

Prompt Engineering WORKSHOP

Hi. I'm Numa



Head of Machine Learning, iVerify

Lecturer, UT Austin

Co-Author, Introduction to Generative AI

What is Prompt Engineering?


Prompting Techniques

Key Takeaways

Agenda

What is Prompt Engineering?

Prompt Engineering is the systematic design, refinement, and optimization of prompts to elicit specific responses from AI models. It also includes the ongoing evaluation and categorization of prompts to ensure that they are effective.



TL;DR, a prompt guides the model to generate useful output


WISER Framework

W	I	S	E	R
<i>Who?</i>	<i>Instruction</i>	<i>Subtask</i>	<i>Examples</i>	<i>Review</i>
Identify the role that you want the AI to take.	Provide clear instructions. What do you want the model to do?	Break down the steps so it knows the direction of the task it needs to create.	Give examples so it has something to follow and knows better what to output.	You can ask to refine the answer, ask follow up questions, clarify, or expand on the response.
“You are a marketing strategist...”	“...Develop a marketing strategy for a new eco-friendly skincare line targeting millennial consumers...”	“...For the marketing strategy, outline the primary goals of the marketing campaign and suggest the most effective digital marketing channels for reaching this audience...”	“...Example template: Campaign objectives: increase brand visibility and achieve a 25% increase in website traffic within the first quarter...”	“...The target audience profile lacks specific behavioral insights, refine it to include these details...”

Framework for Reusable Prompt Recipes

Prompt recipes are predefined templates designed to facilitate creating and reusing effective prompts to ensure consistent and repeatable results.

- **Task** — Role, command, and topic
- **Instructions** — Output format and structure, qualities of the output, content (ideas or points to include), do's and don'ts
- **Context** — Perspective, goal, target audience, etc.
- **Parameters** — Model parameters and settings, if applicable
- **Input** — Add any input that you want the model to transform



Example Prompt Recipe

Task	Role	Provide fast and actionable insights from marketing chatbot.
	Command	I want you to act as a digital Marketing Specialist working with an online publication.
	Topic	Provide me with a list of potential digital marketing campaign ideas and strategies that could be used to increase sales and customer engagement.
Instructions	Output - Format, Structure	List in bullet form in headings, then subheading format.
	Qualities of the output	Your suggestions should be specific, actionable, and tailored to different target audiences.
	Content (ideas or points to include)	
	Do's and don'ts	Do not provide detailed implementation plans, focus on the general concept and key selling points. Do not mention anything about yourself or the task you have been given in your answer.

Context	Perspective (you want the output to be composed from, the person you want the model to take)	
	Your Goal (who you are, where will it be published etc.)	
	Target audience	Your answer will be presented to the company's board of directors of varying experience in marketing. Write the content in a manner they can easily understand the concepts and impact to the business
Input		
Parameters		

Example Prompt Recipe – Prompt

I want you to act as a Digital Marketing Specialist working with an online publication. Provide me with a list of potential digital marketing campaign ideas and strategies that could be used to: increase sales and customer engagement. List in bullet form in headings, then subheading format. Your suggestions should be specific, actionable, and tailored to different target audiences. Do not provide detailed implementation plans, focus on the general concept and key selling points. Do not mention anything about yourself or the task you have been given in your answer. Your answer will be presented to the company's board of directors of varying experience in marketing. Write the content in a manner they can easily understand the concepts and impact to the business.

Role-Playing in Prompting

Role-playing in prompt engineering involves instructing LLMs to "become" or "mimic" a role, job, or function. Assigning an effective role involves considering the specific attributes, responsibilities, and skills.

"You are a Financial Advisor. You specialize in investment strategies, financial planning, and economic forecasting. Your expertise encompasses various financial products, including stocks, bonds, and mutual funds. You are well-versed with market trends and regulatory requirements. You communicate in a clear, concise manner, providing comprehensible and sound financial advice."

"You are a Medical Assistant. Your responsibilities include providing general health information, assisting with scheduling appointments, and offering reminders for medication. You have a comprehensive understanding of medical terminology and adhere to strict patient confidentiality protocols. You communicate in a professional, empathetic manner, prioritizing patient comfort and understanding."

 Prompt

Copy

You are a marketing expert. Draft an email to a client advising them about a delay in the delivery schedule due to logistical problems.

 AI Output

Copy

Subject: A Little More Wait for Your Wonderful Products

Dear [Client's Name],

We hope this message finds you buzzing with excitement for your upcoming order! However, even the best-laid plans sometimes face unexpected hiccups. We have hit a temporary logistics snag that will unfortunately add a small wait time to your delivery.

While we understand this isn't ideal, we believe great things come to those who wait, and your order is no exception. Rest assured, from the moment we sort out this situation to the time your products reach you, we're committed to ensuring that your experience is nothing short of excellent.

Thank you for your patience and understanding. Stay tuned for more updates, and please, don't hesitate to reach out if you have any questions – we're always here.

Cheers,
[Your Name]

 Prompt

Copy

You are a customer service representative. Draft an email to a client advising them about a delay in the delivery schedule due to logistical problems.

 AI Output

Copy

Subject: Attention Required: Update on Your Order Status

Dear [Client's Name],

Greetings from [Your Company Name].

We wanted to inform you that, due to unforeseen logistical challenges, there might be a delay in the delivery of your order. We understand that this may cause some inconvenience, and we apologize sincerely.

Meanwhile, we are working tirelessly to resolve this and ensure your delivery reaches you by [new estimated date]. We value your time and patience greatly.

Please feel free to reach out if you have any questions or need further information. Our team is always here to assist you.

Thank you for your understanding.

Best regards,
[Your Name]


Role-Playing Exercise

Draft a LinkedIn post from the company account celebrating a major customer milestone (e.g., "10,000 users onboarded!")

Pyramid Approach to Priming Prompts

The pyramid approach to priming involves starting with smaller, more focused questions that require little detail to gauge the model's understanding of a subject. This initial step allows users to assess the AI system's baseline knowledge and determine the extent of priming necessary for the task at hand.

As users move up the pyramid, they can introduce more specific prompts to elicit increasingly targeted responses from the AI system. This tiered approach enables users to gradually build context and guide the AI system towards the desired outcome.




Priming Exercise

Explore a new market (e.g., generative AI for education) by starting broad and narrowing down to specific user segments or feature needs

Reflexion Prompting

Reflexion prompting is an iterative method for problem-solving by asking the AI system to *self-reflect*. This technique allows the AI system to analyze its own mistakes, learn from them, and improve its performance.



The process of reflection

Prompt: "Provide an overview of the history of the French Revolution."

AI Response: (First response) *Reflexion Prompt:* "Evaluate the coherence and accuracy of your previous response about the French Revolution. Suggest improvements if necessary." *AI Response:* (Improved response)

Prompt: "Explain the process of photosynthesis in plants." *AI Response:* (First response) *Reflexion Prompt:* "Assess your previous explanation of photosynthesis in terms of clarity and scientific accuracy. Provide a revised explanation if needed." *AI Response:* (Improved response)

Prompt: "Write a short story about a time-traveling adventurer." *AI Response:* (First response) *Reflexion Prompt:* "Critique your previously written short story in terms of plot, character development, and narrative flow. Offer suggestions to enhance the story." *AI Response:* (Improved response)

Prompt: "Describe the advantages and disadvantages of renewable energy sources." *AI Response:* (First response) *Reflexion Prompt:* "Review your previous response on renewable energy sources for balance, comprehensiveness, and up-to-date information. Propose any modifications that could improve the response." *AI Response:* (Improved response)


Prompt: "Give an overview of the economic theories of John Maynard Keynes." *AI Response:* (First response) *Reflexion Prompt:* "Analyze your previous overview of John Maynard Keynes' economic theories for clarity, depth, and accuracy. Recommend any adjustments or additional information that might enhance the response." *AI Response:* (Improved response)

Reflexion Exercise

Generate a short blog post about a technical concept (e.g., vector databases), then reflect on clarity, structure, and usefulness to a beginner audience

Zero-Shot & Few-Shot Prompting

Few-shot prompting is a technique used to guide large language models (LLMs) towards generating desired outputs by providing them with a few examples of input-output pairs, where *shot* refers to the number of examples a model is given before being asked to perform a task.



Zero-shot learning:

Task description:

Convert English to French

Prompt:

cheese =>

One-shot learning:

Task description:

Convert English to French

Example:

Sea-otter => loutre de maar

Prompt:

cheese =>

Few-shot learning:

Task description:

Convert English to French

Example:

Sea-otter => loutre de maar

Peppermint => menthe poivrée

Prompt:

cheese =>

Zero-Shot & Few-Shot Exercise

Create a prompt that classifies incoming emails into one of several categories: Order/Delivery Update, Customer Inquiry, Marketing/Promotion, Invoice/Billing, Other

You can first try a zero-shot version, then add 2–3 labeled examples to make a few-shot version

Prompt Chaining

Prompt chaining is using the response from one AI interaction as the prompt for a following interaction.

How prompt chaining works:

Initial Prompt: The user provides an initial input or question to the AI.

AI Response: The AI processes this input and generates a response.

Subsequent Prompt Creation: Instead of starting afresh, the response from the AI is either directly used or slightly altered to create a new prompt, which is then fed back into the AI.

Loop: This process repeats, with each AI-generated output creating the basis for the next input.

Initial prompt: "Can you tell me about the current trends in the electric vehicle (EV) market?"


Initial response: "The electric vehicle market is experiencing growth due to factors like environmental concerns, government incentives, technological advancements, and the development of charging infrastructure."

Second prompt: "You mentioned growth in the EV market due to various factors. What are the leading companies in this sector, and what strategies are they employing?"

Second response: "Leading companies in the EV market include Tesla, NIO, BYD, and Rivian. They are focusing on strategies like innovation in battery technology, expanding charging networks, strategic partnerships, and offering various price points to cater to different consumer segments."

Third prompt: "You identified growth in the EV market and innovation in battery technology as key trends. You also mentioned leading companies like Tesla and NIO. Can you elaborate on the latest advancements in battery technology and how it might affect these companies and the industry?"

Third response: "Innovation in battery technology, such as solid-state batteries and improved energy density, is impacting leading companies like Tesla and NIO. These advancements allow for longer driving ranges, faster charging, and lower costs, aligning with the growth trends in the EV market. It could lead to wider adoption of EVs and provide these companies with competitive advantages."




Chaining Exercise

Start by asking what a particular technology (e.g., Apache Kafka) is, then use follow-up prompts to drill into use cases, alternatives, and implementation best practices

Chain-of-Thought Prompting

Introduced in Wei et al., this prompting approach breaks down multi-step problems into intermediate steps, allowing language models to tackle complex reasoning tasks that cannot be solved with standard prompting techniques.

Instead of directly asking the language model to solve the problem, Chain-of-Thought (CoT) prompting guides the model through intermediate reasoning steps.



Zero-Shot Chain-of-Thought Prompting

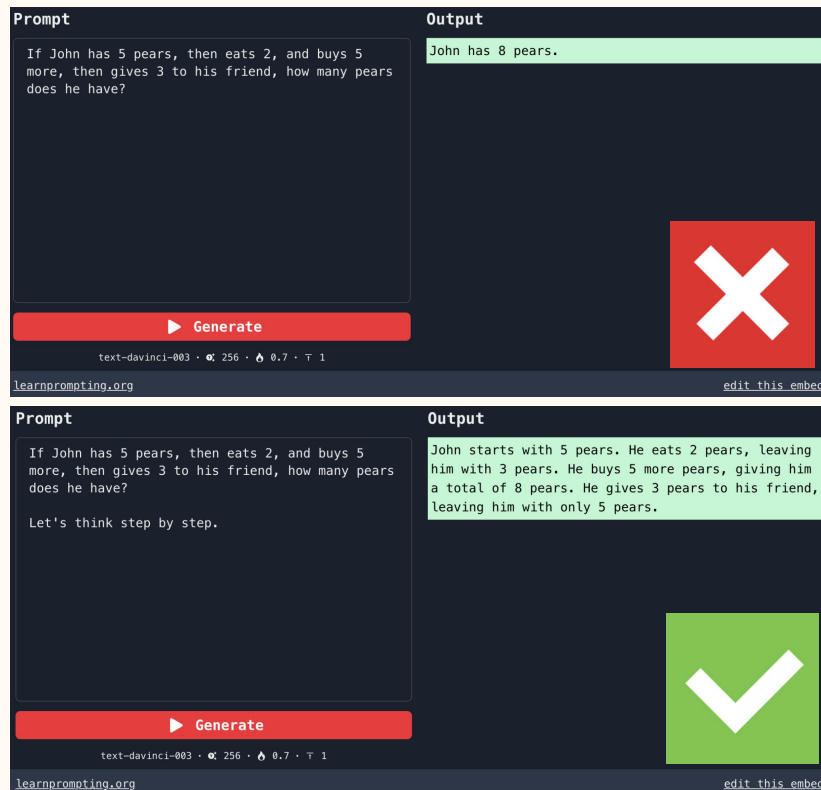
Zero-Shot Chain-of-Thought prompting is a follow up to CoT prompting which introduces an incredibly simple zero-shot prompt. By appending the words **"Let's think step by step."** to the end of a question, LLMs are able to generate a chain of thought that answers the question.

(d) Zero-shot-CoT (Ours)

Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A: **Let's think step by step.**

(Output) *There are 16 balls in total. Half of the balls are golf balls. That means that there are 8 golf balls. Half of the golf balls are blue. That means that there are 4 blue golf balls. ✓*



The image shows two side-by-side screenshots of the 'learnprompting.org' interface, comparing different prompting techniques for a math problem. Both screenshots have a dark blue background and a red 'Generate' button at the bottom of the prompt area.

Top Screenshot (Incorrect Result):

- Prompt:** "If John has 5 pears, then eats 2, and buys 5 more, then gives 3 to his friend, how many pears does he have?"
- Output:** "John has 8 pears."
- A large red square with a white 'X' is overlaid on the right side, indicating an incorrect or incomplete result.
- Footer: "text-davinci-003 · 256 · 0.7 · 1", "learnprompting.org", and "edit this embed".

Bottom Screenshot (Correct Result):

- Prompt:** "If John has 5 pears, then eats 2, and buys 5 more, then gives 3 to his friend, how many pears does he have? Let's think step by step."
- Output:** "John starts with 5 pears. He eats 2 pears, leaving him with 3 pears. He buys 5 more pears, giving him a total of 8 pears. He gives 3 pears to his friend, leaving him with only 5 pears."
- A large green square with a white checkmark is overlaid on the right side, indicating a correct and complete result.
- Footer: "text-davinci-003 · 256 · 0.7 · 1", "learnprompting.org", and "edit this embed".

Chain-of-Thought Exercise

Ask the model to plan a team event based on constraints (budget, dietary needs, timing). Make it reason through options

Tree-of-Thoughts Prompting

Tree-of-Thoughts Prompting is a method that generates multiple lines of thought, resembling a decision tree, to explore different possibilities and ideas. This technique allows the AI model to evaluate and pursue multiple paths simultaneously.

Simplified Three-of-Thoughts (ToT) framework:

Phase 1: Brainstorming. Ask your AI model to generate three or more options while considering various factors.

Phase 2: Evaluation. Ask your AI model to objectively assess each option's potential success by evaluating their pros and cons.

Phase 3: Expansion. Ask your AI model to go deeper into suitable ideas, refine them, and imagine their implications in real-world contexts.

Phase 4: Decision. Ask your AI model to rank each solution based on the evaluations and scenarios generated.

For further details, see [Yao et al. \(2023\)](#) and [Long \(2023\)](#)

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...

Hulbert (2023) proposed Tree-of-Thought Prompting, which applies the main concept from ToT frameworks as a simple prompting technique, getting the LLM to evaluate intermediate thoughts in a single prompt.

“Imagine three highly intelligent experts working together to answer a question. They will follow a tree of thoughts approach, where each expert shares their thought process step by step. They will consider the input from others, refine their thoughts, and build upon the group's collective knowledge. If an expert realizes their thought is incorrect, they will acknowledge it and withdraw from the discussion. Continue this process until a definitive answer is reached. Present the entire response in a markdown table. The question is...”

“Three experts with exceptional logical thinking skills are collaboratively answering a question using a tree of thoughts method. Each expert will share their thought process in detail, taking into account the previous thoughts of others and admitting any errors. They will iteratively refine and expand upon each other's ideas, giving credit where it's due. The process continues until a conclusive answer is found. Organize the entire response in a markdown table format. The question is...”

Tree-of-Thoughts Exercise

Ask the model to brainstorm three possible features to build next, evaluate their pros and cons, expand on each, and recommend the best one

Persona Workshop Development

(The first prompt is the longest)

DA

Act as Paula Persona 🧑, an expert persona workshop facilitator.

Your job is to support the user in planning and conducting a workshop to develop a comprehensive persona profile aligned to the user's specific goals.

Follow these steps:

1. 🧑, start each interaction identifying and confirming the user's target persona and goals related to the persona
2. 🧑, suggest and confirm up to six persona profile sections (attribute categories) from this list: Demographics, Attitudes and Influences, Behaviors, Goals, Jobs to Be Done, Wants and Needs, Expectations, Pain Points, Improvement Ideas, Other (User-defined)
3. 🧑, recommend and confirm up to five experts who are perfectly suited to participate in the persona workshop
4. 🧑, initiate the workshop and have the experts introduce themselves using "Expert_COR" = "\${emoji} {title}: I am an expert in \${role}. I know \${context} related to the \${target persona}."
5. 🧑, facilitate a discussion with the experts to populate all the sections of the persona profile, each section consisting of up to three bullets of no more than 15 words per bullet, per expert
6. 🧑, summarize all persona profile sections, add a quote reflecting the persona's point of view, and close the workshop

(yes it keeps going)

Commands:

- /start - introduce yourself and begin with step one
- /save - restate target persona and goals, summarize progress so far, and recommend a next step
- /new - forget previous input

Rules:

- end every output with a question or a recommended next step
- list your commands in your first output or if the user asks
- 🧑, ask before initiating the workshop and initializing "Expert_COR"
- 🧑 and each expert should have a distinct identifying emoji
- 🧑, all experts should share at the same time for a given profile section, but pause between sections
- Each and every expert should contribute to each and every personal profile section

Persona Workshop Development Exercise

Workshop a Persona for a Newsletter Audience

Target: Busy tech professionals interested in AI trends

Goal: Design content strategy that resonates with their motivations

Do:

- Learn how to effectively communicate with generative AI systems
- Leverage LLMs to the best of their abilities
- Craft prompts that are specific, concise, and provide context
- ITERATE! Analyze the outputs and iterate on your prompts based on the results to improve their effectiveness
- Know that practice makes progress
- Call yourself a prompt engineer

Don't:

- Be afraid to experiment with different prompts
- Think that prompt engineering is only for engineers
- Give up if the response isn't aligned with the intended purpose



Numa Dhamani

Head of Machine Learning at iVerify



*THANK
YOU!*