

Generative AI & Agentic AI — 6-Month Project-Based Curriculum

Code the Future

Format: Live, instructor-led | **Duration:** 6 Months | **Mode:** Online, hands-on, project-driven

Outcome: Job-ready GenAI/Agentic AI Engineer with 6+ deployed projects and 1 capstone.

MONTH 1 — Foundations: Python, ML/DL & NLP for AI

Goal: Build the programming, math, and NLP foundation required for generative and agentic systems.

Week	Topics
Week 1	Python for AI — advanced Python, APIs, async programming, virtual environments
Week 2	ML/DL refresher — neural networks, backpropagation, PyTorch/TensorFlow basics
Week 3	NLP fundamentals — tokenization, embeddings, word2vec, attention mechanism
Week 4	Transformer architecture deep-dive — encoder/decoder, self-attention, positional encoding

Project: Build and train a mini transformer/text classifier from scratch

Tools: Python, PyTorch, Hugging Face Transformers

MONTH 2 — Generative AI Fundamentals & Prompt Engineering

Goal: Understand how LLMs work and learn to control their behavior effectively.

Week	Topics
Week 5	LLM landscape — GPT, Claude, Gemini, Llama, open vs closed models, context windows
Week 6	Prompt engineering — zero-shot, few-shot, chain-of-thought, system prompts
Week 7	Structured outputs, function calling, JSON mode, output parsing
Week 8	Working with APIs — Anthropic/OpenAI SDKs, token management, cost & latency optimization

Project: AI-powered content generation tool (e.g., resume builder, blog writer) using an LLM API

Tools: Anthropic API / OpenAI API, Python, Streamlit

MONTH 3 — RAG (Retrieval-Augmented Generation) & Vector Databases

Goal: Ground LLMs in custom data to build accurate, domain-specific applications.

Week	Topics
Week 9	Embeddings & semantic search, vector databases (Pinecone, Chroma, FAISS)
Week 10	Building RAG pipelines — chunking strategies, retrieval, re-ranking

Week 11	LangChain & LlamaIndex — chains, retrievers, memory, document loaders
Week 12	Advanced RAG — hybrid search, evaluation of retrieval quality, reducing hallucination

Project: RAG-based chatbot over a custom knowledge base (e.g., company docs, PDFs, or a support FAQ)

Tools: LangChain/LlamaIndex, ChromaDB/Pinecone, Python

MONTH 4 — Fine-Tuning & Model Customization

Goal: Learn to adapt and optimize open-source models for specific tasks.

Week	Topics
Week 13	Fine-tuning fundamentals — full fine-tuning vs parameter-efficient methods
Week 14	LoRA & QLoRA — fine-tuning open-source LLMs (Llama, Mistral) on custom datasets
Week 15	Model evaluation — benchmarks, human-in-the-loop evaluation, RLHF basics
Week 16	Model optimization — quantization, distillation, inference optimization

Project: Fine-tuned open-source LLM for a domain-specific task (e.g., customer support or legal document summarization)

Tools: Hugging Face, PEFT/LoRA, Google Colab/local GPU

MONTH 5 — Agentic AI: Building Autonomous Agents

Goal: Move from single-turn LLM calls to autonomous, tool-using, multi-step agents.

Week	Topics
Week 17	Agent fundamentals — reasoning + acting (ReAct), planning, tool use
Week 18	Building agents with LangGraph / CrewAI / AutoGen — single-agent workflows
Week 19	Multi-agent systems — agent collaboration, role-based agents, orchestration
Week 20	Agent memory, guardrails, error handling, and human-in-the-loop control

Project: Multi-agent system that automates a real workflow (e.g., research assistant, data-analysis agent, or customer-support agent team)

Tools: LangGraph, CrewAI/AutoGen, Python, external APIs/tools

MONTH 6 — Deployment, MLOps for LLMs, Ethics & Capstone

Goal: Ship production-ready GenAI/Agentic systems and prepare for industry roles.

Week	Topics
Week 21	Deploying LLM apps — FastAPI/Streamlit, containerization (Docker), cloud hosting basics
Week 22	LLMOps — monitoring, logging, cost tracking, prompt versioning, CI/CD for AI apps
Week 23	AI safety & ethics — hallucination mitigation, bias, responsible AI, data privacy
Week 24	Capstone project development, documentation, portfolio & mock interviews

Project: End-to-end deployed Agentic AI application combining RAG + tool use + multi-agent orchestration, presented to a panel

Tools: Docker, FastAPI, cloud deployment (AWS/GCP/Azure basics), Git

Skills Covered End-to-End

Python | PyTorch/Hugging Face | Prompt Engineering | LLM APIs (Claude, GPT, Gemini) | RAG | Vector Databases | LangChain/LlamaIndex | Fine-Tuning (LoRA/QLoRA) | Agentic Frameworks (LangGraph, CrewAI, AutoGen) | Multi-Agent Systems | LLM Ops | Deployment (Docker, FastAPI, Cloud) | AI Ethics & Safety

Deliverables by Program End

- 6 hands-on projects (portfolio-ready, GitHub hosted)
- 1 industry-grade capstone: deployed Agentic AI application
- Resume, LinkedIn, and GitHub profile positioned for GenAI/Agentic AI roles
- Mock interview experience and placement support

Suggested Weekly Structure

- 2 live instructor-led sessions per week
- 1 doubt-clearing/mentorship session per week
- Weekly assignment + monthly project submission
- Peer code reviews and collaborative build sessions

Note: This curriculum assumes learners have basic Python and ML familiarity (or have completed a foundational Data Science track). Pacing can be adjusted for absolute beginners vs. working professionals, and depth in fine-tuning (Month 4) can be trimmed for a more API/agent-first track.