

Lab 001: What are AI Agents?

Level: L50 Path: All paths Time: ~15 min ? Cost: Free ? No account needed

What You'll Learn

What an AI agent is (and isn't) The four core properties of an agent: perception, memory, reasoning, action How agents differ from simple chatbots and from traditional software

Real-world examples of AI agents

Introduction

The word "agent" is everywhere in AI right now ? but what does it actually mean?

An AI agent is software that uses a Large Language Model (LLM) as its reasoning engine to autonomously pursue a goal, deciding *what to do* and *which tools to call* at each step ? without you pre-programming every possible path.

The key word is autonomous: the agent doesn't just answer a question. It plans, acts, observes the result, and adjusts.

The Four Properties of an Agent

1. ? Perception The agent receives input ? a user message, a file, an API response, an event.
2. ? Memory The agent stores information across turns: Short-term (context window): the current conversation Long-term (vector store / DB): facts, history, retrieved documents
3. ? Reasoning The LLM decides *what to do next*: answer directly, call a tool, ask a clarifying question, or break the goal into sub-steps.
4. ? Action The agent *does something*: calls an API, queries a database, writes a file, sends an email, triggers another agent.

AI Agent Loop ? Perceive, Reason, Act, Observe

question "? Check Your Understanding" In the agent loop, what happens after the agent acts (e.g., calls an API)?

??? success "Answer" The agent observes the result ? the tool output is fed back into context so the LLM can reason over the new information and decide the next step. This closes the loop: perceive ? reason ? act ? observe ? reason again.

Agent vs. Chatbot vs. Traditional Software

	Traditional Software	Chatbot	AI Agent
Logic defined by	Developer	Developer	LLM (at runtime)
Handles new situations	? Only what's coded	?? Within trained patterns	? Adapts dynamically
Uses tools	? Hardcoded	?? Limited	? Discovers and calls tools
Multi-step reasoning	? ? ? ?	? ? ? ?	? Predictability
Predictability	? Very predictable	? ?	? Less predictable

tip "When NOT to use an agent" Agents are powerful but complex. Use a simple LLM call for single-turn Q&A. Use an agent only when the task requires multi-step reasoning, tool use, or dynamic decision-making.

question "? Check Your Understanding" A traditional chatbot follows a pre-programmed decision tree. How does an AI agent differ when it encounters a situation the developer didn't anticipate?

??? success "Answer" An AI agent uses the LLM to adapt dynamically at runtime ? it reasons about the new situation and decides what to do, even if that exact scenario was never coded. A traditional chatbot can only handle what was explicitly programmed.

question "? Check Your Understanding" When should you use a simple LLM call instead of building a full AI agent?

??? success "Answer" Use a simple LLM call for single-turn Q&A tasks that don't require multi-step reasoning, tool use, or dynamic decision-making. Agents add complexity and should only be used when the task truly needs autonomy.

Real-World Examples

Agent | What it does | GitHub Copilot | Reads your code, suggests completions, chats, runs commands | OutdoorGear Sales Agent *(this repo's capstone scenario)* | Queries product and order data, generates charts, interprets sales trends | Microsoft 365 Copilot | Reads emails, calendar, files, drafts replies, summarizes meetings | AutoGen research agent | Spawns sub-agents to search, synthesize, and write a report | question "? Check Your Understanding" Which of the four core agent properties (perception, memory, reasoning, action) is primarily responsible for the agent deciding *what to do next*?

??? success "Answer" Reasoning. The LLM uses reasoning to decide the next step ? whether to answer directly, call a tool, ask a clarifying question, or break the goal into sub-steps.

Perception handles input, memory stores context, and action executes the decision.

Key Terms

Term | Definition | LLM | Large Language Model ? the AI brain (e.g., GPT-4o, Phi-4) | Tool / Function | A function the LLM can call (e.g., search_database, send_email) | Context window | The "working memory" of the LLM ? everything it can see at once | Prompt | The instructions + context sent to the LLM | Token | The unit LLMs process ? roughly 3/4 of a word | Grounding | Connecting agent responses to real, verified data |

? Knowledge Check

question "Q1 (Multiple Choice): Which of the following best describes an AI agent?"

- A) A chatbot that follows a pre-programmed decision tree - B) A machine learning model fine-tuned on your company's data - C) Software that uses an LLM to autonomously pursue a goal, deciding what to do and which tools to call at each step - D) A rule-based keyword matching system that routes users to FAQs

??? success "? Reveal Answer" Correct: C

An AI agent uses an LLM as its *reasoning engine* to autonomously decide what to do at each step ? including which tools to call, when to loop, and when to stop. Option A describes a traditional chatbot. Option B is fine-tuning (changes model behavior, not agent structure).

Option D is a classic NLP routing system.

question "Q2 (Multiple Choice): In the perceive ? reason ? act ? observe loop, what is the purpose of the 'observe' step?"

- A) The agent reformulates the original user query before reasoning - B) The agent receives the result of an action and adds it back to context for the next reasoning step - C) The agent calls the LLM to generate a final answer - D) The agent saves the conversation to long-term memory

??? success "? Reveal Answer" Correct: B

After the agent *acts* (calls a tool, runs code, queries a database), it *observes* the result ? the tool output is added back to the message history. This closes the loop: the LLM now has new information to reason over in the next step. The loop continues until the agent decides it has enough to answer.

question "Q3 (Multiple Choice): Which of the following is NOT one of the four core properties of an AI agent?"

- A) Perception - B) Compilation - C) Memory - D) Action

??? success "? Reveal Answer" Correct: B ? Compilation is not a core agent property

The four core properties are Perception (receives inputs), Memory (retains context), Reasoning (uses LLM to decide next step), and Action (calls tools/APIs/code). Compilation is a programming language concept, not part of the agent architecture.

Summary

An AI agent is an LLM-powered system that perceives, remembers, reasons, and acts to achieve a goal. It differs from traditional software because the logic is not hardcoded ? the LLM decides at runtime. This flexibility is powerful, but requires careful design of instructions and tools.

Next Steps

? Lab 002: AI Agent Landscape 2025 ? Compare all the tools and platforms available today.

Supporting kit

The site build generates stable PDF and ZIP downloads. The ZIP contains the lab guide and every supporting file declared for this lab.

[Download generated PDF](#) [Download generated ZIP](#)

Supporting file: CHEATSHEET.md

Supporting files: Lab 001: What are AI Agents?

Start with the documented local or free path. Record assumptions and validation evidence.

Keep hosted services optional unless the lab marks them required.

Supporting file: README.md

Lab 001: What are AI Agents?

Use these files with the lab guide and generated kit.