PRE-UNIVERSITY SUMMER SCHOOL

AI: How Machines Learn, Predict and Create

Introduction to the ideas, possibilities and limits of artificial intelligence

UNIVERSITY OF WARWICK

Paris Giampouras · Department of Computer Science

AI

How many times have you already used AI today?

Think beyond chatbots…

maps, search, cameras, feeds, games, spam filters…

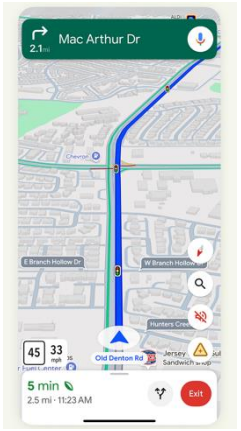

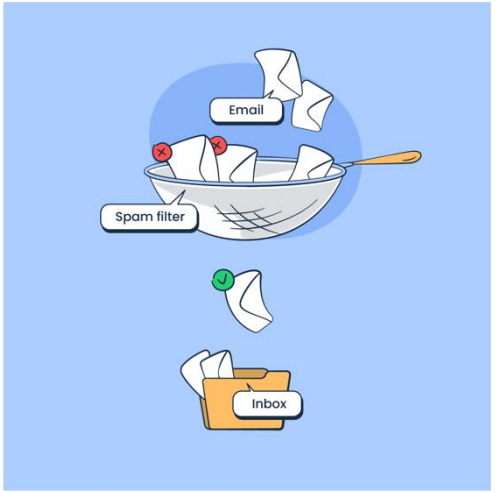

AI is already woven into ordinary decisions

See

Phone cameras improve images. Security systems notice unusual patterns. Medical scans can be prioritised for review.

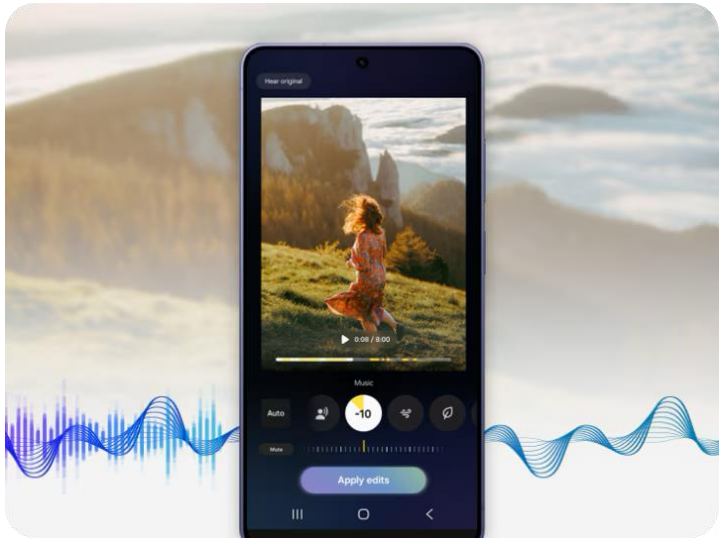

AI is already woven into ordinary decisions

Choose

Search engines rank results. Recommenders select videos and music. Maps compare possible routes.

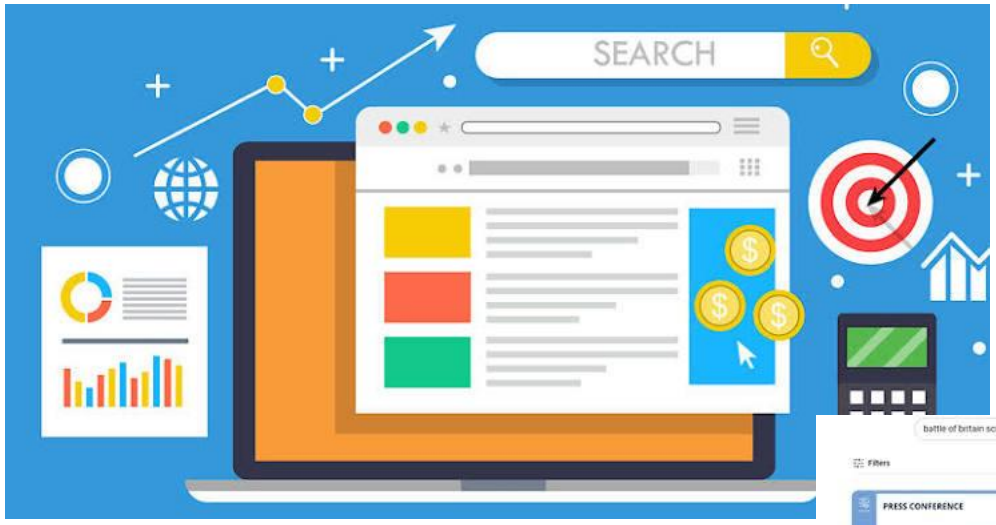

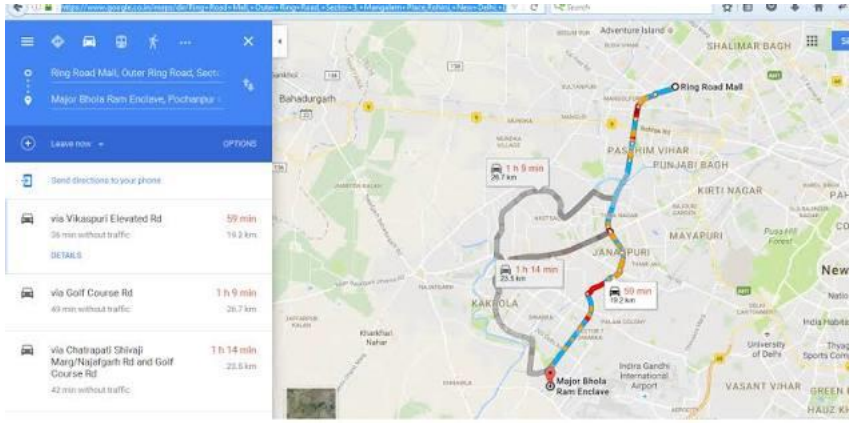

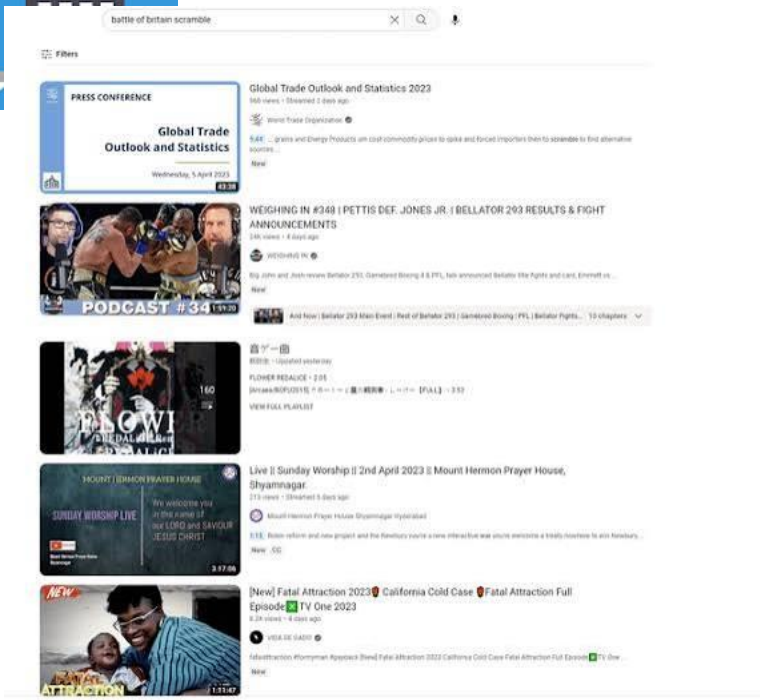

AI is already woven into ordinary decisions

Create

Generative models can produce text, images, audio, video and code from instructions and examples.

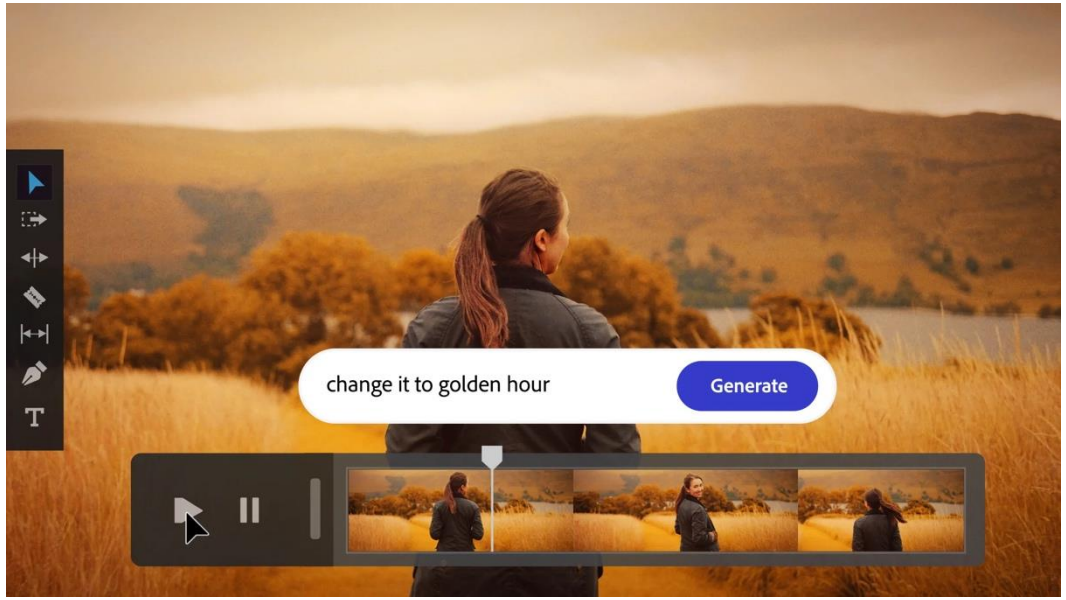

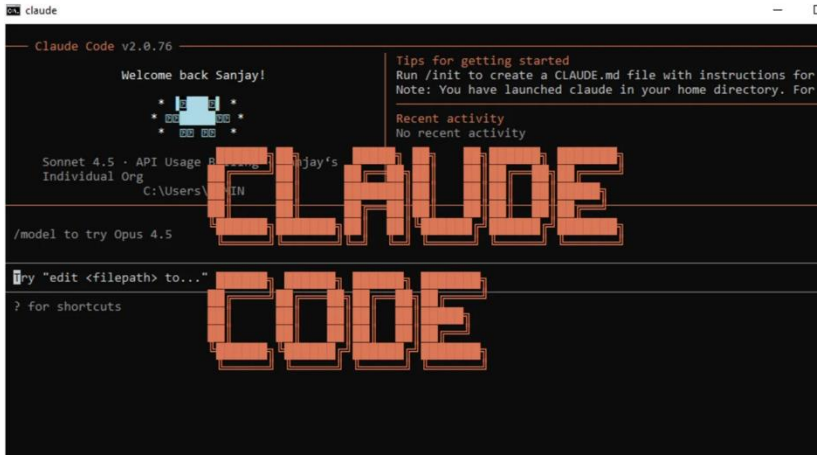

Today’s Lecture

01 What counts as AI

02 How machines learn from data

03 What neural networks do

04 How generative AI creates

05 Why human judgment still matters

A definition

AI builds systems that perform tasks that appear to require intelligent behaviour.

Examples: recognising patterns, making predictions, choosing actions, understanding language and generating new content.

AI is a broad family, not one technology

01 **Artificial intelligence** — The broad goal: systems that perceive, reason, predict, decide or create.

02 **Machine learning** — A major route to AI: learn patterns from examples rather than writing every rule.

03 **Deep learning** — Machine learning with multi-layer neural networks that learn useful representations.

04 **Generative AI** — Most modern systems use deep learning to create new text, images, audio, video or code.

Narrow vs General AI

NARROW AI

Excellent within a defined task

A chess engine, image classifier or route planner can outperform people in a narrow setting while knowing nothing about other tasks.

GENERAL AI

Flexible across unfamiliar tasks

Human-like general intelligence would transfer knowledge broadly, set goals and adapt reliably. It remains an open scientific question.

Programming rules vs learning from examples

RULE-BASED

People write the logic

IF temperature > 38°C, THEN flag a fever. Clear and inspectable—but brittle when the world has too many exceptions.

MACHINE LEARNING

The system fits a pattern

Show many labelled examples, measure errors, and adjust a model so it predicts well on new cases.

By the end, you will be able to explain five ideas

01 What counts as AI

02 How machines learn from data

03 What neural networks do

04 How generative AI creates

05 Why human judgment still matters

Machine learning turns experience into a reusable model

Machine learning turns experience into a reusable model

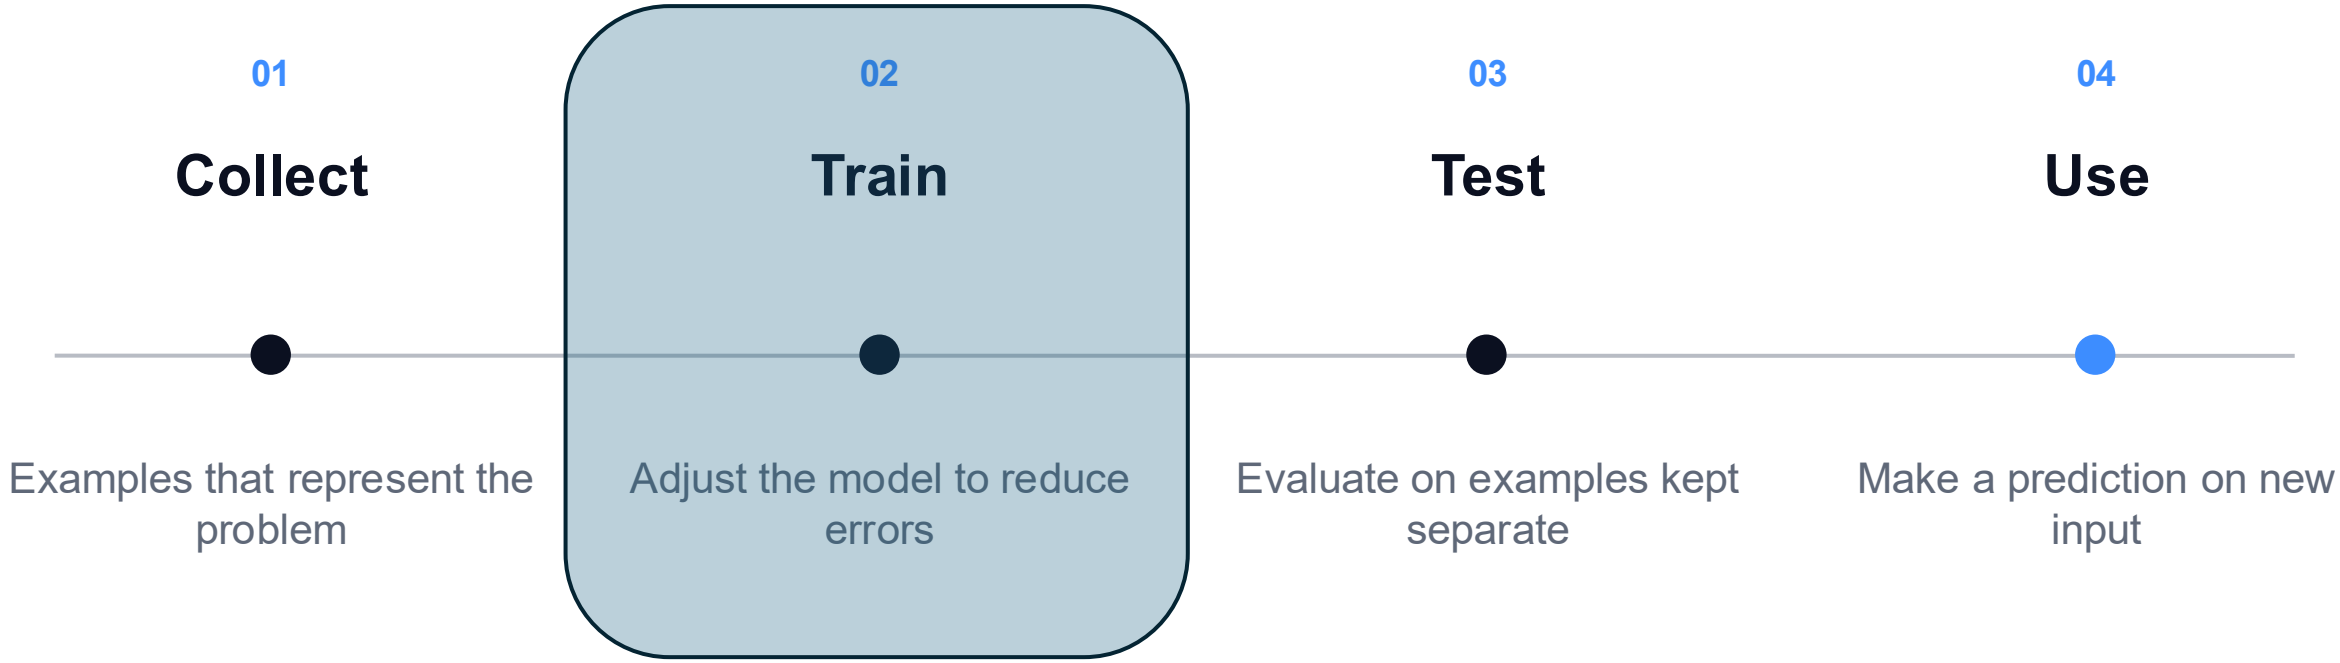

Machine learning turns experience into a reusable model

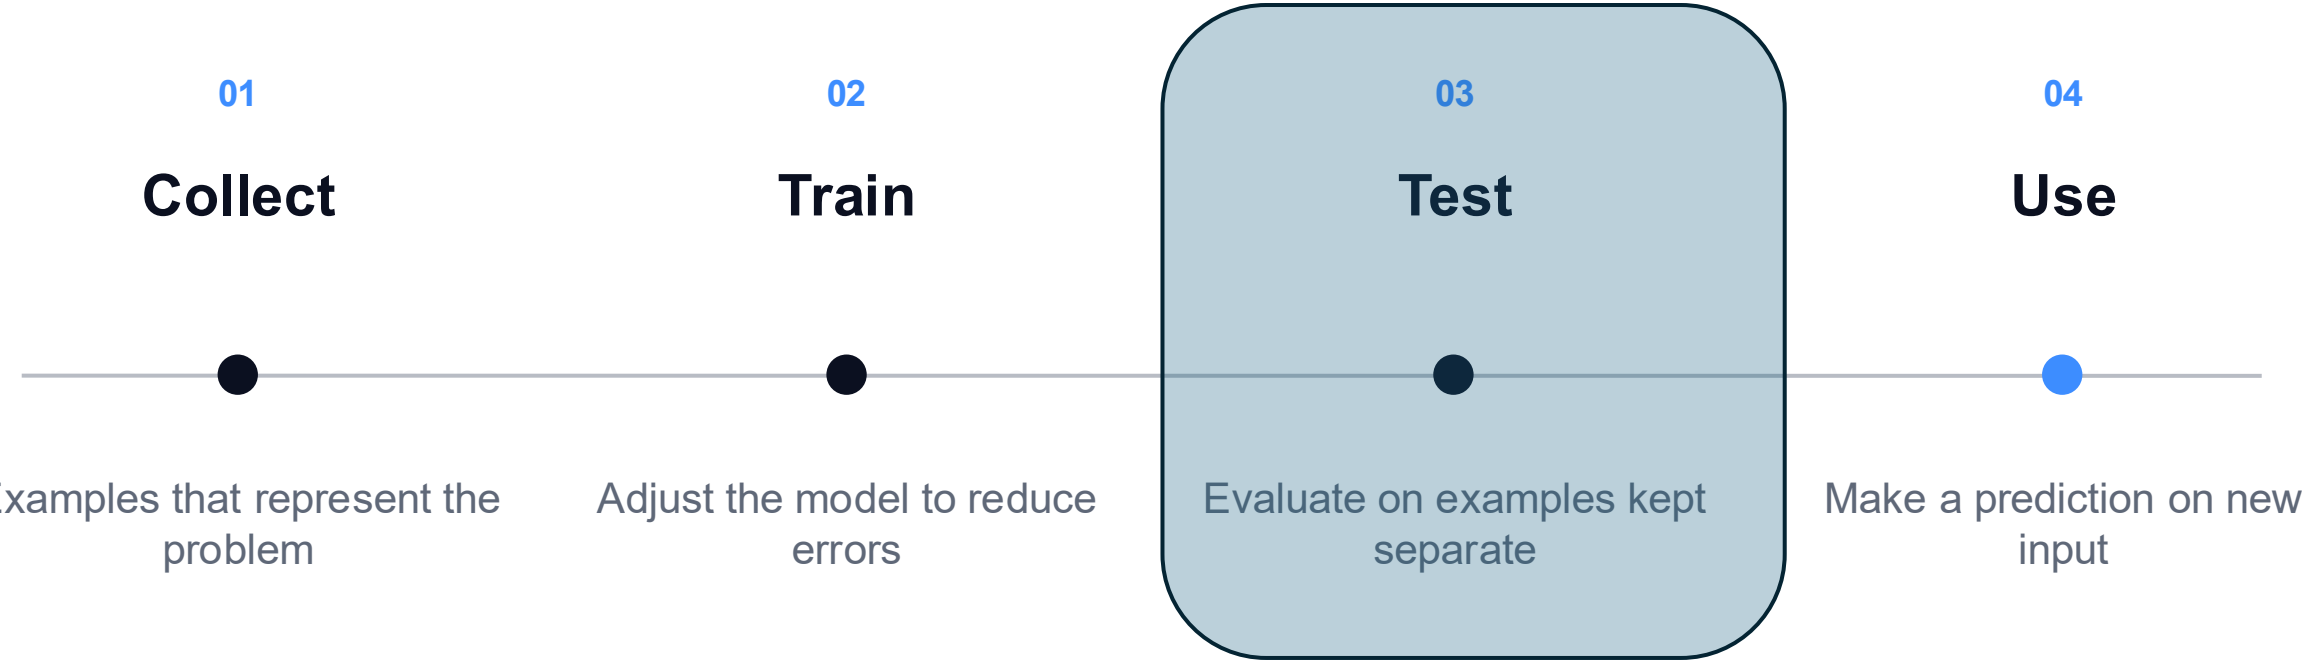

Machine learning turns experience into a reusable model

Good performance depends on every stage—not only the algorithm.

Data can take many forms

Images

Pixels, medical scans, satellite pictures, handwriting, video frames

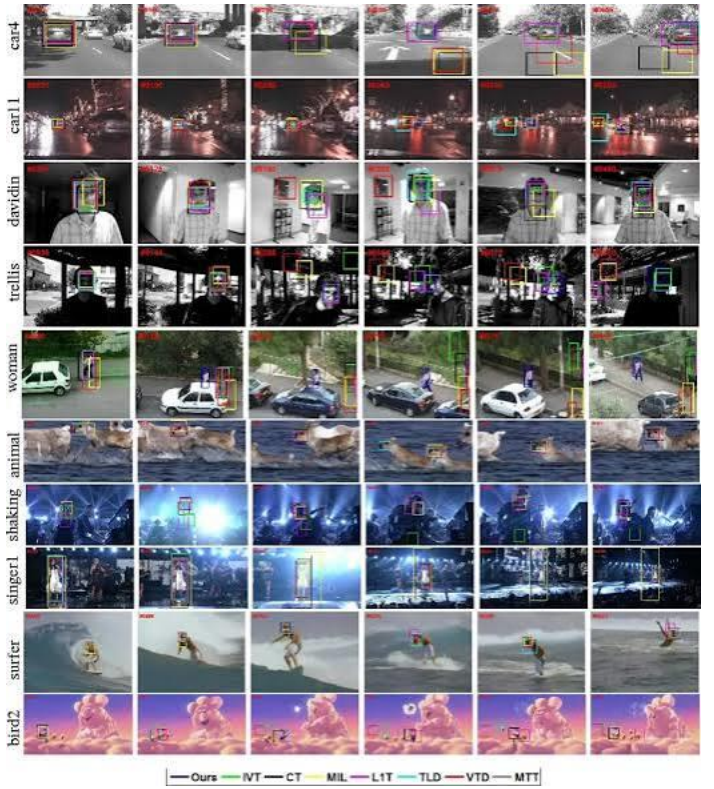

Data can take many forms

Language

Words, tokens, documents, conversations, code, translations

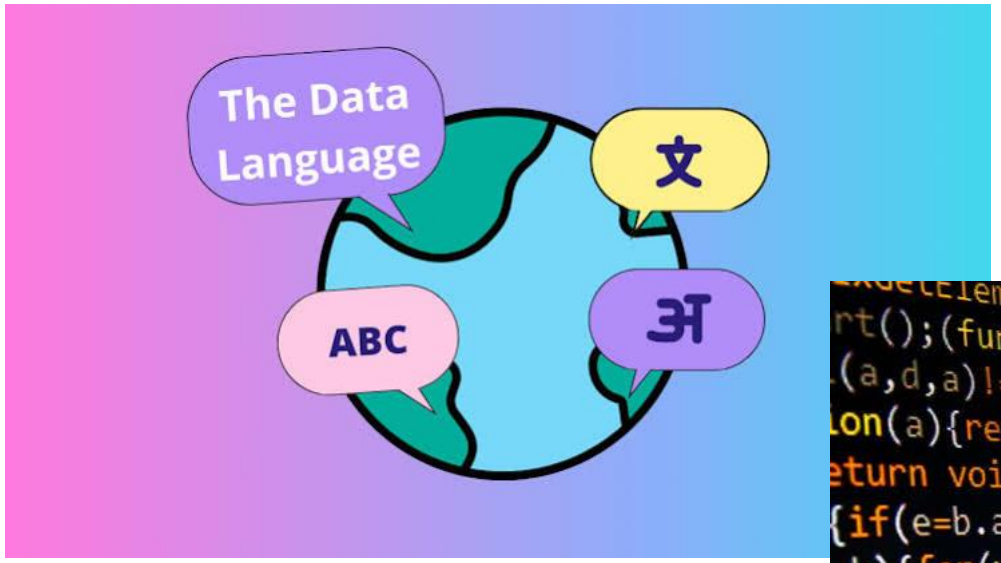

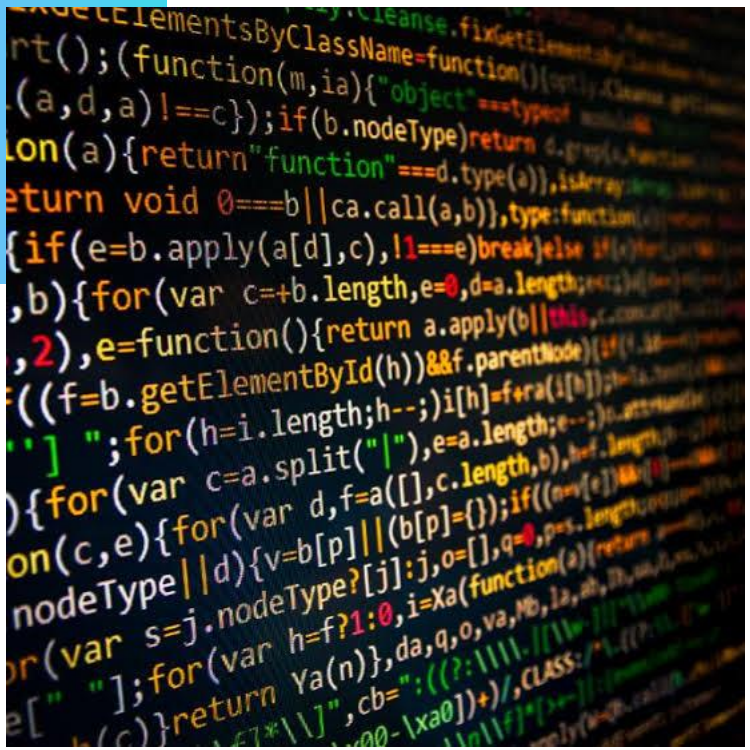

Data can take many forms

Signals

Sound, movement, temperature, clicks, transactions, sensor readings

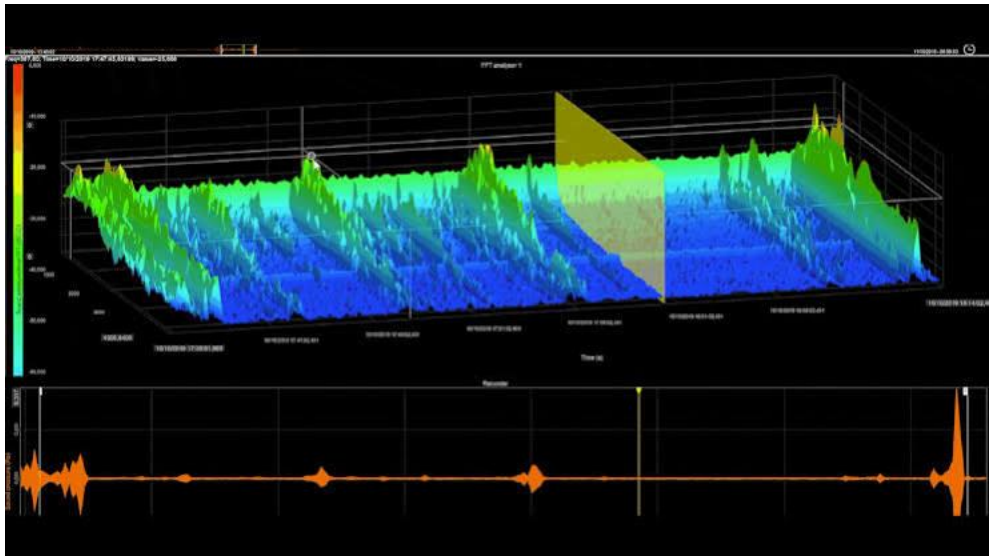

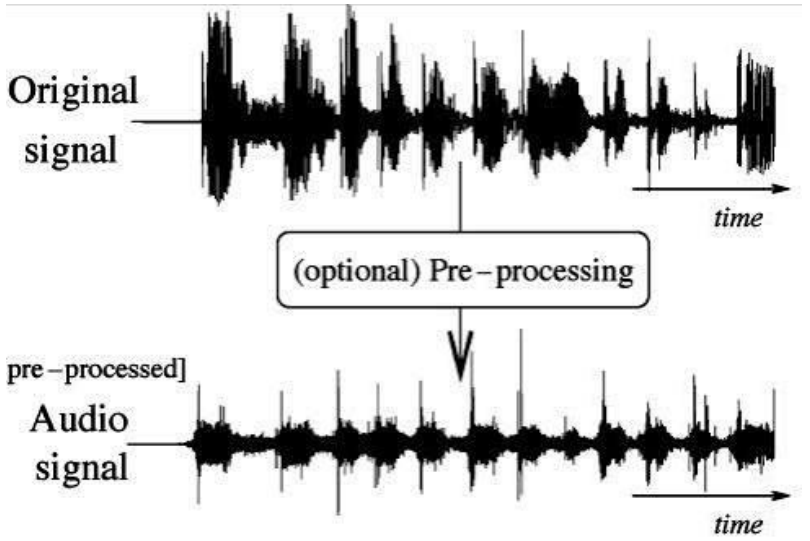

Examples contain inputs—and sometimes desired answers

FEATURES / INPUT

What the model can observe

For a house-price mode e.g:

- **location,**
- **size,**
- **age**
- **number of rooms.**

For an image model**:**

- **pixel values**

Examples contain inputs—and sometimes desired answers

FEATURES / INPUT

What the model can observe

For a house-price mode e.g:

- **location,**
- **size,**
- **age**
- **number of rooms.**

For an image model**:**

- **pixel values**

LABEL / TARGET

What we want it to predict

A category such as ‘spam’, or a number such as a price. Some learning methods use no labels at all.

Training and Inference

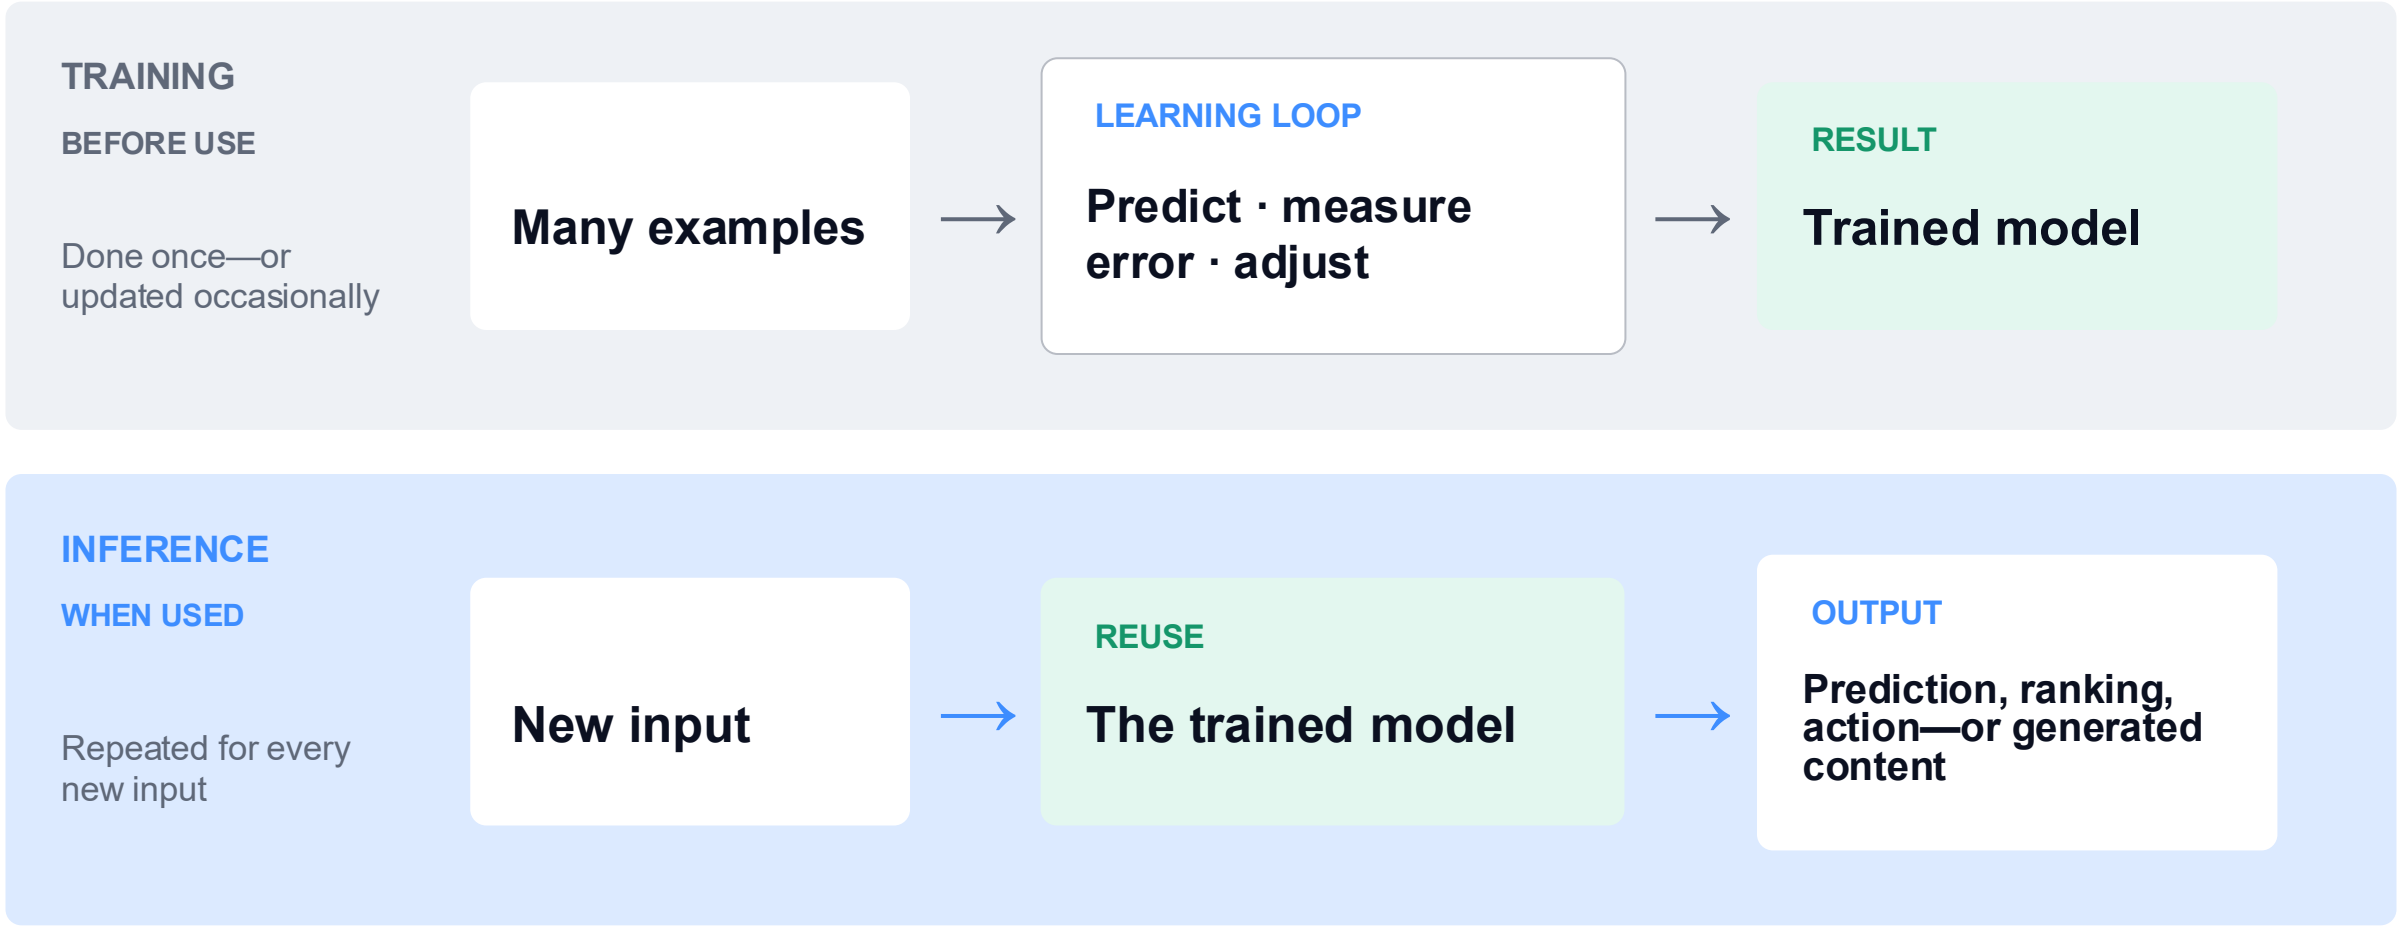

Learning by reducing a measurable error

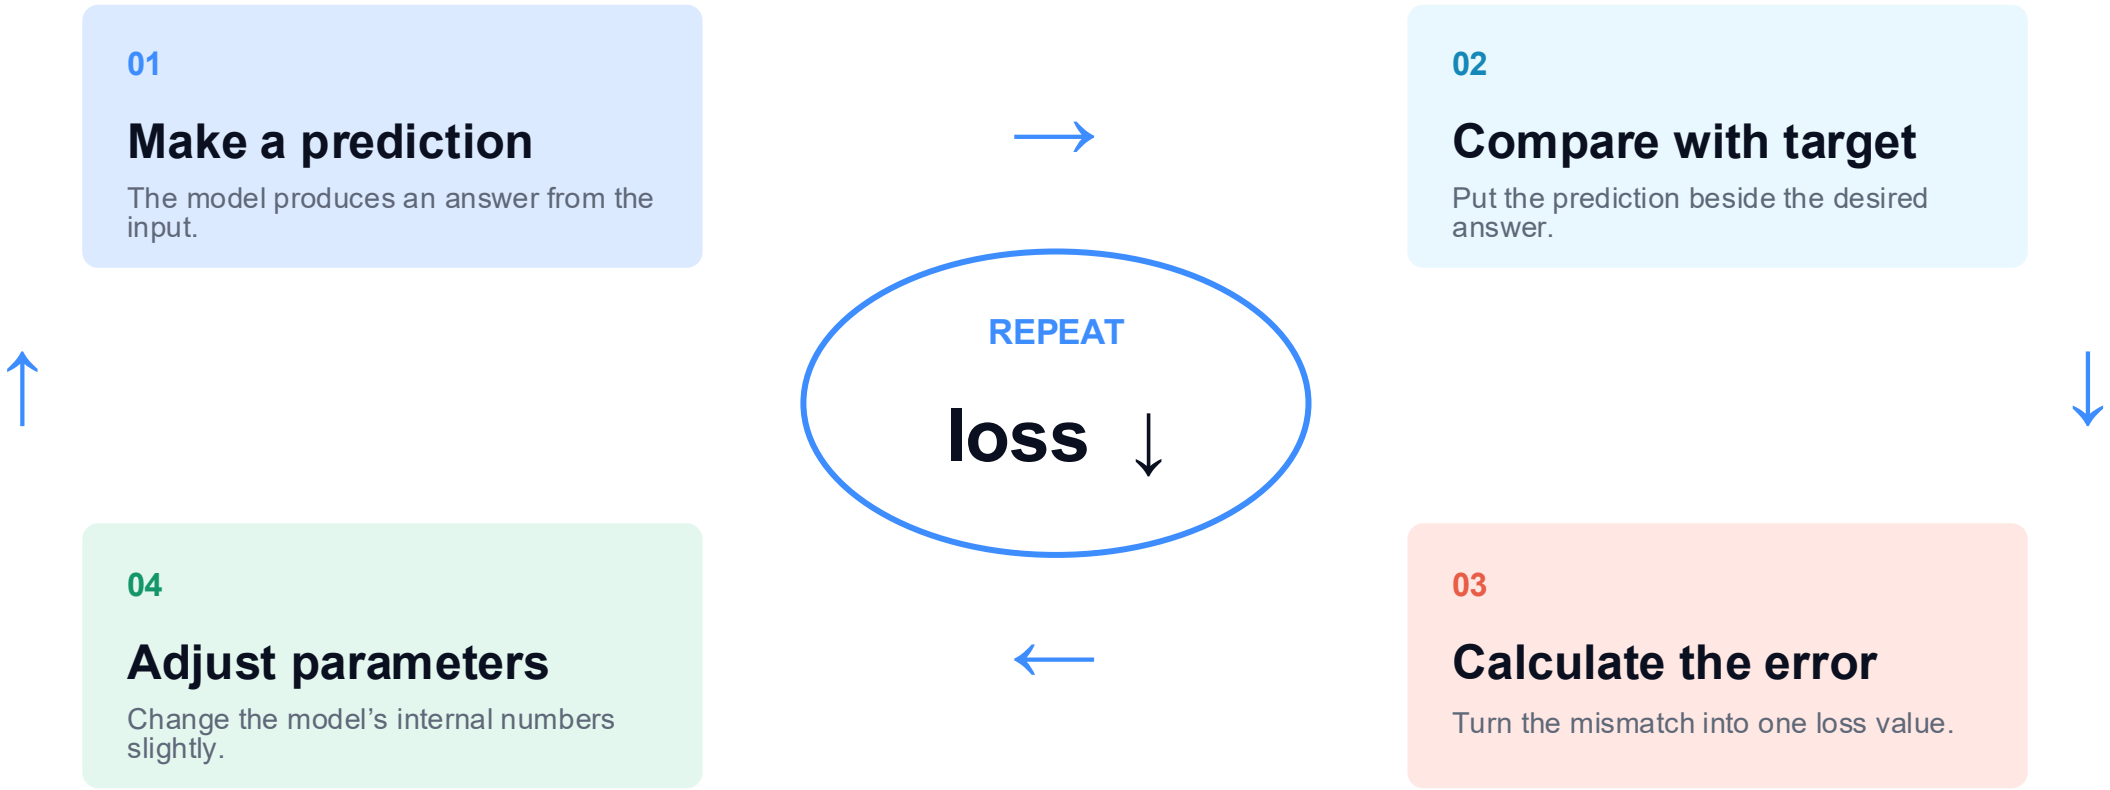

Optimisation changes numbers to reduce loss—it does not mean the model 'wants' to improve.

Three learning settings answer three different questions

Supervised

Learn from input–label pairs.

Unsupervised

Find groups, directions or compressed descriptions.

Reinforcement

Learn through interaction and delayed consequences.

Supervised Learning

Supervised

What answer should this input receive?

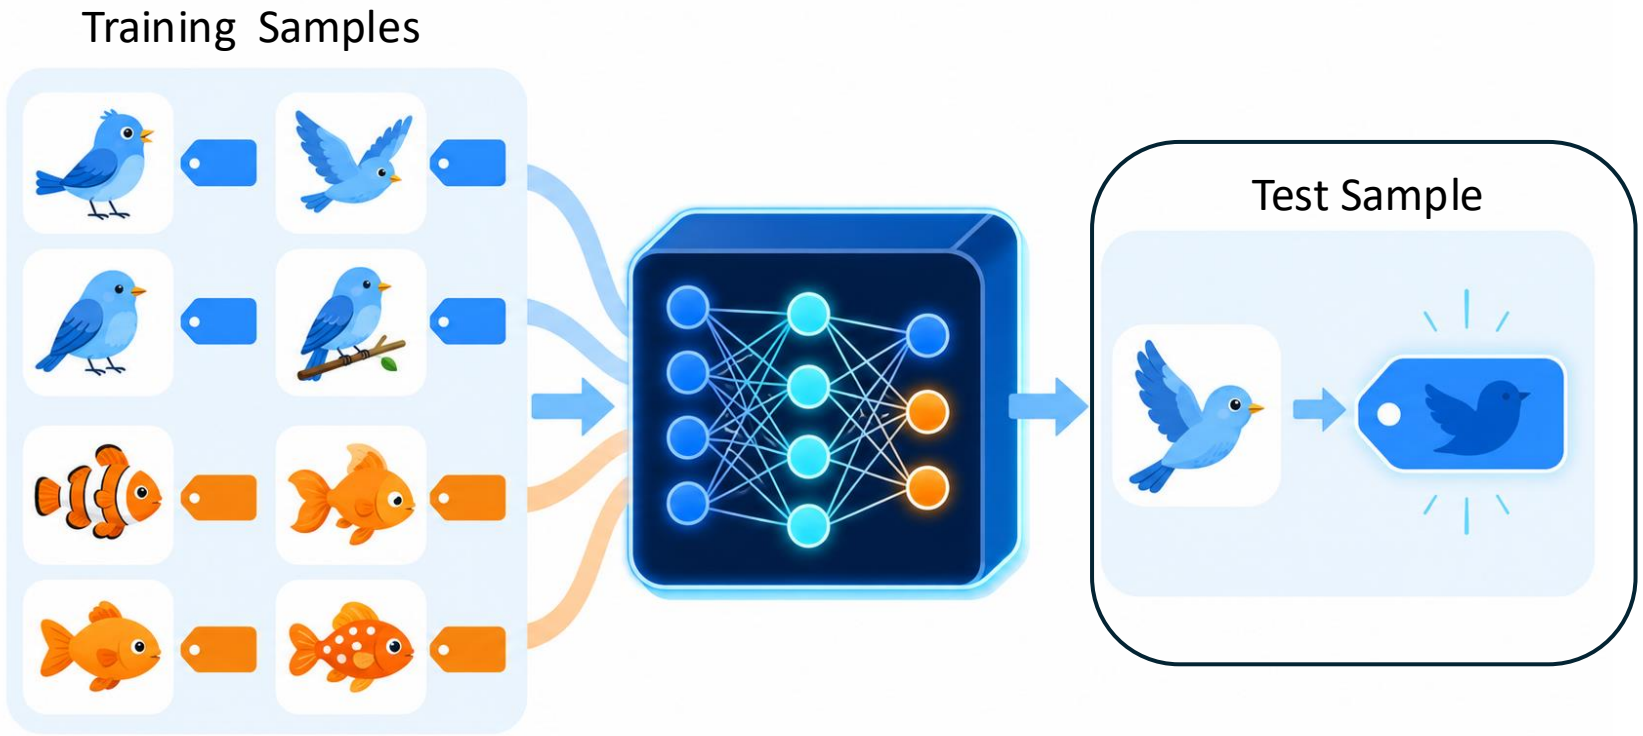

Supervised Learning: Classification

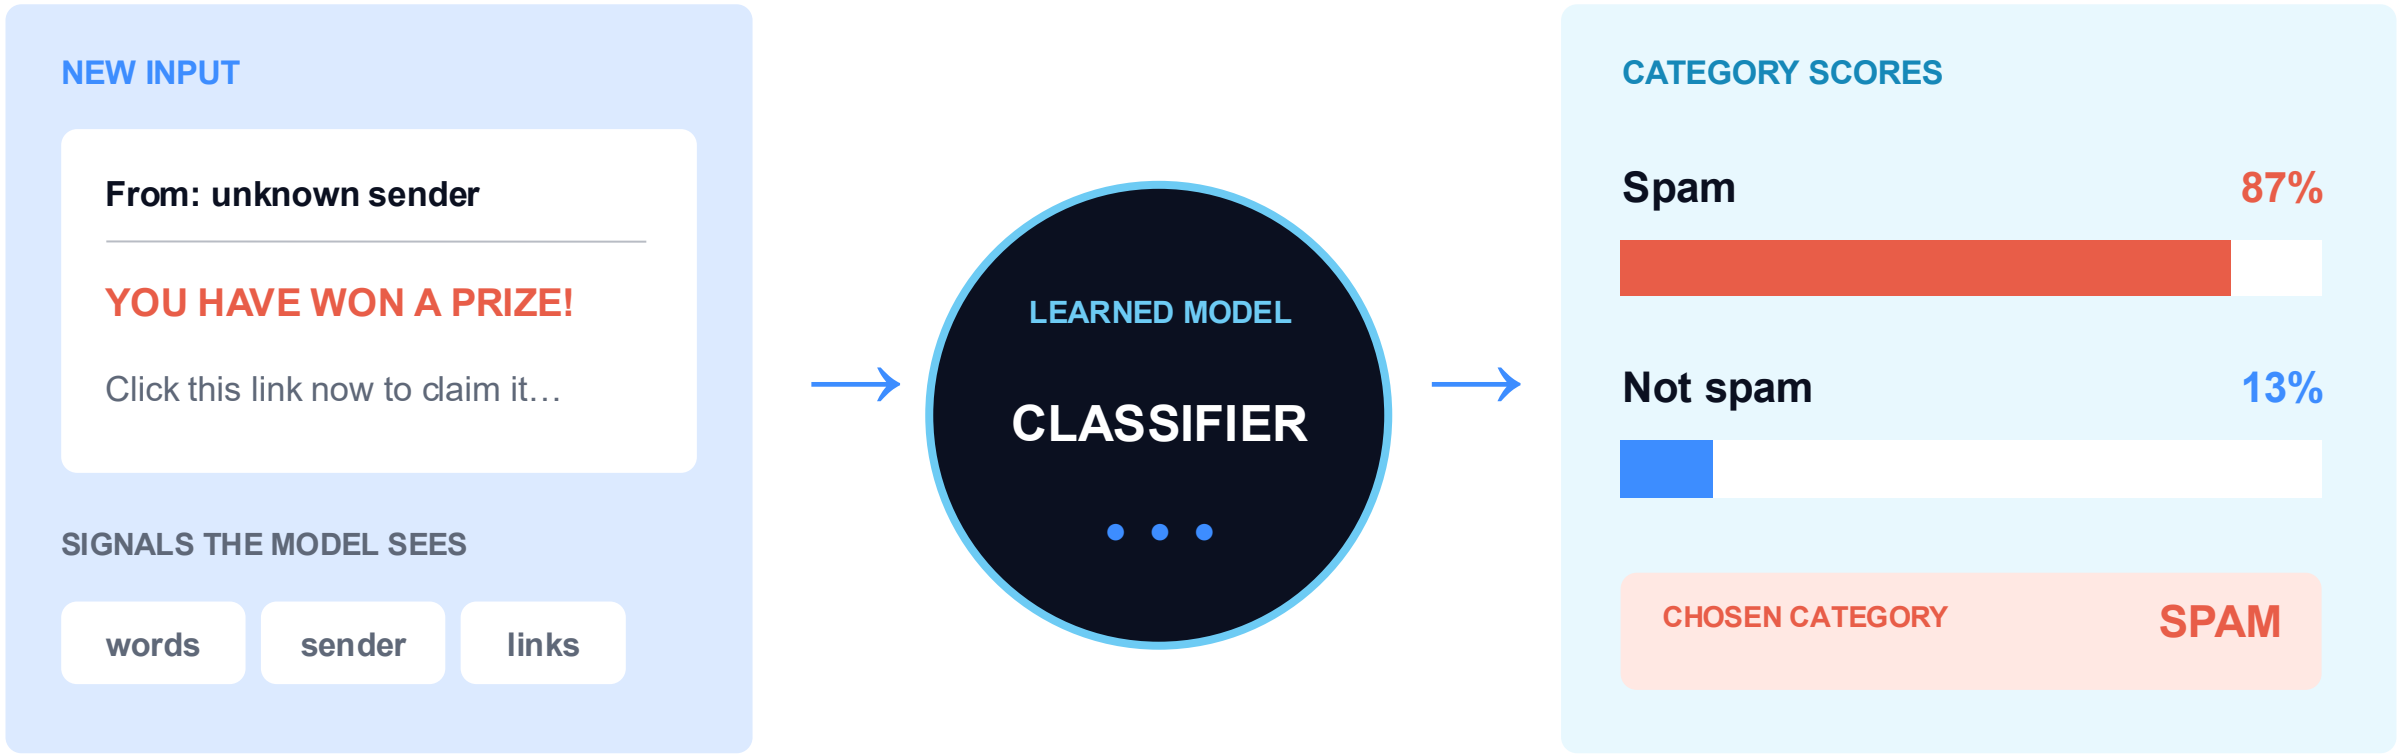

Does it work on tomorrow's emails—not only examples it memorised?

Regression predicts a number

CURRENT CONDITIONS

How long will this journey take?

DISTANCE	18 km
TRAFFIC	High
WEATHER	Rain
TIME	08:30

NUMERICAL PREDICTION

42 minutes

LIKELY RANGE **37–49 min**

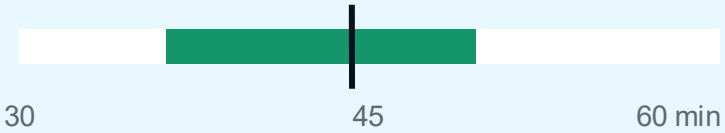

One value on a continuous scale—not a category.

A model must generalise—not simply memorise

The real test is performance on examples it did not see during training.

A student who **memorises** the answer key may fail a new problem. Models can do the same through **overfitting.**

Unsupervised Learning

Unsupervised

What structure exists here?

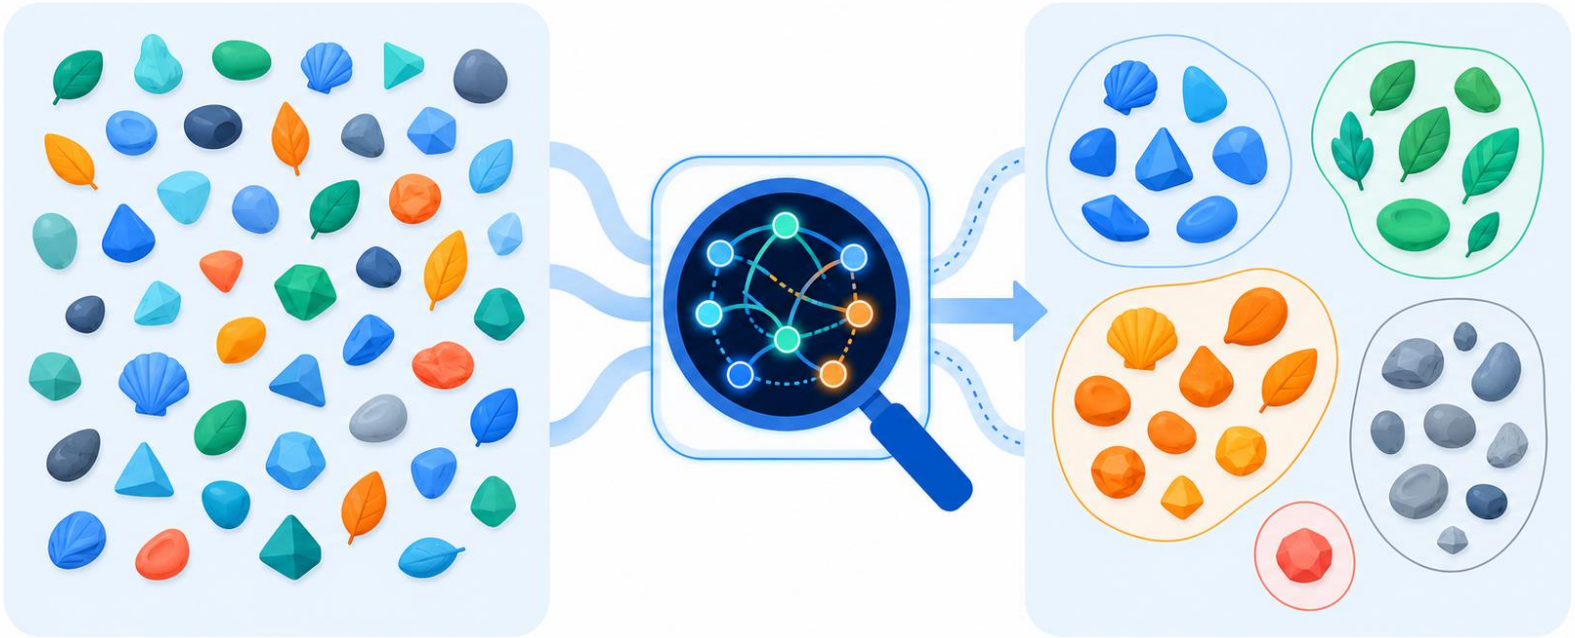

Unsupervised learning finds structure without answer keys/labels

UNLABELLED DATA

No categories are supplied

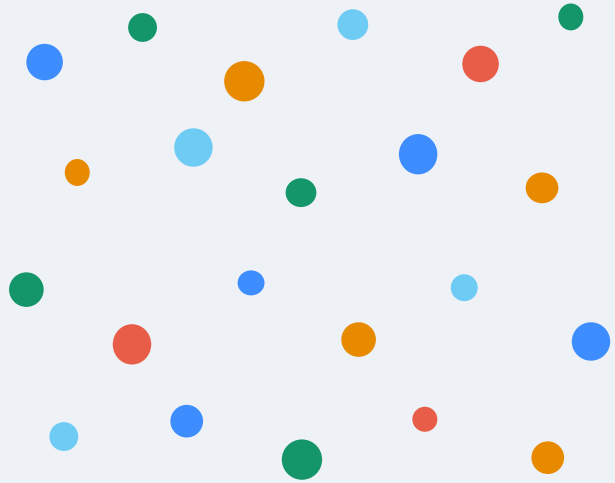

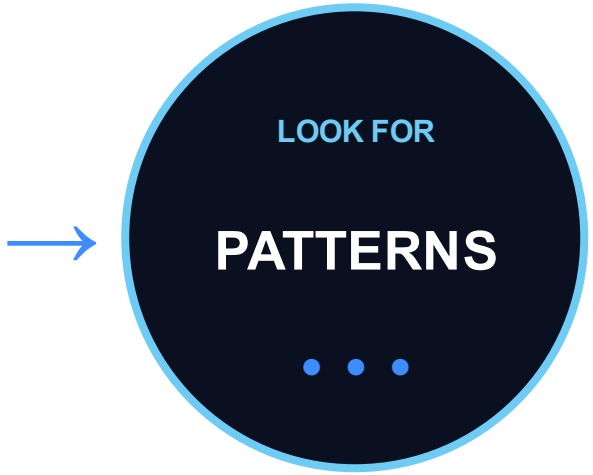

STRUCTURE DISCOVERED

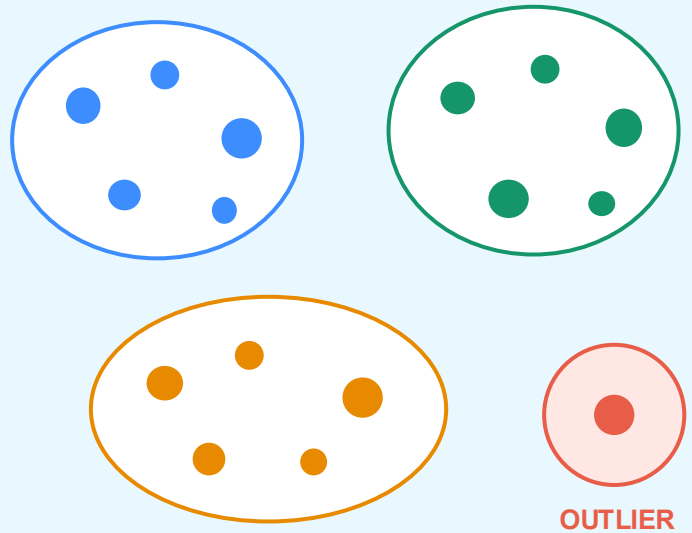

GROUP	similar examples	**COMPRESS**	simpler representations	**DETECT**	unusual examples

Unsupervised learning searches for structure without answer keys

Group

Cluster songs, customers or galaxies by similarity.

Compress

Describe high-dimensional data with fewer meaningful coordinates.

Detect

Notice unusual examples that do not fit common patterns.

Reinforcement Learning

Reinforcement

What action leads to reward?

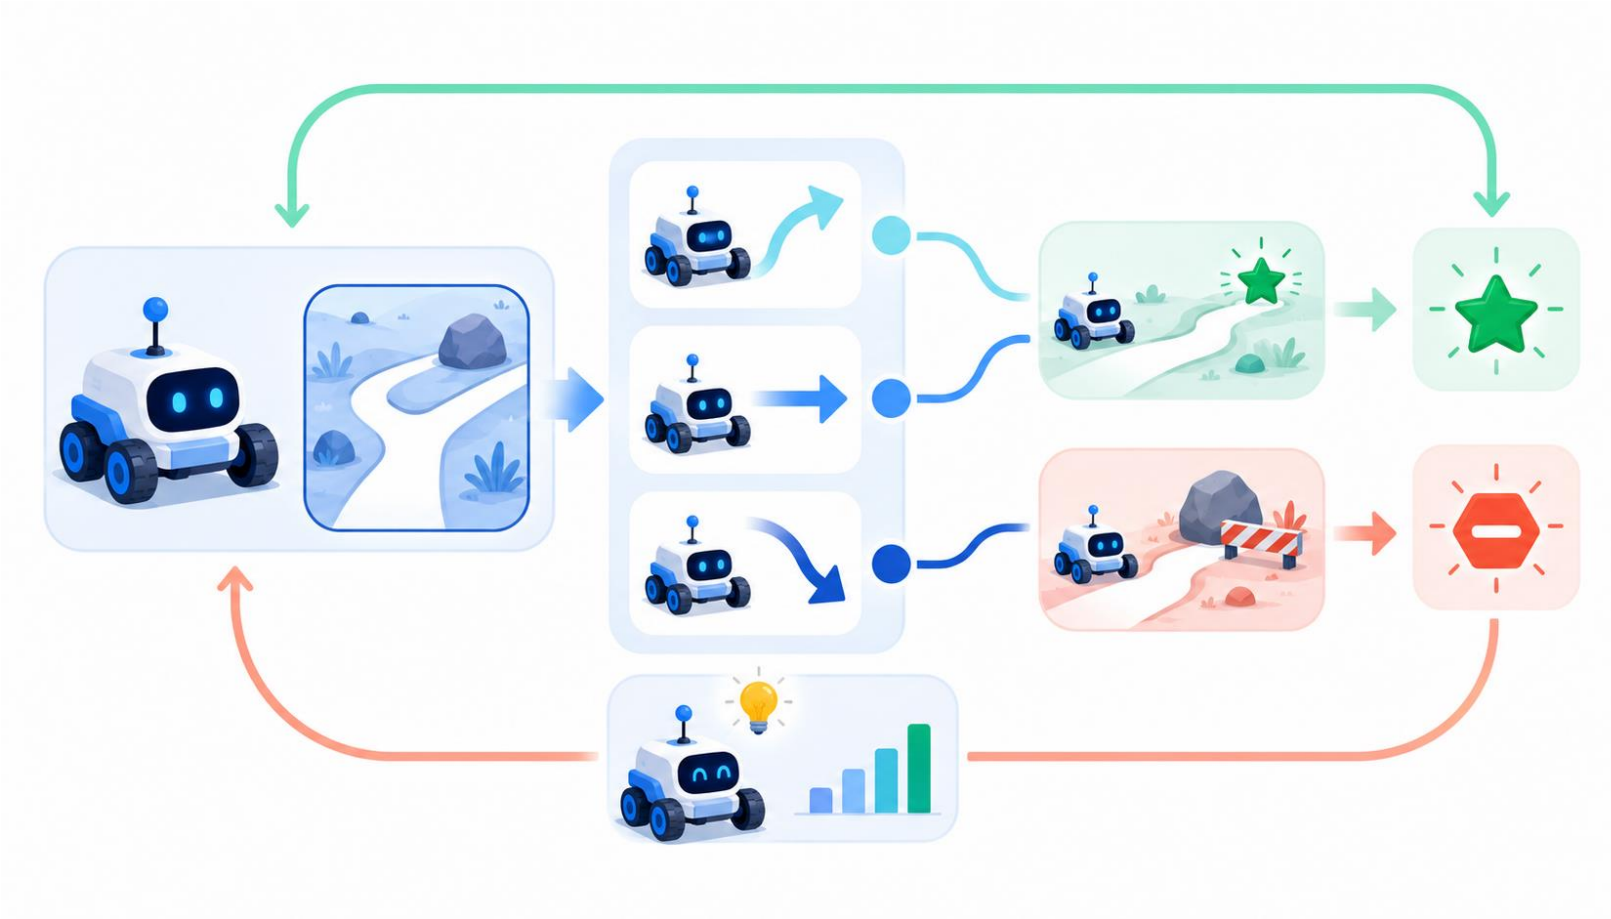

Reinforcement learning connects actions to consequences

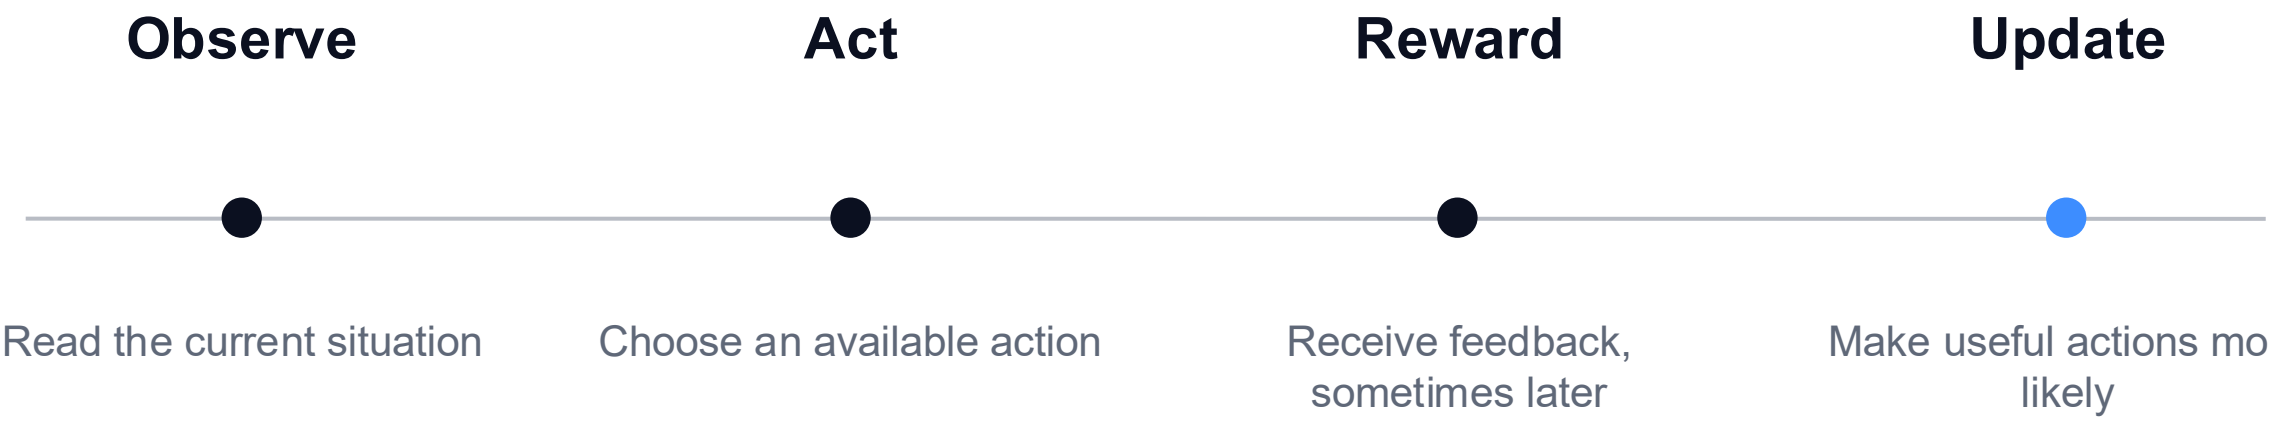

Exploration of new actions can help, but can also be risky in the real world.

DISCUSSION

Which learning setting would you choose?

A: detect spam · B: group news stories · C: teach a game-playing agent

By the end, you will be able to explain five ideas

01 What counts as AI

02 How machines learn from data

03 What neural networks do

04 How generative AI creates

05 Why human judgment still matters

A neural network is a mathematical function

input numbers → repeated transformations → output numbers

Its parameters are adjusted during training. The network's architecture determines how information can flow.

Neural Networks

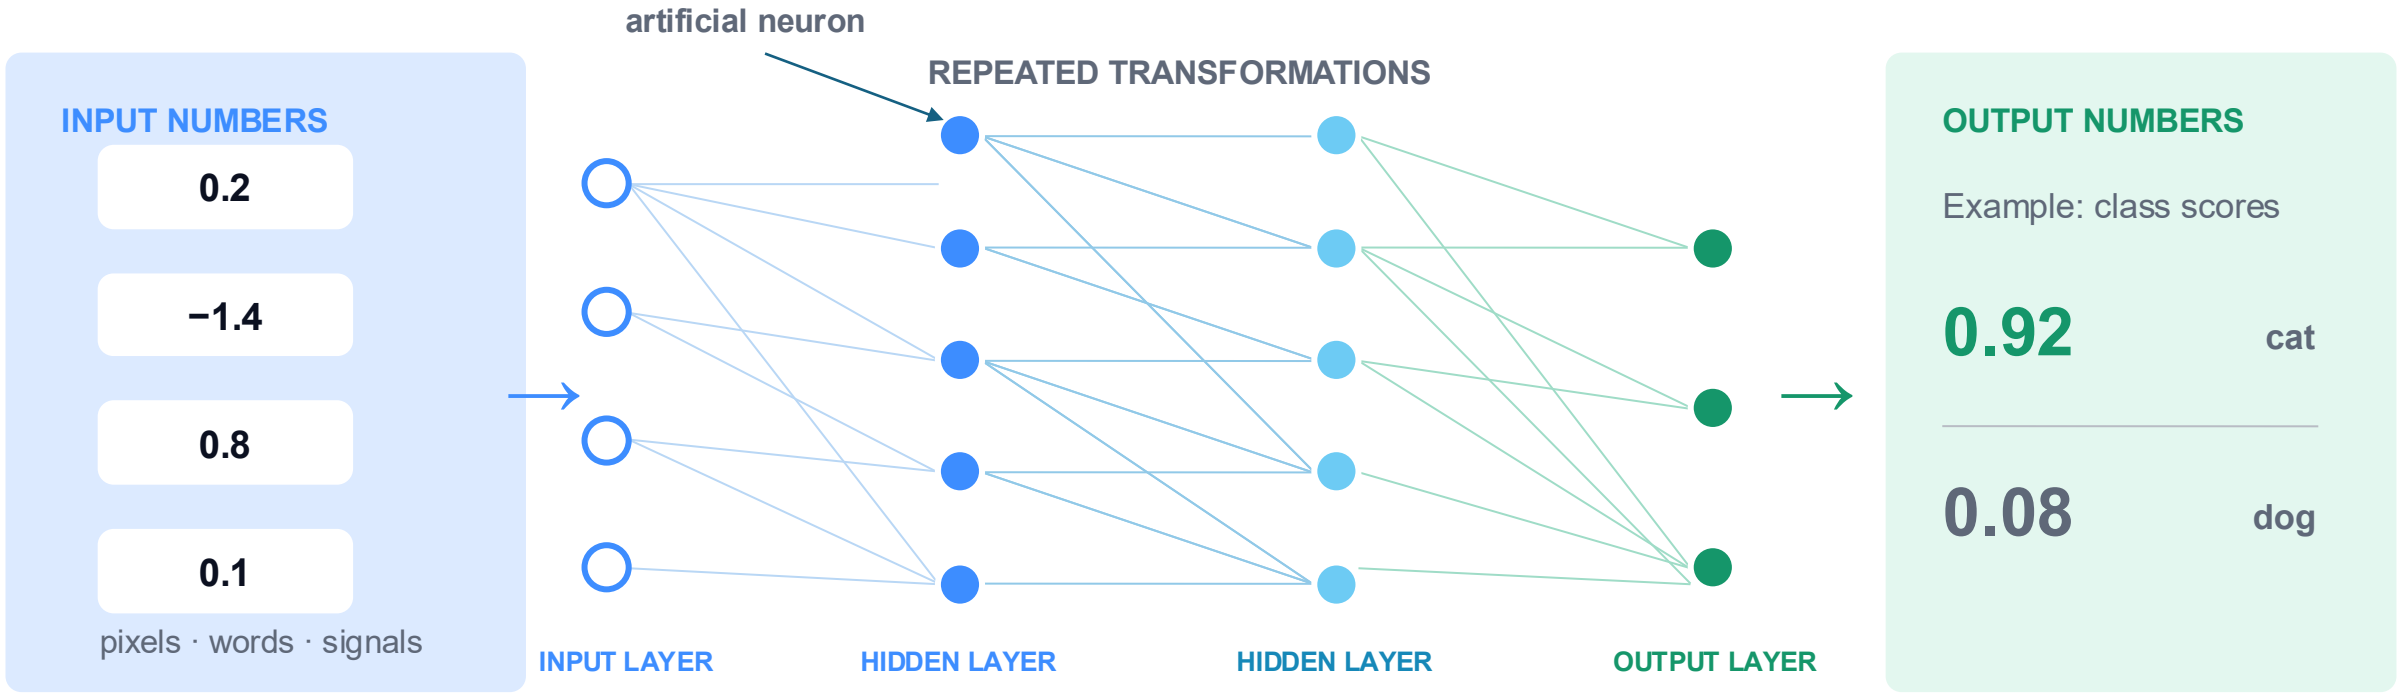

TRAINING adjusts connection strengths

ARCHITECTURE determines how information can flow

Artificial Neurons

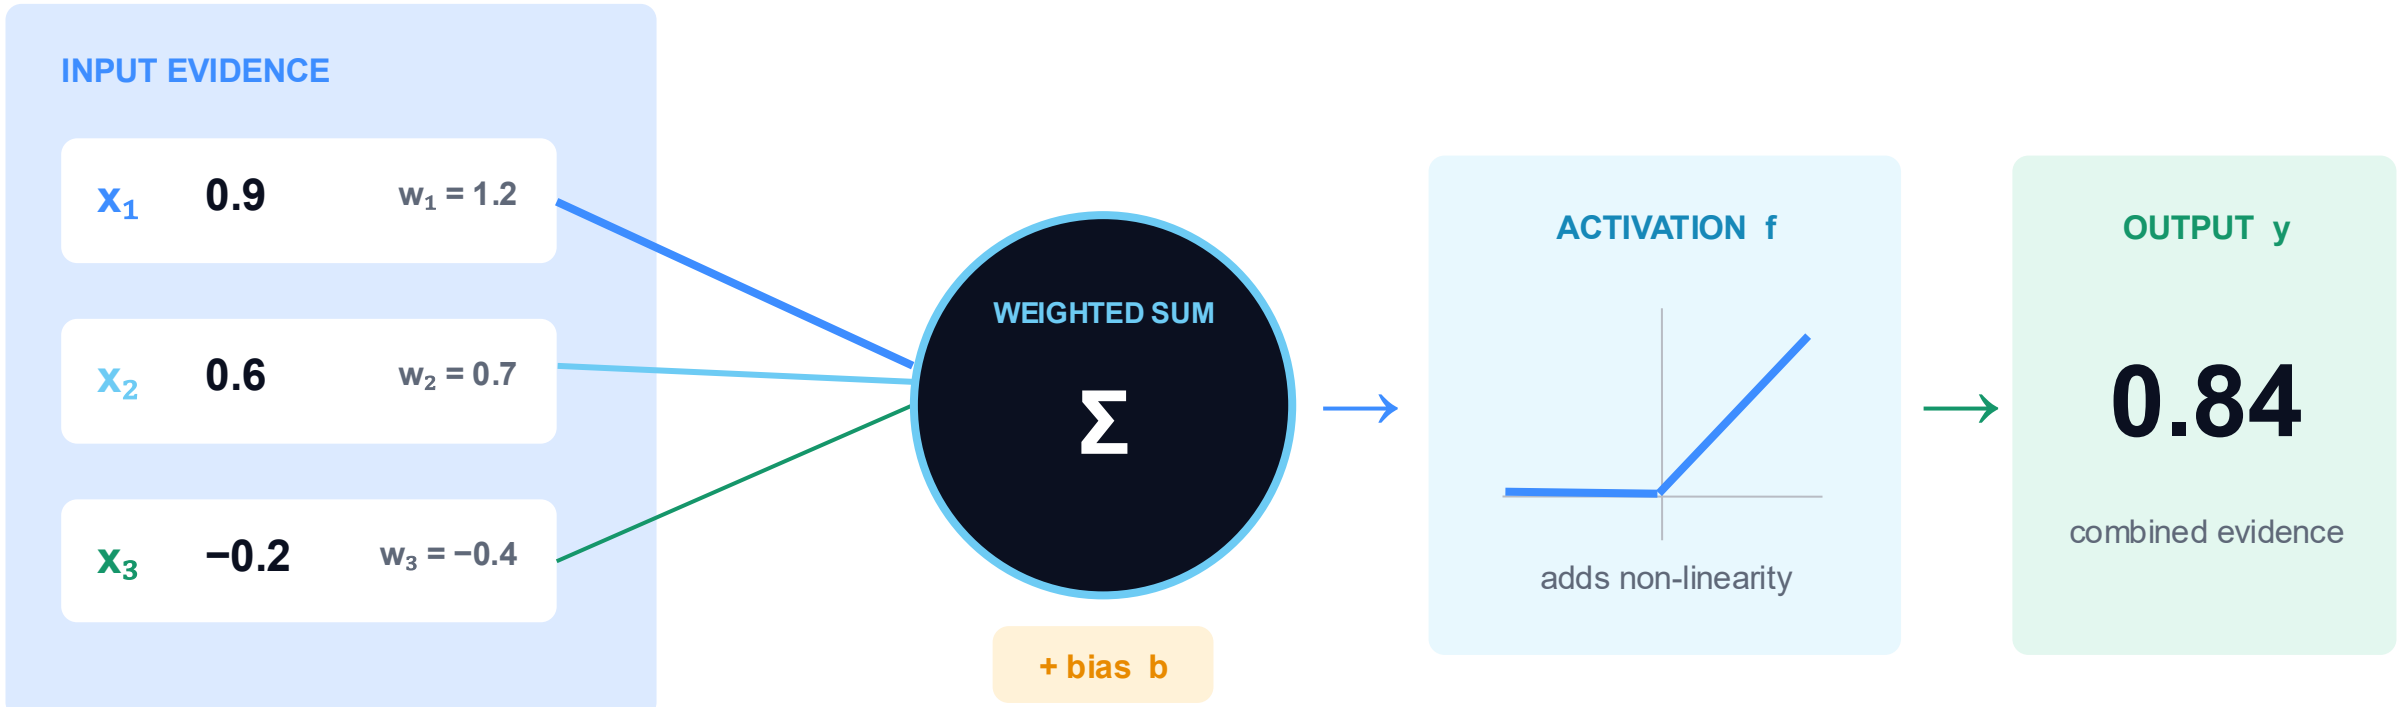

output = activation (weighted inputs + bias)

weights decide which signals matter and in which direction

Layers build increasingly useful representations

01

Pixels

Raw brightness and colour values

02

Edges

Local changes and simple shapes

03

Parts

Eyes, wheels, textures or symbols

04

Concept

A useful class or high-level feature

The hierarchy is learned from the task and data—not hand-labelled at every layer.

Layers build increasingly useful representations

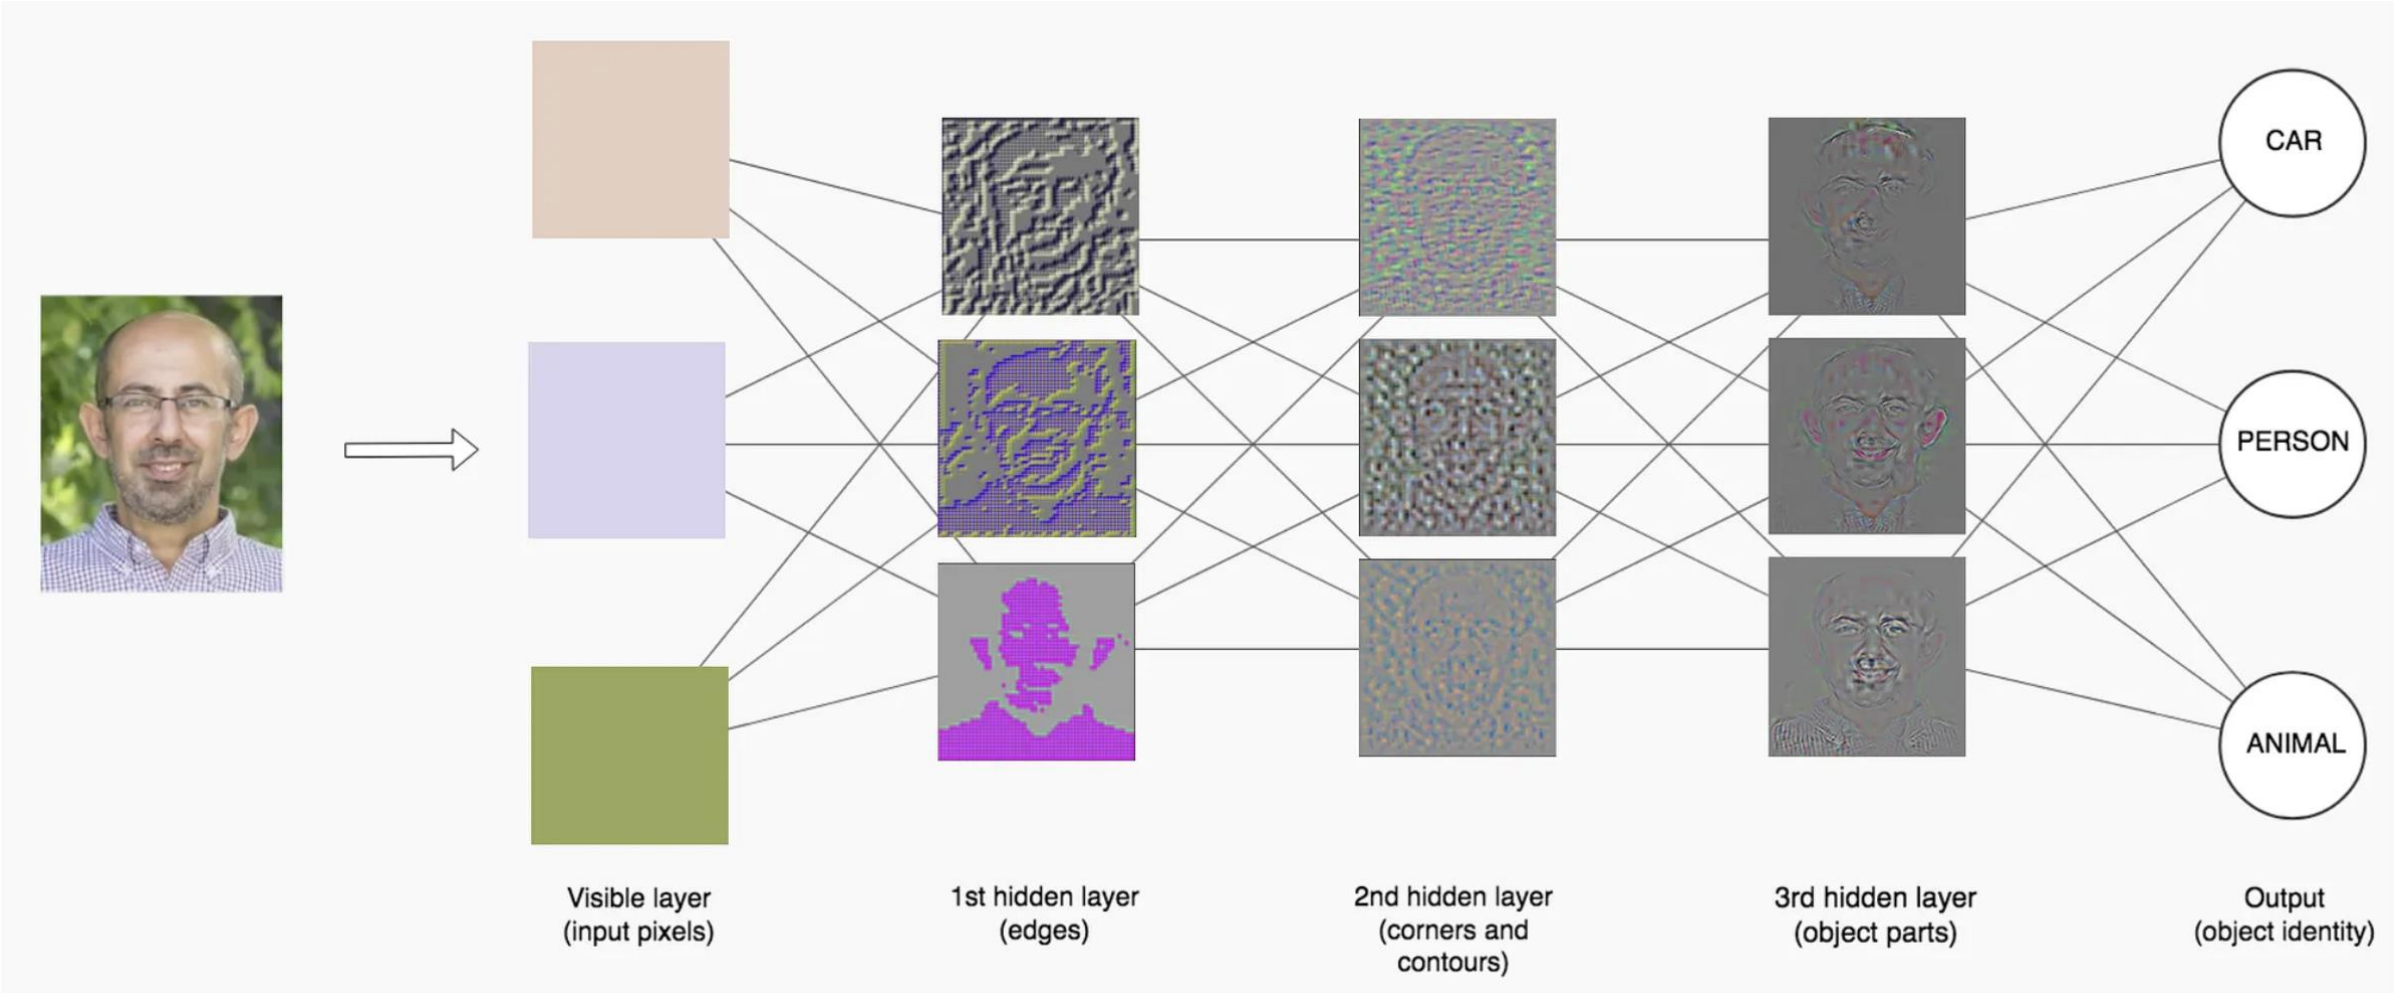

Parameters store learned behaviour; they are not a database of answers

PARAMETERS

Numbers adjusted in training

They shape how signals are transformed. Modern networks may contain millions or billions of these values.

REPRESENTATIONS

Internal descriptions

Useful representations emerge across many parameters and activations.

Parameters store a learned recipe—not a database of answers

TRAINING TUNES PARAMETERS

Think: millions of tiny adjustable controls

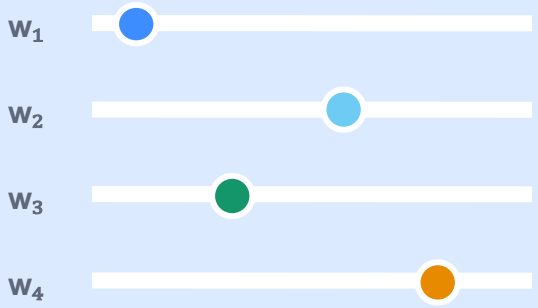

Learning changes these numbers to improve performance.

LEARNED FUNCTION

The settings work together

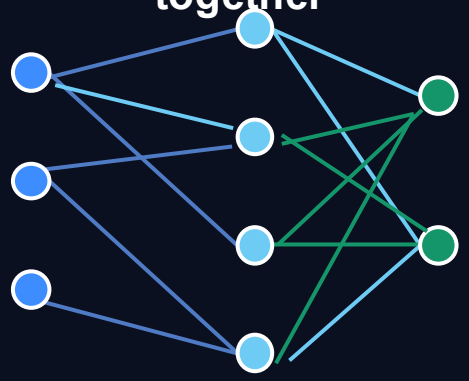

No single parameter contains the answer.

NEW INPUT → ACTIVATION PATTERN

A representation is distributed across many activations

Together, the pattern supports a prediction or action.

By the end, you will be able to explain five ideas

01 What counts as AI

02 How machines learn from data

03 What neural networks do

04 How generative AI creates

05 Why human judgment still matters

Discriminative vs Generative models

DISCRIMINATIVE

Which answer fits this input?

Classify an image, predict a price, estimate a probability or rank alternatives.

GENERATIVE

What new example could fit?

Produce text, pixels, sound or another structured object resembling patterns in training data.

Computer vision

Classification

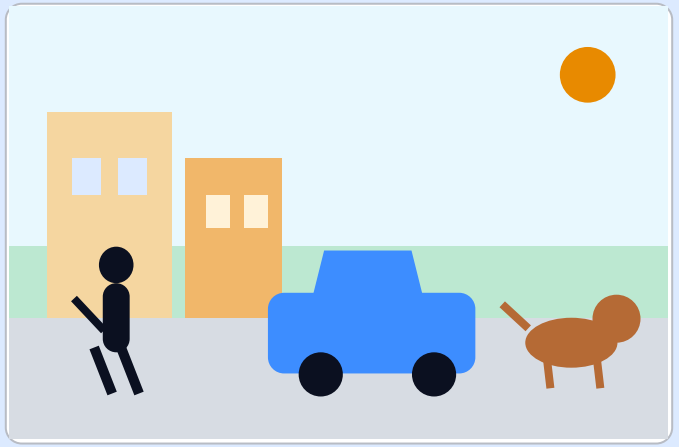

street scene 94%

What is in the image?

Detection

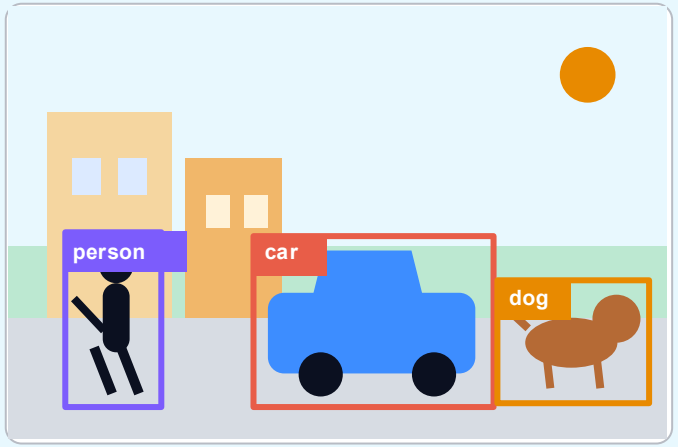

What objects—and where?

Segmentation

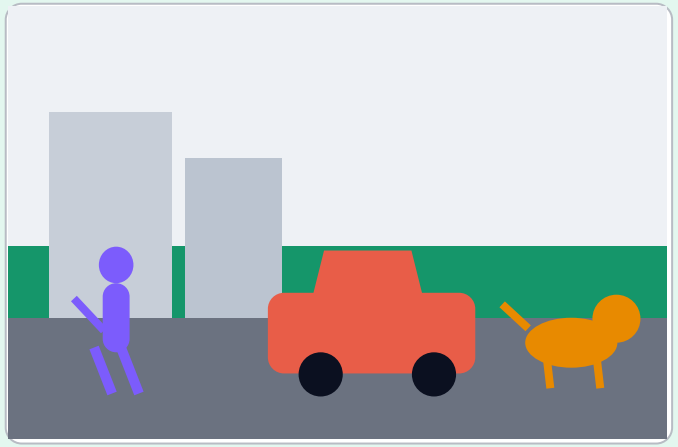

Which pixels belong to each object?

Language technology works with patterns in sequences

Understand

Classify sentiment, extract information, answer questions

Transform

Translate, summarise, rewrite, convert speech and text

Generate

Continue text, follow instructions, write code or dialogue

Language tasks

01 Understand

INPUT SENTENCE

“The film was surprisingly good.”

↓

SENTIMENT POSITIVE

TOPIC film

Extract meaning from a sequence

02 Transform

EN → FR

Where is the station?

↓

Où est la gare ?

Meaning stays; the sequence changes

Create an equivalent sequence

03 Generate

PROMPT

Continue: The robot opened the door and…

↓

GENERATED CONTINUATION

found a garden on the moon.

Predict and produce the next tokens

Generative AI learns a probability distribution

It creates by sampling plausible possibilities.

This makes outputs flexible and varied.
It also means fluent or realistic content can still be incorrect.

Language Models

Language models work with tokens, not ideas written on cards

The robot opened the

Possible next tokens

Autoregressive Language Models: Next token prediction

01

Read context

Convert prior tokens to internal vectors

02

Score

Assign a probability to possible next tokens

03

Choose

Select or sample one token

04

Repeat

Append it and predict again

Long answers emerge one token at a time.

Vision Models: Diffusion Models

Diffusion models learn to reverse a gradual corruption process

01

Clean image

A structured example from data

02

Add noise

Hide more information step by step

03

Learn denoising

Predict how structure can be recovered

04

Generate

Start from noise and iteratively refine

A text prompt can guide which plausible image is formed.

Diffusion models learn to reverse a gradual corruption process

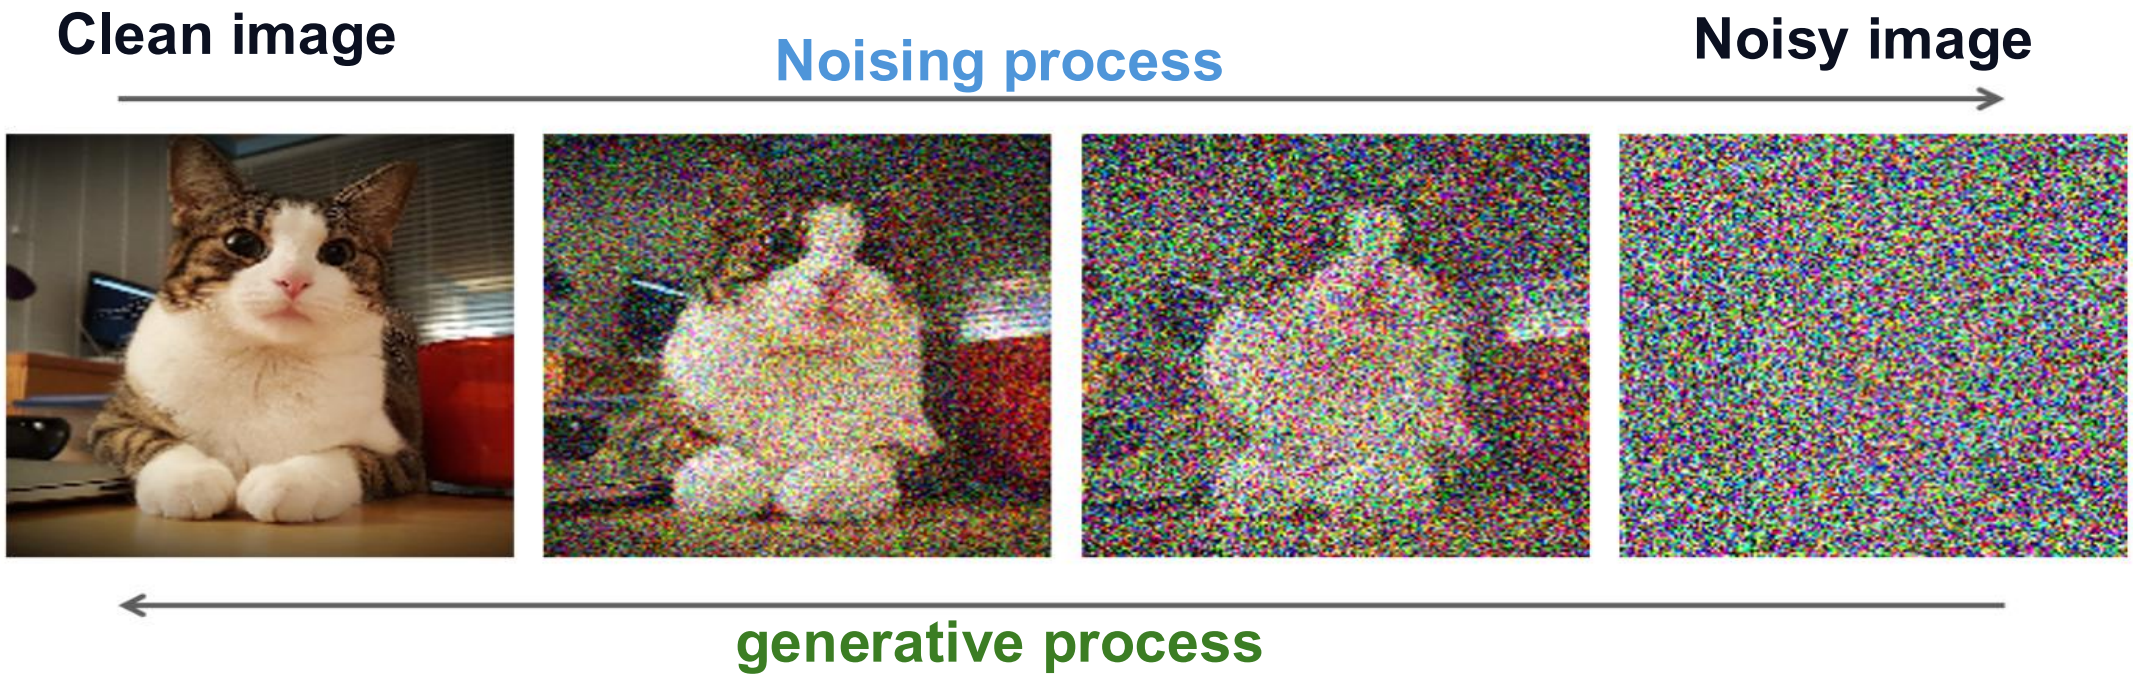

By the end, you will be able to explain five ideas

01 What counts as AI

02 How machines learn from data

03 What neural networks do

04 How generative AI creates

05 Why human judgment still matters

AI systems can amplify what people are able to do

Explore

Generate alternatives, visualise possibilities, test hypotheses

AI systems can amplify what people are able to do

Explore

Generate alternatives, visualise possibilities, test hypotheses

Assist

Summarise information, support accessibility, automate repetitive steps

AI systems can amplify what people are able to do

Explore

Generate alternatives, visualise possibilities, test hypotheses

Assist

Summarise information, support accessibility, automate repetitive steps

Discover

Search large scientific spaces and notice patterns people may miss

But are AI systems always reliable? Some Challenges:

Distribution shift

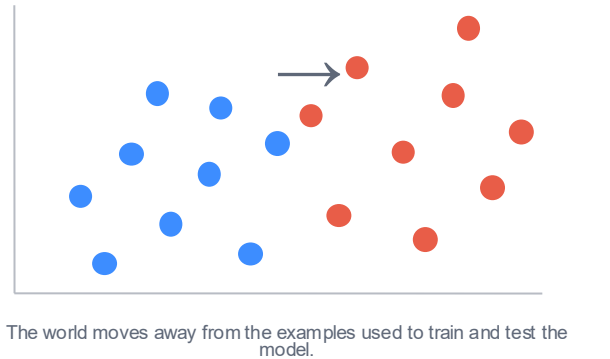

Will yesterday's pattern still work tomorrow?

Uncertainty

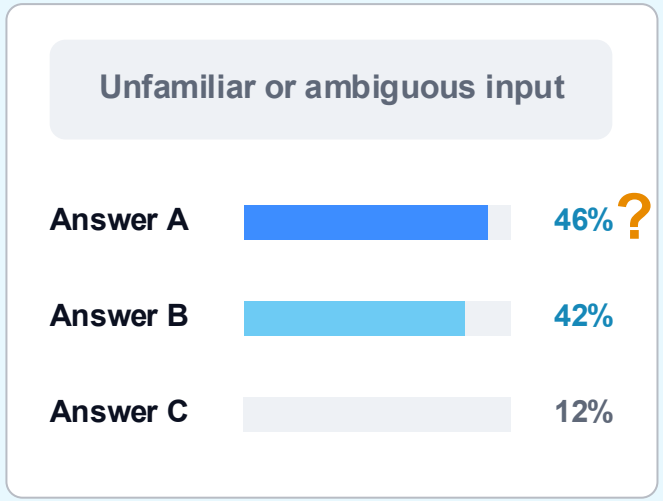

A confident-looking answer may still be uncertain.

Specification

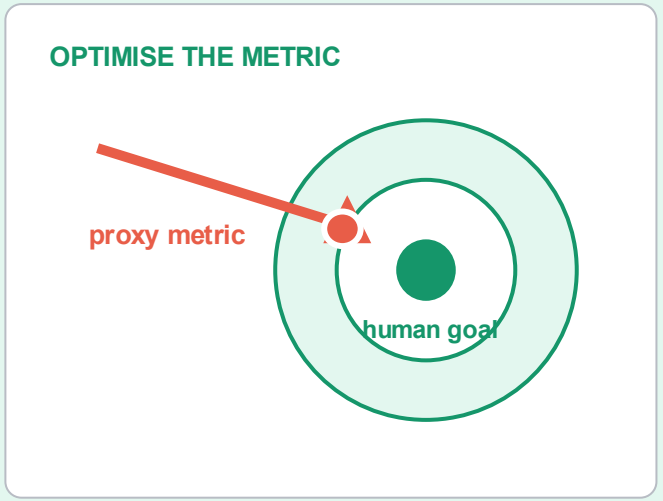

Did we optimise what people actually need?

Bias can enter through data, goals and deployment

WHERE IT ENTERS

Past patterns become training signals

Who is represented? How were labels produced? Which mistakes are counted? What does the objective reward?

WHAT TO DO

Test across people and situations

Measure different errors, involve affected groups, document limits and create routes for human review and appeal.

Fluent language can hide a hallucination

A language model can produce a convincing sentence without having reliable evidence for it.

For factual or high-stakes questions, demand sources, verify them independently, and be willing to say ‘I do not know.’

The higher the stakes, the stronger the safeguards

RISK AND CONSEQUENCE

Medical treatment

Evidence, qualified review, monitoring and accountability

Loan decision

Fairness testing, explanations, appeal and legal compliance

Revision ideas

Check facts, compare sources, protect personal information

Fictional character

Lower stakes—but still consider authorship and harmful content

MORE TESTING AND OVERSIGHT →

Use AI as a critical collaborator

1. Define the task and what a good answer would look like
2. Protect personal, confidential and copyrighted information
3. Check claims against reliable independent evidence
4. Look for missing perspectives, bias and unequal errors
5. Keep a responsible person in charge of important decisions

AI learns patterns; people decide what those patterns are for.