APPLIED AND GENERATIVE AI

Introduction to Agentic AI

Level: Foundation • 2 days (expandable to 3) • Virtual, In-person

Overview

"Agent" is the most used and least defined word in AI right now. Behind the hype is a real and important shift: systems that do not just answer a question but pursue a goal, planning steps, using tools, observing results, and correcting course along the way. The hard part is not understanding the idea, it is developing the judgment to know when autonomy genuinely helps, when it quietly adds risk, and how to tell a real agentic capability from a rebranded chatbot.

This is a hands-on, foundation course. It builds understanding in dependency order: first what actually distinguishes an agent from a chat model, then how agents plan and reason, then tools, the mechanism that lets an agent act on the world. With the mechanics in place, the course turns to the judgment that matters most: where agents help and where they hurt, how oversight and guardrails keep them trustworthy, and how to read today's agent landscape with clear eyes. Rather than survey every framework, the course goes deep on the concepts that outlast them. Every module includes a lab, and each module builds on the one before it.

Who Should Attend

- Professionals who keep hearing "agents" and want a grounded, hype-free understanding
- Developers and analysts sizing up whether agentic AI fits a problem they own
- Managers and decision makers evaluating agent products and vendor claims

Learners new to generative AI itself should take *Introduction to Generative AI* first.

Prerequisites

- Basic familiarity with a generative AI tool such as ChatGPT or Claude
- No programming experience is required
- Curiosity about where AI autonomy helps and where it should stop

What You Will Learn

- Explain what makes an AI system agentic and how agents differ from chat models
- Describe the agent loop: planning, acting through tools, observing, and adjusting
- Explain how tool use works and why it is the key to agents doing real work
- Judge which tasks suit an agent and which are better served by simpler approaches
- Describe the oversight and guardrails that make agent use trustworthy
- Evaluate agent products and claims, and recognize common failure modes

Course Outline

Day one: what agents are and how they work

- From Chatbot to Agent
  - Answering versus pursuing: what changes when a system has a goal

  - The agent loop: plan, act, observe, adjust
  - Agents you may already be using without noticing
  - Lab: give the same task to a chat model and an agent and compare how each got its result
- How Agents Plan and Reason
  - Breaking a goal into steps, and revising the plan as results come in
  - What the agent remembers: context, memory, and their limits
  - Why plans go wrong: drift, dead ends, and loops
  - Lab: watch an agent work through a multi-step task and narrate its plan and corrections
- Tools: How Agents Act on the World
  - What a tool is: search, files, code, and other systems an agent can call
  - How the agent chooses a tool and interprets what comes back
  - Connecting agents to systems, and a first look at standards like MCP
  - Lab: run an agent with different tool sets and observe how its abilities change

Day two: judgment, trust, and the landscape

- When Agents Help and When They Hurt
  - The autonomy spectrum: from suggestion to supervised action to full delegation
  - What makes a task agent-friendly: clear goals, checkable results, low blast radius
  - The honest case against agents: tasks where simpler AI, or no AI, wins
  - Lab: sort a set of real tasks along the autonomy spectrum and defend each placement
- Guardrails, Oversight, and Trust
  - Human in the loop: approval points and knowing what an agent did
  - Permissions and limits: deciding what an agent may touch
  - Recognizing failure early: the signs an agent is stuck or off course
  - Lab: design the guardrails for a realistic agent deployment scenario
- Reading the Agent Landscape
  - Coding agents, workflow agents, research agents, and multi-agent systems
  - Cutting through marketing: questions that reveal what a product actually does
  - Where agentic AI is heading, and what to watch
  - Lab: evaluate a real agent product against a checklist built from the course

Extended Version

The three-day version keeps the same gradient and adds depth and applied practice:

- More hands-on time directing agents on longer, realistic tasks
- A closer look at how agents connect to real systems, as a bridge to *Building Agentic AI with the Model Context Protocol (MCP)*
- Organizational readiness: policy, risk, and where to pilot agents first
- A capstone in which teams design, guardrail, and defend an agent-based solution to a real scenario