AGENTIC AI CURRICULUM

Module 01: Python for AI Agents

Topics Covered

- **Pydantic (data validation):** Defines and enforces structured, validated data models for reliable AI agent inputs and outputs.
- **Asyncio (parallel tasks):** Enables concurrent execution of multiple AI agent tasks using asynchronous programming techniques.
- **API handling with FastAPI:** Builds high-performance APIs to integrate AI agents with external applications.

Resources

- Coursera and EC- Council

Module 02: LLM Reasoning & Memory

Topics Covered

- **Chain-of-Thought (CoT):** Guides AI agents to reason step-by-step while solving complex problems.
- **Tree-of-Thought:** Enables agents to explore multiple reasoning paths before selecting an optimal solution.
- **Context Windows vs. Long-term Memory:** Compares short-term context handling with persistent memory mechanisms for long-running agents.

Resources : DeepLearning.AI: *and ISACA*

Module 03: The Agentic Loop

Topics Covered

- **ReAct Pattern (Reason + Act):** Combines reasoning and action execution to enable goal-driven agent behavior.
- **Reflection:** Reviews prior agent responses to improve future decision-making.
- **Self-Correction loops:** Detects errors and autonomously refines agent behavior over time.

Resources: Coursera and LangChain Academy

Module 04: Tool Calling & MCP

Topics Covered

- **Function calling:** Enables AI agents to invoke external functions and tools dynamically during execution.
- **Model Context Protocol (MCP):** Standardizes secure access to tools, data sources, and execution environments.
- **Connecting to SQL / APIs:** Integrates AI agents with databases and third-party services for real-world automation.

Resources

- Google and IBM

Module 05: Agent Frameworks – I

Topics Covered

- **CrewAI (Role-based orchestration):** Designs role-based multi-agent systems where agents collaborate to complete tasks.
- **AutoGen (Conversational agents):** Builds conversational agents that communicate with each other to solve complex problems.

Resources

- DeepLearning.AI and Microsoft

Module 06: Agent Frameworks – II

Topics Covered

- **LangGraph (Stateful graphs):** Creates stateful, graph-based agent workflows that maintain execution history.
- **Building complex workflows with cycles:** Designs advanced agent workflows supporting loops, retries, and conditional execution.

Resources

- Coursera and LangChain

Module 07: Advanced RAG for Agents

Topics Covered

- **Agentic RAG (Self-RAG):** Builds agents that autonomously decide when and how to retrieve external knowledge.
- **Adaptive Retrieval:** Dynamically adjusts retrieval strategies based on query complexity.
- **Vector DBs (Pinecone / Milvus):** Stores, indexes, and retrieves semantic data using vector databases.

Resources

- DeepLearning.AI and IIIT Hyderabad

Module 08: Memory & State Management

Topics Covered

- **Short-term vs. Long-term memory:** Manages transient and persistent memory for continuous agent operations.
- **SQLite for persistence:** Stores agent state and memory reliably using lightweight databases.
- **Checkpointing:** Saves and restores agent execution states to ensure fault tolerance

Resources

- Coursera and IBM

Module 09: Agentic Security

Topics Covered

- **Prompt Injection:** Identifies and mitigates prompt-based attacks on AI agents.
- **Data Leakage:** Prevents exposure of sensitive information in AI agent interactions.
- **OWASP for LLMs:** Applies security best practices and addresses common vulnerabilities in large language models.
- **Sandboxing Agent Actions:** Restricts agent actions to controlled environments to prevent unintended system impact.

Resources

- ISACA and EC-Council

Module 10: AgentOps & Monitoring

Topics Covered

- **Tracing (LangSmith / Arize):** Tracks and debugs agent execution flows for performance and reliability.
- **Cost tracking:** Monitors and optimizes token usage and operational costs of AI agents.
- **Evaluation Frameworks (G-Eval):** Assesses agent output quality using structured evaluation techniques.

Resources

- Coursera and LangChain Academy

Module 11: Human-in-the-Loop (HITL)

Topics Covered

- **Designing approval gates:** Implements human approval steps for critical agent decisions.

- **Escalation protocols:** Establishes fallback mechanisms when agents encounter uncertainty or risk.
- **Feedback-based learning:** Improves agent performance through structured human feedback loops.

Resources Johns Hopkins University and Google Cloud

Module 12: Capstone Project

Topics Covered

- **Building a "Fully Autonomous Market Researcher":** Designs an end-to-end agent that autonomously gathers, analyzes, and summarizes market data.
- **Building an "Automated IT Support Suite":** Builds an agent system that autonomously handles IT support tasks and user queries.

Resources

- IIT Madras