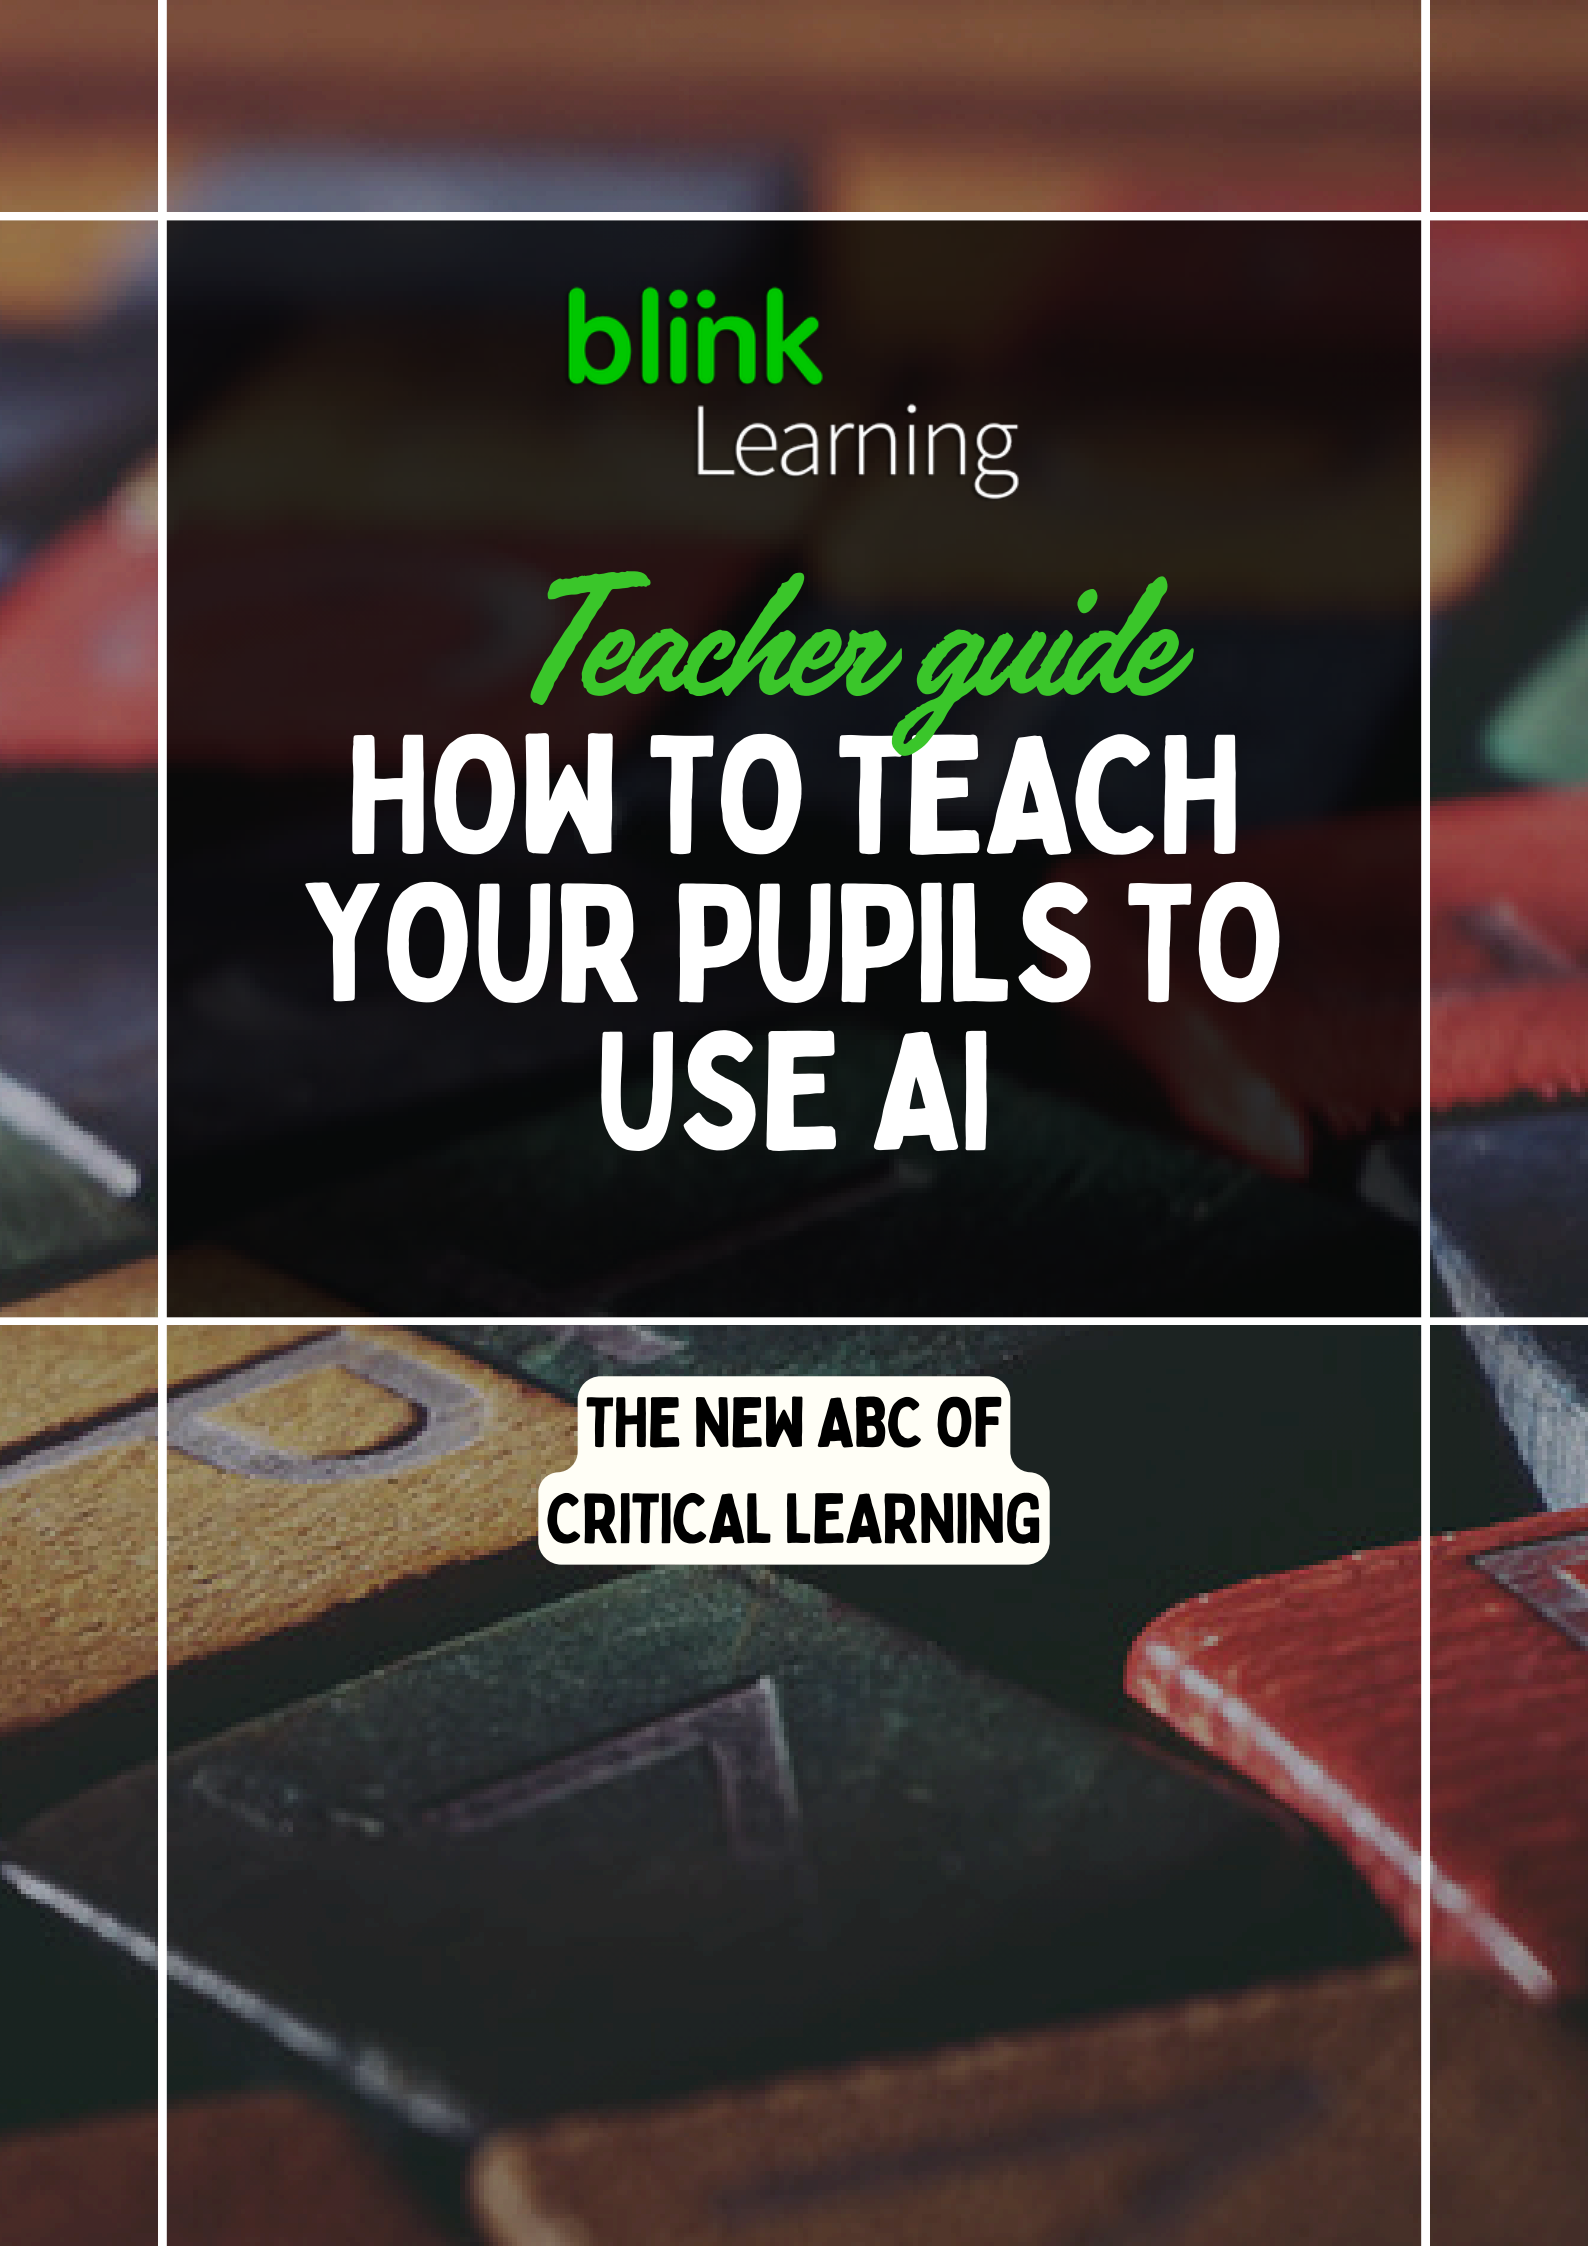
blink
Learning
Teacher guide
HOW TO TEACH YOUR PUPILS TO USE AI
THE NEW ABC OF CRITICAL LEARNING

GUIDE: 'HOW TO TEACH YOUR PUPILS TO USE AI'

blink Learning

TEACHING UNIT

DISCOVERING ARTIFICIAL INTELLIGENCE

blink Learning

CONTENTS

blink Learning

CONTENTS

Learning situation 3: Train the Clumsy Robot

ASSESSMENT OF THE TEACHING PROGRAMME UNIT

blink Learning

INTRODUCTION

This teaching unit, entitled 'Discovering Artificial Intelligence: how machines think, learn and improve', aims to introduce younger pupils to the foundations of machine learning and computational thinking in an accessible, playful and ethical way.

Through three interconnected learning situations, pupils explore three essential processes underpinning artificial intelligence:

1. **Prediction** – how language models choose the next word (The Sad Dragon and Predictive Text activity).
2. **Training** – how models learn to recognise patterns from examples (Train the Mysterious Model activity).
3. **Reinforcement** – how models improve their behaviour through feedback (Train the Clumsy Robot activity).

These experiences require no digital resources and use an unplugged format to turn complex concepts into collaborative, experimental activities.

Each session combines play, observation, hypothesis-building and final reflection, allowing pupils to develop an intuitive understanding of basic AI mechanisms and the human role in their design and use.

Unit structure

The unit is organised around three learning situations, with a **progressive thread moving** from simple information generation to improvement through feedback.

Each situation includes:

- A narrative context or opening game.
- A guided practical experience.
- A shared reflection.
- A final outcome or visible piece of evidence.

blink Learning

INTRODUCTION

Flexibility and adaptability

This proposal is flexible and modular, allowing it to be used:

- **At different educational levels:**
  - It is designed for pupils aged 8–12, with the complexity of language, examples and pupil autonomy adapted accordingly.
  - For younger pupils, activities can be simplified through visual materials and teacher support.
  - For older pupils, they can be extended with ethical reflection, data recording and links to real digital tools.

- **Across different curriculum areas:** Although linked to Computing, the unit has been designed for an interdisciplinary approach and can be used in different educational contexts. Teachers need no prior specialist knowledge or special materials.
  - In **English**, pupils practise precise communication and collaborative storytelling.
  - In **Mathematics**, they identify patterns and record data.
  - In **Science**, they compare human, animal and artificial learning.
  - In **values and citizenship education**, they reflect on ethics and technological responsibility.
  - In **art, drama and movement**, they express simulated processes creatively and physically.

- **Timing:** the teaching programme unit is intentionally flexible. Its three learning situations are internally coherent and can be taught independently or combined into an extended unit. They can be delivered consecutively or spread across the academic year.

This versatility allows the unit to form part of an annual Computing programme, a whole-school interdisciplinary project on innovation, science or critical thinking, or a special themed event such as a cultural week.

blink Learning

INTRODUCTION

Educational purpose

The ultimate aim is for pupils to understand that artificial intelligence is not magic, but learning, and that its development depends on data, rules and human decisions.

Through practical experience, children discover that machines learn in ways that resemble people, but without consciousness or values of their own. Humans are therefore ultimately responsible for their behaviour and ethical use.

This unit teaches more than technology: it develops critical thinking, collaboration and digital responsibility—three essential competencies for twenty-first-century citizenship.

blink Learning

Learning situation 1:

THE SAD DRAGON AND PREDICTIVE TEXT

Title:
The Sad Dragon and Predictive Text

Recommended age:
Pupils aged 8–12; adaptable for younger pupils with more guidance and pedagogical scaffolding.

Subject area(s):
Choose from English, Mathematics, Science—with a focus on technology—and values and citizenship education.

Approximate duration:
45–60 minutes (extendable to two sessions)

Final outcome / challenge:
Collaborative story + poster/record of the class's 'selection algorithm'

Overview:
The class simulates how a large language model (LLM) works by generating a story one word at a time from an initial prompt. The group suggests possible next words, votes or decides how to resolve ties using majority, chance or other rules, and records each decision. This shows that an LLM 'predicts the next word' using context and learnt patterns.

blink Learning

Learning situation 1:

THE SAD DRAGON AND PREDICTIVE TEXT

Learning objectives

- Understand, in accessible language, that an LLM **predicts the next word** based on context and language patterns.
- Experience how different **selection strategies**—majority, chance and limiting options—affect a text's coherence and creativity.
- Recognise basic **elements of narrative**: opening, characters, setting, conflict and ending.
- Develop **participation, reasoning and collective decision**-making skills.
- Introduce functional school-level AI vocabulary: 'prompt', 'next word', 'context' and 'selection rules'.

Key and subject-specific competencies

- **Communication and literacy:** spoken expression, vocabulary development and narrative structure.
- **Critical digital competence:** basic understanding of how AI tools work and developing an appropriate mental model.
- **Mathematics/statistics:** initial understanding of frequency, choice and chance.
- **Learning to learn/citizenship:** collaboration, agreement and metacognition.

Content

- **Storytelling:** opening formulas ('Once upon a time...'), story sequence, cohesion and coherence.
- **Unplugged computational thinking:** rules, procedure, input (prompt) → process (selection) → output (story).
- **Basic AI:** predicting the next word (tokens), context and the influence of prior data—knowledge of how stories usually work.
- **Selection strategies:** majority (more likely/coherent), chance (more creative/risky), top-k (choosing from the two or three most-supported options).

Learning situation 1:

THE SAD DRAGON AND PREDICTIVE TEXT

1. Understanding how a large language model (LLM) works

Large language models (LLMs), such as ChatGPT, Gemini and Claude, are artificial intelligence programs designed to predict the next word in a sequence of text. They do not write complete sentences all at once or know in advance what they will say. They use statistical prediction:

1. **Training:** the model 'reads' millions of online texts, books, news articles and other sources. From these it learns linguistic patterns: how words combine, which expressions tend to follow others and which structures are typical of a story, recipe or conversation.
2. **Input or 'prompt':** the user provides an instruction or starting point ('Write a story about a sad dragon beginning with Once...').
3. **Prediction:** the model analyses the context and calculates which word is most likely to come next, based on what it has learnt and what fits the context.
4. **Selection:** it chooses one possible option. With a low 'temperature', it usually chooses the most probable word, producing more coherent and predictable answers. With a high temperature, it may choose less likely options, producing more creative but potentially less coherent answers.
5. **Iteration:** it repeats the process for each new word, taking account of what it has already written. The text is built word by word, like a story invented as it unfolds.

Put simply, a language model works like a pupil trying to guess the next word in a text from previous experience of other texts.

This learning situation turns that process into a collaborative unplugged activity in which pupils collectively act as a 'human language model'.

Learning situation 1:

THE SAD DRAGON AND PREDICTIVE TEXT

2. How to carry out the activity

Teacher preparation:

- Write the opening prompt on the board:
- 'Write a story about a sad dragon beginning with the word Once.'
- Prepare a visible space on the board or large sheet of paper with columns for:
  - Current context (what has been written so far)
  - Suggested next words
  - Selection rule (majority, chance, etc.)
  - Chosen word

It should look something like this:

What we've written so far	Suggested words	How we make decisions	Chosen word
Once upon a time, there was a...	dragon / castle / child / wizard	Majority of votes	dragon
Once upon a time, there was a dragon...	sad / green / huge / brave	Majority of votes	sad
Once upon a time, there was a sad dragon who...	lived / cried / flew / sang	Tie → chosen at random	was crying
Once upon a time, there was a sad dragon who was crying...	a lot / every night / in his cave / because he was alone	Majority of votes	because he was on his own

Step-by-step activity:

1. **Activate prior knowledge:** briefly discuss stories with the class. How do they begin? What usually happens? Which words are typical? Explain that they will become 'a machine that predicts the next word'.
2. **Story opening:** write 'Once' and ask the group to suggest the next word. Each pupil writes the word they think should follow on an individual whiteboard, sticky note or another suitable medium.

3. Show the options: all pupils reveal their words at the same time. Quickly count the choices. At first, options will be obvious ('upon', 'a'), but they will begin to diverge.
4. Collective decision: when there are several options, the group votes. If there is a tie, they decide how to resolve it—by chance, majority or another rule. This is a key moment, representing different AI-model strategies.
5. Record: note each choice and selection method. This will help pupils later reflect on how decisions changed the story's final outcome.
6. Metacognitive pauses: every few words, stop and ask:

- Does our story make sense? Does it fit the prompt?
- What would happen if we chose another word?
- Are we repeating the same words too often?
- Why is the text moving in this direction?

7. Finish: when the story reaches an ending or time runs out, read it in full. Ask:

- Did we know we were going to write this story?
- Why did it turn out this way?
- What role did our decisions play?
- Would another class write the same story?
- How is this similar to, and different from, the way AI works?

3. Conclusion and possible extensions

Conclusion for the teacher:
This activity provides a **tangible experience of how an AI model 'thinks'** when generating text. Pupils understand that AI:

- Does not invent from nothing: it uses learnt patterns.
- Does not have an overall plan: it chooses one word at a time and does not know the final text until it produces it.
- Can be more or less creative or coherent depending on how choices are made.
- Reflects biases and conventions in its training texts: if every story depicts dragons as bad, the model will tend to repeat that pattern.

Pupils therefore learn through play that AI does not 'think' or 'feel'; it **predicts** based on what it has learnt from data.

Extension opportunities:

- Suggest other prompts unrelated to stories, such as explaining how to make an omelette or solving a maths problem. Ask pupils to reflect on how output quality changes when they have less context or knowledge about what they need to write.
- Introduce further AI vocabulary: instructions given to a model are called a 'prompt'. AI models that generate text are called LLMs (large language models).

4. Interdisciplinary links

1. **Science / Mathematics:**
   - Analyse voting frequencies or the probability of one word rather than another being selected.
   - Graph how often each strategy—majority, chance and so on—was used.
2. **English:**
   - Compare stories created using different rules, such as majority or chance.
   - Identify which choices add coherence and which add originality.
3. **Values education / Digital citizenship:**
   - Reflect on bias: what kinds of dragons appear? Are they all bad?
   - Discuss how the data we use influences what AI 'learns', just as people learn from their environment.
4. **Extension with real technology:**
   - Try the same prompt in an AI assistant such as ChatGPT or Gemini and compare the results.
   - Analyse how the model's choices—length, tone and coherence—are similar to or different from human choices.

Final message

This activity is not intended to teach programming, but critical thinking about AI. Its value lies in pupils experiencing the process from within: deciding, making mistakes, reflecting and understanding that artificial intelligence is not magic, but a form of prediction based on language patterns.
It is an excellent starting point for exploring how machines learn—and how we learn.

blink Learning

Learning situation 2:

TRAIN THE MYSTERIOUS MODEL

Recommended age:
Pupils aged 8–12.

Subject area(s):
Choose from Science, Mathematics, Computing and values and citizenship education.

Approximate duration:
45–60 minutes, extendable through additional or comparative versions

Final outcome / challenge:
Discover the model's 'secret rule' and understand how AI learns from examples and mistakes

Overview
The group faces a challenge: discover the hidden rule the teacher has in mind, just as an artificial intelligence model must learn a pattern from data.
The teacher shows examples that do or do not follow the rule, and pupils formulate, test and refine hypotheses with each new piece of data.
The process mirrors supervised AI training: observe examples → detect patterns → make mistakes → correct → improve.
As the game progresses, pupils understand that AI learns through observation and adjustment, and that its accuracy depends on the quantity and quality of data.

blink Learning

Learning situation 2:

TRAIN THE MYSTERIOUS MODEL

Learning objectives

- Understand that AI learns to recognise patterns from examples.
- Experience learning through trial, error and gradual adjustment.
- Value the importance and diversity of data in machine-learning quality.
- Develop logical thinking, hypothesis formulation and group cooperation.
- Introduce functional AI vocabulary: training, data, model, pattern, error and adjustment.

Key and subject-specific competencies

- Science and technology: observation, experimentation and inductive reasoning.
- Mathematics: pattern identification, classification and rule formulation.
- Digital: conceptual understanding of machine learning.
- Learning to learn: planning, revising hypotheses and self-regulation.
- Social and civic: cooperation, respect for mistakes and collective critical thinking.

Content

- Science / Computing: how machines learn from data.
- Unplugged computational thinking: training, model adjustment and pattern detection.
- Mathematics: classification, logical deduction and hypothesis testing.
- Values and citizenship education: mistakes as part of learning, cooperation and group decision-making.
- AI vocabulary: data, patterns, training, error, model, accuracy and bias.

blink Learning

TEACHER GUIDE

Learning situation 2:

TRAIN THE MYSTERIOUS MODEL

This activity provides an experiential explanation of how artificial intelligence learns. Pupils do not need computers: learning takes place through examples, hypotheses and corrections, just as an AI model learns when trained on data.

1. Core idea

Begin by explaining that artificial intelligence learns by observing examples.

For instance, an AI designed to recognise dogs sees thousands of pictures of dogs and of other things. It discovers patterns such as pointed ears, muzzles, legs and fur. Nobody tells it directly; it learns to distinguish them by adjusting its 'rules' whenever it makes a mistake.

In this activity, the class recreates that process. The teacher has a secret rule that acts as the AI's 'hidden pattern'. Without knowing it, the group must discover it from the examples the teacher shows, just as an algorithm does during training.

2. Preparation

Before beginning, choose a simple but not obvious rule suited to the group's age. Possible examples:

- **Round things**
- **Objects used for writing**
- **Animals that live in water**
- **Pictures containing the letter A**
- **Metal objects**
- **Things used at school**

Prepare 15–20 pictures, cards or real classroom objects.
Also prepare two signs or areas on the board labelled:
Belongs | Does not belong

Learning situation 2:

TRAIN THE MYSTERIOUS MODEL

3. How to carry out the activity

1. Present the challenge:
Explain that you are an 'artificial intelligence' that has already been trained, but nobody knows its rule. They must discover it, just like scientists who study how machines think.

2. First set of examples:
Show five pictures or objects, classifying each one yourself as 'Belongs' or 'Does not belong'.
Do not explain why. Simply show which group is correct.

3. Initial hypotheses:
Ask pupils to identify the pattern. Listen to their hypotheses and write them on the board.
(For example: 'They are round things', 'They are fruit', 'They are small things'.)

4. New examples and adjustments:
Show more objects. After each one, ask:
'Which group do you think it belongs to? Why?'
If they are right, confirm it; if not, say that it does not belong and let them revise their hypothesis.
Pupils will gradually **refine their rule** as they receive more data, just as an AI adjusts its model during training.

5. Testing phase:
When they think they know the rule, show **new examples** that they have not seen before.
They must now predict which group each example belongs to.
Then reveal whether they are correct.

- If they make mistakes, discuss why:
- Was their rule too broad?
- Did they need more examples?
- Was anything confusing?

Learning situation 2:

TRAIN THE MYSTERIOUS MODEL

4. Plenary and reflection

End the session with a short discussion guided by questions such as:

- How did we learn the rule?
- What happened when we made a mistake?
- How is this similar to the way an AI learns?
- Why are examples so important?
- What would happen if we only showed objects of one colour or type?

Help them connect this with the key idea:

Artificial intelligence **does not think as we do, but it learns by observing patterns in data.**
If examples are good and varied, it learns well.
If the data is limited or confusing, it makes mistakes, just as we did at the start of the game. All AI knowledge is based on prior human knowledge.

5. Extension opportunities

- **Different models:** divide the class into two groups and give each group a different sequence of examples. At the end, ask them to guess the other group's rule: they will see that each model learns something different.

- **Bias:** train the model using only examples of one colour or type (for example, only red objects). Then show a green one and ask whether it should belong. This will help pupils understand bias.

- **Comparison with real AI:** if devices are available, try an image-recognition app such as Google Lens and compare its correct answers and mistakes with those of the class model.

Learning situation 2:

TRAIN THE MYSTERIOUS MODEL

6. Interdisciplinary options

- **Natural Sciences**
  - **Biological classification:** use pictures of animals or plants and discover the characteristic that groups them (mammals, egg-laying animals, flowering plants, etc.).
  - → Link this to scientific classification criteria and observing patterns in nature.
  - **Environmental impact:** AI is not just words or images. Reflect on how AI might help—or fail—when classifying waste, detecting species and so on.

- **Mathematics**
  - **Statistics and probability:** count the model's correct answers and errors, calculate success rates and discuss what 'better training' means.
  - **Logic and patterns:** introduce basic conditional logic ('if it has four sides, then...'), linking it to computational thinking.

- **English**
  - **Grammar:** use examples involving word classes or phrases for classification.
  - **Reading comprehension:** consider how machines can 'understand' a text without truly understanding its meaning.
  - **Technical vocabulary:** introduce terms such as pattern, data, error, hypothesis and training through the game.

- **Art**
  - **Recognising styles:** use pictures of artworks or different lines and colours so pupils identify what they have in common.
  - → Example: 'works with curved lines', 'drawings using warm colours'.
  - **Guided creation:** pupils invent their own 'model' with a secret aesthetic rule, which classmates try to discover.
  - → This fosters creativity and metacognition about visual rules.

Learning situation 2:

TRAIN THE MYSTERIOUS MODEL

- **Values and citizenship education**
  - **Data ethics:** discuss the responsibility involved in training AI—what data is used, where it comes from, and whether it may be unfair or incomplete.

- **Error and learning:** connect model errors with human learning and value mistakes as part of the process.

- **Collaborative work:** emphasise that 'human models' work better when they share ideas and combine information, just as AI learns better from more data.

Final message

The aim is not for pupils to learn a formal definition of 'machine learning', but to **experience how learning takes place through observation and correction.**

It is an excellent opportunity to discuss:

- The importance of **data** in artificial intelligence.
- The value of **mistakes** as part of learning.
- The difference between 'thinking' and 'recognising patterns'.

When a pupil exclaims, 'Now I know the rule!', they have experienced on a small scale what AI does: **learn from data.**

blink Learning

Learning situation 3:

TRAIN THE CLUMSY ROBOT

- **Recommended age:** Pupils aged 8–12; it can be adapted for younger children with simpler tasks and visual support.
- **Subject area(s):** Computing, values and citizenship education, Physical Education and Science.
- **Approximate duration:** 45–60 minutes (it can be repeated with different tasks or compared across rounds)
- **Final outcome / challenge:** Help the 'clumsy robot' complete a task using only binary feedback ('right/wrong')

Overview

Pupils experience how artificial intelligence learns through feedback or reinforcement.

One pupil—or the teacher—acts as a 'clumsy robot' that must complete a simple task, such as drawing a shape, moving around the classroom or arranging objects, without knowing the objective.

The rest of the group become its 'trainers': after each action, they may only respond 'right' or 'wrong', without giving instructions or clues.

With each attempt, the robot adjusts its behaviour and gradually improves, just like an AI trained through reinforcement learning: trial, error, correction and improvement.

At the end, the class reflects on how machines learn from human feedback and why they need these signals to adjust their decisions.

blink Learning

Learning situation 3:

TRAIN THE CLUMSY ROBOT

Learning objectives

- Understand that AI can learn through feedback and reinforcement, whether positive or negative.
- Experience trial-and-error learning and understand how performance gradually improves.
- Recognise the importance of human feedback in the development of artificial intelligence.
- Develop communication and group-cooperation skills.
- Introduce functional AI vocabulary: reinforcement, learning, feedback, error, reward and adjustment.

Key and subject-specific competencies

- **Digital:** conceptual understanding of reinforcement learning.
- **Science and technology:** observing processes, formulating and refining hypotheses.
- **Social and civic:** cooperation, empathy and assertive communication.
- **Learning to learn:** self-regulation and valuing mistakes as part of the process.

Content

- Principle of r**einforcement learning**: improvement based on received feedback.
- Mistakes as part of both human and artificial learning.
- The difference between instructions and feedback.
- Core vocabulary: reinforcement, reward, feedback, error, learning, improvement.
- Unplugged computational thinking: action → evaluation → adjustment.

Learning situation 3:

TRAIN THE CLUMSY ROBOT

1. Understanding reinforcement learning

This activity allows pupils to experience, simply and without technology, how artificial intelligence learns through reinforcement—a type of learning based on attempts, mistakes and feedback.

Reinforcement learning (RL) is one of the three main families of machine learning, alongside supervised learning (learning from examples with correct answers) and unsupervised learning (finding patterns without labels).

In reinforcement learning, AI does not receive correct answers, but only evaluation signals indicating whether its actions are good or bad.
It works as follows:
The system, called an agent, acts within an environment.
After each action, it receives a positive or negative reward based on the outcome.
The agent adjusts its behaviour to maximise rewards and avoid penalties.
Over time, it learns which sequence of actions produces the best results.

We can think of a dog learning to sit: nobody explains how to do it; instead, the dog associates the correct action with 'well done!' or a treat, and an incorrect action with 'no'. AI works similarly, learning which behaviours are useful from the consequences it receives.

Real examples of AI trained through reinforcement include:

- **Robots** learning to walk or balance.
- **Chess and video-game programs** such as AlphaGo, which improve by playing thousands of games and correcting their errors.
- **Recommendation systems** that adjust what they display in response to user behaviour, including clicks and viewing time.

In every case, there are no exact instructions—only a system of 'rewards' and 'penalties' that guides learning.
In this activity, the clumsy robot represents the agent (the AI), while the group of pupils acts as its environment and reinforcement system.
The robot does not know the objective, just like an AI starting from scratch.

After each action, the group provides binary reinforcement:

- **'Right'** → the action moves towards the objective (positive reward).
- **'Wrong'** → the action moves away from the objective (negative reward).

At first, the robot acts randomly. Over time **it experiments, observes the responses and improves its strategy** until it completes the task.

This mirrors the technical process of a reinforcement-learning algorithm, which **learns to act correctly without being told the explicit goal** and relies only on the consequences of its actions.

2. Preparation before you begin

Materials:

- Cards or signs displaying the words '**RIGHT**' and '**WRONG**' (or agreed gestures/icons 👍👎).
- Clear space in the classroom or corridor.
- Objects or obstacles such as cones, chairs or sheets of paper on the floor if the task involves movement.
- A board or large sheet of paper to record the 'training rounds' and observe progress.

Choosing the task:

Choose a simple task with several possible steps, suited to the year group, for example:

- Reach a particular classroom object without being told which one.
- Draw a shape on the board, such as a star or triangle.
- Lower a blind.
- Build a tower from cups or blocks.

Arrange the room so that the robot has several available options. For example, it could begin beside a table holding plastic cups, chalk and a pair of compasses.

Essential rules:

- The robot does not know the final objective.
- The group may say only 'right' or 'wrong' after each action.
- Nobody may give direct instructions or explain why.

In this way, pupils faithfully reproduce the principle of **reinforcement learning**, in which the model receives only information about whether its behaviour is moving towards or away from the goal.

Learning situation 3:

TRAIN THE CLUMSY ROBOT

3. Step-by-step activity

1. Introduction (5–10 minutes)

Explain to the class:
'Today you are going to train a clumsy robot. It does not know what to do, but it will gradually learn from your "right" and "wrong" responses. Artificial intelligence learns in the same way: it tries things, makes mistakes and improves through feedback.'

Choose the pupil who will play the clumsy robot; the role can rotate later.
Explain that they may perform only one action per turn, without speaking or asking questions.
The rest of the class are the trainers.

2. First training round (10–15 minutes)

Begin the task.
The robot acts at random ('I take one step', 'I pick up the chalk', 'I turn around') and may perform only one action. The group responds only with:

- '**Right**' if the action moves towards the objective.
- '**Wrong**' if it moves away.

The teacher can record actions on the board to show the learning sequence.
During the first rounds, the robot will probably become confused. This is intentional: at the start, the model knows nothing.

3. Subsequent rounds (15–20 minutes)

The robot adjusts its behaviour in response to the reinforcement.
With each attempt, it learns which actions receive 'right'.
The group also improves the way it gives feedback, avoiding chaos or shouting.

The teacher can add a quantitative dimension:

- Round 1 → 2 correct / 10 attempts
- Round 2 → 5 correct / 10 attempts
- Round 3 → 8 correct / 10 attempts

This makes the **learning curve** visible through data. When the robot completes the task correctly, 'the model has learnt'.

4. Variations or repetitions (10 minutes)
Change the robot and task to compare results:

- How many attempts did each robot need?
- Which feedback helped most: quick, slow, calm or collective?
- These comparisons reinforce the idea that **the quality and clarity of feedback affect learning effectiveness.**

5. Plenary and reflection (10–15 minutes)
Once the task is complete, hold a short group reflection.

Suggested questions:

- How did you know what to do if nobody told you?
- How did you feel when you heard only 'wrong'?
- Did hearing 'right' help?
- What happened when the signals were confusing or contradictory?
- Do you think machines can learn in the same way you did?

Key conclusion:
'Artificial intelligence learns like our robot: it does not know what to do, but adjusts its behaviour through feedback. The clearer and more consistent the reinforcement, the better it learns. If nobody provides feedback, or if signals are confusing, it does not learn effectively.'

4. Practical tips

- Do not allow the group to give disguised clues such as 'Right, but a little further!'
- Only the words '**right**' and '**wrong**', or the agreed gestures, are permitted.
- Adjust task difficulty by age:
  - For younger pupils, use basic movements such as walking or turning.
  - For older pupils, use multi-step tasks or tasks involving objects.
- Use mistakes: every error is an opportunity to explain that learning requires **experimentation and adjustment**.

Learning situation 3:

TRAIN THE CLUMSY ROBOT

5. Extension opportunities

1. **Art version:** the robot must draw an object without knowing what it is; the group reinforces each line.
2. **Logic version:** the robot must arrange cards according to a secret rule such as colour, shape or size.
3. **Comparative version:** two robots work simultaneously with different groups; one receives clear feedback and the other confusing feedback.
4. → **Compare the results:** which one learnt more effectively?
5. **Ethics version:** discuss what would happen if the robot received unfair or contradictory signals, introducing ethical learning or alignment.

6. Connecting with real AI

To finish, show real examples:

- **Robots learning to walk**, using short 'robot learning to walk' videos.
- **AI playing chess or video games** and improving with each match.

Explain that these systems do not know the rules beforehand:

'They learn only by trying, making mistakes and receiving "right" or "wrong", just as you did today.'

Pupils therefore understand that reinforcement learning is one of the main foundations of modern AI and a fundamental technological advance. A person does not need to programme every possibility; instead, the program explores possible options and automates the learning process.

7. Interdisciplinary suggestions – Train the Clumsy Robot

The Train the Clumsy Robot activity offers extensive links with different curriculum areas by combining logical thinking, communication, cooperation, physical control and ethical reflection.

The following suggestions show how it can be adapted or extended collaboratively across subjects.

- **English**
  - **Precise communication:** explore the value of clear words and economy of language. After the activity, pupils can analyse how the lack of explicit instructions forces them to 'communicate with less', just like machines.
  - **Narrative sequences:** write a short account of the 'clumsy robot' experience, highlighting moments of improvement or confusion.
  - **Metalinguistic reflection:** discuss how the words 'right' and 'wrong' change meaning according to context.

- **Mathematics**
  - **Recording data:** record the robot's attempts and successes in a table and represent them in a bar or line graph.
  - **Basic probability and statistics:** analyse whether the robot improves by chance or through learning; estimate success rates.
  - **Sequences and logic:** turn the experience into instructions or pseudocode (if → then).

- **Science**
  - **Animal and human learning:** compare machine reinforcement learning with learning in living beings, such as training a dog or learning to ride a bicycle.
  - **Basic neuroscience:** introduce reward and how the brain responds to success or error.
  - **Scientific simulation:** design similar experiments to observe how behaviour changes after repeated attempts.

Learning situation 3:

TRAIN THE CLUMSY ROBOT

- **Computing and Computational Thinking**
  - **Simple algorithms:** translate the robot's sequence of actions into coded steps (if → then → repeat).
  - **Learning models:** compare this experience with the three types of machine learning: supervised, unsupervised and reinforcement learning.
  - **Model evaluation:** represent the robot's 'learning curve' and discuss how to measure an AI system's improvement.

- **Social and Civic Values**
  - **Ethics of artificial learning:** reflect on human responsibility when training AI. What happens if feedback is unfair or confusing?
  - **Empathy and frustration:** discuss how the robot felt when it received only 'wrong', and relate this to human empathy and motivation.
  - **Collaborative work:** analyse how the group can act more fairly and in a more coordinated way to help the robot improve.

Final message

This activity is a living metaphor for reinforcement learning, one of the foundations of many modern artificial intelligence systems.
The 'clumsy robot' shows pupils that learning is not about knowing from the start, but about trying, making mistakes and adjusting behaviour in response to feedback.

By experiencing it first-hand, pupils understand three fundamental ideas:

1. **Machines do not think independently.** They learn only when someone—a programmer or user—indicates whether their actions are right or wrong.
2. **Mistakes are part of learning.** For both humans and AI, progress arises through correction and positive reinforcement.
3. **The human role is essential.** AI cannot learn ethically without supervision or guidance. People decide what counts as 'right' and 'wrong', and this shapes the system's behaviour.

This conclusion links to cross-curricular themes including **technology ethics, the importance of feedback in human** learning and the value of **collaboration**. Pupils discover that teaching a machine also **teaches us about how we ourselves learn**.

blink Learning

ASSESSMENT OF THE TEACHING PROGRAMME UNIT

General approach

Assessment is conceived as a continuous, formative and competency-based process focused on active participation, conceptual understanding and the practical application of learning.
Technical accuracy is not assessed; instead, emphasis is placed on pupils' ability to observe, infer, communicate, collaborate and reflect on what they have learnt.
Assessment relies mainly on qualitative observation tools—classroom journals, checklists, rating scales and anecdotal records—supplemented by performance outcomes.
A selection of assessment tools is suggested below for teachers to personalise according to their professional judgement.

Subject-specific assessment criteria (adapted from Decree 61/2022 – Community of Madrid)

Official policy	Practical application in the unit
1. Take part in projects that involve problem-solving through computational thinking.	Take an active part in the simulations and suggest strategies as a group during the three learning scenarios.
2. To analyse the effects of technology on society and on everyday life, whilst demonstrating responsible and ethical behaviour.	It explains the role of humans in machine learning and reflects on the fair use of feedback and data.
3. Use experimentation and observation to understand how machines learn from data.	He formulates hypotheses, tests them and revises them as part of the 'mystery model' and 'clumsy robot' process.
4. Work effectively as part of a group, communicating ideas and listening to those of others.	Help with decision-making for the collaborative story, classifying examples and training the robot.

blink Learning

ASSESSMENT OF THE TEACHING PROGRAMME UNIT

Observation tools

1**. Checklist (daily use)**
The teacher records the frequency of, or evidence for, behaviours observed during each session.

Observable indicator	Yes	Sometimes	No
Take an active part in the proposed activity.	☐	☐	☐
Put forward ideas or hypotheses to solve the challenge.	☐	☐	☐
Listen to the group's views and take turns.	☐	☐	☐
Explain in your own words what is happening in the simulation.	☐	☐	☐
Identify errors and suggest improvements.	☐	☐	☐
He links his personal experiences to machine learning.	☐	☐	☐

Complete after each session or at the end of the unit. It may be shared by teachers across different subject areas.

ASSESSMENT OF THE TEACHING PROGRAMME UNIT

2. Rating scale (overall performance)

Performance level	Description of learning outcomes
Excellent (4)	Demonstrates a clear understanding of AI processes (prediction, learning, reinforcement). Explains their ideas accurately and collaborates actively with the group. Links what they have learnt to real-world examples.
Suitable (3)	Take an active part and understand the basic concepts. Apply the rules of the activities correctly and show curiosity and thoughtfulness.
In progress (2)	They participate with some support, identify the main processes, but find it difficult to explain them clearly or to hold the group's attention.
Initial (1)	Shows little engagement with or understanding of the purpose of the activities. Requires constant guidance and verbal or visual support.

This scale may be used as a general rubric or as a qualitative summary in the assessment report.

ASSESSMENT OF THE TEACHING PROGRAMME UNIT

3. Anecdotal record (qualitative observation)

The teacher briefly records significant episodes that demonstrate learning, cooperation or reasoning:

- 'Clara linked the robot's error to mistakes apps make when they do not understand speech correctly.'
- 'Mario's group decided to use chance to break ties and then discussed whether it was fair.'
- 'During the story, Nuria suggested limiting the word options, just as an AI would limit its data.'

These observations are very useful for competency-based assessment and teacher self-evaluation.

Evidence of performance by learning situation

Learning situation	Observable evidence	Performance product
The Sad Dragon and Text Prediction	They take part in decision-making, understand the concept of 'prediction' and recognise patterns in language.	Collective story + record of decisions.
Train the mysterious model	They formulate hypotheses, review their ideas in the face of error, and understand that the data determine the results.	Table of hypotheses and oral/written reflection.
Train the clumsy robot	Understands the role of feedback and works independently to improve their work.	Recording of training sessions + final display: 'This is how machines learn'.

blink Learning

Empower the present, transform your school's future


blink
Learning
www.blinklearning.com