



AI Guide

ONLINE COURSE CONTENT

MODULE 1: TYPES 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

Types of AI..... 2

A. General Information 3

B. Content for STEM HE Teachers..... 3

 Highlights..... 3

 Subtopic Content 3

 Green Decision Context 6

 Real World Examples..... 7

 Example Projects 7

 Recommended Readings..... 7

 Videos/Interactive Resources 7

 References..... 7

C. Assessment and Practice Activities..... 8

 Suggested 60-Minute Sequence 8



Types of AI

A. General Information

Module No	Module 1
Subtopic	Types of AI
AI Literacy Component	Understanding of the different types of AI: Generative AI, predictive AI etc
Decision Making Skills	Evaluation of the tools that are available for each different type of AI
Learning Objectives for STEM HE Teachers	Classify AI tool types and select tools appropriate for specific teaching or green-decision tasks
Learning Objectives for STEM Students	Compare AI tools for a specific educational task connected to green decision-making
Suggested AI Tools	ChatGPT; Canva (Magic Studio); Gradescope; ALEKS
Estimated Duration	60 minutes

B. Content for STEM HE Teachers

Highlights

After this subtopic, teachers should emphasise:

- AI is not one technology but a family of approaches: predictive, generative, classification, and adaptive systems
- Each type of AI has different strengths, data requirements, and limitations
- Predictive AI forecasts outcomes, generative AI creates new content, classification systems sort and detect, adaptive systems combine several techniques
- Selecting the right type of AI for a task is itself a decision-making skill
- Even the most advanced AI type still requires human oversight and interpretation

Subtopic Content

Introduction

This subtopic moves teachers and students from knowing what AI is to understanding that AI is not one technology but a family of approaches, each suited to different tasks. Building on the foundational concepts from Subtopic 1.1, this subtopic introduces four practical categories: predictive AI, which forecasts outcomes from historical and real-time data; generative AI, which creates new content such as text, images, or code; classification and recommendation systems, which sort data or suggest actions; and adaptive AI, which combines several of these techniques and adjusts its behaviour as new data arrives. Drawing on established classification frameworks and the technical literature on machine learning, the goal is for STEM higher education teachers to classify tools accurately, explain the differences to students, and select the right type of AI for a given teaching or green-decision task, rather than treating every AI product as interchangeable.

This distinction matters because expecting the wrong type of tool to do a job it was never built for, such as asking a chatbot to precisely forecast next month's energy demand, leads to both unrealistic expectations and poor decisions. The subtopic walks through how each category emerged, what data and trade-offs it requires, and how to evaluate a tool's maturity before trusting its output, always tying the discussion back to concrete green-decision examples like crop-yield forecasting, environmental report drafting, and wildlife-trade detection. By the end, teachers and students should be able to look at any AI tool and ask "what type of AI is this, and does that match what I actually need?", a habit of critical, tool-appropriate thinking that carries directly into the chatbot-focused subtopic that follows.



Why Classifying AI Matters

The word “AI” is often used as if it described one uniform technology, but in practice it covers a wide range of systems with very different capabilities, data requirements, and limitations. Recognizing this, AI systems can usefully be organized along four practical dimensions: the context in which a system operates, the data and input it uses, the underlying AI model itself, and the task and output it produces (OECD, 2022). Applying a version of this logic in the classroom, without necessarily using the same technical vocabulary, helps teachers and students avoid treating every AI tool as interchangeable, expecting a forecasting tool to be creative, or trusting a generative tool to be factually precise, when in fact each type of AI has its own strengths and blind spots.

Rule-Based Systems versus Machine Learning

Before comparing modern AI types, it helps to distinguish AI from its predecessor, rule-based systems, sometimes called expert systems. Rule-based systems follow explicit instructions written by a human programmer: for example, “if pH is below 6.5, classify the water as acidic.” These systems can be reliable for narrow, well-understood problems, but they cannot adapt to situations the programmer did not anticipate. Machine learning systems, by contrast, are shown large volumes of example data and learn their own internal rules for making predictions or generating content. This flexibility is what allows modern AI to handle complex, real-world tasks such as recognizing pollution patterns across thousands of variables, but it also means the system’s internal reasoning is harder to inspect directly, a trade-off often referred to as the “black box” problem.

Machine Learning, Deep Learning, and Neural Networks

Machine learning is the broad approach of learning patterns from data rather than following fixed rules. Within machine learning, deep learning refers to a family of techniques based on artificial neural networks, layered structures loosely inspired by how neurons connect in the brain, relying on “multiple processing layers to learn representations of data with multiple levels of abstraction” (LeCun, Bengio, & Hinton, 2015). These deep neural networks typically include “an input layer, at least three but usually hundreds of hidden layers, and an output layer, unlike neural networks used in classic machine learning models, which usually have only one or two hidden layers” (IBM, n.d.). Most of the AI systems associated with recent breakthroughs, from image recognition to chatbots, are built using deep learning, and it is precisely this added depth that lets them uncover subtler patterns than earlier, shallower models could, which also helps explain why some AI tools require enormous datasets and computing power while simpler machine learning tools do not.

Watch: [“But What Is a Neural Network? | Deep Learning, Chapter 1” \(3Blue1Brown\)](#) for a visual walkthrough of this layered relationship.

Predictive AI

Predictive AI analyses historical and real-time data to forecast future outcomes. It “blends statistical analysis with machine learning algorithms to find data patterns and forecast future outcomes”, extracting insight from historical data to estimate the most likely upcoming event, result, or trend (IBM, n.d.). Examples relevant to green decision-making include estimating tomorrow’s air pollution levels from current traffic and weather data, forecasting energy demand to help balance electricity grids, or predicting the likelihood of flooding from rainfall and river-level trends. Predictive AI systems typically output a probability or a numeric estimate rather than a definitive yes-or-no answer, for example “68% chance of exceeding safe particulate levels tomorrow afternoon”, which is one reason interpreting their outputs correctly requires understanding basic probability alongside the underlying science.

Watch: [“Supervised Learning #2” \(Crash Course, PBS\)](#) to see how predictive systems like these are actually trained.

Generative AI

Generative AI, by contrast, can be defined as artificial intelligence that “responds to a user’s prompt or request with generated original content, such as audio, images, software code, text or video” (IBM, n.d.), drawing on patterns learned from large training datasets rather than retrieving stored answers or forecasting a numeric outcome. Generative adversarial networks, introduced in 2014, were one of the techniques that made this kind of high-fidelity content generation practical, pairing a generator network that produces new content against a discriminator network that learns to judge its quality, with the two networks improving



each other through competition (Goodfellow et al., 2014). Several distinct techniques fall under this generative umbrella. Some generative models are designed to produce variations of an input, useful for exploring alternative designs or scenarios. Others work by gradually adding, and then removing, random noise from an image until a coherent picture emerges, an approach commonly used in modern image-generation tools. A third major approach, used in most of today's headline-making chatbots, processes sequences of data, such as words in a sentence, to predict and generate extended, coherent sequences of new content. These techniques underpin the AI chatbots examined in detail in the next subtopic.

Watch: [Predictive vs Generative AI: How They Work and When to Use Each](#) (IBM Technology) for a side-by-side comparison of these two types.

Classification and Recommendation Systems

Two further categories are especially common in educational and environmental applications. Classification systems sort data into predefined categories, for example, identifying whether an image shows a clean or polluted environment, or labelling a water sample as safe or unsafe based on its measured properties. Recommendation systems suggest content, products, or actions based on patterns in past behaviour, such as suggesting relevant readings to a student based on topics they have previously engaged with, or prioritizing which neighbourhoods should receive an environmental inspection first based on historical risk patterns. Both categories rely on the same underlying machine learning principles as predictive and generative AI, but are organized around a different kind of output: a category label or a ranked suggestion, rather than a forecast or newly generated content. Recommendation systems illustrate why output type alone is not enough to judge a tool: research on large language model-based recommenders shows that optimising primarily for accuracy can mask fairness problems, since most existing fairness-evaluation frameworks still fail to account for personalisation, allowing biased or unrepresentative suggestions to go undetected (Sah et al., 2023).

Watch: [Computer Vision Explained in 5 Minutes](#) for a quick visual look at how classification systems “see” and sort images.

Data Requirements and Practical Trade-offs

Each type of AI also comes with different practical requirements that matter when selecting a tool for a classroom or research context, closely tracking the “data and input” dimension of AI classification frameworks introduced above (OECD, 2022). Predictive and classification systems generally require structured, labelled historical data specific to the problem at hand, for example years of air-quality readings paired with known pollution events, and tend to perform poorly if that historical data is sparse or unrepresentative. Generative AI systems are typically trained once on very large, general-purpose datasets by a technology company, and then made available for anyone to use through a prompt, which means individual teachers and students do not need to train their own model, but also have limited ability to correct systematic errors or gaps in what the model knows. Recommendation systems depend heavily on having a continuous stream of user interaction data to remain accurate over time, which is why their quality often improves the longer a platform is used. Weighing these practical trade-offs, how much data is available, who controls the training process, and how the tool's accuracy might change over time, is as important as understanding the underlying technique when choosing an appropriate AI type for a given task.

Choosing the Right Type of AI for a Task

Evaluating and comparing tools within each category is itself a decision-making skill this module aims to build. A teacher preparing a lesson on renewable energy optimization, for instance, might compare a predictive model that forecasts solar output against a generative tool that drafts explanatory materials for students, choosing the one that fits the specific learning objective rather than defaulting to whichever tool is most familiar or most talked about. Similarly, a student investigating local air quality trends needs a predictive or classification tool to analyse sensor data, not a generative tool, which would be the wrong category of AI for that task no matter how capable it is at generating text. Building the habit of asking “what type of AI is this, and does that match what I actually need?” prevents both over-reliance on the wrong tool and unrealistic expectations of what a given AI system can do.

Real-World Examples Across AI Types

Seeing these categories applied side by side to the same domain makes the distinctions concrete. In renewable energy, a predictive AI system might forecast how much solar power a grid will generate



tomorrow based on weather data, while a classification system separates faulty solar panels from healthy ones using thermal-camera images, and a generative tool drafts a plain-language summary of the day's energy performance for a public dashboard. In disaster response, a predictive model estimates the path and intensity of an approaching storm, a classification system labels satellite images as flooded or dry, and a recommendation system prioritises which communities should receive emergency alerts first based on historical vulnerability data. Walking through examples like these with students and asking them to identify which type of AI is doing the work in each case reinforces the classification skills this subtopic is designed to build.

Limitations Across AI Types

These categories are not always neatly separated in practice: a single tool may combine predictive and generative functions, and the same underlying technique, such as a neural network, can be applied across categories. What matters for AI literacy is recognizing the primary function of a tool, forecasting, generating, classifying, or recommending, since this shapes what kind of output to expect and how critically to evaluate it. As with all AI systems, each type carries limitations: predictive models are only as reliable as the data they are trained on, generative tools can produce plausible but inaccurate content, and classification systems can misclassify inputs that differ from their training examples. Teachers should model this critical evaluation for students, reinforcing that classifying an AI tool correctly is the first step toward using it responsibly in STEM and green-decision teaching contexts. A large benchmarking effort spanning text, image, and vision-language generative models reaches a similar conclusion at a larger scale: improving one dimension of trustworthiness, such as safety, does not automatically improve another, such as fairness or robustness, so each AI type's limitations still need to be checked on their own terms rather than assumed from a single headline metric (Huang et al., 2026).

How These AI Types Emerged

Understanding roughly when and why each type of AI emerged also helps explain the current landscape. Predictive and classification techniques using structured, numerical data have existed in some form since the earliest days of machine learning and remain the backbone of most environmental monitoring and forecasting systems today, precisely because structured sensor and measurement data lends itself well to this approach. Deep learning-based generative techniques are comparatively recent: generative adversarial networks, capable of generating convincing variations of images, were introduced in 2014 (Goodfellow et al., 2014), image-generating diffusion techniques followed a few years later, and the sequence-based technique underlying most modern chatbots became widely available to the public only in the last several years. This recency is one reason generative AI tools are evolving so quickly and why their capabilities and limitations, unlike those of longer-established predictive systems, are still actively being discovered and debated by researchers, teachers, and students alike.

Assessing the Maturity of an AI Tool

Because different types of AI have different histories, they also differ in how mature and well-tested they are for a given use case. A predictive model used by a national weather service has typically been validated against decades of historical outcomes and refined through repeated real-world use. A newly released generative AI feature, by contrast, may have been tested for only months before public release, and its failure modes may not yet be fully documented. This does not mean newer tools should be avoided, but it does mean teachers should adjust their level of scrutiny accordingly: the more established and rigorously validated a tool is for a specific task, the more its outputs can reasonably be trusted without extensive independent verification, and the newer or more general-purpose a tool is, the more that verification becomes necessary. Structured tools such as the Model Openness Framework support this judgement by rating how completely a released model's code, data, and documentation have been shared, giving teachers a concrete signal of how thoroughly a tool's behaviour can be independently verified before classroom use (White et al., 2024). Having established this classification (OECD, 2022), the next subtopic looks closely at one specific and increasingly common type of generative AI: conversational chatbots, examining how they work, how to prompt them effectively, and how to evaluate their responses critically in educational settings.

Green Decision Context

Different types of AI support green decision-making in different ways:

- predictive AI forecasts environmental risks such as crop yields, floods, or methane leaks



- generative AI drafts environmental reports, policy summaries, and educational content
- classification and computer-vision systems detect and sort, for example identifying wildlife-trafficking listings or recyclable materials
- adaptive, multi-system AI integrates several of these techniques at once, such as a smart city coordinating traffic, energy, and waste data

However:

- classifying an AI tool correctly is the first step to knowing how much to trust its output
- every type of AI reflects the limitations of its training data and design
- human specialists remain responsible for combining AI outputs with local knowledge and values

Real World Examples

- **AGRIST – AI Crop Yield Forecasting for Sustainable Greenhouse Agriculture** (predictive AI): uses a machine-learning pipeline (built with Microsoft’s AI Co-Innovation Lab) to forecast greenhouse crop yields, production-deployed since February 2025, illustrating predictive AI’s role in resource planning.
- **Government of Catalonia – GenAI for Environmental Impact Report Automation** (generative AI): uses generative AI to help draft parts of environmental impact reports, cutting preparation time by 45% while requiring mandatory human review, an example of generative AI applied to regulatory environmental work under continued human oversight.
- **German Environment Agency (UBA) AI Lab – AI for Illegal Online Wildlife Trade Detection** (classification): combines image-classification machine learning with rule-based natural language processing to scan online marketplaces for illegal wildlife-trade listings, flagging over 50,000 protected animals in a 6-month pilot, an accessible example of classification-type AI applied to environmental law enforcement.
- **Greyparrot – AI Waste Analytics for Material Recovery Facilities** (computer vision/classification): uses computer vision to characterize material streams in recycling facilities in real time, generating the data insights that inform sorting and recovery decisions.

Example Projects

- **ECOMOD – Capacity Building in Environmental and Economic Policy Modelling and Data Analysis** (2023-2026): builds researchers’ and students’ skills in programming microsimulation models and analysing complex environmental and economic datasets, directly connecting to how different types of AI and data models are applied to climate-related policy questions.
- [CDOA – Academia for Smart City Chief Data Officers](#) (2025-2028): trains future data officers to combine data analytics, AI-driven decision-making, and sustainability in urban governance, illustrating how multiple AI types work together in real institutional decision-making.

Recommended Readings

- OECD (2022). [OECD Framework for the Classification of AI Systems](#).
- IBM. [Generative AI vs. Predictive AI: What’s the Difference?](#)

Videos/Interactive Resources

- [“But What Is a Neural Network? | Deep Learning, Chapter 1” – 3Blue1Brown](#): a visual explanation of the neural-network techniques underlying most modern AI types.
- [“Supervised Learning #2” – Crash Course: Artificial Intelligence \(PBS\)](#): introduces one of the main learning approaches behind predictive and classification AI.
- [Predictive vs Generative AI: How They Work and When to Use Each](#) – IBM Technology.
- [Computer Vision Explained in 5 Minutes](#) – a concise visual explainer of how classification-type AI “sees” and sorts images.

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).



- Banh, L., & Strobel, G. (2023). Generative artificial intelligence. *Electronic Markets*, 33, Article 63. <https://doi.org/10.1007/s12525-023-00680-1>
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative Adversarial Networks. *Advances in Neural Information Processing Systems*, 27. <https://doi.org/10.48550/arXiv.1406.2661>
- Huang, Y., Gao, C., Wu, S., Wang, H., Wang, X., Zhou, Y., et al. (2026). On the Trustworthiness of Generative Foundation Models – Guideline, Assessment, and Perspective. <https://doi.org/10.48550/arXiv.2502.14296>
- IBM. (n.d.). [Generative AI vs. Predictive AI: What’s the Difference?](#) IBM Think. Retrieved August 26, 2026.
- IBM. (n.d.). [What Is Deep Learning?](#) IBM Think. Retrieved August 26, 2026.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep Learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- OECD. (2022). [OECD Framework for the Classification of AI Systems](#).
- 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/10.48550/arXiv.2401.04057>
- Song, Y., Dhariwal, P., Chen, M., & Sutskever, I. (2023). Consistency Models. *Proceedings of the 40th International Conference on Machine Learning*. <https://doi.org/10.48550/arXiv.2303.01469>
- White, M., Haddad, I., Osborne, C., Liu, X.-Y. Y., Abdelmonsef, A., Varghese, S. M., & Le Hors, A. (2024). The Model Openness Framework: Promoting Completeness and Openness for Reproducibility, Transparency, and Usability in Artificial Intelligence. <https://doi.org/10.48550/arXiv.2403.13784>

C. Assessment and Practice Activities

The following activities give teachers concrete, measurable ways to check whether STEM higher education teachers can classify AI tool types and select tools appropriate for a task, and whether students can compare AI tools for a specific educational task connected to green decision-making, 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 the four categories of AI, predictive, generative, classification, and adaptive, using the “Predictive AI”, “Generative AI”, and “Classification and Recommendation Systems” sections and a linked video.
- 10–20 minutes: guided reading and discussion of data requirements and how to match a tool type to a task.
- 20–35 minutes: Activity (a), the tool-type identification task.
- 35–50 minutes: Activity (b), the mismatched-claim verification exercise, scored with the checklist in Activity (c).
- 50–60 minutes: formative quiz and self-reflection, Activity (d).

(a) Tool-Type Identification Task

Give students a short green-decision task, for example “help our class understand this term’s local air-quality trends”, and have them name the first AI tool that comes to mind, typically a general-purpose chatbot, along with one sentence on why. Then, using the four-category framework and the “data and input” dimension introduced above, students identify which type or types of AI the task actually requires, predictive, generative, classification, or adaptive, and name a tool from that category better suited to the task. As a measurable outcome, students submit their initial and revised tool choices side by side with two or three sentences explaining which classification distinctions changed their answer, directly assessing the learning objective of comparing AI tools for a specific task.



(b) Verifying a Tool-Type Claim

Provide students with three or four short marketing-style descriptions of AI tools, at least one of which makes a claim that does not match its underlying category, for example a generative writing assistant claiming it can “precisely forecast next term’s energy costs”, or a predictive model described as being able to “generate a custom lesson plan from scratch”. Students identify which claims are consistent with the tool’s actual type and which describe a category error, correcting the mismatched claim by naming the type of tool that would genuinely support it. As a measurable outcome, students submit an annotated version of the descriptions marking each claim as consistent, a category error, or unclear, together with a one-sentence correction for any category errors found.

(c) Tool-Type Classification Checklist

Teachers can use or adapt the following short checklist, scored yes / no / unsure, to assess a tool selection made in Activity (a) or (b), or independently:

- Has the primary function of the tool, forecasting, generating, classifying, or recommending, been correctly identified?
- Does the chosen tool’s output format actually match what the task requires: a numeric estimate, new content, a category label, or a ranked suggestion?
- Has the tool’s data requirement, structured historical data, a large general-purpose training set, or ongoing interaction data, been considered?
- Would a different, more specialised tool type produce a more reliable or more directly useful result for this specific task?
- Has the tool’s level of validation and maturity for this specific use case been considered, rather than assuming a newer or more general-purpose tool is automatically better?

(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 what a predictive AI tool and a generative AI tool each produce as output.
- Give one example of a green-decision task and name the type of AI, predictive, generative, classification, or adaptive, best suited to it, with a one-sentence justification.
- Name one reason two different types of AI might both be useful for the same overall project, such as a smart-city energy system.
- Reflect in two or three sentences: describe a situation where you, or someone you know, used an AI tool for a task that may not have matched its actual type, and explain what a better-matched tool might have looked like.

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 compare AI tools for a specific task in their own words, directly evidencing the student learning objective for this subtopic.

