



AI Engineering Bootcamp 2026

Beginner to Advanced AI Engineering Program

Official Training Portal: codeveins01@gmail.com

Year: 2026

COURSE OVERVIEW

This **AI Engineering Bootcamp** is designed for beginners, students, software developers, data professionals, entrepreneurs, and working professionals who want to build real-world AI applications. The course starts with the fundamentals of Artificial Intelligence and Python programming and gradually progresses to advanced topics such as Large Language Models (LLMs), Retrieval-Augmented Generation (RAG), AI Agents, Multi-Agent Systems, Multimodal AI, AI Security, Fine-Tuning, and Production Deployment.

By the end of this bootcamp, learners will be able to design, build, deploy, monitor, and scale enterprise-grade AI applications.

CURRICULUM SYLLABUS

MODULE 01

Introduction to AI Engineering

- What is Artificial Intelligence? (AI, Machine Learning, Deep Learning, Generative AI)
- Evolution of AI, Real-world Applications, & Industry Trends
- AI Engineering Roadmap, Ecosystem, & Product Development Lifecycle
- Development Environment Setup: Python, VS Code, Google Colab, Git & GitHub

MODULE 02

Python Programming for AI

- Python Fundamentals: Variables, Data Types, I/O, Operators, & Documentation
- Control Flow: Conditional Statements, Loops, & Nested Logic
- Functions: Parameters, Arguments, Lambda Functions, & Recursion
- Data Structures: Lists, Tuples, Sets, & Dictionaries
- Object-Oriented Programming (OOP): Classes, Objects, & the 4 Pillars
- Error Handling & File Input/Output (TXT, CSV, JSON)

MODULE 03

APIs and Backend Development for AI

- Understanding APIs: REST Fundamentals, HTTP Methods, Request/Response Structure
- FastAPI Development: Creating Endpoints, Request Validation, & Models
- AI API Integration: OpenAI, Claude, and Gemini APIs
- Authentication, API Keys Management, and Response Handling

MODULE 04

Large Language Models (LLMs)

- Introduction to LLMs: Tokens, Tokenization, Context Windows, & Parameters
- Transformer Architecture: Attention Mechanism, Self-Attention, & Multi-Head Attention
- Popular Models: GPT Models, Claude Models, Gemini, and Open-Source Alternatives
- Model Selection: Performance vs Cost, Speed vs Accuracy, & Production Considerations

MODULE 05

Embeddings and Vector Databases

- Understanding Embeddings: Text Embeddings, Semantic & Similarity Search (Cosine Similarity)
- Vector Databases: Introduction to Qdrant, Storing and Querying Vector Data
- Practical Applications: Document Search, Recommendation Systems, & Knowledge Bases

MODULE 06

Retrieval-Augmented Generation (RAG)

- RAG Fundamentals & Architecture: Data Ingestion, Text Chunking, & Indexing
- Retrieval Techniques: Vector Search, Keyword Search, Hybrid Search, & Re-ranking
- Building Systems: Q&A Systems, Knowledge Bases, & Enterprise Search Applications
- RAG Evaluation: Hallucination Detection, RAGAS Metrics, & Quality Improvement

MODULE 07

LangChain and AI Workflows

- LangChain Fundamentals: Chains, Pipelines, Prompt Templates, & Output Parsers
- LCEL (LangChain Expression Language) & Building Optimized AI Pipelines
- Memory Systems: Conversation Memory, Long-Term Memory, & Knowledge Storage

MODULE 08

AI Agents

- Agent Fundamentals & Architecture: ReAct Framework, Tool Calling, & Reasoning Loops
- Agent Memory: Short-Term, Long-Term, and Knowledge Persistence
- Production Agents: Reliability Challenges, Error Recovery, & Performance Monitoring

MODULE 09

LangGraph and Agent Orchestration

- LangGraph Basics: Graph-Based Workflows, Nodes, Edges, & State Management
- Advanced Agent Workflows: Conditional Routing, Loops, Iterations, & Multi-Step Reasoning
- Enterprise Guardrails: Content Moderation, Compliance, GDPR, & Security Controls

MODULE 10

AI Monitoring, Evaluation and Deployment

- LangSmith: Tracing, Monitoring, and Debugging AI Workflows
- Evaluation Systems: Offline/Online Evaluations & Human Feedback Systems
- AWS Deployment: Cloud Computing Basics, API Hosting, & Production Monitoring

MODULE 11

Model Context Protocol (MCP)

- MCP Fundamentals: Why MCP Matters for AI Interoperability
- Building MCP Servers: Python Implementation & Tool Registration
- Connecting External Systems: Databases, APIs, & Enterprise Tools

MODULE 12

Context Engineering

- Understanding Context Management Strategies & Optimization
- Context Compression, Sliding Windows, & Prompt Caching
- Advanced Prompt Engineering: Structured Prompts, System Prompts, & Few-Shot Learning

MODULE 13

Multi-Agent Systems

- Multi-Agent Concepts: Collaboration, Task Distribution, & Workflow Coordination
- Advanced Architectures: Parallel Agents & Subgraph Systems
- Reliability Engineering: Circuit Breakers, Fallback Systems, & Failure Recovery

MODULE 14

Multimodal AI

- Processing & Understanding: Text, Image, Audio, and Video Understanding
- Document Intelligence: PDF Data Extraction & OCR Techniques
- Multimodal RAG: Text + Image Search & Visual Question Answering

MODULE 15

Text-to-SQL Systems

- Database Fundamentals: SQL Basics & Relational Database Queries
- Text-to-SQL Architecture: Natural Language Queries to SQL Generation
- Advanced Systems: Self-Correcting Queries, Query Validation, & Security

MODULE 16

Local AI and Small Language Models

- Small Language Models (SLMs): Qwen, Gemma, & Phi Models
- Running AI Locally: Ollama, llama.cpp, & vLLM
- Local Deployment: Local APIs, Private AI Systems, & Offline Applications

MODULE 17

Voice AI Engineering

- Speech Technologies: Speech-to-Text (STT) & Text-to-Speech (TTS)
- Voice AI Systems: Real-Time Conversations, Voice Assistants, & Context Management
- Advanced Voice Agents: Interrupt Handling, Low-Latency Systems, & Multi-Turn Conversations

MODULE 18

Advanced AI Engineering

- Cost Optimization: API Cost Management & Efficient Token Optimization
- Semantic Caching: Cache Architectures & Retrieval Optimization
- Model Routing: Intelligent Model Selection & Dynamic Routing Strategies

MODULE 19

Fine-Tuning and Synthetic Data

- Fine-Tuning Fundamentals: When to Fine-Tune vs When to Use RAG
- Parameter-Efficient Training: LoRA & QLoRA Model Adaptation
- Dataset Engineering: Data Cleaning, Deduplication, & Synthetic Data Generation

MODULE 20

Reinforcement Learning and AI Products

- Reinforcement Learning Basics: Agents, Environments, Rewards, and Policies
- Modern RL Systems: GRPO & DeepSeek-R1 Concepts for Agent Optimization
- AI Product Thinking: Identifying Use Cases, Product Design, & Success Metrics

MODULE 21

AI Security and Safety

- AI Threats: Prompt Injection Attacks, Jailbreaking, & Hallucinations
- Security Best Practices: Input Validation, Context Protection, & Tool Security
- Enterprise AI Security: Sandboxed Execution, Data Privacy, & Compliance

MODULE 22

Capstone Project

- End-to-End AI Product Development Lifecycle
- Requirement Gathering, System Design, Testing, & Production Deployment
- Final Project Presentation: Live Demonstration, Architecture Explanation, & Industry Feedback

HANDS-ON PROJECTS PORTFOLIO

BEGINNER PROJECTS

- AI Calculator Assistant
- Smart Resume Analyzer

- AI Chatbot using APIs
- Document Question Answering System

INTERMEDIATE PROJECTS

- Market Cap Calculator
- GitHub Profile Analyzer
- Domain-Specific Q&A API
- Vector Database Recommendation System
- Internal Policy Chatbot using RAG

ADVANCED & INDUSTRY LEVEL

- Agentic Organizational Chatbot
- GitHub Assistant using LangGraph
- Employee Onboarding Agent using MCP
- Multimodal Document Intelligence System
- Voice AI Assistant
- **Industry-Level AI Capstone Project**

CAREER SUPPORT & MENTORSHIP

- Professional Resume Building for AI Roles
- LinkedIn Profile Optimization & GitHub Portfolio Development
- AI Interview Preparation & Comprehensive Mock Interviews
- Capstone Project Technical Review & Industry Placement Guidance

SYSTEM HARDWARE REQUIREMENTS

Minimum Requirements (Course Join Karne Ke Liye)

Operating System	Windows 10/11 (64-bit), Ubuntu 20.04+, or macOS 11+
Processor	Intel Core i3 (6th Generation or Above) or AMD Ryzen 3 (2000 Series or Above)
Memory (RAM)	Minimum 4 GB RAM (Recommended 8 GB RAM or Higher)
Storage	Minimum 20 GB Free Space (SSD Preferred but not mandatory)
Internet	Minimum 5 Mbps Broadband Connection (Recommended 20 Mbps or Higher)
Required Software	Google Chrome, VS Code, Python 3.11+, Git & GitHub, Google Colab Account

Recommended Requirements (For Better Performance)

Processor	Intel Core i5 (10th Generation or Above) or AMD Ryzen 5 (5000 Series or Above)
Memory (RAM)	16 GB RAM
Storage	100 GB+ Free SSD Storage
GPU (Optional)	NVIDIA GTX 1650 or Above (RTX Series Preferred)

Important Note: This bootcamp is designed so that students with older laptops can also learn AI Engineering successfully. Most hands-on projects, AI models, RAG systems, Agents, LangChain, LangGraph, MCP, and Generative AI applications will be built using cloud platforms such as Google Colab, OpenAI, Gemini, Claude, and AWS. Even with 4 GB RAM and an older Core i3 processor, all activities can be fully completed using cloud-based infrastructure.