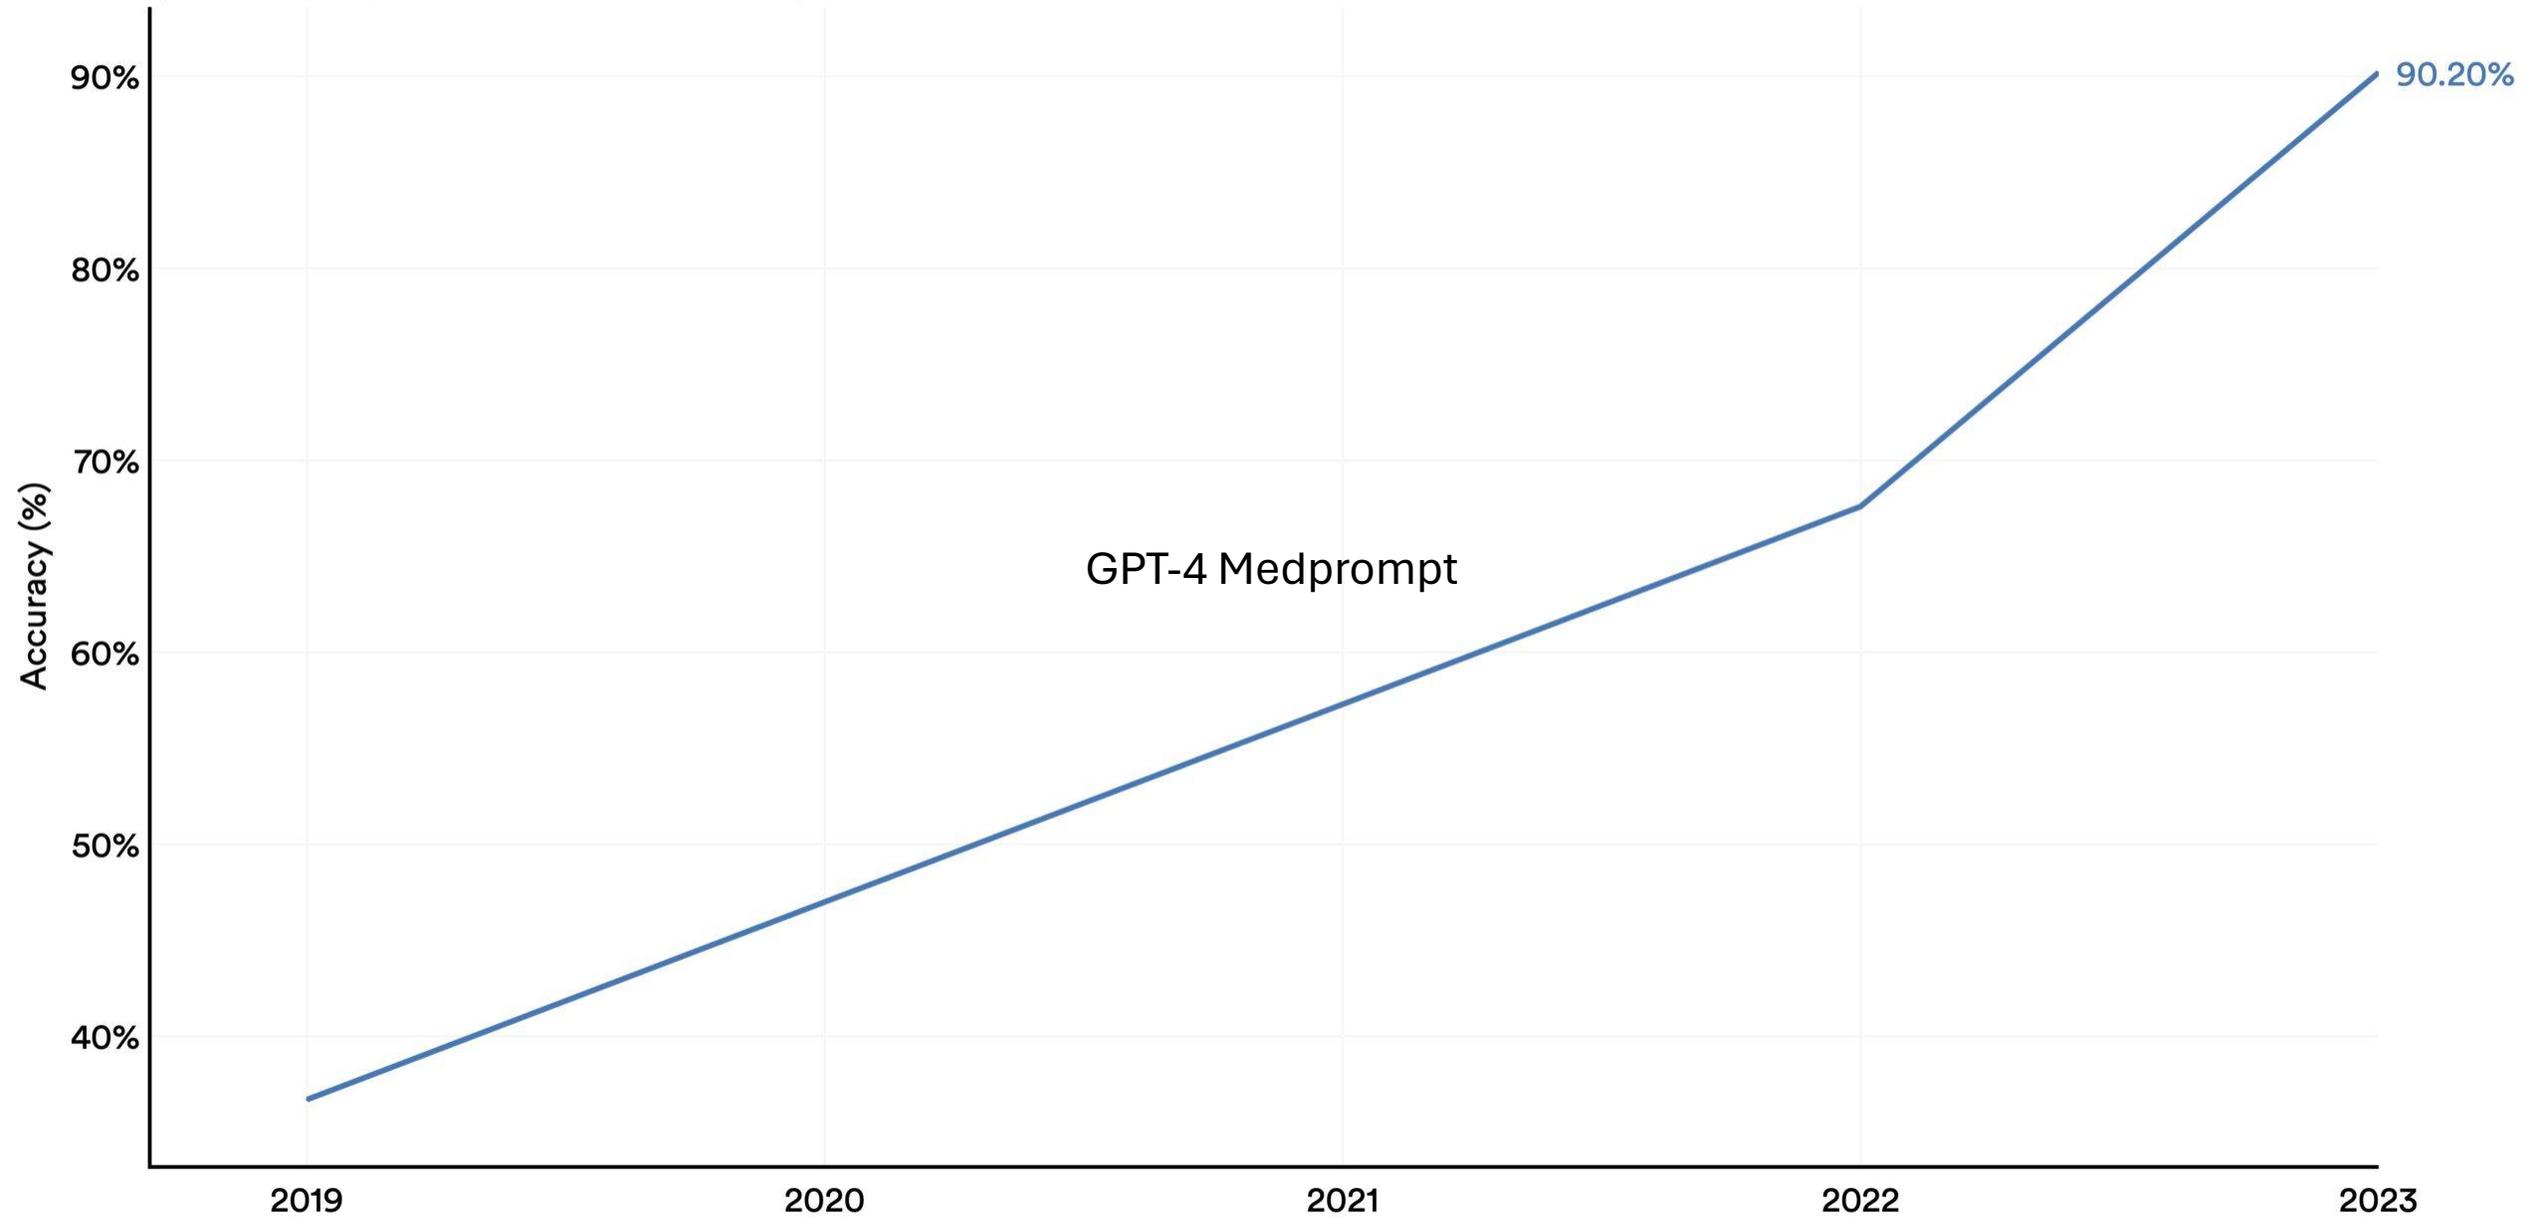
# Select AI Index technical performance benchmarks vs. human performance

Source: AI Index, 2024 | Chart: 2024 AI Index report



# MedQA: accuracy

Source: Papers With Code, 2023 | Chart: 2024 AI Index report



# Timeline of images generated by artificial intelligence

These people don't exist. All images were generated by artificial intelligence.

2014



Goodfellow et al. (2014) - Generative Adversarial Networks

2015



Radford, Metz, and Chintala (2015) - Unsupervised Representation Learning with Deep Convolutional GANs

2016



Liu and Tuzel (2016) - Coupled GANs

2017



Karras et al. (2017) - Progressive Growing of GANs for Improved Quality, Stability, and Variation

2018



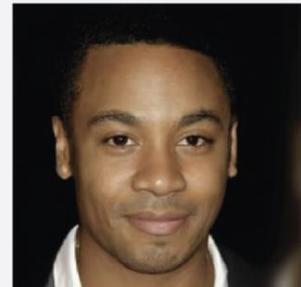
Karras, Laine, and Aila (2018) - A Style-Based Generator Architecture for Generative Adversarial Networks

2019



Karras et al. (2019) - Analyzing and Improving the Image Quality of StyleGAN

2020



Ho, Jain, & Abbeel (2020) - Denoising Diffusion Probabilistic Models

2021 Image generated with the prompt: "a couple of people are sitting on a wood bench"



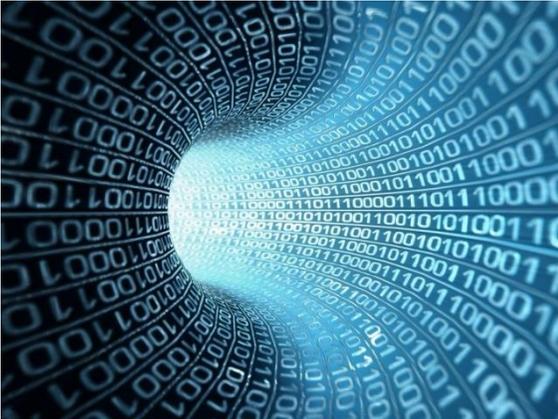
Ramesh et al. (2021) - Zero-Shot Text-to-Image Generation (OpenAI's DALL-E 1)

2022 Image generated with the prompt: "A Pomeranian is sitting on the King's throne wearing a crown. Two tiger soldiers are standing next to the throne."



Saharia et al. (2022) - Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding (Google's Imagen)

# What drive AI's growth?



Big data



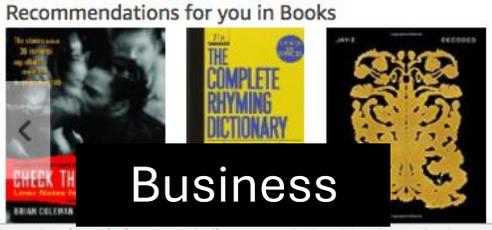
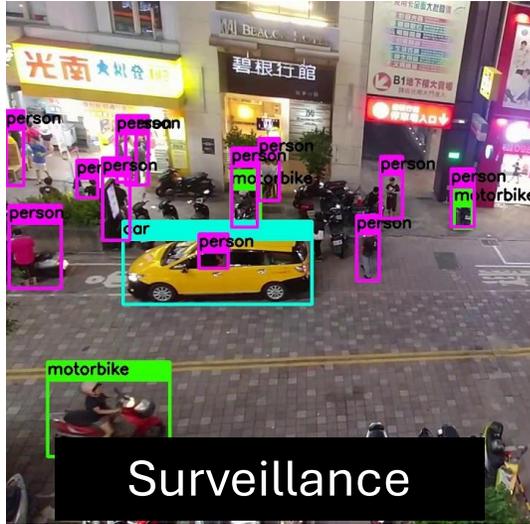
Powerful compute



Advance algorithms

# AI applications

- Healthcare
- Automobile
- Business
- Finance
- Entertainment
- Education
- Gaming
- Manufacturing
- Transportation
- Robotics
- Space Exploration
- Surveillance
- Agriculture



Business

# What is AI?

Intelligence exhibited by artificial systems: machines or software.

The study of AI is about the **creation of intelligence** similar to human (and animals).



Natural Intelligence

# Is calculator AI?

$$(2^3 \cdot 5^2) - \frac{100}{2} + \sqrt{144} + \log_{10}(1000) = ?$$



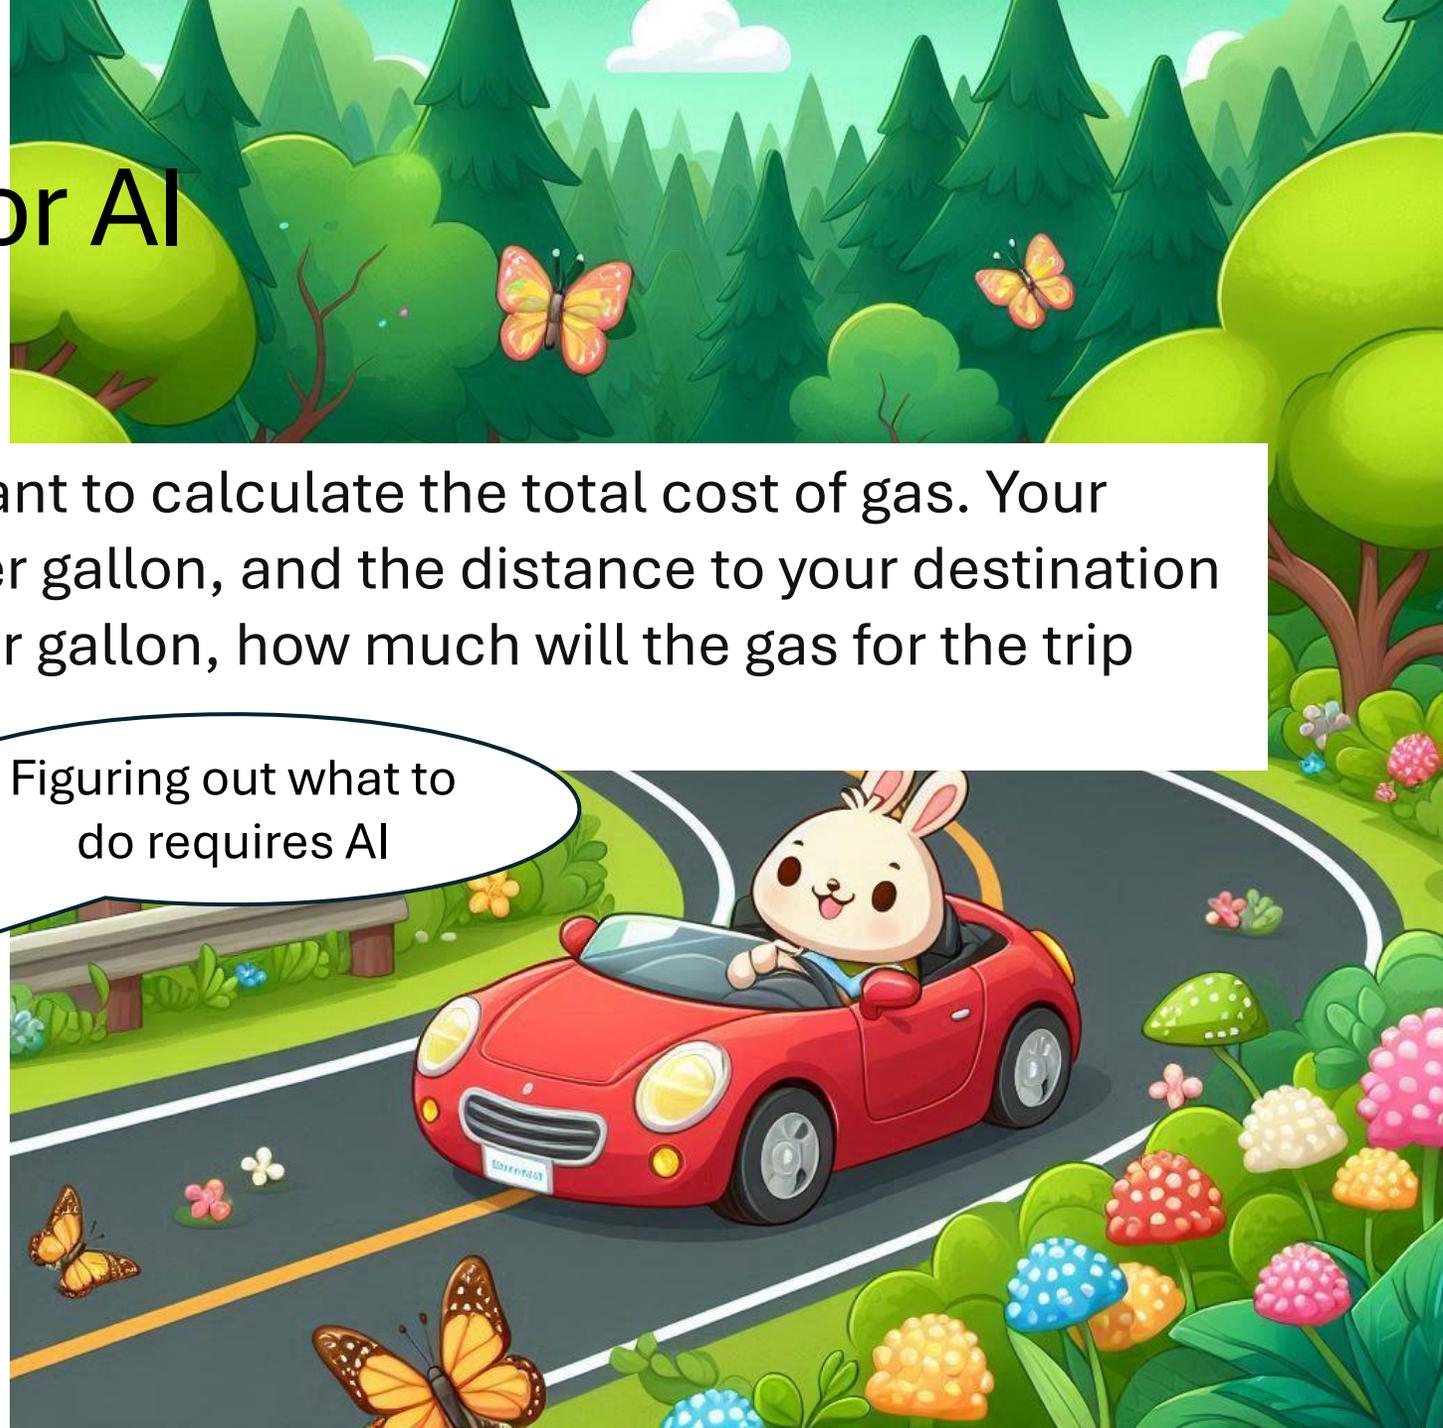
Image credit: Wikipedia

# A simple problem for AI

You're planning a road trip and want to calculate the total cost of gas. Your car's fuel efficiency is 25 miles per gallon, and the distance to your destination is 300 miles. If gas costs \$3.50 per gallon, how much will the gas for the trip cost?

$$\text{Cost} = \$3.50 \times \frac{300}{25}$$

Figuring out what to do requires AI



# More problems for AI

- Prove a theorem
- Play chess
- Plan a surgical operation
- Diagnose a disease
- Navigate in a building
- ...

# Four views of AI

**Thinking Humanly**

**Thinking Rationally**

Unpredictable

Doable

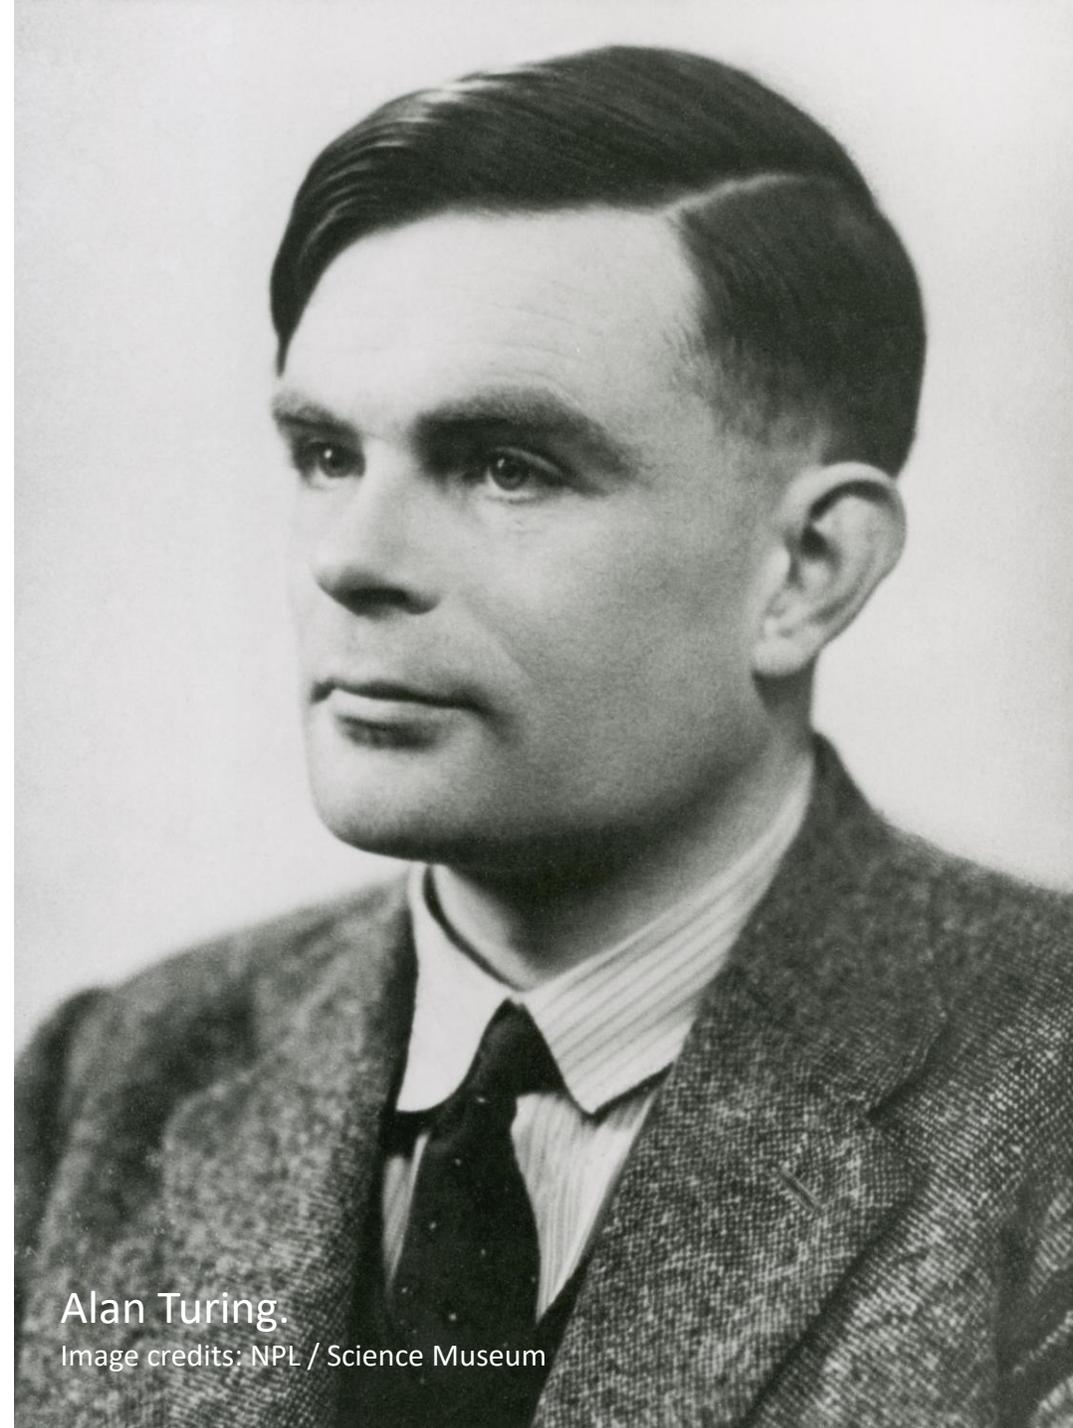
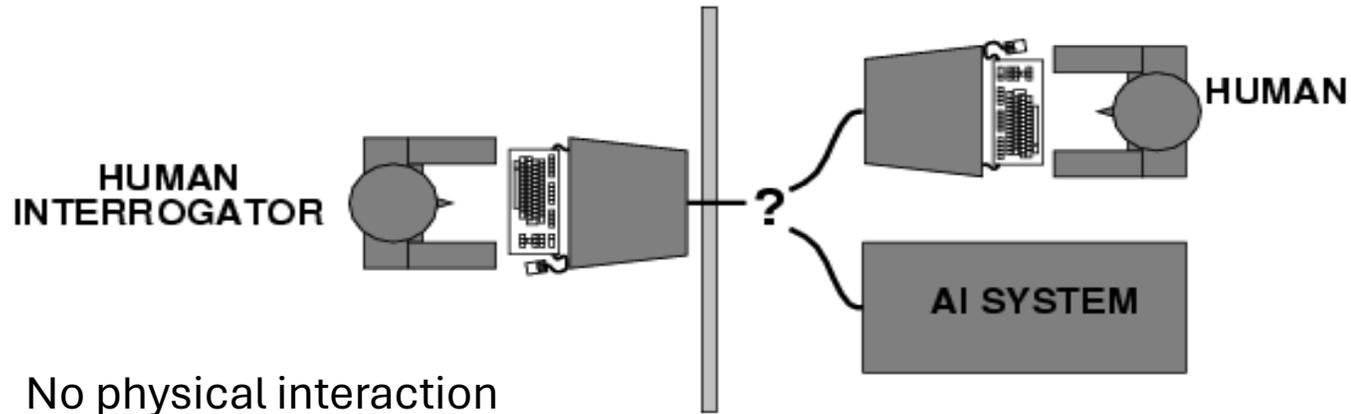
**Acting Humanly**

**Acting Rationally**

1950

# Turing Test

"A computer would deserve to be called intelligent if it could **deceive a human into believing that it was human.**" (Acting humanly)



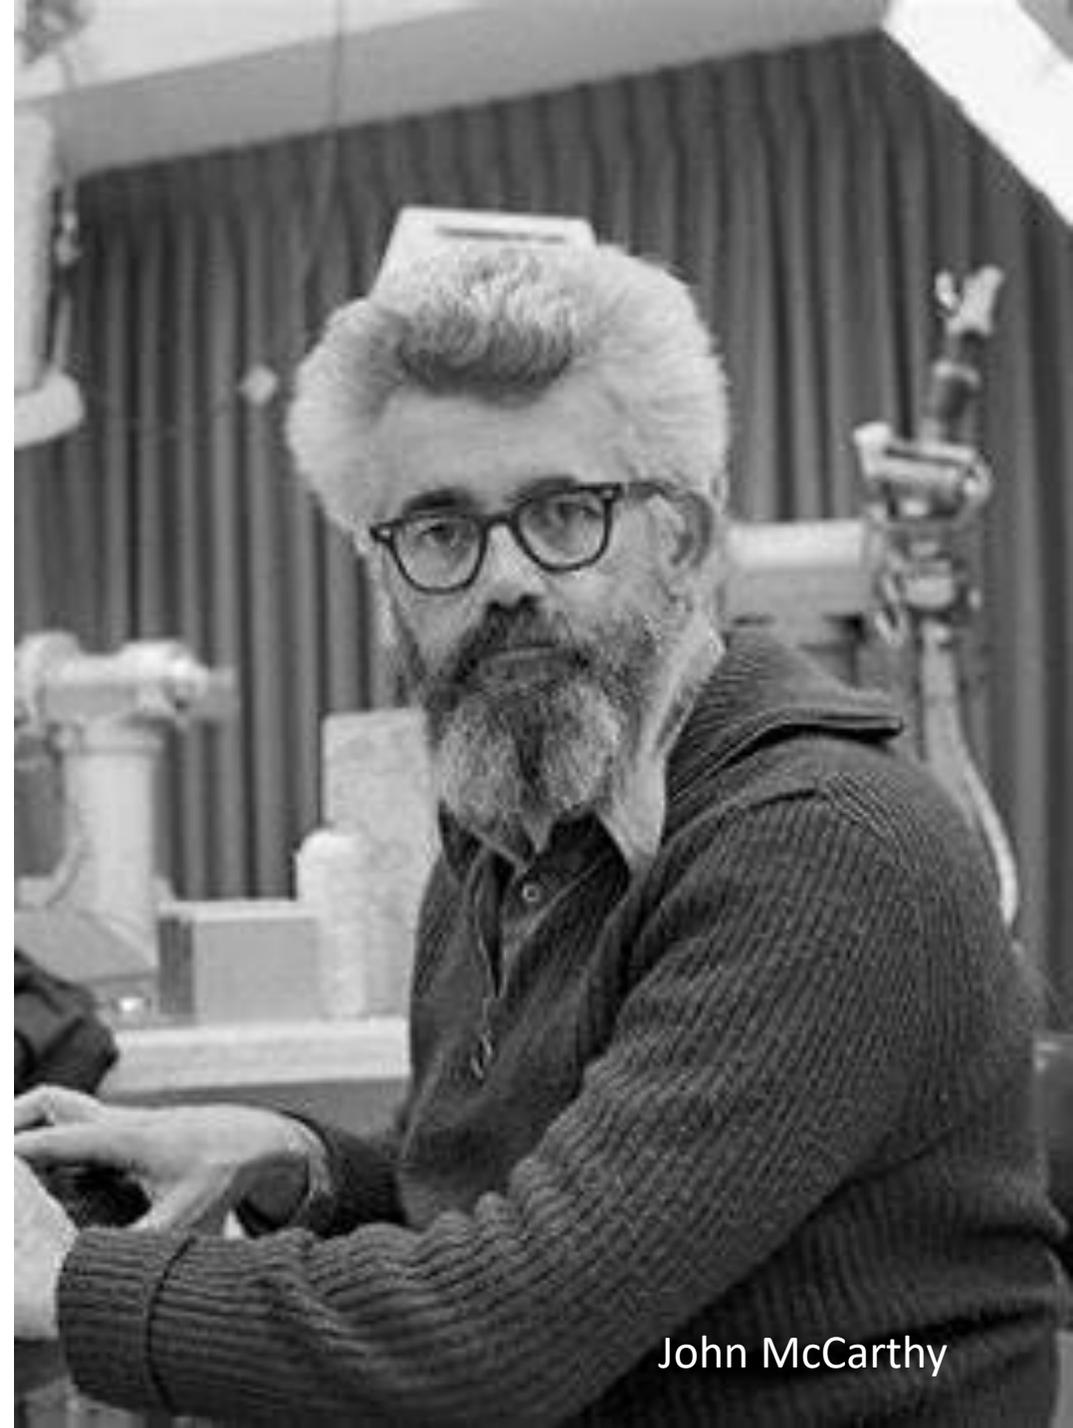
Alan Turing.

Image credits: NPL / Science Museum

1956

# Birth of AI

"The ultimate goal of artificial intelligence is to create computer programs that can **solve problems** and achieve goals **without needing explicit instructions for each step.**"



John McCarthy

# Capabilities of AI

- Knowledge **representation**
- Natural **language** processing (NLP)
- Automated **reasoning** and **problem solving**
- Machine **learning**
- **Perception** (vision, audio, etc)
- **Robotics**

} for intelligence

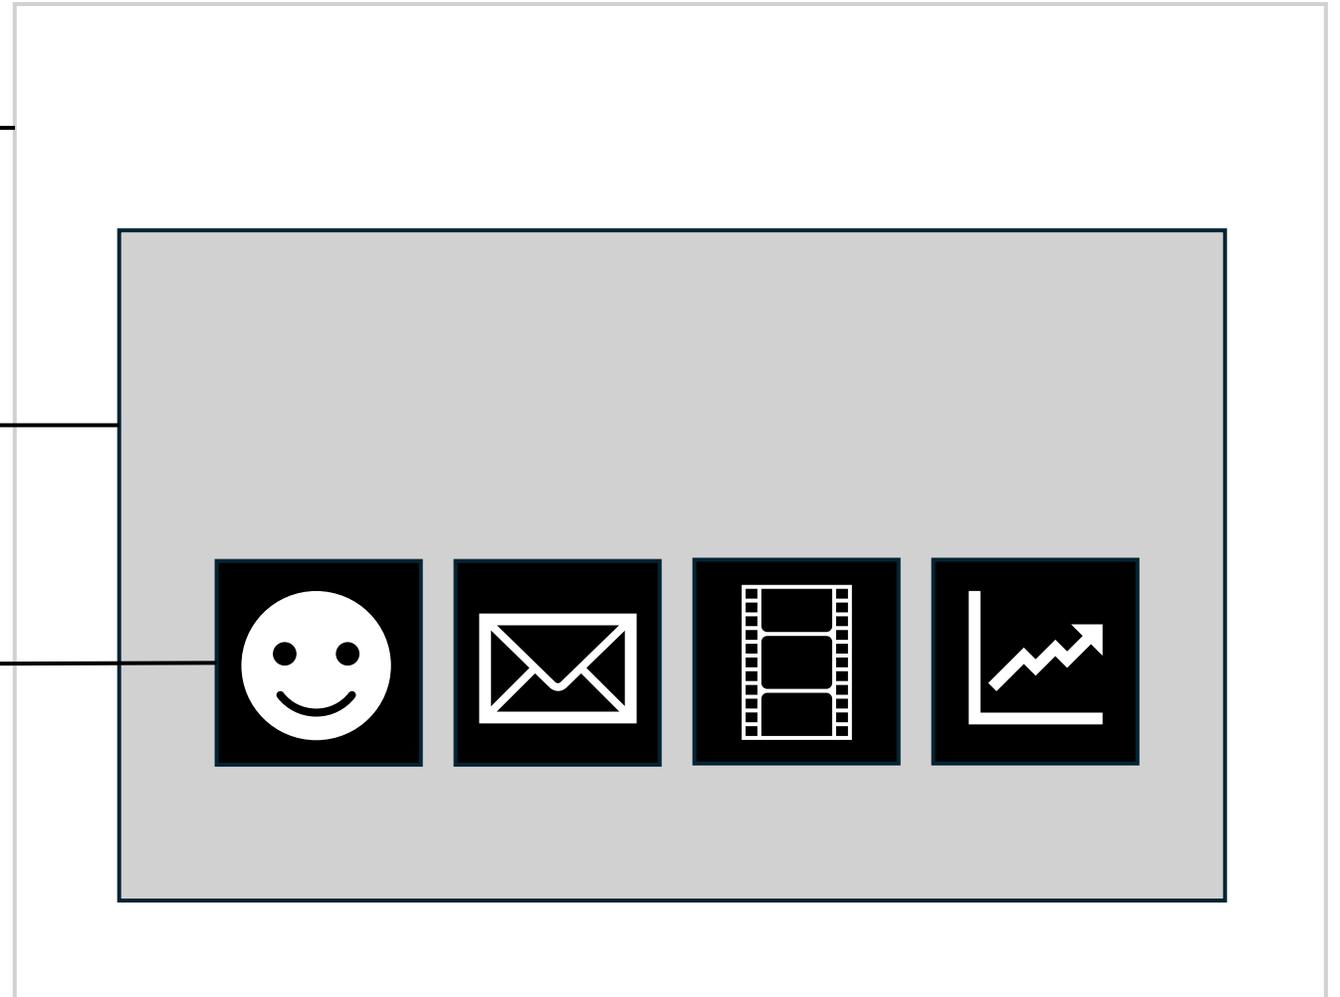
} for interaction  
(Total Turing Test)

# Levels of AI

Artificial  
Superintelligence

Artificial General AI (AGI)  
(Strong AI)

Weak (Narrow) AI



# AI system

Artifact

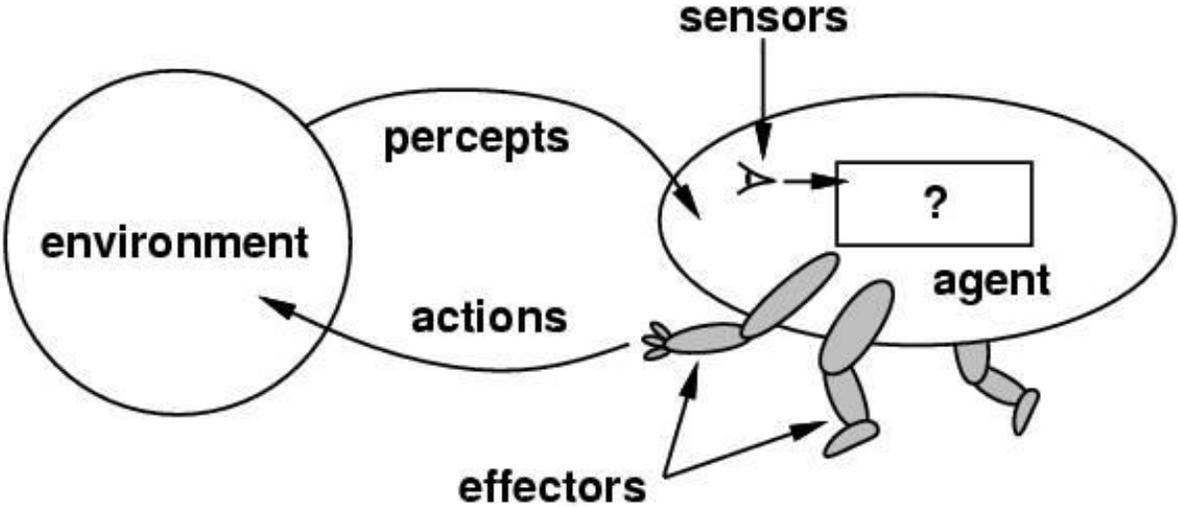
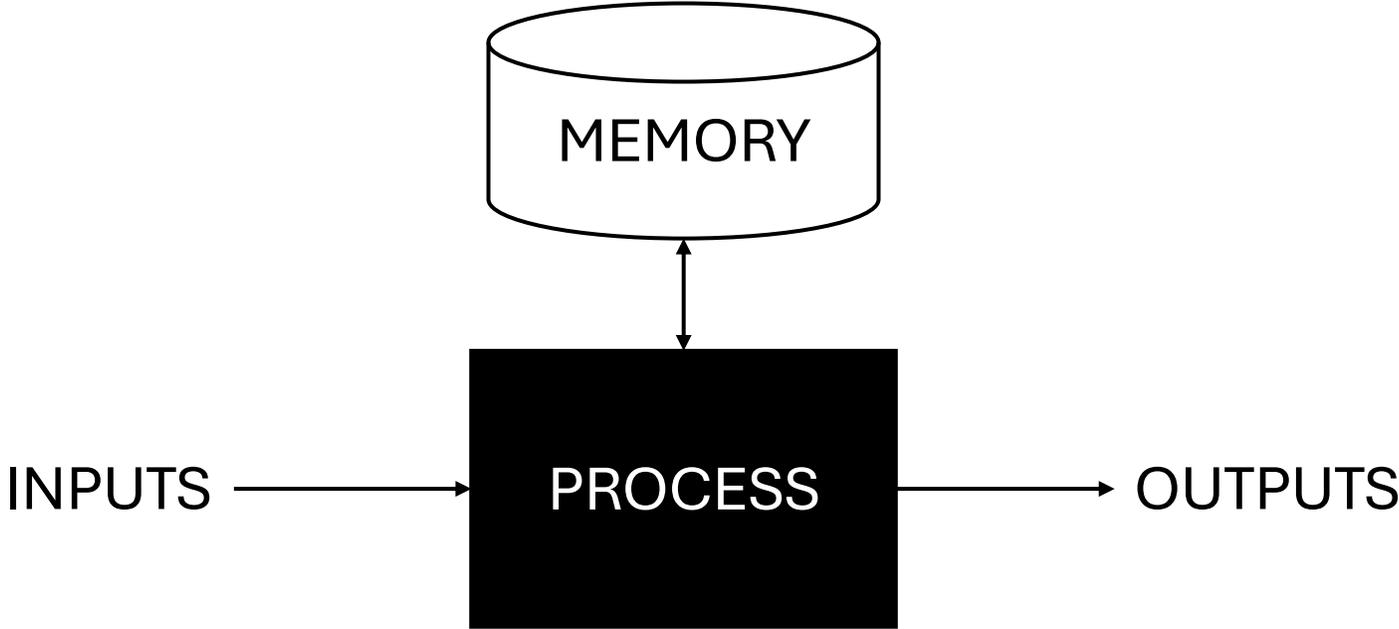


+

Intelligence



AI Programs

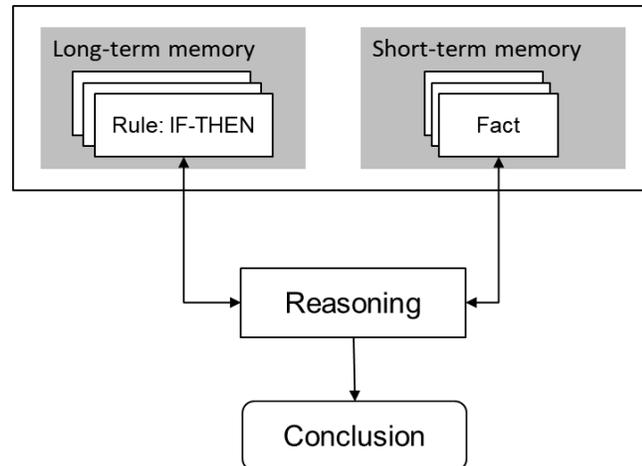
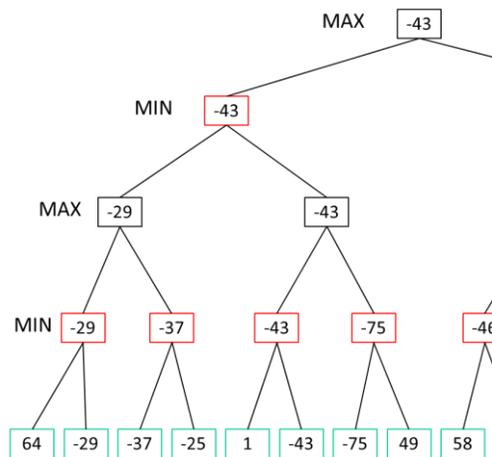


# Two paradigms of AI techniques

Dominates

## Symbolic AI

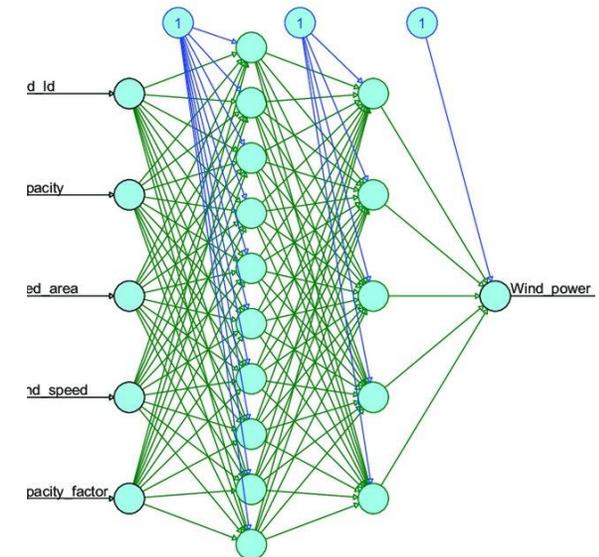
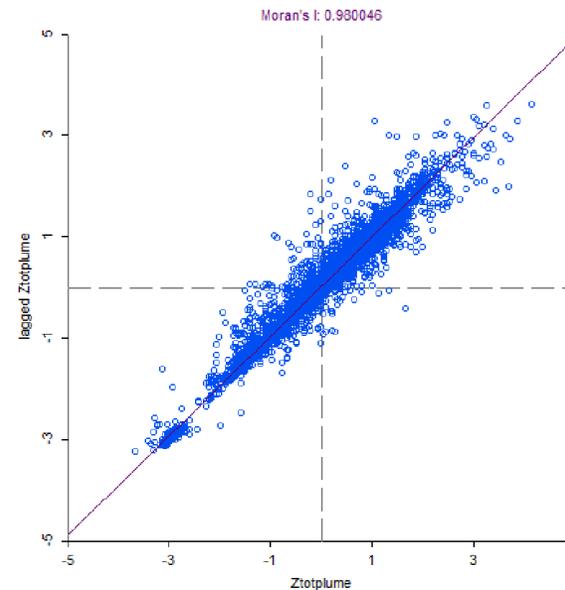
aka Interpretational AI, Good Old Fashion AI (GOFAI), Traditional AI



## Non-symbolic AI

aka Sub-symbolic, Modern AI

- Comprises of **Statistical AI** and **Neural (Connectionist) AI**
- Data driven



# AI techniques

## **Symbolic AI**

aka Interpretational AI, Good Old Fashion AI (GOFAI), Traditional AI

- Problem solving by search
- Knowledge-based system

## **Non-symbolic AI**

aka Sub-symbolic, Modern AI

- Machine Learning
  - Linear regression
  - Logistic regression
  - Artificial Neural Network
  - K-Means clustering



Dominates

# AI techniques

## **Symbolic AI**

aka Interpretational AI, Good Old Fashion AI (GOFAI), Traditional AI

- Problem solving by search
- Knowledge-based system

## **Non-symbolic AI**

aka Sub-symbolic, Modern AI

- Machine Learning
  - Linear regression
  - Logistic regression
  - Artificial Neural Network
  - K-Means clustering

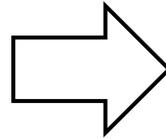
# Problem solving by Search

When we don't know much, we just try

# Problem solving by Search

8	2	
3	4	7
5	1	6

Initial state



1	2	3
4	5	6
7	8	

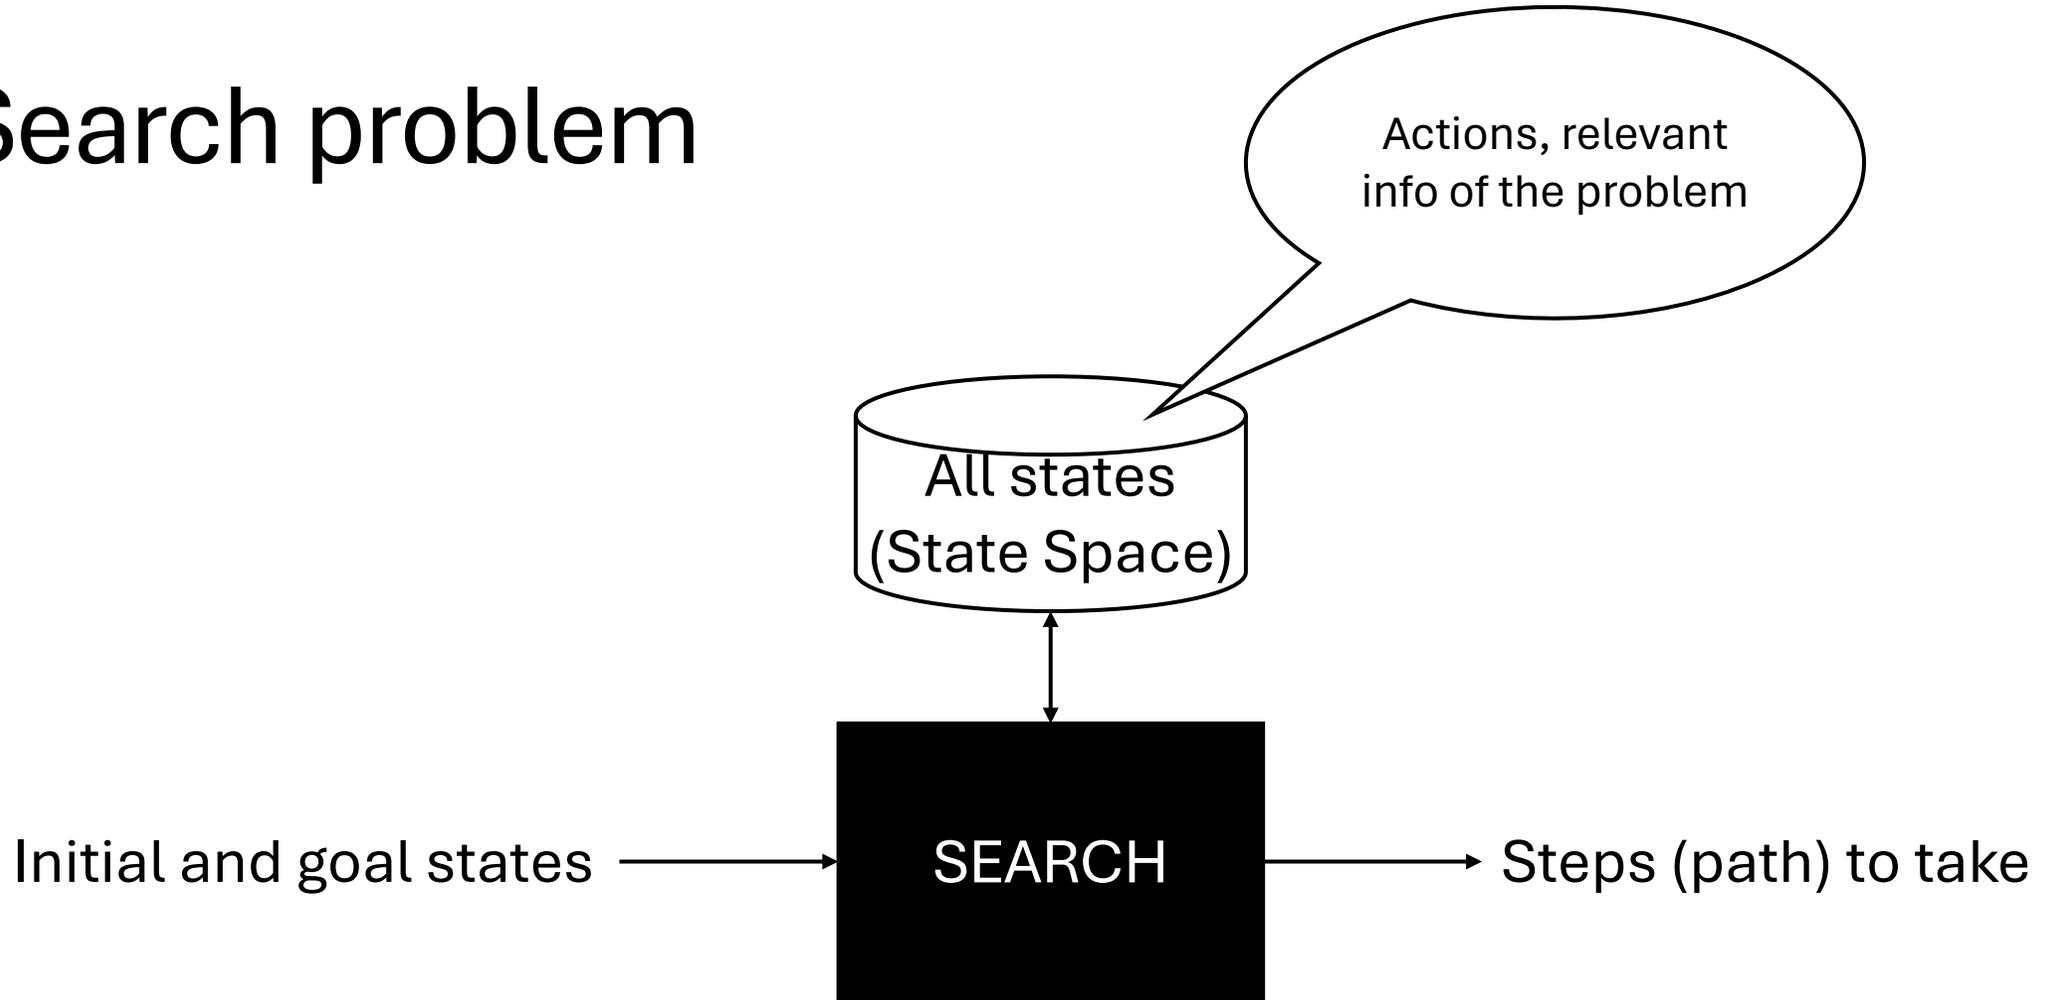
Goal state

Finding answer (to reach a goal) by trials and errors; **exploring alternatives**

# Real-world Search problems

- Route finding: airline travel, computer networks, map navigation, tour plan
- Information search (e.g. web search engine)
- Game playing (e.g. boardgames)
- Electronic ICs (VLSI) layout
- Assembly sequencing
- Packaging arrangement
- Parameters (variables) optimization
- ...

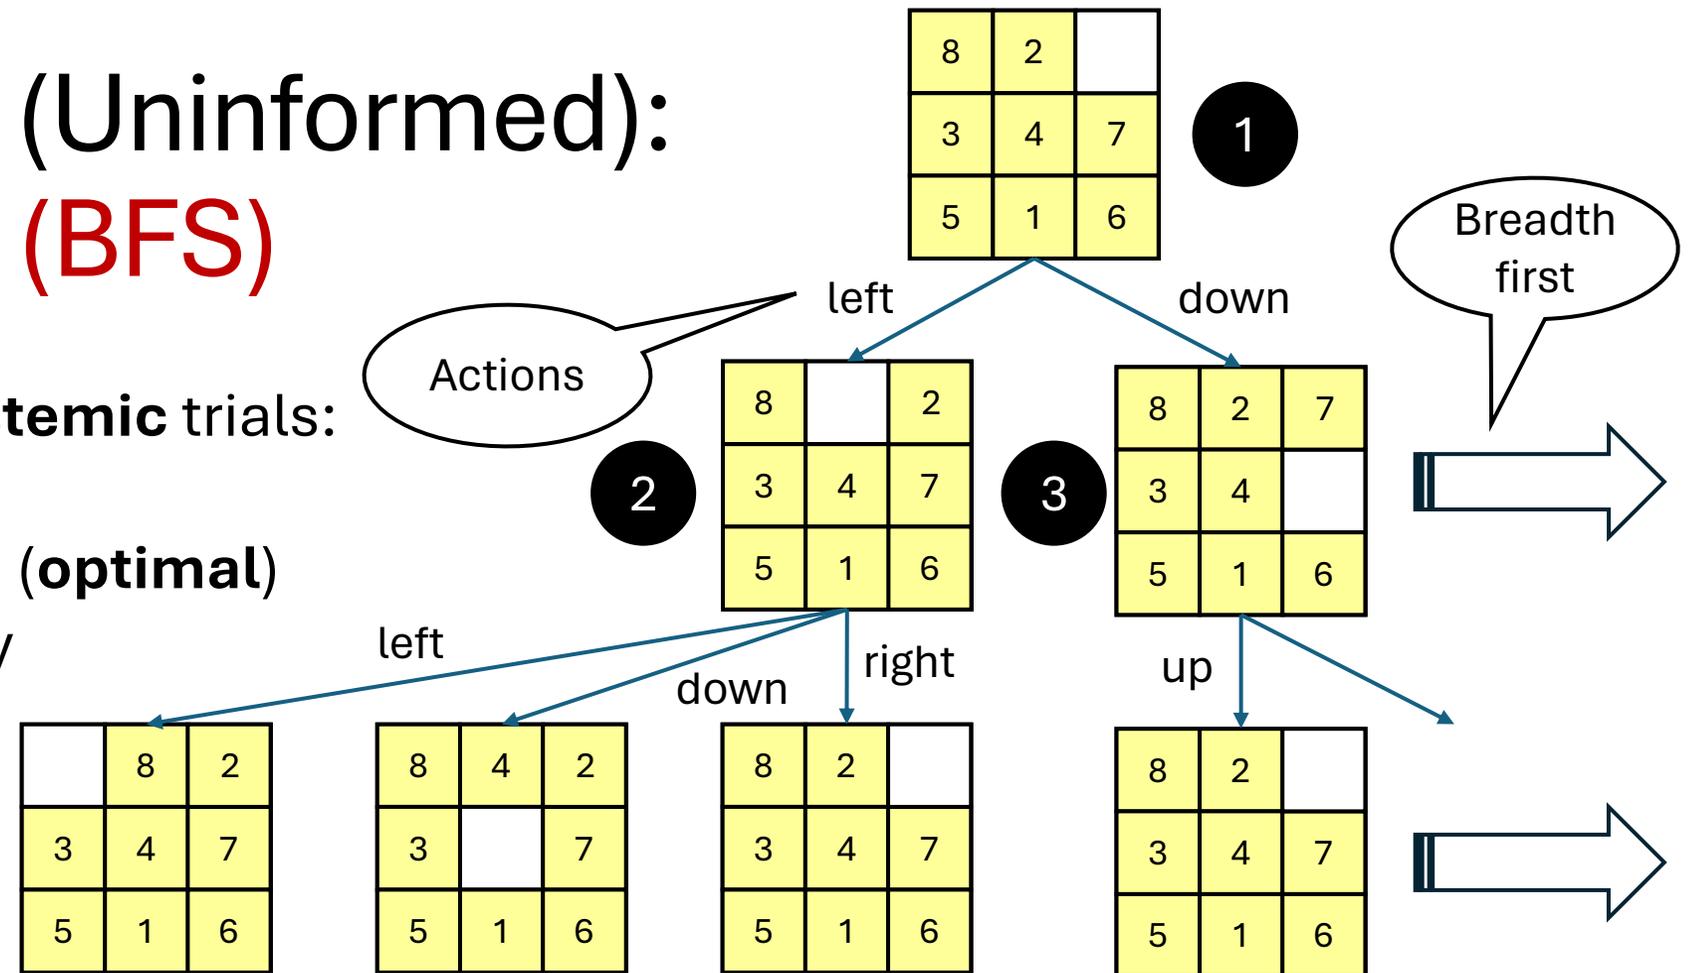
# Search problem



# Blind Search (Uninformed): Breadth-first (BFS)

Finding answer by **systemic** trials:  
try “side” way first

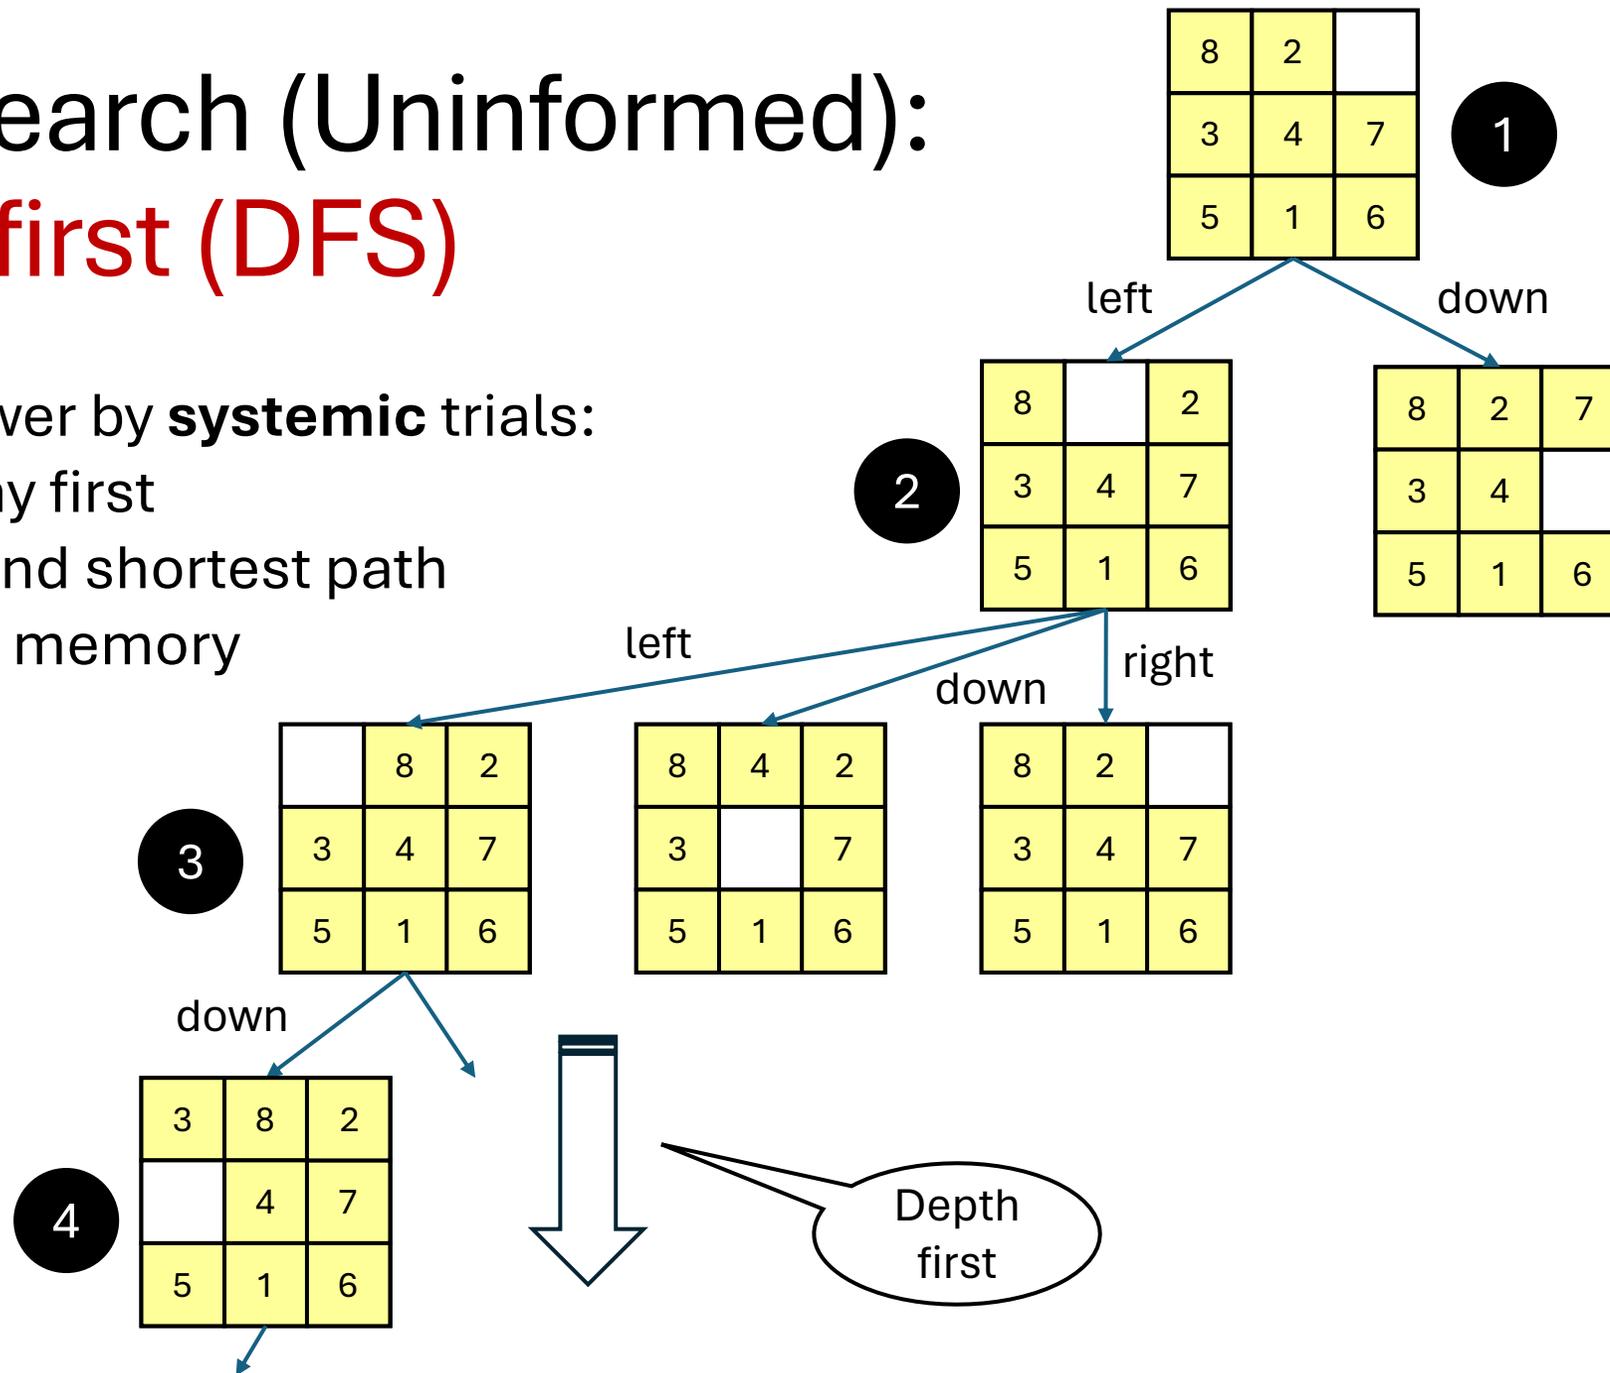
- Finds shortest path (**optimal**)
- Uses more memory



# Blind Search (Uninformed): Depth-first (DFS)

Finding answer by **systemic** trials:  
try “side” way first

- May not find shortest path
- Uses less memory

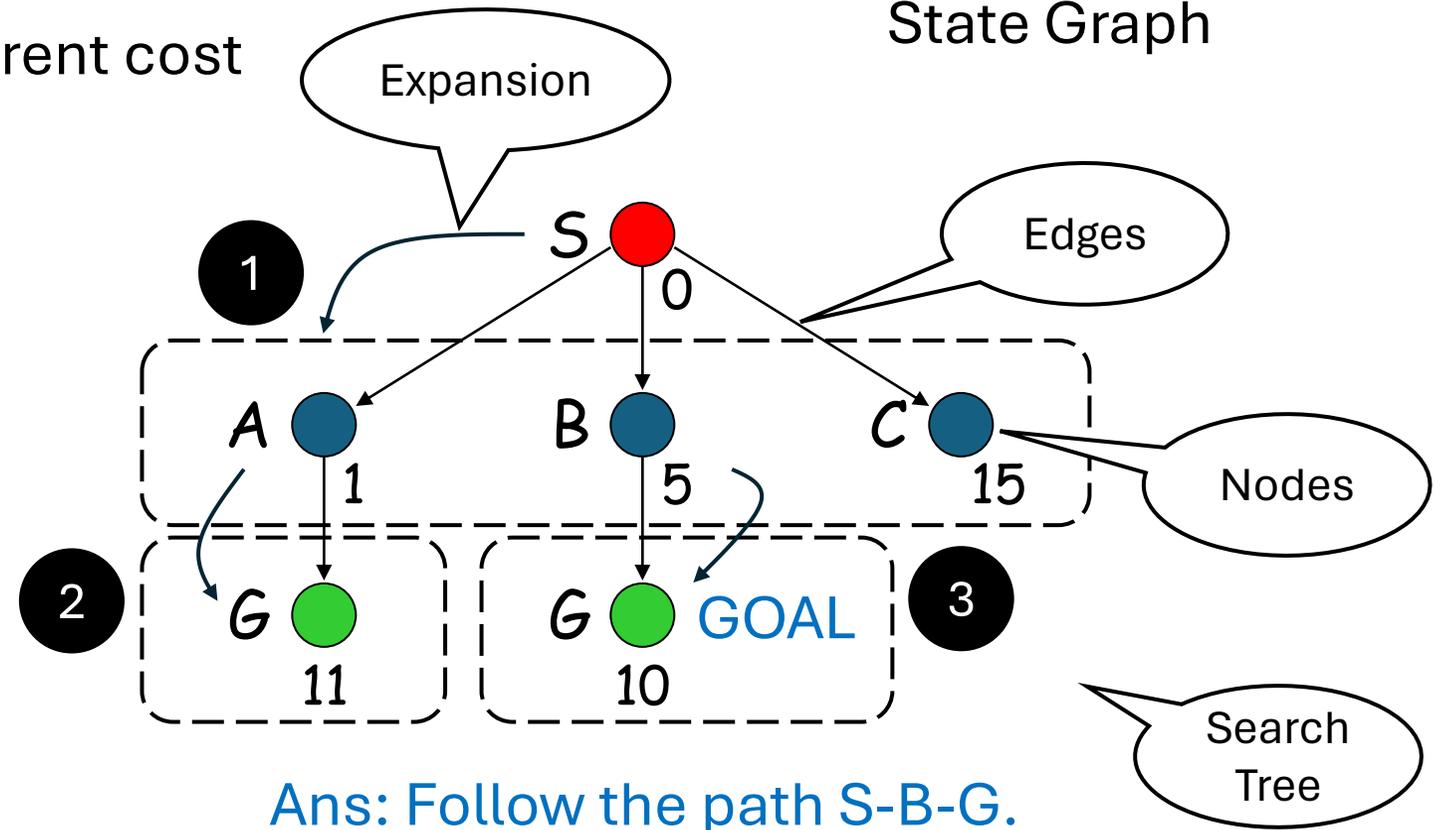
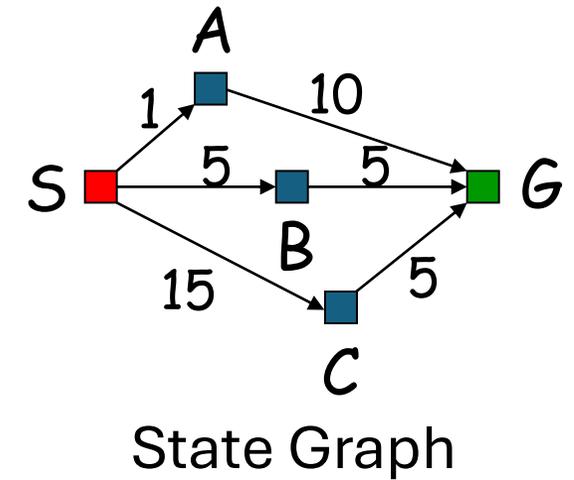


# Blind Search (Uninformed): Uniform-cost (UCS)

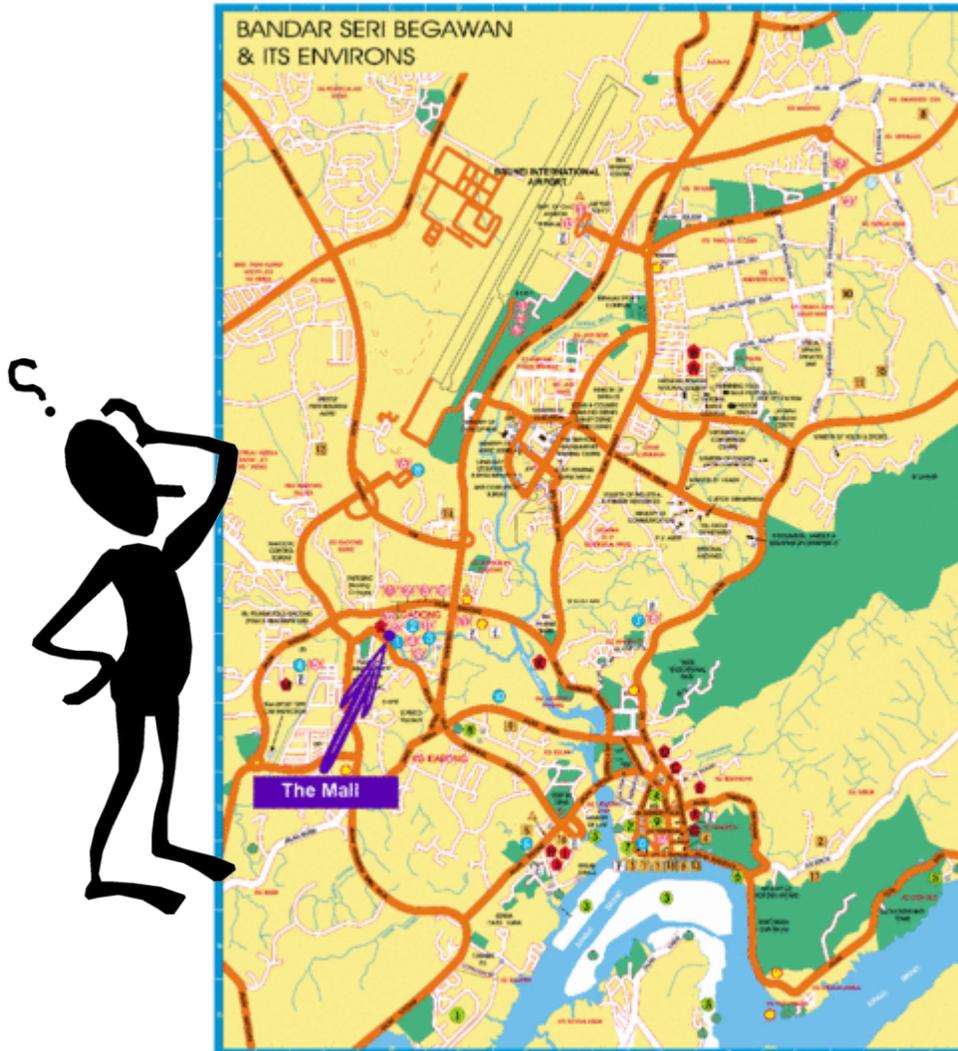
Finding answer by **systemic** trials:  
try “cheapest” (so far) way first

- When each action has different cost
- Finds shortest path

Test goal before **expansion**,  
stop if goal reached  
(likewise, in BFS and DFS)



# Blind vs informed Search



Blind



Informed

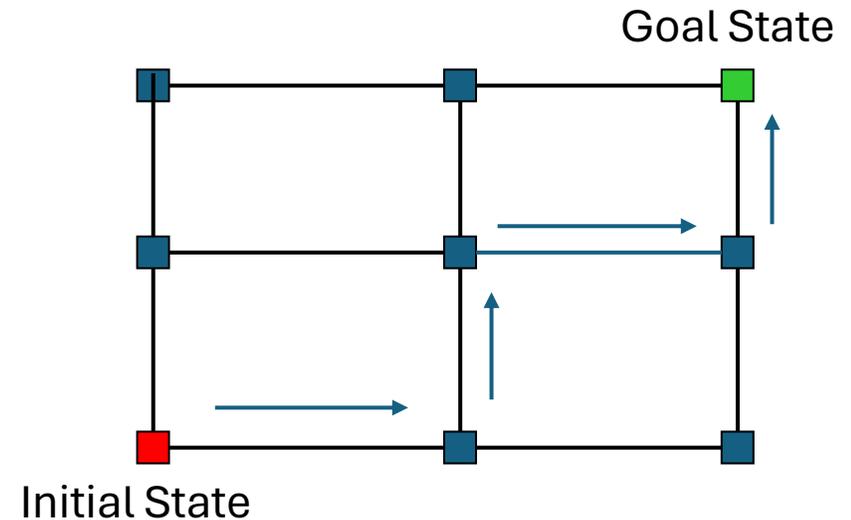
# Heuristic Search (Informed):

## Greedy best-first

Finding answer by **smart** trials: try “closest or best” (estimate) way first

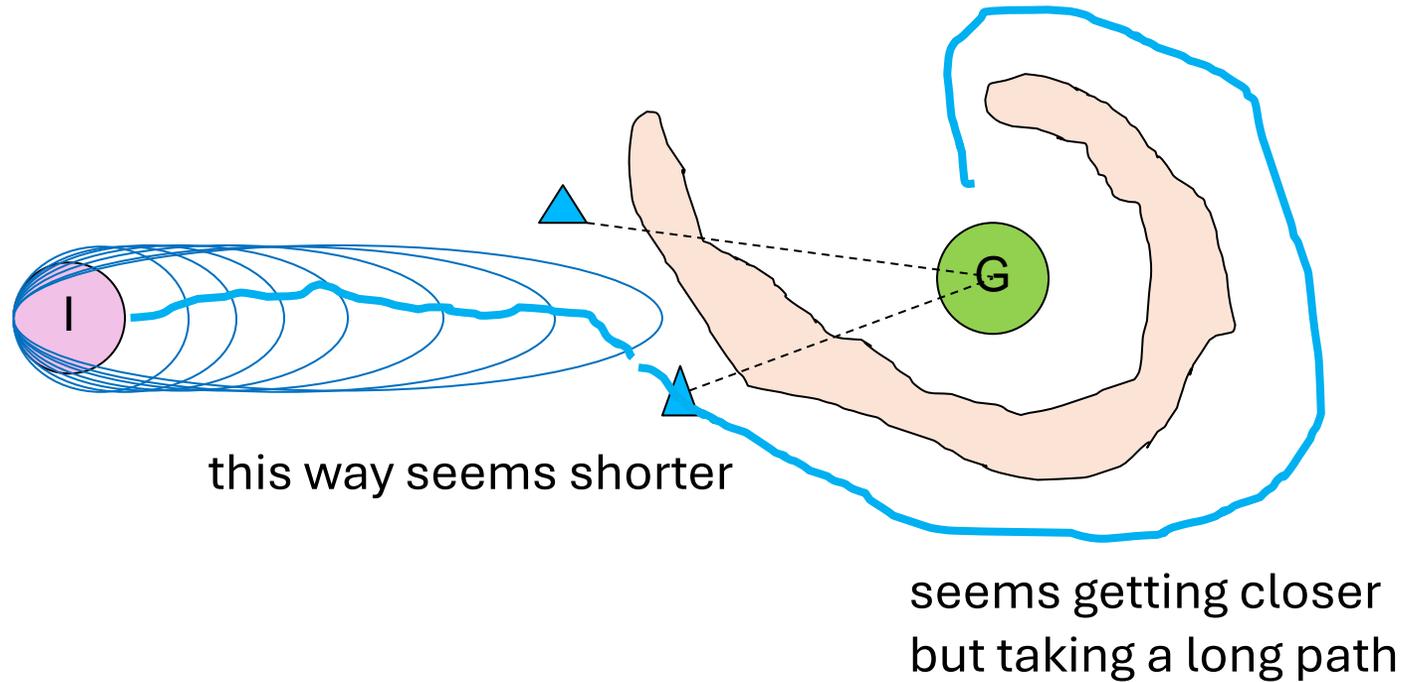
- Uses a **heuristic**  $h(n)$  to estimate “distance” to goal
- Greedy, may reach solution fast
- Performs badly with poor estimates

Heuristic  $h(n)$  = Straight line distance



# Greedy best-first not optimal

**Heuristic:** straight line distance (ignore barriers, actual route/path)

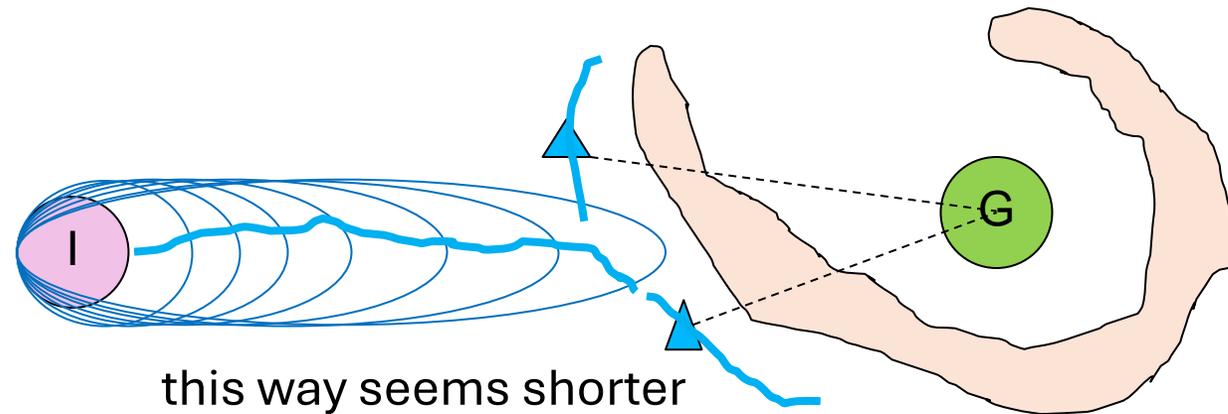


# Heuristic Search (Informed):

## A-star ( $A^*$ )

Finding answer by **smart** trials: try “closest ” (estimate) and “cheapest” (so far) way first

- Uses a heuristic  $h(n)$  to estimate “distance” to goal, and the path cost  $g(n)$  travelled so far
- Optimal, most popular search

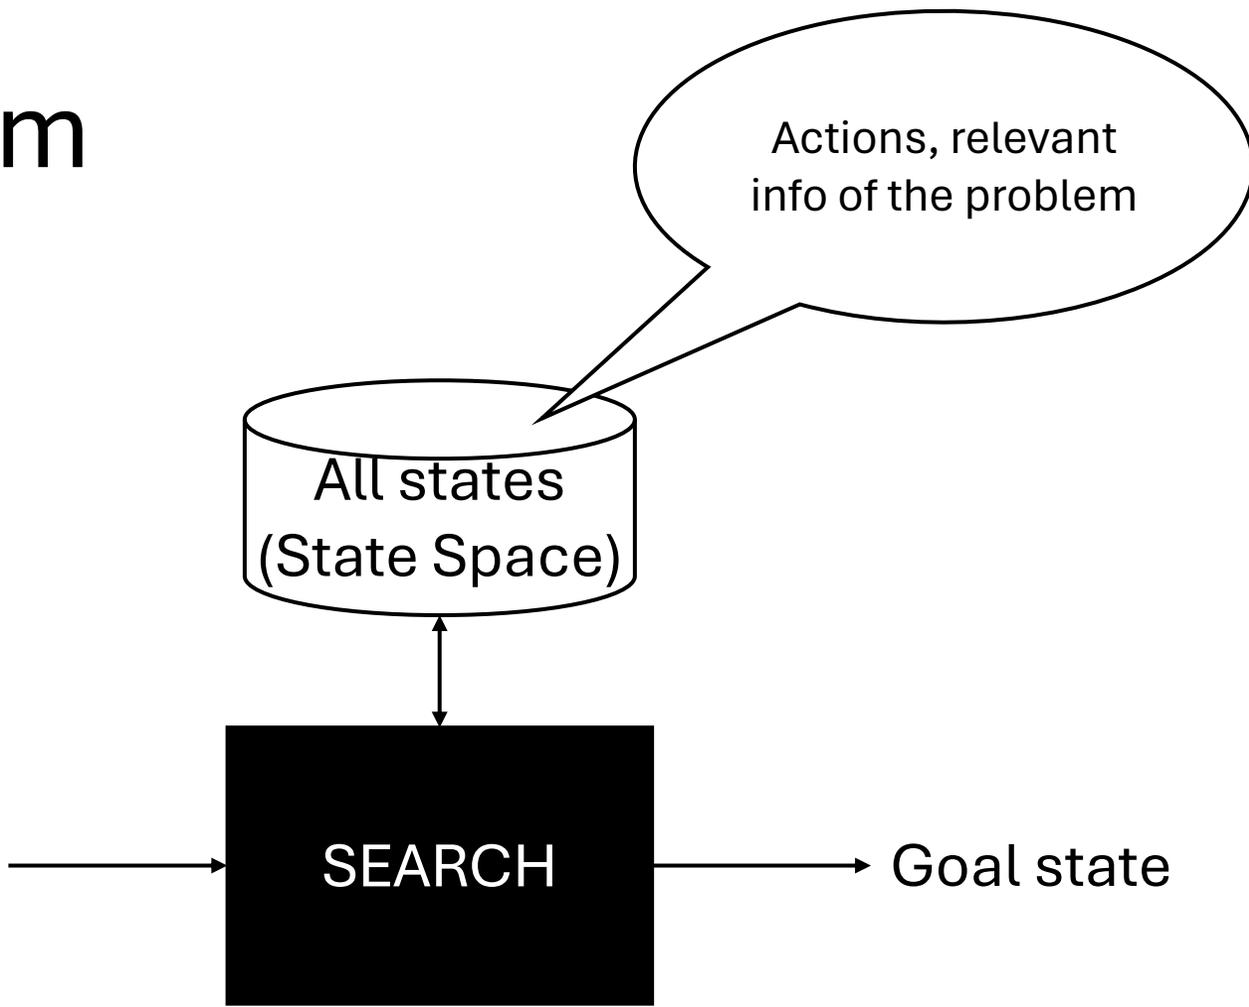


Evaluation of state  $n$ ,  $f(n) = g(n) + h(n)$

seems to be taking long route,  
try alternative

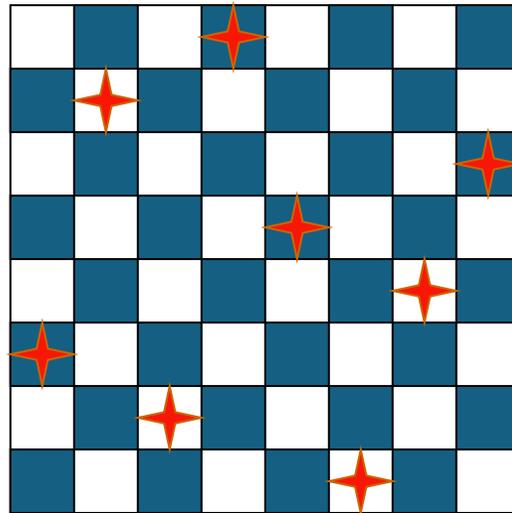
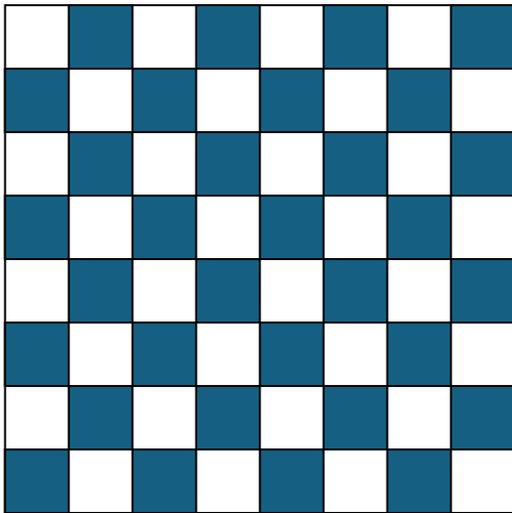
# Search problem

Current state,  
goal state requirements

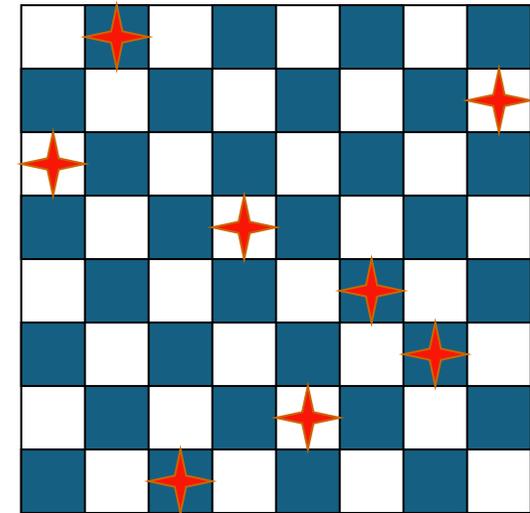


# 8-Queens problem

Problem: Place 8 queens on the chess board without any conflicts



*A solution*



*Not a solution*

# Other search algorithms

- Improvement over basic search algorithms
- Local search
- Beam search
- Adversarial search (game)
- ...

# Knowledge-base systems

When we have knowledge, we reason (think) with our knowledge

# Rules are knowledge

- IF the 'fuel tank' is empty THEN the car is dead
- IF the car is dead AND the 'fuel tank' is empty
- THEN the action is 'refuel the car'
  
- IF the spill is liquid AND the 'spill PH' < 6 AND the 'spill smell' is vinegar
- THEN the 'spill material' is 'acetic acid'

Rules expresses our logics

# Rules are knowledge

- IF the 'fuel tank' is empty THEN the car is dead

- IF the car is dead AND the 'fuel tank' is empty

- THEN the action is 'refuel the car'

**Premise**  
(conditions)

**Consequence**  
(conclusions)

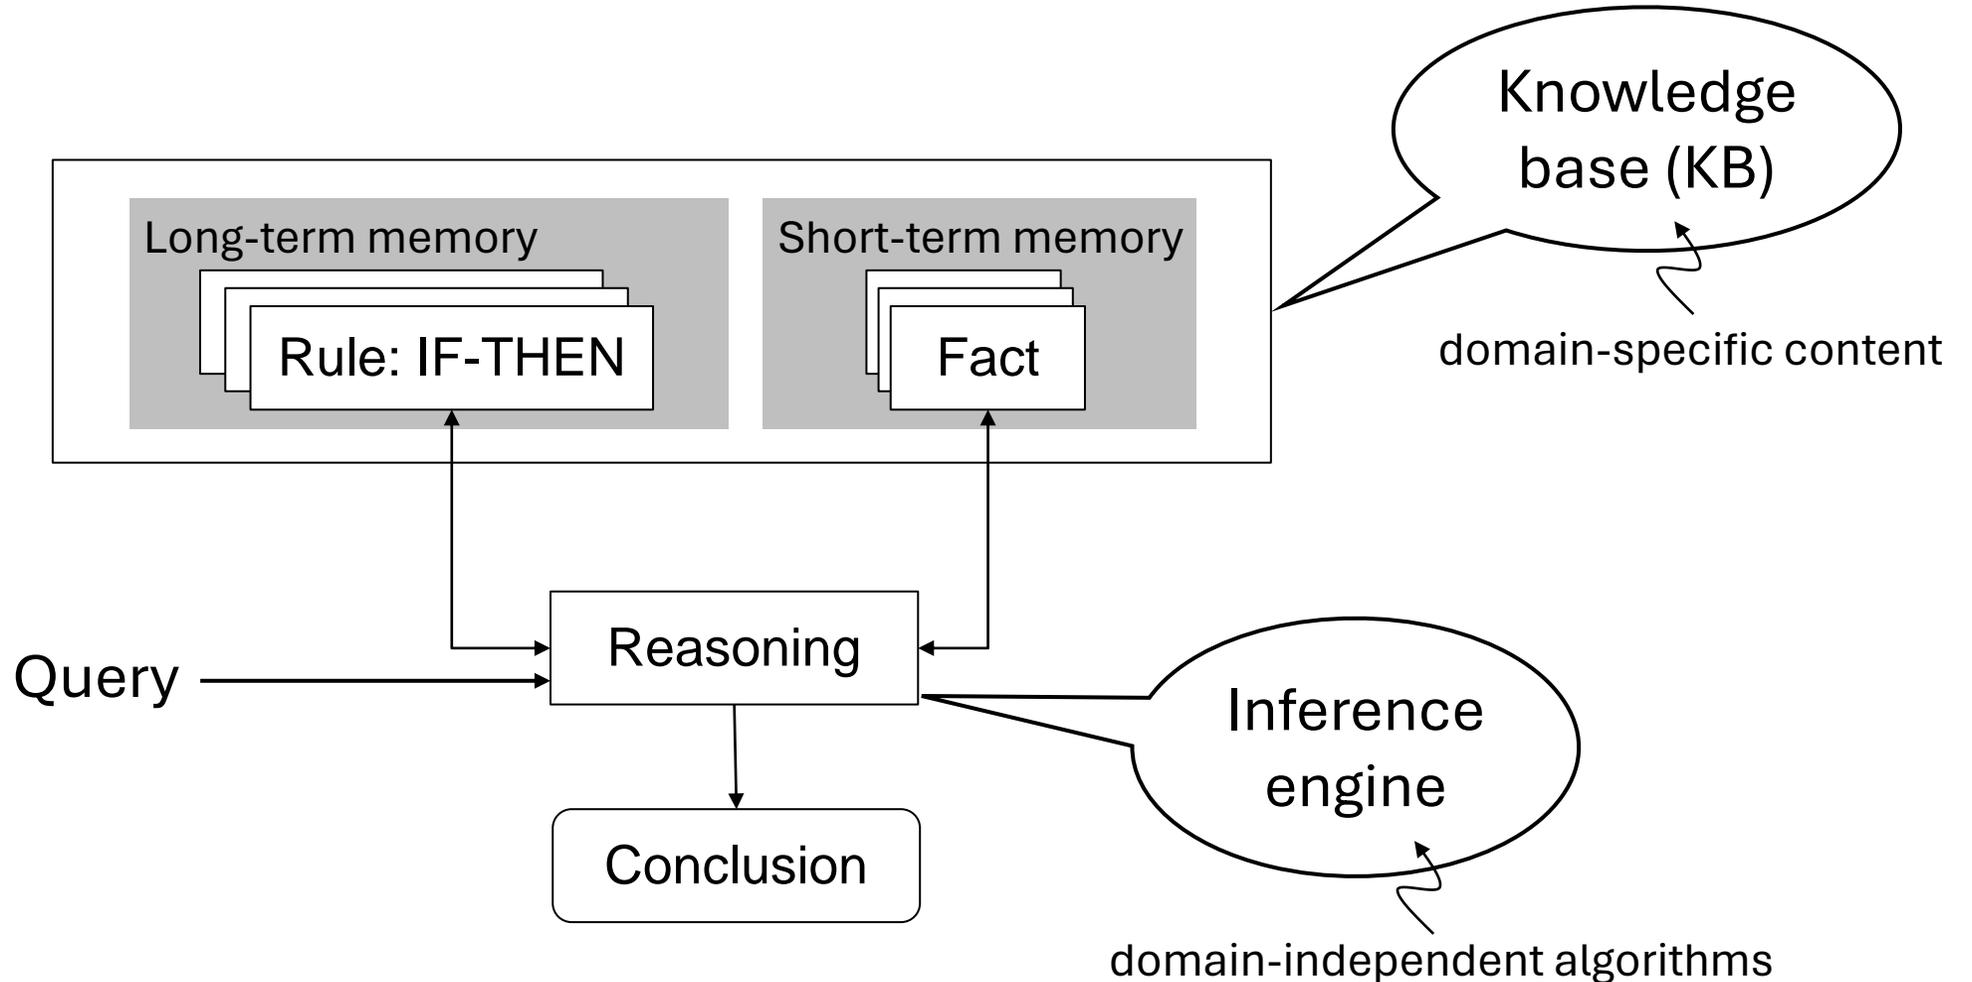
- IF the spill is liquid AND the 'spill PH' < 6 AND the 'spill smell' is vinegar

- THEN the 'spill material' is 'acetic acid'

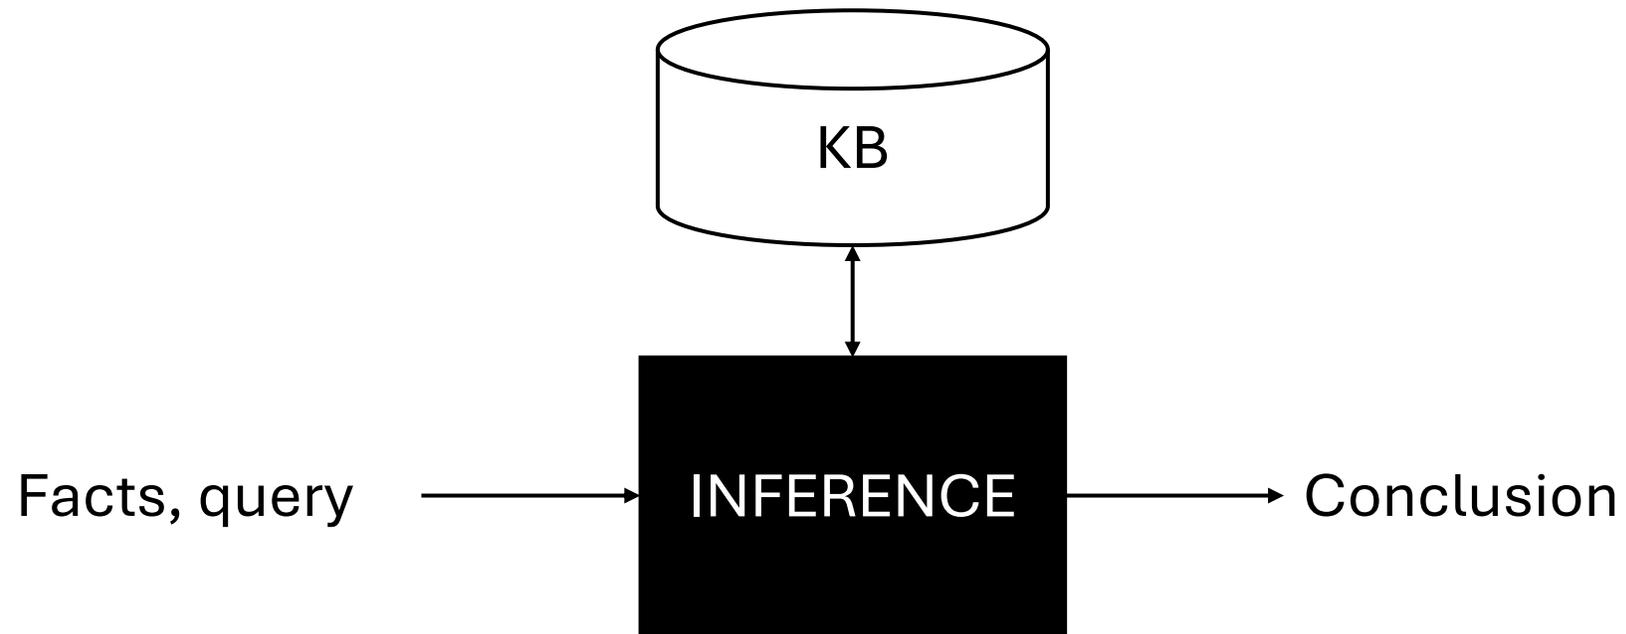
Rules expresses our logics

# Rule-based expert system

Expert system is one type of KB systems, targeting specific domain



# Knowledge-base system



# Inference approaches

- Proof by enumeration (check all rules with all facts)
- Inference rules (algebraic manipulation)
- Forward reasoning
- Backward reasoning
- Proof by resolution

# Forward reasoning

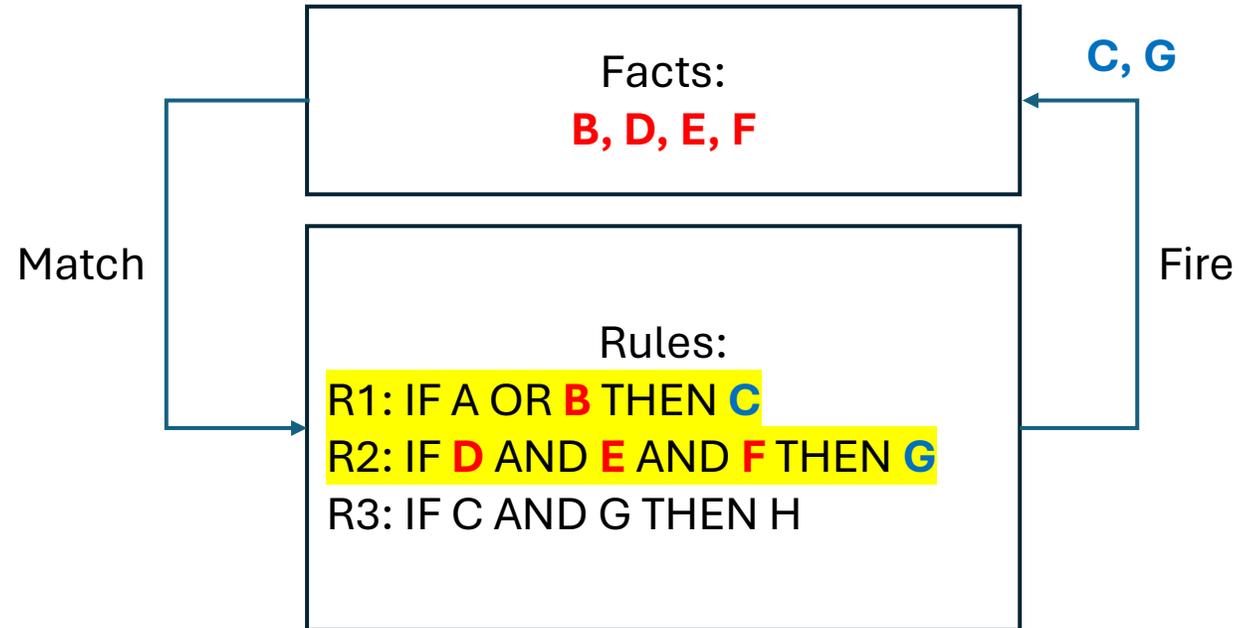
- Given the following rules:
  - R1: IF A OR B THEN C
  - R2: IF D AND E AND F THEN G
  - R3: IF C AND G THEN H
- And the following facts:
  - B, D, E, F
- Goal: Prove H

Facts:  
B, D, E, F

Rules:  
R1: IF A OR B THEN C  
R2: IF D AND E AND F THEN G  
R3: IF C AND G THEN H

# Forward reasoning

- Given the following rules:
  - R1: IF A OR B THEN C
  - R2: IF D AND E AND F THEN G
  - R3: IF C AND G THEN H
- And the following facts:
  - B, D, E, F
- Goal: Prove H



## Steps (match-fire cycles):

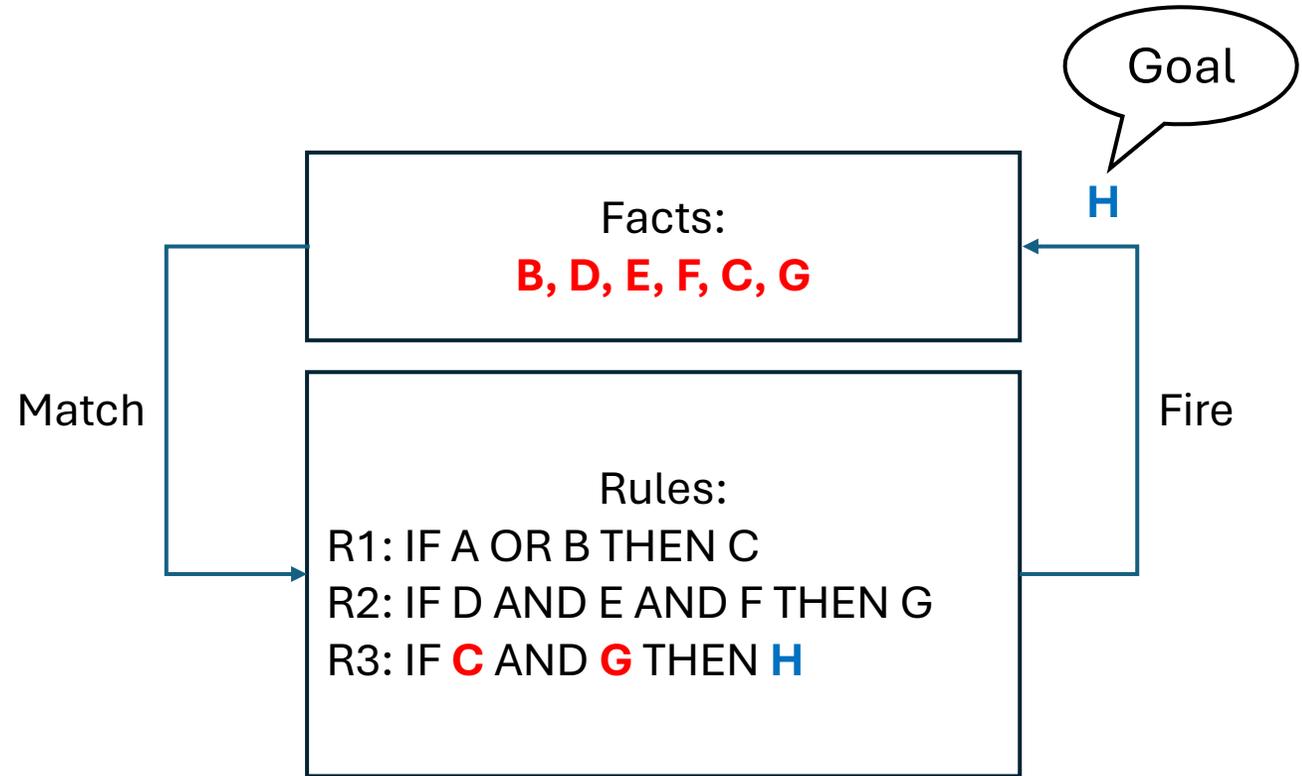
- Check each rule with facts
- If facts **match** premise, **fires** the rule, generate **new** fact(s) from the consequence
- Repeat the match-fire cycle with the updated facts
- Stop when goal achieved

# Forward reasoning

- Given the following rules:
  - R1: IF A OR B THEN C
  - R2: IF D AND E AND F THEN G
  - R3: IF C AND G THEN H
- And the following facts:
  - B, D, E, F
- Goal: Prove H

## Steps (match-fire cycles):

- Check each rule with facts
- If facts **match** premise, **fires** the rule, generate **new** fact(s) from the consequence
- Repeat the match-fire cycle with the updated facts
- Stop when goal achieved



# Knowledge representation

- We can't program the computer in our languages (yet)
- **Formal logic** is designed to representation knowledge (rules and facts)
- Propositional logic
- First order logic
- ...

$$C \wedge G \Rightarrow H$$

IF C AND G THEN H

$$\exists y \forall x (Loves(x, y))$$

Says there is someone who is loved by everyone in the universe.

# Exploring Wumpus world with KB

## Knowledge base:

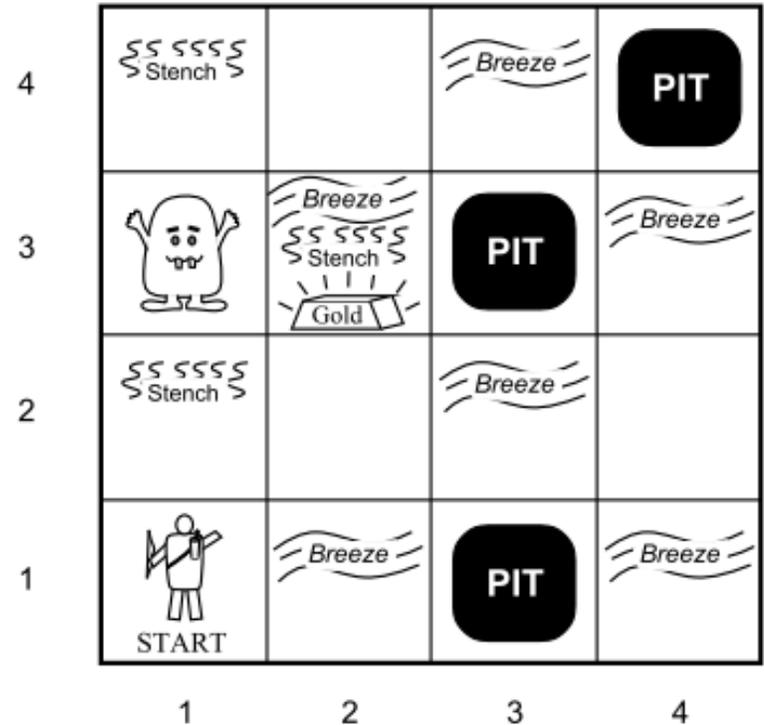
- Squares adjacent to Wumpus are smelly
- Squares adjacent to pit are breezy
- Glitter iff gold is in the same square
- Shooting kills Wumpus if you are facing it
- Shooting uses up the only arrow
- Grabbing picks up gold if in same square
- Releasing drops the gold in same square

## Actions:

- Left turn, Right turn, Forward, Grab, Release, Shoot

## Senses:

- Breeze, Glitter, Smell



Wumpus and Pits don't move  
Agent die if meet Wumpus or Pit  
Agent can sense its current cell

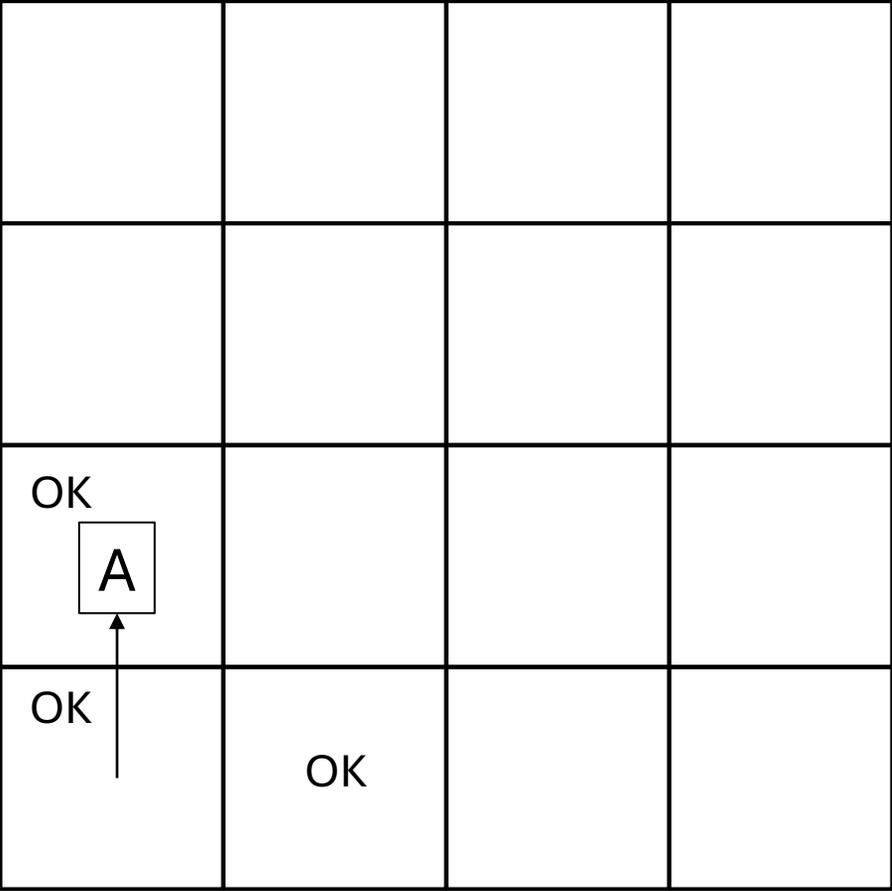
A			

Agent start at (1,1)  
Senses nothing

OK			
OK A	OK		

Agent start at (1,1)  
Senses nothing

Conclude (1,1), (1,2),  
(2,1) are OK (safe)



Agent moves to (1,2)

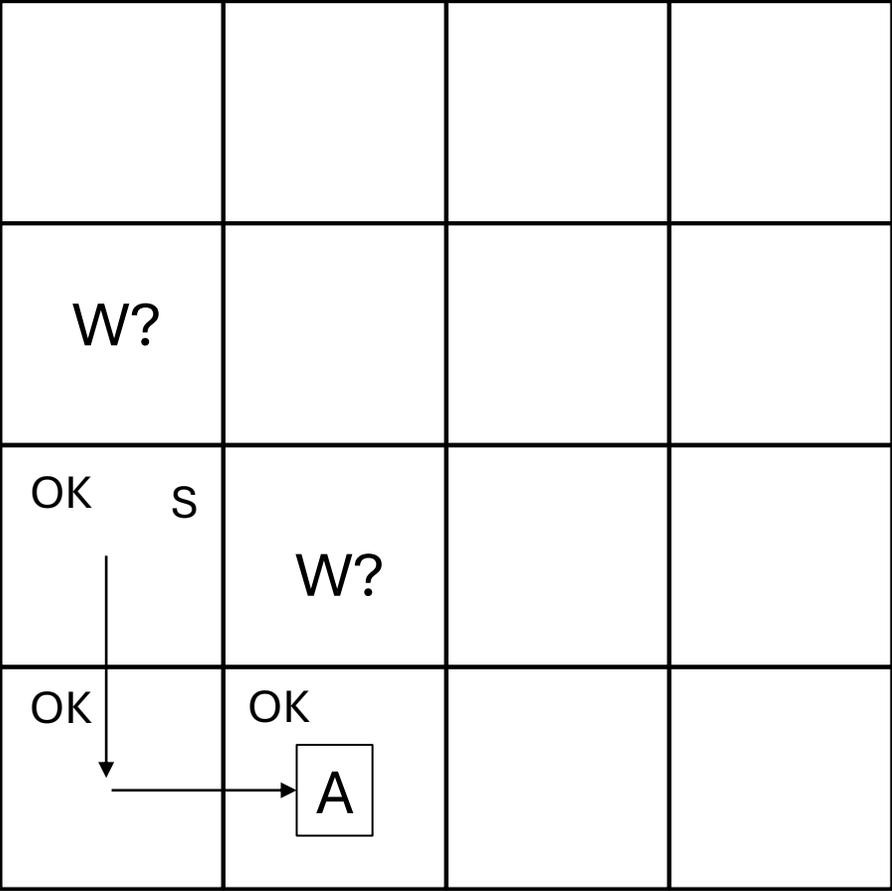
OK S A			
OK	OK		

Agent moves to (1,2)  
Senses stench

W?			
OK A S	W?		
OK	OK		

Agent moves to (1,2)  
Senses stench

Conclude there is  
Wumpus either in  
(1,3) or (2,2) or both.



Agent moves back  
and heads to (2,1)

W?			
OK S	W?		
OK	OK B A		

Agent moves back  
and heads to (2,1)  
Senses breeze

W?			
OK S	W? P?		
OK	OK B A	P?	

Agent moves back  
and heads to (2,1)  
Senses breeze

Conclude that there is  
Pit in either (2,2) or  
(3,1) or both.

W			
OK S	OK		
OK	OK B A	P	

Agent moves back  
and heads to (2,1)  
Senses breeze

Conclude that there is  
Pit in either (2,2) or  
(3,1) or both.

However, if (2,2) has Pit, there  
should be breeze at (1,2), but  
there wasn't.

So, there is Pit in (3,1).

Further, since there is no stench  
in (2,1), there is no Wumpus in  
(2,2), Then the Wumpus must  
be in (1,3).

W			
OK S	OK A		
OK	OK B	P	

Agent moves to (2,2)

W	OK		
OK S	OK A	OK	
OK	OK B	P	

Agent moves to (2,2)  
Senses nothing

Conclude (2,3) and  
(3,2) are safe

W	OK A		
OK S	OK	OK	
OK	OK B	P	

Agent moves to (2,3)

W	OK BSG A		
OK S	OK	OK	
OK	OK B	P	

Agent moves to (2,3)  
Senses breeze, stench  
and glitter

	W?P?		
W	OK BSG G <span style="border: 1px solid black; padding: 2px;">A</span>	W?P?	
OK S	OK	OK	
OK	OK B	P	

Agent moves to (2,3)  
Senses breeze, stench  
and glitter

Conclude that there are  
Wumpus and Pits  
Around (2,3), and  
there is Gold in (2,3)

Gold found

# AI techniques

## Symbolic AI

aka Interpretational AI, Good Old Fashion AI (GOFAI), Traditional AI

- Problem solving by search
- Knowledge-based system

## Non-symbolic AI

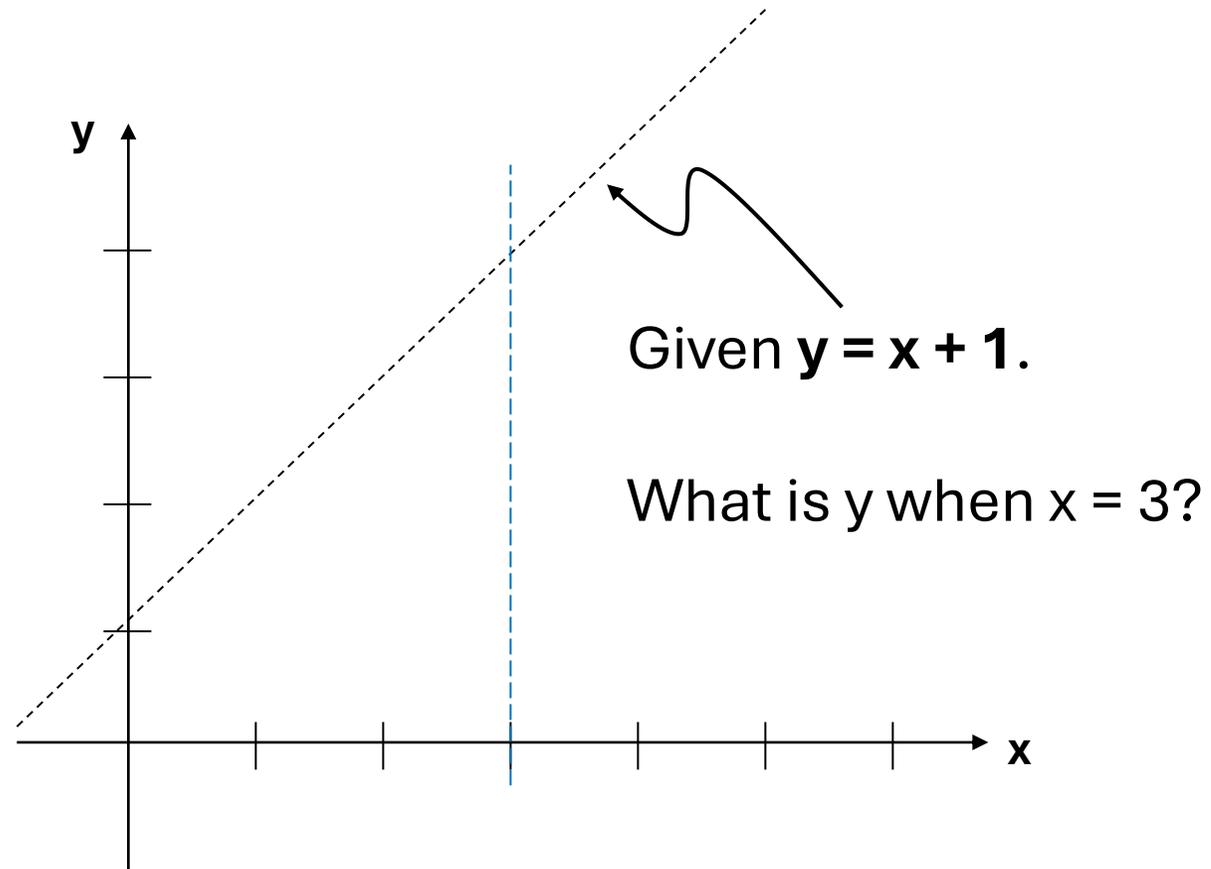
aka Sub-symbolic, Modern AI

- Machine Learning
  - Linear regression
  - Logistic regression
  - Artificial Neural Network
  - K-Means clustering

# Machine Learning

It's impossible to encode all knowledge in the world, can machines learn from observations?

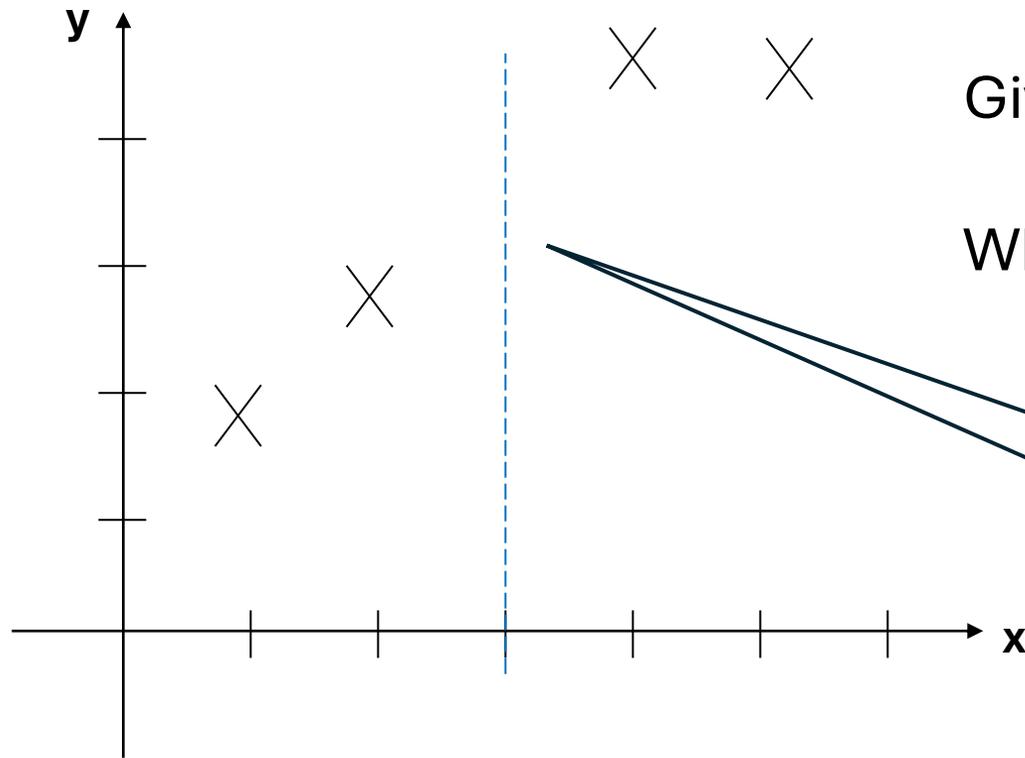
# No learning required



# Traditional computer science



# Learning required

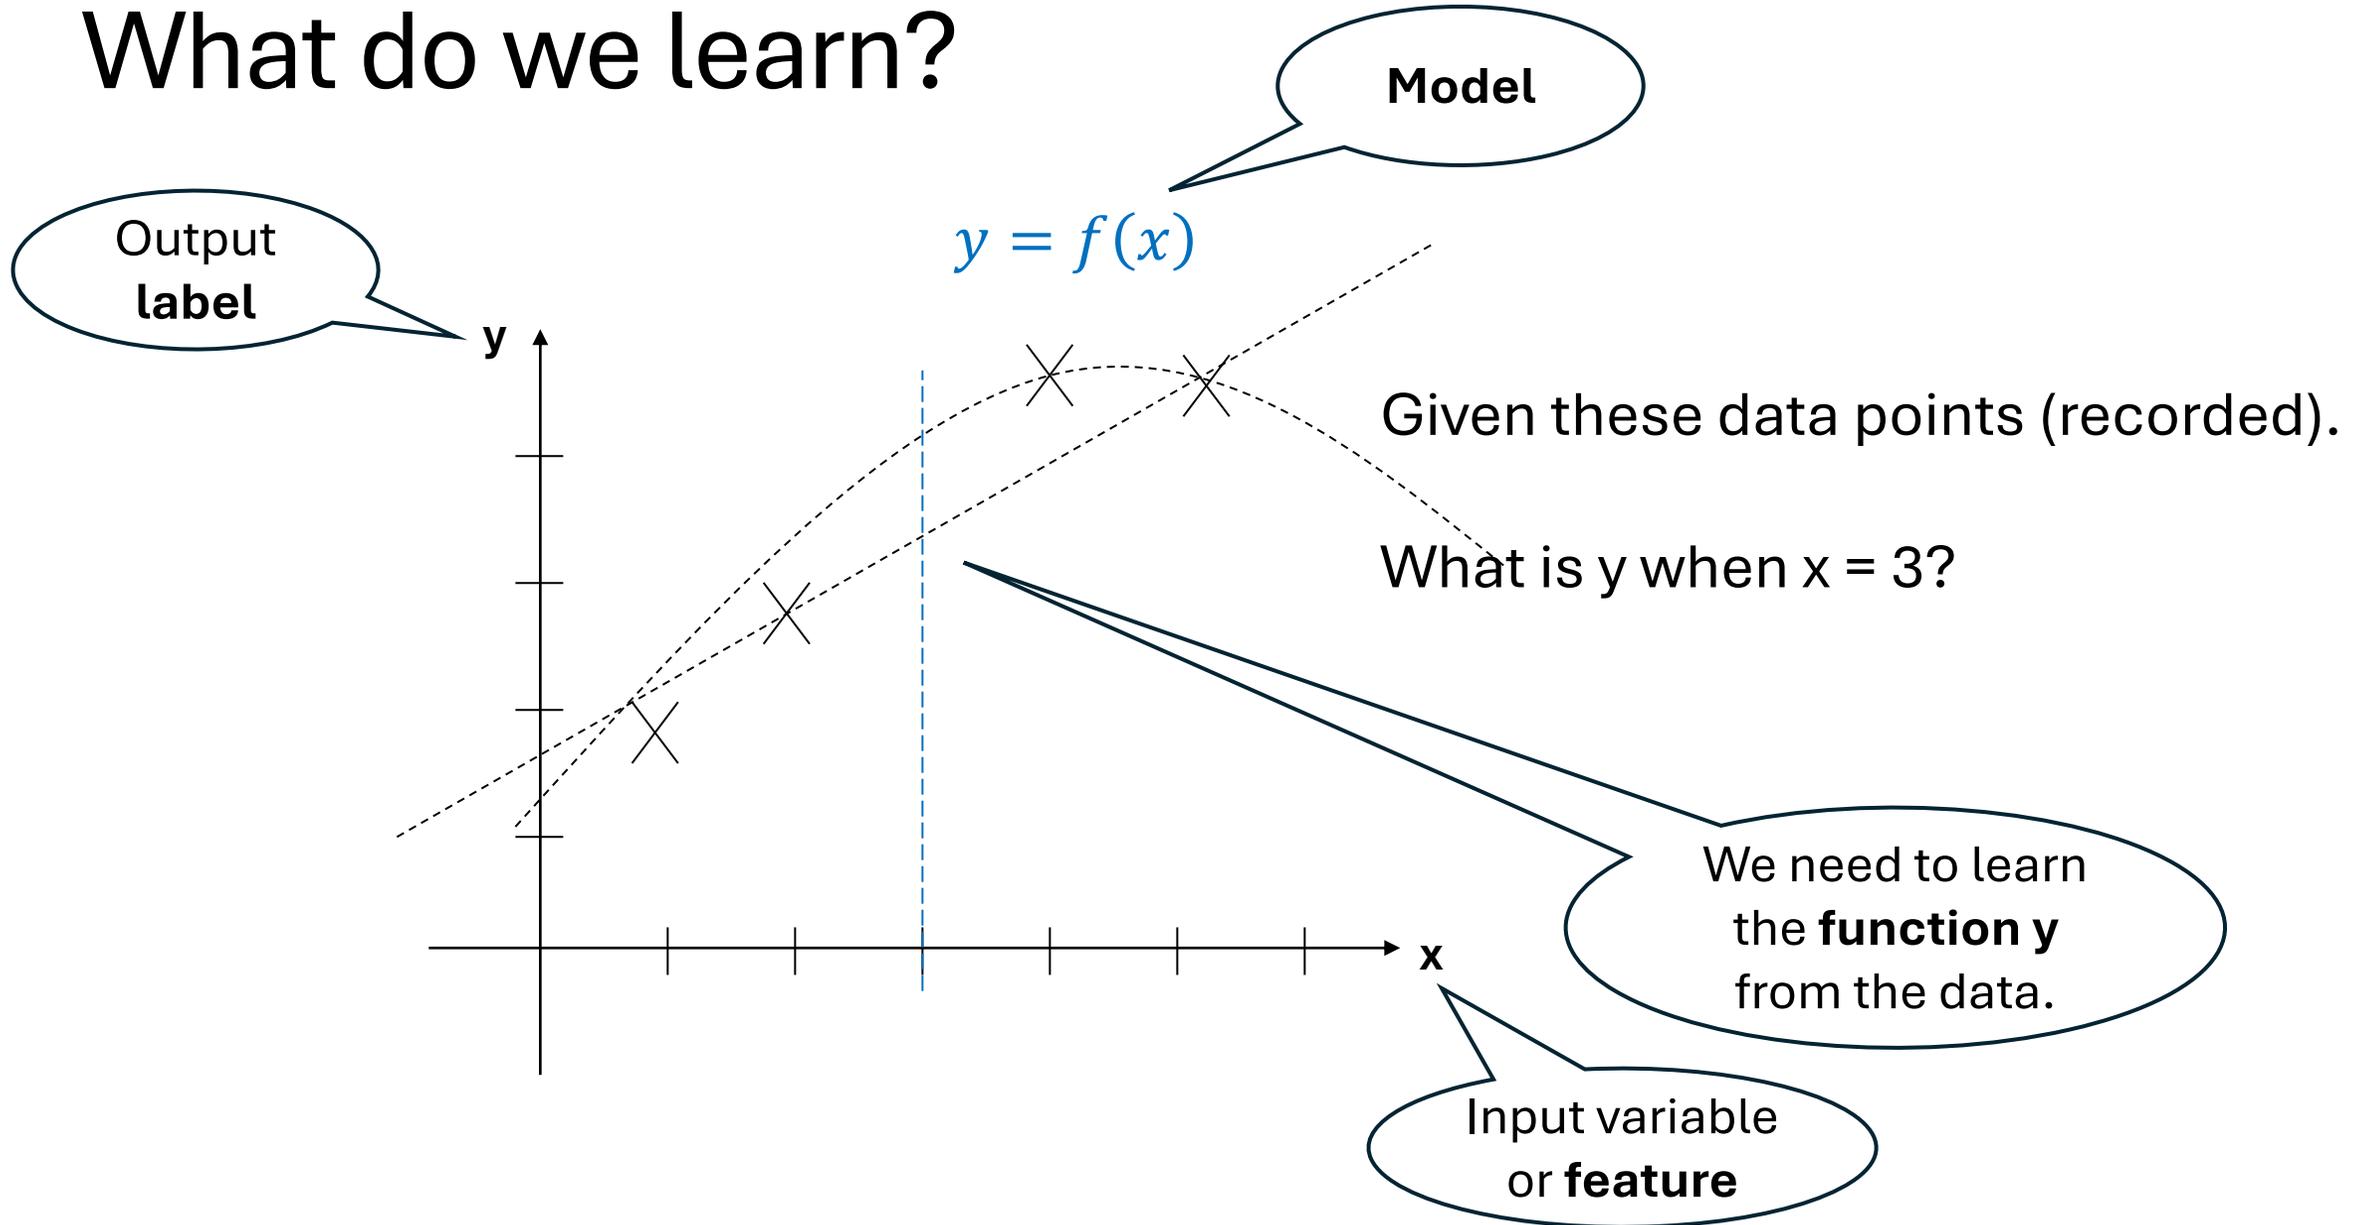


Given these data points (recorded).

What is y when  $x = 3$ ?

We need to figure out from the data.

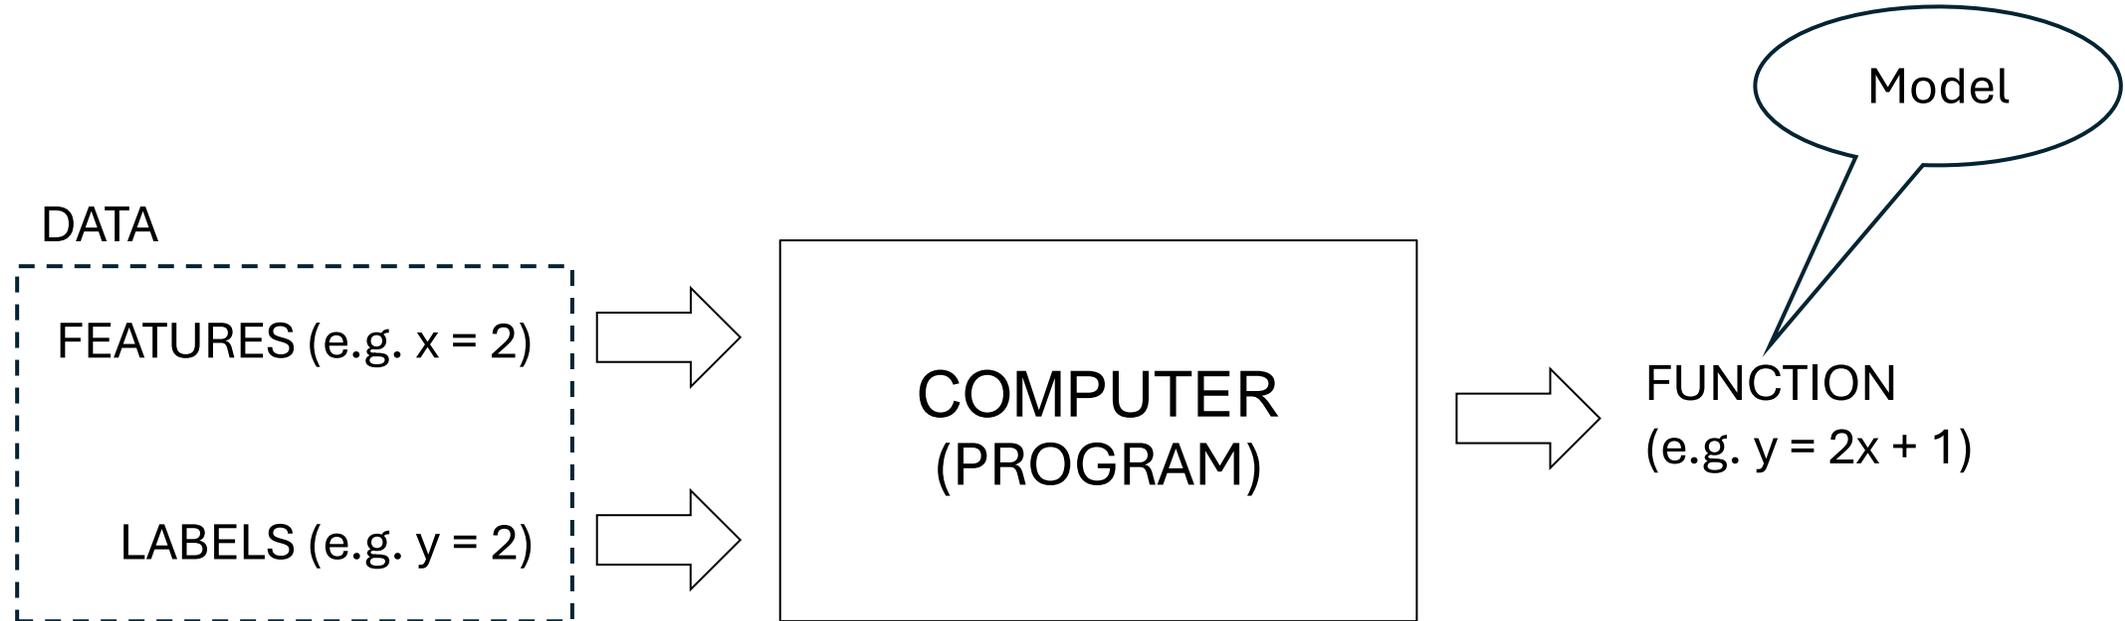
# What do we learn?



# Machine learning



# Machine learning



# Can you figure out?

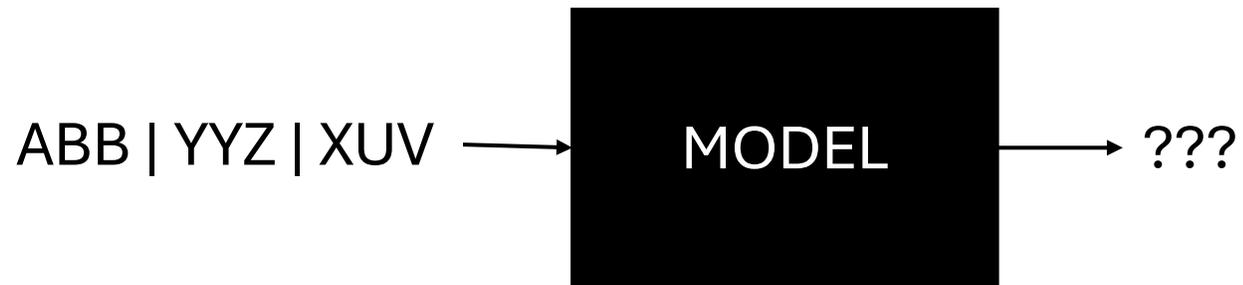


ABB | YYZ | XUV = ???

AAA | BBB | CCC = ABC

AAB | BBC | CCD = ABD

ABB | BCC | CDD = ACD

ABC | DEF | GHI = AEI

# Can you figure out?

And use the model to predict the output for new data

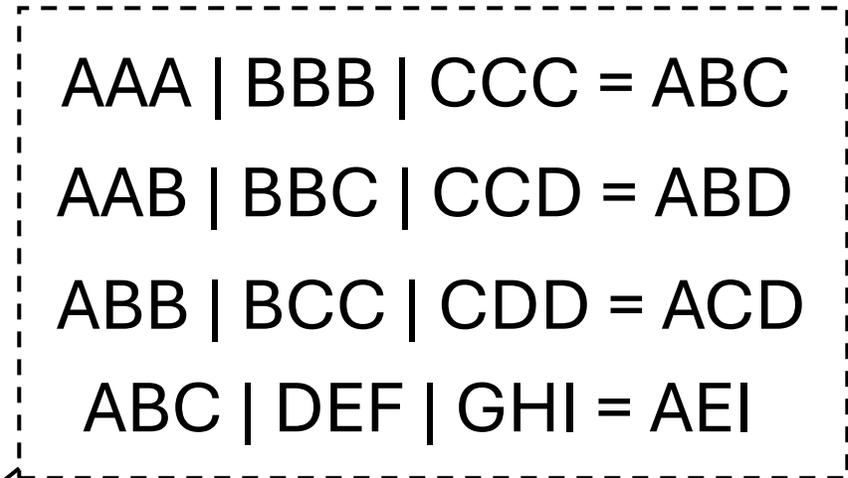
This is **learned** from sample data

ABB | YYZ | XUV = ???

ABB | YYZ | XUV



???



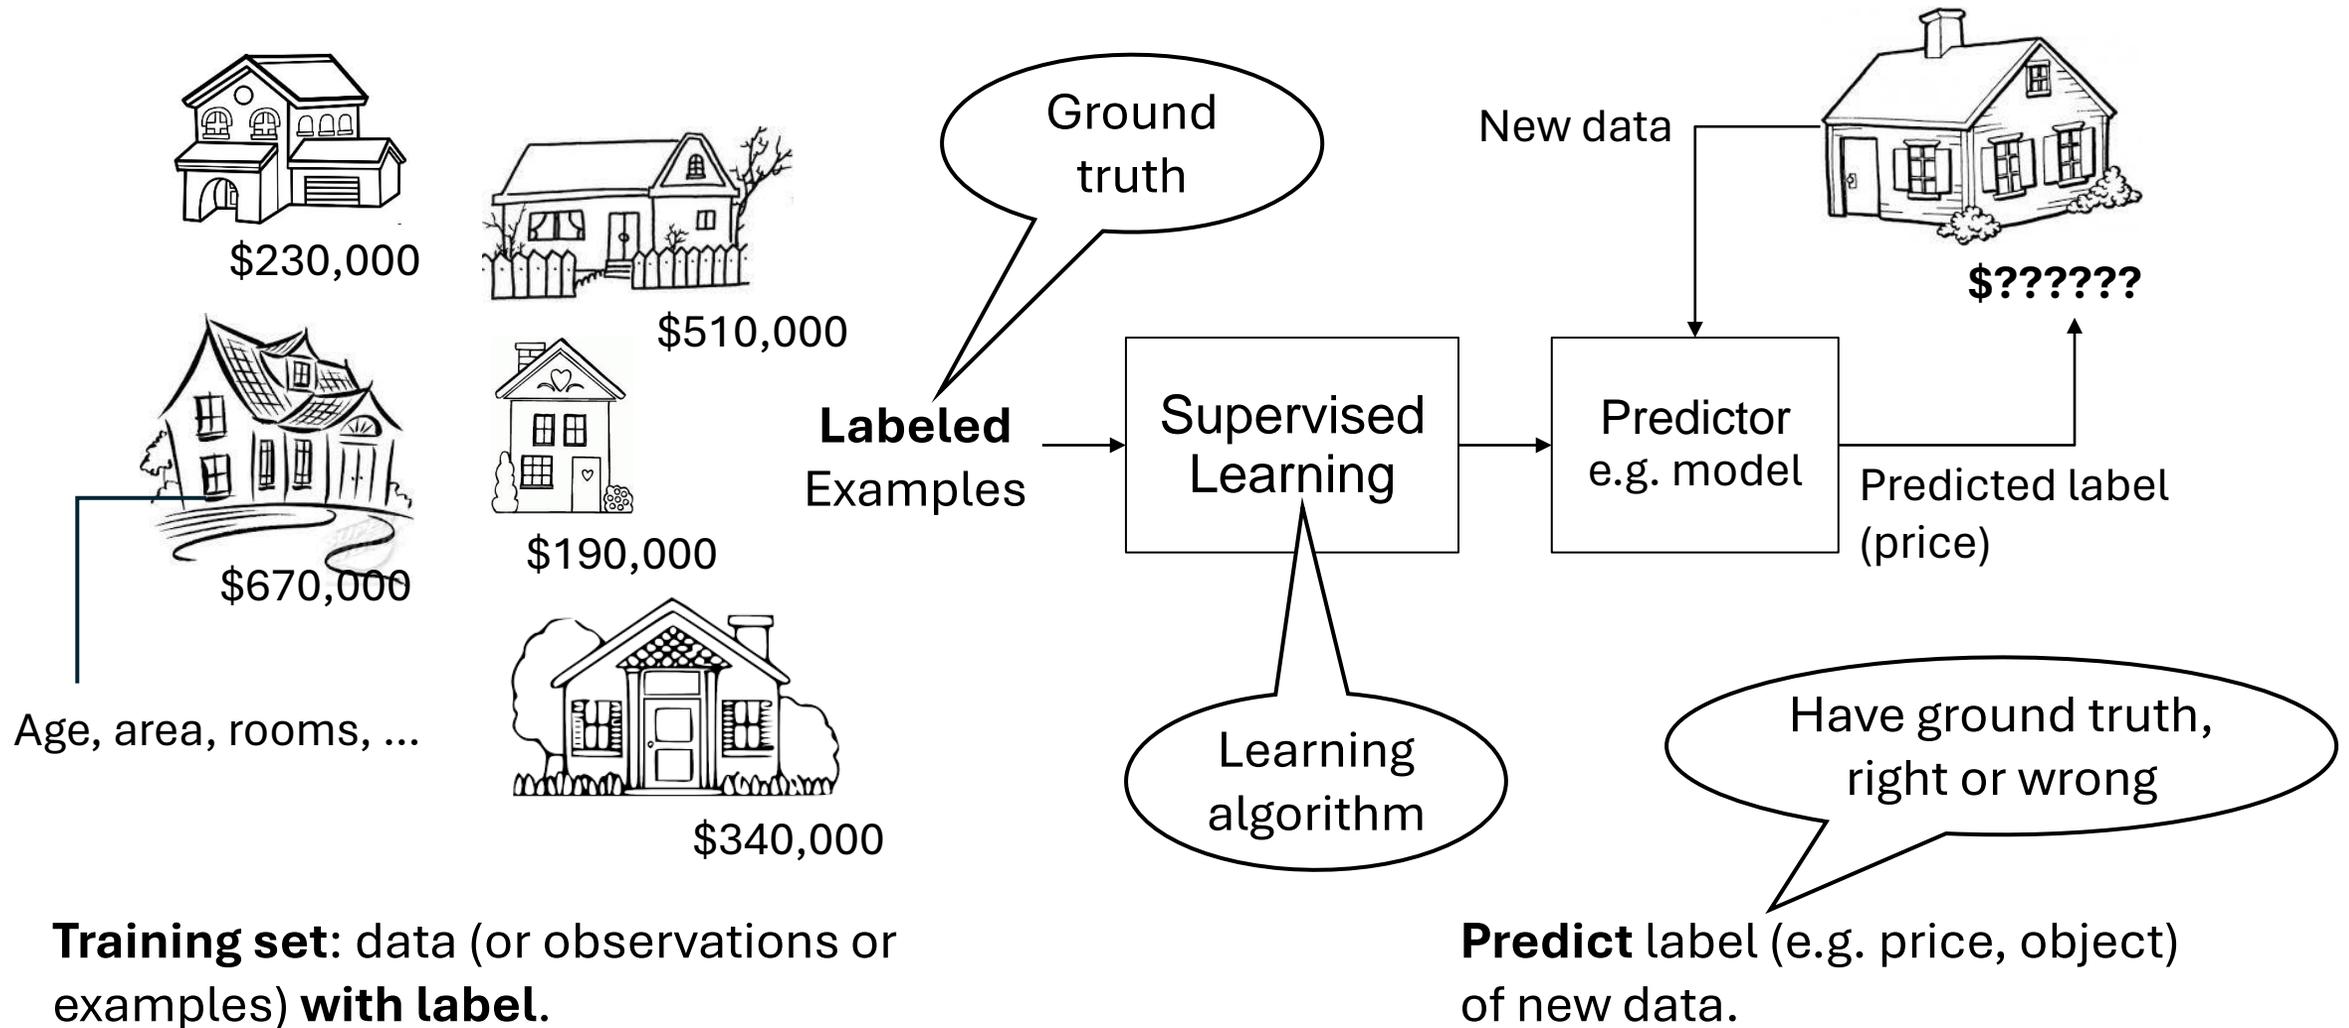
Observations (data)

generalize

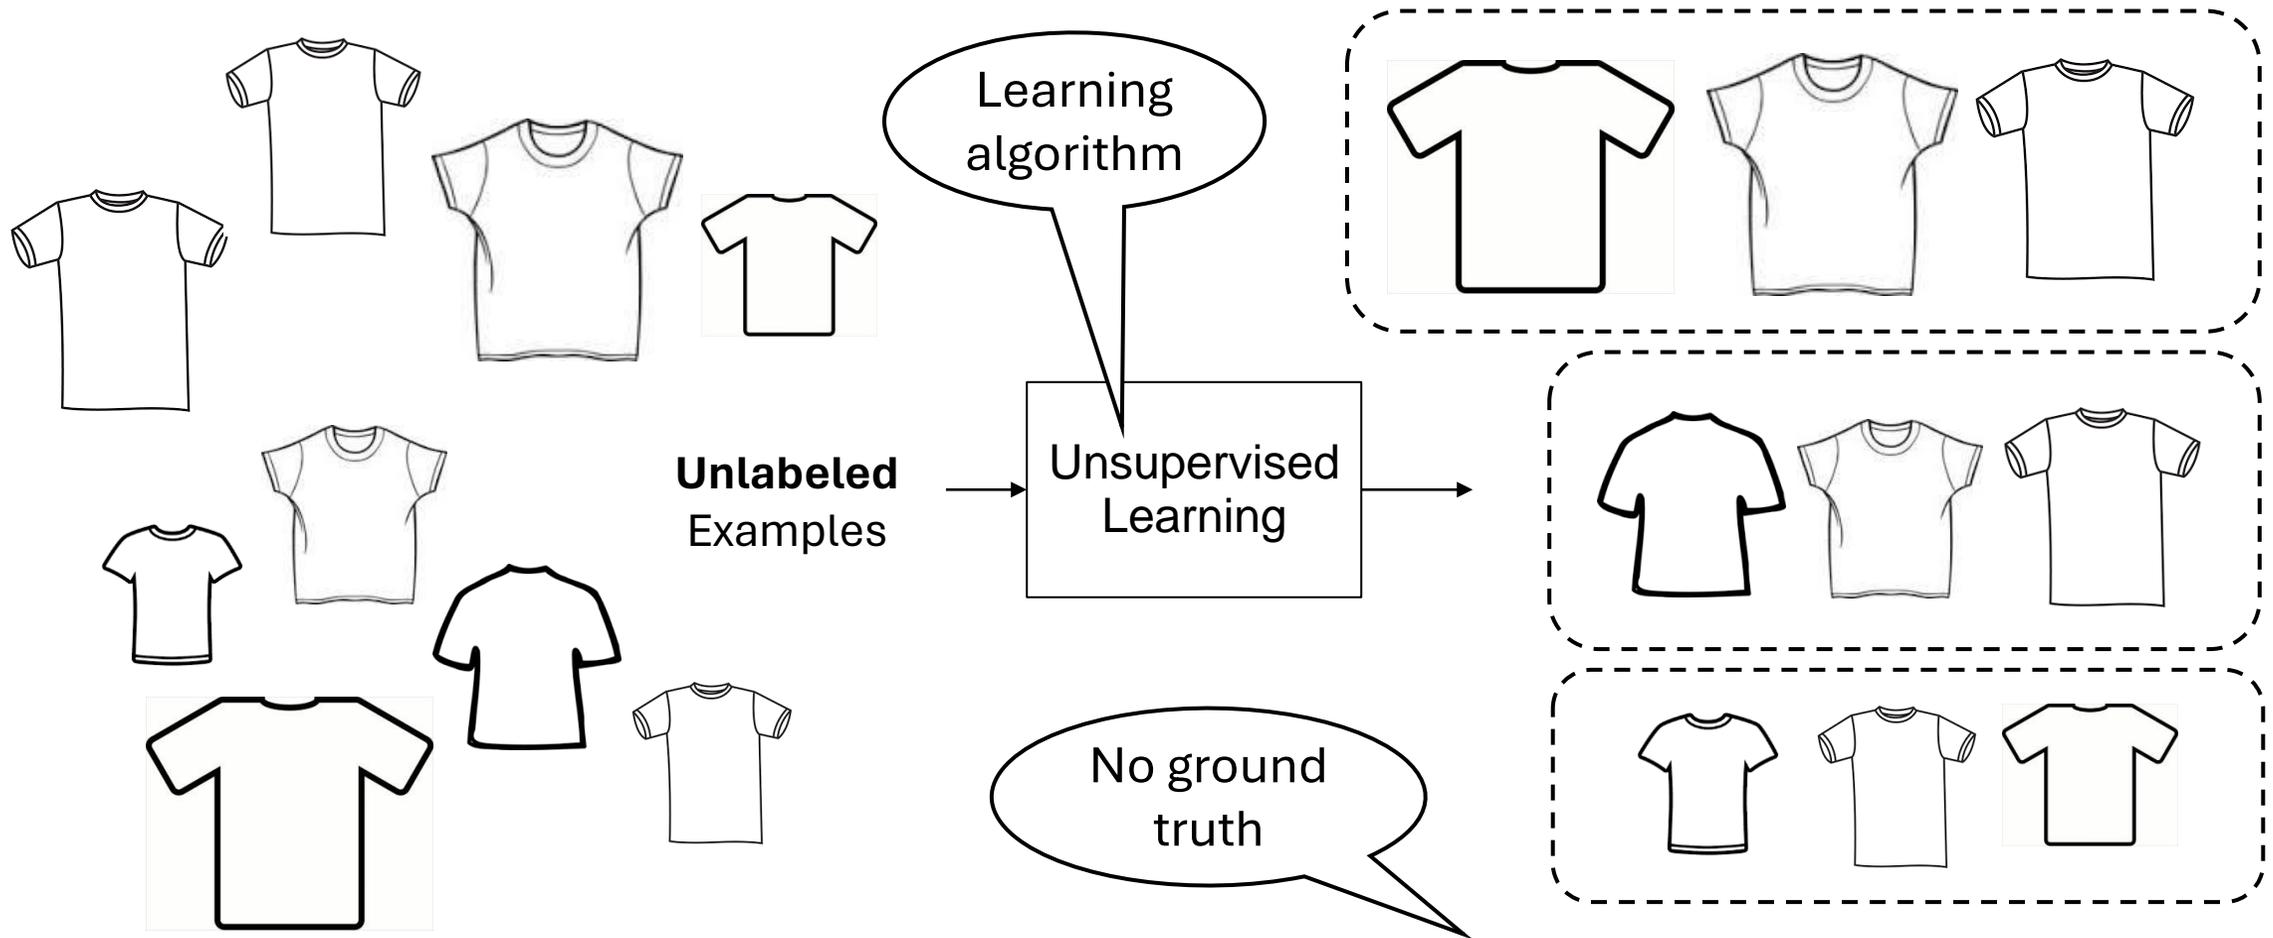
# Ways of machine learning

- Supervised learning
- Unsupervised learning
- Semi-supervised learning
- Self-supervised learning
- Reinforcement learning

# Supervised learning



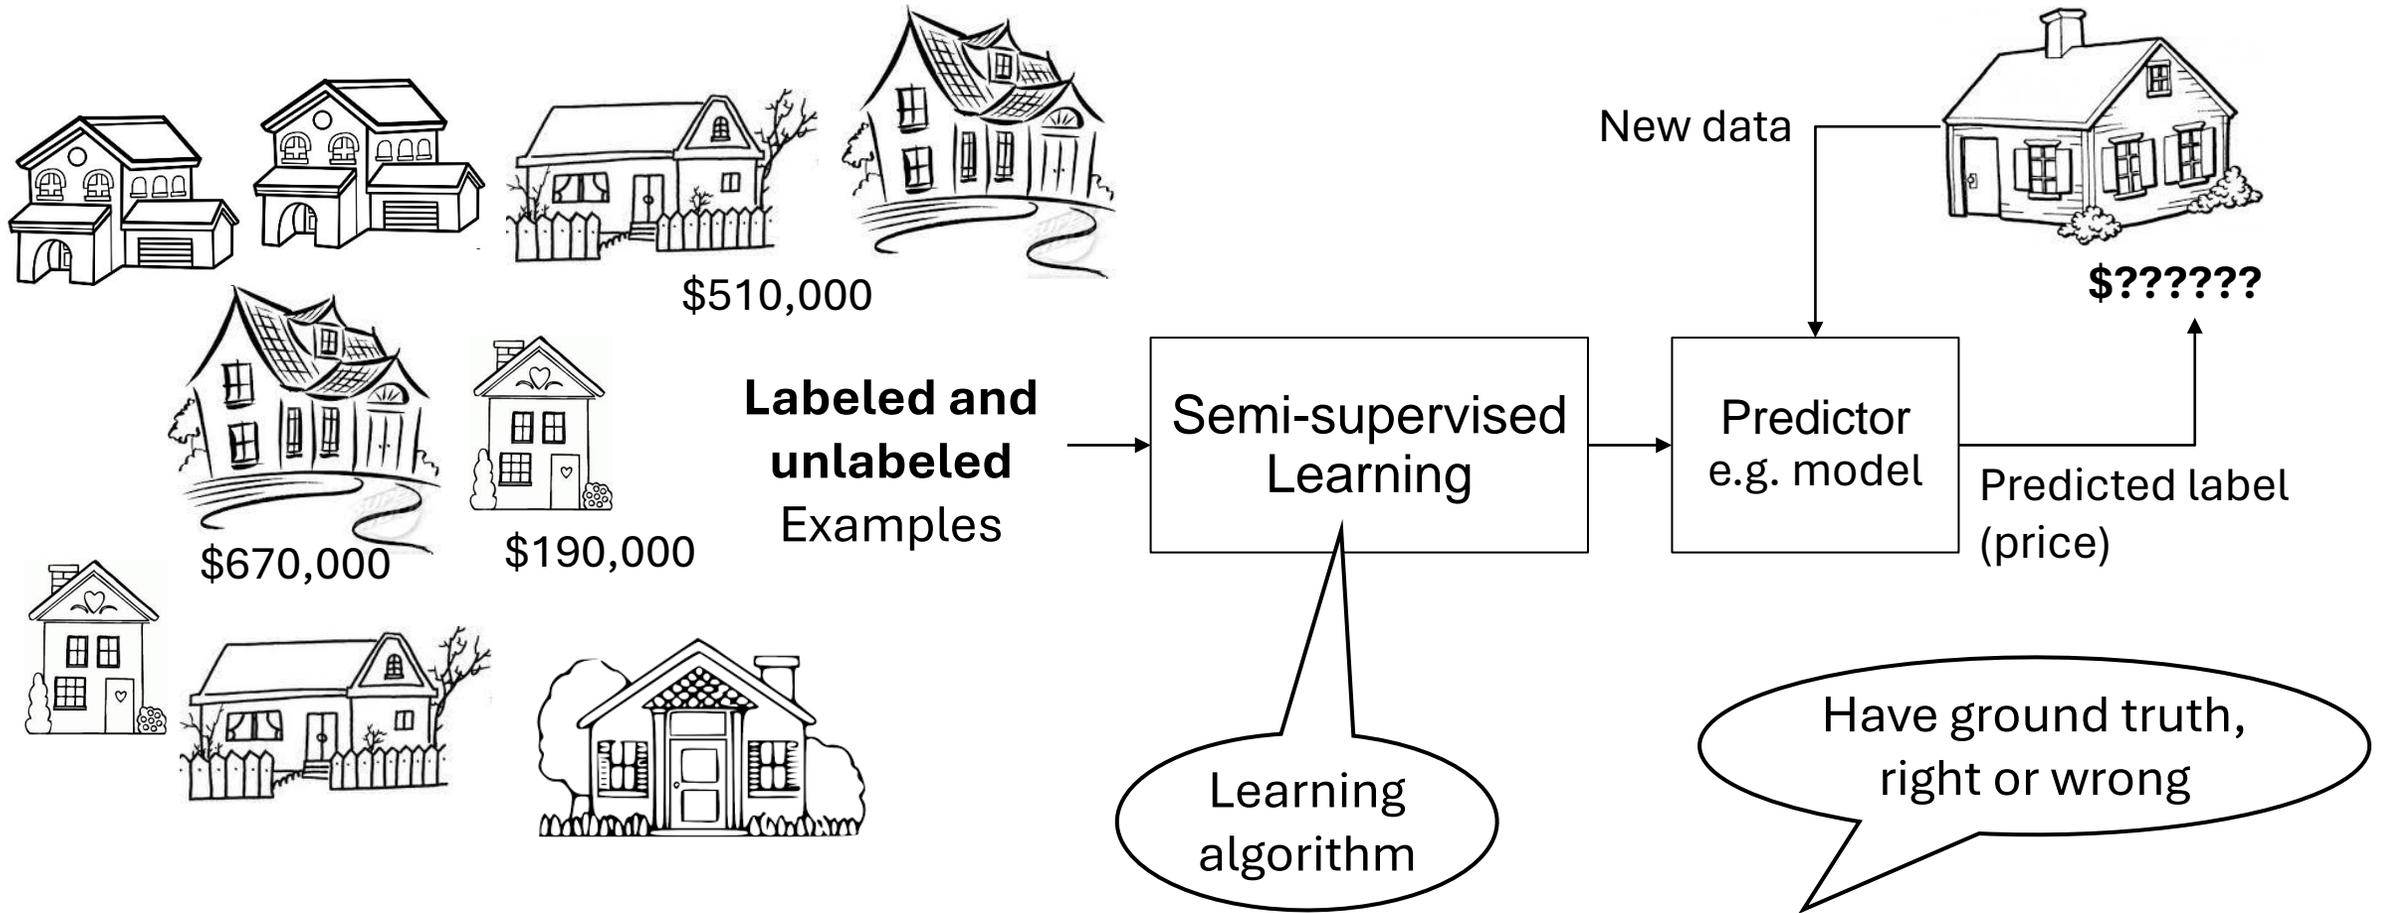
# Unsupervised learning



**Learning set:** data (or observations) **without label.**

**Discover** properties (e.g. size groups) in the data.

# Semi-supervised learning



**Training set:** data (or observations or examples) **with label** and **without label**.

**Predict** label (e.g. price, object) of new data.

# Self-supervised learning

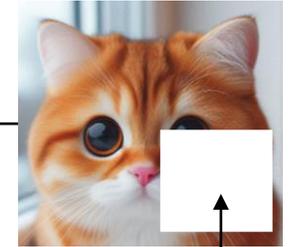


Unlabeled  
Examples

Self-supervised  
Learning

Learning  
algorithm

New data



???

Predictor  
e.g. model

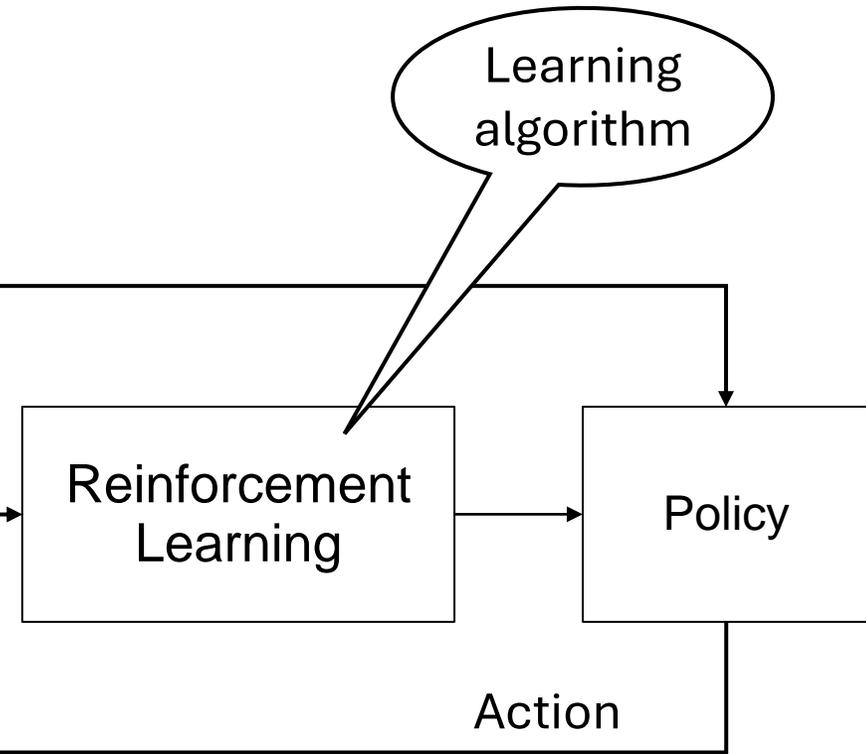
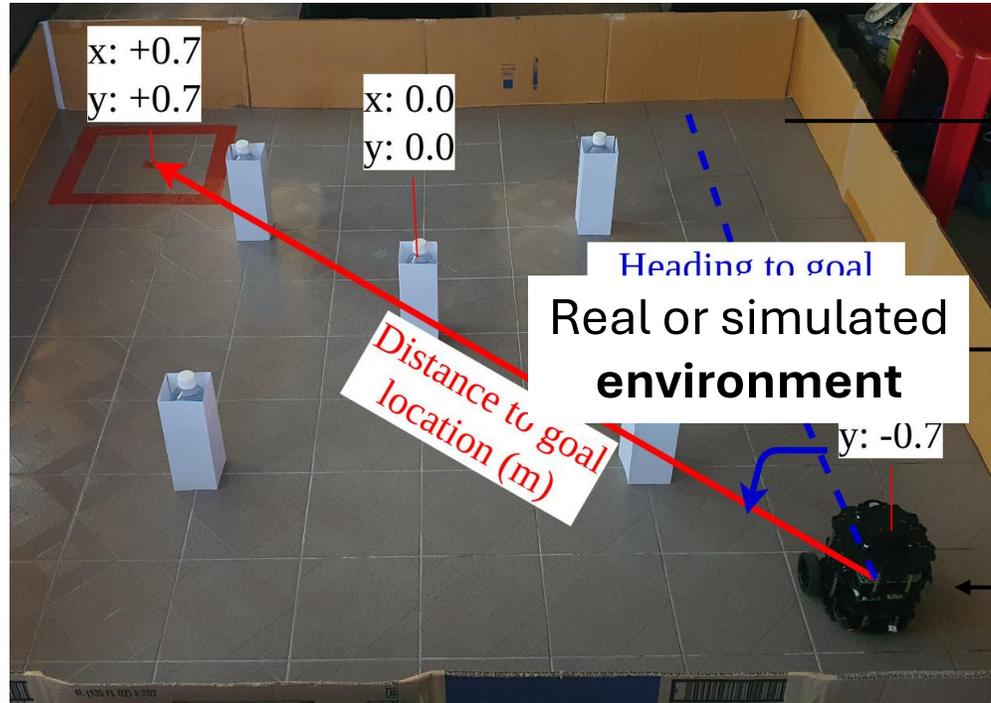
Predicted label

Have ground truth,  
right or wrong

**Training set:** data (or observations or examples) **without label.**

**Predict** “label” (e.g. object) of new data.

# Reinforcement learning



Learning from trials and errors

**Training set:** real or simulated environment, i.e. real-time observations (data)

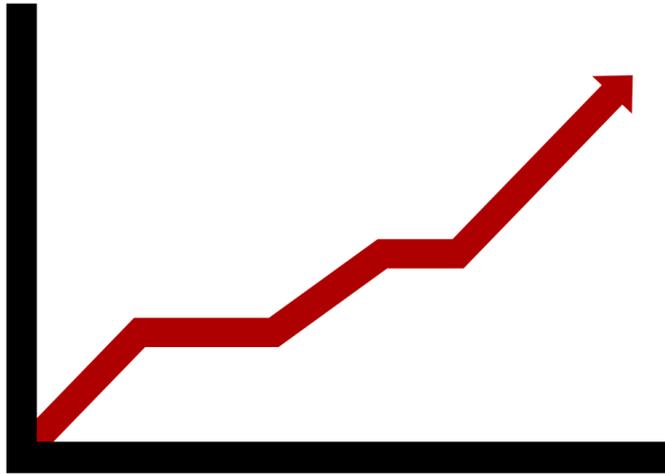
**Policy** maps environment conditions to suitable actions

Supervised machine learning

# Supervised learning: two types of problems

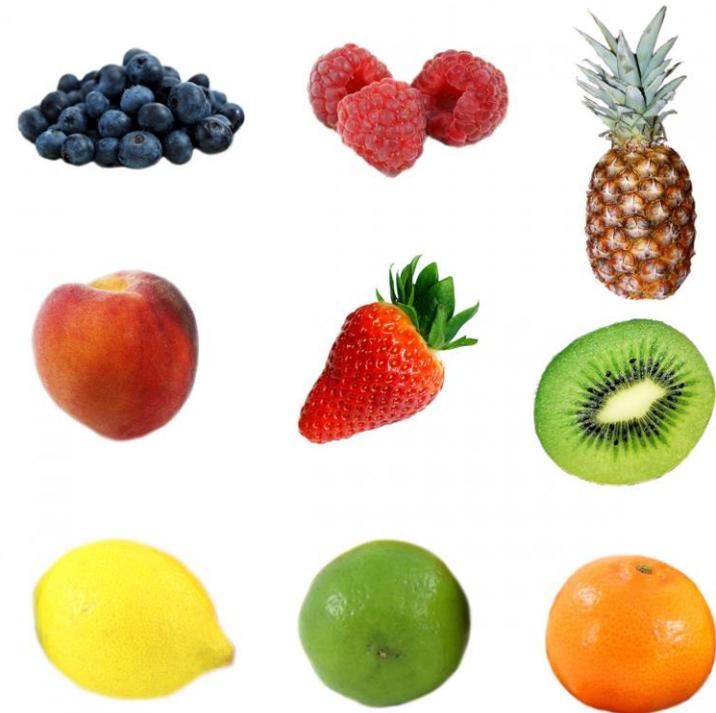
## Regression

- Output Y is **continuous**



## Classification

- Output Y is **discrete**



# Linear regression

Training Set  $D$

Size in feet<sup>2</sup> ( $x$ )

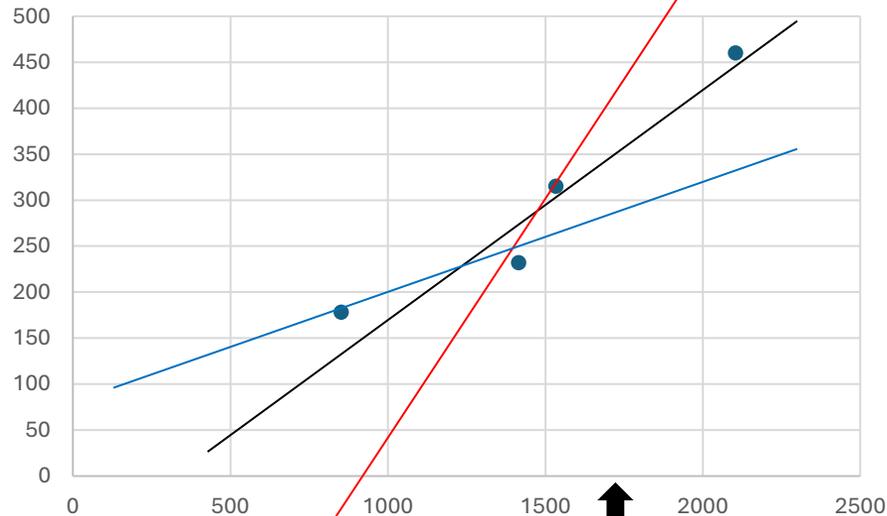
Price (\$) in 1000's ( $y$ )

training data points

2104  
1416  
1534  
852  
...

460  
232  
315  
178  
...

Guess of model type



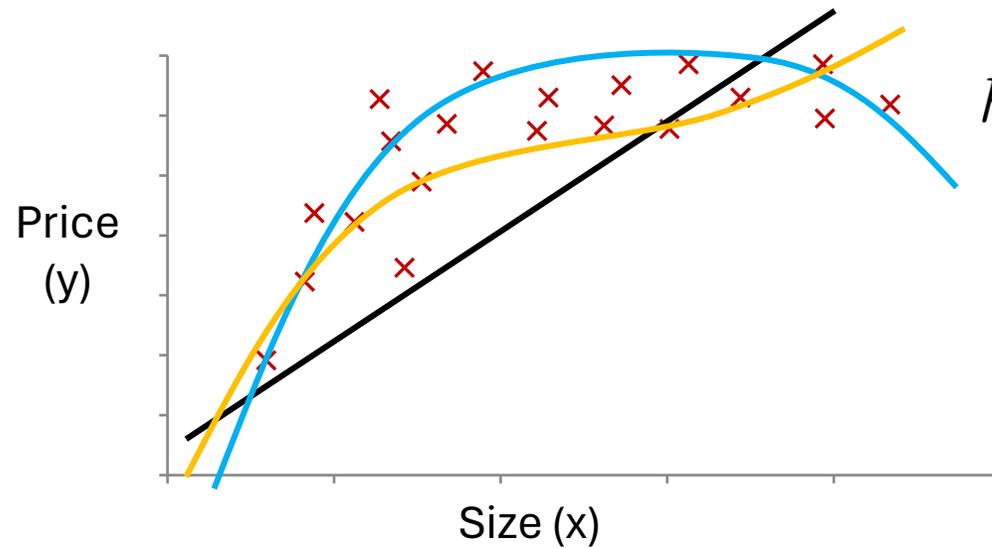
**Hypothesis**  $h: h(x) = h_{\theta}(x) = \theta_0 + \theta_1 x$

**Note:** Straight line eqn  $y = mx + c$

What are the appropriate values of  $\theta_0$  and  $\theta_1$ ?

parameters of  $h$

# Not just straight line: polynomial regression



$$h_{\theta}(x) = \theta_0 + \theta_1 x$$

$$\theta_0 + \theta_1 x + \theta_2 x^2$$

$$\theta_0 + \theta_1 x + \theta_2 x^2 + \theta_3 x^3$$

$$\begin{aligned} h_{\theta}(x) &= \theta_0 + \theta_1 x_1 + \theta_2 x_2 + \theta_3 x_3 \\ &= \theta_0 + \theta_1(\text{size}) + \theta_2(\text{size})^2 + \theta_3(\text{size})^3 \end{aligned}$$

$$x_1 = (\text{size})$$

$$x_2 = (\text{size})^2$$

$$x_3 = (\text{size})^3$$

# Prediction error

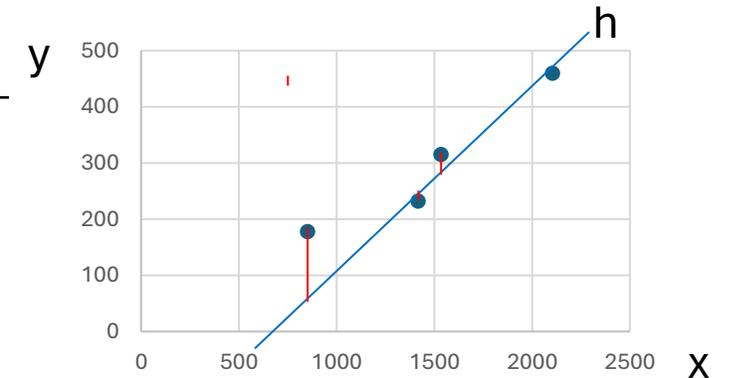
Lets try different values for the parameters and compute the prediction error:

$$error = h(x) - y$$

$$\theta_0 = -234, \theta_1 = 0.33$$

$$h(x) = -234 + 0.33x$$

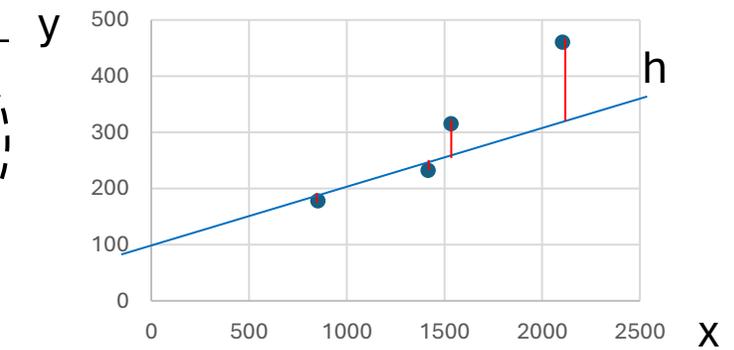
	x	y	h	Error
1	2104	460	460.32	0.32
2	1416	232	233.28	1.28
3	1534	315	272.22	-42.78
4	852	178	47.16	-130.84



$$\theta_0 = 92.8, \theta_1 = 0.1$$

$$h(x) = 92.8 + 0.1x$$

	x	y	h	Error
1	2104	460	303.2	-156.8
2	1416	232	234.4	2.4
3	1534	315	246.2	-68.8
4	852	178	178	0



... keep trying different values.

Cost function = Mean Square Error (MSE)  
(or Root Mean Square Error (RMSE))

# Prediction error

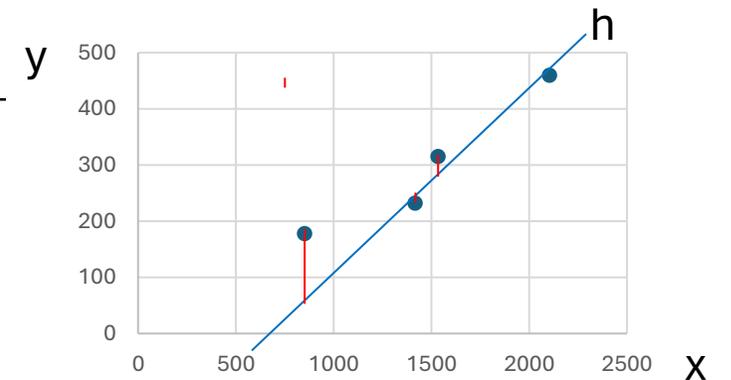
Lets try different values for the parameters and compute the prediction error:

$$error = h(x) - y$$

$$\theta_0 = -234, \theta_1 = 0.33$$

$$h(x) = -234 + 0.33x$$

	x	y	h	Error
1	2104	460	460.32	0.32
2	1416	232	233.28	1.28
3	852	315	272.22	-42.78
4	2104	178	460.32	-130.84

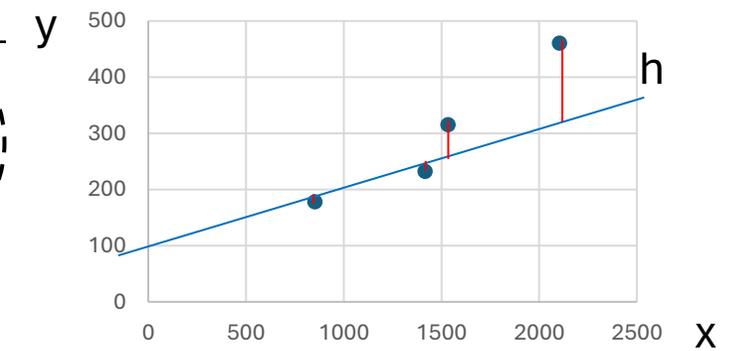


Use **gradient descent** to find parameters that minimizes the cost function

$$\theta_0 = 92.8, \theta_1 = 0.1$$

$$h(x) = 92.8 + 0.1x$$

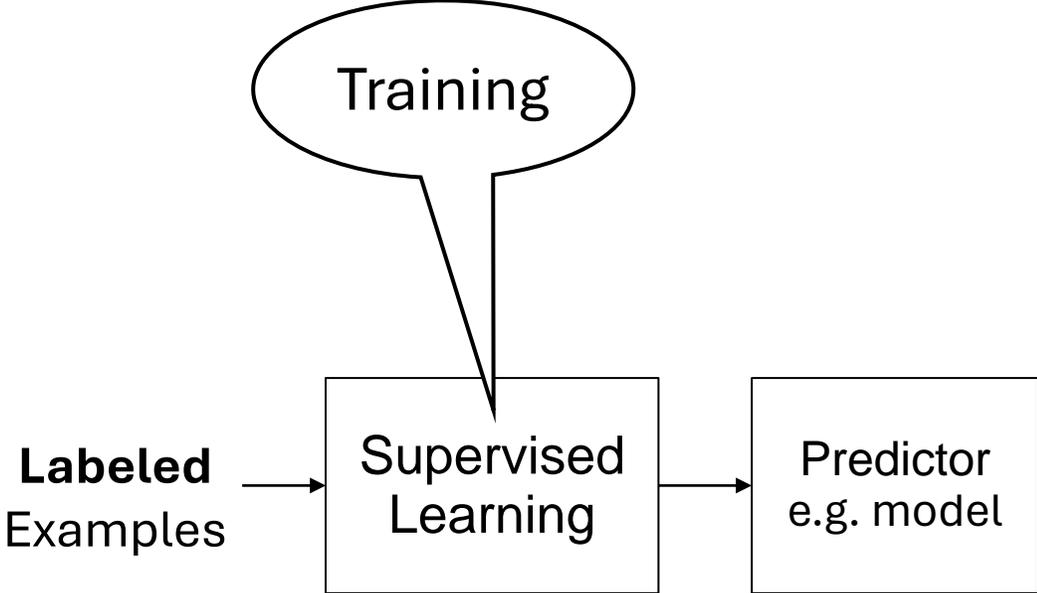
	x	y	h	Error
1	2104	460	303.2	-156.8
2	1416	232	234.4	2.4
3	852	315	246.2	-68.8
4	2104	178	178	0



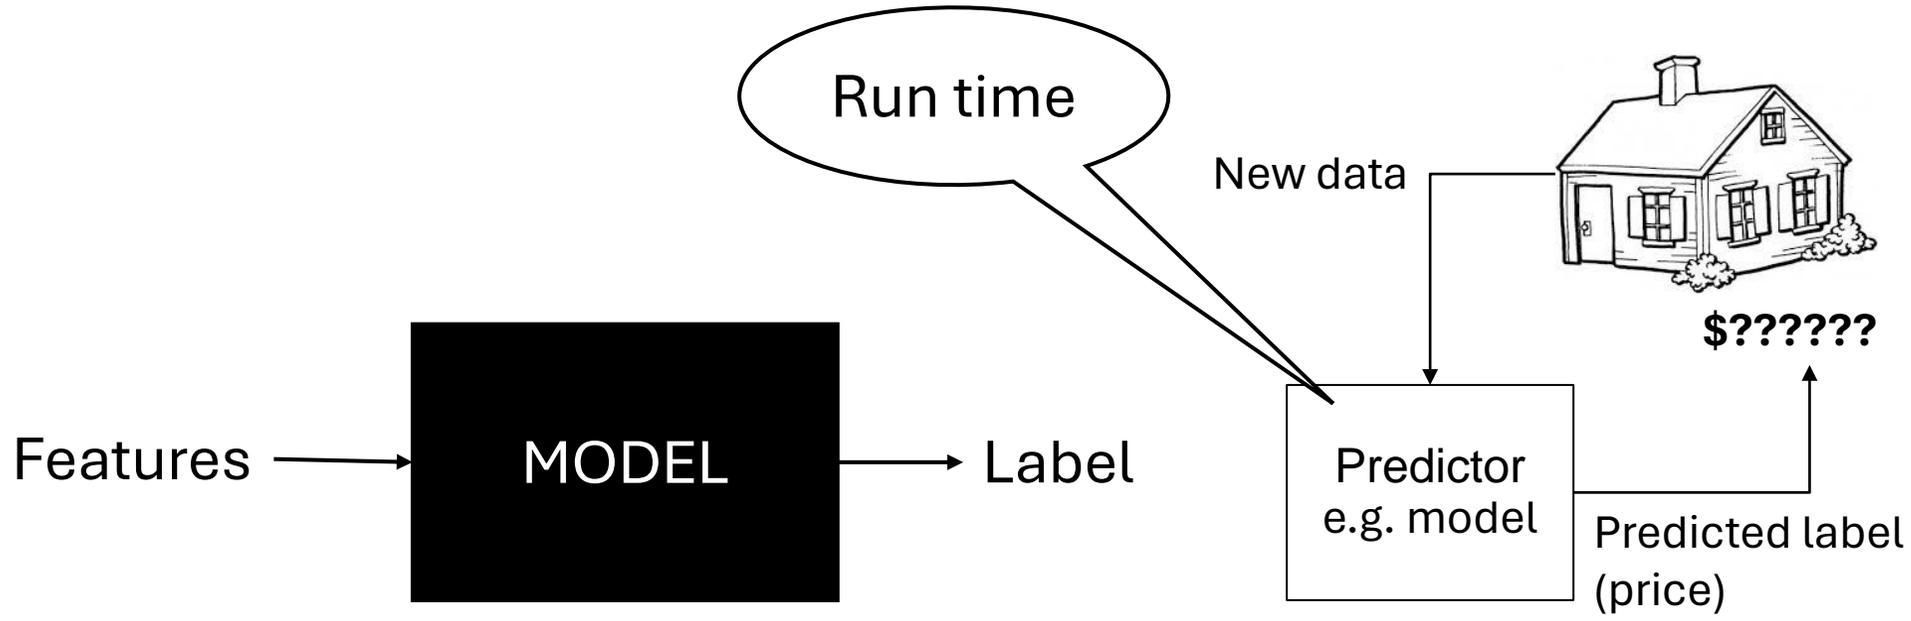
... keep trying different values.

Cost function = Mean Square Error (MSE)  
(or Root Mean Square Error (RMSE))

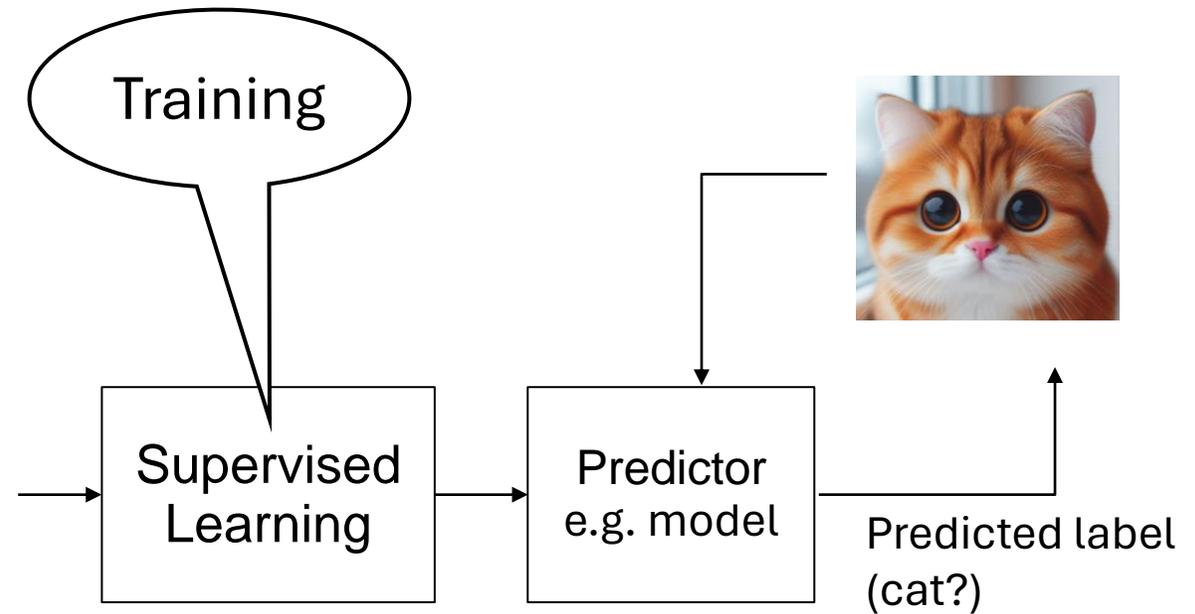
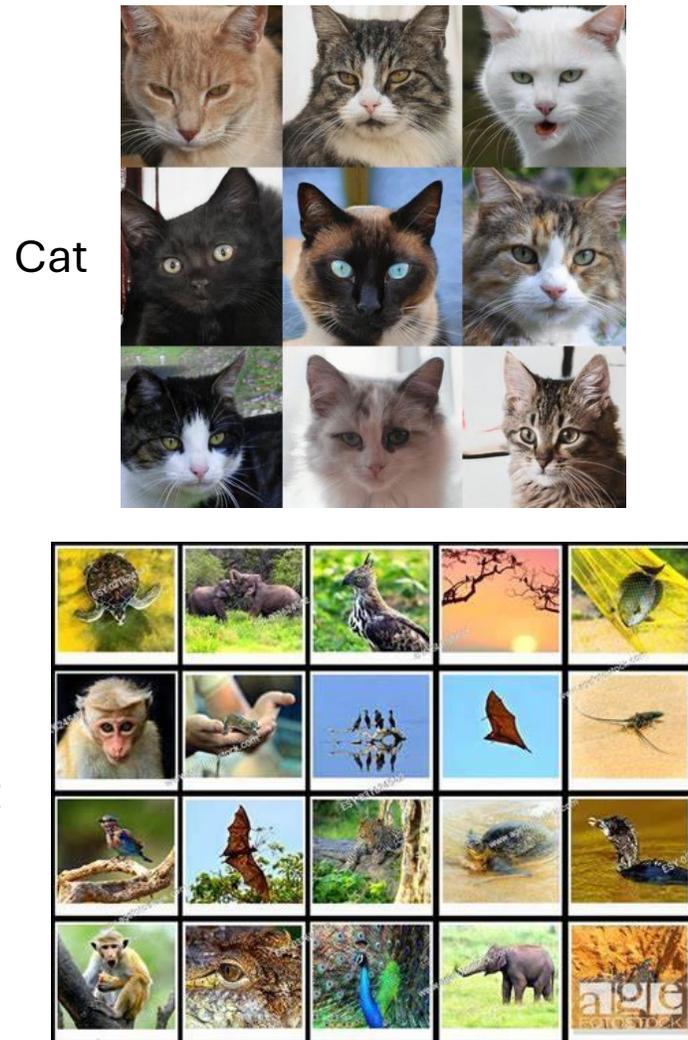
# Two phases of supervised ML



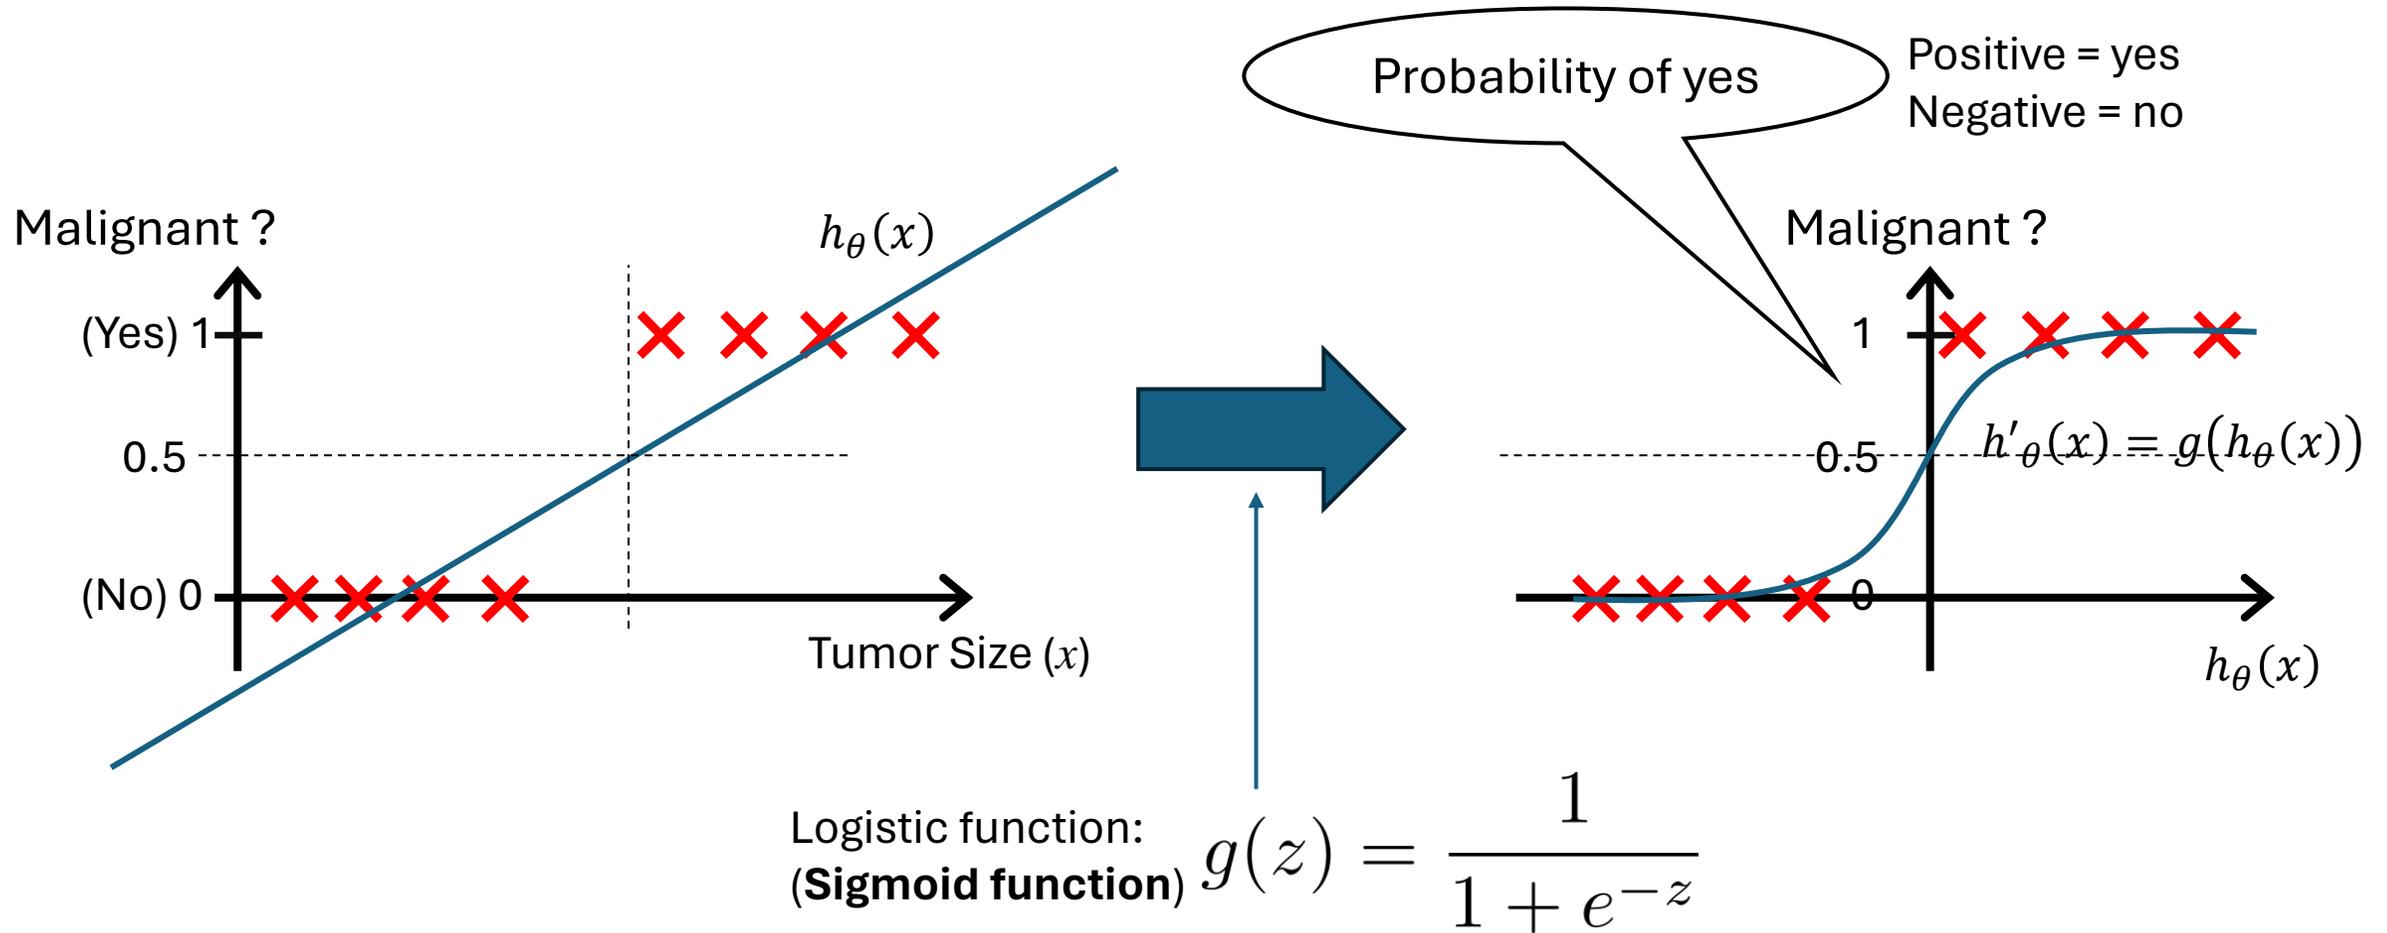
# Two phases of supervised ML



# Classification

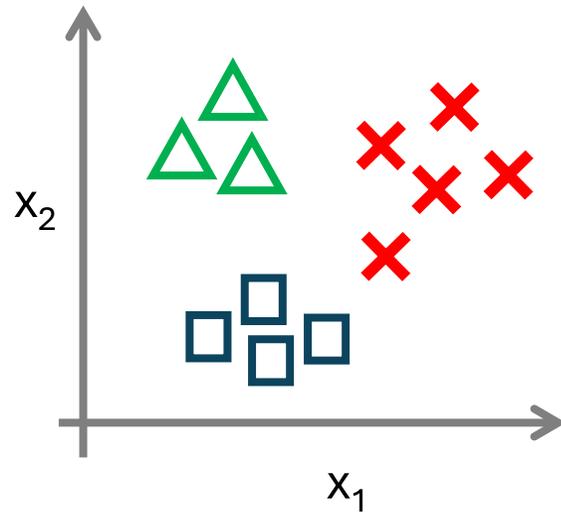


# Classification: logistic regression

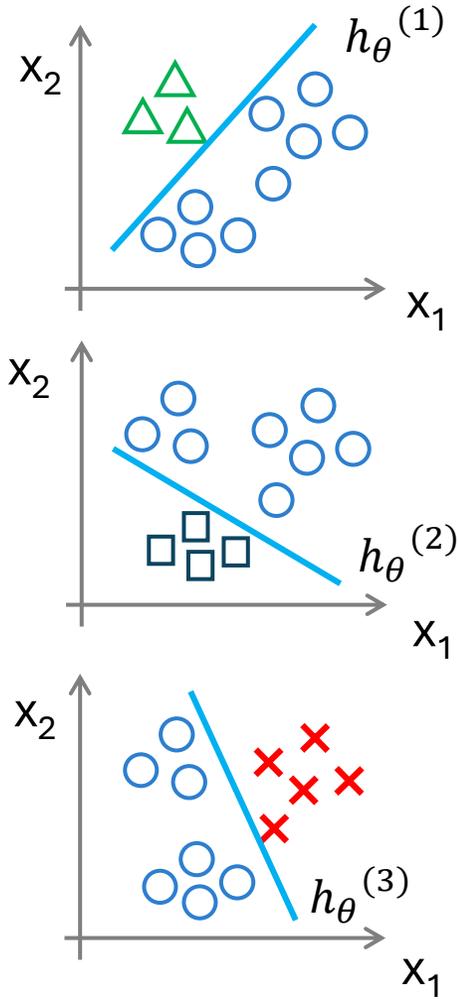




# One-vs-all



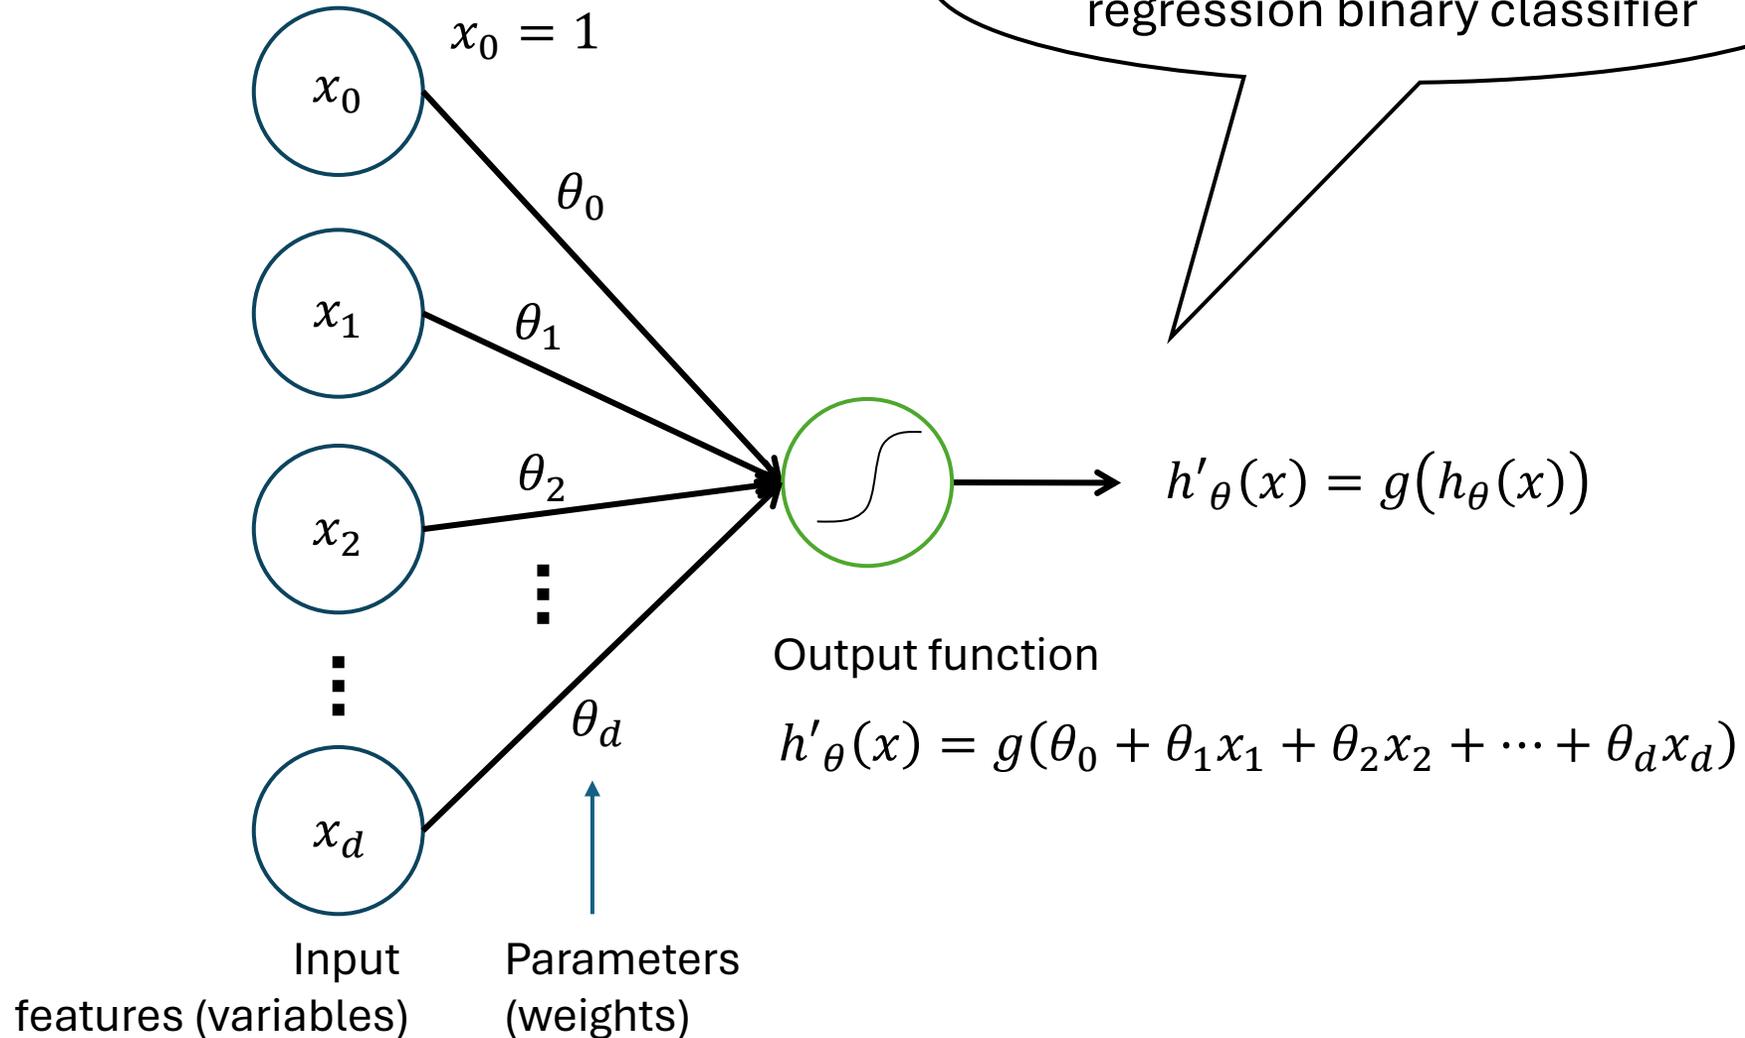
- Class 1: 
- Class 2: 
- Class 3: 



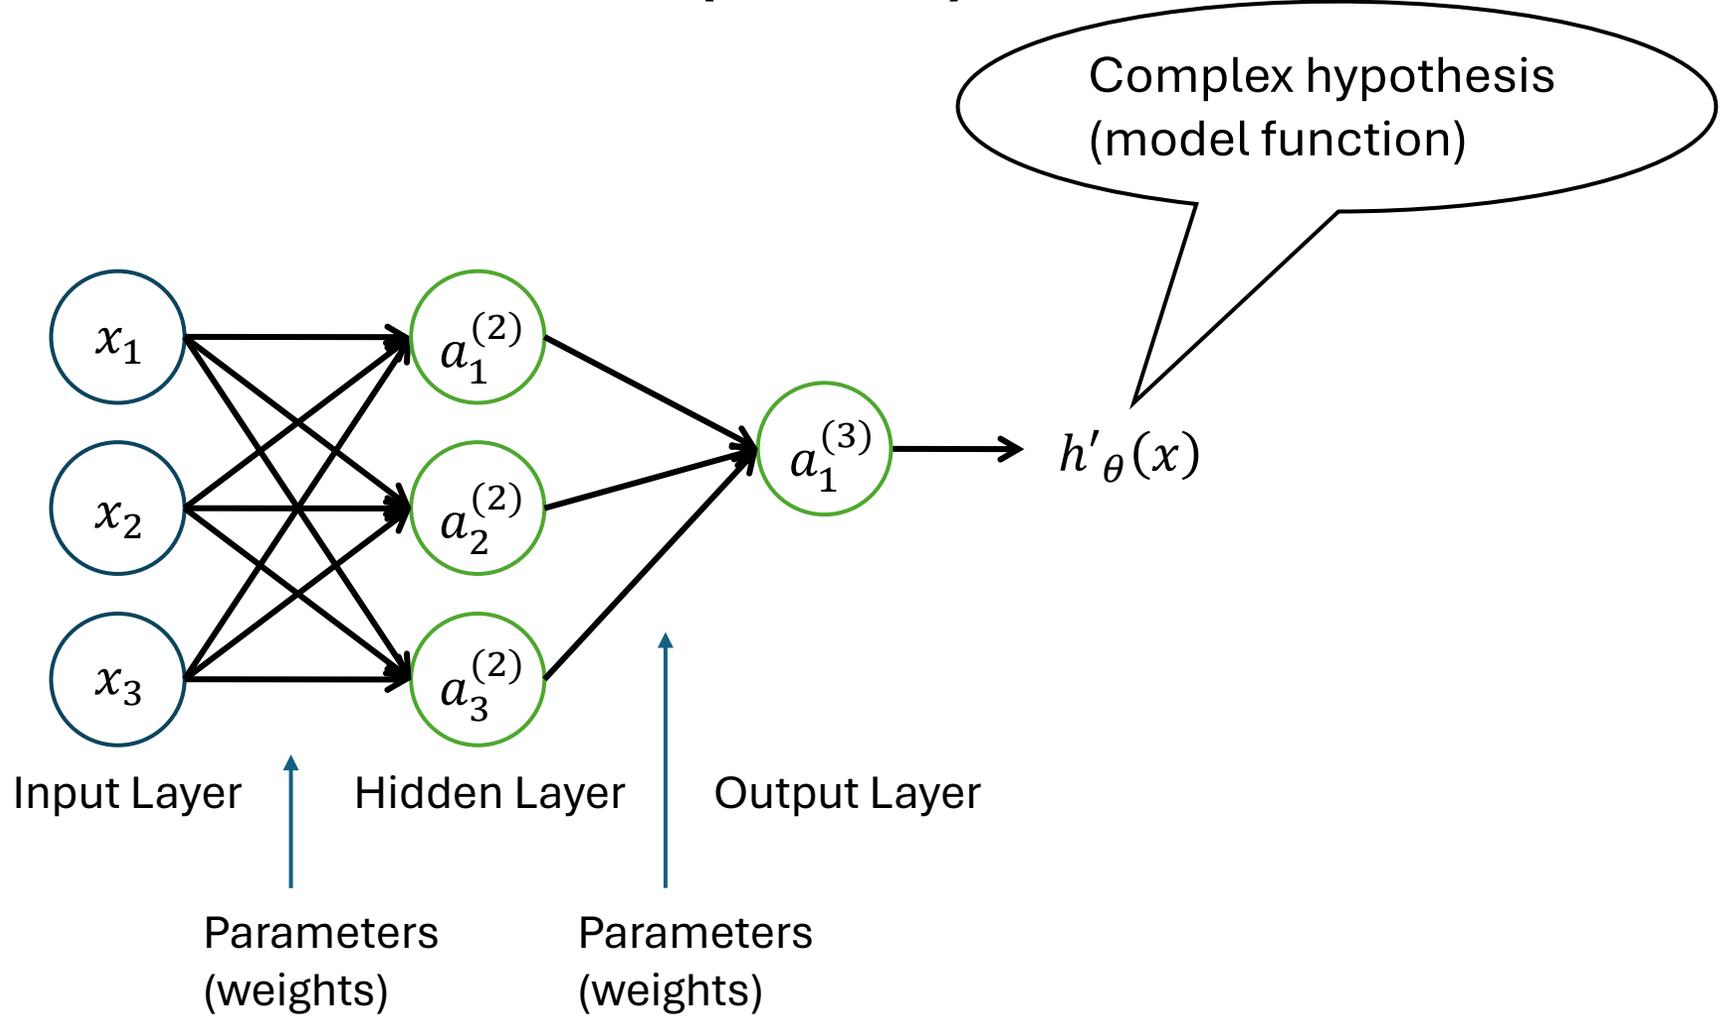
Select label with highest probability

Train the same way as linear regression

# A neuron: perceptron



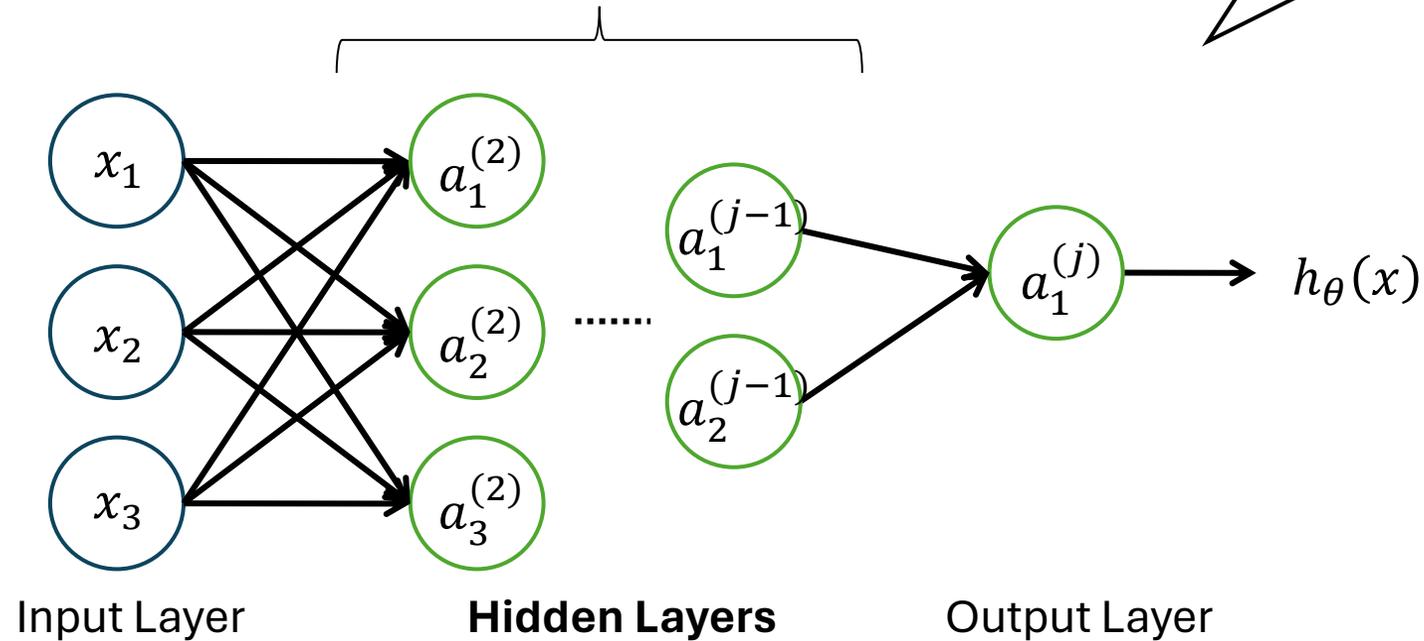
# Artificial Neural Network (ANN)



# Deep Neural Network (DNN)

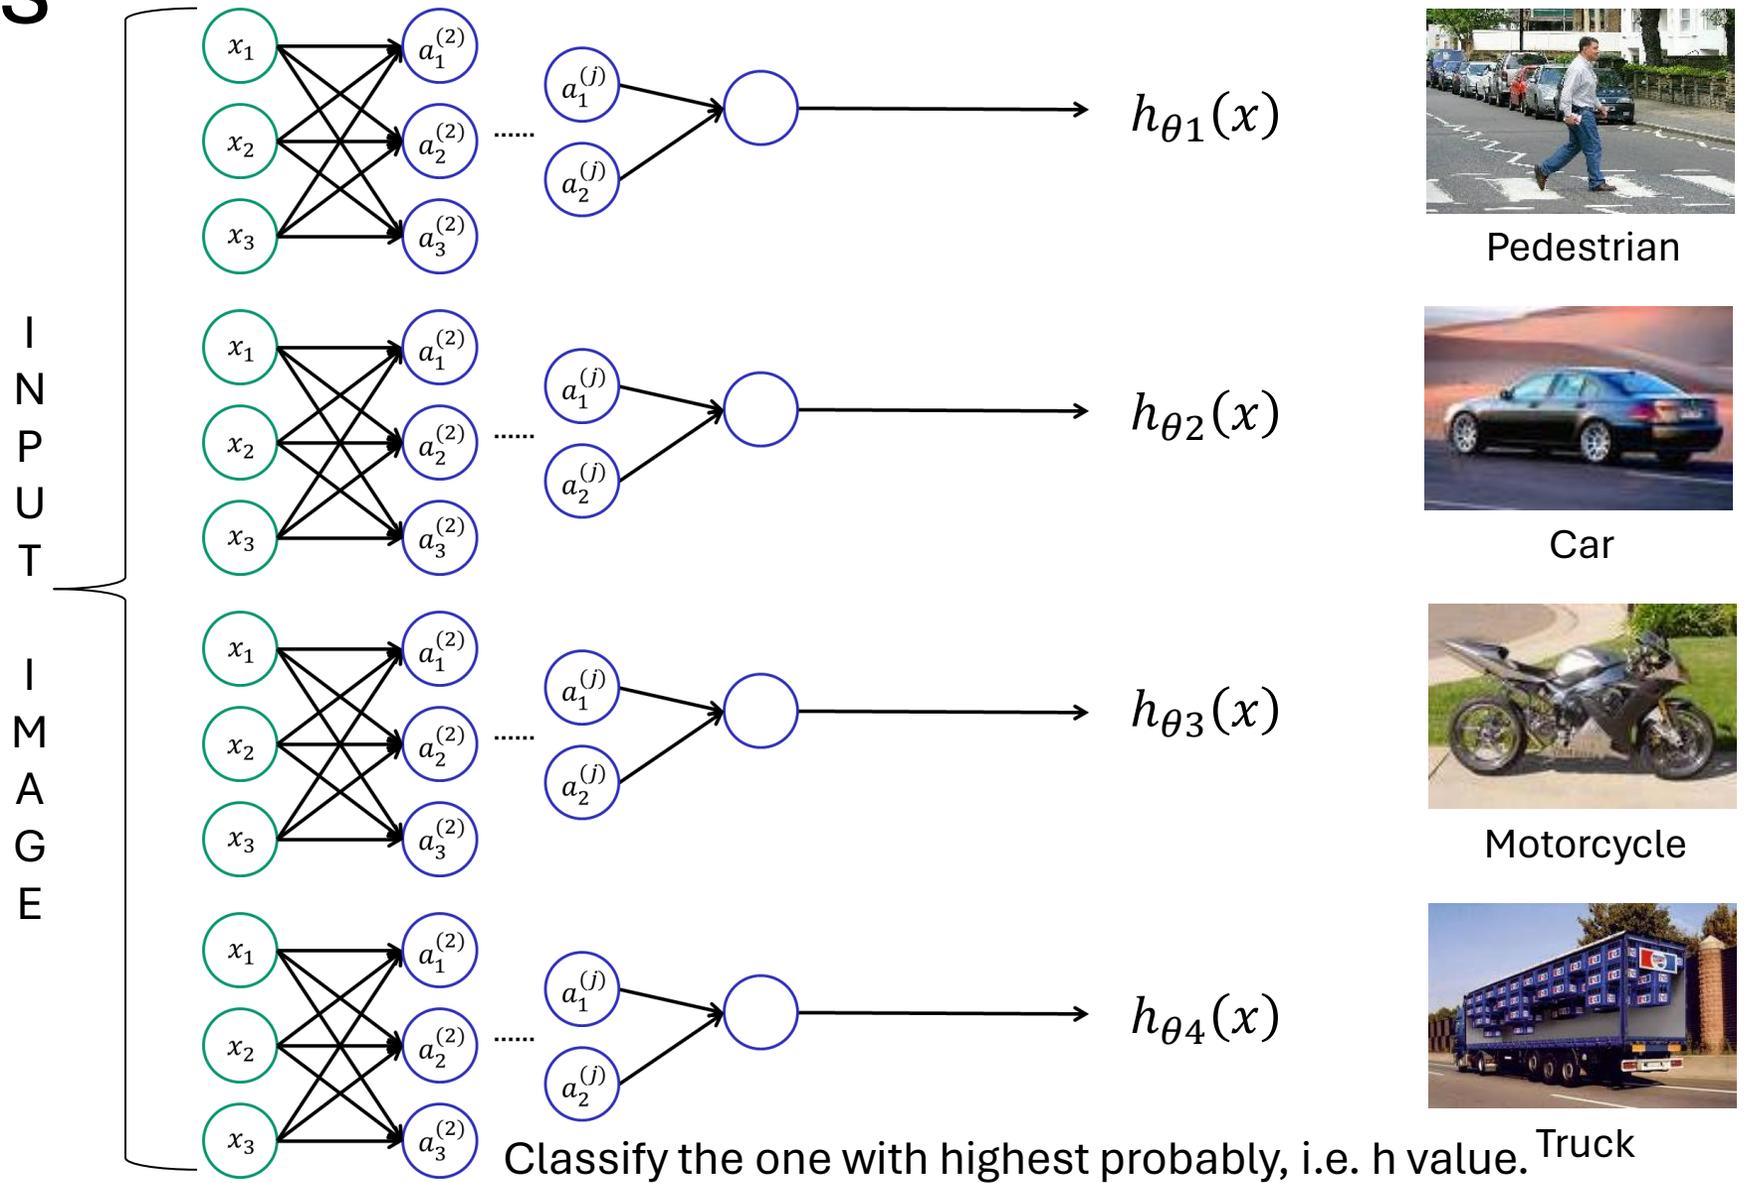
Deep learning (DL)

There may be more than one hidden layer.



Considered **DNN**  
when  $>1$  hidden layer

# Multiclass ANN



# Multiclass ANN



Pedestrian



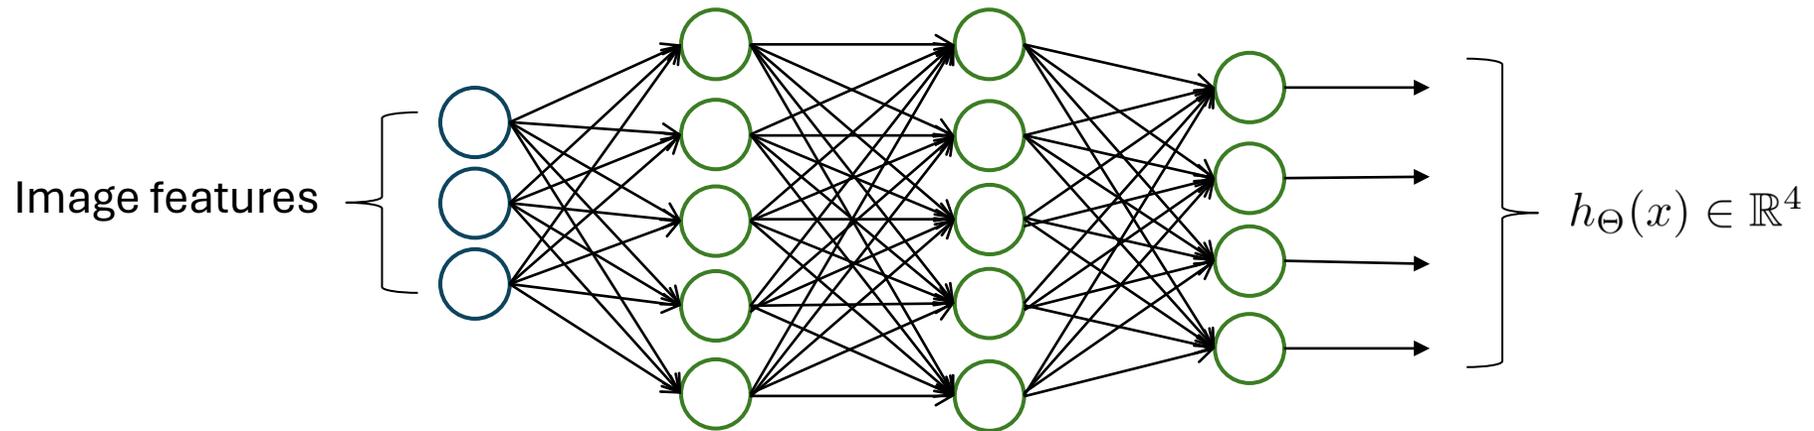
Car



Motorcycle



Truck



We want  $h_{\Theta}(x) \approx \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}$ ,  $h_{\Theta}(x) \approx \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}$ ,  $h_{\Theta}(x) \approx \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}$ , etc.  
when pedestrian      when car      when motorcycle

# Multiclass ANN



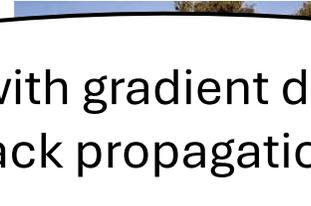
Pedestrian



Car

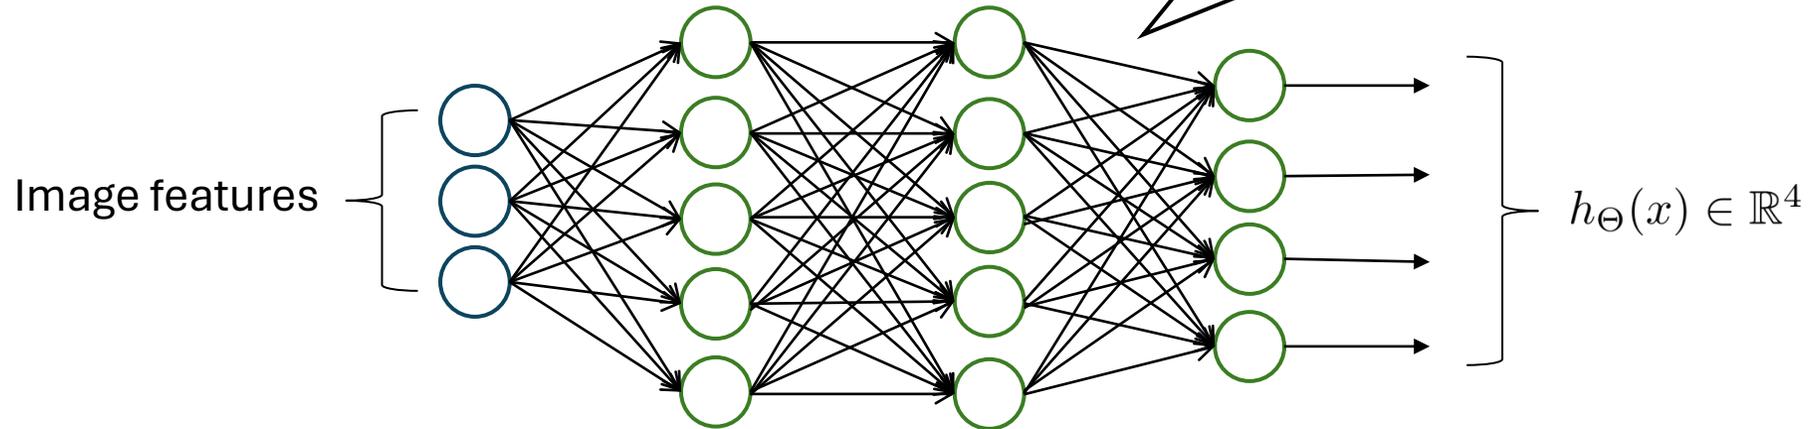


Motorcycle



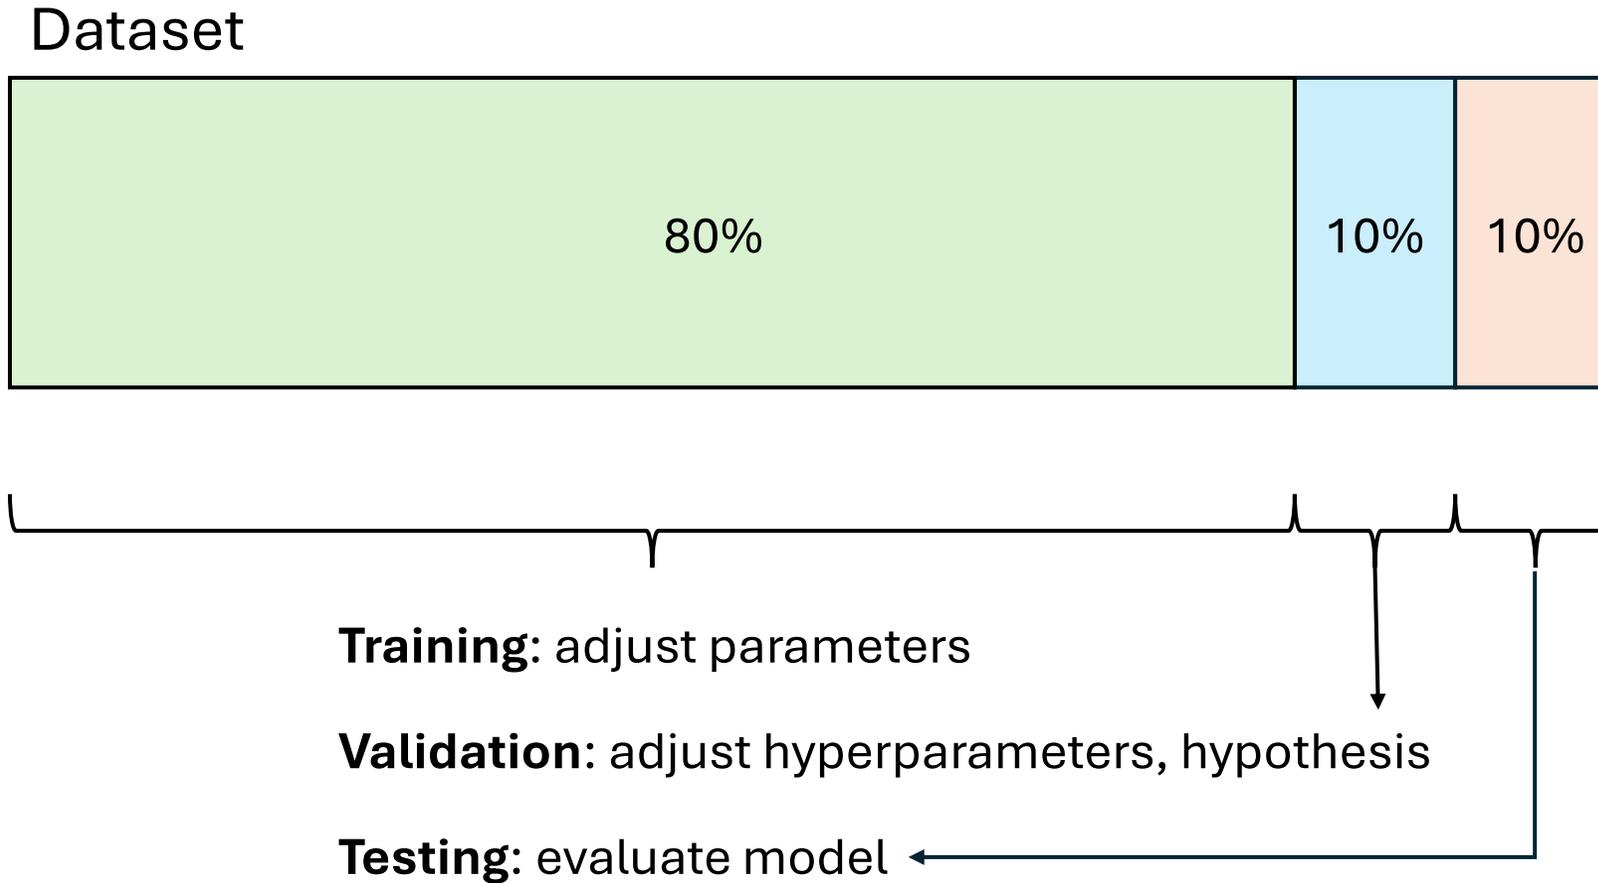
Truck

Train with gradient descent and back propagation



We want  $h_{\Theta}(x) \approx \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}$ ,  $h_{\Theta}(x) \approx \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}$ ,  $h_{\Theta}(x) \approx \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}$ , etc.  
when pedestrian      when car      when motorcycle

# Train, validate and test

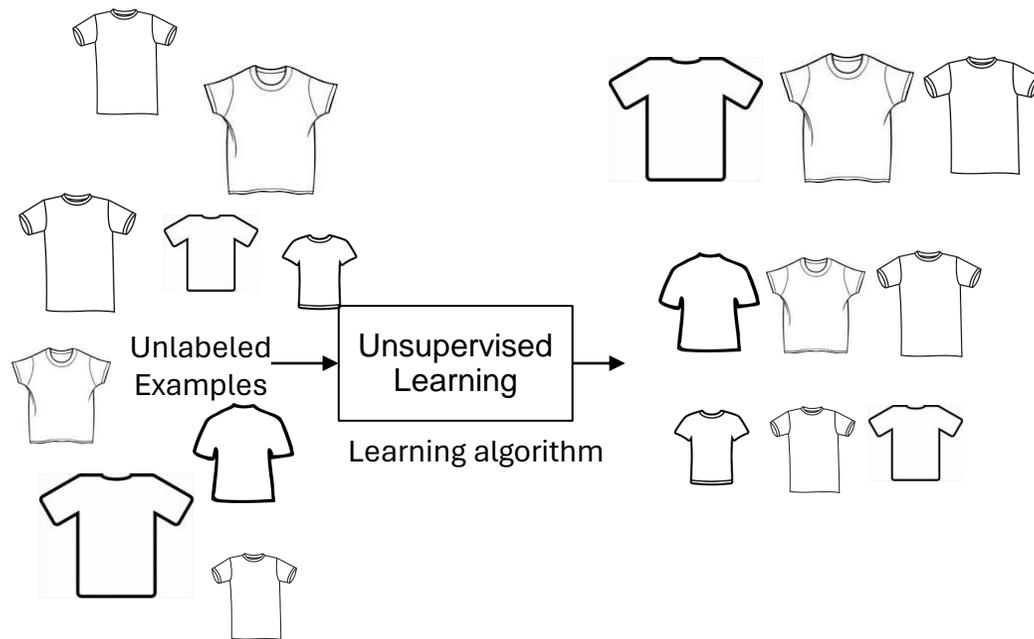


# Unsupervised machine learning

# Unsupervised learning: two types of problems

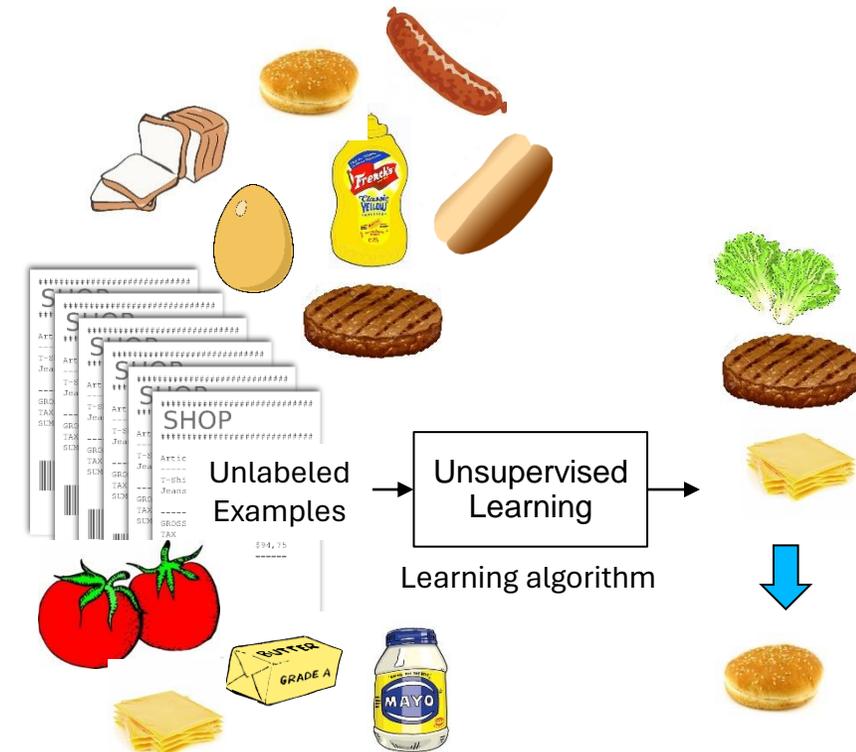
## Clustering

- discover groupings

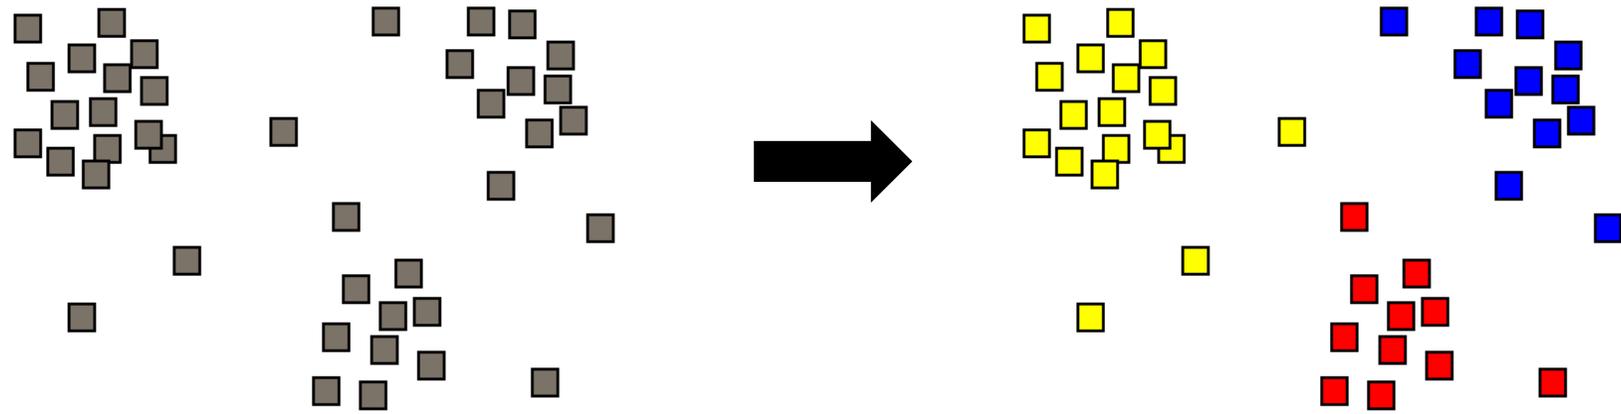


## Association

- discover rules, relations



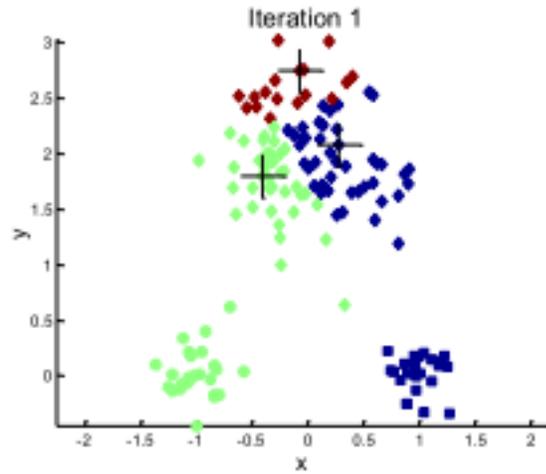
# K-Means clustering



Each group is a **cluster**

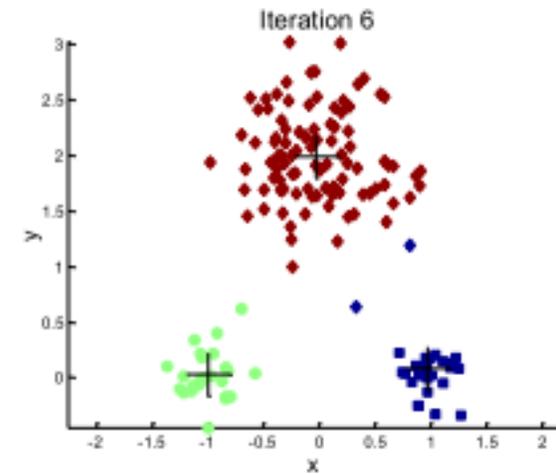
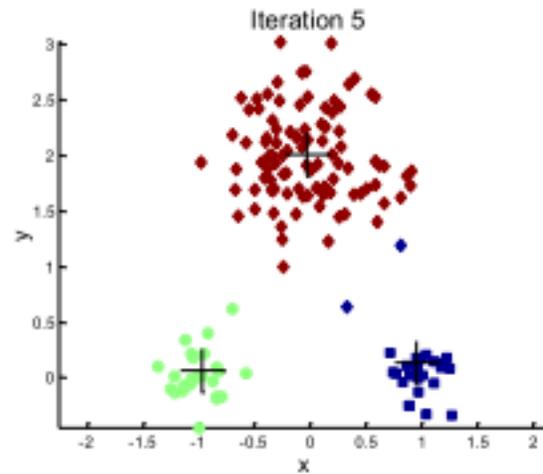
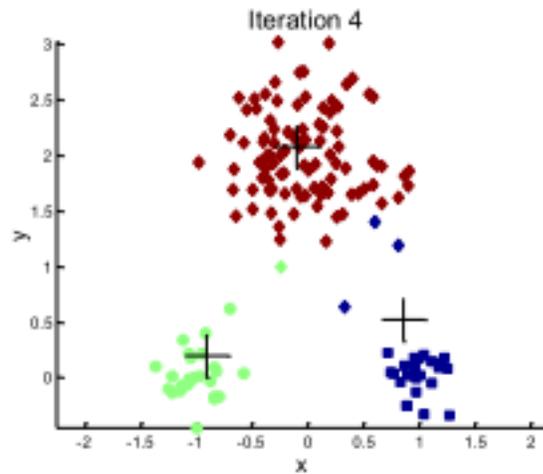
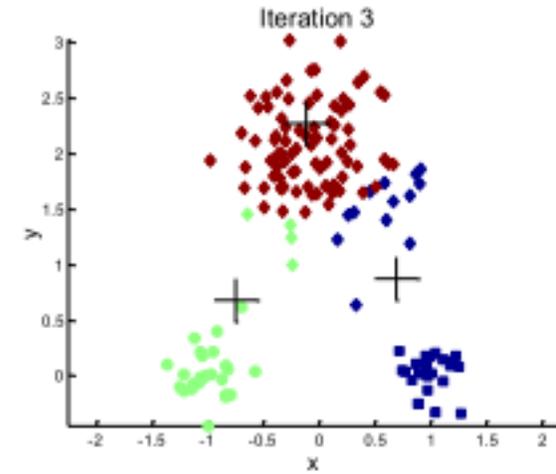
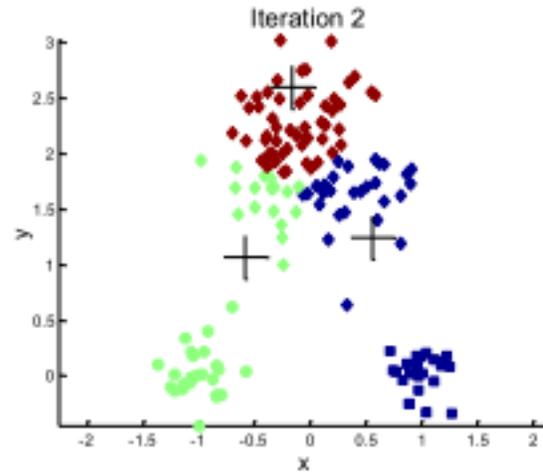
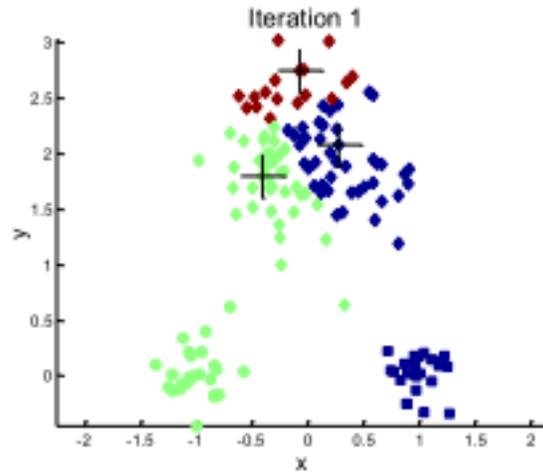
# K-Means clustering steps

K=3



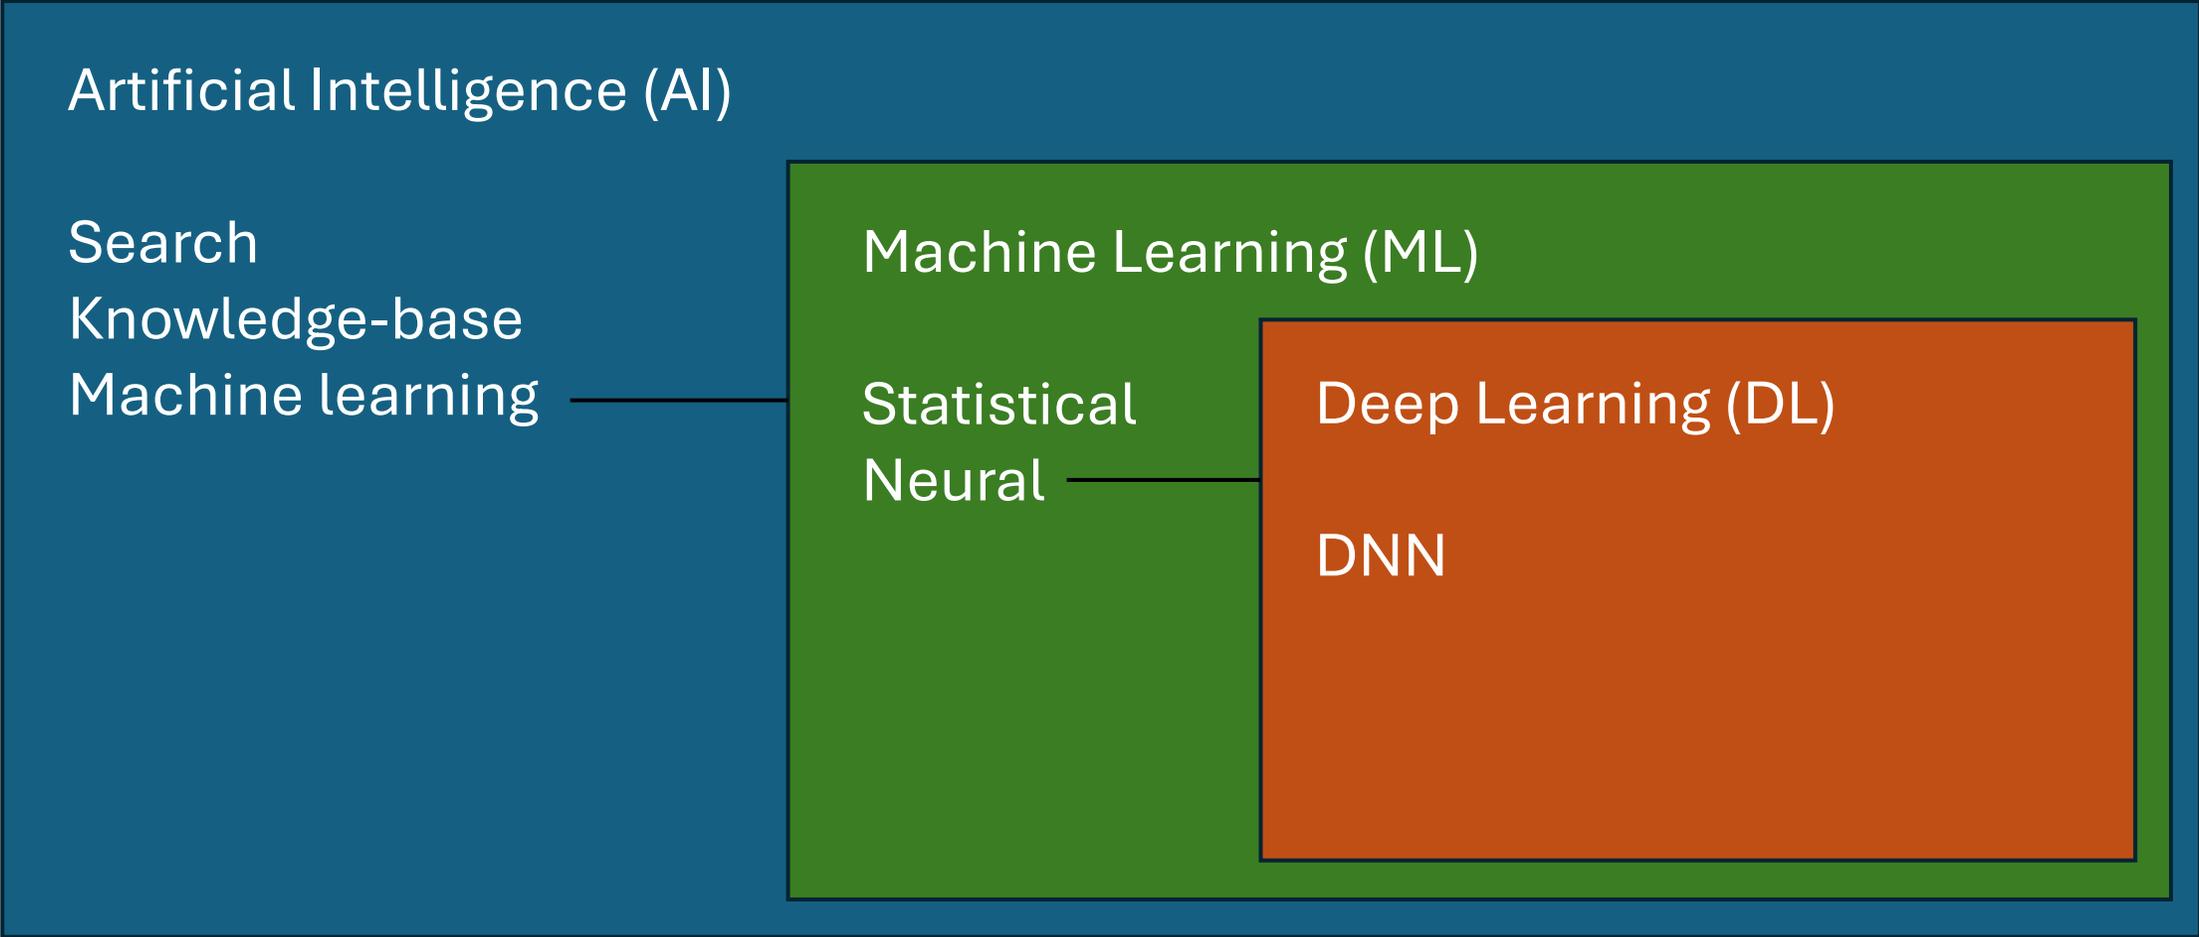
# K-Means clustering steps

K=3



**AI, ML, DL**

# AI, ML, DL



Evaluation



<https://forms.gle/ZWHTNJw3m5qDgR277>

Feedback



<https://forms.gle/o7adMBRjhdnLb2d6A>

Thank you!