SYLLABUS FOR BOARD OF STUDIES ADOPTION — generated from the live catalogue record

Generative AI Application Development — RAG and Fine-Tuning

S-10B · 5 days (configurable 1–5 days) · Informatics · AI · Years: 3rd, PG

CONTENTS

1. Overview 2
2. Hours and credit position 2
3. Course outcomes 2
4. Course content — modules 2
5. Assessment scheme (100 marks) 3
6. CO–PO mapping 3
7. Every participant receives 3
8. Fee 3
9. Designed by 3

Catalogue S-10B · version 2 · generated 2026-10-11

OVERVIEW

Vector databases, retrieval-augmented generation, evaluation and a deployed GenAI application.

Generative AI is reshaping software; students who can build retrieval-augmented and fine-tuned applications hold one of the most sought-after current skill sets.

HOURS AND CREDIT POSITION

Contact hours: 40 Guided project: 20 Self-study: 0 Notional total: 60

Credit claim: 2
60 notional hours ÷ 30 hours per credit (NSQF track, UGC Micro-credentials Guidelines §5.3.3) = 2 credits.

COURSE OUTCOMES

CO1. Explain how retrieval-augmented generation grounds an LLM's answers (K2)
CO2. Build a vector database and retrieval step over a document set (K3)
CO3. Implement a retrieval-augmented application with citations (K3)
CO4. Evaluate a generative-AI application for relevance and hallucination (K4)
CO5. Deploy a generative-AI application and justify its design choices (K5)

COURSE CONTENT — MODULES

Hours shown are the recommended format; module time scales proportionally when the engagement runs 1–5 days.

Module 1 — LLM foundations (5h recommended)
Theory: Capabilities, limits, the application landscape
Lab: Probe an LLM's behaviour and limits
Output: Behaviour log

Module 2 — Prompt and context engineering (5h recommended)
Theory: Patterns, structured output, guardrails
Lab: Engineer prompts for reliable output
Output: Prompt library

Module 3 — Embeddings and vector databases (5h recommended)
Theory: Chunking, indexing, similarity search
Lab: Index a corpus in a vector database
Output: Vector index

Module 4 — Retrieval-augmented generation (5h recommended)
Theory: Grounding answers, citations
Lab: Build a RAG pipeline with citations
Output: RAG pipeline

Module 5 — Evaluation (5h recommended)
Theory: Relevance, faithfulness, hallucination
Lab: Evaluate the RAG application
Output: Evaluation report

Module 6 — Fine-tuning (5h recommended)
Theory: When and how to fine-tune, dataset preparation
Lab: Prepare data and fine-tune a model
Output: Fine-tuned model

Module 7 — Orchestration (5h recommended)
Theory: Chaining, tools, agents
Lab: Orchestrate a multi-step flow
Output: Orchestrated flow

Module 8 — Capstone (5h recommended)
Theory: A deployed GenAI application with evaluation
Lab: Ship and evaluate the capstone
Output: Deployed GenAI application

ASSESSMENT SCHEME (100 MARKS)

Continuous lab exercises — 40 marks
Capstone project & demonstration — 40 marks
Quiz & viva voce — 20 marks

CO–PO MAPPING

CO1 !' PO1 (strength 2)
Retrieval grounding is foundational GenAI knowledge.

CO2 !' PO3 (strength 2)
Building retrieval is solution development.

CO3 !' PO3 (strength 2)
A cited RAG app is a developed solution.

CO4 !' PO4 (strength 2)
Evaluating for hallucination is investigation.

CO5 !' PO10 (strength 2)
Justifying design choices is technical communication.

EVERY PARTICIPANT RECEIVES

- Certificate of completion
- Course material
- LMS access
- Interview question bank
- Mock interview & viva practice
- Optional nasscom NSQF assessment (priced per candidate on your quotation)

FEE

¹1,500 per student, incl. GST
Excludes the optional nasscom assessment where offered. Final pricing on your college's quotation.

DESIGNED BY

Venkat Kurela — Founder & Director, Kurela Cognisive Private Limited