

Generative AI, RAG & Agentic AI Engineering

Build production-grade generative AI and agentic applications with RAG, tools, memory, evaluation, guardrails and deployment.



8 weeks

DURATION

106 hours

GUIDED LEARNING

2 projects

PORTFOLIO



ABOUT THE ACADEMY

Practical learning. Work you can demonstrate.

Brightnest AI Academy brings live instruction, guided practice and career preparation together for individuals and teams.

23

FOCUSED COURSES

6

FLAGSHIP PROGRAMMES

Noida

ACADEMY LOCATION

Our approach

1

Learn

Understand concepts through live instruction and current examples.

2

Build

Apply learning through labs, assignments and guided projects.

3

Explain

Document decisions, results, trade-offs and limitations.

4

Prepare

Connect completed work with CVs, interviews and workplace conversations.

PEOPLE BEHIND THE ACADEMY

Meet Ranjeet Kumar



ACADEMY ADVISER

Ranjeet Kumar

Innovation & Growth Leader

15+ years

Experience in data, artificial intelligence and machine learning leadership.

Areas of expertise

- Data science and machine learning solution leadership
- AI product strategy and scalable data platforms
- Technology team leadership and mentorship
- Practical AI and data adoption for business transformation

Education

Master of Information Systems, Anna University, Chennai
Bachelor of Computer Science Engineering



COURSE OVERVIEW

Build your next set of skills

Build production-grade LLM applications with RAG, AI agents, tool use, orchestration, memory, evaluation and reliable agentic workflows. Progress from LLM and Generative AI Engineering Foundations to Generative and Agentic AI Capstone through guided labs, assessed projects, and portfolio evidence.

Who this is for

Developers, ML engineers and technical professionals building LLM, RAG and AI-agent applications.

Before you begin

Learners should understand Python, APIs, and basic AI/ML concepts. Familiarity with embeddings or cloud services is helpful.

A convincing generative AI portfolio must show grounded retrieval, tool use, agent orchestration, evaluation, guardrails, and reliable deployment-not just prompting.



LEARNING OUTCOMES

What you will be able to do

01

Compare multiple models on a fixed task and create a model-selection scorecard.

02

Create robust structured-output prompts with validation and fallback behaviour.

03

Build a document ingestion pipeline and searchable knowledge index.

04

Improve a baseline RAG system using hybrid retrieval and reranking, then measure quality.

05

Build an assistant that safely calls external tools for search, calculations or business operations.

06

Deploy a RAG/agent application with tracing, cost tracking, caching and error recovery.

07

Build an enterprise-grade AI copilot or agentic workflow with measurable retrieval and task-success metrics.



CURRICULUM

Course curriculum

10 modules · 8 weeks · 106 hours of guided learning

01 LLM and Generative AI Engineering Foundations

6 hours

TOPICS

Transformer/LLM concepts · Tokens · Context windows · Inference parameters · Model selection · Hosted versus open models · Cost/latency/quality tradeoffs

PRACTICAL WORK

Compare multiple models on a fixed task and create a model-selection scorecard.

Tools: Python, Model APIs/open models, Notebooks

Assessment: Model-selection exercise

Portfolio evidence: LLM benchmark sheet

02 Prompt and Context Engineering

8 hours

TOPICS

System/user prompts · Few-shot examples · Structured outputs · Schema validation · Prompt templates · Context management · Prompt injection awareness

PRACTICAL WORK

Create robust structured-output prompts with validation and fallback behaviour.

Tools: Python, Pydantic/JSON Schema

Assessment: Prompt lab

Portfolio evidence: Prompt library with tests



CURRICULUM

Course curriculum

10 modules · 8 weeks · 106 hours of guided learning

03 Embeddings, Vector Search and Knowledge Ingestion

10 hours

TOPICS

Chunking · Metadata · Embeddings · Vector databases · Semantic search · Hybrid retrieval · Document parsing · Indexing · Freshness

PRACTICAL WORK

Build a document ingestion pipeline and searchable knowledge index.

Tools: FAISS/Chroma/Pinecone/Weaviate, Python

Assessment: Retrieval lab

Portfolio evidence: Knowledge ingestion pipeline

04 RAG Engineering and Retrieval Quality

12 hours

TOPICS

Query rewriting · Reranking · Hybrid search · Citations · Context packing · Multi-query retrieval · Evaluation datasets · Grounded answer design

PRACTICAL WORK

Improve a baseline RAG system using hybrid retrieval and reranking, then measure quality.

Tools: LangChain/LlamaIndex or SDKs, Vector DB, Reranker

Assessment: RAG evaluation report

Portfolio evidence: Measured RAG application



CURRICULUM

Course curriculum

10 modules · 8 weeks · 106 hours of guided learning

05 Tool Use, Function Calling and API Integration

10 hours

TOPICS

Tool schemas · Function calling · External APIs · Structured tool results · Error handling · Permissions · Idempotency · Human approval

PRACTICAL WORK

Build an assistant that safely calls external tools for search, calculations or business operations.

Tools: Python, REST APIs, Model SDKs

Assessment: Tool-calling lab

Portfolio evidence: Tool-using AI assistant

06 Agentic AI, Memory and MCP-Style Integrations

12 hours

TOPICS

Agent roles/goals · Planning · Memory types · State · Tool routing · Agent loops · Protocol-based tool integration · MCP concepts · A2A concepts

PRACTICAL WORK

Build a stateful agent connected to tools and a knowledge source with explicit control boundaries.

Tools: Agent framework/SDK, MCP-compatible tooling optional

Assessment: Agent architecture review

Portfolio evidence: Single-agent production prototype



CURRICULUM

Course curriculum

10 modules · 8 weeks · 106 hours of guided learning

07 Multi-Agent Workflows and Enterprise Orchestration

10 hours

TOPICS

Supervisor/worker patterns · Handoffs · Deterministic workflows versus autonomy · Queues · Long-running tasks · Human-in-the-loop · Failure recovery

PRACTICAL WORK

Design and implement a multi-step research-to-action workflow with checkpoints and approvals.

Tools: Agent framework, Workflow engine optional

Assessment: Workflow demo

Portfolio evidence: Multi-agent/workflow demo

08 Evaluation, Guardrails and Responsible GenAI

10 hours

TOPICS

Golden datasets · LLM-as-judge with controls · Hallucination/groundedness · Safety · Prompt injection · PII · Red teaming · Policy checks · Observability

PRACTICAL WORK

Create an automated evaluation suite covering quality, safety, groundedness and tool-use correctness.

Tools: Evaluation framework, Tracing/observability tools

Assessment: Evaluation suite review

Portfolio evidence: GenAI evaluation dashboard/report



CURRICULUM

Course curriculum

10 modules · 8 weeks · 106 hours of guided learning

09 Production Deployment, Observability and Cost

10 hours

TOPICS

API architecture · Streaming · Caching · Rate limits · Async jobs · Secrets · Tracing · Token/cost analytics · Reliability · Scaling · Incident readiness

PRACTICAL WORK

Deploy a RAG/agent application with tracing, cost tracking, caching and error recovery.

Tools: FastAPI, Docker, Cloud, Observability platform

Assessment: Production-readiness checklist

Portfolio evidence: Deployed AI application

10 Generative and Agentic AI Capstone

18 hours

TOPICS

Problem discovery · Architecture · Data ingestion · RAG · Tools · Agents · Evaluation · Security · Deployment · User experience · Business KPI

PRACTICAL WORK

Build an enterprise-grade AI copilot or agentic workflow with measurable retrieval and task-success metrics.

Tools: Full AI application stack

Assessment: Capstone demo and architecture defence

Portfolio evidence: Flagship portfolio GenAI/agentic application

APPLIED PORTFOLIO

Projects you will build

Two portfolio projects supported by guided module work.

PORTFOLIO PROJECT 1

Enterprise Knowledge Copilot

Build a RAG assistant over approved enterprise documents with citations, evaluation and guardrails.

DELIVERABLES

Ingestion pipeline · Retrieval benchmark · App · Traces · Evaluation report

PORTFOLIO PROJECT 2

Agentic Operations Assistant

Build a tool-using agent that retrieves knowledge, calls APIs and requests human approval for sensitive actions.

DELIVERABLES

Agent architecture · Tools · Tests · Guardrails · Observability · Demo

LEARNING EXPERIENCE

Learn, practise and present

- Live instructor-led sessions that connect concepts to real workplace decisions.
- Guided labs and workshops in every module.
- Assignments, checkpoints and practical feedback.
- Portfolio documentation, demonstrations and capstone review.
- Access to recordings and LMS resources according to the published batch policy.
- Career preparation based on completed work and target roles.

Tools you will work with

Python · Model APIs · Open models · Notebooks · Pydantic · JSON Schema · FAISS · Chroma · Pinecone · Weaviate · LangChain · LlamaIndex · SDKs · Vector DB · Reranker · REST APIs

CAREER PREPARATION

Build confidence for interviews

A structured process to connect your learning with your next career step.

01**Role clarity**

Identify roles that fit your background, interests and current skills.

02**Profile readiness**

Review your CV, LinkedIn, GitHub and project summaries.

03**Project proof**

Prepare to demonstrate completed work and explain your decisions.

04**Interview practice**

Practise technical questions, project walkthroughs and interviews.

05**Opportunity updates**

Receive relevant job openings and employer requirements, subject to availability and learner readiness.

Career support helps you prepare. Employment outcomes depend on experience, project quality, interview readiness and market conditions.

LEARNER JOURNEYS

Companies where our learners have been hired

Across the Brightnest AI Academy learner community



Learner employment history across the academy; not a placement guarantee for this course. Names and logos identify employers; they do not imply endorsement or a hiring partnership.



BEFORE YOU ENROL

Your questions, answered

Is the Generative & Agentic AI course suitable for beginners?

This is an advanced-level course. Learners should understand Python, APIs, and basic AI/ML concepts. Familiarity with embeddings or cloud services is helpful.

What will I build during the course?

You will complete guided labs in every module and build two portfolio projects: Enterprise Knowledge Copilot and Agentic Operations Assistant. Deliverables include working files or code, documentation, testing or evaluation evidence, and a final presentation.

Which tools and platforms are covered?

Key tools include Python, model APIs, open models, notebooks, Pydantic, JSON Schema, FAISS, and Chroma. Additional platforms are introduced in relevant modules through practical tasks, and the toolset may evolve as industry practice changes.

How long does the course take?

The course includes approximately 106 guided learning hours across 10 modules, delivered over 8 weeks.

Which career paths can this course support?

The curriculum supports the development of skills relevant to roles such as Generative AI Engineer, Agentic AI Engineer, LLM Application Engineer, and AI Engineer. Career outcomes depend on prior experience, project quality, interview readiness and market conditions; employment is not guaranteed.

BEFORE YOU ENROL

Your questions, continued

Will I receive mentor and career support?

The course includes live instruction, lab support, assignment feedback, project reviews and career preparation covering portfolio development, CV writing, LinkedIn profile improvement, and interview guidance.

FEES & ADMISSIONS

Plan your next step

COURSE FEE

INR 35,000 + GST

EMI and easy instalments available

Duration	8 weeks
Level	Advanced
Learning format	Live mentor-led learning
Batch starts	Expected Saturday / Sunday starts

Check the live Batches page for the next expected weekend. Confirm the exact date, class times and availability with admissions.

Book a free demo class >

+91 73030 11395

brightnestaiacademy.com/batches

A-56, Sector-64, Noida, Uttar Pradesh - 201301