What is AI

This morning your phone unlocked when it saw your face. Your keyboard guessed the next word before you typed it. Your music app put on a song you had never heard, and you liked it. Nobody at Apple, Google or Spotify wrote a rule that says "if it is Marta with her hair up, unlock". So how does the phone know?

That question is the start of this unit. Companies in Mallorca use the same kind of technology to guess how full a hotel will be next month or to sort thousands of support requests. Before you can say where AI helps a company, and where it is only hype, you need to know what is actually inside. Today you will not only read about it: you will train one yourself.

What you will learn

- Define **artificial intelligence** and tell it apart from a program that follows fixed rules
- Place **machine learning** and **deep learning** inside AI, and name the three ways a machine can learn
- Tell **training** from **inference**
- Train an image classifier in the browser with Teachable Machine and explain why it fails
- Avoid the trap of the day: not everything automatic is AI, and a model is only as good as its data

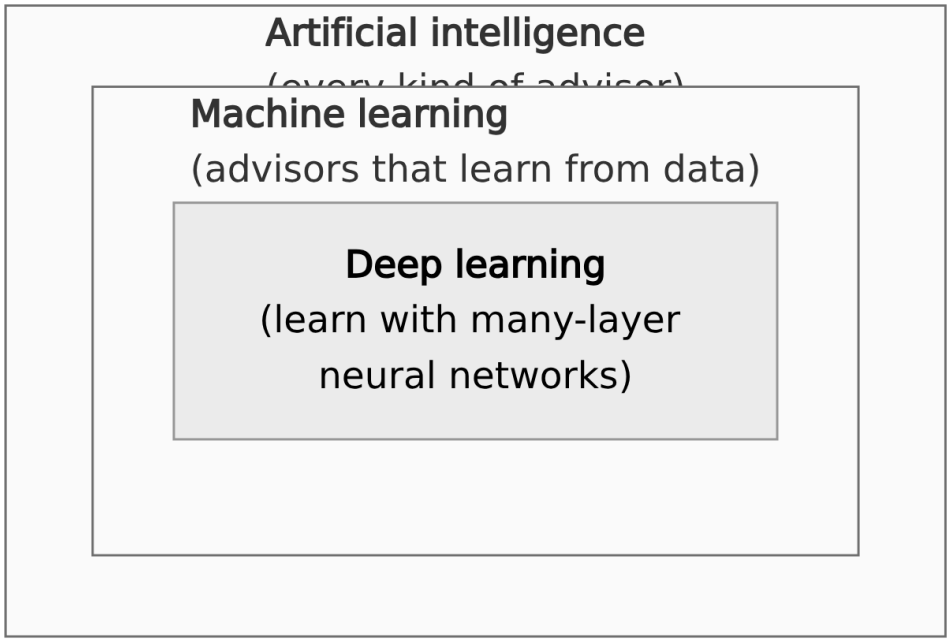
Three circles, one inside the other

> **Artificial intelligence (AI)** is the field of building systems that do tasks we associate with human intelligence: perceiving, reasoning, deciding, generating.

In the city-builder game of this course, AI is the **automatic advisor** that tunes the traffic lights. Some advisors follow rules written by a person; others learn from what they see. That difference gives three circles:

Concept	What it is	Example
Artificial intelligence (AI)	Any system that does an "intelligent" task. Includes rule systems that **do not learn**	A route planner; a medical expert system with rules written by doctors
Machine learning (ML)	Part of AI: the system **learns patterns from data** instead of receiving the rules	A spam filter trained with emails already marked as spam or not
Deep learning (DL)	Part of ML: learning with **neural networks** that have many layers. Good with images, sound and text; needs a lot of data and computing power	Face unlock, automatic translation, chatbots

> A **neural network** is a program made of many small connected units arranged in layers. Each connection has a number (a weight) that gets adjusted while it learns. "Deep" means many layers.



Rules or examples?

This is the key difference, and it is worth drawing:

In classic programming, you write `if temperature < 20 then heating on`. In machine learning, you give thousands of examples with their answers, and the algorithm finds the rules by itself. The result is a **model**: a file full of numbers that captures what it learned.

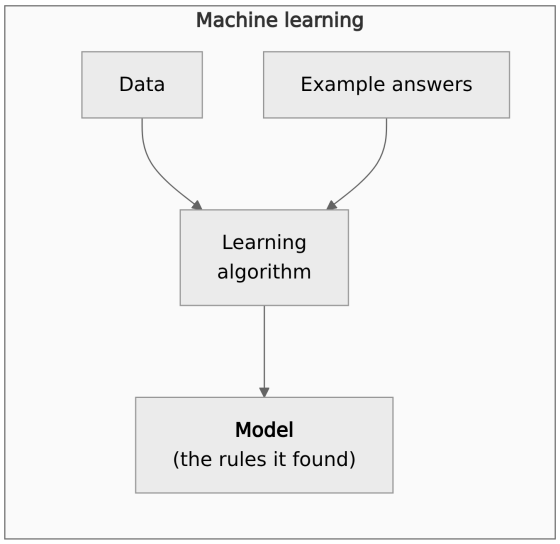
> **Training** is adjusting the model with example data. It is slow and expensive, and it happens a few times. **Inference** is using the trained model on new data to get an answer. It is fast, and it happens millions of times.

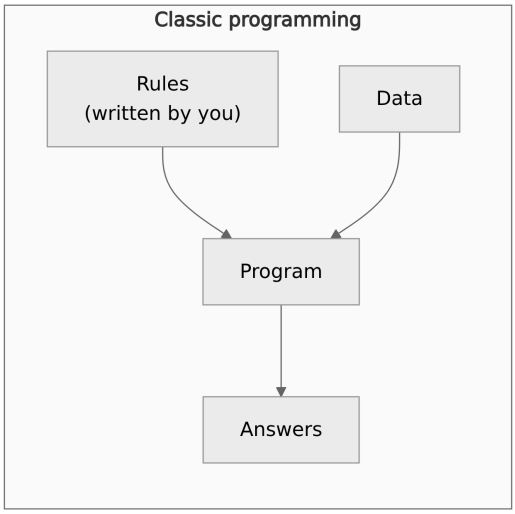
When your phone learned your face, that was training. Every time it unlocks, that is inference.

Three ways to learn

Type	How it learns	Example
Supervised learning	From **labelled** examples: input plus the correct answer. It classifies or predicts a number	Spam yes/no; next month's occupancy of a hotel in Calvià
Unsupervised learning	**No labels**: it looks for groups or structure in the data	Grouping the customers of an online shop into profiles
Reinforcement learning	It tries actions and receives rewards or penalties	A program that learns to play chess; a robot that learns to walk

> **Careful: "everything automatic is AI"**
> A thermostat that turns on the heating below 20 °C is **automation**, not AI: a person wrote the rule and it never learns. A program that calculates VAT is classic programming. Calling these "AI" is the most common mistake in company brochures, and in student projects. Ask: *did it learn from data, or did someone write the rule?*





What AI is not

Two ideas will come back in every lesson of this unit.

A model is only as good as its data. If the examples are few, unbalanced or unfair, the model learns that. This is called **bias**: a model trained mostly with photos of one kind of person recognises that kind of person better. It is not malice; it is the data.

Today's AI is narrow. A model that recognises faces cannot translate a sentence; a translator cannot drive. Each model does one task very well and nothing else. It does not "think" in general.

Why is AI everywhere now, if the ideas are decades old? Three things came together: **lots of data** (everything is digital), **lots of computing power** (graphics cards and the cloud you saw in Unit 3) and **better algorithms** (deep neural networks).

In class: train your first model

You will train an **image classifier** with **Google Teachable Machine**. It runs in the browser, it needs no installation, and you do not need to sign in to train or test. The training happens

on your own computer: the images do not leave it unless you choose to save the project to Google Drive.

The task: a quality check in a factory. Is a part **correct** or **defective**? (A pen with or without its cap, a mug clean or with a paper mark.)

Step	Do this	You should see
1	Open teachablemachine.withgoogle.com → *Get Started* → *Image Project* → *Standard image model*. Name the two classes “correct” and “defective”	Two empty classes
2	Add about **30 images** to each class: hold *Webcam* pressed, or *Upload* the photos your teacher shared on Classroom	A counter of examples in each class
3	Click *Train Model* and do not close the tab	A progress bar, then the *Preview* panel
4	**Inference**: test the model with **10 new cases** (5 per class) that were not used for training. Write down the class it says and its confidence	A table of 10 tests with hits and misses
5	**Break it on purpose**: change the background, the light or the distance, or show an object that is in neither class	Wrong answers given with high confidence
6	**Improve it**: add 20 more varied examples (other backgrounds and light) per class, train again and repeat the 10 tests	More hits, or not: the data decide

Do not record faces or people: this is a practice with objects.

Hand in on Classroom (individual, “S15 · My first model”, by 23:59 the same day): the classes and number of examples with a screenshot; the table of 10 tests and the hit rate before and after improving; the clearest mistake and **why** you think it happened; one thing this model could **never** do even with more training; and two lines on a process of your project company where an image or sound classifier could help. **No generative AI** in this hand-in.

For your project

This lesson feeds Deliverable 4 · AI plan, due **Monday 22 March 2027 at 23:59**:

Today’s work	Becomes	Criterion
Your trained model and its test table	Section D: the working in-browser demo	e
The process where a classifier could help	A first idea for Section A: processes to automate	a
Present and future: why AI is everywhere now, and its limits	Section C: the AI landscape of the sector	c

Key takeaways

- **AI** is the big circle; **machine learning** learns from data; **deep learning** learns with many-layer **neural networks**.
- Classic programming: you write the rules. Machine learning: you give examples and the algorithm finds the rules, stored in a **model**.
- **Training** adjusts the model (slow, rare); **inference** uses it on new data (fast, constant).
- **Supervised** (labelled examples), **unsupervised** (find groups), **reinforcement** (rewards and penalties).
- A model is only as good as its data: **bias** in, bias out. Today’s AI is **narrow**.
- Not everything automatic is AI: a fixed rule is automation.

Quick check

1 Which statement describes the relationship between AI, machine learning and deep learning?

- **A** They are three different technologies with nothing in common
- **B** AI is part of machine learning, which is part of deep learning
- **C** Deep learning is part of machine learning, and machine learning is part of AI
- **D** Machine learning and deep learning are two names for the same thing

2 A thermostat turns on the heating whenever the temperature drops below 20 °C. Is it machine learning?

- **A** Yes, because it decides by itself
- **B** No: a person wrote the rule and it never learns from data
- **C** Yes, it is reinforcement learning
- **D** Only if it is connected to the internet

3 Your Teachable Machine model, already trained, says a new photo is “defective” with 92% confidence. What is happening at that moment?

- **A** Training
- **B** Inference
- **C** Unsupervised learning
- **D** Labelling

4 A shop wants to group its customers into profiles, but nobody has told the system what the profiles are. Which kind of learning fits?

- **A** Supervised
- **B** Unsupervised
- **C** Reinforcement
- **D** None: this cannot be done with AI

5 You trained a classifier with all photos taken on the same white desk. On a wooden desk it fails. What is the most likely reason?

- **A** The model is broken and must be reinstalled
- **B** Deep learning does not work with photos
- **C** The model may have learned the background as well as the object, because all the examples shared it
- **D** The model needs an internet connection to see colours

6 Which two factors explain why deep learning took off in recent years?

- **A** Huge amounts of digital data
- **B** Programmers stopped writing rules
- **C** Much more computing power, such as graphics cards rented in the cloud
- **D** AI became able to think in general, like a person

7 A hotel in Calvià wants to predict how many rooms will be occupied each day next month from its past bookings, flights and holidays. How would you describe the system?

- **A** A fixed rule written by the manager
- **B** Supervised machine learning that predicts a number, trained on historical data
- **C** Reinforcement learning, because it gets a reward
- **D** Not AI, because it uses numbers instead of images

You now know what is inside the advisor and you have trained one. Next, the kind of AI everybody is talking about: models that write. Next: Generative AI and responsible use.

Solutions

1. C · Three circles, one inside the other: **AI** is the largest, **machine learning** is the part that learns from data, and **deep learning** is the part of machine learning that uses neural networks with many layers.

2. B · It is **automation** with a fixed rule. Machine learning needs a model that learned its rules from **data**. Deciding "by itself" is not enough.

3. B · Using a trained model on **new data** to get an answer is **inference**. Training was when you clicked *Train Model* with your labelled examples.

4. B · With **no labels**, the system looks for groups by itself: that is **unsupervised learning**. Supervised learning would need examples already labelled with their profile.

5. C · A model learns whatever pattern separates the examples, including the ones you did not intend. The fix is **more varied data**. A model is only as good as its data.

6. A, C · Data, computing power and better algorithms came together. AI is still **narrow**: it does specific tasks, not general thinking.

7. B · Past days with their **real occupancy** are labelled examples; the model learns from them and **predicts a number** for new days. That is supervised learning, and its quality depends on the history it is trained on.