

GENERATIVE AI AND LLMs

Code: CS33213CS	Course Type: Core	Semester: II
Evolution Scheme (TA-MSE-ESE): 20-30-100	Total Marks: 150	L-T-P (Credits): 2-0-0 (2)

Pre-Requisites: 1. Learners have a basic understanding of NLP.
2. Learners are familiar with basic deep learning concepts.

Course Outcomes: At the end of the course, the student will be able to

CO-1	Explain the fundamentals of generative AI, including transformer architecture, tokenization, embeddings, attention mechanisms, and generative models such as VAE and GAN.
CO-2	Analyze the architecture, training process, and applications of LLMs, including pre training, fine-tuning, instruction tuning, and PEFT techniques such as LoRA and QLoRA.
CO-3	Apply reinforcement learning concepts in the context of LLM alignment, including RLHF, reward modeling, PPO, and KL-divergence for improving model behavior.
CO-4	Develop intelligent AI applications using RAG, vector databases, semantic search, document chunking, and prompt engineering techniques such as zero-shot and few shot prompting.

Course Articulation Matrix:

PO	PEO-1	PEO-2	PEO-3	PEO-4	PEO-5	PO
CO 1	3	2	1	-	-	CO 1
CO 2	3	3	2	1	-	CO 2
CO 3	3	2	3	2	1	CO 3
CO 4	3	3	2	1	2	CO 4

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Unit-1: Generative AI

Introduction to Generative AI & LLMs, Generative AI project lifecycle, Text generation before transformers, Transformers architecture and its components: Self attention, Multi-head attention, Tokenization, embeddings, positional encodings, Decoder-based Vs encoder-decoder based language modelling, Variational Autoencoders (VAE) & Generative Adversarial Networks (GANs).

Unit-2: Large Language Models (LLMs)

Introduction to LLMs, LLM use cases and tasks, Popular LLM families: GPT/Llama/Gemini/Claude, Application of LLMs (GPT, BERT, etc.,) in NLP tasks, Pre-training large language models (GPT,

BERT, BART, and T5), Fine-tuned language models, Computational challenges of training LLMs, Instruction fine tuning, Parameter efficient fine-tuning(PEFT) techniques : LoRA, QLoRA.

Unit-3: Reinforcement Learning

Introduction to reinforcement learning, Reinforcement learning from human feedback (RLHF), RLHF: Reward model, RLHF: Fine-tuning with reinforcement learning, Proximal policy optimization (PPO), KL-divergence.

Unit-4: Retrieval Augmented generation(RAG) and Prompt Engineering

Introduction to RAG, Agentic AI, Introduction to LangChain, Vector databases and indexing, Semantic Search, Document chunking, Introduction to Prompt Engineering, Types of Prompting: Zero-shot and few-shot prompting.

Learning Resources:

Text Books:

1. Pam baker, Generative AI For Dummies, Wiley, 2025
2. Tanmoy Chakraborty, Introduction to Large Language Models Generative AI for Text, Wiley, 2024
3. K.SUBRAMANIAN, Generative AI: Theory Practice and Ethics , Notion Press, 2025