BOARD INFINITY

COURSE 1 OF 3 · AI AGENTS: FROM FOUNDATIONS TO APPLICATIONS

Agent Foundations and Prompt Engineering

Full course syllabus - every module, video lesson, reading and assignment, from the agent loop to research automation.

Board Infinity
Instructor · Offered on Coursera

4	41	16	2 weeks
Modules	Video lessons	Assignments	At 10 hours a week

INSIDE THE COURSE

01	Introduction to AI Agents and Prompt-Driven Systems	6 hours to complete
02	Fundamentals of Prompt Engineering	5 hours to complete
03	Programmatic Prompt Execution and Chaining	5 hours to complete
04	Optimization, Debugging, and Observability	5 hours to complete

Enroll on Coursera

https://www.coursera.org/learn/agent-foundations-and-prompt-engineering

01 ABOUT THE COURSE

How agents work, then how to prompt them.

The course starts with what makes an AI system agentic: how agents perceive, reason and act, and how policy, memory, tools and environment fit together - ending with a minimal agent in Python.

It then covers prompt engineering, turns it into programmatic workflows in Python with chaining, APIs, tools and memory, and finishes with research automation, testing and evaluation - 41 videos, 13 readings and 16 assignments in all.

WHO IT'S FOR

AI learners
Building a structured understanding of how agents work.

Developers
Connecting LLMs, Python, APIs, tools and workflows.

Prompt engineers
A systematic approach to design, evaluation and debugging.

Automation builders
Turning repetitive research tasks into AI-powered systems.

WHAT YOU'LL LEARN

01 **Design AI agents** Module 1
Design and build AI agents using LLM-based architectures and decision flows

02 **Engineer prompts** Module 2
Craft high-quality prompts using few-shot, chain-of-thought, and role-based patterns

03 **Chain workflows** Module 3
Build programmatic prompt chains with Python for automated multi-step workflows

04 **Automate research** Module 4
Develop research automation and summarization pipelines with quality evaluation

02 SYLLABUS

Four modules, agents to automation.

2 weeks to complete at 10 hours a week. Flexible schedule - learn at your own pace.

01 Introduction to AI Agents and Prompt-Driven Systems

6 hours · 9 videos

The agent loop, core components, architectures and a minimal agent

BY THE END OF THIS MODULE YOU CAN

- Describe the structure, purpose and components of AI agents
- Explain the agent loop, and how policy, memory, tools and environment fit
- Compare agent architectures and build a minimal agent in Python

CONTENTS: 9 videos (72 min) · 4 readings (120 min) · 4 assignments (150 min)

PLUGIN: Quick Course Check-In (5 min)

TOPICS: What Are AI Agents? · Agent Architectures and Behaviors · Applications of AI Agents

You'll leave with: a clear model of how agents work, and a minimal agent in Python.

02 Fundamentals of Prompt Engineering

5 hours · 12 videos

Prompt structure, reasoning patterns, evaluation and debugging

BY THE END OF THIS MODULE YOU CAN

- Formulate effective prompts for task-specific outcomes
- Apply role-task-constraint structure, few-shot and chain-of-thought
- Evaluate, debug and organize prompts into reusable libraries

CONTENTS: 12 videos (82 min) · 3 readings (90 min) · 4 assignments (150 min)

TOPICS: Prompt Structures and Templates · Prompt Patterns and Strategies · Evaluating Prompt Quality

You'll leave with: a systematic way to write, evaluate and debug prompts.

03 Programmatic Prompt Execution and Chaining

5 hours
12 videos

Prompt chains in Python, tool-using agents, memory and debugging

BY THE END OF THIS MODULE YOU CAN

- Build chained prompts that pass context between tasks
- Give an agent tools and memory in Python
- Debug chains, then design and run a chained agent end to end

CONTENTS: 12 videos (97 min) · 3 readings (90 min) · 4 assignments (120 min)

TOPICS: Prompt Chaining Basics · Dynamic Prompting with APIs · Error Handling in Prompt Workflows

You'll leave with: a chained agent with tools and memory, built and run in Python.

04 Optimization, Debugging, and Observability

5 hours
8 videos

A research agent project - data layer, summarization and evaluation

BY THE END OF THIS MODULE YOU CAN

- Apply optimization, debugging and observability to agents
- Design a research system's architecture, data layer and data models
- Build a research agent, then test and evaluate its summaries

CONTENTS: 8 videos (55 min) · 3 readings (90 min) · 4 assignments (150 min)

TOPICS: Debugging RAG Pipelines · Observability and Tracing Tools · Response Optimization

You'll leave with: a working research agent, and a way to evaluate what it produces.

03 TOOLS & SKILLS

What you work with, and come away with.

Large language models — Reasoning core

The reasoning core behind the agents and workflows you design.

In this course: From the agent loop to tool calling and evaluation.

Python — Automation

Notebook demonstrations throughout the course.

In this course: A minimal agent, then an end-to-end chained agent.

LLM APIs — Integration

Calling models programmatically.

In this course: Dynamic prompting and API chaining in module three.

LangChain — Framework

Named in the course FAQ among the tools for agents and workflows.

In this course: Context for building agents beyond plain Python.

SKILLS YOU LEAVE WITH

LLMs & agents

- Large Language Modeling
- LLM Application
- Generative AI Agents
- Agentic systems

Build & quality

- Tool Calling
- Debugging
- Verification & Validation
- Quality Assurance

Workflows

- Prompt Engineering
- AI Workflows
- Agentic Workflows
- AI Orchestration

04 CERTIFICATE & CAREER

Finish with a shareable certificate.

Add it to LinkedIn

Add this credential to your LinkedIn profile, resume, or CV.

Print it

Saved to your Coursera Accomplishments page, where you can print it.

Earned, not attended

Awarded on a graded quiz in each of the four modules.

Course 1 of 3

Part of AI Agents: From Foundations to Applications.

POSSIBLE ROLE DIRECTIONS

AI / ML Engineer · Generative AI Developer · AI Automation Developer · Prompt Engineer · AI Solutions Developer · AI Agent Developer

Board Infinity

Instructor · Offered on Coursera · 287 courses

Intermediate level · English + 8 languages · Flexible schedule

Ready to build your first agent?

Enrollment, lessons, grading and the certificate all happen on Coursera.

Enroll on Coursera

Course page: coursera.org/learn/agent-foundations-and-prompt-engineering

Want a coach alongside the course?

Board Infinity 1:1 coaching is a separate, optional Board Infinity offering - it is not part of the Coursera course. Details: https://www.boardinfinity.com/career-coaching/

Board Infinity is an independent education provider. Coursera and the Coursera logo are trademarks of Coursera, Inc. This independent educational resource is not affiliated with, endorsed by, sponsored by, or officially associated with any company, organization or certification body unless explicitly stated. Role directions are not guaranteed outcomes.