Crafting Your Prompts

How to craft prompts to use in Artificial Intelligence (AI)

How can I get the most out of an AI Model?

The purpose of this curated stream is to spend time learning and practising to construct meaningful instructions (prompts) to the computer, to have it work for you to retrieve current knowledge, or create something unique.
Remember - you are in charge.

Sections in this Tutorial:

Section 1: Let's start with ensuring we have common definitions.

Section 2: Structuring your prompt to make it powerful.

Section 3: Guide the model to think through its answer

Section 4: There is no perfect prompt - improving the results by experimenting.

Section 5: Changing the topic - let the LLM know!

Section 6: Summary: Don't assume!

Section 1:

Let's start with ensuring we have common definitions.

Let's start with ensuring we have common definitions.

What is a large language model (LLM)?

A large language model is a type of artificial intelligence program designed to understand and generate human-like text. It learns from vast amounts of data and uses that information to predict and generate coherent sentences, paragraphs on a wide range of topics. It is also referred to as generative AI.

What is a prompt?

LLM software understands language, hence you are 'prompted' to ask the computer to solve a problem. Your prompt instructs the model how to behave and how to research and respond to your request. You are likely familiar with search engines which prompt you for what you wish to search for. LLMs are expecting sufficient information to be able to give you a valuable result.

This tutorial will outline the tricks and tips for getting the most from your LLM engagement.

Section 2:

Structuring your prompt to make it powerful.

Structuring your prompt to make it powerful.

A prompt can be a single question, or a request, or it can be a specific request for information, including necessary background context.

Type	Prompt as keyed
General	What is photosynthesis?
Specific	I am a student of age 12. I am writing a short presentation to teach my fellow students about how photosynthesis works. I would like to create a short presentation to last no longer than 8 minutes. Identify 3 benefits of photosynthesis and include some simple diagrams or pictures.

CIA - Context, Instruction, Additional Considerations

Using the above examples, you know that a specific prompt will yield focussed results, more suited to your needs. To become powerful with prompts, ensure you follow the CIA formula when constructing your requests.
- [Context/Background Information]: Briefly introduce the situation or topic the LLM should consider.
- [Instruction/Task]: Clearly state what you want the LLM to do or generate. Be as specific as possible.
- [Additional Considerations: Optional]: Specify any preferences, constraints, or desired styles for the LLM's results.
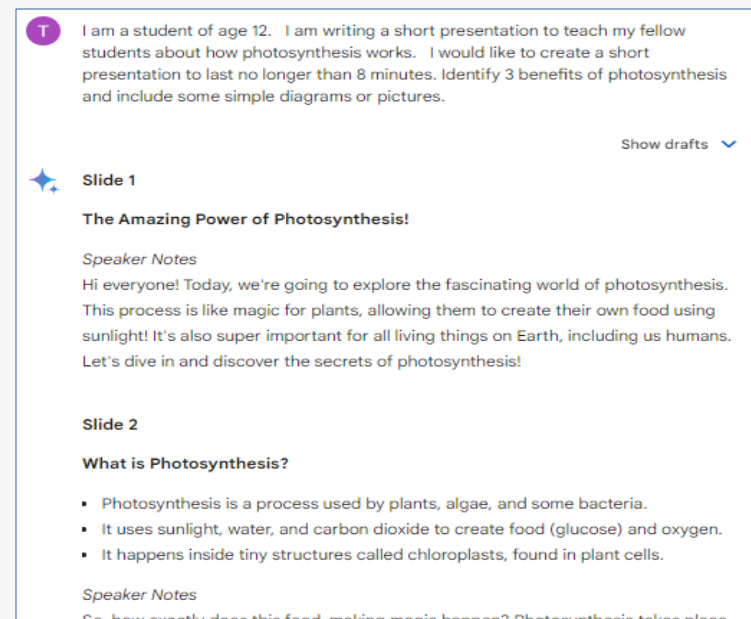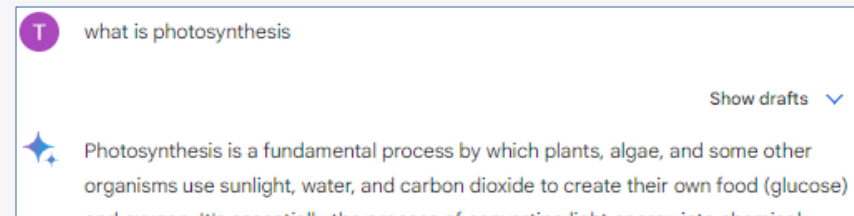
An example of this would be:

● Context/Background Information: I am a student of age 12. I am writing a short presentation to teach my fellow students about how photosynthesis works

● Instruction/Task: I would like to create a short presentation to last no longer than 8 minutes.

● Additional Considerations: Identify 3 benefits of photosynthesis and include some simple diagrams or pictures
Practice:

1. In the LLM of your choice, create prompts using the CIA method.

2. Try prompts with less information, or more information, to see how this impacts results.

The following examples are based on Gemini, however each LLM works similarly. Open up whatever is easiest for you – if you are on an AI Sandbox laptop, there are options on the favourites bar.

Section 3:

Guide the model to think through its answer

Guide the model to think through its answer.

There are benefits in guiding the model to carry out its work step by step.

For instance, if you were wanting to get an itinerary for a trip, you might construct a prompt to instruct the LLM to do the following:

- Identify 5 tourist destinations in northern France
- Then for each destination, recommend a hotel close by
- Provide a map of the destination showing the tourist activity and hotel

The model can then focus, rather than trying to resolve all parts of the problem in a random sequence.

Including specific steps in your prompt also allows you to think about the problems you are looking to resolve, and visualise the format of the answers.

Your prompt could be written as a table, a poem (as a surprise trip for your loved ones), or maps and directions.

Experiment with different types of output, assuming all things are possible. One day they will be.

Practice:
Create a prompt which instructs the LLM to carry out a range of tasks with the results of the first task driving the scope for the second task (as in the 5 tourist destinations chosen leading to the hotel recommendation).

Section 4:

There is no perfect prompt - improving the results by experimenting.

There is no perfect prompt - improving the results by experimenting.

It is unlikely you will always get the results you expect. Where you do, then you have likely written a specific prompt. Sometimes, you may wish to get unexpected results, and you can even ask for that !

Most LLMs will <u>remember your last request</u>.

Hence, you can enter an additional prompts following the LLM results. This is essentially like a human conversation.

Using the tourism example above, I might have forgotten to get information on the cost of the hotel, and the star rating. Hence I could just type in another prompt (straight after the prior results) and the LLM will build on the answer it gave.

- "Include the hotel cost and the star rating of the hotel"

Possible second prompts could be:

- "Please redo the results in table form, including the hotel cost and hotel star rating"

Iterate, Iterate, Iterate!

Don't settle for a sub-optimal result. Ask for the extra information you require, or ask for it to be rewritten with less technical speak, or in a specific tone (business, child like) etc.

Tip: You do not need to reenter the whole prompt again. Assume the LLM knows what you asked and then just request the changes you require. "**Add something to the results** "

Practice:

1. Construct a prompt ensuring the CIA method is followed. Execute this in your LLM.
2. Iterate by asking the LLM to change the answer to include more of this, less of that, or focus on a specific part of the request.
3. Continue iterating to refine the results
4. Introduce a new aspect of the topic (e.g. using the tourist theme, maybe inform the LLM that you don't want to see castles).
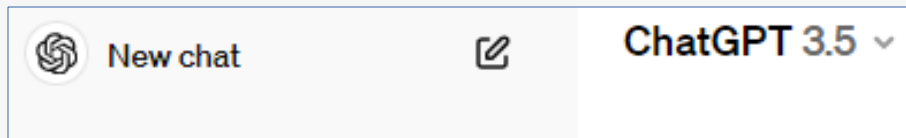
Section 5:

Changing the topic - let the LLM know!
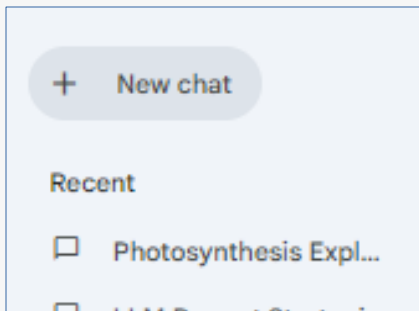
Changing the topic - let the LLM know!

If you have ever changed topic in a conversation without informing people you are doing it, it can have funny or confused circumstances. The LLM is the same.

A discussion about Northern France, and changing to asking a question about restaurants in Nelson may give poor outcomes. Most LLMs have an option to start a new topic, and it forgets the prior context.

- The **"New chat"** buttons shown below are for Chat GPT and Gemini, but every LLM will have one somewhere.

Tip: Your prior prompts (conversation) will in most LLMs be saved under recent or history. You can retrieve that conversation and carry on where you left off later.

Practice:
1. Start a new chat. Enter a prompt and see the result in your LLM.

2. Once done, find your prior chat in the recent (or history), click on that recent topic to reload where you left off. Continue to refine the results by iterating the prompt.

Section 6:

Summary: Don't assume!

Summary: Don't assume!

Whilst the LLM can answer a simple question, it doesn't know whether you want 1 word, 1 line, or chapter and verse. It doesn't know your reading level or experience in the topic.

In order to get the best out of LLM prompts, remember:

1. Powerful prompts follow the CIA format (Context, Instruction, Actions).

2. Guide the model to think through the problem

3. Improve the results by writing new prompts - iterate

4. Start a new chat to avoid topics getting muddled

Lastly, take the time to be clear and specific and your will reap the benefits.