

Mastering Prompt Engineering

Acknowledgement of Country

The University of Queensland (UQ) acknowledges the Traditional Owners and their custodianship of the lands on which we meet.

We pay our respects to their Ancestors and their descendants, who continue cultural and spiritual connections to Country.

We recognise their valuable contributions to Australian and global society.

*The Brisbane River pattern from A Guidance Through Time
by Casey Coolwell and Kyra Mancktelow.*



What we will cover today....



Intro to Large Language Models



Simple Prompt Engineering + Demo



Advanced Prompting Techniques



Data safety and Limitations



Support services and Further training

Introduction to LLMs and Microsoft Copilot

Large Language Models (LLMs):

Large language models (LLMs) are a category of **Artificial Intelligence** (AI) model trained on immense amounts of data, making them capable of understanding and generating natural language and other types of content while performing a wide range of tasks. Examples of popular LLMs include Gemini 2.5 Flash (Google), GPT-4o (OpenAI), Claude 4 (Anthropic) and Copilot (Microsoft, based on GPT-5).

Microsoft Copilot Chat:

You can get free access to Copilot!

Benefits to using Copilot

- Copilot operates on a closed system; this means your data is protected and will not be absorbed as training data for future models.
- Copilot uses models similar to ChatGPT's models and gives you access the same fast and intelligent responses without the paywall of other pro models.

Microsoft Copilot Guide.

More general Copilot information.

Common and acceptable uses of Gen AI

Deciding what is appropriate is up to you but may depend on what is permitted in your course and assessment

- Brainstorming research ideas
- Explaining concepts, clustering themes
- Use it for copy-editing, e.g. reformat a sentence/paragraph, give suggestions on tone, conciseness, writing improvements
- Creating outlines for a letters, presentations, scripts
- Writing code, reviewing code
- Critique and role-playing tool to increase productivity, innovation & learning
- Summarising text input
- Simplifying concepts
- Self-learning – create exam questions/flash cards
- Design feedback and review, e.g. for instructions, methodology, research questions, etc. (don't input sensitive data)
- Synonym suggestions (e.g. for literature searching)
- Translations

What is Prompt Engineering?

Prompt engineering is the process of writing effective instructions for a model, such that it consistently generates content that meets your requirement (OpenAI, n.d). There are different types of prompting techniques for different tasks.

Basic Principles:

Clarity & Specificity

- Avoid vague overly general or broad instructions.
- Clearly state your objective and the context of your question or instructions.

Context Framing

- Provide the right background information to help the AI focus on a specific task.

Structured Instructions

- Use explicit constraints (word count, format, style).

Iterative Refinement

- Refine prompts based on AI output until satisfied.

Quick Example:

You are an undergrad law student and want help with understanding an assignment analysing Rule of Law in Australia.

Weak prompt:

“Tell me about rule of law in Australia.”

Effective Prompt:

“You are a university tutor for undergraduate level law. Analyse the assignment question below and explain what I am being asked to do.

Explain how I should approach this assignment and include a timeline in dot point format for when I should complete each task.

Here is the assignment question

“To what extent is the Rule of Law supported by the Australian Legal System?”

Prompt Engineering in Practice:

Think R.I.C.E.!

R - Role: Assign the AI with a persona or role to adopt within the conversation. This role will influence the tone, style, and knowledge base of the AI.

Example: *"Act as an academic researcher specialising in architecture..." OR "You are an expert in python coding..."*

I - Instruction: The primary task or action you want the AI to perform.

Example: *"...how could I improve the environmental impact of my project..." OR "Analyse this code to identify areas for improvement..."*

C - Context: Provide all necessary background information, parameters, or data the AI needs to complete the instruction.

Example: *"...focus on its implications for sustainable development in urban environments." OR "Provide me with multiple different ways to resolve the coding issues."*

E - Expectation: Clearly define the desired format, length, style, or specific inclusions/exclusions for the output.

Example: *"...provide me with a dot point list and a detailed analysis, use formal academic language, and avoid jargon where possible." OR "Provide me with a list of solutions and a detailed breakdown of the options."*

Prompting Techniques

Simple Prompt or Zero Shot Prompting

- Zero-shot prompting means that the prompt used to interact with the model won't contain examples or demonstrations. The zero-shot prompt directly instructs the model to perform a task without any additional examples to steer it.
- **The demos we completed were Zero shot prompts as we input data without providing it with examples of tasks to complete or orient itself.**

Decomposition Prompting

Decomposition prompting is a powerful approach that breaks down complex problems into simpler, more manageable sub-tasks. This technique is inspired by a fundamental human problem-solving strategy and has shown remarkable success in enhancing AI performance without requiring larger models.

Example: Creating a Plan for a Research Project

This technique is excellent for building a structured plan *before* you start writing. It helps you think through the research process methodically.

Goal: Create a comprehensive research plan to investigate the use of AI in Higher Education.

Zero shot Prompt: "Create a research plan about AI in education."

Note: this type of prompting is especially useful for complex mathematical or engineering tasks that require multiple steps. For these tasks consider LLMs with **large context windows**.

Use a Sequence of Decomposed Prompts.

- **Prompt 1: Identify Key Themes :** "I am planning a research project on the role of AI in Higher Education. Identify 5-7 distinct key themes or sub-topics that are currently being discussed in academic literature on this subject."
- **Prompt 2: Formulate Research Questions :** Based on the theme of 'AI and Academic Integrity' that you identified, generate three potential research questions that would be suitable for an undergraduate project. Ensure the questions are focused and investigable."
- **Prompt 3: Outline a Methodology :** "For the research question: '[Insert one of the generated questions]', suggest a suitable research methodology. Describe the key steps involved, including the type of data to be collected and a possible method for analysis."
- **Prompt 4: Identify Potential Challenges:** "What are three potential challenges or limitations I might face when conducting research based on the methodology you just outlined?"

Socratic Tutor Prompting

You can use AIs to prompt you using them as a Socratic Tutor. This is a process of asking questions to clarify understanding.

You are an expert agricultural sciences tutor.

I want to study the following topics.

- Soil science

- Weeds

Ask me questions that integrate ideas from these topics. After each question that you ask, wait for my response. Then you should respond but do not tell me the answer. Instead, ask me questions and use the Socratic questioning approach to help me arrive at the answer myself.

Chang, E. Y. (2023). Prompting large language models with the socratic method. 2023 IEEE 13th Annual Computing and Communication Workshop and Conference (CCWC)

An example prompt:

<https://thatryanp.medium.com/my-go-to-prompt-for-chatgpt-socratic-coach-7bf0dd2c01ec>

Few Shot Prompting

Few-shot prompting provides examples of how a task is to be solved. It is used to provide guidance for the required output. An easy example is a **sentiment analysis**.

Input:

Here are a few examples of movie reviews and what sentiment they have:

- **Review:** “This movie was absolutely fantastic! I loved every minute of it.”
- **Sentiment:** Positive
- **Review:** “The plot was confusing and the acting was terrible. What a waste of time.”
- **Sentiment:** Negative
- **Review:** “It was okay. Not great, not terrible, just average.”
- **Sentiment:** Neutral

Now, classify the sentiment of this review: “The special effects were amazing, but the story was predictable and the characters were flat.”

Output:

- **Review:** “The special effects were amazing, but the story was predictable and the characters were flat.”
- **Sentiment:** Mixed.
- **Reasoning:** The review contains both positive and negative aspects. While it praises the special effects, it criticizes the story and characters. This balance of contrasting opinions classifies it as a mixed sentiment.

[Prompted 29 November 2024, OpenAI (2024), ChatGPT 4.0 response on topic]

Chain of Thought Prompting

Introduced in [Wei et al. \(2022\)](#), **chain-of-thought (CoT)** prompting enables complex reasoning capabilities through intermediate reasoning steps. You can combine it with few-shot prompting to get better results on more complex tasks that require reasoning before responding.

Input:

Yes or no: would a pear sink in water? **Provide a reasoning for your answer in logical steps before providing an answer.**

Output:

Density determines floating or sinking:

An object floats in water if its density is less than that of water (approximately 1 g/cm^3). If the object's density is greater, it sinks.

Examining a pear's composition:

A pear is composed of water, sugars, fibres, and air pockets within its cellular structure. The presence of air pockets lowers its overall density compared to water.

Considering variations in pears:

While most pears have a density less than water, factors like ripeness, water content, or damage can affect buoyancy. Overripe or denser pears might sink.

Answer: No, a pear would not sink in water under normal conditions. It generally floats because its density is slightly less than water.

[Prompted 29 November 2024, OpenAI (2024), ChatGPT 4.0 response on topic]

Meta Prompting

Meta-prompting is the process of using an LLM to generate or improve a prompt for another task. Instead of manually refining a prompt, you ask an LLM to act as a prompt engineer itself, helping you create more effective instructions. This is often an iterative process where you use an LLM to generate candidate prompts, score them based on performance, and then create variations of the best ones.

Example: A Simple Meta-Prompt

You can use a straightforward template to ask an LLM to refine an existing prompt.

Goal: To improve a prompt designed for summarizing open access research articles.

Input:

Improve the following prompt: "Summarize this article on whales."

Output:

Option 1: Simple and Direct Improvement

This version adds a target audience and a length constraint.

"Summarize the following article about whales for a high school student. The summary should be a single paragraph and no more than 150 words. Focus on the article's main argument and key findings."

Option 2: Structured for Deeper Analysis

This version is ideal for academic or research articles, as it asks for specific components and a structured format.

You are a research assistant. Read the following article and provide a structured summary.

- **Audience:** University undergraduate student.
- **Format:** Use the following headings with bullet points for each section.
- **Research Question(s):** What was the primary question the researchers were trying to answer?
- **Methodology:** Briefly describe how the research was conducted. (e.g., satellite tagging, acoustic monitoring, literature review).
- **Key Findings:** List the 3-5 most important discoveries or results. Use quantifiable data from the article if available.
- **Conclusion & Implications:** What was the final conclusion, and why does it matter for whale conservation, marine biology, or our understanding of the ocean?

Self-consistency prompting

Self-Consistency is an advanced prompting technique that improves the reliability of answers for **complex reasoning tasks**. It is based on the logic that there are often several different ways to reason through a problem that will arrive at the correct answer. Instead of accepting the first answer from a **Chain-of-Thought prompt**, this method generates **multiple different reasoning paths** and then **chooses the most consistent answer among them**.

Instead of just generating one response, **Self-Consistency involves a multi-step process**:

Prompt Multiple Times: Use a Chain-of-Thought (CoT) prompt for a single problem but run it multiple times.

Generate Diverse Reasoning: Critically, use a non-zero **temperature** (e.g., 0.5 or higher) for each run. This introduces randomness, encouraging the model to generate a variety of different reasoning paths.

Aggregate and Vote: Collect all the final answers from the different reasoning paths. The final answer is then selected by taking the majority vote over all the generated responses.

Example: A Multi-Step Reasoning Problem

Problem: "When I was 6, my sister was half my age. Now, I'm 70. How old is my sister?"

Path 1: "When you were 6, your sister was half your age, so she was 3. This means you are 3 years older than her. Now you are 70, so your sister is $70 - 3 = 67$. Answer: 67"

Path 2: "The age difference between you and your sister is $6 - (6 / 2) = 3$ years. If you are 70 now, your sister's age is $70 - 3 = 67$. Answer: 67"

Path 3: "If your sister was half your age, she is still half your age now. Half of 70 is 35. Answer: 35"

Conclusion: By taking the majority vote, you can confidently select 67 as the correct answer, discarding the faulty reasoning path.

Prompt Round Up

This paper contains a round up of all commonly used prompting techniques ranging from basic to more advanced. It includes information on the most highly cited techniques within academic research as well as general definitions and explanations of techniques.

The Prompt Report: A Systematic Survey of Prompt Engineering Techniques

Sander Schulhoff^{1,2*} Michael Ilie^{1*} Nishant Balepur¹ Konstantine Kahadze¹

Amanda Liu¹ Chenglei Si⁴ Yinheng Li⁵ Aayush Gupta¹ HyoJung Han¹ Sevien Schulhoff¹

Pranav Sandeep Dulepet¹ Saurav Vidyadhara¹ Dayeon Ki¹ Sweta Agrawal¹² Chau Pham¹³

Gerson Kroiz Feileen Li¹ Hudson Tao¹ Ashay Srivastava¹ Hevander Da Costa¹ Saloni Gupta¹

Megan L. Rogers⁸ Inna Goncarenco⁹ Giuseppe Sarli^{9,10} Igor Galynker¹¹

Denis Peskoff⁷ Marine Carpuat¹ Jules White⁶ Shyamal Anadkat³ Alexander Hoyle¹ Philip Resnik¹

¹ University of Maryland ² Learn Prompting ³ OpenAI ⁴ Stanford ⁵ Microsoft ⁶ Vanderbilt ⁷ Princeton

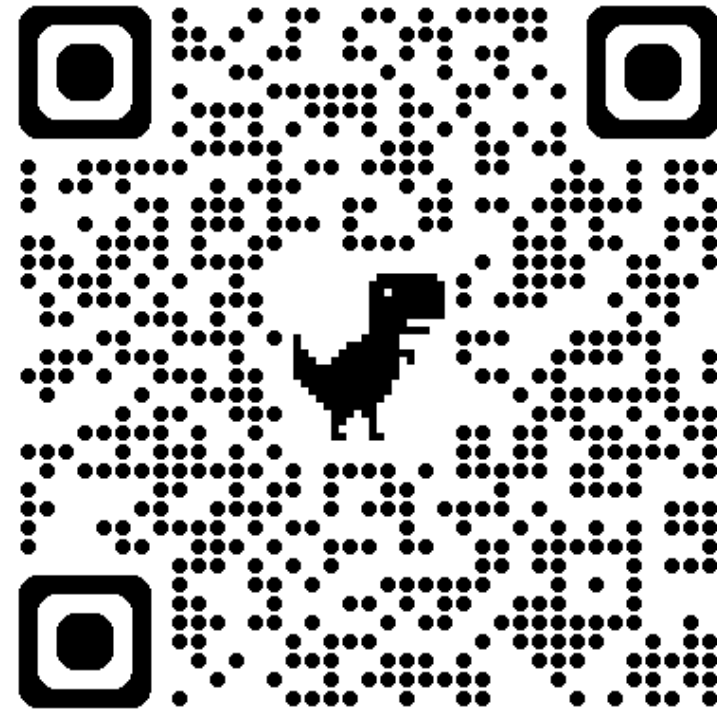
⁸ Texas State University ⁹ Icahn School of Medicine ¹⁰ ASST Brianza

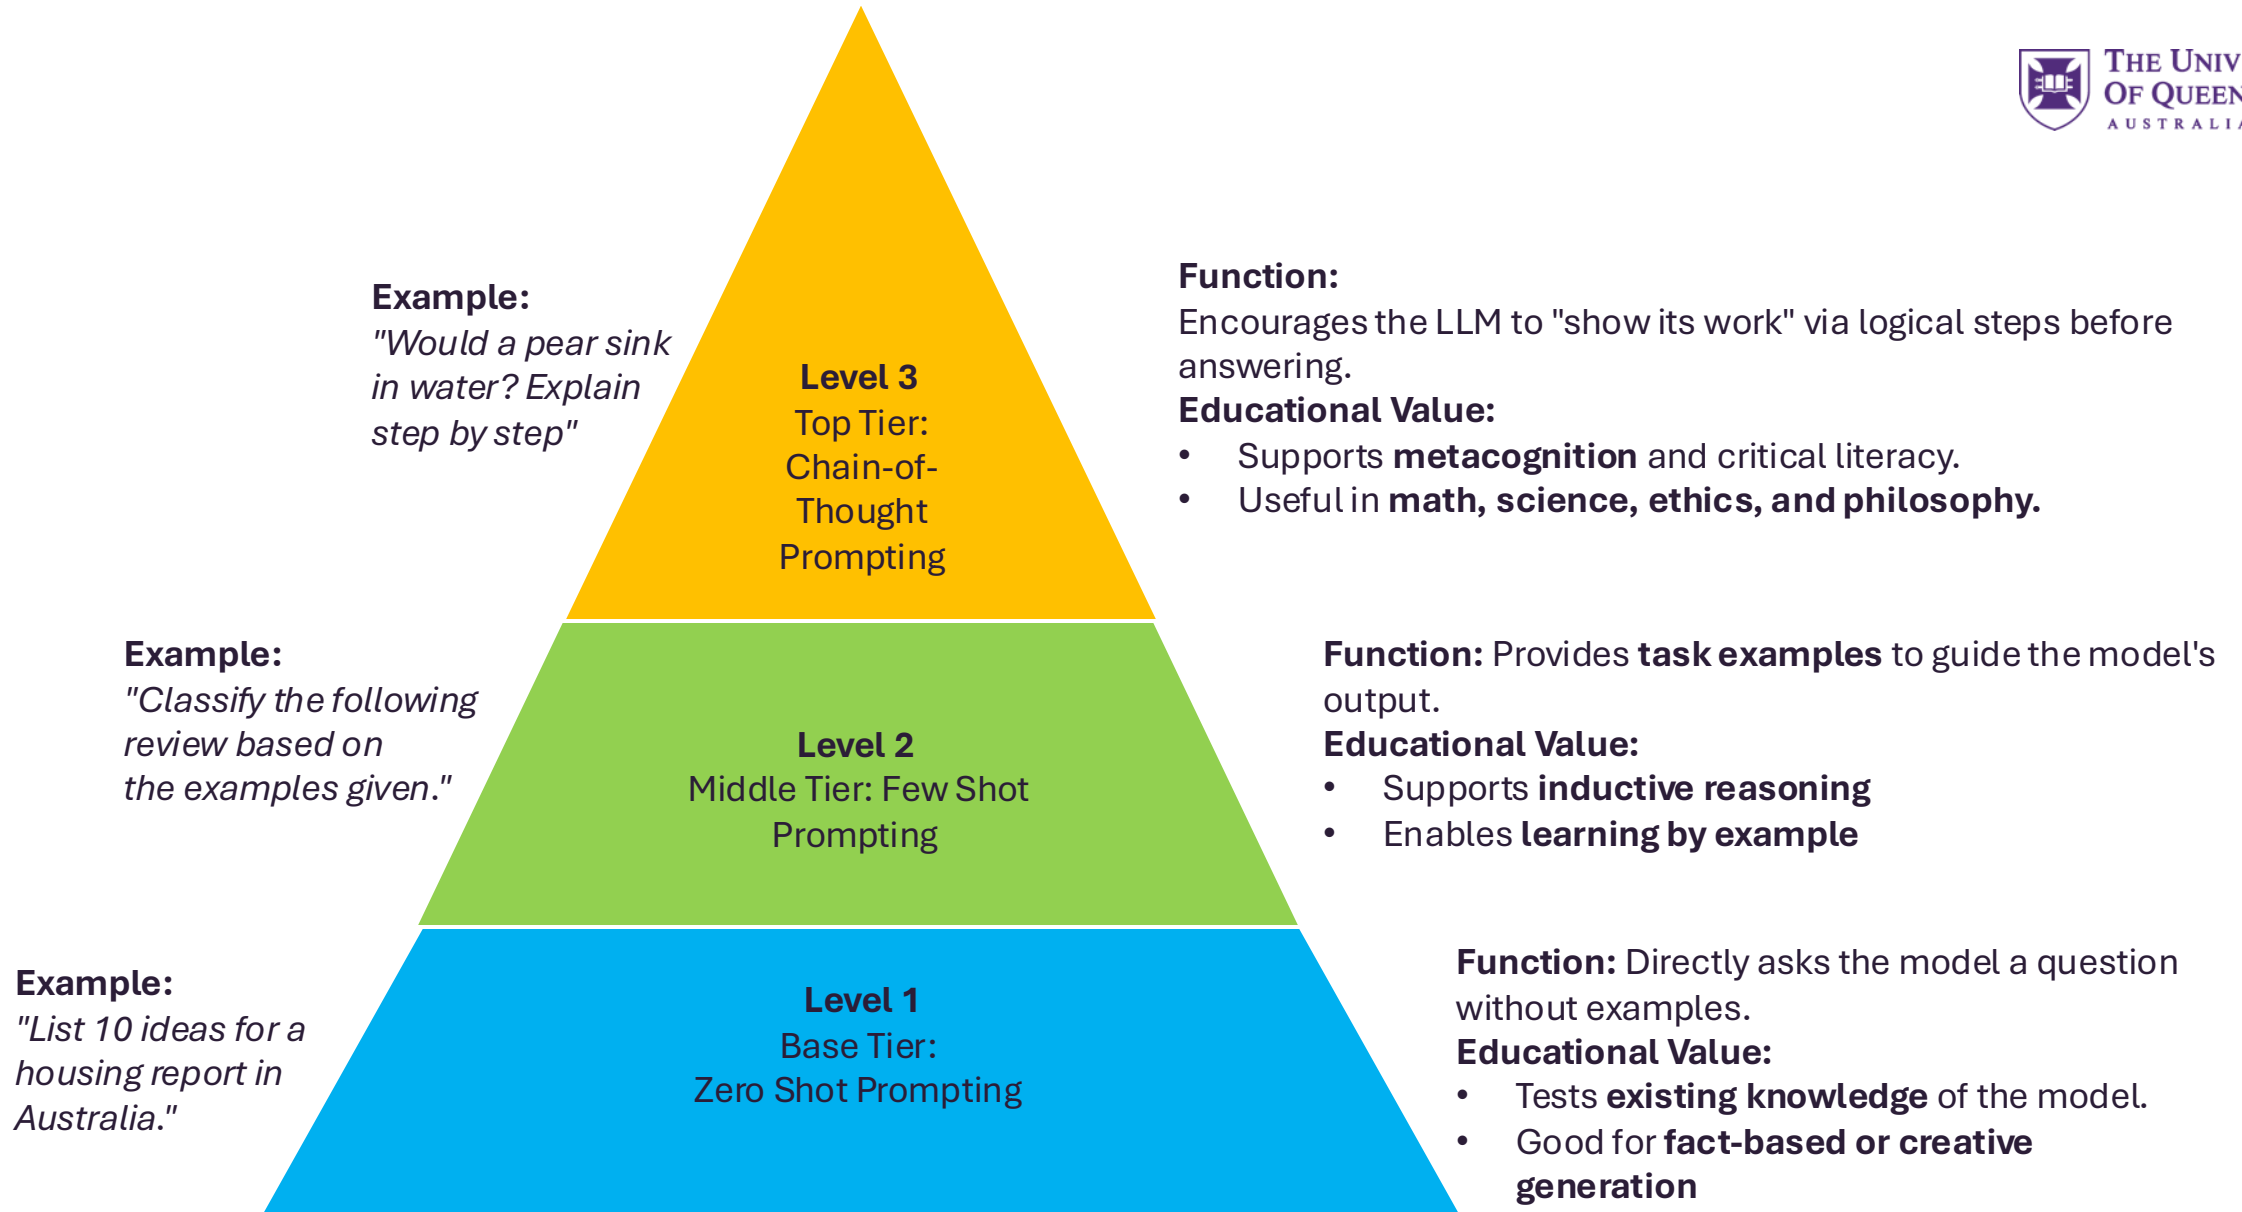
¹¹ Mount Sinai Beth Israel ¹² Instituto de Telecomunicações ¹³ University of Massachusetts Amherst

sschulho@umd.edu milie@umd.edu resnik@umd.edu

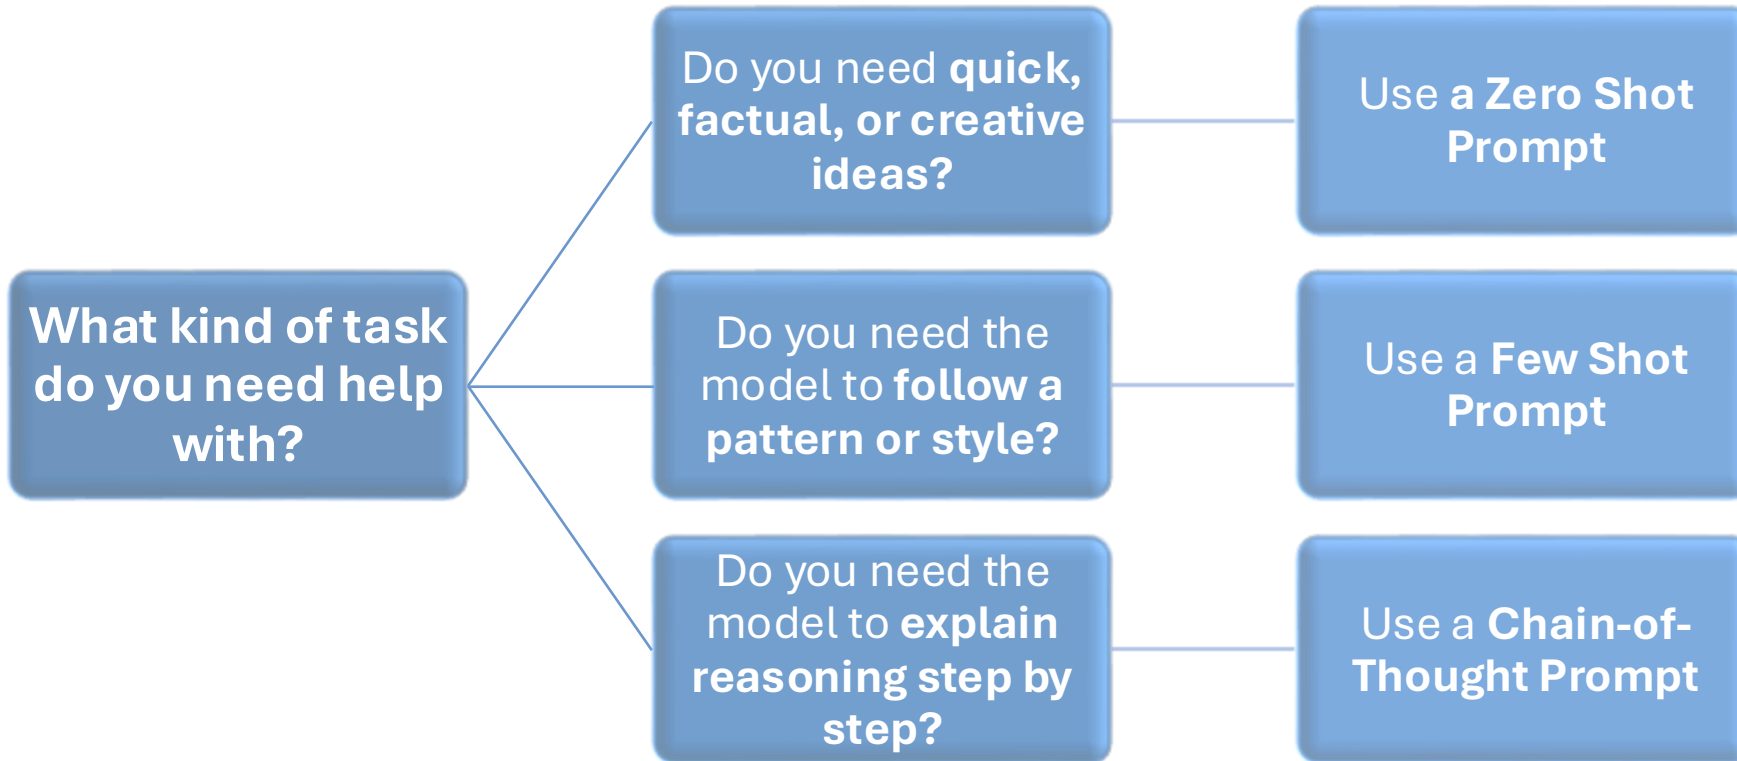
Abstract

Generative Artificial Intelligence (GenAI) systems are increasingly being deployed across diverse industries and research domains. Developers and end-users interact with these systems through the use of prompting and prompt engineering. Although prompt engineering is a widely adopted and extensively researched area, it suffers from conflicting terminology and a fragmented ontological understanding of what constitutes an effective prompt due to its relatively recent emergence. We establish a structured understanding of prompt engineering by assembling a taxonomy of prompting techniques and analyzing their applications. We present a detailed vocabulary of 33 vocabulary terms, a taxonomy of 58 LLM prompting techniques, and 40 techniques for other modalities. Additionally, we provide best practices and guidelines for prompt engineering, including advice for prompting engineering ChatGPT and other state-of-the-art (SOTA) LLMs. We further present a meta-analysis of the entire literature on natural language prefix-prompting. As a culmination of these efforts, this paper presents the most comprehensive survey on prompt engineering to date.





Prompt Selection



Direct question or Task

- No examples provided
 - Fast and general
 - Good for brainstorming, summaries, etc.
- Example: "list 5 causes of climate change"

Provide 2 to 3 Examples

- Helps model mimic structure
 - Useful for formatting, classification, tone.
- Example: "here are examples of positive and negative reviews."

Encourages Transparency and Logic

- Breaks down complex tasks
 - Best for problem-solving and justification
- Example: "Would a pear float in water? Explain why."

Tokenisation and Context Windows

Tokenisation:

LLMs utilise something called tokenisation to process the words within a prompt.

In this process a **word** is broken down into **multiple parts** called **tokens** making them easier for an LLM to process and understand.

Context Window:

A **context window** refers to the **maximum amount of information** (measured in tokens) an **LLM can process** to generate a response. Using a model with a higher context window can be helpful when processing large data sets or when you need the LLM to remember details within a long chat history.

Context Window Comparison

AI Model/Provider	Typical Context Window (Tokens)	Notes on Tokenization Efficiency
Google Gemini 2.5 Flash	Up to 1 million (experimentally up to 10 million)	Known for its very large context window, enabling processing of massive documents/codebases.
OpenAI GPT-4.1	1 million	Very capable model with a substantial context window; often more cost-effective than previous GPT-4 versions. For research such as GPT-4.5.
OpenAI GPT-o3	200,000	
OpenAI GPT-o4	128,000	A highly capable model, widely used
Anthropic Claude 4 Opus	200,000	A very powerful model with a large context window, often used for complex reasoning.
Anthropic Claude 4 Sonnet	200,000	Balances power and affordability.
Mistral Large	128,000	Strong European contender with good performance.

Image prompting

Example Prompt Structure : A raw photo of [subject] [doing something interesting] during [time of day], taken with a [type of camera], using [type of lens] with cinematic lighting -- ar 16:9

Input & Prompt : Generate raw photo of a cat shelving books in a library, taken with a DSLR camera, using a wide angled lens with cinematic lighting -- ar 16:9

Output using Microsoft Copilot:

[Prompted 5 March 2026, Microsoft (2026), "Auto" setting]



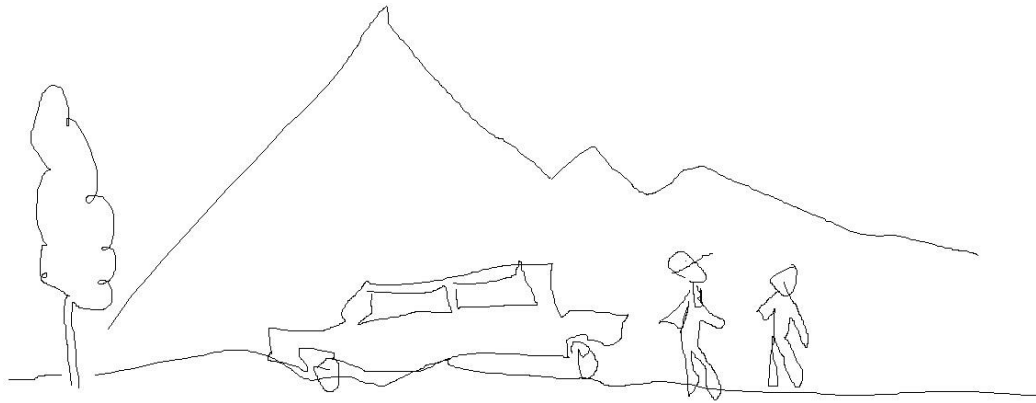
Using an image as a prompt

Output using Copilot

Image prompting in MS copilot

Input & Prompt

Create a photo from this image



Common Pitfalls & Limitations

Pitfalls in Prompt Engineering:

- **Ambiguity and Lack of Context:** Prompts that are vague or lack specific background information can lead to irrelevant or generic AI responses.
- **Overcomplicated Instructions:** Giving the AI too many complex tasks in a single prompt can cause confusion, leading it to ignore parts of the instruction.

LLM Limitations:

- **Hallucination and Misinformation:** AI models can generate incorrect information, facts, or sources that seem plausible but are entirely made up.
- **Bias:** LLMs can reflect inherent biases from their training data, which can manifest in the generated content.
- **Data Privacy:** You should not provide any private or sensitive information when using generative AI tools. This is because your data could be exposed or used for future model training on some platforms.
- **Over-reliance and Failure to Verify:** A major pitfall is not critically assessing the information an AI provides. You should always conduct your own research to verify facts and not rely on the LLM as a sole source of information.

Ethically using AI for study

You need to consider these risks when using generative AI tools for study or work

1. Do not provide any private information when using these tools
2. Verify any information provided by generative AI tools with credible sources and check for missing information.
3. Acknowledge any generative tools that you use for your assignments or work and how you used them. For example, include the name, model or version, date used and how you used it in your assignment or work.
4. Be sure to check with your course coordinators if you plan to use generative AI tools to help you complete assignments.



Referencing Generative AI tools using APA 7th style

Examples from the [APA 7th guide](#) – remember to cite in-text and in the reference list.

Citing ChatGPT and other generative AI in APA 7th

Based on [APA Style](#) guidance.

In-text

Author of generative AI model, Year of version used

Example:

(OpenAI, 2022)

OpenAI (2022)

Reference list

Author of AI model used. (Year of AI model used). *Name of AI model used* (Version of AI model used) [Type or description of AI model used]. Web address of AI model used

Example:

OpenAI. (2022). *ChatGPT* (Dec 20 version) [Large language model]. <https://chat.openai.com/>

The full transcript of a response can be included in an appendix or other supplementary materials.

Acknowledge your use of ChatGPT or other generative AI

If you use ChatGPT or other generative AI to help you **generate ideas or plan your process**, you should still **acknowledge how you used the tool**, even if you don't include any AI generated content in the assignment.

Provide a description of the AI tool used, what you did and the date accessed.

Save a copy of the transcript of your questions and responses from the generative AI tool. You can:

- [Take a screenshot](#)
- Right click and select **Save as** to save the webpage file.

Referencing tips:

- ✓ Reference in-text and in the reference list.
- ✓ Include transcripts of AI responses in an appendix to your assessment.
- ✓ Declare how you have used AI to plan and generate ideas.

App / Software

Microsoft. (Year). *Microsoft Copilot* [Software]. <https://copilot.github.com>

StudentVIP. (2017). *Lost on Campus* (Version 4.1.3) [Mobile app]. Apple App Store. <https://www.apple.com/au/ios/app-store/>

<https://libguides.murdoch.edu.au/APA/internet>

Further reading

Additional resources on prompt engineering

- [Prompt Engineering by OpenAI](#)
- [Prompt Engineering Guide by Learn Prompting](#)
- [Columbia University: How to talk to AIs – Advanced Prompt Engineering](#)
- University of Queensland Library. (2023). "Prompt Engineering." [Prompt engineering - Artificial intelligence - Library Guides at University of Queensland Library](#)
- DAIR.AI. (2024). *Prompt Engineering Guide*. <https://www.promptingguide.ai/>
- Schulhoff, S. (2024). *Prompt Engineering Guide*. Learn Prompting. <https://learnprompting.org/docs/introduction>
- *Include paper of Schulhoff on different prompting techniques*
- Aly, B. (2023). *How to talk to AIs: Advanced Prompt Engineering*. Columbia University. <https://etc.cuit.columbia.edu/news/how-talk-ais-advanced-prompt-engineering>
- Check out [LinkedIn Learning](#) – there's a lot of free prompt engineering classes.

AI Resources



- [AI Student Hub](#)
- [AI Researcher Hub](#)
- [AI Teacher Hub](#)
- [Artificial Intelligence Module](#)
 - The module discusses the use of AI at UQ, including guidelines for students and staff, delves into the legal, ethical, and social implications of AI, and offers practical examples of AI applications. The module also guides readers on how to prompt AI tools, evaluate AI-generated content, and properly cite and publish with AI.
- [Guide to Acknowledging and Referencing AI](#)
- [Machine Translation for Assignments](#)
 - Machine translation is an automated process in which a computer program converts text in one language into another.
- [Microsoft Copilot Chat Guide](#)
 - Copilot Chat is provided to you for free and relies on the GPT-5 large language model. Because Copilot Chat is UQ's enterprise system, it provides a higher level of data security and privacy than other AI tools.



Feedback and Questions

Prompt Engineering - Feedback



Thank you and good luck with your studies!

Faculty Services Librarian Team

The University of Queensland Library

Email: librarians@library.uq.edu.au

<https://web.library.uq.edu.au/library-services/contact-faculty-services-librarians>

