Pearson

EBLANA LEARNING

AI Workbook

for IB DP and CP Students

AI for the IB Innovator







FREE SAMPLE

Rita Bateson

1 What is AI and how it works

Analyze how AI uses mathematical models and training data to recognize patterns and make probabilistic claims about our world.

What this unit is about

This unit explores the architectural foundations of AI, specifically focusing on **Large Language Models (LLMs).** You will define how **algorithms** (sets of step-by-step rules) use **training data** to build mathematical **models.** will investigate **pattern recognition** and how machines generate **predictions** based on **probability** (the likelihood of an event occurring) rather than 'understanding' in a human sense.

Why it matters

When you use a search engine or a revision app, you are interacting with **probabilistic reasoning**. Understanding these systems helps you judge the **reliability** and **accuracy** of the information they provide. Recognizing **statistical bias** (errors caused by unrepresentative data) allows you to use these tools more ethically and effectively in your daily life.

How you'll learn

You will adopt a dual role as both a scientist and a mathematician. You will use **empirical validation** (testing claims through observation) to see if AI outputs are consistent. By testing the **falsifiability** (the ability for a claim to be proven wrong) of AI 'knowledge', you will develop a critical eye for how these **models generalize** information from limited examples.

IB Connections

TOK Areas of Knowledge: Natural Sciences and Mathematics - You will examine how AI mimics the scientific method by identifying patterns in data to form 'laws' or **models**. You will also analyze the mathematical logic of **neural networks** (computational systems inspired by the brain) as a tool for producing knowledge.

ATL Skills: Thinking Skills and Research Skills – You will evaluate the limitations of AI **models** and analyze how data selection impacts the **model's reliability**. You will practice **media literacy** by verifying claims made by generative systems.

TOK Knowledge Questions: To what extent can a mathematical **model** provide us with 'truth'? How do the limitations of **training data** affect the objectivity of knowledge produced by an **algorithm**?

Sample content - uncorrected proof

AI_Workbook_IB_DP_Sample_v8.indd 1 09/09/2026 08:40

1.2 Key terminology

Knowledge question

To what extent does the use of algorithmic knowledge challenge the role of the human 'knower' in justifying truth claims?

Vocabulary exploration: The epistemology of AI

This table will introduce you to the key vocabulary of this unit.

Term and definition	TOK Knowledge question	Prompts
Algorithm: A precise, step-by-step set of rules or instructions followed to solve a specific problem or complete a task.	Can a mathematical rule be considered 'objective' if it is designed by a biased human?	**(AO1)** State one example of an **algorithm** you interact with daily. **(AO2)** Explain how an **algorithm** converts raw data into information. **(AO3)** Discuss whether algorithmic logic is superior to human intuition in decision-making.
Bias: A distortion in data or a programmed preference that results in unfair or unrepresentative outcomes (systemic error).	How does the digital divide influence the 'perspectives' embedded within AI systems?	**(AO1)** Define cognitive **bias** in the context of data collection. **(AO2)** Analyze how historical **bias** in **training data** can reinforce social inequalities. **(AO3)** Evaluate the extent to which a completely 'neutral' AI is possible.
Neural network: A computational **model** inspired by the structure of biological brains, using layers of interconnected nodes to process complex information	If we cannot trace *how* a network reached a conclusion, is that conclusion still 'knowledge'?	**(AO1)** Identify the biological inspiration for **neural networks**. **(AO2)** Describe the relationship between 'hidden layers' and the loss of transparency. **(AO3)** To what extent does the 'black box' nature of AI invalidate its status as a reliable authority?
Probabilistic reasoning: A way of thinking where conclusions are based on the likelihood (chance) of something being true, rather than absolute certainty.	Does 'high **probability**' constitute the same level of justification as 'logical proof'?	**(AO1)** Outline the difference between certainty and **probability**. **(AO2)** Compare how an AI uses **probability** versus how a human uses 'common sense'. **(AO3)** Justify the use of **probabilistic** AI in high-stakes fields like medicine or law.
Generalization: The ability of a **model** to apply what it learned from **training data** to new, unseen situations correctly.	At what point does a **generalization** become a dangerous oversimplification?	**(AO1)** Label a scenario where an AI fails to generalize (overfitting). **(AO2)** Explain why **generalization** is essential for inductive reasoning. **(AO3)** Critically assess the ethical risks of AI generalizing human cultural behaviors.

Sample content - uncorrected proof

AI_Workbook_IB_DP_Sample_v8.indd 4 09/09/2026 08:40

1.7 The Evolution of Machine Logic

Thinking skills | Natural Sciences | Mathematics

Knowledge question

To what extent does the shift from rule-based logic to probabilistic reasoning change our definition of 'intelligence'?

Imagine your music app's 'Discover' playlist or a gaming Non-Player Character that anticipates your movement. These tools do not follow a fixed script written by a human. They rely on **probabilistic reasoning**. AI began very differently. Early scientists tried to build 'expert systems' using **algorithms**. They believed intelligence was like a giant library of 'if-then' statements. The computer followed a rule, for example, 'If the patient has a fever, then check for infection'. However, this approach failed because the world is too messy for perfect rules.

In the late 20th century, the focus shifted toward **neural networks** (computer systems modeled loosely on the human brain's structure). Rather than being told the rules, these systems look at massive amounts of **training data** to find their own patterns. This is called **pattern recognition**. For example, by looking at millions of student essays, a **model** learns which words usually follow each other. It doesn't 'understand' the essay, but it uses **prediction** to suggest the next word with high **accuracy**.

This modern AI makes its **knowledge claims** based on **empirical validation**. However, these **models** have limits. If the **training data** is poor, the AI develops **bias**. Unlike a math formula, an AI's output isn't always 100% certain. It depends on **reliability** over time. In the Natural Sciences, theories require **falsifiability**. Because modern AI is so complex, it can be hard to 'falsify' or explain why it made a specific choice, leading scholars to call it a 'black box'.

Factual Recall (AO1)

Review your learning so far and answer the questions.

1 **Identify the primary difference between early 'expert systems' and modern neural networks.**

2 **State two components that AI uses to make knowledge claims according to the text.**

Sample content - uncorrected proof

AI_Workbook_IB_DP_Sample_v8.indd 14 09/09/2026 08:40

Conceptual Analysis (AO2)

3 Explain how **probabilistic reasoning** differs from fixed algorithmic rules using your social media feed as an example.

4 Analyze the concept of **reliability**. Why might a **model** be consistent but still produce **biased** results?

5 Distinguish between **accuracy** and **empirical validation** in the development of a new AI tool.

Evaluative Discussion (AO3)

6 'If a scientific **model** is not falsifiable, it cannot provide certain knowledge.' To what extent does this statement challenge our trust in AI-driven science?

7 Evaluate the impact of **bias** on the global community. How might **training data** from one culture affect the **accuracy** of an AI used in a different cultural context?

Reflect

Skills Practised:

☐ Thinking ☐ Communication ☐ Social ☐ Self-Management ☐ Research

Select one ATL skill you used today and describe how it helped you navigate the complexity of this AI concept

Sample content - uncorrected proof

AI_Workbook_IB_DP_Sample_v8.indd 15 09/09/2026 08:40

1.10 Machine learning and the anatomy of knowing

Thinking skills | Natural Sciences | Mathematics

Knowledge question

To what extent is the 'knowledge' produced by a probabilistic model comparable to human understanding?

Machine learning: Learning from experience

Most modern AI relies on **Machine Learning (ML)**. In this framework, we can think of the learning process through two primary components:

- The data is the textbook: ML **models** are trained on massive datasets. For example, to teach an AI to identify a cat, you provide millions of images labeled 'cat' or 'not cat'. This is **supervised learning** (a process where every input has a pre-defined correct answer).
- The algorithm is the study method: The **algorithm** (a set of mathematical rules) searches for statistical patterns. It performs complex calculations across internal layers, adjusting its **weights** (the relative importance assigned to specific data features).

The goal: optimization

The **model** adjusts its **weights** until its **predictions** match the correct answers as closely as possible. This is a mathematical process of refinement.

Deep learning: A brain-inspired approach

Deep Learning utilizes artificial **neural networks** (computational systems inspired by the interconnected neurons in a biological brain). These networks feature multiple layers between the initial input and the final output:

- Layer 1: Might detect simple edges and lines.
- Intermediate layers: Combine edges to identify complex shapes.
- Final layers: Combine shapes to recognize high-level concepts, such as a human face.

This organized layering allows **Large Language Models (LLMs)** to process the underlying structure and statistical relationships within human language.

The essential understanding: Prediction *vs* knowing

The power and limitation of AI reside here: it does not 'know' in the human sense. Instead, it excels at **statistical prediction**.

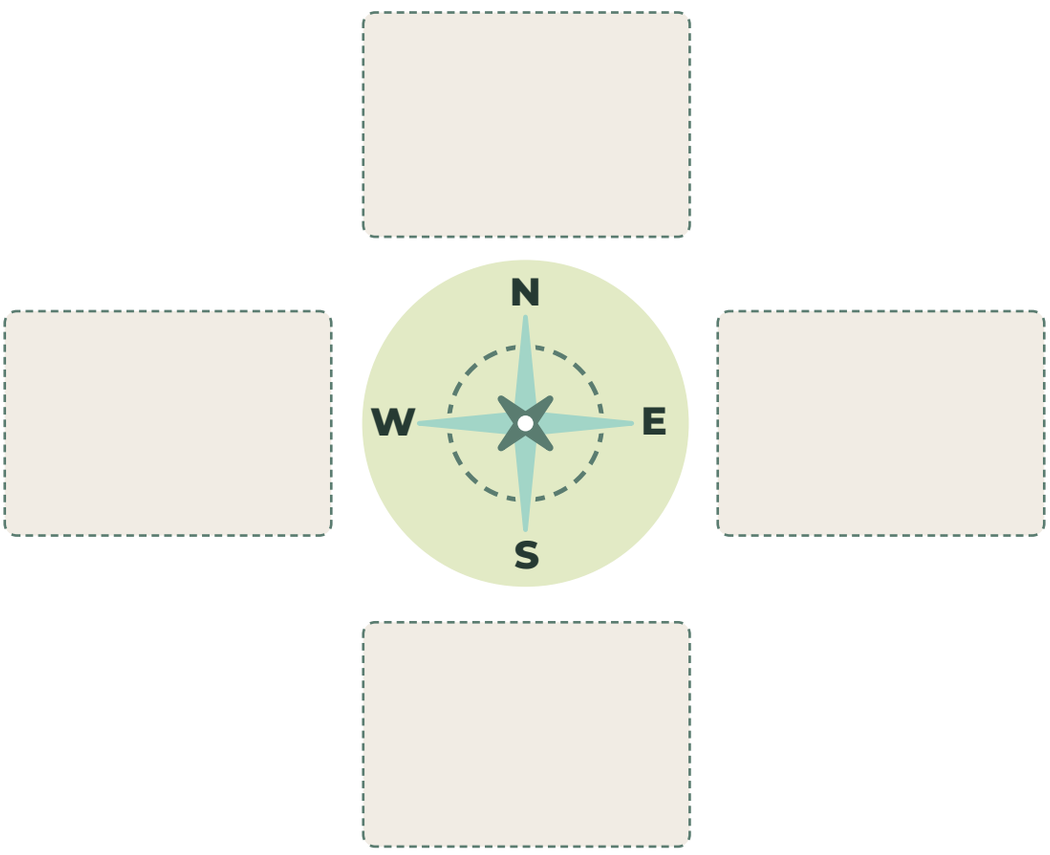
An AI generates a compelling essay because it has calculated the most statistically probable next word in a sequence. Its 'intelligence' reflects patterns in its **training data**. Consequently, AI can appear creative but remains prone to **hallucinations** or systematic **bias**.

Sample content - uncorrected proof

AI_Workbook_IB_DP_Sample_v8.indd 20 09/09/2026 08:40

Compass points activity

1 Reflect on the mechanics of **Machine Learning** described on the previous page. For each compass point, use the guidance below to describe your perspective on AI as a producer of knowledge.



- **NORTH:** Need to know
  What more do we need to know about the **training data** before we can trust an AI's claim?
- **SOUTH:** Stance
  Does **probabilistic reasoning** count as a **justified** way of knowing? Why/why not?
- **EAST:** Excitements
  What are the strengths of using **neural networks** to find patterns humans might miss?
- **WEST:** Worries
  What are the risks of relying on **statistical probability** instead of empirical truth?

Reflect

Skills Practised:

☐ Thinking ☐ Communication ☐ Social ☐ Self-Management ☐ Research

Select one ATL skill you used today and describe how it helped you navigate the complexity of this AI concept.

Sample content - uncorrected proof

AI_Workbook_IB_DP_Sample_v8.indd 21 09/09/2026 08:40

1.12 Building an admissions bot

Thinking skills | Mathematics | Human Sciences

Knowledge question

To what extent can mathematical **models** provide a reliable basis for ethical decision-making in human systems?

In this activity we will manually 'train' a decision tree **algorithm** (a flowchart-like **model** used for **classification**) by identifying which data features provide the highest **information gain**.

The concept: Entropy *vs* purity

AI does not select questions at random. It calculates the feature that splits the data into the 'cleanest' possible groups.

- **High entropy:** A diverse, 'messy' mixture of both 'accepted' and 'rejected' outcomes.
- **Low entropy (purity):** A group that achieves a singular outcome (e.g. 100% 'accepted').

Analyzing the training data

1 Below is a sample of eight students who applied to a selective university last year. Identify the 'root attribute' (the initial question) that best predicts the result.

Student	GPA > 3.8	Extracurriculars	Leadership Role	Result
1	Yes	Yes	Yes	Accepted
2	Yes	No	No	Accepted
3	Yes	Yes	No	Accepted
4	No	Yes	Yes	Accepted
5	No	No	Yes	Rejected
6	No	Yes	No	Rejected
7	No	No	No	Rejected
8	Yes	No	Yes	Accepted

Root atribute: .....

.....

.....

Sample content - uncorrected proof

AI_Workbook_IB_DP_Sample_v8.indd 24 09/09/2026 08:40

Calculating statistical utility

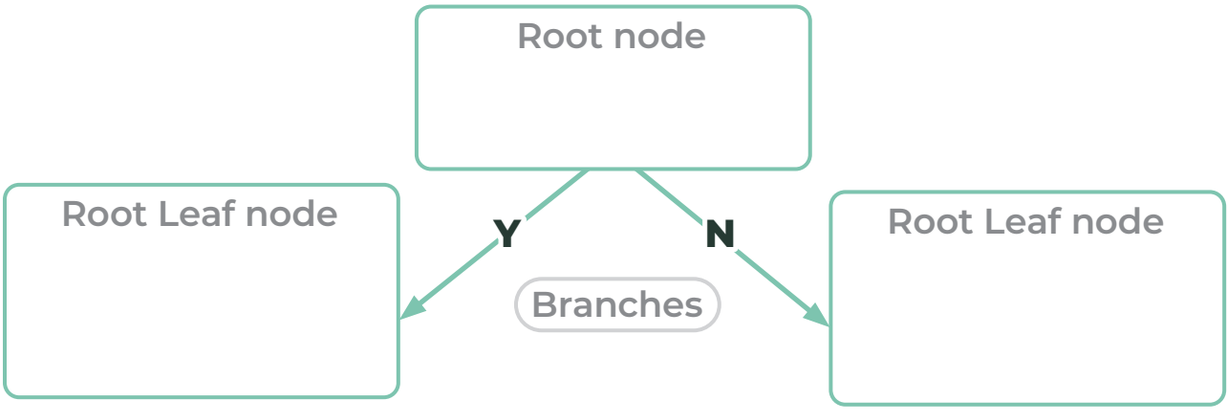
2 Analyze the data on the previous page to fill in the counts for each feature.

Feature	How many accepted?	How many rejected?	Low Entropy? (Yes/No)
GPA > 3.8			
Extracurriculars			
Leadership role			

Building the logic tree

3 Construct your decision tree. Start with the 'cleanest' feature at the top.

- Root node: represents the question with the highest 'accepted' purity in one branch.
- Branches: Two lines labeled 'Yes' and 'No'.
- Leaf nodes: The final outcomes. If a branch is still mixed, you must add another question box!



Empirical validation

4 Apply your completed tree to predict results for new unseen applicants.

Applicant A: GPA 3.9, No Extracurriculars, No Leadership. Prediction: _____

Applicant B: GPA 3.2, Yes Extracurriculars, Yes Leadership. Prediction: _____

Reflect

Skills Practised:

☐ Thinking ☐ Communication ☐ Social ☐ Self-Management ☐ Research

Sample content - uncorrected proof

AI_Workbook_IB_DP_Sample_v8.indd 25 09/09/2026 08:40