GENERATIVE AI COURSE

Prompt Engineering

Lecture Notes for Students

Muhammad Waqas Tahir | Lecturer, The Superior University, Faisalabad Campus

Learning Outcomes

After studying these notes, you should be able to:

- Explain what a prompt is and why prompt quality changes the quality of AI output.
- Describe, in simple terms, how a large language model (LLM) generates text.
- Write clear prompts using the five-part prompt framework.
- Apply core techniques: zero-shot, few-shot, chain-of-thought and iterative refinement.
- Identify the limitations of AI tools and use them responsibly in academic work.

1. Introduction to Prompt Engineering

What is a prompt?

A **prompt** is the instruction, question or context that you give to a generative AI model. It can be a single question, a detailed task description, or even a conversation with examples. The model reads your prompt and produces a response based on patterns it learned during training.

What is prompt engineering?

Prompt engineering is the practice of designing and refining prompts so that an AI model gives accurate, relevant and useful output. It is not about tricking the AI. It is about communicating clearly, the same way you would explain a task to a capable new teammate who knows nothing about your situation.

Why does it matter?

- **Better results:** a clear prompt reduces vague, generic or wrong answers.
- **Saves time:** fewer retries and corrections are needed.
- **Consistency:** well-structured prompts can be reused for repeated tasks.
- **Career skill:** employers increasingly expect graduates to work effectively with AI tools.

> **KEY IDEA**
>
> The quality of the output depends heavily on the quality of the input. A vague prompt gives a vague answer.

2. How Large Language Models Respond

Tools such as ChatGPT, Claude and Gemini are built on **large language models (LLMs)** . An LLM is trained on huge amounts of text and learns to predict what text should come next. A few basic ideas will help you write better prompts.

Concept	Meaning	Why it matters for prompting
Token	A small piece of text (a word or part of a word) that the model reads and writes.	Models can handle only a limited number of tokens. Very long prompts may be cut or partly ignored.
Context window	The total text (your prompt plus the conversation) the model can see at one time.	State important instructions clearly. Do not assume the model remembers everything from a very long chat.
Probabilisticoutput	The model chooses likely next words, so answers can differ each time.	The same question may give different answers. Be specific to get more stable results.
Temperature	A setting that controls randomness. Low is focused, high is more creative.	Use low for facts and code, higher for brainstorming and creative writing.
Hallucination	When the model gives confident but false or invented information.	Always verify facts, numbers, names and references.

3. Anatomy of a Good Prompt

A strong prompt usually contains the following five parts. You do not need all five every time, but the more relevant parts you include, the better the result.

Part	Purpose	Example
1. Role	Tells the AI who it should act as.	Act as a university lecturer in Computer Science.
2. Context	Gives background so the answer fits your situation.	My students are second-semester undergraduates with no prior AI knowledge.
3. Task	States clearly what you want done.	Explain how a neural network learns.
4. Format	Defines the shape of the output.	Use simple language, three short paragraphs and one real-life analogy.
5. Constraints	Sets limits such as tone, length or things to avoid.	Maximum 200 words. Avoid mathematical formulas.

MEMORY TRICK

Role + Context + Task + Format + Constraints. Ask yourself: who is answering, for whom, to do what, in what form, and within what limits?

Weak prompt versus improved prompt

WEAK PROMPT

Tell me about climate change.

IMPROVED PROMPT

Act as a science teacher. I am preparing a 10-minute class talk for first-year university students. Explain the main causes of climate change in simple language. Give five key points as bullet points, each in one sentence, and end with one discussion question.

Notice how the improved prompt answers these questions: who, for whom, what, how and how long.

4. Core Prompting Techniques

Technique	What it means	When to use it
Zero-shot	Ask directly without giving any examples.	Simple, common tasks such as summarizing or translating.
Few-shot	Give two or three examples of the input and the output you expect, then make your real request.	When you need a specific style, format or pattern.
Chain-of-thought	Ask the model to think and explain step by step before giving the final answer.	Logic, maths and multi-step reasoning problems.
Role prompting	Assign a persona or area of expertise to the model.	To control tone, depth and point of view.
Iterative refinement	Improve the answer through follow-up prompts instead of expecting a perfect first reply.	Almost always. Treat the AI as a conversation partner.
Prompt chaining	Break a big task into smaller prompts, using the output of one as the input of the next.	Long or complex work such as reports and projects.
Self-check	Ask the model to review its own answer for errors or missing points.	To improve quality before you rely on the output.

Example: few-shot prompting

> **PROMPT**
>
> Convert each sentence into a polite email opening.
> Sentence: I need the report. Email: Could you please share the report when you have a moment?
> Sentence: Send me the file. Email: I would appreciate it if you could send me the file.
> Sentence: Tell me the meeting time. Email:

The two examples teach the model the pattern, so it can complete the third one in the same style.

Example: chain-of-thought prompting

> **PROMPT**
>
> A shop gives a 20 percent discount on a jacket priced at 4,500 rupees and then adds 5 percent tax on the discounted price. What is the final price? Think step by step and show each calculation before giving the final answer.

5. A Simple Prompt Improvement Cycle

- **Write:** create a first prompt using the five-part framework.
- **Test:** run it and read the output carefully.
- **Evaluate:** check accuracy, relevance, tone and format.
- **Refine:** add missing context, tighten constraints or give an example.
- **Save:** keep the prompts that work well so you can reuse them.

6. Limitations and Risks

- **Hallucinations:**AI can invent facts, quotations, statistics and even references that do not exist.
- **Bias:** models learn from human-written data and may reflect its stereotypes.
- **Outdated knowledge:** a model may not know recent events unless it can search the web.
- **Privacy:** never paste passwords, personal data or confidential documents into public AI tools.
- **Over-reliance:** using AI to replace your own thinking weakens your learning.

7. Responsible and Ethical Use in Academics

- Follow your course and university rules about the use of AI in assignments.
- Use AI to understand, brainstorm and improve your work, not to submit work that is not yours.
- Be open: mention which AI tool you used and for what purpose when asked.
- Verify important information from reliable sources such as textbooks and research papers.
- Remember that you are responsible for everything you submit.

KEY TAKEAWAYS

A prompt is your instruction to the AI. Use Role, Context, Task, Format and Constraints. Use examples and step-by-step thinking for harder tasks. Refine your prompt, verify the output and use AI honestly.

Revision Questions

- **Q1.** Define prompt and prompt engineering in your own words.
- **Q2.** Explain token, context window and hallucination with one example each.
- **Q3.** Write the five parts of a good prompt and give one example for each.
- **Q4.** Differentiate between zero-shot and few-shot prompting.
- **Q5.** Why is chain-of-thought prompting useful for reasoning tasks?
- **Q6.** List three risks of using AI tools and explain how you can reduce each one.

Further reading: free prompting guides are published by OpenAI, Anthropic and Google. Read them to see how different tools recommend structuring prompts.