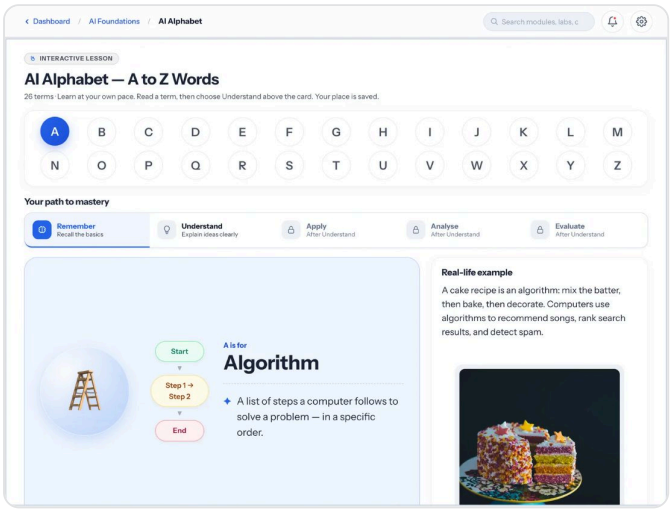
AI Alphabet

A to Z Words

All grades | 60 minutes | Role: **Codebook Keeper** | **Free**, no account needed



The big idea

AI is not one thing but a field of many ideas, and each has a word: algorithm, data, prompt, hallucination. The AI Alphabet takes 26 of them through the same five steps. A learner who can use a word in a real sentence, and spot it used wrongly, can judge claims about AI instead of repeating them.

The learner's job

Learners are the **Codebook Keeper** at Halden Institute · Records Room. People all over town use these words every day. Most of them learned the words from this book. Their task: get every word right. The rest of the building quotes you.

By the end, learners can say

- I can understand 26 foundational AI words in simple language
- I can connect each AI term to an everyday example
- I can use AI vocabulary in short, everyday sentences of their own
- I can distinguish correct from incorrect uses of AI terms in real tools and conversations

Before class

- One computer or tablet for each learner or pair, with EdSutra open. Module 1 is free and needs no account.
- Open the studio once on each device before class, so the Unit still runs if the internet drops.
- Play letter A through all five steps yourself, so you know where learners will click.
- Give each pair four letters to take through all five steps today, for example A, B, H and R.
- Plan return visits: the Unit is finished only when Understand, Apply and Evaluate are passed for all 26 words. Progress is saved.

Words to introduce

Algorithm: A list of steps a computer follows, in order, to solve a problem.

Data: The raw material AI learns from: words, numbers, images, sounds or readings.

Machine learning: AI that gets better at a task by learning from data, not written rules.

Bias: A repeated error that can make AI treat some groups unfairly.

Prompt: The instruction or question you give a generative AI tool.

Hallucination: A fluent, confident answer from an AI that is factually wrong.

Running the lesson 60 minutes

Open 5 min

"Imagine everyone in town learned their AI words from one book. If one meaning in it were wrong, what would happen?"

Before anyone clicks, ask: what do algorithm, bias, hallucination and robot mean? Write the class's meanings on the board and come back to them in Evaluate.

Step	What learners do	Ask the class	Listen for
1. Remember 8 min	For each of their letters, learners read the word card, press Hear it, and read the real-life example beside it. Then they choose Understand on the step bar.	• Which word did you already know, and which one surprised you? • Can you say your word's meaning to your partner without looking?	A short meaning in their own words, tied to an everyday example such as a recipe for algorithm.
2. Understand 10 min	Learners read three definitions and pick the one that best explains the word. A wrong pick shows the key idea, and they can try again.	• Why were the two wrong definitions tempting? • Which wrong definition sounded most like something people say about AI?	The right definition says the same idea in new words; the wrong ones describe a different idea.
3. Apply 10 min	Learners pick the word that fills the blank in a real-life sentence. For A, the phone that finds the fastest route home runs a navigation algorithm.	• What clue in the sentence pointed to the right word? • Where else have you seen this word used this way?	They point to what is happening in the sentence and match it to the word's meaning.
4. Analyze 13 min	Learners read four statements, find the one that does not fit the word, then pick the best reason why. Both parts must be right.	• What do the three statements that fit have in common? • Which reason did you reject, and what was wrong with it?	The shared idea behind the three, such as steps in order for algorithm, and why the odd one lacks it.
5. Evaluate 14 min	Learners judge a claim about the word as True, False or Depends, then pick the best reason. For H: "If an AI sounds fluent and confident, its answer is probably correct."	• Look back at our words on the board. Which meanings would we change now? • When is "Depends" the honest answer, and not a way to dodge?	A verdict with a reason from the word's meaning, such as: sounding sure is not the same as being right.

Steps follow Bloom's revised taxonomy, in the order the Unit presents them.

Deepen the thinking

Critical thinking	Why do short, confident claims about AI, like "more data always makes it better", so often turn out to need "it depends"?
AI thinking	Algorithm, machine learning and neural network are all in the alphabet. How are they related? Is every algorithm machine learning?
Ethics and fairness	Voice recognition often makes more mistakes with some accents than others. Who loses out when that happens, and how could the training data change it?
Connect to life	Which AI word have you heard used wrongly, in an ad, a video or a conversation? What did the speaker really mean?

Misconceptions to listen for

"AI is one thing."
The alphabet shows many separate ideas: data, models, tools, ethics, and you. Each word names a different part of the field.

"These words are too technical for me."
Every word comes with a plain meaning and an everyday example first, like a cake recipe for algorithm. The checks come after.

"If an AI sounds confident, it is probably right."
Hallucination's claim says exactly this, and it is False. Fluent answers can still be wrong; checking them is your job (Y is for You).

In real life the card learners see at the end of the Unit

- News stories, ads and school notices now use words like algorithm, chatbot and deepfake.
- A homework helper or chatbot can give a confident answer that is still wrong: a hallucination.
- A music app picks what plays next with an algorithm: a set of steps that follows your past choices.

Try it this week. Pick three AI words you hear this week on TV, in class or online. Explain each one to someone at home in one plain sentence.

Watch out for. A word can sound smart and still be used wrong. AI isn't always right: chatbots can be confidently wrong, and not every robot uses AI.

Talk about it. If an AI tool gets something wrong, who should check it: the tool, the company, or you?

Exit ticket

1. In one plain sentence, what is an algorithm? Give an everyday example.

   _____

2. Pick one claim you judged. Was it True, False or Depends, and why?

   _____

3. Use one of your words correctly in a sentence about your own life.

   _____

Going further

Module 1's Computer Vision Lab Examples are free, and the camera is optional: learners can spot AI Alphabet words at work there, such as data and algorithm. For homework, use the "Try it this week" activity from the In real life card.

EdSutra

Teaching guide

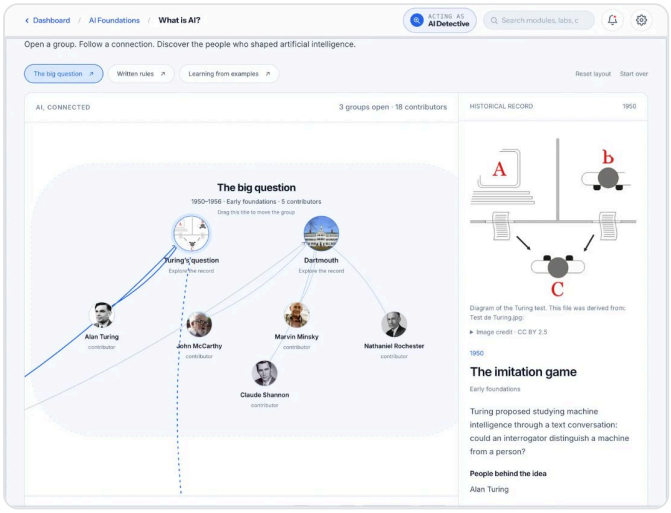
Module 01 · AI Foundations · Unit 02 · Stage 1, Foundations

What is AI?

The Knowledge Map

Grades 6–8 | **40 minutes** | Role: **AI Detective** | **Free**, no account needed

Grade 6: About age 11–12 · Year 7 in the UK · Class 6 in India



The big idea

AI is the field of building machines that do tasks we link with human intelligence, such as using language, recognizing patterns and planning. Some systems follow rules people write; others learn patterns from examples; many combine both. No one person invented AI, and a strong result on one task does not prove a machine can do every task.

The learner's job

Learners are the **AI Detective** at Halden Institute · Intake Desk. AI did not come from one person or one big leap. Follow the map to find the ideas, teams and work behind the tools you see today. Their task: trace an AI idea to the people and work behind it. Show what it can do and where it falls short.

By the end, learners can say

- I can say what AI is in one line: a machine doing a task that usually needs a human mind
- I can tell systems that follow rules from systems that learn from examples
- I can identify key AI milestones, their eras and the people and teams behind them
- I can explain one connection between AI ideas and one limit of a system's results

Before class

- One computer or tablet for each learner or pair, with EdSutra open. Module 1 is free and needs no account.
- Open the studio once on each device before class, so the Unit still runs if the internet drops.
- Play the start yourself: click BEGIN, open a group, and press one "Try it" button, so you know where learners will click.
- Draw two columns on the board, "Written rules" and "Learning from examples". Learners add milestones to them during Understand.
- There is no quiz or step bar: the Unit finishes quietly once all six groups on the map are open.
- Best after AI Alphabet.
- Units open in order: a learner new to EdSutra starts at AI Alphabet. With the full studio, Open every Unit in Settings → Teachers & parents lets a class start here.

Words to introduce

Artificial intelligence: Machines doing tasks that usually need a human mind.

Written rules: Instructions people write for a machine, such as ELIZA's sentence patterns.

Learning from examples: Adjusting a model so it finds a pattern in data, instead of being told.

Search: Looking ahead through possible moves to pick a good one, as Deep Blue did.

Labeled data: Examples that people have tagged with the right answer, like ImageNet's pictures.

Milestone: An important result or event in a field's history, with a date.

Contributor: A person or team whose work helped build an idea, including data workers.

Running the lesson 42 minutes

Open 5 min

"A computer beat the world chess champion in 1997. Does that mean it was intelligent? What would you need to know to decide?"

Before anyone clicks, ask: is AI a machine that follows rules people wrote, or one that learns from examples? Take a show of hands, note the count, and come back to it in Understand.

Step	What learners do	Ask the class	Listen for
1. Remember 15 min	Learners click BEGIN to open the first three groups: the big question, written rules and learning from examples. They open records and find each milestone's date, task and people.	• What does the BEGIN record say AI is, in one line? • Who proposed the name "artificial intelligence", and when?	AI means machines doing tasks linked with human intelligence. The name came from a 1955 proposal by four researchers, not one inventor.
2. Understand 22 min	Learners open the last three groups: learning to see, attention and language, and learning meets search. They press "Try it" on records such as ELIZA and AlphaGo, and read each "Keep in mind" limit.	• Which column does each milestone go in: rules, or learning? Where does AlphaGo go? • Read a "Keep in mind" note. What does that result not prove? • Look back at our show of hands. Was AI rules or learning?	ELIZA and Deep Blue use rules and search; the perceptron, AlexNet and the Transformer learn; AlphaGo combines learning with search.

Steps follow Bloom's revised taxonomy, in the order the Unit presents them.

Deepen the thinking

Critical thinking	ELIZA's replies sounded caring, yet it only matched patterns. How could you test whether a chatbot today understands, or only sounds like it does?
AI thinking	AlphaGo used learning and search together. Why might combining approaches work better than using only rules or only learning?
Ethics and fairness	ImageNet's labels came from people. If those people missed some kinds of pictures, who might a photo app work less well for?
Connect to life	Which AI tool do you use most? Is it more likely using rules, learning from examples, or both? How could you find out?

Misconceptions to listen for

"AI was invented by one person."
The map names 31 people and teams, from Turing to ImageNet's annotators. Even the name "artificial intelligence" came from a 1955 proposal by four researchers.

"All AI learns from data."
ELIZA ran on written rules and Deep Blue searched chess moves. AlphaGo combined learned networks with search. Many systems mix approaches.

"If it sounds human, it understands."
ELIZA's Try it shows four rules turning your words back into a reply. It can sound caring, but you supply the meaning.

In real life the card learners see at the end of the Unit

- A chatbot can sound like a person. ELIZA showed long ago that a convincing reply does not prove understanding.
- Photo apps learn from labeled examples. ImageNet shows how researchers and data workers help build those examples.
- A chess or Go app may combine rules, search and learned patterns, like the approaches you followed on the map.

Try it this week. With an adult, pick one AI tool. Find its team, task and one limit. Explain its story in three short sentences.

Watch out for. A strong result on one task does not prove a machine can do every task. AI has many contributors, not one inventor.

Talk about it. Whose work is easy to miss when a headline names just one AI inventor?

Exit ticket

1. In one line, what is AI?

   _____

2. Name one system that follows written rules and one that learns from examples.

   _____

3. Pick one milestone. What does its result not prove?

   _____

Going further

Module 1's Computer Vision Lab Examples are free, and the camera is optional: learners see a model find hands and faces live, the "learning to see" idea from AlexNet. For homework, use the "Try it this week" activity from the In real life card.

EdSutra

Teaching guide

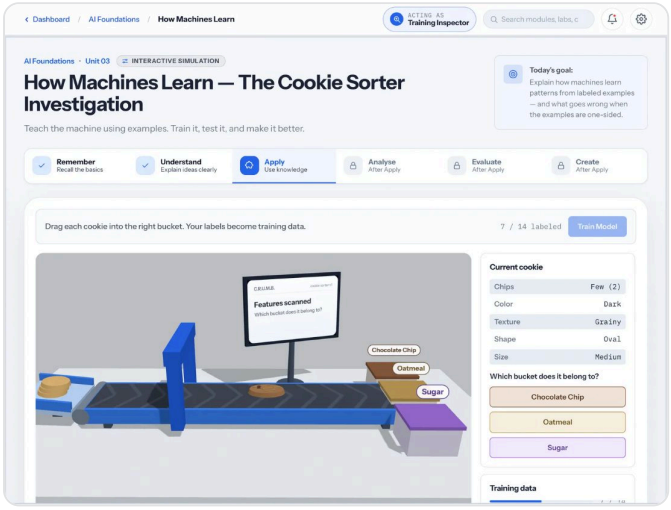
Module 01 · AI Foundations · Unit 03 · Stage 1, Foundations

How Machines Learn

The Cookie Sorter Investigation

Grades 6–9 | **55 minutes** | Role: **Training Inspector** | **Free**, no account needed

Grade 6: About age 11–12 · Year 7 in the UK · Class 6 in India



The big idea

A model learns whatever pattern is easiest in its training examples, which is not always the one we meant. The Cookie Sorter looks perfect, yet on cookies it has never seen it gets about one in four wrong, because it learned color instead of chips. Testing on new examples, and adding counterexamples, is how people catch this and fix it.

The learner's job

Learners are the **Training Inspector** at Halden Institute · Test Bench. The sorter got full marks on its exam. The bakery still gets complaints every week. Their task: find out what the machine really learned, not what it scored.

By the end, learners can say

- I can explain how supervised learning uses labeled examples to learn a pattern
- I can analyze which features a model is using to make predictions
- I can evaluate model behavior with unseen data, accuracy, precision, and recall
- I can diagnose shortcut learning and propose counterexamples that improve generalization
- I can design a stronger feature strategy for a classifier

Before class

- One computer or tablet for each learner or pair, with EdSutra open. Module 1 is free and needs no account.
- Open the studio once on each device before class, so the Unit still runs if the internet drops.
- Play the first two steps yourself, so you know where learners will click.
- Draw two columns on the board, "Training" and "Test". You will fill them in during Analyze.
- Pairs work well: one learner drives, the other is the Training Inspector who notes what the machine leans on.
- Best after AI Alphabet and What is AI?.

Words to introduce

Supervised learning: Learning a pattern from examples that already carry the right answer.

Label: The right answer attached to an example, such as "Chocolate Chip".

Feature: Something the model can measure, such as color or the number of chips.

Training and test data: The examples a model learns from, and the new ones it is checked on.

Shortcut learning: Leaning on an easy pattern that works in training but is not the real reason.

Counterexample: An example added on purpose to break a shortcut.

Accuracy, precision, recall: How often it is right; how many of its picks were right; how many right ones it found.

Running the lesson 55 minutes

Open 5 min

"A machine scored full marks on its exam, but the bakery still gets complaints every week. How can both be true?"

Before anyone clicks, ask: which clue would a machine use to tell a chocolate chip cookie from a sugar cookie? Write the guesses on the board and come back to them in Analyze.

Step	What learners do	Ask the class	Listen for
1. Remember 5 min	Click each part of the sorting machine to learn its name and job, and find all five: input, features, model, decision and buckets.	• Which part turns a cookie into numbers the model can use? • Which part makes the call?	Features are the measurements the model sees, not the cookie itself.
2. Understand 7 min	Press Play and watch cookies get scanned and sorted. Learners notice which features decide each bucket.	• Which feature seemed to decide most of the sorting? • Does the machine know what a cookie is, or only the numbers it was given?	The machine applies a rule to the features it has. It does not understand cookies.
3. Apply 10 min	Learners drag each cookie into the right bucket. Their labels become the training data the model learns from.	• Which cookies were hard to label, and why? • What would happen if one of us labeled carelessly?	Labels are human choices. Mistakes and patterns in them become what the model learns.
4. Analyze 12 min	Run the test batch on cookies the model has never seen. It gets about one in four wrong. Learners inspect each mistake, and can bake their own cookie to try to fool the model.	• What do the wrong cookies have in common? • What did the model lean on: chips, or color? • Look back at our guesses on the board. Were we right?	The model learned color, a shortcut that happened to work in training, instead of chips, the real reason.
5. Evaluate 8 min	Learners read the scorecard (accuracy, precision and recall) and make the call: ship the sorter or send it back.	• Would you ship it? What would the bakery's customers say? • Which number matters most to this bakery, and why?	A decision backed by the test results, not by the training score.
6. Create 8 min	Learners change the features or add new training cookies, retrain, and try to beat Model A.	• Which counterexample did you add, and why did it help? • Did switching a feature off make the model better or worse?	Better features and counterexamples help a model do well on cookies it has never seen.

Steps follow Bloom's revised taxonomy, in the order the Unit presents them.

Deepen the thinking

Critical thinking	A model can be right for the wrong reason. How would you check that a tool marking essays is right for the right reason?
AI thinking	If you could give the sorter one more feature, which would you choose, and what new shortcut might it find?
Ethics and fairness	If a model that approves loans learned a shortcut, who could be harmed, and whose job is it to notice?
Connect to life	Where have you seen a test score that did not show what someone really knew?

Misconceptions to listen for

"If it is accurate on its training data, it is good."

Analyze shows the sorter failing on new cookies. Only test data the model has never seen shows what it really learned.

"The model understands cookies."

It only uses the features it is given. Switch Color off in Create and its decisions change.

"Bias means the machine is being unkind."

Here, bias means learning an easy, wrong pattern from the examples. Counterexamples fix it.

In real life the card learners see at the end of the Unit

- A spam filter learned from emails people sorted into "spam" and "not spam", just like your cookie labels.
- A photo app that finds "dogs" was trained on labeled pictures, then tested on new ones it had never seen.
- At a shop checkout, a card payment can be blocked as "possible fraud" by a model trained on labeled payments.

Try it this week. Sort 10 objects at home into two groups using a secret rule. Ask someone to guess the rule. Did they pick up a shortcut, like color?

Watch out for. A model can be right for the wrong reason, like sorting cookies by color instead of chips. Test it on new examples before you trust it.

Talk about it. Have you ever learned a shortcut that worked in practice but failed on the real test?

Exit ticket

1. What is the difference between training data and test data?

2. Which shortcut did the Cookie Sorter learn, and which clue should it have used?

3. Describe one counterexample that would help, and say why.

Going further

In the full studio, the Neural Playground Lab lets learners train a real network to read handwritten digits and test it on their own writing. For homework, use the "Try it this week" activity from the In real life card.

EdSutra

Teaching guide

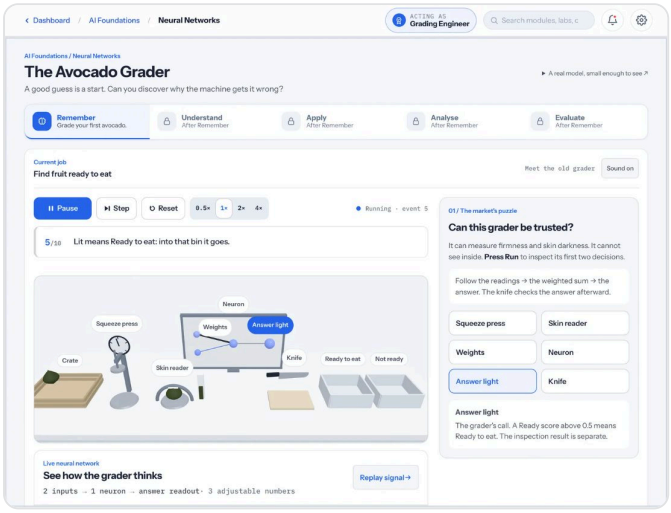
Module 01 · AI Foundations · Unit 04 · Stage 1, Foundations

Neural Networks

The Avocado Grader

Grades 7–10 | **30 minutes** | Role: **Grading Engineer** | **Free**, no account needed

Grade 7: About age 12–13 · Year 8 in the UK · Class 7 in India



The big idea

A neuron multiplies each reading by a weight, adds a bias, and lights up when the total is high enough. One neuron makes one straight cut, so it cannot pick out avocados that are ready to eat, between rock-hard and mushy. Two neurons together can. Training is the computer nudging every weight to make fewer mistakes.

The learner's job

Learners are the **Grading Engineer** at Halden Market · Packhouse. This grader reads two numbers per avocado. Right now it can't tell ready from mushy. Their task: teach the grader to spot ready-to-eat. Then explain what each neuron learned.

By the end, learners can say

- I can name the parts of a neural network (inputs, weights, bias, neuron, output)
- I can work out by hand whether a neuron lights up from its readings, weights and bias
- I can train a network and choose how many neurons it needs
- I can explain what each trained neuron learned, and choose between networks on evidence

Before class

- One computer or tablet for each learner or pair, with EdSutra open. Module 1 is free and needs no account.
- Open the studio once on each device before class, so the Unit still runs if the internet drops.
- Play the first two steps yourself, so you know where learners will click and how the dials feel.
- Draw a squeeze line on the board from 0 (mushy) to 1 (rock-hard). You will use it in Apply to show one cut against two.
- Pairs work well: one learner drives, the other is the Grading Engineer who records each run's settings and score.
- Best after AI Alphabet and What is AI? and How Machines Learn.
- Units open in order: a learner new to EdSutra starts at AI Alphabet. With the full studio, Open every Unit in Settings → Teachers & parents lets a class start here.

Words to introduce

Input: A reading the network gets, here squeeze and skin darkness, from 0 to 1.

Weight: A number that sets how much one reading counts in the sum.

Bias: A number added to the sum. Here it does not mean unfairness.

Neuron: Adds weighted readings and a bias into a signal from 0 to 1.

Hidden neuron: A neuron between the inputs and the output; nobody tells it its job.

Output: The final neuron that makes the call: Ready to eat or Not ready.

Training round: One pass through all 24 labeled practice avocados, nudging every weight.

Running the lesson 39 minutes

Open 5 min

"A machine can only squeeze an avocado and check its skin color. Could it ever tell you which avocados are ready to eat?"

Before anyone clicks, ask: is a ready avocado hard, soft, or in between? Then: could one rule, like "softer than this line", pick out only the ready ones? Note the votes and come back to them in Apply.

Step	What learners do	Ask the class	Listen for
1. Remember 5 min	Learners press Run and follow two avocados through the shop's old one-neuron grader. They predict the second, softer one before the knife shows the truth, then answer: what does a weight do?	• Which two readings does the grader use, and what can it not see? • What did the knife show about the second avocado?	The grader only sees squeeze and skin; it called a mushy one ready. A weight sets how much a reading counts.
2. Understand 8 min	Learners predict whether the old neuron lights up for a rock-hard avocado. Then they turn three dials (firmness weight, skin weight, bias) until it lights for rock-hard ones only, and run the practice crate.	• What happens to the sum when you make the firmness weight bigger? • What does the bias do, if it is not about fairness?	A bigger firmness weight lifts the sum for hard avocados. The bias is added every time, shifting where it lights.
3. Apply 9 min	Learners choose 1 to 3 hidden neurons and 50 to 2,000 rounds, predict how many of 12 new avocados it gets right, and train. Then they compare one neuron with more at the same rounds.	• One neuron trained for 2,000 rounds: did it get there? Where did the mushy avocados go? • Look at our squeeze line. Can one cut keep only the ready avocados on one side?	Ready avocados sit between rock-hard and mushy. One straight cut cannot fence in the middle; two neurons can.
4. Analyze 6 min	Learners probe each neuron of their trained network: the practice crate runs through, and each avocado lands in Lit or Dark by that neuron alone. They name each neuron's job.	• Nobody told the neurons their jobs. What job did each one find? • Why did one hidden neuron miss the middle group?	One lights for any avocado that is not rock-hard, one for mushy ones. The output combines them to find ready-to-eat.
5. Evaluate 6 min	Learners run graders A, B and C on the stall's crate of 12, pick one to send, and back it with one fact from the runs. They also choose what the market should check next.	• B and C both got 12 of 12. Which would you send, and why? • Grader C has the most neurons. Why is that not evidence on its own?	A choice backed by results on the stall's crate, such as B matching C with fewer numbers, not by size.

Steps follow Bloom's revised taxonomy, in the order the Unit presents them.

Deepen the thinking

Critical thinking	A grader got 12 of 12 on one crate. What would you still want to know before trusting it with every delivery?
AI thinking	Face unlock uses a far bigger network than our grader. What stays the same inside it, and what gets harder to explain?
Ethics and fairness	If a grader decided loans or school places instead of avocados, who should be able to open the black box and ask what each part does?
Connect to life	Think of a choice you make by weighing two things, like whether to bike or take the bus. What would your weights be?

Misconceptions to listen for

"A weight is how heavy something is."

The Remember question shows it: a weight sets how much one reading counts in the sum. A weight of 0 ignores that reading.

"Train it longer and it will get there."

One neuron trained for 2,000 rounds still sends every mushy avocado to Ready to eat. Training can only move the one line it has.

"More neurons is always better."

Grader C's three neurons scored no better than grader B's two, and C has more numbers to check. Evaluate refuses "more neurons" as evidence.

In real life the card learners see at the end of the Unit

- Face unlock on many phones uses a neural network that learned which patterns match your face.
- Speech-to-text turns your voice into numbers, passes them through layers, and outputs its best guess at the words.
- At the post office, machines read handwritten addresses. Reading handwriting is a classic neural-network job.

Try it this week. At a shop, gently squeeze a few avocados or peaches. Sort them: rock-hard, just right, mushy. You're checking two things at once, like the grader's two neurons.

Watch out for. A network doesn't understand avocados or faces. It adds up readings times weights. A high score on one crate isn't proof it will work on the next.

Talk about it. Brains inspired neural networks. How is a network like your brain, and how is it different?

Exit ticket

1. What does a weight do in a neuron?

2. Why can one neuron not pick out the ready-to-eat avocados?

3. Which grader would you send to the stall, and what fact backs it?

Going further

In the full studio, the Neural Playground Lab lets learners draw their own digits, train a real network on them, and take it apart neuron by neuron. For homework, use the "Try it this week" activity from the In real life card.