

Document Generated: 04/02/2026

Learning Style: Virtual Classroom

Technology:

Difficulty: Beginner

Course Duration: 2 Days

Next Course Date: **May 14, 2026**

Introduction to Agentic AI (TTAI2900)



About This Course:

This course introduces the principles and practices of Agentic AI, a form of intelligent systems that moves beyond the limits of chat-based AI or productivity assistants. While familiar tools like ChatGPT or Microsoft Office's Copilot respond

to user prompts in the moment, Agentic AI systems are designed to act more autonomously. They can plan steps toward a goal, use tools and data, remember past interactions, and continue working even when not being directly prompted.

Across two days, the course introduces the foundations of agent-based systems, showing how they differ from predictive or reactive AI approaches. It explains the core architectures that drive agent behavior, including cycles of perception, planning, action, and reflection, and highlights how prompts and memory shape their decision-making. We will also look at more advanced setups, such as groups of agents working together, and connects these ideas to real-world applications like research automation, customer support, and business workflows.

Course Objectives:

- Explain what Agentic AI is and how it differs from traditional, predictive, and chat-based AI.
- Describe the core architectures and memory systems that drive agent behavior.
- Recognize how prompts and goal-setting influence agent planning and decision-making.
- Understand how agents make choices, adapt through feedback, and handle uncertainty.
- Identify the features of multi-agent systems and how agents coordinate with each other.
- Relate Agentic AI concepts to real-world application.
- Discuss ethical risks and apply responsible design principles for safe deployment.

Audience:

- This course is designed for business leaders, product managers, consultants, and operations professionals who want to understand how Agentic AI can be applied in practice. The course is intended for non-developers who need to evaluate AI opportunities, guide adoption, or work effectively with technical teams on agent-based systems.

Prerequisites:

- Participants are expected to have a basic understanding of artificial intelligence. Prior exposure to large language model-based tools like ChatGPT or Microsoft Copilot is also necessary, as the course builds

directly on those concepts. While no programming or coding experience is required, learners should be comfortable with general technology concepts and able to follow discussions about how AI systems are structured and applied.

Course Outline:

1. Foundations of Agentic AI

Agentic AI refers to systems that act on their own, keep working toward goals, and adjust as conditions change. It builds on decades of AI evolution, from symbolic reasoning to today's adaptive agents. Core ideas include the types of agents (reactive, deliberative, hybrid), how agents interact with their environment, and the difference between single and multi-agent setups.

- Define Agentic AI and its traits: autonomy, persistence, goal-directedness.
- Differentiate Agentic AI from predictive, reactive, and traditional AI.
- Trace the historical path from symbolic AI to agents.
- Identify agent types and the agent-environment loop.
- Distinguish single-agent vs. multi-agent systems.

2. Prompting for Agentic AI

Prompts are more than questions, they shape how agents plan, reflect, and act. By writing prompts that set goals, reference memory, and encourage reasoning, it's possible to guide agents into more complex behaviors. Prompt engineering links directly to how agents are structured and how they perform in real situations.

- Write prompts that enable planning and tool use.
- Create templates for dynamic goal-setting and reasoning.
- Use prompts that reference memory or past state.
- Simulate decision chains and tool choices through prompting.

3. Architectures of Agentic Systems

The architecture of an agent is built around a loop: perception, planning, action, and reflection. Memory systems extend what agents can do, and external tools or APIs expand their reach. Planning can follow rules or be driven by large language models, and frameworks.

- Recognize the perception ? planning ? action ? reflection cycle.
- Describe short-term, long-term, and episodic memory.
- Explain how agents connect to tools and APIs.
- Compare rule-based vs. LLM-driven planning.
- Understand frameworks like LangChain, ReAct, and AutoGen.

4. Decision-Making and Behavior

Agents make choices by weighing options, following plans, or using heuristics. They learn through feedback, improve over time, and handle uncertainty with retries or fallback logic. Their interactions with the environment often produce complex, sometimes surprising behaviors.

- See how agents select actions with scoring, planning, or heuristics.
- Understand feedback, reflection, and reinforcement.
- Recognize uncertainty handling, retries, and fallbacks.
- Explore how environment and feedback shape behavior.
- Identify cases of emergent behavior.

5. Multi-Agent Systems and Coordination

When agents work together, they may be centrally controlled or act independently while sharing information. Communication methods allow them to negotiate, divide work, and collaborate effectively, though coordination also brings risks. Role specialization makes multi-agent setups resemble real teams.

- Compare centralized and decentralized designs.
- Describe communication and message-passing methods.
- Explain negotiation, delegation, and shared tasks.
- Weigh the benefits and risks of multi-agent systems.
- Understand role specialization in agent teams.

6. Real-World Applications

Agentic AI is already being used in customer support, research, and automation. Some agents handle short, one-off tasks, while others persist in ongoing workflows. Examples like AutoGPT and Adept show different ways agents are applied, and adoption trends reveal how industries are starting to integrate them.

- Categorize uses across domains like support, research, and automation.
- Distinguish between short-term tasks and long-running workflows.
- Examine examples such as AutoGPT, Adept, and AgentOps.
- Match design strategies to specific business needs.
- Note industry adoption patterns and trends.

7. Evaluation and Deployment

Evaluating agents means measuring whether they meet goals and work reliably. Failures like hallucinations or false signals of success are common challenges. Agents can be tested in controlled settings before deployment, then launched in either stateless or stateful forms. Ongoing monitoring keeps their behavior on track once live.

- Define success metrics like task completion and reasoning quality.
- Recognize failure modes including hallucinations and false success.
- Describe sandbox testing approaches.
- Compare deployment models: stateless, stateful, or service-based.
- Apply monitoring strategies in production.

8. Ethics and Responsible Design

Autonomous systems raise ethical concerns and risks, from misuse of resources to unpredictable decisions. Responsible design means adding guardrails, monitoring, and human oversight. Broader principles like transparency, fairness, and accountability guide responsible deployment, while regulation and governance are emerging to address long-term risks.

- Identify risks like autonomy creep and misuse.
- Explain guardrails, monitoring, and human-in-the-loop approaches.
- Consider the balance between autonomy and control.
- Explore strategies for aligning agent behavior with human intent.
- Apply ethical principles of transparency, fairness, accountability, and explainability.
- Anticipate future governance and regulation challenges.