

Multi-Agent AI System using LangGraph

Agentic AI — autonomous systems in which specialized agents collaborate to complete complex goals — has emerged as the defining applied-AI paradigm of 2025–26, yet remains absent from final-year project kit catalogs. This project implements a multi-agent system orchestrated with LangGraph: three specialized agents (Researcher, Planner, Critic) operate as nodes in a stateful graph. The Researcher gathers evidence using web search APIs and retrieval-augmented generation over uploaded documents; the Planner decomposes missions into structured, dependency-aware task plans; the Critic evaluates drafts against explicit quality rubrics and routes failures back for revision, creating autonomous improvement loops bounded by iteration guards.

A Streamlit control room visualizes the live agent graph, inter-agent message flow, tool invocations, token consumption, and intermediate artifacts, with human-in-the-loop checkpoints allowing supervision mid-run. Two complete demo missions — startup idea validation and literature-review synthesis — ship with sample outputs and full execution traces. The system is evaluated on task-completion rubrics (coverage, grounding, revision efficiency), and the report discusses the honest limits of current agent evaluation methodology. The architecture cleanly separates orchestration (LangGraph), tooling (LangChain), memory (vector store), and presentation (control room), giving students four defensible viva layers. Built-to-order with source code, demo missions, full project report, presentation PPT, and viva Q&A; support, this is a first-mover final-year submission in the fastest-growing AI keyword of 2026 — distinct from single-assistant RAG systems.