What is AI?

MODULE 1 | EXCHANGE AI PROJECT

Teaching notes for adult education trainers — **Expeditions for Change: Navigating the Skills & AI Landscape (EXChange AI)**

Ref. No.: 2024-2-BG01-KA210-ADU-000296153

These notes are written as spoken guidance for the educator. They can be copied directly into the Speaker Notes field in PowerPoint and adapted according to audience, time and specific training context.

Co-funded by the European Union

Funded by the European Union. The views expressed are those of the authors and do not necessarily reflect the position of the EU or the granting authority.

Pedagogical Approach & How to Use This Module

Who is this for?

Adult learners and educators who do not need a technical background. The aim is to build a **clear, practical and critical understanding** of AI — not to turn participants into programmers.

Module Structure

1. Machine learning basics
2. Data and its role
3. Key terminology
4. Possibilities and limitations
5. Generative AI

AI may evoke both enthusiasm and anxiety. Good facilitation balances opportunities and limitations — showing useful applications while emphasising verification, human judgement, data protection and ethical use.

> Recommended closing question: *"Which one small task in your teaching practice would you experiment with supporting by AI during the next week?"*

Why This Matters Now — and Demystifying AI

The Scale of Change

Economic forecasts project **$13–22 trillion** in annual value from AI by 2033 — affecting education, healthcare, manufacturing, services, administration and creative industries. These numbers signal direction, not precise prediction: AI is becoming a key **general-purpose technology**.

Generative AI tools such as ChatGPT, Gemini, Claude, Midjourney and DALL-E are currently the most visible for the public because people can interact with them directly.

Three Categories of AI

ANI

Artificial Narrow Intelligence — very good at one specific task: speech recognition, spam filtering, content recommendation. Does not "think" like a human.

AGI

Artificial General Intelligence — a hypothetical system that could do everything a human can do intellectually. Not everyday reality today.

Generative AI

A subgroup that **creates new content**: text, images, code, music, video. Popular because it is easy to use in education and work.

Part 01: Machine Learning

Classical Programming

Humans write explicit rules in advance

Machine Learning

Model learns patterns from many examples

In classical programming, a person describes rules in advance. In machine learning, the approach is different: we give the computer many examples and it discovers patterns in them. For adult learners, an analogy with human learning is helpful — if we want a person to recognise fruit, we show many examples of apples, pears and oranges until they start distinguishing characteristics. Machine learning does something similar, but through data and mathematical models.

Supervised Learning

The machine receives examples with correct answers and learns to connect new inputs with expected outputs.

The A→B Mapping

Input A (email, photo, property data) maps to output B (spam/not spam, cat/not cat, predicted price).

Statistical Relationships

The system does not understand the world in the human sense — it learns a statistical relationship between A and B from examples.

Supervised Learning — Everyday Applications

Many familiar AI applications can be described as A-to-B mappings. The strength of AI does not come from magic, but from a well-defined task.

Spam Filtering

Input: an email → Output: probability it is spam. One of the earliest and most reliable supervised learning applications.

Speech Recognition

Input: audio → Output: text. Powers voice assistants, transcription tools and accessibility features.

Machine Translation

Input: sentence in one language → Output: translation into another. Used in tools like DeepL and Google Translate.

Online Advertising

Input: information about an ad and a user → Output: probability of a click. Drives much of the digital economy.

Ask the group: "In your work, which task could be represented as an input and an output?" For example — input: a learner's text; output: feedback. Choosing the right inputs, outputs and data is more important than the general desire to "use AI".

How Large Language Models Work — and Why Now?

How LLMs Work

LLM stands for **large language model**. These models are trained to predict the next word or element in text. When the model sees "My favourite drink is...", it calculates which words are most likely to follow based on the enormous amount of text on which it was trained.

Next-word prediction, repeated billions of times over enormous text collections, leads to surprisingly complex abilities: writing, summarising, translating, answering questions, generating ideas and dialogue.

⚠ The model does not "know" things the way a person knows through experience. It generates a probable answer based on patterns in data. Results must be checked — especially for facts, numbers, legal requirements, medicine or finance.

Why Progress is Accelerating

Three factors combine to drive rapid progress:

More Data

Vastly larger training datasets than ever before.

Larger Models

Neural networks with billions of parameters can learn finer patterns.

More Powerful Computers

GPU and TPU hardware enables training at unprecedented scale.

For most educators, it is more important to understand how to **use existing tools critically, safely and pedagogically** than to build their own models.

Part 02: Data — The Foundation of AI

Data is the "experience" from which the model learns. If the experience is rich, accurate and diverse, the model has a better chance of working well. If the data is incomplete, wrong or biased, the results will also be problematic.

Structured Datasets

Each row is one example (e.g. one house). Columns are features — square footage, bedrooms — and a target value such as price. The model learns the relationship between features and target.

Manual Labelling

People review examples and assign correct labels. Often expensive and slow, but can provide high quality. Raises questions of ethics, consent and personal data protection.

Observed Behaviour

Platforms record what users click on, how long they view content or what answers they choose. Requires careful attention to ethics and consent.

Existing Datasets

From partners, public repositories or websites. Requires attention to copyright, quality, bias and fit with the purpose of the project.

> Critical questions to always ask: **"What data did this AI learn from?"** and **"For whom might these data not be representative?"**

Data is Messy — Garbage In, Garbage Out

Common Data Problems

Issue	Example	Impact
Erroneous values	Price = 0.001	Model learns wrong relationships
Missing values	"unknown" in a numerical column	Processing errors or bias
Mixed formats	Numbers and text in same column	Makes processing harder
Duplicated records	Same learner entered twice	Skews model training
Incomplete surveys	Only active participants respond	Biased conclusions

In real organisations such problems are normal: different formats, duplicated records, incorrect labels and multiple versions of documents are the rule, not the exception.

The Core Principle

The quality of an AI project starts with **data quality**, not with the choice of a fashionable tool.

Educational Connection

If we use AI to analyse learner feedback but collect only the views of the most active participants, the result may be biased. If we evaluate knowledge through a poorly designed test, the model may learn wrong dependencies.

The Google Cat Example

Even recognising a cat — easy for a human — requires many examples, correct labels and suitable data representation for a machine.

Part 03: The Terminology of AI

Many learners hear "AI", "machine learning", "data science", "deep learning", "neural networks" and "generative AI" at the same time. These terms are related, but they are **not synonyms**.

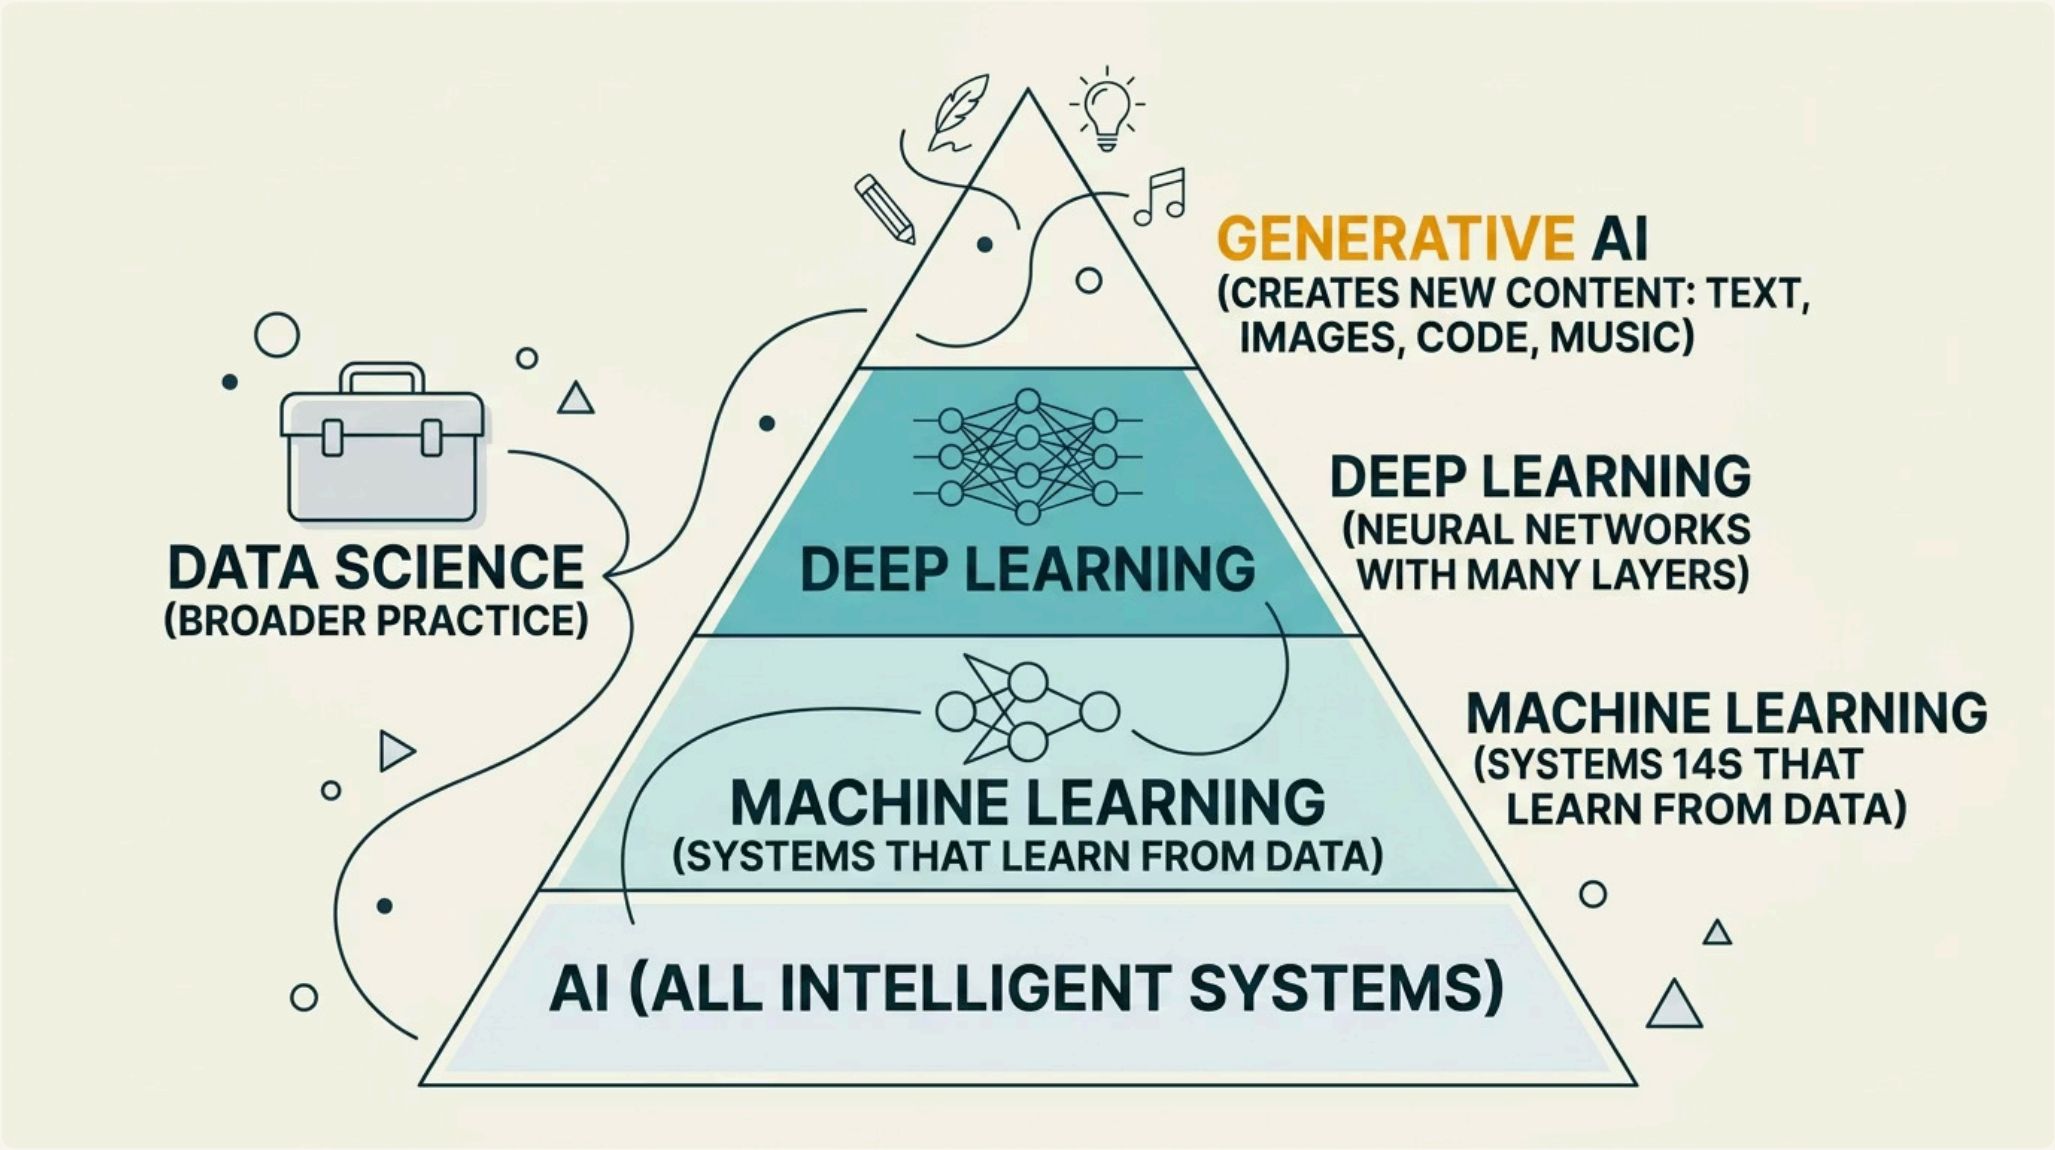

Machine Learning vs. Data Science

Machine learning often ends with a model that automates or supports a decision — for example, a website that predicts house prices. **Data science** often ends with a chart, report, insight or recommendation for people who make decisions — for example, showing which factors are associated with learners dropping out.

Two Complementary Definitions

> **Arthur Samuel:** Machine learning gives computers the ability to learn without being explicitly programmed.

> **Data Science:** The extraction of knowledge and insights from data so that people can act in a more informed way.

Part 04: Deep Learning — A Non-Technical View

Deep learning is a type of machine learning based on artificial neural networks. The term "neuron" is a metaphor — these are not biological cells, but **mathematical functions** that take inputs, combine them and produce outputs. A neural network is not conscious and does not understand like a human. It is a complex mathematical system that adjusts many parameters to connect inputs with outputs.

Inputs

Size, bedrooms, bathrooms, renovation status — raw features fed into the network.

Hidden Layers

Each neuron learns a different part of the dependency. Layers build internal representations of increasing complexity.

Output

The final prediction — for example, a house price or a face identity.

For images: early layers learn edges and lines. Middle layers combine these into parts — eyes, nose, mouth. Later layers combine parts into a full face representation. This is why we use the term "deep" learning — many layers of processing, not deep thinking.

For educators: this explains why deep models can discover complex patterns, but also why they are sometimes difficult to explain. The more layers and parameters there are, the harder it is to trace exactly why the model produced a particular result.

Part 05: What ML Can — and Cannot — Do

Andrew Ng's Rule of Thumb

If an average person can perform a task with about **one second of thought**, there is a strong chance it can be automated or supported through supervised learning.

✓ Defect detection

Does this image contain a defect?

✓ Spam detection

Is this email spam?

✓ Next word

What word is likely to come next?

✓ Pedestrian detection

Is there a pedestrian in this image?

Where ML Struggles

Unpredictable systems

Stock market prices depend on politics, news, psychology and global events. Lots of data does not guarantee reliable prediction when patterns are unstable.

Distribution shift

A model trained on data from one country may work worse in another cultural or linguistic environment.

Few examples

Machine learning struggles when it must learn a complex function from only a few labelled examples.

Human judgement tasks

Context, value judgement, empathy, long-term planning and responsibility require human oversight.

⚠ In education: AI may suggest feedback on a text, but the educator must judge whether it is pedagogically appropriate and fair. AI can detect patterns, but predictions about human behaviour, motivation or "potential" should never be used mechanically for final decisions.

Part 06: Case Studies — Self-Driving Cars & X-Ray Diagnosis

Self-Driving Cars — The Long Tail Problem

In typical situations models may perform well: lanes, other cars, pedestrians, traffic lights. But the problem lies in unusual cases — the so-called **long tail**.

- A stop sign partly hidden by a tree
- A person carrying an unusual sign
- A cyclist giving an unexpected signal
- A construction zone or unexpected pedestrian behaviour

Each case may be rare, but the system must react correctly because the consequences are serious. For high-risk applications, it is not enough for the model to be good "on average".

X-Ray Diagnosis — Quality Over Quantity

Machine learning can be very strong on narrowly defined tasks, such as detecting signs of pneumonia in chest X-rays — **but only if trained on a sufficiently large and high-quality set of labelled images**.

Scenario	Result
~10,000 labelled real X-rays	Model can learn real visual patterns ✓
10 textbook illustrations	Not enough — no diversity of real patients ×

AI in medicine should be seen as an **assistive tool**, not an independent replacement for professional responsibility. The same logic applies in education: AI can support assessment and personalisation, but the educator remains responsible for context and fairness.

Part 07: Strengths, Weaknesses & the AI Toolbox

When ML Works Well

- The task is **stable enough** — small changes in input do not lead to chaotic changes in output
- There is a **lot of high-quality, labelled data** covering different situations including edge cases
- The input-output relationship is **smooth and repeatable**
- The task can be clearly defined as an **A-to-B mapping**

AI is a Toolbox, Not One Tool

Machine learning and deep learning are the most visible parts, but there are also: unsupervised learning, reinforcement learning, planning, knowledge graphs, graphical models and others.

We do not use a hammer for every task. In the same way, not every problem requires generative AI or a neural network. Sometimes a search engine, spreadsheet, rule, database or simple automation is enough.

The First Question

When someone proposes an "AI solution", the first question should always be: **"What problem are we solving?"** Only then should the choice of tool come.

For the EXChange AI project: when using AI in adult education, think about localisation, language, cultural context, digital literacy and the specific needs of learners. A tool that works well in one context is not automatically appropriate in another.

Part 08: Generative AI, Transformers & Broader Considerations

The Virtuous Cycle

More users → more data → better product → more users. This helps explain why AI markets can quickly concentrate around a few strong platforms. Educators should watch not only functions, but also **dependencies on specific providers**.

LLMs Are Not Calculators

Language models generate **probabilistic text**. They are very good at explanations, formulation, summarisation and ideas — but they can make mistakes in arithmetic, precise facts or logical details. For calculations, use a calculator, spreadsheet or Python.

Transformers

The architecture behind GPT, Claude and Gemini. These models process relationships between words, phrases and parts of context very effectively. Participants do not need to memorise technical details, but interactive tools like **Transformer Explainer** build useful intuition.

> Practical advice: Use an LLM to help **formulate a task**, explain a formula or interpret a result — but always **check the numbers** with an appropriate deterministic tool.

When teaching adults about AI, we do not have to start with formulas. We can use visual and interactive resources that build intuition and then, if there is interest, move towards more technical explanations.

Module Summary & Next Steps

KEY TAKEAWAYS

AI is a broad term

It includes different approaches and tools — ANI, AGI, generative AI, deep learning and more. Not all systems work the same way.

Much of modern AI is machine learning

Learning from examples, not from explicitly programmed rules. The A-to-B mapping is the core intuition.

Data is key

Without high-quality, relevant and ethically used data, there is no reliable AI. Garbage in, garbage out.

Generative AI outputs must be checked

LLMs are powerful assistants, but they generate probabilistic text. Always verify facts, numbers and critical information.

AI supports human judgement — it does not replace it

In education, the educator remains responsible for context, fairness and pedagogical appropriateness.

Continue Learning

Visit **exchangeai.eu** for additional resources. Think of small, concrete applications in your own practice: creating examples, adapting text, preparing exercises, translating, summarising feedback or planning lessons.

Closing Question

"Which idea from today's module would you apply first in your own educational practice?"

This project is funded by the European Union. The views expressed are those of the authors and do not necessarily reflect the position of the EU or the granting authority.