



Integration of AI-powered tools and methodologies in the VET sector

BLOCK A

Foundations of AI and Large Language Models

Modules 1–4

Work Package 2 - Training Modules Development

KA210-VET · Erasmus+ · Pr Nr.: 2025-1-DE02-KA210-VET-000364083



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 European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.

Pr Nr.: 2025-1-DE02-KA210-VET-000364083

Foundations of AI and Large Language Models

About the Project

AI in ACTION – Integration of AI-powered tools and methodologies in the VET sector is an Erasmus+ project supporting Vocational Education and Training (VET) professionals in understanding and effectively using Artificial Intelligence. It focuses on making AI practical and accessible for teaching, learning, communication, content creation and everyday professional tasks.

The project combines essential AI knowledge with **hands-on activities and real VET examples**, helping participants use AI tools effectively, understand their possibilities and limitations, critically evaluate their outputs, and integrate AI responsibly into their work.

Who is it for?

The training programme is designed for **VET educators, trainers, mentors and administrators**, including those with little or no previous AI experience. It progressively develops participants' knowledge, practical skills and confidence through tasks connected to real VET contexts.

The Training Programme

The AI in ACTION training programme consists of **16 modules divided into 4 thematic Blocks**, with **each module designed as a 90-minute learning session**. The programme progresses from AI foundations to content creation, professional applications, and the responsible and sustainable use of AI.

Participants combine theoretical understanding with **practical exercises, demonstrations, discussions and hands-on tasks**. They learn how to select and apply AI tools for different purposes, evaluate AI-generated outputs and transfer what they learn to their own professional environments.



The four Blocks are:

Block A – Foundations of AI and Large Language Models

Introduces the foundations of AI, Large Language Models and effective prompting.

Block B – Generative AI Tools for Text, Audio & Visual Content

Moves into hands-on production using AI for professional text, voice and narration, music, video and visual content.

Block C – AI Applications in Professional and Personal Projects

Explores local AI models, agentic workflows, AI-powered data analysis, value creation and the development of participants' own AI implementation projects.

Block D – AI Strategy, Ethics, Funding and Sustainability

Addresses responsible AI use, deepfakes and misinformation, environmental sustainability, and the impact of AI on future skills and career pathways.

Together, the four Blocks create a **step-by-step learning journey from understanding AI to applying it in practice and integrating it responsibly into professional work.**

About Block A

Block A focuses on building the essential foundations needed to understand and work effectively with Artificial Intelligence and Large Language Models. Participants explore the potential of AI in the VET sector, learn key AI concepts and terminology, understand how Large Language Models work and compare different types of AI models. They also develop practical prompting skills that will support their work throughout the rest of the programme.

Block A includes four modules:

Module 1: Potential of AI in the VET Sector

Module 2: Core Concepts and Definitions of Artificial Intelligence

Module 3: Large Language Models (LLMs) - Commercial and Open Source

Module 4: Effective Prompt Engineering and Reverse Prompting



Contents

1

MODULE 1

Potential of AI in the VET Sector**2**

MODULE 2

Core Concepts and Definitions of Artificial Intelligence**3**

MODULE 3

Large Language Models (LLMs) - Commercial and Open Source**4**

MODULE 4

Effective Prompt Engineering and Reverse Prompting

Foundations of AI and Large Language Models

Block Overview

Block A builds the conceptual foundation for the whole training programme. Participants - VET educators, trainers and mentors - arrive with mixed levels of AI experience, from daily ChatGPT users to complete beginners. By the end of this block, every participant can explain what AI is and is not, understands how Large Language Models actually work (and why they fail), can choose between commercial, open-source and local models for a given task, and can write professional-grade prompts using the C.R.A.F.T. framework and advanced techniques such as few-shot, negative and reverse prompting.

Modules	1. Potential of AI in the VET Sector 2. Core Concepts and Definitions of AI 3. Large Language Models - Commercial and Open Source 4. Effective Prompt Engineering and Reverse Prompting
Total duration	Each Module is planned for 90 minutes
Target group	VET educators, trainers, mentors and administrators; no prior AI knowledge required
Materials needed	Laptop per participant (or per pair), internet access, projector, flipchart paper, markers, printed handouts, free accounts on chatgpt.com and/or claude.ai
Block outcome	Each participant leaves with a personal vocabulary of AI terms, a tested comparison of three model types, and a personal prompt library of at least 5 refined prompts

Pedagogical approach for the whole block: roughly 30% input, 70% activity. Every theory segment is immediately anchored in a discussion or hands-on task drawn from participants' own working context.



BLOCK A · FOUNDATIONS OF AI AND LARGE LANGUAGE MODELS

MODULE 1

Potential of AI in the VET Sector

Potential of AI in the VET Sector

Module Overview

Duration	90 minutes
Format	Interactive discussion-led session with icebreaker and mapping exercise
Prerequisites	None - this is the entry point of the programme
Materials	Flipchart, sticky notes in 3 colours, markers, projector, printed “Opportunity–Risk Map” worksheet

Learning Objectives

By the end of this module, participants will be able to:

1. Articulate their own starting point, expectations and concerns regarding AI in vocational education.
2. Identify at least three concrete areas of their daily work where AI could realistically add value (lesson preparation, translation of manuals, drafting safety briefings, differentiated explanations, administrative communication).
3. Identify at least two areas where AI should **not** be used (safety sign-offs, final competence assessment, tool-path validation) and justify why.
4. Describe how AI is changing the skill demands of the industries their trainees are entering.
5. Formulate one personal learning goal for the rest of the training programme.

Session Plan

Time	Segment	Method
0–15 min	Icebreaker: mind map + two truths and a lie	Individual + plenary
15–25 min	Baseline discussion: “Where are we starting from?”	Plenary, show of hands
25–45 min	Input: Why AI matters for VET now	Presentation + examples
45–65 min	Practical Task 1: The Opportunity–Risk Map	Pairs
65–80 min	Practical Task 2: Future skills radar	Small groups
80–90 min	Personal learning goal + wrap-up	Individual + plenary

Icebreaker (15 min)

Each participant draws a quick mind map answering three questions: one thing about yourself nobody in the room knows (your “secret talent”), one thing you do every day, and one thing you can’t live without. Two statements must be true and one a lie - the group guesses the lie..



Potential of AI in the VET Sector

Baseline Discussion - “Where are we starting from?” (10 min)

Run these four questions in plenary. Record answers on flipchart - you will return to them on the final day to make progress visible.

1. When you hear “AI”, what is the first concrete tool that comes to mind?
2. On a scale of 1 to 5 - how confident do you feel explaining what AI is to a trainee or student?
3. Have you used AI yourself for work - a lesson plan, a safety briefing, a CNC program, an email? How often?
4. If generative AI can eventually explain complex technical concepts or write flawless code, what becomes the most irreplaceable part of your role as an educator?

Facilitator note: Question 4 is deliberately provocative. Do not resolve it - let it hang. The whole programme is an extended answer to it.

Theoretical Input: Why AI Matters for VET Now (20 min)

■ The moment we are in

In late 2022, ChatGPT reached 100 million users in two months - the fastest technology adoption in history. Since then, generative AI has moved from novelty to infrastructure: it is embedded in office software, search engines, design tools and, increasingly, in the machines and management systems of the industries VET prepares people for. VET cannot opt out of this shift for three reasons:

Reason 1 - The labour market is changing underneath us. The occupations VET trains for are being reshaped: predictive maintenance in manufacturing, AI-assisted diagnostics in automotive, AI-supported design in construction and creative trades, chat-based customer service in commerce. Trainees who leave VET without basic AI literacy enter the labour market with a gap employers increasingly notice.

Reason 2 - Educators face a productivity revolution. Lesson plans, differentiated worksheets, translated safety instructions, parent emails, quiz banks, assessment rubrics - the routine 80% of an educator’s paperwork can now be drafted in minutes. The 80/20 rule applies: let AI produce the first 80% as a draft; the educator’s expertise supplies the decisive final 20%.

Reason 3 - Our trainees are already using it. Whether or not the institution has a policy, trainees use AI for assignments today. The choice is not “AI or no AI” but “guided, critical use vs. unguided, uncritical use”. VET educators are best placed to teach verification, because verification against physical reality is what vocational education has always done.

■ Where AI realistically helps a VET educator

- **Preparation:** lesson plans, differentiated versions of the same material for stronger and weaker trainees, scenario-based exercises, quiz and exam question banks.



Potential of AI in the VET Sector

- **Language and communication:** translating technical manuals, simplifying dense regulation text, drafting emails to parents, employers and authorities, multilingual materials for international learners.
- **Explanation:** producing alternative explanations and analogies when a trainee doesn't understand the first one - "explain hydraulic pressure using a bicycle pump analogy".
- **Administration:** meeting summaries, report drafting, structuring project applications, documentation.
- **Data:** summarising grades and attendance patterns, spotting outliers, generating charts (covered in depth in Module 10).

Where AI should not go - the red lines

- **Safety-critical sign-off.** AI output is probabilistic (Module 2 explains why). A safety briefing may be *drafted* by AI but must be *validated* by a qualified human. Final tool-path validation, electrical sign-offs, pressure-system checks: never delegated.
- **Final competence assessment.** The judgement that a trainee is competent to work on real equipment carries legal and moral responsibility that cannot be transferred to a statistical model.
- **Personal data.** Trainee names, grades, disciplinary records or ID numbers must never be pasted into free consumer AI tools (developed fully in Modules 3 and 13).

The responsibility question: if a trainee uses ChatGPT to produce work that is wrong, who is responsible? The answer VET has always given: the person who signs off the work. AI does not change accountability; it changes the workflow *before* accountability.

Practical Task 1: The Opportunity–Risk Map (20 min)

Setup: pairs, ideally mixing different subject areas. Each pair receives a large sheet divided into four quadrants and sticky notes in three colours.

Instructions to participants:

1. Draw two axes: horizontal = "Impact on my work" (low → high), vertical = "How soon" (2+ years → this semester).
2. On **green** notes, write concrete tasks from your own training day where AI could realistically help you. Be specific: not "lesson planning" but "generating three difficulty levels of the same welding-theory worksheet".
3. On **red** notes, write tasks where AI should not be used in your context, and *why* - what is the failure cost?
4. On **yellow** notes, write "not sure" cases you want answered during this training week.
5. Place all notes on the map. Discuss placements with your partner - you must agree on each position.

Share-back (5 min of the 20): each pair presents its single most valuable green note and its single most important red note. Facilitator clusters yellow notes on a "parking lot" flipchart - these become a checklist to resolve during the week.

Expected outputs: a photographed map per pair; a plenary parking lot of open questions.



Potential of AI in the VET Sector

Practical Task 2: Future Skills Radar (15 min)

Setup: groups of 3–4 by broad sector (technical, service, creative, administrative).

Instructions:

1. Think about the actual industries your trainees are preparing for. Name **one core skill you teach today that AI might substantially automate within 3–5 years**.
2. Name **two skills that will not go away** - and explain why (physical dexterity, safety judgement, client interaction, accountability).
3. Name **one new skill** your curriculum does not currently teach but should (e.g., verifying AI output against standards, prompt-writing for technical documentation, data literacy).
4. Prepare a 60-second report-back.

Facilitator note: steer groups away from panic (“everything will be automated”) and complacency (“nothing will change”). The realistic pattern is *task* automation within jobs, not wholesale job replacement - the mix of tasks changes, and curricula must follow.

Wrap-up: Personal Learning Goal (10 min)

Each participant completes the sentence on a card: “By the end of this training, I want to be able to ____ so that ____.” Cards are collected and returned on the final day for self-assessment. Close by previewing Module 2: “Before we can use AI well, we need to agree what the word actually means.”



BLOCK A · FOUNDATIONS OF AI AND LARGE LANGUAGE MODELS

MODULE 2

Core Concepts and Definitions of Artificial Intelligence

Core Concepts and Definitions of Artificial Intelligence

Module Overview

Duration	90 minutes
Format	Concept input alternating with quick-check quizzes and two hands-on demos
Prerequisites	Module 1
Materials	Projector, laptops with browser access (teachablemachine.withgoogle.com, quickdraw.withgoogle.com, autodraw.com), printed Vocabulary Cheat-Sheet

Learning Objectives

By the end of this module, participants will be able to:

1. Give a working definition of AI and name its four building blocks (input, output, autonomy, adaptivity).
2. Correctly sort the nested terms: AI $\supset$ Machine Learning $\supset$ Deep Learning $\supset$ Generative AI.
3. Classify everyday systems as “AI”, “not AI” or “it depends”, with justification.
4. Outline the 70-year history of AI, including the two AI winters, and explain why the current wave emerged (data + compute + algorithms).
5. Distinguish narrow AI (ANI), general AI (AGI) and superintelligence (ASI) and state which exists today.
6. Explain what changed in 2022 with generative AI, including the key risk: *fluent but wrong*.
7. Use eight core vocabulary terms correctly: model, parameter, token, prompt, context window, hallucination, fine-tuning, agent.

Session Plan

Time	Segment	Method
0–10 min	Quick Check: “Is it AI - or just clever software?”	Pairs quiz
10–25 min	Input: Working definition + the buzzword hierarchy	Presentation
25–40 min	Input: A short history - winters, why now, ANI/AGI/ASI	Presentation + timeline quiz
40–55 min	Input: What changed in 2022 - Generative AI	Presentation + comparison table
55–75 min	Practical Task: Train your own model (Teachable Machine)	Pairs, hands-on
75–85 min	Vocabulary Cheat-Sheet + AI drawing demos	Plenary
85–90 min	Three ideas for educators + exit ticket	Plenary



Core Concepts and Definitions of Artificial Intelligence

Opening Quick Check: “Is it AI - or just clever software?” (10 min)

In pairs, participants classify each item as YES, NO or IT DEPENDS - and must say why:

System	Answer	Why
CNC machine controller (G-code execution)	NO	Deterministic execution of code; no learning
Google Maps routing	YES	Learns patterns from traffic data
Alexa / smart speaker	YES	Voice recognition via deep learning
CAM software auto-toolpath	DEPENDS	Older versions rule-based; newer ones ML-assisted
Spotify / YouTube recommendations	YES	Collaborative filtering + machine learning
Washing machine timer	NO	Fixed program logic
Facial recognition	YES	Deep learning on image data
ChatGPT writing a tool-list for a job	YES	Generative AI
Spreadsheet formula calculating feed rate	NO	Deterministic calculation

Teaching point: the dividing line is not “computer” vs. “no computer” - it is *learned patterns from data* vs. *hand-coded fixed rules*.

Theoretical Input Part 1: A Working Definition (15 min)

Definition to use all week: *AI is the field that builds machines able to perform tasks that, until recently, required human intelligence - perceiving, reasoning, learning from data, and acting under uncertainty.*

Every AI system has four building blocks:

- **Input** - the data the system sees: text, sensor readings, machine logs, drawings.
- **Output** - what it produces: a label, a number, a sentence, a tool-path, an action.
- **Autonomy** - how much it does without a human in the loop at each step.
- **Adaptivity** - whether it improves with more data and feedback over time.

Sorting the buzzwords (nested circles):

- **Artificial Intelligence** - the umbrella: any machine doing a “smart” task.
- **Machine Learning** - algorithms that learn patterns from data instead of being hand-coded.
- **Deep Learning** - ML built on multi-layer neural networks: image, speech, language.
- **Generative AI** - deep-learning models that *produce new content*: text, image, audio, video.

Ask participants to place the nine quick-check items from 2.4 into the correct circle.



Core Concepts and Definitions of Artificial Intelligence

Theoretical Input Part 2: A Short History and Why Now (15 min)

Seventy years from a question to a chatbot

Year	Milestone	Significance
1950	Turing Test	Alan Turing asks: can a machine convincingly imitate a human?
1956	Dartmouth workshop	The term “Artificial Intelligence” is coined
1966	ELIZA	First chatbot - simulates a therapist using pattern matching
1997	Deep Blue	IBM’s chess engine beats world champion Garry Kasparov
2012	AlexNet	Deep learning wins ImageNet - the modern AI era begins
2017	Transformer	Google paper “Attention Is All You Need” - the engine of modern LLMs
2022	ChatGPT	100M users in 2 months - AI enters everyday life

Interactive moment: “Milestone of the decade?” quiz - participants vote between: SpaceX catches a returning booster; Twitter rebrands to X; Argentina wins the World Cup; the release of ChatGPT. Use the vote to discuss why a chat interface, not a technical breakthrough alone, changed everything: it made AI *usable by everyone*.

The two AI winters - why progress wasn’t linear

AI has over-promised and disappointed twice: 1974–1980 (early symbolic AI couldn’t scale; funding collapsed) and 1987–1993 (“expert systems” proved brittle and expensive). The pattern: breakthrough → over-promising → quiet years of foundational work → next leap. **Lesson for educators:** today’s wave is real, but cycles of correction are normal. Plan AI adoption on a long horizon, not on this month’s hype.

Why now? Three forces converged in the 2010s

- Data** - the internet generated trillions of pages, images and conversations: the fuel.
- Compute** - GPUs (originally built for video games) made training huge models affordable.
- Algorithms** - deep learning and the 2017 Transformer architecture unlocked language at scale.

None of the three would have been enough alone.

Three kinds of AI

- **ANI (Artificial Narrow Intelligence):** excels at one task - spam filter, translator, ChatGPT, a tool-wear detector. *This is every AI that exists today.*
- **AGI (Artificial General Intelligence):** would match a human across virtually any cognitive task. Actively researched; no verified system yet.
- **ASI (Artificial Super Intelligence):** would surpass the best humans across all domains. Hypothetical; subject of long-term safety debate.



Core Concepts and Definitions of Artificial Intelligence

Theoretical Input Part 3: What Changed in 2022 - Generative AI (15 min)

Earlier AI mostly *classified*: “is this email spam?”, “is this chip pattern abnormal?”. Generative AI *creates*: a paragraph, an image, a voice, a video, a draft program. Key properties:

- Trained on huge amounts of text, code, images and audio.
- It doesn’t “look things up” - it **predicts what comes next**, token by token.
- One model, hundreds of tasks: write, translate, summarise, code, draft.
- **Mistakes look fluent.** That is the new risk we must teach.

	Old AI	Generative AI
Task type	Classify, predict	Create new content
Output	Label or number	Text, image, audio, video
Interface	Built into a product	Open chat - prompt anything
Strength	Narrow, precise	Broad, flexible
Failure mode	Wrong but obvious	Wrong but fluent

How an LLM “knows” things - three stages, only one involving you:

1. **Pre-training** - the model reads billions of pages and learns statistical patterns. Done by the AI company; months of compute; the frozen result is the base model.
2. **Fine-tuning** - trainers show the model good and bad answers, teaching it to be useful and safe. Still done by the company; produces the version you use.
3. **Prompting** - you give it context and ask; the model continues your text one token at a time. *This is where your craft as an educator lives* - practised in Module 4.

An AI model is, physically, a large file storing billions of learned numbers (weights/parameters). This demystification matters: there is no consciousness, no database, no understanding - just extraordinarily sophisticated pattern completion.

Practical Task: Train Your Own Model - Teachable Machine (20 min)

Tool: <https://teachablemachine.withgoogle.com> (free, browser-based, no account needed)

Instructions (pairs):

1. Open Teachable Machine → New Image Project → Standard model.
2. Define two classes relevant to your own subject: e.g., “correct hand position / incorrect hand position”, “safety glasses on / off”, “clean weld / dirty weld” (use webcam or objects on the table - pens vs. markers works fine).
3. Record 30+ webcam samples per class. Train the model (one click). Test it live.



Core Concepts and Definitions of Artificial Intelligence

4. Now break it: show the model something it never saw during training (different lighting, angle, a third object). Observe how it *confidently misclassifies* - it must choose one of its known classes.

5. Retrain with more varied samples and observe improvement.

Debrief questions: What did the model actually learn - the object, or the background and lighting? What does this tell you about bias in training data? How does “confident but wrong” here connect to hallucination in chatbots?

Optional lighter demos (10 min or homework): Quick, Draw! (quickdraw.withgoogle.com) - a neural network guesses your sketches; AutoDraw (autodraw.com) - AI-assisted drawing; Animated Drawings (sketch.metademolab.com) - bring a sketch to life. Each shows pattern recognition in a playful, low-stakes way that participants can reuse directly with their own trainees as an AI-literacy icebreaker.

Vocabulary Cheat-Sheet (Handout)

Eight words that will come up every day:

Term	Meaning
Model	A trained mathematical function - the “brain” you query
Parameter	A learned number inside the model; modern LLMs have billions
Token	A chunk of text (~¾ of a word); models read and write in tokens
Prompt	The instruction or context you give the model
Context window	How much text the model can “see” at once
Hallucination	A confident, fluent, but factually wrong answer
Fine-tuning	Extra training to specialise a model for a domain or style
Agent	An LLM given tools and goals to take multi-step actions

Closing: What This Means for Educators (5 min)

Six ideas to carry through the week:

- 1. AI is a tool, not a colleague.** It has no understanding of your trainees, your machines or your safety standards. You provide the judgement; it provides the draft.
- 2. Fluent ≠ correct.** Generative AI sounds confident even when wrong - including about feed rates, materials and procedures. Teaching trainees to verify is now part of vocational literacy.
- 3. Pedagogy and craft still lead.** The best AI-supported lesson is still a good lesson plan. AI changes speed, not principles.
- 4. Use the 80/20 rule.** Let AI handle the initial 80% as your draft; add your expert final 20%.
- 5. Trust your judgement.** AI is a starting point, never the final authority.
- 6. Protect student privacy.** Never include trainee names or personal information in prompts.



Core Concepts and Definitions of Artificial Intelligence

Assessment and Reflection

- **Formative check:** re-run three items from the opening quiz at the end - scores should visibly improve.
- **Exit ticket:** each participant writes one sentence using two vocabulary terms correctly.
- **Transfer task (homework):** try Quick, Draw! with a family member or trainee and note what it teaches about pattern recognition.



BLOCK A · FOUNDATIONS OF AI AND LARGE LANGUAGE MODELS

MODULE 3

Large Language Models (LLMs) - Commercial and Open Source

Large Language Models (LLMs) - Commercial and Open Source

Module Overview

Duration	90 minutes
Format	Concept input + two structured hands-on activities comparing model types
Prerequisites	Modules 1–2 (especially vocabulary: token, parameter, hallucination)
Materials	Laptop per participant, accounts on chatgpt.com and claude.ai, LM Studio or Ollama pre-installed on at least a few machines (or facilitator demo machine), printed Model Selection Guide

Learning Objectives

By the end of this module, participants will be able to:

1. Explain in one sentence what an LLM is and how the Transformer generates text token by token.
2. Explain the difference between deterministic systems (CNC, calculators, spreadsheets) and probabilistic systems (LLMs), and why this matters for verification.
3. Describe what parameters are and why bigger is not always better (dense vs. Mixture-of-Experts architectures).
4. Map the global LLM landscape - US, Chinese, European and Middle-Eastern providers - and name at least two European alternatives relevant to GDPR-sensitive work.
5. Apply GDPR-aware rules before prompting: what never to paste into a free consumer tool; consumer vs. enterprise tiers; training-data opt-outs.
6. Run the same task on a commercial cloud model, an open-source cloud model and a local model, and evaluate the trade-offs.
7. Explain why AI-text detectors are unreliable and what to do instead.

Session Plan

Time	Segment	Method
0–15 min	Input: What is an LLM? The Transformer; “a model is just a file”	Presentation
15–25 min	Input: Probabilistic vs. deterministic - the key concept for VET	Presentation + discussion
25–40 min	Input: Parameters, dense vs. MoE, the global landscape	Presentation
40–50 min	Input: GDPR and data privacy before you prompt	Presentation + settings walkthrough
50–70 min	Practical Task 1: Same task, three places	Individual/pairs
70–85 min	Practical Task 2 + Reality check: AI detectors	Pairs + plenary
85–90 min	Model Selection Guide + wrap-up	Plenary



Large Language Models (LLMs) - Commercial and Open Source

Theoretical Input Part 1: What Is an LLM? (15 min)

One-sentence definition: a Large Language Model is a statistical model trained on huge amounts of text that, given a sequence of words, predicts the most likely next word - and does that one token at a time until the answer is complete.

What makes it work - the Transformer (2017): a neural network architecture from the Google paper “Attention Is All You Need”. The key idea is *attention*: for each new token, the model decides which previous words matter most. The same architecture powers GPT, Claude, Gemini, Llama and Mistral. You don’t need the mathematics; you do need to know that it **predicts - it does not look up**.

How generation actually happens:

- 1. You write: “Explain G02 to a Level-2 apprentice...”
- 2. The model splits your text into tokens: [Explain] [G] [02] [to] [a] [Level] [-2] ...
- 3. It predicts the most likely next token.
- 4. It adds that token and predicts again: → “is” → “a” → “clockwise” → ...
- 5. It stops when the answer is complete.

An AI model is essentially just a file. Under the hype, an LLM is a very large file storing billions of numbers (weights). To use one, three things come together:

- 1. **The weights** - the file itself (e.g., “llama-3.1-8b.gguf”, 4.5 GB).
- 2. **The program** - software that reads the weights, “the engine” (Ollama, LM Studio, llama.cpp).
- 3. **Inference** - running your input through the weights to produce output. Each prompt is one inference run - it costs compute and electricity.

Theoretical Input Part 2: Probabilistic vs. Deterministic (10 min)

This is the single most important concept in the block for VET professionals, because vocational training is built on deterministic systems.

Deterministic - what you know	Probabilistic - what AI is
Same input → same output, always	Same input → similar output, never identical
CNC: G00 X10 Y20 → always the same coordinate	Ask “Explain G02” twice → two different (similar) answers
Calculator: 2 + 2 → 4, forever	The model guesses the next word - like very advanced autocorrect
Spreadsheet: =SUM(A1:A10) → the same total	It does NOT do math, look up facts or query a database
Predictable, repeatable, safe	When wrong, it sounds confident - “hallucination”

Why this matters: once you see AI as a very good *guesser* - not a database - you stop being shocked when it invents a fake safety code or misquotes a standard. It becomes a brainstorming partner, not a calculator. **Treat every output as a draft to verify.**



Large Language Models (LLMs) - Commercial and Open Source

Discussion (5 min): ask each participant to name one deterministic system in their own workshop or office, and describe what would happen if it became probabilistic. (A probabilistic fire alarm? A probabilistic brake test?) This lands the concept permanently.

Theoretical Input Part 3: Parameters, Architectures and the Global Landscape (15 min)

Parameters - why size matters (and when it doesn't)

A parameter is a learned number inside the model, set during training. Roughly: more parameters = more capacity to remember and reason. But bigger is not always better - specialised smaller models often beat huge generalist ones for a specific job, and bigger models cost more to run (compute, memory, electricity, money per question).

The model size ladder (approximate): GPT-2 (2019) - 1.5B, the first headline LLM · Mistral 7B / Llama 3 8B - runs on a strong laptop · Llama 3 70B - needs a serious workstation · GPT-3.5 - 175B, the first ChatGPT, cloud-only · Today's frontier models (GPT, Claude, Gemini) - estimated 1 trillion+ parameters.

Rule of thumb: frontier model for hard tasks, small model for repeatable ones.

Dense vs. Mixture-of-Experts (MoE)

- Dense models:** every parameter fires for every token. Example: Llama 3.1 405B uses all 405B parameters per word. Heavy on memory and compute.
- Mixture-of-Experts:** the model contains many “experts”, and only a few fire per token. Example: DeepSeek-V3 - 671B total parameters, only ~37B active per word. Similar capability, far less compute. This is the dominant modern design - and why headline parameter counts can mislead.

Who builds frontier LLMs - the four regions

Region	Character	Key players
USA	The heavyweights - closed, ultra-large MoE	OpenAI (GPT series), Anthropic (Claude), Google (Gemini); Meta is the open-weights exception (Llama)
China	The efficiency champions - frontier quality at a fraction of the cost, often open-weight	DeepSeek (V3/R1), Alibaba (Qwen series)
Europe	The enterprise contenders - privacy-first, EU hosting, European languages	Mistral AI (FR), Aleph Alpha (DE)
Middle East	The open-source surprise - state-funded, open-weight	TII UAE (Falcon series)

The European ecosystem - your GDPR-friendly alternatives

Provider	Origin	What it offers	Best for
Aleph Alpha	Heidelberg, DE	Explainable AI that cites the exact source paragraph; government/industry focus	Strict GDPR - data never leaves the 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 European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.
Pr Nr.: 2025-1-DE02-KA210-VET-000364083

Large Language Models (LLMs) - Commercial and Open Source

DeepL	Cologne, DE	Translator + DeepL Write for professional writing; localised precision US models can't match	Grant applications, official letters
Nyonic	Berlin, DE	Industrial foundation models for engineering and manufacturing	Industry-specific deployments
Silo AI (AMD)	Helsinki, FI	Poro and Viking - open-source LLMs for EU languages, incl. low-resource ones like Greek	On-premise, multilingual VET
Mistral AI	Paris, FR	Open and commercial models, efficient MoE, EU hosting	Top-tier capability without a US cloud
Kyutai	Paris, FR	Moshi - real-time voice AI that runs locally, open-science	Accessibility, voice-based simulators

Choosing by the job: brainstorming → ChatGPT/Claude · personal or institutional data → Mistral or a local Llama · code → DeepSeek/Claude · EU-language translation → DeepL.

Theoretical Input Part 4: Data Privacy and GDPR - Before You Prompt (10 min)

The core rule: if the tool is free, you are the product. Prompts entered into free consumer tools may be used to train future models.

Two flavours of the same AI:

- **Consumer (free):** free ChatGPT, Claude, Gemini. Your prompts may train the next model. Find the toggle and switch training OFF in Settings (walk participants through it live: ChatGPT → Settings → Data Controls → “Improve the model for everyone” → off).
- **Enterprise / Team:** ChatGPT Team, Claude Team, etc. Contractually not used for training. The right choice for any institutional data.

Never paste into a free tool: trainee names · grades · ID numbers · disciplinary notes · unpublished exam papers · health information · internal institutional documents not cleared for release.

Anonymisation as a working habit: replace names with roles (“Student A, 17, second-year electrical apprentice”), strip identifying details, and reconstruct after generation.

Practical Task 1: Same Task, Three Places (20 min)

Goal: run the *same* prompt in three environments and compare quality, speed, privacy and cost.

The three environments:

1. **Commercial cloud** - chatgpt.com or claude.ai
2. **Open-source in the cloud** - e.g., a hosted open-weight model (Mistral’s Le Chat, DeepSeek chat, or HuggingChat)
3. **Local model** - LM Studio or Ollama on a laptop (facilitator can demo on one machine if installs are limited: ollama run llama3:8b)

Choose ONE of the three prompts (or adapt to your subject):



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 European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.
Pr Nr.: 2025-1-DE02-KA210-VET-000364083

Large Language Models (LLMs) - Commercial and Open Source

- **Technical:** “Explain how a lamp works (or a topic from your own subject) to an 18-year-old apprentice with six weeks of training, using one practical analogy from day-to-day life.”
- **Teaching:** “Draft a short email to a student who failed the last exam. Make it strict but supportive.”
- **Day-to-day:** “Help me explain to a person I meet for the first time what I do for a living.”

Comparison sheet (participants complete per environment): quality of the answer (1–5) · speed · did it follow the instructions? · would you trust it with sensitive data? · what would this cost at scale?

Debrief: typically the frontier commercial model wins on quality, the local model wins on privacy, and the open-source cloud model sits between. The lesson is not “which is best” but “**which is right for which job**”.

Practical Task 2 + Reality Check: The Myth of AI Detectors (15 min)

Mini-experiment (pairs, 8 min): one partner writes 3–4 sentences on a familiar work topic by hand; the other generates a similar paragraph with an LLM. Swap with another pair - can they tell which is which? Tally the guesses.

Input (7 min): the research consensus on AI detectors.

- **They are unreliable.** False positives are common - real student work gets flagged as AI.
- **They are biased.** They disproportionately flag non-native English speakers.
- **They don’t scale.** Even where detection partially works, the cost-benefit isn’t there.

What to do instead: design assignments where the *process* matters, not just the final essay - oral defence of submitted work, in-class application, versioned drafts, personal context that a generic model cannot know, and practical demonstration (VET’s native strength).

Handout: Model Selection Guide (one page)

Your situation	Recommended choice	Why
Brainstorming lesson ideas, no sensitive data	Free ChatGPT / Claude / Le Chat	Best quality-to-effort; no data risk if content is generic
Anything with trainee or institutional data	Enterprise tier, EU provider, or local model	GDPR; contractual no-training guarantees
Offline workshop, confidential industry blueprints	Local model via LM Studio / Ollama	Data never leaves the machine
Translating official documents into/from EU languages	DeepL	Purpose-built, localised precision
Repeatable simple task at volume	Small model (7–8B)	Cheaper, faster, good enough
Hard reasoning, complex drafting	Frontier commercial model	Capability gap is real at the top end



Large Language Models (LLMs) - Commercial and Open Source

Assessment and Reflection

- **Formative check:** collected comparison sheets from Task 1 show whether participants can evaluate models on multiple criteria rather than a single “which is smarter” axis.
- **Exit ticket:** “Name one task from your work you would give to a local model, and one you would only give to a frontier model - and why.”
- **Homework:** install LM Studio or Ollama at home; download one small model; bring one observation to the next module.



BLOCK A · FOUNDATIONS OF AI AND LARGE LANGUAGE MODELS

MODULE 4

Effective Prompt Engineering and Reverse Prompting

Effective Prompt Engineering and Reverse Prompting

Module Overview

Duration	90 minutes
Format	Short mindset input followed by three rounds of intensive pair practice
Prerequisites	Modules 2–3 (vocabulary; probabilistic nature of LLMs)
Materials	Laptop per participant/pair, LLM access, printed C.R.A.F.T. template cards, printed “weak prompts” and “polished outputs” sets for the activities, Prompt Library template

Learning Objectives

By the end of this module, participants will be able to:

1. Explain why prompt quality determines output quality (GIGO) using the Intern analogy.
2. Construct professional prompts with the C.R.A.F.T. framework: Context, Role, Action, Format, Target audience.
3. Apply three advanced techniques: few-shot prompting, negative prompting and reverse prompting.
4. Diagnose a weak prompt and systematically improve it.
5. Reconstruct the probable prompt behind a given output (“reverse engineering”) and write a better one.
6. Start a personal prompt library with at least five tested, refined prompts for their own work.

Session Plan

Time	Segment	Method
0–10 min	Discussion: How do you prompt now? Bad habits, best tricks	Plenary
10–25 min	Input: The mindset shift - the Intern analogy + GIGO	Presentation
25–40 min	Input: The C.R.A.F.T. framework with a full worked example	Presentation
40–50 min	Input: Three advanced techniques	Presentation
50–80 min	Practical Task: The Prompt Workshop - three rounds in pairs	Pairs
80–90 min	Prompt library setup + wrap-up	Individual

Opening Discussion: How Do You Prompt Now? (10 min)

Two flipchart columns filled by the group: “What are the bad practices?” and “What works best for you?” Typical bad practices to surface: one-line vague prompts, no audience specified, accepting the first answer, no format instruction, treating the model as a search engine. Keep the flipchart visible - by the end of the session participants should see everything on the “bad” side has a named remedy.



Effective Prompt Engineering and Reverse Prompting

Theoretical Input Part 1: The Mindset Shift (15 min)

Stop “Googling”. Start delegating.

The Intern analogy: prompting is exactly like delegating to an eager, highly literate, but entirely inexperienced 19-year-old intern.

- Give a vague brief → you get Wikipedia-level fluff.
- Give them audience, format, an example and the safety manual → you get gold.
- Their effort scales with **your** clarity.
- If the output is poor, the prompt was poor.

Garbage In, Garbage Out (GIGO) - the demonstration pair:

- **Bad prompt:** “Make a quiz about electrical circuits.” → generic, off-target, you reject it.
- **Good prompt:** “Act as an electrotechnology lecturer. Build a 10-question quiz on Ohm’s Law for Year-1 apprentices in Germany. 6 multiple-choice + 4 short answer. Include one safety question. Format as a numbered list with the answer key at the end.”

Run both live in front of the group and compare outputs side by side. The difference sells the whole module.

Theoretical Input Part 2: The C.R.A.F.T. Framework (15 min)

Every professional prompt should contain most of these five elements:

Letter	Element	Worked example (student internship support package)
C	Context	“I teach vocational education students aged 16–20. They are preparing for internships and workplace-based learning. The group has mixed levels of confidence, communication skills, and language ability. Some students struggle to reflect on their own progress and explain what they have learned.”
R	Role	“Act as an experienced vocational education teacher, student mentor, and assessment designer.”
A	Action	“Create a student support package that helps learners prepare for, document, and reflect on their internship experience. Include a simple student-friendly explanation of why reflection is important.”
F	Format	“Present the output as structured sections with tables where useful. Keep the sections easy to copy into a handout or worksheet.”
T	Target audience	“The materials should be suitable for vocational students aged 16–20. Use simple, practical, encouraging language. Avoid academic theory and long explanations.”

Teaching notes: not every prompt needs all five at full length - but when output disappoints, the missing letter is usually the diagnosis. Context and Target audience are the two most commonly omitted, and the two with the largest effect.



Effective Prompt Engineering and Reverse Prompting

Theoretical Input Part 3: Three Advanced Techniques (10 min)

- 1. Few-shot prompting - show the model what “correct” looks like before asking.** Example: paste a previous lesson plan and write: “Analyse my style and format. Now write a NEW lesson on [topic] matching this structure.” The model imitates your demonstrated standard instead of guessing at one.
- 2. Negative prompting - tell the model what NOT to do.** Example: “Do not use academic jargon. Do not include equipment that is not listed in my context. Do not invent standards or norm numbers.” Negative constraints are the most direct anti-hallucination tool available at the prompt level.
- 3. Reverse prompting - don’t know what to ask? Have the AI ask you.** Example: “You are an expert curriculum designer. Ask me 5 questions, one at a time, to gather what you need to build my practical exam.” The model becomes the interviewer; you supply expertise; the final output is built from *your* answers rather than the model’s assumptions. This is the highest-leverage technique for complex, underspecified tasks.

Practical Task: The Prompt Workshop - Three Rounds in Pairs (30 min)

■ Round 1 (10 min): Improve a weak prompt using C.R.A.F.T.

Each pair receives one deliberately weak prompt, for example: “Make a lesson about safety.” · “Write a test about metals.” · “Help me with an email to parents.” Run the weak prompt as-is, note the output, then rebuild it with all five C.R.A.F.T. elements and run again. **Deliverable:** before/after outputs plus one sentence naming which letter made the biggest difference.

■ Round 2 (10 min): Role-specific task + one advanced technique

Pick the task matching your role, and apply at least one advanced technique:

- **Technical trainer:** scenario-based diagnostic roleplay - “Act as a broken engine (or faulty circuit, failed weld). Give symptoms only. My students will interview you to find the fault. Reveal the cause only when they ask the right question.”
- **Evaluator/assessor:** a strictly formatted assessment rubric (negative prompting: “do not use vague descriptors like ‘good’ - every level must be observable”).
- **Administrator:** a parent/employer email template for a complex change (few-shot: paste one previous well-received email as the style example).

■ Round 3 (10 min): Reverse engineering

Each pair receives a polished AI output (facilitator prepares 3–4: e.g., a well-structured safety briefing, a differentiated worksheet, a professional rejection email). **Task:** reconstruct the prompt that probably produced it - then write a *better* prompt and run it. Compare your result to the original artefact. **Debrief question:** what did the original prompter forget that you added?



Effective Prompt Engineering and Reverse Prompting

Closing: Build Your Prompt Library (10 min)

Every participant creates a personal prompt library (a simple document or notes file) and stores today's best prompts using this template:

Field	Content
Prompt name	e.g., "Differentiated worksheet generator"
Full prompt text	The complete, final version
Model used	e.g., Claude / ChatGPT / local Llama
What it produces	One-line description
Quality notes	What to check/verify in the output every time

Minimum to leave the room: five prompts - one C.R.A.F.T. showcase, one few-shot, one negative, one reverse-prompting starter, one role-specific. This library grows through Blocks B and C and becomes part of the personal project in Module 12.





Thank you.

Block A · Foundations of AI and Large Language Models

KA210-VET · Erasmus+ · Pr Nr.: 2025-1-DE02-KA210-VET-000364083



Co-funded by
the European Union



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 European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.
Pr Nr.: 2025-1-DE02-KA210-VET-000364083