

MODULE 1 · WEEK 1 · LESSON 3

How Machines Learn: Data, Patterns, Training

The actual mechanics, minus the maths fear

WHAT YOU WILL BE ABLE TO DO

- Explain training, features and labels with a local example
- Distinguish supervised, unsupervised and reinforcement learning
- Describe why 'train vs test' splits prevent self-deception

LESSON CONTENT

Learning = adjusting guesses until errors shrink

Training a model means starting with random guesses, measuring the error against known answers, and nudging internal numbers to reduce that error - millions of times. Nothing mystical: it is guided trial and error at enormous speed.

Features and labels: a spaza-shop example

Suppose you want to predict which customers will buy bread on Fridays. Features are the inputs you measure: day, payday week, weather, past purchases. The label is the answer: bought or not. Supervised learning learns a mapping from features to labels using examples where the answer is already known.

Three families of learning

Supervised learning learns from labelled examples. Unsupervised learning finds structure with no labels - like grouping customers who buy similar baskets. Reinforcement learning learns from reward and penalty, the way game-playing agents learn. And the reason we hold out a test set: a model must be judged on examples it never saw, otherwise it is memorising, not learning.

KEY TERMS

Training	Iteratively adjusting a model to reduce error on known examples.
Features / Labels	The measured inputs / the known answer being predicted.
Supervised learning	Learning a mapping from labelled examples.
Overfitting	Memorising training data so well that new data fails.

HANDS-ON TASK

Teach a human model · Tools: Notebook

1. Pick a simple prediction, e.g. 'will it rain tomorrow'.
2. List 5 features you would use and the label.
3. Explain in 3 sentences how you would check your rule on new days.

KEY TAKEAWAYS

- ✓ Before you go, three takeaways from this lesson.
- ✓ One: Explain training, features and labels with a local example.
- ✓ Two: Distinguish supervised, unsupervised and reinforcement learning.

- ✓ Three: Describe why 'train vs test' splits prevent self-deception.
- ✓ And your hands-on task: teach a human model.
- ✓ Pick a simple prediction, e.g.
- ✓ 'will it rain tomorrow'.
- ✓ When it is done, mark the lesson complete on your dashboard, keep the artifact for your portfolio, and step into the next lesson.
- ✓ I will see you there.

Watch the lesson video on your dashboard, complete the task, then mark the lesson complete. Your artifact belongs in your Portfolio of Evidence.