

Generative AI Course Content

◆ **Module 1: Introduction to Generative AI**

- What is Generative AI?
- Types: Text, Image, Audio, Video, Code generation
- Evolution: From GANs to Diffusion to Transformers
- Key real-world applications (ChatGPT, Midjourney, Runway, GitHub Copilot)

◆ **Module 2: Foundation Models & Architectures**

- Transformers, LLMs (BERT, GPT, LLaMA, Claude, etc.)
- Diffusion Models vs GANs (Stable Diffusion, StyleGAN)
- Contrastive Models (CLIP, DALL·E)
- Open vs Closed Source Models

◆ **Module 3: Prompt Engineering Mastery**

- Zero-shot, few-shot, chain-of-thought prompting
- Role-based prompting
- Tool usage: OpenAI Playground, LangChain PromptTemplates
- Real-world use cases: data extraction, summarization, agents

◆ **Module 4: Working with Text Generation**

- Using APIs: GPT-4, Claude, Gemini
- Building apps with OpenAI SDK, HuggingFace Transformers
- Chatbots, Summarizers, Email & Content Writers
- Project: Custom GPT chatbot for a niche domain

◆ **Module 5: Image Generation & Manipulation**

- Stable Diffusion, Midjourney, DALL·E
- Prompt tuning & style conditioning
- ControlNet, Inpainting, Outpainting
- Project: Build your own AI image studio (text-to-image web app)

◆ **Module 6: Multimodal Generative AI**

- Text-to-video (RunwayML, Sora, Pika Labs)
- Text-to-audio (ElevenLabs, Bark, MusicGen)
- Image Captioning + Visual Question Answering
- Project: AI Video Generator from a Script

◆ **Module 7: Fine-Tuning & RAG Systems**

- Fine-tuning LLMs (LoRA, PEFT)
- Building RAG pipelines using LangChain + FAISS/Weaviate
- Use case: AI assistant on your company or personal data
- Project: Document Q&A system using private PDFs

◆ **Module 8: Generative AI in Production**

- LLMOps overview
- Rate-limiting, caching, and fallback strategies
- Streamlit/Gradio deployment + LangChain Agents
- Project: Build & deploy an AI CoPilot web app

◆ **Module 9: Ethics, Risks & Guardrails**

- Bias, hallucination, safety challenges
- Open-source alignment techniques (ReAct, Constitutional AI)
- Tools: GuardrailsAI, Rebuff, Trulens

◆ **Module 10: Capstone Project + Career Track**

- Capstone: End-to-End Product using Generative AI
- Personal Portfolio creation
- Resume review + job guidance for AI roles
- Bonus: Building a freelancing or SaaS business with Generative AI

Tools Covered

- APIs: OpenAI, Claude, Gemini, Stability AI, HuggingFace
- Frameworks: LangChain, LangGraph, Transformers, Diffusers
- UI: Streamlit, Gradio, Next.js (optional)
- Infra: Weaviate, Pinecone, ChromaDB, n8n
- TTS/STT: ElevenLabs, Whisper
- Versioning: GitHub, Weights & Biases