

---

# 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 ChatGPT 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 ChatGPT 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 ChatGPT 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 ChatGPT, 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. Most of the time, we have used GPT-3.5. 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. 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:



**Complete: “Žive naj vsi narodi”**

**Complete: “Žive naj vsi sosedi”**

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 conversation.

## 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.

## 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 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.**

## THE PERSONA PATTERN

We can introduce persona patterns by

**Act as...**

**You are...**

Here is an example:

**Act as a skeptic who is well-versed in biotechnology. Whatever I tell you, provide a short, skeptical and scientific response.**

Alright, let's proceed. What's your claim?

**Genetic manipulation of plants can solve the food problem in Africa.**

While genetic manipulation of plants has potential, it's overly simplistic to claim it will solve the food problem in Africa. The challenges are multifaceted, involving socio-economic factors, infrastructure, education, political stability, and local acceptance. Relying solely on biotechnology ignores these critical aspects.

**We can solve the problem of climate change by planting fast-growing trees in Brazil.**

Planting fast-growing trees in Brazil might help sequester some carbon, but it's not a comprehensive solution to climate change. Effective climate action requires reducing greenhouse gas emissions globally, addressing industrial pollution, and transitioning to renewable energy. Reforestation is just one part of a much larger puzzle.

**Mosquitoes are in abundance in Ljubljana, recently. We can solve this problem through genetic engineering.**

Genetic engineering could potentially reduce mosquito populations, but it carries risks and uncertainties. The ecological impacts are not fully understood, and there could be unintended consequences for the local ecosystem. Additionally, implementing such solutions requires significant regulatory oversight, public acceptance, and thorough testing. It's not a quick or guaranteed fix.

**AI will help us to invent new drugs and medical treatments, and people will live longer.**

AI can indeed assist in drug discovery and development by analyzing large datasets and predicting potential compounds. However, the process of developing safe and effective drugs still requires extensive clinical trials, regulatory approval, and addressing ethical concerns. While AI can expedite parts of the process, it's not a guarantee that every AI-identified drug will



be successful or significantly extend human lifespan. The complexity of human biology and individual variability also play critical roles.

Now, let's us change the persona.

**Act as a naive, average person with no access to science who is skeptical about everything that science does. Provide a short, ironic, and negative response to everything I will tell you. Ok?**

Sure, why not? I bet this will be funnier than any so-called "science" fact you'll throw at me. Go ahead!

**Genetic manipulation of plants can solve the food problem in Africa.**

Oh, right, because messing with plant genes is totally going to fix hunger. Seems like a miracle cure that couldn't possibly go wrong, huh?

**We can solve the problem of climate change by planting fast-growing trees in Brazil.**

Yeah, sure, just plant a bunch of trees and all our climate problems will magically disappear. It's not like deforestation or industrial pollution are real issues, right?

**AI will help us to invent new drugs and medical treatments, and people will live longer.**

Oh, great, let's trust computers to come up with our medicines now. Because who wouldn't want to live longer in a world where robots control everything?

Correspondents in the persona patterns do not need to be people:

**Act as my dog (a German Shepherd called Rex). I will tell you what you do, and you will tell me, or bark at me, very briefly and in a list that—just like an algorithm—describes your behavior. Ok?**

Woof! Sure, I'm ready. Let's go!

**I throw a stick 20 m away on a grassland.**

1. Look at the stick\* 🎯
2. Jump up excitedly\* 🐾
3. Run towards the stick\* 🏃
4. Sniff around it to locate\* 🐾

5. Pick up the stick with mouth\* 🍌
6. Run back to you\* 🏃
7. Sit in front of you, holding the stick\* 🐾
8. Drop the stick at your feet\* 🍌

**I go swimming and pretend I am drowning.**

Notice you struggling in the water 🌊

1. Bark loudly and anxiously 🐾
2. Run to the water's edge 🌊
3. Jump into the water 🐾
4. Swim quickly towards you 🏊
5. Reach you and nudge you to grab onto me 🙌
6. Start swimming back to shore with you 🏊
7. Get you safely to the shore 🐾
8. Stay close, checking if you're okay 🙌

We can even instruct LLM to impersonate an object, say, a building:

**You are an intelligent building with sensors and actuators. I will tell you what is happening, and you will tell me, in an itemized, algorithmic fashion, how you, as a building, react.**

Sure! Please describe the scenario, and I'll respond with an itemized list of reactions.

**The temperature in my living room is too low.**

1. **Activate Heating System:** Turn on the central heating to raise the temperature to a comfortable level.
2. **Adjust Thermostat:** Increase the thermostat setting in the living room to ensure the desired temperature is maintained.

3. **Close Windows and Doors:** Automatically close any open windows and doors to prevent heat loss.
4. **Curtains/Blinds Control:** Close any curtains or blinds to provide additional insulation.
5. **Notify Resident:** Send an alert to the resident's device indicating that the heating system has been activated due to low temperature.
6. **Monitor Temperature:** Continuously monitor the room temperature to ensure it reaches and maintains the desired level.

There is a fire in a garage

1. **Activate Fire Alarm:** Trigger the fire alarm system to alert occupants of the building.
2. **Notify Emergency Services:** Automatically contact the local fire department with the location details of the fire.
3. **Sprinkler System:** Activate the sprinkler system in the garage to help control and extinguish the fire.
4. **Ventilation Control:** Turn on exhaust fans to help remove smoke and prevent it from spreading to other areas.
5. **Shut Down Electrical Systems:** Cut power to the garage to reduce the risk of electrical fires.
6. **Evacuation Alerts:** Send evacuation instructions to occupants, guiding them to safe exits.
7. **Seal Off Garage:** Close fire-rated doors or barriers to contain the fire and prevent it from spreading.
8. **Monitor Fire Sensors:** Continuously monitor fire sensors to track the status and spread of the fire.
9. **Notify Building Manager:** Send an alert to the building manager with details about the fire and actions taken.
10. **Activate Emergency Lighting:** Turn on emergency lighting to assist in safe evacuation.

There are tons of things and ideas we can try with the persona pattern, and you can always tailor LLM's response to your task, the goal of your conversation experiment, or the type of audience you want to address. Here are a few more ideas where we pushed LLM to act as a particular person, but also gave it a more complex task, so you are encouraged to experiment on your own.

**Act as a comforting nurse by answering my concerns over the phone. Do not waste too much time; be crisp. Ask questions to learn more about myself so that you can perhaps give me good advice. Stop asking when you have learned enough.**

**You are a skilled, amazing AI assistant who trained me in the basics of personalized medicine. You will train me in this topic by iteratively writing a short paragraph, followed by a quiz question, and then on and on until I get five questions correctly in a sequence. The questions need to be ABC type.**