The background features a complex, multi-layered wireframe structure composed of numerous thin, light-colored lines forming various geometric shapes and planes. The structure is semi-transparent and appears to be floating or layered, creating a sense of depth and complexity. The overall aesthetic is futuristic and technical, with a color palette transitioning from dark grey/black at the top to a gradient of purple and blue at the bottom.

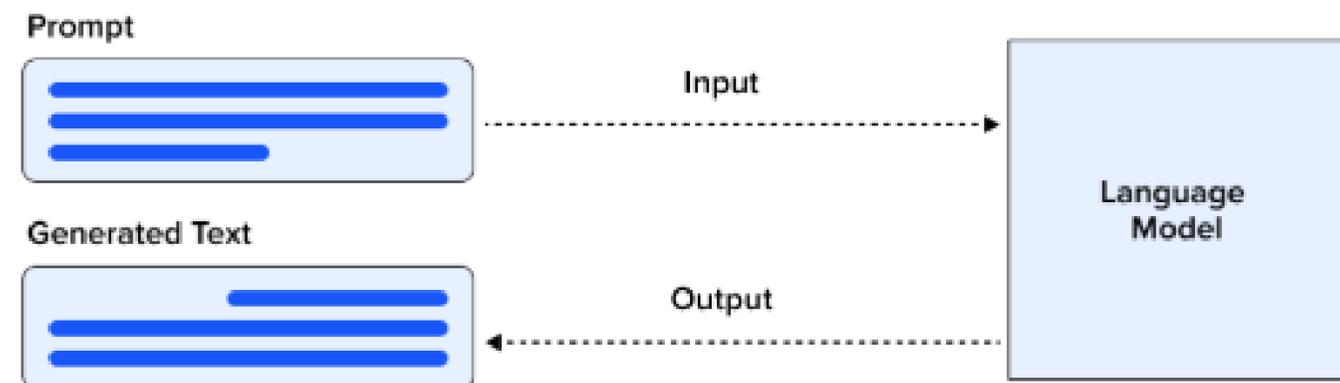
Mind Over Model: Unlocking AI's True
Potential

Prompt Engineering for Large Language Models

H. M. Samadhi Chathuranga Rathnayake

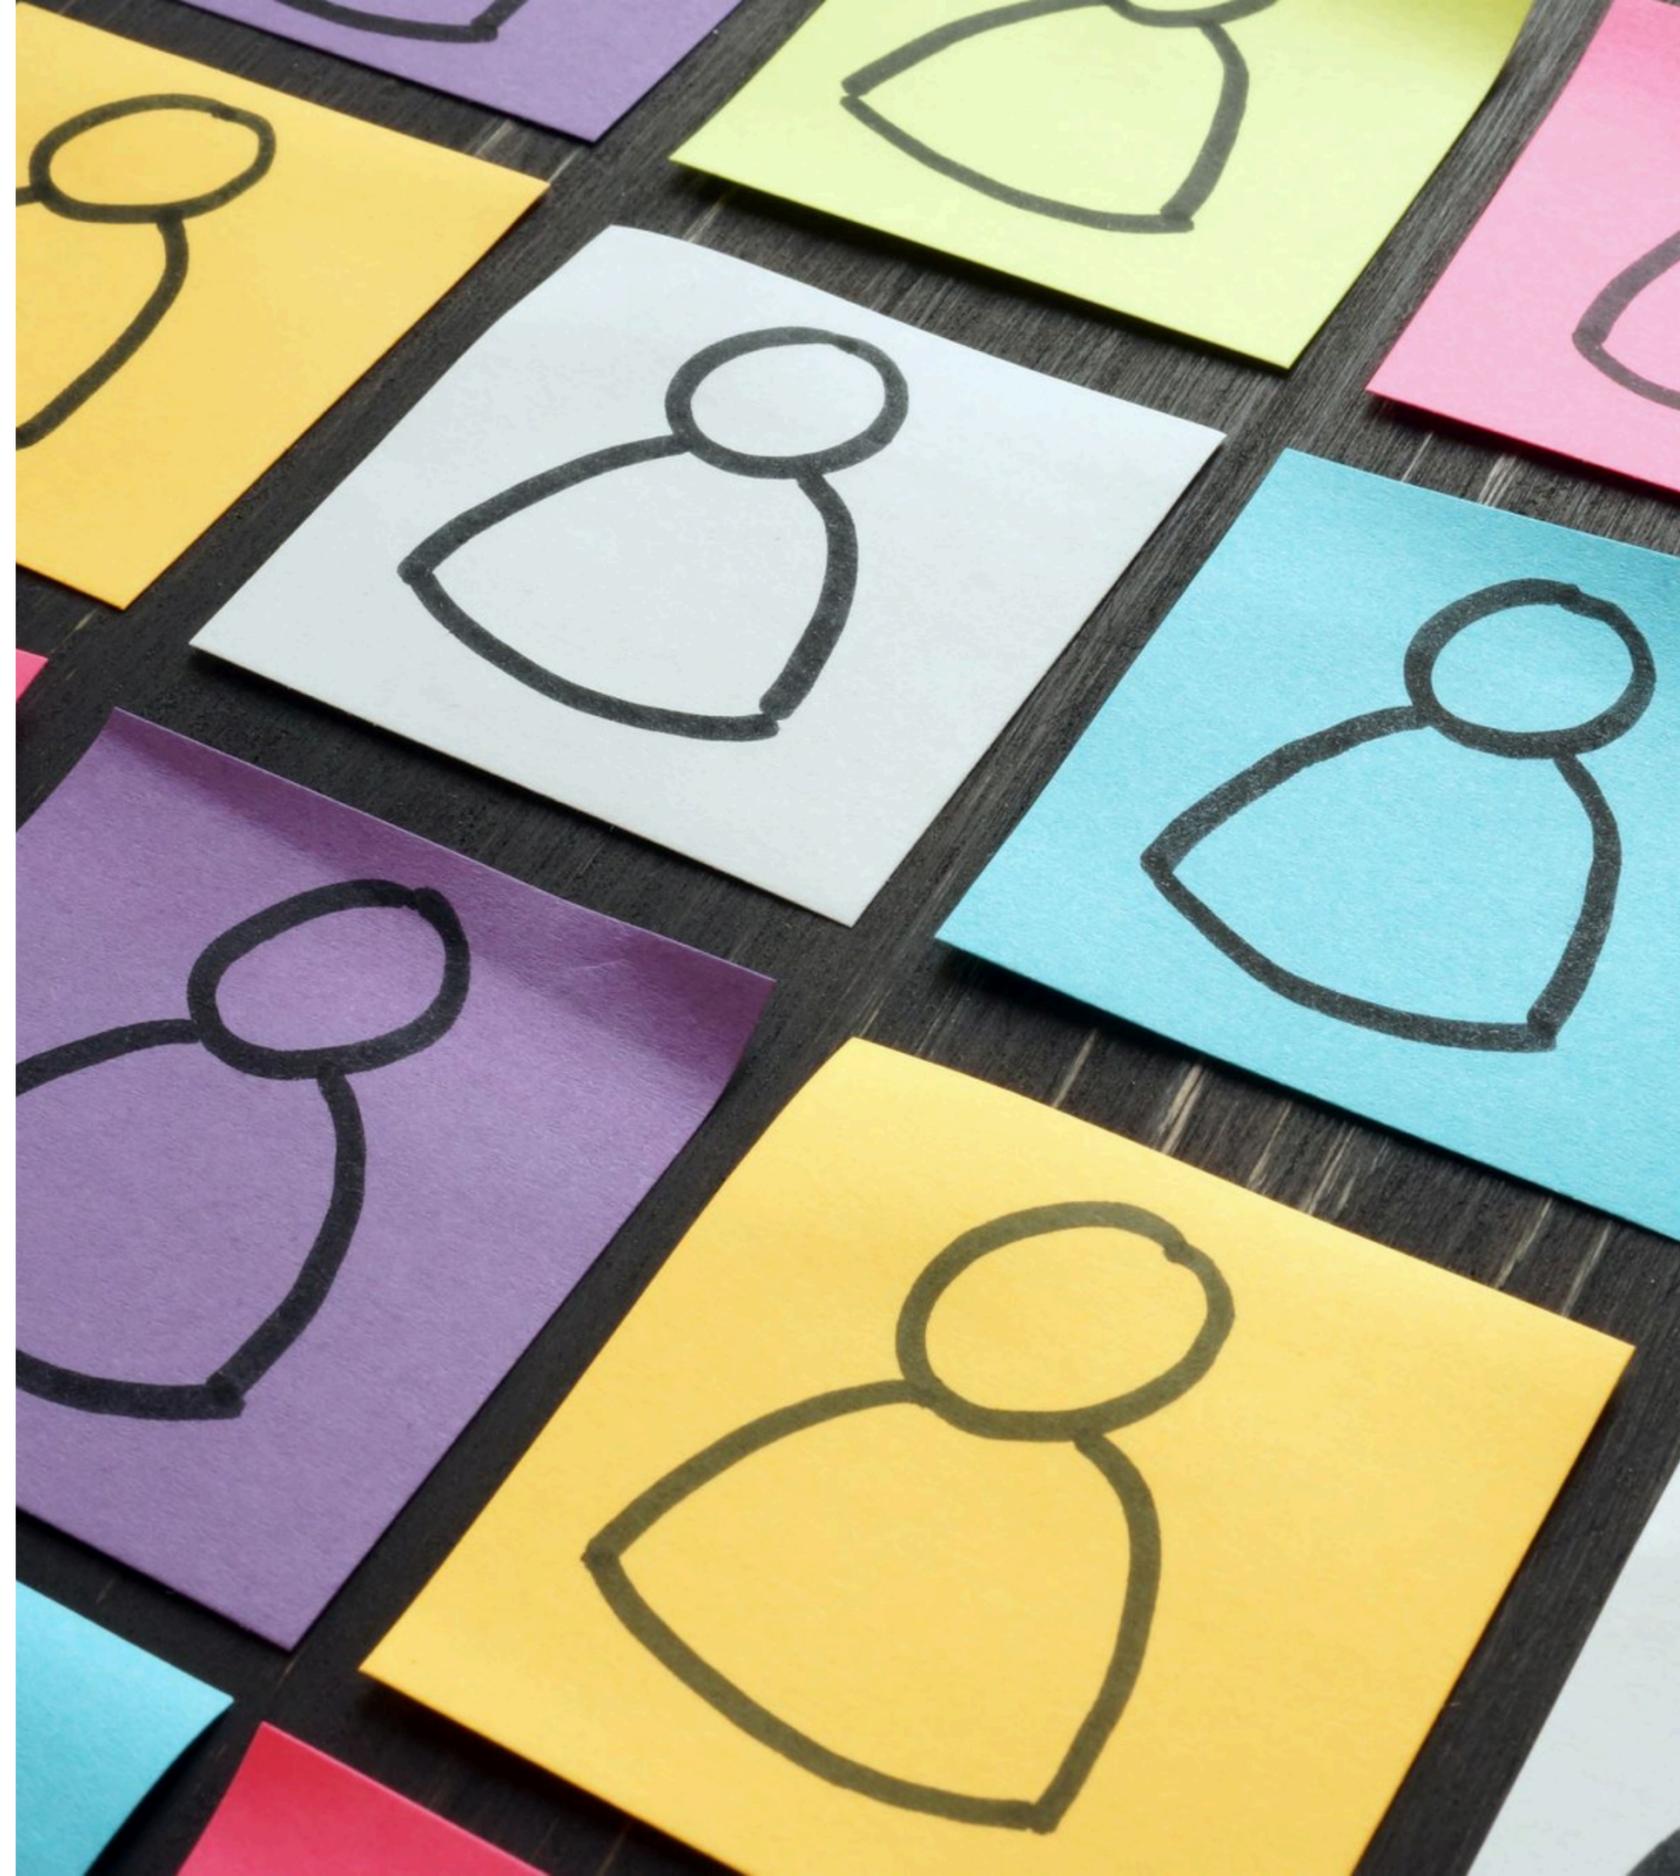
Prompt Engineering

- Prompt engineering is the practice of designing and refining prompts (instructions or questions) to elicit specific, desired responses from AI models, particularly large language models (LLMs), enabling more effective and targeted interactions with these AI systems



Key Parts of a Prompt

- Prompts are made up of several key components that work together to guide the AI.
- While not every prompt will contain all these elements, understanding how each part functions can help you create more targeted and effective inputs.
- The key parts of a prompt are:
 - The Directive
 - Examples
 - Role (Persona)
 - Output Formatting
 - Additional Information



The Directive

- A directive is a concise instruction or question that gives the AI a clear task to perform.
- It can range from a request to generate text, solve a problem, or format information in a specific way.
- Ex- Tell me five good books to read.



Examples

- When the task is more complex, providing Examples can help guide the AI in producing more accurate responses.
- This technique is especially useful in few-shot and one-shot prompting, where a model is given one or more examples of what you expect in the output.
- Ex- Translate the following sentences: Q: I like apples. A: Me gustan las manzanas. Q: I enjoy walking.



Role (Persona)

- Assigning a Role to the AI, also known as a persona, helps frame the response in a specific way.
- By telling the AI to act as an expert, a professional, or a specific character, you can guide the tone, style, and content of the response.
- Ex- You are a doctor. Based on the following symptoms, diagnose the patient.
- Ex- You are a customer service agent. Write an email apologizing for a delayed order.



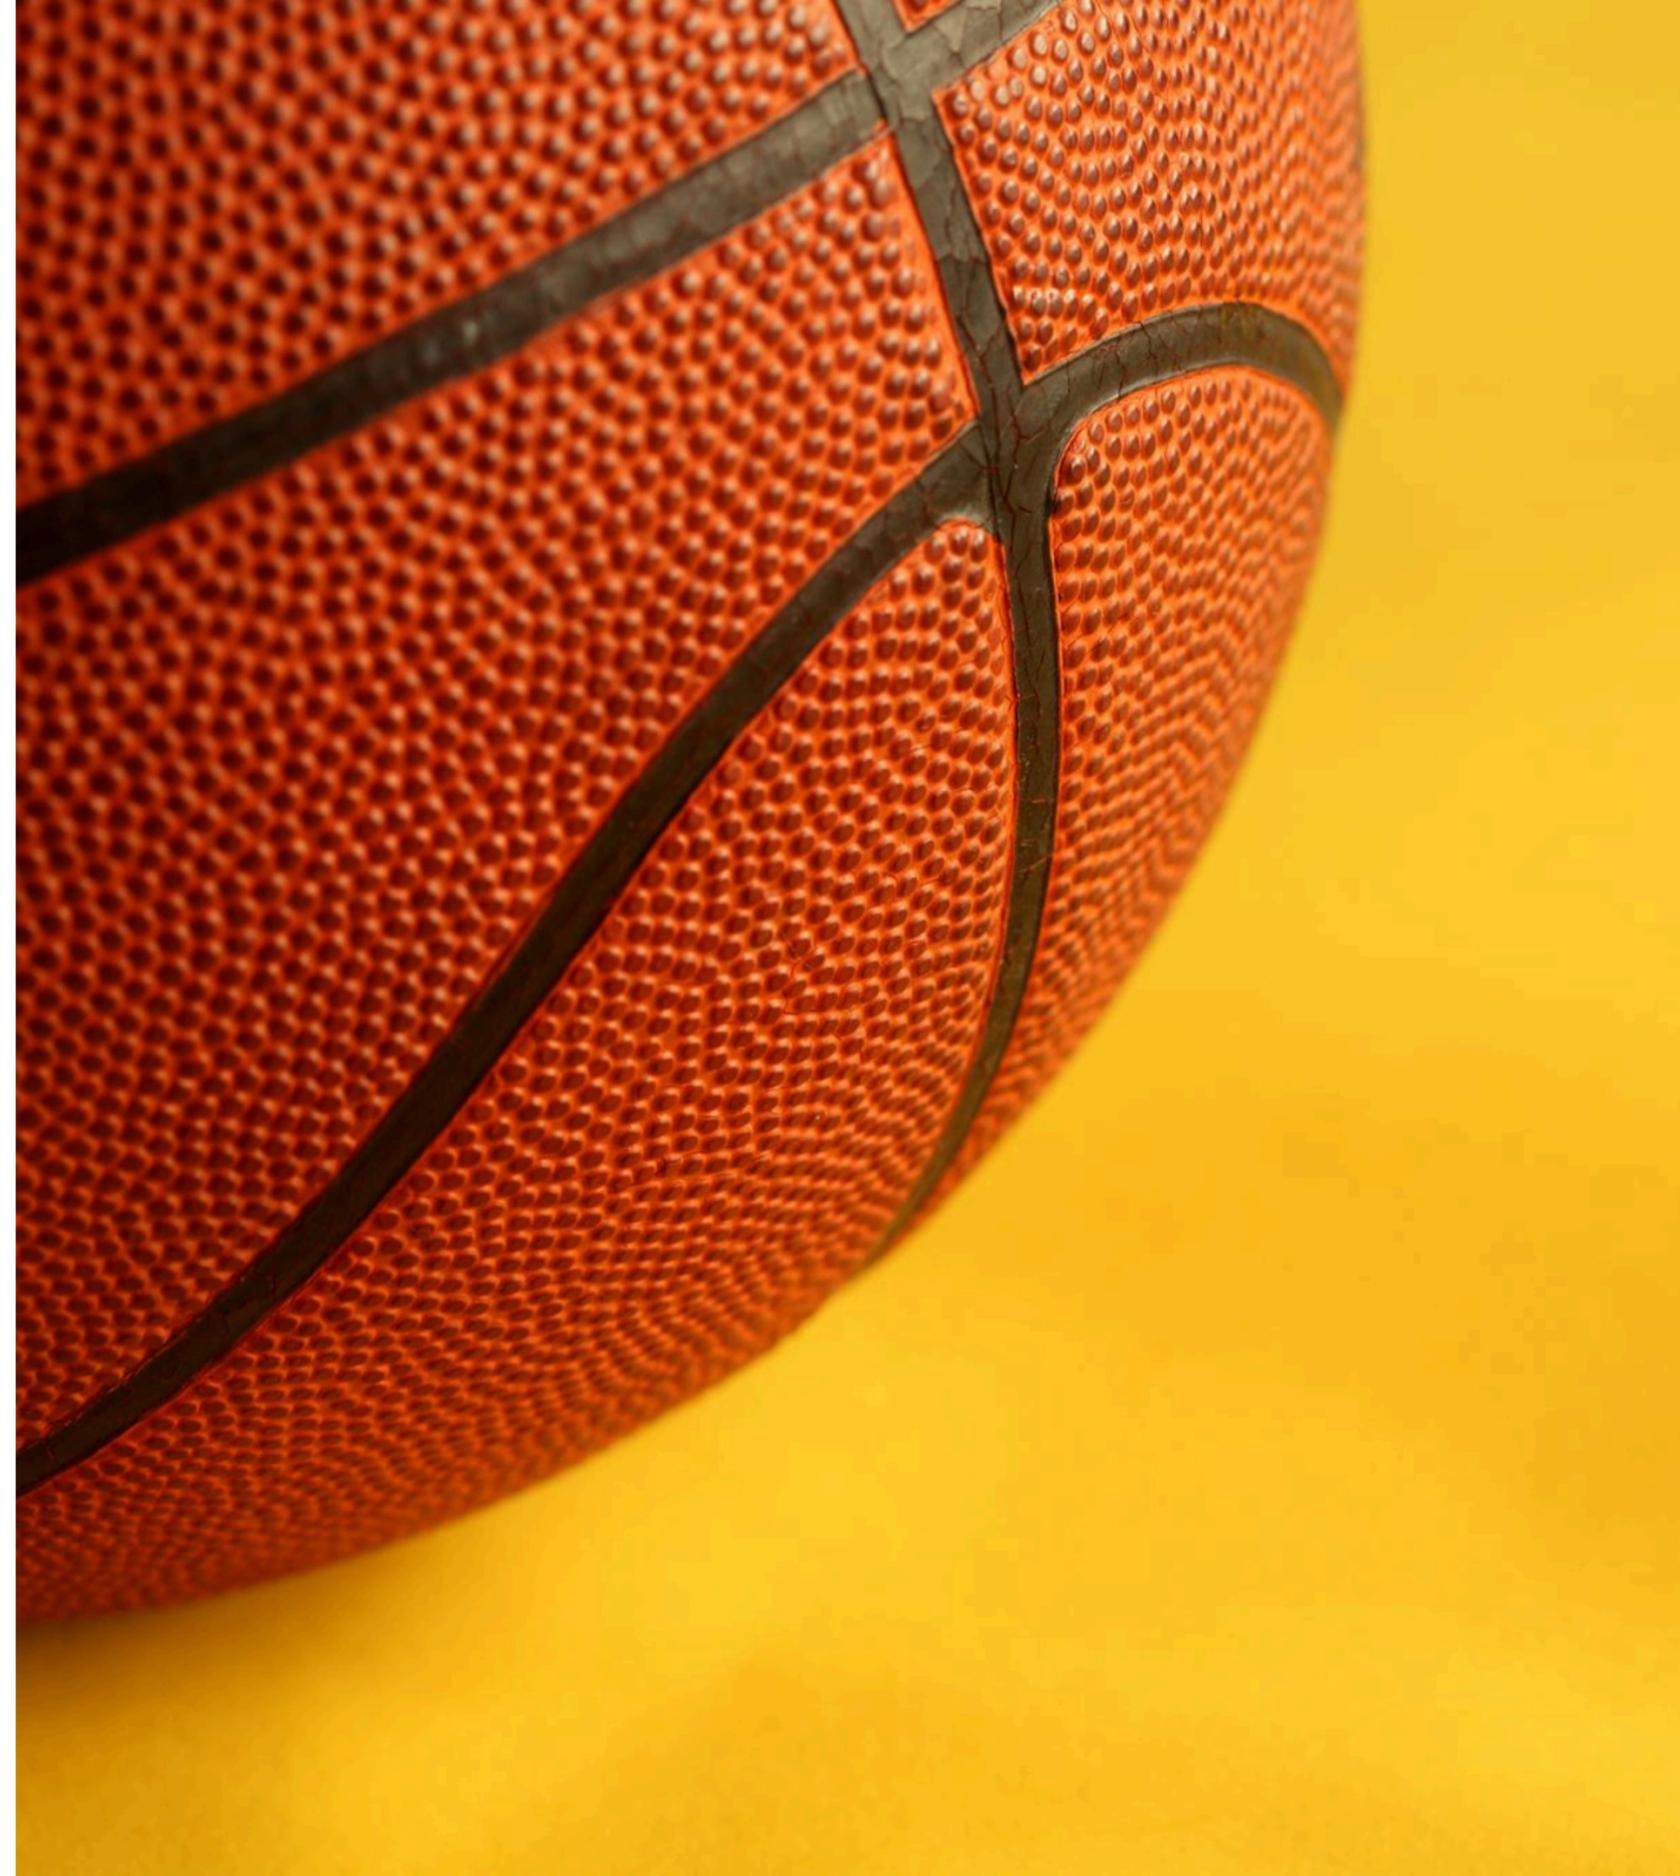
Output Formatting

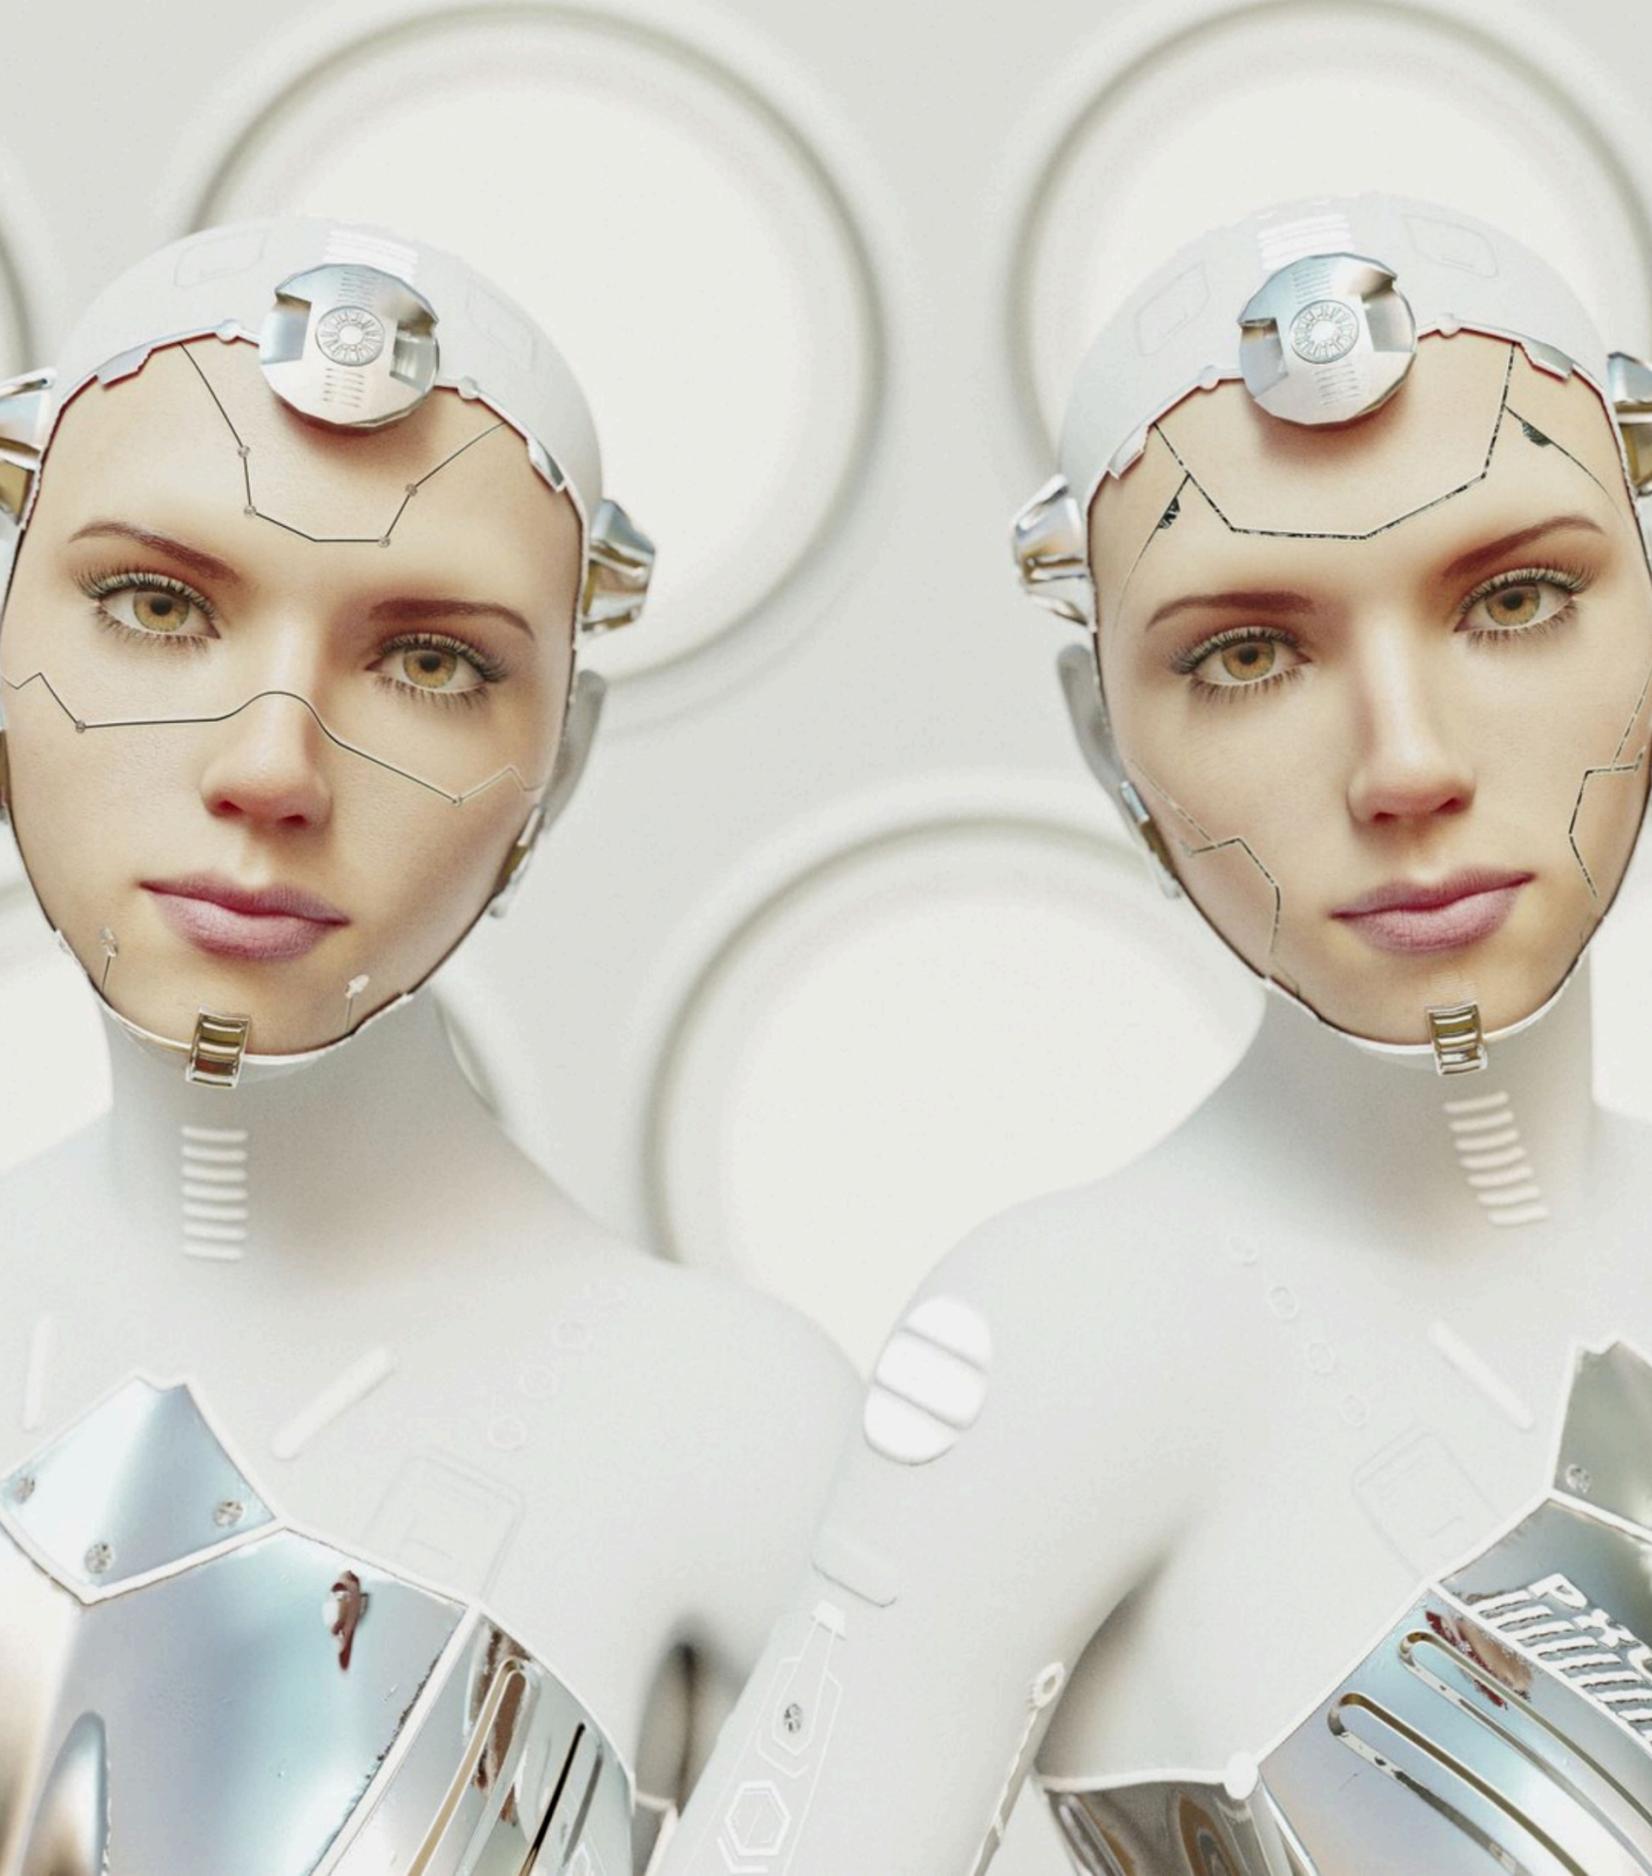
- Sometimes, it's important to specify the format in which you want the AI to present its output.
- Output Formatting ensures that the response follows a particular structure—whether it's a list, a table, or a paragraph.
- Specifying the format can help prevent misunderstandings and reduce the need for additional post-processing.
- Ex- Case: 2024_ABC_International Client: XYZ Corporation Jurisdiction: EU & USA Filed Date: 2024-09-01 Status: Active Lead Attorney: John Doe Next Hearing: 2024-10-15
- Output this information as a CSV.



Additional Information

- Additional information can include relevant facts, data, or other background information that helps the AI generate a more accurate and contextually appropriate response.
- This element is especially important for complex tasks that require specific knowledge.
- Ex- January 1, 2000: Fractured right arm playing basketball. Treated with a cast. February 15, 2010: Diagnosed with hypertension. You are a doctor. Predict the patient's future health risks based on this history.





Shot-Based Prompting

- **Zero-Shot Prompting:** No examples are provided, and the model must rely entirely on its pre-trained knowledge.
- **One-Shot Prompting:** A single example is given to clarify the task for the model.
- **Few-Shot Prompting:** Two or more examples are included, allowing the model to recognize patterns and deliver more accurate responses.

Zero-Shot Prompting

- Zero-shot prompting is the simplest form of prompting.
- Here, we give the model a direct instruction to perform a task without providing any examples or demonstrations.
- This means the model has to rely entirely on its pre-trained knowledge to figure out how to complete the task.
- Ex- Classify the sentiment of the following text as positive, negative, or neutral. Text: I think the vacation was okay. Sentiment:

One-shot Prompting

- One-shot prompting enhances zero-shot prompting by providing a single example before the new task, which helps clarify expectations and improves model performance.
- Ex- Classify the sentiment of the following text as positive, negative, or neutral. Text: The product is terrible. Sentiment: Negative

Few-Shot Prompting

- Few-shot prompting provides two or more examples, which helps the model recognize patterns and handle more complex tasks.
- With more examples, the model gains a better understanding of the task, leading to improved accuracy and consistency.
- Ex- Classify the sentiment of the following text as positive, negative, or neutral. Text: The product is terrible. Sentiment: Negative, Text: Super helpful, worth it Sentiment: Positive



Where to use?

- Zero-shot prompting:
 - Use this when the task is simple, well-understood, or frequently encountered in the model's training data.
 - It's efficient for tasks like basic arithmetic, general queries, or sentiment classification for common phrases.
 - One-shot prompting:
 - This is helpful for tasks that need more specific guidance or when the model struggles with ambiguity.
 - Providing a single example can clarify the task, improving accuracy in tasks like basic classification or structured information extraction.
-



Where to use?

- Few-shot prompting:
 - Best used for complex tasks requiring multiple examples to establish patterns.
 - This technique is ideal for tasks that involve varied inputs, require precise formatting, or demand a higher degree of accuracy, such as generating structured outputs or handling nuanced classifications.
-

Chain-of-Thought Prompting (CoT)

Chain-of-Thought prompting is a technique that improves the performance of language models by explicitly prompting the model to generate a step-by-step explanation or reasoning process before arriving at a final answer.

This method helps the model to break down the problem and not skip any intermediate tasks to avoid reasoning failures.

Chain-of-Thought Prompting (CoT)

Standard Prompting

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The answer is 27. ❌

Chain-of-Thought Prompting

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✅

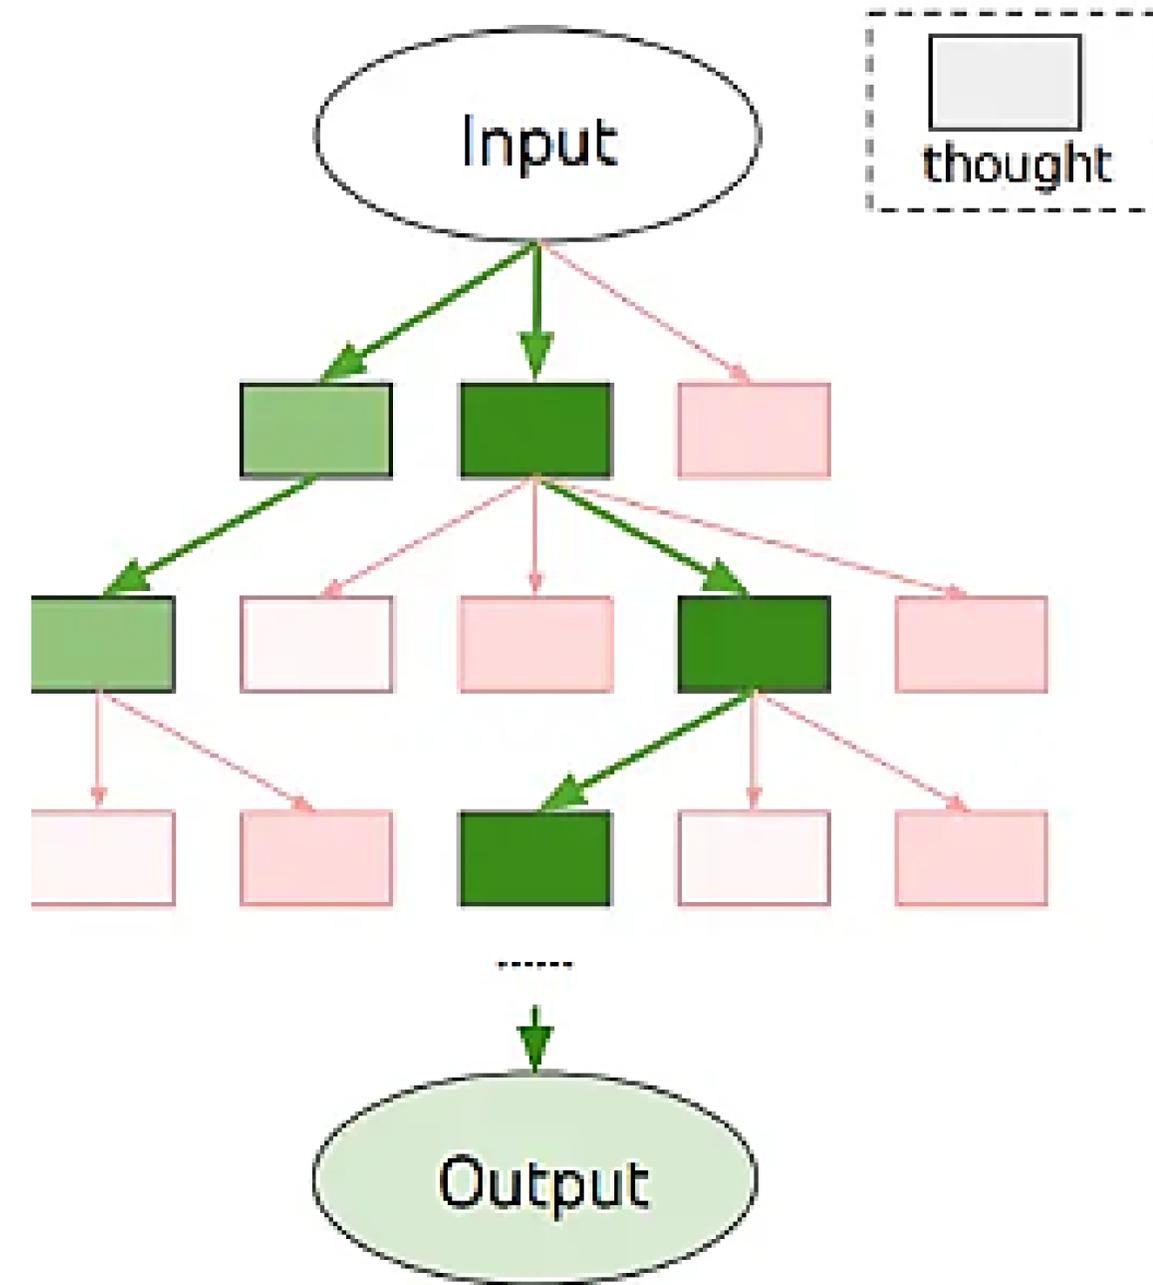
Tree of Thoughts Prompting (ToT)

ToT maintains a tree of thoughts, where thoughts represent coherent language sequences that serve as intermediate steps toward solving a problem.

This approach enables an LM to self-evaluate the progress through intermediate thoughts made towards solving a problem through a deliberate reasoning process.

Tree of Thoughts Prompting (ToT)

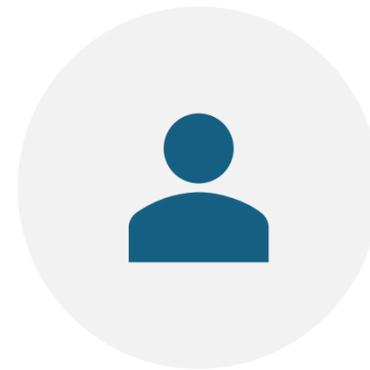
- Ex-
- Imagine three different experts are answering this question.
- All experts will write down 1 step of their thinking,
- then share it with the group.
- Then all experts will go on to the next step, etc.
- If any expert realises they're wrong at any point then they leave.
- The question is...



Other Advanced Prompting Methods



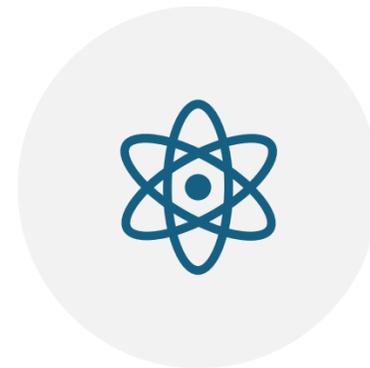
Meta Prompting



Self Consistency Prompting



Graph of Thoughts Prompting (GoT)



Atom of Thoughts Prompting (AoT)

Ethical Prompt Engineering

Bias and Fairness

Privacy and Data Protection

Transparency and Explainability

Inclusivity and Accessibility

User Consent and Control

Social Impact Assessment

Compliance with Regulations

What are AI Hallucinations?



AI hallucinations happen when artificial intelligence gives answers that sound correct but are actually wrong or make no sense.



This happens because the AI, even though it's trained on a lot of data, sometimes misunderstands or misinterprets information.



For example, if an AI is asked about a historical event, it might create a detailed but completely false story.



Another example is an AI giving medical advice that sounds believable but is not accurate.



These mistakes can be either small or very noticeable, depending on the complexity of the question and the context.

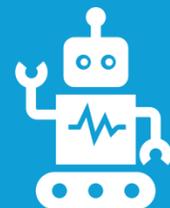
Method 1: “According to...” Prompting



“According to...” prompting is a way to make AI give more accurate answers.



When you start your question with “According to [source],” you help the AI base its answers on real and reliable information.



This makes the AI less likely to make things up.

Method 1: “According to...” Prompting

Before: "Explain the causes of climate change."

After: "According to the Intergovernmental Panel on Climate Change (IPCC), explain the main causes of climate change."

Method 2: Chain-of- Verification Prompting

Chain-of-Verification Prompting is a method where you ask the AI to check the facts it provides step by step.

By verifying each part of its response, the AI can reduce mistakes and improve the accuracy of its answers.

This method involves breaking down the information into smaller parts and confirming each one before moving on to the next.

Method 2: Chain-of- Verification Prompting

Before: "Tell me about the history of the internet."

After: "First, explain when the internet was invented. Then, confirm who were the key people involved."

Method 3: Step-Back Prompting

Step-Back Prompting is a technique where you ask the AI to review its previous response and make sure it is accurate.

By encouraging the AI to take a step back and reassess what it has said, you can help it catch and correct any mistakes.

This method is useful for improving the overall quality and reliability of the AI's answers.

Method 3: Step-Back Prompting

Before: "Explain how photosynthesis works."

After: "Explain how photosynthesis works. Now, step back and check if your explanation covers all the key points accurately."

Method 4: Contextual Anchoring

Contextual Anchoring is a method where you provide the AI with a specific context to guide its responses.

By giving the AI more background information or a clear setting, you help it generate answers that are more relevant and accurate.

This reduces the chance of the AI making things up or going off-topic.

Method 4: Contextual Anchoring

Before: "What are the benefits of exercise?"

After: "In the context of heart health, what are the benefits of regular exercise?"

Method 5: Layered Questioning

Layered Questioning involves asking multiple, related questions in sequence to dig deeper into a topic.

By breaking down a broad question into several smaller, more focused ones, you can help the AI provide more detailed and accurate information.

This method ensures that each part of the answer is thoroughly explored and verified.

Method 5: Layered Questioning

Before: "Tell me about renewable energy."

After: "First, explain what renewable energy is. Next, list the main types of renewable energy sources. Finally, describe the benefits of using renewable energy."



Reflective Prompting is a technique where the AI is asked to reflect on its previous answer and consider if there are any errors or omissions.

This helps the AI to self-correct and improve the accuracy of its responses.

By encouraging the AI to think critically about its output, you can enhance the quality and reliability of the information it provides.



Method 6: Reflective Prompting



Method 6:
Reflective
Prompting

Before: "Explain the process of photosynthesis."

After: "Explain the process of photosynthesis."



Method 7: Scenario-Based Prompting

Scenario-Based Prompting is a method where you place the AI in a specific scenario to guide its responses.

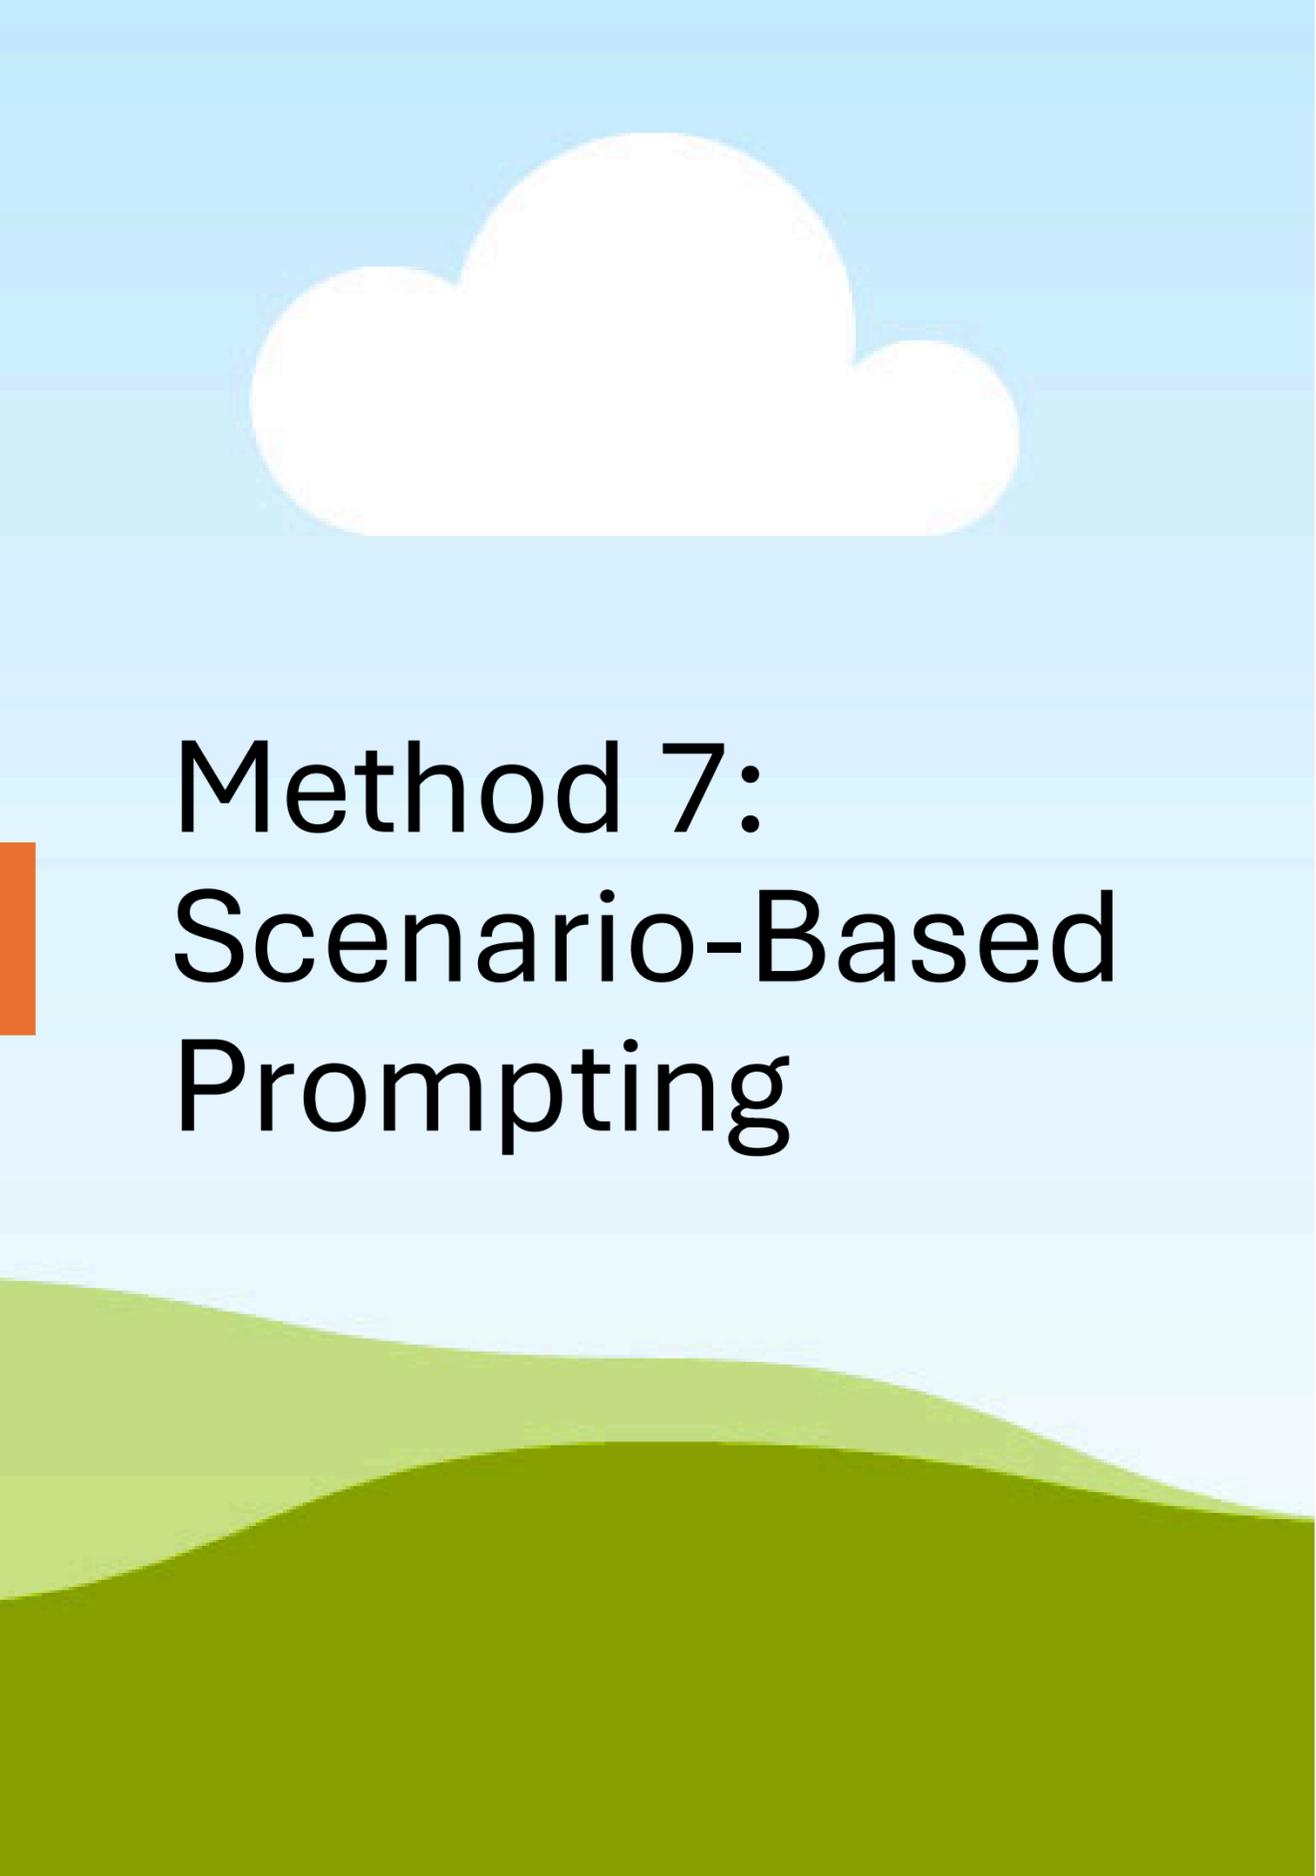
By creating a detailed scenario or setting, you help the AI focus on providing relevant and accurate information tailored to that context.

This approach reduces the likelihood of the AI generating off-topic or incorrect information.



Before: "What should you do in a medical emergency?"

After: "Imagine you are a medical doctor and someone suddenly has a severe allergic reaction. What steps should you take to handle this medical emergency?"



Method 7: Scenario-Based Prompting



Method 8: Feedback Loop Integration

Feedback Loop Integration involves incorporating feedback into the AI's prompt engineering process.

This means you provide feedback on the AI's responses, and use that feedback to improve future interactions.

By continuously refining prompts based on feedback, you can help the AI learn to generate more accurate and reliable information over time.



Method 8: Feedback Loop Integration

Before: "Tell me about the history of the internet."

After: "Tell me about the history of the internet. Your last response missed key events. Please include the invention of the World Wide Web and the rise of social media."



Progressive Prompting is a technique where prompts are gradually detailed to ensure the AI provides accurate and comprehensive answers.

By starting with a broad question and progressively narrowing down the focus, you can guide the AI to cover all necessary aspects of a topic.

This method helps in building up the information step-by-step, making it more reliable.



Method 9: Progressive Prompting



Method 9: Progressive Prompting

Before: "Explain renewable energy."

After: "Explain what renewable energy is. Now, list the main types of renewable energy sources. Finally, describe the benefits of using renewable energy."