
AI CHATBOTS AS RESEARCH ASSISTANTS: SMARTER PROMPTS, BETTER OUTCOMES

PART 2: ADVANCED PROMPT PATTERNS AND CONTROL STRATEGIES

BLAŽ ZUPAN

A HANDS-ON COURSE ON PROMPT ENGINEERING AT BAYLOR COLLEGE OF MEDICINE
HOUSTON, MAY 2025



Welcome to Part 2 of our hands-on course on AI Chatbots as Research Assistants. In the last session, we learned about basic prompt engineering strategies, expressing communication constraints and providing roles, and imposing structure in the AI answers. Here, we upgrade these techniques and explore advanced patterns that give you more control, structure, and creativity in your interactions with large language models. From turning ChatGPT into a tutor, slide designer, or co-author, to using templates, reasoning chains, and flipped roles.

REVERSE ROLES: QUIZ CONSTRUCTION

Consider, for a starting example, turning ChatGPT into a tutor:

I am a natural science student and have just started learning about machine learning, AI, and large language models. Please ask me some questions to see if I understand basic concepts from this field. Turn this Q&A into a game, where the questions must be choice questions, and I answer with ABC. Go one question at a time and adjust the questions according to my knowledge.

Sounds like a fun way to learn! Let's start with a basic question and see how you do. Here we go:

Question 1:

What is a key characteristic that distinguishes machine learning from traditional programming?

- A. Machine learning relies on manually written rules.
- B. Machine learning systems learn from data.
- C. Machine learning models are always interpretable.

What's your answer (A, B, or C)?

B

Correct! Machine learning systems learn from data, which differentiates them from traditional programming where rules are explicitly programmed by humans.

Question 2:

Which of the following is an example of a supervised learning task?

- A. Clustering customers into different segments.
- B. Predicting house prices based on historical data.
- C. Discovering hidden patterns in data without labels.

What's your answer (A, B, or C)?

A

Not quite. Clustering customers into different segments is an example of an unsupervised learning task. The correct answer is B. Predicting house prices based on historical data is a supervised learning task because it involves learning from labeled data.

Let's move on to the next question.

...

This could, of course, go on and on, but we can also change the types of the questions, or difficulty, or focus just by instructing ChatGPT using, say, one of the following inputs:

Please change the questions to focus more on AI and less on machine learning?

Please come up with a simpler set of questions.

Could you focus more on large language models and ethics?

In the class, we also showed that we could turn this Q&A session into a web-based program supported by Python scripts and YAML-formatted files with questions and answers (not included here, but you can see the video in the lecture if you have a particular interest in the program design). Here, we would not like to dive into this, considering computer programming out-of-scope of this course, but here is a prompt if you would like to try this on your own:

Let's create code for this in Python, but I want it to be a web application with FastAPI. It should have a single HTML page that asks multiple-choice questions. Create five sample questions, store them in a YAML file, and allow me to extend the game by editing it.

FLIPPED INTERACTION PATTERN: GRANT PROPOSAL SKELETON

In this interaction pattern, we assign the role of interviewer to the language model. Rather than generating text from a fixed prompt, the model actively gathers information by asking specific, easy-to-answer questions. Once enough data is collected, it uses the responses to generate a structured output—such as a grant proposal skeleton with a title, abstract, and section outline. This example demonstrates how prompting can turn the model into a collaborative assistant that builds content through guided conversation. Here is an example conversion:

I would like you to write a skeleton of a grant proposal. You will ask me questions about my current research work to get the data about it. Questions should be specific and easy to answer. When you have enough information, you will show me the title of the proposal and abstract and provide the outline with one sentence describing the content of each proposal section. Please do not ask me more than seven questions. Now, ask me the first question.

What is the general topic or area of your current research work?

Single-cell data analysis.

Great! Could you briefly describe the specific focus or objective of your research within single-cell data analysis?

I would like to determine the level of synchronicity of the amoeba cells during the development.

Thank you for that information. Could you specify what types of data you are currently collecting or analyzing to study the synchronicity of amoeba cells during development?

Single-cell gene expression data at various stages of development.

Thank you. How are you currently processing or analyzing the single-cell gene expression data in your research?

Using the methods and software I will develop within my project. I need to handle batch effects, dimensionality reduction, data visualization and annotation of visualizations with explanations come from external knowledge bases.

Thank you for the information. How are you planning to validate the level of synchronicity among the amoeba cells based on the processed single-cell gene expression data?

Computationally. We have yet to develop methods to assess and visually depict synchronicity. Our hypothesis is that synchronicity increases through time and during specialization of cells.

Thank you for your responses. Based on the information provided, here is the proposed grant proposal:

Title: Quantifying Synchronicity in Amoeba Cell Development through Computational Analysis of Single-Cell Gene Expression Data

Abstract: This proposal aims to develop novel computational methods to quantify and visualize the level of synchronicity among amoeba cells during developmental stages using single-cell gene expression data. The project will focus on addressing challenges such as batch effects, dimensionality reduction, and integrating external knowledge bases for data interpretation. The hypothesis is that synchronicity increases over time and during cell specialization, which will be validated computationally.

Outline:

1. Introduction

Introduce the significance of studying synchronicity in amoeba cell development.

2. Literature Review

Review current methodologies in single-cell data analysis and existing knowledge about amoeba cell development.

3. Methods

Develop computational methods for handling batch effects in single-cell data.

Implement dimensionality reduction techniques tailored for single-cell gene expression data.

Design algorithms for quantifying synchronicity and visualizing temporal changes.

Integrate external knowledge bases to annotate visualizations and interpret results.

4. Validation

Validate the developed methods using simulated data and benchmark datasets.

Apply the methods to real single-cell gene expression datasets from amoeba cells.

5. Results and Discussion

Present results on the quantification of synchronicity during amoeba cell development.

Discuss findings in relation to the hypothesis and existing literature.

6. Broader Impacts

Discuss potential implications of the research for understanding cellular development and synchronization.

Outline potential applications in biomedical research and beyond.

7. Conclusion

Summarize the significance of the proposed research and future directions.

This skeleton outlines the main sections and content of your grant proposal based on the information provided. If you need further details or adjustments, feel free to let me know!

We may not like the proposed skeleton, but note that this may be just the beginning. We may now request to change it, instruct to refine it, and ask for more data-gathering questions. Or even find a paper like that from the *PLOS Computational Biology* series on “Ten Simple Rules” (on grant writing or structuring our proposal), add it to our prompt, and ask ChatGPT to follow the recommendations.

PERSONAL ASSISTANT PROMPT: SLIDE DESIGN THROUGH CONVERSATION

In a sense, we have already been introducing this type of pattern in the section above, where we are instructing LLMs to ask us about the information that is needed to fulfill a particular task. Here, we combine a prompt of this type with instructions for a particular pattern that the answers should use. Here is an example prompt:

You are my assistant in helping me to design a scientific presentation and related slides. I want to compose up to 5 slides for a topic and mainly include figures and only some text. For each slide, you will help to define the following: **Title****: <TITLE> ****Subtitle****: <ONE SENTENCE SUBTITLE> ****Figure or text****: <SHORT DESCRIPTION OF FIGURE OR TEXT TO INCLUDE> To do this, you'll need to ask me questions about my presentation and research. Ready? Ask one question at a time—no more than seven questions. Propose the slides only after you have finished with the questioning. Now, ask me the first question.**

The text enclosed in double asterisk (“**”) marks the part to be printed in bold in the markdown language. Again, Markdown is a lightweight markup language with plain text formatting syntax designed to be converted to HTML and other formats, and is the language that LLMs like ChatGPT’s models are using for textual output that is then rendered in HTML in a chat window. With the text enclosed in “<” and “>” we, somehow, tell LLM that this should be replaced with generated text. Note that this is not algorithmically “encoded” in LLM, but simply results in this kind of marking that was used in the textual documents that LLM have consumed in training. We here hope that our communication will be understood in this way, since else, we need to change the prompt.

Also notice that our prompt finished with “Now, ask me the first question.” This is to let LLM know that we do not want to print all the questions at once, and to reemphasize that we would like to answer one question at a time.

Here is an example of the conversation that resulted from such a initial prompt:

Sure! Let's get started. What is the main topic or focus of your scientific presentation?

Machine learning and visualization

Great! Could you please specify the particular aspect of machine learning and visualization you would like to focus on? For example, are you concentrating on a specific algorithm, application, dataset, or technique?

Point-based visualizations, machine learning techniques for dimensionality reduction, like t-SNE, PCA and MDS

Got it! What is the primary objective or key message you want to convey through your presentation on point-based visualizations and machine learning techniques for dimensionality reduction?

That point-based visualization are a cool way to presenting multivariate data, especially if we can, in some way, explain what particular groups of points actually mean. Emphasis on explainable machine learning and visualizations! :)

Excellent! Which datasets or specific examples do you plan to use to demonstrate these techniques and their effectiveness in your presentation?

Data from biology and molecular biology, like the machine learning zoo data sets and single-cell gene expression data sets

That sounds very interesting! Could you specify any particular results or key findings from your analysis of these datasets that you would like to highlight in your presentation?

Yes. It turns out that when the data is large and there are many data instances, t-SNE nicely uncovers clustering structure. This needs to be explained by other information sources, like data bases on marker genes.

Great! What types of figures or visualizations do you already have, or plan to create, to support these key findings? For example, do you have specific plots, charts, or diagrams in mind?

I do. t-SNE, PCA and MDS visualizations and annotations of t-SNE plots.

Perfect! Could you also share any preferences for the overall design or style of your presentation slides? For instance, do you have a specific color scheme, layout, or template in mind?

No, just plain black and white. Oh, I forgot something: I will include snapshots of Orange Data Mining toolbox to illustrate visual analytics, and should spend a slide on what visual analytics is all about.

Thank you for the additional information! Here's a quick summary of what we have so far:

- 1. Introduction to Machine Learning and Visualization**
- 2. Techniques for Dimensionality Reduction**
- 3. Visual Analytics Concept**
- 4. Examples and Results from Biological Datasets**
- 5. Using Orange Data Mining Toolbox**

I will now propose the slides based on this information:

Slide 1: Introduction to Machine Learning and Visualization

Title: Introduction to Machine Learning and Visualization

Subtitle: An overview of point-based visualizations for multivariate data

Figure or text: Include a brief introduction text and a visual representation of multivariate data, such as a simple scatter plot showing high-dimensional data points.

Slide 2: Techniques for Dimensionality Reduction

Title: Techniques for Dimensionality Reduction

Subtitle: Exploring t-SNE, PCA, and MDS for data visualization

Figure or text: Display side-by-side visualizations of the same dataset using t-SNE, PCA, and MDS, highlighting the clustering structures.

Slide 3: Visual Analytics Concept

Title: What is Visual Analytics?

Subtitle: Combining data analysis with interactive visualizations

Figure or text: Provide a brief explanation of visual analytics and a snapshot of the Orange Data Mining toolbox interface, showing interactive features.

Slide 4: Examples and Results from Biological Datasets

Title: Examples and Results from Biological Datasets

Subtitle: Clustering structures in single-cell gene expression data

Figure or text: Show t-SNE plots of the single-cell gene expression data with annotations explaining clusters using marker genes and additional information sources.

Slide 5: Using Orange Data Mining Toolbox

Title: Using Orange Data Mining Toolbox

Subtitle: A tool for visual analytics and machine learning

Figure or text: Include screenshots of the Orange Data Mining toolbox in action, demonstrating how to create and interpret visualizations.

Would you like to make any adjustments or add any specific details to these slides?

This is a good start. For brainstorming, I mean. Reading through the text we can spot that it would be, for instance, better to show Orange Data Mining with examples of various visualizations, and correct ChatGPT that we would like to merge this slide with the slide two. This conversation stems from brainstorming the author of these lines had with ChatGPT when preparing the slides for “The Joy of Point-Based Visualizations”, a recently given invited talk presenting the techniques mentioned above. The talk was longer and needed more slides, but this could be done with simple prompting, such as

Please expand the number of slides to 10. Also note that I would like to introduce various visualizations by showcasing Orange on different data sets, including the zoo data set, where I show how similar PCA, MDS, and t-SNE results are, and an example single-cell data set, where it becomes clear that with large data sets with higher number of data instances the resulting visualizations can be very different.

Ok, this is a rather long prompt, but it worked and the resulting proposal for the slides (not shown here) was great. In my presentation, I also needed to show how mathematically very different are the three dimensionality reduction techniques, and asked ChatGPT to render their optimization function.

Please give me the mathematical expression for the optimization function for PCA, MDS and t-SNE. Just math, no explanation.

Certainly! Here are the mathematical expressions for the optimization functions:

PCA (Principal Component Analysis)

$$\arg \max_{\mathbf{w}_k} \mathbf{w}_k^\top \mathbf{S} \mathbf{w}_k \quad \text{subject to} \quad \mathbf{w}_k^\top \mathbf{w}_k = 1$$

where \mathbf{S} is the covariance matrix of the data.

MDS (Multidimensional Scaling)

$$\min_{\mathbf{Y}} \sum_{i < j} (d_{ij} - \|\mathbf{y}_i - \mathbf{y}_j\|)^2$$

where d_{ij} is the distance between points i and j in the original high-dimensional space, and \mathbf{y}_i and \mathbf{y}_j are the corresponding points in the lower-dimensional space.

t-SNE (t-Distributed Stochastic Neighbor Embedding)

$$\min_{\mathbf{Y}} \sum_{i \neq j} P_{ij} \log \frac{P_{ij}}{Q_{ij}}$$

where P_{ij} is the probability that high-dimensional points i and j are neighbors, and Q_{ij} is the probability that their low-dimensional representations are neighbors.

Kind of amazing that this works, right? We have already discussed generation of equations by ChatGPT (see our previous lecture notes) and learned that this is just another sequence of letters and textual commands, this time in LaTeX, that is then conveniently rendered and typeset in the browser, so, technically, no surprises there. We could go on here with personal assistant prompt, but you got the idea, and this section got long already.

TEMPLATE PATTERN: Q&A CONSTRUCTION

Prompt templates are predefined structures or patterns used to guide the language model's text generation. These templates can help in generating more consistent and targeted outputs.

I am going to give you a template for your output. CAPITALIZED WORDS are my placeholders. Fill in my placeholders with your output. Preserve the overall formatting of my template. My template is: **Question: QUESTION **Answer:** ANSWER Do not do anything at this stage. I will give you the data to format in the next prompt. Provided the data, you will create five questions using my template. Ok?**

Sure, I understand! Just let me know when you're ready to proceed with the data.

The data: Impossibly frustrating: why Mission: Impossible 8 was a major letdown Audiences may have flocked to the Final Reckoning but Tom Cruise's latest mission ends up being incredibly underwhelming If the title is sincere, and this really is the final reckoning, then it's been a franchise of two halves. Mission: Impossible diehards tend to underrate the first half (which ran from Brian De Palma's brisk 1996 original to 2011's fun Ghost Protocol) as much as they overrate the second (which launched with 2015's Rogue Nation). Yet the rumbles and grumbles emanating from public screenings suggests a disgruntled consensus is forming around the concluding instalment: that this is an altogether disjointed way to resolve the affairs of Ethan Hunt and his IMF crew, and a shaky

way to ignite the movie summer season. Ninety minutes in which nothing happens over and over again, followed by 70 minutes of M:I B-roll. To better diagnose this latest glitch in the Hollywood machine, we need to return to the relighting of the fuse. This was the franchise to which Tom Cruise retreated in the wake of the commercial underperformance of 1999's *Eyes Wide Shut* and *Magnolia* – the two most rigorous turns of this star's career, films in which Cruise allowed himself to be rattled and be seen to be rattled, only to be met with widespread public and awards-circuit indifference. The *Mission: Impossible*s, by contrast, would be the sort of 4DX-coded sure thing for which audiences have routinely turned out, a creative safe space, even as the films' constituent set pieces pushed their prime mover into performing ever riskier business to ensure bums on seats. In those early films, the character of Hunt was as much martyr and marked man as saviour or secular saint, targeted at every turn by directors with comparably forceful visions. The sensationalist De Palma revelled in the set-up's potential for spectacle; and while, in retrospect, the motorbike-and-mullet combo of 2000's *M:I 2*, directed by John Woo and set to a bruising Limp Bizkit beat, was bound to date rapidly, the sometime animator Brad Bird, in *Ghost Protocol*, had the bright idea of turning the series into a live-action cartoon, with Cruise defying gravity and nature alike by hanging off the side of the Burj Khalifa and personally outrunning a sandstorm. The last four films, however, bear the imprint of screenwriter turned director Christopher McQuarrie, who concluded that what this series needed was a little more conversation, overseeing the construction of a vast story framework for his star to dangle off one-handed. That approach reaches its apotheosis in *The Final Reckoning*, but the scaffolding now overwhelms the spectacle. The attempt to solder eight films together ends in much-rewritten incoherence – see Ving Rhames's confused sendoff – and, worryingly, results in missions being described rather than shown. You wonder whether the insurers blanched after Cruise crocked an ankle shooting 2018's *Fallout*; now we're left with folks talking at length in nondescript rooms. Is this a *Mission: Impossible* movie, as advertised, or some M:I-themed podcast?

Now please create three questions using my template.

Question: What is the main criticism of *Mission: Impossible – Dead Reckoning Part*

Two mentioned in the review?

Answer: The film is described as disjointed and underwhelming, with long stretches where nothing happens and key action replaced by dialogue-heavy scenes.

Question: How do fans typically view the earlier and later halves of the *Mission: Impossible* franchise?

Answer: Fans tend to underrate the first half (1996–2011) and overrate the second half (starting with 2015's *Rogue Nation*).

Question: What shift in tone or focus does the review highlight in the recent *Mission: Impossible* films under Christopher McQuarrie?

Answer: McQuarrie's direction emphasizes intricate story frameworks and conversations over visual spectacle, which dilutes the impact of the action.

While we are almost used that the creation of such Q&A from some short text is possible, the fact that AI can do it looks (to the authors of these lines) immensely impressive.

TEMPLATE PATTERN: ENTITY IDENTIFICATION AND PROFILING

Let us combine the template with further analysis of the data. Suppose we are interested in people involved in the assassination of archduke Franz Ferdinand in Sarajevo but would like to instruct LLM to provide an output in specific form:

I am providing a template for your output. CAPITALIZED words are placeholders. Preserve the overall formatting. My template:

Bio: <NAME>

****Summary:** <ONE SENTENCE SUMMARY>**

****Born:** <YEAR OF BIRTH>, <PLACE>**

****Major impact:** <ONE SHORT PARAGRAPH OF MAJOR ROLE OF THIS PERSON>**

Extract the names of all the people mentioned in the following text and report on them using the template.

<<<In the summer of 1914, the sovereigns of Europe were woven together by treaties, alliances, as well as secret agreements. The Triple Alliance (1882) encompassed the German Empire, Austria, and Italy.[34]

On 28 June 1914, Archduke Franz Ferdinand of Austria, heir presumptive to Emperor Franz Joseph I of Austria, visited Sarajevo, the capital of the recently annexed Bosnia and Herzegovina. Cvjetko Popović, Gavrilo Princip, Nedeljko Čabrinović, Trifko Grabež, and Vaso Čubrilović (Bosnian Serbs) and Muhamed Mehmedbašić (from the Bosniaks community),[35] from the movement known as Young Bosnia, took up positions along the route taken by the Archduke's motorcade, to assassinate him. Supplied with arms by extremists within the Serbian Black Hand intelligence organisation, they hoped his death would free Bosnia from Austrian rule, although there was little agreement on what would replace it.[36]

Nedeljko Čabrinović threw a grenade at the Archduke's car and injured two of his aides, who were taken to hospital while the convoy carried on. The other assassins were also unsuccessful but, an hour later, as Ferdinand was returning from visiting the injured officers, his car took a wrong turn into a street where Gavrilo Princip was standing. He fired two pistol shots, fatally wounding Ferdinand and his wife Sophie.[37] Although Emperor Franz Joseph was shocked by the incident, political and personal differences meant the

two men were not close; allegedly, his first reported comment was "A higher power has re-established the order which I, alas, could not preserve".[38]>>>

Bio: Archduke Franz Ferdinand of Austria

Summary: Archduke Franz Ferdinand of Austria, heir presumptive to Emperor Franz Joseph I, was assassinated in 1914, an event that triggered World War I.

Born: 1863, Graz, Austria

Major impact: Archduke Franz Ferdinand's assassination in Sarajevo by Gavrilo Princip of the Young Bosnia movement on June 28, 1914, sparked a series of events leading to the outbreak of World War I. His death destabilized Europe, setting off a chain reaction of alliances and conflicts that culminated in a global war.

Bio: Gavrilo Princip

Summary: Gavrilo Princip was a Bosnian Serb nationalist who assassinated Archduke Franz Ferdinand of Austria, sparking World War I.

Born: 1894, Obljaj, Bosnia and Herzegovina

Major impact: Gavrilo Princip's assassination of Archduke Franz Ferdinand of Austria on June 28, 1914, was a catalyst for World War I. As a member of the Young Bosnia movement, Princip's actions were driven by a desire for Bosnian independence from Austrian rule, although his act led to widespread and catastrophic consequences.

Bio: Nedeljko Čabrinović

Summary: Nedeljko Čabrinović was one of the assassins involved in the plot to kill Archduke Franz Ferdinand of Austria.

Born: 1895, Sarajevo, Bosnia and Herzegovina

Major impact: Nedeljko Čabrinović attempted to assassinate Archduke Franz Ferdinand of Austria on June 28, 1914, by throwing a grenade at his car. Although unsuccessful in killing the Archduke, Čabrinović's actions contributed to the tensions that led to the outbreak of World War I.

...

The list includes few more people. Amazing, right? In prompt engineering, templates play a crucial role in shaping the way queries or prompts are constructed to elicit specific responses from language models like GPT-3. Templates essentially serve as structured formats or frameworks into which variables or placeholders can be inserted. They can improve standardization and consistency, reproducibility, assist in targeted query generation. Templates can be adapted and scaled across different applications and domains. Once a template structure is defined, it can be reused with minor adjustments for similar tasks or contexts, which saves time and effort in prompt design.

TEMPLATE PATTERN: STRUCTURED SEQUENCES AND PLANNING

For our last example of the template pattern, let us use LLM to compose a workout using a template to list each of the exercises:

Create a random strength workout with complementary exercises. I am going to provide a template for your output, where CAPITALIZED-WORDS are placeholders: replace them with content. Use the template: `##: NAME` <newline> `set and timing:` REPS @ SETS, `ESTIMATED-TIME-IN-MINUTES` <newline> `muscles:` MUSCLE GROUPS WORKED <newline> `difficulty:` DIFFICULTY SCALE 1-5 <newline> `notes:` FORM NOTES. No itemization.

Notice that above we have used several tricks. We embedded the text that should be set in bold with `**`; we already know this stems from the markdown syntax. Placeholders are in capital letters. We have used the “<newline>” tag to tell LLM where we want line breaks. To avoid the use of bullets, we have added “no itemization” at the end of the prompt. The output:

1: Push-ups

set and timing: 15 @ 3, 10 minutes

muscles: Chest, Shoulders, Triceps

difficulty: 3

notes: Maintain a straight body alignment throughout.

2: Dumbbell Shoulder Press

set and timing: 12 @ 3, 12 minutes

muscles: Shoulders, Triceps

difficulty: 4

notes: Keep core engaged to avoid arching the back.

3: Bent-over Dumbbell Rows

set and timing: 12 @ 3, 10 minutes

muscles: Upper Back, Biceps

difficulty: 3

notes: Focus on pulling elbows back and squeezing shoulder blades together.

4: Bodyweight Squats

set and timing: 20 @ 3, 8 minutes

muscles: Quadriceps, Glutes, Hamstrings

difficulty: 2

notes: Keep knees tracking over toes and maintain a neutral spine.

5: Plank

set and timing: 30 seconds @ 3, 6 minutes

muscles: Core, Shoulders, Back

difficulty: 2

notes: Brace your core and avoid sagging hips.

6: Dumbbell Bicep Curls

set and timing: 12 @ 3, 8 minutes

muscles: Biceps, Forearms

difficulty: 2

notes: Control the movement to fully contract the biceps.

7: Glute Bridges

set and timing: 15 @ 3, 8 minutes

muscles: Glutes, Hamstrings, Lower Back

difficulty: 2

notes: Squeeze your glutes at the top of the movement.

Cool-down: Finish with stretching exercises for each major muscle group, holding each stretch for 20-30 seconds.

This workout targets multiple muscle groups and provides a balanced strength training session. Adjust weights and repetitions based on your fitness level and goals.

This is already great, right? Consider AI where together with the text we have a sketch of the exercise, or an instructional video. All this is probably coming, the AI technology of today is still very young. Until then, you can further refine the workout plan using some more prompting and providing constraints to the list:

Group the exercises to the warmup, main part, and cooldown.

Construct the exercise plan to fit within half-an-hour.

Emphasis should be on legs and feet.

Make exercises lighter.

Or anything similar. Try it out!

CHAIN-OF-THOUGHT (CoT) PROMPTING

Chain-of-thought (CoT) prompting is a technique that encourages language models like ChatGPT to break down reasoning tasks step by step. Instead of responding immediately, the model is guided to explain its thought process before arriving at a final answer. To invoke Chain-of-Thought (CoT) prompting, one needs to explicitly instruct the model to break down its reasoning step by step. The simplest way to do this is through direct instruction, e.g., "think step by step." For example, if analyzing the impact of peer review on research quality, a well-structured prompt would be: "How

does peer review improve research quality? Think step by step." This forces the model to structure its reasoning, ensuring logical flow and completeness in its response.

A more advanced method involves prompting for intermediate reasoning before asking for a conclusion. Instead of asking directly whether a study's claim is valid, one can break the problem into sub-questions. For example, if assessing a study that claims AI-generated abstracts are superior to human-written ones, the user might prompt: "Before we conclude, let's consider: How was 'better' defined? What evaluation methods were used? Were there biases in the study design?" This encourages the model to provide a well-reasoned response rather than a simplistic or potentially biased answer.

Note that CoT prompting does not work equally well for all tasks. Vague prompts such as "Explain in detail" may not guide the model effectively, whereas structured instructions yield better responses. Additionally, CoT is not necessary for factual lookups, such as retrieving a Nobel Prize winner's name, where a simple direct question suffices. However, for research tasks requiring critical analysis, multi-step reasoning, and structured argumentation, CoT prompting is a powerful tool.

For instance, compare the output of the following two prompts:

In short, how does one apply single-cell gene expression analysis?

In short, how does one apply single-cell gene expression analysis? Think step-by-step.

ZERO-SHOT PROMPTING

Zero-shot prompting means asking a language model (LLM) to do a task without showing any examples first. The model has to figure out what to do based only on your instructions and what it already knows from its training. This works well for simple or common tasks. For example, the prompt

Summarize the attached research paper in one paragraph

will instruct AI to generate a summary based on its understanding of the text without needing any sample responses. While zero-shot prompting is helpful for quick and general tasks, it has limitations. The output may lack accuracy if the task is complex or ambiguous, and it may not follow a specific format unless explicitly instructed.

ONE-SHOT PROMPTING: SENTIMENT ANALYSIS

There is also one-shot prompting, a technique in prompt engineering where a single example is provided in the prompt to guide the model's response. This example helps the model understand the desired format, style, or reasoning without requiring multiple demonstrations.

Negative

The movie was fantastic! -> Positive
The plot was slow and boring. ->

Positive

Njami, pizza toscana. ->

Negative

Bljak, rotten grapes. ->

FEW-SHOT PROMPTING

In prompt engineering, "few-shot examples" refer to providing a small number of examples or demonstrations within a prompt to help guide the model in generating desired responses. This technique improves the performance of language models by showing them patterns and structures they should follow when producing their outputs.

Let us start with an example:

Statement: The Earth revolves around the Sun.
Question: What does the Earth revolve around?

Statement: Water boils at 100 degrees Celsius.
Question: At what temperature does water boil?

Statement: The Eiffel Tower is located in Paris.
Question: Where is the Eiffel Tower located?

Statement: Photosynthesis occurs in the chloroplasts of plant cells.
Question: Where does photosynthesis occur in plant cells?

Statement: The Great Wall of China can be seen from space.

Question: What can be seen from space?

Statement: Ljubljana is the capital of Slovenia.

Question: What is the capital of Slovenia?

Statement: All organisms inherit genetic information that specifies their structure and function from their parents.

Question: From whom do all organisms inherit the genetic information specifying their structure and function?

Few-shot examples help the LLM understand the context and structure it needs to follow, leading to more accurate and relevant outputs. The technique allows users to tailor the model's responses to specific tasks or styles without extensive retraining or fine-tuning. By seeing a few examples, the model can generalize from them to handle similar tasks, improving its performance. Few-shot examples are particularly useful for niche or specific tasks where large datasets for fine-tuning may not be available.

Another example:

Situation: I wrote a paper and I am happy with the content.

Action: give it to a colleague to make final checks.

Situation: I revised the paper according to my colleague suggestions.

Action: submit to a journal.

Situation: I got the comments from the editor of a journal.

Action: ...

Action: Address the editor's comments and revise the paper accordingly.

Situation: My paper was rejected from the journal.

Action: Review the feedback provided, make necessary revisions, and submit the paper to another journal.

Provide three more examples of this kind.

Certainly, here are three more examples:

Situation: I received the acceptance notification for my paper.

Action: Celebrate the achievement and prepare the final manuscript for publication, following the journal's guidelines.

Situation: I received a request to review a paper for a journal.

Action: Accept the request if it aligns with your expertise and availability, and begin the review process.

Situation: I have an idea for a new research project.

Action: Conduct a literature review to ensure the originality of the idea and then draft a research proposal.

The prompt typically includes few-shot examples in the form of input-output pairs that demonstrate the desired behavior. Here are a few rules on how they can be used effectively:

1. Choose relevant examples: select examples that are representative of the task and cover a range of possible inputs and desired outputs.
2. Clear formatting: format the examples to distinguish them from the actual input the model needs to process. This often involves using a consistent structure or delimiters.
3. Contextual relevance: ensure that the examples are contextually relevant to the task at hand so that the model can infer the correct patterns.
4. Balancing examples: provide enough examples to cover the variety in the task, but not so many that the prompt becomes excessively long or complex.

A note though: the examples above worked with most recent version of ChatGPT, but somehow failed with Copilot, which seems more verbose. The verbosity in Copilot seems to help the novice users but becomes annoying for experienced users. A simple reminder or instruction at the beginning of chat helps:

|

Talk less, please.

FEW-SHOT EXAMPLES WITH INTERMEDIATE STEPS

Few-shot examples can include a more complex structure, that is, not only of the form statement-response. Here is an example:

|

input: fried steak
think: meat
decision: reject

input: brass with potato
think: healthy fish
decision: like

input: stuffed meatballs
think: I do not like meat
decision: reject

input: chicken bowls
think: meat again, no
decision: reject

input: pancake with jam
think: sweet but ok for the end of the day
decision: like

input: veggie chili
think: vegetarian and spicy
decision: like

	input: tuna salad think: fish decision: like
	input: veggie burritos think: vegetarian decision: like
	input: potato salad with mushrooms think: vegetarian decision: like
	input: grilled chicken think: no to meat decision: reject
	input: cooked lamb
think: meat decision: reject	
	input: pasta with gongozola
think: vegetarian output: like	
	input: potica
think: sweet output: like	

This is not much different from our previous examples of few-shot training but includes an extra step (“think”). Notice also that LLM has also followed the style of writing, that is, creating an output that only includes few words.

(MUCH) MORE COMPLEX FEW-SHOT PROMPTING: REASONING CHAINS

Here is a more complex example:

For a given situation, provide a set of actions. No additional comments are necessary.

Situation: I would like to write a research paper.

Think: What part of my current research is publishable?

Action: List three various topics from my current research that deserve attention.

Think: Which of the three topics is the most exciting one?

Action: Pick the most exciting one and write a short summary.

Think: Where to publish?

Action: Search the citation index to find the highest-ranked journal in my field.

Think: Is my research a good match?

Action: Find papers in a similar domain from that journal.

Situation: I need to design an experiment for my research study.

Think: What is the primary research question I want to address?

Action: Formulate a clear and concise research question.

Think: What is the best methodology to answer this research question?

Action: List three different methodologies suitable for this research.

Think: Which methodology will provide the most reliable and valid results?

Action: Select the most appropriate methodology and outline its key steps.

Think: What resources and equipment will I need for this experiment?

Action: Create a list of required resources and equipment needed for the experiment.

Think: Do I have access to all the necessary resources and equipment?

Action: Identify any gaps and plan how to acquire or borrow missing items.

Situation: I need to write a proposal for a Ph.D. thesis.

Think: What is the central research problem I want to address?

Action: Define the primary research problem and its significance.

Think: What are the objectives and aims of my research?

Action: List the main objectives and aims of the proposed research.

Think: What is the existing literature on this topic?

Action: Conduct a literature review and summarize key findings.

Think: What methodology will I use to address my research problem?

Action: Choose a suitable methodology and outline the research design.

Think: What is the potential impact of my research?

Action: Describe the expected contributions to the field and potential applications.

Think: How will I structure my proposal?

Action: Create an outline for the proposal, including sections such as introduction, literature review, methodology, expected results, and timeline.

Think: Who will be my target audience and stakeholders?

Action: Identify and list potential advisors, funding agencies, and collaborators.

Situation: My Ph.D. advisor just resigned from the University and is now a chef in a restaurant.

Think: Who can take over as my Ph.D. advisor?

Action: Identify and list potential faculty members who can serve as your new advisor.

Think: How can I ensure a smooth transition?

Action: Arrange meetings with potential advisors to discuss your research and seek their

guidance.

Think: What administrative steps do I need to take?

Action: Contact the graduate program coordinator to understand the process for changing advisors and complete any necessary paperwork.

Think: How can I maintain momentum in my research during this transition?

Action: Create a plan to stay on track with your research goals, including setting short-term objectives and regular check-ins with your temporary advisor.

Think: What support do I need during this transition?

Action: Reach out to peers, colleagues, and mentors for support and advice on navigating this change.

Few-shot examples with intermediate steps refer to a learning or task-solving approach where a model or algorithm is given a small number of examples (few-shot) and shown the intermediate steps needed to reach the solution rather than just the final answer. This helps the model understand the process and reasoning required to arrive at the solution, improving its performance on similar tasks.

(MUCH) MORE COMPLEX FEW-SHOT PROMPTING: SCORING

Following is another example where a solution also requires numerical grading of an option.

Input: Skaline Luxury rooms Split features accommodation situated in Split, 1 km from Diocletian's Palace and 4.5 km from Znanjan Beach. All units feature air conditioning and a satellite flat-screen TV.

Output: reject

Explanation: too much in the city center, too crowded

Grade: 2/10

Input: Situated in Pučišća, within 400 metres of Sveti Rok Beach and 600 metres of Macel Beach, Bed & Breakfast Blanc offers accommodation with free WiFi, air conditioning, a terrace and a bar. The property features sea views and is 22 km from Olive Oil Museum Brac and 13 km from Gažul. The bed and breakfast has family rooms.

Output: interesting

Explanation: in the village, less crowded, near the sea

Grade: 6/10

Input: Situated within 300 metres of Sveti Rok Beach and 1.5 km of Macel Beach, Lučica, Luxury Rooms By the Sea features rooms with air conditioning and a private bathroom in Pučišća.

Output: this looks great, alone, near the sea, close to the beach, need to check the price

Grade: 9/10

Input: Providing city views and free WiFi, Zara Palace - design rooms provides accommodation conveniently situated in the centre of Zadar, within a short distance of Kolovare Beach, Maestrala Beach and Karma Beach. It is located 28 km from Kornati Marina and features luggage storage space. This 4-star guest house offers private entrance.

Output: reject

Explanation: too much in the city, too crowded

Grade: 3/10

Input: Located at the foot of Vatnajökull National Park, Lilja Guesthouse offers accommodation along Ring Road 1 in Flatey. Höfn is 28.5 km away from the property and Jökulsárlón Glacier Lagoon is 51 km away. Certain units include a seating area for your convenience. You will find a coffee machine in the room. Free WiFi is featured throughout the property. There is a shared lounge at the property. You can view the mountains from the property. Popular activities in the area include glacier tours. Egilsstaðir town centre and airport are 205 km away.

Output: interesting

Explanation: remote location, close to nature, near mountains and glacier

Grade: 7/10

Input: Apartments Tonic is set in the small village of Lumbarda, 5 km from Korčula's Old Town. The property features a spacious terrace overlooking lush Mediterranean vegetation and Adriatic Sea. The nearest beaches are only steps away. All apartments are air-conditioned and feature a patio with sea views and an equipped kitchenette. Free Wi-Fi and free private parking are provided. Lumbarda town centre with restaurants, bars and shops can be reached within 5 minutes on foot from Apartments Tonic. The local bus stop is set in the centre of Lumbarda and every hour a bus passes and takes guests to Korčula Ferry. Guests can rent a motor scooter at the property.

Output: interesting

Explanation: small village, close to the sea, peaceful location, convenient access to town center

Grade: 8/10