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

Generative AI & Prompt Engineering for Every Discipline

C-01 · 3 days (configurable 1–5 days) · Informatics · AI · Years: 1st, 2nd, 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)	2
6. CO–PO mapping	3
7. Every participant receives	3
8. Fee	3
9. Designed by	3

Catalogue C-01 · version 1 · generated 2026-08-24

OVERVIEW

Students build and deploy a working GenAI assistant using an LLM, retrieval and a simple interface. Every branch can take it, it needs no prerequisite beyond basic computing, and students leave with something working that they built themselves.

HOURS AND CREDIT POSITION

Contact hours: 24 Guided project: 6 Self-study: 0 Notional total: 30

Credit claim: 1

30 notional hours ÷ 30 hours per credit (NSQF track, UGC Micro-credentials Guidelines §5.3.3) = 1 credits.

COURSE OUTCOMES

- CO1. Explain how a large language model produces output, naming context, tokens and the causes of hallucination (K2)
- CO2. Construct prompts using role, context and constraint patterns to obtain reliable structured output (K3)
- CO3. Build a retrieval-augmented assistant grounded in a supplied document set (K3)
- CO4. Deploy a working AI assistant behind a simple web interface (K3)
- CO5. Evaluate an AI assistant's answers for accuracy, bias and appropriate use in academic work (K5)

COURSE CONTENT — MODULES

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

Module 1 — How large language models work (4h recommended)

Theory: Tokens, context windows, training vs inference, capabilities and limits, hallucination

Lab: Guided exploration of model behaviour on real prompts

Output: Annotated prompt-behaviour log

Module 2 — Prompt engineering in practice (4h recommended)

Theory: Role, context, constraints; few-shot patterns; structured output; iteration strategies

Lab: Prompt battles — teams improve a weak prompt through five iterations

Output: Personal prompt pattern library

Module 3 — Retrieval-augmented generation (4h recommended)

Theory: Embeddings, vector search, chunking, grounding answers in documents

Lab: Build a Q&A assistant over a supplied document set

Output: Working RAG notebook

Module 4 — Building an AI assistant (4h recommended)

Theory: API calls, system prompts, conversation state, guardrails

Lab: Assemble an assistant with a defined persona and scope

Output: Running assistant script

Module 5 — A simple interface (4h recommended)

Theory: Minimal web UI, input handling, streaming responses, error states

Lab: Wire the assistant to a browser interface

Output: Deployed local demo

Module 6 — Responsible use and capstone (4h recommended)

Theory: Bias, attribution, academic integrity, data privacy; capstone build time

Lab: Capstone: each team ships a discipline-specific assistant

Output: Capstone demo to the cohort

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)

Foundational understanding of the technology every later skill builds on.

CO2 !' PO5 (strength 2)

Prompting is the modern tool-usage skill this workshop exists to teach.

CO3 !' PO3 (strength 2)

Retrieval design is a solution-design exercise against a real corpus.

CO4 !' PO5 (strength 2)

Deployment turns the technique into a demonstrable artefact.

CO5 !' PO12 (strength 2)

Critical evaluation of AI output is the lifelong-learning habit.

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

12,000 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