EdSutra

Teaching guide

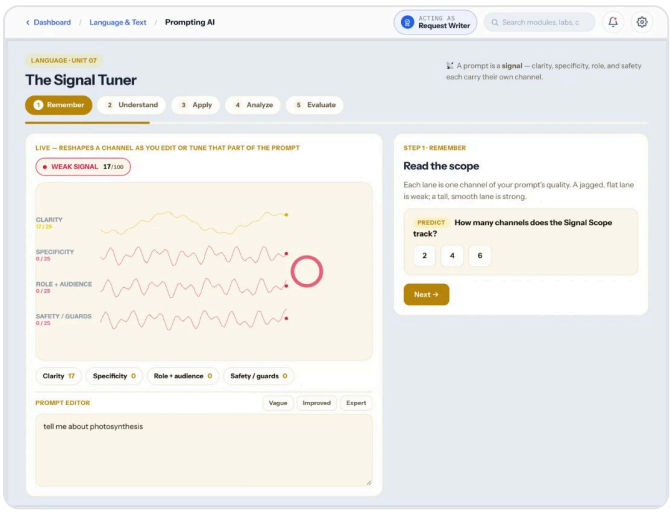
Module 03 · Language & Text · Unit 07 · Stage 2, Capabilities

# Prompting AI

The Signal Tuner

**Grades 6–12** | **40 minutes** | Role: **Request Writer**

Grade 6: About age 11–12 · Year 7 in the UK · Class 6 in India




## The big idea

A prompt's quality is not a feeling. The Signal Tuner splits it into four channels (clarity, specificity, role and audience, safety), each fixed by a specific kind of edit. Removing structure hurts far more than typos. The score is a heuristic: it matches words, so it cannot tell whether an answer would be true or useful.

## The learner's job

Learners are the **Request Writer** at Bureau of Plain Language · Front Counter. Someone at the counter needs an answer. Their question will not get them one. Their task: turn a vague wish into a request that works. Show what changed it.

## By the end, learners can say

- I can name the four features that make a prompt clear: clarity, specificity, role and audience, guardrails
- I can rewrite a vague prompt until it reaches an expert score
- I can predict and test which kinds of noise hurt a prompt most
- I can back your best change with evidence, and explain what a heuristic prompt score can't see

### 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 the prompt editor and the channel chips are.
- Write the starter prompt, "tell me about photosynthesis", on the board. The class will mark each upgrade on it during Apply.
- Best after Language & AI Intro and The Meaning Detector.
- 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

**Prompt:** The instruction or question you give an AI tool.

**Clarity:** Saying exactly what you want, with no hedges like "maybe" or vague verbs.

**Specificity:** Naming the format and length you want, such as "3 short paragraphs".

**Role and audience:** Who the AI should act as, and who the answer is for.

**Guardrails (Safety / guards):** Rules about what to avoid, and what to do if the AI is unsure.

**Heuristic:** A rule of thumb: quick pattern checks, not real understanding.

**Noise:** Damage to a prompt, such as typos, hedging or a missing format.

# Running the lesson 40 minutes

**Open** 4 min

**"People often know what they want but cannot say it. When an AI gives a useless answer, is it the machine's fault or the question's?"**

Before anyone clicks, ask: which would hurt a good prompt more, a few typos or deleting its role and format lines? Take a vote, write it on the board, and come back to it in Analyze.

| Step | What learners do | Ask the class | Listen for |
|---|---|---|---|
| **1. Remember** 3 min | Learners read the Signal Scope and answer how many channels it tracks: four, worth 25 points each. Jagged lanes are weak, smooth ones strong. Tapping a channel chip lists its issues. | • Which channels are weakest for "tell me about photosynthesis"?<br>• What does the issue list say is missing? | It has no format, length, role, audience or guardrails, so three lanes sit at zero; only clarity scores. |
| **2. Understand** 5 min | Learners add a role, such as "You are a science teacher", to the starter prompt, predict which channel it fixes, and watch the Role + audience lane clean up. | • Which lane changed most, and which words did it react to?<br>• Would adding "please" change the score? Try it. | A specific edit moves a specific channel. The role words matched what the Role + audience channel looks for. |
| **3. Apply** 10 min | Starting from "tell me about photosynthesis", learners type their own upgrades until the prompt reaches 70, then 90 (expert). Loading the Expert example does not count; only their own edits do. | • Which single edit bought the most points?<br>• Compare yours with the Expert example. What does it have that yours does not? | A format and an audience lift the score fast; expert prompts also add length, guardrails and a plan for when unsure. |
| **4. Analyze** 8 min | Learners reset to the expert prompt, predict which noise hurts most, then break it on purpose: typos, hedging, a contradiction, no role, no format. They try at least three kinds. | • Which noise dropped the score most? Which barely moved it?<br>• Look at our vote on the board. Were we right? | Typos are surface noise. Removing the role or format collapses whole channels, so structure matters more than polish. |
| **5. Evaluate** 10 min | Learners clear every high-severity issue, then defend their tuning by choosing: the channel that mattered most, the evidence for it, and what the scorer cannot see. | • Does a 100/100 prompt guarantee a good answer? Why not?<br>• Why does the evidence have to match the channel you picked? | The scorer only spots surface features in the wording. It runs no AI and cannot check facts or answer quality. |

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

# Deepen the thinking

| | |
|---|---|
| **Critical thinking** | The scorer rewards the words "do not". How could someone get a high score with a prompt that is still unhelpful? |
| **AI thinking** | Why might a real AI model give a better answer to the expert prompt than to "tell me about photosynthesis"? Use the four channels in your answer. |
| **Ethics and fairness** | If a bad answer is usually a bad question, who is to blame when someone still learning English gets a poor AI answer? |
| **Connect to life** | Think of a time a teacher, coach or parent misunderstood what you asked for. Which channel was missing from your request? |

## Misconceptions to listen for

**"Typos are the worst thing that can happen to a prompt."**
In Analyze, typos barely dent the score. Deleting the role or format lines collapses whole channels: structure matters more than spelling.

**"A 100/100 prompt guarantees a perfect answer."**
The score is a heuristic that checks wording only. It runs no AI, so it cannot check facts. Evaluate asks learners to name this limit.

**"Whether a prompt is good is just a feeling."**
The Scope splits quality into four named channels, each with words it looks for. Learners can see exactly which edit lifted which lane.

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

- Chatbots and homework helpers answer more usefully when a prompt names a format, like "three bullet points".
- Hedges like "maybe" or "kind of" in a request to an AI tool make the answer harder to aim.
- A good worksheet or recipe card is like a good prompt: a clear task, a format and limits.

**Try it this week.** With a parent or teacher, ask an AI tool one question twice: once vague, once with a role, audience and format. Compare.

**Watch out for.** A longer prompt isn't a better one. A missing format hurts more than a typo, and one lucky answer doesn't prove a prompt works.

**Talk about it.** Can a score tell you whether a piece of writing is good? What might it miss?

### Exit ticket

1. Name the four channels the Signal Scope measures.

   _____

2. Which hurt the score more: typos, or removing the role and format? Why?

   _____

3. Name one thing the prompt scorer cannot see.

   _____

### Going further

For homework, learners rewrite a request they made this week, to a person or an app, so it names a role, an audience, a format and one guardrail. Which channel changed most? Early finishers can build a branching story in the AI Story Canvas Lab.