APPLIED TECHNOLOGY ACADEMY

800.674.3550
Info@AppliedTechAc.com
AppliedTechnologyAcademy.com

Introduction to Agentic AI

Applied Technology Academy

A practical introduction to AI systems that can pursue goals, make decisions, use tools and perform multi-step tasks with varying degrees of autonomy. You'll explore how agentic AI differs from traditional and generative AI, examine the essential building blocks of agentic systems — models, goals, instructions, tools, memory, state and decision-making — and learn how single-agent and multi-agent architectures address complex business problems. The course also covers evaluating and operating AI agents in production: testing, observability, AgentOps, security, governance and human oversight.

LEVEL	DURATION	DELIVERY	PREREQUISITES	APPROACH
Introductory	**2 days**	**Instructor-led · live online or classroom**	**None required**	**Technology-neutral**

Course Outline

- **Module 1: Foundations of Agentic AI**
  - Places modern AI agents within the broader evolution of AI. Defines agentic AI and its key characteristics (autonomy, persistence, goal-directed behavior); differentiates it from predictive, reactive and traditional AI; traces the evolution from symbolic AI to LLM-powered agents; covers agent types (reactive, deliberative, hybrid), the agent-environment interaction loop, single- vs multi-agent systems, and how assistants, copilots, workflows and autonomous agents differ.
- **Module 2: Introducing AI Agents**
  - How agents extend generative AI beyond request-and-response by combining models with goals, instructions, tools, memory and decision-making. Covers the core components of an agent; how agents interpret goals and progress toward them; the agent execution loop and the relationship between reasoning and action; tool use for interacting with external systems and data; the role of memory and state; and common patterns for combining models, tools and application logic.
- **Module 3: Prompting for Agentic AI**
  - Prompting techniques designed for agentic workflows. Writing prompts that establish clear goals and responsibilities; designing instructions that guide planning, reasoning and decision-making; reusable prompt templates for dynamic goals and runtime information; prompts that guide tool selection; incorporating memory and prior state; and using constraints and behavioral boundaries to reduce unintended behavior across multi-step tasks.
- **Module 4: Architectures and Patterns for Agentic Systems**
  - The building blocks of agentic architectures and how models, agents, tools, memory and workflows work together. Comparing deterministic workflows with autonomous agents; common patterns (chaining, routing, parallelization, orchestration, agent loops); distinguishing conversational context, agent memory and application state; and selecting an appropriate

architecture based on requirements and risk.

- **Module 5: Tools, Decision-Making and Agent Behavior**
  - How agents decide what actions to take, select tools, evaluate results and adapt. Goal decomposition; tool selection and invocation; distinguishing model reasoning from external tool execution; the roles of planning, execution, reflection and evaluation; handling tool failures; recognizing agent loops, ineffective decisions and false task completion; and balancing autonomy with deterministic controls.
- **Module 6: Multi-Agent Systems and Coordination**
  - Multiple specialized agents collaborating rather than one general-purpose agent. When multi-agent designs help; defining agent roles and responsibilities; communication and coordination patterns; task delegation; combining and evaluating results; sharing context, memory and state; identifying coordination failures; and weighing the benefits against operational complexity.
- **Module 7: Agentic AI in Real-World Applications**
  - Practical enterprise applications and a framework for deciding when an agent-based approach fits. Common use cases; distinguishing problems needing simple LLM integration, deterministic workflows or autonomous agents; how agents interact with APIs, databases, documents and enterprise systems; agentic patterns for research, document processing, customer support, operations and decision support; determining the right level of autonomy; and where human-in-the-loop belongs.
- **Module 8: Evaluating, Deploying and Operating AI Agents**
  - Testing, evaluation, deployment, observability and AgentOps for production. Success metrics (task completion, accuracy, reliability, latency, cost); evaluating agent trajectories (decisions, tool calls, observations, outcomes); recognizing failure modes (hallucinations, agent loops, tool failures, false success signals); sandbox and simulation testing; stateless/stateful/service-based deployment; tracing and observability; auditability; and continuous evaluation.
- **Module 9: Security, Governance and Responsible Agentic AI**
  - Managing the risks agents introduce beyond traditional generative AI. Security and operational risks of autonomous systems; prompt injection; excessive permissions and unrestricted tool access; data security and privacy; least-privilege access; guardrails and validation controls; human-in-the-loop and human-approval controls; balancing autonomy with accountability; transparency, fairness and explainability; and governance, compliance and regulatory considerations.

Prerequisites

Designed for software developers, architects, technical leads, IT professionals, AI practitioners, business analysts and technology decision-makers who want to understand the concepts, architectures and practical applications of agentic AI. A general understanding of AI and generative AI is helpful but not required; no prior experience developing AI agents or with a specific framework is necessary. The course is technology-neutral, focusing on concepts and patterns that apply across modern agentic AI platforms and frameworks.

Learning Objectives

- Explain the fundamental concepts and defining characteristics of Agentic AI
- Differentiate AI agents from traditional AI, generative AI applications, copilots and automated workflows
- Identify the core components of an AI agent — models, goals, instructions, tools, memory and state
- Explain how agents reason about goals, make decisions, use tools and perform multi-step tasks
- Design prompts and instructions that guide agent planning, decision-making and behavior
- Compare common architectures and patterns for building agentic systems
- Distinguish deterministic workflows from autonomous agent-based approaches
- Describe how multiple specialized agents collaborate, delegate tasks and coordinate activities
- Identify enterprise use cases suitable for agentic AI and the appropriate level of autonomy
- Evaluate agent performance using task outcomes, execution trajectories, tool interactions, reliability and cost
- Apply testing, tracing, observability and AgentOps practices to agentic systems

Identify security risks — prompt injection, excessive permissions and tool misuse — and apply guardrails, least-privilege access and human-in-the-loop controls

Evaluate agentic AI solutions from the perspectives of security, governance, accountability and responsible AI

Ready to enroll? Call 800.674.3550 or email Info@AppliedTechAc.com — private team cohorts, GSA MAS purchasing, military credentialing funding, VA GI Bill and VR&E and student financing available.