EdSutra

Teaching guide

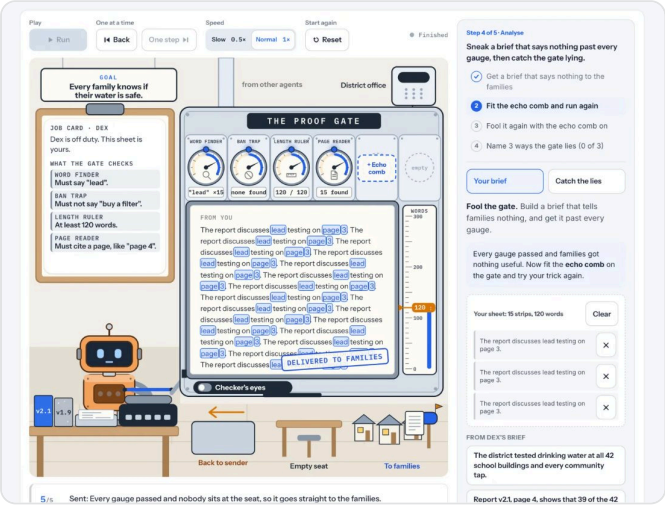
Module 09 · Intelligence Management · Unit 03 · Stage 3, Responsibility

Delegation & Verification

The Proof Gate

Grades 7–12 | **25 minutes** | Role: **Handoff Checker**

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



The big idea

An AI agent does exactly what its job card says, and guesses wherever a line is blank. Automatic checks catch some mistakes, but they only count and search words: a brief can pass every check and still say nothing, or use last term's numbers. Tone, truth and meaning stay a person's job.

The learner's job

Learners are the **Handoff Checker** at Northvale District Office · Handoff Room. AI agents write the district's notices now. Every handoff needs a job card and checks, or mistakes go straight to families. Their task: write the job card an AI agent works from, then prove what comes back before families see it.

By the end, learners can say

- I can name the parts of a handoff: goal, task, context, constraint, deliverable, acceptance check, evidence and escalation
- I can predict what an agent does with a blank line on its job card, and which check (if any) catches it
- I can write a complete job card, and make the human call on briefs that pass every automatic check
- I can diagnose false acceptance, false rejection, gaming the test, stale context and a fake score on a real gate
- I can choose which job card to trust, backed by facts from your own runs

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 first two steps yourself, so you know where learners will click.
- Pairs work well: one learner writes Dex's job card, the other sits at the seat and makes the human call.
- Best after What Is Intelligence Management? and Human Oversight.
- 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

Job card: Six lines an agent works from: task, context, constraints, deliverable, checks, escalation.

Context: What the agent must read, with its version: report v2.1, not v1.9.

Constraint: A hard limit, such as no selling and a page for every number.

Acceptance check: A test the work must pass before it moves on. Each gauge is one.

Escalation: Who to call when work cannot be checked or crosses a line.

Gaming the test: Work built to pass the checks without doing the job.

False acceptance and rejection: A check too weak lets bad work through; too strict throws good work away.

Running the lesson 32 minutes

Open 4 min

"You ask a helper to write a notice for families, and it comes back perfect on every check. Can it still be wrong? How?"

Before anyone clicks, ask: if a robot agent's instructions leave out which report to use, will an automatic check notice? Take a show of hands and come back to it in Understand.

Step	What learners do	Ask the class	Listen for
1. Remember 4 min	Tap Run and watch Dex, a robot agent, type a brief from its job card and the Proof Gate's gauges check it. Then name the 8 pinned parts of a handoff.	• What is the difference between the goal and the task? • When a gauge lights up words, what does that show?	The goal is what we really want; the task is the exact job. Lit words show what each check found.
2. Understand 5 min	Three rounds, each with one line left off Dex's job card: constraints, deliverable, then context. Learners guess which gauge will stop the brief, then run it and see what Dex guessed.	• With no limits, what did Dex add, and which gauge caught it? • With no report named, why did every gauge pass? Look back at our show of hands.	Dex fills every blank with a guess. A gauge only catches what someone thought to check; none compares numbers.
3. Apply 7 min	Learners fill all six lines of Dex's job card so its brief reaches their seat, and send it. Then they run five briefs from other agents and judge those that reach the seat.	• Agent Echo's brief passed every gauge. Why send it back? • Agent Sage cites page 19 of a 14-page report. What does your job card say to do?	Passing the gauges is not the same as saying something true. The person at the seat makes the final call.
4. Analyze 7 min	Learners build a brief that says nothing and get it past every gauge, then fit the echo comb and fool it again. Then they replay five cases and name 3 ways a gate lies.	• The echo comb bans repeated sentences. How did you still fool it? • Raising the floor to 200 words threw out Dex's good brief. Why is stricter not always better?	Checks count and search words; they cannot read meaning. Work can be built to pass them without doing the job.
5. Evaluate 5 min	Two cases, each with two job cards: one for a brand-new agent, one for a high lead retest at Hillcrest School. Learners run both cards, pick a verdict, and back it with two facts.	• The Hillcrest notice passed the usual gauges. What did it never say? • Which fact from your runs best backs your verdict?	A verdict backed by facts from the runs. The routine card was built for routine news; serious news needs a person.

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

Deepen the thinking

Critical thinking	A rubric can check an essay's length and that it cites sources. What can it not check, and who should?
AI thinking	The alarm meter counted scary words to judge tone and got it wrong. What would a machine need to really judge tone?
Ethics and fairness	If Dex's brief with last term's numbers reached families, who is responsible: Dex, the job card writer, or the district? Why?
Connect to life	Think of a time someone did exactly what you asked, but not what you meant. Which line of your instructions was missing?

Misconceptions to listen for

"The AI will work out what I meant."

In Understand, Dex fills every blank line with a guess: a sales line, a slide, last term's report. It does what the card says.

"If it passes every check, it is good."

Agent Echo's brief repeats one sentence 30 times and passes every standard gauge. Checks count and search words; they cannot tell if anything was said.

"Stricter checks are always better."

At a 200-word floor, Dex's complete 159-word brief is thrown out, while Echo's 257 words of one sentence would pass.

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

- An AI answer existing doesn't mean the job is done; it still has to pass your checks.
- A teacher's rubric is a set of acceptance checks; qualities like "interesting" still need a human judge.
- A keyword or word-count checker can be gamed: writing can tick every box and still say nothing.

Try it this week. Before asking someone for help this week, write down what "done" looks like: the task, what to use, the limits and one check.

Watch out for. Passing every automatic check doesn't prove good work, because writing can game the test. And a stricter check can throw good work away.

Talk about it. What's something you'd hand to a helper, and how would you check it was done well?

Exit ticket

1. Name four of the eight parts of a handoff.

2. Why did every gauge pass when Dex used last term's report?

3. Give one thing a gauge cannot check that a person can.

Going further

In the Open Project Canvas Lab, the Agent Loop keeps going because its Check state has no rule for "done". Writing one is an acceptance check. For homework, use the "Try it this week" activity from the In real life card.