Cloud Contraptions
TRAINING & CONSULTING

To discuss this course and customizations:
Call: 434-509-5680 or Email: sales@cloudcontraptions.com

Prompt Engineering for Programmers

Class Duration

21 hours of live training delivered over 3-5 days to accommodate your scheduling needs

Student Prerequisites

- No prompt engineering or generative AI experience is required
- For the programming exercises, experience with selected programming language is needed
- The course can be taught with Java, Python, JavaScript/TypeScript, C#, Go, or Rust

Target Audience

Designed for experienced programmers, ML and platform engineers, architects, and technical leads responsible for bringing LLM capabilities into production systems. Ideal for teams modernizing search, support, and developer productivity tools, as well as managers and L&D partners who evaluate technical upskilling with tangible ROI. Attendees should be comfortable with APIs and software delivery; the course meets them at a professional level and focuses on techniques that improve accuracy, reliability, and time-to-value.

Description

This hands-on course equips software professionals to turn large language models into dependable, revenue-impacting systems. You will master the core and advanced techniques of prompt engineering—zero-shot, few-shot, fine-tuning, RAG, chain-of-thought, prompt chaining, meta-prompting, self-consistency, generate-knowledge prompting, tree-of-thought, PAL, and ReAct—then operationalize them with the Ollama API. Through practical labs, you will iterate prompts, summarize and infer from documents, transform and expand text for real software workflows, and architect a robust RAG pipeline. The capstone builds a production-minded chatbot that integrates MCP servers for secure tool access and orchestration. Expect actionable patterns that cut cycle time, reduce hallucinations, and translate directly into faster delivery and measurable business value.

Cloud Contraptions
TRAINING & CONSULTING

To discuss this course and customizations:
Call: 434-509-5680 or Email: sales@cloudcontraptions.com

Learning Outcomes

- Define prompt engineering, explain its evolution, and identify high-value applications across domains.
- Select among zero-shot, few-shot, fine-tuning, and RAG strategies to elicit accurate, business-ready responses.
- Iteratively refine prompts using structured techniques and implement programmatic prompting via the Ollama API.
- Build extractive and abstractive summarization workflows and automate them with the Ollama API.
- Infer insights and generate recommendations from text using the Ollama API while controlling for accuracy.
- Transform and expand text (tone, voice, format, and context) programmatically to fit software development needs.
- Apply advanced prompt engineering methods: chain-of-thought, prompt chaining, meta-prompting, self-consistency, generate-knowledge prompting, tree-of-thought, PAL, and ReAct.
- Design and ship a domain-specific chatbot that leverages RAG and integrates MCP servers using the Ollama API.

Training Materials

All students receive comprehensive courseware covering all topics in the course. Courseware is distributed via GitHub in the form of documentation and extensive code samples. Students practice the topics covered through challenging hands-on lab exercises. All students receive a downloadable MP4 recording of the training.

Software Requirements

Students will need a web browser to access the cloud-based virtual machine for the class.

Training Topics

Prompt Engineering

- What is Prompt Engineering?
- Definition and Importance
- Historical Context and Evolution
- Applications in Various Domains

To discuss this course and customizations:
Call: 434-509-5680 or Email: sales@cloudcontraptions.com

Understanding and Improving Responses

- Pretrained Models and Prompt Engineering
- Strategies for Effective Responses
- Zero-shot
- Few-shot
- Fine-tuning
- RAG (Retrieval-Augmented Generation)

Iterating to Improve Prompts

- Prompts are rarely Perfect the First Time
- Refining and Improving Prompts
- Zero-shot vs. Few-shot
- Strategies to Iterate and Improve
- Programmatically Prompt the Ollama API

Summarizing Text

- Understanding Summarization
- Types: Extractive vs. Abstractive
- Common Use Cases and Applications
- Programmatically Use Ollama API to Summarize Documents

Inferring from Text

- Basics of Inference
- Techniques for Accurate Inference
- Applications in Various Fields
- Use Ollama API to Make Recommendations

Transforming Text

- Understanding Text Transformation
- Common Strategies and Approaches
- Use Cases in Software Development
- Use Ollama to Transform the Voice and Tone of Responses

Expanding Text

- Basics of Text Expansion
- Key Concepts and Strategies
- Applications in Various Domains
- Use Ollama API to Provide Additional Context

Cloud Contraptions
TRAINING & CONSULTING

To discuss this course and customizations:
Call: 434-509-5680 or Email: sales@cloudcontraptions.com

Prompt Engineering Techniques

- Review Zero-Shot and Few-Shot
- Review Chain-of-Thought
- Review Prompt Chaining
- Meta-Prompting
- Self-Consistency
- Generate Knowledge Prompting
- Tree-of-Thought
- Retrieval Augmented Generation (RAG)
- Directional Stimulus Prompting
- Program-Aided Language Models
- ReAct

RAG (Retrieval-Augmented Generation)

- Introduction to RAG
- Key Concepts and Strategies
- Applications in Various Domains
- Enhance Prompt Engineering with RAG

Build a Chatbot

- Introduction to Chatbots
- Designing Effective Chatbot Prompts
- Historical Context and Evolution
- Build a custom chatbot using Ollama API and Advanced Prompt Engineering Techniques
- Integrate MCP servers into chatbot