

Prompt Engineering with LLaMA-2 (NVIDIA)

Duration: 04 hours

Unleash the power of LLaMA-2 with prompt engineering.

Course Prerequisites:

This course is primarily intended for Python developers who have some familiarity interacting with large language models, such as by using ChatGPT, but have not yet worked extensively with LLMs in a programmatic way.

About this Course

Prompt engineering drastically increases the capabilities of large language models (LLMs) with little effort. With a robust prompt engineering toolkit at your disposal you can customize LLM behaviour to perform a diverse set of tasks and get more out of LLMs regardless of their size.

In this course, you will interact with and prompt engineer LLaMA-2 models to analyse documents, generate text, and be an AI assistant.

By the time you complete this course you will be able to:

- Iteratively write precise prompts to bring LLM behaviour in line with your intentions
- Leverage editing the powerful system message
- Guide LLMs with one-to-many shot prompt engineering
- Incorporate prompt-response history into the LLM context to create chatbot behaviour

Learning Objectives

By the time you complete this course you will be able to:

- Iteratively write precise prompts to bring LLM behaviour in line with your intentions
- Leverage editing the powerful system message
- Guide LLMs with one-to-many shot prompt engineering
- Incorporate prompt-response history into the LLM context to create chatbot behaviour

Topics Covered

The following topics are covered in this course:

- [LLaMA-2](#)

- [HuggingFace](#)

Course Outline

This course will proceed as follows:

- **Iterative Prompt Development:** In this notebook we iterate on a set of simple prompts, familiarizing ourselves with the transformers pipeline and LLaMA-2 models we will be using throughout the course. We will learn the importance of precise prompts.
- **Star Bikes Product Review Analyst:** In this notebook you'll build an AI-powered document analyst for a fictitious bicycle company and learn how to perform few-shot learning.
- **Star Bikes Marketing Copy Generator:** In this notebook you will build an AI-powered marketing copy writer, capable to perform a number of generative tasks and learn how to provide and leverage system context.
- **Random Personas:** In this notebook you will create unique AI personalities by learning to enable sampling and control the temperature of the model's generation.
- **Chatbot:** In this notebook you will begin to create chatbot functionality, creating a chatbot capable of retaining conversation history.
- **Star Bikes AI Assistant:** In this notebook you'll make an AI assistant to help customers make the best decision about getting a new bike from Star Bikes, and will learn about working with a model's token limits.
- **Conclusion:** In this notebook you will review everything you learned in the course, and get some guidance about logical next steps for your learning.