Assignment 12
Agents: Tools, LangGraph, CrewAI, MCP

Part C — Building GenAI Applications

26 questions

★	**Must-know** — asked very frequently in GenAI / ML interviews (14 questions). Prepare these first.
[CODE]	**Hands-on task** — write and run the code (3 tasks).

1. What is an AI agent? How is it different from a chain? ★ · Asked: EPAM, PwC, Google
2. What are tools and how do they improve an LLM's capability? How does tool/function calling work? ★ · Asked: TCS, Google
3. How does a ReAct agent work (Thought → Action → Observation loop)? ★ · Asked: Google
4. Give an LLM the capability to use a calculator. [CODE]
5. Give an LLM the capability to use web search. [CODE]
6. What is a Python REPL tool? What are the security risks of letting an LLM execute code?
7. Create one custom tool and one custom agent. [CODE]
8. What is LangGraph? Why use it instead of a simple agent loop? ★ · Asked: Google
9. What is State in LangGraph? How are nodes and edges added? (add_node, add_edge) · Asked: EPAM, Google
10. How do you add a conditional edge? What are START and END, set_entry_point and set_finish_point?
11. How does LangGraph support memory/checkpointing and human-in-the-loop? · Asked: Google
12. What is CrewAI? How do you create agents, tasks and run a crew? ★
13. Single-agent vs multi-agent systems — when is multi-agent worth the complexity? · Asked: Google
14. What is MCP (Model Context Protocol) and what problem does it solve? ★ · Asked: TCS, EPAM, Infosys, Google
15. How do you evaluate and debug agents (traces, LangSmith/Langfuse, step-level metrics)? · Asked: Google
16. What are common agent failure modes (infinite loops, wrong tool choice, cost blow-ups) and how do you guard against them? · Asked: TCS, Google

17. What is Agentic AI, and how is it different from standard Generative AI or a RAG chatbot? ★ · Asked: TCS, Infosys, PwC
18. Who decides which tool to call? What if the model selects the wrong tool? Why shouldn't an agent have access to every tool? ★ · Asked: Infosys, Google
19. Your agent has 10 tools. The user asks a simple question, but the agent calls 5 different tools. How do you fix this? ★ · Asked: Infosys
20. Would you let an LLM perform a sensitive action automatically, like deleting a customer record or making a financial transaction? Should the LLM decide permissions? ★ · Asked: Infosys, Google
21. What is the difference between normal function calling and MCP? In MCP, what is the difference between tools and resources? ★ · Asked: Infosys, EPAM
22. How would you run three independent agents in parallel and make sure one agent's failure doesn't break the others? How do you add retries with backoff? ★ · Asked: EPAM
23. LangChain vs LangGraph – when would you use each? ★ · Asked: Google
24. What is the planner–executor architecture? How does it differ from a supervisor (orchestrator) agent? · Asked: Google
25. What is agent memory? What are the risks of long-term memory and stale information? · Asked: Google
26. In a Text-to-SQL agent, how do you prevent the LLM from executing destructive or hallucinated SQL queries? ★ · Asked: Deloitte