EdSutra

Teaching guide

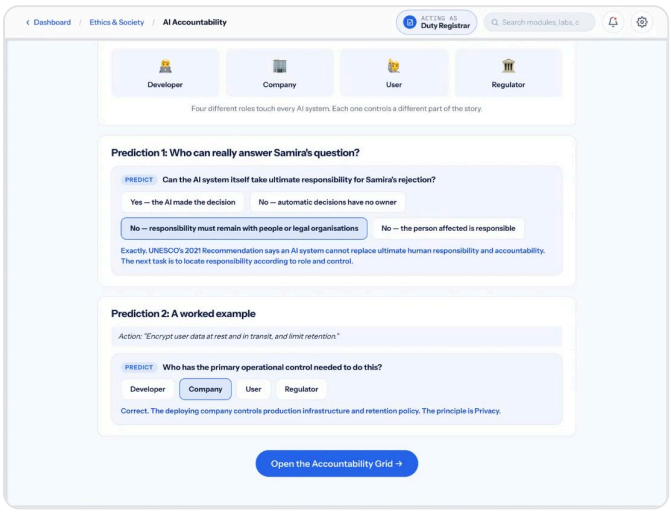
Module 05 · Ethics & Society · Unit 02 · Stage 3, Responsibility

AI Accountability

The Accountability Grid

Grades 7–12 | **45 minutes** | Role: **Duty Registrar**

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



The big idea

When an AI system causes harm, the system itself cannot answer for it. Responsibility stays with people and organizations, according to what each one controls: developers build and test, companies deploy and monitor, users decide how far to rely on it, and regulators set and enforce rules. Each duty has a primary home, but the work is shared.

The learner's job

Learners are the **Duty Registrar** at Northvale Register of Responsibility. Something will go wrong eventually. Today, everyone can still point at someone else. Their task: put a name against every duty. Before anything goes wrong.

By the end, learners can say

- I can differentiate the responsibilities of developers, companies, users, and regulators
- I can apply transparency, fairness, privacy, and safety principles to specific stakeholder actions
- I can identify distractor 'responsibility-washing' actions that deflect rather than allocate accountability
- I can defend a complete accountability map with reasoning grounded in stakeholder capability

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 the intro and the Capability Detective yourself, so you know where learners will click.
- Draw the grid on the board: Developer, Company, User, Regulator down the side; Transparency, Fairness, Privacy, Safety across the top.
- Best after What is AI? and AI Ethics & Bias.
- 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

Accountability: Being answerable for a harm: explaining it, fixing it and making it right.

Stakeholder: A role with a part in an AI system: developer, company, user, regulator.

Capability: What a role can actually do, like edit training data or stop a service.

Principle: One of the four lenses: transparency, fairness, privacy or safety.

Primary home: The role with the most direct capability to carry out an action.

Responsibility-washing: An action that looks like accountability but pushes blame away instead.

Running the lesson 45 minutes

Open 5 min

"When an AI system makes an unfair decision, everyone can point at someone else. So who really answers for it?"

Read Samira's story aloud from the intro, then ask: who is responsible for fixing it: the software, the school, the vendor, the developers, or nobody? Tally the votes and return to them in Evaluate.

Step	What learners do	Ask the class	Listen for
1. Remember 4 min	Read Samira's story, follow the life cycle (build, deploy, use, govern) and meet the four roles: developer, company, user and regulator.	• What are the four roles that touch every AI system? • Which role sets the rules, and which runs the live service?	Developers build, companies deploy, users rely on the output, and regulators set and enforce the rules.
2. Understand 4 min	Answer two predictions: can the AI itself take responsibility for Samira's rejection, and who controls encrypting user data and limiting how long it is kept?	• Why can't the AI system be the one responsible? • Why is encrypting data the company's job, not the developer's?	An AI cannot investigate, fix or make amends. The company runs the servers, so it controls storage and how long data stays.
3. Apply 8 min	In the Capability Detective, pick who has the real capability in four short cases. In the Principle Lens, sort four actions by transparency, fairness, privacy or safety. First answers lock.	• Who decides whether a medical record goes into a prompt? • Comparing results across groups: which principle is that, and why?	The user controls what they type in. Comparing outcomes across groups checks fairness, even if privacy also matters.
4. Analyze 15 min	Solve the grid: place 16 action cards in their one right cell and drop 4 responsibility-washing cards in the Reject tray. Then open Review to read who else shares each duty.	• The card never names who does it. What clue told you? • Why does putting all the blame on the user belong in the Reject tray? • Look at the Review: does a primary home mean only that role is involved?	Read what the action needs, then find who can do it. Rejects push blame away instead of taking on a duty.
5. Evaluate 9 min	In DEFEND, pick the placement that was hardest, say whether who or which principle made it hard, pick the clue that settles it, and say why naming who matters. Then take the quick check.	• Which placement was hardest for you, and what settled it? • Look back at our votes on Samira's case. Would you change yours?	Only people and organizations can investigate and remedy a harm. Naming who acts first shows where to start.

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

Deepen the thinking

Critical thinking	The grid gives each action one primary home. What does that make easier to see, and what does it hide?
AI thinking	A company says its model's behavior just appeared and nobody controls it. What could it still test, watch or limit?
Ethics and fairness	Samira was told the decision was automatic. Is it fair to give no way to appeal? Who should create one?
Connect to life	Think of an app you use. Which of its duties are yours as the user, and which are not?

Misconceptions to listen for

"The AI is responsible because it made the output."

An output cannot investigate harm or offer a remedy. Prediction 1 shows responsibility stays with people and organizations.

"The user is always the one to blame."

Users have real duties, like checking high-stakes answers, but the grid also gives duties to developers, companies and regulators.

"A primary home means only that role is involved."

Review lists who else shares every duty. The grid shows who acts first, not who acts alone.

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

- When a school or company says "the software decided", people still built, chose and ran that software.
- Apps often say "check important answers". That's a real user duty, but users aren't to blame for everything.
- Regulators can investigate when AI causes harm, using rules like data-protection laws, and order fixes.

Try it this week. With a parent or teacher, pick one AI tool you both know. Name who built it, who runs it, who uses it, and who checks on it.

Watch out for. An AI output can't be responsible: it can't fix harm or answer for it. Responsibility is shared among the people and organizations behind it.

Talk about it. If an AI tool made an unfair decision about you, who would you want to answer for it?

Exit ticket

1. Name the four roles in the Accountability Grid.

2. Why can't an AI system itself be held responsible?

3. What is responsibility-washing? Give one example from the Reject tray.

Going further

As a class project, groups pick an AI tool used at school or at home and list one real duty for each of the four roles, plus one responsibility-washing excuse they have heard and why it fails.