



AI Guide

ONLINE COURSE CONTENT

MODULE 1: BASIC CONCEPTS OF AI

2025-1-LT01-KA220-HED-000354784

Artificial Intelligence for Green Understanding, Insights & Decision Empowerment



CAM CONSULTING



SQUAREDEV



Co-funded by
the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the National Agency of Erasmus+ Programme in Lithuania – Education Exchanges Support Foundation (ŠMPF). Neither the European Union nor the ŠMPF can be held responsible for them.

Table of Contents

Basic concepts in AI	2
A. General Information	3
B. Content for STEM HE Teachers.....	3
Highlights.....	3
Subtopic Content	3
Green Decision Context	8
Real World Examples.....	9
Example Projects	9
Recommended Readings.....	9
Videos/Interactive Resources	9
References.....	9
C. Assessment and Practice Activities.....	10
Suggested 60-Minute Sequence	10



Basic concepts in AI

A. General Information

Module No	Module 1
Subtopic	Basic concepts in AI
AI Literacy Component	Core concepts of AI; AI basic terminology; Ability to distinguish between human and AI-generated content
Decision Making Skills	Identify problems; recognise where AI tools might be helpful
Learning Objectives for STEM HE Teachers	Explain fundamental AI concepts and the distinction between AI and non-AI in STEM education
Learning Objectives for STEM Students	Be able to identify whether the use of AI is necessary
Suggested AI Tools	ChatGPT; Gemini for Education; Microsoft Copilot Chat; Claude for Education
Estimated Duration	60 minutes

B. Content for STEM HE Teachers

Highlights

After this subtopic, teachers should emphasise:

- AI is a data-driven pattern-recognition system, not a thinking entity
- AI outputs are probabilistic, not absolute truths
- Not every digital or “smart” tool is actually an AI system
- Distinguishing human-generated from AI-generated content is now a core AI literacy skill
- AI can support environmental decision-making but cannot replace human judgment
- Understanding these basics is necessary before evaluating AI’s role in green decisions

Subtopic Content

Introduction

This content pack introduces the foundational vocabulary and concepts of Artificial Intelligence that STEM higher education teachers need before exploring more advanced applications later in this module. Rather than treating AI as an abstract or purely technical subject, this introduction grounds each concept in real-world and environmental sustainability examples, drawing on established academic and technical sources, so that teachers can build conceptual AI literacy in their students while situating that literacy within socio-scientific environmental decision-making. By the end of this subtopic, teachers should be able to define AI in plain language, use its basic terminology correctly, distinguish AI systems from non-AI systems and AI-generated content from human-generated content, and help students recognise situations in which AI tools might genuinely be useful.

Explore: [How AI Works](#) – Code.org’s free interactive video series companion to this subtopic.

What Is Artificial Intelligence?

There is no single, universally agreed definition of artificial intelligence, and it is worth acknowledging this directly rather than hiding behind one tidy sentence. A foundational 1950 paper sidestepped the definitional question entirely, proposing instead a practical test: rather than asking whether a machine can “think,” it asked whether a machine could produce responses indistinguishable from a human’s in conversation (Turing, 1950). The many competing definitions of AI put forward since can be organized into four families: systems that think like humans, systems that act like humans, systems that think rationally, and systems that act rationally, with most modern AI research pursuing the last of these, building systems that act to achieve good outcomes regardless of whether their internal processing resembles human cognition (Russell & Norvig, 2020). In more applied terms, artificial intelligence can be described as technology that enables computers



and machines to “simulate human learning, comprehension, problem solving, decision making, creativity and autonomy” (IBM, 2024). For the purposes of this module, AI can be understood as computer-based systems designed to perform tasks that typically require human intelligence, such as recognising patterns, making predictions, or generating language and images, while making clear to students that these systems do not “think” or “understand” in a human sense. Instead, they process large amounts of data and produce outputs based on statistical patterns identified during a process called training.

Watch: [“What Is Artificial Intelligence? #1” \(Crash Course, PBS\)](#) for a short visual introduction to this definition.

The term “artificial intelligence” itself dates to 1956, when it was coined at a research workshop that set out to explore whether machines could simulate aspects of human reasoning (McCarthy, Minsky, Rochester, & Shannon, 1955/2006). Early AI systems were largely rule-based: programmers wrote explicit instructions (“if X happens, do Y”) for every situation the system might encounter. Over the following decades, the field shifted toward machine learning, in which systems are instead shown large amounts of example data and learn their own rules, or patterns, from that data. Most of the AI tools teachers and students encounter today, including chatbots, image generators, and recommendation systems, are built on this data-driven, learning-based approach rather than on hand-written rules.

Understanding this shift matters for AI literacy because it explains why modern AI behaves the way it does: its outputs are shaped by the data it was trained on, not by a fixed set of instructions written by a programmer. This is also why AI can generalise to new situations it has never seen before, and why it can also make mistakes or reflect biases that were present in its training data, a risk documented empirically in a widely cited audit of commercial facial-analysis systems, discussed further below (Buolamwini & Gebru, 2018).

Core AI Terminology

Basic AI terminology helps teachers and students communicate precisely about these systems, and the same handful of terms recur throughout this module (Russell & Norvig, 2020).

- An algorithm is a defined procedure or set of rules that a computer follows to solve a problem or process data.
- A model is the specific system that results from training an algorithm on data. It is the model, not the raw algorithm, that is actually used to make predictions or generate content.
- Training data is the collection of examples, such as images, text, or sensor readings, used to teach a model to recognise patterns.
- Parameters are the internal values a model adjusts during training to improve its performance, and
- Inference is the process of applying an already-trained model to new, previously unseen data to produce an output.

Try: [Teachable Machine \(Google\)](#) to train a simple image classifier hands-on and see these terms in action.

An environmental example helps connect these terms: to build a flood-risk prediction system, developers might use an algorithm designed for time-series forecasting, train it on historical rainfall and river-level data (the training data), and let it adjust its internal parameters until it can reliably associate certain rainfall patterns with flooding events. The result is a trained model. When new weather data comes in during an actual storm, the model performs inference, producing an output such as “72% probability of flooding in the river basin within 24 hours”. Recognising this pipeline - data in, training, model, inference, output - helps teachers explain to students what is actually happening inside an AI system, rather than treating it as an unexplainable black box.

Distinguishing AI from Non-AI Systems

A central AI literacy skill is being able to tell whether a given tool is actually AI or simply automation. The key question, consistent with the “acting rationally” framing of AI (Russell & Norvig, 2020), is whether the system learns and adapts its behaviour from data, or whether it only executes fixed, pre-written rules. A household thermostat that switches on the heating whenever the temperature drops below a set point is automation: it follows one fixed rule and never changes its behaviour based on experience. A so-called “smart” thermostat that observes a household’s daily routines over several weeks and begins adjusting the temperature in anticipation of when people wake up or leave the house is exercising a basic form of AI, because its behaviour changes based on patterns it has learned from data.



The same distinction applies to digital tools more broadly. A spreadsheet formula that adds a column of numbers is not AI, no matter how sophisticated the spreadsheet software is marketed to be, but it rather performs a fixed calculation every time. A system that reviews thousands of past transactions and learns to flag transactions that look unusual, without being told in advance exactly what “unusual” means, is exercising AI, because it has derived its own criteria from data. This distinction is important because many consumer products now use the word “AI” or “smart” as a marketing label even when the underlying technology is simple automation. Teaching students to ask “does this system learn from data, or does it just follow fixed instructions?” gives them a durable, technology-independent way to evaluate any tool they encounter, now or in the future.

Human-Generated versus AI-Generated Content

As generative AI tools have become widely available, a second and closely related literacy skill has emerged: the ability to distinguish human-generated from AI-generated content. AI systems can now produce text, images, audio, and video that closely resemble human-created work, which creates both opportunities (drafting assistance, accessibility tools, rapid prototyping) and risks (misinformation, academic integrity concerns, and synthetic media such as deepfakes). This concern is one of the reasons media and information literacy, including the ability to critically evaluate AI-generated content, has been identified as a core competency education systems need to build going forward (UNESCO, 2021).

Rather than relying on any single foolproof test, which quickly becomes outdated as AI tools improve, teachers should model a habit of critical verification. Useful strategies include:

- checking claims against independent, credible sources
- looking for internal inconsistencies or generic, overly uniform phrasing in text
- being alert to subtle visual artefacts in images (unnatural textures, inconsistent lighting, or distorted fine details), and
- asking, when authorship is unclear, whether the source of the content can be verified at all.


For STEM higher education teachers, the goal is not to make students suspicious of all digital content, but to build a habit of verification that applies equally to AI-generated and human-generated claims, both of which can be inaccurate or misleading.

AI as a Decision-Support Tool in Environmental Contexts

In educational and environmental contexts, AI is increasingly used as a decision-support tool rather than a decision-maker. Several examples illustrate this role:

- Traffic-management systems analyse real-time vehicle and sensor data to recommend route changes or restrictions that reduce congestion and emissions.
- Weather and climate models process satellite imagery, atmospheric pressure, and historical records to forecast conditions such as rainfall, heatwaves, or storm paths, supporting disaster preparedness.
- Environmental monitoring networks use sensor data to detect changes in air quality, water quality, or forest cover, flagging areas that may need closer human investigation.
- Smart energy grids analyse consumption patterns to optimise the balance between renewable and conventional energy sources in real time.
- Precision-agriculture systems analyse soil-moisture, weather, and satellite data to recommend irrigation or fertiliser timing, helping reduce water and chemical overuse.
- Wildlife-conservation systems analyse camera-trap and acoustic sensor data to flag likely poaching activity for park rangers to investigate.





AI supports the decision; a human remains responsible for it.

In every one of these examples, the AI system analyses data and generates a recommendation or forecast, but a human, a city planner, an environmental scientist, a grid operator, an agricultural advisor, or a park ranger, ultimately decides what action to take, taking into account context the AI system cannot fully capture, such as political feasibility, community impact, or ethical trade-offs. This division of labour between AI-generated recommendations and human-held decision authority reflects the principle that “human agency and oversight” should be preserved throughout any AI-supported decision process (European Commission High-Level Expert Group on AI, 2019), and is the central idea STEM higher education teachers should reinforce throughout this module.


Problem Identification and Recognising When AI Helps

Before reaching for an AI tool, a critical decision-making skill is identifying the problem clearly and asking whether AI is genuinely needed to solve it. A simple framework can guide this thinking:

1. First, state the problem precisely, rather than describing it vaguely.
2. Second, identify what data is actually available to address it.
3. Third, ask whether the task fundamentally requires recognising complex patterns or making predictions from that data, something AI is suited for, or whether it could be solved with a simple rule, calculation, or lookup table, in which case AI adds unnecessary complexity, and
4. Fourth, consider the stakes of an incorrect output, since higher-stakes problems demand more scrutiny of an AI system’s reliability before its recommendations are acted upon.



A simple framework for deciding whether a task genuinely needs AI



The four-step framework for deciding whether a task genuinely needs AI, from stating the problem through weighing the stakes of a wrong answer.

Applying this framework to a single classroom scenario makes each of the four stages concrete. Suppose a STEM higher education teacher wants to identify students who are at risk of falling behind in a course:

- Step 1 – State the problem: instead of a vague goal such as “help my struggling students,” the teacher states the problem precisely: which enrolled students are at meaningful risk of falling behind this semester, early enough that something can still be done about it?
- Step 2 – Identify the data: the teacher lists what is actually available – attendance records, quiz scores over time, participation logs, and engagement data from the learning-management system.
- Step 3 – Assess the task: if the goal is simply to see who scored lowest on the last quiz, a sorted average answers that without any AI at all. But combining attendance, quiz trends, participation, and engagement into a single early-warning signal is a pattern-recognition task better suited to a predictive AI tool.
- Step 4 – Weigh the stakes: because a wrong flag could mean a student is unfairly labelled “at risk,” or, just as seriously, that a genuinely struggling student is missed, the teacher treats the tool’s output as a prompt for a conversation, not an automatic label, and checks its recommendations against her own knowledge of the class before acting.

Teaching students to work through this kind of framework, rather than assuming AI is always the best or only solution, is itself a core decision-making skill emphasised throughout this module.

Limitations: Probabilistic Outputs and Bias

A key limitation of AI systems is that their outputs are probabilistic rather than absolute truths. This means AI can suggest likely outcomes but cannot guarantee accuracy, and its confidence in a given output does not always correspond to how correct that output actually is. In addition, AI systems can reflect and even amplify biases present in the data they are trained on. This is not a hypothetical risk: a widely cited 2018 audit of three commercial gender-classification systems found error rates below 1% for lighter-skinned men but as high as 34.7% for darker-skinned women, a gap traced directly back to the underlying training data (Buolamwini & Gebru, 2018). The same dynamic applies in environmental contexts: if historical data contains skewed patterns, for example under-monitoring air quality in lower-income neighbourhoods, a model trained on that data may under-predict pollution risks in exactly those areas. Bias of this kind is not confined to computer vision: a critical review of fairness evaluations for large language model-based recommendation systems found that most existing fairness frameworks still overlook personalisation, meaning a system's outputs can look reasonable in aggregate while still under-serving specific groups of users (Sah et al., 2023).

Watch: [AI: Training Data and Bias](#) (Khan Academy / Code.org) for a short walkthrough of exactly this effect.



Therefore, understanding how AI works at a basic level is essential for critically evaluating its role in environmental and educational decision-making. In the context of sustainability, AI has significant potential to support green decision-making, such as optimising energy use, predicting environmental risks, and monitoring pollution levels. However, AI should always be seen as a supportive tool that complements, not replaces, human judgment, especially in complex socio-scientific issues where the costs of an error are unevenly distributed across communities.

Common Misconceptions About AI

As AI has become an everyday topic in the media and in classrooms, several misconceptions have become common, and addressing them directly strengthens the basic literacy this subtopic aims to build. One misconception is that AI is a single, unified intelligence rather than a large collection of separate, narrowly specialised systems built using several distinct approaches (Russell & Norvig, 2020). A chatbot, a flood-prediction model, and an image classifier share underlying techniques, but they are entirely separate systems, each trained for its own purpose, with no shared awareness between them. A second misconception is that AI "understands" the content it processes in the way a human does. In reality, even highly capable systems are manipulating statistical patterns in data, with no awareness of meaning, consequence, or context beyond what can be inferred from that data. A third misconception is that a more advanced or more expensive AI tool is automatically more accurate. In practice, accuracy depends heavily on how well-matched a tool's training data is to the specific task at hand, so a simple, purpose-built model can outperform a large general-purpose one on a narrow problem. Correcting these misconceptions early gives teachers and students a more realistic, and ultimately more useful, mental model of what AI tools can and cannot do.

Data Quality and Its Role in Trustworthy AI

Every example in this subtopic, from flood prediction to pollution monitoring, depends on the quality of the underlying data, and this dependency is worth making explicit. Data quality refers to how accurate, complete, representative, and up-to-date the information used to train and run an AI system actually is. A flood-prediction model trained only on data from wet regions may perform poorly if applied to a historically dry region now experiencing unusual rainfall due to a changing climate. An air-quality monitoring network with sensors concentrated in wealthier neighbourhoods will produce an incomplete picture of pollution citywide, regardless of how sophisticated its underlying AI model is. This dependency directly reflects the broader principle that AI systems should be evaluated for "technical robustness and safety" throughout their lifecycle, including the quality of the data that feeds them (European Commission High-Level Expert Group on AI, 2019). Because of this dependency, evaluating an AI tool responsibly means asking not only "how does this system work?" but also "what data was it built on, and does that data reasonably represent the situation I am applying it to?" This question reappears throughout the module as a practical checkpoint for STEM higher education teachers introducing any AI-based tool to their students. This concern sits within a broader shift in AI research toward trustworthiness as a standard that goes beyond accuracy alone, with researchers arguing that systems should also be assessed for fairness, robustness, privacy, and explainability before being deployed in sensitive decision-making contexts (Wang et al., 2023).

These basic concepts, definitions, terminology, the AI versus non-AI distinction, decision-support framing, problem identification, and awareness of limitations, feed directly into everything that follows in this module. Understanding these basics is the necessary first step before comparing the different types of AI, the tools built on them, and the chatbots students and teachers are most likely to use directly, all of which are addressed in the sections that follow.

Green Decision Context

AI is increasingly used as a decision-support tool in green and environmental contexts, especially in:

- monitoring air, water, and soil quality using sensor data
- forecasting climate risks such as floods, heatwaves, and droughts
- optimizing household and building energy use
- reducing food and material waste through demand and supply prediction
- supporting biodiversity and wildlife-protection efforts

However:

- AI does NOT "understand" sustainability or environmental values
- It identifies statistical patterns in the data it is given
- Human experts and decision-makers remain responsible for interpreting AI outputs and taking final



action

Real World Examples

- **[Winnow](#)** – **AI Food Waste Prevention in Commercial Kitchens:** uses computer vision and smart scales to automatically track discarded food in commercial kitchens, converting waste into measurable data linked to purchasing and menu decisions. It gives students a simple, concrete link between AI-generated data and everyday green decisions, without requiring an advanced technical background.
- **[Voltaware](#)** – **AI Household Energy Intelligence for SME Utilities:** uses a single sensor and machine learning to disaggregate whole-home electricity readings into individual appliance use, illustrating how AI learns patterns from data, rather than following fixed, pre-written rules, to support household energy-saving decisions.
- **[Google Flood Hub](#)** – **AI Flood Forecasting for Climate Adaptation:** combines machine learning and hydrological modelling to issue river flood forecasts up to seven days in advance, supporting preparedness and evacuation decisions. Forecasts remain probabilistic, and local authorities' verification and judgment stay essential, reinforcing the basic concept that AI outputs are likelihoods, not certainties.

Example Projects

- **[AI for Everyone: Bridging the Gap Between Smart Technologies and Society](#)** (2025-2027): builds a MOOC and open educational resources that help general audiences critically understand and responsibly adopt AI technologies, modelling the kind of foundational AI literacy this subtopic introduces.
- **[DigiReady 2.0: Advancing Digital Readiness Across Europe's Universities](#)** (2025-2028): integrates AI-driven data extraction into a shared framework that helps universities assess their own digital readiness, illustrating basic AI concepts, such as automated data extraction and standardized indicators, applied at an institutional level.

Recommended Readings

- Russell, S., & Norvig, P. Artificial Intelligence: A Modern Approach (Pearson)
- UNESCO (2021). [AI and Education: Guidance for Policy-Makers](#).
- European Commission, High-Level Expert Group on AI (2019). [Ethics Guidelines for Trustworthy AI](#).

Videos/Interactive Resources

- **[“What Is Artificial Intelligence? #1” – Crash Course: Artificial Intelligence \(PBS\)](#)**: a short, accessible introduction to core AI concepts.
- **[Teachable Machine](#)** (Google): a free, interactive tool that lets teachers and students train a simple image classifier hands-on, directly reinforcing this subtopic's core concepts.
- **[AI: Training Data and Bias](#)** – Khan Academy / Code.org: a short video on how training data shapes, and can bias, an AI system's outputs.
- **[How AI Works](#)** – Code.org free interactive video series and lessons covering core AI concepts from machine learning to real-world applications.

References

- AIGUIDE. (2026). AIGUIDE WP2-A3 best practices report: AI tools for green decision-making in STEM higher education. Artificial Intelligence for Green Understanding, Insights & Decision Empowerment (2025-1-LT01-KA220-HED-000354784).
- Buolamwini, J., & Geburu, T. (2018). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. Proceedings of Machine Learning Research, 81, 77–91. Available from <https://proceedings.mlr.press/v81/buolamwini18a.html>
- European Commission, High-Level Expert Group on AI. (2019). [Ethics Guidelines for Trustworthy AI](#).
- IBM. (2024). [What Is Artificial Intelligence \(AI\)?](#) IBM Think.



- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2006). A proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955. *AI Magazine*, 27(4), 12. <https://doi.org/10.1609/aimag.v27i4.1904>
- Russell, S., & Norvig, P. (2020). *Artificial Intelligence: A Modern Approach* (4th ed.). Pearson.
- Sah, C. K., Xiaoli, L., & Islam, M. M. (2023). Unveiling Bias in Fairness Evaluations of Large Language Models: A Critical Literature Review of Music and Movie Recommendation Systems. *International Journal of Innovative Science and Research Technology*, 8(12), 2084–2089. <https://doi.org/10.5281/zenodo.10469839>
- Turing, A. M. (1950). Computing Machinery and Intelligence. *Mind*, 59(236), 433–460. <https://doi.org/10.1093/mind/LIX.236.433>
- UNESCO. (2021). [AI and Education: Guidance for Policy-Makers](#).
- Wang, S., Zhang, X., Wang, Y., Liu, H., & Ricci, F. (2023). Trustworthy Recommender Systems: An Overview. *ACM Transactions on Intelligent Systems and Technology*, 1(1), Article 1. <https://doi.org/10.1145/3627826>

C. Assessment and Practice Activities

The following activities give teachers concrete, measurable ways to check whether STEM higher education teachers can explain fundamental AI concepts and the AI-versus-non-AI distinction, and whether students can identify whether the use of AI is necessary for a given task, the two learning objectives stated for this subtopic. They are designed to fit within the subtopic’s 60-minute estimated duration alongside the reading above.

Suggested 60-Minute Sequence

- 0–10 minutes: introduce core definitions and terminology using the “What Is Artificial Intelligence?” and “Core AI Terminology” sections and the linked video.
- 10–20 minutes: guided reading and discussion of the AI-versus-non-AI distinction and common misconceptions.
- 20–35 minutes: Activity (a), the AI-or-automation sorting task.
- 35–50 minutes: Activity (b), the is-AI-needed decision-framework exercise, scored with the checklist in Activity (c).
- 50–60 minutes: formative quiz and self-reflection, Activity (d).

(a) AI-or-Automation Sorting Task

Give students a list of eight to ten familiar systems, for example a household thermostat with a fixed setpoint, a spam filter, a spreadsheet formula, a smart thermostat that learns household routines, a traffic-light timer, a music-streaming recommendation feed, a calculator, and a system that flags unusual banking transactions. Using the test introduced above, whether the system learns and adapts its behaviour from data or only executes fixed, pre-written rules, students sort each item into “AI” or “not AI” and write one sentence justifying each classification. As a measurable outcome, students submit their sorted list with justifications, directly assessing the learning objective of distinguishing AI from non-AI systems.

(b) Applying the Is-AI-Needed Decision Framework

Give students two short scenarios, one educational, for example identifying which students are falling behind based on quiz scores, and one environmental, for example deciding whether to install additional air-quality sensors in a neighbourhood, and have them work through the four-question framework introduced above: state the problem precisely, identify what data is available, decide whether the task needs pattern recognition or could be solved with a simple rule or calculation, and consider the stakes of an incorrect output. As a measurable outcome, students submit a short written recommendation for each scenario stating whether AI is genuinely needed and explaining their reasoning against the four criteria.

(c) AI-Literacy Concept Checklist

Teachers can use or adapt the following short checklist, scored yes / no / unsure, to assess a student’s explanation of an AI system or their classification in Activity (a):

- Does the explanation correctly distinguish whether the system learns from data or follows fixed rules?



- Are basic terms such as algorithm, model, training data, and inference used correctly and not interchangeably?
- Does the explanation avoid claiming the system “understands” or “thinks”, rather than processes statistical patterns?
- Is the possibility of bias in training data acknowledged where relevant?
- Would a data-quality problem, such as incomplete, unrepresentative, or outdated data, plausibly change the system’s output, and has that been considered?

(d) Formative Quiz and Self-Reflection

To close the lesson, students complete a short formative quiz aligned with the subtopic’s learning objectives, for example:

- In one sentence, explain the difference between a rule-based system and a system that uses machine learning.
- Give one example, not used in class, of a tool that is often marketed as “AI” or “smart” but may actually be simple automation, and explain why.
- Name one reason an AI system’s output might be inaccurate even if the system is working exactly as designed.
- Reflect in two or three sentences: describe one task from your own studies or daily life where you would now ask “does this genuinely need AI?” before reaching for an AI tool, and explain what you would check first.

Teachers can score the first three questions against a simple correct / partially correct / incorrect key and use the reflection question as a qualitative check that students can identify whether AI use is necessary in their own words, directly evidencing the student learning objective for this subtopic.

