

15 weekly evening sessions

From Developer to AI Engineer.

Most developers can prompt an LLM. Few can ship one - design it, evaluate it, secure it, deploy it. The market isn't asking who can use AI anymore - it's asking who can engineer it. That's the role that's opening up, and the one you grow into here.

THE REAL PROBLEM

Every developer now has AI in their editor. Far fewer can build the AI system itself.

Calling an API is easy. Retrieval that's actually accurate, agents that don't go off the rails, evals that prove the thing works, and a deployment that survives real traffic and real cost — that's the gap between using AI and engineering with it. That's not hype. It's a skill set — and it's learnable, in the Python you already write.

Course Syllabus

Day 01

LLM Fundamentals & APIs

Transformer architecture (conceptual) — attention, self-attention, encoder/decoder
Working with LLM APIs — OpenAI, Anthropic, Google
Authentication, parameters, streaming & multi-model patterns

Day 02

LLM Fundamentals & APIs: Structured Output

Structured output engineering — JSON mode, function calling, schema-driven responses
Model selection & token economics — commercial vs. open-weight, cost / quality / latency
Context window management
Lab: Multi-Model Comparison Tool

Day 03

Context Engineering

System prompt architecture — persona patterns, instruction hierarchies, prompt-injection defense
Context window management — stuffing, compression, summarization chains
MCP (Model Context Protocol) — tool descriptions & server integration
Lab: Context Engineering Pipeline

Day 04

Python for AI Engineering

Pydantic & type hints — data validation, schema definitions
Async Python patterns for streaming APIs and agent loops
Data manipulation with Pandas & NumPy for AI workflows
SQL essentials for AI & PyTorch conceptual overview



EXPERT TALK
Roey Wullman

DevOps Engineer,
AI expert & Mentor
Develeap

Day 05

RAG End-to-End: Retrieval Foundations

Embeddings & vector databases — Pinecone, pgvector, Weaviate
Chunking & retrieval — fixed, semantic, recursive; hybrid search; re-ranking
Document ingestion patterns — heuristics vs. LLM extraction

Day 06

RAG End-to-End: Evaluation & Build

Text-to-SQL as retrieval — schema injection, safety gates, traceability
RAG evaluation — retrieval quality, answer faithfulness, hallucination detection
Lab: Production RAG System — hybrid search, ingestion, eval suite, monitoring




EXPERT TALK
Shahaf Segev

DevOps Engineer & AI expert
Develeap, Legion Security

Day 07

Agentic AI & Orchestration: Foundations

Agent fundamentals — planning, reasoning loops, memory, tool use
LangChain & LangGraph — chains, agents, state, graph-based flows
MCP in practice — building MCP servers & clients



EXPERT TALK
Elad Cirt

DevOps Engineer & AI expert
Develeap, Philips, Akamai

Day 08

Agentic AI & Orchestration: Multi-Agent

Multi-agent systems with CrewAI — delegation, supervisor patterns, handoffs
Testing agentic workflows — TDD for agents, guardrails, eval loops
Lab: Multi-Agent System — LangGraph + MCP tools with test-driven evaluation

Day 09

Eval, Security & Models: Quality & Safety

Eval-Driven Development (EDD) — design evals first, automated quality gates
LLM evaluation & LangSmith — benchmarks, LLM-as-judge, tracing, cost monitoring
AI security — prompt-injection defense, data-leakage prevention, agent permissions



EXPERT TALK
Yotam Halperin

DevOps Engineer & AI expert
Develeap, XM, Fireblocks

Day 10

Eval, Security & Models: Local & Fine-Tuning

Local & open-weight models — Ollama, Llama, Mistral: when to go local
Fine-tuning fundamentals — SFT, LoRA / QLoRA: when, why, and how
Lab: Eval & Security Suite — quality gates, security tests, observability

Day 11

Production & Deployment: Containers & Cloud

Containerizing AI applications — Docker for model serving, multi-stage builds
AWS Bedrock — managed LLM APIs, model access, inference configuration



EXPERT TALK
Arnon Tzori

DevOps Engineer & AI expert
Develeap, Bluespine

Day 12

Production & Deployment: CI/CD & Monitoring

CI/CD for AI applications — prompt versioning, eval gates, deployment, rollback
Monitoring & cost management — token tracking, latency dashboards, drift detection
Lab: Full Deployment — deploy to AWS with CI/CD, monitoring & cost alerts

Day 13

Capstone: Project Kickoff & Architecture

How to approach a real AI project — scope a problem from your own work
System design — choosing a RAG vs. agentic architecture
Begin building your production-grade capstone



EXPERT TALK
Kobi Avshalom

CTO
Develeap

Day 14

Capstone: Build, Evaluate & Deploy

Implement your RAG or agentic system end-to-end
Add an evaluation suite, security considerations, and AWS deployment
Personalized feedback to strengthen your portfolio project



EXPERT TALK
Inbal Drori

COO , Develeap
CEO , Once

Day 15

Ship & Launch Your AI Career

Final demo — ship a live system used by at least one real user
Interview simulation, tips & guidance; resume writing workshop
Graduation & access to the Develeap AI Alumni Community



EXPERT TALK
Dr. Shani Horowitz-Rozen

Chief Journey Officer,
Develeap

Become an AI Engineer in Just 15 Sessions

Unlock your next career move in AI Engineering

Ready to bring AI engineering

to your team?

Contact us today!



megane.dreyfuss@develeap.com

Contact us now



Megane Dreyfuss

Head of Education & Training
Develeap