



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

BLOCK D

AI Strategy, Ethics, Funding and Sustainability

Modules 13–16

Work Package 2 - Training Modules Development

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



Co-funded by
the European Union



Emphasys
TRAINING CENTRE



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

AI Strategy, Ethics, Funding and Sustainability

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 D

Block D focuses on the responsible, ethical and sustainable use of AI and its future impact on education and careers. Participants explore responsible AI practices, learn to identify and respond to deepfakes and misinformation, examine the environmental impact of AI, and consider how AI is changing skills and career pathways.

By the end of the Block, participants will be better prepared to **use AI responsibly and sustainably**, address emerging AI-related risks, and plan their own future professional development.

Block D includes four modules:

Module 13: AI Ethics and Responsible Use in Education

Module 14: Deepfakes and Fake News - Risks and Detection

Module 15: Environmental Impact and Sustainability of AI

Module 16: Personalized Career Pathways



Contents

13

MODULE 13

AI Ethics and Responsible Use in Education**14**

MODULE 14

Deepfakes and Fake News - Risks and Detection**15**

MODULE 15

Environmental Impact and Sustainability of AI**16**

MODULE 16

Personalized Career Pathways

AI Strategy, Ethics, Funding and Sustainability

Block Overview

Block D is the responsibility and future block. Having learned to use AI (Blocks A–B) and to build with it strategically (Block C), participants now confront the questions that determine whether AI adoption in their institutions is sustainable: What is ethical use in education, and how do we operationalise it? (Module 13). How do we recognise and teach others to recognise synthetic deception? (Module 14). What does AI cost the planet, and how do we use it responsibly? (Module 15). And finally: what does AI mean for each participant's own career and for the career guidance they give trainees? (Module 16).

Modules	13. AI Ethics and Responsible Use in Education 14. Deepfakes and Fake News - Risks and Detection 15. Environmental Impact and Sustainability of AI 16. Personalized Career Pathways
Total duration	Each Module is planned for 90 minutes
Target group	VET educators, trainers, mentors, administrators; Blocks A–C completed
Materials needed	Laptops, projector, printed Responsible AI Checklists, deepfake example set prepared offline, printed Career Pathway Canvas, flipcharts
Block outcome	Each participant leaves with a signed personal Responsible AI Checklist, a deepfake-detection lesson outline usable with trainees, a personal AI-sustainability policy, and a personalised career development plan including funding sources

A note on regulation: this block references the EU AI Act and GDPR as the governing framework for European VET institutions. Regulatory detail evolves; facilitators should verify the current implementation status of AI Act provisions relevant to education before each delivery.



BLOCK D · AI STRATEGY, ETHICS, FUNDING AND SUSTAINABILITY

MODULE 13

AI Ethics and Responsible Use in Education

AI Ethics and Responsible Use in Education

Module Overview

Duration	90 minutes
Format	Risk-mapping discussion + case-based dilemma work + checklist construction
Prerequisites	Blocks A–C (the module operationalises rules previewed throughout)
Materials	Printed dilemma cases (6), Responsible AI Checklist poster + individual copies, flipcharts, sticky notes

Learning Objectives

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

1. Name and illustrate the eight core risks of AI in education: copyright, deepfakes, voice cloning, misinformation, bias, privacy, student dependency, institutional risk.
2. Locate their own AI practice within the EU regulatory frame (GDPR; EU AI Act risk categories as relevant to education).
3. Reason through concrete ethical dilemmas of AI in VET using a structured method (stakeholders → harms → alternatives → decision → disclosure).
4. Apply the ten-point Responsible AI Checklist before any AI-generated content leaves their machine.
5. Distinguish AI as *learning support* from AI as *learning replacement*, and design tasks accordingly.
6. Draft the outline of an institutional AI-use agreement for their own team.

Session Plan

Time	Segment	Method
0–15 min	Risk map: which risks feel real in YOUR work?	Pair task + plenary
15–35 min	Input: The eight risks, the regulatory frame	Presentation
35–60 min	Practical Task 1: The Dilemma Cases	Small groups
60–75 min	Practical Task 2: Build and sign your Responsible AI Checklist	Individual + plenary
75–88 min	Practical Task 3: The institutional agreement sketch	Small groups
88–90 min	Wrap-up	Plenary

Opening Pair Task: The Risk Map (15 min)



AI Ethics and Responsible Use in Education

Before presenting any list, ask pairs (10 min): “Which risks of AI feel most real in your daily life and work? Share one concrete example with your partner.” Then collect on flipchart. Then run the mirror question (educators are not doom-mongers): “Which *benefits* feel most real?” Both lists stay visible - ethics in this module means *keeping the benefits while managing the risks*, not abstinence.

Theoretical Input: The Eight Risks and the Regulatory Frame (20 min)

■ Eight risks every educator using AI should be able to name

1. **Copyright.** Training data may include protected works; output ownership is contested; check tools’ terms for commercial/educational use (practised in Module 7).
2. **Deepfakes.** Realistic synthetic images/videos of real people erode trust and may breach likeness rights and GDPR (dedicated treatment: Module 14).
3. **Voice cloning.** Cloned voices can impersonate; consent is mandatory; several EU jurisdictions are tightening laws fast (consent protocol: Module 6).
4. **Misinformation.** Fluent AI output can spread false claims at scale - including unintentionally, via unverified hallucinations (Module 2’s “fluent but wrong”).
5. **Bias.** Models reflect their training data, including its stereotypes: gendered image generation, cultural defaults, weaker performance in smaller languages. In education, biased materials scale to whole cohorts.
6. **Privacy.** Free tools may train on submissions; personal trainee data never enters consumer AI (Modules 3, 10 rules).
7. **Student dependency.** Over-use erodes foundational skills. The design question: is AI a *tool within* the learning task, or a *shortcut around* it?
8. **Institutional risk.** Reputation, GDPR fines, exam integrity. Document your use; align with institutional policy - and where none exists, help write it (Task 3).

■ The regulatory frame in brief

- **GDPR** governs all personal data of trainees and staff: legal basis, minimisation, the “would I defend this processing to a parent and a regulator?” test. Practical consequences already practised: anonymisation, enterprise/local routes, no personal data in free tools.
- **EU AI Act** takes a risk-based approach. Relevant orientation for education: certain uses (emotion recognition in schools, social scoring) fall under prohibited or high-risk categories; AI systems used in *assessment and admission decisions* attract high-risk obligations; everyday generative-AI assistance for content creation sits in lower-risk territory but carries transparency expectations. **Practical takeaway for VET:** the closer AI moves to *decisions about people* (grading, admission, competence sign-off), the heavier the obligations - which matches the red lines drawn in Module 1 on day one.
- **Institutional layer:** policy, works-council/staff agreements, exam regulations. The law is the floor, not the ceiling.



AI Ethics and Responsible Use in Education

■ Dependency deserves special attention in VET

VET's mandate is *competence* - things a person can reliably do. A trainee who lets AI write every report never learns technical writing; one who lets AI do every calculation never builds estimation sense that catches dangerous errors. **Design heuristic:** AI may support the *practice* of a skill (feedback, examples, variation) but must not replace the *performance* of a skill that is itself the learning objective. Where AI does the performing, the assessment must move to where AI cannot follow (oral defence, live demonstration, in-person practical - Module 3's detector lesson).

Practical Task 1: The Dilemma Cases (25 min)

Groups of 3–4; each group takes two cases. **Method (poster per case):** Stakeholders → possible harms → possible alternatives → your decision → what you would disclose, to whom.

Case 1 - The efficient assessor. A colleague pastes trainee reports (with names) into free ChatGPT to generate feedback, and grades based on it. The feedback quality is visibly good. What do you do?

Case 2 - The perfect portfolio. A trainee submits a flawless project documentation. You suspect heavy AI use, but there is no policy and no reliable detector (Module 3). The practical work itself was adequate. Grade? Confront? Redesign?

Case 3 - The cloned announcer. Your director wants the school's popular retired announcer's voice cloned for all future announcements "as a tribute". The announcer is reachable but hasn't been asked.

Case 4 - The biased poster. The AI-generated recruitment poster series for your technical programme shows, across all 6 iterations, only young men at the machines. Deadline is tomorrow; regeneration costs an evening.

Case 5 - The helpful tutor. A struggling trainee has quietly started using an AI tutor at home and improved markedly. Their written work now exceeds what they can explain orally. Parents are thrilled. Exam is in four weeks.

Case 6 - The data-driven early warning. Your data analysis (Module 10) can flag trainees at risk of dropping out with ~80% accuracy. Management wants the list. False positives get well-meant but stigmatising interventions; the flagged trainees never consented to being scored.

Plenary harvest (8 min of the 25): each group presents ONE case in 90 seconds. Facilitator draws the pattern: nearly every dilemma resolves through the same moves - *consent, disclosure, human judgement retained at the decision point, and redesign of the task rather than policing of the tool*.

Practical Task 2: The Responsible AI Checklist (15 min)

Present the ten-point checklist; each participant then adapts it (add/adjust up to two items for their context), signs it, and keeps a copy. Poster version goes to each institution.

Before AI-generated content leaves your machine:

1. **DISCLOSED** - I have made clear this content is AI-generated (where appropriate).



AI Ethics and Responsible Use in Education

2. **VERIFIED** - I have checked the facts, names, codes, standards myself.
3. **LICENSED** - I have read the tool's terms for commercial/educational use.
4. **NO LIKENESS** - no real person's face or voice is used without written consent.
5. **NO PERSONAL DATA** - no trainee names, grades or IDs were pasted into a public tool.
6. **BIAS-CHECKED** - I have reviewed the output for cultural/linguistic/gender bias.
7. **PROMPT SAVED** - I have saved the prompt and parameters so I can reproduce and explain.
8. **HUMAN IN THE LOOP** - I personally reviewed and approved the final version.
9. **POLICY-ALIGNED** - this use aligns with my institution's AI policy.
10. **DEPENDENCY-SAFE** - this use does not erode a foundational skill my trainees should be building.



BLOCK D · AI STRATEGY, ETHICS, FUNDING AND SUSTAINABILITY

MODULE 14

Deepfakes and Fake News - Risks and Detection

Deepfakes and Fake News - Risks and Detection

Module Overview

Duration	90 minutes
Format	Detection lab + verification-skills training + teaching-transfer workshop
Prerequisites	Modules 6, 8 (participants have themselves produced synthetic voice and video - this module is the mirror)
Materials	Prepared example set: mixed real/synthetic images, audio clips and video snippets (curated in advance, offline copies; use public research examples and self-made samples from Modules 6/8 with consent); laptops; printed SIFT method cards; Detection Lesson template

Learning Objectives

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

1. Explain how deepfakes (face swap, lip-sync, voice clone, full synthesis) and AI-generated text disinformation are produced - connecting to their own Block B production experience.
2. Describe why detection-by-eye is failing as generation quality improves, and name the visual/audio artefacts that still sometimes betray fakes.
3. Apply verification methods that do not depend on artefact-spotting: SIFT (Stop, Investigate the source, Find better coverage, Trace to the original), reverse image search, provenance signals (C2PA/content credentials).
4. Assess the specific risks for their context: scams targeting institutions, fake evidence in disputes, synthetic bullying among trainees, reputational attacks.
5. Design and outline a detection-and-verification lesson for their own trainees.
6. Formulate the institutional response protocol: what to do when a suspected deepfake involving the institution or a trainee appears.

Session Plan

Time	Segment	Method
0–15 min	Detection Lab round 1: real or fake? (baseline test)	Individual voting
15–35 min	Input: How fakes are made; why eyes fail; what still works	Presentation
35–50 min	Detection Lab round 2: verification with method	Pairs
50–65 min	Input + discussion: the risk landscape for VET	Presentation + plenary
65–85 min	Practical Task: Build your detection lesson	Pairs
85–90 min	Institutional response protocol + wrap-up	Plenary



Deepfakes and Fake News - Risks and Detection

Detection Lab Round 1: The Baseline Test (15 min)

Show 8 items (2 images, 2 audio clips, 2 short videos, 2 news-style paragraphs) - half real, half synthetic. Participants vote real/fake for each, privately noting confidence. Reveal answers and score.

Expected and intended result: the room scores near chance on the good fakes, and - importantly - *misclassifies some real items as fake*. Both errors matter: the second is the “liar’s dividend” (once fakes exist, real evidence can be dismissed as fake). This experience, not any lecture, is the module’s foundation.

Theoretical Input: How Fakes Are Made, Why Eyes Fail, What Still Works (20 min)

■ The production side - which you already know

Connect explicitly to Block B: participants have already cloned a voice (Module 6) and generated video (Module 8). Deepfake production is the same technology without consent: **face swap** (a face mapped onto another body), **lip-sync/reenactment** (real footage, new mouth movements matching new audio), **voice cloning** (1–3 minutes of audio suffice - Module 6), **full synthesis** (text-to-video of events that never happened), and **text disinformation at scale** (LLMs producing coordinated fluent content - Module 2’s “fluent but wrong”, weaponised).

■ Why detection-by-eye is a losing game

Every artefact taught in older trainings - weird hands, mismatched blinking, flat lighting - is a generation defect, and generation defects get fixed release by release. Artefacts still worth *glancing* for (hands and jewellery, teeth, hair edges, physics of reflections, audio breathing patterns and room acoustics) can *raise* suspicion - but their absence proves nothing. **Teach the asymmetry: artefacts can help you doubt; only provenance can help you trust.**

■ What still works: verification instead of inspection

- **SIFT - the method to teach everyone.** Stop (don’t share, don’t react). Investigate the source (who published this, what is their track record?). Find better coverage (do reliable outlets report it? lateral reading). Trace to the original (find the earliest version; context is stripped by re-shares).
- **Reverse image/video search** (Google Lens, TinEye): finds earlier versions and original contexts; catches the most common fake - *real* media recaptioned in a false context (“cheapfakes”), which outnumber true deepfakes.
- **Provenance signals:** C2PA / Content Credentials attach signed origin metadata to media; major cameras, phones and AI tools are adopting them. Absence proves nothing (metadata strips easily) - presence is meaningful.
- **The out-of-band rule for high stakes:** any urgent request involving money, credentials or reputation, arriving via any single channel (a voice call “from the director”, a video message “from a parent”) is verified through a *second, independent channel* before action. This one habit defeats nearly all current voice-clone scams.



Deepfakes and Fake News - Risks and Detection

Detection Lab Round 2: Verification with Method (15 min)

Rerun 6 new items - but now pairs may use the full toolkit: SIFT card, reverse image search, source checking, tracing. Compare scores with round 1. **Debrief:** which method contributed most? (Almost always: tracing to the original and lateral reading - not artefact hunting.) Note the time cost per item - verification is minutes, not seconds, which is exactly why “Stop” (don’t insta-share) is the first letter.

Input + Discussion: The Risk Landscape for VET (15 min)

Four risk scenarios, each discussed briefly with the room (“has anything like this touched your institution?”):

1. **Scams against the institution:** voice-cloned “director” requesting urgent transfers or credentials; fake supplier videos. Defence: out-of-band verification as *institutional policy*, not personal habit.
2. **Synthetic bullying among trainees:** fabricated images/audio of classmates or teachers. Defence: clear sanctions, reporting channel, preserved-evidence protocol, and education *before* the first incident.
3. **Fake evidence in disputes:** doctored audio of “what the instructor said”, synthetic proof in exam disputes. Defence: documentation culture; awareness that both fake evidence *and* false “that’s a deepfake” claims (liar’s dividend) will appear.
4. **Reputational attacks and disinformation:** fake news about the institution or sector spreading through parent/community channels. Defence: a designated response owner and a pre-agreed statement workflow.

Practical Task: Build Your Detection Lesson (20 min)

Pairs design a trainee-facing lesson outline (45 or 90 min) using the Detection Lesson template - because the strongest institutional defence is a trained community.

Template sections:

1. **Hook:** a real-or-fake test appropriate to the trainee group (participants may reuse today’s structure; sample materials list included).
2. **One production demo:** show how easily a voice or image is generated (from Block B skills) - demystification kills both naivety (“videos don’t lie”) and panic (“nothing can be trusted”).
3. **The method:** SIFT + reverse search + out-of-band rule, drilled on 3 exercises relevant to the trade (a fake supplier offer for business trainees; a fake safety-recall notice for technical trades; a fake casting call for creative trades).
4. **The norms conversation:** what our class/school does about synthetic content of real people - co-created rules, not imposed ones.
5. **Assessment idea:** trainees verify one viral claim per week for a month (verification journal).

Share-back: two volunteer pairs present their hook + one exercise.



BLOCK D · AI STRATEGY, ETHICS, FUNDING AND SUSTAINABILITY

MODULE 15

Environmental Impact and Sustainability of AI

Environmental Impact and Sustainability of AI

Module Overview

Duration	90 minutes
Format	Evidence-based input + footprint estimation exercise + policy-writing workshop
Prerequisites	Blocks A–C (participants know the difference between training and inference, and between frontier and small models - Modules 2, 3, 9)
Materials	Laptops, printed Footprint Estimation worksheet, calculator, Green AI Decision Tree handout, flipcharts

Learning Objectives

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

1. Distinguish the environmental costs of AI across its lifecycle: hardware manufacturing, model training, and inference (usage) - including electricity, water and e-waste.
2. Put AI's footprint in honest proportion: per-query impact vs. aggregate impact vs. other digital activities, and why numbers vary wildly between sources.
3. Apply footprint-reduction levers available to an ordinary user: right-sizing the model (Module 9's SLMs), batching and reuse (prompt libraries, saved outputs), avoiding regeneration waste, and questioning genuinely pointless uses.
4. Evaluate provider sustainability claims critically (renewable-energy claims, efficiency reporting, greenwashing signals).
5. Draft a proportionate personal/institutional Green AI policy that neither ignores the issue nor collapses into performative abstinence.
6. Integrate AI sustainability into their teaching where their trade touches energy, IT or resource management.

Session Plan

Time	Segment	Method
0–10 min	Opening poll + myth board	Plenary
10–35 min	Input: The real footprint of AI - lifecycle view	Presentation
35–55 min	Practical Task 1: Estimate your own AI footprint	Pairs + worksheet
55–70 min	Input: The reduction levers + reading provider claims	Presentation
70–85 min	Practical Task 2: Write your Green AI policy	Individual/institution groups
85–90 min	Teaching transfer + wrap-up	Plenary



Environmental Impact and Sustainability of AI

Opening Poll and Myth Board (10 min)

Three quick votes (hands): “AI’s environmental impact is (a) negligible, (b) significant but manageable, (c) a dealbreaker.” Then collect claims participants have heard (“every ChatGPT question uses a bottle of water”, “AI will consume a country’s worth of electricity”, “it’s all greenwashing anyway”) on a myth board - to be revisited and sorted (confirmed / exaggerated / it depends) at the end.

Theoretical Input: The Real Footprint - a Lifecycle View (25 min)

■ Three phases, three very different footprints

1. **Hardware manufacturing.** GPUs and data-centre construction: mining of rare materials, chip fabrication (energy- and water-intensive), and a growing e-waste stream as hardware generations turn over quickly. Mostly invisible in “per-query” numbers.
2. **Training.** A one-time (per model) but enormous cost: months of thousands of GPUs. Frontier-model training runs are estimated in the gigawatt-hours - but this cost is *amortised* across billions of subsequent queries. Efficiency note from Module 3: the MoE architecture shift happened largely because compute costs (money *and* energy) forced efficiency.
3. **Inference - where you live.** Each query costs a small amount of electricity (and cooling water, depending on the data centre). Per query: small. At billions of queries per day globally: a rapidly growing share of data-centre demand - and data centres already compete for grid capacity and water in several regions.

■ Honest numbers and why they vary

Teach the *reasoning*, not a single number: per-query energy estimates differ by an order of magnitude between studies because they depend on model size (a 7B local model vs. a frontier MoE), hardware generation and utilisation, data-centre cooling and location, and what is counted (training amortisation in or out?). Useful anchors for classroom proportion: a short LLM text query costs roughly the same order as other everyday digital acts (a search, minutes of streaming) - image and especially **video generation are 1–2 orders of magnitude costlier** - and the *aggregate* problem is real growth in data-centre electricity and water demand, concentrated regionally. **Two errors to avoid:** “one prompt kills a tree” alarmism and “it’s nothing” dismissal. Proportion is the skill.

■ Water and siting

Data-centre cooling consumes water - how much depends heavily on climate and cooling technology. The issue is *local*: a data centre in a water-stressed region is a real conflict; the same centre in a cool, wet region far less so. This is why siting, transparency and regional reporting matter more than global averages.



Environmental Impact and Sustainability of AI

Practical Task 1: Estimate Your Own AI Footprint (20 min)

Pairs complete the Footprint Estimation worksheet:

- 1. Inventory:** list your actual AI usage in a typical week since this training began (text queries, images generated, audio minutes, video seconds - the Block B portfolio counts!).
- 2. Weight it:** apply the worksheet's relative-cost factors (text = 1×, image ≈ 10–50×, video second ≈ 100×+ - order-of-magnitude teaching values, sources noted on the sheet).
- 3. Contextualise:** compare your weighted total against reference anchors on the sheet (hours of video streaming, km of car travel) - *and note the uncertainty ranges*.
- 4. Identify your top lever:** which single change cuts your footprint most? (Almost always: fewer discarded regenerations of images/video - which is also the Module 11 *effort* lever. Efficiency and sustainability point the same way.)

Debrief insight to draw out: for an individual educator, AI's footprint is modest next to commuting or heating - but the *habits* set now scale to whole institutions and cohorts, and educators are multipliers.

Theoretical Input: Reduction Levers and Reading Provider Claims (15 min)

The user's levers (Green AI Decision Tree handout):

- 1. Right-size the model.** Repeatable simple task → small/local model (Module 9), not a frontier model. Same output, fraction of the cost - the decision framework from Module 9 is *also* the green framework.
- 2. Reuse instead of regenerate.** Prompt libraries (Module 4), saved assets (Block B), template agents (Module 9): the greenest query is the one you don't repeat.
- 3. Draft cheap, finalise once.** Iterate wording in text (cheap) before generating the image/video (expensive); storyboard before rendering (Module 8 practice).
- 4. Batch and think first.** One well-crafted C.R.A.F.T. prompt beats five lazy attempts - quality craft *is* efficiency.
- 5. Question the use.** Some uses fail any value test (Module 11 Money Pits) - the Fill-In quadrant deserves a sustainability glance too.

Reading provider claims - three questions: Is the renewable claim *matched annually* (weaker) or *24/7 hourly* (stronger)? Does the provider report water and location-specific data? Are efficiency gains passed to absolute reductions, or eaten by growth (rebound effect)? A claim that survives all three questions is rare - which is itself the lesson in critical reading.



Environmental Impact and Sustainability of AI

Practical Task 2: Write Your Green AI Policy (15 min)

Individually (or in institution groups), draft a 5-line policy - proportionate, practical, defensible:

1. **Default model rule:** e.g., “small/local model for routine tasks; frontier only where quality demonstrably requires it.”
2. **Regeneration rule:** e.g., “text iterations before any image/video generation; max 3 renders per final asset.”
3. **Reuse rule:** e.g., “prompts and assets are saved and shared in the team library.”
4. **Procurement question:** e.g., “sustainability reporting is one criterion in any paid AI tool decision.”
5. **Teaching commitment:** e.g., “AI’s resource cost is mentioned wherever we teach AI use” (one honest sentence, not a guilt lecture).

Volunteers read theirs; the room borrows freely.



BLOCK D · AI STRATEGY, ETHICS, FUNDING AND SUSTAINABILITY

MODULE 16

Personalized Career Pathways

Personalized Career Pathways

Module Overview

Duration	90 minutes (programme finale)
Format	Labour-market input + AI-assisted self-assessment + career canvas + funding clinic
Prerequisites	All Modules 1–15; the Module 12 project as evidence of capability
Materials	Laptops with LLM access, printed Career Pathway Canvas, Funding Source One-Pager (updated per delivery!), participants' Block B/C portfolios, Module 12 canvases

Learning Objectives

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

1. Describe how AI is reshaping career structures in education and in their trainees' industries: which task profiles shrink, which grow, and which new roles emerge.
2. Name emerging AI-related career opportunities and their realistic entry paths (AI engineer, data engineer, ML engineer - and the wider "AI-adjacent" roles more relevant to most: AI-literate trainer, learning technologist, AI coordinator, prompt/workflow specialist within a trade).
3. Use AI itself as a career-development instrument: skills audit, gap analysis, learning-plan generation, CV/profile drafting - with the same verification discipline as everywhere else.
4. Construct a personal career pathway canvas: current strengths → target role/direction → skill gaps → learning actions → evidence portfolio → timeline.
5. Identify concrete funding and support instruments for the next step: Erasmus+ (KA1 mobility, KA2 partnerships), national training funds and vouchers, employer/works-council routes, EU digital initiatives.
6. Advise trainees on AI-era career choices using the same canvas logic, adapted to the trades they teach.

Session Plan

Time	Segment	Method
0–15 min	Input: The AI-era labour market - for educators and their trainees	Presentation + discussion
15–25 min	Input: New and adjacent AI careers	Presentation
25–50 min	Practical Task 1: AI-assisted skills audit and gap analysis	Individual with LLM
50–70 min	Practical Task 2: The Career Pathway Canvas	Individual
70–82 min	Funding clinic: matching pathways to money	Small groups + plenary
82–90 min	Programme close: pathway pitches + certificates	Plenary



Personalized Career Pathways

Theoretical Input Part 1: The AI-Era Labour Market (15 min)

■ The honest pattern: task change, not job apocalypse

Return to Module 1's Future Skills Radar (bring the flipcharts back if kept). The evidence-based pattern: AI automates *tasks* within jobs faster than it eliminates jobs; the value shifts toward tasks AI does poorly - physical skill in unstructured environments, safety judgement and accountability, client trust, supervision of AI output. For VET this is structurally good news: vocational work is disproportionately made of exactly those tasks. But the *mix* inside every occupation changes, and curricula that ignore this produce graduates with yesterday's task profile.

■ What it means for educators specifically

Three shifts in the educator's own profession: **content production is commoditised** (the scarce skill moves from producing materials to curating, verifying and personalising them - the whole arc of this training); **AI literacy becomes a teaching subject embedded in every trade**, making educators who can teach *verification within their trade* newly valuable; and **new institutional roles appear** - AI coordinator, digital learning designer, data-informed quality developer - typically filled by educators who moved early. The participants in this room *are* the early movers; the Module 12 project is their first evidence.

Theoretical Input Part 2: New and Adjacent AI Careers (10 min)

The headline technical roles (relevant context for trainee guidance):

- **AI engineer:** uses AI and machine-learning techniques to develop applications and systems that help organisations become more efficient.
- **Data engineer:** builds systems that collect, manage and convert raw data into usable information for data scientists, analysts and other professionals.
- **Machine learning engineer:** researches, builds and designs the AI responsible for machine learning; maintains and improves existing AI systems.

These need substantial technical education - relevant for *some* trainees (and for VET IT tracks), but they are the visible tip.

The larger, nearer opportunity - AI-adjacent roles: every trade is growing an "AI-competent" variant: the electrician comfortable with predictive-maintenance dashboards, the office administrator who builds the team's agents (Module 9!), the media designer directing generation pipelines (Module 8), the trainer who is the department's verification authority. These roles are reached by *upskilling from an existing trade*, not by career restart - which is precisely VET's home turf, for participants and trainees alike.



Personalized Career Pathways

Practical Task 1: AI-Assisted Skills Audit and Gap Analysis (25 min)

Participants use an LLM as a career-development interviewer - deploying reverse prompting (Module 4) on themselves.

Step 1 - the audit interview (10 min). Prompt: “You are an experienced career coach for education professionals. Interview me, one question at a time (max 8 questions), to build a skills profile: my trade/subject expertise, pedagogical strengths, digital and AI skills (I have just completed a 16-module AI training - I will describe it), languages, and what energises me. Then output the profile as a structured list: Core strengths · Emerging strengths · Underused assets.”

Step 2 - the gap analysis (8 min). Prompt: “Based on my profile, propose 3 realistic development directions for the next 2 years: one ‘deepen current role’, one ‘expand into a new institutional role’, one ‘stretch option’. For each: the 3 most important skill gaps, and how each gap could be closed (course, project, mentoring, certification).”

Step 3 - verification (7 min, in pairs). The Block-wide discipline applies to career advice too: partners cross-examine each other’s AI output - “Which of these gaps is real? Which suggestion is generic fluff? What did the AI miss because it doesn’t know your institution?” Mark the output accordingly. **Career decisions get the same ‘draft to verify’ treatment as everything else.**

Practical Task 2: The Career Pathway Canvas (20 min)

Each participant completes the one-page canvas (using audited, human-verified inputs):

Canvas field	Content
1. Where I stand	Core strengths + the new AI capabilities evidenced this week (portfolio, agent, project)
2. Direction	Chosen development direction (from Task 1, human-verified) - named concretely
3. Gaps	Max 3 skill gaps that matter for the direction
4. Learning actions	One action per gap, each with a format and a date (course/certification, a project, a mentor/community)
5. Evidence portfolio	What will prove the new capability: the Module 12 project outcome, Block B assets, a documented pilot, a workshop I deliver to colleagues
6. Support & funding	Filled during the funding clinic (16.8)
7. Timeline	6-month milestone and 24-month ambition, one line each
8. Trainee transfer	One way I will use this same canvas logic in career guidance for my trainees

Funding Clinic: Matching Pathways to Money (12 min)

Short input, then groups match their canvases to instruments. **Funding Source One-Pager (facilitator updates before every delivery - programmes and calls change):**



Personalized Career Pathways

- **Erasmus+ KA1 (learning mobility):** structured courses and job-shadowing abroad for VET staff - the natural next step after this training; institutions can apply for accreditation for recurring mobility.
- **Erasmus+ KA2 (partnerships, incl. KA210 small-scale like this project):** for participants ready to move from consumer to producer - develop AI-in-VET outputs with European partners; your Module 12 project can be a seed.
- **National instruments (verify per country):** training vouchers/individual learning accounts, digital-pact style programmes for institutions, sectoral funds via chambers and social partners, employer training budgets and works-agreement entitlements.
- **EU digital initiatives:** Digital Europe Programme advanced-digital-skills actions; national Digital Skills and Jobs Coalitions; sector skills alliances.
- **Low/no-cost:** vendor certifications and free courses (note critically: vendor certificates document tool skill, not pedagogy), open university MOOCs, and - highest-value, zero-cost - *internal evidence*: the documented Module 12 pilot presented to management is the strongest funding argument of all.

Group task: each participant writes into Canvas field 6: the best-fit instrument, the concrete next step (call date, application owner, internal approval needed), and asks one question in plenary if blocked. Facilitator collects unanswerable country-specific questions for follow-up.





Thank you.

Block D · AI Strategy, Ethics, Funding and Sustainability

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