

## MODULE 1.3

# Human-Centred Mindset and AI Ethics Essentials

---

*Complete development & upload package for Sikai Chautari (Moodle)*

Stage 1 — Acquire (Foundations) · Strands A1 & A2 — Mindset & Ethics

UNESCO AI CFT 2024 — Human agency · Ethical principles · 12 indicative hours (8 contact + 4 practice)

Grades 9–12 teachers · Kathmandu Valley pilot · Nepali & English

**Version 1.0 · June 2026 · ADB TA 10376-REG**

## Contents

Module overview .....	3
What teachers will be able to do .....	3
How this module fits the programme .....	4
How to use this package .....	4
Four threads that run through every topic .....	4
Part 1 — Moodle course structure (the upload map).....	5
Part 2 — Core content sections (paste-ready).....	8
1 · AI is human-led — agency, hype and the teacher's role .....	8
2 · The core ethical principles for classroom AI.....	9
3 · The anchor rule — AI helps but never certifies.....	9
4 · Privacy and personal data — what never goes into a tool.....	10
5 · Bias and exclusion — gender, caste, language, disability, geography.....	11
6 · Misinformation, fake content and deepfakes .....	12
7 · From principles to practice — your Responsible-AI charter .....	13
8 · The student AI-use agreement — and what's next.....	13
Part 3 — Visual presentation outline (facilitator deck).....	15
Part 4 — Interactive digital elements (H5P & SCORM).....	18
H5P activities (eight) .....	18
SCORM package — Ethical dilemmas: who decides?.....	22
Structured activity — the AI Ethics Debate.....	23
Part 5 — Formative assessment tools.....	24
Module 1.3 knowledge check (14 items) .....	24
Practical-task rubric .....	26
Facilitator observation checklist .....	27
Confidence self-assessment (before & after).....	27
Part 6 — Source mapping & references.....	29
How the sources fed in .....	29
References .....	29

## Module overview

This package contains everything needed to build **Module 1.3 — Human-Centred Mindset and AI Ethics Essentials** as a course on Sikai Chautari (Moodle) and to facilitate it with a cohort of teachers. It is the values floor of the programme: where Module 1.1 made the technology dependable and Module 1.2 made AI honest and clear, Module 1.3 makes its use **principled and accountable** — and turns the programme's two rules into a signed classroom charter.

### AT A GLANCE

<b>Module</b>	1.3 — Human-Centred Mindset and AI Ethics Essentials
<b>Programme &amp; stage</b>	Nepal Teacher AI & Digital Competencies Programme (AIDC) · Stage 1 — Acquire (Foundations)
<b>Competency strands &amp; level</b>	Human-centred mindset (A1) — Human agency · Ethics of AI (A2) — Ethical principles. UNESCO AI CFT 2024. Also seeds the ethics strand of the student AI-literacy curriculum.
<b>Prerequisites</b>	Module 1.2 recommended; can run alongside it. A short self-rating places teachers but excludes no one.
<b>Indicative hours</b>	12 hours (8 contact + 4 supported practice).
<b>Format &amp; platform</b>	Online; hosted on Sikai Chautari (Moodle). <b>Fully unplugged-capable</b> — this module needs no specialised hardware; every activity has an online, an offline and a shared-device route.
<b>Languages</b>	Nepali and English, with locally grounded scenarios. The charter and the student agreement are written bilingually.
<b>Group size</b>	Varies by cohort; the course runs with one facilitator.
<b>Assessment</b>	Auto-marked H5P checks, a structured <b>AI Ethics Debate</b> , a graded quiz (Module 1.3 knowledge check, 14 items), a rubric for the two practical tasks, a facilitator observation checklist, and a confidence self-assessment that feeds the Stage 1 portfolio.
<b>Certificate path</b>	Completion contributes toward the <b>Stage 1 — Acquire</b> certificate; the signed classroom Responsible-AI charter is core Stage 1 portfolio evidence.

## What teachers will be able to do

By the end of Module 1.3, a teacher can:

- Explain **why AI is human-led** and what human agency means when choosing and using AI tools.
- State the **core ethical principles** for classroom AI use, and recognise them in real cases.
- Identify **privacy, fairness and safety risks** in a given AI use, and propose simple mitigations.
- Articulate the rule that **AI may inform a decision but a person remains accountable for it**.
- Recognise **bias, stereotypes and cultural or caste-based mismatch** in AI outputs, and respond well.
- Draft and sign a classroom **Responsible-AI charter**, and extend it into a **student AI-use agreement**.

- Guard against **over-reliance** on AI — keeping pupils' own thinking central — and **verify** AI claims with a simple lateral-reading routine.

## How this module fits the programme

This is the third module of Stage 1 (Acquire → Deepen → Create), and the one that sets the **values and ethics foundation** for everything that follows. Module 1.1 made the technology dependable; Module 1.2 made the *idea* of AI honest; this module makes its *use* principled — so that when teachers begin using AI in real lessons (Module 1.4 onward), they do so safely, fairly and accountably.

It **formalizes the two rules** the earlier modules introduced — **AI helps but never certifies (a person decides)** and the **personal-data red line** — and turns them into a signed classroom charter the teacher actually adopts. The teaching draws on the **UNESCO Recommendation on the Ethics of AI (2022)**, the **UNESCO AI Competency Framework for Teachers (2024)**, and the MIT RAISE **Day of AI** lessons — ‘Deep Fakes’, ‘Truth, Tricks and AI’ (lateral reading), and ‘AI, Effort and Your Brain’ (over-reliance) — and it carries the **CRAFT water-safety example** as its anchor case. It also seeds the ethics strand of the student AI-literacy curriculum that the teacher will later deliver.

## How to use this package

**For the Moodle builder:** Part 1 is your build map — create one course Section per row group and add each item using the Moodle type named. Part 2 supplies the text to paste into each Page or Book. Part 4 specifies the eight H5P activities, the **AI Ethics Debate**, and the ‘Ethical dilemmas — who decides?’ SCORM routine. Part 5 supplies the quiz items, the rubric and the forms.

**For the facilitator:** Part 3 is your slide outline, Part 6 closes with the source mapping, and Part 5 gives you the observation checklist and confidence self-assessment to evaluate participating teachers. This module is discussion-led and fully unplugged-capable; the *Time* column in Part 1 and the slide-by-slide notes in Part 3 carry the pacing.

## Four threads that run through every topic

<b>Inclusion, gender &amp; disability</b>	<b>Bias and exclusion are the explicit content here</b> — gender, caste, language, disability and geography — taught with locally grounded, respectful cases. The module is fully unplugged-capable, so no teacher is excluded by hardware or connectivity.
<b>Multilingual &amp; cultural relevance</b>	Delivered in Nepali and English; scenarios are drawn from Nepali classrooms, not imported. The charter and the student agreement are produced <b>bilingually</b> , and Nepali-language quality is checked before use.
<b>Environmental cost of AI</b>	Named here as part of <b>do-no-harm and responsible use</b> : large AI models carry a real energy and water cost. Teachers meet this as an ethical consideration now and develop it fully in later stages.
<b>Validated tools &amp; verification</b>	This is the <b>ethics-by-design</b> thread: tools are checked before classroom use, and verification is modelled rather than assumed — teachers learn a concrete <b>lateral-reading</b> routine and guard against <b>over-reliance</b> , so pupils' own thinking stays central. The validation burden sits with institutions, not individual teachers.

## Part 1 — Moodle course structure (the upload map)

Build the course in nine sections. For each item below, add a Moodle activity or resource of the **type** named. The **Assessed?** column shows what feeds a grade: *Auto* (H5P / quiz), *Rubric* (a practical task), *Participation* (a forum / submission), or *Self*. The *Time* column doubles as a pacing guide for the facilitator.

### Section 0 — Start here (course header)

Item	Moodle type	Time	Purpose	Assessed?
Welcome to Module 1.3	Label / Page	—	Course banner: welcome, what to expect, time commitment.	—
How this module works (online / offline / shared device)	Page	10 min	Explains the three routes; notes it is fully unplugged-capable.	—
Module 1.3 syllabus (1 page)	File (PDF)	5 min	Downloadable at-a-glance reference for the cohort.	—
Warm-up: does AI pose a threat?	H5P / Feedback	10 min	Three ways AI can do good, three ways it can do harm.	Self
Meet your cohort	Forum	10 min	Introductions and one ethical question each has about AI.	Participation

### Section 1 — AI is human-led

Item	Moodle type	Time	Purpose	Assessed?
1 · AI is human-led — agency, hype and the teacher's role	Page / Book	15 min	Core reading (paste from Part 2).	—
Hype or honest?	H5P — Question Set	8 min	Sort common AI claims into myth and reality.	Auto
Cognitive effort: thinking or outsourcing?	H5P — Drag and Drop	8 min	Sort study scenarios by who is doing the thinking — you or the AI.	Auto
One task I'd hand to AI — and one I never would	Forum	8 min	Surfaces agency through real teaching choices.	Participation

### Section 2 — The core ethical principles

Item	Moodle type	Time	Purpose	Assessed?
2 · The core ethical principles for classroom AI	Page / Book	16 min	Core reading (paste from Part 2).	—
Match the principle to the case	H5P — Drag and Drop	10 min	Connect classroom cases to do-no-harm · fairness · privacy · transparency · human determination.	Auto
The AI Ethics Debate (stakeholder role-play)	Offline + Forum / Files	35 min	Four stakeholder groups weigh AI's risks and benefits; perspective-taking + critical-user reflection.	Participation
Ethics vocabulary	H5P — Dialog Cards	8 min	Plain-language terms with a Nepali gloss.	Auto

### Section 3 — The anchor rule: a person decides

Item	Moodle type	Time	Purpose	Assessed?
3 · The anchor rule — AI helps but never certifies	Page / Book	14 min	Core reading (paste from Part 2); the CRAFT water-safety case.	—
Who owns the decision?	H5P — Question Set	10 min	For each case, who is accountable — the teacher, the tool, or the pupil?	Auto
The anchor rule	Page (teal callout)	5 min	'AI helps, but it never certifies. A person decides.'	—

### Section 4 — Privacy & the personal-data red line

Item	Moodle type	Time	Purpose	Assessed?
4 · Privacy and personal data — what never goes into a tool	Page / Book	15 min	Core reading (paste from Part 2).	—
The personal-data red line	Page (amber callout)	5 min	The one rule we never break — with extra care for data about children.	—
Redact it: spot the personal data	H5P — Question Set	10 min	Identify what must be removed before a task touches a tool.	Auto

### Section 5 — Bias & exclusion

Item	Moodle type	Time	Purpose	Assessed?
5 · Bias and exclusion — gender, caste, language, disability, geography	Page / Book	16 min	Core reading (paste from Part 2).	—
Fair or biased?	H5P — Question Set	10 min	Judge real outputs; name the affected group; propose a fairer fix.	Auto
Bias we've seen — and how we'd respond	Forum	8 min	Pools locally relevant, respectful cases.	Participation

### Section 6 — Misinformation, fakes & deepfakes

Item	Moodle type	Time	Purpose	Assessed?
6 · Misinformation, fake content and deepfakes	Page / Book	15 min	Core reading (paste from Part 2).	—
Real or fake? verify before you share	H5P — Question Set	10 min	Identify deepfakes; tell misinformation from disinformation; choose the right response.	Auto
Lateral reading: spot the red flags	H5P — Question Set	12 min	Check an AI claim against reliable sources, detail by detail.	Auto
Source evaluation checklist	File / Page	6 min	Seven questions for judging whether a source can be trusted.	—
Demo: AI-generated faces & how to check	URL / File	6 min	A short, safe demonstration that fake can look real.	—

## Section 7 — Your Responsible-AI charter & the dilemma routine

Item	Moodle type	Time	Purpose	Assessed?
7 · From principles to practice — your Responsible-AI charter	Page / Book	12 min	Core reading (paste from Part 2).	—
Ethical dilemmas — who decides?	SCORM package	30 min	The module centrepiece; produces a 'dilemma decisions' summary.	Completion + score
Practical Task 1: your classroom Responsible-AI charter	Assignment (file upload)	30 min	A signed one-page charter for the teacher's own class.	Rubric

## Section 8 — The student AI-use agreement, review & assessment

Item	Moodle type	Time	Purpose	Assessed?
8 · The student AI-use agreement — and what's next	Page / Book	12 min	Core reading (paste from Part 2).	—
Can it? Should it? (AI roles by assignment)	H5P — Question Set	10 min	For each step of an assignment, decide if AI can — and should — help.	Auto
Practical Task 2: your student AI-use agreement	Assignment (file upload)	25 min	Allowed · guided · restricted uses for pupils, built with 'can it? should it?'. Allowed · guided · restricted uses for pupils, built with 'can it? should it?'.	Rubric
Module 1.3 knowledge check	Quiz	20 min	Graded check across the whole module (14 items).	Graded (≥70%)
Confidence self-assessment (post) & reflection	H5P / Feedback + Forum	10 min	Growth check; completes the portfolio seed.	Self
Module complete — next steps & certificate path	Page / Label	—	Closure; clear pointer to Module 1.4.	—

## Part 2 — Core content sections (paste-ready)

Full teaching text for all eight topics, written to paste straight into a Moodle **Page** or **Book**. Each topic carries a hands-on **Try it** box, and the module's two rules — the **anchor rule** (a person decides) and the **personal-data red line** — are formalized where they belong. A Nepali-language version sits on the facing Page (bilingual delivery).

### 1 · AI is human-led — agency, hype and the teacher's role

*Bilingual note for the Moodle builder: paste this English text into the Page/Book, and place the Nepali translation on the facing Page (or as the second language block). Keep the headings, the Try it box and the rule callouts in both languages.*

Used well, AI can help a teacher do more of what matters. But it only helps if a person stays firmly in charge of it. This module begins where responsible practice begins: with the idea that **AI in education is human-led**, and the teacher — not the tool — is the professional in the room.

#### Cutting through the hype

AI is surrounded by big claims. Three of them are worth retiring on day one. **'AI will replace teachers'** — it will not; it has no relationship with a child, no judgement about a class, and no accountability. **'AI is magic'** — it is not; as Module 1.2 showed, it predicts patterns from data, which is powerful but ordinary. **'AI is always right'** — it is not; it can be biased and it can be confidently wrong. Seeing AI clearly is the first step to using it ethically.


#### What human agency means

**Human agency** is the teacher's power to make deliberate choices about technology: *whether* to use a tool, *when, how, and who it affects*. Agency is not refusing AI, and it is not accepting it uncritically — it is staying the decision-maker. A teacher with agency asks, before reaching for a tool: does this help my pupils? Is it fair? Is it safe? Could I justify this choice to a parent? The rest of this module gives teachers the principles to answer those questions well.

#### Agency also means protecting thinking

Agency is not only about the teacher's choices — it is also about protecting pupils' thinking. The newest risk of AI in education is **over-reliance**: leaning on a tool so heavily that you stop thinking for yourself. A hard task demands real **cognitive effort** — the productive struggle where learning actually happens — and when that effort is always handed to AI, pupils build up **cognitive debt**: a slow loss of the very skills the work was meant to develop. A pupil who lets AI write every answer may finish the task but learn nothing from it.

The practical tool is one extra question, asked of every task: not just *can AI help?* but **should it?** AI can brainstorm, explain a concept a pupil has already wrestled with, check grammar after they have written, or suggest an approach — but it cannot replace their original ideas, their voice, the struggle through a hard problem, or their responsibility to learn. Teaching pupils to ask **'can it? should it?'** — and to do their own thinking first — is human agency in action, and it returns in the student agreement (Topic 8).

 **Try it — keep your agency** Name one classroom task you might genuinely hand to an AI tool, and one you would never hand over. For each, ask *can it?* and *should it?*, and say *why* in a single sentence. The line between them is your professional judgement — the thing this module sharpens.

## 2 · The core ethical principles for classroom AI


*Bilingual note for the Moodle builder: paste this English text into the Page/Book, and place the Nepali translation on the facing Page (or as the second language block). Keep the headings, the Try it box and the rule callouts in both languages.*

Ethics can sound abstract, but for a classroom it comes down to a short list of principles that each translate into a plain question a teacher can ask. These are drawn from the **UNESCO Recommendation on the Ethics of AI (2022)** and the **UNESCO AI Competency Framework for Teachers (2024)**, kept simple and practical.

### Six principles, and the question each one asks

- **Do no harm (safety).** Use AI only where it helps, and avoid foreseeable harm. *Ask: could this hurt a pupil, or could a simpler approach do the job?*
- **Fairness and non-discrimination.** AI must not disadvantage learners by gender, caste, language, disability or geography. *Ask: is this fair to every pupil in my class?*
- **Privacy and data protection.** Pupils' personal information is protected. *Ask: am I keeping pupils' data safe — and out of public tools?*
- **Transparency.** Be open about when and how AI is used. *Ask: would I be comfortable telling pupils and parents I used this?*
- **Human determination (oversight).** A person makes and owns the decision. *Ask: am I deciding, or is the tool deciding for me?*
- **Accountability.** The teacher remains answerable for the outcome. *Ask: if this goes wrong, who is responsible? (It is me.)*

These six are not a test to pass once; they are a set of habits. The rest of the module takes the three that matter most in daily practice — **human determination, privacy, and fairness** — and works them through in detail.

 **Try it — name the principle** Think of one time you (or a colleague) used, or considered using, a digital or AI tool with a class. Which of the six principles was most at stake? What question would you ask next time before pressing 'go'?

## 3 · The anchor rule — AI helps but never certifies

*Bilingual note for the Moodle builder: paste this English text into the Page/Book, and place the Nepali translation on the facing Page (or as the second language block). Keep the headings, the Try it box and the rule callouts in both languages.*

Of all the principles, one governs the whole programme, because it decides who is in charge: **AI may inform a decision, but a person checks, judges and remains accountable for it.** AI can draft, summarise, suggest, translate and help read evidence — but it does not understand, it is not always right, and it can never be held responsible. So it must never be the one that *decides*.

### The rule, in the words pupils remember


Said most sharply: **'AI helps, but it never certifies. A person decides.'** A teacher may use AI to *draft* a report comment, but the teacher decides the grade. AI might help *organise* evidence, but a person makes the judgement. The tool informs; the human determines. This is not caution for its own sake — it follows directly from what AI is: a pattern-predictor that can be biased (Topic 5) and can hallucinate (Module 1.2), and therefore cannot carry responsibility.

## The CRAFT water-safety example

The programme's flagship pilot makes the rule concrete. In the CRAFT water unit, pupils test water and may use sensors or an app to help read the evidence. But **no AI, sensor or app ever declares water 'safe to drink'** — a person interprets the evidence, and the activity uses pre-selected, non-drinking water sources with clear disclaimers. The technology *informs* the inquiry; it never *certifies* safety. That is the anchor rule in its most important form: where a wrong 'yes' could harm a child, the human decides, every time.

### The anchor rule — the rule that governs the whole programme

**AI helps, but it never certifies. A person decides.** AI may inform a decision — by drafting, summarising or helping read evidence — but a trained person checks it, judges it, and remains accountable for the outcome. Teach pupils the rule in your own words, and model it every time you use a tool: the tool suggests, **you** decide.

 **Try it — rewrite the sentence** Take a sentence like '*the app says the water is safe*' or '*the AI marked this essay a B*'. Rewrite it so a **person** is clearly the one who decides and is accountable. That rewrite is the habit this whole programme is built on.

## 4 · Privacy and personal data — what never goes into a tool

*Bilingual note for the Moodle builder: paste this English text into the Page/Book, and place the Nepali translation on the facing Page (or as the second language block). Keep the headings, the Try it box and the rule callouts in both languages.*

Using AI responsibly with pupils means protecting their personal information carefully. Module 1.1 introduced this as a habit; here it becomes a firm, formalized rule — because generative AI tools are often public, and what you type into them may not stay private.

### What counts as personal data

**Personal data** is any information that identifies a real person: a pupil's full name, photo, address, phone number, marks, health information, and — handled with particular care in Nepal — caste, ethnicity or disability status. The more sensitive the information, and the younger the person, the more protection it deserves.

### The rule, and how to keep it


- **Never type pupils' personal data into a public AI or online tool.** No full names, photos, addresses, marks, health or identity information.
- **Use initials or codes** instead of names; describe a pupil's need without identifying them.
- **Use the school's approved tools** for anything sensitive, and follow the school's data and child-protection rules.
- **Take extra care with data about children** — they cannot consent to how their information is used, so the duty sits with you.

### The personal-data red line — the one rule we never break

Pupils' personal data is **never** typed into a public AI or online tool. A trained person — not the tool — makes the decision; AI may **assist**, but it never decides about a child.

Use initials or codes, use approved tools for anything sensitive, and take extra care with data about children.

**When in doubt, don't** — ask your school's focal point first.

 **Try it — redact before you reach for a tool** Take a real-looking task — a class list, a report comment, a certificate. Cross out everything that identifies a pupil, and rewrite it with initials or codes so it could safely be shown to anyone. This ‘redact-it’ check is the privacy habit, made routine.

## 5 · Bias and exclusion — gender, caste, language, disability, geography

*Bilingual note for the Moodle builder: paste this English text into the Page/Book, and place the Nepali translation on the facing Page (or as the second language block). Keep the headings, the Try it box and the rule callouts in both languages.*

Module 1.2 showed *why* AI is biased — it learns from data, and narrow data makes a narrow machine. This topic asks the question that matters for a teacher: **who gets left out, and what do I do about it?** Fairness is not optional; it is one of the core principles, and in a Nepali classroom it is concrete.

### How AI can exclude Nepali learners

Because most AI is trained on data from elsewhere, its ‘defaults’ can quietly disadvantage particular learners. The dimensions to watch in Nepal are clear:


- **Language.** A tool that assumes English, or handles Nepali and mother-tongue languages poorly, disadvantages pupils who think and write in those languages.
- **Gender.** Outputs that show only men as scientists or leaders, or that stereotype girls, narrow what pupils believe is possible.
- **Caste and ethnicity.** Content that erases, stereotypes or misrepresents particular communities is unfair and harmful; treat this dimension with care and respect.
- **Disability.** Tools that are not accessible — no captions, no screen-reader support, no large text — exclude learners with disabilities by design.
- **Geography.** Examples and assumptions drawn from cities can leave rural learners unseen and unserved.

### Where the bias comes from

Bias does not appear from nowhere — an AI model absorbs it from the human-made sources it learns from. The same patterns are visible all over the everyday internet: opinion pieces that frame an issue one way, online forums where a single viewpoint dominates, reviews left by only some kinds of people, and image or video choices that quietly shape opinion. An **echo chamber** — a space where you mostly meet views like your own — deepens it. When a model learns from skewed sources, it repeats and amplifies the skew. Helping pupils notice bias in ordinary sources is the same skill as noticing it in AI.

### How a teacher responds

Noticing bias is the first act of fairness. When you see it: **name it** (to yourself and, where useful, to pupils); **don't propagate it** — don't hand pupils a stereotyped or exclusionary output as if it were neutral; **choose or adapt fairer materials**; and **teach pupils to spot it too**, so they become critical rather than passive users. Turning bias into a teachable moment is one of the most valuable things this module offers.

 **Try it — spot it, then fix it** Find one example of bias or exclusion in an AI output or tool you've seen — in language, gender, caste, disability or geography. Write one sentence naming who it disadvantages, and one sentence on how you'd make it fairer for your class.

## 6 · Misinformation, fake content and deepfakes

*Bilingual note for the Moodle builder: paste this English text into the Page/Book, and place the Nepali translation on the facing Page (or as the second language block). Keep the headings, the Try it box and the rule callouts in both languages.*

Generative AI can now create images, audio and video that look completely real but are not. For teachers and pupils, this raises a clear ethical and practical challenge: how to respond to a world where seeing is no longer believing. This topic draws on the MIT Day of AI ‘Deep Fakes’ lesson and connects directly to the ‘verify before you trust’ habit from Module 1.2.

### Deepfakes, and two kinds of false information

A **deepfake** is AI-generated media — an image, a voice, a video — that mimics a real person, making it appear they said or did something they did not. The technology is improving quickly, and convincing fakes are easy to make and share. Two related terms matter:

- **Misinformation** — false or misleading information shared *without* the intent to deceive, often by mistake or misunderstanding.
- **Disinformation** — false information spread *deliberately* to mislead, manipulate or harm a person, group or organisation.

The harm is real: deepfakes and disinformation can damage individuals, manipulate opinion, and erode the basic trust a society and a classroom depend on. Pupils are exposed to this content daily.

### How teachers and pupils respond

The response is a habit, not a technology: **pause and verify before you share or believe**. Check the source; ask whether a reputable outlet reports the same thing; be especially sceptical of content that is shocking, perfect, or designed to make you angry. Teach pupils the same reflex — that a convincing image is not proof — and treat ‘is this real?’ as a normal, healthy question. This is the ‘verify before you trust’ rule from Module 1.2, now applied to fake media.


### Lateral reading — a verification routine you can teach

‘Verify before you trust’ becomes teachable when it has a method. **Lateral reading** is the technique professional fact-checkers use: instead of staying on one page trying to judge it, you open new tabs and check its specific claims against two or three independent, reputable sources. The steps are simple enough for any classroom — **spot the surprising claims** in an AI answer; **open new sources** (prefer official, education or news sites); **check each detail** against them, one by one; then **decide** what is supported, what is wrong, and what is missing. AI is confidently wrong often enough that this routine catches real errors, and it turns pupils from passive readers into active checkers.

#### A check you can hand to pupils — can this source be trusted?

Before trusting any AI answer or web source, ask: **Who** made it, and are they qualified? **Why** does it exist — to inform, persuade or sell? **When** was it made — is it current? **Where** is it published — a trustworthy domain? **What** evidence backs the claims, and can you verify it elsewhere? **Is it biased** — one-sided, stereotyped or emotionally charged?

And one extra question for the AI age: **if a model learned from this source, what might it pick up — including its bias?** Verification and fairness turn out to be the same habit, seen from two sides.

 **Try it — read laterally** Take one surprising claim from an AI tool — about Nepali geography, history or science — and read laterally: open two or three reliable sources and check each detail. Note one thing the AI got right and one it got wrong or exaggerated. That side-by-side check is the routine to teach your pupils.

## 7 · From principles to practice — your Responsible-AI charter

*Bilingual note for the Moodle builder: paste this English text into the Page/Book, and place the Nepali translation on the facing Page (or as the second language block). Keep the headings, the Try it box and the rule callouts in both languages.*

Principles only matter when they shape what happens in a real classroom. This module's first practical task turns everything above into a short, signed document the teacher actually adopts: a **classroom Responsible-AI charter**. It is concrete, owned, and ready to enact during the practicum.

### What the charter contains


A good charter fits on one page and is written in the teacher's own words. It should state, plainly: the **anchor rule** (AI helps but never certifies; a person decides); the **personal-data red line**; how the **core principles** apply in this particular class; and a simple summary of **allowed, guided and restricted** AI uses. It is a promise to pupils and parents about how AI will — and will not — be used.

The centrepiece activity feeds straight into it. In the **'Ethical dilemmas — who decides?'** routine (Part 4), teachers work through realistic classroom situations — an AI grade, a privacy request, a biased output, a possible deepfake, the water-safety case — deciding who is accountable each time. The decisions they reach become the clauses of their charter.

### Practical Task 1 — your classroom Responsible-AI charter

Draft and **sign a one-page Responsible-AI charter** for your own classroom. It must include the **never-certify rule**, the **personal-data red line**, the core principles applied to your class, and a clear allowed / guided / restricted summary.

Use your 'dilemma decisions' summary from the centrepiece as your starting point. The signed charter enters your Stage 1 portfolio and is enacted during the practicum.

 **Try it — draft three clauses** Write three clauses for your charter now: one that names the anchor rule, one that states the personal-data red line, and one that addresses fairness for your particular pupils. Three clauses is a charter begun.

## 8 · The student AI-use agreement — and what's next

*Bilingual note for the Moodle builder: paste this English text into the Page/Book, and place the Nepali translation on the facing Page (or as the second language block). Keep the headings, the Try it box and the rule callouts in both languages.*

The charter sets out how *the teacher* will use AI. The second practical task turns the same principles outward, into a clear, fair agreement for *pupils*: a **student AI-use agreement** that tells them, in plain language, how they may use AI for their own work.

### Allowed, guided, restricted

The clearest agreements sort AI use into three simple categories, set out for classwork, homework and projects:

- **Allowed** — uses pupils may make freely (for example, asking an AI tool to explain a difficult idea, or to suggest practice questions).
- **Guided** — uses permitted *with* a teacher's direction or disclosure (for example, using AI to help draft, then showing what was AI-assisted).
- **Restricted** — uses not permitted (for example, submitting AI-written work as their own, or entering anyone's personal data into a tool).

The agreement works best when it is **fair, plainly written, and co-created with pupils** where possible — so it builds honest habits rather than only forbidding things. It respects the same principles as the charter: fairness, privacy, transparency and the anchor rule. (In Stage 2, this grows into a fuller academic-integrity agreement; here it is a clear, firm first version.)

### From 'can it?' to 'should it?' — reasoning it out with pupils

The strongest agreements are not handed down; they are reasoned out. A simple, powerful way to build one is to take a real assignment — a paragraph of writing, a history summary, a science-fair poster, a poem, a multi-step maths problem — and break it into steps. For each step the class asks two questions: **can AI help?** (almost always yes) and, more importantly, **should it?** Drafting an outline, checking grammar after writing, or generating practice questions may be fine; writing the whole draft, or doing the maths pupils are meant to be learning, is not. Sorting each step into **allowed, guided or restricted** — with pupils giving the reason for each — produces an agreement they understand and own, and quietly teaches the over-reliance lesson from Topic 1.

#### Practical Task 2 — your student AI-use agreement

Extend your charter into a one-page **student AI-use agreement** that sets out **allowed, guided** and **restricted** AI uses for classwork, homework and projects, in language your pupils can understand.

Build it with the '**can it? should it?**' method — take one real assignment and reason through it step by step — and, where you can, shape it together with your class. Keep it fair and transparent, protect pupils' personal data, and add it to your Stage 1 portfolio alongside the charter.

### What's next

With dependable habits (Module 1.1), an honest understanding of AI (Module 1.2), and now a principled, accountable mindset (Module 1.3), you are ready to put AI to work in real lessons. **Module 1.4 — Getting Started with AI-Assisted Teaching and Reflective Practice** is where the charter you wrote here governs your first AI-assisted lesson, and where you learn to decide whether a task even needs AI at all.

#### You're ready for Module 1.4

Completing Module 1.3 contributes to your **Stage 1 — Acquire** certificate, and your signed Responsible-AI charter is core portfolio evidence. Next is **Module 1.4 — Getting Started with AI-Assisted Teaching and Reflective Practice**, where this mindset becomes your first cycle of real classroom practice.

## Part 3 — Visual presentation outline (facilitator deck)

A 26-slide deck for the contact sessions. This module is discussion-led — keep each slide light and let the dilemmas, the debate, the verification practice and the drafting carry the learning. The deck follows the eight topics of Part 2; slide numbers double as a pacing guide for the facilitator.

Slide & title	Key points	Suggested visual	Speaker note
<b>S1 · Module 1.3 — Mindset &amp; Ethics</b>	Programme title slide Stage 1 — Acquire · Strands A1 & A2 Facilitator & cohort name	Himalayan Water title; CRAFT strip.	Set the tone: this is the values floor the whole programme stands on.
<b>S2 · Why this module</b>	AI in education is human-led Principled, accountable use The teacher is the professional in the room	Human hand over a tool, not under it.	Frame ethics as practical judgement, not abstract rules.
<b>S3 · What you'll be able to do</b>	The seven module outcomes Agency · principles · the rule · privacy · bias · verify · your charter	Seven simple icons, one per outcome.	Read these as promises; the charter is where they land.
<b>S4 · How today works</b>	Fully unplugged-capable Online / offline / shared device Discussion-led	Three-route diagram.	No hardware needed; this module runs on conversation.
<b>S5 · Warm-up — does AI pose a threat?</b>	Three ways AI can do good Three ways it can do harm Pair, then share	Balance scale: good vs harm.	Surface real hopes and fears before introducing principles.
<b>S6 · AI is human-led</b>	Retire three myths: replaces teachers · is magic · is always right Agency = staying the decision-maker	Myth-busting graphic.	Name the hype openly; clear sight is the first ethical act.
<b>S7 · Don't outsource your thinking</b>	Over-reliance & cognitive debt The struggle is where learning happens Ask: can AI help? should it?	Brain vs crutch graphic.	Agency includes protecting pupils' own thinking — set up Topic 8.
<b>S8 · The core ethical principles</b>	Do-no-harm · fairness · privacy · transparency · human determination · accountability Each is a question you ask	Six-principle ring.	Keep each principle attached to a plain classroom question.
<b>S9 · Activity — match the principle</b>	Connect real cases to the principle at stake Discuss the close calls	Card-sort mat.	Most cases touch several principles — discuss the strongest.
<b>S10 · The AI Ethics Debate</b>	Four stakeholder groups: tech · government · educators · users Prepare · debate · become a critical user	Four role cards in a ring.	The signature group activity — collective judgement, fully unplugged.
<b>S11 · The anchor rule</b>	AI helps but never certifies — a person decides CRAFT water-safety: AI never declares water 'safe'	Teal 'a person decides' callout.	This is the rule that governs the whole programme — say it twice.

Slide & title	Key points	Suggested visual	Speaker note
<b>S12 · Activity — who owns the decision?</b>	For each case, name who is accountable The teacher, the tool, or the pupil?	Decision-owner graphic.	Drill the reflex: the tool suggests, a person decides.
<b>S13 · Privacy &amp; the red line</b>	Pupils' personal data never goes into a public tool Extra care for data about children	Amber 'red line' callout.	The one rule we never break — be firm and concrete.
<b>S14 · Activity — redact it</b>	Cross out everything that identifies a pupil Rewrite with initials or codes	Before/after redacted task.	Make the privacy check a routine, not an afterthought.
<b>S15 · Bias &amp; exclusion</b>	Who gets left out? Language · gender · caste · disability · geography Where bias comes from: skewed sources, echo chambers	Inclusion / exclusion graphic.	Handle caste and identity with care; keep cases local and respectful.
<b>S16 · Activity — fair or biased?</b>	Judge real outputs Name who is disadvantaged Propose a fairer fix	Output-judgement cards.	Turn bias into a teachable moment pupils can spot too.
<b>S17 · Misinformation, fakes &amp; deepfakes</b>	A deepfake mimics a real person Misinformation vs disinformation Seeing is no longer believing	Real vs deepfake face pair.	Connect to 'verify before you trust' from Module 1.2.
<b>S18 · Lateral reading — verify, don't trust</b>	Leave the page; check claims elsewhere Two or three reliable sources, detail by detail The source-evaluation check	New-tabs / open-sources graphic.	Give the verify reflex a method pupils can actually use.
<b>S19 · Activity — real or fake?</b>	Spot the deepfake Choose the right response: pause & verify Check the source	Verify-checklist graphic.	Teach the reflex: a convincing image is not proof.
<b>S20 · Ethical dilemmas — who decides?</b>	The centrepiece routine Work real cases: grade · privacy · bias · deepfake · water-safety Your decisions become charter clauses	Branching dilemma path.	Let teachers feel the weight of deciding — and keep the human accountable.
<b>S21 · Your Responsible-AI charter</b>	One page, signed, in your words The anchor rule · the red line · principles · allowed/guided/restricted Practical Task 1	Charter template card.	Use the dilemma summary as the starting point.
<b>S22 · The student AI-use agreement</b>	Allowed · guided · restricted Built with 'can it? should it?', co-created with pupils Practical Task 2	Three-tier agreement graphic.	Build honest habits, not only prohibitions.
<b>S23 · Two rules to carry forward</b>	A person decides — AI never certifies Pupils' personal data never goes into a public tool	Teal + amber 'two rules' callout.	These two rules govern everything from here on.

Slide & title	Key points	Suggested visual	Speaker note
<b>S24 · Reflection &amp; confidence (post)</b>	Re-rate the seven outcomes One principle I'll act on Monday One thing I'll keep practising	Before/after confidence bars.	Celebrate growth from the morning's warm-up.
<b>S25 · Module 1.3 knowledge check</b>	14 short questions Open book; ≥70% to pass May be set as take-home	Quiz icon.	Reassure: it checks the module, not personal worth.
<b>S26 · You did it</b>	Contributes to the Stage 1 — Acquire certificate Your charter is core portfolio evidence Next: Module 1.4	Certificate ribbon + next-module pointer.	End on momentum and a clear next step.

## Part 4 — Interactive digital elements (H5P & SCORM)

Eight H5P activities, a structured **AI Ethics Debate**, and one SCORM package. H5P items are auto-marked and embed directly in Moodle; the SCORM package is the module's centrepiece — an ethical-dilemma decision routine, complemented by the group debate. Each specification gives the content type, where it sits, its structure, and the exact content to author. The deepfake, verification, bias and over-reliance material is drawn from the MIT Day of AI lessons you provided ('Deep Fakes', 'Truth, Tricks and AI', 'The Brain Behind the Bot', and 'AI, Effort and Your Brain').

### H5P activities (eight)

#### H5P 1 · Hype or honest?

<b>Content type</b>	Question Set
<b>Placement</b>	Section 1
<b>Purpose</b>	Retire the common AI myths and establish that AI is human-led — the foundation of Topic 1.

#### STRUCTURE

- 6 short claims; for each, choose 'myth' or 'honest' with a one-line 'why'.
- Feedback on every option; randomised; re-takeable.

#### EXACT CONTENT TO AUTHOR

- 'AI will replace teachers.' → **Myth** (no relationship, judgement or accountability).
- 'AI predicts patterns from data.' → **Honest** (what it actually does).
- 'AI is always right.' → **Myth** (it can be biased and confidently wrong).
- 'A person should stay in charge of any AI decision.' → **Honest** (human agency).
- 'AI understands pupils the way a teacher does.' → **Myth** (it has no understanding).
- 'AI can help a teacher, if a person checks its work.' → **Honest**.
- Closing feedback: 'AI is a tool; the teacher is the professional in the room.'

#### H5P 2 · Match the principle to the case

<b>Content type</b>	Drag and Drop
<b>Placement</b>	Section 2
<b>Purpose</b>	Connect a real classroom case to the ethical principle most at stake, so the principles become recognisable in practice.

#### STRUCTURE

- Six principle zones; case cards are dragged onto the best-fit zone.
- More than one can fit — teachers pick the strongest; feedback explains why.
- Self-paced; re-takeable.

#### EXACT CONTENT TO AUTHOR

- Zones: **Do-no-harm · Fairness · Privacy · Transparency · Human determination · Accountability**.
- Card: 'A teacher types a pupil's full name and marks into a free tool.' → Privacy.

- Card: ‘An AI tool’s score is entered as the grade without a teacher checking.’ → Human determination.
- Card: ‘A reading tool works well in English but poorly in Nepali.’ → Fairness.
- Card: ‘Pupils aren’t told that their feedback was AI-drafted.’ → Transparency.
- Card: ‘When an AI-assisted decision goes wrong, the teacher takes responsibility.’ → Accountability.
- Card: ‘A tool could be used, but a simpler, safer method would do the job.’ → Do-no-harm.

### H5P 3 · Who owns the decision?

<b>Content type</b>	Question Set
<b>Placement</b>	Section 3
<b>Purpose</b>	Drill the anchor rule — that a person, not a tool, makes and owns the decision.

#### STRUCTURE

- 5 short scenarios; for each, choose who is accountable; feedback names the principle.
- Pass shown at 70%; re-takeable.

#### EXACT CONTENT TO AUTHOR

- ‘An AI tool suggests a grade for an essay.’ → The **teacher** decides and owns the grade.
- ‘An app reads a water sample and shows a result.’ → A **person** interprets it; the app never certifies ‘safe’.
- ‘A chatbot drafts a message to parents.’ → The **teacher** checks, edits and sends it.
- ‘A pupil uses AI to help plan an essay.’ → The **pupil** is responsible for the work they submit.
- ‘A tool flags a pupil as ‘at risk’.’ → A **person** investigates and decides; the flag only informs.
- Closing feedback: ‘AI helps, but it never certifies. A person decides — and is accountable.’

### H5P 4 · Redact it — spot the personal data

<b>Content type</b>	Question Set
<b>Placement</b>	Section 4
<b>Purpose</b>	Make the privacy red line a routine: identify what must be removed before a task ever touches a tool.

#### STRUCTURE

- Short ‘before’ texts (a class note, a certificate request); select the items that are personal data.
- Multiple-response and true/false; feedback on every option; re-takeable.

#### EXACT CONTENT TO AUTHOR

- Item: from a sample line, select the personal data → full name · photo · home address · marks · health note (all are personal data).
- Item (T/F): ‘Initials or a code can safely replace a pupil’s full name in a tool.’ → **True**.
- Item (T/F): ‘It is fine to upload a class photo to a free tool to make a poster.’ → **False**.
- Item: which is the safe way to ask AI for help with a struggling reader? → describe the need **without** naming the pupil.
- Item: caste, ethnicity or disability status is → personal data that needs **particular** care.

- Closing feedback: ‘When in doubt, don't — and ask your school's focal point.’

### H5P 5 · Fair, biased — or fake?

<b>Content type</b>	Question Set
<b>Placement</b>	Sections 5 & 6
<b>Purpose</b>	Spot the two failure modes that demand an ethical response — bias/exclusion, and deepfakes/false information.

#### STRUCTURE

- 6 items mixing short outputs and media descriptions; multiple-choice and true/false.
- Feedback on every option; pass at 70%; re-takeable.

#### EXACT CONTENT TO AUTHOR

- Item: a tool draws every ‘scientist’ as a foreign man → label **bias**; name who it disadvantages; action: choose/adapt fairer materials.
- Item: a reading app handles English but not Nepali well → **exclusion by language**.
- Item: a shocking video of a public figure that may be AI-made → a **deepfake**; action: pause and verify before sharing.
- Item (T/F): ‘Disinformation is spread deliberately to mislead; misinformation is shared by mistake.’ → **True**.
- Item: the safest first step with any surprising image or claim → **check the source / verify before sharing**.
- Closing feedback: ‘Fairness and verification are ethical habits, not technical settings.’

### H5P 6 · Ethics vocabulary

<b>Content type</b>	Dialog Cards
<b>Placement</b>	Sections 2 & 8
<b>Purpose</b>	Front-load the small vocabulary of ethics that recurs across the programme, with a Nepali gloss on each card.

#### STRUCTURE

- A small deck of flip cards (term on front; plain-language meaning + Nepali gloss on back).
- Self-paced; no score; used as a reference.
- 16–18 cards.

#### EXACT CONTENT TO AUTHOR

- Cards: human agency · accountability · do-no-harm · fairness · privacy · transparency · human determination · bias · echo chamber · deepfake · misinformation · disinformation · lateral reading · responsible use · cognitive effort · cognitive debt · over-reliance.
- Example back — *Human determination*: ‘a person makes and owns the decision; AI only informs it.’
- Example back — *Over-reliance*: ‘leaning on AI so heavily you stop thinking — and stop learning — for yourself.’

- Example back — *Lateral reading*: ‘checking a claim by opening other reliable sources, not just trusting the page in front of you.’
- Closing card: ‘AI helps, but a person decides — and pupils’ personal data stays protected.’

### H5P 7 · Cognitive effort — thinking or outsourcing?

<b>Content type</b>	Drag and Drop
<b>Placement</b>	Section 1
<b>Purpose</b>	Make the over-reliance idea concrete: sort study moments by who is doing the thinking — the learner or the AI.

#### STRUCTURE

- Two zones — **You’re doing the thinking** and **AI is doing it for you**; scenario cards are dragged to the best fit.
- Feedback explains the cognitive-effort point on each; self-paced; re-takeable.

#### EXACT CONTENT TO AUTHOR

- Card: ‘You read an article, highlight key points and write your own summary.’ → **You’re thinking**.
- Card: ‘You struggle through a hard problem yourself and finally get it.’ → **You’re thinking**.
- Card: ‘You brainstorm and draft a story first, then ask AI to suggest edits.’ → **You’re thinking** (AI supports).
- Card: ‘You paste the question into a chatbot and copy its answer.’ → **AI is doing it** (cognitive debt).
- Card: ‘You ask AI to write the whole essay and submit it.’ → **AI is doing it**.
- Card: ‘You use AI to check grammar, but decide which suggestions to keep.’ → **You’re thinking** (you stay in charge).
- Closing feedback: ‘Use AI to support your thinking, not to replace it. Do your own thinking first.’

### H5P 8 · Can it? Should it?

<b>Content type</b>	Question Set
<b>Placement</b>	Section 8
<b>Purpose</b>	Drive the student agreement: for each step of an assignment, decide whether AI can help, and whether it should.

#### STRUCTURE

- 6 assignment steps; for each, judge **can AI help?** and **should it?**, then sort into allowed / guided / restricted.
- Feedback on every item; pass at 70%; re-takeable.

#### EXACT CONTENT TO AUTHOR

- Step: ‘Brainstorm a topic for an essay.’ → Can: yes · Should: usually → **allowed / guided**.
- Step: ‘Write the full first draft for you.’ → Can: yes · Should: no → **restricted**.
- Step: ‘Check grammar after you have written.’ → Can: yes · Should: yes → **allowed**.
- Step: ‘Do the multi-step maths you are learning.’ → Can: yes · Should: no → **restricted**.

- Step: ‘Explain a concept you have already tried to understand.’ → Can: yes · Should: yes → **allowed / guided**.
- Step: ‘Suggest sources for you to check yourself.’ → Can: yes · Should: yes, with verification → **guided**.
- Closing feedback: ‘Can is almost always yes. The real question is should — and who is doing the learning.’

## SCORM package — Ethical dilemmas: who decides?

The centrepiece activity. A teacher works through a series of realistic classroom dilemmas, deciding who is accountable each time and watching the consequence of each choice. The goal is to drill the ‘**a person decides**’ reflex until it is automatic — and the decisions a teacher reaches become the clauses of their Responsible-AI charter. Build it as a SCORM package (or, where SCORM authoring is not available, a Moodle Lesson plus the printable scenario cards achieves the same outcome). The dilemmas draw on the agreed programme and the Day of AI scenarios material.

Screen	On-screen content	Interaction	Logic / scoring
1 · Intro	What the routine is and how it runs (≈30 min, re-runnable). The aim: keep a person accountable for every AI-assisted decision.	Continue.	Records start; no score yet.
2 · The AI grade	An AI tool scores a pupil's essay. Do you enter it, use it as one input and decide yourself, or ignore it?	Choose an action → see the consequence.	Best path = teacher decides; illustrates <b>human determination</b> .
3 · The privacy request	A free tool needs pupils' full names and photos to make certificates. What do you do?	Choose → consequence.	Best path = don't enter personal data; illustrates the <b>red line</b> .
4 · The biased output	An image tool draws every ‘leader’ as a man, and a pupil notices. How do you respond?	Choose → consequence.	Best path = name it and adapt; illustrates <b>fairness</b> .
5 · The possible deepfake	A pupil shares a shocking video of a public figure that may be a deepfake. What is your first move?	Choose → consequence.	Best path = pause and verify; illustrates <b>do-no-harm</b> .
6 · The water-safety call	Your sensor / app says a water sample is ‘safe to drink’. Pupils want to act on it.	Choose → consequence.	Best path = a person interprets the evidence; AI never certifies ‘safe’ (CRAFT).
7 · Reflect	Which dilemma was hardest? Which principle was most at stake? What will you write into your charter?	Short free-text reflection.	Must be answered to continue; no pass/fail.
8 · Your dilemma decisions	Auto-built one-screen summary of the decisions you reached and the principle behind each — ready to copy into your charter.	Review + download / print.	Completion = all dilemmas worked + reflection written.

### Completion & scoring rules

**Completion** = all six dilemmas worked through **and** the reflection written.

**Score** = light-touch: the routine rewards keeping a person accountable and naming the principle, not a single 'correct' answer — the point is the reflex, not a grade.

**Re-runnable**, and designed to be re-used: the same dilemmas are ones a teacher can later pose to pupils, unplugged.

## Structured activity — the AI Ethics Debate

Where the SCORM routine drills an *individual* reflex, the **AI Ethics Debate** builds *collective* judgement — and it is the module's most engaging, fully unplugged activity. Teachers experience it as learners first, then learn to run it with their pupils. The class divides into four stakeholder groups, each given a role card; the groups prepare their position, debate AI's risks and benefits, then step back as **critical users**. It delivers the syllabus's perspective-taking and structured-discussion activities directly, and it needs nothing but the printed cards and an organizer sheet.

### THE FOUR STAKEHOLDER ROLES (one card per group)

- **Tech companies** — innovators who build AI tools. Priority: useful, profitable innovation. Tension: balancing progress against bias, misuse and calls for regulation.
- **Governments** — regulators protecting the public. Priority: safety and fairness for all citizens. Tension: how much oversight to impose without stifling progress.
- **Educators** — teachers and school leaders. Priority: AI that helps learning while protecting pupils' privacy and well-being. Tension: when AI helps, and when it replaces thinking or intrudes on pupils.
- **Everyday users** — pupils, parents and consumers. Priority: benefiting from AI while staying safe, fair and in control of their data. Tension: convenience versus privacy and the power to say no.

### HOW IT RUNS

- **Prepare (10 min)**: each group uses a one-page organizer — main priority, ethical concerns, benefits, and the arguments they will make from their stakeholder's view.
- **Debate (15 min)**: groups put their case and respond to others; everyone notes the strongest point from each of the other three groups.
- **Step back — become a critical user (10 min)**: each teacher picks a real AI use (a video-recommendation app, an AI homework grader, a school attendance tracker, pupils using chatbots for homework) and asks the critical-user questions — *Why is the AI giving this answer? Is it fair to everyone? How does it decide? How might it affect others?*
- **Reflect**: which perspective was hardest to argue, and what one takeaway about AI ethics will shape your charter?

### Build & delivery notes

Build in Moodle as a **Page** (instructions + critical-user questions) plus **Files** (the four role cards and the debate organizer) and a **Forum** for groups to post positions and reflections; offline, it runs entirely on printed cards.

Keep it respectful and local: choose AI uses pupils actually meet, and hold the personal-data and fairness lines throughout. Assessed by participation, the reflection feeds Practical Task 1 (the charter).

## Part 5 — Formative assessment tools

Assessment is light-touch and performance-based: auto-marked H5P checks (Part 4), a short graded quiz, a rubric for the two practical tasks, a facilitator observation checklist, and a confidence self-assessment that feeds the Stage 1 portfolio.

### Module 1.3 knowledge check (14 items)

Open-book; pass mark **70%**; re-takeable; may be set as take-home. In Moodle, build as a Quiz using the question type tagged on each item. Answers and feedback are for the builder — show feedback to teachers after submission.

**Q1.** ‘Human agency’ in classroom AI means... *(Multiple choice)*

- A. The teacher must use AI for every task
- B. The teacher deliberately chooses whether, when and how to use AI, and stays in charge
- C. Pupils decide which tools the class uses
- D. AI runs the lesson

**Answer: B**

*Feedback: Agency is staying the decision-maker — not refusing AI, and not accepting it uncritically.*

**Q2.** The programme's anchor rule is... *(Multiple choice)*

- A. AI decides; the teacher follows
- B. AI may inform a decision, but it never certifies — a person checks, judges and remains accountable
- C. Whatever the tool says is final
- D. AI replaces teacher judgement

**Answer: B**

*Feedback: AI helps, but it never certifies. A person decides — and is responsible for the outcome.*

**Q3.** A free tool asks for pupils' full names and photos to make certificates. You should... *(Multiple choice)*

- A. Enter them to save time
- B. Not enter them; use initials or codes, or an approved tool
- C. Ask pupils to enter their own details
- D. Enter only the photos

**Answer: B**

*Feedback: This crosses the personal-data red line. Personal data never goes into a public tool.*

**Q4.** Which of these are core ethical principles for classroom AI? (Select all that apply.) *(Multiple response)*

- A. Do no harm
- B. Always trust the tool
- C. Fairness and non-discrimination
- D. Privacy and data protection

**Answer: A, C and D**

*Feedback: ‘Always trust the tool’ is the opposite of the principles — verification and human oversight are required.*

**Q5.** True or false: an AI water-quality app may declare a sample ‘safe to drink’ and pupils can act on it directly. *(True / False)*

**Answer: False**

*Feedback: AI never certifies. A person interprets the evidence, sources are non-drinking, and disclaimers are shown (the CRAFT rule).*

**Q6.** A tool draws every ‘doctor’ as a man. The fair response is to... *(Multiple choice)*

- A. Use it as is — it's only a picture
- B. Name the bias, don't present it as neutral, and choose or adapt fairer materials
- C. Hide it from pupils
- D. Stop using technology altogether

**Answer: B**

*Feedback: Noticing bias is the first act of fairness; turning it into a teachable moment helps pupils spot it too.*

**Q7.** Misinformation and disinformation differ in that... *(Multiple choice)*

- A. They are the same thing
- B. Disinformation is spread deliberately to mislead; misinformation is false but shared without intent to deceive
- C. Misinformation is always about politics
- D. Disinformation is always harmless

**Answer: B**

*Feedback: Intent is the difference — but both are false, and both call for the verify reflex.*

**Q8.** A pupil shares a shocking video of a public figure. Your responsible first step is to... *(Multiple choice)*

- A. Share it widely so others can judge
- B. Pause and verify — check the source and consider it may be a deepfake — before sharing
- C. Assume it is true because it looks real
- D. Punish the pupil

**Answer: B**

*Feedback: A convincing image is not proof. Verify before you trust, and before you share.*

**Q9.** ‘Transparency’ as an ethical principle means... *(Multiple choice)*

- A. Keeping AI use secret
- B. Being open with pupils and parents about when and how AI is used
- C. Sharing pupils' data openly
- D. Using only free tools

**Answer: B**

*Feedback: Openness builds trust — and lets pupils and parents hold practice to account.*

**Q10.** True or false: if an AI tool makes a decision, the teacher is no longer accountable for it. *(True / False)*

**Answer: False**

*Feedback: Accountability stays with the person. The tool informs; the teacher remains answerable.*

**Q11.** The required artefact that codifies your classroom rules for AI is... *(Multiple choice)*

- A. An exam
- B. A signed Responsible-AI charter (extended into a student AI-use agreement)
- C. A timetable
- D. A textbook order

**Answer: B**

*Feedback: The charter makes the principles concrete and owned; the student agreement turns them toward pupils.*

**Q12.** ‘Lateral reading’ means... (Multiple choice)

- A. Reading the whole page top to bottom before deciding
- B. Opening other reliable sources and checking a claim's details against them
- C. Reading two books at once
- D. Trusting the first result

**Answer: B**

*Feedback: Lateral reading is how fact-checkers verify; leave the page and check the claim against independent, reputable sources.*

**Q13.** A pupil pastes a question into a chatbot and copies the answer without thinking. The main risk is... (Multiple choice)

- A. Nothing — the work is done
- B. Over-reliance and ‘cognitive debt’ — the pupil finishes the task but does not learn
- C. The AI gets tired
- D. It uses too much paper

**Answer: B**

*Feedback: Real learning happens in the effort. Asking ‘can it? should it?’ keeps pupils’ own thinking central.*

**Q14.** This module prepares you for... (Multiple choice)

- A. Leaving teaching to AI
- B. Module 1.4 — Getting Started with AI-Assisted Teaching — putting the mindset and rules into real lessons
- C. Nothing further
- D. An unrelated subject

**Answer: B**

*Feedback: The charter you wrote here governs your first AI-assisted lesson in Module 1.4.*

## Practical-task rubric

Applies to **Practical Task 1** (the signed classroom Responsible-AI charter) and **Practical Task 2** (the student AI-use agreement).

Criterion	Emerging	Developing	Secure
Principles applied	Principles are missing or misunderstood.	Most principles are present but thin or generic.	Do-no-harm, fairness, privacy, transparency and human determination are clearly applied to this class.
The anchor rule is explicit	The never-certify rule is absent or contradicted.	The rule is implied but not stated.	‘AI helps but never certifies; a person decides’ is stated plainly and owned.
Personal-data red line	The red line is missing or weak.	Stated, but vague on what is protected.	Firm and specific, with extra care for data about children.

Criterion	Emerging	Developing	Secure
Fairness & inclusion	Bias and exclusion are not addressed.	Mentioned in passing.	Addresses bias/exclusion across the named dimensions (gender, caste, language, disability, geography).
Clarity & usability for pupils	The student agreement is unclear or unfair.	Readable, but allowed/guided/restricted is fuzzy.	Plain, fair and usable; allowed, guided and restricted uses are clear and co-created where possible.

### Passing

Both artefacts should reach at least **Developing** across all criteria to pass the practical component. **Secure** on 'The anchor rule' and on 'Personal-data red line' is required, because these two rules are non-negotiable and protect pupils directly.

## Facilitator observation checklist

Use this **as you circulate** and during discussion, not only at the end. Mark Y (yes), P (partly) or N (not yet), and note who needs follow-up.

Observable behaviour	Y / P / N	Facilitator note
Explains why AI is human-led and what human agency means	Y / P / N	
States the core ethical principles and recognises them in a case	Y / P / N	
Applies the anchor rule (AI never certifies; a person decides) to a real situation	Y / P / N	
States and respects the personal-data red line, with care for child data	Y / P / N	
Identifies a privacy risk and redacts a task before it touches a tool	Y / P / N	
Recognises bias / exclusion across gender, caste, language, disability or geography	Y / P / N	
Distinguishes misinformation from disinformation and responds by verifying	Y / P / N	
Recognises a likely deepfake and pauses before sharing	Y / P / N	
Verifies an AI claim by reading laterally across reliable sources	Y / P / N	
Asks 'can it? should it?' and guards against over-reliance on AI	Y / P / N	
Drafts and signs a classroom Responsible-AI charter	Y / P / N	
Extends the charter into a fair student AI-use agreement	Y / P / N	

## Confidence self-assessment (before & after)

Teachers rate each statement at the start and again at the end, on a 1–4 scale: **1** Not yet · **2** Getting there · **3** Can do on my own · **4** Can help a colleague. This is for growth, not grading, and it feeds the Stage 1 portfolio.

I can...	Before	After
I can explain why AI is human-led and what human agency means.		
I can state the core ethical principles and recognise them in real classroom cases.		
I can apply the rule that AI helps but never certifies — a person decides.		
I can protect pupils' personal data and redact a task before it touches a tool.		
I can recognise bias, exclusion, misinformation and deepfakes, and respond well.		
I can verify an AI claim by reading laterally, and guard against over-reliance on AI.		
I can write a classroom Responsible-AI charter and a student AI-use agreement.		

**Reflection (write two lines):** one principle I will act on in class on Monday, and one thing I will keep practising.

## Part 6 — Source mapping & references

This module is built on the agreed AIDC programme, the UNESCO ethics frameworks, and the MIT Day of AI student materials you uploaded — ‘Deep Fakes’, ‘Preparing for an AI Ethics Debate’, ‘Truth, Tricks and AI’ (lateral reading), ‘The Brain Behind the Bot’ (exploring bias; AI roles by assignment) and ‘AI, Effort and Your Brain’ (over-reliance). The table shows what each part draws on.

### How the sources fed in

What it informs	Source
The core ethical principles (do-no-harm, fairness, privacy, transparency, human determination, accountability)	UNESCO <i>Recommendation on the Ethics of Artificial Intelligence</i> (2022); UNESCO AI CFT (2024)
The anchor rule (AI helps but never certifies; a person decides) and the CRAFT water-safety example	AIDC programme master syllabus; CRAFT pilot (ADB TA 10376-REG)
The personal-data red line and the ‘redact-it’ privacy walk-through	AIDC programme; UNESCO <i>Guidance for Generative AI in Education and Research</i> (2023)
Bias and exclusion across gender, caste, language, disability and geography	UNESCO AI CFT (2024); Nepal pilot design — GESI (ADB TA 10376-REG)
Misinformation, disinformation and deepfakes; how teachers and pupils respond	Day of AI — <i>How do Machines Create? ‘Deep Fakes’ lesson &amp; Scenarios Activity</i> (uploaded)
Lateral reading and the source-evaluation check (the verification routine)	Day of AI — <i>Truth, Tricks and AI: Learning to Verify Information</i> (uploaded)
The AI Ethics Debate — stakeholder roles, organizer and critical-user reflection	Day of AI — <i>Preparing for an AI Ethics Debate</i> (uploaded)
Over-reliance, cognitive effort / debt and the ‘can it? should it?’ lens	Day of AI — <i>AI, Effort and Your Brain; The Brain Behind the Bot — Lesson 4</i> (uploaded)
‘Where bias comes from’ and the echo-chamber idea	Day of AI — <i>The Brain Behind the Bot — Lesson 2, Exploring Bias on the Internet</i> (uploaded)
The Responsible-AI charter and the student AI-use agreement (the two required tasks)	AIDC programme master syllabus; UNESCO GenAI Guidance (2023)
Competency level, strands (A1, A2) and framework alignment	UNESCO AI CFT (2024) — Human agency · Ethical principles
Platform = Sikai Chautari (Moodle); SCORM / H5P formats; bilingual, low-bandwidth, fully unplugged design	AIDC programme & AI Teacher Training Delivery Specialist ToR

### References

1. UNESCO (2022). *Recommendation on the Ethics of Artificial Intelligence*. Paris, UNESCO.
2. UNESCO (2024). *AI Competency Framework for Teachers*. Paris, UNESCO.
3. UNESCO (2023). *Guidance for Generative AI in Education and Research*. Paris, UNESCO.
4. MIT RAISE — Day of AI. *How do Machines Create?* (ages 14+) — ‘Deep Fakes’ lesson, educator guide and Scenarios Activity. [dayofai.org](http://dayofai.org).
5. MIT RAISE — Day of AI. *Preparing for an AI Ethics Debate; Truth, Tricks and AI: Learning to Verify Information; The Brain Behind the Bot; AI, Effort and Your Brain* — student materials and slide decks. [dayofai.org](http://dayofai.org).

6. UNESCO (2024). *AI Competency Framework for Students*. Paris, UNESCO.
7. Government of Nepal, CEHRD / NCED — Teacher Professional Development framework and professional-conduct expectations; National Curriculum Framework (values & citizenship).

### About this package

The module's **identity, learning outcomes, structure and source mapping** are taken from the agreed AIDC programme (master syllabus), the UNESCO ethics frameworks, and the MIT RAISE Day of AI student materials you uploaded (Deep Fakes; the AI Ethics Debate; Truth, Tricks and AI; The Brain Behind the Bot; AI, Effort and Your Brain).

The detailed **teaching text, slide notes, H5P, debate and SCORM specifications, quiz items and rubric** were authored to fit that structure and are ready for your review, **localization into Nepali**, and adaptation by the delivery team.

**Version 1.0 · June 2026 · Himalayan Water deliverable for ADB TA 10376-REG.**