COMPUTER VISION
MEI/1

University of Beira Interior, Department of Informatics

Hugo Pedro Proença

hugomcp@di.ubi.pt, **2024/25**

Prompt Engineering? What is It?

- *"Prompt engineering is the practice of crafting inputs (prompts) to guide generative AI models toward producing desired outputs."*
  - Reynolds, L., & McDonell, K. (2021). "Prompt Programming for Large Language Models: Beyond the Few-Shot Paradigm."

- *"Prompt engineering involves optimizing the instructions given to AI models to maximize the relevance and accuracy of the generated responses."*
  - Brown, T., et al. (2020). "Language Models are Few-Shot Learners." *NeurIPS*.

- "*The process of designing effective input prompts to ensure large language models generate outputs aligned with user intentions and specific use cases*."
  - Reference: OpenAI. (2021). "Guidelines for Prompt Engineering with GPT-3."

- "*Prompt engineering refers to the process of designing and refining natural language prompts to harness the capabilities of large language models for downstream tasks.*"
  - Google Research. (2022). "Understanding and Improving Few-Shot Prompting."

Prompt Engineering? Purposes

- Essentially, it's about crafting inputs (prompts) that guide the AI to generate the most useful and relevant responses based on the task or goal at hand.

- **Improving Accuracy and Relevance**: By carefully crafting prompts, we can guide the AI to focus on the most relevant information, improving the quality and precision of the answers.

- **Enhancing Creativity**: In tasks where creativity is important (e.g., writing, brainstorming, or generating artwork), prompt engineering can steer the AI towards more imaginative, unique, or out-of-the-box responses.

- **Controlling Output Style and Tone**: Prompt engineering helps to influence the tone and style of the output, such as making it formal, casual, humorous, or professional. This is particularly useful for tasks that require a specific communication style.

- **Reducing Bias or Inappropriate Content**: Well-designed prompts can help mitigate bias, inappropriate, or harmful content generation by setting clearer boundaries and providing context or guidelines within the prompt itself.

- **Task-Specific Optimization**: Different tasks require different kinds of responses. A prompt for summarizing an article will be different from one used for answering a question. By adjusting the prompt for each task, you can optimize the AI's output to match the intended objective.

- **Handling Ambiguity**: When the input is unclear or ambiguous, prompt engineering helps clarify and specify what information the user is looking for, leading to better, more accurate responses.

- **Encouraging Conciseness or Detail**: Depending on the goal, prompt engineering can be used to prompt for short, concise responses or longer, more detailed explanations.

- **Fostering Collaboration**: In collaborative tasks, like group brainstorming or coding, prompt engineering can help AI assist more effectively by focusing on collaborative feedback, suggestions, or adapting to team needs.

Essentials of Prompts

- A **prompt** is the input text or query provided to a language model, which serves as instructions for the model to generate a response. Prompts can vary in length and complexity, ranging from simple phrases to detailed instructions or context.
OpenAI API Documentation, 2021

- A **prompt** is the initial textual input or context provided to a generative language model, designed to specify or guide the task that the model is expected to perform. The quality of the output is often highly dependent on the clarity and specificity of the prompt.
Brown, T. B., et al. (2020). "Language Models Are Few-Shot Learners." *Advances in Neural Information Processing Systems (NeurIPS)*.

- A **prompt** refers to any form of structured or unstructured input data provided to a generative AI system, such as a text snippet, question, or context. It acts as the starting point for the system to produce text, images, or other outputs, depending on the task.
Bommasani, R., et al. (2022). "On the Opportunities and Risks of Foundation Models." *Stanford Institute for Human-Centered Artificial Intelligence (HAI)*.

Source: https://www.techtarget.com/searchenterpriseai/definition/AI-prompt

Explicit Prompts

- Explicit prompts are clear, specific instructions provided to a generative model, detailing what the model is expected to do. These prompts minimize ambiguity by outlining the task clearly.

- **Characteristics**:
    - Direct instructions.
    - Well-defined structure and context.
    - High control over the output.

- **Examples**:
    - **Text Generation (ChatGPT):**
      "*Write a 100-word summary of the novel Pride and Prejudice by Jane Austen.*"
      Output: A concise summary of the book focusing on its main themes and characters.
    - **Image Generation (DALL-E):**
      "Generate an image of a futuristic city with flying cars and skyscrapers in a cyberpunk theme."
      Output: A vivid depiction of a futuristic cyberpunk cityscape.
    - **Code Generation:**
      "Write a Python function to calculate the factorial of a number using recursion."
      Output: A Python function explicitly performing the task.

Implicit Prompts

- Implicit prompts rely on less direct instructions or contextual cues to guide the model. These prompts often mimic human conversational inputs or vague requests, requiring the model to infer the task or context.

- **Characteristics**:
  - Less direct, conversational, or ambiguous.
  - Requires the model to "guess" the intent.
  - Greater reliance on the model's training data and context.

- **Examples**:
  - **Text Generation (ChatGPT):**
    "*What do you think happens in Pride and Prejudice?*"
    Output: A general description or interpretation of the book, potentially omitting specific details.

  - **Image Generation (DALL-E):**
    "*Imagine a futuristic world.*"
    Output: A creative but potentially varied interpretation of a futuristic setting.

  - **Code Generation:**
    "*How can I find the factorial of a number in Python?*"
    Output: An explanation or example of calculating factorials, possibly using loops or recursion.

Explicit/Implicit Prompts

Aspect	Explicit Prompt	Implicit Prompt
Clarity	Highly specific and unambiguous	Conversational or vague; requires inference
Control	Offers more control over the output	Relies more on the model's interpretation
Use Case	Ideal for task-specific outputs	Suited for open-ended or creative tasks
Example Prompt	"Write a poem about autumn in 4 lines."	"Can you write something about autumn?"
Output	A short, structured poem on autumn.	A broader response, possibly less structured.

Explicit prompts are ideal for when we need precise and reliable results. **Implicit prompts** are better for brainstorming, open-ended creativity, or when we don't want to over-constrain the model.

Structure of Prompts

- **A. Context**

  Providing background information or setting the stage for the task.

  "*You are a travel guide specializing in European destinations. Write a 200-word description of Paris.*"

- **B. Task Instruction**

  Clearly specifying what the model is expected to do.

  "*Summarize the following text in three sentences.*"
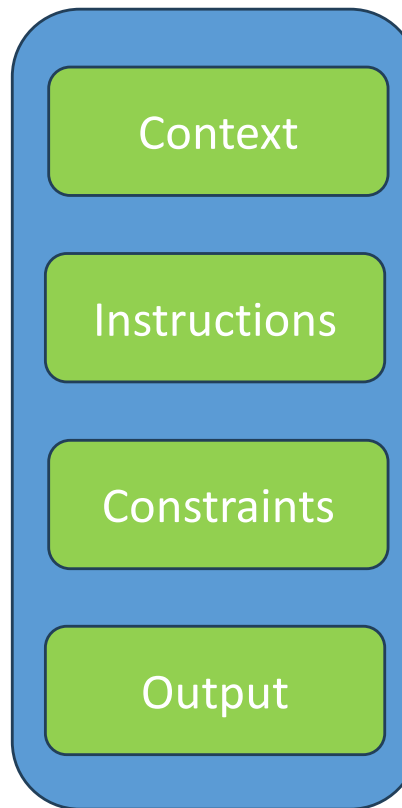
- **C. Constraints**

  Defining rules, boundaries, or limitations for the response.

  "*Generate a response in less than 50 words.*"

- **D. Desired Output Format**

  Specifying the structure, tone, or format of the output.

  "*Write the response as a bullet-point list.*"

Context	Any set up to help the intent clear and objective (role play, background, expertise, examples,…)
Instructions	What to do and how to do it, as if the instructions are given to an 8-year child
Constraints	What to avoid, and what rules to obey, that should be met regardless of the generated output
Output	What to avoid, and what rules to obey, that should be met regardless of the generated output

Good/Bad Prompts

- **Good Prompt**:
  "Write a 150-word summary of the article linked below, focusing on the main argument and supporting evidence."

- **Good Prompt**:
  "You are an academic tutor. Explain the Pythagorean theorem using a step-by-step example with a right triangle where the sides are 3, 4, and 5."

- **Bad Prompt**:
  "Summarize this article."

- **Bad Prompt:**
  "What is the Pythagorean theorem?"

Role Assignment
Assigning a role to the AI to set a specific perspective or tone. "*You are an expert chef. Write a recipe for a simple vegetarian pasta dish.*"

Step-by-Step
Breaking down complex tasks into smaller, clear steps. "*Step 1: Read the text below. Step 2: Identify the main argument. Step 3: Write a 100-word summary.*"

Use Variables
Including placeholders or dynamic variables for reusable prompts.
"Generate a tagline for a company named [CompanyName] that specializes in [Industry]."

Use Examples
Including examples to clarify the task. *"Translate English to French. Example: 'Hello' -> 'Bonjour'. Now translate: 'Goodbye'."*

Chaining Prompts
Using multiple, sequential prompts to guide a multi-step process.
Prompt 1: "Summarize the given text."
Prompt 2: "Rewrite the summary in simpler language for a younger audience."

Types of Prompts – Zero-Shot

- **Definition**: The model is asked to perform a task without being given any examples or demonstrations. The prompt relies entirely on the model's training to understand and execute the task.

- **Characteristics**:
  - No prior examples are provided in the prompt.
  - Used when the model is expected to generalize from its training data.

- **Examples**:
  - **Text Generation**:
    "*Summarize the following paragraph: The quick brown fox jumps over the lazy dog. The dog wakes up startled and chases the fox into the woods.*"
    Output: "A fox jumps over a dog, who wakes up and chases the fox."

  - **Code Generation**:
    "Write a Python function to reverse a string."
    Output: A Python function that performs the string reversal task.

Types of Prompts – One-Shot

- **Definition**: The model is given one example of the task before being asked to perform it. This helps provide minimal context while remaining efficient.

- **Characteristics**:
  - Includes a single example of the task in the prompt.
  - Useful when the model might benefit from seeing one example.

- **Examples**:
  - **Text Generation**:
    "*Here is an example of translating English to French: 'Hello' -> 'Bonjour'. Translate the following: 'Good morning'.*"
    Output: "Bon matin" or "Bonjour."

  - **Code Generation**:
    "*Example: Input: [1, 2, 3] -> Output: [3, 2, 1]. Now reverse this list: [4, 5, 6].*"
    Output: "[6, 5, 4]".

Types of Prompts – Few-Shot

- The model is given multiple examples of the task before being asked to perform it. This provides more context and helps the model better understand the desired format or behavior.

- **Characteristics**:
  - Includes 2–5 examples of the task in the prompt.
  - Useful for complex tasks or when the model needs additional context.

- **Examples**:
  - **Text Generation**:
    "Translate English to French: 'Hello' -> 'Bonjour' Example 2: Translate English to French: 'How are you?' -> 'Comment ça va?' Translate the following: 'Thank you.'
    ***Output**: "Merci."*
  - **Code Generation**:
    "Example 1: Input: [1, 2, 3] -> Output: [3, 2, 1]. Example 2: Input: ['a', 'b', 'c'] -> Output: ['c', 'b', 'a']. Now reverse this list: [4, 5, 6]."
    ***Output**: "[6, 5, 4]".*

Open-Ended Prompts

- **Definition**: Prompts that allow the model to generate creative, unrestricted responses. They encourage exploration and variety in the output.

- **Characteristics**:
  - No rigid format or constraint.
  - Useful for brainstorming, storytelling, or creative tasks.

- **Examples**:
  - **Text Generation**:
    "*Describe a futuristic city.*"
    Output: "A sprawling metropolis with levitating cars, towering skyscrapers made of translucent materials, and a glowing green sky powered by renewable energy."
  - **Image Generation**:
    "Create an artwork that represents peace."

Closed Prompts

- **Definition**: Prompts that restrict the model's response by providing clear guidelines or constraints on the format and content of the output.

- **Characteristics**:
  - Structured and specific.
  - Used for tasks requiring precision, such as question-answering or summarization.

- **Examples**:
  - **Text Generation**:
    "*Summarize this in 20 words: The quick brown fox jumps over the lazy dog. The dog wakes up startled and chases the fox into the woods.*"
    Output: "A fox jumps over a dog. The dog wakes up startled and chases the fox into the woods."
  - **Image Generation**:
    "*Generate an image of a red apple on a wooden table in a minimalist style.*"

Challenges and Limitations

- **Sensitivity to Prompt Wording**

- **Challenge**: Models can be highly sensitive to slight changes in prompt phrasing, which may drastically alter the output.

- **Example**:

  "*Summarize the key points of this article.*"
  Output: A short summary focusing on the most relevant points.

  "What is this article about?"
  ***Output*:** A general and sometimes vague response.

- **Impact**: Users may need to experiment with multiple iterations to achieve the desired output.

Challenges and Limitations

- **Lack of Robustness to Ambiguity**

- **Challenge**: Vague or ambiguous prompts lead to inconsistent or irrelevant outputs. The model cannot always infer user intent without sufficient clarity.

- **Example**:
    "Explain photosynthesis."
    Output: The explanation could be too detailed for a layperson or too simplified for a botanist.

- **Impact**: Requires precise language to ensure the output aligns with user expectations.

Challenges and Limitations

- **Balancing Specificity and Flexibility**

- **Challenge**: Overly specific prompts can restrict creativity, while overly flexible prompts may produce unstructured or irrelevant outputs.

- **Example**:
    Highly Specific: *"Write a 200-word essay about renewable energy, focusing on solar power, using formal language."*

    Too Flexible: *"Write about renewable energy."*

    Balanced: *"Write a short essay about renewable energy, highlighting its benefits and challenges."*

- **Impact**: Finding the right balance often involves trial and error.

Challenges and Limitations

- **Task Complexity**

- **Challenge**: For complex tasks requiring multiple steps or nuanced understanding, a single prompt may fail to guide the model effectively.

- **Example**:

    "Write a report analyzing the economic impacts of climate change and propose solutions."
    Issue: The model may focus more on one part of the task and neglect the other.

- **Solution**: Use chained or step-by-step prompts to address different aspects of the task sequentially.

Challenges and Limitations

- **Prompt Length Limitations**

- **Challenge**: Models like GPT have token limits (e.g., ~4,000 tokens for GPT-3.5 or ~8,000 tokens for GPT-4), which can restrict the length of both prompts and outputs.

- **Example**:
  - A prompt requiring extensive context or multiple examples may exceed the token limit, truncating the response.

- **Impact**: Users must prioritize and condense information while crafting prompts.

Generative AI Applications

- Generative AI is increasingly used in content creation, enabling creators to produce lifelike AI-generated videos, automate scripting, editing, and narration, thereby significantly reducing production time and costs.

- **Example**: Content creators are leveraging AI tools to generate videos, manage e-commerce, and handle social media tasks, enhancing productivity and expanding their reach.

- **Illustrative Video**: "Meet the content creators harnessing AI - and how they use it to make thousands per month"

Generative AI Applications

- In the film industry, generative AI is utilized to modify actors' performances, such as adjusting accents or enhancing specific aspects of dialogue delivery, to achieve a desired level of authenticity.

- **Example**: The film "The Brutalist" employed AI tools to refine the Hungarian accents of its stars, ensuring accurate pronunciation and enhancing the overall authenticity of the performances.

- **Illustrative Video**: "The Brutalist's AI Controversy, Explained"