



GenAI for Workshop

ANI_425 | Expert-Led Live | Automation and Insights | ⚙️⚙️⚙️

Course Duration: 3 days

Unleash the full potential of Large Language Models (LLMs) through hands-on exploration of Prompt Engineering in this workshop. We'll delve into the theoretical and practical aspects of crafting effective prompts, equipping you with the tools and techniques to unlock creative and informative outputs from LLMs like ChatGPT, Copilot for the Web, and Bard. The workshop will begin by demystifying the anatomy of prompts, breaking down their components and exploring different types including open-ended, constrained, and zero-shot prompts. We'll dive into the principles of effective prompting, focusing on techniques like clarity, context, and specificity to steer LLMs towards desired outputs. The use of GenAI by network engineers is then explored. To solidify your understanding, we'll provide opportunities to work on real-world case studies.

Intended Audience

Anyone curious about language models and interested in exploring their potential through effective prompting.

Objectives

After completing this course, the learner will be able to:

- Compare and contrast types of AI
- Sketch the AI application deployment choices
- Sketch high-level LLM architecture
- Provide an intuitive explanation of how LLMs work
- List various popular LLMs and communities
- Engineer prompts for various tasks
- Sketch the process to enhance LLM with RAG and VectorDB

Outline

1. Introduction to AI
 - 1.1 Types of AI
 - 1.2 Gen AI use cases
 - 1.3 Application deployment choices
2. Introduction to LLM
 - 2.1 Architecture
 - 2.2 How It works
 - 2.3 Popular LLMs
 - 2.4 CommunitiesExercise: Explain LLM and how they work
3. Interacting with LLM
 - 3.1 Ways to customize LLM responses
 - 3.2 Incorporating your own data
 - 3.3 APIs for LLMsExercise: Access LLM with APIs
4. Prompt Engineering with playground
 - 4.1 The Magic of PromptsExercise: Basic prompts with public LLM
 - 4.2 Prompt Anatomy
 - 4.3 Zero-shot PromptsExercise: Basic prompts
 - 4.4 Few-shot learningExercise: Multi-shot prompts
 - 4.5 Chain of Thought (CoT)Exercise: CoT Exercise
 - 4.6 Submission Cycle

Exercise: Submission Cycle Exercise

5. Engineering prompts for specific tasks
 - 5.1 Market researchExercise: Content generation

6. RAG
 - 6.1 Why do need RAG?
 - 6.2 RAG Architecture
 - 6.3 VectorDB and tokenization
 - 6.4 Importance of search criteriaExercise: Using RAG with LLM

7. Fine-tuning
 - 7.1 What is fine-tuning?
 - 7.2 How to fine-tune?
 - 7.3 When to use fine-tuning?

8. Putting It All Together
9. Ask us about using your platforms/data