EdSutra

Teaching guide

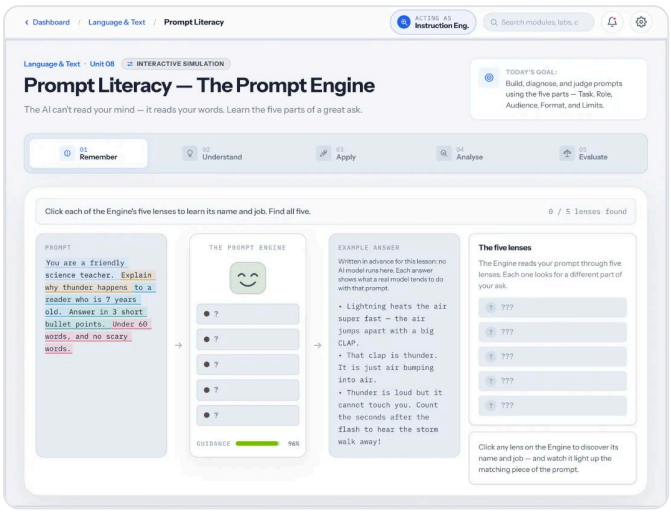
Module 03 · Language & Text · Unit 08 · Stage 2, Capabilities

Prompt Literacy

The Prompt Engine

Grades 7–11 | **45 minutes** | Role: **Instruction Engineer**

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



The big idea

The AI cannot read your mind; it reads your words. A strong prompt has five parts: Task, Role, Audience, Format and Limits, and each one steers a different side of the answer. When an answer disappoints, the cause is usually a part that was vague, missing or fighting another.

The learner's job

Learners are the **Instruction Engineer** at Bureau of Plain Language · Engine Room. A good question you wrote once helps one person. A good instruction helps everyone who uses it after you. Their task: build an instruction that still works when you are not here.

By the end, learners can say

- I can name the five parts of a strong prompt and what each one controls
- I can explain how adding each part changes the same AI answer
- I can build complete prompts for real goals by choosing the right part for each slot
- I can diagnose why a prompt produced a bad answer — vague, missing, or conflicting parts
- I can judge competing prompts against a rubric and choose which one to ship

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 how the lenses light up and what the Guidance meter shows.
- Tell learners up front: the Engine's answers were written in advance. No AI model runs; they show what real models tend to do.
- Best after Language & AI Intro and Prompting AI.
- 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

Task: The action you want, such as explain, list or compare.

Role: Who the AI should act as; it sets the voice.

Audience: Who the answer is for; it sets how hard the words are.

Format: The shape of the answer: a list, steps, a table, a story.

Limits: Rules the answer must follow: how long, what to include or leave out.

Guidance meter: The Engine's gauge of how clearly a prompt steers the answer.

Rubric: A checklist for judging things the same way each time.

Running the lesson 45 minutes

Open 5 min

"If an AI cannot read your mind, what exactly does it need from your words to give you the answer you pictured?"

Before anyone clicks, write "Tell me about black holes." on the board. Ask: what will the answer be like, and what would you add first to improve it? Come back to the guesses in Understand.

Step	What learners do	Ask the class	Listen for
1. Remember 4 min	Learners click each of the Engine's five lenses to learn its name and job, and find all five. Each lens lights up its part of a sample prompt about thunder.	• Which part of the thunder prompt sets who the answer is for? • What would the answer lose without the Limits line?	Each part has its own job: Task is the action, Role the voice, Audience the reader, Format the shape, Limits the rules.
2. Understand 7 min	Learners run the vague prompt "Tell me about black holes.", then add one part at a time: Task, Audience, Format, then Role and Limits. The answer and the Guidance meter change each time.	• Which upgrade changed the answer most, and how? • Look at our board guesses. Did the part we chose first help most?	Same topic, same AI; only the words changed. Each part steered one thing: aim, difficulty, shape, voice, length.
3. Apply 10 min	Two missions: calm a scared 7-year-old about thunder, and find five fast facts for a honeybee poster. Learners pick one chip per slot, run the Engine, and need a score of 4 out of 5.	• Why is "You are an AI assistant" a weak role? • A chip can be specific and still wrong for the job. Which one fooled you?	The best chip fits this goal and this reader. Specific is not enough: a lecture voice is specific but wrong here.
4. Analyze 10 min	Four broken prompts and their bad answers: an essay rewrite, a vaccine talk too hard for a 6-year-old, a pancake wall of text, an impossible summary. Learners name the part that failed.	• In "Make my essay better", better how? • The vaccine prompt had a doctor role. Why did the answer still fail? • Can any AI follow "three words" and "every chapter in detail" at once?	The fault is in the ask: a vague task, a missing audience, a missing format, or limits that fight.
5. Evaluate 9 min	The Prompt-Off: three rival prompts for one job, a volcano study sheet for tomorrow's quiz. Learners tick the rubric for each prompt, run them to see the answers, then ship the best one.	• Prompt B is friendly and clear. Why is it still not the one to ship? • What does Prompt C have that Prompt B is missing?	C fills all five parts. B lacks a format and limits, so it reads like a chat, not a study sheet.

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

Deepen the thinking

Critical thinking	The Engine's answers were written in advance to show what real models tend to do. How could you test whether a real AI behaves the same way?
AI thinking	The AI answered the vaccine question for an average adult because nobody said otherwise. Why would a model fall back on an "average" reader?
Ethics and fairness	Your instruction will be used by people you never meet. What could go wrong for a younger or nervous reader if you leave out the Audience?
Connect to life	Pick a set of instructions you follow often, like a recipe, game rules or a homework sheet. Which of the five parts does it use, and which is missing?

Misconceptions to listen for

"The AI gave a bad answer, so the AI is bad."

In the Autopsy, every bad answer traces back to the ask: a vague task, a missing audience or format, or limits that fight.

"The AI knows what I mean."

The Bedtime Backfire case never mentions the 6-year-old cousin, so the AI writes for an adult. It cannot use information you never gave it.

"Prompting is about typing tricks or magic words."

The five parts are plain, structured communication, the same skill as briefing a person well. The Prompt-Off rewards a complete prompt, not a clever one.

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

- Asking a voice assistant or chatbot for help works best with all five parts: task, role, audience, format, limits.
- Leaving notes for a babysitter or substitute teacher is prompting too: task, audience and limits all matter.
- In a group chat, "make slides" gets guesses, but "five slides for our class by Friday" gets what you need.

Try it this week. Write instructions for making a sandwich using all five parts. Have someone follow them exactly as written. What went wrong?

Watch out for. A bad answer doesn't always mean a bad AI. Check your ask first: it can't use information you never gave it.

Talk about it. Who is the best explainer you know? Which of the five parts do they always include?

Exit ticket

1. Name the five parts of a strong prompt.

2. "Make my essay better." Which part is weak, and how would you fix it?

3. What goes wrong when two limits contradict each other?

Going further

In the AI Story Canvas Lab, learners build a choose-your-own-adventure story and play every path, checking that a stranger could follow each choice. For homework, use the "Try it this week" activity from the In real life card.