

Lesson 05 — How to talk to AI well

The skill that separates “AI is useless” from “AI is incredible.” The model is almost never the problem — the prompt usually is.

AI for Dummies

Module 5

Heavy AI users aren't gifted — they learned one habit: **give the AI everything it needs, right there in your message.**

The one rule that explains everything

Vague in → vague out Specific in → useful out.

✗ VAGUE PROMPT

"recycling lesson ideas"

What you get: a wall of generic bullets that could suit any age, any duration, any classroom in the world. You spend 30 minutes adapting it.

✓ SPECIFIC PROMPT


"You are an experienced kindergarten teacher. I need three hands-on recycling activities for 5-year-olds, 30 min total. List each as a numbered step with materials and timing."

What you get: three print-ready activities tuned exactly to your class. 2 minutes of effort.

Key insight: The AI has no memory of you and no idea what "good" means to you. It only knows what is in your message right now. Give it everything it needs.

The recipe: Role → Context → Task → Format

Four ingredients. Every strong prompt has at least three. All four = near-perfect output.



You don't need all four every time. A quick casual question can survive on Task alone. But the more the output matters, the more ingredients you should add. The sweet spot for real work is all four.

R — Role: put it in the right frame of mind

The Role ingredient sets the perspective, vocabulary, and depth of response. It's the first thing to add when results feel generic.

EXAMPLE ROLES

"You are an experienced kindergarten teacher."

"You are a plain-English solicitor."

"You are a senior data analyst who explains findings to non-technical executives."

Why Role works

Without a role, the AI defaults to a generic helpful voice — useful for no one in particular. Giving a role forces it to adopt the right vocabulary level, the right assumptions, and the right perspective.

"You are a nutritionist" gets different (better) meal advice than no role at all.

The role can be a job title, an expertise level, or a perspective: *"You are a sceptical editor."*

C — Context: give it the background it can't guess

The AI can't ask follow-up questions the way a colleague would. Give it the context *upfront*.

"I'm planning a 30-minute lesson for 5-year-olds about recycling. They have short attention spans and learn best through movement and colour."

"Our SaaS product had a 4-hour outage yesterday. About 2,000 customers were affected. We've fixed the root cause."

"I'm a non-technical finance manager. I need to brief my CEO on blockchain in 2 minutes at tomorrow's board meeting."

What to include in Context

- Who you are and what you're working on
- Who the audience is (age, expertise, role)
- Any relevant constraints (time, budget, length)
- What you already know or have tried
- Any background facts the AI needs to use

Rule of thumb: If a smart new colleague would need to ask you that question on day one, put the answer in Context.

T — Task: be specific about what you want

Vague tasks produce vague output. The Task ingredient specifies *scope*, *quantity*, and *nature* of the output.

❌ VAGUE TASK

"give me ideas"

The AI will produce a wide, shallow list. Could be 3 ideas, could be 20. Could be high-level, could be tactical. You don't know what you'll get.

✅ SPECIFIC TASK

"Give me three hands-on activities."

Quantity: three. Nature: hands-on. Scope: activities. The AI knows exactly what to produce.

Task precision checklist

- **Verb** — write, summarise, compare, explain, fix, translate, list...
- **Quantity** — three ideas, one paragraph, five bullets...
- **Scope** — about X, focusing on Y, limited to Z...
- **Audience** — for a beginner, for my manager, for a press release...

The more of these you specify, the tighter the output.

F — Format: tell it how the answer should look

Without format guidance, the AI defaults to flowing prose paragraphs. **That's rarely the most useful shape.**

"List each as a numbered step with the materials needed and the time it takes."

"Give the answer as a table with columns:
Option | Pros | Cons | Cost."

"100 words max. Plain language. No jargon."

"Format as an email: subject line, 3 short paragraphs, friendly sign-off."

Format options cheat sheet

- **Table** — comparisons, options, specs
 - **Numbered list** — instructions, steps
 - **Bullet list** — quick scanning
 - **JSON / code** — structured data
 - **Word limit** — "50 words max"
 - **Document sections** — intro / body / summary
 - **Email / report / tweet** — named templates
-

All four pieces — assembled into one prompt

This is what a complete, ready-to-paste prompt looks like. Each colour is one ingredient.

You are an experienced kindergarten teacher.

I'm planning a 30-minute lesson for 5-year-olds about recycling. They have short attention spans and learn best through movement and colour.

Give me three hands-on activities.

List each as a numbered step with: the materials needed, the time it takes, and one tip for keeping 5-year-olds engaged. Keep each description under 70 words.

● Role ● Context ● Task ● Format

What you get back: three print-ready activities, each labelled, timed, and adapted for 5-year-olds. No editing required. Copy straight into your lesson plan.

Vague vs. Specific — the AI's actual responses

✗ PROMPT: "RECYCLING LESSON IDEAS"

AI RESPONSE (mock)

Here are some recycling lesson ideas:

- Teach the three Rs: Reduce, Reuse, Recycle.
- Have students sort items into recycling categories.
- Show a video about recycling and discuss.
- Create art projects using recycled materials.
- Discuss why recycling matters for the environment.

These activities can be adapted for any age group.

✓ PROMPT: THE FULL FOUR-PART RECIPE

AI RESPONSE (mock)

1. **The Sorting Game (8 min)**

Materials: box of mixed items, 3 labelled bins.

Each child picks one item and places it in the right bin. Ask "why?" after each. *Tip: use bright, oversized labels.*

2. **Bottle-Cap Mosaic (12 min)**

Materials: bottle caps, glue sticks, card stock.

Kids arrange caps into a shape. *Tip: pre-draw outlines for struggling children.*

3. **Paper-to-Paper Circle (10 min)**

Materials: newspaper, water, mesh screen, sponge. Demonstrate recycled paper; kids press pulp themselves. *Tip: let them feel the wet pulp first.*

Same topic. Wildly different utility. The vague response could have been written by anyone, for anyone. The specific response was written for *you*, for your class, right now.

Habit 1 — Show an example (the most underused trick)

Words like “catchy,” “professional,” or “friendly” mean different things to different people. **An example is unambiguous.**

✗ WITHOUT EXAMPLE

"Write a catchy subject line for our product launch."

The AI guesses what "catchy" means to you. Could be punny. Could be urgent. Could be click-bait. You get a lottery ticket.

✓ WITH EXAMPLE

"Write a subject line like this one: 'Your invoice is ready – 3 things to check.' Now write one for our product launch on Tuesday."

The AI matches your style, tone, and structure exactly.

Where to use examples

- Email subject lines & body copy
- Job ad or social-media post style
- Report or document formatting
- Code style ("write functions like this one")
- Tone of voice ("write in the style of this paragraph")

Why it works: Pattern-matching is what LLMs are best at. Give them a pattern to match.

Habit 2 — Ask for the format you want

AI defaults to paragraphs. **Paragraphs aren't always what you need.** Say so — it complies every time.

"Give the answer as a table with columns:
Tool | Free? | Best for | Biggest downside."

"Format as a numbered step-by-step guide. No
prose. Steps only."

"50 words maximum. Plain English. No sub-
headings."

"Format as a tweet thread — 5 tweets, each
under 280 characters, numbered."

Why explicit format matters

The AI has no idea whether your output will be pasted into a spreadsheet, read aloud, or put on a slide. You do. **Tell it.**

Format guidance also caps length. A prompt without a word limit will often produce more than you need.

Quick win: Add *"Give your answer as a table"* to any comparison question and the result becomes immediately scannable.

Habit 3 — Tell it what *not* to do

Positive instructions alone sometimes aren't enough. A "don't" constraint closes the bad paths before the AI wanders down them.

"Don't use jargon or technical terms."

"Don't make anything up — if you're unsure, say so."

"Don't suggest anything that requires a paid subscription."

"Don't start the email with 'I hope this message finds you well'."

"Don't give me more than five options."

The hallucination stopper

The single most useful "don't" for factual questions:


"If you don't know, say you don't know. Don't guess."

This dramatically reduces hallucinations because it makes the model flag uncertainty rather than filling gaps with confident-sounding nonsense.

Use "don't" for style, scope, and facts. Think of it as building a fence around the output space.

Habit 4 — Iterate, don't restart

Your first prompt is a first draft, not a final request. Stay in the same chat and steer.



Refinement commands that always work

- `"Shorter."` — cuts length
- `"More formal."` — adjusts tone
- `"Put that in a table."` — reformats
- `"Now add a section on cost."` — extends
- `"Remove bullet 3."` — trims

The chat history carries all context. You don't repeat yourself — you build on what's already there. This is faster than starting over *every single time*.

Habit 4 in action — watch the refinement loop

▶ Replay

Three small messages in the same conversation. No rewriting from scratch. Each step builds on the last.

Habit 5 — Let it think for hard problems

Adding “*think step by step*” genuinely improves accuracy on reasoning tasks. It’s the equivalent of asking someone to show their working.

✗ WITHOUT STEP-BY-STEP

Prompt: "If a train leaves at 09:15 and takes 2h 45min, plus a 20-minute connection, what time does it arrive?"

AI response (mock): It arrives at 12:20.

Wrong. The AI skipped the addition and landed on a plausible-sounding answer.

✓ WITH “THINK STEP BY STEP”

Prompt: "...Think step by step."

AI response (mock):

1. Depart 09:15.
2. Add 2h 45min → 12:00.
3. Add 20-min connection → 12:20.
4. Final arrival: **12:20.**

Each step is visible. Errors are easy to spot and correct.

Phrases that trigger step-by-step thinking:

"Think step by step." · "Show your reasoning." · "Walk me through your logic."

When to use it: logic puzzles, multi-step maths, planning decisions, comparing options. Don't use it for simple tasks — it adds unnecessary length.

Common mistake 1 — Being too vague

✗ THE MISTAKE

"help me write something"

"give me ideas"

"explain blockchain"

These prompts work, but they return generic outputs that could apply to anyone, anywhere. You spend more time adapting than you saved by asking.

✓ THE FIX

Add Role + Context + Task. Even Role + Task doubles quality. Even one sentence of context narrows the output dramatically.

"Explain blockchain to a non-technical finance manager who needs to brief their CEO in 2 minutes."

Why vagueness is expensive

When the prompt is vague, the AI has to guess. It guesses middle-of-the-road. Middle-of-the-road output for your specific situation is almost never good enough to use without heavy editing.

Specificity is the cheapest form of editing. Five extra seconds writing the prompt saves five minutes fixing the output.

Rule: If a stranger reading your prompt couldn't tell what a "good" answer looks like, the AI can't either.

Common mistake 2 — Ten things in one prompt

✗ THE MISTAKE

```
"Summarise this article, then rewrite it in a friendlier tone, then translate it to French, then suggest a tweet thread, then give me a 10-word headline, then check it for factual errors."
```

The AI tries to do everything at once and cuts corners on each. The summary is thin, the rewrite is rushed, the translation is mechanical.

✓ The fix: one task per message

Chain the tasks sequentially in conversation instead:

1. Message 1: "Summarise this article."
2. Message 2: "Now rewrite it in a friendlier tone."
3. Message 3: "Translate that to French."

Each task gets full attention. Each output is complete. The iteration habit (Habit 4) handles the rest.

One clear exception: you can ask for multiple *related items of the same type* in one prompt — e.g., "give me five subject line options" is fine. The rule is one *task type*, not one item.

Common mistake 3 — Trusting the first answer on facts

Warning: AI can generate false information with complete confidence. A statistic, a quote, a date, a job title — these can all be *wrong* even when the AI sounds certain. This is called a **hallucination**.

What hallucination looks like

You ask: *"What did Einstein say about imagination?"*

The AI produces a confident-sounding quote with a citation. The quote sounds exactly right. **It may be entirely fabricated.**


Higher risk with: niche topics, specific numbers, recent events, named people.

The fix: verify anything that matters

- Never cite an AI-generated fact without checking it.
 - Add *"if you're not sure, say so"* to factual prompts.
 - Ask for sources — then actually look them up.
 - Treat AI output as a first draft to be verified, not a final answer.
-

Common mistake 4 — Forgetting the knowledge cutoff

Most AI models stopped learning at a specific date — usually six months to a year or more before you're talking to them. **Anything after that date is unknown to the model.**



✗ WILL PRODUCE WRONG ANSWERS

"What is the current price of Bitcoin?"

"Who won the 2025 election?"

"What are the latest AI models available?"


✓ THE FIX

Use a tool with web search (Perplexity, Claude with web browsing, ChatGPT browsing mode).

Or: paste the current information directly into your prompt and ask the AI to work from what you've given it.

The mental model: AI as a smart intern

Think of the AI as a brilliant generalist who just joined your team **this morning** and knows nothing about your project, your company, or your preferences.



What this means in practice

- The intern is *capable of almost anything* — but needs direction.
- A vague “handle this” gets a generic result.
- A clear brief gets exactly what you need.
- You don’t blame the intern for not reading your mind.
- Good directors give specific notes — they don’t just say “be better.”

The shift: You are not a user of a search engine. You are the director of a capable assistant. Your job is to give a good brief.

Try it now — build a real prompt

Pick a task you actually have this week. Write a four-part prompt for it right now.

You are [your role here].

I am [who you are].

I need [the background they need to know].

[Any constraints or audience info].

[Exactly what you want — verb + scope + quantity].

[How you want the answer: table / list / email / word limit / format].

Don't [anything you want excluded].

● Role ● Context ● Task ● Format ● Negation

Starter ideas if you're stuck

- An email you've been putting off writing
- A confusing document you need summarised
- A decision you want to think through
- A meeting agenda for next week
- A job ad, report, or proposal draft

After you get the first response, use **one refinement** — "Shorter," "More formal," or "Add X."
Notice how fast you reach a usable result.

Your pocket reference — the whole lesson on one slide

Ingredient / Habit	What it does	Quick example
Role	Sets perspective & vocabulary	<i>"You are a plain-English lawyer."</i>
Context	Gives necessary background	<i>"I'm reviewing a rental contract for the first time."</i>
Task	Specifies exactly what you want	<i>"Flag any clauses that could cost me money."</i>
Format	Controls output shape	<i>"Bullet list, plain English, under 150 words."</i>
Negation	Closes bad paths	<i>"Don't use jargon. Don't make anything up."</i>
Show example	Pattern-matches style exactly	<i>"Write it like this: [paste example]."</i>
Iterate	Steers without restarting	<i>"Shorter." / "More formal."</i>

Ingredient / Habit	What it does	Quick example
Think step by step	Improves reasoning accuracy	<i>Add to any logic / maths / planning prompt.</i>

When a better prompt isn't the answer

Prompting is powerful, but it can't give a model information it doesn't have.



Current information

The model's knowledge stopped at a cutoff date. No prompt can conjure today's news. **Fix:** use a tool with web search (Perplexity, Claude with web mode).



Precise arithmetic

LLMs reason about maths better than they calculate it. For number-critical work, use a model with a code interpreter or a real calculator. **Always verify numbers.**



Your private documents

The model can't see files it hasn't been given. **Fix:** paste the relevant section, or use a tool with file upload / document Q&A.

Bottom line: prompting is the steering wheel. But you still need to know what the car can and can't do. Know your tool's limits.

What you now know — and what's next

This lesson in one paragraph

The four-part recipe (Role + Context + Task + Format) is your starting point. The five habits (show an example; ask for the format; say what not to do; iterate; let it think) sharpen it. The four mistakes (vague, overloaded, unchecked, outdated) are what to avoid. Treat the AI like a smart intern who needs a clear brief, not a search engine you query with three words.

The skill that compounds: every other AI task in this course gets faster once prompting clicks. This lesson earns its “most important” status.

Theory page →

Interactive prompt builder · Vague vs. specific toggle · Animated iteration demo · Mistakes checklist. Build and test your own prompts right there.

[theory.html →](#)

Homework →

Write real four-part prompts for three scenarios from your own life, run them, and reflect on what worked. The only way to get good is to practise.

[homework.html →](#)