

GENERATIVE AI

COURSE CURRICULUM | BEGINNER TO ADVANCED

LEARN • BUILD • INNOVATE

- Generative AI (GenAI) Course Curriculum

□ Beginner Level (35–40 Hours)

Objective: Build solid foundations in AI, understand GenAI basics, and learn prompt engineering & hands-on APIs.

◆ Module 1: Foundations of AI & GenAI (3 hrs)

- AI vs ML vs DL vs GenAI
- History & breakthroughs in Generative AI
- Key models: GPT, Stable Diffusion, MidJourney
- Real-world applications & industry use cases

◆ Module 2: Introduction to LLMs (5 hrs)

- How LLMs work: tokens, embeddings, transformers
- Pretraining vs fine-tuning
- Popular open-source LLMs (LLaMA, Mistral, Falcon)
- Hosted APIs (OpenAI, Anthropic, Cohere)

◆ Module 3: Prompt Engineering (6 hrs)

- Zero-shot, few-shot prompting
- Instruction-tuning
- Role-based prompts & system instructions
- Structured prompting with templates
- Prompt chaining

◆ Module 4: GenAI Tools & APIs (6 hrs)

- OpenAI API basics (chat completions, embeddings)
- Hugging Face Spaces & Models
- Google Gemini, Anthropic Claude
- VS Code + GitHub Copilot

◆ Module 5: Hands-on with Text, Image, Audio (8 hrs)

- Text: Q&A, summarization, rewriting
- Images: DALL·E, Stable Diffusion, MidJourney basics
- Audio: TTS (text-to-speech), STT (speech-to-text)
- Safety & ethics in GenAI

✓ Beginner Projects

- Build a personal AI note summarizer.
- AI-powered email writing assistant.

- Generate an infographic/poster with Stable Diffusion.
- Create a voice-based chatbot with STT + TTS.

Intermediate Level (45–50 Hours)

Objective: Learn fine-tuning, RAG, multimodal models, and integrating GenAI into apps.

◆ Module 6: Model Fine-Tuning (8 hrs)

- Fine-tuning vs adapter methods (LoRA, PEFT)
- Hugging Face Trainer basics
- Training domain-specific chatbots
- Evaluating model performance

◆ Module 7: Retrieval-Augmented Generation (RAG) (7 hrs)

- Why RAG is needed
- Embeddings for semantic search
- Vector databases: Pinecone, FAISS, Weaviate
- Implementing a RAG-based chatbot

◆ Module 8: Multi-Modal Generative AI (6 hrs)

- Text-to-image (Stable Diffusion, DALL·E)
- Text-to-audio (ElevenLabs, Bark)
- Text-to-video (Pika Labs, RunwayML)
- Vision+Language (CLIP, BLIP-2)

◆ Module 9: GenAI in Software Development (6 hrs)

- Copilot for code generation
- AI-powered debugging
- Test case generation
- AI-assisted documentation

◆ Module 10: Application Development with LangChain (8 hrs)

- LangChain basics (chains, memory, tools, agents)
- Connecting APIs & databases
- Advanced prompting with LangChain
- Deploying apps with Streamlit/Gradio

✓ Intermediate Projects

5. Fine-tune GPT on company FAQs.
6. Build a RAG-based legal/educational Q&A chatbot.
7. AI-powered SQL query generator.
8. Build a multimodal app: text → image + audio + explanation.
9. Code assistant that generates unit tests with Copilot + LangChain.

◆ Module 11: Advanced Transformer Architectures (8 hrs)

- Deep dive into attention mechanisms
- Scaling laws, efficiency (quantization, pruning, distillation)
- LoRA & PEFT methods

◆ Module 12: Scaling & Deployment (7 hrs)

- Serving LLMs with Docker + FastAPI
- Cloud deployment (AWS, GCP, Azure)
- Cost optimization & monitoring

◆ Module 13: AI Safety & Responsible AI (6 hrs)

- Bias in LLMs
- Prompt injection & jailbreaks
- Guardrails (Azure Safety, Llama Guard)
- Responsible AI policies

◆ Module 14: Autonomous AI & Agentic Systems (8 hrs)

- What is Agentic AI (self-directed AI agents)
- LangChain Agents (tool-use, reasoning, planning)
- AutoGPT, BabyAGI, CrewAI
- Multi-agent collaboration (team of agents with roles)
- Use cases: research assistants, workflow automation, orchestration

◆ Module 15: Anthropic Model Context Protocol (MCP) (8 hrs)

- Introduction to MCP and context-sharing
- How MCP allows AI models/tools to communicate
- Structured context passing across agents
- MCP + Agentic AI integration
- Enterprise-grade use cases: knowledge workers, co-pilots, secure tool integration

◆ Module 16: Enterprise GenAI Applications (8 hrs)

- Healthcare: diagnostic copilots
- Finance: fraud/risk copilots
- Retail: AI-powered assistants
- Media: generative pipelines (text → video → marketing)

✦ Advanced + Agentic Projects

10. Deploy a fine-tuned model with Docker + FastAPI on AWS.
11. Build a multi-agent research assistant (Planner + Researcher + Writer).
12. AI-powered workflow automation system with LangChain Agents.
13. RAG + Guardrails chatbot with secure context injection.
14. Real-time voice-to-voice translation app with multimodal input.
15. AI-powered contract compliance checker (legal domain).
16. Anthropic MCP-based Knowledge Assistant (agents sharing context between finance & HR data).
17. Agentic AI + MCP Integration: Build an enterprise helpdesk with multi-agent orchestration (e.g., one agent handles FAQs, one checks HR policies, one generates email replies).

□ Capstone Projects (Choose 2–3)

- Healthcare: AI-powered EHR copilot with Agentic AI + MCP.
- Finance: Risk analysis system with MCP-based secure data exchange.
- Education: Autonomous tutor agent with memory & context-sharing.
- Entertainment: Text-to-video content generator using Agentic AI pipelines.
- Retail: Virtual AI shopping assistant that uses multi-agent + MCP for recommendations + order processing.

✦ BUILD THE FUTURE WITH GENERATIVE AI ✦