



# Large Language Model Prompt Writing Guide

Lapis AI Consults

# Large Language Model Prompting Guide

## Table of Contents

1. [Who is this guide for?](#)
2. [The Basics](#)
  - a. [What is an LLM](#)
  - b. [LLMs Accessing the Internet](#)
  - c. [The Mindset You Need](#)
3. [Prompting Best Practices](#)
4. [Prompting Strategies](#)
  - a. [Priming](#)
    - i. [How to Prime](#)
      1. [Contextual priming](#)
      2. [Expectation Setting](#)
    - ii. [Benefits of Priming](#)
  - b. [Role-Playing](#)
    - i. [How to role-play](#)
    - ii. [Benefits of Role Playing](#)
  - c. [Few-Shot Prompting](#)
    - i. [How to use Few-Shot Prompting](#)
    - ii. [Benefits of Few-Shot Prompting](#)
  - d. [Chain-of-Thought Prompting \(CoT\)](#)
    - i. [How to Prompt a Model for CoT](#)
    - ii. [Benefits of Chain-of-Thought](#)
  - e. [Meta-Prompting](#)
    - i. [How to Use Meta-Prompting](#)
    - ii. [Benefits of Meta-Prompting](#)
5. [Putting it All Together](#)
6. [Wrapping Up: Prompt Mastery](#)
7. [Food for Thought](#)
8. [Moving Forward](#)

## Who is this guide for?

This training manual is for anyone who wants to learn how to get the most out of advanced large language models (LLMs) like ChatGPT (GPT-3.5; GPT-4), Claude, or Gemini.

It's about going beyond basic interactions and learning proven techniques to elicit high-quality, relevant, and actionable outputs from these AI tools.

If any of the following resonate with you, this guide is for you:

1. You want to use LLMs to assist with practical tasks like writing, analysis, research, ideation, or problem-solving, and need to know how to prompt the AI effectively to get reliable and useful results.
2. You're curious about using LLMs for creative projects like storytelling, worldbuilding, or content creation, and want to learn techniques for steering the AI's outputs in desired directions.
3. You're considering incorporating LLMs into your work, projects, or products, and need to understand how to design prompts that produce high-quality, on-target outputs.
4. You've tried asking LLMs basic questions but often get irrelevant, generic, or unhelpful responses, and you want to learn how to get better outputs.
5. You want to dive deeper into the capabilities and limitations of LLMs, and learn strategies for getting the most out of your interactions while avoiding common pitfalls.

Whether you're a complete beginner or have some experience with LLMs, this guide will teach you proven prompting techniques to get higher-quality, more coherent, and more actionable outputs from your AI interactions. No technical expertise is assumed, and the strategies will be explained with clear examples across a range of use cases.

If you want to go from getting okay-ish responses from LLMs to receiving truly impressive, high-quality outputs that actually help you with your tasks and projects, this guide is for you. It's not just about chatting with AI - it's about mastering the art and science of prompting to make AI your partner.

## The Basics

### What is an LLM

LLMs are a type of artificial intelligence trained on vast amounts of text data, allowing them to understand, interpret, and generate human-like text.

Some key strengths of LLMs include:

1. **Broad knowledge:** LLMs are trained on diverse datasets spanning many topics, giving them a wide base of knowledge to draw upon in conversations.
2. **Natural language understanding:** LLMs can comprehend the meaning and context of human language, allowing them to engage in coherent dialogues and provide relevant responses.
3. **Language generation:** Beyond just understanding text, LLMs can also generate human-like text outputs, enabling them to assist with writing, answer questions, and perform other language tasks.
4. **Reasoning abilities:** Advanced LLMs can perform reasoning tasks, such as drawing inferences, making analogies, and engaging in logical deduction. They can use their vast knowledge to make connections and arrive at conclusions.
5. **Learning abilities:** LLMs are capable of learning new tasks from just a few examples, a capability known as few-shot learning. This allows them to adapt quickly to new contexts and requirements.

However, LLMs also have some key limitations:

1. **Lack of real-world experience:** LLMs' knowledge comes from their training data, not real-world experience. They may produce outputs that sound plausible but are factually incorrect or nonsensical.
2. **Inconsistency:** LLM outputs can be inconsistent across conversations or even within the same interaction. They may contradict themselves or forget earlier parts of the dialogue.
3. **Bias and hallucinations:** LLMs can reflect biases present in their training data and are prone to "hallucinating" or making up information to fill in gaps in their knowledge.

4. Limited reasoning: While LLMs can perform impressive reasoning feats, they can struggle with complex, multi-step problems or questions that require deep technical domain expertise (unless you train them on a specific topic)

There are many technical methodologies and programming approaches which can help mitigate these limitations but they require extensive AI software development knowledge and are outside the scope of this guide.

## LLMs Accessing the Internet

The ability for some LLMs to access the internet is a significant enhancement in AI. It allows them to provide information on recent events, evolving situations, and frequently updated topics that might not be covered in their training data. It vastly expands the knowledge available to them, potentially improving the relevance and accuracy of their responses.

Some prominent examples of services you can use which incorporate LLMs accessing the internet include [Microsoft Co-Pilot](#), [Perplexity AI](#) and [ChatGPT Plus](#) subscription.

However, it's crucial to remember that even with internet access, LLMs can still make mistakes or "hallucinate" information. They might give non-existent links as "sources." They may struggle to consistently distinguish reliable from unreliable sources on the web. The most recent information isn't always the most accurate, and LLMs may not always get this balance right.

Therefore, when interacting with LLMs that have internet access, it's important to maintain a critical eye. If the LLM provides a piece of information and cites a web source, it's a good practice to click on that citation and verify the information yourself. Don't take the LLM's output at face value, especially for important or sensitive topics.

In the context of prompting, you can leverage the internet access capability by asking an LLM with access to the internet to seek out specific, current information to inform its responses. You can also instruct it to provide citations for key facts or figures. However, your prompts should also guide the LLM to prioritize reputable sources and to acknowledge when there might be uncertainty or conflicting information on a topic.

## The Mindset You Need

To get the most out of your interactions with LLMs, here are some key points to keep in mind:

- Experiment and explore: Different approaches will work better for different tasks and goals. Be willing to try different prompting strategies, to combine techniques, and to think creatively about how to leverage the LLM's capabilities.
- Fact-check: If you're using the LLM to generate factual content or to inform important decisions, always verify the information against reliable sources.
- Be realistic: An LLM is not magic, though it can appear quite miraculous on what it can do. It is not a silver bullet for every problem you have and is intended to be an aid - not some sentient, omniscient, omnipotent being.
- It's a collaborative process: Think of the interaction with the LLM as a collaboration rather than just giving orders to a machine. You bring your human knowledge, judgment, and objectives to the table. The LLM brings its vast knowledge and language processing capabilities. Work together to achieve your goals.

Remember, prompting is a skill that improves with practice. As you interact more with LLMs, you'll develop an intuition for what works well and what doesn't in different contexts. Stay curious, pay attention to what yields good results, and refine your approach over time.

With the right mindset and techniques, you can unlock the immense potential of LLMs to assist with many different tasks and creative projects. The strategies and examples in the following sections will give you a strong starting point to explore this exciting technology.

## Prompting Best Practices

- Be clear and specific: LLMs are powerful but they're not mind readers. The clearer and more specific you are with your prompts, the better the outputs you'll get. Take the time to think through what exactly you are trying to achieve.

Avoid vague or ambiguous language. Instead of saying "Write about dogs," state "Write a 200-word informative article about the health benefits of owning a dog."

- Iterative refinement: Don't expect perfect results on the first try. Working with LLMs is an iterative process. If the output isn't quite what you're looking for, try rephrasing your prompt, providing more context, or breaking the task down into smaller steps.
- Leverage the LLM's strengths: LLMs excel at tasks involving language understanding, generation, and reasoning. Lean into these capabilities in your prompts. Ask the LLM to summarize, explain, compare, brainstorm, and analyze.
- Provide context: The more context you can provide in your prompt, the better the LLM will be able to assist you. Don't assume it knows what you know. If specific knowledge or constraints are relevant to the task, spell them out explicitly.

This could include key facts, constraints, goals, or perspective you want the model to adopt.

- Use concrete examples to illustrate what you want: Examples help clarify your intent and give the LLM a template to follow. If you want the LLM to write in a specific style, provide a sample of that style. If you're asking it to perform a task, show it an example of the desired output (more on this below in the section "Few-Shot Prompts").

For instance, if you want a professional email response to a customer complaint, you could start with "Write a response to a customer complaint email in the style of this example: [provide example]."

- Break down complex tasks into a series of simpler steps: If your task is complicated or involves multiple parts, guide the LLM through it step-by-step. This makes the process more manageable for the model and allows you to steer the output at each stage.

For example, if you're asking the LLM to write a story, you could start with "Generate an outline for a short story in the mystery genre," then "Develop the main character based on this outline," and so on.

- Specify the format or structure in which you want the output: If you have a specific format in mind for the LLM's output, make that explicit in your prompt.

This could include things like word count, tone of voice, intended audience, or document structure. For instance, "Generate a 500-word blog post aimed at new parents with the following sections: Introduction, Challenges, Solutions, Conclusion."

# Prompting Strategies

## Priming

Priming is a technique used to influence an LLM's output by providing it with a specific context or framework before giving it the main prompt. This can be useful for guiding the model towards a particular style, tone, format, or domain of knowledge.

### How to Prime

There are several ways to prime an LLM:

#### Contextual priming

In this type of priming, you are simply giving the LLM as much background information you want it to take into consideration for it to give you the output you need. Think of it as if you were giving instructions to a person who knows nothing about you, your situation or what you know/don't know.

The example on the next page is a demonstration of this:



You:

Consider the following scenario:

You are a project manager overseeing the development of a new mobile app. The app is designed to help users manage their personal finances. The project is currently in the testing phase, and you've received feedback from beta users indicating some usability issues with the app's navigation.

Now, please outline a plan for addressing these usability issues and improving the app's navigation.

AI:

Here's a plan for addressing the usability issues with the app's navigation:

1. **Gather specific feedback:** Reach out to a subset of beta users to get more detailed feedback on the navigation problems they encountered. Ask for specific examples and pain points.
2. **Conduct usability testing:** Set up user testing sessions where new users interact with the app while thinking aloud. Observe where they struggle with navigation and note any common patterns.
3. **Analyze the issues:** Work with the UX team to analyze the feedback and testing results. Identify the root causes of the navigation problems (e.g., unclear labels, hidden menu items, too many steps).
4. **Brainstorm solutions:** Conduct a design brainstorming session with the UX and development teams. Generate a range of ideas for improving the navigation, such as simplifying the menu structure, adding search functionality, or providing clearer onboarding.
5. **Create new prototypes:** Have the UX team create new navigation prototypes that address the identified issues. Focus on creating a more intuitive and efficient navigation experience.
6. **Conduct follow-up testing:** Test the new prototypes with another round of user testing. Gather feedback to see if the navigation improvements are effective.
7. **Iterate and refine:** Based on the new testing feedback, make further iterations and refinements to the navigation design.
8. **Implement changes:** Once the navigation design is finalized, work with the development team to implement the changes in the app.
9. **Communicate with users:** Let beta users know that their feedback was heard and that improvements have been made based on their input. This helps build goodwill and engagement.
10. **Monitor and adjust:** After launching the updated app, continue to monitor user feedback and analytics related to navigation. Make further adjustments as needed based on real-world usage patterns.

By following this user-centered, iterative approach, we can systematically improve the app's navigation and create a better overall user experience.

As you can see from the output above, the LLM has given a comprehensive response.

However, not all of this might be applicable to your specific situation. For example, you might not have the time or the budget for user testing where they “think out loud.” Use this output as a starting point with your next prompt going deeper to set budget and time constraints.

Remember, the LLM cannot read your mind!

### Expectation Setting

In this form of priming, you are letting the model know ahead of time what will be the purpose of the prompt and what you will expect from the response.

Think of it like a job description, though it does not need to be quite as detailed.

See the example on the next page:

You:

Let's discuss strategies for effective time management in a business setting. Please provide tips that are applicable to a variety of roles and industries, and focus on techniques that are backed by research or expert opinion.

What are some of the most common time management mistakes people make at work, and how can they be avoided?

AI:

Time management is a challenge that many professionals face, regardless of their role or industry. Here are some of the most common time management mistakes and how to avoid them:

1. **Failing to prioritize:** One of the biggest mistakes is treating all tasks as equally important. This leads to wasted time on low-value activities. To avoid this, use techniques like the Eisenhower Matrix to prioritize tasks based on urgency and importance. Focus on tasks that are both urgent and important first.
2. **Multitasking:** Research has shown that multitasking can actually decrease productivity and increase errors. Instead of trying to do multiple things at once, practice single-tasking. Give your full attention to one task at a time, and you'll likely complete it faster and with better quality.
3. **Not taking breaks:** It may seem counterintuitive, but not taking breaks can actually hurt your productivity. Your brain needs time to rest and recharge. Use techniques like the Pomodoro Method - work for 25 minutes, then take a 5-minute break. After 4 cycles, take a longer break. This can help maintain your focus and energy levels.
4. **Failing to delegate:** Trying to do everything yourself can lead to burnout and missed deadlines. Learn to delegate tasks to others when appropriate. This frees up your time for high-priority tasks that require your specific skills or expertise.
5. **Not setting clear goals:** Without clear goals, it's easy to get sidetracked or waste time on unimportant tasks. Set SMART goals (Specific, Measurable, Achievable, Relevant, Time-bound) for your day, week, and month. This helps you stay focused and motivated.
6. **Procrastinating:** Putting off important tasks can lead to last-minute rushes and poor-quality work. To avoid procrastination, break large tasks into smaller, manageable steps. Set deadlines for each step and hold yourself accountable.
7. **Overcommitting:** Taking on too many tasks or projects can lead to stress and decreased productivity. Learn to say no when necessary. Before taking on a new commitment, consider whether you realistically have the time and energy to complete it.

Remember, effective time management is a skill that takes practice. Start by implementing one or two of these strategies and gradually build on them over time. And don't forget to be kind to yourself - no one is perfect, and it's okay if you don't always stick to your plan perfectly. The goal is to make steady progress towards more efficient and effective use of your time.

Here you set the expectation by telling the LLM the purpose of the prompt (time management strategies) and what you will expect from the response (applicable to a variety of roles, etc)

## Benefits of Priming

- Enhanced Relevance: By setting the context or expectations, priming ensures that the model's outputs are more directly applicable to your specific needs or scenarios.
- Building More Robust Dialogue Systems: In conversational AI, priming can help maintain the context over longer interactions, making dialogue systems more robust and capable of handling extended conversations.
- Maintaining focus: LLMs have a tendency to go off on tangents in certain situations when they pick up on secondary threads in the conversation. By narrowing the scope, you reduce the chance of off-topic responses.

## Role-Playing

Role-playing is a powerful strategy for using LLMs to explore different perspectives, generate ideas, or test out scenarios. By asking the model to assume a specific role or persona, you can guide it to provide responses that are tailored to that viewpoint or context.

This can be particularly useful in a business setting, where you might want to consider a problem or situation from multiple angles. By having the model take on the role of different stakeholders or experts, you can get a more diverse range of insights and ideas.

### How to role-play

Role-playing with a language model involves setting up scenarios where the model assumes a specific character or role, tailoring its responses to fit that identity. To engage a model in role-playing, you should clearly define the character's background, personality traits, and the context of the interaction. This sets the stage for the model to adopt the role and generate responses that are consistent with the character's duties and demeanor.

Here's an example of how you might use role-playing in a business context:

You:

Let's imagine you are a seasoned Human Resources Manager at a large tech company. The company is considering implementing a new employee wellness program, but there are concerns about the cost and potential impact on productivity.

From the perspective of an experienced HR Manager, what are some key considerations and recommendations you would provide to the company's leadership regarding this decision?

AI:

As a seasoned Human Resources Manager, I would approach the question of implementing a new employee wellness program with a balance of employee well-being and business objectives in mind. Here are some key considerations and recommendations I would provide to the company's leadership:

1. Evaluate the long-term benefits: While there may be upfront costs associated with implementing a wellness program, it's important to consider the long-term benefits. Healthier employees tend to be more productive, have lower absenteeism rates, and have lower healthcare costs over time. I would recommend conducting a cost-benefit analysis to assess the potential long-term ROI of the program.
2. Align with company culture and values: The wellness program should align with the company's culture and values. If the company places a high value on employee well-being and work-life balance, a wellness program could be a natural fit and a way to demonstrate that commitment to employees.
3. Get employee input: Before designing the program, I would recommend getting input from employees about what kinds of wellness initiatives would be most valuable to them. This could be done through surveys, focus groups, or employee resource groups. Getting employee buy-in early can help with participation rates later.
4. Start small and iterate: To manage costs and assess impact, I would suggest starting with a smaller-scale pilot program. This could focus on a specific department or a limited set of wellness offerings. Based on the results and feedback from the pilot, the program can then be refined and scaled up over time.
5. Measure and communicate impact: It will be important to establish clear metrics for measuring the impact of the wellness program, both in terms of employee health outcomes and business metrics like productivity and healthcare costs. Regularly communicating these impacts to leadership and employees can help build support for the program over time.
6. Consider partnerships: To potentially reduce costs and expand offerings, I would explore partnerships with local gyms, wellness providers, or health insurance companies. These partners may be able to offer discounted services or help with program administration.
7. Ensure equal access: When designing the program, it will be important to ensure that all employees have equal access to the benefits, regardless of their role, location, or health status. This may require offering a range of options or providing accommodations for employees with specific needs.

Overall, while there may be some short-term costs and productivity impacts to consider, I believe that a well-designed employee wellness program can have significant long-term benefits for both employee well-being and the company's bottom line. By starting small, getting employee input, and carefully measuring impact, we can create a program that enhances our culture and supports our most valuable asset – our people.

In this example, the prompt sets up a specific role (a seasoned Human Resources Manager) and a context (considering implementing a new employee wellness program) for the model to respond from. The model then provides a response that considers the situation from the perspective of an HR Manager, focusing on key considerations and recommendations relevant to that role.

## Benefits of Role Playing

- Context-Specific Accuracy: When assuming specific roles, such as a seasoned Customer Service Manager in Retail, the model delivers more accurate and relevant responses tailored to the field and from the viewpoint of that persona.
- Greater Adaptability to Diverse Scenarios: Role playing allows the model to adapt to a variety of contexts. This flexibility is beneficial in multi-functional platforms where the model needs to handle diverse inquiries, ranging from technical support to personal advice.
- Improved Empathy and Emotional Intelligence: The model can produce outputs that are more empathetic and emotionally intelligent. This can lead to more sensitive and appropriate interactions, as well as more human-like responses.

## Few-Shot Prompting

One of the most powerful capabilities of advanced LLMs is their ability to learn from just a few examples. This is known as few-shot learning. In the context of prompting, this means you can teach the LLM to perform a new task by providing just a handful of representative examples in the prompt itself. This is a game-changer for many applications, as it allows you to customize the LLM's behavior without needing large datasets or complex fine-tuning.

### How to use Few-Shot Prompting

Let's take a look at a real-world example:

You want to use an LLM to perform sentiment analysis on customer reviews. You want the model to classify each review as either positive, negative, or neutral. Here's how you could use few-shot prompting to achieve this:

Step 1. Provide a small number of representative examples:

```
1. Review: I absolutely loved this product! It exceeded my expectations
in every way. I would highly recommend it to anyone.

Sentiment: Positive

2. Review: This product was okay, but nothing special. It did the job but
there are better options out there.

Sentiment: Neutral

3. Review: I was extremely disappointed with this product. It broke after
just a few uses and the customer service was unhelpful.

Sentiment: Negative
```

Step 2. Include both the input (the review text) and the ideal output (the sentiment label) for each example. This shows the LLM the task you want it to perform and the format you expect.

Step 3. After providing the examples, give a new input and ask the model to generate the output:

```
4. Review: I was pleasantly surprised by the quality of this product. It works
better than I expected and the price is reasonable.

Sentiment:
```

The LLM, having learned from the examples, will then generate the appropriate sentiment label for this new review:

```
Sentiment: Positive
```

The key to effective few-shot prompting is to choose your examples carefully. They should be representative of the task you want the LLM to perform and be clear in their input-output mapping. The examples should also cover the range of possible outputs. In the sentiment analysis case, we provided an example of each sentiment category (positive, neutral, negative).

The number of examples needed can vary depending on the complexity of the task, but often just 3 well-chosen examples can be enough for the LLM to grasp the pattern. Experiment with different numbers of examples to see what works best for your specific use case.

## Benefits of Few-Shot Prompting

- Rapid Adaptation and Learning: Few-shot prompting enables language models to quickly adapt to new tasks or domains with only a few examples. This flexibility allows users to leverage AI for a wide variety of applications without extensive retraining.
- Data Efficiency: This method is highly efficient in terms of data usage. It requires only a small number of examples to guide the model, making it ideal for situations where data is scarce or expensive to obtain.
- Cost-Effectiveness: By reducing the need for large-scale data collection and lengthy training periods, few-shot prompting lowers the barrier to entry for utilizing advanced AI technologies, making it more accessible and cost-effective for businesses, organizations and researchers.

## Chain-of-Thought Prompting (CoT)

Chain-of-Thought reasoning involves explicitly guiding the model to unpack complex questions into a series of logical, sequential steps. This process helps the model—and the user—navigate from the initial query to the final answer through a clear, traceable path. It's like telling the AI to “show your work” in a math problem, providing not just the answer but also the rationale behind it.

With this strategy, you are guiding the AI to think the way a person might when they are tasked with a challenging problem: take things step by step.

## How to Prompt a Model for CoT

To effectively use Chain-of-Thought prompting (CoT), you should structure your prompt to encourage the model to reveal its thought process explicitly. Here are some tips:

- Start with a clear question: Define the problem or question as clearly as possible.



- Ask for step-by-step reasoning: Explicitly request the model to describe each step of its thought process.
- Encourage detailed explanations: Prompt the model to justify each step with reasons or supporting details.

Here's an example of what you can do with CoT:

Your company is considering whether to implement a four-day workweek. The leadership team is interested in understanding the potential impacts on productivity and employee satisfaction.

The conversation on the next page shows the prompt you might use.

```
You:
Consider a company thinking about switching to a four-day workweek.
Break down the reasoning step-by-step to analyze how this change might
affect productivity and employee satisfaction.

AI:
Step One: Define Productivity and Employee Satisfaction:
Productivity can be measured by the output per hour worked, and employee
satisfaction by employee engagement and retention rates.

Step Two: Consider the Impact on Work Intensity:
With one less workday per week, employees might work more intensely during
the four days, potentially maintaining or even increasing overall productivity.

Step Three: Analyze Employee Satisfaction:
A shorter workweek could lead to higher job satisfaction and well-being,
as employees have more time for personal activities and family, which
could also reduce burnout.

Step Four: Weigh Potential Challenges:
However, there might be challenges in coordinating schedules, meeting client
demands, and maintaining communication. It's important to consider whether
the industry and specific job roles are suitable for this model.

Step Five: Suggest a Trial Period:
Implementing a trial period could help measure the actual impacts on
productivity and employee satisfaction before making a permanent change.
```

The model's response provides a systematic exploration of the question, offering insights into both the potential benefits and the challenges. Each step builds on the previous one, creating a comprehensive analysis that can aid decision-making.

Remember, you need to explicitly tell the model to do this by saying “Think step by step”. You can also give prompts like “take it step by step” or “reason through each step” but you have to be explicit.

## Benefits of Chain-of-Thought

- Greater transparency into the model's thinking: You can see exactly what logic the model used, making the output more interpretable and understandable.
- Easier to identify and debug errors in reasoning: If the model makes a mistake, you can pinpoint exactly where in the reasoning process the error occurred. This makes it easier to provide corrective feedback.
- More trustworthy outputs: When you can see and verify the model's reasoning, you can have more confidence in its conclusions. This is especially important for high-stakes business decisions.

This strategy provides a deeper and more interactive way of solving problems with AI and research shows it significantly increases accuracy, creativity, and reasoning capabilities of an LLM.

## Meta-Prompting (a.k.a. Self Prompting)

Meta-prompting involves a layer of abstraction where the primary interaction with the model is not to get a direct answer but to obtain a well-crafted prompt. This prompt can then be used either by the same model or another system to generate the best possible outputs. It is particularly useful in complex scenarios where the direct question might be too broad or vague, and where a well-defined prompt can lead to more focused and useful responses.

## How to Use Meta-Prompting

To implement meta-prompting effectively:

- Define the Objective Clearly: Clearly articulate the goal or what you hope to achieve with the output. The more specific the objective, the better the model can tailor the prompt.
- Use Iterative Refinement: Often, the first set of generated prompts may need refinement. Use an iterative approach where the model's suggestions are fine-tuned through multiple interactions.
- Evaluate and Select: Once a set of prompts is generated, evaluate their potential effectiveness based on past performance and select the most promising ones for actual use.

Here is an example of how to do this:

You:

```
Generate a prompt that will help a customer service AI consistently deliver empathetic and solution-oriented responses to customers complaining about service delays.
```

AI:

```
Prompt: When a customer expresses frustration about a service delay, acknowledge their feelings, apologize sincerely, and outline the steps being taken to resolve the issue. Ask if they require further assistance or have additional concerns.
```

## Benefits of Meta-Prompting

- Targeted Responses: Meta-prompting helps in crafting prompts that are specifically designed to elicit targeted and high-quality responses.
- Efficiency and Effectiveness: By optimizing the prompts, the model can produce more relevant and effective outputs, reducing the need for further refinements.
- Adaptability: Meta-prompting allows businesses to adapt their AI interactions to different contexts and objectives, enhancing the versatility of their AI tools.

This strategy can be a great starting point for any interaction with an LLM or used at any point throughout the conversation if you're feeling like it is not meeting your needs.

## Putting it All Together

The prompting strategies we've covered - basic prompting, priming, role-playing, few-shot prompting, chain-of-thought prompting, and meta-prompting - are powerful tools on their own. However, the real magic happens when you start combining these techniques to tackle complex, multi-faceted problems.

By leveraging multiple prompting strategies in a single interaction, you can guide the LLM to provide more comprehensive, nuanced, and context-aware responses. This can be especially

valuable in business settings, where problems often require considering multiple perspectives, weighing trade-offs, and synthesizing information from various domains.

Let's walk through an example that showcases how multiple prompting techniques can be combined to address a complex business challenge on the next page:

You:

We will be discussing monetization strategies for Search.com.

**Background information:**

Search.com, is a privacy-focused search engine that summarizes web results using website categories and provides a personalized experience through a combination of search results, apps, and shortcuts. It is designed to put its users first and provide them with the most relevant and up-to-date information available on the web. It also provides users with access to a zero-trace private browser and the ability to customize their search results with 1000 apps.

**Current monetization:**

Search.com currently generates revenue through affiliate links. They earn a part of the profit for some of the products in the Walmart app. Affiliate relationships do not affect the way they display or rank products in any way, and they plan to keep it that way.

Search.com prides itself on not selling any advertisement space as a way to generate revenue. The focus is solely on the end user and this is what makes Search.com unique.

**Prompt:**

Imagine you are the Chief Revenue Officer of Search.com with deep expertise in monetization strategies, understanding business challenges, competitive analysis and data analysis.

What other monetization strategies can Search.com use? In providing recommendations, consider the following:

- The target audience and their needs/preferences
- Search.com's unique features and capabilities
- Different monetization models and potential revenue streams (excluding ads)
- The pros and cons of each proposed strategy
- The feasibility of implementing each strategy

Please provide at least 3 options and explain your reasoning step-by-step. Use your expertise to analyze the business context, competitive landscape, and available data to support your recommendations.

Think through this problem methodically and show your chain of thought.

#### Option 2: Paid API access for businesses

- Offer paid API access to Search.com's search and categorization capabilities
- Target businesses that need curated web data for market research, competitor analysis, etc.
- Charge based on API call volume or offer tiered subscription plans
- Pros: High-margin revenue stream, leverages existing tech assets, B2B model diversifies revenue
- Cons: Requires dedicated sales and support for API customers, may be seen as less aligned with consumer focus
- Feasibility: Medium - would require building out API infrastructure and B2B go-to-market capabilities

#### Option 3: Affiliate revenue expansion

- Expand existing affiliate model beyond Walmart to include more retailers and product categories
- Leverage search and app data to optimize affiliate link placements and targeting
- Potentially charge retailers for prime app placement or "featured" status
- Pros: Builds on existing affiliate competencies, highly targeted based on user data
- Cons: May be seen as compromising search objectivity, revenue potential limited by user purchase intent
- Feasibility: High - could be implemented with current capabilities

#### Option 4: Premium content subscriptions

- Curate or create premium content in high-interest categories (e.g. financial analysis, industry research)
- Offer this content to users on a subscription basis, bundled with the premium search features
- Leverage Search.com's categorization capabilities to surface relevant premium content to users
- Pros: Diversifies revenue beyond search, provides unique value to users, reinforces position as content curator
- Cons: Requires investment in content creation or licensing, may dilute focus on core search product
- Feasibility: Medium - would require building out content curation and subscription management capabilities

#### Step 5: Evaluate and prioritize options

Based on the analysis, I would recommend prioritizing Options 1 (Freemium subscriptions) and 3 (Affiliate expansion) in the near term. These build on Search.com's existing strengths and could be implemented relatively quickly.

Option 2 (Paid API access) is a promising longer-term play to diversify revenue, but would require more upfront investment. Option 4 (Premium content) is intriguing but riskier, and should probably be tested in a limited way before full commitment.

```
The key will be to closely monitor user feedback and engagement metrics as these new monetization strategies are rolled out. We want to ensure we're striking the right balance between revenue generation and staying true to Search.com's user-first ethos. Continual iteration and adjustment will likely be needed.
```

```
Next steps would be to develop detailed financial projections and implementation roadmaps for the prioritized options, and to socialize these recommendations with key stakeholders for input and buy-in. With the right execution, I believe these monetization strategies could significantly accelerate Search.com's revenue growth while preserving its unique value proposition to users.
```

This prompt combines several strategies we discussed.

1. Priming - Setting expectations and giving context
2. Role-playing - Assigning the role of a Chief Revenue Officer
3. Chain of thought - telling it to think step -by-step and show its work

It should be noted that part of the prompt was also created via meta-prompting. Further, you will see that this combination of prompt strategies did not include few-shot prompts. This is because providing example outputs that are very long will get cumbersome for the model, will not enable for as much flexibility and creativity in output and adds to more technical complexity and limitations. Long output examples such as these are better reserved for a method called fine-tuning.

## Wrapping Up: Prompt Mastery

Congratulations on making it through this guide on advanced prompting techniques for large language models! Let's recap some of the key points we've covered:

- LLMs are powerful tools that can be leveraged for a wide range of business applications, from content creation to strategic planning to customer support.
- To get the most out of LLMs, it's essential to move beyond simple question-answer interactions and employ more sophisticated prompting strategies.
- Basic prompting best practices, like being clear and specific, providing examples, and breaking down complex tasks, can significantly improve the quality and relevance of LLM outputs.

- Advanced techniques like priming, role-playing, few-shot prompting, chain-of-thought, and meta-prompting allow you to guide LLMs towards more complex, context-aware, and creative outputs.
- Combining multiple prompting techniques can be especially powerful for tackling multi-faceted business problems that require considering various perspectives, making tradeoffs, and applying domain expertise.

Effective prompting is an iterative process that requires experimentation and continuous learning from interactions with the LLM.

However, it is important to remember that LLMs have limitations and can make mistakes. Think of LLMs as a powerful tool and thought partner, but not a replacement for human judgment and expertise.

As you continue your journey with LLMs, stay curious and keep experimenting. Share your learnings with others, and stay up to date on the latest developments in this rapidly evolving field. With the right approach, LLMs can become an indispensable part of your problem-solving toolkit.

Remember, the key to success with LLMs is not just about crafting the perfect prompt, but about developing a deep understanding of their capabilities and limitations, and continually refining your prompting strategies based on feedback and results.

So go forth and put these advanced prompting techniques into practice! Start small, iterate often, and don't be afraid to try novel approaches. The potential of LLMs is limited by your creativity and willingness to experiment.

## Food for Thought

The applications of LLMs in business are vast and still being discovered. As you start incorporating these advanced prompting techniques into your work, you'll likely find many opportunities to leverage LLMs for automation, insights generation, decision support, and more.

Every day, we all learn something new about this technology which is evolving at an exponential rate. There are many technical capabilities people thought were impossible just months ago that are now becoming best practice.

This is why it's crucial to continue working with AI and do what you can to stay on top of major developments. This will help you evolve as AI evolves.



## Moving Forward

There are many prompting strategies that have been discovered in just the past few months and new ones are being discovered each day. It would be impossible to cover all of them in one introductory guide.

If you're interested in learning more prompt strategies and how to use AI to take your endeavors to unprecedented heights, **contact Lapis** for an informational session on [Personalized Corporate and Individual Trainings](#).

Happy prompting!

*-Olga Topchaya*  
Founder, CEO

Connect on [LinkedIn](#)

---

Ready to advance your company even further with AI?

Book a [Free Consultation](#)  
to explore tailored AI solutions today!

---

[Lapis AI Consults](#) is a full service AI agency, helping businesses innovate their products and workflows to stay ahead in their industry. Our services include:

- ❖ AI Integration
- ❖ Custom AI Software Development
- ❖ AI Business Strategy
- ❖ AI Training