
INTRODUCTION TO PROMPT ENGINEERING

BLAŽ ZUPAN, LJUBLJANA, 2025

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



In this lecture, we start to explore the art of writing prompts. Regarding chatbots and language models (LLMs), a prompt is the input text the user provides to which the model generates a relevant and coherent response. In simple terms, the prompt is whatever goes into LLM so that this generates an output. In a narrow sense, ChatGPT is an app where users write text, and ChatGPT assembles these texts into a prompt, passes it to LLM, and renders the output. Note that in ChatGPT, the prompt is not only what we write in a single message box but the entire text that we wrote from the beginning of our chat, that is, from the top of the current window, provided that this fits within the prompt length limitation.

Users most often use chatbots to answer questions or get explanations, just like in the following example:

In one sentence, what is a large language model?

A large language model is an advanced artificial intelligence system designed to understand, generate, and manipulate human language based on vast amounts of text data.

Where do you get this data to train the model?

The data used to train a large language model comes from diverse and extensive text sources available on the internet, including websites, books, articles, and other publicly accessible written content.

The text we wrote as input to chatbots is marked in bold, and the output is in plain text, both indented and marked with a line. We will use the same notation in the lecture notes below. Notice that the second question above relates to the first one and its answer. The text sent to the LLM to answer, or better, complete the text after the second question, is a concatenation of everything that is included in the chat, that is, in this case, the text of the first question, the answer, and the text of the second question. LLMs do not have any “memory” of our conversation, while chatbots, the application, in a way, does, as for each prompt in an open chat window, sends the entire conversation so far to the LLM.

It is surprising, if not mind-blowing, that a simple LLM can be used in more complex conversations than the one above. Through clever prompting, we can use LLM in much more complex scenarios and for much more elaborate conversations. These lecture notes explore some of these. Before we start, though, a disclaimer: while a vivid user of ChatGPT and LLM, the author of these notes, your lecturer, learned about these tricks from others through experimentation and by watching YouTube videos, reading research papers on LLM, programming with LLM, and lately, through some of the Coursera courses on prompt

engineering. Many of the following examples were inspired by one of these courses, “Prompt Engineering for ChatGPT” by Dr. Jules White, which can be accessed on Coursera, which I also highly recommend for viewing.

The outputs of the prompts below depend, of course, on the LLM used. The chatbot that was used in this lecture notes was ChatGPT. Chatbots are changing almost daily, and there is a good chance that some of the experiments we have done below cannot be replicated, even vaguely. To study and complete the assignments for this course, you may of course use any other chatbot for your experimentation, like Copilot, Claude, or Le Chat. Note that the answers of LLMs are not deterministic (on purpose, as discussed in the class) and are purposely, to some extent randomized, where LLM chooses the text that completes the prompt randomly, among most likely completions (the process is a bit more complex, but we covered that in our previous lectures). With this in mind, the prompts we use here would not have the same result if you had tried them independently, but the results should (hopefully :)) be too different.

RANDOMNESS AND STABILITY OF THE OUTPUTS

Large language models generate more “stable” results for the prompts composed of the text seen in the training data more often. Of course, we have no idea what training data was, but intuitively, the answers to the frequent phrases will be more stable. Here, “stable” means that the output to the same prompt will most often be the same or very similar. For instance, try the following prompts:

Mary had a little...

The Roses are red...

Žive naj vsi narodi...

Choose a prompt and try it several times, preferably starting a new chat each time by clicking on the pencil icon at the top right of the window. (Interestingly, above will no longer not work in MS Copilot, try this is out, though). Now start with another prompt that defines perhaps a less defined context:

My PhD advisor is...

GOOD PROMPTS ARE (VERY) SPECIFIC

Consider the following prompt.

Which scientific journals are good?

In the answer, LLM will remind us that the "goodness" of a scientific journal can depend on various factors and elaborate on what these factors are and how they are measured. Consider that the conversations with ChatGPT are just like those with a person: the more specific the

question, the better the answer. It was such conversations or conversation fragments on which LLMs have been trained and can, in a way, reproduce them when we engage them.

If we are yet to learn how to write specific prompt, there's always a workaround to instruct ChatGPT to ask us for details:

From now on, when I ask a question, you will suggest a better version of my question and ask if I would like to use it. If yes, you will use a refined question, if not, you will suggest another refinement. Ok? I will start with questions in my next prompt.

Sure, that sounds like a helpful approach to refine your questions. Feel free to go ahead with your first question whenever you're ready!

Which scientific journals are good?

Would you like to refine your question to specify a particular field of science or topic?

Molecular biology.

Would you like to know which scientific journals are reputable and widely recognized in the field of molecular biology?

yes

Great! Here's a more refined version: "Which reputable scientific journals are widely recognized in the field of molecular biology?" Would you like to use this version?

yes

In the field of molecular biology, some reputable scientific journals include:

1. Cell
2. Nature Genetics
3. Molecular Cell
4. Journal of Molecular Biology
5. Genes & Development

These journals publish high-impact research articles and are well-regarded in the scientific community for their contributions to molecular biology. If you have any specific questions about these journals or their content, feel free to ask!

Note that regardless of how specific they are, some prompts, again, based on the material, which was used in training of LLM, would produce more and some less stable answers.

Compare, for instance, stability of the answers for the following two prompts:

Žive naj vsi narodi...

Žive naj vsi ljubitelji jazza...

The more specific a prompt is, the more specific the answer will be. Consider the following two prompts:

What kind of city is Ljubljana?

What kind of city with respect to the quality of life is Ljubljana?

The second prompt is more specific and yields a more stable answer from the LLM. Use “with respect to” to provide context and to get more stable answers. Include details.

PROMPTS SPAN TIME: OUR EARLIER PROMPTS AFFECT THE OUTPUT

We have been here already, but it helps to repeat this: prompts are composed by the ChatGPT application, and a single input passed to the LLM is actually – each time we enter a message to ChatGPT – composed from all our previous messages in the chat. By “chat,” we mean a sequence of messages and outputs from a ChatGPT window. Switching the topic or the type of conversation helps to start fresh by starting a new chat by clicking on the pencil icon in ChatGPT’s interface.

Let’s see how we can start a conversation with a prompt that instructs LLM to form a specific type of (very short) answer:

From now on, when I ask a question, answer with a simple sentence or, preferably, with a single word.

Got it!

In which country in the world the people live longest?

Japan.

How long, on the average, they live there?

Around 84 years.

Which country comes next?

Switzerland.

Life expectancy?

Around 83 years.

Notice that all the outputs followed our instructions from the first message. LLM has no memory; ChatGPT composes prompts with a history of our messages, giving us an impression of continuing conversion.

IMPOSING STRUCTURE

We can also propose what should the structure of the answer look like

What kind of city with respect to the quality of life is Ljubljana? Use the following structure and put the headers in bold:

Title: <title of the section>

People: <how many people live there>

Temperature: <what is the average temperature>

Historic buildings: <a short list of historic buildings>

Summary: <two-sentence summary>

Title: Quality of Life in Ljubljana

Population: Approximately 300,000 people live in Ljubljana.

Temperature: Ljubljana experiences a continental climate with warm summers and cold winters. Average summer temperatures range from 20°C to 25°C (68°F to 77°F), while winter temperatures average around 0°C (32°F).

Historic Buildings: Ljubljana boasts several notable historic buildings, including Ljubljana Castle, Triple Bridge, Dragon Bridge, and the Cathedral of St. Nicholas.

Summary: Ljubljana offers a high quality of life with its manageable size, picturesque setting along the Ljubljanica River, and a blend of historic charm and modern amenities. The city's pleasant climate and rich cultural heritage make it an attractive destination for residents and visitors alike.

Same for Trieste.

Title: Quality of Life in Trieste

Population: Trieste is home to approximately 200,000 people.

Temperature: Trieste has a humid subtropical climate influenced by its proximity to the Adriatic Sea. Summers are warm with average temperatures ranging from 24°C to 29°C (75°F to 84°F), while winters are mild with average temperatures around 5°C (41°F).

Historic Buildings: Trieste features a wealth of historic architecture, including Miramare Castle, San Giusto Castle and Cathedral, the Roman Theatre, and the Grand Canal.

Summary: Trieste offers a unique quality of life characterized by its Mediterranean climate, rich history, and cosmopolitan atmosphere. The city's blend of Italian, Slovenian, and Austrian influences, along with its scenic coastal location, makes it a culturally diverse and appealing place to live and visit.

That's great, right? In the prompt, we have instructed ChatGPT to use a specific structure when providing the answer, and it will use this structure if we continue asking about the same topic, that is, in our case, asking about specific cities.

STRUCTURED OUTPUT FOR TEXT FORMATTING

We used the idea of specifying the pattern for text re-formatting. In the past, the author of these lines used macro commands in text editors like emacs for such tasks, but now chatbots have become reliable and simpler option. Ok, here is an example (names used are made up):

Consider the list:

"Marolt, Lina" <lina.marolt@ag.uni-lj.si>,

"Klemenčič, Svit" <svit.klemencic@aluo.uni-lj.si>,

"Drnovšek, Neža" <neza.drnovsek@mf.uni-lj.si>,

"Rihar, Jure" <jure.rihar@ef.uni-lj.si>,

"Zajc, Tisa" <tisa.zajc@fdv.uni-lj.si>,

"Kosec, Lan" <lan.kosec@ff.uni-lj.si>,

"Strnad, Gal" <gal.strnad@fmf.uni-lj.si>,

"Planinšek, Inja" <inja.planinsek@fri.uni-lj.si>

Reshape it into:

Name **Surname** (email)

No other comment, please.

Lina **Marolt** (lina.marolt@ag.uni-lj.si)
Svit **Klemenčič** (svit.klemencic@aluo.uni-lj.si)
Neža **Drnovšek** (neza.drnovsek@mf.uni-lj.si)
Jure **Rihar** (jure.rihar@ef.uni-lj.si)
Tisa **Zajc** (tisa.zajc@fdv.uni-lj.si)
Lan **Kosec** (lan.kosec@ff.uni-lj.si)
Gal **Strnad** (gal.strnad@fmf.uni-lj.si)
Inja **Planinšek** (inja.planinsek@fri.uni-lj.si)

Notice the use of a markdown syntax for bold text?

REVERSE ROLES: TUTORIAL 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 app. 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 something like 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.