AI Literacy at Work — Learning Design Document

Nov 6, 2026 · @vs

Project overview

A 20-minute, self-paced Rise 360 course that teaches office staff to use generative AI tools safely and explains what the EU AI Act expects of their employer. It is a portfolio project: the client context below is a realistic composite, not a real engagement.

Item	Detail
Course	AI Literacy at Work: Using AI Tools Safely under the EU AI Act
Audience	Non-technical office staff in EU organisations that use ChatGPT, Copilot or AI features in their software
Format	Articulate Rise 360, responsive (desktop and phone), web export
Length	About 20 minutes: 4 lessons and a 6-question final quiz
Pass mark	80% (5 of 6), unlimited retakes
My role	Needs analysis, learning design, scriptwriting, legal fact-check, Rise build, image sourcing, QA
Status	Built and exported, November 2026

Needs analysis

The need is legal and behavioural: since 2 February 2025, Article 4 of the AI Act requires organisations that use AI to act on their staff's AI literacy, and staff are already using AI tools without guidance.

Business problem. Any organisation whose staff use ChatGPT or Copilot at work is a *deployer* under the AI Act. Article 4 applies to every deployer, whatever the risk level of its systems. The Digital Omnibus (Regulation (EU) 2026/1744, in force 27 July 2026) reworded the duty to "take measures to support" AI literacy, but a deployer that does nothing still falls short. A short, documented course is one of the simplest measures to evidence.

Performance gap. The risky behaviours are specific and observable, so the course targets behaviour, not legal knowledge:

What staff do now (assumed)	What they should do	Root cause
Paste customer emails or HR data into free public chatbots	Remove personal data; use only approved tools	Don't know that GDPR still applies to AI inputs
Forward AI-written facts, figures and quotes unchecked	Open every source and verify before use	Believe fluent output is accurate
Let a tool rank or reject people	Keep a human deciding and accountable	Don't know hiring is a high-risk use
Stay silent about odd or biased output	Report it to a named contact	No reporting route, no habit

Audience. Office staff across functions (operations, HR, sales, finance), mixed ages, no technical background, English as a working language. They are busy and sceptical of compliance training, so the course must be short, concrete and phone-friendly.

Constraints. 20 minutes maximum; no LMS integration required for the demo; content must be legally accurate and dated, because the AI Act's dates changed in 2026; all images must be free for commercial use with no attribution condition (CC0).

Why e-learning. The content is the same for every learner, must reach many people quickly, and needs a completion record as evidence of an Article 4 measure. Practice of judgement is handled by scenarios rather than live workshops.

Out of scope. Prompt-writing skills, technical AI concepts, and obligations of AI *providers*.

Learning objectives

Four objectives, each written as an observable action and each tested at least once in the final quiz.

#	By the end, learners can…	Taught in	Practised in	Assessed in
1	Explain why a chatbot can give a confident but wrong answer	Lesson 1: flashcards, process	Lesson 1 knowledge check	Quiz Q1, Q5

#	By the end, learners can…	Taught in	Practised in	Assessed in
2	Classify an AI use into the AI Act's four risk levels and state their employer's AI-literacy duty	Lesson 2: tabs, callout, timeline	Lesson 2 sorting activity	Quiz Q2, Q3, Q4
3	Apply five safe-use rules (protect data, check, be open, a human decides, speak up)	Lesson 3: accordion, checklist	Lesson 3 knowledge check	Quiz Q5, Q6
4	Choose the safest action in realistic workplace situations involving AI	Lesson 4 scenario	Lesson 4 scenario (3 decisions)	Quiz Q6

Objective 4 is the one that matters most for behaviour change; objectives 1 to 3 give learners the reasons behind the choices it asks for.

Design approach

The course moves from *why* to *what the law says* to *what to do* to *doing it*: each lesson ends in a check, and the course ends in practice and a quiz.

Instructional strategy. Short lessons of 4 to 6 minutes, one idea each, stated as the lesson's first sentence. Misconceptions are tackled head-on (flashcards pair a myth with the reality) because learners arrive with confident wrong beliefs about AI. Practice comes before the quiz, so the quiz checks transfer rather than recall alone.

Interaction choices and why.

Interaction	Lesson	Why this one
Flashcards (myth / reality)	1	Surfaces and corrects misconceptions quickly
Process (4 steps)	1	Makes "the human checks" a visible step in the workflow
Tabs (4 risk levels)	2	Lets learners compare levels without a wall of text
Sorting activity	2	Classification is the skill objective 2 asks for
Timeline	2	Dates changed in 2026; a timeline shows what applies now and later

Interaction	Lesson	Why this one
Accordion + checklist	3	Five rules, one panel each, then a 3-question pre-flight check to remember
Branching scenario (3 decisions)	4	Judgement under realistic pressure from a manager

Learning-science principles and where each one shows up.

Principle	What it says	Where it is in the course
Pretesting	Trying to answer before you are taught improves later recall, even when the first answer is wrong	Lesson 1 opens with "Before you start: where does a chatbot's answer come from?"
Retrieval practice	Recalling information strengthens memory more than rereading it	A check at the end of Lessons 1 and 3, the sorting activity in Lesson 2, the final quiz
Explanatory feedback	Feedback that says *why* teaches more than right or wrong	Every check, scenario choice and quiz item names the reason or the rule
Segmenting and signaling (Mayer)	Short, labeled chunks lower cognitive load	4-to-6-minute lessons, one idea each; "Note:" and "Why is it important?" labels
Scenario-based practice	Judgement transfers when it is practiced in realistic context	Lesson 4: three decisions under pressure from a manager
Implementation intentions	Planning "when X, I will Y" makes a new habit more likely	Lesson 4 ends with "Your next step": pick one AI task and the rule that applies
Spacing	Retrieval spread over time beats one session	Proposed for rollout: two 1-question reminders at 7 and 30 days (outside Rise)

Scenario design. Mia, an operations coordinator, gets three requests from her manager Tom. Each has a best answer, a safe-but-unhelpful answer and a risky one, because real choices are rarely right versus wrong. Feedback names the rule it applies. The scenario shows dialogue only, with no character photo, to keep every image public domain.

Accessibility. Rise's responsive layout and keyboard-operable blocks; plain language and sentences under 25 words; no meaning carried by colour alone; text alternatives planned for every image; quiz retakes unlimited, with no time limits.

Legal accuracy. Every legal statement was checked against the AI Act and the Digital Omnibus on EUR-Lex (Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744). Lesson 2 ends with "Legal information correct as of 6 November 2026" and a not-legal-advice line, so the content's age is visible to learners and reviewers.

Image licensing. Only CC0 images (no attribution condition, commercial use allowed), with no identifiable people and no visible logos. Each image's source page, author and licence is recorded in a licence log kept with the course files.

Course map and storyboard

Five units run in a fixed order, and every lesson ends in a check. The table below is the screen-by-screen summary. A separate build storyboard holds the full on-screen text for every block.

Unit	Time	Block sequence	Key message on screen
Lesson 1: What AI tools really do	4 min	Text → Flashcards (4 myths) → Process (4 steps, CC0 images) → Knowledge check	"AI predicts, it doesn't know." Step 4 of the process is yours: check the result.
Lesson 2: The EU AI Act in plain language	6 min	Text → Tabs (4 risk levels) → Instruction → Sorting activity → Text (Article 4 duty) → Note callout → Timeline (3 dates) → "Correct as of" note	Using ChatGPT at work makes your employer a deployer; the AI-literacy duty has applied since 2 February 2025.
Lesson 3: Five rules for using AI at work	4 min	Text → Accordion (5 rules) → Checklist note → Multiple-response check	Protect data, check, be open, a human decides, speak up.
Lesson 4: Practice with real situations	4 min	Intro text → Scenario (3 decisions, office background, dialogue labelled "Tom:")	Apply the rules when your manager wants a shortcut.
Final quiz	2 min	6 questions, shuffled, 80% to pass, unlimited retakes	Results screen repeats the five rules.

Timeline dates shown in Lesson 2

Date	What applies
2 February 2025	Bans on unacceptable-risk AI; the AI-literacy duty (Article 4)
2 August 2026	Transparency rules; national authorities supervise AI literacy
2 December 2027	High-risk rules for uses such as hiring and credit (moved from August 2026 by the Digital Omnibus)

Assessment and evaluation

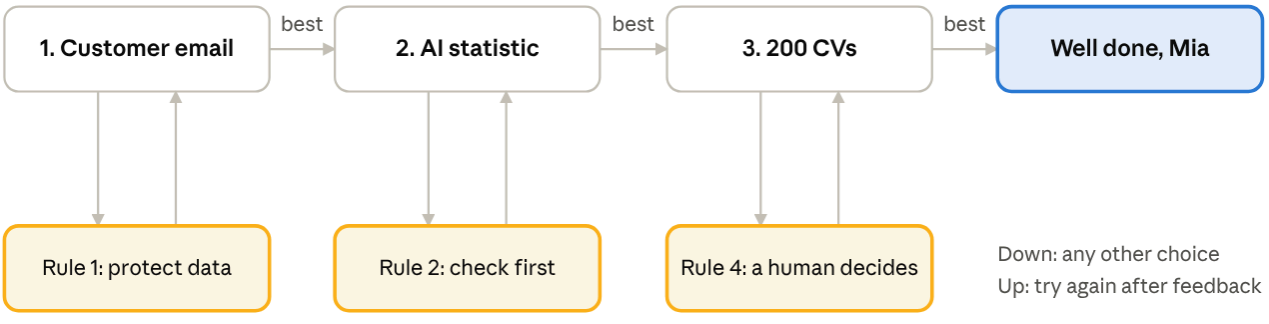
Learners pass at 5 of 6 quiz questions; the organisation judges the course by whether risky AI use falls afterwards, not by completions alone.

Quiz design. Six multiple-choice questions, shuffled, with feedback that names the reason. Four test understanding (Q1 to Q4 cover objectives 1 and 2); two are applied judgement questions (Q5 and Q6). Wrong options are real misconceptions ("It's connected to out-of-date websites", "Certify every employee"), not filler. Unlimited retakes, because the goal is mastery and an evidence record, not ranking.

Scenario as formative practice. In Lesson 4, the best choice moves the story on. A risky or merely-safe choice shows feedback naming the rule, then returns the learner to the same decision. Nothing is scored, so learners can explore the wrong paths safely.

Only the best choice moves Mia on; any other choice loops back

Lesson 4 scenario: three decisions, each with a best, a safe-but-unhelpful and a risky option



Lesson 4 scenario branching · 3 decisions, 1 ending

Each feedback screen names the rule the choice broke or bent, so a wrong turn still teaches.

Evaluation plan (Kirkpatrick levels). A proposed plan for a real rollout:

Level	Question	Measure	When
1 Reaction	Was it relevant and usable?	3-question survey at the end (relevance, clarity, time)	On completion
2 Learning	Did they meet the objectives?	Quiz first-attempt score per question; questions most often missed	On completion
3 Behaviour	Do they use AI more safely?	Requests for approved-tool access; data-loss-prevention alerts for AI sites; reports made to the AI contact	30 and 90 days after rollout
4 Results	Is the organisation safer and compliant?	Completion records as Article 4 evidence; AI-related incidents per quarter	Each quarter

Quiz questions missed by more than 30% of learners on the first attempt would trigger a content review of the lesson that teaches them.

Sources and version notes

Legal content checked 3 and 6 November 2026; dates may change again if the AI Act is amended.

- Regulation (EU) 2024/1689, the AI Act (EUR-Lex) — Article 4 (AI literacy), Article 5 (prohibited practices), Article 50 (transparency), Annex III (high-risk uses)
- Regulation (EU) 2026/1744, the Digital Omnibus on AI (EUR-Lex) — published 24 July 2026, in force 27 July 2026; Annex III high-risk obligations from 2 December 2027
- Images: StockSnap.io, CC0 1.0; licence log kept with the course files

Version	Date	Change
1.1	6 Nov 2026	Replaced all stock images with CC0 images; added sorting instruction and "correct as of" note; scenario dialogue labelled "Tom:"; web export
1.0	5 Nov 2026	All four lessons and quiz built in Rise