
ADVANCED PROMPTING PATTERNS

BLAŽ ZUPAN, LJUBLJANA, 2024

LECTURE NOTES FOR THE UNIVERSITY OF LJUBLJANA'S COURSE ON
CHATGPT FOR RESEARCHERS

These notes continue our writing on the topic of prompt engineering. We start where we have finished with the notes on “Introduction to Prompt Engineering” and continue introducing patterns that guide large language models (LLMs) to produce interesting and useful outputs.

Personal Assistant Prompt

In a sense, we have already been introducing this type of pattern, where we are instructing LLMs to ask us about the information that is needed to fulfill a particular task. Here, we combine this the 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)

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

QUESTION REFINEMENT PROMPT

We can engage language models to help us write better prompts. Here is an example initial prompt:

Whenever I ask a question, suggest a better question and ask me if I want to use it instead.

Or a more specific one:

Whenever I ask a question about molecular biology, suggest a better version emphasizing machine learning and AI. Ask me for the first question to refine.

Or a bit longer one where we also define the structure of LLM's response:

Whenever I ask a question about molecular biology, suggest a better version of a question that involves some computational aspect and specific related algorithm. Present this revised question, and provide an answer (Q: <revised question>**, **A**: <answer>). Be brief, answer with single sentence. Now, ask me for the first question to refine.**

Here is a brief example of the chat that followed our last prompt above:

Sure, what's your initial question about molecular biology that you'd like to refine into a computational aspect?

How do I compare two DNA sequences?

Q: How can I align two DNA sequences using the Needleman-Wunsch algorithm?

A: Use the Needleman-Wunsch algorithm to perform global sequence alignment, which scores matches, mismatches, and gaps to find the optimal alignment between two DNA sequences.

What's the next question you'd like to refine?

Can I infer evolutionary tree given a set of gene sequences from various organisms?

Q: How can I construct an evolutionary tree using Maximum Likelihood method based on a set of gene sequences from various organisms?

A: Utilize the Maximum Likelihood method to infer an evolutionary tree, which calculates the tree that maximizes the probability of observing the given gene sequences under a specified model of sequence evolution.

What other question would you like to refine next?

Here, we have successfully combined several prompting techniques we have already reviewed: programming, that is, instructing LLMs from the start on, question refinement, and constraining the output of the LLM to a specific, pred-defined structure.

COGNITIVE VERIFIER PATTERN

Here, we will ask LLM to help us break our question or problem into a set of subquestions, or subproblems. Our hypothesis is that once we provide an answer to these subquestions, LLM can become more precise and can provide us a better answer. Here is an example of such a conversation, with an all-important initial prompt that specifies the type of the chat we would like to have.

When you are asked a question, follow these rules: (1) generate several additional questions, one by one, that would help you to answer the questions more accurately; (2) after you get my answers, combine them with the individual questions to produce the final answer to the overall question. Number the questions. Ask max 3 subquestions.

Sure, I can follow those rules! Could you please ask your main question so we can begin the process?

Can I use machine learning to analyze single-cell gene expression data?

Great question! Let's break it down to get a comprehensive answer:

1. What are the characteristics of single-cell gene expression data that make it suitable or challenging for machine learning analysis?
2. What are the typical preprocessing steps required before applying machine learning to single-cell gene expression data?
3. What are some specific machine learning techniques commonly used for analyzing single-cell gene expression data?

Once I have your answers to these questions, I can give you a detailed and accurate response about using machine learning for analyzing single-cell gene expression data.

1) large number of features and data instances. 2) normalization and batch-effect correction. 3) dimensionality reduction, t-SNE visualization.

Great, thank you for the answers! Now, let's combine them with the original question:

Main question: Can I use machine learning to analyze single-cell gene expression data?

Answers to subquestions:

1. What are the characteristics of single-cell gene expression data that make it suitable or challenging for machine learning analysis?
 - Single-cell gene expression data typically have a large number of features and data instances, which can pose challenges for traditional machine learning algorithms due to scalability issues and high-dimensional data.
2. What are the typical preprocessing steps required before applying machine learning to single-cell gene expression data?
 - Typical preprocessing steps include normalization of gene expression values to account for differences in sequencing depth and gene expression levels, as well as batch-effect correction to remove technical variations introduced during sample preparation and sequencing.
3. What are some specific machine learning techniques commonly used for analyzing single-cell gene expression data?
 - Specific techniques include dimensionality reduction methods such as principal component analysis (PCA) and t-distributed stochastic neighbor embedding (t-SNE) for visualizing high-dimensional data and identifying clusters of cells with similar gene expression profiles.

Final answer: Yes, machine learning can be effectively used to analyze single-cell gene expression data. However, it requires addressing challenges such as handling a large number of

features and instances, performing necessary preprocessing steps like normalization and batch-effect correction, and employing specific techniques such as dimensionality reduction and t-SNE visualization to extract meaningful insights from the data.

Can you modify the initial prompt for the ChatGPT to be a bit less verbose? Can you modify it so that besides the question we can also get an explanation why answering this question will provide additional insight and help in composing final answer? Can you compose a prompt to request that questions and their motivations are printed out in a structured way?

AUDIENCE PERSONA PATTERN

This pattern instructs LLM to produce the output for a particular audience. As always, this is best illustrated with an example:

In brief, in one paragraph, explain personalized medicine and its principal tools. Assume I have no background knowledge in medicine or molecular biology, and I have yet to work in science.

Personalized medicine refers to a medical approach that customizes healthcare decisions and treatments to each individual patient based on their unique genetic makeup, environment, and lifestyle. This tailored approach aims to enhance treatment effectiveness and reduce side effects by considering genetic variations that influence disease susceptibility and response to therapies. The principal tools of personalized medicine include genetic testing to identify specific genetic mutations or variations, biomarker analysis to predict disease risk or treatment response, and advanced imaging techniques to assess disease progression at a molecular level. These tools enable healthcare providers to deliver precise treatments that are more likely to be effective for each patient, marking a shift towards more targeted and individualized healthcare strategies.

Explain personalized medicine and its principal tools. Assume I am ten years old kid who likes to play with LEGO bricks.

Hey there! Imagine personalized medicine like building a LEGO spaceship that's just for you. It's all about doctors using special tools to understand exactly how your body works and what makes you unique. They look at things like your genes, which are like tiny instructions that tell your body how to grow and work. With genetic testing, they can see if you have any special differences in your genes that might affect your health. Biomarkers are like little signs in your body that show if something might be wrong or how well a treatment is working, kind of like different LEGO pieces showing if your spaceship is strong or needs fixing. And then there are cool machines that take pictures inside your body to see everything up close, like finding the right LEGO pieces to make your spaceship perfect. So, instead of giving everyone the same medicine, doctors can use these tools to give you exactly what you need to stay healthy and strong, like building a LEGO spaceship that's just right for you!

See a huge difference in the output? Above, we requested the explanation to be brief and in one paragraph just in observance of the space for these lecture notes. Audience persona pattern starts with “Assume I am ...”.

Here is another example:

Explain, in short, using max few sentences, what large language models are; assume I have no computer science background.

Large language models are advanced computer programs designed to understand and generate human-like text. They learn from vast amounts of written language to predict and produce coherent sentences, making them capable of tasks like answering questions, summarizing text, and even generating creative writing. These models are built using complex algorithms and require substantial computational power to operate effectively.

Explain large language models and how I can use them. Assume I am a freshmen and I get bored easily.

Large language models are like super-smart assistants that understand and generate text. You can use them in various ways, like getting answers to questions, writing essays or stories, or even just chatting for fun. They learn from lots of text, so they can talk about almost anything. If you get bored easily, they can help keep things interesting by providing new ideas or information on topics you're curious about.

Audience persona patterns are fun, and it is great experimenting with them. They are helpful also in practicing for your talks, presentation or pitches when you are thinking about how to convey your ideas for audience. Say, you invented a new chemical that protects wood, and you are pitching your idea to businesspeople with no knowledge of chemistry. Or you want to sell your AI-based data analysis system to boat manufacturers. Or convince a director of a bank why would they desperately need an anthropologist...

Here are a few more ideas, try them out or better, experiment on your own. I particularly liked the first one below, but the answer was too verbose to include it here.

Explain the process of scientific funding and grants. Assume I am a kid.

Explain the process of scientific funding and grants. I only accept explanations in math.

Explain the process of scientific funding and grants. I am computer scientist and only understand pseudocode and algorithms.

FEW-SHOT EXAMPLES

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

Let us start with an example:

Generate questions for the following statements:

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 a capital of Slovenia.

Question: What is the capital of Slovenia?

Statement: all organisms inherit the genetic information specifying 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 on those tasks. Few-shot examples are particularly useful for niche or specific tasks where large datasets for fine-tuning may not be available.

Another example:

For a given situation provide an action.

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: 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.

Few-shot examples are typically included in the prompt in the form of input-output pairs that demonstrate the desired behavior. Here are few rules 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: clearly 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 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.

Here is another example demonstrating the use of few shot examples for sentiment analysis:

Example 1: Text: "I love this product, it works wonderfully!"

Sentiment: Positive

Example 2: Text: "This is the worst experience I have ever had."

Sentiment: Negative

New Input: Text: "The service was okay, not great but not bad."

Sentiment: ?

New Input: Text: "The service was okay, not great but not bad." Sentiment: Neutral

Text: "After the dinner, I had a strange taste in my mouth."

Text: "After the dinner, I had a strange taste in my mouth." Sentiment: Negative

Text: Sailing on high winds was great.

Text: "Sailing on high winds was great." Sentiment: Positive

Can you change the above prompt to answer without repeating the text input? Can you do this using the programming pattern, that is, using initial prompt for instructions?

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

output: reject

input: brass with potato

think: healthy fish

output: like

input: stuffed meatballs

think: I do not like meat

output: reject

input: chicken bowls

think: meat again, no

output: reject

input: pancake with jam

think: sweet but ok for the end of the day

output: like

	input: veggie chili think: vegetarian and spicy output: like
	input: tuna salad think: fish output: like
	input: veggie burritos think: vegetarian output: like
	input: potato salad with mushrooms think: vegetarian output: like
	input: grilled chicken think: no to meat output: reject
	input: cooked lamb
think: meat output: 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.

Here is a more complex example:

For a given situation provide set of actions. No additional comments 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 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 type of 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.

Following is another example where 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 Zrnan Beach. All units feature air conditioning and a satellite flat-screen TV.

Output: reject

Explanation: too much in the city, 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, Maestrara 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

TEMPLATE PATTERN

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: Julian's Assange's wife has told of her elation that the WikiLeaks founder has been released from Belmarsh prison in London and will soon be a "free man" under a deal in which he will plead guilty to violating US espionage law. Speaking from Australia, where she flew on Sunday to prepare her family's new life, Stella Assange, a human rights lawyer, said she had not told the couple's two young sons, Gabriel and Max, about their father's release after five years in jail for fear of the information leaking. She said: "All I told them was that there was a big surprise. And, on the morning that we left, I told them where we're heading to the airport, and we got on the plane, and I told them that we were going to visit our family, their cousin, their grandfather and so on. "And they still don't know. We've been very careful, because obviously, no one can stop a five- and a seven-year-old from, you know, shouting it from the rooftops at any given moment. Because of the sensitivity around the judge having to sign off the deal, we've been very careful, just gradually, incrementally telling them information." Assange, 52, was woken at 2am in his cell in Belmarsh on Monday before being handcuffed and transported to Stansted airport to take a chartered jet via Bangkok to the island of Saipan in the Northern Mariana Islands, which is US territory.

Now please create three questions using my template.

Question: Where did Julian Assange travel after his release from Belmarsh prison?

Answer: Julian Assange was transported to Stansted airport in London and then took a chartered jet via Bangkok to the island of Saipan in the Northern Mariana Islands, which is US territory.

Question: Why did Stella Assange refrain from informing her sons about Julian Assange's release immediately?

Answer: Stella Assange refrained from immediately informing her sons about Julian Assange's release due to the sensitivity around the judge signing off on the deal. She was concerned that their young age might lead them to inadvertently disclose the information.

Question: How long had Julian Assange been in jail before his release?

Answer: Julian Assange had been in jail for five years before his release from Belmarsh prison in London.

Let us combine this 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.

For our final prompt in this lecture, 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, replaces them with content. This is 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 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!