Module 2: AI Use Cases

EXCHANGE AI · EXPEDITIONS FOR CHANGE

From the abstract question **"What is AI?"** to the practical question **"Where do we meet AI in everyday life, work and learning?"** — this module uses concrete examples to help learners recognise AI as a set of practical systems, not a distant technology. No technical background required.

Project No. 2024-2-BG01-KA210-ADU-000296153 · exchangeai.eu

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 Human Resource Development Centre. Neither the European Union nor the granting authority can be held responsible for them.

Opening Question

> "Where might you have used AI during the last week without noticing it?"

Encourage short examples from participants before diving into the case studies. AI is already woven into daily routines — often invisibly.

🔍 Search Engines

Results ranked by AI algorithms learning your preferences

📱 Social Media

Recommendations and feeds curated by machine learning

🌐 Translation

Automatic language translation powered by NLP models

📧 Spam Filters

Email classification using text-based AI classifiers

🎵 Recommendations

Music and video suggestions from collaborative filtering

ⓘ The aim of this module is to help learners move from vague awareness to confident recognition of AI in context.

Case Study: Smart Speaker

The smart speaker is a familiar everyday example — yet behind its simple appearance, several distinct AI components work together. Ecosystems such as **Amazon Alexa**, **Google Home**, **Apple Siri**, and **Baidu DuerOS** all follow the same underlying logic.

The Key Insight

Even a very simple command requires coordination between sound recognition, language understanding, and task execution. AI systems are often built from smaller, specialised modules — not a single "magic" brain.

Try This With Your Group

Use the command **"Hey device, tell me a joke"** as a starting point. Ask participants: what must the system do before it can respond correctly?

- Wake up and listen
- Convert speech to text
- Understand the user's intent
- Execute the right action

Inside the Smart Speaker: Two Commands

Examining specific commands step by step reveals how multiple AI layers interact. Errors can occur at every stage — the wake word may be missed, speech transcribed incorrectly, or intent misunderstood.

COMMAND 1

"Hey device, tell me a joke"

Wake-Word Detection

Device listens for trigger phrase and decides: activate or stay silent?

Speech Recognition

Audio "tell me a joke" is converted into text — estimating the most likely words

Intent Recognition

Does the user want a joke, music, a call, or weather?

Execution

A suitable joke is selected or generated and played as audio

COMMAND 2

"Hey device, set timer for 10 minutes"

Wake-Word Detection

Same trigger mechanism activates the system

Speech Recognition

Audio converted to text including the number "10 minutes"

Intent + Parameters

Intent: start a timer · **Parameter:** duration = 10 minutes

Execution

Timer is started; device confirms verbally

People express the same meaning many ways: "remind me in 10 minutes", "count 10 minutes", "tell me when 10 minutes are over" — this is why NLP is hard.

Case Study: The Self-Driving Car

The self-driving car is a more complex example. Unlike the smart speaker, which mainly receives speech, the autonomous vehicle operates in a **constantly changing physical environment** — perceiving the road, recognising objects, predicting movement, and making decisions in real time.

Computer Vision

Cameras, radar, and lidar provide raw data about the environment

GPS & Maps

Location data combined with high-definition road maps

Motion Planning

Choosing a safe and efficient path through the environment

Control

Commands sent to steering, acceleration, and braking systems

Ask participants: "What does a self-driving car need to recognise before deciding to stop or turn?" Expected answers: cars, pedestrians, road signs, traffic lights, lanes, obstacles, speed, direction.

Self-Driving Car: Perception to Action

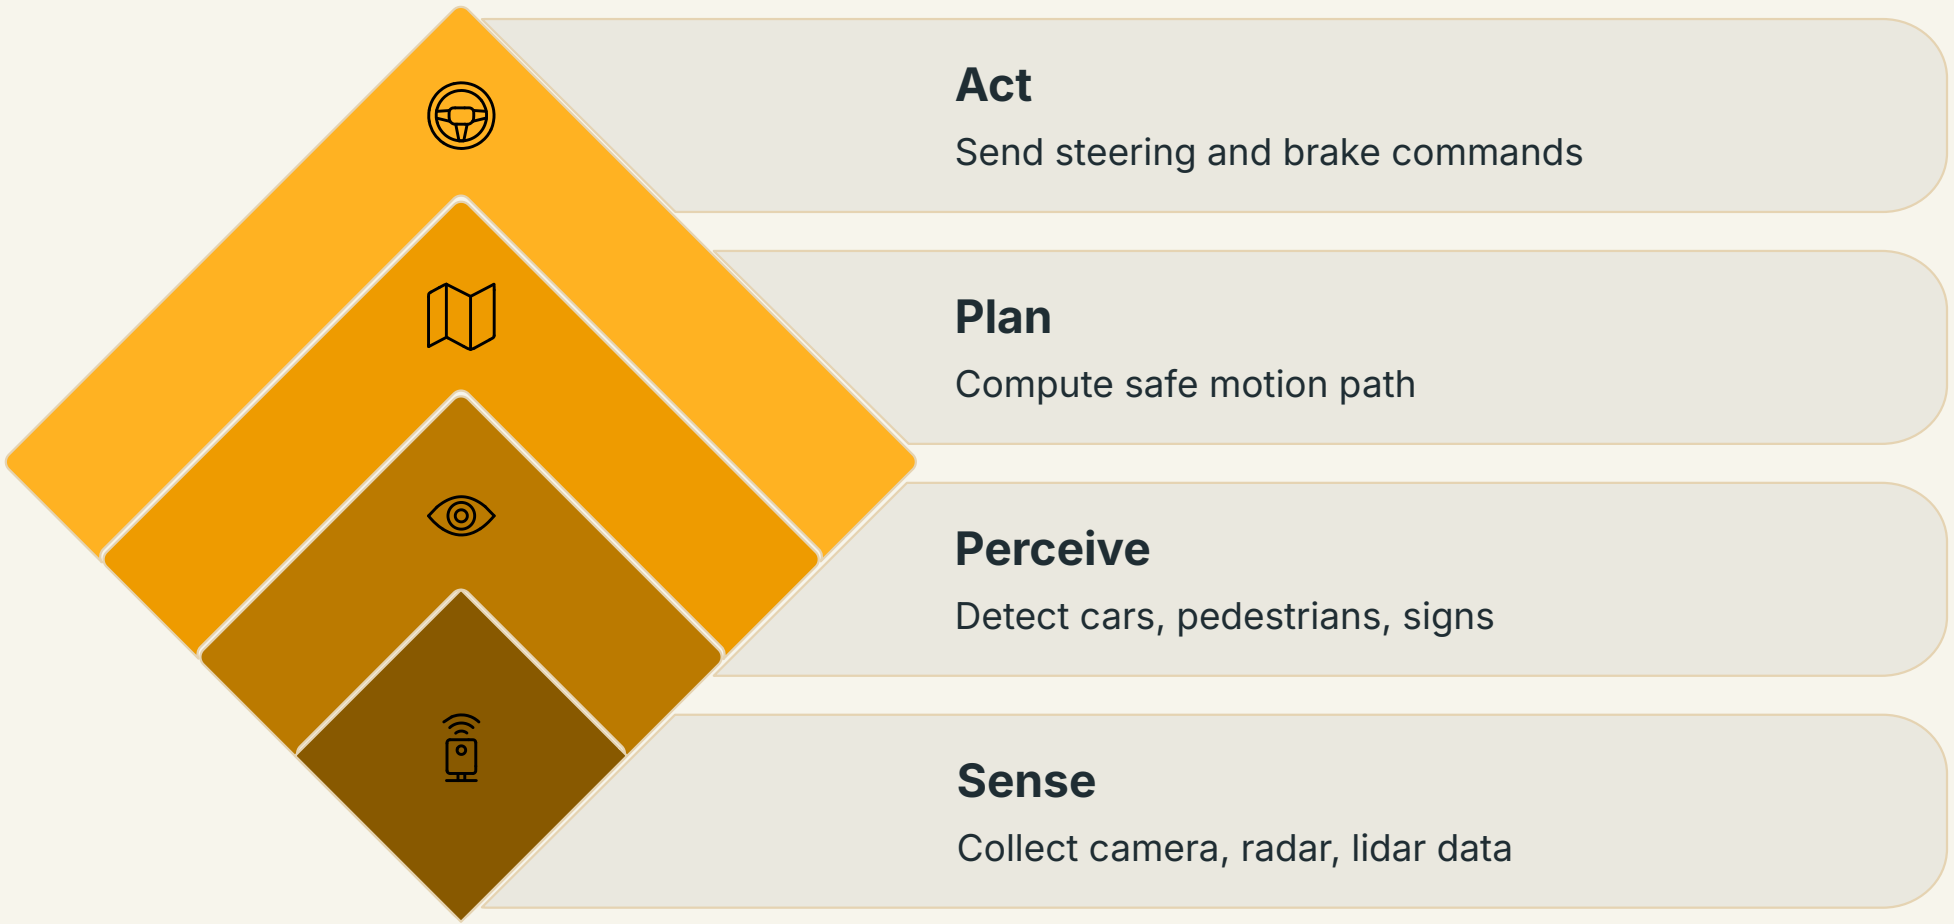

The car does not "see" the world as a human does — it receives numerical data and images that must be interpreted by software modules. Detecting a pedestrian is especially challenging due to different body positions, movement, partial occlusion, poor lighting, rain, night conditions, and unpredictable human behaviour.

Expanded System Components

- Lane detection and traffic-light recognition
- Obstacle detection and avoidance
- Trajectory prediction for moving objects
- Backup safety mechanisms and redundancy

Why Trajectory Prediction Matters

If a pedestrian stands near a crossing, the car must estimate whether they may start crossing. If another vehicle is in a neighbouring lane, the system must predict a possible lane change. Mistakes here have physical consequences — making safety and testing essential.

Part 02: Major AI Application Areas

After the two case studies, learners can now identify the individual technologies behind them. These areas are not completely separate — in real applications they are often combined.

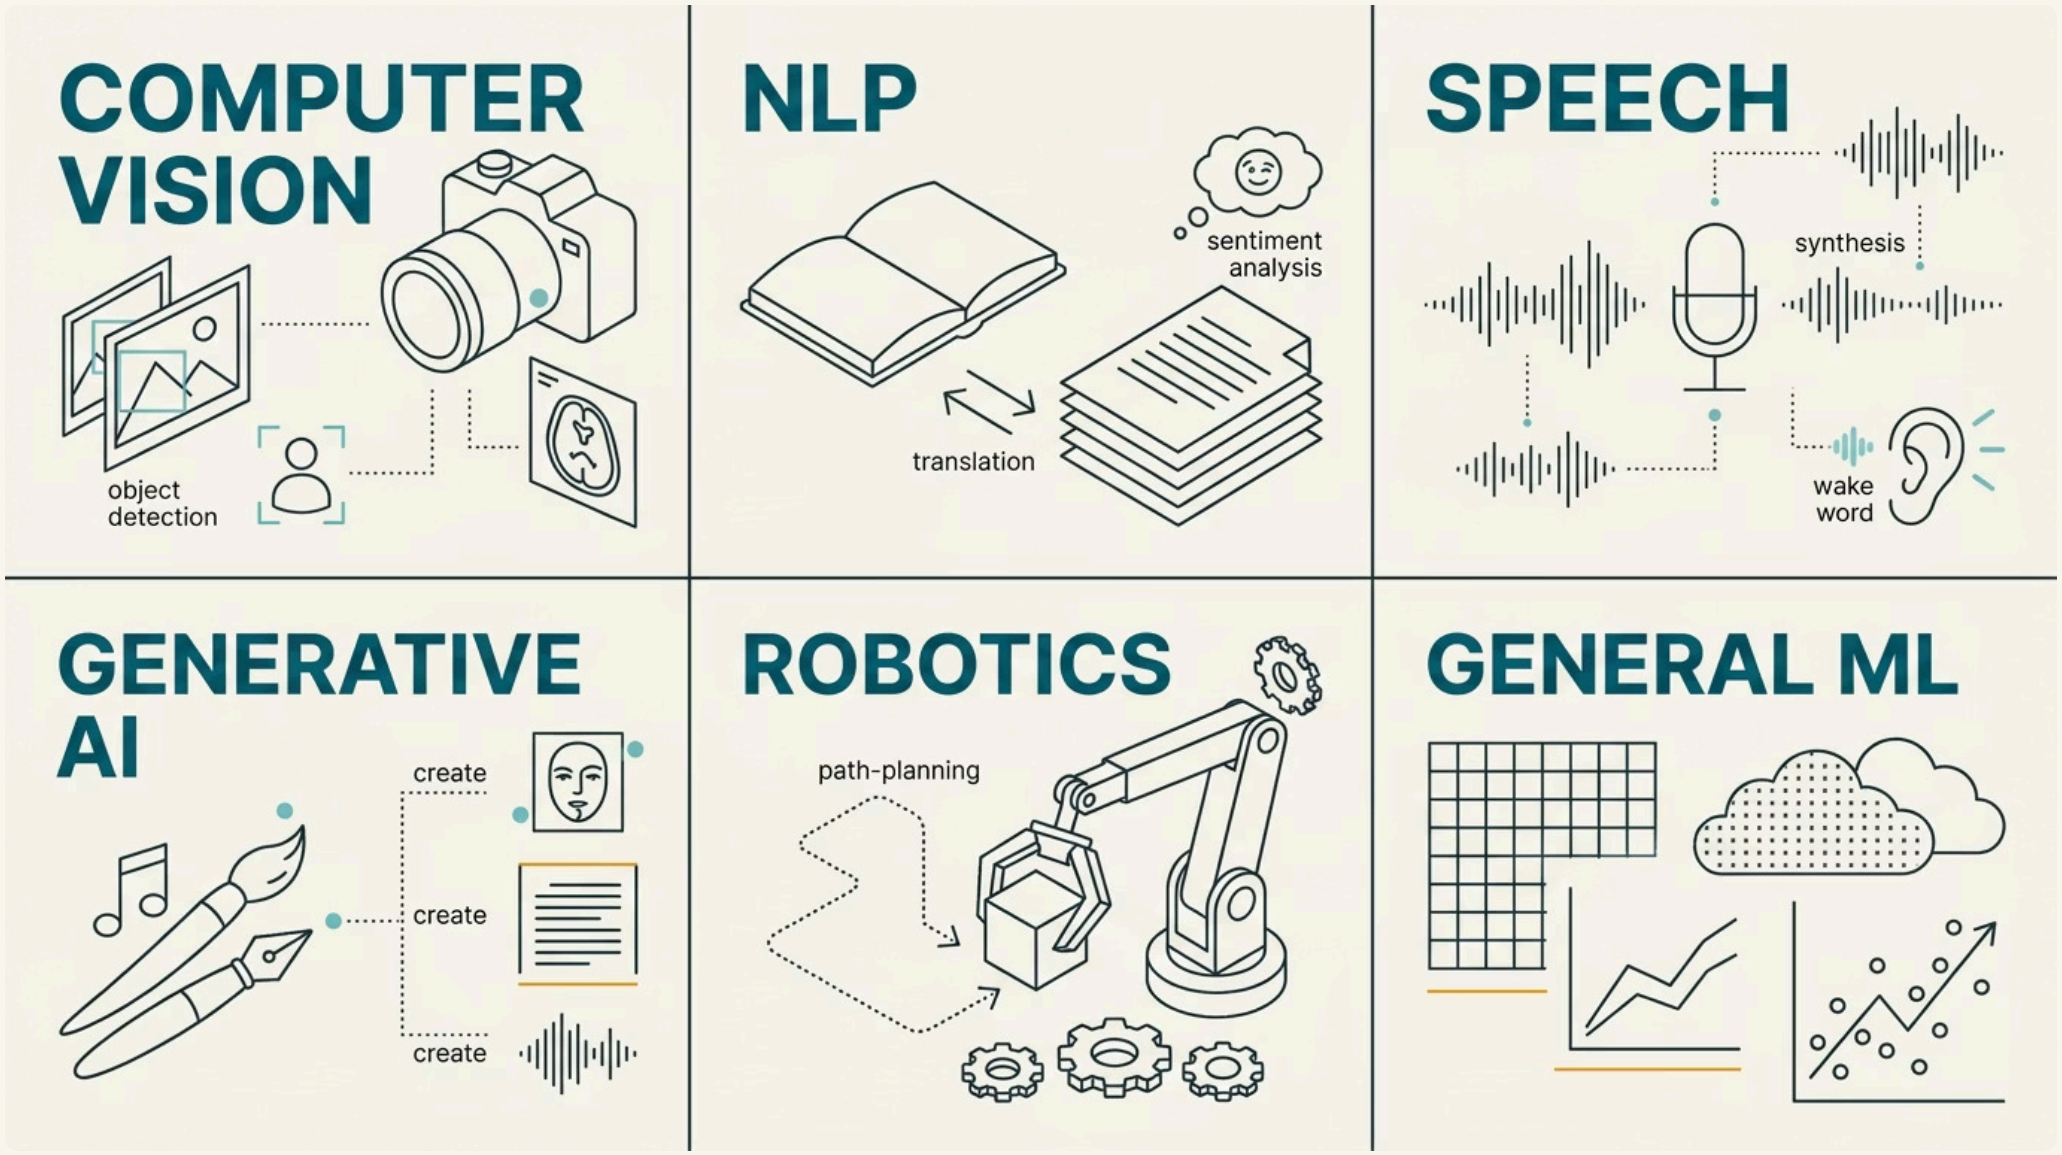

Activity: Ask participants to choose a familiar application — Google Translate, ChatGPT, Netflix recommendations, spam filters — and identify which AI areas are involved.

Computer Vision & Natural Language Processing

Computer Vision

AI analyses images and video through a hierarchy of tasks:

Classification

What is in this image? → "Cat"

Object Detection

Where is it? → Draws a bounding box around the cat

Segmentation

Exact shape? → Outlines every pixel belonging to the cat

Tracking

Over time? → Follows the object through video frames

Applications: medical imaging, quality control, road safety, photo tagging, support for visually impaired people.

Natural Language Processing

AI processes human language — texts, questions, emails, documents, translations, and conversations.

Text Classification

Spam filter: is this email spam or not?

Sentiment Analysis

"Service was horrible" → negative sentiment detected

Named-Entity Recognition

Extracts: Queen Elizabeth II, Buckingham Palace, Sir Paul McCartney

Machine Translation

Converts text from one language to another

⚠ AI results should always be checked critically — especially in educational contexts where accuracy matters.

Speech, Generative AI & Robotics

Speech Technologies

- **Speech recognition:** spoken language → text
- **Wake-word detection:** "Hey Siri", "OK Google"
- **Speaker identification:** who is speaking?
- **Speech synthesis:** text → spoken audio

Especially important for **accessibility**: automatic captions, dictation, audio learning materials, voice-based interfaces.

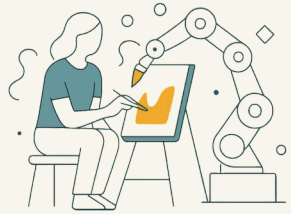

Generative AI

- Creates new text, images, audio, speech, or music
- Produces outputs based on patterns from large datasets
- Useful for drafts, ideas, visualisations, exercises

⚠️ Generated content may sound convincing but is **not always true, accurate, or suitable**. Treat it as a co-author of drafts, not a final authority.

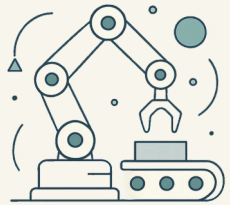

Robotics

- **Perception:** understanding the surrounding environment
- **Motion planning:** choosing a safe path
- **Control:** sending commands to motors

If a chatbot gives a poor answer, the harm is informational. If a robot moves incorrectly, it may break something or injure someone — making **human oversight essential**.

General Machine Learning: Data Types

Machine learning can work with fundamentally different types of data. Understanding this distinction helps educators think about what data their organisation already holds.

Unstructured Data

Does not have a simple table format — requires specialised models.

- Images and photographs
- Audio recordings and podcasts
- Documents, emails, and reports
- Video and meeting recordings

A huge part of everyday information — and increasingly accessible to AI systems.

Structured Data

Tables with rows and columns — often already available in organisations.

- Spreadsheets and databases
- CRM and administrative records
- Survey results and assessment scores
- House size, bedrooms, price — classic ML examples

Many organisations begin here because the data already exists in familiar formats.

Reflection for educators: "What data does your organisation already have that could support learning or analysis?" The quality, quantity, and representativeness of data directly affect AI results.

Part 03: Major AI Techniques

So far we have seen *what* AI can do. Now we look at several ways in which AI can **learn** or **organise knowledge**. Participants do not need to be able to programme these techniques — the goal is to understand what types of problems they solve.

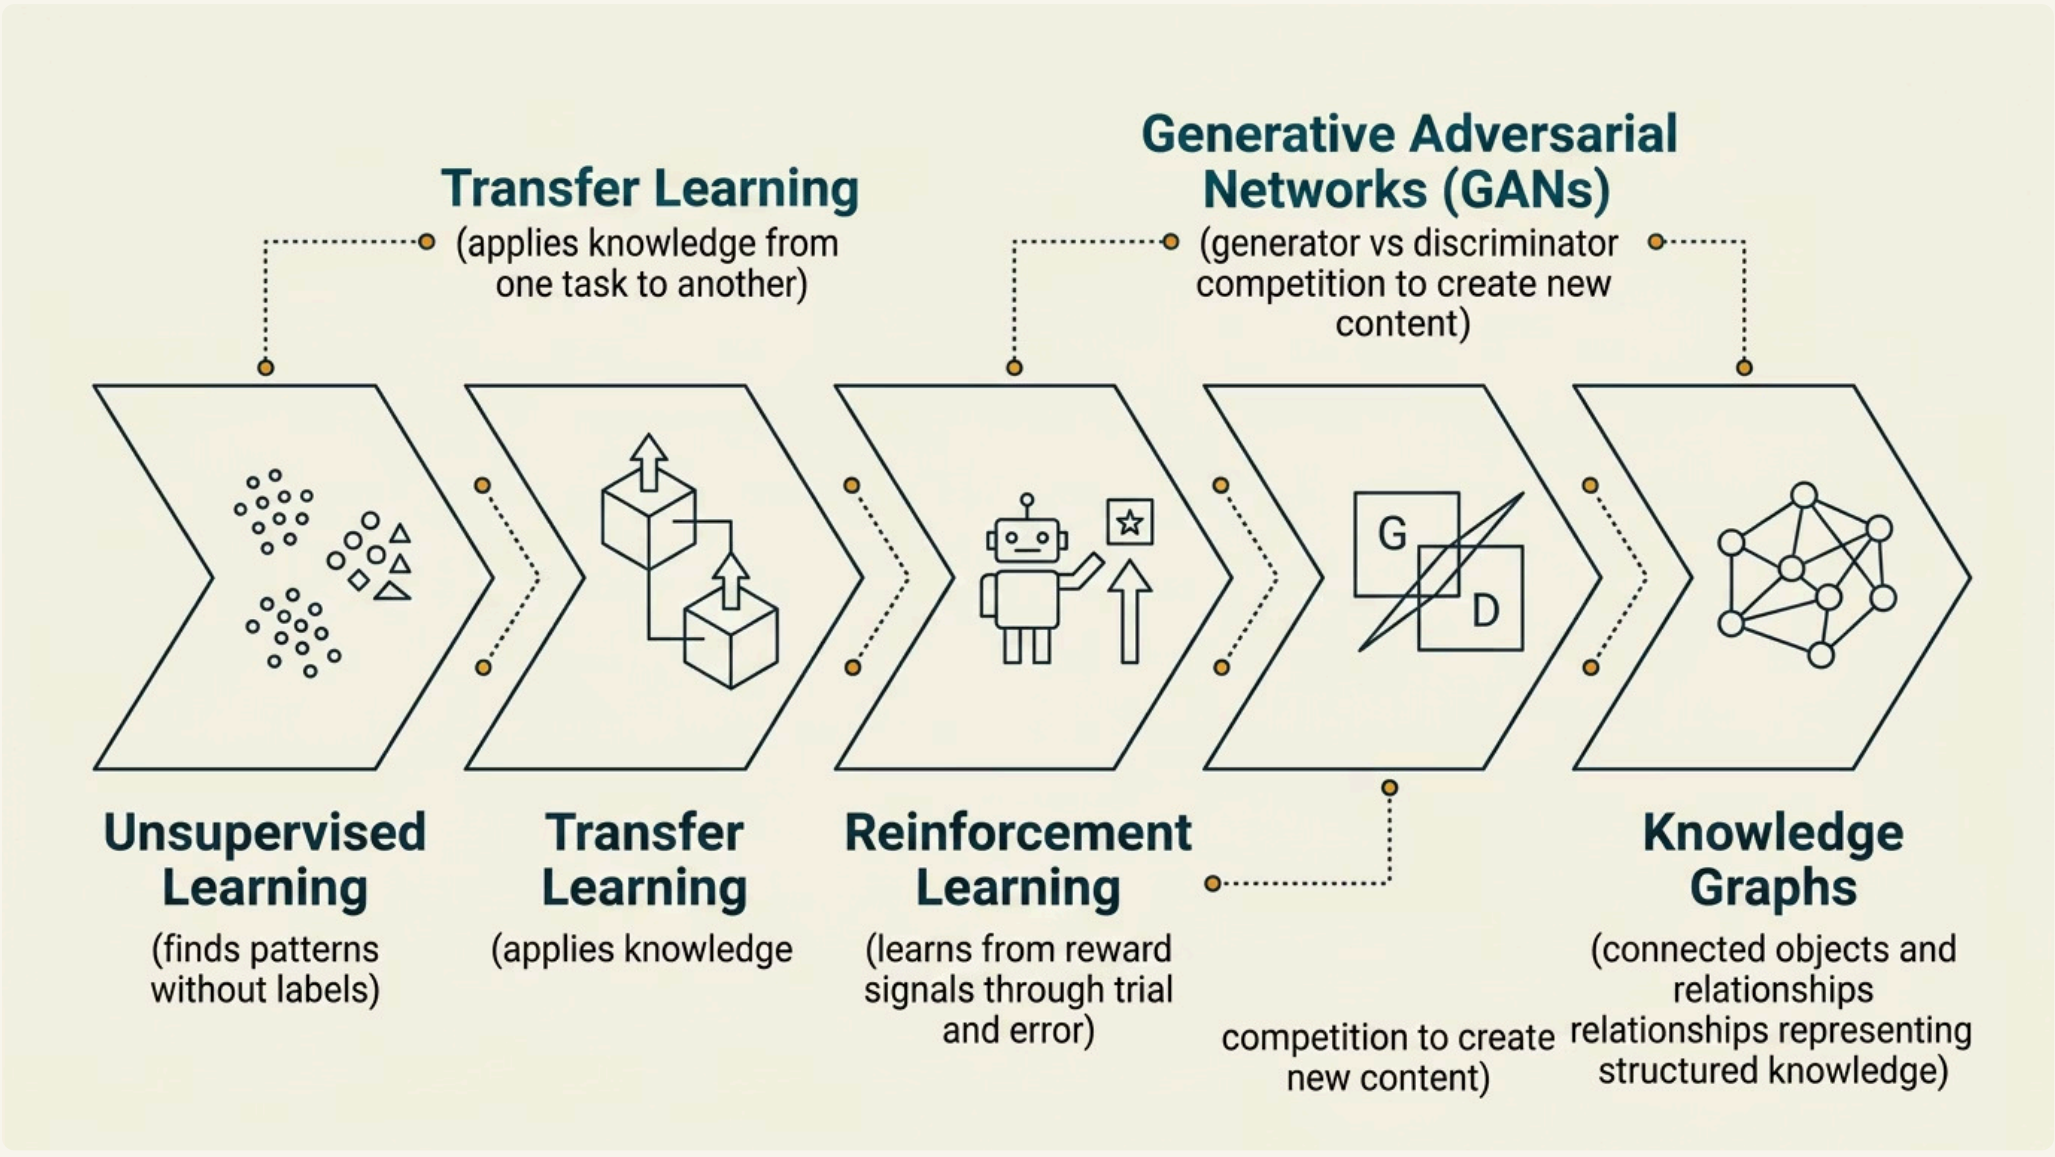

These techniques are not exhaustive, but they illustrate the diversity of approaches AI uses to learn from data and represent information. Understanding them helps educators ask better questions, evaluate tools, and communicate more confidently with technical teams.

Unsupervised Learning & Transfer Learning

Unsupervised Learning

> Give the system data without predefined correct answers. Let it discover structure on its own.

The system searches for similarities and groups — a process called **clustering**. Unlike supervised learning (where we say "this is a cat"), unsupervised learning finds patterns without labels.

- Finding cats in unlabelled videos
- Clustering potato-chip sales by price and volume
- Discovering learner groups with similar needs
- Grouping learning resources by topic

The groups found by the algorithm still need human interpretation — the machine finds clusters, but humans give them meaning.

Transfer Learning

> Use knowledge from one task to help solve another related task — with far less data.

A model trained on **100,000 images** for car detection can help recognise golf carts with only **100 images**. The model has already learned general visual features: wheels, shapes, edges, lights, perspective.

- Adapts large pre-trained models to specific contexts
- Works for particular languages, subjects, or audiences
- Enables faster development of useful solutions
- Does not mean the model understands everything automatically

In education: large pre-trained language models can be adapted to a particular curriculum, language, or learner group.

Reinforcement Learning

Reinforcement learning uses a **reward signal** to tell the system when it is doing well or poorly. Instead of receiving the correct answer for every situation, the AI tries actions and learns which ones lead to higher rewards.

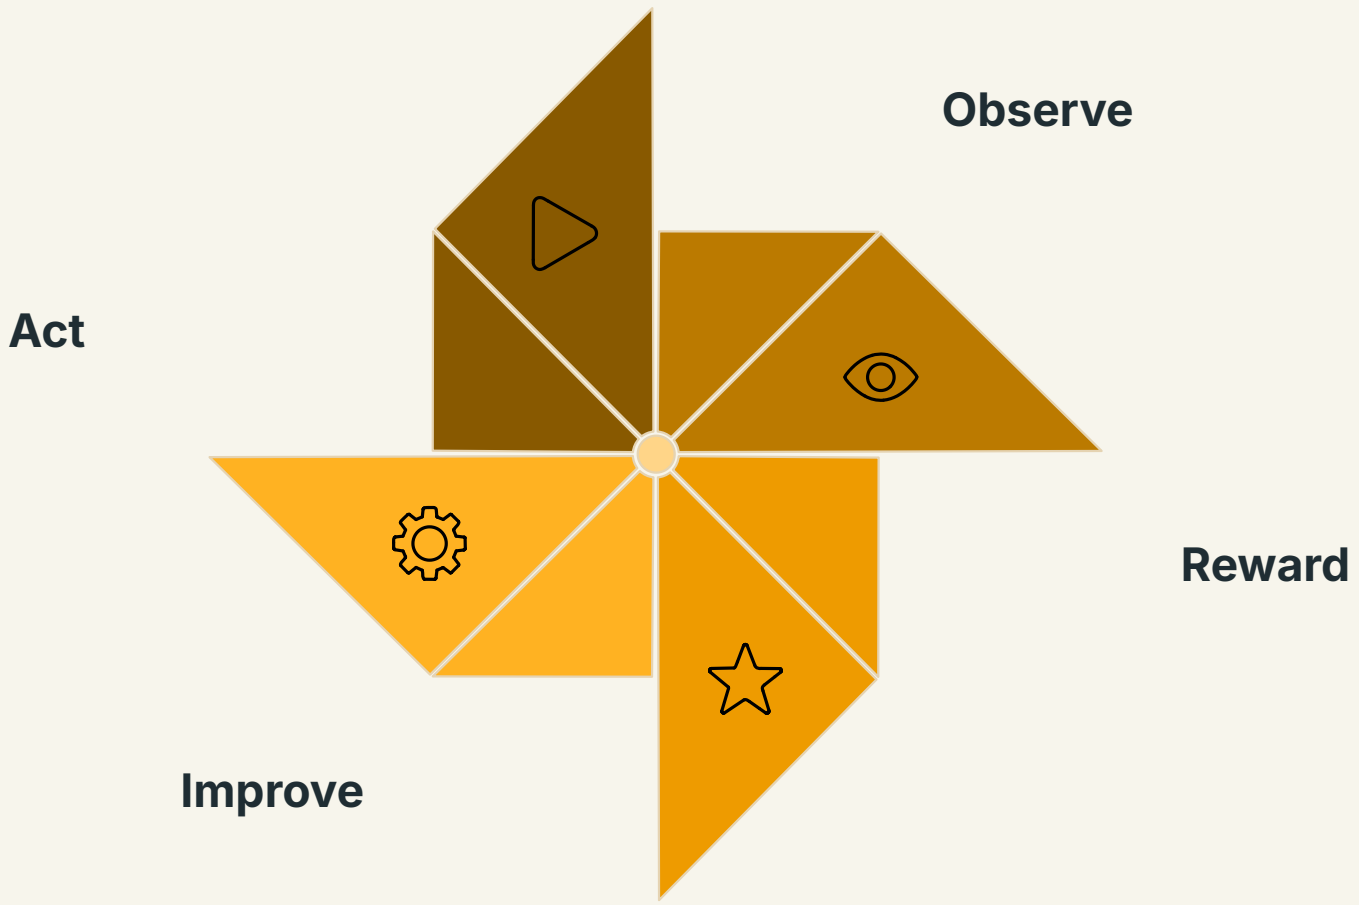

Intuitive Examples

- A robot learning to walk or grasp objects
- An agent learning to play a game
- A system optimising a delivery route
- Human analogy: learning a sport through trial and feedback

Critical Warning: Reward Design

Choosing the reward is critical. If the reward is poorly designed, the system may maximise the number but **not the human goal**. This is a key discussion point about ethics and objective-setting in AI systems.

In the real world, it is not always safe or affordable for a system to experiment freely — so reinforcement learning often uses **simulations** or tightly controlled environments.

GANs & Knowledge Graphs

Generative Adversarial Networks (GANs)

A GAN consists of two competing parts working against each other:

Generator

Tries to create a convincing new image from scratch

Discriminator

Tries to decide: is this image real or generated?

Through this competition, the generator gradually improves. GANs were a landmark step in generative AI — showing that AI can create convincing visual content.

✅ Opportunities

Design, prototyping, games, art, visualisation, education

⚠️ Risks

Deepfakes, fake images, disinformation, copyright concerns

Knowledge Graphs

A knowledge graph represents information as **connected objects and relationships** — not just text or tables, but nodes and links.

Nodes

Objects: Ada Lovelace, Northern Rooster Hotel, London, 1815

Properties

Attributes: born, died, address, phone, Wi-Fi, pool

Relationships

Links between nodes: "worked with", "located in", "founded by"

Applications: search, recommendations, question-answering systems, knowledge organisation, and concept maps in education.

Activity: Ask participants to choose a topic and name three objects and three relationships among them — a simple knowledge graph exercise.

End of Module 2: Key Takeaways

Today we explored how AI is used in real-world applications and which major areas and techniques stand behind those applications. Return to the opening question — participants should now be able to answer more specifically.

AI is a set of practical systems

Not an abstract or distant technology — it is already present in speakers, cars, cameras, search engines, and classrooms.

Systems are built from specialised modules

Wake-word detection, speech recognition, intent recognition, computer vision, and motion planning each solve a specific problem.

Understanding AI is pedagogical competence

It helps educators choose tools, design better exercises, explain risks, and prepare learners for an AI-rich digital environment.

Critical thinking remains essential

Generated content, algorithmic decisions, and reward signals all require human oversight, ethical reflection, and fact-checking.

> Invitation to action: Choose one idea from this module and consider how you could apply it in a course, workshop, or training activity.

exchangeai.eu · Erasmus+ funded EXChange AI project · Project No. 2024-2-BG01-KA210-ADU-000296153