ECHO — LEARN. BUILD. LAUNCH.

LLM Engineering Roadmap

Master the technologies behind modern AI systems. Build real-world applications. Launch your AI Engineering career.



Python



Transformers



RAG



Fine-Tuning



AI Agents



LLMOps



First Job

ABOUT THIS ROADMAP

Designed for Complete Beginners

No previous AI knowledge is required. You will gradually learn how modern LLM systems work and build production-ready AI applications through hands-on projects — all the way to your first AI Engineering position.

Beginner Friendly

Zero experience required. We start from the very basics.

Project Based

Every module ends with a real, deployable project.

Industry Focused

Tools and stacks used by top AI teams worldwide.

Career Oriented

Mentorship, CV, LinkedIn, and interview prep included.

MODULE 01

Python & AI Foundations

Build the programming foundation that every AI Engineer depends on. Learn Python from scratch with a focus on data structures, scripting, and libraries used in real LLM workflows.

Key Topics

Python syntax · NumPy · Pandas · APIs · JSON · Git & GitHub

Hands-on Project

Persona CLI — A conversational command-line AI persona powered by an LLM API

Skills You'll Gain

Python scripting · REST API integration · Prompt engineering basics · Version control

Career Outcome

Ready to work as a Junior AI Developer or pursue any LLM-focused engineering role

MODULE 02

Transformers & LLMs

Understand the architecture powering every modern AI system. Explore how attention mechanisms, tokenization, and transformer models work — and learn to use leading LLM APIs professionally.

Key Topics

- Attention & Transformer architecture
- Tokenization & embeddings
- OpenAI, Anthropic & Mistral APIs
- Prompt engineering patterns

Hands-on Project

- Semantic Matcher — A tool that computes semantic similarity between texts using embeddings

Tools & Technologies

- OpenAI SDK
- Hugging Face Transformers
- LangChain
- Tiktoken

Career Outcome

- Capable of building and integrating LLM-powered features into production applications

MODULE 03

Retrieval-Augmented Generation (RAG)

Give LLMs access to your own data. RAG is one of the most in-demand skills in the AI industry. Learn to build systems that retrieve relevant information and generate grounded, accurate answers.

01

Document Ingestion

Load, split, and preprocess text data from PDFs, URLs, and databases

02

Vector Embeddings

Encode documents into semantic vector representations using embedding models

03

Retrieval Pipeline

Query vector stores to surface the most relevant context for each user question

04

Generation & Grounding

Feed retrieved context to an LLM and produce factual, cited answers

Hands-on Project

Single URL RAG Bot — Ingest any webpage and answer questions from its content

Tools & Technologies

LangChain · ChromaDB · FAISS · OpenAI Embeddings · BeautifulSoup

Career Outcome

Build enterprise-grade RAG systems — one of the top requested skills for AI Engineering roles in 2025

MODULE 04 & 05

Fine-Tuning & AI Agents

Fine-Tuning LLMs

Customize pre-trained models on domain-specific data to dramatically improve task performance without full retraining.

Key Topics

LoRA · QLoRA · PEFT · Supervised fine-tuning · Dataset curation · Hugging Face Trainer

Project

Quick Fine-Tuning — Fine-tune a small open-source LLM on a custom Q&A dataset

Outcome

Adapt any open-source model to a specific business domain or task

AI Agents

Build autonomous systems that plan, reason, use tools, and take multi-step actions to complete complex goals.

Key Topics

ReAct agents · Tool calling · Function calling · Memory · Multi-agent systems · LangGraph

Project

Calculator Tool Calling — An LLM agent that uses tools to solve mathematical problems step by step

Outcome

Design and deploy production-ready AI Agents with persistent memory and multi-tool orchestration

MODULE 06 & 07

AI Applications & LLMOps

AI Web Applications

Build and deploy full-stack AI-powered web applications that real users can interact with.

Key Topics

- FastAPI backend development
- Streaming responses
- Chat UI with Streamlit
- Authentication & sessions

Project

- Web Chat — A full-stack AI chat interface with streaming, history, and multi-turn memory

LLMOps

Move from prototype to production. Learn how to monitor, trace, evaluate, and scale AI systems reliably.

Key Topics

- Observability & tracing
- LangSmith & Weights & Biases
- Evaluation pipelines
- Docker & cloud deployment

Project

- LLMOps Tracing — Instrument a RAG pipeline with full observability, evals, and cost tracking

YOUR CAREER JOURNEY

From Student to AI Engineer

Echo doesn't stop at teaching. We guide you through every step of your career launch.

Student
You begin as a complete beginner

Foundations
Python · LLMs · RAG · Agents

Projects
7 real deployable AI projects

GitHub Portfolio
Showcase your work publicly

LinkedIn & CV
Professional-grade profiles built with mentors

Mock Interviews
Technical & behavioral preparation

First Position
AI Engineer — ready to build on day one

WHY ECHO?

Everything You Need to Launch

We don't stop at teaching. We help you launch your career.



Hands-on Projects

Build 7 real AI applications you can show employers.



Weekly Labs

Live coding sessions applying each module in practice.



GitHub Portfolio

A polished public portfolio that proves your skills.



LinkedIn Optimization

Stand out to AI recruiters with an expert-reviewed profile.



ATS-Friendly CV

A resume engineered to pass automated screening systems.



Interview Preparation

Mock interviews with detailed feedback from working engineers.



Career Mentorship

1-on-1 guidance from experienced AI professionals.



Community Support

A peer network of engineers learning and growing together.

ECHO — LEARN. BUILD. LAUNCH.

Start Building the Future of AI.

> "We don't just teach AI.
> We build AI Engineers."

Python → Transformers → RAG → Fine-Tuning → AI Agents → LLMOps → **Your First AI Engineering Position**

Ready to begin your journey?

Enroll with Echo | **Learn More**