3-MONTH AI MASTERY CURRICULUM

Personal 1-on-1 Program

Python • Machine Learning • Deep Learning • Generative AI

From Zero to Job-Ready in 90 Days

90	12	4	1-on-1
Days	Modules	Phases	Sessions/Week

Samjho Karo Badho | AI, Sabke Liye
genaiwallah.com

TABLE OF CONTENTS

PROGRAM OVERVIEW

This 3-month intensive 1-on-1 curriculum takes you from Python basics to building real Generative AI applications. Every week is structured with theory, hands-on coding, and a mini-project. By the end of 90 days, you will have a portfolio of real AI projects and the skills to compete for AI roles — regardless of your college background.

Phase	Duration	Focus	Outcome
1	Weeks 1–2	Python Foundations	Write clean Python code
2	Weeks 3–4	Data & ML Basics	Understand & use ML models
3	Weeks 5–6	ML Deep Dive	Build & evaluate ML pipelines
4	Weeks 7–8	Deep Learning	Build neural networks
5	Weeks 9–10	Generative AI	Use & fine-tune LLMs
6	Weeks 11–12	Projects & Career	Portfolio + job-ready

Session Structure (Each Week)

Session	Day	Duration	What Happens
Session 1	Tuesday/Wednesday	90 min	Theory + Concept explanation + Live coding
Session 2	Friday/Saturday	90 min	Practice problems + Mini-project work + Q&A
Self Study	Daily	1–2 hrs	Assignments, reading, practice exercises

PHASE 1 — PYTHON FOUNDATIONS

Month 1 | Weeks 1–2 | 4 Sessions

No prior coding experience needed. We start from absolute zero and build a strong Python foundation that everything else depends on. By end of Week 2 you will be writing real programs confidently.

Python	Jupyter Notebook	VS Code	GitHub Basics

MODULE 1 — Python Basics

Week 1 | 2 Sessions

Session 1 — Setup & Fundamentals

- Installing Python, VS Code, Jupyter Notebook
- Variables, data types: int, float, string, boolean
- Basic operators: arithmetic, comparison, logical
- Taking input, printing output, string formatting
- Comments and code readability

Session 2 — Control Flow & Functions

- if / elif / else conditions with real examples
- for loops and while loops — when to use which
- List, tuple, dictionary, set — differences and use cases
- Writing your first functions — parameters, return values
- Mini project: Build a simple calculator

MODULE 2 — Intermediate Python

Week 2 | 2 Sessions

Session 3 — Functions, Files & Modules

- Lambda functions, map, filter, list comprehensions
- Reading and writing files — text and CSV
- Exception handling — try, except, finally
- Importing modules: os, sys, datetime, random
- Creating your own modules

Session 4 — OOP & Libraries Intro

- Object-Oriented Programming — classes, objects, methods
- Inheritance and encapsulation (simple examples)
- Introduction to pip — installing packages
- First look at NumPy and Pandas
- Week project: Student grade tracker program

★ Phase 1 Milestone Project: Build a data processing script that reads a CSV, cleans missing values, and outputs a summary report.

PHASE 2 — DATA & MACHINE LEARNING BASICS

Month 1 | Weeks 3–4 | 4 Sessions

Before training models, you need to understand and prepare data. This phase covers the full data science workflow and introduces your first machine learning models.

NumPy	Pandas	Matplotlib	Scikit-learn	Seaborn

MODULE 3 — Data Science Foundations

Week 3 | 2 Sessions

Session 5 — NumPy & Pandas Mastery

- NumPy arrays — creation, indexing, slicing, operations
- Pandas DataFrames — loading, exploring, filtering data
- Handling missing values — dropna, fillna strategies
- Merging, grouping, aggregating data
- Real dataset practice: Analyse student performance data

Session 6 — Data Visualization & EDA

- Matplotlib — line plots, bar charts, histograms, scatter plots
- Seaborn — heatmaps, pair plots, distribution plots
- Exploratory Data Analysis (EDA) — step by step process
- Finding patterns, outliers, correlations in data
- Mini project: Full EDA on a real Indian dataset (e.g. IPL or job data)

MODULE 4 — Introduction to Machine Learning

Week 4 | 2 Sessions

Session 7 — ML Concepts & First Models

- What is Machine Learning — supervised, unsupervised, reinforcement
- The ML workflow: data → features → model → evaluate → deploy
- Train/test split — why it matters
- Linear Regression — intuition, math, implementation
- Logistic Regression — classification problems
- Scikit-learn API pattern: fit, predict, score

Session 8 — Model Evaluation

- Evaluation metrics: accuracy, precision, recall, F1 score
- Confusion matrix — reading and interpreting
- Overfitting vs underfitting — bias-variance tradeoff
- Cross-validation — k-fold explained simply
- Week project: Predict student placement (Yes/No) from dataset

★ Phase 2 Milestone Project: End-to-end ML pipeline — load data, clean, visualize, train a classifier, evaluate, and present findings.

PHASE 3 — MACHINE LEARNING DEEP DIVE

Month 2 | Weeks 5–6 | 4 Sessions

Now we go deeper. Advanced algorithms, feature engineering, and building production-quality ML pipelines that companies actually use.

Random Forest	XGBoost	SVM	Pipelines	Feature Engineering

MODULE 5 — Advanced ML Algorithms

Week 5 | 2 Sessions

Session 9 — Tree-Based Models

- Decision Trees — how they split, depth, pruning
- Random Forest — ensemble learning, bagging concept
- Gradient Boosting — XGBoost, LightGBM overview
- Support Vector Machines (SVM) — margin, kernel trick
- When to use which algorithm — practical decision guide

Session 10 — Unsupervised Learning

- Clustering — K-Means algorithm step by step
- Hierarchical clustering — dendrograms
- Dimensionality Reduction — PCA intuition
- DBSCAN — density-based clustering
- Mini project: Customer segmentation using clustering

MODULE 6 — Feature Engineering & Pipelines

Week 6 | 2 Sessions

Session 11 — Feature Engineering

- What is feature engineering — why it matters more than model choice
- Encoding categorical variables — label, one-hot, target encoding
- Feature scaling — StandardScaler, MinMaxScaler
- Feature selection — correlation, importance scores
- Handling imbalanced datasets — SMOTE, class weights

Session 12 — ML Pipelines & Hyperparameter Tuning

- Scikit-learn Pipeline — chaining preprocessing + model
- GridSearchCV and RandomizedSearchCV
- Cross-validation inside pipelines
- Saving and loading models — joblib, pickle
- Week project: Salary prediction model with full pipeline

★ Phase 3 Milestone Project: Build a complete ML application with feature engineering, pipeline, hyperparameter tuning, and model export.

PHASE 4 — DEEP LEARNING

Month 2 | Weeks 7–8 | 4 Sessions

Deep Learning is the engine behind modern AI. This phase teaches you neural networks from scratch — how they learn, why they work, and how to build them for real computer vision and NLP tasks.

TensorFlow	Keras	PyTorch	CNN	Transfer Learning

MODULE 7 — Neural Networks Fundamentals

Week 7 | 2 Sessions

Session 13 — How Neural Networks Work

- Perceptron — the single neuron explained
- Multi-layer networks — input, hidden, output layers
- Activation functions — ReLU, Sigmoid, Softmax (when to use)
- Forward propagation — how data flows through the network
- Backpropagation — how the network learns (intuition, not just math)
- Gradient descent — learning rate, epochs, batch size

Session 14 — Building Your First Neural Network

- TensorFlow + Keras setup and API overview
- Building a Sequential model step by step
- Loss functions — MSE for regression, CrossEntropy for classification
- Optimizers — SGD, Adam, RMSprop compared
- Training loop — fit, validate, plot training curves
- Mini project: Handwritten digit classifier (MNIST)

MODULE 8 — CNN & Transfer Learning

Week 8 | 2 Sessions

Session 15 — Convolutional Neural Networks

- Why CNNs for images — convolution, filters, feature maps
- Pooling layers — max pooling, average pooling
- CNN architecture — building blocks explained
- Dropout and Batch Normalization — preventing overfitting
- Building a CNN for image classification

Session 16 — Transfer Learning & NLP Basics

- Transfer Learning — reusing pre-trained models
- Using VGG16, ResNet, MobileNet with Keras
- Fine-tuning — when and how to unfreeze layers
- Introduction to NLP — tokenization, embeddings
- Text classification with simple neural network
- Week project: Image classifier with transfer learning (choose your own topic)

★ Phase 4 Milestone Project: Build and deploy an image classifier — train, evaluate, save, and create a simple web interface using Streamlit.

PHASE 5 — GENERATIVE AI FOUNDATIONS

Month 3 | Weeks 9–10 | 4 Sessions

This is where everything comes together. We move into the world of Large Language Models, prompt engineering, and building real GenAI applications — the skills most in-demand in 2025.

Transformers	LangChain	OpenAI API	RAG	Vector DBs

MODULE 9 — LLMs & Prompt Engineering

Week 9 | 2 Sessions

Session 17 — How LLMs Work

- Transformer architecture — attention mechanism (conceptual)
- How GPT, Claude, Gemini are trained — pre-training, fine-tuning, RLHF
- Tokens, context window, temperature — what they mean practically
- Hallucination — why it happens and how to reduce it
- Difference between open-source and closed models
- When to use GPT-4 vs Claude vs Gemini vs LLaMA

Session 18 — Prompt Engineering Mastery

- Zero-shot, few-shot, and chain-of-thought prompting
- System prompts — how to set AI behavior and persona
- Prompt templates — structure for consistent outputs
- ReAct prompting — reasoning + acting
- Prompt injection attacks — what they are and how to prevent
- Mini project: Build a custom AI assistant with system prompts

MODULE 10 — Building with LLM APIs

Week 10 | 2 Sessions

Session 19 — OpenAI & Anthropic APIs

- API keys, rate limits, pricing — practical guide
- OpenAI Python SDK — chat completions, streaming
- Anthropic Claude API — messages API structure
- Handling API errors, retries, timeouts gracefully
- Building a simple chatbot from scratch with API
- Token counting and cost optimization strategies

Session 20 — RAG Systems & Vector Databases

- What is RAG — Retrieval Augmented Generation explained simply
- Why RAG beats fine-tuning for most use cases
- Vector embeddings — converting text to numbers the AI way
- ChromaDB setup — storing and querying vectors locally
- Building a basic document Q&A; system end-to-end
- Week project: Chat with a PDF using RAG

★ Phase 5 Milestone Project: Build a RAG-powered Q&A; system that answers questions from your own documents with source citations.

PHASE 6 — GENAI PROJECTS & CAREER PREP

Month 3 | Weeks 11–12 | 4 Sessions

The final phase is about building real, impressive projects and preparing yourself for AI roles. You will complete your portfolio, polish your resume, and do mock interviews.

LangChain	Streamlit	GitHub	Portfolio	Interviews

MODULE 11 — Advanced GenAI with LangChain

Week 11 | 2 Sessions

Session 21 — LangChain Framework

- LangChain architecture — chains, agents, tools, memory
- Building chains — LLMChain, SequentialChain
- Memory types — ConversationBufferMemory, SummaryMemory
- Tools and agents — letting AI decide what to do
- LangChain Expression Language (LCEL) — modern approach

Session 22 — Capstone Project Build

- Choose your capstone from 3 options (see projects page)
- Full development session — coding, debugging, testing
- Deploying to Streamlit Cloud or HuggingFace Spaces — free hosting
- Writing a great README for GitHub
- Recording a 2-minute demo video for portfolio

MODULE 12 — Career Preparation

Week 12 | 2 Sessions

Session 23 — Portfolio & LinkedIn

- GitHub profile optimization — pinned repos, README badges
- LinkedIn headline, about section, featured projects
- How to write project descriptions that impress recruiters
- Personal portfolio website — quick setup on Vercel/Netlify
- Resume update — AI skills section, projects, certifications

Session 24 — Mock Interviews & Next Steps

- Top 30 AI/ML interview questions — go through all
- System design question: How would you build X using AI?
- Live mock interview — Harsh interviews you, gives feedback
- Salary negotiation tips for AI roles in India
- Your 30-day post-program action plan
- Graduation — certificate, community access, ongoing mentorship

★ *Final Capstone Project: A full-stack AI application — deployed, documented, and demo-ready for job interviews.*

TOOLS & RESOURCES

Category	Tool / Platform	Purpose	Cost
Code Editor	VS Code + Jupyter	Writing and running code	Free
Python Env	Anaconda / pip venv	Package management	Free
Data Science	NumPy, Pandas, Matplotlib	Data work	Free
ML Library	Scikit-learn	Classical ML models	Free
Deep Learning	TensorFlow / Keras	Neural networks	Free
GenAI	OpenAI API	GPT-4 access for projects	Paid — ~$5/mo
GenAI Alt	Google Gemini API	Free tier available	Free tier
LLM Framework	LangChain	Building AI apps	Free
Vector DB	ChromaDB	Local vector storage	Free
Deployment	Streamlit Cloud	Host AI apps free	Free
Version Control	GitHub	Portfolio + code backup	Free
Compute	Google Colab	Free GPU for training	Free
Practice	Kaggle	Datasets + competitions	Free

ASSESSMENT & PROJECTS

Milestone Projects — One Per Phase

	Data Processing Script Read a CSV of student marks, clean missing values, compute averages, export report. **Placement Predictor** ML model that predicts if a student gets placed based on CGPA, skills, internships. **Salary Estimation Pipeline** Full ML pipeline with feature engineering, XGBoost, hyperparameter tuning, model export. **Image Classifier App** CNN trained on custom images, deployed as Streamlit web app with file upload.
	Chat With PDF RAG system — upload any PDF, ask questions, get answers with page references.
	AI Career Coach Bot LangChain + OpenAI chatbot that gives career advice, reviews resumes, suggests jobs.

Capstone Project Options (Week 12)

Choose one of the three options for your final project:

Option A — AI Study Assistant

A chatbot that answers questions from your college textbooks using RAG. Upload PDFs, ask anything, get cited answers.

Option B — Job Application AI Agent

An AI agent that takes a job description, rewrites your resume to match, drafts a cover letter, and prepares interview questions.

Option C — AI Content Creator

A multi-agent system that generates LinkedIn posts, Instagram captions, and YouTube scripts from a single topic input.

CAREER ROADMAP AFTER COMPLETION

After completing this 3-month program you will have a strong portfolio, real deployed projects, and the skills to apply for these roles:

Role	Salary Range (India)	Top Skills Needed	Difficulty
AI/ML Engineer	₹6–18 LPA	Python, Scikit-learn, LangChain, RAG	Medium
LLM App Developer	₹8–20 LPA	LangChain, OpenAI API, RAG, FastAPI	Medium
Data Scientist	₹5–15 LPA	Python, ML, Statistics, Visualization	Medium
AI Automation Specialist	₹4–10 LPA	LangChain, APIs, Make.com, Agents	Low
Prompt Engineer	₹4–12 LPA	Prompt Engineering, LLM knowledge	Low
AI Freelancer	₹30–80k/month	Any specialization + communication	Low
Remote AI Role (Global)	$600–1500/mo	Strong portfolio + English + LangChain	High

Your 30-Day Post-Program Action Plan

Week 1

- Polish all 6 milestone projects — clean code, good README
- Update LinkedIn headline, about, featured section
- Update resume — add all skills, projects, program completion

Week 2

- Apply to 10 jobs per day on LinkedIn, Naukri, Instahyre
- Send 10 personalised LinkedIn connection requests daily
- Comment on 5 AI posts daily — build visibility

Week 3

- Start preparing for interviews — 30 questions per day
- Do mock interview with a peer or record yourself
- If no responses yet — follow up, improve resume, try different titles

Week 4

- Post one AI learning update on LinkedIn — builds personal brand
- Apply for Kaggle competitions — adds credibility
- Explore freelancing on Upwork/Toptal if no job offers yet

Samjho. Karo. Badho.

AI, Sabke Liye | Made in India

genaiwallah.com