


WALK AWAY WITH A PRODUCTION READY MULTI-AGENT SYSTEM

Join our GenAI course during which you'll learn all about Multi-Agent systems with LLM orchestration and Retrieval Augmented Generation. Come with a project in mind and leave with a production ready and scalable application.

CONTACT US

Email: hello@cellebris.com

Website: www.cellebris.com



INTRODUCTION TODAY



As AI continues to advance, models that can access and integrate external knowledge in real-time are becoming game-changers. Retrieval-Augmented Generation (RAG) is at the forefront of this shift, enabling applications to deliver more accurate, dynamic, and context-rich responses.

In this course, we'll break down the core concepts, practical strategies, and essential tools you need to harness RAG—no matter your level of expertise.

By the end, you'll be equipped with the insights and frameworks needed to move from theory to real-world implementation.

WHAT YOU'LL LEARN TOMORROW

Foundations of RAG: Understand what it is, why it matters, and the core principles behind retrieval-augmented models.

Key Components & Architecture: Learn about vector databases, embeddings, and how RAG pipelines integrate retrieval with generation.

Use Cases & Industry Examples: Explore how leading organizations leverage RAG for customer support, research, knowledge management, and more.

Implementation Steps & Tooling: Discover the best libraries, frameworks, and platform integrations to build your own RAG solution.

Best Practices & Troubleshooting: Get expert tips on model optimization, data curation, and error handling to ensure robust performance.



WHAT YOU'LL LEARN

THE MODULES

Module 1: Define Product Scope and Architecture

1.0 Introduction to Multi-Agent Systems: Offers a broad understanding of how multiple intelligent agents collaborate, setting the stage for an ecosystem where diverse AI components work together seamlessly.

1.1 Introduction to LLMs and Their Limitations: Highlights the boundaries of large language models, helping you understand why relying on a single model often falls short in complex applications.

1.2 Introduction to RAG: Bridging the Gap: Explains how Retrieval-Augmented Generation (RAG) fills the knowledge gaps of LLMs by pulling in external, contextually relevant data, ensuring more accurate and timely responses.

1.3 Defining Key Components of a RAG System: Breaks down the essential building blocks—retrievers, vector databases, embeddings, and LLMs—so you know precisely which pieces must fit together.

1.4 Defining Product Scope: Identifying Use Cases: Guides you to pinpoint where multi-agent RAG solutions can add real value, ensuring that you focus on meaningful, outcome-driven scenarios.

1.5 Designing the Architecture: A Modular Approach: Encourages a layered, modular design so that each component (retrieval, generation, orchestration) can be tweaked, swapped, or scaled independently.

Module 2: Setting Up Development Environment

2.1 Choosing a Platform and Tools: Helps you select the right frameworks, libraries, and infrastructure that can support robust RAG implementations and multi-agent workflows.

2.2 Configuring the Development Environment: Ensures your tools, dependencies, and environments are properly aligned, creating a stable foundation for iterative development and testing.

2.3 Data Acquisition and Preparation: Emphasizes obtaining clean, relevant, and well-structured data—critical for producing accurate retrievals and guiding your agents' decision-making.



WHAT YOU'LL LEARN

THE MODULES

Module 3: Setting Up Hosting and CI/CD

3.1 Introduction to Cloud Platforms for Deployment: Reviews the pros and cons of different cloud providers, ensuring you pick a scalable and secure environment for hosting agents.

3.2 Containerization with Docker: Introduces containerization for consistent deployment, making it easier to replicate development conditions in production.

3.3 Continuous Integration and Continuous Deployment (CI/CD): Automates testing and delivery, allowing faster updates to your multi-agent system without sacrificing reliability.

Module 4: Building a Backend Data and AI Platform

4.1 Designing the Data Model for RAG: Aligns your data structures with retrieval needs, ensuring fast, accurate lookups that inform your agents.

4.2 Building the Retrieval Component: Focuses on implementing a retrieval layer that quickly surfaces relevant documents, guiding the LLM in generating more accurate answers.

4.3 Integrating the LLM for Generation: Explains how to seamlessly merge your chosen language model with retrieved data so the agent's responses are both context-rich and authoritative.

4.4 Building APIs for Data Access and Interaction: Exposes the underlying capabilities through APIs, enabling other services and agents to easily request and use data.

4.5 Integrating Semantic Search and Summarization: Demonstrates how advanced techniques refine retrieval results, making them more understandable, context-aware, and valuable to downstream agents.



WHAT YOU'LL LEARN

THE MODULES

Module 5: Integrating Data and Defining Events

5.1 Understanding Event-Driven Architectures: Introduces the concept of events as triggers for agent actions, ensuring that your system responds dynamically to changing conditions.

5.2 Building Data Pipelines for Real-Time Updates: Sets up flows to keep your agents informed with up-to-the-minute data, preventing stale responses.

5.3 Triggering LLM Actions Based on Events: Connects event signals to LLM operations, allowing agents to automatically generate insights, summaries, or decisions when new information arrives.

Module 6: Building Front-End UX, API, and Event Generation

6.1 Designing an Intuitive User Interface: Emphasizes user-centric design so that human operators can interact smoothly with the multi-agent system's capabilities.

6.2 Building Interactive Components: Details how to create dynamic UI elements that let users guide the multi-agent process and view results in real time.

6.3 Developing a Robust API for Front-End Integration: Ensures front-end tools can seamlessly request data, trigger events, or call AI functions, promoting a flexible and responsive UX layer.

6.4 Generating Events from User Interactions: Treats user actions as signals, allowing the system's agents to respond proactively when someone requests information, uploads data, or alters configurations.



WHAT YOU'LL LEARN

THE MODULES

Module 7: Building an Adaptive UX

7.1 Understanding User Behavior and Preferences: Encourages analyzing user patterns to tailor the experience, helping agents provide more relevant information.

7.2 Personalizing the User Experience (Learning from the User in Real Time): Shows how the system can adapt based on user feedback, making results more personalized and efficient.

7.3 Implementing Adaptive UI Elements: Discusses dynamically altering the interface based on what users need, ensuring that the right tools and data are always at their fingertips.

7.4 Integrating a “Brain” for the Services You Use: Envisions a meta-agent that orchestrates other agents and services, acting as a central intelligence layer that continually refines user experiences.

Module 8: Managing a Production Environment

8.1 Deploying to Production: Guides you through launching your multi-agent RAG solution with minimal downtime and maximum readiness.

8.2 Monitoring and Logging: Stresses the importance of real-time performance data and logs to quickly identify issues and maintain quality.

8.3 Scaling for Increased Traffic: Prepares your system to handle growing user demands, ensuring that retrievals and generations remain fast and accurate.

8.4 Security Best Practices: Protects user data and the integrity of your agents, essential for maintaining trust and compliance.

8.5 Maintaining and Updating Your RAG Solution: Establishes routines for continuous improvement, ensuring your multi-agent application stays relevant, reliable, and aligned with evolving user needs.

PROJECTED NUMBERS

\$632B

Global AI Spending
(By 2028)

80%

Engineering Upskilling
Requirements (By 2027)

72%

Increased AI Adoption Across
Business Functions

\$15.7T

Global Economic Impact Of
AI (By 2030)

SOURCES

- **IDC** - Worldwide Spending on Artificial Intelligence Forecast to Reach \$632 Billion in 2028, According to a New IDC Spending Guide - 19th Aug. 2024
- **Gartner** - Gartner Says Generative AI will Require 80% of Engineering Workforce to Upskill Through 2027 - 3rd Oct. 2024
- **McKinsey** - The state of AI in early 2024: Gen AI adoption spikes and starts to generate value - 30th May. 2024
- **PwC** - PwC's Global Artificial Intelligence Study: Exploiting the AI Revolution -



CONCLUSION

THE FUTURE

Retrieval-Augmented Generation represents a powerful evolution in AI—one that combines the strengths of large language models with the richness of external data. By following the guidance in this PDF, you're now ready to start building your own RAG-driven applications.

To continue your journey, visit www.cellebris.com/course for more in-depth resources, community forums, workshops, and advanced courses that will help you take your RAG implementations even further.

ENTRY RECOMMENDATIONS

THE COURSE

Familiarity with software development

Shell (5%)

JavaScript (5%)

Python (90%)

Front and Back end Python experience:

class based software development

Bonus (but not essential):

Django & Celery

Containerisation using tools like Docker, Kubernetes

Data Integration (ingestion and processing of data)

API integrations