WhatsApp/ Tel: +1-929-672-1814
Email: info@genai-training.com
www.genai-training.com

Foundations of Prompt Engineering

Course ID # Genai-103-Prompt-Engineering

In this course, you will learn the principles, techniques, and the best practices for designing effective prompts. This course introduces the basics of prompt engineering and progresses to advanced prompt techniques. You will also learn how to guard against prompt misuse and to mitigate when interacting with FMs.

Course level: Basics and Intermediate
Duration: 12 hours

Activities

This course includes eLearning interactions, theory and Lab works.

Course objectives

In this course, you will learn to:

- What is prompt engineering, define prompt engineering and apply general best practices when interacting with FMs (Foundation Model)
- Identify the basic types of prompt techniques, including zero-shot and few-shot learning. Multi Shot and COT Chain oh thought
- Apply basic and advanced prompt techniques when necessary for your use case
- Identify which prompt-techniques are best-suited for specific models
- Identify potential prompt misuses
- Analyze potential bias in FM responses and design prompts that mitigate that bias

Intended audience

This course is intended for:

Prompt Engineers, Data Scientists, Developers, Data Engineer and IT admins

Prerequisites

We recommend that attendees of this course have taken the following courses:

- Introduction to Generative AI Essentials – Art of the Possible (14-hour, digital course)
- Planning a Generative AI Project (3-hour, digital course)
- Amazon Bedrock Getting Started (3-hour, digital course)

WhatsApp/ Tel: +1-929-672-1814
Email: info@genai-training.com
www.genai-training.com

Course outline

Introduction

- Introduction
- Basics of Foundation Models
- Fundamentals of Prompt Engineering
- Prompt Types and Techniques

Prompt Types and Techniques

- Basic Prompt Techniques
- Advanced Prompt Techniques
- Model-Specific Prompt Techniques
- Addressing Prompt Misuses
- Mitigating Bias
- Conclusion

Course Summary
- Session descriptions

Lesson descriptions

Session 1: Basics of Large Language Models

In this session, you will learn how to develop a fundamental understanding of foundation models (FMs), including an understanding of a subset of FMs called large language models (LLMs). First, you will be introduced to the basic concepts of a foundation model such as self-supervised learning and finetuning. Next, you will learn about two types of FMs: text-to-text models and text-to-image models. Finally, you will learn about the functionality and use cases of LLMs, the subset of foundation models that most often utilize prompt engineering.

Session 2: Fundamentals of Prompt Engineering

In this lesson, you are introduced to prompt engineering, the set of practices that focus on developing, designing, and optimizing prompts to enhance the output of FMs for your specific business needs. This Session first defines prompt engineering and describes the key concepts and terminology of prompt engineering. Then, the Session uses an example prompt to show the different elements of a prompt. Finally, the Session provides a list of general best practices for designing effective prompts.

Session 3: Basic Prompt Techniques

In this lesson, you will learn about basic prompt engineering techniques that can help you use generative AI applications effectively for your unique business objectives. First, the Session defines zero-shot and few-shot prompting techniques. Then, the Session defines chain-of-thought (CoT) prompting, the building block for several advanced prompting techniques. This Session provides tips and examples of each type of prompt technique.

Session 4: Advanced Prompt Techniques

WhatsApp/ Tel: +1-929-672-1814
Email: info@genai-training.com
www.genai-training.com

In this lesson, you will be introduced to several advanced techniques including Self Consistency, Tree of Thoughts, Retrieval augmented generation (RAG), Automatic Reasoning and Tool-use (ART), ReAct, and LangChain. Examples are provided to show each technique in practice.

Session 5: Model-specific Prompt Techniques

In this lesson, you will learn how to engineer prompts for a few of the most popular FMs including Amazon Titan, Anthropic Claude, and AI21 Labs Jurassic-2. You will learn about the different parameters you can configure to get customized results from the models. Next you will learn about prompt engineering best practices for each of the models.

Session 6: Addressing Prompt Misuses

In this lesson, you will be introduced to adversarial prompts, or prompts that are meant to purposefully mislead models. You will be learning about prompt injection and prompt leaking, two types of adversarial prompts. You will be provided with examples of each.

Session 7: Mitigating Bias

In this lesson, you will learn how bias is introduced into models during the training phase and how that bias can be reproduced in the responses generated by an FM. You will learn how biased results can be mitigated by updating the prompt, enhancing the dataset, and using training techniques.

Session 8: Lab works

Lab Works: Agents perform the following tasks:
- Extend foundation models to understand user requests and break down the tasks that the agent must perform into smaller steps.
- Collect additional information from a user through natural conversation.
- Take actions to fulfill a customer's request by making API calls to your company systems.
- Augment performance and accuracy by querying data sources.

To use an agent, you perform the following steps:
1. (Optional) Create a knowledge base to store your private data in that database. For more information, see Knowledge bases for Amazon Bedrock.
2. Configure an agent for your use case and add at least one of the following components: • At least one action group that the agent can perform.
To learn how to define the action group and how it's handled by the agent, see Create an action group for an Amazon Bedrock agent.
• Associate a knowledge base with the agent to augment the agent's performance. For more information, see Associate a knowledge base with an Amazon Bedrock agent.
3. (Optional) To customize the agent's behavior to your specific use-case, modify prompt templates for the pre-processing, orchestration, knowledge base response generation, and post-processing steps that the agent performs. For more information, see Advanced Prompts in Amazon Bedrock.
4. Test your agent in the Amazon Bedrock console or through API calls to the TSTALIASID. Modify the configurations as necessary. Use traces to examine your agent's reasoning process at each step of its orchestration. For more information, see Test an Amazon Bedrock agent and Trace events in Amazon Bedrock.
5. When you have sufficiently modified your agent and it's ready to be deployed to your application, create an alias to point to a version of your agent. For more information, see Deploy an Amazon Bedrock agent.
6. Set up your application to make API calls to your agent alias.

7. Iterate on your agent and create more versions and aliases as necessary.

Session 9 &10: Projects