


Prompt Engineering for Educational Publishing and Assessment



Niall McNulty
Version 1, September 2023

Introduction	3
Definition of a Prompt: What is Prompting, and Why is it Important?	4
Basic Tasks You Can Get AI to Do	4
Text Summarisation	4
Information Extraction	4
Question Answering	4
Text Classification	5
Conversation	5
Code Generation	5
Content Creation	5
Translation	5
Limitations of AI	5
Structuring a Prompt: Exploring the Components in Detail	6
Instruction	6
Context	6
Input Data	6
Output Indicator	6
How to Improve Your Prompt	6
Grounding the LLM - Using External Sources	6
Using Structured Text - The Role of Markdown and Syntax	6
Step-by-Step - Making ChatGPT Show Its Work	7
Conditional Prompts	7
Multi-turn Prompts	7
Quality Assurance	7
Accessibility	7
Types of Prompting Techniques	8
Zero-Shot Prompting	8
Few-Shot Prompting	8
Chain-of-Thought Prompting	8
Best Practices and Prompt Engineering Tips for Crafting Effective Prompts	9
Clarity and Specificity	9
Planning and Objectives	9
Customisation and Context	9
Testing and Refinement	9
Pitfalls to Avoid	9
Metrics for Evaluation	9
Legal Aspects	9
Conclusion	10
Few-Shot Prompting: A Deep Dive	11
Chain-of-Thought Prompting: A Deep Dive	12
FAQs	13
Glossary	14
About the Author	15

Introduction

Welcome to this guide on Prompt Engineering for Educational Publishing and Assessment. If you've found your way here, you're likely intrigued by AI and large language models' possibilities in educational content creation. Perhaps you've had some experience with AI but are looking for a structured approach to harness its capabilities thoroughly. Look no further; this guide is your springboard into the transformative world of AI-driven educational publishing.

So, why is prompt engineering pivotal? In an age where AI is becoming increasingly proficient at generating human-like text, the adage "garbage in, garbage out" holds truer than ever. The quality of AI output is heavily influenced by the prompts you use. Crafting the perfect prompt is both an art and a science, and mastering it can significantly enhance the quality, relevance, and efficacy of AI-generated educational content.

In the following sections, you'll discover the what, why, and how of prompt engineering. From basic tasks and techniques to best practices and advanced strategies, this guide covers all facets, tailored specifically for educational publishers. Whether you're looking to summarise chapters, craft assessments, or develop interactive study guides, this guide offers actionable insights and examples to get you started on the right foot.

Who is This Guide For?

Are you an educational publisher intrigued by the capabilities of AI and language models like ChatGPT? Have you dabbled in AI but are unsure how to effectively integrate it into your professional workflows? If so, this guide is tailor-made for you. We aim to provide you with basic and advanced prompt engineering techniques. As the educational landscape evolves, mastering the art of effective prompting can be critical when creating high-quality educational content and assessments. This guide serves as your launching pad into the world of AI-aided publishing.

But First, What is ChatGPT?

ChatGPT is an advanced language model designed to process and understand natural language in a manner akin to human interaction. It relies on large-scale machine learning models trained on extensive datasets to generate text. By predicting the likelihood of the next word in a given context, ChatGPT serves as a sophisticated prediction machine, generating responses based on user prompts.

Purpose of the Guide

ChatGPT is far more versatile than its name suggests; it's not just for chat. Its applications are manifold, from sentiment analysis and content drafting to summarising and debugging

code. In educational publishing, it can aid in tasks ranging from blurb-writing and chapter summarisation to crafting reflective exercises.

But here's the catch: ChatGPT—or any AI model's effectiveness depends heavily on the quality of your prompts. This is where prompt engineering comes in. It's the art and science of crafting the perfect query to elicit the desired response from the AI model. Remember the age-old computing principle of "garbage in, garbage out"? It holds for AI as well. A well-designed prompt can lead to insightful and accurate output, whereas a poorly constructed one can result in irrelevant or erroneous responses.

Definition of a Prompt: What is Prompting, and Why is it Important?

Prompting gives an AI model a query or a set of instructions to generate a specific output. Think of it as steering the AI in a particular direction; you're guiding it to produce the desired information or response. This makes effective prompting crucial for obtaining accurate, relevant, and context-appropriate results from AI models like ChatGPT.

While simple prompts can yield good results, the devil is in the details. The quality of your output hinges on the amount and quality of information you include in your prompt. A well-crafted prompt will typically include:

- **Instruction:** The primary question or task you're asking the model to perform.
- **Context:** Additional background information that helps the model understand the query better.
- **Inputs:** Data or parameters that the model might need to generate the desired output.
- **Examples:** Illustrative scenarios that can guide the model towards the type of answer you're looking for.
- **Output:** Examples of what you want the model to generate, and in what format.

By carefully selecting and combining these elements, you can significantly enhance the quality of the model's output.

Basic Tasks You Can Get AI to Do

Text Summarisation

Text summarisation is a staple in natural language generation. It can condense articles or even entire chapters into concise summaries.

Example: Summarising critical takeaways from a chapter to create a revision guide or an overview for teachers.

Information Extraction

Language models are adept at generating text and can classify and extract specific information.

Example: Extracting key facts or dates from a history textbook for flashcards or quizzes.

Question Answering

Improving the format of your prompt can yield specific and accurate answers from the model.

Example: Creating a Q&A section at the end of a chapter by asking the model questions based on the chapter's content.

Text Classification

Beyond simple instructions, text classification allows you to categorise text into predefined labels.

Example: Categorising exercise questions by difficulty level or subject matter.

Conversation

Prompt engineering can guide a language model in conversational tasks, specifying its intent and identity.

Example: Designing an interactive study guide where the model converses with the student, offering explanations and posing questions.

Code Generation

Language models are increasingly effective at generating code, and this goes beyond just software development.

Example: Generating code snippets for a computer science textbook or interactive learning module.

Content Creation

Another task that could be valuable is the generation of original content, such as questions or exercises.

Example: Creating multiple-choice questions or essay prompts for assessments.

Translation

Language models can also assist in translating text between languages.

Example: Translating educational materials for international editions.

Limitations of AI

While AI offers many capabilities, it's crucial to acknowledge its limitations. Language models like ChatGPT might generate incorrect or misleading information, and they do not possess understanding or awareness. It's vital to cross-reference important facts and exercise human judgement in curating AI-generated content.

Structuring a Prompt: Exploring the Components in Detail

Instruction

Start with a specific task or instruction you want the model to perform. Repeating this instruction towards the end of the prompt can exploit the model's recency bias for more accurate results.

Prompt

Summarise the key points of this chapter. Ensure you highlight the most important concepts:

The many slight differences which appear in the offspring from the same parents, or which it may be presumed have thus arisen, from being observed in the individuals of the same species inhabiting the same confined locality, may be called individual differences. No one supposes that all the individuals of the same species are cast in the same actual mould. These individual differences are of the highest importance for us, for they are often inherited, as must be familiar to every one; and they thus afford materials for natural selection to act on and accumulate, in the same manner as man accumulates in any given direction individual differences in his domesticated productions. These individual differences generally affect what naturalists consider unimportant parts; but I could show, by a long catalogue of facts, that parts which must be called important, whether viewed under a physiological or classificatory point of view, sometimes vary in the individuals of the same species. I am convinced that the most experienced naturalist would be surprised at the number of the cases of variability, even in important parts of structure

Output

Key Points Summary:

Individual Differences: The chapter emphasizes that offspring from the same parents often exhibit slight variations, termed as individual differences.

Importance of Variability: These individual differences are crucial as they are often inherited and provide the basis for natural selection to act upon.

Scope of Variation: Contrary to the belief that variations usually occur in unimportant physiological traits, the chapter argues that even important structural parts can vary among individuals of the same species.

Role in Natural Selection: The inherited individual differences are the raw materials that allow for natural selection to accumulate traits, much like how humans selectively breed traits in domesticated animals.

Surprising Extent of Variability: The chapter contends that even seasoned naturalists would be surprised at the extent of variability, especially in important structural aspects of organisms.

Context

Additional context helps steer the model toward more accurate and relevant responses. Assigning a role to the AI model can further refine this context. Note - I won't provide outputs for all the sample prompts but you can copy and paste them into ChatGPT or another AI model and test these for yourself.

Prompt

You are an educational expert. Your task is to generate multiple-choice questions based on the following text:

[insert text]

Before generating questions, identify the main themes in the text.

Input Data

This could be a document or specific data you're interested in working with. You can enrich ChatGPT with additional capabilities, like search queries for more contextual answers. This can be done using plugins or else by copying and pasting data into the prompt.

Prompt

Based on this historical passage, extract key dates and events.

[insert text or data]

Output Indicator

Be explicit about the desired format of the output, whether it's paragraphs, bullet points, or tables. Providing a structural cue can help guide the model's response.

Prompt

Provide the summary in bullet-point form. Start with 'Here's a bulleted list of key points:'

How to Improve Your Prompt

Grounding the LLM - Using External Sources

Incorporate external, up-to-date, or specific information for more grounded responses from the model.

Prompt

Based on the following table of student performance statistics, generate an analysis.

Using Structured Text - The Role of Markdown and Syntax

Structured text formats like Markdown can help the model better understand your prompt, resulting in clearer and more refined output.

Step-by-Step - Making ChatGPT Show Its Work

Requesting the model to outline its thought process can improve reliability and allow for better scrutiny of its output.

Prompt

Create a multiple-choice question based on this text. Take a step-by-step approach, providing reasoning for each option.

[insert text]

Conditional Prompts

Conditional prompts allow you to specify particular conditions for the model's content generation, increasing the relevance of the output.

Prompt

Generate quiz questions based on this text. Limit the questions to topics related to World War II.

[insert text]

Multi-turn Prompts

These prompts involve an iterative conversation with the model, enabling more nuanced and complex tasks.

Prompt

Start by generating a question based on this paragraph. After that, provide four answer options. Then, indicate the correct answer and explain the reasoning.

[insert text]

Quality Assurance

The quality of AI-generated content is paramount, especially in educational settings. Beyond crafting effective prompts, it's advisable to have a review process in place. Use subject matter experts to validate the accuracy of AI-generated content. Running multiple iterations can also help in refining the output.

Accessibility

Inclusivity should be at the forefront of educational content. When using AI-generated content, ensure it's accessible to all, including individuals with disabilities. This could involve generating alternative text for images or ensuring documents are well structured and use correct formatting for headings and content to assist with accessibility.

Types of Prompting Techniques

Zero-Shot Prompting

In zero-shot prompting, the model generates output without prior examples, relying solely on the prompt's clarity and specificity.

Prompt

Create a multiple-choice question on photosynthesis

Output

Which of the following is a by-product of photosynthesis? A) Oxygen B) Carbon dioxide C) Nitrogen D) Hydrogen

Few-Shot Prompting

Few-shot prompting involves providing the model with a few examples to understand the task better. This technique is beneficial for complex tasks.

Prompt

Create a question based on the following example: 'What is the capital of France?'
Answer: 'Paris'. Now, create a question on American history.

Output

Question: Who was the first President of the United States?
Answer: George Washington

Chain-of-Thought Prompting

This technique breaks down complex tasks into smaller, logical steps, often combined with few-shot prompting for better results.

Prompt

Step 1: Identify the main theme of this passage:

To Sherlock Holmes she is always THE woman. I have seldom heard him mention her under any other name. In his eyes she eclipses and predominates the whole of her sex. It was not that he felt any emotion akin to love for Irene Adler. All emotions, and that one particularly, were abhorrent to his cold, precise but admirably balanced mind. He was, I take it, the most perfect reasoning and observing machine that the world has seen, but as a lover he would have placed himself in a false position. He never spoke of the softer passions, save with a gibe and a sneer. They were admirable things for the observer—excellent for drawing the veil from men's motives and actions. But for the trained reasoner to admit such intrusions into his own delicate and finely adjusted temperament was to introduce a distracting factor which might throw a doubt upon all his mental results. Grit in a sensitive instrument, or a crack in one of his own high-power lenses, would not be more disturbing than a strong emotion in a nature such as his. And yet there was but one woman to him.

Step 2: Summarise the theme.

Step 3: Generate a question based on the main theme.

Output

Step 1: The main theme of the passage revolves around the complex relationship between emotion and reason, particularly in the character of Sherlock Holmes. It delves into how Holmes views emotions as a hindrance to his analytical abilities, yet paradoxically, he is shown to have a unique regard for Irene Adler, whom he refers to as "THE woman."

Step 2: The theme explores the tension between emotion and reason in Sherlock Holmes, who is portrayed as a highly rational and analytical individual. While he generally disdains emotion, there is an exception in his subtle, complex feelings for Irene Adler.

Step 3: Question: How does Sherlock Holmes' view on emotions conflict with his regard for Irene Adler?

Best Practices and Prompt Engineering Tips for Crafting Effective Prompts

Clarity and Specificity

- **Be Clear and Specific:** Ensure your instructions are explicit to get focused responses.
- **Writing Clear Instructions:** Specify the desired format or structure for the output.
- **Specificity:** Detailed and descriptive prompts often yield better results.

Planning and Objectives

- **Plan:** Consider the learning objectives and outcomes before crafting prompts.
- **Start Simple:** Begin with straightforward tasks and gradually introduce complexity.

Customisation and Context

- **Customise for Your Students:** Adapt prompts to individual student needs, interests, and abilities.
- **Using Variables:** Utilise variables to make prompts adaptable for various educational settings.

Testing and Refinement

- **Test and Refine:** Experiment with different prompts and adjust based on AI responses.
- **Testing and Iterating:** Constant refinement is vital to achieving desired results.

Pitfalls to Avoid

- **Avoid Impreciseness:** Being too complex with prompts can lead to vague or unclear outcomes.
- **To Do or Not to Do:** Focus on specifying what the model should do rather than what it shouldn't.

Metrics for Evaluation

The efficacy of your prompts can be gauged using various metrics. Depending on the task, these could range from accuracy and precision for information extraction tasks to readability scores for content generation. For educational content, metrics for evaluating the outputs in this context should relate closely to the learning outcomes and intended learner for your content. Consistently evaluate these metrics to refine your prompt engineering efforts.

Legal Aspects

Consider legal implications, such as copyrights and data protection, when incorporating AI-generated content. Generally, the entity that prompts the model owns the generated

content, but it's always good practice to consult legal advice to navigate the complex landscape of AI and intellectual property.

Conclusion

As we conclude this guide, we hope you've found it an enlightening journey into the world of prompt engineering for educational publishing. The era of AI in education is not on the horizon; it's already here. By understanding and mastering the techniques of prompt engineering, you position yourself at the forefront of this educational transformation. You're now equipped with the tools, strategies, and best practices to craft meaningful, accurate, and high-quality educational content using language models like ChatGPT.

Remember, the key to effective AI integration is continual learning and adaptation. As you experiment with different prompts and techniques, don't hesitate to refine, iterate, and even revisit this guide. With your newfound knowledge, the potential applications are vast, limited only by your imagination and creativity.

Here's to crafting a smarter, more effective educational landscape, one prompt at a time!

Few-Shot Prompting: A Deep Dive

Few-shot prompting involves providing examples (known as "shots") to condition the model for a specific task. This technique has shown its effectiveness in a variety of domains and applications:

Why it Matters

While large language models have impressive zero-shot capabilities, they can struggle with more complex tasks. Few-shot prompting acts as in-context learning, steering the model towards better performance by providing demonstrations in the prompt.

Limitations and Variations

Few-shot prompting can vary by the number of shots given—1-shot, 3-shot, 5-shot, and so on. However, this technique has limitations in handling specific reasoning tasks, necessitating advanced prompt engineering for optimal results.

Use Cases in Educational Publishing

- **Natural Language Understanding:** Enhance sentiment analysis in student feedback.
- **Question Answering:** Improve the generation of Q&A sections in textbooks.
- **Summarisation:** Guide the model to produce concise and effective chapter summaries.

Prompt:

"Example 1: Question: 'What is the capital of France?' Answer: 'Paris'

Example 2: Question: 'Who wrote Romeo and Juliet?' Answer: 'William Shakespeare'

Now, create a question based on American history."

Chain-of-Thought Prompting: A Deep Dive

Chain-of-thought (CoT) prompting involves breaking down a complex task into smaller sub-tasks, often within a single or multiple prompts. This technique enhances the model's reasoning capabilities.

Benefits

CoT increases the accuracy and relevance of the model's responses by focusing on logical reasoning steps. This is particularly useful for tasks that require complex reasoning.

How it Works

The primary mechanism of CoT is solving a problem step-by-step, with each step based on logical reasoning. When combined with few-shot prompting, the model learns the task and the reasoning process behind it.

Performance Metrics

The effectiveness of CoT is more apparent in larger models and often scales with the number of parameters. This advanced technique complements standard and few-shot prompting for more complex tasks.

Prompt:

"Step 1: Identify the main themes in this history passage.

Step 2: Generate a question based on the main theme.

Step 3: Provide four answer options for the question.

Step 4: Indicate the correct answer and explain the reasoning."

FAQs

What is Prompt Engineering?

Prompt Engineering is creating effective queries or instructions to get specific responses or actions from an AI model.

How does Few-Shot Prompting differ from Zero-Shot Prompting?

Few-shot prompting involves providing examples to the model to improve its understanding of the task. In contrast, Zero-Shot Prompting relies solely on the prompt's instructions.

Can I use ChatGPT for tasks other than text generation?

ChatGPT can perform various tasks, including text summarisation, question answering, and even code generation.

Why is context important in a prompt?

Context helps the model understand the nuances and specific requirements of the task, leading to more accurate and relevant output.

What are Multi-turn Prompts?

Multi-turn prompts involve a series of interactions with the model to complete more complex tasks that may require clarification or additional inputs.

How can I improve the accuracy of AI-generated content?

Techniques like Few-Shot Prompting, Chain-of-Thought Prompting, and grounding the model with external data can significantly improve output accuracy.

Is it possible to customise prompts for different educational settings?

Yes, variables and conditional prompts can be used to adapt prompts for various educational contexts.

What should I avoid when crafting prompts?

Avoid being too vague or imprecise in your instructions. The more specific and clear you are, the better the output you can expect.

Glossary

Affordances: Additional capabilities or data can be used alongside a language model to improve output.

Chain-of-Thought Prompting: An advanced technique that breaks down complex tasks into smaller sub-tasks to improve the model's reasoning capabilities.

Context: Additional background or information provided in a prompt to guide the model's output.

Few-Shot Prompting: A technique that provides the model with a few examples to understand the task better.

Grounding: Incorporating external, up-to-date, or specific information to enhance the model's output.

LLM (Large Language Model): A machine learning model trained on vast amounts of text data to perform various tasks in natural language processing.

Markdown: A lightweight markup language used for formatting text, often used to structure prompts for better clarity.

Multi-turn Prompts: Prompts that involve back-and-forth conversation with the model for more nuanced tasks.

Output Indicator: Specific instructions about the desired format or structure of the model's output.

Prompt Engineering: The art and science of crafting queries or instructions to guide an AI model in generating specific outputs.

Zero-Shot Prompting: A technique where the model generates output based solely on the prompt without prior examples.

About the Author

With nearly a decade of experience at Cambridge University Press and Assessment, I have been deeply involved in educational publishing, especially developing textbooks for schools, digital educational material, and teacher training courses. My work has consistently focused on aligning educational products with market needs and rigorous academic standards.

Building on this extensive background, I am currently exploring the integration of Generative AI technologies into the educational publishing workflow. The goal is to harness AI's capabilities to enhance content quality and facilitate adaptive learning environments. As I navigate this juncture of traditional educational publishing and emerging AI technologies, I aim to contribute meaningfully to the field's evolution.

Contact

Niall McNulty

Email: niall.mcnulty@gmail.com

Tel: +27 76 739 1687

LinkedIn: <https://www.linkedin.com/in/niallmcnulty/>

Licensing Information

This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/). You are free to share, copy, distribute, and transmit the work under the following conditions:

- **Attribution:** You must give appropriate credit to the author, provide a link to the license, and indicate if changes were made.
- **ShareAlike:** If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.

For more details, please visit the [Creative Commons BY-SA 4.0 License webpage](https://creativecommons.org/licenses/by-sa/4.0/).