EdSutra

Teaching guide

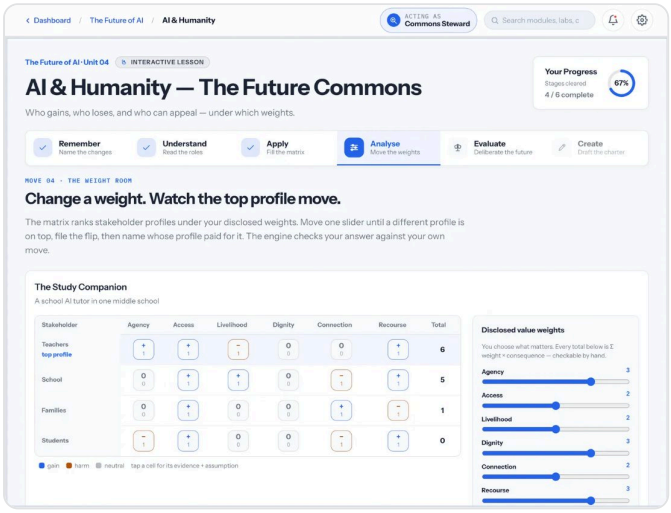
Module 08 · The Future of AI · Unit 04 · Stage 3, Responsibility

AI & Humanity

The Future Commons

Grades 7–12 | 30 minutes | Role: **Commons Steward**

Grade 7: About age 12–13 · Year 8 in the UK · Class 7 in India



The big idea

An AI deployment changes tasks, roles and relationships, and the changes land differently on each person. Which future looks best depends on the value weights you choose, so those weights should be said out loud. Learners show who gains, who loses and who can appeal, then design conditions that keep human choice alive.

The learner's job

Learners are the **Commons Steward** at The Future Commons · Community Table. The town has three futures for its AI tutor and its clinic queue. Your job is not to pick a winner — it is to make sure nobody can claim a future without showing its trade-offs. Their task: hold the table: show who gains, who loses, and who can appeal — under weights you disclose.

By the end, learners can say

- I can distinguish automation, augmentation, substitution, agency, dependency, access, and recourse
- I can predict how one AI deployment changes a task and a human role, then compare stakeholder perspectives
- I can fill a consequence matrix for an authored scenario using provided evidence and explicit assumptions
- I can change value weights and identify which recommendation flips, why, and who bears the cost
- I can deliberate among plausible futures and defend adoption conditions using evidence, distribution, and reversibility
- I can draft a Human Choice Charter: protected decisions, disclosure, opt-out, skill preservation, access, appeal, review triggers

Before class

- One computer or tablet for each learner or pair, with EdSutra open (this Module is part of the full EdSutra studio).
- Open the studio once on each device before class, so the Unit still runs if the internet drops.
- Play Remember and Understand yourself, so you know how the table of gains and harms reads.
- Pairs work well in Analyze and Evaluate: one learner moves the weights, the other says whose total paid for it.
- Best after AI & Climate and AI Safety.
- Units open in order, and Module 1 comes first. To start a class here, switch on Open every Unit in Settings → Teachers & parents.

Words to introduce

Automation: A task is done by the system instead of by a person.

Augmentation: A person's work changes, but the person stays central.

Substitution: A person leaves the loop entirely.

Dependency: Reliance on a system that is hard to reverse.

Recourse: A way to challenge a result and reach a human with authority.

Stakeholder: Anyone a decision affects, such as teachers, families or patients.

Value weights: How much each value, like agency or access, counts in a total.

Running the lesson 40 minutes

Open 4 min

"A school gets an AI tutor that grades homework. Is that good news? Good news for whom?"

Before anyone clicks, ask: when the tutor starts grading, what happens to the teacher's job? Does it disappear, shrink and change, or stay the same? Take a vote and come back to it in Understand.

Step	What learners do	Ask the class	Listen for
1. Remember 4 min	Learners match seven plain-language examples to seven concepts: automation, augmentation, substitution, agency, dependency, access and recourse.	• What is the difference between automation and substitution? • Which example showed a dependency forming?	Automation takes over a task; substitution takes the person out of the loop. They are different changes.
2. Understand 6 min	Learners predict how the Study Companion changes the teacher's role and what its biggest new dependency is, then tap each seat at the table to read its gains and harms.	• Look at our vote. Did the teacher's job disappear, or change? • Why does one tutor look different from each seat?	Grading shrank and coaching grew. Each group gains on some values and loses on others.
3. Apply 7 min	Learners attach the right evidence cards to each of the four stakeholder rows. Some cards, like a vendor's marketing sheet, are distractions.	• Why doesn't the vendor's marketing sheet count as evidence? • Which consequence rests on an assumption, with no evidence at all?	Evidence must be an observation about that group. A seller's claim or a training curve is not.
4. Analyze 7 min	Learners raise or lower one value weight until a different stakeholder's profile, its weighted total, is on top. Then they name whose profile fell the most.	• Which weight did you move, and who came out on top? • Whose profile paid for your move?	Changing a weight is a choice that moves the cost onto someone, and the totals show who.
5. Evaluate 6 min	A clinic considers an AI triage assistant. Learners choose auto-triage, a nurse-confirmed hybrid or a manual queue, back it with at least two evidence chips that fit, and name who bears the cost.	• What did your future gain, and what did it give up? • Who bears the cost of your choice?	No future is correct. A good defense uses evidence that fits the choice and names who pays.
6. Create 6 min	Learners draft a Human Choice Charter for the Study Companion, choosing one condition for each of seven slots, from protected decisions and opt-out to appeal and review.	• Which decision did you keep human-only, and why? • Which slot would you fight hardest to keep?	Keeping human choice is a design decision: named conditions that someone can check and enforce.

Steps follow Bloom's revised taxonomy, in the order the Unit presents them.

Deepen the thinking

Critical thinking	In this Unit the weights were shown. Where in real decisions are the weights hidden, and how could you find them?
AI thinking	If an AI tutor helps you every evening, how could you tell whether it is augmenting you or building dependency?
Ethics and fairness	Auto-triage serves more patients but drops the human check. Who should get the final say: patients, nurses or the clinic's managers?
Connect to life	Think of a self-checkout or an automated phone line. Could you reach a person when something went wrong?

Misconceptions to listen for

"Automation and substitution are the same thing."
In Understand the tutor takes over grading, but the teacher stays and coaches more. The task left; the person did not.

"One AI future is objectively best."
In Analyze, moving one weight puts a different group on top. Which future looks best depends on the weights.

"If it is efficient, it is good for everyone."
The table shows the school's efficiency gain beside lost peer help and connection. A gain on one value is not a total gain.

In real life the card learners see at the end of the Unit

- Self-checkouts automate a task; whether a staff member is still there to help is a separate choice.
- A study app that helps you learn augments you; one that does the work for you can build dependency.
- If an automated system blocks your account, a way to appeal to a person is called recourse.

Try it this week. Pick one change technology brought to your school or home. List who gained, who lost something, and who was never asked.

Watch out for. Efficient isn't the same as good for everyone. Which future looks best depends on what you weigh most, so say your weights out loud.

Talk about it. Which decisions about your life should always stay with a person, never a machine?

Exit ticket

1. What is the difference between automation and augmentation?

2. What is recourse, and why does it matter?

3. Why did changing one weight change who came out on top?

Going further

For a class project, groups write a Human Choice Charter for a real tool the school uses, such as a learning app: one protected decision, an opt-out, an appeal route and a review trigger.