Generative AI Master Syllabus

WEEK 1 — Python & AI Foundations

Learn Python from scratch and understand what AI, ML and Generative AI are before building anything.

Day 01 Course orientation, what AI, Machine Learning, Deep Learning and Generative AI mean, program roadmap and career roles.
Day 02 Set up Python, VS Code/Jupyter and virtual environments; first Python program, variables and data types.
Day 03 Operators, conditional statements and loops in Python; daily coding practice.
Day 04 Functions, lists, tuples, sets and dictionaries; reusable code and modules.
Day 05 File handling, exception handling, JSON and working with packages using pip.
Day 06 NumPy for AI: arrays, indexing, vectorised operations and basic maths.
Day 07 Pandas basics: DataFrames, loading CSV data, cleaning and exploring; weekly assessment.

WEEK 2 — ML, Deep Learning & NLP Basics

Understand how machines learn, how neural networks work and how language is processed.

Day 08 Supervised, unsupervised and reinforcement learning with real-life examples.
Day 09 Machine learning workflow: features, training, testing, overfitting and model evaluation.
Day 10 Neural networks: neurons, layers, activation functions and loss functions.
Day 11 Backpropagation and gradient descent; training a simple neural network.
Day 12 Natural Language Processing basics: text cleaning, tokens, word representations and use cases.
Day 13 Transformer architecture fundamentals: why transformers changed AI.
Day 14 Revision, hands-on mini lab and weekly assessment on AI, ML and DL fundamentals.

WEEK 3 — Generative AI Fundamentals

Understand how Generative AI works and the model types behind text, image and code generation.

Day 15 What Generative AI is and how it works; real-world applications across industries.
Day 16 Discriminative AI vs Generative AI; where each one is used.
Day 17 Model types: autoregressive models (GPT), VAEs, GANs and diffusion models.
Day 18 Tokens, tokenization, embeddings and vector representations (hands-on).
Day 19 Training vs inference, model parameters and how model size affects results.
Day 20 Foundation models and multimodal AI; overview of ChatGPT, Gemini, Claude and open models.
Day 21 Hands-on lab with popular GenAI tools and weekly assessment.

WEEK 4 — Large Language Models & Prompt Engineering

Know how LLMs work and write prompts that give reliable, structured outputs.

Day 22 LLM architecture: transformer concepts, attention and self-attention.
Day 23 Pre-training, instruction tuning and RLHF: how chat models are made.
Day 24 Closed models (GPT-style) vs open-weight models; reasoning models and extended thinking.
Day 25 Context windows, temperature, top-p and output control; hallucinations and limitations; model selection and cost (tokens, latency, pricing).
Day 26 Prompt engineering: zero-shot and few-shot prompting, system, user and assistant messages, role-based and instruction-based prompts.
Day 27 Chain-of-thought, task decomposition, structured outputs with JSON and schemas, reusable prompt templates.
Day 28 Prompt evaluation and optimization, guardrails and prompt-injection awareness; weekly assessment.

Greens Technologies

WEEK 5 — LLM APIs & Multimodal AI

Build real applications on LLM APIs and work with images, speech and audio.

Day 29 OpenAI, Anthropic and Gemini APIs: API keys, environment variables, sending prompts and processing responses.
Day 30 Open-weight models with Ollama and vLLM; streaming responses.
Day 31 Function calling, tool calling and structured output generation.
Day 32 Error handling, retries and rate limits; prompt caching and batch APIs for cost reduction.
Day 33 Multimodal Generative AI: text-to-image generation, image understanding and vision-language models.
Day 34 Speech-to-text, text-to-speech, audio and video generation fundamentals; multimodal prompting; document parsing, OCR and visual Q&A.
Day 35 Project 1 (Beginner): AI Chatbot with prompt templates, conversation history and streaming responses.

WEEK 6 — Retrieval-Augmented Generation (RAG)

Give LLMs your own data and build grounded, accurate answers with sources.

Day 36 What RAG is and why LLMs need external knowledge; RAG architecture overview.
Day 37 Document loading and text extraction from PDF, HTML and tables.
Day 38 Chunking strategies and metadata; embedding models.
Day 39 Vector databases (FAISS, Chroma, pgvector) and similarity search.
Day 40 Semantic and hybrid search (BM25 + vectors); retrieval quality and reranking.
Day 41 Grounded answers with source citations; RAG evaluation and optimization.
Day 42 Hands-on RAG pipeline lab and weekly assessment.

WEEK 7 — Fine-Tuning & Final Project

Customise models, choose the right approach and finish your Generative AI project.

Day 43 Prompting vs RAG vs fine-tuning: when to use which.
Day 44 Preparing datasets and supervised fine-tuning fundamentals.
Day 45 Preference tuning concepts: RLHF and DPO.
Day 46 LoRA and QLoRA; fine-tuning open-weight models with Hugging Face and PEFT.
Day 47 Distillation and smaller models; model evaluation, overfitting and fine-tuning in production.
Day 48 Project 2 (Intermediate): PDF Q&A Assistant with retrieval, grounded answers and citations.
Day 49 Project presentation, final assessment, portfolio review and interview preparation.

Greens Technologies