EdSutra

Teaching guide

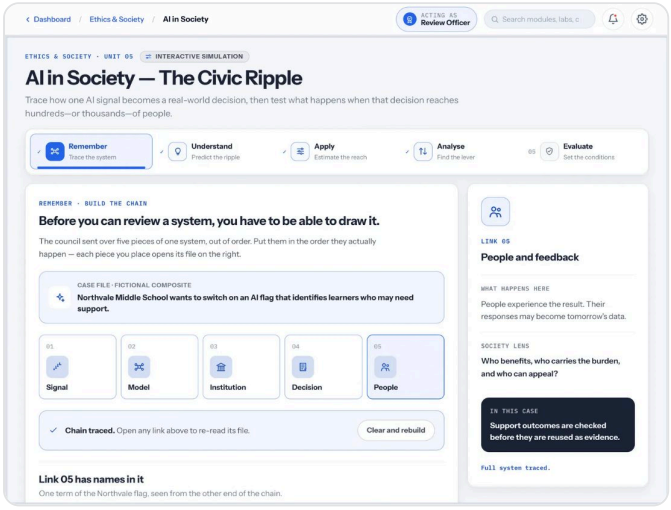
Module 05 · Ethics & Society · Unit 05 · Stage 3, Responsibility

AI in Society

The Civic Ripple

Grades 7–12 | 25 minutes | Role: **Independent Review Officer**

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



The big idea

An AI score changes lives only when an institution turns it into a decision that reaches people. Scale multiplies mistakes: a 6% miss rate across 1,000 decisions is 60 wrong calls. Human review helps only for the share of cases it covers, so a responsible launch matches its reach and safeguards to the stakes.

The learner's job

Learners are the **Independent Review Officer** at Northvale Civic Review Office. Three places want to turn on an AI system. Each one will make choices about real people. Their task: trace one AI decision from data to people. Make a finding you can defend.

By the end, learners can say

- I can identify the five links between an AI signal and a social outcome
- I can explain why institutions, operating rules, and feedback loops shape AI's impact
- I can calculate how reach, observed miss rate, and human review change unresolved misses
- I can compare deployment plans and locate which assumption creates the difference
- I can defend a launch posture with safeguards matched to the stakes

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 Remember step yourself, so you know where learners will click.
- Pairs work well: one learner drives, the other is the Review Officer who writes down each estimate and safeguard.
- Best after How Machines Learn 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

Institution: An organization, like a school, that decides what an AI output means in practice.

Operating rule: The rule that turns a score into an action, such as counselor review.

Reach: How many decisions a system makes in one cycle.

Miss rate: The share of cases the system gets wrong, seen in a trial.

Review coverage: The share of possible misses that a person checks.

Feedback loop: A system's outputs shape the next round of its data.

Safeguard: A protection such as human review, an appeal path or outcome monitoring.

Running the lesson 29 minutes

Open 4 min

"A school's AI flag is right 94 times in 100. That sounds good. So why might the school still need to think hard before switching it on?"

Before anyone clicks, ask: if the flag makes 1,000 decisions a term, how many will it get wrong? And if a counselor reviews half of them, how many slip through? Write the guesses down for Apply.

Step	What learners do	Ask the class	Listen for
1. Remember 5 min	Put the five links of the school's support flag in order, from input signal to people. Each link opens its file. Then read about three learners at the last link.	• Who wrote the rule at the institution link, and who can change it? • Priya, Marcus and Amara had three different outcomes. What were they?	A correct flag, a wrong flag, and a missed learner who needed help. The model is only one link of five.
2. Understand 4 min	Make two predictions before the answer shows: what changed when a school swapped counselor review for automatic removal from an activity, and which part of a feed reinforces itself.	• The model's score stayed the same. Why did the effect on learners change? • How can a feed help create the pattern it claims to observe?	The institution attached a harsher action to the same score. Clicks on what is shown become tomorrow's training data.
3. Apply 6 min	Run the estimator: pick a reach (200, 1,000 or 5,000 decisions) and a review coverage, read the unresolved misses, and save two different settings to compare.	• At 1,000 decisions and no review, how many misses are left? • Why does even full review not bring the number to zero? • Look at our guesses from the start. How close were we?	60, since 1,000 times 6% is 60. Reviewers catch only 75% of what they see, so some misses remain.
4. Analyze 4 min	Compare two pairs of rollout plans side by side and name the input that explains the gap: reach, review coverage or the observed miss rate.	• The district rollout and the pilot share a 6% miss rate. What differs? • Why is one "impact score" less useful than three separate inputs?	Reach and review coverage changed together. Keeping inputs separate shows which lever made the difference.
5. Evaluate 6 min	For at least two of three cases (a school flag, a job screen, a library feed), choose launch, limited pilot or pause, and attach two or three safeguards that fit the stakes.	• Why can you pick at most three safeguards? • Which safeguard does the job screen need that the library feed does not?	Each safeguard costs time or money. High-stakes cases need human review or an appeal path; lower-stakes ones need monitoring or opt-out.

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

Deepen the thinking

Critical thinking	Amara was never flagged, so there was nothing to appeal. How could a school find the people its system never notices?
AI thinking	The flag reads attendance and grades. What could it never see, and how does that shape who it misses?
Ethics and fairness	Marcus spent a term explaining himself because nobody told him a record existed. What should people be told when AI scores them?
Connect to life	Where in your life does one decision reach many people at once, like a timetable or an app update?

Misconceptions to listen for

"A 94% success rate means the errors don't matter."

Apply multiplies the 6% miss rate by reach: 1,000 decisions leave 60 wrong calls before any review.

"Human review fixes every mistake."

Review covers only a share of cases and catches 75% of what it sees. Misses the flag never raised never reach it.

"AI's impact comes from the model alone."

The chain shows the institution's rule matters too: the same score caused more harm once the school attached a harsher action.

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

- When a school or hospital uses an AI score, its rule decides whether a person reviews it first.
- A headline says an AI is "94% accurate". Across 1,000 cases, that's still 60 misses.
- Video apps reach millions of people, so even a tiny error rate can touch huge numbers.

Try it this week. Take a "95% right" claim from an ad or the news. Work out how many people the other 5% would be at your school.

Watch out for. A small error rate stops being small at scale, and human review only helps for the cases it actually covers.

Talk about it. Should a new AI tool start as a small pilot before everyone uses it? Why?

Exit ticket

1. Put the five links in order, from signal to people.

   _____

2. A tool with a 6% miss rate makes 5,000 decisions. How many misses?

   _____

3. Name one safeguard for a high-stakes decision, and say why it fits.

   _____

Going further

As a class project, groups design a pilot for an AI tool their school might use: how many people it reaches first, who reviews it, how people appeal, and what result would make them stop.