CURATED STREAM 5

Prompting Large Language Models

Crafting prompts to get the most out of AI

How to get the most out of an AI model

This curated stream is about learning to craft meaningful instructions — prompts — so an AI model works for you. Remember: you are in charge.

1. **Common definitions**
2. **Structuring a powerful prompt**
3. **Guide the model to reason step-by-step**
4. **There is no perfect prompt — iterate**
5. **Changing topic — let the model know**
6. **Summary — don't assume**

Section 1

Common definitions

Start with shared language: large language models and prompts

What is a large language model?

A large language model (LLM) is a type of AI program designed to understand and generate human-like text.

It learns from vast amounts of text, then predicts and generates coherent sentences and paragraphs on a huge range of topics. You may also hear it called generative AI.

Working with an LLM

FRONTIER LLMs IN 2026

Claude Opus 4.8 (Anthropic), **GPT-5.6** (OpenAI), **Gemini 3.5 Flash** (Google), **Llama 4** (Meta). Most accept million-token prompts and read text, images, and even video.

What is a prompt?

An LLM understands language, so you 'prompt' it — you ask the computer to solve a problem in plain words.

Q

Your prompt

Tells the model how to behave and what to research or write.

i

Your context

Background information, examples, or constraints you provide.

A

The response

Built from your prompt — the more specific you are, the better the result.

Think of it as a smarter search box. Search engines match keywords; an LLM expects enough context to give you a genuinely useful answer.

Section 2

Structuring your prompt

Specific prompts produce focused, useful results

General vs. specific prompts

Compare these two prompts on the same topic. Notice how much more the model has to work with.

GENERAL

“What is photosynthesis?”

You'll get a textbook-style definition. Probably useful, but not tailored to who you are or what you need.

SPECIFIC

“I am a 12-year-old student preparing a short presentation about how photosynthesis works. The talk must be under 8 minutes. Identify 3 benefits and include simple diagrams or pictures.”

Now the model knows the audience, the length, the deliverable, and the supporting visuals — and tailors everything to that.

The CIA formula for powerful prompts

Three building blocks every strong prompt should include:

C

Context

Briefly introduce the situation, role, or background the model should consider.

I

Instruction

Clearly state what you want the model to do or produce. Be as specific as possible.

A

Additional

Optional: preferences, constraints, length, tone, audience, format, or style.

Tip: Even a one-line context can change the whole reply.

CIA in action: a worked example

EXAMPLE PROMPT

C **Context**
I am a 12-year-old student writing a short presentation to teach my class how photosynthesis works.

I **Instruction**
Create a short presentation that lasts no longer than 8 minutes.

A **Additional**
Identify 3 benefits of photosynthesis and include simple diagrams or pictures.

PRACTICE

1. Pick an LLM you already use (ChatGPT, Claude, Gemini, Copilot).

2. Write a prompt using the CIA structure on a topic you care about.

3. Try the same prompt with less context, then more context. Compare the results.

Section 3

Guide the model to reason

Asking for steps unlocks better answers

Ask for steps, not just an answer

Modern LLMs are strongest when you break the problem into ordered steps. The model focuses on one thing at a time instead of trying to solve everything at once.

EXAMPLE — TRIP ITINERARY PROMPT

1. Identify 5 tourist destinations in northern France.
2. For each destination, recommend a hotel close by.
3. Provide a map for each destination showing the tourist activity and hotel.

Try unusual formats too: ask for the itinerary as a table, a poem, a map description, or a packing checklist. Treat the output format as part of the prompt.

Practice: chain the steps

PRACTICE

Write a prompt where the result of step one drives the scope of step two.

For example: choose 5 destinations, then recommend a hotel for each. The hotel recommendations depend on the destinations the model picks.

Try one of these instead:

- Plan a 3-day Nelson-Tasman day-trip itinerary.
- Build a 5-meal weekly menu, then a shopping list from those meals.
- Design a 4-week beginner workout, then the warm-up for each session.

WHY IT WORKS

Focuses the model.
Each step has one job.

Forces you to be clear.
Listing steps makes you visualise the answer first.

Easier to fix.
If a step goes wrong, you can correct it without restarting.

Section 4

There is no perfect prompt

Improve results by iterating — it's a conversation, not a search

Iterating with follow-up prompts

Most LLMs remember the conversation. You don't have to retype the original prompt — just ask for what's missing or what to change.

Your prompt

"Plan 5 destinations in northern France with a hotel for each."

Follow-up 1

"Include the hotel cost and the star rating of each hotel."

Follow-up 2

"Redo that as a table with cost and star rating columns."

LIKE A CONVERSATION

Each follow-up builds on what came before. The model remembers the destinations, the hotels, the constraints — you only need to add what changed.

Iterate, iterate, iterate

Don't settle for a sub-optimal result. Push the model — that's how you find a good one.

Ask for less jargon

"Rewrite that without technical terms" — or for a child audience, a board paper, etc.

Change the tone

"Make it warmer, like talking to a colleague." Switch between formal, casual, persuasive.

Add a constraint

"Keep it under 150 words." "In NZ English." "As a bulleted list with 5 items."

Introduce a new aspect

"Now exclude castles." "Add the carbon footprint of each trip." Steer mid-conversation.

Tip: You don't need to retype the original prompt. Just ask for the change you want.

Section 5

Changing topic

Tell the model when you switch direction — start a new chat

Changing topic? Start a new chat

When you change topic mid-conversation, the model keeps the old context — which can give odd or confused answers.

Example: a long thread about northern France, suddenly switched to "recommend a restaurant in Nelson?" The model may still be thinking 'France'.

START A NEW CHAT

Every LLM has a "New chat" button — usually top-left or in the sidebar.

ChatGPT, Claude, Gemini, Copilot all work the same way — the new chat forgets the prior context, which is what you want.

REOPEN A PAST CHAT

Old conversations are saved under Recent or History.

Click on a prior conversation to pick up where you left off — useful when you want to keep iterating on something from last week.

Practice: new chat and history

PRACTICE

1

Start a new chat and run a fresh prompt

Open your LLM, click New chat, and enter a prompt on any topic. Note that the model has no memory of previous chats.

2

Reopen a prior conversation and iterate

Find an earlier chat under Recent or History. Click it to reload. Continue refining the results by adding follow-up prompts.

Section 6

Summary — don't assume

Be clear, be specific, and reap the benefits

The four habits of strong prompting

The model can answer a one-word question, but it doesn't know your reading level, your audience, or whether you want a sentence or a chapter. Tell it.

1

Follow CIA

Context, Instruction, Additional considerations.

2

Ask for reasoning

Break the request into ordered steps when the task is complex.

3

Iterate

Treat it like a conversation — add, refine, change tone.

4

Start fresh

New topic? New chat. Don't let old context contaminate the answer.

Be clear, be specific — and you'll reap the benefits.

Now go and experiment.

There is no perfect prompt — only the next iteration.

Nelson AI Sandbox · aisandbox.org.nz