

Agentic AI for Multi-Agent Systems (DRAFT)

Mark A. Austin

University of Maryland

austin@umd.edu

ENCE 688P, Spring Semester 2026

October 12, 2026

Overview


- 1 Introduction to Agentic AI
 - Definition, Autonomy, Architecture and Design ...
- 2 Advanced Design Patterns
- 3 Agent Workflow Design Patterns
 - Single-Shot, Chaining, Routing, Orchestrator, Evaluator, Tools, Parallel, Autonomous
- 4 Travel Planning Application
 - System Architecture, Weather Analysis Agent, Itinerary Planner Agent, Travel Planning System Agent, Data Sources
- 5 Appendix A: Cheatsheets
- 6 Appendix B: Software Considerations

A set of small navigation icons typically found in Beamer presentations, including symbols for back, forward, search, and other slide controls.

Agentic AI

Agentic AI systems can act independently, make decisions, and complete multi-step tasks with little/no human guidance.


High-Level Agent Pattern:



Agentic AI

Agentic AI


AI becomes agentic when it can **decompose complex goals**; **adapt** its **plan**; **learn** from experience and feedback; **coordinate** with the **actions** of **others**; **predict** and **trade-off**.



Non-Agentive AI vs Agentive AI

Characteristic	Agentic AI	AI Agent
Operation	Operates independently	Performs specific tasks
Behavior	Intelligent collaborator	Task-focused assistant
Scope	Manages complex workflows	Automates repetitive tasks
Evolution	Updates behavior through experience	Doesn't evolve beyond responsibilities

Non-Agentive AI vs Agentive AI



Characteristics of Non-Agentive AI

Reactive and Stateless Workflow:

- Processes single-pass (input $\rightarrow$ output) without retaining active conversational or execution state across complex loops.

Tool Usage:

- Rare/none (pure text generation).


No Autonomous Planning:

- Lacks capability to break down (decompose) high-level goals into multi-step execution plans or perform self-correction.


Human-in-the-Loop Required:

- Requires explicit human prompting and constant oversight for every decision and subsequent action.


Non-Agentic AI vs Agentic AI




Agentic AI vs Generative AI



Agentic AI Organization of Concepts



Key Traits of Agentic AI



Key Traits of Agentic AI

Operates Independently:

- The system **works on its own** without needing human guidance at every step. It holds **internal state**, evaluates the options in front of it, and uses its own internal logic to pick an action.

Goals and Planning:

- Starting with a big objective (requirements), it uses some form of planning, search, task decomposition, learned policies – to break it into smaller steps and how to finish them.
- Sophisticated **strategies of planning** can work forward and backward from a target state.

Key Traits of Agentic AI

Tool Use and Environment Interaction:

- Agentic AI systems can call external systems (e.g., web databases) and APIs in a structured validated way, and then act on what comes back.
- Agents treats **tool calls** as part of controlled interaction with the world, not just a suggestion box.

Learn from Experience and Feedback:

- **Systems dynamically learn** better scenario selection over time by observing outcomes and previously taken actions.

Key Traits of Agentic AI

Predict and Trade-off:

- Models evaluate the effects of their own actions, weighing risk vs reward and priorities with respect to improving overall performance.

Adaptability:

- It learns from results and changes its tactics when new information or unforeseen factors force it to, thereby making it adaptable.

Evaluation and Safety Layers:

- Support for monitoring, guardrails and validation are built in (not just appendages).

Agentic AI Application Stacks

Trading Stack:

- Watch market conditions, weigh risks against hard constraints, execute within strict boundaries.

Robotics Stack:

- Plan a path, react to sensor input, replan when the world conditions change.

Agentic AI Benefits and Limitations

Benefits of Agentic AI:

- Handles complexity and changes gracefully.
- Can recover from unexpected outcomes through self-correction.
- Reduced need for constant human management.

Limitations of Agentic AI:

- Decision making can be hard to interpret.
- Multi-agent orchestration can cause conflicts.
- Balancing autonomy with (controlled) oversight is delicate.

Autonomy

Autonomy Spectrum

AI Agents:

- Task-driven, low-backbone executors that work within a certain scope based on pre-determined rules or shallow learning.

Agentic AI:

- Systems have the capability to adjust mission plans, coordinate with team members, and together manage flexible workflows.

Autonomous AI:

- Decision making over long horizons, executes plans, and seeks to meet goals with limited or no human integration (or involvement).

Autonomy Spectrum

AI Agents vs. Agentic AI vs. Autonomous AI

AI AGENTS




Task-Specific, Human-Guided.
Follows Instructions.

AGENTIC AI



Goal-Oriented, Plans Steps.
Seeks Information.


AUTONOMOUS AI



Self-Governing, Adapts & Evolves.
Operates Independently.

Levels of Autonomy

From Workflows to Agents



Levels of Autonomy

Level 1: Human-Assisted

- AI support humans to automate certain tasks, but is not able to make decisions by itself, and as such, requires human control and input.

Level 2: Human-Supervised Automation

- System is capable of some independent operations; however, humans are on hand to monitor it and can override operations if needed.

Level 3: Conditional Automation

- AI can operate independently for long stretches but might ask a person to intervene or take charge in tricky or uncertain situations.

Levels of Autonomy

Level 4: Fully Self-Driving in Limited Scenarios

- System can operate all functions unaided, but only under very specific circumstances or limited environments as needed.

Level 5: Completely Self-Sufficient in Any Domain or Task

- AI is completely self-sufficient and can adjust itself without intervention.

Source: testigor.com.

Agentic AI

Architecture and Design

Agentic AI Architecture and Design

Modular Design:

Planning and Meta-Reasoning:


Safety an Guardrails:

Interaction and Feedback:


Trade-Offs, Risk and Governance

Agentic Rag vs Rag

RAG (Retrieval-Augmented Generation)



Agentic RAG (AI-Agent Powered RAG)



Advanced Design Patterns

Advanced Design Patterns

Objectives:

- ...
- ...
- ...
- ...
- ...

Advanced Design Patterns

Memory Management

Memory management covers ...

Advanced Design Patterns

Learning and Adaptation

Learning and adaptation covers ...

Advanced Design Patterns

Model Control Protocol (MCP)

Model control protocol covers ...

Advanced Design Patterns

Goal Setting and Monitoring

Goal setting and monitoring (objective tracking) covers ...

Agent Workflow


Design Patterns

Single-Shot Workflow

Single-Shot Workflow

A [single-shot workflow](#) has one prompt and one response. It's simple and direct.

Workflow:




Chaining Workflow

Chaining Workflow

A [chaining workflow](#) breaks down complex tasks into sequential steps.

Workflow:




Routing Workflow

Routing Workflow

The **routing workflow** intelligently directs requests to the right destination.

Workflow:




Orchestrator Workflow

Orchestrator Workflow

An **orchestrator workflow** coordinates multiple AI agent to solve complex tasks.

Workflow:




Evaluator Workflow

Evaluator Workflow

The **evaluator workflow** generates, evaluates, and refines for higher quality output.

Workflow:




Tools Use Workflow

Tools Use Workflow

The **tools use** workflow lets AI accomplish tasks by loading and executing external tools.

Workflow:




Parallel Workflow

Parallel Workflow

Parallel workflow runs multiple tasks simultaneously for speed and scale.

Workflow:



Autonomous Workflow

Autonomous Workflow

An **autonomous workflow** sets goals, plans, acts, observes and adapts.

Workflow:



Choose the Right Pattern

Guidelines:

- Use simple workflows for quick and direct AI tasks.
- Combine multiple steps when tasks require deeper reasoning.
- Use tools and agent to automate complex real-world actions.
- Choose the autonomous systems pattern for adaptive self-managed workflows.

Travel Planning Application

Travel Planning System

Travel Planning System

System Architecture

Travel Planning System

WeatherAnalysisAgent: ...

Travel Planning System

ItineraryPlannerAgent: ...

Travel Planning System

TravelPlanningSystemAgent: ...

Travel Planning System

Weather Data

Hotel Accommodation Data


References

- Amidi A., and Amidi S., Super Study Guide: Transformers and Large Language Models (First Edition), ISBN 979-8-836693-31-2, 2024.
- Guilli A., Agentic Design Patterns (Hands-on Guide to Building Intelligent Systems), Springer, 2025.
- Lakshmanan V., and Hapke H., Generative AI Design Patterns: Solutions to Common Challenges when building GenAI Agents and Applications, O'Reilly, 2026.
- Raieli S., and Luculano G., Building AI Agents with LLMs, RAG and Knowledge Graphs, Packt Publishing, 2025.


Appendix A:

Cheatsheets

Agentic Design Patterns




Top Agentic Design Patterns



Agentic Workflow Design Patterns

8 AGENTIC AI WORKFLOW PATTERNS



Appendix B: Software Considerations

Building Agents with `create_agent()`

Benefit of `create_agent()`:

- Manages the ReAct loop

Thought $\longrightarrow$ Action $\longrightarrow$ Observation $\longrightarrow$ Repeat

- Handles message history automatically
- Implements iteration limits to prevent infinite loops
- Provides production-ready error handling
- Returns clean, structured responses

Create agent, get response ...

```
agent = create_agent( model, tools=[calculator], )  
response = agent.invoke({"messages": [HumanMessage(content=query)]})
```