



GenAI, ML and Prompt Engineering Workshop

ANI_427 | Expert-Led Live | Automation and Insights | ⚙️⚙️⚙️

Course Duration: 4 days

The use of Artificial Intelligence (AI) is growing in every field, and all businesses are embracing Machine Learning (ML) and Generative AI (GenAI). This workshop covers ML/DL and AI techniques and their applications. We discuss when GenAI is appropriate and when other AI techniques are more suitable for use cases like planning, predictive maintenance, and performance management. We also cover Large Language Models (LLMs), their architecture, and use cases in business. Techniques to adapt or enhance GPT models, such as prompt engineering, Retrieval Augmented Generation (RAG), and fine-tuning, are included. The workshop features exercises to equip you with tools and techniques for creative and informative outputs from LLMs. The workshop can be customized with network-specific use cases.

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 the high-level LLM architecture
- Provide an intuitive explanation of how LLMs work
- List popular LLMs and communities
- Create engineer prompts for various tasks
- Sketch the process to enhance LLM with RAG and Vector databases

Outline

1. Introduction to AI
 - 1.1 Types of AI - Predictive, Generative, etc.
 - 1.2 AI use cases in different areas of business
 - 1.3 Application deployment choices
 2. Non-Generative AI
 - 2.1 ML and DL defined
 - 2.2 Types of learning
 - 2.3 Predictive AI techniques
 - 2.4 Types of ML/DL models
 - 2.5 Use cases and associated techniquesExercise: Match techniques with a business problem
 3. Introduction to GenAI and LLM
 - 3.1 GenAI Application Architecture
 - 3.2 How it works
 - 3.3 Popular LLMs
 - 3.4 Various LLM CommunitiesExercise: Explain LLMs and how they work
 4. Interacting with LLM
 - 4.1 Ways to customize LLM responses
 - 4.2 Incorporating your own data
 - 4.3 APIs for LLMsExercise: Access LLM with APIs
 5. Prompt Engineering with Playground
 - 5.1 The magic of promptsExercise: Basic prompts with public LLM
 - 5.2 Prompt anatomy
 - 5.3 Zero-shot prompts
 - 5.4 Few-shot learning
 - 5.5 Chain of Thought (CoT)
 - 5.6 Submission Cycle
- Exercise: Multi-shot prompts
-
- Exercise: CoT Exercise
-
- Exercise: Submission Cycle Exercise
6. Engineering prompts for specific tasks
- 6.1 Custom Task 1
- 6.2 Custom Task 2
- Exercise: Custom Task 1
-
- Exercise: Custom Task 2
7. Retrieval Augmented Generation (RAG)
 - 7.1 Why do we need RAG?
 - 7.2 RAG Architecture
 - 7.3 VectorDB and tokenization
 - 7.4 Importance of search criteriaExercise: Using RAG with LLM
8. Fine Tuning
 - 8.1 What is fine tuning?
 - 8.2 How and when to fine tune?
9. Putting It All Together
10. Ask us about using your platforms/data