Course Syllabus: Prompt Engineering

Course Title: Prompt Engineering: The Art of Communicating with AI

Target Audience: This course is suitable for a wide range of individuals, including content creators, marketers, technical writers, product managers, data scientists, and developers. It requires a keen interest in AI and an ability to think critically about language and problem-solving.

Course Level: This is a comprehensive program that starts with the basics and progresses to expert-level techniques.

Duration: 10 Weeks

Course Description:

Prompt engineering is the critical skill of the modern AI era. This course will guide you through the systematic process of designing, refining, and optimizing prompts to elicit accurate, relevant, and creative responses from large language models (LLMs). We will move beyond simple queries to explore advanced strategies, including multi-step reasoning, tool integration, and the ethical considerations of working with AI. By the end of this course, you will be able to "program" an LLM with natural language, unlocking its full potential for a wide array of applications.

Learning Objectives

Upon successful completion of this course, students will be able to:

- Understand the core mechanics of Large Language Models (LLMs) and their limitations.
- Master foundational prompt engineering techniques like zero-shot, few-shot, and persona prompting.
- Implement advanced reasoning strategies such as Chain-of-Thought (CoT) and Tree-of-Thought (ToT) to solve complex problems.
- Develop a systematic, iterative process for prompt creation and refinement.
- Apply prompt engineering to a variety of domains, including content creation, coding, data analysis, and creative writing.
- Recognize and defend against common prompt injection attacks and other security vulnerabilities.
- Understand the ethical implications and responsible practices of using generative AI.
- Build a portfolio of advanced prompts and the outputs they generated, demonstrating expert-level skill.

Course Structure: A Step-by-Step Learning Path

Part 1: Foundational Prompting (Weeks 1-3)

This section builds a strong conceptual foundation in how LLMs work and the essential techniques for basic communication.

Week 1: The Foundation of Generative AI

- The history and evolution of LLMs.
- How do LLMs "think"? Understanding the attention mechanism and tokenization.
- Introduction to popular LLMs (e.g., GPT, Gemini, Claude, Llama).
- The concept of the "prompt" and its components (role, task, context, format).
- **Hands-on Lab:** Writing your first few effective prompts.

Week 2: Core Prompting Techniques

- **Zero-shot prompting:** The most basic form of prompting.
- **Few-shot prompting:** Providing examples to guide the model's behavior.
- **Persona prompting:** Assigning a specific role or persona to the AI.
- **Instruction tuning:** How specific instructions dramatically improve output quality.
- **Hands-on Lab:** Use zero-shot and few-shot techniques to guide an LLM for different tasks like summarization and classification.

Week 3: Prompting for Structured & Creative Tasks

- Prompting for specific output formats (e.g., JSON, CSV, Markdown).
- Using prompts for creative applications: songwriting, storytelling, and ad copy.
- The art of providing context to generate highly specific and accurate outputs.
- **Hands-on Project:** Create a series of prompts to generate a structured content plan for a marketing campaign.

Part 2: Advanced Reasoning & Optimization (Weeks 4-6)

This section moves beyond basic prompting to advanced, strategic methods that tap into the LLM's reasoning capabilities.

Week 4: Unlocking Advanced Reasoning

- **Chain-of-Thought (CoT) prompting:** Teaching the model to think step-by-step.
- **Zero-shot CoT:** Using a single phrase to trigger advanced reasoning.
- **Tree-of-Thought (ToT) prompting:** Exploring multiple reasoning paths to find the best solution.
- **Hands-on Project:** Apply CoT and ToT to solve complex analytical problems.

Week 5: The Iterative Prompting Workflow

- The importance of iteration: Why a prompt is never perfect on the first try.
- A/B testing prompts to determine which performs better.
- Evaluating AI outputs: metrics for accuracy, creativity, and fluency.
- Using prompt libraries and templates for efficiency.
- **Hands-on Lab:** Iteratively refine a single prompt to achieve a specific, high-quality outcome.

Week 6: Prompting for Code & Complex Tasks

- Using LLMs for coding assistance, code generation, and debugging.
- Prompting for data analysis, including generating code for data visualization.
- Advanced techniques for eliciting information and synthesizing complex topics.
- **Hands-on Project:** Use prompts to generate code for a specific programming problem and debug a broken snippet.

Part 3: Expert-Level Strategies & Safety (Weeks 7-10)

This final section focuses on the expert-level applications, ethical considerations, and the technical side of prompt engineering.

Week 7: Defending Against Prompt Injection

- Understanding prompt injection, jailbreaking, and other adversarial attacks.
- Techniques for building robust "guardrails" around your prompts.
- Implementing a "separation of instructions" approach.
- **Hands-on Lab:** Test the robustness of a prompt by attempting to jailbreak it.

Week 8: Interacting with External Systems

- Introduction to API-based prompting.
- Using LLMs with tools and APIs to extend their capabilities (e.g., web search, calculators).
- The concept of Retrieval-Augmented Generation (RAG) and its importance.
- **Hands-on Project:** Build a system that uses an LLM to answer questions by searching a provided knowledge base.

Week 9: Ethical Prompting & Responsible AI

- Identifying and mitigating bias in AI-generated content.
- The ethics of deepfakes, plagiarism, and misinformation.
- The role of a prompt engineer in ensuring responsible AI use.
- Best practices for transparency and accountability.

Week 10: Final Project & The Future of Prompting

- **Capstone Project:** Design and build a comprehensive portfolio of advanced prompts, demonstrating mastery of the course material across different domains. You will document your process, iterations, and final results.
- The future of prompt engineering: The shift towards automated prompting and model fine-tuning.
- Career paths and opportunities in the Generative AI space.

Assignments & Grading

- **Weekly Hands-on Labs & Exercises:** 25%
- **Intermediate Projects (Weeks 3 & 6):** 30%
- **Final Capstone Project:** 35%
- **Prompt Iteration & Documentation:** 10%